Forecasting the Impact of IXP Outages Using Anycast

Leandro M. Bertholdo*, João M. Ceron†, Lisandro Z. Granville‡,
Roland van Rijswijk-Deij*§
*University of Twente, Enschede, The Netherlands
{l.m.bertholdo, r.m.vanrijswijk}@utwente.nl
†SIDN Labs, Arnhem, The Netherlands
{joao.ceron}@sidn.nl
‡Federal University of Rio Grande do Sul, Porto Alegre, Brazil
granville@inf.ufrgs.br
§NLnet Labs, Amsterdam, The Netherlands

***Abstract*—Internet eXchange Points (IXPs) play a major role in Internet connectivity. They provide an infrastructure for traffic exchange, attracting Internet Service Providers (ISPs), Content Delivery Networks (CDNs) and cloud providers, leading to a plethora of options for network operators to connect. Alongside, policymakers are interested in understanding how failures of these infrastructures could affect the cyber economy. In this paper, we propose a new method based on anycast to assess the importance of IXPs in terms of coverage and regional representativeness with the goal of predicting the impact of IXP failures. We deployed an anycast infrastructure to connect to major IXPs, simulate outages, and map the impact of such failures. Our experiments show that our methodology can predict how the traffic flows when one IXP goes down, helping operators get prepared to deal with these events, and add more information to better assess the importance of IXPs as a critical infrastructure.**

***Index Terms*—Internet Exchange Point, Internet Topology, Anycast Networks, Outage planning.**

I. Introduction

Internet eXchange Points (IXPs) are an essential element of the Internet and are becoming even more important as a flat Internet topology (free of Tier-X hierarchy) [1] is being chased by content providers [2]. IXPs contribute by providing direct peering between Internet Service Providers (ISPs), telcos, cable and mobile carriers, content providers, Web enterprises, governmental and financial services. A direct relationship between Autonomous Systems (ASes) is a way to reach better response times and decrease interconnection costs.

IXPs rearrange the connection matrix at a national, regional, and international level [1] [3]. With hundreds of IXPs over the world, it is relevant for network operators and policymakers to better know how they work, and especially what the impact is of a major IXP failure. For policymakers, it is important as an outage risk assessment on the country's cyber economy to know whether the remaining Internet infrastructure can support a long-term failure of an IXP. For network operators, it is relevant for capacity planning issues to see beyond the routing table and evaluate how traffic flows between IXPs.

978-3-903176-27-0

Although there are a few sources of information available to determine the importance of each IXP and how they relate to each other, an open issue is the absence of methods to estimate the impact of a failure in an IXP. Examples of critical outages are those that AMS-IX suffered in 2015, and DE-CIX in 2018 [4] - which broadly affected the Internet in Germany [5], or the IX.br/SP case in 2018, which impacted the Brazilian portion of the Internet [6] with users complaining about slowness and failures to open websites. It is not easy to forecast the impact of scenarios as described above. To be able to do so, it is important to better understand the relationships between IXPs as well as their coverage.

In this paper, we propose a method based on anycast active probing to enrich the information available about IXPs. We placed anycast sites connected to the 'Open Peering Policy' at each one of the major IXPs and used these to map the behavior of 6.5 million networks using different strategies. This approach enables to analyze IXPs in real situations, allowing us to propose metrics to evaluate, compare, and quantify the impact of an IXP outage. Anycast sites at IXPs provide us with a distinct view to better qualify and understand how the traffic flows through them, and it enables us to analyze participants' forwarding data planes and compare these to the IXP routing table itself. This method gives us a distinct view to answer our research questions:

- RQ1: How much of the Internet can be reached by connecting to the biggest IXPs?
- RQ2: How representative are the biggest IXPs inside the country/region/continent they are placed in?
- RQ3: How does the network traffic flow when major IXP failures take place?

As contributions, we present our method to forecast the impact of a major IXP outage and to obtain a new set of metrics to qualify IXPs in terms of coverage by traffic direction, IXP hegemony, and country/regional representativeness.

This paper is organized as follows. First, we provide background information and discuss related work that quantifies and qualifies IXPs (Section II); then, we detail our method-

ology and measurement plan (Section III), and present our results in terms of IXP coverage and hegemony (Section IV). In the end, we simulate a major outage on DE-CIX and compare this to analyses made during real-world events, and apply our method to simulate outage of other IXPs (Section IV-F).

II. Background and Related Work

A. Background

Conceptually, IXPs are just a physical and transport infrastructure provided to interconnect ASes. Historically, IXPs surged to solve a simple problem: "keep the local traffic local". If one compares the today's Internet exchanges with the traditional basic infrastructure for old commercial trade routes, the latter used the sea and the roads, while fibers can be seen as the roads for the Internet, and the Internet eXchanges are the cities where everything comes together. The IXPs grew around these main fiber ports, close to overseas fibers or in highly populated regions.

The most important benefit of IXPs is do not charge by traffic volume[1] allowing for rapid growth. The Covid19 pandemic highlighted the importance of IXPs. Some IXPs reported a growth of up to 60% [7], while transit providers, such as Telecom Italia, agreed to adopt an "open peering policy" at IXPs [8]. Having an "open peering policy" means that a participant of the IXP is willing to exchange traffic with all other participants without restrictions. Generally, open peering is facilitated by a route server provided by the IXP operator. This route server maintains information about the routes offered by participants [9]. Other options to connect participants among one another are restricted policies, or private agreements. In this paper, we focus exclusively on open peering participants.

Since their initial deployments, IXPs shortened paths, providing technical and economical improvements. Nowadays, they also attract non-local ASes [10] aiming to connect to specific networks present in an IXP (*e.g.,* looking for Telecom Italia peering), raising the IXP's scope to a national and global coverage level.

There is a consensus on the value of IXPs in the growth of the Internet. Moreover, IXPs are important infrastructures to support new technologies (*e.g.,* IoT and 5G), as discussed in [11]. Meanwhile, there is still an ongoing debate about how critical IXPs are; if they are classified as "critical infrastructure" it might have implications such as governmental compliance. For APNIC, IXPs are important but not critical [12]. However, Evans *et al.* [13] discuss the importance to understand how an outage could affect the regional connectivity in the Ashburn/US area in the context of the region's critical cyber infrastructures. They suggest a study that simulates a disruption of that IXP.

B. Related Work

Chatzis *et al.* [14] have identified IXPs as an excellent vantage point for Internet measurements. They analyzed data flows extracted from one unnamed IXP in Europe showing how rich this data is. More recently, Müller *et al.* [15] also used the same IXP sflow extraction approach to validate traffic source and destination identifying possible spoofed traffic passing through the IXP. While network flow data provides the most complete view of an IXP, it might be affected by local privacy laws, which turns it into a difficult approach to apply worldwide.

Most of the other studies tried to understand IXPs using public data from different sources of information: self-declared IXP information from databases such as Euro-IX [16], PeeringDB [17], and PCH [18], traceroute data, IXP looking-glasses, and route collectors such as RIS and RouteViews.

Klöti *et al.* [19] made the first comparative analysis of IXPs using information from PCH, PeeringDB, and Euro-IX. They spotted the limitations of these databases, which includes cases of outdated information, lack of consistency, and fragmented data. More recently, a new initiative called IX-Federation [20] tries to unify IXP data, but its adoption is still limited.

Beyond IXP databases, other enrichment methods have been used to unveil the impact of IXPs on the Internet. Böttger *et al.* [21] used Planetlab [22] and Ark [23] traceroutes to analyze 10 years of IXP growth and the "flattening" phenomenon.

Another improvement is a new traceroute implementation used by RIPE Atlas. The "Traixroute" tool [24] seeks to detect transverse paths across IXPs by correlating time-to-live, round-trip-time, IXP addressing, and reverse DNS data to identify traffic flowing through IXPs. They also use data from PCH, PeeringDB, and RouteViews. While traceroutes are valuable to infer information about IXPs, they require vantage points (VPs). The largest set of vantage points – over 12,000 worldwide – is provided by RIPE Atlas. Unfortunately, RIPE Atlas has a bias towards Europe and North America, and only has vantage points in about 15% of all ASes.

Other researchers analyzed routing information through IXP looking glasses [25] [26]. However, IXP looking glass usage is not standardized. In some implementations, each AS connects directly, sometimes the IXP's route server is connected, sometimes it is not. The final point is the IXP's routing table itself; it just presents a possible network path but it does not show whether such a path is used by participants.

Routing tables and routing traces have been used as a ground truth by the research community to study and map the Internet topology [27]–[29]. However, for IXPs, both types of measurements do not fit very well; either because IXP network addressing is not globally reachable by traceroutes, or because the AS number used by IXPs is transparent and does not appear in the routing table AS-Path (hidden path problem). These limitations are in place because each IXP is a "closed system" with limited reachability. In our approach, we directly connect at each IXP to obtain more information, such as individual AS forwarding and routing preferences.

Finally, the most closely related work is an approach to assess the impact of IXP outages developed by Giotsas *et al.* [30]. That work proposes a methodology for detecting peering infrastructure outages relying on the observation of BGP

[1]Some IXPs charge by infrastructure usage (access port).

communities from updates on RouteViews and RIPE RIS collectors. Their approach successfully identifies an ongoing outage. In contrast, in our study, we use anycast active probing to map where traffic goes and to forecast the impact of IXP outages that have yet to happen.

III. Methodology and Solution

IXPs provides routing path composed by routing information from all the member. This mean, IXPs *per si* do not provide transit to the Internet, and so all the member should have another way to reach the Internet. When the preferred routing path through IXP is unavailable, the traffic will flow using another – commonly an Internet transit provider or another IXP. By establishing anycast sites inside IXPs and transit providers is possible to map individual AS routing preferences by each IXP.

Anycast networks provide a cheap way to deploy such vantage points, as so allow us to emulate failures situations by just deactivating a specific anycast site. As our goal is to understand how the network traffic flows in an IXP outage, we first measure the coverage of each IXP, quantify the overlap of ASes at multiple IXPs, and identify the preferred paths for ASes connected in more than one IXP. To do so, we deployed anycast sites on 10-major IXPs and advertise anycast prefix for all open peers. Next, we use anycast active measurements to understand how the peers in each IXP exchange traffic with us. The anycast catchment mapping[2] enables us to compare the AS's preferences between transit providers and IXPs, or between two or more IXPs, allowing us to map traffic flows and how our announcement propagates over IXPs.

We validate the anycast active probing approach against data flows extracted from one of the largest IXPs in Brazil, similarly to that performed by [3] [14] [15]. The passive data flow analysis show similar results we got from active probing in terms of networks source and destination. The anycast active investigation approach makes it possible to circumvent local privacy laws related to accessing data flows. This fact allows us to analyze the behavior of IXPs worldwide. The remainder of this section provides details of our approach.

A. Measurement infrastructure

Anycast Network: We have obtained access to Tangled testbed [31], which provides with automated fine-grained routing control over anycast prefixes while collecting catchment data, two main points used in our approach. For this experiment, we selected testbed sites connected to IXPs trying to cover the most important IXPs on each continent. We chose based on two criteria: the proximity to the main fiber infrastructure (Figure 1), and the number of participants, as provided by PeeringDB [17], Euro-IX [16], or IXP's website.

We deployed our infrastructure in eight of the ten biggest IXPs globally, plus the biggest in Australia. The exception is *IX.br/RS*: we mainly use this site for validation purposes, as the IXP team was willing to collaborate with our investigations.

[2]Catchment mapping discovers which networks go to which anycast site.

IXP Name	Rank	ASes	Open Peer	Traffic	Website
IX.br/SP	1	2.048	1.473	10TB	www.ix.br
DE-CIX	2	1.007	799	9TB	www.de-cix.net
AMS-IX	3	881	629	8TB	ams-ix.net
LINX	4	827	720	4TB	www.linx.net
NAPAfrica	5	366	362	1.5TB	napafrica.net
FranceIX	7	432	329	1TB	franceix.net
SIX Seattle	9	331	301	1.5TB	seattleix.net
EQ-Sin	10	–	–	–	ix.equinix.com
IX Australia	29	283	217	0.3TB	ix.asn.au
IX.br/RS	46	242	172	0.4TB	www.ix.br

TABLE I: Selected IXPs by PeeringDB Ranking (Oct 2020)

Table I summarizes the selected IXPs, showing data about how many ASes adopt an "Open Peering Policy" – the focus of this paper. EQ-Sin does not make this data publicly available.

Anycast Active Measurement: Tangled testbed provides the tool named Verfploeter [32] to maps the anycast catchment. This tool uses to actively "ping" 6.5 million of the ICMP responsive IP addresses based on the Internet address census hitlist [33]. Using this approach we mapped around 4 million of /24 networks in 63,052 different autonomous systems from a total of 70k ASes active on the Internet [34].

B. Measurement Plan

In our approach, we deployed anycast sites at each IXP and another one using a regular transit provider on the Internet. Our anycast prefix is propagated to the IXP and to the transit provider as depicted in Figure 2. Our routing policy gives preference to receiving traffic in the IXP over transit (*Drain site*) – the more specific prefix /24 has preference over /23. The active measurement system (*pinger*) is one anycast site generating ICMP packets to every /24 network on the hitlist [32]. Afterwards, if one AS uses the IXP route, the ICMP response packet is sent to the anycast site inside the IXP. When the requested network is not using the IXP route our *Drain* site will receive the reply. We use this approach to obtain the IXP's coverage (Figure 2a), and to identify whether one IXP has the preference over others (Figure 2b).

During our measurement, we regularly need to change routing configurations. When we do so, we respect the times presented in the literature as necessary for routing and forwarding convergence [35] [36] [37]. We had the support of one IXP to validate our process against IXP sflow data. We also take precautions to avoid inaccurate measurements: we

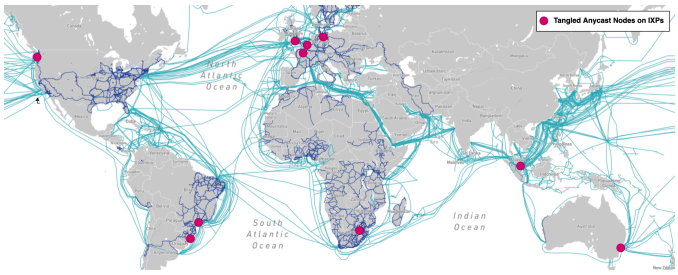

Fig. 1: Tangled IXPs sites over world fiber maps.

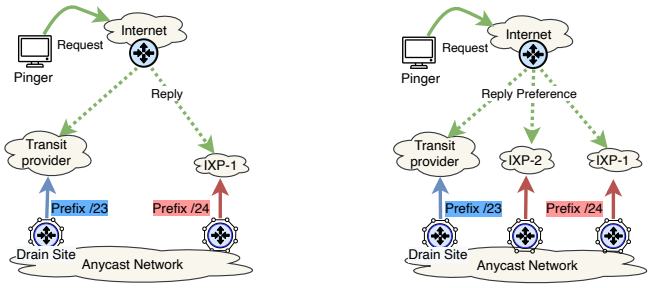

(a) Coverage. (b) Hegemony.

Fig. 2: IXP Experiments.

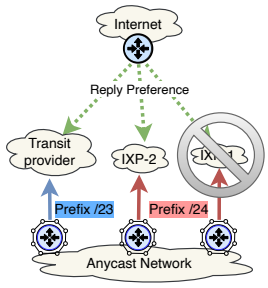

Fig. 3: Modeling IXP Outages.

observed the IXP route-server for neighboring stability issues, to be aware of instabilities on each IXP. We compare our results with the IXP routing table, and several times we contacted IXP administration or individual ASNs to better understand some of our results (*e.g.,* as in §IV-E). As part of our analysis, we started measuring each IXP individually in terms of coverage, hegemony, and representativeness.

Coverage: In this measurement, we explore how many networks and ASes each IXP can reach individually. Figure 2a describes the used setup. We set up two anycast sites at a time: our (*Drain*) and one anycast site in the IXP. If one IXP participant uses the more specific announcement generated on the route-server and spreads the path to other ASes they have a relationship with, we can receive traffic for all the relationships of each participant (peer, customers, and siblings, as defined by CAIDA [28]). The sum of all individual AS relationships will provide the *IXP AS-Cone*. This number represents the total number of ASes able to forward traffic at that IXP.

Representativeness: The data collected in the previous experiment (*Coverage*) is IP geolocated by country code. We also classify the AS origin using the regions of the Regional Information Registries (RIRs). Representativeness is related to networks in same region of the IXPs. This approach is useful for policymakers and government reports and analysis.

IXP Hegemony: In this measurement, we deploy all anycast sites at the same time aiming to identify which IXP has preference to receive the traffic when one AS receives announcements for our more specific prefix from two or more IXPs. Figure 2b shows an overview of the measurement.

IXP Outages: In Figure 3 we depicts the experiment: we keep our prefix /23 announcement in the *Drain site* and the /24 announcement on all other IXPs, in the same way we do the IXP hegemony experiment. Using the IXP hegemony experiment as a baseline, we repeated the experiment several times, turning off one by one each anycast site and mapping this new state. Our objective is to understand how the traffic profile changes in the absence of one IXP. Identifying how big is the portion is redirected to Internet providers (going to our *Drain site*), or forwarded to another IXP. This approach it is possible to map several scenarios, as forecast simultaneous failures (*e.g.,* lose access to all IXPs in Europe). However, in this work, we just consider the case of a single IXP failing.

Here, we also consider that large Internet players – those that own more prefixes or have presence in multiple IXPs – are relatively stable in adopting open policies at IXPs, as well as the average of other participants. Big players are important in the potential of causing large traffic shifts in disruptions. We also believe that the disruptions (*e.g.,* fiber cuts) last short periods (days) before returning to the previous state, making us able to detect it. As a consequence, we can safely map the behavior during an outage by observing the traffic we actively generate. This approach is shown sufficient to map the number of networks redirected between IXPs and those send to transit providers (*i.e.,* Internet).

IV. Results

In this section, we present our results and discuss our findings. First, we examine the IXP's coverage. Then, we evaluate which IXP is preferred to deliver traffic when the participant is connected in multiple IXPs. Lastly, we describe the effects of an IXP outage in terms of traffic shift. To carry out these experiments we perform measurements from Oct-2020 to Mar-2021. Between Oct-Nov we established a traffic baseline between IXPs. In Dec-2020 we improved the stability detection by adding neighboring information. Afterward, we sampled IXPs several times at different intervals excluding periods of instabilities. The outage simulation used a mean of measurements collected on a "stable" day.

A. IXP Participants Stability

The IXP infrastructure is known to be resilient [3]. However, when thousands of ASes are peering together, occasional disruptions are difficult to avoid. When considering an IXP's coverage, a disruption for one individual participant can affect hundreds or thousands of other ASes, causing significant traffic shifts. To avoid that, we first analyzed the stability of participants' connections to the IXP route-server to ascertain whether some IXPs are more stable in terms of participants than others. We consider a participant as stable if it keeps at least one route active on the route-server. Usually, IXPs with more remote peers are more prone to these instabilities.

In Table II, we summarize all peer variations by IXP over 30 days of measurements (from 8-Dec-2020). Unfortunately, AMS-IX, EQ-Singapore, and IX.br/SP do not provide individual peer information. In this evaluation, each IXP shows

IXP	Max	Min	Mean	Std. Deviation
AMS-IX	–	–	–	–
Australia IX	219	216	218.15	0.67
EQ-SIN	–	–	–	–
DE-CIX	808	754	802.59	2.71
FranceIX	335	317	332.68	1.45
IX.br/RS	180	145	177.22	1.58
IX.br/SP	–	–	–	–
LINX	735	716	729.20	4.38
NAPAfrica	366	355	363.57	2.09
SIX	265	259	263.40	0.85

TABLE II: Stability of route-server AS neighbors in each IXP

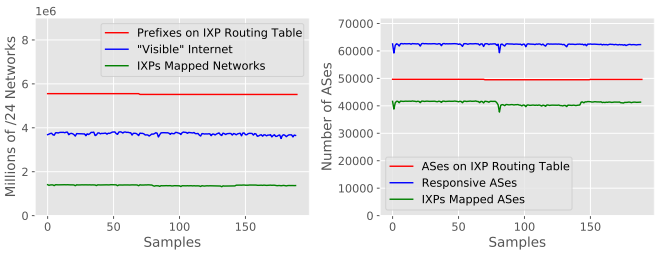

Fig. 4: Internet Coverage by all IXPs together.

around 1% of unstable participants considering a 15-minute sampling window. Here we can identify two main causes of instability: ASes with regional presence connected in just one IXP; or ASes remotely connecting on IXPs. This metric allows us to select a period of stability to deploy our outage test.

B. Coverage

To identify the coverage of each IXP – the maximum reachability of an IXP in terms of networks and ASes – we used two approaches: we performed active measurements with our anycast network and collected route-server routing tables. We compare both in terms of AS and IPv4 address space reachability. When performing active measurements, we consider an AS covered by an IXP if at least one network (/24 prefix) within the AS is successfully mapped to our anycast node inside the IXP. As a result, we successfully mapped 28% of all individual /24 networks from the IPv4 address space, and 89% of all active ASes seen in the Internet global routing table. This limit is related to ICMP responsive addresses. While they do not cover the full IPv4 address space they are considered representative for the Internet [38].When we compared results between routing table collection and our active measurement process, we found a mean difference of 4%, and 12% in the worst case. This difference is quite small – since we are considering only networks that can reply to ICMP requests – and it means that our mapping process fails to cover a maximum of 12% of all ASes in any IXP routing table.

In Figure 4, we summarize how much we could reach in terms of networks and ASes using our infrastructure on IXPs. The red line indicates the number of IPv4 addresses in the entire address space and ASes we have seen in the IXPs routing table respectively. The blue line indicates the total of IPv4 networks and ASes we successfully mapped from our vantage point on the Internet. The green line indicates networks and ASes that reached one of our sites on IXPs - showing that network/AS could see our announcement on that IXP and forward traffic to it.

This experiment revealed that we can reach roughly 38% of all ICMP-responsive IPv4 networks and 57% of all active ASes on the Internet by connecting to the 10 IXPs we selected, which is a considerable number since we are using a small subset of IXPs. Part of those numbers were expected, as we selected big IXPs in terms of prefixes and ASes connected.

We look at two cases in more detail: (1) LINX allows us to reach more than 800k /24 prefixes, covering 17.66% of all responsive IPv4 prefixes in 31,512 ASes. The LINX routing table announces around 2M /24 prefixes with a considerable part not delivering traffic at the IXP. This Internet asymmetry was first identified in [39]; (2) SIX has just 265 ASes using the route-server, but some of them announce a considerable amount of IPv4 address space from eyeball networks, reaching more than 600k /24 prefixes, making them bigger than all other IXPs in our set in terms of address space.

Global Coverage Metric: Considering the autonomous system concept, if one network inside the AS can forward a packet to the IXP, the whole AS can do it as well. If not, this is an AS administrator's preference. For this reason, we evaluate IXP coverage at AS level rather than /24 prefix level.

The regular metric used to evaluate IXPs is the routing table. However, this only evaluates the traffic from the IXP towards other ASes, not the other way. So we split our global coverage metric in two: one for inbound and the other for outbound traffic. This metric represents the percentage of ASes reachable in the IXP using these two views: inbound traffic obtained by active measurements and outbound via the BGP routing table. We consider the total number of active ASes on the Internet (70k in Jan-21) and the individual number of ASes seen in each IXP to determine our Global Coverage Metric.

In Table III we present the results found from the point of view of one AS connected to each IXP. The (In)bound column represents ASes that one can receive traffic from, and the (Out)bound column the ASes one can send traffic to. Values are in percentage of total ASes on the Internet. For example, peers connected to DE-CIX will receive traffic from 46% of all ASes, but can send traffic to 66% of them. This means that, among all measured IXPs, DE-CIX is the most comprehensive for content providers. However, if you are an Internet provider (eyeball network) using just open peering, the best place is AMS-IX for receiving traffic from 50% of all ASes.

IXP	In	Out	IXP	In	Out
AMS-IX	50%	49%	EQ-Sin	21%	23%
DE-CIX	**46%**	**66%**	SIX	20%	21%
LINX	46%	44%	NAPAfrica	17%	19%
FranceIX	26%	25%	IX Australia	16%	18%
IX.br/SP	25%	31%	IX.br/RS	03%	07%

TABLE III: Inbound and outbound Coverage of IXPs.

Interestingly, in some places like AMS-IX, our active measurements can reach more ASes than available in the routing

table – therefore, we double-checked whether our prefix was mistakenly leaking. We concluded that the difference between the active measurements and routing table information pointed to traffic asymmetry. This means that there are some ASes announcing prefixes on the IXP but not using the prefix we announced. We also found the inverse situation, *i.e.*, ASes forwarding traffic to us but not announcing any prefix to the IXP - some of them were other anycast networks. The fact we used an anycast network and also generate traffic out-of-IXP, makes us able to detect such behavior.

Coverage Overlap: Another way to explore IXP coverage is to identify the ASes overlapping on IXPs considering the proportion of total ASes on the Internet. In other words, what fraction of ASes can be reached at more than one IXP. In Figure 5, we show this overlap comparing the IXP AS-Cone of two sets of IXPs using our method. In the first case (Figure 5a), we found that 44% of all ASes of SIX, IX.br/SP, and IX Australia can be reached at any of them. In our second case (Figure 5b), we identified an overlap of 50% between LINX, DE-CIX, and AMS-IX. LINX has 16% exclusive ASes in this set. When we compare our active probe results with routing tables, the first case has very similar results, but in the second case the overlapping goes to 43%. Part of this difference occurs because 2k ASes appearing in the routing table do not deliver any traffic to those IXPs.

After all, IXPs with high coverage overlap could be used to increase the network redundancy. One example is DE-CIX and AMS-IX that have similar coverage. However, a small overlap indicates IXPs that will increase the number of routing paths reachable in the IXP.

C. Representativeness

Our chosen IXPs are known to have a world-wide reach. However, as the main idea behind any IXP is to keep local traffic local, we quantify how representative they are in the region they are located (Figure 6). Our results show AMS-IX as the most representative for European ASes (RIPE region). In Asia (APNIC region), the best coverage is by Equinix Singapore; in North America (ARIN), it is SIX Seattle; and in South America (LACNIC region), it is IX.br/SP. In Africa, we surprisingly found LINX showing the best coverage, more than

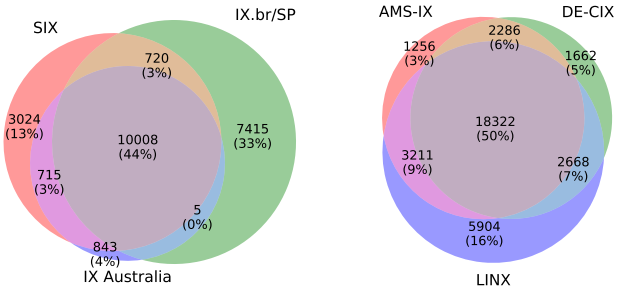

(a) SIX, IX.br/SP and IX Australia. (b) AMS-IX, DEC-IX and LINX.

Fig. 5: IXP's overlap: shared IXP AS-Cone on IXPs.

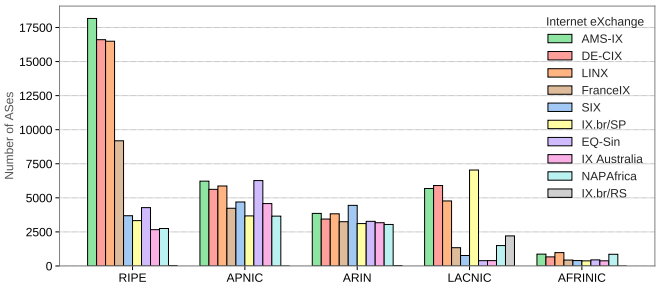

Fig. 6: IXP Coverage by Geographical Region

even NAPAfrica, confirming previous findings of Gupta *et al.* [40], who showed African traffic detouring through Europe.

In Table IV we compare, in terms of percentage, how representative each IXP is for the country and region they are located. We mapped ASes and prefixes using the IP2Location service. Despite being among the ten largest IXPs in terms of traffic volume, FranceIX and SIX show a small regional coverage. FranceIX can exchange traffic for just 37% of all ASes registered in the RIPE NCC region but it is very representative for all French ASes (65%). SIX has low representativeness at region and country levels. In this case, we have to take into account that the US has continental dimensions and SIX is located in an extremity of the country. A similar principle applies to IX.br/RS. Regarding the representativeness criterion, IX.br/SP is the most country- and region-wise representative. The local traffic is the main strength of IX.br/SP by the metrics we analyzed, exchanging traffic for 83% of all Brazil's, and 75% of all LACNIC region ASes.

IXP	Region	Country	IXP	Region	Country
AMS-IX	72%	75%	EQ-Sin	65%	64%
DE-CIX	67%	76%	**SIX**	**28%**	**25%**
LINX	66%	61%	NAPAfrica	67%	83%
FranceIX	**37%**	**65%**	IX Australia	48%	71%
IX.br/SP	75%	83%	IX.br/RS	24%	31%

TABLE IV: Coverage Area: Percentage of traffic exchanged within geographical area (country and RIR region).

D. IXP Hegemony

In Figure 7, we compare the coverage of each IXP with the preferred paths showing how many ASes choose to deliver traffic in one or another IXP. As seen, DE-CIX has IXP hegemony in the RIPE region – it is the preferred one. When comparing with routing table coverage (gray bar), we notice the broader scope of AMS-IX not being fulfilled in terms of preference. The figure also shows the dominance of EQ-Sin, IX.br/SP, NapAfrica, and SIX for the regions they are located in. Specifically looking at IX.br/SP and SIX, we can infer that anyone connecting there will receive traffic from almost all ASes covered by ASes in the same region. This information leads us to conclude that most open policy participants use routing preferences to deliver traffic to the IXP in the same region where they are registered. This is a good sign for

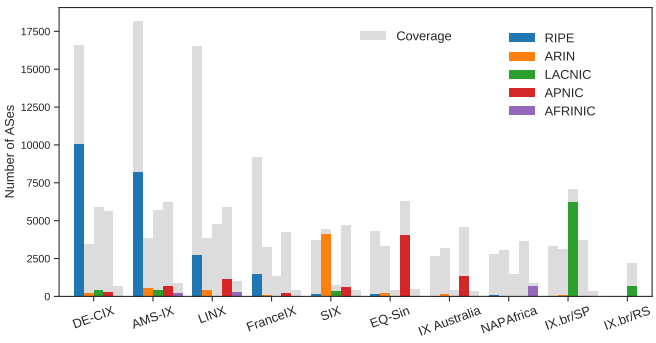

Fig. 7: IXP Dominance by Geographical Region

connectivity as local traffic exchange tends to improve the user experience.

E. IXP Hegemony over Time

Complementary to the previous analysis, we conducted the IXP hegemony experiment over time to observe how steady this behavior is. While the dynamics of the Internet affect routing and causes traffic shifts, most of the time in our observation the traffic going to each IXP is steady. For this experiment, we chose to examine a more "turbulent" period of 20 days, from 24-Oct-2020 to 12-Nov-2020. An investigation of such events helps us to better understand failure scenarios.

In Figure 8, we depict the fraction of networks that each IXP absorbs by using different colors. The traffic is related to the number of /24 prefixes mapped to each IXP. SIX, LINX, AMS-IX, and DE-CIX receive most of the traffic and their distribution is quite representative in terms of /24 networks. France-IX, Australia-IX, and others receive less traffic. The graph only shows IXP traffic, not the portion that is routed through transit over the Internet.

In Figure 8 were highlighted four events that have produced traffic shifting. Event ① shows a traffic transfer from LINX to SIX. While investigating this time window, operators of submarine cables reported equipment problems affecting the TGN Atlantic cable; and, because other links were already overloaded, a cascading effect took place. This had a major impact in England and Spain and we also observe several ASes redirecting traffic to other paths in the following days. The fall-out of this event lasted for weeks.

Events ② and ③ are related to our transit provider. They have confirmed that the cause was a link disruption affecting their site in Frankfurt but did not confirm whether it was related to the TGN Atlantic cable or not. When this event happened, a significant portion of traffic from DE-CIX shifted to AMS-IX and a minor part to transit over the Internet. At event ④, we noticed a traffic shift from LINX to IX.br/SP. This shift was caused by a routing table increase from 150k prefixes to more than 300k on IX.br/SP. In that case, we notice AS6939 (Hurricane Electric) announcing prefixes from 8k ASes at IX.br/SP, growing the coverage of IX.br/SP. They used to announce prefixes for 2k ASes in IX.br/SP and for 8K just in LINX. This manoeuvre shifted traffic from LINX to IX.br/SP. When investigating weeks later, IX.br/SP numbers had returned to 150k prefixes.

F. IXP Outages

In this final subsection, we investigate what happens if a major IXP has an outage. To simulate this, we turned off each anycast node in the target IXP and checked how the traffic shifted to other IXPs and to our transit provider on the Internet. As ground truth to analyze failures in IXPs is still limited, we use the reports available to compare the outages of AMS-IX, DEC-IX, and IX.br/SP with our results.

In Figure 9, we provide a visual summary of how traffic shifts between IXPs. On the left side, we have the IXP origin – those where we simulated a failure. On the right side, we show to where traffic moves. The lines in the middle show the volume of networks that shift when an outage occurs. This figure represents three rounds of measurements in a "stable" day using three distinct pinger sources on different day hours.

In the first experiment, we disconnected our anycast node at IXP.br/SP. As result, we observed a larger amount of traffic flowing to the Internet (37%) and SIX Seattle (23%). Then, in the next experiment, we disabled our site at SIX and kept all others active. We observed almost all traffic (84%) delivered to transit providers in the region. Finally, in the last experiment, we simulate failures on the European continent, where distance is lower and there is plenty of optical capacity, with more ASes multi-connect at several IXPs. This is the case for the three big IXPs (AMS-IX, LINX, and DE-CIX) who share roughly 50% of their IXP AS-cone (Figure 5a). When we disconnected LINX we see the majority of the traffic (48%) being routed to Internet providers, and a smaller fraction going to FranceIX (18%), and AMS-IX (16%). When we simulate an AMS-IX outage, AMS-IX traffic goes 36% to LINX, 24% to transit providers, and 23% to DE-CIX. When we simulate a DEC-IX failure, we observed 48% of networks previously routed there being redirected to AMS-IX, 23% to Internet providers, 17% to LINX. Complementary to this, we also map the impact on networks in Germany and the Netherlands against outages of DE-CIX and AMS-IX. When we simulate an outage on DE-CIX, we see 57% of German ASes sending traffic to AMS-IX, and 37% routing through Internet providers. For networks in the Netherlands, we noticed a larger impact on providers, with 47% of Dutch networks draining to the Internet, 28% to LINX, and just 18% to DE-CIX. This shows more German ASes using AMS-IX as a second IXP while Dutch ASes rely on transit providers, or another IXP that we did not cover.

These experiments provide insights regarding the effects of a possible outage at these IXPs. For example, in the case of an IXP.br/SP outage, some Brazilian ASes prefer to exchange traffic in the US rather than use IX.br/RS, impacting the latency and, ultimately, the user experience, which are facts also reported by AS operators [6] in previous IX.br/SP failures. In the case of SIX, ASes need to provide sufficient contingency bandwidth in their transit providers to receive all traffic previously routed by networks at SIX (84%) or plan to

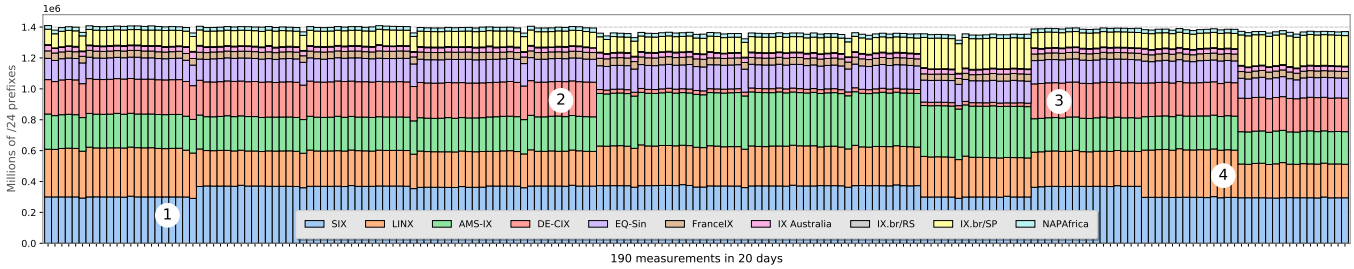

Fig. 8: IXP Hegemony experiment results over 20 days (24-Oct-2020 to 12-Nov-2020)

connect to other IXPs in the US to achieve redundancy for these networks.

We analyse our simulations in other outages: the AMS-IX in 2015 and the DE-CIX in 2018 [4]. In AMS-IX event, the DE-CIX team noticed the flaw affecting its structure [41]. They observed a small traffic decrease in DE-CIX, and losses between 3-5% to some ASes. One cause was the router overload in participants peering in both infrastructures with the same router. The AMS-IX case was a logical outage (loop), causing router overload, and its effect was analyzed using RIPE Atlas. They identified an intersection of 40-50% of direct participants in DE-CIX and AMS-IX, and a remarkable traffic asymmetry of 16% between ASes in both IXPs. When using our methodology, we forecast similar results for traffic asymmetry and ASes overlapping in both IXPs. The traffic shift, in terms of volume, could not be evaluated because we do not have access to the stats per participant from that period.

The DE-CIX outage in 2018 was a power outage, affecting not just the route-servers but one entire IXP datacenter. It registered an impact on big Internet providers such as TeliaNet, Level3, Deutsche Telecom, and a ripple effect on other providers [4]. Here, the outage affected open and private peering infrastructures. We forecast 23% of DE-CIX networks overloading internet providers, but again we do not have information about private peering networks to add to this number. Our method can be applied to better forecast this scenario by establishing private peering agreements and simulating an outage of ASes connected in a specific datacenter.

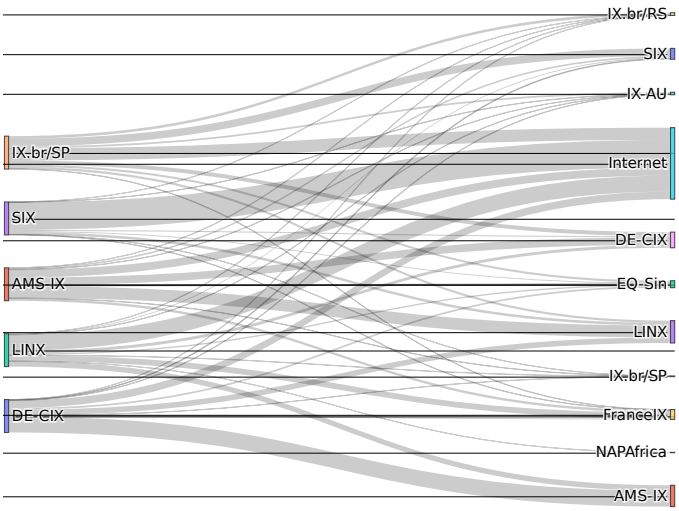

Fig. 9: Mapping traffic shifts on failures.

V. Conclusions and Future Work

In this paper, we introduced a new method using anycast active probing to forecast the impact of IXP outages. This new view enables us to quantify the coverage of each IXP and understand the preference of each participant in using such infrastructures. We determined metrics to evaluate the IXPs by measuring how much of the Internet we can reach by connecting to each one, and how representative they are for the region and country they are placed in. We compare IXP routing tables against our anycast active probe results showing a significant asymmetry in some IXPs.

We show, by using "Open Peering" at our set of 10 IXPs, that we can reach 38% of all 'visible' networks and 56% of all ASes on the Internet. While this is a remarkable number of networks that we can reach only using IXP, the access to most networks still depends on private peering agreements or transit from the Tier hierarchy. Part of this situation is because a huge address space is concentrated in a few ASes. We also quantify the overlapping of prefixes and ASes for our set of IXPs and identified the preferred IXPs for each AS. This information shows that despite remote peering growing at IXPs, participants make use of routing preferences to deliver IXP traffic locally.

Our method was applied to forecast how traffic flows in case of a major IXPs failure. The experiment on IXP outages is also useful for policymakers to better understand the impact of each IXP on the Internet inside a country. Likewise, telco operators and IXPs could use this methodology to identify infrastructure bottlenecks and collect data for a better risk assessment. A similar infrastructure also can be deployed to specific datacenter infrastructures, aiming to simulate individual datacenters outage.

Future work – During our experiments, we detected traffic asymmetry inside IXPs, in distinct grades. Especially some cases of full asymmetry are remarkable, with participants only receiving outbound traffic from the IXP and never delivering inbound traffic. We aim to better investigate these issues, identifying what motivates such behavior, or whether it is just related to misconfiguration issues.

VI. Acknowledgments

This project have the support of SIDNLabs and NSNetLabs and is founded by DHS HSARPA Cyber Security Division (HSHQDC-17-R-B0004-TTA.02-0006-I), Netherlands Organization for scientific research NWO (628.001.029), and CONCORDIA, the Cybersecurity Competence Network supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 830927.

References

[1] A. Dhamdhere and C. Dovrolis, "The internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of the 6th International COnference*, 2010, pp. 1–12.

[2] T. Arnold, J. He, W. Jiang, M. Calder, I. Cunha, V. Giotsas, and E. Katz-Bassett, "Cloud provider connectivity in the flat internet," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 230–246.

[3] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large european ixp," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, 2012, pp. 163–174.

[4] A. Henthorn-Iwane, "Understanding internet exchanges via the de-cix outage," https://www.thousandeyes.com/blog/network-monitoring-de-cix-outage/.

[5] ——, "Capacity media," https://www.capacitymedia.com/articles/3799671/DE-CIX-down-at-Frankfurt-after-power-failure-at-Interxion, 09 2020.

[6] Registro.br, "Lista caiu - mail list archives," https://eng.registro.br/pipermail/caiu/2018-September/, 09 2018, (Accessed on 1/7/2021).

[7] M. Candela, V. Luconi, and A. Vecchio, "Impact of the covid-19 pandemic on the internet latency: A large-scale study," *Computer Networks*, vol. 182, p. 107495, 2020.

[8] M. Candela and A. Prado, "Italian operators' response to the covid-19 pandemic," *ACM SIGCOMM Computer Communication Review*, vol. 51, no. 1, pp. 26–31, 2021.

[9] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger, "Peering at peerings: On the role of ixp route servers," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 31–44.

[10] G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, X. Dimitropoulos, and V. Giotsas, "O peer, where art thou? uncovering remote peering interconnections at ixps," in *Proceedings of the Internet Measurement Conference 2018*, 2018, pp. 265–278.

[11] T. Hoeschele, C. Dietzel, D. Kopp, F. H. Fitzek, and M. Reisslein, "Importance of internet exchange point (ixp) infrastructure for 5g: Estimating the impact of 5g use cases," *Telecommunications Policy*, vol. 45, no. 3, p. 102091, 2021.

[12] Apnic, "Critical infrastructure – faqs – apnic," https://www.apnic.net/get-ip/faqs/critical-infrastructure/#are-ixp-critical-infrastructure.

[13] N. Evans and W. Horsthemke, "Regional critical infrastructure," in *Cyber Resilience of Systems and Networks*. Springer, 2019, pp. 355–380.

[14] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann, "On the benefits of using a large ixp as an internet vantage point," in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 333–346.

[15] L. Müller, "Challenges in inferring spoofed traffic at ixps," in *Conference on Emerging Networking Experiments And Technologies*, 2019.

[16] EuroIX, "Ixps," https://ixpdb.euro-ix.net/en/ixpdb/ixps/, 10 2019, (Accessed On 1/7/2021).

[17] PeeringBD, "Peeringdb database," https://www.peeringdb.com/, 12 2020, (Accessed On 1/7/2021).

[18] P. C. House, "Internet exchange directory," https://www.pch.net/ixp/dir, 10 2020, (Accessed On 1/7/2021).

[19] R. Klöti, B. Ager, V. Kotronis, G. Nomikos, and X. Dimitropoulos, "A comparative look into public ixp datasets," *ACM SIGCOMM Computer Communication Review*, vol. 46, no. 1, pp. 21–29, 2016.

[20] IX-F, "Internet exchange federation," http://www.ix-f.net/, 07 2014, (Accessed On 1/7/2021).

[21] T. Böttger, G. Antichi, E. L. Fernandes, R. di Lallo, M. Bruyere, S. Uhlig, and I. Castro, "The elusive internet flattening: 10 years of ixp growth," *CoRR*, 2018.

[22] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "Planetlab: an overlay testbed for broad-coverage services," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.

[23] "Archipelago (ark) measurement infrastructure," https://www.caida.org/projects/ark/, 01 2007, (Accessed On 1/7/2021).

[24] G. Nomikos and X. Dimitropoulos, "traixroute: Detecting ixps in traceroute paths," in *International Conference on Passive and Active Network Measurement*. Springer, 2016, pp. 346–358.

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

[26] S. H. B. Brito, M. A. S. Santos, R. dos Reis Fontes, D. A. L. Perez, H. D. L. da Silva, and C. R. E. Rothenberg, "An analysis of the largest national ecosystem of public internet exchange points: The case of brazil," *Journal of Communication and Information Systems*, vol. 31, no. 1, 2016.

[27] V. Giotsas, S. Zhou, M. Luckie, and K. Claffy, "Inferring multilateral peering," in *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, 2013, pp. 247–258.

[28] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. Claffy, "As relationships, customer cones, and validation," in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 243–256.

[29] V. Giotsas, M. Luckie, B. Huffaker, and K. Claffy, "Inferring complex as relationships," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 23–30.

[30] V. Giotsas, C. Dietzel, G. Smaragdakis, A. Feldmann, A. Berger, and E. Aben, "Detecting peering infrastructure outages in the wild," in *Proceedings of the conference of the ACM special interest group on data communication*, 2017, pp. 446–459.

[31] L. M. Bertholdo, J. M. Ceron, W. B. de Vries, R. de Oliveira Schmidt, L. Z. Granville, R. van Rijswijk-Deij, and A. Pras, "Tangled: A cooperative anycast testbed," in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 2021, pp. 766–771.

[32] W. B. de Vries, R. de O. Schmidt, W. Hardaker, J. Heidemann, P.-T. de Boer, and A. Pras, "Broad and Load-Aware Anycast Mapping with Verfploeter," in *Proceedings of the 2017 Internet Measurement Conference*, ser. IMC '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 477–488. [Online]. Available: https://doi.org/10.1145/3131365.3131371

[33] L. project, "Lander:internet address history it91w-20200710," 2020.

[34] T. Bates, P. Smith, and G. Huston, "Cidr report," https://www.cidr-report.org/as2.0/, 06 2007, (Accessed On 1/7/2021).

[35] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed internet routing convergence," *ACM SIGCOMM Computer Communication Review*, vol. 30, no. 4, pp. 175–187, 2000.

[36] R. Teixeira, S. Uhlig, and C. Diot, "Bgp route propagation between neighboring domains," in *International Conference on Passive and Active Network Measurement*. Springer, 2007, pp. 11–21.

[37] R. B. d. Silva and E. S. Mota, "A Survey on Approaches to Reduce BGP Interdomain Routing Convergence Delay on the Internet," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, 2017.

[38] X. Fan and J. Heidemann, "Selecting representative ip addresses for internet topology studies," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010, pp. 411–423.

[39] W. de Vries, "How asymmetric is the internet?" in *IFIP Conference on Autonomous Infrastructure, Management and Security*, 2015.

[40] A. Gupta and et. al., "Peering at the internet's frontier: A first look at isp interconnectivity in africa," in *International Conference on Passive and Active Network Measurement*. Springer, 2014.

[41] C. Dietzel, "Investigation of traffic dependencies between ixps in failure scenarios," https://www.menog.org/meetings/menog-16/, 2016.