A comprehensive survey on internet outages

Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, Antonio Pescapé

PII: S1084-8045(18)30113-9

DOI: 10.1016/j.jnca.2018.03.026

Reference: YJNCA 2107

To appear in: *Journal of Network and Computer Applications*

Received Date: 10 January 2018

Revised Date: 6 March 2018

Accepted Date: 22 March 2018

A Comprehensive Survey on Internet Outages

Giuseppe Aceto[a,b], Alessio Botta[a,b], Pietro Marchetta[a], Valerio Persico[a,b], Antonio Pescapé[a,b]

[a] *University of Napoli "Federico II" (Italy)*
[b] *NM2 s.r.l. (Italy)*

Abstract

Internet outages are inevitable, frequent, opaque, and expensive. To make things worse, they are poorly understood, while a deep understanding of them is essential for strengthening the role of the Internet as the world's communication substrate. The importance of research on Internet outages is demonstrated by the large body of literature focusing on this topic. Unfortunately, we have found this literature rather scattered, since many different and equally important aspects can be investigated, and researchers typically focused only on a subset of them. And, to the best of out knowledge, no paper in literature provides an extensive view on this important research topic. To fill this gap, we analyze all the relevant facets of this important research topic, stepping from the critical review of the available literature. Our work sheds light on several obscure aspects such as, for example, the different challenges considered in the literature, the techniques, tools, and methodologies used, the contributions provided towards different goals (e.g., outage analysis and detection, impact evaluation, risk assessment, countermeasures, etc.), the issues that are still open, etc.. Moreover, it provides several innovative contributions achieved analyzing the wide and scattered literature on Internet outages (e.g., characterization of the main causes of outages, general approach for implementing outages detection systems, systematic classification of definitions and metrics for network resilience, etc.). We believe that this work represents an important and missing starting point for academy and industry to understand and contribute to this wide and articulate research area.

Keywords: Outage, Large-scale Outages, Internet Outages, Fault, Detection, Resilience, Earthquake, Tsunami, Hurricane, Cable Cut, DDoS, Network Attack, Security, Outage Impact, Mitigation, Risk Assessment, Survey.

1. Introduction and Motivation

The professional, personal, and political lives of almost two billion users worldwide now critically depend on the Internet. Financial transactions, business operations, and many other applications require high availability and good performance of this critical, highly dynamic, extremely heterogeneous, planetary-scale, and largely opaque ecosystem of networks. However, as any other critical infrastructure, this one trillion-dollar communication system experiences outages.

Internet outages are *inevitable*, *frequent*, *opaque*, *expensive*, and *poorly understood.* They are *inevitable* because the perfect system is not achievable in practice since issues and threats can not be completely prevented, or their prevention can be economically unfeasible. Outages are *frequent*: in just three weeks of monitoring, Katz-Bassett et al. [132] discovered persistent reachability problems involving about $10,000$ distinct prefixes, with one in five of the problems lasting for more than 10 hours. Large-scale Internet outages are also continuously reported in Renesys blog [17] and the outage mailing list [14]. They are also *opaque* since the Internet has good built-in abilities (e.g., IP routing) to limit their impact and thus visibility and traceability. Internet outages are *expensive*: according to Forbes [68], Amazon lost about $66,240 dollar per minute on 19 August 2013 due to a blackout. More in general, outages preventing users to con-

Email addresses: giuseppe.aceto@unina.it (Giuseppe Aceto), a.botta@unina.it (Alessio Botta), pietro.marchetta@unina.it (Pietro Marchetta), valerio.persico@unina.it (Valerio Persico), pescape@unina.it (Antonio Pescapé)

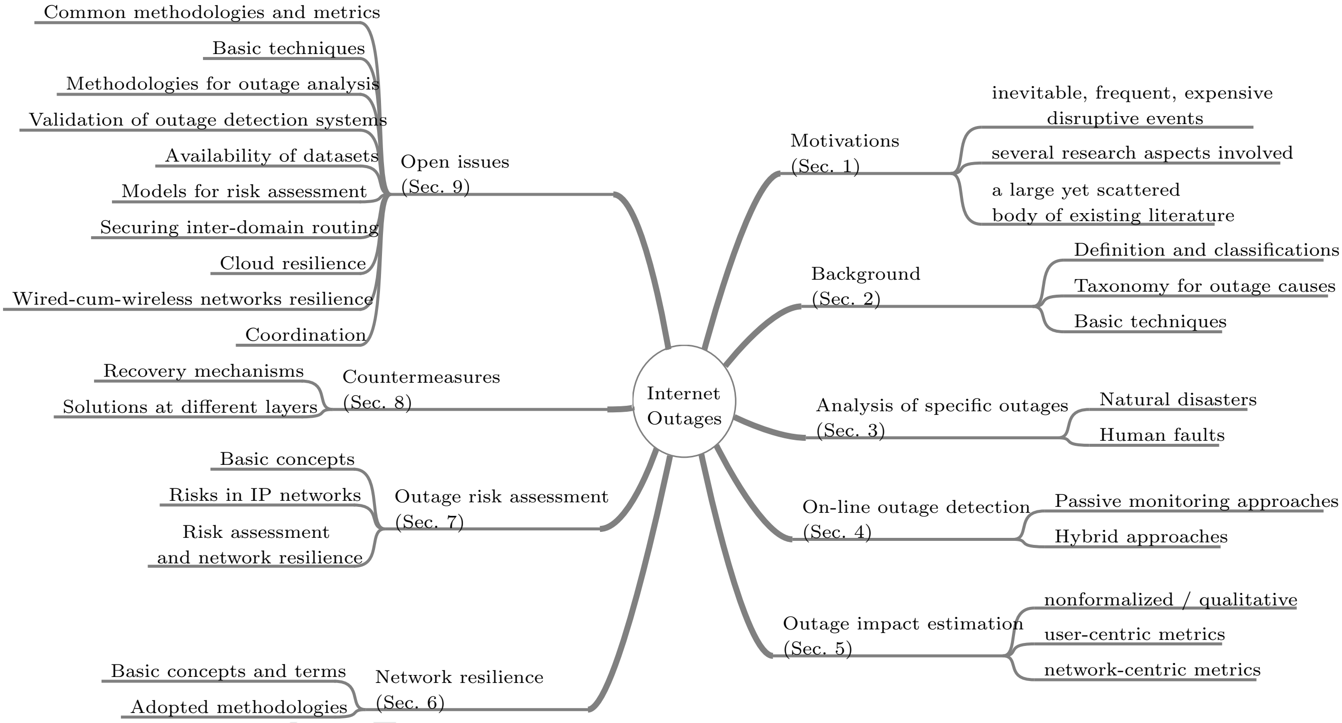
Internet Outages
Motivations (Sec. 1)
inevitable, frequent, expensive disruptive events
several research aspects involved
a large yet scattered body of existing literature
Background (Sec. 2)
Definition and classifications
Taxonomy for outage causes
Basic techniques
Analysis of specific outages (Sec. 3)
Natural disasters
Human faults
On-line outage detection (Sec. 4)
Passive monitoring approaches
Hybrid approaches
Outage impact estimation (Sec. 5)
nonformalized / qualitative
user-centric metrics
network-centric metrics
Network resilience (Sec. 6)
Basic concepts and terms
Adopted methodologies
Outage risk assessment (Sec. 7)
Basic concepts
Risks in IP networks
Risk assessment and network resilience
Countermeasures (Sec. 8)
Recovery mechanisms
Solutions at different layers
Open issues (Sec. 9)
Common methodologies and metrics
Basic techniques
Methodologies for outage analysis
Validation of outage detection systems
Availability of datasets
Models for risk assessment
Securing inter-domain routing
Cloud resilience
Wired-cum-wireless networks resilience
Coordination

Figure 1: Internet outages: an overview of the topics and contributions addressed in this paper.

nect to data-centers cost on average about $5,000 per minute [191]. Finally, our understanding of Internet outages is severely weakened by the lack of data, information and collaboration from the network operators for which network outages represent a sensitive topic critical to their business.

Internet outages may happen for a number of reasons including (i) natural disasters, such as earthquakes and hurricanes; (ii) software attacks, such as worms or prefix hijacking; (iii) physical attacks, such as military attacks, terrorism, or electromagnetic pulse attacks; (iv) accidental misconfiguration; (v) certain forms of censorship such as the one adopted during the Arab spring; (vi) software bugs (e.g., in the routers); (vii) network equipment hardware failures, and many other factors. We thoroughly discuss the literature on network outages in relation with those causes in Section 2.1.

Although the Internet proved to be relatively robust to localized disruptions, Internet outages can still leave large sections of the population without network access for short or large time periods depending on their extent [91]. A deep understanding of Internet outages is essential for strengthening the role of this infrastructure as the world's communication substrate. This is increasingly important as the Internet is becoming the foundation of new physical-world-related applications as *Cyber-Physical Systems* and the *Internet of Things* [26].

To determine a positive future evolution of this communication system, researchers and network operators need to clearly understand how to *prevent*, *detect*, and *mitigate* Internet outages. Preventing outages in the network is the best option to guarantee high availability to the services and applications relying on the Internet. At the same time, fast and accurate detection of an ongoing outage is the essential preliminary step to trigger effective countermeasures whose primary goal is to mitigate as much as possible the impact of the outage as perceived by the final users. All these imperative operations, however, require a deep understanding of Internet outages. Questions including "*why, when, and where do these events occur? what is the expected impact? how are they likely to happen?*" call for the development of theoretical and practical instruments.

The importance of this topic is demonstrated by the large body of literature focusing on Internet outages. We have found that this literature is rather scattered since Internet outages include many different and equally important aspects while researchers typically focused only on a subset of them. Previous works have surveyed this vast topic focusing on specific aspects or viewpoints, failing to provide an overall picture that could inform researchers and practitioners that are new to the topic, or want to expand their specialistic knowledge on Internet outages to other aspects, or simply look for new possible applications for their expertise. We refer to these works in the relevant sections of our survey. With this survey, we provide the research community with a comprehensive view on Internet outages discussing all the relevant aspects, offering a reference starting point for researchers willing to understand and/or contribute to this wide and articulate research area. We also consider an up-to date list of notable Internet outages, on a time span of 17 years, testifying the variety and the frequency (almost one per year) of major disruptions of access to Internet services; Figure 2 reports the timeline of notable outages, that are discussed in Section 3.

1.1. Challenges

Internet outages pose to the research community a number of challenges including the ones we discuss in the following and addressed in this paper.

1.1.1. How to analyze Internet outages

Large-scale disruptive events such as earthquakes or hurricanes have a terrible destructive power. Understanding the effects of these events when referring to roads, buildings, or human beings is quite straightforward. On the other hand, assessing their effects on the Internet infrastructure is much less obvious. Indeed, even in case of link and node failures, the routing might be able to automatically find a new stable configuration, guaranteeing good connections between any pair of nodes in the network. However, this may happen or not depending on whether the underlying physical topology does allow it, while the new stable configuration of the network may not be the optimal one. Understanding how these automatic processes take place when Internet outages arise is important to gather the essential knowledge to (i) model similar events; (ii) clarify how the network reacts in these cases; (iii) recognize ongoing Internet outages; (iv) develop effective solutions.

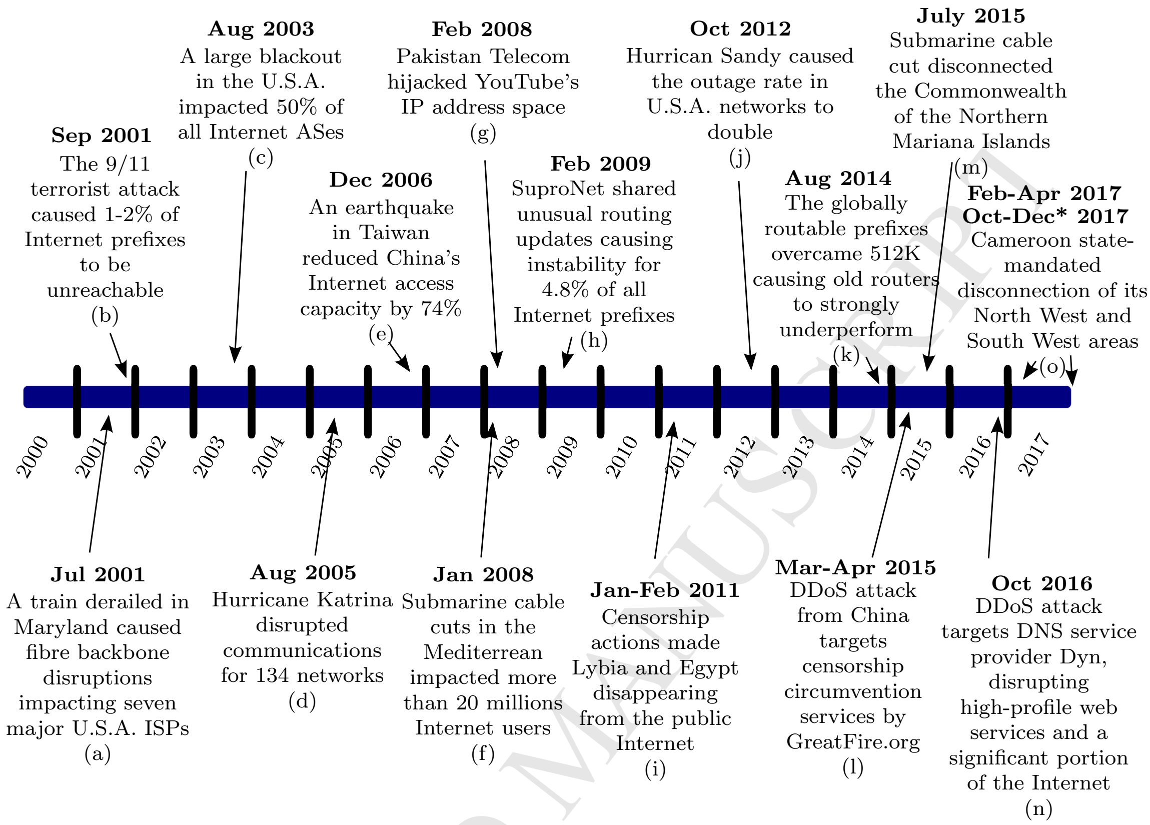
Sep 2001 The 9/11 terrorist attack caused 1-2% of Internet prefixes to be unreachable (b)
Aug 2003 A large blackout in the U.S.A. impacted 50% of all Internet ASes (c)
Dec 2006 An earthquake in Taiwan reduced China's Internet access capacity by 74% (e)
Feb 2008 Pakistan Telecom hijacked YouTube's IP address space (g)
Feb 2009 SuproNet shared unusual routing updates causing instability for 4.8% of all Internet prefixes (h)
Oct 2012 Hurrican Sandy caused the outage rate in U.S.A. networks to double (j)
Aug 2014 The globally routable prefixes overcame 512K causing old routers to strongly underperform (k)
July 2015 Submarine cable cut disconnected the Commonwealth of the Northern Mariana Islands (m)
Feb-Apr 2017 Oct-Dec* 2017 Cameroon state-mandated disconnection of its North West and South West areas (o)
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
Jul 2001 A train derailed in Maryland caused fibre backbone disruptions impacting seven major U.S.A. ISPs (a)
Aug 2005 Hurricane Katrina disrupted communications for 134 networks (d)
Jan 2008 Submarine cable cuts in the Mediterrean impacted more than 20 millions Internet users (f)
Jan-Feb 2011 Censorship actions made Lybia and Egypt disappearing from the public Internet (i)
Mar-Apr 2015 DDoS attack from China targets censorship circumvention services by GreatFire.org (l)
Oct 2016 DDoS attack targets DNS service provider Dyn, disrupting high-profile web services and a significant portion of the Internet (n)

Figure 2: A timeline reporting concrete examples of some of the main and well known Internet outages. References (in chronological order): (a) [225], (b) [178], (c) [73], (d) [11], (e) [9], (f) [59], (g) [52], (h) [256], (i) [72, 71], (j) [74], (k) [221], (l) [165], (m) [27], (n) [159], (o) [29].
* Still ongoing at the time of writing.

1.1.2. How to detect Internet outages

Network operators need systems and systematic approaches to detect in (near) real-time ongoing Internet outages. These early-warning systems are essential to limit the impact of these type of events. The alarms are useful to estimate the frequency, the scope, and the root causes of network outages and trigger (predefined) effective countermeasures as well.

1.1.3. How to quantify the impact of an outage

When an outage affects the infrastructure, a widely accepted set of theoretical and practical instruments is needed to quantify the caused damage. In this way, we can compare different outage episodes, rank them, and focus on the most disruptive ones to learn how to deal with them in the future.

1.1.4. How to quantify network robustness to Internet outages

Any node or link of the network may fail at any time. Similar events may happen simply due to network equipment obsolescence. Thus, quantifying the vulnerability of the network to similar events is an important task also to plan new investments aiming at improving the network infrastructure.

1.1.5. How to assess the risk of disruptive Internet outages

As for any other critical infrastructure, assessing the risk of an outage of the Internet enables network operators to insure their infrastructure with private insurance companies. Unfortunately, the set of theoretical and practical instruments normally used to assess the risk of critical infrastructure cannot be applied as is to the Internet environment for several reasons. For instance, the Internet is still potentially able to perfectly deliver services in case of outages by effectively re-routing the traffic.

1.1.6. How to survive to and mitigate Internet outages

None of the networks part of the Internet is immune from outages. Even great investments in a given network do not prevent this network from being affected by outages occurring in other portions of the Internet. The obvious conclusion is that network operators have to deal with Internet outages. For this reason, we need systems and approaches to recover from outages, prevent them or mitigate as much as possible their impact on the network.

When facing these challenges, the research community must also deal with the complex operational climate of the Internet, where independent networks are forced to collaborate and compete. In this context, it is not surprising that network operators are reluctant to disclose detailed information on the managed infrastructures, thus strongly weakening our understanding of the Internet dynamics and evolution in general, and of Internet outages in particular.

1.2. Methodology and Contribution

In this paper we provide a survey on Internet outages analyzing the articulate state of the art in this field. According to the indications reported in [138], we adopt the research methodology described in the following (see Fig. 1 for an overall view).

- We start providing three main contributions. We identify a non-ambiguous definition of *Internet outage*. We propose a new characterization for the possible causes of these disrupting events learning from previous works classifying network outages. We describe the basic techniques commonly adopted in this research area. These important points are presented and discussed in Sec. 2.
- We critically review the literature on Internet outages, classifying the works according to the specific outage-related aspect they aimed at investigating or the basic principle adopted to reach this goal. We start by focusing on the analysis of specific episodes of large-scale Internet outages (Sec. 3). Then, we discuss systems or systematic approaches to detect these disrupting events in real-time or soft real-time (Sec. 4). Based on what we learned from these two steps, we then analyze the metrics and approaches used in literature to quantify the impact of an Internet outage (Sec 5). Successively, we focus our attention on papers proposing approaches or metrics to assess the robustness of the network (Sec. 6) or the risk associated with these events (Sec. 7). Finally, we examine the countermeasures proposed in literature to recover from Internet outages and prevent or mitigate their consequences (Sec. 8).
- The previous steps are then used as input to derive the open issues in the field of Internet outages (Sec. 9). We provide concluding remarks in Sec. 10.

To the best of our knowledge, this is the first comprehensive survey on Internet outages available in literature. We believe that this paper provides novel and elaborated contributions of interest for the research community, achieved analysing the wide and scattered literature on Internet outages (we have considered more than 210 related works), and sheds light on the current and future research issues in this field. We have proposed several categorisations and taxonomies for the main concepts related to Internet outages, elaborated views and thorough discussions on the main aspects of this important research topic.

2. Background

In this Section, we provide the reader with the necessary background to fully understand the main concepts related to Internet outages. More specifically, we introduce definitions and propose a characterization for the causes of outages. Finally, we discuss the basic techniques commonly used in Internet outage-related works and propose a classification for them as well.

2.1. Definition and classification of Internet outages

We define an Internet outage as *the particular condition in which the network lies when one or multiple network elements located in a specific geographic area either do not work properly or are not reachable due to intentional or accidental events.*

With this definition, we aim at underlying few important concepts. First of all, we want to emphasize the clear separation between the network outage (i.e. a suboptimal network condition) and its original cause (e.g., natural disasters, human errors, etc.): these two aspects are often confused in literature. As the relation between the occurrence of perturbing events (the causes) and the effect of the event on the network status (outage) is not trivial, the analysis and modeling of outages is severely limited when a clear distinction between events and status is not performed. An outage occurs every time the network deviates from the expected operational status. For instance, a network congested by legitimate traffic should not be considered as subjected to an outage since a similar event is somehow expected although undesirable. On the other hand, a network poorly performing due to the presence of malicious traffic (e.g., traffic generated by worm or network attacks) is subjected to an outage since this is a clear deviation from a normal operational status. Finally, network outages are typically well localised and their impact is often stronger in the proximity of the affected portion of the network. In the rest of the paper, we use the term network outages, failures, or disruptions interchangeably.

In literature, classification of network outages is performed based either on the consequences of the disrupting events, or on their causes. We report in Figure 3 the classifications of outages as defined in the papers that addressed outage classification at more general level; in the figure we include our classification, that is described in detail in Section 2.1. In the following we discuss the reported classifications.

One of the first analyses of outages for a massively distributed computerized system has been performed on Public Switched Telephone Network (PSTN) [141]. This work classifies outages based on their cause, and provides an estimation of the impact they had on customers (in terms of lack of service). Considered outages were reported by operators, according to U.S. Federal Communications Commission requirements [36]—i.e. their duration was of 30 minutes or more, or potentially affected more than 30,000 customers, or affected airports, 911 Service (emergency telephone number), nuclear power plants, major military installations, and key government facilities. Implicitly this operatively defined (major) outages based on their consequences, while the subsequent classification and analysis reported in [141] is focused on causes, as derived from the brief descriptions in the reports. Notably, sometimes the classification mixes root causes with intermediate ones, e.g. for the category *Acts of nature*, are considered as subcategories namely *Cable, Power supply, Facility* as damaged from burrowing animals or lightning, and *Natural disasters*, exemplified as *Earthquakes, Hurricanes, or Floods.* Moreover, *Overloads* are also reported as outage category, although considered "somewhat problematic" due to the very definition of outage based on affected customers: in the considered outage reports the number of customers are the served ones, not the ones rejected. We highlight that, specifically with regards to overloading, the Internet differs radically from PSTN: first, the packet-switched and datagram-oriented nature of the Internet, as opposed to circuit-switched PSTN, allows for a smooth degradation of performance instead of service rejection in case of overloading; second, the Internet has been designed as a best-effort service and is used as such by residential and most small enterprise customers, with no pre-allocation of resources before accepting (or rejecting) a communication. These circumstances, and the consideration that lack of prevision or support of user demand is mostly a business-related matter, more than a technical one, led us to exclude overloading (and congestion-related issues) *per se* from the outages, except when they are consequence of intentional action or other failures. We reported [141] for historical background and to introduce significant aspects (overloading, natural causes, issues in classifying outages); in the following we consider works explicitly aimed at Internet outage analysis and classification, and will compare the present work with them.

Wu et al. [247] provided an outage classification based on the impacted *physical* and *logical links.* A logical link is defined as the connection between Autonomous Systems (hereafter simply AS) made possible by several physical links. The authors classify network outages in the following six categories sorted by severity: partial peering teardown (a few but not all of the physical links between two ASes fail); AS partition (internal failure breaks an AS into a few isolated parts); depeering (discontinuation of a peer-to-peer relationship); teardown of access links (failure disconnects the customer from its provider); AS failure (an AS disrupts connection with all of its neighboring ASes); regional failure (failure causes reachability problem for many ASes in a region). Another work based on the consequences of an outage is [66], where a network reliability metric is proposed, based on Mean Time To first Failure and Mean Time To Recovery of devices and links. To this aim, the authors survey different statistical models of network failures derived from empirical data, stressing the general falseness of the assumption of independence between faults. The reasons causing multiple failures are classified in

- *structural*—common services or components, e.g. equipment shared among providers, physical infrastructure shared among carriers;
- *dynamic*—failure of a component or service increases the stress on other ones;
- *epistemic*—a failure is not observed until another occurs, hidden e.g. because of redundancy or automatic recovery mechanisms.

Markopoulou et al. [166] provided an outage classification from the point of view of the backbone service provider. Based on IS-IS protocol messages, simultaneity, and optical-to-IP layer mappings, the authors progressively narrow down the possible causes for the outages highlighted by connectivity status changes, and statistically characterize them according to occurrence frequency and time to repair. The final result is a classification of detected outages in six classes as follows

- maintenance—planned downtime of devices or links;
- router-related—including crash/reboot, linecard failure or reset, CPU overload, human misconfiguration;
- optical-related—including optical device failure and cable cuts;
- other multiple-links—time-overlapping failures with unspecified cause
- single-link high-frequency—including end-of-life deployments suffering from ageing, or prolonged testing/upgrade activities
- single-link low-frequency—single failures with unspecified cause

Finally, Çetinkaya et al. [58] presented a general categorization of network challenges (i.e. potential triggers of faults in networks, thus a concept wider than *causes* for outages). The considered categories are: *large-scale disasters* (e.g., earthquakes, hurricanes, pandemics), *socio-political and economic challenges* (e.g., terrorism, censorship), *dependent failures* (e.g., power shortages), *human errors* (e.g., misconfigurations), *malicious attacks* (e.g., prefix hijacking attacks), *unusual but legitimate traffic* (e.g., crowds looking for information on a breaking news), and *environmental challenges* (e.g., due to the mobility of nodes in an ad-hoc network). In their study, they analyze both spatial and temporal properties of these failures: the former are useful to quantify geographic distances that must be put between data centers so to guarantee that a certain outage is not likely to affect multiple data centers at a time, whereas the latter can be used to characterize recovery times.

Note how researchers adopted different points of view in their outage classification: Markopoulou et al. [166] and Çetinkaya et al. [58] used the causes of the Internet outages to classify them whereas Wu et al. [247] and Cholda et al. [66] focused on the consequences of the outage. If causes at the basis of the outage are ignored, any derived characterization or model is hardly generalizable to other networks or other time spans, as it either implicitly includes extraordinary events, altering the overall statistics, or neglects the impact of such events, limiting the descriptive power of the model. As a consequence, also the predictive power of the model is reduced, and its usefulness for risk assessment, evolutive maintenance, and design becomes limited as well. Therefore in the following section we focus on outage causes, and derive a characterization to describe notable outages analyzed in the literature.

2.2. Causes

Concrete examples of events causing Internet outages are discussed in the following, leading to the characterization we propose at the end of the paragraph. A timeline showing relevant cases is also reported in Fig. 2.

2.2.1. Natural causes

Several studies in literature investigated the impact of natural disasters on IP networks demonstrating how earthquakes, hurricanes, and thunderstorms can cause severe network disruption. Examples are the Taiwan earthquake (2006) [9], the Japan earthquake (2011) [23], and the Katrina (2005) and Sandy (2012) hurricanes [11, 74] and [21]. Besides catastrophic natural events,

			characterization axes					
			origin		intentionality		disruption type	
source	class		natural	human	accidental	intentional	primarily physical	pure logical
Çetinkaya et al. [58]		large scale disasters	-		✗		✗	
		socio-political and economic challenges		✗		✗	-	
		dependent failures	-		-		-	
		human errors		✗	✗		-	
		malicious attacks		✗		✗	-	
		unusual but legitimate traffic*		✗	✗			✗
		environmental challenges**	-		-		✗	
Markopoulou et al. [166]		maintenance*		✗		✗	-	
	shared	router-related	-		✗		-	
		optical-related	✗		✗		✗	
		multiple-links failure	-		-		-	
	single	high failure links	-		-		✗	
		low failure links	-		-		-	

✗	The class is characterized according to the specified cause type.
-	The class does not distinguish inside the specified characterization axis.
*	This class does not fall in our definition for outages.
**	According to our definition, some cases comprised in this class are not considered outages

Table 1: Mapping of cause-based outage classification schemes on our proposed characterization.

also minor—yet powerful—local natural events can cause severe disruption and long lasting outages, if they affect critical parts of Internet infrastructure. As an example, the cable cut occurred on the seabed between the islands Saipan and Tinian (Commonwealth of Northern Mariana Islands, CNMI), in July 2015 was attributed to the strong undersea currents that swipe the shallow sea [27] (and that in 2008 made a boulder roll over and sever the same cable). Nevertheless the outage affected, with Internet, also phone/SMS, banking , airlines, and the Weather Service office, leading the Government of CNMI to declare a "State of Significant Emergency" because of the outage, that lasted 3 weeks.

Network outages can be the direct consequence of power shortages. Examples are the blackouts occurred in Moscow during May 2005 [20] and in Northeastern United States during August 2003. The latter involved almost 50 million people and affected about 50% of all Internet ASes [73]. Blackouts may also be the consequence of other events: for instance, after the Japan earthquake, nuclear power plants were taken offline upon a request of the government, leading to a long lasting power shortage in the country. We refer the reader to [91, 94, 205] for a comprehensive insight into interdependencies among critical infrastructures.

2.2.2. Human-originated causes

Large-scale outages can be due to human action, either intentional, or by mistake. Entire countries may *disappear* from the public Internet as the consequence of censorship measures. This form of censorship was put in action in 2011 in Libya [72] and Egypt [71]. Note that not all the censorship techniques cause network outages, more subtle techniques exist [69, 33]. The reader may refer to [144] for more insights into censorship and co-option in the Internet. Logical or physical disruption can be obviously caused by attacks. Prefix hijacking attacks[42], forcing the traffic sent to a set of destinations to be routed along different routes, cause logical disruption. An example is the prefix hijacking attack on YouTube performed by a Pakistani ISP [52]. An increasingly evident type of attack is constituted by massive Distributed Denial of Service (DDoS), often carried on by means of botnets, such as the attack in October 2016 against the DNS service provider Dyn [159], but also with infrastructures such as the Chinese "Great Cannon" [165].

Physical disruptions instead can be caused by terrorist or military attacks: for instance, about 1% of the globally announced Internet prefixes showed immediate loss of reachability during the World Trade Center attack [178, 181]. Another example is represented by the Thai anti-government protestants temporarily cutting off a large portion of Internet connectivity of their country in 2013 [157].

2.2.3. Logical failures

Device misconfiguration is another source of network outages. For instance, route aggregation in border routers can potentially cause persistent forwarding loops and traffic blackholes [146]. The incidents namely referred to as the Google's May 2005 outage [182] and the "China's 18-Minute Mystery" [70] have been also ascribed to misconfiguration. Rarely, also incorrect de-peering between ISPs can cause disruptions [48]. Legacy network equipments can generate network outages as well. For instance, when the number of globally routable prefixes overcame 512K in August 2014, the Internet suffered significant outages due to old routers limiting their use of a specialized, and expensive, type of memory known as ternary content-addressable memory (TCAM) to 512K prefixes by default [221, 148]. Network outages may also rise when ASes decide to de-peer as the result of a business decision [48].

2.2.4. A cause-based general characterization

The analysis reported so far surfaced a few key aspects of every network outage, and leads us to define the characterization of the Internet outages in Fig. 4. In this characterization, outage causes are classified according to the:

- origin (natural or human),
- intentionality (accidental or intentional),
- type of disruption caused (primarily physical or pure logical).

We use the term *primarily* because physical disruptions may determine or not also logical disruptions as side effect. We refer to pure logical disruptions, instead, as those logical disruption being not the consequence of physical disruptions. For instance, earthquakes and hurricanes are natural disasters that accidentally cause physical disruptions of network routing elements while misconfigurations are human faults that accidentally causes pure logical disruption. Military attacks, instead, are performed by humans to intentionally cause physical disruptions, while maintenance activities can accidentally determine logical disruptions (while planned downtime for maintenance is not to be considered as an outage). Of course *intentionality* is applicable only to events of *human* origin.

Considering the classifications of outages described in Section 2.1, we have proposed a cause-based approach, therefore we compare with [166] and [58] providing a mapping between our characterization and the cited ones in Table 1. It can be noted that the aspects we have highlighted are mostly orthogonal to those adopted in former classifications; for the case of Markopoulou et al. this is highly evident, and reflects the fact that the authors did not intend to present a general characterization, but only to characterize a specific operational network from the observation point of the network operator itself, and their classification is scoped by the adopted detection and analysis methodology (see Section 4 for more detail). Compared with Çetinkaya et al. we focus on *causes*, and characterize them according to few fundamental properties, while the authors present a coarse, partially overlapping grouping of *challenges* aimed at a convenient description of outages types, and relate them to 49 attributes that describe also the potential impact, its extension, and other aspects related to the affected infrastructure. Exploiting these differences, the comparison in Table 1 can be used as a cross reference to further exploring outages from different point of views, and shows interesting equivalences, namely *socio-political and economic challenges* with *malicious attacks*. We argue that our characterization provides practical value when researching network outages, as it reflects more closely the different disciplines that have analyzed network outages. In fact the concepts, methods, and data available for statistical characterization of fault occurrence are highly different among the classes we consider. The validity of the descriptive models can change across different classes according to completely different criteria and time scales (e.g. human causes versus natural ones); as a consequence, also prevention, mitigation, and recovery strategies change accordingly. Moreover, the nature of the disruption (primarily physical versus pure logical) additionally affects the monitoring/detection possibilities (see Section 4), and as a consequence the strategies to cope with outages change as well (see Section 6 for a more in-depth analysis).

2.3. Basic detection and analysis techniques

Researchers working in the field of Internet outages often rely on a common set of basic techniques to perform detection and analysis. Based on our experience and after the analysis we did, we classify these techniques in the taxonomy reported in Fig. 5.

Basic techniques are either directly related or not related to the network traffic. When adopting non-traffic-related basic techniques, researchers commonly inspect (1.) *non-structured data sources* such as technical blogs (e.g., Renesys [17]), mailing lists (e.g., NANOG [12] and outages [14]), alarms raised by the final users of networks and services complaining through microblogging social networks; (2.) *semi-structured data sources* such as device usage and error logs, customer emails, quality alarms, and user activity logs; (3.) *structured data sources* such as network trouble tickets. For instance, Banerjee et al. [43] recently used text mining and natural language processing to analyse the outage mailing list.

Most of the literature, however, is based on traffic-related techniques: these techniques can be divided in *active probing* and *passive monitoring*. Active probing techniques inject into the network purposely forged synthetic measurement traffic: by observing how the network treats the injected traffic, researchers can infer the status of the network under investigation to potentially detect failures. Due to the radically distributed ownership of the Internet among its constituents parts (e.g., ASes), active probing represents a valuable tool to gather knowledge about this ecosystem of networks on which no one has full access nor control. The most used active probing tools adopted by researchers working on Internet outages are *ping* and *traceroute*. Both tools rely on ICMP [194]: the former estimates the round trip time related to a given destination as the time elapsed between sending an ICMP Echo Request and the receiving the corresponding ICMP Echo Reply. The lack of responses when using ping or very high delay may uncover network failures along the path [198, 199, 200]. Traceroute allows operators to troubleshoot network failures and poor network performance by listing the IP addresses owned by the devices traversed by traffic towards a destination. While Traceroute provides important knowledge to detect and locate network outages, this basic diagnostic tool is also known to be affected by several drawbacks, e.g., load balancers [39], anonymous routers [168], hidden routers [192], misleading intermediate delays [164], and third-party addresses [161]. More in general, active probing is not scalable due to the imposed network overhead: for this reason, it can be profitably used only towards a reduced set of destinations periodically probed. Due to these limitations, active probing can mainly expose large-scale, long-lasting outage events.

Differently from active probing, passive monitoring does not inject additional traffic into the network but takes advantage of the traffic-related information already stored by the network devices or by third-parties. These techniques can be further classified according to the type of traffic they are related to, specifically *control-plane* and *data plane* traffic. When focusing on *control-plane* traffic, researchers inspect routing-related information. In case of inter-domain routing, an extremely valuable source of information on global Internet dynamics and outages is represented by BGP [204]. Thanks to BGP route collectors [108], research projects such as Routeviews [19], RIPE RIS [28], PCH [15], and publicly accessible BGP looking glasses [137] allow researchers to monitor the best routes announced by ASes. To some extent, an Internet outage can be visible in the BGP traffic since it potentially causes several best routes to be dropped and a remarkable amount BGP update messages to be exchanged. Unfortunately, despite their numbers and privileged locations across the Internet[1], the BGP route collectors provide a forcedly limited yet valuable information on the global inter-domain routing [108]. Only partial routing information can be gathered about those ASes far from the route collectors available. Also, since BGP only exposes best routes, researchers need to aggregate BGP data over larger periods of time to gather a more complete view of the logic interconnections of an AS: this process is error prone since also stale interconnections might be considered. When studying In-

[1] For instance, PCH BGP route collectors are located at large Internet Exchange Points.

ternet outages, researchers can also take advantage of the type of relations among ASes (e.g., customer-provider or peer-to-peer) inferred and made publicly available by institution such as CAIDA [86, 5]. Note, however, that accurately inferring these relations is an extremely complex task since some of the basic assumptions on which these inferences are made do not always hold in the real Internet [103] and very large ASes, spanning over different countries or continents may have different types of relation depending on the specific geographic area. Also information related to the intra-domain routing may provide useful hints on network outages. For instance, by logging Interior Gateway Protocol (IGP) messages generated by the network routers with tools such as Packet Design Route Explorer [18], researchers can investigate how the network reacts to disruptive events. Unfortunately, the operational climate of Internet where ASes collaborate and compete generally disincentives the sharing of IGP-data with the research community. For this reason, only few researchers had the privilege to access similar data sources.

Finally, as the control plane, also the *data plane* provides useful information on network outages. For instance, several works monitored the aggregated volumes of traffic in the network. Indeed, when multiple network components fails, the total amount of traffic in the network may significantly drop. Typically, the traffic volumes are compared with patterns referred to as “normal” periods in order to evaluate the approximate ratio of traffic drop due to the outage.

3. Analyses of specific outages

The basic techniques described in the previous section have been widely used to analyze specific Internet outage episodes. In this Section we present and discuss these analyses , ordering them according to the *origin*, then by specific event or group of events, as summarized in Table 2.

3.1. The Japan earthquake in March 11, 2011

This disaster [23] is the most investigated natural event impacting IP networks. The works focusing on this event also represent a good example of the heterogeneity of methodologies, data sources, and views adopted by researchers conducting similar studies. Overall, we noticed how each work provided an interesting piece of an overall puzzle: one can achieve a more comprehensive understanding of this single event only by reading all the available works. This result is the direct consequence of the lack of assessed methodologies, best practices or guidelines in literature for analyzing Internet outages. We consider this as one of the most important open issues in this research area, as we further discuss in Sec. 9.

In the following, we examine the main findings of the most relevant works focusing on this specific event. Cho et al. [62] investigated the impact of the earthquake on a Japanese ISP named IIJ, mainly relying on passive monitoring: they analysed (i) inter-domain control plane information, i.e. the BGP traffic exchanged by the ISP with a major peer collected using Quagga MRT [16]; (ii) intra-domain control plane information, i.e. OSPF signalling collected via Packet Design Route Explorer [18]; and (iii) data plane information, i.e. traffic volumes on trans-oceanic links and residential traffic collected via SNMP. Despite the large impact caused by the earthquake on the infrastructure of this ISP (trans-oceanic link failures, connectivity lost in six prefectures in the Tohoku area), the authors noticed how over-provisioning and traffic re-routing strongly limited the visibility of this disruption outside the ISP.

Fukuda et al. [97] investigated the earthquake impact on a different Japanese ISP named SINET4. They also used routing information (BGP and OSPF) and traffic volumes logging the event messages generated by routers. The authors observed (a.) significant traffic drop in the backbone of the ISP; (b.) the lost of connectivity for several university sites; and (c.) traffic congestion on other links due to users accessing realtime streaming sites to obtain emergency information. A full recovery of the network was reached only 5-6 weeks after the event.

Liu et al [154] characterized the inter-domain

Natural disasters	Japan earthquake (2011)	3.1	[23, 62, 97] [154, 47]
	Other earhquakes	3.2	[140, 79, 53, 197]
	Hurricanes	3.3	[74, 115, 30]
Human faults	Intentional disruption	3.4	[78, 79, 219, 46] [165, 159, 29]
	Accidental causes	3.5	[78, 79, 219, 46] [165, 159, 29]

Table 2: Analyses of specific Internet Outages.

rerouting occurred after the Japanese earthquake by using the betweenness centrality metric applied to BGP data. The authors observed that three major providers of inbound traffic to Hong Kong were affected by unstable routing due to a cable fault after the earthquake.

Finally, Bischof et al. [47] gained insight into the impact of this earthquake by mainly relying on data-plane measurements performed by a widely adopted peer-to-peer system (i.e. BitTorrent), identifying the specific regions and network links where Internet usage and connectivity were most affected. The authors used two plugins developed for the Vuze BitTorrent client [237], named Ono [55] and NEWS [64], to anonymously collect usage statistics as well as passive monitoring and active measurements such as Traceroutes towards a subset of connected peers. By leveraging the view of this popular P2P system, the authors documented an overall decrease in the usage of BitTorrent as well as routing changes in the affected area.

Despite thousands of victims and huge destruction, all these works demonstrated that the most powerful earthquake ever recorded to have hit Japan had only a relatively limited and well localised impact on the Internet functionality.

3.2. Other earthquakes

Other earthquakes have been investigated as well. Köpp [140] analysed BGP data to shed light on the consequences of an earthquake hitting New Zealand in 2011. Köpp first identified the network prefixes allocated to the network equipments located in the proximity of the epicenter by using GeoLite City Database of MaxMind [8]. Then, he searched for route-changes or withdrawals related to these prefixes seen at the London Internet Exchange Point (LINX), noticing only a limited impact. Both the Japanese and New Zealand earthquakes have been also investigated in [79]. For sake of completeness, how earthquakes impact the Internet has been also the subject of many works focusing on user activity on social networks. For instance, Bruns et al. [53] investigated user activity on Twitter after the New Zealand earthquake. Similarly, Qu et al. [197] pinpointed the effect of the 2010 Yushu Earthquake on a popular Chinese microblogging site. Other similar studies can be found in [196, 184, 92, 150, 222, 231, 155, 183, 229, 236, 211].

3.3. Hurricanes

Hurricanes also attracted great interest in the community. For instance, hurricane Katrina hitting the Gulf region of US in 2005 has been analysed by Cowie et al [74], that reported great loss in local networks but only very limited consequences on the public Internet which continued to achieve high reliability. The consequence of hurricane Sandy has preliminary been investigated in [115]: Heidemann et al. used large-scale ping-based experimental campaign to assess if steadily reachable destinations were not reachable any more likely due to the hurricane. The authors noticed how the outage rate doubled in the area hit by the hurricane while networks took about four days to fully recover. Aben [30] relied on DNS reverse lookup to inspect Traceroute traces traversing NYC and ASH areas, noticing how most of the paths were rerouted around the areas hit by the hurricane Sandy, making networks still operational and interconnected in spite of the difficult circumstances.

3.4. Intentional disruption

Many researchers dissect censorship actions responsible for large-scale network outages. Dainotti et al. [78, 79] analysed how censorship caused Egypt and Libya to mostly disappear from the public Internet. The authors relied on both active probing and passive monitoring: they exploited control-plane information such as BGP-data and Regional Internet Registries (RIRs) [1] delegation files, but also data plane information such as unsolicited traffic to unassigned address space (commonly referred to as Internet Background Radiation - IBR [185]) and Traceroute measurements. The authors noticed how, during the Arab Spring, the amount of unsolicited data plane traffic coming from these geographic areas significantly dropped. They also argued how analysing IBR can be an effective tool for uncovering large-scale network disruption [46]. Slater [219] used SNMP counters to analyze the problems caused by DDoS attack against Internet Root Servers occurring during October 2002, as well as consequences of router misconfiguration.

Marczak et al. [165] have deeply analyzed a DDoS targeted at the censorship circumvention services offered by GreatFire.org, an organization that aims at monitoring and countering Internet censorship in China. From the inspection of server logs and by active measurements the researchers assessed the

Man-In-The-Middle nature of the attack and inferred several properties of the infrastructure enforcing it (dubbed as "the Chinese Great Cannon"): the most notable ones are its probabilistic nature (impairing detection efforts) and the injection of HTML code causing unaware users to participate in the attack.

A recent DDoS attack with huge impact and visibility has been the one targeting the DNS service provider *Dyn*, because it indirectly affected the many high-traffic websites using the service [159]. In the aftermath of the event, the attack was tracked back to botnets of the *Mirai* type, that leveraged vulnerable IoT appliances in the order of $100,000$. While the technical properties of the botnet were not new per-se, the volume of generated network traffic on the target has reportedly reached unprecedented levels of $665 Gbps$.

When the blocking is intentionally enforced by a state, it is usually enforced at ISPs level [33]. Such is the case for the already mentioned *Arab Spring* in 2011, and recently for Cameroon (still in act at the time of writing) [29]. Such events have occurred in non-democratic countries, in correspondence with elections or social and political unrest, and having a clear impact on freedom of expression, it is often denounced by international human rights organizations [32], backed by anecdotal references or leaked documents (blocking orders to ISPs).

3.5. Accidental causes

A class of outages of human origin that can be ascribed to mistakes or misconfiguration are related with prefix hijacking, such as episodes investigated in [42]. Researchers from RIPE ATLAS [52] used RISwhois[2] and BGPlay to inspect the information stored by the BGP route collectors and uncover how Pakistan aimed at blocking the YouTube website through prefix hijacking. Hiran et al. [116] used control-plane data (BGP updates and AS-level topology) and data-plane measurements (Traceroute traces) to dig into the hijacking of $50,000$ prefixes in April 2010 made by China Telecom. Thanks to these data sources, the authors classified this event as an incident. Wan et al. [182] observed through RouteViews BGP data how, immediately prior to a long lasting Google service outage occurred in May 2007, an AS operated by Cogent started announcing itself as the originator of an IP prefix assigned to Google. Outages in the Domain Name System (DNS) have been analysed as well: Pappas et al. [186, 187] identified delegation inconsistency, lame delegation, diminished server redundancy, and cyclic zone dependency. These operational errors heavily impacted availability and query delays. Chan et al. [59] used active probing to pinpoint the consequence of a submarine cable fault occurred in 2010 on end-to-end network performance. The authors employed Traceroute and other network diagnostic tools to assess the path-quality degradation in terms of delay and asymmetric packet losses.

Regarding more general network operational failures, Markopoulou et al. [166] analyzed IGP messages exchanged in the Sprint backbone to characterize failures that affect IP connectivity. Surprisingly, they noticed that about 20% of failures in the Spring backbone occurred during periods of scheduled maintenance activities: this large fraction of outages seems to be self-inflicted by the network operator. Turner et al. [232] used non-traffic-related data such as email logs, router configurations and syslogs to analyse over five years of failures occurring in the CENIC network serving most of the public education and research institutions in California. Mahajan [158] inspected three weeks of BGP data to discover how BGP configuration errors are pervasive, with 200-1200 prefixes (0.2-1.0% of the BGP table size) suffering from misconfiguration each day. Finally, Huang et al. [122] correlated six months of BGP update streams in the Abilene backbone with a catalogue on known disruptions of nodes and links noticing the importance of simultaneously analysing multiple BGP update streams to detect most of the important events.

3.6. Discussion on outage analyses

Dissecting specific episodes of network outages provides essential knowledge on (i) how the network globally and locally reacts to large scale disruptive events and (ii) which specific approaches and tools we can employ to perform similar studies. Reviewing these works, we noticed how researchers commonly relied on scattered data sources related to both control- and data-plane and different approaches such as active probing and passive monitoring. The main reason behind this choice is the need for a more comprehensive view of the phenomenon that can be achieved only by crossing different sources of information.

[2] http://www.ris.ripe.net/cgi-bin/riswhois.cgi

To this regard, it is important to underline how the adopted basic instruments suffer from severe limitations as we anticipated in the previous section. Indeed, BGP data are well known to provide a heavily incomplete view of the inter-domain routing [108]. Relying exclusively on the control-plane may also generate false alarms [254]. For instance, this may happen due to default routing [54], causing Internet traffic to normally reach its destination even if the corresponding route does not appear at the control-plane level. Also data-plane measurements are not free of limitations. For instance, Traceroute may induce users to incorrectly reverse engineering the network path [192, 161, 162, 168, 39]. Also the lack of replies when using ping does not necessarily implies lack of connectivity [121]. Several concerns also exist on the accuracy of IP Geolocation Databases such as Maxmind as extensively demonstrated in different works (e.g., [193, 255]).

In conclusion, researchers investigating Internet outages should carefully consider the validity of their conclusions in light of the limitations of adopted tools.

4. On-line Outage Detection

Passive monitoring	Control Plane	Profile	[51]
		Time	[228, 81]
		Other	[105, 56, 249, 253]
	Data Plane	Core	[212, 78, 85]
		Edge	[230, 110, 63]
Active probing	Based on Ping and Traceroute		[200, 198, 199, 223]
	Tomography		[89, 75, 82]
Hybrid active-passive			[132, 125, 248, 160]

Table 3: On-line outage detection approaches.

Detecting ongoing network outages is important to qualitatively and quantitatively understand the type, the scope, and the consequences of a disruptive event, as well as to timely activate mitigation and remediation activities. In this Section, we focus our attention on systems and tools designed for systematically detecting network outages in real-time or near real-time. The goal in this case is to understand the frequency and the duration of similar events to also possibly trigger effective countermeasures. These approaches are particularly helpful either in case of physical disruptions (allowing the identification and the possible replacement of the damaged network components), or in the case of logical disruptions (allowing network administrators to quickly restore a satisfying operational status).

We characterize the detection methods in literature based on the adopted monitoring teachiques (see Table 3) and discuss them accordingly.

Moreover, we extrapolate the general approach adopted by these studies and we propose the flowchart reported in Figure 6.

Detecting network outages typically requires four steps: (1) Data collection and preprocessing; (2) Network outage detection; (3) Outage locating; and (4) Root cause analysis. Most systems implement only the first two steps. These systems continuously monitor the network by collecting data, typically with a combination of the basic techniques we introduced in Sec. 2, such as active probing or passive monitoring. During this step, data filtering and sanitization also take place in order to remove as much noise as possible from the data. Different algorithms are then applied on the refined data during the second step, to detect large- and small-scale events potentially related to Internet outages. All these events are properly logged to enable more advanced off-line analysis.

Some advanced works also tried (i) to locate the outage—i.e. identify the network element(s) originating the detected event and/or (ii) to find the root causes of the disruption—i.e. its origin. This approach is indeed adopted in scenarios where most part of the network is under control (e.g. for ISPs), so root cause analysis and mitigation/resolution of the outage is more likely thanks to the detailed knowledge and monitoring/control possibilities: as an example of a fault diagnosis system for a complex access network we refer to [57].

In the following, we discuss outage detection tools or approaches according to which specific mechanism is primarily employed during the data collection process.

4.1. Passive monitoring

Most used passive monitoring techniques are based on control plane information, leveraging data collected by means of the BGP protocol. Other approaches rely on analysis of data plane traffic, mainly based on volume variations related to outage events.

4.1.1. Control Plane

Public BGP repositories proved to be extremely helpful also for systematic outage detection, as for the analysis of specific outage episodes. BGP data publicly available like RIBs and update messages are systematically crawled from public repositories (e.g., Routeviews, RIPE, etc.) during the data collection step. This data is then converted in a suitable format for the subsequent analysis. During the detection step, a common approach is to group all the BGP messages originating by a given event. This procedure can be affected by inaccuracy (i.e. grouping also BGP messages not related to the outage) and incompleteness (i.e. missing part of the BGP messages related to the outage of interest). To a certain extent, such problems are unavoidable due to the forcedly limited visibility of available data sets. BGP data can also be used to locate the ASes responsible for the detected disruption. The main drawback when relying on BGP data to systematically detect network outages is the large number of false alarms since many legitimate events may also determine changes of the path or of the origin prefix. We discuss in the following several approaches adopting this scheme.

Profile-based detection. Li and Brooks [51] developed I-seismograph, a tool for detecting network outages and evaluating their impact. The main idea is modelling the *normal* state of the Internet, and then monitoring the network for a given period to measure if and how the Internet deviates from this status. The data collection process is entirely based on public BGP data while outage detection relies on a two-phase clustering methodology. Given a *normal profile* of BGP, they compare current BGP attribute values with reference values. An outage is detected if there is a significant difference in a fixed time-window. The BGP normal profile is represented as a cluster of normal BGP attribute values. The measured BGP attributes will fall into a separate *abnormal cluster* in case an outage occurs.

Time-based change detection. BGP Eye [228] clusters BGP updates related to the same BGP event and correlates the events across multiple border routers in order to expose anomalies. The system attempts to identify the root-cause of these anomalies. Two different perspectives are considered: (i) an Internet Centric perspective, to track anomalies through an analysis of AS-AS interactions, and (ii) a Home Centric perspective, to provide an insight on how an AS is affected by anomalies originated from external ASes also several hops distant. The authors evaluated the system against the routing outages caused by the Slammer Worm on 2003 and the prefix hijacking by a Turkey Net (AS9121) on 2004. Deshpande et al. [81] proposed an online mechanism for the detection and the analysis of routing instabilities potentially caused also by network outages. The detection step is based on the analysis of time-domain characteristics of BGP update messages. More specifically, filtering and adaptive segmentation techniques are applied on time series of feature data in order to isolate periods of instabilities. BGP route changes are then also used to locate the ASes originating the routing instabilities.

Other detection approaches. Glass et al. [105] also relied on BGP data. The authors adopted tensor factorization for detecting events of interest, and graph-theory analysis to locate the origin ASes. A similar approach is adopted in [56]. Principal Components Analysis [249] and machine learning techniques [253] have been also proposed for anomaly detection in BGP data.

4.1.2. Data Plane

Other outage detection systems do not primarily rely on inter-domain routing data but on other traffic-based data sources. According to the position of the collection points in the network, these approaches can be grouped in *core-based*, i.e. observing data traffic in transit networks, and *edge-based*, i.e. observing data in stub networks or on end hosts. Some approaches rely on different viewpoints both *core-based* and*edge-based*, exploiting *space invariance* of some traffic properties (e.g. see [77] for a characterization of worm traffic in a trans-oceanic link and an edge network).

Core-based. Flow-based Approach for Connectivity Tracking (FACT) [212] proposed by Schatzmann et al. relied on flow-level data exported by all border routers of a network to compare the incoming and outgoing traffic flows. During the data collection process, FACT collects NetFlows records and aggregates flows per remote host, networks, or AS. The key idea behind the detection process is that a network outage is likely to result in (i) an increasing number of unsuccessful one-way connection to a remote destination (a network prefix, an AS, etc.), and (ii) a decreasing number of successful two-way connections. Other researchers proposed to rely on unsolicited data plane traffic to detect network outages [78, 85]. Glatz et al. [85] monitored unsolicited data plane traffic towards a live network to detect and characterize fine-grained outages affecting local networked services.

Edge-based. A completely different approach is represented by the PerfSonar project [230, 110], a collaborative network monitoring platform providing several network troubleshooting tools. The system is deployed in several independent research and educational networks enabling sharing of data like specific SNMP counters useful to expose network outages. Finally, Choffness et al. [63] proposed Crowdsourcing Event Monitoring (CEM). CEM passively monitors and correlates the performance of end-user applications in order to expose network events including outages. The data collection process is implemented as an extension to BitTorrent. Each end node monitors flow and path-quality information such as throughput, loss, and latencies to locally detect an event.

Common challenges. The works cited above share the difficulty to guarantee user privacy. Djatmiko et al. [87] proposed a general approach to overcome this limitation. More specifically, they designed a distributed mechanism based on Secure Multi-Party Computation (MPC) [80] to correlate NetFlow [7] measurements passively collected from multiple ISPs. MPC consists of a set of cryptographic methods allowing different parties to aggregate private data without revealing sensitive information, thus avoiding the aforementioned limitations. The authors integrated MPC in FACT [212] allowing network operators to troubleshoot outages by solely relying on flow-level information.

4.2. Active probing

Many other outage detection systems primarily used active probing during the data collection process. These systems often rely on Ping and Traceroute for periodically probing a number of destinations from several vantage points. Other approaches are based on tomography techniques, i.e. perform targeted end-to-end measurements based on the—possibly limited—knowledge of the topology of the measured networks (often the whole Internet). Both the types of active approaches usually rely on distributed active measurement platforms such as Archipelago [4], Planetlab [45], DIMES [215], etc.

4.2.1. Approaches based on Ping and Traceroute

Quan et al. proposed Trinocular [200], a system probing each IP block with ICMP echo requests (pings) at 11 minute intervals and classifying responses in two main categories: (i) positive, in case of an ICMP reply is received; (ii) negative, in case of ICMP replies indicating network is unreachable (e.g., destination unreachable), or in case of lack of responses. Negative responses potentially expose Internet outages. The authors demonstrate that a single computer can track outages across the (analyzable) Internet by probing a sample of 20 addresses in all 2.5M responsive /24 address blocks, and detect 100% of the outages lasting more than 11 minutes.

The basic principles behind Trinocular were preliminary explored in [198, 199]. The authors used this approach to investigate macro-events (hurricane Sandy in 2012, the Japanese earthquake in March 2012, and the Egyptian revolution in January 2012) as well as micro-events such as classic daily network outages. They exploited this technique to evaluate the availability of the whole Internet: according to their results, the Internet has a "2.5 nines" availability.

Schulman and Spring [223] designed and deployed a measurement system called "ThunderPing" measuring the connectivity of residential Internet hosts before, during, and after forecast periods of severe weather. The data collection process implemented in ThunderPing crawls weather forecast to identify geographic areas interested by extreme weather condition (e.g., thunderstorms) and triggers ping measurements towards residential hosts located in these areas. ThunderPing demonstrated how failures are four times as likely during thunderstorms and two times as likely during rain compared to clear weather.

4.2.2. Tomography-based approaches

A special family of outage detection systems exploits binary tomography. Binary tomography is the process of detecting link failures by sending coordinated end-to-end probes [89]. It only requires network elements to perform classic packet forwarding operations. Relying on Ping and Traceroute, instead, requires the devices to actively collaborate by providing ICMP responses. Hence, binary tomography is particularly helpful when the networks are configured to discard ICMP messages coming from the outside.

Cunha et al. [75] observed that binary tomography is sensitive to the quality of the input: (i) the network topology and (ii) a set of end-to-end measurements, in the form of a reachability matrix, which indicates whether each path is up or down. The authors developed two methods for generating higher quality inputs to improve the accuracy and the efficiency of binary tomography algorithms. They observed that binary tomography algorithms cannot be always directly applied in real networks, because they tend to generate a remarkable amount of false alarms. The authors proposed (i) a probing method for quickly distinguishing persistent path failures from transient congestion, as well as (ii) strategies for aggregating path failures in a consistent reachability matrix.

Dhamdhere et al. [82] proposed a troubleshooting algorithm, called "NetDiagnoser", based on binary tomography. They extended this technique to improve the diagnosis accuracy in the presence of multiple link failures. NetDiagnoser actually relies on Traceroute-like measurements, performed by troubleshooting sensors located at end hosts inside multiple ASes, and it also considers information on paths obtained analyzing BGP and IGP messages after rerouting around a failure.

Network tomography is a powerful tool. However, it is not free of limitations [252]: fast detection of network outages implies high probing rate, practically infeasible in real networks. Also network dynamics may weaken the basic assumption that the injected packets are traversing the same links previously observed. Load balancing [40] further exacerbates this issue.

4.3. Hybrid active-passive approaches

Few outage detection systems jointly used active probing and passive monitoring.

Katz-Bassett proposed Hubble [132], a system detecting Internet reachability problems where routes to the destination network exist at the control plane according to BGP public data but packets do not reach the destination network through the data plane. Data collection relies on BGP data, Ping and Traceroute measurements triggered by changes at the control plane. Hubble proved to discover 85% of the reachability problems that would be found with a pervasive probing approach, but reducing the probing traffic by 94.5%. The authors also studied the trade-off between sampling and accuracy, arguing that their use of multiple samples per destination network greatly reduces the number of false conclusions about outages.

In [125], Javed et al. designed "PoiRoot", a real-time system allowing ISPs to accurately isolate the root causes of any path change affecting their prefixes. PoiRoot exploits BGP data but also combines existing measurement tools (e.g., Traceroute and Reverse Traceroute [133]) to gain higher visibility on the ongoing events.

Another system exploiting both active and passive monitoring is Argus, proposed by Xiang et al. [248] to detect prefix hijacking. Argus relies on live BGP feeds provided by BGPmon [195] and daily Traceroutes archives made available by CAIDA Archipelago [4] and the iPlane project [156]. The key idea behind this system is that routers polluted by a prefix hijacking usually cannot get a reply from the victim prefix. Accordingly, the authors correlate data-plane (un)reachability with control-plane anomaly from a large number of public BGP route-servers and looking-glasses to expose these network outages.

Other similar works exist. For instance, Hu et al. [160] passively collected BGP routing updates and information from the data plane: the basic idea is to use data plane information in the form of edge network fingerprinting to disambiguate potentially numerous suspect IP hijacking incidences based on routing anomaly detection.

4.4. Discussion

Internet outage detection systems that mainly rely on passive monitoring are highly efficient although very prone to (i) false alarms and (ii) non trivial privacy-related concerns. On the other hand, solutions based on active probing are effective although poorly scalable. For instance, by continuously injecting ICMP probes into the network towards a large number of representative destinations, Trinocular [200] increases by almost 0.7% the Internet "background radiation", i.e. the unsolicited traffic that all the networks in the Internet observe. This load imposed on the network might be easily considered unacceptable. In addition, the necessary trade-off between number of targeted destinations and sampling period causes outage detection systems relying exclusively on active probing to likely report only large and long-lasting network outages. The hybrid approaches, instead, seem to represent the best option since they combine the advantages of passive monitoring and active probing. These systems primarily adopt passive monitoring to continuously gather coarse grained information on the network status. Opportunistic measurements based on active probing are triggered to gather additional information only when this lightweight process provides clues of possible Internet outages.

5. Outage impact evaluation

Nonformalized		[62, 97, 78, 59, 17, 31, 49, 124]
Formal metrics	User	[91, 119, 87]
	Network	[51, 247, 154, 153, 44] [34, 59, 79, 46]

Table 4: Approaches for outage impact evaluation.

Measuring the impact of a network outage is a key challenge when either analyzing specific Internet outages or systematically detecting them. In Table 4 we group the literature on outage impact evaluation according to the adopted approaches, discussed heareafter. Most of the scientific works dissecting specific episodes provided only a qualitative evaluation of the impact of the outage, or reported a quantification not based on shared or significant metrics. Other works, instead, formally introduced metrics and approaches adopting either a user- or a network-centric stand point. We discuss both types in the following.

5.1. Nonformalized impact evaluation

Among works focusing on the Japanese earthquake, Cho et al. [62] cites NTT reports with numbers of damaged base stations, transmission lines, and circuits for fixed line services, restrictions to voice calls acceptance nationwide. Regarding the Internet, the impact of the earthquake is found in volumes of traffic seen by Internet Service Providers located in the area of the quake [62, 97]. Evaluating traffic volumes has several drawbacks acknowledged by the authors themselves: part of volume drop was due to *scheduled* power outages due to restoration works, similarly voluntary shutdowns of servers and networks were performed by companies (and users) to reduce power consumption (we do not include such intentional operations in the definition of Internet outages); the reduced usage due to outages was also likely offset by the increased use of the Internet for searching information, and telecommuting. In the analyses also routing information is considered, and related to submarine cable cuts, reporting the earthquake impact in terms of link-state neighbor events per unit time: related but more complex and formal metrics are described in Section 5.3. We also refer to Section 3.1 for the analyses of the specific outage event.

Similar baseline counts relative to outage impact are reported by [78] in detecting network effects of censorship in the Arab Spring, where the existence and extension of outage is measured by number of network prefixes (continuous sets of IP addresses) or number of individual IP addresses affected e.g. by BGP withdrawal or DDoS attacks (see Section 3.4 for the related analysis methodology). In this case the ultimate effect on users was inferred as global disconnection of the whole population of the affected countries, as the network prefixes amounted to the total of addresses available to country ISPs. Similar considerations, with country-wide affected population, are reached regarding submarine cable cuts [59, 17], specially for countries with limited connection options with neighbours.

Finally, consequences of outages (especially caused by censorship) have been qualitatively analyzed by NGOs and not-for-profit institutions in terms of human rights violations (free speech, accountability of governments, discrimination) [31, 49] or in terms of economic costs, and impact on countries growth and development [124].

5.2. User-centric impact evaluation

The impact of outages is measured according to the troubles perceived by the end-users of the network. For instance, in [91], an outage impact is evaluated with respect to the estimated population served by the affected network routes, based on studies (such as [119]) that correlate population density with Internet usage. The interesting aspect of such approach is that it leads to an impact evaluation closer to the user, since it tries to consider the affected population rather than measurable network metrics. Similarly, in [87] Djatmiko et al. relied on the amount of unreachable BGP prefixes to estimate the number of affected clients, in order to evaluate the severity of an outage. All these works implicitly assume that the final mission of the Internet is to guarantee global connectivity to end-users. Accordingly, the larger is the section of population being disconnected, the larger is the impact of the network outage. This approach is as simple as limited: although mere connectivity is the necessary condition for an end-user, also the perceived network performance matters.

5.3. Network-centric impact evaluation

A few other works evaluated the impact of an outage introducing network-related metrics. An approach commonly adopted is comparing the network status and its performance *before*, *during*, and *after* the outage. In this way, researchers aim at identifying, modeling, and quantifying the perturbation caused by the outage. As we discuss in the following, works available in literature may strongly differ in how this goal is achieved in practice.

In I-seismograph, Li et al.[51] relied on BGP data to characterize the network status. They defined the impact of an outage in the Internet as any deviation from BGP normal profile. This deviation can be described in terms of a *magnitude* and a *direction*. The magnitude represents an absolute, quantitative evaluation of the outage intensity. The direction provides deeper insight, since it indicates which BGP attribute(s) deviates the most from normalcy. I-seismograph exploits a vectorial approach to evaluate magnitude and direction of the outage.

To evaluate the impact of an outage, Wu et al. [247] defined two families of network-centric metrics: (i) reachability impact metrics (RIMs) and (ii) traffic impact metrics (TIMs). The former are focused on how many paths are not available any more. The latter are based on the key observation that the traffic traversing failed links is shifted to new paths after an outage, and this may cause network congestion. Specifically, they define two RIMs: an absolute RIM as the number of AS pairs that lose reachability due to the outage, and a relative RIM as the percentage of disconnected AS pairs over the maximum number of AS pairs that could possibly lose reachability. As for traffic impact metrics (TIMs), an issue is the lack of accurate information on actual traffic distribution among ASes. For this reason, the authors introduced the concept of “link degree”, that is, the number of shortest policy-compliant paths traversing a link. Based on this concept, they introduced three TIMs: (1) an absolute TIM, which is the maximum increase of link degree among all links; (2) a relative TIM, which is the ratio between the absolute TIM and the link degree of the *new path* (after rerouting); and (3) an evenness TIM, which is the ratio between the absolute TIM and the link degree of the failed path. This metric captures the evenness of re-distributed traffic for the failed link.

Similarly, Liu et al. [154, 153] proposed an impact evaluation methodology based on the changes of the traffic load of every AS. To estimate the traffic load of an AS, they proposed a betweenness centrality metric computed over the AS-level topology graph of the Internet extracted from BGP data. Betweenness centrality of a node is the number of shortest paths from all vertices to all others that pass through that node, thus measuring the centrality of the node in the network. Betweenness centrality proved to be a more meaningful measure than connectivity for both the load and importance of a node [95]: betweenness centrality is more global to the network, whereas the latter refers to a local effect. The key idea in [154] is to use betweenness centrality to measure how many AS paths traverse a certain AS, since this will be directly related to the amount of traffic transferred by that AS. The authors proposed approaches to evaluate the aggregated time of route changes, the worst affected components, etc. Banerjee et al. [44] also discussed the limitations of using the connectivity to evaluate the impact of an outage: the authors proposed a new metric called *Region-Based Largest Component Size* (RBLCS) related to the size of the largest network connected component once all the nodes of a region fail. Another graph-based impact evaluation approach is proposed in [34], in which two ways of measuring the impact of a failure are presented: the first is based on the number of link failures, the second considers the *terminal reliability* measuring the effect on the global connectivity of the network caused by a cut.

In [59], Chan et al. introduced two kind of metrics for evaluating the impact of submarine cables faults. The first kind focuses on routing dynamics and it is based on similarity metrics (i.e. the Jaccard distance) between the AS-level routes before and after the disruptive event. The second kind of metrics, instead, focuses on path performance degradation and it is based on the correlation between loss and delay. In [79], the impact of an outage is measured by considering how many IP addresses in the affected geographical area likely lost connectivity. This analysis is based on passive monitoring of data plane unsolicited traffic coming from the affected region. Note that this is a relative measure, because it is compared to the number of IP addresses that were visible before the outage. Furthermore, this information can be used to infer another measure, that is, the maximum radius of the impact. This approach adopts both a user- and a network-centric points of view and it was further explored in [46] where other IBR-derived metrics are proposed to gain insight into macroscopic connectivity disruptions.

5.4. Discussion on outage impact evaluation

The vast majority of the outage-related studies only performed a qualitative or nonformalized evaluation of the outage impact providing a quick insight on the consequences of the outage. The largely incomplete and scattered information provided by similar studies only poorly contributes towards an exhaustive and accurate understanding of the possible consequences of network disruptive events. From this point of view, the few works proposing formally defined user- or network- centric metrics appear to be extremely valuable. Indeed, the proposed metrics are non-ambiguous and can be adopted to investigate and compare different network outage events. Unfortunately, we noticed the lack of a widely accepted framework of formally defined metrics: each work proposed or adopted different metrics causing the Internet outages documented in literature to be very hard to compare. As we deepen in Sec. 9, this heterogeneity of metrics represents a key open issue.

6. Network resilience

Internet outages are inevitable, forcing network operators and administrators to understand whether or not their infrastructure is enough resilient. In this section, we first briefly discuss the definitions proposed in literature and the terms commonly adopted when referring to network resilience. Then, we discuss the metrics and approaches proposed to quantify it (summarized in Table 5). Note that a comprehensive survey on network resilience is out the scope of this paper. Our goal is to provide an overview of the approaches and metrics used in the field of Interned outages. For more details on network resilience, we refer the reader to [202, 201].

Qualitative		[62, 224]
Formal	Resilience	[88, 202, 213, 201, 172, 106]
	Resilience-related	[152, 34, 244, 170, 169, 175]
	Data-center	[104, 136, 101, 107]

Table 5: Evaluation approaches for network resilience and related metrics.

6.1. Definitions of Resilience

The term resilience has its origin in the Latin word *resiliere*, which means to *bounce back*. Some studies tried to address the lack of standardization and rigour when quantitatively defining resilience in general by reviewing the different existing definitions, concepts, and approaches. For instance, Henry et al. [201] showed how the general concept of resilience has been variably declined in different contexts. In physics, for instance, resilience is the ability of a material to resume its natural shape after being subjected to a force. In sociology, it is the ability of a social system to respond and recover from disasters [126, 207].

Regarding communication networks, Sterbenz et al. [224] defined resilience as the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation [224]. Whitson et al. [202] provided two complementary definitions (i) a *static resilience*, which is “related to the ability of an entity or system to maintain function when shocked”; and (ii) a *dynamic resilience*, which is “related to the speed at which an entity or system recovers from a severe shock to achieve a desired state.” They argue that the former is related to the time to failure while the latter is related to its time to recovery. Finally, Cholda et al. [65] introduced the concept of “Quality of Resilience” (QoR), to correlate resilience with the methodologies adopted to estimate the Quality of Service (QoS) in a network. The authors also provided a comprehensive survey of resilience differentiation frameworks in communication networks.

Researchers often focus on different aspects of network resilience referred to as fault tolerance, reliability, elasticity, and survivability, and other concepts that overlap or are closely related to these, such as dependability, and security. A tentative solution to this abundance and confusion of terminologies is the comparative analysis performed in [35], where network dependability is compared to fault-tolerance, reliability, security, and survivability, and is related to a number of attributes such as availability, maintainability, etc.

We clarify the main terms and concepts in the following. For further details, we refer the reader to [224, 35].

Fault Tolerance. A fault tolerant network is able to deliver services even in the presence of multiple faults, i.e. flaws potentially causing a deviation from the normal or correct operational status of the network. One way to achieve fault tolerance is the use of redundancy. We discuss redundancy and other outage countermeasures in Section 8.

Reliability. Commonly used in the design, deployment, and maintenance of critical systems, the reliability of a network describes the probability of not observing any failure in a certain time span. Hence, reliability quantifies the continuity of proper service and it is sometimes implicitly used in outage-related studies. For example, in [169, 34] the network resilience under geographically correlated failures is evaluated by calculating the average *two-terminal reliability.* Using the average two-terminal reliability is quite a common approach, especially in graph-theoretical works. However, Segovia et al. [213] argued that, when considering connection-oriented networks, a two-terminal reliability metric is not appropriate to assess the capability of the network to guarantee connections.

Availability. Availability is a concept closely related to reliability, and in [25] is defined as "the ability of a system to be in a state to perform a required function at a given instant of time or at any instant of time within a given time interval; assuming that the external resources, if required, are provided". To tell the difference from reliability and availability, [35] argues that "availability is reliability evaluated at an instant" (instead of over a time interval). Similar metrics are adopted to measure availability and reliability.

Elasticity. Introduced by Sydney et al [227], the elasticity of the network is formally defined as the area under the curve of throughput versus the percentage of remaining nodes in a network under attack. Hence, the elasticity aims at (i) describing the adaptability of a (network) topology to node and link failures, and (ii) capturing the overall percentage of flows rerouted under the aforementioned failures.

Survivability. In [224], survivability is defined as the capability of a system to fulfill its mission, in a timely manner, in the presence of threats such as attacks or large-scale natural disasters. Castet et al.[127] observed how resilience and survivability are interchangeably used according to the specific context of a given study. Usually, in network-oriented studies, survivability is seen as a *static component.* For further details on survivability, the reader may refer to [114] [139].

6.2. Evaluating network resilience

Network resilience has very often been evaluated only qualitatively, i.e. no formal metric was introduced or adopted. This is a common limitation of outage-related studies where qualitative evaluations are much more common than quantitative, structured evaluations. For instance, Cho et al. [62] argued that, during the Japanese earthquake, "despite many failures, the Internet was impressively resilient to the disaster". Interesting, Sterbenz et al. [224] argued that "it is widely recognised that the Internet is not sufficiently resilient, survivable, and dependable, and that significant research, development, and engineering is necessary to improve the situation". Being able to evaluate *how much* the network is resilient against disruptions is an essential next step toward a more comprehensive understanding of network outages and their mitigation.

For this reason, other works formally introduced metrics and approaches to quantify the network resilience.

6.2.1. Works introducing formal metrics for resilience

Dolev et al. [88] proposed to evaluate resilience by using classic graph-based connectivity metrics, such as the average shortest path length, the largest component size, and the number of connected node pairs in the network. The authors also consider the routing policy-compliant directed graph modelling the analyzed network. Whitson et al. [202] proposed a probabilistic model for evaluating the static network resilience described as a Probability Density Function (PDF) of two-terminal network reliability when considering external failures affecting the network. In [213], Segovia et al. measured resilience by considering the number of connections that survive a large-scale failure. Hence, the complement of this metric can be used to evaluate the disruption caused by an outage. The authors applied this approach to GMPLS-based transport networks. In [201], Henry et al. proposed a resilience metric as a time-dependent function describing the ratio of recovery at a given time from an outage suffered by the system at a certain point in the past. In [172], a fuzzy architecture assessment for critical infrastructure resilience is presented. Finally, Gorman et al. [106] proposed distance based approaches for identifying critical nodes and links in communication networks and evaluated them through simulations. In this pioneering work, the authors demonstrated the importance of the structural properties of small world and scale free networks. They also preliminary explored a method for analyzing the interactions of physical and logical networks demonstrating how these are dependent on at both a micro and macro level: although the database of national data carriers adopted in this paper is now largely outdated, this conclusion appears still valid nowadays.

6.2.2. Works focussing on resilience-related metrics

Other works faced the problem of quantifying the resilience of a network relying on operations research and graph theory techniques. They also commonly refer to the *vulnerability* of the network instead of its resilience. To the best of our knowledge, no formal and shared definition of vulnerability exists in literature in the context of outages (and most confusingly, in the context of security it is defined as "an internal fault that enables an external fault to harm the system" [41], i.e. as a sub-type of threat, not as a system attribute). Difficulties and approaches in assessing a network vulnerability to disruptive events are also reported in [83]. We argue that, in the context of outages, vulnerability has been utilized as the complement of resilience, i.e. a minimally vulnerable network is a network with maximum resilience, and vice versa. We discuss some of these works in the following.

Usually, works focussing on evaluation of network vulnerability (i) propose a graph representation of the network, (ii) define an outage model, and (iii) evaluate resilience-related metrics. Typically, an outage is modelled as a circular disk, centred at some point in the network (usually a network node). Any network element intersecting with this disk is destroyed by the outage and it is removed from the graph.

For example, Li et al. [152] assessed the survivability of a network affected by random region failures by examining how the network throughput performance degrades. Moreover, they also discuss the network upgrade issue against such region failures by also addressing the corresponding traffic throughput optimization problem and linear programming formulation.

Similarly, Agarwal et al. [34] modelled an outage as a disk around its epicentre. Their focus is on network vulnerability against intentional outages including military attacks, providing algorithms to find (i) the vulnerable points within the network in case of single and multiple disasters; and (ii) the points responsible for the most significant destruction. Interestingly, they also proposed a simple probabilistic model in which the probability of each network element failure is given. This is due to the fact that a network element does not necessarily fail, even when it is close to the outage epicentre. They proposed two metrics: (i) the number of link failures caused by the outage; (ii) the two-terminal reliability.

Wang et al. [244] also proposed a probabilistic outage model and defined metrics and methods to assess network vulnerability. They argued that "neglecting probabilistic behavior of a region failure may significantly over-estimate or under-estimate its impact on network reliability". The authors proposed three metrics to assess the network vulnerability: (i) the *remaining link capacity*, i.e. the expected capacity of all remaining survived links; (ii) the *pairwise capacity reduction*, i.e. the expected decrease in traffic between a pair of given nodes; and (iii) the *pairwise connecting probability*,i.e. the probability that a pair of given nodes with path protection is still connected.

In [170, 169], Neumayer et al. modelled outages as random circular cuts and random line-cuts. They proposed a method to calculate network performance metrics based on geometric probability. They evaluate network robustness by calculating the average two-terminal reliability of the network nodes. In [175], Neumayer et al. assessed the vulnerability of the fiber infrastructures to disasters, exploiting a graph model in which nodes and links are geographically located on a plane.

We only cited the most significant works that addressed the problem of network outages. A remarkable amount of works exist in literature [214, 99, 175, 44, 83, 243, 170, 90]. These are general studies not strictly focused on IP networks, although some of them apply the proposed methodology to a piece of the Internet. For example, in [152] two real network topologies are adopted for simulation purposes: (i) the USA network and (ii) the NFSNET network. Often, these studies do not consider restrictions imposed by the policy-driven Internet routing. Furthermore, considerations on inaccuracy or incompleteness of the graph representing the network are only rarely addressed. In our opinion, a major effort in applying these approaches to real IP networks should be made.

6.2.3. Resilience of data-center networks

Finally, a topic that has recently gained attention is the resilience of data-center networks. This is mainly due to the strict Service Level Agreements (SLAs) they have to meet, and thus the involved economic aspects. Most of current literature in this field mainly focused on single failures of links, computational elements, network devices, etc. inside the data center network (e.g., [104, 174]). Quantifying the data center network resilience to outages occurring in the public Internet received little attention.

The evaluation of resilience for data center systems is addressed in [136], in which Khalil et al. proposed a general resilience metric framework, based on monitoring efficiency features which would be impacted by an outage.

In [101], Ghosh et al. attempted to quantify resilience of IaaS cloud [13]. Their definition of resilience includes the notion of change. They consider two types of changes: (i) changes in client demand (e.g., job arrival rate), and (ii) changes in system capacity (e.g., the number of available physical machines). Ghosh et al. proposed to quantify cloud resilience in terms of effect of changes on two performance-based quality-of-service (QoS) metrics: the job rejection rate and the provisioning response delay. Their analysis is based on a stochastic reward network approach.

In [107], Greenberg et al. observed that Cloud Service Providers (CSPs) exploit geo-distributed networks of data-centers. Geo-diversity can enhance performance (e.g., service delays) and increase reliability in the presence of an outage impacting an entire site. In this context, they made considerations on optimal placement and sizing of these data-centers.

7. Outage Risk Assessment

Risk assessment is "the determination of quantitative or qualitative value of risk related to a concrete situation and a recognized threat" [129]. A risk is made up of two essential components: (i) the magnitude of the potential loss, and (ii) the probability that the loss will occur. In RFC 4949 [216] the concept of *acceptable risk* is defined. Specifically, a risk is acceptable if it is understood and tolerated, usually because the cost or difficulty of implementing an effective countermeasure for the associated vulnerability exceeds the expectation of loss. Risk assessment methodologies are commonly employed in safety engineering and reliability engineering studies. On the other hand, this concept has not been deeply examined in computer networks engineering so far.

We believe that further investigations should be carried on. There are two main reasons for this. Firstly, the Internet has become a critical infrastructure, which is something we have already discussed in Sec. 1. Many risk assessment studies have been performed on other critical infrastructures (e.g., electric power systems). The same should be done for the Internet. Secondly, risk assessment methodologies can be even more useful in smaller, private networks. Suppose in fact that a company wants to insure its network infrastructure: this can not be done if insurance companies do not know how to perform a risk assessment study on a computer network.

7.1. Risk Assessment in IP networks

The problem of assessing the outage-related risk in IP network is not commonly addressed in literature. In this paragraph, we aim to provide a brief presentation of studies that somehow try to address this problem, so to understand the current state of the art.

An important contribution on risk assessment in networking is provided in [66], where a methodology is defined for assessing risk in networked communications, and designing risk-aware networks. The authors propose a three-dimensional scheme to describe the complexity of Internet networking for risk framing, with the dimensions being

- horizontal sectioning—segmenting the network operator scope in *access, regional, core, and inter-domain*;
- vertical layers of network technology/protocol—fitting the TCP/IP layering, but for the presence of an *Overlay* layer between *Network* and *Application* (service);
- market elements—related to resources shared within the system of different network providers, e.g. *physical infrastructure, equipment, peering, outsourcing, end-user services.*

Moreover, they compare communication networks with networking systems such as aviation and railways, highlighting analogies and differences with impact on risk analysis. The basis for risk analysis is then founded on metrics for network reliability and availability of repairable systems. These are linked with network failure models and loss estimation models. Finally, drawing from finance theoretical models, the risk is expressed in terms of business consequences for different actors, namely providers, users, regulators and researchers. This way the authors of [66] are able to propose a method to map events affecting network functionality onto a quantity expressing the economic risk of the network operator.

In [233], Vajanapoom et al. formulated three risk management techniques for the design of a resilient network, based on (i) the minimization of the maximum damage that could occur in the network, (ii) the minimization of the maximum risk in the network, and (iii) the minimization of the root mean square damage. The paper proposes a risk assessment methodology that is functional to the aforementioned risk minimization techniques. Furthermore, Vajanapoom et al. [233] [234] also adapted the risk concept to networked environments. Specifically, they use the concept of *network state* as a tuple in which the i-th element specifies whether the i-th network component is in a failure state or not. Thus, there are a total of 2^n possible network states (i.e., failure scenarios). Then, the risk associated with a network state s is equal to the product of the probability of the network being in state s and the amount of damage occurring in network state s. Since all network states are mutually exclusive to each other, the overall network risk is equal to the sum of the risks associated with each network state over all states. Then, two pieces of information are needed: (i) the probability of a state and (ii) the amount of damage that corresponds to that state. According to the authors, the evaluation of a state probability can be determined by opportunely multiplying the appropriate failure probabilities of all network components.

The damage evaluation can be measured in different ways. In connection-oriented networks, such as WDM and MPLS networks, they argue that it is natural to consider the amount of damage associated with the loss of each end-to-end connection due to network failures. Therefore, the amount of damage that corresponds to a certain network state is the sum of the amounts of damage of all failed connections in that network state s. They also claim that, if information on the traffic is available, a damage metric associated with each end-to-end connection that incorporates the societal or monetary effects of the loss can be determined.

In [218], Silva et al. proposed an architecture for risk analysis in cloud environments. Specifically, they propose a model in which the cloud consumer (CC) can perform risk analysis on a cloud service provider (CSP) before and after contracting the service. The proposed model establishes the responsibilities of three actors: the CC, the CSP, and Information Security Labs (ISLs). This third actor is an agent that represents a public or private entity specialized on information security (e.g., an academic or private laboratory). The authors claim that inclusion of this actor makes the risk analysis more credible to the CC.

In this architecture, five risk analysis variables are proposed: (i) the Degree of Exposure (DE), which defines how the cloud environment is exposed to certain external or internal threat; (ii) the Degree of Disability (DD), which defines the extent to which the cloud environment is vulnerable to a particular security requirement; (iii) the Probability (P), which defines the probability of an incident occurrence, (e.g., a threat exploiting a vulnerability); (iv) the Impact (I), which defines the potential loss in the event of a security incident; and (v) the Degree of Risk (DR), which defines the degree of risk for a given scenario of a security incident. Their risk analysis works in two well-defined phases: (i) the risk specification and (ii) the risk assessment. The former defines and quantifies threats, vulnerabilities, and information assets that will compose the risk analysis, whereas the latter consists in the quantification of the aforementioned five variables. The architecture provides a language for the specification of risk, the Risk Definition Language (RDL), specified in XML. This language is used by ISL to specify threats and vulnerabilities, and contains information such as the risk ID, the ISL ID, the threat/vulnerability ID, and reference to a Web Service Risk Analyzer (WSRA), which is a web service specified by ISL to perform the quantification of the DD and the DE. Several other studies deal with risk analysis in cloud systems as well; the reader may refer to [171, 242, 98, 238].

A risk assessment model for Optical Backbone Networks is proposed in [84], where the authors develop a probabilistic model to analyze the penalty from service interruption due to a disaster. The outage causes considered in this work are hurricanes, earthquakes and Weapons of Mass Destruction (WMD), covering the categories both natural and human-originated, but limited to physical damages, of type unintentional (the supposed targets of WMD are based on city population and city importance, not intentionally aiming to the distruction of the network). Failures refer to link failures, and in case of network device failure this is translated to failure of the links connected to the device. The risk model is used to inform a preventive deployment of backup links, so that the overall cost of rendundant infrastructure and potential losses/penalties are minimized, leading to a formulation equivalent to the Integer Linear Programming arc-flow multicommodity problem. The risk model is also used to deal with Correlated Cascading Failures (CCFs), where the probability of further failures depend on the first one, to propose a reactive Traffic Engineering solution aimed at mitigating and recovering from the disaster.

7.2. Risk Assessment and Resilience

The concepts of resiliency and risk assessment are related with each other. In [102], a risk assessment methodology for networked critical infrastructure is presented. This methodology is made of two principal components, (i) the modeling at technological level and (ii) the modeling at economical level. Furthermore, this work explores the relationships between risk assessment and resilience. The authors argue that “risk assessment is a function of event likelihood, vulnerability and impact. When it comes to resilience it is necessary to have tools in place that may assess the behaviour of complex systems in terms of propagation of failure and recovery. Clearly this goes a step further with respect to typical risk assessment”. Therefore, as already noticed in Section 6, the concept of resiliency is necessary in order to evaluate the dynamic behaviour and stability of a system. Further considerations on the links between resilience and risk assessment can be found in the “Resilience and Risk assessment for Critical Infrastructures” workshop [94]. In this, it is said that “The concept of resilience can be seen as a superset in which typical risk assessment is a complementary part.”

7.3. Discussion

So far, we have discussed the notion of “risk assessment”. We motivated the need for further investigations in this field and presented the current state-of-the-art. Only few risk assessment studies have been made in the context of computer networks. On the other hand, these studies are quite common in other engineering fields and are usually focused on a specific outage, e.g. earthquakes. In fact, different outages will cause the network components to react in different ways. For example, in case of earthquakes, we are interested in knowing the response to shocks and vibrations of every network component, so to opportunely mix this information with the earthquake propagation data to understand what happens to each component and to the whole network. In case of hurricanes, different figures of *mechanically reliability* will be required. In case of logically-disruptive outages, such physical-level figures would be useless, and different models need to be developed.

Furthermore, when assessing the risk, it should be necessary to also take into account the interdependent response of correlated systems, e.g. telecommunication and electric power systems. An example is provided in [147], in which the focus is on seismic hazards. Models that consider multiple interdependent networks have been introduced in literature and can be used to perform risk assessment: we refer the reader to [241] and works cited therein.

8. Countermeasures

In this section, we aim at providing an overview on the countermeasures proposed in literature to address the problem of network outages. We first provide an overview of network recovery mechanisms. Then, we discuss the specific solutions proposed in literature to prevent or mitigate the consequences of network outages.

8.1. Recovery mechanisms

A resilient network requires the deployment of (fast) recovery mechanisms, to bring the system back to a fully operational state. According to the literature, recovery mechanisms are necessary to ensure network resilience. As we deepen in the following, recovery mechanisms can be either reactive or proactive; they can also be progressive, especially in case of physical disruption of network components. We provide a general overview of recovery mechanisms in the following, referring the readers to [172] for more details about metrics and taxonomies on recovery mechanisms.

Fig. 7 reports a simple “outage model” as introduced by Henry et al. in [201]: a generic system affected by an outage is modelled as a Finite

State Machine (FSM) composed of three states: (i) the original state, (ii) the disrupted state, and (iii) the recovered state; and two transitions: (a) the system disruption (from the original state to the disrupted state), and (b) the system recovery (from the disrupted state to the recovered state). We focus here on the last aspect. As detailed in [118], network recovery mechanisms can be categorized into *reactive* and *proactive mechanisms.* In reactive recovery mechanisms, after a failure is detected, network nodes re-run the routing algorithm and exchange information for the routing to converge. This can take quite a long time, especially in BGP, where convergence times are generally long. On the other hand, in proactive recovery mechanisms, (i) some network failures are assumed, and (ii) corresponding recovery settings are pre-calculated and distributed among network elements, so that, in case one of that failures is detected, the recovery mechanism immediately selects one of the pre-calculated settings (the one that corresponds to the detected failure). This mechanism aims at reducing the convergence time required by routing protocols. However, in [118] Horie et al. claimed that 'when the failure has not been considered in the pre-calculation, the recovery mechanism cannot completely recover from the failure.' Hence, real-time outage detection techniques identifying possible network failures are essential for such mechanisms to properly operate. Sometimes a different terminology is considered: some works, e.g. [112], consider *protection schemes* and *restoration schemes* that are synonymous of proactive and reactive mechanisms, respectively.

Wang et al. [240] observed it may be impossible to repair all the failed elements simultaneously (e.g., for budget constraints) especially in case of physical disruption of network components. For this reason, *progressive network recovery mechanisms* may be necessary. These involve multiple recovery stages, and cause the network capacity to be only progressively restored over time. The authors noticed how different recovery processes will result in different amount of network capacity increase after each stage due to the limited available repair resources. In [112], Hansen et al. proposed a differentiation based on the scope of the recovery: they defined *global recovery mechanisms* covering link and node failures by calculating a new end-to-end path, and *local recovery mechanisms*, in which failures are handled locally by the neighbor nodes. In [201], the recovery action consists in two key aspects: (i) a component recovery mechanism, which describes policies for restoring or repairing a disrupted component, and (ii) an overall resilience strategy, which is related to implementing component recovery mechanisms at the system level.

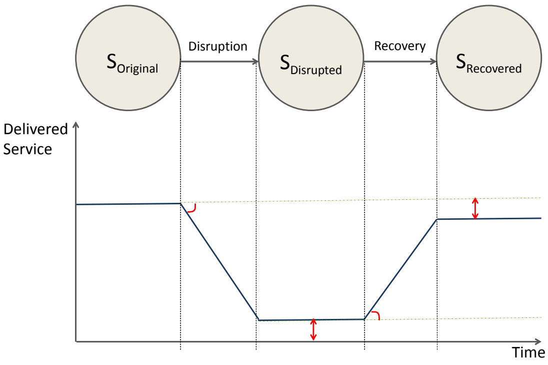

Figure 7: An outage model, as introduced by Henry and Ramirez-Marquez in [201].

An approach encompassing both proactive and reactive aspects is the *autonomic* one, characterized by a control loop including automatic monitoring, analysis, planning, and execution phases [135]. Such approach has been proposed e.g. for threat mitigation in the Internet of Things [38].

8.2. Outage Solutions

In this section, we discuss the solutions proposed to prevent and mitigate outage consequences. We organized them according to (i) the network layer they mainly operate at and (ii) how they perform recovery (in a reactive or proactive way). In Fig. 8 we present the taxonomy we propose in this paper for their solutions. Each solution will be presented and discussed in the following paragraphs.

8.2.1. Solutions at the Physical Layer

In this paragraph, we discuss solutions that have been proposed in literature to deal with outages by mainly operating at the physical layer. These solutions often outline design choices such as redundancy and how to exploit it in an effective manner.

At this level, recovery requires the detection of the damaged components and their replacement or fixing. It is important to define effective strategies to restore the service as soon as possible. For this reason, we will focus on proactive recovery mechanisms (or protection schemes).

Redundancy. Redundancy is a key design feature for fault-tolerant systems. Without a certain degree of (physical and/or logical) redundancy, resiliency can never be achieved.

In [97], Fukada et al. analyzed the impact of the Japanese earthquake on an important national network. They observed that, even though some physical links were damaged, the network connectivity was maintained thanks to two levels of network redundancy, (i) a physical link level redundancy, and (ii) a network topology level redundancy. The former is guaranteed by dual physical links that route along different geographical paths, whereas the latter is provided by redundant multiple loops in the network topology. Cho et al. [62] analyzed the same outage on a different network. They emphasized the importance of redundancy and over-provisioning in the network design as well.

Nonetheless, several works pinpoint limitations of redundancy. In [213], Segovia et al. argued that usually redundancy-based techniques are effective under single-failure scenarios rather than for outages, since the cost of implementing massive redundancy for rarely occurring events is prohibitive. Furthermore, when considering logically-disruptive outages, Dhamdhere et al. [82] claimed that path redundancy does not always guarantee protection, because router misconfigurations could prevent a backup link from coming up. They argue that these failures are non-transient in nature and can only be resolved by the intervention of a network operator. In [247], Wu et al. discovered that, in spite of the apparent physical redundancy, a large number of ASes are vulnerable to a single access link failure; furthermore, BGP policies severely further limit the network resilience under failure. They find that about 35% of the ASes can be disconnected from the rest of the network by a single link failure, which they claim to be the most common failure in today's Internet. An outage can thus simultaneously disrupt a large amount of stub ASes.

Redundancy is also characterized by a trade-off between cost and performance. Therefore, Horie et al. [118] argued that it cannot be applied for outages, because the probability of such failures is quite low and the implementation cost for preparing against such failures is very high. In [213], Segovia et al. reached a similar conclusion: in case of outages, redundancy is not economically sustainable.

To sum up, we argue that *a certain degree of redundancy is a necessary condition to face outages, but it is not sufficient.* This is due to the fact that if a path is disrupted and no other path is available, the only possible solution consists in repairing the disrupted path. If however another path is available, further techniques must be employed to quickly and effectively recover from the outage, because (i) the idea of realizing fully-redundant network systems is not applicable (for economic reasons), and because (ii) redundancy is weak against correlated failures that an outage would be likely to cause.

Link Prioritization. Segovia et al. [213] proposed a protection scheme based on link failures. A network node can fail partially or totally, according to the status of its links. They supposed that all the links in a network have equal probability of being hit by a certain failure, and that it is possible to make them *invulnerable* at a fixed cost per link. In case of outages, several links will be affected at once; assuming that only a fixed budget is available for shielding links, only a limited number of them can be made invulnerable. Accordingly, the authors proposed an optimization model to decide which links should be part of the set of invulnerable links. Obviously, "invulnerability" is only an idealization, that could be achieved, to a certain extent, with redundant techniques. Having a fixed budget helps avoiding the economic limitation described in the previous paragraph. The combinatorial and non-deterministic nature of the problem requires the introduction of approximate solutions. For this reason, in [213] Segovia et al. proposed two heuristic-based approaches to the problem, whose common purpose is to produce a prioritized list of links to make "invulnerable", from which to choose according to the available budget.

Progressive Recovery Mechanisms. The progressive recovery mechanism proposed by Wang et al. in [240] is somehow complementary to link prioritization. The authors proposed an optimal recovery mechanism to progressively restore the network capacity under fixed budget constraints. This is another example of proactive recovery mechanism.

The Link Prioritization scheme focused on selecting "invulnerable" links, whereas Wang et al. assumed that each network element can fail: the focus in this case is mainly on the identification of an optimal recovery scheme. The metric for choosing the recovery order is the flow capacity of the links. The basic problem has been formulated as an optimization problem based on Mixed Integer Programming (MIP) shown to be NP-hard. The authors proposed heuristic algorithms to solve the problem as well.

In [243], Wang proposed two optimization problems: the first problem considers effective connection recovery when a disruptive event happens, whereas the second one studies network augmentation to build a resilient network against any single region failure. Wang showed how also these problems are NP-hard and require heuristic algorithms to be solved.

Risk-based Resilient Network Design. In [234], Vajanapoom et al. proposed a risk-based resilient network design. The main design problem taken into account is: given a working network and a fixed budget, how to best allocate the budget for deploying a survivability technique in different parts of the network based on the risk management. The authors proposed four risk management based approaches for survivable network design: (i) minimum risk; (ii) minimum-maximum damage; (iii) minimum-maximum risk; and (iv) minimum-RMS damage survivable network.

8.2.2. Solutions at the Data-Link Layer

In this paragraph, we present solutions that mainly operate at the data-link layer. We focus on the adaptation SONET/SDH-like resilience techniques to IP networks. These solutions can be set in a reactive or proactive manner.

In [226], Suwala and Swallow argue that IP traffic can be protected using techniques below layer 3 by considering SONET/SDH-like mechanisms. In this paper, descriptions of linear SONET Automatic Protection Switching (APS) for routers, Resilient Packet Ring protection (RPRP), IP interface bundling (IPIB), and MPLS fast reroute (MPLS-FRR) are covered. The key motivation is that layer 3 solutions are limited by the need to communicate among multiple routers, whereas the aforementioned techniques are not subject to this constraint. APS and RPRP are mechanisms that aim at exploiting redundant paths in a fast and efficient way. These are based on the existence of protection links. IP interface bundles are used to group several physical link into a single virtual link, i.e. a logical link. If one or more physical links fails, traffic can be quickly shifted to other links in the bundle. This mechanism is transparent to the routing protocol. Nevertheless, disruptive outages are likely to cause the failure of all the links in the bundle.

MPLS-FRR aims at repairing damaged tunnels by creating a "bypass tunnels" that replace the failed links. Bypass tunnels simply represent other MPLS TE tunnels; these can be set in a reactive or proactive manner, and usually work by adding one more labels to the packets traveling on a primary tunnel, in order to divert traffic onto the bypass tunnel. However, as observed in [245] [24], while MPLS-FRR is the major technique currently deployed to handle network failures, practical limitations still exist, in terms of complexity, congestion, and performance predictability.

8.2.3. Solutions at the Network/Application Layers

In the following paragraphs, we present solutions that directly operate at the network layer or at the application layer, typically making use of overlay networks. Some solutions operate at both layers.

Resilient Routing Reconfiguration (R3). At the end of the previous paragraph we have described challenges that have still be addressed in MPLS-FRR. In [245], Wang et al. argued that, in early 2010, two of the largest ISPs in the world gave instances of severe congestion caused by FRR in their networks. Motivated by the aforementioned limitations, Wang et al. [245] proposed Resilient Routing Reconfiguration (R3), a novel routing protection scheme. R3 can quickly mitigate failures by pre-computing forwarding table updates for the loss of each link. They argue that R3 is (i) congestion-free under a wide range of failure scenarios; (ii) efficient w.r.t. router processing overhead and memory requirements; (iii) flexible in accommodating diverse performance requirements; and (iv) robust to traffic variations and topology failures. This routing scheme exploits an operations research approach: R3 strongly depends on a novel technique for covering all possible failure scenarios with a compact set of linear constraints on the amounts of traffic that should be rerouted. The authors formulate a linear programming model to characterize optimal rerouting, and implement R3 protection using MPLS-ff, a simple extension of MPLS, while the base routing can use either OSPF or MPLS. The authors implemented R3 on Linux and claim that their Emulab evaluations and simulations based on real Internet topologies and traffic traces show that R3 achieves near-optimal performance.

BGP Modifications. Several studies proposed modifications to BGP, based on the observation that it is the de-facto standard for inter-domain routing in the Internet. In the previous sections, we only considered BGP as an analysis tool. In this paragraph we report studies for which BGP is the object of the analysis instead. The perspective is this: after the occurrence of an outage, BGP is responsible for the recovery of interdomain connectivity (exploiting a reactive recovery mechanism); is it effective enough? Is it fast enough? There is a huge amount of studies, in literature, that cover the problem of BGP convergence times after the occurrence of a (large) failure; for example, see [208, 210, 209], [151, 109, 50, 188].

Another issue often raised in literature is related to the fact that BGP is still based on 'the honor system' [70], that is, any organization on the Internet can easily assert that it owns the IP addresses of any other organization, and it is up to the receivers of these BGP updates to decide whether to trust the information or not. Hence, a third question on BGP can be considered: is it secure enough? The answer is easily no. Several BGP security vulnerabilities are presented in RFC 4272 [173], and although there are proposals for secure BGP versions (e.g., BGPSec [149]), it is easy to understand the difficulties that arise when trying to actually introduce them in the Internet. A survey on securing BGP is provided in [123]. Furthermore, illicit prefixes could be imported/exported also as a consequence of misconfigurations rather than intentional attacks. For example, Mahajan et al. [158] showed that misconfiguration errors are pervasive, with 200-1200 prefixes suffering from misconfiguration each day.

In conclusion, BGP has still some problems. It may be slow to converge and it is not secure at all. Modifications to BGP have been often proposed in literature, but the reader will easily understand the difficulty of modifying the way routers work in the whole Internet.

LIFEGUARD. In [134], Katz-Bassett et al. proposed LIFEGUARD, a system for automatic failure localization and remediation. It uses active measurements and a historical path atlas to locate faults. The authors propose an approach to quickly bypass disrupted areas. The key idea is to give data-centers and other well-provisioned edge networks the ability to repair persistent routing problems, regardless of which network along the path is responsible for the outage. If some alternative, working policy-compliant path can deliver traffic during an outage, the data center or edge network should be able to cause the Internet to use it. The interesting characteristic of LIFEGUARD is that it is actually deployable on today's Internet. The authors argue that existing approaches often allow an AS to avoid problems on its forward paths to destinations, but little control over the paths back to the AS is provided. LIFEGUARD provides reverse path control through BGP poisoning. Specifically, the origin AS insert the (partially) disrupted AS into its path advertisements. This way, it appears that the disrupted AS has already been visited. When the announcements reach the disrupted AS, BGP's loop-prevention mechanism will drop the announcement. Hence, networks that would have routed through the disrupted AS will only learn of other paths. The authors show that LIFEGUARD's rerouting technique finds alternate paths 76% of the time. This mechanism seems to provide two main advantages: (i) the networks that use LIFEGUARD are less affected by outages in the rest of the Internet, and (ii) the amount of traffic in the disrupted area decreases, so that the "survived" network capacity can be better exploited by people in the affected area.

RiskRoute. In [91], Eriksson et al. proposed RiskRoute, a routing framework for mitigating network outage threats. The authors introduce the concept of bit-risk miles, the outage risk weighted distance of network routes. This measure is proposed with respect to four properties: (i) geographic distance; (ii) outage impact; (iii) historical outage risk; (iv) immediate/forecast outage risk. RiskRoute is a routing framework based on the definition and opportune use of bit-risk miles. Specifically, the objective is the minimization of the bit-risk miles of routes in a network infrastructure. RiskRoute can be used to reveal the best locations for provisioning additional PoP-to-PoP links, or new AS peering connections that would be advisable to establish, etc.

Eriksson et al. assessed and evaluated RiskRoute, determining the providers that have the highest risk to disaster-based outage events. They are also able to provide provisioning recommendations for network operators that can in some cases significantly lower bit-risk miles for their infrastructures. Under this perspective, it can be used as a high-level resilient routing framework. Further considerations on resilient routing frameworks can be found in [189], in which Pei et al. provided a survey on research efforts in the direction of enhancing the dependability of the routing infrastructure.

Geographically Informed Inter-Domain Routing (GIRO). In [179], Oliveira et al. proposed a new routing protocol and address scheme, called GIRO ("Geographically Informed

Inter-Domain ROuting"). GIRO uses geographic information to assist (and not replace) the provider-based IP address allocation and policy-based routing. The authors argue that, incorporating geographic information into the IP address structure, GIRO can significantly improve the scalability and performance of the global Internet routing system. Within the routing policy constraints, geographic information enables the selection of shortest available routing paths. The authors argue that traversing longer distance (and more routing devices) is likely to increase the chance of outage, as well as other performance metrics. For this reason, we presented GIRO in this context, even though its main focus is not on proposing an outage-aware routing scheme.

Resilient Routing Layers (RRL). Resilient Routing Layers (RRL) were firstly introduced in [112] by Hansen et al. Given a network topology, RRL assumes a certain number of failures in the network node(s). Each failure scenario is associated with a different topology, that can be derived from the original one. On this, RRL pre-calculates an opportune routing table. These are called Routing Layers (RL). Each RL attempts to configure the network topology so to re-route traffic preserving the reachability of other parts of the network. All nodes in the network share the calculated RLs, and select the same single RL when a network failure occurs. RRLs represent an example of proactive recovery mechanism, because "backup routes" are pre-calculated. In [111], Hansen et al. demonstrated how their RRL method can be used as a tool for recovery from outages.

RRL with Overlay. An adaptation of RRLs to accommodate large-scale failures is also provided for example in [118] by Horie et al., who proposed an overlay network approach. In fact, they argue that using an overlay network is convenient for different reasons: first of all, methods based on overlay networks can be easily and quickly deployed, since no standardization process is needed. Furthermore, they claim that the application-level traffic routing performed by overlay routing can overcome the shortcomings in policy-based BGP routing.

Underlays Fused with Overlays (UFO). In [250], Zhu et al. proposed UFO, a Resilient Layered Routing architecture. The authors argue that common routing protocols and overlay networks have their pros and cons. Therefore, they propose an architecture that tries to achieve the best of both worlds. In fact, they argue that common routing protocols scale relatively well, but do not react quickly to changing network conditions. On the contrary, overlay networks can quickly respond to changing network conditions, but they do not scale very well, since they rely on aggressive probing. UFO stands for Underlays Fused with Overlays to stress its two-layered architecture. It provides the abstraction of a subscription service for network events occurring along the underlying paths between the overlay nodes. Explicit cross-layer notifications helps improving the efficiency of reactive routing at the overlay layer without compromising scalability, since notification messages are propagated only to the participating overlay nodes.

Further studies on RRLs have been proposed; the reader may also refer to [142, 143]. In this paragraph we presented the concept of RRL. We have presented a 'native RRL' solution [111], an 'overlay-based' solution [118], and a compromise between the two [250]. Pros and cons of the different approaches were presented.

Network Resilience through Multi-Topology Routing. In [167], Menth and Martin propose the use of Multi-Topology Routing (MTR) to achieve a higher network resiliency. MTR is an optional mechanism within IS-IS, used today by many ISPs. It provides several different IP routing schemes within one network. In [167], Menth and Martin enhance MT routing to provide resiliency. The key idea is simple: under normal network conditions, a basic MTR scheme is used. If a node detects the outage of one of its adjacent links or neighbor node, it deviates all traffic that has to be sent according to the routing table over this failed element to another interface over an alternative routed provided by a another MTR scheme. Under this perspective, MTR is somehow similar to RRL. The authors argue that their solution can guarantee a failover time comparable with MPLS solutions based on explicit routing.

Resilient Overlay Networks (RON). Resilient Overlay Networks (RON) have been first proposed in [37]. They rely on the overlay networks advantages discussed in the previous paragraphs. A RON allows Internet applications to detect and recover from path outages within several seconds, whereas it can require several minutes to current wide-area routing protocols. RONs have a relatively simple conceptual design: several RON nodes are deployed at various locations on the Internet and form an application-layer overlay that cooperates in routing packets. Since a RON provides a classic overlay architecture, its main con is the low scalability [250]. For further information/details about the impact of geographically correlated failures on overlay-based data dissemination, the reader may refer to [235].

Content Delivery Networks (CDN). The primary goal of Content Delivery Networks is to quickly and reliably deliver contents to end-users. At the same time, several studies (e.g., [176]) demonstrated how the use of CDNs can actually help preventing or mitigating pontential Internet outages caused by malicious behaviours. Indeed, CDNs typically replicate contents across different geographically distributed servers. When a user requests a content, the CDN typically redirects this request towards the *best* server according to different policies. In this way, the CDN can dramatically mitigate the consequences of DDOS attack, since the high number of requests generated by the compromised machines would be load balanced across different servers. This complementary service provided by CDN is becoming more and more critical in today networks.

8.2.4. Overall Resilient Strategies

At the highest level, we can define overall resilient strategies, that provide recommendations and guidelines to obtain a more resilient network.

Resilinets. In [224] [22], Sterbenz et al. presented Resilinets, an overall strategy to achieve a higher network resiliency. Resilinets is based on a number of axioms, strategies, and principles. Basically, Resilinets considers a two-phase strategy called D2R2 + DR, as shown in Fig. 9. The first active phase, D2R2, stands for “defend, detect, remediate, recover”; it is the inner control loop and describes a set of activities that are undertaken in order for a system to rapidly adapt to challenges and attacks and maintain an acceptable level of service. The second active phase, DR, stands for “diagnose and refine”; it is the outer loop that enables longer-term evolution of the system in order to enhance the approaches to the activities of phase one.

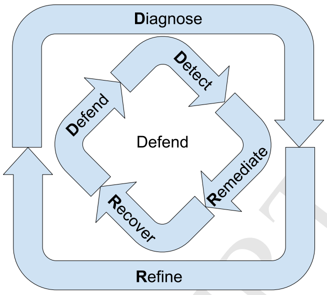

Figure 9: The Resilinets strategy as shown and described in [224].

9. Open issues

In the previous sections, we have discussed the techniques and methodologies that have been proposed to face the main challenges in dealing with Internet outages, i.e. (i) dissect specific outage episodes; (ii) systematically detect them; (iii) quantify their impact on a network; (iv) understand the resilience of the network; (v) assess the overall risk of the infrastructure to this type of disruptive events; (vi) recover from them or prevent and mitigate their impact.

Based on our elaboration of the literature, we highlight here the open issues requiring additional research on the short and long term. In Tab. 6 we map such open issues on the main challenging problems that informed the structure of our literature research. It is evident that some of the open issues span the whole spectrum of Internet outages challenges (namely, the lack of common methodologies and metrics), while most regard one or few aspects alone. In the following we cover each open issue in detail.

9.1. Common definitions and metrics

While a structured approach to failures of digital systems at large has been presented time ago ([217], currently at the 3rd edition, and [41]), still much variability is present in terminology, specifically reagarding failures in networking and Internet outages. The lack of widely accepted definitions of important terms represents an issue that significantly slows down the research in this field. For example, very often in literature resilience and fault-tolerance are considered synonyms and only few studies attempt to make a difference between resilience and survivability. Since metrics are derived from definitions, it is not surprising the lack of a widely accepted framework of formally defined metrics for quantifying the resilience of an IP network to similar disruptive events. Similarly, we also noticed the lack of shared metrics when quantifying the impact of Internet outages: very often only a qualitative evaluation of the impact is carried out. Some other works defined their own metrics, thus

		Challenging problems					
		Analyze	Detect	Quantify Impact	Quantify network robustness	Assess risk	Survive and mitigate
Open issues	Common methodologies and metrics	✗	✗	✗	✗	✗	✗
	Basic techniques		✗				
	Methodologies for outage analysis	✗					
	Validation of outage detection systems		✗				
	Availability of datasets	✗		✗	✗	✗	
	Models for risk assessment	✗				✗	
	Securing inter-domain routing						✗
	Cloud resilience				✗		✗
	Wired-cum-wireless networks resilience	✗			✗		✗
	Coordination	✗	✗				✗

Table 6: Mapping *open research issues* on *challenging problems* addressing Internet Outages.

preventing from a systematic approach to the subject. Clearly, this strongly weakens our ability to compare the several network disruptive events detected or analysed in literature. Finding a convergence point is definitely an important open issue.

Along this direction, we also believe that an effort should be made to define *subjective* metrics when quantifying the impact of an outage. Indeed, in the last years, Quality of Experience is attracting more and more interest compared to Quality of Service. The latter focus on objective parameters that do not necessarily reflect the quality *perceived* by the end-user which is essentially what a service provider really cares about. Similarly, we believe that an effort should be made to evaluate the impact of an outage in terms of how users perceive it. For example, whether or not the performance perceived by end-users significantly degrades should be an important aspect to consider when evaluating the severity of a network outage.

9.2. Basic techniques

In Sec. 2, we provided an overview of the basic techniques and data sources commonly used in the outage-related literature. Unfortunately, these limited tools potentially lead to an incomplete or inaccurate knowledge about the network status. Drawing conclusions based on this knowledge is thus a process prone to errors. A perfect example in this regard is represented by the network topology models. In order to detect outages, evaluate their impact, or estimate the resilience of a network, several studies rely on the knowledge of the underlying network topology. When these studies are applied to real Internet infrastructure, the problem of gathering knowledge about the topology arises [96, 67]. Current data sources (e.g., Traceroute or BGP) cause the obtained topology model to be inaccurate and incomplete. Incomplete topologies do not contain all the nodes or links of the actual network. For instance, BGP-derived AS-level topologies are well known to be incomplete [108, 180, 113, 60]. Inaccurate topologies contain incorrect information such as non-existing nodes or links (due to anonymous routers [168], hidden routers [192], uneven and per-packet load balancing [39], third-party addresses [161], unresolved IP aliasing [163], sampling biases [145], etc.), or incorrect node or link attributes such as the locations of PoPs [93], or the capacity of links. Some studies proposed approaches to reduce the inaccuracies of fault diagnosis algorithms in the presence of partial topology information. For example, in [117] Holbert et al. proposed a strategy to infer the missing portions of a topology, based on the use of UDP datagrams in case some routers in the network drop ICMP messages. Many other countermeasures have been developed in the field of Internet topology discovery [246, 206], but only some of them have been adopted in Internet outage-related works. Based on our studies, we argue that it is better to work on incomplete topologies. This way, if in the real network contains unreported edges or nodes, the actual outage impact and network resilience will be respectively lower and higher than the estimated ones. As for detection, if the topology contains non-existing edges or nodes, conclusions drawn when traffic is re-routed through non-existing links, routers, or ASes are biased. We believe that further studies on the effects of inaccuracy and incompleteness in outage-related works should be made. We also argue that a quantitative evaluation of incompleteness and inaccuracy of the topology models exploited in these works might be very helpful to estimate a sort of “confidentiality interval” for impact and resilience metrics.

9.3. Methodologies for outage analysis

Several works focused on the analysis of specific episodes of large-scale Internet outages. These analyses are of the utmost importance since investigating specific events increases our understanding of the scope and the consequences of similar disruptive events as well as the utility of the instruments for gaining insight on them. However, critically reviewing these works (Sec. 3), we noticed the lack of a widely accepted methodology: each work adopted its own approach built on top of a subset of basic techniques and data sources to derive insights and conclusions. We believe that developing a widely accepted structured approach for this type of analysis is an important future direction in this field. Nowadays, the validity and scope of the findings and conclusions drawn starting from different methodologies are very hard to quantify. “Which specific data source(s) should be taken into account? Which specific tools and how to configure it? Given the measurement outcomes also in the light of the limited visibility on the ongoing events in the Internet, to what extent the achieved conclusions can be considered valid?” are only some of the questions that the research community should address along this direction.

9.4. Validation of outage detection systems

The numerous outage detection systems developed during the last years continuously or opportunistically monitor the network with focused or general purpose measurements reporting thousands of network disruptions over time. These systems are particularly helpful to understand the frequency and location also of the small outages, a type of disruption representing a common threat to network operators. Unfortunately, the general climate of collaboration and competition among the different networks in the Internet greatly disincentives the network operators to share data or knowledge on the threats occurring in their networks. Accordingly, it is very hard to validate the outcomes of these systems: not being able to enumerate and deepen false alarms as well as undetected threats, researchers can not properly evaluate and improve their outage detection systems. This challenge affects the entire research community involved in Internet measurements and represents a severe obstacle to the advancement in the Internet outage-related field.

9.5. Availability of datasets

In analogy to the issue of Section 9.4, accessing detailed data about Internet outages is not an easy task to accomplish. Indeed, focused datasets are rarely available. *The Outages Archive* [14] and the *Internet outage dataset* [3] are the only two relevant examples currently available, to the best of our knowledge. The former collects messages exchanged by operators and practitioners through the Outages mailing list since 2006. The latter contains the results of measurements campaings aiming at investingating generic or specific Internet outages (outages clustering, outages detection, address reachability, studies of hurricanes, etc.). While, in the latter, data is structured and organized at different levels of abstraction, the former barely provides a collection of e-mails, thus requiring additional effort for being processed (e.g., as done by Banerjee et al. [43]). According to the limited information available, often, inquiries about outages are carried out leveraging data that can detect outages indirectly, such as path-tracing data (usually obtained through traceroute) or BGP-related information. Relevant examples are the datasets made available by CAIDA [6] or by the ANT Lab [2]. This kind of datasets proved a valuable source of information, but require full understanding of the mechanism, the procedures adopted for the analyses and the conditions in which they are carried out, thus often needing proper assumptions to be leveraged. Due to the inherent complexity, we refer to the specific studies for more details about these datasets.

Finally, we believe that relevant examples for existing outage monitoring and analysis systems are also worth to be mentioned, although not publicly providing datasets at time of writing. CAIDA recently released a publicly-accessible operational prototype for IODA [10], a system aimed at monitoring the Internet in near-realtime (leveraging information related to BGP, IBR, and active probing), with the goal of identifying macroscopic Internet outages significantly impacting an AS or a large fraction of a country.

9.6. Models for Risk Assessment

In Section 7, we discussed the notion of risk assessment applied to data-network infrastructures. We argue that the main problem is modelling the outage. We reviewed a number of studies that modelled an outage as a disk, which is an oversimplification if compared to the models commonly exploited in civil engineering. The development of appropriate outage models appears to be a hot open issue. Indeed, for instance, being able to perform a comprehensive risk assessment study on a network infrastructure would enable companies insuring them. To this aim, *predictive models* would be of utmost importance, to evaluate the possible evolution of the network infrastructure given the knowledge of its current status, and to assess the (improved) consequences of a risk mitigation strategy. Unfortunately, most of the outages causes are specifically hard to predict. Regarding natural disasters there is a long history of scientific research addressing predictions of earthquakes, hurricanes, and similar catastrophic events [177], and it is a hard scientific quest still open [190]. Regarding human-related causes, the prediction of malicious activities is also specifically hard: some preliminary work has been done modeling the psychology of an unfaithful employee (a so-called *insider threat*) [128], although it is additionally hindered by the real-world applicability and privacy concerns related to the required surveillance of workers. All these difficulties add to the inherently hard problem of modeling a highly dynamic distributed system—the Internet—whose behavior emerges from the communications of millions human users and increasingly hard-to-enumerate machines.

9.7. Securing inter-domain routing

An important future direction is securing the inter-domain routing. Indeed, this solution may effectively prevent Internet outages caused by malicious behaviors such as prefix hijacking attacks or accidental misconfigurations as incorrect prepending. Although BGP is known to be affected by these and other security vulnerabilities (see RFC 4272 [173] and Huston et al. [123]), securing inter-domain routing still represents an open issue. Indeed, modifying the way the whole Internet works proved to be an extremely complex process also often referred to as Internet ossification. Therefore, we highlight a quite obvious yet essential future direction: finding applicable ways to make prefix advertisements secure. We argue that this could lead to a dramatic decrease of outages causing logical disruption.

9.8. Cloud Resilience

Directly related to what we just observed in the previous paragraph, we can reach a less obvious future direction. While modifying the way all routers in the Internet work is very hard, it is relatively easy for data-center managers and cloud service providers to modify the way their border routers work in order to achieve higher resilience against outages occurring in the public Internet. To a certain extent, this is true for ISPs as well. Based on this, routing schemes or other solutions aiming at improving the service resilience could be deployed. An example was presented earlier with LIFEGUARD [134]. The unilateral deployment of failure avoidance techniques would avoid the previously described limitations. We claim that a major effort should be put along this direction, motivated by the strict SLAs that must be met by these service providers. To this aim, promising technologies are emerging related to Software Defined Networks (SDN), enabling the needed experimentation (and evolvability) on operating data center networks [76].

9.9. Wired-cum-Wireless Networks Resilience

Wireless networks and mobile services are becoming the more and more a part of the global communication infrastructure in the *convergence* of networks towards the *Internet of Things*. Despite this, the resilience of wide area networks comprising wireless paths has not been studied extensively as their steadily growing importance would require. Before the integration of mobile telephone networks and the Internet, such analyses have been performed mainly focusing on voice service and measured in terms of *blocked calls*, although (low bandwidth-) data communications were cited but not estimated in the outage analysis: we refer to [220] for an early analysis of outages and wireless network survivability in a pre-convergence scenario. Other works have addressed the performance of wireless metropolitan area networks [130, 131], or evaluated (in simulation) the interdependence between electrical power distribution networks and mobile networks in case of faults [120]. Finally, due to their ease of deployment wireless networks have been considered as means for backup and mitigation in case of disasters [61, 100, 203].

Compared with the scientific corpus available for *wired* wide-area-networks, the paucity of research on outage analysis, mitigation and prevention for networks including wireless paths is evident. In our opinion this can be a symptom that several already mentioned open issues (especially the lack of common definitions and metrics and the lack of models for risk assessment) have so far limited the interest and research of this aspect. Nevertheless we expect more attention to be drawn to this specific topic in sight of the spread of *Industry 4.0* scenarios [239].

9.10. Coordination

Small outages located inside a given network can easily be detected, located, and fixed by the corresponding operator. On the other hand, large-scale outages involving final users, content providers, and multiple transit networks definitely require the coordinated action of different entities forced to share knowledge and data to find and solve the issue. The outage mailing lists [14] represents a common way for network operators to advertise or enquiry about Internet outages. Compared to the highly sophisticated technology employed in their infrastructure, this coordination tool appears somehow anachronistic and largely ineffective. Hence, we noticed large room for improvements on this side that can be achieved with widely accepted (i) best practises and (ii) standard procedures based on (iii) a widespread measurement infrastructure providing all the required information about the network status in real time. Systems like COVE [251] and perfSONAR [230, 110] represent a first important step along this direction.

10. Conclusions

Outages on critical infrastructures such as power grid, water distribution and transportation systems, rise severe concerns due to the potential large impact on our daily life. This is more and more true for the Internet also. Accordingly, it is not surprising that this topic attracted great interest from research and network operators community in the last years. In this paper, we have provided a detailed analysis of the state of the art of research related to Internet outages. Fig. 1 reports a quick snapshot of the main aspects we have considered. Our final goal was providing a comprehensive and elaborated view on Internet outages and all the related aspects relevant for researchers approaching this wide research area.

In more detail, we analysed and systematically revised the large and scattered body of works on this topic, providing an extensive and carefully organized picture of the literature related to Internet outages. Moreover, to the best of our knowledge, we provided several innovative contributions achieved through this study: (i) a road to systematically study Internet outages; (ii) a characterization of the causes of these disruptive events; (iii) a classification of the basic techniques used by researchers; (iv) a systematic analysis of works dissecting specific Internet outage episodes, underlying common practices, weaknesses, and differences of the proposed approaches; (v) a general approach for outage detection; (vi) a classification of approaches for outage impact evaluation; (vii) an apportionment of definitions and metrics for evaluating the resilience of a communication network; (viii) a systematic discussion on the assessment of risk of networks to outages; (ix) an overview of the solutions proposed to prevent, mitigate, or resolve Internet outages and their consequences. (x) a detailed analysis of open issues and future directions in this research field.

The paper constitutes an important starting point for researchers willing to simply understand or to contribute to this wide and articulate research area.

Acknowledgments

We thank the former student Massimiliano Molinari for participating in the collection of the surveyed literature and for his help with this work. This work is partially funded by art. 11 DM 593/2000 for NM2 srl (Italy).

References

[1] Regional Internet Registry. http://en.wikipedia.org/wiki/Regional_Internet_registry.

[2] Ant lab, dataset. https://ant.isi.edu/datasets/index.html, .

[3] Ant lab, internet outage dataset. https://ant.isi.edu/datasets/outage/index.html, .

[4] Archipelago Measurement Infrastructure. http://www.caida.org/projects/ark.

[5] The CAIDA AS Relationships Dataset. http://www.caida.org/data/as-relationships.

[6] Caida data. https://www.caida.org/data/.

[7] Cisco IOS NetFlow. http://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow.

[8] MaxMind GeoLite. http://dev.maxmind.com/geoip/legacy/geolite.

[9] Hengchun earthquake. http://en.wikipedia.org/wiki/2006_Hengchun_earthquake.

[10] Caida ioda, internet outage detection and analysis. http://www.caida.org/projects/ioda/.

[11] Hurricane Katrina. http://en.wikipedia.org/wiki/Hurricane_Katrina.

[12] NANOG Mailing list. https://www.nanog.org.

[13] The NIST Definition of Cloud Computing. http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.

[14] The outages archive. https://puck.nether.net/pipermail/outages.

[15] Packet Clearing House (PCH), Route Views archive. http://www.pch.net/documents/data.

[16] Quagga Routing Suite. https://www.quagga.net/.

[17] Oracle dyn, blog. http://www.renesys.com/blog.

[18] Packet design, route explorer. http://www.packetdesign.com/products/route-explorer, .

[19] University of oregon, route views project. http://www.routeviews.org, .

[20] Blackout-2005: Widespread power outages in Moscow. http://english.pravda.ru/news/russia/25-05-2005/62978-0.

[21] Hurricane Sandy. http://en.wikipedia.org/wiki/Hurricane_Sandy.

[22] ResiliNets: Multilevel Resilient and Survivable Networking Initiative Wiki. Information Technology and Telecommunications Center (ITTC) at the University of Kansas. http://wiki.ittc.ku.edu.

[23] Tohoku earthquake and tsunami. http://en.wikipedia.org/wiki/2011_Tohoku_earthquake_and_tsunami.

[24] http://www.wandl.com/html/support/papers.

[25] *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*. IEEE, 1990.

[26] The Internet of Things. Technical Report 27441, International Telecommunication Union, November 2005.

[27] Timeline for cnmi outage. https://www.subcableworld.com/scw-newsfeed/marine-services/timeline-for-cnmi-outage, 2015.

[28] Ripe routing information service (ris). http://www.ripe.net/data-tools/stats/ris, 2015.

[29] 45 days in cyber darkness: Cameroon switches off internet for second time in the year. https://ifex.org/cameroon/2017/11/20/internet-shutdown, 2017.

[30] E. Aben. Ripe atlas: Hurricane sandy and how the internet routes around damage. https://labs.ripe.net/Members/emileaben/ripe-atlas-hurricane-sandy-global-effects, 2012.

[31] access now. accessnow blog. https://www.accessnow.org/blog.

[32] I. S. F. access now. Submission to the un human rights committee on concerns and recommendations on cameroon (open letter). https://internetwithoutborders.org/fr/wp-content/uploads/sites/2/2017/09/UNHRCommittee-Submission-Cameroon.pdf, sep 2017.

[33] G. Aceto and A. Pescapè. Internet censorship detection: A survey. *Computer Networks*, 83:381–421, 2015. doi: 10.1016/j.comnet.2015.03.008. URL http://dx.doi.org/10.1016/j.comnet.2015.03.008.

[34] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman. Network vulnerability to single, multiple, and probabilistic physical attacks. *IEEE MILCOM*, 2010.

[35] M. Al-Kuwaiti, N. Kyriakopoulos, and S. Hussein. A comparative analysis of network dependability, fault-tolerance, reliability, security, and survivability. *IEEE Communications Surveys & Tutorials*, 11(2):106–124, 2009.

[36] R. F. Albers. Outage reporting and customer notification team. Technical report, Bell Atlantic, jan 1996.

[37] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient overlay networks. In *Proceedings of the ACM SOSP Conference*, 2001.

[38] Q. M. Ashraf and M. H. Habaebi. Autonomic schemes for threat mitigation in internet of things. *Journal of Network and Computer Applications*, 49:112–127, 2015.

[39] B. Augustin, T. Friedman, and R. Teixeira. Measuring multipath routing in the internet.

[40] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with paris traceroute. *ACM IMC*, 2006.

[41] 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*, 1(1):11–33, 2004.

[42] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the internet. In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 265–276. ACM, 2007.

[43] R. Banerjee, A. Razaghpanah, L. Chiang, A. Mishra, V. Sekar, Y. Choi, and P. Gill. Internet outages, the eyewitness accounts: Analysis of the outages mailing list. *Passive and Active Measurement Conference (PAM)*, 2015.

[44] S. Banerjee, S. Shirazipourazad, and A. Sen. Design and analysis of networks with large components in presence of region-based faults. *Communications (ICC), 2011 IEEE International Conference on. IEEE*, 2011.

[45] A. Bavier, M. Bowman, B. Chun, D. Culler, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak. Operating system support for planetary-scale network services. *NSDI*, 2004.

[46] K. Benson, A. Dainotti, and K. Claffy. Gaining insight into as-level outages through analysis of internet background radiation. *CoNEXT Student'12*, 2012.

[47] Z. S. Bischof, J. S. Otto, and F. E. Bustamante. Distributed systems and natural disasters: Bittorrent as a global witness. *ACM SWID*, 2011.

[48] D. G. Blogs. http://www.renesys.com/2005/12/peering-the-fundamental-archit, 2005.

[49] I. S. F. . I. W. Borders. Censure internet archives. https://internetwithoutborders.org/category/censure-internet/, 2018.

[50] A. Bremler-Barr, Y. Afek, and S. Schwarz. Improved bgp convergence via ghost flushing. *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies, vol. 2. IEEE*, 2003.

[51] J. L. S. Brooks. I-seismograph: Observing and measuring internet earthquakes. In *Infocom 2011 proceedings*, 2011.

[52] M. Brown and E. Zmijewski. Pakistan telecom hijacks youtube: Or how to syn-flood dos yourself while annoying everyone on the planet. *Renesys Corporation*, 2008.

[53] A. B. J. E. Burgess. Local and global responses to disaster: #eqnz and the christchurch earthquake. In *Dis-*

aster and Emergency Management Conference, Conference Proceedings. Vol. 2012. AST Management Pty Ltd, 2012.

[54] R. Bush, O. Maennel, M. Roughan, and S. Uhlig. Internet optometry: Assessing the broken glasses in internet reachability. *IMC'09*, 2010.

[55] D. R. C. F. E. Bustamante. Taming the torrent. *SIGCOMM'08*, 2008.

[56] M. Caesar, L. Subramanian, and R. H. Katz. Towards localizing root causes of bgp dynamics. *Computer Science Division, University of California*, 2003.

[57] Á. Carrera, C. A. Iglesias, J. García-Algarra, and D. Kolařík. A real-life application of multi-agent systems for fault diagnosis in the provision of an internet business service. *Journal of Network and Computer Applications*, 37:146–154, 2014.

[58] E. K. Çetinkaya and J. P. Sterbenz. A taxonomy of network challenges. *9th International Conference on Design of Reliable Communication Networks (DRCN), pp. 322–330*, 2013.

[59] E. Chan, X. Luo, W. Fok, W. Li, and R. Chang. Non-cooperative diagnosis of submarine cable faults. *Passive and Active Measurement Conference (PAM)*, 2011.

[60] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the sidewalk ends. *CoNEXT'09*, 2009.

[61] F. Chiti, R. Fantacci, L. Maccari, D. Marabissi, and D. Tarchi. A broadband wireless communications system for emergency management. *IEEE Wireless Communications*, 15(3), 2008.

[62] K. Cho, C. Pelsser, R. Bush, and Y. Won. The japan earthquake: The impact on traffic and routing observed by a local isp. In *Proceedings of the ACM Special Workshop on Internet and Disasters (SWID)*, 2011.

[63] D. R. Choffnes, F. E. Bustamante, and Z. Ge. Crowdsourcing service-level network event monitoring. In *ACM SIGCOMM Computer Communication Review*, volume 40, pages 387–398. ACM, 2010.

[64] D. R. Choffnes, F. E. Bustamante, and Z. Ge. Using the crowd to monitor the cloud: Network event detection from edge systems. *In Proc. of ACM SIGCOMM*, 2010.

[65] P. Cholda, A. Mykkeltveit, B. E. Helvik, O. J. Wittner, and A. Jajszczyk. A survey of resilience differentiation frameworks in communication networks. *IEEE Communications Surveys & Tutorials, 4th Quarter*, 2007.

[66] P. Chołda, E. L. Følstad, B. E. Helvik, P. Kuusela, M. Naldi, and I. Norros. Towards risk-aware communications networking. *Reliability Engineering & System Safety*, 109:160–174, 2013.

[67] K. Claffy, Y. Hyun, K. Keys, and M. Fomenkov. Internet mapping: from art to science. *Conference For Homeland Security, 2009. CATCH'09. Cybersecurity Applications & Technology. IEEE*, 2009.

[68] K. Clay. Amazon.com goes down, loses $66,240 per minute. `http://www.forbes.com/sites/kellyclay/2013/08/19/amazon-com-goes-down-loses-66240-per-minute`, 2013.

[69] R. Clayton, S. Murdoch, and R. Watson. Ignoring the great firewall of china. In *Privacy Enhancing Technologies workshop 2006*, 2006.

[70] J. Cowie. China's 18-Minute Mystery. `http://www.renesys.com/2010/11/chinas-18-minute-mystery`, 2010.

[71] J. Cowie. Egypt leaves the internet. `http://www.renesys.com/2011/01/egypt-leaves-the-internet`, 2011.

[72] J. Cowie. Libyan disconnect. `http://www.renesys.com/2011/02/libyan-disconnect-1`, 2011.

[73] J. Cowie, A. Ogielski, B. Premore, E. Smith, and T. Underwood. Impact of the 2003 blackouts on internet communications. *Preliminary Report, Renesys Corporation (updated March 1, 2004),*, 2003.

[74] J. Cowie, A. Popescu, and T. Underwood. Impact of hurricane katrina on internet infrastructure. *Renesys Report*, 2005.

[75] I. Cunha, R. Teixeira, N. Feamster, and C. Diot. Measurement methods for fast and accurate blackhole identification with binary tomography. *IMC'09*, 2009.

[76] B. Dai, G. Xu, B. Huang, P. Qin, and Y. Xu. Enabling network innovation in data center networks with software defined networking: A survey. *Journal of Network and Computer Applications*, 94:33–49, 2017.

[77] A. Dainotti, A. Pescapè, and G. Ventre. Worm traffic analysis and characterization. In *2007 IEEE International Conference on Communications*, pages 1435–1442, June 2007. doi: 10.1109/ICC.2007.241.

[78] A. Dainotti, C. Squarcella, E. Aben, K. Claffy, M. Chiesa, M. Russo, and A. Pescapè. Analysis of country-wide internet outages caused by censorship. *ACM IMC*, 2011.

[79] A. Dainotti, R. Ammann, E. Aben, and K. Claffy. Extracting benefit from harm: Using malware pollution to analyze the impact of political and geophysical events on the internet. *ACM SIGCOMM Computer Communication Review, Volume 42, Number 1*, 2012.

[80] R. C. I. Damgaård, M. Computation, an Introduction, . in *Contemporary Cryptology and pp. 41–87*, and 2005. *Contemporary Cryptology, pp. 41–87*, 2005.

[81] S. Deshpande, M. Thottan, T. K. Ho, and B. Sidkdar. An online mechanism for bgp instability detection and analysis. *IEEE transactions on computers, Vol. 58, No. 11*, 2009.

[82] A. Dhamdhere, R. Teixeira, C. Dovrolis, and C. Diot. Netdiagnoser: Troubleshooting network unreachabilities using end-to-end probes and routing data. *CoNEXT'07*, 2007.

[83] T. N. Dihn, Y. Xuan, M. T. Thai, P. M. Pardalos, and T. Znati. On new approaches of assessing network vulnerability: Hardness and approximation. *IEEE/ACM Transactions on Networking, vol. 20, no. 2*, 2012.

[84] F. Dikbiyik, M. Tornatore, and B. Mukherjee. Minimizing the risk from disaster failures in optical backbone networks. *Journal of Lightwave Technology*, 32: 3175–3183, 2014.

[85] E. G. X. Dimitropoulos. Classifying internet one-way traffic. In *Proceedings of the 2012 ACM conference on Internet measurement conference, pp. 37-50. ACM,*, 2012.

[86] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Uffaker, Y. Hyun, K. Claffy, and G. Riley. As relationships: Inference and validation. *ACM SIGCOMM Computer Communication Review, vol. 37, no. 1*, 2007.

[87] M. Djatmiko, D. Schatzmann, Z. Dimitropoulos, A. Friedman, and R. Boreli. Collaborative network outage troubleshooting with secure multiparty com-

putation. *IEEE Communications Magazine*, 2013.

[88] D. Dolev, S. Jamin, O. Mokryn, and Y. Shavitt. Internet resiliency to attacks and failures under bgp policy routing. *Comput. Netw., vol. 50, no. 16, pp. 3183–3196*, 2006.

[89] N. Duffield. Network tomography of binary network performance characteristics. *Information Theory, IEEE Transactions on, 52(12), 5373-5388*, 2006.

[90] M. M. A. A. El-semary. Vulnerability assessment for mission critical networks against region failures: A case study. In *Proceedings of the 2nd International Conference on Communications and Information Technology (ICCIT)*, 2012.

[91] B. Eriksson, R. Durairajan, and P. Barford. Riskroute: A framework for mitigating network outage threats. *CoNEXT'13*, 2013.

[92] C. C. G. Eysenbach. Pandemics in the age of twitter: Content analysis of tweets during the 2009 h1n1 outbreak. *PloS one, vol. 5, issue 11, e14118*, 2010.

[93] D. Feldman, Y. Shavitt, and N. Zilberman. A structural approach for pop geo-location. *Computer Networks*, 56(3):1029–1040, 2012.

[94] G. G. R. Filippini. risk assessment and resilience for critical infrastructures. In *Proceedings of Workshop on Risk Assessment and Resilience for Critical Infrastructures*, 2012.

[95] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.

[96] B. D. T. Friedman. Internet topology discovery: a survey. *IEEE Communications Surveys, 4th Quarter 2007, vol. 9, no. 4*, 2007.

[97] K. Fukuda, M. Aoki, S. Abe, Y. Ji, M. Koibuchi, M. Nakamura, S. Yamada, and S. Urushidani. Impact of tohoku earthquake on r&e network in japan. *ACM SWID 2011*, 2011.

[98] M. L. H. R. Gamble. Secagreement: Advancing security risk calculations in cloud services. *2012 IEEE Eighth World Congress on Services*, 2012.

[99] W. N. J. Gaudiot. Network resilience: A measure of network fault tolerance. *IEEE Transactions on Computers, vol. 39, no. 2*, 1990.

[100] B. P. Gautam and K. Wasaki. Using a redundant wi-fi network as an emergency detour route to proactively reduce disaster risk in wakkanai, hokkaido. In *Information Science, Electronics and Electrical Engineering (ISEEE), 2014 International Conference on*, volume 3, pages 1830–1837. IEEE, 2014.

[101] R. Ghosh, F. Longo, V. K. Naik, and K. S. Trivedi. Quantifying resiliency of iaas cloud. *2010 29th IEEE International Symposium on Reliable Distributed Systems*, 2010.

[102] G. Giannopoulos, B. Dorneanu, and O. Jonkeren. risk assessment methodology for critical infrastructure. *JRC Scientific and Policy Reports*, 2013.

[103] P. Gill, M. Schapira, and S. Goldberg. A survey of interdomain routing policies.

[104] P. Gill, N. Jain, and N. Nagappan. Understanding network failures in data centers: Measurement, analysis, and implications. *SIGCOMM'11*, 2011.

[105] K. Glass, R. Colbaugh, and M. Planck. Automatically identifying the sources of large internet events. *Intelligence and Security Informatics (ISI), 2010 IEEE International Conference on. IEEE*, 2010.

[106] S. P. Gorman, L. Schintler, R. Kulkarni, and R. Stough. The revenge of distance: Vulnerability analysis of critical information infrastructure. *Journal of Contingencies and Crisis Management*, 2004.

[107] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. The cost of a cloud: Research problems in data center networks. *ACM SIGCOMM Computer Communication Review, Vol. 39, No 1*, 2009.

[108] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. On the incompleteness of the as-level graph: a novel methodology for bgp route collector placement. In *Proceedings of the 2012 ACM conference on Internet measurement conference, pp. 253-264. ACM*, 2012.

[109] H. Guo, W, Su, H. Zhang, and S. Kuo. On the convergence condition and convergence time of bgp. *Computer Communications 34 (2011) 192–199*, 2010.

[110] A. Hanemann, J. W. Boote, E. L. Boyd, J. Durand, L. Kudarimoti, R. apacz, D. M. Swany, S. Trocha, and J. Zurawski. Perfsonar: A service oriented architecture for multi-domain network monitoring. *Service-Oriented Computing-ICSOC 2005, pp. 241-254. Springer Berlin Heidelberg,*, 2005.

[111] A. F. Hansen, A. Kvalbein, T. Cicic, and S. Gjessing. Resilient routing layers for network disaster planning. In *Proceedings of the International Conference on Networking*, 2005.

[112] A. F. Hansen, A. Kvalbein, T. Cicic, S. Gjessing, and O. Lysne. Resilient routing layers for recovery in packet networks. In *Proceedings of the 2005 International Conference on Dependable Systems and Networks (DSN'05)*, 2005.

[113] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. A systematic framework for unearthing the missing links: Measurements and impact. *NSDI*, 2007.

[114] P. E. Heegaard and K. S. Trivedi. Network survivability modeling. *Computer Networks 53 (2009) 1215–1234*, 2009.

[115] J. Heidemann, L. Quan, and Y. Pradkin. A preliminary analysis of network outages during hurricane sandy. *University of Southern California, Information Sciences Institute*, 2012.

[116] R. Hiran, N. Carlsson, and P. Gill. Characterizing large-scale routing anomalies: A case study of the china telecom incident. *Passive and Active Measurement. Springer Berlin Heidelberg,*, 2013.

[117] B. Holbert, S. Tati, S. Silvestri, and T. L. Porta. On the benefits of network topology inference for fault diagnosis under partial information. *IEEE MILCOM*, 2013.

[118] T. Horie, G. Hasegawa, S. Kamei, and M. Murata. A new method of proactive recovery mechanism for large-scale network failures. *Advanced Information Networking and Applications, International Conference on. Los Alamitos, CA, USA: IEEE Computer Society, 2009, pp. 951–958*, 2009.

[119] J. B. Horrigan. Broadband adoption and use in america", in *Federal Communications Commission*, 2010. *Federal Communications Commission*, 2010.

[120] S. Horsmanheimo, N. Maskey, L. Tuomimäki, H. Kokkoniemi-Tarkkanen, and P. Savolainen. Evaluation of interdependencies between mobile communication and electricity distribution networks in fault scenarios. In *2013 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia)*, pages 1–6, Nov 2013. doi: 10.1109/ISGT-Asia.2013.6698771.

[121] Z. Hu, L. Zhu, C. Ardi, E. Katz-Bassett, H. V. Mad-

hyastha, J. Heidemann, and M. Yu. The need for end-to-end evaluation of cloud availability.” in *Passive and Active Measurement, pp. 119-130. Springer International Publishing*, 2014. *Passive and Active Measurement, pp. 119-130. Springer International Publishing*, 2014.

[122] Y. Huang, N. Feamster, A. Lakhina, and J. Xu. Diagnosing network disruptions with network-wide analysis. *SIGMETRICS'07, June 12–16*, 2007.

[123] G. Huston, M. Rossi, and G. Armitage. Securing bgp - a literature survey. *IEEE Communications Surveys & Tutorials, vol. 13, no. 2, 2nd Quarter*, 2011.

[124] T. B. Institution. Blog posts and reports, search terms: *internet outage censorship*. `https://www.brookings.edu/search/?s=internet+outage+censorship`.

[125] U. Javed, I. Cunha, D. R. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy. Poiroot: Investigating the root cause of interdomain path changes. *SIGCOMM'13*, 2013.

[126] J. Joerin, R. Shaw, Y. Takeuchi, and R. Krishnamurthy. Assessing community resilience to climate-related disasters in chennai, india. *International Journal of Disaster Risk Reduction 1 (2012) 44-54*, 2012.

[127] J. C. J.Saleh. Survivability and resiliency of spacecraft and space-based networks: a framework for characterization and analysis”, in *American Institute of Aeronautics and Astronautics, AIAA Technical Report 2008-7707*, 2008. *American Institute of Aeronautics and Astronautics, AIAA Technical Report 2008-7707*, 2008.

[128] M. Kandias, A. Mylonas, N. Virvilis, M. Theoharidou, and D. Gritzalis. An insider threat prediction model. In *International Conference on Trust, Privacy and Security in Digital Business*, pages 26–37. Springer, 2010.

[129] S. Kaplan and B. J. Garrick. On the quantitative definition of risk. *Risk analysis*, 1(1):11–27, 1981.

[130] R. P. Karrer, I. Matyasovszki, A. Botta, and A. Pescapé. Experimental evaluation and characterization of the magnets wireless backbone. In *Proceedings of the 1st international workshop on Wireless network testbeds, experimental evaluation & characterization*, pages 26–33. ACM, 2006.

[131] R. P. Karrer, I. Matyasovszki, A. Botta, and A. Pescapé. Magnets-experiences from deploying a joint research-operational next-generation wireless access network testbed. In *Testbeds and Research Infrastructure for the Development of Networks and Communities, 2007. TridentCom 2007. 3rd International Conference on*, pages 1–10. IEEE, 2007.

[132] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. Studying black holes in the internet with hubble. *NSDI (pp. 247-262)*, 2008.

[133] E. Katz-Bassett, H. V. Madhyastha, V. K. Adhikari, C. Scott, J. Sherry, P. van Wesep, A. Krishnamurthy, and T. Anderson. Reverse traceroute. *NSDI,*, 2010.

[134] E. Katz-Bassett, C. Scott, D. R. Choffness, I. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy. Lifeguard: Practical repair of persistent route failures. *SIGCOMM'12*, 2012.

[135] J. Kephart and D. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, Jan 2003. ISSN 0018-9162. doi: 10.1109/mc.2003.1160055. URL `http://dx.doi.org/10.1109/MC.2003.1160055`.

[136] Y. H. Khalil, A. Elmaghraby, and A. Kumar. Evaluation of resilience for data center systems. *Proc. ISCC 2008, pp.340-345*, 2008.

[137] A. Khan, T. Kwon, H. Kim, and Y. Choi. As-level topology collection through looking glass servers. In *Proceedings of the 2013 conference on Internet measurement conference, pp. 235-242. ACM*, 2013.

[138] B. Kitchenham. Procedures for performing systematic reviews. *Keele University Technical, Report TR/SE-0401*, 2004.

[139] J. C. Knight, E. A. Strunk, and K. J. Sullivan. Towards a rigorous definition of information system survivability. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX'03). IEEE*, 2003.

[140] C. Köpp. Controlled internet outage monitoring. *Network Architectures and Services*, 2013.

[141] D. R. Kuhn. Sources of failure in the public switched telephone network. *Computer*, 30(4):31–36, 1997.

[142] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne. Fast recovery from link failures using resilient routing layers. In *Proceedings of the 10th IEEE Symposium on Computers and Communications (ISCC 2005)*, 2005.

[143] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne. Multiple routing configurations for fast ip network recovery”, in *IEEE/ACM Transactions on Networking, vol. 17, no. 2*, april 2009. *IEEE/ACM Transactions on Networking, vol. 17, no. 2*, 2009.

[144] M. B. C. Labovitz. Censorship and co-option of the internet infrastructure. *Ann Arbor, 1001:48104*, 2011.

[145] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie. Sampling biases in ip topology measurements. *INFOCOM 2003 vol. 1, pp. 332-341. IEEE*, 2003.

[146] F. Le, G. G. Xie, and H. Zhang. On route aggregation. *SIGCOMM, p. 6. ACM*, 2011.

[147] K. Leelardcharoen. *Interdependent response of telecommunication and electric power systems to seismic hazard*. PhD thesis, Georgia Institute of Technology, 2011.

[148] R. Lemos. Internet routers hitting 512k limit, some become unreliable. `http://ars.to/1r9AbxJ`.

[149] M. Lepinski. BGPSEC protocol specification. `https://tools.ietf.org/html`, 2012.

[150] J. Li, A. Vishwanath, and H. R. Rao. Retweeting the fukushima nuclear radiation disaster. *Communications of the ACM, vol. 57, no. 1*, 2014.

[151] Q. Li, M. Xu, J. Wu, P. P. Lee, and D. M. Chiu. Toward a practical approach for bgp stability with root cause check. *J. Parallel Distrib. Comput. 71 (2011) 1098–1110*, 2011.

[152] R. Li, X. Wang, and X. Jiang. Network survivability against region failure. *Proc. 2011 Int. Conf. Signal Processing, Commun. and Comput. (ICSPCC), pp.1-6, 14-16*, 2011.

[153] Y. Liu, Z. Luo, R. K. C. Chang, and J. Su. Characterizing inter-domain rerouting by betweenness centrality after disruptive events.

[154] Y. Liu, Z. Luo, R. K. C. Chang, and J. Su. Characterizing inter-domain rerouting after japan earthquake. *NETWORKING 2012, Part II, LNCS 7290, pp. 124–135, 2012.*, 2012.

[155] G. Lotan, E. Graeff, M. Ananny, D. Gaffney, I. Pearce, and D. Boyd. The revolutions were tweeted: Informa-

tion flows during the 2011 tunisian and egyptian revolutions. *International Journal of Communication 5 (2011), Feature 1375–1405*, 2011.

[156] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iplane: An information plane for distributed services. *OSDI '06: 7th USENIX Symposium on Operating Systems Design and Implementation*, 2006.

[157] D. Madory. `http://www.renesys.com/2013/12/protests-lead-outage-thailand`, 2013.

[158] R. Mahajan, D. Wetherall, and T. Anderson. Understanding bgp misconfiguration. *SIGCOMM'02*, 2002.

[159] S. Mansfield-Devine. {DDoS} goes mainstream: how headline-grabbing attacks could make this threat an organisation's biggest nightmare. *Network Security*, 2016(11):7 – 13, 2016. ISSN 1353-4858. doi: http://dx.doi.org/10.1016/S1353-4858(16)30104-0. URL `http://www.sciencedirect.com/science/article/pii/S1353485816301040`.

[160] X. H. Z. M. Mao. Accurate real-time identification of ip prefix hijacking. *2007 IEEE Symposium on Security and Privacy (SP'07)*, 2007.

[161] P. Marchetta, W. de Donato, and A. Pescapè. Detecting third-party addresses in traceroute traces with ip timestamp option. *Passive and Active Measurements*, 2013.

[162] P. Marchetta, V. Persico, E. Katz-Bassett, and A. Pescapè. Don't trust traceroute (completely). In *ACM CoNEXT Student workshop,*, 2013.

[163] P. Marchetta, V. Persico, and A. Pescapè. Pythia: yet another active probing technique for alias resolution. *CoNEXT, pp. 229-234.*, 2013.

[164] P. Marchetta, A. Botta, E. Katz-Bassett, and A. Pescapè. Dissecting round trip time on the slow path with a single packet. *Passive and Active Measurement Conference (PAM)*, 2014.

[165] B. Marczak, N. Weaver, J. Dalek, R. Ensafi, D. Fifield, S. McKune, A. Rey, J. Scott-Railton, R. Deibert, and V. Paxson. An analysis of china's "great cannon". In *5th USENIX Workshop on Free and Open Communications on the Internet (FOCI 15)*, Washington, D.C., 2015. USENIX Association. URL `http://blogs.usenix.org/conference/foci15/workshop-program/presentation/marczak`.

[166] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, Y. Ganjali, and C. Diot. Characterization of failures in an operational ip backbone network. *IEEE/ACM transactions on networking, vol. 16, no. 4*, 2008.

[167] M. M. R. Martin. Network resilience through multi-topology routing. *University of Wurzburg, Institute of Computer Science, Tech. Rep. 335*, 2004.

[168] M.Gunes and K. Sarac. Resolving anonymous routers in internet topology measurement studies", in *IEEE INFOCOM*, 2008. *IEEE INFOCOM*, 2008.

[169] S. N. E. Modiano. Network reliability with geographically correlated failures. In *IEEE INFOCOM 2010 proceedings*, 2010.

[170] S. N. E. Modiano. Network reliability under random circular cuts. *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE (pp. 1-6). IEEE*, 2011.

[171] J. Morin, J. Aubert, and B. Gateau. Towards cloud computing sla risk management: Issues and challenges. *2012 45th Hawaii International Conference on System Sciences. IEEE*, 2012.

[172] G. Muller. Fuzzy architecture assessment for critical infrastructure resilience. *Procedia Computer Science 12 (2012) 367–372*, 2012.

[173] S. Murphy. Bgp security vulnerabilities analysis. `https://tools.ietf.org/html/rfc4272`, 2006.

[174] M. Nazari Cheraghlou, A. Khadem-Zadeh, and M. Haghparast. A survey of fault tolerance architecture in cloud computing. *Journal of Network and Computer Applications*, 61:81–92, Feb 2016. ISSN 1084-8045. doi: 10.1016/j.jnca.2015.10.004. URL `http://dx.doi.org/10.1016/j.jnca.2015.10.004`.

[175] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano. Assessing the vulnerability of the fiber infrastructure to disasters. *IEEE/ACM Transactions on Networking, vol. 19, no. 6*, 2011.

[176] E. Nygren, R. K. Sitamaran, and J. Sun. The akamai network: A platform for high-performance internet applications. *SIGOPS OSR, 44:2–19*, 2010.

[177] Y. Ogata. A prospect of earthquake prediction research. *Statistical Science*, pages 521–541, 2013.

[178] A. Ogielski and J. Cowie. Internet routing behavior on 9/11 and in the following weeks. *Renesys Corporation*, pages 5–6, 2002.

[179] R. Oliveira, M. Lad, B. Zhang, and L. Zhang. Geographically informed inter-domain routing. *Network Protocols, 2007. ICNP 2007. IEEE International Conference on. IEEE*, 2007.

[180] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)completeness of the observed internet as-level structure. *IEEE/ACM Transactions on Networking, vol. 18, no. 1*, 2010.

[181] C. on the Internet Under Crisis Conditions: Learning from the Impact of September 11. The Internet Under Crisis Conditions, T. N. A. Press, and 2003.

[182] T. W. P. C. Oorschot. Analysis of bgp prefix origins during google's may 2005 outage. *Parallel and Distributed Processing Symposium*, 2006.

[183] K. S. L. Palen. Pass it on?: Retweeting in mass emergency. In *Proceedings of the 7th International ISCRAM Conference*, 2010.

[184] L. Palen, K. Starbird, S. Vieweg, and A. Hughes. Twitter-based information distribution during the 2009 red river valley flood threat. *Bulletin of the American Society for Information Science and Technology, vol. 36, no. 5*, 2010.

[185] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of internet background radiation. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement (pp. 27-40). ACM,*, 2004.

[186] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang. Impact of configuration errors on dns robustness. *SIGCOMM'04*, 2004.

[187] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang. Impact of configuration errors on dns robustness. *IEEE Journal on Selected Areas in Communications, vol. 27, no. 3*, 2009.

[188] D. Pei, X. Zhao, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Improving bgp convergence through consistency assertions. *IEEE INFOCOM*, 2002.

[189] D. Pei, D. Massey, and L. Zhang. A framework for resilient internet routing protocols. *UCLA, Tech. rep. CSD-TR-030052*, 2003.

[190] A. Peresan, V. G. Kossobokov, and G. F. Panza. Operational earthquake forecast/prediction. *Rendiconti Lincei*, 23(2):131–138, Jun 2012. ISSN 1720-0776. doi: 10.1007/s12210-012-0171-7. URL `https://doi.org/10.1007/s12210-012-0171-7`.

[191] M. Perlin. Downtime, outages and failures - understanding their true costs. `http://www.evolven.com/blog/downtime-outages-and-failures-understanding-their-true-costs.html`, 2012.

[192] P. M. A. Pescapè. Drago: Detecting, quantifying and locating hidden routers in traceroute ip paths. *IEEE Global Internet Symposium*, 2013.

[193] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. Ip geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review 41.2 (2011): 53-56*, 2011.

[194] J. Postel. RFC792. `http://tools.ietf.org/html/rfc792`, 1981.

[195] T. B. project. http://bgpmon.netsec.colostate.edu.

[196] Y. Qu, P. F. Wu, and X. Wang. Online community response to major disaster: A study of tianya forum in the 2008 sichuan earthquake. In *Proceedings of the 42nd Hawaii International Conference on System Sciences*, 2009.

[197] Y. Qu, C. Huang, P. Zhang, and J. Zhang. Microblogging after a major disaster in china: A case study of the 2010 yushu earthquake. *CSCW 2011*, 2011.

[198] L. Quan and J. H. Y. Pradkin. Detecting internet outages with precise active probing (extended). *USC Technical Report*, 2012.

[199] L. Quan, J. Heidemann, and Y. Pradkin. Visualizing sparse internet events: Network outages and route changes. In *Proc. of First ACM Workshop on Internet Visualization*, 2012.

[200] L. Quan, J. Heidemann, and Y. Pradkin. Trinocular: Understanding internet reliability through adaptive probing. *SIGCOMM'13*, 2013.

[201] D. H. J. E. Ramirez-Marquez. Generic metrics and quantitative approaches for system resilience as a function of time. *Reliability Engineering and System Safety 99 (2012) 114–122*, 2011.

[202] J. C. W. J. Ramirez-Marquez. Resiliency as a component importance measure in network reliability. *Reliability Engineering and System Safety 94 (2009) 1685-1693*, 2009.

[203] D. Reina, M. Askalani, S. Toral, F. Barrero, E. Asimakopoulou, and N. Bessis. A survey on multihop ad hoc networks for disaster response scenarios. *International Journal of Distributed Sensor Networks*, 11 (10):647037, 2015.

[204] Y. Rekhter, T. Li, and S. Hares. RFC4271. `http://tools.ietf.org/html/rfc4271`, 2006.

[205] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine 21 (6) (2001) 11–25, doi:10.1109/37.969131. ISSN 0272-1708*, 2001.

[206] M. Roughan, W. Willinger, O. Maennell, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the internet's autonomous systems. *IEEE Journal on Selected Areas in Communications, vol. 29, no. 9*, 2011.

[207] S. A. J. K. Routray. Community resilience framework for an earthquake prone area in baluchistan. *International Journal of Disaster Risk Reduction 2 (2012) 25-36*, 2012.

[208] A. Sahoo, K. Kant, and P. Mohapatra. Characterization of bgp recovery time under large-scale failures. In *IEEE ICC proceedings.*, 2006.

[209] A. Sahoo, K. Kant, and P. Mohapatra. Improving bgp convergence delay for large-scale failures. In *DSN'06 Proceedings*, 2006.

[210] A. Sahoo, K. Kant, and P. Mohapatra. Bgp convergence delay after multiple simultaneous router failures: Characterization and solutions. *Computer Communications 32 (2009) 1207–1218*, 2009.

[211] T. Sakaki, F. Toriumi, K. Uchiyama, Y. Matsuo, K. Shinoda, K. Kazama, S. Kurihara, and I. Noda. The possibility of social media analysis for disaster management. *Humanitarian Technology Conference (R10-HTC), 2013 IEEE Region 10. IEEE*, 2013.

[212] D. Schatzmann, S. Leinen, J. Kögel, and W. Mühlbauer. Fact: Flow-based approach for connectivity tracking. *In Passive and Active Measurement (pp. 214-223). Springer Berlin Heidelberg*, 2011.

[213] J. Segovia, P. Vilà, E. Calle, and J. L. Marzo. Improving the resilience of transport networks to large-scale failures. *JOURNAL OF NETWORKS, VOL. 7, NO. 1*, 2012.

[214] A. Sen, S. Murthy, and S. Banerjee. Region-based connectivity - a new paradigm for design of fault-tolerant networks. *IEEE International Conference on High Performance Switching and Routing (HPSR), pp. 1-7*, 2009.

[215] Y. S. E. Shir. Dimes: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review, vol. 35, no. 5*, 2005.

[216] R. Shirey. RFC4949. `http://tools.ietf.org/html/rfc4949`, 2007.

[217] D. P. Siewiorek and R. S. Swarz. Reliable computer systems: design and evaluation. *AK Peters, Ltd.*, 6, 1998.

[218] P. F. Silva, C. B. Westphall, C. M. Westphall, M. M. Mattos, and D. R. dos Santos. An architecture for risk analysis in cloud. *ICNS 2014, The Tenth International Conference on Networking and Services*, 2014.

[219] W. F. Slater. The internet outage and attacks of october 2002. *Chicago Chapter of the Internet Society*, 2002.

[220] A. P. Snow, U. Varshney, and A. D. Malloy. Reliability and survivability of wireless and mobile networks. *Computer*, 33(7):49–55, 2000.

[221] J. Sobrinho, L. Vanbever, F. Le, and J. Rexford. Distributed route aggregation on the global network. *ACM CoNEXT*, 2014.

[222] E. S. Spiro, J. Sutton, M. Greczek, S. Fitzhugh, N. Pierski, and C. T. Butts. Rumoring during extreme events: A case study of deepwater horizon 2010. *WebSci 2012*, 2012.

[223] A. S. N. Spring. Pingin' in the rain. *IMC'11*, 2011.

[224] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 2010.

[225] H. Styron. Csx tunnel fire: Baltimore. *US Fire Administration Technical Report USFA-TR-140, Federal Emergency Management Administration, Em-*

mitsburg, MD, 2001.

[226] G. S. G. Swallow. Sonet/sdh - like resilience for ip networks: A survey of traffic protection mechanisms. *IEEE Network*, 2004.

[227] A. Sydney, C. Scoglio, P. Schumm, and R. E. Kooij. Elasticity: Topological characterization of robustness in complex networks. *Bionetics'08*, 2008.

[228] S. T. Teoh, S. Ranjan, A. Nucci, and C. Chuah. Bgp eye: A new visualization tool for real-time detection and analysis of bgp anomalies. In *Proceedings of the 3rd international workshop on Visualization for computer security (pp. 81-90). ACM*, 2006.

[229] T. Terpstra, A. de Vries, R. Stronkman, and G. L. Paradies. Towards a realtime twitter analysis during crises for operational crisis management. In *Proceedings of the 9th International ISCRAM Conference*, 2012.

[230] B. Tierney, J. Metzger, J. Boote, E. Boyd, A. Brown, R. Carlson, M. Zekauskas, J. Zurawski, M. Swany, and M. Grigoriev. Perfsonar: Instantiating a global network measurement framework. In *SOSP Workshop. Real Overlays and Distributed System,*, 2009.

[231] F. Toriumi, T. Sakaki, K. Shinoda, K. Kazama, S. Kurihara, and I. Noda. Information sharing on twitter during the 2011 catastrophic earthquake. In *Proceedings of the 22nd international conference on World Wide Web companion (pp. 1025-1028). International World Wide Web Conferences Steering Committee*, 2013.

[232] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage. California fault lines: Understanding the causes and impact of network failures. *SIGCOMM'10*, 2010.

[233] K. Vajanapoom, D. Tipper, and S. Akavipat. A risk management approach to resilient network design. In *2010 International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, 2010.

[234] K. Vajanapoom, D. Tipper, and S. Akavipat. Risk based resilient network design. *Telecommunication Systems, 52(2), 799-811*, 2011.

[235] K. K. N. Venkatasubramanian, A. the impact of geographically correlated failures on overlay-based data dissemination, . in *IEEE GLOBECOM 2010 and pp-1-5*, and D. 2010. *IEEE GLOBECOM 2010, pp- 1-5*, 2010.

[236] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen. Microblogging during two natural hazards events: What twitter may contribute to situational awareness. *CHI 2010*, 2010.

[237] VUZE and I. V. http://www.vuze.com.

[238] P. S. B. Walters. Quirc: A quantitative impact and risk assessment framework for cloud security. *2010 IEEE 3rd International Conference on Cloud Computing*, 2010.

[239] J. Wan, M. Yi, D. Li, C. Zhang, S. Wang, and K. Zhou. Mobile services for customization manufacturing systems: an example of industry 4.0. *IEEE Access*, 4: 8977–8986, 2016.

[240] J. Wang, C. Qiao, and H. Yu. On progressive network recovery after a major disruption. *IEEE INFOCOM*, 2011.

[241] J. Wang, C. Jiang, and J. Qian. Robustness of internet under targeted attack: a cascading failure perspective. *Journal of Network and Computer Applications*, 40: 97–104, 2014.

[242] P. Wang, W. Lin, P. Kuo, H. Lin, and T. C. Wang. Threat risk analysis for cloud security based on attack-defense trees. *2012 8th International Conference on Computing Technology and Information Management (NCM and ICNIT), vol. 1, pp. 106-111*, 2012.

[243] X. Wang. Network recovery and augmentation under geographically correlated region failures. *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE (pp. 1-5). IEEE*, 2011.

[244] X. Wang, X. Jiang, and A. Pattavina. Assessing network vulnerability under probabilistic region failure model. *2011 IEEE 12th International Conference on High Performance Switching and Routing (HSPR)*, 2011.

[245] Y. Wang, H. Wang, A. Mahimkar, R. Alimi, Y. Zhang, L. Qiu, and Y. R. Yang. R3: Resilient routing reconfiguration. *SIGCOMM'10*, 2010.

[246] W. Willinger and M. Roughan. Internet topology research redux. *ACM SIGCOMM eBook: Recent Advances in Networking*, 2013.

[247] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin. Internet routing resilience to failures: Analysis and implications. *CoNEXT'07*, 2007.

[248] . Y. Xiang, Z. Wang, X. Yin, and J. Wu. Argus: An accurate and agile system to detecting ip prefix hijacking. In *Workshop on Trust and Security in the Future Internet*, 2012.

[249] K. Xu, J. Chandrashekar, and Z. Zhang. A first step toward understanding inter-domain routing dynamics. In *SIGCOMM'05 Workshops*, 2005.

[250] Y, Zhu, A. Bavier, N. Feamster, S. Rangarajan, and J. Rexford. Ufo: A resilient layered routing architecture. *ACM SIGCOMM Computer Communication Review, Vol. 38, No. 5*, 2008.

[251] H. Yang and S. S. Lam. Collaborative verification of forward and reverse reachability in the internet data plane. *IEEE 22nd International Conference on Network Protocols (ICNP), 2014.*, 2014.

[252] Y.Huang, N.Feamster, and R.Teixeira. Practical issues with using network tomography for fault diagnosis,", in *Computer Communication Review, 38(5):53–57*, october 2008. *Computer Communication Review, 38(5):53–57*, 2008.

[253] J. Zhang, J. Rexford, and J. Feigenbaum. Learning-based anomaly detection in bgp updates. In *SIGCOMM'05 Workshops*, 2005.

[254] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. An analysis of bgp multiple origin as (moas) conflicts. *IMW'01*, 2001.

[255] Y. S. N. Zilberman. A study of geolocation databases. *arXiv preprint arXiv:1005.5674*, 2010.

[256] E. Zmijewski. Reckless driving on the Internet. `http://www.renesys.com/blog/2009/02/the-flap-heard-around-the-worl.shtml`, 2009.

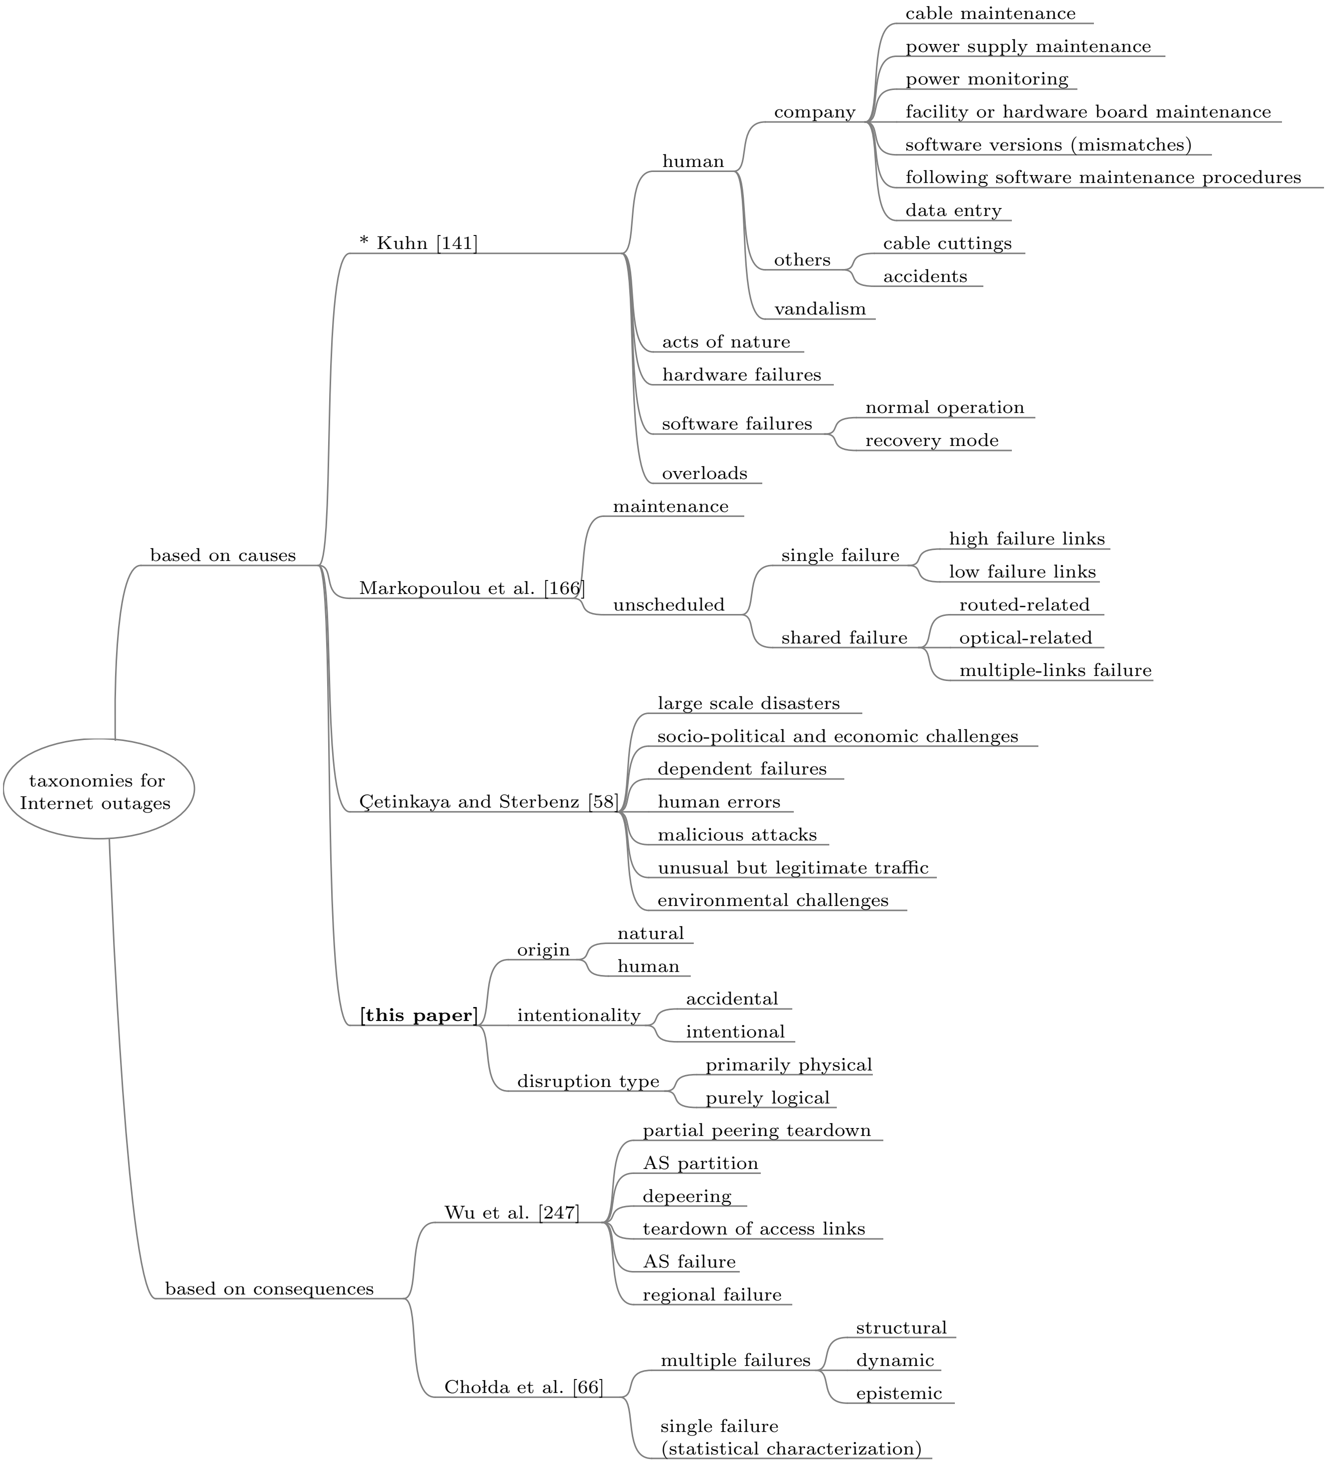

Figure 3: Classification of Internet outages: taxonomies from literature and this paper.
* Considers Public Switched Telephone Network, not the Internet.

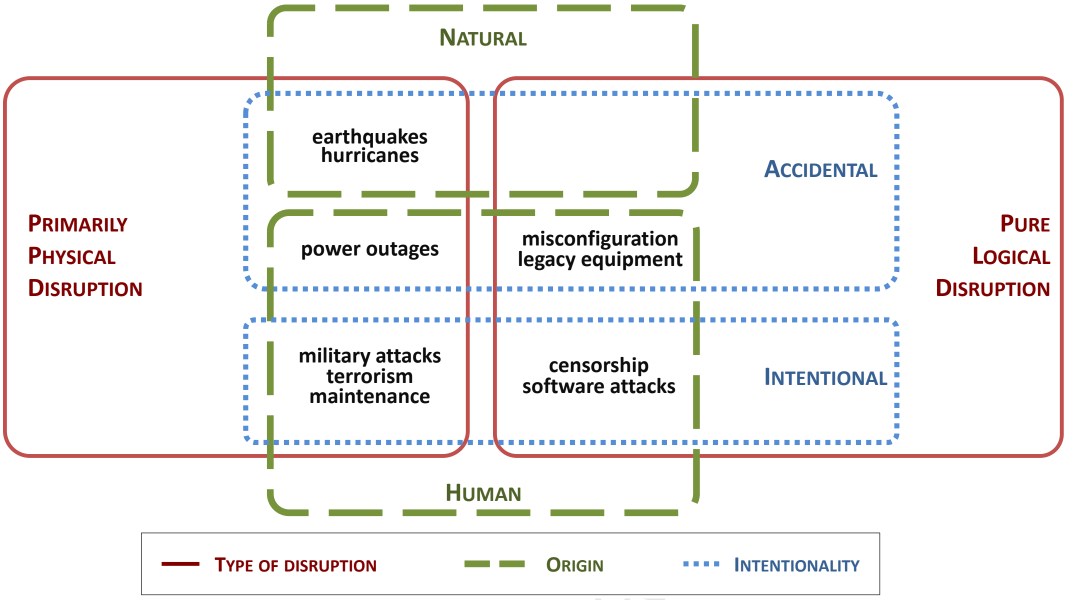

Figure 4: A characterization of the causes of outages based on the origin, the intentionality, and the type of disruption.

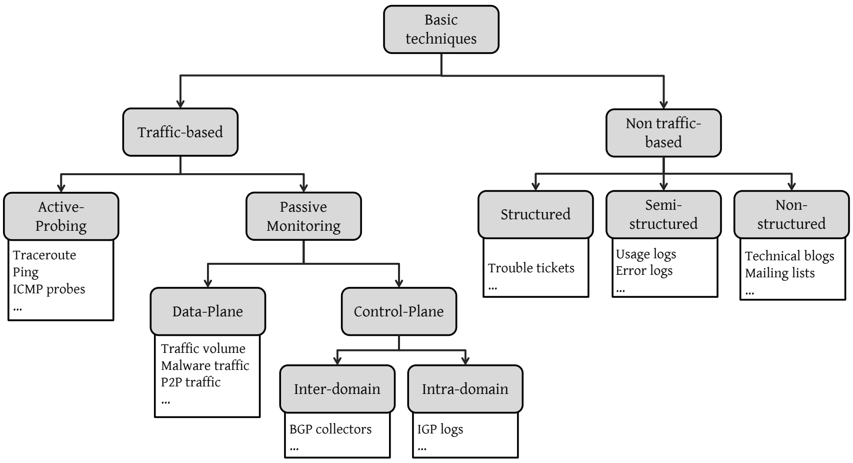

Figure 5: A classification of the basic techniques adopted in outage-related works.

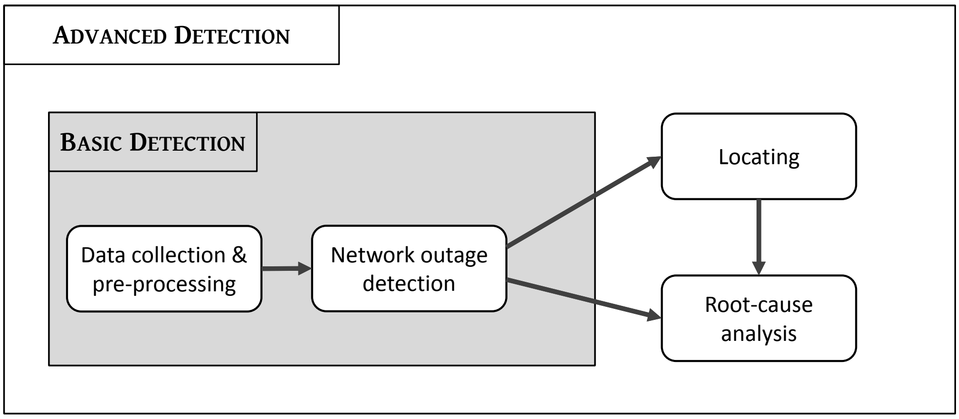

Figure 6: Flowchart of a generic outage detection tool.

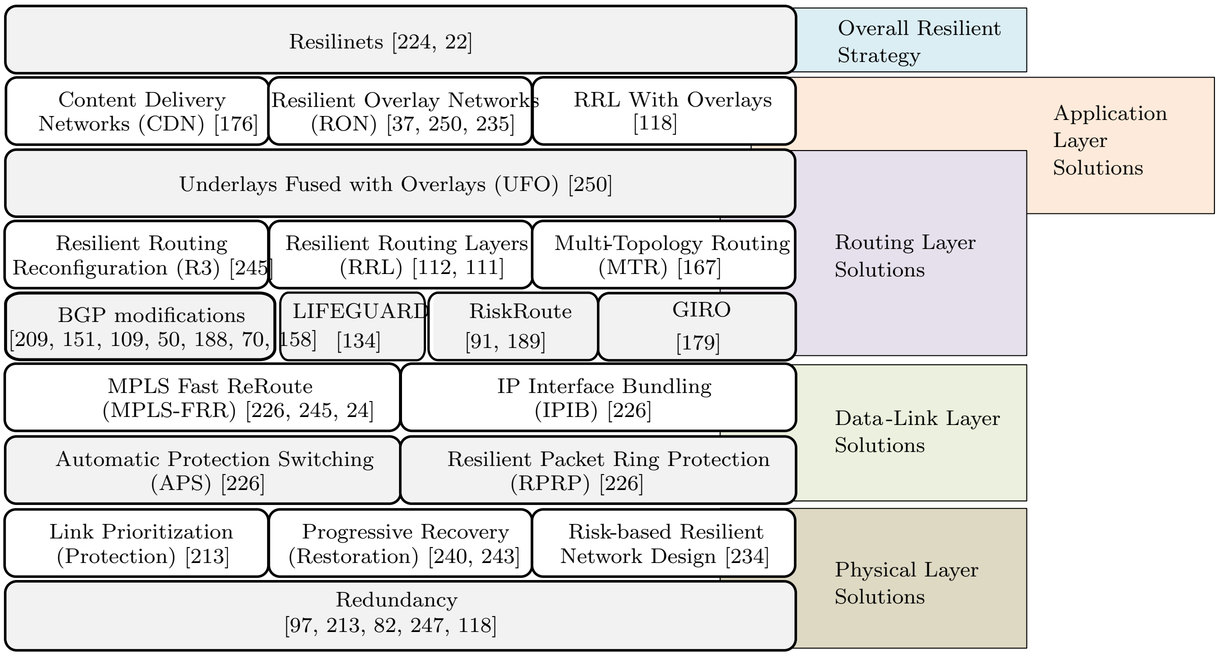
Resilinets [224, 22]
Overall Resilient Strategy
Content Delivery Networks (CDN) [176]
Resilient Overlay Networks (RON) [37, 250, 235]
RRL With Overlays [118]
Application Layer Solutions
Underlays Fused with Overlays (UFO) [250]
Resilient Routing Reconfiguration (R3) [245]
Resilient Routing Layers (RRL) [112, 111]
Multi-Topology Routing (MTR) [167]
Routing Layer Solutions
BGP modifications [209, 151, 109, 50, 188, 70, 158]
LIFEGUARD [134]
RiskRoute [91, 189]
GIRO [179]
MPLS Fast ReRoute (MPLS-FRR) [226, 245, 24]
IP Interface Bundling (IPIB) [226]
Data-Link Layer Solutions
Automatic Protection Switching (APS) [226]
Resilient Packet Ring Protection (RPRP) [226]
Link Prioritization (Protection) [213]
Progressive Recovery (Restoration) [240, 243]
Risk-based Resilient Network Design [234]
Physical Layer Solutions
Redundancy [97, 213, 82, 247, 118]

Figure 8: A classification of the specific solutions to prevent or mitigate the consequences of network outages.

Giuseppe Aceto (giuseppe.aceto@unina.it) is a Post Doc at the Department of Electrical Engineering and Information Technology of University of Napoli Federico II. Giuseppe has a PhD in telecommunication engineering from the University of Napoli Federico II. His work falls in measurement and monitoring of network performance and security, with focus on censorship. He has served and serves as reviewer for several journals and conferences (e.g. IEEE Transactions on Cloud Computing, Elsevier's JNCA, Computer Communications, Globecom, ICC). Giuseppe Aceto is co-author of papers in international publishing venues (IEEE Transactions on Network and Service Management, Elsevier Journal of Network and Computer Applications, Computer Networks, INFOCOM, ACM SAC, etc.) and is co-author of a patent. Giuseppe is the recipient of a best paper award at IEEE ISCC 2010.

Alessio Botta received the M.S. degree in telecommunications engineering and the Ph.D. degree in computer engineering and systems from the University of Naples Federico II, Naples, Italy. He currently holds a post-doctoral position with the Department of Computer Engineering and Systems, University of Naples Federico II. He has co-authored over 50 international journal and conference publications. His current research interests include networking, and, in particular, network performance measurement and improvement, with a focus on wireless and heterogeneous systems. Dr. Botta has served and serves as an independent reviewer of research and implementation project proposals for the Romanian government. He was a recipient of the Best Local Paper Award at the IEEE ISCC 2010. In the research area of networking, he has chaired international conferences and workshops, served and serves several technical program committees of international conferences (IEEE Globecom and IEEE ICC), and acted as a reviewer for different international conferences (the IEEE Conference on Computer Communications) and journals (the IEEE Transactions on Mobile Computing, the IEEE Network Magazine, and the IEEE Transactions on Vehicular Technology).

Pietro Marchetta received his Master degree and PhD degree in Computer Engineering at University of Napoli in 2014. Currently, he holds a postdoctoral position at the Department of Electrical Engineering and Information Technology of the University of Napoli Federico II (Italy). His main research activities focus on methodologies, techniques and large-scale distributed platforms for Internet measurements with a specific focus on Internet topology, routing, and performance. He served and serves a reviewer for a dozen of conferences and journals (e.g. Elsevier's Computer Networks and Future Generation Computer Systems). For his research, he received some awards including the first place at the ACM Student Research Competition at SIGCOMM 2012 and the Best Student Workshop Paper Award in CoNEXT 2013. Pietro Marchetta has also been IT lead for the Smart City project "S2move – Smart and Social Move" financed by MIUR.

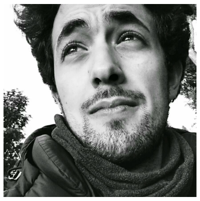

Valerio Persico is a Post Doc at the Department of Electrical Engineering and Information Technology of University of Napoli Federico II. He has a PhD in computer engineering from the University of Napoli Federico II. His work focuses on measurement and monitoring of cloud network infrastructures. Recently, he is working on bioinformatic. Valerio is the recipient of the best student paper award at ACM CoNext 2013.

Antonio Pescapé [SM '09] is a Full Professor atthe Department of Electrical Engineering and Information Technology of the University of Napoli Federico II (Italy). His research interests are in the networking field with focus on Internet Monitoring, Measurements and Management and on Network Security. Antonio Pescapé has coauthored over 180 journal and conference publications and he is co-author of a patent. For his research activities he has received several awards, comprising a Google Faculty Award, several best paper awards and two IRTF (Internet Research Task Force) ANRP (Applied Networking Research Prize).