Following the Data Trail: An Analysis of IXP Dependencies

Malte Tashiro[1,2][0000−0003−1257−9534], Romain Fontugne[2][0000−0002−9816−5625], and Kensuke Fukuda[1,3][0000−0001−8372−2807]

[1] SOKENDAI, Tokyo, Japan
[2] IIJ Research Laboratory, Tokyo, Japan {malte,romain}@iij.ad.jp
[3] NII, Tokyo, Japan kensuke@nii.ac.jp

Abstract. Internet exchange points (IXPs) play a vital role in the modern Internet. Envisioned as a means to connect physically close networks, they have grown into large hubs connecting networks from all over the world, either directly or via remote peering. It is therefore important to understand the real footprint of an IXP to quantify the extent to which problems (e.g., outages) at an IXP can impact the surrounding Internet topology. An IXP footprint computed only from its list of members as given by PeeringDB, or the IXP's website, is usually depicting an incomplete view of the IXP as it misses downstream networks whose traffic may transit via an IXP although they are not directly peering there. In this paper we propose a robust approach that uncovers this dependency using traceroute data from two large measurement platforms. Our approach converts traceroutes to paths that include both autonomous systems (ASes) and IXPs and computes AS Hegemony to infer their interdependencies. This technique discovers thousands of dependent networks not directly connected to IXPs and emphasizes the role of IXPs in the Internet topology. We also look at the geolocation of members and dependents and find that only 3 % of IXPs with dependents are entirely local: all members and dependents are in the same country as the IXP. Another 52 % connect international members, but only have domestic dependents.

1 Introduction

The Internet is continuously growing and its topology is getting increasingly more complex. Originally a clear hierarchy, the structure of the Internet transforms into a flat mesh [31]. This transformation is facilitated by the advance of internet exchange points (IXPs), that establish peering facilities where networks can connect directly. Enticed by the promise of cost reduction and potential latency improvements [16], the number of both IXPs and IXP members has seen consistent growth over the past years [24]. While the original idea of IXPs was to promote connectivity between physically close networks, reduce unnecessary routing detours, and "keep local traffic local" [27], their use has grown beyond that. IXPs can now provide impressive reach [49], they have become important infrastructure for content delivery networks (CDNs) [25], and are used for DDoS mitigation [54, 55].

This evolution makes it paramount to understand the footprint of IXPs in the Internet topology. Like any part of the Internet they are subject to failures and congestion [15, 35]. The goal of this paper is to quantify the importance of IXPs beyond their members. A better understanding of how networks, maybe inadvertently, depend on IXPs, can help system engineers, peering coordinators, and policy makers in their decision process. For example, to increase resilience a network operator might want to avoid using two transit providers that both depend on the same IXP.

The Internet topology has been studied in the past by measuring dependency between autonomous systems (ASes). These studies mostly analyzed AS paths in BGP data [36,38,46], occasionally complementing it with traceroute data [19]. Our proposed method quantifies not only AS-level dependencies, but expands the analysis to IXPs. The study of IXP dependencies allows us to contextualize the role of IXPs for the Internet. Unfortunately, IXPs are generally not visible in AS paths, hence BGP data is unsuitable for our study (Section 3). Instead, we aim to infer dependency only from traceroute data as IXP peering LANs are easily identifiable there. We use IPv4 traceroute data from two large measurement platforms in combination with PeeringDB [9] and looking glass information to detect IXPs and how they are traversed. By transforming traceroutes to `n`-paths, AS paths enhanced with IXP identifiers, we can adapt the AS Hegemony [36] metric to reveal the inter-dependencies between ASes and IXPs (Section 4). We validate this new approach against results obtained from BGP data and find that both methods produce comparable numbers of AS dependencies (Section 5).

Based on the computed dependency values we highlight some of the differences and similarities between transit ASes and IXPs (Section 6). Then we investigate the *topological* footprint of IXPs by comparing the relationship between the number of members and dependent ASes, showing that they are not strictly correlated (Section 7). We also look at the *geographic* footprint of IXPs revealing that some IXPs have dependents in countries not covered by their members. In addition, we infer the *Regionality* of IXPs, i.e., if their members and dependents are from the same country, and show a comparison between all monitored IXPs. Finally, we dissect the dependents of two large IXPs, `DE-CIX Frankfurt` and `IX.br São Paulo`, highlighting the different roles played by these IXPs in the Internet topology; `DE-CIX Frankfurt` is acting mostly as an international hub for dependents from multiple countries, whereas `IX.br São Paulo` dependents are almost entirely from Brazil (Section 8).

In summary, our main contributions are:

- A method to infer IXP and AS dependencies from traceroute data.
- A high-level study of the role of IXPs in the Internet topology compared to ASes by analyzing the number and magnitude of dependencies.
- The analysis of the topological and geographical footprint of IXPs based on the location of dependent networks, as well as insights into the locality of connected networks.

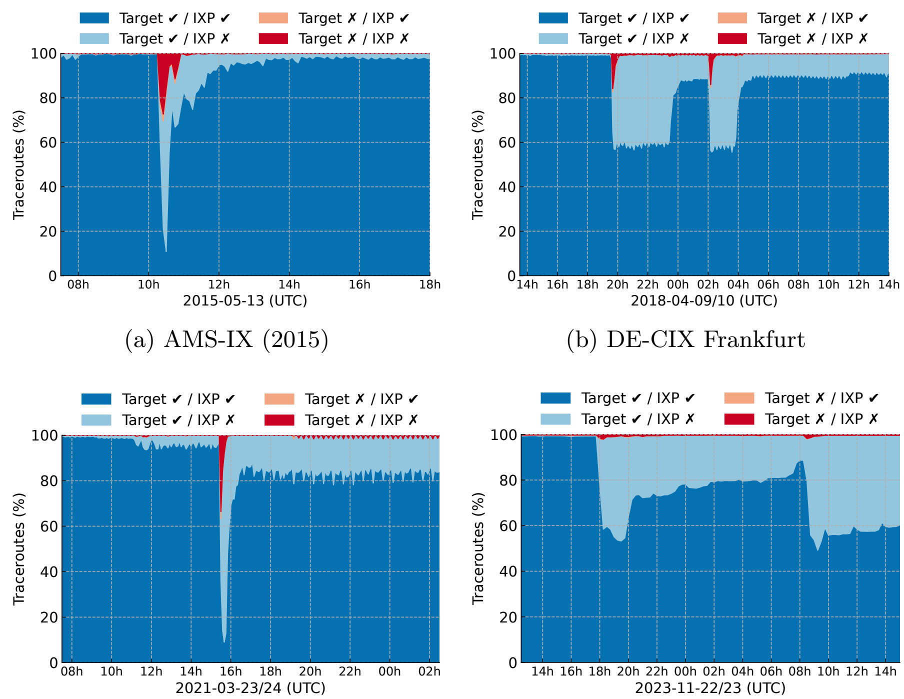

(a) AMS-IX (2015)

(b) DE-CIX Frankfurt

(c) LINX LON1

(d) AMS-IX (2023)

Fig. 1. Impact of partial outages at three IXPs on connectivity. Traceroutes are grouped into four categories, depending on if the target was reached (`Target ✓`) and if the IXP was on the path (`IXP ✓`).

- A case study highlighting the usefulness of dependencies by comparison of two large IXPs, `DE-CIX Frankfurt` and `IX.br São Paulo`, showing disparate topological footprints.
- The publication and periodic update of IXP dependency information [53].

2 Motivation

IXPs play an increasingly vital role in the global Internet topology. However, as any complex system they are not safe from failures caused by human error, power outages, or other sources. Understanding the potential impact of these failures is hindered by the fact that IXPs operate mostly at Layer 2, and thus, they are disregarded by large scale Internet topology studies which are usually based on BGP data [32, 34, 36, 40].

To motivate the importance of understanding the relationship between ASes and IXPs we take a look at four historical outages at large IXPs: A partial outage at AMS-IX in 2015 due to a configuration error during maintenance work [12], a

Table 1. Composition of probe-target pairs used to investigate IXP outages.

	AMS-IX (2015)	**DE-CIX**	**LINX**	**AMS-IX (2023)**
Probe-target pairs	11,655	25,945	25,876	46,377
Atlas probes	3116	4437	4300	6595
Probe ASes	1186	1933	1631	2428
Probes/AS (median)	2	4	3	3
Target IPs	441	561	799	1241
Target ASes	240	358	500	719
IPs/AS (median)	10	20	21	16
RTT via IXP (avg.)	55 ms	87 ms	126 ms	81 ms
RTT bypassing IXP (avg.)	84 ms	93 ms	131 ms	84 ms

power outage at a DE-CIX Frankfurt datacenter in 2018 [41], an outage caused by undisclosed technical reasons at LINX LON1 in 2021 [7], and finally another outage at AMS-IX in 2023 caused by a misconfigured access list [13]. We replicate the methodology of [15] and investigate these outages through the lens of RIPE Atlas [50]. We analyze the incidents with the help of traceroute data, which gives us a fine-grained view of the paths before, during, and after the outage. To infer the impact of these outages on connectivity, we want to inspect traceroutes from probes that consistently traverse the IXP in the absence of failures. We then monitor if these traceroutes continue to reach their targets during the outage, and if the IXP is still traversed.

To select a reliable set of probes and targets, we start with all traceroutes run by Atlas probes during the day before the outage. The traceroutes are grouped by probe-target pairs, and we keep only pairs for which (1) the IXP's peering LAN is visible in *all* traceroutes, (2) the target host responds to the traceroutes, and (3) we see at least 24 traceroutes within the inspected interval (i.e., on average one per hour). This filtering results in a diverse set of pairs as shown in Table 1.

Next, we inspect the behavior of traceroutes between these probes and targets around the time of the outage. Each traceroute is categorized based on the reachability of the target and if the peering LAN of the IXP is present on the path. Fig. 1 shows the category distribution over time. The traceroute results are grouped into five-minute bins. In all cases, the time of the outage (two incidents in the case of DE-CIX and AMS-IX 2023) is clearly visible. The dark blue area represents the normal state in which the traceroute traverses the IXP and the target responds. However, at the time of the outage, a large share of traceroutes does not pass through the IXP anymore. Although the majority still reaches the target through rerouting (light blue), a significant amount loses connectivity to the target temporarily (red). This impact is most visible in the AMS-IX 2015 and LINX LON1 outages, where up to 71 % and 83 % of traceroutes avoided the

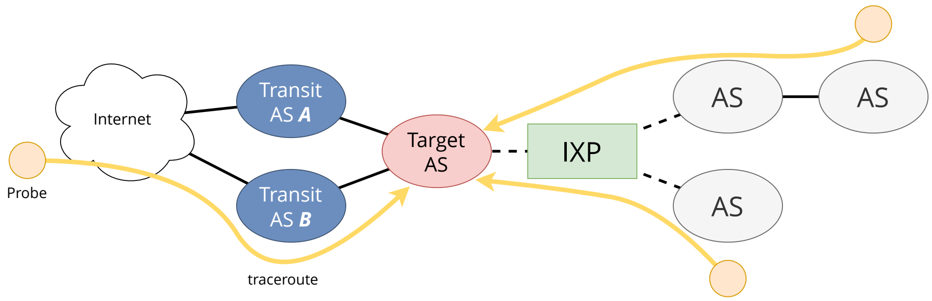

Fig. 2. Typical Internet topology: An AS connects to the Internet via transit providers and might also peer with other ASes via an IXP. Inference of the topology via traceroute may miss ASes that are not traversed (e.g., Transit AS A), but detects IXPs.

IXP and up to 30 % and 37 % failed to reach their target at the peak of the outage.

Although rerouting around the IXP allows some probes to reach the target (light blue), it does come with a price: Comparing the average RTT of a probe-target pair reaching the destination via the IXP to the same target but routing around the IXP we see an increase by 52 % in the case of AMS-IX 2015 (from 55 ms to 84 ms, cf. Table 1). Even though the effect was smaller for DE-CIX (+7 %), LINX (+5 %), and AMS-IX 2023 (+4 %) it is still clear that routing around the IXPs has a negative performance impact. In addition, these transitions between IXP and transit networks may incur additional financial costs for the operators.

This analysis demonstrates that an outage at the IXP does have consequences for networks that use it for connectivity. It is therefore imperative to understand which networks use IXPs, either directly or via transit ASes, and to what degree. The goal of this paper is to analyze the relationship between IXPs and networks and quantify how networks utilize IXPs for connectivity. We call this relationship a *dependency* between an AS and an IXP if we detect that the AS consistently uses the IXP for its global connectivity. A precise definition is given in Section 4.3.

3 Design Decisions

In this section we explain the differences between BGP data and traceroutes for dependency calculations (summarized in Table 2) and emphasize that traceroutes are required to detect dependencies on IXPs.

We identified three key differences that are relevant when inferring dependencies: The scope of the visible topology, the ability to identify IXPs, and the temporal granularity with which dependencies can be calculated. We illustrate these differences based on Fig. 2, which shows a simplified topology around a target AS for which we want to infer dependencies. The first difference is the

Table 2. Property comparison of BGP and traceroute to infer the Internet topology and dependencies. BGP data is available in real time, gives a global view of the topology visible to the peers, but does not contain IXPs. Traceroutes reveal only the parts of the topology they traverse and thus need to be aggregated over a time window, but reveal the locations of IXPs.

Type	**Topological granularity**		**Temporal granularity**
	Scope	IXP visibility	
BGP	Global	No	Real time
traceroute	Per path	Yes	Time window

topological scope of both datasets: BGP peers usually share their full routing tables with route collectors, and combining these tables reveals a large part of the Internet topology including all globally reachable ASes. Traceroutes only reveal ASes that are on the path the packets traverse, hence the number of discovered ASes is mainly governed by the number of collected traceroutes. However, because both BGP and traceroute data are obtained from a limited number of vantage points, both do not provide a complete view on AS links. The second difference is the ability to identify IXPs in the discovered topology. We can infer the presence of IXPs on the path with traceroute by detecting hops that are within the peering LAN of an IXP as shown in Fig. 2. In contrast, IXPs are not visible in BGP data because they are not explicitly involved in the interdomain routing process. The third difference is the temporal granularity with which the data can be retrieved. BGP data is available as snapshots in form of routing information databases (RIBs) complemented by update files, which enable seamless reconstruction of the visible topology at any time. As mentioned above, traceroutes discover ASes by actively traversing them. Since traceroutes obtained from measurement platforms run periodically — and to random targets — it is necessary to aggregate data over a long period of time to obtain a reasonable view of the topology (further discussed in Section 4.1). Therefore, traceroute data is unsuitable for the study of transient changes and should only be used to build a long-term view of the topology.

In summary, while there are differences when inferring dependencies from BGP and traceroute data (see Table 2), our study of IXPs makes the use of traceroutes mandatory. This requires the design of a new data processing pipeline to infer dependencies, which is described in the following section. Further limitations of traceroute are also explained in Section 9.

4 Methodology

Figure 3 illustrates the different parts of the proposed processing pipeline. Our analysis is based on large traceroute datasets retrieved from measurement platforms like RIPE Atlas and CAIDA's Ark. Each traceroute is transformed to a `n`-path, an AS path enhanced with IXP identifiers. Finally, the `n`-paths are fed

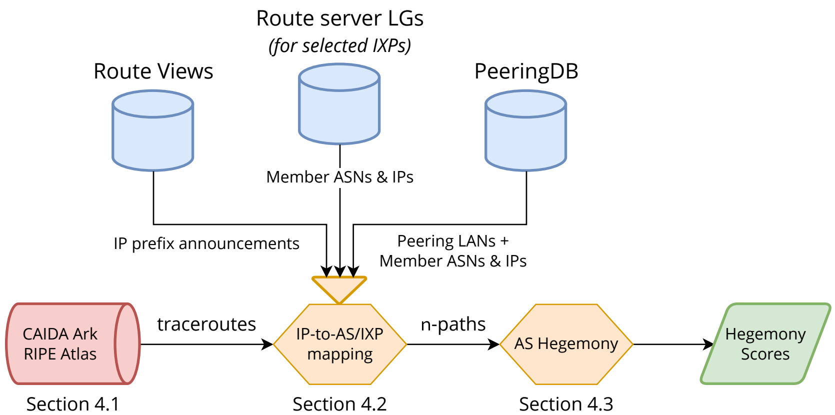

Fig. 3. The processing pipeline: First, traceroutes from CAIDA Ark or RIPE Atlas are converted to `n`-paths. The conversion uses a combination of Route Views, PeeringDB, and IXP looking glass information to map traceroute hops to ASes or IXPs. Finally, the `n`-paths are fed to the AS Hegemony algorithm to compute Hegemony scores.

to the AS Hegemony algorithm, resulting in Hegemony scores from which we infer dependencies.

4.1 Dataset

We consider two traceroute datasets: the topology measurements [14] from RIPE Atlas [50] and the IPv4 Routed /24 Topology Dataset from CAIDA's Ark platform [4]. The Atlas topology measurements[4] employ all Atlas probes but these are not evenly distributed and are known to have location bias [20, 21, 29, 51]. Therefore, we select a sample of 1000 probes based on the approach in [18] that prioritizes AS-path diversity, i.e., the probes are selected in a way that increases the AS-path length between them, thus providing a set of probes that are widespread over the Internet topology. The probe set covers 115 countries and 1000 ASes. CAIDA's dataset consists of 90 probes (a.k.a. monitors) located in 36 countries covering 74 ASes. We select all of them.

Both datasets are attempts to reveal the Internet topology by conducting traceroute measurements from the probes to a very large number of target IP addresses. CAIDA computes its list of target IPs by selecting one random IP from each routed /24 prefix and distributes this list between all probes. It guarantees that the target list is entirely processed before the next round of measurements starts. The Atlas target IP list consists of the first address of each globally routed prefix seen in BGP, which includes large prefixes and is therefore coarser. In addition, the list is reset daily and there is no guarantee that the target IPs have all been processed. Atlas probe scheduling makes the probe assignment

[4] Measurement IDs for IPv4: 5051 and 5151

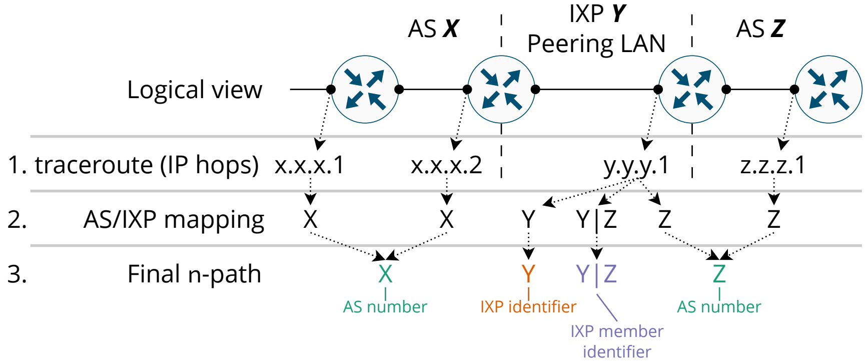

Fig. 4. Transformation of a traceroute to an `n`-path: Individual traceroute hops are mapped to their respective AS (2.). If a hop lies within an IXP peering LAN, it is mapped to an IXP identifier. If there is additional information available that assigns the precise IP to an IXP member, the hop is also mapped to an IXP member identifier and the AS of the member. Finally, duplicate AS numbers are removed (3.).

less deterministic and probes are capped to one traceroute every 15 minutes. As a consequence, a preliminary comparison between the two datasets revealed challenges with the Atlas dataset.

We analyzed data for one week in September (2022-09-12 – 2022-09-19), which contained ≈ 1.8 million traceroutes for Atlas and ≈ 136 million traceroutes for Ark. When computing Hegemony scores for each dataset we found that 67 % of target ASes are probed by RIPE Atlas probes from less than 10 ASes and are ignored as providing unreliable results (as explained in Section 4.3), whereas only 9 % of targets are ignored with the Ark dataset. To increase the reliability of results we increased the size of the data window for Atlas to four weeks (2022-09-05 – 2022-10-03; ≈ 7.4 million traceroutes), which reduced the amount of ignored targets to 24 %. When comparing the results with the previous and following weeks we found that they were similar. Thus, for ease of discussion we present only results based on one week of data for Ark and four weeks for Atlas.

4.2 Translating traceroute data to `n`-paths

The next step of the pipeline is to transform the traceroute results to `n`-paths. A `n`-path looks like an AS path, but contains additional hops that represent IXPs and IXP members. As shown in Fig. 3 we combine different data sources to build a comprehensive IP-to-AS/IXP mapping. Similar to existing approaches [26, 43, 47] we rely on IXP data from the PeeringDB [9] database and BGP data from Route Views [11] to detect IXPs and ASes in traceroutes. We identify IXPs in traceroute using IXP peering LANs from PeeringDB. Well maintained PeeringDB entries also contain interface information for individual IXP members, which enables us to also include the member AS in our paths. PeeringDB was

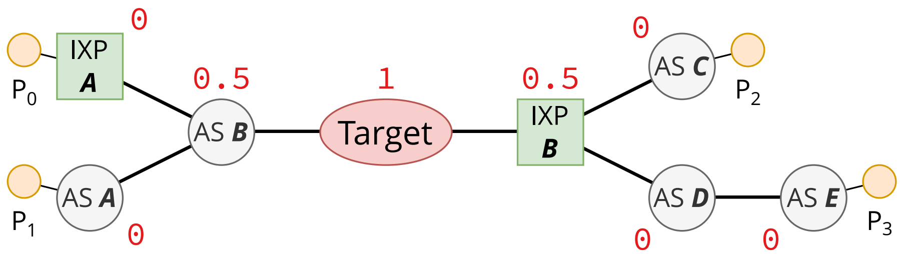

Fig. 5. A simplified view of AS Hegemony applied to n-paths: Hegemony scores are assigned to both ASes and IXPs. Nodes close to the vantage points are ignored to prevent biased scores caused by an uneven probe distribution.

validated in [44], where the authors show that PeeringDB membership is reasonably representative of the Internet's transit, content, and access providers in terms of business types and geography of participants, and data is generally up-to-date. For IXPs that use the Alice-LG [1] looking glass (e.g., DE-CIX, IX.br, LINX), we retrieve additional information to enhance membership data. BGP data from Route Views [11] is used to map other hops to ASes.

Figure 4 highlights the transformation process in more detail. In the first step, each hop of the traceroute is mapped to an AS number, an IXP identifier, or both. If detailed information about an IP is available, an IXP member identifier is added as well. If an IP address can not be mapped, the hop is ignored, which is a known shortcoming of using traceroute to infer AS-paths [42,45,57], but has limited impact on our approach. In our datasets, the IP-to-AS mapping fails for 2.5 % to 3.5 % of hops. Since traceroute sends more than one packet per hop, it is possible to receive replies from different IP addresses. If these IP addresses map to different ASes, they are included as an AS set. After mapping all hops we remove duplicate ASes, keeping only the first occurrence in the `n`-path.

4.3 AS Hegemony

We adapt AS Hegemony, a metric for quantifying AS inter-dependencies [36], to work with traceroute data and measure dependency to IXPs. We selected this metric because it is based on simple principles (i.e., using no BGP-specific heuristics such as AS relationships), it takes only paths as input data, and it has proven to be practical in various use cases [28, 33, 34, 37].

First, we give an intuitive explanation based on a simplified example of AS Hegemony applied to `n`-paths in Fig. 5. Hegemony scores range from 0 to 1 and roughly represent the ratio of paths to the target AS that include the transit node. Since a `n`-path contains IXP identifiers, Hegemony scores are not only assigned to ASes but IXPs as well. For example, in Fig. 5 there are four paths from the probes (P_{0-3}) to the target AS. Two of these paths each pass through nodes AS B and IXP B, resulting in a Hegemony score of 0.5. To prevent biased scores caused by an uneven probe distribution, the calculation dismisses nodes close to the vantage points, which is why nodes IXP A, AS A, and AS C-E have a score of 0. A transit node (AS *or* IXP) with a high Hegemony score is commonly referred to as a *dependency* for the target AS (e.g., AS B and IXP B in Fig. 5), and large transit nodes may have numerous *dependent* target ASes. Even though the peering LAN of an IXP can be the target of traceroutes, we limit the target nodes for our study to ASes and include IXPs only as transit nodes. We give a precise definition of *dependency* below.

Formally, AS Hegemony is mostly based on Betweenness Centrality. For a graph $G = (V, E)$ composed of a set of nodes V and edges E, the betweenness centrality is defined as

$$BC(v) = \frac{1}{S} \sum_{u,w \in V} \sigma_{uw}(v) \tag{1}$$

where $\sigma_{uw}(v)$ is the number of paths from u to w passing through v, and S is the total number of paths. AS Hegemony adapts this method to make it more robust against sampling error incurred by the limited number of vantage points from which the graph is created. For this study we build one graph for each target AS t that consists only of `n`-paths towards t (similar to the local graphs of [36]). Graphs constructed this way are useful to study the dependency of the target on other nodes, which is the focus of our study.

Hence, we slightly modify the notation of the original work and define the Hegemony score of a node v in the graph of t as

$$\mathcal{H}_t(v, \alpha) = \frac{1}{n - (2\lfloor \alpha n \rfloor)} \sum_{j=\lfloor \alpha n \rfloor + 1}^{n - \lfloor \alpha n \rfloor} BC_{(j)}(v) \tag{2}$$

where

$$BC_{(j)}(v) = \frac{1}{S} \sum_{w \in V} \sigma_{jw}(v) \tag{3}$$

is the BC value computed with paths from only one viewpoint j, n is the total number of viewpoints, $\lfloor . \rfloor$ is the floor function, and $0 \leq 2\alpha < 1$ is the ratio of disregarded viewpoints such that the top and bottom $\lfloor \alpha n \rfloor$ viewpoints with the highest/lowest number of paths passing through the node are ignored. As recommended in the original work, we set α to 0.1.

Note that the change in notation — $\mathcal{H}_t$ compared to $\mathcal{H}$ in the original work — only indicates the difference in paths that are used for the calculation. In the original work, $\mathcal{H}$ is used in the context of *global* graphs, which include all paths from vantage points to *any* target AS, whereas $\mathcal{H}_t$ only includes paths going towards AS t. The computation is the same, only the input is different.

Definition: In accordance with past work [37], we consider dependencies with a Hegemony score lower than 0.1 as marginal and only report dependencies with a higher score. Therefore, for the purposes of this work, a node t has a *dependency* on node v if $\mathcal{H}_t(v) \geq 0.1$.

Table 3. The ten ASes with most dependents in BGP compared to traceroute and Spearman's rank correlation for all AS ranks. The rank correlation is calculated between BGP and the respective traceroute dataset.

		Rank			**Dependencies**		
ASN	Name	BGP	Ark	Atlas	BGP	Ark	Atlas
3356	Level 3	1	= 1	= 1	27,026	28,886	23,811
1299	Arelion	2	= 2	= 2	21,866	26,308	18,805
174	Cogent	3	= 3	= 3	19,394	19,405	13,334
6939	Hurricane Electric	4	= 4	= 4	18,883	13,636	13,035
2914	NTT-GIN	5	↓ 7	↓ 10	5819	3615	2904
3257	GTT	6	= 6	↓ 7	5456	4794	3705
6461	Zayo	7	↑ 5	↓ 8	4905	6129	3568
9002	RETN	8	↓ 10	↓ 9	3859	2777	3029
6453	Tata Communications	9	= 9	↓ 11	3345	3161	2573
52320	GlobeNet	10	↓ 12	↓ 14	2890	2375	2012
					Spearman's ρ for all ranks	0.86	0.84

As explained in Section 4.1, in our datasets the number of probes (i.e., vantage points) per target AS may vary and a low number of probes is likely to produce uncertain Hegemony scores. For this reason we avoid unreliable Hegemony scores by studying only scores that have been computed from paths collected by probes from at least 10 different ASes.

Finally, all reported dependencies should be treated as estimates based on the traceroutes captured within the specified interval. However, as described in Section 4.1 we confirmed that results are stable over time at least in close temporal proximity. Furthermore a dependency on an upstream system is based on the visible path usage within the interval. It does not imply that a failure of the upstream leads to a complete disconnect of the dependents as there might be backup connections available that are normally on standby and thus not visible in traceroute. Nonetheless, as motivated in Section 2 we expect that an outage at the upstream has some negative impact on its dependents.

5 BGP Cross-Validation

To cross-validate our methodology with previous work and get an understanding of how dependencies derived from traceroute data compare to the ones computed from BGP data, we inspect the AS dependencies from both approaches. We fetched BGP-based Hegemony results from the publicly available archive [6] for the timestamp `2022-09-19T00:00`.

The BGP dataset contains 262,817 dependencies on 10,243 transit ASes, whereas using traceroute data from Ark (Atlas) we found 227,136 (180,480) dependencies on 6970 (6145) transit ASes. This difference stems mainly from our

Table 4. Dependency Categories

	Ark			Atlas		
Type	Category	Count	Pct.	Category	Count	Pct.
AS	Low	3876	55.61 %	Low	3440	55.98 %
	Medium	2985	42.83 %	Medium	2616	42.57 %
	High	109	1.56 %	High	89	1.45 %
	Total	6970	100 %	Total	6145	100 %
IXP	Low	114	53.52 %	Low	107	48.86 %
	Medium	94	44.13 %	Medium	104	47.49 %
	High	5	2.35 %	High	8	3.65 %
	Total	213	100 %	Total	219	100 %

conservative thresholds on the minimum number of probe ASes, and because some routers do not reply to traceroute probes. Comparing the top ten transit ASes with the largest number of dependents in Table 3 shows that the rankings derived from BGP and traceroute are largely the same, with slight differences in the absolute number of dependents. If we inspect the top 100 (1000) ASes, we see an overlap of 80 % (80 %) with an average difference of 54 (28) dependents for ASes that are contained in both BGP and traceroute for Ark. The overlap is the same for Atlas (81 % and 79 %) although the average difference of dependents increases slightly (95 and 50). Finally, we compare the overall order of AS rankings using Spearman's rank correlation. There are 6420 ASes with dependents in both BGP and Ark, and 5704 ASes for BGP and Atlas. A correlation of $\rho = 0.86$ for BGP/Ark and $\rho = 0.84$ for BGP/Atlas reveals that there is a strong correlation between the order of ASes in terms of number of dependents.

There can be small differences between the control (BGP) and data (traceroute) plane [30], as well as the aforementioned caveats of traceroute. Overall, however, the sets of ASes and their order are very similar and although the number of dependents is not exactly the same, the approach of computing AS Hegemony on traceroute data produces AS dependencies similar to past work. Thus, we expect the IXP dependency results to be sensible.

6 Comparing transit ASes and IXPs

We now focus only on the results we have obtained with the traceroute datasets and put the role of IXPs in perspective by comparing dependency results of IXPs and transit ASes. For this comparison we look at the distribution of two properties: (1) The *number* of dependents relying on ASes and IXPs respectively and (2) the *strength* of the dependencies as indicated by the Hegemony score. For Ark (Atlas) there are 6970 (6145) ASes and 213 (219) IXPs with dependents in our results (see "Total" rows in Table 4), corresponding to ≈ 61 % (54 %) of

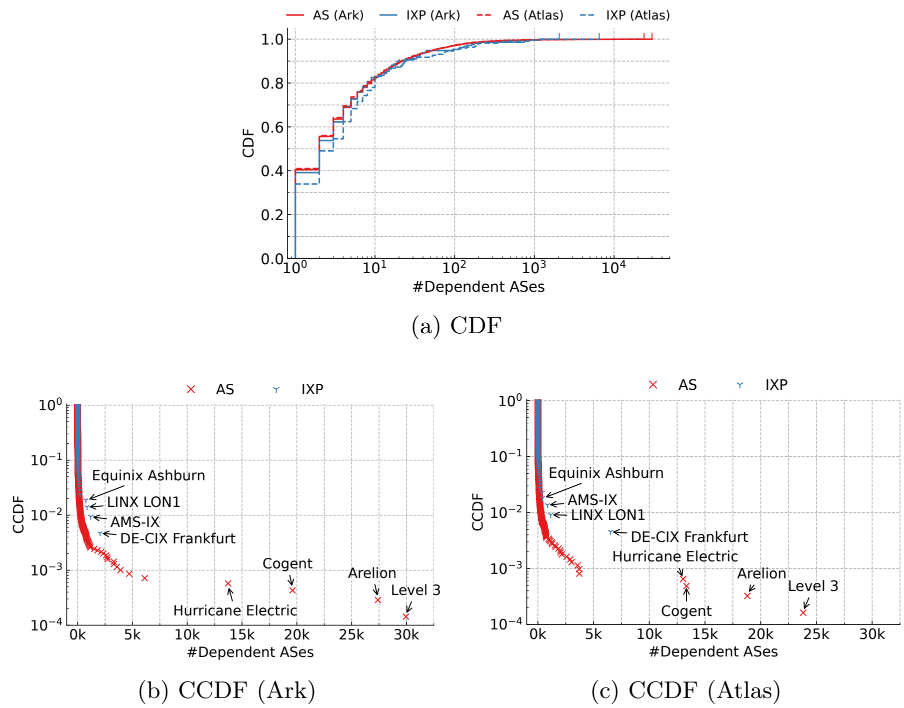

Fig. 6. Dependency distributions for 6970 (6145) ASes and 213 (219) IXPs. (a): The relative distributions are similar for ASes and IXPs. (b) & (c): There are some large ASes with significantly more dependents. This observation is present, although less pronounced, for IXPs as well, with four large IXPs standing out.

the 11,439 transit ASes[5] and ≈ 27 % of the 803 IXPs with at least two members listed in PeeringDB, which is a first indication that networks are relying more on transit ASes rather than IXPs to reach the wide Internet.

The overall distribution of the number of dependents is shown in Fig. 6a. Even though the number of points in each CDF differ by one order of magnitude, the shapes are remarkably similar. Around 40 % of ASes and IXPs have only a single dependent, and around 80 % have less than 10, meaning that the vast majority of networks have rather small number of dependents. However, the tail of the distribution reveals some ASes and IXPs with a large number of dependents (highlighted in Figs. 6b and 6c). In case of ASes, four large Tier 1 providers have several times more dependents than the remaining distribution, emphasizing the still important role of these networks in the Internet. This effect is also present for IXPs, with four IXPs standing out. From these numbers one may think that these large IXPs have similar properties to the Tier 1 providers, but this is not the case, as Hegemony scores for IXPs are significantly lower.

[5] Number of ASes with a customer cone > 1 according to [3].

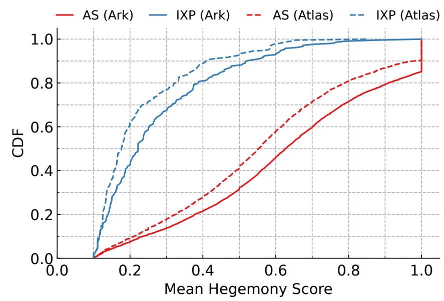

(a) Hegemony Distribution (CDF)

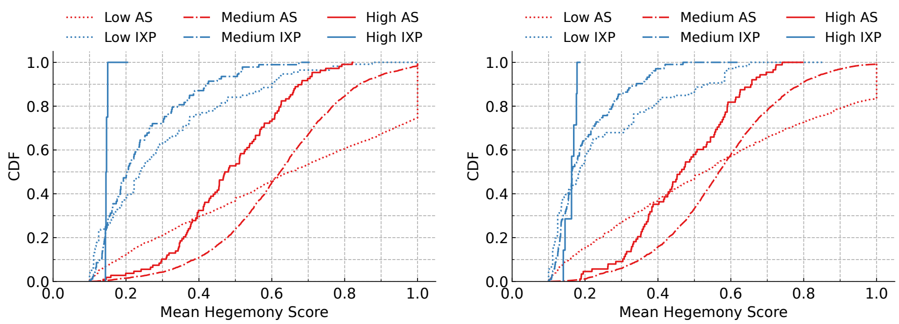

(b) Hegemony Distribution (by Size) Ark (c) Hegemony Distribution (by Size) Atlas

Fig. 7. Hegemony distributions for 6970 (6145) ASes and 213 (219) IXPs. (a): IXPs have weaker dependencies than ASes. (b) & (c): Splitting the CDF by number of dependents (Low ≤ 2; Medium ≤ 180; High > 180) reveals that strong dependencies are concentrated in *Low* ASes/IXPs.

The distribution of mean Hegemony scores in Fig. 7a highlights an important disparity between the strength of the dependencies between ASes and IXPs. The dependencies on IXPs are a lot weaker, with a median value of 0.22 (0.18) compared to 0.63 (0.55) for ASes. This means that a smaller number of paths are going through IXPs, which reveals that IXPs are usually used for a limited number of destinations as opposed to transit ASes that provide routes to the whole Internet. This agrees with the commonly accepted role of IXPs in the topology (see Fig. 2). In particular, 14.8 % (9.7 %) of ASes have a mean Hegemony score of 1, indicating that all their dependents rely fully on them (e.g., a single homed network with all its paths passing through a single upstream AS).

To understand the differences between these distributions in more detail, both ASes and IXPs are classified into categories according to their number of dependents. We take inspiration from [32] and define three categories: *Low* (≤ 2 dependents), *Medium* (≤ 180), and *High* (> 180). Therefore, in the following, the terms *Low IXP* (*High IXP*) refer to an IXP with a low (high) number of dependents. We use this term to distinguish from *Small IXP* (*Large IXP*), which is an IXP with a small (large) number of *members*. The resulting category sizes are shown in Table 4; The relative distribution of ASes is almost the same for Ark and Atlas, whereas a slight trend towards *Low* IXPs is visible for Ark.

The distributions of the six categories are shown in Figs. 7b and 7c which clearly show that the ASes and IXPs with a mean Hegemony score of 1 are almost exclusively those having only a single or two dependents. Conversely, an increasing number of dependents corresponds to weaker average Hegemony scores. This is also expected, as *Low* ASes are usually small Internet service providers (ISPs) with a few local single-homed customers, whereas larger ASes are likely to serve other large multi-homed ASes. Although the strength of dependencies is overall weaker for IXPs, we can assume the same is also true for IXPs.

In summary, while transit ASes have generally more dependents, there are some IXPs that separate themselves from the rest, possibly indicating a progression from their traditional role. However, the strength of dependencies on IXPs still remains relatively low, especially compared to transit ASes.

7 IXP Characterization

We now focus only on IXPs and characterize them based on two major factors: Their *topological* and *geographical* footprints. The topological footprint describes the role of an IXP in the Internet topology, based on its number of dependents. The geographical footprint refers to the physical location of the IXP and its dependents.

First, we take a look at the visibility of IXPs in traceroute (Section 7.1) to clarify the relationship between the presence of IXPs in traceroute and their number of dependent networks. Then, we inspect the connection between the number of members and dependents (Section 7.2). We comment on the general geographic footprint of IXPs (Section 7.3) and finally introduce a metric called *Regionality* (Section 7.4) to quantify the relationship between the location of an IXP and its members or dependents.

Since the results of the Ark and Atlas datasets are very similar, we limit the discussion and plots from now on only to the results obtained with the Ark dataset.

7.1 Topology: IXP Visibility in traceroute

The study of IXP dependencies based on traceroute data comes with the limitation that we can only infer dependencies for an IXP if we see it in traceroute. This raises two questions: (1) What is the general visibility of IXPs and their members in our data and (2) does the visibility of an IXP imply the presence of dependents?

We see that a majority of IXPs is visible and that there is a relationship between the size of the IXP and its visibility. In addition, the majority of visible IXPs does not have any dependents. We found that over 57 % of IXPs with at

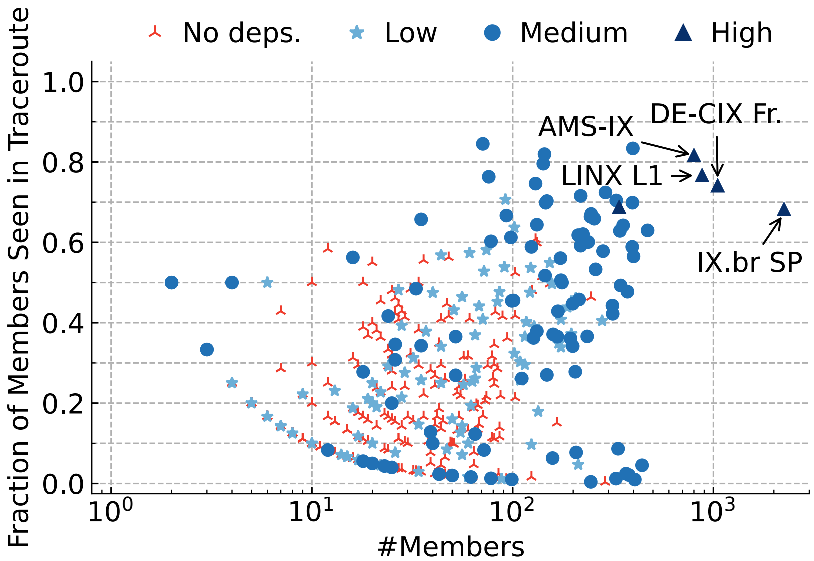

(a) IXP Visiblity in Traceroute

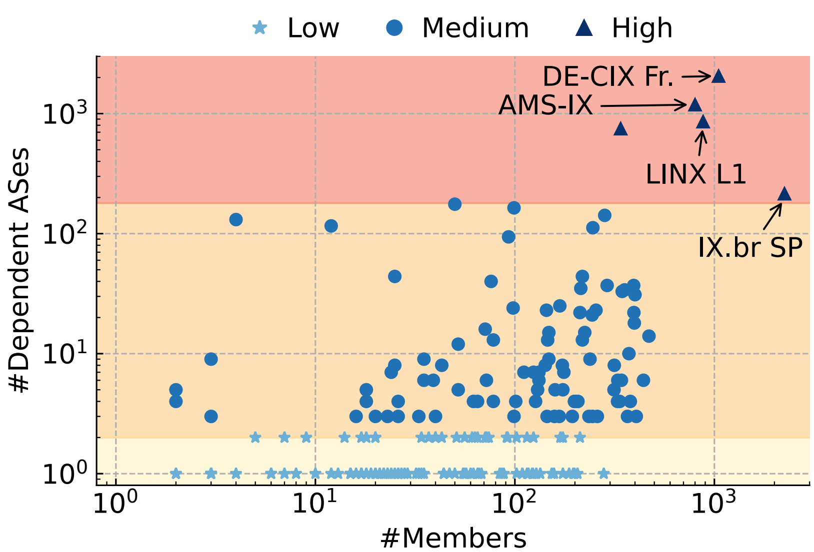

(b) IXP Dependents vs. Members

Fig. 8. Data for 205 IXPs with dependents and 259 IXPs without. Excludes eight IXPs with dependents but no members seen. (a): Large IXPs with many members enjoy good visibility. (b): The number of members and dependents are not strongly connected.

least two members in PeeringDB are visible in our dataset, i.e., we see traceroute responses from within their peering LAN.

To investigate the relationship between the size of an IXP (in terms of its members) and its visibility, we count for each IXP how many members are present in traceroute. For a better comparison of IXPs with different sizes, we compute the number of members seen in traceroute as a fraction. Fig. 8a shows the number of members per IXP against the fraction of members seen in traceroute. We can see that an increasing number of members usually translates into an increase in the fraction of visible members. In particular, large IXPs with around a thousand members have at least 68 % of members visible in traceroute.

Next, to answer if all visible IXPs have dependents, we visualize IXPs with different symbols, based on their dependency category (Table 4). There are 213 IXPs with dependents in our dataset, however an even larger number of 259 IXPs are visible, but have no dependents. These are mostly smaller IXPs with a median number of 33 members of which a median fraction of 0.22 are visible in traceroute. From this figure it is apparent that members at large IXPs have more paths going through the IXP peering LAN than members at smaller IXPs.

In summary, for large IXPs we usually observe numerous dependent ASes, but for the majority of IXPs a more expected pattern is observed: the members are visible in traceroute, however only in a limited number of paths which result in a marginal dependency.

7.2 Topology: Member and Dependent Relationship

To understand the relationship between the number of members and dependents better, we rearrange the points of Fig. 8a to show the absolute number of dependents in Fig. 8b. There is a general trend that seems to link the number of members to the number of dependents: The median number of members for *Low* IXPs is 40, followed by 146 for *Medium*, and 814 for *High* IXPs. This is intuitive: Each member of an IXP possibly connects to several more ASes and therefore increases the chance of incurring a dependency.

However, even though there is a trend, for almost all IXP sizes there are also examples of IXPs with low and medium number of dependents. In an extreme case, the largest IXP in terms of members is `IX.br São Paulo` with 2246 members, but only 215 dependents. This is a clear difference to the likes of `DE-CIX Frankfurt` (1050 members; 2055 dependents) or `AMS-IX` (800 members; 1187 dependents). This reveals that the size of the IXP is not the only factor that drives the number of dependents and it also suggests that IXP may be used differently by their members, which is something we investigate more in Section 7.4.

7.3 Geography: Geographic Footprint of IXPs

Dependencies give us a better view of the topological footprint of an IXP in the Internet. We can also leverage these results to investigate the geographical footprint of IXPs and enhance information from PeeringDB by mapping the dependents of an IXP to countries. Precise geolocation is not trivial [56], which is why we employ a conservative approach to locate ASes. We take all announced prefixes of an AS and lookup their assigned country based on NRO stats files [48]. An AS is mapped to a country only if all its originated prefixes geolocate to the same country. If an AS announces prefixes that map to different countries, we do not geolocate it to a particular country. Instead, we treat this AS as an *international* network.

We find new countries for 48 IXPs, i.e., IXPs that have dependents in countries that are not already covered by their listed members. The most extreme case is `AMS-IX`, which already connects members from 63 different countries, but we see dependencies in 43 additional countries, increasing the geographical footprint by 68 %. The median increase for all 48 IXPs with additional countries is 35 %, revealing the hidden reach that is not obvious from simple membership data.

7.4 Geography: IXP Regionality

Based on the countries of members and dependents, we can also formulate a metric we call *Regionality*. It simply represents the fraction of members (or dependents) that are in the same country as the IXP. Regionality allows us to quickly asses, in a coarse manner, if an IXP is used to reach ASes from the local region, or if it is operating on an international scale.

First, we take a look at all IXPs and plot the regionality based on their members in Fig. 9a. Since there is a large overlap of points, we show the distribution of IXPs with green bars, labeled with a secondary y-axis on the right side. In general, regionality based on members is balanced, with a median value of 0.5. The distribution reveals that 50 IXPs (6 % of 803 IXPs with members) have a regionality of 1, i.e., all members are in the same country as the IXP. These IXPs are mostly small, with a median of four members. The largest IXP with a regionality of 1 is `UNY-IX`, which connects 20 members (mostly universities), all from Indonesia.

The regionality based on dependents, shown in Fig. 9b, is usually higher than the one based on members: 121 IXPs (57 % of 213 IXPs with dependents) have a regionality of 1. If we consider both IXPs members *and* dependents, still 6 IXPs (3 %) are entirely local, and 111 IXPs (52 %) have international members, but only national dependents. Overall, there is a surprising balance of national and international members, but usually ASes that depend on an IXP are solely operating in the same region as the IXP.

However, we observe interesting differences when comparing the regionality of IXPs with a high number of dependents. For four *High* IXPs there is a clear shift towards international dependents. `DE-CIX Frankfurt` with 1050 members and a regionality of 0.24 moves to 2055 dependents with a regionality of 0.07. A similar trend is visible for `LINX LON1` (878 members, 0.31 regionality → 857 dependencies, 0.15 regionality), `AMS-IX` (800, 0.13 → 1185, 0.07), and `Equinix Ashburn` (339, 0.44 → 749, 0.09). One exception is `IX.br São Paulo`, which has the same regionality of 0.89 for both its 2246 members and 215 dependents.

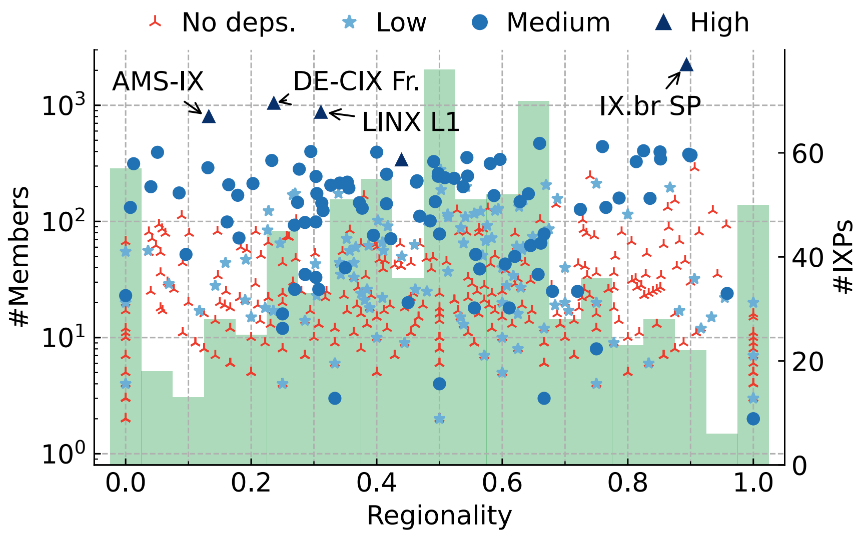

(a) Member Regionality

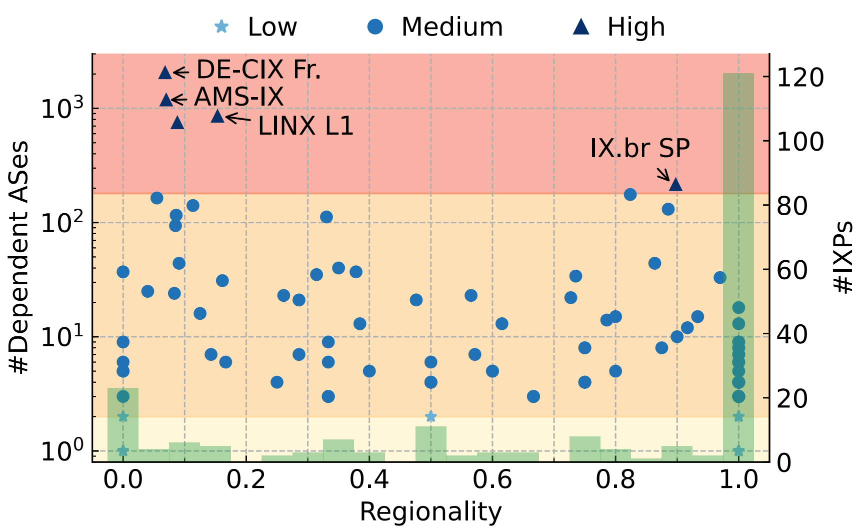

(b) Dependents Regionality

Fig. 9. Regionality distribution of IXPs based on the location of (a) their members, and (b) their dependents. Green bars show the distribution of points, labeled with the right y-axis. Contains data for (a) 803 IXPs with at least two members and (b) 213 IXPs with dependents.

Table 5. Properties of DE-CIX Frankfurt and IX.br São Paulo.

IXP	**DE-CIX Frankfurt**	**IX.br São Paulo**
Members	1050	2246
... visible in traceroute	778	1532
Dependents	2055	215
Member regionality	0.24	0.89
Dependent regionality	0.07	0.89
Countries with dependents	94	16

It is also worth mentioning that not all large European IXPs are extremely international and have many dependents. For example, `EPIX.Katowice` in Poland has 406 members with a regionality of 0.83, but only three dependents. In the grand scheme, very large IXPs are rather the exception than the norm.

8 A Case Study: `DE-CIX` and `IX.br`

In this final section of the analysis we take a closer look at two IXPs and show that numerous dependencies from the same country can be connected to an IXP by a single member. We choose two large IXPs in terms of members: `DE-CIX Frankfurt` (1050 members) and `IX.br São Paulo` (2246 members). Both enjoy good visibility in traceroute, with 778 and 1532 members present for `DE-CIX` and `IX.br` respectively. Table 5 shows additional properties of the IXPs.

Even though `IX.br` has 47 % more members, it has 89 % less dependencies. This may be an indicator that members are using `IX.br` mostly for local connectivity. This is expected as IXPs emerged to reduce hierarchies and connect physically close networks and `IX.br` seems to fulfill this purpose.

To confirm if an IXP provides more regional connectivity, or if it developed into an international transit hub, a simple idea is to look at the countries of the members located at an IXP. Doing this for our two examples reveals that `DE-CIX` connects networks from 68 countries, with 23 % of members originating from Germany, whereas `IX.br` is more local, connecting 20 countries and a majority of 89 % of members from Brazil. However, by taking into account the countries of dependents as well, we can expand our view to reveal connections that are deeper in the topology. For `DE-CIX` this reveals 442 dependents in 34 countries not seen in neither PeeringDB nor the looking glass, whereas `IX.br` only has 11 dependents in nine new countries.

Dependencies also provide a different view on the countries reached via the IXPs. Fig. 10 shows the ten countries with the most dependents for `DE-CIX` and `IX.br`. The country code `**` is introduced by our conservative geolocation: It represents international ASes that can not be assigned to a single country. For `DE-CIX` (Fig. 10a) the country with the most dependents is Russia (331 dependents), even though only 36 members are located there. Even more pronounced is India: There is no member that we can exclusively geolocate to India, but there

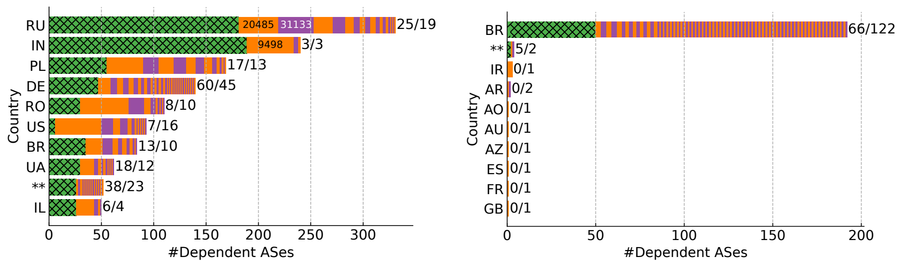

(a) DE-CIX Frankfurt (b) IX.br São Paulo

Fig. 10. Number of dependent ASes per country for two large IXPs. The green hatched bars show dependents that are shared between members, whereas orange and purple bars represent one member each. The label shows the number of members that share dependents and the number of singular members (i.e., orange and purple bars). The country code `**` marks dependents that can not be attributed to a single country. Not shown are 723 dependents from 85 countries for (a) and seven dependents from seven countries for (b).

are 240 dependents. The use of dependencies reveals that even though `DE-CIX` is located in Germany, it is a common mean to reach many Indian ASes. In contrast, looking at the dependency locations of `IX.br` (Fig. 10b) only reinforces the impression that it is used for regional connectivity: 89 % of dependents are in Brazil, the same country as the IXP.

In the final step, we characterize individual members by analyzing through which members the dependents are reached, hence discovering if there are major networks at IXPs responsible for many dependencies. The orange and purple color bars in Fig. 10 represent one member each. The green hatched bars show dependent ASes that are reached via more than one member. Shared dependencies are caused by traceroutes to different prefixes of an AS that are forwarded over separate members of the IXP. Next to each bar is a label that describes the number of members that share dependents, and the number of members with unique dependencies. For example, for Russian dependents at `DE-CIX` there are 25 members that share 181 dependents and 19 members with a total of 150 dependents.

Coming back to the Indian dependents of `DE-CIX`, we can now see that 45 are reached over AS9498/Bharti Airtel, which is the only member registered in India. Another observation is that the majority of dependents is connected by multiple members. However, AS9498 is present for all 189 shared dependents as well. Out of these, 188 are shared with AS6461/Zayo, which does not announce prefixes registered in India, but is a large Tier 1 provider. There is one dependent additionally shared with AS4637/Telstra Global.

For Russia, there are also a number of ASes with presence in Russia (AS20485 / TransTeleCom, AS31133 / MegaFon) connecting 38 and 34 dependents each. In contrast to the Indian dependents, however, there is no single member respon-

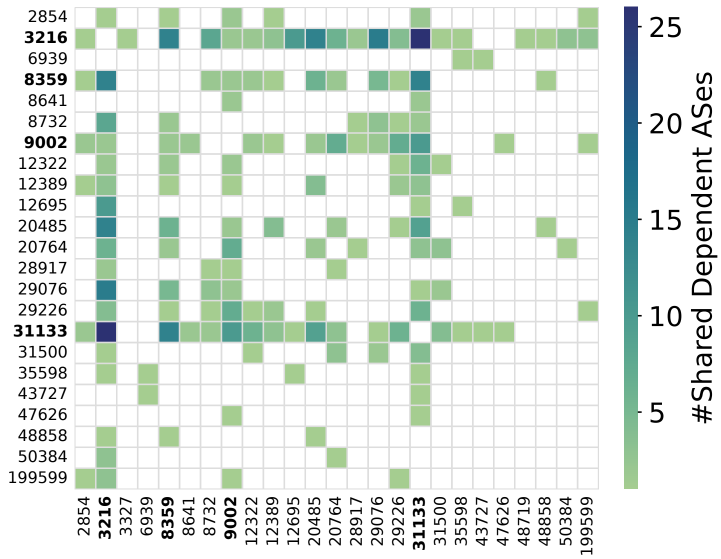

(a) DE-CIX Frankfurt, Russia

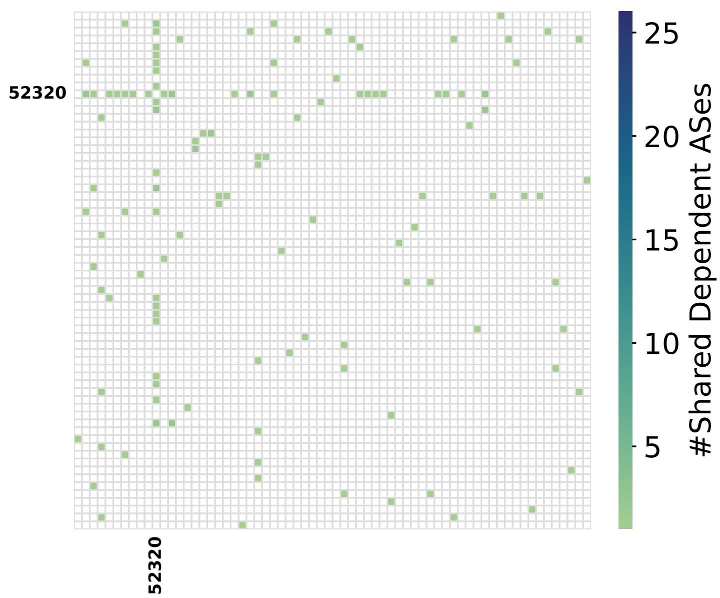

(b) IX.br São Paulo, Brazil

Fig. 11. Detailed view of member ASes with shared dependents located in (a) Russia for DE-CIX Frankfurt and (b) Brazil for IX.br São Paulo. Members that share dependents with at least ten other members are highlighted. To reduce visual noise only highlighted members are shown in (b).

sible for most dependents, which is also apparent when looking at the shared dependents.

While there are only three members that share 189 dependents in India, the 181 dependents in Russia are shared between 25 members. Fig. 11a visualizes the pairings of the individual members. Colored cells indicate shared dependents between a pair of members and the intensity of the color signifies the amount. For example, member AS3216 shares 26 dependents with AS31133. The sparseness of the heatmap shows that there are almost no members contributing to a significant number of dependents. Two exceptions are AS3216/Vimpelcom and AS31133/MegaFon, which are present in 20 and 17 pairs incurring a total of 92 and 66 dependents each. Still, the shared dependents are well distributed when compared to the Indian case.

We see that these "far away" dependencies are reached through a comparatively small number of members, however, the diversity of local members is visible: For Germany (Fig. 10a) we observe less dependents than members, and there are no members with a large number of dependents. Indeed, the median number of dependents for the 45 members with dependents in Germany is 1.

With this in mind, we can dissect the Brazilian dependents of `IX.br` and see that the number of dependents is a lot smaller than the number of members and over 87 % of the 122 members with dependents only lead to a single dependent. This observation also holds for shared dependencies as shown in Fig. 11b. 50 dependents are shared between 66 members and 64 % of members only have a single shared dependent. The absence of any major member is also apparent, as 88 % share dependents with only one or two other members. The only exception is AS52320/GlobeNet, which shares 24 dependents with 20 other members and thus incurs the most dependencies out of all members. Compared to `DE-CIX`, however, dependents are still distributed rather well.

The comparison of two large IXPs gave us several insights: First, regional connectivity is distributed well, with less dependencies than members, and no single member incurring a large number of dependencies. Second, even large IXPs are nowadays still used to "keep local traffic local", as seen at the example of `IX.br`. Third, by looking at the number and locations of dependents we get a better understanding of how IXPs are used for connectivity in today's Internet and found that some large IXPs may be used as international hubs.

9 Limitations

In this section we summarize the general limitations of this study, as well as specific limitations bound to our design choices and their impact on the results.

First, we do not have a perfect view of the Internet topology. As discussed in Section 3, inferring the Internet topology from BGP or traceroute data is always deemed to provide a partial view of the Internet. Reasons for this include the diverse routing policies employed on the Internet and placement of vantage points [40]. Increasing the reach of measurement platforms and improving the distribution of vantage points is an ongoing research topic [17, 18, 51], any advances in that topic would improve this work as well. Hence, the reported dependencies should be interpreted as an *approximation* of the real world. Furthermore, by definition dependencies are derived from the most frequent routes so we expect strong dependencies to be observed even in a partial view of the Internet. Our approach can miss dependencies due to lack of visibility, nevertheless dependencies reported in this study are based on actual observations and should therefore represent a sample of the real world.

Using traceroute data to infer `n`-paths has some common pitfalls [42, 45, 57] including some related to the identification of IXPs. Some routers might be configured to not respond to traceroute probes, others might not respond from the address of the interface on which the probe was received. Both cases are particularly relevant if the affected router connects to the peering LAN of an IXP. As mentioned above, these are factors that can lead to missed dependencies. However, as discussed in Section 5 we expect the results to be reasonable, based on the comparison with existing approaches.

Finally, we obtain peering LAN and member information from PeeringDB. Although we do not expect every IXP to be listed in PeeringDB, we believe this is the most complete data source currently available. In addition, the peering policies of popular content providers require ISPs to maintain up-to-date PeeringDB entries [2, 5, 8, 10], indicating that it is currently the preferred source of information by the industry. We manually confirmed that the peering LAN information for the largest IXPs is accurate and improved the member listing via the route server looking glasses where feasible. It is still possible that we miss an IXP member if they are not listed in PeeringDB and do not peer with the route server. However, this data only affects the detailed analysis in Section 8 for which we have route server information available thus we expect a minimal impact.

In summary, the limitations of our approach generally lead to a reduced number of dependencies. The reported dependencies are an estimate based on available data but should not be interpreted as a ground-truth view of the topology. With these limitations in mind, we demonstrated in this paper that estimated dependencies are a useful tool to study the general importance of IXPs and ASes in the structure of the Internet.

10 Related Work

There are several other works that characterize or analyze IXPs from different points of view. Prehn et al. [49] perform an analysis of the reachability benefit gained by peering at large IXPs, based on route server snapshots as well as inferred routes for bi-lateral and private peering. Giotsas et al. [39] propose a system that detects outages at IXP facilities, relying on the observation of BGP communities, and as part of their evaluation they study the impact of an outage at `AMS-IX` in detail. Brito et al. [23] provide an in-depth study of the `IX.br` peering ecosystem in Brazil based on data obtained from looking glasses at all `IX.br` IXPs.

Most closely related to this work, Bertholdo et al. [22] measure the coverage and representativeness of IXPs as part of their effort to forecast the impact of IXP outages. The focus of our work is not to detect outages, but more generally to provide insights about IXP usage and a long-term dataset of IXP dependency. Bertholdo et al. rely on active measurements and anycast sites deployed at large IXPs for the time of their study, whereas our results are continuously updated, as long as the measurement platforms stay active. Furthermore, they map out IXP reachability using ICMP ping replies, which enables the answer of a yes/no query, but loses path information that the traceroute data in our study retains.

11 Conclusion & Future Work

In this paper, we presented a study on IXP dependency. We gained further insights into the topological footprint of IXPs, showing that some IXPs have a number of additional dependents that are not direct members. We also investigated the geographical footprint of IXPs and found that some of them are indirectly used by ASes in many more countries than what can be inferred from PeeringDB or looking glasses. In addition, we observed that many IXPs are local, connecting to a large degree, or even exclusively, networks from the same country. Finally, for large IXPs we identified specific members that incur a large number of international dependents.

The investigation of IXP dependency opens a new avenue of research that we aim to explore in future work. An additional degree of detail can be achieved by computing per-interface dependency, that may yield further insights into how large members with multiple interfaces operate at IXPs. A different direction is the investigation of country dependencies, i.e., the role of IXPs from the viewpoint of a country. In today's geopolitical climate even internet infrastructure might become a target, making the knowledge of dependency especially valuable.

Public Dataset and Reproducibility To empower network administrators and facilitate future research, we periodically update our results and make them publicly available at https://internethealthreport.github.io/ixp-dependency. On this website, we also publish the analysis code and data to aid in reproducibility of our work. Since all our results are based on open data, users can also run their own analysis on different time windows if they desire.

Finally, we expanded the data pipeline used to analyze the IXP outages in Section 2 to allow a general analysis of any AS or IP prefix and publish the code in a separate repository [52].

Acknowledgments We thank the reviewers for their feedback. We would also like to thank Emile Aben from the RIPE NCC for fruitful discussions and CAIDA for giving us access to their dataset.

References

1. Alice-LG - Your friendly looking glass, https://github.com/alice-lg/alice-lg
2. AWS: Settlement Free Peering Policy, https://aws.amazon.com/peering/policy/
3. CAIDA AS Rank, http://as-rank.caida.org/
4. The CAIDA UCSD IPv4 Routed /24 Topology Dataset - 2022-09-01 – 2022-09-30, https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/
5. Cloudflare Peering Policy, https://www.cloudflare.com/peering-policy/
6. IHR Archive, https://ihr-archive.iijlab.net/
7. LINX LON1 Outage – March 2021, https://web.archive.org/web/20221206083420/https://www.linx.net/incidents-log/
8. Peering with Meta, https://www.facebook.com/peering/
9. PeeringDB, https://www.peeringdb.com/
10. Prerequisites to Peer with Google, https://peering.google.com/#/options/peering
11. University of Oregon Route Views Project, http://www.routeviews.org/
12. Follow-up on previous incident at AMS-IX platform (May 2015), https://web.archive.org/web/20160327075404/https://ams-ix.net/newsitems/195
13. Outage on Amsterdam peering platform (Nov 2023), https://www.ams-ix.net/ams/outage-on-amsterdam-peering-platform
14. Aben, E.: Measuring More Internet with RIPE Atlas (Jan 2016), https://labs.ripe.net/author/emileaben/measuring-more-internet-with-ripe-atlas/
15. Aben, E.: Does The Internet Route Around Damage? - Edition 2021 (Apr 2021), https://labs.ripe.net/author/emileaben/does-the-internet-route-around-damage-edition-2021/
16. Ahmed, A., Shafiq, Z., Bedi, H., Khakpour, A.: Peering vs. Transit: Performance Comparison of Peering and Transit Interconnections. In: International Conference on Network Protocols (ICNP). pp. 1–10. IEEE (2017). https://doi.org/10.1109/ICNP.2017.8117549
17. Alfroy, T., Holterbach, T., Pelsser, C.: MVP: Measuring Internet Routing from the Most Valuable Points. In: Internet Measurement Conference (IMC). pp. 770–771. ACM (2022). https://doi.org/10.1145/3517745.3563031
18. Appel, M., Aben, E., Fontugne, R.: Metis: Better Atlas Vantage Point Selection for Everyone. In: Network Traffic Measurement and Analysis Conference (TMA). IFIP (2022)
19. Arnold, T., He, J., Jiang, W., Calder, M., Cunha, I., Giotsas, V., Katz-Bassett, E.: Cloud Provider Connectivity in the Flat Internet. In: Internet Measurement Conference (IMC). pp. 230–246. ACM (2020). https://doi.org/10.1145/3419394.3423613
20. Bajpai, V., Eravuchira, S.J., Schönwälder, J.: Lessons Learned from using the RIPE Atlas Platform for Measurement Research. ACM SIGCOMM Computer Communication Review **45**(3), 35–42 (Jul 2015). https://doi.org/10.1145/2805789.2805796
21. Bajpai, V., Eravuchira, S.J., Schönwälder, J., Kisteleki, R., Aben, E.: Vantage Point Selection for IPv6 Measurements: Benefits and Limitations of RIPE Atlas Tags. In: IFIP/IEEE Symposium on Integrated Network and Service Management (IM). pp. 37–44. IEEE (2017). https://doi.org/10.23919/INM.2017.7987262
22. Bertholdo, L.M., Ceron, J.M., Granville, L.Z., van Rijswijk-Deij, R.M.: Forecasting the Impact of IXP Outages Using Anycast. In: Network Traffic Measurement and Analysis Conference (TMA). IFIP (2021)
23. Brito, S.H.B., Santos, M.A.S., Fontes, R.d.R., Perez, D.A.L., Rothenberg, C.E.: Dissecting the Largest National Ecosystem of Public Internet eXchange Points

in Brazil. In: Passive and Active Measurement Conference (PAM). pp. 333–345. Springer (2016). https://doi.org/10.1007/978-3-319-30505-9_25
24. Böttger, T., Antichi, G., Fernandes, E.L., di Lallo, R., Bruyere, M., Uhlig, S., Tyson, G., Castro, I.: Shaping the Internet: 10 Years of IXP Growth (2018). https://doi.org/10.48550/ARXIV.1810.10963
25. Böttger, T., Cuadrado, F., Tyson, G., Castro, I., Uhlig, S.: Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN. ACM SIGCOMM Computer Communication Review **48**(1), 28–34 (Apr 2018). https://doi.org/10.1145/3211852.3211857
26. Chang, H., Jamin, S., Willinger, W.: Inferring AS-level Internet Topology from Router-Level Path Traces. In: Scalability and Traffic Control in IP Networks. vol. 4526, pp. 196–207. SPIE (2001). https://doi.org/10.1117/12.434395
27. Chatzis, N., Smaragdakis, G., Feldmann, A., Willinger, W.: There is More to IXPs than Meets the Eye. ACM SIGCOMM Computer Communication Review **43**(5), 19–28 (Nov 2013). https://doi.org/10.1145/2541468.2541473
28. Cho, S., Fontugne, R., Cho, K., Dainotti, A., Gill, P.: BGP hijacking classification. In: Network Traffic Measurement and Analysis Conference (TMA). pp. 25–32. IEEE (2019). https://doi.org/10.23919/TMA.2019.8784511
29. Dang, T.K., Mohan, N., Corneo, L., Zavodovski, A., Ott, J., Kangasharju, J.: Cloudy with a Chance of Short RTTs: Analyzing Cloud Connectivity in the Internet. In: Internet Measurement Conference (IMC). pp. 62–79. ACM (2021). https://doi.org/10.1145/3487552.3487854
30. Del Fiore, J.M., Merindol, P., Persico, V., Pelsser, C., Pescapé, A.: Filtering the Noise to Reveal Inter-Domain Lies. In: Network Traffic Measurement and Analysis Conference (TMA). pp. 17–24. IEEE (2019). https://doi.org/10.23919/TMA.2019.8784618
31. Dhamdhere, A., Dovrolis, C.: The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In: International Conference on emerging Networking EXperiments and Technologies (CoNEXT). pp. 185–198. ACM (2010). https://doi.org/10.1145/1921168.1921196
32. Dhamdhere, A., Dovrolis, C.: Twelve Years in the Evolution of the Internet Ecosystem. IEEE/ACM Transactions on Networking **19**(5), 1420–1433 (Oct 2011). https://doi.org/10.1109/TNET.2011.2119327
33. Du, B., Testart, C., Fontugne, R., Akiwate, G., Snoeren, A.C., kc claffy: Mind your MANRS: Measuring the MANRS Ecosystem. In: Internet Measurement Conference (IMC). pp. 716–729. ACM (2022). https://doi.org/10.1145/3517745.3561419
34. Fontugne, R., Ermoshina, K., Aben, E.: The Internet in Crimea: a Case Study on Routing Interregnum. In: IFIP Networking Conference (Networking). pp. 809–814. IEEE (2020)
35. Fontugne, R., Pelsser, C., Aben, E., Bush, R.: Pinpointing Delay and Forwarding Anomalies Using Large-Scale Traceroute Measurements. In: Internet Measurement Conference (IMC). pp. 15–28. ACM (2017). https://doi.org/10.1145/3131365.3131384
36. Fontugne, R., Shah, A., Aben, E.: The (Thin) Bridges of AS Connectivity: Measuring Dependency Using AS Hegemony. In: Passive and Active Measurement Conference (PAM). pp. 216–227. Springer (2018). https://doi.org/10.1007/978-3-319-76481-8_16
37. Gamero-Garrido, A.: Transit Influence of Autonomous Systems: Country-Specific Exposure of Internet Traffic. Ph.D. thesis, University of California, San Diego, USA (2021), https://www.escholarship.org/uc/item/0hg720zn

38. Gamero-Garrido, A., Carisimo, E., Hao, S., Huffaker, B., Snoeren, A.C., Dainotti, A.: Quantifying Nations' Exposure to Traffic Observation and Selective Tampering. In: Passive and Active Measurement Conference (PAM). pp. 645–674. Springer (2022). https://doi.org/10.1007/978-3-030-98785-5_29
39. Giotsas, V., Dietzel, C., Smaragdakis, G., Feldmann, A., Berger, A., Aben, E.: Detecting Peering Infrastructure Outages in the Wild. In: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM). pp. 446–459. ACM (2017). https://doi.org/10.1145/3098822.3098855
40. Gregori, E., Improta, A., Lenzini, L., Rossi, L., Sani, L.: On the Incompleteness of the AS-level Graph: a Novel Methodology for BGP Route Collector Placement. In: Internet Measurement Conference (IMC). pp. 253–264. ACM (2012). https://doi.org/10.1145/2398776.2398803
41. Henthorn-Iwane, A.: Understanding Internet Exchanges via the DE-CIX Outage (Apr 2018), https://www.thousandeyes.com/blog/network-monitoring-de-cix-outage
42. Hyun, Y., Broido, A., kc claffy: On Third-party Addresses in Traceroute Paths. In: Passive and Active Network Measurement Workshop (PAM) (2003), https://catalog.caida.org/paper/2003_3rdparty
43. Hyun, Y., Broido, A., kc claffy: Traceroute and BGP AS Path Incongruities (2003), https://catalog.caida.org/paper/2003_asp
44. Lodhi, A., Larson, N., Dhamdhere, A., Dovrolis, C., kc claffy: Using PeeringDB to Understand the Peering Ecosystem. ACM SIGCOMM Computer Communication Review **44**(2), 20–27 (Apr 2014). https://doi.org/10.1145/2602204.2602208
45. Mao, Z.M., Rexford, J., Wang, J., Katz, R.H.: Towards an Accurate AS-Level Traceroute Tool. In: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM). pp. 365–378. ACM (2003). https://doi.org/10.1145/863955.863996
46. McQuistin, S., Uppu, S.P., Flores, M.: Taming Anycast in the Wild Internet. In: Internet Measurement Conference (IMC). pp. 165–178. ACM (2019). https://doi.org/10.1145/3355369.3355573
47. Nomikos, G., Dimitropoulos, X.: traIXroute: Detecting IXPs in traceroute paths. In: Passive and Active Measurement Conference (PAM). pp. 346–358. Springer (2016). https://doi.org/10.1007/978-3-319-30505-9_26
48. Number Resource Organisation: NRO Extended Allocation and Assignment Reports, https://www.nro.net/about/rirs/statistics/
49. Prehn, L., Lichtblau, F., Dietzel, C., Feldmann, A.: Peering Only? Analyzing the Reachability Benefits of Joining Large IXPs Today. In: Passive and Active Measurement Conference (PAM). pp. 338–366. Springer (2022). https://doi.org/10.1007/978-3-030-98785-5_15
50. RIPE NCC Staff: RIPE Atlas: A Global Internet Measurement Network. The Internet Protocol Journal **18**(3), 2–26 (Sep 2015)
51. Sermpezis, P., Prehn, L., Kostoglou, S., Flores, M., Vakali, A., Aben, E.: Bias in Internet Measurement Platforms. In: Network Traffic Measurement and Analysis Conference (TMA). pp. 1–10. IEEE (2023). https://doi.org/10.23919/TMA58422.2023.10198985
52. Tashiro, M.: Atlas traceroute outage inspector, https://github.com/m-appel/atlas-traceroute-outage-inspector
53. Tashiro, M., Fontugne, R., Fukuda, K.: Accompanying website and data, https://internethealthreport.github.io/ixp-dependency/

54. Wagner, D., Kopp, D., Wichtlhuber, M., Dietzel, C., Hohlfeld, O., Smaragdakis, G., Feldmann, A.: United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 970–987. ACM (2021). https://doi.org/10.1145/3460120.3485385
55. Wichtlhuber, M., Strehle, E., Kopp, D., Prepens, L., Stegmueller, S., Rubina, A., Dietzel, C., Hohlfeld, O.: IXP Scrubber: Learning from Blackholing Traffic for ML-Driven DDoS Detection at Scale. In: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM). pp. 707–722. ACM (2022). https://doi.org/10.1145/3544216.3544268
56. Winter, P., Padmanabhan, R., King, A., Dainotti, A.: Geo-locating BGP prefixes. In: Network Traffic Measurement and Analysis Conference (TMA). pp. 9–16. IEEE (2019). https://doi.org/10.23919/TMA.2019.8784509
57. Zhang, Y., Oliveira, R., Wang, Y., Su, S., Zhang, B., Bi, J., Zhang, H., Zhang, L.: A Framework to Quantify the Pitfalls of Using Traceroute in AS-Level Topology Measurement. IEEE Journal on Selected Areas in Communications **29**(9), 1822–1836 (Oct 2011). https://doi.org/10.1109/JSAC.2011.111007