

Impact of Resource Sharability on Dual Failure Restorability in Optical Mesh Networks

Chadi Assi¹, Wei Huo¹, and Abdallah Shami²

¹ Concordia Institute for Information Systems Engineering, Concordia University,
Montreal, Quebec, CANADA
assi@ciise.concordia.ca

² Department of Electrical and Computer Engineering, The University of Western
Ontario, London, Ontario, CANADA
ashami@eng.uwo.ca

Abstract. Protection capacity re-provisioning mitigates the impact of double-link failures by provisioning new capacity for unprotected demands in a network designed to achieve 100% restorability under single-link failures. We study the performance of re-provisioning in networks under various resource sharing degrees. Intuitively, the lower is the sharability degree of resources the smaller is the number of unprotected connections resulting after the recovery from a failure. However, we show in this paper that limited resource sharability also implies limited flexibility for the network in finding capacity for unprotected demands after failures; which accordingly limits the capability of re-provisioning schemes in improving the network restorability. We further study the performance of different re-provisioning algorithms under distributed control and we show that contentions severely impact the restorability performance. Finally, we propose a simple mechanism to mitigate the effects of contentions.

1 Introduction

Significant progress has been made towards making optical networks resilient in the event of single link failures. Protection schemes with preplanned spare capacity [1] have been extensively studied in optical mesh networks, where protection capacity can either be dedicated or shared among multiple connections whose primary paths are physically disjoint [2, 3, 4]. Now as the size and the complexity of optical mesh networks continue to grow, dual failures become increasingly probable [6, 7]. Hence designing recovery algorithms to protect against such failure events and ensure service continuity is a paramount concern.

To date, various research efforts have already addressed the problem of routing connections under dual failure assumptions, and findings show that designs offering complete dual-failure restorability require more than double the amount of spare capacity [5, 14]. In order to avoid this excessive deployment of extra spare capacity in the network, *capacity reconfiguration* after the occurrence of and recovery from the first failure has been proposed [6-10]. After the occurrence

II

of the first failure, the failed connections are restored from their working paths into their protection paths. Hence upon complete recovery, backup capacity along now active protection routes can no longer be shared. As a result some of the connections in the network will become indirectly unprotected, therefore increasing the network vulnerability to a subsequent failure. Capacity re-provisioning provides a mechanism by which one can find and allocate new protection capacities for these newly unprotected connections without a priori knowledge of the location of the second failure. Many backup re-provisioning algorithms for handling multiple failures have recently been proposed [6, 7, 8, 10, 14] and comprehensive studies indicate that re-provisioning can dramatically lower network vulnerability.

In this paper we discuss a new re-provisioning scheme to improve the connection restorability in optical shared mesh networks. We assume two near-simultaneous failures, where the second failure occurs after the first failure is recovered from, but before it is physically repaired. A critical objective for re-provisioning is to reduce the total number of connections that have to be re-provisioned. Here the motivations are twofold: (1) to reduce management overheads in simultaneously provisioning a large number of connections, and (2) to lower reservation contention between multiple unprotected connections trying to establish backup capacity. The latter may result in increased blocking rates for re-provisioning, which in turn will increase vulnerability to subsequent failure(s).

To better utilize the network resources, connections may be allowed to share their protection wavelengths if their corresponding working paths are link disjoint. These protection resources are activated upon the occurrence of a failure to restore the affected connections and therefore cannot be shared anymore, leaving some connections unprotected in the network. The number of resulting unprotected connections can be very large if the degree of sharability (also referred to as sharability index) of protection wavelengths is high. Accordingly, limiting the resource sharability in the network will result in reducing the number of unprotected demands. However, since re-provisioning makes use of available resources in the network to provision new protection capacity, limited sharability will yield lower flexibility in finding and assigning resources. Therefore, it is clear that there are two conflicting design constraints: on one hand limited SI may reduce the number of unprotected connections but at the expense of less flexibility in allocating protection capacity for unprotected connections; on the other hand, higher SI may result in larger number of unprotected connections after the first failure with higher degree of flexibility in provisioning protection capacity. Hence, one objective of the paper is to provide a comprehensive study on the performance of capacity re-provisioning under different sharability degrees.

Another factor that may limit the performance of re-provisioning is the underlying implementation of the protocol, namely centralized and distributed control. In this paper, we compare the performances of re-provisioning under the two implementations. Furthermore, we propose a simple and efficient technique to cope with the adverse effects of contentions incurred in distributed re-provisioning. The rest of the paper is organized as follows. In section 2 we

study the impact of resource sharability on network restorability. In section 3 we introduce a new re-provisioning algorithm and in section 4 we discuss the performance under centralized and distributed control. Section 5 presents performance evaluation and comparisons and finally we conclude in section 6.

2 Impact of Sharability on Network Restorability

Under normal conditions, the network is usually protected against all single failures. Namely, when a failure occurs, all connections whose working paths are affected by that failure are re-routed on their corresponding protection paths [2, 3, 4]. Hence, a connection recovery usually requires source node notification and recovery signaling to configure the protection resources (e.g., wavelengths and cross connect switches) along the backup route [4]. However, since these protection resources may also be shared with other unaffected connections, these may become unprotected and vulnerable to the next failure [10]. Overall, to summarize these unprotected connections, one can classify them into three categories:

1) *Indirectly Affected Connections*: Upon failure, shared protection resources are activated by the failed connections which may cause some connections (whose backup lightpaths share these protection resources) to become unprotected.

2) *Directly Affected Working Connections*: A failed demand that is re-routed to its backup path is still vulnerable to a second failure that may affect its active protection path.

3) *Directly Affected Backup Connections*: Demands whose protection connections have failed due to the first failure.

Clearly, increased numbers of unprotected connections in the network can increase its vulnerability to subsequent failures and hence lower the overall network restorability. To improve the service availability, re-provisioning exploits the available capacity in the network to re-establish new backup paths for unprotected connections in advance of a failure (and right after the recovery from the first fault). One such scheme has been proposed in [8, 10] and its performance is evaluated here as well, termed thereafter as Scheme I. Namely, after recovery from the first failure, the algorithm categorizes all unprotected connections into one of the three categories detailed above. Subsequently, it attempts to establish new protection lightpaths for these demands.

Note that the degree of wavelength sharability (i.e., the number of connections allowed to share the same protection resources) directly impacts the number of unprotected connections that needs re-provisioning. Clearly, if this sharability degree is large, the network will potentially admit more connections since more backup paths are packed together; and a large number of connections will become vulnerable after the first failure (i.e., the impact of connections in category 2 to those in category 1). Conversely, limiting the sharability of a protection wavelength will reduce the total number of unprotected connections after the first failure; however that comes at the expense of reduced network performance that limited sharability yields.

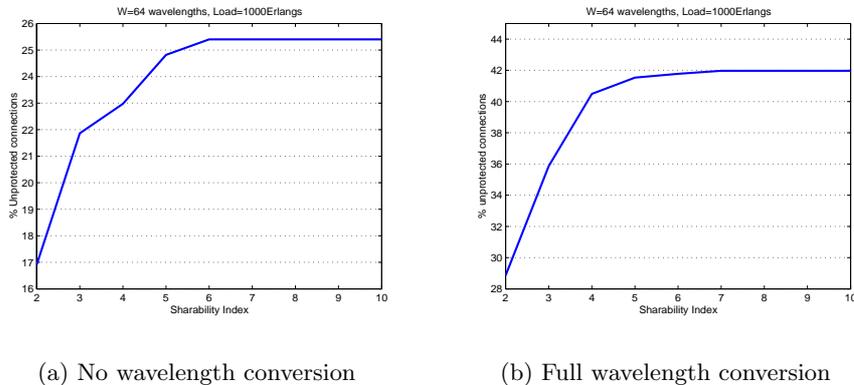


Fig. 1. Percentage of Unprotected Connections before Re-provisioning vs. SI

Fig.1 shows a case on a nation wide network (Fig. 3) where the percentage of unprotected connections is measured after the recovery from the first failure by varying the sharability index, SI. Clearly, the figure shows a substantial increase in the number of unprotected connections (e.g., 9% increase in wavelength continuous network and 13% in a wavelength convertible network [15]) in the network after the first failure if unlimited resource sharability is allowed. Therefore, intuitively this finding suggests that limiting the sharability index will result in a lower number of unprotected demands after recovery and therefore in a less number of connections to be re-provisioned. However, on the other hand, a smaller number of unprotected demands to be re-provisioned does not always guarantee that the percentage (or the number) of unprotected connections in the network after re-provisioning is reduced. The reason for that is that limited resource sharability will effectively limit the performance of the re-provisioning algorithm in finding protection resources for unprotected demands. Moreover, limited sharability will limit the overall performance of the network since fewer connections can be accommodated by the network. Alternatively, unlimited resource sharability may provide the necessary flexibility in allocating protection capacity to unprotected demands after the first failure but, as the figure shows a larger group of demands become unprotected and require re-provisioning. Moreover, if distributed re-provisioning is implemented, larger number of connections attempting to reserve protection capacity will amplify the contentions over resources, yielding higher blocking and leaving more connections unprotected even when resources are available. Therefore, it is clear that there are two conflicting design constraints: on one hand limited SI may reduce the number of unprotected connections but at the expense of limited network performance and less flexibility in allocating protection capacity for unprotected connections. On the other hand, higher SI may yield a larger number of unprotected connections

after the first failure and magnifies the effect of contentions under distributed control but yields higher degree of flexibility in provisioning protection capacity.

3 Network Re-provisioning

3.1 An Alternative Approach

A critical objective for re-provisioning is the reduction of the number of unprotected connections after the first failure. Achieving this objective strongly depends on the resource availability in the network after the failure and the limit imposed on their sharability. Here, when a failed connection is recovered onto its backup lightpath, a large number of unaffected demands, whose backup connections share the same protection capacity on any link along the backup of the failed connection, will be unprotected and hence require re-provisioning. Typically, the larger the sharability of a wavelength, the larger will be the number of indirectly affected connections in the network that become unprotected. Therefore, if only a new working lightpath is setup for each failed connection, then traffic is switched back from the backup lightpath to the new working path and the shared protection capacity becomes available leaving the indirectly affected connections protected; we refer to this methodology thereafter as Scheme II [16].

The effectiveness of Scheme II is best shown via an illustrative example in Fig.2. We assume initially b_1 , b_2 and b_3 are all setup using λ_1 , and b_1 shares λ_1 on link (D-E) with b_2 and on link (E-H) with b_3 . When link (B-F) fails, w_1 is restored to its backup b_1 and as a result, b_2 and b_3 become unavailable since they share protection capacity with b_1 . Hence b_1 , w_2 and w_3 become all unprotected and three new protection paths (or capacity) need to be re-provisioned in order to fully protect the network against a subsequent failure. Under Scheme II however, when w_1 is restored to its backup, connection b_1 , w_2 and w_3 become *temporarily unprotected*. Hence, if we can find a new working path (w_1^{new}) that is link disjoint with b_1 to carry the failed traffic, then b_2 and b_3 can also become available again and their corresponding connections (w_2 , w_3) are fully protected. Note that w_1^{new} may not be disjoint with w_2 and/or w_3 (w_2 in this example). Therefore, b_1 cannot share any protection resource with b_2 . In a wavelength continuous network, a new backup b_1^{new} (and protection wavelength) that is link-disjoint with w_1^{new} has to be provisioned. In a wavelength convertible network, the conflict links are identified (e.g., (D-E)) and a different wavelength is provisioned along those links (e.g., λ_2 can be assigned to b_1 on link (D-E) leaving the rest of the backup lightpath intact). Note that Scheme II differs from Scheme I in that the number of connections to be re-provisioned upon a failure is dramatically reduced, whereas the number of *temporarily unprotected* connections during the re-provisioning time remains the same. Finally, Scheme II has a strong impact on the overall network restorability. Even under higher SI, the number of connections to be re-provisioned can be much reduced. The number of connections to be re-provisioned under Scheme II is substantially lowered re-

sulting in a less management overhead and lighter impacts of contentions under distributed control.

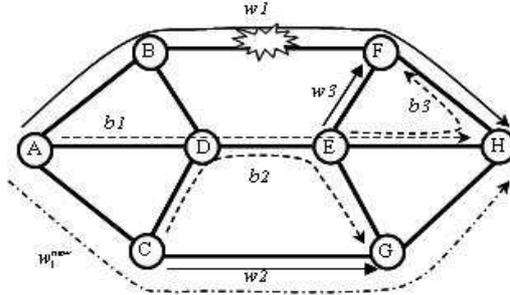


Fig. 2. Example for Re-provisioning

Now clearly, under the premise that only new working path for failed connections are provisioned, resource sharability may have minimal impact on the performance of the re-provisioning algorithm since working capacity can not be shared. Notice, however, that Scheme II does not neglect completely the effect of resource sharability; when a link fails, protection wavelengths on this link will also fail and these wavelengths may be protecting a larger number of connections when the SI is higher. Moreover, if the algorithm does not succeed in provisioning (some) new working connections for failed demands, then (as in Scheme I) the new scheme will identify the unprotected connections (those in categories 1 and 2) and re-provision them. Accordingly, the SI will have an impact on increasing the number of connections to be re-provisioned.

3.2 Implementation Perspective

We have seen that the performance of re-provisioning strongly depends on the algorithm itself and the sharability of protection resources. Another factor that also affects the performance is the implementation of the underlying algorithm; an algorithm typically can either have a centralized or a distributed control [11].

Under a centralized implementation, a central network management system holds the global information of network resources, such as network topology, link states, wavelength usage on each link, sharability information for protection resources, etc., and the corresponding steps of the particular algorithm are executed at this central controller. Here, upon the occurrence of a failure, the network will take the responsibility of recovering the failed connections through a standard signaling recovery protocol [4] and the central controller is informed through an alarm message to initiate the re-provisioning procedure of each unprotected demand. Clearly, centralized re-provisioning of unprotected connections is done sequentially in order to avoid contention for capacity. Contentions

for capacity may lead to increasing the number of unprotected connections in the network and therefore increasing the vulnerability to a subsequent failure.

Alternatively, under distributed implementation, the source node of each unprotected demand is responsible for re-provisioning new protection capacity for its connection. An unprotected demand is typically identified by either the node detecting the link failure or by the source node of a failed connection. We deploy here a distributed provisioning approach with forward reservation [12], whereby the source node of one unprotected demand computes a new path and/or a new protection wavelength. Subsequently the node sends a control message containing the new selected wavelength to reserve resources along the entire path. If at least one node along the route is not successful in reserving the selected wavelength, the reservation fails and the connection is deemed unprotected. Here, unlike the centralized scheme where all connections are re-provisioned sequentially, all unprotected connections attempt to reserve protection capacity simultaneously and therefore contentions [11, 12, 13] may likely occur among connections requesting the reservation of the same resource. Clearly, a connection failing to successfully find new protection capacity will be left unprotected and will ultimately increase the network vulnerability to a subsequent failure.

Note that, if the number of unprotected connections resulting from the first failure and simultaneously attempting to re-provision new protection capacity is quite large, contentions over resources is more likely to increase; thereby, leaving a large number of unprotected demands in the network upon re-provisioning. Hence, a larger sharability index will yield a larger group of unprotected connections and accordingly the intensity of contentions under distributed re-provisioning will be magnified. Therefore, to achieve better network restorability, the effect of contentions will have to be reduced. Here, one advantage Scheme II possesses over Scheme I is that the number of connections to be re-provisioned is potentially much smaller; therefore making the impact of contentions on network restorability less severe. Nonetheless, it is still a concern as it will be prevalent in the next section. To mitigate the impacts contentions may have on the network restorability, we propose that unprotected connections attempting to re-provision and failing to succeed due to contention, be allowed to reattempt after selecting a different wavelength if possible. The advantage of reattempting is that blocking due to contentions may be reduced whereas the drawbacks are increased network re-provisioning times. Later we will see that re-provisioning retries strongly reduce the impact of contentions, and accordingly improves the dual failure network restorability.

4 Simulation Results

In this section, we present some numerical examples to illustrate the performance of capacity re-provisioning on improving the double failure restorability of optical networks and we study the impact that resource sharability may have on the performance. We implement a discrete event simulation tool using C++ to simulate the re-provisioning algorithm. The sample network topology we

VIII

used is shown in Fig.3. and it consists of 24 nodes and 86 unidirectional links. Requests are uniformly distributed between all source-destination pairs and arrive at each node via a Poisson process with a mean arrival rate of λ arrivals/ms. The connection-holding time is exponential distributed and the number of wavelengths (W) per link is 64.

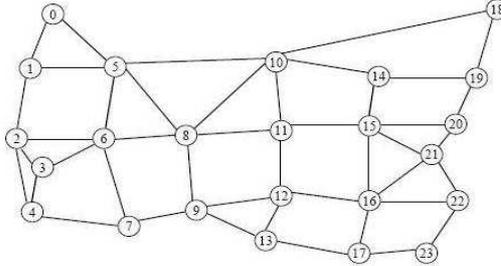


Fig. 3. A Sample Network Topology

We simulate the failure of a unidirectional link and we measure the percentage of the unprotected connections in the network after and before re-provisioning. Fig.4. shows the performance of the two re-provisioning algorithms, namely Scheme I and Scheme II, in a wavelength continuous (Fig.4a) and convertible (Fig.4b) network. As the sharability of protection resources increases, the figure shows that the percentage of unprotected connections in the network before re-provisioning increases clearly because more connections are admitted to the network. The SI here varies between 2 and 10 where a SI = 1 corresponds to the dedicated protection case. Clearly, capacity re-provisioning (under both schemes) improves the network performance by substantially reducing the percentage of unprotected connections (e.g., a decrease from 42% to 6% at higher SI in a wavelength convertible network using Scheme I) and therefore making the network less vulnerable to subsequent failures.

One interesting finding, shown in Fig.4a, is that of the impact of SI on the re-provisioning gain. Namely, the lower percentage of unprotected connections before re-provisioning at lower SI does not necessarily mean a good re-provisioning performance (i.e., a lower percentage of unprotected connections) after re-provisioning. The figure shows that as the SI increases, the percentage of unprotected connections after re-provisioning decreases for Scheme I (10%-6%) while it remains almost constant for Scheme II ($\sim 3\%$) with better performance than that of Scheme I. The reason for this is that a lower SI will limit the flexibility of the re-provisioning algorithm in finding and judiciously allocating protection resources. While on the other hand, a higher SI will allow the network to accommodate more unprotected demands during re-provisioning. Therefore, the figure shows a larger performance gain at higher SI (38% (42%-6%)) than

at lower SI (20% (30%-10%)). Alternatively, Scheme II shows a fixed percentage of unprotected connections at different sharability indices. The reason for that is due to the fact the Scheme II gives preference to provisioning new working capacity for failed demands in order to avoid re-provisioning a larger number of protection connections. Similar results are shown in Fig.4b, except that the percentage of unprotected connections is smaller since in a wavelength continuous network, fewer connections are admitted to the network due to the wavelength continuity constraint.

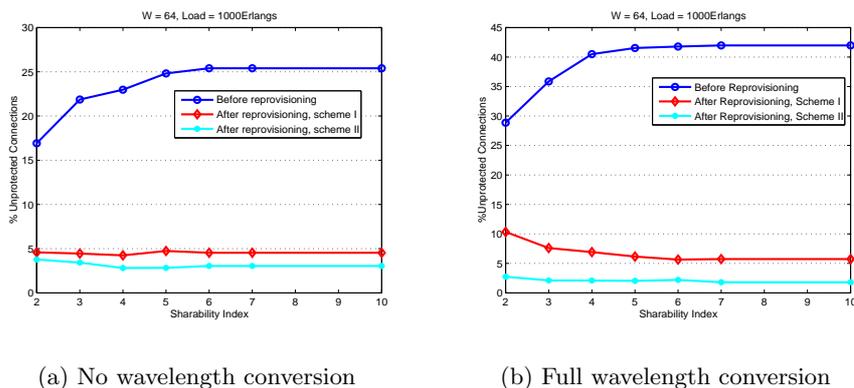


Fig. 4. Percentage of Unprotected Connections - before and after Re-provisioning

Next, we present the performance of capacity re-provisioning under distributed control. As we have already mentioned, when a large number of unprotected connections attempt to re-provision simultaneously, contentions among resources will occur in the network. This will prevent some demands from protecting their connections although resources may be available. Fig.5a shows a comparison of re-provisioning for centralized and distributed implementation of Scheme I in a wavelength continuous network under different sharability degrees. Clearly, the percentage of unprotected connections after distributed re-provisioning is much higher than that under centralized re-provisioning (10%~14% and 5%) which is mainly due to the effects of contentions. Moreover, at higher SI, by comparison to the performance under lower SI, a larger number of connections starts simultaneously the re-provisioning process and contends for the resources; hence stronger contentions may occur and as the figure shows, the percentage of unprotected connections increases. We should note that under distributed re-provisioning, a demand may fail to protect its connection for two reasons: (1) due to unavailable resources or (2) due to contentions with other connections. We measured the impact of contentions on increasing the number of unprotected connections and the results are shown in Fig.5a. Clearly, most of the connections fail to be protected due to contentions while attempting to

reserve capacity. Here, unlike centralized scheme, in distributed re-provisioning the blocking due to insufficient resources decreases while the blocking due to contentions increases substantially due to the latency in receiving resource updates in time. Therefore, a node may attempt to reserve capacity that may already have been reserved by some other connection, leading to blocking due to contentions.

To minimize the impact of contentions, we propose that a connection that is being blocked due to only contention be allowed to select a new wavelength and retry its reservation. Here, when a demand fails to allocate protection capacity (wavelength) to its connection due to contentions, it will be allowed to reselect a different wavelength and re-attempt re-provisioning. Fig.5b shows the improvement of this *re-select and reserve* scheme in reducing the percentage of unprotected connections after re-provisioning. The disadvantage of retrying, however, is the increase in the overall re-provisioning time. Our simulations showed that the total network re-provisioning time is kept below 1 second when the total number of retries is 3.

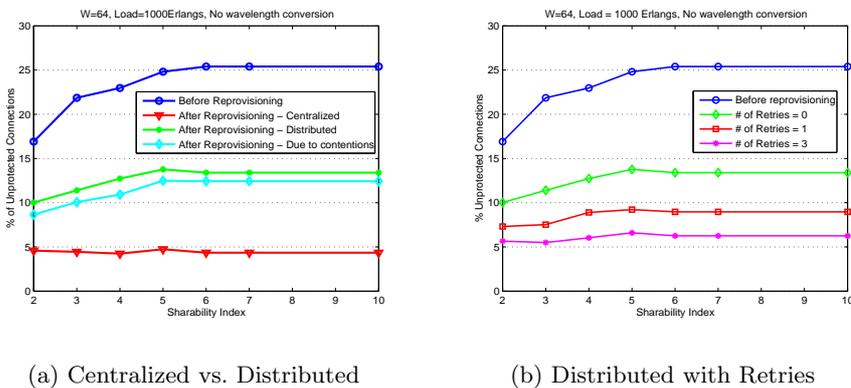
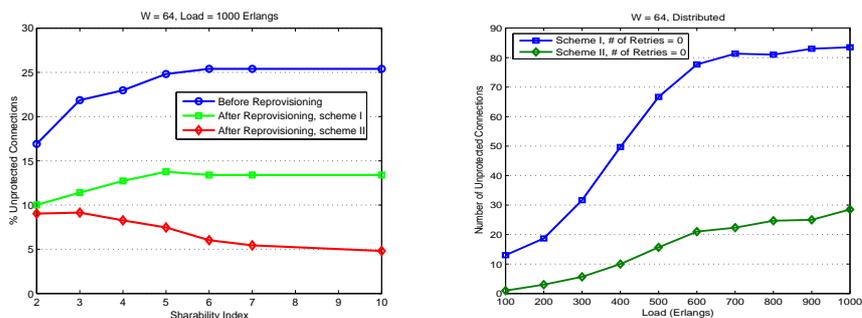


Fig. 5. Effects of Contentions on the Performance of Re-provisioning, Scheme I

Finally, a comparison between Scheme I and Scheme II under distributed re-provisioning is presented in Fig.6a. Under Scheme I, the algorithm identifies all resulting unprotected connections after the recovery of failed connections (i.e., all connections in categories 1, 2 and 3) and those are all re-provisioned simultaneously. However, under Scheme II, the algorithm starts by re-provisioning only a fraction of those connections and namely, new working connections for the failed demands (i.e., those in categories 2 and 3); when unsuccessful, some of the connections in category 1 are identified and re-provisioned accordingly. Moreover, that fraction of demands is not re-provisioned simultaneously, therefore leading to lesser impact of contentions. Now when the SI is low, the performance of both schemes is similar (because fewer connections are packed together, hence

the number of unprotected connections in category 1 is small) with Scheme II slightly outperforming Scheme I since the effect of contentions is not as strong. When the SI increases, the number of unprotected connections increases and accordingly contentions start to have larger impact on Scheme I; hence the increase in the percentage of unprotected connections after re-provisioning. However, under Scheme II the behavior is quite different. As SI increases, the percentage of unprotected connections after re-provisioning decreases (note that this does not necessarily mean that the total number of unprotected connections is reduced, since more unprotected connections may exist before re-provisioning) since the impact of contentions is not as severe and the algorithm judiciously uses the available capacity to protect more connections. As a conclusion, the effects of contentions on Scheme II is much smaller than that on Scheme I. Fig.6b shows the number of unprotected connections after re-provisioning that fail to reserve protection capacity due to only contentions when the sharability of resources is very large. Clearly, the figure shows the strong impact contentions have on degrading the performance of Scheme I.



(a) Percentage of Unprotected Connections

(b) Number of Unprotected Connections due to Contention

Fig. 6. Performance of Distributed Re-provisioning, Scheme I vs, Scheme II

5 Conclusion

In this paper we studied the problem of improving the network restorability through protection capacity re-provisioning for unprotected connections after a failure. We showed that resource sharability plays a significant role in the performance of network re-provisioning. Namely, we showed that although limited resource sharability may yield to a small number of unprotected connections after a failure, it affects the performance of re-provisioning due to the limited

flexibility in allocating protection capacity among unprotected demands. We studied the performance of re-provisioning under distributed control and we showed the strong impacts contentions may have on degrading the performance under different sharability conditions. We proposed a simple mechanism whereby contentions effects can be mitigated to improve the overall network restorability. We also presented a comparison of different re-provisioning schemes and we showed how the algorithm itself may strongly impact the overall performance.

References

1. W. Grover: "Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking", *prentice hall*, 2003.
2. S. Ramamurthy, B. Mukherjee: "Survivable WDM Mesh Networks, Part II - Restoration", *IEEE ICC 1999*, 1999.
3. J. Labourdette: "Shared Mesh Restoration in Optical Networks", *Proc. OFC 2004*, Feb. 2004.
4. J. Yates, G. Li: "Challenges in Intelligent Transport Network Restoration", *Proc. OFC 2003*, March 2003.
5. M. Clouqueur, W. D. Grover: "Mesh-restorable networks with complete dual failure restorability and with selectively enhanced dual-failure restorability properties", *SPIE OPTICOMM*, Boston, MA, July-Aug 2002.
6. D. Schupke, R. Prinz: "Performance of Path Protection and Rerouting for WDM Networks Subject to Dual Failures", *Proc. OFC 2003*, March 2003.
7. S. Kim, S. Lumetta: "Evaluation of Protection Reconfiguration for Multiple Failures in WDM Mesh Networks", *Proc. OFC 2003*, March 2003.
8. R. Ramamurthy, A. Akyamac, J-F. Labourdette, S. Chaudhuri: "Pre-Emptive Re-provisioning in Mesh Optical Networks", *Proc. OFC 2003*, March 2003.
9. P. Charalambous, et.al. "A National Mesh Network Using Optical Cross-Connect Switches", *Proc. OFC 2003*, March 2003.
10. J. Zhang, K. Zhu, B. Mukherjee: "A Comprehensive Study on Backup Re-provisioning to Remedy the Effect of Multiple-Link Failures in WDM Mesh Networks", *ICC04*, Paris, June 2004.
11. L. Shen and B. Ramamurthy: "Centralized vs. Distributed Connection Management Schemes under Different Traffic Patterns in Wavelength-Convertible Optical Networks", *Proc. IEEE ICC02*, NY, 2002.
12. Y. Mei, and C. Qiao: "Efficient Distributed Control Protocols for WDM Optical Networks", *Proc. ICCCN*, September 1997.
13. H. Zang, J. Jue, and B. Mukherjee: "A Review of Routing and Wavelength Assignment Approaches for Wavelength Routed Optical WDM Networks", *Optical Networks*, No.1, January 2000.
14. M. Clouqueur, W.D. Grover: "Availability Analysis of Span-restorable Mesh Networks", *IEEE JSAC*, vol.20, no. 4, May 2002, pp. 810-821.
15. B. Ramamurthy and B. Mukherjee: "Wavelength conversion in WDM networking", *IEEE J. Select. Area. Commun.*, vol. 16, no. 7, pp. 1061-1073, September 1998.
16. C. Assi, W. Huo, A. Shami, and N. Ghani: "Analysis of Capacity Re-Provisioning in Optical Mesh Networks", to appear in *IEEE Communications Letters*, (accepted in January 2005), 2005.