Chapter 1

DESIGN OF LOGICAL TOPOLOGIES FOR WAVELENGTH ROUTED NETWORKS

Rudra Dutta
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534
rdutta@csc.ncsu.edu

George N. Rouskas
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534
rouskas@csc.ncsu.edu

Keywords: Virtual Topology, Logical Topology, Wavelength Division Multiplexing (WDM), Wavelength Routed Network, Lightpath, All Optical Network

Abstract

A *virtual topology* over a wavelength routed WAN consists of clear optical channels between nodes called *lightpaths*. These carry traffic end-to-end without electronic switching, creating an optical layer of the topology. Virtual topology design aims at combining the best of optical switching and electronic routing abilities. Designing a virtual topology on a physical network consists of deciding the lightpaths to be set up in terms of their source and destination nodes and wavelength assignment.

In this chapter we provide a complete formulation of the problem and survey the literature on the topic. We restrict ourselves to transport networks rather than local area networks, and static topology design as opposed to topologies in which individual lightpaths are set up and torn down in response to traffic demands.

1. INTRODUCTION

Wide area "All Optical Networks" with *wavelength division multiplexing* (WDM), using *wavelength routing*, are considered to be candidates for future wide area backbone networks. The ability to tap into attractive properties of optics, including the very high bandwidth potential of optical fiber, makes these networks attractive for backbone transport networks. At the same time, the WDM technique can be used to bridge the mismatch between user and fiber equipment. A fuller discussion of wide area optical networks can be found in (Green, 1992; Mukherjee, 1997; Ramaswami and Sivarajan, 1998; Green, 1996).

In recent times, there has been growing interest in virtual topology design problems on these networks. *Virtual topology* design over a WDM WAN is intended to combine the best features of optics and electronics. The architecture uses clear channels between nodes, called *lightpaths*, so named because they traverse several physical links but information traveling on a lightpath is carried optically from end-to-end. Usually a lightpath is implemented by choosing a path of physical links and reserving a particular wavelength on each of these links for the lightpath. This is known as the *wavelength continuity constraint*, indicating that a lightpath consists of a single wavelength over a sequence of physical links. Because of limitations on the number of wavelengths that can be used, and hardware constraints at the network nodes, it is not possible to set up a clear channel between every pair of source and destination nodes. The particular set of lightpaths we decide to establish on a physical network constitutes the virtual (otherwise called the logical) topology.

The tradeoff involved here is between bandwidth and electronic processing overhead. Forming lightpaths locks up bandwidth in the corresponding links on the assigned wavelength, but the traffic on the lightpath does not have to undergo optoelectronic conversion at intermediate nodes. A good virtual topology trades some of the ample bandwidth inherent in the fiber to obtain a solution that is the best of both worlds.

Optical fiber can be used simply as a point-to-point link carrying only one channel, using one wavelength. The use of WDM increases the bandwidth available and the use of virtual topologies effects reduction of delay, allowing more efficient use of bandwidth by appropriate routing. Fig. 1.1 shows a simple physical network in which lightpaths, indicated by dotted lines, have been set up to allow communication by a clear channel between nodes which are not directly connected by a fiber link. An attractive feature of the process of stepping up from point-to-point fibers to WDM and then virtual topologies is that it can be undertaken in an incremental manner with current networks (Mukherjee et al., 1996).

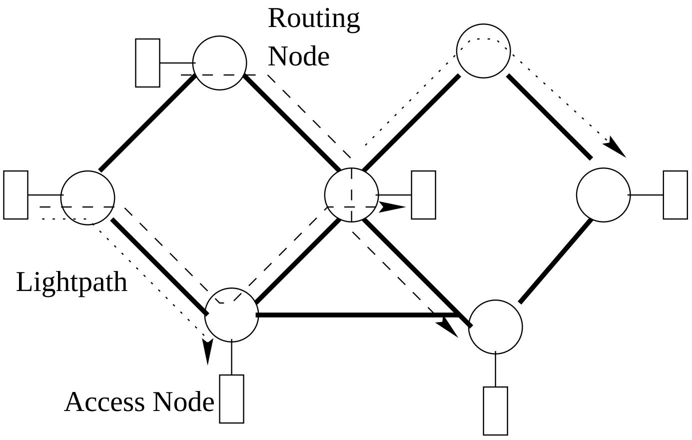

Figure 1.1 A WDM network. The routing nodes are interconnected by point-to-point fiber links and may have access nodes connected to them. The dashed lines and dotted lines show lightpaths.

The virtual topology provides a certain measure of independence from the physical topology, because different virtual topologies can be set up on the same physical topology, and allows us to choose a topology which will result in greater network performance, given network conditions such as average traffic between network nodes.

In general, virtual topology design problems can be formulated as optimization problems aimed at maximizing network throughput or other performance measures of interest. Typically, the exact solution can be shown to be NP-hard, and heuristic approaches are needed to find realistic good solutions. For this purpose, the problem can be decomposed into four subproblems, which we discuss in detail in Section 3.2.

The above discussion focuses on the issue of throughput or delay optimization, which are related to network performance. There are at least two other important related issues. The first is related to the cost required to set up and operate the network, which is an important practical consideration. Thus, a particular virtual topology may result in lower delay and higher throughput than another, but if the latter virtual topology involves the use of fewer expensive network components such as optical switches or converters, resulting in a lower overall implementation cost, then in practice it may well be chosen over the "better" one. This issue is discussed in Section 2.

The second issue relates to the reconfigurability of optical networks using virtual topology. Reconfigurability is seen as one of the strengths of optical networks in general and the virtual topology approach in particular. A virtual topology is designed on the basis of traffic patterns and a physical topology. Being able to implement a new virtual topology provides adaptability (when traffic patterns change), self-healing capability (when the physical topology changes due to failure of network components) and upgradability (when the physical topology changes due to the addition or upgrading of network components). Thus being able to redesign a virtual topology and configure the network to the new one from the old one is of interest to the virtual topology problem in general, and we have considered it within the scope of this survey.

Scope A similar virtual topology design problem exists for broadcast optical networks, used as LANs, also known as multihop networks. Virtual topology design for these networks is a different problem. One of the reasons is that with a broadcast medium, the physical topology does not constrain the virtual topologies that can be implemented. Another reason is that since each lightpath in the network needs a unique wavelength, there is no possibility of wavelength reuse as with WDM WANs. A survey of these problems for multihop networks can be found in (Labourdette, 1998).

The virtual topology design problem outlined has been formulated in terms of static traffic demands. That is, the bandwidth demand from one node to another is considered to be known when designing the virtual topology. This is distinct from topology design problems for networks in which we are interested in estimating and obtaining optimum blocking probabilities under dynamic traffic demands, that is, calls which are established and terminated on demand (Chlamtac et al., 1993). In the present case, if the traffic pattern changes significantly, it would act as the input data for a new virtual topology design. The old virtual topology would be reconfigured to the new virtual topology, a topic we discuss in Section 5.

The advantages of optical technology lie in switching and transmission, not processing or storage. Thus, electronic switching and transmission are more suitable in access networks where the bandwidth requirements are low and processing requirements (as in routing or consolidating) are relatively high. The virtual topology design problem is accordingly defined on transport networks only, not access networks.

The rest of this chapter is organized as follows. In Section 2, the architecture of wavelength routed WANs is described and notations pertaining to these are introduced. Section 3 describes approaches related to network performance optimization, including mathematical formulations and algorithms. Some particular approaches not conforming to any of these categories are described in Section 4. Section 5 addresses the reconfiguration issue. Section 6 concludes the chapter.

2. ARCHITECTURE AND NOTATIONS

Wavelength Division Multiplexing (WDM) refers to the use of distinct wavelengths over an optical fiber to implement separate channels. An optical fiber can carry several channels in parallel, each on a particular wavelength. An *add/drop multiplexer* (ADM) is an optical system that is used to modify the flow of traffic through a fiber at a routing node (Gerstel et al., 1998). An ADM passes traffic on certain wavelengths through without interruption or optoelectronic conversions, while other wavelengths are added or dropped, carrying traffic originating or terminating at the node. A *Wavelength Router* (WR) is a more powerful system than an ADM. It takes in a signal at each of the wavelengths at an input port, and routes it to a particular output port, independent of the other wavelengths (Chlamtac et al., 1993; Ramaswami and Sivarajan, 1996). A WR with N input and N output ports capable of handling W wavelengths can be thought of as W independent $N \times N$ switches. These switches have to be preceded by a wavelength demultiplexer and followed by a wavelength multiplexer to implement a WR. They are sometimes also called Wavelength Routing Switches (WRS) or wavelength crossconnects. A *Wavelength Converter* is an optical device that can be used in an optical router, to convert the wavelength a channel is being carried on without intermediate optoelectronic conversion (Ramaswami and Sasaki, 1997). Wavelength conversion allows a clear optical channel to be carried on different wavelengths on different physical links. Different levels of wavelength conversion capability are possible. *Full wavelength conversion* capability implies that any input wavelength may be converted to any other wavelength. *Limited wavelength conversion* denotes that each input wavelength may be converted to any of a specific set of wavelengths, which is not the set of all wavelengths for at least one input wavelength. If a node has limited or full wavelength conversion capability, then the conversion to be effected can be configured as part of the virtual topology design.

The advantage of wavelength conversion is that the virtual topology that can be implemented is less constrained, since the wavelength continuity constraint is removed. Thus wavelength use is more efficient. However, the use of converters increases cost, as well as the complexity of the problem. The cost can be decreased by using limited conversion rather than full conversion, and assuming a small number of converters rather than conversion capability in every node. But these assumptions introduce the problems of specifying the nature of the limited conversion and placement of converters in the network, which greatly increase the difficulty of topology design.

The virtual topology designed and implemented on a physical network not only determines the performance of the network in terms of metrics like throughput, but also carries a cost associated with the virtual topology, determined by how many and what network components are used to implement that virtual topology. Attempting to model the network cost is a topic related to the virtual topology design problem. The primary goal of such studies is to provide a measure of the relative impact of various system components on system cost, and hence provide guidelines for economically efficient virtual topology design, rather than actually determine the cost of implementing a virtual topology. Comparatively few studies have been undertaken in this area. Guidelines that result from such studies may relate to choosing some initial parameters for the virtual topology, as suggested in (Banerjee and Mukherjee, 1997), or may be integrated into the optimization procedure to find the virtual topology. The latter approach is taken in (Chen and Banerjee, 1995), where a heuristic is designed for the topology design problem with a goal of maximizing wavelength utilization in the Wavelength Routers, which would certainly have an impact on the cost of the virtual topology.

2.1 NOTATIONS

In this section, we define some terminology and notations and introduce some concepts which will be used in the following sections, and which are common to most formulations of the virtual topology problem.

Physical Topology A graph $G_p(V, E_p)$ in which each node in the network is a vertex, and each fiber optic link between two nodes is an arc. Each fiber link is also called a **physical link**, or sometimes just a link. The graph is usually assumed to be undirected, because each fiber link is assumed to be bidirectional. There is a weight associated with each of the arcs which is usually the fiber distance or propagation delay over the corresponding fiber.

Lightpath A lightpath is a clear optical channel between two nodes. That is, traffic on a lightpath does not get converted into electronic forms at any intermediate nodes, but remains and is routed as an optical signal throughout. With the usual wavelength continuity constraint, the lightpath becomes a sequence of physical links forming a path from source to destination, along with a single wavelength which is set aside on each of these links for this lightpath.

Virtual Topology A graph $G_v(V, E_v)$ in which the set of nodes is the same as that of the physical topology graph, and each lightpath is an arc. It is also called the **logical topology**, and the lightpaths are also called **logical links**. This graph is assumed to be directed, since a lightpath may exist from node A to node B while there is none from node B to node A. This graph is also weighted, with the *lightpath distance* of each lightpath (see below) acting as the weight of the corresponding arc.

Link Indicator Whether a physical link exists in the physical topology from a node l to another node m, denoted by p_{lm} which is 1 if such a link exists in the physical topology and 0 if not.

Lightpath Indicator Whether a lightpath exists from a node i to another node j, denoted by b_{ij} which is 1 if such a lightpath exists in the virtual topology and 0 if not.

Lightpath Distance The propagation delay over a lightpath, denoted by d_{ij} for the lightpath from node i to node j. It is the sum of the propagation delays over the physical links which make up the lightpath in the virtual topology.

Physical Degree The physical degree of a node is the number of physical links that directly connect that node to other nodes.

Virtual Degree The virtual (or logical) degree of a node is the number of lightpaths connecting that node to other nodes. The number of lightpaths originating and terminating at a node may be different, and we denote them by *virtual out-degree* and *virtual in-degree* respectively. We speak simply of the virtual degree if these are assumed to be equal, as they often are. If this degree is assumed to be same for all nodes of the network, then this is called the virtual degree of the network. The virtual degree is determined in part by the physical degree, but is also affected by the consideration of what volume of electronic switching can be done at a node (Ramaswami and Sivarajan, 1996).

Physical Hops The number of physical links that make up a lightpath is called the physical hop length of that lightpath.

Logical Hops The number of lightpaths a given traffic packet has to traverse, in order to reach from source to destination node over a particular virtual topology, is called the virtual or logical hop length of the path from that source to that destination in that virtual topology.

Traffic Matrix A matrix which specifies the average traffic between every pair of nodes in the physical topology. If there are N nodes in the network, the traffic matrix is an $N \times N$ matrix $\Lambda = [\lambda^{(sd)}]$, where $\lambda^{(sd)}$ is the average traffic from node s to node d in some suitable units, such as arriving packets per second, or a quantized bandwidth requirement. This matrix provides in numerical terms the nature of how the total network traffic is distributed between different source-destination node pairs, that is, the pattern of the network traffic.

Virtual Traffic Load When a virtual topology is established on a physical topology, the traffic from each source node to destination node must be routed over some lightpath. The aggregate traffic resulting over a lightpath is the load offered to that logical link. If a lightpath exists from node i to node j, the load offered to that lightpath is denoted by λ_{ij}. The component of this load due to traffic from source node s to destination node d is denoted by $\lambda_{ij}^{(sd)}$. The maximum of the logical loads is called the **congestion**, and denoted by $\lambda_{\max} = \max_{i,j} \lambda_{ij}$.

2.2 ARCHITECTURE

In this section we characterize in more detail the WDM wavelength routed network we have been describing above, and which Fig. 1.1 illustrates. The network consists of several routing nodes which are connected to each other by point-to-point optical fibers, specified by the physical topology. Each of the routing nodes may have access nodes connected to it. For the purposes of virtual topology design, however, only the aggregate traffic between routing nodes is important. Thus we can assume that each routing node has exactly one access node connected to it. We concentrate on the routing nodes and refer to them simply as nodes. The traffic matrix specifies the aggregate traffic from every node to each of the other nodes.

The fiber links connecting the nodes each support a specific number of wavelengths, say W. Each of the nodes is equipped with a WR capable of routing these W wavelengths. In general, no wavelength conversion capability is assumed to exist at any of the nodes.

Lightpaths are set up on the physical topology, creating the virtual topology. A lightpath is set up by configuring the source and destination nodes to originate and terminate a specific wavelength, then choosing a path from the source to destination node and configuring the WR at each intermediate node on that path to forward that wavelength optically to the next node. Two lightpaths that share a physical link must be assigned different wavelengths. The total number of wavelengths used on all links must be W or less. It is usually assumed that the numbers of lightpaths terminating and originating at each node are equal, and this number is same for each node. Thus the network is usually assumed to have a unique logical degree.

Traffic is routed from each source to destination node over a single lightpath if one exists for that source and destination, or a sequence of more than one lightpaths or logical hops. It is usually assumed to simplify the optimization problem that traffic for a single source-destination pair may be bifurcated over different virtual routes. The aim of creating the virtual topology is to ensure that more traffic can be carried with fewer optoelectronic conversions along the way. The extreme case of this would be if a lightpath could be set up from each source to each destination; however, the number of wavelengths available is usually too limited to allow this. At the other extreme is a virtual topology which is identical to the physical topology, so that optoelectronic conversion occurs at every intermediate node. With reasonable and achievable virtual topologies, the number of optoelectronic conversions should not be very large. Together with the fact that in high speed wide area networks the propagation delay dominates over the queueing delay (as long as links are not loaded close to capacity), queueing delays are typically neglected in the problem formulation (Ramaswami and Sivarajan, 1996).

The goal of the virtual topology design process is usually to optimize some network performance metric. Thus, a particular formulation of the problem may seek to minimize network congestion, or minimize average packet delay. In the optimization, usually the number of wavelengths available is taken as a constraint. If both minimizations are desired, then one of them is usually expressed as a constraint by relating it to a known physical network characteristic. In general both are important because too little emphasis placed on the congestion aspect usually results in a virtual topology very similar to the physical topology, and too little emphasis placed on the delay aspect can result in virtual topologies which bear little relation to the physical topology, with long lightpaths that increase delay (Ramaswami and Sivarajan, 1996).

3. PERFORMANCE OPTIMIZATION

In this section we provide an exact formulation of the virtual topology design problem using the packet traffic approach, and discuss specific techniques and heuristics used to solve it.

3.1 FORMULATION

The exact formulation of the virtual topology problem is usually given as a Mixed Integer Linear Program. The formulation provided here follows closely that in (Krishnaswamy and Sivarajan, 1998), and also those in (Ramaswami and Sivarajan, 1996; Mukherjee et al., 1994; Mukherjee et al., 1996). The symbols and terminology are as defined in Section 2.1. New terminology is defined as necessary.

Additional Definitions Let $H = [h_{ij}]$ be the **allowed physical hop matrix**, where h_{ij} denotes the maximum number of physical hops a lightpath from node i to node j is allowed to take. This hop matrix is one of the ways to characterize the bounds which lightpaths in the virtual topology must be within. Let $c_{ij}^{(k)}$ be the **lightpath wavelength indicator**, *i.e.*, $c_{ij}^{(k)}$ is 1 if a lightpath from node i to node j uses the wavelength k, 0 otherwise. Let $c_{ij}^{(k)}(l,m)$ be the **link-lightpath wavelength indicator**, to indicate whether the lightpath from node i to node j uses the wavelength k and passes through the physical link from node l to node m. Let Δ_l denote the logical degree of the virtual topology.

Objective: Subject to the constraints below, minimize the congestion of the network, that is,

$$\min \lambda_{\max} \tag{1.1}$$

Degree Constraints

$$\sum_j b_{ij} \leq \Delta_l, \quad \forall i \tag{1.2}$$

$$\sum_j b_{ji} \leq \Delta_l, \quad \forall i \tag{1.3}$$

Traffic Constraints

$$\lambda_{ij} \leq \lambda_{\max}, \quad \forall (i,j) \tag{1.4}$$

$$\lambda_{ij} = \sum_{sd} \lambda_{ij}^{(sd)}, \quad \forall (i,j) \tag{1.5}$$

$$\lambda_{ij}^{(sd)} \leq b_{ij}\lambda^{(sd)}, \quad \forall (i,j),(s,d) \tag{1.6}$$

$$\sum_j \lambda_{ij}^{(sd)} - \sum_j \lambda_{ji}^{(sd)} = \left\{ \begin{array}{ll} \lambda^{(sd)}, & s = i \\ -\lambda^{(sd)}, & d = i \\ 0, & s \neq i, d \neq i \end{array} \right\} \forall (s,d) \tag{1.7}$$

Wavelength Constraints

$$\sum_{k=0}^{W-1} c_{ij}^{(k)} = b_{ij}, \quad \forall (i,j) \tag{1.8}$$

$$c_{ij}^{(k)}(l,m) \leq c_{ij}^{(k)}, \quad \forall (i,j),(l,m),k \tag{1.9}$$

$$\sum_{ij} c_{ij}^{(k)}(l,m) \leq 1, \quad \forall (l,m),k \tag{1.10}$$

$$\begin{aligned}&\textstyle\sum_{k=0}^{W-1}\sum_l c_{ij}^{(k)}(l,m)p_{lm} - \sum_{k=0}^{W-1}\sum_l c_{ij}^{(k)}(m,l)p_{ml} \\ &= \left\{\begin{array}{ll} b_{ij}, & m=j \\ -b_{ij}, & m=i \\ 0, & m \neq i, m \neq j \end{array}\right\} \forall (i,j),m\end{aligned} \tag{1.11}$$

Hop Constraints

$$\sum_{lm} c_{ij}^{(k)}(l,m) \leq h_{ij}, \quad \forall (i,j),k \tag{1.12}$$

Discussion Most of the above constraints are self-explanatory. Many of them enforce the consistency between the various parameters and variables of the formulation. Constraint (1.7) asserts the conservation of traffic at lightpath endpoints. Expression (1.11) asserts the conservation of every wavelength at every physical node for each lightpath.

The parameters, or inputs, to the formulation are the traffic matrix Λ, the hop bound matrix H, the number of wavelengths in a fiber W, the desired logical degree Δ_l, and the details of the physical topology graph. The variables, whose values at optimum are the "output" of the MILP, relate to the virtual topology graph, wavelength assignment in the virtual topology, and the traffic routing over the virtual topology. The lightpath indicators b_{ij} provide the virtual topology graph. The lightpath wavelength and link-lightpath wavelength indicators provide the wavelength assignments to the lightpaths in the virtual topology and also the physical links implementing each lightpath. Lastly, the virtual traffic load variables λ_{ij} and $\lambda_{ij}^{(sd)}$ provide the routing of the traffic between each source and destination on the virtual topology. This formulation allows for no more than one lightpath from one node to another.

Formulations of this problem are possible that address only some and not all of these aspects. In Section 3.2 we discuss such approaches. Even when all these aspects are addressed, or the same aspect is addressed, different formulations of the problem are possible. Specific formulations can be found in the literature but are not discussed here.

This formulation gets quickly intractable with the size of the network. One of the ways it can be made more tractable is to aggregate traffic from a given source node to all destination nodes, that is, not formulate the problem in terms of the traffic components between each source-destination pair $\lambda^{(sd)}$, but traffic components for each source node $\lambda^{(s)}$ only. This results in a more tractable formulation because the number of variables and constraints is smaller, otherwise the formulation is similar. Of course, a solution to such an aggregation does not provide a complete solution, moreover there may be no corresponding complete solution. However, the aggregate problem, being less constrained than the original one, helps set achievability bounds on the full problem, such as lower bounds on the achievable congestion (Ramaswami and Sivarajan, 1996; Krishnaswamy and Sivarajan, 1998). Bounds which can be calculated with significantly lower computational costs than solving the full problem are useful in evaluating heuristics, as discussed in Section 3.2.

Usually, such an aggregate formulation is used after relaxing the MILP above into an LP, that is, allowing the lightpath, lightpath wavelength and link-lightpath wavelength indicator variables to take up values from the continuous interval $[0, 1]$ rather than constraining them to be binary variables. The relaxation, like the aggregate formulation, results in a less constrained formulation, When the MILP is relaxed, an extra “cutting plane” constraint is introduced (Ramaswami and Sivarajan, 1996; Krishnaswamy and Sivarajan, 1998), to ensure that the definition of congestion remains consistent with the MILP formulation when traffic components may be weighted with the “fractional lightpaths” that the relaxation introduces.

3.2 HEURISTICS

The problem, whose exact formulation is given in Section 3.1, and some of its subproblems are known to be NP-hard (Chlamtac et al., 1992; Mukherjee et al., 1994; Krishnaswamy and Sivarajan, 1998; Banerjee and Mukherjee, 1996). Thus for networks of moderately large sizes it is not practical to attempt to solve this problem exactly. Heuristics to obtain good approximations are needed. In the rest of this section we discuss heuristic approaches to the virtual topology design problem or to related subproblems.

Subproblems The full virtual topology design problem can be approximately decomposed into four subproblems. The decomposition is approximate or inexact. Solving the subproblems in sequence and combining the solutions may not result in the optimal solution for the fully integrated problem. It is also possible that some later subproblem may have no solution given the solution obtained for an earlier subproblem, so no solution at all to the original problem may be obtained. Although this decomposition follows (Mukherjee et al., 1996), it is also consistent with the decompositions of (Ramaswami and Sivarajan, 1996; Mukherjee et al., 1994; Krishnaswamy and Sivarajan, 1998; Banerjee and Mukherjee, 1996). The subproblems are as follows.

1. **Topology Subproblem:** Determine the virtual topology to be imposed on the physical topology, that is determine the lightpaths in terms of their source and destination nodes.
2. **Lightpath Routing Subproblem:** Determine the physical links which each lightpath consists of, that is route the lightpaths over the physical topology.
3. **Wavelength Assignment Subproblem:** Determine the wavelength each lightpath uses, that is assign a wavelength to each lightpath in the virtual topology so that wavelength restrictions are obeyed for each physical link.
4. **Traffic Routing Subproblem:** Route packet traffic between source and destination nodes over the virtual topology obtained.

In terms of the formulation provided in Section 3.1, the topology subproblem consists of determining the values of the lightpath indicator variables b_{ij}, the lightpath routing subproblem consists of determining the values of the variables $c_{ij}^{(k)}(l,m)$, the wavelength assignment subproblem consists of determining the values of the variables $c_{ij}^{(k)}$, and the traffic routing subproblem consists of determining the values of the variables $\lambda_{ij}^{(sd)}$. It may be noted that the above description of the lightpath routing subproblem is approximate, since determining $c_{ij}^{(k)}$ would be redundant after determining $c_{ij}^{(k)}(l,m)$.

The traffic routing subproblem may appear to be not essential to the virtual topology design issue. Indeed, once the virtual topology is fixed by solving the first three subproblems, the traffic routing subproblem is the known one of routing traffic over a given topology, for which many algorithms exist. However, it is included in the list of subproblems since in the exact formulation it is an integral part of the problem to determine how traffic flows over the virtual topology being designed, as it should be to optimize network performance metrics.

As we remarked above, the decomposition into subproblems is inexact. Exact solution of all the subproblems is also not possible since some of the subproblems are NP-hard as well. Heuristics must be employed to obtain good solutions to the subproblems. This also leads to the possibility of obtaining no solution to the full problem. Some constraints are usually relaxed so that at least some solution is obtained from the heuristics, which can be then tested for near optimality using achievability bounds as we discuss in the following section. One of the constraints which is commonly relaxed is that of the maximum number of wavelengths that can be carried by a fiber. Sanity checks must be performed at the end to verify that the solution obtained is feasible.

The virtual topology problem can be decomposed into different subproblems than the ones we list above. Such different decompositions are used in many of the studies we survey. However, we consider the above decomposition to be reasonable and fairly consistent with any others proposed in the literature we survey, and we shall refer only to this decomposition while discussing such studies.

Bounds To evaluate an approximate solution produced by a heuristic, we would like to know how close the obtained solution is to the optimal one. Since we are using the heuristic because of the very reason that the optimal solution cannot be obtained in the first place, we must resort to comparing the solution obtained with known bounds on the optimal solutions derived from theoretical considerations. These are the achievability bounds we have mentioned before (they are bounds on what can be achieved in principle) and we discuss them below.

Lower Bounds on Congestion The goal of virtual topology design is often to minimize network congestion, as in our formulation in Section 3.1. A lower bound on the congestion obtained from theoretical considerations allows us to know that an even smaller value of congestion cannot be achieved by any solution, and helps us evaluate the solution produced by some heuristic. We discuss several lower bounds on congestion below. Our discussion follows closely that of (Ramaswami and Sivarajan, 1996), and also that of (Krishnaswamy and Sivarajan, 1998), as well as literature on virtual topology problems in broadcast LAN scenarios as referred to in (Ramaswami and Sivarajan, 1996; Labourdette, 1998). More details can be found in these sources.

Physical topology independent bound: This bound utilizes the fact that the load on each logical link would be the same, and this would be the congestion, if the total traffic in the network were equally distributed among all the lightpaths. The value of this congestion would then act as a lower bound on any virtual topology that could be designed for the network under the given traffic conditions. This bound takes into account the total traffic demand, but not the distribution of total traffic among the different source-destination pairs (that is, the traffic pattern).

As such, it assumes that traffic for any source-destination pair can be assigned to any lightpath in the virtual topology, and hence, it ignores the physical topology.

Let $\overline{H}$ be the traffic weighted average number of logical hops in the virtual topology. If E_l denotes the number of lightpaths in the virtual topology and r denotes the total arrival rate of packets to the network, then it is easy to see that $\lambda_{\max} \geq r\overline{H} \,/E_l$, and setting a lower bound on $\overline{H}$ results in a lower bound on the congestion. For the traffic weighted number of hops to be minimum, source-destination pairs with the largest amount of traffic must be connected by a small number of logical hops. Since there can be a maximum of $N\Delta_l$ node pairs connected by a single logical hop, we assume these are exactly the node pairs with the largest traffic between them, and similarly for two, three, and larger number of logical hops. The traffic weighted average number of logical hops in this case is a lower bound. This is given by the expression $\overline{H} \geq \sum_k kS_k$ where S_k is the sum of the traffic fractions (with respect to total network traffic r) which are assumed to be carried in k hops. These traffic fractions can be determined as follows: arrange the traffic fractions in descending order of magnitude, and divide them in blocks, the i-th block being made up of $N\Delta_l^i$ successive fractions in that list. Thus the first block consists of the first $N\Delta_l$ traffic fractions, the second block consists of the next $N\Delta_l^2$ traffic fractions, and so on. Then the sum S_k is the sum of the traffic fractions which form the k-th block.

Minimum flow tree bound: This bound is derived from similar considerations as above, but on a per node basis. In this bound, we take into account the restriction that each source node can only source Δ_l lightpaths altogether, in addition to the considerations above. Thus this is a stronger bound. The traffic weighted average number of logical hops $\overline{H}$ is bounded by assuming that each node is connected by one logical hop to the Δ_l nodes to which it has the largest amounts of traffic, by two hops to the Δ_l^2 nodes to which it has the next largest amounts of traffic, and so on. We omit the derivation and exact expression of this bound, which can be found in (Ramaswami and Sivarajan, 1996).

Iterative bound: This type of bound is developed in (Ramaswami and Sivarajan, 1996; Krishnaswamy and Sivarajan, 1998) by aggregating and then relaxing the MILP formulation and solving it as mentioned in Section 3.1. The additional constraint imposed on the relaxed aggregate formulation is that the congestion be higher than a lower bound on the congestion known *a priori*, such as the minimum flow tree bound discussed above. To improve the tightness, the value obtained for the congestion by solving the relaxed aggregate LP can be used as a new value of the *a priori* bound and the LP solved again to yield a further tightened bound on the congestion. This iterative process can be carried out repeatedly to improve the tightness of the bound.

Independent topologies bound: This bound is proposed in (Banerjee et al., 1997) as another method of taking into account the physical topology in computing a bound on the congestion. Successive topologies are derived to maximize one-hop, two-hop, three-hop traffic etc., and these topologies are not allowed to constrain each other. Finally the resulting congestion is read out as a bound. The authors note that this bound is only a little tighter than the flow tree based bound if traffic is uniform, but becomes much tighter for highly nonuniform traffic.

Lower Bounds on the Number of Wavelengths It is usually necessary in virtual topology design to complete the design using as few distinct wavelengths as possible, since in practice there is a limit on the number of wavelengths a fiber can carry. This limit may be known and introduced in the exact formulation as in the formulation of Section 3.1, but such a limit is often not included in heuristic approaches. A lower bound on the number of wavelength needed for a particular problem is then useful in evaluating the solution provided by the heuristic. Also, in the presence of practical limitations, an easy-to-compute lower bound on the number of wavelengths can provide a quick negative answer to the question of whether a virtual topology design problem is at all feasible or not. Two such bounds can be found in (Ramaswami and Sivarajan, 1996). The first bound is derived from the simple consideration that the node with the minimum physical degree δ_p must source Δ_l lightpaths. Then the number of wavelengths required is bounded from below by $W \geq \lceil \Delta_l / \delta_p \rceil$. The second bound is derived by assuming that each node sources lightpaths to exactly those nodes it can reach with the minimum number of physical hops. We omit the derivation and exact expression of this bound, which can be found in (Ramaswami and Sivarajan, 1996).

Bounds on the number of wavelengths required can also be found under specific assumptions regarding the solution to the different subproblems. Two such bounds are demonstrated in (Chlamtac et al., 1993), one based on the assumption that the virtual topology being implemented is a hypercube and a specific node mapping algorithm is used, the other relating the number of wavelengths needed for topologies with and without the wavelength continuity constraint.

Heuristic Approaches and Techniques In the design of heuristics or approximate solutions to the virtual topology problem, emphasis is placed on different aspects of the problems by different authors. In the majority of the literature, heuristics are designed for only some and not all the subproblems. Some assumption regarding the nature of the virtual topology to be implemented is often a starting point for heuristic methods. Below we discuss heuristics found in the literature surveyed under three different categories. In the first, it is assumed that the virtual topology to be implemented is a well-known regular topology, such as a hypercube or a shufflenet. In the second, the lightpaths of the virtual topology are assumed to be already known in terms of sources and destinations for each instance of the problem, and the lightpath routing and wavelength assignment subproblems are addressed. No particular assumption is made regarding the virtual topology in the last category. Some of the interest in the study of regular topologies in the context of virtual topologies for WANs came from the assumption that to some extent the virtual topology could dictate the physical topology, that is, fibers could be laid to supplement a physical topology before implementing a virtual topology. As more and more fiber has been laid in practice and has become part of single wavelength optical networks utilizing the fibers as point-to-point links, the concern has shifted to extracting more utilization out of these fibers using WDM and virtual topologies, rather than having to lay more fibers. Thus studies relating to arbitrary physical topologies have attracted more interest in recent times.

Regular Topologies Regular topologies such as hypercubes or shufflenets have several advantages as virtual topologies. They are well understood, and results regarding bounds and averages are comparatively easier to derive. Routing of traffic on a regular topology is usually also simpler and results are available in the literature, so the traffic routing subproblem usually becomes trivial. Also, regular topologies possess inherent load balancing characteristics.

Once a regular topology is decided upon as the one to be implemented as a virtual topology, it remains to decide which physical node will realize each given node in the regular topology (this will be referred to as the node mapping subproblem) and which sequence of physical links between two physical nodes will be used to realize each given edge in the regular topology, that is, lightpath (this will be called the path mapping subproblem). This procedure is also called embedding a regular topology in the physical topology. In terms of the subproblems introduced in Section 3.2, the choice of the regular topology together with the node mapping problem make up the virtual topology subproblem, and the path mapping problem corresponds to the lightpath routing subproblem. Obviously, the number of nodes in the (regular) virtual topology may not be chosen with complete freedom, instead it must obey the constraints of the regular topology. In case the physical topology has a few nodes less than the regular topology, this can usually be circumvented by adding fictitious nodes to it before embedding (Mukherjee et al., 1994). If a few more nodes are present in the physical topology, then some of the ones with less traffic may be combined for the purpose of embedding, though this introduces further approximations to the virtual topology solutions. In general, the node mapping and path mapping problems leave out of consideration the network traffic pattern, and utilize metrics such as fiber distance and wavelength reuse to route lightpaths over the physical topology. Thus there is a tacit assumption of a uniform or close to uniform traffic pattern in the use of regular topologies as virtual topologies. The mappings must also be free of wavelength clashes and must obey any predefined limit on the number of wavelengths. These problems are themselves known or conjectured to be NP-hard (Chlamtac et al., 1993; Mukherjee et al., 1994), hence heuristics are needed for them.

We omit a detailed discussion of the studies available in the literature on this topic and only mention a few of them. In (Chlamtac et al., 1993), the physical topology is first mapped into an “equivalent string” preserving some characteristics of the topology, and then the selected regular topology is embedded into this string. Three different regular topologies are compared for their suitability as virtual topologies in (Marsan et al., 1993). The comparison is based on the number of logical hops between nodes and the number of wavelengths required. In (Mukherjee et al., 1994), heuristics are developed for the node mapping problem for regular topologies (specifically, hypercubes), and the solution is carried through to path mapping as well as wavelength assignment. Thus, the node mapping part of the virtual topology subproblem, the lightpath routing, and the lightpath wavelength assignment subproblems are addressed. A greedy algorithm and a simulated annealing algorithm are specified for obtaining an initial mapping and then refining the mapping based on considerations of minimizing overall network message delay. In (Mukherjee et al., 1996), a very similar simulated annealing heuristic is presented, with the difference that the traffic routing subproblem is assumed to be solved using the flow deviation method. The flow deviation method is a good heuristic alternative to an exactly optimal linear programming routing flow solution. The literature in which it was developed is referred to in (Mukherjee et al., 1996) and also (Labourdette, 1998). This method starts from an initial flow assignment, and iteratively deviates flows over alternate paths, avoiding links carrying the largest amounts of traffic.

Pre-specified Topologies We now discuss studies which focus on the lightpath routing subproblem, and possibly the wavelength assignment and traffic routing subproblems. In other words, the virtual topology in terms of a list of lightpaths with their source and destination nodes is supposed to be given for each instance of the problem.

The traffic pattern in the network would certainly have been taken into account when the lightpaths were decided upon. Because of this, in some approaches we discuss below, the traffic pattern is not taken into account, though it is also possible to utilize the traffic pattern information in routing lightpaths. The lightpath routing and wavelength assignment subproblems can then be viewed as having goals defined purely in terms of the lightpaths, such as minimization of the number of distinct wavelengths needed.

We mention a few of the studies in the literature, omitting detailed discussion as before. In (Chlamtac et al., 1992), not only the source and destination, but also the routing of the lightpaths is assumed to be given. That is, the lightpath wavelength assignment subproblem is addressed, and it is called the Static Lightpath Establishment (SLE) problem. It is proved that SLE as stated is equivalent to the n-graph-colorability problem, and hence NP-complete. A greedy heuristic algorithm to assign wavelengths to a given set of lightpaths with the aim of using as few wavelengths as possible is presented. The study presented in (Chen and Banerjee, 1995) assumes that the virtual topology subproblem has been solved and the set of lightpaths to be established is available in terms of the source and destination nodes of the lightpaths. The objective for the routing and wavelength assignment problem presented is to maximize wavelength utilization at the switches. This objective is presented in terms of the utilization of the Wavelength Routers at each network node. To formally define the problem, the concept of a "Latin Square" is introduced. Lightpath routing and wavelength assignment is posed in terms of completion of partial Latin Squares. Two heuristic algorithms are presented to complete partial Latin Squares at individual network nodes, and then a scheme is specified to use these in combination to solve the lightpath routing and wavelength assignment problem at the network level. In (Banerjee and Mukherjee, 1996), the virtual topology is assumed to be given in terms of a list of lightpaths with their source and destination nodes for each instance of the problem. The lightpath routing problem is formulated in terms of lightpath traffic as a multicommodity flow problem which is known to be NP-complete. It is suggested that the problem size can be reduced considerably by customizing the formulation for each instance of the problem, by pruning the search tree for lightpath routes and relaxing integer constraints. The study employs known heuristic methods, including randomized rounding and graph coloring, with provably good characteristics, to address the lightpath routing and wavelength assignment problems.

Arbitrary Topologies There are various studies proposing heuristic methods for arbitrary virtual topologies. These studies address the virtual topology subproblem itself, as well as some or all of the subsequent subproblems of virtual topology design. Most of these methods take into account the effect of the network traffic pattern, since arbitrary virtual topologies are usually called for in response to non-uniform traffic patterns and irregular physical topologies. Some of the heuristics proposed are similar to each other.

In (Zhang and Acampora, 1995), the problem is looked upon as the establishment of an optical connection graph over a WAN based on the average traffic demand, and then using demand based routing on this connection graph, that is, dynamic virtual circuits, which allocate whole lightpaths at a time. The connection graph subproblem presented is therefore identical to the first three subproblems of the virtual topology problem as presented in Section 3.2. The problem is formulated as a nonlinear integer programming problem, and an approximate decomposition is presented. The heuristic algorithm is then presented, which is based on a greedy approach, and attempts to utilize the number of wavelengths used for a maximum number of lightpaths.

Several different heuristics are presented in (Ramaswami and Sivarajan, 1996), including one which also attempts to create lightpaths between nodes in order of decreasing traffic demands. Lightpaths are established between the nodes that have the maximum amount of traffic between them. If all traffic is accounted for but each node does not have the required degree, the rest of the lightpaths are placed at random obeying the constraints. A modified version of this heuristic is also presented in which a pair of lightpaths in opposite directions is initially set up for each physical edge, then the original algorithm is exactly followed. This ensures that traffic can always be routed on the shortest physical path between any two nodes and hence can satisfy any physically realizable delay constraints. Another heuristic depends on the iterative bound developed in this study by relaxing the MILP formulation as described in Section 3.2, and rounding off the lightpath indicators. Finally, a heuristic is presented that does not take into account the traffic pattern at all, but concentrates on creating lightpaths that use only a few physical edges, since this should conserve wavelengths.

A similar heuristic maximizing one logical hop traffic is briefly described in (Banerjee and Mukherjee, 1997), but a heuristic with the opposite objective is also suggested. This heuristic aims at maximizing multihop traffic, since concentrating only on single hop traffic can lead to congestion due to multihop traffic. Some results are provided in which the two approaches appear to perform very similarly to each other. Details of wavelength assignment are not discussed.

The study in (Banerjee et al., 1997) also suggests that attempts to maximize one logical hop traffic concentrate on the comparatively larger traffic components, and may cause the smaller traffic components to be routed unreasonably and cause congestion on some physical links. A scheme involving mapping the network to a bipartite graph is specified to avoid unbalanced loading. A known graph algorithm is then specified as a heuristic for wavelength assignment.

In (Krishnaswamy and Sivarajan, 1998), a heuristic algorithm following the LP relaxation heuristic from (Ramaswami and Sivarajan, 1996), but more complete, is presented. The lightpath wavelength indicator variables are also rounded and then a least resistance algorithm followed to choose a single routing for each lightpath. A final phase of the design eliminates wavelength clashes.

4. RELATED APPROACHES

In this section we discuss some techniques and algorithms that are different from those described in Section 3.2, but which are related to the problem of virtual topology design for wavelength routed networks.

Incremental Benefit Analysis: In (Mukherjee et al., 1996), a study of the incremental benefits of introducing a virtual topology over optical WANs is undertaken. A realistic traffic pattern is obtained from the T1 NSFNET backbone data of January 1992. Three schemes are applied on this data to scale up the traffic pattern. The first scheme merely used efficient routing to establish a baseline and the other two used WDM and virtual topology with WDM respectively. The most dramatic result was in the increase of the scaleup factor, from 49 and 57 in the first two schemes to 106 in the third, and the link utilization which went from 32% and 23% in the minimum loaded link to 71% in the last one (the maximum link load remained 99% in all three schemes). Thus this analysis provides demonstration of the benefits of implementing a virtual topology, as well as the incremental nature in which it may be undertaken.

Limited Conversion: The motivation for the study presented in (Ramaswami and Sasaki, 1997) is the lower cost associated with limited conversion of wavelengths at Wavelength Routers as opposed to full conversion, as we remarked in Section 2. The wavelength assignment subproblem is the focus of this study. Several results are obtained in theoretical terms about ring networks with specific wavelength conversion capabilities. These results are followed by constructive proofs rather than simply existence proofs, so that a blueprint is provided for the actual construction of such ring networks. Some results are derived for more general physical network topologies.

Traffic Grooming: As we have remarked in Section 1, each lightpath has a high bandwidth and it may not be possible for single users to utilize this bandwidth. Lightpaths must be viewed as transport channels in the backbone network, in which traffic from multiple user applications is multiplexed in by access networks. In a sense, this is the justification for including the traffic routing subproblem in the virtual topology design problem, since traffic for individual applications must be routed onto the virtual topology provided, so that lightpaths carry traffic obtained by aggregating lower speed traffic streams. The pattern of multiplexing traffic onto lightpaths affects the efficiency of optical forwarding of information through Wavelength Routers, since all information in an entire lightpath will need to undergo optoelectronic conversion and electronic routing at an intermediate node if even one lower speed traffic stream from that lightpath has to be terminated at the intermediate node. This also reflects in the cost (in numbers and capabilities) of network components needed. (Gerstel et al., 1998) addresses some of these issues. Different ring architectures are specified and compared on the basis of results derived regarding the average number of transceivers at the nodes, number of wavelengths, average number of physical hops and characterization of traffic patterns on which they perform best. Similar issues arising in arbitrary physical topologies and extension to actual grooming methods would appear to be areas worth further investigation.

Generalized Lightpaths: In (Sahasrabuddhe and Mukherjee, 1999), the concept of a lightpath is generalized into that of a *lighttree*, which, like a lightpath, is a clear channel implemented with a single wavelength with a given source node. But unlike the lightpath, a lighttree has multiple destination nodes, thus a lighttree is a point-to-multipoint channel. It is emphasized that topologies using lighttrees would be more optimal for any given situation since the lighttree is a more general construct, and there may be possibilities of optical multicast.

As we already know, optimal solutions are not practically obtainable, and with a more general construct and hence a much larger search space this is going to be even more true. Heuristic solutions will have to be designed to obtain good solutions, and must be tailored to suit the larger search space. With unicast traffic problems, the lighttree approach trades off more bandwidth to further improve delay, congestion, and physical hop characteristics than the lightpath approach. This is the tradeoff we mentioned in Section 1. The challenge in this case will be to design heuristics that can cope with the increased complexity of the problem and yet produce solutions in which a good tradeoff is achieved.

5. RECONFIGURATION ISSUES

As we have already remarked, the problem of reconfiguring a network from one virtual topology to another is a related problem to virtual topology design. Two possible approaches to this problem are discussed in this section.

Cost Approach In this approach, it is assumed that the current virtual topology as well as the new virtual topology that the network must be reconfigured to are known, together with the physical topology details. The concern is to minimize the cost of the reconfiguration. The cost can be expressed in terms of the number of Wavelength Routers that need to have their optical switching reprogrammed, or the total number of optical switchings that need to be changed to implement the new lightpaths and eliminate old ones. These metrics are appropriate since they reflect the amount of time the network must be taken off line to make the changes, as well as the reprogramming effort for the reconfiguration. Other similar metrics may also be applicable. It may be the case that the network cannot be taken off line at all, but that a succession of intermediate virtual topologies have to be designed to eliminate single, or groups of, routers which can be reconfigured and put back in operation. Much more complicated metrics reflecting total time taken to reconfigure as well as the effort to redesign the intermediate topologies need to be developed in this case.

We have not found any study of these reconfiguration problems in the literature for wavelength routed WANs, though studies involving the reconfiguration of virtual topologies for broadcast LANs exist, as detailed in the survey of related literature carried out in (Labourdette, 1998). These studies involve link-exchange and branch-exchange techniques to minimize the cost of converting one virtual topology into another, and similar methods may be possible for wavelength routed network which are the topic of this survey.

Optimization Approach Another approach is to assume that only the current virtual topology is given, together with the changed traffic pattern and/or physical topology that makes reconfiguration necessary. This is the approach taken in (Banerjee and Mukherjee, 1997). The reconfiguration algorithm proposed involves solving the new virtual topology problem on its own without reference to the current virtual topology to obtain a new optimal solution, with a new optimal value for the objective function which is noted. The virtual topology design problem is then reformulated with an additional constraint that constrains the old objective function to this noted value, and a new objective function that involves minimizing the number of lightpaths that must be either added or removed.

While this method is guaranteed to find a solution that results in a virtual topology that is optimal for the new conditions, it does not achieve a balance between finding an optimal new virtual topology and one that involves as little change from the old one as possible. It is possible that a very costly reconfiguration will be undertaken for only a slight gain in network performance. More balanced formulations of this problem may be possible, and heuristics designed on such formulations are likely to perform better in practice.

6. CONCLUDING REMARKS

The problem of virtual topology design for wide area wavelength routed optical networks covers a considerable area, and many approaches to this and related problems have been taken in the literature. The lightpaths of a virtual topology are set up to trade off the ample bandwidth available in the fiber with the optoelectronic conversion and electronic processing at intermediate nodes. Exact formulations of the problem are known to be computationally intractable, so heuristics for determining and implementing a virtual topology have been proposed. Most heuristics attempt to address parts of the problem rather than the whole, by decomposing the problem approximately into subproblems, or address special cases of network topology.

Virtual topology design is a growing research area. New areas of investigation include extending results obtained for special cases to broader context, and extending the freedom allowed in formulating the problem to take advantage of improving equipment capabilities.

References

Banerjee, D. and Mukherjee, B. (1996). A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. *IEEE JSAC*, 14(5):903–908.

Banerjee, D. and Mukherjee, B. (1997). Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study. *Proc. IEEE INFOCOM*, pages 269–276.

Banerjee, S., Yoo, J., and Chen, C. (1997). Design of wavelength-routed optical networks for packet switched traffic. *Journal of Lightwave Technology*, 15(9):1636–1646.

Chen, C. and Banerjee, S. (1995). Optical switch configuration and lightpath assignment in wavelength routing multihop lightwave networks. *Proc. IEEE INFOCOM*, pages 1300–1307.

Chlamtac, I., Ganz, A., and Karmi, G. (1992). Lightpath communications: An approach to high bandwidth optical wans. *IEEE Trans. Communications*, 40(7):1171–1182.

Chlamtac, I., Ganz, A., and Karmi, G. (1993). Lightnets: Topologies for high-speed optical networks. *Journal of Lightwave Technology*, 11(5/6):951–961.

Gerstel, O., Ramaswami, R., and Sasaki, G. (1998). Cost effective traffic grooming in WDM rings. *Proc. IEEE INFOCOM*, pages 69-77.

Green, P. (1992). *Fiber optic network*. Englewood Cliffs, NJ Prentice Hall.

Green, P. (1996). Optical networking update. *IEEE JSAC*, 14(5):764–779.

Krishnaswamy, R. and Sivarajan, K. (1998). Design of logical topologies: a linear formulation for wavelength routed optical networks with no wavelength changers. *Proc. IEEE INFOCOM*, pages 919–927.

Labourdette, J.-F. P. (1998). Traffic optimization and reconfiguration management of multiwavelength multihop broadcast lightwave networks. *Computer Networks and ISDN Systems*, 30(9-10):981–998.

Marsan, M., Bianco, A., Leonardi, E., and Neri, F. (1993). Topologies for wavelength-routing all-optical networks. *IEEE/ACM Trans. Networking*, 1(5):534–546.

Mukherjee, B. (1997). *Optical communication networks.* McGraw-Hill.

Mukherjee, B., Ramamurthy, S., Banerjee, D., and Mukherjee, A. (1994). Some principles for designing a wide-area optical network. *Proc. IEEE INFOCOM*, pages 110–119.

Mukherjee, B., Ramamurthy, S., Banerjee, D., and Mukherjee, A. (1996). Some principles for designing a wide-area optical network. *IEEE/ACM Trans. Networking*, 4(5):684–696.

Ramaswami, R. and Sasaki, G. (1997). Multiwavelength optical networks with limited wavelength conversion. *Proc. IEEE INFOCOM*, pages 489–498.

Ramaswami, R. and Sivarajan, K. (1996). Design of logical topologies for wavelength-routed optical networks. *IEEE JSAC*, 14(5):840–851.

Ramaswami, R. and Sivarajan, K. (1998). *Optical networks: a practical perspective.* Morgan Kaufmann Publishers.

Sahasrabuddhe, L. and Mukherjee, B. (1999). Light-trees: Optical multicasting for improved performance in wavelength-routed networks. *IEEE Communications Magazine*, pages 67–73.

Zhang, Z. and Acampora, A. (1995). A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength re-use. *IEEE/ACM Trans. Networking*, 3(3):281–288.