Optimized IP-Based vs. Explicit Paths for One-to-One Backup in MPLS Fast Reroute

Michał Pióro*†, Artur Tomaszewski*, Cezary Żukowski*, David Hock‡, Matthias Hartmann‡, Michael Menth‡
*Institute of Telecommunications
Warsaw University of Technology, 00-665 Warsaw, Poland
Email: {mpp,artur}@tele.pw.edu.pl
†Department of Electrical and Information Technology
Lund University, 221-00 Lund, Sweden
Email: mpp@eit.lth.se
‡Institute of Computer Science
University of Würzburg, D-97074 Würzburg, Germany
Email: {hock,hartmann,menth}@informatik.uni-wuerzburg.de

***Abstract*—Primary and backup paths in MPLS fast reroute (FRR) may be established as shortest paths according to the administrative link costs of the IP control plane, or as explicitly calculated arbitrary paths. In both cases, the path layout can be optimized so that the maximum link utilization for a set of considered failure scenarios is minimized. In this paper, we propose a linear program for the optimization of the path layout for explicitly calculated paths, which can either produce single paths and route entire traffic along those paths, or generate multiple paths and spread the traffic among those paths providing load balancing. We compare the resulting lowest maximum link utilization in both cases with the lowest maximum link utilization that can be obtained by optimizing unique IP-based paths. Our results quantify the gain in resource efficiency usage provided by optimized explicit multiple paths or explicit single paths as compared to optimized IP-based paths.**

I. Introduction

Multiprotocol label switching (MPLS) enables connection-oriented communication in connectionless communication networks. Virtual connections can be established between any two points in the network by adding labels to each packet and forwarding the packets according to these labels on so-called label-switched paths (LSPs). To be able to quickly react to failures in the network, MPLS provides fast reroute (FRR) capabilities. These are local mechanisms that enable the failure-detecting router to switch packets to preconfigured backup LSPs. This yields faster reaction than end-to-end protection where the source node detects a failure along the primary path and then switches the traffic to the backup path.

To establish primary and backup LSPs, the routing of an associated IP control plane may be used. As an alternative, primary and backup LSPs may be set up according to arbitrary explicit paths that are pre-calculated, e.g., by some path computation element (PCE) [1]. Traffic engineering in terms of routing optimization is possible with both approaches. In this paper, we consider the maximum link utilization for the failure-free case and for a set of considered failure scenarios as an important performance metric that should be minimized.

If the MPLS path layout is based on IP routing, in the failure-free case primary LSPs follow the least-cost paths (according to the administrative link costs) from the source to the destination, and in the failure situation backup paths follow the least-cost paths from the failure-detecting router to the destination. Adjusting the administrative link costs is the only way to influence the routing. In [2], we have shown that it is an $\mathcal{NP}$-hard problem to find optimal link costs even for the failure-free case. Therefore, often heuristic methods are used for routing optimization. We proposed such a heuristic in [3]. When several equal-cost paths exist between two nodes, it is uncertain which of the paths is actually selected. This uncertainty can be avoided by using link cost settings that lead to unique shortest paths (USP). We extended our heuristic to generate optimized routing with USP in [4].

If LSPs follow explicit paths, arbitrary paths can be chosen as primary and backup paths. In this paper, we propose a mathematical formulation for optimizing the path layout in two steps. First, a number of possible paths are calculated. Then, the best set of these paths is chosen by solving an appropriate linear program (LP). Integer solutions of the LP produce single paths while non-integer solutions yield multiple paths over which traffic is spread according to a load balancing function.

IP-based paths are a proper subset of explicit single paths, which in turn are a proper subset of explicit multiple paths. Therefore, we compare the quality of those distinct path layouts after optimization.

This paper is structured as follows. Section II gives an overview of MPLS one-to-one backup. In Section III, we briefly explain our heuristic for the link cost optimization of IP routing, summarize previous work on this topic, and give an overview of related work in the area. In Section IV, we provide the proposed mathematical formulations for the optimization of explicit paths for MPLS one-to-one backup and differentiate our approach from related work. Section V compares the performance of the optimized IP-based layout of primary and backup paths with the optimal path layouts of explicit single paths and explicit multiple paths. Finally, Section VI concludes this work.

©IEEE, 14th International Telecommunications Network Strategy and Planning Symposium (NETWORKS 2010), Warsaw, Poland, September 2010

II. MPLS Fast Reroute: One-to-one Backup

In this section, we explain the basics of the MPLS fast reroute (MPLS-FRR) mechanism. It is a local backup mechanism, i.e., the routers adjacent to a failure act as so-called points of local repair (PLRs) and redirect packets over alternative local backup LSPs to the destination.

MPLS has two different backup mechanisms: facility backup with link and router bypasses provided by a number of backup LSPs routed around the failed component, and one-to-one backup with link and node detours that are specific for each of the LSPs' destinations. In this paper, we focus on the one-to-one backup option. Further details on the facility backup and other resilience mechanisms can be found in [5].

In the case of the MPLS one-to-one backup, for each flow individual backup LSP tunnels are installed from every possible PLR on the flow's path to the flow destination. Depending on the failure type, two different types of protection tunnels are used, as shown in Figure 1: link detour tunnels and router detour tunnels. As indicated in Figure 1(a), the failing link is protected by using a backup LSP from the PLR to the flow destination r_{tail} that does not contain the link. If a node fails, as indicated in Figure 1(b), the local backup path must not include the next node on the flow's primary path either. Therefore, a backup path from the PLR to r_{tail} that does not contain the next node is installed.

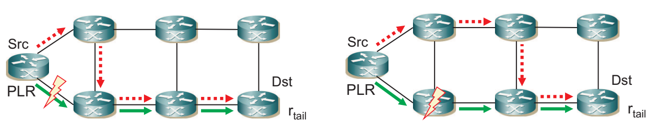

(a) $LinkDetour(PLR, r_{tail})$. (b) $RouterDetour(PLR, r_{tail})$.

Fig. 1. One-to-one backup uses detour tunnels.

When a failure occurs, it is difficult for a router to detect whether the adjacent router or the connecting link has failed. Therefore, we assume router failures whenever possible and use link detour only when the last link of the primary path fails (as the next node is just the destination, router detour cannot be applied in that case).

III. Heuristic Optimization for IP Link Cost based LSP Layout

In this section, we briefly introduce intra-domain IP routing and the unique shortest path (USP) routing property. Then, we explain our heuristic optimizer and summarize the previous work on link cost optimization.

A. Intra-domain IP Routing and USP

IP based intra-domain routing follows least-cost paths that are determined according to administrative link costs. Between two nodes in the network, e.g., between the PLR and the destination of an LSP, there may exist more than one shortest path with the minimal cost. When the backup LSP is installed, only one of the paths is arbitrarily chosen by a so-called tiebreaker. The criteria used to select this path are not standardized and might even change over time. This leads to uncertainties in the path layout and can cause unexpected load shifts on different links. To avoid this problem, link cost settings can be chosen so that neither in the failure-free case nor in the considered failure scenarios multiple least-cost paths between any two destinations exist. A path layout meeting this requirement is called a unique shortest path (USP) routing.

B. A Heuristic Link Cost Optimizer

The path layout in intra-domain IP networks can only be influenced by appropriately choosing the administrative link costs. Obtaining the link cost setting for a given network and a given traffic-matrix that leads to the lowest maximum link utilization over all links is an $\mathcal{NP}$-complete problem [2]. Therefore, heuristics are used to improve the routing. In [3], we presented the basic algorithm of the heuristic that is used in this paper. That heuristic is similar to the threshold accepting heuristic. It tries to find the link cost setting that leads to the routing with the lowest maximum link utilization over the failure free situation and a set of failure scenarios. Other objective functions also exist and are used throughout literature; an overview and comparison of those objective functions can be found in [6]. In [4], we extended our heuristic to optimize administrative link cost settings leading to USP routing. That publication also provides further information on the problems caused by unpredictable tiebreaker decisions.

To obtain good results, our heuristic is run several times with random initial link cost settings. We take the best solution of all optimization runs as the final best link cost setting.

C. Related Work

There exists vast literature regarding link cost optimization. Different heuristic and mathematical approaches are used to solve the problem. Furthermore, different objective functions are being addressed [6]. Some work considers only the failure free routing, while other work has been extended to include the optimization of link costs for a certain set of failure scenarios. Due to the limited space, we do not provide an in-depth discussion of that literature but refer to the extensive summaries in our previous publications [4], [6].

IV. Mathematical Program Formulation

A. Preliminaries

An MPLS/IP network is modeled using a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with the set of nodes $\mathcal{V}$ and the set of directed links (i.e., arcs) $\mathcal{E} \subseteq \mathcal{V}^2 \setminus \{(v,v) : v \in \mathcal{V}\}$. The nodes correspond to the MPLS/IP routers while the arcs correspond to the directed IP links. For each $e \in \mathcal{E}$, $a(e)$ will denote the originating node, $b(e)$ the terminating node, and c_e the capacity of link e. For each $v \in \mathcal{V}$, the sets of arcs outgoing from node v and incoming into node v will be denoted by $\delta^+(v)$ and $\delta^-(v)$, respectively.

Let $\mathcal{D}$ be the set of demands. For each $d \in \mathcal{D}$, let $o(d)$ and $t(d)$ denote, respectively, the originating node and the terminating node of demand d, and let h_d denote the volume of demand d, which is the required bandwidth of the corresponding LSP connection.

Let $\mathcal{S}$ be the set of failure states. In the paper we consider all single, complete failures of the links, thus $\mathcal{S} \equiv \mathcal{E}$.

For each $e \in \mathcal{E}$, let variable X_e define the total load of link e in the normal state, i.e., the state with all links operating. Further, for each $e \in \mathcal{E}$ and $s \in \mathcal{S}$, let variable Y_e^s define the total load of link e in the network state corresponding to the failure of link s. The objective of the problem of optimizing the routing of LSP connections and the layout of their primary and backup paths can now be shortly defined as follows:

$$\min \quad Z \tag{1a}$$

$$s.t. \quad Z \geq X_e/c_e \qquad e \in \mathcal{E} \tag{1b}$$

$$Z \geq Y_e^s/c_e \qquad e \in \mathcal{E}, s \in \mathcal{S} \setminus \{e\}. \tag{1c}$$

We now introduce a formal model of the problem, (FRR-LP), which is based on the link-path formulation (for the notion of link-path formulation see [7]).

B. Link-path Formulation

Link-path (LP) formulation considers the paths of LSP connections explicitly, both in the normal state (primary paths) and in each failure state (backup paths). The paths are handled through appropriate path lists that are predefined in any instance of (FRR-LP). In effect, the path lists define an instance of (FRR-LP) when other input parameters (as the network's graph, demand, etc.) are fixed and given.

For each $d \in \mathcal{D}$, let $\mathcal{P}_d$ be a set of candidate primary paths from $o(d)$ to $t(d)$ for the LSP connection of demand d, and for each $p \in \mathcal{P}_d$, let x_{dp} be a binary path variable that is equal 1 if, and only if, path p is selected for demand d as its primary path. For each $e \in \mathcal{E}$, let P^e denote the set of all candidate primary paths that use link e, and for each $d \in \mathcal{D}$ and $e \in \mathcal{E}$, let $\mathcal{P}_d^e \subseteq \mathcal{P}_d$ denote the set of all candidate primary paths from $\mathcal{P}_d$ that use link e.

Further, for each $s \in \mathcal{S}$, $d \in \mathcal{D}$ and $p \in \mathcal{P}_d^s$, let $\mathcal{Q}_{dps}$ be a set of candidate backup paths for the primary path p in failure state s, i.e., when link s fails; each path $q \in \mathcal{Q}_{dps}$ starts in node $a(s)$ and terminates in node $t(d)$. Note that each primary path $p \in \mathcal{P}_d$ has its own set of candidate backup paths for each failure state s. For each $q \in \mathcal{Q}_{dps}$, let y_{dpsq} be a binary path variable that is equal 1 if, and only if, for demand d path p has been selected as the primary path, and path q is selected as the backup path for p in the network state corresponding to the failure of link s. For each $s \in \mathcal{S}$, $d \in \mathcal{D}$, $p \in \mathcal{P}_d^s$ and $e \in \mathcal{E}$, let $\mathcal{Q}_{dps}^e \subseteq \mathcal{Q}_{dps}$ denote the set of all candidate backup paths from $\mathcal{Q}_{dps}$ that use link e. Finally, let $\mathcal{W}_d^{es}$ denote the set of such paths p in $\mathcal{P}_d^e$ that link e belongs to that part of p which is not affected by the failure of link s. Note that if $p \in \mathcal{P}_d^e$ and $s \notin p$, then $p \in \mathcal{W}_d^{es}$. The path selection constraints of (FRR-LP) are as follows:

$$\sum_{p \in \mathcal{P}_d} x_{dp} = 1 \qquad d \in \mathcal{D} \tag{2a}$$

$$\sum_{q \in \mathcal{Q}_{dps}} y_{dpsq} = x_{dp} \qquad s \in \mathcal{S}, d \in \mathcal{D}, p \in \mathcal{P}_d. \tag{2b}$$

Constraint (2a) says that for each demand d exactly one primary path must be selected, and constraint (2b) says that exactly one backup path must be selected for primary path p if p is selected as the primary path of demand d. Then, variables X_e and Y_e^s specifying the total load of every link in each network state (normal or failure) are quite straightforwardly defined by the path variables:

$$X_e = \sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{P}_d^e} h_d x_{dp} \qquad e \in \mathcal{E} \tag{3a}$$

$$Y_e^s = \sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{W}_d^{es}} h_d x_{dp} + \sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{P}_d^s} \sum_{q \in \mathcal{Q}_{dps}^e} h_d y_{dpsq} \qquad e \in \mathcal{E}, s \in \mathcal{S}. \tag{3b}$$

Altogether constraints (1) and (2)-(3) define the link-path formulation (FRR-LP) of the FRR design problem.

In the link-path formulation all structural properties of the paths are controlled by the path generation method (for explanation of path generation in network design problems see [7]). Basically, the paths are shortest paths with respect to the optimal values of the dual variables corresponding to the constraints of the (FRR-LP) (see Section IV-D), but they must also satisfy additional conditions implied by the nature of the studied problem, mainly that a backup path may not go through the end node of the failed link: for $p \in \mathcal{P}_d$ such that $s \in p$, a path $q \in \mathcal{Q}_{dps}$ must start in $a(s)$, end in $t(d)$, and omit node $b(s)$ ($a(s)$ is the point of local repair in this context). As explained in the Section IV-D, the most difficult task is generating primary paths.

C. Discussion

The basic problem dealt with in this paper could also be formulated in the node-link formulation [7]. We do not present this formulation for the lack of space. The advantage of such a formulation, referred to as (FRR-NL), is that it is compact. The number of its variables is proportional to $|\mathcal{E}|^2|\mathcal{D}|$, and not exponential with the size of the problem as in (FPR-LP). (Observe that the number of variables in (FRR-LP) is exponential as the number of paths is exponential.) This allows for solving the problem directly; the disadvantage is that we can do that only for tiny networks (e.g., to have a benchmark solution for another method). Clearly, the linear relaxation of the node-link formulation (FRR-NL) provides a lower bound of the considered objective function value, but, because of the $\mathcal{NP}$-hardness of the linear relaxation, it does not provide the solution of the problem in terms of multiple (bifurcated) explicit paths.

On the contrary, the linear relaxation of the (FRR-LP) formulation does provide a valid optimal solution in the case of multiple explicit paths, and at the same time delivers a lower bound of the considered single-path objective function value; this lower bound is in general better than the bound provided by the liner relaxation of (FRR-NL). The linear relaxation of (FRR-LP) is solved with the path generation method. While the resulting set of paths is usually very small (as already mentioned, the set of all paths is extremely large since the number of paths is exponential in the size of the graph), this set is, in general, not sufficient to solve the (FRR-LP) to optimality. Thus solving the MIP of the (FRR-LP) with that set of paths provides an upper bound (hopefully of good quality, i.e., near-optimal) of the objective function value.

When the MIP of the (FRR-LP) is solved by the branch and bound approach, in each node of the branch and bound tree (such a node is a linear relaxation of the problem with additional constraints stating which paths on the current lists must be used and which must not be used; the linear relaxation of the original problem is the root node of the tree) an attempt must be made to generate additional paths, resulting in the so called branch and price method. Such a capability is not supported by every commercial solver (in particular it is not supported by the CPLEX solver), and potentially must be implemented by hand.

The process of generating new paths at the node of the branch and bound tree (in particular at the root node of the tree) is iterative: the current solution is used to provide the values of the dual variables that correspond to the constraints of the formulation (the values measure how tightly the constraints are satisfied, or what is the impact of the constraints on the objective function). Those values of the dual variables are the metrics for path computation. The intuition behind that process is as follows: if a constraint (e.g. a capacity/utilization constraint of a link) has low impact, then there is fair amount of capacity that could be used by new paths and the value of the metric is low. It should be noted, that those metrics are solely used in the process of generating candidate paths (also explicit paths) and have nothing to do with IP link metrics.

D. Solution Methods

As already mentioned, (FRR-NL) is a compact formulation of the FRR design problem, which could be used to solve problem FRR directly with a commercial MIP solver. Since the number of binary variables in such a MIP is huge, problem (FRR-NL) can be efficiently solved only for small network instances. For medium- and large-size networks only the LP relaxation of (FRR-NL) can be solved efficiently, providing a lower bound for the optimal objective function value.

Formulation (FRR-LP) is not compact because the number of paths grows exponentially with the size of the network graph. The formulation is based on pre-specified path sets $\mathcal{P}_d$ and $\mathcal{Q}_{dps}$, and as such it should be solved by column (path) generation and the branch and price approach (see [7]).

As stated in Section IV-C, the branch and price approach is difficult to implement. Hence, we propose a simplified method. The idea of our method is first to solve the linear relaxation (LR) of problem (FRR-LP) through path generation and then to use the (short) path lists $\mathcal{P}_d, d \in \mathcal{D}$, $\mathcal{Q}_{dps}, d \in \mathcal{D}, p \in \mathcal{P}_d, s \in \mathcal{S}$ defined by non-zero flows in the optimal solution of the considered LR. Then, in the second phase, we solve problem (FRR-LP) in binary path variables using a MIP solver. The rationale behind this approach is that the total number of binary variables in an instance of (FRR-LP) is equal to $\sum_{d\in\mathcal{D}} |\mathcal{P}_d| + \sum_{d\in\mathcal{D}} \sum_{p\in\mathcal{P}_d} \sum_{s\in\mathcal{S}} \left|\mathcal{Q}_{dps}\right|$ which can be a reasonable number provided the path sets are small.

Path generation required to solve the linear relaxation of (FRR-LP) is done in the following way (for explanation of path generation see [7], and for more complex cases [8]). Let λ_d, σ_{dps}, π_e and φ_{es} be optimal dual variables corresponding to constraints (2a), (2b), (1b) and (1c), respectively. (Observe that the values of optimal dual variables are readily obtained from the LP solver while solving an instance of LR for given path sets.) Then, to generate a new backup path for each demand $d \in \mathcal{D}$, path $p \in \mathcal{P}_d$ and failed link $s \in \mathcal{S}, s \in p$, we find a shortest path q' between nodes $a(s)$ and $t(d)$ with respect to link metrics φ_{es} (this is done easily, e.g., with Dijkstra's algorithm). If the length of q' is strictly smaller than σ_{dps}/h_d then we add path p' to the list $\mathcal{Q}_{dps}$.

Generating primary paths is more difficult. For each demand $d \in \mathcal{D}$, we have to find a path p' between $o(d)$ and $t(d)$ shortest with respect to the path length defined as:

$$\sum_{e\in p'} \pi_e + \sum_{s\in p'} \alpha_s + \sum_{s\in S} \sum_{e\in p'(s)} \varphi_{es} \tag{4}$$

where α_s is the length of the shortest, with respect to link metrics φ_{es}, path q from $a(s)$ to $t(d)$ not containing node $b(s)$, and $p'(s)$ denotes the set of arcs that form the part of path p' from $s(d)$ to $a(s)$ (if $s \notin p'$ then $p'(s) = p'$). Note that the values of α_s have been found while generating backup paths as described above. Certainly, path p' is added to set $\mathcal{P}_d$ only if its length (4) is strictly smaller than λ_d/h_d. In fact, the problem of generating such a path p' is difficult, most likely $\mathcal{NP}$-hard. Still, in practice it can be solved pretty effectively by means of a specialized binary programming problem (BP) with a reasonable number of binary flow variables.

In the BP for generating a primary path p' for a fixed demand $d \in \mathcal{D}$ we use binary variables $x_e, y_{es}, z_{es}, e, s \in \mathcal{E}$. Their meaning is as follows: $x_e = 1$ if, and only if, link e belongs to the path p' that we are looking for; $y_{es} = 1$ if, and only if, e belongs to a shortest backup path (from $a(s)$ to $t(d)$ omitting node $b(s)$) of path p' in the case of the link s failure ($s \in p'$); $z_{es} = 1$ if, and only if, e belongs to the part of path p' that is not affected by the failure of link s. The BP in question is as follows:

$$\min \quad \sum_{e\in\mathcal{E}} \pi_e x_e + \sum_{s\in\mathcal{E}} \sum_{e\in\mathcal{E}} \varphi_{es} y_{es} + \sum_{s\in S} \sum_{e\in\mathcal{E}} \varphi_{es} z_{es} \tag{5a}$$

$$s.t. \quad \sum_{e\in\delta^+(v)} x_e - \sum_{e\in\delta^-(v)} x_e = 0, \;\; v \in \mathcal{V} \setminus \{o(d), t(d)\} \tag{5b}$$

$$\sum_{e\in\delta^+(v)} x_e - \sum_{e\in\delta^-(v)} x_e = 1, \quad v = o(d) \tag{5c}$$

$$\sum_{e\in\delta^+(v)} x_e - \sum_{e\in\delta^-(v)} x_e = -1, \quad v = t(d) \tag{5d}$$

$$\sum_{e\in\delta^+(v)} y_{es} - \sum_{e\in\delta^-(v)} y_{es} = 0, \quad s \in \mathcal{E}, \; v \in \mathcal{V} \setminus \{a(s), t(d)\} \tag{5e}$$

$$\sum_{e\in\delta^+(v)} y_{es} - \sum_{e\in\delta^-(v)} y_{es} = x_s, \quad s \in \mathcal{E},\ v = a(s) \tag{5f}$$

$$\sum_{e\in\delta^+(v)} y_{es} - \sum_{e\in\delta^-(v)} y_{es} = -x_s, \ s \in \mathcal{E},\ v = t(d) \tag{5g}$$

$$y_{es} \le 1 - x_e, \quad e,s \in \mathcal{E} \tag{5h}$$

$$y_{es} = 0, \quad e,s \in \mathcal{E},\ b(s) \ne t(d),\ b(e) = b(s) \tag{5i}$$

$$\sum_{e\in\delta^+(v)} z_{es} - \sum_{e\in\delta^-(v)} z_{es} = 0$$

$$s \in \mathcal{E},\ v \in \mathcal{V} \setminus \{o(d), a(s)\} \tag{5j}$$

$$\sum_{e\in\delta^+(v)} z_{es} - \sum_{e\in\delta^-(v)} z_{es} = x_s \quad s \in \mathcal{E},\ v = o(d) \tag{5k}$$

$$\sum_{e\in\delta^+(v)} z_{es} - \sum_{e\in\delta^-(v)} z_{es} = -x_s \quad s \in \mathcal{E},\ v = a(s) \tag{5l}$$

$$z_{es} \le x_e, \quad e,s \in \mathcal{E}. \tag{5m}$$

The solution of the above problem delivers a shortest path $p' = \{e \in \mathcal{E} :\ x_e = 1\}$, possibly after elimination of loops which can happen when some π_e and φ_{es} are equal to 0. (Path elimination is effective.)

V. Comparison of different optimized path layouts for MPLS One-to-One Backup

In this section, we compare different optimized path layouts for MPLS fast reroute. We investigate explicit multipath (EXPLICIT-MP) and single path (EXPLICIT-SP) path layouts as well as the single path layout based on IP link costs (IP-SP). First, we explain the experimental setup, then, we discuss the complexity of the linear programs, and finally, we provide numerical results.

A. Experimental Setup

The networks under study are displayed in Table I. These include the research networks Cost239 [9], Geant [10], and Labnet [11] as well as the popular Rocketfuel topologies [12]. The traffic matrices used for the optimization were created resembling real-world data with the method proposed in [13] and extended in [14]. All entries in the traffic matrices except for the diagonal are strict positive, i.e., there is a demand d with volume $h_d > 0$ between each arbitrary pair of nodes $v,w \in (V), v \ne w$ in the network. Thus, the total number of demands in a network is always $|\mathcal{D}| = |\mathcal{V}|\,|\mathcal{V} - 1|$.

The maximum link utilization values provided for EXPLICIT-MP and EXPLICIT-SP were obtained using path generation methods as explained in Section IV. Depending on the network instance at least one and at most 42 iterations of the PG algorithm were performed including both the generation of backup and primary paths. The values for IP-SP were obtained using the heuristic optimizer for link cost optimization presented in Section III. For each topology the heuristic was run at least 50 times with random initializations. Depending on the topology the average number of evaluations during an optimization run is between 100,000 and 600,000.

B. Complexity of the Linear Programs

The complexity of the linear programs significantly increases with the number of demands and links in the network.

The path generation providing EXPLICIT-MP could be optimally solved for the smallest considered networks, Cost239, Geant, and Labnet, and the optimal objective value represents a lower bound for the optimal solution of EXPLICIT-SP. Abovenet (AB) has the same number of nodes and demands as Labnet, still due to the much larger number of links an optimal solution could not be reached in acceptable computation time. The path generation algorithm was stopped after a preconfigured time limit set at most to several days for the largest networks. The same holds for the even larger networks (AT&T, EB, EX, SP and TI). Therefore, for these larger networks the provided EXPLICIT-MP solutions are not guaranteed to represent a lower bound to the maximum link utilization for EXPLICIT-SP.

All results obtained for EXPLICIT-SP are based on the paths provided by EXPLICIT-MP. Thus, in general, they are suboptimal and provide an upper bound to the optimal objective value for the single path MPLS-FRR one-to-one backup problem.

C. Numerical Results

First, the optimized maximum link utilization is compared for the IP link cost based single path LSP layout and for the explicit single path LSP layout. In general, the results show that their relation depends on the network topology. In some cases, both approaches lead to routing solutions of equal quality. The fact that for AB and EB networks the result is slightly better for IP-SP than for EXPLICIT-SP illustrates the fact that EXPLICIT-SP results represent only suboptimal solutions. In the other networks, IP-SP is not more than 20% worse than EXPLICIT-SP. A prominent exception to this behavior is the Labnet network: in this case the optimized IP-based path layout results in a degradation of the maximum link utilization by more than 50%. In previous work [15], we analyzed the necessary backup capacity for MPLS-FRR also in terms of the necessary configuration overhead. This overhead can be compared for EXPLICIT-SP and IP-SP by considering the number of LSP paths. Obviously, the number of primary paths is identical for both layouts and equals the number of demands $|\mathcal{D}| = |\mathcal{V}|\,|\mathcal{V} - 1|$ because the routing uses only single paths in both cases. But the number of backup paths is often much higher for EXPLICIT-SP than for IP-SP. This is probably due to the fact that while IP-SP is based on the shortest path principle, EXPLICIT-SP tends to accept longer paths (with more necessary backup paths) if they provide smaller maximum link utilization.

Next, we compare the maximum link utilization of explicit multipath and single path LSP layouts. The results show that using multiple paths can significantly reduce the load of the networks. For AB and EB networks the maximum resource

TABLE I
PERFORMANCE METRICS FOR OPTIMIZED PRIMARY AND BACKUP PATH LAYOUT USING EXPLICIT MULTIPATHS (EXPLICIT-MP), EXPLICIT SINGLE PATHS (EXPLICIT-SP), AND IP-BASED SINGLE PATHS (IP-SP).

Network				Number of primary paths $\|\mathcal{P}_d\|$			Number of backup paths $\|\mathcal{Q}_{dps}\|$			Maximum link utilization		
ID	Name	$\|\mathcal{V}\|$	$\|\mathcal{E}\|$	EXPL.-MP	EXPL.-SP	IP-SP	EXPL.-MP	EXPL.-SP	IP-SP	EXPL.-MP	EXPL.-SP	IP-SP
CO	Cost239	11	52	127	110	110	324	226	174	64.19%	83.7%	87.60%
GE	Geant	19	60	355	342	342	1005	958	874	71.64%	79.1%	92.93%
LA	Labnet	20	53	483	380	380	1928	1012	878	38.79%	45.4%	68.93%
AB	Abovenet	20	156	479	380	380	2917	865	728	22.99%	90.6%	90.31%
AT	AT&T	28	120	803	756	756	3013	2233	1982	47.78%	73.4%	87.72%
EB	Ebone	25	126	690	600	600	2614	1586	1353	30.92%	65.4%	64.55%
EX	Exodus	22	102	538	462	462	2033	1156	1041	33.35%	66.6%	68.52%
SP	Sprintlink	33	190	1127	1056	1056	4183	3679	2613	53.71%	65.3%	71.03%
TI	Tiscali	38	232	1422	1406	1406	4259	4214	3091	71.73%	79.22%	85.52%

utilization can be improved by more than 50%[1], for the other networks by at least 10%. However, the use of multipath LSP layouts can also significantly increase the number of primary and backup paths, e.g., in AB even by more than 200%. This means that the operators can have a significant gain in resource efficiency but they have to accept an increased control plane complexity. In future work, it would be interesting to find out if for EXPLICIT-SP better resource utilization values can be obtained that are much closer to the values obtained for EXPLICIT-MP.

VI. CONCLUSION

In this paper, we have discussed two alternatives for establishing primary and backup paths for the one-to-one backup option in MPLS fast reroute. One alternative uses explicit arbitrary paths and the other uses the paths that are induced by the algorithms of an IP control plane. To minimize the maximum link utilization for a set of considered failure scenarios, the layout of the explicit paths can be directly optimized. In contrast, the layout of the paths depending on the IP control plane can be optimized only indirectly by setting appropriate administrative link costs.

We presented a linear program for obtaining (1) optimal explicit primary and backup paths if flows may be bifurcated. We used those multipath structures as input for another linear program providing (2) optimized unique explicit primary and backup paths. And we used the heuristic from [4] to obtain (3) optimized unique primary and backup paths that satisfy IP routing constraints. We produced optimized paths according to (1) – (3) for various networks and traffic matrices.

A comparison has shown that (i) multiple explicit primary and backup paths often allow for significantly lower maximum link utilization than unique explicit paths, and that (ii) unique primary and backup paths satisfying IP routing constraints may lead to significantly higher maximum link utilization. On the other hand, the use of explicit path layouts may significantly increase the number of backup paths and thereby the resulting configuration effort. Thus, a considerable improvement of the resource efficiency usage in protected MPLS networks as compared to the simple setup of primary and backup paths with the IP control plane can be obtained for the price of increased control plane complexity required for establishing optimized explicit paths and potential load balancing.

[1] Combined with the investigations above, this suggests that especially the upper bounds for AB and EB are not very tight, as they are far worse than the EXPLICIT-MP values and even worse than the IP-SP values.

ACKNOWLEDGMENT

This work was supported by the Euro-FGI FP6 NoE as well as the Euro-NF FP7 NoE. The Polish authors were funded by Polish Ministry of Science and Higher Education under research grant N517 397334. The German authors were funded by Deutsche Forschungsgemeinschaft under grant TR257/23-2. They also thank Prof. P. Tran-Gia for the stimulating environment which was a prerequisite for this work. The Polish authors thank Prof. P. Pavón Mariño from University of Cartagena (Spain) for his work related to this paper.

REFERENCES

[1] A. Farrel, J.-P. Vasseur, and J. Ash, "RFC4655: A Path Computation Element (PCE)-Based Architecture," Aug. 2006.
[2] M. Pióro, Á. Szentesi, J. Harmatos, A. Jüttner, P. Gajowniczek, and S. Kozdrowski, "On Open Shortest Path First Related Network Optimisation Problems," *Performance Evaluation*, vol. 48, pp. 201 – 223, 2002.
[3] M. Menth, M. Hartmann, and R. Martin, "Robust IP Link Costs for Multilayer Resilience," in *IFIP-TC6 Networking Conference (Networking)*, Atlanta, GA, USA, May 2007.
[4] D. Hock, M. Hartmann, M. Menth, and C. Schwartz, "Optimizing Unique Shortest Paths for Resilient Routing and Fast Reroute in IP-Based Networks," in *IEEE Network Operations and Management Symposium (NOMS)*, Osaka, Japan, Apr. 2010.
[5] M. Menth, R. Martin, M. Hartmann, and U. Spoerlein, "Efficiency of Routing and Resilience Mechanisms in Packet-Switched Communication Networks," *European Transactions on Telecommunications (ETT)*, vol. 21, no. 2, 2010.
[6] M. Hartmann, D. Hock, C. Schwartz, and M. Menth, "Objective Functions for Optimization for Resilient and Non-Resilient IP Routing," in 7^{th} *International Workshop on Design of Reliable Communication Networks (DRCN)*, Washington, D.C., USA, Oct. 2009.
[7] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. San Francisco, CA, USA: Morgan Kaufman, 2004.
[8] S. Orlowski and M. Pióro, "On the complexity of column generation in survivable network design with path-based survivability mechanisms," in *International Network Optimization Conference (INOC)*, 2009.
[9] P. Batchelor et al., "Ultra High Capacity Optical Transmission Networks. Final Report of Action COST 239," Jan. 1999.
[10] "The GEANT website," http://www.geant.net/, 2008.
[11] M. Menth, "Efficient Admission Control and Routing in Resilient Communication Networks," PhD thesis, University of Würzburg, Faculty of Computer Science, July 2004.
[12] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," in *ACM SIGCOMM*, Pittsburgh, PA, Aug. 2002.
[13] A. Nucci, A. Sridharan, and N. Taft, "The Problem of Synthetically Generating IP Traffic Matrices: Initial Recommendations," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 35, no. 3, pp. 19–32, Jul. 2005.
[14] M. Roughan, "Simplifying the Synthesis of Internet Traffic Matrices," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 35, no. 5, pp. 93 – 96, 2005.
[15] R. Martin and M. Menth, "Backup Capacity Requirements for MPLS Fast Reroute," in 7^{th} *ITG Workshop on Photonic Networks*, Leipzig, Germany, Apr. 2006, pp. 95–102.