 The definitive version of this paper has been published in 2nd Conference on Next Generation Internet Design and Engineering (NGI), 2006, 10.1109\/ngi.2006.1678255.

Comparison of Capacity Requirements for the Self-Protecting Multipath and Similar Mechanisms in Resilient Packet Networks

Michael Menth†, Andreas Reifert††, and Jens Milbrandt†
† Dept. of Distributed Systems, Inst. of Computer Science, University of Würzburg, Germany
Email: {menth,reifert,milbrandt}@informatik.uni-wuerzburg.de
††Institut für Kommunikationsnetze und Rechnersysteme (IKR), University of Stuttgart, Germany
Email: reifert@ikr.uni-stuttgart.de

***Abstract*— We compare the capacity requirements of two new end-to-end (e2e) protection switching mechanisms: the self-protecting multipath (SPM) and several (multi-)path protection (PP) methods. Their structure consists of disjoint parallel e2e paths and the traffic is distributed over these paths according to a load balancing function. If one of the paths fails, the traffic is redistributed to the working paths according to a path failure specific load balancing function. The contribution of this work is the calculation of the path layout and the load balancing functions for both the PP and the SPM method. We use exact optimizations and simple heuristics for that objective and take a dimensioning approach to compare the capacity requirements of the different mechanisms. Our results illustrate, e.g., that the savings potential depends on the network topology and that 17% additional capacity can be sufficient for full resilience against all single router and link failures in well designed networks.**

Keywords: protection and restoration, load balancing

I. Introduction

Carrier grade networks can not afford outages due to internal link or router failures that compromise the Quality of Service (QoS) perceived by their customers. Therefore, backup mechanisms are required to detour affected traffic aggregates around the outage location. In contrast to IP rerouting, such mechanisms must react fast and they must control the deviation paths. Fast failure detection and fast reaction is achieved by exchanging periodic "Hello" messages and switching the traffic onto pre-computed and pre-installed backup paths as soon as these periodic messages do not arrive anymore. 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 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. In this study, we strive to minimize the required backup capacity and take it as the measure for the performance comparison of different backup mechanisms. In [6], [7] the optimum path layout and load balancing for the primary and backup paths is computed for a given network topology and traffic matrix. This optimal solution leads to complex multipath structures that may branch and join at interior nodes which makes them hard to configure. Furthermore, in case of a network failure, the relocation of unaffected aggregates to deviation paths is sometimes needed, which imposes heavy signaling load on the network in a critical situation.

The contribution of this paper is the optimization of recently proposed simple protection switching mechanisms [8] that may be implemented by mechanisms like MPLS that support explicit routing. We take advantage of the load balancing potential of multipath forwarding and minimize the required extra capacity by polynomial-time optimization algorithms. Our multipath structures are significantly simpler than general multipaths 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 – 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 II we point out the difference between our work and other routing optimization approaches. In Section III we explain the optimization of the primary and backup paths and the load balancing to minimize the required extra capacity. The numerical results in Section IV demonstrate the performance of the recently proposed protection switching mechanisms. Section V summarizes this work and gives some outlook on further work.

II. 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.

A. 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 [9]. 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 [10]. Pure IP and MPLS solutions may also be combined [11]. These approaches require the knowledge of the traffic matrix which can be well obtained in MPLS networks [12]. The solution in [13] is based on a stable closed loop solution using multipath structures and it renounces on the knowledge of the traffic matrix. 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 [14], [15] achieves that goal very well.

The authors of [16] 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. A distributed protocol solution for GMPLS is given in [17]. If backup capacity sharing is allowed, the 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. The minimization of backup resources can also be done for pure IP routing [18], [19]. However, it is less effective because destination-based routing allows for more powerful traffic engineering than source-*and*-destination-based forwarding (e.g. MPLS).

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 a minimum amount of 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 [20] 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 multipath 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. In [6], [7] multipath routing is used and the required network resources are minimized by calculating the optimum path layout and routing independently for each failure scenario. Although these backup solutions lack technical constraints that make them feasible for real-world systems, they present lower bounds for the required backup capacity.

B. Restrictions for Path Layout

We consider the independent path layout calculation based on general multipaths for the normal operation mode and for each failure scenario like in [6], [7]. 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. We explain why these results can not be implemented as restoration mechanisms and derive technical side constraints for feasible backup solutions.

Firstly, a failure-specific protection mechanism requires that the information about the exact location of the failure is propagated to all ingress routers to trigger the activation of their backup paths. This entails extensive signaling in a critical system state at a time for which the long distance connectivity in terms of hops is corrupted.

Secondly, the relocation of unaffected primary paths must be done first if required. Then, backup paths can be activated for affected primary paths. Otherwise, the simultaneous relocation of primary paths and the activation of backup paths might lead to transient overload on some network elements. Hence, deflecting more paths than necessary requires a coordinated switching order of distributed switching locations. This problem is avoided if the relocation of unaffected primary paths is not required.

Thirdly, a failure-specific backup design requires potentially a separate alternate path for each primary path in each protected failure scenario. This leads to a large amount of backup paths which must be pre-installed and administered. This makes the path configuration very complex and a tremendously large number of paths is a problem for the state maintenance of today's core network routers.

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

Another restriction for path layout are Shared Risk Link Groups (SRLGs) [21], [22], [23] 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.

C. Proposal of New Protection Switching Mechanisms

Based on the previous insights, we present two fundamentally different protection switching mechanisms. As outlined above, only multipath 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], [24] 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 multipath routing that allows for load distribution in failure cases.

Our first studied alternative is e2e path protection (PP) for a single primary path with a multipath as a backup path which is composed of link or node disjoint parallel paths. The second alternative is the e2e self-protecting multipath (SPM) which we have originally suggested in [8]. It consists of link or node disjoint parallel paths and does not differentiate explicitly between primary and backup paths. In the failure-free scenario, the traffic is distributed over all parallel paths according to the load balancing function for the failure-free case. If one of these disjoint paths fails, the traffic is redistributed onto the remaining active paths according to a path failure specific load balancing function. In the next section, we describe a computation for a suitable path layout and an optimization for multipath load balancing in connection-oriented networks.

III. Optimization

In this section, we explain first our notations taken from basic linear algebra to represent flows and paths. We describe various side conditions as linear inequalities. Mostly, we get linear programs (LPs) that can be solved by standard software like *ILOG CPlex* [25] or the *GNU Linear Programming Kit* [26]. We first formulate a general optimal primary and backup path solution, which turns out to be computationally infeasible. Then, we propose methods to calculate the primary and the backup path structure separately. Finally, we compute the load balancing functions for (multi-)path protection (PP) and the self-protecting multipath (SPM).

A. 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 $\mathbf{u}^\top\mathbf{v}=\sum_{i=1}^{n}u_iv_i$. Binary operators $\circ\in\{+,-,\cdot\}$ are applied component-wise, i.e. $\mathbf{u}\circ\mathbf{v}=(u_0\circ v_0,\ldots,u_{n-1}\circ v_{n-1})^\top$. The same holds for relational operators $\circ\in\{<,\leq,=,\geq,>\}$, i.e. $\mathbf{u}\circ\mathbf{v}$ equals $\forall\ 0\leq i<n\colon u_i\circ v_i$. For simplicity reasons we define special vectors $\mathbf{0}=(0,\ldots,0)^\top$ and $\mathbf{1}=(1,\ldots,1)^\top$ with context specific dimensions.

B. Formulation of Networking Concepts and Side Conditions

1) 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 $(a_\omega)_{ij}=\begin{cases}0 & \omega(e_j)\neq v_i\\ 1 & \omega(e_j)=v_i\end{cases}$. The incidence matrix $\mathbf{A}\in\{-1,0,1\}^{n\times m}$ is defined as $\mathbf{A}=\mathbf{A}_\omega-\mathbf{A}_\alpha$. The product $\mathbf{Ae_j}$ yields a node vector. The i-th position of $\mathbf{Ae_j}$ contains –1 if v_i is the source node of link e_j and +1 if v_i is the target node; otherwise, it contains zero. The product $\mathbf{v_j}^\top\mathbf{A}$ yields a link vector. The i-th position of $\mathbf{v_j}^\top\mathbf{A}$ contains –1 if e_i is an outgoing link of node v_j and +1 if e_i is an incoming link; otherwise, it contains zero. Loops cannot be expressed by that formalisms.

2) Demands, Traffic Matrix, Paths, and Flows:

a) Demands and Traffic Matrix: We define the demand of a flow between routers $\mathbf{v_i}$ and $\mathbf{v_j}$ by $d=(i,j)$ 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 $c(d)$ and corresponds to an entry in the traffic matrix.

b) 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 multipath $\mathbf{p_d}\in[0,1]^m$, which is more general since the traffic may be split into several partial paths carrying a non-integer 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

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

Cycles within a path containing only inner nodes can be easily removed from a potential solution. In contrast, 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:

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

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

3) 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.

4) 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.

a) 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 $\mathcal{D}_\mathbf{s} = \mathcal{D}$.

b) Source Traffic Reduction: An aggregate flow is removed from the traffic matrix if the source node v_i of demand $d = (i,j)$ fails. Hence, we get $\mathcal{D}_\mathbf{s} = \mathcal{D} \setminus \{(i,j) : \mathbf{v_i}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\}$. If the failed node is the destination of a flow, "server push" traffic can still be transported through the network although it cannot be delivered correctly.

c) Full Traffic Reduction: We assume that traffic with a failed source or destination node is stalled. Hence, we get $\mathcal{D}_\mathbf{s} = \mathcal{D} \setminus (\{(i,j) : \mathbf{v_i}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\} \cup \{(j,i) : \mathbf{v_i}^\top \mathbf{s}_\mathcal{V} = 1, 1 \leq j \leq n, i \neq j\})$.

5) Failure Indication Function: The failure indication function $\phi(\mathbf{p}, \mathbf{s})$ indicates whether a path p is affected by a failure pattern $\mathbf{s}$ [27]. 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

$$\phi(\mathbf{p}, \mathbf{s}) = \begin{cases} 1 & {\mathbf{s}_\mathcal{E}}^\top \mathbf{p} + {\mathbf{s}_\mathcal{V}}^\top \mathbf{ti}(\mathbf{p}) > 0 \\ 0 & \text{otherwise.} \end{cases}$$

6) 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

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

7) 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

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

where $\mathbf{w} \in (\mathbb{R}_0^+)^m$ is a vector of weights, which 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 $d \in \mathcal{D}_\mathbf{s}$.

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

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

b) 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 resources of a redirected connection and to make them available for others. This is respected by the following capacity constraints:

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

C. Optimum Primary and Backup Path Solution

We summarize the above derived formalism to compute the optimum primary and backup path solution. 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 $\mathbf{p_d}$ and $\mathbf{q_d}$ such that $\mathbf{q_d}$ protects $\mathbf{p_d}$. 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 and this description cannot be solved by LP solvers. In addition, the failure indication function $\phi(\mathbf{p}, \mathbf{s})$ cannot 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 $\mathbf{p_d}, \mathbf{q_d} \in [0,1]^m$ is allowed, the primary and the backup paths are general multipaths which are too difficult to administrate. However, the restriction $\mathbf{p_d}, \mathbf{q_d} \in \{0,1\}^m$ leads to single path structures only. This is too restrictive since it prohibits cost-effective backup solutions and the computation becomes more difficult due to a required integer solution for $\mathbf{p_d}$ and $\mathbf{q_d}$.

D. Path Layout Heuristics

In the following, we calculate the primary paths to fix $\mathbf{p_d}$ in Equation (3) and Equation (4) such that we get rid of the quadratic expressions of free variables. As another consequence, the failure indication function $\phi(\mathbf{p}, \mathbf{s})$ depends then only on constant values and becomes also a pre-computable constant. First, we propose two heuristics to find single primary paths for PP. Then, the backup paths for PP are computed either by the above presented LP or by another heuristic. This heuristic also yields the path layout for SPM.

1) Primary Path Computation: Minimum Traffic Routing (MT): 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 path layout inducing a minimum traffic load on each network element.

a) Minimum Traffic Constraints: The traffic in all nodes is given by the auxiliary vector $\mathbf{a^V} \in (\mathbb{R}_0^+)^n$ and it is computed by $\mathbf{a^V} = \sum_{d \in \mathcal{D}} c(d) \cdot \mathbf{ti}(\mathbf{p_d})$. The idea is to minimize the maximum traffic through all nodes to a value a_{max}^V such that $\mathbf{a^V} \leq a_{max}^V \cdot \mathbf{1}$ holds. To avoid very long paths, the objective function takes also the overall required node capacity $\mathbf{1}^\top \mathbf{a^V}$ into account:

$$M^V \cdot a_{max}^V + \mathbf{1}^\top \mathbf{a^V} \rightarrow \min. \tag{9}$$

The constant $M^V \in \mathbb{R}_0^+$ controls the tradeoff between the conflicting goals "little maximum traffic per node a_{max}^V" and "little overall node capacity $\mathbf{1}^\top \mathbf{a^V}$" that have both to be minimized. A small M^V favors little overall node capacity while a large M^V favors little maximum traffic per node. We have chosen a value of $M^V = (|\mathcal{V}| + 1) \cdot |\mathcal{D}| = n \cdot (n+1)^2$ in our experiments.

b) 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, which takes 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 multipath structure.

2) Primary Path Computation: Shortest of k-Disjoint Shortest Path (kDSP): With the primary paths computed by the MT method, a link and node disjoint backup path cannot always be found although two disjoint paths exist in the network [28]. To guarantee the existence of k disjoint backup paths if topologically feasible, we propose to take the shortest path of a k (node and link) disjoint shortest paths solution (kDSP) with $k \geq 2$ [29], [30].

3) Backup Path Computation: Optimum Calculation (OPT): The optimum backup path solution for given primary paths can be obtained by a slight modification of the LP formulation in Section III-C. The primary paths $\mathbf{p_d}$ are removed from the set of free variables. This yields a LP formulation which can be solved efficiently and the corresponding results are the path layout and a load balancing function for all locations where the backup paths fork. However, the structure of the resulting backup path is potentially very complex since the partial paths are not necessarily disjoint. The following heuristic solves this problem.

4) Backup Path Computation: k-Disjoint Shortest Path (kDSP): We remove the links and (possibly) the inner nodes of the primary paths $\mathbf{p_d}$ from the network and calculate again a k-disjoint shortest paths solution that we use as for backup purposes. The results are at most k disjoint single paths, however, without a load balancing function which is calculated in Section III-E.3.

5) Path Layout for SPMs: We determine the disjoint parallel path for an SPM also by a k-disjoint shortest paths solution. There is no distinction between primary and backup paths and the corresponding load balancing function is calculated in Section III-E.

6) Adaptation to SRLGs: For the computation of disjoint multipaths we use the kDSP algorithm which is simple and efficient to compute. However, it does not take general SRLGs into account which is a different and NP-hard problem [31]. Basically, our kDSP heuristic can be substituted by any other routing scheme yielding disjoint multipaths.

E. Calculation of Load Balancing Functions

An SPM for a demand d 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.

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

$$\mathbf{f_d(s)} = \left(\phi(\mathbf{p_d^0}, \mathbf{s}), ..., \phi(\mathbf{p_d^{k_d-1}}, \mathbf{s}) \right)^\top. \tag{10}$$

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

b) 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

$$\mathbf{1}^\top \mathbf{l_d^f} = 1. \tag{11}$$

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

$$\mathbf{f}^\top \mathbf{l_d^f} = 0. \tag{12}$$

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

1) Load Balancing Heuristics for Disjoint Paths: There are many possibilities for load balancing over multipaths.

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

b) 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 $(l_d^f)_i = \frac{1 - f_i}{\mathbf{1}^\top (\mathbf{P_d})_i} / \left(\sum_{0 \leq j < k_d} \frac{1 - f_j}{\mathbf{1}^\top (\mathbf{P_d})_j} \right)$.

2) Optimized Load Balancing: Load balancing is optimal if the required capacity $\mathbf{b}$ is minimal to cover all demands $d \in \mathcal{D}$ for all protected failure scenarios $\mathbf{s} \in \mathcal{S}$. We formulate a LP to describe the solution. The free variables are

$$\mathbf{b} \in (\mathbb{R}_0^+)^m, \quad \forall d \in \mathcal{D} \; \forall \mathbf{f} \in \mathcal{F}_d : \mathbf{l_d^f} \in (\mathbb{R}_0^+)^{k_d}. \tag{13}$$

The objective function is given by Equation (5). It must be minimized under load balancing and bandwidth constraints. The load balancing constraints in Equations (11) and (12) must be respected by all $\mathbf{l_d^f}$ and the bandwidth constraints are newly formulated.

a) Bandwidth Constraints with Capacity Reuse:

$$\forall \mathbf{s} \in \mathcal{S} : \sum_{d \in \mathcal{D}_\mathbf{s}} \mathbf{P_d}^\top \mathbf{l_d^{f_d(s)}} \cdot c(d) \leq \mathbf{b}. \tag{14}$$

b) Bandwidth Constraints without Capacity Reuse: Releasing capacity unnecessarily leads to a waste of bandwidth if it cannot be reused by other connections. Therefore, load balancing factors $\mathbf{l_d^f}$ 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

$$\mathbf{l_d^f} + \mathbf{f} \geq \mathbf{l_d^{f_d(0)}}, \tag{15}$$

where $\mathbf{l_d^{f_d(0)}}$ 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

$$\forall \mathbf{s} \in \mathcal{S}: \underbrace{\sum_{d \in D_\mathbf{s}} c(d) \cdot \mathbf{P_d}^\top \mathbf{l_d^{f_d(s)}}}_{\text{used capacity}} +$$

$$\underbrace{\sum_{d \in D_\mathbf{s}} c(d) \cdot \mathbf{P_d}^\top (\mathbf{f_d(s)} \cdot \mathbf{l_d^{f_d(0)}})}_{\text{inactive partial paths}} + \underbrace{\sum_{d \in D \setminus D_\mathbf{s}} c(d) \cdot \mathbf{P_d}^\top \mathbf{l_d^{f_d(0)}}}_{\text{failed connections}} \leq \mathbf{b}. \tag{16}$$

Note that the term $\mathbf{f_d(s)} \cdot \mathbf{l_d^0}$ expresses an element-wise multiplication of two vectors. Hence, if bandwidth reuse is possible, Equation (14) is used as bandwidth constraint, otherwise Equations (15) and (16) must be respected. Neither protection constraints (Equations (3) and (4)) nor path constraints (Equations (1) and (2)) apply.

3) Adaptation to Path Protection: The adaptation of the above explained load balancing scheme to path protection mechanisms is simple. We describe the primary paths $\mathbf{p_d}$ 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 $\mathbf{f_d^{PP}(s)}$ is described by

$$\mathbf{f_d^{PP}(s)} = \begin{cases} \mathbf{u^0} & \phi(\mathbf{p_d}, \mathbf{s}) = 0 \\ \mathbf{f_d(s)} & \phi(\mathbf{p_d}, \mathbf{s}) = 1 \end{cases} \tag{17}$$

with $\mathbf{u^0} = (0, 1, ..., 1)^\top$. By substituting the inactivity pattern in Equation (10) by Equation (17), the load balancing optimization in Section III-E.2 can also be applied to path protection schemes.

IV. Numerical Results

In this section we compare the extra capacity required for resilience purposes of the presented protection mechanisms. We determine the required network capacity, i.e. the sum of all link bandwidths, which is needed 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.

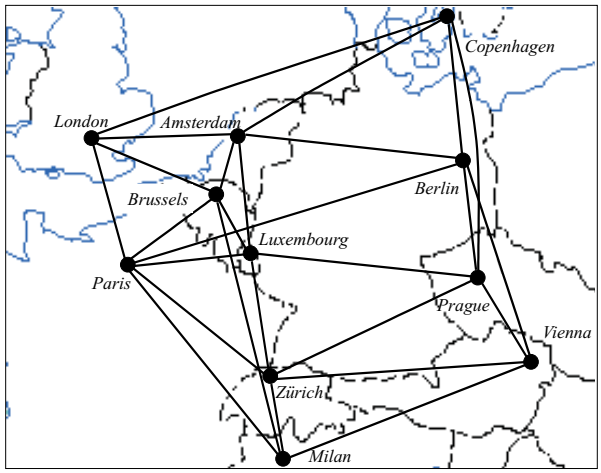

Fig. 1. The backbone topology of the COST239 network.

We compare the backup performance of the path protection schemes and the self-protecting multipath in the COST 239 core network [32] (11 routers, 26 bidirectional links in Europe, cf. Figure 1) and in the Labnet [33] (20 routers, 53 bidirectional links in US, cf. Figure 2).

- At the beginning, we briefly recall all discussed protection methods.
- We test the influence of multipath routing and load balancing together with different alternatives for primary and backup paths layout in the above networks.
- We study then the most promising mechanisms in additional sample networks of the literature.
- To relate the performance of the SPM to other mechanisms, we compare it with the backup requirements for p-cycles.
- We use homogeneous traffic matrices, full traffic reduction, bandwidth reuse, and the protection of single router and link failures as default since 30% of all network failures are due to router failures and 70% of them are due to link failures [34].
  These side conditions have of course a significant impact on the required backup capacity. We have investigated them in other papers and summarize finally their results very shortly to give a complete picture of the backup performance regarding SPM, PP, and shortest path routing.

A. Overview of Investigated Protection Mechanisms and Abbreviations

For the sake of an easier understanding, we recall the discussed protection switching mechanisms and define abbreviations. With path protection (PP), a primary single-path is protected by a backup multipath. 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

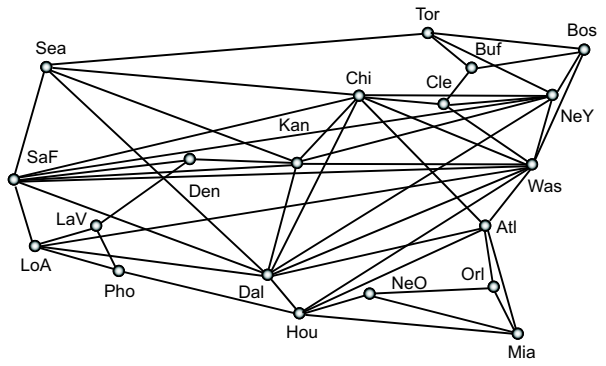

Fig. 2. Topology of the Labnet network.

through each router (MT). The backup multipath may be computed together with an appropriate load balancing scheme by a LP optimization (OPT) which does not necessarily yield disjoint paths. As an alternative, the $(k{-}1)$-disjoint shortest paths ($(k{-}1)$DSP) may be taken. 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). The self-protecting multipath (SPM) is different from PP. It consists of disjoint parallel path that are obtained by a k-disjoint shortest paths solution (kSPM). In single failure scenarios, a kSPM leads to at most $k{+}1$ 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 chosen like above (E, R, O).

In the following, we mainly use these abbreviations to refer to specific protection mechanisms. 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 multipath together with a load balancing scheme is computed by a LP for PP. Finally, 5SPM-E signifies a SPM consisting of up to 5 disjoint paths with equal load balancing. Shortest paths routing as used in OSPF or IS-IS is denoted by "OSPF".

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.

B. Impact of Multipath Routing and Load Balancing

We investigate the impact of multipath routing and load balancing on the backup performance. First, we consider path protection schemes and then we study the self-protecting multipath.

a) PP Schemes: Figures 3 and 4 show the required backup capacity in the COST239 and the Lab03 network for all path protection schemes ({kDSP,MT}-{$(k-1)$DSP-{E,R,0},OPT} with $2 \leq k \leq 5$). The following observations are valid for primary paths found by MT and by kDSP.

For $k=2$, only one backup path is available. If a primary path fails, 100% of the traffic is transported over the remaining path, i.e., the performance of all load balancing alternatives (E, R, O) coincides. For larger k, more disjoint backup paths are available and the traffic can be better distributed in a failure case. Therefore, less extra capacity is required on the backup links. The most striking performance gain is achieved for taking $k=3$ instead of $k=2$. The reason is the following. Even for $k \geq 4$, only 3 disjoint path can be found because of the network topology. Therefore, the reduction of the required backup capacity by multipath routing is limited.

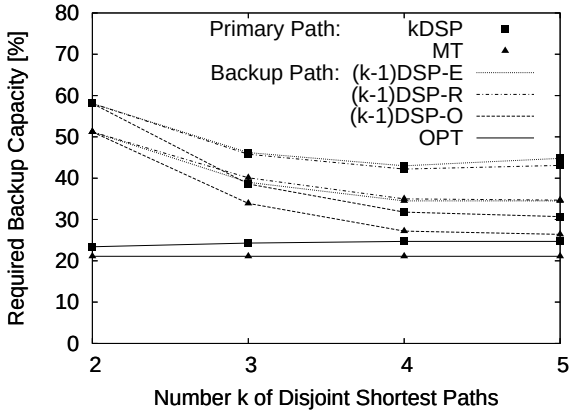

Fig. 3. Impact of multipath routing and load balancing for path protection methods in the COST-239 network.

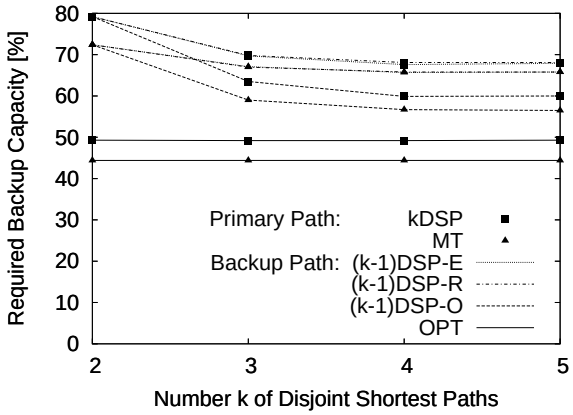

Fig. 4. Impact of multipath routing and load balancing for path protection methods in the Lab03 network.

The layout of the primary path depends on the heuristic (MT or kDSP for a specific k). It 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 {MT,kDSP}-OPT PP mechanisms are most efficient because the backup path is not limited by k disjoint shortest paths. As a consequence, the performance of {MT,kDSP}-

OPT is almost independent of k. However, complex multipath structures are hard to deploy and to manage in failure cases. In addition, the backup path computation is very time consuming. Therefore, disjoint multipaths are desired for backup purposes although they require significantly more capacity. Equal and reciprocal load balancing for the backup multipath lead approximately to the same results. Optimization of the load balancing function reduces the required extra capacity by about 10 percent points.

If a large k effects a longer primary path, more capacity is required for normal operation without failure. In contrast to the load balancing options R and O, the load balancing option E cannot compensate the increased capacity requirements by load distribution because it is insensitive to the length of the primary path. As a result, slightly more capacity is required for 5DSP-4DSP-E than for 4DSP-3DSP-E in the COST-239 network.

b) SPM: Figures 5 and 6 show the required backup capacity in the COST 239 and the Lab03 network for various SPMs (kSPM-$\{$E,R,O$\}$) in comparison with the best PP schemes (MT-($k-1$)DSP-O and MT-OPT).

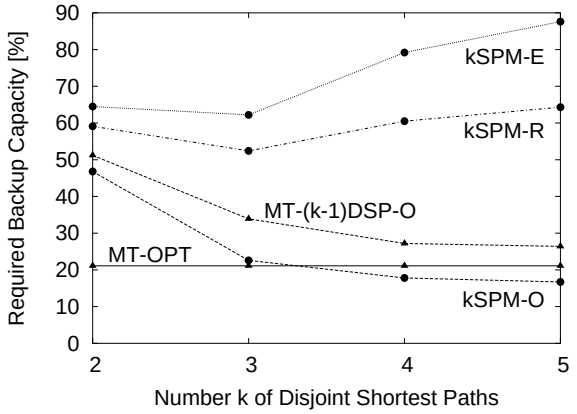

Fig. 5. Impact of multipath routing and load balancing for the self-protecting multipath in the COST-239 network.

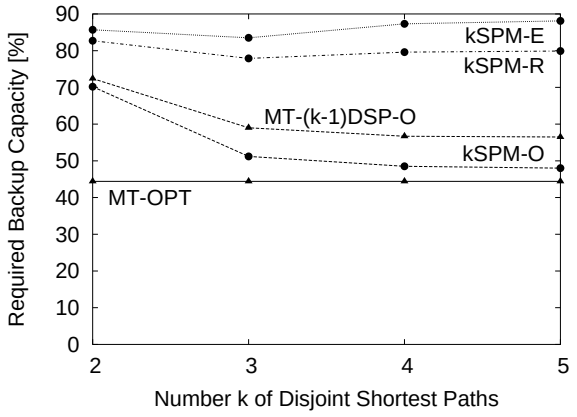

Fig. 6. Impact of multipath routing and load balancing for the self-protecting multipath in the Lab03 network.

In contrast to the PP methods, the load balancing function (E, R, O) has a greater impact on the backup performance of SPMs than for PP methods and their impact increases with the maximum number of parallel paths k. Although a capacity reduction is expected due to increased path diversification in failure cases, the backup performance of kSPM-E and kSPM-R degrades considerably with increasing k in the COST-239 network. In the Lab03 network, it stays about constant. If k increases, longer paths join the SPM. The SPM with equal or reciprocal load balancing (kSPM-E or kSPM-R) cannot avoid their extensive use which leads to more required network capacity. Hence, multipath routing for SPM with only simple load balancing schemes reveals no or minor benefits.

Optimized load balancing reduces the required backup capacity of the SPM considerably and the potential savings increase with path diversification. 5SPM-O is about 10 percent points superior to MT-4DSP-O in both networks, which has been proven to be the best feasible PP solution. In the COST 239 network, 5SPM-O is even better than MT-OPT. It requires only 17% additional capacity to protect the network against all link and router failures.

We motivate the superiority of the SPM by the following explanation. With the multipath routing of the SPM, each link carries traffic from more aggregates than with single primary path routing of PP, but it carries only a fraction of their traffic. In case of a link failure, the traffic of more aggregates is affected such that the load of the failed link can be spread out over more backup paths than with single primary paths and PP. As a consequence, less shareable backup capacity is required on the individual links.

Like above, there is only a single backup path for $k=2$ 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 such a way that strong traffic concentrations are prevented in any network element. This avoids that a large traffic rate must be redirected if this element fails. This idea is similar to the MT heuristic, which helps to find suitable primary paths.

C. Impact of Network Topologies

Figure 7 shows the required backup capacity for various protection mechanisms in various example networks. A point in the figure stands for a certain network and protection mechanism. The x-axis indicates the average number of disjoint parallel paths k^* between any two nodes in the respective network and the y-axis indicates the required backup capacity. The studied protection switching mechanisms are simple OSPF rerouting, 5DSP-4DSP-O, MT-4DSP-O, 5DSP-OPT, MT-OPT, and 5SPM-O, and their corresponding required backup capacities 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. Therefore, they have a higher average node degree than their lowercase counterparts. Note that the MT-5DSP and MT-OPT protection mechanisms

are missing for some networks because for some failure cases no backup paths 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-O according to an exponential function. Furthermore, the relative savings compared to OSPF rerouting increase with increasing k^*. The SPM is superior to all feasible PP schemes. That can be explained as follows. A kDSP-$(k-1)$DSP-O is structurally very similar to a kSPM because they use the same disjoint paths of a kDSP computation. But due to the limitation of Equation (17), the optimization of the load balancing function for path protection methods has fewer degrees of freedom, so, comparable SPMs require less 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 multipath backup structure. Hence, the SPM is the best of all feasible solutions in all investigated networks.

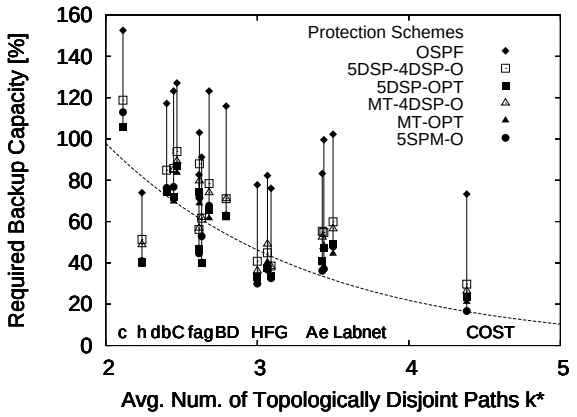

Fig. 7. Comparison of protection switching mechanisms in example networks.

D. Comparison of the SPM with p-Cycles

In [5], [35] the p-cycle concept has been investigated. An optimal p-cycle layout has been found to protect the network with the least capacity possible using a maximum cycle length as side constraint. The experiments were also conducted with the COST-239 network but with the original and partly asymmetric traffic matrix which is given in [36]. The most effective solution required 44% more backup-capacity-related to the capacity requirements for shortest path routing based on the hop count without resilience. For comparison reasons, we calculate the performance value for the 5SPMO and get an additional bandwidth of 23.4%.

E. Impact of Other Parameters

We have seen above that the required capacity depends significantly on the network topology (cf. Section IV-C) and on the traffic matrix (cf. Section IV-D). To complete this study, we have investigated these issues in [8] and [37] report briefly on our findings.

1) Impact of the Network Topology: In [8] we investigated the influence of the network topology. We simulated random networks and controlled their size and node degree. The average node degree has the major impact on the required backup capacity because it limits the number of disjoint paths. This number is crucial as the superiority of the SPM is due to multipath forwarding and load balancing. Also PP mechanisms require only little extra capacity but they are mostly worse than SPM. Unlike the SPM, OSPF rerouting is not able to profit from a well connected network and, therefore, the superiority of the SPM over OSPF rerouting grows with increasing node degree. The size of the network had no significant impact on the required extra capacity.

2) Impact of the Traffic Matrix: In [37] we showed that the traffic matrix has a tremendous impact on the required backup capacity. We also considered the topologies of the COST-239 and the Labnet networks. The required backup capacity for the SPM amounts to 17% for a homogeneous traffic matrix, 23% for a realistic traffic matrix, and 67% for an extremely skewed and, therefore, also unrealistic traffic matrix. The extra capacity for OSPF rerouting was 72%, 78%, and 114%, respectively. Hence, the SPM saves in all cases at least 55% backup capacity. This shows that the superiority of the SPM over shortest path rerouting remains for various traffic matrices. The PP mechanisms behave similarly, they are worse than SPM, but better than OSPF rerouting.

3) Impact of the Traffic Reduction, Protection, and Bandwidth Reuse Options: In [37] we have investigated the traffic reduction, protection, and bandwidth reuse options for the calculation of the required backup capacities for the SPM mechanism. The traffic reduction options have hardly any impact on the required network capacity. In sufficiently large networks, link protection is less demanding than router and full protection. The bandwidth reuse saved about 5% backup capacity in the studied networks.

V. Conclusion

If a link or node failure occurs in a resilient network, the traffic is quickly deviated around the outage location by protection switching mechanisms. In this paper, we have described the self-protecting multipath (SPM) and some variants for that purpose. They are based on multipath structures consisting only of disjoint paths and corresponding path failure specific load balancing functions. They are simple to implement because they do not require the notification of the explicit failure location and do only redirect traffic aggregates that are affected by the failure.

The objective of our work was the calculation of the path layout and the load balancing functions for these mechanisms such that the required backup capacity is minimized for a given network topology and traffic matrix. The optimization is based on heuristic algorithms and polynomial-time linear programs (LP). Our numerical results showed that our LP-optimization of the load balancing functions reduces the backup capacity significantly together with multipath routing.

In contrast, simple load balancing heuristics do not help much. The SPM is the simplest and most efficient one of our investigated backup solutions. It requires only 17% backup capacity to protect all single link and node failures in the COST239 network for a homogeneous traffic matrix while conventional shortest path (re)routing needs about 72% in this case. Of course, the backup capacity depends on the network topology and the traffic matrix but we have shown that the superiority of the SPM over shortest path (re)routing remains for other networks and heterogeneous traffic matrices. The p-cycle approach is another well known protection switching mechanism which has also been recommended to save extra capacity for protection purposes. Its minimum backup capacity in the COST239 network is 44% for a "real life" traffic matrix [5], [35] while our calculation for the SPM approach requires only 23.4% backup capacity under the same conditions.

Currently, we are working on the configuration of the SPM for existing networks with given link bandwidths. We would like to investigate the impact of multiple failures on the QoS degradation in networks that are resilient to single failures. Suitable network structures are a prerequisite for cheap backup capacities and should be further identified.

REFERENCES

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

[2] Ping Pan, George Swallow, and Alia Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels," http://www.ietf.org/internet-drafts/draft-ietf-mpls-rsvp-lsp-fastreroute-07.txt, Aug. 2004.

[3] Achim Autenrieth and Andreas Kirstädter, "Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS," *IEEE Communications Magazine*, vol. 40, no. 1, pp. 50–57, Jan. 2002.

[4] Wayne D. Grover, "Cycle-Oriented Distributed Preconfiguration: Ring-Like Speed with Mesh-Like Capacity for Self-Planning Network Restoration," in *IEEE International Conference on Communications (ICC)*, Jun 1998, pp. 537–543.

[5] Claus G. Gruber and Dominic A. Schupke, "Capacity-Efficient Planning of Resilient Networks with p-Cycles," in *International Telecommunication Network Strategy and Planning Symposium (Networks)*, June 2002, pp. 389–395.

[6] K. Murakami and H. S. Kim, "Comparative Study on Restoration Schemes of Survivable ATM Networks," in *IEEE Infocom*, Kobe City, Japan, April 1997, pp. 345 – 352.

[7] Kazutaka Murakami and Hyong S. Kim, "Optimal Capacity and Flow Assignment for Self–Healing ATM Networks Based on Line and End-to-End Restoration," *IEEE/ACM Transactions on Networking*, vol. 6, no. 2, pp. 207–221, Apr. 1998.

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

[9] Michal Pióro and Deep Medhi, *Routing, Flow and Capacity Design in Communication and Computer Networks*, Morgan and Kaufman, June 2004.

[10] Bernard Fortz and Mikkel Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights," in *IEEE Infocom*, Tel-Aviv, Israel, 2000, pp. 519–528.

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

[12] Stefan Schnitter and Martin Horneffer, "Traffic Matrices for MPLS Networks with LDP Traffic Statistics," in *International Telecommunication Network Strategy and Planning Symposium (Networks)*, Vienna, Austria, June 2004, pp. 231 – 236.

[13] Ivan Gojmerac, Thomas Ziegler, Fabio Ricciato, and Peter Reichl, "Adaptive Multipath Routing for Dynamic Traffic Engineering," in *IEEE Globecom*, San Francisco, CA, Nov. 2003.

[14] Zhiruo Cao, Zheng Wang, and Ellen Zegura, "Performance of Hashing-Based Schemes for Internet Load Balancing," in *IEEE Infocom*, Tel Aviv, Israel, 2000.

[15] Gero Dittmann and Andreas Herkersdorf, "Network Processor Load Balancing for High–Speed Links," in *International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*, San Diego, CA, 2002, pp. 727–735.

[16] Murali S. Kodialam and T. V. Lakshman, "Minimum Interference Routing with Applications to MPLS Traffic Engineering," in *IEEE Infocom*, Mar 2000, vol. 2, pp. 884–893.

[17] Guangzhi Li, Dongmei Wang, Charles Kalmanek, and Robert Doverspike, "Efficient Distributed Path Selection for Shared Restoration Connections," in *IEEE Infocom*, 2002.

[18] L. Sahasrabuddhe, S. Ramamurthy, and B. Mukherjee, "Fault Tolerance in IP-Over-WDM Networking: WDM Protection vs. IP Restoration," *IEEE Journal on Selected Areas in Communications(Special Issue on WDM-Based Network Architectures)*, vol. 20, no. 1, pp. 21–33, Jan. 2002.

[19] A. Nucci, Bianca Schroeder, S. Bhattacharyya, Nina Taft, and Christophe Diot, "IGP Link Weight Assignment for Transient Link Failures," in 18^{th}*International Teletraffic Congress (ITC)*, Berlin, Sept. 2003.

[20] Rainer R. Iraschko, M. H. MacGregor, and Wayne D. Grover, "Optimal Capacity Placement for Path Restoration in STM and ATM Mesh-Survivable Networks," *IEEE/ACM Transactions on Networking*, vol. 6, no. 3, pp. 328 – 336, June 1998.

[21] John Strand, Angela L. Chiu, and Robert Tkach, "Issues for Routing in the Optical Layer," *IEEE Communications Magazine*, vol. 39, no. 2, pp. 81–87, Feb 2001.

[22] B. Rajagopalan, J. Luciani, and D. Awduche, "RFC3717: IP over Optical Networks: A Framework," ftp://ftp.isi.edu/in-notes/rfc3717.txt, Mar. 2004.

[23] Kireeti Kompella and Yakov 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.

[24] Jiang Wang, Laxman Sahasrabuddhe, and Biswanath Mukherjee, "Path vs. Subpath vs. Link Restoration for Fault Management in IP-over-WDM Networks: Performance Comparisons Using GMPLS Control Signaling," *IEEE Communications Magazine*, vol. 40, no. 11, pp. 80–87, Nov. 2002.

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

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

[27] Hiroyuki Saito and Makiko Yoshida, "An Optimal Recovery LSP Assignment Scheme for MPLS Fast Reroute," in *International Telecommunication Network Strategy and Planning Symposium (Networks)*, June 2002, pp. 229–234.

[28] D. Anthony Dunn, Wayne D. Grover, and Mike 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.

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

[30] Jack Edmonds and Richard M. Karp, "Theoretical Improvements in the Algorithmic Efficiency for Network Flow Problems," *Journal of the ACM*, vol. 19, no. 2, pp. 248–264, Apr. 1972.

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

[32] P. Batchelor et al., "Ultra High Capacity Optical Transmission Networks. Final Report of Action COST 239," Jan. 1999.

[33] Michael Menth, Stefan Kopf, and Jens Milbrandt, "A Performance Evaluation Framework for Network Admission Control Methods," in *IEEE Network Operations and Management Symposium (NOMS)*, Seoul, South Korea, Apr. 2004, pp. 307 – 320.

[34] Athina Markopoulou, Gianluca Iannaccone, Supratik Bhattacharyya, and Chen-Nee Chuah, "Characterization of Failures in an IP Backbone," in *IEEE Infocom*, Hongkong, Mar. 2004.

[35] Dominic A. Schupke, Claus G. Gruber, and Achim Autenrieth, "Optimal Configuration of p-Cycles in WDM Networks," in *IEEE International Conference on Communications (ICC)*, New York, 2002.

[36] Christian Mauz, "Mapping of Arbitrary Traffic Demand and Network Topology on a Mesh of Rings Network," in *IFIP Working Conference on Optical Network Design and Modelling*, Feb. 2001.

[37] Michael Menth, Jens Milbrandt, and Andreas Reifert, "Sensitivity of Backup Capacity Requirements to Traffic Distribution and Resilience Constraints," in 1^{st}*Conference on Next Generation Internet Networks Traffic Engineering (NGI)*, Rome, Italy, Apr. 2005.