High Availability in the Future Internet

Levente Csikor, Gábor Rétvári, and János Tapolcai

MTA-BME Future Internet Research Group
High Speed Networks Laboratory
Dept. of Telecommunications and Media Informatics
Budapest University of Technology and Economics,
Magyar tudósok körútja 2., H-1117 Budapest, Hungary
`{csikor,retvari,tapolcai}@tmit.bme.hu`

Abstract. With the evolution of the Internet, a huge number of real-time applications, like Voice over IP, has started to use IP as primary transmission medium. These services require high availability, which is not amongst the main features of today's heterogeneous Internet where failures occur frequently. Unfortunately, the primary fast resilience scheme implemented in IP routers, Loop-Free Alternates (LFA), usually does not provide full protection against failures. Consequently, there has been a growing interest in LFA-based network optimization methods, aimed at tuning some aspect of the underlying IP topology to maximize the ratio of failure cases covered by LFA. The main goal of this chapter is to give a comprehensive overview of LFA and survey the related LFA network optimization methods, pointing out that these optimization tools can turn LFA into an easy-to-deploy yet highly effective IP fast resilience scheme.

Keywords: IP Fast ReRoute, network optimization, remote loop-free alternates, reliable networks, protection, failures.

1 Introduction

Current Internet has reached the level of reliability, where Internet and cloud services are widely spreading among users. This gives an increasing push on the service providers to operate the Internet without any interruption and slowly win the trust of most of the potential users. We expect that the reliability of IP networks will further improve in the future, and Internet will become a critical infrastructure for the society. Reliability means that at any given time the connection is ensured throughout the network, and a failure is handled so fast that virtually no packet loss is noticed by the end-users.

Nowadays, not just Internet Service Providers (ISPs) and end-users are concerned, but many other multimedia suppliers started to gain a foothold in this field and broadcast digital content over the IP (Internet Protocol). Moreover, traditional telephony is already being replaced by IP based telephony in order to reduce costs and provide more options, e.g., send text, media and data simultaneously during the phone call. Due to this continuous technical change and digital convergence (collectively referred to as information and communication

technologies [42]), a huge number of (real-time) multimedia applications use IP network as primary transmission medium, which is the first driving force for a more reliable IP communication infrastructure. Furthermore, not just the scope of contents is growing but the number of the newly connected consumers and terminal equipments as well. The population in the world is currently growing at a rate of around 1.10% per year, and the expected population will be 8 billion in 2025[1]. Today, about two billion people are using the Internet, while six billion people are already a subscriber of some mobile services at the end of 2011 [25]. What is more, in the near future not only human beings will be connected to the Internet all the time, but many machines used day by day will have unique IP addresses and will be accessible from anywhere. This will lead to an emerging communication form, called M2M (Machine-to-Machine), which will likely generate a huge part of total traffic. Moreover, due to the evolution of mobile infrastructure, most of the new users will access all the digital content through their smartphones, increasing the traffic that has to be delivered at the same time. Through the development of already used entertainment services, for instance, television broadcasting and Video on Demand (VoD), two-thirds of the world's mobile traffic will be video by 2016. Note that not just mobile phones account for mobile traffic, since in 2011 the number of mobile-connected tablets reached 34 million, and each tablet generated 3.4 times more traffic than smartphones [10]. Cisco has fore-casted that, if not just mobile users are considered, then *almost 90% of all consumer traffic will account for real-time video broadcasting*, e.g., IPTV, P2P streaming, VoD. Since the improving quality of the content (e.g., High Definition movies, 4K, lossless audio coding) involves a growing size of media streams, the aforementioned proportion will likely grow even further. Necessarily, the Internet has to keep up with these real-time applications, which require continuous and reliable connections. Consequently, if a link or node fails in the network, it does not only cause routing instability, but a whacking portion of the traffic will be lost. Note that low latency is substantial for the standard Internet applications as well [50].

Therefore, high availability has become an all-important factor in operational networks. However, studies over the past decade have shown that the reliability of Internet falls short of the five nines (99.999%) availability, which is readily available in the public-switched telephone network (PSTN) [28]. Note that reliability plays an important role in pure mobile networks too, since an average user will be more satisfied with a constant guaranteed access speed than a, however high, abruptly fluctuating service in densely populated areas [18]. If high availability were always provided, this would open the door for completely new opportunities. For instance, imagine that in the future a qualified doctor could easily and uninterruptedly treat a patient thousands of miles away through remote surgery, by the means of the basic communication network.

Availability in the Internet is severely plagued by component failures that occur frequently due to various reasons, such as physical interruptions, flapping

---

[1] `http://www.worldometers.info/world-population/` (accessed in Jan 2013)

interfaces[2], etc. To overcome these issues, the IETF (Internet Engineering Task Force) defined a native IP fast resilience framework, called IP Fast ReRoute (IPFRR [44]), and a basic IPFRR specification called Loop-Free Alternates (LFA [4]). LFA, however, has a substantial drawback in that it cannot protect every possible failure case. Recently, a generalization of LFA has been published, called Remote Loop-Free Alternates (rLFA [5]), which can provide higher protection, but still has limitations evading the possibility to have an ultimate solution. Many IPFRR proposals have appeared to vanquish this, but due to the additional complexity implementation and deployment of these proposals have been rather scarce so far (see Sec. 2 for details). This leads ISPs to rely completely on the only IPFRR scheme available in IP routers today, LFA, and investigate network optimization techniques to adapt their topology to provide the best protection with LFA possible.

The aim of this chapter is to overview the evolution of Internet from reliability perspective. In particular, we focus on the intriguing questions of fast IP resilience. In Section 2, we review former IPFRR proposals and their disadvantages. Afterwards, in Section 3 we show how LFA works and then, in Section 4, we present a deeper analysis of failure case coverage provided by LFA and we discuss how this can be improved with emerging LFA network optimization techniques. We extend these results to the emerging Remote LFA IPFRR specification in Section 5. At last but not least, in Section 6 we conclude our message and sketch up further questions that need to be investigated and solved in the future.

2   Related Work and Background

In operational networks more than 85% of unplanned failures affect only links and almost the half of them are transient [35], i.e., 50% of all failures last less than a minute [23]. In order to reduce the latency and increase the reliability, additional network functionality is essential to recognize the failure and reroute the affected packets rapidly around the failed component.

Formerly, failures were handled by the intra-domain routing protocols, such as OSPF (Open Shortest Path First [38]) or IS-IS (Intermediate System To Intermediate System [24]). After a failure, the information about it was distributed throughout the network so that every router can recalculate the shortest paths with the failed component removed. This process is called re-convergence, and, depending on network size and routers' shortest path calculation efficiencies, it can take between 150 ms and a couple of seconds [30,26]. It is obvious that this is beyond what real-time applications can afford, even if a small delay can be tolerated via buffering.

To overcome these issues, the IETF (Internet Engineering Task Force) defined a framework, called IP Fast ReRoute (IPFRR [44]), for native IP protection in order to reduce failure reaction time to tens of milliseconds. It tries to avoid the global re-convergence with *local rerouting* and *pre-computed detours*, converting

---

[2] A hardware or software failure of an interface can cause the router to announce it alternately as "up" and "down".

the reaction scheme of standard IP networks into faster proactive protection mechanisms [46]. The matter of these approaches is that the router adjacent to the failure tries to solve the problem individually by means of pre-computed alternate routing tables, which are installed long before any failure occurs.

As one of the first approach, a basic specification for IPFRR was defined by the IETF, called Loop-free Alternates (LFA [4]), which is simple, standardized and already available in today's routers [48,27]. In LFA, when a failure occurs, the adjacent router tries to pass the packet to an alternate neighbor, who still has a functioning path to the destination. However, such neighbor does not always exist, evading LFA for providing 100% failure case coverage in every network.

Therefore, in the past few years many IPFRR proposals have appeared, but the majority of them require additional management burden, complexity and non-standard functionality in IP forwarding and routing protocols. Some of them change the traditional IP's destination based forwarding [32,51,16,3], while others use signaling to indicate that a packet is on detour [22,11,1,49,31,33]. On the other hand, there are methods, which use explicit tunneling to avoid the failed components [7,17,36,6]. Proposals in [43,37] have topological requirements, whilst the mechanism proposed in [29] uses a central server to pre-compute forwarding decisions for common failure scenarios and download them into the routers. Accordingly, at the moment none of these proposals is available in IP routers, since they need modifications to the existing protocols, making LFA the only deployable option for ISPs.

In order to improve the level of fast protection provided by LFA, the IETF has published a generalization called Remote LFA (rLFA) [5]. This method provides additional backup connectivity when none can be provided by the basic mechanism. However, even though rLFA can provide higher reliability, it still inherits topology dependence from pure LFA, and thus providing 100% failure case coverage with pure IP remains to be an open question.

3   Providing Fast Protection with LFAs

Probably, the easiest way to understand how basic LFA and remote LFA work, is through an example. Consider the network depicted in Fig. 1, where all links have unit costs. Suppose that node e wants to send a packet to node d' and its default next-hop[3] f is unreachable, since the link between them went down. In this case, e has to find an alternate neighbor, which will not pass the packet back, i.e., which still has a path to the destination unaffected by the failure. Fortunately, node b fulfills this requirement, so e can reroute the traffic destined to d, towards b. Next, suppose that node s wishes to send packets to node d and eventually link (s, a) fails. Now, s can only reach node b. However, since node b has an ECMP (Equal Cost Multiple Path) to node d and it does not know about the failure, it can pass the packet back to s causing a loop. Therefore, this failure case cannot be protected with simple LFA. However, if a tunnel existed between node s and

---

[3] In IP routing, the next router along the shortest path to a destination is called *next-hop.*

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

e (marked by a black dashed line in Fig. 1), then node e, now being an indirect neighbor of s, would become a LFA for d, thereby protecting the link (s, a). Consequently, when a link cannot be entirely protected with local LFA neighbors, the protecting routers try to seek the help of a remote LFA staging point. This repair tunnel endpoint needs to be a node in the network reachable from the source without traversing the failed component. Furthermore, the repair tunnel endpoint needs to be a node from which packets will normally flow towards their destinations without being attracted back to the failed component. These supplementary logical links, used be remote LFA, are provided by tunnels.

Accordingly, after the remote node receives the package, it sends it towards the primary destination. Note that these tunnels are provided by a simple label stack in an MPLS/LDP (Multiprotocol Label Switching/Label Distribution Protocol) [2] enabled network, which is practically relevant nowadays. However, there exist MPLS networks with RSVP–TE (Reservation Protocol–Traffic Engineering) extension, wherein IPFRR is not the only option for fast protection [39,21]. On the other hand, suppose now that node s wants to send a packet to node d', and the link (s, b) fails. Then, (s, b) cannot be protected for a lack of a suitable tunnel since all nodes, whose shortest paths are unaffected by the failure, can only be reached from s through the failed (s, b) itself. This suggests that while the use of rLFA can definitely provide higher protection level than pure LFA, it still does not facilitate full protection for all failure cases in a general.

4   Analyzing and Improving LFA-Based Fast Protection

The most important question concerning LFA is to *analyze how it performs in different network topologies*, what are the fundamental lower and upper bounds of failure case coverage, and *how protection coverage could be improved*.

The authors of [41] made the first steps in this direction, in that they gave a graph-theoretical characterization of LFA protection efficiency. To measure this LFA efficiency in an arbitrary network graph G, the following LFA failure coverage metric is defined [4,41]:

$$\eta(G) = \frac{\text{\#LFA protected source-destination pairs}}{\text{\#all source-destination pairs}} \tag{1}$$

The rLFA failure case coverage $\mu(G)$ can be defined in a similar way.

Table 1. LFA graph extension results for link protection

Topology			Uniform cost		Weighted	
Name	n	m	η_0	Gr_η	η_0	Gr_η
AS1221	7	9	0.833	1	0.809	2
AS3257	27	64	0.946	3	0.923	11
AS6461	17	37	0.919	2	0.933	4
Abilene	12	15	0.56	6	0.56	8
Italy	33	56	0.784	13	0.784	20
Germany	27	32	0.695	5	0.695	12
AT&T	22	38	0.823	6	0.823	13
Germ_50	50	88	0.801	22	0.9	21

The authors in [41] observe that the quintessential worst-case graphs for IPFRR are rings, i.e., cycle graphs wherein all nodes are of degree two [16,8]. In particular, the LFA failure case coverage in an arbitrary 2-connected network is bounded by $\frac{1}{n-1} \leq \eta(G) \leq 1$, and the lower bound is attained for even rings. In a special case, where all links have unit costs, full LFA protection, i.e., $\eta(G) = 1$ can only be reached if every link is contained in at least one triangle[4]. This suggests that complete graphs, chordal graphs [20] and maximal planar graphs have full LFA coverage in the uniform cost case. If arbitrary link costs are taken into account, then the aforementioned condition is even more stricter [41]. However, the latter condition is only sufficient, but not necessary.

As a way of an improvement, the so called *LFA graph extension* problem was also studied in [41], which asks for augmenting the network with the fewest number of new links in order to reach 100% LFA coverage. For example, if the network depicted in Fig. 1 is augmented with 4 new links (namely, (s,d), (b,d), (d,e), and (d',e)), then every link will be contained in at least one triangle, i.e., $\eta(G)$ will be 1.

Unfortunately, adding unit cost links to the network cannot always be afforded by network operators, since it will definitely change at least one shortest path, which might have been previously tweaked with great accuracy to match the needs of the network in terms of load balancing, traffic engineering, etc. [19,47]. In order to prevent this, it was suggested in [41] that the cost of the new link should be larger than the length of the longest shortest path. With this in mind, the example network (Fig. 1) should be augmented with 6 new links with sufficiently high costs in order to attain full protection. Finding the smallest number of additional new links proved a very hard problem, i.e., it is NP-complete [41]. Therefore, an Integer Linear Program (ILP) and an approximating greedy heuristic were developed to solve this problem.

The algorithms were studied in real-world network topologies inferred from Rocketfuel dataset [34] and SNDLib [45]. Succinct results are shown in Table 1,

---

[4] Triangle is a cycle of length 3.

where n denotes the number of nodes, while m indicates the number of links in the network. The initial LFA coverage is marked by η_0, and Gr_η denotes the number of new links added by the greedy heuristic[5] in order to attain full LFA coverage. First observation is that in each link cost case, on average the initial LFA coverage is about 80% and never reaches 100%. For smaller topologies, a few number of new links have to be added to reach full protection, while in larger and sparser networks this number is significantly more. In particular, in the German backbone ca. one fourth of the number of links originally existing have to be added. Additionally, more links are needed for full coverage in weighted networks than in uniform cost graphs. The main conclusion of these results is that some networks readily lend themselves to LFA graph extension, and in many cases adding only 2-4 new links can boost up the LFA failure case coverage close to 100%.

In a subsequent study [40], another improving approach, called the *LFA cost optimization* problem, was examined. The problem asks for finding an optimal cost setting in a network, which produces 100% failure case coverage. This problem proved NP-complete as well. The complexity directly comes from the fact that shortest paths change during the optimization process, therefore it is possible that altering a link cost can provide protection for a certain source-destination pair, but it may eliminate LFAs to other destinations. For instance, consider network depicted in Fig. 1 again and suppose that node c wants to send a packet to node e and the link between them fails. This failure cannot be protected in the current link cost setting, since node d, the only available neighbor of c, will pass the packet back. However, if the cost of the link (c, d) would be, say 4, then d would route the affected traffic to a detour, since its next-hop towards e would be a. By means of this small modification, the initial LFA failure coverage can be improved by 15%. Additionally, it was also proved in [40] that the average node degree Δ plays an important role to determine the attainable upper bound of LFA failure coverage. As a matter of fact, for any connected graph G with $n > 2$:

$$\eta(G) \leq \tfrac{n}{n-1}(\Delta - 2) + \tfrac{2}{n-1} \ .$$

This suggests that in a large but sparse networks (i.e., where $\Delta < 3$), the protection provided by LFA can be very poor.

In order to solve LFA cost optimization problem, a computationally expensive optimal algorithm as well as efficient approximating heuristics were proposed [40,15]. A brief result can be seen in Table 2, where n and m denote the number of nodes and the number of links in the network, respectively. The column Δ indicates the average node degree of the network, while $\eta(G, c)$ marks the initial LFA coverage. At last but no least, $\eta(G, c*)$ represents the LFA coverage reached by cost optimization. One can easily observe that in most cases, chiefly when the average node degree is higher than 3.5, close to perfect LFA coverage can be attained. There were, however, some exceptional topologies where LFA

---

[5] Since the greedy heuristic is faster and it performs almost the same as the ILP, in this paper we concentrated on the greedy approach exclusively.

Table 2. LFA cost optimization results for link protection

Topology				Cost optimization	
Name	n	m	Δ	$\eta(G,c)$	$\eta(G,c*)$
AS1221	7	9	2.57	0.809	0.833
AS3257	27	64	4.74	0.923	1
AS6461	17	37	4.35	0.933	1
Abilene	12	15	2.50	0.56	0.674
Italy	33	56	3.39	0.784	0.944
Germany	27	32	2.94	0.695	0.911
AT&T	22	38	3.45	0.823	0.987
Germ_50	50	88	3.52	0.9	0.966

cost optimization was less appealing. For such networks, combining LFA cost optimization with LFA graph extension could be a viable option.

Since both methods are effective ways of improving LFA coverage in operational networks, the question promptly arises as to what extent the combination of these two approaches can be effective for LFA-based network optimization. However, it is not obvious how these methods should be set together, in particular, how many links should be added and when should cost optimization be executed. The authors in [12] investigated just these questions. They showed by extensive simulations that the combination of the approaches performs the best if we only add 1 new link at a time and then execute a round of cost optimization. These two phases should follow each other until 100% LFA failure case coverage is attained. Their results suggest that the combined algorithm can significantly reduce the number of additional links (on average by more than 50%) necessary for reaching full protection with LFA providing an intermediate solution.

5   Improving Fast Protection with Remote LFA

The wide spectrum of LFA network optimization strategies presented so far provide a rich set of options for operators to choose from, according to their own preference on whether it is economically more feasible to add new links, change link costs, or do both in order to reach high LFA-protection in their network. Nevertheless, in the near future operators should think about upgrading to the remote LFA specification instead, since it has become available in commercial routers [9] and can definitely provide higher protection.

The authors of [13] spearheaded the research to determine the topological requirements and the protection efficiency of remote LFA as well as to find optimization methods to tweak the network for 100% rLFA coverage. They showed that if the implementations support the extended version of rLFA, then every unit cost network is fully protected out of the box. Furthermore, it turned out that, unlike pure LFA, rLFA provides almost full protection in ring topologies [13]. Moreover, there is no 2-edge-connected network, which would not have

Table 3. rLFA graph extension results for link protection

Topology			Uniform cost	
Name	n	m	μ_0	Gr_μ
AS1221	7	9	0.833	1
AS3257	27	64	0.954	1
AS6461	17	37	1	0
Abilene	12	15	0.833	1
Italy	33	56	0.951	2
Germany	27	32	0.882	1
AT&T	22	38	0.8875	2
Germ_50	50	88	1	0

at least 33% of rLFA coverage, while almost the half of the source-destination pairs are protected in every 2-node-connected network.

In order to provide higher protection when rLFA coverage is small, LFA graph extension was adapted from [41] to rLFA. The resultant *rLFA graph extension* problem asks how many links one must add in a real-world network topology to achieve full rLFA coverage. At the moment, it is unclear whether this problem is also NP-complete, although it seems likely that it actually is. The greedy heuristic from [41] was adopted to solve this problem and in [14] more approximating heuristics are examined. A brief view of the results can be found in Table 3, where the notations are the followings: n and m denote the number of nodes and the number of already existing links in the network, respectively. The initial rLFA coverage is indicated by μ_0, whilst Gr_μ marks the number of new links that have to be added to achieve full rLFA protection.

The first observation is that there were two networks, which are initially fully protected, while every other network required less than 3 new links to reach 100% failure case coverage. Furthermore, the number of links have to be added is much less than when only simple LFA capable routers are present, especially, in the Italy and Germ_50 topologies. The results also indicate that on average 3.6 new links are necessary to attain full rLFA protection, while this number is 14.5 in the case of pure LFA.

At the moment, the results for rLFA only cover the case of unit cost networks. A comprehensive investigation of the case of arbitrary cost networks is currently an open problem.

6   Conclusions and Future Work

Due to the increasing number of end-users and real-time services, one of the most important challenges of the future Internet architecture is to be resilient

against failures. Since failures occur frequently, IP networks should be able to survive component failures and ensure the service continuity.

In the past few years, a plethora of proposals appeared on how to modify the current IP routers in order to overcome the fast resilience problem, but none of them became industry standard due to their substantial added complexity. As of today, only the Loop-Free Alternates and the Remote LFA specifications have found their ways to operational IP and MPLS/LDP networks, thanks to their simplicity and deployability, and it seems highly unlikely that this situation will change in the future. Since the protection coverage of LFA and rLFA crucially depends on both the underlying network topology and the link costs, the tasks to uncover the intricacies of this dependence as well as to optimize a network for high LFA/rLFA protections have become compelling. This chapter intended to give a comprehensive survey on the state-of-the-art on these pressing problems.

It turned out that there exist many real-world network topologies, where LFA and rLFA can only protect a fraction of possible failures. Fortunately, with LFA and rLFA the protection coverages can be often boosted close to 100% in these networks just by cleverly adding a few new links. However there are still some cases, where LFA/rLFA graph extension cannot be afforded due to limited resources, but in these cases optimizing link costs is still a good approach to increase the coverage. Moreover, the most efficient approach is the combination of the above two. In our experience just adding one or two links and tuning the link costs carefully always resulted in a high increase in failure coverage.

There are still many challenges, which should be solved before it can be attained in most real-world IP networks. In the aforementioned optimization methods the traffic engineering and load balancing issues were not considered at all. Thus, rerouting the traffic after a failure may not results service continuity because some links in the network become congested. Besides, current LFA was developed to protect single failure cases only. As future work we also plan to deal with multiple failures.

With all these in mind, researchers and the industry are facing with an intriguing and complex challenge of converting IP network to a reliable and highly available architecture in the future.

References

1. Amund, K., Fosselie, H.A., Čičic, T., Stein, G., Olav, L.: Multiple routing configurations for fast IP network recovery. IEEE/ACM Trans. Netw. 17(2), 473–486 (2009), doi:http://dx.doi.org/10.1109/TNET.2008.926507
2. Andersson, L., Minei, I., Thomas, B.: LDP specifiaction. RFC 5036 (October 2007)

3. Antonakopoulos, S., Bejerano, Y., Koppol, P.: A simple IP fast reroute scheme for full coverage. In: 2012 IEEE 13th International Conference on High Performance Switching and Routing (HPSR), pp. 15–22 (2012), doi:10.1109/HPSR.2012.6260822
4. Atlas, A., Zinin, A.: Basic specification for IP fast reroute: Loop-Free Alternates. RFC 5286 (2008)
5. Bryant, S., Filfils, C., Previdi, S., Shand, M., So, N.: Remote LFA FRR. IETF DRAFT (December 2012)
6. Bryant, S., Filsfils, C., Previdi, S., Shand, M.: IP fast reroute using tunnels. Internet Draft (2007)
7. Bryant, S., Shand, M., Previdi, S.: IP fast reroute using Not-via addresses. Internet Draft (2010)
8. Čičic, T.: An upper bound on the state requirements of link-fault tolerant multi-topology routing. IEEE ICC 3, 1026–1031 (2006)
9. Cisco Systems: IP Routing: OSPF Configuration Guide, Cisco IOS Release 15.2S - OSPF IPv4 Remote Loop-Free Alternate IP Fast Reroute (downloaded: April 2012)
10. Cisco VNI: Global mobile data traffic forecast update, 2011-2016 (February 2012)
11. Császár, A., Enyedi, G., Tantsura, J., Kini, S., Sucec, J., Das, S.: IP fast re-route with fast notification. Internet Draft (June 2012)
12. Csikor, L., Nagy, M., Rétvári, G.: Network optimization techniques for improving fast IP-level resilience with Loop-Free Alternates. Infocommunications Journal 3(4), 2–10 (2011)
13. Csikor, L., Rétvári, G.: IP fast reroute with remote loop-free alternates: the unit link cost case. In: Proc. RNDM, pp. 16–22 (2012)
14. Csikor, L., Rétvári, G.: On providing fast protection with remote loop-free alternates: Analyzing and optimizing unit cost networks. Submitted to Telecommunication Systems Journal (2013)
15. Csikor, L., Rétvári, G., Tapolcai, J.: Optimizing igp link costs for improving IP-level resilience with loop-free alternates. Computer Communications (2012), doi:10.1016/j.comcom.2012.09.004
16. Enyedi, G., Rétvári, G., Cinkler, T.: A novel loop-free IP fast reroute algorithm. In: Pras, A., van Sinderen, M. (eds.) EUNICE 2007. LNCS, vol. 4606, pp. 111–119. Springer, Heidelberg (2007)
17. Enyedi, G., Szilágyi, P., Rétvári, G., Császár, A.: IP Fast ReRoute: lightweight Not-Via without additional addresses. In: INFOCOM Mini-Conf. (2009)
18. Ericcson Consumer Lab: Smartphone usage experience. Ericsson Consumer Insight Summary Report (2013)
19. Fortz, B., Rexford, J., Thorup, M.: Traffic engineering with traditional IP routing protocols. IEEE Comm. Mag. 40(10), 118–124 (2002)
20. Golumbic, M.C.: Algorithmic Graph Theory and Perfect Graphs, 2nd edn. Elsevier Science (2004)
21. Hock, D., Hartmann, M., Menth, M., Pioro, M., Tomaszewski, A., Zukowski, C.: Comparison of ip-based and explicit paths for one-to-one fast-reroute in MPLS networks. Springer Telecommunication Systems Journal, 1–12 (2011), doi:10.1007/s11235-011-9603-4
22. Hokelek, I., Fecko, M., Gurung, P., Samtani, S., Cevher, S., Sucec, J.: Loop-free IP fast reroute using local and remote LFAPs. Internet Draft (February 2008)
23. Iannaccone, G., Chuah, C.N., Mortier, R., Bhattacharyya, S., Diot, C.: Analysis of link failures in an IP backbone. In: ACM SIGCOMM Internet Measurement Workshop, pp. 237–242 (2002)

24. ISO: Intermediate ststem-to-intermediate system (is-is) routing protocol. ISO/IEC 10589 (2002)
25. ITU-T: ICT facts and figures (2011), `http://www.itu.int/ITU-D/ict/facts/2011/material/ICTFactsFigures2010.pdf` (downloaded: January 2013)
26. Iyer, S., Bhattacharyya, S., Taft, N., Diot, C.: An approach to alleviate link overload as observed on an IP backbone. In: INFOCOM (2003)
27. Juniper Networks: Junos 9.6 routing protocols configuration guide (2009)
28. Kuhn, D.R.: Sources of failure in the public switched telephone networks. IEEE Computer 30(4), 31–36 (1997)
29. Kwong, K.W., Gao, L., Guerin, R., Zhang, Z.L.: On the feasibility and efficacy of protection routing in IP networks. In: INFOCOM, long version is available in Tech. Rep. 2009. University of Pennsylvania (2010)
30. Labovitz, C., Malan, G.R., Jahanian, F.: Internet routing instability. IEEE/ACM Transactions on Networking 6(5), 515–528 (1998)
31. Lakshminarayanan, K., Caesar, M., Rangan, M., Anderson, T., Shenker, S., Stoica, I.: Achieving convergence-free routing using failure-carrying packets. In: Proc. SIGCOMM (2007)
32. Lee, S., Yu, Y., Nelakuditi, S., Zhang, Z.L., Chuah, C.N.: Proactive vs reactive approaches to failure resilient routing. In: INFOCOM (2004)
33. Li, A., Yang, X., Wetherall, D.: Safeguard: safe forwarding during route changes. In: CoNEXT, pp. 301–312 (2009)
34. Mahajan, R., Spring, N., Wetherall, D., Anderson, T.: Inferring link weights using end-to-end measurements. In: ACM IMC, pp. 231–236 (2002)
35. Markopoulou, A., Iannacone, G., Bhattacharyya, S., Chuah, C.N., Diot, C.: Characterization of failures in an IP backbone. In: Proc. IEEE Infocom (March 2004)
36. Menth, M., Hartmann, M., Martin, R., Čičic, T., Kvalbein, A.: Loop-free alternates and not-via addresses: A proper combination for ip fast reroute? Computer Networks 54(8), 1300–1315 (2010), doi:10.1016/j.comnet.2009.10.020
37. Merindol, P., Pansiot, J.J., Cateloin, S.: Providing protection and restoration with distributed multipath routing. In: International Symposium on Performance Evaluation of Computer and Telecommunication Systems, SPECTS 2008, pp. 456–463 (2008)
38. Moy, J.: OSPF version 2. RFC 2328 (April 1998)
39. Pan, P., Swallow, G., Atlas, A.: Fast reroute extensions to RSVP-TE for LSP tunnels. RFC 4090 (2005)
40. Rétvári, G., Csikor, L., Tapolcai, J., Enyedi, G., Császár, A.: Optimizing igp link costs for improving IP-level resilience. In: Proc. of DRCN, pp. 62–69 (October 2011)
41. Rétvári, G., Tapolcai, J., Enyedi, G., Császár, A.: IP Fast ReRoute: Loop Free Alternates revisited. In: INFOCOM, pp. 2948–2956 (2011)
42. Sallai, G.: Defining infocommunications and related terms. Acta Polytechnica Hungarica 9(6), 5–15 (2012)
43. Schollmeier, G., Charzinski, J., Kirstädter, A., Reichert, C., Schrodi, K., Glickman, Y., Winkler, C.: Improving the resilience in IP networks. In: Proc. HPSR (2003)
44. Shand, M., Bryant, S.: IP Fast Reroute framework. RFC 5714 (2010)
45. SNDLib: Survivable fixed telecommunication network design library, `http://sndlib.zib.de` (downloaded: April 2012)
46. Sterbez, J., Cetinkaya, E.K., Hameed, M.A., Jabbar, A., Qian, S., Rohrer, J.P.: Evaluation of network resilience, survivability, and disruption tolerance: Analysis, topology generation, simulation and experimentation. Springer Telecommunication Systems Journal, 1–32 (2011), doi:10.1007/s11235-011-9573-6

47. Swallow, G., Bryant, S., Andersson, L.: Avoiding equal cost multipath treatment in MPLS networks. RFC 4928 (June 2007)
48. Systems, C.: Cisco IOS XR Routing Configuration Guide, Release 3.7 (2008)
49. Čičic, T., Hansen, A., Apeland, O.: Redundant trees for fast IP recovery. In: Broadnets, pp. 152–159 (2007)
50. Vulimiri, A., Michel, O., Godfrey, P.B., Shenker, S.: More is less: reducing latency via redundancy. In: Hotnets (2012)
51. Zhong, Z., Nelakuditi, S., Yu, Y., Lee, S., Wang, J., Chuah, C.N.: Failure inferencing based fast rerouting for handling transient link and node failures. In: INFOCOM (2005)