University of Würzburg
Institute of Computer Science
Research Report Series

Failure-Specific Self-Protecting Multipaths – Increased Capacity Savings or Overengineering?

Michael Menth, Rüdiger Martin, Ulrich Spörlein

Report No. 420 April 2007

University of Würzburg
Institute of Computer Science
Department of Distributed Systems
Am Hubland, 97074 Würzburg, Germany
{menth,martin,spoerlein}@informatik.uni-wuerzburg.de

Failure-Specific Self-Protecting Multipaths – Increased Capacity Savings or Overengineering?

Michael Menth, Rüdiger Martin, Ulrich Spörlein
University of Würzburg
Institute of Computer Science
Department of Distributed Systems
Am Hubland, 97074 Würzburg, Germany
{menth,martin,spoerlein}@informatik.uni-wuerzburg.de

Abstract

The self-protection multipath (SPM) is a simple protection switching mechanism. It distributes the traffic over several disjoint paths from source to destination according to a traffic distribution function. When a path fails, the traffic is redistributed to the working paths according to another traffic distribution function, i.e., the traffic distribution function depends on the failed path. The contribution of this work is the introduction of a failure-specific traffic distribution function for the SPM that depends on the exact failure of the paths. We present a linear program for the global optimization of the traffic distribution function of all SPMs in all protected failure scenarios. Finally, we compare the amount of protected traffic that can be transported in the network for the conventional SPM and the new failure-specific SPM (FSPM).

1 Introduction

Protection switching methods are used to deviate affected traffic quickly in case of network failures. For instance, a primary path may be protected by a disjoint backup path such that if a network element of the primary path fails, the source router can quickly redirect the traffic to the backup path. The backup paths require backup capacity to carry the deviated traffic in failure cases. In packet-switched networks, several backup path can share this capacity if they are activated in different failure scenarios. This reduces the required backup capacity and thereby the entire capital expenses for the network. The self-protecting multipath (SPM) [1] is a protection switching algorithm that maximizes the sharing of backup capacities and needs, therefore, relatively little backup capacity. The structure of the SPM consists of several disjoint paths from source to destination over which the traffic is distributed according to a traffic distribution function. If one of the paths fails, the traffic is distributed over the working paths according to another traffic distribution function. The traffic distribution functions of all SPMs in the network can be optimized such that the maximum link utilization in any protected failure scenario is minimal. As a consequence, more traffic can be transported.

The contribution of this paper is the extension of the conventional SPM to a failure-specific SPM. The traffic distribution function of the SPM depends on the pattern of working and non-working paths. As they present several degrees of freedom, they can be globally optimized to

This work was funded by Siemens AG, Munich. The authors alone are responsible for the content of the paper.

minimize the maximum link utilization in any protected failure scenario. We propose now to make the traffic distribution function specific to the failed network element within the path instead of to only the failed path itself. This extension increases the number of different traffic distribution functions that are evoked in the failure scenarios and thereby the degrees of freedom are also increased. The new failure-specific SPM (FSPM) is clearly more complex. Therefore, it is crucial to assess its savings potential since only significant savings may outweigh the additional operational and implementation complexity. This is the objective of this paper.

The paper is structured as follows. Section 2 explains the SPM and the FSPM and points out their differences. Section 3 presents a linear program for the optimization of both mechanisms. Section 4 shows the performance of both mechanisms compared to simple rerouting based on the shortest path and analyzes the complexity of the optimization algorithm for both the SPM and the FSPM. Finally, we summarize this work and draw our conclusions in Section 5.

2 Overview on Resilience Mechanisms

In this section we give a short overview on various resilience mechanisms to contrast the SPM against other approaches.

2.1 Restoration Mechanisms

Restoration mechanisms take actions only after a network failure. They try to find new routes or set up explicit backup paths when the traffic cannot be forwarded anymore due to link or node failures. The disadvantage of such methods is obvious: they are too slow (0.5 - 30s) for realtime communication. However, the re-convergence of the IP routing algorithm is a very simple and robust restoration mechanism [2, 3]. Another example are backup paths in MPLS that are set up after a network failure.

2.2 Protection Switching Mechanisms

The authors of [4] give a good overview on different protection switching mechanisms for MPLS.

2.2.1 End-to-End Protection with Primary and Backup Paths

Backup paths are set up simultaneously with primary paths and in case of a failure, the traffic is just shifted at the path ingress router of a broken primary path to the corresponding backup path. This is called end-to-end protection. It is faster than restoration methods but the signalling of the failure to the path ingress router takes time and traffic being already on the way is lost.

2.2.2 Fast Reroute Mechanisms

MPLS fast reroute (FRR) tackles the problem of lost traffic in case of end-to-end protection. Backup paths towards the destination are set up not only at the ingress router of the primary path but at almost every node of the path [5]. Then, a backup path is immediately available if the primary path breaks at some location. Currently, fast reroute mechanisms are also discussed

for IP networks. Several solutions are being discussed, but a preferred method is not established yet [6–9].

2.2.3 Self-Protecting Multipath

The self-protecting multipath (SPM) has been presented first in [1, 10] and can be implemented, e.g., with MPLS. Its path layout consists of disjoint paths and the traffic is distributed over all of them according to a traffic distribution function (see Figure 1). If a single path fails, the traffic is redistributed over the working paths according to another traffic distribution function such that no traffic is lost. Thus, a specific traffic distribution function is required for every pattern of working paths.

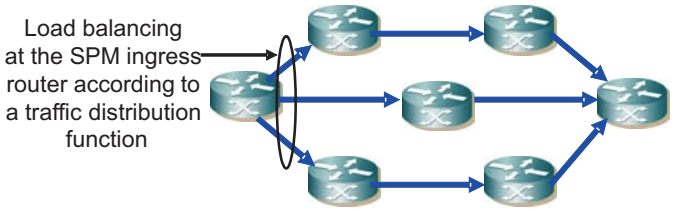

Figure 1: The SPM performs load balancing over disjoint paths according to a traffic distribution function which depends on the working paths.

Load balancing algorithms [11, 12] finally distribute the traffic over the multipaths according to the respective traffic distribution function. In [13] the SPM has been adapted to the so called Integer SPM (iSPM) that performs path selection instead of traffic distribution, i.e., it picks only one suitable path from the multipath between source and destination according to the current failure pattern. This requires a heuristic solution since the computational complexity of the optimal solution becomes infeasible. A comparison of the SPM to other resilience mechanisms based on multipath structures like demand-wise shared protection (DSP) can be found in [14].

2.2.4 Failure-Specific Self-Protecting Multipath

The contribution of this paper is the introduction and investigation of the FSPM. It works like a normal SPM, i.e., the FSPM redistributes the traffic only if one of its paths is hit by a failure. However, the FSPM provides a specific traffic distribution function for each failure location. Thus, the FSPM has more different traffic distribution functions than the normal SPM. This leads to more degrees of freedom in the optimization process and yields thereby a potential for performance improvements. However, there are also some disadvantages from the technical point of view. Firstly, the FSPM requires the knowledge of the exact failure location which must be signalled in a situation when the operation of the network is corrupt. Secondly, the increased number of different traffic distribution functions complicates the implementation of the SPM and, thirdly, it makes the optimization more difficult and time consuming. Therefore, we quantify the potential for performance improvements of the FSMP to assess whether its benefit outweighs its drawbacks.

2.3 Routing Optimization

The traffic matrix and the paths of the flows determine the resource demands on the links. The layout of the paths may be optimized to minimize either the link utilization or the required network capacity. In the following, we address briefly different optimization objectives to distinguish our optimization problem from others.

2.3.1 Routing Optimization in Combination with Network Dimensioning

In not yet provisioned networks, the network capacity and the routing may be determined. If failure scenarios are not protected, shortest path routing requires the least capacity. With resilience requirements, however, backup resources must be provided and may be shared by different flows in different failure scenarios. Routing optimization can reduce the required network capacity considerably by maximizing the capacity sharing. This has been exemplified by [1] and [15].

2.3.2 Routing Optimization for Legacy Networks

In already provisioned networks or legacy networks, the capacity of the links is fixed. If the traffic matrix is given, the maximum link utilization in the network under failure-free conditions can be minimized by a suitable routing. This has been done for IP networks [16], for MPLS networks, and for hybrid networks [17]. If restoration or protection switching is applied, the target is the minimization of the maximum link utilization in any protected failure case. This has been done for IP networks [2, 3, 18] and for MPLS networks [19]. Thereby, backup capacities may be shared by different flows and in different failure scenarios. The objective of this work is to optimize the SPM in such a way that the maximum link utilization in any protected failure scenario is minimized. This is equivalent to a maximization of the amount of transportable traffic with resilience requirements by scaling up the traffic matrix up to the point where traffic is lost in at least one failure scenario.

3 Optimization of the Normal SPM and the FSPM for Deployment in Legacy Networks

The SPM consists of parallel paths over which the traffic is distributed according to a traffic distribution function. A suitable choice of the multipath layout and the optimization of the failure specific traffic distribution function can minimize the maximum link utilization ρ_{max} in any protected failure scenario. First, we address the path layout, then we explain the linear program for the optimization of the traffic distribution functions, and, finally, we analyze the complexity of the linear program.

3.1 Path Layout

First we consider algorithms to find disjoint parallel paths and then we address the problem of SRLGs.

3.1.1 Algorithms for Disjoint Parallel Paths

The SPM consists of disjoint parallel paths. This is not a requirement, however, for disjoint paths the remaining paths are still working if a single path fails due to the failure of a single network element. Some network topologies do not allow to find disjoint paths, but we do not consider that case in this investigation and there are workarounds to cope with that problem. A very intuitive method to find link or node disjoint paths in a network is based on the shortest path algorithm. The disjoint paths are obtained iteratively: once a shortest path between a pair of nodes is found, its links and interior nodes are removed from the topology. When no additional path can be found, the algorithm stops. This simple approach cannot always find disjoint paths although a disjoint paths solution exist, or it may not always find the shortest disjoint paths. Bhandari's book [20] gives a good overview on different algorithms to find disjoint paths in networks and we use them in our software. In this work, we try to find at most 5 link and node disjoint paths for the path layout of the SPMs.

3.1.2 Adaptation to SRLGs

Shared risk link groups (SRLGs) are sets of links in a network that may fail simultaneously. Reasons may be, e.g., links on different wavelengths within a common fiber or links on different fibers within a common duct – they fail together in case of an electronic device failure or fiber cut. Another frequent reason for SRLGs are router failures. To work with SRLGs, the disjoint paths of SPMs should not contain links of the same SRLGs; otherwise, several paths of the SPM fail simultaneously and they do not protect each other anymore. Therefore, an adaptation of the paths layout to SRLGs must avoid links of common SRLGs on disjoint paths. This is a difficult NP-hard problem [21] which cannot be solved efficiently for general SRLGs. However, specific SRLGs can be respected efficiently, e.g. by node disjoint paths like in this work. The path layout for SPMs in case of SRLGs is not the focus of our work but rather the optimization of the path failure specific traffic distribution functions for SPMs in the next section.

3.2 Optimization of the Load Balancing Functions

The objective of this section is the optimization of the path failure specific traffic distribution functions for SPMs. First, we explain our notation of path concepts, then we introduce implications of failure scenarios, and finally, we propose two simple heuristics and an exact optimization for the traffic distribution functions to minimize the maximum link utilization of all protected failure scenarios.

3.2.1 Notation of Network Concepts

We introduce some basic notation from linear algebra that we use to model links, traffic aggregates, single paths, and multipaths.

Let $\mathbb{X}$ be a set of elements, then $\mathbb{X}^n$ is the set of all n-dimensional vectors and $\mathbb{X}^{n\times m}$ the is set of all $n \times m$-matrices with components taken from $\mathbb{X}$. Vectors $\mathbf{x} \in \mathbb{X}^n$ and matrices $\mathbf{X} \in \mathbb{X}^{n\times m}$ are written bold and their components are written as $\mathbf{x} = \begin{pmatrix} x_0 \\ \cdot \\ x_{n-1} \end{pmatrix}$ and $\mathbf{X} = \begin{pmatrix} x_{0,0} & \cdots & x_{0,m-1} \\ \cdot & & \cdot \\ x_{n-1,0} & \cdots & x_{n-1,m-1} \end{pmatrix}$. The

scalar multiplication $c \cdot \mathbf{v}$ and the transpose operator $^\top$ are defined as usual. The scalar product of two n-dimensional vectors $\mathbf{u}$ and $\mathbf{v}$ is written with the help of matrix multiplication $\mathbf{u}^\top \mathbf{v} = \sum_{i=1}^{n} u_i \cdot v_i$. Binary operators $\circ \in \{+, -, \cdot\}$ are applied component-wise, i.e. $\mathbf{u} \circ \mathbf{v} = (u_0 \circ v_0, \ldots, u_{n-1} \circ v_{n-1})^\top$. The same holds for relational operators $\circ \in \{<, \leq, =, \geq, >\}$, i.e. $\mathbf{u} \circ \mathbf{v}$ equals $\forall\, 0 \leq i < n\colon u_i \circ v_i$. For simplicity reasons we define special vectors $\mathbf{0} = (0, \ldots, 0)^\top$ and $\mathbf{1} = (1, \ldots, 1)^\top$ with context specific dimensions.

A network $\mathcal{N} = (\mathcal{V}, \mathcal{E})$ consists of $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ unidirectional links. The links are represented as unit vectors $\mathbf{e_i} \in \{0,1\}^m$, i.e. $(e_i)_j = 1$ if $i = j$, and $(e_i)_j = 1$ if $i \neq j$ for $0 \leq i, j < m$. We denote traffic aggregates between routers $\mathbf{v_i} \in \mathcal{V}$ and $\mathbf{v_j} \in \mathcal{V}$ by $d = (i, j)$ and the set of all aggregates by $\mathcal{D} = \{(i, j) : 0 \leq i, j < n \text{ and } i \neq j\}$. A single path p between two distinct nodes is a set of contiguous links represented by a link vector $\mathbf{p} \in \{0,1\}^m$. 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 $\mathbf{P_d} = (\mathbf{p_d^0}, ..., \mathbf{p_d^{k_d - 1}})$. Thus, a multipath is described by a matrix $\mathbf{P_d} \in \{0,1\}^{k_d \times m}$.

3.2.2 Implications of Failure Scenarios

A failure scenario s is given by a set of failing 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 $\mathbf{p}$ is affected by a failure scenario s; otherwise, it yields 0. The failure symptom of a multipath $\mathbf{P_d}$ is the vector $\mathbf{f_d}(\mathbf{s}) = \left(\phi(\mathbf{p_d^0}, \mathbf{s}), ..., \phi(\mathbf{p_d^{k_d - 1}}, \mathbf{s})\right)^\top$ 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. In this work, we take the protection of all single link or node failures into account such that at most one single path of an SPM multipath fails. The set of all different failure symptoms for the SPM $\mathbf{P_d}$ is denoted by $\mathcal{F}_d = \{\mathbf{f_d}(\mathbf{s}) : s \in \mathcal{S}\}$.

Normally, all traffic aggregates $d \in \mathcal{D}$ are active. If routers fail, some demands disappear which leads to a traffic reduction that is expressed by the failure scenario specific set of aggregates $\mathcal{D}_s$.

- *No Traffic Reduction (NTR):* We assume hypothetically that failed routers lose only their transport capability for transit flows, but they are still able to generate traffic. Therefore, we have $\mathcal{D}_s = \mathcal{D}$.
- *Source Traffic Reduction (STR):* If a certain router fails, all traffic aggregates with this source node disappear.
- *Full Traffic Reduction (FTR):* We assume that traffic aggregates with failed source *or* destination are stalled.

We use FTR for the computation of the results in this paper, but we considered all options for network dimensioning for the normal SPM in [22] and analyzed their impact.

3.2.3 The Traffic Distribution Function and Simple Heuristics

There is one SPM for each traffic aggregate $d \in \mathcal{D}$. The SPM has a traffic distribution function to distribute the traffic over its k_d different paths. If certain paths fail due to the failure s which is indicated by the symptom $\mathbf{f_d}(\mathbf{s})$, the traffic distribution function shifts the traffic to the remaining working paths. The operation of the normal SPM and the FSPM differs only with regard to this traffic distribution function.

- The normal SPM needs a traffic distribution function $\mathbf{l_d^f}$ for each symptom $\mathbf{f} \in \mathcal{F}_d$ that results from any protected failure scenarios $s \in \mathcal{S}$.
- The FSPM needs a traffic distribution function $\mathbf{l_d^s}$ for each protected failure scenarios $s \in \mathcal{S}$ that affects one of its paths, i.e., if $\mathbf{f_d}(\mathbf{s}) \cdot \mathbf{1} > 0$ is true.

Since the traffic distribution function $\mathbf{l_d^x} \in (\mathbb{R}_0^+)^{k_d}$ describes a distribution, it must obey

$$\mathbf{1}^\top \mathbf{l_d^x} = 1. \tag{1}$$

Furthermore, failed paths must not be used, i.e.

$$\mathbf{f}^\top \mathbf{l_d^x} = 0. \tag{2}$$

A simple example for a traffic distribution function is equal traffic distribution over all working paths, i.e., $\mathbf{l_d^x} = \frac{1}{\mathbf{1}^\top(\mathbf{1}-\mathbf{f})} \cdot (\mathbf{1}-\mathbf{f})$. Another relatively simple option is distributing the load over the partial paths $\mathbf{p_d^i}$ indirectly proportionally to their length $(\mathbf{1}^\top \mathbf{p_d^i})$. This can be computed by $(l_d^x)_i = \frac{1-f_i}{\mathbf{1}^\top \mathbf{p_d^i}} / \left(\sum_{0 \le j < k_d} \frac{1-f_j}{\mathbf{1}^\top \mathbf{p_d^j}} \right)$. Both heuristics require a lot of backup capacity [10]. Therefore, optimization of the traffic distribution function is required.

3.2.4 Optimization of the Load Balancing Function

The optimization configures the traffic distribution functions in such a way that the maximum link utilization ρ_{max} is minimal in any failure scenario $s \in \mathcal{S}$ for given link capacities and a given traffic matrix.

The traffic rate associated with each traffic aggregate $d \in \mathcal{D}$ is given by $c(d)$ and corresponds to an entry in the traffic matrix. We describe the network capacity by a bandwidth vector $\mathbf{b} \in (\mathbb{R}_0^+)^m$, which carries a capacity value for each link. Similarly, the vector indicating the traffic rates on all links, which are induced by a specific SPM $\mathbf{P_d}$ and a specific failure scenario $s \in \mathcal{S}$ or the resulting failure symptom $f_d(s) \in \mathcal{F}_d$, is calculated by $\mathbf{P_d} \cdot \mathbf{l_d^x} \cdot c(d)$ with $x \in \{f_d(s), s\}$.

We now formulate constraints for the traffic transport over the network in all protected scenarios $s \in \mathcal{S}$ under the side constraint that all links have a maximum utilization of ρ_{max}. In packet switched networks, resources are not physically bound to traffic aggregates. If traffic is rerouted due to a local outage, the released resources can be immediately reused for the transport of other traffic. Under this assumption, the capacity constraints are

$$\forall s \in \mathcal{S} : \sum_{d \in \mathcal{D}_s} \mathbf{P_d} \cdot \mathbf{l_d^x} \cdot c(d) \le \mathbf{b} \cdot \rho_{max}. \tag{3}$$

with $x \in \{f_d(s), s\}$. In [10, 22], we have also proposed constraints that apply when capacity cannot be reused and we have investigated them in the context of network dimensioning.

The objective of the optimization is the minimization of the maximum link utilization ρ_{max}. The free variables, which must be set in the optimization process, are the traffic distribution functions

- $\mathbf{l_d^f} \in (\mathbb{R}_0^+)^{k_d}$ for all $d \in \mathcal{D}$ and for all $\mathbf{f} \in \mathcal{F}_d$ for the normal SPM,
- $\mathbf{l_d^s} \in (\mathbb{R}_0^+)^{k_d}$ for all $d \in \mathcal{D}$ and for all $\mathbf{s} \in \mathcal{S} : \mathbf{f_d}(\mathbf{s}) \neq \mathbf{0}$ for the FSPM,

and the maximum link utilization ρ_{max} itself. The following constraints must be respected in the optimization process to obtain valid traffic distribution functions and to avoid overload on the links.

- **(C0)**: Equation (1) assures that the traffic distribution function is a distribution.
- **(C1)**: Equation (2) assures that failed paths will not be used.
- **(C2)**: Equation (3) assures that the bandwidth suffices to carry the traffic in all protected failure scenarios.

3.3 Analysis of the Linear Program Complexity

We estimate the number of free variables and the number of constraints of the LP depending on the network size since they influence its computation time and memory consumption.

3.3.1 Number of Free Variables

The maximum link utilization ρ_{max} is just a single free variable. The consideration of the traffic distribution functions $\mathbf{l}_d^{f_d(s)}$ is more complex and we address first the normal SPM only. One SPM exists for each traffic aggregate $d \in \mathcal{D}$ and for each SPM a traffic distribution function $\mathbf{l_d^f}$ is needed for every SPM failure symptom $f \in \mathcal{F}_d$. A traffic distribution vector has an entry for each of the k_d paths of the SPM. There is one traffic distribution vector for each SPM failure symptom. We take all single link and node failures into account in addition to the working scenario, so we have exactly $|\mathcal{F}_d| = k_d + 1$ different failure symptoms. We use a full traffic matrix in our study, thus, the number of traffic aggregates is $|\mathcal{D}| = n \cdot (n-1)$. We denote the average number of outgoing links per node by the average node degree deg_{avg} which can be calculated by $deg_{avg} = \frac{m}{n}$. The average number of disjoint paths for all SPMs is given by $k^* = \frac{1}{|\mathcal{D}|} \cdot \sum_{d \in \mathcal{D}} k_d$ and it is smaller than the average node degree deg_{avg}. Taking this into account, the overall number of free variables is $\sum_{d \in \mathcal{D}} k_d \cdot (k_d + 1) \approx n \cdot (n-1) \cdot k^* \cdot (k^* + 1) \leq m^2$. Thus, the number of free variables scales quadratically with the number of links in the network. In case of the FSPM, this number is multiplied by the average length $\bar{l}$ of all partial paths in terms of links and nodes which depends again on the network size.

3.3.2 Number of Constraints

We calculate the number of constraints resulting from (C0), (C1), and (C2) of the previous section. Both (C0) and (C1) require for each path failure specific traffic distribution function one constraint such that we get $n_{C0} = n_{C1} = \sum_{d \in \mathcal{D}} (k_d + 1) \approx n \cdot m$ different equations. For the FSPM this is again multiplied by $\overline{l}$. Constraint type (C2) requires an equation for each link and for each protected failure scenario, i.e. for the working scenario and all single link and node failures. Therefore, the number of constraints for (C2) is exactly $n_{C2} = m \cdot (1 + m + n)$. Thus, the overall number of constraints is roughly $m^2 + 3 \cdot m \cdot n + m$ for the normal SPM. Hence, the number of constraints also scales about quadratically with the number of links in the network. For the FSPM, this sum increases to $m^2 + m \cdot n \cdot (2 \cdot \overline{l} + 1) + m$ which is also affected by the average path length $\overline{l}$ in terms of links and nodes.

4 Results

In this section we illustrate the benefits and shortcomings of the new FSPM in comparison with the normal SPM.

4.1 Performance Improvement through FSPM Compared to the Normal SPM

First we consider the resource efficiency of the SPM compared to shortest path routing (SPR) and compare then the resource utilization of the new FSPM to the one of the normal SPM.

4.1.1 Resource Efficiency of the SPM Compared to Shortest Path Routing

We show by means of a multitude of sample networks that the SPM is a very efficient protection switching mechanism. The degree of a network node is the number of its outgoing links. We construct sample networks for which we control the number of nodes $n \in \{10, 15, 20, 25, 30, 35, 40\}$, the average node degree $deg_{avg} \in \{3, 4, 5, 6\}$, and the deviation of the individual node degree from the average node degree $deg_{dev}^{max} \in \{1, 2, 3\}$. We use the algorithm of [10] for the construction of these networks since we cannot control these parameters rigidly with the commonly used topology generators [23–27]. We sampled 5 networks for each of the 84 different network characteristics and tested altogether 420 different networks.

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 SPM $\left(\rho_{max}^{SPM}\right)$ and shortest path rerouting $\left(\rho_{max}^{SPR}\right)$. We define the protected capacity gain $\gamma = \rho_{max}^{SPR} / \rho_{max}^{SPM}$ to express how much more traffic can be transported by SPM than by SPR at the same maximum link utilization in the network. Figure 2 shows the protected capacity gain for these networks under the assumption of a homogenous traffic matrix and homogeneous link bandwidths. Each point in the figure stands for the average result of the 5 sample networks with the same characteristics. The shape and the size of the points determine the network characteristics, the corresponding x-coordinates indicate the average number of disjoint paths k^* for the SPMs in networks, and the y-coordinates show the protected capacity gain of the SPM. The figure reveals an obvious trend: the protected capacity gain of the SPM increases significantly with an increasing number of disjoint parallel paths k^* in the networks.

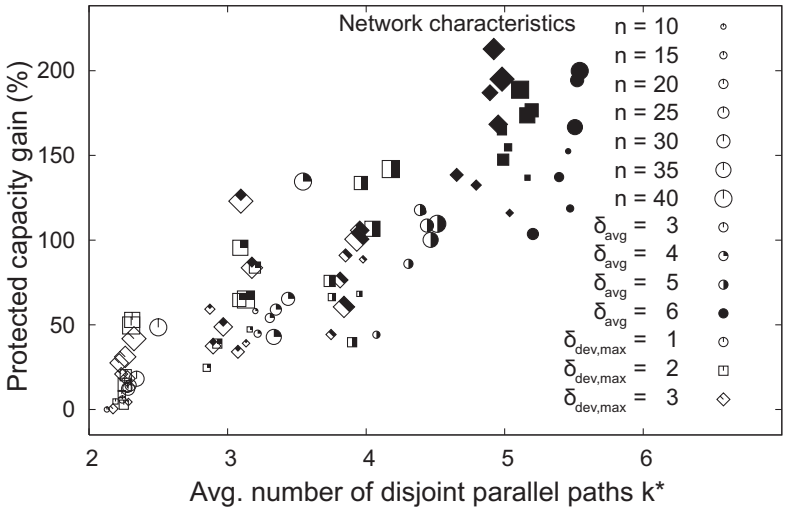

Figure 2: Protected transmission gain of the normal SPM relative to SPR for random networks.

Networks with the same average node degree deg_{avg} are obviously clustered since the average node degree deg_{avg} and k^* are strongly correlated. Networks with a small deviation deg_{dev}^{max} regarding their average node degree (circles) have a larger k^* than those with a large deg_{dev}^{max} (diamonds). Large networks lead to a slightly larger protected capacity gain than small networks, however, this trend is not so obvious. After all, the SPM is quite efficient since it can carry 50% to 200% more protected traffic than SPR in sufficiently meshed networks.

4.1.2 Comparison of the Maximum Link Utilization by SPM and FSPM

We conduct the same experiments like above for the FSPM to assess the performance improvement of the FSPM compared to the SPM. In 416 out of 420 networks the maximum link utilization of the FSPM was exactly the same. In only 4 networks the maximum link utilization was reduced by FSPM by 2.7%, 1.9%, 0.7% and 0.5% relative to the maximum link utilization of the normal SPM. These networks were 20 to 40 nodes large and had a rather small average node degree of about $deg_{avg} \approx 2.5$. Thus, the performance improvement of the FSPM compared to the normal SPM is rather negligible in most cases. This shows that the normal SPM is already a very flexible concept that does not need to be further complicated by failure-specific traffic distribution functions. We still compare the complexity of the SPM and the FSPM in the following.

4.2 Complexity Increase through FSPM Compared to the Normal SPM

We compare the complexity of the FSPM and the normal SPM. To that end we first investigate the number of different required load balancing functions for each end-to-end SPM, and then we study the computation time of the optimization problem.

4.2.1 Number of Required Traffic Distribution Functions

The normal SPM requires one load distribution function for the failure-free scenario and one load distribution function for the failure of each of its path if only a single element failure is protected. Thus, the average number of required traffic distribution functions per end-to-end SPM scales with k^*+1 with k^* being the average number of disjoint parallel paths. This is also illustrated in Figure 3. However, the number of traffic distribution functions never exceeds 6 since we limited the number of parallel path in the SPM to 5. For the FSPM the number of different traffic distribution functions scales both with k^*+1 and the length of these paths that increases with the network size. As a consequence, the FSPM requires between 7.5 and 25 different traffic distribution functions per end-to-end SPM.

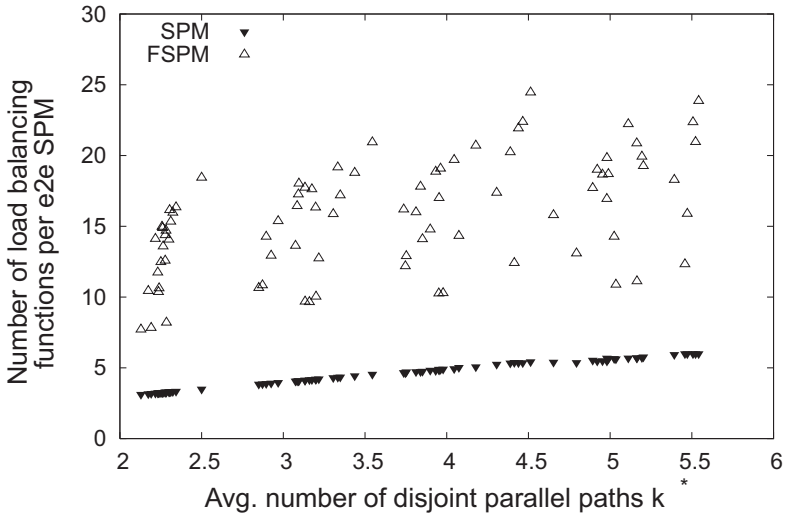

Figure 3: The FSPM requires more traffic distribution function per ingress-egress pair than the normal SPM.

4.2.2 Computation Time for the Solution of the Optimization Problem

We have experimented with several LP-solvers such as the GNU Linear Programming Kit [28], the BPMPD [29], the LPsolve [30], and the CLP [31] and the latter one turned out to yield the fastest program and the smallest program sizes for the class of our optimization problems [13]. Figure 4 shows the computation times averaged over all 420 random networks. The computation time of the optimization process depends primarily on the network size in terms of links, but also on the number of nodes in the network, which explains the jerky curves. The computation time for the FSPM is in all cases about 16% larger than the one for the normal SPM which is actually surprisingly little since a multiple number of traffic distribution functions is optimized for the FSPM.

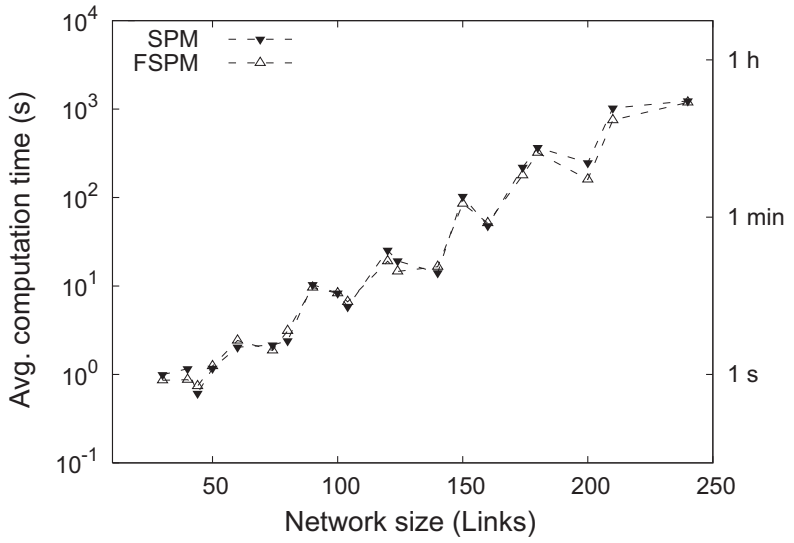

Figure 4: Average computation time for the optimization of the traffic distribution function for the normal SPM and the FSPM depending on the network size.

5 Conclusion

We have reviewed the self-protecting multipath (SPM) that distributes its traffic over several disjoint paths. As it is a protection switching mechanism, it has a load distribution function for the failure of each of its path. These load distributions provide degrees of freedom that can be used to minimize the maximum link utilization in failure scenarios [13]. In this paper we proposed and investigated an obvious extension of the SPM with more flexibility: the failure-specific SPM (FSPM) provides a traffic distribution function for each single failure that affects its disjoint paths.

We have described a linear program formulation both for the SPM and the FSPM for the following problem. A network topology together with its link capacities, the disjoint paths of the SPMs, and a traffic matrix are given; the traffic distribution functions of all SPMs must be optimized such that the maximum link utilizaiton is as small as possible in all protected single link and node failure scenarios $\mathcal{S}$. The solution provides an optimal configuration of all the end-to-end SPMs in the network. We provided numerical results for 420 sample networks and showed that 50 to 200% more protected traffic can be transported with the SPM than with shortest path rerouting (SPR) in networks with a sufficiently high connectivity. The FSPM lead in just 4 out of 420 networks to only marginal capacity savings. The computation time for the optimization program of the FSPM takes on average 16% longer than for the normal SPM. In addition, the FSPM needs to store up to 5 times more traffic distribution functions than the SPM and it requires the signaling of the exact failure location.

After all, we do not recommend the use of the FSPM for application in practice. However, this is good news for the SPM: Its efficiency is so high that it cannot be improved by complex extensions.

References

[1] M. Menth, A. Reifert, and J. Milbrandt, "Self-Protecting Multipaths - A Simple and Resource-Efficient Protection Switching Mechanism for MPLS Networks," in *3rd IFIP-TC6 Networking Conference (Networking)*, (Athens, Greece), pp. 526 – 537, May 2004.

[2] A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, and C. Diot, "IGP Link Weight Assignment for Transient Link Failures," in *18th International Teletraffic Congress (ITC)*, (Berlin), Sept. 2003.

[3] B. Fortz and M. Thorup, "Robust Optimization of OSPF/IS-IS Weights," in *International Network Optimization Conference (INOC)*, (Paris, France), pp. 225–230, Oct. 2003.

[4] A. Autenrieth and A. Kirstädter, "Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS," *IEEE Communications Magazine*, vol. 40, pp. 50–57, Jan. 2002.

[5] P. Pan, G. Swallow, and A. Atlas, "RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels," May 2005.

[6] M. Shand and S. Bryant, "IP Fast Reroute Framework." `http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-06.txt`, Oct. 2006.

[7] A. Atlas and A. Zinin, "Basic Specification for IP Fast-Reroute: Loop-free Alternates." http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-spec-base-06.txt, Mar. 2007.

[8] S. Bryant, S. Previdi, and M. Shand, "IP Fast Reroute Using Not-via Addresses." http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-notvia-addresses-00.txt, Dec. 2006.

[9] M. Gjoka, V. Ram, and X. Yang, "Evaluation of IP Fast Reroute Proposals," (Bangalore, India), Jan. 2007.

[10] M. Menth, *Efficient Admission Control and Routing in Resilient Communication Networks*. PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland, July 2004.

[11] R. Martin, M. Menth, and M. Hemmkeppler, "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), Oct. 2006.

[12] R. Martin, M. Menth, and M. Hemmkeppler, "Accuracy and Dynamics of Multi-Stage Load Balancing for Multipath Internet Routing," in *IEEE International Conference on Communications (ICC)*, (Glasgow, Scotland, UK), June 2007.

[13] R. Martin, M. Menth, and U. Spoerlein, "Integer SPM: Intelligent Path Selection for Resilient Networks," in *IFIP-TC6 Networking Conference (Networking)*, (Atlanta, GA, USA), May 2007.

[14] M. Menth, R. Martin, A. M. C. A. Koster, and S. Orlowski, "Overview of Resilience Mechanisms Based on Multipath Structures," in *International Workshop on Design of Reliable Communication Networks (DRCN)*, (La Rochelle, France), Oct. 2007.

[15] K. Murakami and H. S. Kim, "Optimal Capacity and Flow Assignment for Self–Healing ATM Networks Based on Line and End-to-End Restoration," *IEEE/ACM Transactions on Networking*, vol. 6, pp. 207–221, Apr. 1998.

[16] B. Fortz, J. Rexford, and M. Thorup, "Traffic Engineering with Traditional IP Routing Protocols," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 118–124, 2002.

[17] S. Köhler and A. Binzenhöfer, "MPLS Traffic Engineering in OSPF Networks - A Combined Approach," in 18^{th}*International Teletraffic Congress (ITC)*, (Berlin, Germany), Sept. 2003.

[18] 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.

[19] M. Pióro and D. Medhi, *Routing, Flow and Capacity Design in Communication and Computer Networks*. Morgan and Kaufman, June 2004.

[20] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Norwell, MA, USA: Kluwer Academic Publishers, 1999.

[21] J. Q. Hu, "Diverse Routing in Optical Mesh Networks," *IEEE/ACM Transactions on Networking*, vol. 51, no. 3, pp. 489 – 494, 2003.

[22] M. Menth, J. Milbrandt, and A. Reifert, "Sensitivity of Backup Capacity Requirements to Traffic Distribution and Resilience Constraints," in 1^{st}*Conference on Next Generation Internet Design and Engineering (NGI)*, (Rome, Italy), Apr. 2005.

[23] B. M. Waxman, "Routing of Multipoint Connections," *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 9, pp. 1617–1622, 1988.

[24] E. W. Zegura, K. L. Calvert, and M. J. Donahoo, "A Quantitative Comparison of Graph-Based Models for Internet Topology," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 770–783, 1997.

[25] C. Jin, Q. Chen, and S. Jamin, "t'Inet: Internet Topology Generator," Tech. Rep. CSE-TR-433-00, Department of EECS, University of Michigan, USA, 2000.

[26] A. Medina, I. Matta, and J. Byers, "BRITE: An Approach to Universal Topology Generation," in *International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, (Cinncinnati, Ohio, USA), Aug. 2001.

[27] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network Topology Generators: Degree-Based vs. Structural," in *ACM SIGCOMM*, Aug. 2002.

[28] The Free Software Foundation, “The GNU Linear Programming Kit (GLPK) Version 4.8.” http://www.gnu.org/software/glpk/glpk.html, 2005.

[29] C. Mészáros, “BPMPD Web Site.” http://www.sztaki.hu/ meszaros/bpmpd/, 1998.

[30] M. Berkelaar, K. Eikland, and P. Notebaert, “Web Site for lp_solve.” http://groups.yahoo.com/group/lp_solve, 2005.

[31] J. Forrest, “The COIN-OR Linear Program Solver (CLP).” http://www.coin-or.org, 2005.