

Detecting a Crisis: Comparison of Self-Reported vs. Automated Internet Outage Measuring Methods

Denis Orlov
Simon Möller
Sven Düfer
Steffen Haesler
Christian Reuter

denis.orlov@stud.tu-darmstadt.de
simon.moeller@stud.tu-darmstadt.de
sven.duefer@stud.tu-darmstadt.de
haesler@peasec.tu-darmstadt.de
reuter@peasec.tu-darmstadt.de

Science and Technology for Peace and Security (PEASEC), Technical University of Darmstadt
Germany

ABSTRACT

Every day, there are internet disruptions or outages around the world that affect our daily lives. In this paper, we analyzed these events in Germany in recent years and found out how they can be detected, and what impact they have on citizens, especially in crisis situations. For this purpose, we take a look at two different approaches to recording internet outages, namely the self-reporting of citizens and automatic reporting by algorithmic examination of the availability of IP networks. We evaluate the data of six major events with regard to their meaningfulness in quality and quantity. We found that due to the amount of data and the inherent imprecision of the methods used, it is difficult to detect outages through algorithmic examination. But once an event is publicly known by self-reporting, they have advantages to capture the temporal and spatial dimensions of the outage due to its nature of objective measurements. As a result, we propose that users' crowdsourcing can enhance the detection of outages and should be seen as an important starting point to even begin an analysis with algorithm-based techniques, but it is to ISPs and regulatory authorities to support that.

KEYWORDS

Internet Outages, Germany, Self-Organization, Crowdsourcing, Crisis Informatics

1 INTRODUCTION

In the history of the global internet, there have been frequent disruptions or failures of the network infrastructure [1], including

in Germany. These are due to man-made, technical, or force majeure causes. According to an International Telecommunication Union (ITU) report, 4.9 billion people, or 63% of the world population, use the internet, with a higher rate of 76% in urban areas [24]. While internet access makes information and skills accessible, and let people develop routines and everyday behavior on a functioning internet, it often breaks down during crises and disasters, such as the July 2021 floods in Western Europe, where it is so urgently needed and becomes critical. In a previous study, it has already been shown that in the event of a power outage where also internet communication is lost, recovery is more difficult because the communication of the emergency response teams to fix the problem also takes place over the internet [46]. There is an increasing demand for such crisis related applications from the citizens' point of view [26]. However, they usually require an existing internet connection in order to retrieve information from servers. Consequently, such applications may not be usable in crisis situations and therefore cannot fulfill their purpose. Even in the absence of another causal crisis, an internet outage can be perceived as a crisis in itself and have serious consequences where people have strategies to cope [29].

There are several ways to address the problem, either by strengthening the resilience of the Information and Communications Technology (ICT) infrastructure to establish higher availability of internet access, even in times of crisis, or by focusing on decentralized, local first approaches that can cope with disrupted or fluctuating internet access [18]. From a HCI perspective, developers and designers are encouraged to integrate the offline state to consider which features of the app should still work, rather than disabling the app as a whole. Regardless which way to go, it is important to better understand the vulnerabilities and know the statistics, mechanisms and detection of internet outages. An often underestimated factor in this field is the perspective of citizens, who are either users or detectors, and who have an interest in functioning internet access in the location where they live and not just where outages in data centers or carrier backbones can be measured. Therefore, in this study, we compare self-reported internet outages against automated methods and additionally look at time periods of six

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Veröffentlicht durch die Gesellschaft für Informatik e.V.

in K. Marky, U. Grünefeld & T. Kosch (Hrsg.):

Mensch und Computer 2022 – Workshopband, 04.-07. September 2022, Darmstadt

© 2022 Copyright held by the owner/author(s).

<https://doi.org/10.18420/muc2022-mci-ws10-321>

known events in both methods. We follow two different approaches (self-reporting and three algorithm-based methods) to detect internet outages and, to our knowledge, compare both directly for the first time. The question arises as to the differences between these methods regarding accuracy in the sense of *qualitative* and *quantitative* measurements and the possible *combination* of each. In this work *qualitative* means in terms of internet outages, how accurate each outage is described by reported or measured data, whereas *quantitative* is the amount of data that can be analyzed and if in consequence how many outages can be detected.

- **RQ1:** What are the *qualitative* differences between self-reporting and automated measurements in terms of accuracy and reliability of the data provided by each method?
- **RQ2:** What are the *quantitative* differences between the methods regarding the amount of data available for each?
- **RQ3:** Which methods are best suited for capturing disturbances in terms of *quality* and *quantity*, and how can they be *combined* for a better result, if necessary?

As a contribution this work shows that both approaches have weaknesses that can be minimized by a combination. Where humans are better in outages that are not known initially or dependent to crisis events, automated tools allow a more accurate analysis of the dimensions and duration of the outage.

In the next section, we will highlight related work regarding outages and citizens' perspectives and give a background of the automated reporting methods in section 2 followed by section 3 explaining the methods of this work. Then we show the results for self-reporting, and automated reporting and compare them on six distinct events in section 4, followed by the discussion and conclusion in section 5, closing with limitations and outlook in section 6.

2 RELATED WORK AND BACKGROUND

In the following, we present work on the causes and consequences of internet outages, as well as on users' perspectives on them. We highlight the role of connectivity in disaster situations, which is often the cause of outages in a cascading effect, and finally give a theoretical and empirical background on how outages can be detected by both citizens and algorithms and what weaknesses or strengths they have.

If the access to the internet is disrupted, the reasons can be manifold and related to disasters such as earthquakes [9], floods [10], volcanic eruptions [47], hurricanes [34], solar storms [25], lightning strikes [28], or forest fires [11], and purposeful or unintentional actions such as network attacks [31], terrorist attacks [32], undersea cable cuts [7], censorship [12], governmental actions against misinformation [44] or misconfiguration [30]. For this reason, Aceto et al. [1] proposed three characteristics "origin" (natural vs. human), "intentionality" (accidental vs. intentional) and "disruption type" (primarily physical vs. pure logical) to classify the cause of an outage. Even though the characteristic "natural vs. human" is controversial [8], we will reference to the work and classification of Aceto et al. because it's contribution is so central to the topic of crisis related outages.

While the potential causes of an outage are manifold, the consequences for users are almost the same: Apps are not working,

people are cut from information and cannot use the tools they know and use in everyday life. While short time being offline could also be refreshing, it is a problem in the long run [17, 29]. As internet outages are often linked to other disruptions that trigger a cascading effect, people do not only have to deal with the information and communication cutoff, but also tackle the crisis in this situation. Power outages, for example, lead to an increased need for information [21, 39] and to massive psychological stress [41]. Even though there are approaches to ensure that even collaborative apps still work in the event of an outage [18, 35], collaboration usually stops [2].

2.1 Measuring Internet Outages

Because of the potential consequences, detecting outages is necessary, and the picture is not always so clear that there is another crisis as a triggering event. While there are several technical possibilities to detect outages, which mostly rely on measuring network devices or signals in central network infrastructure [15], many causes are hard to detect but are often related to the actions of people who, for example, carry out maintenance work [40]. People can also play a substantial role, e.g. by analyzing social network activities to detect power and communication outages [36]. It is important to answer the question of how best to measure outages while taking into account user connectivity by determining what is relevant from the perspective of citizens and not just through the lens of an Internet Service Provider (ISP) or the availability of a distinct service in a data center. And even if ISPs are able to measure these outages or are forced to report outages, the problem remains that this data is not publicly available. Another aspect concerns the fact that people can also report outages in residential areas or on a very small level, where algorithm-based measurements could be blind or generate a lot of traffic just for the measurement, but would lack in accuracy, reliability and the willingness to report.

With the method of *Self-Reporting* citizens autonomously report an internet outage by providing information about the outage such as duration, ISP, type, and location, e.g. by using a mailing list as Banerje et al. did [4]. They say that having users who report (instead of algorithms) opens the chance to gather semantic context and user issues. At this point, two of this method's disadvantages can already be determined. Firstly, an internet connection is required to report an outage, and secondly, reports can be error-prone due to incorrect or subjective entries.

With *Automated Reporting*, there are several technologies to detect outages on the basis of an algorithm which use active or passive technologies by sending requests to devices, listening to network traffic [43], or using big data analysis [5]. Of all technologies, we further provide details about most common technologies for monitoring (1) *Active Probing*, (2) *Border Gateway Protocol*, and (3) *Network Telescope*, as they are consequently used in our analysis.

(1) *Active Probing* (AP) is based on a methodology developed by the University of Southern California called Trinocular [38]. Here, monitoring agents ping routable /24 IPv4 address ranges in 10 minutes intervals and check their availability. Outages with a duration at least equal to the interval length should be reliably detected [19]. A network has to consist of at least 15 pingable devices to be monitored. The failure of multiple devices is necessary

for registering a network outage. By using Bayesian statistics, the responses are interpreted and the number of messages required for interpretation purposes is minimized [19]. Yet, the Trinocular method is limited to only IPv4 reachable network devices and it may falsely register IPv4 networks as unavailable if dynamic allocation of IPv4 addresses leads to address changes on monitored devices. Finally, Domain Name System (DNS) errors are not detected, since IP addresses are pinged directly and not through DNS entries.

(2) The Border Gateway Protocol (BGP) is a dynamic routing protocol that manages internet traffic worldwide and it works for both IPv4 and IPv6. It is possible to audit the reachability of IP networks through BGP and thus its availability. Routing information is publicly available and outages can be easily detected and traced. Yet, this can only be used to detect outages that affect the routing via BGP. Outages that go beyond routing are not detected using this method [19].

(3) With Network Telescope (NT) the so-called background noise of internet traffic is scanned which mostly consists of communication from malware infected computer systems such as botnets. NT scans how many unique IP addresses are involved in this communication. In the event of an internet disruption, the number of IP addresses decreases as a consequence of inaccessibility. NT can complement AP, as it can also capture devices that are located behind a firewall. In reality though, only a small portion of the total background noise is actually captured. Distinguishing between traffic in the background noise and traffic that has simply been modified using IP spoofing can also be difficult. Accordingly, short term outages or those small in scope can be hard to detect [19]. An observation by the University of Southern California found that the number of active IP addresses involved in this background noise was roughly constant over the course of a day in the U.S. and Europe, but fluctuated significantly in other countries such as Russia, India, and China [19]. As a result, monitoring with AP and NT in the latter countries has become more difficult, since comparative values must be determined by these methods in order to reliably detect a disturbance.

3 METHODS

To answer the research questions mentioned in section 1, we will consider heise iMonitor [20], a German internet outage self-reporting portal, and "Internet Outage Detection and Analysis" (IODA) [14], as a data source for automated reporting. Heise iMonitor offers access to a huge dataset of internet outages since 2001 and the possibility to crawl the raw data which other platforms like downdetector [33] do not offer. IODA was chosen over "Analysis of Network Traffic" (ANT) because ANT only offers AP for the detection of outages, while IODA offers AP, BGP, and the NT. Moreover, the visualization of the data at IODA offered a greater level of detail. Our goal is to provide an overview of internet disruptions and outages in Germany over the last 20 years for Self-Reporting and over the last 5 years for IODA. The analysis will be in terms of the number of affected households, duration, frequency of occurrence, but also by geographic coverage. From this data, we will filter out six known disruptions and create an overview to examine how reliable internet availability was during the time period mentioned.

3.1 Self-Reporting with iMonitor

First, the data analysis deals with the data provided in iMonitor with datasets optimized for web viewing, which we crawled into CSV files using a python script. Then, all the resulting CSV files were merged into one all-encompassing file, which contains the following columns: *Area Code*, *City*, *District*, *ISP*, *Type of Access*, *Start of outage*, *End of outage*, *Duration*, *Type of Fault*, *Id*, and *Comment*. The data was collected at 2021-11-16 4:56 p.m. and includes 80,317 entries of outages in Germany, back to 2001-08-17 6:00 p.m.. We then used Visual Basic macros in Microsoft Excel to format the data, correct failures, and sum up outages per day and location.

The duration was double checked by calculating the value *Time Diff* on the basis of the values *End of outage* and *Start of outage*. These values were thoroughly checked as the website allows values for the field *End of outage* to be inserted with a date before the field *Start of outage*, which can give incorrect results. In addition, there are records without values in the field *End of outage*, which means that the *Time Diff* cannot be calculated. These incorrect reports are grouped together and have been excluded, which were about 14 percent of the dataset. To determine the number of reports for a reported outage duration, we count the number of occurrences by the different time intervals in the *Time Diff* field. We determined the interval as duration of disturbance up to 1 hour: $0 > \text{time diff} \geq 60$, duration of disturbance up to 2 hours: $60 > \text{time diff} \geq 120$ and so on.

3.2 Automated Reporting with IODA

The second data source used for analysis of automated reporting is the IODA platform. It uses data from about 500 monitoring agents of the RouteViews and RIPE Routing Information Service (RIS) projects, which check in real time which IPv4 /24 network blocks are routable over the BGP. Also 4.2 million /24 networks are monitored by AP alone. Its data can be filtered and displayed on a country-wide and state-wide basis, as well as in time intervals ranging from one day to one month. Data records are available retroactively from 2016-07-23 12:00 p.m. UTC in 10 minute intervals for AP and in 5 minute intervals for BGP and NT.

To avoid *false positive* and *false negative* entries, outages are categorized as such if they are visible by at least AP and NT or by BGP alone where signal measurement errors are less likely.

IODA itself states that only a drop in all three measured values should be meaningfully considered an outage, otherwise the interference would have to be investigated further [6]. More precisely, due to the nature of the measurement with AP, as described in section 2.1, a reliable statement cannot necessarily be made about the occurrence of an outage with this method alone. For this reason, an outage will only be considered by us if a clear drop is also visible in the NT or the BGP signals. Since BGP also monitors more /24 networks than AP, especially if the outage is relatively small, it is possible for an outage to be visible in BGP but not in AP. If an outage was detected in the dashboard, the start and end times are determined by the interval of a deviation from a normal value, consequently the value before the outage (Fig. 1). The start time is set to the data point that is located directly before the start of the outage and the end time is set to the data point at which a normal value could be measured again. Thus the maximum length




Figure 1: Outage capture method for data in IODA Dashboard. The start time is set to the data point that is located directly before the start of the outage and the end time is set to the data point at which a normal value could be measured again. Thus the maximum length of the disturbance can be determined. To determine the severity of the outage, we compare the average value during the outage (disturbance value) to the normal value that was measured directly before the outage and express it in a loss in percentage.

of the disturbance can be determined. To determine the severity of the outage, we compare the average value during the outage (disturbance value) to the normal value that was measured directly before the outage and express it in a loss in percentage. It should be noted that this value can fluctuate over the course of a day, so we will calculate the average of the last hour to get a more precise approximation. As NT does not provide an average value due to the strong fluctuations, we provide the highest percentage deviation in the signal. Because of the temporal resolution, we archive an accuracy of 5 to 10 minutes. Regional filtering in all 16 states allows the geographic impact of a disturbance to be determined directly. We have additionally excluded incorrectly captured outages by IODA in our analysis because of the centralized measurement of IODA from their location in south California. Or if data for a certain period of time is implausible, such as during a complete, worldwide internet outage unmentioned in any news agency. As an example, on 2016-12-29 from 04:00 a.m. to 05:00 a.m. the AP signal dropped to 4% and the NT signal even dropped to 0%, which happened worldwide according to IODA. Since we did not find any evidence that this outage really occurred, we assume that this is a false positive. Another cause of implausible data is related to patch days, which take place every second Tuesday of

the month and can also lead to a potential drop in the readings. An example for a patch day could be seen on 2019-02-13, most notably in states that are located in the "Blue Banana", a dense area in Western and Central Europe: AP as well as NT signals dropped by about 5% for three hours during the night because critical security updates from Microsoft and SAP were distributed. To determine the possible maximum number of devices for the outages, we performed a calculation based on the following formula for AP and BGP: $Max(N_{Devices}) = (S_{normal} - S_{interference}) \times 254$

NT records unique IP addresses that are involved in darknet traffic which mainly affects devices that are behind a firewall and cannot otherwise be detected by AP and BGP. The problem here, however, is that only a subset of darknet traffic is actually scanned, the data varies greatly and, to our knowledge, there is no information on how many devices are involved in darknet traffic on average. Therefore, unfortunately, no reliable statement can be made here about the total number of devices involved for NT.

3.3 Corresponding Events

Using keywords such as "internet disruption", "Germany", "internet outage", a search engine was used to search for internet disruptions with strong media presence and significant impact to identify outages, which we use to compare the methods. We identified six events of relevance (Tab 1). Subsequently, both the iMonitor data and the IODA dashboard were searched and documented for anomalies within the disruption period. The results were substantiated by iMonitor using comments from the disruption reports, visually identified in the IODA dashboard graphs, and reviewed for prominent irregularities in the data.

4 RESULTS

In the following, we provide results for each method (Self-Reporting and Automated Reporting) and then compare each other directly in six corresponding events.

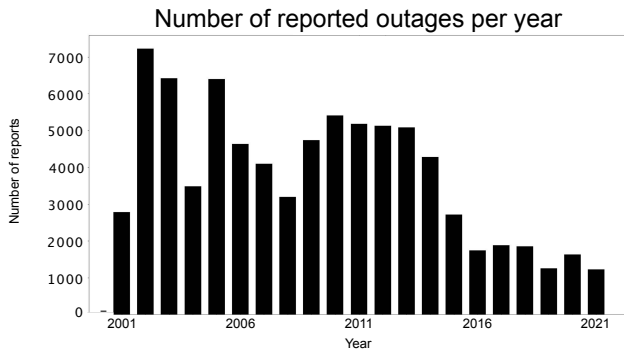
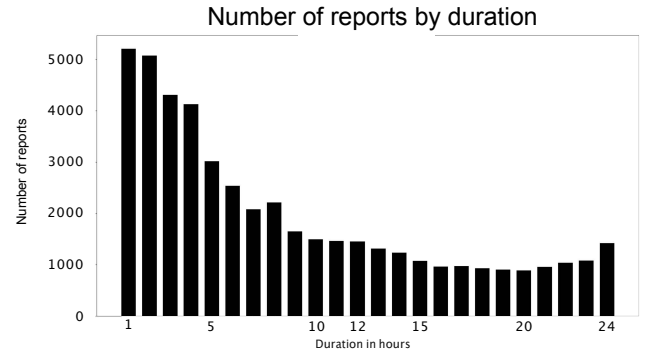
4.1 Self-Reporting with iMonitor

We report the data and thereby consider (1) the number of reported outages per year, have an eye on (2) the 10 biggest outages, (3) show the geographical mapping, and document (4) frequency and duration of all outages.

- (1) In the first step of the analysis of the iMonitor data set, the data are sorted by frequency of reports per date. It can be seen that the number of reports decreases significantly over the last years (Fig. 2). The exact reason for this is not known to us, but possible reasons are that either fewer incidents occurred over time, they were no longer reported, or they were not noticed.
- (2) Considering the largest number of reports, the largest number was reached on 2005-01-10 with 289 reports, followed by 2002-05-06 with 202 and 2005-10-01 with 157. Only two of the ten largest counts could be linked to corresponding outage events, that is on 2005-10-01 (157 reports) with a misconfiguration at the largest ISP in Germany and on 2009-12-15 (134 reports) where a fiber cable was cut during construction work of the Autobahn 45.

Table 1: The six corresponding events used in our analysis, classified by us on the proposed characterization by Aceto et al. [1]

Date	Event	Class	Origin		intentionality		disruption type	
			natural	human	accidental	intentional	primarily physical	pure logical
2021-07-14	Flooding summer 2021	Crisis event	X		X		X	
2021-07-29	BGP Hijacking	Accidental hacking		X	X			X
2021-02-18	ISP disruption Vodafone	Unknown		X	X			
2020-07-17	Cloudflare outage	Misconfiguration		X	X			X
2020-08-18	Equinix power outage	Power disruption	X	X	X		X	
2020-08-13	CenturyLink Lvl(3) routing failures	Misconfiguration		X	X			X

**Figure 2: Amount of outage reports per year between 2001 to 2021 in self-reporting in heise iMonitor.****Figure 3: Amount of reports from 2001 to 2021 by duration within a range of up to 24 hours.**

- (3) In the geographical coverage of the disturbance report, it can only be seen that dense metropolitan areas stand out, whereas the eastern and rural regions have smaller numbers of reports.
- (4) Faults with the same or similar duration are grouped and counted to obtain an overview of the number of faults per duration. As can be seen, there is a higher frequency of faults with a short duration (Fig. 3). Furthermore, it can be noted that 68.53% of all reported faults are resolved within the first 24 hours. Of all disturbances within the first 60 minutes, those with a duration of more than 30 minutes tend to be reported more frequently (59.25%) than shorter disturbances and people tend to round up the reported duration to values in 5 minute time periods (Fig. 4).

4.2 Automated Reporting with IODA

Consequently, we present the results of automated reporting following the methodology described in chapter 3.2. We first show

- (1) 32 distinct outages we identified, give an overview on (2) the number of measured networks, then report (3) the geographical mapping and finally consider the (4) frequency and duration of outages measured with these technologies.

- (1) We went through all available data sets at IODA for each state and were able to record 32 outages according to the criteria we defined (Tab. 2). We assigned an ID to each outage and recorded the values for each state if the outage was visible in multiple ones. At this point it must be mentioned that false positives may still be included in the sample. In contrast to the data analysis in self-reporting, where the presence of data indicates a disturbance, in the case of the IODA data sources an interpretation of the absence of data is necessary to be able to detect an outage. Due to the accompanying imprecision of the methods used to collect data, measurement errors that lead to false positives or false negatives can easily occur. In addition, disruptions that only have a small impact on the

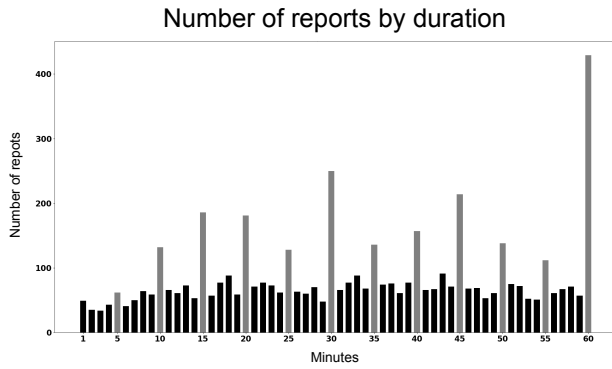


Figure 4: Data analysis of outage duration up to 1 hour from 2001 to 2021, where peaks show the tendency of users' to report striking time frames.

internet availability in a state are not visible in IODA. This makes it difficult to find outages.

- (2) In the collected data, it is noticeable that more /24 networks were measured by BGP between 2016 and 2018 than in the following years. In 2018, there were still 1.6 million networks recorded in Germany, while in 2017 there were only around 470,000. Germany does not have 1.6 million /24 IPv4 networks, as this would correspond to 409,600,000 IPv4 addresses. However, according to RIPE, only 84,432,869 IPv4 addresses are currently assigned to Germany [45]. These 470,000 /24 IPv4 networks would correspond to 120,320,000 IPv4 addresses, which is at least closer to the actual value. This dip in the measured values can be observed primarily in Europe for reasons unknown to us. Since AP and BGP always involve IPv4 /24 subnets, the number of routable devices affected can be calculated. There are up to 254 devices in a subnet. However, a routable device does not necessarily have to be a router in a household, so this value represents only a possible maximum number of affected network devices that have a unique IPv4 address. It should also be mentioned here that the figures refer to the subnets measured by AP and BGP, so there may be a larger number of unreported cases. It is obvious that, for example, the BGP value for outage 10 is much higher than the AP value (Tab 2). This is due to the fact that BGP covers many more networks than AP. The most serious outage in recent years were probably the BGP disruptions in June and July 2020 in Baden-Württemberg, when the signal collapsed by more than 50%. The duration of the disruptions ranged from 10 to 105 minutes. It is also visible that even with a drop of only 1%, several tens of thousands of devices can already be affected. However, this also illustrates that IODA is too inaccurate for measuring small outages, as these would not show up meaningfully in the measurement data.
- (3) Geographic coverage can be determined very accurately at the federal state level. However, in small federal states such as Berlin or Hamburg, a more precise statement can be made about how a disturbance has affected the area. In Bavaria, for

example, a large-scale disturbance in Munich would cause a drop in the graphs, even though the rest of Bavaria might not be affected. Thus, the smaller the state, the more accurate the measurements.

- (4) The duration can be determined very precisely by IODA's real-time measurements, which allows the exact time and duration of the disturbance to be determined. The actual period can range from 10 minutes to 10,095 minutes in our recorded outages. However, if the three disturbances longer than 10,000 minutes are excluded, then the average disturbance duration is 133.04 minutes. It can be seen that BGP outages are comparatively frequent. Especially in Rhineland-Palatinate, BGP outages occur almost monthly, but always with a drop of about 1%. These particular incidents were only recorded on a spot check basis. According to the IODA data, few disturbances occur in Germany. It can be interpreted from this that large disruptions affecting several federal states rarely occur in Germany and that internet availability in this country is stable in most cases. This only applies to disruptions that can be detected by IODA. This will be further elaborated in the discussion.

4.3 Corresponding Events

In order to be able to legitimize the two approaches (Self-Reporting and Automated Reporting), we examined six media-present outages in the iMonitor and IODA sources and evaluated for detectability and traceability (Tab. 1).

4.3.1 Flooding summer 2021. Heavy rainfall in Western Europe caused a flood that largely destroyed buildings and the surrounding infrastructure in Germany and caused 184 casualties [27]. Electricity and transmission towers as well as a large number of households were affected. In iMonitor, there are only isolated fault reports, such as the fault from Eschweiler on 2021-07-15, 07:00 a.m., with the comment: "Complete failure, according to Vodafone hotline due to flooding." Presumably, the few reports are due to the already existing media presence of the incident and the long period of time following the flood. Therefore citizens did not see a reason to report the incident on iMonitor. In IODA, the affected regions were North Rhine-Westphalia and Rhineland-Palatinate in contrast to Lower Saxony. Here, a decrease of AP (blue bar) becomes apparent in the flooded regions on 2021-07-14 12:00 p.m. For example, the "normal value" of AP before the flood in North Rhine-Westphalia was 32,878 and the "disturbance value" was 31,868, which means a decrease of 3.2% (Fig. 5a). In Rhineland-Palatinate, the "normal value" was 8,398 and the "disturbance value" was 8,194, which is a drop of 2.5%. In comparison, no drops were observed in Lower Saxony during the same period.

4.3.2 BGP Hijacking. In this case, IP address blocks were falsely advertised by a provider as their own [22]. As a result, traffic directed to the affected address blocks was redirected to a false destination via the hijacker, namely a Bulgarian provider, according to the analysts. Messages in iMonitor vary from reports such as "Some websites and services are unreachable, but DNS resolution works" to speculations like "Seems to be a routing glitch min. between Vodafone/cable and Telekom/fiber." The IODA dashboard also indicates

an outage during the period because the BGP signal (green line) is no longer present (Fig. 5b). But no data is available for analysis.

4.3.3 ISP disruption at Vodafone. In this case [3], the ISP Vodafone admitted that a large internet outage, for publicly unknown reasons, had occurred. Several Vodafone customers complained in the heise iMonitor about "very slow data transmission and disconnections" and "downstream seems to be okay, but no data is sent via upstream". The IODA graph also shows an outage in the region Berlin for NT (red graph) (Fig. 5c). The disruption primarily affected metropolitan areas, but is not seen outside Berlin.

4.3.4 Cloudflare outage. This [16] was a major DNS configuration error on the part of Cloudflare, which, however, is not recognizable in either the iMonitor or the IODA dashboard between 21:12 and 21:39 UTC, since websites and services were still accessible and pingable via IP (Fig. 5d).

4.3.5 Equinix power outage. Equinix's data center was affected by a power failure, which meant that many network components could no longer be operated [13]. The power failure occurred in the morning hours and there are no reports in the heise iMonitor. However, in the IODA dashboard, the power outage and thus a disruption of the customers served by Equinix, such as the ISPs Sky and Virgin Media, cloud and IT service providers, as well as enterprises, can be seen, as the BGP signal (green line) is missing as of 06:00 p.m. (Fig. 5e).

4.3.6 CenturyLink / Level(3) routing failures. Since CenturyLink / Level(3) is one of the largest network operators in the world, many operators could not be reached, although backup routes from other providers were available, since most of these are supplied by CenturyLink [37]. Sporadic reports on this can be found in the iMonitor. The graph of the IODA dashboard clearly shows a drop in AP (blue curve) and thus shows the outage we are looking for (Fig. 5f). At 10:00 am, the value drops by 39.9% and subsequently by 27.0% on average.

5 DISCUSSION AND CONCLUSION

In the following, we will contrast the methods used and answer research questions per method individually and close with an overall conclusion:

- **RQ1:** What are the *qualitative* differences between self-reporting and automated measurements in terms of accuracy and reliability of the data provided by each method?
- **RQ2:** What are the *quantitative* differences between the methods regarding the amount of data available for each?
- **RQ3:** Which methods are best suited for capturing disturbances in terms of *quality* and *quantity*, and how can they be *combined* for a better result, if necessary?

5.1 RQ1: Qualitative Differences - one human is not so good as automated reporting, but many are.

For research question 1, we state for self-reporting that the recording of faults by end users is subjective and therefore has various qualitative problems. To record the duration of a malfunction via

the heise iMonitor, the start and end of the malfunction must be entered manually in a data record. It is the responsibility of the person reporting the fault to enter this correctly. During our analysis, we found that the end time of the fault was not entered in various data records. Furthermore, the duration is often rounded or estimated by the reporter. The actual start and end of the disturbance may also differ from the values entered, as these may only be perceived with a time delay. In addition, the reported cause is often vague as it cannot be fully surveyed by the one person reporting. The more data records are available, the more accurate an estimate of the actual values can be made. However, in the case of disturbances with only a few recorded data records, no reliable conclusion can be drawn about the actual duration of the disturbance. This can also be seen in the analysis of the corresponding events in section 3.3, where users tend to interpret a possible cause without having the ground truth or any process that e.g. ISPs moderate the self-reported outages. It is a mayor problem, that users' and ISPs do not share common public data but either do it on there own or evidence is hidden in customer service hotlines. In addition, there is no qualitative control of the data records entered, so that any person can also report non-existent faults, which can falsify existing data accordingly. This is not necessarily due to malice, but may be due to the users' lack of knowledge about how to correctly identify an internet malfunction. From this point of view, we found that a qualitative analysis of data collected by self-reporting is often difficult. No qualitative statement can be made about the number of affected households and the total geographic coverage of a disturbance. They can only be put in relation to other reported disturbances in order to obtain an estimate of the severity in relation to each other.

The data of automated reporting methods is objective and recorded as actually measured and thus correspond to the real situation at the time. The actual duration of the disturbance can thus be quantified with an accuracy of up to 15 minutes. Moreover, a more precise estimate of the geographically affected areas can be provided, but it is limited to federal states. Lastly a statement can be made about the frequency of major disturbances in Germany. This makes the data qualitatively meaningful in these respects. However, it is not possible to draw an exact conclusion about the affected households, but only to estimate an upper limit of affected network devices. In particular, measuring the reachability of networks over BGP is reliable and comparatively easy to access. While the NT is not informative on its own, it can strengthen the informative value of the other two methods to more reliably detect failures. But according to the official IODA site, false positives and false negatives exist in the data, which means that a drop in the measured values – especially with AP – does not necessarily indicate a fault.

5.2 RQ2: Quantitative differences - humans detect outages better, if they report.

There are also differences in quantitative manner, which we report with research question 2, which we first highlight for self-reported outages and then for automated measurements. Since faults are explicitly recorded in self-reporting, they do not have to be filtered out of a larger data set and are therefore clearly identifiable. Websites such as the heise iMonitor can be used to record a fault at a fine granular level and thus, for example, the geographical coverage

can be determined down to area codes. Furthermore, in the case of heise iMonitor the data is publicly accessible. The quantitative availability of data strongly depends on the willingness of affected persons to report and may be subject to fluctuations. Due to population distribution, disturbances are more frequently reported in large cities than in rural areas, which could easily lead to the conclusion that disturbances occur less often in rural areas than in large cities. However, this statement cannot be proven by the data. Additionally, the absence of reported faults does not indicate the absence of actual faults, as they may simply not have been reported. Because of this, no estimation for the frequency of outages can be given. Due to the fact that an active internet connection is required to report an internet malfunction, malfunctions may also only be recorded if any other source of connection e.g. mobile internet is available. Actual disruptions of only short duration may not be reported if the effort involved is not proportionate to the actual disruptions. On the other hand, a single report is not meaningful enough to identify an actual fault. This is the reason why several reports are needed to be certain that there has actually been a malfunction. This quantitative requirement is not always given.

Real-time measurement at IODA can be used to provide information timely about possible internet disruptions or outages. IODA is not well suited to reliably detect small disruptions due to the inaccuracy of the data. The detection of interference is made more difficult by dynamic IP addresses, the day-night cycle, public holidays and weekends, as this can cause the measured values to fluctuate without the occurrence of any interference. This leads to the fact that a quantitative availability of data for large disturbances exists, but no meaningful data for small disturbances can be determined. In addition, as mentioned previously false negatives and false positives exist in the data. Combined with the fact that the absence of data in a large data-set must be interpreted to be able to detect an outage, this makes the process of recording them difficult. This makes the detection of interference via IODA difficult. Therefore, IODA is better suited to analyze disturbances whose existence is already certain. As already stated in chapter 2, AP currently only works via IPv4. As things stand in 2021, a significant amount of network traffic on the internet still takes place over IPv4, but IPV6 continues to spread and will probably replace IPv4 in the future. According to Google IPv6 statistics, 56.41% of all devices on the internet in Germany are accessible via IPv6 (as of 2021-12-27) [23]. This change primarily affects end users, which means that citizens in particular are less confronted with IPv4. The integration of IPv6 for AP is an active research topic (as of 2018) [19], but to our current knowledge not yet implemented. Furthermore, as mentioned in 4.2, only problems in OSI layers 1 to 4 can be detected. None of the methods investigated can detect errors in the higher layers, which mainly concerns errors in DNS resolution. Local interferences existing within IODA can affect the worldwide-measured data. This leads to non-existing disturbances displayed for Germany while existing disturbances from that period are overwritten by these false disturbances. This can be observed especially in the months of February 2021 and May 2020. Further, we found large gaps in the measurement data, especially in 2021 for which we don't know the reason of. Together with the points described above, this increasingly limits the quantitative availability of data on faults.

5.3 RQ3: Best Approach and Combinability - humans give the hint, automated reporting the details

One approach to challenge the weaknesses discussed in Research Question 1 and 2, may be a combined approach, which we discuss in Research Question 3. Self-reporting has major qualitative and quantitative deficiencies, which makes it worthwhile to consider it primarily in conjunction with other data sources. In particular, it can be used to gain an impression of users on the ground and to narrow down the spatial effects more precisely. On its own, however, it is not meaningful. The automated measurements of IODA, on the other hand, offer many qualitative strengths, allowing disruptions to be precisely defined in time and spatially determined to a single state. Quantitatively, however, there are also deficiencies. IODA is well suited for a more detailed analysis of already known failures in the past and thus to be able to draw conclusions about future failures if necessary. Furthermore, the data from self-reporting can be used to gain insight into how users perceived a known outage and to get a more detailed overview of the geographical spread.

5.4 Conclusion

As a conclusion, looking at all methods together, in terms of quality, they all unfortunately have gaps that make a comprehensive investigation of internet disruptions of every possible kind unfeasible. The quantitative availability of data is also insufficient in all cases. Particularly in the case of self-reporting, a qualitative examination of disruptions is often not possible, partly due to the quantitative lack of data. IODA has to deal with IPv4 and measurement inaccuracies, which make it difficult to capture disturbances. In summary, all of the methods presented are not satisfactory for reliably capturing internet faults as well as outages and keeping statistics on them with reliable conclusions. But as we have seen in section 4.3, IODA is better suited for analyzing known outages than for finding them. Humans can put a lens on specific outages by reporting them, but automated reporting is more accurate to gain the temporal and spatial dimension of these specific outages.

6 LIMITATIONS AND OUTLOOK

One limitation lies in the nature of the network protocols: If the measurement gaps of 2021 continue and the integration of IPv6 is not successful in the methods used for data collection, then IODA will no longer be relevant in the future. Another limitation concerns the decreased amount of reports in the self-reporting database over the last years and the scope covering outages only in Germany, which might not be transferable to all regions. Last, to mention even with having the six corresponding events to compare with ground truth, there is still information missing knowing exactly all details of an outage to cover the whole story for deep analysis. Since automated reporting does not directly measure mobile networks, this work focuses on outages in the general internet infrastructure or broadband connections rather than mobile internet. But several problems like outages based on routing failures affect users regardless of how they connect.

Regarding *automated reporting*, there are some projects to mention, that can improve or complement the methods used by IODA but are currently on a more experimental level:

- Thunderping is a variant of Trinocular that only activates during extreme weather events, such as hurricanes, and analyzes the internet accessibility in the affected regions. The project is limited to the USA but raw data can be requested for analysis, also there already exist visualizations for hurricane sandy (2012) and hurricane irene (2011) [42].
- Disco [43] is also a variant of Trinocular, but its measurement is provided peripheral through worldwide distributed agents, which collect data locally send them to a central collecting point [43]. The data is thereby less susceptible to disturbances of central monitoring sites. This provides a more accurate perspective of how the end-user perceives internet outages. However, this approach is limited to the distribution of the (Ripe Atlas Probes) Agents.
- Passive from Content Delivery Network (CDN) is an approach to analyze the CDN for drops in its data volume. CDN provides large media data for end-users such as for streaming. Detecting irregularities in the amount of transferred data volume can help in detecting outages. The plan here is to provide a better alternative for NT.

There might be new approaches for *self-reporting* using more advanced technologies like app based local measurements and guided reporting instead of manual action of filling an online form. There might also be efforts of ISPs to be more transparent towards outages by integrating users on a more open level like working together and using the potential of crowdsourcing. This may also be done by more regulations from the authorities. However, it remains to be seen what a combined approach of self-reporting and automated reporting might look like. A pure reduction of self-reporting for outage detection and automated reporting for outage analysis is certainly too short-sighted. It would make sense to carry out a evaluation with experts and users how both phases (detection and analysis) can benefit from the other approach.

ACKNOWLEDGMENTS

Major thanks to Anja-Liisa Gonsior who supported the overall writing process. This work has been funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY center and by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

REFERENCES

- [1] Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, and Antonio Pescapé. 2018. A comprehensive survey on internet outages. *Journal of Network and Computer Applications* 113 (2018), 36–63. <https://doi.org/10.1016/j.jnca.2018.03.026>
- [2] Tooba Ahsen, Zi Yi Lim, Aaron L. Gardony, Holly A. Taylor, Jan P de Ruiter, and Fahad Dogar. 2021. The Effects of Network Outages on User Experience in Augmented Reality Based Remote Collaboration - An Empirical Study. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 313 (oct 2021), 27 pages. <https://doi.org/10.1145/3476054>
- [3] Katarina Amtmann and Momir Takac. 2021. Vodafone-Störung in Deutschland: Massive Probleme brachten Kunden zum Toben - „Der größte Witz“. <https://www.merkur.de/verbraucher/probleme-vodafone-stoerung-deutschland-unitymedia-kunden-internet-telefon-news-probleme-zr-90208844.html>. *Merkur* (2021). [Online; accessed December-2021].
- [4] Ritwik Banerjee, Abbas Razaghpahan, Luis Chiang, Akash Mishra, Vyas Sekar, Yejin Choi, and Phillipa Gill. 2015. Internet Outages, the Eyewitness Accounts: Analysis of the Outages Mailing List. In *Passive and Active Measurement*, Jelena Mirkovic and Yong Liu (Eds.). Springer International Publishing, Cham, 206–219. https://doi.org/10.1007/978-3-319-15509-8_16
- [5] Ryan Bogutz, Yuri Pradkin, and John Heidemann. 2019. Identifying Important Internet Outages. In *2019 IEEE International Conference on Big Data (Big Data)*. 3002–3007. <https://doi.org/10.1109/BigData47090.2019.9006537>
- [6] CAIDA. 2019. IODA Dashboard screencast. <https://www.youtube.com/watch?v=VzOv7g1Xy3k>. 6:50 min [Online; accessed December-2021].
- [7] Edmond WW Chan, Xiapu Luo, Waiting WT Fok, Weichao Li, and Rocky KC Chang. 2011. Non-cooperative diagnosis of submarine cable faults. In *International Conference on Passive and Active Network Measurement*. Springer, 224–234.
- [8] Ksenia Chmutina and Jason Von Meding. 2019. A dilemma of language: “Natural disasters” in academic literature. *International Journal of Disaster Risk Science* 10, 3 (2019), 283–292.
- [9] Kenjiro Cho, Cristel Pelsser, Randy Bush, and Youngjoon Won. 2011. The Japan Earthquake: The Impact on Traffic and Routing Observed by a Local ISP. In *Proceedings of the Special Workshop on Internet and Disasters (Tokyo, Japan) (SWID '11)*. Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. <https://doi.org/10.1145/2079360.2079362>
- [10] European Commission, Joint Research Centre, L Feyen, E Krausmann, G Karagiannis, and Z Turksezer. 2017. *Climate change and critical infrastructure : flood*. Publications Office. <https://doi.org/10.2760/437836>
- [11] North Bay/North COast Broadband Consortium. 2019. Telecommunications Outage Report: Northern California Firestorm 2017. (2019). <https://ecfsapi.fcc.gov/file/1053130424752/EAS-1-NBNCBC-Telecommunications-Outage-Report-2017-Firestorm.pdf>
- [12] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. 2014. Analysis of Country-Wide Internet Outages Caused by Censorship. *IEEE/ACM Trans. Netw.* 22, 6 (dec 2014), 1964–1977. <https://doi.org/10.1109/TNET.2013.2291244>
- [13] Caroline Donnelly. 2020. Equinix confirms review under way into UPS failure at Docklands datacentre. <https://www.computerweekly.com/news/252487861/Equinix-confirms-review-under-way-into-UPS-failure-at-Docklands-datacentre>. *Computer Weekly IT* (2020). [Online; accessed December-2021].
- [14] Center for Applied Internet Data Analysis (CAIDA). 2021. Internet Outage Detection and Analysis (IODA). <https://ioda.caida.org/ioda>. [Online; accessed December-2021].
- [15] Vasileios Giotas, Christoph Dietzel, Georgios Smaragdakis, Anja Feldmann, Arthur Berger, and Emile Aben. 2017. Detecting Peering Infrastructure Outages in the Wild. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (Los Angeles, CA, USA) (SIGCOMM '17)*. Association for Computing Machinery, New York, NY, USA, 446–459. <https://doi.org/10.1145/3098822.3098855>
- [16] John Graham-Cumming. 2020. Cloudflare outage on July 17, 2020. <https://blog.cloudflare.com/cloudflare-outage-on-july-17-2020/>. [Online; accessed December-2021].
- [17] Sukeshini A Grandhi, Linda Plotnick, and Starr Roxanne Hiltz. 2020. An Internet-Less World? Expected Impacts of a Complete Internet Outage with Implications for Preparedness and Design. *Proc. ACM Hum.-Comput. Interact.* 4, GROUP, Article 03 (jan 2020), 24 pages. <https://doi.org/10.1145/3375183>
- [18] Steffen Haesler, Ragnar Mogk, Florentin Putz, Kevin T. Logan, Nadja Thiessen, Katharina Kleinschnitger, Lars Baumgärtner, Jan-Philipp Stroscher, Christian Reuter, Michele Knott, and Matthias Hollick. 2021. Connected Self-Organized Citizens in Crises: An Interdisciplinary Resilience Concept for Neighborhoods. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing (Virtual Event, USA) (CSCW '21)*. Association for Computing Machinery, New York, NY, USA, 62–66. <https://doi.org/10.1145/3462204.3481749>
- [19] John Heidemann, Yuri Pradkin, and Guillermo Baltra. 2018. The Policy Potential of Measuring Internet Outages. In *Proceedings of the TPRC, the Research Conference on Communication, Information and Internet Policy* (johnh: pafile). Washington, DC, USA.
- [20] heise Medien GmbH & Co. KG. 2021. iMonitor – Internet-Störungen. <https://www.heise.de/netze/netzwerk-tools/iMonitor-internet-stoerungen/>. [Online; accessed December-2021].
- [21] Ira Helsloot and Ralf Beerens. 2009. Citizens’ Response to a Large Electrical Power Outage in the Netherlands in 2007. *Journal of Contingencies and Crisis Management* 17, 1 (2009), 64–68. <https://doi.org/10.1111/j.1468-5973.2009.00561.x>
- [22] Martin Holland. 2021. BGP-Hijacking: Massive Internet-Störungen im Festnetz der Telekom. <https://www.heise.de/news/Internet-Massive-Probleme-im-Festnetz-der-Telekom-6150438.html>. *heise online* (2021). [Online; accessed December-2021].
- [23] Google IPv6. 2022. Statistics. <https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption>. [Online; accessed December-2021].
- [24] International Telecommunication Union (ITU). 2021. Measuring digital development. Facts and figures 2021. <https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2021.pdf>. [Online; accessed December-2021].

- [25] Sangeetha Abdu Jyothi. 2021. Solar Superstorms: Planning for an Internet Apocalypse. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) (SIGCOMM '21). Association for Computing Machinery, New York, NY, USA, 692–704. <https://doi.org/10.1145/3452296.3472916>
- [26] Marc-André Kaufhold, Jasmin Haunschild, and Christian Reuter. 2020. Warning the Public: A Survey on Attitudes, Expectations and Use of Mobile Crisis Apps in Germany. In *Proceedings of the European Conference on Information Systems (ECIS)*. AIS.
- [27] Kreienkamp, Frank and Philip, Sjoukje Y. and Tradowsky, Jordis S. and Kew, Sarah F. and Lorenz, Philip and Arrighi, Julie and Belleflamme, Alexandre and Bettmann, Thomas and Caluwaerts, Steven and Chan, Steven C. and Ciavarella, Andrew and De Cruz, Lesley and de Vries, Hylke and Demuth, Norbert and Ferrone, Andrew and Fischer, rich M. and Fowler, Hayley J. and Goergen, Klaus and Heinrich, Dorothy and Henrichs, Yvonne and Lenderink, Geert and Kaspar, Frank and Nilson, Enno and Otto, Friederike E L and Ragone, Francesco and Seneviratne, Sonia I. and Singh, Roop K. and Skålevåg, Amalie and Termonia, Piet and Thalheimer, Lisa and van Aalst, Maarten and Van den Bergh, Joris and Van de Vyver, Hans and Vannitsem, Stéphane and van Oldenborgh, Geert Jan and Van Schaeybroeck, Bert and Vautard, Robert and Vonk, Demi and Wanders, Niko. 2021. Rapid attribution of heavy rainfall events leading to the severe flooding in Western Europe during July 2021. , 51 pages.
- [28] Zheng (Eddie) Li, Mingfei Liang, Liam O'Brien, and He Zhang. 2013. The Cloud's Cloudy Moment: A Systematic Survey of Public Cloud Service Outage. *International Journal of Cloud Computing and Services Science (IJ-CLOSER)* 2 (12 2013). <https://doi.org/10.11591/closer.v2i5.5125>
- [29] Nicole Lupien, Sukeshini A. Grandhi, Linda Plotnick, and Star Roxanne Hiltz. 2017. Wait, Did You Say No Internet? An Exploratory Study of the Perceived Impact of Internet Outage. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17 Companion). Association for Computing Machinery, New York, NY, USA, 231–234. <https://doi.org/10.1145/3022198.3026332>
- [30] Ratul Mahajan, David Wetherall, and Tom Anderson. 2002. Understanding BGP Misconfiguration. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (Pittsburgh, Pennsylvania, USA) (SIGCOMM '02). Association for Computing Machinery, New York, NY, USA, 3–16. <https://doi.org/10.1145/633025.633027>
- [31] Steve Mansfield-Devine. 2016. DDoS goes mainstream: how headline-grabbing attacks could make this threat an organisation's biggest nightmare. *Network Security* 2016, 11 (2016), 7–13.
- [32] A Ogielski and J Cowie. 2002. Internet routing behavior on 9/11 and in the following weeks. *Renesis Corporation* (2002), 5–6.
- [33] LLC. Downtdetector Ookla. 2021. Downtdetector. <https://downtdetector.com/>. [Online; accessed December-2021].
- [34] Gerard O'Reilly, Ahmad Jrad, Ramesh Nagarajan, Theresa Brown, and Stephen Conrad. 2006. Critical Infrastructure Analysis of Telecom for Natural Disasters. In *Networks 2006. 12th International Telecommunications Network Strategy and Planning Symposium*. 1–6. <https://doi.org/10.1109/NETWKS.2006.300396>
- [35] Partha Sarathi Paul, Bishakh Chandra Ghosh, Ankan Ghosh, Sujoy Saha, Subrata Nandi, and Sandip Chakraborty. 2020. *Disaster Strikes! Internet Blackout! What's the Fate of Crisis Mapping?* Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3379503.3403532>
- [36] Udit Paul, Alexander Ermakov, Michael Nekrasov, Vivek Adarsh, and Elizabeth Belding. 2020. #Outage: Detecting Power and Communication Outages from Social Networks. Association for Computing Machinery, New York, NY, USA, 1819–1829. <https://doi.org/10.1145/3366423.3380251>
- [37] Matthew Prince. 2020. August 30th 2020: Analysis of CenturyLink/Level(3) Outage. <https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage/>. [Online; accessed December-2021].
- [38] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In *Proceedings of the ACM SIGCOMM Conference* (johnh: pafille). ACM, Hong Kong, China, 255–266. <https://doi.org/10.1145/2486001.2486017>
- [39] Christian Reuter. 2014. Communication between power blackout and mobile network overload. *International Journal of Information Systems for Crisis Response and Management (IJISCRAM)* 6, 2 (2014), 38–53.
- [40] Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. 2018. Advancing the Art of Internet Edge Outage Detection. In *Proceedings of the Internet Measurement Conference 2018* (Boston, MA, USA) (IMC '18). Association for Computing Machinery, New York, NY, USA, 350–363. <https://doi.org/10.1145/3278532.3278563>
- [41] G James Rubin and M Brooke Rogers. 2019. Behavioural and psychological responses of the public during a major power outage: A literature review. *International Journal of Disaster Risk Reduction* 38 (2019), 101226. <https://doi.org/10.1016/j.ijdr.2019.101226>
- [42] Aaron Schulman and Neil Spring. 2011. Pingin' in the Rain. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference* (Berlin, Germany) (IMC '11). Association for Computing Machinery, New York, NY, USA, 19–28. <https://doi.org/10.1145/2068816.2068819>
- [43] Anant Shah, Romain Fontugne, Emile Aben, Cristel Pelsser, and Randy Bush. 2017. Disco: Fast, good, and cheap outage detection. In *2017 Network Traffic Measurement and Analysis Conference (TMA)*. 1–9. <https://doi.org/10.23919/TMA.2017.8002902>
- [44] Nishant Shah. 2021. Digital Infrastructure, Liminality, and World-Making Via Asia| (Dis)information Blackouts: Politics and Practices of Internet Shutdowns. *International Journal of Communication* 15, 0 (2021). <https://ijoc.org/index.php/ijoc/article/view/13977>
- [45] Regional Internet Registries Statistics. 2022. Germany (DE) - IPv4 address statistics. https://www-public.imtbs-tsp.eu/~maigron/RIR_Stats/RIPE_Allocations/IPv4/ByNb/DE.html. [Online; accessed December-2021].
- [46] Murray Turoff, Victor Bañuls, Linda Plotnick, and Starr Hiltz. 2014. Development of a dynamic scenario model for the interaction of critical infrastructures. In *Proceedings of the 11th International ISCRAM Conference*.
- [47] T Wilson, G Kaye, C Stewart, and J Cole. 2007. Impacts of the 2006 eruption of Merapi volcano, Indonesia, on agriculture and infrastructure. *GNS Science Report* (2007), 69.

Table 2: The results of our analyses for AP, BGP, and NT based on 32 identified outages with date, duration and measurements for affected states in Germany.

ID	Start	Duration	State	AP				BGP				NT
				normal	new	↓%	devices	normal	new	↓%	devices	
1	2016-12-20 5:35 a.m.	2 h 35 m	BY	-	-	-	-	946064	932753	1,4	3380994	-
			BE	-	-	-	-	832233	817897	1,8	3641344	-
2	2017-03-08 9:00 a.m.	1 d 3 h 35 m	BY	-	-	-	-	945404	900116	5	11503152	-
			HB	-	-	-	-	317690	275309	15,4	10764774	-
			HE	-	-	-	-	1227905	1179198	4,1	12371578	-
			NW	-	-	-	-	1205295	1158889	4	11787124	-
3	2017-09-12 3:55 a.m.	15 m	BW	-	-	-	-	830864	815696	1,9	3852672	-
			BY	-	-	-	-	902360	879381	2,6	5836666	-
			HE	-	-	-	-	1226435	1195105	2,6	7957820	-
4	2017-10-21 6:30 p.m.	7 d 1 h 0 m	SL	-	-	-	-	202268	199960	1,2	586232	-
5	2017-12-10 9:05 a.m.	7 d 0 h 5 m	BW	-	-	-	-	851437	778341	9,4	18566384	-
6	2018-02-17 10:55 a.m.	30 m	BW	-	-	-	-	832471	787116	5,8	11520170	-
			HH	-	-	-	-	732842	694562	5,5	9723120	-
			BB	-	-	-	-	252774	236392	6,9	4161028	-
			HE	-	-	-	-	1234151	1185943	4,1	12244832	-
			NW	-	-	-	-	1125101	1080735	4,1	11268964	-
7	2018-02-22 7:50 a.m.	2 d 14 h 20 m	BE	-	-	-	-	707522	682946	3,6	6242304	-
8	2018-12-16 2:00 a.m.	7 d 0 h 15 m	BB	-	-	-	-	12250	11917	2,8	84582	-
			MV	-	-	-	-	6421	6343	1,2	19812	-
9	2019-02-06 1:50 a.m.	3 h 10 m	MV	3724	3571	4,2	38.946	-	-	-	-	6,2
			ST	5193	4749	9,3	112.756	-	-	-	-	6,1
			BW	23173	22562	2,7	155.385	-	-	-	-	1,6
10	2019-02-07 11:25 p.m.	15 m	HE	27301	27123	0,9	45.423	57678	56638	1,8	264160	-
11	2019-02-25 3:05 p.m.	1 h 20 m	RP	-	-	-	-	24771	24515	1,04	65024	-
12	2019-02-28 3:05 p.m.	3 h 5 m	RP	-	-	-	-	24771	24523	1,01	62992	-
13	2019-03-27 7:20 a.m.	20 m	RP	-	-	-	-	24914	24658	1,03	65024	-
14	2019-04-02 6:45 a.m.	8 h 55 m	TH	-	-	-	-	8295	8158	1,01	34798	-
15	2019-05-05 1:45 a.m.	1 h 35 m	BW	-	-	-	-	122867	118373	0,7	1141476	-
16	2019-05-20 3:35 a.m.	4 h 45 m	HB	-	-	-	-	4752	4684	1,6	17272	-
17	2019-05-22 12:55 a.m.	2 h 35 m	TH	4615	4439	4	44.592	-	-	-	-	4
18	2019-05-23 12:50 a.m.	3 h 10 m	TH	4605	4408	4,5	50.002	-	-	-	-	1,3
19	2019-06-19 5:00 p.m.	1 h 0 m	HB	-	-	-	-	4769	4512	0,5	65278	-
20	2019-11-04 10:00 p.m.	5 h 20 m	RP	-	-	-	-	19180	18932	1,3	62992	-
21	2019-12-09 10:10 p.m.	5 h 15 m	RP	-	-	-	-	18774	18526	1,3	62992	-
22	2019-12-09 9:00 p.m.	4 h 30 m	TH	-	-	-	-	8234	8095	1,7	35306	-
23	2020-03-21 8:05 a.m.	1 h 30 m	BY	-	-	-	-	88532	87504	1,1	261112	-
24	2020-06-20 1:10 p.m.	1 h 45 m	BW	-	-	-	-	124033	58512	52,9	16642334	-
25	2020-06-26 11:00 p.m.	50 m	TH	3998	3928	1,8	17.991	8516	8271	2,9	62230	-
26	2020-07-18 7:10 a.m.	45 m	BW	-	-	-	-	124221	58698	52,8	16642842	-
27	2020-10-23 11:05 p.m.	10 m	BW	-	-	-	-	124135	58606	52,8	16644366	-
28	2021-06-29 6:05 a.m.	1 h 55 m	TH	4208	4158	1,2	12.751	-	-	-	-	4
29	2021-09-22 12:50 p.m.	3 h 50 m	ST	4296	4166	3,1	33.099	10261	10051	2,1	53340	-
30	2021-09-22 4:50 a.m.	40 m	ST	4234	4142	2,2	23.454	10261	10051	2,1	53340	-
31	2021-09-22 6:00 p.m.	25 m	ST	4253	4182	1,7	17.950	10261	10051	2,1	53340	-
32	2021-10-13 7:10 a.m.	1 h 40 m	SL	2492	2186	13,9	77.681	-	-	-	-	17



Figure 5: IODA Dashboard views of the selected cases. Blue = AP, Green = BGP, Red = NT