Alternative Schemes for Proactive IP Recovery

Audun Fosselie Hansen*†, Tarik Cicic* and Stein Gjessing*‡
*Networks and Distributed Systems Group, Simula Research Laboratory, Oslo, Norway
Email: {audunh, tarikc, steing}@simula.no
† Telenor R&D, Oslo, Norway
‡ Department of Informatics, University of Oslo, Norway

***Abstract*— Recovery at the IP layer has originally been handled by the slow process of IP re-convergence. As the dependence on the Internet broadens and real time applications like VoIP become a common service of the Internet, fast proactive recovery becomes an important property of the communication protocols. There are currently two IETF initiatives for proactive recovery drawing considerable attention, IP Fast Reroute and Multi-Topology Routing using Multiple Routing Configurations. In this paper we evaluate and compare these approaches.**

***Index Terms*— IP Fast Reroute, loop-free proactive IP recovery, multi-topology routing**

I. Introduction

Internet communications increasingly affect our lives. To answer this trust, the Internet must emerge as a robust and reliable platform for future communication services. Robustness against component failures has always been one of the design requirements of the Internet [1], and distributed IP re-convergence has been the implementation. However, the time-scale of this approach is not compatible with the services emerging today.

A great effort has been taken to improve the time-scale of IP re-convergence. The detection time has been improved by using shorter hello-message intervals [2]. Studies have showed that there is a lower limit for this hello-message interval, else it may provide instabilities in the network [3]. Another approach could be failure notification from the physical layer. However, the physical layer will not be aware of router failures. Others approach the time-scale by improving the effectiveness of routing information dissemination, either by reducing the amount of updates in stable periods [4] or prioritizing the update messages [5], [6]. Also the speed of calculating new shortest path trees has been improved by using an incremental approach using the old trees as input to the new calculation [7].

Although all this effort has been spent, the time-scale is still not suitable for new real-time services. Francois and others have recently presented results demonstrating that even in a controlled lab environment, 0.3 - 1 seconds is the best achievable recovery time for IP re-convergence [8]. Hence, they conclude that network wide dynamic IP recovery will never provide recovery in a time-scale suitable for real-time services. Instead, they suggest that a proactive approach for IP rerouting should be developed. A proactive scheme also has the advantage that the IP re-convergence process can be put on hold preventing unnecessary instabilities. Most failures seem to be transient [9] and hence IP re-convergence may not be invoked at all.

Within IETF there are currently activities focusing on proactive IP fast reroute [10], [11]. The focus is on how to provide connectionless proactive recovery using backup next-hops. They suggest to solve failure scenarios that have no loop-free backup next-hop by multi-hop repair paths, i.e. there is a router more than one hop away that can provide a loop-free path to the destination.

Multi-Topology (MT) routing is another activity with the potential of solving connectionless proactive recovery [12], [13], [14]. These drafts describe extensions to ISIS and OSPF in order to define independent IP topologies that can be used to compute different paths for unicast traffic, multicast traffic, different classes of service, or an in-band network management topology. This scheme can also be used to implement proactive IP recovery based on Multiple Routing Configurations (MRC) [15], [16], [17].

Since MRC has been thoroughly evaluated in existing papers, this paper will focus on IP Fast Reroute evaluations including a comparison with MRC. Francois and Bonaventure have recently presented some figures evaluating IP Fast Reroute for link failures [18]. In this paper we present a more thorough evaluation and we also include node failures.

The rest of this paper is organized as follows. Section II will describe the concept of IP Fast Reroute as specified in IETF. Section III will give an evaluation of IP Fast Reroute. In section IV we will describe MT-routing and how it can be used with MRC providing proactive

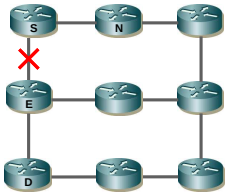

Fig. 1. Illustrates how a loop can occur when using an alternative next hop.

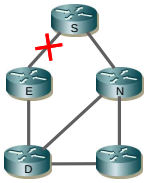

Fig. 2. Illustrates how ECMP provides loop-free alternative next hop.

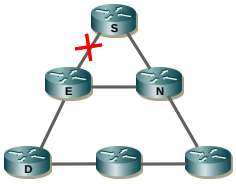

Fig. 3. Illustrates a failure scenario where the condition for Link-protecting alternates is fulfilled.

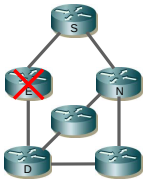

Fig. 4. Illustrates a failure scenario where the condition for Node-protecting alternates is fulfilled.

IP recovery. Section V will discuss some pros and cons of the two schemes, and section VI will conclude the paper.

II. IP Fast Reroute Framework

The IP Fast Reroute work in IETF is motivated by pointing at new applications, like VoIP, requiring milliseconds recovery [10]. In addition, they emphasis the problems with instability and micro-loops during IP reconvergence. The goal of the framework is to prevent such instabilities and provide fast recovery. The main idea behind this approach is to allow a router whose neighbor link/node has failed to forward traffic to a precomputed alternative next-hop until the router installs the new primary next-hops based upon the changed network topology.

Due to the connectionless nature of IP, such hop-by-hop backup next hops may cause looping of packets. Figure 1 shows an example. For simplicity we let all link weights be 1, and we assume shortest path routing. Let S be the local router that detects the failure on the interface towards node E, E be the node downstream of the failed interface, D the destination and N the alternative next-hop router. Suppose S forwards the packet to the alternative next hop N. N will then forward the packet back to S since the cost from N to D is shortest via S. Hence, a loop has occurred.

The IP Fast Reroute framework specifies three basic categories of pre-calculated repair paths that can avoid loops as described above. These are Equal cost multi-paths (ECMP), loop-free alternates and Multi-hop repair paths.

A. Equal cost multi-paths (ECMP)

Figure 2 gives an example of ECMP. Equal cost multi-paths exist when the cost from S via E to D equals the cost from S via N to D:

$$cost(S \vdash E, D) = cost(S \vdash N, D)$$

Figure 2 illustrates how ECMP provides a loop-free alternative path when using shortest path routing. The path cost from N directly to D is lower than the cost from N to D via S.

If the link weights do not equal the hop count, node E may be included in the alternative ECMP. If the network is supposed to handle node failures, one must also make sure that node E is not included in the ECMP.

B. Loop-free alternates

These are alternative paths that are longer than the primary path, but still provide loop-free routing to the destination. Such a path exists when a direct neighbor (N) of the detecting node (S) has a path to the destination which can be guaranteed to not traverse the failure, i.e. the failed link or node is not included in the alternative path.

IP Fast Reroute specifies a condition for Link-protecting alternates and a more restrictive condition for Node-protecting alternates [11].

1) Link-protecting alternates: To guarantee loop-free alternates for link failures, the following condition must hold:

$$cost(N, D) < cost(N, S) + cost(S, D)$$

Figure 3 shows a failure scenario where this condition holds. In this scenario node N would not route the packets back to the failure.

2) Node-protecting alternates: Alternate next-hops for node failures require a stronger condition than what is the case for link failures. If node E failed in figure 3, node N would choose node E as next hop towards destination D, and hence node N can not be used as a backup next hop to protect the failure of node E.

To guarantee loop-free alternates for node failures, the following condition must hold:

$$cost(N, D) < cost(N, E) + cost(E, D)$$

Figure 4 gives an example of a failure scenario where the condition holds for a failure of node E.

[11] also gives a special condition for broadcast and non-broadcast multiple access (NBMA) links. In this paper we do not consider such links.

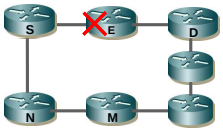

Fig. 5. Illustrates a failure scenario which can be covered using the U-turn strategy.

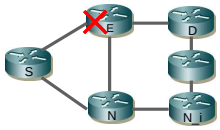

Fig. 6. Illustrates a failure scenario that can be solved by tunneling.

C. Multi-hop repair paths

Normally, one can not expect that there exist ECMP or other loop-free alternates for all failure scenarios in a network. For such failure scenarios, [10] points to multi-hop repair paths, which means that there is a router more than one hop away from the detecting node (S), from which traffic will be forwarded to the destination without traversing the failure. The following subsections will overview some alternative multi-hop repair path strategies.

1) U-turn Alternates: U-turn alternates [19] are best explained through an example. In figure 5, when node S uses node N as backup next hop, node N must not use the primary next hop S towards D, but rather use the loop-free node protecting alternate (node M) towards D. This means that node N is not allowed to give packets from S a u-turn back to S. Such alternates may be used when no other loop-free alternates exist, and when the following conditions are fulfilled:

1) N must have S as its primary next hop towards D.
2) N must have a node-protecting alternate M for the destination D with respect to node S.
3) If covering node failures, the path from M to D must not contain the failed node E.

2) Multi-hop Tunneling: Tunneling can be used to steer the packets to a node N_i that is i hops away from S and that has a loop-free path to the destination D without traversing the failure. [20] suggest using IP in IP tunneling for this purpose [21]. Tunneling is performed without signaling, using only packet encapsulation. With respect to figure 6, S would encapsulate the packets affected by the failure in a header using node N_i as destination. N_i would decapsulate the packets and forward them according to normal routing towards the destination D. This approach can only be used when the packets tunneled from S to N_i do not traverse the failure.

3) Multi-hop Tunneling using Not-via Addresses: None of the approaches listed above can guarantee full coverage. Recently, a full coverage tunneling strategy using Not-via addresses has been proposed [22]. The semantics of a Not-via address are that a packet addressed to a Not-via address must be delivered to the router with that address, not via the neighboring router on the interface to which that address is assigned. In other words, one must ensure that the packets affected by the failure of router E are delivered to router M that according to the primary route to destination D is downstream of E (figure 7). Routers advertise Not-via addresses for all their neighbor components. A Not-via address is used by other routers when the corresponding component has failed. Each router in the network must calculate the best path to each Not-via address or group of addresses. The path is calculated without the component that the Not-via address is meant to protect. The router S that detects the failure will then encapsulate the packets and address them to the Not-via address that router M has advertised for the particular failure (Me in figure 7). The routing table of router S will have a destination address Me which have been calculated on a topology without router E.

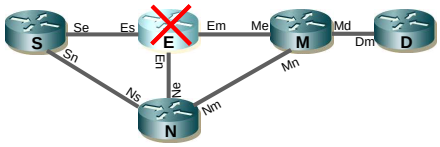

Fig. 7. Gives an example of how the Not-via approach could be configured. Node S is the node detecting a failure of node E, node M is the tunnel end-point and node D is the destination.

III. Evaluation of Fast Reroute

As described above, not all failure scenarios can be covered by simple ECMPs or other loop-free alternates. The IP Fast Reroute approach suggest to obtain full coverage by using different types of multi-hop repair paths, i.e. U-turns, multi-hop tunnels of different lengths or tunneling using Not-via addresses as described in section II-C. From a management and control point of view, such a mix of mechanisms is not favorable, due to complexity and potential mis-configuration. In this section we will investigate to what extent only ECMP and loop-free alternates can provide failure coverage, and in addition we will see the type and amount of multi-hop repair paths needed.

A. Method

For given topologies we have calculated the number of failure scenarios that are covered by ECMP, other loop-free alternates and different multi-hop repair path strategies. For multi-hop tunnels, it also calculates the length of the tunnels.

We have calculated these properties on a wide range of both synthetic and real-world topologies. The synthetic topologies have been generated using the Brite topology tool [23] with both the Waxman model [24] and the

Generalized Linear Preference (GLP) model [25]. The number of nodes has varied from 32 to 128. The average node degree has been 4 or 6 for Waxman and 3.6 for GLP. For the combinations of model, number of nodes and node degree we have generated 100 random topologies. The real world topologies are 60 POP-level topologies collected from Rocketfuel [26] and 8 POP-level topologies from Oliver Heckmann [27]. These topologies include among others AT&T, German Telecom and Sprint.

As routing function we use shortest path. For simplicity we configure all link weights to 1, which means that the cost from one node to another equals the number of hops. It should be noted that this is a favorable condition for IP Fast Reroute, since more ECMPs and other loop-free alternates exist when only considering hop count.

We have calculated the IP Fast Reroute approach as follows. For each node in the network (detecting node) we let each neighbor component (link or node) fail. For each failure we then calculate the relevant rerouting alternative to each destination which would in the normal case have traffic routed through the failure. We do not look at the scenarios where the failed node is also the destination.

IP Fast Reroute configuration strategies: We envision 3 different configuration scenarios of IP Fast Reroute. Link failures are the most common failure, and hence a strategy for only providing link failure coverage may be an alternative (1). If, however there is a requirement for handling node failures as well, configuring for covering node failures is an alternative (2). Such an alternative will also cover link failures. In both strategy 1 and 2 we allow the use of multi-hop repair paths. Configuration strategy 3 represents a scenario where the routers or the operator does not support the use of multi-hop repair paths, e.g. due to the complexity.

1. Covering link failures: In this case we configure IP Fast Reroute according to the condition for link protecting alternates. When no loop-free link protecting alternates exist, we configure u-turns, multi-hop tunneling or tunneling using Not-via addresses, respectively.

2. Covering node failures: In this case we configure IP Fast Reroute according to the condition for node protecting alternates. When no loop-free node protecting alternates exist, we configure u-turns, multi-hop tunneling or tunneling using Not-via addresses, respectively. Link failures will also be covered with this configuration strategy.

3. Loop-free alternates only: In this case we configure IP Fast Reroute according to the condition for node protecting alternates. If the condition for node protection alternates is not satisfied, we try to configure according to the less restrictive condition for link protecting alternates. This strategy will use no multi-hop repair paths, and hence some failure scenarios may not be covered.

B. Results

Figures 8 to 10 and table I and II present results from our calculations. In the figures, the x-axis denotes the fast reroute alternative used (ECMP, other loop-free alternates, U-turn, tunnels of different lengths and Not-via). The y-axis denotes in percent the amount of failure scenarios that required a particular fast reroute alternative. The graph (line) names indicate ”configuration strategy - topology model - number of nodes - average node degree”. Also note that if a value is plotted in a figure, the value is higher than 0.

In figure 8 we observe that configuration strategy 1 (Link protection) has more other loop-free alternates (“Others”) than configuration strategy 2 (node protection). This is due to a less restrictive condition for link protecting alternates. Configuration strategy 2 requires slightly more U-turns, more tunnels, longer tunnels and more Not-via tunnels than strategy 1. We also note that both methods need U-turns, tunnels and Not-via tunnels to obtain full coverage.

Figure 9 illustrates that networks with higher average node degrees provide more ECMPs and loop-free alternates and less U-turns, tunnels and Not-via tunnels than networks with lower node degrees. Still, high-degree networks need U-turns, tunnels and Not-via tunnels to obtain full coverage.

In figure 10 we compare the effect of different network sizes. We have that they differ slightly with respect to ECMPs and other loop-free alternates, and that the network size does not seem to influence the amount and length of U-turns, tunnel and Not-via tunnels. Larger networks have longer maximum path lengths and hence a greater variation of path lengths than smaller networks. Therefore, the probability for having ECMPs is less for larger networks than for smaller networks when using our failure model. This is reflected in figure 10.

Table I shows the results for configuration strategy 3 (loop-free alternates only) for different types of networks. The table denotes the amount (in percent) of failure scenarios that could be covered by ECMP (ECMP Node/Link) and node protecting alternates (Others Node/Link), the amount of node failure scenarios that could not be covered (No Node), the amount of link failure scenarios that could be covered in addition to the node failure scenarios (Others Link Only) and the number of

link failure scenarios that could not be covered at all (No Link). Both ECMP and node protecting alternates cover node and link failures. “Others Link Only” is then the amount of the unsuccessful node scenarios that could be covered by link protecting alternates. From the table we can conclude that a great amount of the node failure scenarios cannot be covered, and also that not all link failure scenarios can be covered. For the Heckmann topologies, the real world topologies with lowest node-degree, the amounts are 23.7 % for links and 44.5 % for nodes.

Table II shows the results for configuration strategy 1 and 2 for real world networks. We see that the tendencies encountered for the synthetic networks are also valid for the real world networks, and that full coverage can not be obtained without Not-via tunneling.

This section has showed that not all failure scenarios can be covered by a simple configuration of ECMPs and other loop-free alternates (configuration strategy 3). One alternative to obtain full coverage is to successively try to configure ECMP, other loop-free alternates, U-turns, general tunnels and Not-via tunnels (configuration strategy 1 and 2). From a management point of view this alternative provides a mix of relatively complex mechanisms to implement and configure. The figures also show that using the U-turn strategy does not increase the coverage considerably. Tunneling using Not-via addresses is the only mechanism that stand-alone provides full coverage. The results have shown that no networks can be fully covered without Not-via tunneling, and since Not-via tunneling is inevitable, the best approach may be to use Not-via tunneling only.

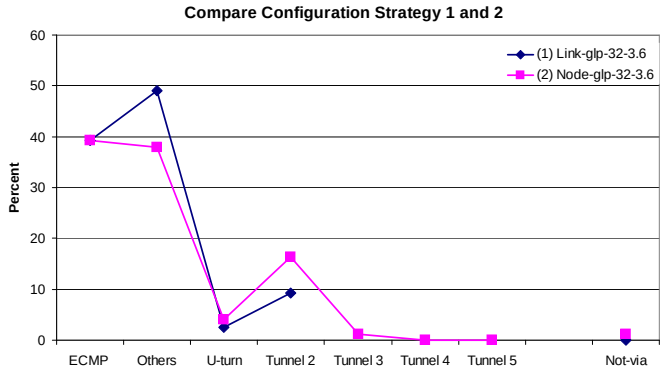

Fig. 8. Comparison of the amount of failure scenarios that can be covered by different IP Fast Reroute alternatives when using configuration strategies 1 (Link) and 2 (Node).

IV. Multi-Topology (MT) routing

Within IETF, initiatives have been taken to specify Multi-Topology routing, both for ISIS [12] and OSPF

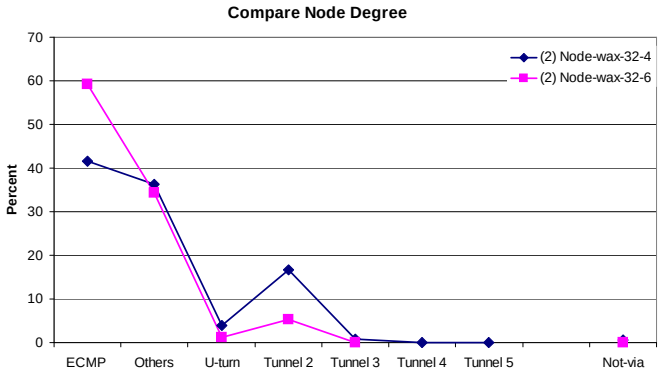

Fig. 9. Comparison of the amount of failure scenarios that can be covered by different IP Fast Reroute alternatives when using different average node degrees.

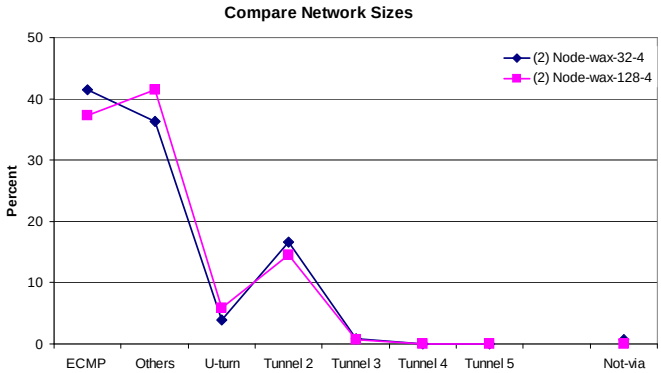

Fig. 10. Comparison of the amount of failure scenarios that can be covered by different IP Fast Reroute alternatives when using networks of different sizes.

[13], [14]. These drafts describe extensions to ISIS and OSPF in order to define independent IP topologies called Multi-Topologies (MTs). They propose that the MT extensions can be used for computing different paths for unicast traffic, multicast traffic, different classes of service, or an in-band network management topology.

These independent topologies are maintained separately, i.e. the routers exchange information of each topology in independent link state advertisements. An MT may contain all or only some of the nodes, all or only some of the links, and the same or different link weights than the original topology.

A router must be able to decide the topology to use

TABLE I

Success Percentage of loop-free alternates only

Topology	ECMP Node/Link	Others Node/Link	No Node	Others Link Only	No Link
wax-32-4	41.5	36.3	22.2	13.5	8.7
wax-32-6	59.2	34.2	6.6	5.8	0.8
wax-128-4	37.3	41.5	21.2	4.7	16.5
Rocketfuel	66.6	10.7	22.7	17.3	5.4
Heckmann	28.9	26.4	44.7	23.7	21.0

TABLE II

REAL WORLD TOPOLOGIES FOR CONFIGURATION STRATEGY 1 (LINK) AND 2 (NODE)

Topology	Conf strategy	ECMP	Others	Uturn	Tunnel 2	Tunnel 3	Tunnel 4	Tunnel 5	Not-via
Rocketfuel	Link (1)	66.6	28.0	1.5	3.1	0.1	0	0	0.7
Rocketfuel	Node (2)	66.6	10.7	1.2	16.0	0.8	0.07	0.005	4.6
Heckmann	Link (1)	28.9	50.1	3.8	10.5	3.6	1.8	0.4	0.9
Heckmann	Node (2)	28.9	26.3	3.1	26.8	5.3	2.0	0.5	7.0

for a particular packet. The MT-routing drafts suggest to solve this by marking of packets, e.g. setting specific bits in the header.

A. *Multi-topology routing and recovery*

Multiple Routing Configurations (MRC) has been developed as a concept for proactive recovery in connectionless networks [15], [16]. The main idea of this concept is to build backup topologies (called configurations) of the original topology in such way that each link and/or node is isolated in at least one of the backup topologies. This means that shortest path routing in a backup topology will not select a path through a node or a link that is isolated in that topology. A node that is isolated can still be reached, but it can not be used as transit node.

Figure 11 shows an example of how to isolate each node once using 2 backup topologies. In the failure-free case, the full topology is used for routing. When a router detects that its neighbor has failed, it starts routing the traffic according to the topology where the failed router is isolated. For the routing to be correct, the other nodes in the network must also be informed so that traffic will be routed according to this backup topology all the way to the egress (destination) node. We suggest that the node that detects the failure marks the packets with the identifier of the backup topology. The other nodes in the network will then be aware of what topology to route according to.

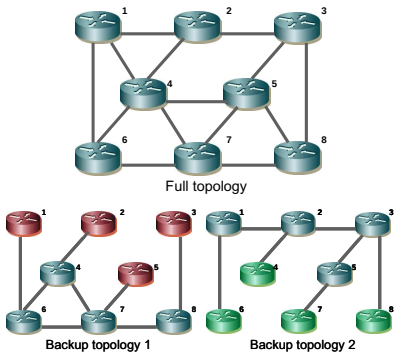

Fig. 11. Shows the original topology and how each node can be isolated once using 2 backup topologies.

In previous papers, we have demonstrated that MRC scale with respect to the number of backup topologies needed [16], [28]. Even for very large networks 3-5 backup topologies were enough to cover all nodes or links. Our evaluations has been performed on the same topologies as we have used for the IP Fast Reroute evaluations in section III.

The routing in a backup topology is restricted, due to several isolated nodes or links in each topology. For this reason the backup paths lengths will be slightly longer than what are the lengths for full IP re-convergence. [16], [28] show that these are indeed within acceptable bounds, i.e. 5-16 % longer dependent on the number of backup configurations used. In [17], we also demonstrated that the traffic load in the network after failure is quite similar to IP full re-convergence. In [29], Gjessing details how our scheme can be implemented using MT routing [12], [13], [14] and stub routers [30]. Also Menth and Martin [31] state that MT-routing can be used for recovery, however they are not detailing any approach for generating suitable backup topologies.

Although, the MT-routing approach seems very suitable for recovery, there exist some open issues. The first is how to mark the packets with the identifier of the backup topology. The MT-routing drafts [12], [13], [14] suggest to use either the DSCP field or subnetting for IPv4 or an extension header for IPv6. The second open issue relates to the fact that the IETF drafts recommend that a packet should be routed according to the same topology end-to-end in a network, i.e. they do not recommend to change topology in intermediate nodes. This recommendation is meant to prevent looping between topologies. If we restrict our method to only change topology once, we can guarantee loop-free routing. Both marking of packets and changing topology are open issues that can be supported in the current MT-standards. Using MT-routing for recovery will however impose a third requirement that is currently not supported. A router detecting a failure needs to know what backup topology should be used when a particular neighbor link or node fails. This requires that each node has a simple table that denotes what topology isolates each of its neighbors. Our goal is that this should be supported in future MT-routing specifications.

V. Discussion and Comparison

An efficient recovery approach should at least handle any single link and node failure in a network. In addition, complexity and overall performance are important aspects to consider. Section III-B has shown that IP Fast Reroute does not provide complete failure coverage using loop-free alternates only ([11]). To obtain full coverage an alternative is to complete the coverage by also using U-turns, general tunnels or Not-via tunneling as evaluated in section III. Such an IP Fast Reroute approach would create a a complex mix of different mechanisms. Since tunneling using Not-via addresses is the only strategy providing complete failure coverage, we find an approach based on only the Not-via strategy as the most viable approach. This section will therefore discuss and conceptually compare IP Fast Reroute using Not-via addresses and Multi-topology Routing using MRC ([15]).

From a management point of view, the Not-via strategy is considered rather complex. The IGP must handle an increased address space, and the changes to the IGP are considered extensive. MT-routing with MRC provides a simple global view of the isolated components. Each backup configuration is maintained in a separate configuration (topology) exclusively used for recovery.

Both methods will require some extra state to be stored in the routers. The absolute value of this extra state will be heavily dependent on the implementation choices. Some intuitive strategies could be as follows. Not-via would need one extra destination address for each Not-via address in the network in the original routing table. MRC would need one additional routing table per configuration. These additional routing tables may not contain all global IP addresses if a local address scheme could be implemented for the backup configurations.

Both Not-via and MRC require that each router perform more than one shortest path tree (SPT) calculation. MRC will require one SPT calculation per configuration while Not-via will require one SPT calculation per failed component. Not-via may decrease the number of SPT calculation by using Shared Risk Groups and hence calculating one SPT per group. Both methods may decrease the number and complexity of the SPT calculation by using an incremental approach. Note that the number of SPT calculations may be similar if the SRG for Not-via equals the isolated components for MRC.

Compared to a full IP re-convergence, Not-via and MRC will provide somewhat longer backup paths. Figure 12 compares the path lengths of normal failure-free routing, OSPF full re-convergence (reroute), Optimal local (the best achievable path length for local recovery), Not-via and MRC with 3 and 7 backup configurations. We observe that Not-via provide shorter paths than MRC with 3 backup configurations and a bit longer paths than MRC with 7 backup configurations. It should be noted that these results show the best case for Not-via, i.e. that only one component is isolated at the same time (no SRGs). MRC isolates more than one component in each configuration, and hence the number of links available for backup routing is restricted. Increasing the number of backup configurations improves the path lengths.

Easy support for shared risk groups (SRGs) is considered an important property of a recovery scheme. Not-via can support SRGs, however the simplicity of this support can be disputed due to a complex address scheme. MRC provides global common backup topologies that isolates a group of components simultaneously, and hence may provide more easy SRG support.

Tunneling of packets in a network requires an additional header, and hence the packet overhead will increase. Adding a header may also enforce fragmentation and defragmentation of packets due to the MTU. Although this may not degrade the performance considerably, a scheme without tunneling would be preferable. Not-via tunneling have no other option than adding a header, while MRC may be content by a few bits in the existing header.

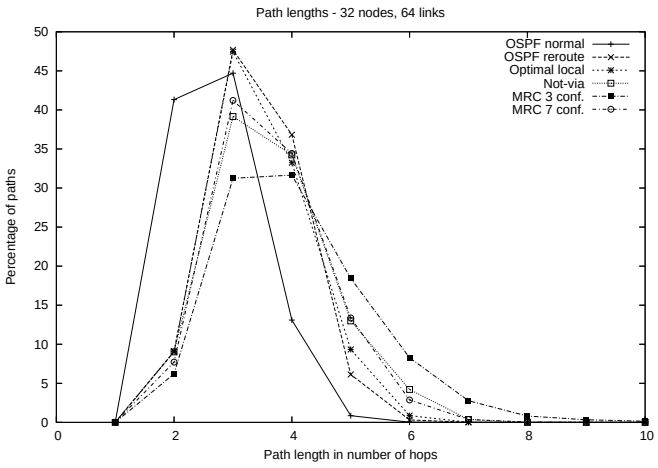

Fig. 12. Compares the backup path lengths of normal failure-free routing, OSPF full re-convergence (reroute), Optimal local (the best achievable path length for local recovery), Not-via and MRC with 3 and 7 backup configurations. We have used 100 Brite generated topologies with 32 nodes and 64 links.

VI. Conclusion

In this paper we have discussed two different approaches for handling proactive recovery in IP networks.

Both IP Fast Reroute and MT-routing using MRC can be supported by standardization processes within IETF. The authors have previously published performance evaluations on MRC, and this paper has therefore focused on IP Fast Reroute evaluations. For IP fast reroute to offer full coverage, Not-via tunneling seems inevitable, and hence we believe a scheme based on Not-via tunneling only as the most viable. We have argued that both Not-via and MT-routing using MRC are viable candidates for solving proactive connectionless IP recovery, but also how MRC seems to provide a simpler scheme, particularly from a management point of view. For these reasons we strongly encourage future MT-routing standards to fully support this recovery approach.

REFERENCES

[1] D. D. Clark, "The Design Philosophy of the DARPA Internet Protocols," *SIGCOMM, Computer Communications Review*, vol. 18, no. 4, pp. 106–114, aug 1988.

[2] C. Alaettinoglu *et al.*, "Towards milli-second igp convergence," IETF Internet Draft, Nov. 2000, draft-alaettinoglu-ISIS-convergence-00.txt.

[3] A. Basu and J. G. Riecke, "Stability Issues in OSPF Routing," in *Proceedings of SIGCOMM 2001*, San Diego, California, USA, August 2001, pp. 225–236.

[4] P. Pillay-Esnault, "Ospf refresh and flooding reduction in stable topologies," IETF Internet Draft, 2003, draft-pillay-esnault-ospf-flooding-07.txt.

[5] A. Shaikh, A. Varma, L. Kalampoukas, and R. Dube, "Routing stability in congested networks: Experimentation and analysis," in *Proceeding of SIGCOMM 2000*, 2000.

[6] G. L. Choudhury, "Prioritized treatment of specific ospf version 2 packets and congestion avoidance," IETF Internet Draft, 2004, draft-ietf-ospf-scalability-09.txt.

[7] P. Narvaez, K.-Y. Siu, and H.-Y. Tzeng, "New Dynamic SPT Algorithm based on a Ball-and-String Model," *IEEE/ACM Trans. Netw.*, vol. 9(6), pp. 706–718, 2001.

[8] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving sub-second IGP convergence in large IP networks," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 2, pp. 35 – 44, July 2005.

[9] G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot, "Analysis of link failures in an ip backbone," in *2nd ACM SIGCOMM Workshop on Internet Measurement*, Nov. 2002, pp. 237–242.

[10] M. Shand, "IP Fast Reroute Framework," IETF Internet Draft, 2005, draft-ietf-rtgwg-ipfrr-framework-04.txt.

[11] A. Atlas and A. Zinin, "Basic specification for IP fast reroute: Loop-free alternates," Internet Draft, 2005, draft-ietf-rtgwg-ipfrr-spec-base-04.txt.

[12] T. Przygienda, N. Shen, and N. Sheth, "M-ISIS: Multi topology (MT) routing in IS-IS," Internet Draft, 2005, draft-ietf-isis-wg-multi-topology-11.txt.

[13] P. Psenak, S. Mirtorabi, A. Roy, L. Nguen, and P. Pillay-Esnault, "MT-OSPF: Multi topology (MT) routing in OSPF," IETF Internet Draft, Apr. 2005, draft-ietf-ospf-mt-04.txt.

[14] S. Mirtorabi and A. Roy, "Multi-topology routing in OSPFv3 (MT-OSPFv3)," Internet Draft, 2005, draft-ietf-ospf-mt-ospfv3-00.txt.

[15] A. Kvalbein, A. F. Hansen, T. Čičić, S. Gjessing, and O. Lysne, "Fast IP network recovery using multiple routing configurations," in *Proceedings of INFOCOM*, Apr. 2006.

[16] A. F. Hansen, T. Čičić, S. Gjessing, A. Kvalbein, and O. Lysne, "Resilient routing layers for recovery in packet networks," in *Proceedings of International Conference on Dependable Systems and Networks (DSN)*, June 2005.

[17] A. F. Hansen, A. Kvalbein, T. Čičić, S. Gjessing, O. Lysne, T. Jensen, and O. N.Østerbø, "Fast, effective and stable IP recovery using resilient routing layers," in *The 19th International Teletraffic Congress (ITC19)*, 2005.

[18] P. Francois and O. Bonaventure, "An evaluation of ip-based fast reroute techniques," in *CoNext'05*, Toulouse, France, 2005.

[19] A. Atlas, "U-turn alternates for IP/LDP local protection," Expired Internet Draft, 2077, draft-atlas-ip-local-protect-uturn-03.txt.

[20] S. Bryant, C. Filsfils, S. Previdi, and M. Shand, "IP fast reroute using tunnels," IETF Internet Draft, Apr. 2005, draft-bryant-ipfrr-tunnels-02.txt.

[21] W. Simpson, "IP in IP tunneling," IETF RFC 1853, 1995.

[22] S. Bryant, M. Shand, and S. Previdi, "IP fast reroute using not-via addresses," Internet Draft, 2005, draft-bryant-shand-IPFRR-notvia-addresses-01.txt.

[23] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *Proceedings of IEEE MASCOTS*, Aug. 2001, pp. 346–353.

[24] B. M. Waxman, "Routing of multipoint connections," *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 9, pp. 1617–1622, Dec. 1988.

[25] T. Bu and D. Towsley, "On distinguishing between internet power law topology generators," in *Proceedings of IEEE INFOCOM*, New York, June 2002, pp. 638–647.

[26] "Rocketfuel topology mapping," WWW, http://www.cs.washington.edu.

[27] O. Heckmann, M. Piringer, J. Schmitt, and R. Steinmetz, "On realistic network topologies for simulation," in *ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, 2003.

[28] A. Kvalbein, A. F. Hansen, T. Čičić, S. Gjessing, and O. Lysne, "Fast recovery from link failures using resilient routing layers," in *Proceedings 10th IEEE Symposium on Computers and Communications (ISCC)*, June 2005.

[29] S. Gjessing, "Implementation of two resilience mechanisms using multi topology routing and stub routers," in *to appear in Advanced International Conference on Telecommunications (AICT'06)*, Feb. 2006.

[30] A. Retana, L. Nguyen, R. White, A. Zinin, and D. McPherson, "OSPF stub router advertisement," IETF RFC 3137, June 2001.

[31] M. Menth and R. Martin, "Network resilience through multi-topology routing," in *Proceedings of the 5th International Workshop on Design of Reliable Communication Networks (DRCN)*, Oct. 2005.