

Performance Comparison of Not-Via Addresses and Maximally Redundant Trees (MRTs)

Michael Menth and Wolfgang Braun
University of Tübingen, Department of Computer Science, Germany

Abstract—Maximally redundant trees (MRTs) have been suggested in IETF as a new promising method for fast rerouting in IP networks. Unlike loop-free alternates (LFAs), they can protect against all link and node failures if topology allows. In this work, we compare MRTs and not-via addresses, which also guarantee full-failure coverage. We evaluate path lengths and relative link loads for single link and node failures in various test networks. The performance of MRTs significantly depends on the chosen root node so that the analysis presented in this paper should be performed before the roll-out of MRTs in operational networks.

Index Terms—Resilience, IP fast reroute, resource management

I. INTRODUCTION

When link or node failures happen in IP networks, routing tables are updated by distributed protocols (routing reconvergence) so that traffic affected by the failure is forwarded over backup paths. As this process may take in the order of several seconds, manufacturers and network operators work on fast reroute (FRR) solutions that quickly bypass the affected traffic around the failure location. A node uses fast reroute techniques as soon as it detects that a specific next-hop is no longer reachable until its routing table contains again working entries for the affected destinations. Such a node is called point of local repair (PLR). IP FRR solutions minimize traffic loss during the reconvergence time of IP routing. Furthermore, they may be used to delay IP routing reconvergence and may even avoid it if the failure is only short-lived; this saves the network from potential routing instabilities when a component fails and when it becomes operational again.

In the last decade, various FRR mechanisms have been proposed, but only loop-free alternates (LFAs) [1] are available in modern routers as they are simple and do not require coordination efforts among routers. Their drawback is that they cannot protect all destinations against failures. Therefore, LFAs improve the state of the art, but they do not constitute a perfect solution for FRR in IP networks.

Not-via addresses [2] constitute an alternative solution to provide backup for all destinations within a network provided the network is still connected. Not-via addresses use IP-tunnels to bypass traffic around an unreachable hop on the shortest path to the next-next-hop towards the destination. The forwarding logic is an extension of IP forwarding. Recently, the not-via draft status has been changed to informational due to complexity and state issues that can occur in practice.

The authors acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG) under grant ME2727/1-1. The authors alone are responsible for the content of the paper.

Maximally redundant trees (MRTs) [3] present another option. They basically provide two disjoint trees from any node to all other nodes in the network. They may use either IP-in-IP or LDP tunnels, can be implemented in a very scalable way, and their forwarding logic is independent of IP routing. Extensions for multicast exist. Available Routing Constructs (ARC) [4] have been presented only recently in IETF and are very similar to MRTs. In contrast to not-via addresses, we are not aware of any performance evaluations of MRTs. Therefore, we compare MRT variants with different algorithmic complexity in backup path length and link utilization. In order to classify the results, we compare MRTs to not-via addresses which are already well-investigated.

The remainder of the paper is structured as follows. In the next section, we give a brief overview of existing FRR proposals. In Section III we describe the concepts of not-via addresses and MRTs in more detail and compare them in Section IV. Finally, Section V summarizes this work and gives conclusions.

II. RELATED WORK

The ability of IP routing for sub-second reaction to failures was studied in [5], [6] as well as stability issues when performing optimizations to accelerate IP reconvergence. In contrast, the authors of [7] proposed loop-free convergence which delays the reconvergence process to update routing tables in an order that avoids transient loops which requires coexistence of IP FRR techniques to prevent traffic loss.

FRR techniques provide alternative forwarding solutions in IP networks during the IP reconvergence process. The IETF has defined loop-free alternate (LFAs) for that purpose [1]. If the next-hop is not reachable, traffic is forwarded to an alternate next-hop provided that this action does not create a loop. LFAs are simple because they require neither cooperation among nodes in the network nor tunneling mechanisms. However, appropriate LFAs are not available for all destinations and failures so that some traffic cannot be protected [8]–[11]. The failure coverage of LFAs can be improved by optimizing IGP link costs [12], [13] or by adding additional links to the network [14]. Another possibility is the use of remote LFAs. In the absence of a local LFA, a remote LFA can tunnel backup traffic to another node that is able to forward the traffic to the destination [15]. Remote LFAs improve the failure coverage of LFAs, but are neither able to provide backup paths for all destinations in case of any single link or node failure.

The IETF proposed not-via addresses [2] as an alternative to achieve 100% failure coverage by design in two-connected networks. We provide details in Section III-A. Variants of not-via addresses were suggested in [16], [17]. The authors of [18] proposed the usage of not-via addresses for disruption-free green traffic engineering. A comparison of LFAs and not-via addresses was given in [19].

Multiple routing configurations (MRCs) define multiple virtual topologies for IP networks over which traffic is forwarded. If a failure in one topology occurs, traffic can be switched to another topology that does not suffer from this failure. Several variants of MRCs have been described in [20]–[23]. The authors of [24] proposed an extension called 2DMRC to handle concurrent multi-failures with MRCs. MRTs are similar to MRCs and will be presented in detail in Section III-B.

Failure-inferencing-based FRR (FIFR) for IP networks was proposed in [25]. Routers detect packets arriving at other interfaces than usual and infer failure conditions. FIFR is able to avoid loops in case of failures by means of interface-specific forwarding tables.

Failure-carrying packets (FCP) were presented in [26]. All routers in the network share a common network map which does not change in case of a failure. In failure cases, packets are equipped with appropriate information which helps to deliver them on loop-free paths.

FRR concepts were first developed for MPLS technology and standardized in [27]. The authors of [28] give an extensive overview on MPLS and IP FRR mechanisms including LFAs and not-via addresses.

III. DESCRIPTION OF NOT-VIA ADDRESSES AND MRTS

In this section we provide a condensed description of not-via addresses and MRTs. We explain the operation and path calculation for the MRT approach in more detail because MRTs are relatively new and seem to be more complex to understand than not-via addresses.

A. Not-Via Addresses

A not-via address B_P is an IP address that is forwarded to node B not via node P . Such an address is helpful to implement FRR if P is a neighbor node of B . Therefore, the number of additional not-via addresses equals the number of unidirectional links in a network.

1) *Protection of Non-Last Hops*: Consider Figure 1(a). Node P is the next-hop towards destination D at node S and node B is the next-next hop. If S detects that P is no longer reachable, it encapsulates packets destined towards D with the not-via address B_P so that these packets are carried around the unreachable node P to the next-next-hop B . Node B decapsulates such packets and forwards them as usual to their destinations.

2) *Protection of Last Hops*: Consider Figure 1(b). Now, node P is already the destination of the packet. If P is no longer reachable from S , the above illustrated approach is not applicable as a next-next-hop does not exist. In such a case, S encapsulates packets destined to P with the not-via address

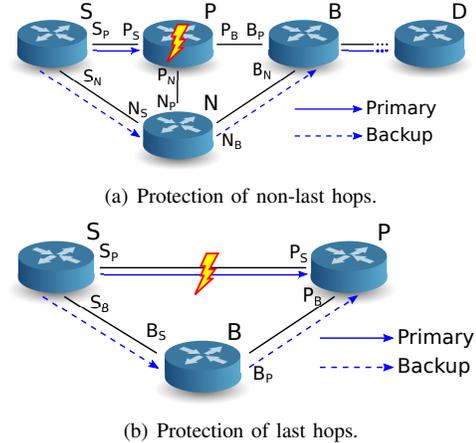


Fig. 1. Use of not-via addresses.

P_S and forwards it to any neighbor. When forwarding such packets, the potentially failed link from S to P will be avoided. However, if node P is down, none of its neighbors will be able to deliver such packets. If a node detects a failed next-hop for a packet with a not-via address, then it drops the packet to avoid forwarding loops.

3) *Modification of the Routing Protocols*: Routing protocols must be adapted to calculate entries in the routing tables for not-via addresses. To calculate the next-hop information for not-via addresses B_P , node P needs to be removed from the topology.

B. Maximally Redundant Trees (MRTs)

We explain the use of MRTs in a network that is at least two-connected, i.e., a single link or node failure cannot divide it into disconnected islands. However, the use of MRTs is not limited to such networks. We first sketch how pairs of disjoint backup paths are computed, then we explain which of them is selected, and finally we illustrate how this mechanism can be implemented with existing protocols.

1) *Construction of Backup Paths*: In any node, two disjoint backup paths towards any other node are constructed. The construction of backup paths needs to be the same in all nodes to assure consistency. First, a root node R is chosen which must be the same for all nodes. Then, all (bidirectional) links of a network topology are given a direction in such a way that all of them are part of a cycle through the root R . If a link belongs to several cycles, its orientation must be the same in all cycles. Moreover, only a single link l_R must enter the root node R so that removing that link leaves a directed acyclic graph (DAG). The directed graph including l_R is called an almost DAG (ADAG).

Various approaches for the construction of ADAGs are described in [29]. The “Lowpoint” (LP) variant uses a simple depth first search (DFS) algorithm, is very fast, and does not take link costs into account. It may lead to larger cycles. The “Shortest Path First” (SPF) variant uses Dijkstra’s algorithm, uses link costs as input, requires more computation overhead, and leads to shorter cycles.

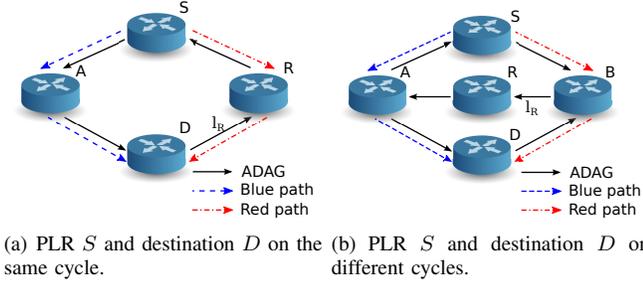


Fig. 2. Construction of disjoint backup paths.

The ADAG structure provides for two backup paths that are disjoint but do not minimize path length. Consider a PLR S and a destination D . If they are located on the same cycle, the two parts of the cycle from S to D provide two disjoint paths. This is depicted in Figure 2(a). The backup path in cycle orientation is denoted blue, the one against cycle orientation is denoted red. If several cycles exist that contain both S and D , the shortest one of them in terms of link cost is taken for the computation of the backup path. Therefore, the layout of the backup paths depends on the link costs. If S and D are not located on the same cycle, then they are located on two different cycles that intersect at least in the root node R and link l_R due to the construction of the ADAG. Then we find again two backup paths as depicted in Figure 2(b). The blue backup path goes from the PLR S against cycle orientation to a intersection node with the cycle on which D is located; from there, it follows that cycle in orientation to destination D . The red backup path goes from the PLR S in cycle orientation to another intersection node with the cycle on which D is located; from there, it follows that cycle against orientation to destination D . Figure 2(b) also shows that backup paths are constructed without the overlapping parts of the two cycles (between A and B) which keeps the backup paths short.

2) *Use of the Backup Paths:* If a node detects a failure of the normal forwarding plane, it determines whether the packet is to be sent over the blue or the red backup path and adds a hint that the packet is on the blue or red backup path [29, Section 4.7].

In case of a failure, the choice whether traffic is forwarded over the blue or the red backup path is limited by the location of the failure: if the failed element is part of the red backup path, the blue backup path needs to be chosen and vice-versa. If the failed element is not part of both backup path, any one of them could be chosen. However, in any case the existing MRT algorithm chooses a feasible backup path without taking the backup path length into account which makes it computationally efficient.

Thus, there is an obvious improvement of the existing MRT algorithm: choose the shorter backup paths when both are feasible! We investigate that alternative backup path selection in Section IV-C6.

3) *Implementation of MRTs:* The blue/red backup paths are constructed in such a way that a blue/red backup path from S to D containing node A implies complete overlap with the

blue/red backup path from A to D . Thus, the blue/red backup paths run along sink trees and can be implemented as such. Three methods have been proposed using either IP or label distribution protocol (LDP) tunnels. More detailed information is provided in [3].

a) *IP-in-IP Tunneling:* In pure IP networks, two additional IP addresses (blue and red) are associated with each node in the network. In case of a failure, traffic encapsulated with the appropriate colored address and forwarded by nodes accordingly along the sink tree. This approach is similar to tunneling traffic with not-via addresses.

b) *Classic LDP Tunneling:* In LDP networks, one label is associated with each node in the network. A label switched path (LSP) is set up from any other node in the network to the destination following normal IP routing. MRTs are implemented as follows. Additional blue and red labels are associated with each node in the network and additional blue and red LSPs are established along the MRT sink trees.

c) *LDP Tunneling Using Extension Labels:* Another variant using extension labels [30] saves LDP labels for endpoints. A single label is associated with each node in the network. To mark a packet as belonging to the blue/red backup path, an extension label is inserted between the IP and the LDP address of the destination node. A blue and red sink tree still need to be established to carry traffic with blue and red extension labels. The technique to forward packets over different topologies is called multi-topology routing and may also be used for traffic engineering purposes [31].

IV. PERFORMANCE EVALUATION

In this section we compare not-via addresses and MRTs with regard to backup path length and relative link load. We first explain our methodology and performance metrics, present the networks under study, and then provide performance results.

A. Methodology

We study on which paths traffic is forwarded in IP networks under failure-free conditions (\emptyset) and for a set of failure scenarios that contain all bidirectional single link failures and all single node failures. We denote this set of considered scenarios by \mathcal{S} which also contains the failure-free scenario \emptyset . We construct the path layout for the following rerouting methods:

- IP rerouting,
- IP FRR using not-via addresses (towards next-next-hop),
- IP FRR using MRTs towards destination,
- IP FRR using MRTs towards next-next-hop.

We investigate multiple MRTs using different ADAGs that result from different root nodes and ADAG construction methods (LP and SPF variant). We compare path lengths and relative link loads. In the following, we define the exact metrics used for evaluation.

1) *Quantification of Path Lengths:* All traffic from a node A to a node B is subsumed by the traffic aggregate $g(A, B)$ and the set of all traffic aggregates in the network is denoted by \mathcal{G} . To facilitate our analysis we assume that ECMP is not

enabled so that all flows of an aggregate share the same path. Then, the path length of a single aggregate $g \in \mathcal{G}$ in a specific failure scenario $s \in \mathcal{S}$ is unique and denoted by $L(g, s)$. We define four different metrics for path lengths.

- Average path length under failure-free conditions

$$L_{\emptyset}^{avg} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} L(g, \emptyset) \quad (1)$$

- Maximum path length under failure-free conditions

$$L_{\emptyset}^{max} = \max_{g \in \mathcal{G}} (L(g, \emptyset)) \quad (2)$$

- Average path length for the set of considered scenarios \mathcal{S}

$$L_{\mathcal{S}}^{avg} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \max_{s \in \mathcal{S}} (L(g, s)) \quad (3)$$

- Maximum path length for the set of considered scenarios \mathcal{S}

$$L_{\mathcal{S}}^{max} = \max_{g \in \mathcal{G}, s \in \mathcal{S}} (L(g, s)) \quad (4)$$

2) *Quantification of Link Loads:* We calculate the (relative) load $\rho(l, s)$ of a specific link $l \in \mathcal{E}$ in a particular scenario $s \in \mathcal{S}$ by summing up the rates of all traffic aggregates forwarded over that specific link and divide that sum by the capacity of that link. We consider two different metrics.

- Maximum (relative) link load under failure-free conditions

$$\rho_{\emptyset}^{max} = \max_{l \in \mathcal{E}} (\rho(l, \emptyset)) \quad (5)$$

- Maximum (relative) link load for the set of considered scenarios \mathcal{S}

$$\rho_{\mathcal{S}}^{max} = \max_{l \in \mathcal{E}, s \in \mathcal{S}} (\rho(l, s)) \quad (6)$$

B. Networks under Study

In our study we investigate the 11 networks described in Table I. Most of them are taken from the ‘‘topology zoo’’ [32] and three networks are taken from [33], [34], and [35]. As only two-connected topologies are resilient against single link and node failures, we work only with such structures to simplify our analysis. As the original networks are not two-connected, we strip off a minimum number of nodes to make them two-connected. Table I indicates the number of nodes $|\mathcal{V}|$ and the number of links $|\mathcal{E}|$ in the original topologies and in the two-connected topologies used in our study.

The number of neighbors of a node is denoted as its node degree. The table also indicates the maximum node degree δ_{max} as well the source of the network information. Some networks have heterogeneous link bandwidths, others have homogeneous link bandwidths. In most cases, administrative link costs for routing purposes are not given so that we work with uniform link costs. Due to the lack of further information, we assume homogeneous traffic matrices for all networks. For each network the traffic matrix was scaled such that the maximum link load under failure-free conditions is 1.0. This approach is unrealistic, but it does not affect path lengths.

TABLE I
NETWORKS UNDER STUDY

Network	Original topology		2-connected topology			Source
	$ \mathcal{V} $	$ \mathcal{E} $	$ \mathcal{V} $	$ \mathcal{E} $	δ_{max}	
AGIS	25	60	16	42	5	[32]
GEANT	40	122	30	100	7	[32]
GARR	48	124	20	60	7	[32]
InternetMCI	19	66	18	64	7	[32]
PionierL3	27	64	24	58	5	[32]
Rediris	19	62	18	60	10	[32]
Uninett2011	66	186	57	168	7	[32]
UUnet	42	154	38	146	11	[32]
Nobel	-	-	28	82	5	[33]
Labnet03	-	-	20	106	10	[34]
COST239	-	-	11	52	6	[35]

It does affect relative link loads but all investigated reroute mechanisms are affected in the same way so that this approach produces comparable results for not-via addresses and MRTs.

C. Analysis of Backup Path Lengths

We investigate the (backup) path lengths that are obtained with IP routing and the various IP FRR algorithms.

1) *Average Backup Path Length $L_{\mathcal{S}}^{avg}$ in the AGIS Network:* We first analyze the average backup path length $L_{\mathcal{S}}^{avg}$ for IP rerouting, not-via addresses, and MRTs in the AGIS network with uniform link costs. Figure 3(a) shows these values and also indicates the average path length for IP routing L_{\emptyset}^{avg} under failure-free conditions for comparison purposes. The latter is 2.6 hops long and the shortest by construction. IP rerouting produces the shortest backup path length with 4.7 hops on average. Not-via addresses lead to clearly longer backup paths which are on average 6.3 hops long.

For MRTs the figure shows the average backup path lengths $L_{\mathcal{S}}^{avg}$ for all possible root nodes and for the ADAG construction methods LP and SPF. The average path length for MRTs depends significantly on the root node. It is shortest for root node 5 and ADAG construction method SPF; under these conditions it is 6.8 hops long and relatively close to the backup path length of not-via addresses.

Looking at the topology of the AGIS network in Figure 4, we realize that node 5 is located very centrally in the network and has the maximum node degree. The latter affects that node 5 is part of many relatively short cycles which allows for relatively short backup paths. Other root nodes lead to significantly longer backup paths. For instance, root node 7 and ADAG construction method LP cause the largest average backup path length of 9.8 hops. Node 7 is located such that cycles containing that node are relatively long; therefore, the backup paths are also quite long.

Looking at the ADAG construction methods, we observe that the LP method mostly leads to average backup path lengths that are marginally larger than those constructed by the SPF method, but outliers exist. For root node 11, the average backup path length for SPF is more than one hop larger than the one for LP, while for root nodes 6 and 14

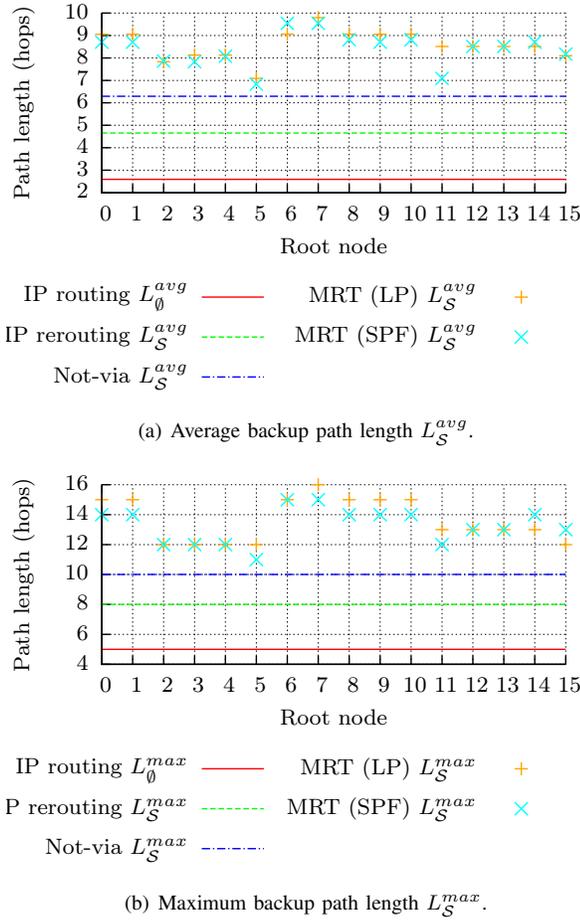


Fig. 3. Backup path lengths in the AGIS network with uniform link costs.

the ADAG construction method LP leads to visibly shorter average backup paths than the SPF method. Below the line, the choice of the ADAG construction method (LP or SPF) has clearly less impact on the average backup path length than the choice of the root node.

2) *Maximum Backup Path Length L_S^{max} in the AGIS Network:* Comparing Figure 3(b) with Figure 3(a) shows that the longest backup paths can be significantly larger than the presented averages. They are 8 instead of 4.7 hops with IP rerouting, 10 instead of 6.3 hops with not-via addresses, and 11 and 12 instead of 6.8 and 7.0 hops for the best MRTs.

With root node 7 and ADAG construction method LP, the largest observed backup path is 16 hops long. This backup path is illustrated in Figure 4. The situation occurs if node 10 sends traffic to node 0 and the link $1 \rightarrow 0$ fails. The traffic reaches the PLR (node 1) after 4 hops and is then redirected over the red backup path, which is 12 hops long, to the destination. Thereby, the backup path traverses parts of the primary path again.

3) *Impact of Backup Path Variants for the AGIS Network:* FRR methods provide a backup path from the PLR either to the destination (e.g. LFAs) or to the next-next-hop (e.g. not-via addresses). With MRTs both alternatives can be supported.

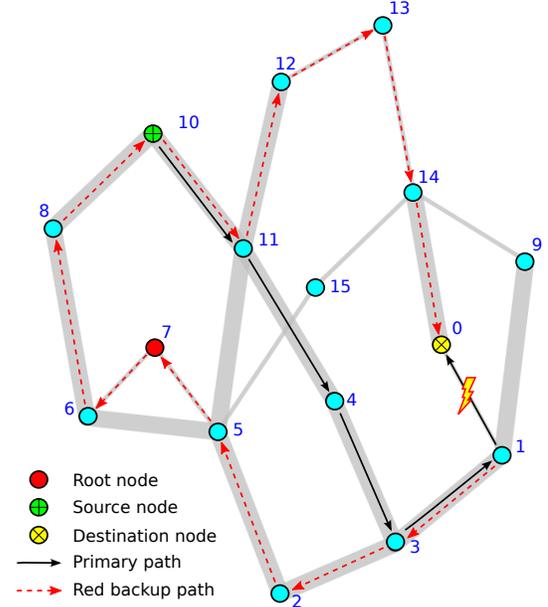


Fig. 4. The AGIS network with 16 nodes and 42 links; the width of the links is proportional to their bandwidths; node 10 sends traffic to node 0, but node 1 redirects that traffic over its red backup path (root node 7, ADAG construction method LP) due to the failure of link $1 \rightarrow 0$.

Above we provided results only for the detour option from the PLR to the destination. We also studied the bypass option from the PLR to the next-next-hop for all possible root nodes and for the two ADAG construction methods LP and SPF. We observed that the detour and bypass variants have an even lower impact on the backup path length than the ADAG construction method whose impact is clearly dominated by the choice of the root node. In any case, the average and maximum backup path lengths are longer than those of not-via addresses.

4) *Impact of Link Costs for the AGIS Network:* The path layout significantly depends on link costs. In our study we have used uniform link costs because real link costs were not available for the networks under study. To evaluate the path length for other link costs, we choose link costs that are inverse proportional to the link bandwidth. The AGIS network has two types of significantly different link bandwidths (see Figure 4) which produce link costs that are clearly different from uniform link costs.

As a result, average path lengths under failure-free conditions increase from 2.6 hops for uniform link costs to 2.8 hops. Average backup path lengths for IP rerouting increase from 4.7 to 5.2 hops, for not-via addresses from 6.3 to 6.9 hops, and for MRT with root node 5 and ADAG construction method SPF from 6.8 to 7.0 hops. Thus, link costs have only a minor impact on the path lengths with MRTs. Moreover, the path layout of the ADAG construction method LP is even independent of the link costs.

5) *Confirmation of the Results by Other Investigated Networks:* Figure 5 shows the average backup path length L_S^{avg} for all investigated FRR methods and networks. For MRT the root nodes with the shortest maximum backup path lengths are chosen depending on the construction methods LP and SPF; we observe that the best root node is mostly the same

for LP and SPF. The average path length of IP routing under failure-free conditions is a lower bound of the average backup path lengths. Backup paths for IP rerouting are the shortest; in some networks L_S^{avg} for IP rerouting is 100% larger than the average path length, in some other networks it is only 30% larger. In all networks, backup paths for not-via addresses are longer than those for IP rerouting but to a different degree. With MRTs, average backup path lengths are mostly one or two hops longer than with not-via addresses. In all investigated networks, the impact on maximum path lengths L_S^{avg} of the ADAG construction methods (LP and SPF), the backup path variant (detour or bypass), and the link costs (uniform link costs and inverse proportional to link bandwidths) are also very small. The Uninett constitutes an exception because for most root nodes the ADAG construction method SPF leads to clearly shorter backup paths than the ADAG construction method LP.

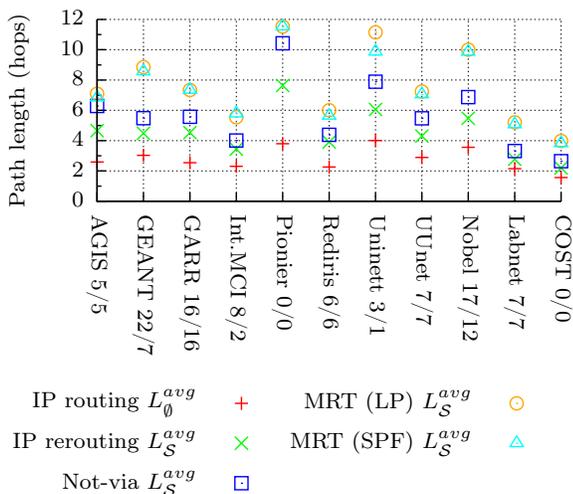


Fig. 5. Comparison of average backup path length L_S^{avg} for all investigated FRR methods and networks; for MRT the root node with the shortest backup paths is chosen and indicated with the network name (LP/SPF).

6) *Alternative Backup Path Selection*: The MRT construction algorithm as described in [29, Section 4.7] determines whether the blue or the red backup path is to be taken (see Section III-B2). We consider an alternative path selection algorithm that obviously leads to shorter backup path lengths: if neither the blue nor the red backup path are affected by the failure, the shorter one is taken. This improvement makes the path selection algorithm more complex.

We implemented the alternative path selection and evaluated the path lengths. As we found hardly any improvement for the best root nodes, in particular no reduction of the maximum backup path length, we evaluated in how many cases the alternative path selection algorithm has effect and compiled the results in Table II for the ADAG construction method SPF an for all investigated networks and the best root nodes. In most cases (89.5% – 99.6%), either the blue or the red backup path is affected so that the path selection algorithm has no choice. If there is a choice, the existing path selection algorithm already takes the shorter backup path quite often, so

TABLE II
EFFECT OF ALTERNATIVE PATH SELECTION IN THE TEST NETWORKS WITH UNIFORM LINK COSTS AND ADAG CONSTRUCTION METHOD SPF; THE ROOT NODE WITH THE SHORTEST MAXIMUM PATH LENGTH IS CHOSEN.

Network, root	Not applicable	No improvement	Improvement
AGIS, 5	96.4%	1.5%	0.7%
GEANT, 7	95.6%	1.5%	2.9%
GARR, 16	98.4%	0.4%	1.2%
InternetMCI, 2	96.2%	1.5%	2.3%
PionierL3, 0	99.6%	0.2%	0.2%
Rediris, 6	98.9%	0.1%	1.0%
Uninett2011, 1	96.2%	2.1%	1.7%
UUnet, 7	93.6%	3.4%	3.0%
Nobel, 12	93.0%	3.6%	3.4%
Labnet03, 7	89.5%	8.2%	2.3%
COST239, 0	91.2%	6.6%	2.2%

that the backup path length can be reduced only rarely (0.2% – 3.4% of the cases). We obtained similar results for ADAG construction method LP. Though the advantage of the new backup path selection is obvious, it rarely has effect so that improvements are hardly measurable in our test networks. This result is probably due to the fact that we chose the best root nodes for evaluation which already minimize the maximum backup path length.

D. Analysis of Relative Link Loads

We study the relative link loads that are increased by the backup traffic for the various rerouting algorithms.

1) *Maximum Relative Link Load L_S^{max} in the AGIS Network*: Figure 6 shows the maximum relative link load in the AGIS network with uniform link costs. The value is 1.0 for IP routing under failure-free conditions according to our construction of the traffic matrices. With IP rerouting, an increased load of 1.83 occurs, and with not-via addresses an increased load of 1.92. For MRTs the maximum relative load is below, between, or above these values depending on the root node. We observe that the ADAG construction method LP tends to cause lower maximum link loads than the SPF method, but exceptions to that rule exist. The lowest maximum load values are achieved with root node 4 which is even significantly lower than the one of IP rerouting. For most root nodes the LP construction method produces lower maximum or equal link loads compared to the SPF construction method, but for a few others, the SPF construction method provides better results.

The reason why MRTs can lead to lower maximum link loads than IP rerouting or not-via addresses is that some of the constructed MRTs carry backup traffic on less loaded links. This is facilitated through the diverse sets of backup paths – some path layouts are just better, some are worse. Another aspect is that not-via addresses locally bypass the traffic around the unreachable hop and create hot spots. In contrast, MRTs tend to deviate backup traffic away from the failure before delivering it to the destination. Thereby MRTs implicitly avoid local hot spots near next-next-hops.

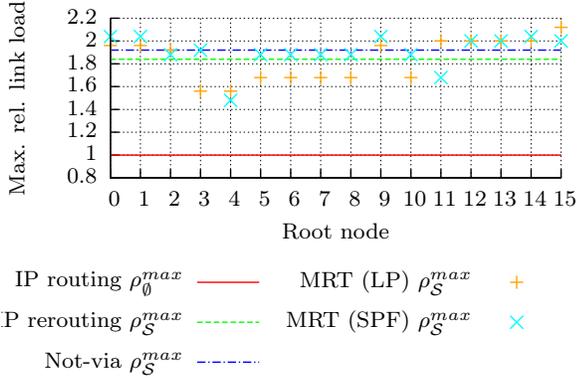


Fig. 6. Maximum relative link load in the AGIS network with uniform link costs.

Unfortunately, the root node minimizing the backup path length does not minimize the maximum link load. Thus, network operators need to decide which of the two metrics is more important for their business before choosing an appropriate root node.

2) Impact of Backup Path Variants in the AGIS Network:

The bypass variant for MRTs constantly leads to higher link loads than the detour variant. Traffic is concentrated on operational links towards the next-next-hops before being forwarded towards the destination. This increases the load on these links.

3) *Impact of Link Costs in the AGIS Network:* Link costs have a tremendous impact on maximum link loads. For uniform link costs, IP rerouting has a maximum link load of 1.83 while it reaches a value of 2.55 for link costs inverse proportional to link capacities. The maximum link loads for not-via addresses increase from 1.92 to 2.95. Also the maximum link loads for MRTs increase so that most of them are between the one of IP rerouting and the one of not-via addresses. This is not surprising as the fact that link costs provide a large optimization potential for minimization of maximum link utilization – also in combination with not-via addresses – is well known [36]. Link cost optimization for MRTs is possible in a similar way.

4) *Confirmation of the Results by Other Investigated Networks:* Figure 7 shows the maximum relative link loads ρ_S^{max} for all investigated rerouting methods and networks. The maximum path length for IP routing under failure-free conditions is 1.0 according to our construction of traffic matrices. Maximum link loads for IP rerouting are mostly significantly larger, followed by the one of not-via addresses and MRTs. However, exceptions to that rule exist. In the UUnet network, MRTs lead to lower maximum link loads than IP rerouting or not-via addresses. Also in the GARR and Pionier network MRTs lead to lower maximum link loads than not-via addresses. The maximum link loads for MRTs can be further reduced by choosing more appropriate root nodes, namely those that minimize the maximum link load instead of the maximum path length. However, maximum link loads can be effectively reduced by optimizing link costs [36] so that

we do not consider MRTs as a good option to avoid overload due to backup traffic.

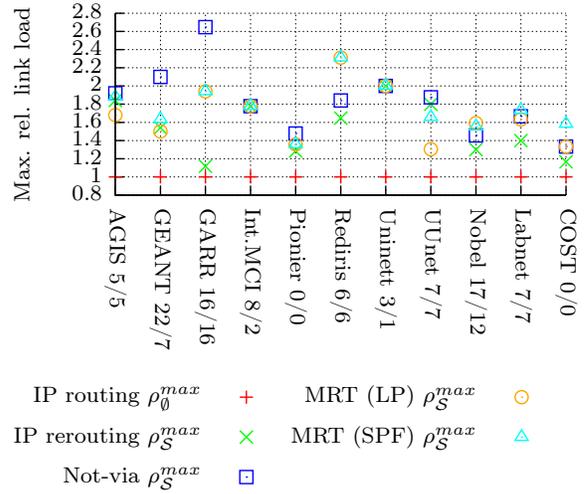


Fig. 7. Maximum relative link loads ρ_S^{max} for all investigated rerouting methods and networks; for MRT the root node with the shortest backup paths is chosen and indicated with the network name (LP/SPF).

E. Summary

We have shown that MRTs lead to significantly longer backup paths than IP rerouting and also to clearly longer paths than not-via addresses. The difference becomes evident when considering maximum path lengths instead of averaged path lengths. The backup path lengths for MRTs depend significantly on the root node. Thus, an a priori analysis can help to configure MRTs such that backup path lengths are minimized. Nevertheless, the backup path lengths are still clearly longer than those of not-via addresses. The impact of the ADAG construction methods LP and SPF on backup path length was significant only in the Uninett network. For the LP method, backup path lengths are hardly influenced by IP link costs because LP's ADAG construction does not consider them. With the SPF method we observed slightly more impact of IP link costs on backup path lengths. In contrast, IP link costs clearly influence the backup path lengths with not-via addresses.

Depending on the network and IP link costs, maximum relative link loads with MRTs can be better than with IP rerouting but also worse than with not-via addresses. However, the use of MRTs to reduce link loads is not recommendable as other methods exist.

V. CONCLUSION

We have evaluated and compared MRTs and not-via addresses for various test networks. With not-via addresses, one additional IP address is needed for every unidirectional link while MRTs require two additional addresses per node. Thus, MRTs require less state inside routers than not-via addresses.

We have shown that MRTs can lead to significantly longer backup path lengths than IP rerouting and not-via addresses. However, this can be mitigated through an appropriate choice

of the root node. Then, MRTs have in some networks only slightly longer backup paths than not-via addresses. We witnessed only minor impact on backup path length by ADAG construction methods (LP or SPF), link costs, and backup path variants (detour or bypass). We investigated an obvious method to reduce backup path lengths, showed that it has only little effect in realistic settings, and explained the reasons.

Reroute mechanisms increase maximum relative link loads through backup traffic. MRTs can lead to larger but also to smaller values than not-via addresses, depending on the network and IP link costs.

Although MRTs cause excessive path lengths in many networks, standardization recently focuses on MRTs due to less required state and computational complexity. This is a difficult task because savings in routing table size and routing calculation effort depend not only on the considered network topology but also on implementation specifics.

ACKNOWLEDGEMENTS

The authors thank Matthias Hartmann and David Hock for their code base and support and Alia Atlas, Andras Csaszar, and Gabor Enyedi for their insightful comments.

REFERENCES

- [1] A. Atlas and A. Zinin, "RFC5286: Basic Specification for IP Fast Reroute: Loop-Free Alternates," Sep. 2008.
- [2] S. Bryant, S. Previdi, and M. Shand, "IP Fast Reroute Using Not-via Addresses," <http://tools.ietf.org/id/draft-ietf-rtgwg-ipfr-notvia-addresses>, Dec. 2012.
- [3] A. Atlas, R. Kebler, M. Monstantynowicz, G. Enyedi, A. Csaszar, R. White, and M. Shand, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees," <http://tools.ietf.org/html/draft-ietf-rtgwg-mrt-fr-architecture>, Mar. 2012.
- [4] P. Thubert and P. Bellagamba, "Available Routing Constructs," <http://tools.ietf.org/html/draft-thubert-rtgwg-arc>, Oct. 2012.
- [5] A. Basu and J. G. Riecke, "Stability Issues in OSPF Routing," in *ACM SIGCOMM*, San Diego, CA, USA, Aug. 2001, pp. 225–236.
- [6] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving Sub-Second IGP Convergence in Large IP Networks," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 2, pp. 35 – 44, Jul. 2005.
- [7] P. Francois and O. Bonaventure, "Avoiding Transient Loops during the Convergence of Link-State Routing Protocols," *IEEE/ACM Transactions on Networking*, 2007.
- [8] —, "An Evaluation of IP-Based Fast Reroute Techniques," in *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Toulouse, France, Oct. 2005, pp. 244–245.
- [9] A. F. Hansen, T. Cicic, and S. Gjessing, "Alternative Schemes for Proactive IP Recovery," in *2nd Conference on Next Generation Internet Design and Engineering (NGI)*, Valencia, Spain, Apr. 2006.
- [10] M. Gjoka, V. Ram, and X. Yang, "Evaluation of IP Fast Reroute Proposals," in *IEEE International Conference on Communication System Software and Middleware (COMSWARE)*, Bangalore, India, Jan. 2007.
- [11] C. Filsfils, Ed., P. Francois, Ed., M. Shand, B. Decraene, J. Uttaro, N. Leyman, and M. Horneffer, "RFC6571: Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks," Jun. 2012.
- [12] H. Trong Viet, P. Francois, Y. Deville, and O. Bonaventure, "Implementation of a Traffic Engineering Technique that Preserves IP Fast Reroute in COMET," in *Rencontres Francophones sur les Aspects Algorithmiques des Tlcommunications (ALGOTEL)*, Carry Le Rouet, France, Jun. 2009.
- [13] G. Retvari, L. Csikor, J. Tapolcai, G. Enyedi, and A. Csaszar, "Optimizing IGP Link Costs for Improving IP-level Resilience," in *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Krakow, Poland, Oct. 2011.
- [14] G. Retvari, J. Tapolcai, G. Enyedi, and A. Csaszar, "IP Fast ReRoute: Loop Free Alternates Revisited," in *IEEE Infocom*, Shanghai, China, Apr. 2011.
- [15] S. Bryant, C. Filsfils, M. Shand, and N. So, "Remote LFA FRR," <http://tools.ietf.org/html/draft-ietf-rtgwg-remote-lfa>, Dec. 2012.
- [16] A. Li, P. Francois, and X. Yang, "On Improving the Efficiency and Manageability of NotVia," in *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, New York, NY, Dec. 2007.
- [17] G. Enyedi, G. Retvari, P. Szilagy, and A. Csaszar, "IP Fast ReRoute: Lightweight Not-Via," in *IFIP-TC6 Networking Conference (Networking)*, Aachen, Germany, May 2009.
- [18] N. Wang, C. Michael, and K. H. Ho, "Disruption-Free Green Traffic Engineering with NotVia Fast Reroute," *IEEE Communications Letters*, vol. 15, no. 10, Oct. 2011.
- [19] R. Martin, M. Menth, M. Hartmann, T. Cicic, and A. Kvalbein, "Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute?" *Computer Networks*, vol. 54, no. 8, pp. 1300 – 1315, Jun. 2010.
- [20] M. Menth and R. Martin, "Network Resilience through Multi-Topology Routing," in *5th International Workshop on Design of Reliable Communication Networks (DRCN)*, Island of Ischia (Naples), Italy, Oct. 2005, pp. 271 – 277.
- [21] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne, "Fast IP Network Recovery Using Multiple Routing Configurations," in *IEEE Infocom*, Barcelona, Spain, Apr. 2006.
- [22] G. Apostolopoulos, "Using Multiple Topologies for IP-only Protection Against Network Failures: A Routing Performance Perspective," Institute of Computer Science (ICS) of the Foundation for Research and Technology – Hellas (FORTH), Heraklion, Crete, Greece, Tech. Rep. TR377, 2006.
- [23] T. Cicic, A. F. Hansen, A. Kvalbein, M. Hartmann, R. Martin, M. Menth, S. Gjessing, and O. Lysne, "Relaxed Multiple Routing Configurations: IP Fast Reroute for Single and Correlated Failures," *IEEE Transactions on Network and Service Management (IEEE TNSM)*, vol. 6, no. 1, pp. 1 – 14, Mar. 2009.
- [24] A. F. Hansen, O. Lysne, T. Cicic, and S. Gjessing, "Fast Proactive Recovery from Concurrent Failures," in *IEEE International Conference on Communications (ICC)*, Glasgow, UK, Jun. 2007.
- [25] S. Nelakuditi, S. Lee, Y. Yu, Z.-L. Zhang, and C.-N. Chuah, "Fast Local Rerouting for Handling Transient Link Failures," *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, pp. 359–372, Apr. 2007.
- [26] K. Lakshminarayanan, M. Caesar, M. Rangan, T. Anderson, S. Shenker, and I. Stoica, "Achieving Convergence-Free Routing using Failure-Carrying Packets," in *ACM SIGCOMM*, Kyoto, Japan, Aug. 2007.
- [27] P. Pan, G. Swallow, and A. Atlas, "RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels," May 2005.
- [28] A. Raj and O. Ibe, "A Survey of IP and Multiprotocol Label Switching Fast Reroute Schemes," *Computer Networks*, vol. 51, no. 8, pp. 1882–1907, 2007.
- [29] A. Atlas, G. Enyedi, and A. Csaszar, "Algorithms for Computing Maximally Redundant Trees for IP/LDP Fast-Reroute," <http://tools.ietf.org/html/draft-enyedi-rtgwg-mrt-fr-algorithm>, Oct. 2012.
- [30] Q. Zhao, L. Fang, C. Zhou, L. Li, N. So, and R. Torvi, "LDP Extensions for Multi Topology Routing," <http://tools.ietf.org/html/draft-ietf-mppls-ldp-multi-topology>, Dec. 2012.
- [31] A. Kvalbein and O. Lysne, "How can Multi-Topology Routing be used for Intradomain Traffic Engineering?" in *SIGCOMM Workshop on Internet Network Management*, Pisa, Italy, 2006.
- [32] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The Internet Topology Zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, pp. 1765 –1775, Oct. 2011.
- [33] "SNDlib 1.0 – Survivable Network Design Data Library," <http://sndlib.zib.de>, 2005.
- [34] M. Menth, "Efficient Admission Control and Routing in Resilient Communication Networks," PhD thesis, University of Würzburg, Faculty of Computer Science, July 2004. [Online]. Available: {<http://www.opus-bayern.de/uni-wuerzburg/volltexte/2004/994/pdf/Menth04.pdf>}
- [35] C. Mauz, "Mapping of Arbitrary Traffic Demand and Network Topology on a Mesh of Rings Network," in *IFIP Working Conference on Optical Network Design and Modelling*, Feb. 2001.
- [36] D. Hock, M. Hartmann, M. Menth, and C. Schwartz, "Optimizing Unique Shortest Paths for Resilient Routing and Fast Reroute in IP-Based Networks," in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Osaka, Japan, Apr. 2010.