

Self-Protecting Multipaths - A Simple and Resource-Efficient Protection Switching Mechanism for MPLS Networks

Michael Menth¹, Andreas Reifert², Jens Milbrandt¹

¹ Department of Distributed Systems, Institute of Computer Science,
University of Würzburg, Am Hubland, 97074 Würzburg, Germany
{menth,milbrandt}@informatik.uni-wuerzburg.de

² Institut für Kommunikationsnetze und Rechnersysteme (IKR),
University of Stuttgart, Pfaffenwaldring 47, D-70569 Stuttgart, Germany
reifert@ikr.uni-stuttgart.de

This work was done while the author was with¹

Abstract. In this paper we propose the concept of an end-to-end (e2e) Self-Protecting Multi-Path (SPM) as a protection switching mechanism that may be implemented, e.g., in Multiprotocol Label Switching (MPLS) networks. In case of local outages, resilient networks redirect the traffic from a failed link over an e2e backup path to its destination. In this case, Quality of Service (QoS) can only be provided if sufficient extra capacity is available. If backup capacity can be shared among different backup paths, multi-path routing allows for considerable savings regarding this extra capacity. The SPM consists of disjoint paths that carry the traffic both in normal operation mode and during local outages. If a partial path is affected by a network failure, the traffic is just distributed to the remaining working paths. This structure is easy to configure and the switching to failure mode operation is simple since no signalling is required. Based on analytical results, we show that load balancing of the traffic across the disjoint paths can reduce the required backup capacity significantly. The backup performance depends strongly on the network topology, and the SPM outperforms simple Open Shortest Path First (OSPF) rerouting by far.

1 Introduction

Carrier grade networks can not afford outages due to internal link or router failures that are visible to their customers. Therefore, they require mechanisms to deviate affected traffic aggregates around the outage location. In contrast to Open Shortest Path First (OSPF) rerouting, these mechanisms have to react fast and they have to provide control over redirected traffic. Fast failure detection is achieved by frequently exchanged “Hello” messages and fast reaction is done by switching the traffic onto pre-computed and pre-installed backup paths. This is called protection switching [1]. In contrast, rerouting denotes the convergence of routing protocols in a narrow sense, i.e., reachability information is exchanged and the routing tables are calculated anew. Since we focus only on path layout and load distribution and not on signaling details, we use the terms rerouting and protection switching synonymously in this work.

Traffic rerouting to maintain pure connectivity does not suffice in carrier grade networks since Quality of Service (QoS) must be maintained. Our objectives are resilient networks, i.e., the customer should not perceive an internal outage by service interruptions or degraded QoS due to bottlenecks on backup paths. Therefore, resilient networks need some extra capacity which is the difference between the required network capacity with and without resilience requirements. Extra capacity is needed for backup purposes, however, it is costly and should be small, so we take it as a performance measure in our study.

Many different rerouting approaches have been proposed in the literature [2, 3]. For example, the traffic may be rerouted only locally or to a different end-to-end (e2e) backup path. However, the backup capacity has not been considered. An optimum path layout and load balancing that requires a minimum backup capacity is computed in [4, 5] for a given network topology and traffic matrix. This optimal solution leads to complex multi-paths that may branch and join at interior nodes, i.e. they are hard to configure. Furthermore, it makes the reorganization of unaffected paths necessary in case of a network failure, which imposes heavy signaling load on the network in a critical situation.

The contribution of this paper is the proposal of a new and simple e2e protection switching mechanisms – called Self-Protecting Multi-Path (SPM) – that may be implemented by explicit routing mechanisms like MPLS. We take advantage of the load balancing potential of multi-path forwarding and minimize the required extra capacity by a polynomial-time optimization algorithm. Our multi-path structures are significantly simpler than general multi-paths since they consist only of disjoint paths. Only traffic shifting of affected traffic aggregates onto backup paths is needed. The minimization of the extra capacity is still very effective such that – depending on the network topology – 20% additional transmission capacity is sufficient to provide full resilience against all single node and link failures.

Given this result, resilience can be implemented at lower cost on the network layer than on the physical layer where fault tolerance is achieved by resource duplication. An exception is the concept of p -Cycles [6, 7] which allows for a more economic protection. It also achieves savings in backup capacity by shared protection and implicit multi-path routing in failure cases. The path layout must adhere to physical layer restrictions and its optimization is more difficult than the one for SPM.

The paper is organized as follows. In Section 2 we point out the difference between other routing optimization approaches and our work. In Section 3 we explain the SPM together with its load balancing options to minimize the required extra capacity for network resilience. The numerical results in Section 4 demonstrate the performance of the SPM. Section 5 summarizes this work and gives some outlook on further work.

2 Related Work

This work is about routing optimization and load balancing in a very broad sense. To avoid any confusion, we delimit it from other network optimization studies.

2.1 Routing Paradigms

There are two major forwarding paradigms: destination based forwarding and connection oriented forwarding.

Destination Based Forwarding In pure Internet Protocol (IP) technology, routers identify the corresponding output interface based on the destination address in the packet header according to their routing tables. The routes in IP forwarding are usually set up by means of routing protocols like the Open Shortest Path First (OSPF) protocol [8]. They exchange reachability information associated with link costs based on which the output ports for the shortest paths to certain destinations are computed. By manipulating the link costs, the routing can be influenced which gives room for traffic engineering. Load balancing over multiple paths is possible if several paths to the same destination have equal costs. This Equal Cost Multi-Path (ECMP) is implemented, e.g., in OSPF.

Connection Oriented Forwarding MPLS is a connection oriented switching technology, i.e., traffic is forwarded along virtual connections that build an overlay network. Packets matching a set of attributes in a router create a Forwarding Equivalent Class (FEC). A so-called LSP Ingress Router (LIR) identifies them and groups them together into a single traffic aggregate by assigning the packets a common label on top of their header. This traffic aggregate is forwarded along a Label Switched Path (LSP) to the LSP Egress Router (LER) that pops the label. The intermediate routers of the LSP forward the packets by label swapping corresponding to the information in their label information base (LIB). The LIB holds a table about incoming LSPs that are identified by their ingress interface and their ingress label and maps them to their egress interface and their egress label. In contrast to routing tables, the information in the LIBs is provided at connection setup. At that occasion, the path of an LSP may be determined automatically by routing protocols or it may follow a pre-computed explicit route.

The routing granularity and the forwarding resolution in MPLS is much finer than in IP because the attributes of a FEC may be, e.g., source *and* destination address. Traffic to a same destination may be carried over different paths that have completely different costs by using explicit routes in MPLS. Explicit routing can be mimicked by source routing in IP technology but this is not advisable since it slows down the forwarding speed of routers considerably. In addition, explicit routing along multiple paths is restricted to ECMP. Therefore, connection oriented technologies like MPLS allow for more powerful traffic engineering than destination based forwarding.

2.2 Routing Optimization

A well investigated problem is routing optimization in the presence of limited link capacities to maximize the supportable traffic intensity whose structure is given by a traffic matrix. This is a multi-commodity flow problem and its solution can be implemented, e.g., by LSPs. For IP routing, a similar approach can be done by setting the link cost appropriately such that all traffic is transported through the network and

that the mean and maximum link utilization is minimized [9]. Pure IP and MPLS solutions may also be combined [10]. These approaches require the knowledge of the traffic matrix which is usually not known for best effort traffic. This problem is tackled by [11] presenting a stable closed loop solution using multi-path structures. Load balancing should be done on a per flow basis and not on a per packet basis to avoid packet reordering which has a detrimental effect on the TCP throughput. The hash based algorithm in [12] achieves that goal very well. The authors of [13] present an online solution for routing with resilience requirements. They try to minimize the blocking probability of successive path requests using suitable single-paths as primary paths and backup paths. The backup bandwidth may be shared or dedicated.

Routing with resilience requirements can also be considered under a network dimensioning aspect, i.e. the traffic matrix is given and the link capacities must be set. This problem is trivial without resilience requirements since a suitable bandwidth assignment for the shortest paths is already an optimum solution. It becomes an optimization problem if capacity sharing for backup paths is allowed. The routing must be designed and the capacity must be assigned such that primary paths and shared backup paths require minimal network capacity while the backup mechanisms provide full resilience for a given set of protected failure scenarios. This is fundamentally different from the above problem since both the routing and the link bandwidth are optimized simultaneously. Note that the results of such calculations depend on the capabilities of the applied restoration schemes. The results of [14] can be well implemented since this work applies only single-paths for both primary and backup paths and relocates only affected primary paths. However, they renounce on multi-path routing and load distribution for path restoration purposes. This is especially important in outage scenarios because traffic diverted over several different paths requires only a fraction of the backup capacity on detour links. If backup capacity sharing is allowed, this backup capacity may be used in different failure scenarios by different rerouted traffic aggregates, which leads to increased resource efficiency since less additional resources must be provisioned in the network. In [4, 5] multi-path routing is used. The required network resources are minimized by calculating the optimum path layout and routing independently for each failure scenario. These backup solutions are too difficult for implementation but they present lower bounds for the required backup capacity.

2.3 Restrictions for Path Layout

We explain why the results in [4, 5] can not be implemented as restoration mechanisms and derive technical side constraints for feasible backup solutions. The path layout and the load balancing is calculated for the normal operation mode and for each failure scenario independently and general multi-path structures are allowed. In an outage case, broken paths must be rerouted but aggregates that are not affected by the failure might also need to be shifted to implement the resource minimal solution.

Firstly, the knowledge of the specific location of the failure is required to apply the optimized path layout and load balancing. Therefore, the exact outage information must be propagated to all ingress routers to trigger protection switching for a specific outage scenario. This entails extensive signaling in a critical system state where the reachability is corrupted.

Secondly, the relocation of the paths can not be done simultaneously. Deflecting more paths than necessary might lead to transient overload on some network elements and can be avoided if only broken paths are redirected.

Thirdly, if each connection holds a backup path for each protected failure scenario, a large amount of paths must be pre-installed and administered. This makes the path configuration very complex and the large number of paths is a problem for the state maintenance of today's core network routers.

Fourthly, to keep the fault diagnostics and the reaction to failures simple, the ingress router should be able to detect a failure and to react locally by switching the traffic to another path. With general multi-path structures, paths may fork and join in transit routers. If a partial path fails, the entire multi-path loses some packets and can not be used anymore. Implementing general multi-paths as a superposition of overlapping single-paths prevents that problem because only some partial paths may fail in case of a local outage. However, this increases the number of parallel LSPs and makes the state management more complex. Hence, only disjoint paths should be used to achieve simple fault diagnostics for multi-path forwarding.

Another restriction for path layout are Shared Risk Link Groups (SRLGs) [15–17] which group network elements together that may fail simultaneously with a high probability. For instance, all links originating at the same router fail if the router goes down. SRLGs are motivated by optical networking where a single optical fiber duct accommodates several logically separate links. In our work, we consider only the first scenario and the second one in a trivial way by excluding parallel links. However, we do not take general SRLGs into account because our focus is the performance evaluation of the basic SPM and not its adaptation to SRLGs.

3 Self-Protecting Multi-Path for Simple Protection Switching

The experiments in [4] have shown that e2e protection mechanisms require less backup capacity than local detours because the traffic of failed paths is redirected early at the source avoiding bottlenecks or much backup capacity around the outage region. Therefore, we focus only on e2e protection switching. We use e2e multi-paths routing because it allows for load distribution in failure cases. As outlined above, only multi-path structures consisting of disjoint paths should be applied and only traffic from paths that are affected by a failure should be rerouted.

The basic structure of an e2e Self-Protecting Multi-Path (SPM) for a single e2e aggregate d consists of parallel disjoint paths. We compute them using a k (link and node) Disjoint Shortest Paths (k DSP) algorithm [18, 19] whose calculation is fast. However, it does not take general SRLGs into account, which is a NP hard problem. The SRLGs are not focus of this work but they can be easily integrated into SPMs by substituting the k DSP heuristic by any other calculation yielding a link and node disjoint multi-path. The SPM sends traffic over all its partial paths. If the LIR recognizes that a partial path fails, it simply redistributes the traffic onto the working paths.

Fault tolerance depends on the set of considered failure scenarios $s \in \mathcal{S}$ (s signifies failed links or routers), including the working mode, for which resilience is guaranteed. The LIR is not aware of the exact failure scenario s but loss of light (LoL) or missing

keep-alive or “Hello” messages [1] for a partial path indicate a failure symptom $f_d(s)$. It consists of the failed and working paths of the SPM that carries the traffic aggregate d . For every aggregate d and for every failure symptom $f_d(s)$ a load balancing function I_d^f is configured. If the LIR diagnoses the failure symptom $f_d(s)$, it redistributes the traffic of d according to I_d^f on the working paths. For example, the equal distribution of the traffic aggregate d onto all working paths is a very simple load balancing function.

The traffic matrix specifies the rates for all traffic aggregates. They have to be supported in all protected scenarios \mathcal{S} which entails a lower bound on the link bandwidths. The sum of all link capacities is the required overall capacity. It should be as small as possible because it represents capital or operational costs. The load balancing function provides some degrees of freedom for the minimization of backup capacities. A simple optimization approach is the assignment of a large portion of d to short partial paths and of a small portion of d to long partial paths. Mathematically speaking, we distribute the rate of a traffic aggregate onto the working paths of an SPM reciprocally to the lengths of these paths. The load balancing function of an SPM can also be exactly optimized. In [20] we modelled technical constraints by linear equations and used linear programming as optimization method. As the solution for I_d^f consists of real values, the computation can be performed in polynomial time.

4 Backup Efficiency of Self-Protecting Multi-Paths

In this section we evaluate the performance of the SPM both in example and random networks using homogeneous traffic matrices. The impact of heterogeneous traffic matrices is investigated in [21]. We determine the required network capacity, i.e. the sum of all link bandwidths, which is required to accommodate the traffic matrix without resilience if shortest path routing (OSPF) is used based on the hop count metric. We take it as a reference value since it is a lower bound for the required network capacity. Then we calculate the required capacity for a given protection scheme to meet the resilience requirements. The resulting extra capacity is the performance measure in our studies. Note that this extra capacity is not always used for backup purposes only because protection mechanisms require sometimes longer paths than the shortest paths for normal operation. However, we use the term extra capacity and backup capacity exchangeably since the extra capacity is required to provide resilience with the respective protection mechanism. For resilience purposes, we take all single link and router failures in the set of protected failure scenarios \mathcal{S} into account.

The calculations for the routing and the load balancing were carried out on a Pentium IV 1.5 GHz standard PC and took some seconds for small networks and some minutes for large networks.

4.1 Impact of Path Layout and Load Balancing on the Required Backup Capacity

We investigate the impact of path layout and load balancing for SPM on the backup performance in two test networks. The Lab03 network in Figure 1(b) is taken from the testbed of the KING project [22]. It is a modification of the UUNET in 1994 where

all nodes with a node degree of at most 2 are successively removed. The network in Figure 1(a) is the optical core of the infrastructure in the COST-279 project [23]. The project was part of the “European Co-operation in the Field of Scientific and Technical Research” and concentrated on ultra-high capacity optical transmission networks. We use both networks in our performance evaluation because they have different properties.

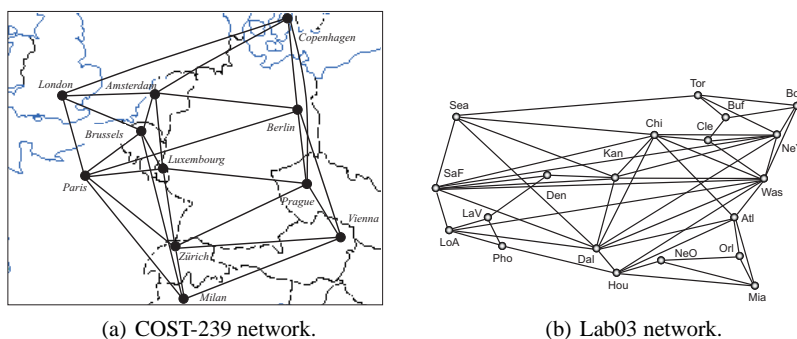
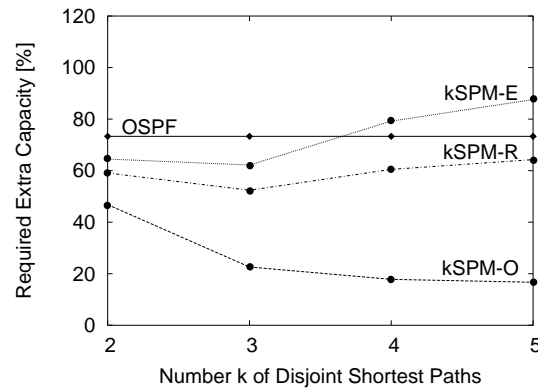


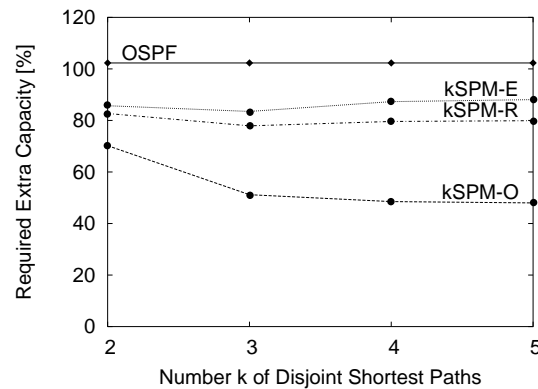
Fig. 1. Test networks.

Figures 2(a) and 2(b) show the backup performance for different SPMs in the COST-239 and in the Lab03 network. The x-axis shows the parameter k for the k DSP calculation for the path layout. The load balancing options are given by different curves. The additional capacity for SPM with equal load balancing is marked by k SPM-E, for SPM with reciprocal load balancing it is marked by k SPM-R, and for SPM with optimized load balancing it is marked by k SPM-O. In addition, the backup capacity for OSPF is given. The SPM require clearly less capacity than OSPF rerouting. The k SPM-O is most economic and its efficiency increases with increasing k . As there are more disjoint backup paths available for larger k , the traffic can be better redistributed in a failure case and less extra capacity is required. The most articulate performance gain is achieved for taking $k = 3$ instead of $k = 2$ disjoint paths. Due to the network topological restrictions, only 4 disjoint paths can be found mostly even for $k = 5$. Therefore, the backup capacity can not be arbitrarily reduced.

In the COST-239 network, the performance of k SPM-E and k SPM-R degrades for increasing k and more extra capacity is needed. The same effect can also be observed to a minor extent in the Lab03 network. If an SPM consist of more disjoint shortest paths, some of them are significantly longer than the shortest one. Their extensive use can not be avoided with k SPM-E or k SPM-R which leads to an increased required network capacity. Hence, SPM with simple load balancing schemes reveal only minor benefits and the optimization is worthwhile.



(a) COST-239 network.



(b) Lab03 network.

Fig. 2. Impact of multi-path routing and load balancing on the backup capacity of SPMs.

4.2 Impact of Network Topology Characteristics

To study the impact of the network topology in more detail, we conduct studies based on random networks and take for 5SPM-O as protection switching mechanism. At first, we describe our algorithm for the construction of random networks. Then we illustrate the impact of the network topology on the backup performance of SPMs both in absolute values and in comparison to the backup performance of OSPF rerouting.

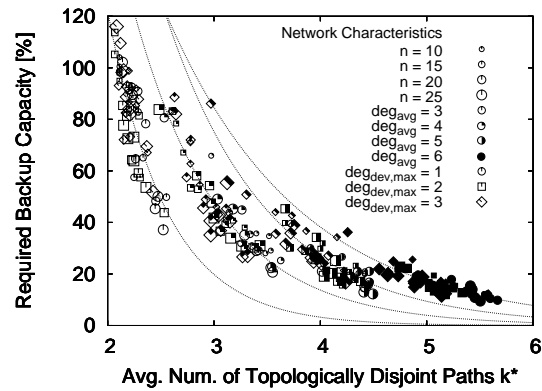
Construction of Random Networks We construct random networks and control some of their essential characteristics. One of them is the degree $deg(v)$ of a node v , which is the number of links v is connected with. We briefly explain our network construction

method that incorporates features of the well know Waxman model [24, 25]. It is an efficient algorithm that provides control over the minimum, the average, and the maximum node degree (deg_{min} , deg_{avg} , deg_{max}), and avoids loops and parallels.

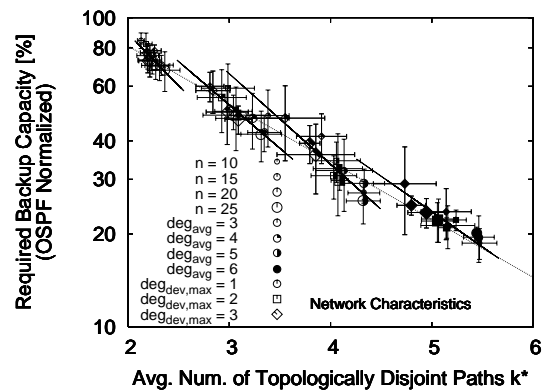
The algorithm starts with an empty link set $\mathcal{E} = \emptyset$ and defines a single arbitrary node $v_{start} \in \mathcal{V}$ connected. Then, $\frac{|\mathcal{V}| \cdot deg_{avg}}{2}$ links are added successively to \mathcal{E} by connecting suitable nodes v_α and v_ω . An arbitrary node v_α is chosen from a set of preferred nodes \mathcal{V}_α with the following properties. All $v \in \mathcal{V}_\alpha$ are connected and have $deg(v) \leq deg_{max}$. If a node $v \in \mathcal{V}$ exists with $deg(v) < deg_{min}$, all $v \in \mathcal{V}_\alpha$ must have $deg(v) < deg_{min}$. The set of potential neighbor nodes \mathcal{V}_ω obeys the following requirements: Loops and parallels must be avoided, i.e. $v_\alpha \notin \mathcal{V}_\omega$ and $(v_\alpha, v_\omega) \notin \mathcal{E}$. Furthermore, if an unconnected node $v \in \mathcal{V}$ exists, all $v \in \mathcal{V}_\omega$ must be unconnected. The node $v_\omega \in \mathcal{V}_\omega$ is chosen according to a probability distribution which depends on v_α and \mathcal{V}_ω . Here, the Waxman model comes into play. Each node has a position in the plane. The Euclidean distance $d(v, w)$ induces a weight $P(v, w) = a \cdot e^{-\frac{d(v, w)}{b \cdot d_{max}}}$ with $d_{max} = \max_{v, w \in \mathcal{V}} d(v, w)$, and $P(v, w)$ produces the probability distribution $p_{v_\alpha}(w) = \frac{P(v_\alpha, w)}{\sum_{v \in \mathcal{V}_\omega} P(v_\alpha, v)}$. Given a maximum node degree deviation deg_{dev}^{max} , the minimum node degree is set to $deg_{min} = \max(deg_{avg} - deg_{dev}^{max}, 2)$ and the maximum node degree is set to $deg_{max} = deg_{avg} + deg_{dev}^{max}$.

Absolute Backup Performance We investigate the required backup capacity for 240 random networks of different size, different average node degree deg_{avg} , and different maximum node degree deviation deg_{dev}^{max} . There are 5 random networks for each topology description. In Figure 3(a), the x-axis indicates the average number of disjoint parallel paths k^* that are found for all source–destination pairs in a network and the y-axis shows the required backup capacity. In general, we observe that the required backup capacity decreases with increasing k^* . We identify four clusters of networks that are marked by dashed lines which are least square interpolations among the points of these clusters according to an exponential function. It turns out that all networks of a cluster have the same average node degree deg_{avg} . The dashed lines make the clusters more visible, however, the extrapolation of those curves does not make sense since deg_{avg} is a trivial upper bound on k^* . Within a cluster, the network size n seems to be irrelevant. A small maximum deviation deg_{dev}^{max} of the node degrees $deg(v)$ from the average node degree deg_{avg} seems to increase k^* , and leads to more efficient backup solutions within a cluster. Therefore, resilience can be achieved at lower cost if the network topology is symmetric.

Backup Performance Relative to OSPF Rerouting Figure 3(b) shows the backup capacity for the same networks in relation to the backup capacity for OSPF rerouting. The OSPF normalization dampens the influence of topological characteristics and shows clearly the benefits of the SPM approach in comparison with conventional rerouting. For all 5 random networks with the same topological characteristics, we build the mean of their k^* and the mean of their ratios of the SPM and OSPF rerouting backup capacity. The horizontal and vertical lines provide the 90% confidence intervals. The data are plotted on a logarithmic scale to make exponential trends better visible.



(a) Absolute backup capacity.



(b) Backup capacity in relation to OSPF routing.

Fig. 3. Required extra capacity for SPM in random networks.

The dashed line is the least square interpolation of all experiments and the solid lines are the interpolations within a cluster of networks with the same average node degree deg_{avg} . The four clusters confirm the above observation that deg_{avg} of a network is strongly correlated with k^* . Increasing the average node degree deg_{avg} shifts the exponential trend slightly towards larger backup capacity. Again, we observe an exponential decay with regard to an increasing k^* , i.e., the superiority of the SPM over OSPF rerouting increases with a larger average number of disjoint paths k^* because SPM reduces the required backup capacity by multi-path forwarding.

5 Conclusion

In this paper we have proposed the Self-Protecting Multi-Paths (SPM) as an e2e protection switching mechanism for MPLS networks. It is used to achieve network resilience in a single autonomous system. We started with an overview of related work and argued for a simple backup solution like the SPM. The SPM carries the traffic of a single e2e traffic aggregate over disjoint parallel paths that may have different length. The traffic of an aggregated may be distributed over these paths according to a load balancing function. If a single partial path fails, the traffic is redirected to the other working paths according to another load balancing function. This action can be performed without signalling across the network because traffic aggregates are only shifted to parallel paths and only aggregates that are affected by a network failure are relocated.

The network capacity dimensioned for OSPF routing without fault tolerance is the reference case in our performance evaluation study. The performance measure for a certain protection switching mechanism is the additional network capacity that is required to achieve resilience. Our experiments showed that the backup performance benefits from the transmission over disjoint parallel paths. Optimized load balancing leads to enormous bandwidth savings for the SPM such that it can provide full resilience against all single link and node failures with less than 17% backup capacity in the COST239 network. This makes failure protection on the network layer significantly cheaper from a resource point of view than on the physical layer if resource doubling is applied.

We constructed random networks and controlled some of their fundamental network characteristics. The amount of required extra capacity depends on the network topology and, in particular, on the average number k^* of disjoint paths in the network whereas the network size has no influence on the required extra capacity. Since OSPF rerouting can also achieve network resilience, we compared the backup capacity required for SPMs and OSPF rerouting. Our simulations revealed that the amount of extra capacity for SPMs decays exponentially with k^* compared to OSPF routing. Only 20% of the OSPF extra bandwidth is needed in suitable networks. In addition, the reaction time of SPM is faster than OSPF rerouting because no reachability information must be exchanged if a failure occurs.

As a challenge remain, e.g., fast heuristics for the calculation of an optimized load balancing are needed for large networks. Suitable network structures are a prerequisite for cheap backup capacities and should be further identified. The optimization of the SPM must be adapted to networks with given link capacities and a structure of their traffic matrix to maximize their throughput while meeting resilience requirements. Moreover, the underlying path computation for the SPMs should be extended towards SRLGs and the impact of multiple failures on the QoS degradation is to be investigated in networks that are resilient against single failures.

References

1. Sharma (Ed.), V., Hellstrand (Ed.), F.: RFC3469: Framework for Multi-Protocol Label Switching (MPLS)-based Recovery. <http://www.ietf.org/rfc/rfc3469.txt> (2003)
2. Pan, P., Gan, D.H., Swallow, G., Vasseur, J.P., Cooper, D., Atlas, A., Jork, M.: Fast Reroute Extensions to RSVP-TE for LSP Tunnels. <http://www.ietf.org/internet-drafts/draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt> (2003)

3. Autenrieth, A., Kirstädter, A.: Engineering end-to-end ip resilience using resilience-differentiated qos. *IEEE Communications Magazine* **40** (2002) 50–57
4. Murakami, K., Kim, H.S.: Comparative Study on Restoration Schemes of Survivable ATM Networks. In: *IEEE INFOCOM'97, Kobe City, Japan* (1997) 345 – 352
5. 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 of Networking* **6** (1998) 207–221
6. Grover, W.D.: Cycle-Oriented Distributed Preconfiguration: Ring-like Speed with Mesh-like Capacity for Self-planning Network Restoration. In: *Proceedings of IEEE ICC '98*. (1998) 537–543
7. Gruber, C.G., Schupke, D.A.: Capacity-efficient Planning of Resilient Networks with p -Cycles. In: *Proceedings of Networks 2002*. (2002) 389–395
8. Moy, J.: RFC2328: OSPF Version 2. <ftp://ftp.isi.edu/in-notes/rfc2212.txt> (1998)
9. Fortz, B., Thorup, M.: Internet traffic engineering by optimizing OSPF weights. In: *IEEE INFOCOM'00*. (2000) 519–528
10. Köhler, S., Binzenhöfer, A.: MPLS traffic engineering in OSPF networks - a combined approach. In: *18th International Teletraffic Congress (ITC18), Berlin* (2003)
11. Gojmerac, I., Ziegler, T., Ricciato, F., Reichl, P.: Adaptive Multipath Routing for Dynamic Traffic Engineering. In: *GLOBECOM'03, San Francisco* (2003)
12. Dittmann, G., Herkersdorf, A.: Network Processor Load Balancing for High-Speed Links. In: *SPECTS 2002, San Diego, CA* (2002) 727–735
13. Kodialam, M.S., Lakshman, T.V.: Minimum Interference Routing with Applications to MPLS Traffic Engineering. In: *Proceedings of IEEE INFOCOM 2000. Volume 2*. (2000) 884–893
14. Iraschko, R.R., MacGregor, M.H., Grover, W.D.: Optimal Capacity Placement for Path Restoration in STM and ATM Mesh-Survivable Networks. *IEEE/ACM Transactions on Networking* **6** (1998) 328 – 336
15. Strand, J., Chiu, A.L., Tkach, R.: Issues For Routing In The Optical Layer. *IEEE Communications Magazine* **39** (2001) 81–87
16. Rajagopalan, B., Luciani, J.V., Awduche, D.O.: IP over Optical Networks: A Framework. <http://www.ietf.org/internet-drafts/draft-ietf-ipo-framework-05.txt> (2003)
17. Kompella, K., Rekhter, Y.: Routing Extensions in Support of Generalized Multi-Protocol Label Switching. <http://www.ietf.org/internet-drafts/draft-ietf-ccamp-gmpls-routing-09.txt> (2003)
18. Suurballe, J.W.: Disjoint Paths in a Network. *Networks* **4** (1974) 125–145
19. Edmonds, J., Karp, R.M.: Theoretical Improvements in the Algorithmic Efficiency for Network Flow Problems. *Journal of the ACM* **19** (1972) 248–264
20. Menth, M., Reifert, A., Milbrandt, J.: Optimization of End-to-End Protection Switching Mechanisms for MPLS Networks. Technical Report, No. 320, University of Würzburg, Institute of Computer Science (2004)
21. Menth, M., Milbrandt, J., Reifert, A.: Sensitivity of Backup Capacity Requirements to Traffic Distribution and Resilience Constraints. Technical Report, No. 322, University of Würzburg, Institute of Computer Science (2004)
22. Hoogendoorn, C., Schrodi, K., Huber, M., Winkler, C., Charzinski, J.: Towards Carrier-Grade Next Generation Networks. In: *ICCT 2003, Beijing, China* (2003)
23. Batchelor et al., P.: Ultra High Capacity Optical Transmission Networks. Final report of Action COST 239. <http://barolo.ita.hsr.ch/cost239/network/> (1999)
24. 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** (1997) 770–783
25. Waxman, B.M.: Routing of Multipoint Connections. *IEEE Selected Areas in Communications* **6** (1988) 1617–1622