Evolution of IP Fast-Reroute Strategies

Andrzej Kamisiński
Department of Telecommunications
AGH University of Science and Technology
Kraków, Poland
andrzejk@agh.edu.pl

***Abstract*—Due to increasing requirements related to quality of service and network dependability, several IP Fast-Reroute mechanisms have been designed to limit the consequences of one or more simultaneous failures in routed computer and communication networks. However, the existing mechanisms vary in terms of the underlying concept, the maximum number of handled failures in the network, the employed signaling methods, and the operation mode. In this paper, a high-level overview of the main common features and conceptual differences of the selected IP Fast-Reroute solutions and other related mechanisms is presented and discussed, as the first step towards a more extensive and in-depth analysis. In addition, future research directions are proposed in the context of the selected challenges identified in existing networks.**

***Index Terms*—Network dependability, Network resilience, IP Fast Reroute**

I. Introduction

Considering the increasing reliance of users, service providers, and several institutions on the communication infrastructure, dependability of networks has become one of the critical factors influencing many decisions made by network operators. At the same time, failures of network components are inevitable and effective protection methods are necessary to avoid service downtime, the resulting loss in revenue, and the potential damage to a company's reputation. In routed networks without additional protection mechanisms, failures often result in service disruptions that may last for periods of hundreds of milliseconds or several seconds, depending on the underlying technology and specific configuration of network devices. Every such disruption may lead to traffic losses or forwarding loops occurring while the network is re-converging on the new topology. Thus, increasing requirements with respect to quality of service and dependability of computer and communication networks have led to the development of IP/MPLS Fast-Reroute mechanisms [1], [2].

The Loop-Free Alternates (LFA) scheme [3] is the basic specification for implementing the IETF IP Fast-Reroute standard [2] in computer and communication networks. Due to its simplicity and transparent operation with respect to other routing solutions, it has been implemented in many existing network devices. At the same time, LFA does not guarantee full failure coverage in the network. Thus, one of the goals of the following research efforts was to improve the failure coverage in different scenarios assuming single failures of links or nodes, and then also multiple simultaneous failures of network components. Fig. 1 presents the related timeline showing when the selected proposals were published.

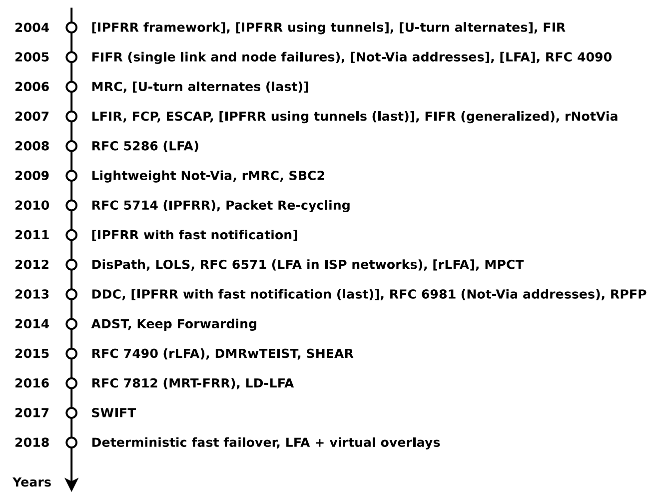

Fig. 1. The Selected Proposals Related to IP Fast Reroute — Timeline (FIR: Failure Insensitive Routing, LFIR: Loop-Free FIR, FIFR: Failure Inferencing-based Fast Rerouting, FCP: Failure-Carrying Packets, ESCAP: Efficient SCan for Alternate Paths, LFA: Loop-Free Alternates, rLFA: Remote LFA, MRC: Multiple Routing Configurations, rMRC: Relaxed MRC, LOLS: Localized On-Demand Link State Routing, MPCT: Minimum Protection Cost Tree, DDC: Data-Driven Connectivity, RPFP: Reverse Path Forwarding Protection, ADST: Arc-Disjoint Spanning Trees, DMRwTEIST: Disjoint Multipath Routing with Three Edge-Independent Spanning Trees, LD-LFA: OpenFlow-Based LFAs with Loop Detection). Names in the square brackets refer to the first published IETF drafts, unless stated otherwise.

In this paper, a high-level overview of the main characteristic features and conceptual differences of the selected IP Fast-Reroute solutions and other related mechanisms is presented and discussed. Considering the large number of solutions in this area, the proposed classification is intended to be the first step towards a more extensive and in-depth analysis covering both the early proposals and the most recent ideas. The basic concepts and definitions related to dependable systems are introduced in [4], [5].

The remainder of this paper is structured as follows. In Section II, the general classification of the selected existing IP Fast-Reroute strategies is introduced and discussed with respect to the related criteria, as well as advantages and disadvantages. Section III presents the possible research directions leading to future IP Fast-Reroute solutions. Finally, Section IV concludes the paper.

II. Classification of Existing Strategies

In this section, the selected existing IP Fast-Reroute methods are classified with respect to the following criteria: the underlying concept, the maximum number of simultaneous failures, the employed signaling mechanisms, and the operation mode. The corresponding sections contain brief explanations of the selected criteria, as well as references to example proposals. For a survey of the selected early IP Fast-Reroute and MPLS Fast-Reroute strategies, including the discussion of micro-loops, avoidance methods, and some design guidelines, the reader is referred to [6]. Additional evaluation results are provided in [7].

A. Underlying Concept

The existing IP Fast-Reroute proposals are based on different underlying ideas. Tab. I presents the selected concepts shared by particular solutions.

TABLE I
Underlying Concept — Examples

Description	References
Forwarding to a loop-free alternate	[3], [8]–[10]
Input interface-aware routing	[8], [11]–[17]
Additional/extended forwarding tables and other data structures	[16], [18]–[24]
Tunneling	[9], [25]–[28]
Redundant spanning trees	[14], [17], [22], [26], [29]–[31]
Failure coverage analysis and improvements (e.g., network graph modifications, virtual overlays, optimization of link costs, loop detection extensions)	[10], [32]–[40]

The early concepts following the IETF IP Fast-Reroute standard [2] include solutions which forward packets to available loop-free alternates. One definite advantage of this approach is its simplicity. However, depending on network topology, it may not be able to provide full coverage of different single-failure scenarios, which was one of the important issues addressed in more recent proposals.

The second group of solutions was designed around the idea of making rerouting decisions based on information about the input network interfaces on which packets arrive. If a packet is received on a different interface than expected, it is a potential sign of a detour. In particular, if the packet arrives on an interface being also its primary output interface leading to the corresponding destination, a forwarding loop is detected and the packet should be rerouted due to possible failure experienced by the downstream nodes. This method can also be used to deal with issues related to asymmetric routing schemes, which is a significant advantage.

The solutions belonging to the third group either rely on additional forwarding tables (for instance, see [18]), or they assume that the primary FIB is extended, so that it can store some necessary information used by the protection mechanism (an example is presented in [41]). This information may also be stored in other data structures, depending on specific requirements. The immediate consequence is that more system resources are used, which might not be acceptable in some deployments, especially such that involve resource-constrained devices. In addition, modifications of existing devices would also be necessary to be able to benefit from this approach.

IP Fast-Reroute methods relying on tunneling encapsulate the received packets and forward them to an alternative next-hop router, whenever the corresponding primary next-hop router is not directly reachable. The new destination address for each encapsulated packet may be determined either based on information disseminated by the routing protocol with appropriate extensions (see, e.g., [25]) or with the aid of another signaling mechanism.

Several existing solutions are based on the concept of redundant spanning trees. Protection against failures is achieved by identifying multiple spanning trees within the network graph, and using the suitable trees to forward packets to the respective destinations. The trees may be determined according to some additional constraints — for example, they should have a common root node and they can be edge/arc- or node-disjoint to ensure that in specific failure scenarios, at least one of the available trees can be used to deliver packets to the destination. In some cases, it might be necessary to pass additional information to the downstream nodes about the trees that have already been selected by the upstream nodes to forward a packet (for instance, see [22]), which requires an additional signaling mechanism, such as reserved bits in the packet header. Consequently, existing network devices would need to be updated to support the new method.

Considering the limited failure coverage of many existing solutions, as well as varying difficulty and cost of their deployment, one of the goals of the following research efforts was to provide the necessary tools to help network operators decide which methods they should deploy in their environments. Another objective was to determine whether any modification of the network would be necessary to increase the number of covered failure scenarios. In particular, several related analyses have been made in the context of the widely-supported LFA scheme, leading to new insights and important solutions [32], [34]–[40].

B. Maximum Number of Simultaneous Failures

As single-link failures represent the majority of failures in typical operational IP backbone networks [42], several IP Fast-Reroute strategies have been designed to provide adequate protection. At the same time, most of the existing solutions cannot deal with multiple simultaneous failures effectively, limiting the ability of networks to recover from failures caused by such events as natural disasters or area-based attacks. Thus, the corresponding methods have been designed to help solve this issue, while in the selected cases, the strategies proposed for single-link failures also cover failure scenarios affecting an entire Shared Risk Link Group (SRLG). Tab. II presents the classification of the selected solutions with respect to the maximum number of simultaneous failures that may still be handled properly, as long as the network graph remains connected.

TABLE II
MAXIMUM NUMBER OF SIMULTANEOUS FAILURES

Description	References
Single link failures[a]	[2], [3], [8]–[13], [15], [18]–[21], [25], [29], [30], [43]–[52]
Single node failures	[2], [3], [8], [10], [11], [15], [18]–[21], [25], [30], [45]–[47], [50]–[53]
Dual link failures	[14], [23], [53], [54]
Multiple failures[b]	[16], [17], [22], [24], [28], [55]–[58]

[a]Including failure of an SRLG: [8], [25], [49], [51].
[b]Full coverage only for single failures: [28].

C. Signaling

Some existing IP Fast-Reroute strategies rely on additional information being exchanged between nodes participating in the recovery procedure. Thus, the corresponding signaling mechanisms are required and should be taken into account in new network designs, as well as during preparations for deployment in existing environments. On the other hand, several methods are available that do not need any signaling scheme. The selected solutions are classified in Tab. III according to the corresponding signaling requirements.

TABLE III
SIGNALING

Description	References
Packet header[abc]	[10], [14], [19], [20], [22], [25], [30], [48], [53], [55]–[58]
Control messages	[23], [30], [47], [49], [51], [58]
Not required[de]	[2], [3], [8], [9], [11]–[13], [15]–[18], [24], [28], [43]–[46], [50], [54]
Other	[29]

[a]Not necessary if virtual links are used: [58].
[b]Required only if used together with Loop-Free Alternates: [25].
[c]Note that [10] also requires a functional SDN network using the OpenFlow protocol.
[d]In the case of SRLG protection, [8] requires an additional signaling mechanism between neighboring routers.
[e]Note that [9] assumes that the network provides a functional tunneling mechanism.

In the case of the first group of solutions, signaling relies on additional bits which are added to the IP header of forwarded packets. The number of necessary bits varies, depending on the specific way in which each method operates. For example, the solution presented in [19] marks packets with an identifier of the selected routing configuration and the authors recommend that the IPv4 DSCP field be used for this purpose. Alternatively, other suggested methods include IPv6 extension headers or tunneling with the use of the private address space. Another proposal, Failure-Carrying Packets [55], stores within each packet additional information about failed links required for routing that packet. Although the authors argue that the number of necessary entries is generally expected to be small, it is also important to note that it may depend on network topology and specific failure scenarios.

The second important signaling scheme is based on control messages and is typically associated with the corresponding protocol. It may either assume that messages are exchanged between two nodes, or it may rely on controlled flooding to deliver messages from the source to the other nodes (for example, see [47]). In fact, the necessary information may also be carried in the control messages of existing routing protocols, such as the OSPF protocol. An example solution which could benefit from this method is presented in [30][1].

D. Operation Mode

The centralized or distributed operation of IP Fast-Reroute methods has important implications on their performance. In the centralized case, routing decisions are usually made by a logically-centralized network controller having a complete view of the network topology and the state of particular devices. Two strong advantages of this approach are flexibility and the ability to make informed decisions. At the same time, the controller needs to be designed in such a way that allows it to take appropriate action without significant delay while guaranteeing both high availability and consistency among the controller units (for a detailed discussion in the context of Software-Defined Networks, the reader is referred to [59]). Considering the failure detection time, the transmission delay between the controller and other nodes, as well as potentially large amount of information that needs to be collected and processed before making a decision, this set of requirements has already been identified as a challenge [60]–[62]. On the other hand, distributed solutions are able to react almost instantly by triggering local protection mechanisms, but they often have limited knowledge about the most recent changes in the network topology. Thus, other ideas have also been proposed that assume a mixed approach, combining the advantages of the centralized and distributed approaches [10], [23], [62], [63]. Tab. IV summarizes the selected solutions with respect to their operation mode.

TABLE IV
OPERATION MODE

Description	References
Distributed	[2], [3], [8], [9], [11]–[22], [24], [25], [28], [30], [43]–[49], [51], [53], [54], [56]–[58]
Centralized[a]	[21], [29]
Mixed	[10], [23], [50], [55], [62]

[a]Note that the solutions proposed in [21], [29] can also be implemented in a distributed fashion, provided that specific conditions are satisfied.

III. FUTURE RESEARCH DIRECTIONS

In the case of typical computer and communication networks, several IP Fast-Reroute mechanisms have been proposed and the related problems have been well studied over the last years. The known solutions vary in terms of the failure coverage and the maximum number of handled failures in the network. Some of them can deal with multiple

[1]Note that the authors modified the original concept of Not-Via addresses in such a way that the number of required IP addresses is significantly smaller.

simulatenous failures, yet still, there are not many that can be easily deployed in existing environments. Thus, one possible direction of the ongoing research efforts might be to provide tools and solutions which are simple to manage, simple to implement, and compatible with existing network hardware and software, depending on the general use case. At the same time, increasing requirements related to network resilience, as well as increasing adoption of SDN and virtualization technologies, call for further improvements of existing resilience provisioning methods, while also addressing several important design challenges, such as: logically-centralized control plane and its potential reliance on the same data plane it manages, failure detection time, transmission delay, and implementation of the necessary logic in generic switching devices. While benefiting from the flexibility offered by SDN, future fast-reroute solutions should be designed in such a way that does not increase the complexity of SDN considerably.

IV. Conclusion

In this paper, a high-level overview of the main characteristic features and conceptual differences of the selected IP Fast-Reroute solutions and other related mechanisms has been presented and discussed. In addition, the related timeline was shown. Although numerous solutions have been proposed in the past, the vast majority of them have not been implemented in off-the-shelf network devices. Among the likely reasons are both the overall complexity of the design and additional requirements which made the solutions incompatible with existing devices and network protocols. Considering the selected challenges identified in existing networks, the corresponding future research directions were outlined. The proposed classification is intended to be the first step towards a more extensive and in-depth analysis.

V. Acknowledgments

The author would like to thank Stefan Schmid and Gábor Rétvári for early discussions and valuable comments related to the scope of the paper. This article is based upon work from the COST Action CA15127 ("Resilient Communication Services Protecting End-user Applications from Disaster-based Failures — RECODIS") [64], supported by COST (European Cooperation in Science and Technology).

RECODIS
Resilient communication services protecting end-user applications from disaster-based failures

References

[1] P. Pan, G. Swallow, and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels," Internet Requests for Comments, RFC Editor, RFC 4090, May 2005. [Online]. Available: https://tools.ietf.org/html/rfc4090

[2] M. Shand and S. Bryant, "IP Fast Reroute Framework," Internet Requests for Comments, RFC Editor, RFC 5714, Jan 2010. [Online]. Available: https://tools.ietf.org/html/rfc5714

[3] A. Atlas and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates," Internet Requests for Comments, RFC Editor, RFC 5286, September 2008. [Online]. Available: https://tools.ietf.org/html/rfc5286

[4] A. Avizienis, J. C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, pp. 11–33, 2004.

[5] J. Rak, *Resilient Routing in Communication Networks*, 1st ed. Springer Publishing Company, Incorporated, 2015.

[6] A. Raj and O. C. Ibe, "A Survey of IP and Multiprotocol Label Switching Fast Reroute Schemes," *Comput. Netw.*, vol. 51, no. 8, pp. 1882–1907, Jun. 2007. [Online]. Available: http://dx.doi.org/10.1016/j.comnet.2006.09.010

[7] M. Gjoka, V. Ram, and X. Yang, "Evaluation of ip fast reroute proposals," in *2007 2nd International Conference on Communication Systems Software and Middleware*, Jan 2007, pp. 1–8.

[8] A. Atlas, "U-turn Alternates for IP/LDP Fast-Reroute," Network Working Group, Internet-Draft, IETF, Internet-Draft, Feb 2006. [Online]. Available: https://tools.ietf.org/pdf/draft-atlas-ip-local-protect-uturn-03.pdf

[9] S. Bryant, C. Filsfils, S. Previdi, M. Shand, and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)," Internet Requests for Comments, RFC Editor, RFC 7490, April 2015. [Online]. Available: https://tools.ietf.org/html/rfc7490

[10] W. Braun and M. Menth, "Loop-free alternates with loop detection for fast reroute in software-defined carrier and data center networks," *Journal of Network and Systems Management*, vol. 24, no. 3, pp. 470–490, Jul 2016. [Online]. Available: https://doi.org/10.1007/s10922-016-9369-9

[11] K. Xi and H. Chao, "ESCAP: Efficient SCan for Alternate Paths to Achieve IP Fast Rerouting," in *Global Telecommunications Conference, 2007. GLOBECOM '07. IEEE*, Nov 2007, pp. 1860–1865.

[12] 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, April 2007.

[13] G. Enyedi and G. Rétvári, "A Loop-Free Interface-Based Fast Reroute Technique," in *2008 Next Generation Internet Networks*, April 2008, pp. 39–44.

[14] A. Gopalan and S. Ramasubramanian, "Multipath routing and dual link failure recovery in ip networks using three link-independent trees," in *2011 Fifth IEEE International Conference on Advanced Telecommunication Systems and Networks (ANTS)*, Dec 2011, pp. 1–6.

[15] S. Antonakopoulos, Y. Bejerano, and P. Koppol, "A simple ip fast reroute scheme for full coverage," in *2012 IEEE 13th International Conference on High Performance Switching and Routing*, June 2012, pp. 15–22.

[16] B. Yang, J. Liu, S. Shenker, J. Li, and K. Zheng, "Keep Forwarding: Towards k-link failure resilient routing," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 1617–1625.

[17] K.-T. Foerster, Y.-A. Pignolet, S. Schmid, and G. Tredan, "Local Fast Failover Routing With Low Stretch," *SIGCOMM Comput. Commun. Rev.*, vol. 48, no. 1, pp. 35–41, Apr. 2018. [Online]. Available: http://doi.acm.org/10.1145/3211852.3211858

[18] J. Wang and S. Nelakuditi, "IP Fast Reroute with Failure Inferencing," in *Proceedings of the 2007 SIGCOMM Workshop on Internet Network Management*, ser. INM '07. New York, NY, USA: ACM, 2007, pp. 268–273. [Online]. Available: http://doi.acm.org/10.1145/1321753.1321764

[19] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne, "Multiple Routing Configurations for Fast IP Network Recovery," *IEEE/ACM Transactions on Networking*, vol. 17, no. 2, pp. 473–486, April 2009.

[20] D. Imahama, Y. Fukushima, and T. Yokohira, "A reroute method using multiple routing configurations for fast ip network recovery," in *2013 19th Asia-Pacific Conference on Communications (APCC)*, Aug 2013, pp. 433–438.

[21] B. Zhang, J. Wu, and J. Bi, "RPFP: IP Fast ReRoute with Providing Complete Protection and without Using Tunnels," in *2013 IEEE/ACM 21st International Symposium on Quality of Service (IWQoS)*, June 2013, pp. 1–10.

[22] T. Elhourani, A. Gopalan, and S. Ramasubramanian, "Ip fast rerouting for multi-link failures," *IEEE/ACM Transactions on Networking*, vol. 24, no. 5, pp. 3014–3025, October 2016.

[23] P. Thorat, S. Jeon, S. M. Raza, and H. Choo, "Pre-provisioning of local protection for handling dual-failures in openflow-based networks," in *2017 13th International Conference on Network and Service Management (CNSM)*, Nov 2017, pp. 1–6.

[24] T. Holterbach, S. Vissicchio, A. Dainotti, and L. Vanbever, "SWIFT: Predictive Fast Reroute," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '17.

New York, NY, USA: ACM, 2017, pp. 460–473. [Online]. Available: http://doi.acm.org/10.1145/3098822.3098856
[25] S. Bryant, S. Previdi, and M. Shand, "A Framework for IP and MPLS Fast Reroute Using Not-Via Addresses," Internet Requests for Comments, RFC Editor, RFC 6981, August 2013. [Online]. Available: https://tools.ietf.org/html/rfc6981
[26] A. Atlas, C. Bowers, and G. Enyedi, "An Architecture for IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)," Internet Requests for Comments, RFC Editor, RFC 7812, June 2016. [Online]. Available: https://tools.ietf.org/html/rfc7812
[27] S. Bryant, C. Filsfils, S. Previdi, and M. Shand, "IP Fast Reroute using tunnels," Network Working Group, Internet-Draft, IETF, Internet-Draft, November 2007. [Online]. Available: https://tools.ietf.org/pdf/draft-bryant-ipfrr-tunnels-03.pdf
[28] M. Xu, Q. Li, L. Pan, Q. Li, and D. Wang, "Minimum protection cost tree: A tunnel-based ip fast reroute scheme," *Computer Communications*, vol. 35, no. 17, pp. 2082 – 2092, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366412002137
[29] M. Medard, S. G. Finn, R. A. Barry, and R. G. Gallager, "Redundant trees for preplanned recovery in arbitrary vertex-redundant or edge-redundant graphs," *IEEE/ACM Transactions on Networking*, vol. 7, no. 5, pp. 641–652, Oct 1999.
[30] G. Enyedi, P. Szilágyi, G. Rétvári, and A. Császár, "IP Fast ReRoute: Lightweight Not-Via without Additional Addresses," in *INFOCOM 2009, IEEE*, April 2009, pp. 2771–2775.
[31] A. Gopalan and S. Ramasubramanian, "IP Fast Rerouting and Disjoint Multipath Routing With Three Edge-Independent Spanning Trees," *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, pp. 1–14, 2015.
[32] C. Filsfils, P. Francois, M. Shand, B. Decraene, J. Uttaro, N. Leymann, and M. Horneffer, "Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks," Internet Requests for Comments, RFC Editor, RFC 6571, June 2012. [Online]. Available: https://tools.ietf.org/html/rfc6571
[33] D. Hock, M. Hartmann, C. Schwartz, and M. Menth, "Effectiveness of link cost optimization for ip rerouting and ip fast reroute," in *Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance*, B. Müller-Clostermann, K. Echtle, and E. P. Rathgeb, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 78–90.
[34] G. Rétvári, J. Tapolcai, G. Enyedi, and A. Császár, "IP fast ReRoute: Loop Free Alternates revisited," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 2948–2956.
[35] G. Rétvári, L. Csikor, J. Tapolcai, G. Enyedi, and A. Császár, "Optimizing igp link costs for improving ip-level resilience," in *2011 8th International Workshop on the Design of Reliable Communication Networks (DRCN)*, Oct 2011, pp. 62–69.
[36] L. Csikor, M. Nagy, and G. Rétvári, "Network optimization techniques for improving fast ip-level resilience with loop-free alternates," *Infocommunications Journal*, vol. 3, no. 4, pp. 2–10, December 2011. [Online]. Available: http://eprints.gla.ac.uk/131074/
[37] L. Csikor and G. Rétvári, "Ip fast reroute with remote loop-free alternates: The unit link cost case," in *2012 IV International Congress on Ultra Modern Telecommunications and Control Systems*, Oct 2012, pp. 663–669.
[38] L. Csikor, J. Tapolcai, and G. Rétvári, "Optimizing igp link costs for improving ip-level resilience with loop-free alternates," *Computer Communications*, vol. 36, no. 6, pp. 645 – 655, 2013, reliable Network-based Services. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366412003167
[39] M. Nagy, J. Tapolcai, and G. Rétvári, "Optimization methods for improving ip-level fast protection for local shared risk groups with loop-free alternates," *Telecommunication Systems*, vol. 56, no. 1, pp. 103–119, May 2014. [Online]. Available: https://doi.org/10.1007/s11235-013-9822-y
[40] ——, "Node Virtualization for IP Level Resilience," *IEEE/ACM Transactions on Networking*, vol. 26, no. 3, pp. 1250–1263, June 2018.
[41] T. Elhourani, A. Gopalan, and S. Ramasubramanian, "IP Fast Rerouting for Multi-Link Failures," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 2148–2156.
[42] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. N. Chuah, Y. Ganjali, and C. Diot, "Characterization of Failures in an Operational IP Backbone Network," *IEEE/ACM Transactions on Networking*, vol. 16, no. 4, pp. 749–762, Aug 2008.
[43] S. Nelakuditi, S. Lee, Y. Yu, and Z.-L. Zhang, "Failure insensitive routing for ensuring service availability," in *Quality of Service — IWQoS 2003*, K. Jeffay, I. Stoica, and K. Wehrle, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 287–304.
[44] S. Lee, Y. Yu, S. Nelakuditi, Z.-L. Zhang, and C.-N. Chuah, "Proactive vs reactive approaches to failure resilient routing," in *IEEE INFOCOM 2004*, vol. 1, March 2004, p. 186.
[45] Z. Zhong, S. Nelakuditi, Y. Yu, S. Lee, J. Wang, and C.-N. Chuah, "Failure Inferencing Based Fast Rerouting for Handling Transient Link and Node Failures," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 4, March 2005, pp. 2859–2863 vol. 4.
[46] K. Xi and H. Chao, "IP Fast Rerouting for Single-Link/Node Failure Recovery," in *Fourth International Conference on Broadband Communications, Networks and Systems, 2007. BROADNETS 2007.*, Sept 2007, pp. 142–151.
[47] A. Li, P. Francois, and X. Yang, "On Improving the Efficiency and Manageability of NotVia," in *Proceedings of the 2007 ACM CoNEXT Conference*, ser. CoNEXT '07. New York, NY, USA: ACM, 2007, pp. 26:1–26:12. [Online]. Available: http://doi.acm.org/10.1145/1364654.1364688
[48] S. Kamamura, T. Miyamura, C. Pelsser, I. Inoue, and K. Shiomoto, "Scalable backup configurations creation for ip fast reroute," in *2009 7th International Workshop on Design of Reliable Communication Networks*, Oct 2009, pp. 312–318.
[49] Y. Li and M. G. Gouda, "Ip fast reroute in networks with shared risk links," in *NETWORKING 2010*, M. Crovella, L. M. Feeney, D. Rubenstein, and S. V. Raghavan, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 213–226.
[50] K. W. Kwong, L. Gao, R. Guerin, and Z. L. Zhang, "On the feasibility and efficacy of protection routing in ip networks," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, pp. 1543–1556, Oct 2011.
[51] A. Csaszar, G. Enyedi, J. Tantsura, S. Kini, J. Sucec, and S. Das, "IP Fast Re-Route with Fast Notification," Network Working Group, Internet-Draft, IETF, Internet-Draft, February 2013. [Online]. Available: https://tools.ietf.org/html/draft-csaszar-rtgwg-ipfrr-fn-01
[52] S. Antonakopoulos, Y. Bejerano, and P. Koppol, "Full Protection Made Easy: The DisPath IP Fast Reroute Scheme," *IEEE/ACM Transactions on Networking*, vol. 23, no. 4, pp. 1229–1242, Aug 2015.
[53] S. Kini, S. Ramasubramanian, A. Kvalbein, and A. Hansen, "Fast Recovery From Dual-Link or Single-Node Failures in IP Networks Using Tunneling," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1988–1999, Dec 2010.
[54] ——, "Fast Recovery from Dual Link Failures in IP Networks," in *INFOCOM 2009, IEEE*, April 2009, pp. 1368–1376.
[55] K. Lakshminarayanan, M. Caesar, M. Rangan, T. Anderson, S. Shenker, and I. Stoica, "Achieving Convergence-free Routing Using Failure-carrying Packets," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '07. New York, NY, USA: ACM, 2007, pp. 241–252. [Online]. Available: http://doi.acm.org/10.1145/1282380.1282408
[56] S. S. Lor, R. Landa, and M. Rio, "Packet Re-cycling: Eliminating Packet Losses Due to Network Failures," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, ser. Hotnets-IX. New York, NY, USA: ACM, 2010, pp. 2:1–2:6. [Online]. Available: http://doi.acm.org/10.1145/1868447.1868449
[57] G. Robertson and S. Nelakuditi, "Handling multiple failures in ip networks through localized on-demand link state routing," *IEEE Transactions on Network and Service Management*, vol. 9, no. 3, pp. 293–305, September 2012.
[58] J. Liu, A. Panda, A. Singla, B. Godfrey, M. Schapira, and S. Shenker, "Ensuring Connectivity via Data Plane Mechanisms," in *NSDI*, 2013, pp. 113–126.
[59] A. J. Gonzalez, G. Nencioni, B. E. Helvik, and A. Kamisiński, "A Fault-Tolerant and Consistent SDN Controller," in *2016 IEEE Global Communications Conference (GLOBECOM)*, Dec 2016, pp. 1–6.
[60] J. Kempf, E. Bellagamba, A. Kern, D. Jocha, A. Takacs, and P. Skldstrm, "Scalable fault management for openflow," in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 6606–6610.
[61] S. Sharma, D. Staessens, D. Colle, M. Pickavet, and P. Demeester, "Openflow: Meeting carrier-grade recovery requirements," *Computer Communications*, vol. 36, no. 6, pp. 656 – 665, 2013, reliable

Network-based Services. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366412003349

[62] N. L. M. v. Adrichem, B. J. v. Asten, and F. A. Kuipers, "Fast recovery in software-defined networks," in *2014 Third European Workshop on Software Defined Networks*, Sept 2014, pp. 61–66.

[63] M. Markovitch and S. Schmid, "Shear: A highly available and flexible network architecture marrying distributed and logically centralized control planes," in *2015 IEEE 23rd International Conference on Network Protocols (ICNP)*, Nov 2015, pp. 78–89.

[64] J. Rak, D. Hutchison, E. Calle, T. Gomes, M. Gunkel, P. Smith, J. Tapolcai, S. Verbrugge, and L. Wosinska, "RECODIS: Resilient Communication Services Protecting End-user Applications from Disaster-based Failures," in *2016 18th International Conference on Transparent Optical Networks (ICTON)*, July 2016, pp. 1–4.