Partial Protection in Networks with Backup Capacity Sharing

Greg Kuperman[†], Eytan Modiano[†], and Aradhana Narula-Tam[‡]

[†]*MIT LIDS*
77 Massachusetts Ave., Cambridge, MA 02139, USA
email: {gregk, modiano}@mit.edu

[‡]*MIT Lincoln Laboratory*
244 Wood Street, Lexington, MA 02420 USA
email: arad@ll.mit.edu

Abstract: We develop a novel network protection scheme that guarantees a minimum grade of service upon a link failure by utilizing backup capacity sharing between demands. Our scheme achieves up to 83% reductions in protection resources.
© 2011 Optical Society of America

1. Introduction

As data rates continue to increase, the failure of a network element can cause catastrophic service disruptions, leading to millions of dollars in lost revenue. Full protection schemes are used to protect against such failures; often doubling the amount of resources needed to route a connection. Consequently, many service providers do not offer any protection whatsoever. We propose a "partial protection" scheme, which offers quantifiable grades of protection at a fraction of the cost of full protection.

Full protection schemes have been studied extensively [1–3]. The most common scheme in backbone networks today is guaranteed path protection [2] (also known as $1+1$ protection), which provides an edge-disjoint backup path for each working path, resulting in 100% service restoration after any single link failure. In this paper, we consider an alternate form of guaranteed protection, where a fraction of a demand is guaranteed in the event of a link failure. Since link failures are relatively uncommon and are typically repaired quickly [2], users may be willing to tolerate a reduced flow if data rate guarantees can be made at a substantially reduced cost.

With our "partial protection" approach, a minimum fraction q of the demand is guaranteed to remain available between the source and destination after any single link failure. When q is equal to 1, the service is fully protected, and when q is 0, the service is unprotected. Our work is inspired by a quantitative framework for deterministic partial protection developed in [4]. In [5], a "theory" of partial protection was presented for a single commodity, and optimal algorithms for capacity allocation were presented. In [6], the savings that can be achieved by guaranteeing part of the demand in the event of a link failure, as opposed to full protection, are examined. Further, it was shown in [6] that the amount of protection that can be guaranteed depends on the topology of the network. However, the multiple commodity case, where backup capacity can be shared between demands, has not been previously investigated.

In this work, we introduce novel routing strategies that allocate working and spare capacity to meet partial protection requirements when backup resources can be shared between demands. In fact, we show that depending on the value of q, it may be possible to provide protection at minimal allocation cost. We consider two cases: preemptive and non-preemptive partial protection. For the preemptive case, primary resources available prior to a link failure may be preempted to provide backup for other demands, as long as all protection requirements are met after the failure. For the non-preemptive case, only demands that are directly affected by the link failure drop to the rates guaranteed under partial protection. We develop an optimal linear program formulation, as well as a time-efficient algorithm for dynamic (one-at-a-time) arrivals.

2. Partial Protection with Backup Capacity Sharing

Each demand needs to be routed from its source to destination such that upon a link failure, at least a fraction q of that demand is guaranteed to remain. Both working traffic and protection flows (defined as the flow after a failure)
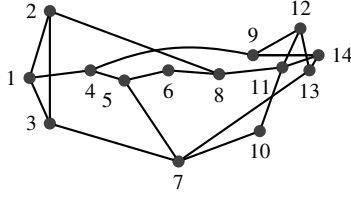
Fig. 1: 14 Node NSFNET backbone network

can be bifurcated to traverse multiple paths between the source and destination. The objective is to minimize the total capacity needed to support the primary flows before a link failure and the partial protection requirements afterwards.

Backup capacity sharing for $1+1$ protection has been previously explored in [1, 3]. To briefly summarize, if two primary demands are edge-disjoint from each other, then under a single-link failure model, at most one can be disrupted at any given point in time. Hence, the two demands can share backup capacity since only one will need to be restored after a link failure. Our partial protection framework allows the primary flow to be split along multiple paths, thus reducing the cost of providing protection. This is due to the fact that splitting the primary flow distributes the risk so that upon an edge failure, less primary flow is disrupted, which then requires fewer protection resources.

Restoring a fraction q of any demand, as opposed to restoring all of it, allows for protection sharing that is not possible when the full demand needs to be restored. In fact, sometimes it is possible to meet partial protection requirements using only the shortest path routing. If demands are allowed to drop down to their partial protection flows, then primary allocation from the demands that did not fail can be "preempted" to protect other demands, thus significantly reducing the need for spare capacity allocation. While preemptive protection may achieve the minimum possible allocation, it comes at the cost of exposing connections to more frequent failures. As an alternative, we consider the non-preemptive case, where only demands that are directly affected by the link failure have their flow reduced.

As opposed to single-commodity formulations, sharing backup capacity for multi-commodity flows is significantly more complex [1]. We develop a linear program (LP) for the preemptive and non-preemptive case when bifurcation is allowed. A non-bifurcating protection scheme similar to $1+1$ protection called $1+q$ protection was presented in [5]. This $1+q$ protection scheme routes one unit along the primary path and q along a disjoint protection path. A mixed integer linear program (MILP) is developed for $1+q$ protection with backup capacity sharing. The LP's and MILP are not presented in this paper due to space constraint, but can be found in [7].

We first consider the batch case where demands are routed jointly. The various partial protection schemes are compared via simulation using the NSFNET topology (Fig. 1) with 100 random unit demands. The protection requirement, q, for each demand has a truncated normal distribution with standard deviation $\sigma = \frac{1}{2}$. The mean of q is varied between 0 and 1 for each iteration. Routing solutions under the various protection schemes were determined using the CPLEX optimization solver.

The average costs to route the demand and protection capacity using the different routing strategies are plotted in Fig. 2 as a function of the expected value of q. Shortest path routing without protection considerations is used as a lower bound for the allocation cost. The cost of providing partial protection q is the difference between the cost of the respective protection strategies and the shortest path routing. In this simulation, preemptive partial protection is able
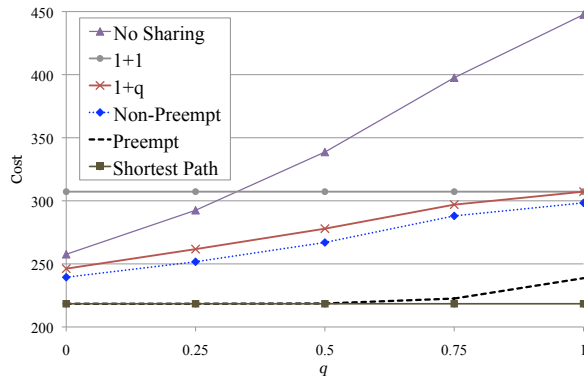

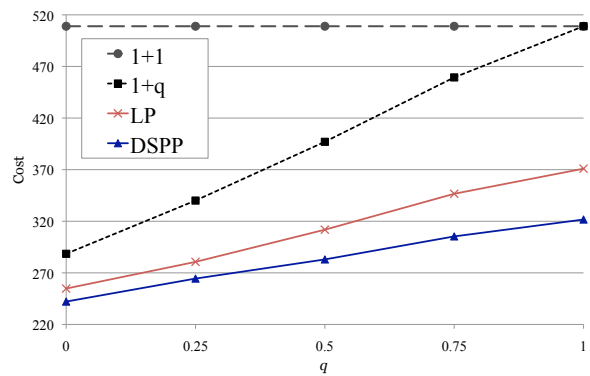Fig. 2: Capacity cost vs. q for batch routing

Fig. 3: Capacity cost vs. q for dynamic arrivals

to meet requirements using only the capacity needed for the shortest path routing for $q \leq \frac{1}{2}$, and only an additional increase in total capacity of 2% for $q \leq \frac{3}{4}$. When considering savings in excess resources, preemptive partial protection achieves reductions of 83% at $q = 1$ over both $1 + 1$ protection and $1 + q$ protection, and 92% over the non-sharing protection. Non-preemptive shared partial protection, at $q = \frac{1}{2}$, achieves reductions in excess resources of 59% over $1 + 1$ shared protection and 19% over $1 + q$ shared protection. When compared to the non-sharing protection scheme, non-preemptive partial protection achieves reductions of 59% in excess resources at $q = \frac{1}{2}$ and 66% at $q = 1$.

We next consider the case of dynamic (one-at-a-time) arrivals, and present a time-efficient algorithm to find a minimum-cost capacity routing. Determining how much backup capacity can be shared for 1+1 guaranteed protection with dynamic arrivals was investigated in [1, 3]. In those papers, conflict sets were used to determine potential backup sharing on an edge by examining of how much backup capacity was allocated on one edge to protect against the failure of another. If more backup capacity is already allocated on some edge than is needed to protect for the failure of another edge, then that edge's backup capacity can be shared. This model can be extended to partial protection by guaranteeing that any particular demand has its partial flow requirement met after a failure.

For the case of one-at-a-time routing, previous works offer heuristics to jointly optimize the primary and backup path for each incoming demand, as was done in [1, 3]. We instead choose a simple strategy of using the shortest path for the primary route. Our simulations show that using the shortest path for the primary route in fact performs better than jointly optimizing the primary and backup paths for each incoming demand. We call our algorithm Dynamic Shared Partial Protection (DSPP).

We compare, via simulation, DSPP to $1 + 1$, $1 + q$ and the non-preemptive LP, each of which jointly optimizes the primary and backup paths for each incoming demand (a "greedy optimal" approach). Demands arrive dynamically and are served one-at-a-time in the order of their arrival. Performance of the strategies is compared using the NSFNET topology (Fig. 1) with 100 random unit demands. The protection requirement, q, for each demand has a truncated normal distribution with a standard deviation $\sigma = \frac{1}{2}$. The mean of q is varied between 0 and 1 for each iteration.

The costs to route the demand and protection capacity are plotted in Fig. 3 as a function of the expected value of q. It is seen that the partial protection scheme offers significant savings over $1 + 1$ routing over a wide range of q, with the LP achieving even greater gains than $1 + q$ protection because of its use of flow bifurcation. With dynamic arrivals, we find that DSPP performs *better* than the greedy schemes that jointly optimize the primary and backup paths. The greedy optimal approach of jointly optimizing the primary and backup routes will often take a longer primary path for an incoming demand in order to take advantage of backup sharing. The longer primary path makes it more difficult for future demands to find disjoint primary routes, thus lowering their ability to share protection resources. While other works have focused on finding heuristics to jointly optimize the primary and backup paths for each incoming demand, it appears to be a better approach to simply take the shortest path for the primary route.

3. Conclusion

In this paper, a novel protection scheme to provide shared partial protection for a multi-commodity setting was introduced. Simulations show that these protection schemes offer significant savings over the most common protection schemes used today. A reduced complexity algorithm (DSPP) was developed, and for dynamic arrivals, DSPP actually performs better than jointly optimizing the primary and backup paths for each incoming demand.

References

[1] C. Ou, J. Zhang, H. Zang, L. Sahasrabuddhe, and B. Mukherjee, "New and Improved Approaches for Shared-Path Protection in WDM Mesh Networks," Journal of Lightwave Technology **22** (2004).

[2] W. Grover, *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET, and ATM Networking* (Prentice Hall, 2004).

[3] G. Mohan, S. Murthy, and A. Somani, "Efficient algorithms for routing dependable connections in WDM optical networks," Networking, IEEE/ACM Transactions on **9**, 553–566 (2002).

[4] O. Gerstel and G. Sasaki, "Quality of Protection (QoP): a Quantitative Unifying Paradigm to Protection Service Grades," Optical Networks Magazine **3**, 40–49 (2002).

[5] G. Kuperman, E. Modiano, and A. Narula-Tam, "Analysis and algorithms for partial protection in mesh networks," in "INFOCOM, 2011 Proceedings IEEE," (IEEE, 2011), pp. 516–520.

[6] A. Das, C. Martel, and B. Mukherjee, "A Partial-Protection Approach Using Multipath Provisioning," in "ICC '09. IEEE International Conference on Communications. Proceedings," (2009).

[7] G. Kuperman, E. Modiano, and A. Narula-Tam, "Partial Protection in Networks with Backup Capacity Sharing," Tech. rep., MIT, http://web.mit.edu/gregk/www/papers/PPNBCS.pdf (2011).