



Regular paper

Communication networks

Efficiency of routing and resilience mechanisms in packet-switched communication networks

Michael Menth*, Rüdiger Martin, Matthias Hartmann, and Ulrich Spörlein

University of Würzburg, Institute of Computer Science, Germany

SUMMARY

In this work we compare the efficiency of various routing and resilience mechanisms. Their path layout determines the utilization of links in the network under normal operation and in failure scenarios. For the comparison, the performance measure is the maximum utilization ρ_S of all links for a set of protected failures \mathcal{S} . A routing mechanism is considered more efficient than another if it leads to a lower maximum link utilization ρ_S . We consider standard and optimized versions of IP routing and rerouting, optimized routing using explicit paths and end-to-end protection switching, as well as standard and optimized versions of MPLS fast reroute. The results show that routing optimization reduces the maximum link utilization significantly both with and without failure protection. The optimization potential for resilient routing is limited by the applied mechanism and depends heavily on the network structure and the set of protected failure scenarios \mathcal{S} . Copyright © 0000 AEIT

1. Introduction

Network failures occur frequently. They lead to end-to-end disconnection and potential overload on backup paths through rerouted traffic. This is not tolerable for customers of Internet service providers (ISPs) and hence service availability and quality of service are crucial parts of service level agreements (SLAs). As a consequence, network providers use protection switching and restoration mechanisms to guarantee service continuation even in the presence of failures.

Operators wish to reduce the risk of overload in a network and minimize QoS violation at lowest possible cost. They want to keep the utilization of their links low without new investments into infrastructure. Hence,

they must make best use of existing network resources. Therefore, routing or resilience mechanism X should carry the traffic on links with sufficient bandwidth to minimize the maximum utilization ρ_S^X of all links and in all failure scenarios \mathcal{S} against which protection is required. We call this set the set of protected failures \mathcal{S} . We use the maximum link utilization ρ_S^X as performance measure in our work since it quantifies the efficiency of a routing or resilience mechanism.

The contribution of this paper is a comprehensive study regarding the efficiency of optimized and non-optimized routing and resilience mechanisms. We look at several variants of IP routing and rerouting, optimized routing based on explicit single paths and end-to-end protection switching (primary/backup paths and self-protecting multipath), and various versions of MPLS fast reroute. A compact overview of the mechanisms under study is given in Section 3.4. We quantify their efficiency,

*Correspondence to: Lehrstuhl für Informatik III, Am Hubland, 97074 Würzburg. E-mail: menth@informatik.uni-wuerzburg.de

the impact of the network topology, and the impact of the set of protected failures \mathcal{S} (e.g. single link and/or node failures).

Section 2 gives an introduction to routing and resilience mechanisms and to optimization objectives. Section 3 explains the resilience mechanisms under study in more detail as well as their path layout. Section 4 compares the efficiency of routing and resilience mechanisms in different network topologies and with different resilience requirements. Finally, we summarize this work and draw conclusions in Section 5.

2. Overview: Routing, Resilience, and Optimization Objectives

In this section we provide a brief overview of routing and resilience mechanisms and discuss two different objectives for routing optimization.

2.1. Routing Mechanisms

Routing determines the layout of the paths in a network. In connectionless networks, e.g. IP networks, traffic is forwarded according to the destination addresses given in the packet headers and the forwarding information given in the forwarding tables of the routers. Thus, modifying the forwarding tables affects the path layout of all paths to a specific destination. This is different in connection-oriented networks like MPLS networks. An explicit path can be set up by adding appropriate per-connection information in the forwarding tables of the intermediate switches. After connection setup, the switches can forward packets according to their connection number and the information given in the forwarding tables. The layout of the paths can be determined either by connectionless routing in the network, e.g. label switched paths (LSPs) in MPLS may be set up on the paths on which IP routing carries the setup messages, or signalling messages are forced to set up the connection along an optimized explicit route that has possibly been calculated offline before.

2.2. Resilience Mechanisms

We now review various classes of wide-spread resilience mechanisms and discuss their pros and cons.

2.2.1. Restoration Mechanisms Restoration mechanisms establish backup paths after a failure has occurred. Therefore, they are too slow to protect traffic of real-time applications [1]. However, they are robust and can survive

multiple network failures. IP routing and rerouting is an example for restoration. They restore the connectivity as long as the network is physically connected. However, restoration can be applied both in connectionless and connection-oriented networks.

2.2.2. End-to-End Protection Switching End-to-end (e2e) protection switching mechanisms can be applied only in connection-oriented networks. They protect primary paths by disjoint backup paths. Both the primary and the backup paths are established upon connection setup. During failure-free operation, traffic is carried on the primary path. When the primary path fails, the head end router switches the traffic to the backup path. E2E protection switching is significantly faster than restoration, but requires link management protocols [2] to recognize path failures. The detection of the failure and the triggering of the failover takes some time during which traffic is still lost. Possibly several backup paths may be used. However, if the primary and all backup paths of a connection simultaneously fail, protection switching can no longer maintain connectivity.

2.2.3. Segment Protection Segment protection mechanisms are also applicable only in connection-oriented networks. They divide a primary path into multiple overlapping segments, each of which is protected by a node-/link-disjoint backup segment. Segment protection is considered to be fast and efficient in terms of backup capacity requirements in optical networks [3]. In a similar way, line- and end-to-end restoration were compared in [4] in the context of ATM networks. In contrast to segment protection, line restoration protects just single links.

2.2.4. Fast Reroute Mechanisms Fast reroute (FRR) mechanisms recognize failures directly at the outage locations and redirect the traffic from there to minimize the reaction time. Multiprotocol label switching (MPLS) offers two options for FRR [5] and, currently, FRR mechanisms are also intensively discussed for IP routing [6,7]. Thus, FRR mechanisms exist for both connectionless and connection-oriented networks.

2.3. Optimization Objectives

The path layout in networks is determined by routing and resilience mechanisms. It can be modified by appropriate configuration which is an important means for traffic engineering. For explicit paths and e2e protection switching, paths are directly computed and provided to the routing system. In IP networks, one can modify administrative link costs based on which least-cost paths

are constructed. Details will be presented in Section 3.1. In the following we discuss two different optimization objectives.

2.3.1. Optimization of Network Configuration In networks with already provisioned link capacities, the risk of congestion should be minimized. Therefore, traffic should be carried on links with sufficient bandwidth. This can be achieved by computing and configuring routing such that the maximum utilization of all links in the network is minimized for an anticipated traffic matrix as this leaves room to compensate traffic fluctuations in situations with increased user activity. The maximum link utilization presents just one objective function but many others are possible. Multiple papers have addressed this problem for non-resilient networks [8–10]. In networks with resilience requirements, routing optimization becomes more complex. The routing and resilience mechanism should be configured in such a way that the maximum utilization ρ_S of all links in the network is minimized during failure-free operation and in all protected failure scenarios \mathcal{S} . Considerably fewer papers and books have addressed this problem for resilient networks [4, 11–14].

2.3.2. Optimization of Network Dimensioning In non-provisioned networks, only the topology and the anticipated traffic matrix are given. Again, traffic should be carried on links with sufficient bandwidth. However, in contrast to above, routing must be computed that the traffic can be carried without QoS violations and link bandwidths must be dimensioned that the installation costs of the network are also minimized [15]. Thus, this requires a joint optimization of network provisioning and routing configuration. It is also known as the network design problem. It is quite hard when link capacities are available only in fixed quantities or capacity costs are non-linear [16] as it is the case in optical networks. Also resilience requirements make this problem hard since the given traffic matrix must be supported for a given set of protected failure scenarios \mathcal{S} [17, 18].

3. Routing and Resilience Mechanisms under Study

In this section, we present the routing and resilience mechanisms we consider in the performance comparison of Section 4. We explain their basic operation and the optimization of their path layout for network configuration as explained in Section 2.3.1. Finally, we present a short overview of the mechanisms under study.

3.1. IP Routing and Rerouting

IP routers forward data packets using destination-based routing using forwarding tables. They map address prefixes to outgoing interfaces. A router determines the appropriate outgoing interface for a packet by a longest prefix match for its destination in the forwarding table. A prefix can be associated with more than one interface if multiple equivalent paths to the destination exist. Single path routing forwards the traffic only to the next hop with the lowest device ID while multi-path routing splits the traffic equally among all possible next hops [19, Section 7.2.7].

The routing tables are usually constructed in a distributed manner by routing protocols like OSPF or IS-IS. They use administrative link costs to calculate the next hops based on least-cost paths which are also called shortest paths when administrative link costs are interpreted as distances. Single shortest path (SSP) routing is default, but we also consider the equal-cost multipath (ECMP) option, which allows multipath routing over all least-cost paths. More precisely, the traffic is equally distributed over all interfaces that are on a shortest path to the destination. ECMP makes the routing independent of device IDs and spreads the traffic over multiple paths which often leads to more balanced link utilizations. In [20] load balancing strategies for ECMP routing are developed and investigated.

A salient feature of IP rerouting is its robustness against network failures. Topology information is broadcast in regular intervals by link state advertisements (LSAs) which implicitly inform all routers about failures. The routing protocols adapt the routing tables to the working topology and restore the connectivity of the network as long as it is physically connected. This rerouting may take seconds, but currently new mechanisms for IP fast rerouting are investigated [6, 7].

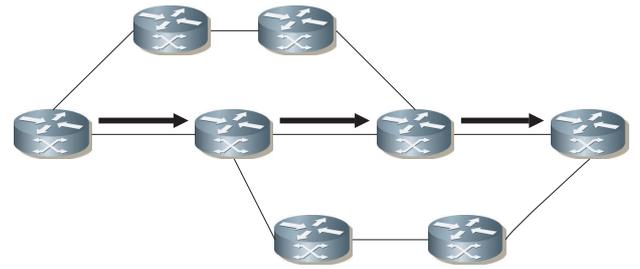
Frequently used configurations of IP routing use either a multiple of the inverse link capacity as virtual link costs or the hop count metric, i.e., the cost for any link is set to 1. However, the link costs can be adjusted by heuristic algorithms in such a way that the maximum link utilization ρ_S of the network is minimized for all protected failure scenarios \mathcal{S} . For the numerical results in Section 4 we use the method from [21] for the optimization of IP link costs both for SSP and ECMP routing. We refer to these options by optSSP and optECMP. Similar objective functions are used in [11–13, 22].

3.2. Routing and End-to-End Protection Switching Using Explicit Paths

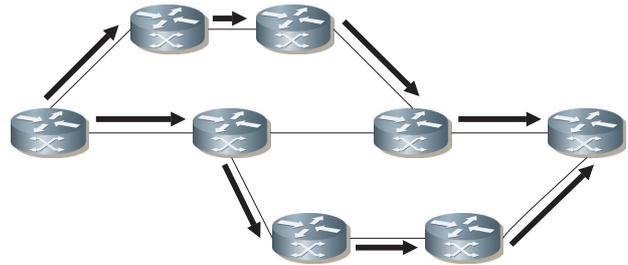
Explicit paths are usually calculated by a path computation element (PCE) or a similar device. Based on a path layout, a connection through the network is set up. Today's packet-switched networks often use label switched paths (LSPs) using MPLS technology for that purpose. We study single explicit paths (SEP) for networks without resilience requirements. For networks with resilience requirements, the simple primary/backup (PB) path concept is considered and the more complex self-protecting multipath (SPM) which is basically a generalization of the primary/backup path concept. First we describe the operation and optimized path layout of the SPM. Then we derive an optimized path layout for primary/backup paths and single explicit paths as a special case of SPM optimizations.

3.2.1. Self-Protecting Multipath (SPM) The self-protecting multipath (SPM) is an e2e protection switching mechanism and can be considered as a generalization of the primary/backup path concept. Its path layout consists of up to k link- or node-disjoint partial paths that can be calculated from a k -disjoint-shortest-path (k -DSP) computation according to [23]. The DSP computation is required since in some "trap topologies" the shortest path prohibits a disjoint backup paths (cf. Figures 1(a) and 1(b)).

The path layout of a 3-SPM is depicted in Figure 2. All partial paths are established during the connection setup. The traffic is distributed over the disjoint paths according to a load balancing function that depends on the pattern of working and broken paths of the SPM. To protect against single failures, the 3-SPM requires 4 different traffic distribution functions: one for the failure-free scenario and one for the failure of each of its partial paths. The traffic distribution functions can be optimized that the maximum link utilization is minimal for a set of protected failure scenarios \mathcal{S} . It is numerically well tractable for networks with a size of up to 60 nodes and can improve the protected throughput to a large extent [24]. However, real load balancing can be problematic due to distribution inaccuracies [25, 26]. Without losing the savings potential of the SPM, heuristics can optimize the load balancing functions of the SPM in such a way that its paths carry either 0% or 100% of the traffic. That means, the iSPM transmits traffic only over a single path both under failure-free conditions and in failure scenarios and the load balancing function acts as a path selection function. These heuristics are very fast and can optimize the SPM for large networks of up to 200 nodes within



(a) Single shortest path routing prohibits the existence of a disjoint backup path.



(b) The disjoint-shortest-path computation finds disjoint paths.

Figure 1. Path layouts in the trap topology.

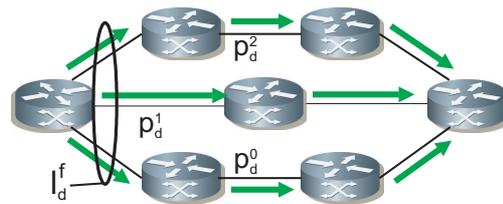


Figure 2. The k -SPM distributes the traffic of a demand d over up to k disjoint paths p_d^0, \dots, p_d^{k-1} according to a traffic distribution function l_d^f which depends on the pattern f of working and non-working paths.

several minutes. We call this method integer SPM (iSPM) [27] and use it as default for the SPM throughout this paper.

3.2.2. Primary/Backup Paths (PB) The simplest form of e2e protection switching is the primary/backup (PB) path concept. A primary and a backup path are established during the connection setup. They are link- or possibly also node-disjoint to protect against single link or single node failures. The 2-iSPM is a good approximation for an optimized primary/backup path concept. Essentially, two disjoint shortest paths are calculated for all ingress-egress

relations of the network and the 2-iSPM optimization tells which of both paths is primary and backup path.

3.2.3. Single Explicit Paths (SEP) The path layout for single explicit paths (SEP) in networks without protection requires for each ingress-egress pair a path such that the maximum link utilization is minimal when the expected traffic matrix is carried. We derive a path layout for SEPs as follows. The *i*-iSPM connects each ingress-egress pair using a disjoint multipath with up to k partial paths. We optimize the path selection function for the failure-free scenario that the maximum link utilization is minimized. Thus, the iSPM selects one path per ingress-egress pair in the failure-free scenario and we use it as a simple approximation of an optimized SEP.

3.3. MPLS Fast Reroute

MPLS fast reroute (MPLS-FRR) is a protection switching mechanism implementing the local repair principle [5]. It provides a point of local repair (PLR) at any router within a label switched path (LSP) such that the traffic can be rerouted at any possible failure location. The advantage of fast reroute methods in general is that PLRs can recognize the failure faster than the head end router of the path and, therefore, the reaction time of fast reroute mechanisms is shorter than the one of e2e protection mechanisms.

Although MPLS supports the setup of explicit paths, LSPs are in practice often set up along the shortest paths of the IP control plane. This is what we assume for the primary paths and also for all backup paths around failed elements in this section. Its advantage compared to explicit paths is that connectivity may be restored after some time through reconvergence when both primary and backup paths fail.

MPLS-FRR offers two backup options that are presented in the following with simple optimization methods. The optimization methods increase the spreading of the backup traffic and decrease thereby the required backup capacity. More efficient path layouts can certainly be found, but they are more complex, require explicit paths, and only a few research papers address this issue [28–31].

3.3.1. One-to-One Backup (Detour) One-to-one backup provides for any path at any PLR a separate backup path that redirects the traffic towards its destination r_{tail} . Figures 3(a) and 3(b) illustrate the standard path layout of these backup paths. They follow the shortest paths from the PLR to the respective destination r_{tail} and avoid the potentially failed elements, i.e. the link and the node after the PLR, because these network elements must not

be contained in the backup paths. The backup paths are called detours. To reduce the complexity of the state maintenance, detour LSPs towards the same destination may be merged to a single LSP when they meet on the way to the destination. However, this does not impact the path layout.

The backup capacity requirements for one-to-one backup can be reduced by modifying the link detours as shown in Figure 3(c). All link detours except for the first link within a path go one hop upstream within the path and then take essentially the router detour at this location [32]. We call this a push-back detour and refer to the optimized one-to-one backup by *optDetour*.

3.3.2. Facility Backup (Bypass) Facility backup provides protection switching for every network element. The standard path layout uses shortest paths without the failed network elements to set up so-called link and router bypasses. Figure 4(a) illustrates a link bypass. A link failure is protected by a backup path around this link, i.e., the backup path starts at the PLR and ends at the next hop (NHOP). This backup path is used as deviation around the failed link for all flows that are usually carried over this link and acts like a tunnel. Similarly, a router failure is protected by a backup path from the PLR to the next next hop (NNHOP) of the respective path (cf. Figure 4(b)). Note that several backup paths are required to protect a single router failure since traffic comes from and leaves for different interfaces of the protected router.

The backup capacity requirements for the facility backup can be reduced by modifying the link backup as follows. Flows use router bypasses instead of link bypasses wherever possible. The last link of a flow is protected by a push-back bypass which is illustrated in Figure 4(c). This backup path sends the traffic one hop upstream and takes the router bypass at this location. If a flow contains only a single link, this one-link path is protected by a normal link bypass [33]. We refer to the improved facility backup by *optBypass*.

3.3.3. Difference between One-to-One and Facility Backup The backup paths for one-to-one backup start at the PLR and end at the tail router of the path while the backup path for facility backup just bypasses a single resource. Figure 5(a) shows that with one-to-one facility backup each potential PLR within a path has its own detour towards the destination. In contrast, Figure 5(b) shows that there is only one bypass from the PLR to the NHOP of the failed element from the perspective of a flow and this bypass is used by multiple flows. The path layout

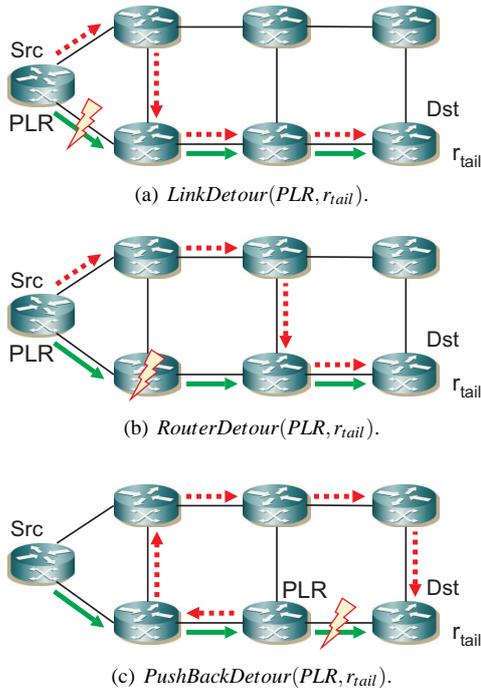


Figure 3. MPLS FRR: one-to-one backup.

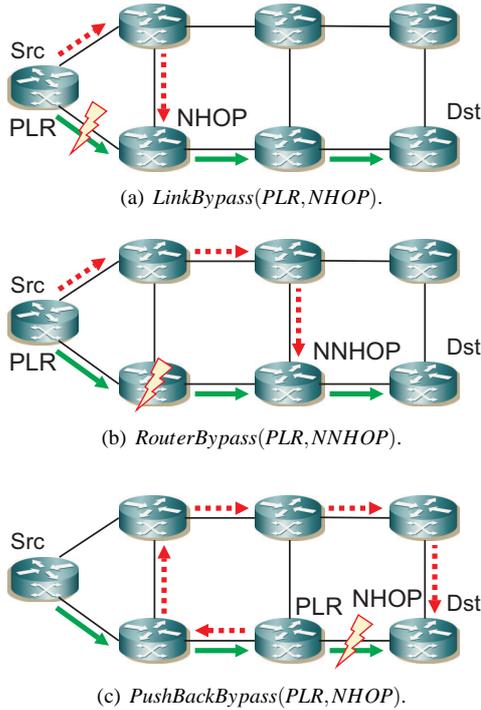


Figure 4. MPLS FRR: facility backup.

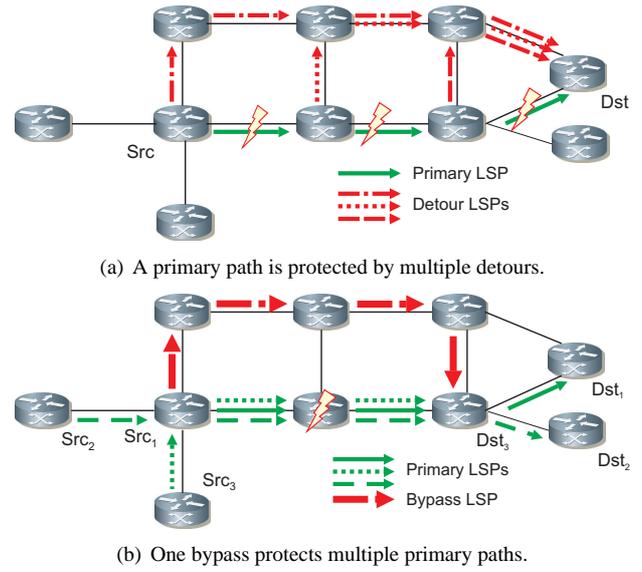


Figure 5. Comparison of MPLS-FRR one-to-one and facility backup.

in Figures 3(a)–3(c) and Figures 4(a)–4(c) seems to be the same for flows using detours and bypasses. However, this is only true if the merge point of the bypass with the unaffected downstream part of the primary path lies on the shortest path from the PLR to the destination in the affected topology. This is in general not true. Then, the path layouts of link and router detours coincide.

3.4. Overview of the Mechanisms under Study

Table 1 summarizes the mechanisms under study. We consider IP routing and rerouting which can send the traffic according to virtual link costs either along single shortest path (SSP) or equally along equal-cost multipaths (ECMP). It is a restoration mechanism and does not require the definition or setup of special backup paths. Optimized path layout is abbreviated by optSSP or optECMP.

The path layout may be completely given in form of explicit paths which can be implemented, e.g., by MPLS. Without resilience requirements, an optimized set of single explicit paths (SEP) is set up between all ingress-egress pairs of a network. With resilience requirements, optimized primary and backup paths (PB) may be used. The integer self-protecting multipath (iSPM) consists of up to k disjoint paths and we set $k = 5$ in our experiments. An optimized path selection function chooses a single partial

Table 1. Overview of routing and corresponding resilience mechanisms under study.

Paradigm	Routing	Restoration / protection
IP routing	SSP/optSSP ECMP/optECMP	SSP/optSSP ECMP/optECMP
Explicit paths	SEP	PB SPM
MPLS FRR	SSP	Detour/optDetour Bypass/optBypass

for transmission of the traffic depending on the pattern of failed and working paths within the multipath structure.

For MPLS fast reroute (MPLS-FRR) we assume that primary paths are set up according to the shortest path principle. With the one-to-one backup option (Detour), the point of local repair (PLR) provides for each primary path a separate detour path along the shortest path in the working topology to its destination. With the facility backup option (Bypass), the PLR bypasses just the failed network element along a shortest path in the working topology using a tunnel. The resource efficiency of both mechanisms can be improved by local modifications of the path layout (optDetour, optBypass).

4. Results

In this section, we first explain the general experiment setup and the performance measure for the subsequent investigations. Then, we assess the relative efficiency of different routing mechanisms by studying the maximum utilization of all links in the failure-free scenario. We extend these experiments towards resilience mechanisms and consider the maximum utilization of all links for all single link failures. We illustrate the impact of the network structure on the ability of different resilience mechanisms to keep the maximum link utilization low. Finally, we show how the set of protected failure scenarios \mathcal{S} influences the maximum link utilization.

4.1. Performance Measure and Experiment Setup

We apply a routing or resilience mechanism X to a network with given link capacities and traffic matrix. If applicable, we perform routing optimization for network configuration (cf. Section 2.3.1). Then, we calculate the maximum utilization ρ_S^X of all links and for all protected failure scenarios \mathcal{S} including the failure-free case. Certainly other performance metrics could be

used: overall capacity consumption in the network, the objective function proposed by Fortz [8] for failure-free conditions and its extension for failure scenarios [11] or modifications thereof. We also experimented with these performance metrics and obtained very similar findings. However, we think that the maximum link utilization is a clearer and more challenging goal, and its extension to failure scenarios is straightforward; therefore, we base the presentation of our results only on this performance metric.

The mechanisms under study $X \in \{\text{SSP, optSSP, ECMP, optSSP, SEP, PB, iSPM, Bypass, Detour, optBypass, optDetour}\}$ have been presented in Section 3 and are summarized in Section 3.4 for quick reference. The protected failure scenarios \mathcal{S} are the failure-free case \emptyset , the set of all single link failures L , the set of all single node failures R , or the set of all single link and node failures LR . If mechanism X can be optimized, we optimize it for the same set of failures \mathcal{S} that is the basis for the calculation of the considered maximum link utilization ρ_S^X .

The maximum link utilization ρ_S^X is an indicator for the absolute efficiency of X with protection of \mathcal{S} . However, the absolute value is not very expressive for comparison purposes as it depends on the link capacities and the traffic matrix. Therefore, we rather consider the *efficiency ratio* $f_S^X(Y) = \rho_S^X / \rho_S^Y$, and compare the relative efficiency of different resilience mechanisms X and Y for the same set of protected failure scenarios \mathcal{S} . The value $f_S^X(Y)$ indicates how much traffic can be transported with routing or resilience mechanism Y in comparison to X while causing the same maximum link utilization. Similarly, we use the efficiency ratio $f_S^X(\mathcal{S}') = \rho_S^X / \rho_{\mathcal{S}'}^X$ and compare the impact of different sets of protected failures \mathcal{S} and \mathcal{S}' on the efficiency of X . Its interpretation is analogous to the one of $f_S^X(Y)$.

Our comparison is based on a large set of random networks. A resilient network topology must be at least 2-connected, i.e., any node in the network can fail without partitioning the topology into disconnected subgraphs. Such structures are found in the core of wide area networks, but usually not in access networks. In typical Internet topologies, the number of links connected to a node, i.e. the node degree, follows a power law distribution as some few core nodes connect many satellite nodes. This, however, does not lead to a resilient network structure. Therefore, we do not use standard topology generators such as BRITTE [34], but we use our own topology generator [35] that generates only at least 2-connected random graphs and also allows to control other network parameters quite strictly. The random networks in our evaluation have a fixed size in terms of nodes $n \in$

$\{10, 15, 20, 25, 30, 35, 40, 45, 50\}$ and a given average node degree $\delta_{avg} \in \{3, 4, 5, 6\}$ which is the fraction $\delta_{avg} = \frac{m}{n}$ of the number of unidirectional links m and the number of nodes n . Furthermore, the degree of individual nodes may deviate by at most $\delta_{dev}^{max} \in \{1, 2, 3\}$ from the average node degree. We use 15 instances of each possible combination which yields 1620 different random networks that were evaluated for each routing or resilience mechanism X . We present the results in a very condensed form that accounts only for the most relevant topological characteristics. More detailed resilience analyzes of specific networks can be obtained using the methodology of [36].

We assume that all links of a network have the same capacity and that the corresponding traffic matrices are homogeneous, i.e., the same traffic rate is exchanged between any two nodes. This is certainly not a realistic assumption since the network capacities are not tailored according to the traffic demands. However, this constitutes difficult networking conditions and serves our goals for several reasons. First, the experimental design is simple and easy to understand. Second, the maximum link utilization ρ_S^X in a network heavily depends on the absolute values of the link capacities and the traffic matrix. However, the efficiency ratios $f_S^X(Y) = \rho_S^X / \rho_S^Y$ or $f_S^X(S') = \rho_S^X / \rho_{S'}^X$ that are used in our evaluations are independent of the scaling of the link bandwidths and traffic matrix. This eliminates the dependency of their absolute values. Third, the problem of badly provisioned transmission capacities challenges the ability of the routing and resilience mechanisms to carry traffic where capacities are and makes differences in this ability more visible.

4.2. Efficiency of Routing Mechanisms without Failure Protection

In this section, we consider networking under failure-free conditions. The routing mechanisms are optimized only for the failure-free scenario and also the maximum link utilization ρ_S^X is calculated only for the failure-free case, i.e., we have ρ_0^X . We compare the efficiency of different routing mechanisms Y relative to the one of standard SSP routing using the efficiency ratios $f_0^{SSP}(Y)$. Figure 6 shows their average from all sample networks depending on the network size. Each point in the figure is an average value from 180 different networks. At first sight, we observe that the efficiency ratios for all routing mechanisms are larger than 1.0, i.e., their maximum link utilization is smaller than the one of SSP routing. Thus, SSP routing is less efficient than the other routing algorithms. Standard ECMP routing is 35–40% better than standard SSP routing

because it leads to a better traffic distribution in the network. Optimized single explicit paths (SEP) are most efficient. They increase the transmission capacity of the network by 60–140% compared to SSP routing and give thereby a lower bound on the potential for optimized IP routing. The efficiency of optimized ECMP routing is similar to the one of SEP for small networks, but for large networks it is about 20% less efficient. Optimized SSP routing is about 10% less efficient than optECMP. Looking at all curves, we realize that the difference among the optimized routing algorithms is clearly visible, but the difference between optimized and unoptimized routing algorithms is larger. Thus, the routing efficiency can be significantly improved by optimization while the choice of the specific routing mechanism is secondary for networks without resilience requirements.

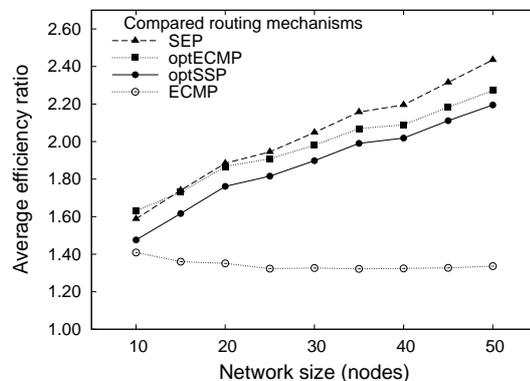
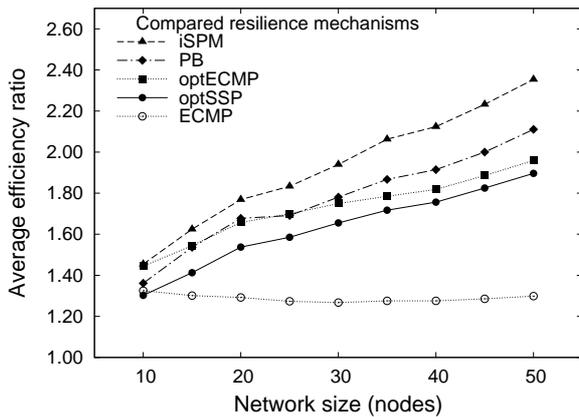


Figure 6. Efficiency ratios $f_0^{SSP}(Y)$ of various routing methods Y compared to default SSP routing without any failure protection.

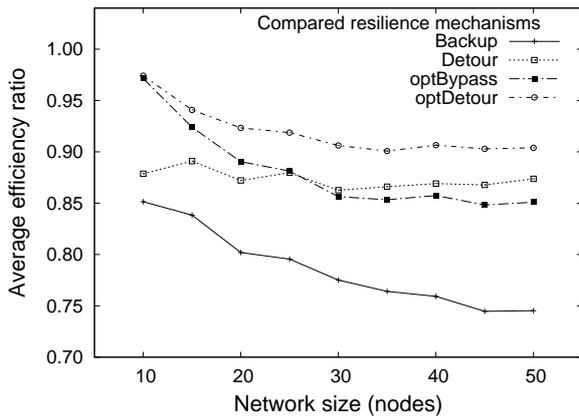
The efficiency of optimized routing mechanisms clearly increases with the network size. We explain that phenomenon in the following. Ideally, link bandwidths are dimensioned for the expected traffic. However, we used equal link bandwidths for our experiments. This leads to mismatches between the bandwidth and the traffic rate on links. As the possibility for strong mismatches increases with the network size, the potential to reduce the maximum link utilization ρ_0^{SSP} by routing optimization also increases. Hence, although random networks are not realistic examples, they help to illustrate how well routing algorithms can exploit increasing optimization potentials.

4.3. Efficiency of Resilience Mechanisms with Protection against Single Link Failures

We conduct the same experiments as in Section 4.2 but now with protection against single link failures. We consider the maximum link utilization during failure-free operation and in all single link failure scenarios and calculate the efficiency ratios $f_L^{SSP}(Y)$ of the resilience mechanism Y relative to SSP (re)routing.



(a) IP restoration and e2e protection switching.



(b) MPLS fast reroute.

Figure 7. Efficiency ratios $f_L^{SSP}(Y)$ of various resilience mechanisms Y compared to standard SSP (re)routing with protection against single link failures.

Figure 7(a) shows the efficiency ratios of the resilience mechanisms that correspond to the routing mechanisms studied in Figure 6. In contrast to Figure 6, the efficiency ratio of a mechanism Y in Figure 7(a) is calculated by

$f_L^{SSP}(Y) = \frac{\rho_L^{SSP}}{\rho_L^Y}$, i.e., single link failures are considered for the maximum link utilization of Y and SSP which serves as a reference. At first sight, Figure 7(a) is very similar to Figure 6 since the qualitative behavior of the efficiency ratios is the same for all mechanisms. However, the efficiency ratios for protection against link failures are about 5–30% lower than without any protection. In large networks, iSPM is about 25% more efficient than optimized primary/backup paths (PB). Thus, iSPM can profit more from the optimization potential than PB since iSPM has more degrees of freedom than PB due to multiple paths. Thus, good traffic distribution in failure cases is very important for efficient routing with resilience requirements. SEP is basically iSPM without failure protection. The gap between iSPM and optimized IP routing with resilience requirements is much larger than the gap between SEP and optimized IP routing without resilience requirements. This shows that the iSPM becomes really efficient and superior to other mechanisms when resilience is required. The efficiency ratios for optimized IP routing (optSSP, optECMP) are about 20% smaller for link protection than without any protection. With link protection, the difference of the efficiency ratios between optimized and unoptimized resilience mechanisms is again very large. The difference in efficiency among optimized mechanisms is larger for link protection than without protection. Thus, the choice of the optimized resilience mechanism does matter.

The path layout for SSP routing and MPLS FRR is the same for the failure-free case but differs for protected failure scenarios. Figure 7(b) shows the efficiency ratios for MPLS FRR mechanisms relative to SSP (re)routing. They are all smaller than 1.0, i.e., the maximum link utilizations for MPLS FRR mechanisms are larger than those for SSP routing. Thus, SSP rerouting is more efficient than MPLS FRR. The standard facility backup (Bypass) has the smallest efficiency ratios between 0.75 and 0.85, followed by the standard one-to-one backup (Detour) with ratios between 0.87 and 0.89. The improved bypass achieves values between 0.85 and 0.97 and the improved detour lies between 0.90 and 0.97. Thus, facility backup requires more backup capacity than one-to-one backup and the improved path layout for both MPLS FRR options leads to significantly larger efficiency ratios. We explain these findings in the following.

With the standard facility backup, the point of local repair (PLR) intentionally redirects all backup traffic over the same bypass tunnel when a link fails. As a consequence, the utilization of the corresponding backup links is very high in that case such that the maximum

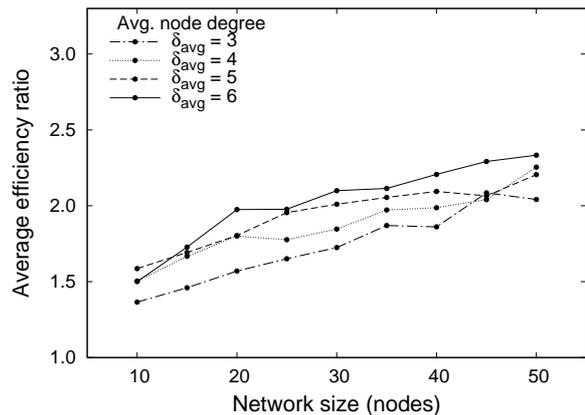
link utilization of SSP routing is exceeded by far. With one-to-one backup, the PLR distributes the traffic over different paths towards the destination. This leads to some distribution of the backup traffic and to lower utilization values of the backup links in failure cases. The improved versions for facility and one-to-one backup differ from the standard versions by the substitution of link bypasses through router bypasses or push-back bypasses and the substitution of link detours through push-back detours. These mechanisms lead to a better distribution of the backup traffic and, thereby, to a lower utilization on the backup links in failure cases. Similar results in a different context can be found in [32, 33]. We considered only simple improvements for MPLS FRR whose paths can still be set up with a distributed routing algorithm similar to IP routing. We expect that explicit paths for primary and backup paths increase the efficiency of MPLS FRR significantly, but they cannot be set up in a distributed manner.

Note that Figures 7(a) and 7(b) do not inform about how much backup capacity is required. This issue is addressed in Section 4.5.

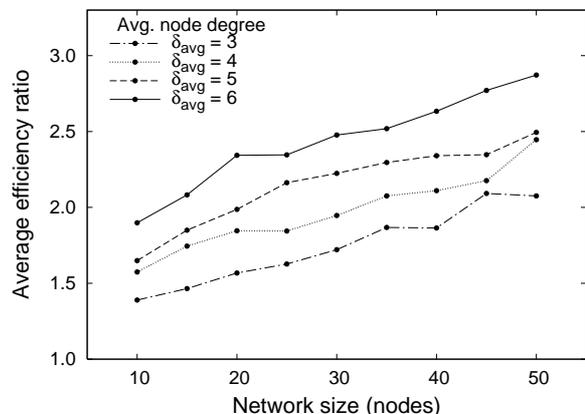
4.4. Impact of the Network Structure on the Efficiency of Routing and Resilience Mechanisms

We study the impact of the network structure and in particular the impact of the node degree on the efficiency in networks with and without resilience requirements.

Figures 8(a) and 8(b) illustrate the efficiency of optimized SSP routing and optimized single explicit paths (SEP) relative to standard SSP routing without protection of any failures. They show that the efficiency ratios increase not only with the network size but also with the average node degree, i.e., highly meshed networks have a larger potential for routing optimization than networks with a rather low average node degree. In sparsely meshed networks, SEP is hardly better than optimized SSP routing. Its optimization essentially selects best paths from a set of disjoint multipaths whose number is low in networks with small node degrees. More disjoint paths can be found in networks with large node degrees which also increases the optimization potential for SEP. Our results show that SEP has clearly larger efficiency ratios than optimized SSP routing under these conditions. Obviously, the network itself has a large optimization potential, but IP routing can take only rather little advantage of highly meshed topologies even with optimization. Reason for that is the destination-based routing principle of IP routing. Traffic for the same destination but from different sources is



(a) Efficiency ratio $f_0^{SSP}(optSSP) = \rho_0^{optSSP} / \rho_0^{SSP}$ for optimized SSP routing.



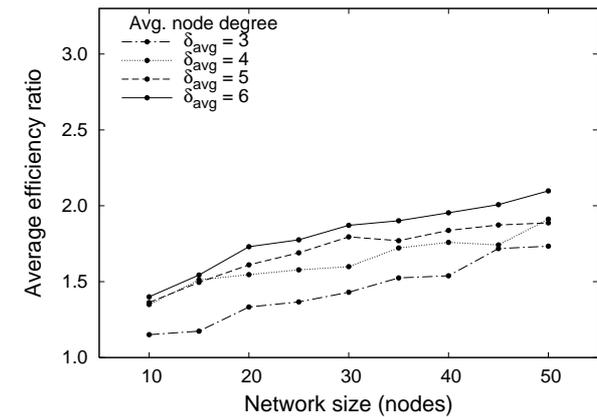
(b) Efficiency ratio $f_0^{SSP}(SEP) = \rho_0^{SEP} / \rho_0^{SSP}$ for SEP.

Figure 8. Efficiency ratios for optimized SSP and SEP relative to unoptimized SSP *without protection of any failures* (\emptyset).

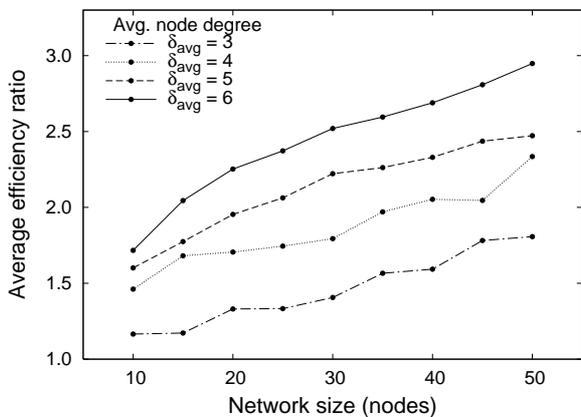
carried on the same downstream path once their paths share a common node. This clusters flows over a few links and increases their utilization.

Figures 9(a) and 9(b) illustrate the efficiency ratios of optimized SSP routing and iSPM compared to default SSP (re)routing with protection against single link failures.

In sparsely meshed networks, optimized SSP routing and iSPM need about the same backup capacity while in well meshed networks, the iSPM is significantly more efficient than optSSP. Obviously, the constraints for destination based routing also prohibit an effective optimization of SSP routing with resilience requirements



(a) Efficiency ratio $f_L^{SSP}(optSSP) = \rho_L^{optSSP} / \rho_L^{SSP}$ for optimized SSP routing.



(b) Efficiency ratio $f_L^{SSP}(iSPM) = \rho_L^{iSPM} / \rho_L^{SSP}$ for iSPM.

Figure 9. Efficiency ratios for optimized SSP and iSPM relative to unoptimized SSP with protection of single link failures (L).

in well meshed networks. Comparing Figures 9(a) and 9(b) with Figures 8(a) and 8(b) we realize that the efficiency ratios are smaller in networks with resilience requirements than in networks without resilience requirements. In addition, the efficiency ratio of iSPM depends more on the average node degree with protection than without protection. In other words, the superiority of explicit paths over IP routing is more visible when protection is required and increases with the average node degree.

4.5. Impact of the Protected Failure Scenarios on the Efficiency of Resilience Mechanisms

In this section, we study the impact of various protected failure scenarios \mathcal{S} on the efficiency of the resilience mechanisms. We use the iSPM and the facility backup option of MPLS FRR as candidates for end-to-end and local protection mechanisms because iSPM is most efficient and facility backup has the least configuration overhead among MPLS FRR options. We consider the following protection variants: no protection (\emptyset), protection against single link failures (L), protection against single router failures (R), and protection against single link and single router failures (LR). We calculate the efficiency ratios $f_0^{iSPM}(Y) = \frac{\rho_Y^{iSPM}}{\rho_0^{iSPM}}$ and $f_0^{Bypass}(Y) = \frac{\rho_Y^{Bypass}}{\rho_0^{SSP}}$ for the protection variants $Y \in \{L, R, LR\}$. We use standard SSP routing as the unprotected baseline for facility backup because standard MPLS FRR also takes the shortest paths in the failure-free scenario. The results are compiled in Figures 10(a) and 10(b).

The curves for L , R , and LR -protection are clearly below 1.0. Networks with protection need some of their capacity to carry backup traffic and lead, therefore, to a larger maximum link utilization than networks without protection. This decreases the efficiency ratios $f_0^Y(X)$ below 1.0 for any protection mechanism Y . For iSPM the efficiency ratios increase with increasing network size from 0.6 to 0.73 and from 0.66 to 0.8 depending on the failure protection. We already observed similar effects for iSPM in Sections 4.2, 4.3, and 4.4, but they were based on a comparison with SSP. In contrast, the results in Figure 10(a) show a comparison with iSPM in the failure-free scenario, i.e. SEP, which is also an optimized mechanism. The reason for the increase of efficiency with the network size is that in large networks backup capacity can be shared among a larger number of protected aggregates that need it in different failure scenarios than in small networks. As a consequence, backup capacity can be shared more effectively in larger networks and, therefore, less backup capacity is required. Figure 10(a) shows also that for the protection against single link failures less backup capacity is required than for the protection against single node failures or single link and node failures. The failure of a node is more severe than the failure of a link as it also implies the failure of its adjacent links. Hence, traffic must be carried by fewer resources than in case of a link failure. This leads to larger link utilization values compared to link failures. Small networks with only 10 nodes are an exception from that rule. When a node fails, traffic from and to that node is removed from the network. Thus, node

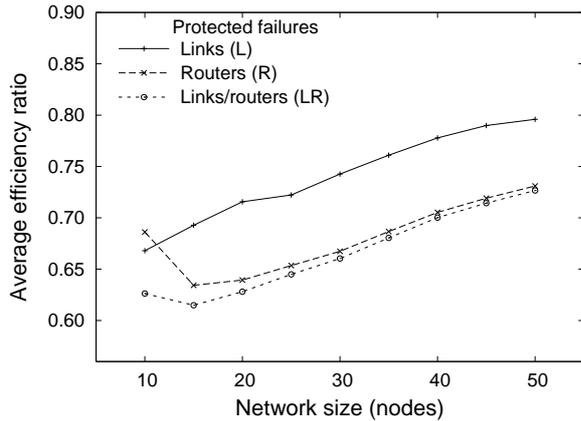
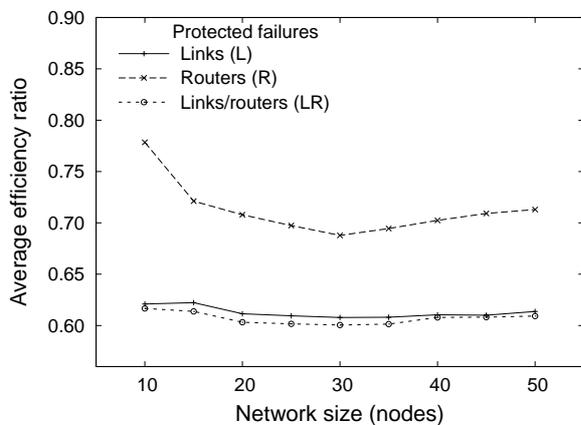
(a) Efficiency ratio $f_0^{iSPM}(Y) = \rho_Y^{iSPM} / \rho_0^{iSPM}$ of iSPM.(b) Efficiency ratio $f_0^{Bypass}(Y) = \rho_Y^{Bypass} / \rho_0^{Bypass}$ of the MPLS facility backup.

Figure 10. Efficiency ratio for iSPM and MPLS FRR facility backup for different protection variants Y relative to the unprotected variant \emptyset .

failures also modify the traffic matrix. The traffic reduction due to single node failures is about 20% in networks with 10 nodes and 13.3% in networks with 15 nodes. This leads to lower maximum link utilizations ρ_R^{iSPM} in case of node failures than in case of link failures (ρ_L^{iSPM}). However, this effect decreases with increasing network size and is, therefore, visible only in small networks.

Figure 10(b) shows the efficiency ratios for MPLS Bypass. They are almost independent of the network size. Thus, unlike iSPM, MPLS FRR cannot take advantage of the increased sharing potential for larger networks. The

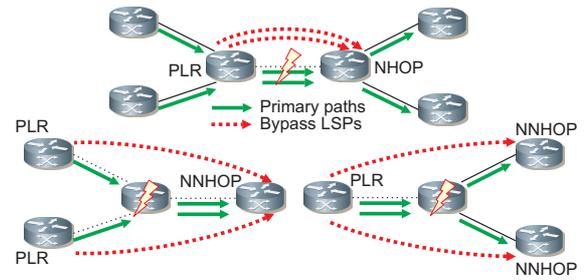


Figure 11. In contrast to link bypasses, router bypasses distribute the traffic over possibly different backup paths to the NNHOP of the PLR.

efficiency ratios are about 0.61 for the protection against single link or single link and node failures and 0.72 for the protection against single node failures. This is rather surprising as iSPM is most efficient with protection against link failures only. With the facility backup, the point of local repair (PLR) intentionally redirects all backup traffic over the same link bypass tunnel when a link fails. As a consequence, the utilization on some backup links is possibly very high in particular failure scenarios such that the maximum link utilization ρ_L^{Bypass} is also very high. The effect of this problem is reduced for router bypasses. They carry the traffic from possibly different PLRs to possibly different NNHOPs. As a consequence, different backup paths are used. This reduces the overall amount of backup traffic on individual links. Figure 11 illustrates this phenomenon.

5. Conclusion

In this work we investigated how well various routing and resilience mechanisms can avoid overload situations under failure-free conditions and in failure cases. For some mechanisms, the path layout is fixed (single shortest path and equal-cost multipath IP routing and rerouting with hop-count metric, MPLS one-to-one and facility backup with standard path layout) while the path layout for some other mechanisms (optimized versions of IP routing and rerouting, optimized single explicit paths, optimized explicit primary and backup paths, and self-protecting multipaths) can be optimized to minimize the maximum utilization of all links in protected failure scenarios \mathcal{S} . Our extensive experiments with 1620 randomly constructed networks showed that optimized routing and resilience mechanisms can carry up to two or even three times more traffic than mechanisms with fixed paths. The potential for improvement obviously depends on the exact setting

including the traffic matrix and the link capacities, but topological characteristics and the set of protected failure scenarios \mathcal{S} have also a significant effect. Moreover, optimized explicit paths (in particular self-protecting multipaths) can better balance traffic in networks than optimized IP routing especially if resilience is required and provided that sufficiently many routing alternatives exist in a network. The results of our study are rather simple to understand and quite intuitive. Nevertheless, to the best of our knowledge this is the first study compiling these important findings and quantifying them with extensive numerical results.

REFERENCES

- G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, C. Diot, Feasibility of IP Restoration in a Tier-1 Backbone, *IEEE Network Magazine* (Special Issue on Protection, Restoration and Disaster Recovery).
- J. Lang (Ed.), RFC4204: Link Management Protocol (LMP) (Oct. 2005).
- C. S. Ou, S. Rai, B. Mukherjee, Extension of Segment Protection for Bandwidth Efficiency and Differentiated Quality of Protection in Optical/MPLS Networks, *Optical Switching and Networking: A Computer Networks Journal* 1 (1) (2005) 19 – 33.
- K. Murakami, 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* 6 (2) (1998) 207–221.
- P. Pan, G. Swallow, A. Atlas, RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels (May 2005).
- S. Rai, B. Mukherjee, O. Deshpande, IP Resilience within an Autonomous System: Current Approaches, Challenges, and Future Directions, *IEEE Communications Magazine* 43 (10) (2005) 142–149.
- A. Raj, O. Ibe, A Survey of IP and Multiprotocol Label Switching Fast Reroute Schemes, *Computer Networks* 51 (8).
- B. Fortz, M. Thorup, Internet Traffic Engineering by Optimizing OSPF Weights, in: *IEEE Infocom*, Tel-Aviv, Israel, 2000, pp. 519–528.
- B. Fortz, J. Rexford, M. Thorup, Traffic Engineering with Traditional IP Routing Protocols, *IEEE Communications Magazine* 40 (10) (2002) 118–124.
- M. Pióro, Á. Szentesi, J. Harmatos, A. Jüttner, P. Gajowniczek, S. Kozdrowski, On Open Shortest Path First Related Network Optimisation Problems, *Performance Evaluation* 48 (2002) 201 – 223.
- B. Fortz, M. Thorup, Robust Optimization of OSPF/IS-IS Weights, in: *International Network Optimization Conference (INOC)*, Paris, France, 2003, pp. 225–230.
- A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, C. Diot, IGP Link Weight Assignment for Transient Link Failures, in: *18th International Teletraffic Congress (ITC)*, Berlin, 2003.
- D. Yuan, A Bi-Criteria Optimization Approach for Robust OSPF Routing, in: *3rd IEEE Workshop on IP Operations and Management (IPOM)*, Kansas City, MO, 2003, pp. 91 – 98.
- M. Pióro, D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*, Morgan Kaufman, 2004.
- C. Pluntke, M. Menth, M. Duelli, CAPEX-Aware Design of Survivable DWDM Mesh Networks, in: *IEEE International Conference on Communications (ICC)*, Dresden, Germany, 2009.
- M. Scheffel, R. G. Prinz, C. G. Gruber, A. Autenrieth, D. A. Schupke, Optimal Routing and Grooming for Multilayer Networks with Transponders and Muxponders, in: *IEEE Globecom*, San Francisco, CA, USA, 2006.
- G. Willems, P. Arijs, W. V. Parys, P. Demeester, Capacity vs. Availability Trade-offs in Mesh-Restorable WDM Networks, in: *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Budapest, Hungary, 2001.
- M. Menth, R. Martin, U. Spoerlein, Network Dimensioning for the Self-Protecting Multipath: A Performance Study, in: *IEEE International Conference on Communications (ICC)*, Istanbul, Turkey, 2006.
- D. Oran, RFC1142: OSI IS-IS Intra-Domain Routing Protocol (Feb. 1990).
- T. W. Chim, K. L. Yeung, K.-S. Lui, Traffic Distribution over Equal-Cost-Multi-Paths, *Computer Networks* 49 (4) (2005) 465–475.
- M. Menth, M. Hartmann, R. Martin, Robust IP Link Costs for Multilayer Resilience, in: *IFIP-TC6 Networking Conference (Networking)*, Atlanta, GA, USA, 2007.
- A. Sridharan, R. Guerin, Making IGP Routing Robust to Link Failures, in: *IFIP-TC6 Networking Conference (Networking)*, Ontario, Canada, 2005.
- R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*, Kluwer Academic Publishers, Norwell, MA, USA, 1999.
- M. Menth, R. Martin, U. Spoerlein, Optimization of the Self-Protecting Multipath for Deployment in Legacy Networks, in: *IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, UK, 2007.
- R. Martin, M. Menth, M. Hemmkepler, 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, 2006.
- R. Martin, M. Menth, M. Hemmkepler, Accuracy and Dynamics of Multi-Stage Load Balancing for Multipath Internet Routing, in: *IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, UK, 2007.
- R. Martin, M. Menth, U. Spoerlein, Integer SPM: Intelligent Path Selection for Resilient Networks, in: *IFIP-TC6 Networking Conference (Networking)*, Atlanta, GA, USA, 2007.
- H. Saito, M. Yoshida, An Optimal Recovery LSP Assignment Scheme for MPLS Fast Reroute, in: *International Telecommunication Network Strategy and Planning Symposium (Networks)*, 2002, pp. 229–234.
- G. Li, D. Wang, C. Kalmanek, R. Doverspike, Efficient Distributed Path Selection for Shared Restoration Connections, in: *IEEE Infocom*, 2002.
- D. Wang, G. Li, Efficient Distributed Solution for MPLS Fast Reroute, in: *4th IFIP-TC6 Networking Conference (Networking)*, Waterloo, Ontario, Canada, 2005, pp. 502 – 513.
- G. Li, D. Wang, R. Doverspike, Efficient Distributed MPLS P2MP Fast Reroute, in: *IEEE Infocom*, 2006.
- R. Martin, M. Menth, K. Canbolat, 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 Jose, CA, USA, 2006.
- R. Martin, M. Menth, K. Canbolat, Capacity Requirements for the Facility Backup Option in MPLS Fast Reroute, in: *IEEE Workshop on High Performance Switching and Routing (HPSR)*, Poznan, Poland, 2006, pp. 329 – 338.
- A. Medina, I. Matta, J. Byers, BRITE: An Approach to Universal Topology Generation, in: *International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, Cincinnati, Ohio, USA, 2001.
- M. Menth, Efficient Admission Control and Routing in Resilient Communication Networks, PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland (July 2004).
- M. Menth, M. Duelli, R. Martin, J. Milbrandt, Resilience Analysis of Packet-Switched Communication Networks, accepted for *IEEE/ACM Transactions on Networking*.