

Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute?

Michael Mentha^a, Matthias Hartmann^a, Rüdiger Martin^a, Tarik Čičić^b, Amund Kvalbein^c

^aUniversity of Würzburg, Institute of Computer Science, Germany

^bUniversity of Oslo, Department of Informatics, Norway

^cSimula Research Laboratory, Oslo, Norway

Abstract

The IETF currently discusses fast reroute mechanisms for IP networks (IP FRR). IP FRR accelerates the recovery in case of network element failures and avoids micro-loops during re-convergence. Several mechanisms are proposed. Loop-free alternates (LFAs) are simple but cannot cover all single link and node failures. Not-via addresses can protect against these failures but are more complex, in particular, they use tunneling techniques to deviate backup traffic. In the IETF it has been proposed to combine both mechanisms to merge their advantages: simplicity and full failure coverage.

This work analyzes LFAs and classifies them according to their abilities. We qualitatively compare LFAs and not-via addresses and develop a concept for their combined application to achieve 100% single failure coverage, while using simple LFAs wherever possible. The applicability of existing LFAs depends on the resilience requirements of the network. We study the backup path length and the link utilization for both IP FRR methods and quantify the decapsulation load and the increase of the routing table size caused by not-via addresses. We conclude that the combined usage of both methods has no advantage compared to the application of not-via addresses only.

Key words: IP networks, routing, resilience

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

Preprint submitted to Computer Networks

August 3, 2009

1. Introduction

Failures of network elements are common and inevitable in the operation of communication networks [1]. Therefore, resilience mechanisms are required to maintain the connectivity in failure cases. Re-convergence of the routing tables is a simple restoration mechanism in IP networks. It is robust [2], but slow [3]. New emerging services such as voice over IP, virtual private networks for finance, and other real-time business applications require stringent service availability and reliability. Their demand for a very fast reaction to failures led to the development of fast reroute (FRR) techniques where backup paths are available at each intermediate node of a primary path for immediate local failover. For multiprotocol label switching (MPLS) technology, two different FRR approaches have already been standardized [4].

Pure IP networks also need fast resilience. Current IETF drafts and other publications propose various methods for IP FRR [5, 6, 7, 8, 9]. Besides quick failure recovery, IP FRR is helpful to prevent packet loss caused by micro-loops which possibly occur in the routing re-convergence phase of IP networks. Local failure recovery suppresses network-wide failure notification and thereby global re-convergence. This avoids micro-loops for short-lived failures which is a big advantage since 50% of all failures last less than a minute [1, 10]. In case of long-lived failures, IP FRR is useful to gain time for ordered loop-free convergence as suggested in [11]. It is widely believed that IP FRR mechanisms should protect against all probable failures, e.g., all single link failures and possibly also all single node failures. Moreover, fast protection mechanisms should not make difficult situations more critical, in particular, they should not lead to routing loops in case of unanticipated multiple failures.

In this paper we focus on two IP FRR mechanisms: loop-free alternates (LFAs) and not-via addresses. LFAs redirect traffic to neighboring nodes that still have a shortest path towards the destination avoiding the failed element [6]. LFAs are simple but cannot protect all single failures. Some LFAs are able to protect only link failures, others protect also router failures. Some lead to routing loops in case of multiple failures, others are safe. Not-via addresses provide local IP-in-IP tunnels to the next-next-hop (NNHOP) around the failed element [7]. They are more complex. Forwarding tables require additional entries for the new not-via addresses and the associated path calculation implies significantly more computation effort than normal addresses. Tunneling might lead to packet fragmentation due to MTU limitations and it requires decapsulation at the tunnel egress router which possibly reduces its forwarding speed. However, not-via addresses offer

100% failure coverage. Thus, it has been proposed in the IETF to repair failures with LFAs wherever possible and use not-via addresses only for the remaining scenarios [5, 7].

This paper has several contributions. First, we provide a new classification for LFAs with respect to their failure protection capabilities. Second, we discuss the pros and cons of LFAs and not-via addresses. Third, we present a concept for the combined application of LFAs and not-via addresses. Fourth, we study the backup path length and the link utilization for both IP FRR methods and quantify the decapsulation load and the increase of the routing tables caused by not-via addresses. Fifth, we conclude that the combined usage of both methods to achieve 100% single failure coverage has no advantage compared to the application of not-via addresses only.

The paper is structured as follows. Section 2 introduces a new classification of LFAs. Section 3 explains not-via addresses. In Section 4 we qualitatively compare both mechanisms and propose a concept for their combined application to fulfill various resilience requirements. Section 5 presents and interprets the results of our performance evaluation. After a short discussion of related work in Section 6, we summarize our conclusions in Section 7.

2. Classification of Loop-Free Alternates

In this section we review the definition of LFAs and classify them according to their failure protection capabilities.

2.1. Definition of LFAs

We consider a *source node* and a *protected next hop* on a shortest path towards *destination*. Another *neighbor node* of provides a loop-free alternate (LFA) when it has a shortest path towards which does not contain and [6]. If link or node fails, forwards the traffic destined to over instead of, and from the deviated packets take the shortest path towards. Thus, LFA provides at node for destination protection against the failure of link or node. LFAs for each destination are pre-computed and installed in the forwarding information base (FIB) of a router. The RFC 5286 [6] specifies three criteria for LFAs to guarantee different levels of protection quality and loop avoidance. We illustrate these conditions and provide a classification of neighbor nodes as LFAs with respect to their failure protection capabilities.

2.2. Loop-Free Condition (LFC)

We consider source s and destination d in Figure 1. The numbers associated with the links are the link metrics taken into account for shortest path computation. When link $s-p$ fails, packets can only be rerouted over neighbor n . However, this creates a forwarding loop because the shortest path from s to d leads over n . Therefore, n cannot be used as LFA by s to protect the failure of link $s-p$. To avoid such loops, the following loop-free condition (LFC) must be met:

$$\text{dist}(s, n) + \text{dist}(n, d) < \text{dist}(s, p) + \text{dist}(p, d) \quad (1)$$

In Figure 2 both neighbors n and p of source s fulfill this condition with regard to destination d . The example in Figure 1 illustrates that there are certain single link or node failures that cannot be protected by LFAs.




Figure 1: The neighbor n of s cannot be used as LFA towards d because it does not meet the loop-free condition (LFC).

2.3. Node-Protection Condition (NPC)

We consider the failure of node p in Figure 2. When LFAs are installed that meet the LFC, s reroutes traffic to neighbor n where the next hop is again p . p reroutes the traffic back to s and a routing loop occurs. Therefore, n cannot be used as LFA by s to protect the failure of node p . However, p can be used for that objective. A neighbor node n must meet the following node-protection condition (NPC) to protect the failure of a node p :

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

An LFA meeting the LFC only is called link-protecting while an LFA also meeting the NPC is called node-protecting. Since the NPC implies the LFC¹, every node-protecting LFA is also link-protecting, but not vice-versa.

¹ $\text{dist}(s, n) + \text{dist}(n, p) < \text{dist}(s, p)$ – (a) follows from the triangular equation, (b) holds since the shortest path from s to p leads via n .




Figure 2: Only the node-protecting LFA can be used to protect against the failure of node P .

2.4. Downstream Condition (DSC)

We consider source S and destination D in Figure 3. N provides a node-protecting LFA for S . If two nodes P_S and P_N fail simultaneously, S reroutes its traffic to N . N cannot forward the packets either, and reroutes them to D which is a node-protecting LFA for S in that case. Thus, a routing loop occurs. Such loops which can appear in case of multiple failures can be avoided if only LFAs are used that comply with the downstream condition (DSC):

$$\text{dist}(N, D) < \text{dist}(S, D) \quad (3)$$

An LFA fulfilling this condition is called downstream LFA. Allowing only downstream LFAs guarantees loop avoidance for all possible failures because packets get always closer to the destination. In Figure 3, N can be used as downstream LFA for S but not vice-versa which avoids the routing loop in our example. S must use another neighbor – if available – to protect against the failure of N .




Figure 3: Neighbor N is a downstream LFA of S but not vice-versa. The use of only downstream LFAs avoids loops in the presence of multiple failures.

2.5. Equal-Cost Alternates (ECAs)

A special case of LFAs are equal-cost alternates (ECAs), i.e., alternative next hops which provide an alternative path that is not longer than the primary path. An example is depicted in Figure 4. Source S knows several equal-cost paths towards D . If its next hop N_1 fails, it can use any of the remaining equal-cost paths as LFA that do not contain the failed element. Thus, either N_2 or P may be used as ECA and even both may be used at the same time. In particular, if the standard routing uses the equal-cost multipath (ECMP) option, the traffic affected by the failure is equally redistributed over the remaining paths. It is easy to see that ECAs cannot create loops in case of multiple failures as they are always downstream LFAs. They are link-protecting but not necessarily node-protecting (see N_1 in Figure 4). This also shows that downstream LFAs are not necessarily node-protecting.




Figure 4: The equal-cost alternates (ECAs) N_1 and P provide alternate paths of the same length as the primary path. N_1 is just link-protecting while P is node-protecting.

2.6. Classification of LFAs

The conditions above limit the number of neighbor nodes which can be used as potential LFAs and thereby create sets of neighbors with different abilities to protect against failures and to avoid loops. The Venn diagram in Figure 5 partitions the set of neighbor nodes into 7 different categories. Equal-cost alternates (ECAs) are always downstream LFAs (fulfill DSC). Downstream LFAs are always loop-free (fulfill LFC). Some neighbor nodes do not meet any of the corresponding conditions. Thus, the set of ECAs is contained in the set of downstream LFAs which is part of the set of general LFAs which are a subset of all neighbor nodes. The node-protecting property of LFAs is orthogonal to the other conditions. There are representatives for every proposed category. Both neighbor nodes in Figure 4 are ECAs, but only P is node-protecting. N_1 in Figure 2 and N_2 in Figure 3 are both downstream LFAs, but only N_2 is node-protecting. N_1 in Figure 2 is a non-downstream LFA and node-protecting, and examples for non-downstream LFAs

which are not node-protecting can also be constructed. Node 6 in Figure 1 does not meet any condition and cannot be used as LFA.

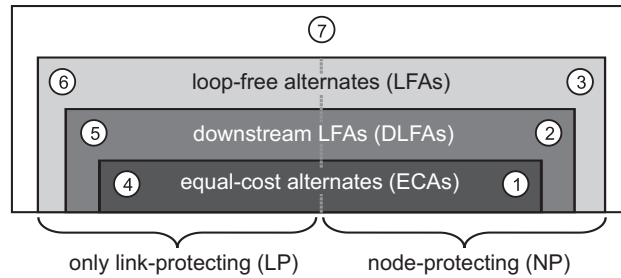


Figure 5: Classification of neighbor nodes with regard to their ability as forwarding alternates to protect failures and to prevent loops.

We order the LFA categories in Figure 5 according to a possible preference for their usage as LFAs (the ultimate preference is the network operator’s decision [6]):


1. node-protecting ECAs,
2. node-protecting downstream LFAs,
3. node-protecting LFAs that do not fulfill the downstream condition,
4. ECAs that are just link-protecting,
5. downstream LFAs that are just link-protecting,
6. LFAs that are just link-protecting and do not fulfill the DSC.

Class (7) contains neighbors that cannot be used as LFAs as they create routing loops.


3. IP Fast Reroute Using Not-Via Addresses

Not-via addresses provide explicit protection tunnels from a source node around a protected next hop (NHOP) towards the next-next hop (NNHOP) that all lie on a primary path from to . This tunnel is implemented using IP-in-IP encapsulation. Figure 6(a) illustrates this concept. The backup path goes from via to where primary and backup paths merge. However, when addresses encapsulated packets to the normal address , they are carried from over to as lies on the shortest path from to . Thus, a mechanism is needed to carry backup traffic from to not via . To that end, a so-called

not-via address “ not via ” (or short:) is introduced and packets destined to that address are never routed over . If NHOP is not reachable from due to a link or node failure, encapsulates packets destined to in another IP packet addressed to NNHOP using its not-via address . The packets are carried from not-via to , decapsulated at , and from there the original packets are further forwarded to .



(a) S detects failure of next hop : packets are encapsulated and carried to not via using the not-via address .



(b) Next hop is the destination, failure of last link : encapsulates packets with address and forwards them to one of its neighbors. From there, they are forwarded to not via avoiding the failed link.

Figure 6: Use of not-via addresses to protect the failure of intermediate nodes and links, and the last link.

IP FRR using not-via addresses requires additional entries in the forwarding tables for not-via addresses. Not-via addresses have the form where can be any node and can be any of its neighbors. Therefore, the number of not-via addresses equals the number of unidirectional links in the network. The forwarding entries for the not-via addresses can be constructed by distributed routing algorithms [7]. Essentially, the path computation for is based on the topology where is removed.

Figure 6(b) shows how not-via addresses can be used to protect the last link, i.e. when the NHOP is already the destination. In contrast to above, node S assumes that only the next link instead of the NHOP has failed; otherwise, the packet could not be delivered anyway. S encapsulates the packet and addresses it towards A . The meaning of A at node S is that the direct link $S \rightarrow A$ must not be used. Instead, the packet is forwarded to a neighbor other than A that passes it on towards A . Since the packet is sent to A , it cannot loop back to S . Finally, A decapsulates the packet and the original packet has reached its destination. If indeed not only link $S \rightarrow A$ but also node A has failed, the packet is discarded as soon as it reaches another neighbor of A .

IP FRR using not-via addresses guarantees 100% failure coverage for single link and node failures unless there is an articulation point in the network that splits the network into two disconnected parts. The concept is very similar to the MPLS FRR facility backup option which installs local bypasses to every NNHOP [12]. However, the backup paths in MPLS may follow explicit routes, therefore, the path layout of MPLS-FRR has more degrees of freedom than the one of IP FRR using not-via addresses.




Figure 7: Unnecessarily long backup paths occur if the tunnel from S to the NNHOP A intersects with the downstream paths from S to D .

Not-via detour paths can be unnecessarily long and waste capacity, but they do not create loops. In the example in Figure 7, packets are normally carried from S to D over P , M , and A . If $S \rightarrow P$ fails, these packets are tunneled to A such that they take the long path $S \rightarrow A \rightarrow M \rightarrow P \rightarrow A \rightarrow D$. However, it is theoretically possible to perform an operation analogous to penultimate hop popping in MPLS. When a packet arrives at a router whose path to D does not traverse the failure and the next hop to D is the same as the next hop to A , the encapsulation can be removed and backtracking can be avoided.




Figure 8: Routing loops can occur if packets are recursively tunneled to not-via addresses in case of multiple failures. Therefore, recursive tunneling to not-via addresses is prohibited.

To prevent routing loops after simultaneous multiple failures, recursive tunneling using not-via addresses is prohibited [7]. In the example in Figure 8, cannot deliver packets to if nodes and fail. encapsulates packets to with the not-via address and sends them to . cannot forward the packets to because also fails. If recursive tunneling was allowed, would encapsulate the packets with the not-via address and return them to creating a routing loop.

4. Comparison of LFAs, Not-Via Addresses, and their Combined Application

In this section, we qualitatively compare LFAs and not-via addresses and propose a concept for their combined application in networks with different resilience requirements.

4.1. Pros and Cons of LFAs and Not-Via Addresses

We discuss LFAs and not-via addresses with respect to various properties that are important for FRR mechanisms.

4.1.1. Backup Path Length

With LFAs, traffic is carried from the LFA directly to the destination along a shortest path. This is different with not-via addresses. They deviate the traffic around the failed element and merge the backup and primary path at the NNHOP. The example in Figure 7 shows that this can lead to unnecessarily long paths. We quantitatively evaluate this issue in Section 5.3.

4.1.2. Failure Coverage

The example in Figure 1 shows that some single link or node failures cannot be protected by LFAs. This is not due to a pathological construction but common observation [13, 14, 15]. In contrast, not-via addresses can protect against all single link and node failures if such failures do not disconnect the network.

4.1.3. Compatibility with Loop-Free Re-Convergence Schemes

The computation of the not-via tunnels can be temporally decoupled from the computation of the basic routing. Thus, during routing re-convergence, the tunnels remain stable making not-via addresses compatible with additional mechanisms for loop-free re-convergence [11, 16]. This is also fulfilled for LFAs. When an LFA is used, it has a path to the destination that does not use the failed network element. During re-convergence to the failure topology, paths can get only longer and, thus, the path from to remains stable.

4.1.4. Protection of Multicast Traffic

Protection of multicast traffic is an issue and currently investigated by the IETF [17]. Not-via addresses deviate the traffic to the NNHOP through tunnels. Thus, the NNHOP can infer the usual interface from the not-via address and run the reverse path forwarding (RPF) check required for multicast traffic correctly [7]. Protection of multicast traffic with LFAs seems complex and is currently not discussed.

4.1.5. Adaptability to SRLGs

The functionality of not-via addresses can be easily adapted to shared-risk link groups (SRLGs). If SRLGs are known, the shortest path computation for the respective not-via address is simply performed in the topology with all elements from the SRLG removed. This is more complicated for LFAs [6] due to the distributed and uncoordinated nature of LFAs.

4.1.6. Complexity of Path Computations

The complexity of backup path computation is in general higher than the complexity of primary path computation because the failure of each link and node must be taken into account. In case of not-via addresses, a router in the network must remove any other node to compute shortest paths trees (SPT) towards the not-via addresses of 's neighbors. Incremental SPT (iSPT) computations reduce this effort that is proportional to the number of nodes in the network to an equivalent of 5 to 13 SPT computations in real world networks with 40 to 400

nodes [7]. For LFAs, the computational cost of determining individual repair paths for all destinations can be very high as well. The computation of ECAs is very easy since ECAs are obtained for free from the usual shortest path calculations. In general, the computational routing complexity and its assessment is hardware- and implementation-dependent.

4.1.7. Forwarding Complexity

In case of protection by LFAs, the FIB of router r provides an LFA for the protection of d against the failure of r . This LFA is used as an alternative next hop. In case of not-via addresses the FIB of router r provides a not-via address n for the protection of d against the failure of r . This not-via address is used for tunneling the packet to its NNHOP when r fails. If the NHOP n is already the destination d , an alternative next-hop is provided to which the packet is forwarded after tunneling towards n . So far, the forwarding complexity of LFAs and not-via addresses is similar. However, LFAs are used only locally while protection by not-via addresses introduces new addresses in the network and the routing tables must hold additional entries for them. There is one not-via address for each uni-directional link in the network. As not-via addresses are used only for local bypass, a router r needs to know a next-hop for a not-via address only if r is on at least one backup path from some other node towards the destination of that not-via address. Entries for all other not-via addresses are basically superfluous at r and could be removed from its routing table. We investigate this further in Section 5.6.

4.1.8. Tunneling Complexity

Not-via addresses fully rely on IP tunneling. Tunneling involves en- and decapsulation of tunneled packets. Encapsulation prepends an additional IP header to the packet. Thus, it leads to increased packet lengths inside tunnels and may result in packet fragmentation due to maximum transmission unit (MTU) limitations. Furthermore, tunneling may have a performance impact on the forwarding speed of routers. Most current hardware can achieve encapsulation without performance degradation. Packet decapsulation at the tunnel endpoint, however, requires two lookup operations. The first one to recognize the tunnel endpoint, the second for further forwarding with the inner IP address. Most modern hardware is designed to perform this also at line rate. On legacy hardware this can slow down the handling of decapsulation traffic to almost half the line rate depending on the router load. So the major disadvantage caused by tunneling stems from packet decapsulation on legacy hardware. In Section 5.5 we study the maximum amount

of traffic that needs to be decapsulated by each node in failure cases. In contrast to protection by not-via addresses, protection by LFAs does not use tunneling.

4.2. Combined Application of LFAs and Not-Via Addresses

We consider three different resilience requirements with different sets of protected elements and different demands for loop avoidance:

- (i) protection against all single link failures,
- (ii) protection against all single link and all single router failures,
- (iii) protection against all single link and all single router failures with loop avoidance in the presence of multiple failures.

Not-via addresses fulfill the strictest resilience requirement (iii). LFAs alone cannot even meet the loosest one because they cannot achieve 100% failure coverage. However, protection by LFAs is simpler than protection by not-via addresses because LFAs do not require new addresses in the network, it does not lead to performance issues due to tunneling, and it possibly leads to shorter backup paths. This motivates the combined application of LFAs and not-via addresses as proposed in the IETF [5, 7]: use LFAs where possible and not-via addresses where needed to achieve 100% failure coverage. As LFAs have different properties, only certain LFA types can be used in the above cases in the following order of preference:

- (i) (1), (4), (2), (5), (3), (6), and not-via.
- (ii) (1), (2), (3), and not-via; (4), (5), and not-via to protect the last link.
- (iii) (1), (2), and not-via; (4), (5), and not-via to protect the last link.

The numbers correspond to the LFA types in Figure 5. Note that only-link-protecting LFAs (type 6) cannot be used for the protection of the last link for (ii) and (iii) since they may create loops if the destination node is down.

5. Analysis of the Combined Application of LFAs and Not-Via Addresses

We study the availability of different LFA types in resilient network structures and illustrate how many of them can be used for the resilience requirements defined in Section 4.2. Then we investigate performance measures for not-via addresses and their combined application with LFAs. We compare the path prolongation on backup paths and the link utilization to those of slow IP restoration. Then, we quantify the decapsulation traffic from not-via tunnels and the minimum number of not-via addresses in the router FIBs.

5.1. Networks under Study

For our analysis we use the topologies of the COST239 [18] and GEANT [19] networks (see Figures 9(a) and (b)). We also examined other networks but these results are not presented here since they do not yield additional insights. For the generation of the traffic matrices we use the method proposed in [20] and enhanced in [21] to generate synthetic traffic matrices resembling real-world data. We use different sets of link weights for the networks. One option is to set all link weights to one and perform simple hop count routing (HC) as often used in unoptimized networks. For the COST239 network, we also use link weights which are inverse proportional to the link capacities (INVCAP) as recommended by Cisco [22]. For the GEANT network, we additionally use the real link weights (REAL) that are based on the inverse metrics with some modifications. They can be obtained from the data of [23]. We perform single shortest path first (SPF) routing. The equal-cost multi path (ECMP) option is not used for our analysis. It has only little influence on our results and it entails several decisions regarding path splitting, backup path splitting, or protection utilization that are not yet specified in the not-via drafts. When multiple equal-cost paths towards a destination are available, the interface with the lowest ID is installed as the active interface as specified for IS-IS [24, Sect. 7.2.7]. When evaluating the mechanisms we consider a set of scenarios \mathcal{S} . We use the set \mathcal{S}_1 that contains the failure-free scenario and all single link failures for the evaluation of resilience requirement (i), and the set \mathcal{S}_3 with the failure-free scenario and all single link and node failures for the evaluation of resilience requirement (iii). The networks under study have heterogeneous link capacities and we scaled the traffic matrices so that the maximum link utilization does not exceed 100% for IP restoration and failure scenarios \mathcal{S} .

5.2. Protection of Destinations by LFAs and Not-Via Tunnels

We evaluate how many destinations can be protected by various LFA types and how many require not-via tunnels for their protection. The LFAs are chosen according to the recommendations given in Section 4 and the results of this analysis depend on the desired resilience levels (i) – (iii) because some LFAs cannot be used for stricter requirements. The results are presented in Figures 10 and 11 for the COST 239 and the GEANT network. The x-axes show the node IDs and the y-axes the percentage of destinations per node that are covered by the respective LFA or not-via tunnels. We label the LFA types according to the classification in Section 2.6. We differentiate between not-via tunnels protecting intermediate nodes and not-via tunnels protecting the last hop (LH) since they are used differently (cf. Section 3).

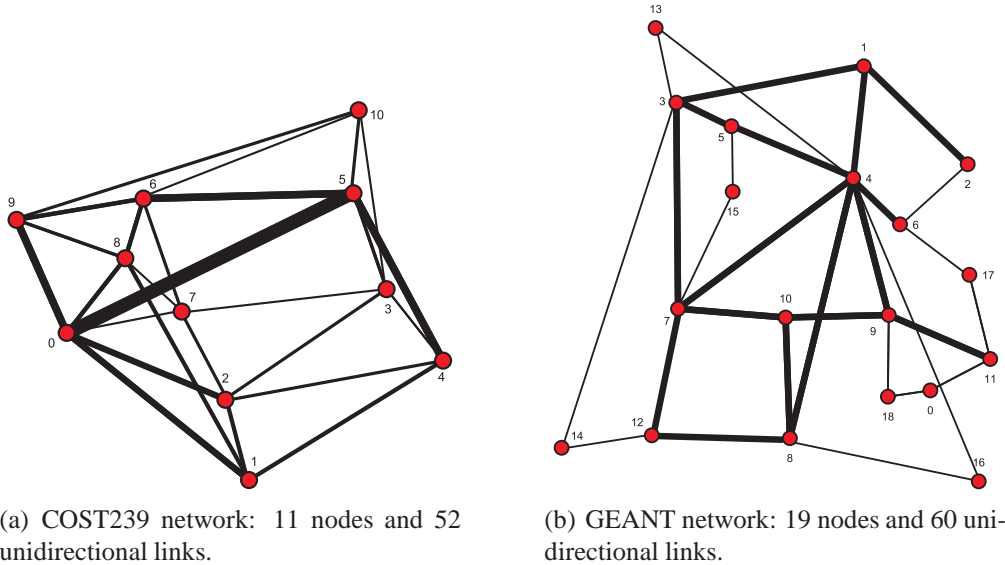



Figure 9: Networks under study.


We observe that for hop count routing only three out of six LFA types exist. This is due to the uniform link cost metric where all the links have the same costs and hence this finding can be generalized to all networks using hop count routing. We briefly explain this finding. ECAs that are only-link-protecting (type 4) do not exist since there are no parallel links. Downstream LFAs (type 2 & 5) other than ECAs do not exist either. The downstream criterion requires that the alternate neighbor is closer to the destination than the deviating node. Since the distance $\text{dist}(u, v)$ from u to its neighbor v is always 1, this can only be true for equal-cost paths. Thus, all downstream LFAs are also ECAs. This has another implication. If loop avoidance in general failure cases is required (iii), LFAs other than ECAs cannot be used in networks that use simple hop count routing (cf. Figures 10(e) and 11(e)).

Now we study the availability of LFAs in the COST239 network. Its topology represents a class of networks whose nodes are well connected among each other. Most nodes can reach any other node within two hops. Figures 10(a) and (b) show the percentage of destinations protected by various LFA and not-via tunnel types when only link protection is required (i). Alternates that can be used at each protection level are placed at the bottom in the figure and not-via addresses at the top. LFAs that must be replaced by not-via addresses at higher protection levels are placed in between. When hop count routing is used (Figure 10(a)), ECAs (type 1)

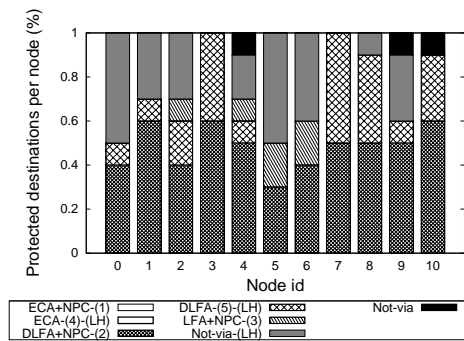


(a) HC: link protection only (i)


(b) INVCAP: link protection only (i)




(c) HC: link + node protection (ii)



(d) INVCAP: link + node protection (ii)



(e) HC: link + node protection - no loops during multiple failures (iii)



(f) INVCAP: link + node protection - no loops during multiple failures (iii)

Figure 10: Protection of destinations by LFAs and not-via tunnels in the COST239 network under different resilience requirements, using different link metrics (hop count (HC), inverse-capacity (INVCAP)).

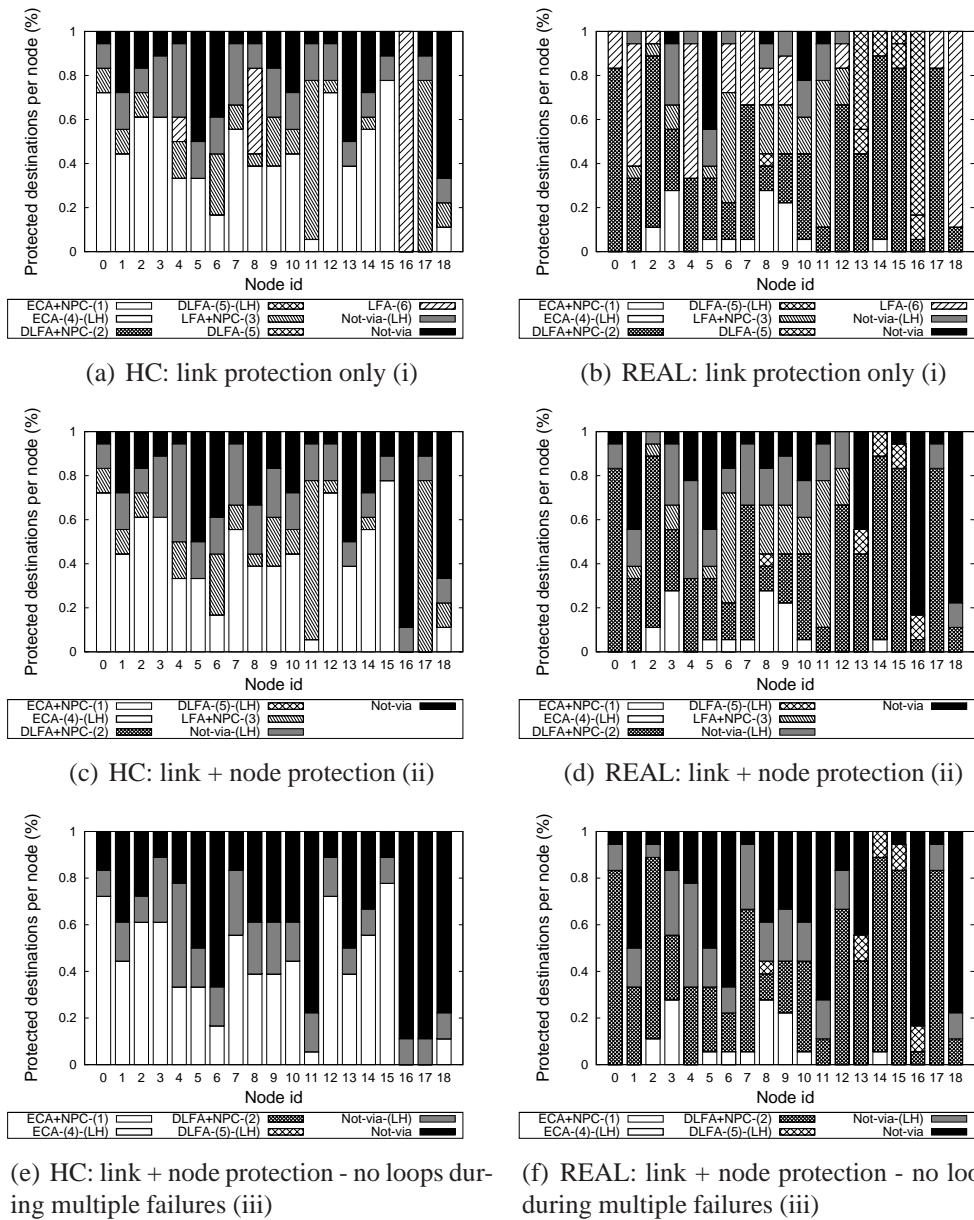


Figure 11: Protection of destinations by LFAs and not-via tunnels in the GEANT network under different resilience requirements, using different link metrics (hop count (HC), real (REAL)).

protect between 20% – 50% of the destinations and node-protecting LFAs (type 3) vary from 0% – 30%. Only-link-protecting LFAs (type 6) are applicable for 40% – 50% of the destinations, mainly to protect the last hops of the relatively short paths. Links belonging to cycles with 3 nodes can be easily protected by only-link-protecting LFAs. Only the links `0-1` and `5-6` are not part of such a cycle and need to be protected by not-via tunnels at node 0 and 5, respectively. When using INVCAP link weights, the high diversity of link capacities leads to a path layout without equal-cost paths. Thus, no ECAs are available here. Between 30% – 60% of the destinations are protected by node-protecting downstream LFAs (type 2). Node-protecting LFAs that do not fulfill the downstream condition (type 3) protect only 0% – 20% of the destinations. The remaining destinations are mainly last-hops and are protected with link-protecting LFAs that comply with the downstream condition (type 5) or not (type 6). Figures 10(c) and (d) show the results when all single link and node failures are protected (ii). Compared to (i), all only-link-protecting LFAs (type 5 & 6) are replaced with not-via tunnels. Figures 10(e) and (f) show the results for link and node protection with general loop avoidance (iii). Now, even node-protecting LFAs (type 3) are not sufficient as the downstream condition must be fulfilled for LFAs, too. Therefore, the node-protecting LFAs (type 3) are replaced by not-via tunnels. For HC routing, now only ECAs and not-via tunnels are in use as already concluded above.

The GEANT network represents a more sparsely connected class of topologies. The paths between node pairs are significantly longer than in the COST239 network since the nodes lie on circles of three to five nodes. Comparing Figure 11(a) with Figure 10(a) we observe for HC routing that even for link protection only (i), many nodes require not-via tunnels to protect intermediate hops as well as last hops. Nodes 4, 8, 16 create the only cycle with 3 nodes in the network, so they are the only nodes having only-link-protecting LFAs. Node 16 is special as it protects all its destinations by only-link-protecting LFAs. When node protection is required (ii), it thus must use not-via tunnels for the protection of all destinations. In a similar way, node 17 protects all its destinations by non-downstream LFAs and requires not-via tunnels for the protection of all destinations when loop avoidance is required (iii). Thus, the existence of suitable LFAs depends on the network topology and the link costs. If routing loops must be avoided in case of single node or multiple other failures, only a fraction of all LFAs can be used. Then, some nodes cannot protect even a single destination by LFAs in certain topologies. Hence, not-via addresses are not only necessary to achieve 100% failure coverage, at some nodes they are the only protection option. When REAL link weights are used, the available protection options are more diverse. In contrast to

the INVCAP weights in the COST239 network, the hand-tuned REAL metrics still provide equal-cost paths (Figure 11(b)). Apart from that, the characteristics of the resulting protection options are similar to those of the COST239 network. In both networks we observe that using alternative link weights instead of HC routing, a larger number of destination can be protected with LFAs. Nonetheless, there are still many destinations that require not-via tunnels for protection.

5.3. Path Prolongation

Delay sensitive applications require short paths also in failure cases. Long backup paths should be avoided and, therefore, the length of the backup paths is an important property that should be analyzed. We assess the path prolongation, i.e. the difference between primary and backup path length and compare IP restoration, protection by not-via addresses only, and their combined application with LFAs. We consider link protection (i) and node protection with loop avoidance (iii) and use S_L and S_{LR} for their evaluation, respectively. Figure 12(a) shows the CCDF of the path prolongation in the GEANT network with HC metrics. The x-axis shows the path prolongation in hops and the y-axis shows the conditional probability that a path affected by a failure increases by more than x hops.

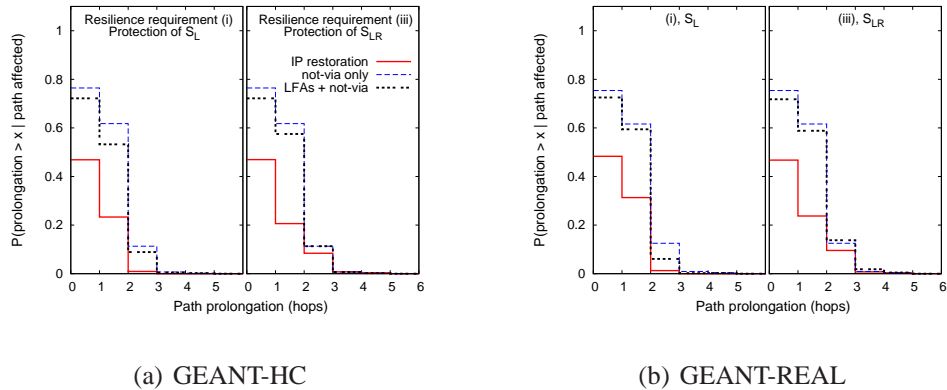


Figure 12: Path prolongation in the GEANT network for resilience requirements (i) and (iii).


The length of the primary path is determined by IP routing which follows the shortest paths, at least when hop count routing is used. IP restoration leads to the shortest backup paths possible and serves as a comparison baseline. For both sets of considered failure scenarios S_L and S_{LR} , about 50% of the backup paths for IP restoration have the same length as their primary paths. This is possible

when IP restoration finds an end-to-end equal-cost path when an element of the primary path fails. In case of IP FRR – no matter which type of protection – only 25% of the backup paths have the same lengths as corresponding primary paths because fewer ECAs are available for local repair at intermediate nodes. Also the probability that backup paths are at least one or two hops longer than primary paths is significantly larger compared to IP restoration. However, there is only a small difference in backup path length between protection by not-via addresses only and by their combined application with LFAs although pathological examples like in Figure 7 lead to significantly longer backup paths for not-via addresses. The difference even decrease for the stricter resilience requirement (iii). The same analysis for the REAL link weights (Figure 12(b)) and for the COST239 network with different link weights leads to very similar results that show hardly any difference in backup path prolongation.


5.4. Maximum Link Utilization

The maximum link utilization U_{max} is the maximum utilization of a link over all considered failure scenarios. We study the maximum link utilizations caused by IP restoration, not-via addresses, and their combined application with LFAs. The traffic matrices (as described in Section 5.1) are normalized so that the maximum of the maximum link utilizations is 1.0 for IP restoration when all single link and node failures are considered. Figures 13(a)–(d) show the fraction of links whose maximum link utilization exceeds a certain utilization value U . The left part of the figures presents an evaluation based on single link failures () and resilience requirement (i) while the right part is based on single link and node failures () and resilience requirement (iii). According to our construction, the maximum utilization value for IP restoration is 1.0 when link and node failures are considered and HC routing is used.


The maximum link utilization for IP FRR mechanisms is clearly larger than the one for IP restoration for a large fraction of links. When HC routing is used, the results for not-via only and for the combined approach are almost identical for both networks and both resilience requirements. But when INVCAP link metrics are used, we observe a very unfavorable effect with LFAs. Especially when only link failures are protected, the link utilizations are higher when not-via addresses are used in combination with LFAs than for single use of not-via. We explain this phenomenon using the COST239 network (Figure 9(a)). With INVCAP link metrics, the small link $l_{1,2}$ has a very high link weight, while all other links adjacent to router r_1 have relatively small link costs. When link $l_{1,2}$, the largest link in




(a) COST239-HC



(b) GEANT-HC



(c) COST239-INVCAP



(d) GEANT-REAL

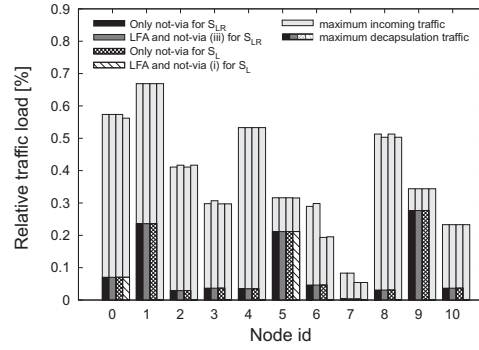
Figure 13: Fraction of links with maximum link utilization greater than a given value for different resilience requirements and resilience mechanisms in the COST239 and the GEANT network.

the network, fails, router r can only use node s as LFA because all other adjacent routers would route traffic back to r . Then, all traffic over r destined to s is using the small link rs as backup link which leads to a (theoretical) link utilization of 275%. This simple example shows that with INVCAP metrics, the only possible LFAs are often those with poor connectivity. In our example networks, this happened only for resilience requirement (i), but examples for requirements (ii) and (iii) can also easily be constructed. In the COST239 network, IP FRR leads to a maximum utilization of 1.85 while in the GEANT network, the maximum value is about 1.85. Thus, IP FRR can lead to heavy traffic concentration and overload on some backup links. Note that utilization values larger than 1.0 are only theoretical and translate into packet loss in real networks. It can also be seen that the intelligent selection of the REAL link weights in the GEANT network leads to very good results, even for IP-FRR.

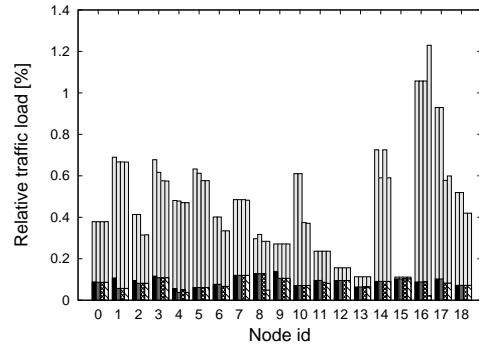
5.5. Decapsulated Traffic from Not-Via Tunnels

In Section 4.1.8 we have argued that tunneling possibly leads to a slowdown of the forwarding speed at the decapsulating router in case of old hardware. Therefore, we investigate the amount of incoming traffic which must be decapsulated from not-via tunnels. We define the *capacity of a node* as the sum of its incoming interface capacities, the *incoming load of a node* is the sum of its incoming traffic rates, and the *decapsulation load of a node* is the sum of its incoming traffic rate in terminating not-via tunnels. The metrics of interest are the node capacities, the maximum incoming traffic load per node, and the maximum decapsulation load per node whereby the maximum is calculated either over r or s . We look at protection by not-via addresses only and their combination with LFAs whereby resilience requirement (i) is used with r and (iii) with s . We calculate the performance metrics and normalize the load and the decapsulation load of a node by its capacity because the traffic rates per node differ by orders of magnitude.


Figures 14(a)–(d) show the results for the COST239 and the GEANT network. The fact that the incoming load of node 16 in Figure 14(b) is larger than 1.0 is theoretical because we do not drop packets and instead allow link utilizations > 1.0 . It shows that the links of this node are heavily overloaded in some failure scenarios. Nodes 1, 5 and 9 of the COST239 network have a maximum decapsulation load between 21% and 29% relative to their capacity in the worst case. All other nodes have values below 10% relative to their capacity. In the GEANT network, all nodes have a decapsulation load of at most 16%, no matter if they carry a lot or rather little other traffic apart from decapsulation traffic. Nodes 12, 13, and 15 in Figure 14(b) are interesting as their major load can consist of decapsulation



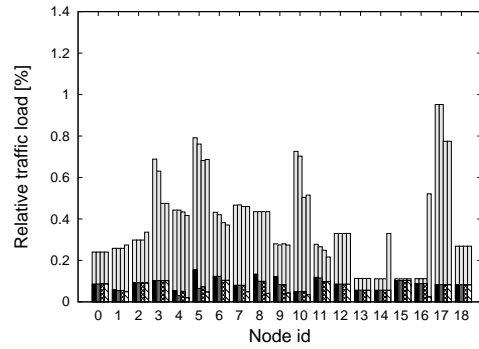
(a) COST239-HC



(b) GEANT-HC



(c) COST239-INVCAP



(d) GEANT-REAL

Figure 14: Amount of decapsulated traffic per node relative to maximum node capacity for COST239 and GEANT.

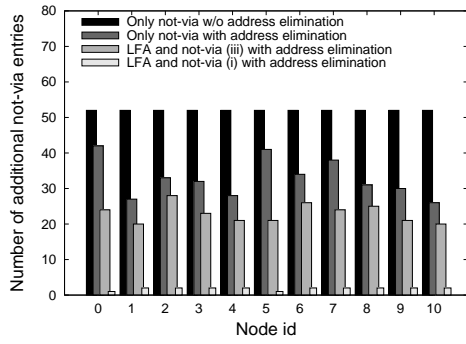
traffic. It shows that a large percentage of a node’s incoming traffic can be subject to decapsulation. This is not a problem in that particular case as only a small percentage of the node’s capacity is used. Looking at the different protection options and resilience requirements, we observe that the maximum decapsulation load is roughly the same for all of them. The COST239 network is an exception if only soft resilience (i) is needed because then most nodes do not need to decapsulate traffic at all. Only very few nodes show a clearly higher decapsulation load when only not-via is used (e.g. node 5 in Figure 14(d)). The conclusion is that the combined application of not-via addresses and LFAs possibly reduces the decapsulated traffic in many scenarios, but it hardly reduces the maximum decapsulation load, at least if strict resilience (iii) is required.

5.6. Number of Not-via Addresses per FIB


Not-via addresses create new entries in the routing and forwarding tables. We evaluate how many of them must be known to each node in particular when not-via addresses are only used to complement LFAs. For each link (u, v) , a not-via address $u \rightarrow v$ is required, i.e., the number of not-via addresses in a network equals its number of unidirectional links. Figures 15(a)–(d) show that each node in the COST239 handles 52 not-via addresses and each node in the GEANT network handles 60 not-via addresses. These numbers of additional not-via addresses are probably not a heavy burden for routing protocols as well as routing and forwarding tables because these entities usually support also many external prefixes. In contrast to ordinary addresses, only nodes along all possible (equal-cost) paths of a specific not-via tunnel can encounter the corresponding not-via addresses. Therefore, this not-via address could be removed from the FIBs of all other nodes. Even when the ECMP option is not used (as in our study here), the decision which equal-cost path is actually taken is not deterministic. Therefore, the not-via addresses must not be removed on all possible paths.

The figures show that the fraction of removable not-via addresses is significant. However, the number of remaining not-via addresses greatly varies among different nodes.


Using LFAs wherever possible and not-via tunnels only when needed further reduces the number of not-via addresses that need to be supported by each node. Resilience requirement (iii) is strict and allows only a few LFA types to be used whereas resilience requirement (i) allows all LFA types to be used. As a consequence, the fraction of remaining not-via addresses per node in Figures 15(a)–(d) is smaller for the combined application (iii) than for protection by not-via addresses only and even smaller for the relaxed resilience requirement (i). With




(a) COST239-HC



(b) GEANT-HC



(c) COST239-INVCAP



(d) GEANT-REAL

Figure 15: Number of additional entries required in the forwarding tables of individual nodes for the COST239 and the GEANT network.

(i) only the not-via addresses \mathcal{N} and \mathcal{L} need to be supported in the COST239-HC network. However, any node lies on a 3-hop shortest path from 0 to 5 (5 to 0) when \mathcal{L} (\mathcal{N}) is removed. Therefore, these not-via addresses create 2 entries in all nodes of the network except for nodes 0 and 5. In contrast, in the GEANT-HC network the combined application of not-via addresses and LFAs saves only a small number of additional not-via addresses per node compared to protection by not-via addresses only when unused not-via addresses are removed from the FIBs. One reason for the low number of removable not-via addresses per node is that a not-via address becomes obsolete only if all the traffic protected by its not-via tunnel can be protected by LFAs that fulfill the resilience requirements (i) or (iii), respectively. When INVCAP or REAL link weights are used (Figures 15(c) and (d)), more destinations can be protected with LFAs (cf. Figures 10 and 11) compared to HC routing and, therefore, more not-via addresses can be eliminated.

Hence, this analysis showed that the number of not-via addresses that need to be supported by each node can be significantly decreased by elimination of unused not-via addresses at individual nodes. However, deciding whether a not-via address is needed or not is not a simple task and requires substantial computation effort. This holds for the elimination of unused not-via addresses when only not-via addresses are used for protection or when they are used in combination with LFAs. As the combined application of not-via addresses and LFAs does not save many entries per node compared to protection by not-via addresses only with elimination of unused not-via addresses, and given the fact that the number of not-via addresses is not dramatically high within a network, we conclude that the number of additional addresses is not a reasonable driver for the combined application of both IP-FRR methods.

6. Related Work

The work in [25] gives a survey on various approaches for IP resilience including early ideas of IP FRR. This is done at a very early stage so that LFAs and not-via addresses have not yet appeared. [5] provides a framework for IP FRR currently under development by the IETF routing working group (RTGWG). This group also published an RFC for LFAs [6] and an Internet draft proposing not-via addresses [7]. Improvements to not-via addresses have been proposed in [26]. The authors of [27] give an extensive overview on MPLS and IP FRR mechanisms including LFAs and not-via addresses, but they neither provide a classification nor a quantitative evaluation with regard to their applicability. First insights into the failure coverage of these IP FRR mechanisms have been given in [13, 14, 15].

However, only average values over all nodes in the network [13, 14] or cumulative distribution functions for the number of alternate nodes offering a specific repair mechanism [15] were reported. The detailed study of our work, i.e. the classification of LFAs and the availability analysis of different LFA types, is new. Results for backup path lengths were also presented in [14] but only for LFAs alone. None of the other studies has looked at the performance of the combined application of not-via addresses with LFAs.

LFAs cannot protect against all single link and node failures. In contrast to this inability of LFAs, resilience differentiation intentionally protects only some traffic in the network [28]. This rather depends on the desire of customers than on the basic ability of the network.

Fast reroute (FRR) concepts were first developed for MPLS technology and standardized in [4]. Currently, extensions for point-to-multipoint are under discussion to protect multicast traffic [29, 30]. The ability of IP routing for sub-second reaction to failures was studied in [3, 31] as well as stability issues when performing such optimizations.

Multiple routing configurations (MRC) provide a different FRR concept for IP networks. Various flavors of MRCs have been described in [32, 8, 33, 34]. MRCs create a small set of backup routing configurations which are used in failure cases. They complement each other in the sense that at least one valid route remains operational in each single link or node failure scenario for each pair of nodes in at least one configuration. MRCs can be implemented using the multi-topology extensions for OSPF and IS-IS [35, 36, 37]. [38] proposed an extension called 2DMRC to handle concurrent multi-failures with MRCs. The technique of multi-topology routing has also been used for improved service differentiation [39].

Failure inferencing based fast rerouting (FIFR) is another FRR concept for IP networks. It exploits the fact that packets arrive at routers through other interfaces than usual if rerouting is applied during network element failures. It computes interface-specific forwarding tables where the next hop of a packet does not only depend on its destination address but also on the incoming interface. Transient link failures [9] and transient node failures [40] failures can be handled. The original mechanism had problems with asymmetric link weights, but this has been fixed in [41] where extensions for inter-AS failures have also been developed. [42] suggested a modification called blacklist-based interface-specific forwarding (BISF) that avoids routing loops also in case of multiple failures.

The authors of [43] developed a method to achieve fast recovery of BGP peering link failures. Important are also concepts for loop-free re-convergence that can be used in combination with IP FRR mechanisms in case of long-lived fail-

ures [11]. One possible suggestion for loop-free reconvergence specifies an order in which nodes are allowed to update their forwarding tables in case of outages and after failure repair or installation of new network elements [44, 45].

Failure-carrying packets (FCP) constitute a completely different approach. All routers in the network have the same network map which does not change in case of a failure. Instead, packets are equipped with information about failures which helps to forward them on loop-free paths in case of failures [46].

7. Conclusion

In this work we provided a classification for loop-free alternates (LFAs) and proposed orders of preference for their application depending on the desired resilience level. LFAs cannot protect against all single link and node failures in a network. In particular, if LFAs must not lead to routing loops in case of single node failures or multiple other failures, only a subset of the existing LFAs can be used. As a result, only a fraction of all destinations can be protected. Therefore, it is proposed in the IETF [5, 7] to use LFAs where possible and to complement them by not-via addresses where needed to achieve full failure coverage. The motivation for this idea is the fact that LFAs seem to be simpler. We elaborated a concept for the combined application of LFAs and not-via addresses and compared it to protection with not-via addresses only. While LFAs provide slightly shorter backup paths, they tend to overload small links when the default INVCAP link weights are used. The maximum amount of decapsulated traffic from not-via tunnels in failure cases is mostly rather small compared to the overall traffic load and it cannot be effectively reduced by the combined application of not-via addresses and LFAs compared to protection by not-via addresses only. The amount of additional entries in the FIBs equals the number of links in the network and should not be a problem for routing protocols or FIB size. This amount can be significantly reduced since only the nodes which are on the path of a particular not-via tunnel need to know the corresponding not-via address. However, the combined approach cannot further decrease this number effectively. Hence, we have not found any significant advantages of the combined application of LFAs and not-via addresses compared to the protection by not-via addresses only. Therefore, we recommend to use either pure protection by LFAs and tolerate the partial failure coverage or pure protection by not-via addresses and tolerate the decapsulation traffic and the additional addresses in the FIBs. This has the advantage of a homogeneous protection mechanism which is easier to manage.

References

- [1] G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, C. Diot, Analysis of Link Failures in an IP Backbone, in: ACM SIGCOMM Internet Measurement Workshop, Marseille, France, 2002, pp. 237 – 242.
- [2] B. Fortz, M. Thorup, Robust Optimization of OSPF/IS-IS Weights, in: International Network Optimization Conference (INOC), Paris, France, 2003, pp. 225–230.
- [3] A. Basu, J. G. Riecke, Stability Issues in OSPF Routing, in: ACM SIGCOMM, San Diego, CA, USA, 2001, pp. 225–236.
- [4] P. Pan, G. Swallow, A. Atlas, RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels (May 2005).
- [5] M. Shand, S. Bryant, IP Fast Reroute Framework, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-rtgwg-ipfrr-framework-11> (Jun. 2009).
- [6] A. Atlas, A. Zinin, RFC5286: Basic Specification for IP Fast Reroute: Loop-Free Alternates (Sep. 2008).
- [7] M. Shand, S. Bryant, S. Previdi, IP Fast Reroute Using Not-via Addresses, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-rtgwg-ipfrr-notvia-addresses-05> (Jul. 2009).
- [8] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, O. Lysne, Fast IP Network Recovery Using Multiple Routing Configurations, in: IEEE Infocom, Barcelona, Spain, 2006.
- [9] S. Nelakuditi, S. Lee, Y. Yu, Z.-L. Zhang, C.-N. Chuah, Fast Local Rerouting for Handling Transient Link Failures, IEEE/ACM Transactions on Networking 15 (2) (2007) 359–372.
- [10] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, Characterization of Failures in an IP Backbone, in: IEEE Infocom, Hongkong, 2004.
- [11] M. Shand, S. Bryant, A Framework for Loop-free Convergence, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-rtgwg-lf-convfrmwk-05> (Jun. 2009).

- [12] R. Martin, M. Menth, K. Canbolat, Capacity Requirements for the Facility Backup Option in MPLS Fast Reroute, in: IEEE Workshop on High Performance Switching and Routing (HPSR), Poznan, Poland, 2006, pp. 329 – 338.
- [13] P. Francois, O. Bonaventure, An Evaluation of IP-Based Fast Reroute Techniques, in: ACM Conference on Emerging Network Experiment and Technology (CoNEXT), Toulouse, France, 2005, pp. 244–245.
- [14] A. F. Hansen, T. Cicic, S. Gjessing, Alternative Schemes for Proactive IP Recovery, in: Conference on Next Generation Internet Design and Engineering (NGI), Valencia, Spain, 2006.
- [15] M. Gjoka, V. Ram, X. Yang, Evaluation of IP Fast Reroute Proposals, in: IEEE International Conference on Communication System Software and Middleware (COMSWARE), Bangalore, India, 2007.
- [16] P. Francois, M. Shand, O. Bonaventure, Disruption-Free Topology Reconfiguration in OSPF Networks, in: IEEE Infocom, 2007.
- [17] D. Papadimitriou, P. Francois, IP Multicast Fast Reroute Framework, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-dimitri-rtgwg-mfrr-framework-00> (Feb. 2008).
- [18] P. Batchelor et al., Ultra High Capacity Optical Transmission Networks. Final Report of Action COST 239 (Jan. 1999).
- [19] The GEANT website, <http://www.geant.net/> (2009).
- [20] A. Nucci, A. Sridharan, N. Taft, The Problem of Synthetically Generating IP Traffic Matrices: Initial Recommendations, ACM SIGCOMM Computer Communications Review 35 (3) (2005) 19–32.
- [21] M. Roughan, Simplifying the Synthesis of Internet Traffic Matrices, ACM SIGCOMM Computer Communications Review 35 (5) (2005) 93 – 96.
- [22] Cisco Systems, OSPF Design Guide, <http://www.cisco.com/application/pdf/paws/7039/1.pdf> (2005).
- [23] B. Quoitin, S. Uhlig, S. Balon, J. Leprope, Providing Public Intra-Domain Traffic Matrices to the Research Community, ACM SIGCOMM Computer Communications Review 36 (1) (2006) 83–86.

- [24] ISO, ISO/IEC 10589:2002: Intermediate System to Intermediate System Intra-domain Routing Information Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473) (Nov 2002).
- [25] S. Rai, B. Mukherjee, O. Deshpande, IP Resilience within an Autonomous System: Current Approaches, Challenges, and Future Directions, *IEEE Communications Magazine* 43 (10) (2005) 142–149.
- [26] A. Li, P. Francois, X. Yang, On Improving the Efficiency and Manageability of NotVia, in: *ACM Conference on Emerging Network Experiment and Technology (CoNEXT)*, New York, NY, 2007.
- [27] A. Raj, O. Ibe, A Survey of IP and Multiprotocol Label Switching Fast Reroute Schemes, *Computer Networks* 51 (8) (2007) 1882–1907.
- [28] P. Cholda, A. Mykkeltveit, B. E. Helvik, O. J. Wittner, A. Jajszczyk, A Survey of Resilience Differentiation Frameworks in Communication Networks, *IEEE Communications Surveys & Tutorials* 9 (4) (2007) 32–55.
- [29] R. Aggrawal, D. Papadimitriou, S. Yasukawa, RFC4875: Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) (May 2007).
- [30] J. L. L. Roux, R. Aggarwal, J. Vasseur, M. Vigoureux, P2MP MPLS-TE Fast Reroute with P2MP Bypass Tunnels, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-mpls-p2mp-te-bypass-02> (Mar. 2008).
- [31] P. Francois, C. Filsfils, J. Evans, O. Bonaventure, Achieving Sub-Second IGP Convergence in Large IP Networks, *ACM SIGCOMM Computer Communications Review* 35 (2) (2005) 35 – 44.
- [32] M. Menth, R. Martin, Network Resilience through Multi-Topology Routing, in: *International Workshop on Design of Reliable Communication Networks (DRCN)*, Island of Ischia (Naples), Italy, 2005, pp. 271 – 277.
- [33] G. Apostolopoulos, Using Multiple Topologies for IP-only Protection Against Network Failures: A Routing Performance Perspective, Tech. Rep. TR377, Institute of Computer Science (ICS) of the Foundation for Research and Technology - Hellas (FORTH), Heraklion, Crete, Greece (2006).

- [34] T. Cicic, A. F. Hansen, A. Kvalbein, M. Hartmann, R. Martin, M. Menth, S. Gjessing, O. Lysne, Relaxed Multiple Routing Configurations: IP Fast Reroute for Single and Correlated Failures, accepted for IEEE Transactions on Network and Service Management (IEEE TNSM).
- [35] P. Psenak, S. Mirtorabi, A. Roy, N. L., P.-E. P., RFC4915: Multi-Topology (MT) Routing in OSPF (Jun. 2007).
- [36] N. Rawat, R. Shrivastava, D. Kushi, OSPF Version 2 MIB for Multi-Topology (MT) Routing, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-ospf-mt-mib-03> (Nov. 2008).
- [37] T. Przygienda, N. Shen, N. Sheth, RFC5120: M-ISIS: Multi Topology (MT) Routing in IS-ISs (Feb. 2008).
- [38] A. F. Hansen, O. Lysne, T. Cicic, S. Gjessing, Fast Proactive Recovery from Concurrent Failures, in: IEEE International Conference on Communications (ICC), Glasgow, UK, 2007.
- [39] K.-W. Kwong, R. Guérin, A. Shaikh, S. Tao, Improving Service Differentiation in IP Networks through Dual Topology Routing, in: ACM Conference on Emerging Network Experiment and Technology (CoNEXT), New York, NY, USA, 2007.
- [40] Z. Zhong, S. Nelakuditi, Y. Yu, S. Lee, J. Wang, C.-N. Chuah, Failure Inferencing based Fast Rerouting for Handling Transient Link and Node Failures, in: IEEE Global Internet Symposium, Miami, FL, USA, 2005.
- [41] Wang, Junling and Nelakuditi, Srihari, Ip fast reroute with failure inferencing, in: ACM Sigcomm Workshop on Internet Network Management (INM), Kyoto, Japan, 2007.
- [42] J. Wang, Z. Zhong, S. Nelakuditi, Handling Multiple Network Failures through Interface Specific Forwarding, in: IEEE Globecom, San Francisco, CA, 2006.
- [43] O. Bonaventure, C. Filsfils, P. Francois, Achieving Sub50 Milliseconds Recovery Upon BGP Peering Link Failures, in: ACM Conference on Emerging Network Experiment and Technology (CoNEXT), Toulouse, France, 2005.

- [44] P. Francois, O. Bonaventura, M. Shand, S. Bryant, S. Previdi, Loop-Free Convergence Using Ordered FIB Updates, Internet-Draft (work in progress), <http://tools.ietf.org/html/draft-ietf-rtgwg-ordered-fib-02> (Feb. 2008).
- [45] P. Francois, O. Bonaventure, Avoiding Transient Loops during IGP Convergence in IP Networks, in: IEEE Infocom, Miami, Florida, 2005.
- [46] K. Lakshminarayanan, M. Caesar, M. Rangan, T. Anderson, S. Shenker, I. Stoica, Achieving Convergence-Free Routing using Failure-Carrying Packets, in: ACM SIGCOMM, Kyoto, Japan, 2007.