

University of Würzburg  
Institute of Computer Science  
Research Report Series

**Sensitivity of Backup Capacity  
Requirements to Traffic Distribution  
and Resilience Constraints**

Michael Menth, Jens Milbrandt<sup>1</sup> and Andreas  
Reifert<sup>2</sup>

Report No. 322

February 2004

<sup>1</sup> 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

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



# Sensitivity of Backup Capacity Requirements to Traffic Distribution and Resilience Constraints

## 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 work, we present several end-to-end protection switching mechanisms for application in Multiprotocol Label Switching (MPLS). In case of local outages in the network, they deviate the traffic around the failed element over backup paths. They are easy to implement and reduce the additional capacity to maintain the Quality of Service (QoS) on the backup paths. We study the capacity savings of the presented methods for various protection schemes with different traffic matrices. We further test the influence of different resilience constraints such as the set of protected failure scenarios, bandwidth reuse restrictions due to optical communication, and traffic reduction due to failed border routers.

**Keywords:** Resilience, Protection Switching, Network Dimensioning

## 1 Introduction

Carrier grade networks are expected to provide Quality of Service (QoS) in terms of packet loss and delay, to offer 99.999% availability, and high reliability even in case of network failures. This challenge arises, e.g., for virtual private networks (VPNs) or in the terrestrial radio access network (UTRAN) of the Universal Mobile Telecommunication System (UMTS). Today's Internet Protocol (IP) technology enables a global and best effort interconnection of remote hosts and servers. The best effort service does not meet the requirements of carrier grade networks but the wide deployment and the simple operation of current IP networks calls for a *Next Generation Networks* (NGN) based on IP to substitute frame relay and ATM solutions.

Conventional telephone networks achieve the five nines reliability by massive redundancy of hardware provisioning. With packet-switched networks, a similar reliability could be achieved by traffic rerouting in case of local outages, e.g., link or router failures. With connection-oriented technologies like Multiprotocol Label Switching (MPLS) route pinning is possible which provides the setup of label switched paths (LSPs) with explicit routes. This

enables a finer control of routing in outage scenarios. Such mechanisms are called protection switching. In [1] we have presented different end-to-end (e2e) or border-to-border (b2b) protection switching methods that work with minimum information about network failures, i.e., an ingress router switches traffic of a failed path locally onto backup paths. Multi-path structures and load balancing provide degrees of freedom for the minimization of required backup bandwidth if backup capacity sharing is possible. In [2], we have investigated the performance of these mechanisms with regard to required backup capacity in different network topologies. The Self-Protecting Multi-paths (SPM) is the most efficient solution for protection switching and its performance depends on the network topology.

The contribution of this paper is the investigation of the proposed protection switching mechanisms for different traffic matrices and under various resilience constraints such as the set of protected failure scenarios, bandwidth reuse restrictions due to optical communication, and traffic reduction due to failed border routers.

This paper is organized as follows. Section 2 describes the protection switching mechanisms in detail, the applied traffic model, and various resilience constraints. The numerical results in Section 3 demonstrate the performance of our protection switching mechanisms under the discussed side conditions. Section 4 summarizes this work.

## 2 Protection Switching Mechanisms

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 LSPs. For IP routing, a similar approach can be done by setting the link cost appropriately such that all traffic is transported through the network and that the mean and maximum link utilization is minimized [3]. Pure IP and MPLS solutions may also be combined [4]. These approaches require the knowledge of the traffic matrix which is usually not known for best effort traffic. This problem is tackled by [5] 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 [6] achieves that goal very well. The authors of [7] 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 [8] 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 [9, 10] 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 [9, 10] 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 [9, 10]. 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 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) [11, 12, 13] which group network elements together that may fail simultaneously with a high probability. For instance, all links originating at the same router fail if the router goes down. SRLGs

are motivated by optical networking where a single optical fiber duct accommodates several logically separate links. In our work, we consider only the first scenario and the second one in a trivial way by excluding parallel links. However, we do not take general SRLGs into account because our focus is the performance evaluation of basic protection switching mechanisms and not their adaptation to SRLGs.

### 2.3 Protection Switching Mechanisms for Backup Capacity Reduction

We present several protection switching mechanisms and explain the idea of our algorithms for their layout.

Path protection (PP) mechanisms transport the traffic usually on a primary single-path and use the multi-path backup structure only in case that the primary path fails due to some failure in the network. We have proposed two different solutions to calculate the primary path. The  $k$ DSP approach takes the shortest path of a  $k$  disjoint shortest path solution [14, 15]. This guarantees that link and node disjoint backup paths can be found afterwards if they are topologically feasible. Another routing approach is minimum traffic (MT) routing, i.e., the primary paths are chosen in such a way that the maximum traffic rate traversing a node is kept small. An optimum multi-path backup structure (OPT) together with a load balancing function for each node can be derived by a general, computation intensive linear program (LP). However, the result is a multi-path where partial paths may fork and join, hence, it is not suitable for implementation purposes. The computation of up to  $k - 1$  disjoint shortest paths like above is the preferred solution for practical application. Based on this structure, an optimized load balancing for the multi-path is computed for the case that the primary path fails. This optimization is based on a non-integer LP that runs quite fast.

The Self-Protecting Multi-Path (SPM) consists of up to  $k$  paths of a  $k$ DSP computation. In contrast to PP, the traffic is distributed onto all paths in the no failure case, too. If one partial path fails, the traffic is redistributed onto the working paths by a *path failure* specific load distribution function. Like above, the load distribution functions are derived based on a non-integer LP.

In the following, we mainly use abbreviations to refer to specific protection mechanism. For example, 5DSP-4DSP-O means that the single primary b2b path is chosen as the shortest path from a 5-disjoint shortest path solution and the other (at most) 4 are taken for path protection. Load balancing is done in an optimal way by a non-integer LP. 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-O signifies an SPM consisting of up to 5 disjoint paths with optimal load balancing.

### 2.4 Resilience Constraints

Network resilience is a soft expression as it means fault tolerance against a set of faulty networking scenarios that adhere to some assumptions. We present them in the following.

### 2.4.1 Protected Failure Scenarios

The optimization of protection switching mechanisms requires a set of protected failure scenarios  $\mathcal{S}$  which contains by default the working scenario. We consider three different options. “Link protection” takes only all single bidirectional link failures into account, “router protection” respects only single router failures, and we call the consideration of both single bidirectional link and router failures “full protection”.

### 2.4.2 Traffic Reduction

In normal operation without any failures, all b2b aggregates are active. If routers fail, some of them may disappear. We consider several options. If network nodes lose only their capability to transport transit flows but if they are still able to generate traffic, then we talk about “no traffic reduction”. If failed nodes stop sending traffic, we talk about “source traffic reduction”. We have “full traffic reduction” if traffic is stalled if either its source or destination node does not work.

### 2.4.3 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. Hence, “bandwidth reuse” is possible. 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 the “no bandwidth reuse” option because network resources allocated by failed paths can not be reused for backup purposes.

## 3 Experimental Setup

We shortly describe the studied network topologies and the model for the traffic matrices that are used for an analytical calculation of the required capacity based on the equations in [1].

### 3.1 Test Networks

We want to evaluate the performance of protection switching mechanism in the context of carrier grade networks. Therefore, we focus on core network structures by means of two topologies taken from operational networks.

The network structure is described by graph notation, i.e., the topology is given by  $\mathcal{N} = (\mathcal{V}, \mathcal{E})$  where the set of vertices  $\mathcal{V}$  contains all routers and the set of edges  $\mathcal{E}$  contains all bidirectional links. The number of bidirectional edges leaving a node  $v$  is called node degree  $deg(v)$ . The average node degree can be computed by  $deg_{avg} = \frac{2 \cdot |\mathcal{E}|}{|\mathcal{V}|}$ .

The network in Figure 1 is the optical core of the infrastructure in the COST-239 project [16]. The project was part of the “European Co-operation in the Field of Scientific and Technical Research” and concentrated on ultra-high capacity optical transmission networks. The

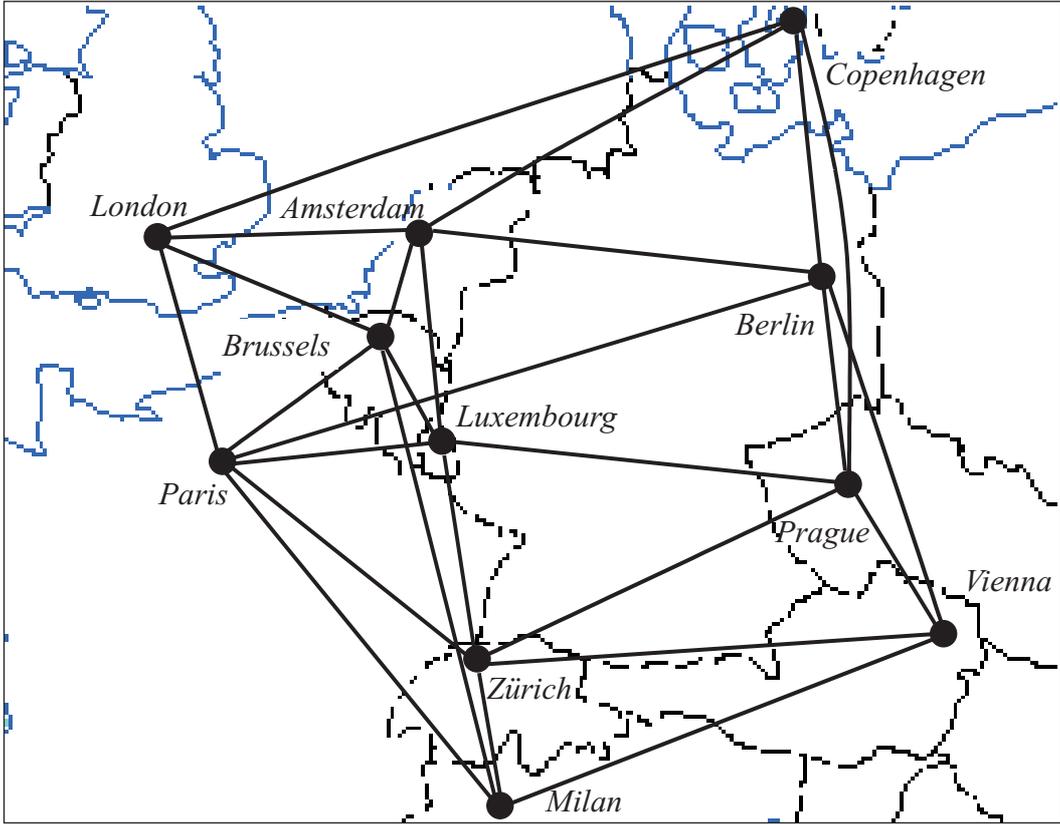


Figure 1: The COST-239 core network.

Lab03 network in Figure 2 is taken from the testbed of the KING project [17]. It is a modification of the UUNET in 1994 where all nodes with a node degree of at most 2 are successively removed. We use both networks in our performance evaluation because they have different structural properties.

### 3.2 Traffic Matrices

The overall offered traffic by all b2b aggregates is denoted by  $c_{tot}$ . We create a traffic matrix proportional to the populations  $\pi(v)$  associated with the respective nodes  $v$ . The calculation of the b2b aggregate traffic  $c(g_{v,w})$  is based on the populations given in Tables 1 and 2 and the following equation:

$$c(g_{v,w}) = \begin{cases} \frac{c_{tot} \cdot \pi(v) \cdot \pi(w)}{\sum_{x,y \in \mathcal{V}, x \neq y} \pi(x) \cdot \pi(y)} & \text{for } v \neq w, \\ 0 & \text{for } v = w. \end{cases} \quad (1)$$

As the traffic matrix has possibly a significant impact on the required backup capacity, we distort its structure. The distortion is still based on  $c_{tot}$  and the node populations  $\pi$  according

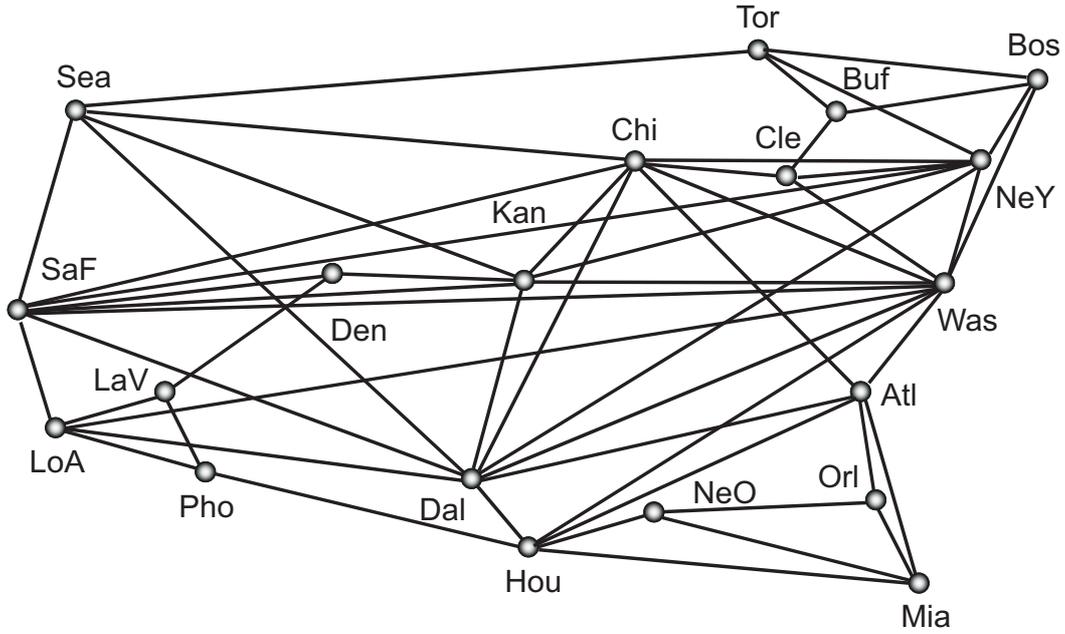


Figure 2: The Lab03 network.

to Equation (1) but we modify  $\pi$  using an exponential extrapolation with parameter  $t$ :

$$\pi(v, t) = |\mathcal{V}| \cdot \bar{\pi} \cdot \frac{\exp(\delta(v) \cdot t)}{\sum_{v \in \mathcal{V}} \exp(\delta(v) \cdot t)}, \quad (2)$$

with  $\bar{\pi} = \sum_{v \in \mathcal{V}} \pi(v)$ . The value  $\delta(v)$  is determined by  $\pi(v, 1) = \pi(v)$ , i.e.  $\delta(v) = \ln(\frac{\pi(v)}{\bar{\pi}})$ . According to that construction, the traffic matrix based on the original population  $\pi$  and  $\pi(\cdot, 1)$  are the same.

Table 3 describes the effect of the extrapolation on the city sizes. All city sizes are equal for  $t = 0$ . As a consequence, all b2b aggregates carry the same traffic. If a city is larger than the average city size ( $\pi(v) > \bar{\pi}$ ), it is scaled down by a negative value of  $t$  and it is scaled up for a positive value of  $t$ . With increasing  $|t|$ , the number of cities below the average size increases and the number of cities above the average size decreases. Therefore, the coefficient of variation of the city sizes increases. As a consequence, most of the traffic flows among fewer cities, which impacts the coefficient of variation of the entries in the traffic matrix. We consider the average path length ( $len_{path}^{avg}$ ) weighted by the corresponding traffic. Large cities (for  $t = 1$ ) are usually connected closer among each other than smaller cities. If they grow in size, the traffic aggregates among them gain in proportion of the overall traffic and the hop distance among them dominates the average path length. Thus, the average path length in Table 3 decreases with increasing  $t$  for the Lab03 network and it shows two local minima for COST-239. These observations show that the traffic matrix has changed considerably.

Table 1: Population of the cities and their surroundings for the Lab03 network.

$name(v)$	$\pi(v) [10^3]$	$name(v)$	$\pi(v) [10^3]$
Atlanta	4112	Los Angeles	9519
Boston	3407	Miami	2253
Buffalo	1170	New Orleans	1338
Chicago	8273	New York	9314
Cleveland	2250	Orlando	1645
Dallas	3519	Phoenix	3252
Denver	2109	San Francisco	1731
Houston	4177	Seattle	2414
Kansas	1776	Toronto	4680
Las Vegas	1536	Washington	4923

Table 2: Population of the respective countries for the COST-239 network.

$name(v)$	$\pi(v) [10^3]$	$name(v)$	$\pi(v) [10^3]$
Amsterdam (NL)	16101	Paris (F)	59343
Berlin (D)	82360	Prague (CZ)	10300
Bruxelles (B)	10292	Rome (I)	58018
Copenhagen (DK)	5363	Vienna (A)	8141
London (UK)	60075	Zurich (CH)	7261
Luxembourg (L)	447		

## 4 Results

We compute optimized path protection structures and OSPF rerouting for both example networks. If not mentioned differently, we use a homogeneous traffic matrix (i.e.  $t = 0$ ), full protection, no traffic reduction, and bandwidth reuse as resilience constraints. We dimension the links appropriately and  $c(\mathcal{N}_S)$  denotes the sum of their capacities. The sum of all link capacities  $c(\mathcal{N})$  for OSPF routing without resilience requirements is the reference case. We express the required backup capacity relative to this reference case by  $(\frac{c(\mathcal{N}_S)}{c(\mathcal{N})} - 1) \cdot 100\%$ .

### 4.1 Impact of Traffic Matrices

Figures 3 and 4 show the required backup capacity in the COST-239 and the Lab03 network for different traffic matrices that are obtained as described in Section 3. All previously discussed protection switching mechanisms are compared. Solid lines stand for well implementable solutions. The dashed lines refer to general multi-paths as backup paths and are rather of theoretical interest. The curves in both figures distinguish significantly in their absolute shape but all have a minimum of required backup capacity in common. The required backup capacity in the COST-239 network increases for all mechanisms clearly with the variation of the traffic matrix whereas the minimum capacity for the Lab03 network is obtained for  $t = 1$  which corresponds to the most realistic traffic matrix. Hence, both the network topology and the

Table 3: Properties of extrapolated node populations.

$t$	-3	-2	-1	0	1	2	3
<b>Lab03</b>							
$c_{var}[\pi(\cdot, t)]$	7.88	2.62	0.78	0	0.69	2.02	5.29
$len_{path}^{avg}$	2.91	2.68	2.43	2.15	1.91	1.77	1.72
<b>COST-239</b>							
$c_{var}[\pi(\cdot, t)]$	3.16	3.10	2.28	0	0.98	1.36	1.56
$len_{path}^{avg}$	1.64	1.54	1.48	1.56	1.53	1.54	1.57

traffic matrix have an impact on the required capacity.

The major difference in required backup capacity results from the protection switching mechanisms. The difference in required backup capacity between OSPF and the protection switching mechanisms is evident for all traffic matrices. The Self-Protecting Multi-Path (SPM) is by far the most efficient well implementable one and outperforms often even MT-OPT and 5DSP-OPT. It requires only 17% additional capacity to protect the network against all link and router failures for  $t = 0$  in Figure 3. The curves for MT-OPT and MT-4DSP stop at  $t = 1.5$  in Figure 4 because MT routing yields for larger values of  $t$  some primary paths that prohibit the existence of a disjoint backup path. Moreover, the difference between the required backup capacity of SPM and OSPF is almost constant, i.e. the absolute capacity savings of about 60% do not depend on the traffic matrix. Hence, the performance of the SPM is very attractive for all traffic matrices in our investigated networks and outperforms clearly other feasible mechanisms.

## 4.2 Impact of Various Resilience Constraints

We investigate the traffic reduction, protection, and bandwidth reuse options for the calculation of the required backup capacities.

Figures 5 and 6 show the required backup capacity for the 5SPM-O protection scheme in the COST-239 and in the Lab03 network. The traffic reduction has no effect if 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 Lab03 network.

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 Lab03 network. 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

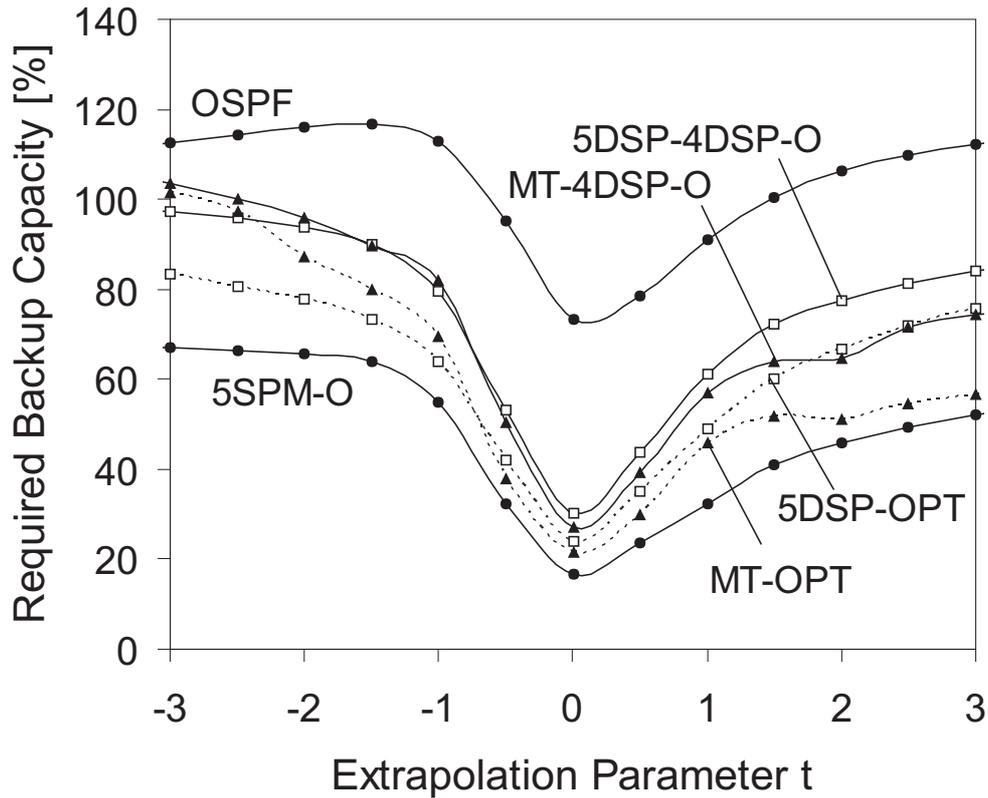


Figure 3: The required backup capacity depending on the traffic matrix for the COST-239 network.

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.

## 5 Conclusion

In this paper we have presented various end-to-end protection switching mechanisms that are simple and easy to implement. In case of a local outage, the source router deviates affected traffic by switching it onto different end-to-end paths such that no additional signaling is required in failure scenarios. Backup capacity sharing among different flows in different failure scenarios allows for considerable capacity savings. Multi-path structures together with load balancing offer degrees of freedom for capacity minimization which is described in [1].

We have investigated the suitability of these backup structures with regard to different traffic matrices, different protected failure scenarios, different traffic reduction models for border router failures, and different capacity reuse options. The results show that the optimized

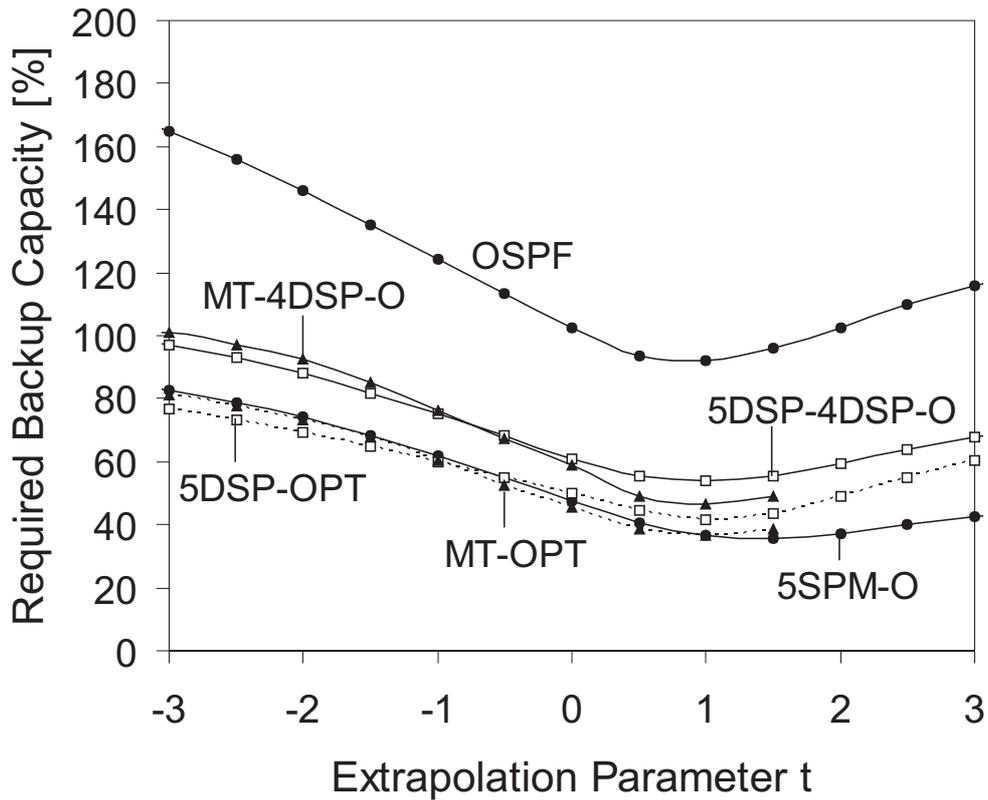


Figure 4: The required backup capacity depending on the traffic matrix for the Lab03 network.

Self-Protecting Multi-Path (SPM) requires the least backup capacity and the difference to other solutions increases with realistic traffic matrices. The required capacity depends on the protected failure scenarios but for large networks, the protection of all router failures is similarly expensive as the protection of all router and link failures. If border routers fail, the traffic load in the network might be partly reduced. The removal of flows with (1) only a failed ingress router, (2) only a failed egress router, (3) and with both ingress and egress routers failed has a significant impact on the traffic in the network but does not influence the required backup capacity. “Bandwidth reuse” on working elements of failed path for backup purposes reduced the required capacity in all experiments by about 5%. The “no bandwidth reuse” option applies, e.g., in optical networks.

In conclusion, the SPM is an attractive backup solution whose benefit against alternative solutions is increased by realistic networking scenarios and it works well under all considered resilience constraints.

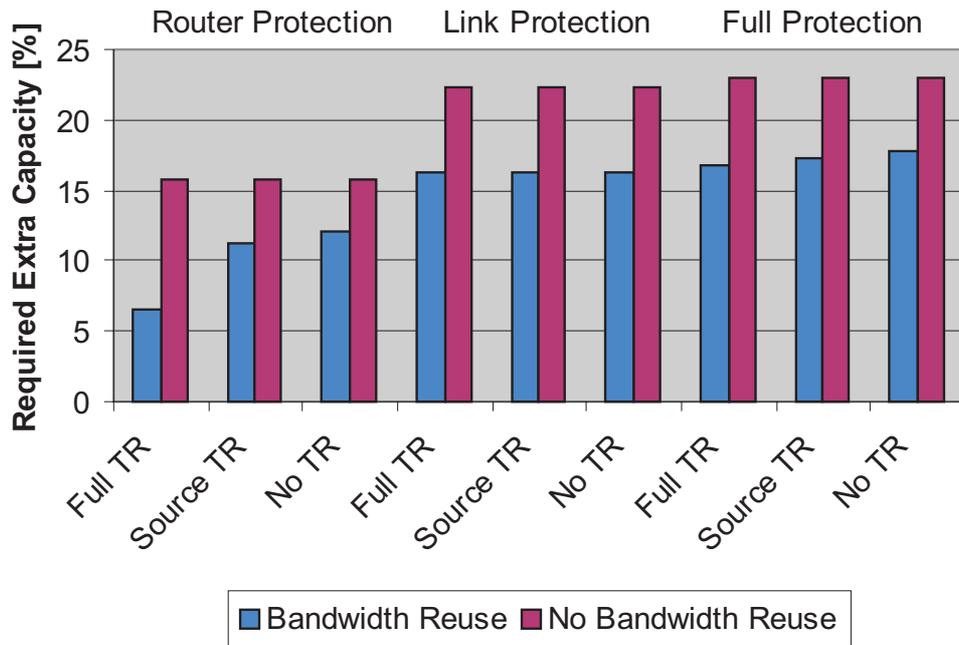


Figure 5: Impact of protected failure scenarios, traffic reduction, and bandwidth reuse in the COST-239 network.

## References

- [1] M. Menth, A. Reifert, and J. Milbrandt, "Optimization of End-to-End Protection Switching Mechanisms for MPLS Networks," Technical Report, No. 320, University of Würzburg, Institute of Computer Science, Feb. 2004.
- [2] M. Menth, A. Reifert, and J. Milbrandt, "Self-Protecting Multipaths - A Simple and Resource-Efficient Protection Switching Mechanism for MPLS Networks," in *3<sup>rd</sup> IFIP-TC6 Networking Conference*, (Athens, Greece), pp. 526 – 537, May 2004.
- [3] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *IEEE INFOCOM'00*, pp. 519–528, 2000.
- [4] S. Köhler and A. Binzenhöfer, "MPLS traffic engineering in OSPF networks - a combined approach," in *18th International Teletraffic Congress (ITC18)*, (Berlin), 9 2003.
- [5] I. Gojmerac, T. Ziegler, F. Ricciato, and P. Reichl, "Adaptive Multipath Routing for Dynamic Traffic Engineering," in *GLOBECOM'03*, (San Francisco), Nov 2003.
- [6] G. Dittmann and A. Herkersdorf, "Network Processor Load Balancing for High-Speed Links," in *SPECTS 2002*, (San Diego, CA), pp. 727–735, 2002.

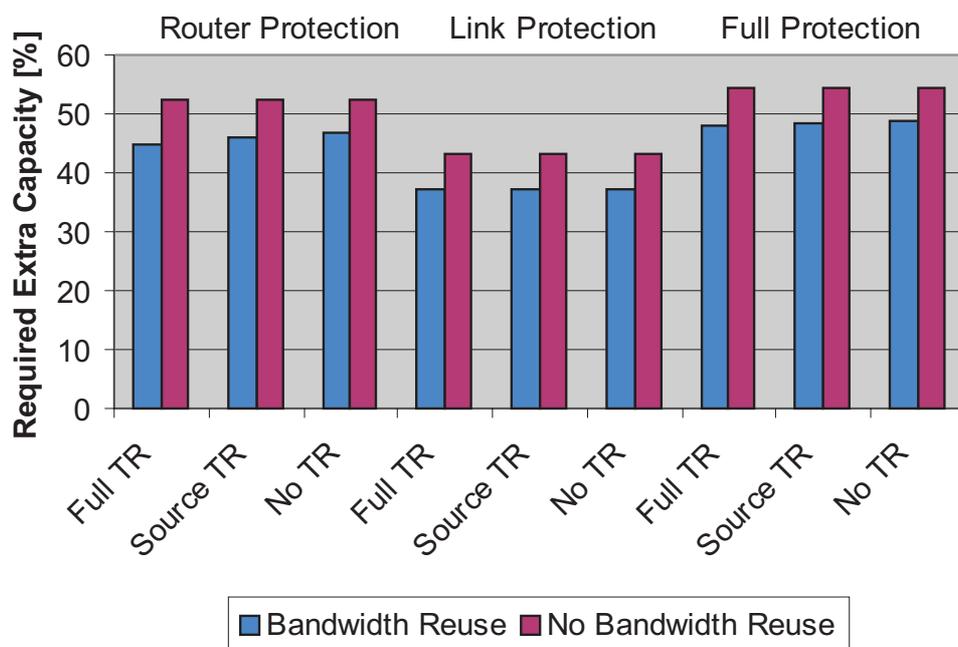


Figure 6: Impact of protected failure scenarios, traffic reduction, and bandwidth reuse in the Lab03 network.

- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.

- [13] 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.
- [14] J. W. Suurballe, “Disjoint Paths in a Network,” *Networks*, vol. 4, pp. 125–145, 1974.
- [15] 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.
- [16] P. Batchelor et al., “Ultra High Capacity Optical Transmission Networks. Final report of Action COST 239.” <http://barolo.ita.hsr.ch/cost239/network/>, 1999.
- [17] C. Hoogendoorn, K. Schrodi, M. Huber, C. Winkler, and J. Charzinski, “Towards Carrier-Grade Next Generation Networks,” in *ICCT 2003*, (Beijing, China), April 2003.