

Addressing node failures in all-optical networks

Sun-il Kim and Steven S. Lumetta

*Departments of Computer Science and Electrical and Computer Engineering,
University of Illinois at Urbana-Champaign, Coordinated Science Laboratory
1308 W. Main Street Urbana, Illinois 61801
sunilkim@uiuc.edu; lumetta@uiuc.edu*

Received 8 November 2001; revised manuscript received 15 March 2002

We investigate the effectiveness of link-protection schemes in terms of their ability to handle node failures in all-optical networks. We focus on node recovery, using a ring-based scheme (double-cycle cover; DCC) and generalized loopback. Comparisons are made by introduction of metrics that measure a protection scheme's effectiveness in handling node failures. We also introduce measures that quantify *failure impact*, the degree to which a network is left vulnerable to a failure after a previous failure has been successfully recovered. Comparison with DCC for five sample networks shows that generalized loopback is more robust against node failures and can provide the same level of reliability while using 16% less capacity on average over five sample networks.

© 2002 Optical Society of America

OCIS codes: 060.4250, 060.4510.

1. Introduction

Reliability of the communications infrastructure is becoming increasingly important. We must understand the issues in the design of reliable, high-speed networks to support our future communication needs. Algorithms must be easily scalable, since networks change constantly, and must be able to perform recovery on a diverse set of topologies. We must be able to handle changes locally, without global coordination, which means that network reliability must be provided through distributed operations. Algorithms must also be efficient in terms of their capacity, in order to reduce operation costs. Finally, recovery speed is becoming increasingly important, since disruptions become more severe with increases in network capacity.

All-optical networks (AONs) offer several advantages over optically opaque networks. AONs offer faster switching with the absence of electronic and photonic processing delays that pose a bottleneck on the time of the total transmission. AONs can also handle signals with different data rates, protocols, and formats and are not restricted to certain standards as in opaque networks. However, they have limited functionality in wavelength conversion, signaling capability, and detailed performance monitoring.^{1,2} These limitations complicate the application of many efficient recovery schemes studied in the literature.³

To provide a desired level of reliability in AONs and to address the important issues discussed above, we must be able to analyze quantitatively the advantages and the trade-offs of different protection schemes. In this paper we investigate the effectiveness of link-protection schemes in terms of their ability to handle node failures. The degree to which a network is left vulnerable to a failure after a successful recovery of the first failure is also investigated. We use the term *failure impact* to refer to this concept. Node recovery is generally much more complicated than link recovery, and the differences must be considered carefully in the analysis.

We investigate node recovery, using generalized loopback, which can be implemented in AONs and meets the stated goals. To aid this evaluation, we introduce measures that quantify the effects of providing node protection on routing path lengths and network con-

nectivity. We study failure impact by considering *double failures*. A double failure consists of a combination of two independent single failures. We assume that the second failure occurs long enough after the first to allow the normal recovery to be completed, but before any physical repair can be made. We develop intuitive measures of failure impact and provide comparisons between a ring-based link-protection scheme called double-cycle cover (DCC)⁴ and the generalized loopback scheme.⁵

For the remainder of the paper, we first proceed with the discussion of different protection schemes. Node failures are discussed in Section 3. Section 4 discusses the effectiveness of link-protection algorithms in terms of handling node failures. We also present our measures of failure impact and the simulation results. Generalized loopback incurs a penalty in terms of bandwidth efficiency when providing protection against node failures. We quantify this penalty in Section 5. Finally, we provide our conclusions.

2. Background

Significant amounts of data and voice traffic can be lost to network failure. Such failures include channel failures, link failures, and failures of nodes (optical switches). Channel failures caused by card failures at a port of an optical switch are the most common type of failure in optical networks. Link failures (fiber cuts caused by wayward backhoes, amplifier failures, and the like) are also common and can result in failures of all the channels that are carried on the fiber. We ignore cases in which logically separate links are placed in the same conduit; a cut through such a conduit can result in two links failing simultaneously.⁶ Node failures are less common but can cause failures of all the links that are adjacent to the node.

The implications of data losses that result from such failures grow more severe as more and more crucial applications are deployed onto wavelength-division-multiplexing (WDM) networks. Effective methods to address failures are therefore necessary in order to minimize the effect of failures and to provide network services with a certain level of guaranteed reliability. Protection and restoration are the two main approaches that address failures in fiber-optic networks.^{7,8}

Restoration, which addresses failures by locating free links for backup after a failure occurs, is not considered in our study, for two reasons. First, the signaling required for dynamical discovery of a backup path makes restoration slow; synchronous optical network (SONET) digital cross-connect systems, for example, often target a 2-s restoration time.^{9–11} Second, the limitation of signaling capabilities makes application of restoration techniques difficult in AONs.³

Protection can be divided into two types: path protection and link protection. Path protection requires knowledge of the whole path and selection of a backup path that is disjoint from the primary path. In dedicated path protection the traffic is simply switched over to the backup path in the event of a failure. In shared path protection, channels used for backup paths are shared among different lightpaths. Shared path protection is much more efficient than dedicated path protection in terms of capacity usage. However, its application is difficult in AONs because of the signaling that is required for setting up the intermediate switches in backup paths after a failure.³ This signaling also makes shared path protection slower.

Link protection, in contrast, just routes around the failed link instead of finding an entirely different path from source to destination. Link protection is faster than path protection, since the time it takes to discover a failure is the propagation delay on a single link for link protection, whereas for path protection it is the delay on the entire path in the worst case. For this reason we focus on link protection. Capacity efficiency is also an important issue. Note that there is a trade-off between link protection and path protection in terms of the speed of recovery and capacity efficiency. Link protection provides faster recovery, and




Fig. 1. (a) DCC, (b) generalized loopback.

path protection is more capacity efficient.⁷

Shared link protection is similar to shared path protection, and for the same reasons it is not an attractive solution for AONs. Most link-protection algorithms use rings to design part or all of the network.^{12–15} Ring covers for meshes are widely used for their simplicity in operation, low management overhead, and fast recovery, but they have drawbacks such as restricted traffic routing, difficult optimization problems, and the possibility of a full network reconfiguration in response to a minor extension.^{12, 16–18} The p -cycles approach extends ring covers¹⁹ to provide more efficient link protection but focuses on optically opaque networks with electronic (optical–electronic–optical) regeneration or wavelength conversion and low-granularity variable-capacity links, whereas we study optically transparent networks with full-fiber or full-wavelength granularity with equal capacity links. Restricting the p -cycles approach to this context reduces its efficacy and does not produce a fair comparison; thus a full comparison is outside the scope of this paper. We next discuss two protection schemes, generalized loopback and the ring-based scheme DCC, which meet our stated goals.

2.A. Double-Cycle Cover

A DCC covers every link in a network with exactly two rings with opposite directions, thus solving several of the problems with earlier ring cover approaches.⁴ DCCs, however, retain the scalability problems of previous ring-based approaches.

Figure 1(a) shows how a DCC is set up. Each link is covered by two protection cycles (or rings) shown in solid and dashed lines with an arrow to indicate the direction of the traffic flow. DCCs are simple to implement and manage and provide rapid recovery through the use of an automatic protection switching (APS) mechanism. Failures are recovered by means of rerouting signals from the broken, working channels to protection channels. The following example illustrates how APS in DCCs works. Consider the failure of the link [V, Z] in Fig. 1(a). Channels carrying lightpaths that traverse the broken link are switched over to the protection fibers. We assume that each connection is bidirectional and that the two protection cycles that cover the broken link are used to recover each direction.

2.B. Generalized Loopback

The generalized loopback algorithm also provides protection through an APS-like mechanism and is applicable to any two-node redundant graph.⁵ Like DCCs, protection fibers (channels) are reserved. In generalized loopback, a pair of conjugate digraphs are used for routing primaries and protection reservation. Routing is performed with the *primary digraph*, and protection is provided through the use of its conjugate, the *secondary (or backup) digraph*. These digraphs are calculated once for a network before the network is put on line for the first time. Upon detection of a failure, nodes adjacent to the failure simply flood the preestablished backup digraph with backup traffic.

The following example illustrates how generalized loopback works. In Fig. 1(b) the

primary digraph is shown in solid lines and the secondary digraph is shown in dashed lines. This figure, for simplicity, shows only a single pair of conjugate digraphs that allow unidirectional communication. Two sets of these are used for bidirectional communication. Consider the failure of link [V, Z]. Channels carrying lightpaths that traverse the broken link are switched over to the channels on the backup digraph. The backup digraph is flooded with the backup traffic by node V, and the backup traffic finds its way to node Z. Note the two possible paths that can be taken by the backup traffic: (V-U-W-X-Z) and (V-U-W-Y-Z). The protocol⁵ used ensures that only the traffic that arrives at a node first is forwarded to the out ports. Traffic that arrives subsequently is simply discarded, and the node that sent out this traffic is notified by means of a negative acknowledgment message (NACK). On receiving a NACK, the node stops forwarding traffic to the corresponding out port. Therefore only one of the two possible paths is actually established.

Generalized loopback requires only a subset of the links in the backup digraph for protection against single failures⁵ and allows tuning between the capacity of a network and its ability to minimize the failure impact. A *minimal backup digraph* uses only the minimum number of links necessary to guarantee full robustness from single failures. Links in the backup digraph that are not included in a minimal backup digraph can either carry unprotected working traffic (in the event of a failure, recovery is not effected for this type of traffic) or be reserved for protection along with the minimal backup digraph. We call the links for which unprotected traffic can be carried on the backup digraph *noncritical links*. For the five networks used in our study, the minimal backup digraphs are shown in solid lines and noncritical links are shown in dashed lines in Figs. 2 and 3. There is a trade-off between using noncritical links and minimizing failure impact. This trade-off is examined in Subsection 4.B.

3. Node Failures

Node recovery is generally much more complicated than link recovery. We introduce the term *stream* for better illustrating the effect of node failures. A stream of a node represents a lightpath traversing that particular network node. Recovery algorithms must consider all streams of a node in the event of a node failure. A stream that starts or terminates at the failed node cannot be protected without having node (hardware) redundancy.

There may be cases in which only parts of the node fail and some parts remain functional. For example, only some of the links on an optical switch may cease to function because of some hardware failure in the switch. In such cases the protection scheme may be able to treat them as individual link failures and take appropriate measures to recover them. However, determining the nature and the extent of the failure may be difficult. There is another problem with allowing different failure modes for nodes: Most protection schemes do not support simultaneous recovery from several failures. Therefore we assume that none of the channels on the failed node remain active.

Link-protection algorithms typically protect only individual links. Node failures, although addressed by some ring-based algorithms, usually require complex synchronization for ring hopping in order to guarantee full network connectivity. Generalized loopback can handle node failures effectively, and the protocol requires no differentiation between link and node failures. In providing protection against node failures, both DCCs and generalized loopback can protect only a single stream per wavelength per node. Attempts to recover multiple streams of a failed node on the same wavelength can result in contention for backup capacity. Routing must therefore be done carefully in order to provide protection from all single-node failures. Consequently, there is a penalty in terms of capacity usage (network overbuild) when considering node protection. We quantify this penalty for generalized loopback in Section 5.




Fig. 2. (a) New Jersey local-access transport area (N.J. LATA). (b) Advanced Research Projects Agency network (ARPANET). (c) COST 239 [a European backbone network, Coopération Européenne dans le Domaine de la Recherche Scientifique et Technique (European Cooperation in the Field of Scientific and Technical Research)].

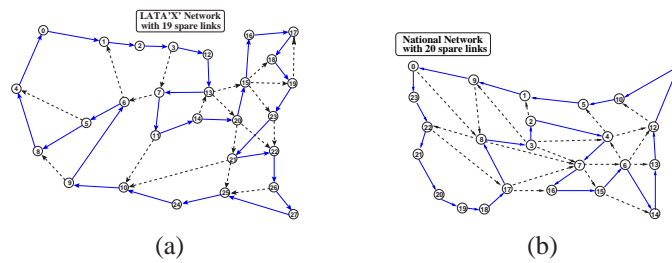


Fig. 3. (a) LATA “X”. (b) National.

4. Evaluation Measures

4.A. Protected Connectivity and Path-Length Expansion

In this section we introduce two measures that aid the study of node protection. *Robust connectivity* measures the global end-to-end protection capability of a network. Robust connectivity of a network is the average percentage of nodes to which connections from a node can be routed with full robustness to any single failure. Intuitively we want an algorithm to be able to support connections between any pair of source and destination nodes with full robustness to any single failures. *Robust path-length expansion* is used to understand the penalty of restricting routing in order to provide robustness to failures in terms of the expected increase in number of hops for end-to-end paths. Path-length expansion for a given algorithm measures the average over all pairs of nodes of the ratio of the shortest robust path between the nodes to the shortest unprotected path between nodes. Node pairs that cannot be connected via a path robust to single failures are ignored in the calculation; expansion averages only those paths that can be made robust.

Table 1 shows the results for robust connectivity and path-length expansion ratio. DCC has low robust connectivity compared with that of generalized loopback. Path-length expansion for DCC varies dramatically with the average and the maximum length of the ring used in the cover. The data in Table 1 show an interesting trade-off between the robust connectivity measure and the path-length expansion ratio. Because the source and the destination nodes of a robust path must lie on the same ring with DCCs, longer rings result in better end-to-end connectivity. Longer rings also increase path-length expansion, however, because the length of a robust path for a DCC depends on the length of the ring that covers both the source and the destination nodes.

A second trade-off related to cycle length involves the ability to address failure impact. After failure and recovery of a node in a ring, another node failure within the same ring

Table 1. Robust Connectivity–Path-Length Expansion^a

Graph	Node Degree	Generalized		
		Loopback	DCC (Longest Cycle /Average Cycle)	
N.J. LATA	4.2	100%/1.000	49%/1.000	(4/3.28)
COST 239	3.9	100%/1.045	34.3%/1.040	(5/3.5)
National	3.7	100%/1.063	76.7%/1.618	(18/3.95)
LATA “X”	3.4	100%/1.086	Not available	
ARPANET	3.2	100%/1.096	84.7%/1.326	(14/6.4)

^aFor DCC the parenthesized values represent the length of the longest cycle and the average length of cycles.

cannot be recovered and breaks the recovery path for the first failure. Creating longer rings thus leaves the network more vulnerable to double failures but improves robust connectivity for single failures.

For the generalized loopback algorithm, providing robust end-to-end connections has only a minor effect on the path lengths. Path-length expansion is some function of the node degree, and higher average node degree gives better path-length expansion ratios, but the difference is small.

4.B. Vulnerability and Exposure

Generalized loopback is more attractive than DCC in terms of supporting node recovery. In this section a set of metrics that can be used to understand the trade-off between capacity efficiency and the ability to minimize failure impact is discussed.

We next introduce metrics that are independent of the traffic model for networks to evaluate the two link-protection schemes. Using these traffic-independent metrics, we seek a general understanding of protection schemes’ ability to address failure impact. To avoid using traffic models, our measures evaluate an average of all possible streams per wavelength incident upon each node in the network.

A quantitative measure of the degree to which a single link failure leaves the network vulnerable to another link failure can be found in the literature.²⁰ This measure is extended to two-node failures and link-node failures. A *two-node failure* (or *link-node failure*) consists of two independent single failures (two nodes or a link and a node) in a network graph.

A stream s in node N is said to be *vulnerable* to a second stream t in node M if failure of N prevents recovery of s after failure of M and recovery of t . The average of vulnerabilities of all streams in a node then yields *node vulnerability* of a node, which represents the number of nodes to which the node is vulnerable. The node vulnerability of a network is the average node vulnerability over its nodes. Equation (1) is used to calculate the node vulnerability of a network:

$$\frac{1}{|V|} \sum_{n \in V} \left[\sum_{m \in V} \frac{\sum_{s \in S(n)} \sum_{t \in S(m)} f(s, t)}{|S(n)||S(m)|} \right],$$

$$f(s, t) = \begin{cases} 1 & \text{if } s \text{ or } t \text{ cannot be recovered} \\ 0 & \text{otherwise} \end{cases},$$

$$S(n) = \text{set of valid streams in node } n. \quad (1)$$

For link-node failures, *node vulnerability* and *link vulnerability* are defined similarly to the vulnerability discussed with two-node failures. We can derive link vulnerability by

simply multiplying node vulnerability by the ratio of number of nodes to number of links. For the rest of the paper we use the term vulnerability to mean node vulnerability. Vulnerability for link-node failures can be calculated with Eq. (2). Time ordering of double failures between a link and a node affects the vulnerability of a network. The failure function $f(s, l)$ reflects time ordering of failures: $f(s, l)$ denotes s failing first and l failing second:

$$\frac{1}{|E|} \sum_{l \in E} \left[\sum_{n \in V} \frac{\sum_{s \in S(n)} f(s, l)}{|S(n)|} \right],$$

$$f(s, l) = \begin{cases} 1 & \text{if } s \text{ or } l \text{ cannot be recovered} \\ 0 & \text{otherwise} \end{cases},$$

$$S(n) = \text{set of valid streams in node } n. \quad (2)$$

We measure vulnerability for connections with lengths of two or more hops. Consider a light path p_1 that traverses nodes x , y , and z . Node x may be the source of the connection, and node z may be the destination of the connection. After the failure of node y , the stream from x to z could be recovered by the protection algorithm. If either node x or z fails after the failure and recovery of node y , the connection cannot be restored, and p_1 is broken. This results in a minimum bound on vulnerability measure at 2.0.

Table 2. Vulnerability/First-Failure Exposure/Second-Failure Exposure^a

Graph	DCC	Generalized Loopback	
		(% Backup Links Used)	
N. J. LATA	3.05/2.24/2.82	2.71/2.40/2.36 (100)	3.01/2.49/2.36 (91.3)
COST 239	3.61/2.70/2.21	3.51/2.97/2.72 (100)	3.56/3.03/2.79 (94.6)
National	7.56/5.79/5.12	4.69/3.72/3.57 (100)	7.52/5.12/5.52 (79.5)
LATA "X"	Not avail.	6.73/4.32/4.78 (100)	—
ARPANET	10.87/7.57/8.89	7.03/4.34/5.17 (100)	9.49/6.39/7.24 (71.9)

^aParentthesized values show the percentage of network links that were reserved for protection.

We introduce another quantitative measure of failure impact that captures a notion of time ordering of two independent failures. *First-failure exposure* of a node represents cases in which a node fails first and all broken streams are not fully restored after a second failure. *Second-failure exposure* of a node represents cases in which the node fails second and the broken streams cannot be restored. First- and second-failure exposure of links are defined in the same manner.

Only streams that are protected from single failures are considered in order to get a fair comparison between different protection schemes. For example, DCCs have low connectivity, and the total number of protected streams is much less than with generalized loopback. See Table 1 for connectivity measures.

Measured values of vulnerability and exposures for two-node failures are shown in Table 2. Our results show that generalized loopback provides a comparable degree of failure impact while using 16% less capacity on average. Capacity is calculated by means of counting the number of links in a backup graph used by an algorithm. With equal capacity, generalized loopback improves vulnerability by an average of 22%. Characteristics of link-node failures are similar to those of two-node failures.

Figures 4(a) and 4(b) show that first-failure exposure and second-failure exposure cross for the N.J. LATA network. In a network with a high average node degree (essentially meaning more links and number of alternate paths available for recovery), if a second




Fig. 4. N.J. LATA: (a) two-node failure, (b) link-node failure.




Fig. 5. ARPANET: (a) two-node failure, (b) link-node failure.

failure breaks the first failure’s recovery path, the second failure is recovered more successfully. This is because generalized loopback reclaims links in the broken recovery path, which in turn can be used to recover the second failure. As more links are used for backup, the chance of second failure hitting first failure’s recovery path decreases, improving first-failure exposure. This is the result of a decrease in number of hops for recovery paths; fewer nodes and links in a recovery path implies a smaller exposure. Second-failure exposure also improves as more links become available for its own recovery. The same effect was seen on the COST 239 network.

In Figs. 5(a) and 5(b), second-failure exposure is always greater than first-failure exposure. The topology of the network affects how the two measures relate to each other. Since the minimal backup digraph of the ARPANET network consists of a Hamiltonian cycle, most second failures break the recovery path of the first failure and first failures prevent recovery of the second failure. Therefore, reclaiming links from a broken recovery path does not necessarily improve recovery for the second failure. The National network and the LATA “X” network have properties that are similar to those of the ARPANET network.

A sharp increase in measured values near the right-hand side of the figures shows an interesting change in properties of failure impact. Additions of the first few noncritical links to the minimal backup digraph creates a topology of several rings with few shared nodes, allowing failures to be recovered independently.

5. Discussion

The generalized loopback scheme is limited to a single stream per wavelength per node in routing when node protection is addressed. This restriction has an adverse effect on capacity efficiency and results in overbuilding of the network. We quantify this penalty by looking at the number of wavelengths used in a network by simulation of routing uniformly distributed full-mesh traffic demands. We use an on-line algorithm in the sense that we assume no knowledge of future demands and cannot reroute existing connections on the network to optimize provisioning. Each request is assumed to be bidirectional with a uniformly distributed demand of one wavelength between each source and destination. We route using a random order of $[N \times (N - 1)]/2$ bidirectional requests to simulate an on-line provisioning process.

Table 3. Considering Link-Only Protection Allows Routing of Multiple Streams

Graph	Number of Connections Routed	# λ Used for Link-Only Protection	# λ Used for Link and Node Protection
N.J. LATA	110	20	38
COST 239	342	70	88
National	552	110	138
LATA "X"	756	172	220
ARPANET	380	66	98

Link-only protection allows for routing arbitrary number of streams per wavelength per node, and therefore the number of wavelengths used for the primary routes is comparable with that for path protection. Results shown in Table 3 show that restricting routing to a single stream per wavelength per node results in a significant increase in the number of wavelengths used in the network. The reported values have a variance of $\sim 5\%$ over different random ordering of requests.

6. Conclusions and Future Research

Node protection is more complicated than link protection and requires a careful analysis. We have investigated node failures and presented reliability measures that effectively address node protection for link-protection schemes. Results of our investigation show significant advantages of generalized loopback in terms of robustness and capacity efficiency in addressing node recovery compared with DCC in AONs. Our results also show that generalized loopback can better address failure impact.

Another advantage of generalized loopback is that a good portion of backup capacity ($\sim 16\%$ on average) can be used to carry unprotected traffic while providing the same level of reliability as DCCs. This reassignment of capacity allows for efficient use of existing capacity in networks.

Protocol extensions to allow routing of multiple streams per wavelength per node in generalized loopback can improve its capacity efficiency and is left for future investigation. Quantifying the effect of reconfiguring the network after a failure is also interesting and can improve generalized loopback's ability to minimize failure impact. To study failure impact in other protection schemes, such as path protection, metrics that consider traffic models are needed.

References and Links

1. T. E. Stern and K. Bala, *Multiwavelength Optical Networks: A Layered Approach* (Prentice-Hall, Upper Saddle River, N.J., 2000).
2. Z. Zhang, J. Fu, D. Guo, and L. Zhang, "Lightpath routing for intelligent optical networks," *IEEE Netw.* (July/Aug. 2001), pp. 28–35.
3. S. Chaudhuri, E. Bouillet, and G. Ellinas, "Addressing transparency in DWDM mesh survivable networks," in *Optical Fiber Communication Conference*, Vol. 4 of OSA Trends in Optics and Photonics Series (Optical Society of America, Washington, D.C., 2001), paper TuO5.
4. G. Ellinas, A. G. Hailemariam, and T. E. Stern, "Protection cycles in mesh WDM networks," *IEEE J. Sel. Areas. Commun.* **1**, 152–156 (2000).
5. M. Médard, S. G. Finn, R. A. Barry, W. He, and S. S. Lumetta, "Generalized loop-back recovery in mesh networks," *IEEE Trans. Netw.* **10**, 153–164 (2002).

6. P. Sebos, J. Yates, G. Hjalmtysson, and A. Greenberg, "Auto-discovery of shared risk link groups," in *Optical Fiber Communication Conference*, Vol. 4 of OSA Trends in Optics and Photonics Series (Optical Society of America, Washington, D.C., 2001), paper WDD3-1.
7. S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks. I. Protection," in *Proceedings of the Conference on Computer Communications (Infocom)* (Institute of Electrical and Electronics Engineers, New York, 1999), Vol. 2, pp. 744–751.
8. S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks. II. Restoration," in *Proceedings of the Conference on Computer Communications (Infocom)* (Institute of Electrical and Electronics Engineers, New York, 1999) Vol. 3, pp. 2023–2030.
9. W. D. Grover, "The SelfHealing network," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)* (Institute of Electrical and Electronics Engineers, New York, 1987), Vol. 2, pp. 1090–1095.
10. J. Sosnosky, "Service application for SONET DCS distributed restoration," *IEEE J. Sel. Areas Commun.* **12**, 59–68 (1994).
11. T. H. Wu, "A passive protected self-healing mesh network architecture and applications," *IEEE/ACM Trans. Netw.* **2**(1), 40–52 (1994).
12. W. D. Grover, "Case studies of survivable ring, mesh and mesh-arc hybrid networks," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)* (Institute of Electrical and Electronics Engineers, New York, 1992), Vol. 1, pp. 633–638.
13. E. L. Hahne and T. D. Todd, "Fault-tolerant multimesh networks," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)* (Institute of Electrical and Electronics Engineers, New York, 1992), Vol. 1, pp. 627–632.
14. T. H. Wu, *Fiber Network Service Survivability* (Artech House, Norwood, Mass., 1992).
15. J. Shi and J. P. Fonseka, "Hierarchical self-healing rings," *IEEE/ACM Trans. Netw.* **3**(6), 690–697 (1995).
16. T.-H. Wu, D. J. Kolar, and R. H. Cardwell, "High-speed self-healing ring architectures for future interoffice networks," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)* (Institute of Electrical and Electronics Engineers, New York, 1989), Vol. 2, pp. 801–807.
17. O. J. Wasem, "Optimal topologies for survivable fiber optic networks using SONET self-healing rings," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)* (Institute of Electrical and Electronics Engineers, New York, 1991), Vol. 3, pp. 2032–2038.
18. O. J. Wasem, "An algorithm for designing rings for survivable fiber networks," *IEEE Trans. Reliab.* **40**, 428–432 (1991).
19. W. D. Grover and D. Stamatelakis, "Cycle-oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network reconfiguration," in *Proceedings of the Conference on Computer Communications (Infocom)* (Institute of Electrical and Electronics Engineers, New York, 1998), Vol. 1, pp. 537–543.
20. S. S. Lumetta, and M. Médard, "Towards a deeper understanding of link restoration algorithms for mesh networks," in *Proceedings of the Conference on Computer Communications (Infocom)* (Institute of Electrical and Electronics Engineers, New York, 2001), Vol. 1, pp. 367–375.