Catastrophic Failures in a Backbone Network

Jane M. Simmons, *Fellow, IEEE*

***Abstract*—When developing a backbone network protection scheme, one important consideration is the number of concurrent link failures against which protection should be provided. A carrier typically is concerned with links going down due to independent failure events such as fiber cuts or amplifier failures. When considering just these factors, a carrier has little incentive to offer protection against more than two concurrent link failures. However, when catastrophic events are also considered, it is shown that it is worthwhile to protect mission-critical connections against up to three concurrent link failures. It is also shown that a newly developed connection setup protocol can be effective in rapidly recovering from multiple link failures.**

***Index Terms*—Catastrophic failures, correlated failures, dynamic networking, multiple failures, protection, restoration.**

I. Introduction

PROVIDING the proper amount of protection in a network is necessary to meet customer expectations for service while judiciously allocating network capacity and cost. Network customers typically specify the desired availability for a given connection when negotiating the service level agreement (SLA) with the carrier, where availability is defined as the probability that a connection is in a working state at a given instant of time. The carrier analyzes the expected failure rates and associated repair rates in the network, and determines an appropriate protection mechanism to meet the SLA.

One of the most significant failures to consider is that of a total link failure. (A link is the fiber that runs between two network nodes, where the network nodes are the sites at which traffic is sourced, terminated, and switched.) A link failure typically brings down all connections traversing that link. This disruptive effect is magnified when a network suffers multiple concurrent link failures.

Carriers must determine the likelihood of multiple concurrent link failures and design the protection scheme accordingly to meet their contractual SLAs. In calculating the availability of a connection, SLAs frequently do not include downtime caused by catastrophic events such as earthquakes, hurricanes, tornadoes, or terrorist attacks or other hostile acts. Under these stipulations, the downtime that is important is that due to equipment failure, random fiber cuts (e.g., due to a backhoe), or maintenance events. When link failures are limited to such events, then it does not make sense for a carrier to allocate resources to protect against more than two concurrent link failures. Furthermore, depending on the network topology and availability targets, providing protection against just a single link failure may be sufficient for many connections.

Manuscript received March 30, 2012. The associate editor coordinating the review of this letter and approving it for publication was X. Cao.

J. M. Simmons is with Monarch Network Architects, Holmdel, NJ, USA.

Digital Object Identifier

However, there is likely to be a small subset of traffic that can be considered mission critical, where any downtime, regardless of the source, is potentially harmful; e.g., traffic vital for national defense. For this type of traffic, one also needs to consider catastrophic events not typically covered by an SLA, even though the occurrence of such events may be rare. When the effects of such events are included in the analysis, using a simple correlated failure model, we show that providing protection against up to three concurrent link failures can be warranted. As was shown in [1], designing this level of protection for a small subset of the traffic can be accomplished with relatively little extra spare capacity.

There has been much research already on protection from multiple failures (e.g., [2] [3]), however, the bulk of the work considers only two concurrent failures and/or only random link failures. One purpose of this paper is to examine how the consideration of catastrophic failures yields different results as to the number of concurrent failures that deserve attention in designing a network protection scheme.

Section II presents the three backbones network that were studied, and the assumptions regarding failure and repair rates. Section III looks at the likelihood of multiple link failures, both with and without catastrophic events being modeled. It also illustrates the limitations of a protection scheme solely based on pre-calculated link-diverse paths when concurrent link failures occur. Section IV proposes a more dynamic protection scheme based on a newly developed rapid connection setup protocol, which improves upon the performance.

Note that multiple concurrent link failures does not imply that the failures occur simultaneously. It is more likely that one link fails, and prior to this link being repaired, another link fails. Additionally, note that providing protection against catastrophic failures is different from providing protection against failures to shared-risk link groups (SRLGs). SRLGs in backbone networks typically consist of a small number of links where the associated fiber partially resides in the same conduit, such that the links are vulnerable to a single cut. SRLGs can be taken into account when determining diverse paths for protection (e.g., [4] [5]). In contrast, during a catastrophe, *any* of the links within a geographic area may fail, which is more challenging to address.

II. Modeling Assumptions

A. Network Topologies

To examine the likelihood of multiple concurrent link failures in a backbone network, three networks were studied. The largest of the networks, shown in Fig. 1(a), is the baseline continental United States network used in the Core Optical Networks (CORONET) program [1] [6] [7]. The network is composed of 75 nodes and 99 links. The average nodal

(a) (b) (c)

Fig. 1. Network topologies studied. Node locations are similar to those of existing carriers, however, none of the networks represent an actual carrier network.

TABLE I
NETWORK CONNECTIVITY STATISTICS

	Network 1	Network 2	Network 3
Only 2 link-diverse paths	2145	1493	427
Only 3 link-diverse paths	575	271	8
Only 4 link-diverse paths	55	5	0
Only 5 link-diverse paths	0	1	0

degree of 2.6 is in line with that of most US backbone networks. This network topology was specifically designed to be capable of providing a high degree of protection. For example, four completely link-diverse cross-continental paths exist in this network, which is not a common feature in US carrier networks. The second network [8], shown in Fig. 1(b), is somewhat more representative of current carrier networks. This network has 60 nodes and 77 links, and provides three link-diverse cross-continental paths. Finally, the third network [8], shown in Fig. 1(c), is representative of a relatively small carrier, with 30 nodes and 36 links, and two link-diverse cross-continental paths. Table I shows the connectivity statistics for the source/destination pairs in each network. For example, 575 of the source/destination pairs in Network 1 have only three link-diverse paths between them.

B. Failure Rates and Repair Rates

The most common cause of link failures is a fiber cut. We assume that the fiber-cut rate is 2 cuts per 1000 miles per year [9] [10], and that the time to repair a fiber cut is uniformly distributed between 6 and 10 hours. Network links are also susceptible to optical amplifier failures. It is assumed that optical amplifiers have a FIT (failures in 10^9 hours) rate of 2000, with the repair rate uniformly distributed between 3 and 5 hours.

There are also likely to be planned maintenance events, but we assume that a carrier can exert control over when these occur. We also assume that optical switch failures (other than the nodal amplifiers), and phenomena such as PMD (polarization mode dispersion) degradation bringing down a link, are relatively infrequent events. In addition to link failures, individual connections are vulnerable to component failures, most notably transponder failures. However, carriers typically employ 1:M transponder protection, which minimizes the need for re-routing the affected connection. Transponder failures are not included in our model.

Once a link has failed, it is assumed that path-based restoration is invoked to restore any failed connection that requires protection. The connection is rerouted on a new path from source to destination, without needing to first isolate where the failure has occurred. This is important in networks with optical bypass, where fault isolation can be slower [8]. (Segment-based protection has also been proposed to protect against multiple failures, as it can typically recover from one failure per segment. However, with geographically correlated failures, it would not provide a significant benefit over path-based protection, as the failures are likely to occur in the same segment.)

In the study, we first assume that 1+N shortest link-diverse paths are pre-calculated for each source/destination node pair, where 1+N is the number of such paths that exist between the two nodes. This provides protection against any combination of N link failures. The protection mechanism can be either dedicated or shared; in either case, for simplicity, we refer to the scheme as 1+N protection.

C. Catastrophic Failures

We model catastrophic failures assuming correlated link failures, which is reasonable, though somewhat arbitrary. (Correlated link failures were also assumed in [11].) We assume that a catastrophe hits, on average, one node of Network 1 each year; the rate is correspondingly lower for Networks 2 and 3, which have fewer nodes. Each node is assumed to have an equal probability of being afflicted. With probability 5%, we assume that the catastrophe results in the whole node failing, which is modeled as all of its incident links failing. For the remaining catastrophes, we assume that each link incident to the afflicted node fails independently with probability 35%. Furthermore, any non-incident link that passes within 35 km of the afflicted node is assumed to fail with probability 10%. (Thus, we are modeling a catastrophic failure over a geographic area of radius 35 km.) These assumptions resulted in an average of approximately one failed link per catastrophe. It is further assumed that the links fail any time from the onset of the catastrophe to 30 minutes later, with uniform probability (the assumption of 30 minutes is not critical; the salient point is that the failures are not simultaneous). The time to repair a link that has failed due to a catastrophe is uniformly distributed between one and three days. If multiple links fail due to a catastrophe, they are repaired independently.

III. Results

Two sets of simulations were run for each network. In the first set, only fiber cuts and equipment failure were considered; in the second set, catastrophic failures were added to the model. Each simulation run modeled a duration of 100,000 years, which resulted in the variance of the statistics being less than 0.1% of the mean. For each network, the number of failed links in the network was tracked, along with the number of failed paths for each possible source/destination pair.

Table II shows the average amount of time per year that there are N concurrent link failures in the network, for N ranging from 1 to 3, both with and without catastrophes. When catastrophes are not considered, providing 1+3 protection provides virtually no benefit. Thus, protecting against three failures would not be a judicious use of resources if meeting an SLA were the main objective. Providing 1+2 protection may be beneficial. Almost half of the source/destination pairs with at least three diverse paths would need 1+2 protection to achieve 99.999% (i.e., "platinum-level") availability, although none would need it if the availability target were only 99.99%.

When catastrophes are added to the model, the fraction of time with three concurrent link failures increases by an order of magnitude. For those source/destination pairs with four link-diverse paths, roughly 70% of them need 1+3 protection to meet 99.999% availability. Given the rarity of catastrophes, per-annum averages may not be the most relevant statistic. For mission-critical connections, one needs to consider the ramifications *given that a catastrophe has occurred*. The duration of the three-failed-path event for those source/destination pairs with four diverse paths averages approximately 20 hours. Providing a fourth path such that the service can remain working during this length of time could be crucial.

A number of source/destination pairs cannot meet 99.999% availability even with 1+N protection. When catastrophes are not modeled, this percentage is 50%, 60%, and 75%, for Networks 1 through 3, respectively. With catastrophes, the percentages are 90%, 90%, and 99%. Most of these vulnerable source/destination pairs have just two diverse paths between the endpoints. Given this vulnerability, it is worthwhile considering alternative protection schemes, as described next.

TABLE II
Average Time Per Year with N Failed Network Links

		Network 1	Network 2	Network 3
1 Failed Link in Network	No Catastrophes	404 hours	361 hours	263 hours
	With Catastrophes	425 hours	378 hours	272 hours
2 Failed Links in Network	No Catastrophes	10 hours	8 hours	4 hours
	With Catastrophes	19 hours	15 hours	7 hours
3 Failed Links in Network	No Catastrophes	0.1 hours	0.1 hours	0.04 hours
	With Catastrophes	3 hours	2 hours	0.6 hours

IV. Protection Using Rapid Connection Setup

It is known that incorporating a dynamic aspect into the restoration process, where restoration paths are computed after failures occur, results in higher availability when dealing with multiple failures [2] [3] [12] [13]. We consider a protection scheme where two diverse paths are pre-calculated for each connection to enable immediate recovery from a first failure. If both of these pre-calculated paths fail, then another path is *dynamically* searched for *at the time of the second failure* by issuing a new connection request. If this path subsequently fails, another connection request is issued. (This is in contrast to schemes such as in [12], where the path to use for protection against a second failure is calculated after the first failure occurs. Such a scheme is viable for the more highly connected network considered in [12], where at least three link-diverse paths exist for most source/destination pairs.)

The drawback of dynamically searching for a protection path at the time of failure has previously been the relative slow speed of recovery. However, using the newly developed connection setup protocol of [1], which enables a connection to be set up in less than 100 ms, dynamism is a viable alternative that can meet stringent restoration time requirements. This protocol relies on sending probes over potential paths, where the set of paths to probe is calculated by a Path Computation Element (PCE). Under no failures, the path set is periodically recomputed to take into account current network utilization levels. In addition, in the protection scheme proposed here, whenever a link fails (or is restored), the PCE re-computes the path set in preparation for the next failure. It is assumed that there is enough time between failures to allow for this calculation. Thus, the path set that is calculated upon the i^{th} failure in the network is used as the paths to probe when (and if) the $(i+1)^{st}$ failure occurs.

The PCE cannot predict where the $(i+1)^{st}$ failure will occur; thus, probes need to be sent on a number of paths to increase the probability that a new path will be found. Under no failures, on the order of three probes may be sent, where the destination chooses the best path to use based on cost and available resources. Under failure conditions, more probes would be sent, to better ensure at least one makes it through to the destination. The destination does not need to wait for all probes to arrive before selecting the new path. Also, in contrast with many schemes that incorporate a dynamic aspect, e.g., [14], there is no need for the fault to be isolated prior to initiating recovery, resulting in more rapid recovery times.

The main advantage that the connection setup protocol of [1] has over other setup protocols, e.g., GMPLS (Generalized Multi-Protocol Label Switching), is that it more effectively deals with simultaneous connection requests, where contention for resources may arise. This is especially important when used for protection, as there are likely multiple failed connections issuing a setup request. (There are unlikely to be a large number, though, because the dynamic aspect is only invoked when a connection suffers multiple failures and that connection requires high availability.)

The simulations were rerun using this 1+1+Dynamic protection scheme rather than the pre-calculated 1+N protection scheme. A limit of five probes was arbitrarily imposed for the dynamic setup process (to limit the control-plane traffic). The paths to probe for each source/destination pair are based on eliminating all links that are down at the time of path computation plus one or more links on the current path. In the case without catastrophes, the percentage of source/destination pairs not meeting 99.999% availability is reduced to 0% in Networks 1 and 2, and 1% in Network 3. (If three probes were used instead of five, this percentage would increase to 6% for Network 3.) When catastrophes are considered, 80 to 90% of the source/destination pairs still do not meet 99.999% availability, mainly due to failures that result in there being no possible path between source and destination (i.e., no protection scheme could recover from such failures). However, dynamic protection still provides benefits. As compared to pre-calculated 1+N protection, the total downtime decreases by roughly 60% in all three networks. Thus, adding dynamism is beneficial, even for a sparse network such as Network 3.

V. Conclusion

It was shown that the number of concurrent link failures to protect against depends on whether or not catastrophes are considered. With catastrophes, three concurrent link failures warrant attention, especially for mission-critical applications. It was also shown that a new connection setup protocol makes dynamic protection a practical option for multiple failures.

References

[1] A. L. Chiu, et al., "Architectures and protocols for capacity efficient, highly dynamic and highly resilient core networks," *J. Opt. Commun. Netw.*, vol. 4, no. 1, pp. 1–14, Jan. 2012.

[2] D. A. Schupke, A. Autenrieth, and T. Fischer, "Survivability of multiple fiber duct failures," *Third International Workshop on the Design of Reliable Communication Networks (DRCN)*, Oct. 2001, Budapest.

[3] S. Kim and S. S. Lumetta, "Evaluation of protection reconfiguration for multiple failures in WDM mesh networks," *Proc. OFC 2003*, Mar. 23-28, 2003, Atlanta, GA, vol. 1, pp. 210–211.

[4] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*, Boston, MA: Kluwer Academic Publishers, 1999.

[5] D. Xu, Y. Xiong, C. Qiao, and G. Li, "Trap avoidance and protection schemes in networks with shared risk link groups," *IEEE/OSA J. Lightw. Technol.*, vol. 21, no. 11, pp. 2683–2693, Nov. 2003.

[6] A. A. M. Saleh, "Dynamic multi-terabit core optical networks: architecture, protocols, control and management (CORONET)," DARPA BAA 06-29, Proposer Information Pamphlet.

[7] Sample Optical Network Topology Files, Available at: http://www.monarchna.com/topology.html.

[8] J. M. Simmons, *Optical Network Design and Planning*, New York, NY: Springer, 2008.

[9] R. Feuerstein, "Interconnecting the Cyberinfrastructure," *CyberInfrastructure 2005*, Aug. 15-16, 2005, Lincoln, NE.

[10] B. Manseur and J. Leung, "Comparative analysis of network reliability and optical reach," *Proc. NFOEC 2003*, Sept. 7-11, 2003, Orlando, FL.

[11] H.-W. Lee, E. Modiano, and K. Lee, "Diverse routing in networks with probabilistic failures," *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1895–1907, Dec. 2010.

[12] J. Zhang, K. Zhu, and B. Mukherjee, "Backup reprovisioning to remedy the effect of multiple link failures in WDM mesh networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 57–67, Aug. 2006.

[13] Y. Li, et al., "Availability analytical model for permanent dedicated path protection in WDM networks," *IEEE Commun. Lett.*, vol. 16, no. 1, pp. 95–97, Jan. 2012.

[14] Y. Sone, W. Imajuku, and M. Jinno, "Multiple failure recovery of optical paths using GMPLS based restoration scheme escalation," *Proc. OFC/NFOEC 2007*, Mar. 25-29, 2007, Anaheim, CA.