

Resilience Analysis of Packet-Switched Communication Networks

Michael Menth, *Member, IEEE*, Michael Duelli, Ruediger Martin, and Jens Milbrandt

Abstract—In the presence of local network outages, restoration and protection switching mechanisms redirect the traffic over alternative paths to mitigate the effect of failures. However, some failure combinations still lead to loss of ingress-egress connectivity within a network or to severe congestion due to rerouted traffic. Congestion may also be caused by unexpected traffic shifts due to changed user behavior or due to changes of interdomain routing. This paper presents a framework for the analysis of ingress-egress unavailability and link congestion due to (1) failures, (2) changes of user behavior, and (3) changed interdomain routing. It proposes algorithms to find the most probable combinations of (1), (2), and (3) according to some models, and to evaluate the connectivity and the relative link load of the network under these conditions. We have implemented this concept in a software tool and its visualization of the results leads to a comprehensive view of the network’s resilience. It helps to anticipate potential ingress-egress disconnection and congestion before failures and overload occur or before planned modifications (new infrastructure, new routing, new customers) take effect. Thus, it detects weak points in a network, predicts the effectiveness of potential upgrades, and thereby supports careful bandwidth overprovisioning.

Index Terms—Resilience, network reliability and survivability, robustness, capacity (over-)provisioning

I. INTRODUCTION

Internet service providers (ISPs) promise a certain availability and quality of service (QoS) to their customers. The negotiated values are part of service level agreements (SLAs). The ingress-egress unavailability of a network is the expected time fraction during which the network connectivity for a certain ingress-egress pair is lost. The availability is compromised by failures if the border-to-border (b2b) connectivity between two endpoints in the network is lost [1]. QoS in terms of packet loss and delay degrades if congestion occurs in the network. It is mostly caused by additional backup traffic in case of failures [2] (1), overload due to exceptional user behavior (2), or interdomain rerouting [3] (3). Restoration and protection switching mechanisms redirect the traffic over alternative paths in case of failures to mitigate their effect. Capacity overprovisioning addresses the problem of fluctuations of the traffic matrix over time and it can also reduce overload due to redirected traffic [4]. However, some failure combinations and load situations do still lead to b2b unavailability and to severe congestion.

Resilience is the ability of a network to provide a good service also under exceptional conditions. It is an important issue in carrier grade networks and comprises the maintenance of both connectivity and QoS in case of failures

and extraordinary load. The intention and contribution of this paper is the efficient calculation and visualization of a network’s resilience. For our analysis, we define networking scenarios z which have effect on the network’s availability and the utilization of its links. They are characterized with respect to (1), (2), and (3), and have a certain probability. As an exhaustive investigation of all possible scenarios is not feasible, we first provide an efficient algorithm to generate the set of most probable networking scenarios \mathcal{Z} . This is a great difference to most other resilience studies that consider only a certain type of failures, e.g. all single link and/or router failures. Correlated failures, i.e. shared risk groups (SRGs) [5], can be modelled and multiple independent failures are respected automatically. Furthermore, the effect on QoS caused by simultaneous redirected traffic in failure scenarios, fluctuations of the traffic matrix, and extra traffic due to interdomain rerouting can be investigated. We evaluate the b2b connectivity and the utilization of the links for each considered networking scenario $z \in \mathcal{Z}$ and collect these data in a statistic. We propose several graphical representations of these data to give a comprehensive view of a network’s resilience. They are intuitive and help service providers with the definition of appropriate SLAs and network upgrades as they can visualize their impact in advance.

This paper is structured as follows. Section II reviews related work regarding network resilience. Section III explains the framework for resilience analysis. Section IV illustrates the results of the analysis using various visualization approaches, performs sensitivity studies, and motivates the application of the analysis for network upgrades. Finally, Section V summarizes this work and draws conclusions.

II. NETWORK FAILURES AND RESILIENCE

In this section, we review fundamentals about network failures and resilience mechanisms that deviate the traffic around outage locations in the network. We give an overview of related work and clarify our contribution.

A. Network Failures

A good overview and characterization of network failures is given in [6], [7]. We can distinguish planned outages and unplanned failures. Planned outages are intentional, e.g. due to maintenance, and operators can take measures in advance. Unplanned outages are hard to predict and can be further subdivided into failures with internal causes (e.g. software bugs, component defects, etc.) and those with external causes (e.g. digging works, natural disaster, etc.).

Quantitative analyses and statistics about frequency and duration of failure events that occur in operational networks like the Sprint IP backbone are given in [8], [9]. They show that link failures are part of common network operation and the majority of them is short-lived, i.e., their duration is shorter than 10 minutes. Moreover, they indicate that 20% of all failures are due to planned maintenance activities. Almost 30% of the unplanned failures are shared by multiple links and can be attributed to router-related and optical equipment-related problems, while 70% affect only a single link at a time.

The mean time between failures (*MTBF*) and the mean time to repair (*MTTR*) are used to characterize the unavailability of a network element which is $p = \frac{MTTR}{MTBF}$. Different values for *MTBF* and *MTTR* can be found in the literature for nodes and for links [6], [7], [10]–[12]. In this study, we choose $MTTR = 2$ h and $MTBF = 2 \cdot 10^6$ h for nodes, i.e., each node v has the same unavailability $p(v) = 10^{-6}$. The unavailability of a link increases with its length. We assume $MTTR = 12$ h and a mean distance per cable cut and year of $MDCCY = 800$ km to calculate the $MTBF(l) = \frac{MDCCY}{L(l)} \cdot 365 \cdot 24$ h for a link l with length $L(l)$. Thus, a link l with a length of $L(l) = 100$ km has an unavailability of $p = \frac{100 \text{ km} \cdot MTTR}{MDCCY \cdot 365 \cdot 24 \text{ h}} = 1.71 \cdot 10^{-4}$.

B. Resilience Mechanisms

In case of a network failure, resilience mechanisms redirect the affected traffic around the failure location. They can be classified into protection switching and restoration mechanisms. Protection switching establishes backup paths in advance while restoration finds a new path only after a failure occurs. Therefore, protection switching reacts faster than restoration and is usually applied by lower layers. A good overview can be found in [6], [7]. In this study, we use IP rerouting for illustration purposes, but our framework does not depend on any specific routing or resilience mechanism.

IP networks implement destination based routing and calculate the routing tables in a distributed manner according to the shortest path principle. If several shortest paths exist towards a destination, the traffic may be forwarded to a suitable interface with the lowest ID [13, Section 7.2.7]¹, which is single shortest path (SSP) routing, or it may be split equally among all interfaces of the shortest paths, which is called equal-cost multipath (ECMP) routing. If a link or node fails, the routing tables are automatically recalculated and the traffic follows the next shortest paths after some time required for signalling and path calculations [14]. Thus, the b2b IP connectivity is maintained as long as the network is physically connected. In our study, we use ECMP with the standard hop count metric, i.e., all link costs are set to one. Link costs may be manipulated for traffic engineering purposes, e.g., to minimize the link utilization under normal conditions [15] or to make the network robust against link failures [16]–[20], but this is not the focus of this paper.

¹This rule does not hold for OSPF and not all routers running IS-IS implement it.

C. Related Work Regarding Resilience Analysis

The authors of [21] present calculations for the b2b availability of various resilience mechanisms, e.g. dedicated and shared primary and backup path concepts or restoration methods. When rerouting in networks is considered, multiple failures affect the availability which leads to complex calculations. Therefore, either a limited number of the most probable failure scenarios is taken into account [22] or the analysis is limited to single or double failures only. In [23]–[26] the impact of double failures is analyzed in networks that are resilient to single failures. Most papers regarding resilience issues consider only the availability of the network in terms of disconnection probability [27], [28], but some other studies also take the expected lost traffic (ELT) as a performance measure into account to quantify the missing capacity during failures [10], [12], [29]. To reduce the ELT, backup capacity is required that may be used by low priority traffic during failure-free operation of the network [30]. In fact, not all services require the same network reliability [31]. Resilience can also be considered on the application layer, e.g., the availability of services can be improved by alternative servers and caching techniques [32]. NetScope is a tool to calculate the load on the links of a network to predict the effect of various traffic matrices, special failure scenarios, or alternate routing [33]. Our approach can be viewed as a statistical analysis of this idea regarding multiple networking conditions. The authors of [34] consider the completion time of IP reroutes. Within that interval routing loops can occur that lead to temporary loss of connectivity and transient SLA violations. The study provides various statistics based on simulation experiments and shows that networks with similar topological properties can lead to significantly different “goodness factors”. In contrast to our work this study addresses only temporary service disruptions due to routing dynamics but not due to topological disconnection.

D. Contribution of this Work

The framework presented in the next section calculates the b2b availability of a network and the complementary cumulative distribution function (CCDF) of the link load depending on network failures and traffic fluctuations. It is a unification of the methods in prior work [28], [35] where traffic variations have not been considered.

Most approaches in literature are static in the sense that they respect only explicitly specified failures of one or two network elements. However, the probability of multiple network failures grows with increasing network size. We respect these failures if their probability is large enough. Failure combinations with lower probability are neglected in the analysis, but we give bounds on the uncertainty of our results. In addition, our software tool is able to model failures on a finer scope than links and nodes since line cards and router interfaces can also be represented. However, in this study, we model the network only with links and nodes.

Network providers need to know the availability of their networks. Different views on the availability help them to discover different weaknesses. Our tool provides statistics

for the consequences of the network availability on specific b2b aggregates, on the connectivity of a specific node to all other nodes in the network, and on the overall traffic in the network. For most other studies the availability of individual b2b aggregates is the sole result. Our tool is not limited to a specific resilience mechanism which is unlike many other studies. In particular, it handles IP rerouting, the most widely used restoration mechanism, which is more difficult to analyze than simple primary-backup path structures. Regarding potential overload, we provide the CCDF of the load on every link relative to its bandwidth instead of the overall ELT. This information helps network providers to decide whether individual link capacities suffice to provide QoS also in the presence of the most probable failures and changes of the traffic matrix [4]. We propose different performance measures that map the complex information of the CCDFs into a single number. This is useful in practice because it helps to quickly identify the links with the largest risk to be overloaded.

III. RESILIENCE ANALYSIS

Network failures and abnormal traffic matrices lead to unavailability of the network for ingress-egress pairs and to overload on links. The analysis assigns reasonable probabilities to failure scenarios and abnormal traffic matrices, and identifies the most relevant combinations of them in order to analyze them and to derive statistical measures for unavailability and overload in the network.

A. General Notation

A network topology is given by a graph consisting of a set of nodes \mathcal{V} and links \mathcal{E} . We use the Nobel network [36] in Figure 1 with 28 nodes and 41 bidirectional links for illustration purposes in Section IV. The bandwidth of a link $l \in \mathcal{E}$ is denoted by $b(l)$. The network is expected to carry a set of traffic aggregates $\mathcal{G} = \{d = (v, w) : v, w \in \mathcal{V}\}$. The traffic matrix h determines the rate $c_h(g)$ of each aggregate $g \in \mathcal{G}$. A network failure is characterized by its set of failed elements $s \subseteq (\mathcal{V} \cup \mathcal{E})$. Thus, the empty set (\emptyset) stands for the failure-free scenario. We characterize a networking scenario $z = (s, h)$ by its failed elements s and its traffic matrix h . There is a multitude of failure scenarios s and traffic matrices h with a very low probability $p(s)$ and $p(h)$. They lead to an even larger set of combined networking scenarios with even lower probabilities $p(z) = p(s) \cdot p(h)$. Therefore, an exhaustive investigation of all possible networking scenarios is not feasible. Finally, the function $u(g, l, s)$ describes the routing and indicates the fraction of the aggregate g using link l in failure scenario s .

B. Generation of Relevant Networking Scenarios

A networking scenario z needs to be considered only if its probability is sufficiently high, i.e. it meets a certain threshold $p(z) \geq p_{min}$. We identify a set of relevant failure scenarios \mathcal{S} and find then for each $s \in \mathcal{S}$ a set of traffic matrices $\mathcal{H}(s)$ that leads to the set of relevant networking scenarios $\mathcal{Z} = \{(s, h) : s \in \mathcal{S}, h \in \mathcal{H}(s), p(s) \cdot p(h) \geq p_{min}\}$ for our resilience analysis. We first show how the set of relevant failures

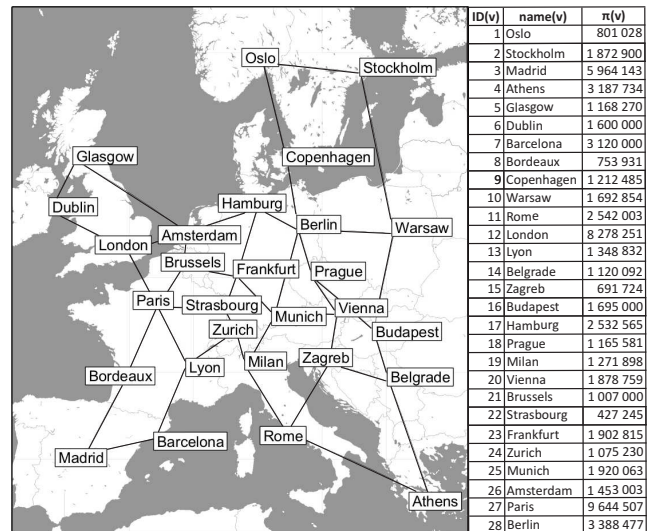


Fig. 1. European Nobel network and the populations of the corresponding cities and surrounding areas.

\mathcal{S} can be obtained efficiently. Then we present exemplary models for traffic matrices h that capture exceptional user behavior and interdomain routing, and use them to generate the networking scenarios \mathcal{Z} .

1) Constructing the Set of Relevant Failure Scenarios \mathcal{S} :

We assume that single link and node failures can occur as independent failure events and denote this set by $\hat{\mathcal{S}}$. The probability $p(\hat{s})$ of a failure event $\hat{s} \in \hat{\mathcal{S}}$ is the unavailability $p(l)$ or $p(v)$ of the corresponding link or node in Section II-A. Our framework can model failure events on a finer scale such as fiber cuts, line card failures, or other internal failures of a router, but we decided to stick with this level of abstraction for our study. The failure events $\hat{s} \in \hat{\mathcal{S}}$ are assumed to be independent of each other, but shared risk groups (SRGs) such as shared risk link or node groups (SRLG, SRNG) [5] can be modelled by virtual elements \hat{s} indicating the simultaneous failure of several resources. This technique has been applied in [37]. We define a compound failure scenario $s \subseteq \hat{\mathcal{S}}$ as a subset of independent failure events $\hat{s} \in \hat{\mathcal{S}}$ that occur simultaneously. Its probability is

$$p(s) = \left(\prod_{\hat{s} \in s} p(\hat{s}) \right) \cdot \left(\prod_{\hat{s} \in \hat{\mathcal{S}} \setminus s} (1 - p(\hat{s})) \right). \quad (1)$$

The set \mathcal{S} contains all (compound) failure scenarios $s \subseteq \hat{\mathcal{S}}$ with probability $p(s) \geq p_{min}$ where p_{min} is the probability threshold for relevant networking scenarios.

In the following we present an algorithm that recursively creates the set of all relevant failure scenarios $s \in \mathcal{S}$. To that end, we number the events \hat{s}_i in an ascending order with regard to their probability $p(\hat{s}_i)$. Algorithm 1 (RFS) constructs the set of relevant failure scenarios \mathcal{S} starting with $\mathcal{S} = \emptyset$. The recursive procedure is invoked with $\text{RFS}(\emptyset, \emptyset, 1)$, i.e. the initial independent failure event to be considered is \hat{s}_0 , the initial partial failure scenario is $s^* = \emptyset$, and its preliminary probability is $p(s^*) = 1$. The algorithm recursively steps through the set of independent failure events $\hat{s}_i \in \hat{\mathcal{S}}$. It constructs a compound failure scenario s^* incrementally and the recursion ends either if the probability $p(s^*)$ of the partial compound failure sce-

nario s^* is lower than p_{min} or if all independent failure events $\hat{s}_i \in \hat{\mathcal{S}}$ have been considered as potential members of s^* . In the latter case, the failure scenario s^* joins \mathcal{S} at the end of each recursion. At program termination, the set \mathcal{S} contains all compound failure scenarios with a probability of at least p_{min} .

```

Input: failure event number  $i$ , partial scenario  $s^*$ ,
          and its probability  $p(s^*)$ 
if ( $p(s^*) \geq p_{min}$ ) then {partial scenario  $s^*$  still probable
enough}
  if ( $i = |\hat{\mathcal{S}}|$ ) then {all independent failure events  $\hat{s}_i$ 
have been considered}
     $\mathcal{S} \leftarrow \mathcal{S} \cup \{s^*\}$ 
  else
    RFS( $i+1, s^* \cup \hat{s}_i, p(s^*) \cdot p(\hat{s}_i)$ )
    RFS( $i+1, s^*, p(s^*) \cdot (1 - p(\hat{s}_i))$ )
  end if
end if

```

Algorithm 1: RFS: constructs the set of relevant failure scenarios \mathcal{S} .

2) *Generation of Traffic Matrices:* Traffic matrices of real networks fluctuate over time in a 24 h and 7 day period. However, when we talk about fluctuations of the traffic matrix, we understand the deviation from the usual maximum of the 7 day period, i.e. the busy hour. We use the simple gravity model [38] to generate traffic matrices for the illustration of our analysis. It can be easily replaced by other, more sophisticated models in the future. Then, we extend it towards overload due to local hot spots and overload due to interdomain rerouting. Finally, we explain how to determine those traffic matrices that need to be considered in our analysis.

a) *Normal Traffic Matrix:* For the illustration in Section IV we use the model from [39] where the rate $c(v, w)$ of a b2b traffic aggregate between node v and w is proportional to the population $\pi(v)$ and $\pi(w)$ in the area of v and w . The population numbers for our test network are given in Figure 1 [40]. Finally, the rate $c(g)$ of a b2b aggregate $g = (v, w)$ can be calculated based on a given overall traffic rate C_{tot} by the following equation:

$$c(g = (v, w)) = \begin{cases} \frac{\pi(v) \cdot \pi(w) \cdot C_{tot}}{\sum_{x, y \in \mathcal{V}, x \neq y} \pi(x) \cdot \pi(y)} & \text{if } v \neq w \\ 0 & \text{if } v = w \end{cases} \quad (2)$$

b) *Changed Traffic Matrix Structure due to Local Hot Spots:* Increased load in networks can occur locally. We capture this by hot spots according to [4]. A hot spot is a node with traffic attraction increased by a factor f_{HS} . We use $f_{HS} = 2$ in our study. We model a hot spot by modifying its population using

$$\pi_{hotspot}^v(w) = \begin{cases} \pi(w) & \text{if } w \neq v \\ f_{HS} \cdot \pi(w) & \text{if } w = v \end{cases} \quad (3)$$

before Equation (2) is applied. Every node is a potential hot spot with a probability of p_{HS} and even several hot spots may occur simultaneously. Therefore, we characterize simultaneous hot spots by the set of routers with increased attractiveness

$\mathcal{V}_{HS} \subseteq \mathcal{V}$. The normal scenario without hot spots is described by $\mathcal{V}_{HS} = \emptyset$. The probability of a specific hot spot scenario is

$$p(\mathcal{V}_{HS}) = (\prod_{v \in \mathcal{V}_{HS}} p_{HS}) \cdot (\prod_{v \notin \mathcal{V}_{HS}} (1 - p_{HS})). \quad (4)$$

This overload model is conservative since it does not increase the overall traffic in the network. It causes a traffic shift and changes the structure of the traffic matrix. As a consequence, an increased or decreased load on the links can be observed.

c) *Increased Traffic Rates due to Interdomain Rerouting:* Due to BGP misconfiguration or other failures, interdomain routing may change, and specific border routers may temporarily receive increased traffic rates. We call this increased load due to interdomain rerouting. It is a rather complex phenomenon [3], but we want to keep things simple to study only fundamental effects in Section IV. We model an interdomain rerouting location v by adding the f_{IR} -fold to the rates of all aggregates starting or terminating in v , which are calculated according to Equation (2) and Equation (3). Thus, in contrast to hot spots, interdomain rerouting increases the traffic rate in the network. Basically, changes of interdomain routing can also reduce the received rate of a border router, but this is not of interest in our study. We assume that a node receives additional traffic from outside its domain with a probability of p_{IR} and with an additive interdomain rerouting factor of $f_{IR} = 1$. This can happen to every border node and also simultaneously to several nodes. This general situation can be denoted by a set $\mathcal{V}_{IR} \subseteq \mathcal{V}$ containing the interdomain rerouting locations, i.e., $\mathcal{V}_{IR} = \emptyset$ describes the normal scenario. The probability of a specific interdomain rerouting event is

$$p(\mathcal{V}_{IR}) = (\prod_{v \in \mathcal{V}_{IR}} p_{IR}) \cdot (\prod_{v \notin \mathcal{V}_{IR}} (1 - p_{IR})). \quad (5)$$

If both endpoints of an aggregate are interdomain rerouting locations, its traffic rate is $(1 + 2 \cdot f_{IR})$ times larger than normal.

d) *Construction of Relevant Traffic Matrices $\mathcal{H}(s)$:* A traffic matrix $h = (\mathcal{V}_{HS}, \mathcal{V}_{IR})$ is characterized by the set of hot spots and the set of nodes that are overloaded due to interdomain rerouting. Its probability is $p(h) = p(\mathcal{V}_{HS}) \cdot p(\mathcal{V}_{IR})$. The set of relevant traffic matrices $\mathcal{H}(s)$ for a relevant failure scenario s comprises all traffic matrices h with $p(s) \cdot p(h) \geq p_{min}$. It can be efficiently computed by an algorithm similar to Algorithm 1. To guarantee that each relevant failure scenario $s \in \mathcal{S}$, i.e. $p(s) \geq p_{min}$, is combined with at least one traffic matrix, we define that the normal traffic matrix $h = (\emptyset, \emptyset)$ is also contained in any set of relevant traffic matrices $\mathcal{H}(s)$. However, its probability is only $p(h = (\emptyset, \emptyset)) = (1 - p_{HS})^{|\mathcal{V}|} \cdot (1 - p_{IR})^{|\mathcal{V}|}$ which is close to 1, but still smaller than 1. As a result, some networking scenarios $z = (s, h)$ with $h \in \mathcal{H}(s)$ can have a probability slightly smaller than p_{min} . Thus, the set of all relevant networking scenarios $\mathcal{Z} = \{z = (s, h) : s \in \mathcal{S}, h \in \mathcal{H}(s)\}$ we consider in our analysis can be slightly larger than expected. However, this has no impact on the correctness of the analysis.

C. Calculation of the Ingress-Egress Unavailability of the Network

The calculation of the exact network unavailability $p_{dis}(g)$ for a b2b aggregate $g = (v, w)$ is in general too complex since

it requires the consideration of all possible failure scenarios. It can be approximated based on the set of relevant failure scenarios \mathcal{S} whose overall probability is $p(\mathcal{S}) = \sum_{s \in \mathcal{S}} p(s)$ by the conditional probability

$$p_{dis}^{\mathcal{S}}(v, w) = \frac{1}{p(\mathcal{S})} \cdot \sum_{s \in \mathcal{S}} p(s) \cdot \text{DISCONNECTED}(v, w, s) \quad (6)$$

which respects only the relevant failure scenarios \mathcal{S} . The function $\text{DISCONNECTED}(v, w, s)$ yields 1 if nodes v and w are disconnected in the presence of failure scenario s ; otherwise it yields 0. The values $p_{dis}^{\mathcal{S}}(v, w)$ are not exact since only the relevant failure scenarios are considered in Equation (6). We get a lower and an upper bound for $p_{dis}(v, w)$ by

$$p_{dis}^{\min}(v, w) = p(\mathcal{S}) \cdot p_{dis}^{\mathcal{S}}(v, w) \text{ and} \quad (7)$$

$$p_{dis}^{\max}(v, w) = p(\mathcal{S}) \cdot p_{dis}^{\mathcal{S}}(v, w) + (1 - p(\mathcal{S})) \cdot 1. \quad (8)$$

The upper bound is exact if v and w are disconnected in all unconsidered failure scenarios. Likewise, the lower bound is exact if they are connected.

D. Calculation of the Relative Link Load

Another objective of our analysis is the assessment of potential overload. To that end, we first calculate the load $c_z(l)$ for every link l for all relevant networking scenarios $z \in \mathcal{Z}$. We consider this load $c_z(l)$ of a link l relative to its bandwidth $b(l)$ and call it the relative link load $\rho_z(l) = \frac{c_z(l)}{b(l)}$. We then compute the complementary cumulative distribution function (CCDF) of the relative link loads which is induced by the probabilities $p(z)$.

Algorithm 2 computes for all $s \in \mathcal{S}$ the routing function $u(g, l, s)$ that determines the fraction of the traffic aggregates g which use link l in the presence of the failure scenario s . Then, it computes for all relevant networking scenarios $z \in \mathcal{Z}$ and for all links $l \in \mathcal{E}$ in the network a load set $\mathcal{L}(l)$ with tuples $(z, c_z(l))$. The load $c_z(l)$ on the link l in networking scenario z is the sum of the traffic contributions from all traffic aggregates $g \in \mathcal{G}$ to that link. Note that this algorithm is fast because the routing function is calculated once for every relevant *failure* scenario $s \in \mathcal{S}$ in the outer loop of Algorithm 2 and not in the inner loop for every relevant *networking* scenario $z \in \mathcal{Z}$.

The load set $\mathcal{L}(l)$ depends on \mathcal{Z} and its information helps to derive the conditional CCDF of the relative link load by

$$P(\rho(l) > r | \mathcal{Z}) = \frac{1}{p(\mathcal{Z})} \cdot \sum_{\{(z, c_z(l)) \in \mathcal{L}(l) : \rho_z(l) > r\}} p(z). \quad (9)$$

Note that $\rho(l)$ can be viewed as link utilization when its value is below 1. The CCDF for the relative link load of Equation (9) is only an approximation because the probability $p(\mathcal{Z}) = \sum_{z \in \mathcal{Z}} p(z)$ of all considered networking scenarios \mathcal{Z} is usually smaller than 1. However, we can give a lower and an upper bound for the unconditioned CCDF $P(\rho(l) > r)$ of $\rho(l)$ by

$$P_{\min}(\rho(l) > r) = P(\rho(l) > r | \mathcal{Z}) \cdot p(\mathcal{Z}) \text{ and} \quad (10)$$

$$P_{\max}(\rho(l) > r) = P(\rho(l) > r | \mathcal{Z}) \cdot p(\mathcal{Z}) + (1 - p(\mathcal{Z})). \quad (11)$$

```

Input: set of relevant failure scenarios  $\mathcal{S}$ 
for all  $s \in \mathcal{S}$  do
   $u(\cdot, \cdot, s) \leftarrow \text{CALCULATEROUTING}(s)$ 
  for all  $h \in \mathcal{H}(s)$  do
     $z \leftarrow (s, h)$ 
    for all  $l \in \mathcal{E}$  do
       $c_z(l) \leftarrow 0$  {initialization}
      for all  $g \in \mathcal{G}$  do
         $c_z(l) \leftarrow c_z(l) + c_h(g) \cdot u(g, l, s)$ 
      end for
       $\mathcal{L}(l) \leftarrow \mathcal{L}(l) \cup (z, c_z(l))$ 
    end for
  end for
end for

```

Algorithm 2: CALCULATELOAD: calculates the load $c_z(l)$ for each link $l \in \mathcal{E}$ for all relevant networking scenarios $z \in \mathcal{Z}$.

IV. ILLUSTRATION OF THE RESILIENCE ANALYSIS

This section illustrates the application of the presented resilience analysis in the Nobel network in Figure 1. Single shortest path routing based on the hop count metric is applied and in case of network failures, the traffic is rerouted. We consider in our analysis independent link and node failures, independent hot spots, and independent extra traffic at border routers. First, we illustrate the analysis of the network availability and then the analysis of the relative link load.

A. Analysis of the Network Availability

In this subsection, we calculate the network unavailability with a minimum probability threshold for relevant failure scenarios of $p_{\min} = 10^{-10}$. This threshold is computationally well feasible and covers a large set of multiple failure scenarios. We choose this very low value to obtain very accurate results for the unavailability and low upper bounds. We present the unavailability from different perspectives and then show how incremental upgrades of the network improve its availability. Finally, we conduct a sensitivity analysis concerning the assumed link and node failure probabilities.

1) *Network Availability from Different Perspectives:* We propose three views on the conditional disconnection probabilities $p_{dis}^{\mathcal{S}}(v, w)$ (cf. Equation (6)) with different aggregation levels. This helps to identify aggregates and points of presence (PoP) with a high unavailability or give an overall impression of the network's unavailability for all aggregates. We consider first the network unavailability in the default network as depicted in Figure 1.

a) *Network Availability for Specific Aggregates:* Figure 2(a) illustrates the network's unavailability for the bidirectional aggregates between router Madrid (ID 3) and any of the 27 other routers in the network. The height of the columns shows the unavailability of the network on a logarithmic y-axis for the aggregates between Madrid and the peer routers indicated by the column numbers. The white columns correspond to the network as presented in Figure 1 while the gray shaded columns correspond to improved network topologies that are discussed later. The aggregates are arranged along the x-axis

in descending order of their unavailability. The column widths are proportional to the traffic volume of the aggregates from the normal traffic matrix such that their relative importance is revealed.

Figure 2(a) shows that only low availability can be guaranteed especially for the aggregates from Madrid to routers 1, 2, 4, 5, 6, and 9 which should be improved or respected in the SLAs.

b) Network Availability for Specific Points of Presence:

We average the network's conditional unavailability for all b2b aggregates of a router by a weighted sum and obtain the conditional average unavailability of the network from the view of a single router by

$$p_{dis}^S(v) = \frac{\sum_{w \in \mathcal{V}, w \neq v} (p_{dis}^S(v, w) \cdot c(v, w) + p_{dis}^S(w, v) \cdot c(w, v))}{\sum_{w \in \mathcal{V}, w \neq v} (c(v, w) + c(w, v))} \quad (12)$$

under the condition that only the relevant failure scenarios \mathcal{S} are respected. Figure 2(b) shows the average unavailability of the network from the perspective of each router. The x-axis indicates the node IDs. For the sake of easier readability, we have arranged the node IDs in Figure 1 according to the descending order of the network's unavailability from their perspective. The figure quickly shows that the network's unavailability from the perspective of routers 1 – 9 is rather large and should be improved or respected in SLAs.

c) Network Availability for the Overall Traffic:

To characterize the unavailability of the network relative to the overall traffic, we calculate the CCDF regarding the unavailability p_{dis}^S . To that end, we consider the unavailability of all b2b aggregates and weight them with their rates:

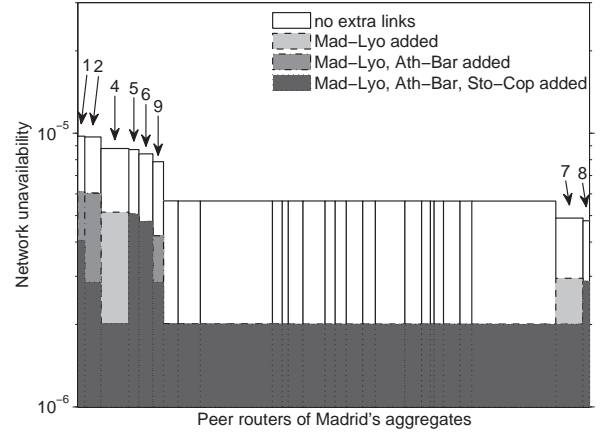
$$P(p_{dis}^S > x) = \frac{\sum_{g \in \mathcal{G}: p_{dis}^S(g) > x} c(g)}{\sum_{g \in \mathcal{G}} c(g)} \quad (13)$$

Figure 2(c) shows that for about 45% of the traffic the network unavailability is about $2 \cdot 10^{-6}$, and for about 55% of the traffic it is larger than this value. The value $2 \cdot 10^{-6}$ is a lower bound for the minimum b2b network unavailability for an aggregate because the network is unavailable if the source or destination router fails and the unavailability of a node is assumed to be 10^{-6} . Figure 2(c) shows that this minimum value can be reached for about 45% of the overall traffic. Figure 2(b) illustrates that the network unavailability averaged over the traffic of a single router is at least $3 \cdot 10^{-6}$ because every router has some aggregates with rather low availability.

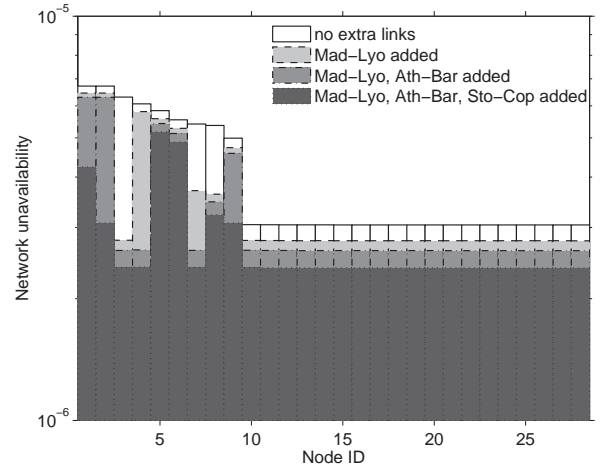
2) Improving the Network Availability: Figure 2(b) suggests that the availability of routers (1) – (9) should be improved. Therefore, we successively add additional links and visualize their impact on the different unavailability reports. We insert links from Madrid to Lyon, from Barcelona to Athens, and from Stockholm to Copenhagen and discuss how they influence the network's availability. We assume that these links do not share common risks with other links.

a) Adding the Link Madrid \leftrightarrow Lyon (3 \leftrightarrow 13):

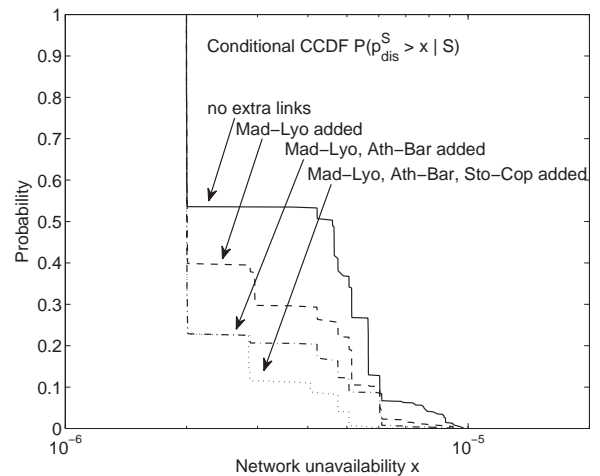
Figure 2(a) shows that the network unavailability for all aggregates starting and ending in Madrid is strongly reduced when the link from Madrid to Lyon is added. The unavailability of most aggregates



(a) B2B network unavailability perceived by all aggregates of PoP Madrid.



(b) Average b2b network unavailability perceived by all routers.



(c) Conditional CCDF of the b2b network unavailability perceived by the overall traffic.

Fig. 2. B2B network unavailability from different perspectives.

approaches even the theoretical value of $2 \cdot 10^{-6}$. Figure 2(b) illustrates that the average unavailability of PoP Madrid (3) decreases from more than $6 \cdot 10^{-6}$ to less than $3 \cdot 10^{-6}$ and those of PoPs Barcelona (7) and Bordeaux (8) decrease from more than $5 \cdot 10^{-6}$ to less than $4 \cdot 10^{-6}$. This is because the new link reduces the number of double link failures that disconnect these cities from most of the other cities. The unavailability for other PoPs is only slightly reduced. The CCDF in Figure 2(c) also shows that the b2b unavailability for a significant amount of traffic is reduced partly to a lower level of $3 \cdot 10^{-6}$ and partly even to the theoretical lower bound of $2 \cdot 10^{-6}$. However, a good portion of the traffic still faces a large network unavailability. Especially the large cities Barcelona and Athens contribute to that effect because they are connected to the network with only two links.

b) *Adding the Link Athens↔Barcelona (4↔7)*: Figure 2(b) illustrates that adding the link between Athens and Barcelona decreases the average unavailability of the corresponding PoPs from about $4 \cdot 10^{-6}$ and $6 \cdot 10^{-6}$ down to a bit more than $2 \cdot 10^{-6}$ while the unavailability for other PoPs is hardly reduced. The CCDF in Figure 2(c) shows that the b2b unavailability for a significant amount of traffic decreases to the theoretical lower bound of $2 \cdot 10^{-6}$ – this is traffic starting or ending in Athens or Barcelona.

c) *Adding the Link Stockholm↔Copenhagen (2↔9)*: Finally, we add a new link between Stockholm and Copenhagen. Figure 2(a) shows that the network availability for Madrid’s aggregates to these two cities is significantly improved, but also the availability for the aggregate from Madrid (3) to Oslo (1) improves notably. Figure 2(b) illustrates that the availability of Stockholm (2), Copenhagen (9), and Oslo (1) is visibly improved, but the unavailability of Stockholm and Copenhagen remains larger than the one of the majority of other routers since simultaneous failures of the links Copenhagen↔Berlin (9↔28) and Stockholm↔Warsaw (2↔10) still disconnect these cities from the rest of the network. Therefore, the improvement by the new link is rather limited. In addition, the Scandinavian cities are rather small such that only a minor fraction of the overall traffic benefits from the new link. This can be nicely observed in Figure 2(c).

Our resilience analysis and the graphical summary reports of the results help to get a quick impression of the network availability for the overall traffic (Figure 2(c)), PoPs with only little network availability can be easily found (Figure 2(b)), and individual aggregates with a large network unavailability can be identified (Figure 2(a)). This knowledge provides suitable availability values for SLAs, it gives hints where to upgrade the network to increase its availability, and supports network planners in strategic decisions. The network availability can also be improved by providing alternative border routers for interdomain traffic. This aspect can be well integrated in our framework by modifying the routing function in such a way that traffic is carried in case of an egress node failure to an alternative egress node.

3) Sensitivity Analysis w.r.t. Unavailability Assumptions:

The above results are based on assumed unavailability values for nodes and links as described in Section II-A, i.e. $p(v) = 10^{-6}$ for nodes, a mean distance per cable cut and

year of $MDCCY = 800$ km and a mean time to repair of $MTTR = 12$ hours. In our sensitivity analysis we consider $p(v) = 10^{-6}$ and $p(v) = 10^{-5}$ and a mean distance per cable cut and year of $MDCCY = 400$ km and $MDCCY = 1600$ km. We perform the experiments for the base network without additional links. Figure 3 shows the impact of different link and node unavailabilities on the network unavailability seen by the overall traffic. The minimum network unavailability is about $2 \cdot p(v)$ and the logarithmic x-axis makes this large impact of the node availability on the network availability of the overall traffic obvious. The link availability also has a significant influence. Thus, a careful assessment of the availability parameters is required before applying this availability analysis for practical purposes.

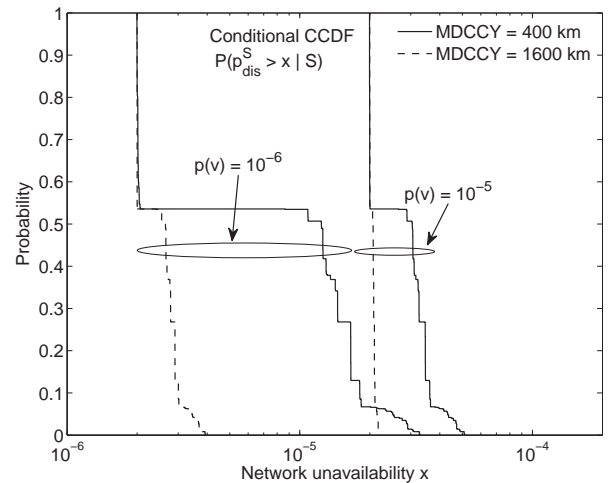


Fig. 3. Conditional CCDF of the b2b network unavailability seen by the overall traffic: sensitivity against assumptions about link and node unavailability.

B. Analysis of Potential Overload

We analyze potential overload in networks which is caused by redirected traffic, traffic hot spots, or by extra traffic due to interdomain rerouting. As “overload” is not well defined, we look at the CCDF of the relative link load. We study the impact of the probability threshold p_{min} that controls the size of the set of relevant networking scenarios \mathcal{Z} in our analysis and the influence of the assumed link and node unavailabilities. Then, we consider the impact of additional hot spots and interdomain rerouting on the relative link load. As the information given by the CCDFs is too complex for practical applications, we finally propose several functions that map the CCDFs to simple numbers characterizing the risk of overload on a link.

For the analysis of the relative link load we dimension the link capacities of our test network such that 20% of their capacity is utilized in case of failure-free operation and the normal traffic matrix. This dimensioning rule disregards available capacity granularities, but we use this setting only for the illustration of our framework and in particular to facilitate the interpretation of the presented results.

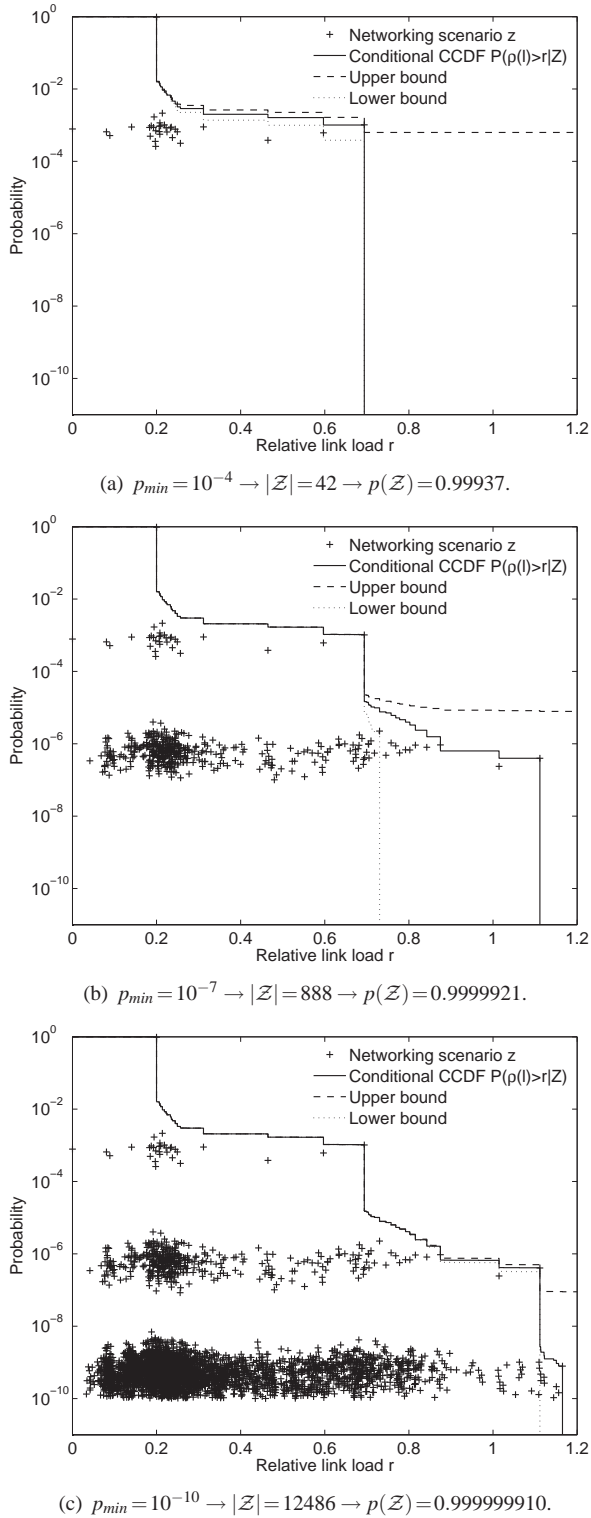


Fig. 4. Conditional CCDF of the relative link load $\rho(l)$ for the link between Brussels and Frankfurt together with a lower and an upper bound for the unconditioned CCDF.

1) *Impact of the Probability Threshold for the Relevant Networking Scenarios:* We consider the resilience of the Nobel network for the normal traffic matrix without hot spots and rerouted interdomain traffic. Figures 4(a)–4(c) show the CCDF of the relative load for the link from Brussels to Frankfurt on a logarithmic scale. Thus, the curves show the probability that the relative link load is larger than a certain value r . The considered networking scenarios $z \in \mathcal{Z}$ are illustrated by crosses ‘+’ and their positions indicate their relative link load and probability $(\rho_z(l), p(z))$. They cause the decay of the CCDF. In our software tool the crosses are interactive such that the respective networking scenarios are displayed when the mouse is dragged over them.

The curve in Figure 4(a) is calculated based on a threshold of $p_{min} = 10^{-4}$ which leads to a set of $|\mathcal{Z}| = 42$ relevant networking scenarios with an overall probability of $p(\mathcal{Z}) = 0.99937$. The solid line is the conditional CCDF of the relative link load based only on the set of relevant networking scenarios \mathcal{Z} . The graph also shows a lower and an upper bound for the unconditioned CCDF. The distance between the curves of these bounds is exactly $1 - p(\mathcal{Z})$, but it looks wider for smaller probability values due to the logarithmic scale of the y-axis. The value $p_{min} = 10^{-4}$ is rather large and leaves a high uncertainty regarding the unconditioned CCDF in the range of interest where the link tends to be overloaded.

We plot the CCDF for $p_{min} = 10^{-7}$ and $p_{min} = 10^{-10}$ in Figures 4(b) and 4(c). Their corresponding sets of relevant networking scenarios are significantly larger with $|\mathcal{Z}| = 888$ and $|\mathcal{Z}| = 12468$ elements such that they cover a probability of $p(\mathcal{Z}) = 0.9999921$ and $p(\mathcal{Z}) = 0.999999910$, respectively. As a consequence, Figure 4(a) contains only one cluster of networking scenarios z which are only single link failures. Figure 4(b) contains two clusters as it also includes double link and single node failures. Figure 4(c) contains even three clusters because also triple link and combined double link and single node failures are respected. Furthermore, the conditional CCDFs have different shapes in the right part of the graph and the distance between the upper and lower bound for the conditioned CCDF becomes smaller with decreasing p_{min} .

2) *Sensitivity Analysis w.r.t. Unavailability Assumptions:* The results presented above are based on assumed unavailability values for nodes and links as described in Section II-A, i.e. $p(v) = 10^{-6}$ for nodes, a mean distance per cable cut and year of $MDCCY = 800$ km with a mean time to repair of $MTTR = 12$ hours. We perform a sensitivity analysis similar to Section IV-A3. We consider $p(v) = 10^{-6}$ and $p(v) = 10^{-4}$ and a mean distance per cable cut and year of $MDCCY = 400$ km and $MDCCY = 1600$ km and take the Nobel network without any additional links for our experiments. Figure 5 shows the impact of different link and node unavailabilities on the overload probabilities of the link between Brussels and Frankfurt in analogy to Figure 4(c). The major difference is observed between the curves with different link availabilities while there is only a small difference between curves with different node availabilities. Curves for $p(v) = 10^{-6}$ and $p(v) = 10^{-5}$ almost completely coincide, therefore, we used $p(v) = 10^{-4}$ instead. This is in contrast to Section IV-A3 which shows that node availabilities have the larger impact on

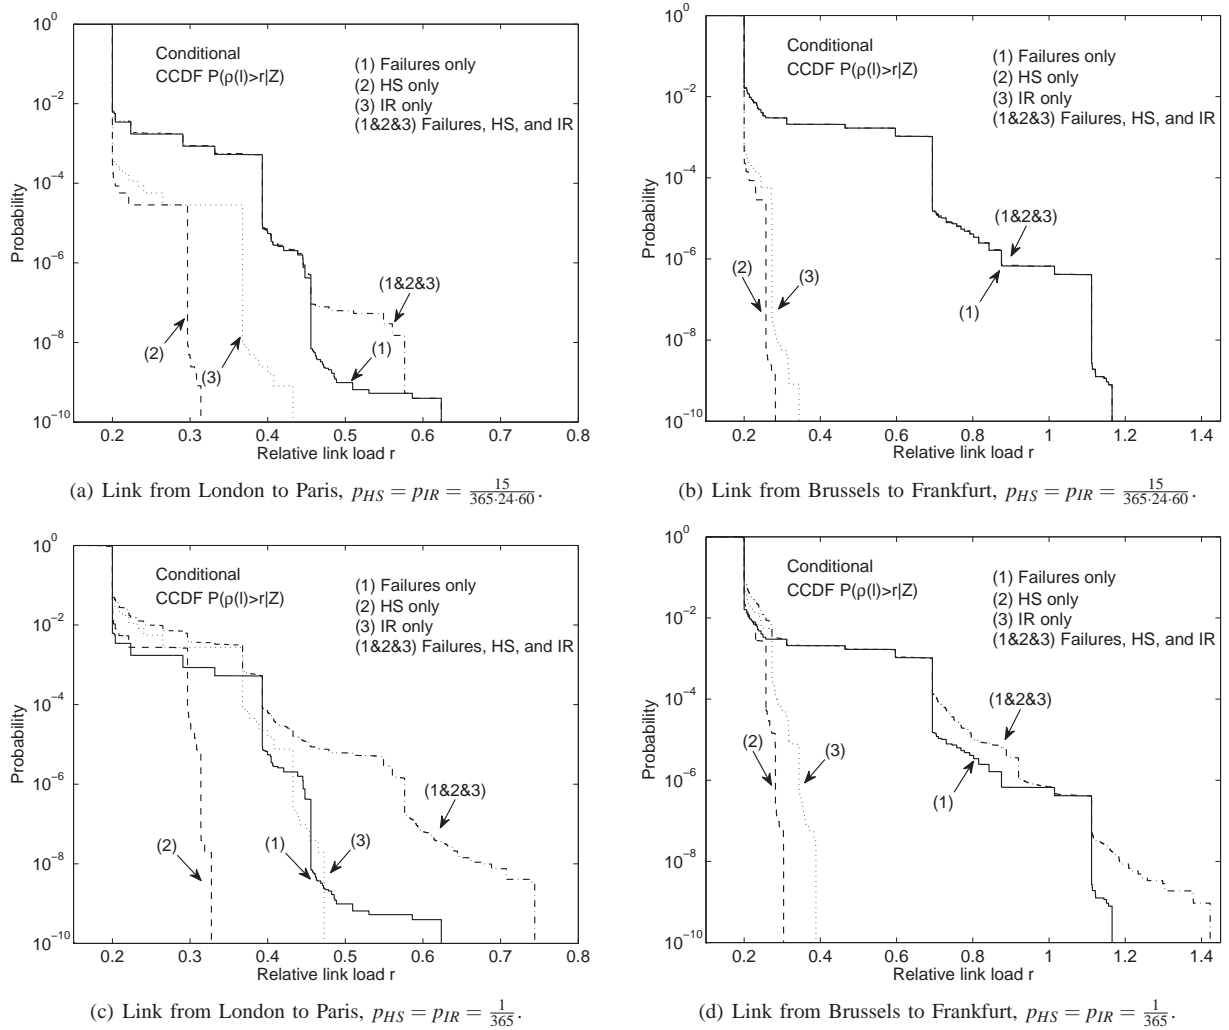


Fig. 6. Impact of failures (1), hot spots (2), interdomain rerouting (3) on the conditional CCDF of the relative link load for $p_{min} = 10^{-10}$.

the network unavailability for the overall traffic. This can be explained as follows. Most (even multiple) link failures do not disconnect the network such that link failure probabilities do not impact the network availability. However, every link failure leads to traffic rerouting, therefore, it is evident that link failure probabilities have an impact on overload probabilities. The impact of node failures on the CCDF is smaller than the one of link failure because node failure probabilities are smaller than link failure probabilities. Hence, for potential overload analysis in practice, it is important to have a good estimate for link failure probabilities while exact values for node failures are secondary as long as they are significantly smaller than link failure probabilities.

3) Impact of Hot Spots (HS) and Interdomain Rerouting (IR): In the following we investigate and compare the impact of hot spots (HS), interdomain rerouting (IR), and failures. To that end, we assume that a node is a HS for 15 minutes per year, i.e. $p_{HS} = \frac{15}{365 \cdot 24 \cdot 60}$, with a multiplicative HS factor of $f_{HS} = 2$. The same holds for IR locations, i.e. $p_{IR} = \frac{15}{365 \cdot 24 \cdot 60}$, and an additive IR factor of $f_{IR} = 1$. Figures 6(a)–6(b) show the conditional CCDFs of the relative load of the links from

London to Paris and from Brussels to Frankfurt for this HS and IR model. They take into account the effect of (1) failures, (2) HS with a multiplicative HS factor of $f_{HS} = 2$, (3) IR with an additive IR factor of $f_{IR} = 1$, and a combination of all three possibilities (1), (2), and (3).

We first look at Figure 6(a) which shows the CCDFs from London to Paris. As the network is dimensioned in such a way that all links are 20% utilized under normal operation, all CCDF curves strongly decay at a relative link load of $r = 0.2$.

The dashed curve (2) shows the impact of only HS. It already decays at relative link loads lower than $r = 0.2$ because some HS divert the traffic from the link and reduce its load. However, such HS scenarios have a very low probability. Therefore, their impact is not visible on the logarithmic y-axis. The relative link load increases mainly in two different probability ranges that correspond to single and double HS. Single HS have a probability of $2.85 \cdot 10^{-5}$ and double HS have a probability of $8.14 \cdot 10^{-10}$. Triple HS have a probability smaller than $p_{min} = 10^{-10}$ and are not considered. A HS factor of $f_{HS} = 2$ at most doubles the rate of an aggregate in case of a single HS and it at most quadruples the rate

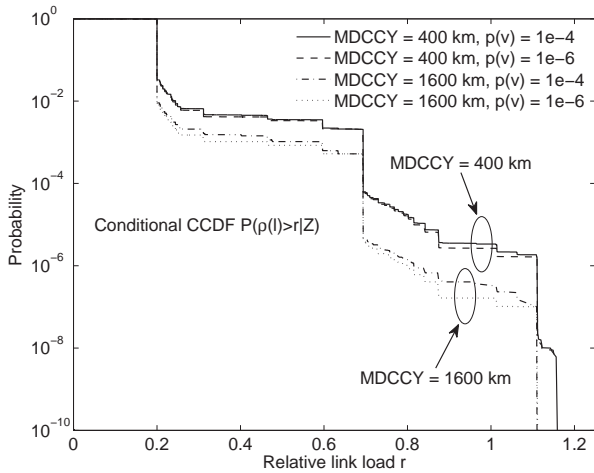


Fig. 5. Conditional CCDF of the relative link load $\rho(l)$ for the link between Brussels and Frankfurt and $p_{min} = 10^{-10}$: sensitivity against assumptions about link and node unavailability.

of an aggregate between two HS. Therefore, the relative link load for networking scenarios with only HS is bounded by the theoretical value 0.8, but the maximum observed relative link load is only 32%. The upper bound is not reached due to the fact that at most one aggregate quadruples its rate in case of a double HS. Although single HS double the rate of several aggregates, the maximum link utilization stays even below 40% because a single link carries multiple aggregates and not all of them have increased rates.

The dotted curve (3) shows potential overload only due to IR. The relative link load increases in the same probability ranges like the dashed curve because the probability model for interdomain overload is the same as for HS. If either v or w is an IR location, the extraordinary traffic matrix h increases the traffic rate $c_h(v, w)$ of an aggregate to the $(1 + f_{IR})$ -fold (i.e. 2) compared to the normal rate $c(v, w)$ or even to the $(1 + 2 \cdot f_{IR})$ -fold (i.e. 3) if both v and w are IR locations. In contrast to HS, the traffic volume of the entire traffic matrix increases. As a consequence, IR causes larger load increases than HS. The maximum observed relative link load is 44% instead of the theoretical upper bound of 60% when both ends of an aggregate are IR locations. The reason for this phenomenon is the same as in the case of HS.

The solid curve (1) shows the potential overload only due to failures. The relative link load increases at higher probabilities compared to the curves for only HS and only IR because failures are more likely in our model than HS and IR. The three main decay steps of the curve correspond to single, double, and triple link failures. Since the solid line is mostly above the dashed and the dotted line, failures are likely to cause stronger load increases than HS and IR. The maximum observed relative link load is about 62%.

The dashed-dotted curve (1&2&3) shows the CCDF of simultaneous failures, HS, and IR. The figure shows that it can be above the curves for failures, i.e., failures and simultaneous HS or IR can cause higher potential overload than just failures. However, this happens only with a very low probability. The maximum observed link utilization due to possibly combined

failures, HS, and IR is also 62%, i.e. the same value as for failures only.

These findings are extremely link-specific. Figure 6(b) shows the corresponding data for the link between Brussels and Frankfurt. The impact of HS is about the same, the impact of IR is weaker, but the impact of failures is significantly larger. Instead of 62% maximum link utilization we observe about 116%. Thus, a common overprovisioning factor for all links is not appropriate, but the results of this analysis can be used for advanced capacity overprovisioning. Some links require 3 times more capacity than under normal operation to be safe against overload, some other links require 6 times more capacity. Our analysis helps to quantify this amount.

As an alternative to adding more capacity on a link, routing optimization may be applied. IP link weights can be set in such a way that the relative loads of all links are as low as possible both in the failure-free scenario and in probable failure cases. This has been studied in [20]. Overload due to HS or IR was not considered.

4) *Sensitivity Analysis w.r.t. Hot Spot (HS) and Interdomain Rerouting (IR) Assumptions:* A prerequisite for application of the overload analysis in practice is an appropriate overload model for HS, IR, and failures. While the impact of different failure probabilities was already shown in Section IV-B2, we illustrate the impact of different HS and IR probabilities in the following. Figures 6(c) and 6(d) present the CCDFs for the links from London to Paris and from Brussels to Frankfurt in analogy to Figures 6(a) and 6(b), but their underlying overload model assumes that a node experiences a HS or a IR event for 24 hours per year each instead of just 15 minutes. The CCDFs for HS (2), IR (3), and failures, HS, and IR (1&2&3) now reflect the impact of single, double, and triple HS or IR locations as their probabilities are $2.54 \cdot 10^{-3}$, $6.99 \cdot 10^{-6}$, and $1.92 \cdot 10^{-8}$.

Comparing Figures 6(a) and 6(c), we see the impact of the modified overload model for HS and IR. The curves for HS and IR are lifted towards higher probability ranges. As triple HS or IR locations are now also considered, we see a third region at small probabilities where the relative link load of the dashed and the dotted curves increase. The impact of triple IR locations is more visible than the impact of triple HS. Figure 6(c) shows that the potential overload caused only by IR can be similarly high or even higher than the potential overload caused only by failures. In addition, simultaneous failures, HS, and IR can cause visibly higher relative link load than failures alone. Moreover, the curve (1&2&3) in Figure 6(c) reveals higher probabilities for large link loads than the one in Figure 6(a) where HS and IR are less likely.

These observations are link-specific. Figures 6(b) and 6(d) show that for some links such as from Brussels to Frankfurt the relative link load due to failures is still a multiple of the relative link load for HS and IR in spite of their larger probability. Also the effect of combined failures, HS, and IR is almost the same as for failures only, and the curve for (1&2&3) is not changed that much due to increased probabilities for HS and IR.

Hence, if HS and IR are rather seldom and not stronger than the multiplicative HS factor $f_{HS} = 2$ or the additive IR factor $f_{IR} = 1$, the relative link load due to failures is for

most links an upper bound for overload due to other reasons. This suggests that it is most important to look at the required backup capacity to carry rerouted traffic in failure cases in order to safely overprovision a network with capacity. This capacity also suffices to accommodate traffic fluctuations due to HS and IR. Similar conclusions were also obtained in [4]. If the probability for HS and IR is larger, this result cannot be generalized. Apart from that, more research regarding overload models is required and empirical evidence is needed.

5) *Comparison of the CCDFs for Different Links:* The conditional CCDF of the relative load $\rho(l)$ of a link l contains the complete information about its potential overload. If the CCDF $P(\rho(l_0) > r | \mathcal{Z})$ of a link l_0 lies for all utilization values below the one of another link l_1 , then the risk of overload for l_0 is clearly smaller than for l_1 . However, Figure 7 shows that this is not a monotone relation. It shows the CCDF of the link utilization for the links from Munich to Vienna and from Athens to Belgrade considering only network failures using $p_{min} = 10^{-10}$. For some utilization values r , the link from Munich to Vienna has a larger CCDF value than the link from Athens to Belgrade (e.g. $r=0.5$) and for some other relative link load values this is vice-versa (e.g. $r=0.75$). Therefore, the CCDFs are difficult to compare. As a consequence, the CCDF is not a suitable means to identify the links with the highest risk to be overloaded in practice. Hence, a function is required to map the information given by the CCDF into a single real number representing the risk of a link to be overloaded. This facilitates a simple comparison of links with regard to their potential overload. In the following section we discuss different mapping functions for that purpose.




Fig. 7. Conditional CCDF of the relative load for the links from Munich to Vienna and from Athens to Belgrade for $p_{min} = 10^{-10}$.

6) *Mapping Functions for Simple Overload Metrics:* One objective of our resilience analysis is to identify links that are most likely to be overloaded. We propose three different functions $R(l)$ mapping the complex information of the CCDF of the relative link load for a link l to real numbers that characterize its potential overload. These values may be used to compare the potential overload of different links and to identify those with the largest risk to be overloaded.

a) *Mapping Functions Based on Overload Probabilities:* The network provider may use an overload threshold r that should not be exceeded by the load of a link. Thus, we define the assessment function for potential overload on link l by $R_r(l) = P(\rho(l) > r | \mathcal{Z})$. Note that this function depends on the value of the overload threshold r . Table I presents the mapping results for $r \in \{0.3, 0.6, 0.9\}$ and shows that the value r indeed influences the ranking order for a few links. This is similar to the phenomenon in Figure 7.

TABLE I
MAPPING FUNCTIONS BASED ON THE OVERLOAD PROBABILITY $R_r(l)$.

link id	$R_r(l)$, $r=0.3$	link id	$R_r(l)$, $r=0.6$	link id	$R_r(l)$, $r=0.9$
Rom-Zag	0.0089	Rom-Zag	0.0045	Osl-Sto	0.0042
Fra-Str	0.0088	Osl-Sto	0.0042	Rom-Zag	0.0031
Osl-Sto	0.0069	Fra-Str	0.0013	Fra-Str	0.0013

b) *Mapping Functions Based on Relative Link Load Percentiles:* The relative link load percentiles $R_q(l) = \text{argmin}(r : P(\rho(l) \leq r | \mathcal{Z}) \geq q)$ help to create a mapping function which depends on the percentile parameter $0 \leq q \leq 1$. Table II shows the mapping results for $q \in \{0.999, 0.99999\}$ and makes the dependency of R_q on the percentile parameter q obvious.

TABLE II
MAPPING FUNCTIONS BASED ON THE RELATIVE LINK LOAD PERCENTILE $R_q(l)$.

link id	$R_q(l)$, $q=0.999$	link id	$R_q(l)$, $q=0.99999$
Bud-War	0.858	Zag-Vie	1.387
Cop-Osl	0.729	Bud-War	1.267
Zag-Vie	0.425	Cop-Osl	0.923

c) *Mapping Functions Based on Weighted Relative Link Loads:* The above overload measures consider only a single point of the conditional CCDF of the relative link load $\rho(l)$, but operators might wish to take the information of the entire CCDF into account. We achieve this by weighting the CCDF with a suitable weight function $w(r)$:

$$R_w(l) = \int_0^{r_{max}} P(\rho(l) > r | \mathcal{Z}) \cdot w(r) dr \quad (14)$$

and we choose $w(r) = 10^{e_{mlwd} \cdot \frac{r}{r_{max}}}$ whereby e_{mlwd} is the maximum logarithmic weight difference which is an arbitrary parameter. This assessment function respects all relative link load values up to r_{max} in the diagram. Thus, the ranking depends on r_{max} and e_{mlwd} . Table III shows the rankings for $r_{max} = 1$ and $e_{mlwd} \in \{2, 4, 6\}$ and makes the influence of the latter parameter explicit.

TABLE III
MAPPING FUNCTIONS BASED ON WEIGHTED RELATIVE LINK LOADS $R_w(l)$.

link id	$R_w(l)$, $e_{mlwd}=2$	link id	$R_w(l)$, $e_{mlwd}=4$	link id	$R_w(l)$, $e_{mlwd}=6$
Mun-Vie	0.065	Mun-Vie	0.094	Fra-Str	0.550
Fra-Str	0.034	Cop-Osl	0.053	Mun-Vie	0.141
Cop-Osl	0.034	Fra-Str	0.052	Cop-Osl	0.126

The three proposed mapping functions define metrics for the risk of overload on a link. The ranking result of a single

mapping function depends on its typical parameter(s). Likewise, different mapping functions also lead to different ranking results. However, all ranking results are similar and show the same trend. Nevertheless, an operator needs to choose the most appropriate mapping function and the corresponding parameter(s) to define overload for his purpose.

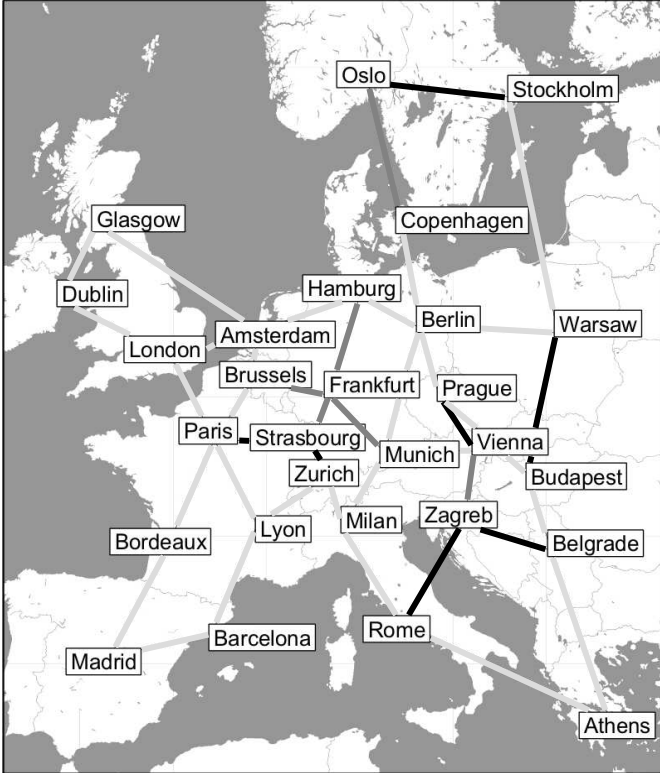


Fig. 8. The colors of the links in the Nobel network indicate their potential overload due to network failures: dark links are more likely to be overloaded; in this example, overload is defined as the probability for relative link load $\rho(l) > 0.6$.

d) Potential Overload at a Glance: The risk of overload in a network can be shown at a glance based on the overload metrics of the mapping functions. Our software tool translates the result of the mapping function $R(l)$ into a color value which is used to display the corresponding links in the topology. Our tool allows to choose any of the above proposed mapping functions and the corresponding parameters r , q , or e_{mlwd} . Changing them does not require any further time-consuming analysis because the stored CCDF of the relative link loads are sufficient to calculate new link colors.

Figure 8 shows an example using the mapping function based on overload probabilities taking only network failures into account. Overload is defined as relative link load larger than $r = 0.6$. For better readability, we have discretized the colors into only 3 values: light gray for $p(\rho(l) > 0.6) < 0.0001$, medium gray for $0.0001 \leq p(\rho(l) > 0.6) < 0.002$, and dark gray for $0.002 \leq p(\rho(l) > 0.6)$. Thus, 7 links have a high overload probability, another 6 links have a medium overload probability, and all other 28 links have only a low overload probability. While changes of the critical relative link load r

can change the overload order of some links (cf. Table I), we obtain at least similar plots for different mapping functions and parameters as long as we look at typical overload values.

V. CONCLUSION

In this paper, we proposed an analysis to assess potential network unavailability and link overload due to exceptional events. In case of network failures, restoration and protection switching mechanisms can reroute traffic from broken paths to backup paths provided that the network is still physically connected; otherwise the network becomes unavailable for some ingress-egress pairs. Traffic redirection increases the load and utilization on the links of the backup paths and may cause overload. This happens similarly in the presence of local traffic shifts within a network (hot spots) or in case of additional transit traffic due to interdomain rerouting. The contribution of this paper is twofold.

The first contribution is an algorithmic framework for the assessment of network unavailability and overload. The idea of our analysis is to derive statistical results for network unavailability and link overload from a probabilistic description of (1) network failures, (2) local hot spots, and (3) interdomain rerouting. The presented analysis is very general as it copes with different routing and resilience mechanisms and arbitrary shared risk groups. Our approach requires the analysis of different networking scenarios $z = (s, h)$ consisting of failure scenarios s and traffic matrices h . With a certain probability, each link and node of a network can fail, different hot spots can occur, and extra traffic can enter the network at any border router. However, an exhaustive analysis of all possible networking scenarios is prohibitive due to limited calculation time. We solved this problem mainly by two approaches. First, only networking scenarios with a minimum probability of p_{min} are respected in the analysis of potential overload such that computation speed can be traded for accuracy, and upper and lower bounds are given for the approximated results. Second, we designed efficient algorithms that reuse intermediate results in different computations.

The second contribution is the graphical summary of the vast amount of statistical data in order to make them comprehensible which is necessary for resilience analysis in practice. To that end, we illustrated the application of the analysis for the topology of the European Nobel network. We proposed different summary reports for network unavailability that easily show the impact of topology changes on the network resilience. Potential overload of a link is presented by the complementary cumulative distribution function (CCDF) of its relative load. Sensitivity studies showed that the results for the unavailability analysis significantly depend on estimates for link and node unavailabilities while the results for the overload analysis mainly depend on estimates for link unavailabilities. Failures have probably a larger impact on the expected overload than increased traffic rates due to hot spots or interdomain rerouting, but more research on overload models is required and empirical data are needed. The CCDF of the relative link load carries the full information about potential link overload. However, CCDFs are difficult to understand and compare, so a

single number indicating the risk of a link to be overloaded is desired. We proposed different mapping functions that serve as definitions for the risk of a link to be overloaded. The individual methods and their parameterizations have an impact on the exact overload order of links. The result of these mapping functions allows to draw network maps where the color of the links characterizes their potential overload which can be easily interpreted by network engineers.

After all, the proposed resilience analysis and the visualization of its results can assist network and service providers with the operation of their networks. Our software tool gives hints for appropriate availability values which are useful for the definition of feasible SLAs. It detects underprovisioned links before overload occurs, it supports economic capacity overprovisioning, and it predicts the impact of potential infrastructure changes or upgrades on the resilience of the network.

ACKNOWLEDGEMENTS

The authors would like to thank Matthias Hartmann, David Hock, Florian Hoehn, Jan Junker, Matthias Koller, Frank Lehrieder, Christian Schwartz, and David Stezenbach for their programming efforts.

REFERENCES

- [1] S. Bhattacharyya, C. Diot, G. Iannaccone, A. Markopoulou, and C. Chuah, "Service Availability in IP Networks," Sprint, ATL Research Report RR03-ATL-071888, July 2003.
- [2] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, "An Approach to Alleviate Link Overload as Observed on an IP Backbone," in *IEEE Infocom*, San Francisco, CA, April 2003.
- [3] T. Schwabe and C. G. Gruber, "Traffic Variations Caused by Inter-domain Re-routing," in *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Ischia Island, Italy, Oct. 2005.
- [4] M. Menth, R. Martin, and J. Charzinski, "Capacity Overprovisioning for Networks with Resilience Requirements," in *ACM SIGCOMM*, Pisa, Italy, Sep. 2006.
- [5] L. Shen, X. Yang, and B. Ramamurthy, "Shared Risk Link Group (SRLG)-Diverse Path Provisioning under Hybrid Service Level Agreements in Wavelength-Routed Optical Mesh Networks," *IEEE/ACM Transactions on Networking*, vol. 13, no. 4, pp. 918–931, Aug. 2005.
- [6] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery*, 1st ed. Morgan Kaufmann / Elsevier, 2004.
- [7] B. Mukherjee, *Optical WDM Networks*, 2nd ed. Springer, 2006.
- [8] G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot, "Analysis of Link Failures in an IP Backbone," in *ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, Nov. 2002, pp. 237–242.
- [9] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, and C.-N. Chuah, "Characterization of Failures in an IP Backbone," in *IEEE Infocom*, Hongkong, Mar. 2004.
- [10] G. Willems, P. Arijis, W. V. Parys, and P. Demeester, "Capacity vs. Availability Trade-offs in Mesh-Restorable WDM Networks," in *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Budapest, Hungary, Oct. 2001.
- [11] H. C. Cankaya, A. Lardies, and G. W. Ester, "A Methodology for Availability-Aware Cost Modelling of Long-Haul Networks," in *International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*, San Jose, CA, Jul. 2004.
- [12] S. D. Maeschalck, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jaeger, R. Inkret, B. Mikac, and J. Derkacz, "Pan-European Optical Transport Networks: an Availability-Based Comparison," *Photonic Network Communications*, vol. 5, no. 3, pp. 203–225, 2005.
- [13] D. Oran, "RFC1142: OSI IS-IS Intra-Domain Routing Protocol," Feb. 1990.
- [14] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot, "Feasibility of IP Restoration in a Tier-1 Backbone," *IEEE Network Magazine (Special Issue on Protection, Restoration and Disaster Recovery)*, March 2004.
- [15] B. Fortz, J. Rexford, and M. Thorup, "Traffic Engineering with Traditional IP Routing Protocols," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 118–124, 2002.
- [16] B. Fortz and M. Thorup, "Robust Optimization of OSPF/IS-IS Weights," in *International Network Optimization Conference (INOC)*, Paris, France, Oct. 2003, pp. 225–230.
- [17] A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, and C. Diot, "IGP Link Weight Assignment for Transient Link Failures," in *18th International Teletraffic Congress (ITC)*, Berlin, Sep. 2003.
- [18] D. Yuan, "A Bi-Criteria Optimization Approach for Robust OSPF Routing," in *3rd IEEE Workshop on IP Operations and Management (IPOM)*, Kansas City, MO, Oct. 2003, pp. 91–98.
- [19] A. Sridharan and R. Guerin, "Making IGP Routing Robust to Link Failures," in *IFIP-TC6 Networking Conference (Networking)*, Ontario, Canada, May 2005.
- [20] M. Menth, M. Hartmann, and R. Martin, "Robust IP Link Costs for Multilayer Resilience," in *IFIP-TC6 Networking Conference (Networking)*, Atlanta, GA, USA, May 2007.
- [21] P. Cholda and A. Jajszczyk, "Availability Assessment of Resilient Networks," in *12th GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems (MMB) together with 3rd Polish-German Teletraffic Symposium (PGTS)*, Dresden, Germany, Sep. 2004, pp. 389–398.
- [22] V. O. K. Li and J. A. Silvester, "Performance Analysis of Networks with Unreliable Components," *IEEE Transactions on Communications*, vol. 32, no. 10, pp. 1105–1110, Oct. 1984.
- [23] M. Clouqueur and W. D. Grover, "Computational and Design Studies on the Unavailability of Mesh-restorable Networks," in *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Munich, Germany, Apr. 2000, pp. 181–186.
- [24] —, "Availability Analysis of Span-Restorable Mesh Networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, pp. 810–821, 2002.
- [25] D. A. Schupke and R. G. Prinz, "Capacity Efficiency and Restorability of Path Protection and Rerouting in WDM Networks Subject to Dual Failures," *Photonic Network Communications*, vol. 8, no. 2, Sep. 2004.
- [26] M. Menth, R. Martin, and U. Spoerlein, "Impact of Unprotected Multi-Failures in Resilient SPM Networks: a Capacity Dimensioning Approach," in *IEEE Globecom*, San Francisco, California, USA, Nov. 2006.
- [27] W. Najjar and J.-L. Gaudiot, "Network Resilience: A Measure of Network Fault Tolerance," *IEEE Transactions on Computers*, vol. 39, no. 2, pp. 174–181, Feb. 1990.
- [28] J. Milbrandt, R. Martin, M. Menth, and F. Hoehn, "Risk Assessment of End-to-End Disconnection in IP Networks due to Network Failures," in *6th IEEE Workshop on IP Operations and Management (IPOM)*, Dublin, Ireland, Oct. 2006, pp. 181–192.
- [29] P. Cholda and A. Jajszczyk, "Reliability Assessment of Optical p-Cycle," *IEEE/ACM Transactions on Networking*, vol. 15, no. 6, Dec. 2007.
- [30] M. Durvy, C. Diot, N. Taft, and P. Thiran, "Network Availability Based Service Differentiation," in *11th IEEE International Workshop on Quality of Service (IWQoS)*, Berkeley, CA, USA, 2003, pp. 305–324.
- [31] P. Cholda, A. Mykkeltveit, B. E. Helvik, O. J. Wittner, and A. Jajszczyk, "A Survey of Resilience Differentiation Frameworks in Communication Networks," *IEEE Communications Surveys & Tutorials*, vol. 9, no. 4, 2007.
- [32] M. Dahlin, B. B. V. Chandra, L. Gao, and A. Nayate, "End-to-End WAN Service Availability," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, pp. 300–313, April 2003.
- [33] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "NetScope: Traffic engineering for IP Networks," *IEEE Network Magazine*, pp. 11–19, Mar. 2000.
- [34] R. Keralapura, A. Moerschell, C.-N. Chuah, G. Iannaccone, and S. Bhattacharyya, "A Case for Using Service Availability to Characterize IP Backbone Topologies," *Journal of Communications and Networks*, vol. 8, no. 2, Jun. 2006.
- [35] J. Milbrandt, M. Menth, and F. Lehrieder, "A Priori Detection of Link Overload due to Network Failures," in *ITG/GI Conference on Communication in Distributed Systems (KiVS)*, Bern, Switzerland, Feb. 2007.
- [36] "SNDlib 1.0 – Survivable Network Design Data Library," <http://sndlib.zib.de>, 2005.
- [37] M. Duelli, M. Hartmann, M. Menth, R. Hülsermann, and M. Düser, "Performance Evaluation of IP over Cost-Optimized Optical Multilayer Networks with SRLGs," in *ITG Workshop on Photonic Networks*, Leipzig, Germany, Apr. 2008.

- [38] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic Matrix Estimation: Existing Techniques and New Directions," in *ACM SIGCOMM*, Pittsburgh, USA, Aug. 2002.
- [39] M. Menth, "Efficient Admission Control and Routing in Resilient Communication Networks," PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland, July 2004.
- [40] T. Brinkhoff, "Population of the Major Cities and Agglomerations for Each Country," <http://www.citypopulation.de/>, 1998-2006.



Michael Menth studied computer science and mathematics at the University of Wuerzburg/Germany and Austin/Texas. He worked at the University of Ulm/Germany and Wuerzburg and obtained his PhD in 2004. Currently, he is assistant professor and heading the research group "Next Generation Networks" at the Institute of Computer Science in Wuerzburg. His special interests are performance analysis, optimization of communication networks, resource management, resilience issues, and Future Internet. Dr. Menth holds numerous patent applications and received various scientific awards for innovative work.



Michael Duelli studied computer science and mathematics at the University of Wuerzburg/Germany. He received his diploma degree in computer science in 2007. Since then he is a researcher at the Institute of Computer Science in Wuerzburg and pursuing his PhD. His current research focuses on optimizing optical transport networks – especially Carrier Ethernet – in combination with performance evaluation and resilience analysis.



Ruediger Martin studied computer science and mathematics at the University of Wuerzburg/Germany and SUNY Albany/New York. He worked at the University of Wuerzburg as a research assistant with focus on load balancing mechanisms, load models for capacity overprovisioning, and analysis of resilience mechanisms in communication networks. He received several IEEE best paper awards for his work. After receiving his PhD in 2008, he joined the European Commission in Brussels as a scientific officer in the area of Future and Emerging Technologies (FET) which fosters novel non-conventional approaches and foundational research.



Jens Milbrandt studied computer science and economics in Wuerzburg/Germany and received his diploma degree in 2001. He finished his PhD in computer science on performance evaluation and resource management in next generation networks at the University of Wuerzburg in 2006. Since then he works as a scientist and network engineer for the Bell Labs Germany as part of Alcatel-Lucent in Stuttgart. His interests are now focused on the evaluation and analysis of new network technologies for future packet transport infrastructure.