

Introduction

Rudra Dutta, Ahmed E. Kamal, George N. Rouskas

Computer networking has certainly changed the face of business, education, and society in general in the last half a century. The increasing impact of computer networks can be traced to the twin advances of *where* the network can provide service, and *how much* the network can do. The first advance was started by the initial deployment of the Internet in its various incarnations, and carried on more recently by the advent of ubiquitous networking, utilizing wireless technology in various forms. The second has undergone an equally dramatic change, not just in terms of the volume of bits that can be transferred, but also the predictability with which this can be accomplished. The concepts of QoS and SLAs have emerged as the new metrics with which to measure the usefulness of the network, in addition to more fundamental metrics such as throughput. In providing very high performance in both throughput and QoS, optical networking technology has come to the fore in the last two decades.

The development of communication technology using fiber-optic links and lasers for signal transmission increased the quality of data communication dramatically on several fronts. The effect was to make much higher bitrates achievable than before. However, the term “networking” is usually reserved for algorithms and protocols that enable connectivity over a large set of stations which are not all connected directly to each other; that is, operating at the third layer of the OSI model, performing forwarding and routing. The term “optical networking”, therefore, is sometimes a little confusing, since it might appear that the impact of optical transmission technology is limited to the physical layer only. However, for several reasons that have become apparent over the years, considerations of physical layer characteristics cannot be ignored in network layer design. At the core of such considerations is the fact that the hallowed layering principle, although a very powerful and useful one, serves best when viewed as general guidelines rather than inflexible shackles. Like all good things, layer abstraction can be overdone.

This has become more generally recognized and acknowledged in recent years, so much so that “cross-layer” has become a generally understood term

in network design. Today, nobody would suggest designing a network layer or a transport protocol for use in a multihop wireless network without taking wireless characteristics into account. Similarly, the optical physical layer forces itself into many networking concerns. Examples include the design of logical topologies with constraints imposed by impediments specific to optical transmission systems, the design of link layer protocols and architectures with the knowledge that the ratio of transmission delay to propagation delay is dramatically smaller than in previous wired networks, and the design of network survivability with the disparity in speed and granularity of optical layer versus higher layer restoration. Such issues, recognized now as cross-layer design issues, have been pursued by the optical networking research and development community for years. Another excellent example of such “typically optical” networking research area is traffic grooming, the topic of this book.

Traffic grooming first came to be recognized as a research area in the mid-1990’s; possibly the first use of the term “grooming” in this connection was by Sasaki *et al.* in 1998, but earlier work by the same researchers as well as others can be retrospectively seen to address the same general design problem. While many variations of the problem have been investigated and many different flavors propounded, it is possible to recognize a core problem, in the light of which all others can be discussed. This core problem is essentially a network design problem of resource allocation.

In the essential grooming problem, the resource to be allocated is the switching capacity at each node of a network. The network is assumed to be composed of physical links of optical fiber that are already in place, and the demand on the network is in the form of a static traffic demand matrix that specifies a required rate of flow of traffic for each pair of nodes in the network. Naturally, to operate the network to satisfy the traffic demands, a routing of the traffic flows onto the physical topology of fibers must be obtained. However, the grooming problem is a multi-level routing problem, and one in which the goal reflects an emergent cost function rather than a direct one, as we explain below.

With Wavelength Division Multiplexing (WDM), it is possible to multiplex several optical channels of different wavelengths onto the same fiber. Optical switching technology allows an optical switch, otherwise known as an Optical Cross-Connect (OXC), to forward the optical signal arriving on an incoming fiber at the physical layer, without the use of digital electronics, onto an outgoing fiber. Further, such optical switching can be performed in a wavelength selective manner; two optical channels from the same incoming fiber can be routed, at the physical layer, onto different outgoing fiber links. Such optical switching over a series of switches forms *lightpaths* or clear optical channels. The collection of such lightpaths formed in the network is called a *virtual topology*, otherwise known as *logical topology* (because the lightpaths are *logical links*, as opposed to the physical fiber links). Specifying the sequence of physical fibers that each lightpath traverses, and specifying a specific wavelength channel that it will use on these fibers, is called Rout-

ing and Wavelength Assignment (RWA) of the virtual topology. Unless nodes are equipped with wavelength converters, optical signals can only be switched on the same wavelength channel; this is known as the wavelength continuity constraint on lightpaths.

To inject traffic into the lightpath at the source node and to extract it from the lightpath at the destination node requires the use of digital logic, and thus the nodes must be equipped with Digital Cross-Connects (DXC). With DXCs, it is also possible to extract traffic from a lightpath, and instead of forwarding it, injecting it to another lightpath. Thus a second, electronic, level of traffic forwarding is possible. Such forwarding is referred to as Opto-Electro-Optic (OEO) switching, tracing the traffic through the switching node. OEO switching should not be needed if traffic between any pair of nodes is routed on its own lightpath. However, there is very likely to be a rate mismatch between the traffic demands and the lightpath capacity that provided one of the initial drivers for traffic grooming.

With currently available optical technology, the data rate of each wavelength is on the order of 2.5-10 Gbps, while channels operating at 40 Gbps and beyond will be commercially available in the near future. These rates are large for single source-to-destination channels even for backbone networks. To utilize efficiently this capacity, a number of independent lower-rate traffic streams must be multiplexed into a single lightpath. In that case, all traffic components cannot be carried on a single lightpath from source to destination, and some electronic routing of traffic with OEO conversion becomes inevitable. However, a backbone node utilizing OEO introduces delay, variability of delay, loss of throughput, and increased probability of errors, thus there is good reason to keep electronic routing to a minimum. Moreover, with the deployment of commercial WDM systems, it has become apparent that the cost of the DXCs used for OEO switching can be one of the dominant costs in building optical networks. The amount of traffic routed electronically, and thus the above direct and indirect costs to the network, will depend on the exact virtual topology is formed and the mapping of lower-rate traffic components onto the lightpaths. Hence the concept of *grooming* traffic into available wavelengths arises to meet network cost or other design goals.

If the optical network is viewed as a directed graph $G_p = (V, E_p)$ where each edge represents an optical fiber link between its endpoints, then the traffic grooming problem, given a number of wavelengths W supported by each fiber and a grooming factor C , asks for the solution which optimizes some cost metric (such as the total DXC capacity over all nodes) to enable the routing of a given static traffic demand matrix $T = [t_{ij}]$. The t_{ij} are integer values expressing the aggregate traffic from node i to j in multiples of some base rate, and C is the bandwidth of a single wavelength in multiples of the same unit. Thus source-to-destination traffic components are in general *sub-wavelength* in nature. The solution consists of:

1. *specifying the lightpath set*, a set of lightpath requests, denoted by $G_v(V, E_v)$; the virtual connection graph on the same set of nodes as the physical topology,
2. *solving the RWA problem*, providing the routing \mathcal{R} mapping each lightpath in G_v to a sequence of links of G_p forming a path from the source of that lightpath to its destination, and the wavelength assignment \mathcal{L} , assigning a wavelength to each lightpath out of a set of W wavelengths, such that no directed link is traversed by more than W lightpaths, no directed link is traversed by more than one lightpath of every given wavelength, and (optionally) the wavelength continuity constraint is obeyed,
3. *routing or grooming traffic*, specifying a grooming solution \mathcal{G} mapping each element of T to a sequence of lightpaths in G_v from the source of that traffic component to its destination, such that no lightpath is assigned total traffic more than C ,

such that the requirement on the cost metric is met. We represent this pictorially in Figure 1.1. Thus the input to the problem is T and G_p , and the problem is to map T onto G_p . The solution requires specifying \mathcal{R} , \mathcal{L} , G_v , and \mathcal{G} , which complete the mapping.

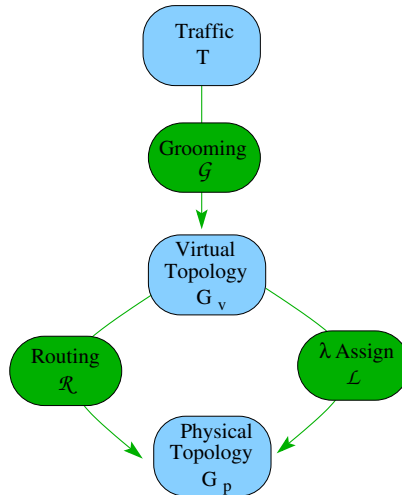


Fig. 1.1. A schematic representation of the grooming problem

Many variations of the original grooming problems have been studied in the literature. Early work concentrated on static versions of the problem, such as we have presented above, where all traffic demands are known to be reasonably unvarying over time, and embodied in a traffic matrix. In most of these variations, the problems are known to be or conjectured to be NP-hard. The literature also includes many heuristic algorithms for traffic grooming.

Nor has research been confined entirely to better algorithm design for the original problem. As the field of research has matured and responded to the changing needs as seen and foreseen by the practitioner community, the concept of traffic grooming has been carried to different arenas, and literature in this area has diverged. Comparatively recently, lower level networks in which a dynamic traffic model is more realistic have gained attention from the grooming community. It is important to note that the focus of grooming traffic shifts as a consequence of the above change in capabilities and design strategy. Reduction of OEO interchanging costs may continue to be an objective of traffic grooming. But the primary objective may now well be a minimization of the blocking behavior of the network; this is not particularly relevant in static traffic grooming because with good planning the entire traffic matrix is expected to be carried by the network. But making a similar 100% guarantee under statistically described dynamic traffic may be prohibitive in cost and not desirable. Similarly, the consideration of fairness is not relevant for the static problem, but may become an important one for the dynamic case.

For dynamic traffic, the grooming problem must be seen as one of supplying a *policy design* for the network, that is an algorithm that the network control plane can employ to make decisions in response to traffic change events. The state space consists of the current virtual topology, RWA solution, and subwavelength traffic routing. The action space ranges from admission control, to network layer routing, virtual topology modification, and subwavelength traffic component rearrangement.

Several traffic variation models are worth considering in dynamic traffic grooming, and the adoption of different models has marked the further evolution of the field, as researchers sought to understand and reflect real-world grooming scenarios. Subwavelength traffic components can be viewed simply as calls, using an arrival/departure model, Poisson or otherwise. It is also possible to view end-to-end traffic demands which are long-lived and do not completely depart, but undergo changes in magnitude from time to time, giving rise to an increment/decrement variation model. Another model that has been promising is the scheduled window model, in which traffic demands arrive or are scheduled, but provide a window within which they must be served, rather than immediately or as soon as possible. In an echo of the busy hour model of yesteryear's telephony networks, it is possible to model variation by providing entire traffic matrices. Network traffic demand is always one of these matrices, but changes from one to the other at unpredictable or statistically predictable times. The times may be completely predictable, in which case all traffic can be seen as pre-defined as scheduled, and the problem can in fact be modeled as a static problem.

The connection with the work in the protocol and signaling community is worth remarking upon. MPLS/GMPLS or ASON frameworks provide mechanisms which can be used to perform static or dynamic grooming. However, these developments have focused (as appropriate for the role of protocol standardization bodies) on enabling technology rather than design strategies. The

network administrator is provided mechanisms to set up TE or QoS actions; but what actions are to be taken is left up to the administrator, who must look elsewhere for algorithms that provide policy or strategy decisions. It is in this sense that research work such as traffic grooming provides a necessary complement to the development of enabling technology; in turn, the grooming researcher must understand what tools are realistically available now or likely to be in the near future.

In this book, we have attempted a gathering together of many of these threads that connect to the central concern of network design and resource allocation with both electronic and optical switching, and typically sub-wavelength demands. We cannot claim that we have covered the entire ground in such a divergent field, but we hope that we have come very close to it; the final judgement is up to the reader. At the risk of a little overlap, we have allowed each chapter to articulate its context within this backdrop, so that each may be read by itself without too much reference back and forth. This flow of individual topics along the same channel before diverging to their separate goals seemed appropriate for the field of traffic grooming.