Analysis of Capacity Re-provisioning in Optical Mesh Networks

Chadi Assi, *Member, IEEE*, Wei Huo, Abdallah Shami, *Member, IEEE*, and Nasir Ghani, *Senior Member, IEEE*

***Abstract*— We consider optical mesh networks and we study the problem of improving service availability under dual near-simultaneous failures. We propose a re-provisioning algorithm that improves the robustness of optical mesh networks and we compare its performance with another scheme under various degrees of resource sharing. The effectiveness of our proposed approach in improving restorability is demonstrated through simulation analysis.**

***Index Terms*— Optical networks, protection, capacity re-provisioning, robustness.**

I. Introduction

As THE size and the complexity of optical networks continue to grow, dual failures become increasingly probable. Hence designing recovery algorithms to protect against such failures is a paramount concern. Normally, the network is protected against single failures using one of the protection schemes (e.g., dedicated/shared, path/link/segment [1,2,6]). Whenever a failure occurs, all affected connections are re-routed on their corresponding protection paths [1,2]. However, since protection resources may also be shared with other unaffected connections, these may become unprotected and vulnerable to subsequent failures [4,5]. Generally, unprotected connections can be classified into three types:

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 is still vulnerable to another failure that may affect its protection path.

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

Clearly, larger numbers of unprotected connections in the network increases its vulnerability to subsequent failures. To improve the overall service availability, re-provisioning [3,4,5] exploits the available capacity in the network to re-establish new backup paths for unprotected connections right after the recovery from the first fault without a priori knowledge of the location of the second failure. Now, as more connections are allowed to share their protection resources (e.g., to achieve better network utilization), more demands will be packed together. Hence a recovery from a failure will leave a larger number of demands in the network unprotected. Alternatively, limiting the level of sharability of resources will reduce the number of unprotected demands; however this evidently comes at the expense of reduced network performance since bandwidth will not be used efficiently.

Additionally, since re-provisioning makes use of available resources in the network to provision new protection capacity, limited level of resource sharability will yield a lower flexibility in finding and assigning resources. Therefore, it is clear that there are two conflicting design constraints: on one hand limited sharability 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 sharability may result in larger number of unprotected connections after the first failure with higher degree of flexibility in provisioning protection capacity. One objective of this work is to provide a study on the performance of capacity re-provisioning under different levels of sharability for protection resources.

We also propose a new re-provisioning algorithm and contrast its performance with a conventional scheme. Here, the objective of the algorithm is to reduce the total number of connections that have to be re-provisioned. The motivations are twofold: (1) reduce management overheads in simultaneously provisioning a large number of connections, and (2) to lower reservation contention between multiple unprotected connections trying to simultaneously establish backup capacity.

II. Network Re-provisioning

A re-provisioning algorithm typically takes several inputs including network topology/usage information and a list of unprotected demands (as classified earlier) that require re-provisioning. The algorithm then tries to establish backup lightpaths [5] for unprotected demands using the available capacity in the network (this algorithm is referred to as Scheme I). Clearly, when the level of sharability of protection resources is high, more connections are packed together and as a result a failure in the network will result in a larger number of unprotected connections. These connections, in turn, will require protection capacity re-provisioning in order to improve the network restorability. Note, the larger the number of unprotected connections in the network, the higher is the management overhead and may result in excessive contentions

Manuscript received December 8, 2004. The associate editor coordinating the review of this letter and approving it for publication was Prof. Maode Ma.

Chadi Assi and Wei Huo are with the Concordia Institute for Information System Engineering, Concordia University, Montreal, Quebec, Canada (e-mail: assi@ciise.concordia.ca).

Abdallah Shami is with the Department of Electrical and Computer Engineering, the University of Western Ontario, London, Ontario, Canada.

Nasir Ghani is with the Department of Electrical and Computer Engineering, Tennessee Tech. University, USA.

Digital Object Identifier 10.1109/LCOMM.2005.07004.

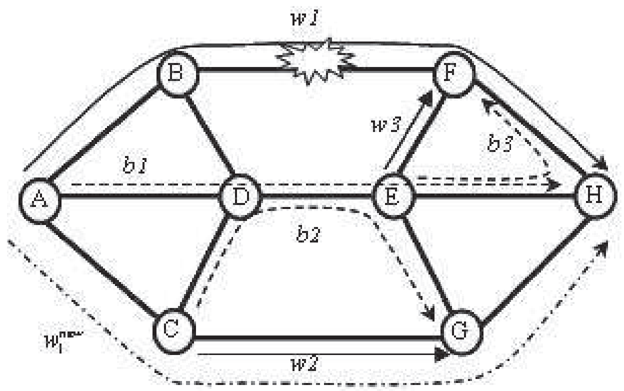

Fig. 1. Example of re-provisioning.

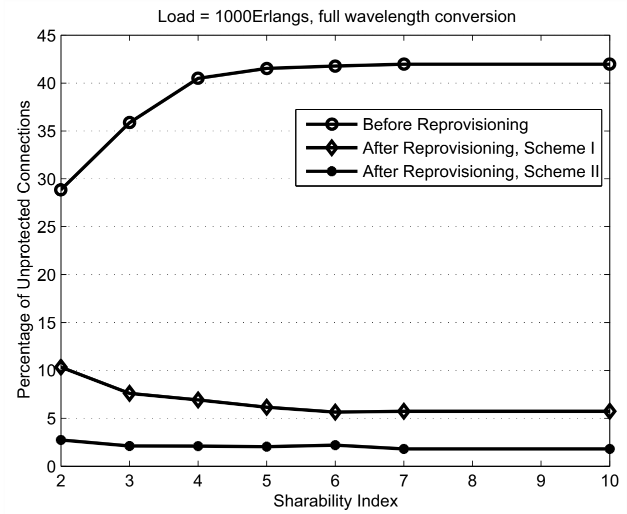

Fig. 2. Percentage of UC before/after re-provisioning.

among simultaneous demands attempting to reserve resources if re-provisioning is done under distributed control.

When a connection (w_i) is restored onto its backup (b_i), shared protection capacity along b_i becomes *temporarily unavailable* for other demands whose backup routes share that capacity. Instead of provisioning new backup capacity for these newly unprotected demands (whose number may be large when the level of sharability is high), a new working path w_i^{new} may be provisioned for each failed lightpath, w_i, that is link-disjoint with b_i. Upon successfully completing the provisioning of w_i^{new}, the traffic is simply reverted back from b_i to w_i^{new} leaving the rest of unaffected connections intact (this scheme is referred to as scheme II). Here, traffic is switched back to w_i^{new} upon successfully provisioning the required resources, thus avoiding any major traffic disruptions.

Note that protection capacity along b_i may not preserve its sharability status as w_i^{new} could be non link-disjoint with (some) demands whose protection lightpaths share protection capacity with b_i. In such a case, a new pair (w_i^{new}, b_i^{new}) are re-provisioned and traffic is reverted from b_i to w_i^{new}. Finally, if this step is not successful, the algorithm computes the set of unprotected connections resulting from the recovery of w_i and re-provision them accordingly (similar to scheme I). Note that, when wavelength conversion is deployed, only the links along b_i where protection wavelength(s) cannot be shared are identified and new protection wavelength(s) on those links are provisioned. Here even when the resource sharability degree is high, scheme II still outperforms scheme I since only a fraction of unprotected demands are re-provisioned.

The effectiveness of Scheme II is best shown via an illustrative example in Fig. 1. 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. Furthermore, when the resource sharability degree is very large, this number of connections to be re-provisioned under scheme II is substantially lowered resulting in a less management overhead and lighter impacts of contentions under distributed control.

III. Numerical Examples

We study the performance of lightpath re-provisioning in a sample core topology [5]. Requests are uniformly distributed between all source-destination pairs. The number of wavelengths per link is 64 and the load in the network is fixed at 1000 Erlangs; we simulate a failure on a unidirectional link.

Fig. 2 shows the percentage of unprotected connections (UC) in the network before and after re-provisioning as the level of resource sharability (SI) varies. Here, the level of resource sharability of a wavelength on a link indicates the number of connections allowed to be protected by this wavelength link. As the level of sharing increases, the percentage of UC (before re-provisioning) in the network after a failure dramatically increases. This is due to the fact that as a wavelength link protects more demands, the recovery of a single connection to this wavelength link will leave a larger number of connections unprotected. Clearly, capacity re-provisioning improves the network performance by substantially reducing

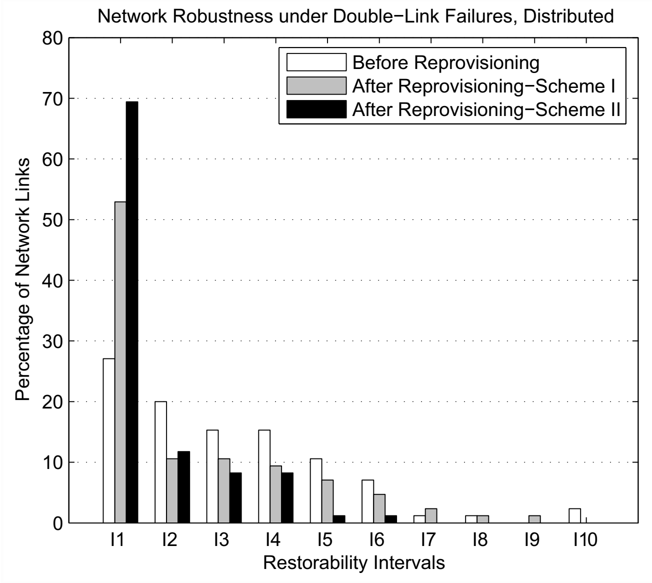

Fig. 3. Network robustness.

the percentage of UC (e.g., a decrease from 42% to 6% of UC at higher SI using scheme I) and therefore making the network less vulnerable to subsequent failures.

Note that a lower percentage of UC before re-provisioning at lower SI does not necessarily mean a good re-provisioning performance (i.e., a lower percentage of UC after re-provisioning). The reason is that *a lower SI will limit the flexibility of the re-provisioning algorithm in finding and judiciously allocating protection resources* among UC. The figure shows as the SI increases, the percentage of UC after re-provisioning slightly decreases for scheme I ($10\% - 6\%$) while it remains almost constant for scheme II ($\sim 3\%$) with better performance than that of scheme I.

Conversely, higher SI will allow the network to re-provision more UC by sharing the limited available resources. Therefore, the figure shows a larger performance gain at higher SI ($\sim 38\%(42\% - 6\%)$) than at lower SI ($\sim 20\%(30\% - 10\%)$ for scheme I. Scheme II shows a fixed percentage of UC at different SI. This 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, and as a result gains marginally less from a higher level of sharability. Note however that the re-provisioning gain is improved ($\sim 40\%$ at higher SI vs. 26% at lower SI) since more connections are admitted to the network at higher SI. Moreover, scheme II exhibits a superior performance over scheme I since the algorithm effectively re-provisions less UC.

Next, we study the impact of re-provisioning on increasing network robustness [6] under unlimited SI. Robustness is defined as the capability of the network to maintain high restorability of its connections (e.g., $\geq 95\%$) when two links are randomly taken down (one after the other). We measure the robustness before/after re-provisioning. Our evaluation is based upon measuring the percentage of links in the network that yields higher restorability after the first failure. The larger the fraction of links that yield higher restorability, the better is the overall robustness. In other words, given equal failure probability on all links, if restorability is kept at a desirable level for the majority of these links, then the network is said to be more robust.

Fig. 3 shows 10 different intervals for the network restorability ranging from $0 - 100\%$. Namely, one large interval (I_{10}) is chosen to cover a relatively low restorability range $0 - 73\%$ (interval 10) and the remaining intervals (I_1-I_9) are chosen in increments of 3% to cover higher ranges above 73% (e.g., $I_{10} = [0\% - 73\%[, \ldots, I_1 = [97\% - 100\%]$). Furthermore, the restorability, $R(i,j)$, of a double failure (i,j) is defined as the portion of all working paths $w_i + w_j$ on links i and j that are simultaneously affected and survive the failures. Fig. 3 shows the distribution of the number of links (percentage) with regard to restorability. Namely, it shows the robustness of the network as the probability of having the restorability (R) within a certain interval. When the network does not use re-provisioning, the 90% restorability is defined: $Pr(R \geq 90\%) = Pr(R \in I_1) + Pr(R \in I_2) + Pr(R \in I_3) = 0.32$. After re-provisioning using Scheme I, this value increases to $Pr(R \geq 90\%) = 0.865$ and even further to $Pr(R \geq 90\%) = 0.93$ using Scheme II. The results show that robustness improves substantially after re-provisioning. Prior to re-provisioning, the probability that the restorability is above 90% under any double link failure scenario is only 0.32 (i.e., only 32% of the network links yields restorability above 90% after first recovery). The results also show that Scheme II achieves significantly better robustness since the percentage of unprotected connections after re-provisioning ($\sim 3\%$) is much smaller than that under Scheme I ($\sim 6\%$), see Fig. 2.

IV. Conclusion

We studied the problem of improving robustness in optical networks with various resource sharability degrees under dual near-simultaneous failures. We showed that re-provisioning substantially improves the restorability and hence the robustness of optical networks after a failure and we compared the performance of two different algorithms.

Acknowledgment

The authors would like to thank the anonymous reviewers for their constructive comments. This work was supported by Natural Sciences and Engineering Research Council of Canada (NSERC), and FQRNT.

References

[1] J. Labourdette, "Shared mesh restoration in optical networks," in *Proc. OFC*, Mar. 2004.

[2] S. Ramamurthy *et al.*, "Survivable WDM mesh networks, part II - restoration," in *Proc. IEEE ICC*, vol. 3, pp. 2023-2030, June 1999.

[3] S. Kim *et al.*, "Evaluation of protection reconfiguration for multiple failures in WDM mesh networks," in *Proc. OFC*, vol. 1, pp. 210-211, Mar. 2003.

[4] R. Ramamurthy *et al.*, "Pre-emptive re-provisioning in mesh optical networks," in *Proc. OFC*, vol. 2, pp. 785-787, Mar. 2003.

[5] J. Zhang *et al.*, "A comprehensive study on backup re-provisioning to remedy the effect of multiple-link failures in WDM mesh networks," in *Proc. IEEE ICC*, vol. 3, pp. 1654-1658, June 2004.

[6] M. Clouqueur *et al.*, "Availability analysis of span-restorable mesh networks," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 810-821, May 2002.