

Integer SPM: Intelligent Path Selection for Resilient Networks

Rüdiger Martin, Michael Menth, and Ulrich Spörlein

University of Würzburg, Institute of Computer Science, Germany
{martin,menth,spoerlein}@informatik.uni-wuerzburg.de

Abstract. The self-protecting multipath (SPM) is a simple and efficient end-to-end protection switching mechanism. It distributes traffic according to a path failure specific load balancing function over several disjoint paths and redistributes it if one of these paths fails. SPMs with optimal load balancing functions (oSPMs) are unnecessarily complex because traffic aggregates potentially need to be split which is an obstacle for the deployment of SPMs in practice. The contribution of this paper is the proposal of an integer SPM (iSPM), i.e., the load balancing functions take only 0/1 values and effectively become path selection functions. In addition, we propose a greedy heuristic to optimize the 0/1 distributions. Finally, we show that the iSPM is only little less efficient than the oSPM and that the computation time of the heuristic for the iSPM is clearly faster than the linear program solver for the oSPM such that the iSPM can be deployed in significantly larger networks.

1 Introduction and Related Work

Carrier grade networks typically require high availability in the order of 99.999% such that restoration or protection switching is needed. Restoration mechanisms, e.g. shortest path rerouting (SPR) in IP networks, try to find new routes after a network element fails. Such methods are simple and robust [1, 2] but also slow [3]. Protection switching pre-establishes backup paths for fast switch-over in failure cases [4]. The classical concept is end-to-end (e2e) protection with primary and backup paths. In case of a failure, the traffic is just shifted at its path ingress router from the primary to the backup path. The switching is fast, but the signalling of the failure to the ingress router takes time and traffic already on the way is lost. Therefore, fast reroute (FRR) mechanisms provide backup alternatives not only at the ingress router but at almost every node of the primary path. Fast reroute mechanisms are already in use for MPLS [5, 6] and are currently also discussed for IP networks [7–10].

In this context, the self-protecting multipath (SPM) has been proposed in previous work [11, 12] as an e2e protection switching mechanism. Its path layout consists of disjoint parallel paths and the traffic is distributed over all of them

This work was funded by Siemens AG, Munich, and by the Deutsche Forschungsgemeinschaft (DFG) under grant TR257/18-2. The authors alone are responsible for the content of the paper.

according to a traffic distribution (or load balancing) function (see Figure 1). If a single path fails, the traffic is redistributed over the working paths according to another traffic distribution function. Thus, a specific traffic distribution function \mathbf{I}_d^f is required for each demand d and for every pattern \mathbf{f} of working and non-working paths. Opposed to the conventional primary and backup paths concept, the SPM does not distinguish between a dedicated primary and backup paths. Both under failure-free conditions and in case of network failures, the traffic may be spread over several of the disjoint paths. And in contrast to optimum primary and backup paths [13], the SPM performs a traffic shift only if at least one of its disjoint paths is affected by a failure. Thus, the reaction is based on local information and signalling of remote failures across the network is not required. This is important as the connectivity in such a situation is compromised.

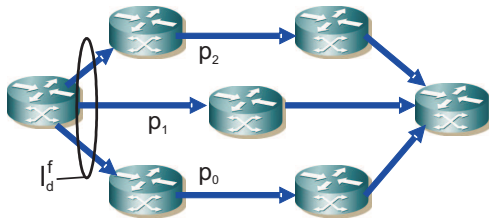


Fig. 1. The SPM distributes the traffic of a demand d over disjoint paths $\mathbf{P}_d = (p_d^0, \dots, p_d^{k_d-1})$ according to a traffic distribution function \mathbf{I}_d^f which depends on the pattern \mathbf{f} of working and non-working paths.

When a network is given with link capacities, traffic matrix, and the path layout for the disjoint paths of the SPMs, the traffic distribution functions \mathbf{I}_d^f can be optimized. Optimization means that the maximum utilization of any link in the network is minimized for a set of protected failure scenarios \mathcal{S} . Optimum traffic distribution functions \mathbf{I}_d^f can be calculated by linear programs (LPs) [14] and may split the demands for transmission over different paths. A comparison with other resilience mechanisms showed that this optimal SPM (oSMP) is very efficient [15] in the sense that it can carry more primary traffic to achieve the same maximum utilization values than optimized single shortest path (SSP) and equal-cost multipath (ECMP) IP (re)routing, variants of MPLS FRR, and various e2e protection mechanisms based on the primary and backup path principle.

However, the oSMP has three major drawbacks. Firstly, optimal traffic distribution functions require that traffic aggregates are potentially split and carried over different paths. Thus, load balancing techniques are needed for the implementation of the SPM, which makes the SPM unnecessarily complex and which is a major obstacle for its deployment. Secondly, the LPs for the optimization of the oSMP become computationally infeasible for large networks. Thirdly, load balancing techniques required for traffic distribution are problematic due to inaccuracies caused by stochastic effects [16].

The contribution of this work is the definition of the integer SPM (iSPM) that allows only 0/1 values in the traffic distribution function \mathbf{I}_d^f . This abandons the problems induced by fractional load balancing, but thereby the traffic distribution function \mathbf{I}_d^f effectively becomes a path selection function. The 0/1 constraints make the optimization more difficult. Therefore, we develop a powerful heuristic for that problem. We show that the iSPM is only little less efficient than the oSPM and that the heuristics are much faster than the LPs such that the iSPM can be applied in significantly larger networks than the oSPM.

This paper is organized as follows. Section 2 reviews the superiority of the oSPM over SSP (re)routing in small and medium-size networks and analyzes the values of the optimal traffic distribution functions. Section 3 describes the heuristic for the optimization of the 0/1 traffic distribution functions \mathbf{I}_d^f for the iSPM. Section 4 compares the efficiency of oSPM and iSPM, it studies the efficiency of the iSPM in large networks, and it compares the time for the optimization of the traffic distribution functions for the oSPM and iSPM. Finally, the conclusion in Section 5 summarizes this work.

2 The Optimal Self-Protecting Multipath (oSPM)

The configuration of the SPM in existing networks is a two-stage approach. First, the k-shortest paths algorithm from [17] finds a suitable node and link disjoint multipath \mathbf{P}_d for each demand d . Then, the traffic distribution functions \mathbf{I}_d^f must be assigned for all demands d and their respective failure patterns \mathbf{f} of working and non-working paths. In this section we briefly review the optimal assignment for the distribution functions \mathbf{I}_d^f by linear programs (LPs) [14] and show the superiority of this optimal SPM (oSPM) over single shortest path (SSP) (re)routing in small and medium size networks.

2.1 Measuring and Comparing the Efficiency of Resilience Mechanisms

We perform a parametric study to measure and compare the efficiency of resilience mechanisms. The degree $deg(v)$ of a network node v is the number of its outgoing links. We construct sample networks for which we control the number of nodes n in the range from 10 to 200, the average node degree $deg_{avg} \in \{3, 4, 5, 6\}$, and the maximum deviation of the individual node degree from the average node degree $deg_{max} = \{1, 2, 3\}$. We use the algorithm of [12] for the construction of these networks since we cannot control these parameters rigidly with the commonly used topology generators [18–22]. We sampled 5 random networks for each combination of network characteristics and tested altogether 1140 different networks. This is a huge amount of data and for the sake of clarity we restrict our presentation to a representative subset thereof. However, all statements made also hold for the larger data set. We consider the maximum link utilization of a network in all single link and router failure scenarios $s \in \mathcal{S}$ and compare it for the optimized oSPM assignment (ρ_{max}^{oSPM}) and unoptimized SSP (re)routing (ρ_{max}^{SSP}). We use the unoptimized SSP (re)routing as our comparison baseline since it is the most widely used in today’s Internet. A comparison of the oSPM to optimized SSP (re)routing can be found in [15]. We use the protected capacity gain $\gamma_{SSP}^{oSPM} = (\rho_{max}^{SSP} - \rho_{max}^{oSPM}) / \rho_{max}^{oSPM}$ as performance measure to express

how much more traffic can be transported by oSPM than by SSP with the same maximum link utilization. All figures in this paper are based on the assumption of a homogeneous traffic matrix and homogeneous link bandwidths, i.e., the entries of the traffic matrix are all the same and all links of a network have the same bandwidth. This, however, is not a major restriction as the topologies are random.

2.2 Superiority of the oSPM over SSP (Re)Routing

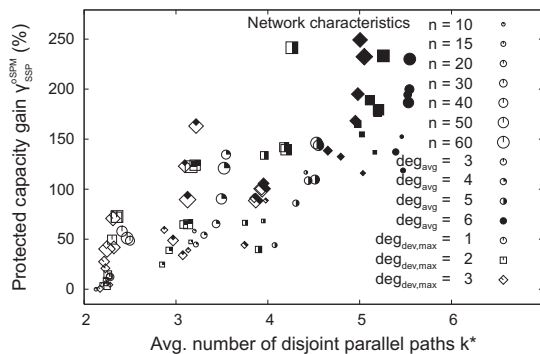


Fig. 2. Protected transmission gain γ_{SSP}^{oSPM} of the oSPM compared to SPR for random networks depending on their average number of parallel paths.

Figure 2 shows the protected capacity gain γ_{SSP}^{oSPM} for the oSPM for small to medium size networks. Each point in the figure stands for the average result of the 5 sample networks with the same characteristics. The shape, the size, and the pattern of the points determine the characteristics of these networks, the corresponding x-coordinates indicate the average number of disjoint paths k^* that could be found in the networks for the SPM structures. The protected capacity gain increases significantly with an increasing number of disjoint parallel paths k^* . More parallel paths increase the traffic distribution over the network and, thus, the capacity sharing potential for different failure scenarios. Networks with the same average node degree deg_{avg} are clustered since there is a strong correlation between k^* and deg_{avg} . Finally, large networks lead to a significantly larger protected capacity gain γ_{SSP}^{oSPM} than small networks. Ideally, link bandwidths are dimensioned for the expected traffic based on the traffic matrix and the routing. In our study, we have random networks with equal link bandwidths. Thus, there are mismatches between the bandwidth and the traffic rate on the links. As the possibility for strong mismatches increases with the network size, the potential to reduce the maximum link utilization by optimized resiliency methods also increases. Although random networks are not realistic, they help to illustrate how well routing algorithms can exploit the optimization potential.

2.3 Analysis of the oSPM Traffic Distribution Functions

# of active paths	1	2	3	4	5	Path number	1	2	3	4	5
Traffic distribution functions \mathbf{I}_d^f (%)	60	33	6.5	0.5	0	Average traffic share of a demand (%)	88.5	10	1.0	0.5	0

Table 1. Number of traffic distribution functions \mathbf{I}_d^f that use a given number of active paths for the COST239 network and the traffic share of demand d carried over the up to five possible paths in this network averaged over all traffic distribution functions and failure scenarios.

The analysis of the oSPM traffic distribution functions leads to two observations. First, most traffic distribution functions use one active path only and very few use more than two at the same time. Second, even if more than one path is active, almost all load is carried by a single active path. We exemplify these observations for the European research network COST239 in Table 1. It shows the percentage of traffic distribution functions \mathbf{I}_d^f that effectively use a certain number of active paths in the left part.

We sort the paths of an SPM in a specific failure scenario $s \in \mathcal{S}$ according to the proportion of the traffic they carry and number them. The right part shows the average proportion of the traffic carried by each of the paths. The values in the table show that the optimal traffic distribution function carry most of the traffic over a single link although more alternatives exist. These observations motivate the key idea to restrict the traffic distribution functions to 0/1 values without significantly losing the increased efficiency of the SPM.

3 The Integer SPM (iSPM)

The integer SPM (iSPM) allows only 0/1 values for the traffic distribution functions \mathbf{I}_d^f which makes the optimization even more difficult. This section first clarifies some notation and then presents a greedy heuristic to optimize iSPM configurations.

3.1 Concept and Basic Notation

To formalize the SPM concept, we explain our basic notation, introduce implications of failure scenarios, and describe the concept of path failure specific traffic distribution functions.

General Nomenclature A network $\mathcal{N} = (\mathcal{V}, \mathcal{E})$ consists of $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ unidirectional links. A single path \mathbf{p} between two distinct nodes is a set of contiguous links represented by a link vector $\mathbf{p} = \begin{pmatrix} p_0 \\ \vdots \\ p_{m-1} \end{pmatrix} \in \{0, 1\}^m$. If and only if $p_i = 1$ holds, path p contains link i . We denote traffic aggregates between routers $\mathbf{v}_i \in \mathcal{V}$ and $\mathbf{v}_j \in \mathcal{V}$ by $d = (i, j)$. The basic structure of an SPM for a traffic aggregate d is a multipath \mathbf{P}_d that consists of k_d paths \mathbf{p}_d^i for $0 \leq i < k_d$ that are link and possibly also node disjoint except for their source and destination nodes. It is represented by a vector of single paths $\mathbf{P}_d = (\mathbf{p}_d^0, \dots, \mathbf{p}_d^{k_d-1})$.

Implications of Failure Scenarios A failure scenario s is given by a set of failed links and nodes. The set of protected failure scenarios \mathcal{S} contains all outage cases including the normal working case for which the SPM should protect the traffic from being lost. The failure indication function $\phi(\mathbf{p}, s)$ yields 1 if a path p is affected by a failure scenario s ; otherwise, it yields 0. The failure symptom of a multipath \mathbf{P}_d is the vector $\mathbf{f}_d(\mathbf{s}) = \left(\phi(\mathbf{p}_d^0, \mathbf{s}), \dots, \phi(\mathbf{p}_d^{k_d-1}, \mathbf{s}) \right)^\top$ and indicates its failed single paths in case of failure scenario s . Thus, with a failure symptom of $\mathbf{f}_d = \mathbf{0}$, all paths are working while for $\mathbf{f}_d = \mathbf{1}$ connectivity cannot be maintained. The set of all different failure symptoms for the SPM \mathbf{P}_d between \mathbf{v}_i and \mathbf{v}_j is denoted by $\mathcal{F}_d = \{\mathbf{f}_d(\mathbf{s}) : s \in \mathcal{S}\}$.

Traffic Distribution Functions There is one SPM for each traffic aggregate d . This specific SPM has a general traffic distribution function to distribute the traffic over its k_d different paths. While the oSPM implements fractional traffic distribution and can use all working paths in parallel, the iSPM selects only a single path due to the restriction to 0/1 values. Thus, the iSPM uses the traffic distribution function as a path selection function. If certain paths fail, which is indicated by the symptom $\mathbf{f}_d(\mathbf{s})$, the traffic distribution function shifts the traffic to one (iSPM) or several (oSPM) of the remaining working paths. Thus, the SPM needs a traffic distribution function $\mathbf{I}_d^{\mathbf{f}}$ for each symptom $\mathbf{f} \in \mathcal{F}_d$ that results from any protected failure scenarios $s \in \mathcal{S}$. In this work, we take the protection of all single link or node failures into account such that at most one single path of a disjoint SPM multipath fails. This implies k_d different traffic distribution functions $\mathbf{I}_d^{\mathbf{f}}$ for every traffic aggregate d . Since the general traffic distribution function $\mathbf{I}_d^{\mathbf{f}} \in (\mathbb{R}_0^+)^{k_d}$ describes a distribution, it must obey $\mathbf{1}^\top \mathbf{I}_d^{\mathbf{f}} = 1$. Furthermore, failed paths must not be used.

3.2 A Greedy Algorithm for Optimizing iSPM Configurations

An iSPM configuration can be described by the following set $\mathcal{L} = \{\mathbf{I}_d^{\mathbf{f}} = \binom{n_0}{n_{k_d-1}} : d \in \mathcal{D}, f \in \mathcal{F}_d, \mathbf{I}_d^{\mathbf{f}} \in \{0, 1\}^{k_d}, \mathbf{1}^\top \mathbf{I}_d^{\mathbf{f}} = 1\}$ and comprises all traffic distribution functions of the network. A neighboring iSPM configuration \mathcal{L}' differs from \mathcal{L} by exactly one traffic distribution vector $\mathbf{I}_d^{\mathbf{f}}$. In the following $\rho_{max}^{\mathcal{S}, \mathcal{E}}(\mathcal{L})$ denotes the global maximum link utilization for a iSPM configuration \mathcal{L} over all scenarios \mathcal{S} and all links \mathcal{E} . Opposed to that, the local maximum link utilization for a iSPM configuration \mathcal{L} in scenario $s \in \mathcal{S}$ and the links of path \mathbf{p}_d^i is denoted by $\rho_{max}^{s, \mathcal{E}(\mathbf{p}_d^i)}(\mathcal{L})$. Since $\{s\} \subseteq \mathcal{S}$ and $\mathcal{E}(\mathbf{p}_d^i) \subseteq \mathcal{E}$, the inequality $\rho_{max}^{\mathcal{S}, \mathcal{E}}(\mathcal{L}) \leq \rho_{max}^{s, \mathcal{E}(\mathbf{p}_d^i)}(\mathcal{L})$ holds, i.e. the local value is only a lower bound for the global value.

```

Require: network  $\mathcal{N} = \{\mathcal{V}, \mathcal{E}\}$ , traffic demands  $\mathcal{D}$ , multipath  $\mathbf{P}_d$  for each aggregate
 $d \in \mathcal{D}$ , and initial traffic distribution functions  $\mathcal{L}$ 
1: calculate  $\rho_{max}^{new} \leftarrow \rho_{max}^{S, \mathcal{E}}(\mathcal{L})$ 
2: repeat
3:    $\rho_{max} \leftarrow \rho_{max}^{new}$ 
4:   identify scenario  $s_{max} \in \mathcal{S}$  and link  $l_{max} \in \mathcal{E}$  where  $\rho_{max}^{S, \mathcal{E}}(\mathcal{L})$  is reached
5:   for all traffic aggregates  $d$  carrying traffic over  $l_{max}$  in  $s_{max}$  do
6:     identify single path  $\mathbf{p}_d^i$  of multipath  $\mathbf{P}_d$  with  $l_{max} \in \mathbf{p}_d^i$ 
7:     for all single paths  $\mathbf{p}_d^j$  ( $j \neq i$ ) of  $\mathbf{P}_d$  do
8:       set  $\mathcal{L}(d, j)$ :  $\mathbf{p}_d^j$  carries demand  $d$  in  $s_{max}$  instead of  $\mathbf{p}_d^i$ 
9:       calculate  $\rho(d, j) \leftarrow \rho_{max}^{s_{max}, \mathcal{E}(\mathbf{p}_d^j)}(\mathcal{L}(d, j))$  with  $\mathcal{E}(\mathbf{p}_d^j) = \{l : l \in \mathbf{p}_d^j\}$ 
10:      insert  $(d, j)$  into sorted list  $Q$  according to ascending  $\rho(d, j)$ 
11:    end for
12:  end for
13:  repeat
14:    remove first tuple  $(d, j)$  from  $Q$ 
15:    calculate  $\rho_{max}^{new} \leftarrow \rho_{max}^{S, \mathcal{E}}(\mathcal{L}(d, j))$ 
16:    if  $\rho_{max}^{new} < \rho_{max}$  then
17:       $\mathcal{L} \leftarrow \mathcal{L}(d, j)$ 
18:    end if
19:  until  $\rho_{max}^{new} < \rho_{max} \vee Q = \emptyset$ 
20: until  $\rho_{max}^{new} \geq \rho_{max}$ 

```

Algorithm 1: Heuristic algorithm for the optimization of the load balancing functions of the iSPM.

Algorithm 1 describes the heuristic for the optimization of the iSPM configuration. It follows a greedy approach to keep the computational complexity low. Initially, we choose a iSPM configuration \mathcal{L} where every traffic distribution function \mathbf{I}_d^f sends the traffic for demand $d \in \mathcal{D}$ over a shortest working path for the respective failure pattern $\mathbf{f} \in \mathcal{F}$. Then, in each traversal of the outer loop (line 2-20), the algorithm basically chooses a neighboring iSPM configuration \mathcal{L}' with a lower maximum link utilization $\rho_{max}^{S, \mathcal{E}}(\mathcal{L}')$.

This is done in two steps. First, we identify the bottleneck link l_{max} and the bottleneck scenario s_{max} (line 4). Then we consider the following neighboring iSPM configurations $\mathcal{L}(d, j)$ (line 5-12). The demand d must be carried by the current configuration \mathcal{L} over the bottleneck link l_{max} (line 5) and configuration $\mathcal{L}(d, j)$ differs from \mathcal{L} only in such a way that d is relocated from the bottleneck path \mathbf{p}_d^i containing l_{max} to another path \mathbf{p}_d^j within its multipath \mathbf{P}_d (line 8). These neighboring iSPM configurations $\mathcal{L}(d, j)$ potentially improve the utilization of the bottleneck link in the bottleneck scenario. We assess their quality by the computationally less expensive local maximum utilization value $\rho(d, j) = \rho_{max}^{s_{max}, \mathcal{E}(\mathbf{p}_d^j)}(\mathcal{L}(d, j))$ (line 9) and rank them according to this value (line 10). Then, the neighboring iSPM configuration $\mathcal{L}(d, j)$ with the best local maximum utilization value $\rho(d, j)$ is chosen that also improves the overall maximum utilization value $\rho_{max}^{S, \mathcal{E}}(\mathcal{L}(d, j))$ (line 13-19).

We chose this simple version of our algorithm for presentation because it nicely shows the key concept and because it produced very good results in all our experiments. However, in pathological cases with two independent bottlenecks links l_{max} and bottleneck scenarios s_{max} the algorithm might have problems. Such cases require more enhanced methods that we cannot present here due to lack of space.

4 Results

In this section, we first show that the path selection functions of the iSPM lead to almost the same efficiency as the load balancing functions of the oSPM. Then we compare the empirical computation time for the configuration of the iSPM and the oSPM depending on the network size. Finally, we show the benefit of the iSPM with respect to single shortest path (SSP) (re)routing in large networks.

4.1 Comparison of the Efficiency of iSPM and oSPM in Small and Medium-Size Networks

Figure 3 shows the relative deviation $\Delta_{oSPM}^{iSPM} = (\rho_{max}^{iSPM} - \rho_{max}^{oSPM}) / \rho_{max}^{oSPM}$ of the maximum link utilization of the iSPM (ρ_{max}^{iSPM}) from the one of oSPM (ρ_{max}^{oSPM}). Again, each point in the figure stands for the average result of the 5 sample networks with the same characteristics. The figure reveals an obvious trend: the maximum link utilizations ρ_{max}^{iSPM} of the iSPM are larger than those of the oSPM and the difference increases with an increasing number of parallel paths k^* .

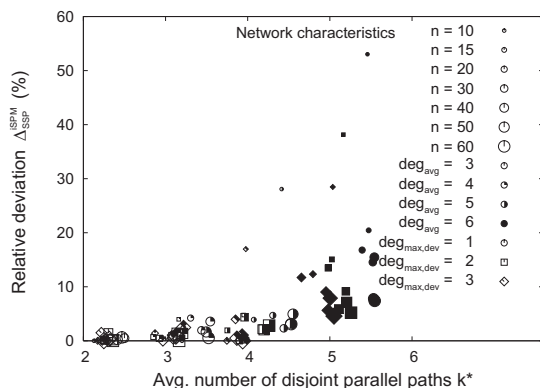


Fig. 3. Relative deviation Δ_{oSPM}^{iSPM} of the maximum link utilization of the iSPM (ρ_{max}^{iSPM}) from the one of the oSPM (ρ_{max}^{oSPM}).

The iSPM heuristic reaches deviation values of up to 50% for very small networks with $n=10$ nodes, but for large networks the deviations are rather small. We explain this observation in the following. The number of demands in the network scales quadratically with the number of nodes. Since the iSPM heuristic is restricted to integer solutions, it can shift only entire traffic aggregates to

alternate paths while the oSPM is not restricted to any traffic granularity. In particular, for $n=10$ nodes this granularity is too coarse for the iSPM to achieve similarly good maximum link utilizations as the oSPM.

For networks with at least $n \geq 30$ nodes, the deviations fall below 15%. And for networks with at least $n \geq 15$ nodes and a moderate number of disjoint parallel paths ($2 \leq k^* \leq 4.5$), the deviation is smaller than 5% compared to the one of the oSPM. Considering the fact that large values of $k^* \approx 5$ are rather unrealistic in real networks, the approximation of the oSPM by the iSPM yields very good results for realistic networks. In addition, the oSPM requires additional bandwidth to compensate load balancing inaccuracies which is not accounted for in this comparison.

As the traffic distribution function of the oSPM effectively degenerates to a path selection function in case of the iSPM, the iSPM cannot distribute the traffic of a single aggregate over different paths. However, we observe that the iSPM is still almost as efficient as the oSPM and so its efficiency also increases with an increasing number of disjoint parallel paths k^* . We explain that phenomenon as follows. The k^* disjoint paths serve as local sensors and indicate remote failures. Thus, more paths imply more accurate information about the network health that leads to a more efficient path selection in failures cases. In addition, more paths also provide more alternatives to reduce the maximum link utilization in Algorithm 1.

4.2 Comparison of the Computation Time for iSPM and oSPM

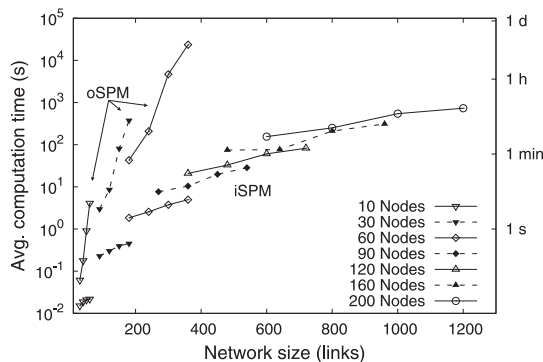


Fig. 4. Average computation time for the optimization of the iSPM and the oSPM.

Figure 4 shows the average computation time of the iSPM heuristic and the oSPM optimization depending on the network size in links and in nodes. For the iSPM, values for network sizes between 10 and 200 nodes are provided while for the oSPM, values are only available for networks of up to 60 nodes because the memory requirements of the LPs exceed the capabilities of our machines for larger networks.

The type of LP solver has a large impact on the computation time for the oSPM. The presented data in Figure 4 stem from our analysis in [14] with the COmputational INfrastructure for Operations Research (COIN-OR) solver [23] which turned out to be the fastest freely available solver for this problem formulation. While the optimization of the oSPM already reaches values in the order of a day for $n = 60$ nodes, the heuristic runs clearly below 1 h even for very large networks with $n = 200$ nodes. The computation time of the iSPM heuristic is clearly sub-exponential and neither dominated by the number of nodes nor the number of links. With an increasing number of nodes, more traffic demands are possible candidates for reallocation to alternative paths in Algorithm 1 while with an increasing number of links, the computation of the global $\rho_{max}^{\mathcal{S},\mathcal{E}}$ -value becomes more time intensive.

4.3 Efficiency of the iSPM in Large Networks

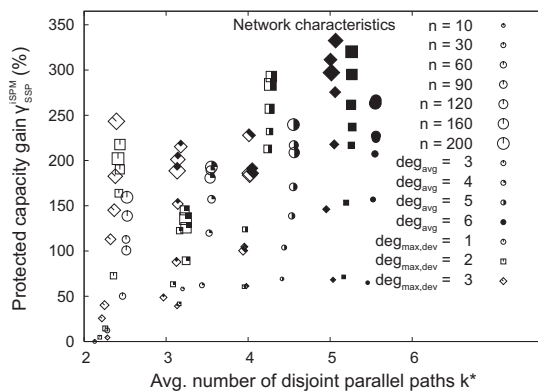


Fig. 5. Protected capacity gain γ_{SSP}^{iSPM} of the iSPM compared to SSP routing.

While Figure 2 shows the protected capacity gain γ_{SSP}^{oSPM} of the oSPM compared to single shortest path (SSP) (re)routing for random networks with 10 – 60 nodes, Figure 5 shows the gain γ_{SSP}^{iSPM} of iSPM compared to SSP routing for random networks with 10 – 200 nodes because the heuristic for the configuration of the iSPM can cope with larger networks than the LP-based optimization for the oSPM. We observed in Figure 2 that the protected capacity gain of the oSPM increases with increasing network size and this trend continues with the iSPM for larger networks in Figure 2. As a result, the iSPM can carry between 150% and 330% more protected traffic than SSP routing.

5 Conclusion

The SPM is a simple end-to-end protection switching mechanism that distributes the traffic of a single demand over several disjoint paths and it redistributes it if one of its disjoint paths fails. Thus, it is basically quite simple, but optimal path failure (\mathbf{f}) specific traffic distribution functions $\mathbf{I}_d^{\mathbf{f}}$ require that traffic aggregates

d may be split. This makes the simple mechanism unnecessarily complex and the accuracy of practical load balancing algorithms suffers from stochastic effects. In addition, the configuration of such optimal SPMs (oSPMs) in large networks is a time-consuming process that prevents its deployment in large networks.

To get rid of these problems, we suggested in this work the integer SPM (iSPM) that uses only 0/1 traffic distribution functions which effectively become path selection functions. As the restriction to 0/1 values makes the optimization problem more complex, we proposed a simple greedy heuristic to optimize the configuration of the iSPM such that the maximum link utilization of all protected failure scenarios \mathcal{S} is minimized. We showed that the iSPM is only little less efficient ($< 5\%$) than the oSPM in medium-size or large networks. Furthermore, the optimization of the configuration takes about one hour for the iSPM in networks with 200 nodes while it takes about one day for the oSPM in networks with 60 nodes. And finally, the iSPM can carry between 150% and 330% more protected traffic than hop count based single shortest path routing in large networks with 160 – 200 nodes. After all, this work brings the SPM a major step forward to deployment in practice.

Acknowledgment

The authors would like to thank Prof. Tran-Gia for the stimulating environment which was a prerequisite for that work.

References

1. Nucci, A., Schroeder, B., Bhattacharyya, S., Taft, N., Diot, C.: IGP Link Weight Assignment for Transient Link Failures. In: 18th International Teletraffic Congress (ITC), Berlin (September 2003)
2. Fortz, B., Thorup, M.: Robust Optimization of OSPF/IS-IS Weights. In: International Network Optimization Conference (INOC), Paris, France (October 2003) 225–230
3. Francois, P., Filsfils, C., Evans, J., Bonaventure, O.: Achieving Sub-Second IGP Convergence in Large IP Networks. *ACM SIGCOMM Computer Communications Review* **35**(2) (July 2005) 35 – 44
4. Autenrieth, A., Kirstädter, A.: Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS. *IEEE Communications Magazine* **40**(1) (January 2002) 50–57
5. Pan, P., Swallow, G., Atlas, A.: RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels (May 2005)
6. Saito, H., Yoshida, M.: An Optimal Recovery LSP Assignment Scheme for MPLS Fast Reroute. In: International Telecommunication Network Strategy and Planning Symposium (Networks). (June 2002) 229–234
7. Kvalbein, A., Hansen, A.F., Cicic, T., Gjessing, S., Lysne, O.: Fast ip network recovery using multiple routing configurations. In: IEEE Infocom, Barcelona, Spain (April 2006)
8. Shand, M., Bryant, S.: IP Fast Reroute Framework. <http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-06.txt> (October 2006)

9. Francois, P., Bonaventure, O.: An Evaluation of IP-Based Fast Reroute Techniques. In: CoNEXT (formerly QoFIS, NGC, MIPS), Toulouse, France (October 2005)
10. Bryant, S., Shand, M., Previdi, S.: IP Fast Reroute Using Not-via Addresses. <http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-06.txt> (October 2006)
11. Menth, M., Reifert, A., Milbrandt, J.: Self-Protecting Multipaths - A Simple and Resource-Efficient Protection Switching Mechanism for MPLS Networks. In: 3rd IFIP-TC6 Networking Conference (Networking), Athens, Greece (May 2004) 526 – 537
12. Menth, M.: Efficient Admission Control and Routing in Resilient Communication Networks. PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland (July 2004)
13. Murakami, K., Kim, H.S.: Optimal Capacity and Flow Assignment for Self-Healing ATM Networks Based on Line and End-to-End Restoration. *IEEE/ACM Transactions on Networking* **6**(2) (April 1998) 207–221
14. Menth, M., Martin, R., Spoerlein, U.: Optimization of the Self-Protecting Multipath for Deployment in Legacy Networks. In: IEEE International Conference on Communications (ICC). (June 2007)
15. Menth, M., Martin, R., Hartmann, M., Spoerlein, U.: Efficiency of Routing and Resilience Mechanisms. In: currently under submission. (2007)
16. Martin, R., Menth, M., Hemmkeppler, M.: Accuracy and Dynamics of Hash-Based Load Balancing Algorithms for Multipath Internet Routing. In: IEEE International Conference on Broadband Communication, Networks, and Systems (BROADNETS), San Jose, CA, USA (October 2006)
17. Bhandari, R.: Survivable Networks: Algorithms for Diverse Routing. Kluwer Academic Publishers, Norwell, MA, USA (1999)
18. Waxman, B.M.: Routing of Multipoint Connections. *IEEE Journal on Selected Areas in Communications* **6**(9) (1988) 1617–1622
19. Zegura, E.W., Calvert, K.L., Donahoo, M.J.: A Quantitative Comparison of Graph-Based Models for Internet Topology. *IEEE/ACM Transactions on Networking* **5**(6) (1997) 770–783
20. Jin, C., Chen, Q., Jamin, S.: tInet: Internet Topology Generator. Technical Report CSE-TR-433-00, Department of EECS, University of Michigan, USA (2000)
21. Medina, A., Matta, I., Byers, J.: BRITE: An Approach to Universal Topology Generation. In: International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), Cincinnati, Ohio, USA (August 2001)
22. Tangmunarunkit, H., Govindan, R., Jamin, S., Shenker, S., Willinger, W.: Network Topology Generators: Degree-Based vs. Structural. In: ACM SIGCOMM. (August 2002)
23. Forrest, J.: The COIN-OR Linear Program Solver (CLP). <http://www.coin-or.org> (2005)