[1], we can assume that fairly large numbers of connections can be signaled together in a message. Finally, the third assumption will not be valid if the restoration paths are selected so as to increase the resource sharing among connections [7]. In this case, connections with a common source and destination may not share the same restoration path which will presumably have a significant effect on the performance of the *aggregation over common path* scheme. The performance of other signaling aggregation schemes will also be affected since the restoration path will no longer be the "shortest" available path. Figure 3 illustrates the combined effect on the performance of signaling aggregation schemes in comparison to per-connection signaling when the above assumptions do not hold and instead 1) the message processing time increases with the number (n) of signaled connections - specifically the message processing time is ($A + B \times n$) where A is the fixed per-packet overhead and B (= $A/10$) is the time required to process each connection signaled by the message; 2) one message can signal at most 10 connections and 3) the restoration paths are calculated so as to increase restoration resource sharing using the scheme described in [7]. Note that the recovery times obtained with the signaling aggregation schemes under these conditions remain less than one-third of the recovery times obtained with per-connection signaling. Additionally, there appears to be little difference in the recovery times achieved with different signaling aggregation schemes.

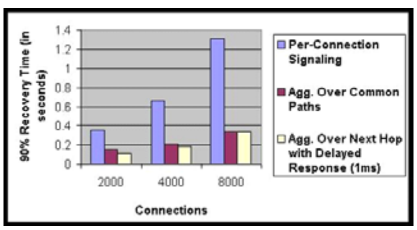

Figure 3 Performance of Signaling Aggregation with message processing times dependent on number of connections being signaled, limited numbers of connections signaled per message and non-shortest path routing.

4. Conclusions

In this paper, we proposed and evaluated several signaling aggregation schemes that reduce the number of signaling messages, thus avoiding long queuing delays during restoration signaling for large numbers of connections. By incorporating the proposed aggregation mechanisms, restoration signaling can continue to provide fast recovery from network failures even for very large number of connections in the network.

5. References

1. J.P. Lang *et al.*, "Generalized MPLS recovery functional specification," draft-bala-gmpls-recovery-functional-00.txt, Internet Draft - Work in Progress, August 2002.
2. M. Goyal, G. Li, and J. Yates, "Shared mesh restoration: a simulation study," *Optical Fiber Commun. Conf.*, March 2002.
3. The Network Simulator - ns-2, http://www.isi.edu/nsnam/ns/.
4. D. Wang, G. Li, J. Yates, P. Sebos, C. Kalmanek, and A. Greenberg, "Optical NNI interworking demonstration," *IEEE/LEOS Summer Topical Meeting,* July 2002.
5. K.C. Lee and V. Li, "A Wavelength-Convertible Optical Network," *IEEE/OSA Jour. Lightwave Tech.,* Vol. 11, No. 5/6, pp. 962-970, May/June 1993.
6. A. Shaikh, and A. Greenberg, "Experience in black-box OSPF measurement," *ACM SIGCOMM Internet Measurement Workshop,* November 2001.
7. G. Li, D. Wang, C. Kalmanek, and R. Doverspike, "Efficient distributed path selection for shared restoration connections," *IEEE Infocom 2002,* June 2002.

TuI3 **2:45 PM**

Limiting Sharing on Protection Channels in Mesh Optical Networks

R. Ramamurthy, J-F. Labourdette, A. Akyamac, S. Chaudhuri, *Tellium Inc., Oceanport, NJ, Email: jlabourdette@tellium.com*.

We examine two approaches in shared mesh restoration to limiting the number of lightpaths protected by a shared-channel. The goal is to prevent shared-channels from protecting a large number of lightpaths without significantly increasing protection capacity.

1. Introduction and Problem Definition

In shared-mesh restoration [1,2,3,4], each working path has a diverse backup path. In one restoration architecture [1,2], backup routes are pre-computed, and protection channels on the backup path are pre-assigned at the time of path provisioning. Channels on the backup path may be shared between primary paths whose working paths are diverse. Each shared-channel protects a set of lightpaths that share that channel on their backup paths. If the routing algorithm does not discriminate between shared-channels while routing a lightpath, some shared-channels may protect a large number of lightpaths (although on average the number of lightpaths protected by a shared-channel is small). Figure 1 illustrates the distribution of the number of lightpaths protected by shared-channels for a typical mesh optical network with 45 nodes and a demand of 80 lightpaths. The routing algorithm does not discriminate between shared-channels, and as a result, there is a shared-channel that protects 18 lightpaths, although, on the average, a shared-channel protects about 6 lightpaths.

# LPs using a shared channel	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
# of occurrences	10	12	6	6	18	7	1	8	4	4	2	0	1	4	2	2	0	1

Figure 1: Distribution of number of lightpaths protected by shared-channels (av. no. of LPs per shared channel is ~ 6).

If a shared-channel protecting a large number of lightpaths fail, then those lightpaths are at risk upon a subsequent failure on their primary paths. Upon a second failure on their primary path, they would have to be re-provisioned. By limiting the sharing on protection channels, we limit the number of lightpaths that would experience re-provisioning (order of seconds) as opposed to restoration [6] (order of 10s to 100s of msec) in cases of double failures. The impact on service availability is being further studied. In this work, we examine two approaches to limiting the number of lightpaths protected by a shared-channel. The goal is to eliminate the extreme cases of shared-channels protecting a large number of lightpaths, while at the same time ensuring that the protection capacity does not increase significantly.

2. Solution Alternatives

We will examine two approaches to limiting the sharing. In (1) capping, we set a hard limit on the number of lightpaths using a shared-channel. The routing algorithm considers only those shared-channels that have not exceeded the limit (which we call the Cap) in the number of lightpaths that they protect. In (2) load balancing, the routing algorithm does not discriminate between shared-channels during routing, however, during channel selection, the channel that protects the least number of links is selected. In this fashion, it is intended that, each shared-channel protects about the same number of links.

Analysis of Capping Approach

In capping, a limit is placed on the number of lightpaths using a shared-channel. The routing algorithm considers only those shared-channels that have not exceeded the limit in the number of lightpaths that they protect. We first show that a well-chosen limit can be robust to the network topology and demand pattern. Let,

- R = # protection channels / # working channels in the network
- L_{av} = average # lightpaths using a shared-channel
- L_{max} = maximum # lightpaths using a shared-channel
- h_w = average working hops of all lightpaths
- h_p = average backup hops of all lightpaths

Then, by definition [5], $R = (h_p / h_w)(1 / L_{av})$ is inversely proportional to the average number of lightpaths using a shared-channel. Then $R \geq (h_p / h_w)(1 / L_{max})$. A limit of L_{max} =1 in which each protection channel can protect at most 1 lightpath, is equivalent to dedicated mesh (1+1) protection. Since lower-bound on R is inversely proportional to L_{max}, we can expect that a with a sufficiently large choice of L_{max}, changes in the value of L_{max} will cause small changes in R. We thus expect that the specific choice of the sharing limit will not impact the ratio R of protection channels to working channels as long as it is sufficiently large.

We conducted a set of experiments to study the impact and sensitivity of the sharing limit on a variety of network parameters. The set of networks and demands in our experiments is a mix of representative real networks and demands, and randomly generated networks and demands. We considered 4 representative real networks: netA (45 nodes), netB (17 nodes), netC (50 nodes), and netD (100 nodes), and 3 randomly generated networks: net50 (50 nodes), net100 (100 nodes), and net200 (200 nodes). Figure 2 illustrates the ratio R of protection channels to working channels for the different networks for different values of the sharing limit. We observe that as the sharing limit L_{max} increases, R decreases sharply at first, and then decreases gradually, and then remains flat. Most of the sharing gains are obtained when the sharing limit is below about 5 lightpaths. In all the

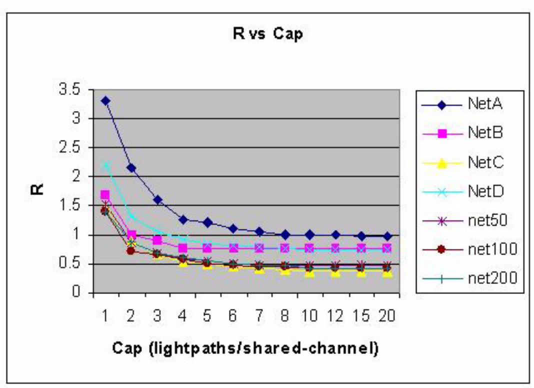

Figure 2: Ratio of protection capacity to working capacity (R) against the value of sharing limit.

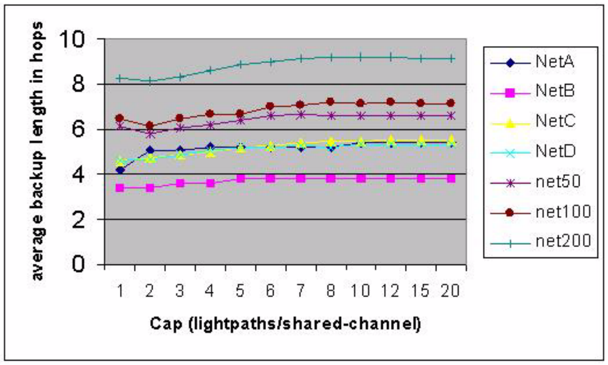

Figure 3: Average backup path hops versus sharing limit.

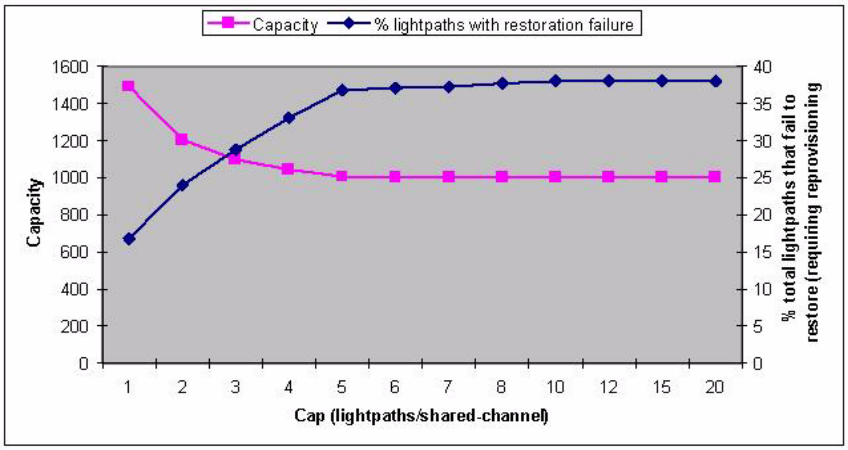

Figure 4: Impact of the sharing limit on restoration upon double failures.

networks, beyond a sharing limit of 10 lightpaths, there are no incremental gains in protection capacity. The figure also illustrates that R is inversely proportional to the sharing limit. We observe that for larger networks, sharing saturates at larger values of the sharing limit. Based on these samples, a sharing limit of around 6 lightpaths is robust across a variety of networks.

Figure 3 illustrates the average number of hops on the backup path as the sharing limit varies. Backup path hops directly influences the restoration time upon a failure, with larger number of backup hops generally consuming more restoration time. We observe that as the sharing limit increases, the average number of backup hops, and thus restoration times, increase marginally. This is because, as the sharing limit increases, there are more opportunities for sharing, and the backup path may traverse a longer distance trying to use links with sharable channels. The routing algorithm can assign a cost-model for sharable channels to achieve a trade-off between the amount of sharing and the length of the backup path [7].

Figure 4 plots the impact of the sharing limit on the percentage of lightpaths that are restored upon double failures for a given network. The values are averaged over all double failures. Also plotted is the capacity requirement for each value of the sharing limit. We observe that as the sharing limit increases, the required capacity decreases (due to a decrease in protection capacity because of better sharing), however, due to better sharing, there are more contentions for capacity upon double failures, and as a result, the percentage of lightpaths whose restoration fails under double failures increases. The figure indicates that to achieve better resiliency against double failures, the sharing limit needs to be decreased below 5, consequently the capacity requirement increases.

Analysis of Load-balancing Approach

In load balancing, the routing algorithm does not discriminate between shared-channels during routing, however, during channel selection, the channel that protects the least number of links is selected. In this fashion, it is intended that each shared-channel protects about the same number of links. Table 1 illustrates the ratio between the capacity required with load balancing and the capacity required without load balancing. This ratio is illustrated with and without imposing a sharing cap of 10. We observe that there is negligible and inconsistent difference between the capacity requirements with and without load balancing, i.e., load balancing on shared-channels does not affect the capacity requirement. This is also true when there is a sharing limit

.

	Without Cap	**With Cap (10)**
NetA	1.002	1.002
NetB	1	0.992044064
NetC	1.004280831	0.999302575
NetD	1.003536902	1.003298775
Net50	1.003287929	1.000468384
Net100	0.996265441	0.99453709
Net200	0.997869426	0.989348935

Table 1: Ratio of capacity requirement with load balancing and without load balancing.

	Without Cap	**Without Cap**	**With Cap (10)**	**With Cap (10)**
	no load-balancing	**with load-balancing**	**no load-balancing**	**with load-balancing**
NetA	9	6	8	8
NetB	13	11	8	8
NetC	17	12	8	8
NetD	13	10	8	8
Net50	13	9	10	9
Net100	14	12	10	10
Net200	17	17	10	10

Table 2: Maximum number of lightpaths shared by some shared-channel with and without load balancing.

Table 2 illustrates the maximum number of the lightpaths that are shared by some shared-channel with and without load balancing. We observe that load-balancing does have a limiting effect on sharing, i.e., load-balancing does reduce the maximum number of lightpaths that are shared by some shared-channel. However, we observe that reduction in maximum lightpaths shared by some shared-channel is not deterministic, i.e., it is not possible to guarantee that a certain limit will not be exceeded. When there is a sharing limit, we observe that load balancing does not have any effect on the maximum number of lightpaths that are shared by some shared-channel.

3. Conclusion

Based on the experiments we conclude that capping is preferable to load balancing and achieves the following:

- Imposing a sharing limit on shared-channels eliminates the cases of shared-channels protecting a large number of lightpaths. When the sharing limit is chosen well, there is no capacity penalty.
- A sharing limit of around 6 lightpaths is robust across a variety of topologies and demand sets.
- Load-balancing on shared-channels does not provide quantifiable benefits in limiting sharing especially if a sharing limit is already imposed.

References

[1] J-F. Labourdette, et. al., *"Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks"*, Photonic Network Communications, 4:3/4, pp. 219-235, July-December 2002.

[2] G. Ellinas, et. al., "*Routing and Restoration Architectures in Mesh Optical Networks*", Optical Networks Magazine, Issue 4:1, January/February 2003.

[3] B. Doshi, et al, "Optical Network Design and Restoration", Bell-Labs Technical Journal, Vol. 4, No. 1, pp. 58-84, January-March 1999.

[4] J. Doucette, W. Grover, T. Bach, "Bi-Criteria Studies of Mesh Network Restoration: Path Length vs. Capacity Tradeoffs", in Proc. of OFC 2001, Anaheim, CA, March 2001.

[5] J-F. Labourdette, et. al., "*Approximate Dimensioning and Performance Analysis of Mesh Optical Networks*", submitted for publication.

[6] A. A. Akyamac, et. al., "*Optical Mesh Network Modeling: Simulation and Analysis of Restoration Performance*", NFOEC 2002, Sept. 2002, Dallas, TX.

[7] E. Bouillet, et. al., *"Enhanced Algorithm Cost Model to Control tradeoffs in Provisioning Shared Mesh Restored Lightpaths*", OFC 02, April 2002, Anaheim, CA.

TuI4 **3:00 PM**

Computing Optically Disjoint Paths for Survivable All-Optical Networks

N. Garg, R. Simha, *The George Washington University, Washington, DC, Email: nikhil@gwu.edu.*

We consider the problem of finding two paths for given wavelength channels in the network such that one of them remains available in the event of an optical attack by addition of minimal opto-electronic equipment.

Introduction

Networks can be designed to recover from various failure and attack scenarios by assigning additional resources in anticipation of failures or attacks. Such pre-planned allocation of resources for protection has received much attention in the routing and wavelength assignment literature for optical networks. In general, past work has focused on finding link-disjoint paths for primary and secondary paths of a connection. However, this body of work has focused on optical networks with opto-electronic conversion at the switches. In an all-optical network, however, it is not sufficient for the paths to be merely link-disjoint: an attack or malfunctioning light source on a primary path may in fact corrupt the secondary path via