

Optimizing Unique Shortest Paths for Resilient Routing and Fast Reroute in IP-Based Networks

David Hock, Matthias Hartmann, Michael Menth, Christian Schwartz
 University of Würzburg, Institute of Computer Science
 Am Hubland, D-97074 Würzburg, Germany
 Email: {hock,hartmann,menth,schwartz}@informatik.uni-wuerzburg.de

Abstract—Intradomain routing in IP networks follows shortest paths according to administrative link costs. When several equal-cost shortest paths exist, routers that use equal-cost multipath (ECMP) distribute the traffic over all of them. To produce single-shortest path (SSP) routing, a selection mechanism (tie-breaker) chooses just one of the equal-cost paths. Tie-breakers are poorly standardized and use information that may change over time, which makes SSP routing unpredictable. Therefore, link costs producing unique shortest paths (USP) are preferred.

In this paper, we show that optimized SSP routing can lead to significantly higher link utilization than expected in case of non-deterministic tie-breakers. We investigate the impact of the allowed link cost range on the general availability of USP routing. We use a heuristic algorithm to generate link costs for USP routing and to minimize the maximum link utilization in networks with and without failures.

Fast reroute (FRR) mechanisms can repair failures faster than conventional IP rerouting by pre-computing shortest backup paths around failed network elements. However, when multiple equal-cost paths exist, the backup path layout is unpredictable. We adapt our heuristic to optimize USP routing for IP-FRR using not-via addresses and MPLS-FRR with facility and one-to-one backup. Finally, we compare the performance of USP with various other routing schemes using realistic Rocketfuel topologies.

Keywords: IP routing optimization, resilience, IP and MPLS fast reroute

I. INTRODUCTION

The routing in today's intradomain routing protocols (OSPF [1], IS-IS [2]) is determined by administrative link costs. Routers forward packets along cost-minimal paths towards their destinations. Often, there exist several shortest paths which have equal costs. Routers that use the equal-cost multipath (ECMP) option distribute traffic over all available shortest paths. When single shortest path (SSP) routing is required, each router uses a tie-breaker to select just one of the equal-cost shortest paths. However, tie-breakers are not properly standardized and might even use non-deterministic information like, e.g., router-internal interface numbers whose order can change after restart or the current link load, or even select a random next hop. This makes the paths of general SSP routing hard to predict. Traffic engineering techniques like routing

carry which traffic. Therefore, link cost settings are preferred that lead to a unique shortest path (USP) routing without equal-cost paths so that no tie-breakers are required.

When links or routers in an IP network fail, the information about the topology change is broadcast through the network and triggers the recalculation of the forwarding tables in all routers. This IP reconvergence ensures that routers can again reach each other as long as the topology remains physically connected. But it is slow and can take up to several seconds. To minimize traffic loss and allow quick local reaction to outages, fast reroute (FRR) mechanisms have been developed. Not-via is such a mechanism for pure IP networks. It precalculates backup paths along shortest paths around a failed node to the corresponding next-next-hop. For MPLS, there are two different FRR options. Facility backup uses local bypasses for backup traffic in a similar way as not-via. One-to-one backup redirects traffic directly to its destination. If the paths are established along shortest IP paths, primary and backup LSPs can be automatically set up and do not need to be configured with explicit paths. To predict the impact of backup traffic for resource management, the backup paths must be unique in order to be predictable. Therefore, the underlying IP routing should produce only unique backup paths.

One objective of traffic engineering is to minimize the maximum link utilization over all links in the network. When resilient routing is required, this also includes the link utilizations that occur in a certain set of protected failure scenarios. The optimization is performed by adjusting the administrative link costs of all links. In this paper, we show that optimized SSP routing can lead to significantly higher load than expected when tie-breakers work differently than assumed. This cannot happen when the routing produces USP in all protected scenarios. However, this puts more constraints on the link cost settings. We use a heuristic algorithm to find link cost settings that result in USP routing and optimize them in order to minimize the maximum link utilization. We illustrate that the fraction of available USP routings as well as the quality of the optimized routings depend on the maximum allowed link costs. We adapt our heuristic to produce optimized USP link cost settings for the FRR mechanisms mentioned above. Finally, we compare the performance of optimized ECMP, SSP, USP routing in failure-free IP networks as well as for IP reconvergence and various FRR mechanisms for a limited set of protected failure cases.

Funded by Deutsche Forschungsgemeinschaft (DFG) under grant TR257/23-2. The authors alone are responsible for the content of the paper. optimization or admission control need to know which links

The paper is structured as follows. Section II motivates the need for USP in IP networks. In Section III we quantify the impact of wrong tie-breaker assumptions, describe our USP routing optimization, and discuss some examples. Section IV extends the routing optimization towards FRR mechanisms and compares the performance of various routing and resilience mechanisms. In Section V we review related work and explain how it differs from our work. Finally, Section VI concludes this work.

II. THE NEED FOR USP IN IP NETWORKS

We discuss some scenarios where USP routing is desired.

A. Classic IP Networks

Shortest path routing is unpredictable when equal-cost paths exist. Using ECMP traffic is equally distributed over all shortest paths. This load balancing is done on a per-flow basis so that an equal distribution of the traffic is subject to statistical variations [3], [4]. Also the prospective path of a new flow is unpredictable which leads to problems with path-dependent flow control functions like admission control and flow termination [5]. SSP does not have these issues because only a single path for each source-destination pair is used. However, the use of non-deterministic tie-breaker decisions leads to an unpredictable path layout. This can be avoided by using USP.

B. IP Fast Reroute

When a network element in a classic IP network fails, the outage is signaled to all routers and each of them recalculates its routing table. This can take a significant amount of time, during which packets are possibly lost. Fast reroute (FRR) mechanisms provide faster reaction through local pre-computed backup paths. The most promising IP FRR mechanism is not-via, which is currently being standardized in the IETF [6], [7]. *Not-via addresses* are used to reach the next-next hop towards the destination not via the failed router or link. This way, each router is not only reachable by its regular address but also with several not-via addresses that indicate the failures of its direct neighbors. The routes to these not-via addresses are pre-computed by all routers in a distributed way based on the normal IP link costs but without using the indicated not-via-router. Then, they are stored in the forwarding tables of the routers. A comprehensive description of the mechanism can be found in [8].

When SSP routing is used and several equal-cost shortest paths around a failed next-hop exist, the paths of the not-via bypasses are hard to predict. Unlike in classic IP networks, this problem cannot be solved by ECMP. Since all rerouted packets are encapsulated, they appear to come from the failure-detecting node and are addressed to the next-next hop after the failure. Therefore, hash-based load balancing algorithms treat them like a single flow, i.e., they are all carried over a single backup path. Hence, there is no way to balance the backup traffic over equal-cost paths and the chosen backup path is also hard to predict if the details of the load balancing algorithms are not known. Thus, for IP FRR using not-via, USP routing is even more desired than in classic IP networks.

C. MPLS Fast Reroute

In MPLS networks, label switched paths (LSPs) can be set up either using pre-computed, explicit paths or along the shortest paths in an underlying IP network. In this paper, we focus on the second option, which is a wide-spread practice due to its simplicity. To set up a LSP, signaling packets are forwarded from the ingress router to the egress router of the LSP and, thereby, determine its path layout. When equal-cost paths exist, it is almost impossible to predict the exact path layout. As this holds for both SSP and ECMP routing, USP routing is a desired feature also in IP-based MPLS networks. MPLS comes with two fast reroute options. We briefly describe the structure of their backup paths. A more illustrative explanation is available in [9].

1) *Facility Backup*: The facility backup option bypasses traffic around a failed network element. These bypasses are established from the point of local repair (PLR) where the failure is detected along a shortest path around the failure towards the merge point (MP) where they rejoin the primary LSP. When the MP is the next-next-hop after the failure, not-via and facility backup use the same backup path layout.

2) *One-to-One Backup*: Instead of a single bypass for all LSPs, one-to-one backup creates an individual backup path for each flow. These paths are established along shortest paths from the PLR directly to the egress router of the LSP, which leads to a different path layout than facility backup.

III. ROUTING OPTIMIZATION FOR USP

One main traffic engineering objective is the minimization of the maximum link utilization in the network. To that end, the path layout needs to be modified. We introduce our optimization method and quantify for optimized SSP routing the impact of unknown tie-breakers on the link utilization. We show that the maximum link costs must be large enough so that link cost settings with USP routing can be efficiently found and effectively optimized. Finally, we determine optimal link cost ranges also for the optimization of other routing mechanisms.

A. Heuristic Optimizer

We use a heuristic optimizer from previous work [10], [11], which is only briefly described here. It works on a directed graph $G = (V, E)$ with routers V and links E , and minimizes a given optimization objective Ψ . Every link $e \in E$ is assigned an administrative cost value $k \in [1 : k_{\max}]$ and, thus, the search space consists of $(k_{\max})^{|E|}$ different link cost settings. The optimizer implements a threshold accepting heuristic. It starts with a random link cost configuration \mathbf{k} , and uses neighborhoods where up to 25% of the link weights are randomly changed. When a new neighbor link cost configuration is better or not worse than a threshold above the current best value, it is accepted as new current configuration and the search continues. If no strictly better solution is found after a previously configured number of iterations, the currently best value is returned as the final optimization value. The

heuristic can be restarted several times with different random start configurations \mathbf{k} and the best link cost setting of all runs is returned as final result.

As objective function, we use the maximum link utilization $\Psi(\mathbf{k}) = \max_{s \in \mathcal{S}, l \in \mathcal{E}} (\rho(l, s, \mathbf{k}))$, which is the maximum utilization $\rho(l, s, \mathbf{k})$ over all links l and all considered failure scenarios $s \in \mathcal{S}$, when link costs configuration \mathbf{k} is used in the network. Routing optimization can be done for failure-free routing only (\mathcal{S}_0) or also for a limited set of failures \mathcal{S} , as for instance all single link failures (\mathcal{S}_L).

For our experiments, we use the Rocketfuel topologies [12] in addition to the well known Cost239 [13] and Geant [14] networks. We remove all nodes that are connected by only a single link because they are never used to forward other traffic and, thus, are irrelevant for traffic engineering. The sizes and the node degrees of the networks are compiled in Table I. To generate synthetic traffic matrices resembling real-world data we use the method proposed in [15] and enhanced in [16]. The traffic matrices are scaled so that the maximum link utilization over all single link failures equals 100% for ECMP routing based on the hop-count metric, i.e. all link costs are set to 1. For our comparisons, we allow theoretical link utilizations above 100% without packet drops.

TABLE I
NETWORKS UNDER STUDY

ID	Name	$ V $	$ E $	$AvgDeg$
AB	Abovenet	20	156	7.8
AT	AT&T	28	120	4.28
CO	Cost239	11	52	4.73
EB	Ebone	25	126	5.04
EX	Exodus	22	102	4.64
GE	Geant	19	60	3.16
SP	Sprintlink	33	190	5.78
TI	Tiscali	38	232	6.11

B. Impact of Non-Deterministic Tie-Breakers

Link cost optimization is an offline process, i.e., an external tool takes a representation of the network and its traffic matrix as input and returns an optimized link cost setting. These link costs are then configured in the real network. For SSP routing, the tools have to assume certain tie-breakers though the routers in the network might use different ones. When link costs are optimized for wrong tie-breakers, routers possibly send traffic to already highly loaded links. This has a devastating effect on the link utilization. We quantify this problem in the following.

First, we optimize SSP routing and use the lowest next-hop port number as tie-breaker. We obtain an optimized link cost setting \mathbf{k}_0 that results in a maximum link utilization Ψ_0 . We apply this link cost setting \mathbf{k}_0 to the network. Then we randomly jumble the port-numbers in each router which leads to different tie-breaker decisions and calculate Ψ_1 for the new routing. We repeat this experiment 100 times and record the worst values $\Psi_{\text{worst}} = \max_{(i=1..100)} (\Psi_i)$ for each network. The experiment is conducted both under failure-free conditions and including single link failures. The results for each network are present in Figure 1, which shows the possible

increase in maximum link utilization ($\frac{\Psi_{\text{worst}}}{\Psi_0} - 1$) in percent. Different tie-breakers can significantly increase the maximum link utilization both under failure-free conditions and in case of single link failures. For example in the Tiscali network with failure free routing, the link utilization is over 180% higher. The presented results can differ for other initial (optimized) link cost settings \mathbf{k}_0 , and many tie-breaker changes that affect only slightly loaded paths have no effect at all on the maximum link utilization. Therefore, these results cannot be generalized, but they show that traffic engineering is useless when relying on unknown tie-breakers.

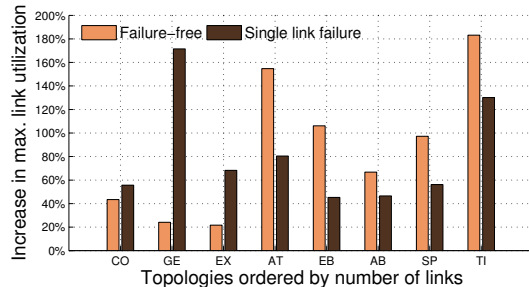
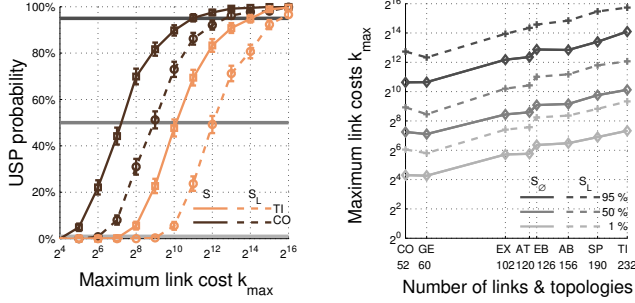


Fig. 1. Possible increase of max. link utilization with different tie-breakers

C. USP Probability

The allowed maximum link cost k_{max} constrains optimized path layout. When k_{max} is small, equal-cost paths cannot be avoided whereas with a larger link cost range, the routing can be configured so that all split points can be eliminated. To illustrate this, we empirically quantify the USP probability for different k_{max} values, i.e., the probability that a random routing configuration with link costs $k \in [1 : k_{\text{max}}]$ has only unique shortest paths. For each evaluated k_{max} , we generate 1000 random link cost configurations \mathbf{k} and count how many of these settings produce USP routing. Figure 2(a) shows the resulting estimated USP probabilities in the smallest (Cost239) and the largest (Tiscali) evaluated network both for failure-free and resilient routing. The 99% confidence intervals are indicated to show the accuracy of the results. The USP probability is close to zero for small k_{max} . At a certain value it clearly increases and approaches 100% when k_{max} is large enough. For resilient routing, the USP probability is much smaller than for failure-free routing. This is due to the fact that resilient routing requires USPs for all protected failure scenarios \mathcal{S} while failure-free routing requires USPs only for the failure-free case. Hence, any link cost setting producing USPs for resilient routing produces USPs also for failure-free routing but not vice-versa. The differences in USP probability can be rather large. For example in the Tiscali network with $k_{\text{max}} = 2^{10}$, the USP probability is below 10% for resilient routing and already around 50% for failure-free USP routing.

The shape of the USP probability curves is almost identical for all evaluated network topologies but their positions on the x-axis differ. Figure 2(b) describes the location of the curve for all investigated networks. It shows the k_{max} values



(a) USP probabilities for selected networks (Cost239 and Tiscali). (b) Maximum link costs for which the USP probability is 1%, 50%, and 95%.

Fig. 2. Dependence of the USP probability on the maximum link cost and the network size.

for which the USP probability is approximately 1%, 50%, and 95%. The topologies are placed on the logarithmic x-axis according to their number of links. For better readability, we connect the values of the different topologies. The lines have an approximately linear shape, i.e., k_{\max} seems to grow polynomially with the number of links in the network, but in fact k_{\max} also highly depends on the structure of the topology. The curves for 1%, 50%, and 95% USP probability are far apart from each other. That means, to clearly increase the USP probability, the maximum link cost k_{\max} needs to be significantly increased. This holds for any network size. To obtain the same USP probability for resilient routing as for failure-free routing, the maximum link cost settings need to be about four times larger.

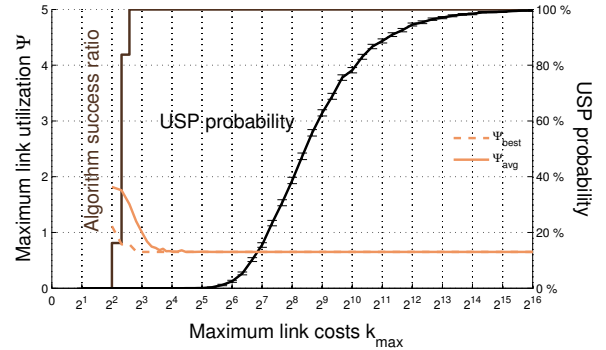
D. Routing Optimization for USP

It is easy to manually create a USP routing for failure-free conditions. Simply, all links on a spanning tree in the network are set to link cost 1 and all other links get a very high value so that only links on the tree are used for (USP-)routing. Another option which even works during any failure scenario is to number the links consecutively and assign a link cost of 2^i to the link with the number i . This way, it is impossible to get the same cost sum over different paths and the USP property is ensured. The latter method only works for networks with up to 16 links. Both methods lead to very bad load balancing and high link loads as only few links are used. A good USP routing distributes traffic so that the maximum link utilization is low. This can be achieved by routing optimization. We extend our heuristic to generate USP routings and then show under which circumstances we receive good results.

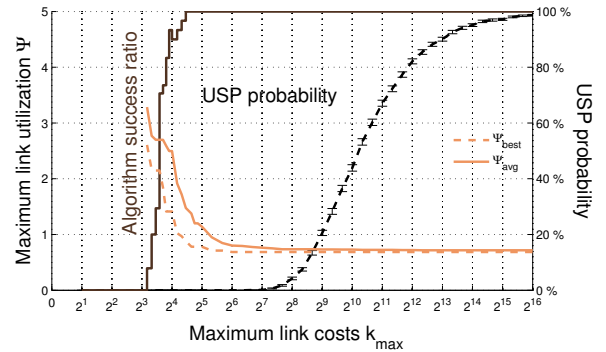
1) *Optimizer*: For USP routing, equal-cost paths must not exist. Inspired by [17], we extend our objective function Ψ and penalize each equal-cost path split with a high penalty value 50. Good Ψ values should be smaller than 1.0 (100% link utilization) and even bad configurations can never lead to (theoretical) link loads higher than 5000%, therefore, this penalty term ensures that even the worst USP solution is preferred to any non-USP solution. Using this simple extension, the heuristic searches systematically for USP routings. However,

it might not find a USP solution at all if k_{\max} was not chosen large enough.

2) *Impact of Link Cost Ranges on Max. Link Utilization*: We optimize the routing in the Exodus network for different k_{\max} values. For each k_{\max} , we perform 50 optimizations for failure-free and for resilient routing. Figure 3 shows the best and the average maximum link utilization values of the 50 experiments. Because the runtime of the heuristic is limited, the obtained results are not necessarily optimal. The values are connected with a solid line for better readability. The success ratio in the figure indicates the percentage of optimizations where the heuristic could find a USP solution. For very small costs, no solutions were found. For $k_{\max} \in [4 : 7]$ in failure-free routing (Figure 3(a)) and for $k_{\max} \in [9 : 20]$ in resilient routing (Figure 3(b)), the heuristic is more and more successful in finding USP solutions. With larger allowed costs, the heuristic always finds a USP routing. For small k_{\max} , the only possible USP routings that are found lead to a high maximum link utilization. Good link utilizations could be achieved only with higher maximum link costs k_{\max} . Analog to the success ratio, resilient USP routing requires a higher k_{\max} than failure-free routing to achieve good results. The figure also shows the USP probability from Section III-C for the different k_{\max} values. Surprisingly, good results are already found in cost ranges where the USP probability is still smaller than 0.1%.



(a) Failure-free scenario: S_0 .



(b) Failure-free and all single link failure scenarios: S_L .

Fig. 3. Success ratio of the optimizer and maximum link utilization of optimized link cost settings depending on the maximum link costs k_{\max} in the Exodus network.

E. Optimal Cost Ranges

To determine the best k_{\max} for different routing mechanisms, we optimize the maximum link utilizations of resilient routing for ECMP, SSP, and USP routing. Figure 4 presents the best and average results out of 50 optimizations for each k_{\max} in the Cost239 network. For USP, the average and best utilization values decrease with increasing k_{\max} . For SSP, the average and best quality of the optimized routing solutions are almost independent of the maximum link costs. For ECMP, the best link cost settings are found for small k_{\max} and the quality of the resulting routing deteriorates with increasing maximum link costs. At first sight, this is surprising since all routings with small maximum link cost can equally be configured with a larger k_{\max} . However, larger maximum link costs increase the search space so that our heuristic optimizer cannot find the best equal-cost path solutions anymore. This observation probably also holds for other similar optimization algorithms, but this is beyond the scope of this paper. The described phenomena are very similar for all examined network topologies, and are also true for the FRR mechanisms. Therefore, in the following optimization, we use a small $k_{\max} = 8$ for ECMP and SSP and a large $k_{\max} = 2^{16}$ for USP.

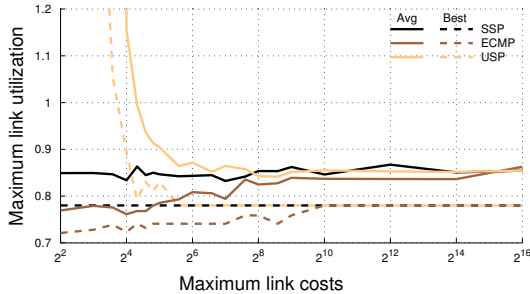


Fig. 4. Maximum link utilization for optimized resilient routing depending on the maximum link costs k_{\max} .

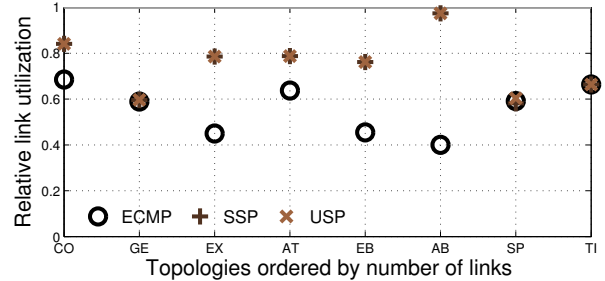
IV. PERFORMANCE COMPARISON

In this section we study link utilization and path length for optimized IP routing. First, we show the performance of ECMP, SSP, and USP routing under failure-free conditions. Then, we investigate resilient routing and compare IP routing reconvergence based on ECMP, SSP, and USP with the fast reroute schemes not-via and MPLS one-to-one backup. We perform 50 optimization runs for each of the various routing options and select the link cost setting with the lowest maximum link utilization for our analysis.

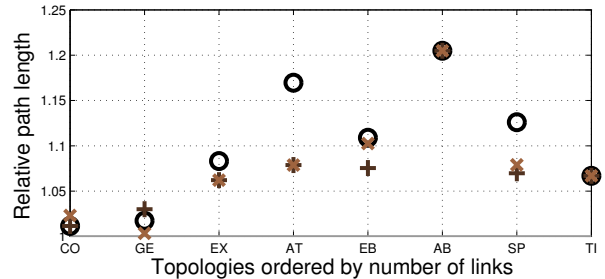
A. Optimized Routing under Failure-Free Conditions

We study the performance of optimized USP routing under failure-free conditions and compare it with optimized ECMP and SSP. Figure 5(a) shows the maximum link utilization for several networks relative to the one of unoptimized hop-count routing. For almost all reported networks, ECMP clearly achieves the lowest maximum link utilization, and SSP and USP routing mostly lead to higher utilization values. While

the actual performance of the different routing options strongly depends on the network topology, USP routing is always as good as regular SSP routing. Figure 5(b) shows the average path length of the resulting routings, normalized by the average path length of hop-count routing which produces shortest paths. The path lengths for SSP and USP routing hardly differ. The ECMP paths are often slightly longer because we only include the longest partial path in our calculation, and routing optimization artificially prolongs some partial paths by assigning them the same costs as to shorter paths in order to balance traffic over all of them.



(a) Maximum link utilization relative to hop-count routing.



(b) Average path length relative to hop-count routing.

Fig. 5. Comparison of optimized ECMP, SSP, and USP routing under failure-free conditions.

B. Optimized Resilient IP Routing and its Applications

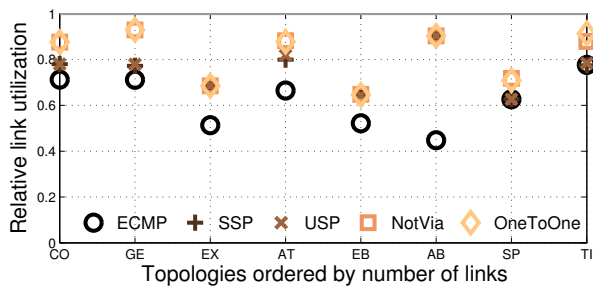
To study the performance of resilient routing, we optimize the IP link costs to minimize the maximum link utilization over all single link failure scenarios S_L . We investigate the ECMP, SSP, and USP routing options with reconvergence in failure cases, the not-via IP FRR mechanism which has the same path layout as the facility backup method for MPLS FRR, and the one-to-one backup method for MPLS FRR.

Figure 6(a) shows the maximum utilization of all links in all considered failure scenarios S_L relative to the one of hop-count routing. The maximum link utilization is the same for SSP and USP. This holds also when FRR mechanisms are used. Thus we omitted the SSP FRR results in Figure 6. ECMP leads to the lowest maximum utilization and the differences to USP are smaller than under failure-free conditions. Fast reroute mechanisms mostly cause larger maximum link utilizations because rerouted traffic continues to be carried close to the outage location which experiences an increased load of backup

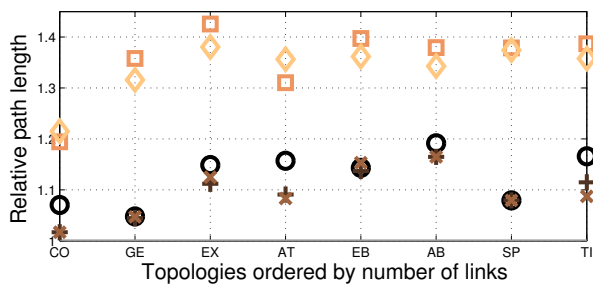
traffic. In contrast, IP reconvergence mechanisms can reroute the traffic more evenly through the network.

Figure 6(b) compares the path length prolongation of the different routing mechanisms compared to hop-count routing. For each ingress-egress aggregate, we take the longest path that occurs in any failure scenario and calculate the average length of these longest paths. The path lengths for ECMP, SSP, and USP are quite similar. The FRR mechanisms prolong the backup paths dramatically because traffic is forwarded on detours around the failure and not on the second-best paths starting at the source as in normal IP reconvergence.

Not-via (MPLS facility backup) usually has the longest paths because it takes a bypass to the next-next-hop router instead of tunneling directly to the destination like MPLS one-to-one backup.



(a) Maximum link utilization relative to hop-count routing.



(b) Maximum path length relative to hop-count routing.

Fig. 6. Comparison of optimized ECMP, SSP, and USP routing as well as fast reroute mechanisms under all single link failures S_L

We have shown that the performance of optimized SSP and USP is quite similar in all studied experiments and that it does not differ a lot from the performance of ECMP, especially for resilient routing. This is also true for FRR mechanisms, which in general lead to increased maximum link utilization and significantly longer average maximum path lengths.

V. RELATED WORK

We briefly review existing work regarding the optimization of IP routing in general, the existing work on USP routing, and optimization of fast reroute mechanisms.

A. General Optimization of IP Routing

In previous work [11] we provided a review of different objective functions that were used for the optimization of IP

routing. The maximum utilization of all links [18]–[23] and the continuous, piece-wise linear, monotonically increasing penalty function proposed by Fortz et al. [23]–[27] are most frequently used. In [27] the authors propose an optimization of the link weights by iteratively modifying the slopes of the Fortz objective function. Fortz et al. were the first to consider reactions to changed networking conditions like sudden congestion due to increased user activity or backup traffic after rerouting. They proposed to modify only a few link costs to adapt the routing to the new conditions [28], [29].

Objective functions can also take a limited set of “protected” failure situations into account. In case of failures, IP routing reconverges which leads to changed load conditions. Optimization of resilient routing attempts to improve routing also under failure conditions using the same set of link cost settings. It has been first presented in [30]–[32]. Later contributions look at faster heuristics [33] or alternative objective functions for special application scenarios [10], [34].

Optimizing IP routing is an NP-complete problem even in the failure-free case [35]–[37]. It can be optimally solved by (integer, mixed integer) linear programs (ILPs, MIPs, LPs) [21], [35], [38]–[44]. Due to the large solution space, these methods are mostly applicable only to small networks because of long runtimes. Therefore, faster heuristics are frequently used. Local search techniques [24] have been applied, genetic algorithms [18]–[20], [25], [45]–[47], simulated annealing, or other heuristics [26], [48]. The efficiency of heuristic methods based on tabu search and steepest ascent are compared with limits obtained by MIPs [49]. Zuo and Pitts [50] investigated the influence of link cost ranges on optimization results but not in the context of USP routing. Riedl et al. increase the search space by considering non-additive link costs [22]. Xu et al. developed a new routing protocol based on link costs PEFT which enables optimal traffic engineering [23].

B. Tiebreakers and Unique Shortest Paths

In case of ECMP, traffic is only approximately evenly split over equal-cost paths towards a desired destination [3]. Thorup and Roughan [4] investigated this problem. To account for the load fluctuations, they added a multiplicative penalty of 20% more load on the links when a traffic aggregate is split over equal-cost paths and respected that in the objective function for routing optimization. This already led to a reduction of path splits. They integrated this concept in the local search technique of [24] and compared ECMP and USP performance considering the failure-free case in the same networks as used in [24]. Lucraft et al. [17] also generated USP routing solutions based on the local search technique of [24] and used multiplicative and additive penalties in the objective function for equal-cost paths. In our algorithm we adopt only the additive approach. Petterson et al. create symmetric USPs without any optimality goal using constraint programming [51]. The paths just needed to be able to carry a given traffic matrix. Zhang [52] proposes a new mathematical formulation of the USP problem which can be solved with constraint generation algorithms. In [53] we compared different exact

and heuristic solution methods based on ILP formulations to obtain USP solutions for optimal IP routing under failure-free conditions and showed that the exact mathematical methods can solve only small problem instances. A simple heuristic based on [10] specialized for failure-free routing mostly led to better solutions than upper bounds received by the exact methods. Bley proposes a decomposition approach to find optimal USPs [54]. In a first step an optimal path routing is computed using integer programming techniques and in the second step link cost settings are determined that induce this routing. In [55] he considers the cost-optimal design of IP networks with USP routing. To the best of our knowledge, none of the papers has tried to produce optimal USP solutions for resilient IP network which is the objective of our work.

C. Fast Reroute Mechanisms

A framework for IP FRR [6] is currently under development in the IETF routing working group. The not-via mechanism is introduced in [7] and improvements have been proposed in [56]. The authors of [57] give an extensive overview on MPLS and IP FRR mechanisms. First insights into the failure coverage of these IP FRR mechanisms have been given in [58]–[60]. Fast reroute concepts were first developed for MPLS technology and standardized in [61]. Currently, extensions for point-to-multipoint are under discussion to protect multicast traffic [62], [63]. The ability of plain IP routing for sub-second reaction to failures was studied in [64], [65] as well as stability issues when performing such optimizations.

VI. CONCLUSION

When several equal-cost shortest paths exist between two nodes in IP networks, tie-breakers decide which of the paths is chosen for single shortest path (SSP) routing. However, the tie-breakers are not clearly defined and possibly do not work in a deterministic way. We have shown that this can lead to unexpected high link loads in the examined networks in spite of routing optimization. To avoid problems with non-deterministic tie-breakers, we proposed to use only IP link costs that lead to unique shortest paths (USP), i.e., only a single shortest path exists between a source and destination pair. We have shown that the fraction of (random) link cost settings that produce USP solutions depends on the maximum allowed link costs and on the set of protected failure scenarios.

One motivation for USP routing is traffic optimization. We presented a heuristic algorithm to find link cost settings that produce USP routing and to minimize the maximum link utilization in all protected failure scenarios for classic IP networks, not-via IP fast reroute, and MPLS networks with facility or one-to-one backup. Our performance comparison shows that, in comparison to regular SSP, the additional path-uniqueness constraint of USP does not deteriorate the routing quality with respect to maximum link utilization and path length. It is a remarkable result that this holds even for the more constraint cases of resilient IP routing and fast reroute. Equal-cost multipath (ECMP) often achieves the lowest link-utilization but at the price of uncertain routing decisions

for specific packets. These results hold in all considered Rocketfuel topologies. Hence, USP routing solutions for IP-based networks can be effectively found and optimized for various reroute techniques without impairing the performance of single path routing. Fast reroute mechanisms in general lead to larger maximum link utilization values and to longer backup paths.

REFERENCES

- [1] J. Moy, "RFC2328: OSPF Version 2," Apr. 1998.
- [2] ISO, "ISO/IEC 10589:2002: Intermediate System to Intermediate System Intra-domain Routing Information Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473)," Nov 2002.
- [3] R. Martin, M. Menth, and M. Hemmkepler, "Accuracy and Dynamics of Multi-Stage Load Balancing for Multipath Internet Routing," in *IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, UK, Jun. 2007.
- [4] M. Thorup and M. Roughan, "Avoiding Ties in Shortest Path First Routing," AT & T, Shannon Laboratory, Florham Park, NJ, Technical Report, 2001.
- [5] M. Menth, F. Lehrieder, B. Briscoe, P. Eardley, T. Moncaster, J. Babiarz, A. Charny, X. J. Zhang, T. Taylor, K.-H. Chan, D. Satoh, R. Geib, and G. Karagiannis, "PCN-Based Admission Control and Flow Termination," accepted for *IEEE Communications Surveys & Tutorials*, 2010.
- [6] M. Shand and S. Bryant, "IP Fast Reroute Framework," <http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-11.txt>, Jun. 2009.
- [7] S. Bryant, S. Previdi, and M. Shand, "IP Fast Reroute Using Not-via Addresses," <http://tools.ietf.org/id/draft-ietf-rtgwg-ipfrr-notvia-addresses-04.txt>, Jul. 2009.
- [8] R. Martin, M. Menth, M. Hartmann, T. Cicic, and A. Kvalbein, "Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute?" *currently under submission*, 2009, <http://www3.informatik.uni-wuerzburg.de/staff/menth/Publications/papers/Menth07-Sub-25.pdf>.
- [9] M. Menth, R. Martin, M. Hartmann, and U. Spoerlein, "Efficiency of Routing and Resilience Mechanisms in Packet-Switched Networks," accepted for *European Transactions on Telecommunications (ETT)*, 2009.
- [10] M. Menth, M. Hartmann, and R. Martin, "Robust IP Link Costs for Multilayer Resilience," in *IFIP-TC6 Networking Conference (Networking)*, Atlanta, GA, USA, May 2007.
- [11] M. Hartmann, D. Hock, C. Schwartz, and M. Menth, "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, Oct. 2009.
- [12] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," in *ACM SIGCOMM*, Pittsburgh, PA, Aug. 2002.
- [13] P. Batchelor et al., "Ultra High Capacity Optical Transmission Networks. Final Report of Action COST 239," Jan. 1999.
- [14] "The GEANT website," <http://www.geant.net/>, 2008.
- [15] A. Nucci, A. Sridharan, and N. Taft, "The Problem of Synthetically Generating IP Traffic Matrices: Initial Recommendations," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 3, pp. 19–32, Jul. 2005.
- [16] M. Roughan, "Simplifying the Synthesis of Internet Traffic Matrices," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 5, pp. 93 – 96, 2005.
- [17] D. Lucraft, A. Eremin, and F. Ajili, "Enforcing Path Uniqueness in Internet Routing, booktitle = sac, year = 2006."
- [18] E. Mulyana and U. Killat, "An Alternative Genetic Algorithm to Optimize OSPF Weights," in *15th ITC Specialist Seminar*, Würzburg Germany, Jul. 2002.
- [19] A. Riedl, "A Hybrid Genetic Algorithm for Routing Optimization in IP Networks Utilizing Bandwidth and Delay Metrics," in *IEEE Workshop on IP Operations and Management (IPOM)*, Dallas, TX, Oct. 2002.
- [20] C. Reichert and T. Magedanz, "A Fast Heuristic for Genetic Algorithms in Link Weight Optimization," in *5th International Workshop on Quality of future Internet Services (QofIS)*, Barcelona, Spain, Sep. 2004, pp. 144–153.

- [21] M. P. Petterson, R. Szymanek, and K. Kuchcinski, "A CP-LP Hybrid Method for Unique Shortest Path Routing Optimization," in *International Network Optimization Conference (INOC)*, 2007.
- [22] A. Riedl and D. A. Schupke, "Routing Optimization in IP Networks Utilizing Additive and Concave Link Metrics," *IEEE/ACM Transactions on Networking*, vol. 15, no. 5, pp. 1136–1148, Oct. 2007.
- [23] D. Xu, M. Chiang, and J. Rexford, "Link-State Routing with Hop-by-Hop Forwarding can Achieve Optimal Traffic Engineering," in *IEEE Infocom*, Phoenix, AZ, Apr. 2008.
- [24] B. Fortz and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights," in *IEEE Infocom*, Tel-Aviv, Israel, 2000, pp. 519–528.
- [25] M. Ericsson, M. Resende, and P. Pardalos, "A Genetic Algorithm for the Weight Setting Problem in OSPF Routing," *Journal of Combinatorial Optimization*, vol. 6, pp. 299–333, 2002.
- [26] S. Balon and G. Leduc, "Combined Intra- and Inter-domain Traffic Engineering using Hot-Potato Aware Link Weights Optimization," in *ACM SIGMETRICS (short paper)*, 2008.
- [27] S. Srivastava, G. Agrawal, and D. Medhi, "Dual-Based Link Weight Determination Towards Single Shortest Path Solutions for OSPF Networks," in *International Teletraffic Congress (ITC)*, 2005.
- [28] B. Fortz, J. Rexford, and M. Thorup, "Traffic Engineering with Traditional IP Routing Protocols," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 118–124, 2002.
- [29] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 756 – 767, May 2002.
- [30] —, "Robust Optimization of OSPF/IS-IS Weights," in *International Network Optimization Conference (INOC)*, Paris, France, Oct. 2003, pp. 225–230.
- [31] D. Yuan, "A Bi-Criteria Optimization Approach for Robust OSPF Routing," in *3rd IEEE Workshop on IP Operations and Management (IPOM)*, Kansas City, MO, Oct. 2003, pp. 91 – 98.
- [32] A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, and C. Diot, "IGP Link Weight Assignment for Transient Link Failures," in *18th International Teletraffic Congress (ITC)*, Berlin, Sep. 2003.
- [33] A. Sridharan and R. Guerin, "Making IGP Routing Robust to Link Failures," in *IFIP-TC6 Networking Conference (Networking)*, Ontario, Canada, May 2005.
- [34] M. Menth and M. Hartmann, "Threshold Configuration and Routing Optimization for PCN-Based Resilient Admission Control," *Computer Networks*, vol. 53, no. 11, pp. 1771 – 1783, Jul. 2009.
- [35] M. Pióro, Á. Szentési, J. Harmatos, A. Jüttner, P. Gajowniczek, and S. Kozdrowski, "On Open Shortest Path First Related Network Optimisation Problems," *Performance Evaluation*, vol. 48, pp. 201 – 223, 2002.
- [36] B. Fortz and M. Thorup, "Increasing Internet Capacity Using Local Search," *Computational Optimization and Applications*, vol. 29, no. 1, pp. 13–48, 2004.
- [37] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufman, 2004.
- [38] S. Köhler, D. Staehle, and U. Kohlhaas, "Optimization of IP Routing by Link Cost Specification," in *15th ITC Specialist Seminar*, Würzburg, Jun. 2002, pp. 167 – 175.
- [39] A. Sridharany, R. Guérin, and C. Diot, "Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks," in *IEEE Infocom*, San Francisco, CA, Apr. 2003.
- [40] C. Lopes and A. de Sousa, "Heuristics for the MPLS Network Design with Single Path Minimum Weight Routing," in *International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks (HETNET)*, 2005.
- [41] Y. Wang, Z. Wang, and L. Zhang, "Internet Traffic Engineering without Full Mesh Overlaying," pp. 565 – 571, 2001.
- [42] K. G. Ramakrishnan and M. A. Rodrigues, "Optimal Routing in Shortest-Path Data Networks," *AT&T Bell Laboratories Technical Journal*, pp. 117–138, Spring 2001.
- [43] A. Tomaszewski, M. Pióro, M. Dzida, and M. Zagozdzon, "Optimization of Administrative Weights in IP Networks Using the Branch-and-Cut Approach," in *International Network Optimization Conference (INOC)*, 2005.
- [44] A. Tomaszewski, M. Pióro, M. Dzida, M. Mycek, and M. Zagozdzon, "Valid Inequalities for a Shortest-Path Routing Optimization Problem," in *International Network Optimization Conference (INOC)*, Spa, Belgium, 2007.
- [45] É. Mulyana and U. Killat, "A Hybrid Genetic Algorithm Approach for OSPF Weight Setting Problem," in *2nd Polish-German Teletraffic Symposium (PGTS)*, Gdansk Poland, Sep. 2002.
- [46] A. Riedl, "Optimized Routing Adaptation in IP Networks Utilizing OSPF and MPLS," in *IEEE International Conference on Communications (ICC)*, Anchorage, May 2003.
- [47] N. Wang, K.-H. Ho, , and G. Pavlou, "Adaptive Multi-Topology IGP Based Traffic Engineering with Near-Optimal Network Performance," in *IFIP-TC6 Networking Conference (Networking)*, Singapore, May 2008.
- [48] J. Harmatos, "A Heuristic Algorithm for Solving the Static Weight Optimisation Problem in OSPF," in *IEEE Globecom*, Nov. 2001.
- [49] M. Dzida, M. Zagozdzon, and M. Pióro, "Optimization of Resilient IP Networks with Shortest Path Routing," in *International Workshop on the Design of Reliable Communication Networks (DRCN)*, La Rochelle, France, 2007.
- [50] Y. Zuo and J. Pitts, "Impact of Link Weight Ranges on OSPF Weight Solutions," in *2nd*, Shanghai, China, Aug. 2007.
- [51] M. P. Petterson, R. Szymanek, and K. Kuchcinski, "A CP-LP Approach to Network Management in OSPF Routing," in *SAC*, 2007.
- [52] C. Zhang, "A Novel Mathematical Model for the Unique Shortest Path Routing Problem," ArXiv e-prints, Technical Report, No. 2008arXiv0807.0038Z, Jun. 2008. [Online]. Available: <http://arxiv.org/abs/0807.0038>
- [53] M. Dzida, M. Zagozdzon, M. Zotkiewicz, M. P. Petterson, M. Pióro, M. Duelli, and M. Menth, "Three Methods for Optimizing Single-Shortest Path Routing," in *4th Conference on Next Generation Internet Networks (NGI)*, Kraków, Poland, Apr. 2008.
- [54] A. Bley, "An Integer Programming Algorithm for Routing Optimization in IP Networks," in *16th Annual European Symposium on Algorithms (ESA)*, Karlsruhe, Germany, 2008, pp. 198–209.
- [55] —, "A Lagrangian Approach for Integrated Network Design and Routing in IP Networks," in *International Network Optimization Conference (INOC)*, 2003, pp. 107–113.
- [56] A. Li, P. Francois, and X. Yang, "On Improving the Efficiency and Manageability of NotVia," in *ACM Conference on emerging Networking Experiments and Technologies (CoNEXT)*, New York, NY, Dec. 2007.
- [57] A. Raj and O. Ibe, "A Survey of IP and Multiprotocol Label Switching Fast Reroute Schemes," *Computer Networks*, vol. 51, no. 8, pp. 1882–1907, 2007.
- [58] P. Francois and O. Bonaventure, "An Evaluation of IP-Based Fast Reroute Techniques," in *ACM Conference on emerging Networking Experiments and Technologies (CoNEXT)*, Toulouse, France, Oct. 2005, pp. 244–245.
- [59] A. F. Hansen, T. Cicic, and S. Gjessing, "Alternative Schemes for Proactive IP Recovery," in *2nd Conference on Next Generation Internet Design and Engineering (NGI)*, Valencia, Spain, Apr. 2006.
- [60] M. Gjoka, V. Ram, and X. Yang, "Evaluation of IP Fast Reroute Proposals," in *IEEE International Conference on Communication System Software and Middleware (COMSWARE)*, Bangalore, India, Jan. 2007.
- [61] P. Pan, G. Swallow, and A. Atlas, "RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels," May 2005.
- [62] R. Aggrawal, D. Papadimitriou, and S. Yasukawa, "RFC4875: Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)," May 2007.
- [63] J. L. L. Roux, R. Aggarwal, J. Vasseur, and M. Vigoureux, "P2MP MPLS-TE Fast Reroute with P2MP Bypass Tunnels," <http://tools.ietf.org/id/draft-ietf-mpls-p2mp-te-bypass-02.txt>, Mar. 2008.
- [64] A. Basu and J. G. Riecke, "Stability Issues in OSPF Routing," in *ACM SIGCOMM*, San Diego, CA, USA, Aug. 2001, pp. 225–236.
- [65] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving Sub-Second IGP Convergence in Large IP Networks," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 2, pp. 35 – 44, Jul. 2005.