ELSEVIER

Information Sciences 149 (2003) 183–195

INFORMATION SCIENCES
AN INTERNATIONAL JOURNAL

www.elsevier.com/locate/ins

Performance evaluation of optical mesh restoration schemes

Sunggy Koo *, Suresh Subramaniam

Department of Electrical and Computer Engineering, George Washington University, Washington, DC, USA

Received 27 September 2001; accepted 8 May 2002

Abstract

This paper investigates the performance of various schemes for restoration of lightpaths in an optical mesh network. We consider schemes in which backup paths for links or paths are pro-computed, but wavelength assignment is either pro-computed (static) or done dynamically upon failure (dynamic). In static restoration we compare the performance of link and path restoration in terms of capacity requirement and restoration time. For the dynamic restoration, we consider several dynamic wavelength reservation schemes such as backward and forward reservation with different portion of wavelength to be reserved in both directions. Their performance is evaluated in terms of restoration speed and restorability. Simulation results indicate that path protection requires less capacity than link protection for 100% restorability, and dynamic restoration with backward wavelength reservation shows better restorability than static restoration schemes, when the extra capacity is relatively small. Forward wavelength reservation shows faster restoration speed with lower restorability than the backward wavelength reservation. In general, the optimal restoration scheme depends on extra capacity available and restoration time limit.

Keywords: Restoration; Protection; Optical mesh network; Restorability; Restoration speed; Restoration capacity

* Corresponding author.
E-mail addresses: sunggy@seas.gwu.edu (S. Koo), suresh@seas.gwu.edu (S. Subramaniam).

0020-0255/02/$ - see front matter
PII: S0020-0255(02)00256-6

1. Introduction

To accommodate the exponential growth of the Internet, transport networks based on wavelength-division multiplexing (WDM) technology [9] are increasingly being deployed in carrier networks. A WDM network provides *lightpaths*, high-capacity virtual links, to higher layer client networks. Each lightpath in a WDM optical network may operate at data rates of up to several gigabits per second. Since WDM networks carry such high volumes of traffic, a single failure of any of the network components may have severe consequences.

In order to maintain a high level of service availability, carriers and/or service providers usually over-provision their network with extra capacity to protect their working components. When a network failure occurs the restoration procedure has to be invoked by a pre-assigned restoration source node according to a pre-planned restoration mechanism depending on the availability of network resources and the quality of service requirements. In this paper restoration and protection are used interchangeably.

In this performance evaluation study, we consider schemes in which backup paths for links or paths are pre-computed, but wavelength assignment is either pre-computed (static) or done dynamically upon failure (dynamic). In static restoration we compare the performance of link and path restoration in terms of capacity requirement and restoration time. For the dynamic restoration, we consider several dynamic wavelength reservation schemes such as backward and forward reservation. Their performance is evaluated in terms of restoration speed and restorability.

Many papers on optical layer restoration have been published in the literature [1–5,10]. Link-based protection methods and end-to-end path protection and restoration are described in [2–5]. In [1], the authors propose a distributed method to do restoration in mesh networks of Optical Cross-Connects. They assume full wavelength conversion where a lightpath can use a different wavelength on each link along its path, whereas in our model we assume an optical network with no wavelength conversion capability. Working and restoration capacity assignment problems were addressed in [2]. In [7], the authors present capacity-assignment algorithms and signaling protocols for various restoration methods and evaluate their performance in terms of capacity efficiency, restoration speed and proportion of connections that can be restored successfully. In our study, we use the same framework as in [6,7] with more restoration schemes such as link-based restoration and more dynamic wavelength reservation methods such as single/partial/full forward wavelength reservation. Some analysis on restoration time in optical networks is considered in [3] where a model for the restoration time calculation is presented but the model does not consider any message queueing delays. In our study as well as in [1,6,7], however, a significantly more detailed model is used, where a failure situation is simulated with consideration of possible queueing delays in each node.

The rest of the paper is organized as follows. In Section 2, we discuss restoration strategies with different options for wavelength assignment and reservation in optical WDM networks. Section 3 describes the network architecture including our basic node-link model. In Section 4, we explain restoration signaling protocols that are based on recent proposals of Generalized Multi-protocol Label Switching (GMPLS), enhanced with our extensions specifically designed for restoration purposes. In Section 5, we present simulation results on the performance of various restoration schemes in terms of capacity requirement, restoration speed and restorability, focusing on the trade-offs between them. Conclusions are given in Section 6.

2. Protection/restoration strategies

2.1. Protection

There are two major categories for protection schemes. One is path protection where the restoration processes are done between end-nodes of lightpaths. The other is link protection where the two end-nodes of the failed link take care of the restoration procedures. Fig. 1 shows an example of each protection scheme. In path protection, each lightpath affected by the failure is re-routed on another link disjoint path by the source and destination node of that lightpath. In link protection, the two nodes adjacent to the failed link reroute all the failed lightpaths so that the failed link is bypassed, but the original paths are otherwise unchanged.

End-to-end path protection many take longer time to restore affected traffic since the failure has to be notified to one or both of two end-nodes. On the other hand, link protection may require more protection capacity since all lightpaths are forced to use unaffected portions of their routes. Note that a

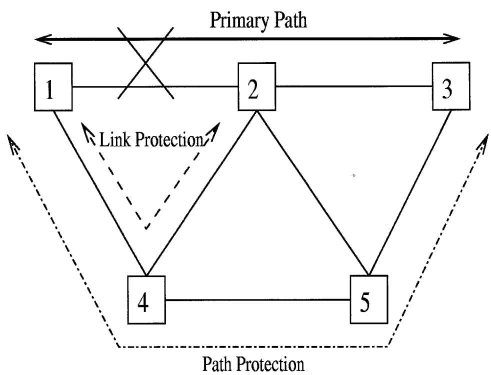

Fig. 1. Path and link protection.

separate backup lightpath may be used to restore each wavelength on the failed links, in general. When their corresponding primary paths will not be affected by any single link failure on the network, protection capacity on a link can be shared by more than one backup paths in all the schemes we consider.

2.2. Wavelength assignment (WA)

In this study, we consider a static traffic pattern in which a set of working lightpaths is given and routing and wavelength assignment is done off-line. It is assumed that no node has wavelength conversion capability, which means that the wavelength continuity constraint applies to every lightpath. We also assume that the protection paths (for the lightpaths in the case of path protection, and for the wavelengths in the case of link protection) are also precomputed.

However, we consider two different strategies for assigning wavelengths. In static WA (ST-WA), a wavelength is pre-assigned for the backup paths so that it is independent of the failed link, whereas in dynamic WA (DY-WA), wavelengths are assigned to the protection paths dynamically through a distributed signaling procedure. In the dynamic case, it is possible that one or more lightpaths fail to be restored due to contention among those lightpaths simultaneously trying to reserve resources (wavelengths) for their restoration paths upon a network failure. It is expected that DY-WA will take more time for restoration due to the extra messages for finding a free wavelength and reserving it. However DY-WA has the potential for using lower capacity than ST-WA because it is invoked *after* the failure. Note that DY-WA makes sense only for path protection schemes.

2.3. Wavelength reservation in DY-WA

We now present two schemes for a lightpath to find and reserve a wavelength in DY-WA. These are based on the reservation schemes proposed for lightpath provisioning in [8] and used for restoration in [6,7]. Upon a failure, one of the two end nodes of a failed lightpath or failed link is designated as the source. DY-WA method can be classified as forward reservation (FR) or backward reservation (BR), depending on the direction of the message that makes the actual wavelength reservation as it traverses hop by hop. In forward reservation, the available wavelengths are reserved on a forward pass to the destination in a greedy manner. On the other hand, in backward reservation, the availability of wavelengths along a backup path is determined using a probe message in a forward pass from the source to the destination, and the wavelengths are actually reserved on the backward pass from the destination to the source. Even with backward reservation, there is no guarantee for the wavelengths available during the probing period to remain available until the actual reservation takes place due to the distributed nature of resource reservation. We assume that all wavelengths used on the links of the unaffected portions of the primary paths will be released and available for other backup paths immediately after each node comes to know of the failure.

3. Network architecture

In our network architecture that is same as the one found in [6,7], each *node* of the network includes the units shown in Fig. 2. The message processor, optical cross-connect (OXC) controller, and message transmitter are represented as First In First Out (FIFO) queues with deterministic service times. All of the service times for these queues are set as input parameters to our simulation model since these rates correspond to computing power of CPU, transmission rate and speed of optical switch, which depend on fast-changing technology. (More detailed descriptions of service rates are given in Section 5.)

As in [6,7], we assume that the network has an OAM (operation, administration, and maintenance) wavelength dedicated for exchanging control messages for signaling. We also assume that the links are bi-directional, either consisting of a single fiber carrying light in two directions, or two uni-directional fibers carrying light in opposite directions. Without loss of generality, we assume the latter case here.

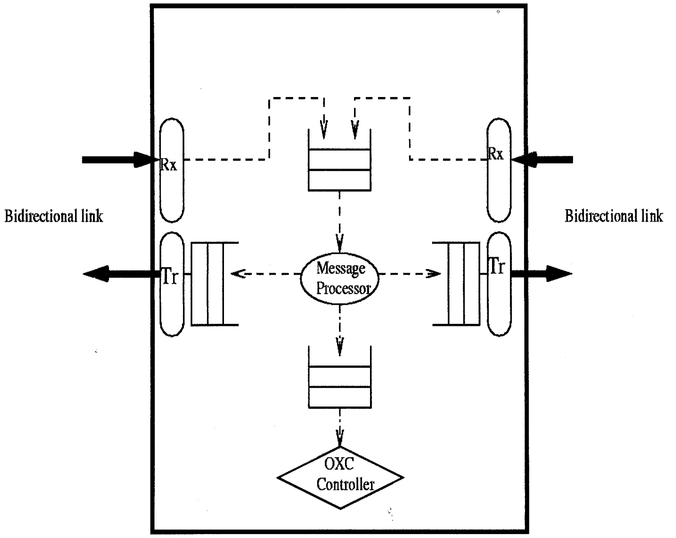

Fig. 2. Node and link model for restoration signaling.

Each node keeps a database of the status of wavelengths on (uni-directional) links outgoing from the node. The message processor of a node processes a received message based on a pre-defined signaling protocol (see Section 4 for more details). It has access to information regarding the availability of wavelengths on the outgoing links from the node, and it can change their status when needed. The algorithmic decisions by the message processor are based on the information available in the message, such as message type and node type (source/destination/intermediate).

The message processor is connected to the OXC/switch controller, which controls the node's switching process to cross-connect the lightpaths. When the message processor decides that OXC/switch configuration needs to be changed, it issues a switch re-configuration command to the switch controller. The switch controller, which has a buffer to store multiple switch re-configuration commands if needed, implements the commands sequentially or in a batched way. Similarly, the message processor and message transmitter buffer messages when they receive new messages while still handling other messages.

4. Restoration signaling: an extension to GMPLS

In order to implement various restoration schemes we propose a simple extension to one of the signaling protocols of GMPLS, namely, resource reservation protocol-traffic extension (RSVP-TE) [11,12]. RSVP-TE is a signaling protocol for GMPLS control plane proposed by IETF. In static wavelength assignment, our extension to RSVP-TE is to add a two-way hand shake between the end-points of a backup path to make sure that the restoration path is successfully activated. For dynamic wavelength assignment, however, we add a *probing/reserve* message exchange in addition to the hand shaking process.

In ST-WA, the only signaling that is necessary upon link failure is to configure the OXC/switch fabrics on the backup paths. As soon as a source node notices a failure, it sends out a SET_PATH message to its destination node along the pre-calculated backup path with the pre-assigned wavelength as in [6,7]. Every node on the backup path re-configures its OXC appropriately when it receives the SET_PATH message. Upon completion of an OXC re-configuration, a node sends an OXC_DONE message to the source to inform that it is ready for restoration. When OXC_DONE messages have arrived at the source from all nodes on the backup path, a 'hand-shaking' process takes place between the source and destination to finalize the restoration procedure.

Since no wavelength is pre-assigned in DY-WA, a source sends messages towards its destination in an effort to find a wavelength on the backup path, before configuring the backup path as described in ST-WA. The actual message exchange depends on the specific wavelength reservation (WR) method: backward/forward.

In backward reservation (BR) as in [6,7], the source node sends out a PROBE_W message that collects information about free wavelengths on the backup path as it traverses toward the destination hop by hop. If any of the intermediate nodes finds there is no continuous wavelength available, it sends back a PROBE_NACK message to the source node. If a non-empty set of wavelengths arrives at the destination node, then the destination node selects one (Single:S) or some (Partial:P) or all (Full:F) wavelength(s) out of the set and sends a RESERV_W message to the source node with the selected wavelength(s). Every node on the way back to the source tries to reserve the selected wavelength(s). If it fails to reserve the wavelength(s) due to the contention from other backup paths, then it sends back a RESERV_NACK message to the destination node. Once a wavelength is successfully reserved, the same activation of backup path as in ST-WA will take place.

On the other hand, in forward reservation (FR), the source node sends RESERV_W message with one or more wavelengths (depending on S, P, F) available on its outgoing link. If the destination receives a non-empty set of reserved wavelengths, it selects one and activates the backup path with the selected wavelength.

Because many backup paths may invoke the distributed WR procedure when a link fails, it is possible that a reservation attempt ends in failure due to reservation conflicts. Each path is allowed to re-try WR after a random back-off time. The maximum number of re-tries allowed is an important factor in determining restorability and restoration time.

5. Simulation and numerical results

We used the NSFNET topology that has 14 nodes and 21 bi-directional (42 uni-directional) links. For a given number of lightpaths N, source and destination nodes were randomly chosen. Shortest path routing was used to select the primary path for each lightpath, and the first-fit WA was used for ST-WA.

The service rates are: message processing rate (P messages/s), cross-connect configuration rate (X cross-connects/s) and message transmission rate (T message/s). We used the following values for the service rates are: $P = 3000$, $X = 1000$ and $T = 500,000$. We assume that each link is 400 km long. In DY-WA, the back-off time is randomly chosen between 1 and 5 ms.

5.1. Restoration capacity vs. speed

We first obtain the capacity (the highest indexed wavelength that is used on a link) for different cases. Fig. 3 shows how capacity scales with N for no protection, path protection (ST-WA), and link protection (LP). We next show the average restoration times for path protection (PP) with ST-WA and

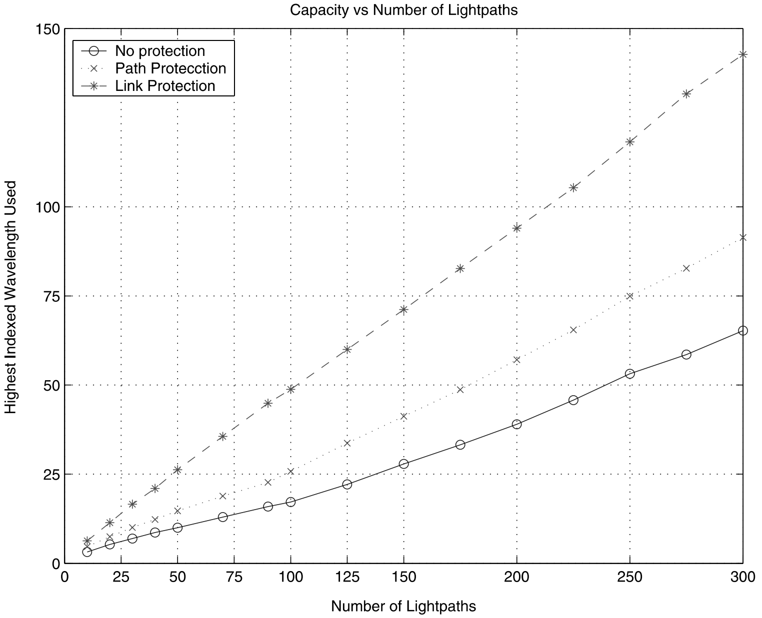

Fig. 3. Capacity vs. number of lightpaths, *N*.

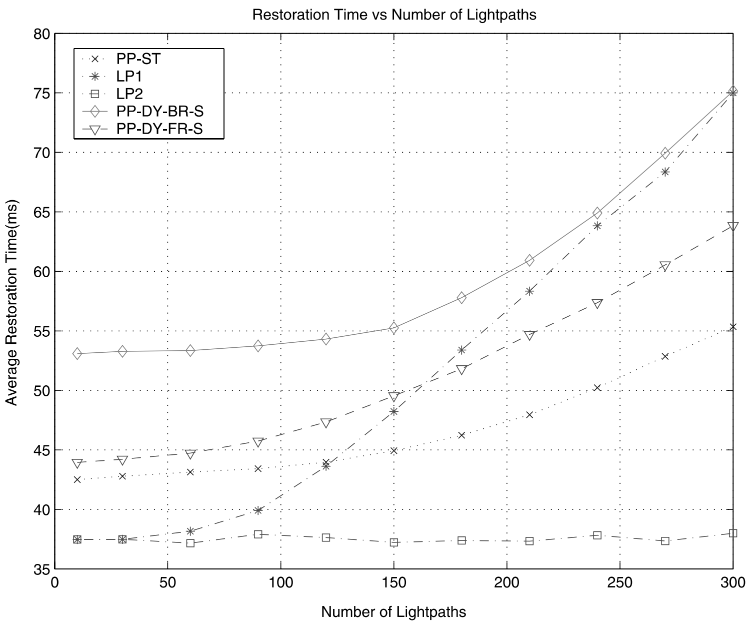

Fig. 4. Average restoration time vs. number of lightpaths, *N*.

DY-WA (using BR-Single and FR-Single), and link protection (LP1 and LP2) as a function of N in Fig. 4. In LP1, we assume that all lightpaths on the failed link will be re-routed along the same backup route and an OXC re-configuration command for each path is processed one by one, experiencing possible queuing delays in the OXC controller queue at a node.

As shown in Fig. 4, the average restoration time for LP1 sharply increases with N due to the queuing delay. If, however, all of the OXC commands at a node are batched up and processed at once, the queuing delay can be reduced significantly (LP2 in Fig. 4), and restoration time does not scale up with N. The difference in capacity requirement between PP and LP is significant for all N. When $N = 300$, for example, link protection requires 188% extra capacity (relative to working capacity) for restoration, whereas path protection needs only 40%. One trade-off between capacity requirement and recovery speed that we can observe from Figs. 3 and 4 is that, for $N = 300$, LP2 performs faster recovery by 20 ms with about 70% more capacity than PP. In Fig. 4, we also note that BR-S has a longer restoration time than FR-S due to the extra message exchanges in BR.

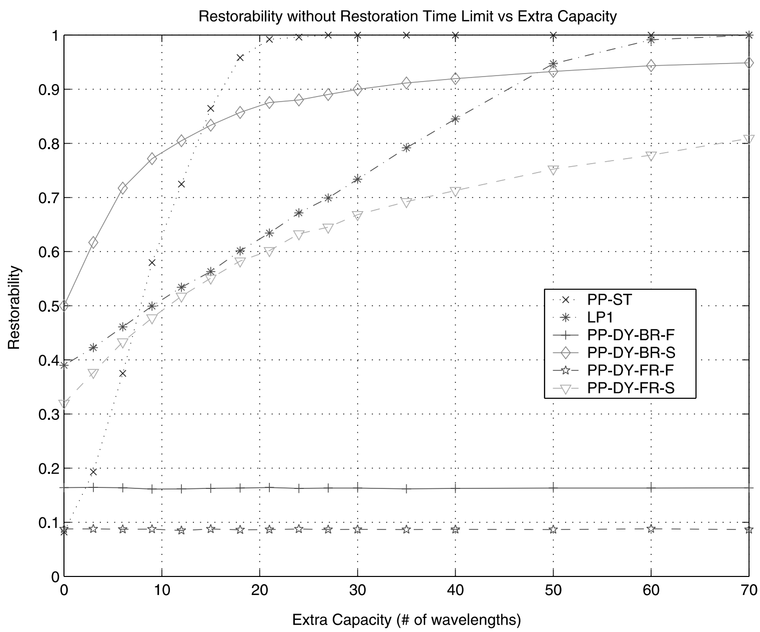

Fig. 5. Restorability vs. capacity, no restoration time limit, no re-tries allowed for DY-WA ($N = 200$).

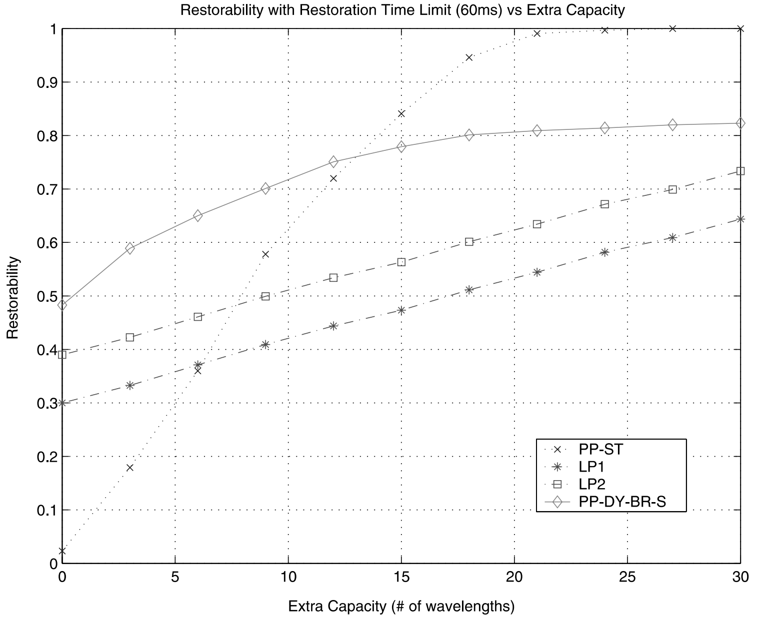

Fig. 6. Restorability vs. capacity with restoration time limit of 60 ms, re-tries allowed for DY-WA ($N = 200$).

5.2. Restorability vs. restoration (extra) capacity

In Fig. 5, we observe that when $N = 200$, PP-DY-BR-S outperforms PP-ST with up to 35% (approximately 15 more wavelengths) of extra capacity. Another trade-off between restorability and restoration capacity is shown in Fig. 6. An important observation here is, even with a tight limit (60 ms) on restoration time, that observation from Fig. 5 still holds.

We also note that link protection with batched process of OXC re-configuration (LP2) shows about 10% higher restorability than with non-batched process (LP1) with restoration time limit.

5.3. Effects of number of re-tries and restoration (extra) capacity

Finally, we present the effect that both number of re-tries and extra capacity have on restorability and restoration time in Figs. 7 and 8. It is interesting that, in PP-DY-BR-S, a large number (more than 10) of re-tries does not improve restorability significantly. We also note that average restoration time, when $N = 200$, also remains steady with a large number of re-tries (more than 10)

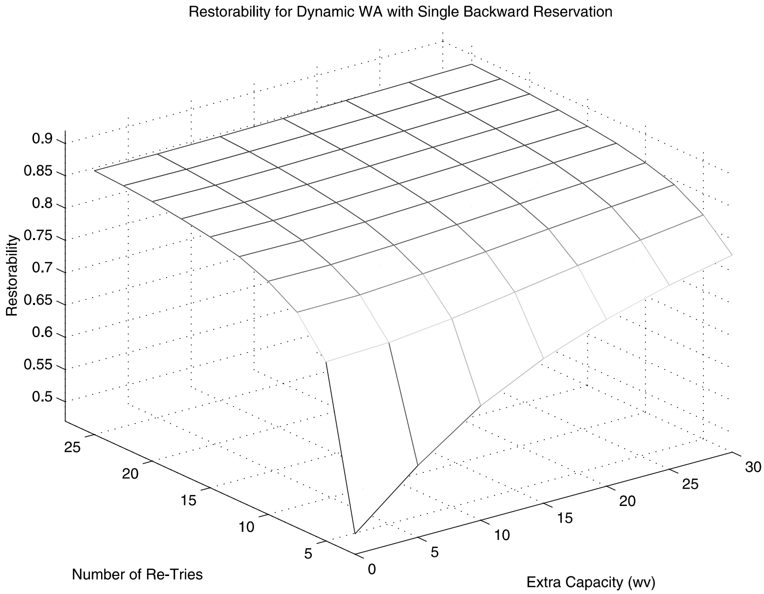

Fig. 7. The effect of number of re-tries and extra capacity on restorability with PP-DY-BR-S ($N = 200$).

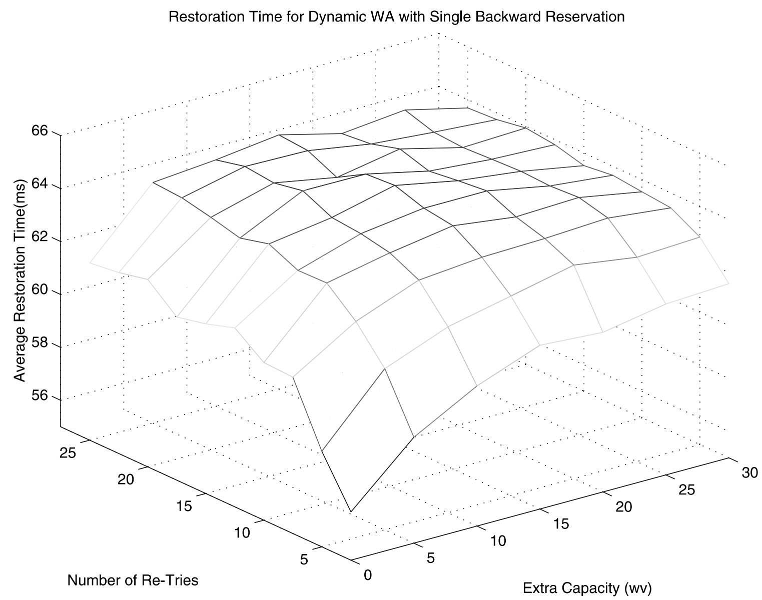

Fig. 8. The effect of number of re-tries and extra capacity on restoration time with PP-DY-BR-S ($N = 200$).

and large amounts of extra capacity (more than 15 wavelengths). This indicates that restorability is limited by capacity rather than number of re-tries for moderately large extra capacities.

6. Conclusions

We have investigated the trade-offs between speed, capacity and restorability for various mesh restoration methods with the goal of providing guidelines for choosing the optimal restoration scheme. Link based protection may provide a faster restoration of affected traffic at the cost of more capacity required for protection purpose. On the other hand, end-to-end path protection requires less capacity than link protection. However, since restoration processes are managed by the two end nodes of a failed lightpath, it takes longer restoration time than link-based protection. Our results show that, with relatively small extra capacity, there is a role for dynamic wavelength assignment in restoring failed lightpaths within a limited restoration time.

Acknowledgements

This work was supported in part by NSF Grant ANI 9973111 and by DARPA Grant N66001-00-18949 (co-funded by NSA).

References

[1] R. Doverspike, G. Sahin, J. Strand, R. Tkach, Fast restoration in a mesh network of optical cross-connects, in: Proceeding of OFC, February 1999.
[2] S. Ramamurthy, B. Mukherjee, Survivable WDM mesh networks, part I – Protection, in: Proceeding of INFOCOM, 1999, pp. 744–751.
[3] S. Ramamurthy, B. Mukherjee, Survivable WDM mesh networks, part II – Restoration, in: Proceedings of ICC, 1999, pp. 2023–2030.
[4] O. Gerstel, R. Ramaswami, Optical layer survivability: a service perspective, IEEE Communications Magazine March (2002) 104–113.
[5] O. Gerstel, R. Ramaswami, Optical layer survivability: an implementation perspective, IEEE JSAC 18 (10) (2002) 1885–1899.
[6] G. Sahin, Service and restoration routing and wavelength assignment and restoration signaling in optical networks, Ph.D. Thesis, University of Washington, August 2001.
[7] G. Sahin, S. Subramaniam, M. Azizoglu, Signaling and capacity assignment for mesh-based restoration schemes in optical networks, Journal of Optical Networking 1 (5) (2002) 188–205.
[8] Y. Mei, C. Qiao, Efficient distributed control protocols for WDM all-optical networks, in: Proceedings of ICCCN, 1997, pp. 150–153.

[9] T.E. Stern, K. Bala, Multiwavelength Optical Networks: A Layered Approach, Prentice-Hall, Englewood Cliffs, NJ, 1999.
[10] D. Zhou, S. Subramaniam, Survivability in optical networks, IEEE Network 14 (6) (2000) 16–23.
[11] A. Banerjee et al., Generalized multiprotocol label switching: an overview of routing and management enhancements, IEEE Communications Magazine January (2001) 144–150.
[12] A. Banerjee et al., Generalized multiprotocol label switching: an overview of signaling enhancements and recovery techniques, IEEE Communications Magazine July (2001) 144–151.