

Multiple Failure Survivability in WDM Mesh Networks

Sun-il Kim

University of Illinois at Urbana-Champaign
 CS Department
 Coordinated Science Laboratory
sunilkim@crhc.uiuc.edu

Steven S. Lumetta

University of Illinois at Urbana-Champaign
 ECE Department
 Coordinated Science Laboratory
lumetta@uiuc.edu

Abstract—Survivability is important in designing reliable optical networks and has been widely studied in the literature. Most practical studies of high-speed recovery algorithms assume a single-link failure model. As networks grow in size and complexity, both the likelihood and the impact of failures increase. In order to guarantee that a failure in one part of a network does not affect the operation of the entire global network, we must understand the impact of multiple failures and provide efficient solutions to address them. In this paper, we study the impact of non-simultaneous, two-link failures on different protection schemes and provide an understanding of the different failure scenarios that arise when a multiple failure model is considered. We then discuss tradeoffs between different solutions that may be employed to address multiple failures, and present a simple dynamic protection reconfiguration technique, which can be used to address multiple failures.

We quantify the tradeoffs between different protection schemes in terms of capacity cost and survivability from two sequential link failures. In dynamic protection reconfiguration, protection routes are quickly recomputed and resources are reassigned in the event of a failure, optimizing robustness from additional failures. With the use of a simple reconfiguration technique, a network's survivability is limited only by topological constraints, where only complete network disconnections result in broken connections. For five sample networks, the dynamic protection reconfiguration technique requires little additional capacity while significantly improving the survivability.

Keywords: optical fiber communication, WDM network, optical network, network survivability, network reliability, protection, restoration, multiple failures,

I. INTRODUCTION

Optical transports have allowed us to meet the increasing demands for bandwidth driven by the exponential growth of data traffic through the use of wavelength division multiplexing (WDM) technology. With the extremely high volume of traffic being carried on WDM networks, failures such as fiber cuts can result in a loss of enormous amount of data (on the order of several terabits per second) and revenue [1]. It is imperative to maintain a high level of reliability in order to support the growing number of critical applications that utilize these networks. Survivability—the ability to recover from network failures—is therefore an important aspect of optical networking.

The material presented in this paper is based in part upon work supported by National Science Foundation grants ANI 01-21662 ITR and ACI 99-84492 CAREER. The content of the information does not necessarily reflect the position or the policy of that organization.

There are many survivability techniques that offer tradeoffs between recovery speed, protection capacity, and management overhead and complexity (a brief survey of existing algorithms is provided in Section II). Because varying network environments have different reliability needs, a good understanding of these tradeoffs is necessary in order to efficiently manage and operate a network. In addition, as both the size and complexity of networks continue to increase, and as large networks become interconnected, the ability to gracefully degrade in the event of a failure also becomes important. To this end, multiple failure survivability is critical. A *network must be able to handle multiple failures so that a failure in one part of a network does not affect the entire network's ability to recover from subsequent failures*. A few studies dealing with double-link failures have been presented in the literature, but this problem still receives relatively little attention. In this article, we first present an overview of the failure models and the existing failure management schemes. We introduce a classification scheme for multiple failure management techniques and provide an overview of the topic.

A few intuitive techniques to solve the problem of multiple failure survivability exist and offer tradeoffs in terms of recovery time (which in turn translates to data loss), cost overhead, and management complexity. For example, dynamic recovery techniques can best handle multiple failures, but usually take much longer to recover broken connections compared to preplanned recovery techniques. On the other hand, *multiple failure protection planning* may be used to protect a network from multiple failures using preplanned recovery. This method allows fast recovery, but may incur inefficiently high cost in terms of network resources. Also, some combination of the two approaches may be implemented, but the increase in protocol, management and hardware complexity renders the solution less attractive. Further discussion of this issue is presented in Section 4.

We present *dynamic protection reconfiguration*, which is used to perform fast reconfiguration of recovery routes and reallocate protection wavelengths after a failure (online) to allow maximum recoverability from multiple failures. Note that we focus on recovery techniques that guarantee less than 100ms recovery time and, therefore, consider protection rather than (dynamic) restoration [2]. Dynamic protection reconfiguration, or DPR in short, is a more capacity-efficient alternative to multiple failure protection planning. A network dynamically adapts to failures under DPR, and can achieve the maximum

robustness bounded by the topology.

We assume two-link connected mesh networks in this paper, but the technique is not limited by the connectivity of a network. Most WDM wide area networks found in practice and in the literature are two-link connected. We also introduce a multiple failure classification scheme for WDM mesh protection algorithms, which aids in better understanding multiple failure survivability. We also introduce two metrics that capture a network's ability to handle multiple failures, which is used to evaluate different protection schemes and the effectiveness of the dynamic protection reconfiguration technique. Dynamic protection reconfiguration is evaluated using four common protection schemes—dedicated path protection, shared path protection, dedicated link protection, and shared path protection.

We find that two classes of failures that affect all protection algorithms—recovery paths used for the first failure are hit by the second failure, or connections affected by the second failure cannot be recovered because recovery paths are broken by the first failure—contribute to a large number of unsuccessful multiple failure recovery. Additionally, capacity-efficient algorithms that allow sharing of protection resources can dramatically improve operation costs, but significantly amplify the impact of failures. Our results show that shared protection is about 46% more vulnerable to unrecoverable two-link failures compared to dedicated protection. Protection reconfiguration can significantly raise the *recovery ratio*, the average percentage of successfully recovered connections from two-link failures, allowing over 98.9% of the affected connections to be protected. Protection reconfiguration can be used in conjunction with most protection schemes to achieve the maximum robustness for a given topology leaving the network vulnerable only to complete partitioning. Finally, we propose the use of DPR with *maximally disjoint paths* (MDP). When the number of failures is greater than or equal to the degree of connectivity, it may not be possible to reallocate a backup route that is completely disjoint with the live route. Using MDP's instead of leaving the connections vulnerable in such cases allow the network achieve maximum robustness under multiple failures.

The remainder of the paper is organized as follows. The next section provides a brief background on optical layer survivability. In Section 3, we introduce normalized vulnerability and recovery ratio—metrics used to evaluate a network's susceptibility to two-link failures. We then discuss our failure classification scheme for two-link failures. A detailed description of dynamic protection reconfiguration is provided in Section 4. In Section 5, we apply our metrics under the classification scheme to sample networks from the literature, and evaluate the impact of dynamic protection reconfiguration on four different protection schemes (dedicated/shared link/path protection). Finally, Section 6 presents our conclusions.

II. BACKGROUND

A. Failure Models

The high data rates in modern and future networks—on the order of several Gb/s in an optical network—exacerbate

the severity of failures. Failure modes in optical networks consist of channel failures, link failures, and failures of optical crossconnects (OXCs). Channel failures are the most common, and are often caused by the failure of a card or cards at a port of an OXC. Link failures—fiber cuts caused by wayward backhoes, amplifier failures, etc.—are also common, and lead to failure of all channels on all fibers in the link. Node (OXC) failures are less common, but can cause failure of all channels that originate, pass through, or terminate at the node. We focus on link failures in this paper.

B. Survivability Schemes

Protection and Restoration are the two main approaches that address link failures in optical networks [2]. Restoration addresses failures by locating free wavelength channels for backup after a failure occurs. Protection preplans backup routes that are used in the event of a failure. Protection and restoration offer a tradeoff between the speed of recovery and efficiency in terms of the use of spare capacity [3], [4]. Restoration schemes are more efficient in terms of capacity requirements, and offer better multiple failure survivability because it dynamically finds backup paths after a failure. However, protection can be implemented in a capacity efficient manner [5], [6], [7], [8], [9] and offer much faster recovery with the absence of the excess signaling delay needed for dynamic route discovery [10], [11], [12]. Restoration based recovery takes about 2 seconds, whereas protection schemes can achieve complete recovery in the order of tens of milliseconds [1]. Given the fact that many transport infrastructures require rapid recovery, we focus on protection algorithms.

Protection schemes can be generally classified into four types: dedicated path protection (DPP), shared path protection (SPP), dedicated link protection (SLP), and shared link protection (SLP). These techniques assume two-link connectivity, and guarantee recovery under single failures. Path protection requires the knowledge of the whole path and selection of a backup path that is link-disjoint (link failure survivable) or node-disjoint (node failure survivable) from the primary path. In DPP, a dedicated backup path is setup for each connection, and to achieve recovery in the event of a failure, this alternate path is used. DPP offers fast recovery because no signaling between the source and the destination nodes is required, but is inefficient in terms of capacity requirements. In SPP, the end nodes of a lightpath signal the intermediate nodes to establish the backup route. Capacity reserved for backup can be shared among different connections that do not share any common failure modes, or can be used to carry low priority (unprotected) traffic, which is preempted in the event of a failure. The signaling and configuration of the intermediate OXCs render SPP slow compared to DPP [13].

In link protection, nodes that are adjacent to the failure initiate recovery using reserved protection capacity. SLP allows sharing of protection capacity among different connections, whereas resources are dedicated in DLP. SLP is thus more cost-efficient compared to DLP. However, SLP is slower compared to DLP due to signaling and configuration of intermediate nodes required after a failure. With link protection, failure

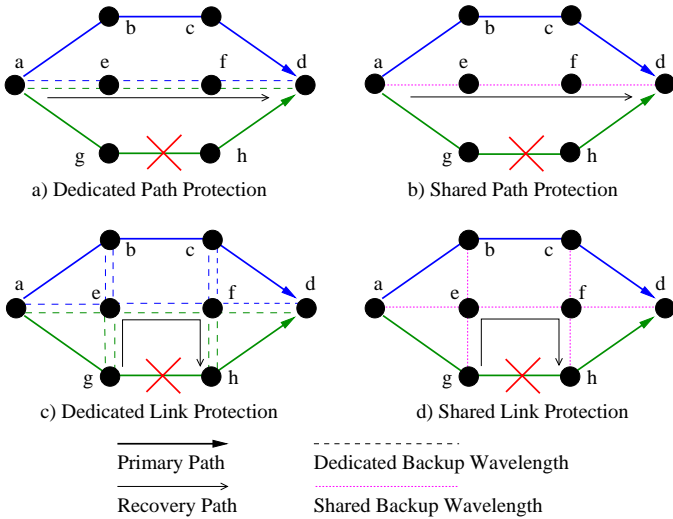


Fig. 1. Illustration of different protection approaches.

recovery usually involves the use of more local resources compared to path protection (as protection routes are found for each link rather than for the entire lightpath). Recovery is usually faster than path protection because recovery is initiated more quickly and fewer OXCs need to be signaled and configured (shorter, more local recovery routes), but it is more inefficient in terms of protection capacity [6], [8], [14].

Figure 1 illustrates the four different protection schemes. All four examples show two routed connections from node a to node d . We denote the paths p_1 ($[a-b-c-d]$) and p_2 ($[a-g-h-d]$). In DPP, p_1 and p_2 share the same recovery route $[a-e-f-d]$, but separate wavelengths are reserved. In SPP, p_1 and p_2 share both the recovery route and a backup wavelength channel. DLP and SLP protects each individual link along the primary path using separate recovery routes. Link $[a,b]$ is protected by path $[a-e-b]$ (Similarly- $[b,c]$ - $[b-e-f-c]$, $[c,d]$ - $[c-f-d]$, $[a,g]$ - $[a-e-g]$, $[g,h]$ - $[g-e-f-h]$, $[h,d]$ - $[h-f-d]$). DLP allocates separate wavelength channels for the recovery routes, whereas SLP allows sharing of wavelengths.

These protection schemes capture the characteristics of most protection algorithms found in the literature. For example, DLP based on selecting shortest recovery paths possible for each link provides the fastest recovery among all protection algorithms, including mesh protection [2], [5], [9], streams protection [15], flooding-based mesh restoration [16], generalized loopback [17], p-cycles [18], ring-based schemes such as cycle double covers [19], and other protection methods that attempt to offer advantages of both PP and LP [20], [21]. On the other hand, SPP is considered the most capacity-efficient protection algorithm [13].

III. MULTIPLE FAILURE SURVIVABILITY

All of the survivability techniques discussed in the previous section are designed to handle single link or single node failures. However, the ability to better handle multiple failures is an important aspect of operating high performance networks and has recently started to become a topic of interest [22], [23], [24], [25], [26], [27], [28], [29].

Although this work focuses on protection, it is important to understand that restoration techniques are inherently capable of handling multiple failures as it dynamically establishes new routes for connections that are affected by failures. In terms of survivability, restoration techniques are optimal in that they are limited only by the network topology and the available capacity. The downside of restoration, as mentioned in the previous section, is the recovery speed. Because all affected connections must be rerouted dynamically, the online computation time and recovery protocol renders restoration less attractive for networks that require high availability. In contrast, network management may decide to precompute multiple recovery routes and reserve enough wavelength channels to allow multiple failure recovery. Such techniques, which we refer to as *multiple failure protection planning* can be expensive in terms of protection capacity, and therefore requires careful planning. In [25], an optimization technique for offline double-link failure protection planning for dedicated and shared link protection was presented. An optimization technique to reduce capacity is presented, but in general multiple failure protection planning techniques are inefficient in terms of network resources usage.

Also, some combination of the two approaches may be implemented, but the increase in protocol, management and hardware complexity in implementing two different recovery protocols render the solution less attractive.

We next discuss multiple failure models and the different approaches to managing multiple failures.

A. Multiple Failure Models

Multiple failures can be separated into two categories, differentiated by the temporal relationship between failures. *Simultaneous failures* refer to cases where multiple components fail at the same time (from a recovery algorithm's point of view). In other words, two failures happen close enough in time to disrupt recovery. One example of simultaneous failures is shared risk link group (SRLG) failures, such as a cut through a conduit shared by several topologically diverse links. SRLG requires special attention, as most recovery techniques cannot guarantee recovery from such failures, and poses complicated management issues that are out of the scope of this article [30]. *Sequential failures* are failures of multiple components separated by enough time for the recovery algorithm to complete recovery of each failure, but before they can be physically repaired. Recovery times are on the order of milliseconds in the optical layer, whereas physical repair of failures may take several hours or even days, exposing networks to sequential failures. Given the relative time scales for recovery and physical repair, sequential failures are much more likely to occur compared to simultaneous failures (with the exception of SRLG failures). We focus on addressing sequential multiple failures, and for the rest of the paper we use the terms multiple failures to mean sequential failures.

In this paper, to study a network's ability to support graceful degradation multiple failures, we consider two-link failures. A *two-link failure* consists of two independent link failures




Fig. 2. Illustration of different types of failures.

in a network graph. The second failure occurs long enough after the first to allow normal recovery to complete (on the order of milliseconds) but before any physical repair can be accomplished (possibly up to several hours or days). The two-link failure model effectively captures the characteristics of sequential failures, and aids in understanding the critical issues in multiple failure survivability ¹.

B. Failure Classification

Recovery algorithms can fail for many reasons, ranging from physical disconnection of a network to resource unavailability due to sharing of protection channels. In [22], a hierarchical failure classification was presented for fiber-switched link protection schemes. In this section, we present an intuitive failure classification scheme for two-link failures to capture the characteristics of connection-based protection schemes. The classes are organized into two types—fundamental and algorithmic—and we use these types to structure the text. For each class, we provide an intuitive rationale for the problems that lead to such failures, and explain the relative importance of the failure classes for typical networks.

1) *Fundamental failures*: Fundamental failures consist of network *disconnection failures* and *capacity failures* that occur as a result of network structure rather than as a result of the properties of recovery algorithms. No protection algorithm can possibly recover from disconnection failures. In a two-link redundant mesh network, failure of two links can partition the network, disconnecting some nodes from others and rendering recovery impossible for any protection scheme. The primary source of disconnection failures is nodes of degree two, as any such node becomes disconnected when both links to the node fail. Examples of each type appear on the upper left in Figure 2.

In the National network six nodes of degree two as well as six two-link cuts that partition more than one node from the rest of the network lead to disconnection failures. As shown in Figure 3, examples of the former include nodes 11 and 23, while examples of latter occur in the long chain of nodes in




Fig. 3. The National network.

the lower left of the network. In total, the network has 24 pairs of links that cause disconnection failures.

Capacity failures also result from the fact that many mesh networks are two-link connected. In protection reconfiguration, which is discussed in detail in Section IV, a secondary backup path is computed and reserved after a network is recovered from a failure. Pre-planning for two-link failure, as in [25], also requires a secondary backup path to be computed. However, in two-link connected networks a third, disjoint path (secondary backup path) does not exist for many node pairs, leaving some connections vulnerable to second link failures. Like disconnection failures, nodes of degree two pose a problem and capacity failures may also be disconnection failures.

2) *Algorithmic failures*: There are three algorithmic failures—*path hit*, *broken path*, and *blocked shared path*—that correspond to common aspects of protection algorithm design. All protection algorithms that are found in the literature are affected by the first two types of failures, and blocked shared path affects all algorithms that allow sharing of backup resources among different connections. It is important to note that some algorithmic failures may also be fundamental failures. In this section, we highlight the circumstances in which each class applies.

The first class of algorithmic failures arises from two-link failure scenarios where the second failure breaks the recovery path for the first. We term this class of failures *path hit failures*. Another class of algorithmic failures that all protection algorithms encounter results from the first failure breaking the recovery path of the second failure. These failures, which we term *broken path failures*, arise from the possibility that a failed link is a part of assigned backup paths for some lightpaths. It is purely an effect of pre-planning protection channels, which is inherent in all rapid recovery schemes found in the form of protection in the literature. Broken path failures and path hit failures occur as a result of the same phenomenon differentiated only by the time ordering of the two failures. Therefore, vulnerabilities resulting from broken path failures and path hit failures are expected to be exactly the same. If reconfiguration (or static reservation of a secondary backup path) is used, a connection may experience slight data loss due to jitter caused by switching of a failed live path to a reserved one. With broken path failures, the live data stream is not affected.

¹The significance of the two-link failure model, as intended in the original contribution in [22], is that it can be used to *understand* the impact of failures that can not be captured using single failure models. Our intent in focusing the two-link failure model for this paper is to understand the effect of multiple failures on a network rather than to study a double link failures scenarios in specific. Thus, it is imperative to note that the ideas and findings in this paper are not limited to two-link failures, but can be applied to n-link sequential failures in general.

Implementing protection algorithms in a capacity-efficient manner leads to the third class of algorithmic failures. This type of failure, which we term *blocked shared path failures*, occurs in cases where shared backup resources are used to recover the first link failure, which leaves some lightpaths without protection. This class of failure affects all algorithms that address capacity-efficiency through sharing of protection wavelengths among different connections. Therefore, there is an inherent tradeoff between the degree to which the protection is optimized for capacity and the network's ability to minimize failure impact.

Other than multiple failures

C. Multiple Failure Management

We classify multiple failure management techniques based on protection schemes into three categories—offline multiple failure protection planning, online multiple failure protection planning, and dynamic protection reconfiguration (online). offline multiple failure protection planning techniques [25], [22], [24], [23], [31] pre-allocate backup wavelengths using pre-computed recovery routes based on a static traffic demand. Online multiple failure protection planning performs optimization (recovery path computation and resource allocation) *online* at the time of provisioning. A network reoptimized each time a new connection request arrives or leaves the network, and therefore supports dynamic traffic demands. However, it can introduce a considerable amount of management complexity as a difficult optimization problem may need to be solved very frequently. To the best of our knowledge, online multiple failure protection planning has not been studied. Dynamic protection reconfiguration adapts dynamically to failures by computing and allocation resources for new recovery routes only for the connections that were affected by a failure, after the failure has been successfully recovered by a protection scheme.

D. 1+N Pre-allocation

Pre-allocation involves setting up 1+N diverse paths (where N is the number of failures that the network must handle) and assigning wavelengths for all connection when the network is initially provisioned. This scheme best supports static traffic, and thus allows for offline capacity optimization of routing and wavelength assignment (RWA). Integer Linear Programming (ILP) approaches are most commonly used to optimize for capacity, but often are slow for large networks. Full optimization may not be possible for many practical networks as the complexity of optimization problems grows exponentially with the size of the network and the demand as well a change in N. The optimization process, however, does not pose a serious problem in terms of management overhead, as it is performed only once for the network. Heuristics, such as genetic algorithms and simulated annealing, can be used, but designing efficient heuristics may also be difficult. Protection capacity requirement is high for pre-allocation. Existing techniques have focused on optimizing capacity for double link failures and require over 200% capacity for recovery.

E. Static Reconfiguration

In static reconfiguration, two diverse paths are computed and allocated for each connection, and some buffer wavelengths based on a computation of all possible second failure scenarios are reserved. Optimization for RWA is done offline using ILP or other heuristics similar to pre-allocation. In the event of a failure, after the affected connections are recovered, the network is reconfigured according to the precomputed second failure scenario using the buffer wavelengths. Management can then choose to reiterate the buffer computation and reservation process in order to support recovery from additional failures. During this reoptimization process, either ILP or some heuristic can be used. Capacity cost is lower compared to pre-allocation, but still may be expensive depending on the efficiency of the optimization technique used. To the best of our knowledge, no static reconfiguration allocation algorithm has been evaluated, but may be useful in designing future networks.

F. Dynamic Reconfiguration

Since most protection schemes are designed to handle one failure at a time, they can be naturally extended to handle sequential failures using reconfiguration. In the event of a failure, reconfiguration identifies and protects connections that are affected by the failure and the connection that are left vulnerable to additional failures. Reconfiguration information is then dynamically computed. Because the network dynamically adapts to a specific failure, dynamic reconfiguration can handle an arbitrary number of sequential failures (as long as the topology permits), and requires little additional capacity. Dynamic adaptation requires online allocation of new wavelengths, however, which may not be possible if not enough capacity is available in the network due to a high load.

The speed of the reconfiguration process is important, as the network is left vulnerable to additional failures during the computation. Therefore, fast heuristics are more attractive compared to full optimization techniques such as ILP due to its run-time. A study on dynamic reconfiguration showed that less than 10% additional capacity is required to support two link failures [26] with a fast heuristic. Dynamic provisioning of lightpaths and the use of fast heuristics for RWA result in sub-optimal capacity utilization. Reoptimization similar to the technique used in static reconfiguration can be used to periodically redesign protection for existing traffic. There is a tradeoff in terms of management overhead and capacity cost based on the frequency of this reoptimization process.

IV. DYNAMIC PROTECTION RECONFIGURATION

In this section, we describe our simple dynamic protection reconfiguration algorithm that can be used in conjunction with most existing protection schemes to provide optimal multiple failure protection. The goal is to maximize the network's ability to handle multiple failures (and therefore minimize service disruptions) while taking into account capacity costs and management complexity.

A more capacity-efficient alternative to multiple failure protection planning (1+N Pre-allocation) is to dynamically adapt

	1+N Pre-allocation	Static Reconfiguration	Dynamic Reconfiguration
Protection Capability	Limited by topology & choice of N	Limited by topology	Limited by topology & available capacity
Traffic Demand	Static	Static	Dynamic
Routing and Wavelength Assignment	Offline	Offline w/ online reoptimization	Online
Algorithm Complexity	High (ILP)	Moderate (heuristic) to High (ILP)	Low (fast heuristic)
Capacity Cost	High (1+N paths)	Moderate (2 paths+ Δ)	Low (~ 2 paths)
Management Overhead	Low	Low to moderate	Moderate

Fig. 4. Multiple Failure Management Schemes.

to failures. A small portion of our results were first published in [26] where we introduced the idea of reconfiguration and hinted at the different tradeoffs in using such technique.

Our work is different from a few related works published since the introduction of the multiple failure classification in [22]. First, [25], [23], [24], [28] solves two-link survivability in the context of link protection, but the dynamic protection reconfiguration technique is orthogonal to the choice of protection schemes. In this paper, we apply the technique to both link protection and path protection. Furthermore, our goal is not to provide a detailed optimization technique for a particular protection scheme, but rather study the impact of multiple failures and understand the tradeoffs involved in different classes of algorithms. Second, most assume that the network is three-edge connected. Three-edge connected topologies allow complete recovery under any two link failures because three edge-disjoint paths can be found for every node pair. Dynamic protection reconfiguration does not limit the number of failures up to which the network can survive. Survivability depends on the topology and the choice of the routing algorithm. For instance, a network can fully handle k sequential failures, if the network is at least $k + 1$ -connected. In most cases, employing the dynamic protection reconfiguration technique allows a network to reach the upper bound in survivability (limited by the topology). Last, others consider a static traffic demand and perform offline optimization. We consider an online routing model where offline optimization is not applicable. Many emerging WDM mesh networks will require fast setup and teardown of lightpaths under dynamic traffic demands which may render offline techniques less attractive [14], [32], [33].

Based on the many common networks found in practice and in the literature, dynamic protection reconfiguration assumes two-edge connectivity, and therefore assumes that the network can be fully protected against any single failure.

Preplanned double link failure survivability (1+N Pre-allocation, N=2) requires a primary and two disjoint backup paths. We term these paths *disjoint path triples*. Disjoint triples may not exist for some node pairs since we assume two-connected networks. Dynamic protection reconfiguration can be used with existing protection schemes with a small additional routing constraint to work around this problem. At the time of provisioning, disjoint primary/backup path pairs must be carefully selected to allow a possible third path,

which is disjoint from the original backup path, to be routed. This constraint does not significantly affect capacity utilization in protection schemes (evaluation details appear in Section V). Therefore, different path selection techniques and routing strategies can be used in conjunction with dynamic protection reconfiguration with minimal impact on the efficacy of the routing algorithm.

The failure classification scheme presented in the previous chapter provides useful insights towards better addressing multiple failures. One intuitive and effective way to achieve higher multiple failure survivability is to first target algorithmic failures. Algorithmic failures can be prevented by reconfiguring the network (recomputing recovery paths and re-allocating protection resources for affected connections). Blocked shared path failures are the easiest to avoid as it simply requires allocation additional wavelengths based on the same recovery routes. Path hit and broken path failures are slightly more difficult because they require computation of new recovery routes.

The goal is to reconfigure the affected connections as quickly as possible while minimizing capacity usage. Optimally, reconfiguration can complete close to recovery time. Given the multiple failure model described in Section II, it is important to have the network reconfigured in time to handle additional failures. Figure 5 shows an outline of the steps involved in dynamic protection reconfiguration. Re-computed paths are found by using the same path selection technique used for lightpath provisioning in the network. Reconfiguration starts immediately after a failure is detected. We assume that the centralized manager keeps track of all the connections and that the affected connections are easily identified once the failure is recovered and localized. Keeping a data structure that tracks all the primary paths and the recovery paths provisioned on each link is useful in quickly identifying connections that are affected by the failure. Once the failure is detected, table lookup can be performed using the data structure.

A. Reconfiguring for Optimal Vulnerability

In two-connected networks (n-connected), there may not exist a disjoint path pair after failure and recovery of a link (or n-1 links). Therefore, the path pairs must share some links. The links that must be utilized and shared in order to establish a connection between two nodes are termed critical links and

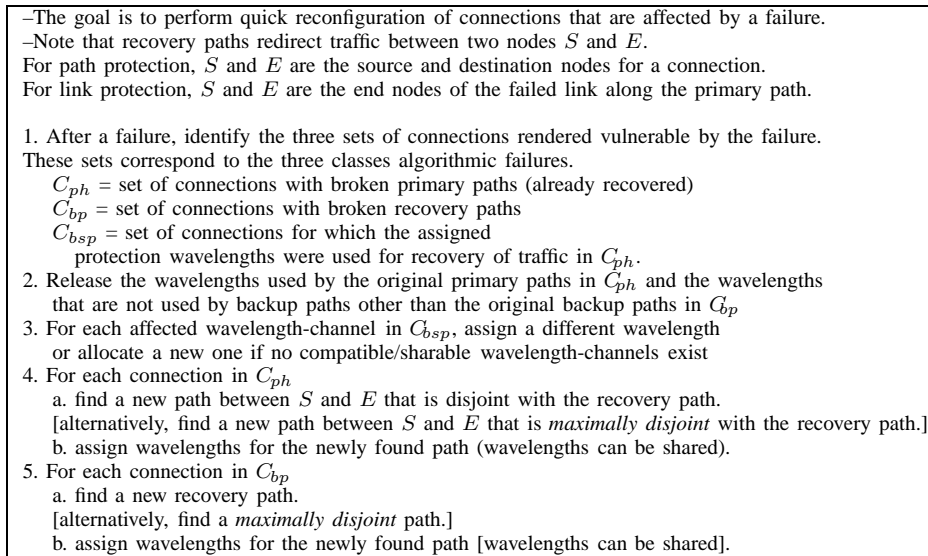


Fig. 5. Outline of the Dynamic Protection Reconfiguration Technique.

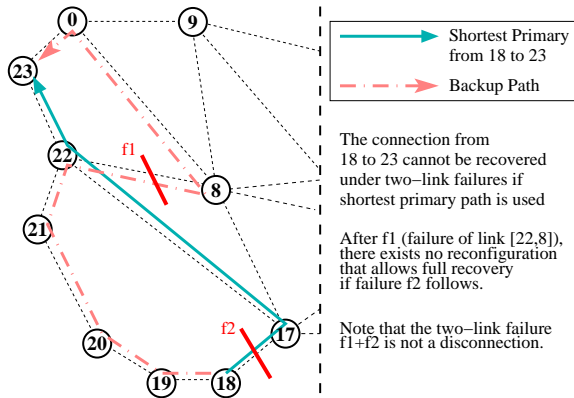


Fig. 6. Example of non-critical link failure affecting two-link survivability due to poor choice of routes.

the goal is to reconfigure with minimum number of critical links.

Since many mesh networks contain nodes of degree two, completely disjoint path (CDP) triples may not exist between some node pairs (node pairs with one or both end nodes with degree two), which leads to capacity failures (reconfiguration failures). In the National network, there are 108 unordered node pairs for which disjoint triples cannot be found. For these connections, CDP reconfiguration will fail. However, capacity failures that are not also disconnection failures can be addressed by finding what we term *maximally disjoint paths* (MDP). Figure 7 illustrates MDP selection. Dynamic protection reconfiguration with MDP allows the network to support optimal multiple failure survivability. The algorithm for computing MDP is simple and is outlined in Figure 8.

It is also interesting to note that, depending on the network topology and the routing choices made, failures of non-critical links can also impair the network's ability to recover from multiple failures. Figure 6 illustrates the effect of choosing the shortest path for the connection from node 18 to 23. Notice

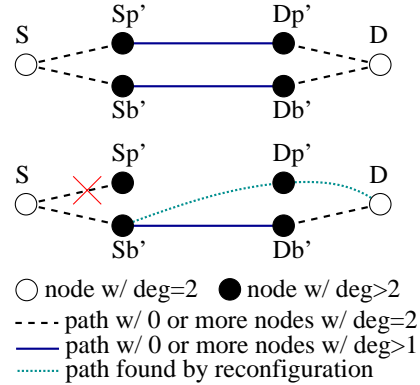
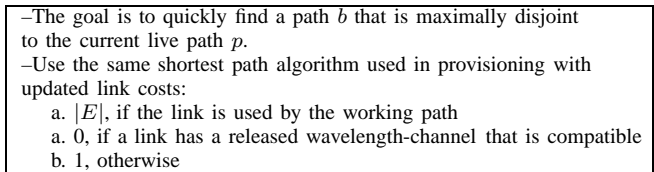
Fig. 7. Maximally disjoint path selection. After a link failure and dynamic protection reconfiguration, only link(s) between S and Sp' is(are) left vulnerable to a second link failure.

Fig. 8. Simple algorithm for computing min cost - Maximally Disjoint Paths

that f_1 occurs on a link that is not critical to the flow between nodes 18 and 23. If an alternate route is chosen for the primary and backup pairs (18-17-8-0-23 and 18-19-20-21-22-23), only critical link failures impair the network's survivability (in the two-link failure scenario).

V. EVALUATION

This section provides the details and the results from our investigation of multiple failure survivability and dynamic protection reconfiguration. We study four protection schemes—DPP, SPP, DLP, and SLP—using five representative networks shown in figures 3 and 9. Before discussing the results, we




Fig. 9. Sample Networks.

first provide the evaluation details. In this section only, we use DPR to refer to dynamic protection reconfiguration for brevity.

Since the results on different networks were similar, we present the results of the National network in this section, and, for clarity of presentation, we show the results of the other networks in the Appendix. The Arpanet network is different because it is three link connected.

A. Network and Traffic Model

To some extent, failure impact depends on the network traffic conditions. Therefore, study of different protection algorithms requires a fair and consistent basis for comparison. We assume uniform traffic demands, which can effectively aid in capturing the different characteristics of protection algorithms.

We consider on-line provisioning with uniformly distributed full-mesh traffic demands, and we assume that the network is optically opaque and capable of full wavelength conversion. On-line provisioning means that we have no knowledge of future demands, and cannot reroute existing connections on the network to optimize provisioning upon receipt of a new request. Each request is assumed to be a bidirectional connection with a uniformly distributed demand of one connection between each source and destination. $(N \times (N-1))/2$ bidirectional requests are routed in random order to simulate an on-line provisioning process. Although, in practice, the demands may not be uniformly distributed among different requests, we believe that studying uniformly distributed traffic demands is sufficient in that it shows the characteristics of different protection schemes for the purpose of studying multiple failure survivability. We assume that each λ -channel has a cost of 1 in terms of calculating capacity. The total cost of capacity is therefore the sum of the overall of working paths and the total number of the reserved protection wavelengths. Although a uniformly distributed traffic demand is assumed for evaluation in this work, our metrics, the failure classification scheme, and DPR assume nothing about the traffic model.

B. Protection Routing

The details of our routing algorithms for the four protection schemes are described in this section. We evaluated the protection schemes without DPR, with DPR using CDP selection (DPR-CDP), and finally with DPR using MDP (DPR-MDP) selection. Our primary goal in path selection is to reduce the capacity cost, therefore protection capacity is used as the primary metric. When applicable, path length is also used to break a tie between choices that have the same capacity cost. Shorter paths are preferred over longer ones. Path hit failures and broken path failures are a function of the length of the recovery routes. Therefore, selecting shorter recovery routes may allow some connections to avoid path hit and broken path failures (without DPR). However, for shared protection schemes, allowing longer paths may reduce protection capacity because average sharing is increased. Therefore, there is a tradeoff between capacity and the multiple failure survivability. For our experiments, we first find the shortest paths available in the network. Then, we let the algorithm consider paths that are some number hops longer than the length of a shortest path. We report the results on using some extra number of hops that allowed the most reduction in capacity with least impact on multiple failure survivability. Because allowing extra hops in dedicated protection is meaningless, only shortest paths were used.

First, we route the full mesh demand. Then, to calculate the average normalized vulnerability and recovery ratios as well as DPR capacity cost, we simulate all possible two-link failure scenarios. DPR is used after a first link failure, and the average DPR capacity cost is the average of the additional cost incurred by DPR over all first link failures in the network. Vulnerabilities and recovery ratios can be calculated after simulating the second link failure.

1) *Path Protection*: For path protection, we select a link-disjoint path pair, between the source and destination nodes, that minimizes capacity cost. For DPP, we select a path pair that has the shortest path lengths. For SPP, since protection paths can be shared, we perform local optimization on sharing protection capacity. In other words, paths are chosen such that the network is optimized assuming no knowledge of future connection requests. With no information about future routing demands in the on-line routing model, local optimization is optimal. Wavelength assignment for a backup route, therefore, is determined by evaluating all possible available wavelengths to maximize sharing (minimize cost). With the on-line routing model, it is also assumed that no previously routed lightpaths can be disrupted to perform rerouting optimizations. The path selection process is similar to the joint selection algorithm presented in [5].

As discussed in Section IV, DPR places a small constraint on routing. Because we hope to restore lightpaths after two links are cut, we must choose the primary/backup path pair that allows another link-disjoint path. Again, this constraint virtually has no effect on routing. Only a single connection in the LATA 'X' network was affected (provisioned with a different choice of paths).

Connections with either degree two source or destination

N=24 E=44	DPP	SPP	DLP	SLP
# extra hops	0	3	0	4
primary/backup path length	2.899	2.971	2.899	2.899
avg. link load	4.167	5.312	2.641	3.868
capacity cost	88.64	103.91	132.41	177.0
<i>disconnection</i>	0.013			
<i>no reconfiguration</i>				
<i>vulnerability</i>	<i>0.650</i>	<i>1.0</i>	<i>0.089</i>	<i>0.697</i>
<i>recovery ratio</i>	<i>0.913</i>	<i>0.900</i>	<i>0.966</i>	<i>0.963</i>
<i>path hit</i>	<i>0.464</i>	<i>0.799</i>	<i>0.078</i>	<i>0.289</i>
<i>broken path</i>	<i>0.464</i>	<i>0.799</i>	<i>0.078</i>	<i>0.289</i>
<i>blocked shared path</i>	–	<i>1.0</i>	–	<i>0.697</i>
<i>dynamic protection reconfiguration-CDP</i>				
<i>vulnerability</i>	<i>0.507</i>	<i>0.781</i>	<i>0.022</i>	<i>0.061</i>
<i>recovery ratio</i>	<i>0.947</i>	<i>0.940</i>	<i>0.993</i>	<i>0.993</i>
<i>capacity</i>	<i>0.259</i>	<i>0.304</i>	<i>0.018</i>	<i>0.043</i>
reconfiguration cost	38.64	142.36	50.27	77.86
	1%	5.6%	0.9%	2.8%
<i>dynamic protection reconfiguration-MDP</i>				
<i>vulnerability</i>	0.013			
<i>recovery ratio</i>	0.993	0.994	0.995	0.996
reconfiguration cost	80.18	175.96	123.82	135.82
	2%	6.9%	2.1%	4.9%

TABLE I

NORMALIZED VULNERABILITY VALUES FOR THE NATIONAL NETWORK. VULNERABILITY MEASURES ARE IN ITALICS.

nodes, or both, do not have three CDPs. For DPR-CDP, these connections result in capacity failures. For DPR-MDP, we use the MDP selection technique presented in the previous section.

2) *Link Protection*: DLP and SLP were also set up in a manner akin to DPP and SPP. For DLP, shortest paths for both primary and all backup paths were chosen to achieve efficiency in both recovery speed and capacity cost. For SLP, backup paths for individual links in the primary paths were found in a manner akin to SPP. Results show that our implementation is consistent with results found in the literature [8], [34], [35], with SPP providing 20%~50% savings in capacity over SLP while using, on average over the five sample networks, less than 50% of the capacity used for primary channels.

CDP triple selection is different from path protection in that, with link protection, protection routes protect each individual link in a lightpath. Therefore, CDPs or MDPs, are found for each link using instead of source and destination node pairs.

C. Results

First, Table I and Tables IV~VI (presented in the Appendix) show the average primary and protection path lengths, capacity cost for each protection scheme without DPR, as well as the average link load. Average backup path lengths are presented in terms of hop counts. Average link load is the average number of both primary and protection paths provisioned on a link. This measure shows how many connections are affected, on average, by link failures. We also show the vulnerability measure applied to the failure classification scheme and the recovery ratios for each protection scheme (without DPR, with DPR-CDP, and with DPR-MDP). Average DPR cost is also reported. We use the National network to discuss the results.

	National	ARPANET	N.J. LATA	COST 239	LATA 'X'
SPP	1.27	1.07	1.17	1.40	1.11
SLP	1.46	1.41	1.14	1.31	1.93

TABLE II

AVERAGE BACKUP PATH LENGTH OVERHEAD IN TERMS OF NUMBER OF HOPS FOR SHARED PROTECTION ALGORITHMS OVER DEDICATED PROTECTION.

Fundamental failures make up a small portion of network vulnerabilities at 0.013 for the National network, which is most heavily affected by physical dissections among the five sample networks.

The results indicate that most unrecoverable second link failures are caused by path hit failures and broken path failures. Vulnerability measures resulting from path hit failures and broken path failures have the exact same value confirming our intuition that these failures are only different in the ordering of the failure of the two links. These failures are directly affected by the lengths (in terms of number of hops) of the recovery paths as more link failures affect protection paths of longer length, leaving a network more susceptible to additional failures. Similarly, average link load shows the number of primary and protection paths that are affected when a link fails. More paths are affected by a failure in a network with a higher average link load, and the lengths of protection paths have a direct impact on this measure. Therefore, path hit failures and broken path failures are some function of the lengths of the protection paths. Many two link failures can also cause both path hit and broken path failures. There is an interesting difference in algorithmic properties between PP and LP highlighted by the difference in vulnerabilities from algorithmic failures. DPP and SLP for the National (and also N.J. LATA) network, despite similar average backup path lengths, have different path hit/broken path vulnerabilities. The impact of a two-link failure is smaller for LP because it depends on the operation of links that are more local to the point of failure whereas in PP, recovery may depend on backup paths that require the use of links across the entire network.

Intuitively, optimizing protection capacity cost can significantly degrade multiple failure survivability for a network. The effect of sharing optimization is two fold. First, in shared protection schemes, protection paths are elongated by a factor 1.2 on average over the five networks for PP and an average of 1.45 for LP. Penalties incurred by SPP and SLP for each network is shown in table II. In efforts to achieve more efficient use of wavelength channels in the network, longer backup paths may be provisioned to allow more efficient sharing between different connections. As discussed above, longer recovery paths expose the network to more path hit failures and broken path failures. Allowing longer protection paths has a huge effect on multiple failure survivability. Second, blocked shared path failures make the network more susceptible to additional failures. SPP suffers significantly more from blocked shared path failures compared to SLP. One very interesting result is that for SPP, the network is vulnerable to every second link failure (network vulnerability of 1.0)

without DPR. Blocked shared paths dominate the network's susceptibility to second failures. For SPP, blocked shared path failure vulnerability is 1.0, which shows that at least one protection wavelength is shared between connections that are affected by a two-link failure. In contrast, SLP does not suffer as much from blocked shared path failures. The difference, again, is due to the fact that in link protection, recovery routes are closer to the point of failures. LPs have larger link loads compared to PPs, but a much smaller fraction of the connections that are affected by the first failure are affected by the second failure. Generally, LP handles multiple failures better than PP due to the fact that recovery is dependent on a smaller section of the network that is more local with respect to the failed links.

DPR maximizes multiple failure survivability of a network with little additional capacity. For the National network, DPR-CDP requires no more than 5.6% additional capacity with the full mesh demand routed, and DPR-MDP requires up to 6.9%. With DPR, a protection scheme can optimize for capacity without affecting multiple failure survivability. With DPR-CDP, an average recovery ratio of over 99% can be achieved regardless of the choice of protection schemes. The optimal vulnerability of 0.013 for the National network, can be achieved with DPR-MDP. DPR significantly reduces vulnerability for all protection schemes, which allows network management to consider other tradeoffs between protection schemes, such as recovery speed, cost and management overhead, without worrying about multiple failure survivability.

1) *Discussion:* A significant advantage of DPR over multiple failure protection planning techniques is that DPR is not limited to handling a fixed number of link failures (for example, two-link failures) and can cover arbitrary number of failures without reserving excess protection wavelengths. It is important to note that with a failure scenario involving three or more links, the impact of disconnection failures starts to dominate and it becomes increasingly difficult to recover lightpaths. The two-link failure model is very useful in the sense that it captures the necessary details and the characteristics of multiple failure protection techniques and different protection schemes under multiple failures.

It is interesting to note that investigating node failures in our evaluation of multiple failure survivability provides little additional insight. The key difference between link and node failures is that more lightpaths are affected by a failure of a node because a node failure is equivalent to a failure involving failures of all the links adjacent to the node. Therefore, node failures incur higher capacity cost, but the recovery ratio remains more or less the same (average of the recovery ratio of the adjacent links). Protection capacity may also be slightly affected by node failures because fewer node-disjoint paths exist compared to link-disjoint paths in mesh networks. Quantifying the effect of the choice over different number of paths in shared protection schemes is out of the scope of this work.

VI. CONCLUSION

In order to further guarantee high quality of services for the increasing communications demands, we must be able to

maximize network utilization even in the event of failures. Graceful degradation of services must be considered in designing more robust and dependable future networks. Multiple failure survivability is a direct measure of a network's ability to operate effectively under failures.

Pre-planning multiple recovery routes and reserving enough wavelengths to allow a network to recover from multiple failures can be expensive in terms of protection capacity. In contrast, dynamic protection reconfiguration can allow a network to achieve high survivability under multiple failures while utilizing little additional capacity.

Our results showed that a large number of two-link failures that result in complete outages for some connections were caused by path hit and broken path failures (a two-link failure affects both the primary and backup routes for some connections). Shared protection schemes are much more capacity-efficient compared to dedicated protection schemes—by over 30% for path protection and over 45% for link protection—but suffer significantly in terms of multiple failure survivability. Shared path protection is vulnerable to roughly 35% more links compared to dedicated path protection, and shared link protection is vulnerable to about 57% more links compared to dedicated link protection. This difference is due to blocked shared failures and increased protection path lengths in shared protection schemes. Shared protection schemes are much more capacity-efficient however.

On five representative networks, dynamic protection reconfiguration required less than 4.4% additional capacity for dedicated path protection, less than 10% for shared path protection, less than 4.3% for dedicated path protection, and less than 7.5% for dedicated path protection. Compared to networks configured without dynamic protection reconfiguration, we found that the network vulnerability was reduced by over 90%. Dynamic protection reconfiguration allows networks to operate with the maximum multiple failure survivability for a given network topology where only disconnections (topological separation) can leave a lightpath unrecoverable. The average number of successfully recovered channels under two-link failures was over 99%.

The choice of protection schemes offer tradeoffs in terms of recovery speed, capacity and equipment cost, management overhead, and multiple failure survivability. Dynamic protection reconfiguration allows different protection schemes to achieve about the same level of multiple failure protection with little additional capacity. With DPR, tradeoffs between protection schemes, such as recovery speed, cost and management overhead, can be considered without affecting multiple failure survivability.

APPENDIX

Tables IV~VI show the evaluation results. The Arpanet network is three-link connected. Therefore, we can get 100% recovery ratio (it is not affected by fundamental failures). Maximally disjoint path formulation is not used with the Arpanet network because completely disjoint path triples can always be found between all node pairs.

N=20 E=32	DPP	SPP	DLP	SLP
# extra hops	0	1	0	4
primary/backup path length	2.763	2.805	2.753	2.753
avg. link load	4.121	4.411	3.505	4.945
capacity cost	81.75	85.69	147.25	194.31
disconnection	0.0			
<i>no reconfiguration</i>				
<i>vulnerability</i>	<i>0.812</i>	<i>1.00</i>	<i>0.202</i>	<i>0.812</i>
recovery ratio	88.9%	88.5%	94.6%	04.9%
<i>path hit</i>	<i>0.696</i>	<i>0.783</i>	<i>0.135</i>	<i>0.393</i>
<i>broken path</i>	<i>0.696</i>	<i>0.783</i>	<i>0.135</i>	<i>0.393</i>
<i>blocked shared path</i>	-	0.998	-	0.8
<i>dynamic protection reconfiguration-CDP</i>				
<i>vulnerability</i>	0.0			
recovery ratio	100%			
<i>capacity</i>	0.0			
reconfiguration cost	115.06	174.88	201.88	147
	4.4%	10.5%	4.3%	7.5%

TABLE III

EVALUATION RESULTS FOR THE ARPANET NETWORK. VULNERABILITY MEASURES ARE IN ITALICS. CDPs CAN BE FOUND FOR ALL NODE PAIRS BECAUSE ARPANET IS THREE-LINK CONNECTED.

N=11 E=23	DPP	SPP	DLP	SLP
# extra hops	0	1	0	1
primary/backup path length	1.745	1.745	1.745	1.745
avg. link load	2.309	2.691	2.0	2.281
capacity cost	19.39	21.22	25.04	27.39
disconnection	0.012			
<i>no reconfiguration</i>				
<i>vulnerability</i>	<i>0.431</i>	<i>0.905</i>	<i>0.194</i>	<i>0.727</i>
recovery ratio	90.6%	89.6%	93.7%	93.5%
<i>path hit</i>	<i>0.273</i>	<i>0.372</i>	<i>0.15</i>	<i>0.251</i>
<i>broken path</i>	<i>0.273</i>	<i>0.372</i>	<i>0.15</i>	<i>0.251</i>
<i>blocked shared path</i>	-	0.846	-	0.692
<i>dynamic protection reconfiguration-CDP</i>				
<i>vulnerability</i>	<i>0.328</i>	<i>0.427</i>	<i>0.03</i>	<i>0.077</i>
recovery ratio	94.1%	93.4%	98.7%	98.4%
<i>capacity</i>	<i>0.206</i>	<i>0.245</i>	<i>0.024</i>	<i>0.055</i>
reconfiguration cost	5.04	13.57	3.04	9.13
	1.1%	4.3%	0.5%	2.8%
<i>dynamic protection reconfiguration-MDP</i>				
<i>vulnerability</i>	0.012			
recovery ratio	98.9%	99%	99.2%	99.3%
reconfiguration cost	5.04	13.57	13.91	10.17
	1.1%	4.3%	2.4%	3.1%

TABLE IV

EVALUATION RESULTS FOR THE N.J. LATA NETWORK. VULNERABILITY MEASURES ARE IN ITALICS.

N=19 E=37	DPP	SPP	DLP	SLP
# extra hops	0	3	0	2
primary/backup path length	2.24	2.304	2.24	2.24
avg. link load	3.164	4.345	2.196	2.885
capacity cost	49.95	61.46	66.16	80.43
disconnection	0.011			
<i>no reconfiguration</i>				
<i>vulnerability</i>	<i>0.592</i>	<i>0.998</i>	<i>0.096</i>	<i>0.670</i>
recovery ratio	92.3%	90.8%	96%	95.7%
<i>path hit</i>	<i>0.42</i>	<i>0.626</i>	<i>0.083</i>	<i>0.218</i>
<i>broken path</i>	<i>0.42</i>	<i>0.626</i>	<i>0.083</i>	<i>0.218</i>
<i>blocked shared path</i>	-	0.998	-	0.662
<i>dynamic protection reconfiguration-CDP</i>				
<i>vulnerability</i>	<i>0.464</i>	<i>0.7</i>	<i>0.023</i>	<i>0.056</i>
recovery ratio	94.8%	94%	99.2%	99.1%
<i>capacity</i>	<i>0.293</i>	<i>0.35</i>	<i>0.019</i>	<i>0.042</i>
reconfiguration cost	16.32	62.05	27.14	34.32
	0.9%	5.2%	1.1%	2.6%
<i>dynamic protection reconfiguration-MDP</i>				
<i>vulnerability</i>	0.011			
recovery ratio	99.1%	99.4%	99.5%	99.5%
reconfiguration cost	24.49	76	56.11	45.19
	1.3%	6.3%	2.3%	3.4%

TABLE V

EVALUATION RESULTS FOR THE COST 239 NETWORK. VULNERABILITY VALUES ARE IN ITALICS.

N=28 E=47	DPP	SPP	DLP	SLP
# extra hops	0	2	0	6
primary/backup path length	3.286	3.505	3.286	3.286
avg. link load	4.669	5.188	2.861	4.415
capacity cost	127.96	139.83	204.04	286.17
disconnection	6014	4040	9590	4434
disconnection	0.006			
<i>no reconfiguration</i>				
<i>vulnerability</i>	<i>0.747</i>	<i>1.0</i>	<i>0.09</i>	<i>0.605</i>
recovery ratio	90.8%	89.9%	96.7%	96.5%
<i>path hit</i>	<i>0.586</i>	<i>0.8</i>	<i>0.072</i>	<i>0.265</i>
<i>broken path</i>	<i>0.586</i>	<i>0.8</i>	<i>0.072</i>	<i>0.265</i>
<i>blocked shared path</i>	-	1.0	-	0.605
<i>dynamic protection reconfiguration-CDP</i>				
<i>vulnerability</i>	<i>0.644</i>	<i>0.789</i>	<i>0.026</i>	<i>0.090</i>
recovery ratio	95.3%	95.1%	99.3%	99.3%
<i>capacity</i>	<i>0.297</i>	<i>0.329</i>	<i>0.019</i>	<i>0.062</i>
reconfiguration cost	109.7	266.89	180	169.53
	1.8%	6.6%	1.9%	3.8%
<i>dynamic protection reconfiguration-MDP</i>				
<i>vulnerability</i>	0.006			
recovery ratio	99.7%	99.8%	99.7%	99.8%
reconfiguration cost	144.51	285.96	249.75	254.17
	2.4%	7.1%	2.6%	5.7%

TABLE VI

EVALUATION RESULTS FOR THE LATA 'X' NETWORK. VULNERABILITY MEASURES ARE IN ITALICS.

REFERENCES

- [1] T. E. Stern and K. Bala, *Multiwavelength Optical Networks; A Layered Approach*, Prentice-Hall, Upper Saddle River, NJ, 2000.
- [2] S. Ramamurthy, L. Sahasrabudde, and B. Mukherjee, "Survivable WDM mesh networks," *IEEE/OSA JLT*, vol. 21, no. 4, pp. 870–83, April 2003.
- [3] R. S. K. Chang, C. P. Botham, D. Johnson, G. N. Brown, M. C. Sinclair, M. J. O'Mahony, and I. Hawker, "A multi-layer restoration strategy for reconfigurable networks," in *Proc. of IEEE GLOBECOM*, 1994, vol. 3, pp. 1872–78.
- [4] H. Kobriniski and M. Azuma, "Distributed control algorithms for dynamic restoration in dcs mesh networks: Performance evaluation," in *Proc. of IEEE GLOBECOM*, 1993, vol. 3, pp. 1584–8.
- [5] C. Xin, Y. Ye, S. Dixit, and C. Qiao, "A joint working and protection path selection approach in WDM optical networks," in *Proc. of IEEE GLOBECOM*, 2001, vol. 4, pp. 2165–8.
- [6] A. Fumagalli and L. Valcarenghi, "The preplanned weighted restoration scheme," in *IEEE Workshop on High Performance Switching and Routing*, 2001, pp. 36–41.
- [7] X. Su and C. Su, "An online distributed protection algorithm in WDM networks," in *Proc. of IEEE ICC*, 2001, vol. 5, pp. 1571–5.
- [8] B. Caenegem, B. Wauters, and P. Demeester, "Spare capacity assignment for different restoration strategies in mesh survivable networks," in *Proc. of IEEE ICC*, 1997, vol. 1, pp. 288–92.
- [9] R. Ramamurthy, Z. Bogdanowicz, S. Samieian, D. Saha, B. Rajagopalan, S. Sengupta, S. Chaudhuri, and K. Bala, "Capacity performance of dynamic provisioning in optical networks," *IEEE/OSA JLT*, vol. 19, no. 1, pp. 40–8, 2001.
- [10] H. Fujii and N. Yoshikai, "Double search self-healing algorithm and its characteristics," *Electronics and Communications in Japan-Part 1*, vol. 77, no. 3, pp. 75–8, 1994.
- [11] W. D. Grover, "The selfhealing network," in *Proc. of IEEE GLOBECOM*, 1987, vol. 2, pp. 1090–5.
- [12] T. H. Wu, "A passive protected self-healing mesh network architecture and applications," *IEEE/ACM TON*, vol. 2, no. 1, pp. 40–52, February 1994.
- [13] E. Bouillet, K. Kumaran, G. Liu, and I. Saniee, "Wavelength usage efficiency versus recovery time in path-protected DWDM mesh networks," in *Proc. of IEEE/OSA OFC*, 1998.
- [14] M. Kodialam and T.V. Lakshman, "Dynamic routing of locally restorable bandwidth guaranteed tunnels using aggregated link usage information," in *Proc. of IEEE INFOCOM*, 2001, vol. 1, pp. 376–85.
- [15] S. Kim and S. S. Lumetta, "Capacity-efficient protection with fast recovery in optically transparent mesh networks," in *Proc. of IEEE BROADNETS*, Oct. 2004.
- [16] S. Kim and S. S. Lumetta, "Restoration of all-optical mesh networks with path-based flooding," *IEEE/OSA JLT*, vol. 21, no. 11, November 2003.
- [17] M. Médard, S. G. Finn, R. A. Barry, W. He, and S. S. Lumetta, "Generalized loop-back recovery in mesh networks," *IEEE Trans. on Networking*, vol. 10, no. 1, pp. 153–64, February 2002.
- [18] W. D. Grover and D. Stamatelakis, "Cycle-oriented distributed pre-configuration: Ring-like speed with mesh-like capacity for self-planning network reconfiguration," in *Proc. of IEEE ICC*, 1998, vol. 1, pp. 537–43.
- [19] G. Ellinas, A. G. Hailemariam, and T. E. Stern, "Protection cycles in mesh WDM networks," *IEEE JSAC*, vol. 18, no. 10, Oct. 2000.
- [20] P. Ho and H.T. Mouftah, "Slsp: A new path protection scheme for the optical internet," in *Proc. of IEEE/OSA OFC*, 2001.
- [21] D. Xu, Y. Xiong, and C. Qiao, "Novel algorithms for shared segment protection," *IEEE JSAC*, vol. 21, no. 8, pp. 1320–31, Oct. 2003.
- [22] S. S. Lumetta and M. Médard, "Towards a deeper understanding of link restoration algorithms for mesh networks," in *Proc. of IEEE INFOCOM*, April 2001.
- [23] H. Choi, S. Subramaniam, and H. A. Choi, "On double-link failure recovery in WDM optical networks," in *IEEE Infocom*, June 2002.
- [24] M. Clouqueur and W. D. Grover, "Mesh-restorable networks with complete dual failure restorability and with selectively enhanced dual-failure restorability properties," in *SPIE Opticomm*, July 2002.
- [25] W. He, M. Sridharan, and A. K. Somani, "Capacity optimization for surviving double-link failures in mesh-restorable optical networks," in *Proceedings of OptiComm*. SPIE, July 2002, vol. 4874, pp. 13–24.
- [26] S. Kim and S. S. Lumetta, "Evaluation of protection reconfiguration for multiple failures in optical networks," in *Proc. of IEEE/OSA OFC*, 2003.
- [27] W. He and A. Somani, "Path-based protection for surviving double-link failures in mesh-restorable optical networks," in *IEEE Globecom*, Dec. 2003.
- [28] D. A. Schupke and R. G. Prinz, "Capacity efficiency and restorability of path protection and rerouting in WDM networks subject to dual failures," *Photonic Network Communication*, vol. 8:2, 2004.
- [29] M. Fredericks, P. Datta, and A. K. Somani, "Evaluating dual-failure restorability in mesh-restorable WDM optical networks," in *IEEE ICCCN*, Oct. 2004.
- [30] P. Sebos, J. Yates, G. Hjalmtysson, and A. Greenberg, "Auto-discovery of shared risk link groups," in *Proc. of IEEE/OSA OFC*, 2001.
- [31] S. Kim and S. S. Lumetta, "Addressing node failures in all-optical networks," *OSA Journal of Optical Networking*, vol. 1, no. 4, pp. 154–63, April 2002.
- [32] L. Ling and A. K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE/ACM TON*, pp. 779–86, 1999.
- [33] L. Li, S. Wang, and S. Xu, "Dynamic routing and assignment of wavelength algorithms in multi-fiber wavelength division multiplexing networks," 1999.
- [34] Jeyakesavan Veerasamy, S. Venkatesan, and Jay C. Shah, "Spare capacity assignment in telecom networks using path restoration and further improvement using traffic splitting," *The Journal of Systems and Software*, vol. 47, no. 1, pp. 27–33, 1999.
- [35] R. R. Iraschok, M. H. MacGregor, and W. D. Grover, "Optimal capacity placement for path restoration in stm or atm mesh-survivable networks," *IEEE/ACM TON*, vol. 6, no. 1, 1998.