Survivable WDM Mesh Networks

S. Ramamurthy, Laxman Sahasrabuddhe, *Member, IEEE*, and Biswanath Mukherjee, *Member, IEEE*

***Abstract*—In a wavelength-division-muliplexing (WDM) optical network, the failure of network elements (e.g., fiber links and cross connects) may cause the failure of several optical channels, thereby leading to large data losses. This study examines different approaches to protect a mesh-based WDM optical network from such failures. These approaches are based on two survivability paradigms: 1) path protection/restoration and 2) link protection/restoration. The study examines the wavelength capacity requirements, and routing and wavelength assignment of primary and backup paths for path and link protection and proposes distributed protocols for path and link restoration. The study also examines the protection-switching time and the restoration time for each of these schemes, and the susceptibility of these schemes to multiple link failures. The numerical results obtained for a representative network topology with random traffic demands demonstrate that there is a tradeoff between the capacity utilization and the susceptibility to multiple link failures. We find that, on one hand, path protection provides significant capacity savings over link protection, and shared protection provides significant savings over dedicated protection; while on the other hand, path protection is more susceptible to multiple link failures than link protection, and shared protection is more susceptible to multiple link failures than dedicated protection.**

We formulate a model of protection-switching times for the different protection schemes based on a fully distributed control network. We propose distributed control protocols for path and link restoration. Numerical results obtained by simulating these protocols indicate that, for a representative network topology, path restoration has a better restoration efficiency than link restoration, and link restoration has a faster restoration time compared with path restoration.

***Index Terms*—Capacity requirement, failure, lightpath, optical network, optimization, protection, protection-switching time, restoration, survivability, wavelength routing, wavelength-division multiplexing (WDM).**

I. Introduction

WAVELENGTH-DIVISION multiplexing (WDM) divides the tremendous bandwidth of a fiber into many nonoverlapping wavelengths (WDM channels) [1], which can be operated at any desirable speed, e.g., peak electronic speed of a few gigabytes per second. An access station may transmit signals on different wavelengths, which are coupled into the fiber using wavelength multiplexers. An a optical cross-connect (OXC) can route an optical signal from an input fiber to an output fiber without performing optoelectronic conversion.

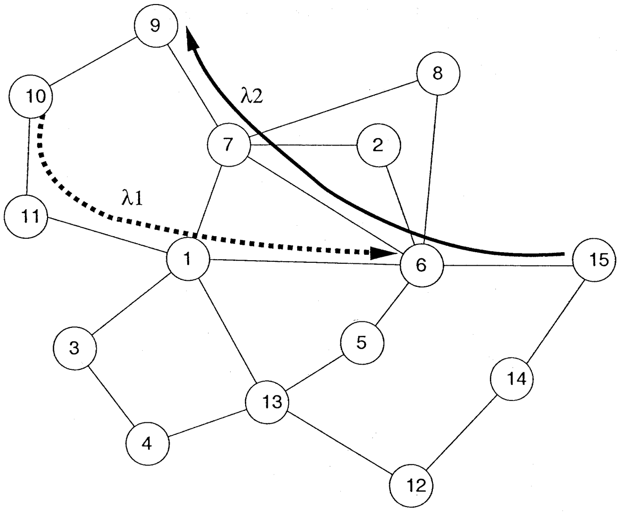

Fig. 1. Architecture of a wavelength-routed optical network.

A wavelength-routed optical network, shown in Fig. 1, consists of OXCs (labeled 1 through 15) interconnected by communication links. Each communication link consists of a pair of unidirectional fiber links. We assume that an access station is connected to each OXC. For clarity of exposition, we will consider the access station/OXC combination as an integrated unit, which we will refer to as a network node. In this work, we assume that all OXCs are wavelength selective, and there is no wavelength conversion in the network; the approaches to accommodate wavelength conversion are relatively straightforward.

In a wavelength-routed network, a connection between a source node and a destination node is called a *lightpath*. A lightpath is an optical channel that may span multiple fiber links to provide an all-optical connection between two nodes. In the absence of wavelength converters, a lightpath would occupy the same wavelength on all fiber links that it traverses. Two lightpaths on a fiber link must be on different wavelength channels to prevent the interference of the optical signals. Fig. 1 shows the following wavelength-continuous lightpaths: a) between Nodes 10 and 6 on wavelength λ_1 and b) between Nodes 15 and 9 on wavelength λ_2. In this work, we assume that all connection requests are unidirectional.

The failure of a network component such as a fiber link can lead to the failure of all the lightpaths that traverse the failed link. Since each lightpath is expected to operate at a rate of several gigabytes per second, a failure can lead to a severe data loss. Although higher protocol layers [such as asynchronous transfer mode (ATM) and Internet protocol (IP)] have recovery procedures to recover from link failures, the recovery time is still sig-

Manuscript received November 27, 2001; revised July 18, 2002. This work was supported in part by the National Science Foundation under Grants NCR-9508239 and ANI-9805285, as well as by Pacific Bell and the UC MICRO Program. Portions of this paper were presented at IEEE INFOCOM'99 (protection) and IEEE ICC '99 (restoration) conferences.

S. Ramamurthy was with the Department of Computer Science, University of California, Davis, CA 95616 USA. He is now with CIENA Corporation, Linthicum, MD 21090-2205 USA.

L. Sahasrabuddhe was with the Department of Computer Science, University of California, Davis, CA 95616 USA. He is now with SBC Communications, Inc., San Ramon, CA USA (e-mail: ls9526@sbc.com).

B. Mukherjee is with the Department of Computer Science, University of California, Davis, CA 95616 USA (e-mail: mukherje@cs.ucdavis.edu).

Digital Object Identifier 10.1109/JLT.2002.806338

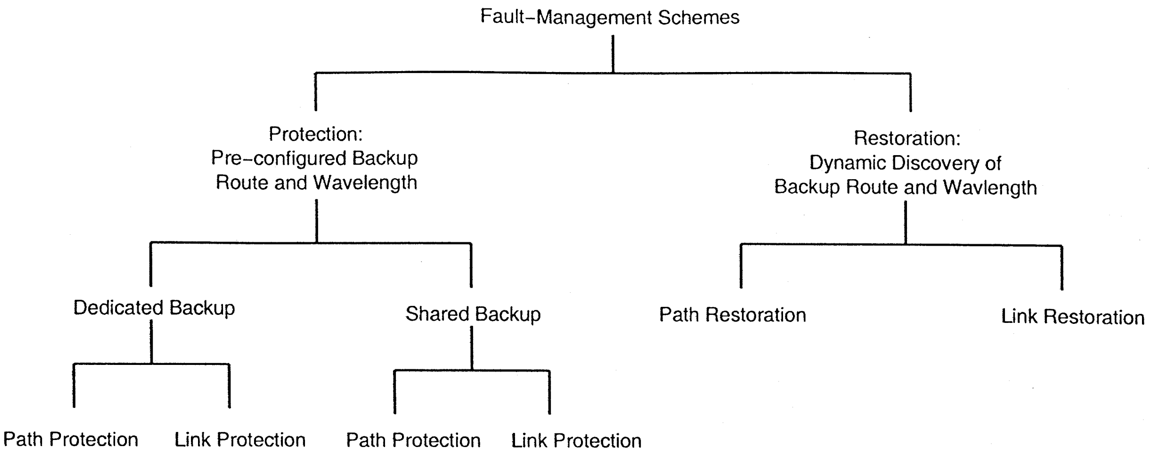

Fig. 2. Different schemes for surviving link failures.

nificantly large (on the order of seconds), whereas we expect that restoration times at the optical layer will be on the order of a few milliseconds to minimize data losses [2]. Furthermore, it is beneficial to consider restoration mechanisms in the optical layer for the following reasons [3]: 1) the optical layer can efficiently multiplex protection resources (such as spare wavelengths and fibers) among several higher layer network applications, and 2) survivability at the optical layer provides protection to higher layer protocols that may not have built-in protection.

There are several approaches to ensure fiber network survivability [4]–[7]. Survivable network architectures are based either on dedicated resources or on dynamic restoration. In dedicated-resource protection (which includes automatic protection switching (APS) and self-healing rings), the network resources may be dedicated for each failure scenario, or the network resources may be *shared* among different failure scenarios. In dynamic restoration, the spare capacity available within the network is utilized for restoring services affected by a failure. Generally, dynamic restoration schemes are more efficient in utilizing capacity due to the multiplexing of the spare-capacity requirements and provide resilience against different kinds of failures, while dedicated-resource protection schemes have a faster restoration time and provide guarantees on the restoration ability.

This study examines different approaches (illustrated in Fig. 2) to survive link failures.[1] These approaches are based on two basic survivability paradigms: 1) path protection/restoration and 2) link protection/restoration.

- Path protection/restoration:

  In path protection, backup resources are reserved during connection setup, while in path restoration, backup routes are discovered dynamically after the link failure. When a link fails [illustrated in Fig. 3(a)], the source node and the destination node of each connection that traverses the failed link are informed about the failure via messages from the nodes adjacent to the failed link, as illustrated in Fig. 4.

  — **Dedicated-path protection:** In dedicated-path protection (also called 1:1 protection), the resources along a backup path are dedicated for only one connection and are not shared with the backup paths for other connections.

  — **Shared-path protection:** In shared-path protection, the resources along a backup path may be shared with other backup paths. As a result, backup channels are multiplexed among different failure scenarios (which are not expected to occur simultaneously), and therefore, shared-path protection is more capacity efficient when compared with dedicated-path protection.

  — **Path restoration:** In path restoration, the source and destination nodes of each connection traversing the failed link participate in a distributed algorithm to dynamically discover an end-to-end backup route. If no routes are available for a broken connection, then the connection is dropped.

- Link protection/restoration:

  In link protection, backup resources are reserved around each link during connection setup, while in link restoration, the end nodes of the failed link dynamically discover a route around the link. In link protection/restoration [illustrated in Fig. 3(b)], all the connections that traverse the failed link are rerouted around that link, and the source and destination nodes of the connections are oblivious to the link failure.

  — **Dedicated-link protection:** In dedicated-link protection, at the time of connection setup, for each link of the primary path, a backup path and wavelength are reserved around that link and are dedicated to that connection. In general, it may not be possible to allocate a dedicated backup path around each link of the primary connection and on the same wavelength as the primary path. For example, Fig. 5 shows a bidirectional ring network with one connection request be-

[1] In this work, we focus primarily on single-link failures, because they are the predominant form of failures in optical networks. Fiber cuts, although rare, must be dealt with effectively. They have been reported to occur with a FIT (failure-in-time: number of failures in 10^9 h) value of approximately 11 000 FIT per 10 km of fiber, in typical telecom networks, i.e., for every 10 km of fiber, a cut is experienced approximately once every 12 years [8]. Time to repair the cuts varies from a few hours to a few days. Thus, we design fault-management techniques to combat single-fiber failures. Although multiple fiber failures are extremely rare, we also evaluate the performance of some of our designs in case of two fiber cuts.

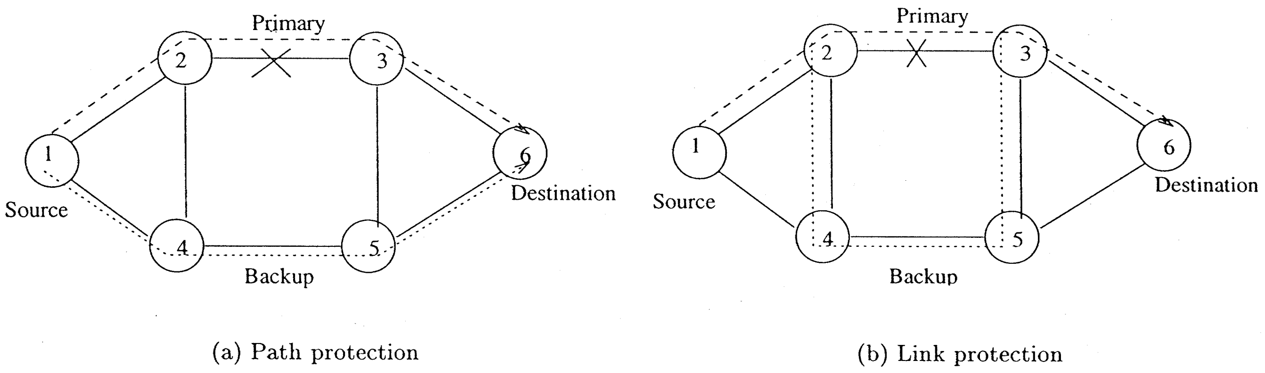

(a) Path protection

(b) Link protection

Fig. 3. Protection schemes.

tween Node 1 and Node 5. The backup path around link (2,3), viz. (2,1,8,7,6,5,4,3), and the backup path around link (3,4), viz. (3,2,1,8,7,6,5,4), share links in common and hence cannot be dedicated the same wavelength.[2] Since our experience indicates that dedicated-link protection utilizes wavelengths very inefficiently, we will not consider dedicated-link protection in this work.

- **Shared-link protection:** In shared-link protection, the backup resources reserved along the backup path may be shared with other backup paths. As a result, backup channels are multiplexed among different failure scenarios (which are not expected to occur simultaneously), and therefore shared-link protection is more capacity-efficient when compared with dedicated-link protection.
- **Link restoration:** In link restoration, the end nodes of the failed link participate in a distributed algorithm to dynamically discover a route around the link. If no routes are available for a broken connection, then the connection is dropped.

In this paper, we investigate the wavelength capacity requirements, routing and wavelength assignment of primary and backup paths, and protection-switching time requirements for path- and link-protection schemes. We also propose distributed protocols and study the restoration time requirements for path- and link-restoration schemes.

The design of a survivable optical network has been studied in [2], [3], [6], [7], and [9]–[17]. The work in [2] and [3] addresses the issues in designing a survivable optical layer. In [6] and [7], the authors examine 1 + 1 protection and other WDM network architectures with optical protection. In [9], the authors propose physical protection schemes and a path-restoration scheme based on 1 + 1 protection. The work in [10] considers different approaches for fault-tolerant design of optical ring networks. In [11], the authors propose an algorithm that protects optical mesh networks from link and node failures. In [12], the authors propose analytical methods to estimate capacity utilization in optical networks that are resilient against single-link failures. The work in [13] compares different approaches for restoring lightpaths. In [14], the authors review various protection and restoration techniques in an IP-over-WDM network. The work in [15] compares different fault-management techniques in an IP-over-WDM network. The work in [16] presents efficient routing algorithms for computing the primary and backup routes in an optical network. In [17], the authors present connection management algorithms for a survivable WDM optical network.

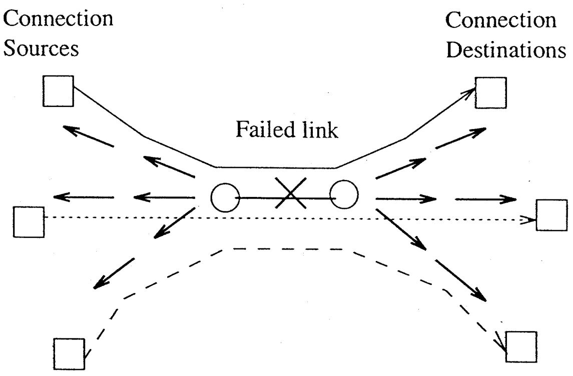

Fig. 4. Link-fail messages sent to all source and destination nodes of connections traversing a failed link.

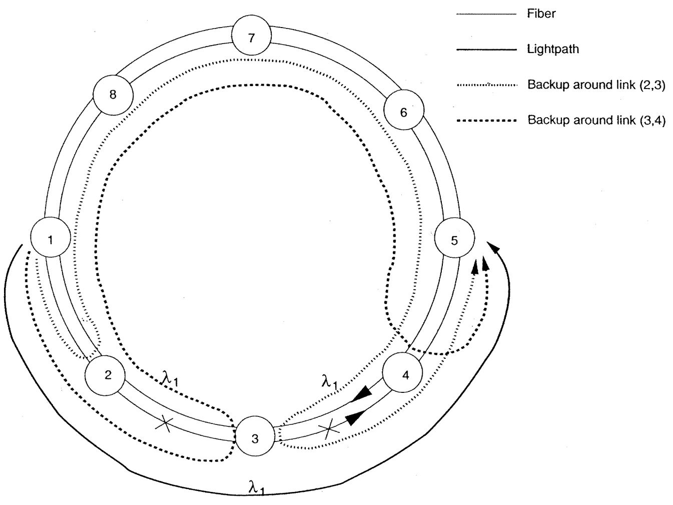

Fig. 5. Illustrative example showing that dedicated-link protection is not possible in a bidirectional ring network.

[2] If wavelength converters are present in the network, then dedicated-link protection is possible by having the backup paths around (2,3) and (3,4) on different wavelengths.

Path- and link-restoration schemes have been extensively researched in circuit-switched transport networks [4], [5], [18]–[20]. In [18], the authors report that path restoration provides about 19% improvement in spare-capacity utilization over link restoration in circuit-switched transport networks. Distributed protocols for restoration have been extensively researched in circuit-switched transport networks [21]–[24] and in ATM networks [5], [20]. Our study borrows appropriate techniques from previous work, develops new techniques, and applies them to the optical network setting.

In Section II, we develop Integer Linear Program (ILP) formulations for the routing and wavelength assignment problem, and wavelength utilization for a static traffic demand, for each of the different protection schemes. Section III presents numerical results for wavelength utilization, and protection-switching times on a representative network topology for different protection schemes. Section IV presents distributed restoration protocols for path restoration and link restoration and numerical results for restoration times on a representative network topology. Section V concludes this work with a discussion of its main contributions.

II. Problem Formulation

In this section, we develop ILP formulations of path- and link-protection schemes to protect against single-link failures. We assume that the network topology and a demand matrix (consisting of the number of connections to be established between each node pair) are given. We assume that the set of alternate routes[3] (which are used to satisfy any demand) between each node pair can be precomputed or is given. Our objective is to minimize the total number of wavelengths used on all the links in the network (for both the primary paths and backup paths). The ILP solution also determines the routing and wavelength assignment of the primary and backup paths. Generally, capacity efficiency can be measured in two ways: 1) given a certain capacity, maximize the protected carried demand [9], or 2) given a certain demand and given a 100% restoration requirement, minimize the total capacity used. In our formulations, we require that all demands should be protected, and we minimize the total capacity used. ILPs 1, 2, and 3 minimize the capacity utilizations for dedicated-path protection, shared-path protection, and shared-link protection, respectively.

A. Notation

We define the notation employed to formulate the ILPs. We are given the following: 1) the network topology represented as a directed graph G, 2) a demand matrix, i.e., the number of lightpath requests between node pairs, and 3) alternate routing tables at each node. Also given are the following.

- N: Nodes in the network (numbered 1 through N). (Node pairs are numbered 1 through $N \times (N-1)$.)
- E: Links in the network (numbered 1 through E).
- W: Number of wavelengths on a link.
- R^i: Set of alternate routes for node pair i.
- $M^i = |R^i|$: Number of alternate routes between node pair i. Let M be the maximum number of alternate routes between any node pair, i.e., $M = \max_i M^i$.
- R^i_j: Set of eligible alternate routes between node pair i after link j fails.
- End nodes(j): The set of alternate routes between the node pair adjacent to link j.
- d_i: Demand for node pair i, in terms of number of connection requests. (Each connection requires the bandwidth of a full wavelength channel.)

We require the ILPs to solve for the following variables.

- w_j: Number of wavelengths used by primary lightpaths on link j.
- s_j: Number of *spare* wavelengths used on link j.
- $\gamma^{i,r}_w$ takes on the value of 1 if the r^{th} route between node pair i utilizes wavelength w before any link failure; 0 otherwise. These variables are employed in all ILPs.
- $\alpha^{i,b}_{w,p}$ takes on the value of 1 if the dedicated backup route b on wavelength w is employed for protecting a primary route p between node pair i; 0 otherwise. These variables are employed only in ILP1.
- $\delta^{i,b}_{w,p}$ takes on the value of 1 if the shared backup route b on wavelength w is employed for protecting a primary route p between node pair i; 0 otherwise. These variables are employed only in ILP2.
- $l^{j,r}_w$ takes on the value of 1 if wavelength w is utilized on restoration route r between the node pair that is adjacent to j, when link j breaks; 0 otherwise. These variables are employed only in ILP3.
- m^j_w takes on the value of 1 if wavelength w is utilized by some restoration route r that traverses link j; 0 otherwise. These variables are employed in ILP2 and ILP3.

B. ILP Formulations

1) ILP1: Dedicated-Path Protection: Minimize the total capacity used

$$\text{Minimize} \sum_{j=1}^{E} (w_j + s_j). \qquad (1)$$

The number of lightpaths on each link is bounded

$$(w_j + s_j) \leq W, \qquad 1 \leq j \leq E. \qquad (2)$$

The demand between each node pair i is satisfied as

$$d^i = \sum_{r=1}^{M_1} \sum_{w=1}^{W} \gamma^{i,r}_w, \qquad 1 \leq i \leq N(N-1). \qquad (3)$$

The number of primary lightpaths traversing link j is written as

$$w_j = \sum_{i=1}^{N(N-1)} \sum_{r \in R^i, j \in r} \sum_{w=1}^{W} \gamma^{i,r}_w, \qquad 1 \leq j \leq E. \qquad (4)$$

[3] In alternate routing, each network node has a routing table that contains a list of a limited number of fixed routes to each destination node. The list of routes can be based on one or more constraints, such as shortest path, shared-link risk groups (SLRGs) [25], etc. A connection request arriving at the node utilizes one of the routes to the destination node from the list of available routes.

The number of spare channels utilized for link j is written as

$$s_j = \sum_{i=1}^{N(N-1)} \sum_{b\in R^i, j\in b} \sum_{p\in R^i, p\neq b} \sum_{w=1}^{W} \alpha_{w,p}^{i,b}, \qquad 1 \le j \le E. \tag{5}$$

The wavelength-continuity constraint, i.e., only one primary or backup lightpath, can use wavelength w on link j, written as

$$\sum_{i=1}^{N(N-1)} \sum_{r\in R^i : j\in r} \gamma_w^{i,r} + \sum_{i=1}^{N(N-1)} \sum_{b\in R^i : j\in b} \sum_{p\in R^i, p\neq b} \alpha_{w,b}^{i,p} \le 1, \quad 1\le w\le W, \quad 1\le j\le E. \tag{6}$$

Due to a link failure, if route p fails between node pair i, then the demand between node pair i should still be satisfied as

$$\sum_{w=1}^{W} \gamma_w^{i,p} = \sum_{b\in R^i, b\neq p} \sum_{w=1}^{W} \alpha_{w,p}^{i,b}, \qquad p\in R^i, 1\le i\le N(N-1). \tag{7}$$

2) ILP2: Shared-Path Protection: Minimize the total capacity used, written as

$$\text{Minimize} \sum_{j=1}^{E} (w_j + s_j). \tag{8}$$

The number of channels on each link is bounded, written as

$$w_j + s_j \le W, \qquad 1\le j\le E. \tag{9}$$

Demand between each node pair is satisfied, written as

$$\sum_{r=1}^{M_i} \sum_{w=1}^{W} \gamma_w^{i,r} = d^i, \qquad 1\le i\le N(N-1). \tag{10}$$

The definition of the number of primary lightpaths traversing a link is

$$w_j = \sum_{i=1}^{N(N-1)} \sum_{r\in R^i, j\in r} \sum_{w=1}^{W} \gamma_w^{i,r}, \qquad 1\le j\le E. \tag{11}$$

The definition of the spare capacity required on link k is

$$s_k = \sum_{w=1}^{W} m_w^k, \qquad 1\le k\le E. \tag{12}$$

Constraints to indicate whether wavelength w is reserved for some restoration path on link k are

$$m_k^w \le \sum_{i=1}^{N(N-1)} \sum_{p,b\in R^i, k\in b} \delta_{w,p}^{i,b}, \qquad 1\le k\le E, 1\le w\le W \tag{13}$$

and

$$N(N-1)\times E\times M\times m_k^w \ge \sum_{i=1}^{N(N-1)} \sum_{p,b\in R^i, k\in b} \delta_{w,p}^{i,b}, \quad 1\le k\le E, \ 1\le w\le W. \tag{14}$$

The wavelength-continuity constraint, i.e., only one primary or backup lightpath can use wavelength w on link j, is written as

$$\left(\sum_{i=1}^{N(N-1)} \sum_{r\in R^i : j\in r} \gamma_w^{i,r} \right) + m_w^j \le 1, \quad 1\le j\le E, 1\le w\le W. \tag{15}$$

Constraints to ensure that two backup lightpaths can share wavelength w on link k only if the corresponding primary paths are fiber disjoint are written as

$$\sum_{i=1}^{N(N-1)} \sum_{p\in R^i : f\in p} \sum_{b\in R^i : k\in b} \delta_{w,p}^{i,b} \le 1 \quad 1\le f\le E, 1\le k\le E, 1\le w\le W. \tag{16}$$

The constraints to ensure that every primary lightpath is protected by a back-up lightpath are written as

$$\sum_{w=1}^{W} \gamma_w^{i,p} = \sum_{b\in R^i, b\neq p} \sum_{w=1}^{W} \delta_{w,p}^{i,b} \quad 1\le i\le N(N-1), \forall p\in R^i, 1\le w\le W. \tag{17}$$

3) ILP3: Shared-Link Protection: The total capacity used should be minimized as

$$\text{Minimize} \sum_{j=1}^{E} (w_j + s_j). \tag{18}$$

The number of lightpaths on each link is bounded as

$$s_j + w_j \le W, \qquad 1\le j\le E. \tag{19}$$

The demand between each node pair i is satisfied mas

$$\sum_{r=1}^{M_i} \sum_{w=1}^{W} \gamma_w^{i,r} = d^i, \qquad 1\le i\le N(N-1). \tag{20}$$

The definition of the number of primary lightpaths traversing each link is

$$\sum_{i=1}^{N(N-1)} \sum_{r\in R^i, j\in r} \sum_{w=1}^{W} \gamma_w^{i,r} = w_j, \qquad 1\le j\le E. \tag{21}$$

The definition of the spare capacity required on link k is

$$s_k = \sum_{w=1}^{W} m_w^k, \qquad 1\le k\le E. \tag{22}$$

The constraints indicating if wavelength w is utilized for some restoration path on link j are written as

$$m_w^j \le \sum_{k=1}^{E} \sum_{r\in \text{endnodes}(k), j\in r} l_w^{k,r}, \quad 1\le j\le E, 1\le w\le W \tag{23}$$

and

$$EMm_w^j \ge \sum_{k=1}^{E} \sum_{r\in \text{endnodes}(k), j\in r} l_w^{k,r}, \quad 1\le j\le E, 1\le w\le W. \tag{24}$$

The wavelength-continuity constraint, i.e., only one primary or link restoration lightpath can use a wavelength w on link j, is written as

$$\left(\sum_{i=1}^{N(N-1)} \sum_{r\in R^i:j\in r} \gamma_w^{i,r}\right) + m_w^j \le 1, \qquad 1\le j\le E,\ 1\le w\le W. \tag{25}$$

The link restoration demands are met after link j fails for each wavelength w, written as

$$\sum_{r\in \mathrm{endnodes}(j)} l_w^{j,r} = \sum_{i=1}^{N(N-1)} \sum_{r\in R^i:j\in r} \gamma_w^{i,r}, \qquad 1\le w\le W,\quad 1\le j\le E. \tag{26}$$

C. Example ILP Solutions

In this section, we present examples carried out to illustrate the problems and understand the solutions provided by the ILPs. Consider the network in Fig. 1. Assume that the demand consists of two connections: a) the first from Node 10 to Node 6 and b) the second from Node 15 to Node 9. The routes and wavelengths of primary and backup lightpaths for dedicated-path protection (as solved by ILP1) are shown in Table I. The total capacity utilization of this solution is 16 wavelength links (where one wavelength link is a wavelength used on a link): six wavelength links for the primary lightpaths, and ten wavelength links for the backup lightpaths.

The routes and wavelength assignments for the primary and backup lightpaths as produced by the shared-path protection in ILP2 are shown in Table II. We note that this solution utilizes six wavelength links for primary paths and nine wavelength links for backup paths for a total of 15 wavelength links. We note that the two working lightpaths—(10,11,1,6) and (15,6,7,9)—are link-disjoint. As a result, upon any single-link failure, at most one of the two lightpaths can fail, i.e., both lightpaths cannot fail simultaneously upon any single-link failure. Therefore, the backup lightpaths can share wavelengths, since they will not be activated simultaneously. This observation leads to the routes and wavelength assignments for the working and backup lightpaths shown in Table II. We note that wavelength λ_1 is shared by *both* of the backup routes on link (10,9).

The routes and wavelength assignments for the primary and backup lightpaths as produced by ILP3 (shared-link protection) are shown in Table III. The solution utilizes a total of 24 wavelength links comprised of six wavelength links for primary paths and 18 wavelength links for backup paths. We note that backup wavelength links are not dedicated, and hence, for example, the wavelength λ_2 is shared by backup paths on links (10,11), (11,1), and (1,7).

D. Solution Approach

The routing and wavelength assignment (RWA) problem (with no protection for any demands) has been shown to be NP-complete [26]. We anticipate that the problems formulated in ILP's 1–3 are NP-complete as well. We utilized the CPLEX 6.5 software package to solve the instances of the ILPs generated for a representative network topology. We note that the number of variables and the number of equations for the ILPs grow rapidly with the size of the network; therefore, the ILP formulations are practical only for small networks (a few tens of nodes). For larger networks (a few hundreds of nodes), we need to employ heuristic methods.

TABLE I
ROUTES AND WAVELENGTHS OF PRIMARY AND BACKUP LIGHTPATHS UNDER DEDICATED-PATH PROTECTION

Connection	*Primary Lightpath*	*Backup Lightpath*
10 → 6	(10,11,1,6) on λ_1	(10,9,7,6) on λ_1
15 → 9	(15,6,7,9) on λ_1	(15,14,12,13,1,11,10,9) on λ_2

TABLE II
ROUTES AND WAVELENGTHS OF PRIMARY AND BACKUP LIGHTPATHS UNDER SHARED-PATH PROTECTION

Connection	*Primary Lightpath*	*Backup Lightpath*
10 → 6	(10,11,1,6) on λ_1	(10,9,7,6) on λ_1
15 → 9	(15,6,7,9) on λ_2	(15,14,12,13,1,11,10,9) on λ_1

TABLE III
ROUTES AND WAVELENGTHS OF PRIMARY AND BACKUP LIGHTPATHS UNDER SHARED-LINK PROTECTION

Connection	*Primary Lightpath*	*Failed link*	*Restoration Lightpath*
10 → 6	(10,9,7,6) on λ_2		
15 → 9	(15,6,7,9) on λ_1		
		(6,7)	(6,1,7) on λ_1
		(7,6)	(7,8,6) on λ_2
		(7,9)	(7,1,11,10,9) on λ_1
		(9,7)	(9,10,11,1,7) on λ_2
		(10,9)	(10,11,1,7,9) on λ_2
		(15,6)	(15,14,12,13,5,6) on λ_1

III. ILLUSTRATIVE EXAMPLES AND DISCUSSION

We performed our studies on an example 16-wavelength network of interconnected rings shown in Fig. 1. This topology was chosen to be representative of typical mesh topologies employed in telecommunications networks. We chose a set of four alternate routes between each node pair, ensuring that all link-disjoint routes between the node pair are included. For this network topology, we ran ILPs 1–3 on random demands, where each random demand had between 10 and 35 connection requests.

A. Results

We tabulate the results from our ILPs for the interconnected-rings network in Table IV. The first column indicates the number of connections in the demand matrix. The second column indicates the capacity utilization of the optimal routing and wavelength assignment of the lightpaths obtained from the RWA ILP formulation without any protection [26]. The third, fourth, and fifth columns indicate the capacity utilization for ILPs 1–3, respectively.[4] Note that shared-path protection utilizes the network capacity more efficiently than the other two protection schemes.

If the connections in the network are protected by employing path- or link-protection schemes, then no connections will be

[4] Numbers that are asterisked indicate the best solution reported by CPLEX running for 10 h on an otherwise unloaded 1-GHz Pentium-4 Linux workstation.

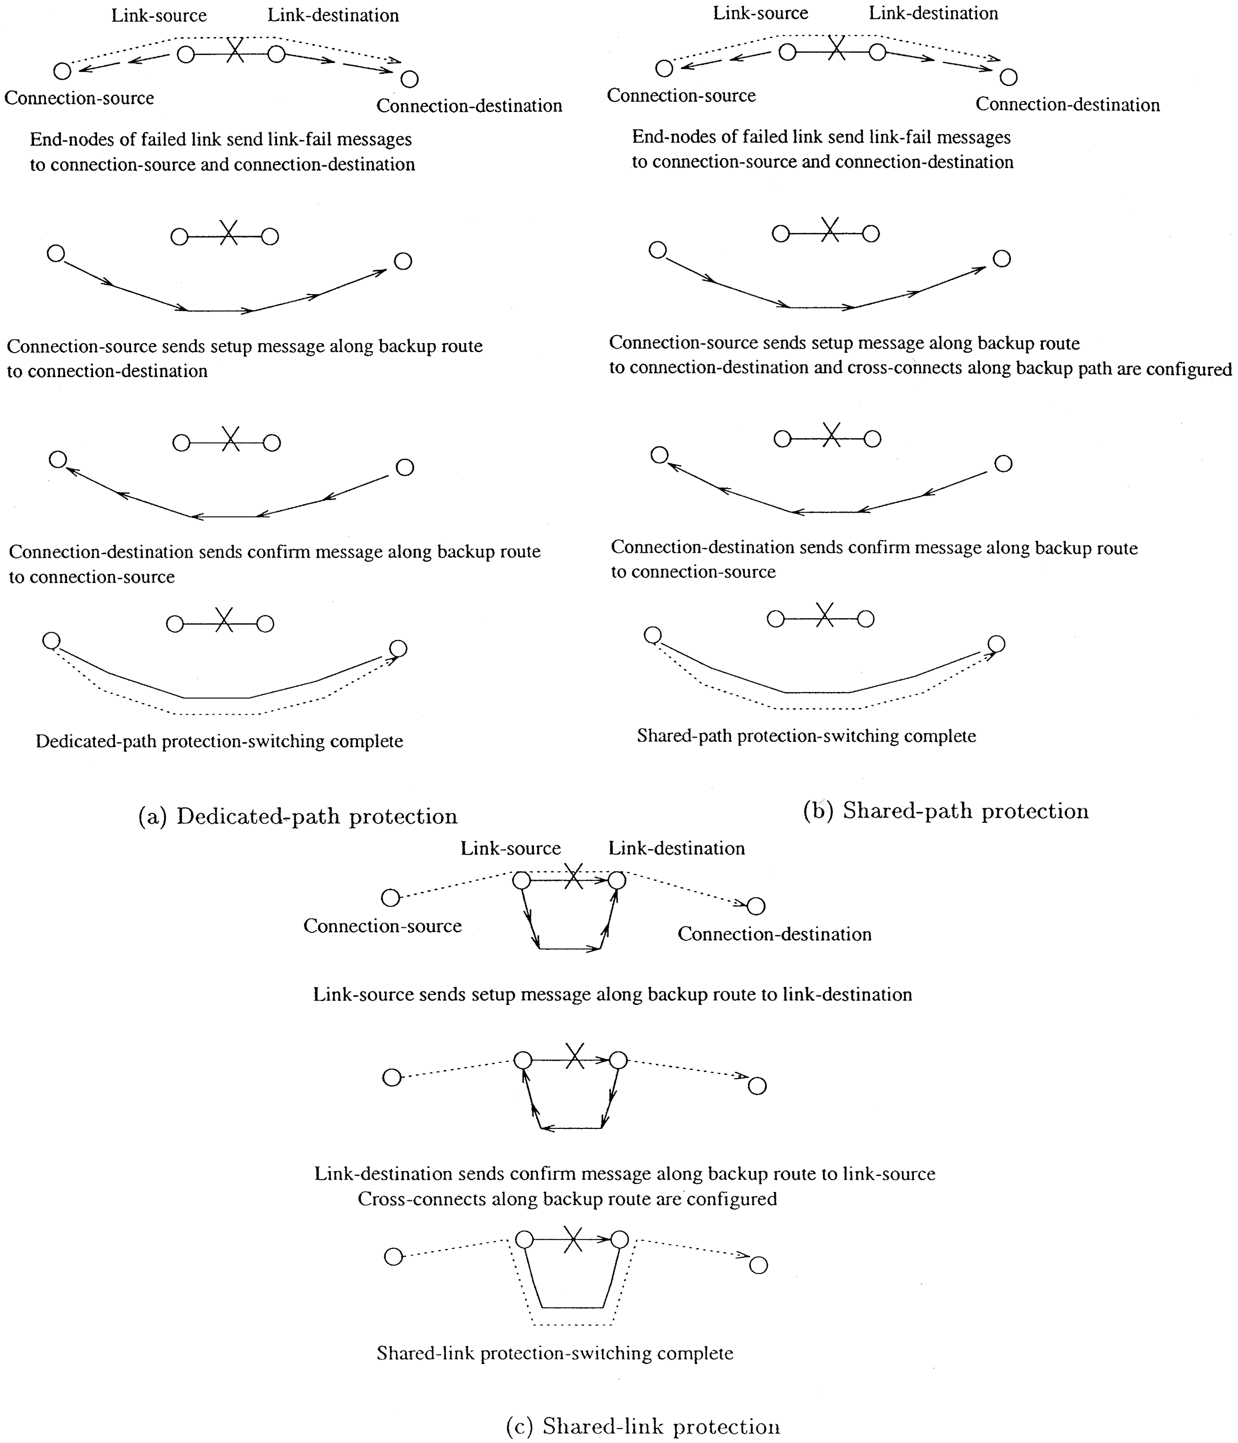
Link-source
Link-destination
Connection-source
Connection-destination
End-nodes of failed link send link-fail messages to connection-source and connection-destination
Connection-source sends setup message along backup route to connection-destination
Connection-destination sends confirm message along backup route to connection-source
Dedicated-path protection-switching complete
Link-source
Link-destination
Connection-source
Connection-destination
End-nodes of failed link send link-fail messages to connection-source and connection-destination
Connection-source sends setup message along backup route to connection-destination and cross-connects along backup path are configured
Connection-destination sends confirm message along backup route to connection-source
Shared-path protection-switching complete
Link-source
Link-destination
Connection-source
Connection-destination
Link-source sends setup message along backup route to link-destination
Link-destination sends confirm message along backup route to link-source
Cross-connects along backup route are configured
Shared-link protection-switching complete

(a) Dedicated-path protection

(b) Shared-path protection

(c) Shared-link protection

Fig. 6. Illustration of protection-switching procedure for each protection scheme. (a) Dedicated-path protection. (b) Shared-path protection. (c) Shared-link protection.

dropped under a single-link failure scenario, but some connections may be dropped if more than one link fails at the same time. Table V shows the average number of connections that are dropped when two links fail in the network simultaneously. The average is computed over all possible two-link failure scenarios. The first column indicates the number of connections in the demand matrix, and the second, third, and fourth columns indicate the average number of connections that are dropped when two links fail for dedicated-path, shared-path, and shared-link protection, respectively.

In summary, for our network topology, and for the random traffic demands we considered, shared-path protection provides significant savings in capacity utilization over dedicated-path and shared-link protection schemes, and dedicated-path protection provides marginal savings in capacity utilization over shared-link protection. On the other hand, shared-path

TABLE IV
CAPACITY UTILIZATION (IN WAVELENGTH LINKS) FOR PATH- AND LINK-PROTECTION SCHEMES FOR THE NETWORK IN FIG. 1, WITH RANDOM TRAFFIC DEMAND

Number of Connections	No Protection	Dedicated path	Shared path	Shared link
10	29	69	63*	86*
15	31	80	66*	89*
20	45	119	95*	120*
25	59	165	136*	192*
30	74	200	168*	229*
35	84	233	196*	249*

TABLE V
AVERAGE NUMBER OF CONNECTIONS THAT ARE DROPPED WHEN TWO SIMULTANEOUS LINK FAILURES OCCUR FOR PATH- AND LINK- PROTECTION SCHEMES FOR THE NETWORK IN FIG. 1, WITH RANDOM TRAFFIC DEMAND

Number of Connections	Dedicated path	Shared path	Shared link
10	0.189	0.221	0.137
15	0.118	0.214	0.142
20	0.285	0.417	0.235
25	0.392	0.527	0.333
30	0.57	0.729	0.473
35	0.609	0.815	0.484

protection is a little more susceptible to two-link failures than dedicated-path and shared-link protection schemes, and dedicated-path protection is more susceptible to two-link failures than shared-link protection.

B. Protection-Switching Time

The time taken from the instant a link fails to the instant the backup path of a connection traversing the failed link is enabled is defined to be the protection-switching time for the connection. In this section, we shall estimate the protection-switching times for the different protection schemes. We assume that a link failure is detected by the network nodes adjacent to the link, and that all network nodes participate in a distributed protocol outlined below to perform protection switching. We also assume that the control network is reliable, i.e., does not incur message losses, and is fully distributed, and we assume that the transmission time for control messages can be neglected in comparison to the link propagation delay.[5] In our calculations, we employ "typical" values for the various parameters, such as the propagation delay, fault-detection time, switch-configuration time, etc., which, to the best of our knowledge, are representative of emerging network technologies. (We remark that the values for some of these terms can change and evolve as various component technologies continue to mature.) Several assumptions were found.

- The message-processing time at a node D is 10 μs, corresponding to the execution of 10 000 instructions on a 1-GHz CPU. The queuing delays of control messages at a node are assumed to be included in the message-processing time.
- The propagation delay on each link P is 400 μs, corresponding to a link length of 80 km.
- The time to configure, test, and set up an OXC is C. Since we do not have a good estimate of C at this time, we will study the impact of C on the protection-switching time by allowing it to take on values of 10 ns, 10 μs, 500 μs, and 10 ms.
- The time to detect a link failure is F. Our estimate of F is 10 μs, which is based on the feedback received from experts on the subject.
- The number of hops from the link source[6] to the source node of the connection is n.
- In path (link) protection, m is equal to the number of hops in the backup route from the source (link-source) node to the destination (link-destination) node.

Under these assumptions, we outline the protection-switching procedures for the different schemes as follows.

- *Dedicated-path protection*: Fig. 6(a) illustrates the steps in the protection-switching procedure for dedicated-path protection. First, the end nodes of the failed link, upon detecting a link failure, send *link-fail* messages to the source node and the destination node of the connection. Then, the source node sends a *setup* message to the destination node along the backup route (which is determined in advance at the time of connection setup). The destination node, upon receiving the setup message, sends a *confirm* message back to the source node, thus completing the protection-switching procedure. The total time for dedicated-path protection-switching is

$$F + n \times P + (n+1) \times D + 2 \times m \times P \\ +2 \times (m+1) \times D.$$

  We note that the OXCs along the backup path are configured at the time of the connection setup and hence do not need to be configured during the protection-switching procedure.
- *Shared-path protection*: Fig. 6(b) illustrates the steps in the protection-switching procedure for shared-path protection. First, the end nodes of the failed link, upon detecting a link failure, send *link-fail* messages to the source node and the destination node of the connection. Then, the source node sends a *setup message* to the destination node along the backup route (which is determined in advance at the time of connection setup) and configures the OXCs at each intermediate node along the backup path (in shared protection, at the time of connection setup, wavelengths are reserved in advance for backup paths but OXCs are not configured to allow for sharing of backup wavelengths). The destination node, upon receiving the setup message, sends a *confirm* message back to the source node, thus completing the protection-switching procedure. The total time for shared-path protection switching is

$$F + n \times P + (n+1) \times D + (m+1) \times C + 2 \times m \times P \\ +2 \times (m+1) \times D.$$

[5] Since the size of a control message is expected to be at most a few thousand bits, and since the transmission rate on a wavelength channel is expected to be at least a few gigabytes per second, we expect the transmission time for a control message to be at most a few microseconds. On the other hand, since the length of a typical link in a telecom network can easily be a few tens of kilometers, the link propagation delay is expected to be a few hundred microseconds, or higher.

[6] In this work, we use the terms "link source" and "link destination" to refer to the source-end and the destination-end, respectively, of a unidirectional link.

- *Shared-link protection*: Fig. 6(c) illustrates the steps in the protection-switching procedure for dedicated-link protection. First, upon detecting a link failure, the link source of the failed link sends a *setup* message to the link destination along the backup route (which is determined in advance at the time of connection setup) and configures the OXCs at each intermediate node along the backup path (in shared protection, at the time of connection setup, wavelengths are reserved in advance for backup paths but OXCs are not configured to allow for sharing of backup wavelengths). The link destination, upon receiving the setup message, sends a *confirm* message back to the link source, thus completing the protection-switching procedure. The total time for shared-link protection switching is

$$F + (m+1)\times C + 2\times(m+1)\times D + 2\times m\times P.$$

The average protection-switching time for a single-link failure is the protection-switching time averaged over all the connections that traverse the failed link, i.e., the expected time to restore a connection traversing a failed link (or the expected "downtime" for a connection traversing a failed link). The network-wide average protection-switching time is the weighted average of the protection-switching time averaged over all single-link failures, and weighted by the number of connections traversing a failed link. The network-wide average protection-switching time is indicative of the expected data losses due to a link failure.

Table VI shows the average protection-switching times for all single-link failures for different protection schemes when there is a random demand of 30 connections in the network, and the OXC configuration time is 10 μs. The routing and wavelength assignment for the primary and backup lightpaths for different protection schemes are performed according to the ILPs 1–3. Tables VII–X show the network-wide average protection-switching times for each of the protection schemes for random demands, when the OXC configuration times are 10 ns, 10 μs, 500 μs, and 10 ms, respectively. We note that the network-wide average protection-switching time for dedicated-path protection-switching remains the same in Tables VII–X.

Summary: When the OXC configuration time is low (10 ns), the protection schemes in increasing order of average protection-switching times are as follows: a) shared link, b) dedicated path, and c) shared path. When the OXC configuration time is high (10 ms), the protection schemes in increasing order of average protection-switching times are as follows: a) dedicated path, b) shared link, and c) shared path.

The backup paths in shared-link protection tend to have fewer hops than the backup paths in path protection. Also, in shared-link protection, the end nodes of the failed link do not send messages to the source node and destination node of each connection that traverses the failed link. Therefore, when the OXC configuration time is low (10 ns) and the propagation delays of control messages (that establish the backup path for a connection) dominate the protection-switching time, shared-link protection has a better protection-switching time than the path-protection schemes. However, when the OXC configuration time is high (10 ms), the time required to configure OXCs on the backup path in shared-protection schemes dominate their protection-switching times (we note that in

TABLE VI
AVERAGE PROTECTION-SWITCHING TIMES IN MILLISECONDS FOR DIFFERENT PROTECTION SCHEMES, WHEN THERE ARE 30 CONNECTIONS IN THE NETWORK AND THE OXC CONFIGURATION TIME IS 10 μs. THE ENTRY "-" INDICATES THAT THE CORRESPONDING PROTECTION SCHEME DID NOT UTILIZE THE CORRESPONDING LINK TO ROUTE ANY CONNECTIONS

Failed link	*Dedicated-path*	*Shared-path*	*Shared-link*
(1,3)	3.73	6.46	2.53
(1,6)	2.70	5.07	2.25
(1,7)	3.32	3.94	2.53
(1,11)	3.11	7.09	3.36
(1,13)	4.14	5.95	2.53
(2,6)	2.91	2.43	1.7
(2,7)	1.68	3.44	1.7
(3,1)	2.5	5.96	2.53
(3,4)	2.70	-	2.53
(4,3)	3.32	3.44	2.53
(4,13)	2.99	3.44	2.53
(5,6)	2.29	6.47	2.53
(5,13)	2.91	4.45	-
(6,1)	4.55	6.46	1.7
(6,2)	4.14	6.71	1.7
(6,5)	3.32	5.95	2.53
(6,7)	2.5	2.93	1.7
(6,8)	5.37	6.21	1.7
(6,15)	3.83	6.20	4.19
(7,1)	2.91	4.61	3.36
(7,2)	-	2.93	1.7
(7,6)	4.96	5.70	-
(7,8)	3.73	-	1.7
(7,9)	3.56	4.78	3.36
(8,6)	4.14	5.46	1.7
(8,7)	2.09	2.93	1.7
(9,7)	3.73	6.21	3.36
(9,10)	3.73	3.93	3.36
(10,9)	3.62	5.46	3.36
(10,11)	3.52	3.44	3.36
(11,1)	3.11	4.95	3.36
(11,10)	3.01	6.96	3.36
(12,13)	3.73	4.95	4.19
(12,14)	-	4.69	4.19
(13,1)	2.91	5.46	2.53
(13,4)	3.73	4.95	2.53
(13,5)	3.52	-	2.53
(13,12)	5.37	4.44	4.19
(14,12)	3.11	5.45	4.19
(14,15)	4.14	4.95	4.19
(15,6)	4.82	6.3	4.19
(15,14)	4.41	6.70	4.19

TABLE VII
NETWORK-WIDE AVERAGE PROTECTION-SWITCHING TIMES IN MILLISECONDS FOR DIFFERENT PROTECTION SCHEMES WHEN THE OXC CONFIGURATION TIME IS 10 ns

Connections	*Dedicated-path*	*Shared-path*	*Shared-link*
10	3.33	4.86	2.85
15	2.65	3.92	2.84
20	2.92	4.61	2.81
25	3.32	5.04	2.75
30	3.49	5.02	2.79
35	3.33	4.78	2.81

TABLE VIII
NETWORK-WIDE AVERAGE PROTECTION-SWITCHING TIMES IN MILLISECONDS FOR DIFFERENT PROTECTION SCHEMES WHEN THE OXC CONFIGURATION TIME IS 10 μs

Connections	*Dedicated-path*	*Shared-path*	*Shared-link*
10	3.33	4.91	2.90
15	2.65	3.96	2.88
20	2.92	4.66	2.85
25	3.32	5.09	2.80
30	3.49	5.07	2.83
35	3.33	4.83	2.85

TABLE IX
NETWORK-WIDE AVERAGE PROTECTION-SWITCHING TIMES IN MILLISECONDS FOR DIFFERENT PROTECTION SCHEMES WHEN THE OXC CONFIGURATION TIME IS 500 μs

Connections	*Dedicated-path*	*Shared-path*	*Shared-link*
10	3.33	7.32	5.08
15	2.65	5.99	5.05
20	2.92	7.00	5.00
25	3.32	7.64	4.92
30	3.49	7.59	4.97
35	3.33	7.16	5.01

TABLE X
NETWORK-WIDE AVERAGE PROTECTION-SWITCHING TIMES IN MILLISECONDS FOR DIFFERENT PROTECTION SCHEMES WHEN THE OXC CONFIGURATION TIME IS 10 ms

Connections	*Dedicated-path*	*Shared-path*	*Shared-link*
10	3.33	53.88	47.34
15	2.65	45.30	47.14
20	2.92	52.29	46.72
25	3.32	57.01	46.04
30	3.49	56.43	46.45
35	3.33	52.31	46.75

dedicated-path protection, the OXCs on the backup paths are preconfigured at the time of connection setup). As a result, when the OXC configuration time is high, dedicated-path protection has a better protection-switching time than the shared protection schemes (shared-path and shared-link).

IV. DISTRIBUTED RESTORATION PROTOCOLS

In this section, we examine dynamic-restoration schemes to protect against link failures. In dynamic-restoration schemes, the backup path for a connection is not determined in advance at the time of connection setup (like in the protection schemes), but it is determined dynamically (from the available spare capacity) upon a failure. We study dynamic restoration algorithms, and examine their performance and restoration-time requirements.

A. Dynamic Restoration

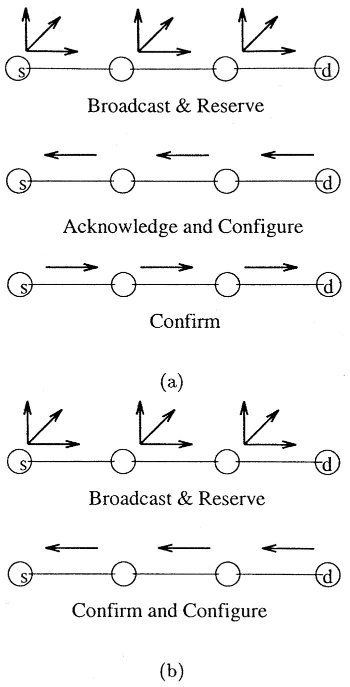

Fig. 7. Restoration process.

Distributed network restoration protocols have been researched in the literature [17], [21], [22], [24]. Distributed-restoration protocols discover backup paths dynamically upon the failure of a network component. In order to find a backup path for a connection, most distributed algorithms utilize the three-phase restoration process illustrated in Fig. 7(a): 1) the source node seeking a backup path sends out broadcast messages on all outgoing links with available capacity; 2) when a broadcast message reaches the destination node, the destination sends an acknowledgment message along the path traversed by the broadcast message, and simultaneously configures OXCs along the way; and 3) when the acknowledgment message reaches the source node, it sends a confirm message to the destination, thereby completing the connection setup on the backup path. Such control messages are exchanged on the control network, and the control network is assumed to be reliable.

We have adapted the three-phase protocol described previously for restoration of failed connections in an optical network. Since each connection on a failed link is on a different wavelength channel, the restoration process can be performed in parallel on different wavelengths, and connections that are broken due to a link failure do not contend for wavelengths for backup paths. Our restoration algorithm for discovering a lightpath on a given wavelength between a node pair is based on a two-phase process illustrated in Fig. 7(b).

1) The source node of the lightpath sends broadcast messages (with each broadcast message having a maximum hop limit[7]) on all outgoing links, and simultaneously reserves the wavelength on them. Intermediate nodes forward the broadcast message while reserving the wavelength on outgoing links.
2) When a broadcast message reaches the destination node, it sends a confirm message along the path to the source node. Upon receiving a confirm message, intermediate

[7] In our implementation, we assumed that the maximum hop limit is the diameter of the graph representing our network topology.

nodes configure their OXCs and forward the message. When the confirm message reaches the source node, a lightpath is established between the source and the destination nodes. Wavelengths that were reserved during the broadcast process but were not utilized for the backup path are released by cancel messages that are sent by a node upon a timeout[8] or upon exceeding the number of hops that a broadcast message can traverse.

B. Link Restoration

In link restoration, when a link (s, d) fails, note that each connection that traverses the link is on a different wavelength. Node s performs a restoration path search for each connection on the same wavelength that the connection utilizes. The restoration-path search for all connections can be performed in parallel on different wavelengths, utilizing the two-phase restoration process outlined previously. Restoration path searches for different connections do not contend for network resource, since they are on different wavelengths. When a restoration path for a connection is found, the connection is switched to the restoration path. Let m be the number of hops in the restoration path. The restoration time for the connection is

$$F + (m + 1) \times C + 2 \times (m + 1) \times D + 2 \times m \times P.$$

C. Path Restoration

The algorithm for path restoration upon a link failure is as follows.

1) The nodes adjacent to the failed link send *link-fail* messages to all the source and the destination nodes of all connections that traverse the failed link (illustrated in Fig. 4). As this message propagates to a source (and to the destination) node of a connection, the wavelength allocated for that connection may be released for use by other connections. Let n be the number of hops from the source-end node of the failed link to the source node of the connection.
2) When a source node of a connection receives a *link-fail* message, it initiates a restoration-path search on a certain set of wavelengths. All of the free wavelengths on the failed link may be partitioned[9] into sets of wavelengths, one set for each of the connections that traverses the failed link, and the source nodes of connections initiate a search on its designated set of wavelengths. This partitioning of wavelengths may be performed by the source-end node of the failed link and can be included in the link-fail message. This partitioning of wavelengths ensures that different connections that traverse the failed link do not contend for wavelength resources when they search for backup paths. The restoration-path search is performed on each wavelength in parallel, utilizing the two-phase restoration process illustrated in Fig. 7(b). If a restoration path is found, the connection is setup on the restoration path. If more than one restoration path is found for a connection, the first one found is utilized, and the others are released. Since sources of different connections are searching for restoration routes on different wavelengths, they do not contend for network resources. Let m be the number of hops in the restoration route from the source node to the destination node.

The restoration time for the connection is

$$F + n \times P + (n{+}1) \times D + (m{+}1) \times C + 2 \times m \times P + 2 \times (m{+}1) \times D.$$

The link-restoration efficiency is the ratio of the number of connections that are restored after the link failure to the total number of connections that traverse the failed link. The network-wide restoration efficiency is the weighted average of the link-restoration efficiency, weighted by the number of connections that traverse a failed link, averaged over all single-link failures. The average restoration time for a single-link failure is the restoration time averaged over all the connections that traverse the failed link. The network-wide average restoration time is the weighted average of the restoration time averaged over all single-link failures, and weighted by the number of connections traversing a failed link.

We have simulated the path- and link- restoration procedures on the 16-wavelength network in Fig. 1 to understand their behavior. We assume that connections (with uniformly distributed source-destination pairs) arrive as a Poisson process and are active for an exponentially distributed holding time with a mean of 1 (normalized) unit. We assume fixed-alternate routing with four alternate routes and first-fit wavelength assignment[10] [27]. In addition, we do not set aside spare wavelengths on any links in advance. We assume that the network parameters take on the values as specified in Section III-B. In the results shown here, the OXC configuration time is assumed to be 10 μs.

After 100 000 connection arrivals, we freeze the network state, simulate the failure of each link in the network, and record the number of restored connections for path and link restoration. Table XI illustrates the restoration performance of path and link restoration at a load of 60 Erlangs. The network-wide restoration efficiencies for path and link restoration were 65% and 49%, respectively. The network-wide average restoration time for path and link restoration were 3.55 and 2.78 ms, respectively. Table XII illustrates the restoration efficiency and restoration time for path and link restoration at different loads. *Generally, path restoration has a better restoration efficiency than link restoration, and link restoration has a better restoration time compared with path restoration*. Path restoration performs a search for a backup path on an end-to-end basis (the backup path could possibly be on a different wavelength), whereas link restoration is constrained to find backup paths around the failed link on the same wavelengths as that of the failed connections. As a result, path restoration performs better in finding wavelength-continuous backup paths. The backup paths in link restoration tend to have fewer hops than the backup paths found in path restoration. In addition, in link restoration, the end nodes of the failed link do not send messages to the source node and destination node of each connection that traverses the failed link. Therefore, link restoration has a

[8] In our implementation, we set the value of the timeout to be infinity.

[9] In our implementation, we assume that the free wavelengths are partitioned equally among all connections that traverse the failed link.

[10] In first-fit wavelength assignment, the lowest numbered wavelength among the set of free wavelengths on a route is chosen.

TABLE XI
RESTORATION PERFORMANCE OF PATH AND LINK RESTORATION FOR THE 16-WAVELENGTH NETWORK IN FIG. 1 WITH A LOAD OF 60 ERLANGS. RESTORATION TIMES ARE IN MILLISECONDS

Failed link	Connections lost	*Path restoration*		*Link restoration*	
		# restored	restoration time	# restored	restoration time
(1,6)	5	5	3.19	5	2.03
(1,7)	1	1	4.19	1	1.7
(1,11)	8	7	4.48	7	3.36
(1,13)	6	4	3.36	3	2.53
(2,6)	1	0	-	0	-
(2,7)	1	0	-	0	-
(3,1)	5	3	3.36	3	2.53
(4,3)	1	1	3.36	0	-
(4,13)	4	1	3.36	1	2.53
(5,6)	2	2	2.94	1	2.53
(5,13)	5	5	2.98	3	3.08
(6,1)	1	1	3.78	1	1.7
(6,2)	4	3	2.80	2	2.25
(6,5)	4	3	2.53	2	2.53
(6,7)	5	5	2.86	5	1.7
(6,8)	1	1	2.95	0	-
(6,15)	6	4	4.4	3	4.19
(7,1)	2	2	3.78	2	2.11
(7,2)	1	0	-	0	-
(7,6)	5	5	2.78	5	1.7
(7,8)	3	3	3.36	2	2.11
(7,9)	2	1	5.02	1	3.36
(8,6)	1	1	2.53	1	2.53
(9,7)	5	2	2.74	2	3.36
(10,9)	2	1	2.53	1	3.36
(10,11)	1	0	-	0	-
(11,1)	6	1	5.02	1	3.36
(11,10)	4	3	4.60	1	3.36
(12,13)	5	4	3.88	3	4.74
(13,1)	4	3	4.47	3	2.53
(13,4)	8	7	3.65	7	2.76
(13,5)	3	2	3.57	2	3.77
(13,12)	3	1	3.78	0	-
(14,12)	2	1	2.53	1	4.19
(14,15)	2	1	3.36	0	-
(15,6)	4	2	4.6	1	4.19
(15,14)	5	2	4.39	1	4.19

TABLE XII
RESTORATION PERFORMANCE OF PATH AND LINK RESTORATION FOR THE 16-WAVELENGTH NETWORK IN FIG. 1

Load	*Path restoration*		*Link restoration*	
(Erlangs)	Efficiency (%)	Time (ms)	Efficiency (%)	Time (ms)
10	96	3.89	89	2.99
20	87	3.66	79	3.12
30	85	3.66	69	2.97
40	73	3.63	59	3.19
50	66	3.85	47	2.94
60	65	3.55	49	2.78

better restoration time than path restoration. The restoration efficiency for path and link restoration decreases as the load increases, because there are fewer spare wavelengths available in the network.

V. CONCLUSION

Optical networks based on WDM technology can potentially transfer several gigabytes per second of data on each fiber link in the network. However, the high capacity of a link has the drawback that a link failure can potentially lead to the loss of a large amount of data (and revenue). Thus, all such failures must be dealt with quickly and efficiently.

This study examined different approaches to survive link failures in an optical network. These approaches are based on two basic survivability paradigms: 1) path protection/restoration, and 2) link protection/restoration. In path- and link-protection schemes, backup paths and wavelengths are reserved in advance at the time of connection setup. Path- and link-restoration schemes are dynamic schemes in which backup paths are discovered (from the spare capacity in the network) upon the occurrence of a failure. We formulated ILPs to determine the capacity utilization for the protection schemes discussed previously for a given traffic demand. The numerical results obtained for a representative network topology and for random demands (between 10 and 35 connections) indicate that shared-path protection provides significant savings in capacity utilization over dedicated-path and shared-link protection schemes, and dedicated-path protection provides marginal savings in capacity utilization over shared-link protection. On the other hand, if two fiber links fail in the network at the same time, then the number of connections that are dropped under shared-path or dedicated-path protection schemes is more than those that are dropped under shared-link protection. Thus, we observe that, in each protection scheme, there is a tradeoff between the capacity utilization and the susceptibility to multiple fiber failures.

We formulated a model of protection-switching times for the different protection schemes, based on a fully distributed control network. Based on our assumptions, we find that when the OXC configuration time is low (10 ns), the protection schemes in increasing order of average protection-switching times (for a random demand of 30 connections) are as follows: 1) shared-link—2.79 ms, 2) dedicated-path—3.49 ms, and 3) shared-path—5.02 ms. When the OXC configuration time is high (10 ms), the protection schemes in increasing order of average protection-switching times (for a random demand of 30 connections) are as follows: 1) dedicated-path—3.49 ms, 3) shared-link—46.45 ms, and 4) shared-path—56.43 ms.

We proposed distributed control protocols for path and link restoration. Numerical results obtained from simulation experiments on these protocols indicate that path restoration has a better restoration efficiency than link restoration, and link restoration has a faster restoration time than path restoration.

References

[1] B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, July 1997.

[2] P. Bonenfant, "Optical layer survivability: A comprehensive approach," in *Proc. OFC '98*, vol. 2, San Jose, CA, Feb. 1998, pp. 270–271.

[3] O. Gerstel and R. Ramaswami, "Optical layer survivability: A services perspective," *IEEE Commun. Mag.*, vol. 38, pp. 104–113, Mar. 2000.

[4] T. Wu, *Fiber Network Service Survivability*. Norwood, MA: Artech House, 1992.

[5] ——, "Emerging technologies for fiber network survivability," *IEEE Commun. Mag.*, vol. 33, pp. 58–74, Feb. 1995.

[6] O. Gerstel and R. Ramaswami, "Optical layer survivability—An implementation perspective," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1885–1899, Oct. 2000.

[7] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, pp. 16–23, Nov.–Dec. 2000.

[8] *Tutorial Sessions, Optical Fiber Communications Conference (OFC '99)*, San Diego, CA, Feb. 1999.

[9] O. Crochat, J.-Y. Le Boudec, and O. Gerstel, "Protection interoperability for WDM optical networks," *IEEE/ACM Trans. Networking*, vol. 8, pp. 384–395, June 2000.

[10] O. Gerstel, R. Ramaswami, and G. Sasaki, "Fault tolerant multiwavelength optical rings with limited wavelength conversion," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1166–1178, Sept. 1998.

[11] M. Medard, S. G. Finn, R. A. Barry, and R. G. Gallager, "Redundant trees for preplanned recovery in arbitrary vertex-redundant or edge-redundant graphs," *IEEE/ACM Transa. Networking*, vol. 7, pp. 641–652, Oct. 1999.

[12] E. Limal, S. L. Danielsen, and K. E. Stubkjaer, "Capacity utilization in resilient wavelength-routed optical networks using link restoration," in *Proc., OFC '98*, vol. 2, San Jose, CA, Feb. 1998, pp. 297–298.

[13] E. Karasan and E. Goldstein, "Optical restoration at the wavelength-multiplex section level in WDM mesh networks," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 1343–1345, Sept. 1998.

[14] A. Fumagalli and L. Valcarenghi, "IP restoration vs. WDM protection: Is there an optimal choice?," *IEEE Network*, vol. 14, pp. 34–41, Nov.–Dec. 2000.

[15] L. Sahasrabuddhe, S. Ramamurthy, and B. Mukherjee, "Fault management in IP-over-WDM networks: WDM protection vs. IP restoration," *IEEE J. Select. Areas Commun.*, to be published.

[16] G. Mohan, C. S. R. Murthy, and A. K. Somani, "Efficient algorithms for routing dependable connections in WDM optical networks," *IEEE/ACM Trans. Networking*, vol. 9, pp. 553–566, Oct. 2001.

[17] H. Zang and B. Mukherjee, "Connection management for survivable wavelength-routed WDM mesh networks," *Optical Networks Mag.*, vol. 2, no. 4, pp. 17–28, July–Aug. 2001.

[18] R. R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal capacity placement for path restoration in mesh survivable networks," in *Proc. ICC '96*, Dallas, TX, June 1996, pp. 1568–1574.

[19] M. Herzberg, S. J. Bye, and A. Utano, "The hop-limit approach for spare-capacity assignment in survivable networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 775–784, Dec. 1995.

[20] J. Anderson, B. T. Doshi, S. Dravida, and P. Harshavardhana, "Fast restoration of ATM networks," *IEEE J. Select. Areas Commun.*, vol. 12, pp. 128–138, Jan. 1994.

[21] W. D. Grover, "The self-healing network: A fast distributed restoration technique for networks using digital crossconnect machines," in *Proc. IEEE GLOBECOM '87*, Tokyo, Japan, Nov. 1987, pp. 28.2.1–28.2.6.

[22] C. E. Chow, J. Bicknell, S. McCaughey, and S. Syed, "A fast distributed network restoration algorithm," in *Proc. IEEE IPCCC '93*, Tempe, AZ, Mar. 1993, pp. 261–267.

[23] H. Sakauchi, Y. Nishimura, and S. Hasegawa, "A self-healing network with an economical spare-channel assignment," in *Proc. IEEE GLOBECOM '90*, San Diego, CA, Dec. 1990, pp. 438–443.

[24] H. Komine, T. Chujo, T. Ogura, and T. Soejima, "A distributed restoration algorithm for multiple-link and node failures of transport networks," in *Proc. IEEE GLOBECOM '90*, San Diego, CA, Dec. 1990, pp. 459–463.

[25] S. Ramamurthy and B. Mukherjee, "Fixed alternate routing and wavelength conversion in wavelength routed optical networks," in *Proc. IEEE GLOBECOM '98*, vol. 4, Sydney, Australia, Nov. 1998, pp. 2295–2302.

[26] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 489–500, Oct. 1995.

[27] H. Zang, J. P. Jue, L. Sahasrabuddhe, R. Ramamurthy, and B. Mukherjee, "Dynamic lightpath establishment in wavelength routed WDM networks," *IEEE Commun. Mag.*, vol. 39, pp. 100–108, Sept. 2001.

S. Ramamurthy received the B.Tech degree in computer science from Indian Institute of Technology, Chennai, and the M.S., and Ph.D. degrees in computer science from University of California, Davis.

Previously, he worked at Tellium and at Telcordia Technologies. He is currently with CIENA Corporation, Linthicum, MD, working on control plane protocols for intelligent optical networks.

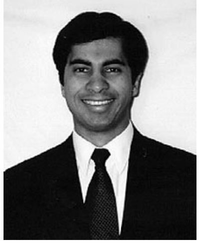

Laxman Sahasrabuddhe (S'94–M'00) received the B.Tech. degree from the Indian Institute of Technology, Kanpur, in 1992, the M.Tech. degree from the Indian Institute of Technology, Madras, in 1994, and the Ph.D. degree from the University of California, Davis, in 1999.

From 1999 to 2000, he was an Embedded Software Engineer at Amber Networks, which was acquired by Nokia in July 2001. Currently, he is a Principal Member of Technical Staff at SBC Communications, Inc., San Ramon, CA.

Dr. Sahasrabuddhe is the recipient of the "Best Doctoral Dissertation Award," from the College of Engineering, University of California, for his research on wavelength-division-multiplexing (WDM) optical networks.

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

He held a GTE Teaching Fellowship and a General Electric Foundation Fellowship with the University of Washington. In July 1987, he joined the University of California, Davis, where he has been Professor of Computer Science since July 1995 and served as Chairman of Computer Science from September 1997 to June 2000. He serves or has served on the editorial boards of the *ACM/Baltzer Wireless Information Networks* (WINET), *Journal of High-Speed Networks*, *Photonic Network Communications*, and *Optical Network Magazine*. He is author of the textbook *Optical Communication Networks* (New York: McGraw-Hill, 1997), a book that received the Association of American Publishers, Inc.'s 1997 Honorable Mention in Computer Science. He is a Member of the Board of Directors of IPLocks, Inc., a Silicon Valley startup company. He has consulted for and served on the Technical Advisory Board of a number of startup companies in optical networking. His research interests include lightwave networks, network security, and wireless networks.

Dr. Mukherjee is co-winner of paper awards presented at the 1991 and the 1994 National Computer Security conferences. He serves or has served on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING and IEEE *Network*. He also served as Editor-at-Large for optical networking and communications for the IEEE Communications Society and as the Technical Program Chair of the IEEE INFOCOM '96 conference.