Classification of Two-Link Failures for All-Optical Networks

Steven S. Lumetta
ECE Department, U. of Illinois, Coordinated Science Lab, 1308 W. Main, Urbana, IL 61801
Phone: (217) 244-5564, FAX: (217) 244-5685, email: lumetta@uiuc.edu

Muriel Médard
EECS Department, MIT, Lab. for Information and Decision Systems, 35-212, Cambridge, MA 02139
Phone: (617) 253-3167, FAX: (617) 253-3578, email: medard@mit.edu

Abstract: This paper studies the relationship between failure localization and the properties of link restoration algorithms for mesh networks through an examination of recovery from multiple failures.

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

1. Introduction

As all-optical networking becomes feasible, a need arises to understand the relationship between failure modes and appropriate recovery algorithms. Failure of a link is a common failure mode in such networks. Links fail in many ways, ranging from amplifier failure to power outages to wayward backhoes. Generally, the failure of a link entails the failure of all fibers traversing that link, and algorithms for recovery must not assume that some part of the link remains intact. Algorithms addressing link failures must provide high-speed recovery through distributed operations on flexible topologies, and must simultaneously balance efficiency against efficacy.

In this paper, we examine the relationship between the design of restoration algorithms and their ability to gracefully recover from multiple link failures. In particular, we consider two-link failures, which consist of two independent link failures in a network graph. The second failure occurs long enough after the first to allow normal recovery to complete but before any physical repair can be accomplished. We discuss only algorithms that recover from all one-link failures, thus we assume that recovery from the first failure is always successful.

Preplanned link restoration algorithms strike a good balance between speed and efficiency. SONET bi-directional line-switched rings (BLSR's), for example, utilize this approach. Work for more general topologies has led to two algorithms. The first, a ring-based approach, uses double cycle covers [2]. In a double cycle cover, each link falls on two rings with opposing directions; each fiber is backed up BLSR-style onto one of the rings. The second algorithm, termed generalized loopback, requires selection of a digraph, called the primary, such that the conjugate digraph, called the secondary, can be used to carry backup traffic for any link failure in the primary [3]. Construction of a primary involves selection of a single direction for each link in the network. When a link fails, traffic from the primary fibers in the link is broadcast across the corresponding backup digraphs. Unnecessary backup paths are torn down after restoration completes.

2. Failure classification

We present a classification scheme for two-link failures. For each class, we provide an intuitive rationale for the problems that lead to such failures and describe the capabilities necessary for a recovery algorithm to avoid them. We report an algorithm's ability to recover from each class (and lower-level classes) in terms of vulnerability, a global notion of recovery. Vulnerability is calculated by counting the number of ordered pairs of links, i.e., two-link failures, for which the algorithm fails to completely recover and dividing the count by the number of links in the network minus one. To illustrate the relative importance of the classes for typical networks, we use a running example based on the U.S. National (USN) network shown in Figure 1. The graph has twenty-four nodes and forty-four edges.

For many classes of two-link failures, the vulnerability of a network depends on the choice of directions in the digraphs. The problem of selecting directions can be challenging, as the space of possibilities is large and sparsely populated. The USN network has $1.8\text{x}10^{13}$ possible primaries, of which roughly $1.1\text{x}10^{10}$ support recovery from any one-link failure. We apply genetic algorithms to place good upper bounds on vulnerability. In many cases, these algorithms find optimal solutions.

2.1 Fundamental failures

Two fundamental classes of failures occur as a result of network structure. The first, from which no algorithm can recover, is disconnection failures. In a two-link redundant mesh network, the failure of two links can partition the network, disconnecting some nodes from others and rendering recovery impossible. The USN network has six nodes of degree two as well as six two-edge cuts that partition more than one node from the rest of the network. In total, 24 ordered pairs of links cause disconnection failures, resulting in a vulnerability of 0.56 edges.

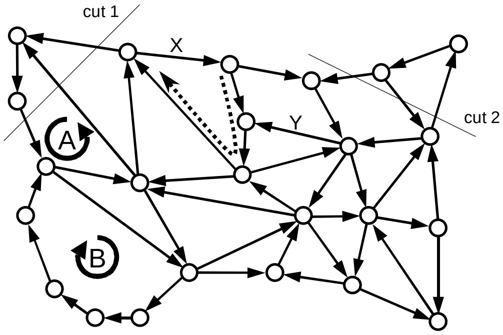

Fig. 1. USN network used as example of failure.

Link-restoration schemes switch whole fibers of traffic to effect recovery, exposing the second fundamental class, capacity failures. A capacity failure occurs when a network does not contain link-disjoint paths between the endpoints of two failed links. Capacity failures occur with any network cut of three edges: no two of the three links can recover simultaneously, as recovery of either failed link requires the use of the third, intact link. More subtle cases also arise, such as failure of links [A,B] and [C,D] in the clique on nodes A, B, C, and D. Finding two link-disjoint paths in a network is known as two-commodity integral flow. For both directed and undirected graphs, the problem is NP-complete [1], but in practice it is solved readily for small networks. The USN network has eight nodes of degree three and two cuts of three edges; the latter are marked in Figure 1. Each node contributes six ordered pairs to capacity failures. Each cut lies near nodes of degree three and contributes four ordered pairs each, summing to 56 pairs. USN's vulnerability to fundamental failures is thus 1.86 edges.

2.2 Basic algorithmic failures

The next three classes of failures correspond to three common but orthogonal aspects of recovery algorithm design. We discuss the impact of each class in terms of the USN network.

The first class arises from the logical partitioning of secondary fibers in each link into distinct backup networks, as typified by the selection of secondary digraphs for generalized loopback. Such separation reduces the number of links carrying traffic in a particular direction across a cut. As such failures arise from the selection of link directions, we term them directional failures. The number of directional failures in a network depends on the choice of directions, and we apply a genetic algorithm to place an upper bound on the number. We assume the use of backup arcs complementary to the primary arcs, which limits the impact of directional failures. For the USN network, the digraph shown in Figure 1 has no directional failures.

A second class of basic algorithmic failures stems from physical limitations of optical technology. Issues of signal regeneration and jitter restrict the length of feasible restoration paths. A long restoration path may not achieve restoration at all. We term such failures path length failures. Path length failures primarily effect algorithms that select restoration routes dynamically. The lengths of preplanned routes are fixed in advance and do not change in response to previous failures. The USN network requires paths of at least five hops to recover from the failure of any single link. For 14 ordered pairs of links, such as the long arc in cycle C in combination with any other arc in the cycle, a backup path length of six hops is required. With five-hop paths, vulnerability rises to 2.19 edges. If six or more hops are allowed, vulnerability remains at 1.86 edges.

Path hit failures, typically the most important class of basic algorithmic failures, arise when a second failure breaks the restoration path for a first. Recovering quickly from a path hit failure is difficult, as both the primary and secondary fibers in the second link must be restored. In the absence of ternary fibers reserved for restoration of secondary fibers, a concurrent distributed search to restore both fibers must solve the two-commodity integral flow problem. The use of preplanned restoration routes for the two fibers can eliminate the need for a dynamic solution to the problem, but only at the cost of encoding a significant amount of network structure into the restoration algorithm. Furthermore, achieving complete recovery is complicated by the issue of the unbroken links in the restoration path of the first fiber. If link-based restoration is used for both fibers, the presence of these links can prevent recovery unnecessarily by blocking potential restoration paths for one or both fibers. Alternatively, an algorithm can apply path-based restoration for the broken restoration path, but only at the cost of speed.

We count the number of path hit failures by selecting for each link a restoration path that minimizes the number of path hit failures that are not also fundamental failures. The minimum number of path hit failures for the USN network is 59 ordered pairs, and vulnerability at this level is 3.23 edges.

2.3 Practical algorithmic failures

The remaining three failure classes correspond to constraints in known restoration algorithms. In particular, they represent the effect of selecting a minimum-hop restoration route for the first failure rather than the best route in terms of vulnerability, the impact of using preplanned rather than dynamic routes, and the outcome of imposing topological constraints on route selection.

Blocked path failures occur when the restoration route of a first failure prevents recovery of a second failure in spite of the existence of compatible recovery routes. In Figure 1, for example, recovering link X (primary direction is reverse of backup) along the dotted route prevents subsequent recovery of link Y, and the ordered pair (X,Y) presents a blocked path failure. Blocked path failures are the last class of failures relevant to generalized loopback. Genetic algorithms provide upper bounds on the number of such failures in USN. When routes of ten or more hops are allowed, the genetic algorithm produces only 7 ordered pairs of blocked path failures (vulnerability of 3.40 edges). If paths are restricted to eight or nine hops, vulnerability rises to 3.42 edges; six or seven hops, 3.53 edges; and five hops, 3.91 edges.

Preplanning failures, the second class of practical algorithmic failures, occur in algorithms that select per-link restoration routes in advance, such as the use of preplanned shortest paths. A genetic algorithm identified a preplanned solution with a vulnerability of 4.74 edges using backup paths of at most five hops. Path lengths in general affect preplanned solutions only indirectly by restricting the choice of directions. In USN, path lengths have no effect on preplanning.

Topological constraint failures occur due to the enforcement of topological constraints on restoration routes, such as the use of rings in double cycle covers. Double cycle covers tend to have single long cycles, necessitating fairly high limits on restoration route hop count in order to make recovery feasible. We create a double cycle cover for the USN network by forming cycles in a single direction within each face and forming a final cycle in the opposite direction along the perimeter. To shorten the length of the perimeter cycle, we first fold the extrusions in the lower-left and in the upper-right of the network into the interior. The resulting double cycle cover requires twelve-hop restoration routes to allow recovery of any single link. The total vulnerability is 8.14 edges.

3. Conclusion

We described a hierarchical classification scheme for two-link failures, identified associated aspects of recovery algorithm design for each class, and applied the scheme to the USN network. For USN, vulnerability with a double cycle cover [2] is more than twice that of generalized loopback [3]. Use of preplanned shortest paths is more competitive, but is still significantly more vulnerable. We have applied the same techniques to two other networks, NJ LATA and Arpanet, and found similar results. These networks are similar in size but fairly different in structure from the USN network.

Looking to the future, improvements on generalized loopback may be able to reduce vulnerability by another factor of two (for any of the networks studied), but will require significant algorithmic advances. Addressing path hit failures seems most promising, as the majority of failures between the fundamental level and that of current algorithms fall into that class.

References

[1] S. Even, A. Itai, A. Shamir, “On the Complexity of Timetable and Multicommodity Flow Problems,” in SIAM J. Computing, 5, pp. 691-703.

[2] G. Ellinas, T. E. Stern, “Automatic Protection Switching for Link Failures in Optical Networks with Bi-Directional Links,” in Proceedings of Globecom 1996, pp. 152-6.

[3] M. Médard, S. G. Finn, R. A. Barry, “WDM Loop-back Recovery in Mesh Networks,” in Proceedings of INFOCOM, 1999.