A Survey of Virtual Topology Design Algorithms for Wavelength Routed Optical Networks

RUDRA DUTTA
GEORGE N. ROUSKAS
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534

ABSTRACT

In the past few years, there has been growing interest in wide area "All-Optical Networks" with wavelength division multiplexing (WDM), using wavelength routing. Due to the huge bandwidth inherent in optical fiber, and the use of WDM to match user and network bandwidths, the wavelength routing architecture is an attractive candidate for future backbone transport networks.

A virtual topology over a WDM WAN consists of clear channels between nodes called lightpaths, with traffic carried from source to destination without electronic switching "as far as possible", but some electronic switching may be performed. 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 survey we first describe the context and motivations of the virtual topology design problem. We provide a complete formulation of the problem, and describe and compare the formulations and theoretical results as well as algorithms, heuristics, and some results in the current literature in the field. The reconfigurability issue, which is another attractive characteristic of optical networks, is also discussed and the literature surveyed.

This survey is restricted to transport networks with wavelength routing. Similar virtual topology problems also arise in multihop broadcast local area optical networks, but this work does not directly apply to them and corresponding literature is not included in this survey. This survey also relates to the design of a static topology, not one in which individual lightpaths are set up and torn down in response to traffic demand.

1 INTRODUCTION

1.1 Context

In the past few years, there has been growing interest in wide area "All-Optical Networks" with wavelength division multiplexing (WDM), using wavelength routing. These 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 [8, 13, 18, 9].

Virtual topology design over a WDM WAN is intended to combine the best features of optics and electronics. This type of architecture has been called "almost-all-optical" because traffic is carried from source to destination without electronic switching "as far as possible", but some electronic switching may be performed. 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. This constraint can be relaxed by assuming the availability of wavelength converters at intermediate nodes. However, this involves not only expensive equipment but further complications relating to the tuning delay of converters and the issue of converter placement, and in this survey we treat the wavelength continuity constraint as part of the problem, for the most part. 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 consists of the virtual (otherwise called the logical) topology.

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

The use of WDM allows the utilization of the large bandwidth inherent in optical fiber. In some cases, the fiber has been used as a simple alternative to copper wire. This means that only a single wavelength is used to carry information over a fiber and the fiber then acts as a point-to-point link of a given bandwidth. With WDM, each wavelength can utilize bandwidths comparable to that which the entire fiber was providing. With the further use of wavelength routing or virtual topologies, the bandwidth available to traffic goes up further. Figure 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 that are not directly connected by a fiber link. Two lightpaths can share a physical link by using different wavelengths. 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 [15]. 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, though the set of all virtual topologies that can be set up is constrained by the physical topology. In setting up a virtual topology, the usual considerations are delay, throughput, equipment cost and reconfigurability.

The virtual topology design problem bears certain similarities to other topology design problems in networks. At a rudimentary level, routing problems on any packet-switched network are an attempt toward imposing a topology on the network. In fact, it is helpful to think of the virtual topology as being an optical layer which provides a topology distinct from the physical topology on which to route packet traffic [10,18]. This is appropriate because optical technology, allowing more capacity than electronic technology in the field of transmission but not nearly as developed in the information processing and storage fields, is better suited for application at the data transport level, that is, nearer to the physical layer. However, the virtual topology is designed with the express intention of

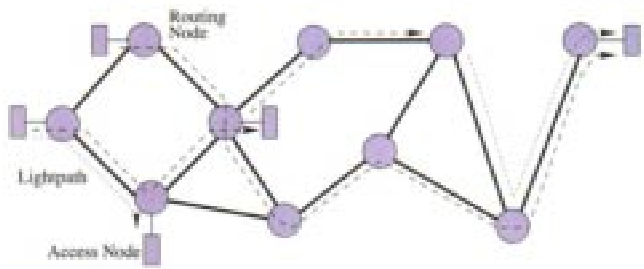

FIGURE 1: *A WDM network. The routing nodes are interconnected by point-to-point fiber links and may have access nodes connected to them. The dotted lines show lightpaths.*

managing network performance and improving congestion or throughput metrics, so traffic routing over the available topology has been considered as part of the virtual topology problem [14, 17, 15]. The Virtual Path (VP) allocation problem in ATM networks also concerns setting up a topology to be imposed on the physical topology available with a view to improving network performance metrics. With complete relaxation of the wavelength continuity constraint, the virtual topology design problem becomes more like, but not quite identical to, the ATM VP allocation problem.

Virtual topologies can also be designed on broadcast LAN (so-called multihop) lightwave networks. This problem is similar to the virtual topology problem for all-optical WANs, but with sufficient differences for it to be considered a different problem, as we remark in Section 1.2.

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 easily 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. The first is to decide what virtual topology to embed on a given physical topology, that is, what are the lightpaths to be implemented. Routing of these lightpaths on the physical topology and the assignment of wavelengths to them are the next two problems. The routing of packet traffic on the lightpaths is also usually seen to be a part of the virtual topology problem.

The above discussion focuses on the issue of throughput or delay optimization, which is 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 2.1.

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.

1.2 Scope

A similar virtual topology design problem exists for broadcast optical networks, used as LANs. In these networks, there is a single broadcast medium that is accessed by all nodes in the network. Lightpaths are set up by assigning a wavelength to a source-destination pair, which then acts as a clear channel between them. Traffic is sent from one node to another using a lightpath if one is available, or a sequence of lightpaths if a direct lightpath is not available. For this reason, these networks are called multihop networks. These networks are distinct from the wavelength routed WANs that are the subject of this survey. This survey does not deal with virtual topology problems in broadcast optical networks. One of the reasons the virtual topology problem is different in that case 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 [11].

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 or the average traffic flow 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 designing topologies and algorithms that will allow us to estimate and obtain optimum blocking probabilities under dynamic traffic demands, that is, calls which are established and terminated on demand [6]. This is not to say that traffic demands are assumed never to change in an actual network. However, such changes are not visible to a single instance of our static virtual topology design problem. If the traffic pattern changes significantly, it would act as the input data for a new virtual topology design, and the old virtual topology would be reconfigured to the new virtual topology, a topic we discuss in Section 5. Consequently, each virtual topology is designed only on the basis of a single average traffic demand pattern. Of course, it is possible that such a traffic pattern is itself made up of a combination of different estimates of network traffic, but this is not part of the virtual topology problem. Another way to make this distinction is to state that the virtual topologies we consider are sets of static lightpaths, not slowly varying lightpaths as would be the case if lightpaths were set up and torn down in response to user demands.

The virtual topology design encompasses only the transport network and not the access network. As we remarked above, the advantages of optical technology lie in switching and transmission, not processing or storage. Thus, electronic switching and transmission (or similar protocols over a physical fiber medium) 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 or backbone networks only, not access networks.

1.3 Structure of the Survey

This survey 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 survey.

2 Architecture and Notations

2.1 Network Components

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. The number of wavelengths that each fiber can carry simultaneously is limited by the physical characteristics of the fiber and the state of optical technology used to combine these wavelengths onto the fiber and isolate them off the fiber. This limit was of the order of 10 in past years and is currently of the order of 100 and growing. WDM has been seen as not only an obvious multiplexing method for the optical medium, but as a technique vital to utilizing the huge bandwidth of the fiber medium, because it can be used to correct the mismatch between the bandwidth available in the fiber and the bandwidth requirement of end users [14].

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 [7]. An ADM passes traffic on certain wavelengths through without interruption or optoelectronic conversions (conversions to electronic form and back to optical form), while traffic on other wavelengths is terminated optically, that is, converted to electronic form (the wavelength is dropped). Some wavelengths can also be added, that is, traffic is injected at this node using those wavelengths.

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 [6, 17]. 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, as shown in Figure 2. Thus a WR can cross-connect the different wavelengths from the input to the output, where the connection pattern of each wavelength is independent of the others. For this reason, it is sometimes also called a wavelength cross-connect. This description highlights the routing

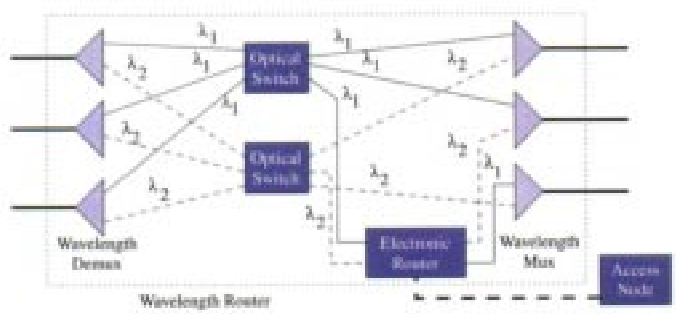

Figure 2: *A wavelength router.*

function of the WR. Of course, some of the wavelengths on some of the input ports may be carrying signals that are destined for an access node directly connected to the WR in question. In this case, that signal has to be extracted from the optical medium at this WR. To do this, it is necessary to terminate that particular wavelength, convert the data into electronic form, and deliver it to a higher layer. There may also be signals on some wavelengths that need to be forwarded to other nodes on a different wavelength. The wavelength has to be terminated in this case as well, the signal extracted to electronic form, converted back to optical form in the other wavelength and injected to an output port. To highlight the fact that the WR does both forwarding entirely in the optical domain (optical switching) as well as forwarding via conversion to electronic form and back to optical more like conventional routing), they are sometimes also called Wavelength Routing Switches (WRS) [14, 15].

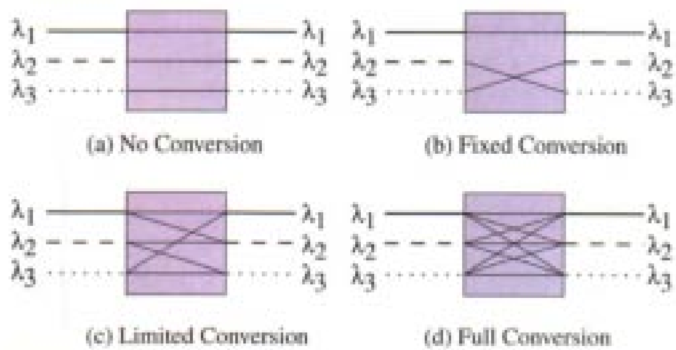

FIGURE 3: *Wavelength conversion.*

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 [16]. Without wavelength conversion, an incoming signal from port p_i on, say, the wavelength λ_1 can be optically switched (without intermediate optoelectronic conversions) to any port p_j, but only on the wavelength λ_1. With wavelength conversion capability, this signal could be optically switched to any port p_j on any wavelength λ_k. That is, 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. Figure 3 illustrates the differences for a single input and single output port situation; the case for multiple ports is more complicated but similar. 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. A special case of this gives us fixed wavelength conversion, where each input wavelength is converted to exactly one wavelength. If each wavelength is "converted" only to itself, then we have no conversion. If a node has limited or full wavelength conversion capability, then the conversion to be affected 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 increase can be minimized 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 related field to the virtual topology design problem. The primary goal of such studies is to provide an idea of the comparative impacts 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; see [2] for such a study. Guidelines that result from such studies may relate to choosing some initial parameters for the virtual topology, as suggested in [2], or may be integrated into the optimization procedure to find the virtual topology. The latter approach is taken in [4], 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.

2.2 Notations

In this section, we define some terminology and notations and introduce some concepts that will be used in the following sections, and that 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 that is usually the fiber distance or propagation delay over the corresponding fiber.

Lightpath: A lightpath, as we remarked above, 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 that 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. Usually 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 that 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 effected by the consideration of what volume of electronic switching can be done at a node [17].

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 that 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 x 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 $\lambda_{max} = \max_{i,j} \lambda_{ij}$.

2.3 Architecture

In this section we characterize in more detail the WDM wavelength routed network we have been describing above, and which Figure 1 illustrates. The network consists of several routing nodes that are connected to each other by point-to-point optical fibers. The nodes and their physical connections are 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. Every physical link carries at most one channel (lightpath) in each direction on each of the W wavelengths.

Lightpaths are set up on the physical topology, creating the virtual topology. Each arc of the virtual topology graph is a lightpath. 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. Thus at the intermediate nodes the traffic is not converted to electronic form, and the lightpath acts as a single-hop path from source to destination, or a pipe, with no queuing delay. Two lightpaths that share a physical link must be assigned different wavelengths. The total number of wavelengths used on a certain link must be W or less. The logical in-degree and out-degree are usually equal for each node. The logical degree of every node is usually assumed to be the same and this is called the logical degree of the network.

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 [17].

The goal of the virtual topology design process is usually to optimize network performance, such as by minimizing network congestion or minimizing 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 that bear little resemblance to the physical topology, with convoluted lightpaths that increase delay [17].

3 NETWORK 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 [10], and also those in [17, 14, 15]. The symbols and terminology are as defined in Section 2.2. 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 that 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:

Minimize the congestion of the network, that is,

$$\min \lambda_{max} \tag{1}$$

Subject to:

Degree Constraints

$$\sum_j b_{ij} \le \Delta_l, \forall i \tag{2}$$

$$\sum_j b_{ji} \le \Delta_l, \forall i \tag{3}$$

Traffic Constraints

$$\lambda_{ij} \le \lambda_{max}, \forall (i,j) \tag{4}$$

$$\lambda_{ij} = \sum_{sd} \lambda_{ij}^{(sd)}, \forall (i,j) \tag{5}$$

$$\lambda_{ij}^{(sd)} \le b_{ij} \lambda^{(sd)}, \forall (i,j),(s,d) \tag{6}$$

$$\sum_j \lambda_{ij}^{(sd)} - \sum_j \lambda_{ji}^{(sd)} = \begin{cases} \lambda^{(sd)}, s = i \\ -\lambda^{(sd)}, d = i \\ 0, s \ne i, d \ne i \end{cases} \forall (s,d) i \tag{7}$$

Wavelength Constraints

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

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

$$\sum_{ij} c_{ij}^{(k)}(l,m) \le 1, \forall (l,m),k \tag{10}$$

$$\sum_{k=0}^{W-1} \sum_l c_{ij}^{(k)}(l,m) p_{ml} - \sum_{k=0}^{W-1} \sum_l c_{ij}^{(k)}(l,m) p_{ml} = \begin{cases} b_{ij}, m=j \\ -b_{ij}, m=i \\ 0, m \ne i, m \ne j \end{cases} \forall (i,j) m \tag{11}$$

Hop Constraints

$$\sum_{lm} c_{ij}^{(k)}(l,m) \le h_{ij}, \forall (i,j),k \tag{12}$$

Discussion

The degree constraints (2) and (3) constrain the virtual topology to a given logical degree. Among the traffic constraints, (4) defines the network congestion. Expression (5) asserts that the total traffic on a lightpath is the sum of the traffic components on that lightpath due to all the different pairs of source and destination nodes. Constraint (6) captures the fact that the component of traffic on a lightpath due to a particular source-destination pair can be present only if the lightpath exists in the virtual topology, and cannot be more than the total traffic for that source-destination pair. Constraint (7) is an expression of the conservation of traffic flow at lightpath endpoints. All but one of the remaining constraints relate to the allocation of wavelengths to lightpaths. Constraint (8) ensures that a lightpath, if it exists in the virtual topology, has a unique wavelength out of the available ones. Constraint (9) enforces the consistency of the lightpath wavelength indicators and the link-lightpath wavelength indicators, and expression (10) enforces that a wavelength can be used at most once in every physical link, avoiding a wavelength clash. Expression (11) asserts the conservation of every wavelength at every physical link endpoint for each lightpath. The last remaining constraint, expression (12), enforces the bounds on the number of physical hops each lightpath is allowed.

The parameters, or inputs, to the formulation are the traffic matrix Λ, the hop bound matrix H, the number of wavelengths supported by 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 used to implement 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.

Formulations of this problem are possible that address only some and not all of these aspects. In Section 3.2.3 we discuss such approaches. Even when all these aspects are addressed, or the same aspect is addressed, different formulations of the problem are possible.

The exact formulation in [14] seeks to minimize the average message delay. For this purpose fiber distance metrics are introduced as parameters in the formulation. Throughput parameters are not considered, and the traffic parameters are used only to weigh the delay variables. Multiple and single logical hop paths are explicitly distinguished in the formulation. Unlike the above formulation, in [14], as well as [10, 15], the given logical in-degree and out-degree of each node (the number of receivers and transmitters available at that node) is considered to be a separate parameter, not a single parameter for the network. The formulation in [14] does not address the wavelength assignment issue, and there is no constraint corresponding to the physical hop constraint on the lightpaths, as given by (12) in the above formulation.

In [17], the bound on the number of wavelengths a single fiber can carry is ignored in the exact formulation. The wavelength assignment problem is also not addressed. The physical hop bound is not used to constrain the virtual topology; instead there is an average delay constraint on each source-destination pair. The maximum propagation delay between any source-destination pair in the physical topology is denoted by $d^{(max)}$ and introduced in the formulation, together with a tuning factor α that determines how tightly the virtual topology should be constrained by $d^{(max)}$. The average packet delay for each source-destination pair is constrained to be less than or equal to the product of α and $d^{(max)}$. Thus the delay performance metric is addressed using a constraint such as the following:

Delay Constraints

$$\sum_{ij} \lambda_{ij}^{(sd)} \leq \lambda^{(sd)} \alpha d^{(max)}, \forall s,d \quad (13)$$

The throughput performance metric is addressed by the goal of the optimization, which is to minimize congestion as in the formulation provided above.

The exact formulation in [15] follows the exact formulation in [14], but adds some elements. The capacity of each lightpath C is introduced as a parameter and it is used to refine the delay minimization goal and to set up an alternative optimization goal of maximizing the offered load to the network. The second goal is nonlinear in nature. The wavelength assignment problem is made part of the exact formulation.

In [2], the formulation is similar to the one we have provided above, but a simplification is introduced as part of the formulation. The goal of the optimization is to minimize the average physical hop lengths of the lightpaths in the virtual topology. The wavelength continuity constraint is not included. The lightpath distance is constrained to be bounded by the fiber distance of the shortest path in the physical topology between the endpoints of that lightpath, together with a factor α introduced as a parameter. The simplification consists of pruning the search space based on the fiber length of paths: lightpaths are only allowed to use physical paths that are among the K shortest such paths between any pair of nodes, K being a parameter in the problem formulation.

Some common features of the exact formulations from various studies in the literature are also shared by the formulation provided above. Bifurcation of traffic for a source-destination pair over different virtual paths is allowed. In view of the predominance of propagation delay over queueing delay in high speed wide area networks, the queueing delay is neglected in almost every formulation of the problem. A special parameter β is introduced in [2] to bound the congestion to values at which it is reasonable to neglect queueing delays. The formulation provided above also has a specific feature not shared by all the others in that it does not allow for more than one lightpath from one node to another.

As we discuss in Section 3.2, this formulation gets quickly intractable with size. 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 lower; otherwise the formulation is similar. Of course, the solution to the problem so formulated does not provide a complete solution to the full problem in terms of routing traffic over the virtual topology designed; moreover there may not be a feasible solution to the original problem corresponding to a solution for the aggregate problem. 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 [17, 10]. Bounds that can be calculated with significantly lower computational costs than solving the full problem are useful in evaluating heuristics employed to obtain good solutions to the full problem, 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, and hence is suitable for deriving achievability bounds. When the MILP is relaxed, an extra "cutting plane" constraint is introduced [17, 10], 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

An exact formulation of this problem such as the one given in Section 3.1 quickly grows intractable with increasing size of the network. In fact, this problem and some of its subproblems are known to be NP-hard [5, 14, 10, 1]. 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.

3.2.1 Subproblems

The full virtual topology design problem can be approximately decomposed into four subproblems. The decomposition is approximate or inexact in the sense that solving the subproblems in sequence and combining the solutions may not result in the optimal solution for the fully integrated problem, or some later subproblem may have no solution given the solution obtained for an earlier subproblem, so no solution to the original problem may be obtained. Although this decomposition follows [15], it is also consistent with the decompositions of [17, 14, 10, 1]. The subproblems are as follows.

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.

Lightpath Routing Subproblem: Determine the physical links that each lightpath consists of, that is, route the lightpaths over the physical topology.

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.

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 fixing the values of the variables $c_{ij}^{(k)}$ (l, m) would entail specifying wavelengths for each lightpath, that is, solving the wavelength assignment subproblem as well. The lightpath routing subproblem actually consists only of determining values of link-lightpath indicators (not included in the above formulation) and does not refer to wavelength assignment in any way.

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 and hence attempting to solve the subproblems and combine the solutions may result in a suboptimal solution or even no solution to the full problem. 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 that is commonly relaxed is that of the maximum number of wavelengths that can be carried by a fiber. Of course, if after the solution is obtained we find that some of the relaxed constraints have been violated (for example, if we have assigned a link to carry more lightpaths than the maximum number of wavelengths it can carry, or the total number of unique wavelengths used in the virtual topology is more than the maximum number of wavelengths a fiber can carry), then we would need to abandon that solution and search for a different one, possibly after modifying the heuristics used.

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.

3.2.2 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 (so called because 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 [17], and also that of [10], as well as literature on virtual topology problems in broadcast LAN scenarios as referred to in [17, 11]. 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 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

$$\lambda_{max} \geq rH'/E_l \tag{14}$$

Thus, setting a lower bound on H' results in a lower bound on the congestion. This is done from the following consideration. For N nodes in the network and a logical degree Δ_l, the maximum number of source-destination node pairs that can be connected by only a single hop is $N\Delta_l$. Similarly the maximum number of source-destination pairs that can be connected by only two hops is $N\Delta_l^2$, by only three hops is $N\Delta_l^3$, and so on. 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. Accordingly, assume that the $N\Delta_l$ source-destination pairs with the largest traffic between them are each connected by a single hop path (that is, there is a lightpath between each source and the corresponding destination). The $N\Delta_l^2$ source-destination pairs with the next largest traffic should be connected by two hop paths, and so on. The traffic weighted average number of logical hops in this case is a lower bound, in other words

$$H' \geq \sum_k kS_k \tag{15}$$

where S_k is the sum of the traffic fractions (with respect to total network traffic r) which consist of the k-th block when the traffic fractions are arranged in descending order of magnitude, and the i-th block is made up of $N\Delta_l^i$ successive elements in that list.

Minimum flow tree bound: This bound is derived from similar considerations as above, but on the basis of each source node rather than the network as a whole. In the above we assumed that the $N\Delta_l$ source-destination pairs are all connected by single hop paths, and so on, but this is impossible if the $\Delta_l + 1$ top traffic components all have the same source node, for example. 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 calculation of this bound is done by assuming that for each source, the source is connected by one logical hop to the Δ_l destinations to which it has the largest amounts of traffic, by two hops to the Δ_l^2 destinations to which it has the next largest amounts of traffic, and so on. We then form the sum of these traffic components weighted by the appropriate number of hops, as in the above scheme, and obtain the bound for the traffic weighted average number of logical hops H' similar to (15). The bound on congestion is then obtained from (14), as before. We omit the derivation and exact expression of this bound, which can be found in [17].

Iterative bound: This type of bound is developed in [17, 10] 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 of the bound, 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 improved bound on the congestion. This iterative process can be carried out repeatedly to improve the tightness of the bound. It is remarked in [17] that about 25 iterations result in a bound that is improved very little by further iterations.

Independent topologies bound: This bound is proposed in [3] as another method of taking into account the physical topology in computing a bound on the congestion. This bound is also based on relaxing the MILP formulation to obtain a linear problem. First a logical feasible virtual topology is obtained that maximizes the one-hop traffic. In the next stage, the traffic carried by this topology is eliminated from the total traffic and another virtual topology is designed that carries the maximum possible two-hop traffic out of the residual traffic. This procedure is repeated for successive number of hops until all the traffic is routed. The value of congestion can now be extracted from the several topologies. If this procedure were carried out exactly, it would be an exact solution and would not be more easily computable than a solution to the problem itself. However, the several logical topologies are not allowed to constrain each other, so that the topologies for more than one hop are not necessarily feasible. Thus the topologies are mutually independent. The authors of [3] 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 wavelengths 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, following [17], are discussed below.

Physical topology degree bound: This bound is derived from the simple consideration that each node in the virtual topology must source a number of lightpaths equal to the logical degree of the virtual topology Δ_l. Considering the node with the minimum physical degree δ_p in the physical topology, there must be sufficient number of wavelengths to allow Δ_l lightpaths to be realized over δ_p physical links, that is, the number of wavelengths required is bounded from below by $W \geq (\Delta_l / \delta_p)$.

Physical topology links bound: Another consideration is that each physical link is traversed in general by multiple lightpaths in each direction. If we count each physical link once every time a lightpath traverses it, then the total will be greater than the number of directed physical links in the topology (we multiply the number of links in the undirected physical topology by two to get this latter number). This is possible because different lightpaths on the same physical link (in the same direction) use different wavelengths. Thus the average number of lightpaths traversing a physical link can be obtained by dividing the total number of physical links traversed by all the lightpaths by the number of directed physical links in the topology, and this is a bound on the number of wavelengths required.

For the number of physical links traversed by all the lightpaths, we use the following argument. For each source node, reaching every other node requires a minimum number of physical hops, using the shortest physical path. Let us list the $N - 1$ possible destination nodes for a source node n_i in increasing order of the number of physical hops needed to reach that destination from n_i, and then assume that the Δ_l lightpaths sourced from n_i traverse the Δ_l physical paths that lead to the first Δ_l destinations on that list. We assume this is true for each source node n_i . This assumption results in a lower bound on the total number of physical links traversed, and hence on the number of wavelengths. We omit the derivation and exact expression of this bound, which can be found in [17].

Another bound, which is derived with respect to the lightpath routing and wavelength assignment subproblems and not the complete problem, may nevertheless be useful in some cases. This is the NWC bound [5] (for Non-Wavelength Continuous), which states that given a physical topology and a set of lightpaths to be established, the number of wavelengths needed to establish a virtual topology obeying the wavelength continuity constraint is not less than the number needed to establish the same virtual topology without the wavelength continuity constraint. In general, it is not clear how tight this bound is, and it is not easily computable. However, it can be shown that for networks with topologies that are acyclic, this bound is not only tight but exact [5]. Thus it could be useful if after finding a solution to the lightpath routing subproblem, we find an acyclic topology and want a bound to evaluate a heuristic solution to the wavelength assignment subproblem.

Bounds on the number of wavelengths required can also be found under specific assumptions regarding the solution to the different subproblems. For example, in [6] it is demonstrated that if the virtual topology being implemented is decided to be a hypercube (as part of solving the topology subproblem), then for a specific proposed algorithm to map the hypercube nodes to physical network nodes (solving the rest of the topology subproblem), the number of wavelengths required cannot be less than $2n/3$, where n is the number of nodes in the physical network. A similar result for a torus embedding, as well as a general result in terms of the physical topology and the topology being embedded, is also provided in [6].

3.2.3 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 on 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 topology may not be chosen with complete freedom; instead it must obey the constraints of the regular topology. For example, the number of nodes in a torus must obey $n = m^2$ for some m. 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 [14]. 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, although this introduces further approximations to the virtual topology solutions. In general, the node mapping and path mapping problems leave out of consideration the traffic patterns in the network, and utilize metrics such as fiber distances and wavelength reuse to route lightpaths over the physical topology. Thus there is a tacit assumption of reasonably 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 [6, 14], hence heuristics are needed for them. Below we discuss some approaches taken in the literature with regular topologies as candidates for virtual topologies.

Embedding via strings: In [6], the authors propose a two-phase heuristic for performing the mappings. In the first phase, an "equivalent" string is obtained for the physical network. In the second phase, the selected regular topology is embedded into the string. The elements listed in the string represent the nodes of the physical topology. Each edge connecting two successive elements of the string corresponds to a path between the corresponding nodes. The string also has the property that any two paths that are edge disjoint in the string are also edge disjoint in the physical topology, so that path mapping and wavelength allocation on the string can be translated in a straightforward manner to the physical topology.

Three methods of obtaining such a string representation from a physical topology are discussed in [6]. The first two involve finding a Hamiltonian path and an Eulerian path in the physical topology. However, there is no guarantee that such paths can be found in an arbitrary physical topology, and hence these methods are not completely general, though they have other attractive characteristics. The third method involves finding a spanning tree in the physical topology, which only requires that the topology graph be connected. Algorithms for embedding a torus topology and a hypercube topology into these string representations are then presented. A general result regarding a lower bound on the number of wavelengths needed to embed an arbitrary topology G_e on another topology G_g is then derived in terms of characteristics of G_e and G_g, and this result is used to derive lower bounds for the torus and the hypercube cases. The resulting virtual topologies are compared in terms of logical degree of the network, average number of logical hops, and the number of wavelengths needed. The hypercube embedding is seen to have the better (i.e., smaller) number of logical hops, but it requires a larger network degree and more wavelengths. It appears that algorithms can be similarly developed for embedding other regular topologies in the string representations.

The approximation in this approach is introduced at the time of obtaining a string representation. The string obtained is constrained by only the criteria mentioned above, and this does not guarantee that the particular string we obtain for a given instance will lead us to an optimal or near-optimal virtual topology. However, given the regular topology to be embedded and the embedding function, it is shown in [6] that the particular algorithms presented are near-optimal. In this approach, the traffic pattern in the network is ignored.

Comparison of topologies: Three different regular topologies are compared for their suitability as virtual topologies in [12]. The comparison is based on the number of logical hops between nodes and the number of wavelengths required. No traffic pattern considerations are included. In fact this comparison does not refer to a physical network at all. The regular topologies are looked upon only as candidates for embedding into physical topologies, and it is assumed that the required node mapping and path mapping may be carried out. In this sense this study is from the perspective of the first part of the virtual topology subproblem, in that the virtual topology is proposed but node mappings are not performed. However, this study addresses the issue of wavelength routing given the virtual topology, which provides the path mapping.

The three topologies compared are the *K*-grid topology (an extension to *K* dimensions of the Manhattan Street topology), the twin shuffle topology (a shufflenet with twice the connectivity, that is, two shufflenets in parallel), and a modified de Bruijn graph (de Bruijn graphs with connectivity multiplied by a factor). The basis of comparison is the degree of the network. However, this is expressed as the number of bundles of fiber leaving each node, where each bundle carries fibers to exactly one other node but may contain more than one fiber. Thus, the allowed number of wavelengths along a connection between two nodes would be larger with this scheme than with a single fiber scheme. This difference is expressed by a difference between the number of wavelengths required to implement the virtual topology and a normalized version of the same number.

The authors point out that the de Bruijn graph topology has the advantage of a constant number of hops between nodes, and this number is smaller than either the average or maximum number of hops for the other two topologies. The number of wavelengths required appear to be similar for the three topologies.

Simulated Annealing: In [14], 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. Moreover, the nature of the physical topology is taken into account to some extent. The goal of each heuristic is to minimize overall average message delay in the virtual topology, which involves the fiber distances and the traffic pattern in the network. Two heuristic approaches are developed, one based on a greedy algorithm and the other on simulated annealing.

In both approaches, fictitious nodes are created in the physical topology if it has fewer nodes than the chosen hypercube. The greedy approach starts with a reasonable initial node mapping obtained by mapping the nodes with the highest physical degrees first and attempting to map nodes so that logical nodes that are neighbors in the logical topology are mapped to physical nodes that are neighbors in the physical topology as far as possible. This node mapping is then refined by traversing the list of nodes in reverse order of initial embedding and swapping each with a different node if this would reduce the average message delay. Once the final node mapping is obtained, the path mapping is obtained by shortest path routing on the physical topology to realize each edge of the hypercube, that is, lightpath. A remark is made regarding the traffic routing subproblem that shortest path routing on the hypercube virtual topology has been assumed.

To assign wavelengths, physical links are ordered in decreasing order of number of lightpaths passing through them. For each physical link, each lightpath is assigned a different wavelength if it has not already been assigned one through some other physical link. Wavelength conflicts are avoided while assigning wavelengths. This algorithm is heuristic in nature and in particular is not guaranteed to find a solution if the number of wavelengths is bounded, even if a solution exists. With unbounded number of wavelengths, this algorithm does not impose an obvious non-trivial bound on the number of wavelengths it uses, and such a bound is not discussed in [14].

In the simulated annealing heuristic, the node mapping process starts with an initial random mapping. The perturbation is provided by swapping the mapping of two nodes in the virtual topology. The acceptance criterion is the average message delay in the network; the new mapping is always accepted if the resulting delay decreases, and is accepted if the delay increases with a probability that goes down as the simulated temperature falls. Once the mapping "freezes" the path mapping and wavelength assignment are carried out as for the greedy algorithm.

In [15], 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 [15] and also [11]. 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: In this section we discuss studies that 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, as we have seen, can be used to determine what lightpaths to set up. The traffic pattern can also be useful in routing the lightpaths and the traffic, because the network performance metrics which are the goals of the virtual topology design are related to traffic, such as message delay or congestion. Thus the traffic pattern may be utilized even in an approach focused on the lightpath routing subproblem only. However, in some approaches, it is considered that the traffic pattern has been properly taken into consideration while solving the virtual topology subproblem that resulted in the given set of lightpaths to be implemented, and thus the lightpaths take care of the traffic characteristics of the network. 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.

SLE: In [5], not only the source and destination, but also the routing of the lightpaths, are also assumed to be given, and the problem is seen to be the assignment of wavelengths to these lightpaths. That is, the lightpath wavelength assignment subproblem is addressed, and it is called the Static Lightpath Establishment (SLE) problem. SLE as posed includes a bound on the number of wavelengths that can be used to establish the lightpaths. It is proved that SLE as stated is equivalent to the n-graph-colorability problem, and hence is NP-complete.

A heuristic algorithm to assign wavelengths to a given set of lightpaths with the aim of using as few wavelengths as possible is presented. This algorithm is based on a greedy allocation heuristic that iteratively assigns a wavelength to as many edge disjoint lightpaths as possible before going on to the next wavelength. Longer lightpaths are allocated a new wavelength earlier. The algorithm terminates when all lightpaths have been assigned a wavelength. There is no obvious non-trivial bound to the number of wavelengths needed by this algorithm, and none is discussed in [5]. A modified version of this algorithm that allocates wavelengths only until a given maximum number of wavelengths have been allocated, and then stops, is also given. The use of such an algorithm is in deriving values of blocking probabilities than can be compared with values encountered in the dynamic case.

The problem of dynamic lightpath establishment, in which lightpaths are set up and torn down on demand, and the goal is to provide the minimum blocking probability seen by new lightpath demands, is also discussed in [5]. However, this is outside the scope of this survey.

Wavelength utilization: The study presented in [4] 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. Thus the lightpath routing and wavelength assignment subproblems are addressed, and an integrated approach is taken for these two subproblems. It is assumed that traffic related objectives have been addressed while obtaining the set of lightpaths to establish.

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 (WR) at each network node. It is assumed that enough distinct wavelengths are available so that every WR can switch a lightpath from each of its input ports to each of its output ports. If each node has a physical degree of N, then at least N wavelengths are needed to switch these N^2 lightpaths at the WR. However, with certain wavelength allocations to some of the lightpaths, it may not be possible to switch every wavelength at every input port to some output port, while a different wavelength assignment would allow all the lightpaths to be set up. Thus wavelength utilization is defined for a WR in terms of the number of input port wavelengths that are not "blocked" as described above.

To formally define the problem, the concept of a "Latin Square" is introduced. An $N \times N$ Latin Square is filled with N distinct elements such that no two elements in the same row or column are the same. This is seen to correspond to a wavelength assignment at a node such that wavelength utilization is maximum. A partial Latin Square is one in which not all entries are present but those that are present obey the constraint above. This represents a node at which some lightpaths have already been assigned wavelengths. Thus for achieving maximum wavelength utilization, a routing and wavelength assignment algorithm would have to ensure that the solution results in a wavelength assignment at each node that is as close as possible to a full Latin Square. The algorithm can start with a partial Latin Square at each network node and proceed to fill them by routing and assigning wavelengths to lightpaths such that this goal is achieved. The initial partial Latin Square at each node will normally be completely empty, or some entries may be prefilled if some of the lightpath routing or wavelength assignments are constrained by some factor outside the scope of this problem. It is remarked that not all partial Latin Squares can be completed into full Latin Squares, and deciding whether a given partial Latin Square can be completed is an NP-complete problem.

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. The first algorithm is based on backtracking. To reduce the number of backtracking steps, the concept of a degree of freedom is introduced for each entry in a partial Latin Square. For an empty entry, the degree of freedom is the number of different values that can be assigned to that entry that would still result in a partial Latin Square, while the degree of freedom is zero for an entry that already has a value assigned to it. The algorithm iteratively assigns values to empty entries, each time assigning a value to the empty entry with the least degree of freedom, and assigning the particular value that would result in the minimum reduction of the total degree of freedom of other entries in the same row or the same column. If an empty entry is seen to have a degree of freedom of zero, the algorithm backtracks to the last entry and assigns a different value. The algorithm terminates when all entries are assigned values or it is known that the square cannot be completed. At worst this corresponds to an exhaustive search. The second algorithm is based on converting the problem into an edge coloring problem in a bipartite graph, but with this algorithm the solution may violate the wavelength continuity constraint; that is, the WR must have wavelength conversion capability.

When routing a lightpath over the network nodes and assigning a wavelength, there may be no routing that allows the lightpath to occupy the entry with the minimum degree of freedom at each intermediate node. Thus the scheme for the overall problem at the network level involves a search for the k best shortest routes for each lightpath, then picking the one with the minimum total degree of freedom over each intermediate node. Similarly the wavelength is assigned by choosing a value that results in the minimum reduction of the total degree of freedom over each intermediate node. A remark is made to the effect that the goal is to maximize the traffic carried in one logical hop. No further detail is provided as to how to successively choose lightpaths for routing and wavelength assignment.

Randomized rounding and graph coloring: In [1], 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 and wavelength assignment subproblems are addressed. The traffic pattern in the network is not considered, and it may be assumed that this was taken into consideration when obtaining the set of lightpaths, so that each lightpath carries traffic nearly to its capacity. Thus lightpaths themselves are used as units of traffic and congestion.

The lightpath routing problem is formulated in terms of lightpath traffic as a multicommodity flow problem that 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, in terms only of the lightpaths that are given, and also by pruning the search tree by assuming that the optimal routing of a lightpath can always be found among a few alternate shortest path routing of the lightpath on the physical topology. The integer constraints of the formulation may also be relaxed.

The technique of randomized rounding is used as a heuristic algorithm to determine lightpath routing. The goal is to minimize the number of wavelengths needed to establish all the lightpaths. First, the integer constraints on the flows representing lightpaths are relaxed, creating a non-integral multicommodity flow problem, and this is solved by some linear programming method. Then, a phase called path stripping is carried out, in which a set of possible paths is created for each lightpath. Successive paths for the commodity representing that lightpath are found from among the links that carry any part of the flow of that commodity in the solution to the relaxed flow problem. Each path is given a weight equal to the minimum fraction of the commodity carried by a link in the path, and this minimum value is subtracted from each link participating in the path. Once the set of paths has been obtained for all the lightpaths, a single path is chosen for each lightpath randomly, using the weights assigned during path stripping.

The wavelength assignment is presented as a separate subproblem once the lightpath routing has been carried out. A transformation of this problem to a graph coloring problem which is known to be NP-complete is specified. An efficient well-known sequential graph coloring algorithm called smallest-last coloring is specified as the chosen method due to its simplicity.

The study employs known heuristic methods with provably good characteristics to address the lightpath routing and wavelength assignment problems. It may be noted that [1] demonstrates the use of the algorithms specified for dynamic as well as static lightpath establishment.

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 this section we discuss such heuristic approaches.

In [20], 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. The algorithm iteratively attempts to create as many lightpaths as possible using each wavelength without violating the wavelength clash and continuity constraints. Lightpaths are assigned between source-destination node pairs in descending order of the amount of average traffic flowing between them, which favors one logical hop traffic. By assigning a wavelength to as many lightpaths as possible before going on to the next wavelength, the attempt is made to utilize the number of wavelengths used for a maximum number of lightpaths. Only a predefined number of wavelengths can be used, and the algorithm terminates when no more lightpaths can be set up using these wavelengths. The authors remark that there is no guarantee that the virtual topology obtained would be connected. To allow the routing of traffic from any node to any other on this topology, the suggestion is made that the procedure be stopped when one wavelength still remains, and then this wavelength be used to connect any disconnected subnetworks that may have been formed before using it to form any other lightpaths that may be possible. The study goes on to describe a routing scheme that dynamically allocates and deallocates these lightpaths on demand, but this is outside the scope of this survey.

Several different heuristics are presented in [17]. The first one is simply called "heuristic logical topology design algorithm", and it also attempts to create lightpaths between nodes in order of decreasing traffic demands. A network degree is assumed to be given as part of the problem. Each lightpath is established between the nodes that have the maximum amount of traffic between them that is not already carried by some lightpath, provided wavelength clash, continuity, and degree constraints are obeyed. 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 that is only applicable if the logical degree is greater than the physical degree. In this algorithm, 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 3.2.2. The higher values of the lightpath indicator variables are taken to represent actual lightpaths and the lower values are discarded, obeying the degree constraints, yielding a virtual topology. 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. Thus the heuristic first creates lightpaths between all nodes that are one physical hop apart, then between all nodes that are two physical hops apart, and so on, while the degree constraints are not violated. Wavelength assignment algorithms for the last two heuristics are not discussed, and lightpath routing for the LP relaxation algorithm is not discussed.

A similar heuristic maximizing one logical hop traffic is briefly described in [2], but a heuristic with the opposite objective is also suggested. Since in a virtual topology some traffic will always be carried in multiple logical hops because of constraints on number of wavelengths and node degrees, a heuristic approach must account for multihop traffic and not only concentrate on maximizing single hop traffic. Accordingly, this heuristic aims at maximizing 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 [3] 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 is presented to avoid this. A complete bipartite graph is created in which each partition contains all nodes of the physical topology. The edges are weighted with traffic demands between corresponding nodes. Each edge represents the shortest physical path between corresponding nodes. Now a minimum weighted perfect matching is identified and the corresponding edges are eliminated from this graph. The traffic carried by the eliminated edges is rerouted over the least congested of the remaining edges. This is repeated until the number of edges connected to each node has been reduced from N to Δ_l, the desired logical degree. Now this graph represents a logical topology with each edge representing a lightpath. The lightpath routing and wavelength assignment can be done arbitrarily and the congestion will be the same, but a different number of wavelengths will be required. Since minimizing the number of wavelength is known to be NP-hard, a known path-graph based algorithm is specified for path embedding and wavelength assignment. This algorithm is described in literature referred to in [3]. If the number of wavelengths is more than the number desired, extra wavelengths may be eliminated by choosing wavelengths that implement the least number of lightpaths, and rerouting traffic carried by those lightpaths along others with least load.

In [10], a heuristic algorithm following the LP relaxation heuristic from [17], but more complete, is presented. After rounding the lightpath indicator variables obtained by a suitable number of iterations of the LP obtained by relaxing the exact formulation, the virtual topology subproblem has been solved. The constraint on the number of physical hops for lightpaths, which is a part of the exact formulation provided in [10], may be lost in the relaxation, in the sense that some lightpaths may be formed with larger number of hops than was allowed for in the exact formulation. Now the lightpath wavelength indicator variables are also rounded using a specified rounding algorithm that sets the highest of the alternative values to 1 and the rest to 0, maintaining consistency with the lightpath indicators as rounded previously. This results in a set of lightpath routings for each lightpath. Then a path of least resistance is followed for each lightpath from among the possible choice of paths to pick the routing for the lightpath. It appears that backtracking may be necessary at this stage. At the end of this stage, a tentative wavelength assignment is also obtained, but this assignment may not be free of clash. The last stage of the procedure is to resolve wavelength clashes. Different approaches to this are discussed in [10], and the one specified involves coloring a path-graph by first listing nodes in order of decreasing degree, removing one node each time, and then sequentially assigning the first available free color to the nodes in this order. A remark is made that the results obtained using the other approaches specified are similar. This results in a wavelength assignment on the original virtual topology. It appears that the bound on the number of wavelengths, which is part of the exact formulation provided, may not be strictly obeyed by the final solution obtained by following this procedure.

4 RELATED APPROACHES

In this section we discuss some techniques and algorithms that are different from those described in Section 3.2.3, but that are related to the problem of virtual topology design for wavelength routed networks.

Incremental benefit analysis: In [15], a study of the incremental benefits of introducing a virtual topology over optical WANs is undertaken. Several simplifying assumptions such as infinite buffers and adequate number of wavelengths are made. The authors obtained data regarding the traffic pattern in the T1 NSFNET backbone network. The data were collected in January 1992. At that time, the NSFNET backbone used optical fibers as physical medium, but only as T1 links, without the use of either WDM or wavelength routing. This data serve to establish the pattern of traffic only, that is, the proportion of total traffic in the network flowing between each source-destination node pair.

The goal of implementing techniques such as WDM or virtual topologies in such a network would be to scale up this traffic pattern. Thus we would like to increase the overall traffic flowing in the network without changing the relative quantities of different source-destination traffic. This is the equivalent of designing a virtual topology for minimum congestion as we discussed before. The pattern is scaled up in three different ways. First, the method of flow deviation is used to scale up the traffic pattern without the use of either WDM or wavelength routing, but merely by routing traffic efficiently. This serves as the baseline with which to compare scaleup benefits obtained with the use of the other two techniques. A maximum scaleup factor of 49 was observed. This scheme uses two lightpaths in each direction corresponding to each fiber, and is thus equivalent to a virtual topology that is the same as the physical topology, with a single wavelength.

The second scheme uses WDM, but no wavelength routing. In other words, each fiber now acts as not one but multiple point-to-point lightpaths, but there are no lightpaths spanning more than one physical link. The nodal degree was limited to 4 and lightpaths were added by inspection. The best scaleup factor was now observed to be 57.

The last scheme applies wavelength routing as well as WDM, to implement arbitrary lightpaths and virtual topologies, though it appears that the virtual topologies are all chosen to be hypercubes. Simulated annealing as discussed previously in [15] was applied to find the best possible virtual topologies in terms of scaleup. The maximum scaleup factor was observed to be 106.

It was noted that the average packet delay, as well as propagation delay and queueing delay, both of which were modeled, increased somewhat over the three schemes. However, the average number of logical hops decreased. The most dramatic result is in the increase of the scaleup factor, and the link utilization, which went from 32% and 23% in the minimum loaded link in the first two schemes to 71% in the last one, while 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 [16] is the lower cost associated with limited conversion of wavelengths at wavelength routers as opposed to full conversion, as we remarked in Section 2. In this study, some terms related to limited conversion are first defined. Theoretical results are derived regarding the virtual topologies that networks with limited conversions can support. A network node is said to have a wavelength degree k if each wavelength at an input port can be switched to one of a maximum of k wavelengths at the output port. Not all wavelengths may be switchable to k different wavelengths. In terms of our earlier terminology, full wavelength conversion would correspond to a wavelength degree of W, with every wavelength switchable to every other, and fixed and no wavelength conversions would both be represented by a wavelength degree of 1. The set of lightpaths to be established, in terms of their source and destination nodes, as well as the route to be followed, are assumed to be given. Thus 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. For example, it is shown that a ring network with full wavelength conversion capability at one node and no wavelength conversion at the others can be used to assign clash-free wavelengths to any set of lightpaths, as long as the maximum number of lightpaths assigned a path over a single physical link is no more than W (which is a physical bound). A similar result is also derived with a ring network that has two nodes of wavelength degree 2, and no wavelength conversion at the others. All 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.

Similar results are derived for more general physical network topologies. In particular, results are derived about a star network where limited wavelength conversion is only employed at the single hub node. There is a corollary that can be used to extend this result to arbitrary topologies, with the (somewhat severe but not entirely unreasonable) restriction that the number of physical hops is no more than 2 for any lightpath in the virtual topology. A result removing this restriction for the case of physical topologies in the form of specially constructed tree networks is stated but no constructive proof is supplied. It is stated that a proof is possible, but it is not mentioned whether the proof is constructive.

All the results in this study are exact and none depend on heuristic methods. It is not entirely clear how to integrate this with heuristic techniques, which can solve other subproblems, or how to extend this work to arbitrary topologies, but the issue of limited wavelength conversion capabilities would appear to be worth further investigation.

Traffic grooming: As we have remarked in Section 1, each lightpath has a high bandwidth and this bandwidth may not be possible to be utilized by single users. 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 traffic routing subproblem may also be viewed as including this multiplexing problem, though we have not viewed it this way in our discussion so far.

The study in [7] is motivated by consideration of this issue. 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 electro-optic 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. Thus, traffic has to be "groomed", or efficiently multiplexed onto lightpaths, based on this consideration. In addition, it is recognized that the manner in which such grooming is done will have an effect on the cost of the network in terms of the number of transceivers that must be placed as part of the add-drop multiplexers and wavelength routers at each node, which loosely corresponds to the logical degree of the node as previously defined. It is suggested that this is a dominant cost in the network, in addition to the number of wavelengths used and the average number of physical hops in lightpaths. These latter metrics are the only ones that have been largely addressed in most studies on wavelength routed optical networks to date.

Only some of the issues raised by these considerations are addressed in [7]. Ring architectures are considered for the physical topology since rings are expected to be of more interest in optical networks in the near future due to availability of ring-type protocols and architectures. Several different ring architectures are specified. In terms of our earlier terminology, one of these implements the physical topology as the virtual topology, and another implements a fully connected virtual topology, that is, a lightpath between every node pair with traffic between them. The other ring architectures proposed implement specific virtual topologies. The ring architectures are compared on the basis of results derived regarding the average number of transceivers at the nodes, number of wavelengths, and average number of physical hops. These parameters reflect the suitability of these architectures for traffic grooming. The conditions of traffic (such as static or dynamic, uniform or non-uniform), under which the different architectures are most useful, are derived. Actual methods of assigning the traffic multiplexing are not discussed. Similar issues arising in arbitrary physical topologies and extension to actual grooming methods would appear to be areas worth further investigation.

Generalized lightpaths: In [19], 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. The physical links implementing a lighttree form a tree rather than a path in the physical topology, hence the name. The points that are emphasized in this study are the following. A lighttree is a more general representation of a lightpath, hence the set of virtual topologies that can be implemented using lighttrees is a superset of the virtual topologies that can be implemented only using lightpaths. Thus for any given virtual topology problem, an optimal solution using lighttrees is guaranteed to be at least as good and possibly an improvement over the optimal solution obtained using only lightpaths. Another attractive feature of lighttrees is the inherent capability for optical multicasting. The current study refers to only unicast and broadcast traffic problems and identifies the multicast problem as an area of ongoing study.

An illustrative example is given, and the mathematical formulation of the problem is outlined. Optical switch architectures involved in networks based on lighttrees are also reviewed.

As pointed out previously, the optimal solution using a more general construct is certain to improve on the optimal solution with a less general one. However, 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 trade-off 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 trade-off is achieved. This appears to be an area worth further investigation.

5 RECONFIGURABILITY CONSIDERATIONS

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.

5.1 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 that 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 [11]. 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 to exploit for the wavelength routed network, which are the topic of this survey.

5.2 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 [2]. The reconfiguration algorithm proposed in [2] 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 that 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

As we have seen, 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. It is not easy or useful to attempt to compare results from the different approaches. In conclusion, we summarize the approaches taken to this problem and qualitative results obtained in the literature.

Virtual topology design over a wide area wavelength routed optical WDM network is an attempt to use the best of both the optical and electronic world. In wide area backbone networks, the lightpaths of a virtual topology are set up to trade off the ample bandwidth available in the fiber with the electro-optic conversion and electronic processing time at intermediate nodes. Studies have been made regarding the benefit obtained by using a virtual topology over using the same fiber network without WDM, or with WDM but no wavelength routing. These studies indicate that using virtual topologies can improve network characteristics, and virtual topology design is an important research area.

We have seen that the usual goals of virtual topology design are to improve some network performance measure such as congestion. Exact formulations of the problem are available in the literature, but they are known to be, or conjectured 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. Heuristic solutions have been proposed that provide tractable approximate solutions to the subproblems. Studies have been undertaken to determine which of several alternative heuristics can be expected to work better under different given network and traffic conditions. Heuristics have also been developed for special cases of network topology and theoretical results have been derived for such special cases.

Virtual topology design is a growing research area. New areas of investigation have been suggested and explored in the literature. These 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. It is expected that many new results will be obtained in this field of research in the future.

REFERENCES

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

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

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

[4] C. Chen and S. Banerjee. Optical switch configuration and lightpath assignment in wavelength routing multihop lightwave networks. Proc. IEEE INFOCOM, pages 1300-1307, 1995.

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

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

[7] O. Gerstel, R. Ramaswami and G. Sasaki. Cost effective traffic grooming in WDM rings. Proc. IEEE INFOCOM, 1998.

[8] P.E. Green. *Fiber optic network*. Englewood Cliffs, NJ Prentice Hall, 1992.

[9] P.E. Green. Optical networking update. IEEE JSAC, 14(5):764-779, Jun 1996.

[10] R.M. Krishnaswamy and K.N. Sivarajan. Design of logical topologies: a linear formulation for wavelength routed optical networks with no wavelength changers. Proc. IEEE INFOCOM, pages 919-927, 1998.

[11] J.-F. P. Labourdette. Traffic optimization and reconfiguration management of multiwavelength multihop broadcast lightwave networks. Computer Networks and ISDN Systems, 30:981-998, May 1998.

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

[13] B. Mukherjee. *Optical communication networks*. McGraw-Hill, 1997.

[14] B. Mukherjee, S. Ramamurthy, D. Banerjee, and A. Mukherjee. Some principles for designing a wide-area optical network. Proc. IEEE INFOCOM, pages 110-119, 1994.

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

[16] R. Ramaswami and G.H. Sasaki. Multiwavelength optical networks with limited wavelength conversion. Proc. IEEE INFOCOM, pages 489-498, 1997.

[17] R. Ramaswami and K.N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. IEEE JSAC, 14(5):840-851, June 1996.

[18] R. Ramaswami and K.N. Sivarajan. *Optical networks: a practical perspective*. Morgan Kaufmann Publishers, 1998.

[19] L.H. Sahasrabuddhe and B. Mukherjee. Light-trees: Optical multicasting for improved performance in wavelength-routed networks. IEEE Communications Magazine, pages 67-73, Feb 1999.

[20] Z. Zhang and A. Acampora. A heuristic wavelength assignment algorithm for multihop wdm networks with wavelength routing and wavelength re-use. IEEE/ACM Trans. Networking, 3(3):281-288, Jun 1995.

This work was supported by the NSF under grant NCR-9701113.

Rudra Dutta

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