A Priori Detection of Link Overload due to Network Failures

Jens Milbrandt, Michael Menth, and Frank Lehrieder

Department of Distributed Systems, Institute of Computer Science
University of Würzburg, Am Hubland, D-97074 Würzburg, Germany
Phone: (+49) 931-888 6644, Fax: (+49) 931-888 6632
{milbrandt,menth,lehrieder}@informatik.uni-wuerzburg.de

Abstract. Restoration or protection switching mechanisms are triggered by link or node failures to redirect traffic over backup paths. These paths then carry the normal and the backup traffic which may lead to overload and thereby to quality of service (QoS) violations, i.e. to excessive packet loss and delay. In this paper, we present a method to assess the potential overload of the links due to network failures. We calculate the complementary cumulative distribution function (CCDF) of the relative load for each link in the network. We discuss various performance measures that condense this information to a single value per link which is suitable for a link ranking. This helps to identify weak spots of the network and to appropriately upgrade the bandwidth of links although they are not overloaded during normal operation. We implemented the concept in a software tool which helps network providers to anticipate the potential overload in their networks prior to failures and intended modifications (new infrastructure, new routing, new customers, ...) and to take appropriate actions.

1 Introduction

Network resilience is an important issue in carrier grade networks. It comprises the maintenance of both the connectivity and the quality of service (QoS) in terms of packet loss and delay during network failures. To maintain connectivity, restoration or protection switching mechanisms are triggered by link or node failures and redirect traffic over backup paths. These paths then carry the normal and the backup traffic which may lead to overload and thereby to QoS violations. While the service availability in terms of connectivity has been studied quite well [1], the a priori detection of overload due to network failures has not yet been investigated in depth.

In this paper, we present a concept to calculate the risk of overload due to network failures and implement it in a software tool. Basically, failures are associated with certain probabilities and the triggered traffic redirection causes a specific relative link load on all links. Considering all possible network failures, the law of total probability allows the derivation of the complementary cumulative distribution function (CCDF) of

This work was funded by the Bavarian Ministry of Economic Affairs and the German Research Foundation (DFG). The authors alone are responsible for the content of the paper.

the relative load of all links. However, this simple principle comes with several problems that are solved in this paper.

Firstly, the evaluation of all possible failure scenarios is computationally not feasible for medium size or large networks. Therefore, the analysis must limit its scope to the most relevant failure scenarios. In most previous studies, only the impact of single failures has been considered as the probability of multiple failures is rather low. However, the number of multiple failures increases strongly with the network size such that their overall probability cannot be neglected for the approximation of the CCDF. We define all (multi-)failures with a probability larger than p_{min} as relevant – irrespective of their number of failed components – and consider them for our analysis. We present an algorithm to efficiently find the set of all relevant failure scenarios. The framework is designed in such a way that both independent and correlated multiple failures can be respected. The latter ones are better known as shared risk resource groups (SRRGs) [2]. For instance, shared risk link groups (SRLGs) consist of IP links that are logically distinct on the network layer but share a common resource on the link layer; the failures of this common resource makes all links of the SRLG inoperative.

Secondly, the CCDF of the relative link load is the most detailed information we obtain from the analysis, but it is not suitable to compare the risk of link overload of several links because the CCDF does not define a strict order relation among them. To facilitate link rankings, we discuss various performance measures that condense the information of the CCDF into a single value. These condensed values help operators to quickly identify the most critical links and to find weak spots in their networks. It allows them to prevent congestion due to redirected traffic by upgrading the bandwidth of the most jeopardized links appropriately although they are not overloaded during normal operation.

This paper is structured as follows. In Section 2 we review related work regarding network resilience. Section 3 explains our algorithms to calculate the CCDFs of the relative link loads that are caused by the most relevant failures. Section 4 visualizes these CCDFs and develops simple assessment functions to condense their information into a single value; furthermore, additional features of our software tool are illustrated to provide a quick overview of the potential overload in the entire network. Finally, Section 5 summarizes this work and gives an outlook on possible extensions.

2 Network Failures and Resilience

In this section, we review basics about network failures and resilience mechanisms that deviate the traffic around an outage location in the network. We give an overview on similar work and comment on our contribution.

2.1 Network Failures

A good overview and characterization of network failures is given in [3, 4]. We can distinguish planned and unplanned failures. Planned outages are intentional, e.g. due to maintenance, and operators can take countermeasures 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 disasters, 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 [5, 6]. They show that link failures are part of everyday's 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. Of the unplanned failures, almost 30% are shared by multiple links and are related to problems with routers or optical equipment, 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 by $p = \frac{\text{MTTR}}{\text{MTBF}}$. Different values for MTBF and MTTR can be found in the literature for nodes and for links [3, 4, 7–9]. In this study, we choose MTTR = 2 h and MTBF = $2 \cdot 10^6$ h for nodes, i.e. $p_{node} = 10^{-6}$. Furthermore, we use MTTR = 12 h and MTBF = $\frac{300 \text{ km}}{L(l)} \cdot 365 \cdot 24$ h for links where $L(l)$ denotes the length of the link l such that we get $p_{link} = L(l)/219000$ km.

2.2 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. Protection switching establishes backup paths in advance while restoration finds a new path after a failure has occurred. Therefore, protection switching reacts faster than restoration. A good overview of resilience mechanisms can be found in [3, 4]. 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 a link or node fails, the routing tables are automatically recalculated and the traffic follows the next shortest paths after some time [10]. Thus, e2e IP connectivity is maintained as long as the network is physically connected. If several shortest paths exist towards a destination, the traffic may be forwarded to the interface with the highest priority, 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. In our study, we use ECMP with the standard hop count metric, i.e., all link costs are set to one. However, the link costs may be manipulated for traffic engineering purposes, e.g., to minimize the link utilization under normal conditions [11] or to make the network robust against link failures [12–15].

2.3 Related Work Regarding Resilience Analysis

The authors of [16] present calculations for e2e availability of various resilience mechanisms, e.g. dedicated and shared primary and backup path concepts or restoration methods. When rerouting in networks is considered, many multiple failures affect the availability which leads to complex calculations. Therefore, either a limited number of most probable failure scenarios is taken into account [17] or the analysis is limited to only single or double failures. In [18–21] the impact of double failures is analyzed in networks that are resilient to single failures. Most papers regarding resilience issues consider only e2e availability [1], but some other studies also take the expected lost traffic (ELT) as a performance measure into account to quantify the missing capacity

during failures [7, 9]. To reduce ELT, backup capacity is required that may be used by low priority traffic during failure-free operation of the network [22]. Resilience can also be considered on the application layer, e.g., the availability of services can be improved by alternative servers and caching techniques [23]. 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, and can be used for the inner loop of routing optimization [24]. Our tool is basically an extension of that approach towards statistical results.

2.4 Contribution of this Work

The above mentioned studies are static in the sense that they respect only explicitly specified failures of (single) network elements. This is a reasonable start for resilient QoS provisioning, but the probability of multiple network failures grows with increasing network size. Therefore, multiple failures need to be taken into account if the network size increases. The objective is to make networks resilient to the majority of likely failure scenarios in the sense that no overload occurs due to redirected traffic. Hence, the impact of the majority of likely failure scenarios is more important for the network resilience than the impact of a few devastating but very unlikely multiple failures.

The novelty of this work is the assessment of the potential overload in a network. We present a framework that yields a distribution of the link load caused by redirected traffic in failures scenarios. This helps Internet service providers (ISPs) (1) to detect weak spots in their network and (2) to improve the resilience of their network systematically without general overprovisioning [25]. The improvement can be achieved (a) by improved routing and rerouting in failure cases, (b) by the upgrade of existing links, or (c) by the introduction of new infrastructure. We currently develop a tool that predicts the resilience of the network after such modifications to support the ISP with his decision process.

3 Calculation of the Relative Link Load

We assess the potential overload of the links by analyzing the impact of failure scenarios on their relative link load. As not all failure scenarios can be covered by the analysis due to computational complexity, we determine the most relevant ones and take only them into account. Various failure scenarios lead to the same so-called “effective” working topology. We take advantage of that fact for the calculation of the traffic rates on the links. They are needed to derive the link-specific conditional complementary cumulative distribution function (CCDF) of the relative link load.

3.1 Relevant Failure Scenarios

In our study, we consider a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consisting of routers $v \in \mathcal{V}$ and directed links $l \in \mathcal{E}$. A simple application of our framework assumes link and node failures as basic and independent failure events. However, it is possible to model failure events on an even lower level, e.g., individual interfaces and line cards. The set $\hat{\mathcal{S}}$ represents the set of all independent failure events. Note that this set may also contain shared risk resource groups (SRRGs) such as shared risk link or node groups (SRLG, SRNG) [2].

Each of these failure events occurs with probability $p(\hat{s})$ and we number the events $\hat{s}_i$ $(0 \leq i < |\hat{\mathcal{S}}|)$[1] in a descending order according to $p(\hat{s}_i)$. A (compound) failure scenario $s \subseteq \hat{S}$ consists of several independent failure events $\hat{s} \in \hat{\mathcal{S}}$ that incidentally occur at the same time. Its probability is $p(s) = (\Pi_{\hat{s} \in s} p(\hat{s})) \cdot (\Pi_{\hat{s} \in \hat{\mathcal{S}} \setminus s}(1 - p(\hat{s})))$ and it is relevant if it has a probability of at least $p(s) \geq p_{min}$. Finally, the set of relevant failure scenarios $\mathcal{S} = \{s \in \mathcal{P}^{\hat{\mathcal{S}}} : p(s) \geq p_{min}\}$ comprises all relevant failure scenarios and is a subset of the power set $\mathcal{P}^{\hat{\mathcal{S}}}$. In particular, the failure-free scenario $s = \emptyset$ is relevant and part of $\mathcal{S}$.

Algorithm 1 constructs $\mathcal{S}$. At the beginning, the global set of relevant failure scenarios is initialized with $\mathcal{S} = \emptyset$. The recursive procedure RELEVANTSCENARIOS$(i, s, p(s))$ is called with arguments $(0, \emptyset, 1)$. The algorithm steps recursively 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 scenario 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 $(i = |\hat{\mathcal{S}}|)$ **then** {all independent failure scenarios have been considered}
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$
  else {partial scenario s is probable enough to be relevant}
    if $(p(s) \cdot p(\hat{s}_i) > p_{min})$ then RELEVANTSCENARIOS$(i+1, s \cup \hat{s}_i, p(s) \cdot p(\hat{s}_i))$
    if $(p(s) \cdot (1 - p(\hat{s}_i)) > p_{min})$ then RELEVANTSCENARIOS$(i+1, s, p(s) \cdot (1 - p(\hat{s}_i)))$
  end if

Algorithm 1: RELEVANTSCENARIOS: constructs the set of relevant scenarios $\mathcal{S}$.

3.2 Effective Topologies

The effective topology $T(s)$ caused by a compound failure scenario s is characterized by its set of working links and nodes. A link works if and only if itself and its adjacent routers do not fail. A router works if and only if itself and at least one of its adjacent links do not fail. Thus, all scenarios containing the failure of a router and some of its adjacent links lead to the same effective topology T. We subsume all of these scenarios in the set $\mathcal{S}(T)$ and the probability of T is inherited by $p(T) = \sum_{s \in \mathcal{S}(T)} p(s)$. The set $\mathcal{T} = \bigcup_{s \in \mathcal{S}} T(s)$ denotes the set of all relevant effective topologies.

3.3 Calculation of the CCDF of the Relative Link Load

Network failures lead to rerouting and changed load situations on many links. To detect the risk of potential overload due to failures, we derive the CCDF of the relative loads on all links based on the relevant failure scenarios $s \in \mathcal{S}$ and their probabilities $p(s)$. An aggregate g symbolizes the traffic between two specific routers and $\mathcal{G}$ is the set of all aggregates. The rate of a single aggregate is $c(g)$ and it is given by the traffic matrix. The routing function $u(T, l, g)$ determines the fraction of the rate $c(g)$ of aggregate g

[1] $|\mathcal{X}|$ denotes the cardinality of a set $\mathcal{X}$.

that flows over link l in the effective topology T. It allows the computation of the traffic rate on link l in the presence of effective topology T by $c(T,l) = \sum_{g \in \mathcal{G}} c(g) \cdot u(T,l,g)$.

Input: set of effective topologies $\mathcal{T}$
  for all $T \in \mathcal{T}$ **do**
    CALCULATEROUTING(T)
    for all $l \in \mathcal{E}$ **do**
      $c(T,l) \leftarrow 0$ {initialization}
      for all $g \in \mathcal{G}$ **do**
        $c(T,l) \leftarrow c(T,l) + c(g) \cdot u(T,l,g)$
      end for
      $\mathcal{L}(l) \leftarrow \mathcal{L}(l) \cup (T, c(T,l))$
    end for
  end for
Output: link-specific load sets $\mathcal{L}(l)$ for $l \in \mathcal{E}$

Algorithm 2: CALCULATELOAD: calculates the load $c(T,l)$ for each link $l \in \mathcal{E}$ and for all considered effective topologies $T \in \mathcal{T}$.

The load set $\mathcal{L}(l)$ contains all tuples $(T, c(T,l))$ consisting of the effective topologies $T \in \mathcal{T}$ and the corresponding traffic rates $c(T,l)$ on link l. Algorithm 2 computes the load sets $\mathcal{L}(l)$ for all links $l \in \mathcal{E}$ in an efficient way. The number of tuples $(T, c(T,l)) \in \mathcal{L}(l)$ depends on the set of relevant failure scenarios $\mathcal{S}$ which further depends on the threshold p_{min}. In particular, the probability of all relevant failure scenarios $p(\mathcal{S}) = \sum_{s \in \mathcal{S}} p(s)$ is smaller than 1 in most practical scenarios. Based on the traffic rates $c(T,l)$ and the capacity $c(l)$ of link l, its relative load $U(T,l) = \frac{c(T,l)}{c(l)}$ can be calculated. This link load $U(T,l)$ differs from a link utilization by the fact that it can take values larger than 1, neglecting that traffic may be lost due to congestion. It is useful to estimate the link bandwidth required to carry the traffic without any loss.

Finally, we can calculate the *conditional* CCDF of the relative link load by

$$p(U(l) > u | \mathcal{S}) = \frac{1}{p(\mathcal{S})} \cdot \sum_{\{s : s \in \mathcal{S} \wedge U(T(s),l) > u\}} p(s). \tag{1}$$

We can also give a lower and upper bound for the *unconditioned* CCDF of $U(l)$ by $p^{lower}_{bound}(U(l) > u) = p(U(l) > u | \mathcal{S}) \cdot p(\mathcal{S})$ and $p^{upper}_{bound}(U(l) > u) = p(U(l) > u | \mathcal{S}) \cdot p(\mathcal{S}) + (1 - p(\mathcal{S}))$.

4 Application Examples

We illustrate the above presented concept in an example network. We show the impact of the probability threshold p_{min} on the trustworthiness of the obtained CCDFs of the relative link load and demonstrate that the CCDFs do not establish an absolute order among the links. Therefore, we introduce different assessment functions that condense the information of the CCDF into a single value to get a simple result for the risk of

overload and to facilitate a comparison among links. The assessment functions also help to visualize the potential overload of the entire network at a glance. Finally, we give some examples for the applicability of the analysis in practice.

4.1 Test Environment

In the following, we apply the above presented analysis to the topology depicted in Figure 3 which is the basic structure of a typical core network in the U.S. There is one traffic aggregate $g=(v,w)$ for each pair of nodes v and w, and we define a static aggregate rate

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

where $\pi(v)$ is the population of city $v\in\mathcal{V}$ and C is the rate of the overall network traffic. The populations for all cities associated with the nodes in our test network are taken from [26].

We assume hop-count based shortest path routing and rerouting using the equal-cost multipath (ECMP) option. We dimension the link capacities of our test network such that they are utilized by 20% in the non-failure case. Therefore, the choice of the overall rate C is irrelevant as we look only at the relative link load. This is an artificial scenario as it disregards available granularities for link capacities. However, we use this artificial setting only to illustrate our framework and we do not derive any results that are biased by this simplifying assumption.

For the sake of simplicity, we limit our view to bidirectional link failures and node failures as basic failure events $\hat{s}$. We use the unavailability values given in Section 2.1 as failure probabilities $p(\hat{s})$. However, our tool is able to handle also more detailed failures such as those of line cards or single interfaces. In addition, SRRGs can also be modelled.

4.2 CCDF of the Relative Link Load

We first study the impact of the probability threshold p_{min} that controls the set of relevant failure scenarios $\mathcal{S}$ taken into account in the analysis. Then, we compare the CCDF for different links and show that it is not possible to establish a link order based on the CCDF of the relative link load.

Impact of the Probability Threshold p_{min} on the CCDF Figure 1(a) shows the conditional CCDF of the relative link load for link Dal→Was. The x-axis indicates the relative link load u and the logarithmic y-axis the probability $p(U(l)>u|\mathcal{S})$ that this value is exceeded. The points below the curve represent the effective topologies $T\in\mathcal{T}(\mathcal{S})$ that result from the relevant scenarios $\mathcal{S}$ and that cause the decay of the CCDF. Their coordinates consist of the relative link loads $U(T,l)$ and the probability $p(T)$. In our software the points for the effective topologies are sensitive such that the set of subsumed failure scenarios is displayed when the mouse is dragged over them.

The curve in Figure 1(a) is calculated based on a threshold of $p_{min}=10^{-6}$ which leads to a set of $|\mathcal{S}|=3260$ relevant failure scenarios with an overall probability of $p(\mathcal{S})=0.994504$. The graph also shows the lower and upper bound for the unconditioned CCDF. The probability threshold $p_{min}=10^{-6}$ leaves a large uncertainty regarding

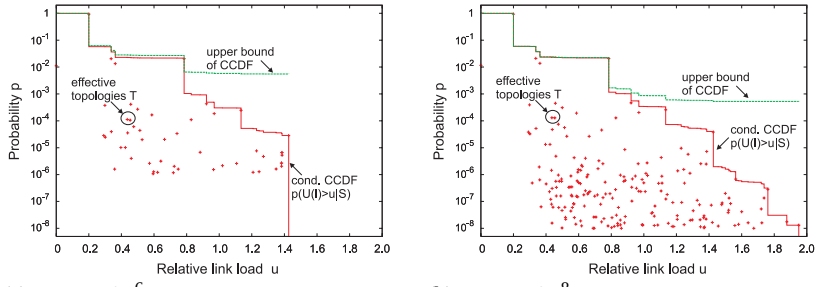

(a) $p_{min} = 10^{-6} \rightarrow$
$|\mathcal{S}| = 3260, \quad p(\mathcal{S}) = 0.994504.$

(b) $p_{min} = 10^{-8} \rightarrow$
$|\mathcal{S}| = 42030, \quad p(\mathcal{S}) = 0.999475.$

Fig. 1. Conditional CCDF of the relative link load $U(l)$ for link Dal→Was together with the lower and upper bound for the unconditioned CCDF.

the unconditioned CCDF in the range of interest where the link tends to be overloaded. Therefore, we plot the CCDF for $p_{min} = 10^{-8}$ in Figure 1(b). As a consequence, the set of relevant failure scenarios is now significantly larger such that it covers a probability of $p(\mathcal{S}) = 0.999475$. The curve has now a different shape in the right part of the graph and the distance between upper and lower bound for the conditioned CCDF is significantly smaller. In the following we use $p_{min} = 10^{-8}$.

Comparison of the CCDFs for Different Links The conditional CCDF of the relative link load $U(l)$ of a link l contains the maximum information about its overload probability. If the CCDF $p(U(l_0) > u|\mathcal{S})$ of the relative load of a link l_0 lies for all 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 2 shows that link Dal→Was has a lower CCDF value than link SaF→Was for some relative link load values u (e.g. $u = 0.4$), and for some other values this is vice-versa (e.g. $u = 0.6$). Therefore, the CCDF does not provide an order relation for links and it is not appropriate for rankings of the links according to their potential overload.

4.3 Simple Assessment Functions for Potential Overload

The objective of our resilience analysis is to identify links that are most likely to be overloaded, but the CCDF cannot achieve that goal. Therefore, we propose three different assessment functions $R(l)$ that condense the information of the CCDF to a single value. They may be used for link rankings and to identify the most critical links.

Assessment Function Based on Overload Probabilities The network provider can define a critical relative link load value u_c that should not be exceeded. Thus, we define the assessment function for potential overload on link l by $R_{u_c}(l) = p(U(l) > u_c|\mathcal{S})$. Note that this ranking depends on the critical relative link load value u_c. Table 1 presents the rankings for $u_c \in \{0.3, 0.6, 0.9\}$ and shows that the value u_c indeed influences the ranking for the selected links.

Assessment Function Based on Relative Link Load Percentiles Another assessment function uses percentiles of the relative link load, i.e. the relative link load values

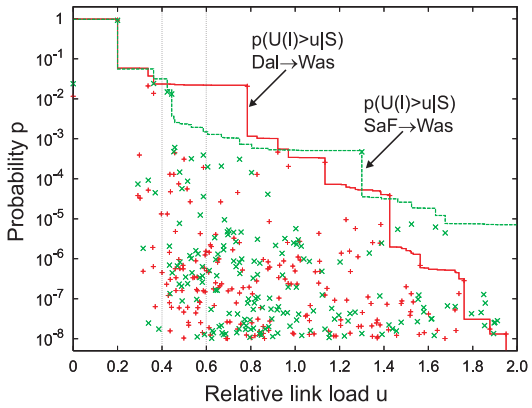

Fig. 2. Conditional CCDF of the relative link load for link Dal→Was and SaF→Was for $p_{min} = 10^{-8}$.

Table 1. Links ranked according to the overload probability $R_{u_c}(l)$.

link id	$R_{u_c}(l)$, $u_c = 0.3$	link id	$R_{u_c}(l)$, $u_c = 0.6$	link id	$R_{u_c}(l)$, $u_c = 0.9$
SaF-Sea	0.0734	LoA-SaF	0.0329	LoA-SaF	0.0324
LoA-SaF	0.0600	Den-SaF	0.0318	SaF-Sea	0.0285
Den-SaF	0.0540	SaF-Sea	0.0291	Den-SaF	0.0252

$R_q(l) = min(u : p(U(l) \leq u|\mathcal{S}) \geq q)$. It depends on the percentile parameter $0 \leq q \leq 1$. Table 2 shows the rankings for $q \in \{0.999, 0.99999\}$ and makes the dependency of R_{u_c} on the percentile parameter q obvious for selected links.

Table 2. Links ranked according to the relative link load percentile $R_q(l)$.

link id	$R_q(l)$, $q = 0.999$	link id	$R_q(l)$, $q = 0.99999$
Sea-Tor	1.512	NeO-Orl	8.761
Kan-SaF	1.3	Kan-SaF	2.628
NeO-Orl	0.2	Sea-Tor	2.198

Assessment Function Based on Weighted Relative Link Loads The above overload measures consider only a single point within the conditional CCDF of the relative link load $U(l)$, 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 $w(u)$:

$$R_w(l) = \int_0^{u_{max}} p(U(l) > u|\mathcal{S}) \cdot w(u) du \tag{3}$$

and we choose $w(u) = 10^{e_{mlwd} \frac{u}{u_{max}}}$ whereby e_{mlwd} is the maximum logarithmic weight difference. This assessment function respects all relative link load values up to u_{max} in the diagram. Thus, the ranking depends on u_{max} and e_{mlwd}. Table 3 shows the rankings

for $u_{max} = 1$ and $e_{mlwd} \in \{2,4\}$ and makes the influence of the latter parameter explicit for selected links.

Table 3. Links ranked according to their weighted relative link load $R_w(l)$.

link id	$R_w(l)$, $e_{mlwd} = 2$	link id	$R_w(l)$, $e_{mlwd} = 4$
Dal-Was	0.514	Mia-Orl	6.7
Sea-Tor	0.511	Sea-Tor	5.926
Mia-Orl	0.478	Dal-Was	4.526

4.4 Potential Overload at a Glance

We show the potential overload in the network at a glance by displaying its topology with gray shaded links. The different intensities indicate the risk of overload due to network failures and, therefore, the value of the assessment function $R(l)$ is translated into an intensity value. A slide bar for the parameter u_c, q, or e_{mlwd} allows to change the intensities to get the suitable contrast. This is well feasible in realtime since the calculation of the above assessment functions needs only the stored CCDF but no further time-consuming analysis. Figure 3 illustrates the concept for the assessment function based on overload probabilities with $u_c = 0.6$.

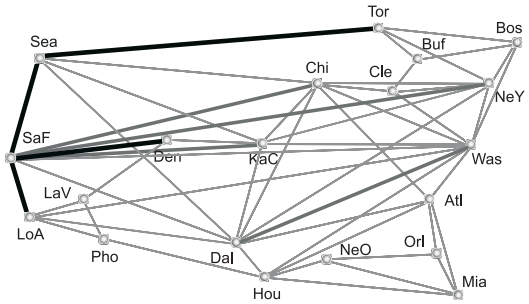

Fig. 3. The contrast of the links indicates their risk of overload due to network failures: dark links are more likely to be overloaded than light links.

4.5 Application Scenarios for A Priori Detection of Link Overload

Our tool allows a network operator to detect link overload a priori, i.e. before congestion occurs in the network due to redirected traffic in failure cases, and helps them to dimension the link bandwidths large enough to support QoS also during local network outages. The tool may be applied before the network configuration is changed to anticipate the impact of the changes on the potential overload. We list some examples of such changes.

- The tool helps to analyze whether a planned bandwidth upgrade of a link is sufficient to improve its potential overload or whether the upgrade is even wasteful.
- When new links or nodes are added to the network, the routing changes in the failure-free scenario and also in failure scenarios which impacts the potential overload on the links.

– When link metrics are changed, the routing changes both in the failure-free scenario and in failure scenarios which also impacts the potential overload on the links. It also helps to optimize link metrics manually.
– When new customers are added, the traffic matrix changes which impacts the relative load of the links and thereby their potential overload in the network.
– SRLGs of leased lines are not always known in advance. If new knowledge about them is available, a new failure event $\hat{s}$ is added that entails the simultaneous failure of the links in the SRLG. This may have a tremendous effect on the potential overload.

5 Conclusion

Network failures trigger restoration or protection switching mechanisms that redirect traffic onto backup paths which increases the relative load of their links. In this paper, we have proposed a framework to assess the risk of link overload in a network due to failures that occur with certain probabilities. We implemented a tool that derives an approximative complementary cumulative distribution function (CCDF) of the relative link load for all links in the network. The analysis considers only the set of relevant failure scenarios which have a probability larger than a minimum threshold p_{min}.

As the full information of the CCDF is often too complex for practical applications, we proposed to condense it to a single value by so-called risk assessment functions for which we discussed three basically different approaches. Their outcome can be used to rank links according to their risk of overload. These values lead to a presentation of the analysis results that is useful for network operators to quickly identify the links that have the highest risk to be overloaded although they do not reveal problems in the failure-free case. These links may be upgraded with additional bandwidth to prevent congestion due to redirected traffic in advance.

Currently, we extend our tool to a priori detect overload which may also be due to other causes such as local hot spots [25] or inter-domain rerouting [27].

References

1. Milbrandt, J., Martin, R., Menth, M., Hoehn, F.: Risk Assessment of End-to-End Disconnection in IP Networks due to Network Failures. In: 6^{th}IEEE Workshop on IP Operations and Management (IPOM), Dublin, Ireland (2006)
2. Datta, P., Somani, A.K.: Diverse Routing for Shared Risk Resource Groups (SRRG's) in WDM Optical Networks. In: 1^{st}IEEE International Conference on Broadband Communication, Networks, and Systems (BROADNETS). (2004) 120 – 129
3. Vasseur, J.P., Pickavet, M., Demeester, P.: Network Recovery. 1. edn. Morgan Kaufmann / Elsevier (2004)
4. Mukherjee, B.: Optical WDM Networks. 2 edn. Springer (2006)
5. Iannaccone, G., Chuah, C.N., Mortier, R., Bhattacharyya, S., Diot, C.: Analysis of Link Failures in an IP Backbone. In: ACM SIGCOMM Internet Measurement Workshop, Marseille, France (2002) 237 – 242
6. Markopoulou, A., Iannaccone, G., Bhattacharyya, S., Chuah, C.N.: Characterization of Failures in an IP Backbone. In: IEEE Infocom, Hongkong (2004)
7. Willems, G., Arijs, P., Parys, W.V., Demeester, P.: Capacity vs. Availability Trade-offs in Mesh-Restorable WDM Networks. In: International Workshop on the Design of Reliable Communication Networks (DRCN), Budapest, Hungary (2001)

8. Cankaya, H.C., Lardies, A., Ester, G.W.: 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 (2004)
9. Maesschalck, S.D., Colle, D., Lievens, I., Pickavet, M., Demeester, P., Mauz, C., Jaeger, M., Inkret, R., Mikac, B., Derkacz, J.: Pan-European Optical Transport Networks: an Availability-Based Comparison. Photonic Network Communications **5** (2005) 203–225
10. Iannaccone, G., Chuah, C.N., Bhattacharyya, S., Diot, C.: Feasibility of IP Restoration in a Tier-1 Backbone. IEEE Network Magazine (Special Issue on Protection, Restoration and Disaster Recovery) (2004)
11. Fortz, B., Rexford, J., Thorup, M.: Traffic Engineering with Traditional IP Routing Protocols. IEEE Communications Magazine **40** (2002) 118–124
12. Fortz, B., Thorup, M.: Robust Optimization of OSPF/IS-IS Weights. In: International Network Optimization Conference (INOC), Paris, France (2003) 225–230
13. Nucci, A., Schroeder, B., Bhattacharyya, S., Taft, N., Diot, C.: IGP Link Weight Assignment for Transient Link Failures. In: 18^{th} International Teletraffic Congress (ITC), Berlin (2003)
14. Yuan, D.: A Bi-Criteria Optimization Approach for Robust OSPF Routing. In: 3^{rd}IEEE Workshop on IP Operations and Management (IPOM), Kansas City, MO (2003) 91 – 98
15. Sridharan, A., Guerin, R.: Making IGP Routing Robust to Link Failures. In: IFIP-TC6 Networking Conference (Networking), Ontario, Canada (2005)
16. Cholda, P., Jajszczyk, A.: Availability Assessment of Resilient Networks. In: 12^{th}GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems (MMB) together with 3^{rd}Polish-German Teletraffic Symposium (PGTS), Dresden, Germany (2004) 389–398
17. Li, V.O.K., Silvester, J.A.: Performance Analysis of Networks with Unreliable Components. IEEE Transactions on Communications **32** (1984) 1105–1110
18. Clouqueur, M., Grover, W.D.: Computational and Design Studies on the Unavailability of Mesh-restorable Networks. In: International Workshop on the Design of Reliable Communication Networks (DRCN), Munich, Germany (2000) 181 – 186
19. Clouqueur, M., Grover, W.D.: Availability Analysis of Span-Restorable Mesh Networks. IEEE Journal on Selected Areas in Communications **20** (2002) 810 – 821
20. Schupke, D.A., Prinz, R.G.: Capacity Efficiency and Restorability of Path Protection and Rerouting in WDM Networks Subject to Dual Failures. Photonic Network Communications **8** (2004)
21. Menth, M., Martin, R., Spoerlein, U.: Impact of Unprotected Multi-Failures in Resilient SPM Networks: a Capacity Dimensioning Approach. In: IEEE Globecom, San Francisco, California, USA (2006)
22. Durvy, M., Diot, C., Taft, N., Thiran, P.: Network Availability Based Service Differentiation. In: 11^{th}IEEE International Workshop on Quality of Service (IWQoS), Berkeley, CA, USA (2003) 305–324
23. Dahlin, M., Chandra, B.B.V., Gao, L., Nayate, A.: End-to-End WAN Service Availability. IEEE/ACM Transactions on Networking **11** (2003) 300–313
24. Feldmann, A., Greenberg, A., Lund, C., Reingold, N., Rexford, J.: NetScope: Traffic engineering for IP Networks. IEEE Network Magazine (2000) 11–19
25. Menth, M., Martin, R., Charzinski, J.: Capacity Overprovisioning for Networks with Resilience Requirements. In: ACM SIGCOMM, Pisa, Italy (2006)
26. Menth, M.: Efficient Admission Control and Routing in Resilient Communication Networks. PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland (2004)
27. Schwabe, T., Gruber, C.G.: Traffic Variations Caused by Inter-domain Re-routing. In: International Workshop on the Design of Reliable Communication Networks (DRCN), Ischia Island, Italy (2005)