TECHNICAL REPORT 05-EMIS-03

Grooming Telecommunications Networks: Optimization Models and Methods

Richard S. Barr[*] M. Scott Kingsley[†]
Raymond A. Patterson[‡]

JUNE 22, 2005

Abstract

Grooming has emerged as an active area of research within the operations research and telecommunications fields and concerns the optimization of network transmissions that span multiple distinct transmission channels, protocols, or technologies. This study explores the meaning of grooming, the technical context in which it can be applied, and example situations. A new taxonomy captures key aspects of grooming problems and is used to summarize over 50 key publications on this important traffic-engineering and optimization problem class.

Keywords: Grooming, network design, aggregation, channel assignment, multiplexing, bundling.

---

[*] Department of Engineering Management, Information, and Systems, Southern Methodist University, Dallas, TX 75205 USA, email: barr@engr.smu.edu

[†] OptionTel, LLC, Dallas, TX 75252 USA, email: scottkingsleyemail@yahoo.com

[‡] School of Business, University of Alberta, Edmonton, Alberta T6G 2R6 Canada, email: Ray.Patterson@ualberta.ca

Network grooming is an industry term that has been adopted in the academic literature and applied to a variety of optimization problems. The grooming of a network to optimize utilization of its traffic-carrying capacity can have a significant impact on its profitability, reliability, and availability.

This chapter explores the meaning of grooming within various technological contexts and provides illustrative examples. A new grooming taxonomy, PACER, is introduced and key publications in the area are summarized for quick reference.

1 What is grooming?

Grooming has come to encompass a variety of meanings within the telecommunications industry and literature. The following examples reflect this diversity.

- Cinkler (2003) defines grooming as combining traffic streams to carry a more data and distinguishes between end-to-end (sub-rate) and intermediate (core) grooming. He further examines traffic grooming and λ, or wavelength, grooming and defines hierarchical grooming as the combination of both.
- Dutta and Rouskas (2002b) define grooming in WDM (wavelength-division-multiplexing) networks as "techniques used to combine low-speed traffic streams onto high-speed wavelengths in order to minimize the network-wide cost in terms of line terminating equipment and/or electronic switching."
- Weston-Dawkes and Baroni (2002) talk about grooming architectures, in the context of mesh networks, using optical switches (OSs) and optical cross connects (OCXs) as "a strategy for the placement of intermediate grooming sites, routing of traffic, and rules for how often traffic is groomed as it traverses the network."
- Zhu et al. (2003b) investigate "next-generation optical grooming switches" and their impact on network throughput and network resource efficiency. They define traffic grooming as "a procedure of efficiently multiplexing/demultiplexing and switching low-speed traffic streams onto / from high-capacity bandwidth trunks in order to improve bandwidth utilization, optimize network throughput, and minimize network cost." They further assert that "traffic grooming is an extremely important issue for next-generation optical WDM networks to cost-effectively perform end-to-end automatic provisioning."
- Zhu and Mukherjee (2003) define traffic grooming in WDM optical networks as the bundling of "low-speed traffic streams onto high-capacity optical channels" and concur with Barr and Patterson (2001) that "grooming is a term used to describe the optimization of capacity utilization in

> transport systems by means of cross-connections of conversions between different transport systems or layers within the same system."

Thus, a consensus is forming that telecommunications grooming is the optimization of network transmissions that span multiple distinct transmission channels or methods. Grooming can occur within multiple layers of the same technology or between technologies. Grooming can be performed when signals are bundled for extended-distance transmission and when cross-connection equipment converts signals between different wavelengths, channels, or time slots.

Hence, grooming is more than assigning time-slots or optimizing traffic routing (Bennett, 2002). Grooming is complex routing, and often implicitly assumes bundling or multiple capacities or multiple layers of transmission. In an abstract sense, grooming is a complex multicommodity network flow problem with multiple transmission layers, each having its own set of constraints related to hops, distance travelled, speed of travel, capacity, etc.

Network grooming is often described in the context of a particular technology, such as SONET, WDM rings, and WDM mesh networks (Zhu and Mukherjee, 2002). The essential ingredient for a telecom network transport problem to be called "grooming" is that there are multiple layers of transport within the system. The tell-tale signs of multiple layers are cross-connections or conversions between different transport systems or layers within the same system, which may involve time-slot or frequency conversion equipment to increase a network's efficiency and effective capacity.

The results of grooming efforts include:

- Changing circuits' channels and time-slot assignments,
- Eliminating wavelength-continuity and distinct-channel assignment constraints on some or all circuits,
- Improving capacity utilization,
- Increasing the number of utilizable routing possibilities, and
- Simplifying the problem by decomposing it into easier to solve subproblems (Barr and Patterson, 2001).

2 A taxonomy of grooming problems

Grooming of telecommunications networks, then, involves optimizing an interrelated set of functions or problems, and each definition above focuses on some subset. These signal-routing and traffic-engineering functions are summarized in the following *PACER taxonomy* of grooming activities:

- *Packing* or grouping lower-speed signal units into higher-speed transport units. Examples: aggregating a set of T-1 demands into ATM cells at an access node and combining a set of ATM cells into a SONET frame.

- *Assigning* of demand units to transmission channels (e.g., time-slots, frequencies, wavelengths) within a given transport layer. Examples: assigning demands to SONET time-slots (TSA) and assigning WDM lightpaths to specific wavelengths (or λ) on each span of a given mesh or ring network.

- *Converting* signals between channels in the same transport layer. Examples: employing Time-Slot Interchange (TSI) within a SONET ADM to reshuffle time slots of transitioning traffic and using an optical crossconnect (OXC) to change lightpaths' wavelengths at a transitioning node or to switch both time-slot and wavelength for a given signal.

- *Extracting/inserting* lower-speed signal units to/from transitioning higher-speed units. For example, using an add/drop multiplexer (ADM) or B-DCS to terminate a lower-rate (sub-wavelength) SONET demand.

- *Routing* demand units between their origins and destinations. Examples: determining the OD path that each OC-3 or DS-0 demand will follow and creating a set of lightpaths in an optical network for a given demand matrix.

Grooming research addresses different PACER subsets, as detailed in Section 7, where the taxonomy is used to categorize the various models and approaches. In all cases, however, grooming involves multiplexing or bundling functions.

3 Multiplexing and bundling

Multiplexing and *bundling* methods combine multiple streams into composite streams that travel at higher speeds and with higher capacities (Doverspike, 1991). Multiplexing simultaneously transmits different messages over a communication network by partitioning the available bandwidth or other resource.

Telecommunication networks employ three types of multiplexing, each of which partitions one resource into a distinct set: space-division, frequency-division, and time-division multiplexing (Stern and Bala, 1999). The partitioning of physical space to increase transmission bandwidth is called *space-division multiplexing* (SDM); examples include the bundling of multiple fibers into a cable or using multiple fibers within a network link. Partitioning the frequency spectrum into independent channels is called *frequency-division multiplexing* (FDM) and examples are wavelength-division multiplexing, which enables a given fiber to carry traffic on many distinct wavelengths by dividing the optical spectrum into wavebands with embedded channels. The division of bandwidth time into repeated time-slots of fixed length is called *time-division multiplexing* (TDM), which allows wavelength sharing of non-overlapping time slots. Any combination of these multiplexing and bundling approaches can be utilized.

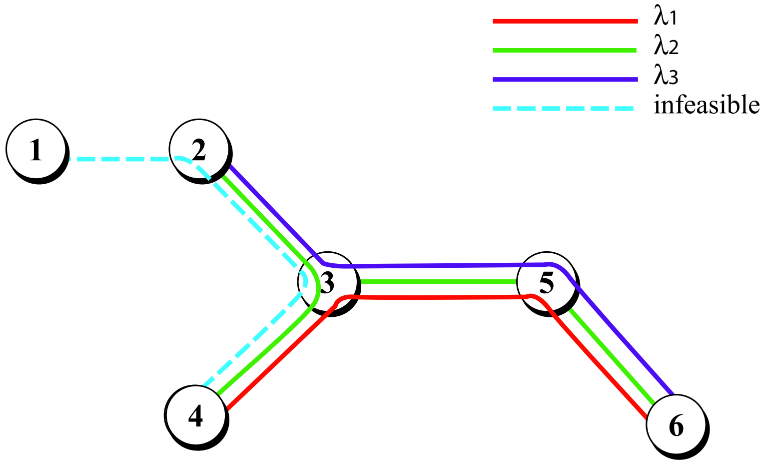

Figure 1: Unroutable 1-4 traffic due to wavelength-continuity constraints (Barr and Patterson, 2001; Betts, 1998)

4 Routing and channel assignment for lightpaths

Per Barr and Patterson (2001), routing and channel assignment in lightpaths often take the form of a multicommodity network flow problem with multiple layers of routing demand from point-to-point. Each commodity represents origin-destination connections (O-Ds) that are transported over the network, just as is done in the multicommodity flow problem. Without conversion equipment, lightpaths must be assigned to distinct channels or wavelengths referred to as a *distinct channel assignment* (DCA). Similar considerations and constraints apply to other types of bandwidth partitioning in FDM, TDM, and λ-channels in waveband-routed networks (Barr and Patterson, 2001). Day and Ester (1997) and Mukherjee (1997) illustrate the difficulty of routing and channel assignment due to *wavelength-continuity* and DCA constraints, which lead to channel conflict and contention resulting in misallocation of bandwidth capacity and limitations on wavelength reuse. Grooming in this context attempts to better utilize the wavelength capacities.

The hypothetical optical network model shown in Figure 1 demonstrates the potential impact of the wavelength continuity and DCA constraints. There is a unit demand between origin-destination node pairs 1-4, 2-4, 4-6, 3-5, 6-5, and 2-6. Each span has an capacity of three units. The tree topology creates unique routings for all demands and there is sufficient aggregate span capacity to accommodate all O-D circuits. However, not all of the demands can be assigned to one of the three λ-channels available on each link. Since O-D 1-4 cannot be accommodated, the routing and channel-assignment problem—and the overall design—is infeasible. One solution is to add one unit of capacity on links $(1, 2)$, $(2, 3)$, and $(3, 4)$ to accommodate this circuit, giving an overall

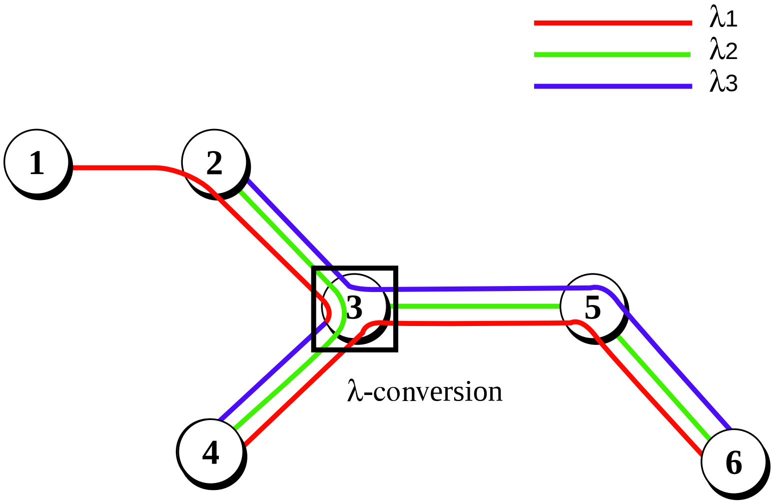

Figure 2: RCA feasibility achieved through λ-conversion grooming at node 3 (Barr and Patterson, 2001; Betts, 1998)

capacity utilization of $13/18 = 72.2\%$.

The addition of grooming equipment to this example can enable the routing of all traffic without increasing link capacities. As shown in Figure 2, if wavelength conversion is introduced at node 3, all demands can be routing on existing wavelengths. Such grooming equipment eliminates lightpath 1-4's wavelength conflict with demand 4-6 on link $(3, 4)$ by switching the connection's frequency from λ_2 to λ_1 at node 3. This increases the usable capacity of the network, permits λ_2 to be used on every network link, avoids the need for a fourth wavelength on three spans, and yields a 13/15 or 86.7% capacity utilization (an improvement of 14.5%). Cost figures are needed to evaluate the tradeoff between the costs of additional grooming equipment and the added wavelengths (Barr and Patterson, 2001).

5 Grooming within layered technologies

Gilder (2000) received much attention for proclaiming that "bandwidth is free." While this statement may have been true regarding the potential bandwidth available in installed fiber-optic networks., to be used such bandwidth must be "lit" by expensive equipment and managed by systems and people—all hardly free.

Moreover, today's telecommunications industry is complicated, capital-intensive, constantly changing, and highly competitive. With high profit margins being rare and brief, economies of scale must be achieved for a service provider to survive (Aidarous and Plevyak, 1994). Unfortunately, the provider's longevity increases this challenge, as it must also accommodate its older “legacy” equipment and technologies. The process of designing networks for adequate but not excessive capacity and assigning traffic on actual existing capacity is an important and challenging necessity.

5.1 Network planning and traffic engineering

Two important categories of telecommunications network design are capacity planning and traffic engineering. Network *capacity planning* determines a physical network’s transmission capacity, based on historical traffic levels and anticipated future demand. Primary considerations include cost, quality of service, reliability, survivability, and availability. The time between planning and availability of such capacity is often years (Sharma, 1997).

On the other hand, *traffic engineering* determines how to place actual demands on the installed capacity. The time frame for this activity is much shorter, involves traffic routing, and has the goal of an on-demand service network.

Grooming is an extension to traffic engineering whereby traffic is assigned to capacity within the network to minimize the resources (e.g., capacity, facilities) of the network in reaction to changes that occur. Grooming is primarily a local operation performed at network nodes. Although routing of traffic is not grooming, it heavily influences what and where grooming can occur. Therefore, when demand routing is considered, grooming becomes a system-level problem.

Considering that millions of traffic demands can be quickly created, these processes are daunting tasks and are often performed inefficiently. Increasingly, the success of service providers is dependent on how quickly and effectively they are accomplished.

To illustrate the necessity of grooming, and how it is accomplished, a review of conceptual network modelling is presented . These concepts are applied to specific technologies to illustrate grooming at different levels of telecommunications networks.

5.2 Telecommunications network models

Telecommunications technologies and systems use layered models to structure and group required functions, enabling the modular design of networks and equipment. Each layer of a model has specific required tasks and functions, and performs services for the adjacent layers. Complications arise from the variety of competing models and layers’ duplication of responsibility in practice. Discussion of the most prevalent models follow.

- **OSI.** One of the first layered models developed is the International Standards Organization Open Systems Interconnection model (ISO OSI). As shown in Figure 5.2, it has seven layers, with the highest being the Application Layer and the lowest the Physical Layer. Like most models, each layer has a client-server relationship with layers above and below it.

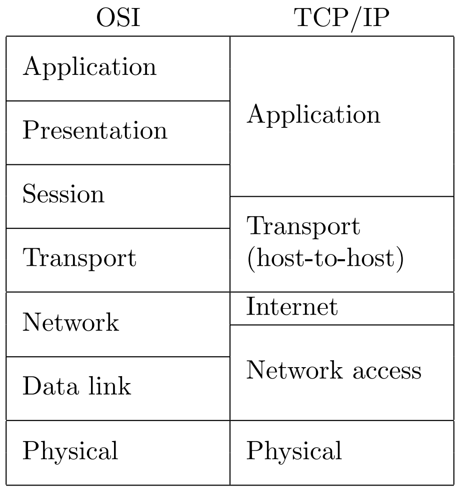

Figure 3: Comparison of OSI and TCP/IP architectures (Stallings, 2002)

The four upper layers ensure that information is delivered in correct and understandable form, end-to-end (Spragins, 1994) The lower layers provide transparent connections between users and operate primarily on a hop-by-hop basis on individual links between network nodes. This model in modern communications in many ways is most theoretical since even though many of the protocols that were defined were developed and are used are not as prevalent as those used in the current Internet.

- **TCP/IP.** The TCP/IP model developed by the Internet Engineering Task Force (IETF) uses five layers, the most recognizable of which is the Internet Protocol (IP). The designation “IP network” may sound like the only component is the IP, but the IP packets require the services of other layers, such as the Transmission Control Protocol (TCP) above it and a method of transport (such as T-1 or SONET) below it on a physical medium such as copper or fiber. Hence, an IP network requires many functions provided by other layers to be useful (Tanenbaum, 1996). Figure 5.2 shows the correspondences between the TCP/IP and OSI models’ layers.
- **IEEE.** Other organizations have also developed important widely used protocols such as the International Electrical and Electronic Engineers (IEEE), which developed popular local-area-network protocols, including Ethernet.

To make an end-to-end connection in the network requires specific functions be performed at each of the major layers and there are many combinations that can be used. The layers can be viewed as a funnel whereby traffic of different types (voice, data, and video) and formats from many different users is channelled into a stream of bits for transport to other locations in the network.

An important element is *encapsulation*, whereby the user data is packaged with overhead information, such as origin and destination addresses, that the layer is responsible for providing (Schwartz, 1987). This header information is often combined with similar formations from other users via multiplexing, as detailed in section 3.

In most models, the information (e.g., voice, data, video) starts at the top of the layer "stack." Each layer attaches its overhead and passes the information to the next layer, which encapsulates it again with more overhead and instructions to be performed, possibly multiplexing it again with streams using the same or different formats. While network design and grooming would be simpler if these encapsulation and multiplexing functions were performed by one type of equipment, that is seldom the case. In the next section, a few examples further illustrate the concept of layering and multiplexing.

5.3 Layering examples

A company's end user sends an e-mail from his computer. The e-mail is formatted using higher-level protocols with overhead information describing its format, then further encapsulated with the TCP and IP containing addressing and other information. The TCP/IP combination is then transported across the company's local-area network, encapsulated within an Ethernet frame. Emails leaving the company are switched outside the company using a router. To be transported, these packets are multiplexed with others into some type of protocol and medium. Although they could be transported using Ethernet, they more commonly use a T-1- or SONET-formatted stream that encapsulates all the packets with additional overhead information before delivery to the Internet service provider (ISP).

The ISP receives many types of information (voice, video and data) from many sources. It may use an ATM system (described below) to create virtual circuits that combine traffic arriving from many locations with traffic with similar destinations. his process involves reversing much of the encapsulation to retrieve, for example, the IP packets from a T-1 frame, to determine their destinations. Once ATM then groups packets with similar destinations into cells for transport using T-1 or SONET. The switching of the ATM cells in the network is also decided on a per-node basis. Similarly, this also cannot be done without extracting ATM cells from any encapsulation attached by subsequent layers. Therefore, there are many stages of multiplexing, transporting, de-multiplexing, switching, re-multiplexing, and transporting before each email reaches its destination.

Note that each layer affects system efficiency by introducing complexity, delay, and overhead. In addition, each layer imposes its own set of constraints on the size, structure, and amount of information that it can carry.

An analogy for this process is a postal system, wherein a letter is placed in an envelope with a destination and return address. When it reaches the post office it is combined with other letters and packages of differing sizes, delivery locations, and requirements. These are groomed by sorting and combining into larger boxes and containers to be sent to different locations. They may be transported by truck for regular delivery or sent via airplane for overnight service. There may be intermediate destinations, at which containers must be opened, the letters and packages resorted for delivery, placed in other containers, and forwarded on. The transport units (boxes, containers, trucks, and airplanes) may be completely full or nearly empty depending on the number of letters and packages involved. With such a system, it is possible that a box or airplane could carry only one letter—an enormous waste of transport capacity—or be given more than it could deliver, thus causing message delays or losses.

These same principles apply to capacity planning and traffic engineering in telecommunications networks. Hence, determining the locations of network equipment (post offices), the protocols (boxes and packages), and the transport types (trucks and airplanes) is vitally important. Even if all the capacity were perfectly utilized, the efficiency can be less than 50% due to the overhead required at the various layers. Strategies include: optimizing capacity utilization through grooming, reducing the number of layers, process automation, and increasing transport speed.

6 Grooming specifics by technology

This section describes grooming from the perspective of key transport technologies, addresses the major areas to be considered, and highlights the complexity of the planning, routing and grooming problem. It also sets the stage for discussion of work in progress to overcome these problems. Key to the understanding of the grooming problem is hidden within multiplexed systems and label switching.

6.1 Multiplexing and label switching

Telecommunications relies on multiplexing to combine many lower-speed signals onto higher-speed lines. TDM time-division multiplexing converts analog signals (such as voice phone calls) to digital form and mixed together in a single digital bit stream of separate fixed-length channels for each input analog signal. Two methods are prevalent: T-1/T-3 systems, based on the North American Digital Hierarchy, and the newer SONET standard, which converts the bit streams to light waves for transmission over distance. WDM combines many wavelengths by assigning them to different frequencies for simultaneous transmission on fiber optic strands. Used primarily in data networks, statistical multiplexing breaks input data streams into variable-sized containers (such as IP packets and ATM cells) for transmission and routing. These systems operate on a "first-come, first-served" basis. They are usually oversubscribed in that they assume that not everyone will be sending at the same time. If they did, the network would saturate. The network would buffer, or temporarily save, a limited amount of data but beyond that limit the data would be lost. Determining how much capacity is required as not to have too much (and therefore excess expense) or too little (resulting in lost data) capacity is determined using statistical calculation methods of queuing theory (Betsekas and Gallager, 1992).

Modern networks require the use of multiple types of multiplexing to move different types of information (voice, video, and data) across the network, creating a hierarchical, multi-commodity scenario. The placement of data within individual layers as well as nesting them together creates a hierarchical multi-commodity-flow grooming problem. Although routing is not specifically grooming, it does have a significant impact on the problem since it dictates what can be groomed into or from.

An improvement on conventional multiplexing has been the introduction of label switching. Using this technique, many of the layers of a network can be masked by assigning labels with additional information that simplifies the multiplexing and routing introduced in other layers. Label switching was introduced in asynchronous transfer mode and improved in multiprotocol label-switching networks. The following section gives more detailed descriptions of multiplexing and label switching by specific technology, to provide a context for network grooming.

6.2 PDH and DS-1/DS-3 multiplexing

Before the advent of the modern Internet data, the public communication network was intended primarily for transferring analog voice calls. To do the same for digital data requires the use of modems that convert the data to analog and transfer it, usually on dedicated, wired connections between endpoints. Since analog phone calls require several wires, a method of multiplexing calls onto a few wires was needed to reduce the number of cables required between telephone switches. The most important method converts 24 voice calls to digital and multiplexed them together for transmission on two copper wires.

This system is called Digital Signal Level 1 (DS-1) and operates at a speed of 1.544 Mbps; when applied to copper wire it is called T-1. Higher digital signal levels were developed to form the North American Plesiochronous Digital Hierarchy (PDH), along with a similar system in Europe (Bellamy, 1991; Grover, 2004). Originally used only in service-provider transport networks and large private networks for voice calls, small businesses now use T-1s for voice calls and Internet data.

Conventional TDM formats (DS-1, DS-3) are point-to-point connections with equipment only at the endpoints. because the T-1 or T-3 channels cannot not be dynamically accessed, to drop or add a channel between the endpoints requires the entire stream to be de-multiplexed, the channels extracted and inserted, and the new channel stream re-multiplexed before continuing. This results in added equipment and expense.

Grooming such networks with digital cross-connect switch (DCS) systems improves this situation. The DCS is a large machine placed at high-traffic locations to connect to large numbers of T-1s or T-3s. The equipment de-multiplexes and "cross-connects" or switches individual channels from incoming connections, the re-multiplexes back into outgoing T-1s and T-3s.

The DCS is an important grooming location in networks. It also made possible the mesh network architecture. However, the mesh also resulted in "backhauling," which occurs when a channel has to go to a hub location where it is taken out of the group it is in (say a T-1) and placed into another group (again maybe another T-1) that is going to another location before it can go to the location it is intended. This is analogous to a trip from Dallas to New York but having to go to Houston for a connecting flight.

6.3 Synchronous optical network

The advent of fiber optics and light waves greatly expanded the ability to transfer large amounts of data at high speeds. The Synchronous Optical Network (SONET) was developed to take advantage of this capability. Initially, SONET equipment initially had T-1 low-speed inputs which were time-division multiplexed, or mapped, into a SONET frame. Part of the SONET frame is reserved for extensive overhead information to allow monitoring and to identify where the T-1s were mapped into the SONET frame since they can be placed there at different starting locations. However, the size of the areas where T-1s could be mapped was rigid.

SONET has seven sections called Virtual Tributaries (VTs) that can contain four T-1s, each within a single Synchronous Transfer Signal Level 1 (STS-1) which operates at 54Mbps. Many STS-1s can be further multiplexed together to create higher-level STSs (Goralski, 1997). For instance, an STS-3 contains three STS-1s. The STS-1 is the basic rate of SONET and is electrical because SONET machines work internally on digital pulses. However, before the electrical STS-1 signals can be transmitted on fiber they must be converted to a light wave. At this point it is an Optical Channel to create an OC-3. Typical optical SONET rates are OC-3, OC-12, and OC-48.

An advantage of SONET is that it is synchronous, or precisely timed. This allows individual channels to be accessed and manipulated without breaking down the SONET frame. An initial problem with SONET, though, was that it used fixed-sized VTs were often a mismatch to formats other than T-1, resulting insufficient or unused capacity in the SONET frame when lower-speed signals were mapped. Fortunately SONET evolved to include enhanced mappings of other formats (such as ATM) by cleverly working around the fixed boundaries. The first was virtual concatenation whereby VTs could be connected to carry higher-speed inputs. The Generic Framing Procedure defines how to more efficiently map and groom a wide variety of lower-speed signals, including ATM and Ethernet, into a SONET frame. The Link Capacity Adjustment Scheme defines how to dynamically adjust the capacity on SONET links dynamically.

SONET also uses DCSs that allowed for regional grooming (see Figure 4) but also introduced a new type of equipment called the add-drop multiplexer (ADM). The ADM allowed adding or dropping lower-speed signals at intermediate locations along a route with less equipment than previously required and by function also became a grooming point.

It also introduced an important new architecture based on closed-loop rings

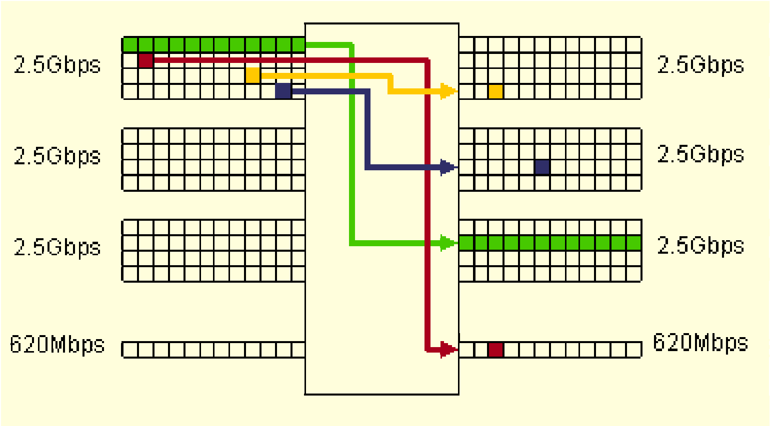

Figure 4: Digital cross-connect switch with STS-1 grooming granularity (Bennett, 2002)

which, when a failure occurred on the ring, would be “self-healing” in that only affected individual channels or the entire SONET frame will be re-rerouted around the side of the ring that is unaffected. However, the addition of protection and restoration capability as well as complexities introduced by the ring architecture resulted in much more complex planning, routing, and grooming problems.

One drawback of TDM networks is they are not historically routed adaptively like data networks are. In other words, IP and ATM networks have functionality built in whereby they route packets and cells “on the fly.” These data routes can be quickly changed as needed. In contrast, TDM networks are “nailed up” or difficult to reconfigure. Channel placement and grooming in data networks is automatic. In contrast, it often takes weeks or months to provide a TDM path. Although many TDM systems such as SONET DCSs can and do select paths and groom channels for lower level connections in the network, the methods are often proprietary and require human intervention.

6.4 (Dense) Wavelength-division multiplexing

Fiber-optic systems originally carried only one wavelength of light per fiber-optic strand. However, the available spectrum is typically large enough to potentially carry thousands of wavelengths. Therefore, much research was focused on closely packing wavelengths while avoiding interference. The multiplexing, or combining many wavelengths within a single fiber is called *wavelength-division multiplexing* or WDM. When the wavelengths are tightly placed frequency-wise it is termed *dense wavelength-division multiplexing* or DWDM (Kartalopoulos, 2000).

Low-speed inputs to DWDM systems are usually SONET, such as sixteen OC-48s. The DWDM system assigns a different wavelength or frequency to each of the OC-48s and all sixteen wavelengths would be simultaneously transmitted on the fiber for an accumulated throughput from the WDM system of OC-768. In a point-to-point configuration this provides a tremendous gain in capacity while reducing the amount of expensive fiber that is needed.

However, DWDM introduces traffic-engineering complexities. For example, once a wavelength's path has been assigned in the network, other demand paths are blocked from using that frequency on those links. What this means is the wavelength must be available on each link from the origination to the intended destination. It may be such that the wavelength is already being used in one or more of the most directly connected links of the network that are needed to complete the connection. Therefore, the wavelength that needs to terminate at the destination location must be routed on other fibers that have that wavelength available or must converted to another available wavelength on the most direct path.

This led to the introduction of the *wavelength cross-connect* (WXC), which functions much the same as DCSs described above. The WXC can switch, or groom, wavelengths from one fiber onto another at hub locations. Since the wavelengths-in-use problem may still exist, wavelength converters were developed to convert one wavelength frequency to another. The problem in wavelength grooming is routing and assigning wavelengths in the network while minimizing the high conversion-equipment expense (Ramaswami and Sivarajan, 2002).

6.5 Internet protocol (statistical multiplexing)

IP networks use statistically multiplexed packets to transfer data. They are basically on "first come, first served" basis and no guarantees are provided that information sent is actually received. IP packets are adaptively routed, meaning that each node in the packet path determines what the packet's next "hop," or routing location, will be. Each packet must be analyzed and the next hop determined.

Little grooming is accomplished in IP networks even though packets can be prioritized by assigned cost parameters on links. Therefore, certain links in the network may be heavily used while others operate well below their capacities. The tremendous growth of IP-based Internet traffic and the migration to other technologies such as Voice-over-IP has shown that much more functionality and control of large IP streams (and hence grooming) is needed. This is accomplished through label-switching, as described in the next sections.

6.6 Asynchronous transfer mode

Asynchronous Transfer Mode (ATM) is a technology and related protocols that accepts multiple streams of data in many different formats, converts them to packets and statistically multiplexing them together into *cells.* Placing the different streams of information within an ATM cell is accomplished using *virtual circuits* whereby permanent or semi-permanent logical connections are set up in the network which are used as long as they are needed and then automatically torn down. In other words, it is a "soft" connection (rather than hard-wired) through the network for transferring multiple formats simultaneously.

To avoid individual packet inspection for routing, ATM is based on *label switching*, where cells with a common destination are assigned a label that the ATM switch uses to reference a routing table to determine the outgoing port and associated link on which the cells will be transferred. ATM uses encapsulation of input data streams (such as IP) and provides the overhead with labels to communicate the information that it provides.

ATM uses two labels: the Virtual Circuit Connection (VCC), which identifies the connection endpoints and the Virtual Path Connection (VPC), which identifies a bundle of VCCs with the same endpoints. It is at the VPC level where grooming is accomplished (Grover, 2004). When an ATM switch receives cells, it looks at the cells' VPC/VCCs as well as incoming streams from other ATM switches, combines those with common endpoints (labels) and assigns new labels to new cells that are going to other switches, a technique referred to as *label swapping*. ATM is "virtual" in that streams going to many locations are aggregated using labels and share common links while being switched to other locations when the commonality diverges.

6.7 Multiprotocol label switching

IP network routing and assignment is simple, effective, but not efficient. Hence ATM's label-switching technique was adapted to IP networks via *Multiprotocol Label Switching* (MPLS) whereby packets with common characteristics could be grouped and routed together without inspecting individual packet. MPLS can groom upper-layer traffic streams from many sources and formats into similar label-switched paths (LSPs). Functionally, MPLS simply appends labels to the existing IP Packets without fully encapsulating them.

MPLS networks use *Label Edge Routers* (LERs) and *Label-Switched Routers* (LSRs). IP packets enter the network via a LER where, like ATM, they are multiplexed and assigned to *Label-Switched Paths* (LSPs) based on *Forwarding Equivalency Classes* of packets that share the same destination. This is a important step because it effectively masks the requirement to analyze the layers below in the core of the MPLS network requiring this information only be necessary at the network edges where lower-level streams ingress and egress. The labels and associated data are then transferred to LSRs in the network core where other MPLS streams are also moving across the network. Much like ATM, the MPLS LSR selects labels with common destinations and, unlike ATM, can stack labels rather than just change them. This enables LSPs with common destinations from that point to be assigned another common label over the others underneath.

MPLambdaS is an extension to MPLS for wavelength routing and grooming, in which each wavelength is assigned a label. By encoding wavelengths with labels, all the MPLS functionality can be attributed at the wavelength level. This protocol also enables dynamic, on-demand set up and tear down of wavelength routes in real or near-real time.

6.8 Generalized multiprotocol label switching

The IETF has proposed the *Generalized Multiprotocol Label Switching* (GMPLS) protocol as a framework for automating and simplifying the routing and grooming problem. The intent is that, for example, when a lower-level T-1 demand arrives it will automatically trigger its optimal placement in a new or existing SONET ST-1 and further into a new or existing STS-1 and wavelength. The hierarchical structure and the ability to stack labels allows for the extension of label switching across multiple layers of network functionality. Ideally this would result in automatic network routing and grooming at all levels. However, as in MPLS, the routing and many other important functions are not defined within GMPLS, leaving additional work to be completed in this area.

GMPLS defines five layers whereby labels and paths could be assigned: (1) the Packet Switching Capable layer, for IP, ATM, MPLS and similar streams; (2) the TDM-Capable layer, for older and SONET TDM systems components; (3) the Lambda (wavelength) Switching Capable layer, for wavelengths in WDM equipment and MPLambdaS-type systems; (4) the Waveband Switching Capable layer, for grouping and assignment of multiple wavelengths; and (5) the Fiber Switching Capable layer, for assigning wavelengths to groups of fibers.

6.9 Assembling multiple technologies

Figure 5 shows a hierarchy of transport technologies found in current telecommunications networks. The nesting effect denotes possible encapsulations (Cinkler, 2003); the directed arrows indicate typical technology up-conversions available from today's equipment manufacturers (Grover, 2004) and the associated multiplexing technique.

Many are the opportunities for grooming in such multi-tiered systems. The next section summarizes key grooming research efforts using the various technologies explored above.

7 Optimization-based grooming

Grooming has emerged as an active area of research within the operations research and telecommunications fields. Grooming problems can usually be represented as mixed-integer linear-programming and graph problems, are typically NP-hard, and must be solved with heuristic techniques for instances of any significant size (Chu and Modiano, 1998; Gerstel et al., 2000). In the past, practitioners have typically dealt with these problems by partitioning the problem into sub-problems that can be more readily solved. The partitioning point is typically at the switching (or cross-connection) point. When these problems are not dealt with holistically, the solution is almost always sub-optimal. However, these problems tend to be so complex that the only realistic way to solve

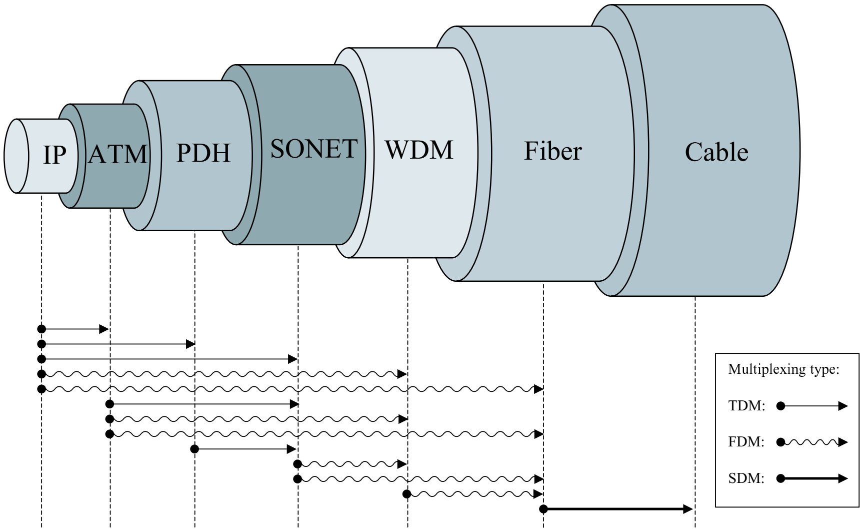

Figure 5: Optical transport layers, with typical up-conversion options. (Cinkler, 2003; Grover, 2004)

them has been to consider the sub-problems separately. Recent advances in the areas of solution techniques, commercial software, and computer processing capabilities enable us to grapple with these extremely large and difficult problems in a holistic manner (Cosares et al., 1995; Reingold, 1999).

While the basic problems and concepts of grooming telecommunications networks are known, the underlying combinatorial problems are computationally daunting. Effective algorithms and realistic models are starting to emerge and this section summarizes the current literature on grooming techniques, from an optimization perspective.

From the growing number of research publications on grooming techniques, Tables 7 and 7 summarize the problem assumptions and the optimization approach(es) reported in 50 key papers. Table 7 cites grooming research on static problems (with deterministic traffic assumptions), sorted by type of technology, and Table 7 reports on dynamic grooming methods for restructuring and re-grooming networks in response to changes in demand patters, and lists additional survey articles on grooming. These tables contain the following information for each source:

- **Citation.** Bibliographic reference.
- **Technology.** The assumed network technology—PDH, SONET, WDM, GMPLS, IP, waveband switching (WBS), and IP over optical (IPOO)—and structure (mesh, tree, ring).
- **S/D.** Traffic/demand assumptions: static (S), dynamic (D), or ordered arrival sequence (O).

- **PACER.** The grooming functions addressed in the model presented, denoted as $xx\ldots$, where x is a letter in the PACER acronym. For example, the classic paper by Doverspike (1991) on static bundling algorithms is classified PAER, since it addresses the packing, assignment, extraction/insertion, and routing functions, but not conversion between channels.

- **Model/algorithm.** The type of grooming model, solution algorithm, or theory presented, where

  - ILP: integer or mixed-integer linear programming model (possibly solved with a generic optimizer),
  - ACO: ant colony optimization (Dorigo and Di Caro, 1999),
  - GP: GRASP heuristic (Feo and Resende, 1995),
  - HEU: specialized heuristic,
  - LB: theoretical lower or upper bound,
  - MS: multi-start heuristic,
  - SP: shortest path,
  - SS: scatter-search heuristic (Glover, 1999), and
  - TS: tabu-search heuristic (Glover and Laguna, 1997),

- **Objective/criteria.** Optimization objective function or quality evaluation criteria used, where

  - Min: minimize;
  - $: cost or cost of;
  - A: number of add/drop multiplexers;
  - B: blocking probability;
  - D: total distance, delay, or hops in routings;
  - F: fiber cables;
  - L: number of wavelengths, channels, or time slots;
  - O: number of optical cross-connects;
  - P: number of paths;
  - p: count of equipment ports;
  - R: total routed traffic;
  - T: count of transponders or repeaters;
  - U: capacity utilization;
  - X: number of conversions; and
  - Δ: number of changes to existing network.

Table 1: Static grooming references, by technology and network structure.

Static-grooming: Citation	Technology	S/D	PACER	Model/algorithm	Objective/criteria	Computation?
Doverspike (1991)	PDH, mesh	S	PAER	HEU, IP	Min $	y
Betts (1998)	PDH, mesh tree	S	PACE	ILP	Min $, L	y
Carpenter et al. (1997)	SONET rings	S	AR	ILP	Min L	n
Chow and Lin (2004)	SONET rings	S	PER	LB, AA	Min A	n
Hu (2001a)	WDM ring	S	ER	Mixed ILP	Min A	y
Hu (2001b)	WDM ring	S	ER	ILP	Min A	y
Dutta and Rouskas (2002a)	WDM rings	S	PACER	ILP, LB	(other)	y
Chiu and Modiano (2000)	WDM SONET rings	S	AE	HEU	Min A	y
Wan et al. (2000)	WDM SONET rings	S	PER	AA	Min $ A	n
Zhang and Qiao (2000)	WDM SONET rings	S	PAER	HEU	Min A,L	y
Battiti and Brunato (2001)	WDM SONET rings	S	PAER	TS	Min A	y
Ghafouri-Shiraz et al. (2001)	WDM SONET rings	S	E	HEU	Min A	y
Wang et al. (2001)	WDM SONET rings	S	PACE	ILP, HEU	Min A	y
Bermond and Ceroi (2003)	WDM SONET rings	S	AE	LB	Min A	n
Zhang and Ramamurthy (2003)	WDM SONET rings	S	E	ILP, HEU, TS	Min A, ...	y
Liu and Tobagi (2004)	WDM SONET rings	S	AE	NLP, ILP	Min $ A	y
Birman and Kershenbaum (1995)	WDM mesh	S	AR	HEU	Min B	y
Fang and Somani (2003)	WDM mesh	S	PACE	ILP, Cplex	Min L	y
Zhemin and Hamdi (2003)	WDM mesh	S	PAER	LP, ILP, HEU	Min $ PL	y
Cavendish et al. (2004)	WDM mesh	S	ACR	HEU	Min LX	n
Houle et al. (2004)	WDM mesh	S	R	TS	Min $	y
Hu and Leida (2004)	WDM mesh	S	PAER	ILP, Cplex	Min L	y
Kennington and Olinick (2004)	WDM mesh	S	PACER	ILP, TS	Min $	y
Zymolka and Koster (2004)	WDM mesh	S	ACR	ILP	Min O	y
Melián et al. (2005)	WDM mesh	S	PECR	ILP, SS+TS+MS	Min $ PLF	y
Prathombutr et al. (2005)	WDM mesh	S	PAER	HEU	Max t, Min RT	y
Strand et al. (2001)	WDM mesh	SO	ACR	HEU,kSP	Min L, LD	y
Cox, Jr. and Sanchez (2001)	WDM SONET mesh	S	PACER	GA,GP	Min $ FO	y
Lardies et al. (2001)	WDM SONET mesh	S	PACER	ILP	Min $ P	y
Zhu and Mukherjee (2002)	WDM SONET mesh	S	PACER	ILP, HEU	Max t	y
Dutta and Rouskas (2002b)	WDM ring, mesh	S	PACER	ILP	Min P	n
Brunato and Battiti (2003)	WDM ring, mesh	S	PAR	GP	Min O	y
Cao et al. (2003b)	WDM WBS mesh	S	ACER	ILP, HEU	Min $	y
Kuri et al. (2004)	WDM WBS mesh	S	CER	ILP, TS	Min $	y
Parthiban et al. (2003)	GMPLS, IP, WDM	S	PE	HEU	Min O	n
Chigan et al. (2003)	IPOO	S	PACER	ILP, Cplex	Min $ pL	y

Table 2: Dynamic grooming references, by technology and network structure, and grooming survey articles.

Dynamic-grooming: Citation	Technology	S/D	PACER	Model/algorithm	Objective/criteria	Computation?
Gerstel et al. (2000)	WDM rings	SD	PAE	na	Min \$ ALT	n
Berry and Modiano (2000)	WDM SONET rings	D	PAER	Lemma, algo	Min \$ A	n
Garlick and Barr (2002)	WDM mesh	D	AR	ACO	Min B	y
Grover and Doucette (2002)	WDM mesh	SD	PACER	ILP	Min \$	y
Zhu et al. (2003a)	WDM mesh	SD	PACER	HEU, SP	Min L, D	y
Thiagarajan and Somani (2001)	WDM mesh	D	PACE	Rules	Fairness	y
Sreenath et al. (2001)	WDM mesh	D	AER	HEU	Min D, Δ	y
Ou et al. (2003)	WDM mesh	D	PCER	HEU	Min U	y
Zhu et al. (2003b)	WDM mesh	D	PACER	HEU	B, U	y
Zhu et al. (2003c)	WDM mesh	D	AC	SP	Min B	y
Comellas et al. (2003)	GMPLS, IP, WDM	SD	AER	Rule, sim	Min B	y

Surveys: Citation	Technology	S/D	PACER	Model/algorithm	Objective/criteria	Computation?
Chan et al. (1998)	WDM ring, mesh	–	–	na	na	n
Modiano and Lin (2001)	SONET, Mesh, IP	S	PACER	na	na	n
Weston-Dawkes and Baroni (2002)	Mesh	S	–	na	na	n
Cinkler (2003)	WDM	SD	–	na	na	n
Cao et al. (2003a)	Waveband	SD	–	na	na	n
Zhu and Mukherjee (2003)	WDM SONET, mesh	S	–	na	na	n

- **Computation?** Were the results of computational experimentation reported [per Barr et al. (1995)]?

Clearly, static grooming has received the earlier and greater attention by researchers and, as evidenced in Table 7, has been studied over a wide range of technologies. The introduction of WDM highlighted the value of this problem class and energized work in this area. Computationally, general-purpose optimization methods can only address relatively small problem instances of the integer-programming grooming formulations, hence a variety of heuristic approaches have been developed to meet the need for solutions to problems of more realistic dimensions.

The stochastic and temporal nature of demand has motivated the study of dynamic traffic engineering methods, as reflected in Table 7. Since the frequency of significant shifts in demand affects algorithmic options, techniques range from on-line algorithms (for real-time grooming) to integer linear programs (for reconfiguring networks experiencing more fundamental demand shifts.)

For additional viewpoints on telecommunications grooming, Table 7 also cites six excellent survey articles on the subject. Each addresses different aspects of this important and varied category of network optimization.

8 Summary and conclusions

Telecommunications network grooming is an active area of research spanning multiple disciplines and technologies. The concepts and solution methods developed for telecommunications grooming problems apply to many other industries where network transportation is optimized over multiple distinct transport methods or capacities. The authors have limited their discussion to telecommunication-specific instances of grooming.

The term "grooming" has come to encompass a variety of meanings in the telecommunications literature. A consensus definition of telecommunications network grooming is developed in this paper. Grooming is usually examined in the context of a specific technology. The technical contexts in which telecommunications network grooming can be applied are also discussed.

A new taxonomy called PACER is developed to capture five key aspects of telecommunication network grooming problems: Packing, Assigning, Converting, Extracting/inserting, and Routing. Key telecommunications network grooming papers are analyzed using PACER, along with the type of technology addressed, the model or algorithmic solution approach, the objectives used by the authors, whether or not computations were used by the authors, and the traffic demand assumptions in the paper. This analysis of the literature provides a good sense of the type and nature of research of research being conducted in the area of telecommunications network grooming.

References

S. Aidarous and T. Plevyak, editors. *Telecommunications Network Management into the 21st Century.* IEEE Press, New York, 1994.

R. Barr, B. Golden, J. Kelly, M. Resende, and W. Stewart. Designing and reporting on computational experiments with heuristic methods. *Journal of Heuristics*, 1(1):9–32, 1995.

R. Barr and R. Patterson. Grooming telecommunications networks. *Optical Networks Magazine*, 2(3):20–23, May/June 2001.

R. Battiti and M. Brunato. Reactive search for traffic grooming in WDM networks. In S. Palazzo, editor, *Evolutionary Trends of the Internet*, pages 56–66. Springer, 2001.

J. Bellamy. *Digital Telephony.* John Wiley, New York, second edition, 1991.

G. Bennett. Tutorial on grooming switches. *LightReading*, February 4 2002. www.lightreading.com.

J.-C. Bermond and S. Ceroi. Minimizing SONET ADMs in unidirectional WDM rings with grooming ratio 3. *Networks*, 41(2):83–86, 2003.

R. Berry and E. Modiano. Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic. *IEEE Journal on Selected Areas in Communications*, 18(10):1961 – 1971, Oct. 2000.

D. Betsekas and R. Gallager. *Data Networks.* Prentice Hall, Englewood Cliffs, NJ, second edition, 1992.

D. Betts. *Optimization Models for Telecommunications Networks Planning and Design: Channel Assignment and Equipment Selection.* praxis, Southern Methodist University, School of Engineering, Dallas, Texas, 1998.

A. Birman and A. Kershenbaum. Routing and wavelength assignment methods in single-hop all-optical networks with blocking. In *INFOCOM '95: Proceedings of the Fourteenth Annual Joint Conference of the IEEE Computer and Communication Societies (Vol. 2)-Volume*, page 431, Washington, DC, USA, 1995. IEEE Computer Society.

M Brunato and R. Battiti. A multistart randomized greedy algorithm for traffic grooming on mesh logical topologies. In A. Bianco and F. Neri, editors, *Next Generation Optical Network Design and Modelling.* Kluwer, 2003.

X. Cao, V. Anand, and C. Qiao. Waveband switching in optical networks. *IEEE Communications Magazine*, 41(4):105 – 112, April 2003a.

X. Cao, V. Anand, Y. Xiong, and C. Qiao. A study of waveband switching with multilayer multigranular optical cross-connects. *IEEE Journal on Selected Areas in Communications*, 21(7):1081 – 1095, Sept. 2003b.

T. Carpenter, S. Cosares, and I. Saniee. Demand routing and slotting on ring networks. DIMACS Technical Report 97-02, Rutgers University, New Brunswick, NJ, 1997.

D. Cavendish, A. Kolarov, and B. Sengupta. Minimizing the number of wavelength conversions in WDM networks with hybrid optical cross-connects. In *Proceedings of the Seventh INFORMS Telecommunications Conference*, Boca Raton, FL, March 2004. INFORMS.

V.W.S. Chan, K.L. Hall, E. Modiano, and K.A. Rauschenbach. Architectures and technologies for high-speed optical data networks. *IEEE Journal of Lightwave Technology*, 16(12):2146 – 2168, Dec. 1998.

C. Chigan, G.W. Atkinson, and R. Nagarajan. Cost effectiveness of joint multilayer protection in packet-over-optical networks. *IEEE Journal of Lightwave Technology*, 21(11):2694 – 2704, Nov. 2003.

A.L. Chiu and E.H. Modiano. Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks. *IEEE Journal of Lightwave Technology*, 18(1):2 – 12, Jan. 2000.

T. Y. Chow and P. J. Lin. The ring grooming problem. *Networks*, 44(3):194–202, 2004.

A. L. Chu and E. H. Modiano. Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming. In *Proceedings of GLOBECOM 1998*, pages 322–327. IEEE, 1998.

T. Cinkler. Traffic and λ grooming. *IEEE Network*, 17(2):16 – 21, March-April 2003.

J. Comellas, R. Martinez, J. Prat, V. Sales, and G. Junyent. Integrated IP/WDM routing in GMPLS-based optical networks. *IEEE Network*, 17(2): 22 – 27, March-April 2003.

S. Cosares, D. Deutsch, I. Saniee, and O. Wasem. SONET toolkit: A decision support system for designing robust and cost-effective fiber-optic networks. *Interfaces*, 25(1):20–40, 1995.

L. A. Cox, Jr. and J. R. Sanchez. Cost savings from optimized packing and grooming of optical circuits: Mesh versus ring comparisons. *Optical Networks Magazine*, 2(3):72–90, May/June 2001.

M. Day and G. Ester. Mesh and ring architectures for the optical network layer. In *Proceedings of NFOEC 97*, 1997.

M. Dorigo and G. Di Caro. The ant colony optimization meta-heuristic. In David Corne, Marco Dorigo, and Fred Glover, editors, *New Ideas in Optimization*, pages 11–32. McGraw-Hill, London, 1999.

R. D. Doverspike. Algorithms for multiplex bundling in a telecommunications network. *Operations Research*, 39(6):925–944, 1991.

R. Dutta and G.N. Rouskas. On optimal traffic grooming in WDM rings. *IEEE Journal on Selected Areas in Communications*, 20(1):110 – 121, Jan. 2002a.

R. Dutta and G.N. Rouskas. Traffic grooming in WDM networks: past and future. *IEEE Network*, 16(6):46 – 56, Nov.-Dec. 2002b.

J. Fang and A.K. Somani. Enabling subwavelength level traffic grooming in survivable WDM optical network design. In *Proceedings IEEE Global Telecommunications Conference (GLOBECOM 2003)*, volume 5, pages 2761 – 2766. IEEE, Dec. 2003.

T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133, 1995.

R. Garlick and R. Barr. Dynamic wavelength routing in wdm networks via ant colony optimization. In M. Dorigo, G. Di Caro, and M. Sampels, editors, *Ant Algorithms*, volume 2463 of *Lecture Notes in Computer Science*, pages 250–255. Springer, 2002.

O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in WDM rings. *IEEE/ACM Transactions on Networking*, 8(5):618–630, 2000.

H. Ghafouri-Shiraz, Guangyu Zhu, and Yuan Fei. Effective wavelength assignment algorithms for optimizing design costs in SONET/WDM rings. *IEEE Journal of Lightwave Technology*, 19(10):1427 – 1439, Oct. 2001.

G. Gilder. *TELECOSM: How Infinite Bandwidth will Revolutionize Our World.* Free Press, New York, 2000.

F. Glover. Scatter search and path relinking. In M. Dorigo D. Corne and F. Glover, editors, *New ideas in optimization*, pages 297–316. McGraw-Hill Ltd., UK, Maidenhead, UK, England, 1999.

F. Glover and M. Laguna. *Tabu Search.* Kluwer Academic Publishers, Boston, 1997.

W. Goralski. *SONET: A Guide to Synchronous Optical Network.* McGraw-Hill, New York, 1997.

W. D. Grover. *Mesh-Based Survivable Networks.* Prentice Hall PTR, Upper Saddle River, NJ, 2004.

W.D. Grover and J. Doucette. Design of a meta-mesh of chain subnetworks: enhancing the attractiveness of mesh-restorable WDM networking on low connectivity graphs. *IEEE Journal on Selected Areas in Communications*, 20(1): 47 – 61, Jan. 2002.

A. Houle, B. Jaumard, and Y. Solari. Dimensioning wdm optical networks with minimum mspp configuration. In *Proceedings of the Seventh INFORMS Telecommunications Conference*, Boca Raton, FL, March 2004. INFORMS.

J. Q. Hu. Optimal traffic grooming for wavelength-division-multiplexing rings with all-to-all uniform traffic. *Journal of Optical Networking*, 1(1):32–42, 2001a.

J. Q. Hu. Traffic grooming in wavelength-division-multiplexing ring networks: a linear programming solution. *Journal of Optical Networking*, 1(11):397–408, 2001b.

J. Q. Hu and B. Leida. Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks. In *Proceedings of IEEE INFOCOM 2004*, 2004.

S. Kartalopoulos. *Introduction to DWDM Technology.* IEEE Press, Bellingham, WA, 2000.

J. L. Kennington and E. V. Olinick. Wavelength translation in WDM networks: optimization models and solution procedures. *INFORMS Journal on Computing*, 16(2):174–187, 2004.

J. Kuri, N. Puech, and M. Gagnaire. Assessing the economic benefit of introducing multi-granularity switching cross-connects in optical transport networks. Technical Report TR-ENST-2004C003, ENST Paris, Paris, July 2004.

A. Lardies, R. Gupta, and R. Patterson. Traffic grooming in a multi-layer network. *Optical Networks Magazine*, 2(3):91–99, May/June 2001.

H. Liu and F. Tobagi. A novel efficient technique for traffic grooming in WDM SONET with multiple line speeds. In *ICC '04*, 2004.

B. Melián, M. Laguna, and J. Moreno-Pérez. Minimizing the cost of placing and sizing wavelength division multiplexing and optical crossconnect equipment in a telecommunications network. *Networks*, 45(4):199–209, July 2005.

E. Modiano and P. Lin. Traffic grooming in WDM networks. *IEEE Communications Magazine*, 39(7):124 – 129, July 2001.

B. Mukherjee. *Optical Communication Networks.* McGraw-Hill, New York, 1997.

C. Ou, K. Zhu, H. Zang, L.H. Sahasrabuddhe, and B. Mukherjee. Traffic grooming for survivable WDM networks - shared protection. *IEEE Journal on Selected Areas in Communications*, 21(9):1367 – 1383, Nov. 2003.

R. Parthiban, R.S. Tucker, and C. Leckie. Waveband grooming and IP aggregation in optical networks. *IEEE Journal of Lightwave Technology*, 21(11): 2476 – 2488, Nov. 2003.

Passakon Prathombutr, Jerry Stach, and E. K. Park. An algorithm for traffic grooming in wdm optical mesh networks with multiple objectives. *Telecommunication Systems*, 28(3-4):369–386, 2005.

R. Ramaswami and K. N. Sivarajan. *Optical Networks: A Practical Perspective.* Morgan Kaufmann, San Francisco, second edition, 2002.

N. Reingold. Telcordia's SONET design tool. presentation at INFORMS Telecommunications Conference, March 1999.

M. Schwartz. *Telecommunication Networks: Protocols, Modeling and Analysis*, volume Reading, MA. Addison-Wesley, 1987.

R. Sharma. *Network Design Using EcoNets.* International Thomson Publishing, Boston, 1997.

J. Spragins. *Telecommunications: Protocols and Design.* Addison-Wesley, Reading, MA, 1994.

N. Sreenath, B. H. Gurucharan, G. Mohan, and C. Siva Ram Murthy. A two-stage approach for virtual topology reconfiguration of WDM optical networks. *Optical Networks Magazine*, 2(3):58–71, May/June 2001.

W. Stallings. *High-Speed Networks and Internets: Performance and Quality of Service.* Prentice Hall, Upper Saddle River, NJ, second edition, 2002.

T. Stern and K. Bala. *Multiwavelength Optical Networks: A Layered Approach.* Addison-Wesley, Reading, MA, 1999.

J. Strand, R. Doverspike, and G. Li. Importance of wavelength conversion in an optical network. *Optical Networks Magazine*, 2(3):33–44, May/June 2001.

A. Tanenbaum. *Computer Networks.* Prentice Hall PTR, Upper Saddle River, NJ, third edition, 1996.

S. Thiagarajan and A.K. Somani. A capacity correlation model for WDM networks with constrained grooming capabilities. In *ICC 2001*, volume 5, pages 1592 – 1596. IEEE International Conference on, June 2001.

P.-J. Wan, G. Calinescu, L. Liu, and O. Frieder. Grooming of arbitrary traffic in SONET/WDM BLSRs. *IEEE Journal on Selected Areas in Communications*, 18(10):1995 – 2003, Oct. 2000.

J. Wang, W. Cho, V. R. Vemuri, and B. Mukherjee. Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections. *Lightwave Technology, Journal of*, 19 (11):1645 – 1653, Nov. 2001.

J. Weston-Dawkes and S. Baroni. Mesh network grooming and restoration optimized for optical bypass. In *Technical Proceedings of the National Fiber Optic Engineers Conference (NDOEF 2002)*, pages 1438–1449, 2002.

S. Zhang and B. Ramamurthy. Dynamic traffic grooming algorithms for reconfigurable SONET over WDM networks. *IEEE Journal on Selected Areas in Communications*, 21(7):1165 – 1172, Sept. 2003.

X. Zhang and C. Qiao. An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings. *IEEE/ACM Transactions on Networking*, 8(5):608 – 617, Oct. 2000.

D. Zhemin and M. Hamdi. Traffic grooming in optical WDM mesh networks using the blocking island paradigm. *Optical Networks Magazine*, 4(6):7–13, Nov/Dec 2003.

Hongyue Zhu, Hui Zang, Keyao Zhu, and Biswanath Mukherjee. A novel generic graph model for traffic grooming in heterogeneous wdm mesh networks. *IEEE/ACM Trans. Netw.*, 11(2):285–299, 2003a.

K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1):122 – 133, Jan. 2002.

K. Zhu and B. Mukherjee. A review of traffic grooming in WDM optical networks: architectures and challenges. *Optical Networks Magazine*, 4(2):55–64, March/April 2003.

K. Zhu, H. Zang, and B. Mukherjee. A comprehensive study on next-generation optical grooming switches. *IEEE Journal on Selected Areas in Communications*, 21(7):1173 – 1186, Sept. 2003b.

K. Zhu, H. Zhu, and B. Mukherjee. Traffic engineering in multigranularity heterogeneous optical WDM mesh networks through dynamic traffic grooming. *IEEE Network*, 17(2):8 – 15, March-April 2003c.

A. Zymolka and A. Koster. Design of survivable optical networks by demand-wise shared protection. In *Proceedings of the Seventh INFORMS Telecommunications Conference*, Boca Raton, FL, March 2004. INFORMS.