

Effectiveness of Link Cost Optimization for IP Rerouting and IP Fast Reroute

David Hock, Matthias Hartmann, Christian Schwartz, and Michael Menth

University of Würzburg, Institute of Computer Science
Am Hubland, D-97074 Würzburg, Germany

Abstract. In this paper, we bring together resilience analysis and routing optimization for IP-based intra-domain networks. When link, node, or multiple failures occur, traffic is rerouted which increases the link load on backup paths and possibly causes congestion. Resilience analysis detects the risk of overload situations a priori based on a large set of most likely failure scenarios. To counteract, the routing can be optimized and configured that such bottlenecks are avoided at least for a smaller set of failure scenarios. In this paper, we demonstrate the effectiveness of this routing optimization in IP networks. We use resilience analysis with suitable aggregate views on relative link loads. Furthermore, we compare conventional IP rerouting with IP fast reroute (IP-FRR) and show that IP-FRR can also significantly profit from routing optimization. This paper reviews major parts of previous publications and presents a new method to visualize and compare the resilience of different routing schemes.

1 Introduction

Outages in communication networks like link and node failures are a matter of fact and cannot be avoided. However, the network can be prepared for such conditions by using self-healing routing mechanisms. When elements on the primary path fail, traffic is rerouted to a backup path. This mechanism alone just assures the connectivity of the network provided that such a backup path exists and can be activated by the protection mechanism. There is another aspect: capacity. Rerouted traffic causes increased load on backup paths so that overload and traffic loss possibly occur. This can be avoided by carefully choosing the layout of primary and backup paths.

In this work, we bring together three issues that have recently attracted attention in the area of fault-tolerant networking. Resilience analysis is an efficient means to quantify the risk of overload in networks due to failures. Optimization of resilient IP routing improves the load conditions in IP networks at least for a small set of likely failure scenarios. Recently developed IP fast reroute (IP-FRR) mechanisms quickly switch traffic to preconfigured backup paths instead of running into transient forwarding loops during the IP rerouting process. We use resilience analysis to demonstrate the effectiveness of

This work is funded by Deutsche Forschungsgemeinschaft (DFG) under grant TR257/23-2. The authors alone are responsible for the content of the paper.

routing optimization in IP networks. We compare the likelihood of overload for unoptimized conventional IP rerouting and for IP-FRR. Finally, we illustrate the impact of routing optimization also for IP-FRR.

The remainder of this work is structured as follows. In Section 2 we explain the fundamentals of IP routing and introduce IP fast reroute. In Section 3, we give an overview of resilience analysis and link cost optimization. In Section 4 we study the effectiveness of routing optimization for IP rerouting and IP fast reroute. Section 5 concludes this work.

2 Fundamentals of IP Routing

We explain IP routing which follows the principle of least-cost (shortest) paths. We show how ambiguities arising from several least-cost paths can be handled. Finally, we review mechanisms for IP-FRR.

2.1 Conventional IP Routing and Reconvergence

In intra-domain IP networks, routers exchange information about the topology and administrative link costs with each other. Based on these routing messages, each node obtains a full view of the link topology including administrative link costs. It uses this information to set up the routing table whereby it associates any destination in the network with the interface leading towards a least-cost path to the destination. Thus, the routing table helps to look up onto which outgoing interface packets destined to a certain node in the network should be forwarded.

In case of a modification of the topology, e.g., due to a link or router failure, a reconvergence process is invoked. The change is broadcast through the entire local network and routers recalculate the outgoing interface mapping in their routing tables based on the new topology. As long as the network is physically connected, IP routing finds new routes for all source-destination pairs. This makes it very robust against network failures.

2.2 Handling Ambiguities due to Several Least-Cost Paths

Depending on the link cost settings, possibly several least-cost paths exist between pairs of nodes in a network. In that case the routing is undefined at first step. However, routers use tie-breakers to decide which of the paths to prefer for routing. E.g., the interface towards a least-cost path with the smallest port number may be chosen [1, Sect. 7.2.7]. However, port numbers within routers are not necessarily predictable. Therefore, it is hard or even impossible to predict the route in case of several least-cost paths a priori. In previous work [2], we quantified that optimized routing can lead to significantly larger relative link loads than expected if traffic is forwarded on other least-cost paths than assumed. Hence, predictable load distribution is important for routing optimization, network planning, and traffic engineering in general.

One solution to that problem is equal-cost multipath (ECMP) routing. It splits the traffic equally among all interfaces towards a least-cost path. As packet-by-packet load

balancing possibly causes packet reordering, load-balancing is done on the flow level. To that end, hash-based load balancing is used, i.e., typical data of a flow like source and destination IP and port numbers are hashed to some value based on which the packet is forwarded to one of the potential interfaces.

Finally, it is possible to choose link costs such that several least-cost paths do not exist. In [2] we implemented that objective as part of IP routing optimization and showed that so-called unique shortest paths (USP) can be efficiently obtained.

2.3 IP Fast Reroute (IP-FRR)

The reconvergence process in IP networks can take up to several minutes. During this time, forwarding loops can appear when some of the routers have updated their routing tables earlier than others. As a consequence, the affected traffic cannot be delivered to its destination, looping the traffic causes high load on the respective links which causes additional overload. To avoid this phenomenon, IP-FRR has been proposed. Routers detecting a failure immediately switch the affected traffic to preestablished backup paths that are likely to be unaffected by the observed failure. There are multiple proposals for the implementation of IP-FRR [3].

With Loop-Free Alternates (LFAs) [4], routers store alternative next-hops in their routing tables which are used when the primary next-hop fails. However, it is not always possible to find neighbor hops that do not loop back the traffic or create routing loops when more than a single link has failed. Therefore, LFAs cannot always provide 100% failure coverage.

A promising alternative are not-via addresses which are currently being standardized in the IETF [5,6]. For any node N there is a not-via address N_F and packets addressed to N_F are forwarded to N while node F is avoided on the path. Hence, the routing tables in the network require additional entries for these not-via addresses. They are used for IP-FRR as follows. We assume that a node A receives a packet that is normally forwarded over F and the next-next-hop N to its destination, but the next-hop F has failed. Then the node A encapsulates this packet towards the not-via address N_F to tunnel it to N . N decapsulates the packet and forwards it to the destination. If the next-hop F is already the destination, the packet can be delivered if only the link from A to F is down but not F itself. Then, A encapsulates the packet to F_A and forwards it to some of its neighbor nodes so that the packet is carried towards F avoiding the link from A to F . Hence, the not-via mechanism leads the traffic on the shortest path according to administrative link costs around the next-hop to the next-next-hop or around the next-link to the next-hop if the next-hop is the destination node. If due to an additional network failure, traffic encapsulated with a not-via address is tunneled again, this can lead to traffic loops in the network. To avoid this problem, already not-via encapsulated traffic must not be tunneled to not-via addresses again but be dropped instead. In [2], we have argued that IP-FRR needs USP to create a predictable backup path layout. We have also shown that such IP link costs can be efficiently found while optimizing the path layout for IP-FRR.

3 Resilience Analysis and IP Link Cost Optimization

In the following, we review resilience analysis and IP link cost optimization.

3.1 Resilience Analysis

Link and router failures may lead to disconnection of nodes within a network and to rerouted traffic causing increased load on backup paths. The resilience analysis in [7] quantifies the disconnection probability of nodes due to failures and the potential overload caused by backup traffic or abnormal traffic demands.

The resilience analysis requires the network topology, the routing and rerouting model, the link capacities, an availability model for network elements indicating failure probabilities as well as a model of the traffic matrix indicating the probability and the structure of abnormal traffic demands. We define networking scenarios $z = (s, h)$ consisting of a failure scenario s and a traffic matrix h . Failure scenarios and traffic matrices are associated with probabilities $p(s)$ and $p(h)$. We assume independence so that the probability of a networking scenario can be calculated by $p(z) = p(s) \cdot p(h)$. The idea of the analysis is to investigate the disconnection of nodes and relative link loads for individual networking scenarios z and these results contribute with a probability weight of $p(z)$ to the final result. Due to computational limitations, it is not possible to consider all possible failure scenarios and traffic matrices. Therefore, the analysis considers only networking scenarios with a probability of at least p_{min} and this set is denoted by \mathcal{Z} . The final results of the analysis are probabilities for the disconnection of a given node pair due to failures and complementary cumulative distribution functions (CCDFs) of the relative load for each link in the network. Both the disconnection probabilities and the CCDF of the relative link load values are conditional in the sense that they refer only to the set of investigated scenarios \mathcal{Z} , but upper and lower bounds on the true value are given. In the following we omit this aspect for the sake of simplicity. In this paper, we consider only network element failures as source for increased traffic on links and use only a single standard matrix without anomalies.

Several aggregated views have been developed in [7] to visualize unavailability. CCDFs of relative link loads are displayed per link. However, it is desirable to have a visualization of potential overload in the entire network at a glance. To that end, the information of the CCDF of the relative link loads can be condensed into a single overload value by various mapping functions. These values can be used to color links in a topological representation of the network.

There are several possible applications of resilience analysis. Using this technique, operators can, e.g., check if the network's current state is sufficient to allow additional clients, to sell better Service Level Agreements, or to deal with the traffic increase arising in the next few months. If this is not the case, the resilience analysis can help to decide where to add new links or routers. Furthermore, resilience analysis can be used to study the influence of a new routing or to investigate the effectiveness of routing optimization on potential overload. The latter application is the one addressed in this publication.

Further details to our framework for resilience analysis together with an overview on related work in this area including examples of resilience analysis, can be found in

our previous publication [7]. Our framework has been implemented as a software tool. It is presented in [8].

3.2 IP Link Cost Optimization

IP routing follows the least-cost paths according to administrative link costs. Traffic engineering is possible by appropriately choosing those link costs that lead to a good load distribution on the links. An objective function defines what is understood by a good load distribution and is later discussed in more detail. Searching for good IP link costs can be automated which is called link cost optimization, sometimes also referred to as link weight optimization.

The input are a network topology, link capacities, a traffic matrix, and a given set of so-called protected failure scenarios \mathcal{S} for which the routing should be optimized. The output of the process are administrative costs for all links in the network. The set \mathcal{S} usually comprises all single link and/or node failures ($\mathcal{S}_L, \mathcal{S}_R, \mathcal{S}_{RL}$). The failure-free state $s = \emptyset$ is always part of this set. For computational reasons, the set of protected failure scenarios \mathcal{S} is usually smaller than the set of considered networking scenarios \mathcal{Z} that is used as a base for resilience analysis.

Finding optimum IP link costs for a given objective function is usually an NP-hard problem even when only considering the failure-free case \mathcal{S}_\emptyset . Therefore, heuristic algorithms are used to search good link costs. An overview of related work including different objective functions and heuristic approaches can be found in [2, 9]. The heuristic we apply for this work is described in [9, 10]. It is similar to the threshold accepting heuristic proposed in [11]. We perform multiple optimization runs with our heuristic and take the result of the best run as final result.

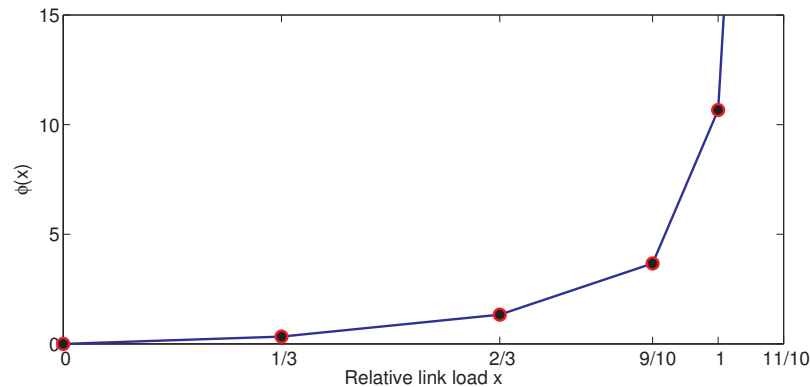


Fig. 1. Fortz's utilization-dependent penalty function ϕ .

In [9] we have studied different objective functions for resilient and non-resilient IP routing which can be used for different application scenarios. Two of them are explained

in more detail here. Both take the relative link load as a parameter. The relative load $\rho(l)$ of a link l is calculated as the quotient of the total traffic on a link and the link's capacity. To illustrate the severeness of possible overload, relative link loads larger than 100% are allowed in the computation.

- U_S^{max} is the maximum relative link load of all links in all protected failure scenarios \mathcal{S} . It is a good choice, if routing optimization is used to guarantee that certain constraints on the relative link load are kept.
- $F_S^{weighted}$ sums up penalties over all links and all protected failure scenarios whereby these penalties increase with increasing relative link load. The penalties are calculated with Fortz's continuous, piecewise linear, monotonically increasing penalty function ϕ [12], which is illustrated in Figure 1. The objective function $F_S^{weighted}$ is good if the main focus of the optimization lies on the overall link loads and the average path lengths.

Different objective functions lead to significantly different optimization results. To visualize that, we consider routing based on the hop-count metric and optimized routing based on objective function U_S^{max} and $F_S^{weighted}$ whereby \mathcal{S} comprises all single link failures.




Fig. 2. CCDF of the maximum relative link load over all single link failure scenarios (COST239 network).

Figure 2 shows the maximum relative link load of all links in the COST239 network [9] in all protected failure scenarios \mathcal{S}_L . The x-axis indicates the relative link load and the y-axis the fraction of links whose maximum relative link load exceeds the value on the x-axis. Hop-count (HC) routing leads to the highest relative link loads, optimized routing based on U_S^{max} leads to the lowest maximum relative link loads. Objective function $F_S^{weighted}$ achieves a compromise. The drawback of U_S^{max} is that it cannot improve the second-worst link when the worst link cannot be improved further. Therefore, we proposed in [9] to combine both objective functions, i.e., we first minimize U_S^{max} and

then $F_S^{weighted}$. This leads to the lowest maximum relative link loads and reduces also the load on other highly loaded links. This is the objective function we use also in this study. Additional constraints can be used, e.g., in [2], we accepted only link cost settings where several least-cost paths are avoided. This is a valuable feature for traffic engineering when ECMP is not used or also for IP-FRR based on not-via addresses. The optimization of IP-FRR has been developed in that work, too.

4 Results

In the following, we study the effectiveness of routing optimization for IP routing and IP-FRR. Therefore, we first analyze unoptimized hop-count routing and then compare it to optimized USP routing. We show that even the link cost optimization with a small set of protected failure scenarios \mathcal{S}_L leads to routings that significantly improve the overall resilience of the network. In a second step, we investigate the difference between unoptimized and optimized routing using not-via IP-FRR techniques.

4.1 Networks under Study

We have run our experiments for different networks including the Rocketfuel topologies [13]. All topologies yield similar results. Here, we show only the results of the Exodus network. The geographical topology of this network is depicted in Figure 3. It is not suitable to add link or node related information, because some nodes are so close to each other that they cannot be differentiated and links overlap. Therefore, we propose another representation of the same topology in Figure 5(a), that will be explained later.




Fig. 3. Exodus network, 22 nodes, 51 links.

The used traffic matrix (TM) has been created resembling real-world data according to the method proposed in [14] and enhanced in [15]. All links were expected to have

identical capacity and the TM was scaled so that the worst relative load experienced by a link in case of single link failures and hop-count routing is 75%. However, relative link loads larger than 100% can be achieved in single node and multiple failure scenarios.

Based on [7], we chose an unavailability of 10^{-6} for all nodes. Each link is unavailable with the same probability of 10^{-4} . The set of investigated scenarios \mathcal{Z} has been calculated for $p_{min} = 10^{-15}$. This results in a number of $|\mathcal{Z}| = 51577$ considered scenarios, about a thousand times more, than the number of single link failures considered for the link cost optimization $|\mathcal{S}_L| = 52$. \mathcal{Z} consists of the failure patterns $\emptyset, L, R, LL, LR, RR, LLL, LLR$, where L denotes a single link and R a single router failure. This way, a resilience analysis with \mathcal{Z} reaches very high precision, while still being computationally feasible. On a "Intel(R) Core(TM)2 Duo CPU E8500 @ 3.16GHz" a resilience analysis of a single routing in the Exodus network with \mathcal{Z} using our software tool [8] took about 300 seconds. The link cost optimization to obtain the best USP routing solution used in this paper took about 66 hours and involved a total number of 18,654,149 routing evaluations with \mathcal{S}_L , the best not-via solution was obtained in about 205 hours and 21,816,259 routing evaluations. However good optimization results can already be achieved after some minutes of optimization¹.

4.2 IP Routing and Rerouting Based on the Hop-Count Metric

In the following, we analyze the potential overload in a network when hop-count routing is used. We investigate the relative load for the link from Palo Alto to Santa Clara because its potential overload is especially high in some failure scenarios. Figure 4 shows the CCDF of the relative link load $\rho(l)$ for this link. The CCDF illustration simplifies the observation of the potential overload for a single link. The probability $P(\rho(l) > x)$ that a relative link load $\rho(l)$ exceeds a certain value x is directly displayed in the graph. In this case, e.g., the probability that relative link loads higher than 60% occur from Palo Alto to Santa Clara is about 0.06% $P(\rho(l) > 0.6) \approx 0.06\%$. This value is later referred to as $R_r^{0.6}$. On the other hand, in at least 99.999999% of all scenarios the relative link load is not larger than about 116%, $P(\rho(l) \leq 116\%) > 99.999999\%$. This value is later referred to as $R_q^{0.99999999}$. In particular, this is true for all single and double link failures as well as single node failures.

If CCDFs are used, a complete figure is necessary to visualize the probabilistic load condition on a link. Monitoring such information for all links in the network becomes more difficult with an increasing network size. Therefore, in [7] we presented various mapping functions to aggregate the information of the per link CCDF into one per link value. Two of those functions are used in this work.

- Mapping function $R_r^x(l) = P(\rho(l) > x)$ is based on overload probabilities. It returns the probability with which the relative load $\rho(l)$ of link l exceeds the relative load value x . Figure 4 illustrates $R_r^{0.6}$.
- Mapping function $R_q^y(l) = \inf(x : P(\rho(l) \leq x) \geq y)$ is based on relative link load quantiles. This mapping function returns the smallest relative link load value x

¹ The routing optimization was parallelized on several CPUs so that the effective computation time could be significantly reduced.

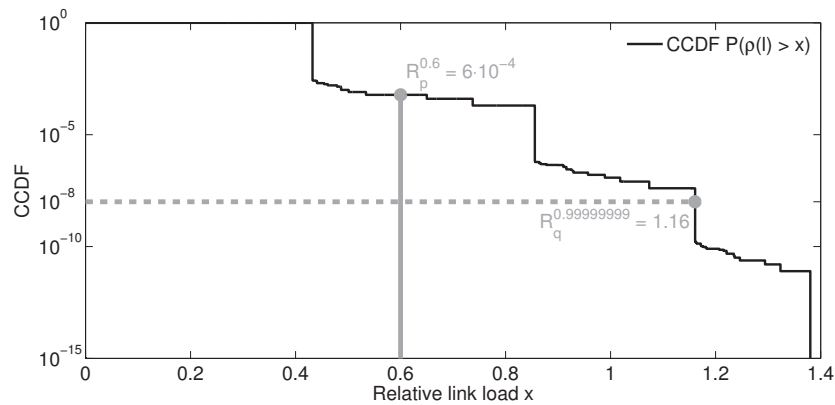


Fig. 4. CCDF of the relative link load $\rho(l)$ for the link between Palo Alto and Santa Clara.

which is not exceeded by a fraction of at least y of all considered network scenarios. Figure 4 depicts $R_q^{0.99999999}$.

We use the mapping functions to convert the CCDF of each link to a single value. Then, we map those values to a color scale indicating the severeness of the potential overload.


The geographical view in Figure 3 is not suitable to add link or node related information. Some nodes are so close to each other that they cannot be differentiated. Forward and backward directions of links cannot be distinguished, either. Therefore, we propose an adjacency matrix to represent the network topology as in Figure 5. The cell of row i column j in the adjacency matrix corresponds to the link between nodes i and j .

Figure 5 shows the adjacency matrix of the Exodus network colored according to the quantile based mapping function $R_q^{0.99999999}$ for unoptimized hop-count routing and optimized USP routing. This illustration shows the potential overload of the whole network and the link with the risk of highest overload can be directly recognized. The colors in the tiles can be converted to numerical relative load values using the color bar on the right side of the graph.


4.3 Optimized IP Routing and Rerouting

In the following, we show the impact of routing optimization on the potential overload.

Figure 6(a) shows the CCDF of the relative load on the link from Palo Alto to Santa Clara for hop-count routing and optimized USP routing. The curve of the optimized USP routing is at all values smaller than the one for hop-count routing. Thus, the routing optimization indeed reduces the overload risk on this particular link. As a consequence, all mapping functions yield smaller values for optimized USP routing than for hop-count routing. This findings hold only for this particular link which was the worst for hop-count routing. The link, depicted in Figure 6(b), between Santa Clara and Miami



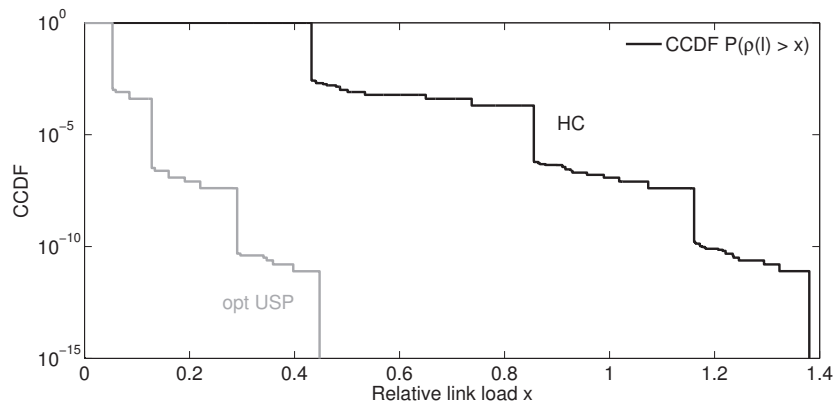
(a) Hop-count routing.



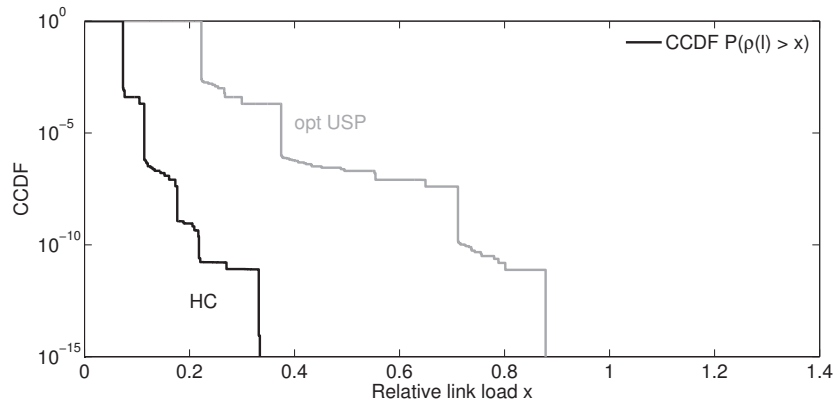
(b) Optimized USP routing.

Fig. 5. Adjacency matrix of the Exodus network colored according to the potential overload risk for different routings. The color of a link corresponds to the 99.99999% quantile of its CCDF of the relative link load. Darker colors indicate higher overload values.

presents an interesting counter example. Here, the risk of overload is larger after routing optimization. An optimized path layout does not decrease the total amount of traffic in



(a) Link from Palo Alto to Santa Clara.



(b) Link from Santa Clara to Miami.

Fig. 6. CCDF of the relative link load $\rho(l)$ for hop-count routing and optimized USP routing.

the network but just distributes it differently over the links. However, Figure 6(b) shows that the resulting load increase on some links does not cause any real problems because the relative link loads still remain relatively low.

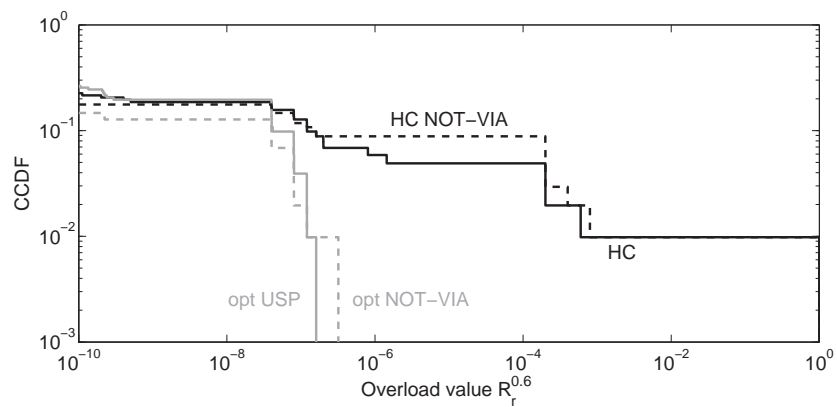
To visualize the impact of routing optimization on the potential overload, we need to take all links of the network into account. Therefore, we calculate the overload values according to any mapping function R_p^x or R_q^y based on the CCDFs for all links. Then, we specify the fraction of links, whose potential overload exceeds a certain value. This leads to a CCDF of the overload values of the chosen function R_p^x or R_q^y .

Figure 7 shows CCDFs of overload values according to both mapping functions for hop-count routing and optimized USP routing as solid lines. Routing optimization redistributes the traffic in the network. On the one hand, this leads to a reduction of the worst overload values in the network. On the other hand, on some links with lower potential overload the values lightly increase. This effect is clearly visible in both graphs.

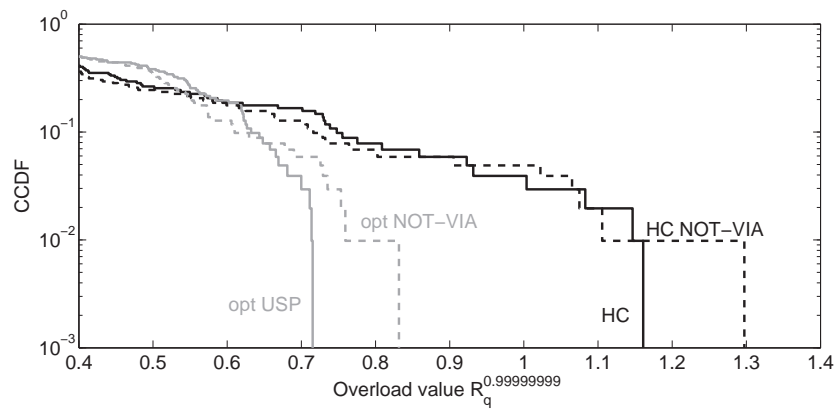
It is an interesting finding, that this result holds for both mapping functions. This shows that the link cost optimization on a small set of protected failure scenarios \mathcal{S}_L is very effective because it significantly improves the resilience calculated on a large set of scenarios \mathcal{Z} .

4.4 IP Fast Reroute Method Not-Via

We investigate not-via IP-FRR based on hop-count routing and based on optimized USP routing in comparison to conventional IP rerouting.



(a) CCDF over all links of the probability that a relative link load exceeds 60%.



(b) CCDF over all links of the 99.999999% quantile of the relative link load.

Fig. 7. Comparison of the CCDFs of the potential overload for IP rerouting and not-via IP-FRR.

We compare the overload values of the entire network for hop-count routing and optimized USP routing to unoptimized and optimized not-via IP-FRR. Figure 7 displays

the overload values of not-via IP-FRR in dashed lines. The potential overload in case of unoptimized not-via FRR is even higher than for conventional IP hop-count routing. Routing optimization significantly improves these values. Optimized not-via IP-FRR reaches overload values of similar quality as optimized USP routing. This holds for both mapping functions $R_p^{0.6}$ and $R_q^{0.99999999}$. The overload values caused by not-via IP-FRR are higher than for conventional IP routing especially due to the increased load on backup paths and the longer average path lengths due to local repair. However, routing optimization can reduce the risk of overload to a secure level also for not-via IP-FRR.

We have shown that not-via IP-FRR based on hop-count routing leads to even higher potential overload than conventional hop-count routing. Therefore, routing optimization is even more beneficial for not-via IP-FRR.

5 Conclusion

Resilience analysis evaluates the load conditions in communication networks for a large set of likely failure scenarios \mathcal{Z} whose probabilities are at least p_{min} . Routing optimization is usually applied to improve load conditions only for a set of most likely failure scenarios \mathcal{S} which is up to a thousand times smaller than \mathcal{Z} . Despite of this big difference in size of the considered failure sets, we have shown that routing optimization significantly reduces potential overload in networks with conventional IP routing and rerouting. Furthermore, we illustrated that without routing optimization IP fast reroute (IP-FRR) possibly causes even more overload than conventional routing and rerouting. However, routing optimization is again very effective for IP-FRR in avoiding potential bottleneck situations and thus even more beneficial for this case. Moreover, it is needed for IP-FRR anyway because the link cost values should be chosen in such a way that equal-cost paths are avoided in order to obtain unambiguous backup paths.

Acknowledgments

The authors thank David Stezenbach for his programming efforts and Prof. Phuoc Tran-Gia for the stimulating environment which was a prerequisite for this work.

References

1. ISO: ISO 10589: Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (1992/2002)
2. Hock, D., Hartmann, M., Menth, M., Schwartz, C.: Optimizing Unique Shortest Paths for Resilient Routing and Fast Reroute in IP-Based Networks. In: IEEE Network Operations and Management Symposium (NOMS), Osaka, Japan (2010)
3. Martin, R., Menth, M., Hartmann, M., Cicic, T., Kvalbein, A.: Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute? accepted for Computer Networks (2010)
4. Atlas, A., Zinin, A.: RFC5286: Basic Specification for IP Fast Reroute: Loop-Free Alternates (2008)

5. Shand, M., Bryant, S.: IP Fast Reroute Framework. <http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-12.txt> (2009)
6. Bryant, S., Previdi, S., Shand, M.: IP Fast Reroute Using Not-via Addresses. <http://tools.ietf.org/id/draft-ietf-rtgwg-ipfrr-notvia-addresses-04.txt> (2009)
7. Menth, M., Duelli, M., Martin, R., Milbrandt, J.: Resilience Analysis of Packet-Switched Communication Networks. accepted for IEEE/ACM Transactions on Networking (2010)
8. Hock, D., Menth, M., Hartmann, M., Schwartz, C., Stezenbach, D.: ResiLyzer: A Tool for Resilience Analysis in Packet-Switched Communication Networks. In: GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems (MMB) and Dependability and Fault Tolerance (DFT), Essen, Germany (2010)
9. Hartmann, M., Hock, D., Schwartz, C., Menth, M.: Objective Functions for Optimization for Resilient and Non-Resilient IP Routing. In: 7th International Workshop on Design of Reliable Communication Networks (DRCN), Washington, D.C., USA (2009)
10. Menth, M., Hartmann, M., Martin, R.: Robust IP Link Costs for Multilayer Resilience. In: IFIP-TC6 Networking Conference (Networking), Atlanta, GA, USA (2007)
11. Dueck, G., Scheuer, T.: Threshold Accepting: a General Purpose Optimization Algorithm. *Journal of Computational Physics* **90** (1990) 161–175
12. Fortz, B., Thorup, M.: Internet Traffic Engineering by Optimizing OSPF Weights. In: IEEE Infocom, Tel-Aviv, Israel (2000) 519–528
13. Spring, N., Mahajan, R., Wetherall, D.: Measuring ISP Topologies with Rocketfuel. In: ACM SIGCOMM, Pittsburgh, PA (2002)
14. Nucci, A., Sridharan, A., Taft, N.: The Problem of Synthetically Generating IP Traffic Matrices: Initial Recommendations. *ACM SIGCOMM Computer Communications Review* **35** (2005) 19–32
15. Roughan, M.: Simplifying the Synthesis of Internet Traffic Matrices. *ACM SIGCOMM Computer Communications Review* **35** (2005) 93 – 96