Analysis and Algorithms for Partial Protection in Mesh Networks

Greg Kuperman
MIT LIDS
Cambridge, MA 02139
gregk@mit.edu

Eytan Modiano
MIT LIDS
Cambridge, MA 02139
modiano@mit.edu

Aradhana Narula-Tam
MIT Lincoln Laboratory
Lexington, MA 02420
arad@ll.mit.edu

Abstract—This paper develops a mesh network protection scheme that guarantees a quantifiable minimum grade of service upon a failure within a network. The scheme guarantees that a fraction q of each demand remains after any single link failure. A linear program is developed to find the minimum-cost capacity allocation to meet both demand and protection requirements. For $q \leq \frac{1}{2}$, an exact algorithmic solution for the optimal routing and allocation is developed using multiple shortest paths. For $q > \frac{1}{2}$, a heuristic algorithm based on disjoint path routing is developed that performs, on average, within 1.4% of optimal, and runs four orders of magnitude faster than the minimum-cost solution achieved via the linear program. Moreover, the partial protection strategies developed achieve reductions of up to 82% over traditional full protection schemes.

I. INTRODUCTION

Mesh networks supporting data rates of multiple gigabytes per second are being deployed to meet the increasing demands of the telecom industry [1]. As data rates continue to increase, the failure of a network line element or worse, a fiber cut, can result in severe service disruptions and large data loss, potentially causing millions of dollars in lost revenue [2]. Currently, there exist few options for protection that offer less than complete restoration after a failure. Due to the cost of providing full protection, many service providers offer no protection whatsoever. By defining varying and quantifiable grades of protection, service providers can protect essential traffic without incurring the cost of providing full protection, making protection more affordable and better suited to user/application requirements. The protection scheme developed in this paper provides "partial protection" guarantees, at a fraction of the cost of full protection, with each session having its own differentiated protection guarantee.

There are a variety of protection strategies available [3], [4], [5]. The most common in backbone networks today is guaranteed path protection [6], which provides an edge-disjoint backup path for each working path, resulting in 100% service restoration after any link failure. Best effort protection is still loosely defined, but generally it offers no *guarantees* on the amount of protection provided. A service will be protected, if

possible, with any unused spare capacity after fully protecting all guaranteed services [2]. Best effort protection can also be referred to as partial capacity restoration, since a service will be restored within existing spare capactiy, typically resulting in less than 100% restoration.

Many users may be willing to tolerate short periods of reduced capacity if data rate guarantees can be made at a reduced cost, especially since link failures are relatively uncommon and are on average repaired quickly [2]. 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. If provided at a reduced cost, many users may opt for partial protection guarantees during network outages.

A quantitative framework for deterministic partial protection in optical networks was first developed in [7]. In this work, a minimum fraction q of the demand is guaranteed to remain available between the source and destination after any single link failure, where q is between 0 and 1. When q is equal to 1, the service is fully protected, and when q is 0, the service is unprotected. More recently, [8] examines the savings that can be achieved by guaranteeing part of the demand in the event of a link failure, as opposed to full protection. It shows that the amount of protection that can be guaranteed depends on the topology of the network. In [9], the partial protection problem on groomed optical WDM networks is studied, under the assumption that flows must traverse a single path.

In this work, we further expand upon the framework developed in [7] and [8]. We develop a "theory" for partial protection that includes optimal algorithms for capacity allocation, and explicit expressions for the amount of required spare capacity. Routing strategies that allocate working and spare capacity to meet partial protection requirements are derived. Similar to [8], flow bifurcation over multiple paths is allowed. Bifurcation reduces the amount of spare capacity needed to support the QoP requirements. In fact, we show that depending on the value of q, it may be possible to provide protection without any spare capacity.

We develop a linear program to find the optimal minimum-cost capacity allocation needed to guarantee partial protection in the event of a link failure. Furthermore, a routing and capacity assignment strategy based on shortest paths is shown to be optimal for $q \leq \frac{1}{2}$. For $q > \frac{1}{2}$, a reduced complexity algorithm based on disjoint path routing is shown to have a cost

that is at most twice the optimal minimum-cost solution, and in practice only slightly above optimal. Simulations over many random network topologies show that this disjoint path routing algorithm performs on average within 1.4% of the minimum-cost solution and leads to as much as 82% savings as compared to traditional full protection schemes.

In Section II, the partial protection model is described. In Section III, the partial protection problem is formulated as a linear program with the objective of finding the minimum-cost allocation of working and backup capacity. In Section IV, a simple path based solution for $q \leq \frac{1}{2}$ is developed. In Section V, properties of the minimum-cost solution for $q > \frac{1}{2}$ are determined and used to develop a time-efficient heuristic algorithm. The results of the algorithm are compared to the optimal solution and to traditional protection schemes. All proofs are omitted for brevity and can be found in [10].

II. Partial Protection Model

The objective of partial protection is to find an allocation that ensures that enough capacity exists to support the full demand before a link failure and a fraction q of the demand afterward. We assume that the graph G, with a set of vertices V and edges E, is at least two-connected. Each link has a fixed cost of use: c_{ij} for each edge $\{i, j\} \in E$. We consider only single link failures. We assume that demands do not share protection capacity with one another. Both working traffic and protection flows (defined as the flow after a failure) can be bifurcated to traverse multiple paths between the source and destination. After the failure of a link, a network management algorithm reroutes the traffic along the allocated protection paths. Without loss of generality, we assume unit demands, unless noted otherwise.

For now assume that link costs are all 1; in the next section we will consider non-uniform link costs. With uniform link costs, the objective is to minimize the total capacity needed to support the flow and the partial protection requirements.

(a) $1 + 1$ protection

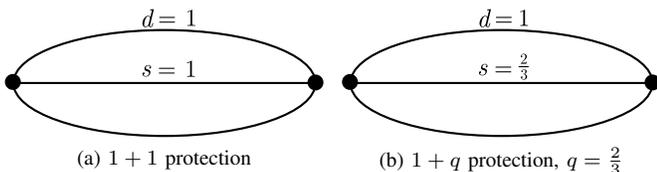(b) $1 + q$ protection, $q = \frac{2}{3}$

Fig. 1: Standard protection schemes

One routing strategy for providing this backup capacity is to use a single primary and a single backup path similar to the $1+1$ guaranteed path protection scheme. Consider the network shown in Figure 1. With $1+1$ protection, one unit of capacity is routed on a primary path and one unit of capacity on a backup, as shown in Figure 1a. Upon a link failure, 100% of the service can be restored on the backup path. Now, consider a partial protection requirement to provide a fraction $q = \frac{2}{3}$ of backup capacity in the event of a link failure. A naïve protection scheme similar to $1+1$ protection would be to route one unit along the primary path and $\frac{2}{3}$ along a disjoint protection path, as shown in Figure 1b. This simple protection scheme will be referred to

as $1 + q$ protection. If the primary path fails, sufficient backup capacity remains to provide service for $\frac{2}{3}$ of the demand.
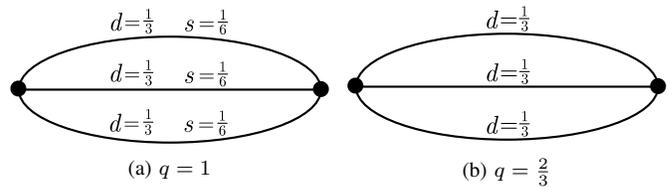

(a) $q = 1$

(b) $q = \frac{2}{3}$

Fig. 2: Protection using risk distribution

For both partial and full protection requirements, in many cases capacity savings can be achieved if the risk is distributed by spreading the primary allocation across multiple paths. For example, by spreading the primary allocation across the three available paths, as shown in Figure 2a, any single link failure results in a loss of at most $\frac{1}{3}$ of the demand. To fully protect this demand against any single link failure (i.e. $q = 1$), additional spare allocation of $s = \frac{1}{6}$ needs to be added to each link. With this strategy, a total of 1.5 units of capacity are required, as opposed to the total of 2 units needed by $1 + 1$ protection. If instead the protection requirement was $q = \frac{2}{3}$, no spare allocation is needed since after any failure $\frac{2}{3}$ units are guaranteed to remain. By spreading the primary and backup allocation across the multiple paths between the source and destination, the risk is effectively distributed and the fraction of primary allocation lost by a link failure is reduced.

III. Minimum-Cost Partial Protection

In this section, a linear program is developed to achieve an optimal minimum-cost solution to the partial protection problem. The objective of the linear program is to find a minimum-cost routing strategy to meet demand d and partial protection requirement q between two nodes s and t. In particular, the full demand must be met before any failure, and in the event of any link failure, a fraction q of that demand must remain. The linear program to solve for the optimal routing strategy, denoted LP_{PP}, is defined below.

A. Linear Program to Meet Partial Protection: LP_{PP}

The following values are given:
- $G = (V, E)$ is the graph with its set of vertices and edges
- (s, t) is the source and destination, respectively
- d is the total demand between the source and destination
- q is the fraction of the demand that must be supported on the event of a link failure
- c_{ij} is the cost of link $\{i, j\}$

The LP solves for the following variables:
- w_{ij} is the working flow assigned on link $\{i, j\}$, $w_{ij} \geq 0$
- s_{ij} is the spare allocation assigned on link $\{i, j\}$, $s_{ij} \geq 0$
- f_{kl}^{ij} is the protection flow assigned on link $\{i, j\}$ after the failure of link $\{k, l\}$, ≥ 0

The objective of LP_{PP} is to:
- Minimize the cost of allocation over all links:

$$\min \sum_{\{i,j\} \in E} c_{ij}(w_{ij} + s_{ij}) \tag{1}$$

Subject to the following constraints:

- Route working traffic between s and t to meet demand d:

$$\sum_{\{i,j\}\in E} w_{ij} - \sum_{\{j,i\}\in E} w_{ji} = \begin{cases} d & \text{if } i = s \\ -d & \text{if } i = t \\ 0 & \text{o.w.} \end{cases}, \ \forall i \in V \quad (2)$$

- Route flow to meet partial protection requirement q after failure of link $\{k,l\}$ between s and t:

$$\sum_{\substack{\{i,j\}\in E \\ \{i,j\}\neq\{k,l\}}} f_{kl}^{ij} - \sum_{\substack{\{j,i\}\in E \\ \{j,i\}\neq\{k,l\}}} f_{kl}^{ji} = \begin{cases} dq & \text{if } i = s \\ -dq & \text{if } i = t \\ 0 & \text{o.w.} \end{cases},$$
$$\forall i \in V, \ \forall\{k,l\} \in E \quad (3)$$

- Working and spare capacity assigned on link $\{i,j\}$ meets partial protection requirements after failure of link $\{k,l\}$:

$$f_{kl}^{ij} \le w_{ij} + s_{ij}, \quad \substack{\forall\{i,j\}\in E \\ \forall\{k,l\}\in E} \quad (4)$$

A minimum-cost solution will provide a flow to meet the demand before a link failure and a flow to meet the partial protection requirement after any single-link failure. As we allow bifurcation, each of these flows may be routed over multiple paths. An interesting characteristic of the optimal solution given by the linear program is that, at each node, flow conservation for the working flow is maintained, but the total allocation for working plus spare capacity, given by $(w_{ij} + s_{ij})$ for edge $\{i,j\}$, does not necessarily maintain flow conservation.

B. Comparison to Standard Protection Schemes

We compare the optimal solution computed by the above linear program to the standard scheme of $1+1$ protection, as well as $1+q$ protection, on 1000 random graph topologies, each containing 50 nodes with an average node degree of 3.1 and random link costs. Two nodes were randomly chosen from each graph to be the source and destination. The minimum-cost allocation for values of q between 0 and 1 was determined by the linear program using CPLEX. The $1+1$ and $1+q$ protection schemes were solved using the Bhandari algorithm for shortest pair of disjoint paths [11].

The average cost to route the demand and protection capacity using the different routing strategies are plotted in Figure 3 as a function of q. The top line, showing capacity requirements under $1+1$ protection, remains constant for all values of q. The next two lines from the top are $1+q$ and LP_{PP}, respectively. As expected, both meet demand and protection requirements using fewer resources than $1+1$, however, the minimum-cost solution produced by the partial protection linear program uses significantly less capacity. A lower bound on the capacity requirement is the minimum-cost routing that provides no protection ($q = 0$), shown in the bottom line of the figure. The cost of providing partial protection q is the difference between the cost of the respective protection strategies and the minimum-cost routing with no protection. Partial protection
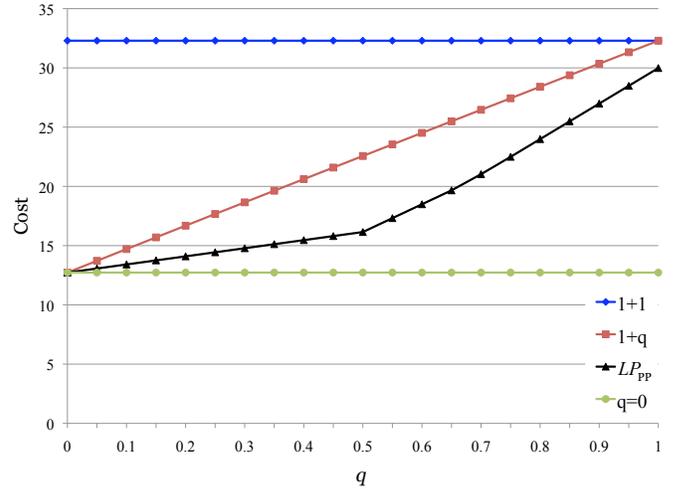

Fig. 3: Capacity Cost vs. q

achieves reductions in excess resources of 82% at $q = \frac{1}{2}$ to 12% at $q = 1$ over $1+1$ protection, and 65% at $q = \frac{1}{2}$ to 12% at $q = 1$ over $1+q$ protection.

IV. SOLUTION FOR $q \le \frac{1}{2}$

In this section we provide insights on the structure of the solution to the minimum-cost partial protection problem on general mesh networks. When $q \le \frac{1}{2}$, we are able to derive an exact algorithmic solution to the partial protection problem. We show that all minimum-cost solutions for $q \le \frac{1}{2}$ will never need spare allocation, allowing us to formulate the partial protection problem using standard network flow conservation constraints. A simple path-based algorithmic solution is then derived. The difficulty in obtaining further insights into the optimal solution for $q > \frac{1}{2}$ stems from the fact that, as mentioned in Section II, the total working and spare allocation does not necessarily meet flow conservation requirements at each node. Without this property, most network flow algorithms do not apply [12] and analysis of the linear program becomes difficult.

Lemma 1 demonstrates that spare capacity is not needed if and only if the working capacity on an edge is less than or equal to $(1-q)$, because that means that any time a link is lost, at least q remains in the network.

Lemma 1. *Given a partial protection requirement q between nodes s and t, the spare capacity needed to satisfy demand and protection requirements is zero if and only if the working capacity on each link is $w_{ij} \le (1-q)$, $\forall\{i,j\} \in E$.*

In Section V, we show routings with zero spare allocation are not necessarily lowest cost for all values of q. However, Lemma 2 shows that when $q \le \frac{1}{2}$, the minimum-cost solution will never use spare allocation.

Lemma 2. *Given a demand between nodes s and t with protection requirement $q \le \frac{1}{2}$, all minimum-cost solutions use no spare capacity: $s_{ij} = 0$, $\forall\{i,j\} \in E$.*

Combining Lemmas 1 and 2, it can be seen that an optimal solution exists that does not use any spare allocation for $q \le \frac{1}{2}$

with $w_{ij} \leq (1 - q)$, $\forall \{i, j\} \in E$. For the case of $q \leq \frac{1}{2}$, no spare allocation is needed and flow conservation constraints are met. Therefore, the linear program can be written using a standard flow formulation without using spare allocation, maintaining flow conservation at each node. The modified linear program, referred to as $LP_{q \leq .5}$, routes the flows on the paths in a manner that minimizes total cost and ensures that no edge carries more than $(1 - q)$.

$$LP_{q \leq .5} : \min \sum_{\{i,j\} \in E} c_{ij} w_{ij} \tag{5}$$

$$\sum_{\{i,j\} \in E} w_{ij} - \sum_{\{j,i\} \in E} w_{ji} = \begin{cases} 1 & \text{if } i = s \\ -1 & \text{if } i = t \\ 0 & \text{o.w.} \end{cases}, \forall i \in V \tag{6}$$

$$w_{ij} \leq (1 - q), \quad \forall \{i, j\} \in E \tag{7}$$

The above linear program achieves a minimum-cost routing in a network by using only working allocation to meet the demand. $LP_{q \leq .5}$ is a network flow problem with directed and capacitated edges, which is recognized as a minimum-cost flow problem [12], for which algorithmic methods exist for finding an optimal solution. In Theorem 1, we show that an optimal solution for $q \leq \frac{1}{2}$ uses at most three paths with allocation q on each of the shortest pair of disjoint paths and allocation $(1 - 2q)$ on the shortest path.[1]

Consider a directed graph $G = (V, E)$ with a source s and destination t. Let p_0 be the cost of the shortest path, p_1 and p_2 be the costs of the two shortest pair of disjoint paths, f_0 be the flow on the shortest path, f_1 and f_2 be the flows on each of the two shortest pair of disjoint paths, respectively, and $\mathcal{T}_{st}(q)$ be the cost of the allocation needed to meet demand and protection requirements between s and t for a value of q.

Theorem 1. *Given a source s and destination t in a two-connected directed network $G = (V, E)$ with $q \leq \frac{1}{2}$, there exists an optimal solution meeting working and partial protection requirements with $f_0 = (1 - 2q)$ and $f_1 = f_2 = q$, giving a total cost $\mathcal{T}_{st}(q) = (1 - 2q)p_0 + q(p_1 + p_2)$, where path 0 is the shortest path and paths 1 and 2 are the shortest pair of disjoint paths.*

V. SOLUTION FOR $q > \frac{1}{2}$

When $q \leq \frac{1}{2}$, no spare allocation is needed and the minimum-cost routing to meet partial protection requirements can be found for any general mesh network. When $q > \frac{1}{2}$, it may be necessary to use spare allocation to meet protection requirements. Since the overall allocation of working plus spare does not necessarily meet flow conservation at any particular node, it may not be possible to provide a simple flow-based description of the optimal solution on general mesh networks. If we consider N disjoint paths between the source and destination, with the i^{th} path having cost p_i, we see that this is

---

[1]It is possible that the shortest path is one of the pair of disjoint paths, in which case $f_0 = (1 - q)$ and $f_2 = q$.

equivalent to a two-node network with N links where the i^{th} link has cost p_i. Hence, we start by investigating the properties of minimum-cost solutions for two-node networks in order to gain insight on solutions for general networks. Using these insights, a heuristic algorithm is developed in Section V-B for general mesh networks.

A. Results for Two-Node Networks

A two-node network is defined as having a source and destination node with N links between them. Each link has a fixed cost of use, c_i. We first note that a solution that uses no spare allocation is not necessarily a minimum-cost allocation when unequal link costs are considered. Consider the example in Figure 4 and let $q = \frac{2}{3}$. Allocating a capacity of $\frac{1}{3}$ onto each link does not use any spare capacity and has total cost of $\frac{1}{3}(1 + 2 + 6) = 3$. In contrast, consider using the two lowest cost links, each with allocation $\frac{2}{3}$. Clearly, the protection requirement is met, and the total cost is reduced to $\frac{2}{3}(1+2) = 2$, which is less than the cost of the zero spare capacity allocation.
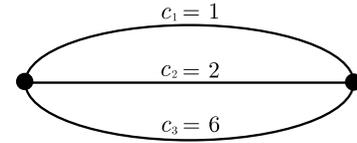

Fig. 4: Two-node network with link costs

For two-node networks, order the edges such that $c_1 \leq c_2 \leq ... \leq c_N$. Define x_i as the allocation on the i^{th} edge. From our analysis (see [10]), we are able to define a value K, which will be important for evaluating two-node networks: $K = argmax_{K=2..N}(c_K \leq \frac{1}{K-1} \sum_{i=1}^{K} c_i)$. K is the maximum number of links such that the incremental cost of using an additional link would not improve the solution.

For $q > \frac{1}{2}$, spare allocation may or may not be needed. Lemma 3 shows when spare allocation is necessary.

Lemma 3. *A minimum-cost allocation for a two-node network uses spare allocation if and only if $q > \frac{K-1}{K}$.*

Next, the exact edge allocations for a minimum-cost solution to meet partial protection requirements on a two-node network are defined. Lemma 4 states that when spare allocation is needed, edges 1 to K will have an equal allocation, and edges $K + 1$ to N will have no allocation. Lemma 5 shows that when spare allocation is not needed, the solution will use $J \leq K$ edges. Lemma 5 also provides the allocations across the J edges.

Lemma 4. *A minimum-cost allocation when $q > \frac{K-1}{K}$ will be an even allocation of $q \frac{1}{K-1}$ on the K lowest cost edges, and no allocation on the remaining edges.*

Lemma 5. *The minimum-cost allocation when $q \leq \frac{K-1}{K}$ will have non-zero allocation on edges 1 to J, where J is the integer satisfying $\frac{J-2}{J-1} < q \leq \frac{J-1}{J}$. Moreover, the minimum-cost allocation when $q \leq \frac{K-1}{K}$ is: $x_i = (1-q)$, $\forall i = 1..(J-1)$; $x_J = (J - 1)q - (J - 2)$; $x_i = 0$, $\forall i = (J + 1)..N$.*

B. Time-Efficient Heuristic Algorithm

Consider a mesh network with N disjoint paths between the source and destination, and let p_i be the cost of the i^{th} path. As discussed in the beginning of the section, by treating the N paths in the general mesh network as a two-node network with N links, the results from Section V-A can be applied to develop a heuristic solution for general mesh networks for the case of $q > \frac{1}{2}$. Recall that for $q \leq \frac{1}{2}$, the optimal minimum-cost solution for general mesh networks was derived in Section IV.

The heuristic algorithm is based on finding the k-shortest edge-disjoint paths for $k = 2$ to $k = N$, where N is the maximum number of edge-disjoint paths and the length of each path is its cost. These paths can be found using the Bhandari algorithm [11]. For each set of k disjoint paths, we look to see if spare allocation is needed, i.e. $q > \frac{k-1}{k}$, and use the minimum-cost allocation given by Lemmas 4 and 5. From the different possible disjoint path routings, the allocation of minimum-cost is chosen. We call this algorithm the Partial Protection Disjoint Path Routing Algorithm (PP-DPRA). Theorem 2 gives a bound on PP-DPRA's performance.

Theorem 2. *PP-DPRA produces a routing with a cost that is at most twice the optimal minimum-cost.*

C. Comparison of PP-DPRA to the Minimum-Cost Solution

The PP-DPRA solution is compared to $1 + 1$, $1 + q$, and LP_{PP}. The simulation is identical to the one run in Section III-B, with PP-DPRA being implemented in C. The average costs to meet demand and protection requirements over all random graphs are plotted in Figure 5. Simulation results show that for $q \leq \frac{1}{2}$, the routing as given by Theorem 1 matches the optimal routing produced by LP_{PP}, and for $q > \frac{1}{2}$, the average cost is greater than optimal by 1.4% on average. Additionally, on average, the running time for routing a demand with PP-DPRA was 0.001 seconds, while with the linear program LP_{PP} it was 22 seconds. This reduction in running time of four orders of magnitude makes the algorithm suitable for networks that require rapid setup times for incoming demands.

We note that [8] also developed an algorithm for meeting partial protection requirements by spreading capacity across disjoint paths. However, the algorithm in [8] was designed to minimize capacity over multiple connections, whereas the algorithms in this paper were designed to minimize costs for one connection at a time, making a direct comparison of the algorithms difficult.

VI. CONCLUSION AND FUTURE DIRECTIONS

In this paper we have developed a mathematical model to provide deterministic partial protection for a single commodity. A linear program was formulated to find a minimum-cost solution in general mesh networks. Simulations show this LP offers significant savings over the most common protection schemes used today. A heuristic algorithm, PP-DPRA, was developed. Simulation results show that this algorithm comes within 1.4% of optimal on average and runs four orders of magnitude faster than the linear program.
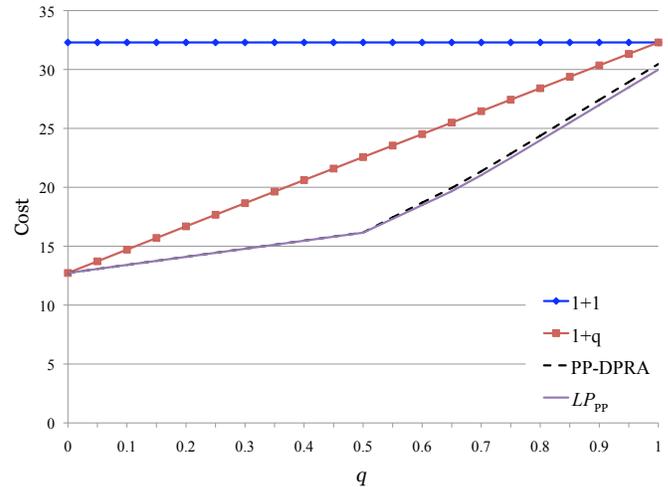

Fig. 5: Cost vs. q for various protection schemes

An important direction for future research will be to consider the additional savings in protection capacity that can be achieved by resource sharing between demands. Currently, resources that could be used to protect multiple demands are potentially being underutilized by being dedicated to only one demand. Preliminary results show that significant savings can be achieved with protection resource sharing. A better understanding of how resources are shared will help develop more efficient algorithms for the partial protection problem.

REFERENCES

[1] B. Mukherjee, "WDM Optical Communication Networks: Progress and challenges," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1810–1824, 2000.
[2] W. Grover, *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET, and ATM Networking*. Prentice Hall, 2004.
[3] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM Mesh Networks," *Journal of Lightwave Technology*, vol. 21, no. 4, p. 870, 2003.
[4] R. Iraschko, M. MacGregor, and W. Grover, "Optimal Capacity Placement for Path Restoration in STM or ATM Mesh-Survivable Networks," *IEEE/ACM Transactions on Networking*, vol. 6, no. 3, pp. 325–336, 1998.
[5] W. Yao and B. Ramamurthy, "Survivable Traffic Grooming with Path Protection at the Connection Level in WDM Mesh Networks," *Journal of Lightwave Technology*, vol. 23, no. 10, p. 2846, 2005.
[6] A. Saleh and J. Simmons, "Evolution Toward the Next-Generation Core Optical Network," *Journal of Lightwave Technology*, vol. 24, no. 9, p. 3303, 2006.
[7] O. Gerstel and G. Sasaki, "Quality of Protection (QoP): a Quantitative Unifying Paradigm to Protection Service Grades," *Optical Networks Magazine*, vol. 3, no. 3, pp. 40–49, 2002.
[8] A. Das, C. Martel, and B. Mukherjee, "A Partial-Protection Approach Using Multipath Provisioning," in *ICC '09. IEEE International Conference on Communications. Proceedings*, 2009.
[9] J. Fang and A. Somani, "On Partial Protection in Groomed Optical WDM Mesh Networks," in *Proceedings of the 2005 International Conference on Dependable Systems and Networks*, 2005, p. 237.
[10] G. Kuperman, E. Modiano, and A. Narula-Tam, "Analysis and Algorithms for Partial Protection in Mesh Networks," MIT, Tech. Rep., 2010. [Online]. Available: http://web.mit.edu/~gregk/www/papers/AAPPMN-TR2010.pdf
[11] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Kluwer Academic Pub, 1999.
[12] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, New Jersey, 1993.