

Network Protection with Multiple Availability Guarantees

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—We develop a novel network protection scheme that provides guarantees on both the fraction of time a flow has full connectivity, as well as a quantifiable minimum grade of service during downtimes. In particular, a flow can be below the full demand for at most a maximum fraction of time; then, it must still support at least a fraction q of the full demand. This is in contrast to current protection schemes that offer either availability-guarantees with no bandwidth guarantees during the downtime, or full protection schemes that offer 100% availability after a single link failure. We develop algorithms that provide multiple availability guarantees and show that significant capacity savings can be achieved as compared to full protection. If a connection is allowed to drop to 50% of its bandwidth for 1 out of every 20 failures, then a 24% reduction in spare capacity can be achieved over traditional full protection schemes. In addition, for the case of $q = 0$, corresponding to the standard availability constraint, an optimal pseudo-polynomial time algorithm is presented.

I. INTRODUCTION

As data rates continue to rise, a network failure can cause catastrophic service disruptions. To protect against such failures, networks typically use full protection schemes, often doubling the cost of resources needed to route a connection. An alternative approach is to provide a guarantee on the maximum time a connection can be disrupted. This is known as an “availability guarantee”, and it is a bound on the fraction of time or probability that a connection can be disrupted. However, these disruptions (downtimes) may be unacceptably long; thus, many service providers opt for the more resource intensive full protection. In this paper, we propose a novel protection scheme with multiple availability guarantees. In addition to the traditional availability guaranteed protection, which maintains the full demand for at least a guaranteed fraction of time, we guarantee partial connectivity at all times. Thus, our approach is a hybrid between the traditional availability guarantees and full protection schemes.

Full protection schemes have been studied extensively [1–6]. The most common scheme in backbone networks today is 1 + 1 guaranteed path protection [7], which provides an edge-disjoint backup path for each working path, and guarantees the full demand to be available at all times after any single link

failure. There has also been a growing body of literature for backup provisioning to meet availability guarantees [8–14]. In all of these, primary and backup flows are allocated such that the connection is disrupted for at most a specified fraction of time or probability. During these down-states, the service is completely disrupted. A version of availability guarantees is considered in [15], where an end-to-end flow having a certain expected capacity, based on link availabilities, is found. In our paper, a flow is guaranteed to be at least a fraction q of the full demand at all times, which is known as “partial protection”. Our novel approach is the first to combine the traditional availability guarantee and partial protection guarantee to allow the user to specify flows with different availability guarantees. Moreover, it is particularly applicable to IP-over-WDM networks where MPLS tunnels are used to provision resources.

The partial protection framework was first developed in [16]. More recently, [17, 18] developed a “theory” of partial protection such that after *any* single link failure, the flow can drop to the partial protection requirement. In [17, 18], a fraction q of the demand is guaranteed to remain available between the source and destination after any failure, where q is between 0 and 1. When q is equal to 1, the service will have no disruptions after any failure, and when q is 0, there will be no flow between the two nodes during the down state. In our work, flows can drop below the full demand for at most a specified fraction of time, and maintain at least q of that demand at all times. Similar to [10–14], the probability of simultaneous failures is assumed to be negligible and we only consider single-link failures.

The novel contributions of this paper include a framework for Multiple Availability Guaranteed Protection (MAGP) and providing associated algorithms for allocating resources. Moreover, in the $q = 0$ case, corresponding to the previously studied scenario where full availability is guaranteed for a fraction of time, we develop an optimal pseudo-polynomial algorithm.

This paper is outlined as follows. In Section II, the model for MAGP is described. In Section III, MAGP is shown to be NP-Hard, and the minimum-cost solution to MAGP is formulated as an MILP. In Section IV, optimal solutions and algorithms for MAGP are developed.

II. MULTIPLE AVAILABILITY GUARANTEED PROTECTION

In this paper, routing strategies are developed and analyzed to minimize the total cost and capacity allocation required to

This work was supported by NSF grants CNS-1116209 and CNS-0830961, by DTRA grants HDTRA1-07-1-0004 and HDTRA-09-1-005, and by the Department of the Air Force under Air Force contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

satisfy each demand's protection and availability requirements. A demand needs to be routed from its source s to destination t such that the flow can drop below the full demand for at most some specified downtime, and maintain at least a fraction q of that full demand at all times. To simplify the analysis, a "snapshot" model is used: the network state is considered after a failure has occurred. Let p_{ij} be the conditional probability that edge $\{i, j\}$ failed given that a network failure has occurred. For simplicity of exposition, instead of a maximum downtime, the Maximum Failure Probability (MFP) is considered, and its value is denoted by P . After some network failure, the flow can be below the full demand, but at least a fraction q of the demand, with at most probability P . The maximum failure probability is the conditional probability that the network is in a downstate given some link disruption, and can be related to the maximum downtime by accounting for expected time between failures and mean time to repair (e.g., see [15]).

We assume that the graph G , with a set of vertices V , edges E , and edge failure probabilities \mathcal{P} , is at least two-connected. Since only single-link failures are considered, edge failures are disjoint events, which gives $\sum_{\{i,j\} \in E} p_{ij} = 1$. Similar to previous works, the primary flow is restricted to a single path. After the failure of a link, a network management algorithm reroutes the traffic along the allocated protection paths. Due to space constraints, in this paper we do not consider backup resource sharing. A full treatment of the MAGP problem with resource sharing can be found in [19]. Without loss of generality, for the remainder of this paper unit demands are assumed.

Consider the example in Fig. 1, with link failure probabilities and flow allocations as labeled (p and f respectively). A unit demand needs to be routed from s to t with $P = \frac{1}{4}$ and partial protection requirement q . In [17], a simple partial protection scheme called $1 + q$ protection was developed, which routes the primary demand on one path and the partial protection requirement onto another edge-disjoint path. After any failure along the primary path, the partial protection requirement is met. This is shown in Fig. 1a with the solid line carrying the primary flow of 1 and the dotted line carrying the protection flow of q . However, in this example the maximum failure probability is exceeded for the $1 + q$ routing: after a failure, the flow drops below the unit demand between s and t with a probability of $\frac{1}{2}$ (because the failure of either of the primary links would drop the demand below its full capacity). A naive alternative would be to simply allocate another path for protection, which would be identical to the $1 + 1$ full protection scheme (shown in Fig. 1b), and utilize 4 units of capacity. After any failure, the full demand of 1 is maintained; thus, the user will face no downtime, which meets and exceeds the maximum probability of failure requirement of $\frac{1}{4}$.

If we allow different levels of protection on different segments of the primary path, then a more resource efficient allocation is possible. Consider keeping the primary flow on the same bottom two edges as before, but instead of allocating an end-to-end backup path along the top two edges, 1 unit of flow

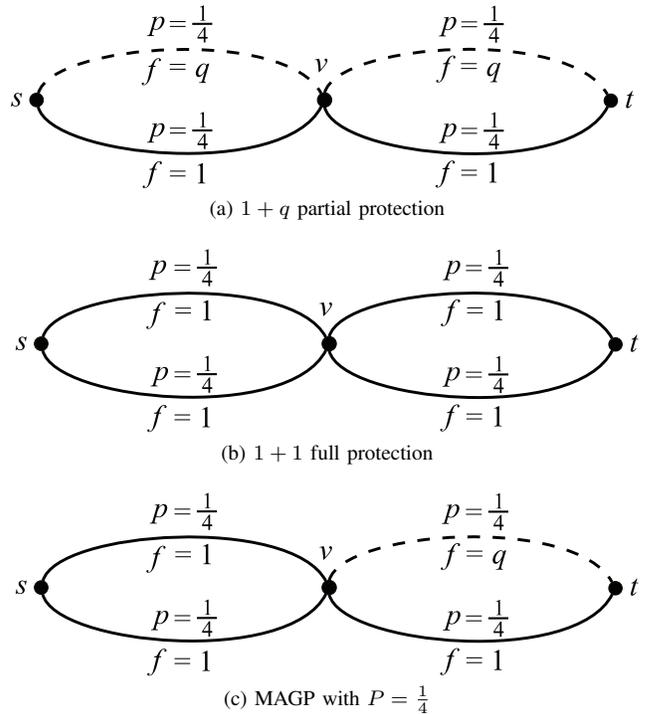


Fig. 1: Comparison of MAGP and traditional protection schemes

is allocated to protect against the failure of $\{s, v\}$ and q units of flow to protect against the failure of $\{v, t\}$ (shown in Fig. 1c). If after some disruption either of the $\{s, v\}$ edges fail, 1 unit of flow will still remain from s to t . By fully protecting the primary $\{s, v\}$ edge, there is zero probability that its failure will cause the flow to drop below the full demand. The probability that the flow will drop below 1 after some failure is $\frac{1}{4}$, which meets the MFP requirement. This routing only needs $3 + q$ units of capacity, as opposed to the 4 units that full protection requires.

III. MINIMUM-COST MULTIPLE AVAILABILITY GUARANTEED PROTECTION

This section investigates minimum-cost allocations for multiple availability guaranteed protection. Each edge $\{i, j\}$ is assumed to have an associated cost c_{ij} . Section III-A begins by demonstrating the MAGP problem to be NP-Hard. Subsequently, an MILP is formulated to find a minimum-cost routing that meets a demand's partial protection and availability requirements. In Section III-B, MAGP is compared to the $1 + 1$ full protection scheme.

A. Mixed Integer Linear Program to Meet Multiple Availability Guaranteed Protection

For a connection request between two nodes s and t , the flow can drop to a fraction q of the demand with at most probability P . Again, the snapshot model is used, and the set of link failure probabilities \mathcal{P} are conditional given a network failure has occurred. Before proceeding with the MILP, we determine the complexity of MAGP to be NP-Hard.

Theorem 1. *Minimum-cost Multiple Availability Guaranteed Protection is NP-Hard.*

Proof: To demonstrate NP-Hardness of MAGP, a reduction from the 1-0 knapsack problem [20] is performed. The proof is omitted for brevity, and can be found in [19]. ■

For the MILP, the following values are given:

- $G = (V, E, C, \mathcal{P})$ is the graph with its set of vertices, edges, costs, and edge failure probabilities
- q is the fraction of the demand between s and t that must be supported in the event of a link failure
- c_{ij} is the cost of link $\{i, j\}$
- p_{ij} is the probability that link $\{i, j\}$ has failed given a network failure has occurred
- P is the maximum probability that the service falls below its full demand

The MILP solves for the following variables:

- x_{ij} is primary flow on link $\{i, j\}$, $x_{ij} \in \{0, 1\}$
- f_{kl}^{ij} is the protection flow on link $\{i, j\}$ after the failure of link $\{k, l\}$, $f_{kl}^{ij} \geq 0$
- z_{kl} is 1 if the failure of link $\{k, l\}$ causes the flow to drop below the primary demand of 1; 0 otherwise
- s_{ij} is total spare allocation on link $\{i, j\}$, $s_{ij} \geq 0$

The objective is to:

- Minimize the cost of allocation over all links.

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

Subject to the following constraints:

- Flow conservation constraints for primary flow: route primary traffic to meet demand.

$$\sum_{\{i,j\} \in E} x_{ij} - \sum_{\{j,i\} \in E} x_{ji} = \begin{cases} 1 & \text{if } i = s \\ -1 & \text{if } i = t, \forall i \in V \\ 0 & \text{o.w.} \end{cases} \quad (2)$$

- Full demand availability constraint: The probability that the flow drops below 1 after a failure is simply the sum of the failure probabilities of the individual edges causing the flow to drop below 1. The sum of these failure probabilities cannot exceed P .

$$\sum_{\{k,l\} \in E} p_{kl} z_{kl} \leq P \quad (3)$$

- Flow conservation constraints for partial service: if the failure of link $\{k, l\}$ causes the flow to drop below 1, route q from s to t ; otherwise, maintain a flow of 1. Let \mathcal{F}_{kl} be the expression $(1 - z_{kl}) + qz_{kl}$.

$$\sum_{\substack{\{i,j\} \in E \\ \{i,j\} \neq \{k,l\}}} f_{kl}^{ij} - \sum_{\substack{\{j,i\} \in E \\ \{j,i\} \neq \{k,l\}}} f_{kl}^{ji} = \begin{cases} \mathcal{F}_{kl} & \text{if } i = s \\ -\mathcal{F}_{kl} & \text{if } i = t, \\ 0 & \text{o.w.} \end{cases} \quad \forall i \in V, \forall \{k, l\} \in E \quad (4)$$

- Capacity allocation: primary and spare capacity assigned on link $\{i, j\}$ meets flow requirements after the failure of link $\{k, l\}$.

$$f_{kl}^{ij} \leq x_{ij} + s_{ij}, \quad \forall \{i,j\} \in E, \forall \{k,l\} \in E \quad (5)$$

A minimum-cost solution will provide an edge capacity allocation such that the flow between s and t drops to a fraction q of the demand with at most probability P .

B. Comparison to Full Protection

Multiple availability guaranteed protection is compared to the 1+1 full protection scheme via simulation. The performance of the strategies is compared using the NSFNET topology (Fig. 2) with 100 random unit demands. The protection requirement q is set to $\frac{1}{2}$ for all demands. All link costs are set to 1, and the probability of failure for any link is proportional to its length, which is reasonable since a longer fiber will have a higher likelihood of being accidentally cut. The maximum failure probability P is varied from 0 to .3 by .05 increments. While the main focus of this paper is the case where the primary flow is restricted to a single path, this simulation also considers allowing the primary flow to be bifurcated. Bifurcation reduces the loss of flow after any edge failure, thereby reducing the total allocation needed to meet requirements. Relaxing the integer constraint on the primary flow variables in the MILP corresponds to enabling bifurcation of the primary flow. Routing solutions for MAGP were determined using CPLEX to solve the MILP. The shortest pair of disjoint paths were used for 1 + 1 protection [21].

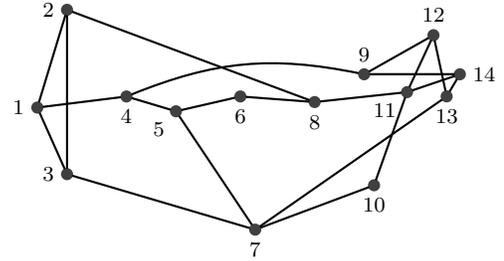


Fig. 2: 14 Node NSFNET backbone network

The average cost to route the demand and protection capacity using the different routing strategies is plotted in Fig. 3 as a function of the maximum failure probability P . Shortest path routing without protection considerations is used as a lower bound. The cost of providing incremental protection with parameters q and P is the difference between the cost of the respective protection strategies and shortest path routing.

Note that allowing the primary flow to bifurcate allows requirements to be met using a lower cost allocation. This is because splitting the primary flow distributes the risk so that upon an edge failure, less primary flow is disrupted, which then requires less protection resources. If the flow is allowed to drop to $\frac{1}{2}$ for 1 out of 20 failures (5% of the time), then a savings of 24% in protection capacity is realized for the case with bifurcation, and 17% without bifurcation. As the flow is allowed to drop more often to its partial protection

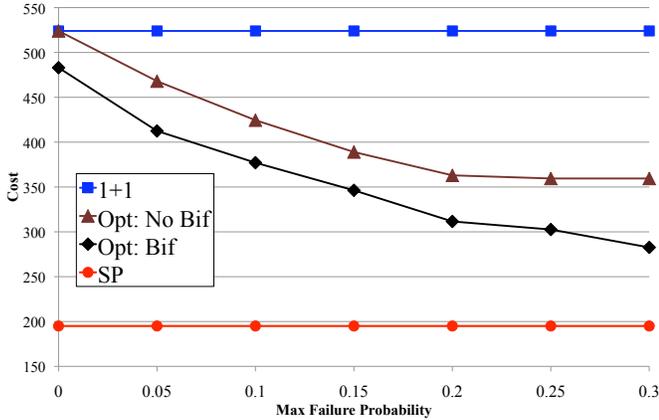


Fig. 3: Capacity cost vs. MFP with $q = \frac{1}{2}$

requirement after a failure, savings increase. For $P = 0.1$, a savings of 45% and 30% is seen for MAGP with and without bifurcation, respectively. For $P = 0.2$, the savings are 65% and 49%. Further increases in P result in only small additional savings; hence, the simulations were stopped at $P = 0.3$.

IV. OPTIMAL SOLUTION AND ALGORITHMS

While the MILP presented in the previous section finds an optimal solution to the multiple availability guaranteed protection problem, it is not a computationally efficient method of finding a solution, nor does it provide insight into why a solution is optimal. In this section, we analyze the MAGP problem to help develop efficient algorithms and heuristics for finding a minimum-cost routing. The MAGP problem requires identifying a primary path such that segments of it are protected in a way that after a link failure, the flow drops to q with probability of at most P .

The case of $q = 0$ is explored in Section IV-A. When $q = 0$, there is no partial protection requirement, so there is only a single availability guarantee. This is the traditional availability guarantee, which has been examined in previous works. An optimal pseudo-polynomial algorithm is presented to solve MAGP with $q = 0$, which to the best of our knowledge is the first such algorithm. In Section IV-B, the case of $q > 0$ is examined. We show that finding a feasible solution to the the closely related problem of Singly Constrained Shortest Pair of Disjoint Paths is strongly NP-Hard (there exists no pseudo-polynomial or ϵ -approximation algorithm), and conjecture that the MAGP problem with $q > 0$ is also strongly NP-Hard. Hence, a heuristic for solving MAGP with $q > 0$ is developed.

A. Availability Guarantees with $q = 0$

When $q = 0$, the partial protection requirement is removed and no flow is needed during the downtime. To solve this problem, a primary path needs to be found such that segments of it are protected, and after a failure, the flow can drop to 0 with probability of at most P . First, a restricted version of the problem is considered where we try to meet availability requirements without the use of spare allocation. It can be shown that the solution to the restricted problem is the constrained shortest

path (CSP) problem [22]. Next, the problem without restrictions on spare allocation is studied. We transform this unrestricted problem to an instance of the restricted one, and use CSP to find an optimal pseudo-polynomial algorithm for MAGP when $q = 0$.

1) *Availability Guarantees Without Spare Allocation:* First, we consider finding the lowest-cost path between s and t that meets the availability guarantee without the use of spare allocation. In other words, we want to find the lowest-cost path such that the sum of all the failure probabilities in that path are less than P . This problem is recognized to be the constrained shortest path problem (CSP) [22], which is NP-Hard. A dynamic program exists that finds the minimum-cost solution to CSP in pseudo-polynomial time [23], with a running time of $O(n^2P)$, where n is the number of nodes in the network; the P factor is what makes this running time pseudo-polynomial. CSP assumes all inputs are integer, so instead of the failure probabilities being between 0 and 1, we multiply P and all p_{ij} values, $\forall \{i, j\} \in E$, by the smallest factor that makes all the values integer (all inputs are assumed to be rational). Thus, for the remainder of this section, P and p_{ij} are assumed to be integer.

In general, a path may not exist from the source to the destination that can meet the availability requirement. Furthermore, if one exists, it is not necessarily of lowest cost. We next examine augmenting the flow with spare allocation to find a minimum-cost solution that meets requirements.

2) *Availability Guarantees With Spare Allocation:* We now examine allowing the use of spare allocation to protect segments of the primary path in a manner that ensures the entire end-to-end path meets availability guarantees. If a failure of a segment in the primary path does not cause a disruption in the end-to-end flow, then that segment is considered protected. A routing meeting guarantees will be a concatenation of protected and unprotected segments. Fig. 4 shows a sample solution for a unit demand between v_1 and v_6 with $P = 0.2$ that illustrates how the use of spare allocation enables meeting availability guarantees. The probabilities of link failure are as labeled, and all lines represent a unit flow.

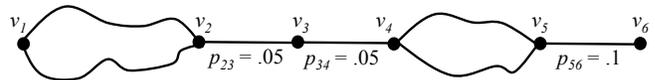


Fig. 4: Routing to meet $P = 0.2$ with $q = 0$ from v_1 to v_6

The primary segments between node pairs (v_2, v_4) and (v_5, v_6) are unprotected, and their total probability of failure must be at most the maximum failure probability of $P = 0.2$. The primary segments between node pairs (v_1, v_2) and (v_4, v_5) are completely protected with the primary path segment being protected by a disjoint backup path. After a failure of either of these protected primary segments, one unit of flow still remains; they contribute a total failure probability of zero to the routing. In this example, disjoint paths were used for the protected segments between node pairs (v_1, v_2) and (v_4, v_5) .

In fact, the lowest cost allocation to form a protected segment between any two nodes i and j is the minimum-cost pair of disjoint paths between the two, as demonstrated in Lemma 1.

Lemma 1. *The minimum-cost protected segment between nodes i and j is the minimum-cost pair of disjoint paths.*

Proof: For a segment between i and j to be protected, 1 unit of flow must remain between the source and destination after any edge failure in the primary path. No backup edge will have an allocation greater than 1 because the primary flow will have 1 unit, and exactly 1 unit will need to be restored after any primary failure. An equivalent problem is to find the lowest-cost routing for 2 units of flow between i and j in a network where every edge has a maximum capacity of 1. After any edge failure, at least 1 unit of flow will remain. This is a minimum-cost flow problem [22], whose solution has integer flows when given integer inputs. Since every edge has a capacity of 1, there will be two distinct edge-disjoint flows of 1 unit each. Clearly, the lowest-cost solution has these flows routed on the minimum-cost pair of disjoint paths. ■

Using Lemma 1, every possible protected segment between any two nodes can be transformed to a single edge with a failure probability of 0 and a cost equivalent to the minimum-cost pair of disjoint paths between the nodes. We denote the cost and probability of the minimum-cost pair of disjoint paths between nodes i and j as \hat{c}_{ij} and $\hat{p}_{ij} = 0$, respectively. Now, any protected segment between some node pair (i, j) can be represented as a single edge between i and j in the network. This edge contains the primary and spare allocation that would be used if a protected segment between i and j was needed. Adding an edge for every possible protected segment transforms the problem back to the restricted version where no spare allocation was allowed. This problem can now be solved using the constrained shortest path algorithm.

Our proposed algorithm is as follows. First find the minimum-cost pair of disjoint paths between each pair of nodes; there are $O(n^2)$ such pairs. Augment the original graph with an edge between every node pair (i, j) that has cost \hat{c}_{ij} and failure probability $\hat{p}_{ij} = 0$, where \hat{c}_{ij} is the cost of the shortest pair of disjoint paths between i and j . Run the constrained shortest path algorithm on the transformed graph to find the minimum-cost solution. We call this algorithm the Segment Protected Availability Guaranteed Algorithm (SPAG).

Theorem 2. *SPAG will return a minimum-cost routing, if one exists, and has a running time of $O(n^4 \log(n) + n^2 P)$.*

Proof: To meet availability requirements, a solution will have a primary path that consists of a combination of protected and unprotected segments. As shown in Lemma 1, a protected segment between any two nodes is the shortest pair of disjoint paths between those nodes. Using the above graph transformation, an edge is added for every feasible protected segment. The constrained shortest path algorithm then evaluates every possible combination of protected and unprotected segments to find the lowest cost solution between the source and destination.

For the running time, the $O(n^4 \log(n))$ component comes from $O(n^2)$ iterations of the shortest pair of disjoint paths algorithm, which takes $O(n^2 \log(n))$ time per iteration [21]. The recursion for the constrained shortest path problem runs in $O(n^2 P)$ time. ■

A simulation similar to that of Section III-B was used to compare SPAG to the optimal solution without bifurcation for $q = 0$. Simulation results show that SPAG is in fact optimal for all tested demands.

B. Meeting Availability Requirements with $q > 0$

Next, the case of $q > 0$ is considered. The problem now has multiple availability guarantees: after an edge failure in the primary path, the flow either remains at 1 or, with at most a probability of P , drops to q . We demonstrate the complexity of a related disjoint paths problem, previously unexplored in the literature, to be *strongly* NP-Hard, which means that there exists no pseudo-polynomial or ϵ -approximation algorithm [20]. Using this result, we conjecture that the multiple availability guaranteed protection problem with $q > 0$ is also strongly NP-Hard, and present an efficient heuristic.

A sample solution is shown in Fig. 5, which consists of alternating fully-protected and q -protected segments (the dotted line being the q flow).

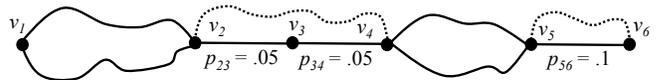


Fig. 5: Routing to meet $P = 0.2$ and $q > 0$ from v_1 to v_6

For the $q = 0$ case, unprotected segments did not use any spare allocation. To find an unprotected segment between some pair of nodes i and j that has a failure probability of at most P , a constrained shortest path was found, for which a pseudo-polynomial time algorithm exists. For a q -protected segment, instead of a constrained shortest path, a pair of disjoint paths between i and j needs to be found such that one of them is constrained: the primary path is constrained to have a probability of failure of at most P . We call this problem the Singly Constrained Shortest Pair of Disjoint Paths (SCSPD). There has been work trying to find the shortest pair of disjoint paths such that each path is constrained by the same parameter [24]. The authors of [24] found that this doubly constrained problem, while NP-Hard, has an ϵ -approximation algorithm. Their problem is distinct from ours in that SCSPD only constrains one of the two paths. Clearly, a solution to the doubly constrained problem is a feasible solution to the singly constrained one, but it is not necessarily optimal, and a lack of a solution to the former does not imply the non-existence of a solution to the latter. In fact, we show that when the constraint is relaxed for one of the paths, simply finding a feasible solution to the SCSPD problem becomes *strongly* NP-Hard, which means that a solution cannot be ϵ -approximated, nor can there be any pseudo-polynomial algorithm for optimality [20].

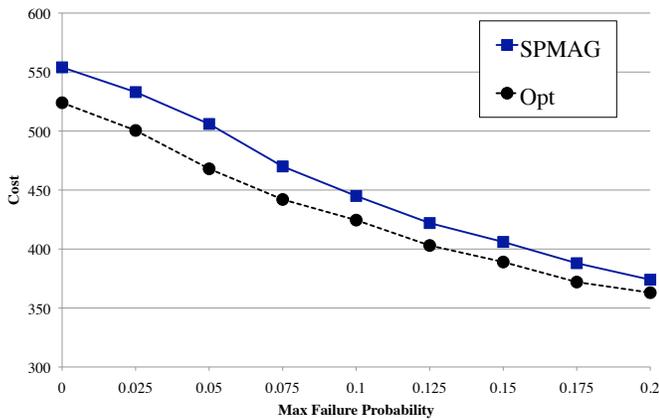


Fig. 6: SPMAG capacity cost vs. MFP with $q = \frac{1}{2}$

Theorem 3. Finding a feasible solution to SCSPD is strongly NP-Hard.

Proof: To demonstrate strong NP-Hardness of SCSPD, a reduction from the 3SAT problem [25] is performed. The proof is omitted for brevity, and can be found in [19]. ■

Since SCSPD is strongly NP-Hard, the dynamic programming approach used to solve for $q = 0$ does not work when $q > 0$. We conjecture that the multiple availability guaranteed protection when $q > 0$ is in fact also strongly NP-Hard, thereby necessitating a heuristic approach to solve the problem. Our proposed heuristic augments the $q = 0$ algorithm: after an optimal solution for $q = 0$ is found, find the shortest disjoint path for the unprotected segments and allocate a flow of q to them. We call this algorithm the Segment Protected Multiple Availability Guaranteed Algorithm (SPMAG).

A simulation similar to that of Section III-B was run comparing SPMAG and the optimal solution to MAGP without flow bifurcation. On average, SPMAG performs within 6% of the optimal solution to the multiple availability guaranteed protection problem.

V. CONCLUSION

In this paper, a novel network protection scheme with multiple availability guarantees was introduced. In particular, the multiple availability guarantees will maintain the full demand for at least a guaranteed fraction of time and guarantee a partial flow during the downtime. If the demand is allowed to drop to 50% of its flow for only 1 out of every 20 failures, a 24% reduction in excess resources can be realized over the traditional 1 + 1 full protection scheme. For the $q = 0$, which corresponds to the previously studied scenario where full availability is guaranteed for a fraction of time, we developed an optimal pseudo-polynomial algorithm. For the case of $q > 0$, we developed a time-efficient heuristic (SPMAG) that performs within 6% of the optimal solution to the multiple availability guaranteed protection problem.

REFERENCES

[1] S. Ramamurthy, L. Sahasrabudde, and B. Mukherjee, "Survivable WDM Mesh Networks," *Journal of Lightwave Technology*, vol. 21, no. 4, p. 870, 2003.

[2] C. Ou, J. Zhang, H. Zang, L. Sahasrabudde, and B. Mukherjee, "New and Improved Approaches for Shared-Path Protection in WDM Mesh Networks," *Journal of Lightwave Technology*, vol. 22, no. 5, 2004.

[3] B. Mukherjee, "WDM Optical Communication Networks: Progress and challenges," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1810–1824, 2000.

[4] 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.

[5] H. Wang, E. Modiano, and M. Médard, "Partial path protection for WDM networks: End-to-end recovery using local failure information," in *Computers and Communications, 2002. Proceedings. ISCC 2002. Seventh International Symposium on*. IEEE, 2002, pp. 719–725.

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

[7] A. Saleh and J. Simmons, "Evolution Toward the Next-Generation Core Optical Network," *Journal of Lightwave Technology*, vol. 24, no. 9, p. 3303, 2006.

[8] J. Zhang, K. Zhu, H. Zang, N. Matloff, and B. Mukherjee, "Availability-aware provisioning strategies for differentiated protection services in wavelength-convertible wdm mesh networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 15, no. 5, pp. 1177–1190, 2007.

[9] C. Saradhi, M. Gurusamy, and L. Zhou, "Differentiated qos for survivable wdm optical networks," *Communications Magazine, IEEE*, vol. 42, no. 5, pp. S8–14, 2004.

[10] M. Tacca, A. Fumagalli, A. Paradisi, F. Unghváry, K. Gadhiraaju, S. Lakshmanan, S. Rossi, A. Sachs, and D. Shah, "Differentiated reliability in optical networks: theoretical and practical results," *Journal of lightwave technology*, vol. 21, no. 11, p. 2576, 2003.

[11] R. Banner and A. Orda, "The power of tuning: A novel approach for the efficient design of survivable networks," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 4, pp. 737–749, 2007.

[12] L. Song, J. Zhang, and B. Mukherjee, "Dynamic provisioning with availability guarantee for differentiated services in survivable mesh networks," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 3, pp. 35–43, 2007.

[13] W. Yao and B. Ramamurthy, "Survivable traffic grooming with differentiated end-to-end availability guarantees in wdm mesh networks," in *Local and Metropolitan Area Networks, 2004. LANMAN 2004. The 13th IEEE Workshop on*. IEEE, 2004, pp. 87–90.

[14] J. Zhang, K. Zhu, H. Zang, and B. Mukherjee, "Service provisioning to provide per-connection-based availability guarantee in wdm mesh networks," in *Optical Fiber Communication Conference*. Optical Society of America, 2003.

[15] S. Rai, O. Deshpande, C. Ou, C. Martel, and B. Mukherjee, "Reliable multipath provisioning for high-capacity backbone mesh networks," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 4, pp. 803–812, 2007.

[16] 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.

[17] 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.

[18] —, "Partial protection in networks with backup capacity sharing," in *Optical Fiber Communication Conference*. Optical Society of America, 2012.

[19] —, "Network Protection with Multiple Availability Guarantees," MIT, Tech. Rep., 2011. [Online]. Available: <http://web.mit.edu/gregk/www/papers/NPMAG.pdf>

[20] M. Garey and D. Johnson, *Computers and intractability: A Guide to the Theory of NP-Completeness*. Freeman San Francisco, CA, 1979.

[21] J. Suurballe and R. Tarjan, "A quick method for finding shortest pairs of disjoint paths," *Networks*, vol. 14, no. 2, 1984.

[22] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, New Jersey, 1993.

[23] H. Jöksch, "The shortest route problem with constraints," *Journal of Mathematical analysis and applications*, vol. 14, pp. 191–197, 1966.

[24] A. Orda and A. Sprintson, "Efficient algorithms for computing disjoint qos paths," in *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 1. IEEE, 2003.

[25] D. Xu, Y. Chen, Y. Xiong, C. Qiao, and X. He, "On finding disjoint paths in single and dual link cost networks," in *INFOCOM 2004*, vol. 1. IEEE, 2004.