Backup Capacity Requirements for MPLS Fast Reroute

Rüdiger Martin and Michael Menth
Department of Distributed Systems, Institute of Computer Science, University of Würzburg
Am Hubland, D-97074 Würzburg, Germany
Phone: (+49) 931-8886651, Fax: (+49) 931-8886632, E-mail:{martin|menth}@informatik.uni-wuerzburg.de

Abstract

MPLS fast reroute (MPLS-FRR) mechanisms deviate the traffic in case of network failures at the router closest to the outage location to achieve an extremely fast reaction time. We review and compare the one-to-one backup and the facility backup concept that are options for MPLS-FRR to deviate the traffic via a detour or a bypass, respectively. Basically, the backup paths can take the shortest path that avoids the outage location from the point of local repair to the tail-end router or to the merge point. We then evaluate the backup capacity requirements and the configuration overhead in terms of the number of backup paths per primary path in a parametric study depending on the network characteristics. While the facility backup concept imposes clearly less configuration overhead than the one-to-one backup, its standard path layout requires more capacity. This can be reduced by a simple modification.

1 Introduction

The operations for multiprotocol label switching (MPLS) fast reroute (MPLS-FRR) mechanisms have been standardized recently by the IETF [1].In case of a network failure, they deviate the traffic at the router closest to the failure location. This can be done by two basically different mechanisms: one-to-one and facility backup. The one-to-one backup deviates the traffic directly from the outage location to its destination while the facility backup just bypasses the traffic around the outage location to repair the original primary paths. The facility backup concept deviates several label switched paths (LSPs) by a single backup path around the failure location while the one-to-one concept needs a private backup path for each LSP. Thus, the facility backup leads to a lower configuration overhead, but it introduces other configuration problems.

The standards provide only the protocol mechanisms to implement a detour or a bypass, but the path layout is not determined. Thus, operators have many degrees of freedom to set up the backup paths. Usually, the default path layout for the backup paths follows the shortest path that avoids the outage location [2]. The authors of [3] suggest a mixed integer linear program (MILP) formulation to find optimum backup paths for the one-to-one mechanism. However, the solution of MILPs is complex and it may be difficult and very time consuming for medium-size or large networks. The authors of [4] present a distributed online algorithm for the one-to-one backup path layout that can be used when LSPs are set up and torn down on demand. It aims at optimized backup capacity sharing depending on the current network state. Simple mechanisms to configure a resilient path layout offline for planned end-to-end demands are still required.

In this paper, we discuss the standard path layout for both one-to-one and facility backup paths. Based on a network dimensioning approach, we consider different outage scenarios, i.e., single router failures only, single link failures only, and single link or router failures. We calculate the required capacity for the standard MPLS-FRR one-to-one and facility backup concepts considering various networks and outage scenarios and compare it to other protection methods.

This paper is structured as follows. Section 2 gives a brief overview on resilience mechanisms in general. Section 3 describes protocol issues specifically for MPLS-FRR, it reviews the default layout of the backup paths and suggests modifications. Section 4 compares the required backup capacity and the required configuration overhead for both options depending on the protected failures. Finally, Section 5 summarizes this work and gives an outlook on further research.

2 Overview on Resilience Mechanisms

In this section we give a brief overview on resilience mechanisms to classify MPLS-FRR. A broader and more complete overview can be found, e.g., in [5]. Resilience mechanisms can be divided into restoration and protection schemes. Restoration sets up a new path after a failure while protection switching pre-establishes backup paths in advance.

2.1 IP Restoration

Usually, restoration is applied by IP rerouting. IP networks have the self-healing property, i.e., their routing re-converges after a network failure by exchanging link state advertisements (LSAs) such that all but the failed nodes can be reached after a while Ű if a working path still exists. In addition, the equal cost multipath (ECMP) option of the most widely used interior gateway routing protocols (IGPs) OSPF [6] and IS-IS [7] distributes traffic over several alternative paths of equal cost to destinations. Another example for restoration besides IP rerouting are backup paths in MPLS that are set up after a network failure.

The disadvantage of such methods is obvious: they are slow. In particular, the interval length to exchange the LSA updates cannot be reduced to arbitrarily small values [8] and the computation of the shortest paths that are needed to construct the routing tables based on the new LSAs requires a substantial amount of time. This time overhead is tolerable for elastic traffic but not for realtime traffic or even high-precision telematic or telesurgery applications. However, the reconvergence of the IP routing algorithm is a very simple and robust restoration mechanism [9, 10].

2.2 Protection Switching Mechanisms

Protection addresses the problem of slow reconvergence speed. It is usually implemented by multiprotocol label switching (MPLS) technology due to its ability to pre-establish explicitly routed backup paths in advance. Depending on the place where the reaction to failures is done, protection switching mechanisms can be distinguished into end-to-end and local protection.

2.2.1 End-to-End Protection Switching

In case of end-to-end protection switching the reaction to a failure along a path is executed at the path ingress router.

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.

Self-Protecting Multipath (SPM) The self-protecting multipath (SPM) [11] consists of disjoint label switched paths (LSPs) and provides at the source several alternatives to forward the traffic to the destination. The traffic is distributed over all alternative paths according to a traffic distribution function (see Figure 1). If one of the paths fails, the traffic is transmitted over the working paths according to another precomputed traffic distribution function. Thus, traffic distribution functions can be optimized a priori to minimize the required backup capacity in the network.

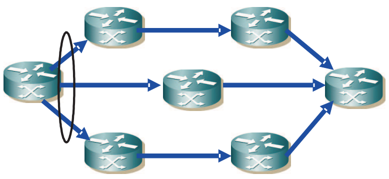

Figure 1: The SPM performs load balancing over disjoint paths according to a traffic distribution function which depends on the working paths.

End-to-end protection switching is faster than restoration methods, but the signalling of the failure to the path ingress router takes time within which traffic is lost.

2.2.2 Local Protection Switching

Local protection schemes tackle 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. Then, a backup path is immediately available if the path breaks at some location. Local protection switching can be implemented by MPLS-FRR [1]. Currently, fast reroute mechanisms are also discussed for IP networks. Several solutions are being discussed, but a preferred method is not yet established [12, 13, 14, 15].

3 Mechanisms for MPLS Fast Reroute

MPLS fast reroute mechanisms protect primary LSPs by local repair methods. A primary LSP is said to be protected at a given hop if it has one or multiple associated backup tunnels originating at that hop. In this work, we want to protect the primary LSP along all intermediate routers of its path. Thus, each intermediate router is a so-called point of local repair (PLR) that serves as head-end router for at least one backup path. MPLS FRR offers two basically different methods for local repair: one-to-one backup and facility backup. In the following, we review these concepts and explain simple standard options for the layout of the backup path.

3.1 Local Repair Options in the MPLS Fast Reroute Framework

We briefly introduce the one-to-one backup and the facility backup together with mandatory conditions regarding the path layout for the protection of link or router failures.

3.1.1 One-to-One Backup Using Detour LSPs

The one-to-one backup sets up a backup path from the PLR to the tail-end of the protected LSP. This backup path is called detour LSP. Each detour LSP protects exactly one primary LSP, but the primary LSP may be protected by several detour LSPs starting at different PLRs. If a detour LSP intersects its protected path further upstream, it may be merged with the primary path at a so-called detour merge point (DMP) to reduce the LSP states in the routers further downstream. However, we disregard this possibility in the following since we focus on the path layout and not on configuration details. Modes are defined in which detour LSPs may contain elements of the protected LSP and others are defined in which such elements are forbidden. In the following, we point out only mandatory constraints to protect against link or router failures.

Link Detour To protect a primary path against a link failure, the router preceding the failed link acts as PLR by redirecting the traffic onto a detour LSP towards the tail-end router r_{tail} of the primary path. The backup path must not contain the failed link, but it may contain the adjacent routers of the failed link. We call this type of backup path $LinkDetour(PLR, r_{tail})$.

Router Detour To protect a primary path against a router failure, the router preceding the failed router acts as PLR by redirecting the traffic onto a detour LSP towards the tail-end router r_{tail} of the primary path. The backup path must not contain the failed router and all its adjacent links. We call this type of backup path $RouterDetour(PLR, r_{tail})$. Note that the primary path cannot be protected against the failure of its head-end or tail-end label switched router (LSR).
Figures 2(a) and 2(b) show that the backup path $LinkDetour(PLR, r_{tail})$ and $RouterDetour(PLR, r_{tail})$ from the same PLR within the same flow can take different shortest paths due to their specific requirements.

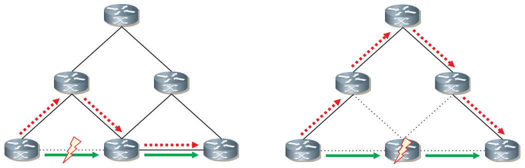
(a) $LinkDetour(PLR, r_{tail})$. (b) $RouterDetour(PLR, r_{tail})$.

Figure 2: One-to-one backup using detours.

3.1.2 Facility Backup Using Bypass LSPs

The facility backup sets up a backup path from the PLR to an upstream router of the protected LSP. This router is called merge point (MP) as it merges the backup path with the protected LSP. Since the backup path bypasses the failure location, it is called bypass LSP. Unlike detour LSPs, a bypass LSP can protect multiple primary LSPs that share the same PLR and MP. In the following, we point out the placement of the MP to protect against link or router failures.

Link Bypass To protect a primary path against a link failure, the router preceding the failed link acts as PLR by redirecting the traffic onto a bypass LSP towards the next hop (NHOP) LSR of the PLR. Thus, the adjacent routers of the link are the head-end and the tail-end LSRs of the bypass LSP which must not contain the failed link. We call this type of backup path $LinkBypass(PLR, NHOP)$.

Router Bypass To protect a primary path against a router failure, the router preceding the failed router acts as PLR by redirecting the traffic onto a bypass LSP towards the next-next hop (NNHOP) LSR of the PLR. Thus, the neighboring routers of the failed router within the primary path are the head-end and the tail-end LSRs of the bypass LSP which must not contain the failed router and all its adjacent links. We call this type of backup path $RouterBypass(PLR, NNHOP)$. Like above, the primary path cannot be protected against the failure of its head-end or tail-end LSR.
The $LinkBypass(PLR, NHOP)$ in Figure 3(a) and the $RouterBypass(PLR, NNHOP)$ in Figure 3(b) from the same PLR within the same flow take different paths due to their specific requirements.

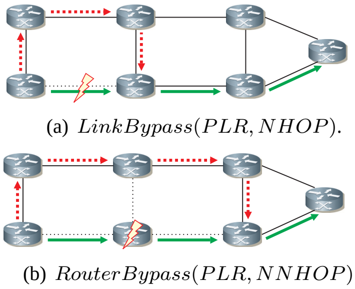
(a) $LinkBypass(PLR, NHOP)$.

(b) $RouterBypass(PLR, NNHOP)$.

Figure 3: Facility backup using bypasses.

3.2 Backup Path Configuration

An intuitive standard approach is characterized by setting up backup LSPs according to the shortest path principle [2]. Each potential PLR, i.e. each intermediate LSR of an LSP, needs a backup path for the protection against the failure of the next link and the next router, respectively. We are trying to assess now the number of required backup paths for the one-to-one backup and the facility backup. We assume n routers and m bidirectional links in the network as well as a fully meshed LSP overlay, i.e., there are $n \cdot (n-1)$ protected LSPs. The length of a specific primary path p is given by $len(p)$ in terms of links and the average number of links per primary path is denoted by $\overline{len}$. The number of adjacent links of router r is given by its node degree $deg(r)$. The average node degree in a network is $deg_{avg} = \frac{2 \cdot m}{n}$

Number of Required Detour LSPs If the one-to-one backup concept uses separate backup paths for the protection against the failure of link and router failures, it requires $len(p)$ link detour LSPs to protect it against all link failures and $len(p)-1$ detour LSPs to protect it against all router failures of the primary path. Thus, $2 \cdot len(p) - 1$ detour LSPs are required altogether for its protection. As a consequence, $n \cdot (n-1) \cdot (2 \cdot \overline{len} - 1)$ detours are needed in the network. The authors of [2] suggest that a link failure can be protected by a $LinkDetour(PLR, r_{tail})$, but it can also be protected by the $RouterDetour(PLR, r_{tail})$. The latter one has just more stringent requirements for the layout of its backup path. Such backup paths exist for all links except the last one within the primary path. Thus, $len-1$ link failures can be protected by a $RouterDetour(PLR, r_{tail})$ and the failure of the last one must be protected by a $LinkDetour(PLR, r_{tail})$. This reduces the number of detours in the network to $n\cdot(n-1)\cdot\overline{len}$ and is the proposed standard path layout for the one-to-one backup concept.

Number of Required Bypass LSPs The network requires $2\cdot m$ link bypasses to protect against the failures of m different links since these backup LSPs can protect multiple primary paths. In addition, router bypass LSPs are needed for the protection against the failure of each of the n routers. We consider a specific router r with $d = deg(r)$ adjacent bidirectional links, from which traffic can be received and to which traffic can be forwarded by that router. If all combinations are possible, $d\cdot(d-1)$ different backup paths are needed to protect possible LSPs carried over r. Thus, $d\cdot(d-1)$ different router bypass LSPs are required for the protection against the failure of this router. As a consequence, a rough guess for the number of required backup path is $2\cdot m+n\cdot deg_{avg}\cdot(deg_{avg}-1)=2\cdot m+2\cdot m\cdot(deg_{avg}-1)=2\cdot m\cdot deg_{avg}=n\cdot deg_{avg}^2$. This expression proposes that considerably fewer bypasses than detours are required to protect the network against all single link and router failures.

4 Performance Comparison

In this section we investigate the performance of the above discussed options for MPLS-FRR by parametric studies regarding different network characteristics. First, we explain our evaluation methodology, then we study the required backup capacity and the number of backup paths per primary path before we compare their efficiency with other well known resilience mechanisms.

4.1 Evaluation Methodology

We explain the network dimensioning approach that we use to calculate the required backup capacity. We also describe the foundation of our parametric study which is based on artificially generated random networks.

4.1.1 Calculation of the Required Backup Capacity

The required backup capacity is the major performance measure in this study. We obtain it as follows for a given network topology, a given traffic matrix, and a given resilience mechanism. The network topology is given by a graph $\mathcal{N} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the set of routers and $\mathcal{E}$ is the set of links. We first compute the capacity $c(l)$ of all links $l \in \mathcal{E}$ in the network that is required to carry the traffic according to the shortest path principle. The sum of these capacities yields the required network capacity $C_\emptyset = \sum_{l\in\mathcal{E}c(l)}$ for the failure-free scenario $\emptyset$. The network must be protected against the failures of a set of failure scenarios $\mathcal{S}$ that contains always the failure-free scenario $\emptyset$. Resilience mechanisms require sufficient backup capacity on the links to carry the traffic in each protected failure scenario. We first determine the link capacity $c(s,l)$ that is required to carry the traffic in each protected failure scenario $s \in \mathcal{S}$ according to the routing which has been changed according to the resilience mechanism. We use it to calculate the required capacity for the resilient network by $C_\mathcal{S} = \sum_{l\in\mathcal{E}} max_{s\in\mathcal{S}}(c(s,l))$. Note that traffic aggregates are inactive in failure scenarios if their source or destination node fails. We express the required backup capacity relative to the capacity needed for shortest path routing by $B = \frac{C_\mathcal{S}-C_\emptyset}{C_\emptyset}$. This methodology can be viewed as a network dimensioning approach. Another option is to calculate, e.g., blocking or QoS violation probabilities for networks with given link capacities. However, we take the network dimensioning approach since it is simpler than the second one, and it is fairer since it grants the capacity to the links where it is needed by the considered resilience mechanism.

4.1.2 Parametric Study

In our parametric study we assume that every network node serves as border router with transit capabilities. We have a fully meshed overlay network and a homogenous traffic matrix. We showed in [16] that the heterogeneity of the traffic matrix has a significant impact on the required backup capacity, but an investigation of this issue in this context is beyond the scope of this paper. We consider three different failure scenarios: all single router failures, all single bidirectional link failures, and all single router and bidirectional link failures. Mainly we use the latter one since 30% of all network failures are due to router failures and 70% are due to link failures [17].

We use sample networks in our study. Most important network characteristics for resilient networks are the network size in terms of nodes $|\mathcal{V}|=n$ and in terms of links $|\mathcal{E}| = m$. They define the average node degree $deg_{avg} = \frac{2\cdot m}{n}$ that indicates the average number of adjacent links of the nodes and is thereby an indirect measure for the network connectivity. In addition, the minimum and the maximum node degree deg_{min} and deg_{max} are also important measures. Since today's well established topology generators cannot control deg_{min} and deg_{max}, we use our own topology generator which is described in [11] and which incorporates features of the well known Waxman model [18, 19]. It allows direct control over n, deg_{avg}, and the maximum deviation deg_{dev}^{max} of the individual node degrees from their predefined average value. It generates connected networks and avoids loops and parallels. We consider networks of size $n \in \{10, 15, 20, 25, 30, 35, 40\}$ nodes with an average node degree $deg_{avg} \in \{3, 4, 5, 6\}$ and a maximum deviation from the average node degree of $deg_{dev}^{max} \in \{1, 2, 3\}$. We generate 5 networks of each combination randomly. This leads in sum to 420 sample networks. For each of them we calculate the required backup capacity for each of the following resilience mechanisms.

- To provide only link protection (LP), we used the $LinkDetour(PLR, r_{tail})$ and $LinkBypass(PLR, NHOP)$, respectively; the set of protected failure scenarios comprises only single link failures. *(LP Detour/Bypass)*
- To provide only router protection (RP), we used the $RouterDetour(PLR, r_{tail})$ and $RouterBypass(PLR, NNHOP)$, respectively; the set of protected failure scenarios comprises only single router failures. *(RP Detour/Bypass)*
- To provide link and router protection (LRP), we used both the $LinkDetour(PLR, r_{tail})$ and the $RouterDetour(PLR, r_{tail})$ or the $LinkBypass(PLR, NHOP)$ and $RouterBypass(PLR, NNHOP)$ for the link and router failures respectively; the set of protected failure scenarios comprises both single link and single router failures. *(LRP Detour/Bypass)*
- As an alternative to the previous scenario, we substitute the link backup paths through existing router backup paths whereever possible. *(LRP-SL Detour/Bypass)*

In the following these abbreviations indicate the protected failures and the applied method. Note that LRP-SL Detour and LRP Bypass are standard path layout approaches as proposed in [2].

4.2 Backup Capacity Requirements

We compare the backup capacity requirements depending on the network size for the 8 above defined protection options. The protection option also determines the set of protected failure scenarios for which the networks are dimensioned. Each point in Figure 4 represents the average backup capacity for all 60 networks of a specific size and the respective investigation scenario. The chosen resilience mechanism has a significant impact on the required backup capacity.

We first compare the capacity requirements for LP, RP, LRP, and LRP-SL without differentiation between the one-to-one and the facility backup option where possible. This makes it easy to understand the reasons for the different capacity requirements of both concepts afterwards. For both the one-to-one and the facility backup, router protection only requires the least resources and in particular less resources than link protection only. This is at first counterintuitive since router failures affect also several adjacent links, but there are two reasons to explain that phenomenon. Inactive aggregates whose source or destination failed because of the router failure make the curves for RP increase significantly for small networks in Figure 4. However, this affects only $\frac{2}{n}$ of the entire traffic and, therefore, this effect shrinks with an increasing network size. Another reason for the reduced backup capacity requirements of RP compared to LP is the improved traffic distribution around the outage location. In case of the one-to-one backup concept, the $LinkDetour(PLR, r_{tail})$ backup paths for LP have a single point of local repair (PLR) while the $RouterDetour(PLR, r_{tail})$ backup paths have different PLRs. Thus, the traffic is deviated over a larger number of different links starting from different locations in the network, which is illustrated in Figure 5. As a consequence, the required backup capacity for the same scenario is distributed over a larger number of different links in the network which increases the potential for backup capacity sharing for different scenarios. In case of the facility backup concept, the $LinkBypass(PLR, NHOP)$ backup paths of LP have a single point of local repair (PLR) and a single merge point (MP) at the next hop (NHOP) while the $RouterBypass(PLR, NNHOP)$ backup paths have different PLRs or different MPs at the next next hop (NNHOP) as illustrated in Figure 6. This leads to an even stronger diversification of the rerouted traffic and, therefore, the gap between RP Bypass and LP Bypass is larger than for the detour concept.

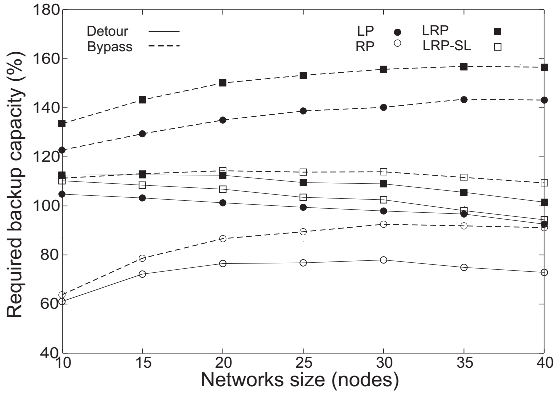

Figure 4: Impact of the network size, the protected failures, and the resilience mechanism on the required backup capacity for MPLS-FRR.

LRP uses the backup paths of both RP and LP and requires clearly more capacity than their maximum. Hence, LP allocates its capacity at different locations compared to RP. As a consequence, the substitution of the $LinkDetour(PLR, r_{tail})$ backup paths through suitable $RouterDetour(PLR, r_{tail})$ backup

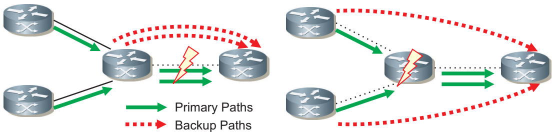

Figure 5: In case of router failures the one-to-one backup concept deviates the traffic from different locations in the network, which leads to a better traffic distribution than in case of link failures.

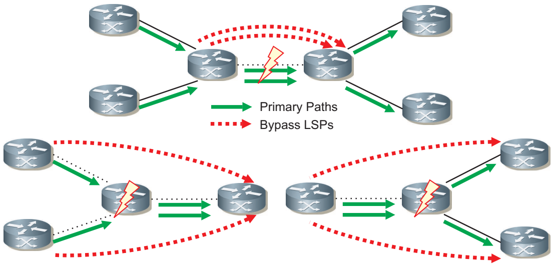

Figure 6: In case of router failures the facility backup concept deviates the traffic from and also to different locations in the network, which leads to an even better traffic distribution.

paths for the one-to-one concept and the substitution of the $LinkBypass(PLR, r_{tail})$ backup paths through suitable $RouterBypass(PLR, r_{tail})$ where possible (LRP-SL) leads to a notable reduction of required backup capacity in Figure 4. However, the last links of the primary paths cannot be protected by suitable router backup paths. This backup capacity reduction technique is very effective for the facility backup. Since one $LinkBypass(PLR, NHOP)$ is used by many primary LSPs, the subsitution of link bypasses through router bypasses increases the traffic spreading and reveals an enormous capacity savings potential.

In all considered investigation scenarios, the facility backup concept always requires more capacity than the one-to-one backup concept. This is clearly due to the reduced capacity sharing potential: it uses a single path to carry the traffic of many affected primary LSPs whose traffic is spread over many links by detour LSPs.

4.3 Configuration Overhead: Number of Backup Paths

As mentioned above, resilience mechanisms differ regarding their configuration overhead. The number of the paths contributes to the number of connection states in the network. Therefore, we compare this measure for the investigated scenarios.

Figure 7 shows the average number of backup-LSPs per primary LSP depending on the network size. For the one-to-one backup concept, the number of backup paths scales with the average path length in the network as follows. The number of backup paths for LP is exactly the average path length. The number of intermediate routers along a path is smaller by one than the number of links and, thus, the number of backup paths for RP is smaller by exactly one than for LP. LRP uses both all link detours from LP and all router detours from RP and, therefore, its number of backup paths is their sum. LRP-SL uses all router detours from RP to substitute appropriate link detours in LP which leads to a protection of link and node failures while keeping the number of backup paths as low as for LP-Standard. For the facility backup concept, the number of backup paths for LP is exactly the number of links m. This yields an average per aggregate of $\frac{m}{n \cdot (n-1)}$ and decreases with the network size. We approximated the number of backup paths for RP as $n \cdot deg_{avg}^2$. The results of our evaluation in Figure 7 reveal that this is an upper bound only. Not all routers serve as transit nodes and not all combinations to tranport transit traffic over the adjacent links of a node are actually used. This effect is extremely strong for small networks where many aggregates are direct connections between neighboring nodes. The number of backup paths for LRP is as before the sum of LP and RP. LRP-SL substitutes link bypasses with router bypasses for all links within a primary LSP except for the last one. Since each link is at least once the last link within the primary LSP that consists of exactly one link connecting neighboring nodes, this does not reduce the number of required backup paths.

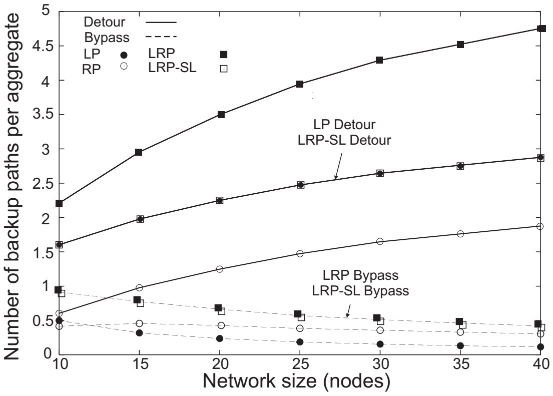

Figure 7: Impact of the network size, the protected failures, and the protection method on the configuration overhead.

The facility backup clearly requires less backup paths than the one-to-one backup since one bypass can protect several primary LSPs. While LRP requires less than one bypass per primary LSP for the facility concept, it requires almost $2-5$ detours per primary LSP for the one-to-one concept.

4.4 Comparison of the Required Backup Capacity for Restoration, End-to-End Protection, and Local Protection

In previous work [11] we investigated the backup capacity requirements for the self-protecting multipath (SPM) and for shortest path rerouting (SPR). We consider them and also equal-cost multipath (ECMP) rerouting for a comparison with both MPLS-FRR concepts regarding the required backup capacity. All single link and node failures are protected.

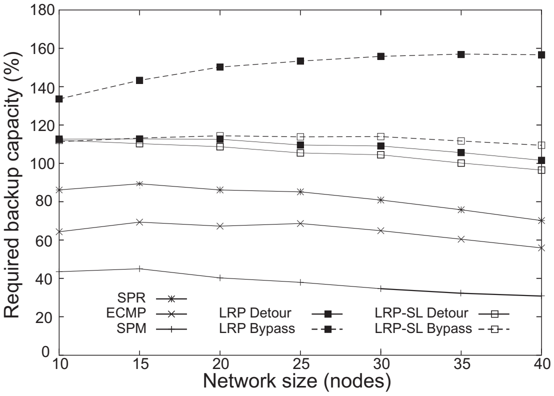

Figure 8: Impact of the network size, the protected failures, and the resilience mechanism on the required backup capacity for restoration, end-to-end protection, and local protection.

Figure 8 shows the averages of their required backup capacity depending on the network size. The SPM requires by far the least capacity, followed by ECMP, and SPR. LRP-SL, proposed as the standard path layout for the one-to-one backup, requires $21-26\%$ more capacity than SPR. LRP, proposed as the standard path layout for the facility backup, requires $47-86\%$ more capacity than SPR. However, MPLS-FRR reacts within tens of milliseconds while SPR as a simple restoration mechanism reacts only within seconds. The reduced configuration overhead of the facility backup concept comes at the expense of additional required capacity if LRP is used as proposed in the standard. Here, LRP-SL helps to reduce the capacity requirements to almost the same level as for the one-to-one backup concept while keeping the configuration overhead low. The SPM seems to be the most attractive resilience mechanism since it requires the least capacity and it is relatively fast as it implements end-to-end protection. However, in contrast to MPLS-FRR, it needs load balancing capabilities and it is not a standardized approach.

5 Summary and Conclusion

We first gave an overview of restoration and protection mechanisms for packet-switched networks. Then we explained the MPLS-FRR framework introduced by the IETF which is intended for fast local protection. It only standardizes protocol signaling issues and the behavior of the label switched routers (LSRs) in case of network element failures. It does not recommend the layout of the backup paths themselves, which is still an open research issue. These backup paths should be short, easy to configure, easy to calculate, and they should require only little additional backup capacity when backup capacity sharing is possible.

The MPLS-FRR framework specifies two different protection types: one-to-one backup using detour LSPs and facility backup using bypass LSPs. We first clarified the requirements for these backup structures and considered a simple mechanism that sets up link and router detours and bypasses, respectively, for all single link and node failures (LP,RP,LRP Detour/Bypass). The backup paths take the shortest paths that avoid the outage location. We further considered a link detour and link bypass substitution (LRP-SL Detour/Bypass) approach.

We evaluated the required backup capacity and the number of backup paths per primary path for the above mentioned backup options. To that end, we conducted a parametric study taking into account 420 artificial networks of different size in terms of nodes, different node degree, and different regularity (deg_{dev}^{max}). Our results show that the facility option in conjunction with the LRP path layout proposed as standard for detours [2] requires more backup capacity than the one-to-one option in conjunction with the LRP-SL path layout proposed as standard for bypasses [2]. This is due to a better distribution of the backup traffic in failure cases achieved by LRP-SL. However, the configuration overhead of the facility backup concept is clearly smaller. If the LRP-SL path layout is also used for the facility option, the backup capacity requirements become nearly as low as for the one-to-one concept. Since the configuration overhead remains low, this is advisable.

Finally, a comparison with other resilience mechanisms showed that the self-protecting multipath (SPM) requires by far the least backup capacity, followed by shortest equal-cost multipath rerouting (ECMP), and single shortest path rerouting (SPR).

Our findings have shown that very simple heuristics for the layout of the backup paths already yield a significant reduction of the required backup capacity which may stimulate more complex and efficient heuristics to obtain further savings. Currently, we work on further simple modifications for the facility and the one-to-one backup that increase the spreading of deviated traffic [20, 21].

References

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

[2] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery*. Elsevier, 2004, pp. 397–422.

[3] H. Saito and M. Yoshida, “An Optimal Recovery LSP Assignment Scheme for MPLS Fast Reroute,” in *International Telecommunication Network Strategy and Planning Symposium (Networks)*, June 2002, pp. 229–234.

[4] D. Wang and G. Li, “Efficient Distributed Solution for MPLS Fast Reroute,” in 4^{rd}*IFIP-TC6 Networking Conference (Networking)*, May 2005.

[5] A. Autenrieth and A. Kirstädter, “Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS,” *IEEE Communications Magazine*, vol. 40, no. 1, pp. 50–57, Jan. 2002.

[6] J. Moy, “RFC2328: OSPF Version 2,” April 1998.

[7] ISO, “ISO 10589: Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service,” 1992.

[8] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot, “Feasibility of IP Restoration in a Tier-1 Backbone,” *IEEE Network Magazine(Special Issue on Protection, Restoration and Disaster Recovery)*, March 2004.

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

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

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

[12] M. Shand and S. Bryant, “IP Fast Reroute Framework,” http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-04.txt, October 2005.

[13] 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-04.txt, July 2005.

[14] G. Schollmeier, J. Charzinski, A. Kirstädter, C. Reichert, K. J. Schrodi, Y. Glickman, and C. Winkler, “Improving the Resilience in IP Networks,” in *IEEE High Performance Switching and Routing (HPSR)*, Torino, Italy, June 2003.

[15] M. Menth and R. Martin, “Network Resilience through Multi-Topology Routing,” in *The* 5^{th} *International Workshop on Design of Reliable Communication Networks*, Island of Ischia (Naples), Italy, Oct. 2005.

[16] 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 Networks Traffic Engineering (NGI)*, Rome, Italy, Apr. 2005.

[17] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, and C.-N. Chuah, “Characterization of Failures in an IP Backbone,” in *IEEE Infocom*, Hongkong, Mar. 2004.

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

[19] 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.

[20] R. Martin, M. Menth, and K. Canbolat, “Capacity Requirements for the Facility Backup Option in MPLS Fast Reroute,” in *IEEE High Performance Switching and Routing (HPSR)*, Poznan, Poland, June 2006.

[21] ——, “Capacity Requirements for the One-to-One Backup Option in MPLS Fast Reroute,” in *IEEE International Conference on Broadband Communication, Networks, and Systems (BROADNETS)*, San José, California, USA, October 2006.