

IP Fast Reroute with Remote Loop-Free Alternates: the Unit Link Cost Case

Levente Csikor, Gábor Rétvári

High Speed Networks Laboratory, Department of Telecommunications and Media Informatics

Budapest University of Technology and Economics

Email: {csikor, retvari}@tmit.bme.hu

Abstract—Up to not so long ago, Loop-Free Alternates (LFA) was the only viable option for providing fast protection in pure IP and MPLS/LDP networks. Unfortunately, LFA cannot provide protection for all possible failure cases in general. Recently, the IETF has initiated the Remote Loop-Free Alternates (rLFA) technique, as a simple extension to LFA, to boost the fraction of failure cases covered by fast protection. Before further standardization and deployment, however, it is crucial to determine to what extent rLFA can improve the level of protection in a general IP network, as well as to find optimization methods to tweak a network for 100% rLFA coverage. In this paper, we take the first steps towards this goal by solving these problems in the special, but practically relevant, case when each network link is of unit cost. We also provide preliminary numerical evaluations conducted on real IP network topologies, which suggest that rLFA significantly improves the level of protection, and most networks need only 2 – 3 new links to be added to attain 100% failure case coverage.

Index Terms—IP Fast ReRoute, Remote Loop-Free Alternates, link protection, heuristics, unit link costs

I. INTRODUCTION

In recent years, high availability has become an important factor in operational networks, since this is required by the increasing number of real-time applications (VoIP, IPTV, online-gaming). The aim is to recognize the failure and reroute the packets rapidly avoiding the failed component. The matter of this is, how fast a router can recognize a failure and can reroute packets avoiding the failed component.

Formerly, the intra-domain routing protocols (Open Shortest Path First or Intermediate System To Intermediate System) handled the failures. The failure information was distributed throughout the network in order to notify each router to recalculate shortest paths with the failed component removed from the topology. This process can take between 150 ms and a couple of seconds, depending on network size and routers' shortest path calculation efficiency [1], [2]. Clearly, this recovery time is beyond what real-time applications require.

Therefore, the IETF defined a framework, called IP Fast ReRoute (IPFRR [3]), for native IP protection in order to reduce failure reaction time to tens of milliseconds in an intra-domain, unicast setting. In order to achieve this goal, the IPFRR techniques are based on *local rerouting* and *pre-computed detours* [3]. This allows instant reaction to the failure and enables the routing protocol to converge in the background.

In the past few years, many IPFRR proposals have appeared to solve this problem. Unfortunately, the majority of them require additional management burden, complexity, non-standard IP forwarding functionality [4]–[9] to existing routing protocols, evading the possibility to be eventually applied in commercial routers.

There is one method, called Loop-Free Alternates (LFA) [10] which made its way into commercial routers [11], [12]. LFA is simple, standardized and already implemented. However, it has a significant drawback: it does not guarantee protection for all possible failure cases, due to strong dependency on actual topology and link costs. Extensive simulations and numerical studies have shown that LFA can only protect 75–85% of the link failures and 50–75% of the node failures, respectively.

To improve the level of fast protection provided by LFA without any modifications to the method or to the underlying network topology, the IETF has published a generalization of LFA called Remote LFA (rLFA) Fast ReRoute framework [13]. This method is an extension to the basic LFA that provides additional backup connectivity when none can be provided by the basic mechanisms.

However, even if it provides higher failure coverage there still exist networks, which are not sufficiently protected by rLFA. Nevertheless, as of now there is no information available about how it performs in different network topologies, what are the fundamental lower and upper bounds on failure case coverage, or how this can be improved [14].

In this paper, we make the first steps in this direction. As a first approach, we shall limit our attention to the special case when link costs are uniform. Our earlier studies on LFA [14] showed that the efficiency of LFA in protecting most failure scenarios crucially depends on both the graph topology and the link costs of the underlying network. Unfortunately, it turned out that it is extremely difficult to consider both at the same time, due to the complexity of the related graph theoretical questions. Therefore, it has proven beneficial to study graph topological concerns separately from the effects of link costs. In the present paper, we follow the same course: first, we initiate the analysis for remote LFA in graphs with unit costs, and in a subsequent study we shall attempt to generalize our results to arbitrary weighted graphs. Considering unweighted graphs is fruitful for a number of further reasons, for instance, this case is highly relevant in real-world networks and, as

shall be shown, results for LFA can only be generalized to rLFA if costs are uniform. Finally, we also found this problem particularly appealing from a theoretical point of view.

In the first part of this paper, we provide the first ever basic graph theoretical toolset for analysing rLFA failure case coverage and we establish a sufficient and necessary condition for a unit cost network to have 100% rLFA failure coverage. We study the “bad cases” for rLFA in which coverage is particularly poor and we shall show that there are 2-node-connected graphs in which rLFA cannot protect more than 50% of the possible failure scenarios, while in 2-edge-connected networks this can go down to 33%. Furthermore, we disclose the connectivity between LFA and rLFA and show the conclusions that can be drawn if information about one of them exists. Building on [13] we distinguish the term of *plain* and *extended remote LFA*, and we show that the usage of extended rLFA makes any network 100% protected, in case of uniform link costs. However, if only plain rLFA capable routers are present, complete failure coverage cannot be attained in real-world network topologies.

Our analysis shows that many practically important graph topologies do not admit 100% rLFA failure coverage, especially with plain rLFA. Recently, LFA network optimization methods were proposed [14]–[16] to optimize certain aspects of the network to obtain maximal failure coverage. These works, however, are restricted to LFA. In the second part of the paper, we generalize these methods to rLFA, in particular, we study the problem of optimizing a network topology for better rLFA protection and we introduce an algorithm for modifying the network, by adding the smallest number of new links, to improve coverage to 100%.

At last but not least, we give a numerical evaluation of the general protectability of real-world ISP networks. We study the performance of our network optimization algorithm in real topologies. As shall be shown, it responds much better than simple LFA, as in numerous cases less than two additional links were necessary to achieve 100% failure coverage. Moreover, we find that some networks have full protection without any modifications.

The rest of the paper is organized as follows: Section II and Section III gives an overview of remote LFA and provides a useful mathematical model. Section IV and Section V are devoted to a graph theoretical remote LFA failure coverage analysis of different notable networks and Section VI discusses the remote LFA graph extension problem where numerical results of many real-world network topologies are described as well. Finally, in Section VII we conclude our work and sketch future research directions.

II. ASSUMPTIONS AND MODEL FORMULATION

In Loop-Free Alternates, the backup routes (repair paths) are precomputed and installed in the router as the backup for the primary paths. Once a router detects a link or adjacent node failure, it switches to the backup path to avoid traffic loss.

Remote LFA allows the backup next-hop¹ to be more than one hop away. It means, that after a failure an adjacent node recognizes it and tries to find a (remote) node whose shortest path to the destination is not affected by the failed component. If such node is found then packets will be forwarded to it. Remote LFA relies on tunnels to provide additional logical links towards backup next-hops. After the remote node receives the package it sends it towards the primary destination. Note that the tunnelled traffic is restricted to shortest paths just like “normal” traffic, hence the tunnel must avoid the failure as well.

Perhaps the easiest way to understand remote LFA, and how it differs from basic LFA, is through an example. Consider network depicted in Fig. 1 and suppose that router s wishes to send a packet to a destination d .




Figure 1. A sample network topology with uniform link costs. Solid lines mark the IP network topology, while black dashed line marks the tunnel

The next-hop of s along the shortest path towards d is a . If, however, the link (s, a) fails, then node s has to find an alternative neighbor to pass on the packet to. It cannot send the packet to, say b , as b has an ECMP (Equal Cost Multiple Path) to destination d and, as it does not know about the failure, it can send the packet back to s causing a loop. Therefore, s has no neighbour that would not pass the packet back to it if chosen as a bypass, so in this case the given source-destination pair cannot be protected via standard LFA. However, if a tunnel is created between s and e (marked by black dashed line in Fig. 1, then e , now being a direct neighbor of s , would become an LFA for d , thereby protecting the link (s, a) .

Consequently, when a link cannot be entirely protected with local LFA neighbors, the protecting router seeks the help of a remote LFA staging point. Note that this tunnel is only used as a detour, so it does not affect the normal flow of traffic in any ways. There are numerous tunnelling mechanisms which fulfill the requirements of this design. In an MPLS/LDP (Multiprotocol Label Switching-Label Distribution Protocol [17]) enabled network, for instance, a simple label stack can be used to provide the required tunnel without any additional packet header modification or any possible IP packet fragmentation due to increased packet size.

Next, consider node d' as destination and link (s, b) as the failed link. Then (s, b) cannot be protected for a lack of a suitable tunnel since all nodes, whose shortest path does not go through (s, b) , can only be reached from s through (s, b) itself. For a formal definition, see the next section. This suggests that while the use of rLFA definitely can provide higher protection

¹In routing, the address of the next device along the shortest path to a destination is called *next-hop*

level than pure LFA, it still does not facilitate full protection for all failure cases in a general topology.

Our mathematical model for studying rLFA is as follows. We model the network topology by a *simple, undirected graph* $G(V, E)$, with V being the set of nodes and E the set of edges. Let $n = |V|$ and $m = |E|$, and denote the complement edge set with \bar{E} . We assume that links are bidirectional and point-to-point. As mentioned earlier, we further assume that each link in G is of the same unit cost, as this assumption allows us to study the purely graph theoretical aspects of rLFA separately from the effect of link weights. In a subsequent paper, we plan to relax this assumption. Furthermore, we presume that each node has a well-defined next-hop towards each destination even if more than one equal cost shortest paths exist. Since, according to [18], most unplanned failures in an operational IP network are transient link failures (about 70%), we shall limit our attention to single link failures exclusively. Moreover, since an arbitrary link can only be protected if the graph describing the network is *2-edge-connected*, we also assume this minimum topological requirement. In this paper, we use the notation $\text{dist}(u, v)$ for any $u, v \in V$ to describe the length of the shortest path from u to v . Let $\text{neigh}(s)$ denote the set of nodes which are the neighbors of an arbitrary node s . Furthermore, $\text{LFA}(x, y)$ denotes the set of nodes protecting the (x, y) source-destination pair.

During a link failure, the repair tunnel endpoint needs to be a node in the network reachable from the source without traversing the failed component. In addition, the repair tunnel endpoint needs to be a node from which packets will normally flow towards their destinations without being attracted back to the failed component. Therefore, the set of routers which can be reached from the source without traversing the failed link is termed the *P-space* [19] of the source 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 usage of extended P-space may enable the source router to reach potential repair tunnel endpoints that were otherwise unreachable. The set of routers from which the destination can be reached, without traversing the failed link is termed the *Q-space* of the destination respect to the failed link (see Fig. 1 and consider s as the source and d as the destination).

The intersection of the P-space of s and the Q-space of d defines the viable repair tunnel endpoints, known as *PQ nodes*, which are practically the remote LFAs. As can be seen, for the case of the example network depicted above there is only one node protecting the link (s, a) . However, considering d' as the destination the P-space and Q-space turn out otherwise. Now, there is no intersection of s' P-space and Q-space of d' , thus viable PQ nodes do only exist if extended P-space is used.

Therefore, in this work we slightly diverge from the specification [13] and we say that a node is remote LFA if it is in the intersection of the "simple" P-space and Q-space

and we shall use the term "*Extended remote LFA*" henceforth when extended P-space is also considered for defining PQ nodes. "Plain" rLFA can be easily implemented and deployed since it does not require any profound modifications to the forwarding plane, while extended rLFA requires sophisticated functionality. Thus, we expect, implementations to provide only the plain rLFA initially.

In the rest of the paper, $\text{rLFA}(x, y)$ denotes the set of nodes which protects source x and destination y with remote LFA under the assumption that only "simple" P-space is used.

From the above discussion, it is clear that in general not all nodes have LFA or even remote LFA protection to every other node. To measure rLFA coverages in a graph G , we adopt the simple metric from [10]:

$$\mu(G) = \frac{\#\text{rLFA protected } (s, d)\text{pairs}}{\#\text{all } (s, d)\text{ pairs}} \quad (1)$$

For LFA, the coverage $\eta(G)$ can be defined in a similar way.

III. PRELIMINARIES

An arbitrary failed link along the shortest path between a source and a destination can only be protected if the intersection of P-space of the source and the Q-space of the destination is not empty. At first, we show an alternative characterization for rLFAs that, as shall be seen, is more amenable to theoretical analysis, if we reformulate this requirement in terms of the shortest path distance function dist .

Observation 1. *For each source s and next-hop e some $n \in V$ is in $\text{P-space}(s, e)$ if and only if*

$$\text{dist}(s, n) < \text{dist}(s, e) + \text{dist}(e, n) \quad , \quad (2)$$

and for each source s and destination d , some $n \in V$ is in $\text{Q-space}(s, d)$ if and only if

$$\text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, d) \quad . \quad (3)$$

One can easily see, that (3) is the basic loop-free criterion of link-protecting LFAs [10], while (2) means that the repair tunnel cannot traverse the failed link. The notion of extended P-space could also be expressed with distance functions:

Observation 2. *For each source s and next-hop e , some $n \in V$ is in the extended $\text{P-space}(s, e)$ if and only if $\exists v \in \text{neigh}(s) : \text{dist}(v, n) < \text{dist}(v, s) + \text{dist}(s, e) + \text{dist}(e, n)$.*

It should be noted that the conditions above hold for arbitrary weighted graphs as well.

Next, we formulate an important corollary of the previous observations. In particular, we show that if an arbitrary node on the shortest path between a source and a destination is rLFA protected, then every further node along that shortest path is protected as well.

Lemma 1. *Let (s, d) be a source-destination pair and let q be a node along the default shortest path from s to d . If $\text{rLFA}(s, q) \neq \emptyset$, then $\text{rLFA}(s, d) \neq \emptyset$.*

Proof: Suppose that $\exists n \in rLFA(s, q)$ and consider the source-destination pair (s, d) . For n to be $n \in rLFA(s, d)$, it has to fulfill Obs. 1. First, it satisfies (2) for (s, d) since P-space does not depend on the destination node. Additionally, it only needs to satisfy (3), notably $\text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, d)$. We know that $\text{dist}(n, q) < \text{dist}(n, s) + \text{dist}(s, q)$ and due to triangle inequality² $\text{dist}(n, d) \leq \text{dist}(n, q) + \text{dist}(q, d)$. Therefore, $\text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, q) + \text{dist}(q, d) \Rightarrow \text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, d)$. ■

An important consequence of Lemma 1 is the simple observation that a graph has full rLFA protection, if and only if each node has an rLFA to each of its next-hops.

Corollary 1. *Let G be a graph with unit link costs. Then, $\mu(G) = 1$, if and only if for each $(u, v) \in E$, u has an rLFA to v and v has an rLFA to u .*

Next, we show that there is a deep connection between basic LFAs and remote-LFAs in unit cost networks.

Theorem 1. *Let $G(V, E)$ be a graph with unit link costs. If an arbitrary node $u \in rLFA(s, d)$ and $u \in \text{neigh}(s)$, then $u \in LFA(s, d)$.*

Proof: First, verify the forward direction. It is easy to observe that $u \in rLFA(s, d)$ implies u is in the Q-space, which precisely coincides with the condition for u to be a link-protecting LFA. Second, if $u \in LFA(s, d)$ then u fulfills (3) and since link costs are uniform, when $u \in \text{neigh}(s)$ then $u \in rLFA(s, d)$ as well because from an arbitrary node a to $b \in \text{neigh}(a)$ the default shortest path traverses the neighboring link. ■

IV. ANALYSIS OF EXTENDED REMOTE-LFA

First, we show that extended rLFA (usage of extended P-space) ensures 100% failure coverage in every network.

Theorem 2. *Let G be an arbitrary 2-edge-connected graph with uniform link costs and suppose that remote-LFA can use the extended P-space option. Then, $\mu(G) = 1$.*

Proof: We show that for each edge $(u, v) \in E$, u has an rLFA to v (and vice versa). This will mean that every node has an rLFA to each of its next-hops, which guarantees $\mu(G) = 1$ by Corollary 1. Since G is 2-edge-connected, we know that (u, v) is contained in at least one chordless cycle. Let the length of this cycle be k . If k is odd, then the single node at distance $\frac{k-1}{2}$ from v along the cycle is an rLFA to u . If, on the other hand, k is even, then the P-space $(u, (u, v)) \cap$ Q-space (u, v) is empty. Observe, however, that the single node of distance $\frac{k}{2}$ from u is contained both in Q-space (u, v) and the extended P-space $(w, (u, v))$, where w is the neighbor of u other than v along the cycle, and so it is an rLFA in terms of the extended P-space option. This completes the proof. ■

²The triangle inequality states that for any triangle, the sum of the lengths of any two sides must be greater than or equal to the length of the remaining side. It is one of the defining properties of the distance function which is used in shortest path routing.

Consequently, if remote-LFA implementations support extended P-space, then unit cost networks have full failure coverage. This may be an important factor to consider by an operator willing to deploy rLFA and to an IP device vendor to implement extended rLFA in its router products. If using the extended P-space is not an option due to the implementation complexity or the lack of support in the routers, however, then we can attain substantially lower protection as shall be seen in the next section.

V. ANALYSIS OF “PLAIN” REMOTE LFAS

In this section, we turn to plain remote LFAs i.e. we consider only standard P-spaces and Q-spaces. We give a graph-theoretical characterization of rLFA coverage, as measured by $\mu(G)$. Our main aim is to identify the attainable *lower and upper bounds* of plain rLFA failure coverage. We describe some methods to easily calculate failure coverages in different notable graph topologies depending only on the number of nodes. Nevertheless, we generalize previous propositions of LFAs from [14]–[16] to rLFA as well.

A. Graphs with good coverage

Network operators facing with the challenge of deploying rLFA need to ask the question, whether their current network topology is amenable to rLFA or not. Therefore, it is crucial to separate graph topologies that are “good” for rLFA (i.e., the ones with $\mu(G) = 1$) away from those that attain a particularly low coverage. First, we characterize the good cases for rLFA.

Theorem 3. *Let G be an undirected, simple graph with uniform link costs. Now, $\mu(G) = 1$, if and only if for each $(i, j) \in E : \exists n \neq i, j$ so that $\text{dist}(i, n) = \text{dist}(j, n)$.*

Proof: The result comes from applying (2) and (3) directly to (i, j) . ■

Further notable graph topologies with 100% failure coverage include chordal graphs [20], infinite grids and “Möbius ladder”. The latter two are similar to grids with the difference that first and last nodes are directly connected.

B. LFA worst-case graphs with rLFA

In the following, we turn to discuss *lower bounds* for rLFA, that is, we seek worst-case graphs in which coverage is particularly poor. It has been observed previously that quintessential worst-case graphs for IPFRR are rings, i.e. cycle graphs in which all nodes are of degree two [6], [21].

Consequently, we consider odd rings first, and then we shall treat even rings. Before that, we repeat a previous proposition from [14], which proved the lower bounds of LFA, measured by $\eta(G)$:

Proposition 1. *For an even ring on n nodes $\eta(G) = \frac{1}{n-1}$, and for an odd ring on n nodes $\eta(G) = \frac{2}{n-1}$.*

Next, we generalize these results to rLFA. In fact, we shall do a bit more, as our analysis will account for the length of

the repair tunnel, which is an important factor in provisioning remote-LFA³.

Theorem 4. *Let C_n be an odd ring on n nodes with $n \geq 3$, and let $1 \leq k \leq \frac{n-1}{2}$ denote an upper bound on the length of the tunnel from the source node to its rLFA. Then, $\mu(C_n) = \frac{2k}{n-1}$.*

Proof: Consider a ring topology on n nodes, n odd, let $(s, d) \in E$ be a *neighboring* source-destination and suppose that the link between them went down. In this case s needs to find a possible remote loop-free alternate since it cannot use its other neighbor because it will pass back the packet. Thus, the possible repair tunnel endpoints are situated on the other side of the ring regarding to the failed link, i.e. if an arbitrary node $u \in \text{rLFA}(s, d)$ then $\text{dist}(s, u) \leq \frac{n-1}{2}$ which is tight if $d \in \text{neigh}(s)$. One can observe that if maximal tunnel length is permitted i.e. $k = \frac{n-1}{2}$ then such kind of repairing node always exists ($\mu(C_n) = \frac{\frac{n-1}{2}}{\frac{n-1}{2}} = 1$). However, if the tunnels need to be shorter than an arbitrary node u can only be an rLFA is $\text{dist}(s, u) \leq \frac{n-1}{2} - l$, where l is the tunnel shortening coefficient, i.e. the greater the l , the shorter the tunnel. Trivially, shortening the tunnel with l dissolves the protection among $\forall(s, d)$ pairs, where $\text{dist}(s, d) = l$. Therefore, rLFA failure coverage can be modified as follows: $\mu(C_n) = \frac{\frac{n-1-2l}{2}}{\frac{n-1}{2}}$. Now, consider $\text{dist}(s, u) \leq k$, where k represents the length of the tunnel. In this manner $l = \frac{n-1}{2} - k$ meaning that $\mu(C_n) = \frac{\frac{n-1-n+1+2k}{2}}{\frac{n-1}{2}} = \frac{2k}{n-1}$. ■

Note that $k = 1$ means that only neighboring nodes can be used as repair tunnel endpoints, which essentially corresponds to LFA. In this case, Theorem 4 yields the same result as Prop. 1 stated for LFA for odd rings.

Theorem 5. *Let C_n be an even ring on n nodes with $n \geq 4$, and let $1 \leq k \leq \frac{n-2}{2}$ denote an upper bound on the length of the tunnel from the source node to its rLFA. Then, $\mu(G) = \frac{2k-1}{n-1}$.*

Proof: Consider a ring on n nodes, n even, and suppose that link between an arbitrary *neighboring* (s, d) source-destination pair went down. According to the case of odd ring, s need to pass the packet to the other side of the ring, however, the possible repair tunnel endpoints cannot be reached without traversing the failed component. Thus, for $\forall(s, d)$ pairs, where $d \in \text{neigh}(s)$: the link (s, d) cannot be protected. One can observe, if $\text{dist}(s, d) \geq 2$ then tunnels, avoiding the failed link exist. Therefore, for an arbitrary source s has remote LFAs to $\forall d$ destination excluding its neighbors ($\mu(C_n) = \frac{n(n-3)}{n(n-1)}$). However, assuming shorter tunnels results that for possible $u \in \text{rLFA}(s, d)$: $\text{dist}(s, u) \leq \frac{n}{2} - l$, where l is a shortening coefficient as it was in the case of odd rings. Now, $l = \frac{n}{2} - k$ meaning that $\mu(C_n) = \frac{\frac{n-1-n+2k}{2}}{\frac{n-1}{2}} = \frac{2k-1}{n-1}$. ■

As before, supposing $k = 1$ results the corresponding statement in Prop. 1 for LFA for even rings. In this regard, rLFA can be seen as a natural generalization of LFA.

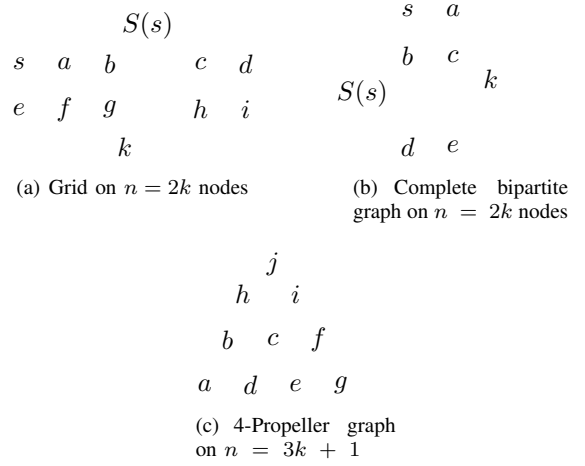


Figure 2. Illustration topologies for Theorem 6 and Theorem 7

C. rLFA worst-case scenarios: 2-node-connected graphs

Below, we continue our analysis towards finding 2-node-connected graphs with low rLFA failure coverage. In what follows, we suppose that there is no constraint on the length of the rLFA tunnel.

Since the simplest 2-node-connected network with low failure coverage is a 4-cycle ($\mu(C_4) = \frac{1}{3}$), we examined notable networks that contain a large number of 4-cycles as subgraphs. We considered the networks depicted in Fig. 2(a) where k denotes the number of 4-cycles, and Fig. 2(b) where k marks the number of node pairs. The following theorem concludes the results:

Theorem 6. *For any $k > 2$ there is a 2-node-connected graph G on $n = 2k$ nodes with $\mu(G) = \frac{k-1}{2k-1}$.*

As a proof, we show that grids (G_k) and complete bipartite graphs ($K_{k,k}$) attain this limit. In grids, $\forall(s, d)$ pairs: $d \in \text{neigh}(s)$ or $d \in S(s)$ cannot be protected, where $S(s)$ denotes the set of nodes situated on the same side. 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 according to Lemma 1.

Similar is the case for $K_{k,k}$ as well. Each $d \in S(s)$ are protected while $\forall d' \notin S(s)$ are neighbors of s and, due to the property of bipartite graphs that every cycle is even, neighbors cannot be protected either.

D. 2-edge-connected graphs

Theorem 7. *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, we show that the so called “4-propeller graph” (P_k) attains this limit. Thus, consider (P_k) depicted in Fig. 2(c) where k denotes the number of blades. One can see that the nodes on the pitch of the propeller blades have remote LFAs to every destination except the neighbors, since they are on an even cycle. Nodes on the side of the blades considered as

³See the `remote-lfa maximum-cost` option on [22]

Table I
LOWER BOUNDS IN WORST-CASE GRAPHS ON n NODES

n	$\mu(G)_{2e}$	B_l	$\mu(G)_{2n}$	B_l
3	1	1	1	1
4	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
5	$\frac{2}{5}$		$\frac{2}{5}$	
6	$\frac{2}{5}$		$\frac{2}{5}$	$\frac{2}{5}$
7	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{3}{7}$	
8	$\frac{19}{56}$		$\frac{3}{7}$	$\frac{3}{7}$
9	$\frac{1}{3}$		$\frac{31}{72}$	

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,

$$\mu(G) = \frac{k(3k-1)+2k+k}{3k(3k+1)} = \frac{k+1}{3k+1} = \frac{1}{3}.$$

So far, we have sought for graphs with low rLFA coverage. Our aim has been to give a tight characterization for the lower bound on $\mu(G)$ for any unweighted graph G . At the moment, we do not have clear answers to these intriguing but hard graph-theoretical questions. What we could prove, however, is that in certain 2-node-connected unweighted graphs $\mu(G)$ can be as low as $\frac{1}{2}$, and in 2-edge-connected graphs an even lower threshold of $\frac{1}{3}$ is also realizable. So far, we have not been able to identify any 2-node-connected or 2-edge-connected graph with smaller rLFA coverage. Thus, we conjecture that $\frac{k-1}{2k-1}$ is an actual lower bound on $\mu(G)$ for 2-node-connected graphs, while $\frac{1}{3}$ is a lower bound on $\mu(G)$ for 2-edge-connected graphs. However, we do not have any formal proofs of these claims as of now.

E. Computational study

It turned out that finding a universal lower bound on rLFA coverage is a hard problem. Clearly, a computational approach might be instructive to support or refute our conjectures. Hence, we generated all non-isomorphic networks on n nodes where $n \in \{1, 2 \dots 9\}$. Note that the generation is very time consuming even if only non-isomorphic graphs are created. Table I summarizes the lower bounds with the following notations: n denotes the number of nodes, $\mu(G)_{2e}$ and $\mu(G)_{2n}$ notes the failure coverage in case of 2-edge-connected and 2-node-connected networks, columns marked by B_l denote the conjectural lower bounds. It can be seen that until $n \leq 4$ results are the same, and if $n \geq 5$ coverages start to increase. One can observe that in the case of $n = 7$ the given failure coverage equals to the coverage attained by 4-propeller graphs mentioned above. It also shows that lower bounds of 2-edge-connected networks are the lowest.

VI. REMOTE-LFA GRAPH EXTENSION

As observed, there exist a lot of graphs with small rLFA coverage. Hence, in this section, we ask to what extent we need to intervene at the graph topology to improve coverage to

100%. This approach is important since (i) this would answer how “far” are poorly protected networks from perfect rLFA coverage, and (ii) would provide an easy way for operators to boost the protection in their networks. We adapt the formal description of the LFA graph extension problem from [14] to rLFA as follows:

Definition 1. *rLFA graph extension problem:* Given a graph $G(V, E)$, find the smallest subset F of the complement edge set \overline{E} of G such that $\mu(G(V, E \cup F)) = 1$.

At the moment, we do not know the complexity of this problem but, based on our former experience with similar network optimization problems for LFA, it seems highly probable that it is also NP-complete. Therefore, we adopt the greedy heuristic from [14] to the rLFA graph extension problem as follows:

Algorithm 1 Greedy rLFA graph extension for graph $G(V, E)$

- 1: **while** $\mu(G(V, E)) < 1$
 - 2: $(u, v) \leftarrow \operatorname{argmax}_{(i,j) \in \overline{E}} \mu(G(V, E \cup \{(i, j)\}))$
 - 3: $E \leftarrow E \cup \{(i, j)\}$
 - 4: **end while**
-

Theorem 8. *Let $G(V, E)$ be a graph with unit link costs. Alg. 1 terminates with full rLFA coverage regardless of the input graph.*

Proof: Alg. 1 surely terminates when all complement links are added but at this point $\mu(G) = 1$ as the complete graphs has full rLFA coverage. ■

Next, we examine how many links one must add in realistic graphs to achieve full rLFA coverage. We chose existing real-world topologies inferred from Rocketfuel [23] database, SNDLib [24] and Topology Zoo [25]. In all topologies, we set all link costs uniformly to 1. Note, there were some networks where inferred link costs were unit costs. The details are in Table II with the following notations: n is the number of nodes, m is the number of links, $\eta(G)$ is the initial LFA coverage, $\mu(G)$ is the initial rLFA coverage, and *Gr.ext* denotes the number of new links added to achieve full failure coverage.

The first conspicuous observation is that there were five networks, which were fully protected with rLFA even without the need of any graph extension. Second, the number of links that have to be added to reach full coverage is much less than when only simple LFA capable routers are present. Such huge improvements can be seen easily in networks where $\eta(G) < 0.9$ e.g. in the Geant topology. However, in the Deltacom topology installation of 79 new links was necessary to achieve full LFA coverage, while with only 4 additional links full rLFA coverage is attainable. The results also indicate that (i) more than 50% of the networks lend themselves to rLFA extension since the maximum number of links needed is less than 2; (ii) on average 3.6 new links are necessary to attain 100% rLFA coverage while in case of simple LFA this number is 14.5.

Table II
REMOTE LFA GRAPH EXTENSION RESULTS

Topology	n	m	$\eta(G)$	Gr.ext	$\mu(G)$	Gr.ext
AS1221	7	9	0.833	1	0.833	1
AS1239	30	69	0.898	6	1	0
AS1755	18	33	0.889	4	1	0
AS3257	27	64	0.946	3	0.954	1
AS3967	21	36	0.864	7	0.969	1
AS6461	17	37	0.919	2	1	0
Abilene	12	15	0.56	6	0.833	1
Arnes	41	57	0.595	18	0.731	6
AT&T	22	38	0.823	6	0.8875	2
Deltacom	113	161	0.542	79	0.885	4
Gambia	28	28	0.037	16	0.111	8
Geant	37	55	0.646	20	0.827	4
Germ_50	50	88	0.801	22	1	0
Germany	27	32	0.695	1	0.882	1
InternetMCI	19	33	0.877	3	0.888	2
Italy	33	56	0.784	12	0.951	2
NSF	26	43	0.86	9	1	0

VII. CONCLUSIONS

Currently, Loop-Free Alternates is the best choice for providing fast protection in pure IP and MPLS/LDP networks. It is a well-known fact that LFA cannot protect every single failure. In our previous works, we showed that improvements can be made by altering the existing network topology. If modifying the network is not an option, remote LFAs may be a better approach.

As in the case of LFA, the number of failure cases protected by rLFA crucially depends on both the graph topology and the link costs. As it seems difficult to consider both at the same time, we studied graph topological concerns separately from the effect of link costs in this paper. This restriction is plausible as a first approach, and we definitely will generalize our results to weighted graphs.

We showed that, under the unit-cost assumption, “extended P-space” results full rLFA failure coverage in every network. This can be an important pointer for operators, currently in the position to deploy rLFA, on how to actually choose link costs. We gave a sufficient and necessary condition for a unit cost graph to be 100% protectable with rLFA. Then, we studied general lower bounds for rLFA coverage. We found that for 2-node-connected graphs on $2k$ nodes the value $\frac{k-1}{2k-1}$ is realizable by grids and complete bipartite graphs and we confirmed computationally that this is a valid lower bound as long as the number of nodes n is smaller than 10. We also found that for 2-edge-connected graphs, this “conjectured” lower bound is $\frac{1}{3}$. We defined the rLFA graph extension problem as the task to augment an unweighted graph with the fewest new links to obtain 100% rLFA protection. We gave a simple greedy algorithm, and we found that, even in very big real-world ISP topologies, adding only 2-3 new links is enough to attain 100% failure coverage.

In the future, we plan to study further remote-LFA related network optimization questions. For instance, in the unweighted case improving rLFA coverage is possible with modifying link costs as well, which looks another intriguing,

and practically relevant, network optimization problem.

VIII. ACKNOWLEDGEMENT

The authors thank the support of High Speed Networks Laboratory at BME. This project was supported by TÁMOP 4.2.2.B-10/1-2010-0009 and OTKA-PD 104939 grants.

REFERENCES

- [1] C. Labovitz, G. R. Malan, and F. Jahanian, “Internet routing instability,” in *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, 1998, pp. 515–528.
- [2] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, “An approach to alleviate link overload as observed on an IP backbone,” in *INFOCOM*, 2003.
- [3] M. Shand and S. Bryant, “IP Fast Reroute framework,” RFC 5714, Jan 2010.
- [4] S. Lee, Y. Yu, S. Nelakuditi, Z.-L. Zhang, and C.-N. Chuah, “Proactive vs reactive approaches to failure resilient routing,” in *INFOCOM*, 2004.
- [5] Z. Zhong, S. Nelakuditi, Y. Yu, S. Lee, J. Wang, and C.-N. Chuah, “Failure inferencing based fast rerouting for handling transient link and node failures,” in *INFOCOM*, 2005.
- [6] G. Enyedi, G. Rétvári, and T. Cinkler, “A novel loop-free IP fast reroute algorithm,” in *EUNICE*, 2007.
- [7] K. Amund, H. A. Fossellie, T. Čičić, G. Stein, and L. Olav, “Multiple routing configurations for fast IP network recovery,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 2, pp. 473–486, 2009.
- [8] S. Bryant, M. Shand, and S. Previdi, “IP fast reroute using Not-via addresses,” Internet Draft, March 2010.
- [9] G. Enyedi, P. Szilágyi, G. Rétvári, and A. Császár, “IP Fast ReRoute: lightweight Not-Via without additional addresses,” in *INFOCOM Mini-conf*, 2009.
- [10] A. Atlas and A. Zinin, “Basic specification for IP fast reroute: Loop-Free Alternates,” RFC 5286, 2008.
- [11] Cisco Systems, “Cisco IOS XR Routing Configuration Guide, Release 3.7,” 2008.
- [12] Juniper Networks, “JUNOS 9.6 Routing protocols configuration guide,” 2009.
- [13] S. Bryant, C. Filfils, M. Shand, and N. So, “Remote LFA FRR,” IETF DRAFT, Jun. 2012.
- [14] G. Rétvári, J. Tapolcai, G. Enyedi, and A. Császár, “IP Fast ReRoute: Loop Free Alternates revisited,” in *INFOCOM*, 2011, pp. 2948–2956.
- [15] G. Rétvári, L. Csikor, J. Tapolcai, G. Enyedi, and A. Császár, “Optimizing igp link costs for improving IP-level resilience,” in *Proc. of DRCN*, Oct. 2011, pp. 62–69.
- [16] L. Csikor, M. Nagy, and G. Rétvári, “Network optimization techniques for improving fast ip-level resilience with loop-free alternates,” *Infocommunications Journal*, vol. 3, no. 4, pp. 2–10, 2011.
- [17] L. Andersson, I. Minei, and B. Thomas, “Ldp specification,” RFC 5036, Oct. 2007.
- [18] 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.
- [19] S. Bryant, C. Filfils, S. Previdi, and M. Shand, “IP fast reroute using tunnels,” IETF DRAFT, Nov. 2007.
- [20] M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, 2nd ed. Elsevier Science, 2004.
- [21] T. Čičić, “An upper bound on the state requirements of link-fault tolerant multi-topology routing,” *IEEE ICC*, vol. 3, pp. 1026–1031, 2006.
- [22] C. Systems, “Ip routing: Ospf configuration guide, cisco ios release 15.2s - ospf ipv4 remote loop-free alternate ip fast reroute,” downloaded: Apr. 2012.
- [23] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, “Inferring link weights using end-to-end measurements,” in *ACM IMC*, 2002, pp. 231–236.
- [24] SNDLib, “Survivable fixed telecommunication network design library,” <http://sndlib.zib.de>, downloaded: Apr. 2012.
- [25] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, “The internet topology zoo,” <http://www.topology-zoo.org>, downloaded: Apr. 2012.