Using Local Information for WDM Network Protection

Hungjen Wang, Eytan Modiano, Muriel Médard
{hjwang, modiano, medard} @mit.edu
Laboratory for Information and Decision Systems, MIT

Abstract

Path protection and link protection schemes are the main means of protecting wavelength-division multiplexed (WDM) networks from the losses caused by a link failure such as a fiber cut. We propose a new protection scheme, which we term partial path protection (PPP), to select end-to-end backup paths using local information about network failures. PPP designates a different restoration path for every link failure of every primary path. PPP allows the re-use of operational segments of the original primary path in the protection path. A novel approach used in this paper is that of a dynamic call-by-call model with blocking probability as the performance metric. This is in contrast with traditional approaches to restoration, which consider capacity-efficiency for batch call arrivals. Since optimizing the blocking probability is a large dynamic optimization problem, we present two heuristics for implementing PPP. We show that a simple method based on shortest path routing for which primary paths are selected first is more effective than a greedy approach that minimizes, for each call arrival, the number of wavelengths used by the primary and backup path jointly.

Protection Schemes

Path protection (PP) and link protection schemes are the current main approaches of protecting wavelength-division multiplexed (WDM) networks against the losses caused by a link failure such as a fiber cut [1,2,3,4,5,6]. Basically, PP requires the protection path of a request to be completely link-disjoint from the corresponding primary path, while the link protection scheme reroutes all affected requests over a set of predetermined paths between the two nodes terminating the failed link. In general, PP is more capacity efficient than link protection [4]. An intermediate approach is span protection, in which portions of paths are protected.

In this paper, we propose a new protection scheme, the partial path protection (PPP) scheme. In PPP, the system specifies a specific end-to-end protection path for *each link* along the primary path. Thus, just like PP, PPP also assigns “end-to-end” protection paths to primary paths, however, in PPP, a single protection path protects only one specific link failure on a primary path, instead of the whole primary path. Thus, PPP is a very special case of span protection, in which the spans vary according to the location of the failure. For example, in Fig.1, a call with source node 1 and destination node 6 has a primary path 1-3-2-5-6.

As illustrated in Table 1, the system applying PPP specifies alternative restoration paths to protect the network from the losses caused by a link failure. Notice that each of these protection paths needs to be link-disjoint only from the link it protects. On the other hand, when applying PP, the network cannot find a protection path for the primary path shown, since there exists no complete link-disjoint path from the primary path connecting the source-destination pair. In short, comparing PPP with PP, we see that the former is more flexible than the latter. Indeed, any path protection scheme is a valid PPP, whereas the reverse does not hold. We expect, therefore, that PPP will enhance system ability to provide protection over the traditional path protection.

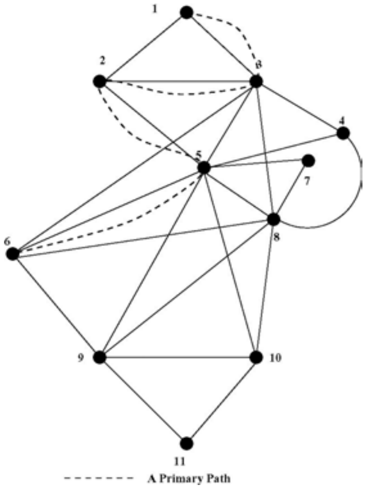

Figure 1. The 11 node, 23 links New Jersey LATA Network

For path protection, since we assume that only a single link failure can occur at a time, a system can allow primary paths with no link in common to share protection bandwidth. We call this protection sharing.

In addition to protection sharing, PPP further allows a protection path to share bandwidth with segments of the primary path that remain operational after link failure. To differentiate both protection schemes in protection sharing, we consider the network in Fig. 2 and assume the network now serves two call requests, (1,5) and (5,4), in sequence. Table 2 shows the resource assignments for primary and protection paths under the PP and the PPP respectively. By exercising protection sharing, the system reserves only one wavelength for protection on link (3,4), which is the key factor for PPP to have a better performance as calls accumulate.

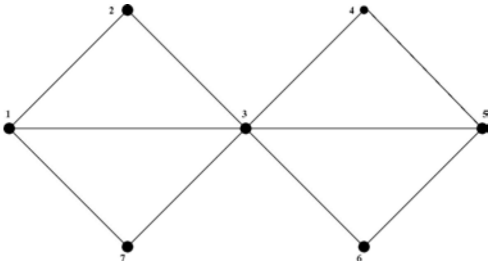

Figure 2. A network illustrating PPP and PP in protection sharing

Link on Primary Path 1-3-2-5-6	Corresponding Protection Path
(1,3)	1-2-5-6
(3,2)	1-3-5-6
(2,5)	1-3-5-6
(5,6)	1-2-3-6

No backup path found for system with PP

Table 1. Backup paths for the primary path in Fig. 1

	SD Pair	Primary	Protection Path (protected link)
Path Protection Scheme	(1,5)	1-3-5	1-2-3-4-5 (1-3)
			1-2-3-4-5 (3-5)
	(5,4)	5-4	5-3-4 (5-4)
Partial Path Protection Scheme	(1,5)	1-3-5	1-2-3-5 (1-3)
			1-3-4-5 (3-5)
	(5,4)	5-4	5-3-4 (5-4)

Table 2. Resource allocation for SD pair (1,5) and (5,4) in Fig. 2

For path protection, since we assume that only a single link failure can occur at a time, a system can allow primary paths with no link in common to share protection bandwidth. We call this protection sharing.

In addition to protection sharing, PPP further allows a protection path to share bandwidth with segments of the primary path that remain operational after link failure. To differentiate both protection schemes in protection sharing, we consider the network in Fig. 2 and assume the network now serves two call requests, (1,5) and (5,4), in sequence. Table 2 shows the resource assignments for primary and protection paths under the PP and the PPP respectively. By exercising protection sharing, the system reserves only one wavelength for protection on link (3,4), which is the key factor for PPP to have a better performance as calls accumulate.

Path Assignment Approaches

We consider two approaches to implement PP and PPP in this paper. The first approach is a *greedy* approach. For each call request, the system uses the fewest previously unused wavelengths to establish the primary and protection paths jointly. Wavelengths already used for protection paths can be used for new protection paths as long as a single link failure does not entail the activation of more than one protection path on that wavelength on any link. The problem formulation is an integer linear program (ILP), a common approach to network routing [3,4,5]. Owing to the space limitation, we omit the ILP formulation here. The second approach first selects the primary path, using a shortest path route. It then selects the protection paths using a shortest path algorithm in which wavelengths already assigned for protection can be used at no additional cost. In PP, the system picks only one backup path for a primary path, while, in PPP, the system selects specific backup path for each link along a primary path. We term this method the *shortest path* approach (SP).

From a computational complexity perspective, the greedy approach is much more complex than the SP solutions. The main reason is that the greedy approach essentially solves a discrete optimization problem, which consumes intensive computing power in most cases, whereas the SP approach can apply polynomial-time algorithms, such as Dijkstra's algorithm, to search for shortest paths for primary and backup paths rapidly. From the perspective of resource efficiency, we note that while the SP approach may require more resources for a given call initially; however, we observe from simulations that over a sequence of calls, the SP approach results in more efficient bandwidth utilization. One explanation for this occurrence is that the greedy approach happens to choose paths with no potential for protection sharing, harming network resource utilization; in contrast, though the SP is not optimal at first, it performs better over time, by encouraging protection sharing.

Simulations and Results

To investigate the protection schemes, we simulate PP and PPP schemes using both the greedy approach and the SP approach. We consider a dynamic call-by-call system with random arrivals, and the system dynamically allocates network resources for primary and restoration paths for a call request. In our call-by-call model, we focus on the problem of whether an available wavelength exists on a link, regarding the network as a circuit-switched network. In effect, we assume full wavelength conversion at all nodes. We also assume that the calls arrive according to a Poisson process and that calls have an Exponentially distributed service time. The traffic load refers to the product of the arrival rate and the average service time.

In our simulations, we consider the New Jersey LATA Network (NJLATA) in Fig. 1 and the NSFNET network in Fig. 3. The main measurements here are the network resource utilization and the steady state blocking probability. The network utilization refers to the number of wavelengths occupied by connections. Blocking probability is related to opportunity cost, referring to the additional revenue available if certain customers were not turned away. Fig. 4 and Fig. 6 present the simulation result for the blocking probability measurement. Fig. 5 and Fig. 7 show the results for network resource utilization.

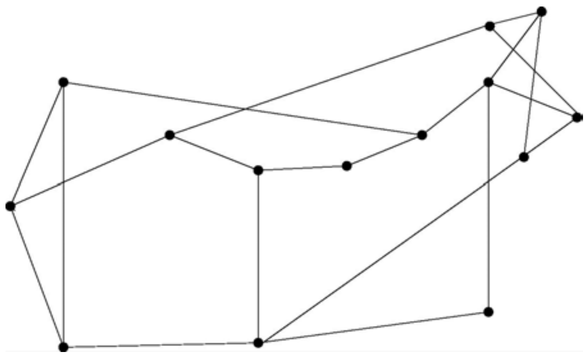

Figure 3. The NSFNET network

The major conclusion made from the simulations is that, as shown in Fig. 4 to 7, the PPP when implemented by the SP approach has the best overall performance. For example, in Fig. 6, with the blocking probability fixed at 0.01, PPP with SP approach made an improvement in traffic load around 20% as compared to PP. This implies that this SP-PPP approach can support 20% more traffic than the conventional PP approach for a 1% blocking probability. These observations meet our expectations. First, since the PPP is more flexible and efficient than PP as discussed, one can see that PPP outperforms PP in both implementations. Second, from the nature of the greedy algorithm, which attempts to occupy the minimum number of wavelengths to serve a call, PPP implemented by the greedy approach will use the fewest wavelengths for backup paths to protect the primary path. Consequently, one single backup path is typically used for each primary path, even though the PPP scheme does not require all the backup paths to be the same. Hence, with greedy approach, PPP does not take full advantage of the potential protection sharing.

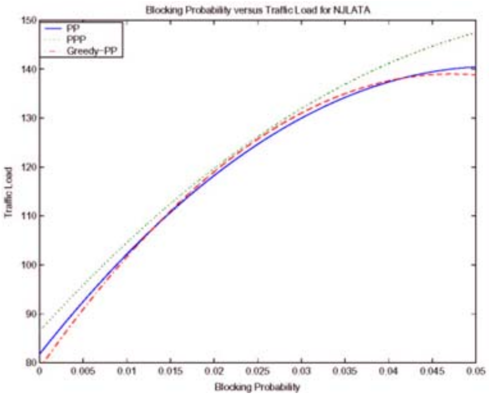

Figure 4. Traffic Load vs. Blocking Probability in NJLATA Network

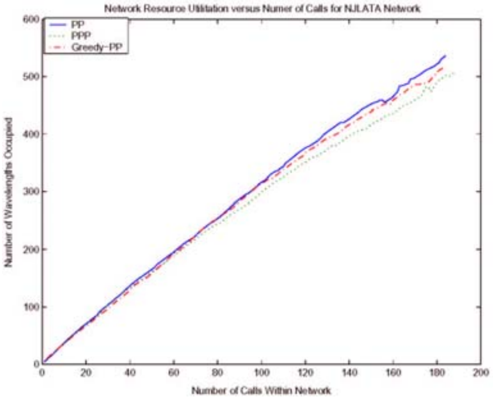

Figure 5. Network Resource Utilization vs. Blocking Probability in NJLATA Network

In contrast, the SP approach dynamically assigns a backup path to each link on a primary path without the constraint of being link-disjoint from the whole primary path, but from the link it protects. This weaker constraint encourages SP to promote potential protection sharing, not only because some segments of primary path could be used for backup purpose but also some wavelengths which protect some links on a primary path now can be assigned to protect other links which the wavelengths has not protected. As a result, the potential protection sharing is encouraged.

Conclusions

We have introduced a novel protection scheme, PPP. Moreover, instead of considering traditional static capacity-efficiency measures for evaluating the efficiency of protection schemes, we considered a dynamic call-by-call model. To avoid the complexity of dynamic optimization, we presented two heuristics to implementing path protection and PPP. These approaches, which we termed greedy and SP, were compared to each other for both path protection and PPP. We have demonstrated that PPP is superior to path protection and that SP is superior to the greedy approach. As expected from the fact that PPP is more general and flexible than path protection, PPP outperforms path protection in terms of resource utilization and blocking probability. Moreover, the SP approach performs better than the greedy approach. It is the dynamic nature of our problem that renders SP superior to the greedy approach. Indeed, SP emphasizes reducing resource use among primary paths, since their bandwidth cannot be shared.

The advantages of PPP over path protection have certain implications in the area of network management. Path protection only requires that the source and destination node be aware that a failure occurred somewhere along the primary path. Localization of the failure is unimportant, since protection takes place in the same way regardless of where the failure occurs. Thus, once the protection path has been set up, the network management does not need to have detailed knowledge of the nature of the failure to effect protection. Path protection can then be handled by higher layer mechanisms. For link protection, local information is needed by the nodes adjacent to the failure, but there is no need to manage protection on a path-by-path basis. Lower layers can therefore ensure link protection. PPP, on the other hand, requires on the part of the network management effecting protection knowledge of the path and of the location of the failed link. Our results point to the fact that visibility by the network management system across layers may be useful for performing protection efficiently.

There are several further research directions for our work. One such direction is to consider the case of batch arrivals rather than dynamic call-by-call arrivals. We expect that the preferable approach in the static batch case is to solve some ILP similar to the one set up for our greedy approach. Another area of further research is the generalization of our PPP algorithm to the case where failures are localized to segments, possible comprising several links. Such a generalization would allow us to study the effect upon blocking probability of different granularities of failure localization.

References

[1] Muriel Medard, Steven G. Finn, Richard A. Barry, and Robert G. Gallager, *Redundant Trees for Prepalnned Recovery in Arbitrary Vertex-Redundant Graphs*, IEEE/ACM Transactions on Networking, vol. 7, no. 5, PP. 641-652, Oct. 1999.

[2] Steven S. Lumetta, Muriel Medard, and Yung-Ching Tseng, *Capacity versus Robustness: A Tradeoff for Link Restoration in Mesh Networks,* Journal of Lightwave Technology, vol. 18, Issue 12, pp. 1765-1775, Mar. 2000.

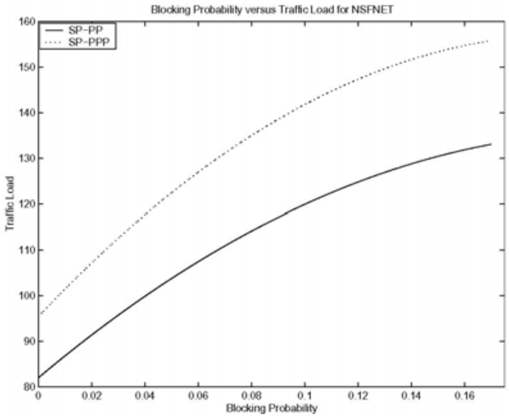

Figure 6. Traffic Load vs. Blocking Probability in NSFNET Network

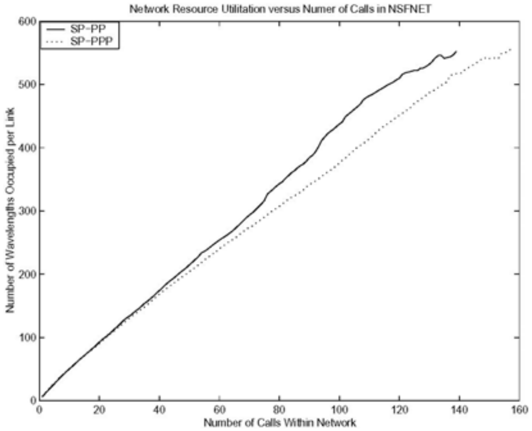

Figure 7. Network Resource Utilization vs. Blocking Probability in NSFNET Network

[3] Eytan Modiano and Aradhana Narula, Survivable Routing of Logical Topologies in WDM Networks, Infocom 2001, Anchorage, Proceedings, IEEE, vol. 1, pp. 348-357, 2001.

[4] S. Ramamurthy and B. Mukherjee, *Survivable WDM mesh Networks, Part I - Protection*, INFOCOM `99, Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, vol. 2, pp. 744-751, 1999.

[5] Murali Kodialam and T.V. Lakshman, *Dynamic Routing of Bandwidth Guaranteed Tunnels with Restoration,* Infocom 2000, April 2000.

[6] T. Wu, *Fiber Network Service Survivability*, Norwood, MA: Artech House, 1992.