

Survivable Lightpath Routing: A New Approach to the Design of WDM-Based Networks

Eytan Modiano, *Senior Member, IEEE* and Aradhana Narula-Tam

Invited Paper

Abstract—Network restoration is often done at the electronic layer by rerouting traffic along a redundant path. With wavelength-division multiplexing (WDM) as the underlying physical layer, it is possible that both the primary and backup paths traverse the same physical links and would fail simultaneously in the event of a link failure. It is, therefore, critical that lightpaths are routed in such a way that a single link failure would not disconnect the network. We call such a routing *survivable* and develop algorithms for survivable routing of a logical topology. First, we show that the survivable routing problem is NP-complete. We then prove necessary and sufficient conditions for a routing to be survivable and use these conditions to formulate the problem as an integer linear program (ILP). Due to the excessive run-times of the ILP, we develop simple and effective relaxations for the ILP that significantly reduces the time required for finding survivable routings. We use our new formulation to route various logical topologies over a number of different physical topologies and show that this new approach offers a much greater degree of protection than alternative routing schemes such as shortest path routing and a greedy routing algorithm. Finally, we consider the special case of ring logical topologies for which we are able to find a significantly simplified formulation. We establish conditions on the physical topology for routing logical rings in a survivable manner.

Index Terms—Lightpaths, network design, network survivability, routing, wavelength-division multiplexing.

I. INTRODUCTION

THIS PAPER DEALS with the problem of routing logical links (lightpaths) on a physical network topology in such a way that the logical topology remains connected in the event of single physical link failures (e.g., fiber cut). This is a relatively new view on the routing and wavelength assignment (RWA) problem, that we believe to be critical to the design of wavelength-division-multiplexing (WDM)-based networks. We call this version of the RWA problem *survivable RWA*. In a WDM network, the logical topology is defined by a set of nodes and lightpaths connecting the nodes while the physical topology is defined by the set of nodes and the fiber connecting

Manuscript received March 8, 2001; revised December 20, 2001. This work was supported by the Defense Advanced Research Projects Agency (DARPA) under the Next Generation Internet (NGI) initiative. The work of E. Modiano was supported by the National Science Foundation (NSF) under Grant ANI-0073730.

E. Modiano is with the Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, Cambridge, MA 02139 USA.

A. Narula-Tam is with the Massachusetts Institute of Technology, Lincoln Laboratory, Lexington, MA 02173 USA.

Publisher Item Identifier S 0733-8716(02)05071-0.

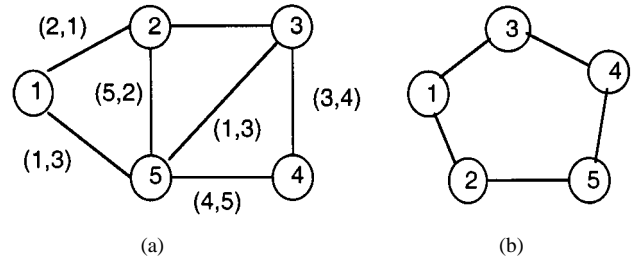


Fig. 1. Survivable routing of a logical topology on a physical topology. (a) Physical topology. (b) Logical topology.

them. Given the logical and physical topologies of the networks, one important question is how to embed the logical topology onto the physical topology. This leads to a static version of the routing and wavelength assignment (RWA) problem. In this version of the problem, the set of lightpaths, defined by the logical topology, is known in advance. In this context various researchers have developed RWA algorithms with the goal of minimizing network costs, including number of wavelengths required, number of wavelength converters, fiber use, etc. [1]. Since with WDM each physical fiber link can support many lightpaths (as many as there are wavelengths on the fiber), once the lightpaths are routed on the physical topology, it is possible (or in fact, likely) that two or more lightpaths would share the same physical link. Hence, the failure of a single physical link, can lead to the failure of multiple links in the logical topology. Since protected logical topologies are often designed to withstand only a single link failure, it is possible that a single physical link failure could leave the logical topology disconnected.

As a simple illustrative example, consider the logical and physical topologies shown in Fig. 1. The logical topology is a ring with nodes ordered 1-3-4-5-2-1. Clearly, such a ring topology is 2-connected, and would remain connected if one of its links failed. The five logical links of this ring can be routed on the physical topology as shown in Fig. 1(a), where each physical link is labeled with the logical link that traverses it. For example logical link (1,3) traverses physical links (1,5) and (5,3). As can be seen from the figure, no physical link supports more than one logical link. Hence, the logical ring would remain protected even in the event of a physical link failure.

Alternatively, had we routed logical link (1,3) on physical links (1,2) and (2,3) the routing would no longer be survivable because physical link (1,2) would have to support both logical links (1,3) and (2,1), hence, its failure would leave the logical

topology disconnected. Furthermore, for many logical topologies, no survivable routings can be found. For example, if the logical topology was a ring with nodes ordered 1-4-2-3-5-1 then it can be easily seen that no routing exists that can withstand a physical link failure. Hence, it is clear that although the logical topology of the network may be connected, once it is embedded on top of a WDM physical network, it may no longer withstand a physical link failure (e.g., fiber cut).

In the context of virtual private networks, the customer might request from the network provider that their lightpaths be routed in such a way that no single physical link failure would leave their VPN disconnected. One simple way to achieve this goal is to route the lightpaths so that no two lightpaths share a physical link. This seemingly simple solution by itself is difficult to obtain. In fact, it was shown in [15] that the related problem of finding disjoint paths for a collection of k source-destination pairs is NP-complete.¹ Furthermore, this simplified solution can be wasteful of resources. For many logical topologies, some of the lightpaths can be routed together while maintaining survivability.

Of course, there has been a significant body of work in the area of optical network protection [2]–[7], [14], [16]. Most previous work in WDM network protection is focused on restoration mechanisms that restore all lightpaths in the event of a physical link failure. Link based restoration recovers from a link failure by restoring the failed physical link, hence, simultaneously restoring all of the associated lightpaths [2], [3], [6]. This is often done using optical loop-back protection [2], [3], [5]. In contrast, path based protection restores each of the lightpaths independently, by finding an alternative end-to-end path for each lightpath [2], [3], [14]. In many cases, it is indeed necessary to restore all failed lightpaths. However, in other cases some level of protection is provided in the electronic layer and restoration at the physical layer may not be necessary. For example, when the electronic layer consists of SONET rings, single link failures can be recovered through loopback protection at the electronic layer. In this case, providing protection at both the optical and electronic layers is somewhat redundant. Another less obvious example is that of packet traffic in the internet where multiple electronic layer paths exist between the source and destination and the Internet protocol (IP) automatically recovers from link failures by rerouting packets.

In such cases, a less stringent requirement may be imposed on the network—for example that the network remain connected in the event of a physical link failure. This approach, of course, is not suitable for all situations. For example, when a network is carrying high priority traffic with quality-of-service (QoS) and protection guarantees, it may still be necessary to provide full restoration. However, when a network is used to support best effort internet traffic, guaranteeing connectivity may suffice. This approach, on which we first reported in [17] and [18], is relatively new in the field of WDM network protection. A similar design goal was considered in [7], where heuristic algorithms were developed in order to minimize the number of source destination pairs that would become disconnected in the event of a

physical link failure. The algorithm in [7] uses tabu search procedures to find disjoint alternate paths for all of the lightpaths.

In this paper, we present a new approach for investigating the problem of routing lightpaths of a logical topology on a given physical topology so that the logical topology can withstand a physical link failure. In Section II, we formulate the problem and establish a necessary and sufficient condition to ensure survivable routing. This condition, leads to some interesting insights into the survivable routing problem and prove that the survivable lightpath routing problem is NP-complete. In Section III, we give an integer linear program (ILP) formulation for the survivable routing problem. We also present low complexity heuristics for the survivable routing problem and compare the performance of these heuristics to that of the full ILP solution. In Section IV, we focus our attention on establishing bidirectional logical rings on the physical topology. The logical ring case leads to a simplified ILP formulation that more easily renders a solution. We also develop necessary conditions on the physical topology for enabling survivable routings for logical rings. Finally, we use our ILP formulation to solve the survivable routing problem for some example networks and compare our results to alternative approaches.

II. PROBLEM FORMULATION

The physical topology of the network consists of a set of nodes $N = \{1 \dots N\}$ and a set of edges E where (i, j) is in E if a link exists between nodes i and j . We assume a bidirectional physical topology, where if (i, j) is in E so is (j, i) . Furthermore, we assume that a failure (cut) of (i, j) will also result in a failure in (j, i) . This assumption stems from the fact that the physical fiber carrying the link from i to j is typically bundled together with that from j to i . Furthermore, in some systems the same fiber is used for communicating in both directions. Finally, we assume that WDM is employed and each physical link (fiber) is capable of supporting W wavelengths in each direction.

The logical topology of the network can be described by a set of logical nodes N_L and logical edges E_L , where N_L is a subset of N and an edge (s, t) is in E_L if both s and t are in N_L and there exists a logical link between them. Given a logical topology, we wish to find a way to route the logical topology on the physical topology such that the logical topology remains connected even in the event of a physical link failure.

In order to route a logical link (s, t) on the physical topology one must find a corresponding lightpath on the physical topology between nodes s and t . Such a lightpath consists of a set of physical links connecting nodes s and t as well as wavelengths along those links. If wavelength changers are available then any wavelength can be used on any link. However, without wavelength changers, the same wavelength must be used along the route. In this paper we assume that either wavelength changers are available or that the number of wavelengths exceeds the number of lightpaths. This assumption allows us to ignore the wavelength continuity constraints and focus only on survivable design.

Let $f_{ij}^{st} = 1$ if logical link (s, t) is routed on physical link (i, j) and zero, otherwise. Now in order to find a routing for the

¹In [15] it was shown that the problem of finding node disjoint paths is NP-complete. This result can be easily extended to link disjoint paths in directed graphs.

logical topology, we must find a route for every logical link (s, t) in E_L . In this paper, we consider bidirectional logical topologies where if $(s, t) \in E_L$ so is (t, s) . Furthermore, we assume that (s, t) and (t, s) follow the same route. That is, if (s, t) traverses physical link (i, j) then (t, s) traverses link (j, i) . For simplicity of notation we describe the logical topology as a set of unordered node pairs representing the bidirectional logical links. Therefore, implicit in finding a route from s to t is also the route from t to s . For simplicity, we present this paper in the context of bidirectional physical and logical topologies; however, it is straightforward to generalize our results to directed topologies.

In this work, we are concerned with finding routings that are survivable. We call a routing *survivable* if the failure of any physical link leaves the (logical) network connected. Of course, a routing cannot possibly be survivable if the underlying logical topology is not redundant. The logical topology is redundant (i.e., 2-connected) if the removal of any logical link does not cause the topology to be disconnected. The following theorems, give some simple yet useful necessary and sufficient conditions for survivability in a network. First, we must define the following notions.

A *cut* is a partition of the set of nodes N into two parts S and $\bar{S} = N - S$. Each cut defines a set of edges consisting of those edges in E with one endpoint in S and the other in $N - S$. We refer to this set of edges as the cut-set associated with the cut $\langle S, N - S \rangle$, or simply $CS(S, N - S)$. Let $|CS(S, N - S)|$ equal the size of the cut-set $\langle S, N - S \rangle$; that is, the number of edges in the cut-set. The following Lemma, also known as Menger's Theorem [12], relates the connectivity of a network to the size of its cut-sets.

Lemma 1: A logical topology with set of nodes N_L and set of edges E_L is 2-connected if and only if every nontrivial cut $\langle S, N_L - S \rangle$ has a corresponding cut-set of size greater than or equal to 2.

Proof: (See [12]) Necessity is due to the fact that if any cut-set consists of only a single link, removal of that link would leave the topology disconnected. Sufficiency is a direct result of the max-flow min-cut theorem. ■

Consider a routing for a logical topology given by the assignment of values to the variables f_{ij}^{st} for all physical links (i, j) and logical links (s, t) , which correspond to the physical links used to route the various logical links. The following Theorem gives a necessary and sufficient condition for a routing of a logical topology to be survivable.

Theorem 1: A routing is survivable if and only if for every cut-set $CS(S, N_L - S)$ of the logical topology the following holds. Let $E(s, t)$ be the set of physical links used by logical link (s, t) , i.e., $E(s, t) = \{(i, j) \in E \text{ for which } f_{ij}^{st} = 1\}$. Then, for every cut-set $CS(S, N_L - S)$,

$$\bigcap_{(s,t) \in CS(S, N_L - S)} E(s, t) = \emptyset$$

Proof: The above condition requires that no single physical link is shared by all logical links belonging to a cut-set of the logical topology. In other words, not all of the logical links belonging to a cut-set can be routed on the same physical link.

This condition must hold for all cut-sets of the logical topology. To prove the theorem, we must show that the above condition is both necessary and sufficient. Necessity is obvious because if there exists a physical link that carries all of the logical links belonging to a cut-set, failure of that link would leave the network disconnected. To see that the condition is also sufficient, notice that the removal of any physical link leaves at least one logical link in each cut-set of the logical topology connected. Hence, the network must still be connected. ■

Notice that it is a direct result of the above theorem that if the logical topology was not redundant then no routing could be survivable. This is because if the logical topology was not redundant then at least one cut-set must exist with size equal to one. The failure of the corresponding link would leave the topology disconnected. Theorem 1 can be generalized to directed topologies, by applying the directed version of Menger's Theorem and considering cut-sets on a directed topology.

Theorem 2: The survivable routing problem is NP-complete.

Proof: The problem is clearly in NP because we can always check that a given routing is survivable in polynomial time by successively removing links and checking for connectedness. To show that the problem is NP-complete we provide a simple transformation from the undirected two commodity integral flow problem ([19], p. 217) to the survivable routing problem. Since the former is a known NP-complete problem, a polynomial time solution for the survivable routing problem would also give a polynomial time solution to the two commodity integral flow problem.

The two commodity integral flow problem considers two source destination pairs (s_1, t_1) and (s_2, t_2) on a graph and an integer traffic requirement R_1 and R_2 between each of the pairs, respectively. Each link in the graph has an integer capacity and the problem is to find a routing for the traffic between the two node pairs without violating the capacity constraint. This problem is known to be NP-complete. An instance of this problem where $R_1 = R_2 = 1$ and $C = 1$ for all links is also known to be NP-complete. This instance of the problem, in fact, corresponds to finding link disjoint paths between the two node pairs. This problem can be easily transformed into an instance of the survivable routing problem as follows. Consider the graph $G = (V, E)$ for the problem and form a new graph G' by adding two nodes s^* and t^* and bidirectional links $\{(s_1, s^*), (s^*, s_2), (t_1, t^*), (t^*, t_2)\}$. Let G' be the physical topology. We form the following ring logical topology, G_L , shown in Fig. 2, consisting of 6 nodes $N_L = \{s_1, s_2, s^*, t_1, t_2, t^*\}$ and bidirectional links $E_L = \{(s_1, t_1), (s_2, t_2), (s_1, s^*), (s^*, s_2), (t_1, t^*), (t^*, t_2)\}$. Now, it is easy to see that finding a survivable routing for G_L on G' , would solve the link disjoint path problem. First notice that no two links of the logical topology can share a physical link, because failure of that link will disconnect the network (see Corollary 1, Section IV). Also note that logical links $(s_1, s^*), (s^*, s_2), (t_1, t^*), (t^*, t_2)$ must be routed using the corresponding physical links since there is no way to get to s^* and t^* without using physical links $(s_1, s^*), (s^*, s_2), (t_1, t^*), (t^*, t_2)$. Hence, the solution to the survivable routing problem must yield disjoint paths between (s_1, t_1) and (s_2, t_2) in the original graph G . ■

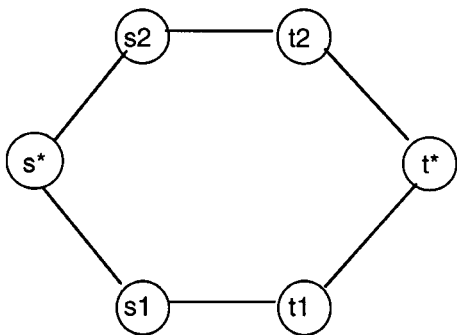


Fig. 2. Logical topology for NP-completeness proof.

III. INTEGER LINEAR PROGRAMMING FORMULATION

Using Theorem 1, we are able to formulate the problem of survivable routing of a logical topology on a given physical topology as an Integer linear program (ILP). Given a physical topology and a corresponding logical topology, we wish to find a way to route the logical topology on the physical topology such that the logical topology remains connected even in the event of a physical link failure.

In order to route a logical link (s, t) on the physical topology one must find a corresponding path on the physical topology between nodes s and t . Such a lightpath consists of a set of physical links connecting nodes s and t as well as wavelengths along those links. Let $f_{ij}^{st} = 1$ if logical link (s, t) is routed on physical link (i, j) and 0 otherwise. Clearly $f_{ij}^{st} > 0$ implies that there exists a physical link between nodes i and j . When the logical links are bidirectional, implicit in finding a route from s to t is also the route from t to s . Using standard network flow formulation finding a route from s to t amounts to routing a unit of flow from node s to node t [10]. This can be expressed by the following set of constraints on the value of the flow variables associated with the logical link (s, t) :

$$\sum_{j \text{ s.t. } (i,j) \in E} f_{ij}^{st} - \sum_{j \text{ s.t. } (j,i) \in E} f_{ji}^{st} = \begin{cases} 1, & \text{if } s = i \\ -1, & \text{if } t = i \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

$\forall i \in N.$

The set of constraints above are flow conservation constraints for routing one unit of traffic from node s to node t . Equation (1) requires that equal amounts of flow due to lightpath (s, t) enter and leave each node that is not the source or destination of (s, t) . Furthermore, node s has an exogenous input of 1 unit of traffic that has to find its way to node t . There are many possible combinations of flow variable values that can satisfy the constraint of (1). Any feasible solution to (1) has a route from s to t embedded in it. It is easy to see that if in addition we required that the path length be minimized (i.e., $\min \sum_{(i,j) \in E} f_{ij}^{st}$ subject to (1)), the solution would also be the unique shortest path ([11], p. 73). Now in order to find a survivable routing for the logical topology, we must find a route for every logical link (s, t) in E_L . Our problem formulation is such that we only need to find the route for logical link (s, t) in one direction; implicit in that route is the route from (t, s) that will follow the same physical links in the opposite direction. Now, using Theorem 1,

the connectivity requirement can be expressed by the following constraint:

$$\forall (i, j) \in E, \quad \forall S \subset N_L, \quad \sum_{(s,t) \in \text{CS}(S, N_L - S)} f_{ij}^{st} + f_{ji}^{st} < |\text{CS}(S, N_L - S)|.$$

The above constraint simply states that for all proper cuts of the logical topology, the number of cut-set links flowing on any given physical link, in either direction², is less than the size of the cut-set. This implies that not all logical links belonging to a cut-set can be carried on a single physical link and immediately satisfies Theorem 1.

If the number of wavelengths on a fiber is limited to W , a capacity constraint can be imposed as follows,

$$\forall (i, j) \in E, \quad \sum_{(s,t) \in E_L} f_{ij}^{st} \leq W.$$

There are a number of objective functions that one can consider. Perhaps the simplest is to find a survivable routing that uses the least capacity. That is, minimize the total number of wavelengths used on all links. An alternative formulation goal may be to minimize the total number of physical links used. Such an approach would lend itself to solutions that maximize physical link sharing by the lightpaths. Here, we focus on the first objective of minimizing total number of wavelengths used and the optimal survivable routing problem can be expressed as the following integer linear program:

$$\text{Minimize} \quad \sum_{\substack{(i,j) \in E \\ (s,t) \in E_L}} f_{ij}^{st}.$$

Subject to:

a) Connectivity constraints: for each pair (s, t) in E_L :

$$\sum_{j \text{ s.t. } (i,j) \in E} f_{ij}^{st} - \sum_{j \text{ s.t. } (j,i) \in E} f_{ji}^{st} = \begin{cases} 1, & \text{if } s = i \\ -1, & \text{if } t = i \\ 0, & \text{otherwise} \end{cases}$$

$\forall i \in N.$

b) Survivability constraints:

$$\forall (i, j) \in E, \quad \forall S \subset N_L, \quad \sum_{(s,t) \in \text{CS}(S, N_L - S)} f_{ij}^{st} + f_{ji}^{st} < |\text{CS}(S, N_L - S)|$$

c) Capacity constraints:

$$\forall (i, j) \in E, \quad \sum_{(s,t) \in E_L} f_{ij}^{st} \leq W$$

d) Integer flow constraints: $f_{ij}^{st} \in \{0, 1\}$

The above ILP can now be solved using a variety of techniques. We implemented this ILP using the CPLEX software package. CPLEX uses branch and bound techniques for solving ILPs and is capable of solving ILPs consisting of up to one million variables and constraints [13]. To illustrate the utility of this

²Again, this is due to the fact that if logical link (s, t) occupies physical link (i, j) in one direction; it also implicitly occupies physical link (j, i) in the other direction.

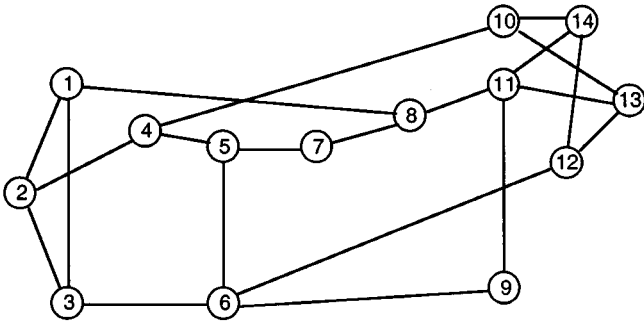


Fig. 3. The 14-node, 21 link NSFNET.

approach, we implemented the ILP for the NSFNET physical topology shown in Fig. 3. We attempted to embed random bidirectional logical topologies of degree-3, 4, and 5, where we define a logical topology of degree k to be logical topology where every node has degree k .

For each, we generated 100 random, 2-connected,³ logical topologies and used the ILP to find optimal survivable routing on the NSFNET. Since we are mainly concerned with the survivable routing, in our implementation we ignored the capacity constraint. Obviously, if needed, the capacity constraints can be easily incorporated into the solution. We also compare our approach to the survivability provided by shortest path routing for the same random logical topologies. In each case we checked to see if the shortest path solution yields a survivable routing. This was accomplished by individually removing each physical link and verifying that the resulting topology remained connected.

Our results are summarized in Table I(a)–(c). Shown in the table are results for both the shortest path solution (labeled Short Path) and the ILP solution (labeled ILP). As can be seen from the table, the ILP was able to find a protected solution for every one of the random logical topologies. In contrast, the shortest path approach resulted in 86 out of 100 of the degree-3 logical topologies being unprotected. With higher degree logical topologies, shortest path was able to protect more of the topologies, still 38 and 27 of the random degree-4 and 5 topologies, respectively, remained unprotected. However, as expected, the ILP solution on average required both more physical links and more total wavelengths (wavelength * links). This difference in link requirements appears to be small and well justified by the added protection that it provides.

Due to the large number of constraints, solving the ILP for large networks can be very difficult. Hence, it is interesting to explore possible relaxations of the ILP formulation that yield survivable routings with reduced complexity. The most obvious relaxation is the linear programming (LP) relaxation where the integer constraints are removed. Unfortunately, however, the LP relaxation leads to many noninteger solutions. Alternatively we can explore relaxations that enforce only some of the survivability constraints. For example, a simple relaxation that applies the survivability constraints only to cuts that include a single node, would prevent a single node from getting disconnected in the event of a fiber cut. With this relaxation, only N survivability constraints are needed and the ILP can be solved easily.

³After generating a random logical topology we first verified that it is at least 2-connected before attempting to find a survivable routing for it.

TABLE I(a)
EMBEDDING RANDOM DEGREE 3 LOGICAL TOPOLOGIES ON THE
NSFNET OF FIG. 3

	Logical Top's	Unprotected solution	Ave. links	Ave. λ *links
ILP	100	0	19.76	46.07
Short Path	100	86	19.31	45.25
Relax - 1	100	10	19.78	46.03
Relax - 2	100	0	19.78	46.07

TABLE I(b)
EMBEDDING RANDOM DEGREE 4 LOGICAL TOPOLOGIES ON THE
NSFNET OF FIG. 3

	Logical Top's	Unprotected solution	Ave. links	Ave. λ *links
ILP	100	0	20.30	60.64
Short Path	100	38	20.17	60.47
Relax - 1	100	0	20.30	60.64
Relax - 2	100	0	20.30	60.64

TABLE I(c)
EMBEDDING RANDOM DEGREE 5 LOGICAL TOPOLOGIES ON THE
NSFNET OF FIG. 3

	Logical Top's	Unprotected solution	Ave. links	Ave. λ *links
ILP	100	0	20.56	75.40
Short Path	100	27	20.48	75.31
Relax - 1	100	0	20.56	75.40
Relax - 2	100	0	20.56	75.40

TABLE I(d)
RUN-TIMES OF ALGORITHMS ON SUN SPARC ULTRA-10

	ILP	Relax-1	Relax-2
Degree - 3	8.3 s	1.3 s	1.3 s
Degree - 4	2 min. 53 sec.	1.5 s	1.5 s
Degree - 5	19 min. 17 sec.	2.0 s	2.0 s

We examined the performance of this rather naïve algorithm and found it to be surprisingly effective. The results of this simple relaxation are shown in Table I (labeled Relax-1). As can be seen from the table, this simple relaxation found a survivable routing for all but ten of the degree-3 logical topologies and all of the degree-4 and 5 topologies. In addition, due to the reduced complexity of this relaxation we were able to run this relaxation for 1000 randomly generated logical topologies of degrees 3, 4, and 5. Again, the relaxation found survivable routings for all but 7% of the degree-3 topologies, and for all 1000 of the degree-4 and 5 logical topologies. This result can be explained by noting that for densely connected networks, the “weakest”

cuts are the single node cuts. Enforcing the survivability constraints for those cuts has a high likelihood of resulting in a survivable routing. This also helps explain the fact that for the degree-3 logical topologies some of the routings found by the relaxation were not survivable, since the degree-3 topology is not as densely connected as degree-4 and 5. Hence, we conclude that this relaxation is effective when the node degrees are large but would probably not be effective for routing degree-2 logical topologies (i.e., ring logical topologies).

A simple extension of this approach enforces the survivability constraints for small cut-sets only. The intuition is that the smallest cut-sets are the most vulnerable and, hence, protecting them will result in a survivable routing with high likelihood. For example, we implemented a relaxation where the survivability constraints were enforced only for cut-sets of size less than or equal to degree of the logical topology plus one. This set of constraints clearly includes all those cuts included in the previous relaxation, since all the single node cuts are of size equal to the degree. With this new relaxation survivable routings were found for all 100 degree-3, 4, and 5 logical topologies; a noticeable improvement over the previous relaxation where survivable routings were not found for ten of the degree-3 logical topologies. Furthermore, when we examined 1000 randomly generated degree-3, 4, and 5 logical topology, the new relaxation found survivable routings for all but three degree-3 logical topologies. The results for this relaxation are labeled Relax-2 on Table I.

In Table I(d), we compare the run-times of the different relaxations when run on a Sun Sparc Ultra-10 computer. Shown in the table are the average run-times for embedding one logical topology. As can be seen from the table, the ILP resulted in relatively large run-times. More importantly, the run-time of the ILP increased dramatically with the degree of the logical topology. While degree-3 logical topologies required only a few seconds, a degree-5 topology required nearly 20 minutes to solve the ILP. Since this is a static design problem, 20 minutes may not be prohibitive. Nonetheless, this dramatic increase in run-time confirms our suspicion that the ILP approach may not scale well to larger networks. In contrast, both relaxations show a dramatic improvement in run-times and, more importantly, the run-times increase only minimally with the degree of the logical topology. This, of course, can be attributed to the fact that the relaxations, for the most part, consider only single node cuts and, hence, only depend on the number of nodes in the topology.

IV. RING LOGICAL TOPOLOGIES

We can gain some additional insight into the survivable routing problem by considering special forms of the logical topology. For example, the ring logical topology, which is the most widely used protected logical topology has a special structure that leads to a simpler problem formulation. In this section, we discuss the special case of embedding ring logical topologies on arbitrary physical topologies.

A unidirectional ring logical topology is an ordered set of nodes $(n_1 \dots n_L)$ where (n_i, n_{i+1}) is in E_L for $0 < i < L$ and (n_L, n_1) is also in E_L . In a bidirectional ring, the reverse connections (n_{i+1}, n_i) and (n_1, n_L) are also in E_L . Since we

focus on protected topologies, here we only consider bidirectional rings. Hence, for simplicity, we assume that all links are bidirectional and refer to the pair of links connecting nodes n_i and n_{i+1} as (n_i, n_{i+1}) . Implied in this notation is that the pair of links between two nodes are treated as a single bidirectional link. Again, as in the previous section we assume that both directions of the logical links utilize the same physical links, but in opposite directions. Recall that a routing of the logical topology is *survivable* if the failure of any physical link leaves the (logical) network connected. The following corollary to Theorem 1 gives a necessary and sufficient condition for a routing of a bidirectional logical ring to be survivable.

Corollary 1: A bidirectional logical ring is survivable if and only if no two logical links share the same physical link.

Proof: It can be easily seen that every cut-set of the ring logical topology contains exactly two links and every pair of logical links share a cut-set, hence, by Theorem 1 no two logical links can share a physical link. ■

Corollary 1 leads to a significant simplification of the survivability constraints. In the general logical topology case the survivability constraints were expressed in terms of constraints on all of the cut-sets (notice that there can be as many as 2^{N-1} such cut-sets). For the ring topology the survivability constraint can be simply replaced by a capacity constraint on the physical links. Specifically, we require

$$\sum_{(s,t) \in E_L} f_{ij}^{st} + \sum_{(s,t) \in E_L} f_{ji}^{st} \leq 1 \\ \forall (i,j) \in E.$$

That is, there can be at most one logical link routed on any given physical link. Note that since the logical links are bidirectional, when route (s,t) uses physical link (i,j) , implicitly it uses the link in both directions. Also, note that since no two lightpaths can share a physical link, the objective of minimizing the total number of physical links and that of minimizing the total number of wavelength * links used are in fact the same (in contrast to the general case where a physical link may be used by multiple logical links). The optimal survivable routing problem for logical rings can be expressed as the following integer linear program:

$$\text{Minimize} \quad \sum_{\substack{(i,j) \in E \\ (s,t) \in E_L}} f_{ij}^{st}$$

Subject to:

a) Connectivity constraints: for each pair (s,t) in E_L :

$$\sum_{j \text{ s.t. } (i,j) \in E} f_{ij}^{st} - \sum_{j \text{ s.t. } (j,i) \in E} f_{ji}^{st} = \begin{cases} 1 & \text{if } s = i \\ -1 & \text{if } t = i \\ 0 & \text{otherwise} \end{cases} \\ \forall i \in N.$$

b) Survivability constraints:

$$\sum_{(s,t) \in E_L} f_{ij}^{st} + \sum_{(s,t) \in E_L} f_{ji}^{st} \leq 1 \\ \forall (i,j) \in E$$

c) Integer flow constraints: $f_{ij}^{st} \in \{0,1\}$.

Again, the above ILP can now be solved using a variety of search techniques. While general ILPs can be rather difficult to solve, this particular ILP is relatively simple. First notice that without the survivability constraints the ILP amounts to solving a shortest path problem. The addition of the survivability constraints makes the solution more difficult to obtain. However, the total number of constraints used is small, relative to the exponential number of constraints used in the general case, hence, the above ILP can be solved very quickly. We were able to solve this ILP using the CPLEX software package running on a SUN SPARC Ultra 10 machine for 10 node rings in less than a second.

A. Necessary Conditions for Survivable Routing

In this section, we develop some necessary conditions on the physical topology to ensure survivable routing of ring logical topologies. Clearly, it is not always possible to route a logical topology on a given physical topology in a manner that preserves the survivability of the logical topology. For example, in the case of a ring, there may be instances where we cannot find disjoint paths for all of the links. In such cases, some of the light-paths will have to share a physical link and the ring would not be survivable. It is interesting to determine under what circumstances it will be possible (or not possible) to find survivable routings. Consider any random ring logical topology. For any cut $\langle S, N-S \rangle$ of the physical topology, let $|CS_P(S, N-S)|$ be the number of physical links along this cut and $|CS_L(S, N-S)|$ be the number of logical links traversing the same cut. Clearly, in order to be able to route the logical links along disjoint physical paths, $|CS_P(S, N-S)|$ must be greater than or equal to $|CS_L(S, N-S)|$. Hence, for a given logical topology one requirement is that for all possible cuts of the physical topology $\langle S, N-S \rangle$, the following must hold:

$$|CS_P(S, N-S)| \geq |CS_L(S, N-S)|.$$

The above condition is necessary, but not sufficient to insure that a survivable routing exists for a particular ring logical topology.

There are situations where one may want to design a physical topology that can support all possible ring logical topologies. One such example may be a service provider that regularly receives requests for ring topologies. Such a service provider may want to design the physical topology of his network so that it can support all possible rings in a survivable manner. Another possible situation is when the logical topology can be dynamically reconfigured [8], [9] for the purpose of load balancing. Here, again, one may want to ensure that the resulting topology can be routed in a survivable manner. The following theorem provides a necessary condition on the physical topology for supporting all possible ring logical topologies in a survivable manner.

Theorem 3: In order for a physical topology to support any possible ring logical topology in a survivable manner the following must hold. For any cut of the physical topology $\langle S, N-S \rangle$

$$|CS_P(S, N-S)| \geq 2 \min(|S|, |N-S|).$$

Theorem 3 says that for all cuts of the physical topology, the number of physical links in the cut set must be greater

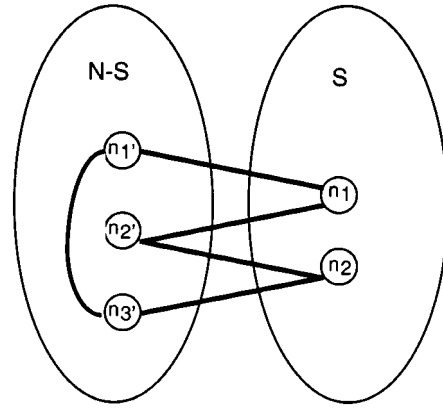


Fig. 4. A logical ring that requires the maximum number of cut-set links.

than or equal to twice the number of nodes on the smaller side of the cut. The condition of Theorem 3 is only a necessary condition. To prove its necessity we must show that there exists a ring logical topology that requires $2 * \min(|S|, |N-S|)$ physical links along the given cut. To show the existence of such a topology we construct the following ring. Suppose without loss of generality that S achieves the minimum of $(|S|, |N-S|)$ and let S contain nodes $n_1 \dots n_s$. Now, construct a logical ring consisting of the following links: $\{(n'_1, n_1), (n_1, n'_2), (n'_2, n_2) \dots (n_s, n'_s), (n'_s, n_1)\}$, where $n_i \in S$ and $n'_i \in (N-S)$. Since $|N-S| \geq |S|$, such a construction always exists. Fig. 4 shows an example where S contains 2 nodes and $|N-S| = 3$. A ring with four links traversing the cut-set is constructed using the above procedure.

Theorem 3 gives necessary conditions for embedding all possible logical rings on a physical topology, including rings of size N . In general, one may want to embed rings of size smaller than N . In this case, the required number of links in the physical topology may be significantly reduced. The following Corollary generalizes Theorem 3 to account for embedding all possible logical rings of size $K \leq N$ in a survivable manner.

Corollary 2: For a physical topology to support any possible K node ring logical topology in a survivable manner the following must hold. For any cut of the physical topology $\langle S, N-S \rangle$

$$|CS_P(S, N-S)| \geq 2 \min(|S|, |N-S|, \lfloor K/2 \rfloor).$$

The $\lfloor K/2 \rfloor$ term accounts for the number of nodes of the logical ring that can be on either side of the cut. Proof of this corollary is essentially identical that of Theorem 3.

Shortest path bound: Another simple yet useful lower bound on the number of links that the physical topology must contain is obtained by observing that each link in the logical topology will use at least as many physical links as would be required if it were routed along the shortest path. Since no two logical links can share a physical link, the number of physical links in the physical topology must obey the following inequality:

$$|E| \geq \sum_{(s,t) \in E_L} |SP(s,t)|.$$

where, $|SP(s,t)|$ is the length, in physical links, of the shortest path from s to t .

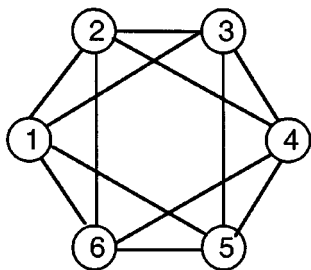


Fig. 5. 6 node degree-4 physical topology.

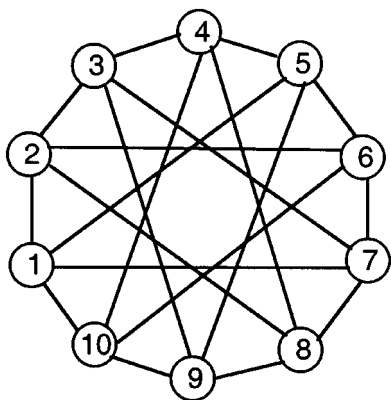


Fig. 6. 10 node degree-4 physical topology.

B. Logical Ring Results

We implemented the ILP for embedding ring logical topologies using the CPLEX software package. We know from the previous section that in order to embed randomly ordered logical rings on a physical topology the physical topology must be densely connected. Hence, for the analysis in this section we consider the 6 and 10 node physical topologies of Figs. 5 and 6. Both of these topologies obey the conditions of Theorem 3 and every node is of degree-4. Furthermore, it can be shown that both topologies are 4-connected. We, therefore, believe that we should be able to find survivable routings for most logical rings.

We attempted to embed all possible 6 and 10 node logical rings on the 6 and 10 node physical topologies. Notice that there are 120 (5!) 6-node logical ring topologies and 362 880 (9!) 10-node logical ring topologies. We used the ring ILP to determine survivable routings for all of these topologies. In addition, we also considered two simple heuristic algorithms for routing the lightpaths.⁴ The *shortest path* solution where each lightpath of the logical topology is routed along the shortest path. Of course, in the case of shortest path, some lightpaths may be routed along the same physical link. In such cases, the shortest path approach would result in an unprotected routing. A somewhat more sophisticated approach is a greedy algorithm that routes lightpaths sequentially using the shortest available path. In order to prevent two lightpaths from sharing a physical link, whenever a physical link is used for routing a lightpath, it is removed from the physical topology so that no other lightpaths can be routed through it. Note that this greedy algorithm

⁴Notice that the relaxations developed in the previous section cannot be applied here because a ring is not densely connected and all cuts are of size 2. Hence, relaxation 1 would be ineffective and relaxation 2 would enforce all cut-set constraints.

TABLE II
EMBEDDING RING LOGICAL TOPOLOGIES ON 6 AND 10 NODE 3-CONNECTED
PHYSICAL TOPOLOGIES

	Logical Top's	No protected solution	Ave. links	Ave. λ *links
6 node-ILP	120	0	7.4	7.4
6 node - SP	120	64 (53%)	6.4	7.2
6 node - GR	120	0	8.1	8.1
10 node-ILP	362880	33760 (9%)	17.8	17.8
10 node - SP	362880	358952 (99%)	11.8	15.5
10 node - GR	362880	221312 (61%)	18.4	N/A

is useful for embedding ring logical topologies since rings require that no two logical links share a physical link. Unfortunately, a similar approach cannot be used to embed arbitrary logical topologies since the connectivity of the logical topology cannot be easily determined by inspecting the routing of individual lightpaths.

Our results are summarized in Table II. For the 6-node physical topology, our ILP was able to find a survivable routing for all 120 logical ring orders. The average number of physical links used to route a logical topology was 7.4. Also, since each physical link supports at most one lightpath, the average number of wavelength * links used was also 7.4. For the 10-node physical topology, our ILP was not able to find a survivable routing for 9.3% of the 362 880 logical ring topologies. When a routing was found, the average number of links used to route a logical topology was 17.8. The greedy algorithm also found a survivable routing for all 6 node logical topologies, but it could not find a survivable routing for 61% of the 10 node rings. With shortest path routing, 53% of the 6-node ring logical topologies were left unprotected and 99% of the 10-node rings were left unprotected. As expected, the ILP was able to protect many more of the logical topologies. Of course, this added protection comes at a price. Shortest path routing used an average of 7.2 wavelengths * links for the 6-node rings and 15.5 wavelengths * links for the 10 node rings, only slightly less than the number of links used by the ILP solution. However, shortest path routing used significantly fewer physical links than the ILP solution. This is, of course, because shortest path routing allows lightpaths to share a physical link, while the ILP does not. Also shown in the table is the number of links used by the greedy algorithm. By definition, the greedy algorithm does not yield a routing when a protected solution is not found, thus, the number of links used can only be calculated when a protected solution is obtained. As expected, the greedy solution used more links than both the ILP and the shortest-path solutions.

Next, we consider the 10-node physical topology of Fig. 6 and attempt to embed random logical ring topologies of various sizes. We attempted to embed 10 000 random logical rings of each size between 5 and 10 nodes. For each ring the set of nodes and their order was chosen at random. Again, we compare the results of our ILP to those obtained using the shortest path routing algorithm and the greedy algorithm. In Fig. 7, we plot the percent of logical topologies for which we failed to obtain a protected routing. As can be seen from the figure, when we

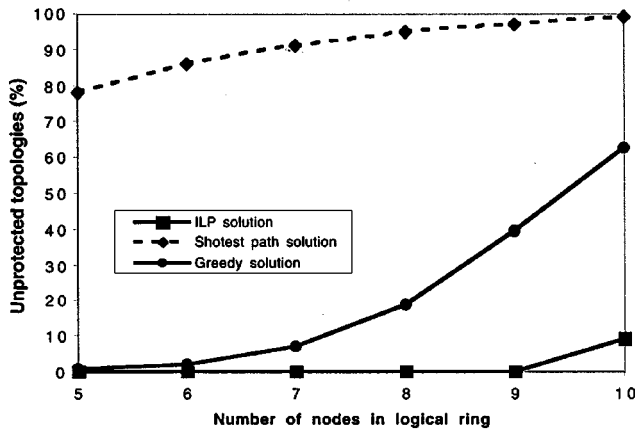


Fig. 7. Fraction of logical ring topologies that cannot be protected on the 10-node physical topology of Fig. 6.

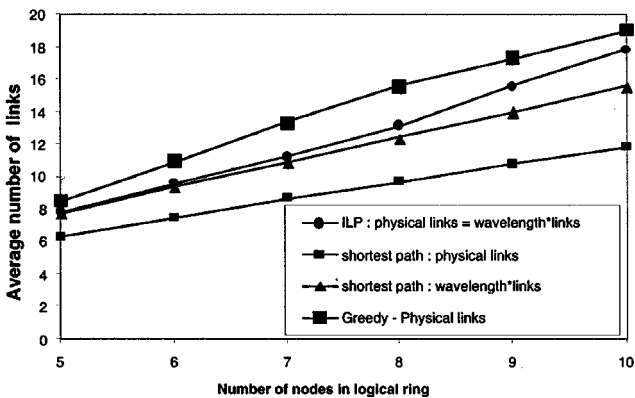


Fig. 8. Average number of links used to embed ring logical topologies on the 10-node physical topology of Fig. 6.

used the ILP we were able to find a protected routing for 100% of the logical rings of size 5 to 9, and fewer than 10% of the 10 node rings were left unprotected. Notice that this latter number is consistent with the results in Table II. However, when shortest path routing was used, the majority of the logical topologies were left unprotected. The greedy approach was able to protect more of the topologies, but not nearly as many as the ILP. In Fig. 8, we plot the average number of physical links used per logical topology. As can be seen from the figure, the shortest path approach indeed uses fewer physical links. However, at a relatively small cost in number of physical links, the ILP solution is able to offer a much greater level of protection. Also, notice that the number of wavelengths * links used with the ILP solution is the same as the number of physical links used. In contrast the shortest path solution uses more wavelength * links than physical links because some physical links were used to support multiple lightpaths. As expected, the greedy approach used the most links. Also, notice that in the case of the greedy approach, the average number of links represents only those topologies for which a protected routing was found. Hence, for those cases the number of physical links is the same as the number of wavelengths * links.

For extremely large topologies, solving the Integer Linear Program may become difficult. Thus, it is interesting to understand what can be obtained from the LP relaxation of the

problem. The LP relaxation will either find: 1) that no solution exists; 2) a solution with integer flows; or 3) a solution with non-integer flows. If the LP relaxation results in no solution, this is a simple way to determine that there is no solution to the ILP either. If the LP relaxation finds an integer solution, then this solution will also be the solution for the ILP. In the third case where the LP relaxation finds a noninteger solution, one must solve the ILP to determine a survivable routing. We solved the LP relaxation for the 6-node and 10-node cases described above to determine the effectiveness of the LP relaxation in solving the integer problem. In the 6-node case, 11.6% of the logical topologies resulted in a noninteger solution. The remaining logical topologies produced integer solutions. In the 10-node case, 97% of the logical topologies for which the ILP was unable to find a survivable routing were also found to be infeasible by the LP relaxation. Unfortunately, 57% of the ring logical topologies produced noninteger solutions to the LP relaxation. As mentioned above, determining a survivable routing for these logical topologies requires solving the ILP.

V. CONCLUSION

This paper considers the problem of embedding protected logical topologies on a WDM physical topology so that the resulting network remains connected in the event of a physical link failure. We proved necessary and sufficient conditions for the survivable routing of the logical topology and used these conditions to develop an ILP formulation for the problem. We used the new ILP formulation to find survivable routings for a variety of network topologies. Our results show that this new formulation is able to offer a much greater degree of protection when compared with shortest path routing. This added protection, of course, comes at the expense of additional network resources. However, it appears from our examples that the additional number of links and wavelengths needed is rather small.

Since solving the ILP for large networks can be difficult, we examined relaxations to the ILP that find survivable routings with reduced complexity. The basic idea behind these relaxations is to enforce only a subset of the cut-set constraints. For example, enforcing the survivability constraints only for single node cuts requires only N constraints rather than the exponential number of constraints required by the cut-set formulation. We found that this approach yields survivable routings with very high probability, especially for densely connected networks. Furthermore, when survivable routings are not found by this "single node cuts" relaxation, additional cut-set constraints can be added until a survivable routing is found. An important direction for future research is to explore alternative relaxations for this problem.

Since this problem is relatively new, many important extensions are possible. For example, this approach can be used to design a network to various degrees of protection. While here we focused on single link failures, multiple failures can be captured in a similar manner. Also, while we focused on minimizing the total number of wavelength * links used, other objective functions, such as total number of physical links used, can also be minimized. Lastly, while here we focused on the survivability constraints only, future work could also consider wavelength limitations and enforcing the wavelength continuity constraint.

REFERENCES

- [1] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Networks Mag.*, pp. 47–60, Jan. 2000.
- [2] S. Ramamuthy and B. Mukherjee, "Survivable WDM mesh networks: Part I—Protection," presented at the Infocom'99, New York, Mar. 1999.
- [3] —, "Survivable WDM mesh networks: Part II—Restoration," presented at the ICC'99, Vancouver, CA, June 1999.
- [4] O. Gerstel, R. Ramaswami, and G. Sasaki, "Fault tolerant multiwavelength optical rings with limited wavelength conversion," presented at the Infocom'97, Kobe, Japan, Apr. 1997.
- [5] M. Medard, S. Finn, and R. A. Barry, "WDM loop-back recovery in mesh networks," presented at the Infocom'99, New York, Mar. 1999.
- [6] A. Fumagalli *et al.*, "Survivable networks based on optimal routing and WDM self-healing rings," presented at the Infocom'99, New York, Mar. 1999.
- [7] O. Crochat and J. Y. Le Boudec, "Design protection for WDM optical networks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1158–1165, Sept. 1998.
- [8] A. Narula-Tam and E. Modiano, "Load balancing algorithms for WDM-based IP networks," presented at the Infocom, 2000, Tel Aviv, Israel, Apr. 2000.
- [9] —, "Dynamic load balancing in WDM packet networks with and without wavelength constraints," *IEEE J. Select. Areas Commun.*, Sept. 2000, to be published.
- [10] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [11] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Englewood Cliffs, NJ: Prentice-Hall, 1982.
- [12] E. Lawler, *Combinatorial Optimization: Networks and Matroids*. New York: Holt, Rinehart and Winston, 1982.
- [13] *ILOG CPLEX 6.5 User's Manual*, 2000.
- [14] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang, "Optical network design and restoration," *Bell Labs Tech. J.*, pp. 58–84, Jan.-Mar. 1999.
- [15] R. M. Karp, "On the computational complexity of combinatorial problems," *Networks*, vol. 5, pp. 45–68, Jan. 1975.
- [16] M. Kodialam and T. V. Lakshman, "Minimum interference routing with application to MPLS traffic engineering," presented at the Infocom, 2000, Tel Aviv, Israel, Apr. 2000.
- [17] E. Modiano and A. Narula-Tam, "Designing survivable networks using effective routing and wavelength assignment (RWA)," presented at the OFC 2001, Anaheim, CA, Mar. 2001.
- [18] —, "Survivable routing of logical topologies in WDM networks," presented at the Infocom 2001, Anchorage, AK, Apr. 2001.
- [19] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman, 1979.



Eytan Modiano received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, CT, in 1986, and the M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, MD, in 1989 and 1992, respectively.

He was a Naval Research Laboratory Fellow between 1987 and 1992, and a National Research Council Postdoctoral Fellow during 1992–1993 while he was conducting research on security and performance issues in distributed network protocols.

Between 1993 and 1999, he was with the Communications Division at Massachusetts Institute of Technology (MIT), Lincoln Laboratory, Cambridge, where he designed communication protocols for satellite, wireless, and optical networks and was the Project Leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999, he has been a member of the faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high speed networks.



Aradhana Narula-Tam received the B.S.E. degree in electrical engineering from the University of Pennsylvania, Philadelphia, in 1990, and the S.M. and Ph.D. degrees, also in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 1992 and 1997, respectively.

From 1997 to 1998, she was in the Research and Advanced Development Group at Motorola Information Systems Group, Mansfield, MA. Since 1998, she has been with the Advanced Networking Group at MIT Lincoln Laboratory. Her research interests include communication theory and optical networks.