A Survey of Fast Recovery Mechanisms in the Data Plane

Marco Chiesa¹ Andrzej Kamisiński^{2,*} Jacek Rak³ Gábor Rétvári⁴ Stefan Schmid⁵

¹ KTH Royal Institute of Technology, Sweden ² AGH University of Science and Technology, Kraków, Poland

³ Gdańsk University of Technology, Poland ⁴ MTA-BME Information Systems Research Group, Hungary

⁵ Faculty of Computer Science, University of Vienna, Austria

Abstract—In order to meet their stringent dependability requirements, most modern communication networks support fast-recovery mechanisms in the data plane. While reactions to failures in the data plane can be significantly faster compared to control plane mechanisms, implementing fast recovery in the data plane is challenging, and has recently received much attention in the literature. This survey presents a systematic, tutorial-like overview of packet-based fast-recovery mechanisms in the data plane, focusing on *concepts* but structured around different networking technologies, from traditional link-layer and IP-based mechanisms, over BGP and MPLS to emerging software-defined networks and programmable data planes. We examine the evolution of fast-recovery standards and mechanisms over time, and identify and discuss the fundamental principles and algorithms underlying different mechanisms. We then present a taxonomy of the state of the art and compile open research questions.

Index Terms—Fast Reroute, Network Resilience, Data Plane

I. INTRODUCTION

Communication networks (datacenter networks, enterprise networks, the Internet, etc.) have become a critical backbone of our digital society. Today, many applications, e.g., related to health, business, science, or social networking, require always-on network connectivity and hence critically depend on the availability and performance of the communication infrastructure. Over the last years, several network issues were reported that led to major Internet outages in Asia [1], resulted in huge losses in revenues [2], affected thousands of airline passengers [3], or even disrupted the emergency network [4]. Many applications already suffer from small delays: A 2017 Akamai study shows that every 100 millisecond delay in website load time can lead to a significant drop in sales [5], and voice communication has a tolerable delay of less than 150 ms; for games it is often less than 80 ms [6].

In order to meet their stringent dependability requirements, communication networks need to be able to deal with *failures* which are becoming more likely with increasing network scale [7]. Today, link failures are by far the most frequent failures in a network [8], [9], and handling failures is a fundamental task of any routing scheme. Historically, resilience to network failures was implemented in the *control plane*: ensuring connectivity was considered the responsibility of the

control plane while the data plane was responsible for forwarding packets at line-speed. Widely deployed routing schemes like OSPF [10] and IS-IS [11], hence include control plane mechanisms which leverage global message exchanges and computation to determine how to recover from link failures.

However, the slow reaction times of control plane mechanisms is becoming increasingly unacceptable [12]–[14]. Indeed, many studies have shown that control plane-based resilience can noticeably impact performance [7], also because of the high path re-computation time [15]. In basic solutions where the network can recover from failure only after the control plane has computed a new set of paths and installed the associated state in all routers, the disparity in timescales between packet forwarding in the data plane (which can be less than a microsecond) and control plane convergence (which can be as high as hundreds of milliseconds) can lead to long outages [16]. While recent centralized routing solutions based on Software-Defined Networks (SDNs) [17]–[19], where all routing computation is performed by a controller which then pushes the results to the affected routers, are faster, there is still an inevitable delay of at least the round trip time between the routers and the controller.

Motivated by these performance issues, we currently witness a trend to reconsider the traditional separation of concerns between the control plane and the data plane of a network. In particular, given that the data plane typically operates at timescales several orders of magnitude shorter than the control plane [16], [20], moving the responsibility for connectivity to the data plane where failure recovery can in principle occur at the speed of packet forwarding is attractive.

Indeed, most modern networks support different kinds of *fast-reroute (FRR)* mechanisms which leverage *pre-computed* alternative paths at any node towards any destination. When a node locally detects a failed link or port, it can autonomously remove the corresponding entries from the forwarding table and continue using the remaining next hops for forwarding packets: a fast local reaction [21]. In FRR, the control plane is hence just responsible for *pre-computing* the failover paths; when a failure occurs, the data plane utilizes this additional state to forward packets. For example, many data centers use ECMP [22] (a data plane algorithm that provides automatic failover to another shortest path), WAN networks leverage *IP Fast Reroute* [23]–[25] or *MPLS Fast Reroute* [26] to deal with failures on the data plane, SDNs provide FRR functionality in terms of *OpenFlow fast-failover groups* [27], and BGP relies

* Corresponding author: Andrzej Kamisiński, andrzejk@agh.edu.pl

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

on BGP-PIC [28] for quickly rerouting flows, to just name a few.

Implementing FRR mechanisms, however, is challenging and requires careful configuration, as the failover behavior needs to be *pre-defined*, before the actual failures are known. Additional challenges are introduced due to the limited functionality in the data plane as well as the stringent latency requirements which do not allow for sophisticated reactions. Configuring FRR is particularly tricky under multiple and correlated failures [29]–[33]. FRR mechanisms in the data plane are hence often seen as a “first line of defense” and lack guarantees: they support a fast but perhaps suboptimal reaction (due to the limited information about the actual failure scenario). In a second phase, the control plane may re-establish the routes more effectively, meeting advanced traffic engineering criteria.

A. Our Contributions

This survey provides an overview of fast-recovery mechanisms with a focus on *concepts*: our primary goal is to familiarize the reader with selected concepts, in a tutorial-like manner, which together form a representative set of fast-recovery methods in the data plane. Indeed, as we will see, different approaches have different advantages and disadvantages, and a conceptual understanding is hence useful when reasoning about which technology to deploy. The topic is timely, not only because of the increasing dependability requirements on communication networks and the resulting trend to move the responsibility for connectivity to data plane mechanisms, but also due to the emergence of programmable data planes which introduce new functionality in the data plane, allowing to implement different approaches. We believe that this provides a need and opportunity for a structured survey.

To be concrete, we structure our discussion around the underlying technologies (e.g., Ethernet, MPLS, IP, SDN) as well as the related use cases (e.g., failure scenarios, inter-domain routing, intra-domain routing, data centers) and technological constraints. We then highlight algorithmic aspects (e.g., complexity) and performance issues (e.g., the total time needed to complete the recovery process). Throughout the paper, we provide clear explanations of the related technical terms and operation principles and use illustrations based on simple examples.

In this paper, we focus on data plane mechanisms, and in particular, packet-based fast-recovery mechanisms. At the same time, we do not discuss control plane mechanisms nor the orthogonal issue of how to *detect* failures in the first place. We concentrate on the most common case, the case of simple best-effort unicast packet communication, and we barely touch on issues related to multicast, QoS, and traffic engineering. Furthermore, given the focus on concepts and the tutorial-style nature of this article, we do not aim to provide a complete survey of the research literature, which is beyond the scope of this paper.

This paper is the first one to investigate fast-recovery mechanisms in the data plane from the perspective of a comprehensive set of available technologies mentioned above as well as a broad range of characteristics, including:

- the underlying general concept (e.g., loop-free alternate next-hop, input interface-aware routing, additional/extended FIBs and other data structures, encapsulation/tunneling, redundant spanning trees, failure coverage analysis and improvement);
- improvements of other solutions including the purpose of changes and the illustration of the evolution of ideas;
- the selected algorithmic aspects, performance;
- the maximum number of simultaneous failures handled by the scheme;
- mode of operation (i.e., distributed, centralized, hybrid);
- signaling aspects (whether the use of packet header/dedicated control messages for signaling is needed, or no signaling is required);
- reliance on existing routing protocols;
- technological constraints;
- deployment considerations including the required modifications in the data/control plane, gradual deployment, possible difficulties and estimated cost;
- the use cases presented by the proposers of schemes and the related assumptions/requirements (time scale, convergence, or investigated failure models).

B. Novelty and Target Audience

There already exist many good surveys on reliable routing in specific communication technologies, such as Ethernet [34], IP [23], [25], [35], [36] (and more recently [23], [25]), MPLS [37], [38], BGP [39], or SDN [40]. However, to the best of our knowledge, this is the first complete and up-to-date tutorial-like survey on the timely topic of fast-recovery mechanisms in the data plane. We believe that only such an overarching approach can highlight the conceptual similarities and differences, and hence help choose the best approach for a specific context.

Indeed, our goal is *not* to provide a comprehensive survey of the literature. Rather, in order to provide an understanding of the underlying key concepts, our paper focuses on the fundamental mechanisms across a spectrum of different technologies and layers.

Our paper hence targets students, researchers, experts, and decision-makers in the networking industry as well as interested laymen.

C. Organization

The organization of this paper adopts the traditional protocol-layer-based view of communication networks. As mentioned above, the perspective in this paper is motivated by emerging technologies such as SDNs and programmable data planes which bring together and unify specific layers and technologies. For this, however, we first need an understanding of the technologies on the individual layers. Therefore, after Section II, which provides a through overview of the fundamental concepts related to network resilience, in each section we review the most important ideas, methods, and standards related to the different layers in the Internet protocol stack in a bottom-up order (see Fig. 1). Whenever a protocol layer spans multiple operational domains, the organization follows

Fig. 1: Organization of the paper.

the traditional edge/access/core separation principle. Inside each section then, the review takes a strict chronological order, encompassing more than 25 years of developments in data-plane fast-recovery methods.

Section III commences with the overview of fast recovery in the link layer, mostly used in local and metro area networks. Then, Section IV reviews fast recovery for MPLS, operating as the “Layer 2.5” in the Internet protocol stack, deployed in access networks and, increasingly, in provider core networks. Section V presents IP Fast ReRoute, the umbrella network-layer fast-recovery framework designed for a single operator domain (or Autonomous System), and then Section VI discusses the selected fast-recovery solutions designed for wide-area networks and, in particular, the Internet core. The separation emphasizes the fundamental differences between the cases when we are in full control of the network (the intra-domain case, Section V) and when we are not (Section VI). Next, we review the methods that do not fit into the traditional layer-based view: Section VII discusses the fast recovery mechanisms proposed for the emerging Software-Defined Networking paradigm, while Section VIII summarizes generic, technology-agnostic solutions.

Finally, we cast the existing fast-recovery proposals in a common framework and discuss some crucial related concepts. In particular, in Section IX we give a comprehensive family of taxonomies to classify fast-recovery schemes along several dimensions, including data-plane requirements, operation mode, and resiliency guarantees. In Section X we briefly discuss some critical issues regarding the interaction of the control-plane and the data-plane schemes during and after a recovery. Last, in Section XI we conclude the paper and identify the most compelling open issues for future research.

II. FUNDAMENTAL CONCEPTS

The undoubtedly large amount of data transmitted by communication networks makes the need to assure their fault-tolerance one of the fundamental design issues. This is justified by a variety of reasons for single and multiple failures of network nodes and links. In particular, failures of physical links, being a dominant scenario of outages in wide-area networks [41], are mostly a result of random activities such as

a cable cut by a digger. According to [42], [43], there is one failure every 12 months related to each 450 km of links, while the average repair time of a failed element is 12 hours. Other events show that the frequency of multiple failures – especially those occurring in a correlated manner, e.g., as a result of malicious human activities, or forces of Nature (tornadoes, hurricanes, earthquakes, etc.) is raising [44].

Conventional routing protocols used in IP networks such as Border Gateway Protocol (BGP) [45] or Open Shortest Path First (OSPF) [10] are characterized by a slow process of a post-failure routing convergence which can even take tens of seconds [42], [46] acceptable only for delay-tolerant applications (provided that the failures are indeed not frequent).

Concerning the set of four intrinsic parameters (i.e., related to network performance) of Quality of Service (QoS) [47]–[49] defined for IP networks as: the maximum IP Packet Transfer Delay (IPTD), IP Delay Variation (IPDV), IP packet Loss Ratio (IPLR) and IP packet Error Ratio (IPER) – see ITU-T recommendations Y.1540 and Y.1541 in [50], [51], such a slow convergence is unacceptable for a wide range of applications with stringent QoS requirements presented in [49] and summarized in Table I. As shown in Table I, for applications belonging to the first four Classes of Service (CoS), i.e., classes 0–3, the considered values of transmission delay (undoubtedly impacted after a failure) should not be higher than 100–400 ms. Similarly, the values of delay variation for classes 0–1 should be at most 50 ms, which during the network recovery phase is even around two-three orders of magnitude less than what the conventional IP routing protocols can offer.

It is important to note that the majority of failure events in IP networks is transient [46], [52], [53] and, therefore, may lead to a lack of routing convergence during a significant time. Mechanisms of fast recovery of flows are thus crucial to assure **network resilience**, i.e., the ability to provide and maintain an acceptable level of service in the presence of various faults and challenges to normal operation [54]–[56].

To explain the fundamental concepts of network resilience to assure fast recovery of the affected flows, in Section II-A we first present the set of relevant disciplines of network resilience followed by definitions of resilience measures. In particular, the set of characteristics analyzed in Section II-A is essential for evaluation of **fast recovery** which we define in this paper as *the ability of a network to recover from failures according to the time-constrained recovery plan to meet the QoS-related requirements of applications*. Next, Section II-B provides a detailed taxonomy of network resilience mechanisms based on utilization of the alternate paths with a particular focus on service recovery time. Finally, in Section II-C, information on the steps of the network recovery procedure is given.

A. Disciplines and Measures

As discussed in [54], resilience disciplines can be broadly divided into two categories: *challenge tolerance* comprising the design approaches to assure the continuity of service and *trustworthiness* focusing on the measurable characteristics of network resilience shown in Fig. 2.

Following [54], challenge tolerance includes disciplines addressing the network design issues to provide service conti-

TABLE I: Classes of Service defined by ITU-T

Class of Service	Description of applications	IPTD	IPDV	IPLR	IPER
Class 0	Real-time, jitter-sensitive, highly interactive (e.g., VoIP, video teleconference)	100 ms	50 ms	1×10^{-3}	1×10^{-4}
Class 1	Real-time, jitter-sensitive, interactive (e.g., VoIP, video teleconference)	400 ms	50 ms	1×10^{-3}	1×10^{-4}
Class 2	Transaction data, highly interactive (e.g., signaling)	100 ms	undefined	1×10^{-3}	1×10^{-4}
Class 3	Transaction data, interactive	400 ms	undefined	1×10^{-3}	1×10^{-4}
Class 4	Tolerating low loss (e.g., short transactions, bulk data, video streaming)	1 s	undefined	1×10^{-3}	1×10^{-4}
Class 5	Typical applications of IP networks	undefined	undefined	undefined	undefined

Fig. 2: Classification of network resilience disciplines based on [54].

nity in the presence of challenges. Among them, *survivability* denotes the capability of a system to fulfil its mission in the presence of threats including natural disasters and attacks. *Traffic tolerance* refers to the ability to tolerate the unpredictable traffic, which may be a remarkable challenge in a multiple-failure scenario implied, e.g., by a disaster (such as a tsunami after an earthquake), or in other situations including, e.g., DoS attacks where the traffic volume is raised unexpectedly far over the expected peak value for the normal operational state. *Disruption tolerance* focuses on the aspects of connectivity among network components primarily in the context of nodes mobility, energy/power challenges, or weak/episodic channel connectivity [57].

Trustworthiness, in turn, comprises the measurable characteristics of the resilience of communication systems to evaluate the assurance that the system will perform as expected [58]. It includes three disciplines, namely: dependability, security and performability. *Dependability* is meant to quantify the level of service reliance. It includes:

- *reliability* $R(t)$ being a measure of service continuity (i.e., probability that a system remains operable in a given $(0, t)$ time period,
- *availability* $A(t)$ defined as the probability that a system is operable at time t .

Its particular version is the *steady-state availability* defined as a fraction of a system lifetime during which the system is accessible, as given in Eq. 1.

$$A = \frac{MTTF}{MTTF + MTTR} \quad (1)$$

where:

$MTTF$ denotes the mean time to failure,

$MTTR$ is the mean time to repair.

- *maintainability* being the predisposition of a system to updates/evolution,
- *safety* being a measure of system dependability under failures,
- *integrity* denoting protection against improper alterations of a system.

Security denotes the ability of a system to protect itself from unauthorized activities. It is characterized by both joint properties with dependability (i.e., by availability, and integrity) as well as individual features including authorisability, auditability, confidentiality, and nonrepudiability [59]. *Performability* provides measures of system performance concerning the Quality of Service requirements defined in terms of transmission delay, delay variation (jitter), throughput/goodput, and packet delivery ratio [54].

As the impact of resilience on the Quality of Service has been recognized as evident, a concept of *Quality of Resilience* (*QoR*) has been introduced in [60] to refer to service resilience characteristics. In contrast to QoS characteristics being rather short-term, the majority of attributes of resilience relating to service continuity, downtime, or availability are long-term by nature [60]. Indeed, most of resilience metrics proposed by International Telecommunication Union–Telecommunication Standardization Sector (ITU-T) summarized in [60] and shown in Table II can only be measured in the long term based on end-to-end evaluations. It is also worth noting that, unlike many QoS measures, QoR characteristics cannot be perceived by users in a direct way, for whom it is not possible to distinguish between a network element failure and network congestion when noticing the increased transmission delay/packet losses.

In spite of the existence of a number of resilience measures, in practice only two of them are widely used, i.e., the mean time to recovery (MTTR) and the steady-state availability defined by Eq. 1, which is also impacted by the recovery time (MTTR factor) [60]. Therefore, in the following part of this section presenting the taxonomy of recovery schemes and recovery procedure, a particular focus in these parts is on the recovery time issues.

B. Taxonomy of Recovery Methods

The need for fast rerouting has its roots in an undoubtedly slow process of post-failure routing convergence of contemporary schemes such as BGP or OSPF which can even take tens

TABLE II: Selected ITU-T Resilience Metrics

ID	Area	Metric
E.800	General (e.g., Internet access)	Instantaneous (un)availability – probability for a network element of being in an up (down) state at a given time
E.802		Mean time between failures (MTBF) – mean time between two consecutive failures of a repaired element
E.820		Mean time between interruptions (MTBI) – mean time between the end of one interruption and start of the next one
E.850		Mean time to failure (MTTF) – mean value of time since the last change of state from down to up until the next failure
E.855		Mean time to recovery (MTTR) – mean value of time when a network element is in down state due to a failure
E.860		Mean up time / mean down time (MUT/MDT) – mean value of time when a network element is in up/down state
E.862		Reliability function $R(t)$ – probability for a network element of being in up state in $(0, t)$ interval
E.880		Retainability – probability that a service will continue to be provided
Y.1540	IP	IP packet loss ratio (IPLR) – the total number of lost packets divided by the total number of transmitted packets
Y.1541		Service availability – a share of the total service time classified as available using the threshold on IPLR
Y.1542		IP service (un)availability (PIU/PIA) – part of time of (un)available IP service based on IP service (un)availability function
Y.1561	MPLS	Packet loss ratio (PLR) – similar to IPLR
		Severe loss block (SLB) – an event at the ingress node for a block of packets with packet loss ratio above the upper bound
		Recovery time – time for recovery operations based on the number of successive time intervals of SLB outcomes
		Service availability, PIU, PIA – defined in a similar way as in Y.1540 but in the context of SLBs
Y.1562	Higher layer protocols	Service availability – similar to Y.1540 but related to the transfer delay and service success ratio

of seconds. The key objective is to reduce the convergence time to the level of less than several tens of milliseconds [35]. In this context, IP network resilience is often associated with the path-oriented Multiprotocol Label Switching (MPLS)-based recovery schemes. IP-MPLS Recovery mechanisms aim to redirect the traffic served by the affected working paths onto the respective alternate (backup) routes [60].

A common observation is that different traffic flows are characterized by differentiated resilience requirements. Therefore, to prevent the excessive use of network resources, it is reasonable to decide on the application of specific recovery schemes on a per-flow basis [61]. In particular, this would mean that only those flows requiring a certain level of service availability in a post-failure period need to be provided with a recovery mechanism.

Following [56], [62], recovery methods can be classified based on several criteria, the most important ones including the backup path set up method, the scope of recovery procedure, the usage of recovery resources, the domain of recovery operations, or the layer of recovery operations shown in Fig. 3.

Concerning the backup path set up method, the alternate paths can be either configured in advance (pre-computed) at the time of setting up the primary path (known as *pre-planned (protection)* switching scheme, or established dynamically after detection of a failure (referred to as the *restoration/rerouting* concept) [61], [63]. Preplanned protection provides faster recovery of the affected flows as alternate paths are pre-established before the failure. However, a disadvantage is its high cost due to the increased level of resource consumption for backup paths set up well before the failure and often used only for a relatively short post-failure period. Restoration methods in practice are considered as a default solution for contemporary IP networks [60]. They are remarkably more capacity-efficient but do not provide 100% of restorability, as the appropriate network resources are not committed before a failure and may not be available for the alternate paths after a failure. Another disadvantage of the restoration schemes is

Fig. 3: Classification of recovery methods.

their lower time-efficiency, as recovery procedure in their case also involves the phase of a backup path calculation.

Protection switching approaches based on preplanned protection are reasonable for IP flows requiring service recovery time below 100 ms [61]. For flows with restoration time between 100 ms and 1 s, it is reasonable to apply a restoration scheme such as MPLS restoration [64]. In the context of the other flows able to tolerate the recovery switching time over 1 s, a conventional Layer 3 rerouting is commonly used. Flows with no resilience requirements are usually not recovered, and their resources are freed (i.e., preempted) to enable recovery of other flows with resilience requirements [61].

Concerning the scope of a backup path [63], [65], recovery

schemes can be divided into:

- *global* schemes where a single backup path protects the entire working path (Fig. 4a). Global schemes are commonly associated with reservation of resources before a failure (i.e., resources pre-reserved and alternate paths pre-established). Such backup paths can be used either after failure only (the so-called 1:1 path protection) or in parallel with the working path to carry the traffic in a normal state (1+1 protection model). Under path protection, switching the traffic onto a backup path is done at the ingress-egress pair of nodes (i.e., the end nodes of a path) as shown in Fig. 4a,
- *local* schemes (with backup paths being either set up in advance, or dynamically after a failure) enabling short detours over the failed link of the affected path (Fig. 4b) or over two consecutive links in the case of a failure of a node (Fig. 4c),
- *segment* schemes with backup paths protecting certain segments of backup paths (Fig. 4d).

Fig. 4: Scope of the backup path: (a) path protection, (b) local protection against a link failure, (c) local protection against a node failure, (d) segment protection.

A general observation is that with a decrease of the scope of backup path protection (i.e., from global toward local protection schemes), the time of service restoration decreases. This can be explained by the fact that for local protection schemes the redirection of the affected flows onto the backup paths is done closer to the failure as well as because backup paths under local protection are remarkably shorter than the respective ones for global protection. Therefore, recovery schemes based on local detours (especially local protection methods involving pre-reservation of backup path resources) are often referred to as *Fast Reroute* concepts in the literature [65].

However, fast switching of flows onto backup paths for local protection schemes is achieved at the increased ratio of network redundancy (denoting the capacity needed to establish backup paths) when compared to global protection schemes. This is due to the fact that total length of backup paths protecting a given primary path for local protection is greater than the length of a backup path providing the end-to-end protection in the global protection scheme. Local recovery

schemes are therefore seen as fast but often more expensive (in the case of local protection) than the respective global schemes, which are, in turn, more capacity-efficient, easier to optimize, but slower concerning recovery time [60].

Network resources (capacity of links) reserved for backup paths can be either dedicated (i.e., reserved for those backup paths exclusively), or *shared* among a set of backup paths. To provide the network resources for backup paths after a failure, sharing the link capacity by a set of backup paths is possible if these backup paths protect mutually disjoint parts of working paths (i.e., to guarantee that after a failure, there would be a need to activate at most one of these backup paths) [56]. Using dedicated backup paths may result in faster recovery (as dedicated backup paths can be fully pre-configured before a failure), but is undoubtedly expensive, as backup paths often require even more network resources than the corresponding working paths (by default, they are longer than the parts of working paths they protect). Another observation is that a classification of backup resource reservation schemes into dedicated and shared is characteristic strictly to protection methods [42], while reactive recovery schemes often involve reservation of dedicated resources for the alternate paths (as these paths are merely the only ones operating after a failure).

For networks divided into multiple domains (each domain often managed by a different owner), recovery actions are performed in these domains separately. In single-domain networks, one recovery scheme can be, in turn, applied in an end-to-end manner for all flows.

Architectures of communication networks are inherently multilayer meaning that one transmission technology such as Optical Transport Network (OTN) using optical (Wavelength Division Multiplexing – WDM) links serves as a carrier for another transfer architecture such as, e.g., IP network [42]. In a layered architecture, an IP network is typically considered as the uppermost layer. Therefore, failures seen by the IP network can happen at several layers and for different reasons making network recovery a challenging issue [46]. In particular, a failure of an optical link or of an optical equipment in the WDM layer serving many higher-layer paths (in the IP layer), if not handled on time by the WDM layer may manifest itself as a simultaneous failure of a number of IP links and thus initiate numerous recovery actions in the IP layer [61].

To mitigate this issue, schemes of coordinated recovery were proposed with the sequence of recovery operations being either bottom-up, top-down, or integrated [66]. Concerning the reduction of the number of recovery operations and a decrease in the recovery time, the bottom-up scheme seems to be the best solution. In this case, recovery operations are performed first at the coarsest granularity in the lower layer (which is usually fast [60]), e.g., for failures of an optical link O_3-O_5 or an optical node O_6 in Fig. 5. The recovery procedure is next continued in the upper (IP) layer only concerning those flows which could not be restored at the lower layer. For instance, in the case of a failure of IP router R_B in Fig. 5, the only possibility is to restore the affected flows in the IP layer.

It is worth noting that recovery operations in the IP layer although being slower, are less expensive (as they are done

Fig. 5: An example scenario of a multilayer recovery. Failures of an optical link O_3 – O_5 or an optical node O_6 can be handled either in the lower (WDM), or in the upper (IP) layer. However, WDM recovery is faster, and, thus preferred. For a failure of an IP router R_B , the only possibility is to proceed with the recovery at the IP layer, as a failure of an IP router, in fact, is equivalent to a failure of the end node of WDM paths.

at finer granularity) than in the optical layer [60]. Therefore, as requirements of flows concerning resilience are indeed differentiated (and there are not many flows requiring high resilience), it may be more resource-efficient to proceed with recovery of these flows at the IP layer. Indeed, the number of such flows may also be not high enough to justify from the cost point of view the application of a recovery procedure at the lower layer operating on the aggregate flows.

Fig. 6 presents relations among the recovery methods for IP networks discussed in this section concerning the time of recovery.

	faster ←.....→ slower		
Backup path setup method	Preplanned (resources pre-reserved)		Reactive (restoration / rerouting)
	faster ←.....→ slower		
Scope of recovery procedure	Local	Segment	Global
	faster ←.....→ slower		
Use of recovery resources	Dedicated		Shared

Fig. 6: Relations among the recovery methods for IP networks concerning the recovery time.

C. Recovery Procedure

As explained in [60], [62], [64], [65], fast recovery of IP flows should comprise the following phases: fault detection, fault localization, hold-off period, fault notification and recovery switching shown in Fig. 7.

The objective of *fault detection* phase is to notice the failure in the network (time T_1 in Fig. 7). a failure can be detected either by a management or a transport (data) plane [42], [67],

Fig. 7: Restoration time components based on [62], [64], [65].

e.g., at the level of optical paths forming the IP virtual links or in the IP network.

Concerning the management plane, failures can be identified by network elements close to the faulty item using the Loss of Clock, Loss of Modulation, Loss of Signal, or degradation of signal quality (e.g., increased signal-to-noise ratio – SNR) [62]. For instance, failure detection in optical networks makes use of information on the optical power or the temperature at the transmitter, or the input optical power at the receiver, power distribution of carriers over the full bandwidth/channel wavelength, or crosstalks [68]. Other hardware components which can generate alarms in optical networks include optical regenerators/reshapers/retimers – 3Rs (e.g., when it is not possible for them to lock to the incoming signal) or switches when they cannot establish a new connection [69].

In the data plane, a fault can be detected by observing a degraded quality in the context of an increased Bit Error Ratio – BER, e.g., by CRC computation (Ethernet), TCP/IP checksum verification [68], and by noticing the increased values of the end-to-end quality parameters such as lower throughput or increased data transmission delay [42]. In multilayer networks, the upper-layer algorithms referring to failure detection in IP and overlay networks can be broadly classified into active and passive schemes [70]. In the active schemes, periodic keep-

alive messages are exchanged between neighbouring nodes. In this case, fast detection of failures comes at the price of an increased amount of control overhead. One of the related examples is the Bidirectional Forwarding Detection (BFD) mechanism [71]. On the other hand, passive schemes only make use of data packet delivery to confirm correct operation of nodes and links (however, if data packets are not sent frequently enough, passive schemes are hardly useful). The status of a given node can then be determined by other nodes either independently, or collaboratively [72].

It is worth noting that, failure detection in multilayer networks is often one of the most redundant tasks, as it is commonly performed at multiple layers [69].

Fault localization (represented by time T_2 in Fig. 7) means identification of the faulty element (point of failure) necessary to determine the network element at which the transmission should be suspended [42], [62]. Precise localization of a faulty element as well as identification of the element type (node/link) is crucial especially in the context of local repair methods, where redirection of the affected flows is performed just around the failed node/link.

In any layered structure, where the IP layer is typically considered as the uppermost one (as in the IP-over-WDM architecture [56], [66]), there is a need to decide on the sequence of layers to perform the recovery operations. In general, recovery operations at the lower (e.g., WDM) layer are executed at coarser granularity (due to the aggregation of IP flows onto optical lightpaths), which reduces the number of recovery operations. Also, if WDM recovery operations are performed fast enough, they can be entirely transparent to the IP layer. Only those failures that cannot be recovered in the WDM layer (e.g., failures of IP routers) need to be handled in the IP layer. Therefore, in the context of the bottom-up sequence of recovery actions, the *hold-off* period (time T_3 in Fig. 7) is used to postpone the recovery operations in the IP layer (e.g., related to failures of the IP layer nodes) until the respective recovery operations are performed first by the lower (WDM) layer [66].

The objective of *fault notification* initiated by a node neighboring to the faulty element is to inform the intermediate nodes along the primary path about a failure and the end nodes of the protected segment of the working path (referred to as ingress and egress nodes in the case of global/segment schemes) to trigger the activation of the backup path. It is important to notice that fault notification time (T_4 in Fig. 7) can be neglected for local repair schemes such as link protection (see e.g., Fig. 4c), where the node detecting the failure is, in fact, the one to redirect the affected flow (i.e., the ingress node of the affected part of a working path).

During *recovery operation (switching)* interval (time T_5 in Fig. 7), reconfiguration of switching at network nodes takes place to redirect the affected flows onto the respective alternate paths. This stage is more time-consuming for restoration schemes than for protection strategies, as it also includes calculation and installation of the alternate path [65].

The recovery procedure is completed after the verification and synchronization phase (given by time T_6 in Fig. 7) when the alternate path is verified and synchronized, as well as after

the traffic is next propagated via the alternate route and reaches the end node of the backup path [42], [65].

After the traffic is switched onto backup paths, the process of manual repair of failed network nodes/links by the personnel is initiated. In practice may take at least several hours [42]. The objective of this *normalization* phase (T_7 in Fig. 7) is to restore the original characteristics of the network from the period before the failure. In particular, normalization also includes the need to update the communication paths, as the recovery paths used after a failure are non-optimal in the physically repaired network (these paths are commonly longer and use more network resources than the respective working paths do before a failure). Therefore, the normalization phase is also to revert the traffic onto the original set of communication paths utilized before the failure.

An important conclusion following from this subsection and the previous one is that fast recovery of the affected flows (especially essential for real-time services with stringent requirements on availability and reliability such as Voice-over-IP [52], [73]) is possible by the application of local protection schemes. These methods involve short backup paths installed before the failure and eliminate the need to send the end-to-end notifications along the working path [65]. Therefore, proactive IP fast reroute schemes investigated in the remaining part of this paper are based on the local protection concept.

III. LINK-LAYER FAST RECOVERY

Resilient routing in Ethernet networks is challenging. On one hand, fast recovery mechanisms are necessary to minimize message losses and increased transmission delay due to one or more failures. However, this is challenging in Ethernet, as frames do not include a Time-to-Live (TTL)-like field known from the network-layer IPv4 protocol, and as even transient forwarding loops can cause significant problems. The first major solution designed to avoid forwarding loops while providing basic restoration capabilities in Ethernet networks was the IEEE 802.1D Spanning Tree Protocol (STP) introduced in [74]. In STP, a single spanning tree is established across the network to assure that there is a unique path between any two nodes. However, despite being simple and easy to scale, it has not been designed to provide fast failure recovery. As mentioned in [75], it is characterized by a remarkably slow convergence time of even up to 50 s [76], which is not acceptable for many applications and gets magnified in networks consisting of hundreds of switches.

Based on the scope of recovery, we can distinguish between *global recovery* schemes to protect any node/link except for the source/destination node of a demand, and *local recovery* schemes [77] to protect against a failure of the incident node/link, minimizing the time necessary for recovery. The global recovery is initiated by the source/destination node, while local recovery is triggered by the immediate upstream node of the failed network element.

Three IEEE Ethernet spanning tree protocols include the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP) [74], [78], and Multiple Spanning Trees Protocol (MSTP) [77], [79]. STP is considered to be the first spanning

tree protocol for Ethernet with resilience functionality, which, upon a failure, triggers the spanning-tree reconstruction procedure. However, links which do not belong to the spanning tree cannot be used to forward traffic, which might lead to increased resource utilization and local link congestions in other areas of the network. Since the introduction of STP, several mechanisms have been proposed to solve this issue (see the related evolution timeline shown in Fig. 8). We discuss the selected representative examples in the following sections. For a discussion of different solutions related to optical networks, the reader is referred to [56], [62], [66].

A. Solutions Based on a Single Spanning Tree

The Rapid Spanning Tree Protocol (RSTP, IEEE 802.1D [74], [78]) was proposed to reduce the negative impact of a long convergence time of a single spanning tree on Ethernet network performance. It operates in a distributed fashion and relies on the proposal-agreement handshaking algorithm to provide synchronization of the state by switches. RSTP also introduces the concept of specific Port Roles related to recovery processes. In particular, as shown in Fig. 9, the *Alternate Port* and *Backup Port* roles are assigned to such ports of a bridge which can be used to provide connectivity in the event of failure of other network components [78].

RSTP not only prevents forwarding loops, but it also enables transmission on the redundant links in the physical network topology. The protocol was evaluated in real test networks [87] in the context of access and metro Ethernet networks by investigating the failure detection time and additional delays introduced by hardware. The reported results confirm that RSTP can converge within milliseconds. At the same time, hardware delays may extend the observed recovery time considerably.

Another mechanism to reduce the reconfiguration time of STP was proposed in [85]. The main idea behind this scheme is to avoid using conventional timeouts to inform about the local completion of a tree formation phase. Instead, the approach uses the additional explicit termination messages which are sent backwards from nodes to their parent nodes in the newly formed tree. Therefore, the scheme is able to converge in a shorter time (based on evaluation results presented in [85], the tree recovery time can be remarkably reduced even to less than 50 ms for a moderate-size network).

In some cases, it is possible to shorten the recovery time after single link failures by constructing the spanning tree in such a way that replacing one failed link of the tree with a different link excluded from that tree will result in another valid spanning tree, as shown in Fig. 10.

To this end, a distributed mechanism was proposed in [88] based on the following three steps:

- *Failure detection*: detecting a failure in the physical layer;
- *Failure propagation*: broadcasting failure information in a “failure information propagation frame” which is assigned the highest priority;
- *Reconfiguration*.

The proposed solution not only avoids constructing and maintaining multiple spanning trees to protect against single link

Fig. 8: Timeline of the selected documents and fast-recovery solutions operating at the link layer (entries marked in gray provide the general context).

Fig. 9: Illustration of the example configuration of a spanning tree for RSTP rooted at node R_5 : An incoming frame is forwarded by R_1 by default in the direction of the root node (R_5 here) via the respective *Root Port* outgoing from R_1 towards R_3 . In the case of a failure of the primary link R_1-R_3 , the frame can be forwarded along a duplicate link via the *Backup Port*. However if both direct links between R_1 and R_3 are not available (e.g., due to their failure), the frame can be forwarded via the *Alternate Port* towards node R_2 .

Fig. 10: Illustration of a self-protected spanning tree (thick black lines) and the new tree returned by the single link switching mechanism after a single link failure (the newly attached link is marked in orange).

failures in two-connected networks, but it also creates new opportunities in terms of load balancing. Besides, it detects failures much faster than STP and RSTP, achieving recovery time values less than 50 ms, while Viking [86] relies on a slower mechanism based on SNMP traps [88].

Another example of this approach, called Fast Spanning Tree Reconnection (FSTR), was described in [91]. This distributed mechanism relies on an Integer Linear Program (ILP). It is executed offline to determine the best set of links that can be attached to the spanning tree to reconnect it after any single link failure. Whenever the failure occurs, switches incident to the failed link send notification messages containing the identifier of the failed link to the preconfigured switches which can activate one of the available *reconnecting* links. The related advantage is that only the directly affected switches and the intermediate switches need to update their tables, while the other devices do not have to. In this case, the recovery time consists mainly of the switch reconfiguration delay. On the other hand, existing Ethernet switches need to be modified to support the proposed solution. In particular, each compatible switch is required to maintain the following tables:

- *Failure notification table*: contains the MAC addresses of the switch interfaces where to send notification messages;
- *Alternate output port table*: includes the output port selected based on the identifier of the failed link (the table is computed offline).

After successful reconfiguration of switches, traffic forwarded via the failed link is instantly redirected, and the backward learning procedure is avoided. The proposed mechanism also considers backup capacity guarantees.

As switches may not be aware of other failed links immediately, an improved distributed mechanism capable of dealing with double-link failures in Metro Ethernet Networks was proposed in [93]. The underlying concept remains similar to their previous works. However, the main difference is that the reconnect links are selected in a way that loops are naturally avoided for double failures (even though each failure is handled independently). The proposed ILP formulation includes an extension to determine the best set of reconnect-links that can reconnect each affected spanning tree, while minimizing the backup capacity reserved in the network and satisfying the preferred protection grade for each connection. The new solution still deals with any single failure successfully. At the same time, only partial protection can be provided for double failures.

The concept of partial spatial protection (PSP) together with a mixed integer linear programming (MILP) was proposed in [96] as an extension of the two schemes ([91] and [93]) described above. In particular, as not all flows require full protection (against a failure of any possible link on a way), the extension described in [96] is to update the FSTR concept in a way to provide protection of flows against failure of a link from a specific subset of links only (to satisfy a given protection grade required by a demand).

B. Solutions Based on Multiple Spanning Trees

The MSTP protocol is based on RSTP and it supports multiple spanning trees [74], [79], [104]. To this end, it partitions the network into multiple *regions* and calculates independent Multiple Spanning Tree Instances (MSTIs) within each of the regions based on the parameters conveyed in Bridge Protocol Data Unit (BPDU) messages which are exchanged between the involved network bridges. The MSTIs are assigned unique Virtual LAN (VLAN) identifiers within the region, so that frames with a given VLAN identifier are forwarded consistently by all bridges within that region. Consistent assignment is achieved based on MST Configuration Identifiers included in BPDUs that are transmitted and received by adjacent bridges in the same region. Such a mechanism is critical for the correct forwarding of frames within the region, as otherwise, some frames might be duplicated or even not delivered to the destination LANs. Note that no LAN can belong to two or more regions simultaneously.

Similarly to RSTP, MSTP defines a set of Port Roles which includes the *Alternate Port* and *Backup Port* roles, assigned to such ports of a bridge that can be used to provide connectivity in the event of failure of other network components, or when other bridges, bridge ports, or LANs are removed from the

network. In particular, an Alternate Port provides an alternate path to the one offered by the Root Port, in the direction of the Root Bridge. A Backup Port, however, can be used whenever the existing path offered by a Designated Port towards the leaves of the spanning tree becomes unavailable [104].

Initially, Alternate and Backup Ports are quickly transitioned to the Discarding Port state to avoid data loops. In the case of bridge or LAN failure, the fastest local recovery is possible when the Root Port can be substituted for the Alternate Port. The related advantage is that as long as the Root Port Path Cost is equal for both ports, bridges located further from the Root Bridge will not see a network topology change [104]. However, in the other case, MSTP reconfigures the topology based on the spanning tree priority vectors, behaving like a distance vector protocol. Note that during the reconfiguration phase, old information related to the prior Root Bridge may still circulate in the network until it ages out. For details related to the operation principles of MSTP, the reader is referred to [104].

A similar approach, which is also the main underlying concept of several other mechanisms, was proposed in Viking [86]. It was designed for a wide range of networking technologies, such as Local-Area Networks, Storage-Area Networks, Metropolitan-Area Networks, and Cluster networks, to provide fast recovery, high throughput, and load balancing over the entire network. Viking is based on multiple spanning trees covering the same network topology. It relies on the VLAN technology widely supported in enterprise-grade Ethernet switches to control how packets are forwarded towards their destinations. In Viking, each packet carries a VLAN identifier associated with one of the available spanning trees. Based on that identifier, downstream switches forward the packet along the path in the corresponding spanning tree, as shown in Fig. 11. While relying on existing failure detection mechanisms implemented in modern network switches, Viking assumes that all end hosts run a local instance of the *Viking Node Controller* responsible for load measurements and VLAN selection. Whenever a link or node becomes unavailable, the centralized *Viking Manager* instructs the Viking Node Controllers using the affected VLANs to change the VLAN identifier carried in subsequent packets, effectively redirecting the related flows onto the precomputed alternative paths.

Compared to the Ethernet architecture relying on a single spanning tree, Viking offers higher aggregate network throughput and the overall fault tolerance. Such an improvement is possible by reducing the downtime to sub-second values and preserving the existing higher-layer connections [86]. It operates in a semi-distributed way¹ and requires the SNMP protocol for internal signalling. The prototype developed by the authors relies on the Per-VLAN Spanning Tree (PVST) implementation by Cisco. As no firmware modifications are necessary, Viking can be deployed on many off-the-shelf Ethernet switches [86].

A distributed recovery scheme for a single-link failure

Fig. 11: Illustration of the spanning-tree switching process based on the VLAN identifier stored in each forwarded Ethernet frame.

scenario, which involves the “a priori” configuration of the alternative trees and the use of VLAN IDs to switch the traffic onto the alternative tree after a failure, was presented in [75]. In this scheme, restoration is performed locally by a node upstream to the failed element. Two mechanisms of recovery are proposed, namely connection-based and repair-based recovery. In connection-based recovery, packets are assigned the backup VLAN ID based on the source node, destination node, and the primary VLAN ID. It means that traffic from different connections can be switched at a given node onto different backup spanning trees. At the same time, in the destination-based approach, packets are assigned the backup VLAN ID only based on the primary VLAN ID and the destination node. As a result, flows destined at a given node from different connections are switched at a given node onto the same backup spanning tree. The latter approach is less complicated and involves a shorter computation time. However, as presented in the evaluation section of the paper, it is less capacity-efficient than the former scheme which determines the transmission paths based on a broader set of input parameters. To avoid forwarding loops, switching the traffic between spanning trees is allowed only once, and controlled by setting one bit in the Class of Service (CoS) field as the recovery bit².

Another local protection scheme called EFUS based on multiple trees and the use of VLANs has been presented in [77], [89]. Its advantage is the ability to provide recovery also in the case of failure of a single *node* if only the network topology is at least 2-connected. It is possible by the utilization of a pair of spanning trees for each network node, and by switching the flow onto the respective protection tree by the immediate upstream node onto the alternate spanning tree.

¹Note that the authors planned to design and evaluate a fully distributed version in their second prototype.

²The CoS field is included in Ethernet frames under 802.1Q VLAN tagging [79]

C. Solutions Based on Recovery Tables

In [92], a method to reduce service recovery time after a single link failure is proposed that is not based on spanning trees, but uses the recovery tables storing the alternate next-hop information for each entry in the conventional forwarding table. As proposed in [92], entries in the recovery table form the respective detours calculated based on information on the shortest paths maintained in the control plane provided by a routing protocol such as OSPF. The scheme uses the concept of VLANs, and changes the uppermost bit of the VLAN ID from 0 to 1 when redirecting the traffic based on the recovery forwarding table, or, to avoid loops, discards the packet if trying to redirect an already redirected one. The use of the uppermost bit of the VLAN ID for recovery indication limits the number of VLANs which can be established.

A resilience scheme for a spanning tree which is based on the preconfigured protection paths and tunnelling was introduced in [99]. To provide fast recovery after a link failure, protection paths are established before the occurrence of the failure by means of protection cycles (*p-cycles*, [105]) defined for each link in the spanning tree. After a failure, the node detecting the failure proceeds with the encapsulation of packets and forwarding them via the the respective protection cycle to detour the failed link.

D. Solutions Based on Message Flooding and Deduplication

As the performance of the link layer had already been identified as a growing problem, the AXE scheme was proposed in [97], [100] as a solution that not only retains the plug-and-play behaviour of the Ethernet, but also ensures near-instantaneous recovery from failures and supports general network topologies. What distinguishes this fast-recovery mechanism from the widely-used Ethernet is that all network links can be used to forward packets, instead of using only a subset of links forming a spanning tree. In particular, AXE is based on the flood-and-learn approach and employs an AXE packet header containing the following four additional fields: the *learnable* flag *L*, the *flooded* flag *F*, a hopcount *HC*, and a nonce used by the deduplication algorithm. At the same time, AXE takes an orthogonal approach to the traditional flood-and-learn Layer-2 mechanisms such as STP in that it does not rely on any control plane at the link layer to compute the spanning tree. Instead, when an AXE node does not know how to reach a destination, it encodes a hop count in the packet header, sets the *learnable* flag in the packet header, and floods the packet throughout the entire network. While receiving multiple copies of the flooded packet, each node learns only the shortest path towards the destination. To this end, since flooding may reduce bandwidth resources, each node maintains a packet deduplication filter implemented as a cache, which avoids flooding packets in cycles — a catastrophic event in a network.

It is worth noting that AXE remains compatible with existing failure detection techniques, such as BFD and different hardware-based mechanisms. According to [97], even while using relatively small filters, AXE can scale to large Layer-2 topologies and can quickly react to failures by avoiding the STP-related computations.

E. Summary

In this section, we highlighted the mechanisms introduced in the literature to recover from failures of network elements at the link-layer level. Particular focus was on presenting the design characteristics contributing to the reduction of the recovery time. In this context, we first discussed the aspects of the conventional the IEEE 802.1D Spanning Tree Protocol (STP), and in particular its remarkably slow convergence time measured even in terms of several tens of seconds. Next, we analyzed the representative schemes aimed at shortening the time of recovery grouped by us into four categories: (a) solutions based on a single spanning tree, (b) schemes utilizing multiple spanning trees, (c) techniques using the recovery tables, and (d) strategies based on message flooding and deduplication. In particular, schemes utilizing a single spanning tree such as, e.g., IEEE 802.1D Rapid Spanning Tree Protocol (RSTP) were found to require significantly less time to complete the recovery procedure of a spanning tree after a failure. By sending, e.g., explicit termination messages instead of using timeouts (as in [85]), or by executing the offline an Integer Linear Program to determine the best set of links for any single link failure scenario in advance ([91]), the recovery of a spanning tree could be completed by those methods only within several tens of milliseconds.

We then discussed the schemes utilizing multiple spanning trees, which can result in the partitioning of a network into subareas with one spanning tree installed in each such region and identified by a given VLAN ID (see, e.g. [104]). Such schemes were also shown in the respective literature to reduce the time of the recovery phase to sub-second values.

Schemes based on recovery tables (e.g., [92]) were next shown to be an essential alternative to tree-based techniques able to recover quickly from failures due to the pre-planned calculation of the alternate next-hops stored in recovery tables at each network node.

Finally, techniques based on message flooding and deduplication such as, e.g., AXE ([100]) relied on utilizing all network links (as opposed to schemes using only links belonging to spanning trees) to forward packets, as well as to recover quickly from failures.

Despite a broad set of link-layer fast recovery mechanisms introduced in the literature, their limited implementation in practice remains an open issue. Indeed, apart from the standard solutions such as IEEE RSTP or MSTP, it is still rare to find the other fast recovery mechanisms implemented in commercially available switches. Also, as link-layer recovery mechanisms were primarily designed for scenarios of single (link/node) failures, another open issue refers to their ability to recover from simultaneous failures of multiple network elements (i.e., due to an attack or another disaster-induced event).

IV. MPLS FAST RECOVERY

The architecture of Multiprotocol Label Switching (MPLS), introduced in [109], relies on Label Switching Routers (LSRs) capable of forwarding packets along Label Switched Paths (LSPs) based on additional labels carried in a packet header.

Fig. 12: Timeline of the selected documents and solutions related to MPLS Fast Reroute (entries marked in gray provide the general context related to the evolution of MPLS).

Fig. 13: Illustration of the two basic local protection methods defined in [113]: (a) *One-to-one backup* and (b) *Facility backup*. In the first figure, the thick black path leading from R1 to R5 represents the primary LSP, the thick blue path leading from R1 to R3 via R6 — the backup LSP protecting node R1, and the thick dashed red path — the backup LSP protecting node R2. In the second figure, the thick solid black path and the blue dotted path represent the primary LSPs, while the thick red path leading from R2 to R4 via R7 is the shared backup LSP which protects both primary LSPs if node R3 or links between R2-R4 fail.

Each of the labels assigns the packet to the corresponding Forwarding Equivalence Class (FEC) that defines a group of IP packets to be forwarded by one or more LSRs in a consistent manner. To be able to distribute information about the assignment of labels to the corresponding FECs among the LSRs in an automated way, a label distribution protocol may be deployed in the network. One of the available implementations is the Label Distribution Protocol (LDP) defined in [116].

In the following sections, different approaches in terms of the design of Fast-Reroute mechanisms for MPLS networks are presented and discussed. In addition, the evaluation results are also reported for the selected solutions deployed in real test network environments, to provide valuable context related to their expected performance.

A. MPLS Fast-Reroute Extensions

As MPLS had many advantages and was recognized to be a promising solution, it was soon extended to support such important functionalities as traffic engineering [110] (further developed in the inter-domain context in [117]) and Fast-Reroute mechanisms to protect LSP tunnels [113]. In particular, the Fast-Reroute extensions enabled local repair of LSP tunnels within 10s of milliseconds based on the established backup LSP tunnels. The following two local protection methods were defined (see Fig. 13):

- *One-to-one backup*: one backup LSP is established for each protected LSP in such a way that it intersects the primary path at one of the downstream nodes;
- *Facility backup*: one backup LSP is established to protect a set of primary LSPs in such a way that it intersects each of the primary paths at one of its downstream nodes.

It is possible to deploy the two listed methods either together or alone. Furthermore, each of them is suitable for protection of links and nodes in the event of network failure.

The general approach to using the Not-Via Fast-Reroute concept in MPLS networks was proposed in [124]. It is based on a single-level encapsulation and forwarding of packets to specifically reserved IP addresses which are also advertised by the IGP. In this way, it is possible to protect unicast, multicast, and LDP traffic against failure of a link, a router, and a shared-risk group. For a detailed discussion of the Not-Via mechanism in the context of IP networks, the reader is referred to Section V-B.

Finally, the Resource Reservation Protocol — Traffic Engineering (RSVP-TE) Fast-Reroute procedures previously defined in [113] have recently been updated in [127] to support Packet Switch Capable (PSC) Generalized MPLS (GMPLS) LSPs. In particular, new signaling procedures and a new `BYPASS_ASSIGNMENT` subobject in the RSVP `RECORD_ROUTE` object are defined to coordinate the assignment of bidirectional bypass tunnels which protect common facilities in both directions along the corresponding co-routed bidirectional LSPs. Consequently, bidirectional traffic can be redirected onto bypass tunnels in such a way that the resulting data paths are co-routed in both directions. The proposed Fast-Reroute strategy may be used either with GMPLS in-band signaling or with GMPLS out-of-band signaling. It also allows to avoid the RSVP soft-state timeout in the control plane which would typically occur after one of the downstream nodes stopped receiving the RESV messages associated with the protected LSP from the upstream Point of Local Repair (PLR) — for example, as a result of unidirectional link failures.

Interested readers looking for structured information covering general terminology and schemes related to protection and restoration techniques in the context of MPLS and Generalized MPLS (GMPLS) may be referred to [114], [115], [121].

B. Methods Based on Optimization

An important subgroup of MPLS Fast-Reroute solutions relies on mathematical optimization to improve the overall performance with respect to different factors, such as: resource utilization, coverage of failure scenarios, and scalability. We summarize the selected approaches below.

To be able to perform the failover from the primary LSP to a precomputed backup LSP in a given failure scenario, the network requires additional resources, such as available bandwidth on network links along the related LSPs. On the other hand, the resources that remain reserved specifically for the purpose of the possible recovery cannot be assigned to other primary LSPs, and thus should be minimized. However, they can be shared by two or more backup LSPs, as long as the corresponding primary LSPs do not fail simultaneously. An illustration of this approach in the context of distributed bandwidth management supporting MPLS Fast Reroute is provided in [118]. The authors focus on the one-to-one protection method and single link or node failures. Although additional information required by the proposed backup path selection algorithm is distributed between adjacent routers using three

different signaling messages, the authors suggest that some information be embedded in the existing `PATH` and `RESV` messages. In such a case, only one new signaling message would need to be introduced in real deployments.

Resilient Routing Reconfiguration (R3) was proposed to address the long-standing shortcomings of earlier fast recovery techniques with respect to insufficient performance predictability and missing protection mechanisms against possible network congestion, especially in multi-failure scenarios [119]. R3 is based on the following two main steps:

- 1) *Offline precomputation*: find a suitable routing scheme and protection (rerouting) scheme for a given traffic matrix, to minimize the maximum link utilization over the entire network;
- 2) *Online reconfiguration*: in the case of failure, reroute traffic via alternative paths and adjust the routing and protection schemes to exclude the failed link from the set of candidate links in the event of subsequent failures.

The first step is not time critical and relies on linear programming duality to convert the related primary optimization problem with an infinite number of constraints into a simpler form containing a polynomial number of constraints³. Although the corresponding model contains $O(|V|^4)$ variables and $O(|V|^4)$ constraints, where V denotes the set of nodes in the network graph, the authors emphasize that their estimation is much lower than for the other existing approaches focused on oblivious routing [133], [134].

In the second step, it is important that the rerouting decision be made as soon as possible, to avoid packet losses and increased delay. Thus, the related operations are designed not to be computationally intensive. In addition, the routing and protection schemes at all involved routers may be updated after the upstream node adjacent to the failed link has started rerouting packets via the selected detour, without affecting the recovery process.

The typical problem of primary and backup path computation in the context of MPLS Fast Reroute was also discussed in [106], [126] where the authors proposed the corresponding nonlinear integer programming models related to link, node, or path protection schemes. To reduce the expected computational complexity, subsequent efforts were made to prepare the linear variants of the related optimization problems which also take into account the available resources on network links [128], [131].

C. Fast Verification of MPLS Networks

Testing and debugging data plane configurations is generally considered a difficult manual task, yet, a correct configuration of the data plane is mission critical to provide the required properties in terms of policy compliance and performance. Reasoning about the data plane behavior *subject to failures* is particularly challenging, as it seemingly introduces a combinatorial problem: it seems unavoidable that one has to test each possible failure scenario individually, and simulate the

³Interested readers may learn the basics of optimization theory from [132].

resulting rerouting, in order to verify that the network behaves correctly under all failure scenarios.

Interestingly, this intuition is wrong: it has recently been shown that in the context of MPLS networks, it is possible to conduct *what-if analyses* in *polynomial time*. In [129], an approach is presented to collect and transform MPLS forwarding and failover tables into a prefix rewriting system formalism, to which automata-theoretical algorithms can be applied to test a wide range of properties, related to reachability (e.g., can A always reach B , even if there are up to 3 link failures?) or network policy (e.g., is it ensured that traffic from A to B traverses a firewall along the way?).

The fast verification is enabled by the nature how labels are organized and manipulated in the MPLS packet header: the labels are organized in a *stack*, and operations limited to *push*, *pop*, and *swap*. This makes it possible to describe the system as a push-down automaton. In [130] a tool called *P-Rex* is presented which realizes the theoretical concepts in [129].

To the best of our knowledge, no polynomial-time solutions exist for conducting what-if analyses in polynomial time for other types of networks which rely on more complex rules, e.g., [135], [136].

D. Other Approaches

Beyond the groups of solutions summarized in the sections above, each of them sharing a common design feature, there are also other fast recovery concepts related to MPLS which are based on interesting ideas, and thus are also worth mentioning. We discuss the selected proposals in this section.

The idea of sharing some backup resources in the network by multiple backup bandwidth-guaranteed LSPs has already been considered in [111] where the authors discuss the trade-off between complete knowledge of the routing scheme at the time of a path setup event and partial knowledge based on the aggregated link utilization information. Maintenance of non-aggregated per-path information may cause potential scalability issues and requires more storage and processing resources. On the other hand, the solution proposed by the authors relies on such information as fraction of bandwidth used on each link by the active LSPs and separately by all backup LSPs. Based on that, the proposed routing algorithms are able to determine the candidate backup paths for local recovery to protect single link or node failures.

A fast recovery method based on diverting traffic back towards the upstream LSR and selecting one of the predefined alternative LSPs was proposed in [108] and improved further in [112], [137]. While the former method already allowed for a significant reduction of the overall path computation complexity and signaling requirements⁴, the latter method was designed to eliminate packet reordering which was one of the disadvantages of the earlier proposals. The average delay during the restoration period has also been improved. The underlying idea of the related fast recovery strategy is illustrated in Fig. 14.

⁴In addition, note that the computations related to primary and alternative paths may be performed at a single switch, to avoid possible issues resulting from distributed computation.

Fig. 14: Illustration of the underlying idea of the improved protection methods described in [112], [137].

Once a failure is detected by one of the LSRs on the primary LSP, packets are sent along the *backward* LSP towards the upstream LSR. The upstream LSR recognizes the backward flow, marks the last packet sent along the broken LSP using one bit of the `Exp` field of the MPLS label stack (note that no overhead is introduced at this point), and stores the subsequent packets received from the upstream LSR in a local buffer to avoid packet reordering. As soon as the previously marked packet is received again from the downstream LSR, all related packets stored in the buffer are forwarded to the upstream LSR. Eventually, the source LSR of the protected LSP redirects packets onto the predefined *alternative* LSP. It is worth noting that this recovery strategy may be used to achieve one-to-one and many-to-one protection faster as well as to avoid network congestion, allowing for better QoS control and reduced packet losses [112], [137]. More importantly, by introducing relatively small packet buffers at LSRs to be able to store copies of the limited number of forwarded packets, it is possible to eliminate packet losses entirely and thus improve the observed TCP performance during the failover considerably [137]. At the same time, the proposed two improved approaches do not specify any particular method for the effective selection of alternative LSPs — an important factor having significant influence on the QoS and the observed path stretch. Interestingly, the authors of [108] already anticipated the potential problems in the context of delay-sensitive network services and outlined the concept of *restoration shortcuts* as one of the possible ways to counteract the expected increase of transmission delay as a result of the failover. In particular, while using a restoration shortcut, traffic is rerouted over an alternative *shortcut* LSP established between the LSR on the primary path (upstream of the failed link) and the destination of the primary LSP, potentially merging into the other existing backup LSPs.

Resilience in MPLS networks can also be provided by multipath structures being a type of protection-switching mechanisms. One of the representative schemes of this kind is the self-protecting multipath (SPM) concept proposed for MPLS networks in [138]. As presented in Fig. 15, for a given end-to-end demand d , in the normal operation mode, SPM uses all pre-established paths for data transmission. The multipath structure should include paths being mutually node-disjoint, to ensure protection against a failure of a single node/link as well as to simplify the procedure to establish the multipath. Load balancing among several paths of the multipath is also helpful while dealing with failures. In the case of any failure affecting a given path in the multipath structure, traffic is redistributed across all the other working paths. Therefore, for a multipath consisting of k paths, SPM needs $k + 1$ different

traffic distribution functions: one for operation in the failure-free scenario, and the next k functions to cover the failure of any of the k paths. The evaluation presented in [138] shows that the approach is very efficient in protecting against failures of single nodes and links, as it requires only about 20% of additional transmission capacity for this purpose.

Fig. 15: Example configuration of a multipath.

One of the variants of SPM has been presented in [139], with the extension to introduce the traffic distribution function specific to the failed network element along the path, instead of the former traffic distribution function defined in the context of a given path.

The SPM concept has been enhanced in [140] with the proposal of a linear program (LP) to optimize the SPM load balancing function to maximize the amount of traffic transported with resilience requirements for legacy networks (i.e., already deployed networks). The objective is achieved by solving the corresponding problem of minimization of the maximum link utilization in any protected failure scenario.

E. Experimental Evaluation

Some of the existing solutions have been evaluated in real test network environments. As the corresponding results may provide a valuable context with respect to the expected performance, we summarize below the selected reports related to fast recovery mechanisms.

The first paper discussed in this section presents the comparison of the following two mechanisms with respect to packet losses and duration of the failover phase [120]:

- MPLS Traffic Engineering (TE) Fast Reroute [107], [113], [123];
- IP Fast Reroute Framework [141].

The evaluation scenarios assumed single link failures and protection based on either predefined backup paths using MPLS TE Fast-Reroute tunnels or IP Fast-Reroute Loop-Free Alternates (for a detailed discussion of IP Fast-Reroute solutions, the reader is referred to Section V). The results presented by the authors suggest that with increasing number of rerouted LSPs in MPLS networks, the failover time increases exponentially (even beyond 50 ms) and packet losses are also higher. In the case of the investigated IP network, increasing the number of IP prefixes also causes a non-linear increase of the failover time, while traffic losses remain almost unchanged.

In contrary to the experiments reported in [120], the second analyzed paper presents the results from a large and geographically-distributed production backbone network in which every major routing node consisted of several core, aggregation, and edge routers [122]. The observed parameters

included TTL changes, packet losses, packet reordering, and one-way delay changes. The experiments were performed over a 14-month period and the two considered restoration methods (OSPF reconvergence and MPLS-TE Fast Reroute) were analyzed during seven consecutive months each. The evaluation results have shown that the MPLS-TE Fast-Reroute mechanism reduces packet losses and packet reordering during link failure events significantly, also suppressing the possible micro-loops.

The third considered paper compares the effectiveness of MPLS-TE Fast Reroute and IP OSPF Fast Reroute based on Loop-Free Alternates in terms of packet losses and network convergence time, for different numbers of LSPs and IP prefixes [125]. Both mechanisms were configured in a WDM test network. The reported results confirm that with increasing number of rerouted LSPs in MPLS networks, the convergence time increases significantly and the observed packet losses are also higher. Again, in the case of the investigated IP network, increasing number of IP prefixes caused a non-linear increase of the convergence time, while traffic losses remained almost unchanged.

Fast recovery mechanisms designed for MPLS have been shown to improve network operation and performance in different failure scenarios. At the same time, there are still some related open challenges. In particular, fast restoration relying on backup LSPs requires that additional LSPs be configured and established in advance and in an effective way, taking into account the trade-off between the coverage of failure scenarios, the available network resources, expected traffic demands, key performance indicators from the perspective of users, and the overall complexity of the system. Moreover, considering the variety of failure detection and mitigation mechanisms running concurrently on different layers of networked systems, another challenge is how to avoid packet losses and possible packet reordering during the failover.

V. INTRA-DOMAIN NETWORK-LAYER FAST RECOVERY

In this and the next sections, we review the representative schemes for fast recovery that operate at the network layer (Layer 3) of the Internet protocol stack. In particular, this section is dedicated to the intra-domain setting, while the next section will discuss the solutions that work across Autonomous Systems (inter-domain). Since the prevailing network-layer protocol today is the Internet Protocol suite, namely the two versions IPv4 and IPv6, most of the fast recovery schemes we discuss here are specifically designed for IP (IP Fast ReRoute, see below). However, the main ideas will be reusable in any network layer that provides connectionless, unreliable datagram service [142].

In the context of a single provider network, the aim of the operator is to achieve the highest possible level of failure resilience possible at the smallest resource footprint and the minimum disruptions to the large deployed base of operational IP networking hardware and software. The main difficulty here is that the packet transmission service provided by IP is fundamentally connectionless; this inherently interferes the operator's intention to "route around" failures, given that

Fig. 16: A taxonomy of the most important general concepts in IP Fast ReRoute.

the underlying network layer does not allow packets to be “pinned” to a detour. Correspondingly, most IP fast recovery schemes need to tediously work around the limitations of the underlying network layer.

As outlined in Section II, the lowest service recovery time and the highest Quality of Resilience (QoR) can be obtained with *preplanned* protection schemes (with backup paths established before the failure), providing detours over small parts of working paths (i.e., *local detours*) with backup capacity reserved in advance to ensure an undisrupted flow of traffic after a failure (*dedicated* protection). For standard IP networking, however, it is not common to see these objectives fulfilled jointly, since the proactive determination of communication paths with the advance reservation of the respective network resources is difficult given the connectionless nature of IP. Therefore, IP Fast Reroute (IPFRR) mechanisms typically rely on shared detours and provide “best-effort” protection only. In other words, in IP failure recovery, usually there is no guarantee that the necessary capacity and network resources remain available along the backup paths after the occurrence of a failure (see, e.g., [143], [144]).

Below, we review the most important intra-domain fast IP recovery schemes following a rough chronological order and we provide a simple taxonomy for classifying the schemes. We note, however, that our coverage of IP Fast ReRoute and related concepts is intentionally noncomprehensive. In particular, we concentrate on *shared*, *preplanned*, *local protection schemes* exclusively (recall Fig. 6), which either work on top of the unmodified connectionless destination-based unicast IP data plane service and a distributed intra-domain IP control-plane protocol (like OSPF or IS-IS, [10], [11]) or require only minimal extension to the bare-bone IP specification, like non-disruptive modifications of the routing protocols or minimal central coordination [145]–[147]. For fast IP-based *restoration* schemes (in contrast to *protection*), see [15], [148], [149]; for pointers on IP multicast fast recovery see, e.g., [150]; and for schemes that do not fit into this pure IP data-plane “ideal”, like O2 routing [151], Failure-carrying Packets (FCP) [152], Protection routing [143], [144], or KeepForwarding [153], see Section VII. We deliberately ignore the intricate issues related to IP fault detection and fault localization, noting that IP recovery usually relies on (from the slowest to the fastest technique) control plane heartbeats, Bidirectional Forwarding Detection, and link-layer notification, see Section II-C. For further detail on IPFRR refer to [23], [24], for the algorithmic aspects see [154]–[158], and for a comprehensive evaluation

and comparison of different IPFRR techniques see [35], [159]. For a taxonomy of the most important concepts in IPFRR, see Fig. 16, and a timeline of related research and standards, see Fig. 17.

A. The IP Fast ReRoute Framework

When a link or node failure occurs in a routed IP network, there is inevitably a period of disruption to the delivery of traffic until the network re-converges on the new topology. Packets for destinations that were previously reached by traversing the failed component may be dropped or may suffer looping [160]–[162]. Recovery from such failures may take multiple seconds, or even minutes in certain cases, due to the distributed nature of the IP control plane [15], [148], [149]. In a typical IP control plane configuration, one or more shortest-path-based intra-domain link-state Interior Gateway Protocols (IGPs), like OSPF (Open Shortest Path First [10]) or IS-IS (Intermediate System-to-Intermediate System [11]), and a policy-based inter-domain path-vector Exterior Gateway Protocol (EGP), like the Border Gateway Protocol (BGP [45]), interact in complex ways in order for routers to converge along stable routes [15], [148], [149]. Many Internet applications (multimedia, VPN) have not been constructed to tolerate such long disruptions.

Reports suggest that roughly 70% of outages in operational IP backbones are *local*, affecting only a single link and only a single Autonomous System (AS) at a time, and *transient*, lasting only a couple of seconds [149], [163], [164]. For such transient local failures, it is not worth executing two full IGP re-convergence processes in a short period of time.

In this section, we review the *IP Fast ReRoute framework* (IPFRR [141]), a mechanism to provide fast data-plane-driven failure mitigation in an intra-domain unicast setting. The first specification for IPFRR was drafted in 2004, reaching the status of an Informational RFC to 2010 [141]. The main goal of IPFRR is (1) to handle short-term disruptions efficiently and (2) remain fully compatible with the IP control plane and data plane, allowing for the incremental deployment with no flagship date. The framework rests on two main design principles: *local rerouting* and *precomputed detours* (recall the general taxonomy in Section II-B). Local rerouting means that only routers directly adjacent to a failure are aware of it, which eliminates the most time-consuming steps of IGP-based restoration, the global flooding of failure information. Additionally, IPFRR mechanisms are proactive in that detours are computed, and installed in the data plane, before a failure occurs. Thus, when a failure eventually shows up, the affected

Fig. 17: Timeline of selected documents and solutions related to IP Fast Reroute.

routers can switch to an alternate path instantly, letting the IGP to converge in the background.

The IPFRR framework distinguishes *local repair paths* (ECMP and LFA), the cases when a router has an immediate IP-level neighbor with a path that is still functional after the failure, and *multi-hop repair paths*, whenever the closest router with a functional repair path is multiple IP links away and, therefore, is not available directly via a local interface (rLFA).

B. Shortest-path Fast Reroute

Most intra-domain IP routing protocols (IGPs) rely on a flooding mechanism to distribute network topology information across the routers inside an AS and a shortest-path-first algorithm, i.e., Dijkstra’s algorithm [32], to calculate the best route, and the primary next-hop router(s) along these routes, to be loaded into the data-plane forwarding table. The idea in all LFA-extensions discussed in this section is to *leverage the IGP link-state database*, readily available and synchronized at all routers, to compute not only a primary next-hop to each destination, but also to *obtain one or more secondary next-hops* as well that can be used as bypass whenever the failure of the primary next-hop is detected. As such, shortest-path-based IPFRR is generally easy to implement and incrementally deploy in an operational network. This, however, comes at the price of a major limitation: namely, as *the bypass paths themselves must also be (at least partially) shortest paths*, it may happen that a proper shortest bypass path does not happen to be available in a given topology for a particular failure case. Correspondingly, in general it is very difficult to achieve 100% protection against all possible failure cases with shortest-path-based IPFRR methods.

Loop-free Alternates (LFA) is the basic specification for the IP Fast ReRoute framework to provide *single-hop repair paths*. In LFA, the emphasis is on simplicity and deployability,

instead of full coverage against all transient failures [32], [172]. Drafted in the IETF Routing Working Group in 2004, LFA reached Standards Track RFC status in 2008 (two years before the actual IPFRR framework specification was finalized [141], see Fig. 17).

As mentioned above, the idea in LFA is to exploit the information readily available in the IGP link-state database to compute secondary next-hops, or “alternates” as per [32], that can be used as a bypass whenever the failure of the primary next-hop is detected. Computing these secondary next-hops must occur in a “loop-free” manner so that the bypass neighbour, which, recall, will not be explicitly notified about the occurred failure and hence will not be aware that it should apply special forwarding decisions to pin the packets to the detour, will not loop the packets back to the originating router. LFA uses some basic graph-theory and simple conditions based on the shortest-path distances calculated by the IGP to ensure that the calculated alternate routes are indeed loop-free.

The conditions based on which routers can choose “safe” alternate next-hops are as follows. Given router s , destination prefix/router d , let e be a shortest-path next-hop of s towards d (there can be more than one next-hop, see below). In addition, let $\text{dist}(i, j)$ denote the shortest-path distance between router i and j . Then, from s to d with respect to the next-hop e , a neighbour $n \neq e$ of s is

- a *link-protecting LFA* if

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

- a *node-protecting LFA* if, in addition to (2),

$$\text{dist}(n, d) < \text{dist}(n, e) + \text{dist}(e, d) , \quad (3)$$

- a *downstream neighbour LFA* if

$$\text{dist}(n, d) < \text{dist}(s, d) , \quad (4)$$

- and an *Equal Cost MultiPath (ECMP)* alternate if

$$\text{dist}(s, n) + \text{dist}(n, d) = \text{dist}(s, d) . \quad (5)$$

The definitions follow each other in the order of “strength” and generality: for instance, an ECMP alternate is always a downstream neighbour, and a node-protecting LFA is also link-protecting. For the exact relations among different LFA types, see [32], [172], and for an illustration of the main concepts in LFA, refer to Fig. 18. In general, a “stronger” notion of LFA alternate next-hop should always be preferred over a “weaker” one whenever multiple choices are available. An algorithm to choose the best option in such cases is specified in [32].

The main advantages of LFA are that it is simple to implement and fully-compatible with the deployed IP infrastructure. Correspondingly, LFA is readily available in most major router products [173]–[175]. Nevertheless, LFA comes with a number of distinct disadvantages as well. First, it needs an additional run of Dijkstra’s algorithm from the perspective of each LFA candidate neighbour to obtain the shortest path distances, causing the extra control CPU load. Second, LFA does not guarantee full failure coverage: depending on the topology and link costs LFA can protect about 80% of single link failures

Fig. 18: Illustration of Loop-free Alternates and ECMP alternates with link costs as marked on the edges. For source router R1 towards destination router R6, both R2 and R3 are ECMP next-hops (marked by orange arrows in the figure), each one providing a node-protecting ECMP LFA with respect to the case if the other one fails, and router R4 is a link- and node-protecting LFA protecting against the (potentially simultaneous) failure of R2 and/or R3 (the backup route is marked by a dashed red arrow in the figure).

and 40–50% of node failures in general. Correspondingly, various optimization methods are available in the literature to improve failure-case coverage in operational networks [176]–[181]. Third, LFA is prone to forming transient “micro-loops” during recovery, caused by certain routers still using the “normal” routes while others already switching to the recovery routes [32], [182], and possibly “LFA loops” as well that may show up after all routers have switched. For instance, using a link-protecting LFA to protect a node failure may generate an LFA loop that will persist until the IGP finishes full reconvergence in the background [32], [183], [184]. In general, there is a trade-off between loop-avoidance and failure-coverage. For example, using only the strong notion of a downstream neighbour (cf. Eq. (4)) as an LFA eliminates both micro- and persistent LFA loops, but the failure case coverage attained this way may be poor in certain provider networks [172].

U-turn Alternates is an extension to the basic LFA specification providing multi-hop repair paths, on top of the local-repair mechanism implemented by LFA, in order to improve the failure case coverage [185]. The U-turn alternates specification never reached RFC status.

The main observation in the U-turn alternates specification is that the only way for a router to remain without an LFA is if it is itself one of the default next-hops of each of its neighbours. Then, a U-turn alternate is a neighbour that has a further neighbour (at most two hops away) that still has a functional path to the destination after the failure. Consequently, that “next-next-hop” neighbour of the U-turn alternate can be used as a bypass whenever the primary next-hop fails and no LFA is available. Unfortunately, this requires some external mechanism to prevent the U-turn alternate to loop the packet back to the originating router (which would be the “default” behaviour). See Fig. 19 for an illustration.

Technically, for router s , destination prefix/router d , a neighbour n of s is a *U-turn alternate* next-hop from s to d if (1) s is the primary next-hop from n to d and (2) n has a node-protecting loop-free LFA to d by (3). In such cases, s can send a packet to n , which, detecting that a packet was received from its primary next-hop (this can be detected by, e.g., Reverse Path Filters [186]), can send it along its

node-protecting LFA. Alternatively, a packet sent to a U-turn alternate can be explicitly marked to signal that it should not be looped back to the originating router.

The U-turn alternates specification is relatively simple to implement. On the other hand, it needs RPF, per-packet signalling, tunnelling, or interface-specific forwarding (see below) to indicate that packet is travelling to a U-turn alternate, plus an additional run of Dijkstra’s algorithm from each neighbour of every U-turn candidate. Still, depending on the topology U-turn alternates cannot guarantee full failure coverage, not even for single link failures. To account for non-complete coverage, [187] presents an extension where multiple U-turns may come one after the other. This provides 100% protection at the cost of worsening the issues related to the signalling of U-turns.

Remote Loop-free Alternates (rLFA) is another extension of the basic LFA specification to extend the scope to multi-hop repair paths [157], [158], [170], [188]. Again, the intention is to improve LFA failure case coverage.

The most important observation that underlies rLFA is that any remote router may serve as an alternate, provided that (1) the remote router has a functional path to the destination after a failure, (2) the shortest path from the originating router to the alternate does not traverse the failed component, and (3) the originating router has some way to tunnel packets from itself to the alternate. In such cases, whenever a router needs to find a bypass to divert traffic away from a failed primary next-hop, it can encapsulate these packets to tunnel the diverted traffic to the remote loop-free alternate (using, e.g., IP-IP, GRE, or MPLS), so that the router at the tunnel endpoint will pop the tunnel header from the packets and use its default next-hop to reach the destination. See Fig. 19 for a sketch of the main ideas in rLFA.

In order to check whether a remote router is a valid rLFA candidate, the shortest path segment from the originating router to the rLFA and from the rLFA to the destination router both must avoid the failed component. The second condition is compatible with the LFA loop-free condition, whereas the first condition is needed because the encapsulated packets, travelling from the originating router to the rLFA, will also follow the shortest path and, as such, may also be affected by the failure.

Formally, for router s , destination router/prefix d , and next-hop e from s to d , some $n \in V$ (not necessarily an immediate neighbor of s) is a *link-protecting remote loop-free alternate* (rLFA) if the below two conditions hold:

$$\text{dist}(s, n) < \text{dist}(s, e) + \text{dist}(e, n) \quad \text{and} \quad (6)$$

$$\text{dist}(n, d) < \text{dist}(n, s) + \text{dist}(s, d) . \quad (7)$$

The node-protecting case can be defined similarly [157].

The pros of rLFA are that it remains largely compatible with IP and it is also straight-forward to implement on top of MPLS/LDP and segment routing [189].

On the negative side, rLFA may cause extra control CPU load as it needs each router to execute a shortest-path computation from potentially each other router, and extra data-plane burden by routers having to maintain possibly a huge

Fig. 19: Illustration of U-turn Alternates and Remote Loop-free Alternates with link costs as marked on the edges. Source router R1 has router R4 as its default shortest-path next-hop towards router R6 (marked by an orange array in the figure) and there is no available LFA candidate neighbour that would provide protection against the failure of the link R1–R4 (marked by a red cross), or the next-hop router R4 itself. However, router R5 is both a U-turn alternate and a Remote Loop-free Alternate in this case: note that both R2 and R3 are possible U-turn alternate next-hops that can pass the packet to R5, but in the case of rLFA the bypass path will be provided along the shortest R1–R5 route exclusively (i.e., by router R2), yielding the new path R1–R2–R5–R6 (marked by dashed red arrays).

number of tunnels to reach each rLFA. Crucially, rLFA may still not provide full protection, not even in the unit-cost case. Papers [157], [158] provide analytical and algorithmic tools to characterize rLFA coverage in general topologies. To address the inadequacy of the failure case coverage, the originating router may use *directed forwarding* to instruct the rLFA to send a packet through a specific neighbour. This modification guarantees full coverage for single-link failures [188], [190]. Interestingly, this use case served as one of the precursors for the development of segment routing [189] and the main motivation to develop a family of LFA extensions that provide complete failure case coverage in the context of this emerging segment routing framework [191]. We note that the rLFA specification largely replaced U-turn alternates.

IP Fast ReRoute using Not-via Addresses (Not-via) is a specification for fast IP failure protection that addresses the limitations of LFA and rLFA. Drafted in 2005, the specification reached RFC status in 2013 [124] but major vendor adoption and large-scale deployments did not ensue.

The main drawback of rLFA is that even if a suitable remote LFA candidate is available after a failure the originating router may not have a way to send bypass packets there, since all its shortest paths to rLFA candidates may converge along a single, possibly failed, next-hop. Not-via overcomes this problem by tunnelling/encapsulating packets through a special “not-via” IP address that explicitly identifies the network component that the repair must avoid. Hence, each router can pre-compute backup paths covering each individual failure case, by taking a modified graph from which the failed component was explicitly removed and then re-computing the shortest paths in this modified graph. In operation, any router along the detour, receiving a packet destined to a not-via address, immediately (1) knows that this packet is currently travelling on a detour and therefore the default routing table next-hop should not be

applied, and (2) identifies the failed component associated with the not-via address and switches to the pre-computed backup shortest path. See Fig.20 for an demonstration of the main concepts in Not-via.

Suppose that router s detects that its primary next-hop e towards router/prefix d has failed. In such cases, s encapsulates the packet in a new tunnel header and into the outer IP header it sets the destination address as a not-via address. The not-via address is an ordinary IP address that was administratively configured, and advertised into the IGP to mean “an address of router d that is reachable not via node e ”. In order to cut down the length of the bypass paths, in the original not-via specification the destination of the repair path is not immediately d but rather the “next-next-hop” of s to d (i.e., the next-hop of e to d). This behaviour has been questioned several times since then [192]–[194]. Routers in the IGP domain will advertise not-via addresses alongside the standard IP prefixes and store a next-hop in the forwarding table for each “ x not-via y ” address by (1) temporarily removing component y from the topology and (2) calculating the shortest path to x in the modified topology. As long as a failure does not partition the network, this mechanism ensures that each single component failure can be protected.

Notably, Not-via was the first viable IPFRR specification to guarantee full failure coverage against single-component failures. Yet, it remains fully compatible with the IP data plane (but not with the control plane, see below), requiring no expensive modification to IP router hardware. Unfortunately, the resource footprint may still be substantial. First, it requires maintaining additional not-via addresses, introducing significant management issues (no standard exists as to how to associate a not-via address with a concrete router/failure case), control-plane burden (possibly thousands of not-via addresses must be advertised across the IGP [195] and each potential failure case requires an additional run of Dijkstra’s algorithm, and data-plane load (not-via addresses appear in the IP forwarding table, which already contains possibly hundreds of thousands of routes). In addition, tunnelling schemes, like rLFA and Not-via [124], [170], [196], [197] may raise unexpected and hard-to-debug packet loss or latency when the size of the encapsulated packets exceeds the MTU (Maximum Transfer Unit), causing a packet drop (when the “Don’t fragment” bit is set) or a costly and timely fragmentation/reassembly process at the tunnel endpoints (in IPv4).

The Not-via specification sparked significant follow-up work. [198] introduces aggregation and prioritization techniques to reduce the computational costs of Not-via and the forwarding table size. The paper also proposes an algorithm (rNotVia) that allows a router to efficiently determine whether it is on the protection path of a not-via address and cuts down unnecessary calculations. Lightweight Not-via [192], [193] aims to break down the management burden, decrease the computational complexity, and reduce the number of not-via addresses needed, based on the concept of (maximally) redundant trees [194], [199]. This modification allows to cover multiple failure scenarios using a single not-via address, which reduces the number of necessary not-via addresses to 2 per router. The question whether the combined use of LFA and

Fig. 20: Illustration of Not-via Addresses and the Failure Insensitive Routing with link costs as marked on the edges. Not-via: when link R2–R5 fails (marked by a red cross in the figure) along the shortest R1–R6 path (marked by an orange arrow), R2 will tunnel packets destined to R6 to the address “R6 not via R5”. The packet will travel along the shortest R2–R6 path with the default next-hop R5 removed from the topology (marked by a dashed dark red arrow) immediately to the next-next-hop (i.e., R6). FIR: when link R2–R5 fails, R2 recomputes its routing table but will suppress the notification for the rest of the routers. Rather, it sends the packet back to R1, which, having received it “out-of-order” from R2, infers that link R2–R5 and/or R5–R6 has failed and hence forwards it via R4 to the destination.

Not-via results in operational benefits is asked in [195]; the answer is generally negative.

Failure Insensitive Routing (FIR) using interface-specific forwarding is an IPFRR proposal originating from the academia, which gained significant following in the research community [165]. This is the first, and possibly the most elegant, IPFRR method proposed (the original paper appeared as early as 2003, predating even the first IPFRR draft, see Fig. 17), providing full protection against single link failures. Later versions and extensions address various shortcomings of FIR, e.g., provide protection against node failures ([168], [200]–[204], see below). Currently, we know of no off-the-shelf router products that implement FIR.

FIR is similar to Not-via in the sense that routers along a detour can identify the failed component. However, instead of using Not-via addresses that explicitly communicate the identity of the failure, in FIR routers rather *infer* it autonomously from packets’ flight: when a packet arrives through an “unusual” interface, through which it would never arrive under normal operation, the set of potential links whose failure could lead to this event can be inferred. Using this inferred information, routers can divert packets to a new next-hop that avoids the failed links. This provides full failure case coverage, at the cost of the modification of the default IP data plane: forwarding decisions in FIR are made not only based on the destination address in a packet, but are also specific to the ingress interface the packet was received on. See Fig. 20 for a sample topology demonstrating interface-specific routing.

The Failure Insensitive Routing (FIR) [165] scheme revolves around two basic concepts: *key links*, used to identify the potential failure cases that could lead to a router receiving a packet on a particular interface, and *interface-specific forwarding tables*, which assign a next-hop to each destination router/prefix separately for each ingress interface. Formally,

given router s , destination d , and next-hop e from s to d , the key links at s with respect to d for the interface $e-s$ are exactly the links along the $s \rightarrow d$ shortest path except $s-e$. The interface specific forwarding table with respect to this interface is obtained simply by removing the key links from the topology and running Dijkstra’s algorithm from s . As long as the topology is 2-connected, removing the key links will not partition the network and full failure coverage is attained.

Once computed, a router will install the per-interface forwarding tables into the data plane and, in normal operation, use standard shortest-path routing to forward packets. To deal with remote failures, the router does not need to react as the interface-specific routes are specific enough to handle all detours, while for local failures it recomputes its shortest paths but suppresses the IGP failure notification procedure from this point in line with the requirements of the IPFRR framework.

In a nutshell, in FIR a router detects that a packet is on a detour from that it is received from its primary next-hop. In this regard, FIR resembles U-turn alternates but it is much more generic than that (works for alternates more than two-hops away). As we have seen, it also generalizes Not-via in that it does not need additional not-via addresses to identify the failure, it can infer this information. It is remarkable how FIR predates both these proposals, yet the main ideas keep on being rediscovered in various forms until today [187].

FIR is remarkably elegant, simple, and fully distributed, and it also provides full failure case coverage, including node failures (see below). On the negative side, interface-specific forwarding is still not available in the standard IP data plane: while most 3rd generation routers store a separate forwarding table at each line card, management and monitoring APIs to these per-interface forwarding tables have never been standardized. In addition, FIR may create persistent loops when more than one link simultaneously fail. To address these issues, [205] extends FIR to handle node failures, [168], [201] generalizes this method to asymmetric link costs and inter-AS links, [202] provides a version that is guaranteed to be loop-free at the cost of increasing forwarding path lengths somewhat, while [200], [203], [206] present further modifications to the backup next-hop selection procedure based on interface-specific forwarding to handle, among others, double link failures.

C. Overlay-based Reroute

Achieving full failure case coverage in IPFRR and adhering to IP’s default connectionless destination-based hop-by-hop routing paradigm at the same time seems a complex problem. Each shortest-path-based IPFRR method we described above suffers from one or more significant shortcomings due to this fundamental contradiction, in terms of failure case coverage or deployability (or both).

The IPFRR methods based on the *fault isolation* technique described below offer a way out of this problem: they implement certain overlay networks on top of the physical topology so that for every possible local failure there will be at least one overlay unaffected by the failure that can be used as a backup. Of course, this may break IP compatibility, as some way to

force bypass traffic into the proper backup overlay is needed: most proposals use some additional bits in the IP header (e.g., in the ToS field), but tunnels, not-via addresses, and multi-topology routing could also be reused for this purpose. This allows bypass paths to no longer be strictly shortest, which lend more degrees of freedom to assign detours for different failure cases and leads to higher failure case coverage.

Below, we summarize the most important fault-isolation techniques proposed so far, with the emphasis on the methodology to obtain the overlay topologies themselves. In the taxonomy of Section II-B, the fault-isolation-based IPFRR mechanisms discussed below (not to confuse with fault detection and fault localization, see Section II-C) belong to the class of *preplanned, shared or dedicated, global protection* schemes.

General Resilient Overlays is a class of IPFRR methods that use different, non-shortest-path-based algorithmic techniques to compute the backup overlays. These methods will be called “general” to stress that the backup can be an arbitrary topology, differentiating this class from the subsequent class where the backups will be strictly limited to tree topologies.

The authors in [207] trace back general fault isolation to FRoots, which considered the problem in the context of high-performance routing [208]. The idea in that paper, and in essentially all extensions [166], [167], [169], [207], [209]–[213], is the same: find a set of overlay topologies so that each router will have at least one overlay it can use as a backup whenever its primary next-hop fails. The difficulty in this problem is to consider arbitrary, possibly highly irregular network topologies and/or link costs [208], and to minimize the number of the resultant overlays to cut down control, management, and data-plane burden.

Key to all IPFRR methods described in this class is the notion of *failure isolation* and *backup overlays*. We say that a topology (e.g., an overlay) isolates a failure when the topology, and the routes provisioned on top of the topology, have the property that there is a certain set of source-destination pairs (the “covered pairs”) and failure cases (the “isolated failures”) so that if we inject a packet into the topology at a covered source towards a covered destination, the packet will avoid the isolated failure and flow undisrupted to the required destination. The set of backup overlays then forms a family of overlay topologies, with the property that for any pair of a source and a destination router and any failure there is at least one backup overlay that isolates the failure and covers the source-destination pair. Then, if the source selects the proper backup overlay for the given failure case(s) and destination node and it has a way to “pin” the packet into that overlay, then it is ensured that the packet will reach the destination node avoiding the failure(s).

Provided that a way exists to pin packets into an overlay, this family of IPFRR mechanisms does not require major modification to the IP control plane and data plane and hence is readily deployable, although not in an incremental way. Also, fault isolation may yield 100% failure case coverage, even for multi-failures [210]. On the other hand, depending on the network topology and link costs the number of backup topologies may be large (in the order of dozens), raising a

substantial management issue. This is especially troublesome when the task is to protect against multiple failures [211]. Additionally, pinning packets to the overlay may not be simple, as the IP header space is a valuable yet scarce resource and using some bits for tagging the backup overlay may interfere with other uses of these bits (e.g., tagging in the ToS field clashes with the use of the same field in DiffServ). Also, some non-standard technique is needed by routers to inject packets into the correct overlay (but see VRO below). Finally, finding the fewest possible backup overlays is usually a computationally intractable problem, which calls for approximations that may deteriorate efficiency.

Starting from FRoots [208], there are many realizations of the basic general fault-isolation technique. Below, we survey the most representative ones. The earliest proposal (following FRoots) seems to be Resilient Routing Layers [166], where the backup overlays are obtained by explicitly removing certain links from the overlay topologies. If the overlays are such that each link is isolated by at least one overlay, then full protection against single-link failures can be attained. Multiple Routing Configurations (MRC) and the related line of works [167], [207], [209], [210] took another approach: instead of effectively removing links from the overlays they reach isolation by setting link costs separately in each overlay so that shortest paths will avoid a certain set of links, effectively isolating these failure cases. This method is easy to adopt for link-, node-, and multiple-failure protection scenarios [210]. Perhaps the most viable implementation avenue was outlined in [214] (but see also [211], [215]): the use of multi-topology routing. Multi-topology routing is a feature that is to become available in IGP [216] to maintain multiple independent overlay routing topologies on top of the default shortest paths, with each topology using an independent set of link costs. One routing topology is devoted to support IP forwarding under normal networking conditions, and when a next-hop of some router becomes unavailable it simply deviates all traffic that would have been sent via that next-hop to another interface, over an alternative routing topology. With a clever optimization of the link costs per topology, full failure case protection against link and node failures can be achieved [211], [214], [215]. The Independent Directed Acyclic Graphs proposal [212] uses a set of DAGs for similar purposes.

Most recently, *Virtual Routing Overlays* (VRO [169], [213]) solves two problems inherent in failure isolation, namely, the lack of standards for injecting packets into overlays and for pinning a packet into the an overlay until it reaches the destination. VRO uses LFA and virtual routers for this purpose: it provisions multiple virtual routers on top of each physical router, each instance with its own set of local link costs, so that whenever a packet encounters a failure it will automatically fall back to a virtual router that is provisioned as an LFA and will flow within the virtual overlay until it reaches the destination. The authors in [213] are able to show that per each physical router at most 4 virtual router instances in the worst case, and at most 2 instances in the average case, suffice to reach 100% failure case coverage.

Independent/redundant trees are a class of IPFRR methods based on the fault isolation technique that restrict the backup

overlays to be trees [145]–[147], [199], [204], [217]–[220].

The idea of searching for backups in the form of simple tree topologies traces back to redundant trees, red-blue trees, and ST-numberings [194], [199]. Consider the example of *redundant trees*, for instance: here, [199] shows a polynomial time algorithm to find a pair of mutually link-disjoint directed trees towards any node (the “red” and the “blue” tree) and claims that such pair of paths can be found in each 2-connected graph. Adopting this idea to IPFRR is then straightforward: whenever the primary next-hop is on the red tree and this next-hop fails then we pin packets to the blue tree (and *vice versa*): the art is then to implement this scheme given the limitations of the IP data plane [193]. *Independent trees*, in contrast, are generally undirected, serving as a backup to potentially multiple destination nodes [219], but in general they may be more difficult to compute [221].

Tree topologies have the property that there is at most one path between any pair of nodes in the tree. Searching the backup overlay in the form of a tree therefore solves several issues in one turn: it prunes the search space significantly, yielding that in most cases the backup trees can be found in polynomial time (recall, the general question is usually NP-complete) and makes routing on top of the overlay trivial. The price is, of course, an increased number of backup overlays needed to be provisioned: while the number of general overlays grows very slowly with the number of routers (close to logarithmically, at least by experience [207]), redundant trees are per-node and hence the number of overlays scales linearly with the number of routers in general [145].

The pros are that independent trees are easy to calculate, can be generalized for multiple failures, and forwarding along the tree-shaped overlay is trivial. On the negative side, we may still need to pin packets to the correct overlay some way, which may cause compatibility issues. Furthermore, circular dependencies between backup trees may make it difficult to find the correct backup in case of multiple failures [194].

Thanks to the simplicity of this approach, there are many versions, extensions, and even standards that use independent/redundant trees for IPFRR [145], [146], [171], [192]–[194], [204], [217]–[220]. The underlying mathematical notions, revolving around red-blue trees, ST-numberings, and independent trees, and the conditions under which they exist, are generally well-understood [171], [194], [221]. Proposals then differ based on whether they use directed (redundant) trees or independent (undirected) trees, and the different ways they implement the resultant overlays.

The earliest proposal to use redundant trees for IPFRR is Lightweight Not-via [192], [193], whereby packets are pinned into the proper (red or blue) overlay tree using not-via addresses. After multiple modifications, this proposal, and the related concept of maximally redundant trees for non-2-connected topologies, reached RFC status and became the MRT-FRR standard [145], [146], [217]. Recently, [171] proposes various heuristics to cut down the length of the resultant paths and [222] compares the performance of MRT-FRR to that of not-via addresses. Similar ideas exist to reach the same goals using undirected independent trees for handling multi-failures [204], [218]–[220].

D. Summary

In this section, we reviewed the most important ideas, proposals, and standards for providing fast network-layer failure recovery in the intra-domain setting. We reviewed IPFRR, the umbrella framework for fast local IP-layer protection, in Section V-A. Then, we discussed the prevalent schemes in this context, namely, the methods that reuse the IGP’s link state database maintained and shortest-path routing for failure recovery (Section V-B) and the schemes that revolve around virtual overlay networks on top of the default forwarding topology for this purpose (Section V-C).

In general, fast IP failure recovery in the intra-domain unicast setting is a well-understood problem area, with several well-established standards [32], [124], [141], [145], [146], [160], [170], [182], [217], inter-operable implementations in off-the-shelf routers [173]–[175], [223], operational experience [172], [195], and extensive research literature [165], [176] available. Open problems remain, however: even after several iterations it is still not entirely clear how to reliably avoid intermittent micro-loops without slowing down re-convergence [182] and provide the required quality of service even during routing transients [143], [144].

VI. INTER-DOMAIN NETWORK-LAYER FAST RECOVERY

Realizing FRR in the inter-domain routing, *i.e.*, across network domains administered by independent entities, entails solving an additional set of challenges compared to intra-domain FRR. The Border Gateway Protocol (BGP) is today’s de-facto inter-domain routing protocol that dictates how network organizations, henceforth referred to as *Autonomous Systems* (ASes), must exchange routing information in the Internet to establish connectivity. Intuitively, since networks do not have control and/or visibility into other network domains, the proposed FRR techniques often achieve much lower resiliency to failures than intra-domain mechanisms. Moreover, inter-domain FRR often requires some sort of coordination among network domains in order to guarantee connectivity, which has been an obstacle for most of the proposed schemes, preventing them from being standardized. In this section, we discuss challenges around the problem of realizing FRR in the wider Internet (*i.e.*, in BGP) and we describe different approaches proposed to tackle this problem. We present a timeline with a selection of work on the topic in Fig. 21, starting from the initial work in the early 2000s exposing the problem of slow BGP convergence, and then covering the sequence of work on improvements for the detection of failures and fast restoration of connectivity in the Internet.

A. Background on BGP Route Computation

BGP is a per-destination policy-based path-vector routing protocol. The owner of an IP prefix π configures its BGP routers to announce a route destined to π to (a subset of) its neighbors according to its routing policies. A BGP route contains the destination IP prefix destination, the sequence of traversed ASes (*i.e.*, the `AS_PATH`) and some additional information that is not relevant for this section. When an AS receives a new BGP route, it extracts the destination IP

Fig. 21: Timeline of the selected documents and solutions related to inter-domain network-layer fast recovery (entries marked in gray provide the general context).

prefix destination π and checks whether the currently selected route is worse than the new one based on its routing ranking policies. In this case, it updates its new best route towards π and announces it to (a subset of) its neighbors according to its routing policies. In the case of link/node failure, reconvergence is triggered by the node adjacent to the failure, which either selects an alternative best route towards the destination or withdraws the existing route by informing its neighbors.

B. Challenges

Realizing Internet-wide FRR requires taking into consideration three unique aspects of the inter-domain routing setting. First, the Internet is not administered by a single organization, but it rather consists of thousands of independent interconnected network organizations with possibly conflicting goals. Second, BGP convergence is significantly slower than intra-domain routing protocols (*i.e.*, on the order of minutes compared to hundreds of milliseconds/seconds for intra-domain protocols) [224], [229]–[233]. Finally, Internet routing protocols must guarantee connectivity towards hundreds of thousands of destination IP prefixes. There are four main consequences related to the design of FRR mechanisms that we can draw from the above considerations. First, failures of inter-domain links can lead to long and severe disruptions, as the BGP control-plane may take minutes to reconverge. FRR at the inter-domain level is therefore critical for restoring Internet traffic. As the number of prefixes can be on the order of tens of thousands, being able to update the forwarding tables quickly is of paramount importance. For example, the reported results of previous measurements indicate that it can take several seconds to update tens of thousands of prefixes [229]. Second, one cannot assume that all ASes will deploy FRR mechanisms. This makes clean-slate FRR approaches very difficult to realize in practice. Third, a network operator has no visibility of the entire Internet network. Some ASes may end up (mistakenly or on purpose) detouring their traffic away

from their announced BGP routes without communicating this information to their BGP neighbors [234]. Fourth, BGP does not carry any root-cause analysis information within its messages (*i.e.*, no explicit link failure notification is available), since all networks must jointly agree on the format of BGP messages, making it hard to be modified. Fifth, inter-domain routing messages are disseminated throughout a domain on top of intra-domain routing protocols. Such an overlay makes the prompt restoration of inter-domain connectivity even more complex [228].

Given the above challenges, we first describe two critical problems concerning the support for FRR on the Internet that go beyond the computation of backup routes: detection of possible remote failures and quick updates of the forwarding planes. We then discuss a set of improvements to BGP that would allow network operators to precompute backup Internet BGP routes.

C. Detection of Local and Remote Failures

Detecting failures in the inter-domain routing entails solving different challenges than in intra-domain routing. First, to implement FRR in BGP, beyond detecting local failures, an AS needs to detect and localize remote Internet failures, as *i)* downstream ASes may not perform any FRR and *ii)* BGP is slow at recomputing a new route, thus leading to traffic disruptions lasting even several minutes. Failures of adjacent peering links between two ASes can be detected using traditional techniques (*e.g.*, BFD [71]) and are not discussed here. Detection of remote link failures can be performed both at the control plane and at the data plane, and we discuss these two different approaches in the following parts of this section.

At the control-plane level, a variety of techniques to detect and localize network failures using BGP data have been proposed. Some techniques require coordination from a geographically-distributed set of trusted vantage points, *e.g.*, PoiRoot [225], [235], [236]. They are typically general enough to detect a variety of anomalies, including link failures and routing instabilities, and root-cause analysis. On the other side of the design space, we have techniques that do not require any coordination among ASes and simply attempt to infer remote link failures from the received bursts of BGP messages describing the route changes/withdrawals at one single location, *e.g.*, SWIFT [229]. SWIFT performs the root-cause analysis on a sudden burst of BGP updates using a simple intuition: it checks whether a certain peering link has suddenly disappeared from the `AS_PATH` of a certain number of BGP announcement messages received close to each other. Internally, SWIFT relies on a metric called *Fit Score* and infers the failed link as the one maximizing its value.

Consider the example shown in Fig. 22 and Fig. 23 inspired by the original paper, where we have seven AS nodes numbered from 1 to 7. AS 3, AS 5, and AS 7 originate ten thousand distinct IP prefixes each, marked in orange, green, and blue, respectively. The converged state of BGP is shown in Fig. 22 (orange, green, and blue arrows represent the BGP paths associated with the corresponding IP prefixes). We assume that AS 2 does not announce any green route

Fig. 22: An illustration of SWIFT: the pre-failure state.

towards AS 5 to its neighbor AS 3. Let us now consider AS 1 which is sending traffic towards AS 7 along path (1 3 5 7). When link (3, 5) fails (see the post-convergence state in Fig. 23), AS 1 starts receiving a sequence of at least twenty thousand updates for all the routes towards AS 5 and AS 7 that traverse link (3, 5). AS 1 therefore infers that one of the peering links from among (1, 3), (3, 5), and (5, 7) must have failed. To determine which link is down, AS 1 waits a little bit to see which paths are initially being modified by BGP. Since AS 1 does not see any path change for the IP prefixes destined to AS 3, it quickly infers that link (1, 3) remains fully operational. Moreover, since AS 1 will soon receive a new route from AS 3 towards AS 7 via link (5, 7), it infers that the failed link must be (3, 5)⁵. After identifying this link, SWIFT quickly reconfigures all the routes traversing link (3, 5), so that they follow a backup route (see the following subsection). We note that the inference algorithm used in SWIFT is more sophisticated than the simplified version that we have just presented here. In particular, for each link l at time t , SWIFT computes the value of the *Fit Score* metric which tracks *i)* the number of prefixes whose paths traverse link l and have been withdrawn at time t , divided by the total number of withdrawal messages received until t , and *ii)* the number of prefixes whose paths traverse link l and have been withdrawn at time t , divided by the total number of prefixes whose paths included link l at time t . Based on the *Fit Score* metric, SWIFT is able to infer the status of a remote link.

One definite advantage of fast recovery mechanisms inferring the status of the network from control-plane messages is that they significantly speed up the recovery upon a failure. At the same time, one of the main disadvantages is that tuning the parameters of these mechanisms for each and every network is often difficult. In particular, general and comprehensively-validated guidelines are still missing. Moreover, deployment of SWIFT-like mechanisms in independent network domains may lead to severe transient forwarding loops due to the uncoordinated nature of SWIFT in situations when multiple failures arise⁶. Finally, control-plane approaches are inherently limited by the slow reconvergence of BGP.

At the data-plane level, detection of remote failures can be

⁵Note that the link failure inference algorithm implemented in SWIFT runs on a per-BGP-session basis and it does not combine BGP updates received from different sessions.

⁶Note that SWIFT has some guaranteed properties — see the original paper for a detailed discussion [229].

Fig. 23: An illustration of SWIFT: the post-failure state.

performed by *i)* monitoring explicit application performance-related information (e.g., similarly to Google Espresso [237]), *ii)* sending active probes (similarly to BFD) [236] to verify the route followed by a packet and to check whether the route is valid, or *iii)* passively inspecting the transported traffic to detect if a subset of the flows are retransmitting a non-negligible number of packets, as described in the Blink system [230]. In particular, Blink uses properties of the TCP implementation/TCP stack which retransmits a non-ACK'ed packet exactly after 200 ms.⁷

One advantage of data-plane approaches is the speed of the failure detection process, as traffic travels order of magnitudes faster than control-plane messages. However, one clear disadvantage is that identifying the location of a failure is a much harder problem, which requires active probing (e.g., traceroutes) and cannot be inferred solely from data traffic.

D. Updating the Forwarding Tables

Single link failures between two ASes may affect hundreds of thousand of destination IP prefixes. Updating the forwarding plane to reroute all these prefixes is therefore a critical operation to achieve fast restoration. For example, assuming 100 μs to update a forwarding entry [238], it takes roughly 10 seconds to update a forwarding table containing 100 thousand prefixes.

Consider the example shown in Fig. 24a in which we depict the forwarding table of a BGP router located in AS 1. The router has installed forwarding rules for the ten thousand IP prefixes originated by AS 3 (the first 10K prefixes), AS 5 (the second 10K prefixes), and AS 7 (the third 10K prefixes). For the sake of simplicity, we denote BGP next hops using AS identifiers, e.g., the BGP next hops of AS 1 are AS 2, AS 3, or AS 4.

When a packet must be forwarded from AS 1, a lookup is performed to extract the next hop towards which the packet is forwarded. This data-plane design clearly has the benefit of achieving low packet processing time (one lookup) and low memory occupancy (one entry per prefix) and has long been used as the reference data-plane architecture by the main router vendors such as Cisco [239], [240] in the early 2000s. Yet, this approach is problematic upon peering link failures/flapping. When the link (1, 3) fails, the router must rewrite the first, third, and fourth 10K thousand forwarding entries in its table to steer traffic away from AS 3. We assume that traffic

⁷Common value used in most implementations on Linux systems.

Fig. 24: Fast data-plane convergence: (a) classical data-plane architecture and (b) the equivalent simplified BGP-PIC data-plane architecture. We indicate updates of single forwarding entries with arrows depicted inside the field of the forwarding entry.

destined to AS 3 will be rerouted through AS 4 while the remaining traffic destined to AS 5 and AS 7 will be rerouted through AS 2. Rewriting 30K entries is an operation that takes non-negligible time and during which traffic will be dropped. The update time in this case grows linearly in the number of prefixes.

This problem is less severe in intra-domain routing where the number of destinations is on the order of few hundred.⁸ Solutions to speed-up the update of the forwarding tables upon a failure have been proposed and we divide them into those targeting failures of adjacent (local) AS-to-AS links and remote ones.

BGP-PIC [28], [238]. In the case of a local peering link failure, one can associate in the data plane a Virtual Next Hop (VNH) with each destination IP prefix, and the VNH with the actual BGP next hop. This VNH simply acts as an *indirection* between the IP destination and the real next hop. Two IP prefixes are associated with the same VNH if and only if they share the same primary and backup AS next hops. Intuitively, one can update many prefixes associated with the same VNH by simply updating the assignment of the VNH to the real AS next hop. This technique has been presented in 2005 [238] and has been incorporated in Cisco devices with the name BGP PIC [28] years later. We now provide a simplified example.

Consider the example shown in Fig. 24b. For each prefix, the control plane at AS 1 computes both the best and backup BGP best routes (and BGP next hops). It then groups IP prefixes that have the same forwarding primary and backup routes in Virtual NextHops (VNHs). In our example, we have just two VNHs: VNH_1 which contains prefixes $P_1, \dots, P_{10K}$ whose primary route is through AS 3 and backup route through AS 4, VNH_2 which contains prefixes $P_{10K+1}, \dots, P_{30K}$ whose primary route is through AS 3 and backup route through AS 2. When a packet must be forwarded at AS 1, the first lookup determines the VNH of the packet, while the second lookup determines the actual BGP next hop. The main benefit of this approach is that the time needed to update the forwarding table upon the failure of link (1,3) is greatly reduced. In fact, as soon as AS 1 learns that link (1,3) is

down, it simply updates the next-hop of VNH_1 to AS 4 and it updates the next-hop of VNH_2 to AS 2: two single updates of the forwarding table. In the worse case, in a large network with N BGP border routers, this approach may require $n - 1$ data-plane updates, which may be performed in roughly 1 ms when $n = 100$.

SWIFT [229]. In the case of a remote link failure, one can generalize the above approach, *i.e.*, create a different VNH for each different AS_PATH. Upon identification of a failure, the control plane can update the VNH mapping based on the computed backup information and only update the VNH traversing the failed link. The number of forwarding table modifications is linear in the number of VNH traversing the failed links, possibly a large number. A different approach has been presented in SWIFT, whose remote failure detection mechanism has already been discussed in Section VI-C. We present a simplified description below, focusing only on the essential insights related to SWIFT.

The control plane associates a backup AS next hop with each destination IP prefix and an inter-AS link traversed on the way towards the destination. The alternative next hop can be quickly activated upon detection of a remote link failure using a carefully designed packet processing pipeline like the one described in the SWIFT system [229] and summarized below.

Fig. 25 shows the forwarding table based on the example from Fig. 22. The SWIFT forwarding pipeline consists of two tables, T_0 and T_1 , arranged in a sequence. Table T_0 maps each IP address to a sequence of actions that fetch information about the traversed ASes as well as the primary and backup links for the associated IP address. In the example, we only protect against failures of the (local) first AS link and the (remote) second AS link. One can protect against a larger number of links by adding more remote links. For each link, we store the backup AS next hop. The first table attaches all the backup options to the packet header. For instance, a packet destined to IP_{10K+1} will have AS 2 as a backup if link (1,3) or link (3,5) fails. Table T_1 is used to determine the correct next-hop AS by verifying whether an entry exists in T_1 that matches any of the backup information attached to the packet. More specifically, when a link fails, the control plane installs two entries in Table T_1 , both matching the failed link and instructing the device to forward the packet to the

⁸Hierarchical intra-domain routing is a common solution to scale both the routing computation and the time needed to update the forwarding state. In contrast, the only hierarchical optimization in BGP is the aggregation of the IP addresses to compute BGP routes and the usage of Longest-Prefix-Matching techniques to forward packets at the data-plane level.

table T_0					
ip_dst	prm_nh	link ₁	1st_bkp	link ₂	2nd_bkp
IP_1	3	(1, 3)	4	n.a.	n.a.
...	...	...	...	...	...
IP_{10K}	3	(1, 3)	4	n.a.	n.a.
IP_{10K+1}	3	(1, 3)	2	(3, 5)	2
...	...	...	...	...	...
IP_{20K}	3	(1, 3)	2	(3, 5)	2
IP_{20K+1}	3	(1, 3)	2	(3, 5)	2
...	...	...	...	...	...
IP_{30K}	3	(1, 3)	2	(3, 5)	2

→

table T_1		
link ₁	link ₂	action
(3, 5)	*	fwd to 1st_bkp
*	(3, 5)	fwd to 2nd_bkp
*	*	fwd to prm_nh

Fig. 25: Simplified SWIFT data-plane forwarding table at AS 1.

corresponding backup AS. For instance, if link (3, 5) fails, the control plane will install two rules that match all packets traversing link (3, 5) at either the first or the second hop, so that the matching packets can be forwarded towards the corresponding backup AS next hop. The implementation of the action in Table T_1 depends on the expressiveness of the data plane. For instance, in OpenFlow, one needs to add rules matching each possible backup next hop to the forwarding action towards that backup next hop. In P4, however, this action could be expressed in a simpler way by extracting the backup next hop from the metadata and finding the egress port based on this next hop.

One advantage of SWIFT is that it requires few updates to the forwarding table to detour the forwarded traffic to a different next hop. More specifically, the number of updates to the forwarding table is linear in the length of the longest path that a network operator wants to protect, *e.g.*, two updates to protect the first and second AS links for all the IP prefixes. At the same time, updating the backup information may still require a large number of updates (these updates do not affect the forwarded traffic, though), thus increasing the overhead on the operating system of the switch. Another limitation of SWIFT is that it only provides a fast reroute alternative for single-link failures assuming there exists an already available BGP route that is not traversing the failed link. Computing BGP routes that are robust to any link failure is beyond the scope of SWIFT and will be discussed in the next section.

E. Fast-Reroute Mechanisms

BGP is a per-destination routing protocol and, as such, comes with limitations similar to IP FRR techniques discussed in Sect. V. Specifically, for some network topologies, it is impossible to find per-destination FRR mechanisms that are robust to even single link failures. Consider the forwarding state shown in Fig. 26 where the green solid arrows represent the forwarding before link (3, 5) fails while the red dashed arrows represent the forwarding at AS 1 and AS 3 after the link has failed. Note that we assume AS 2 does not announce any route towards AS 3. This means that, if the link between AS 3 and AS 5 fails, AS 3 is left with no safe neighbor to whom it could forward its traffic (a forwarding loop would be created otherwise by forwarding traffic back to AS 1).

Inter-domain FRR mechanisms, such as Blink [230], SWIFT [229], and Google Espresso [237], simply reroute traffic along any of the available BGP routes. While these mechanisms cannot guarantee robustness to even single link failures, these approaches are legacy-compatible with BGP

Fig. 26: R-BGP: AS 1 advertises a backup route to AS 3, offering to transit its backup traffic to AS 5 through AS 2.

and, at the same time, provide some minimal degree of robustness. R-BGP [227] is an enhancement of BGP that allows ASes to compute backup paths between ASes in a distributed manner. In R-BGP, ASes announce one additional backup route towards their downstream neighbors, which, in turn, can use this backup route in the event of a downstream failure.

Consider the example in Fig. 26. AS 1 now advertises a backup route to AS 3, offering to transit its backup traffic through AS 2 for all the prefixes announced from AS 5. When link (3, 5) fails, AS 3 simply sends all its traffic to AS 1, which in turn detects from the fact that the traffic is received from an “outgoing” direction that the traffic should be forwarded along the backup path. It therefore immediately sends traffic to AS 2. This mechanism, described in detail in the original paper, guarantees connectivity for any single link/node failure at the inter-domain level. The authors of the paper describe a sequence of optimizations to reduce the amount of additional information piggy-backed on BGP and how to handle spurious withdraw messages (see Section X for more information). Special care during the BGP reconvergence process must be taken into account by incorporating information about the root-cause of a failure. The main disadvantage of R-BGP is that it requires modifications to the BGP protocol that must be adopted by a large number of networks, a cumbersome operation in practice.

F. Summary

In this section, we reviewed the most prominent techniques to deal with failures at the Internet level. We first discussed the main differences in dealing with inter-domain fast reroute compared to intra-domain: (1) the lack of control and visibility into the entire Internet topology and (2) the large-scale amount of destination IP prefixes to be handled. We therefore

discussed the main approaches to *quickly* detect failures at the inter-domain levels by inferring such failures from both BGP-based control-plane signals as well as TCP-based data-plane ones [229], [230]. We then discussed techniques to quickly update the forwarding plane when thousands of forwarding rules have to be modified in response to a link failure. Such techniques either rely on (1) what we referred as “indirection” tables for mapping virtual identifiers to a specific forwarding actions [28], [238] or (2) labeling destination IP prefixes with their explicit path and matching that path with the failed link to compute the backup next-hop [229]. We concluded the section by discussing the currently available yet simple FRR mechanisms in BGP, *i.e.*, reroute on any alternative existing path, as well as mechanisms that require modifications to BGP but achieves guaranteed resiliency for every single link failure [227].

Many open research problems still require to be addressed: despite decades of academic and industrial efforts in improving the resiliency of BGP, the current status-quo is still quite alarming. Substantial efforts must be targeted to the problem of detecting failures at the inter-domain level as both Swift [229] and Blink [230] suffer from false positives and false negatives. Moreover, data-plane approaches such as Blink are highly dependent on the specific configuration of RTO timeouts, which makes it both inaccurate when different congestion control mechanisms will be deployed, but also vulnerable to malicious attacks. Finally, supporting backup paths in BGP seems to be a non-trivial challenge because of the inherent needs to preserve legacy-compatibility, privacy, and the sheer size of routing information currently exchanged on the Internet to glue almost hundreds of thousands of networks and almost one million IP prefixes.

VII. FAST RECOVERY IN PROGRAMMABLE NETWORKS

In this section, we discuss advanced fast-rerouting mechanisms as they are enabled by emerging programmable networks. Indeed, a simpler and faster failure handling was one of the reasons behind Google’s move to Software-Defined Networks (SDNs) [241]. In an SDN, the control over the network devices (e.g., OpenFlow switches) is outsourced and consolidated to a logically-centralized controller. This decoupling of the control plane from the data plane allows to evolve and innovate the former one independently of the constraints and lifecycles of the latter one.

However, it also introduces new challenges. If a link failure occurs, it needs not only be detected but also communicated to a controller which then reconfigures affected paths. This indirection does not only introduce delays, but if based on in-band signaling, the network elements and the controller may even be disconnected due to the failure. For example, controller reaction times on the order of 100 ms have been measured in [242]: the restoration time also depends on the number of flows to be restored, path lengths, traffic bursts in the control network, and may take even longer for larger networks.

Failover in the data plane is hence an even more attractive alternative in the context of SDNs. Local fast-reroute allows

Fig. 27: Illustration of OpenFlow group table: fast-failover (ff) group table for $s1$.

an SDN switch (or “point of local repair”) to locally detect a failure and deviate affected traffic so that it eventually reaches its destination. In the following, we first discuss solutions based on OpenFlow [18], the de-facto standard of SDN. Subsequently, we discuss solutions for programmable dataplanes, such as P4.

A. OpenFlow

OpenFlow supports basic primitives to implement fast failover functions. In particular, OpenFlow 1.1.0 provided a fast-failover action which was not available before: it incorporates a fast failover mechanism based on so-called *groups* allowing to define more complex operations on packets that cannot be defined within a flow alone [243]. And in particular, to predefine resilient and in-band failover routes which *activate* upon a topological change. Before the introduction of such FRR primitives, researchers relied on ad-hoc non-standardized extensions to OpenFlow, *e.g.*, flow priorities, timers, automatic deletion of rules forwarding on failed interfaces, to implement FRR primitives [244].

Figure 27 illustrates the OpenFlow model and its FRR mechanism: a controller (e.g., Floodlight) can populate the forwarding tables of the different switches from a logically centralized perspective. To support fast-failover to alternative links without control plane interaction, OpenFlow uses group tables: the controller can pre-define groups of the type fast-failover (FF in the figure, the group table of $s1$ is shown). A group contains separate lists of actions, referred to as *buckets*. The fast-failover group is used to detect and overcome port failures: each bucket has a watch port and/or watch group to watch the status of the indicated port/group. The bucket will not be used if the element is down. That is, a bucket in use will not be changed unless the liveness of the currently used bucket’s watch port/group is updated. In this case, the group will quickly select the next bucket in the bucket list with a watch port/group that is up. The failover time here hence depends on the time to find a watch port/group that is up.

In principle, most of the FRR mechanisms discussed earlier in this paper, e.g., for MPLS or IP networks, could be ported to SDN. However, without additional considerations, this approach could lead to an unnecessarily large number of flow table entries or to suboptimal protection only [253]. The first is problematic, as the flow table size of OpenFlow switches are

Fig. 28: Timeline of the selected documents and Fast-Reroute solutions related to Software-Defined Networks (entries marked in gray provide additional context related to the evolution of OpenFlow).

often small so that only a few additional forwarding entries can be accommodated for protection purposes. The second is not acceptable for SDN, because unprotected flows remain disconnected or FRR-caused loops persist until the controller comes to rescue.

Several solutions based on OpenFlow have been proposed in the literature so far, e.g., based on ideas from LFA [184], [255] or MPLS [245], by encoding primary and backup paths in the packet header [246], or by extending to OpenFlow's fast failover action based on additional state in the OpenFlow pipeline: for example, *SPIDER* [251] leverages packet labels to carry reroute and connectivity information.

Kempf et al. [245] propose an alternative OpenFlow switch design that allows integrated operations, administration and management (OAM) execution, foremost connectivity monitoring, in MPLS networks by introducing logical group ports. However, the introduced logical group ports remain unstandardized and are hence not part of shipped OpenFlow switches.

Borokhovich et al. [249] showed that it is possible to implement a fast failover mechanism in OpenFlow which always finds a route as long as the underlying network is physically connected, assuming that packet headers can be used to carry information. That is, in this approach, the packet carries a counter for each switch, allowing the packet to explore different paths during its traversal through the network. The authors show that the traversal of the graph can be performed in a variety of ways including Depth-

Fig. 29: Depth-first search example.

First Search (DFS) and Breath-First Search (BFS). This is hence reminiscent of the Failure-carrying Packets (FCP) [152] approach, used for convergence-free routing (for a detailed discussion of this technique, see Section VIII).

We now describe in more details the DFS traversal with the help of an example. We consider the network shown in Fig. 29 with 5 nodes *a*, *b*, *c*, *d*, and *e*, where node *d* is the destination node and *a* is the sender of a packet. The DFS mechanism requires to compute an ordering of all the ports at each switch as shown in the top-left network of Fig. 29. The ordering is used to drive the DFS exploration by trying each outgoing port in the given order until the destination is reached. The first port in the ordering is the default port that is used to send packets in the absence of failures and the set of first ports form a tree rooted at the destination node. The DFS approach requires to store a non-negligible amount of information in the packets header. Namely, each packet header contains the *currently* explored neighbor at a node (called *cur*) and the *parent* (called *par*) from which a packet has been first received, which are shown below each network in Fig. 29. Initially, all the current and parent nodes are not initialized, i.e., set to zero. When a node must send a packet, it follows the port ordering starting from the *cur* index of that node. a node always skips a port in the ordering if the port has failed or it is the parent port. If a node has tried all the ports, only then it sends the packet back to the parent node. This way of forwarding packets is reminiscent of a DFS traversal, which first explores the children nodes and only then it performs backtracking to the parent node. In Fig.29, node *a* initializes *cur* to 1, which however maps to a failed port. It therefore increments *cur* by 1, i.e., it forwards the packet to the next port in the ordering, which is node *b* (top-center network in the figure). When node *b* receives the packet, it remembers that node *a* is the parent by setting port 1 as the parent. It then increases *cur* from 0 to 1, which however is exactly the parent port so it skips it. The next port leads to node *c*, which in turn sets node *b* (port 2) as the parent port and forwards the packet on port 1 back to node *a* - a potential forwarding loop! When node *a* receives the packet, it first sets node *c* (port 3) as the parent node. It then retrieves from the packet header the last explored port at node *a*, $cur[a] = 2$,

and increments it to 3, which however is now leading to its parent node so it skips it. Node a increments again by 1 in the circular ordering, which leads again to the failed port. Node a therefore increments again the current counter and forwards the packet again to node b (bottom-center network in the figure). Node b retrieves from the packet header the last explored port, $\text{cur}['b'] = 2$, and increments it by 1 to 3. This port now leads to node e , which can use its first port to reach the destination.

One disadvantage of the DFS technique (but also for the other graph-exploration techniques) is the huge packet header overhead, where the packet must remember the state of the currently explored node and parent node for each node in the network. We will see in the next subsection how more programmable paradigms allow to dramatically reduce this overhead without sacrificing the level of resilience achievable by FRR techniques.

We now conclude the OpenFlow subsection by briefly discussing some additional challenges in detecting failures and devising languages for building robust network configurations.

Van Adrichem et al. [250] argued that one of the key limitations of achieving a high availability and fast failure detection (below 50 ms) in OpenFlow implementations is that these networks often rely on Ethernet, which in turn relies on relatively infrequent heartbeats. A faster alternative is to use BFD [71], in addition to combining primary and backup paths configured by a central OpenFlow controller.

There also exists interesting work on languages for writing fault-tolerant network programs. Reitblatt et al. [256] developed FatTire, a language which is based a new programming construct that allows developers to specify the set of paths that packets may take through the network as well as the degree of fault tolerance required. Their compiler targets the OpenFlow fast-failover mechanism and facilitates simple reasoning about network programs even in the presence of failures.

B. Programmable Data Planes and P4

Lately, programmable data planes [248] emerged which further enrich the capabilities of networks by allowing to deploy customized packet processing algorithms. While several interesting use cases are currently discussed, *e.g.*, related to monitoring or traffic load-balancing, still little is known today about how to implement FRR mechanisms in such systems. In particular, the P4 programming language [248], one of the emerging languages for programming the data plane forwarding behaviour, does not provide built-in support for FRR.

An interesting approach, ahead of its time, is DDC (for Data-Driven Connectivity) by Liu et al. [16], which moves the responsibility for connectivity to the data plane: DDC maintains forwarding-connectivity via simple changes in forwarding state predicated only on the destination address and incoming port of an arriving packet, *i.e.*, DDC is a *dynamic* FRR mechanism that needs to modify the forwarding function at the speed of the data-plane. The required state and its corresponding modifications are simple enough to be amenable to data-plane implementations with revised hardware. Specifically, a DDC node only needs to store three bits of information

for each destination and update these based on the incoming packets towards that destination. At the high level, DDC aims at building a Directed Acyclic Graph (DAG) for each destination. When failures occur, the DAG is recomputed using information gathered from the data-plane packets based on the well-known *link-reversal* algorithm introduced by Gafni and Bertsekas [257] in their seminal paper. The advantage of the DDC paradigm is that it leaves the network functions that require global knowledge (such as optimizing routes, detecting disconnections, and distributing load) to be handled by the control plane, and moves connectivity maintenance, which has simple yet crucial semantics, to the data plane.

We show an example of the DDC link-reversal algorithm in Fig. 30. We consider the same network used for the graph-exploration FRR technique explained in Fig. 29 consisting of five nodes. We highlight with blue arrows the directed acyclic routing graph towards the top-left destination node d . When the link connecting the bottom-left node a to the destination fails, DDC triggers the link-reversal algorithm: if all the link directions are incoming, the direction of all the links are reversed and a packet is sent on any of these outgoing links. The bottom left node adjacent to the failure is the first node to reverse its link directions and forward the packet to its neighbor b . Consequently, now b has all its links in the “incoming” direction. It therefore reverses their directions and forwards the packet to node c . Node c also has all its links in the incoming direction so it reverses the link directions and forwards the packet back to node a - a potential loop. Node a reverses again all its incoming links and chooses b to forward its packet. Node b now does not have to reverse its links since the link towards node e is already in the outgoing direction. This breaks the potential forwarding loop as the packet is now forwarded to node e and consequently to its destination d . At this point, a new directed acyclic graph spanning all the nodes have been recomputed.

Compared to graph-exploration approaches, DDC achieves identical levels of resiliency, *i.e.*, guaranteed connectivity for any number of failures as long as a physical path exists, but does not incur the exorbitant packet overheads of graph-exploration techniques. DDC only requires to store at each node, for each destination and each port, the current direction of the link plus (1 bit) plus two additional auxiliary bits that are used to synchronize the direction of the link between two nodes in case of multiple link reversal operations.

DDC also describes how to implement some optimizations on the number of reversed links during the reconvergence process so as to speed up reconvergence. In the absence of failures & congestion and assuming an immediate detection of the failure, none of the packets towards a destination get dropped as long as a physical path exists.

Another approach is to draw from the insights on link layer mechanisms such as AXE [97], [100] (see the related discussion in Section III) which was designed to improve the performance of recovery in Ethernet-based networks. AXE takes a similar approach to DDC by moving the responsibility of maintaining connectivity at the data-plane level.

Plinko [247], [252] requires some additional features not available on OpenFlow switch hardware (*i.e.*, testing the

Fig. 30: DDC example.

status of a port in a TCAM match rule), but provides high resiliency: the only reason packets of any flow in a Plinko network will be dropped are congestion, packet corruption, and a partitioning of the network topology. Unlike prior work on data-plane resilience, Plinko takes a simple exhaustive approach: in the case of a failure, the switch local to the failure replaces the old route of a packet with a backup route, effectively bouncing the packet around in the network until it either reaches the destination or is dropped because no path exists. When implemented naively, this approach does not scale; however, Plinko can compress multiple forwarding rules into a single TCAM entry, which renders the approach more efficient, allowing Plinko to scale up to ten thousand hosts.

Most recently, in the context of P4, Chiesa et al. [254], [258] suggested an FRR primitive for programmable dataplanes based on P4. Their solution requires just one lookup in a TCAM, and hence outperforms naive implementations as it avoids packet recirculation. This in turn improves latency and throughput. This approach is essentially a “primitive” which can be used together with many existing FRR mechanisms (which is also demonstrated, e.g., for [259]), allowing them to benefit from avoiding packet recirculation.

C. Summary

This section provided an overview of the fast recovery mechanisms provided by programmable networks, and in particular, in OpenFlow and programmable dataplanes. We have discussed the basic principles underlying these mechanisms, such as group tables, and how the traditional fast failover approaches discussed in the earlier sections can be implemented as well in programmable networks, also pointing out limitations.

Fast recovery in programmable networks is the most recent application domain considered in this tutorial, and potentially the most powerful one, given the flexibilities provided in programmable networks. It is hence also likely the domain which still poses the most open research questions. In particular, we only have a very limited understanding of the algorithmic problems underlying the implementation of a highly resilient and efficient failover in programmable networks, e.g., how many failures can be tolerated and how short the resulting failover rules can be kept, e.g., depending on whether and

how packet headers can be changed. These algorithmic problems are further complicated by the specific technology and data structures that are used to implement recovery in programmable networks, e.g., related to recirculation overheads. Another important issue concerns the development of intuitive high-level concepts for the network programming language itself.

VIII. TECHNOLOGY-AGNOSTIC FAST RECOVERY METHODS

There exists a number of interesting FRR mechanisms which do not exactly fit any of the specific layers and technologies discussed above. In this section, we introduce the reader to some selected solutions which are based on fundamental techniques one should be familiar with.

A. Rerouting Along Arborescences

In order to achieve a very high degree of resilience, several previous works [154]–[156], [220] introduced an algorithmic approach based on the idea of covering the network with arc-disjoint directed *arborescences* rooted at the destination. Network decompositions into arc-disjoint arborescences can be computed in polynomial time, and enable a simple forwarding scheme: when encountering a link failure along one arborescence, a packet can simply follow an alternative arborescence. This approach comes in different flavors, such as deterministic [156] and randomized [155], or without [155] and with [156] packet header rewriting. In general, the approach requires *input port matching* to distinguish on which arborescence a packet is being routed. This however is practical, since many routers maintain a routing table at each line card of the interface for look-up efficiency. In the most simple case, a network can be decomposed into arc-disjoint Hamilton cycles. For example, Fig. 31 shows an example for the case of 2-dimensional torus graphs. In this case, the solution is known to achieve the maximum robustness: a network of degree k can tolerate up to $k - 1$ failures, without losing connectivity. A more general example is given in Fig. 32: here, arborescences rooted at the destination can be of higher degree. While such decompositions always exist, the challenge introduced in the general setting regards the order in which alternative arborescences are tried during failover. Interestingly, today, it is still an open question whether the maximum robustness can be achieved in the general case as well. However, Chiesa et al. showed in [154] that generally, the approach can tolerate at least half of the maximally possible link failures. They also showed that a resilience to $k - 1$ link failures can be tolerated using 3 bits in the packets header or creating $r - 1 < k$ copies of a packet, where r is the number of failed links hit by a packet. While these arborescence-based approaches provide a high resilience, one disadvantage of the approach is that it can lead to long failover paths (consider again the torus example in Fig. 31), introducing latency and load. The latter has been addressed in a line of work by Pignolet et al. [260], [261], e.g., relying on combinatorial designs or postprocessing of the arborescences [262]–[264].

Fig. 31: Decomposition of a 2-dimensional torus into arc-disjoint Hamilton cycles.

Fig. 32: Decomposition of a general graph into arc-disjoint arborescences.

B. Techniques Beyond Arborescences

The use of spanning trees, directed acyclic graphs, or arborescences for resilient routing, may come with the limitation that they only select links from such subgraphs. An interesting more general approach is *Keep Forwarding (KF)* [153]. KF is based on a “Partial Structural Network (PSN)”: the idea is to utilize *all* links and only determine link directions for a *subset* of links to improve resilience (somehow similar to DDC [16] which we discussed in Section VII-B). As such, KF can handle multiple failures with only small path stretch, and does not require packet labeling or state recording. To achieve these properties, KF needs to be “input-port aware”: forwarding depends not only on the destination but also on the input port. The authors also derive a *Graph Imperfectness Theorem*, showing that for an arbitrary graph if any component of it is “imperfect”, there will be no static rule-based routing guaranteeing the “perfect” resilience, even if the graph remains connected after failures. Nevertheless, the authors show in experiments that KF provides a high resilience.

A higher resilience can be achieved by methods which exploit header rewriting, such as Failure-Carrying Packets (FCP) [152]: in FCP, packets can autonomously discover a working path without requiring completely up-to-date state in routers. FCP takes advantage of the fact that typically a network topology does not undergo arbitrary changes, but there is a well-defined set of “potential links” that does not change very often: while the set of the potential links that are *actually functioning* at any particular time can fluctuate, e.g. depending on link failures and repairs, the set of *potential* links is governed by much slower processes (i.e., decommissioning a link, installing a link, negotiating a peering relationship). Thus, one can use fairly standard techniques to give all routers a consistent view of the potential set of links, which is called the Network Map. Motivated by this observation, FCP adopts a link-state approach in that every router has a consistent network map. Since all routers have the same network map, packets only need to carry information about which of these links have failed at the current instant. This “failure-carrying packets approach” ensures that when a packet arrives at a router, that router knows about any relevant failures on the packet’s previous path, and can hence determine the shortest remaining path to the destination (which may also be precomputed in the data plane). This eliminates the need for the routing protocol to immediately propagate failure information to all routers, yet allows packets to be routed around failed links in a consistent loop-free manner. There is also a Source-Routing FCP variant that provides similar properties even if the network maps are inconsistent, at the expense of additional overhead in packet headers. More concretely, in the source-routing variant, like in the basic FCP, a node adds the failed link to the packet header, but replaces the source route in the header with a newly computed route, if any exists, to the destination.

There exist several additional interesting technology-agnostic approaches. To just give one more example: An elegant solution to provide a certain degree of resilience is described in the O2 (“out-degree 2”) paper [151]: in order to reduce the outage time compared to traditional IP networks, O2 provides each node with at least two disjoint next hops towards any given destination, allowing each node to locally and quickly re-distribute the traffic to the remaining next hop(s) if a route fails. The resulting paths are loop-free but may increase link loads slightly, due to increased path lengths.

C. Topological Support for Fast Recovery

Fast rerouting can also be supported at the network topology level, which introduces an interesting additional dimension to the fast recovery space. Likely, a more redundant topology can tolerate more link failures; however, while this is obvious for centralized routing algorithms, it is not necessarily clear how to exploit redundancy in local fast rerouting algorithms. An interesting approach in this context is F10 [259]: a novel network topology reminiscent of traditional fat trees, with better fault recovery properties. In particular, by a clever re-wiring, F10 topologies support the design of failover protocols which can reestablish connectivity and load balancing more generally and quickly.

Fig. 33: Timeline of the selected documents and solutions related to technology-agnostic fast-recovery methods.

D. Summary

We discussed several fast recovery methods which are independent of technologies, from failure-carrying packets over arborescence-based approaches (which do not require packet header modifications), to the design of alternative topologies which explicitly account for the fast failover performance.

There are several interesting open research issues. In particular, the question of how much information needs to be carried in packets to achieve certain resilience guarantees remains an open problem. For example, it is not clear whether the capability of input port matching is sufficient to guarantee connectivity in a k -connected graph under up to $k - 1$ failures, using arborescence-based approaches or in general. Another interesting open problem regards the design of network topologies for efficient fast rerouting beyond fat-trees.

IX. CLASSIFICATION SUMMARY

We now take a step back and identify and classify the principles underlying the specific approaches discussed above for different technologies.

The most simple form of fast recovery is based on *static* failover routing: the forwarding behavior is statically predefined, e.g., using conditional per-node failover rules which only apply in case of failures, and cannot be changed during the failover. For example, it is not possible to implement link reversal algorithms as they require dynamic routing at the nodes.

Considering that a variety of parameters may influence the decision about the preferred **alternative output interface**, fast-recovery mechanisms can come in different flavors (see the overall classification shown in Fig. 34), for example:

- *Destination address matching*: Can forwarding depend on packet header fields other than the destination address? Solutions based solely on the destination address such as [266] are attractive, as they may require less forwarding rules. More general solutions which, e.g., also depend on the source address [260] as well as the source and destination port numbers used by transport-layer protocols [119], [230], may enable a more fine-grained traffic engineering scheme and thus reduce network load during the failover.
- *With or without input network interface matching*: Can the forwarding action to be applied to a packet depend

on the incoming link on which it arrived? Input interface matching can improve the resilience and quality of fast rerouting (in particular, by detecting and avoiding forwarding loops) [204], [267], but may render the forwarding logic more complex.

- *Stack-label matching*: Can messages be forwarded based on the label currently occupying the top position in a stack embedded in the message header? Stack-label matching enables flexible forwarding along pre-established paths in the network, without performing additional routing table lookups based on the values of the primary fields describing the source and the destination of the message [109], [110], [113]. Subsequent detours may be initiated simply by pushing a different label on the stack when necessary. At the same time, stack-label matching requires that the involved devices support the related extensions as well as a label distribution protocol maintaining consistency of the mapping between labels and the corresponding paths.
- *VLAN identifier matching*: Can the forwarding decision depend on the VLAN identifier stored in the message header? Unless the limited range of allowed values is likely to become an issue in specific deployments (especially those involving legacy network devices), VLAN identifiers might be used for fast recovery purposes as a convenient signaling channel leveraging the widely-supported network standard. Whenever a failure is encountered, the local Ethernet switch would typically set the VLAN identifier to a different value associated with one of the alternative spanning trees and then it would forward the message along the selected tree. Downstream switches would forward the message along the same tree until the destination is reached, or until another failed link is detected on the intended path towards the destination. Note that this method may not be used together with the typical VLAN functionality, as it would allow messages assigned to one VLAN to leak into a different VLAN, bypassing the security policy defined on routers.
- *Register/Ad-hoc field matching*: Can programmable network devices make forwarding decisions based on additional sources of information, for example, values stored in hardware registers or in ad-hoc fields associated with the processed packets? Considering an increasing range of potential applications involving programmable network devices as well as substantial efforts supporting the development of future self-driving networks, forwarding decisions may also be influenced by external factors represented by the current values of internal registers and ad-hoc fields. Consequently, the flexibility offered by the underlying systems may lead to a better integration with specific environments and to the development of unique fast-recovery solutions. The related advantage is that custom-designed packet processing pipelines may be evaluated and deployed much faster and potentially at a lower cost, compared to the equivalent proprietary off-the-shelf solutions. At the same time, new designs will remain constrained by the limitations of programmable devices and by increasing performance requirements.

Fig. 34: Classification of the presented fast-recovery mechanisms with respect to the match-action operations performed in the data plane.

A good example illustrating the trade-off between the resource utilization and performance (in terms of latency and throughput) has been presented in [258].

Based on the specific subset of **parameters** which are used by an algorithm to determine the preferred alternative output network interface, the corresponding actions may also be triggered if necessary, both in the context of messages as well as the entire forwarding pipelines maintained by network devices, for example:

- *Packet-header rewriting*: Can nodes rewrite packet headers depending on failed links, e.g., to store information about the failures encountered by a packet along its route? Packet-header rewriting is known to simplify the design of highly resilient fast rerouting algorithms [156], [249]. A well-known example are fast-reroute mechanisms based on failure-carrying packets [152]. Rewriting may also be performed on other objects, such as the internal registers used by programmable devices.
- *Action rewriting*: Can nodes rewrite the intended action assigned to subsequent matching messages, based on the detected signs of failure? To reduce the negative consequences (such as forwarding loops) further, or to optimize the resource utilization in the network, programmable devices may change the preferred action performed on the matching messages. However, an important related concern is preserving network stability following the

failure.

Existing mechanisms also differ in terms of their **objectives** and the operation mode (see Fig. 35):

- *Connectivity*: The most basic and most studied objective is to preserve connectivity. Ideally, a fast rerouting mechanism should always find a route to the destination, as long as the network remains physically connected. Especially the design of mechanisms without header rewriting has received much attention [20], [156], [220], [247], [265], [266], [268].
- *Distributed operation*: The majority of the solutions discussed in this paper have been designed to operate in a distributed fashion. The key advantage of this approach is that network devices can collect the necessary information, develop their internal state, and prepare for future failures without relying on the other devices in the network. At the same time, the recovery decisions may not always be optimal in the case of multiple failed elements, as the involved devices do not coordinate their response with each other.
- *Centralized operation*: Centralized fast-recovery approaches are still expected to be able to make local decisions without significant delay. However, in this scenario, forwarding devices usually depend on a central unit with respect to other key tasks, such as precomputation of the preferred (if necessary, optimal) network-wide recovery.

Fig. 35: Classification of the presented fast-recovery mechanisms with respect to the operation mode.

Fig. 36: Classification of the presented fast-recovery mechanisms with respect to the implementation and operation overheads.

ery strategy taking into account additional performance-critical factors (e.g., network load). As the centralized unit may need to collect the necessary information from the entire network or domain, process it, and then update the fast-recovery rules on forwarding devices, it may quickly become the bottleneck. Further, as the centralized unit and its connections may also fail or be subject to targeted attacks, the capability of the network to respond effectively to subsequent failures might become severely limited in such cases, unless additional measures are deployed to counteract this issue.

Deployment of fast-recovery solutions is always associated with additional **costs** (see Fig. 36), for example:

- **Precomputations:** The internal *knowledge* of the preferred recovery actions often results from precomputations which may be performed either locally or by an external unit (in the case of the centralized operation mode). It needs to be emphasized that this process is initiated during the normal network operation period and

in most cases should have completed by the time of the next failure event. Consequently, at the time of failure, the involved forwarding devices can redirect messages almost instantly. The related cost depends on the algorithmic complexity of the design, ranging from polynomial-time solutions to NP-hard problems.

- **Memory requirements:** To be able to perform the failover within milliseconds, forwarding devices need to develop and maintain internal state information that will define the preferred alternative network interfaces in different failure scenarios. In the context of modern programmable switches relying on expensive Ternary Content-Addressable Memory (TCAM) modules that are able to perform wildcard matches, it is desired that only the minimum required amount of information be stored in such memory modules to preserve space. Depending on how particular fast-recovery solutions have been designed, the memory-related overhead may either be constant or be bound to some parameters of the network

Fig. 37: Classification of the presented fast-recovery mechanisms with respect to the maximum number of failures they were designed to deal with. Note that under some specific circumstances, fast-recovery mechanisms designed to handle single failures might still be able to deal with multiple failures effectively (for example, when the failed components are located in different regions of a well-connected network).

such as number of destination prefixes.

- *Signaling*: Some fast-recovery designs rely on additional signaling to carry important information between the involved nodes. In particular, the selected bits of a message header may be used to indicate the preferred path in the network between the source and destination nodes. Alternatively, if the expected modifications of the message header would disrupt the operation of other network protocols and mechanisms such as VLANs, dedicated control messages may be exchanged by forwarding devices rather than a common protocol. Depending on the specific method in use, the related cost may result from the data overhead or from the additional operation limitations.
- *Incurred path length*: Many existing fast-recovery solutions cannot deal with simultaneous failures of multiple network elements effectively. Even those having such a capability may still be unable to forward messages around the failed components along the shortest possible paths. Consequently, they may not provide any guarantees regarding the observed path length (or stretch). However, the recent advancements in static fast-recovery mechanisms relying on modified arborescence-based network decompositions provide one of the possible solutions to this issue [262], [269].
- *Load*: Another important objective is to avoid congestion during and after failover, see e.g., [260], [270].

Finally, the effectiveness of existing fast-recovery mechanisms in terms of the maximum offered **resilience capabilities** is also diverse (see Fig. 37):

- *Single failures*: Early fast-recovery solutions, especially those operating in the link layer, have been designed to respond to single link or node failures. On one hand, simultaneous failures of two or more network elements are less likely to happen than a failure of a single component [164], which means that being able to restore

connectivity just in the case of single failures already covers the most frequent scenario. At the same time, in large networks involving numerous different devices and subsystems undergoing regular maintenance activities, such events are not uncommon and single-failure-recovery strategies may not always be successful, leading to increased packet losses, disruptive delay, and even persistent or transient forwarding loops. Indeed, according to [164], scheduled maintenance activities alone may have caused 20% of all observed failures in an operational IP backbone network, while almost 30% of all unplanned failures affected multiple links at a time. In this context, it needs to be emphasized that both the resilience requirements as well as the overall complexity of networked systems have been constantly increasing over time.

- *Double link failures*: To extend the fast-recovery capabilities of computer and communication networks beyond the single-failure scenario, improved designs were developed that were able to deal with double link failures. Considering the evolution of fast-recovery strategies, this was an intermediate step towards more general strategies handling simultaneous failures of multiple network components.
- *Multiple link failures*: Dealing with multiple link failures effectively depends not only on the design of a recovery mechanism, but also on the physical network topology. One of the key related parameters coming from graph theory is the edge connectivity of the network topology, further referred to as k . In particular, if up to $k - 1$ links in a given k -connected network suddenly become unavailable, the remaining links can still be used to reach any destination in the network⁹. An example group of the related recent solutions is focused on static fast-reroute

⁹Note that this condition is only related to graph connectivity, while in real networked environments, several additional factors need to be considered as well, such as traffic characteristics, the load of particular network components, and the relevant traffic engineering policies.

mechanisms. The underlying **algorithmic techniques** include:

- *Arborescence-based network decompositions*: By decomposing the network into a set of arc-disjoint arborescences which are rooted at the destination [266], [267] (such a decomposition can be computed in polynomial time), a high degree of resilience can be achieved: when encountering a link failure along one arborescence, a packet can simply follow an alternative arborescence.
- *Combinatorial block designs*: Pignolet et al. [260], [270] observed and exploited a connection of fast rerouting problems to combinatorial block designs: static resilient routing boils down a subfield in distributed computing which does not allow for communication.

Ideally, a static fast rerouting algorithm ensures connectivity whenever the underlying physical network is connected. Feigenbaum et al. [20] showed that without packet header rewriting, such an ideal static resilience cannot be achieved. A weaker notion of resilience was introduced by Chiesa et al. [266], [267]: the authors showed that there exist randomized static rerouting algorithms which tolerate $k - 1$ link failures if the underlying network is k -edge connected, even without header rewriting. At the same time, a fundamental open problem is whether for any k -connected graph, one can find deterministic failover routing functions that are robust to any $k - 1$ failures.

X. DISCUSSION

While our main focus is on data plane fast-recovery mechanisms, we point out that additional challenges may arise from the interaction between the control plane and the data plane during recovery. In particular, data plane mechanisms are often seen as a “first line of defense”: a way to react *quickly*, but possibly *suboptimally*, to failures. For instance, a fast-recovery path may allow to bypass a failed network component instantaneously but, by directing additional traffic to links that in normal circumstances would not have to cope with that traffic, it may cause unexpected link capacity overload, network congestion, packet losses, and other service disruptions. Worse yet, such phenomena may spread to remote parts of the network, harming flows that would not have been affected by the failure otherwise.

These routing transient states may lead to just minor performance degradation as long as they are short-lived and the network can quickly restore the default routing configuration when the failed component eventually comes back online. In this context, a failure is “short-lived” if it does not last much longer than the full *restoration completion time* (recall the recovery procedure timeline in Fig. 7, Section II) and it disappears before the network would enter the *normalization phase*. Typical examples are a quick router reboot event due to a software bug or a flapping interface going through an up-down-up cycle, which usually last only a couple of milliseconds or seconds at the worst. If the failure proves long-lived so that the normalization phase is initiated (recall

again Fig. 7, Section II), then in the second stage the control plane must reconverge to a new static route allocation that is optimized for the changed network configuration, with the failed component permanently removed from the topology. Another transient phase takes place after the normalization process terminates and the outage gets fixed, during which the network reconverges to the original state that existed before the failure. Unfortunately, in each of these transient phases further service disruptions may occur due to certain types of inadvertent routing transients, like short-lived forwarding loops (so called microloops) and routing blackholes, caused by the inconsistent timing of updating the data-plane configuration at different routers.

In order to avoid performance degradation during routing transients, there is a need to carefully orchestrate the way the control plane interacts with the data plane, both in the context of fast recovery and traffic engineering. In this section, we briefly cover some of the related issues and we carefully refer the reader to additional in-depth discussion on each covered topic.

A. Reconvergence and Transient Loops

Different approaches in the literature deal with the interaction between data-plane fast reroute and the control-plane reconvergence in different ways. We discuss some examples in the context of intra-domain and inter-domain routing, and programmable networks.

It is perhaps the context of *shortest-path-based distributed intra-domain routing* where routing transients manifest their adverse effect most visibly. This is on the one hand due to the inherent distributed nature of IGP, where there is minimal or no central oversight on when and how data-plane updates are performed, and on the other hand because of the fundamental properties of shortest-path routing that make it possible for two routers, e.g., one router that is aware of a failure and another one that is not, to appoint each other as the next-hop towards a particular IP destination prefix (leading to a microloop) or failing to synchronize at a consistent forwarding path (leading to a transient or permanent blackhole).

To better understand how much time is required for a link-state IGP to perform different actions following a change of a network topology, the reader is referred to [15] where the authors discuss the related challenges based on detailed measurements and a simulation model. In particular, the IGP convergence time on a single router has been characterized as $D + O + F + SPT + RIB + DD$ where D represents the failure detection time, O — the LSP origination time, F — the LSP flooding time, SPT — the shortest-path tree computation time, RIB — time required to update the local instances of the Routing Information Base (RIB) as well as the Forwarding Information Base (FIB), and DD — the LSP distribution delay. While D , O , and DD are relatively small, the remaining parameters depend either on the network topology or on the number of network prefixes for which the next-hop will change, and thus they become the dominant factors. To limit the undesired transient effects to some degree, link-state IGPs can rely on incremental SPF algorithms and

incremental FIB updates. Note that the IGP convergence time can be improved even further by careful tuning, as long as it does not affect the stability of the network considerably [15], [271]. At the same time, even though the network is expected to converge faster as a result of those improvements, routing transients may still be observed in this case.

A cornerstone contribution for avoiding routing transients during updating the FIBs of different routers in the context of IP routing protocols is made in [272]. The paper proposes a mechanism to control OSPF administrative link weights for adding links to, or removing links from, the network, by progressively changing the shortest-path link weight metric on certain “key” links to reach the required modification, ensuring that in each step of the progression the topology change occurs in a loop-free manner. Note that this mechanism needs a central entity to plan the progression and drive the process; hence, it can be used in the case of a non-urgent (management action) link or node shutdown or restart as well as link metric changes.

The Ordered FIB update (oFIB) proposal, reaching the Informational RFC status in 2013, brings this work further [273]. The idea is to correctly sequence the FIB updates on the routers by each router computing a rank separately that defines the time at which it can safely update its FIB. This approach can be used to administratively add or remove links (similarly to [272]), when there is sufficient time to pace out FIB updates, but it is also useful in conjunction with a fast-reroute mechanism when a link or node failure persists and the task is to re-converge the transient routes created by FRR to the static IGP shortest-path routes.

Another example is situated in the context of *inter-domain fast IP recovery*, and more specifically, BGP and R-BGP [227]. During BGP reconvergence, BGP withdrawals may be sent from an AS A adjacent to a link failure to all its neighbors. These neighbors may, in response to these withdrawals, stop forwarding traffic to A and, consequently, drop traffic. This is undesirable especially in the case when A can still use a fast-recovery path to forward packets to the destination that does not traverse the failed link¹⁰. In fact, the fast-reroute path is a temporary solution to avoid dropping packets and should not interfere with the post-convergence state computed by the control-plane.

R-BGP uses two simple ideas to ensure this. First, a node withdraws a route to a neighboring AS only when it is sure it will not offer it a “valley-free” route in the post-convergence state. Second, an AS keeps forwarding traffic on a withdrawn route until it either receives a new BGP route that does not traverse the failed link or it understands it will not receive a valid BGP route in the post-convergence state. We refer the reader to the original paper [227] for the details.

Similar problems related to the interaction between the control plane and the data plane arise in the context of *programmable networks and software-defined networks*. Even if the control plane is centralized (as it is the case in the standard SDN setup) the data plane is not, and therefore it is essential for the former to carefully orchestrate the update of

the latter, so that at any instance of time, the actual data-plane configuration is consistent. This is a surprisingly complex problem that has received much attention in the research community lately; we refer the reader to the recent survey [274] for an in-depth coverage of the related challenges. In the DDC approach [16], which we covered in Section VII-B, the authors also explain how to reconcile the output of the fast data-plane reconvergence, at any possible transient stage, with a centralized control plane that concurrently computes a global optimum and modifies the data-plane forwarding rules without creating any micro loops during this transition. An orthogonal approach is taken in [275] where the FRR mechanism in the data plane is used just to obtain additional failure information for the (centralized) control plane, for better failover planning; see also [230].

Control-plane–data-plane interaction becomes even more problematic when distributed and a centralized control planes coexist in the same network, both modifying the FIBs of the same set of network switches without synchronizing. Recently, Vissicchio et al. [276] developed a general theory that characterizes the kinds of routing and forwarding anomalies that may occur if the two (or more) control planes act independently from each other.

B. Traffic Engineering and Fast Reconvergence

As mentioned earlier, careful preparation is needed during the switch from the default forwarding paths to the fast-recovery paths after a failure, as well as during switching from the recovery paths to the new forwarding paths, should the failure prove permanent. This is in order to avoid that certain links or entire regions of the network be overwhelmed with traffic bypassing the failure and preclude cascading failures due to the resultant congestion. Accomplishing this traffic optimization task usually requires a concerted effort on the part of the data plane and the control plane, so that the recovery paths fulfill the traffic engineering goals even under failures. There is a breadth of standards, operational best practices, and research proposals to reach this goal; below we give a non-exhaustive overview of some of the most typical approaches; for more detail, refer to the original papers and the related surveys [274].

A significant number of fast-recovery mechanisms do not natively support fine-grained traffic engineering during recovery. In these cases, recovery occurs on a “best-effort” basis, merely hoping that the failover paths will be “good enough” in that there is enough over-provisioned spare capacity available in the network to avoid congestion. This approach is adopted most often for data-plane technologies that otherwise provide very little in the way of optimizing forwarding paths, like L2 protocols (e.g., AXE [97], [100], see Section III), intra-domain IP fast reroute (e.g., the original LFA proposal [32], see Section V) or inter-domain IP routing where BGP does not provide fine-grained mechanisms for optimizing AS-AS paths (e.g., [227], see Section VI).

The second approach is to compute recovery paths on-demand, in a way to ensure that the recovery paths will have enough capacity to handle the failover traffic. This approach

¹⁰Note that BGP withdrawals have to be sent out even in this case.

is used for data-plane technologies that provide means for the control plane to tailor and optimize the forwarding and the recovery paths with respect to arbitrary traffic engineering goals. In this case, the selected TE designs could be implemented on top of the default fast-recovery scheme, overriding the decision made by the underlying fast-recovery mechanism based on additional TE metrics (provided that multiple backup network interfaces are available).

An example showing how traffic engineering mechanisms co-exist with built-in recovery solutions is reflected by the design of MPLS and its extensions (see Section IV). In this context, the functional capabilities allowing for an implementation of traffic engineering policies have been defined in [107]. The related *resilience attributes* may either be limited to just indicate which recovery procedure should be applied to traffic trunks affected by failures (*basic* resilience attribute) or they can also specify detailed actions that should be taken in the case of failure (*extended* resilience attribute). In particular, the extended resilience attribute may define a set of backup paths as well as the rules which control the relative preference of each path. To be able to impose traffic engineering policies, MPLS relies on close interaction with routing. Further extensions to MPLS provide a way to define label-switched tunnels that can be automatically routed away from failed network components, congested links, and bottlenecks [110]. The LSP tunnels are established and torn down by RSVP which takes into account the available resources and existing traffic engineering policies. Meanwhile, establishing new tunnels at the time of failure may lead to packet losses and increased transmission delay. For this reason, the additional extensions proposed in [113] introduced the much-needed capability of MPLS to establish backup LSP tunnels in advance, to be able to perform a fast failover locally within tens of milliseconds following the failure. See e.g., [131] for a linear TE-aware optimization model for link, node, and bandwidth protection in the context of MPLS.

The third approach encompasses a broad set of models, methodologies and algorithms for the co-design of the default, failure-free forwarding paths and the failover paths, in a way as to minimize the adverse effects of the routing transients and service disruptions that occur during the fast-recovery process. This problem is again posed in different forms depending on the data-plane technology used. In intra-domain shortest-path-based IP routing any non-trivial change in the IGP link weights applied to bypass a failed network component will necessarily reroute some, or even all, traffic instances that would otherwise not be affected by the failure [272], [273]. To prevent such cascading service disruptions, [277, Section 4] presents a set of local search algorithms to co-design the default IGP weights and the failover IGP weights so that the number of necessary weight changes, and hence (heuristically) the number of traffic instances being rerouted and thereby suffering service disruptions, is minimal. They show that in many cases as few as 3 weight changes is enough to attain “good enough” performance.

Of course, the best results with the co-design approach can be achieved only with data-plane technologies that, in contrast to shortest-path-based IP routing, allow the control plane

to adjust and fine-tune the forwarding path of each source-destination pair separately. In this context, [143], [144] present heuristic algorithms for the protection routing scheme (which can be used even in a plain IP data plane with a centralized control plane on top, see Section V) and R3 [119] shows an elegant extension of oblivious routing to convert “traffic uncertainty” to “topology uncertainty” that may result from failures (R3 can be used with e.g., MPLS, see Section V); refer to the previous sections for a detailed overview of these and similar techniques.

The general observation is that, even though *survivability and performance are fundamentally at odds* [143], [144] by the need to stack valuable transmission capacity to accommodate bypass paths that could otherwise be used to increase failure-free throughput, simple algorithmic techniques can generally find good trade-offs to reconcile these two conflicting requisites. For this, however, a sufficient data-plane technology and a careful control-plane-data-plane co-design approach is necessary.

C. Multicast Fast Reroute

Fast reroute mechanisms also exist for multicast communication scenarios. A prominent example is based on BIER (Bit Index Explicit Replication) [278], defined in the IETF. BIER is essentially a point-to-multipoint tunnel without explicit tunneling states in the network, that is, BIER does not require states and signaling for multicast in core networks. The 1+1 protection mechanism based on maximally redundant trees used for BIER is described in [279], and a 1:1 FRR scheme based on point-to-multipoint reroute tunnels in [280].

XI. CONCLUSION

This paper presented a structured overview of the concepts and state-of-the-art methods for providing fast recovery in the data plane. By going through the layers, we have discussed a variety of solutions which come with different features and requirements. Given the emerging flexibilities of software-defined networks, one may wonder whether there are some methods which are preferable over others. In general, we can conclude that there is no free lunch. For example, there can be tradeoffs between resilience and efficiency: in scenarios where it is sufficient to provide resilience against one failure, it may be preferable to choose simple solutions, as state-of-the-art mechanisms for many failures can lead to long routes, even under a small number of failures (e.g., approaches based on arborescences). We have also seen that the ability to change the header depending on the failures can greatly improve resilience: it may be possible to achieve a perfect resilience, i.e., to stay connected as long as the underlying network is connected; this is known to be impossible without header rewriting. An example where a “technology detail” can make a big difference regards the ability to match the input port: destination-based-only routing is often less powerful (e.g., if headers are immutable) than routing which can also depend on the source, and if additionally the input port can be matched, the resilience may be improved even further.

We hope that our paper can help practitioners, decision makers, and theoreticians to understand the tradeoffs behind state-of-the-art and future fast recovery methods. We also hope that our paper will be useful for students and researchers interested in this fascinating topic between theory and practice, and help them bootstrap quickly in this field.

ACKNOWLEDGMENTS

This research is partially supported by the WWTF project, *Fast and Quantitative What-if Analysis for Dependable Communication Networks* (WHATIF), ICT19-045, 2020-2024, and by the European Cooperation in Science and Technology (COST) Action CA15127, *Resilient communication services protecting end-user applications from disaster-based failures — RECODIS*. G.R. is also with the Budapest University of Technology and Economics and the MTA-BME Momentum Network Softwarization Research Group. This work was partially supported by Ericsson Research, Hungary.

REFERENCES

- [1] R. Chirgwin, “Google routing blunder sent japan’s internet dark on friday,” in https://www.theregister.co.uk/2017/08/27/google_routing_blunder_sent_japans_internet_dark/, 2017.
- [2] D. Tweney, “5-minute outage costs Google \$545,000 in revenue,” in <http://venturebeat.com/2013/08/16/3-minute-outage-costs-google-545000-in-revenue/>, 2013.
- [3] G. Corfield, “British Airways’ latest total inability to support upwardness of planes caused by Amadeus system outage,” in https://www.theregister.co.uk/2018/07/19/amadeus_british_airways_outage_load_sheet/, 2018.
- [4] C. Gibbs, “ATT’s 911 outage result of mistakes made by ATT, FCC’s Pai says,” in <https://www.fiercewireless.com/wireless/at-t-s-911-outage-result-mistakes-made-by-at-t-fcc-s-pai-says>, 2017.
- [5] J. Young and T. Barth, “Web performance analytics show even 100-millisecond delays can impact customer engagement and online revenue,” *Akamai Online Retail Performance Report*, 2017.
- [6] J. Saldan, “Delay limits for real-time services,” *IETF Draft*, 2016.
- [7] P. Gill, N. Jain, and N. Nagappan, “Understanding network failures in data centers: measurement, analysis, and implications,” *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 350–361, 2011.
- [8] C. Labovitz, G. R. Malan, and F. Jahanian, “Internet routing instability,” *IEEE/ACM transactions on Networking*, no. 5, pp. 515–528, 1998.
- [9] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, “An approach to alleviate link overload as observed on an IP backbone,” in *Proc. IEEE INFOCOM*, vol. 1. IEEE, 2003, pp. 406–416.
- [10] J. Moy, “OSPF Version 2,” Internet Requests for Comments, RFC Editor, RFC 2328, Apr 1998. [Online]. Available: <https://tools.ietf.org/html/rfc2328>
- [11] ISO, “Intermediate System-to-Intermediate System (IS-IS) Routing Protocol,” ISO/IEC 10589, 2002.
- [12] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan, “Data center TCP (DCTCP),” *ACM SIGCOMM computer communication review*, vol. 41, no. 4, pp. 63–74, 2011.
- [13] B. Vamanan, J. Hasan, and T. Vijaykumar, “Deadline-aware datacenter TCP (D2TCP),” *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 115–126, 2012.
- [14] D. Zats, T. Das, P. Mohan, D. Borthakur, and R. Katz, “Detail: reducing the flow completion time tail in datacenter networks,” in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 2012, pp. 139–150.
- [15] P. Francois, C. Filsfil, J. Evans, and O. Bonaventure, “Achieving Sub-second IGP Convergence in Large IP Networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 3, pp. 35–44, Jul. 2005. [Online]. Available: <http://doi.acm.org/10.1145/1070873.1070877>
- [16] J. Liu, A. Panda, A. Singla, B. Godfrey, M. Schapira, and S. Shenker, “Ensuring connectivity via data plane mechanisms,” in *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. Lombard, IL: USENIX, 2013, pp. 113–126. [Online]. Available: https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/liu_junda
- [17] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, “A clean slate 4d approach to network control and management,” *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 41–54, Oct. 2005. [Online]. Available: <http://doi.acm.org/10.1145/1096536.1096541>
- [18] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, “Openflow: Enabling innovation in campus networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1355734.1355746>
- [19] H. Yan, D. A. Maltz, T. E. Ng, H. Gogineni, H. Zhang, and Z. Cai, “Tesseract: A 4d network control plane,” in *NSDI*, vol. 7, 2007, pp. 27–27.
- [20] J. Feigenbaum, B. Godfrey, A. Panda, M. Schapira, S. Shenker, and A. Singla, “Brief announcement: On the resilience of routing tables,” in *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, 2012, pp. 237–238.
- [21] D. Stamatelakis and W. D. Grover, “IP layer restoration and network planning based on virtual protection cycles,” *IEEE Journal on selected areas in communications*, vol. 18, no. 10, pp. 1938–1949, 2000.
- [22] A. Kabbani, B. Vamanan, J. Hasan, and F. Duchene, “FlowBender: Flow-level adaptive routing for improved latency and throughput in datacenter networks,” in *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT ’14. New York, NY, USA: ACM, 2014, pp. 149–160. [Online]. Available: <http://doi.acm.org/10.1145/2674005.2674985>
- [23] J. Papan, P. Segeč, M. Moravcik, M. Kontsek, L. Mikus, and J. Uramova, “Overview of IP Fast Reroute solutions,” in *2018 16th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, Nov 2018, pp. 417–424.
- [24] A. Jarry, “Fast reroute paths algorithms,” *Telecommunication Systems*, vol. 52, no. 2, pp. 881–888, 2013.
- [25] A. Kamisiński, “Evolution of IP Fast-Reroute Strategies,” in *2018 10th International Workshop on Resilient Networks Design and Modeling (RNDM)*, Aug 2018, pp. 1–6.
- [26] P. Pan, G. Swallow, and A. Atlas, “Fast reroute extensions to RSVP-TE for LSP tunnels,” in *Request for Comments (RFC) 4090*, 2005.
- [27] Switch Specification 1.3.1, “OpenFlow,” in <https://bit.ly/2VjOO77>, 2013.
- [28] Cisco, “Configuring BGP PIC Edge and Core for IP and MPLS,” Oct. 2017. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-3s/irg-xe-3s-book/irg-bgp-mp-pic.html
- [29] D. Xu, Y. Xiong, C. Qiao, and G. Li, “Failure protection in layered networks with shared risk link groups,” *IEEE Network*, vol. 18, no. 3, pp. 36–41, 2004.
- [30] P. Sebos, J. Yates, G. Hjalmtysson, and A. Greenberg, “Auto-discovery of shared risk link groups,” in *Proc. Optical Fiber Communication Conference and Exhibit (OFC)*, vol. 3, 2001.
- [31] M. Menth, M. Duelli, R. Martin, and J. Milbrandt, “Resilience analysis of packet-switched communication networks,” *IEEE/ACM Transactions on Networking*, vol. 17, no. 6, pp. 1950–1963, 2009.
- [32] 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>
- [33] T. Elhourani, A. Gopalan, and S. Ramasubramanian, “IP fast rerouting for multi-link failures,” in *Proc. IEEE INFOCOM*, 2014.
- [34] M. Golash, “Reliability in ethernet networks: A survey of various approaches,” *Bell Labs Technical Journal*, vol. 11, no. 3, pp. 161–171, 2006.
- [35] 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.
- [36] J. Papán, P. Segeč, and P. Palúch, “Analysis of existing IP fast reroute mechanisms,” in *2015 International Conference on Information and Digital Technologies*, July 2015, pp. 291–297.
- [37] A. Raj and O. C. Ibe, “A survey of IP and multiprotocol label switching fast reroute schemes,” *Computer Networks*, vol. 51, no. 8, pp. 1882–1907, 2007.
- [38] L. Jorge and T. Gomes, “Survey of recovery schemes in mpls networks,” in *2006 International Conference on Dependability of Computer Systems*. IEEE, 2006, pp. 110–118.

- [39] R. B. da Silva and E. S. Mota, "A survey on approaches to reduce BGP interdomain routing convergence delay on the Internet," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2949–2984, 2017.
- [40] A. S. da Silva, P. Smith, A. Mauthe, and A. Schaeffer-Filho, "Resilience support in software-defined networking: A survey," *Computer Networks*, vol. 92, pp. 189–207, 2015.
- [41] M. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, "Disaster survivability in optical communication networks," *Computer Communications*, vol. 36, no. 6, pp. 630–644, 2013.
- [42] A. J. Piotr Cholda, "Recovery and its quality in multilayer networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 28, no. 4, pp. 372–389, 2010.
- [43] S. D. Maesschalck, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jaeger, R. Inkret, B. Mikac, and J. Derkacz, "Pan-european optical transport networks: An availability-based comparison," vol. 5, no. 3, pp. 203–225, 2003.
- [44] 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)*, 2016, pp. 1–4.
- [45] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," Internet Requests for Comments, RFC Editor, RFC 4271, Jan 2006. [Online]. Available: <https://tools.ietf.org/html/rfc4271>
- [46] O. D. Smita Rai, Biswanath Mukherjee, "IP resilience within an autonomous system: Current approaches, challenges, and future directions," *IEEE Communications Magazine*, vol. 43, no. 10, pp. 142–149, 2005.
- [47] W. C. Hardy, "QoS" *Measurement and Evaluation of Telecommunications Quality of Service*. John Wiley & Sons, 2001.
- [48] J. Gozdecki, R. Stankiewicz, and A. Jajszczyk, "Quality of service terminology in IP networks," *IEEE Communications Magazine*, vol. 41, no. 3, pp. 153–159, 2003.
- [49] R. Stankiewicz, P. Cholda, and A. Jajszczyk, "QoX: What is it really?" *IEEE Communications Magazine*, vol. 49, no. 4, pp. 148–158, 2011.
- [50] "ITU-T Recommendation Y.1540," ITU-T, Tech. Rep.
- [51] "ITU-T Recommendation Y.1541," ITU-T, Tech. Rep.
- [52] A. F. Hansen, T. Cicic, and S. Gjessing, "Alternative schemes for proactive IP recovery," in *Proc. of 2006 2nd Conference on Next Generation Internet Design and Engineering, 2006. NGI '06*, 2006, pp. 1–8.
- [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] J. Sterbenz, D. Hutchison, E. Cetinkaya, A. Jabbar, J. Rohrer, M. Schoeller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles and survey of disciplines," *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.
- [55] J. Sterbenz, C. E.K., M. Hameed, A. Jabbar, S. Qian, and J. Rohrer, "Evaluation of network resilience, survivability, and disruption tolerance: Analysis, topology generation, simulation, and experimentation," *Telecommunication Systems*, vol. 52, no. 2, pp. 705–736, 2013.
- [56] J. Rak, *Resilient Routing in Communication Networks*, 1st ed., ser. Computer Communications and Networks. Springer Publishing Company, Incorporated, 2015.
- [57] M. Khabbaz, C. Assi, and W. Fawaz, "Disruption-tolerant networking: a comprehensive survey," *IEEE Communication Surveys & Tutorials*, vol. 14, no. 2, pp. 607–640, 2012.
- [58] A. Avizienis, J.-C. Laprie, and B. Randell, "Dependability and its threats: A taxonomy," in *Building the Information Society*, J. R., Ed. IFIP International Federation for Information Processing, Springer, 2004, vol. 156, pp. 91–120.
- [59] 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, no. 1, pp. 11–33, 2004.
- [60] P. Cholda, J. Tapolcai, T. Cinkler, K. Wajda, and A. Jajszczyk, "Quality of resilience as a network reliability characterization tool," *IEEE Network*, vol. 23, no. 2, pp. 11–19, 2009.
- [61] A. K. A. Autenrieth, "Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS," *IEEE Communications Magazine*, vol. 40, no. 1, pp. 50–57, 2002.
- [62] P. Cholda, A. Mykkeltveit, B. Helvik, O. Wittner, and A. Jajszczyk, "A survey of resilience differentiation frameworks in communication networks," *IEEE Communication Surveys*, vol. 9, no. 4, pp. 32–55, 2007.
- [63] C. Huang, V. Sharma, K. Owens, and S. Makam, "Building reliable MPLS networks using a path protection mechanism," *IEEE Communications Magazine*, vol. 40, no. 3, pp. 156–162, 2002.
- [64] V. S. et al., "Framework for MPLS-based recovery," IETF, <https://tools.ietf.org/html/rfc3469>, RFC, Informational Standard 3469, February 2003.
- [65] A. Autenrieth, "Recovery time analysis of differentiated resilience in mpls," in *Proc. of DRCN 2003 - Design of Reliable Communication Networks, Banff, Alberta, Canada, October 19-22, 2003*, pp. 333–340.
- [66] P. D. Jean-Philippe Vasseur, Mario Pickavet, *Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Morgan Kaufmann, 2004.
- [67] A. S. S. Ayush Dusia, "Recent advances in fault localization in computer networks," *IEEE Communications Surveys and Tutorials*, vol. 18, no. 4, pp. 3030–3051, 2016.
- [68] "G.975: Forward error correction for submarine systems," ITU-T, Tech. Rep., 2000.
- [69] C. M. Machuca and P. Thiran, "An efficient algorithm for locating soft and hard failures in WDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1900–1911, 2000.
- [70] S. Zhuang, D. Geels, I. Stoica, and R. Katz, "On failure detection algorithms in overlay networks," in *Proc. of IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 3, 2005, pp. 2112–2123.
- [71] D. Katz and D. Ward, "Bidirectional Forwarding Detection (BFD)," Internet Requests for Comments, RFC Editor, RFC 5880, Jun 2010. [Online]. Available: <https://tools.ietf.org/html/rfc5880>
- [72] D. G. R. Steinert, "Towards distributed and adaptive detection and localisation of network faults," in *Proc. of 2010 Sixth Advanced International Conference on Telecommunications*, 2010, pp. 384–389.
- [73] G. N. Reuven Cohen, "Maximizing restorable throughput in MPLS networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 2, pp. 568–581, 2010.
- [74] "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges," *IEEE Std 802.1D-2004 (Revision of IEEE Std 802.1D-1998)*, pp. 1–281, June 2004, https://standards.ieee.org/standard/802_1D-2004.html.
- [75] J. Qiu, M. Gurusamy, K. C. Chua, and Y. Liu, "Local restoration with multiple spanning trees in metro ethernet," in *2008 International Conference on Optical Network Design and Modeling*, March 2008, pp. 1–6.
- [76] K. Elmeleegy, A. L. Cox, and T. S. E. Ng, "On count-to-infinity induced forwarding loops ethernet networks," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, April 2006, pp. 1–13.
- [77] Li Su, Wentao Chen, Haibo Su, Zhenyu Xiao, Depeng Jin, and Lieguang Zeng, "Ethernet Ultra Fast Switching: A tree-based local recovery scheme," in *2008 11th IEEE Singapore International Conference on Communication Systems*, Nov. 2008, pp. 314–318.
- [78] "Part 3: Media Access Control (MAC) Bridges: Amendment 2 - Rapid Reconfiguration," *IEEE Std 802.1w-2001 (Amendment to IEEE Std 802.1d and 802.1t)*, pp. 1–116, July 2001, https://standards.ieee.org/standard/802_1w-2001.html.
- [79] "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks," *IEEE Std 802.1Q-2003 (Incorporates IEEE Std 802.1Q-1998, IEEE Std 802.1u-2001, IEEE Std 802.1v-2001, and IEEE Std 802.1s-2002)*, pp. 1–322, May 2003.
- [80] A. Gopalan and S. Ramasubramanian, "Fast recovery from link failures in Ethernet networks," in *2013 9th International Conference on the Design of Reliable Communication Networks (DRCN)*, Mar. 2013, pp. 1–10.
- [81] —, "Fast Recovery From Link Failures in Ethernet Networks," *IEEE Transactions on Reliability*, vol. 63, no. 2, pp. 412–426, Jun. 2014.
- [82] "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges," *IEEE Std 802.1D-1990*, pp. 1–176, March 1991.
- [83] "IEEE Standard for Local Area Network MAC (Media Access Control) Bridges," *ANSI/IEEE Std 802.1D, 1998 Edition*, pp. 1–373, Dec 1998.
- [84] "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks," *IEEE Std 802.1Q-1998*, pp. 1–214, March 1999.
- [85] S. Varadarajan and T. Chiueh, "Automatic fault detection and recovery in real time switched ethernet networks," in *IEEE INFOCOM '99. Conference on Computer Communications. Proceedings. Eighteenth Annual Joint Conference of the IEEE Computer and Communications*

- Societies. The Future is Now (Cat. No.99CH36320)*, vol. 1, March 1999, pp. 161–169 vol.1.
- [86] S. Sharma, K. Gopalan, S. Nanda, and T. Chiueh, “Viking: a multi-spanning-tree Ethernet architecture for metropolitan area and cluster networks,” in *IEEE INFOCOM 2004*, vol. 4, Mar. 2004, pp. 2283–2294 vol.4.
 - [87] R. Pallos, J. Farkas, I. Moldovan, and C. Lukovszki, “Performance of rapid spanning tree protocol in access and metro networks,” in *2007 Second International Conference on Access Networks Workshops*, Aug. 2007, pp. 1–8.
 - [88] Depeng Jin, Wentao Chen, Zhenyu Xiao, and Lieguang Zeng, “Single link switching mechanism for fast recovery in tree-based recovery schemes,” in *2008 International Conference on Telecommunications*, Jun. 2008, pp. 1–5.
 - [89] D. Jin, Y. Li, W. Chen, L. Su, and L. Zeng, “Ethernet ultra-fast switching: a tree-based local recovery scheme,” *IET Communications*, vol. 4, no. 4, pp. 410–418, Mar. 2010.
 - [90] J. Qiu, M. Gurusamy, K. C. Chua, and Y. Liu, “Local Restoration With Multiple Spanning Trees in Metro Ethernet Networks,” *IEEE/ACM Transactions on Networking*, vol. 19, no. 2, pp. 602–614, Apr. 2011.
 - [91] J. Qiu, Y. Liu, G. Mohan, and K. C. Chua, “Fast Spanning Tree Reconnection for Resilient Metro Ethernet Networks,” in *2009 IEEE International Conference on Communications*, Jun. 2009, pp. 1–5.
 - [92] M. Terasawa, M. Nishida, S. Shimizu, Y. Arakawa, S. Okamoto, and N. Yamanaka, “Recover-Forwarding Method in Link Failure with Pre-Established Recovery Table for Wide Area Ethernet,” in *2009 IEEE International Conference on Communications*, Jun. 2009, pp. 1–5.
 - [93] J. Qiu, G. Mohan, K. C. Chua, and Y. Liu, “Handling Double-Link Failures in Metro Ethernet Networks Using Fast Spanning Tree Reconnection,” in *GLOBECOM 2009 - 2009 IEEE Global Telecommunications Conference*, Nov. 2009, pp. 1–6.
 - [94] J. Farkas and Z. Arató, “Performance analysis of shortest path bridging control protocols,” in *GLOBECOM 2009 - 2009 IEEE Global Telecommunications Conference*, Nov. 2009, pp. 1–6.
 - [95] “IEEE Standard for Local and metropolitan area networks—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks,” *IEEE Std 802.1Q-2011 (Revision of IEEE Std 802.1Q-2005)*, pp. 1–1365, Aug 2011.
 - [96] D. M. Shan, C. K. Chiang, G. Mohan, and J. Qiu, “Partial Spatial Protection for Differentiated Reliability in FSTR-Based Metro Ethernet Networks,” in *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, Dec. 2011, pp. 1–5.
 - [97] J. McCauley, A. Sheng, E. J. Jackson, B. Raghavan, S. Ratnasamy, and S. Shenker, “Taking an axe to L2 spanning trees,” in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XIV. New York, NY, USA: ACM, 2015, pp. 15:1–15:7.
 - [98] “IEEE Standard for Local and metropolitan area networks — Bridges and Bridged Networks - Amendment 24: Path Control and Reservation,” *IEEE Std 802.1Qca-2015 (Amendment to IEEE Std 802.1Q-2014 as amended by IEEE Std 802.1Qcd-2015 and IEEE Std 802.1Q-2014/Cor 1-2015)*, pp. 1–120, March 2016.
 - [99] M. Santos and J. Grégoire, “Improving carrier ethernet recovery time using a fast reroute mechanism,” in *2016 23rd International Conference on Telecommunications (ICT)*, May 2016, pp. 1–7.
 - [100] J. McCauley, M. Zhao, E. J. Jackson, B. Raghavan, S. Ratnasamy, and S. Shenker, “The deforestation of L2,” in *Proceedings of the 2016 ACM SIGCOMM Conference*, ser. SIGCOMM ’16. New York, NY, USA: ACM, 2016, pp. 497–510.
 - [101] “IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks,” *IEEE Std 802.1Q-2005 (Incorporates IEEE Std 802.1Q1998, IEEE Std 802.1u-2001, IEEE Std 802.1v-2001, and IEEE Std 802.1s-2002)*, pp. 1–300, May 2006.
 - [102] “IEEE Standard for Local and metropolitan area networks—Media Access Control (MAC) Bridges and Virtual Bridges [Edition],” *IEEE Std 802.1Q-2012 (Incorporating IEEE Std 802.1Q-2011, IEEE Std 802.1Qbe-2011, IEEE Std 802.1Qbc-2011, IEEE Std 802.1Qbb-2011, IEEE Std 802.1Qaz-2011, IEEE Std 802.1Qbf-2011, IEEE Std 802.1Qbg-2012, IEEE Std 802.1aq-2012, IEEE Std 802.1Q-2012)*, pp. 1–1782, Dec 2012.
 - [103] “IEEE Standard for Local and metropolitan area networks—Bridges and Bridged Networks,” *IEEE Std 802.1Q-2014 (Revision of IEEE Std 802.1Q-2011)*, pp. 1–1832, Dec 2014.
 - [104] “IEEE Standard for Local and Metropolitan Area Network—Bridges and Bridged Networks,” *IEEE Std 802.1Q-2018 (Revision of IEEE Std 802.1Q-2014)*, pp. 1–1993, July 2018.
 - [105] W. Grover and D. Stamatelakis, “Cycle-oriented distributed pre-configuration: ring-like speed with mesh-like capacity for self-planning network restoration,” in *1998 IEEE International Conference on Communications ICC’98*, 1998.
 - [106] O. Lemesko and K. Arous, “Fast reroute model for different backup schemes in MPLS network,” in *2014 First International Scientific Practical Conference Problems of Infocommunications Science and Technology*, Oct 2014, p. 39 41.
 - [107] J. McManus, J. Malcolm, M. D. O’Dell, D. O. Awduche, and J. Agogbua, “Requirements for Traffic Engineering Over MPLS,” RFC 2702, Sep. 1999. [Online]. Available: <https://tools.ietf.org/html/rfc2702>
 - [108] D. L. Haskin and R. Krishnan, “A Method for Setting an Alternative Label Switched Paths to Handle Fast Reroute,” Internet Engineering Task Force, Internet-Draft draft-haskin-mpls-fast-reroute-05, Nov. 2000, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-haskin-mpls-fast-reroute-05>
 - [109] E. Rosen, A. Viswanathan, and R. Callon, “Multiprotocol Label Switching Architecture,” Internet Requests for Comments, RFC Editor, RFC 3031, January 2001. [Online]. Available: <https://tools.ietf.org/html/rfc3031>
 - [110] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, “RSVP-TE: Extensions to RSVP for LSP Tunnels,” Internet Requests for Comments, RFC Editor, RFC 3209, December 2001. [Online]. Available: <https://tools.ietf.org/html/rfc3209>
 - [111] M. Kodialam and T. V. Lakshman, “Dynamic routing of locally restorable bandwidth guaranteed tunnels using aggregated link usage information,” in *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213)*, vol. 1, April 2001, p. 376385 vol.1.
 - [112] L. Hundessa and J. D. Pascual, “Fast rerouting mechanism for a protected label switched path,” in *Proceedings Tenth International Conference on Computer Communications and Networks (Cat. No.01EX495)*, Oct 2001, pp. 527–530.
 - [113] 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>
 - [114] E. Mannie and D. Papadimitriou, “Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS),” Internet Requests for Comments, RFC Editor, RFC 4427, March 2006. [Online]. Available: <https://tools.ietf.org/html/rfc4427>
 - [115] D. Papadimitriou and E. Mannie, “Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration),” Internet Requests for Comments, RFC Editor, RFC 4428, March 2006. [Online]. Available: <https://tools.ietf.org/html/rfc4428>
 - [116] L. Andersson, I. Minei, and B. Thomas, “LDP Specification,” Internet Requests for Comments, RFC Editor, RFC 5036, October 2007. [Online]. Available: <https://tools.ietf.org/html/rfc5036>
 - [117] J. V. A. Farrel, A. Ayyangar, “Inter-Domain MPLS and GMPLS Traffic Engineering – Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions,” Internet Requests for Comments, RFC Editor, RFC 5151, February 2008. [Online]. Available: <https://tools.ietf.org/html/rfc5151>
 - [118] D. Wang and G. Li, “Efficient distributed bandwidth management for MPLS fast reroute,” *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, p. 486 495, April 2008.
 - [119] Y. Wang, H. Wang, A. Mahimkar, R. Alimi, Y. Zhang, L. Qiu, and Y. R. Yang, “R3: resilient routing reconfiguration,” *ACM SIGCOMM CCR*, vol. 40, no. 4, pp. 291–302, 2010.
 - [120] A. Hassan, M. Bazama, T. Saad, and H. T. Mouftah, “Investigation of fast reroute mechanisms in an optical testbed environment,” in *7th International Symposium on Highcapacity Optical Networks and Enabling Technologies*, Dec 2010, p. 247251.
 - [121] N. Sprecher and A. Farrel, “MPLS Transport Profile (MPLSTP) Survivability Framework,” Internet Requests for Comments, RFC Editor, RFC 6372, Sep 2011. [Online]. Available: <https://tools.ietf.org/html/rfc6372>
 - [122] G. Ramachandran, L. Ciavattone, and A. Morton, “Restoration measurements on an IP/MPLS backbone: The effect of fast reroute on link failure,” in *2011 IEEE Nineteenth IEEE International Workshop on Quality of Service*, June 2011, p. 16.
 - [123] K. Koushik, R. Cetin, and T. Nadeau, “Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base for Fast Reroute,” RFC 6445, Nov. 2011. [Online]. Available: <https://tools.ietf.org/html/rfc6445>

- [124] 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>
- [125] T. Benhcine, H. Elbiaze, and K. Idoudi, "Fast reroute based network resiliency experimental investigations," in *2013 15th International Conference on Transparent Optical Networks (ICTON)*, June 2013, p. 1 4.
- [126] O. Lemeshko, A. Romanyuk, and H. Kozlova, "Design schemes for mpls fast reroute," in *2013 12th International Conference on the Experience of Designing and Application of CAD Systems in Microelectronics (CADSM)*, Feb 2013, p. 202 203.
- [127] M. Taillon, T. Saad, R. Gandhi, Z. Ali, and M. Bhatia, "Updates to the Resource Reservation Protocol for Fast Reroute of Traffic Engineering GMPLS Label Switched Paths (LSPs)," Internet Requests for Comments, RFC Editor, RFC 8271, October 2017. [Online]. Available: <https://tools.ietf.org/html/rfc8271>
- [128] O. S. Yermenko, O. V. Lemeshko, and N. Tariki, "Fast reroute scalable solution with protection schemes of network elements," in *2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON)*, May 2017, pp. 783–788.
- [129] S. Schmid and J. Srba, "Polynomial-time what-if analysis for prefix-manipulating MPLS networks," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, April 2018, pp. 1799–1807.
- [130] J. S. Jensen, T. B. Krøgh, J. S. Madsen, S. Schmid, J. Srba, and M. T. Thorgeren, "P-Rex: Fast Verification of MPLS Networks with Multiple Link Failures," in *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '18. New York, NY, USA: ACM, 2018, pp. 217–227. [Online]. Available: <http://doi.acm.org/10.1145/3281411.3281432>
- [131] O. Lemeshko and O. Yermenko, "Linear optimization model of MPLS traffic engineering fast reroute for link, node, and bandwidth protection," in *2018 14th International Conference on Advanced Trends in Radioelectronics, Telecommunications and Computer Engineering (TCSET)*, Feb 2018, p. 1009 1013.
- [132] J. S. Arora, *Introduction to Optimum Design*, 4th ed. Boston: Academic Press, 2017.
- [133] D. Applegate and E. Cohen, "Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs," in *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '03. New York, NY, USA: ACM, 2003, pp. 313–324. [Online]. Available: <http://doi.acm.org/10.1145/863955.863991>
- [134] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg, "Cope: Traffic engineering in dynamic networks," in *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '06. New York, NY, USA: ACM, 2006, pp. 99–110. [Online]. Available: <http://doi.acm.org/10.1145/1159913.1159926>
- [135] C. J. Anderson, N. Foster, A. Guha, J.-B. Jeannin, D. Kozen, C. Schlesinger, and D. Walker, "Netkat: Semantic foundations for networks," *ACM SIGPLAN Notices*, vol. 49, no. 1, pp. 113–126, 2014.
- [136] P. Kazemian, G. Varghese, and N. McKeown, "Header space analysis: Static checking for networks," in *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, 2012, pp. 113–126.
- [137] L. Hundessa and J. Domingo-Pascual, "Reliable and fast rerouting mechanism for a protected label switched path," in *Global Telecommunications Conference, 2002. GLOBECOM '02. IEEE*, vol. 2, Nov 2002, pp. 1608–1612 vol.2.
- [138] M. Menth, A. Reifert, and J. Milbrandt, "Self-protecting multipaths – a simple and resource-efficient protection switching mechanism for MPLS networks," in *NETWORKING 2004. Lecture Notes in Computer Science*, vol. 3042, 2004.
- [139] M. Menth, R. Martin, and U. Sporlein, "Failure-specific self-protecting multipaths — increased capacity savings or overengineering?" in *2007 6th International Workshop on Design and Reliable Communication Networks*, 2007, pp. 1–7.
- [140] —, "Optimization of the self-protecting multipath for deployment in legacy networks," in *2007 IEEE International Conference on Communications*, 2007, pp. 421–427.
- [141] 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>
- [142] R. Callon, "TCP and UDP with bigger addresses (TUBA), a simple proposal for internet addressing and routing," Internet Requests for Comments, RFC Editor, RFC 1347, Jun 1992. [Online]. Available: <https://tools.ietf.org/html/rfc1347>
- [143] K. W. Kwong, L. Gao, R. Guerin, and Z. L. Zhang, "On the feasibility and efficacy of protection routing in IP networks," in *2010 Proceedings IEEE INFOCOM*, March 2010, pp. 1–9.
- [144] —, "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.
- [145] G. Enyedi, A. Csaszar, A. Atlas, C. Bowers, and A. Gopalan, "An Algorithm for Computing IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)," IETF, RFC 7811, Jun. 2016. [Online]. Available: <http://tools.ietf.org/html/rfc7811>
- [146] 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>
- [147] T. Čičić, A. F. Hansen, and O. K. Apeland, "Redundant trees for fast IP recovery," in *Broadnets*, 2007, pp. 152–159.
- [148] C. Alaettinoglu and V. Jacobson, "Towards Milli-Second IGP Convergence," Internet Engineering Task Force, Internet-Draft draft-alaettinoglu-isis-convergence-00, Nov. 2000, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-alaettinoglu-isis-convergence-00>
- [149] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb, "A case study of OSPF behavior in a large enterprise network," in *Proc. ACM IMW*, 2002.
- [150] J. Papan, P. Segec, and P. Paluch, "Utilization of PIM-DM in IP fast reroute," in *2014 IEEE 12th IEEE International Conference on Emerging eLearning Technologies and Applications (ICETA)*, Dec 2014, pp. 373–378.
- [151] G. Schollmeier, J. Charzinski, A. Kirstadter, C. Reichert, K. J. Schrodi, Y. Glickman, and C. Winkler, "Improving the resilience in IP networks," in *Workshop on High Performance Switching and Routing, 2003. HPSR.*, Jun. 2003, pp. 91–96.
- [152] 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>
- [153] B. Yang, J. Liu, S. Shenker, J. Li, and K. Zheng, "Keep Forwarding: Towards k-link failure resilient routing," in *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, Apr. 2014, pp. 1617–1625.
- [154] M. Chiesa, I. Nikolaevskiy, S. Mitrovic, A. Gurtov, A. Madry, M. Schapira, and S. Shenker, "On the resiliency of static forwarding tables," *IEEE/ACM Transactions on Networking*, vol. 25, no. 2, pp. 1133–1146, 2017.
- [155] M. Chiesa, A. Gurtov, A. Madry, S. Mitrovic, I. Nikolaevskiy, M. Schapira, and S. Shenker, "On the resiliency of randomized routing against multiple edge failures," in *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, 2016.
- [156] M. Chiesa, I. Nikolaevskiy, S. Mitrović, A. Panda, A. Gurtov, A. Madry, M. Schapira, and S. Shenker, "The Quest for Resilient (Static) Forwarding Tables," in *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, April 2016, pp. 1–9.
- [157] 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.
- [158] —, "On providing fast protection with remote loop-free alternates," *Telecommunication Systems*, vol. 60, no. 4, pp. 485–502, Dec 2015. [Online]. Available: <https://doi.org/10.1007/s11235-015-0006-9>
- [159] P. Francois and O. Bonaventure, "An evaluation of IP-based fast reroute techniques," in *Proceedings of the 2005 ACM International Conference on emerging Networking EXperiments and Technologies*. ACM, 2005, pp. 244–245.
- [160] M. Shand and S. Bryant, "A framework for loop-free convergence," RFC 5715, January 2010.
- [161] R. Teixeira and J. Rexford, "Managing routing disruptions in internet service provider networks," *Comm. Mag.*, vol. 44, no. 3, pp. 160–165, Mar. 2006. [Online]. Available: <http://dx.doi.org/10.1109/MCOM.2006.1607880>

- [162] F. Clad, P. Merindol, J.-J. Pansiot, P. Francois, and O. Bonaventure, "Graceful convergence in link-state IP networks: A lightweight algorithm ensuring minimal operational impact," *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–1, 2013.
- [163] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot, "Characterization of failures in an IP backbone," in *Proc. IEEE INFOCOM*, 2004.
- [164] 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.
- [165] 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.
- [166] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne, "Fast recovery from link failures using resilient routing layers," in *10th IEEE Symposium on Computers and Communications (ISCC'05)*, June 2005, pp. 554–560.
- [167] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne, "Fast IP network recovery using multiple routing configurations," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, April 2006, pp. 1–11.
- [168] J. Wang and S. Nelakuditi, "IP Fast Reroute with Failure Inference," 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>
- [169] J. Tapolcai and G. Rétvári, "Router virtualization for improving ip-level resilience," in *2013 Proceedings IEEE INFOCOM*, April 2013, pp. 935–943.
- [170] 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>
- [171] J. Tapolcai, G. Rétvári, P. Babarcsi, and E. R. Bérczi-Kovács, "Scalable and efficient multipath routing via redundant trees," *IEEE Journal on Selected Areas in Communications*, pp. 1–1, 2019.
- [172] 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>
- [173] Cisco Systems, "Cisco IOS XR Routing Configuration Guide for the Cisco CRS Router, Release 4.2," 2012.
- [174] Hewlett-Packard, "HP 6600 Router Series: QuickSpecs," 2008, available online: http://h18000.www1.hp.com/products/quickspecs/13811_na/13811_na.PDF.
- [175] Juniper Networks, "JUNOS 12.3 Routing protocols configuration guide," 2012.
- [176] 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.
- [177] 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. Ehtle, and E. P. Rathgeb, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 78–90.
- [178] 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/>
- [179] 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>
- [180] 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>
- [181] M. Hartmann, D. Hock, and M. Menth, "Routing optimization for IP networks with loop-free alternates," *Computer Networks*, vol. 95, pp. 35 – 50, 2016.
- [182] S. Litkowski, B. Decraene, C. Filsfils, and P. Francois, "Micro-loop Prevention by Introducing a Local Convergence Delay," Internet Requests for Comments, RFC Editor, RFC 8333, Mar 2018. [Online]. Available: <https://tools.ietf.org/html/rfc8333>
- [183] G. Enyedi, G. Rétvári, and T. Cinkler, "A novel loop-free IP fast reroute algorithm," in *Dependable and Adaptable Networks and Services*, A. Pras and M. van Sinderen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 111–119.
- [184] 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>
- [185] 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-urn-03.pdf>
- [186] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704, Mar. 2004. [Online]. Available: <https://tools.ietf.org/html/rfc3704>
- [187] B. Zhang, J. Wu, and J. Bi, "RPPF: 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.
- [188] P. Francois, "Improving the convergence of IP routing protocols," Ph.D. Thesis, biblio.info.ucl.ac.be/2007/457147.pdf, 2007.
- [189] C. Filsfils, S. Previdi, L. Ginsberg, B. Decraene, S. Litkowski, and R. Shakir, "Segment Routing Architecture," Internet Requests for Comments, RFC Editor, RFC 8402, Jul 2018. [Online]. Available: <https://tools.ietf.org/html/rfc8402>
- [190] 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>
- [191] S. Litkowski *et al.*, "Topology Independent Fast Reroute using Segment Routing," Internet Engineering Task Force, Internet-Draft draft-ietf-rtgwg-segment-routing-ti-lfa-03, Mar. 2020, work in Progress. [Online]. Available: <https://tools.ietf.org/html/draft-ietf-rtgwg-segment-routing-ti-lfa-03>
- [192] G. Enyedi, G. Rétvári, P. Szilágyi, and A. Császár, "IP Fast ReRoute: Lightweight Not-Via," in *NETWORKING 2009*, L. Fratta, H. Schulzrinne, Y. Takahashi, and O. Spaniol, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 157–168.
- [193] 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.
- [194] G. Enyedi, "Novel algorithms for IP fast reroute," Ph.D. Thesis, 2011.
- [195] M. Menth, M. Hartmann, R. Martin, T. Čičić, and A. Kvalbein, "Loop-free alternates and not-via addresses: A proper combination for IP fast reroute?" *Computer Networks*, vol. 54, no. 8, pp. 1300 – 1315, 2010, resilient and Survivable networks. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128609003491>
- [196] J. Papán, P. Segeč, and P. Palúch, "Tunnels in IP fast reroute," in *The 10th International Conference on Digital Technologies 2014*, July 2014, pp. 270–274.
- [197] 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>
- [198] 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>
- [199] 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.
- [200] 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.
- [201] 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.
- [202] G. Enyedi and G. Rétvári, "A Loop-Free Interface-Based Fast Reroute Technique," in *2008 Next Generation Internet Networks*, April 2008, pp. 39–44.

- [203] 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.
- [204] 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>
- [205] 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.
- [206] 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.
- [207] 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.
- [208] I. Theiss and O. Lysne, "Froots – fault handling in up*/down* routed networks with multiple roots," in *High Performance Computing - HiPC 2003*, T. M. Pinkston and V. K. Prasanna, Eds., 2003, pp. 106–117.
- [209] 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.
- [210] T. A. Kumar and M. H. M. K. Prasad, "Enhanced multiple routing configurations for fast IP network recovery from multiple failures," *CoRR*, vol. abs/1212.0311, 2012. [Online]. Available: <http://arxiv.org/abs/1212.0311>
- [211] T. Cicic, A. F. Hansen, A. Kvalbein, M. Hartmann, R. Martin, M. Menth, S. Gjessing, and O. Lysne, "Relaxed multiple routing configurations: IP fast reroute for single and correlated failures," *IEEE Transactions on Network and Service Management*, vol. 6, no. 1, pp. 1–14, March 2009.
- [212] S. Cho, T. Elhourani, and S. Ramasubramanian, "Independent directed acyclic graphs for resilient multipath routing," *IEEE/ACM Transactions on Networking*, vol. 20, no. 1, pp. 153–162, Feb 2012.
- [213] M. Nagy, J. Topolcai, and G. Rétvári, "Node Virtualization for IP Level Resilience," *IEEE/ACM Transactions on Networking*, vol. 26, no. 3, pp. 1250–1263, June 2018.
- [214] M. Menth and R. Martin, "Network resilience through multi-topology routing," in *DRCN 2005). Proceedings.5th International Workshop on Design of Reliable Communication Networks, 2005.*, 2005, pp. 271–277.
- [215] T. C. Cicic, A. F. Hansen, A. Kvalbein, M. Hartman, R. Martin, and M. Menth, "Relaxed multiple routing configurations for IP fast reroute," in *NOMS 2008 - 2008 IEEE Network Operations and Management Symposium*, 2008, pp. 457–464.
- [216] P. Psenak, S. Mirtorabi, A. Roy, L. Nguyen, and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF," IETF, RFC 4915, Jun. 2007. [Online]. Available: <http://tools.ietf.org/rfc/rfc4915.txt>
- [217] A. Atlas, K. Tiruveedhula, C. Bowers, J. Tantsura, and I. Wijnands, "LDP Extensions to Support Maximally Redundant Trees," IETF, RFC 8320, Feb. 2018. [Online]. Available: <http://tools.ietf.org/rfc/rfc8320>
- [218] 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.
- [219] —, "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.
- [220] 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.
- [221] E. Palmer, "On the spanning tree packing number of a graph: a survey," *Discrete Mathematics*, vol. 230, no. 1, pp. 13 – 21, 2001.
- [222] M. Menth and W. Braun, "Performance comparison of not-via addresses and maximally redundant trees (MRTs)," in *2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, 2013, pp. 218–225.
- [223] Cisco Systems, "IP Routing: OSPF Configuration Guide, Cisco IOS Release 15.2S - OSPF IPv4 Remote Loop-Free Alternate IP Fast Reroute," downloaded: Apr. 2012.
- [224] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed internet routing convergence," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '00. New York, NY, USA: ACM, 2000, pp. 175–187. [Online]. Available: <http://doi.acm.org/10.1145/347059.347428>
- [225] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating internet routing instabilities," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 4, pp. 205–218, Aug. 2004. [Online]. Available: <http://doi.acm.org/10.1145/1030194.1015491>
- [226] J. Chandrashekar, Z. Duan, Z. Zhang, and J. Krasky, "Limiting path exploration in BGP," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies, 13-17 March 2005, Miami, FL, USA.* IEEE, 2005, pp. 2337–2348. [Online]. Available: <https://doi.org/10.1109/INFCOM.2005.1498520>
- [227] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs, "R-BGP: Staying connected in a connected world," in *Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation*, ser. NSDI'07. Berkeley, CA, USA: USENIX Association, 2007, pp. 25–25. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1973430.1973455>
- [228] N. Gvozdiev, B. Karp, and M. Handley, "LOUP: The principles and practice of intra-domain route dissemination," in *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*, ser. nsdi'13. USA: USENIX Association, 2013, p. 413–426.
- [229] 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, SIGCOMM 2017, Los Angeles, CA, USA, August 21-25, 2017*, 2017, pp. 460–473. [Online]. Available: <https://doi.org/10.1145/3098822.3098856>
- [230] T. Holterbach, E. C. Molero, M. Apostolaki, A. Dainotti, S. Vissicchio, and L. Vanbever, "Blink: Fast connectivity recovery entirely in the data plane," in *16th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2019, Boston, MA, February 26-28, 2019*, 2019, pp. 161–176. [Online]. Available: <https://www.usenix.org/conference/nsdi19/presentation/holterbach>
- [231] N. Feamster, D. G. Andersen, H. Balakrishnan, and M. F. Kaashoek, "Measuring the effects of internet path faults on reactive routing," *SIGMETRICS Perform. Eval. Rev.*, vol. 31, no. 1, pp. 126–137, Jun. 2003. [Online]. Available: <http://doi.acm.org/10.1145/885651.781043>
- [232] B. Premore, "An experimental analysis of BGP convergence time," in *Proceedings of the Ninth International Conference on Network Protocols*, ser. ICNP '01. Washington, DC, USA: IEEE Computer Society, 2001, pp. 53–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=876907.881566>
- [233] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan, "BGP Beacons," in *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '03. New York, NY, USA: ACM, 2003, pp. 1–14. [Online]. Available: <http://doi.acm.org/10.1145/948205.948207>
- [234] T. Griffin and G. T. Wilfong, "On the correctness of IBGP configuration," in *Proceedings of the ACM SIGCOMM 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, August 19-23, 2002, Pittsburgh, PA, USA*, 2002, pp. 17–29. [Online]. Available: <https://doi.org/10.1145/633025.633028>
- [235] M. Caesar, L. Subramanian, and R. H. Katz, "Towards localizing root causes of BGP dynamics," EECS Department, University of California, Berkeley, Tech. Rep. UCB/CSD-03-1292, 2003. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2003/6364.html>
- [236] U. Javed, I. Cunha, D. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy, "Poiroot: Investigating the root cause of interdomain path changes," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 183–194. [Online]. Available: <http://doi.acm.org/10.1145/2486001.2486036>
- [237] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat, "Taking the edge off with espresso: Scale, reliability and programmability for global internet peering," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '17. New York, NY, USA: ACM, 2017, pp. 432–445. [Online]. Available: <http://doi.acm.org/10.1145/3098822.3098854>

- [238] O. Bonaventure, C. Filsfil, and P. Francois, "Achieving sub-50 milliseconds recovery upon BGP peering link failures," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1123–1135, Oct. 2007. [Online]. Available: <http://dx.doi.org/10.1109/TNET.2007.906045>
- [239] M. Kopka, "IP Routing Fast Convergence," 2013, available online: https://www.cisco.com/c/dam/global/cs_cz/assets/ciscoconnect/2013/pdf/T-SP4-IP_Routing_Fast_Convergence-Miloslav_Kopka.pdf.
- [240] C. Filsfil, "BGP Convergence in Much Less than a Second," 2007, available online: <http://newnog.net/meetings/nanog40/presentations/ClarenceFilsfil-BGP.pdf>.
- [241] D. Clark, J. Rexford, and A. Vahdat, "A purpose-built global network: Google's move to SDN," *Communications of the ACM*, 2016.
- [242] 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.
- [243] O. N. Foundation, "OpenFlow Switch Specification, Version 1.1.0 Implemented (Wire Protocol 0x02)," ONF TS-002, Feb. 2011. [Online]. Available: <https://www.opennetworking.org/wp-content/uploads/2014/10/openflow-spec-v1.1.0.pdf>
- [244] A. Sgambelluri, A. Giorgetti, F. Cugini, F. Paolucci, and P. Castoldi, "OpenFlow-based segment protection in Ethernet networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 5, no. 9, pp. 1066–1075, Sep. 2013.
- [245] J. Kempf, E. Bellagamba, A. Kern, D. Jocha, A. Takács, and P. Sköldström, "Scalable fault management for OpenFlow," in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 6606–6610.
- [246] R. M. Ramos, M. Martinello, and C. E. Rothenberg, "Slickflow: Resilient source routing in data center networks unlocked by OpenFlow," in *Proc. 38th annual IEEE conference on local computer networks (LCN)*. IEEE, 2013, pp. 606–613.
- [247] B. Stephens, A. L. Cox, and S. Rixner, "Plinko: Building provably resilient forwarding tables," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, ser. HotNets-XII. New York, NY, USA: ACM, 2013, pp. 26:1–26:7.
- [248] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming protocol-independent packet processors," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 87–95, Jul. 2014.
- [249] M. Borokhovich, L. Schiff, and S. Schmid, "Provable Data Plane Connectivity with Local Fast Failover: Introducing Openflow Graph Algorithms," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 121–126.
- [250] 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*, Sep. 2014, pp. 61–66.
- [251] C. Cascone, L. Pollini, D. Sanvito, A. Capone, and B. Sanso, "SPIDER: Fault resilient SDN pipeline with recovery delay guarantees," in *2016 IEEE NetSoft Conference and Workshops (NetSoft)*. IEEE, 2016, pp. 296–302.
- [252] B. Stephens, A. L. Cox, and S. Rixner, "Scalable multi-failure fast failover via forwarding table compression," in *Proceedings of the Symposium on SDN Research*, ser. SOSR '16. New York, NY, USA: ACM, 2016, pp. 9:1–9:12.
- [253] D. Merling, W. Braun, and M. Menth, "Efficient data plane protection for SDN," in *IEEE NetSoft*, 2018.
- [254] R. Sedar, M. Borokhovich, M. Chiesa, G. Antichi, and S. Schmid, "Supporting emerging applications with low-latency failover in P4," in *Proceedings of the 2018 Workshop on Networking for Emerging Applications and Technologies*, ser. NEAT '18. New York, NY, USA: ACM, 2018, pp. 52–57.
- [255] M. Menth, M. Schmidt, D. Reutter, R. Finze, S. Neuner, and T. Kleefass, "Resilient integration of distributed high-performance zones into the belvue network using openflow," *IEEE Communications Magazine*, vol. 55, no. 4, pp. 94–99, 2017. [Online]. Available: <https://doi.org/10.1109/MCOM.2017.1600177>
- [256] M. Reitblatt, M. Canini, A. Guha, and N. Foster, "Fattire: Declarative loop-free tolerance for software-defined networks," in *Proc. 2nd ACM SIGCOMM workshop on Hot topics in software defined networking*. ACM, 2013, pp. 109–114.
- [257] E. M. Gafni and D. P. Bertsekas, "Distributed algorithms for generating loop-free routes in networks with frequently changing topology," *IEEE Transactions on Communications*, vol. 29, pp. 11–18, 1981.
- [258] M. Chiesa, R. Sedar, G. Antichi, M. Borokhovich, A. Kamisiński, G. Nikolaidis, and S. Schmid, "PURR: A Primitive for Reconfigurable Fast Reroute: Hope for the Best and Program for the Worst," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, ser. CoNEXT '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–14. [Online]. Available: <https://doi.org/10.1145/3359989.3365410>
- [259] Y. Liu, D. Halperin, A. Krishnamurthy, and T. Anderson, "F10: A fault-tolerant engineered network," in *Proc. 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2013, pp. 399–412.
- [260] Y.-A. Pignolet, S. Schmid, and G. Tredan, "Load-optimal local fast rerouting for dependable networks," in *Proc. 47th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2017.
- [261] M. Borokhovich, Y.-A. Pignolet, G. Tredan, and S. Schmid, "Load-optimal local fast rerouting for dense networks," in *Proc. IEEE/ACM Transactions on Networking (ToN)*, 2018.
- [262] K.-T. Foerster, A. Kamisiński, Y.-A. Pignolet, S. Schmid, and G. Tredan, "Bonsai: Efficient fast failover routing using small arborescences," in *Proc. 49th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2019.
- [263] —, "Improved fast rerouting using postprocessing," in *38th International Symposium on Reliable Distributed Systems (SRDS)*, 2019.
- [264] K.-T. Foerster, Y.-A. Pignolet, S. Schmid, and G. Tredan, "Local fast failover routing with low stretch," in *Proc. ACM SIGCOMM Computer Communication Review (CCR)*, 2018.
- [265] M. Borokhovich and S. Schmid, "How (not) to shoot in your foot with SDN local fast failover: A load-connectivity tradeoff," in *Proc. 17th International Conference on Principles of Distributed Systems (OPODIS)*, December 2013.
- [266] M. Chiesa, A. Gurtov, A. Madry, S. Mitrovic, I. Nikolaevskiy, M. Shapira, and S. Shenker, "On the Resiliency of Randomized Routing Against Multiple Edge Failures," in *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi, Eds., vol. 55. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016, pp. 134:1–134:15. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2016/6269>
- [267] M. Chiesa, I. Nikolaevskiy, S. Mitrovic, A. Gurtov, A. Madry, M. Shapira, and S. Shenker, "On the Resiliency of Static Forwarding Tables," *IEEE/ACM Trans. Netw.*, vol. 25, no. 2, p. 1133–1146, Apr. 2017. [Online]. Available: <https://doi.org/10.1109/TNET.2016.2619398>
- [268] M. Chiesa, A. Gurtov, A. Madry, S. Mitrovic, I. Nikolaevskiy, A. Panda, M. Shapira, and S. Shenker, "Exploring the limits of static failover routing," *CoRR*, vol. abs/1409.0034v4, 2014. [Online]. Available: <https://arxiv.org/abs/1409.0034v4>
- [269] K.-T. Foerster, Y.-A. Pignolet, S. Schmid, and G. Tredan, "Local fast failover routing with low stretch," in *Proc. ACM SIGCOMM Computer Communication Review (CCR)*, 2018.
- [270] —, "Casa: Congestion and stretch aware static fast rerouting," in *Proc. IEEE INFOCOM*, 2019.
- [271] H. Villför, "Operator experience from ISIS convergence tuning," 2004, presented at the RIPE 47 meeting. [Online]. Available: <https://meetings.ripe.net/ripe-47/presentations/ripe47-routing-isis.pdf>
- [272] P. Francois, M. Shand, and O. Bonaventure, "Disruption free topology reconfiguration in ospf networks," in *IEEE INFOCOM 2007-26th IEEE International Conference on Computer Communications*. IEEE, 2007, pp. 89–97.
- [273] M. Shand, S. Bryant, S. Previdi, C. Filsfil, P. Francois, and O. Bonaventure, "Framework for Loop-Free Convergence Using the Ordered Forwarding Information Base (oFIB) Approach," RFC 6976, Jul. 2013. [Online]. Available: <https://rfc-editor.org/rfc/rfc6976.txt>
- [274] K.-T. Foerster, S. Schmid, and S. Vissicchio, "Survey of consistent software-defined network updates," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1435–1461, 2018.
- [275] M. Markovitch and S. Schmid, "Shear: A highly available and flexible network architecture: Marrying distributed and logically centralized control planes," in *Proc. 23rd IEEE International Conference on Network Protocols (ICNP)*, 2015.
- [276] S. Vissicchio, L. Cittadini, O. Bonaventure, G. G. Xie, and L. Vanbever, "On the co-existence of distributed and centralized routing control planes," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 469–477.
- [277] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, pp. 756–767, 2002.
- [278] Z. Zhang and A. Baban, "Bit index explicit replication (bier) forwarding for network device components," Jul. 11 2017, uS Patent 9,705,784.

- [279] W. Braun, M. Albert, T. Eckert, and M. Menth, "Performance comparison of resilience mechanisms for stateless multicast using bier," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*. IEEE, 2017, pp. 230–238.
- [280] D. Merling, S. Lindner, and M. Menth, "Comparison of fast-reroute mechanisms for BIER-based IP multicast," *Proc. International Conference on Software Defined Systems*, 2020.

Marco Chiesa is an Assistant Professor at the KTH Royal Institute of Technology, Sweden. He received his Ph.D. degree in computer engineering from Roma Tre University in 2014. His research interests include Internet architectures and protocols, including aspects ranging from network design and optimization to security and privacy. He received the IEEE Communication Society William R. Bennett Prize in 2020, the IEEE ICNP Best Paper Award in 2013, and the Applied Network Research Prize in 2012. He has been a distinguished TPC member at

IEEE Infocom in 2019 and 2020.

Andrzej Kamisiński is an Assistant Professor in the Department of Telecommunications at the AGH University of Science and Technology in Kraków, Poland. He received his B.Sc., M.Sc., and Ph.D. degrees from the same University in 2012, 2013, and 2017, respectively. In 2015, Andrzej Kamisiński was a Visiting Ph.D. Student at NTNU (Trondheim, Norway) where he worked with Prof. Bjarne E. Helvik and with Telenor Research on dependability of Software-Defined Networks. In summer 2018, he was a Visiting Research Fellow in the Commu-

cation Technologies group led by Prof. Stefan Schmid at the Faculty of Computer Science, University of Vienna, Austria. Between 2018 and 2020, he was a member of the Management Committee of the *Resilient Communication Services Protecting End-User Applications From Disaster-Based Failures* European COST Action, and in 2020, a Research Associate in the Networked Systems Research Laboratory at the School of Computing Science, University of Glasgow, Scotland. His research interests span dependability and security of computer and communication networks.

Jacek Rak (M'08, SM'13) is an Associate Professor and the Chair of Department of Computer Communications at Gdańsk University of Technology, Gdańsk, Poland. He received his MSc, PhD and DSc (habilitation) degrees from the same university in 2003, 2009, and 2016, accordingly. He has authored over 100 publications, including the book *Resilient Routing in Communication Networks* (Springer, 2015). Between 2016 and 2020 he was leading the COST CA15127 Action *Resilient Communication Services Protecting End-user Applications from Disaster-based Failures* (RECODIS) involving over 170 members from 31 countries. He has also served as a TPC member of numerous conferences and journals. Recently, he has been the General Chair of ITS-T'17 and MMM-ACNS'17, the General Co-Chair of NETWORKS'16, the TPC Chair of ONDM'17, and the TPC Co-chair of IFIP Networking'19. Prof. Rak is the Member of the Editorial Board of Optical Switching and Networking, Elsevier and the founder of the International Workshop on Resilient Networks Design and Modeling (RNDM). His main research interests include the resilience of communication networks and networked systems.

Gábor Rétvári received the M.Sc. and Ph.D. degrees in electrical engineering from the Budapest University of Technology and Economics in 1999 and 2007. He is now a Senior Research Fellow at the Department of Telecommunications and Media Informatics. His research interests include all aspects of network routing and switching, the programmable data plane, and the networking applications of computational geometry and information theory. He maintains several open source scientific tools written in Perl, C, and Haskell.

Stefan Schmid is a Professor at the Faculty of Computer Science at the University of Vienna, Austria. He received his MSc (2004) and PhD degrees (2008) from ETH Zurich, Switzerland. In 2009, Stefan Schmid was a postdoc at TU Munich and the University of Paderborn, between 2009 and 2015, a senior research scientist at the Telekom Innovations Laboratories (T-Labs) in Berlin, Germany, and from the end of 2015 till early 2018, an Associate Professor at Aalborg University, Denmark. His research interests revolve around fundamental and algorithmic problems arising in networked and distributed systems. He is currently leading the ERC Consolidator project *AdjustNet*.