Maximizing Restorable Throughput in MPLS Networks

Reuven Cohen   Gabi Nakibly
Technion – Israel Institute of Technology, Computer Science, Haifa, Israel

***Abstract*—MPLS recovery mechanisms are increasing in popularity because they can guarantee fast restoration and high QoS assurance. Their main advantage is that their backup paths are established in advance, before a failure event takes place. Most research on the establishment of primary and backup paths has focused on minimizing the added capacity required by the backup paths in the network. However, this so-called Spare Capacity Allocation (SCA) metric is less practical for network operators who have a fixed capacitated network and want to maximize their revenues. In this paper we present a comprehensive study on restorable throughput maximization in MPLS networks. We present the first polynomial-time algorithms for the splittable version of the problem. For the unsplittable version, we provide a lower bound for the approximation ratio. We present efficient heuristics which are shown to have excellent performance. One of our most important conclusions is that when one seeks to maximize revenue, local recovery should be the recovery scheme of choice.**

I. Introduction

IP networks should support real-time applications that require stringent availability and reliability, such as Voice over IP and virtual private networks. Unfortunately, failures are still common in the daily operation of networks, for reasons such as improper configuration, faulty interfaces, and accidental fiber cuts [1] . Therefore, mechanisms that restore the flow of traffic quickly and efficiently after a failure are essential.

Many network operators employ recovery mechanisms in Layer 1 and Layer 2 protocols such as WDM, SONET/SDH, and MPLS. These recovery mechanisms guarantee fast restoration and high QoS assurance because they establish backup paths in advance, before a failure event takes place. Such recovery mechanisms are usually referred to as "protection" mechanisms.

In this paper we focus on MPLS-based protection mechanisms [2], [3]. In keeping with MPLS terminology, we refer to the path that carries the traffic before a failure as a primary LSP, and the path that carries the traffic after the primary LSP fails as a backup LSP. Throughout the paper we consider only bandwidth guaranteed protection. For this kind of protection, the backup LSP must be able to provide the same amount of guaranteed bandwidth provided by the primary LSP. To this end, resources should be reserved upon the establishment of each backup LSP, to be used only when the protected element – link or node – fails.

Many MPLS recovery schemes have been proposed. We classify these schemes as follows:

1) Global recovery (GR) schemes [3] (Figure 1(a)): In this class, each primary LSP has one backup LSP. The primary and backup LSPs share the same end nodes. The backup LSP protects against all link/node failures along the primary LSP, and it does not share any link/node with the primary LSP.
2) Local recovery (LR) schemes [4], [3] (Figure 1(b)): In this class, a separate backup LSP is constructed to protect against a possible failure of each element along the primary LSP. Each backup LSP starts at the immediate upstream node of the protected element, and ends at the tail of the primary LSP. A backup LSP may share links with the primary LSP upstream of the failure.
3) Restricted local recovery (RLR) schemes (Figure 1(c)): As in the LR scheme, a backup LSP starts at the immediate upstream node of the protected element but ends at the immediate downstream node.
4) Facility local recovery (FLR) schemes [4] (Figure 1(d)): Backup LSPs are constructed as in the RLR schemes. However, a single backup LSP protects all the primary LSPs that traverse the protected element. This makes the process of restoring the traffic to the backup LSP simpler, using MPLS label stacking [2].
5) Extended k-facility local recovery (EkFLR) schemes [5] (Figure 1(e)): Backup LSPs are constructed as in RLR. However, there might be up to k backup LSPs that protect each element. Obviously, this scheme is more flexible than FLR, and permits the preferred trade-off between higher routing efficiency (k is larger) and lower administration overhead (k is smaller).
6) Unrestricted recovery (UR) schemes (Figure 1(f)): In this class, each primary LSP may be protected by any number of backup LSPs. Moreover, each backup LSP may start and end at any point along the primary LSP, and may protect against failures of any number of elements.

A failure is frequently limited to a single network element – a link or a node. Hence, it is customary to compare recovery schemes by measuring their performance under the assumption that a failure may occur only after the network has recovered from the previous failure. An important implication of this assumption is that two backup LSPs protecting against different failures may share their reserved bandwidth.

Most past research on the selection of backup LSPs is directed at minimizing the total bandwidth reserved for the backup LSPs. To this end, backup LSPs are routed to maximize their bandwidth sharing. This optimization metric is usually referred to as Spare Capacity Allocation (SCA). Models that seek to optimize SCA usually consider a network whose links have unbounded capacity, and a cost function associated

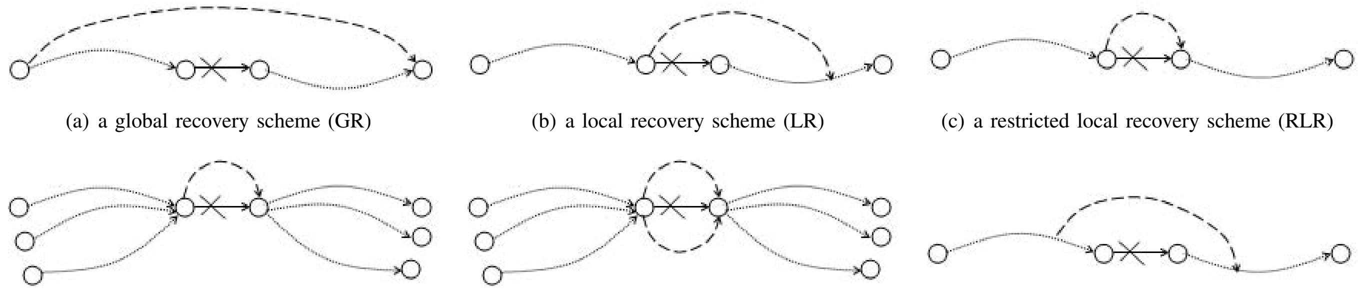

(a) a global recovery scheme (GR)

(b) a local recovery scheme (LR)

(c) a restricted local recovery scheme (RLR)

(d) a facility local scheme (FLR)

(e) an extended k-facility local recovery scheme (EkFLR, k=2)

(f) an unrestricted recovery scheme (UR)

Fig. 1. Illustrations of the various recovery schemes

with bandwidth usage. However, while minimizing the cost of building the backup LSPs is an important goal, network operators usually face a different optimization problem. They have a network with finite link capacities and seek to maximize their revenue by maximizing the traffic the network can accommodate. Another drawback of the SCA optimization for network operators is that the cost associated with an established LSP does not depend on the load imposed on the selected route. In other words, there is no incentive for load balancing.

In light of the above, SCA is not the most practical criterion for network operators. Hence, in this paper we present a comprehensive study of the problem of constructing primary and backup LSPs while *maximizing throughput*. The main contributions of paper are as follows:

1) We show that the splittable version of the problem is in $\mathcal{P}$ and we offer the first polynomial time algorithm for it. In particular, we improve the results presented in [6], where only an FPTAS was shown.
2) We show that the unsplittable version of the problem is $\mathcal{NP}$-complete and has no approximation algorithm with a ratio of $|E|^{1/2-\epsilon}$.
3) We present efficient heuristics that are shown to have excellent performance.
4) We compare the various recovery schemes with respect to the throughput maximization criterion. We show that UR, GR and, LR differ only marginally in their performance. Since LR has the fastest restoration time of the three schemes, it should be the scheme of choice.

The rest of the paper is organized as follows. In Section II we formally define the problems addressed in the paper and discuss their computational complexity. Section III presents algorithms for the problems. In section IV we conduct a simulative comparison of algorithm performance for the various recovery schemes. Finally, Section V concludes the paper.

II. Problem Definition and Computational Complexity

In this paper we focus on link failures. In [7] we extend our results to node failures as well.

We refer to the problem where both primary and backup LSPs should be established as the *Restorable Flow Problem* (RFP). The problem where primary LSPs are given and only backup LSPs must be established is referred to as the *Primary-restricted Restorable Flow Problem* (PRFP). For each of the two problems we study the splittable and the unsplittable variants. In the next subsections we formally define these problems and address their computational complexity. Table I summarizes our main results in this section.

A. The Splittable Primary-restricted Restorable Flow Problem (S-PRFP)

We now define the *Splittable Primary-restricted Restorable Flow Problem* (*S-PRFP*) with respect to each recovery scheme. For simplicity, we assume throughout the paper that only one primary LSP is established for each flow. However, the results of the paper can be easily extended for the case where every flow has several primary LSPs. Let $G = (V, E)$ be a directed graph. Let u_e be the bandwidth capacity of edge $e \in E$. Let $F \subseteq V \times V$ be a set of source-destination pairs representing traffic flow demands. For every traffic flow $f = (s_f, t_f) \in F$, let s_f be the source node, t_f the target node, d_f the bandwidth demand, P_f the sequence of edges along the primary LSP, and w_f the profit for f. A feasible solution is one that admits some of the traffic flows into the network while meeting the edge capacity constraints. Each admitted flow is routed on its primary LSP and must be fully restorable in the face of any single link failure. Hence, for every admitted flow f and edge $e \in P_f$, there must exist a set of backup LSPs that satisfies the constraints of the considered recovery scheme and can accommodate the admitted traffic of f when e fails. The objective is to maximize the total profit of the admitted traffic flows.

Note, the traffic demand of each admitted flow need not be fully satisfied. Moreover, following a failure, the admitted traffic of a flow may be split among several backup LSPs.

Theorem 1: S-PRFP is in $\mathcal{P}$ for all recovery schemes discussed in Section I.

Proof

To show this, we formulate the problem as a linear program. We first present the constraints of the problem that are common to all recovery schemes. Then, we present additional constraints for each individual scheme. We define the following variables:

- $y_{fe}^{\bar{e}}$ – the fraction of d_f routed over edge e when edge $\bar{e}$ fails; when no edge fails, $\bar{e} = \phi$.
- x_f – the total routed fraction of d_f.

	Recovery schemes					
	GR	**LR**	**RLR**	**FLR**	**EkFLR**	**UR**
S-*P*RFP (Sec. II-A)	$\mathcal{P}$ (Theorem 1)	$\mathcal{P}$ (Theorem 1)	$\mathcal{P}$ (Theorem 1)	$\mathcal{P}$ (Theorem 1)	$\mathcal{P}$ (Theorem 1)	$\mathcal{P}$ (Theorem 1)
U-*P*RFP (Sec. II-B)	$\mathcal{NP}$-C (Theorem 2)	$\mathcal{NP}$-C (Theorem 2)	$\mathcal{NP}$-C (Theorem 2)	$\mathcal{NP}$-C (Theorem 2)	$\mathcal{NP}$-C (Theorem 2)	$\mathcal{NP}$-C (Theorem 2)
	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 3)	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 3)	?	?	?	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 3)
S-RFP (Sec. II-C)	?	?	$\mathcal{P}$ (Theorem 6)	$\mathcal{P}$ (Theorem 6)	$\mathcal{P}$ (Theorem 6)	?
U-RFP (Sec. II-C)	$\mathcal{NP}$-C (Theorem 4)	$\mathcal{NP}$-C (Theorem 4)	$\mathcal{NP}$-C (Theorem 4)	$\mathcal{NP}$-C (Theorem 4)	$\mathcal{NP}$-C (Theorem 4)	$\mathcal{NP}$-C (Theorem 4)
	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 5)	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 5)	?	?	?	no $\|E\|^{\frac{1}{2}-\epsilon}$-apx. (Theorem 5)

TABLE I
SUMMARY OF THE COMPUTATIONAL COMPLEXITIES OF THE PROBLEMS

The target function is to maximize the total gained profit:

Maximize $\sum_f w_f \cdot x_f$

subject to the following constraints:

$$\text{(C-1)}\quad \sum_{e=(u,v)} y^{\bar{e}}_{fe} - \sum_{e=(v,u)} y^{\bar{e}}_{fe} = \begin{cases} x_f & v = t_f \\ -x_f & v = s_f \\ 0 & \text{else} \end{cases}$$

$$\forall v \in V, \forall f \in F, \forall \bar{e} \in \{E, \phi\}$$

$$\text{(C-2)}\quad \sum_f d_f \cdot y^{\bar{e}}_{fe} \le u_e \qquad \forall e \in E, \forall \bar{e} \in \{E, \phi\}$$

$$\text{(C-3)}\quad y^{\phi}_{fe} = 0 \qquad \forall f \in F, \forall e \notin P_f$$

$$\text{(C-4)}\quad y^{e}_{fe} = 0 \qquad \forall f \in F, \forall e \in E$$

$$\text{(C-5)}\quad 0 \le x_f \le 1,\ 0 \le y^{\bar{e}}_{fe} \le 1 \qquad \forall e \in E, \forall \bar{e} \in \{E, \phi\}, \forall f \in F.$$

The set (C-1) of constraints ensures flow conservation. The set (C-2) ensures that no edge carries more than its capacity. The set (C-3) ensures that when no failure occurs, each flow is routed only along its primary LSP. The set (C-4) ensures that no flow is routed over a failed link. Finally, the set (C-5) of constraints ensures that the total routed bandwidth of each flow and the routed bandwidth on each backup LSP do not exceed flow demand.

Each of the recovery schemes presented in Section I imposes a set of additional constraints on the backup LSPs. We now present the specific set of constraints for each recovery scheme.

The specific set of constraints for the LR scheme is:

$$\text{(LR-1)}\quad y^{\bar{e}}_{fe} \ge y^{\phi}_{fe} \qquad \forall f \in F, \forall \bar{e} \in E, \{e | e \in E, e \ne \bar{e} \text{ and } e \text{ is not a downstream edge of } \bar{e} \text{ along } P_f\}.$$

The above set of constraints ensures that the backup LSP of f for $\bar{e} = (u, v)$, assuming $\bar{e} \in P_f$, will follow the primary LSP all the way from the source to u. From node u to the destination node, the backup LSP is not constrained.

The specific set of constraints for the RLR scheme is:

$$\text{(RLR-1)}\quad y^{\bar{e}}_{fe} \ge y^{\phi}_{fe} \quad \forall f \in F, \forall \bar{e} \in E, \{e | e \in E, e \ne \bar{e}\}.$$

RLR-1 is similar to LR-1, except that it also ensures that if a backup LSP protects against a failure of edge $\bar{e} = (u, v)$, it will follow the primary LSP not only from the source to u but also from v to the destination.

Since S-PRFP allows the traffic of the failed primary LSP to be split between several backup LSPs, RLR may use an unbounded number of backup LSP for each link failure. Hence, it is easy to see that an optimal solution for the FLR scheme and for the EkFLR scheme can be produced from an optimal solution for the RLR scheme. Hence, there is no need to specify special constraints for these two recovery schemes.

The specific set of constraints for the GR scheme is:

$$\text{(GR-1)}\quad y^{\bar{e}}_{fe} \begin{cases} = \ 0 & \forall \bar{e} \in E, \{f | f \in F, \bar{e} \in P_f\}, e \in P_f \\ \ge \ y^{\phi}_{fe} & \text{otherwise} \end{cases}$$

$$\text{(GR-2)}\quad y^{\bar{e}}_{fe} - y^{\phi}_{fe} = \Delta y^{\bar{e}}_{fe} \qquad \forall f \in F, \forall \bar{e}, e \in E$$

$$\text{(GR-3)}\quad \Delta y^{\bar{e}_1}_{fe} = \Delta y^{\bar{e}_2}_{fe} \qquad \forall f \in F, \forall \bar{e}_1, \bar{e}_2 \in P_f \text{ where } \bar{e}_2 \text{ immediately follows } \bar{e}_1 \text{ on } P_f, \forall e \in E.$$

The set (GR-1) of constraints ensures that the backup LSPs of every flow whose primary LSP crosses the failed edge must be edge disjoint with the primary LSP. The set (GR-2) introduces auxiliary variables $\Delta y^{\bar{e}}_{fe}$. Each of these variables represents the difference between the bandwidth of f routed on the primary LSP, and the bandwidth of f to be routed on the backup LSPs that protect the flow against the failure of edge $\bar{e}$. The set (GR-3) of constraints ensures that for each flow the same set of backup LSPs is used to protect all the edges along the primary LSP.

Finally, the set of specific constraints for the UR scheme is:

$$\text{(UR-1)}\quad y^{\bar{e}}_{fe} \ge y^{\phi}_{fe} \quad \forall \bar{e}, e \in E, \forall f \in F, \bar{e} \notin P_f.$$

This set ensures that if the failed link is not included in the primary LSP of a flow, then the backup LSP is identical to the primary LSP. Otherwise, the set of backup LSPs has no constraint. □

B. The Unsplittable Primary-restricted Restorable Flow Problem (U-PRFP)

We now address the *Unsplittable Primary-restricted Restorable Flow Problem* (*U-PRFP*). There are two differences between U-PRFP and S-PRFP. First, in U-PRFP, a profit can be obtained for a flow only when its entire demand is satisfied. Second, in U-PRFP, the traffic of each flow can be restored using only a single backup LSP. We now address the computational complexity of U-PRFP.

Theorem 2: U-PRFP is $\mathcal{NP}$-complete for all recovery schemes.

Theorem 3: U-PRFP for GR, LR, or UR schemes cannot be approximated within $|E|^{1/2-\epsilon}$ unless $\mathcal{P} = \mathcal{NP}$.

The proofs of both theorems use reductions from U-FP [8]. The proofs are omitted for lack of space, but are presented in [7].

C. The Unsplittable and Splittable Restorable Flow Problems (U-RFP and S-RFP)

We now address the *Unsplittable Restorable Flow Problem* (U-RFP) and the *Splittable Restorable Flow Problem* (S-RFP). Recall that the goal of these problems is to establish not only the backup, but also the original (primary) LSPs. U-RFP establishes one primary LSP for every flow, and one backup LSP for every failure event along the selected primary LSP. A profit is obtained for an admitted flow only if its entire demand is satisfied. In contrast, S-RFP can split the traffic over several primary LSPs. Every edge along these LSPs can be protected by several backup LSPs. The demand of every flow can be partially satisfied, in which case only part of the profit is obtained.

Theorem 4: U-RFP is $\mathcal{NP}$-complete for all recovery schemes discussed in Section I.

This is a trivial consequence of Theorems 2.

Theorem 5: U-RFP for GR, LR, or UR cannot be approximated within $|E|^{1/2-\epsilon}$ unless $\mathcal{P} = \mathcal{NP}$.

This can be shown using a similar reduction to the one presented in the proof of Theorem 3. The details are presented in [7].

Theorem 6: S-RFP is in $\mathcal{P}$ for RLR, FLR, and EkFLR schemes.

Proof sketch

The linear program constraints for S-PRFP with RLR do not depend on the primary LSPs. Thus, we can use this linear program for S-RFP without the set of constraints (C-3) that restricts the primary LSPs. Hence, we get that S-RFP with RLR can be solved in polynomial time. As noted in Section II-A, it can be easily shown that the optimal solutions for S-RFP with RLR, FLR and EkFLR are the same. □

Note that Theorem 6 improves the results presented in [6], where only an FPTAS was shown for S-RFP with RLR.

III. Algorithms for U-PRFP and U-RFP

We present in this section heuristics for U-PRFP. These heuristics can be extended to U-RFP in a straightforward manner. The first heuristic begins by solving the problem called the *Splittable Primary-restricted Flow Problem* (S-PFP). In this problem each flow can only be routed along a primary LSP given in advance, and restorability is not considered. We then sort the flows in a non-increasing order of w_f/d_f. Then, for each flow, we apply the randomized rounding procedure. If the flow is selected by the randomized rounding procedure, we verify that (a) the flow can be routed on its primary LSP without violating the capacity constraints, and (b) for the chosen recovery scheme, feasible backup LSPs that do not violate the capacity constraints also exist. If both conditions hold, the flow is admitted along with its backup LSPs. If there are several feasible backup LSPs, the shortest one is selected.

The second heuristic is based on a well-known algorithm for U-FP [8]. In this algorithm the flows are considered sequentially. A flow f is admitted only if there is a feasible path P for which the following ratio exceeds a predefined threshold.

$$H(f,P) = \frac{w_f}{d_f \sum_{e \in P} 1/u(e)}$$

Our second heuristic adapts this algorithm by admitting a flow f only if

1) The primary LSP P_f is feasible and $H(f, P_f)$ exceeds a given threshold.
2) There are backup LSPs that protect the above primary LSP which are feasible, and their H ratio exceeds the given threshold.

The heuristics are presented in more detail in [7].

IV. Simulation Study

In this section we evaluate the performance of the algorithms presented in Sections II and III for the various recovery schemes. We use the BRITE simulator [9] to simulate MPLS domain topologies according to the Barabasi-Albert model [10].

For each topology, we generate a set of flows according to a power-law distribution. A network topology and a set of flows form together one simulation instance. In the case of PRFP, the simulation instance also contains the primary LSP for each flow. For the primary LSP of each flow, we select the shortest path.

We start by evaluating the various recovery schemes using the optimal algorithm for S-PRFP as presented in Section II-A (OPT-S-PRFP). Figure 2 depicts the results for an MPLS domain with 40 LSRs whose average node degree is 3. To compare the performance of the various recovery schemes, we use a relative performance metric: the ratio between the profit of flows admitted by OPT-S-PRFP and the profit of flows that can be admitted when no backup LSPs have to be established (OPT-S-PFP). This relative performance metric indicates the "penalty" incurred by the restoration requirement. This metric is represented by the y-axis in Figure 2, while the offered load is represented by the x-axis. The value of the offered load is the average number of flows originated by each router.

As expected, it is evident from the graph that UR yields the best performance while RLR yields the worst. In addition, Figure 2 shows that GR yields higher profit than LR. However, the performance of UR, GR and LR differs only marginally (5% on the average), whereas RLR lags behind by about 15%.

Next, we evaluate the performance of the two heuristics (denoted by U-PRFP-1 and U-PRFP-2) for different levels of offered load. Figure 3 depicts the average performance of U-PRFP-1 and U-PRFP-2 over all recovery schemes as a function of the offered load. In this case, we use a different relative performance metric: the ratio between the profit of flows admitted by the U-PRFP heuristic and the profit of flows admitted by OPT-S-PRFP using the same recovery scheme. This relative performance metric indicates the penalty incurred due to the inability to split the traffic following a failure. The

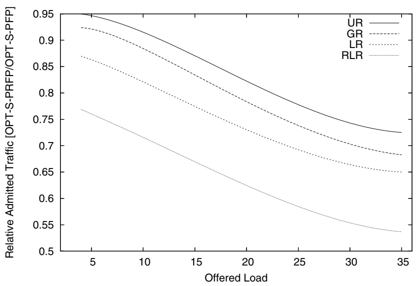

Fig. 2. Relative performance for OPT-S-PRFP

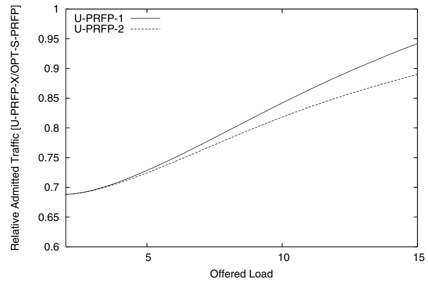

Fig. 3. Relative performance of the heuristics for U-PRFP as a function of the offered load

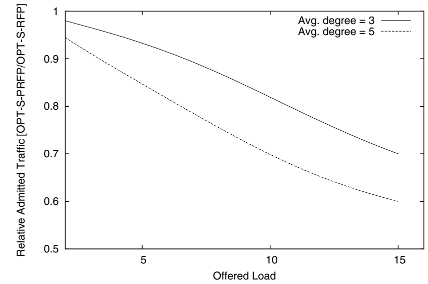

Fig. 4. The relative performance of S-PRFP using RLR

metric is represented by the y-axis The x-axis is the offered load. First, note that as the offered load increases, so does the performance difference between the two algorithms. Note also that algorithm U-PRFP-1 outperforms U-PRFP-2 by up to 7%. This can be attributed to the fact that U-PRFP-1 bases the decision whether to admit a flow on the optimal solution for S-PFP, which takes into account all of the flows. In contrast, U-PRFP-2 bases the same decision on the current network condition only. It is also evident from Figure 3 that as the offered load increases, the penalty from the requirement not to split the flow decreases, and the performance of the heuristics approaches the performnace of OPT-S-PRFP. This result can be explained by the fact that for all flows we use the shortest path as a primary LSP. Due to the power-law distribution of the flows in the network, the primary LSPs create bandwidth bottlenecks when the network becomes congested. This means that even though the optimal algorithm for S-PRFP might be able to back up more flows, these flows cannot be admitted because their primary LSP is saturated.

We now evaluate the penalty of using a single primary LSP for each flow. To this end we use the following relative performance metric: the ratio between the profit of flows that are admitted by OPT-S-PRFP and the profit of flows that are admitted by OPT-S-RFP using the RLR scheme. This relative performance metric indicates the penalty incurred when using a single primary LSP set in advance. This metric is represented by the y-axis of the graph in Figure 4, which depicts the results for two types of 20-LSR MPLS domains: one with average degree of 3 and another with average degree of 5. It is evident that as the offered load increases, so does the penalty for using a single primary LSP set in advance. This relation is not surprising since a highly loaded network requires the traffic to be split into several paths in order to maximize the admitted traffic. It is also evident that the penalty increases for a network with a higher average degree. This is because a higher network degree gives more options for splitting the traffic between two end nodes.

To summarize, the main conclusions we draw from the simulation study are:

- The performance differences between UR, GR, and LR are only marginal while RLR is considerably worse. Hence, LR should be the recovery scheme of choice due to its short restoration time (Fig. 2).
- Heuristics U-PRFP-1 and U-PRFP-2 achieve close to optimal profit (Fig. 3).
- In congested networks, U-PRFP-1 outperforms U-PRFP-2 (Fig. 3).
- When the primary LSPs are set in advance in congested networks, splitting the backup LSPs yields only a small added profit (Fig. 3).
- In non-congested networks, the added profit is small for joint optimization of primary and backup LSPs (Fig. 4).

V. Conclusions

We presented the first comprehensive study of maximizing restorable throughput in MPLS networks. We focused on the establishment of backup LSPs when the primary LSPs are already set. We showed that the splittable version of the problem is in $\mathcal{P}$ while the unsplittable version is $\mathcal{NP}$-complete and cannot be approximated within $|E|^{1/2-\epsilon}$. We developed two practical and efficient heuristics that were shown to achieve excellent performance. Using simulation, we compared the performance of the various MPLS recovery schemes. We showed that LR should be the scheme of choice since it has the fastest restoration time and almost the same performance as the best (UR) scheme.

References

[1] A. Markopulu, G. Iannaccone, S. Bhattacharya, C.-N. Chuah, and C. Diot, "Characterization of failures in an IP backbone," in *Proceedings IEEE INFOCOM*, March 2004.

[2] E. Rosen, "Multiprotocol label switching architecture (MPLS)," RFC 3031, January 2001.

[3] V. Sharma and F. Hellstrand, "Framework for multi-protocol label switching (MPLS)-based recovery," IETF RFC 3469, February 2003.

[4] P. Pan and et. al., "Fast reroute extensions to RSVP-TE for LSP tunnels," IETF RFC 4090, May 2005.

[5] M. Alicherry and R. Bhatia, "Simple pre-provisioning scheme to enable fast restoration," *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, pp. 400–412, 2007.

[6] R. Bhatia, M. S. Kodialam, T. V. Lakshman, and S. Sengupta, "Capacity allocation and routing of locally restorable bandwidth guaranteed connections," in *Proceedings of the IEEE INFOCOM*, 2005.

[7] R. Cohen and G. Nakibly, "Maximizing restorable throughput in MPLS networks," *Technion-Israel Institute of Technology Technical Report CS-2007-12*, jun 2007.

[8] Y. Azar and O. Regev, "Strongly polynomial algorithms for the unsplittable flow problem," in *Proceedings of the 8th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, 2001, pp. 15–29.

[9] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *Proceedings of MASCOTS*, 2001.

[10] A. Barabasi, R. Albert, and H. Jeong, "Scale-free characteristics of random networks: the topology of the world-wide web," in *Physica A: Statistical Mechanics and Its Applications*, vol. 281, June 2006, pp. 69–77.