Risk Assessment of End-to-End Disconnection in IP Networks due to Network Failures

Jens Milbrandt, Rüdiger Martin, Michael Menth, and Florian Höhn

Dept. of Distributed Systems, Inst. of Computer Science, University of Würzburg
Am Hubland, 97072 Würzburg, Germany
Phone +49 931 8886631, Fax +49 931 8886632
{milbrandt,martin,menth,hoehn}@informatik.uni-wuerzburg.de

Abstract. When network failures occur, resilience mechanisms can restore the connectivity as long as the network is physically connected. In this paper, we assess the risk of the physical disconnection of a network which reduces the service availability that is part of service level agreements (SLA). Our analysis takes into account the most probable failures and provides upper bounds on the resulting end-to-end disconnection probabilities. The analysis helps to visualize the overall service availability, to identify routers or points of presence (PoPs) with a high risk to get disconnected, and to evaluate the impact of new links on the service availability. Thus, it assists network planers in upgrading their networks. Furthermore, the analysis verifies whether individual aggregates of already well connected routers still have a high risk to get disconnected, i.e. whether they are likely to break the SLA.

1 Introduction

Network resilience in carrier grade networks comprises the maintenance of both connectivity and quality of service (QoS) during network failures. To maintain logical connectivity in the presence of link or node failures, restoration or protection switching mechanisms redirect traffic over backup paths. If the nodes of a network are only sparsely interconnected or if multiple network failures occur simultaneously, the risk of physical disconnection increases. As a consequence, resilience mechanisms are no longer effective if a network becomes disconnected after a failure. This leads inevitably to violations of service level agreements (SLAs) with customers or peering network providers. While the service availability in cases where logical connectivity can still be maintained is frequently in the focus of network analysis [1–3], the risk assessment of physical network disconnection due to network failures has not yet been investigated in depth.

In this paper, we present a method to assess the risk of end-to-end disconnection in IP networks due to network failures. Network disconnection can occur due to multiple independent or correlated network failures. Correlated failures are due to shared

This work in cooperation with Infosim GmbH & Co.KG was funded by the Bavarian Ministry of Economic Affairs. The authors alone are responsible for the content of the paper.

risk resource groups (SRRGs) [4]. For instance, links that are logically distinct on the network layer but share common resources on the link layer fail simultaneously if their common resource fails. We consider both types of multiple failures. The probability for multiple independent failures is rather low in a small network while in large networks their impact cannot be neglected. Our disconnection analysis therefore considers all relevant single and multi-failures that occur with a probability larger than a minimum threshold. Based on the occurrence probability of the relevant failure scenarios and their impact on the network connectivity, we calculate the disconnection probabilities for all pairs of nodes in the network. We consider the results from different perspectives to asses the overall service availability, to identify points of presence (PoPs) at risk due to insufficient connectivity, and finally to assess the SLA compliance of individual aggregates. This helps to identify weak spots of the network and to appropriately upgrade its infrastructure with additional links. We implemented the concept in a software tool which helps network providers to assess the risk of end-to-end disconnection in their networks prior to any failure and to take appropriate actions.

This paper is structured as follows. In Section 2, we review related work regarding network resilience. Section 3 explains our algorithms to determine a set of relevant network failures and to calculate thereon the aggregate-specific disconnection probabilities. Section 4 presents numerical results for an example network and considers them from different perspectives. It illustrates the potential of our analysis tool. Finally, Section 5 concludes this work.

2 Network Failures and Network Survivability

In this section, we review causes for network failures. Some failures have only a local impact and the corresponding network outage can be compensated by resilience mechanisms. However, failures of larger degree and simultaneous multi-failures jeopardize the physical connectivity of a network and might cause network partitioning. We give an overview of current resilience mechanisms, discuss related work on network survivability, and comment on our contribution to assess the risk of physical network disconnection.

2.1 Network Failures

A good overview and characterization of causes for network failures is given in [5]. Basically, network failures with internal causes (e.g. software bugs, component defects, etc.) can be distinguished from those with external causes (e.g. digging works, natural disaster, etc.). Furthermore, planned failure causes can be distinguished from unplanned failure causes. Planned outages are normally due to network maintenance and, since they are intentional, operators can take preventive measures. Planned network outages should not lead to physical network disconnections whereas unplanned outages jeoparidze the network connectivity. The latter are difficult to predict and, therefore, operators must rely on resilience mechanisms in their networks. However, these mechanisms are useless if a network becomes disconnected. To avoid physical disconnection, operators have to construct their network topology carefully.

Quantitative analyses and statistics about frequency and duration of failure events that occur in an operational network like the Sprint IP backbone are given in [6, 7]. The authors detect all failures affecting the (logical) IP connectivity by analyzing the link state advertisements (LSAs) of interior gateway routing protocols. The results show that link failures are part of everyday's network operation and the majority of them is short-lived (less than 10 minutes). Moreover, they indicate that 20% of all failures are due to planned maintenance activities. Among the unplanned failures, almost 30% hit multiple links and can be attributed to router-related and optical equipment-related problems, while 70% affect only a single link at a time.

2.2 Resilience Mechanisms

Resilience mechanisms can be divided into restoration and protection schemes. Restoration sets up a new path after a failure while protection switching pre-establishes backup paths in advance. A good overview can be found in [5].

Usually, restoration is applied by IP rerouting. IP networks have the self-healing property, i.e., their routing re-converges after a network failure through the exchange of LSAs such that all but failed nodes can be reached after a while if the network is still physically connected. The reconvergence of the IP routing algorithm is a very simple and robust restoration mechanism [8, 9]. However, a clear disadvantage of IP routing is its slow convergence which is tolerable for elastic traffic but not for realtime traffic with strict QoS constraints.

Protection switching mechanisms address the problem of slow reconvergence speed. They can be implemented, e.g. in multi-protocol label switching (MPLS) technology by explicitly routed and pre-established backup paths. Depending on the location of the reaction to a failure, protection switching mechanisms can be distinguished into end-to-end and local protection. In case of end-to-end protection switching, the reaction to a failure within a path is executed at the head end router. Local protection schemes deviate the traffic at the router which is closest to the outage location to achieve an extremely fast reaction time [5, 10].

2.3 Network Survivability

The previosuly described resilience mechanisms are effective only if a network does not break apart after a failure, i.e., if all network nodes are still physically interconnected. Physical connectivity is a basic and most critical requirement for network survivability which can also be considered from different point of views [11]. Network survivability in terms of physical connectivity is a matter of topological network design and is thus subject to optimization [12]. Different network topologies have different survivability characteristics and require different strategies to improve survivability. Self-healing ring networks [13] represent a popular topology for metropolitan area networks (MANs) and have a good survivability potential since they operate with full hardware redundancy. Optical mesh networks [14, 15] characteristic for wide area networks (WANs) and their survivability can be improved by the installation of additional links. Network survivability frameworks (cf. e.g. [11, 16]) define assessment and analysis models as well as

different performance measures for the evaluation of multi-layer network survivability. These frameworks are also applied in network analysis and network design [17–19].

2.4 Tools for Network Analysis and Network Design

We implemented our approach to assess the risk of end-to-end disconnection in IP networks due to network failures in software since we are not aware of any standard tool for that purpose. Generally, standard simulation software like the OPNET Modeler [20] or the Network Simulator (ns-2) [21] provide means for analyzing the resilience of a network. However, these simulators focus on the dynamics of network resilience mechanisms with static sets of (single) network failures. They do not provide appropriate means for the calculation of probabilities for physical disconnection in a network for which a large amount of failures must be considered.

Other software products like e.g., the library of test instances for Survivable fixed telecommunication Network Design (SNDlib) [22], the TOolbox for Traffic Engineering Methods (TOTEM) [23], or the NetScope tool for traffic engineering in IP networks [24] focus on the evaluation of traffic engineering algorithms for network optimization. These tools cover only the usual network design and optimization problems such as routing, load balancing, flow allocation, and network dimensioning. However, they have no advanced functions to assess the risk of network disconnection.

2.5 Contribution of this Work

The above mentioned approaches 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, simultaneous multi-failures need to be taken into account if the network size increases. Our objective is to assess the risk of end-to-end disconnection in IP networks due to network failures and to improve network survivability in terms of physical connectivity through identification of weak spots in the network.

The novelty of this work is the integration of outage probabilities in the assessment of physical survivability of a network. We present an assessment method that yields histograms of disconnection probabilities from the perspectives of the set of b2b traffic aggregates, a single network node, or the set of all network nodes. This helps Internet service providers (1) to detect weak spots in their network and (2) to improve the survivability of their network by the systematical installation of additional links. We currently develop a tool that predicts the risk of end-to-end disconnection before and after such a network modification to support the ISP in his decision process.

3 End-to-End Disconnection due to Network Failures

We analyze the impact of potential failure scenarios on the physical connectivity of the network. As not all failure scenarios can be covererd for that analysis, we determine only the most relevant. Some failure scenarios lead to the same working ("effective") topology which, in turn, leads to the same end-to-end disconnections in the network. We handle them jointly for the calculation of the disconnection probabilities.

3.1 Relevant Failure Scenarios

We first identify the relevant failure scenarios for our resilience analysis. To that aim, we collect all independent failure events $\hat{s} \in \hat{\mathcal{S}}$, $|\hat{\mathcal{S}}| = n$. Note that this set may also contain shared risk resource groups (SRRGs) such as shared risk link or node groups (SRLG, SRNG) [4]. Each of these failure events occurs with probability $p(\hat{s})$ and we number the events $\hat{s}_i$ in an descending order according to $p(\hat{s}_i)$. We define a compound failure scenario $s \subseteq \hat{\mathcal{S}}$ as a subset of independent failure events that occur simultaneously with $p(s) = (\Pi_{\hat{s} \in s} p(\hat{s})) \cdot \left(\Pi_{\hat{s} \in \hat{\mathcal{S}} \setminus s} (1 - p(\hat{s}))\right)$. 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 failure scenarios. Algorithm 1 constructs the set $\mathcal{S}$ starting with $\mathcal{S} = \emptyset$ at the beginning. The recursive procedure is invoked with RELEVANTSCENARIOS$(0, \emptyset, 1)$. The algorithm steps recursively through the set of independent failure events $\hat{s}_i \in \hat{\mathcal{S}}$ for $0 \leq i \leq n$. It constructs a compound failure scenario s incrementally and the recursion ends either if all n independent failure events $\hat{s}_i$ have been considered as potential members of s or if the probability $p(s)$ of the partial compound failure scenario s is too low. In either case, 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 larger than the threshold p_{min}.

Input: failure event number i, partial scenario s, and its probability $p(s)$
  if $(i = n)$ **then** {all independent failure events $\hat{s}_i$ have been considered}
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$
  else if $(p(s) > p_{min})$ **then** {partial scenario s is probable enough}
    RELEVANTSCENARIOS$(i+1, s \cup \hat{s}_i, p(s) \cdot p(\hat{s}_i))$
    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 only if itself and its adjacent routers do not fail. A router only works 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 Disconnection Probabilities

For the calculation of disconnection probabilities, we assume a network with node set $\mathcal{V}$. The disconnection probability $p_{dis}^{\mathcal{S}}(v,w)$ of a single aggregate $g(v,w)$ between two

network nodes v and w is calculated as

$$p_{dis}^{\mathcal{S}}(v,w) = \frac{1}{p(\mathcal{S})} \cdot \sum_{T \in \mathcal{T}} p(T) \cdot \text{ISDISONNECTEDIN}(v,w,T) \tag{1}$$

under the condition that only the relevant failure scenarios $\mathcal{S}$ and their corresponding effective topologies $\mathcal{T}$ are respected. The function ISDISCONNECTEDIN() yields 1 if nodes v and w are disconnected in the topology T or 0 otherwise. The values $p_{dis}^{\mathcal{S}}(v,w)$ are underestimated since not all possible failure scenarios are considered in Equation (1). To get an upper bound for $p_{dis}^{\mathcal{S}}(v,w)$, we can calculate the unconditioned disconnection probability $p_{dis}^{max}(v,w)$ for the aggregate $g(v,w)$ as

$$p_{dis}^{max}(v,w) = p(\mathcal{S}) \cdot p_{dis}^{\mathcal{S}}(v,w) + (1 - p(\mathcal{S})) \cdot 1 \tag{2}$$

under the assumption that the nodes v and w are disconnected in all unconsidered failure scenarios $s \in \hat{\mathcal{S}}^n \setminus \mathcal{S}$. To illustrate the application of our concept in Section 4, we will use, nevertheless, the conditioned disconnection probabilities $p_{dis}^{\mathcal{S}}(v,w)$. This is justified by chosing a small threshold p_{min} which covers a large set of failure scenarios and decreases the uncertainty inherent to Equation (1).

4 Application of the Concept

This study is limited to link or node failures only. However, our software tool is able to handle correlated elemental failures of general shared risk resource groups, as well.

4.1 Test Environment

To give a numerical example for our end-to-end disconnection analysis, we use the NOBEL network topology depicted in Figure 1. The set $\mathcal{V}$ comprises all network nodes each associated with a european city. For each pair of ingress/egress nodes v and w, we define a static aggregate rate

$$r(v,w) = \begin{cases} \frac{\pi(v) \cdot \pi(w) \cdot \mathcal{R}}{\sum_{x,y \in \mathcal{V}, x \neq y} \pi(x) \cdot \pi(y)} & \text{if } v \neq w \\ 0 & \text{if } v = w \end{cases} \tag{3}$$

where $\pi(v)$ is the population of city $v \in \mathcal{V}$ and surroundings and $\mathcal{R}$ is the rate of the overall network traffic. The populations of all cities represented by nodes in our test network are shown in Figure 1 and are used to calculate the traffic matrix according to Equation (3).

The probability $p(\hat{s})$ of a failure event $\hat{s}$ depends on the availability of the corresponding network element. For the sake of simplicity, we set the node failure probabilities to $p_{node} = 10^{-6}$ and compute the link failure probabilities p_{link} according to [5] as $p_{link} = \frac{\text{MTTR}}{\text{MTBF}}$ where MTTR = 24 h is the mean time to repair and MTBF = $\frac{450 \cdot 365 \cdot 24}{L}$ h is the mean time between failures for a link with length L in kilometers. We set the minimum probability threshold for relevant failure scenarios to $p_{min} = 10^{-14}$. This value covers a large set of $|\mathcal{S}| = 513957$ relevant failure scenarios with very small uncertainty of $1 - p(\mathcal{S}) = 3.3 \cdot 10^{-9}$. Since this error is negligible, we do not show $p_{dis}^{max}(v,w)$ and concentrate on $p_{dis}^{\mathcal{S}}(v,w)$ only in the following sections instead.

ID(v)	name(v)	π(v)
1	Oslo	801.028
2	Stockholm	1.872.900
3	Madrid	5.964.143
4	Athen	3.187.734
5	Glasgow	1.168.270
6	Dublin	1.600.000
7	Barcelona	3.120.000
8	Bordeaux	753.931
9	Copenhagen	1.212.485
10	Paris	9.644.507
11	London	8.278.251
12	Berlin	3.388.477
13	Rome	2.542.003
14	Hamburg	2.532.565
15	Warsaw	1.692.854
16	Munich	1.920.063
17	Vienna	1.878.759
18	Frankfurt	1.902.815
19	Budapest	1.695.000
20	Lyon	1.348.832
21	Amsterdam	1.453.003
22	Belgrade	1.120.092
23	Milan	1.271.898
24	Prague	1.165.581
25	Zurich	1.075.230
26	Brussels	1.007.000
27	Zagreb	691.724
28	Strasbourg	427.245

Fig. 1. European NOBEL test network and associated city populations.

4.2 Overall Service Availability

Figure 2 shows the disconnection probabilities $p^{\mathcal{S}}_{dis}(v,w)$ in logarithmical scale for all 756 ($= 28 \cdot 27$) aggregates between any source-destination pair (v,w) calculated by Equation (1). For each curve, the aggregates are sorted along the x-axis according to their disconnection probability in descending order for better readability. The upmost curve corresponds to the original network. More than half of all aggregates are disconnected with a probability of $p^{\mathcal{S}}_{dis} \in [10^{-5}, 3 \cdot 10^{-5}]$ while the probability for the remaining aggregates is one order of magnitute lower with $p^{\mathcal{S}}_{dis} = 2 \cdot 10^{-6}$. The lower curves correspond to the network with additional links installed and we explain their meaning later in this work.

This presentation from the perspective of the entire set of aggregates provides a good overview regarding the overall availability of the network. However, Figure 2 does not reveal particular nodes whose network connectivity should be improved.

4.3 Detection of PoPs at Risk

To detect network nodes with potentially unsufficient connectivity we summarize the disconnection probabilities for individual routers. To that aim, we calculate the overall disconnection probability $\overline{p}^{\mathcal{S}}_{dis}(v)$ of all aggregates of a single router v weighted by their traffic rate by

$$\overline{p}^{\mathcal{S}}_{dis}(v) = \sum_{w \in \mathcal{V}, w \neq v} p^{\mathcal{S}}_{dis}(v,w) \cdot \frac{r(v,w)}{\sum_{x,y \in \mathcal{V}, x \neq y} r(x,y)} \tag{4}$$

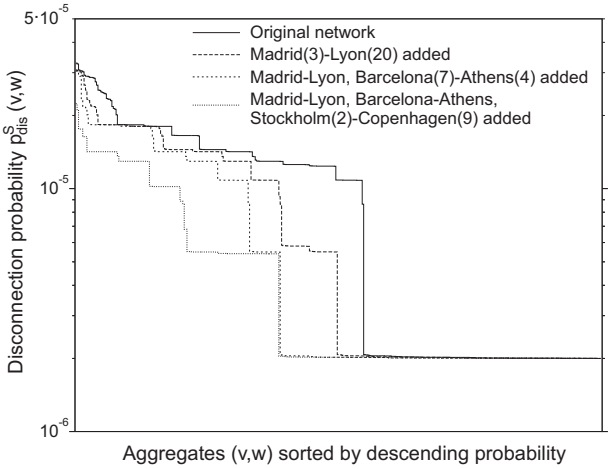

Fig. 2. Disconnection probabilities from the perspective of the entire set of aggregates.

under the condition that only the relevant failure scenarios $\mathcal{S}$ are respected. The term $\overline{p}_{dis}^{\mathcal{S}}(v)$ expresses the probability that traffic originating at or destined to router v is disconnected.

Figure 3 shows $\overline{p}_{dis}^{\mathcal{S}}(v)$ for all 28 routers of the NOBEL network. The nodes are arranged along the x-axis according to their node ID from Figure 1. We assigned their IDs so that the routers appear in descending order of their disconnection probability $\overline{p}_{dis}^{\mathcal{S}}(v)$ in the orignal network for easier readability. Each column corresponds to a single router v and its width is proportional to the traffic volume $\frac{\sum_{w \in \mathcal{V}} r(v,w)}{\sum_{x,y \in \mathcal{V}, x \neq y} r(x,y)}$ transported from and to v. The tall white columns represent the probabilities of the routers to get physically disconnected from other nodes in the original network. The grey columns represent the disconnection probabilities after inserting additional links as described later. Figure 3 reveals nine routers with higher disconnection probabilities than other nodes. Those are routers 1 through 9 that, e.g., already suffer from disconnection if a double link failure occurs since they have only two adjacent links. This finding also explains the probabilities $p_{dis}^{\mathcal{S}}(v,w) \in [10^{-5}, 3 \cdot 10^{-5}]$ in Figure 2. These probabilities can be observed for aggregates originating at or destined to one of these nine nodes.

Having identified the nodes with unreliable network connectivity, the network operator can improve these probabilities, e.g., by installing additional links. For demonstration purposes, we first add the relatively short link Madrid-Lyon since Madrid (ID 3) is the largest one among those nine nodes as indicated by the column width and therefore causes high traffic losses in case of disconnection from the network. As expected, Madrid profits most from this additional link, but also Barcelona (ID 7) and Bordeaux (ID 8) clearly benefit from the new link in Figure 3. The overall network connectivity increases as seen from the second line from top in Figure 2.

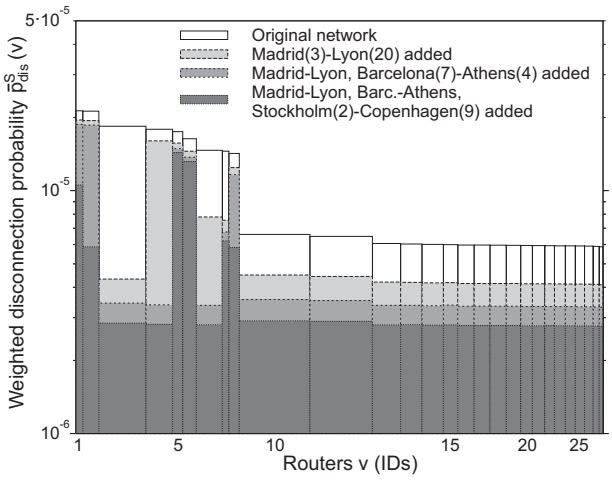

Fig. 3. Disconnection probabilities from the perspective of the set of routers.

Next we add an additional link Barcelona-Athens since the two cities (IDs 7,4) are now the major cities with unsufficient connectivity. The disconnection probabilities decrease in Figures 2 and 3 as intended. Thus, we select the link Stockholm-Copenhagen as third additional link to improve the availability of both medium size cities (IDs 2,9) since they are close to each other. However, none of Figures 2 and 3 shows the same decrease in disconnection probability as before. This is due to the possible double link failure Berlin-Copenhagen and Warsaw-Stockholm that disconnects the triangle Oslo-Stockholm-Copenhagen completly from the network. Thus, our evaluation quickly shows whether a new link improves the network connectivity significantly or not.

Figure 2 also shows that the disconnection probability cannot be reduced below $2 \cdot 10^{-6}$ for individual aggregates. At $p^{\mathcal{S}}_{dis}(v,w) = 2 \cdot 10^{-6}$ the probabilities are dominated by the failure of either the source or the destination router of the respective aggregates. Only increasing the node reliability, e.g., by installing redundant hardware, further improves this value.

In practice, our methodology to assess the risk of possible end-to-end disconnection enables network operators to identify points of presence (PoPs) in the network where topological changes are required. Our software evaluates whether these changes improve the connectivity significantly or not. It confirms intuition and gives additional hints since it considers much more relevant scenarios than can be seen at first sight. Both perspectives presented above support network planners in strategic decisions.

Disconnection Probabilities for individual Aggregates To fulfill service level agreement (SLA) requirements for specific end-to-end connections, information about the availability of single aggregates are required. We use the disconnection probabilites seen from the perspective of a single router for individual aggregates considering their

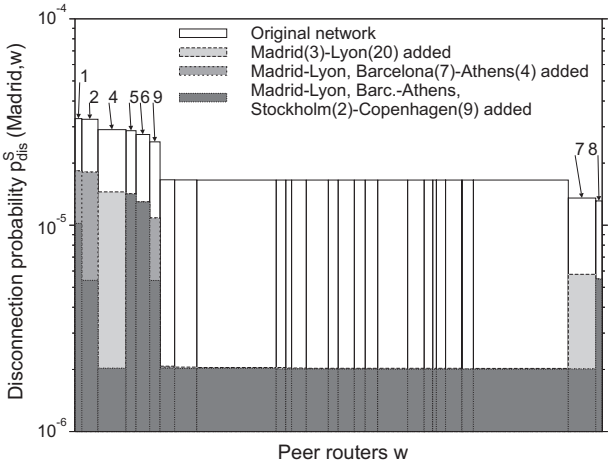

Fig. 4. Disconnection probabilities from the perspective of router *Madrid*.

size. Figure 4 shows the disconnection probabilities $p_{dis}^{\mathcal{S}}(Madrid, w)$ of all aggregates (cf. Equation (1)) from router Madrid to any other router w in the network. These 27 peer routers are sorted along the x-axis in ascending order of their disconnection probability in the original network. Each column corresponds to a single bidirectional aggregate between Madrid and a single peer router w. The IDs for the routers at risk from the previous section are annotated above the respective columns. The column width is proportional to the traffic volume $\frac{r(\text{Madrid},w)}{\sum_{v \in \mathcal{V}} r(\text{Madrid},v)}$ transported in both directions between router Madrid and its peer w and represents the lost traffic in case of disconnection.

From Figure 4 we clearly observe that the first additionally installed link Madrid-Lyon directly impacts the disconnection probabilities of all aggregates since Madrid serving as source or destination is now more reliably connected to the network. Further additional links improve the connectivity for a subset of the aggregates only. This perspective helps to identify where such an improvement is required to meet end-to-end reliability for certain aggregates.

5 Conclusion

In this paper, we proposed a method to assess the risk of end-to-end disconnection in IP networks due to network failures. To that aim, we calculate the disconnection probabilities for all end-to-end aggregates for a set of considered failures $\mathcal{S}$. In contrast to a static analysis that considers only fixed sets of failures (e.g. all single and double link failures), we dynamically construct the set $\mathcal{S}$ for our analysis in such a way that the probability for the set of neglected failure scenarios is low. That allows to give upper bounds on the calculated values. In addition, our tool integrates the option to model correlated link failures such that complex real world scenarios with shared risk groups (SRGs) can be well modelled.

We gave several examples for the application of our analysis. It provides a view on the overall service availability, it helps to identify routers or points of presence with the highest risk to get disconnected, and it tells whether adding a new link improves the service availability significantly or not. We have seen that the introduction of new links can improve the service availability only to a certain degree as long as the availability of border routers or exchange points is not further increased. The analysis from the perspective of a specific router illustrated the fraction of the traffic with high and low service availability and showed that even PoPs with relatively high availability have aggregates that are likely to miss the objectives negotiated in the service level agreements (SLA). Thus, our tool provides crucial information for network planers in upgrading networks and for service providers in specifying feasible SLAs.

Currently, we extend our tool towards overload that may be caused by traffic variations in the network which are due to inter-domain rerouting [25] or due to traffic hot spots.

References

1. Chandra, B., Dahlin, M., Gao, L., A.Nayate: End-to-End WAN Service Availability. In: Proc. of 3^{rd} USENIX Symposium on Internet Technology & Systems (USITS), San Francisco, USA (2001) 97–108
2. Bhattacharyya, S., Diot, C., Iannaccone, G., Markopoulou, A., Chuah, C.: Service Availability in IP Networks. ATL Research Report RR03-ATL-071888, Sprint (2003)
3. Keralapura, R., Chuah, C., Iannaconne, G., Bhattacharrya, S.: Service Availability: A New Approach to Characterize IP-Backbone Topologies. In: Proc. of 12^{th} International Workshop on Quality of Service (IWQoS), Montreal, Canada (2004) 232–241
4. Datta, P., Somani, A.K.: Diverse Routing for Shared Risk Resource Groups (SRRG's) in WDM Optical Networks. In: 1^{st} International Conference on Broadband Communication, Networks, and Systems (BROADNETS). (2004) 120 – 129
5. Vasseur, J., Pickavet, M., Demeester, P.: Network Recovery. Elsevier (2004)
6. Iannaccone, G., Chuah, C., Mortier, R., Bhattacharyya, S., Diot, C.: Analysis of Link Failures in an IP Backbone. In: Proc. of ACM SIGCOMM Internet Measurement Workshop, Marseille, France (2002) 237–242
7. Markopoulou, A., Iannaccone, G., Bhattacharyya, S., Chuah, C., Diot, C.: Characterization of Failures in an IP Backbone. In: Proc. of 23^{rd} IEEE Infocom, Hong Kong, China (2004)
8. 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)
9. Fortz, B., Thorup, M.: Robust Optimization of OSPF/IS-IS Weights. In: International Network Optimization Conference (INOC), Paris, France (2003) 225–230
10. Pan, P., Swallow, G., Atlas, A.: RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels. `http://www.ietf.org/rfc/rfc4090.txt` (2005)
11. Zolfaghari, A., Kaudel, F.J.: Framework for Network Survivability Performance. IEEE Journal on Selected Areas in Communications **12**(1) (1994) 46–51
12. Boorstyn, R.R., Frank, H.: Large-Scale Network Topological Optimization. IEEE Transactions on Communications **25**(1) (1977) 29–47
13. Towster, H., Stephenson, R.W., Morgan, S., Keller, M., Mayer, R., Shalayda, R.: Self-healing ring networks: Gateway to public information networking. IEEE Communications Magazine **28**(6) (1990) 54–60

14. Ramamurthy, S., Mukherjee, B.: Survivable WDM Mesh Networks, Part I - Protection. In: Proc. of IEEE Conference on Computer Communications (INFOCOM). (1999) 744 – 751
15. Ramamurthy, S., Mukherjee, B.: Survivable WDM Mesh Networks, Part II - Restoration. In: Proc. of IEEE International Conference on Communications (ICC). (1999) 2023 – 2030
16. Medhi, D.: A Unified Approach to Network Survivability for Teletraffic Networks: Models, Algorithms and Analysis. IEEE Transactions on Communications **42**(2) (1994) 534–548
17. Cardwell, R.H., Monma, C.L., Wu, T.H.: Computer-Aided Design Procedures for Survivable Fiber Optic Networks. IEEE Journal on Selected Areas in Communications **7**(8) (1989) 1188–1197
18. Gavish, B., Trudeau, P., Dror, M., Gendreau, M., Mason, L.: Fiberoptic Circuit Network Design Under Reliability Constraints. IEEE Journal on Selected Areas in Communications **7**(8) (1989) 1181–1187
19. Cankaya, H., Lardies, A., Ester, G.: Network Design Optimization from an Availability Pespective. In: International Telecommunication Network Strategy and Planning Symposium (Networks), Vienna, Austria (2004)
20. OPNET Incorporated: OPNET Modeler. `http://opnet.com/products/modeler/` (1987)
21. University of Southern California - Information Sciences Institute (USC-ISI): Network Simulator (ns-2). `http://www.isi.edu/nsnam/ns/` (1995)
22. Zuse-Institute Berlin (ZIB): SNDlib 1.0 – Survivable Network Design Data Library. `http://sndlib.zib.de` (2005)
23. Balon, S., Lepropre, J., Monfort, G.: TOTEM 2.0 - TOolbox for Traffic Engineering Methods. `http://totem.run.montefiore.ulg.ac.be` (2005)
24. Feldmann, A., Greenberg, A., Lund, C., Reingold, N., Rexford, J.: NetScope: Traffic Engineering for IP Networks. IEEE Network Magazine (2000) 11–19
25. 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)