

B. Mukherjee, D. Banerjee and S. Ramamurthy are with the Department of Computer Science, University of California, Davis. A. Mukherjee is with the College of Computing, Georgia Institute of Technology.

An earlier version of this paper was presented at IEEE INFOCOM '94.

B. Mukherjee, D. Banerjee, and S. Ramamurthy were supported in parts by NSF Grant Nos. NCR-9205755 and NCR-9508239 and ARPA Contract No. DABT63-92-C-0031. A. Mukherjee was supported by NSF Grant No. NCR-9116177/NCR-9396299.

Some Principles for Designing a Wide-Area WDM Optical Network

Biswanath Mukherjee, Dhritiman Banerjee, S. Ramamurthy and Amarnath Mukherjee

Abstract— We explore design principles for next-generation optical wide-area networks, employing wavelength-division multiplexing (WDM) and targeted to nationwide coverage. This optical network exploits wavelength multiplexers and optical switches in routing nodes, so that an arbitrary virtual topology may be embedded on a given physical fiber network. The virtual topology, which is used as a packet-switched network and which consists of a set of all-optical “lightpaths,” is set up to exploit the relative strengths of both optics and electronics — viz. packets of information are carried by the virtual topology “as far as possible” in the optical domain, but packet forwarding from lightpath to lightpath is performed via electronic switching, whenever required.

We formulate the virtual topology design problem as an optimization problem with one of two possible objective functions: (a) for a given traffic matrix, minimize the network-wide average packet delay (corresponding to a solution for *present traffic demands*), or (b) maximize the scale factor by which the traffic matrix can be scaled up (to provide the maximum capacity upgrade for *future traffic demands*). Since simpler versions of this problem have been shown to be NP-hard, we resort to heuristic approaches. Specifically, we employ an iterative approach which combines “simulated annealing” (to search for a good virtual topology) and “flow deviation” (to optimally route the traffic — and possibly bifurcate its components — on the virtual topology). In this paper, we do not consider the number of available wavelengths to be a constraint, i.e., we ignore the routing of lightpaths and wavelength assignment for these lightpaths. We illustrate our approaches by employing experimental traffic statistics collected from NSFNET.

1 Introduction

New research directions are needed for the development of nationwide “optical” networks. These networks should be able to tap into the tremendous transmission bandwidth potential of a single strand of fiber (on the order of a few terabits per second) as well as exploit other attractive properties of optics. However, they must also allow mechanisms by which end-users operating with slower (opto)electronic components (most likely operating in the gigabits-per-second range) can interface with the network. To correct this bandwidth mismatch, wavelength-division multiplexing (WDM) is a key technique for allowing all of the end-users’ equipment to operate at only electronic speeds. By using WDM, the fiber can carry a high aggregate system capacity of multiple parallel wavelength channels, with each channel operating at much slower speed [11, 5, 9, 13, 10, 19, 20, 23].

We explore new architectures for the next generation of networks employing optical technology for wide-area (nationwide) coverage. We examine an “optical” wide-area

WDM network which utilizes wavelength multiplexers and optical switches in a routing node, so that an arbitrary virtual topology can be embedded on a given physical fiber network. The virtual topology, which is packet switched and which consists of a set of “all-optical lightpaths,” is set up to exploit the relative strengths of both optics and electronics — viz. packets of information are carried by the virtual topology “as far as possible” over the same wavelength in the optical domain (i.e., there is no wavelength conversion [16] in a lightpath), but packet forwarding from lightpath to lightpath is performed via electronic switching, whenever required. Optical switching at a node is achieved by using a wavelength-routing switch (WRS), which is capable of optically bypassing a lightpath from an input fiber to an output fiber, without any electronic processing. Because there is no wavelength conversion in the WRS, the wavelength of the lightpath stays the same in the output fiber as it was in the input fiber.

This architecture is a combination of the well-known “single-hop” [19] and “multihop” [2, 7, 8, 20] approaches, and it attempts to exploit the characteristics of both. A “lightpath” in this architecture provides “single-hop” communication. However, by employing a limited number of wavelengths, it may not be possible to set up “lightpaths” between all user pairs; as a result, “multi-hopping” between “lightpaths” may be necessary. In addition, when the prevailing traffic pattern changes, a different set of “lightpaths” forming a different “multihop” virtual topology may be more desirable. A networking challenge is to perform the necessary reconfiguration with minimal disruption to the network’s operations [14, 15, 26]. In this architecture, using wavelength multiplexers provides the advantage of higher aggregate system capacity due to spatial reuse of wavelengths and supports a large number of users, given a limited number of wavelengths. Specifically, we intend to investigate the overall design, analysis, and optimization of a nationwide WDM network consistent with device capabilities, e.g., aimed at upgrading the NSFNET.

We formulate the problem as an optimization problem which can optimally select a virtual topology subject to transceiver (transmitter and receiver) and wavelength constraints, with one of two possible objective functions: (a) for a given traffic matrix, minimize the network-wide average packet delay (corresponding to a solution for *present traffic demands*), or (b) maximize the scale factor by which the traffic matrix can be scaled up (to provide the maximum capacity upgrade for *future traffic demands*). Since

the objective functions are non-linear and since simpler versions of the problem have been shown to be NP-hard, we resort to heuristic approaches. Specifically, we employ an iterative approach which combines “simulated annealing” (to search for a good virtual topology) and “flow deviation” (to optimally route the traffic – and possibly bifurcate its components – on the virtual topology). We illustrate our approaches by employing experimental traffic statistics collected from NSFNET.

Section 2 explains the system architecture including motivation, general problem statement, and an illustrative example. The problem is formulated as a combinatorial optimization in Section 3. Since the problem is NP-hard, heuristic solutions are developed in Section 4. Results of applying experimental data (collected from NSFNET) to our algorithms are discussed in Section 5. Section 6 concludes the paper.

2 System Architecture

2.1 Motivation

Consider the NSFNET T1 backbone in Fig. 1. Store-and-forward packet switching is performed at the network nodes. Although fiber is employed to connect the nodes, the fiber’s tremendous transmission bandwidth is not exploited since data transmission on each fiber link is performed only at T1 (1.544 Mbps) rate, and on a single wavelength.

Our research on optical networks is fueled by the following criteria and observations. (1) Any future technology must be incrementally deployable. That is, instead of scrapping an existing and operational wide-area fiber-based network such as NSFNET, we examine how such networks can be upgraded to support (and exploit the capabilities of) WDM. In this regard, we show later how one might replace (or upgrade) the existing packet switches in NSFNET with wavelength-routing switches (WRS); or alternately, how one might embellish an electronic switch with an optical component to transform it to a WRS. (2) We must also explore how the WDM solution can be used to upgrade an existing ATM solution. (3) In addition, we must explore how the WDM solution can support both ATM and non-ATM services, based on WDM’s data and protocol transparency property¹. However, the last point is beyond the scope of this paper. In the following subsections, we provide a general problem statement, followed by an illustrative example using the NSFNET T1 topology.

¹Since multiple WDM channels on the same fiber can be operated independently, they can carry data at different rates and formats (including some analog and some digital channels, if desired); also, the protocols controlling the data transfers over different channels can be different, so that one can establish independent subnetworks operating on different sets of WDM channels over the same fiber plant. This is referred to as WDM’s *transparency* property.

2.2 General Problem Statement

The problem of embedding a desired virtual topology on a given physical topology (fiber network) is formally stated below. We are given the following inputs to the problem.

1. A physical topology $G_p = (V, E_p)$ consisting of a weighted undirected graph, where V is the set of network nodes, and E_p is the set of links connecting the nodes. Undirected means that each link in the physical topology is bidirectional. Nodes correspond to network nodes (packet switches) and links correspond to the fibers between nodes; since links are undirected, each link may consist of two fibers or two channels multiplexed (using any suitable mechanism) on the same fiber. Links are assigned weights, which may correspond to physical distances between nodes. A network node i is assumed to be equipped with a $D_p(i) \times D_p(i)$ wavelength-routing switch (WRS), where $D_p(i)$, called the physical degree of node i , equals the number of physical fiber links emanating out of (as well as terminating at) node i ².
2. Number of wavelength channels carried by each fiber = M .
3. An $N \times N$ traffic matrix, where N is the number of network nodes, and the (i, j) -th element is the average rate of traffic flow from node i to node j . Note that the traffic flows may be asymmetric, i.e., flow from node i to node j may be different from the flow from node j to node i .
4. The number of wavelength-tunable lasers (transmitters) and wavelength-tunable filters (receivers) at each node.

Our goal is to determine the following.

1. A virtual topology $G_v = (V, E_v)$ as another graph where the outdegree of a node is the number of transmitters at that node and the indegree of a node is the number of receivers at that node. The nodes of the virtual topology correspond to the nodes in the physical topology. Each link between a pair of nodes in the virtual topology corresponds to a direct all-optical “lightpath” between the corresponding nodes in the physical topology. (Noting that each such link of the virtual topology may be routed over one of several possible paths on the physical topology, an important design issue is “optimal routing” of *all lightpaths* so that the constraint on having a limited number of wavelengths per fiber is satisfied.)
2. A wavelength assignment for lightpaths, such that if two lightpaths share a common physical link, they must necessarily employ different wavelengths.
3. The sizes and configurations of the WRSs at the intermediate nodes. Once the virtual topology is determined and the wavelength assignments have been

²Note that $D_p(i)$ includes the fiber(s) corresponding to local connections, i.e., for attaching an electronic router to the WRS (shown later in this paper). There are wavelength-related and cost-related issues which affect the decision on the number of fibers chosen to connect the local node to the local switch, and these issues are discussed in Section 2.3.

performed, the switch sizes and configurations follow directly.

Communication between any two nodes now takes place by following a path (a sequence of lightpaths) from the source node to the destination node on the virtual topology. Each intermediate node in the path must perform (1) an opto-electronic conversion, (2) electronic routing (possibly ATM switching, if needed), and (3) electro-optic forwarding onto the next lightpath.

2.3 An Illustrative Example

We employ an illustrative example to demonstrate how WDM can be used to upgrade an existing fiber-based network. Using a slightly modified version of the NSFNET in Fig. 1 as an example, we demonstrate how a hypercube (Fig. 2) can be embedded as a virtual topology over this physical topology (to be discussed shortly). We also assume an undirected virtual topology comprising of bidirectional lightpaths in this example. In general, the virtual topology may be a directed graph, as our formulation in Section 3 will assume.

The nodal switching architecture consists of an optical component and an electronic component. The optical component is a wavelength-routing switch (WRS), which can optically bypass some lightpaths, and which can locally terminate some other lightpaths by directing them to node's electronic component. The electronic component is an electronic packet router (which may be an ATM switch), which serves as a store-and-forward electronic overlay on top of the optical virtual topology. Fig. 3 provides a schematic diagram of the architecture of the Utah node in our illustrative example.

The design of a WRS can take several forms [23]. An attractive choice is to employ an array of optical space-division switches, one per wavelength, between the demux and mux stages. (See Fig. 3.) These switches can be re-configured under electronic control, e.g., to adapt the network's virtual topology on demand. One approach would be to have the local lasers/filters normally operated on fixed wavelengths, but a facility to tune them to different wavelengths must be provided.

The virtual topology chosen for our illustration is a 16-node hypercube (e.g., Fig. 2), although algorithms for arbitrary virtual topology embedding will be studied later in this paper. Two embeddings which result from the optimization criteria for hypercube embedding studied in [21] are shown in Figs. 4 and 5. One of these embeddings (Fig. 5) assumes that all of the local laser-filter pairs at a node operate on different wavelengths, while the other (Fig. 4) does not.

Note that each virtual link in the virtual topology of Fig. 2 is a "lightpath" (or a "clear channel") with electronic terminations at its two ends only. For example, the CA1-NE virtual link in Fig. 5 could be set up as an all-optical channel on one of several possible wavelengths on one of several possible physical paths, e.g., CA1-UT-CO-NE, or CA1-WA-IL-NE, or others (see Fig. 5). According to the

solution in [21], the first path is chosen on wavelength 2 for this CA1-NE lightpath. This means that the WRSs at the UT and CO nodes must be properly configured to establish this CA1-NE lightpath. For example, the switch at UT must have wavelength 2 on its fiber to CA1 connected to wavelength 2 on its fiber to CO. Since we have assumed in this example that connections are bidirectional (note that the virtual topology in this illustration is also an undirected graph), the CA1-NE connection implies two directed lightpaths, one from CA1 to NE and the other from NE to CA1.

The solutions in Figs. 4 and 5 require a maximum of 5 and 7 wavelengths per fiber, respectively, by employing shortest-path routing of lightpaths on the physical topology. If only one fiber is used to connect the local node to the local WRS, each of the lightpaths emanating from (and terminating at) that node would need to be on a different wavelength to avoid wavelength conflicts on the local fiber (as in the solution in Fig. 5 for the entire network, and in Fig. 3 for the UT node); accordingly, this solution needs more wavelengths to embed a virtual topology. If multiple fibers are used to connect the local node to the local WRS, multiple lightpaths may emerge from a node on the same wavelength, and hence fewer wavelengths would be needed, e.g., see corresponding solution in Fig. 4 for the same hypercube virtual topology. However, in both solutions, not all wavelengths on all fibers may be utilized, e.g., in Fig. 5, only wavelengths 2 and 4 are used on the CA1-UT fiber, while wavelengths 1, 2, 3, and 5 are used on the UT-CO fiber.

For the solution in Fig. 5, the details of one of the nodes, viz. the one at UT, are shown in Fig. 3. Note that this switch has to support four incoming fibers plus four outgoing fibers, one each to nodes AB, CA1, CO, and MI, as dictated by the physical topology. In general, each switch also interfaces with four lasers (inputs) and four filters (outputs), with each laser-filter pair dedicated to accommodate each of the four virtual links which a node has to support on the virtual topology. The labels "1l 2b 3d 5l" on the output fiber to CO indicate that the UT-CO fiber uses four wavelengths 1, 2, 3, and 5, with wavelengths 2 and 3 being "clear channels" through the UT switch and directed to the physical neighbors CA1 and MI, respectively, while wavelengths 1 and 5 connect to two local lasers. However, in this example, the virtual topology embedded is an "incomplete" hypercube with nodes AB and XY considered non-existent. Hence some nodes, including UT, have fewer than four neighbors. For this illustration, three laser-filter pairs at UT need to be operated – one on wavelength 1 (for connection to physical neighbor CO, to support the lightpath UT-TX); another on wavelength 4 (for connection to physical neighbor CA1, to support the lightpath UT-CA1); and a third on wavelength 5 (for connection to physical neighbor CO, to support the lightpath UT-CO) (see Fig. 5). Labels on the switch's output ports indicate which wavelength on which input fiber or local laser is connected to which wavelength on which output fiber or local filter.

Fig. 3 shows a few more interesting issues of this architecture. The box labelled “Router” is an electronic switch which takes information from terminated lightpaths (1c 4b 5c) as well as a local source (labelled “Workstation”), and routes them to the local lasers (lightpath originators) and the local destination. We reiterate that the “Router” can be any electronic switch, including an ATM switch. Also, the “Router” could have been the existing electronic switch in an electronic network, with its input and output ports connected directly to the incoming and outgoing fibers, respectively, at the switching node. The non-router portions of the node architecture in Fig. 3 are the optical embellishments that may be incorporated to upgrade the electronic switch to incorporate a WRS.

The WRS associated with the Utah switch would be different in the solution corresponding to Fig. 4. Since we would have multiple fibers connecting the electronic router to the WRS, the size of the WRS would need to be a 7×7 switch instead of the 4×4 switch used in the solution corresponding to Fig. 5. However, since the solution corresponding to Fig. 4 requires fewer wavelengths, the number of space division switches inside the largest WRS would reduce from 7 to 5. There are cost implications associated with the WRS design and these costs may influence the solution that is adopted.

3 Formulation of the Optimization Problem

We formulate the problem as an optimization problem, using principles from multicommodity flow for physical routing of lightpaths and traffic flow on the virtual topology, and using the following notation:

1. s and d used as subscript or superscript denote *source* and *destination* of a packet, respectively,
2. i and j denote *originating and terminating nodes, respectively, in a lightpath*, and
3. m and n denote *endpoints of a physical link that might occur in a lightpath*.

βGiven:

- Number of nodes in the network = N .
- Maximum number of wavelengths per fiber = M (a system-wide parameter).
- Physical topology P_{mn} , where $P_{mn} = P_{nm} = 1$ if and only if there exists a direct physical fiber link between nodes m and n , where $m, n = 1, 2, 3, \dots, N$; $P_{mn} = P_{nm} = 0$ otherwise (i.e., fiber links are assumed to be bidirectional).
- Distance matrix, viz. fiber distance d_{mn} from node m to node n . For simplicity in expressing packet delays, d_{mn} is expressed as a propagation delay (in time units). Note that $d_{mn} = d_{nm}$ since fiber links are bidirectional, and $d_{mn} = 0$ if $P_{mn} = 0$.
- Number of transmitters at node $i = T_i$ ($T_i \geq 1$). Number of receivers at node $i = R_i$ ($R_i \geq 1$).

- Traffic matrix λ_{sd} which denotes the average rate of traffic flow from node s to node d , with $\lambda_{ss} = 0$ for $s, d = 1, 2, \dots, N$. (Additional assumptions are that packet interarrival durations at node s and packet lengths are exponentially distributed, so standard M/M/1 queuing results can be applied to each network link (or “hop”) by employing the independence assumption on interarrivals and packet lengths due to traffic multiplexing at intermediate hops. Also, by knowing the mean packet length (in bits/packet), the λ_{sd} can be expressed in units of packets/second.)
- Capacity of each channel = C (normally expressed in bits/second, but converted to units of packets/second by knowing the mean packet length).

βVariables:

- Virtual topology: The variable $V_{ij} = 1$ if there exists a lightpath from node i to node j in the virtual topology; $V_{ij} = 0$ otherwise. Note that this formulation is general since lightpaths are not necessarily assumed to be bidirectional, i.e., $V_{ij} = 1 \not\Rightarrow V_{ji} = 1$.
- Traffic routing: The variable $\lambda_{ij}^{s,d}$ denotes the traffic flowing from node s to node d and employing V_{ij} as an intermediate virtual link. Note that traffic from node s to node d may be “bifurcated” with different components taking different sets of lightpaths.
- Physical topology route: The variable $p_{mn}^{ij} = 1$ if the fiber link P_{mn} is present in the lightpath for virtual link V_{ij} ; $p_{mn}^{ij} = 0$ otherwise.
- Wavelength color: The variable $c_k^{ij} = 1$ if a lightpath from originating node i to terminating node j is assigned the color k , where $k = 1, 2, \dots, M$; $c_k^{ij} = 0$ otherwise.

βConstraints:

- On virtual topology connection matrix V_{ij} :

$$\sum_j V_{ij} \leq T_i \quad \forall i \quad (1)$$

$$\sum_i V_{ij} \leq R_j \quad \forall j \quad (2)$$

The equalities in (1)-(2) hold if all transmitters at node i and all receivers at node j are in use.

- On physical route variables p_{mn}^{ij} :

$$p_{mn}^{ij} \leq P_{mn} \quad (3)$$

$$p_{mn}^{ij} \leq V_{ij} \quad (4)$$

$$\sum_m p_{mk}^{ij} = \sum_n p_{kn}^{ij} \quad \text{if } k \neq i, j \quad (5)$$

$$\sum_n p_{in}^{ij} = V_{ij} \quad (6)$$

$$\sum_m p_{mj}^{ij} = V_{ij} \quad (7)$$

- On virtual topology traffic variables $\lambda_{ij}^{s,d}$:

$$\lambda_{ij}^{sd} \geq 0 \quad (8)$$

$$\sum_j \lambda_{sj}^{sd} = \lambda_{sd} \quad (9)$$

$$\sum_i \lambda_{id}^{sd} = \lambda_{sd} \quad (10)$$

$$\sum_i \lambda_{ik}^{sd} = \sum_j \lambda_{kj}^{sd} \quad \text{if } k \neq s, d \quad (11)$$

$$\sum_{s,d} \lambda_{ij}^{sd} \leq V_{ij} * C \quad (12)$$

- On coloring of lightpaths c_k^{ij} :

$$\sum_k c_k^{ij} = V_{ij} \quad (13)$$

$$\sum_{ij} p_{mn}^{ij} \cdot c_k^{ij} \leq 1 \quad \forall m, n, k \quad (14)$$

Objective: Optimality Criterion

- (a) Delay minimization:

$$\text{Minimize } \sum_{ij} \left(\sum_{sd} \lambda_{ij}^{sd} \left[\sum_{mn} p_{mn}^{ij} * d_{mn} + \frac{1}{C - \sum_{sd} \lambda_{ij}^{sd}} \right] \right) \quad (15)$$

- (b) Maximizing offered load (equivalent to minimizing maximum flow in a link):

$$\min \left(\max \left[\sum_{sd} \lambda_{ij}^{sd} \right] \right) \equiv \max \frac{C}{\min (\max [\sum_{sd} \lambda_{ij}^{sd}])} \forall i, j \quad (16)$$

Explanation of Equations The above equations are based on principles of conservation of flows and resources (transceivers, wavelengths, etc.) as well as on conflict-free routing, e.g., two lightpaths that share a fiber should not be assigned the same wavelength. Equations (1)-(2) ensure that the number of lightpaths emanating out of and terminating at a node are at most equal to that node's out-degree and in-degree, respectively. Equations (3)-(4) constrain the problem so that p_{mn}^{ij} can exist only if there is a physical fiber and a corresponding lightpath present. Equations (5)-(7) are the multicommodity equations that account for the routing of a lightpath from its origin to its termination. Equations (8)-(12) are responsible for the routing of packet traffic on the virtual topology, and they take into account the fact that the combined traffic flowing through a channel cannot exceed the channel capacity. Equation (13) requires that a lightpath be of one color only. Equation (14) ensures that the colors used in different lightpaths are mutually exclusive over a physical link.

Equations (15) and (16) represent two possible objective functions. In Equation (15), in the innermost brackets, the first component corresponds to the propagation delays on the links mn which form the lightpath ij , while the second component corresponds to delay due to queueing and

packet transmission on lightpath ij (using a M/M/1 queueing model for each lightpath). If we assume shortest-path routing of the lightpaths over the physical topology, then the p_{mn}^{ij} values become deterministic. If, in addition, we neglect queueing delays, the optimization problem in Equation (15) reduces to minimizing $\sum_{sd} \sum_{ij} \sum_{mn} \lambda_{ij}^{sd} \cdot p_{mn}^{ij} \cdot d_{mn}$ which is a mixed-integer linear program in which the V_{ij} and the c_k^{ij} variables need to have integer solutions, while the λ_{ij}^{sd} variables do not.

The objective function in Equation (16) is also nonlinear and it minimizes the maximum amount of traffic that flows through any lightpath. This corresponds to obtaining a virtual topology which can maximize the offered load to the network if the traffic matrix is allowed to be scaled up. We choose this optimization for our algorithms in Section 4 because our purpose is to demonstrate how to upgrade the capacity of existing fiber-based networks by employing WDM.

4 Algorithms

4.1 Subproblems

The optimization problem in Section 3 is NP-hard, since several sub-problems of this problem are NP-hard. The problem of optimal virtual topology design can be partitioned into the following four subproblems, which are not necessarily independent:

1. determine a good virtual topology, viz. which nodal transmitter should be directly connected to which nodal receiver,
2. route the lightpaths over the physical topology,
3. assign wavelengths optimally to the various lightpaths (this problem has been shown to be NP-hard in [8]), and
4. route packet traffic on the virtual topology (as in any packet-switched network).

Subproblem 1 addresses how to properly utilize the limited number of available transmitters and receivers. Subproblems 2 and 3 deal with proper usage of the limited number of available wavelengths, and has been addressed in [3, 25]. Subproblem 4 minimizes the effect of store-and-forward (queueing plus transmission) delays at intermediate electronic hops.

4.2 Previous Work

The problem of designing optimal virtual topologies has been studied before [4, 14]. Our formulation is more general in the sense that we accommodate many of the physical connectivity constraints which were not considered earlier. In general, the optimal virtual topology problem has been conjectured to be NP-hard, which means that the problem cannot be solved optimally for large problem sizes, unless one resorts to some form of exhaustive search. One instance of this problem has been formulated as a mixed integer linear program which gets difficult to solve with increasing

problem size [14]. Accordingly, heuristic approaches have been employed to solve these problems [4, 14].

Related work on these problems can be found in [8, 21, 24, 29]. Embedding of a packet-switched virtual topology on a physical fiber plant in a switched network was first introduced in [8], and this network architecture was referred to as a *lightnet*. Some algorithms to embed a hypercube virtual topology were provided in [8, 21]. The work in [24] proposes a virtual topology design for packet-switched networks. The average hop distance is minimized which automatically increases the network traffic supported. The work in [24] uses the physical topology as a subset of the virtual topology, employing algorithms for maximizing the throughput subject to bounded delay characteristics.

The Routing and Wavelength Assignment (RWA) problem (Subproblems 2 and 3) for routing connections in circuit-switched networks has been studied in [3, 25, 29]. In [29], a heuristic algorithm is presented for effectively assigning a limited number of wavelengths among the access stations of a multihop switched network. In [25], an upper bound is derived on the carried traffic for any RWA algorithm, while in [3], a sequence of algorithms is assembled to efficiently solve large instances of the RWA problem. Conflict-free wavelength routing in wide-area optical networks employing the concept of Latin Squares for configuring the wavelength switches is studied in [6]. Some theoretical results on the bounds for minimum congestion routing have also been studied in [28] by using the concept of flow trees to study minimum and maximum congestion for a network.

Some approaches also employ wavelength conversion at intermediate nodes [16]. These approaches mitigate the need for having wavelength continuity, i.e., the fact that a lightpath needs to use the same wavelength on each fiber that it passes through. Recoloring of lightpaths in order to accommodate new connections in a dynamic environment has been examined in [17], while the use of wavelengths outside the Erbium-doped region for routing of short lightpaths is studied in [18]. In addition to the above theoretical studies, a number of comprehensive proposals [11, 9, 10] have appeared on the design of optical WDM WANs including the development of devices, protocols, architectures, and applications.

4.3 Our Solution Approach

To obtain a thorough understanding of the problem, we concentrate on Subproblems 1 and 4 above, i.e., *for the purposes of this paper, we do not consider the number of available wavelengths to be a constraint*. In the expanded problem, both the number of wavelengths and their exact assignments are critical, and this is a topic of our ongoing study. Specifically, we employ an iterative approach consisting of “simulated annealing” to search for a good virtual topology (Subproblem 1), in conjunction with the “flow deviation” algorithm for optimal (possibly “bifurcated”) routing of packet traffic on the virtual topology (Subproblem 4). Also, although the virtual topology can be an undirected graph, we consider lightpaths to be bidi-

rectional in our solution here since most (Internet) network protocols rely on bidirectional paths and links. In addition, we consider Optimization Criterion (b) of Equation (16) (maximizing offered load) for our illustrative solution below, mainly because we are interested in upgrading an existing fiber-based network to a WDM solution.

We start with a random configuration (virtual topology) and try to find a good virtual topology through simulated annealing by using node-exchange (similar to branch-exchange [14]) techniques. Then, we *scale up* the traffic matrix to ascertain the *maximum throughput* that can be accommodated by the virtual topology, using flow deviation for packet routing over the virtual topology. For a given traffic matrix, the flow-deviation algorithm minimizes the network-wide packet delay by properly distributing the flows on the virtual links (to reduce the effect of large queueing delays).

We have used measured data over the NSFNET backbone as our sample traffic matrix. We scale up each entry in the traffic matrix by a constant scaleup factor and verify if the offered load from the scaled-up traffic matrix can be accommodated by the virtual topology. Our goal is to design the virtual topology that can accommodate the maximum traffic scaleup. This provides an estimate of the maximum throughput we can expect from the current fiber network if it were to support WDM, and if future traffic characteristics were to model present-day traffic characteristics except for the traffic intensities to grow by a constant scale factor. While it is difficult to predict future traffic characteristics, we believe that our approach provides a reasonable framework for analysis and design.

4.4 Simulated Annealing

Simulated annealing (along with genetic algorithms) has been found to provide good solutions for hard optimization problems [1]. In the simulated annealing process, the algorithm starts with an initial random configuration for the virtual topology. Node-exchange operations are used to arrive at neighboring configurations. In a node-exchange operation, adjacent nodes in the virtual topology are examined for swapping, e.g., if node i is connected to nodes j , a , and b , while node j is connected to nodes p , q , and i in the virtual topology, after the node-exchange operation between nodes i and j , node i will be connected to nodes p , q , and j , while node j will be connected to nodes a , b , and i . Neighboring configurations which give better results (lower average packet delay) than the current solution are accepted automatically. Solutions which are worse than the current one are accepted with a certain probability which is determined by a system control parameter. The probability with which these failed configurations are chosen, however, decreases as the algorithm progresses in time so as to simulate the “cooling” process associated with annealing. The probability of acceptance is based on a negative exponential factor and is inversely proportional to the difference between the current solution and the best solution obtained so far.

4.5 Flow Deviation Algorithm

By properly adjusting link flows, the flow deviation algorithm [12] provides an optimal algorithm for minimizing the network-wide average packet delay. However, traffic from a given source to a destination may be bifurcated, i.e., different fractions of it may be routed along different paths in order to minimize the packet delay. If the flows are not balanced, then excessively loading of a particular channel may lead to large delays on that channel and thus have a negative influence on the network-wide average packet delay. The algorithm is based on the notion of *shortest-path flows* which first calculates the linear rate of increase in the delay with an infinitesimal increase in the flow on any particular channel. These “lengths” or “cost rates” are used to pose a shortest-path flow problem (which can be solved using one of several well-known algorithms such as Dijkstra’s algorithm, Bellman-Ford algorithm, etc. [27]) and the resulting paths represent the “cheapest” paths on which some of the flow may be deviated. An iterative algorithm determines *how much* of the original flow needs to be deviated. The algorithm continues until a certain performance tolerance level is reached.

5 Experimental Results

The simulated annealing algorithm as well as the flow deviation algorithm were both used to derive results for the virtual topology design problem, viz. to study Subproblems 1 and 4 outlined in Section 4.1. The traffic matrix employed for this mapping is an actual measurement of the traffic on the T1 NSFNET backbone for a 15-minute period (11:45pm to midnight on January 12, 1992, EST). The raw traffic matrix showing traffic flow in bytes per 15-minute intervals between network nodes is shown in Fig. 6.³ Nodal distances used are the actual geographical distances and they are not shown here to conserve space. Initially, it was assumed that each node could set up at most four bidirectional lightpath channels, but later more experiments were conducted to study the effect of having higher nodal degree. The number of wavelengths per fiber was assumed to be large enough so that all possible virtual topologies could be embedded⁴.

Some of our results are tabulated in Table 1. For each experiment, the maximum scaleup achieved is tabulated along with the corresponding individual delay components, the maximum and minimum link loading as well as the av-

³NSFNET backbone data was collected by the nstat program and made available to us by the National Science Foundation (NSF) through its Merit partnership. The traffic matrix shown in Fig. 6 is not exactly the same as the one used in our previous work [21]. Discrepancies can be attributed to different ways of filtering the raw data. The matrix we are currently using is more accurate. Also, the traffic matrix in Fig. 6 shows that several nodes have “self traffic”. This is due to the fact that nodes in the NSFNET are actually gateways connecting to regional networks, so some of the intra-regional traffic at each node also showed up in our measurements.

⁴If we limit the number of wavelengths supported on each fiber, it might not be possible to set up all possible lightpaths in a virtual topology. There are established bounds on the minimum number of wavelengths required to set up arbitrary virtual topologies for a given number of nodes [22].

erage hop distance⁵. Since the aggregate capacity for the carried traffic is fixed by the number of links in the network, reducing the average hop distance can lead to higher values of load that the network can carry. The queuing delay was calculated using a standard M/M/1 queuing system with a mean packet length calculated from the measured traffic (133.54 bytes/packet) and a link speed of 45 Mbps. *We assume infinite buffers at all nodes.* The “cooling” parameter for the simulated annealing is updated after every 100 acceptances using a geometric parameter of value 0.9. A state is considered “frozen” when there is no improvement over 100 consecutive trials.

§Physical Topology as Virtual Topology (No WDM) Our goal was to obtain a fair estimate of what optical hardware can provide in terms of extra capabilities. In this experiment, we start off with just the existing hardware (as in any electronic, packet-switched network), comprising of fiber and point-to-point connections using a single bidirectional lightpath channel per fiber link, i.e., *no WDM is used.* Using flow deviation, the maximum scaleup that could be achieved was found to be 49 (only integer values of the scaleup were considered). The link with the maximum traffic (WA-IL) was loaded at 98%, while the link with the minimum traffic (NY-MD) was at 32%. (These numbers are truncated to show their integer part only.) These values serve as a basis for comparison as to what can be gained in terms of throughput by adding extra WDM optical hardware, viz. tunable transceivers and wavelength routing switches at nodes.

§Multiple point-to-point links (No WRS) The goal of the next experiment was to determine how much throughput we could obtain from the network without adding any photonic switching capability at a node, but by adding extra transceivers (up to four) at each node, i.e., *WDM is used on some links, but no WRS capability is employed at any node.* The initial network had 21 links in the physical topology (see Fig. 1). Using extra transceivers at the nodes, we set up extra links on the paths NE-CO, NE-IL, WA-CA2, CA1-UT, MI-NJ, and NY-MD. These lightpaths are chosen manually. Different combinations were considered and the choice of channels which provided the maximum scaleup was chosen. Given 14 nodes, each with a nodal degree of four, we should have been able to have 28 channels. However, the GA node is connected only to TX and PA, both of which are physically connected to four nodes already; hence, they do not have any free transceivers to create an extra channel to GA. Thus, we could add only six new channels. In this case, the maximum scaleup was found to be 57. We found that the two NY-MD channels had a load of only 23%, while the UT-MI channel had a load of 99%.

§Arbitrary Virtual Topology (Full WDM) In the final experiment, we assumed *full WDM with all nodes equipped with WRSs*, i.e., it may now be possible to set up light-

⁵The average hop distance for packets is an important figure of merit for the topology chosen. It not only has a direct bearing on the queuing delays that a packet suffers, but more importantly, it determines the maximum offered load for the network. The product of the average hop distance and the offered load in the network equals the aggregate traffic load on all the links.

paths between any two nodes. We did not constrain the number of wavelengths supported in each fiber, so that all graphs of degree four were candidates for possible virtual topologies; also, all lightpaths were assumed to be routed over the shortest path on the physical topology. Starting off with a random initial topology, we used simulated annealing to get the best virtual topology. The experiment ran on an unloaded Sparc 10 for approximately three days. The best virtual topology, which is shown in Table 2, provided a maximum scaleup of 106. Clearly, the increased scaleup demonstrates the benefits of the WDM-based virtual-topology approach. Now, the minimum loading was on link UT-TX at 71%, while all the other links were above 98% loading.

Comparisons The various delay characteristics (viz. overall average packet delay (FD), average propagation delay encountered by each packet (PD), average queueing delay experienced by each packet (QD), and the mean hop distance (HD)), as functions of the scaleup (throughput), for the above three experiments are shown in Figs. 7 through 9. The scaleup provides an estimate of the throughput in the network. We note from these figures, as well as from Table 2, that the propagation delay is the dominant component of the packet delay. Also, at light loads, the average propagation delay faced by packets in NSFNET is a little over 9 ms (for the given traffic matrix), and this serves as a lower bound on the average packet delay. As a basis for comparison, the coast-to-coast one-way propagation time is nearly 23 ms. Thus, on an average, each packet travels about 40% of the coast-to-coast distance.

Note that the average queueing delay increases slightly with increasing traffic until the scaleup nearly reaches its maximum value, after which there is a sharp increase in queueing delay. The propagation delay also increases with increasing scaleup as more traffic is deviated away from shortest path routes by the flow deviation algorithm. One interesting feature is that the average hop distance decreases as the traffic load is increased; this is again because of the flow deviation algorithm which will deviate traffic onto longer links, which might increase the propagation delay encountered by a packet (compared to shortest-path routing), but helps to decrease the average hop distance.

Fig. 10 plots the different aggregate average packet delays for the three different schemes. The throughput advantage of having no WDM vs. using WDM on a few links but no WRS vs. employing full WDM and WRSs is again clear from this figure. But we notice that the delay in the first two cases is lower than that when using a virtual topology. This is because, in the full WDM case, the shortest path along the physical topology cannot always be chosen because of the virtual topology, so that some packets may have to travel longer distances, in general. However, the scaleup in the virtual topology is much more than that for the other two schemes; so addition of switches at intermediate nodes to do wavelength routing provides a significant improvement in throughput for the network.

Effect of Nodal Degree and Wavelength Requirements

So far, the nodal degree (P) was four. Now, we consider

full WDM (with a WRS at each node), and increase the nodal degree to 5 and 6. We find that the maximum scaleup increased nearly proportionally with increasing nodal degree. Actually, with the scaleup of 106 for $P=4$ as a baseline, proportional increase in scaleup for $P=5$ and 6, would yield 132.5 and 159, respectively. However, in our experiments, the observed maximum scaleups for $P=5$ and 6 were higher, viz. 135 and 163, respectively (refer to Table 3). This is due to the fact that, as the nodal degree is increased, the average hop distance of the virtual topology is reduced, which provides the extra improvement in the scaleup. Minimizing hop distance can be an important optimization problem, and has been studied in [24].

Although no constraints on wavelengths per fiber were imposed in this study, we also examined the wavelength requirements to set up a virtual topology using shortest-path routing of lightpaths on the physical topology. Assuming no limit on the supply of wavelengths, but with the wavelength constraints as outlined before (Section 3), the maximum number of wavelengths required for embedding the best virtual topology (which provided the maximum scaleup) with degree $P=4, 5,$ and 6 were found to be 6, 8, and 8 wavelengths, respectively. The corresponding distributions of the number of wavelengths used in each of the 21 fiber links of the NSFNET (see Fig. 1) are plotted in Fig. 11. We find that, with increasing nodal degree, i.e., with an increasing number of lightpaths to be supported, the average number of wavelengths a fiber needs to support increases. However, due to the combination of reasons such as desired virtual topology, shortest-path routing of lightpaths, and wavelength constraints, it may so happen that there is no link on the physical topology that employs all of the required wavelengths. This happened for our $P=6$ experiment, i.e., although 8 wavelengths were required to embed the virtual topology, no physical link carried all 8 wavelengths.

6 Discussion

We explored design principles for next-generation optical wide-area networks, employing wavelength-division multiplexing (WDM) and targeted to nationwide coverage. We showed that such a WDM-based network architecture can provide a high aggregate system capacity due to spatial reuse of wavelengths. Our objective was to investigate the overall design, analysis, upgradability, and optimization of a nationwide WDM network consistent with device capabilities.

The virtual topology optimization problem discussed in this paper serves as an illustration, and it is a first step towards a robust and versatile WDM WAN solution. We believe that a significant amount of room exists for developing improved approaches and algorithms, and a number of research issues must be addressed in this regard. An interesting avenue of research is to study how routing and wavelength assignment of lightpaths can be combined with the choice of virtual topology and its corresponding packet

routing in order to arrive at an optimum solution. Dynamic establishment and reconfiguration of lightpaths is an important issue which needs to be thoroughly studied.

Acknowledgments

We acknowledge Eric Aupperle, Bilal Chinoy, Elise Gerich, Susan Horvath, and Mark Knopper from NSF/Merit Partnership for providing us with the NSFNET data. We thank Rajiv Ramaswami, Tony Acampora, Zhensheng Zhang, and the reviewers for several helpful comments.

<i>Parameter</i>	<i>Physical Topology</i>	<i>Many pt – pt links</i>	<i>Arbitrary v.t.</i>
Maximum scale up	49	57	106
Avg. Pkt. Delay	11.739 ms	12.5542 ms	17.8146 ms
Avg. Prop. Delay	10.8897 ms	10.9188 ms	14.4961 ms
Queueing Delay	0.8499 ms	1.6353 ms	3.31849 ms
Avg. Hop Distance	2.12186	2.2498	1.71713
Max. Link Loading	98%	99%	99%
Min. Link Loading	32%	23%	71%

Table 1: Summary of experimental results.


<i>Src</i>	<i>Neighbours</i>
WA	CA1, CA2, MI, UT
CA1	WA, CO, IL, TX
CA2	WA, PA, NE, GA
UT	WA, TX, IL, MD
CO	CA1, MD, NE, GA
TX	CA1, UT, GA, NJ
NE	CA2, CO, IL, MI
IL	CA1, UT, NE, PA
PA	CA2, IL, NY, NJ
GA	CA2, CO, TX, NY
MI	WA, NE, NY, NJ
NY	PA, GA, MI, MD
NJ	TX, PA, MI, MD
MD	CO, NY, NJ, UT

Table 2: Virtual topology for nodal degree = 4 and best scaleup (106).

Transceivers/node	Scaleup
4	106
5	135
6	163

Table 3: Traffic scaleups for different nodal degrees.

Figure 1: NSFNET T1 backbone.



Hardwired solution of the switch at UTAH

Figure 3: Details of the Utah (UT) node.

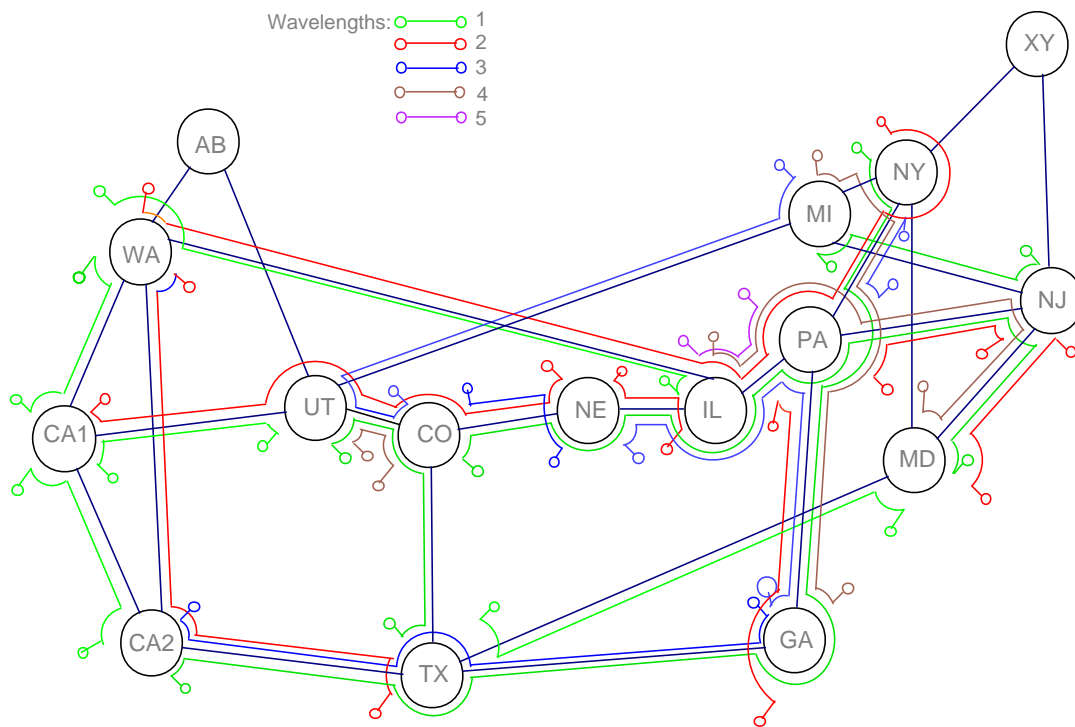


Figure 4: The physical topology with embedded wavelengths corresponding to an optimal solution (more than one transceiver at any node can tune to the same wavelength).

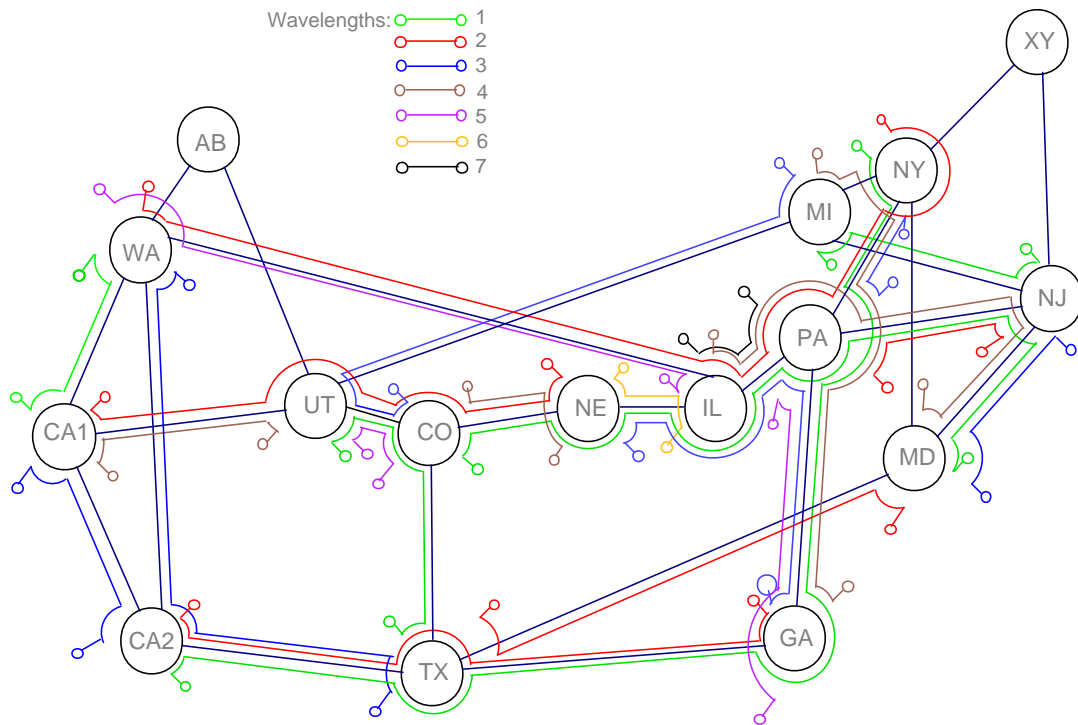


Figure 5: The physical topology with embedded wavelengths corresponding to an optimal solution (all transceiver pairs at any node must be tuned to different wavelengths).

		<i>Traffic Matrix</i> (multiply by 10 to get bytes per 15-minute interval)													
	WA	CA1	CA2	UT	CO	TX	NE	IL	PA	GA	MI	NY	NJ	MD	
WA	53065	268225	117095	27185	196575	8795	53825	249035	34220	18515	311785	96725	44195	191410	
CA1	719135	39095	610070	301290	586365	261795	398780	1549725	114465	214105	799325	1031395	552420	775865	
CA2	109160	475695	360	466130	85075	363705	86375	856715	100325	46200	516350	62095	139155	215780	
UT	70165	62050	136445	0	19090	6090	7015	28815	20030	32600	131120	121575	23925	69745	
CO	1227715	1599940	190245	34365	3595	40365	107790	622270	240235	179245	721075	1185620	131820	217565	
TX	18415	165355	34265	55215	34030	0	26145	26850	8780	38720	60580	48245	15425	69555	
NE	370065	620060	1023115	44755	220365	79000	0	1141795	198280	219565	1540270	933255	236695	1638750	
IL	149525	2345475	2103480	85225	282150	26675	970820	3185	439515	330060	900570	711585	202005	889010	
PA	849300	199420	373500	60070	249915	68140	250690	610200	0	396225	1106855	1476115	456745	631390	
GA	18630	419260	102640	37395	223405	94815	49890	570845	68455	1400	363215	261715	126950	143675	
MI	111680	376125	582980	50665	94450	129940	187900	378915	204770	251215	454995	596745	322770	371920	
NY	312275	1318380	198710	146225	429985	71520	173240	573215	396000	294260	2116365	742465	279960	659740	
NJ	393745	553420	186005	75395	84165	8515	44905	243960	1176820	356930	691765	792130	70690	521995	
MD	819050	2270095	542870	229610	892785	318230	327035	918520	306075	16635	1296980	1375980	627470	1216290	

Figure 6: Traffic matrix (in bytes per 15-minute interval).




Figure 7: Delay vs. throughput (scaleup) characteristics with no WDM, i.e., physical topology as virtual topology.




Figure 8: Delay vs. throughput (scaleup) characteristics with WDM used on some links, but no WRSs, i.e., multiple point-to-point links are allowed on the physical topology.




Figure 9: Delay vs. throughput (scaleup) characteristics with full WDM on some links and a WRS at each node, i.e., arbitrary virtual topologies are allowed.

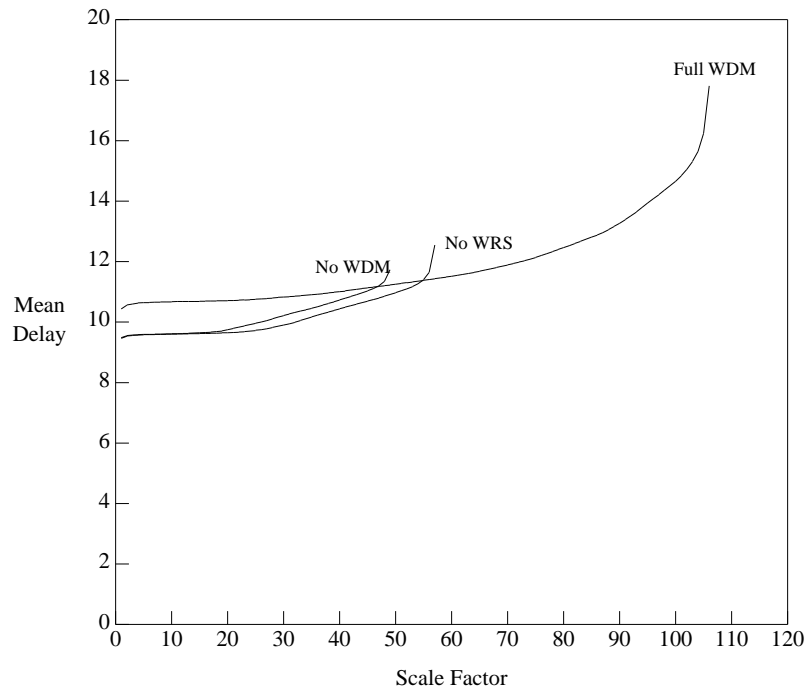


Figure 10: Delay vs. throughput (scaleup) characteristics for different virtual topologies.

References

- [1] A. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines*. John Wiley and Son, 1989.
- [2] A. S. Acampora. A multichannel multihop local lightwave network. In *Proceedings, IEEE GlobeCom '87*, Tokyo, Japan, November 1987.
- [3] D. Banerjee and B. Mukherjee. Practical approaches for routing and wavelength assignment in large all-optical wavelength-routed networks. *IEEE Journal on Selected Areas in Communications*, 14(5), June 1996.
- [4] J. A. Bannister, L. Fratta, and M. Gerla. Topological design of the wavelength-division optical network. In *Proceedings, IEEE INFOCOM '90*, pages 1005–1013, San Francisco, CA, June 1990.
- [5] C. A. Brackett. Dense wavelength division multiplexing networks: Principle and applications. *IEEE Journal on Selected Areas in Communications*, 8(6):948–964, August 1990.
- [6] C. Chen and S. Banerjee. Optical switch configuration and lightpath assignment in wavelength routing multihop lightwave networks. In *Proceedings, IEEE INFOCOM '95*, pages 1300–1307, Boston, MA, June 1995.
- [7] I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications : An approach to high bandwidth optical WANs. *IEEE Transactions on Communications*, 40(7):1171–1182, July 1992.
- [8] I. Chlamtac, A. Ganz, and G. Karmi. Lightnets: Topologies for high speed optical networks. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):951–961, May/June 1993.
- [9] C. A. Brackett et. al. A scalable multiwavelength multihop optical network : A proposal for research on all-optical networks. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):736–753, May/June 1993.
- [10] G. Hill et. al. A transport network layer based on optical network elements. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):667–679, May/June 1993.
- [11] S. B. Alexander et. al. A precompetitive consortium on wide-band all-optical networks. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):714–735, May/June 1993.
- [12] L. Fratta, M. Gerla, and L. Kleinrock. The flow deviation method: an approach to store-and-forward communication network design. *Networks*, 3:97–133, 1973.
- [13] P. E. Green. *Fiber Optic Networks*. Prentice Hall, 1993.
- [14] J. F. P. Labourdette and A. S. Acampora. Reconfiguration algorithms for rearrangeable lightwave networks. *IEEE Transactions on Communications*, 39(8):1223–1230, August 1991.
- [15] J. F. P. Labourdette, G. W. Hart, and A. S. Acampora. Branch-exchange sequences for reconfiguration of lightwave networks. *IEEE Transactions on Communications*, 42(10):2822–2832, October 1994.
- [16] K. C. Lee and V. O. K. Li. A wavelength convertible optical network. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):962–970, May/June 1993.
- [17] K. C. Lee and V. O. K. Li. A circuit rerouting algorithm for all-optical wide-area networks. In *Proceedings, IEEE INFOCOM '94*, pages 954–961, Toronto, Canada, June 1994.
- [18] K. C. Lee, V. O. K. Li, S. M. Hwang, and A. E. Willner. Routing for all-optical networks using wavelengths outside erbium-doped fiber amplifier bandwidth. *Journal of Lightwave Technology*, 13(5):791–801, May 1995.
- [19] B. Mukherjee. WDM-based local lightwave networks – part I: Single-hop systems. *IEEE Network Magazine*, 6(3):12–27, May 1992.
- [20] B. Mukherjee. WDM-based local lightwave networks – part II: Multihop systems. *IEEE Network Magazine*, 6(4):20–32, July 1992.
- [21] B. Mukherjee, S. Ramamurthy, D. Banerjee, and A. Mukherjee. Some principles for designing a wide-area optical network. In *Proceedings, IEEE INFOCOM '94*, pages 110–119, Toronto, Canada, June 1994. (Also, expanded version submitted to *IEEE/ACM Transactions on Networking*).
- [22] R. K. Pankaj. *Architectures for Linear Lightwave Networks*. PhD thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, September 1992.
- [23] R. Ramaswami. Multiwavelength lightwave networks for computer communication. *IEEE Communications Magazine*, 31(2):78–88, February 1993.
- [24] R. Ramaswami and K. Sivarajan. Design of logical topologies for wavelength-routed all-optical networks. In *Proceedings, IEEE INFOCOM '95*, pages 1316–1325, Boston, MA, April 1995.
- [25] R. Ramaswami and K. Sivarajan. Optimal routing and wavelength assignment in all-optical networks. *IEEE/ACM Transactions on Networking*, 3(5):489–500, October 1995.
- [26] G. N. Rouskas and M. H. Ammar. Dynamic reconfiguration in multihop WDM networks. *Journal of High Speed Networks*, 4(3):221–238, 1995.
- [27] W. Stallings. *Data and Computer Communications*. Macmillan Pub. Co., New York, NY, fourth edition, 1994.
- [28] B. Yener and T. Boult. A study of upper and lower bounds for minimum congestion routing in lightwave networks. In *Proceedings, IEEE INFOCOM '94*, pages 138–148, Toronto, Canada, June 1994.
- [29] Z. Zhang and A. Acampora. A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse. *IEEE/ACM Transactions on Networking*, 3(3):281–288, June 1995.

Biswanath Mukherjee (S '82 - M '87) received the B.Tech. (Hons) degree from Indian Institute of Technology, Kharagpur (India) in 1980 and the Ph.D. degree from University of Washington, Seattle, in June 1987. At

Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship. In July 1987, he joined the University of California, Davis, where he became Professor of Computer Science in July 1995. He is co-winner of prize paper awards presented at the 1991 and the 1994 National Computer Security Conferences. He serves on the editorial boards of the IEEE/ACM Transactions on Networking and Journal of High-Speed Networks. He co-guest-edited a special issue of the Journal of High-Speed Networks on the topic of WDM Networks, published in 1995. He also served as Technical Program Chair for the IEEE INFOCOM '96 Conference. His research interests include lightwave networks, high-speed local and metropolitan area networks, and network security. Readers may contact him via e-mail at mukherje@cs.ucdavis.edu.

Dhritiman Banerjee received the B.Tech. (Hons) degree from Indian Institute of Technology, Madras (India) in 1992 and the M.S. degree from the University of California, Davis in 1994.

He is currently a Research Assistant with the Networks Research Lab at the University of California, Davis, where he is working towards the Ph.D. degree. His research interests include architectures and protocols for WDM wide-area optical networks. Readers may contact him via email at banerjed@cs.ucdavis.edu.

S. Ramamurthy received the B. Tech. (Hons) degree in Computer Science from the Indian Institute of Technology, Madras (India) in 1991. He is currently working towards his doctoral degree in the Department of Computer Science at the University of California, Davis, and is currently employed at Bay Networks, Santa Clara, CA. His research interests include design, analysis, and modeling of computer networks, particularly optical networks, and study of network algorithms.

Amarnath Mukherjee received his Ph.D. degree in Computer Science from the University of Wisconsin at Madison in 1990, and a Bachelor of Technology (Honors) degree in Computer Science and Engineering from the Indian Institute of Technology at Kharagpur in 1984. He works on network workload modeling and resource-management algorithms. Earlier, he worked on error recovery protocols.




Figure 11: Distributions of the number of wavelengths used in each of the 21 fiber links of the NSFNET for the virtual topology approach with nodal degree $P = 4, 5,$ and 6 .