BUDAPEST UNIVERSITY OF TECHNOLOGY AND ECONOMICS
FACULTY OF ELECTRICAL ENGINEERING AND INFORMATICS
DOCTORAL SCHOOL OF COMPUTER SCIENCE

NEW NETWORK OPTIMIZATION METHODS FOR FAST PROTECTION IN IP NETWORKS

Levente Csikor
M.Sc. in Computer Engineering

Summary of the Ph.D. Dissertation

Supervised by

Dr. Gábor Rétvári, Ph.D.
Senior Research Fellow
High Speed Networks Laboratory (HSN Lab)
Dept. of Telecommunications and Media Informatics
MTA-BME Future Internet Research Group
Budapest University of Technology and Economics

Budapest, Hungary
2015

1 Introduction

Nowadays, many commercial telecom and multimedia providers have started to broadcast their contents over the Internet in order to reduce costs and reach new users. This continuous convergence is producing a more and more heterogeneous traffic with different demands as many real-time applications such as VoIP, IPTV, online gaming, etc. have become available. However, the integration has happened so fast that the network community could not adapt the IP protocol suite fast enough to keep up with the new requirements and there still exist missing components to facilitate the desired transmission quality.

One of the main concerns is that an operational network frequently suffers component failures [18, 10, 28, 32, 22] and the currently used reactive techniques [4, 2] take too much time to recover the appropriate paths and bypass a failed component. Instead, faster local rerouting-based proactive approaches can be used, whereby after a failure, the traffic could be immediately switched to an alternate route by means of precalculated backup paths. Local rerouting means that the adjacent router tries to solve the problem locally without notifying every other router about the failure.

Although there have been numerous proposals to provide faster proactive protection in IP networks [1, 17, 12, 31, 6, 16], so far only the Loop-Free Alternates (LFA, [3]) technique has been widely implemented in today's commercial routers. The main idea in LFA is to ensure that, in case of the failure on the primary forwarding path, there be a suitable backup neighbor whose path to the destination is unaffected by the failure. The strength of LFA lies in its simplicity, however, this simplicity comes at the price that not all networks could be protected. For instance, consider the network depicted in Fig. 1, and suppose node s

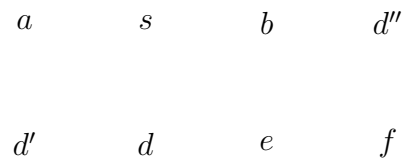

Figure 1: A simple network where certain failure cases can be protected neither by LFA nor by Remote LFA. The arrows indicate the shortest paths from s to d' and d''

wants to send a packet to node d', but the link between node s and its next-hop a has failed. In this case, s cannot pass the packet to its other neighbor, as node b's shortest path to d' may go through nodes s and a, this way causing a loop.

Recently, a generalization of LFA, called Remote LFA (rLFA, [5]), has been defined where, in contrast to simple LFA, not just direct neighbors but remote nodes too can be considered when searching for the backup path. In our example above, if a tunnel[1] was created between s and e, then e, now being a direct neighbor of s, would become an LFA as node e's shortest path to d' bypasses the failed link (s, a). However, rLFA is equally plagued by LFA's main weakness, namely, that there are many networks that cannot be *completely* protected, leaving these networks vulnerable to certain failure scenarios. Suppose that if

[1]Note that the tunneled traffic is restricted to shortest paths just like "normal" traffic, hence the tunnel must avoid the failure as well.

node s wishes to send a packet to node d'' and the link (s, b) fails. Then, (s, b) cannot be protected since all nodes whose shortest path does not go through (s, b) can only be reached from s through (s, b) itself. This suggests that while the use of rLFA definitely can provide higher protection level against link failures than pure LFA, it still does not facilitate full protection for all failure cases in a general topology.

2 Research Goals

The objective of my Dissertation is to give a comprehensive mathematical analysis of the different forms of LFAs and show to what extent a network can be protected by these mechanisms when single link or node failures occur. I study the aspects of a network's topology that make it unsuitable to attain full protection with LFA over that network and I also show best-case scenarios. Furthermore, I investigate whether the failure case coverages could be improved by optimizing the network topology.

In particular, in the first part I investigate to what extent a network can be optimized to improve the level of protection provided by LFA with only minor changes to the topology. I examine the complexity of these problems, and I give algorithms to solve them. I demonstrate that the protection provided by LFA can be boosted close to 100% failure case coverage against single link failures, while I also show that my methods are effective against the rarer but relevant case of single node failures as well.

In the second part, I extend the failure case coverage analysis and network optimization techniques, so far only available to LFA, to the upcoming rLFA technique, and I show that the level of protection can be improved significantly by Remote LFA compared to the protection provided by pure LFA, and I analyze which particular topological properties bound the failure case coverage of Remote LFA.

Last but not least, I show network optimization techniques to provide 100% rLFA failure case coverage irrespectively of whether link or node protection is considered.

I believe that my results may contribute to the general understanding of the level of protection that the LFAs can provide, and may help hesitating network operators to reach a verdict of using any of the available approaches.

3 General Assumptions

In my Dissertation, I am dealing with LFA and rLFA intended to provide fast failure recovery in an intra-domain setting. Consequently, I suppose that the analyzed network is an autonomous system, where all routers have a complete view of the network, i.e., a *link-state routing protocol* such as OSPF [24] or IS-IS [14] is deployed. I also assume that the calculation of next-hops is *based on only the destinations' addresses* and no other information (e.g., source routing, additional notifications) is taken into account. Moreover, *each node has a well-defined next-hop* to each destination, even if Equal Cost Multiple Paths (ECMPs) exist, i.e., when the number of shortest paths from a (source) node to a destination is more than one, then one of them is selected, exclusively. This is mostly a

technical assumption to ease the mathematical treatment and it is easy to relax if necessary.

I further assume that only interior destinations must be protected, thus every failure outside the autonomous system should be treated by inter-domain routing protocols such as BGP, or by the IGPs of other autonomous systems. Since in BGP, the routing is not solely based on shortest paths, in this Dissertation, I do not deal with failures among ASes.

Furthermore, since the most common failures are transient single failures [21, 13], I only deal with *single link or node failure*. Note, however, that in case of node protection special care must be taken to handle the so called *last-hop problem*, which arises when destination d is an immediate neighbor of source s and the default shortest path between them is exactly the link (s, d) (see, for instance, the case of the source-destination pair (b, d'') in Fig. 1). In such cases, the node failure we want to protect is exactly the failure of the destination d itself, which case is hardly protectable by LFA. Therefore, to resolve this issue and increase the overall level of protection, it is recommended to support fallback to link protection in these cases [19]. However, from a traffic engineering point of view maximizing the coverage in such a way may not be the right approach, since the affected traffic may be dropped on congested backup links. So, there is no standard on how this case should be treated. Therefore, in my Dissertation I consider both cases. During my LFA analysis, for such source-destination pairs I only require that the link (s, d) be protected by a link-protecting LFA, and I ignore the node-protection requirement. On the other hand, in my Remote LFA analysis I consider node protection as *undefined* between an arbitrary neighboring node pair. However, if there is a significant difference between the two approaches, I shall examine both cases.

Furthermore, it should be noted that in case of multilayered networks many overlay (IP) links use one physical link as the transmission medium. The failure of such a particular link (e.g., a cable cut-off) may cause failures in a set of virtual links in the overlay topologies. Dealing such cases of multiple failures is out of scope of my Dissertation.

From graph topological aspect, I assume that the network is a *simple, undirected, weighted graph* and the links are bidirectional and point-to-point, i.e., no SRLGs (Shared Risk Link Group) or LANs (Local Area Network) exist throughout the network. Furthermore, the costs of the links are symmetric. In case of Remote LFA, I initiate the analysis in graphs with unit costs. Note that unweighted graphs are highly relevant in real-world networks, and as shall be shown, results for LFA can only be generalized to rLFA under the unit cost assumption.

Since an arbitrary link can only be protected if the graph of the network is *2-edge-connected*, I assume this minimum topological requirement for link-protecting case. For the case of node protection, I also assume the graph to be *2-node-connected*. Connectedness is important, since if after a failure the network fell into two distinct parts, there is no technique that can ensure a working path between an arbitrary source-destination pair.

4 Methodology

According to the general assumptions, I always model a network as a *simple, undirected, weighted graph* $G(V, E)$ with V being the set of nodes and E the set of links. Let $n = |V|$

and $m = |E|$, and denote the complement link set with $\overline{E}$. A particular link is denoted by (i, j), where i and j are nodes from V. Link costs are represented by a cost function $c : E \mapsto N$. The cost of a link (i, j) is denoted with $c(i, j)$. Since link costs are assumed to be symmetric $c(i, j) = c(j, i)$. In order to describe the length of the shortest path between an arbitrary node pair (u, v), I use the notation $\text{dist}(u, v)$. As a consequence of considering bidirectional and symmetric links, $\text{dist}(u, v) = \text{dist}(v, u)$.

In my model, I always sought graph-theoretical *lower and upper bounds* on failure case coverages in order to study the performance of the different protection mechanisms in artificial and in real-world network topologies. In the former case, I generated different networks known from the literature such as topologies used in telecommunications (e.g., rings, grids) and well-known topologies from graph theory such as Möbius ladder and bipartite graphs. On the other hand, real-world topologies could be inferred from several existing ISP (Internet Service Provides) datasets such as the Rocektfuel [20], SNDLib [29], and Topology Zoo [15].

I was always careful to study networks with different topological properties, e.g., size of the network, density, average node degree, connectivity. After I found that in an arbitrary network the failure case coverages can be particularly poor, I tried to find *network optimization methods* to improve them. In these cases, my aim was solely to *improve the failure case coverages* of the different forms of LFAs, and I did not deal with traffic engineering or load balancing related issues, or possible congestion that could occur when the protected traffic is directed to a (possibly) crowded link. To validate proposed theorems mathematical proofs were made. In general, the results indicated that *most of the problems are NP-complete*, so *exact algorithms (ILPs)* were formulated to solve them. However, in most of the cases (algorithms mainly run on real network topologies), *approximating heuristics* were necessary to obtain reasonable results within a limited time frame. In the latter cases, I used *greedy algorithm* as a first approach. Then, a *simulated annealing based heuristic framework* was developed with many *tunable parameters* to avoid getting stuck in local optima that may happen by the greedy approach. However, as shall be shown, the greedy approach outperformed the simulated annealing based framework in many cases.

Last but not least, the simulation tools were implemented particularly in C++/ LEMON[2] and the results were conducted through various BASH[3] scripts. In some cases, the networks were described in Geographic Markup Language (GML), therefore in order to analyze them with LEMON a conversion was needed, which was carried out by my own software, called GML2LGF converter [9].

[2]LEMON stands for Library for Efficient Modeling and Optimization in Networks. It is a C++ template library providing efficient implementations of common data structures and algorithms with focus on combinatorial optimization tasks connected mainly with graphs and networks (http://lemon.cs.elte.hu/trac/lemon accessed in October 2013).

[3]Bourne Again SHell - it is a command-line interface for interacting with Unix/Linux based operating systems.

5 New Results

5.1 Loop-Free Alternates

As it was briefly shown in Section 1, in LFA, when a failure occurs, the adjacent router tries to pass the packet to an alternate neighbor that still has a functioning path to the destination.

Next, I show the conditions and different protection schemes of LFA formally. Consider the network depicted in Fig. 1. Suppose that node s wants to send a packet to node d and its default next-hop a is unreachable, since the link (s, a) between them went down. In this case, s has to find an alternate neighbor, which will not pass the packet back. Fortunately, node b fulfills this requirement, so s can reroute the traffic destined to d towards b. Here, we say that b is a link-protecting LFA for node s towards destination node d [3].

Definition 5.1. *Given a connected graph $G(V, E)$, some source s and destination d, let e be the default next-hop of s towards d. Then, some neighbor n of s is a* link-protecting LFA *for s to d if*

(i) *$n \neq e$, and*

(ii) *the loop-free condition applies:*

$$\text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, d) \ . \tag{1}$$

One can easily verify this condition. If the distance from node n to d is strictly less than the distance from n back to the source node s then to node d means that node n will not pass the packet back to s, this way bypassing the failed link (s, e).

Similarly to link protection, node-protecting LFA can also be defined.

Definition 5.2. *For some source s and destination d, let e be the default next-hop of s towards d. Then, some neighbor n of s is a* node-protecting LFA *for s to d if, in addition to* (i) *and* (ii) *in Definition 5.1, the node-protection condition also applies:*

$$\text{dist}(n, d) < \text{dist}(n, e) + \text{dist}(e, d) \ . \tag{2}$$

5.2 Improving LFA Coverage by Optimizing IGP Link Costs

As observed above, there are certain failure cases that can be protected neither by LFA nor by Remote LFA. This calls for developing network optimization tools to tune the network topology in a way as to increase the number of failure cases protectable by LFA. There are various approaches to reach this end. One way is *LFA Network Design*, which aims to design LFA-friendly network topologies right from the outset [11]. Another approach is *LFA Graph Extension*, where the task is to augment the network topology with a few number of new links to boost LFA coverage [25, 27]. On the other hand, *LFA Cost Optimization* asks to construct IGP link costs in a way as to maximize the number of possible failure cases protectable by LFA [30, 23, 26, 8]. Furthermore, *Combined LFA network optimization* [7]

asks for both adding new links and optimizing link costs to the same end, since a negative side effect caused by one of them can be compensated by the other one.

In this section, I investigate *LFA Cost Optimization problem*, which again asks to optimize link costs in a way as to maximize the level of protection provided by LFA.

By analyzing the LFA Cost Optimization problem, I have obtained the following results.

Thesis 1. *I have defined the LFA Cost Optimization problem formally. I have proven that this problem is NP-complete both in the link-protecting and node-protecting cases, and I have proposed a heuristic framework to obtain approximate solutions. By extensive numerical evaluations, I have found that my algorithms yield at least* 10% *improvement in LFA coverage on many real-world network topologies.*

5.2.1 Problem Formulation

Formally, *LFA Cost Optimization problem for the link-protecting case* can be defined as follows:

Definition 5.3. *LFACostOptLP(G, $\mathcal{S}$): Given a graph $G(V, E)$ and a set of source-destination pairs $\mathcal{S}$, is there a cost function c so that $\eta_{\mathcal{S}}^{LP}(G, c) = 1$?*

Easily, a similar problem formulation LFACostOptNP(G, $\mathcal{S}$) exists for the node-protecting case as well. When no ambiguity arises, I shall refer to both problems simply as LFA-CostOpt. In addition, I shall in many cases treat the optimization version of these problems, that is, I shall seek the costs that maximize network-wide LFA coverage instead of merely asking whether or not a cost setting for full protection exists.

5.2.2 The Potential of LFA Cost Optimization

The question immediately arises as to whether it is worth optimizing costs for LFA at all. Easily, readjusting costs in most of the cases alters, possibly in a negative way, default shortest paths, which might have been previously tweaked with great accuracy to match the needs of the network. As the following claim indicates, the wins achievable with optimizing link costs for LFA can be substantial (more than 50%), which might compensate for the losses in forwarding efficiency in certain cases.

Thesis 1.1. *[J2, C1] For any $k > 0$ integer, there is a graph G on $n = 4k + 2$ nodes with two cost functions c_1 and c_2, so that $|\eta(G, c_1) - \eta(G, c_2)| \geq \frac{1}{2}$.*

Consider the so called "Möbius ladder" topology with unit costs depicted in Fig. 2(a), wherein the link-protecting LFA failure coverage $\eta^{\mathrm{LP}}(G, c) = 0.4$. However, in Fig. 2(b), the costs of diagonals are chosen so that the paths between any two nodes are shorter around the ring than through it via a diagonal. This way, as one easily verifies, every source-destination pair is protected, i.e., $\eta^{\mathrm{LP}}(G, c) = 1$. Moreover, since the paths from the LFAs reached by the diagonals bypass the failed next-hops as well in every case, $\eta^{\mathrm{NP}}(G, c) = 1$.

This graph construction can be generalized to arbitrary Möbius ladder topology on $n = 4k + 2$, where $k \in \{1, 2, 3 \ldots\}$, and one can always choose the above cost setting strategy to achieve complete LFA protection.

(a) Wheel graph with unit cost links

(b) Wheel graph with optimized link costs

Figure 2: Illustration for Thesis 1.1

5.2.3 Computational Complexity

I studied the computational complexity of the LFA Cost Optimization problem in the link-protecting case.

Thesis 1.2. *[J2, C1] I have proven that LFACostOptLP(G, $\mathcal{S}$) is NP-complete.*

The complexity directly comes from the fact that shortest paths change during the optimization process, therefore it is possible that altering a link cost can provide protection for a certain source-destination pair, but it may eliminate LFAs to other destinations. Moreover, if only one destination node is considered at a time, the complexity still remains NP-complete. This problem can be mapped to Protection Routing problem (PR, [17]), wherein the aim is to provide protection against single node failures by means of different directed spanning DAGs (Directed Acyclic Graph). It was proven that PR is NP-complete, and I have shown in [C1] that LFACostOptLP(G, $\mathcal{S}$) can be (Karp-)reduced to PR, since a cost function uniquely determines a certain DAG and vice versa. I have observed that the proof remains valid if we treat link failures only and disregard node failures.

Since the above-mentioned PR problem essentially treats only node failures, the reduction given for the link-protecting case in Thesis 1.2 remains valid for node-protecting case as well.

Thesis 1.3. *[J2] I have shown that LFA Cost Optimization problem remains NP-complete in the node-protecting case.*

5.2.4 Numerical Evaluation

In [J2, C1], an exact algorithm was given to solve this problem, but in most of the cases it could not finish in finite time. Therefore, developing approximating heuristics seems the only viable option in order to solve this problem.

Thesis 1.4. *[J2, C1] I have developed a family of fast and efficient heuristics for solving the LFACostOptLP(G, $\mathcal{S}$) problem. A series of extensive simulation studies I have conducted on real and artificial network topologies suggest that on average at least 10% LFA coverage improvement can be attained in real-world network topologies.*

Table 1: Brief results for the LFA cost optimization heuristics in some real topologies.

Topology	n	m	Δ	$\eta_{LP}(G,c)$	$\eta_{LP}(G,c^*)$	$\eta_{NP}(G,c)$	$\eta_{NP}(G,c^*)$
AS1221	7	9	2.57	0.809	0.833	0.452	0.523
AS1239	30	69	4.60	0.873	0.963	0.757	0.937
AS1755	18	33	3.66	0.872	0.993	0.764	0.941
AS3257	27	64	4.74	0.923	1	0.726	0.938
AS3967	21	36	3.42	0.785	0.953	0.642	0.897
AS6461	17	37	4.35	0.933	1	0.738	0.886
Abilene	12	15	2.5	0.56	0.674	0.515	0.606
AT&T	22	38	3.45	0.822	0.987	0.58	0.82
Deltacom	113	161	2.85	0.577	0.662	0.488	0.581
Geant	37	55	2.97	0.69	0.76	0.41	0.622
Germany	17	25	2.94	0.695	0.911	0.562	0.727
InternetMCI	19	33	3.47	0.904	0.932	0.704	0.809
Italy	33	56	3.39	0.784	0.944	0.57	0.803

I have defined a family of heuristics with tunable performance and running time, which facilitates for picking the right approximation algorithm for the particular topology and network optimization goals under consideration. By extensive simulations conducted on a wide selection of realistic network topologies, I have found that in many real network topologies close to perfect LFA coverage can be achieved with cost optimization. In particular, I have found that the denser the network, the higher the LFA coverage. For instance in networks, where the average node degree is greater than 3 are amenable to LFA, but when the degree falls below 3, the chances of getting a high LFA coverage rapidly vanish. Brief results of link-protecting case attained in real topologies are presented in Table 1, where n and m mark the number of nodes and the number of links, while Δ denotes the average node degree. $\eta_{LP}(G,c)$ and $\eta_{LP}(G,c^*)$ represents the initial link-protecting LFA coverage and the link-protecting LFA coverage obtained by LFA Cost Optimization, respectively.

My results indicate that LFA Cost Optimization greatly helps the operator to provide higher availability in the network, even when adding new connectivity to the topology is not feasible. In addition, I have found that the running time of the approximation algorithm strongly depends on the topology, in particular, on the number of links and nodes. In some cases, all 500 rounds terminate in just a couple of minutes, but sometimes it takes a couple of hours. Note that running time is not that important in practice, because LFA Cost Optimization is performed offline only once, before the final deployment of a network.

Thesis 1.5. *[J2] I have extended the algorithmic framework to cover the node-protecting case as well, and I have conducted simulations indicating that the algorithms increase LFA protection against single node failures by* 10-20 *%.*

By means of extensive simulations, I have found that the extended algorithm could improve node-protecting LFA coverage by about $10-20\%$, so the average initial 60% failure case coverages could be boosted up to $75-80\%$. As in the case of link protection, succinct results of the node-protecting case obtained in real topologies can be seen in Table 1, where $\eta_{NP}(G,c)$ denotes the initial node-protecting LFA coverage, whilst $\eta_{NP}(G,c^*)$ marks the node-protecting LFA coverage attained by LFA Cost Optimization.

5.3 Combined LFA Network Optimization

So far, the LFA Graph Extension problem and the above LFA Cost Optimization problem have been studied separately in the literature. However, in certain cases a network operator may not be able to rely on solely one of them. The cost of augmenting the network with new physical links can be reasonable expensive, and in certain cases only one additional link would be required to attain the desired reliability. On the other hand, when other objectives, such as traffic engineering or load balancing play a more important role, altering the link costs and by this way affecting certain shortest paths may cannot be tolerated. Moreover, an additional link may alters existing shortest paths, which can be compensated by optimizing link costs at the same time. In these cases, an efficient combination of the aforementioned two methods could help network operators to find a compromise solution. By analyzing Combined LFA Network Optimization, I have attained the following results.

Thesis 2. *I have defined the Combined LFA Network Optimization problem formally. I have proven that this problem is NP-complete in both the link-protecting and node-protecting cases, and I have shown how the LFA Graph Extension and the LFA Cost Optimization methods should be combined. I have also given simulations indicating that both adding new links to a network and optimizing link costs at the same time are effective ways to improve the LFA coverage in operational networks.*

5.3.1 Problem Formulation

The *Combined LFA Network Optimization problem* can be formulated as follows:

Definition 5.4. *LFACombinedOptLP(G, $\mathcal{S}$, k): Given a simple, undirected, weighted graph $G(V, E)$, a set of source-destination pairs $\mathcal{S}$, and a positive integer k, is there a set $F \subseteq \overline{E}$ with $|F| \leq k$ and properly chosen cost function c, so that $\eta_{\mathcal{S}}^{LP}(G(V, E \cup F), c) = 1$?*

The difference from the LFA Graph Extension problem is that in the above formulation we allow for link costs and, consequently, the shortest paths to change.

5.3.2 Computational Complexity

Thesis 2.1. *[J1] I have shown that the Combined LFA Network Optimization problem is NP-complete.*

From all the optimization problems treated so far, Combined LFA Network Optimization is the most difficult, since both subproblems it covers are NP-complete. Hence, this combined problem is NP-complete as well.

Therefore, instead of aiming for an optimal solution, I propose a heuristic algorithm based on the consecutive application of the heuristics presented for the individual aforementioned subproblems. In particular, in every iteration I executed an LFA Graph Extension phase followed by an LFA Cost Optimization phase. In the first phase, I add one new link to the network. In the LFA Cost Optimization phase, I compute a link cost setting that approximately maximizes the LFA coverage on the augmented graph obtained in the previous phase. The two phases are applied iteratively one after the other, until the LFA coverage reaches 1.

Thesis 2.2. *[J1] I have shown that the combined algorithm significantly reduces (on average by more than* 50*%) the number of additional links necessary for reaching* 100*% LFA coverage.*

So far, the combined algorithm was found to produce the best results, indicating that in many real networks complete LFA-based protection is attainable with adding only a few new links.

5.4 Analysis of Remote LFA Failure Case Coverage

Remote LFA (rLFA) is an extension to LFA that provides additional backup connectivity when none can be provided by the basic mechanisms. As it was shown in Fig. 1, in Remote LFA when a link cannot be entirely protected with local LFA neighbors, the protecting router seeks the help of a remote LFA staging point by means of tunneling. Note that such tunnels are only used as detours, so they do not affect the normal flow of traffic in any ways. There are numerous tunneling mechanisms that fulfill the requirements of this design. In an MPLS/LDP (Multiprotocol Label Switching-Label Distribution Protocol) enabled network, for instance, a simple label stack can be used to provide the required tunnel.

We have seen as well, however, that while the use of rLFA definitely can provide higher protection level than pure LFA, it still does not facilitate full protection for all failure cases in a general topology.

Below, I show when can be a certain source-destination pair protected via Remote LFA. In case of link failure, the set of routers, which can be reached from a source without traversing the failed link is termed the $P-$*space* of the source with respect to the failed link (hereafter $\mathcal{P}_{\rm LP}$, where $_{\rm LP}$ refers to link-protecting case). The set of routers from which the destination can be reached without traversing the failed link is termed the $Q-$*space* (hereafter $\mathcal{Q}_{\rm LP}$) of the destination with respect to the failed link. Since the source router will only use a repair path when it has detected the failure of the link, the initial hop of the repair path needs not be subject to the source's normal forwarding decision process. Therefore, the term "extended P-space" was also defined, which is the union of the P-spaces of each of the source's neighbors. The intersection of the source's $\mathcal{P}_{\rm LP}$ (or $\mathcal{P}^{\rm e}_{\rm LP}$) and the destination's $\mathcal{Q}_{\rm LP}$ with respect to the failed link defines the viable repair tunnel endpoints, known as $\rm PQ_{LP}$-nodes, which are practically the Remote LFAs. When extended P-space is also taken into account, the $\rm PQ_{LP}$-nodes are considered as extended Remote LFAs. In case of node protection, terms ($\mathcal{P}_{\rm NP}$, $\mathcal{P}^{\rm e}_{\rm NP}$ $\mathcal{Q}_{\rm NP}$, $\rm PQ_{NP}$-nodes) can be defined in a similar way.

In the following, I give new tools for a graph-theoretical analysis of rLFA coverage, as measured by $\mu_{\rm LP}(G)$ and $\mu_{\rm NP}(G)$, and I point out some of the intricate relations of basic LFA and Remote LFA.

Thesis 3. *I have performed analytical and numerical study of the level of protection achievable by Remote LFA in both link- and node-protecting cases. I have given alternative sufficient and necessary conditions for a node pair to be rLFA-protected in an arbitrary network. I have found a deep connection between basic LFAs and Remote LFAs, and I have proven that in unit cost networks extended Remote LFA provides* 100*% failure case coverage against single link failures. I have also proven that this statement does not apply for node protection.*

5.4.1 Link-protecting Case

First, I have reformulated the notions of P-space and Q-space in terms of shortest path distances, similar to the way in which pure LFAs were defined in [3] (see Eq. 5.1 and 5.2). This brings us to the following alternative sufficient and necessary conditions for link-protecting rLFAs:

Thesis 3.1. *[C2, J3] For source s, destination d, and next-hop e, some node $n \neq s, d$ is a link-protecting Remote LFA for the $s-d$ pair (i.e., $n \in \mathrm{rLFA}_{\mathrm{LP}}(s,d)$) if and only if*

$$\mathrm{dist}(s,n) < \mathrm{dist}(s,e) + \mathrm{dist}(e,n) \tag{3}$$

$$\mathrm{dist}(n,d) < \mathrm{dist}(n,s) + \mathrm{dist}(s,d) \ . \tag{4}$$

One can easily see that (3) in essence states that for source s and next-hop e some $n \in V$ is in $\mathcal{P}_{\mathrm{LP}}(s,e)$. In addition, (4) means that for source s and destination d some $n \in V$ is in $\mathcal{Q}_{\mathrm{LP}}(s,d)$ and so the repair tunnel cannot traverse the failed link. This condition also corresponds to the basic loop-free criterion of link-protecting LFAs.

Second, I have reformulated the notions of "extended P-space" in terms of shortest path distances. Again, this leads to the following alternative sufficient and necessary condition for link-protecting extended rLFA:

Thesis 3.2. *[C2, J3] For source s, destination d, and next-hop e, some node $n \neq s, d$ is an extended link-protecting Remote LFA for the $s-d$ pair if and only if*

$$\exists v \in neigh(s) : \mathrm{dist}(v,n) < \mathrm{dist}(v,s) + \mathrm{dist}(s,e) + \mathrm{dist}(e,n) \tag{5}$$

$$\mathrm{dist}(n,d) < \mathrm{dist}(n,s) + \mathrm{dist}(s,d) \ . \tag{6}$$

In the following, I give the relation between simple LFA and Remote LFA, in particular I show the consequence that can be drawn if one of them exists.

Thesis 3.3. *[C2, J3] I have proven the following equivalence conditions for basic LFA and rLFA in both link- and node-protecting cases:*

- *For an arbitrary node $u \in \mathrm{rLFA}(s,d)$ and $u \in \mathrm{neigh}(s) \Rightarrow u \in \mathrm{LFA}(s,d)$, where $\mathrm{neigh}(s)$ is the set of nodes which are directly connected to node s.*
- *For an arbitrary node $u \in \mathrm{rLFA}(s,d)$ and $u \in \mathrm{neigh}(s) \Leftrightarrow u \in \mathrm{LFA}(s,d)$ under the assumption that costs are uniform.*

5.4.2 Node-protecting Case

Similarly to the case of link protection, I have reformulated the notations of P-space and Q-space in terms of shortest path distances, which lead to the following alternative sufficient and necessary condition for node-protecting rLFA:

Thesis 3.4. *[J3] For source s, destination d, and next-hop e, some $n \neq s, d$ is a node-protecting Remote LFA for the $s-d$ pair (i.e., $n \in \mathrm{rLFA}_{\mathrm{NP}}(s,d)$) if and only if*

$$\mathrm{dist}(s,n) < \mathrm{dist}(s,e) + \mathrm{dist}(e,n) \tag{7}$$

$$\mathrm{dist}(n,d) < \mathrm{dist}(n,e) + \mathrm{dist}(e,d) \ . \tag{8}$$

Again, (7) in essence states that for source s and next-hop e some $n \in V$ is in $\mathcal{P}_{\mathrm{NP}}(s,e)$. Moreover, one can easily observe that the condition of $\mathcal{P}_{\mathrm{NP}}$ remains the same as for $\mathcal{P}_{\mathrm{LP}}$. Furthermore, the statement of (8) is the basic loop-free criterion of node-protecting LFAs.

I have also reformulated the notation of "extended P-space" in terms of shortest path distances for the case of node protection, which again brings us to the following alternative sufficient and necessary condition for extended node-protecting rLFA.

Thesis 3.5. *[J3] For source s, destination d, and next-hop e, some node $n \neq s, d$ is an extended node-protecting Remote LFA for the $s-d$ pair if and only if*

$$\exists v \in \mathrm{neigh}(s) : \mathrm{dist}(v,n) < \mathrm{dist}(v,e) + \mathrm{dist}(e,n) \tag{9}$$

$$\mathrm{dist}(n,d) < \mathrm{dist}(n,e) + \mathrm{dist}(e,d) \ . \tag{10}$$

Note that the reformulated conditions in Thesis 3.1, Thesis 3.2, Thesis 3.4, and Thesis 3.5 are true for any arbitrary weighted network.

In the next two theses, I compare the rLFA coverages attainable with plain and extended rLFAs.

Thesis 3.6. *[C2, J3] I have shown that extended Remote LFA provides complete protection against single link failures in any unit cost network.*

In an arbitrary 2-edge-connected graph with unit link costs, $\mu_{\mathrm{LP}} = 1$ if and only if for each $(u,v) \in E, u$ has a $\mathrm{rLFA}_{\mathrm{LP}}$ to v and vice versa. We know that (u,v) is contained in at least one chordless cycle, which has a length of k. One can easily observe that if k is odd, then $\mathrm{PQ}_{\mathrm{LP}}$-space is not empty, while if k is even, then by means of extended P-space possible repair tunnel endpoints can be still obtained.

In the following, I show that Thesis 3.6 does not apply for node-protection.

Thesis 3.7. *[J3] I have shown that, in general, extended remote LFA does not necessarily provide complete protection against single node failures, not even in unit cost networks.*

Consider the case of Fig. 5.4.2 where the next-hop c went down and node s wishes to send a packet to node d. The shortest path between them is denoted by the solid arrow. The potential repair tunnel endpoints are in $\mathrm{PQ}_{\mathrm{NP}}$-nodes, which is empty in this case. Unfortunately, this remains the case even if using the $\mathcal{P}^{\mathrm{e}}_{\mathrm{NP}}$ would be an option. This means that there are networks that cannot be 100% protected against node failures by nor "plain" neither "extended" Remote LFA.

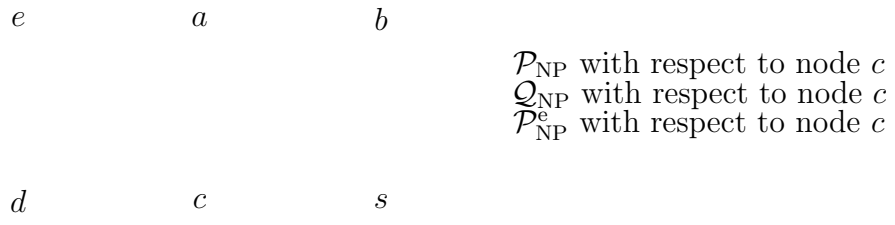

Figure 3: Extended P-space cannot guarantee full protection against single node failures

5.5 Lower Bounds and Network Optimization for Remote LFA

In the following, I return to the case of "plain" Remote LFAs and I suppose that only this option is available in seeking possible repair tunnel endpoints.

My first goal is to give a graph-theoretical characterization of rLFA coverage. In particular, my aim is to identify the attainable lower bounds of plain rLFA failure coverage against both link and node failures in *unit cost* networks. Along the way, I give some methods to easily calculate $\mu(G)$ in different families of graph topologies notable in building resilient networks.

As it turns out, there are some networks in which Remote LFA failure coverage is particularly low. Therefore, my second goal is to define special *network optimization* algorithms, purposed at attaining full rLFA coverage in every network by augmenting the network with suitably chosen new links.

Seeking lower bounds is a recurring scheme in computer science. In our case, a properly chosen worst-case graph, i.e., one over which the rLFA coverage either for link or node failures is particularly poor, will exhibit the pathologic cases that should be avoided in the network optimization phase. In line with this argumentation, I have extended the worst case protection analysis for LFA coverage given in [C2] to rLFA. In particular, I sought graphs G for which $\mu_{\mathrm{LP}}(G)$ or $\mu_{\mathrm{NP}}(G)$ is minimal.

Thesis 4. *I have found that in certain unit cost networks remote LFA coverage can be as low as* 33% *for single link failures, and* 0% *for single node failures. I have also proposed heuristic algorithms to approximately solve the rLFA Graph Extension problem for both single link and node failures. I have given extensive simulations indicating that in the case of link protection on average* 3.6 *new links are enough to attain* 100% *rLFA coverage, in case of node protection on average* 4.05 *links are necessary, while for extended node-protecting rLFA this number is* 3.3.

5.5.1 Link-protecting Case

First, I show how low the rLFA coverage could be in 2-edge-connected networks, then I continue my analysis with 2-node-connected networks.

Thesis 4.1. *[C2, J3] I have shown that for any* $k \geq 1$ *there is a 2-edge-connected graph* G *on* $n = 3k + 1$ *nodes with* $\mu(G) = \frac{1}{3}$.

As a proof, I show that the so called "4-propeller graph" (P_k, see Fig. 4(e)) attains this limit. Thus, consider a propeller graph with k number of blades. One can see that the nodes

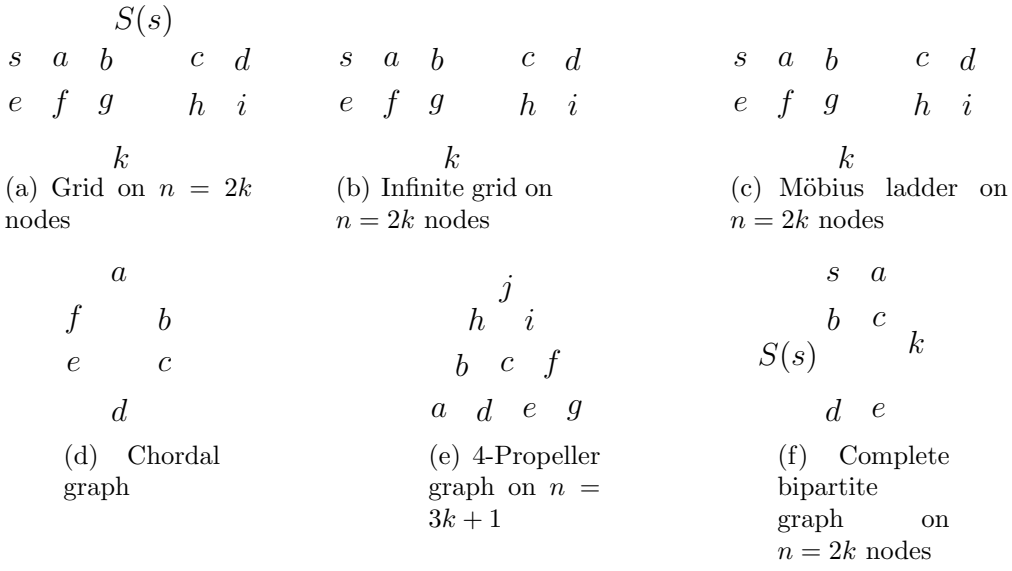

(a) Grid on $n = 2k$ nodes

(b) Infinite grid on $n = 2k$ nodes

(c) Möbius ladder on $n = 2k$ nodes

(d) Chordal graph

(e) 4-Propeller graph on $n = 3k + 1$

(f) Complete bipartite graph on $n = 2k$ nodes

Figure 4: Illustration topologies

on the pitch of the propeller blades have rLFAs to every destination except the neighbors, since they are on an even cycle. Nodes on the side of the blades considered as sources can only protect adjacent link failures if the nodes in the face of them are considered as destinations. Finally, the node in the middle has Remote LFAs only for destination nodes situated on the pitch of the blades. Thus, $\mu(G) = \frac{1}{3}$.

Next, I turn to the lower bounds of 2-node-connected networks. The following thesis concludes the results:

Thesis 4.2. *[C2, J3] I have found that for any $k > 2$ there is a 2-node-connected graph G on $n = 2k$ nodes with $\mu_{\mathrm{LP}}(G) = \frac{k-1}{2k-1} < 0.5$.*

As a proof, I show that grids (G_k (see Fig. 4(a))) and complete bipartite graphs ($K_{k,k}$, see Fig. 4(f)) attain this limit. In grids, every source-destination pair, where the destination is the neighbor of the source, or the destination is situated on the same side (denoted by $S(s)$), cannot be protected. It is easy to see, that every node is in a 4-cycle wherein neighbors as destinations are not protectable and the shortest paths to every node on the same side traverses one of the neighbors. Thus, such nodes are unprotected. Similar is the case for $K_{k,k}$ as well. Each destination, which is on the same side of the source is protected. However, every other node that is the neighbor of the source, due to the property of bipartite graphs that every cycle is even, cannot be protected either.

5.5.2 Node-protecting Case

In the following, I show that, contrary to link-protecting case where $\mu_{\mathrm{LP}}(G)$ was lower bounded by a constant, for the node-protecting case $\mu_{\mathrm{NP}}(G)$ can become arbitrarily small, and this can be attained even in a not so complicated network topology. As mentioned

in Section 3, during my Remote LFA analysis the node protection between an arbitrary neighboring node pair is considered as undefined. Therefore, I only took into account graphs in which at least one non-adjacent node pair exists (i.e., non-complete graphs). Even in these graphs, the question is only interesting when single node failures, at least theoretically, can be repaired, so I had to focus only on 2-node-connected graphs.

The following thesis concludes the results:

Thesis 4.3. *[J3] I have found that for any $n \geq 4$, there is a 2-node-connected graph G on n nodes with $\mu_{\mathrm{NP}}(G) = \frac{2(n-3)}{n^2-5n+6} \leq \frac{4}{n}$.*

Again, as a proof I show a particular graph on n nodes, hereafter denoted by $\mathcal{L}_n$, that attains this limit. An example for $\mathcal{L}_n$ for the case when $n = 6$ is depicted in Fig. 5. The

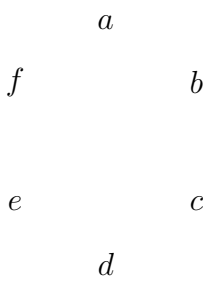

Figure 5: Worst-case graph for $\mathrm{rLFA}_{\mathrm{NP}}$ on $n = 6$ nodes

main topological characteristic of $\mathcal{L}_n$ is that there is one node on the top with degree of $n-1$, there are two nodes with degree of 2, while the remaining $n-3$ nodes have a degree of 3. Correspondingly, the number of non-adjacent source-destination pairs is $n^2 - 5n + 6$. One easily sees, in addition, that only those node pairs can be protected that are in opposite in the 4-cycles. The number of such node pairs equals twice the number of 4-cycles in the graph (i.e., $n-3$), and therefore there are $2(n-3)$ protected node pairs. Consequently, we have $\mu_{\mathrm{NP}}(\mathcal{L}_n) = \frac{2(n-3)}{n^2-5n+6}$. Observe that, in the limit, this bound tends to zero, meaning that in very large $\mathcal{L}_n$ graphs the fraction of rLFA node-protected source-destination pairs diminishes.

I have established using numerical evaluations that these are indeed lower bounds for every graph with $n < 10$, and I have conjectured that these are general lower bounds.

5.6 Remote LFA Graph Extension

As it was shown, there exist arbitrary large networks with particularly poor rLFA coverage. Therefore, the question immediately arises what extent we need to intervene at the graph topology to improve coverage to 100% in both link- and node-protecting cases. This problem is important since *(i)* this would answer how "far" are poorly protected networks from perfect rLFA failure coverage and *(ii)* would provide an easy way for operators to boost the protection in their networks.

Table 2: Brief results of Remote LFA Graph Extension for link-protecting case

Topology	η_{LP}	$Gr_{\eta_{\mathrm{LP}}}$	μ_{LP}	$Gr_{\mu_{\mathrm{LP}}}$	η_{NP}	$Gr_{\eta_{\mathrm{NP}}}$	μ_{NP}	$Gr_{\mu_{\mathrm{NP}}}$	μ^{e}_{NP}	$Gr^{e}_{\mu_{\mathrm{NP}}}$
AS1221	0.833	1	0.833	1	0.083	3	0.083	1	0.083	1
AS1239	0.898	6	1	0	0.658	16	0.843	1	0.928	1
AS1755	0.889	4	1	0	0.704	7	0.912	1	1	0
AS3257	0.946	3	0.954	1	0.521	20	0.702	5	0.866	3
AS3967	0.864	7	0.969	1	0.715	10	0.896	2	0.994	1
AS6461	0.919	2	1	0	0.505	8	0.596	3	0.747	2
Abilene	0.56	6	0.833	1	0.608	3	0.725	2	0.872	1
AT&T	0.823	6	0.888	2	0.565	12	0.684	4	0.849	2
Deltacom	0.542	79	0.885	4	0.436	113	0.818	9	0.868	9
Geant	0.646	20	0.827	4	0.411	30	0.676	5	0.74	5
Germany	0.695	1	0.882	1	0.599	8	0.77	2	0.955	2
InternetMCI	0.877	3	0.888	2	0.558	9	0.837	3	0.916	1
Italy	0.784	12	0.951	2	0.574	24	0.839	3	0.926	2

5.6.1 Problem Formulation

I adapted the LFA Graph Extension problem from [27] to *rLFA Graph Extension* problem, which asks for augmenting the network with a suitably chosen new link.

Definition 5.5. *$rLFAGraphExtensionLP(G)$: Given a graph $G(V,E)$, find the smallest subset F of the complement edge set $\overline{E}$ of G such that $\mu_{\mathrm{LP}}(G(V,E\cup F)=1$.*

Similarly, in the case of node protection this definition can be formalized as follows:

Definition 5.6. *$rLFAGraphExtensionNP(G)$: Given a graph $G(V,E)$, find the smallest subset F of the complement edge set $\overline{E}$ of G such that $\mu_{\mathrm{NP}}(G(V,E\cup F)=1$.*

5.6.2 Numerical Evaluation

Since all previously studied LFA network optimization problem turned out NP-complete and so there were no viable optimal algorithms that could solve the problems in viable time, I have immediately turned to approximation algorithms in this case as well. I have defined a family of heuristics to solve $rLFAGraphExtensionLP(G)$ and $rLFAGraphExtensionNP(G)$, respectively, which implement a greedy approach and simulated annealing based methods as well. In all cases, however, the greedy method produced the best results. Furthermore, since extended P-space cannot guarantee full protection against single node failures, I have taken into account this option as well.

The following thesis concludes the results:

Thesis 4.4. *[C2, J3] I have developed a family of fast and efficient heuristics for solving $rLFAGraphExtension(G)$ problems. By extensive numerical evaluations, I have found that for more than 50% of the examined unit cost networks only 2 new links are needed to attain 100% $\mathrm{rLFA}_{\mathrm{LP}}$ coverage, and on average the number of new links is only 3.6. In the node-protecting case, on average 4.05 additional links are necessary, which drops to 3.3 in case of using extended rLFA.*

Succinct results can be found in Table 2. For the link-protecting case, η_{LP} is the initial link-protecting LFA coverage, $Gr_{\eta_{\mathrm{LP}}}$ denotes the number of new links added by the LFA

Graph Extension algorithm to reach 100% link-protecting LFA coverage. μ_{LP} is the initial link-protecting rLFA coverage, $Gr_{\mu_{\mathrm{LP}}}$ marks the results of the rLFA Graph Extension algorithm. For the node-protecting case, η_{NP} and μ_{NP} are the initial node-protecting LFA coverage and rLFA coverage, respectively; while column $Gr_{\eta_{\mathrm{NP}}}$ and $Gr_{\mu_{\mathrm{NP}}}$ mark the number of new links added by the LFA Graph Extension algorithm and rLFA Graph Extension algorithm, respectively. The last two columns cover the node-protecting case, where extended rLFA was also considered.

The first observation is that there were some networks that were fully protected with link-protecting rLFA right away, even without the need of any graph extension. Second, the number of links that have to be added to reach full coverage with link-protecting rLFA is much less than when only simple LFA capable routers are present. Analogously, in case of node protection much less additional link are needed for 100% rLFA coverage than when only simple node-protecting LFAs are only available. The results indicate that for more that 50% of the networks only 2 new links are necessary to attain 100% link-protecting rLFA coverage, and on average the number of new links is only 3.6, while in case of simple LFA this number is 14.5. In case of node-protection, if extended rLFA is considered then on average only 3.3 new links are necessary to reach 100% rLFA coverage. Note that this number is nut much greater (4.05) in case of plain rLFA.

6 Application of Results

I believe that my results may have a significant impact not just in an academic setting but for the industry as well. I hope that my results may contribute to the general understanding of the failure case coverage provided by LFAs, and may help hesitating network operators to reach a verdict of using any of the available approaches. In particular, my research in LFA based optimization was conducted with Ericsson Research, Hungary, TrafficLab as an industrial partner, and my optimization algorithms were implemented and used in an internal application with graphical user interface in which network operators can analyze and optimize their network(s).

Regarding my Remote LFA analysis, I have been in contact with the IETF to make Remote LFA draft as an RFC. My alternative definitions of P-,Q- and extended P-spaces in terms of costs (Eq. 3, Eq. 4 and Eq. 5) provided an easier understanding whether a node is an rLFA or not[4], since the basic loop-free conditions of simple LFA [3] were also defined in terms of costs. Moreover, I believe that analyzing Remote LFA first in the literature and showing its advantages and network optimization approaches will give an increasing push on service providers to operate the Internet without any interruption and indeed win the trust of most of the potential users.

[4] `http://tools.ietf.org/html/draft-ietf-rtgwg-remote-lfa-06#page-25`

Acknowledgements

First of all, I would like to thank my supervisor, Gábor Rétvári for all the time and energy he has put into supervising and guiding my research. His encouragement, guidance and support were indispensable to becoming a researcher in the field of telecommunications. He taught me a lot of things in the way of thinking, developing and presenting. I want to highlight that without his useful comments on all of my paper preparations, I would not have been able to write any substantial research paper. He always gave me directions and thoughts when I felt myself lost.

I am also indebted to Zalán Heszberger, who was my supervisor in my M.Sc. years and helped me make my way towards being a PhD student. During my PhD studies, he was always able to help me when I was a bit trembled.

I would like to thank to my colleagues and roommates, Zoltán Fehér, Balázs Lajtha, Attila Kőrösi, Márton Csernai, Péter Babarczi, Felicián Németh, Balázs Sonkoly, András Gulyás and all the people whose names would not fit on this paper for all their help and the excellent social events we spent together.

My work was carried out in the Department of Telecommunications and Media Informatics at Budapest University of Technology and Economics, in the Hungarian research group, called MTA-BME Lendület and in the High Speed Networks Laboratory (HSNLab), therefore I would like to thank to János Tapolcai, Róbert Szabó, Attila Vidács, Sándor Molnár for their support and advices and to Erzsébet Gyõri for all her help.

Last but not least, I would like to thank to my parents for supporting me through all my student years, and made it possible to me to concentrate on my studies. I also would like to thank to my brother and to all my relatives for their love and patience, without the atmosphere of such a great family, my dreams would have never come true.

Publications

Journal Papers

[J1] **L. Csikor**, M. Nagy, and G. Rétvári, "Network Optimization Techniques for Improving Fast IP-level Resilience with Loop-Free Alternates," Infocommunications Journal, vol. 3, iss. 4, pp. 2-10, 2011. ($6/2 = 3$)

[J2] **L. Csikor**, G. Rétvári, and J. Tapolcai, "Optimizing IGP Link Costs for Improving IP-level Resilience with Loop-Free Alternates," Computer Communications Journal, Special Issue on Reliable Network-based Services, vol. 36, iss. 6, pp. 645-655, 2013. ($6/2 = 3$)

[J3] **L. Csikor** and G. Rétvári, "On Providing Fast Protection with Remote Loop-Free Alternates," Telecommunication Systems Journal, vol. 60, iss. 4, pp. 485-502, 2015. ($6/1 = 6$)

Conference Papers

[C1] G. Rétvári, **L. Csikor**, J. Tapolcai, G. Enyedi, and A. Császár, "Optimizing IGP Link Costs for Improving IP-level Resilience," in Proc. International Workshop on Design Of Reliable Communication Networks (DRCN), Krakow, Poland, pp. 62-69, 2011. (winner of Best Paper Award) ($3/4 = 0.75$)

[C2] **L. Csikor** and G. Rétvári, "IP Fast Reroute with Remote Loop-Free Alternates: the Unit Link Cost Case," in Proc. RNDM, pp. 16-22, 2012. ($3/1 = 3$)

Book Chapters

[B1] **L. Csikor**, G. Rétvári and J. Tapolcai. "High Availability in the Future Internet", Future Internet Assembly 2013: Validated Results and New Horizons, Lecture Notes in Computer Science, *The Future Internet*, vol. 7859, pp. 64-76, 2013. ($6/2 = 3$)

Other Publications

[OC2] **L. Csikor** and Z. Fehér. "Exploring Hidden Relations in Moving Human Groups". In *Proc., POSTER 2010*, Prague, Czech Republic, 6. May 2010. ($3/2 = 1.5$)

[OC1] **L. Csikor** and Z. Fehér. "Information Spreading in Self Oragnizing Mobile Networks". In *Proc., MACRo Conference 2010*, Tirgu Mures, Romania, 14-15. May 2010. ($3/2 = 1.5$)

[OJ1] P. Babarczi, F. Tanai, **L. Csikor**, J. Tapolcai and Z. Heszberger. "Útvonalválasztás késleltetés-toleráns hálózatokban". *Híradástechnika*, vol. 66, no. 1, pp. 23-31, 2011. ($1/5 = 0.2$)

[OC3] F. Németh, B. Sonkoly, A. Gulyás, **L. Csikor**, J. Tapolcai, P. Babarczi and G. Rétvári. "Improving resiliency and throughput of transport networks with OpenFlow and Multipath TCP: Demonstration of results over the Géant OpenFlow testbed". *Open Networking Summit*, Flyer and Demo, Santa Clara, USA, Apr. 2013.

[OC4] F. Németh, B. Sonkoly, **L. Csikor**, and A. Gulyás, "A Large-Scale Multipath Playground for Experimenters and Early Adopters," in ACM SIGCOMM (DEMO), Hong Kong, China, pp. 482-483, 2013.

[OC5] B. Sonkoly, F. Németh, **L. Csikor**, L. Gulyás, and A. Gulyás, "SDN based testbeds for evaluating and promoting multipath TCP," in Proc. IEEE International Conference on Communications (ICC), pp. 3044-3050, June 2014. (3/5 = 0.6)

[OC6] A. Csoma, B. Sonkoly, **L. Csikor**, F. Németh, A. Gulyás, W. Tavernier and S. Sahhaf "ESCAPE: Extensible Service ChAin Prototyping Environment using Mininet, Click, NETCONF and POX". In *Proc., ACM SIGCOMM 2014*, Demo, Chicago, Aug. 2014.

[OC7] A. Csoma, B. Sonkoly, **L. Csikor**, F. Németh, A. Gulyás, D. Jocha, J. Elek, W. Tavernier and S. Sahhaf "Multi-layered Service Orchestration in a Multi-Domain Network Environment". In *Proc., EWSDN 2014*, Demo, Budapest, Sep. 2014.

References

[1] A. Atlas. U-turn alternates for IP/LDP fast-reroute. Internet-Draft, draft-atlas-ip-local-protect-uturn-03, February 2006.

[2] C. Alaettinoglu, V. Jacobson, and H. Yu. Towards milli-second IGP convergence. Internet-Draft, draft-alaettinoglu-isis-convergence-00.txt, IETF, (downloaded: Oct 2013), 2000.

[3] A. Atlas and A. Zinin. Basic specification for IP fast reroute: Loop-Free Alternates. RFC 5286, 2008.

[4] A. Basu and J. G. Riecke. Stability issues in ospf routing. In *ACM SIGCOMM*, pages 225–236, San Diego, CA, USA, August 2001.

[5] S. Bryant, C. Filfils, S. Previdi, M. Shand, and N. So. Remote LFA FRR. Internet-Draft, Dec. 2012.

[6] S. Bryant, M. Shand, and S. Previdi. IP fast reroute using Not-via addresses. Internet-Draft, March 2010.

[7] L. Csikor, M. Nagy, and G. Rétvári. Network optimization techniques for improving fast IP-level resilience with Loop-Free Alternates. *Infocommunications Journal*, 3(4):2–10, 2011.

[8] L. Csikor, G. Rétvári, and J. Tapolcai. Optimizing IGP link costs for improving IP-level resilience with loop-free alternates. *Computer Communications*, Sep 2012.

[9] Levente Csikor. GML 2 LGF converter. http://csikor.tmit.bme.hu/GML2LGF.

[10] N. Feamster and H. Balakrishnan. Detecting bgp configuration faults with static analysis. In *NSDI*, pages 43–56, 2005.

[11] C. Filsfils, P. Francois, M. Shand, B. Decraene, J. Uttaro, N. Leymann, and M. Horneffer. Loop-free alternate (LFA) applicability in service provider (SP) networks. RFC 6571, June 2012.

[12] P. Francois, M. Shand, and O. Bonaventure. Disruption-free topology reconfiguration in ospf networks. In *IEEE INFOCOM*, Anchorage, USA, May 2007. INFOCOM 2007 Best Paper Award.

[13] G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot. Analysis of link failures in an IP backbone. In *ACM SIGCOMM Internet Measurement Workshop*, pages 237–242, 2002.

[14] ISO. Intermediate ststem-to-intermediate system (is-is) routing protocol. ISO/IEC 10589, 2002.

[15] Simon Knight, Hung X. Nguyen, Nick Falkner, Rhys Bowden, and Matthew Roughan. The internet topology zoo. *Selected Areas in Communications, IEEE Journal on*, 29(9):1765–1775, 2011.

[16] A. Kvalbein, A. F. Hansen, T. Čičic, S. Gjessing, and O. Lysne. Fast IP network recovery using multiple routing configurations. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pages 1–11, 2006.

[17] K.-W. Kwong, L. Gao, R. Guerin, and Z.-L. Zhang. On the feasibility and efficacy of protection routing in IP networks. In *INFOCOM, long version is available in Tech. Rep. 2009*, University of Pennsylvania, 2010.

[18] C. Labovitz, A. Ahuja, and F. Jahanian. Experimental study of Internet stability and backbone failures. In *FTCS*, pages 278–285, 1999.

[19] S. Litkowski, B. Decraene, C. Filsfils, and K. Raza. Operational management of loop free alternates. Internet-Draft, draft-litkowski-rtgwg-lfa-manageability-01, February 18, 2013.

[20] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. Inferring link weights using end-to-end measurements. In *ACM IMC*, pages 231–236, 2002.

[21] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, Y. Ganjali, and C. Diot. Characterization of failures in an operational IP backbone network. *IEEE/ACM Transactions on Networking*, 16(4):749–762, 2008.

[22] A. Markopoulou, G. Iannacone, S. Bhattacharyya, C.-N. Chuah, and C. Diot. Characterization of failures in an IP backbone. In *Proc. IEEE Infocom*, Mar. 2004.

[23] M. Menth, M. Hartmann, and D. Hock. Routing optimization with IP Fast Reroute. Internet-Draft, July 2010.

[24] J. Moy. OSPF version 2. RFC 2328, Apr. 1998.

[25] M. Nagy, J. Tapolcai, and G. Rétvári. Optimization methods for improving IP-level fast protection for local shared risk groups with loop-free alternates. *Springer Telecommunication Systems Journal*, 2012.

[26] G. Rétvári, L. Csikor, J. Tapolcai, G. Enyedi, and A. Császár. Optimizing IGP link costs for improving IP-level resilience. In *Proc. of DRCN*, pages 62–69, Oct. 2011.

[27] G. Rétvári, J. Tapolcai, G. Enyedi, and A. Császár. IP Fast ReRoute: Loop Free Alternates revisited. In *INFOCOM*, pages 2948–2956, 2011.

[28] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb. A case study of OSPF behavior in a large enterprise network. In *IMC*, pages 217–230, 2002.

[29] SNDLib. Survivable fixed telecommunication network design library. http://sndlib.zib.de, downloaded: Apr. 2012.

[30] H. T. Viet, P. Francois, Y. Deville, and O. Bonaventure. Implementation of a traffic engineering technique that preserves IP Fast Reroute in COMET. In *Rencontres Francophones sur les Aspects Algorithmiques des Telecommunications, Algotel*, 2009.

[31] J. Wang and S. Nelakuditi. IP fast reroute with failure inferencing. In *ACM SIGCOMM Workshop on Internet Network Management – The Five-Nines Workshop*, 2007.

[32] D. Watson, F. Jahanian, and C. Labovitz. Experiences with monitoring OSPF on a regional service provider network. In *ICDCS*, pages 204–212, 2003.