On Large-scale IP service Disruptions Dependencies

Alfred Arouna*†, Ioana Livadariu†, Azan Latif Khanyari*†, Ahmed Elmokashfi†
* *Oslo Metropolitan University*, Oslo, Norway
† *Simula Metropolitan*, Oslo, Norway

***Abstract*—Large part of the human activities has become digitalized as many of the daily activities rely on the Internet. However, this critical infrastructure is subject to disruption that can be caused by malicious actors or misconfiguration. Network operators are expected to implement community driven best-practices that contribute to the Internet's resilience. In practice many operators fail to implement them due to the high complexity of these tasks. Consequently, large-scale disruption of IP services can impact numerous customers through a cascading effect. In this work, we assess how three major Internet disruption impact their connectivity. We found that these disruptions cascade from the affected operators towards large part of their customers. However, customers that follow best practices have a short recovery time and rely on rerouted paths through a myriad of topologically close backups links.**

***Index Terms*—BGP, dependency, IP service disruption**

I. Introduction

Connecting a large part of humanity, the Internet represents one of the most critical infrastructures in use today. The importance of the Internet resilience to the daily lives of billions cannot be overstated, yet parts of this infrastructure keep failing every now and then [1], [2]. Many factors like misconfigurations or failures in the Internet Service Provider (ISP) networks and temporary degradation in performance can contribute to such disruptions.

Internet has thus far withstood major outages as well as unprecedented sudden changes in traffic patterns and volumes. Depending on their root cause, such disruptions can potentially impact numerous users. Moreover, outages that occurs in the Internet core can potentially propagate to customer networks, which can lead to service unavailability for users [3]. Despite the increasing frequency of such events, our understanding of the cascading effects of such events is mostly provided by commercial monitoring platforms [4]. Due to highly distributed and interconnected nature of the Internet, it is challenging to build a holistic view that encompass the underlying dependencies that cause the Internet disruption propagation towards different networks. This paper aims to bring us closer to understanding such dependencies by looking at how the Internet managed to absorb three major Internet disruptions. The first incident was a routing leak by Verizon. The second incident was caused by a misconfiguration in CenturyLink network that effectively disconnected this organization from the Internet. The third incident was a major failure at the London INternet eXchange (LINX).

We use both IPv4 control plane (routing table dumps and update messages) and data plane (end-to-end traceroutes) measurements to track and understand the impact of these incidents. We observe that large part of ASes use backup and unseen/new links to recover shortly after experiencing network instability. However, due to their chain of dependency, small ASes are more prone to failure, regardless of the type of incident. Furthermore, existing tools do not seem to capture significantly inter-connectivity changes for small ASes. Therefore, we encourage collaborative approach to strengthen best practices implementation and monitoring. We also suggest additional data plane measurements to complement control plane partial visibility of the cascading effect.

II. Background & Related Work

A. Related work

Several papers focus on detecting and classifying BGP vulnerabilities by leveraging control [5] or data plane [6]. Several studies propose hybrid approaches by relying on both data and control plane [7]. However, the scope of these studies is limited to the classification and/or detection of specific types of malicious activity, while another body of work studies large-scale Internet failures [8]–[13].

Our goal, however, is to characterize the impact of BGP incidents from major networks on both control and data plane, as well as to understand the lack and presence of impact across different type of networks. We focus in this study on three major BGP incidents: Verizon leak in 2019, CenturyLink outage in 2020 and LINX outage in 2021. Several online articles quickly reported on these major incidents [3], [14], [15], but to achieve our goal, further analysis is needed.

B. BGP Incidents

Verizon leak: On Monday, June 24th, 2019, Verizon leaked a series of prefixes hijacked by a smaller provider, i.e., DQE Communications (AS33154). Lasting more than three hours, the incident affected CDNs and Alexa top websites [14].

CenturyLink outage: At the end of August 2020, CenturyLink experienced an outage that affected multiple services and spanned more than 8 hours [3]. The outage was traced back to a misconfigured Flowspec announcement [16]. Forcing CenturyLink to request other tier-1 providers to de-peer.

LINX outage: On 23 March 2021, the LINX, suffered a major outage for more than 6 hours [17]. Four hours after the incident began, a faulty switch at the LINX LON1 facility lost significant traffic, resulting in a shift to alternative routes to AMS-IX and DE-CIX [15].

978-3-903176-59-1 ©2023 IFIP

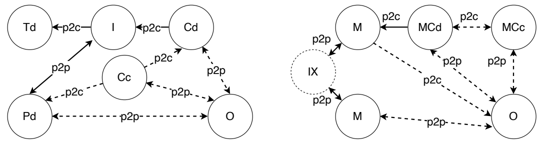

(a) Verizon and CenturyLink (b) LINX

Fig. 1: AS categories from the *I*(ncident) and IXP ASes perspective. *M*(ember) are IXP members. *Cd* (*MCd*) and *Cc* (*MCc*) ASes are customers of *I* and *M*. *Pd* are direct peers of *I*. *O*(thers) are the remaining ASes.

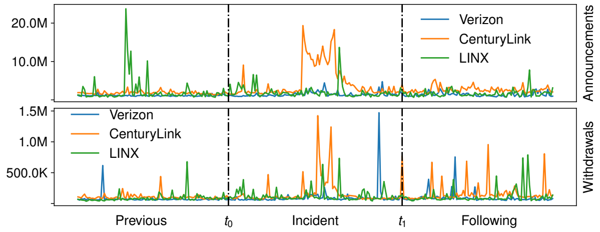

Fig. 2: BGP updates of prefixes per 15 minutes.

C. AS Incidents Categories

We devise six classes of ASes based on the business relations of the AS and the topological distance to the incident AS(es)/IXP AS members as in Figure 1. We identify the incident ASes (*I*) and the LINX AS members (M) from the CAIDA's AS Rank [18] and Internet eXchange Points (IXPs) datasets [19]. Using the AS Relationship dataset [20], we extract customers ASes directly connected to the incident ASes (IXP members), i.e., *Cd* (*MCd*). *Cc* (*MCc*) are customers of customers following p2c relationships from *Cd* (*M*). *Td* are transit ASes for any *I* and *Pd* are direct peers of incident ASes. We group the unclassified ASes into the *O*(thers) category.

III. Control Plane Impact

A. Data collection and processing

For each incident, we use BGPStream [21] to collect IPv4 BGP routing tables and updates from full-feed monitors from 17 BGP collectors over the span of three days centered in the incident day. We consider prefixes seen by at least 10 monitors per day as this eliminates prefixes observed locally by a single collector and/or within a limited region. Using the BGP AS-PATH attribute collected from the monitors, we built prefix-to-origin mappings that we use to enhance CAIDA's public prefix-to-AS mappings dataset [22].

Figure 2 shows the number of aggregated updates every 15 minutes. We find a significant increase for both types of updates. *Verizon leak* has the lowest scale for both number of updates and duration. Comparing the BGP activity during the incident to the previous day shows an 137% increase in the number of updates. Analyzing *CenturyLink outage*, we find that the announcements and withdrawals increase with more than 19M and 1.3M, respectively. Compared with the previous day, this amounts to 475% and 224% increase for announcements and withdrawals, respectively. Moreover, we observe an increased number of BGP updates during the day following the event. The impact of the LINX outage is also visible during the incident day as well as the following day. We observe an increase of more than 211% in announcements and 8.6% withdrawals. Note that the prefixes having high number of announcements prior to the outage day are not involved in the LINX outage, and thus we exclude them from our analysis.

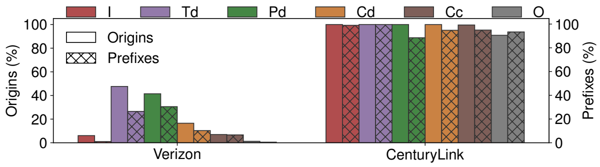

Fig. 3: Percentage of impacted AS origins and prefixes per AS category by the CenturyLink outage and Verizon leak.

B. Impact on Small ASes

Next, we investigate how different networks were impacted. To this end, we first consider prefixes for which the incident AS appears on the AS-PATH prior the incident day and analyze the BGP announcements during the incident. Using the prefix-to-AS dataset we extract the origin AS of the impacted prefixes and classify each AS into one of the incident categories. Figure 3 shows the percentage of ASes and prefixes per category impacted by the CenturyLink and Verizon incidents.

The *Verizon leak* impacted more than 8.25% of the 800K routed prefixes. These prefixes were advertised by 5K ASes which were dominated by Pd and Td categories. However, indirect customers (Cc) ASes dominated the 43% of ASes who's entire IPv4 space was affected. These ASes include the hijacking source AS33154 and its customer Allegheny Technologies Inc (AS396531). The ratio of fully impacted ASes was higher than 50% during *CenturyLink outage*. Most of these networks were customers of customers of CenturyLink (Cc). Specially, of CenturyLink's ASes, AS3356 experienced the highest impact. *LINX* members (M) were highly impacted (91%), but only 21 of these impacted networks have all their advertised prefixes affected. We hypothesize that these members were connected to the faulty switch [17].

C. Traffic Rerouting

To understand how networks reacted to the incidents, we track prefixes activity during our measurement period. To minimize routing table dump on the available prefixes, we use an eight-hour granularity. We also exclude from this analysis prefixes advertised by the incident ASes and monitors from these ASes, as we seek to understand the incidents impact on other networks. Since LINX is an IXP, we use the BGP Route Server AS8714. For the other two incidents we consider the most impacted ASes – Verizon's AS701 and CenturyLink's AS3356 [3], [14]. Using the topology snapshot of all the first hops per prefix/origin on the day before the incident as a baseline, we track for each prefix the first hop during and

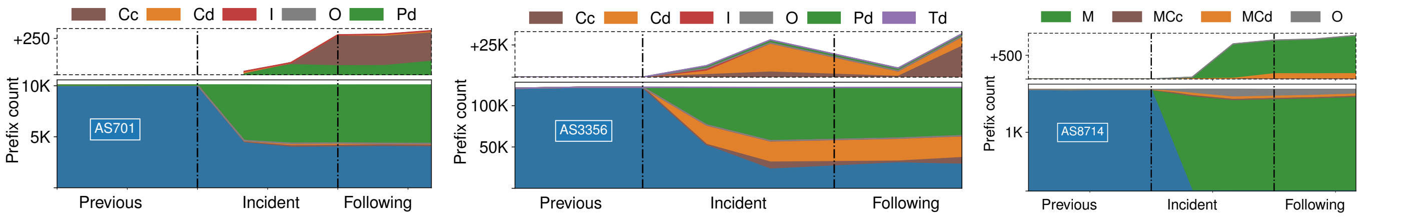

(a) AS701 for Verizon leak (b) AS3356 for CenturyLink outage (c) AS8714 for LINX outage

Fig. 4: Number of prefixes subject to different routing practices on the incident and following day. First hops are grouped into AS categories. Exiting and backup are plotted in the bottom part which new/unseen links are in the top part.

after the incident day. Depending on the observed changes in the first hop after the start of the incidents, we classify the prefixes into the following three classes. *Same link* are prefixes that use the incident AS as first hop. *Backup link* refer to prefixes that start using a different link. *Unseen/New link* are prefixes that use a first hop that was not seen by any of the full-feed monitors before the incidents. Figure 4 shows the traffic rerouting for each BGP event. The bottom part illustrates the backup and previously visible links while the new/unseen links are above the dotted line.

During the *Verizon leak*, 78% of the 1K ASes use AS701 to reroute their traffic. Large part of these ASes used *backups* links while only 1% used *unseen/newly* visible BGP sessions. Approximately 2.3% of the ASes relied on both links. However, most of the ASes of the *same link* class were direct customer ASes (Cd), while 60% of ASes using *backup links* were Pd ASes. Upstreams of these Pd ASes were dominated by tier-1 networks like CenturyLink and AT&T. Some of the Pd ASes were also initially used for *new links* before being replaced the following day by Cc ASes.

During the *CenturyLink outage*, 88% of the ASes changed their upstream. Most of the rerouted traffic relied on *backup links*, while a small percentage used *newly visible links*. Apart from CenturyLink's own ASes, ASes of the *same link* class were 98% dominated by a myriad of Cd ASes. In addition to Cd, Pd ASes dominated the 5K *backup* upstreams at 28% and 37% respectively. However, Cd (53%) largely dominated *new links*, followed by Cc ASes (17%), with traffic rerouted through regionally ASes.

After the start of the *LINX outage* 57% of members rerouted using only *backup links*, while the remaining ASes relied on both *backup* and *unseen/new links*. Members using *backup links* rerouted their traffic through other members (36%) and members direct customers (34%). Similarly, 75% of *unseen links* appeared between members and their direct customers. Note that the heavily used *new links* appeared after a significant loss of traffic from LON1 [17].

IV. Data Plane Impact

A. Data collection

We collect and analyze IPv4 traceroute data extracted from the public RIPE Atlas measurements [23]. Contrary to our control-plane analysis, we limit our measurements period to the incident day, i.e., before, during and after the incident.

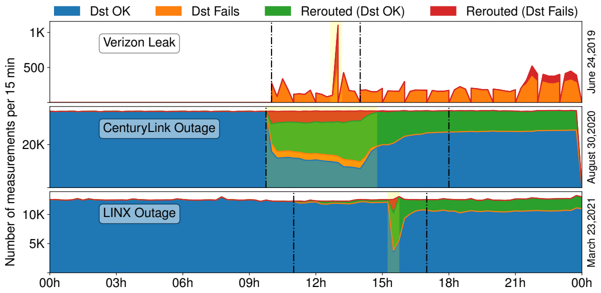

Fig. 5: **Filtered traces on incident day per incident**. Traceroutes can fail (all Dst Fails) or be successful (all Dst OK).

In the case of *CenturyLink and LINX outages*, we filter measurements that traverses ASes owned by these two organizations using our constructed prefix-to-AS mappings. This yields approximately 4M traceroutes for the former and 1.5M for the latter incident, from 48K and 16K source/destination (src/dst) pairs respectively. From CenturyLink measurements we filter out IP pairs between which we identify load balancing[1] prior to the incident. This step reduces our collected CenturyLink traceroutes to 3.4M. For the *Verizon leak* we follow a different approach to gauge its impact on the data plane. A prefix hijack typically results in traffic blackholing, which will lead to failed traceroutes. Thus, we track traceroutes from 192K IP src/dst pairs that include any of the 65K leaked prefixes.

B. Performance impact

We started our performance analysis by evaluating the selected traceroutes in terms of *reachability* and *AS-level path stability*. We rely on the RIPE Atlas Sagan[2] traceroute module to classify traceroutes into *successful* or *failed* depending on whether the measurement reach their destination or not. Next, we check whether we observe any change in the AS path. Specifically, we consider a path *rerouted* if the incident AS disappears from the AS path after the start of the incident. We thus divide out selected measurements into

[1] An IP pairs can use more than one path between the source and the destination, with one of the path relying on CenturyLink.

[2] https://ripe-atlas-sagan.readthedocs.io/en/latest/

failed/successful traceroute measurements with path reroutes (`Rerouted (Dst OK)`/ `Rerouted (Dst Fails)`) and without reroutes (`Dst OK`/ `Dst Fails`). We plot in Figure 5 the evolution of each of these classes for the considered events. We color-code each type of traceroute and mark with dashed lines the start and end of each incident. For each incident, we observe periods of time with significant increases in failed measurements, which we highlight with yellow and further refer as critical periods. Note that, for Verizon leak, we consider only failed measurements and filter out src/dst pairs traceroutes that were failing prior to the incident. Most of the selected traceroutes successfully reach their destination without any reroute, meaning that these traceroutes cross parts of the topology that are not impacted by the respective incident. However, 20% and 10% of the CenturyLink and LINX successful traceroutes, respectively, appear to be rerouted after the start of the incidents indicating the impact of these outages on path stability. Moreover, a small percentage of our selected traceroutes fail to reach their destination. For Verizon and LINX, rerouting affected 16% of the failed traceroutes, while this ratio increased to 58% for CenturyLink failed traceroutes.

Focusing on the performance impact evolution, we find that most of the failed measurements are limited to the incident periods. Moreover, a significant part of the traceroutes affected by the outages are rerouted after the incident starts. The number of the rerouted traceroutes, however, decreases progressively to further stabilize towards the end of the incident.

To understand the network recovery time, we focus on the critical period. Thus, we group our selected measurements per AS source/destination (src/dst) pairs and consider one such pair successful if all the traceroutes that start and finish within the AS source and AS destination, respectively, are successful. A total of 8% of the AS pairs failed during the critical period of the CenturyLink outage, while 80% of the 27K ASes recovered within 15 minutes. A closer inspection of the fast-recovering pairs revealed that src/dst ASes pairs are dominated by direct customers (44.4%) and customers of customers (51.5%) of the incident ASes. For *LINX outage*, approximately 85% of ASes src/dst pairs recovered shortly. However, we found that 83% of src/dst pairs successfully rerouted around LINX LON1 4 hours later as recorded in the LINX incident log, which also reported a significant traffic loss [17] at the same time. Most of these ASes are customers of customers.

C. Small ASes Failure

We investigate which networks are prone to failures due to BGP incidents from the data plane perspective. Specifically, we seek to understand whether the network class influences the failure rate of a network. To reduce sporadic measurements, we impose a minimum of two measurements per AS. We consider *stable*, all measurements where we do not observe any rerouting (`Dst OK` traces) and label as *affected* all the other collected traces. Hence, we compute the ratio of *affected* versus *stable* measurements per AS during the critical period.

We generate the distribution of the affected measurements for each of the three incidents. While we considered only failed traceroutes for *Verizon leak*, 94% of the 340 ASes have contributed for less than 5/1000 of the traces. Using this ratio, we divide the ASes into *less* and *highly* affected during the leak. We observed failed measurement for 2K and 1K ASes during the *CenturyLink and LINX outage*, respectively. Compared with the Verizon leak, we found a higher ASes failure rate variability. Using the affected ratio value distribution for these two events, we derive the failure thresholds as 0.4 and 0.5 for CenturyLink and LINX outage, respectively. Networks with a failure rate value lower/higher than these values are considered *less/highly* affected.

V. AS-dependent Vulnerability

A. Variability of AS Inter-connectivity

We further analyze the long-term effect of such incidents in terms of resilience. Thus, we evaluate the evolution of the AS hegemony [24] for AS pairs that are connected through the incident AS. Intuitively, the hegemony score measures how networks depend on other networks on the Internet. Thus, this metric should capture the paths diversity changes between ASes pairs over time. We collected publicly available local AS hegemony scores for each AS towards the incident ASes on the following days: on the incident day, one day before and after the incident, one and three months after the incident.

We extract the daily median values of hegemony scores by AS pairs and plot in Figure 6 the distribution per day. The hegemony scores for indirect and direct customers (Cc and Cd) appear to be relatively consistent per AS categories. We note high variability in the incident day for direct peers. This observation is valid during the *Verizon leak*. However, we note a significant change in hegemony scores for the incident ASes during *CenturyLink outage*. Direct peers (Pd) exhibited the highest variability. Moreover, the median hegemony scores for these networks are still slightly increased after the incident. For the *LINX outage* we compute the median hegemony scores towards LON1 members. Contrary to the previous incidents, the scores were relatively consistent and low across different network classes. Members (M) scores are less than 0.1, which is most likely due to the high diversity of member types and rerouting policies through other IXPs [15].

B. AS Vulnerability

For small ASes, we seek to understand whether their impact is influence by other networks that appear on the path towards the incident. Recall that a successful trace can still traverse both *high (H)* and *less (L)* affected networks. Considering only the classified intermediate ASes, we identify the following AS vulnerability relationships: *HH*, *LL*, *LH* and *HL*. Intuitively, we expect to observe mostly the former. After extracting the AS vulnerability relationships for each incident, we compute the contribution of intermediate AS vulnerability class to each trace outcome and plot in figure 7 these values. *HH* contribute the most to failed measurements for *Verizon leak* and successful traces for both *CenturyLink* and *LINX* outages.

For the *Verizon leak*, the root cause of the failed *HH* relationship is linked to one network which in turn relied

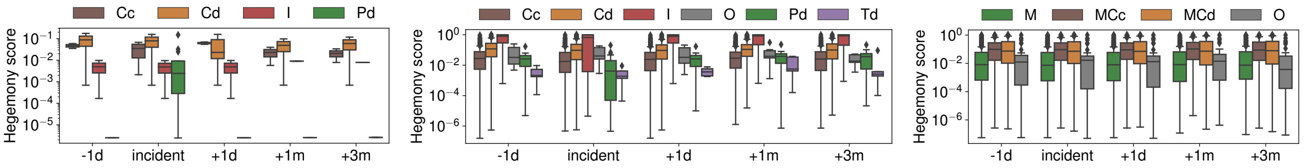

(a) Verizon leak (b) CenturyLink outage (c) LINX outage

Fig. 6: Variability of hegemony scores (log scale) per incident and the selected days: day before (-1d), day after (+1d), month after (+1m) and 3 months after (+3m). Direct peer hegemony scores show a decline in Verizon and CenturyLink incidents.

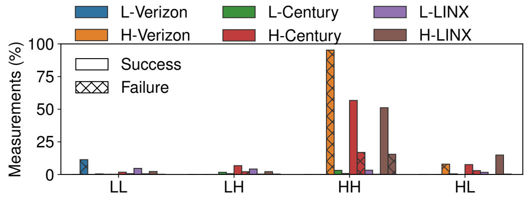

Fig. 7: Contribution of consecutive intermediate ASes classes to Cc (MCc) ASes traces outcomes.

on one upstream. For *CenturyLink outage*, 90% of Cc ASes were highly affected (H-Century) and handled 65% of successful traces. Similarly, 74% of the 152 indirect customers highly affected by the *LINX outage* handle 62% of successful measurements. The contribution of *HH* to successful traces is correlated with the source upstreams diversity and topological proximity to the incident.

VI. Conclusions

Internet resilience has been subject to numerous studies, but little work focus on evaluating the cascading effect of Internet incidents. In this work, we analyze how three major Internet disruptions impact its resilience. Our results indicate that in these cases, networks start to rely on backup and unseen/new links. As a result, a large proportion of the affected ASes use multihoming to recover shortly after experiencing network instability. However, we also observe that this robustness varies across networks. Hence, small networks are more likely to fail, regardless of the type of incident, due to their dependency chain to major ASes. Furthermore, the metric for studying inter-dependability networks directly depend on whether these networks peer with the disrupted network. For small ASes, this control plane-based metric remains mostly consistent before and after the incidents, confirming the need for additional data plane measurements. Indeed, the control plane has partial visibility and is limited to gauge the magnitude of data plane reachability. Therefore, we recommend that a) small ASes should participate in the collaborative MANRS project and b) MANRS could integrate additional metric based on data plane measurements to strengthen global routing security.

VII. Acknowledgements

We would like to thank Romain Fontugne and the anonymous reviewers for their insightful comments. This work was funded by the Norwegian Research Council grant # 288744.

References

[1] M. Hicks, “Seven Outages That Shook Up 2021.” [Online]. Available: https://dcweb.thousandeyes.com/blog/seven-outages-shook-up-2021

[2] A. Medina, “Looking Back at the Most Disruptive Internet Outages of 2020.” [Online]. Available: https://www.thousandeyes.com/blog/most-disruptive-internet-outages-2020

[3] M. Prince, “August 30th 2020: Analysis of CenturyLink/Level(3) Outage,” 2020. [Online]. Available: https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage/

[4] D. Belson, “Internet disruptions overview for Q2 2022.” [Online]. Available: https://blog.cloudflare.com/q2-2022-internet-disruption-summary/

[5] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, “Artemis: Neutralizing bgp hijacking within a minute,” *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, dec 2018.

[6] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, “ispy: Detecting ip prefix hijacking on my own,” *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1815–1828, 2010.

[7] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, “Detecting prefix hijackings in the internet with argus,” in *ACM IMC*, 2012.

[8] L. Quan, J. Heidemann, and Y. Pradkin, “Trinocular: Understanding internet reliability through adaptive probing,” *SIGCOMM CCR*, 2013.

[9] A. Schulman and N. Spring, “Pingin’ in the rain,” in *ACM IMC*, 2011.

[10] R. Padmanabhan, A. Schulman, D. Levin, and N. Spring, “Residential links under the weather,” in *ACM SIGCOMM*, 2019.

[11] M. Luckie and R. Beverly, “The impact of router outages on the as-level internet,” in *ACM SIGCOMM*, 2017.

[12] G. Baltra and J. Heidemann, “Improving the optics of active outage detection (extended),” USC/ISI, Tech. Rep., 2019.

[13] P. Richter, R. Padmanabhan, N. Spring, A. Berger, and D. Clark, “Advancing the art of internet edge outage detection,” in *ACM IMC*, 2018.

[14] T. Strickx, “How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today,” 2019. [Online]. Available: https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-today/

[15] E. Aben, “Does The Internet Route Around Damage? - Edition 2021,” 2021. [Online]. Available: https://labs.ripe.net/author/emileaben/does-the-internet-route-around-damage-edition-2021/

[16] P. R. Marques, J. Mauch, N. Sheth, B. Greene, R. Raszuk, and D. R. McPherson, “Dissemination of Flow Specification Rules,” RFC 5575, Aug. 2009. [Online]. Available: https://www.rfc-editor.org/info/rfc5575

[17] “Incidents Log,” 2021. [Online]. Available: https://www.linx.net/incidents-log/

[18] “CAIDA AS Rank,” 2022. [Online]. Available: http://as-rank.caida.org/

[19] “The CAIDA UCSD IXPs Dataset, 202101,” 2022. [Online]. Available: https://www.caida.org/catalog/datasets/ixps

[20] “The CAIDA AS Relationships Dataset,” 2022. [Online]. Available: https://www.caida.org/catalog/datasets/as-relationships/

[21] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, “Bgpstream: a software framework for live and historical bgp data analysis,” in *ACM IMC*, 2016.

[22] “Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6 ,” 2022. [Online]. Available: https://www.caida.org/catalog/datasets/routeviews-prefix2as/

[23] “Ripe atlas,” 2021. [Online]. Available: https://atlas.ripe.net/

[24] R. Fontugne, A. Shah, and E. Aben, “The (thin) bridges of as connectivity: Measuring dependency using as hegemony,” in *PAM*, 2018.