University of Würzburg
Institute of Computer Science
Research Report Series

End-to-End Protection Switching Mechanisms for MPLS Networks

Michael Menth, Jens Milbrandt[1] and Andreas Reifert[2]

Report No. 320 February 2004

[1] Department of Distributed Systems
Institute of Computer Science, University of Würzburg
Am Hubland, D-97074 Würzburg, Germany
phone: (+49) 931-8886644, fax: (+49) 931-8886632
{menth|milbrandt}@informatik.uni-wuerzburg.de

[2] Institut für Kommunikationsnetze
und Rechnersysteme (IKR)
University of Stuttgart, Germany
{reifert}@ikr.uni-stuttgart.de

End-to-End Protection Switching Mechanisms for MPLS Networks

Michael Menth, Jens Milbrandt
Department of Distributed Systems
Institute of Computer Science, University of Würzburg
Am Hubland, D-97074 Würzburg, Germany
phone: (+49) 931-8886644, fax: (+49) 931-8886632
{menth|milbrandt}@informatik.uni-wuerzburg.de

Andreas Reifert
Institut für Kommunikationsnetze
und Rechnersysteme (IKR)
University of Stuttgart, Germany
{reifert}@ikr.uni-stuttgart.de

Abstract

In this paper we propose and investigate several end-to-end (e2e) protection switching mechanisms with application in Multiprotocol Label Switching (MPLS) networks. In case of local outages, the traffic of affected paths is switched to other e2e paths. In this case, Quality of Service can only be provided if sufficient extra capacity is available on the backup paths. If backup capacity can be shared among different backup paths, multi-path routing allows for considerable savings regarding this extra capacity. We propose simple e2e backup mechanisms based on multi-path routing and optimize the load balancing by polynomial-time algorithms to minimize the extra capacity. The mechanisms are simple because the multi-paths consist of disjoint paths that are easy to configure and only traffic of failed paths is switched onto backup paths. Our results illustrate that the savings potential depends on the network topology and that 20% additional resources are sufficient for full resilience against all single router and link failures in well designed networks.

1 Introduction

Carrier grade networks can not afford outages due to internal link or router failures that are visible to their customers. Therefore, they require mechanisms to detour affected traffic aggregates around the outage location. In contrast to IP rerouting, these mechanisms have to react fast and they have to provide control over redirected traffic. Fast failure detection is achieved by frequently exchanged "Hello" messages and fast reaction is done by switching the traffic onto pre-computed and pre-installed backup paths. This is called protection switching [1]. In contrast, rerouting denotes the convergence of routing protocols in a narrow sense. However, as we focus only on the path layout, we use the terms rerouting and protection switching synonymously in this work.

Many different rerouting approaches have been proposed in the literature [2, 3], e.g. the traffic may be rerouted only locally or to a different end-to-end (e2e) backup path, but the backup capacity has not been considered. In [4, 5] the concept of p-cycles is investigated. Traffic rerouting to maintain pure connectivity does not suffice in carrier grade networks since Quality of Service (QoS) must be maintained. Our objectives are resilient networks, i.e. the customer should not perceive an internal outage by service interruptions or degraded QoS due to bottlenecks on backup paths. Therefore, resilient networks need some extra capacity which is the difference between the required network capacity with and without resilience requirements. Extra capacity is needed for backup purposes, however, it is costly and should be small, so we take it as a performance measure in our study. In [6, 7] the optimum path layout and load balancing is computed for a given network topology and traffic matrix. This optimal solution leads to complex multi-paths that may branch and join at interior nodes, i.e. they are hard to configure. Furthermore, it makes the reorganization of unaffected paths necessary in case of a network failure, which imposes heavy signaling load on the network in a critical situation.

The contribution of this paper is the proposal of new simple protection switching mechanisms that may be implemented by explicit routing mechanisms like MPLS. We take advantage of the load balancing potential of multi-path forwarding and minimize the required extra capacity by polynomial-time optimization algorithms. Our multi-path structures are significantly simpler than general multi-paths since they consist only of disjoint paths. Only traffic shifting of affected traffic aggregates onto detour paths is needed. The minimization of the extra capacity is still very effective such that – depending on the network topology and the traffic matrix – 20% additional transmission capacity is sufficient to provide full resilience against all single node and link failures. Given this result, resilience can be implemented at lower cost on the network layer than on the physical layer where fault tolerance is achieved by resource duplication.

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

2 Related Work

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

2.1 Routing Optimization

A well investigated problem is routing optimization in the presence of limited link capacities for a given traffic matrix. This is a multi-commodity flow problem and its solution can be implemented, e.g., by Label Switched Paths (LSPs) in MPLS. For IP routing, a similar approach can be done by setting the link cost appropriately such that all traffic is transported through the network and that the mean and maximum link utilization is minimized [8]. Pure IP and MPLS solutions may also be combined. These approaches require the knowledge of the traffic matrix which is usually not known for best effort traffic. This problem is tackled by [9] presenting a stable closed loop solution using multi-path structures. Load balancing should be done on a per flow basis and not on a per packet basis to avoid packet reordering which has a detrimental effect on the TCP throughput. The hash based algorithm in [10] achieves that goal very well. The authors of [11] present an online solution for routing with resilience requirements. They try to minimize the blocking probability of successive path requests using suitable single-paths as primary paths and backup paths. The backup bandwidth may be shared or dedicated.

Routing with resilience requirements can also be considered under a network dimensioning aspect, i.e. the traffic matrix is given and the link capacities must be set. This problem is trivial without resilience requirements since a suitable bandwidth assignment for the shortest paths is already an optimum solution. It becomes an optimization problem if capacity sharing for backup paths is allowed. The routing must be designed and the capacity must be assigned such that primary paths and shared backup paths require minimal network capacity while the backup mechanisms provide full resilience for a given set of protected failure scenarios. This is fundamentally different from the above problem since both the routing and the link bandwidth are optimized simultaneously. Note that the results of such calculations depend on the capabilities of the applied restoration schemes. The results of [12] can be well implemented since this work applies only single-paths for both primary and backup paths and relocates only affected primary paths. However, they renounce on multi-paths routing and load distribution for path restoration purposes. This is especially important in outage scenarios because traffic diverted over several different paths requires only a fraction of the backup capacity on detour links. If backup capacity sharing is allowed, this backup capacity may be used in different failure scenarios by different rerouted traffic aggregates, which leads to increased resource efficiency since less additional resources must be provisioned in the network. In [6, 7] multi-path routing is used and the required network resources are minimized by calculating the optimum path layout and routing independently for each failure scenario. However, feasible backup solutions require additional technical constraints that are missing in [6, 7] but these results present lower bounds for the required backup capacity.

2.2 Restrictions for Path Layout

We consider the independent path layout calculation based on general multi-paths for the normal operation mode and for each failure scenario like in [6, 7]. We explain why these results can not be implemented as restoration mechanisms and derive technical side constraints for feasible backup solutions. In an outage case, the broken paths are definitely rerouted but paths that are not affected by the failure might also need to be shifted to obtain a resource minimal solution. First, this requires that the information about the specific location of the failure is propagated to all ingress routers to trigger protection switching for a specific outage scenario. This entails extensive signaling in a critical system state at a time for which the long distance connectivity in terms of hops is corrupted. Second, the relocation of the paths can not be done simultaneously. If more paths than necessary are deflected, this might lead to transient overload on some network elements that can be avoided if only broken paths are redirected. Third, if each aggregate connection holds a backup path for each protected failure scenario, a large amount of paths must be pre-installed and administered. This makes the path configuration very complex and the large number of paths is a problem for the state maintenance of today's core network routers. Fourth, to keep the fault diagnostics and the reaction to failures simple, the ingress router should be able to detect a failure and to react locally by switching the traffic to another path. With general multi-path structures, paths may fork and join in transit routers. If a partial path fails, the whole multi-path can not be used anymore. Implementing general multi-paths as a superposition of overlapping single-paths prevents that problem because only some paths may fail in case of a local outage. However, this increases the number of parallel LSPs and makes the state management more complex. Finally, only disjoint paths are left as transport alternatives for multi-paths.

Another restriction for path layout are Shared Risk Link Groups (SRLGs) [13, 14, 15] which group network elements together that may fail simultaneously with a high probability. For instance, all links originating at the same router fail if the router goes down. SRLGs are motivated by optical networking where a single fiber duct accommodates several logically separate links. In our work, we consider only the first scenario and the second one in a trivial way by excluding parallel links. However, we do not take general SRLGs into account as our focus is the investigation of basically different backup mechanisms and not their adaptation to SRLGs.

2.3 Proposal of New Protection Switching Mechanisms

Based on the previous insights, we derive two fundamentally different protection switching mechanisms. As outlined above, only multi-path structures consisting of disjoint paths should be applied and only traffic from paths that are affected by a failure should be rerouted. The experiments in [6] have also shown that e2e protection mechanisms require less backup capacity than local detours because the traffic of the failed paths is redirected early at the source avoiding bottlenecks around the outage region. Therefore, we focus only on e2e protection switching and use multi-path routing that allows for load distribution in failure cases.

The first alternative we propose is e2e path protection for a single-path as primary path and a multi-path as a backup path composed of link or node disjoint paths. We propose two different methods for primary path computation. The second alternative is an e2e self-protecting multi-path that consists of link or node disjoint paths. If one of these paths fails, the traffic is redistributed onto the remaining active paths. In the next section, we describe a suitable path layout computation and an optimization for multi-path load balancing.

3 Optimization

In this section we formulate the optimization problem by linear equations. We use basic notations from linear algebra to represent flows and paths. We describe the problem solutions as linear programs (LPs) that can be solved by standard software like *ILOG CPlex* [16] or the *GNU Linear Programming Kit* [17]. We adapt this formulation to various protection mechanisms.

3.1 Optimum Primary and Backup Path Solution

3.1.1 Basic Notation

Let $\mathbb{X}$ be a set of elements, then $\mathbb{X}^n$ is the set of all n-dimensional vectors and $\mathbb{X}^{n\times m}$ the set of all $n \times m$-matrices with components taken from $\mathbb{X}$. Vectors $\mathbf{x} \in \mathbb{X}^n$ and matrices $\mathbf{X} \in \mathbb{X}^{n\times m}$ are written bold and their components are written as $\mathbf{x} = \begin{pmatrix} x_0 \\ \cdot \\ x_{n-1} \end{pmatrix}$ and $\mathbf{X} = \begin{pmatrix} x_{0,0} & \cdots & x_{0,m-1} \\ \cdot & & \cdot \\ x_{n-1,0} & \cdots & x_{n-1,m-1} \end{pmatrix}$.The scalar multiplication $c\cdot\mathbf{v}$ and the transpose operator $^\top$ are defined as usual. The scalar product of two n-dimensional vectors $\mathbf{u}$ and $\mathbf{v}$ is written with the help of matrix multiplication $\mathbf{u}^\top\mathbf{v} = \sum_{i=1}^{n} u_i \cdot v_i$. Binary operators $\circ \in \{+, -, \cdot\}$ are applied component-wise, i.e. $\mathbf{u} \circ \mathbf{v} = (u_0 \circ v_0, \ldots, u_{n-1} \circ v_{n-1})^\top$. The same holds for relational operators $\circ \in \{<, \leq, =, \geq, >\}$, i.e. $\mathbf{u} \circ \mathbf{v}$ equals $\forall\ 0 \leq i < n{:}\ u_i \circ v_i$. For reasons of simplicity, we define special vectors $\mathbf{0} = (0, \ldots, 0)^\top$ and $\mathbf{1} = (1, \ldots, 1)^\top$ with context specific dimensions.

3.1.2 Links and Nodes

The network $\mathcal{N} = (\mathcal{V}, \mathcal{E})$ consists of $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ unidirectional links that are represented as unit vectors $\mathbf{v_i} \in \{0,1\}^n$ and $\mathbf{e_i} \in \{0,1\}^m$, i.e. $(v_i)_j = \begin{cases} 0 & i \neq j \\ 1 & i = j \end{cases}$ for $0 \leq i, j < n$ and $(e_i)_j = \begin{cases} 0 & i \neq j \\ 1 & i = j \end{cases}$ for $0 \leq i, j < m$. The links are directed and the operators $\alpha(e_i)$ and $\omega(e_i)$ yield the sending and the receiving router of a link. The outgoing and incoming incidence matrices $\mathbf{A}_\alpha$ and $\mathbf{A}_\omega$ describe the network connectivity, i.e. $(a_\alpha)_{ij} = \begin{cases} 0 & \alpha(e_j) \neq v_i \\ 1 & \alpha(e_j) = v_i \end{cases}$ and $(a_\omega)_{i,j} = \begin{cases} 0 & \omega(e_j) \neq v_i \\ 1 & \omega(e_j) = v_i \end{cases}$. The incidence matrix $\mathbf{A} \in \{-1, 0, 1\}^{n\times m}$ is defined as $\mathbf{A} = \mathbf{A}_\omega - \mathbf{A}_\alpha$. The j-th column of $\mathbf{A}$ indicates the source and target of link e_j. The vector $\mathbf{Ae_j}$ yields a node vector. It has a -1 in the i-th row if the source node of e_j is v_i, it has a 1 in the i-th row if the target node of e_j is v_i, and there are zeroes in all other positions. The j-th row of $\mathbf{A}$ indicates the outgoing and incoming links of node v_j. The link vector $\mathbf{v_j}^\top\mathbf{A}$ has a -1 for all outgoing links, a 1 for all incoming links, and zeroes in all other positions. Loops can not be expressed by that formalisms.

3.1.3 Demands, Traffic Matrix, Paths, and Flows

Demands and Traffic Matrix We consider a network with traffic flowing from each ingress router to every other. We call this b2b relationship between $\mathbf{v_i}$ and $\mathbf{v_j}$ a demand $d = (i, j)$ and denote the set of all demands by $\mathcal{D} = \{(i, j) : 0 \leq i, j < n \text{ and } i \neq j\}$. The associated traffic rate is given by $d.c$ and corresponds to an entry in the traffic matrix.

Paths A path p_d of a demand $d \in \mathcal{D}$ between distinct nodes v_α and v_ω is a set of contiguous links represented by a link vector $\mathbf{p_d} \in \{0,1\}^m$. This corresponds to a single-path. However, we usually apply the concept of a multi-path $\mathbf{p_d} \in [0,1]^m$, which is more general since the traffic may be split into several partial paths carrying a real fraction of the traffic. A path follows conservation rules, i.e. the amount of incoming traffic equals the amount of outgoing traffic in a node which is expressed by

$$\mathbf{A}\mathbf{p_d} = (\mathbf{v}_\omega - \mathbf{v}_\alpha). \tag{1}$$

While cycles containing only inner nodes can be easily removed, cycles containing the start or end node of a path are more problematic. Therefore, we formulate a condition preventing this case. The expressions ${\mathbf{v}_\alpha}^\top \mathbf{A}_\omega$ and ${\mathbf{v}_\omega}^\top \mathbf{A}_\alpha$ yield the incoming edges of start node v_α and all outgoing edges of end node v_ω of a path p_d. Hence, cycles containing the start or end node can be prevented if the following equations hold:

$$({\mathbf{v}_\alpha}^\top \mathbf{A}_\omega)\mathbf{p_d} = 0 \text{ and } ({\mathbf{v}_\omega}^\top \mathbf{A}_\alpha)\mathbf{p_d} = 0. \tag{2}$$

Flows Given a cycle-free path p_d, the corresponding flow $d.c \cdot \mathbf{p_d}$ takes the traffic rate into account.

3.1.4 Protected Scenarios

A protected failure scenario is given by a vector of failed nodes $\mathbf{s}_\mathcal{V} \in \{0,1\}^n$ and a vector of failed links $\mathbf{s}_\mathcal{E} \in \{0,1\}^m$. We denote a failure pattern shortly by $\mathbf{s} = \left(\begin{smallmatrix}\mathbf{s}_\mathcal{V}\\ \mathbf{s}_\mathcal{E}\end{smallmatrix}\right)$. The set $\mathcal{S}$ contains all protected outage scenarios including $\mathbf{s} = \mathbf{0}$, i.e. the no failure case.

3.1.5 Traffic Reduction

In normal operation without any failures, all demands $d \in \mathcal{D}$ are active. If routers fail, some demands may disappear. We consider several options.

No Traffic Reduction We assume that failed routers lose only their transport capability for transit flows but are still able to generate traffic. Therefore, we have $\mathcal{D}_\mathbf{s} = \mathcal{D}$.

Source Traffic Reduction An aggregate flow is removed from the traffic matrix if the source node v_i of demand $d = (i,j)$ fails. If a failed node is the destination of a flow, “server push” traffic may still be transported through the network, hence $\mathcal{D}_\mathbf{s} = \mathcal{D} \setminus \{(i,j) : {\mathbf{v_i}}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\}$.

Full Traffic Reduction In contrast to above we assume that the traffic with a failed destination is stalled. An aggregate flow is removed from the traffic matrix if a node fails which is either the source or the destination of a flow, hence $\mathcal{D}_\mathbf{s} = \mathcal{D} \setminus (\{(i,j) : {\mathbf{v_i}}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\} \cup \{(j,i) : {\mathbf{v_i}}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\})$.

3.1.6 Failure Indication Function

The failure indication function $\phi(\mathbf{p},\mathbf{s})$ indicates whether a path p is affected by a failure pattern s [18]. Path p is affected by a link failure pattern $\mathbf{s}_{\mathcal{E}}$ if $\mathbf{s}_{\mathcal{E}}{}^{\top}\mathbf{p} > 0$. To formulate this analogously for node failures we define traces. The α-trace is $\mathbf{tr}_{\alpha}(\mathbf{p_d}) = \mathbf{A}_{\alpha}\mathbf{p_d}$ and the ω-trace is $\mathbf{tr}_{\omega}(\mathbf{p_d}) = \mathbf{A}_{\omega}\mathbf{p_d}$, respectively. We obtain the interior trace $\mathbf{ti}$ by excluding the corresponding end or the start node of the α- or ω-trace, respectively, i.e. $\mathbf{ti}(\mathbf{p_d}) = \mathbf{A}_{\alpha}\mathbf{p_d} - \mathbf{v}_{\alpha} = \mathbf{A}_{\omega}\mathbf{p_d} - \mathbf{v}_{\omega}$. Path p is affected by a node failure pattern $\mathbf{s}_{\mathcal{V}}$ if $\mathbf{s}_{\mathcal{V}}{}^{\top}\mathbf{ti}(\mathbf{p}) > 0$. Finally, the failure indication function is $\phi(\mathbf{p},\mathbf{s}) = \begin{cases} 1 & \mathbf{s}_{\mathcal{E}}{}^{\top}\mathbf{p} + \mathbf{s}_{\mathcal{V}}{}^{\top}\mathbf{ti}(\mathbf{p}) > 0 \\ 0 & \text{otherwise.} \end{cases}$

3.1.7 Protection Alternatives

A path restoration scheme introduces a backup path q_d which is activated if the primary path fails. This backup path protects against link and/or node failures of each primary path p_d depending on the required type of resilience. A backup path q_d is link protecting if

$$\mathbf{q_d}{}^{\top}\mathbf{p_d} = 0 \tag{3}$$

and it is both link and node protecting if the following holds

$$\mathbf{ti}(\mathbf{q_d})^{\top}\mathbf{ti}(\mathbf{p_d}) = 0. \tag{4}$$

3.1.8 Objective Function and Capacity Constraints

We describe the capacity of all links by a vector of edges $\mathbf{b} \in (\mathbb{R}_0^+)^m$. The overall capacity in the network is the objective function that is to be minimized. It can be computed by

$$\mathbf{w}^{\top}\mathbf{b} \to \min \tag{5}$$

where $\mathbf{w} \in (\mathbb{R}_0^+)^m$ is a vector of weights, that is normally set to $\mathbf{w} = \mathbf{1}$. If the connectivity is maintained by a backup path in case of a failure pattern $\mathbf{s} \in \mathcal{S}$, the following bandwidth constraints guarantee that enough capacity is available to carry the traffic generated by the demands $d \in \mathcal{D}_{\mathbf{s}}$.

Bandwidth Reuse In packet switched networks, no resources are physically dedicated to any flows. If traffic is rerouted due to some outages, the resources can be automatically reused for transporting other traffic. Under this assumption, the capacity constraints are

$$\forall \mathbf{s} \in \mathcal{S} : \sum_{d \in \mathcal{D}_{\mathbf{s}}} d.c \cdot ((1 - \phi(\mathbf{p_d},\mathbf{s})) \cdot \mathbf{p_d} + \phi(\mathbf{p_d},\mathbf{s}) \cdot \mathbf{q_d}) \leq \mathbf{b}. \tag{6}$$

No Bandwidth Reuse In optical networks connections are bound to physical resources like fibers, wavelengths, or time slots. If a network element fails, there might not be enough time to free the resource of a redirected connection. This is respected by the following capacity constraints:

$$\forall \mathbf{s} \in \mathcal{S} : \sum_{d \in \mathcal{D}} d.c \cdot \mathbf{p_d} + \sum_{d \in \mathcal{D}_\mathbf{s}} d.c \cdot \phi(\mathbf{p_d}, \mathbf{s}) \cdot \mathbf{q_d} \leq \mathbf{b}. \tag{7}$$

3.1.9 Optimal Solution Summary

We summarize the above derived formalism. The free variables to be set by the optimization are

$$\mathbf{b} \in (\mathbb{R}_0^+)^m \text{ and } \forall d \in \mathcal{D} : \mathbf{p_d}, \mathbf{q_d} \in [0,1]^m \tag{8}$$

Both the primary paths $\mathbf{p_d}$ and the backup paths $\mathbf{q_d}$ conform to the conservation rule Equation (1) and exclude start and end nodes explicitly from cycles by Equation (2). The capacity constraints have to be respected either with or without bandwidth reuse (Equation (6) and Equation (7)). Equation (3) and/or Equation (4) may be respected to design p_d and q_d such that q_d protects p_d. The objective function Equation (5) is to be minimized while these constraints are respected.

Unfortunately, the path protection constraints (Equation (3) and Equation (4)) are quadratic with respect to the free variables. Therefore, this description can not be solved by LP solvers. In addition, the failure indication function $\phi(\mathbf{p}, \mathbf{s})$ can not be transformed into a linear mapping. Therefore, we have no efficient algorithm to compute the desired structures $\mathbf{p_d}$ and $\mathbf{q_d}$. If the complexity of the primary and backup multi-paths is restricted, e.g. to disjoint single- or multi-paths, the computation becomes even more difficult due to a required integer solution for $\mathbf{p_d}$ and $\mathbf{q_d}$. Therefore, we use heuristics in the following.

3.2 Separate Primary Path Calculation for Path Protection (PP)

Due to the computational problems and due to the difficulty of controlling the structure of the multi-paths we propose to first calculate a suitable primary path and then to find an appropriate backup path. Given the primary path, the quadratic conditions in terms of free variables in Equations (3) and (4) disappear. In addition, the failure indication function $\phi(\mathbf{p_d}, \mathbf{s})$ is independent of any free variables. Due to the simplicity of the structure, we want the primary path to be a single-path.

3.2.1 Primary Path Computation: Minimum Traffic (MT) Routing

If a network element carries a large amount of traffic and fails, this traffic has to be redistributed and requires a lot of backup capacity near the outage location. Therefore, we construct a routing inducing a minimum traffic load on each network element.

Minimum Traffic Constraints The overall traffic on all links is given by the auxiliary vector $\mathbf{a}^{\mathbf{E}} \in (\mathbb{R}_0^+)^m$ and the overall traffic in all nodes is given by the auxiliary vector $\mathbf{a}^{\mathbf{V}} \in (\mathbb{R}_0^+)^n$, respectively.

$$\mathbf{a}^{\mathbf{E}} = \sum_{d \in \mathcal{D}} d.x \cdot \mathbf{p_d} \,\text{and}\, \mathbf{a}^{\mathbf{V}} = \sum_{d \in \mathcal{D}} d.x \cdot \mathbf{ti}(\mathbf{p_d}) \tag{9}$$

$$\mathbf{a}^{\mathbf{E}} \leq a^E_{max} \cdot \mathbf{1} \,\text{and}\, \mathbf{a}^{\mathbf{V}} \leq a^V_{max} \cdot \mathbf{1}. \tag{10}$$

The value $d.x$ may be set to 1 if only the number of connections is to be minimized or it may be set to $d.c$ if their rate should be taken into account. We use $d.x = d.c$.

Objective Function To avoid very long paths, the objective function takes also the overall required link and node capacity $\mathbf{1}^\top \mathbf{a}^{\mathbf{E}}$ and $\mathbf{1}^\top \mathbf{a}^{\mathbf{V}}$ into account:

$$M^X \cdot a^X_{max} + \mathbf{1}^\top \mathbf{a}^{\mathbf{Y}} \rightarrow \min \tag{11}$$

In our experiments, we set $X = V$ and $Y = V$. The constants $M^E, M^V \in \mathbb{R}_0^+$ control the tradeoff between the conflicting goals "little maximum traffic per network element a^E_{max}, a^V_{max}" and "little overall capacity $\mathbf{1}^\top \mathbf{a}^{\mathbf{Y}}$" that have both to be minimized. A small M^X favors little overall capacity while a large M^X favors little maximum traffic per network element.

Path Constraints Like above, the flow conservation rule (Equation (1)) and the exclusion of start and end nodes from cycles (Equation (2)) have to be respected. Since we are interested in single-path solutions, $\mathbf{p_d} \in \{0,1\}^m$ is required. This, however, leads to a mixed integer LP taking long computation times. Therefore, we relax this condition to $\mathbf{p_d} \in [0,1]^m$ to get a non-integer LP. To obtain a desired single-path as primary path, we take the strongest single-path of the calculated multi-path structure.

3.2.2 Shortest of k-Disjoint Shortest Path (kDSP)

With the primary paths computed by the MT method, a link and node disjoint backup path can not always be found although two disjoint paths exist in the network [19]. To guarantee the existence of k disjoint backup paths if topologically possible, we propose to take the shortest path of a k (node and link) disjoint shortest paths solution (kDSP) with $k \geq 2$ [20, 21].

3.2.3 Optimum Backup Path Calculation (OPT)

The optimum backup path solution for given primary paths can be obtained by a slight modification of the LP formulation in Section 3.1.9. The primary paths $\mathbf{p_d}$ are removed from the set of free variables. This yields a LP formulation which can be solved efficiently. However, the structure of the resulting backup path is potentially very complex since the partial b2b paths are not necessarily disjoint. The following heuristic solves this problem.

3.2.4 Backup Path Heuristic (kDSP)

To have a switching structure that is easy to configure, we desire a multi-path backup solution that consists of disjoint single-paths. We obtain them by a k-disjoint shortest paths solution. However, this provides only the structure of the backup path in terms of several disjoint single-paths. The load balancing of that multi-path is still to be computed. Since this is a special case of the optimization of the load balancing for the self-protecting multi-paths, we refer to Section 3.3.3.

3.2.5 Adaptation to SRLGs

For the computation of disjoint multi-paths we use the kDSP algorithm which is simple an efficient to compute. However, it does not take general SRLGs into account which is a different and NP hard problem. Basically, our kDSP heuristic can be substituted by any other routing scheme yielding disjoint multi-paths.

3.3 Self-Protecting Multi-Paths (SPM)

A self-protecting multi-path consists of k_d link and (not necessarily) node disjoint paths (except for source and destination) $\mathbf{p_d^i}$ for $0 \leq i < k_d$. It is represented by a vector of single-paths $\mathbf{P_d} = (\mathbf{p_d^0}, ..., \mathbf{p_d^{k_d-1}})^\top$. These paths are equal in the sense that they all may be active even without any network failure.

Inactivity Pattern $\mathbf{f_d}(\mathbf{s})$ If only a single link or router fails, at most one of the disjoint paths $\mathbf{p_d^i}$, $0 \leq i < k_d$, for a demand d is affected unless the source or destination node fails. In general, the inactivity pattern $\mathbf{f_d}(\mathbf{s}) \in \{0,1\}^{k_d}$ indicates the failed paths of the SPM depending on the failure pattern $\mathbf{s}$. It is computed by

$$\mathbf{f_d}(\mathbf{s}) = \Big(\phi(\mathbf{p_d^0}, \mathbf{s}), ..., \phi(\mathbf{p_d^{k_d-1}}, \mathbf{s})\Big)^\top. \tag{12}$$

With an inactivity pattern of $\mathbf{f_d} = \mathbf{0}$ all paths are working while for $\mathbf{f_d} = \mathbf{1}$ connectivity can not be maintained. The set of all different failures for SPM $\mathbf{P_d}$ is denoted by $\mathcal{F}_d = \{\mathbf{f_d}(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$.

Load Balancing Function $\mathbf{l_d^f}$ For all demands $d \in \mathcal{D}$ and for all inactivity patterns $\mathbf{f} \in \mathcal{F}_d$, a load balancing function $\mathbf{l_d^f} \in (\mathbb{R}_0^+)^{k_d}$ must be found with

$$\mathbf{1}^\top \mathbf{l_d^f} = 1. \tag{13}$$

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

$$\mathbf{f}^\top \mathbf{l_d^f} = 0. \tag{14}$$

Finally, the vector indicating the transported traffic of demand d over all links is calculated by $\mathbf{P_d}^\top \mathbf{l_d^f} \cdot d.c$.

3.3.1 Load Balancing Heuristics for Disjoint Paths

There are many possibilities for load balancing over multi-paths.

Equal Load Balancing The traffic may be distributed equally over all working paths, i.e. $\mathbf{l}_\mathbf{d}^\mathbf{f} = \frac{1}{\mathbf{1}^\top(\mathbf{1}-\mathbf{f})} \cdot (\mathbf{1} - \mathbf{f})$.

Reciprocal Load Balancing The load balancing factors may be indirectly proportional to the length of the partial paths ($\mathbf{1}^\top\mathbf{p}$). This can be computed by $(l_d^f)_i = \frac{1-f_i}{\mathbf{1}^\top(\mathbf{P_d})_\mathbf{i}} / \left(\sum_{0\leq j<k_d} \frac{1-f_j}{\mathbf{1}^\top(\mathbf{P_d})_\mathbf{j}}\right)$.

3.3.2 Optimized Load Balancing

Load balancing is optimal if the capacity $\mathbf{b}$ required to cover all demands $d \in \mathcal{D}$ and all failure patterns $\mathbf{s} \in \mathcal{S}$ is minimal. We formulate a LP to describe the solution. The free variables are

$$\mathbf{b} \in (\mathbb{R}_0^+)^m, \quad \forall d \in \mathcal{D}\ \forall \mathbf{f} \in \mathcal{F}_d : \mathbf{l}_\mathbf{d}^\mathbf{f} \in (\mathbb{R}_0^+)^{k_d}. \tag{15}$$

The objective function is given by Equation (5). It must be minimized under load balancing and bandwidth constraints. The load balancing constraints in Equations (13) and (14) must be respected by all $\mathbf{l}_\mathbf{d}^\mathbf{f}$ and the bandwidth constraints are newly formulated.

Bandwidth Constraints with Capacity Reuse

$$\forall \mathbf{s} \in \mathcal{S} : \sum_{d \in \mathcal{D}_\mathbf{s}} \mathbf{P_d}^\top \mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{s})} \cdot d.c \leq \mathbf{b}. \tag{16}$$

Bandwidth Constraints without Capacity Reuse Releasing capacity unnecessarily leads to a waste of bandwidth if it can not be reused by other connections. Therefore, load balancing factors $\mathbf{l}_\mathbf{d}^\mathbf{f}$ of active paths must only increase in an outage scenario, except for failed paths for which they are zero. This quasi monotonicity can be expressed by

$$\mathbf{l}_\mathbf{d}^\mathbf{f} + \mathbf{f} \geq \mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{0})}, \tag{17}$$

where $\mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{0})}$ is the load balancing function without failures. The bandwidth $\mathbf{b}$ must take the unused primary bandwidth of failed paths into account as well as the primary bandwidth of connections that are removed due to a router failure. Therefore, we get as bandwidth constraints

$$\forall \mathbf{s} \in \mathcal{S} : \underbrace{\sum_{d \in D_\mathbf{s}} d.c \cdot \mathbf{P_d}^\top \mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{s})}}_{\text{used capacity}} + \\ \underbrace{\sum_{d \in D_\mathbf{s}} d.c \cdot \mathbf{P_d}^\top (\mathbf{f_d}(\mathbf{s}) \cdot \mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{0})})}_{\text{inactive partial paths}} + \sum_{d \in D \setminus D_\mathbf{s}} \underbrace{d.c \cdot \mathbf{P_d}^\top \mathbf{l}_\mathbf{d}^{\mathbf{f_d}(\mathbf{0})}}_{\text{failed connections}} \leq \mathbf{b}. \tag{18}$$

Note that the term $\mathbf{f_d}(\mathbf{s}) \cdot \mathbf{l_d^0}$ expresses an element-wise multiplication of two vectors. Hence, if bandwidth reuse is possible, Equation (16) is used as bandwidth constraint, otherwise Equations (17) and (18) must be respected. Neither protection constraints (Equations (3) and (4)) nor path constraints (Equations (1) and (2)) apply.

3.3.3 Adaptation to Path Protection

The adaptation of the above explained load balancing scheme to path protection mechanisms is simple. We denote the primary paths p_d together with its disjoint backup single-paths as SPM $\mathbf{P_d}$ with $\mathbf{p_d} = (\mathbf{P_d})_\mathbf{0}$. The essential difference between the path protection scheme and the SPM is the inactivity pattern if the primary path is working. For path protection schemes, the inactivity pattern $\mathbf{f_d^{PP}}(\mathbf{s})$ is described by

$$\mathbf{f_d^{PP}}(\mathbf{s}) = \begin{cases} \mathbf{u^0} & \phi(\mathbf{p_d}, \mathbf{s}) = 0 \\ \mathbf{f_d}(\mathbf{s}) & \phi(\mathbf{p_d}, \mathbf{s}) = 1 \end{cases} \tag{19}$$

with $\mathbf{u^0} = (0, 1, ..., 1)^\top$. By substituting the inactivity pattern in Equation (12) by Equation (19), the load balancing optimization in Section 3.3.2 can be applied to path protection schemes, too.

4 Numerical Results

In this section we evaluate the performance of the presented protection mechanisms both in sample and random networks using homogeneous traffic matrices. Our evaluation methodology is the following. We determine the required network capacity, i.e. the sum of all link bandwidths, which is required to accommodate the traffic matrix without resilience if shortest path routing (OSPF) is used based on the hop count metric. We take it as a reference value since it is a lower bound for the required network capacity. Then we calculate the required capacity for a given protection scheme to meet the resilience requirements. The resulting extra capacity is the performance measure in our studies. Note that this extra capacity is not always used for backup purposes only, because sometimes protection mechanisms require longer paths than the shortest one in normal operation. However, we use the term extra capacity and backup capacity exchangeably since the extra capacity is required to provide resilience with the respective protection mechanism.

We compare the backup performance of the path protection schemes and the self-protecting multi-paths in the COST 239 core network [22] (11 routers, 26 bidirectional links in Europe) and in the Labnet [23] (20 routers, 53 bidirectional links in US). We test the different alternatives for primary and backup paths layout, and for load balancing. We use full traffic reduction, bandwidth reuse, and the protection of single router and link failures as default but we investigate the influence of alternative settings, too. Finally, we study the impact of topological network characteristics on the required backup capacity for self-protecting multi-paths.

4.1 Overview of Investigated Protection Mechanisms and Abbreviations

We have discussed various backup possibilities that we summarize at the end of this section. A primary single-path may be protected by a backup $k-1$-multi-path (PP). The primary path may be determined by a k-disjoint shortest paths solution (kDSP) or by a single-path routing that minimizes the transit traffic through each router (MT). The backup multi-path may be computed together with an appropriate load balancing scheme by a LP optimization (OPT) which does not necessarily yield disjoint paths for the multi-path. The $(k–1)$-disjoint shortest paths ($(k–1)$DSP) may also be taken as a backup path solution. In that case, a load balancing scheme is needed. The load may be balanced equally over all parallel paths (E), reciprocally to the length of the disjoint parallel paths (R), or according to an optimized solution computed by a LP (O). As an alternative to the PP solution, the k-disjoint shortest paths are taken as a self-protecting multi-path (kSPM). A kSPM leads to $k+1$ different and easy to diagnose path failure symptoms (including the normal operation). Each of these symptoms requires an own load balancing scheme that may be again set like above (E, R, O).

In the following, we mainly use these abbreviations to refer to specific protection mechanism. E.g., 5DSP-4DSP-R means that the single primary path is chosen as the shortest from a 5-disjoint shortest path solution and the other (at most) 4 are taken for path protection. Load balancing is done reciprocally to the respective path lengths. With MT-OPT the primary path is found by a MT routing solution and the backup multi-path together with a load balancing scheme is computed by a LP. Finally, 5SPM-E signifies a self-protecting multi-path consisting of up to 5 disjoint paths with equal load balancing.

The calculations for the routing and the load balancing were carried out on a Pentium IV 1.5 GHz standard PC and took for the kSPM-O and $\{$MT,kDSP$\}$-$(k–1)$DSP-O some seconds for small and some minutes for large networks. The $\{$MT,kDSP$\}$-OPT computation is more complex and took up to hours.

4.2 Impact of Protection Schemes on the Required Backup Capacity

First, we test the different protection schemes with regard to their required backup capacity.

4.2.1 Performance of Path Protection Mechanisms

We start with the path protection schemes. Figure 1 and Figure 2 show the required backup capacity in the COST239 and the Labnet network for all path protection schemes ($\{k$DSP,MT$\}$-$\{(k-1)$DSP-$\{$E,R,0$\}$,OPT$\}$ with $2 \leq k \leq 5$). We observe the following. The choice of the primary path has a significant influence on the required extra capacity. Throughout all experiments, the results for minimum traffic (MT) routing yields by 5-10 percent points better results than taking the shortest path of kDSP as primary path. The following holds both for primary path found by MT and by kDSP. For $k = 2$ there is only one backup path available. In this case, 100% of the traffic is transported over the remaining path if a path fails, i.e. the performance of all load balancing alternatives (E, R, O) coincides. As there are more disjoint backup paths available for larger k, the traffic can be better distributed in a failure case and less extra capacity is required on the backup links. The most articulate performance gain is achieved for taking $k = 3$ instead of $k = 2$. Due to the network topology, only 3 disjoint path can be found mostly even for $k = 4$. This is a reason why the backup capacity can not be arbitrarily reduced.

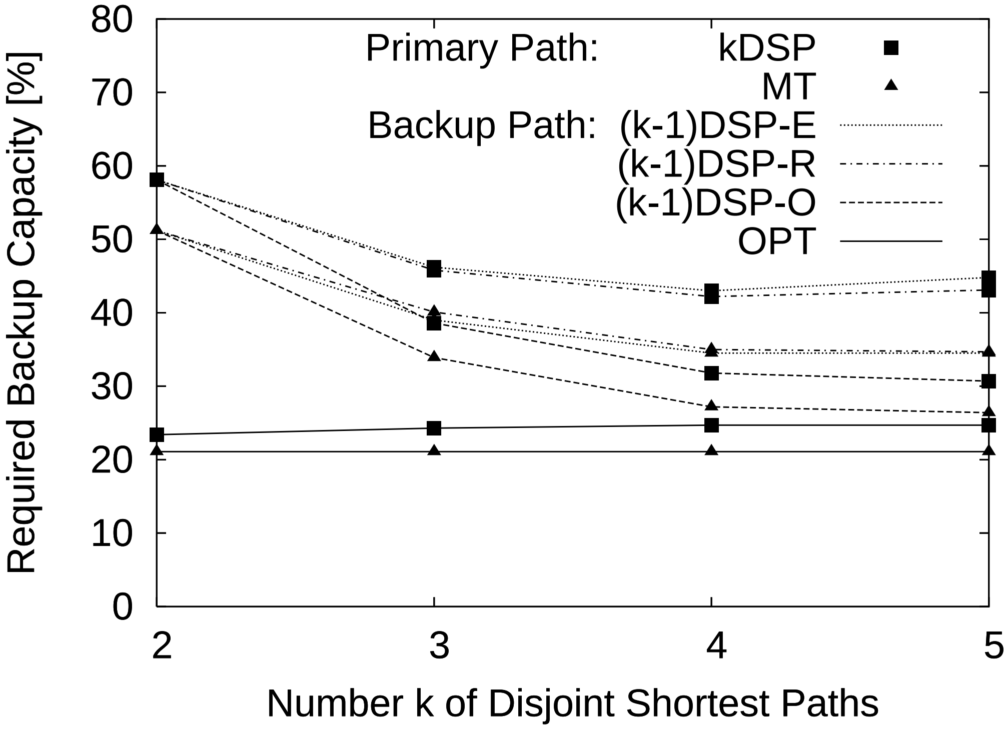

Figure 1: Impact of load balancing for path protection in the COST239 network.

Equal and reciprocal load balancing for the backup multi-path lead approximately to the same results. The optimization of load balancing reduces the required extra capacity by about 10 percent points and the combined optimized backup path and load balancing computation yield another 5-10 percent points capacity reduction. Optimized backup paths (OPT) are more efficient since they do not need to consist of disjoint parallel path. This, however, is also the reason why we consider them problematic. Complex multi-path structures are hard to deploy and hard to manage in failure cases. In addition, the backup path computation is very time consuming. Depending on the heuristic (kDSP or MT) and k, a different primary path may be found. The larger k, the longer the shortest path of kDSP may be. The primary path of minimum traffic routing is independent of k. The outcome for the combined optimization of backup path layout and load balancing only depends on the primary path layout. Therefore, MT-OPT is completely independent of k.

4.2.2 Performance of Self-Protecting Multi-Paths

Figures 3–4 show the required backup capacity in the COST 239 and the Labnet network for the self-protecting multi-paths (kSPM-{E,R,O}) in comparison with the best path protection

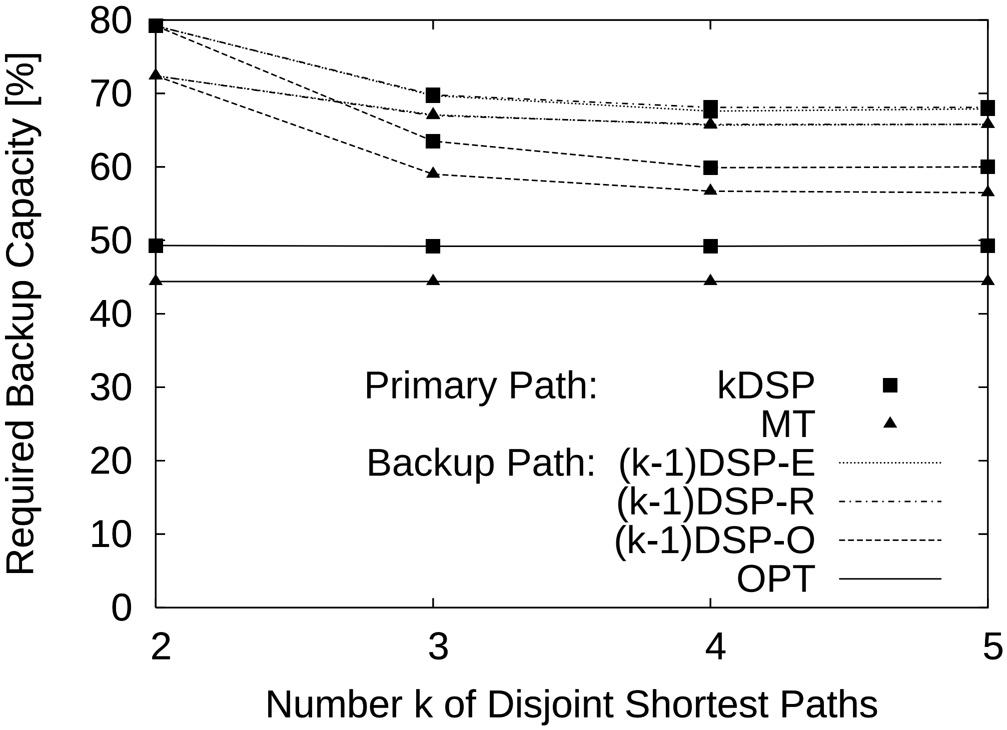

Figure 2: Impact of load balancing for path protection in the Labnet network.

schemes (MT-($k-1$)DSP-O and MT-OPT).

In contrast to the path protection scheme, the performance for kSPM-E and kSPM-R degrades for increasing k and more extra capacity is needed in the COST239 network. The same effect can also be observed to a minor extent in the Labnet network. As the SPM consists of more disjoint path for larger k, these paths are longer. Their extensive use can not be avoided with kSPM-E or kSPM-R and so they lead to an increased required network capacity. Hence, SPM with simple load balancing schemes reveal only minor benefits. The SPM become more economic with optimized load balancing and with increasing k. SPM-O is about 10 percent points superior to MT-4DSP-O in both networks, which has been proven to be the best feasible path protection solution. In the COST 239 network, SPM-O is even better than MT-OPT. We motivate the superiority of the SPM by the following explanation. In contrast to a single primary path, a SPM distributes the traffic from a single source through the network over several disjoint path. In case of a link failure, the affected traffic stems from more different demands and only a fraction of each of their traffic $d.c$ is carried over the failed link. As more demands are affected than with a single primary paths, the load of the failed link can be spread out over more backup paths. As a consequence, less shareable backup capacity is required on these backup paths. Like above, for $k = 2$ there is only a single backup path in a failure case but the corresponding extra capacities for 2SPM-$\{$E,R,O$\}$ do not coincide in the figure, i.e. load balancing does matter. The optimized load balancing distributes the traffic in a way to prevent strong traffic concentrations in any network element. This avoids that a large traffic rate must

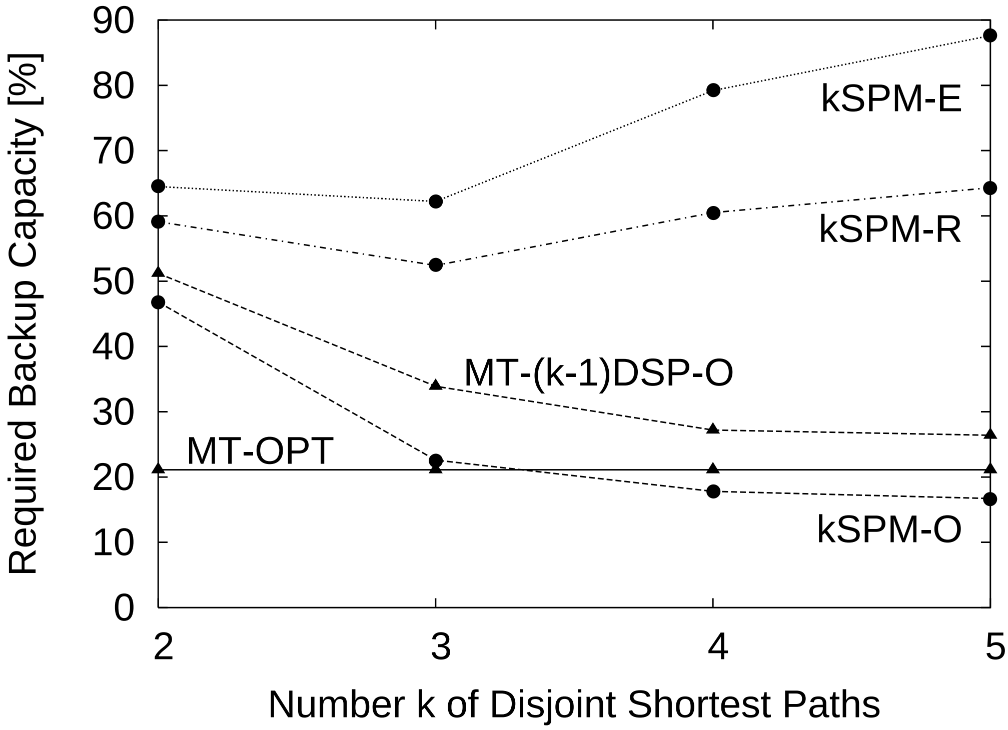

Figure 3: Impact of load balancing for self-protecting multi-paths in the COST 239 network.

be redirected if this element fails. This idea is similar to the MT heuristic for finding suitable primary paths.

4.3 Impact of the Traffic Reduction, Protection, and Bandwidth Reuse Options

We investigate the traffic reduction, protection, and bandwidth reuse options for the calculation of the required backup capacities. First, if a router fails, either flows originating or ending at that router are removed (*full traffic reduction (full TR)*), or only flows originating from there are removed (*source TR*), or no flows are removed (*no TR*), i.e. we consider the router failure only for transit flows. This has some impact on the traffic volume in the network. Second, different types of failures can be considered. We consider only single router failures (*router protection*), single link failures (*link protection*), or both single router and link failures (*full protection*). This influences the number of protected failure scenarios $\mathcal{S}$ for which the network capacity must be provisioned. Third, in optical networks, link bandwidth must be released explicitly before it is reused. This is probably not possible in a failure scenario. Therefore, we consider the case that link bandwidth of a failed path may be reused for backup paths like in packet-switched networks (*bandwidth reuse*) and the case that this link bandwidth can not be reused by backup paths (*no bandwidth reuse*).

Figure 5 and Figure 6 show the required backup capacity for the 5SPM-O protection scheme in the COST 239 and in the Labnet network. The traffic reduction has no effect if

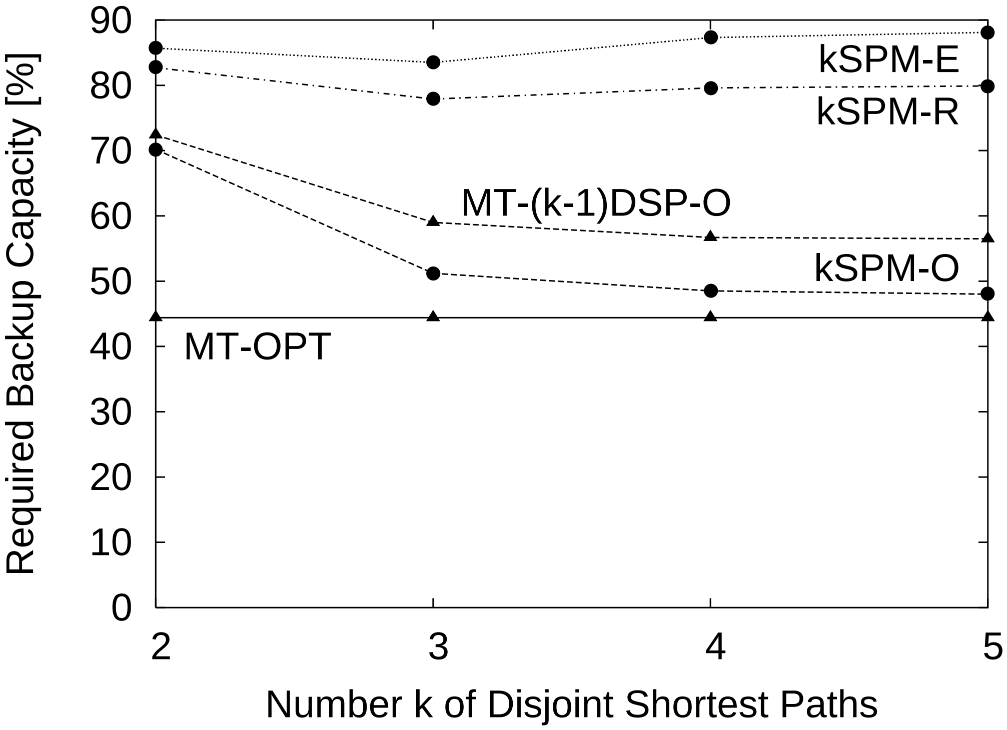

Figure 4: Impact of load balancing for self-protecting multi-paths in the Labnet network.

only link failures occur; otherwise it has hardly effect except for router failures in the COST 239 network. Due to the small size of that network, the proportion of the reduced traffic is large related to the overall traffic and, therefore, the impact of full traffic reduction is significantly larger than in the Labnet.

If both single link and router failures are protected, slightly more capacity is required than just for single link or router failures, respectively. In the COST 239 network, single router failure protection needs the least backup capacity while single link failure protection needs the least backup capacity in the Labnet. The reason for that contradictory result is again the network size. In networks with a small average shortest path length, there are only a few flows traversing transit routers. Only these flows are redirected if a router fails, other flows originating or ending at that router are either removed or stay unchanged depending on the considered traffic reduction option. In medium size networks, this effect vanishes and router failure protection requires almost as much backup capacity as full protection. The mere link failure protection is about 10 percent points cheaper which is quite significant.

Throughout all experiments the “no bandwidth reuse” restriction leads to about 5 percent points more backup capacity compared to bandwidth reuse by backup paths. In our investigations, we use as default options full traffic reduction, full protection, and bandwidth reuse.

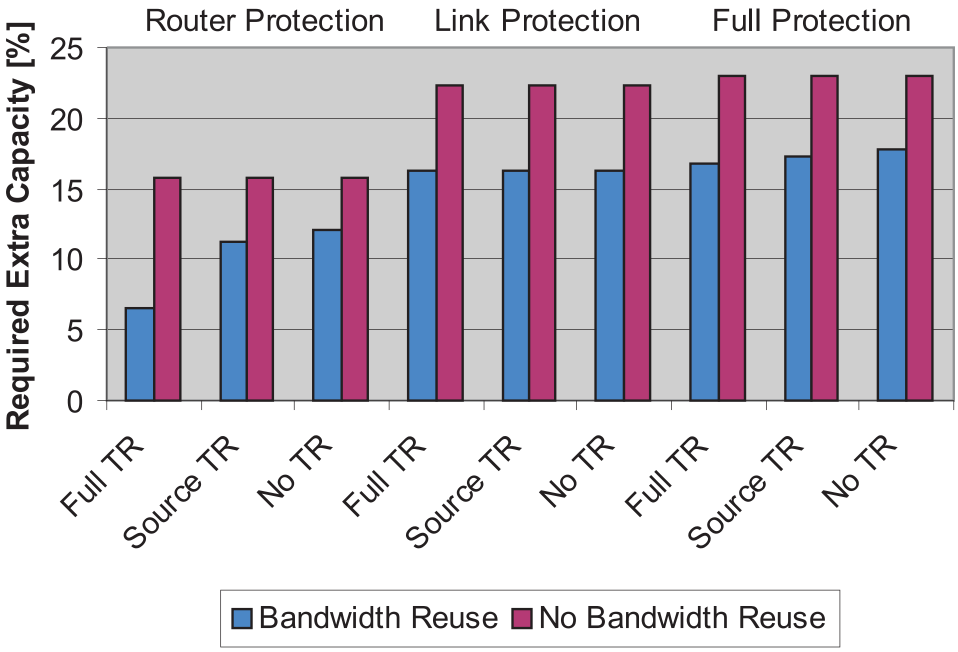

Figure 5: Impact of protection level, traffic reduction, and bandwidth reuse in the COST239 network.

4.4 Impact of the Traffic Matrix

The default traffic matrix in our studies consists of a homogeneous traffic distribution. To study the impact of heterogenous traffic distributions, we assign the traffic proportionally to the corresponding population of the nodes in the network [23]. To increase the variance of the traffic matrix, we manipulate the population by an exponential extrapolation. The traffic matrix is homogeneous for the extrapolation parameter $t=0$, it is normal for $t=1$, and it has extreme variance for $t=2$.

We dimension the COST239 network and the Labnet for 5SPM-O and OSPF routing with full protection and for OSPF without protection. Table 1 shows the additional capacity for 5SPM-O and SP with full protection in percent compared to OSPF without protection. In addition, it shows the fraction of both values, which is the OSPF-normalized backup capacity for 5SPM-O. The values show that the required backup capacity increases in the COST239 network if the variance of the traffic matrix increases (t). For OSPF holds the same, therefore, the advantage of 5SPM-O remains. In the Labnet we observe the opposite. With increasing t, the required capacity for the 5SPM-O reduces significantly but the one for OSPF only slightly. Therefore, the advantage of SPM becomes more clear for heterogeneous traffic matrices. Hence, we have shown that the backup performance depends also on the traffic matrix. However, there is no obvious general trend how realistic traffic matrices influence the required backup capacity of optimized protection switching mechanisms. This depends on the network

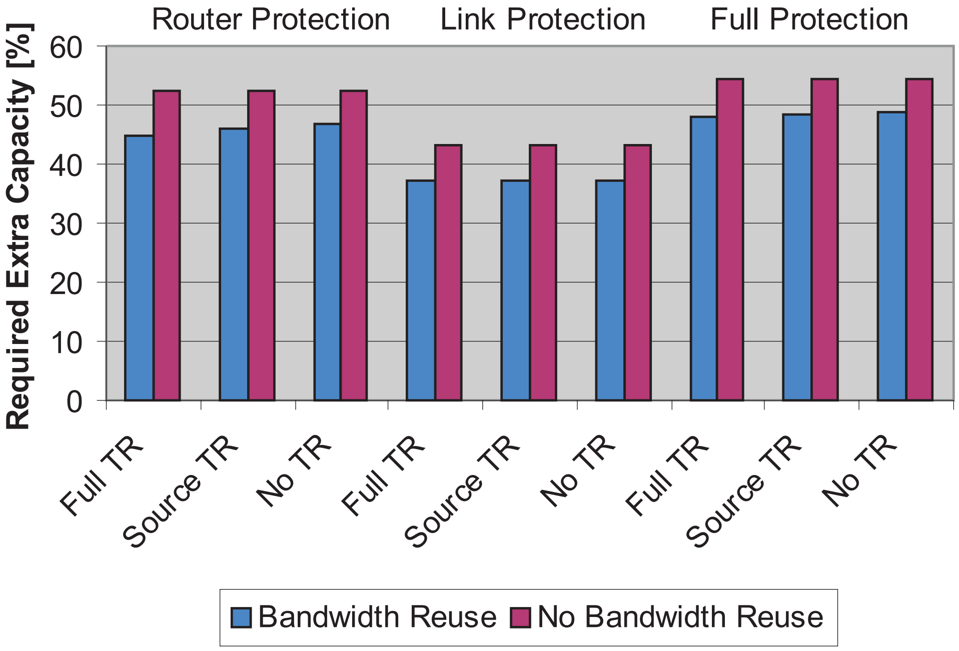

Figure 6: Impact of protection level, traffic reduction, and bandwidth reuse in the Labnet network.

topology and the traffic matrix itself.

4.5 Performance Comparison of Protection Mechanisms in Various Network Topologies

In Figure 7 the required backup capacity is given for various protection mechanisms in various sample networks. The protection mechanisms are simple OSPF rerouting, 5DSP-4DSP-O, MT-4DSP-O, 5DSP-OPT, MT-OPT, and 5SPM-O. The x-axis indicates the average number of disjoint parallel paths k^* in a network and the y-axis shows the required backup capacity. The different protection mechanisms are distinguished by the point shape. Symbols belonging to the same network are grouped together by a vertical line. The sequence of these vertical lines maps the sequence of the letters in the figure. Lowercase letters correspond to networks taken from [6] while uppercase letters correspond to these networks with the modification that nodes with a node degree of at most 2 are successively removed. Note that the MT-5DSP and MT-OPT protection mechanisms are missing for some networks because no backup path could be found due to the choice of the primary path.

In general we observe that the required backup capacity decreases with increasing k^* for all protection mechanisms. The dashed line shows the least square interpolation of the results for 5SPM-0 according to an exponential function. Furthermore, the relative savings compared to OSPF rerouting increase. The self-protecting multi-paths are superior to path

Table 1: Required backup capacity for different traffic matrices in [%].

extr-pol. par. t	0	1	2
COST239			
5SPM-O	16.7	32.4	45.9
shortest path	73.3	91.1	106.2
normalized	22.9	35.6	43.3
Labnet			
5SPM-O	47.5	36.6	37.2
shortest path	102.3	91.9	102.6
normalized	46.4	39.8	36.3

protection schemes which can be explained as follows. A kDSP-$k-1$DSP-O is structurally very similar to a kSPM because they use the same disjoint path of a kDSP computation. But due to the limitation of Equation (19), the optimization of the load balancing for path protection methods has fewer degrees of freedom and so, comparable self-protecting multi-paths are more economic regarding backup capacity. The 5SPM-0 clearly outperforms mostly all other protection mechanisms, only the optimized backup paths 5DSP-OPT and MT-OPT lead sometimes to less backup capacity at the expense of a complex backup multi-path structure. Hence, the self-protecting multi-path is the best feasible solution for all investigated networks.

4.6 Impact of Network Topology Characteristics

To study the impact of the network topology in more detail, we conduct studies based on random networks. First, we describe our algorithm for the construction of random networks and then we present the results of a parameter study.

4.6.1 Random Network Construction

We construct random networks controlling some of their essential characteristics. One of them is the degree $deg(v)$ of a node v, which is the number of links v is connected with. We briefly explain our network construction method that incorporates features of the well know Waxman model [24, 25]. It is an efficient algorithm that provides control over the minimum, the average, and the maximum node degree (deg_{min}, deg_{avg}, deg_{max}), and avoids loops and parallels.

The algorithms starts with an empty link set $\mathcal{E} = \emptyset$ and defines a single arbitrary node $v_{start} \in \mathcal{V}$ connected. Then, $\frac{|\mathcal{V}|\cdot deg_{avg}}{2}$ links are added successively to $\mathcal{E}$ by connecting suitable nodes v_α and v_ω. An arbitrary node v_α is chosen from a set of preferred nodes $\mathcal{V}_\alpha$ with the following properties. All $v \in \mathcal{V}_\alpha$ are connected and have $deg(v) \leq deg_{max}$. If a node $v \in \mathcal{V}$ exists with $deg(v) < deg_{min}$, all $v \in \mathcal{V}_\alpha$ must have $deg(v) < deg_{min}$. The set of potential neighbor nodes $\mathcal{V}_\omega$ obeys the following requirements: Loops and parallels must be avoided, i.e. $v_\alpha \notin \mathcal{V}_\omega$ and $(v_\alpha, v_\omega) \notin \mathcal{E}$. Furthermore, if an unconnected node $v \in \mathcal{V}$ exists,

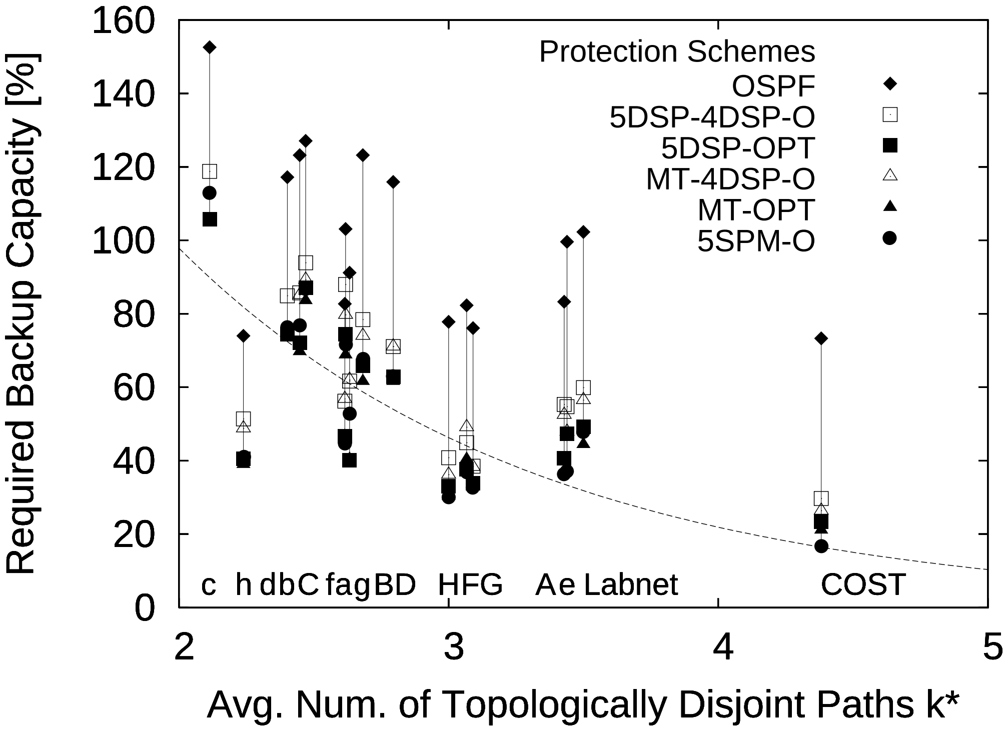

Figure 7: Comparison of protection mechanisms in sample networks.

all $v \in \mathcal{V}_\omega$ must be unconnected. The node $v_\omega \in \mathcal{V}_\omega$ is chosen according to a probability distribution which depends on v_α and $\mathcal{V}_\omega$. Here, the Waxman model comes into play. Each node has a position in the plane. The Euclidean distance $d(v,w)$ induces a weight $P(v,w) = a \cdot e^{-\frac{d(v,w)}{b \cdot d_{max}}}$ with $d_{max} = \max_{v,w \in \mathcal{V}} d(v,w)$, and $P(v,w)$ produces the probability distribution $p_{v_\alpha}(w) = \frac{P(v_\alpha,w)}{\sum_{v \in \mathcal{V}_\omega} P(v_\alpha,v)}$. Given a maximum node degree deviation deg_{dev}^{max}, the minimum node degree is set to $deg_{min} = \max(deg_{avg} - deg_{dev}^{max}, 2)$ and the maximum node degree is set to $deg_{max} = deg_{avg} + deg_{dev}^{max}$.

4.6.2 Parameter Study

Figure 8 shows the required backup capacity for 240 random networks of different size, different average node degree deg_{avg}, and different maximum node degree deviation deg_{dev}^{max} as topology characteristic. There are 5 random networks for each topology description. Like above, the correlation between k^* and the required backup capacity is clearly visible. We identify four clusters of networks and it turns out that they have the same average node degree deg_{avg}. The dashed lines are least square interpolations among the points of these clusters. This makes the clusters more visible, however, the extrapolation of those curves does not make sense since $k^* \leq deg_{avg}$ holds. Within a cluster, the network size n seems to be irrelevant. A small maximum deviation deg_{dev}^{max} of the node degrees $deg(v)$ from the average node degree deg_{avg} seems to increase k^*, leading to more efficient backup solutions within a cluster.

Therefore, resilience can be achieved at lower cost if the network topology is symmetric.

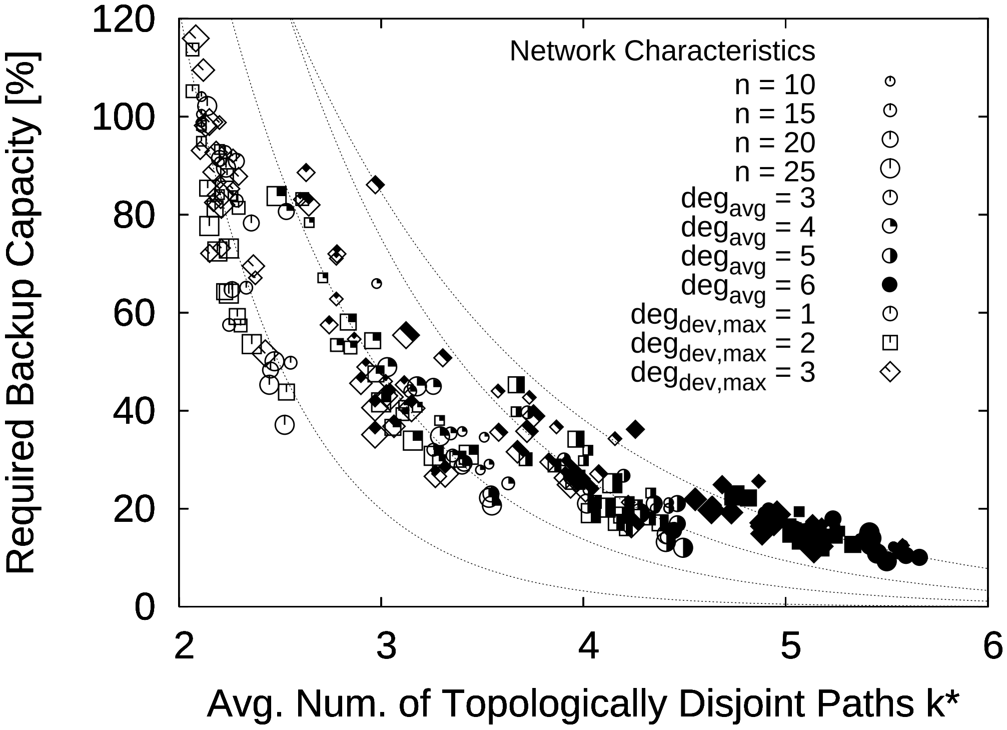

Figure 8: Required extra capacity for self-protecting multi-paths in random networks.

4.6.3 Backup Performance Relative to OSPF Rerouting

In Figure 9 the same data are presented in a different way. For all 5 random networks with the same topological characteristics, we build the mean of their k^* and the mean of their ratios of the SPM and OSPF rerouting backup capacity. The horizontal and vertical lines provide the 90% confidence intervals. The data are plotted on a logarithmic scale to make exponential trends better visible. The dashed line is the least square interpolation of all experiments and the solid lines are the interpolations within a cluster of networks with the same average node degree deg_{avg}. The four clusters confirm the above observation that deg_{avg} of a network is strongly correlated with k^*. Increasing the average node degree deg_{avg} shifts the exponential trend slightly towards larger backup capacity. Again, we observe an exponential decay with regard to an increasing k^*. The average number of disjoint parallel path k^* has also a clear impact on the OSPF normalized backup capacity. The OSPF normalization dampens the influence of topological characteristics and shows clearly the benefits of the SPM approach in comparison with conventional rerouting. A larger number of disjoint paths increases the superiority of the SPM over OSPF rerouting.

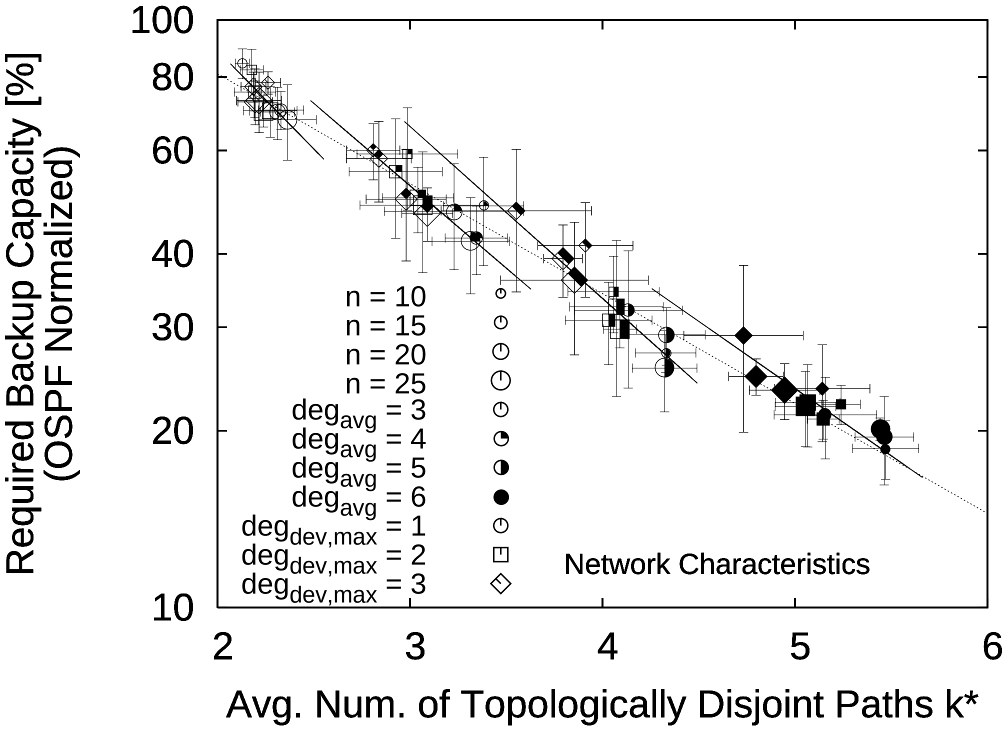

Figure 9: Required extra capacity for self-protecting multi-paths in random networks as a fraction of OSPF extra capacity.

5 Conclusion

In this paper we have considered the capacity requirements for e2e backup mechanisms in autonomous systems, that are necessary for resilient networks. The routing together with an optimum load balancing and bandwidth provisioning in the network is calculated for the working case and all single failure cases such that the overall capacity is minimized. So far, either simple and little efficient restoration schemes [12] have been optimized or the optimization was done for technically not feasible mechanisms [6]. We tried to close that gap by defining two different resource-efficient and simple to implement protection mechanisms. Only traffic of failed paths is shifted, so the signalling overhead in a failure case is low. We use multi-path routing to minimize the extra capacity, however, we use only multi-path structures consisting of disjoint paths to keep the configuration and the failure diagnostics simple. The first type is path protection based on a suitable primary single-path and a multi-path for backup purposes. The second type is a so-called self-protecting multi-path. The required extra capacity is minimized by means of load balancing optimization.

We evaluated the performance of our proposed mechanisms in different existing networks as well as in random networks by adapting their structure and by dimensioning their links such that sufficient capacity is available in normal operation and in all protected failure scenarios. Our results showed that OSPF rerouting often requires more than 100% extra capacity to provide resilience while for the new mechanisms only a fraction of that capacity is needed. We also studied the influence of various side conditions like failure types, traffic reduction due to failed routers, or bandwidth reuse restrictions. The structure of the traffic matrix has a significant impact on the backup performance. We illustrated that the amount of required extra capacity depends on the network topology, and in particular on the average number k^* of disjoint paths in a network but not on the network size. Our network simulations revealed that the amount of extra capacity decays exponentially with k^* compared to the additional resources for OSPF routing. Hence, our new mechanisms lead to very cheap network resilience in suitable network topologies. As an example, only 17% extra capacity is required in the COST 239 network to provide full resilience against all single link and router failures. This makes protection mechanisms on the network layer significantly cheaper than on the physical layer from a resource point of view.

As a challenge remain, e.g., fast heuristics for the calculation of an optimized load balancing for large networks. Suitable network structures are a prerequisite for cheap backup capacities and should be further identified. Backup mechanisms for networks with given link capacities should be adapted and optimized to maximize the throughput of a network while meeting resilience requirements. The impact of multiple failures on the QoS degradation is to be investigated in networks that are resilient against single failures and, finally, inter-area fault tolerance is still to tackle.

References

[1] V. Sharma (Ed.) and F. Hellstrand (Ed.), “RFC3469: Framework for Multi-Protocol Label Switching (MPLS)-based Recovery.” http://www.ietf.org/rfc/rfc3469.txt, Feb. 2003.

[2] P. Pan, D.-H. Gan, G. Swallow, J. P. Vasseur, D. Cooper, A. Atlas, and M. Jork, “Fast Reroute Extensions to RSVP-TE for LSP Tunnels.” http://www.ietf.org/internet-drafts/draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt, June 2003.

[3] A. Autenrieth and A. Kirstädter, “Engineering end-to-end ip resilience using resilience-differentiated qos,” *IEEE Communications Magazine*, vol. 40, pp. 50–57, Jan 2002.

[4] W. D. Grover, “Cycle–Oriented Distributed Preconfiguration: Ring–like Speed with Mesh–like Capacity for Self–planning Network Restoration,” in *Proceedings of IEEE ICC ’98*, pp. 537–543, Jun 1998.

[5] C. G. Gruber and D. A. Schupke, “Capacity–efficient Planning of Resilient Networks with p–Cycles,” in *Proceedings of Networks 2002*, pp. 389–395, Jun 2002.

[6] K. Murakami and H. S. Kim, “Comparative Study on Restoration Schemes of Survivable ATM Networks,” in *IEEE INFOCOM’97*, (Kobe City, Japan), pp. 345 – 352, April 1997.

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

[8] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *IEEE INFOCOM'00*, pp. 519–528, 2000.

[9] I. Gojmerac, T. Ziegler, F. Ricciato, and P. Reichl, "Adaptive Multipath Routing for Dynamic Traffic Engineering," in *GLOBECOM'03*, (San Francisco), Nov 2003.

[10] G. Dittmann and A. Herkersdorf, "Network Processor Load Balancing for High–Speed Links," in *SPECTS 2002*, (San Diego, CA), pp. 727–735, 2002.

[11] M. S. Kodialam and T. V. Lakshman, "Minimum Interference Routing with Applications to MPLS Traffic Engineering," in *Proceedings of IEEE INFOCOM 2000*, vol. 2, pp. 884–893, Mar 2000.

[12] R. R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal Capacity Placement for Path Restoration in STM and ATM Mesh-Survivable Networks," *IEEE/ACM Transactions on Networking*, vol. 6, pp. 328 – 336, June 1998.

[13] J. Strand, A. L. Chiu, and R. Tkach, "Issues For Routing In The Optical Layer," *IEEE Communications Magazine*, vol. 39, pp. 81–87, Feb 2001.

[14] B. Rajagopalan, J. V. Luciani, and D. O. Awduche, "IP over Optical Networks: A Framework." http://www.ietf.org/internet-drafts/draft-ietf-ipo-framework-05.txt, Sep 2003.

[15] K. Kompella and Y. Rekhter, "Routing Extensions in Support of Generalized Multi–Protocol Label Switching." http://www.ietf.org/internet-drafts/draft-ietf-ccamp-gmpls-routing-09.txt, Oct 2003.

[16] ILOG, Inc., www.cplex.com, *CPLEX*.

[17] A. Makhorin, *GNU Linear Programming Kit Reference Manual Version 4.0*. Free Software Foundation, Inc., 59 Temple Place — Suite 330, Boston, MA 02111, USA, May 2003.

[18] H. Saito and M. Yoshida, "An optimal recovery LSP assignment scheme for MPLS fast reroute," in *Proceedings of Networks 2002*, pp. 229–234, Jun 2002.

[19] D. A. Dunn, W. D. Grover, and M. H. MacGregor, "Comparison of k–Shortest Paths and Maximum Flow Routing for Network Facility Restoration," *IEEE Journal on Selected Areas in Communications*, vol. 2, no. 1, pp. 88–99, 1994.

[20] J. W. Suurballe, "Disjoint Paths in a Network," *Networks*, vol. 4, pp. 125–145, 1974.

[21] J. Edmonds and R. M. Karp, "Theoretical Improvements in the Algorithmic Efficiency for Network Flow Problems," *Journal of the ACM*, vol. 19, pp. 248–264, Apr 1972.

[22] P. Batchelor et al., "Ultra High Capacity Optical Transmission Networks. Final report of Action COST 239." http://barolo.ita.hsr.ch/cost239/network/, 1999.

[23] M. Menth, S. Kopf, and J. Milbrandt, "A Performance Evaluation Framework for Network Admission Control Methods," in *IEEE Network Operations and Management Symposium (NOMS)*, (Seoul, South Korea), April 2004.

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

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