Detecting Peering Infrastructure Outages in the Wild

Vasileios Giotsas
CAIDA/TU Berlin
vgiotsas@ucsd.edu

Christoph Dietzel
TU Berlin/DE-CIX
christoph@inet.tu-berlin.de

Georgios Smaragdakis
MIT/TU Berlin
gsmaragd@csail.mit.edu

Anja Feldmann
TU Berlin
anja@inet.tu-berlin.de

Arthur Berger
MIT/Akamai
awberger@csail.mit.edu

Emile Aben
RIPE NCC
emile.aben@ripe.net

ABSTRACT

Peering infrastructures, namely, colocation facilities and Internet exchange points, are located in every major city, have hundreds of network members, and support hundreds of thousands of interconnections around the globe. These infrastructures are well provisioned and managed, but outages have to be expected, e.g., due to power failures, human errors, attacks, and natural disasters. However, little is known about the *frequency* and *impact* of outages at these critical infrastructures with high peering concentration.

In this paper, we develop a novel and lightweight methodology for detecting peering infrastructure outages. Our methodology relies on the observation that BGP communities, announced with routing updates, are an excellent and yet unexplored source of information allowing us to pinpoint outage locations with high accuracy. We build and operate a system that can locate the epicenter of infrastructure outages at the level of a building and track the reaction of networks in near real-time. Our analysis unveils four times as many outages as compared to those publicly reported over the past five years. Moreover, we show that such outages have significant impact on remote networks and peering infrastructures. Our study provides a unique view of the Internet's behavior under stress that often goes unreported.

CCS CONCEPTS

• **Networks** → **Network components**; **Network measurement**; **Network structure**;

KEYWORDS

Outages, Colocation, Interconnection Facility, IXP, Peering, BGP Community, Resilience.

1 INTRODUCTION

Today, our economy as well as our social life, rely on the smooth and uninterrupted operation of the Internet. While the Internet has shown an amazing resilience as a whole, even short outages can have a significant impact on a subset of the Internet user population. Past major Internet outages have been studied in depth, including outages due to network component failure, e.g., hardware, software, and configuration failures in routers [98], optical layer outages [47], natural disasters [20, 23, 35, 56, 84], and nation-wide censorship [23, 24, 83]. Most of these events affected either individual networks or entire regions. This can be attributed to the fact that the Internet's architecture used to be quite hierarchical. Thus, most *local outages* were expected to have a *local impact*.

During recent years the Internet infrastructure has changed significantly, a phenomenon that is referred to as the "flattening" of the Internet's hierarchy. In this setting, the majority of Internet inter-domain traffic flows *directly* between edge networks, bypassing transit providers [62]. For example, eyeball networks reduce their transit costs and improve end-to-end performance [41, 49] by directly peering with content providers, content distribution networks, and cloud providers, which are now a major source of traffic [32, 46, 82]. Direct peering is enabled by *third party peering infrastructures* (also referred as carrier-neutral peering infrastructures), such as *colocation facilities* and *Internet Exchange Points* (IXPs). These infrastructures are increasingly deployed in cities around the globe [50], and their members are growing constantly [61, 68] supporting hundreds of thousands of peerings [100].

Given the high concentration of peerings established at colocation facilities and via IXPs, many government bodies consider them critical infrastructures [30, 39, 64, 96]. Unfortunately, little is known about outages at these peering infrastructures, i.e., outages due to interruption, misconfiguration, and failure in the supply of power, the hardware, or the software that supports the operation of the peering facility. Such outages are affecting multiple networks, thus have different characteristics than those due to faulty operation or failure of an individual router or a single network provider. To the best of our knowledge, the only detailed study of such an outage is about the World Trade Center after the September 11 attack [13]. The report concludes that the catastrophic failure had "little effect on the Internet as a whole" but "a major effect on the services offered by some information and service providers". However, these infrastructures have gained an increasingly international set of network members in the last 15 years [16, 18]. Thus, it is quite possible that a *local outage* at one of these infrastructures today has a more *global effect*.

Figure 1: Detected and reported infrastructure outages per semester since 2012. The peak in the 2012/12 bin is due to Hurricane Sandy.

Unfortunately, a system that can detect and report on peering infrastructure outages in an automated fashion is not available. Such a system would be of increasing interest for many Internet stakeholders. Network operators can be informed, in real-time, about ongoing outages, which today mainly happens via out-of-band communication after the event, if at all. Timely detection of outages based on routing data can help operators optimize their mitigation strategies and the communication with their customers. Policy makers can make use of such a system to improve their situational awareness regarding the threats to critical infrastructures. Finally, researchers can understand how the evolving Internet behaves under stress. To enable the above capabilities, we build *Kepler*, a system that detects peering infrastructure outages with the aim to understand the externalities of such outages, improve current monitoring practices, and potentially help in improving the resilience of the Internet at the regional and global level.

By extracting location meta-data encoded in BGP messages, *Kepler* can detect 159 facilities and IXP outages over the last 5 years, four times more than publicly reported in popular operators and outage mailing lists [25, 26, 67, 74]. Figure 1 shows the number of facility and IXP outages we detect per 6-months since 2012, compared to the number of facility and IXP outages reported. Surprisingly, even infrastructure outages with large effects are not necessarily communicated via these mailing lists.[1] One alternative communication channel of outage events is through social media, where operators often resort to seeking answers on network disruptions. However, extracting this information remains a manual and error-prone search process [9].

To develop *Kepler* we have to tackle the following challenges:
Identify Outages: How to detect outages at peering infrastructures given that previous work has illustrated that even identifying the AS responsible for major routing events is a challenging task [42, 94, 99].
Characterize Outages: The next challenge is to assess the *start*, *duration*, *impact*, and *frequency* of an outage. Often, public information, such as press releases after an outage, are of questionable accuracy and detail, and there is limited transparency on what actually happened and which parts of the Internet were affected.
Locate Outages: The third challenge is to detect the *exact location of an outage*. While a map of the U.S. long-haul fiber-optic

infrastructure including some of the carrier facilities of major U.S. ISPs was released last year [34], we lack a detailed map of peering infrastructures. Two recent works attempt to tackle this problem by using large-scale active traceroute campaigns to infer the IP-level connectivity at colocation facilities [50, 72]. However, these methods scale only for a limited number of ASes or a limited number of facilities. This is due to the scale of the required active queries and the resource limitations of the available measurement platforms, such as RIPE Atlas and Looking Glasses [48, 91].

Our Approach: We introduce a novel methodology to reliably detect peering infrastructure outages in the wild and investigate their impact. Our detection mechanism relies on the observation that BGP is no longer purely an "information hiding protocol" [92]. The *BGP Communities attribute*, introduced with RFC1997 [17] in 1996, provides meta-information regarding prefixes announced to customer and peer networks, and is used for traffic engineering [85], traffic blackholing to mitigate attacks [31], and network troubleshooting [44]. Their use has become quite popular in recent years (Section 3.2) allowing us to use them as a *crowd-sourcing mechanism* for acquiring accurate location information for about 50% of all BGP IPv4 updates (Section 5.2).

While BGP routing updates have been used to detect outages limited to the AS and prefix granularities [8, 20, 24, 60], our novel insight is that *Communities with location information* in BGP updates can reveal the occurrence and location of peering infrastructure outages. Our methodology relies on location-based BGP Community values and allows us to pinpoint the exact location as well as the starting time and duration of the outage at high accuracy. To assess the impact of an outage, we track the changes in the use of the Communities by the members of the affected facility. However, since the semantics of the Communitiy attributes vary in geolocation granularity, from facility or IXP to city or metropolitan area, and Communities are not attached in every BGP update, monitoring Communities alone is not sufficient. To address these limitations, we augment our analysis with a physical map of facilities which allows us to correlate location-specific routing changes with the colocation of ASes in common peering infrastructures (Section 3.3). Moreover, we use archived and a small number of targeted traceroute measurements to confirm our inferences (Section 6.3).

In summary, our contributions are the following:
- A novel lightweight methodology for detecting, localizing, and tracking outages at peering infrastructures through passive monitoring of BGP data, by combining location-tagging BGP Communities with colocation data in facilities and IXPs.
- We instantiate our methodology in an operational monitoring system, *Kepler*, and we use it to study infrastructure outages visible in public BGP data between 2012 – 2016. We unveil four times as many outages at major peering infrastructures as compared to those previously reported in major networking mailing lists and news websites.
- We augment our analysis with targeted and archived traceroute measurements, and traffic data to further investigate the impact of the detected outages. We find that a large number of the affected links with remote networks can be hundreds or even thousands of miles away from the location of the incident, challenging the mental model that local outages have only local impact. Our study

---

[1]For example, the May 2015 outage at AMS-IX was discussed in the Austrian ATNOG mailing list [4] but not the more popular NANOG and outages mailing lists.

reveals that interconnection strategies such as remote peering and the colocation of ASes at multiple diverse locations create unexpected interdependencies among peering infrastructures that remain largely unnoticeable during normal operation, but disrupt connectivity in counter-intuitive ways during outages.

The rest of the paper is organized as follows. Section 2 discusses the changing interconnection landscape. Section 3 introduces our methodology and the datasets we compile to make it feasible. Section 4 explains how we develop *Kepler* to implement the proposed methodology, which we evaluate in Section 5. Finally, Sections 7 and 8 discuss the implications of our work and summarize our contributions respectively.

2 BACKGROUND

Networks often interconnect through multiple *physical links* established over peering facilities, sometimes even in different locations in the same city [73, 92]. While in the past the majority of facilities were maintained by individual transit providers to interconnect with their customers, the advent of IXPs and the flattening of the Internet hierarchy led to the increasing popularity of *carrier-neutral* facilities, such as colocation facilities, which allow connectivity independent of specific providers [54, 70].

Colocation facilities offer the hosting of servers and network equipment to facilitate networks' interconnections, typically via *cross-connects* or Private Network Interconnects (PNI), i.e., a point-to-point circuit [12]. Facilities are mainly concentrated in metropolitan areas, with major telecommunication hubs like London and New York hosting dozens of facilities [50]. While it is common practice among facility operators not to publish the number of PNIs, there are indications that their number is continuously growing. Equinix reports more than 188K cross-connects over its 145 facilities (Q3/2016) [37]. Moreover, high-profile acquisitions suggest a highly dynamic sector, including the acquisition of Telecity by Equinix for $3.8 Billion [36], and Telx by Digital Reality for $1.9 Billion [97]. Interconnection paradigms such as remote peering and tethering are increasingly deployed, allowing networks in remote sites of the same facility to exchange traffic directly [77].

An IXP is a physical infrastructure composed of layer-2 Ethernet switches which interconnect edge routers of members [18]. Once a physical connection is established, ASes can chose between different flavors of peering: (*i*) bilateral public peering, (*ii*) bilateral private peering via a virtual local network, similar to PNIs in colocation facilities, (*iii*) multilateral public peering over IXP route servers [52, 89], or (*iv*) remote peering with the members of affiliated IXPs [16]. Today, there are more than 300 IXPs in the world [81], particularly in Europe, but their popularity also increases in other regions, including the USA [61], Latin America [11], and Africa [40]. The number of members varies from tens to multiple hundreds, e.g., DE-CIX Frankfurt and AMS-IX Amsterdam have over 700 members [2, 28]. Moreover, IXPs are not just local interconnection points but they are becoming international hubs, through the use of layer-2 carriers and Virtual PoPs (vPoPs). For instance, LINX London interconnects networks from more than 72 countries [65, 66]. It is also increasingly popular for IXPs to form conglomerates by interconnecting with each other [45], while distributed IXPs, such as NL-IX, interconnect their remote sites to offer virtual backbone and

remote access to their network members. Studies show that IXPs enable hundreds of thousands of peerings [1], the large majority being multi-lateral peerings [52, 89]. Traffic exchanged at IXPs has increased significantly in recent years [18], exceeding 5 Tbps at large IXPs.

With the advent of Content Distribution Networks (CDNs) and the placement of data caches close to the users, the interconnection landscape has become increasingly clustered in large metropolitan hubs [50, 70]. The geographic agglomeration of the peering activity has led to an increasingly *symbiotic* relationship between IXPs and colocation facilities: IXPs benefit from placing their switches in locations where ISPs can easily install their network equipment, while facility operators often subsidize the presence of IXPs in their space to increase the attractiveness of their colocation ecosystem [12, 78]. These mutual interconnection incentives create tight physical interdependencies between IXPs and facilities. For example, DE-CIX has distributed its peering fabric among 12 different facilities in the greater Frankfurt metropolitan area [29], while the Equinix Frankfurt KleyerStrasse (FR5) colocation facility hosts 10 different IXPs [81].

3 METHODOLOGY

In this section, we describe our methodology for detecting and localizing peering infrastructure outages.

3.1 Challenges and Concept

Recall that the main purpose of BGP is to provide *reachability* information and not *connectivity* information [92]. Thus, relying on the BGP path or the AS-level topology of the Internet is not sufficient to detect the physical location of a peering, and the location of the underlay interconnection infrastructure. To illustrate the challenges in detecting and pinpointing the exact *physical* location of a peering outage consider the topology of Figure 2. It consists of four ASes (AS_i), four colocation facilities (F_j), and two IXPs (IX_k). Figure 2(b) and 2(c) are the results of two different outages, at colocation facility F_2 and at IXP IX_1, respectively. Initially, AS_1 reaches AS_2 via private peering over facility F_2; AS_2 reaches AS_4 via public peering over the IXP IX_1; and AS_3 reaches AS_4 via IX_1. Note, some paths involve multiple facilities, e.g., from AS_2 to AS_4 via IXP IX_1, F_2, and F_4, and from AS_3 to AS_4 via IX_1, F_3, and F_4.

The failure of F_2, Figure 2(b), affects both private and public interconnections at this facility. The private ones are affected directly, the public ones only indirectly since F_2 hosts part of IXP IX_1's switching fabric. In our example, two paths change: AS_1 switches to its backup path via F_1, and AS_2 switches to its backup path to AS_4 over F_4. Note that the AS paths do not change. However, the involved facilities and IXPs do change. Likewise, the failure of IX_1, (Figure 2(c)), partially affects the paths of F_2, F_3, and F_4, since the new routes have to bypass IX_1. This can cause a large number of BGP updates. Yet, the AS paths themselves again do not necessarily change. Both scenarios illustrate the increasing symbiotic relationships between colocation and IXP peering infrastructures. Such inter-dependencies have already led to confusion when locating and reporting the cause of outages [3, 87].

Our examples show that it is not sufficient to track AS-level changes to determine the outage location, we need to monitor
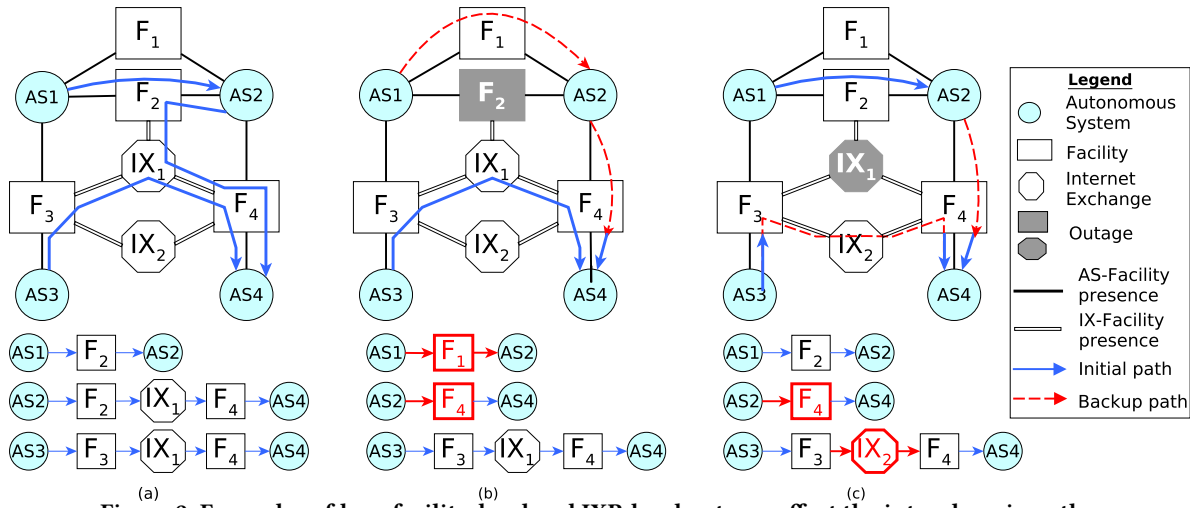
Figure 2: Examples of how facility-level and IXP-level outages affect the inter-domain paths.

facility-level paths and correlate them across *multiple route changes*. In Figure 2(b), the fact that F_2 disappears from all paths, while IX_1 disappears only from the path through F_2, is sufficient to infer that the outage occurred at F_2. Similarly, for Figure 2(c) the outage can be localized at IX_1 and not F_1, since the AS_1–AS_2 path through the facilities/IXP remains unchanged, while the AS_3–AS_4 path is re-routed via IX_2 concurrently with a path change from AS_2 to AS_4.

The example above allows us to derive the following insights about infrastructure-level outage detection:

Facility-level Inter-domain Hops: The four ASes appear to exchange traffic directly when observing only the AS-level paths. However, the physical paths involve multiple intermediate facility-level and IXP-level infrastructures that introduce externalities in the resilience of the AS interconnections. We need to capture these infrastructures to accurately localize outages.

Path Correlation: To uncover the failure location within the complex infrastructure of today's Internet, we have to correlate path changes across multiple vantage points with colocation data at facilities and IXPs.

Before and After Comparison: To understand the source and impact of an outage, one needs to compare routes during an outage to those before the outage—the "healthy" state. Therefore, we need the ability to continuously monitor the routing system.

A major challenge is how to get sufficiently fine-grained facility information. A key insight of our approach is that we can extract facility information per routing update through the analysis of BGP communities. Moreover, it is feasible to collect detailed facility maps from various public sources using techniques described in [50, 68], thanks to the increasing openness in the sharing of colocation data to support a more flexible peering setup process or even automate it altogether [7, 63]. Indeed, today the large majority of peerings are multilateral peerings that do not involve formal contractual agreements [100].

3.2 BGP Community Dictionary

BGP Communities have the format X:Y, where X, Y are two 16-bit values (extended communities use four octets [93]). By convention, the first two octets encode the ASN of the operator that sets the community, while the next two octets encode an arbitrary value that is used by the operator to denote specific information such as the ingress location of a route. There are two types of communities: (i) *inbound communities* that are applied when an operator receives a prefix advertisement at an ingress peering point, and (ii) *outbound communities* that are applied when an operator sends a prefix advertisement at an egress peering point.

The Rise of BGP Communities: Between 2010 and 2016 the visible number of networks using BGP Communities has more than doubled from 2,500 to 5,500, and the number of unique community values has tripled to more than 50K in 2016 (Figure 3). Moreover, the number of Community values per prefix announcement has increased from an average of 4 to 16. These communities encode a wealth of routing meta-data, but unfortunately, the community is possibly the only BGP attribute with no specific semantics and values that are neither standardized nor have a uniform encoding [33]. Consequently, extracting meaningful information from the communities is not possible without additional sources of interpretation.

Location-Encoding Ingress Communities: Each operator uses different values to encode *location* information at various granularities. For example, in Figure 4 the BGP collector receives routes for prefixes 184.84.242.0/24 and 2.21.67.0/24 with a common AS subpath 13030 20940. The first route is tagged with community 13030:51904. The value 13030 in the top 16 bits indicates that AS13030 has applied the community. The value 51904 in the bottom 16 bits, indicates that this community is used to tag routes received at the Coresite LAX-1 facility [59]. Similarly, the second route is tagged with two communities from AS13030. The value 51702 means that the route's ingress point was the Telehouse East London facility, and the value 4006 means that the route was received by a public peer at the LINX IXP Juniper LAN.

Figure 3: Number of unique BGP Communities values (left y-axis), compared to unique top two octets.

Figure 4: Inferring physical locations from BGP Communities.

While the community values are not standardized, many operators publicly document their community schemes either in their Internet Routing Registry (IRR) records or in their support Web pages. However, the documentation is in natural text and lacks a standardized structure and terminology, therefore its parsing necessitates significant manual work that is unsustainable given the large number of BGP Communities. To tackle this problem we develop a web-mining tool that enables the automatic compilation of a community dictionary. We first use a Web Scraper to extract the text from the *remarks* sections of IRR records and from ASes' web pages. Then, a text parser analyzes the extracted text using the Natural Language ToolKit [10] to discover infrastructure-related communities. We identify sub-strings that include community values using regular expression matching, on which we use Stanford's Named Entity Recognizer (NER) [43] to identify named entities, focusing on entities that pertain to locations or infrastructure operators. To improve the accuracy of NER for network-related entities, we adopt the techniques proposed by Banerjee et al. [5] and we search PeeringDB [81], Euro-IX [38], and IRR records, for organization names that match capitalized words encountered in communities documentation. These sources also enable us to determine the network type of the identified entities. For our community dictionary, we only keep communities that tag three types of Named Entities: (i) city-level locations, (ii) IXPs, and (iii) colocation facilities.

Then, using syntactic analysis we filter-out outbound communities that define location-specific traffic engineering actions. In particular, we perform Part-of-Speech tagging to distinguish verbs in passive voice used for documenting inbound communities (e.g., "received", "learned", "exchanged"), and ones in active voice that define actions (e.g., "announce", "block"). Finally, we assign a single location identifier to all entities related to a common location. Different operators use different naming, such as city names ("New York City"), city initials ("NYC"), or IATA airport codes ("JFK"). To determine if the different location identifiers refer to the same location we query the Google Maps Geocoding API [53] to obtain the coordinates for each identifier, and we group together identifiers that are within 10 km from each other.

IXP Path Redistribution Communities: We augment our dictionary with path redistribution communities used by IXP route servers. IXP route servers often use communities to aid their members in controlling how their prefixes are advertised to other route server members [57], e.g., *advertise to all*, and *advertise to selected peers*. Thus, a route server community on a BGP route indicates that
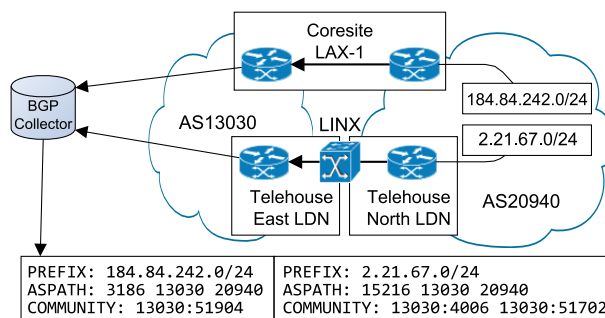
the route traversed the IXP and the first 16 bits of the community value indicates the IXP ASN.

Dictionary Statistics: As of December 2016, our community dictionary includes 5,284 communities by 468 ASes and 48 route servers, and covers 288 cities in 72 countries, 172 IXPs, and 103 facilities. While 468 ASes is a small fraction of the ASes, it includes all but two Tier-1 ASes and most major peering ASes. Note that for the two Tier-1 ASes (XO Communications and Verizon) missing from our dictionary we observed less than 20 different community values in the public BGP data, which indicates that they either do not use communities to annotate their PoPs, or they do not propagate such communities outside their domain and do not provide publicly accessible community documentations. Figure 5 shows the geographical coverage of locations we extract from the communities. The majority of the communities (66%) tag a location in Europe, followed by North America (24.5%), while only 2% of the communities cover locations in Africa and South America. Although the interconnection ecosystem in these regions is indeed relatively underdeveloped [55, 71], the difference in coverage can be also explained by biases in the underlay documentation sources, such as the completeness of the different Internet Routing Registries [6], and the fact that our natural language parser works only with English text. As we elaborate in Section 5.2, location BGP Communities included in our dictionary are present in about half of all BGP IPv4 updates. To ensure freshness we recompute our dictionary every two weeks and always use the dictionary from the corresponding time period for route processing. To validate the correctness of our automatically-generated community dictionary, we compared it against a manually-constructed dictionary. Due to the overhead of manually parsing community documentations, we limited the validation to the 25 ASes in our dictionary with the highest number of BGP paths annotated. We did neither find a false positive nor a false negative.

Attrition of BGP Communities: To understand the attrition rate of location-encoding communities we study the communities classified either as "geographical location" or as "interconnection point" by Donnet and Bonaventure in 2008 [33]. Only 552 of the 2,980 communities in their dictionary are visible in the aggregated Route-Views/RIS BGP data across 2016, while the rest appear not to be used anymore. On the other hand, of the 5,284 communities in our dictionary, only 471 (9%) are also in the 2008 dictionary. However, only 7 (1.5%) of the common community values changed meaning after almost a decade, indicating that the semantics of communities

Figure 5: The geographic spread of trackable infrastructure.

within an AS change rarely. Since location-encoding communities are used for operational purposes, such as troubleshooting and traffic engineering, the stability of community semantics minimizes the risk of misconfigurations when setting these communities on prefix advertisements.

The above findings highlight the value of our automated community interpretation to enable a frequent extension of the community dictionary with new values, to remove stale entries, and to maintain a high-degree of coverage of the active communities. Moreover, the risk of misinterpreting the community values due to stale entries is small even in the time span of years.

3.3 Colocation Map

The majority of the communities annotate routes at city-level granularity, which is too coarse to pinpoint a peering infrastructure outage at the facility-level or IXP-level. To achieve the intended detection granularity, we complement the BGP communities with a high-resolution colocation map that includes three types of interconnections: (i) ASes to IXPs, (ii) ASes to facilities, and (iii) IXP to facilities. For each facility we also record the building-level address, so that we know which facilities, IXPs and ASes operate at the cities annotated by our community dictionary. To this end, we mine the colocation data from PeeringDB [81] and DataCenterMap [27], as well as individual AS websites. Since names of facilities and facility operators are not standardized, we use the facility address (postcode and country) to identify common facilities among the different data sources. We then merge the tenants listed in each data source for the same facility to increase the completeness of our colocation map. Similarly, IXP names also differ between datasets. To identify and merge the records that refer to the same IXP we use the URLs of the IXP websites, and the location (city/country) where the IXP operates. We use the constructed colocation map in the city-level outage signal arbitration to de-correlate the "fate" of various ASes in the same city during an incident, based on their presence or absence at facilities. Thus, we can pinpoint the likely facility-level or IXP-level location of incidents and increase the coverage of our outage detection capabilities to physical locations beyond those explicitly encoded in BGP communities.

3.4 Detection Methodology Overview

To detect and localize peering infrastructure outages we propose Algorithm 1. Its input is a stream of BGP data, the BGP Community dictionary, the colocation map, as well as targeted active measurements for incident investigation.

The first step is to parse the BGP Communities attribute of the collected BGP routes and find paths annotated with the traversed *Points-of-Presence* (PoPs). We use these paths to analyze the PoP-level routing dynamics. When we use the term "PoP" without any other qualification, we refer to any of city, IXP, or facility. We filter-out transient paths to ensure that we have a stable baseline of the routing system, and we update the set of stable paths periodically to account for path changes after the start of our detection process.

Next, we start monitoring the incoming BGP updates for PoP-level deviations from the stable baseline. Instead of checking for AS path changes, we check if the relevant community values change. When we observe a large enough fraction of paths that deviates from the baseline PoP within the same time frame, we call it *outage signal*. An outage signal corresponds to a spike in localized routing activity and indicates that a routing incident affected a specific PoP. Yet, it does not indicate if the incident is due to an outage.

Link-level events such as the de-peering of two large peers, or AS-level incidents such as the disconnection of an IXP member, can also lead to such an outage signal. To determine the source of the signal, we trigger a detailed *signal investigation* process that classifies the signal as link-level, AS-level, or PoP-level based on the number and disjointedness of the affected ASes.

If the signal is classified as a PoP-level outage, the algorithm proceeds to explore the granularity of the PoP. Here, we combine the colocation map with active traceroute measurements that we collect either opportunistically by mining public traceroute repositories, such as those provided by PathCache [95], or by executing our own targeted traceroute campaigns. The traceroute paths help us to validate the outage and eliminate false positives by mapping the IP-level hops to IXPs and facility interfaces using the techniques described in [50, 76]. When the data-plane and control-plane inference identify the same PoP as the source of the outage, we consider the outage as validated. We determine the length of the outage (i) by actively probing the involved interfaces and (ii) by monitoring BGP messages for changes in the communities that indicate that the paths have returned to the baseline PoP. Since we mainly rely on passive measurements via BGP, our active monitoring is rather selective and does not rely on greedily probing all infrastructure addresses. Therefore, our approach is practical and conforms to the resource limitations of publicly available measurement platforms, including RIPE Atlas [90] and Looking Glasses [48].

4 THE KEPLER SYSTEM

In this section, we present the design and implementation of *Kepler*[2], a system that relies on our methodology to detect outages in the wild and investigate them. While the analysis of BGP data is lightweight, our experience with operating *Kepler* shows that the efficient design of different modules is critical to make the system practical and accurate. Figure 6 illustrates the architecture of *Kepler*.

4.1 Input Module: Data Preprocessing

The first part of *Kepler* preprocesses all data sources. First, it generates the BGP Community dictionary and the colocation map. For the continuous BGP data we use BGPStream [79] to decouple *Kepler*

---

[2]Data and additional technical details are available at http://kepler.inet.tu-berlin.de

---

Algorithm 1: Overview outage detection and investigation

Input: (BGP paths, BGP Community Dictionary, Colocation Map, Targeted Active Measurements)

Output: Location, Time and Duration of a PoP-level Outage

$Paths_{mapped} \longleftarrow$ Map BGP paths to traversed PoPs based on the attached Communities meta-data;

$Paths_{mapped}^{stable} \longleftarrow$ Filter-out transient paths;

for *BGP updates in new measurement interval* **do**

    $Paths_{mapped}^{diverted} \longleftarrow$ calculate how many paths diverted from the PoP in the stable baseline;

    if $\frac{Paths_{mapped}^{diverted}}{Paths_{mapped}^{stable}} > T_{fail}$ **then**

        Signal investigation

        $Signal_{type} \longleftarrow$ Infer the type of outage signal based on the number of affected ASes and AS links;

        if $Signal_{type}$ *is PoP* **then**

            $POP_{type}^{BGP} \longleftarrow$ Determine the type of PoP based on the colocation map;

            $POP_{type}^{trace} \longleftarrow$ Confirm the affected PoP through traceroute queries;

            if $POP_{type}^{BGP} \equiv POP_{type}^{trace}$ **then**

                while $Outage_{state}$ *is True* **do**

                    $duration \longleftarrow$ record the duration of the outage

            return $Outage(time, POP, duration)$

---

from the sources of BGP feeds, and thus, obtain a unified feed of sorted BGP records. In addition, *Kepler* sanitizes the collected paths by discarding paths with AS loops, private ASNs, or special-purpose ASNs [22]. Currently, we use all RouteView and RIPE RIS collectors. For every BGP update with attached BGP community values, *Kepler* uses the dictionary to infer which physical infrastructure a route traverses. Hereby, *Kepler* also infers which location-based BGP community refers to which hop of the BGP path, either by mapping the first two octets of the community to the same ASN hop in the path, or by applying the methodology in [51] in the case of IXP route server communities.

4.2 Monitoring Module: Outage Detection

Kepler's monitoring module identifies all the PoPs P for which we have physical location information from the community dictionary. These are the PoPs that we monitor in detail. Then, *Kepler* periodically computes a set of stable routes that involve p for all $p \in P$. A prefix route is stable if it traverses P for a period of d_s consecutive days (the default value is 2 days). Thereafter, we check for PoP-level routing changes vs. the baseline stable path. Hereby, we consider the following change to a route from s to d involving PoP $p \in P$: (i) an *explicit withdrawal*, (ii) another AS path not involving PoP p, and (iii) an announcement with another community—an *implicit withdrawal*. In addition, we check for *BGP State messages* to detect potential disruptions in the BGP feed that can cause gaps in our BGP stream and disregard updates due to it. Note, if the AS path changes but the community tag involving p remains the same, we

do not consider the update a route change for p. However, we consider changes to the community tag as route change even if the AS path remains unchanged.

We bin routing updates in time intervals to correlate path changes with routing incidents. Since most of the ASes that set the ingress Communities are close to one of our BGP collectors it suffices to use a relatively short time interval. We use a binning interval of 60 seconds (twice the default MRAI time [88]). At the end of each binning interval we compare the paths from the baseline to the paths in the current bin and determine the fraction of paths that continues to traverse p. If this fraction is below a threshold of T_{fail} we may have an outage signal. However, an aggregated comparison of all the paths can be biased by ASes that account for a disproportionately large number of paths. For instance, if a partial outage in p affects the paths of many regional ASes but not the paths of a large Tier-1 AS, then the total fraction of paths may not fall below the detection threshold T_{fail} causing a false-negative. Therefore, we group the paths based on the ASes that are involved in the tagged links and determine outages per AS. If the fraction of paths of an AS a involving p falls below the threshold T, we say that a is subject to an outage signal in the current binning interval. An outage signal is an indicator of a *possible* outage event but the definite inference is a task of the signal investigation module. After each binning interval, we remove the changed paths from the set of stable paths. We also refresh the set of stable paths every 2 days to account for new paths and new community values. Note, the focus of this module is to detect the start of an outage.

4.3 Outage Signal Investigation

Kepler's outage signal investigation considers all outages signaled within a time interval and determines the granularity of the triggering event. We distinguish four incidents: (i) *link-level*, (ii) *AS-level*, (iii) *operator-level*, and (iv) *PoP-level* outages. For PoP-level events we identify the physical location. *Kepler* also *tracks* the new physical location after the rerouting of a stable path and the time it takes for a path to return back based on the same principle to estimate the duration of the outage. To increase the confidence for the duration of each outage and the reaction of network operators, *Kepler* relies on targeted active measurements.

We distinguish four different granularities of outage signals.
Link-level: Changes to an AS-link with a large number of prefixes, can cause an outage signal, e.g., a de-peering or even a MED change between two Tier-1 ASes. Since such link-level incidents are not the focus of this paper we require that more than three different ASes have to be affected to trigger an investigation.
AS-level: Changes in the availability of a densely connected AS can cause multiple of its peers to change their paths away from a specific location concurrently. For instance, if an IXP member decides to terminate its membership, it will terminate all public peering BGP sessions at that IXP. If all affected links intersect at a single common AS, either as near-end or as far-end neighbor, we classify the signal as AS-level.
Operator-level: We combine multiple AS-level outages to an operator-level outage, if all of the affected links include ASes that belong to the same operator. Our motivation is that operators often administer multiple sibling ASes each with different functions but
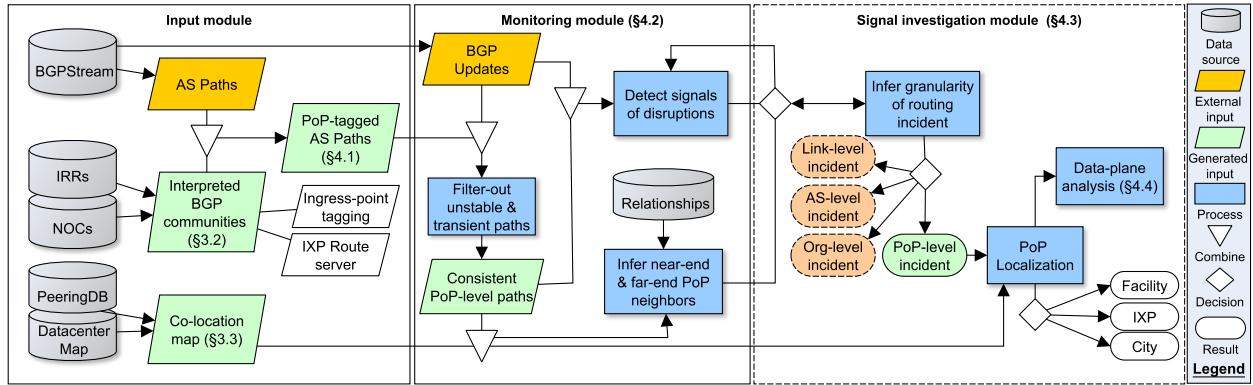
Figure 6: Flowchart of *Kepler*'s outage signal detection and investigation.

often hosted on the same infrastructures. For instance, the Equinix Ashburn Exchange hosts three different sibling ASes operated by Bell Canada. An organizational-wide policy or network change will effect all sibling ASes. We map ASes to organizations using the methodology from [14].

PoP-level: When a signal involves multiple AS links with disjoint near-end and far-end ASes and organizations, we classify it as PoP-level. In particular, we require that the set of affected links includes at least three different non-sibling near-end ASes and three-different non-sibling far-end ASes that are disjoint. From that, we infer a PoP-level incident if at least three different AS-level and operator-level incidents occur in the same binning interval at the same PoP. Next, we refine our localization for PoP-level outages.

Disambiguation of Outage Signals: Recall, from Figure 2, that the physical connectivity between two ASes can involve multiple physical PoPs. With ingress communities we can only identify PoPs at the *near-end* of an AS pair. However, depending on the peering strategy, which includes private peering and local or remote public peering, there may be up to four facilities between the ingress PoP and the far-end AS. A failure in any of them will trigger an outage signal at the near-end facility. To disambiguate such signals we correlate outage signals from multiple PoPs, combined with our colocation map. Our assumption is that outages at the near-end facility, the one identified by the ingress community of an AS, should affect all paths tagged with this community that involve links with far-end ASes co-located in the same facility. More specifically, we infer the outage in the near-end facility if at least 95% of the paths with co-located ASes are affected. We allow for a 5% margin to account for possible inaccuracies in the colocation map, such as spurious AS-to-facility presences, based on the results in [50].

If this is not the case, we check if the outage location is among the facilities where the affected far-end ASes have a presence. Accordingly, we repeat the above process for all facilities where any of the remote ASes has a presence and for which we have an outage signal in this binning interval. Figure 2(c) illustrates this process. When we infer that the near-end facility is not the outage epicenter, and the far-end peers have no facility in common (after checking the colocation map) we increase the PoP granularity to IXP-level and we repeat the same process. Namely, we collect the common IXPs among the near-end and the far-end peers and we check if all the common IXP members have been affected, e.g., in Figure 2(c)

the outage source is IX_1 and not F_3 or F_4. If we fail to converge to a single IXP as the outage source, we cannot make an inference and resort to targeted traceroute queries to discover the outage source. If during a binning interval we successfully converge to a facility/IXP for multiple outage signals, and all the facilities/IXPs operate in the same city, we abstract the granularity of the incident to city-level.

Increasing Signal Resolution: Unfortunately, communities are not always PoP specific but coarser, e.g., only at the IXP level (colocated IXP). To further refine our inference, we utilize again the colocation maps. For outage signals with IXP communities we check if all IXP peers or only IXP peers in specific facilities are not reachable. Thus, we check for each facility that the IXP is involved only if all routes of that facility are affected. If this is the case, we can infer that the outage is at the facility rather than the IXP, e.g., at F_2 and not IX_1 in Figure 2(b). We follow a similar methodology for outage signals detected using city-level communities, with the additional step of checking for IXP-level failures, if we infer that the outage did not occur in a facility.

4.4 Data-Plane Analysis

Kepler validates the occurrence and determines the outage duration via data-plane analysis, using both archived and targeted traceroute queries. We again initialize the analysis with a set of stable paths, whereby, we focus on paths that cross the monitored facilities and IXPs. To construct an extensive set of stable paths without incurring high measurement cost, we follow an approach similar to PathCache [95] and consume the publicly available traceroute paths collected by RIPE Atlas [90], CAIDA's Ark [15], and iplane [69]. *Kepler* also has an interface to initiate traceroute campaigns using public probing platforms [48, 90]. For mapping the traceroutes to ASes, IXPs, facilities, and data sanitization, we use techniques proposed in [19, 50, 76]. The facility mapping part is the only one that requires active measurement. To keep the number of required active measurements low, we focus on the ASes that are not covered by our community dictionary, yet are colocated at the facilities of interest.

Since we use opportunistic measurements for our baseline set of paths, we have to focus on subpaths. Namely, if an AS pair appears to consistently interconnect over the same IXP or facility hops in the traces of the last four consecutive weekly path dumps, we include

(a) Outages at PoP-levels with diff. thresholds.

(b) Detection sensitivity of our system.

(c) Fraction of IPv4 and IPv6 BGP path in 2016 with at least one location community.

Figure 7: Tuning parameters for *Kepler* (a) and (b), and Fraction of updates with location communities (c).

the corresponding paths in our baseline dataset. This approach may remove some AS pairs with very diverse interconnection footprint which is desirable for the purpose of confirming outages, since path changes between AS pairs with low path diversity are less likely to reflect intra-domain routing changes.

When *Kepler* detects an outage for a PoP, it identifies the baseline paths of AS pairs that interconnect over the PoP. Next, it selects the same sources and destinations and repeats the traceroute queries. If the fraction of baseline paths that continues to cross the PoP is below a threshold T_{fail}, we confirm the outage and continue probing to determine the duration of the outage. Otherwise, we either have a false-positive in our outage inference, or the service was restored in the mean time. Unfortunately, there is a 5 to 15 minute lag in receiving BGP updates. To eliminate false positives, we continue the traceroute analysis until the next set of BGP updates. If the outage signal is still in the BGP data, but the traceroutes did not confirm the outage, we conclude that we have a false-positive and disregard it.

When over 50% of the paths (traceroute if available/BGP otherwise) return to the baseline we consider the outage as restored. However, for a number of outages we observe periods of oscillations. When two consecutive outages for the same PoP are separated by less than 12 hours, we conclude that they are part of the same incident. Its downtime is the sum of the individual outage durations.

5 KEPLER EVALUATION

In this section we present a data-driven evaluation of *Kepler*'s capabilities. We first analyze the detection sensitivity of our algorithm, and how we tune *Kepler* to optimize the detection of PoP-level outages. We then discuss the reach of *Kepler* and its limitations, and we present our validation efforts to understand its accuracy and precision.

5.1 Sensitivity and Calibration

Kepler has two main parameters: (i) the time window for determining stable paths and (ii) the threshold which triggers an outage signal (T_{fail}). For the stable paths, a window smaller than 1 day would include transient paths, while windows higher than 5 days yield small sets of stable paths that restrict *Kepler*'s coverage. Therefore, we use two days to obtain a stable yet extensive baseline of paths. *Kepler* is more sensitive to the threshold parameter, as shown in Figure 7a. For 2016, it shows the number of detected outage signals

at link-level, AS-level, and facility/IXP-level for thresholds ranging from 2% to 50%. We assess the efficiency of the different threshold levels by validating the control-plane outage signals against the data-plane measurements for each signal. The number of detected facility/IXP-level outages, which is our focus, remains stable for thresholds from 2% to 15%. Higher thresholds lead to missing outage signals that have been confirmed by concurrent traceroute path changes. The missed outages are partial, i.e., outages limited to certain systems of a facility/IXP and affect a subset of its members. On the other hand, thresholds below 2% increase the number of outages that have to be investigated, and lead to mis-classification of AS-level and link-level outages as PoP-level. Note, that some of the additional outage signals raised for low thresholds may capture partial outages of limited impact that traceroute measurements fail to detect. We select a threshold of 10% to be relatively conservative and minimize wrong inferences, while still being able to capture medium-scale partial outages.

5.2 Data Analysis Reach and Coverage

A natural question is what fraction of BGP paths, can be analyzed with *Kepler*. Figure 7c shows the fraction of IPv4 and IPv6 BGP updates per month in 2016 with at least one location-encoding community. About 50% of the IPv4 and 30% of the IPv6 paths include such communities and, thus, are usable by *Kepler*. Moreover, *Kepler*'s communities consistently tag over 35% of the IPv4 and 28% of the IPv6 AS links across every BGP snapshot. One reason for the larger fraction of IPv4 paths/links compared to IPv6 is that ISPs still focus less on optimizing IPv6 traffic flows.

The next question is at what fraction of the facilities can *Kepler* uncover outages. We define a facility as *trackable* if it has a minimum number of networks whose interconnections can be located by the communities in *Kepler*'s dictionary so that our methodology is applicable. To distinguish PoP-level from AS-level or link-level incidents, we rely on correlation of updates from multiple members and we require that we have at least *six* different members that can be located through communities, 3 at the near-end of a link, and 3 at the far-end. The colocation databases we mined in Section 3.3 include 1, 742 facilities with at least one AS member. For each of the 1, 742 facilities, Figure 7b shows the total number of their members compared to the ones that are trackable. There are 1, 209 facilities with less than 6 members, thus, in principle, we can track 533 facilities. Of these we miss 130 (24%) for which we currently have

Table 1: Facilities coverage per continent

Continent	Facilities		
	All	>5 members	Trackable
Europe	878	305	243
North America	529	132	105
Asia/Pacific	233	70	46
South America	76	19	11
Africa	26	6	4

less than 6 trackable members. Therefore, the detected outages by *Kepler* are a *lower bound* of all possible outages. Note that while for the trackable facilities we are able to detect all full outages, it is possible that some partial outages may be undetected depending on the number of affected trackable facility members. Given the increase in the community usage and in member ASes we expect these numbers to increase over the next years. Importantly, we are able to cover 180 out of 183 (98%) facilities with at least 20 members which are the most prominent interconnection hubs.[3] Table 1 breaks down the covered facilities per continent. *Kepler*'s coverage is better for Europe and North America, while Africa and South America have the smallest fraction of trackable facilities. Note over 80% of all the facilities included in the colocation datasets (PeeringDB, DataCenterMap) are located either in Europe or in North America. While the low number of facilities in the other regions may indicate a geographical bias in the available colocation databases, the European and North American peering ecosystems are significantly more developed, with 73% of all the ASNs and 70% of all IPv4 addresses assigned to countries in the RIPE and ARIN zones.

5.3 Validation

We first check the accuracy and completeness of our PoP inference via communities, by obtaining ground-truth data of the facility-level interconnections from three large ISPs and one major CDN via private communication that use BGP location communities. Each gave us their list of facilities with neighbor ASes—in total location information for roughly 5K AS pairs. We find that our community-based localization is correct in every case, which is not surprising given the operational importance of communities. From Figure 8a, which plots the fraction of AS links vs. the number of facilities (the main plot is zoomed-in for AS links with more than 1 PoP), we see that we are missing less than 5% of the interconnections. On the side, we find that a large fraction of AS pairs only involves a single physical location. 60% are multilateral peerings between networks co-located at a single IXP, while the rest are interconnections between stub ASes and their transit providers. Still, a significant number of AS pairs involves many physical locations, in particular, if the two ASes are tier-1 or tier-2 ASes and peer with each other.

We then check the accuracy of *Kepler*'s inferences. We consider as *true-positive* any inferred outage for which we find an external data source that confirms the outage occurred in the same facility/IXP at the same time. Validating *false-positives*, i.e., inferred outages that did not happen, is more challenging since it is possible that an outage was not publicly reported, or that it was reported in a source that we could not discover. Nonetheless, we consider as

---

[3] Two of the non-trackable facilities with more than 20 members are in India and the other in Argentina.

false-positives incidents that happened in the same location/time as an inferred outage, but affected different infrastructures from the inferred one. We consider as *false-negative* any outage reported by an external data source which affected a trackable facility, but for which *Kepler* did not infer the outage. To collect validation data we parsed messages in the NANOG and Outages mailing lists [67, 74], news articles from specialized websites [25, 26], incident reports from 18 Network Operating Centers (NOCs) [75], as well as privately shared information. We successfully validated 53 out of the 159 detected outages (Figure 1) as true positives. We also found 6 cases of false positives. In these cases, *Kepler* determined the correct location of the incident but in reality the root cause of the outages were fiber cuts that affected multiple co-located ASes. In terms of false-negatives, *Kepler* did not miss any full outage that affected trackable facilities. However, we found 4 undetected small-scale partial outages that affected facilities with less than 30 tenants and were mis-classified as AS-level incidents.

6 RESULTS

Next, we use *Kepler* to detect and assess the impact of peering infrastructure outages during the past five years. To this end, we provide a detailed analysis of sample incidents enabled by *Kepler* and underpin our findings with active measurements to (*i*) confirm outages, (*ii*) track path changes, (*iii*) measure rerouting times and RTT increase, and (*iv*) infer the impact on traffic at seemingly unrelated locations.

6.1 Detected Facility Outages

The passive detection capabilities of *Kepler* allow us to conduct a historical analysis of archived BGP stream and PeeringDB data from 2012 to 2016. Overall, we detect 159 outages that include 103 outages among 87 facilities, and 56 outages in 41 IXPs, as shown in Figure 1. To contextualize the completeness of our findings we collect facility and IXP outages, reported in two popular mailing lists, NANOG [74] and Outages [67], plus two specialized data center and colocation websites [25, 26]. They only report 24% of the detected outages, missing most of the incidents that occur outside the US and the UK.

We find that 53% of the outages are in Europe, 31% in the US, and the remaining ones in the other regions. The median outage duration is 17 minutes and 40% of the outages exceed 1 hour (see Figure 8b). With regards to frequency we find that the number of detected outages is not increasing drastically over the last five years, see Figure 1. In general, we find that IXP outages last longer than facility outages. One reason may be in the possible causes of outages. Most facility outages are due to basic infrastructure failures, e.g., power or fiber cuts. Hence, restoring service mainly depends on infrastructure recovery. IXPs also suffer from software and/or configuration failures which apparently take longer to resolve.

To correlate the duration of each outage with general service availability, we add support lines for 99.9%, 99.99%, and 99.999% uptime. This is slightly optimistic since 5 IXPs had multiple outages in the same year. Still, 5% of the monitored 403 facilities fail to meet the 99.99% uptime mark and 18% the 99.999% uptime mark. Consequently, to provide services with availability beyond 99.999% service providers must use redundant facilities.
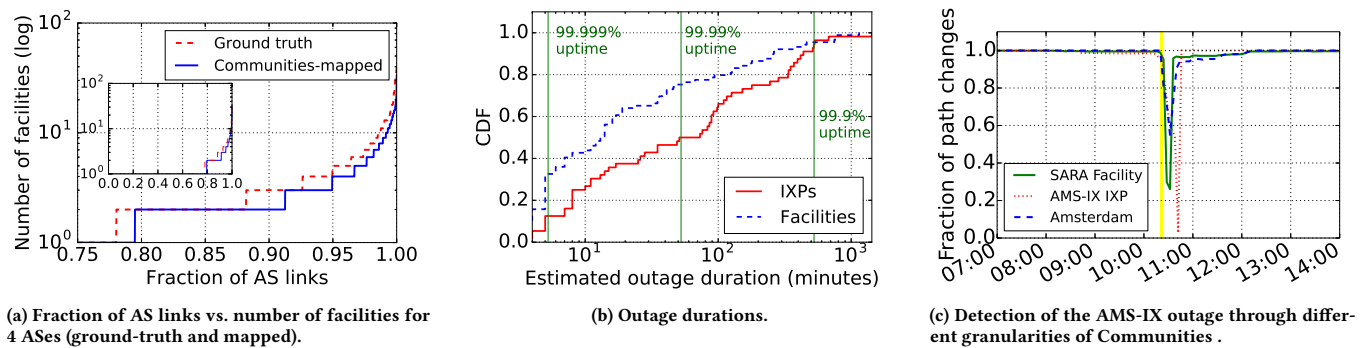
(a) Fraction of AS links vs. number of facilities for 4 ASes (ground-truth and mapped).

(b) Outage durations.

(c) Detection of the AMS-IX outage through different granularities of Communities .

Figure 8: Physical location deciphered by BGP updates (a), Outage durations (b), and AMS-IX outage case study (c).

6.2 Outages in Depth: Case Studies

To demonstrate *Kepler*'s capability to investigate outages we now focus on three outages in detail. The first one occurred at AMS-IX, a major IXP in Amsterdam, NL, at 2015-05-13. The outage was caused by a loop in the switching fabric during planned maintenance [18]. Figure 8c plots the path change fractions for three different aggregation levels. The outage caused the IXP to loose almost all routes and more than 90% of the exchanged traffic for about 10 minutes. It took about 15 minutes for the traffic to recover. The incident is clearly visible in all aggregation levels, but the paths tagged with the AMS-IX Communities show the largest drop indicating the actual source of the outage.

However, the changes in aggregated paths can be misleading. Indeed, Figure 9a shows the effect of two independent facility-level outages in London [3], on the co-located London-IXP LINX, and a third facility, TH East. At time *A*, when the first outage occurs, we see almost no change at the city level aggregation, while both LINX and TH East are affected. At time *B* we observe a city level signal, which does not correspond to a facility outage but rather a re-routing of paths from a major Tier-1 AS. At time *C*, when the second outage happens, we witness a major drop only through TH East. *Kepler* identifies correctly the *A* and *C* signals as PoP-specific and the *B* signal as AS-specific, and instead of inferring either LINX or TH East as the potential sources of the outages it proceeds to the signal localization by examining the impact of the outage on the far-end ASes against the facilities where these ASes are co-located.

This process is illustrated in Figure 9b, where it becomes clear that at time *A* and *C* two major subsets of the ASes at TC HEX8/9 and at TH North are affected. The far-end ASes in other London facilities (not depicted) show no concurrent signs of outage, allowing us to identify correctly TC HEX8/9 and TH North as the outage sources. Also note that at time *B* only a single AS is affected. This plot also highlights that ASes handle outages differently. Some return to their "stable" path once the outage is over, while other remain with their new path. This set of outages illustrates *Kepler*'s ability to disambiguate the source of outage signals to facilities.

6.3 Outages in Depth: Active Measurements

Next, we highlight the benefits of incorporating active measurements in *Kepler*. We focus on the outage at AMS-IX.

Backup paths: Figure 10a shows the BGP path changes while Figure 10b shows the traceroute path changes. While the overall path changes follow the same trend, the backup paths that are activated differ. The BGP Communities are mainly provided by large ASes with very diverse peering connectivity, allowing them to activate alternative peering links at remote IXPs. On the other hand, the majority of traceroute probes are hosted in local ASes, and so are the targets, and thus most, 75% of the alternate routes are via the transit interconnections.

Path restoration time: BGP path re-convergence took about 4 hours until 95% of the paths returned. Approximately 5% of the paths did not return even days after the outage. Such permanent route changes are either due to manual intervention or by the BGP decision process that prefers the newest path to break ties and minimize route flapping. Although, 85% of the traceroute paths return within one hour back to AMS-IX, a significant fraction continues to cross transit links. These results show that the actual impact of an outage on both control-plane and data-plane routing paths significantly outlasts the root cause of the outage, possibly necessitating a review of the SLAs provided by infrastructure operators.

Impact on End-to-end Delays: *Kepler* uses the traceroute data to assess the impact of an outage in round-trip time (RTT). While we acknowledge that RTTs from traceroute may not reflect RTTs as seen by TCP, they serve as indications of performance changes. Figure 10c shows the empirical cumulative distribution function of the RTT delays for the paths that traverse AMS-IX. We distinguish three time periods: (i) 20 minutes before the outage, (ii) during the outage, and (iii) 20 minutes after the outage, i.e., 10 minutes after the operation returned to normal. Moreover, we separated the paths into those that use AMS-IX (AMS-IX) and those that do not (No AMS-IX) during and after the outage. During the outage the median RTT rises by more than 100 msec for rerouted paths. For unchanged paths, the median RTT increase is moderate, and while some experience significant increase the tail does not grow as much. After the outage, this RTT increase disappears. Moreover, paths that returned to AMS-IX within 20 minutes experience roughly the same RTT as before the outage. However, 30% of the paths that still use the alternative interconnections continue to see increased RTTs of about 40 msecs due to sub-optimal routing in terms of distance.
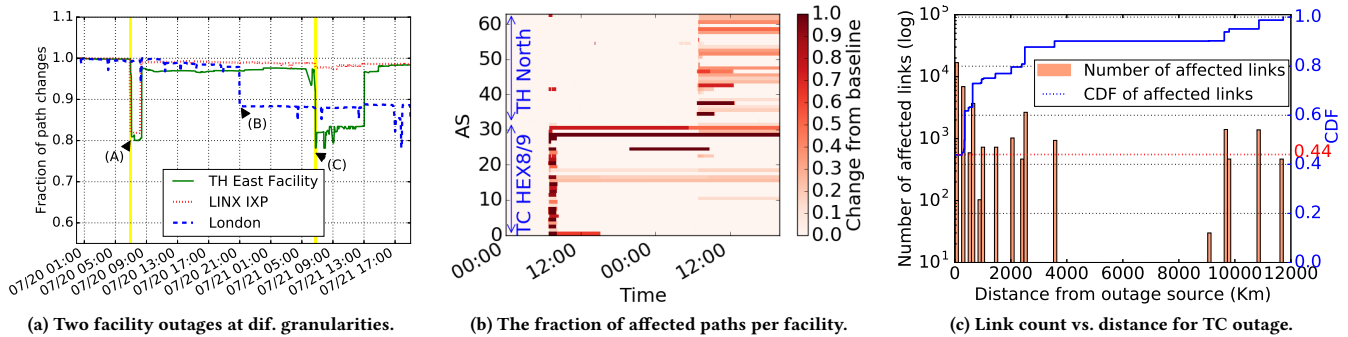
(a) Two facility outages at dif. granularities.   **(b) The fraction of affected paths per facility.**   **(c) Link count vs. distance for TC outage.**

**Figure 9: Two outages at different colocation facilities in London (TC HEX8/9 and TH North) at July 20 & 21, 2016.**

(a) BGP path changes.   **(b) Traceroute path changes.**   **(c) Effect on RTT.**   **(d) IPv4 traffic at EU-IXP.**

Figure 10: AMS-IX outage seen by *Kepler* (a)-(c), and by IPFIX traffic at EU-IXP (d).

6.4 Outages: Remote Impact

Remote Networks: To understand the impact of infrastructure outages we study the locations of the ASes that have been affected by the two London outages. We localize the IPs of the far-end interfaces of the affected ASes identified by *Kepler*'s traceroute using DRoP [58]. Figure 9c plots the distance from London in km vs. the number of affected ASes. Surprisingly, only 44% of the far-end interfaces are also in London. More than 45% of the interfaces are in a different country with more than 20% outside Europe. The main reason for such a widespread impact of localized failures is the increasing popularity of remote peering, an interconnection practice that allows ASes without physical presence at a peering hub to interconnect through layer-2 transport providers that resell peering ports across remote facilities [65]. Castro et al. estimated that 20% of the members in large IXPs connect remotely [16], which is consistent with our findings. This underlines the importance of understanding the facility-level topology when analyzing the impact of an outage.

Remote Infrastructures: To further challenge the expectation that a "local outage has only local impact", we complement *Kepler* with passive measurements from a major European IXP (EU-IXP). Figure 10d depicts the traffic volumes in Gbps at EU-IXP during the AMS-IX outage based on IPFIX data collected at its switching fabric with sampling rate 1/10K [21]. The two IXPs are 360 kilometers away. During the AMS-IX outage (t_0), we notice a sharp drop in IPv4 traffic—about 10% (215 Gbps). After 10 minutes – when the AMS-IX outage stopped (t_1) the traffic is rising above the expected average volume. This lasts for approximately 15 minutes. After the outage was restored (t_2), the traffic returns to normal levels.

To scrutinize this counter-intuitive observation we study the per member traffic. Only 136/533 members have a significant reduction in traffic (mean loss is 1 Gbps, max. loss 25 Gbps), with the rest seeing a mean growth of 188 Mbps (max.12 Gbps). However, traffic losses dominate even though moderate traffic increases are typical during this time of the day. The top 25 ASes with a traffic decrease account for 83% of the total loss. The outage above is no singular event. During other outages we observe similar traffic reductions, albeit smaller as the distance increases.

Remote Impact Explained: The conventional wisdom is that network operators should use separate edge routers for each colocation facility or IXP. However, due to the high cost of edge routers, operators often use a single router for multiple facilities introducing interdependencies among peering infrastructures, especially when they have common members. In addition, operational best practice prioritizes peering over upstream to keep traffic local and to reduce cost. Thus, prefixes reachable via IXPs will use the IXP rather than an upstream provider. Consider a scenario where an ISP uses two IXPs and the capacity of neither is sufficient for the total traffic of the ISP. While using peering links beyond 50% violates best practices, price pressure may force operators to ignore such guidelines. During outages ISPs may rely on their upstream causing a traffic drop at the other IXP, without extra cost for short outages due to the 95th percentile [86] pricing.

The most important reason is asymmetric paths [80], which are common in today's Internet. For peering infrastructures we call a path asymmetric if one direction only involves facility A and the reverse path only involves facility B. An outage at either of the two facilities causes a reduction of traffic at the other. We find that more

than 10% of all (source, destination) combinations between AMS-IX and EU-IXP members have asymmetric routes, which account for most of the traffic losses at EU-IXP.

7 IMPLICATIONS

Implications for Policymakers: The operation of our system, *Kepler*, and our analysis increases the transparency in Internet infrastructure outages. This can inform best practices for improving resilience, and would be of use to regulators and policy makers given the critical role of such infrastructures [30, 39]. In addition, with *Kepler* one can provide testimony based on hard evidence to assess the degree of violation of service level agreements, e.g., the 5 nines reliability, and to characterize an outage as full or partial, and to assess the impact on the operation of a network.

Implications for Peering: Our analysis shows that redundant peering strategies may increase the resilience to outages. Still, in some cases, we observed peering disruptions even when redundant peering was available. We argue that there is significant space for improving peering resilience by taking into account the physical isolation of peering infrastructures. Unfortunately, the interdependency among the various peering infrastructures is often not well known, and thus greater resilience might be achieved with more collaboration between peering infrastructure operators and network operators.

Implications for Operation: Our study shows that an increasing number of networks tag their BGP announcements with communities, and that about half of the prefix announcements include location-based communities. This practice is of great help for detecting outages. However, we should point out that the propagation of location-based communities has a downside. The leakage of this information enables easier detection by third parties of the location where two networks establish interconnections. This leakage can be used for business intelligence, and for targeted attacks. Hence, we will make *Kepler* available via an interactive interface. But we will only share our dictionary of location-based communities on request.

8 CONCLUSION

Outages at colocation facilities and IXPs affect the operation of hundreds of networks. In this paper, we show that control-plane messages provide an excellent, yet unexplored source of information that can be utilized to detect peering infrastructure outages in the wild. We develop a methodology to analyze the values of the BGP Communities attribute to accurately detect the location of a peering outage at the level of a building. While our method is general enough to be applied to any stream of BGP data, we show that the implementation is far from trivial. Based on our methodology, we built and operate *Kepler* for detecting peering infrastructure outages. Over a 5-year period, we detected about 160 colocation facility or IXP outages, which is four times what could be discerned from operator mailing lists and related sources. Our results show that local outages at these peering infrastructures can have an impact on remote networks and seemingly unrelated remote peering infrastructures. Thus, *Kepler* can provide feedback to operators, researchers, and policy makers alike to improve the understanding of the Internet's resilience.

REFERENCES

[1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. 2012. Anatomy of a Large European IXP. In *ACM SIGCOMM*.
[2] AMS-IX. 2016. Connected Parties. https://ams-ix.net/connected_parties. (2016).
[3] Ars Technica. 2016. BT, other ISPs hit by second major Internet outage-power failure blamed. goo.gl/fGx6nF. (July 2016).
[4] ATNOG Mailing List. 2015. AMSIX Heute. https://atnog.at/pipermail/atnog/2015-May/000022.html. (May 2015).
[5] R. Banerjee, A. Razaghpanah, L. Chiang, A. Mishra, V. Sekar, Y. Choi, and P. Gill. 2015. Internet Outages, the Eyewitness Accounts: Analysis of the Outages Mailing List. In *PAM*.
[6] G. D. Battista, T. Refice, and M. Rimondini. 2006. How to extract BGP peering information from the internet routing registry. In *SIGCOMM workshop on Mining network data*. ACM, 317–322.
[7] A. Beccaris, D. Quinn, D. Barroso, H. Adollarsson, and M. Walster. 2016. PINDER: peer speed-dating. RIPE NCC IXP tools hackathon. http://accel.waffle.sexy/pinder.pdf. (October 2016).
[8] K. Benson, A. Dainotti, K. Claffy, and E. Aben. 2012. Gaining Insight into AS-level Outages through Analysis of Internet Background Radiation. In *TMA*.
[9] R. Beverly and L. Alt. 2014. On the Potential for Mining Unstructured Public Data to Aid Network Intelligence. (2014).
[10] S. Bird. 2006. NLTK: The Natural Language Toolkit. In *COLING-ACL*.
[11] S. Brito, M. Santos, R. Fontes, D. Perez, and C. Rothenberg. 2016. Dissecting the Largest National Ecosystem of Public Internet eXchange Points in Brazil. In *PAM*.
[12] Broadband Internet Technical Advisory Group Report (BITAG). 2014. Interconnection and Traffic Exchange on theInternet. (2014).
[13] C. Partridge, P. Barford, D. D. Clark, S. Donelan, V. Paxson, J. Rexford, and M. K. Vernon. 2003. *The Internet Under Crisis Conditions: Learning from September 11.* The National Academy Press.
[14] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger. 2010. Towards an AS-to-Organization Map. In *ACM IMC*.
[15] CAIDA. 2016. Archipelago (Ark) Measurement Infrastructure. http://www.caida.org/projects/ark/. (2016).
[16] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. 2014. Remote Peering: More Peering without Internet Flattening. In *CoNEXT*.
[17] R. Chandra, P. Traina, and T. Li. 1996. BGP Communities Attribute. IETF RFC 1997. (1996).
[18] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. 2013. There is More to IXPs than Meets the Eye. *ACM CCR* 45, 5 (2013).
[19] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao. 2009. Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users. In *CoNEXT*.
[20] K. Cho, C. Pelsser, R. Bush, and Y. Won. 2011. The Japan Earthquake: the impact on traffic and routing observed by a local ISP. In *ACM CoNEXT SWID workshop*.
[21] B. Claise, B. Trammell, and P. Aitken. 2013. RFC 7011: Specification of the IPFIX Protocol for the Exchange of Flow Information. (2013).
[22] Cymru. 2016. BGP Bogon Refence. https://goo.gl/An2cdU. (2016).
[23] A. Dainotti, R. Amman, E. Aben, and kc claffy. 2012. Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet. *ACM CCR* 42, 1 (2012).
[24] A. Dainotti, C. Squarcella, E. Aben, KC Claffy, M. Chiesa, M. Russo, and A. Pescape. 2011. Analysis of Country-wide Internet Outages Caused by Censorship. In *ACM IMC*.
[25] Data Center Dynamics. 2016. http://www.datacenterdynamics.com/. (2016).

[26] Data Center Knowledge. 2016. http://www.datacenterknowledge.com/. (2016).
[27] Data Center Map. 2016. http://www.datacentermap.com/. (2016).
[28] DE-CIX. 2016. Connected networks in DE-CIX Frankfurt. https://goo.gl/DnPz6s. (2016).
[29] DE-CIX. 2016. Frankfurt enabled sites. https://goo.gl/BG2yWv. (August 2016).
[30] Department of Homeland Security. 2010. Communications Sector-Specific Plan: An Annex to the National Infrastructure Protection Plan. (2010).
[31] C. Dietzel, A. Feldmann, and T. King. 2016. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In PAM.
[32] F. Dobrian, A. Awan, D. Joseph, A. Ganjam, J. Zhan, V. Sekar, I. Stoica, and H. Zhang. 2011. Understanding the Impact of Video Quality on User Engagement. In ACM SIGCOMM.
[33] B. Donnet and O. Bonaventure. 2008. On BGP communities. ACM CCR 38, 2 (Mar 2008), 55–59.
[34] R. Durairajan, P. Barford, J. Sommers, and W. Willinger. 2015. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In ACM SIGCOMM.
[35] Z. Durumeric, E. Wustrow, and J. A. Halderman. 2013. ZMap: Fast Internet-Wide Scanning and its Security Applications. In USENIX Security Symposium.
[36] Equinix. 2016. Equinix Expands Data Center Leadership Position with Close of Telecity Acquisition. https://goo.gl/vbXDCv. (2016).
[37] Equinix. 2016. Investor Relations: Annual and Quarter Results. http://investor.equinix.com/. (2016).
[38] European Internet Exchange Association. 2016. https://www.euro-ix.net. (2016).
[39] European Union Agency for Network and Information Security. 2010. Critical Infrastructures and Services, Internet Infrastructure: Internet Interconnections. https://goo.gl/SJMJfJ. (2010).
[40] R. Fanou, P. Francois, and E. Aben. 2015. On the Diversity of Interdomain Routing in Africa. In PAM.
[41] P. Faratin, D. Clark, S. Bauer, W. Lehr, P. Gilmore, and A. Berger. 2008. The Growing Complexity of Internet Interconnection. Communications and Strategies (2008).
[42] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. 2004. Locating Internet routing instabilities. In ACM SIGCOMM.
[43] J. R. Finkel, T. Grenager, and C. Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In Annual Meeting on Association for Computational Linguistics.
[44] K. Foster. 2003. Application of BGP Communities. The Internet Protocol Journal 6, 2 (Sep 2003).
[45] France-IX. 2017. Interconnection with other IXPs. https://goo.gl/gBYcRH. (2017).
[46] A. Gerber and R. Doverspike. 2011. Traffic Types and Growth in Backbone Networks. In OFC/NFOEC.
[47] M. Ghobadi and R. Mahajan. 2016. Optical Layer Failures in a Large Backbone. In IMC.
[48] V. Giotsas, A. Dhamdhere, and k. claffy. 2016. Periscope: Unifying Looking Glass Querying. In PAM.
[49] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy. 2014. Inferring Complex AS Relationships. In IMC.
[50] V. Giotsas, G. Smaragdakis, B. Huffaker, M. Luckie, and k. claffy. 2015. Mapping Peering Interconnections at the Facility Level. In CoNEXT.
[51] V. Giotsas and S. Zhou. 2013. Improving the Discovery of IXP Peering Links through Passive BGP Measurements. In Global Internet.
[52] V. Giotsas, S. Zhou, M. Luckie, and kc claffy. 2013. Inferring Multilateral Peering. In CoNEXT.
[53] Google. 2016. Google Maps Geocoding API. https://goo.gl/mvDy17. (2016).
[54] S. P. Gorman and E. J. Malecki. 2002. Fixed and Fluid: Stability and Change in the Geography of the Internet. Telecommunications Policy 26, 7-8 (2002), 389–413.
[55] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett. 2014. Peering at the internet's frontier: A first look at isp interconnectivity in africa. In PAM.
[56] J. Heidemann, L. Quan, and Y. Pradkin. 2012. A Preliminary Analysis of Network Outages During Hurricane Sandy. Technical Report ISI-TR-2008-685b. USC/Information Sciences Institute.
[57] N. Hilliard, E. Jasinska, R. Raszuk, and N. Bakker. 2016. Internet Exchange BGP Route Server Operations. IETF RFC 7948. (September 2016).
[58] Bradley Huffaker, Marina Fomenkov, et al. 2014. DRoP: DNS-based router positioning. ACM SIGCOMM Computer Communication Review 44, 3 (2014), 5–13.
[59] Init7 NOC. 2016. BGP Communities For Init7 customers. http://as13030.net/static/pdf/as13030_bgp_communities.pdf. (January 2016).
[60] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. 2008. Studying Black Holes in the Internet with Hubble. In NSDI.
[61] R. Kloti, B. Ager, V. Kotronis, G. Nomikos, and X. Dimitropoulos. 2016. A Comparative Look into Public IXP Datasets. ACM CCR 46, 1 (2016).
[62] C. Labovitz, S. Lekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. 2010. Internet Inter-Domain Traffic. In ACM SIGCOMM.
[63] C. Lees, J. Paussa, and A. Fenioux. 2016. Peer Match-making. RIPE NCC IXP tools hackathon. https://goo.gl/g1uZKA. (October 2016).

[64] C. Lévy-Bencheton, L. Marinos, R. Mattioli, T. King, C. Dietzel, and J. Stumpf. 2015. Threat Landscape and Good Practice Guide for Internet Infrastructure. EU Agency for Network and Information Security (ENISA) (2015).
[65] LINX. 2016. LINX From Anywhere. https://goo.gl/gHhfn1. (2016).
[66] LINX. 2016. LINX Virtual PoP Programme. https://goo.gl/eABGhC. (2016).
[67] Outage Reporting List. 2016. https://puck.nether.net/pipermail/outages/. (2016).
[68] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and K. Claffy. 2014. Using PeeringDB to Understand the Peering Ecosystem. ACM CCR 44, 2 (2014).
[69] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Kirshnamurthy, and A. Venkataramani. 2006. iPlane: An information plane for distributed systems. In ACM OSDI.
[70] E. J. Malecki. 2002. The Economic Geography of the Internet's Infrastructure. Economic Geography 78, 4 (2002), 399–424.
[71] E. J. Malecki. 2012. Internet networks of world cities: agglomeration and dispersion. International Handbook of Globalization and World Cities (2012), 117.
[72] R. Motamedi, B. Chandrasekaran, B. Maggs, R. Rejaie, and W. Willinger. 2014. On the Geography of X-Connects. Technical Report CIS-TR-2014-02. University of Oregon.
[73] W. Muhlbauer, A. Feldmann O. Maennel, M. Roughan, and S. Uhlig. 2006. Building an AS-Topology Model that Captures Route Diversity. In ACM SIGCOMM.
[74] NANOG. 2016. Mailing List and Archives. https://www.nanog.org/list/archives. (2016).
[75] NOC Incident Reporting websites. 2017. http://goo.gl/iNRlRY. Version 1. (2017).
[76] G. Nomikos and X. Dimitropoulos. 2016. traIXroute: Detecting IXPs in traceroute paths. In PAM.
[77] W. B. Norton. 2010. The Art of Peering: The Peering Playbook. (2010).
[78] W. B. Norton. 2014. European vs. U.S. Internet Exchange Points. https://goo.gl/qvsgG4. (2014).
[79] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti. 2016. BGPStream: a software framework for live and historical BGP data analysis. In IMC.
[80] V. Paxson. 1997. End-to-End Routing Behavior in the Internet. IEEE/ACM Transactions on Networking 5, 5 (1997), 601–615.
[81] PeeringDB. 2016. IXPs and colocation database. https://www.peeringdb.com. (2016).
[82] I. Poese, B. Frank, G. Smaragdakis, S. Uhlig, A. Feldmann, and B. Maggs. 2012. Enabling Content-aware Traffic Engineering. ACM CCR 42, 5 (2012).
[83] L. Quan, J. Heidemann, and Y. Pradkin. 2012. Detecting Internet Outages with Precise Active Probing. Technical Report ISI-TR-701. USC/Information Sciences Institute.
[84] L. Quan, J. Heidemann, and Y. Pradkin. 2013. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In ACM SIGCOMM.
[85] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig. 2003. Interdomain Traffic Engineering with BGP. IEEE Communications Magazine 41, 5 (2003), 122–128.
[86] V.R. Raja, A. Dhamdhere, A. Scicchitano, S. Shakkottai, and S. Leinen. 2014. Volume-Based Transit Pricing: Is 95 the Right Percentile?. In PAM.
[87] The Register. 2016. BT internet outage was our fault, says Equinix. https://goo.gl/YBNYwF. (July 2016).
[88] Y. Rekhter, T. Li, and S. Hares. 2006. A Border Gateway Protocol 4 (BGP-4). IEFT RFC 4271. (2006).
[89] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. 2014. Peering at Peerings: On the Role of IXP Route Servers. In ACM IMC.
[90] RIPE NCC. 2016. RIPE Atlas. https://atlas.ripe.net/. (2016).
[91] RIPE NCC. 2016. RIPE Atlas Rate Limits. https://goo.gl/8Wr6S6. (2016).
[92] M. Roughan, W. Willinger, O. Maennel, D. Pertouli, and R. Bush. 2011. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. IEEE J. on Sel. Areas in Comm. 29, 9 (2011).
[93] S. Sangli, D. Tappan, and Y. Rekhter. 2006. BGP Extended Communities Attribute. IETF RFC 4360. (2006).
[94] A. Sapegin and S. Uhlig. 2013. On the extent of correlation in BGP updates in the Internet and what it tells us about locality of BGP routing events. Computer Communications 36, 15–16 (2013).
[95] R. Singh and P. Gill. 2016. PathCache: A Path Prediction Toolkit. In SIGCOMM Poster and Demo Session.
[96] R. Stapleton-Gray and W. Woodcock. 2011. National Internet Defense – Small States on the Skirmish Line. Comm. of the ACM 54, 3 (2011).
[97] Fortune Tech. 2015. Digital Realty Trust to buy Telx in $1.88 billion deal. https://goo.gl/Bq4nJj. (2015).
[98] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage. 2010. California Fault Lines: Understanding the Causes and Impact of Network Failures. In ACM SIGCOMM.
[99] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. 2002. Observation and Analysis of BGP Behavior under Stress. In ACM IMW.
[100] B. Woodcock and M. Frigino. 2016. Survey of Internet Carrier Interconnection Agreements. Packet Clearing House.. (November 2016).