

Understanding Failure Localization in Mesh Networks

Sun-il Kim and Steven S. Lumetta

CS and ECE Department, University of Illinois, Coordinated Science Lab
1308 W. Main, Urbana, IL 61801

Phone:(217)244-5564, FAX:(217)244-5685
email: sunilkim@uiuc.edu, lumetta@uiuc.edu

Abstract: This paper investigates the effectiveness of link protection schemes in terms of their ability to handle node failures and to localize multiple failures of links and nodes.

© 2002 Optical Society of America

OCIS codes: (060.4250) Fiber optics and optical communications/networks; (060.4510) Fiber optics and optical communications/optical communications

1 Introduction

Recovering from node failures is generally much more complicated than link recovery. Node failures have not been given much attention in the literature and the characteristics need to be studied carefully. In this paper, we examine and quantify the effectiveness of protection schemes in terms of their ability to handle node failures. We also investigate protection algorithms' ability to recover from multiple failures of links and nodes, extending the study of multiple link failures in [1] by introducing a more complicated failure model. We investigate issues that did not exist for failure models that deal with only link failures. The tradeoff between capacity efficiency and robustness to multiple failures is demonstrated.

Link protection preplans recovery from failures using only the portion of a network in the vicinity of a failure and can be executed in a distributed fashion as soon as failure information reaches nodes that are adjacent to the failed link or node [2]. Path protection can also be implemented in a distributed manner, but usually requires the use of a centralized manager in order to efficiently utilize capacity. We focus on link protection schemes and restrict our study to full-fiber or full-wavelength granularity with equal capacity links with no OEO regeneration (All-Optical).

We compare two link protection algorithms. The first approach, called double cycle cover, covers every link in the network with exactly two rings with opposing directions, and recovery is effected in a manner similar to the SONET BLSR [3]. The second algorithm, termed generalized loopback, applies to any two-node connected graph and requires a selection of a digraph, called the primary, and its conjugate, called the secondary [4]. Generalized loopback provides both link and node protection by flooding the secondary digraph with traffic from the node upstream from the failure. The protocol requires no differentiation between link and node failures, and allows dynamic path selection in order to route around previous failures [1]. Generalized loopback only requires a subset of the backup fibers to protect against single failures [4], and allows tuning between the capacity of a network and its robustness to multiple failures.

2 Failure Models

Link failures (fiber cuts) are the most common failure modes in optical networks. Generally, link failures represent a cut through an entire conduit and recovery algorithms must not assume that some part of the link remains intact. We ignore cases where logically separate links are placed in the same conduit; a cut through such a conduit can result in two links failing simultaneously [5].

Node failures are different from link failures and node recovery is generally much more complicated than link recovery. Many light paths pass through a network node. From the local node protection algorithm's point of view, all light paths that have the same 2-hop source and destination are grouped into what we term *streams*. All streams traversing through a node must be considered for protection in the event of failure of the node. We assume complete node failure; none of the links connected to a node remain active when the node fails. A stream that starts or terminates at a failed node cannot be restored without having node (hardware) redundancy.

We consider *multiple failures* to study failure localization. A multiple failure consists of a combination of two independent single failures. We assume that the second failure occurs long enough after the first to allow normal recovery to complete, but before any physical repair can be made. If recovery is effected through the use of only a small portion of the network (in the vicinity of the failure), multiple failures can be handled better.

3 Evaluation Measures

3.1 Protected Connectivity and Path Length Expansion

Before quantifying failure localization, it is important to understand how effectively protection algorithms address node recovery. Intuitively, we want an algorithm to be able to support connections between any pair of source and destination nodes with full robustness to single link and node failures.

We introduce two measures. *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. *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. Robust connectivity/Path length expansion. No double cycle cover was available for LATA 'X'.

graph	node degree	generalized loopback	double cycle cover (longest cycle / average cycle)
NJ 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 Avail.
ARPANET	3.2	100%/1.096	84.7%/1.326 (14/6.4)

Table 1 shows the results for robust connectivity and path length expansion ratio for five sample networks. Double cycle cover has very low robust connectivity compared to that of the generalized loopback. The average path length expansion for double cycle cover varies dramatically with the average and max length of the ring used in the cover. An interesting tradeoff between robust connectivity and path length expansion is shown. As the source and destination nodes of a robust path must lie on the same ring with double cycle covers, 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 double cycle cover depends on the length of the ring that covers both the source and destination nodes. A second tradeoff related to cycle length involves the ability to address multiple failures. After failure and recovery of a node in a ring, another node failure within the same ring cannot be recovered, and breaks the recovery path for the first failure. Creating longer rings thus leaves the network more vulnerable to multiple failures, but improves robust connectivity for single failures.

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

3.2 Vulnerability and Exposure

Generalized loopback is more attractive than double cycle cover in terms of supporting node recovery. In this section, a set of metrics that can be used to understand the tradeoff between capacity-efficiency and the ability to localize failures is discussed.

A quantitative measure of failure localization in terms of two-link failures was analyzed in [1]. This measure is extended to two-node failures and link-node failures as follows. 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 the *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. For link-node failures, *node-vulnerability* and *link-vulnerability* are defined similarly to the vulnerability discussed with two-node failures. Time ordering of multiple failures between a link and a node affects the vulnerability of a network.




Fig. 1. (a) Vulnerability for five sample networks. (b) Two-node vulnerability and exposures for NJ LATA. In both graphs, solid sections of the lines are superior to a double cycle cover; dashed lines are inferior. No double cycle cover was found for LATA 'X'

We introduce another quantitative measure of failure localization 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 is not fully restored after a second failure. *Second-failure exposure* of a node represents cases where the node fails second and cannot be restored. First-failure exposure and second-failure exposure of links are defined in the same manner.

Our failure localization metric assumes that only one stream can be protected because most link protection schemes that provide node protection can only protect one stream per node for all nodes. Only streams that are protected from single failures are considered in order to get a fair comparison between different protection schemes. Our measures indicate the degree to which a network allows failures to be localized. A high value indicates that the reliability of a typical node or a link depends on the continued operation of a large section of the network, whereas a low value indicates more local dependencies.

The measurements represent data based on the number of links removed from the secondary digraph. Measured values of vulnerability for the five sample networks are shown in Figure 1(a). Comparison with double cycle cover for the five different networks shows that generalized loopback provides the same level of failure localization while using 20% less capacity on average. With equal capacity, generalized loopback improves failure localization by an average of 22%.

In Figure 1(b), first-failure exposure and second-failure exposure cross several times. In a network with a low ratio of number of links to nodes, if a second failure breaks the first failure's recovery path, the second failure has a better chance of recovering. This is because generalized loopback reclaims links in the broken recovery path of the first failure and uses them 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 effect is the result of a decrease in number of hops for recovery paths; fewer nodes and links in a recovery path implies a smaller chance for a second failure to hit the path. Second-failure exposure also improves as more links become available for its own recovery. The same effect was seen on the COST 239 network. For the other three networks, second-failure exposure is always greater than first-failure exposure.

4 Conclusion

We introduced a model for node failures and presented measures that effectively address node protection and failure localization. Results of our investigation show advantages of generalized loopback over ring-based algorithms in robustness and spare capacity in addressing node recovery and multiple failures compared to double cycle cover.

References

1. S. S. Lumetta, and M. Médard "Towards a Deeper Understanding of Link Restoration Algorithms for Mesh Networks," in *Proc. INFOCOM '01*.
2. S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part I: Protection," in *Proc. INFOCOM '99*, vol. 2, pp. 744-51.
3. G. Ellinas and T. E. Stern, "Automatic Protection Switching for Link Failures in Optical Networks with Bi-directional Links," in *Proc. GLOBECOM '96*, vol. 1, pp. 152-6.
4. M. Médard, S. G. Finn, R. A. Barry, W. He, and S. S. Lumetta, "Generalized Loop-back Recovery in Mesh Networks," in *IEEE Transactions on Networking*, to appear.
5. Panagiotis Sebos, Jennifer Yates, Gisli Hjalmtýsson and Albert Greenberg, "Auto-discovery of Shared Risk Link Groups," in *Proc. OFC '01*, WDD3-1.