

Analyzing Remote Peering Deployment and its Implications for Internet Routing

Fabrizio Mazzola, Augusto Setti, Pedro Marcos, Marinho Barcellos

Abstract—Internet eXchange Points (IXPs) have significantly transformed the structure and economics of the Internet by allowing many nearby networks to connect directly, avoiding the need for service providers. These large IXPs are so beneficial that they are not just used by local networks, but also by Autonomous Systems (AS) worldwide. This is made possible by Remote Peering (RP), which typically involves the use of RP resellers to access remote IXPs. In this paper, we evaluate the effects of RP on different routing aspects, using a representative group of IXPs located on three continents. First, we compare the growth of RP deployment over one and a half years on seven IXPs. We identify that, overall, the growth is smaller than previously thought and depends on the size, age, and location of the IXP. Then, we reveal substantial presence of route announcement mispractices (when networks prioritize the remote IXP over the local IXP), which are associated to routing anomalies. Third, we analyse the reliability of RP interfaces and discover that they have more undesirable BGP session state changes and stay less time continuously with their BGP session active compared to traditional connections. Finally, we see that the adoption of RP-related BGP communities, i.e. to perform traffic engineering to remote peers, is not what could be inferred and far from desirable. We make our data and results available to the community via a web portal.

I. INTRODUCTION

Internet eXchange Points increase connectivity by concentrating together a large number of networks. IXPs shorten Internet paths, improve performance, and reduce interconnection cost to its members [1]–[4]. As of February 2023, there were over 1,000 IXPs deployed worldwide [5]–[7], with the largest IXPs having surpassed 1,000 members [8]–[10] and 10 Tbps of peak traffic [8], [11]–[13].

An original motivation of IXPs was to keep local traffic local by having ASes physically present (deploying a router) at an IXP facility. However, IXPs are no longer restricted to members physically present at IXP facilities. *Remote Peering* allows an AS to reach a remote IXP, typically replacing the physical presence at the IXP facility with the services of a layer-2 provider. RP benefits include a quicker setup, no additional hardware, and lower installation costs when compared to local peering [14]–[16]. Its benefits have made RP quite popular; for example, as of February 2023, ASes from 80 different countries connect to LINX remotely [8]. To cope with the demand for remote peering, IXPs and RP resellers have expanded their offerings [17]–[19]. As a result, some IXPs have more than 50 official partners selling remote peering services [10], [20], [21].

Despite being widely deployed at IXPs around the world, the community is still trying to understand the usage of remote peering, how/if it is still expanding, and the implications to Internet routing [22]–[28]. The precursory investigations

about RP focused on methods to identify which networks remotely connect at different IXPs [16], [29]. Later studies have provided preliminary insights about the impacts of RP, such as introducing undesired latency to anycast [30], making it harder to detect peering infrastructure outages [31] and performing worse than local routes to reach specific prefixes on IXPs [32].

In this work, we perform a comprehensive measurement study of different routing impacts of RP in the wild, including its impacts on IXP growth, prefix announcement mispractices, IXP member’s connectivity and operation time, and usage of BGP communities to perform traffic engineering (TE). For three months, we infer RP interfaces and announced prefixes on seven IXPs (identified in Table I), including five of the world’s ten-largest IXPs by membership, to shed light on how RP evolves in different IXPs. Using these inferences, we investigated how prefix announcement mispractices for ASes connected on multiple IXPs via RP and local peering could introduce routing problems, inadvertently steering their traffic to the remote IXP. Next, we looked at the BGP session status of ASes in IXP route servers (RS) and evaluated the reliability and stability of remote connections. Lastly, we analyze the popularity and usage patterns of BGP communities explicitly created by three IXPs to perform TE on networks connecting via RP. Our contributions are as follows.

- 1) We investigate the evolution of RP deployment on IXPs with varying characteristics, such as the number of members and traffic exchanged (§III). To enable the study, we perform active latency measurements and collect BGP data from seven IXPs worldwide between August and October 2022. To widen our observation window, we compare our results with [32]. Our results show that the number of remote interfaces has expanded faster than local ones since 2021. We also observed that newer, less established IXPs tend to lead the growth of RP use. We also noted that although the number of remote members grew over time, there was an unexpected decrease in the number of remote prefixes announced.
- 2) We select ASes connecting to IXPs via different connection methods (i.e., RP and local peering) to further investigate their prefix announcement practices (§IV). We identified more than 60 distinct ASes announcing their most specific prefixes (therefore preferential) on the remote IXP instead of prioritizing their local IXP connection, accounting for more than 37000 affected prefixes. These practices are highly likely to introduce *trombone paths*, impairing peering performance and connectivity.

- 3) We investigate the connection characteristics of remote and local IXP members on a subset of IXPs that enabled data collection for every IXP interface individually (§V). Our fine-grained data collection of the ASes connected to the IXP RS (every 15 minutes over one and a half months) allows us to evaluate the connection stability and reliability between remote and local interfaces through different metrics, such as the uninterrupted time they stay with an active BGP session and the number of BGP session state changes. We found that RP connections to the IXP RS are consistently less stable than local ones for all IXPs and that BGP sessions in RP interfaces not only stay online less time in total but also change state more frequently.
- 4) We collect and analyze the usage of BGP communities specifically created to perform TE to remote peers on three IXPs (§VI). Our findings show that, despite created and made available for more than one year in the three IXPs, the usage of these communities is still negligible, with less than 1.11% of IXP members applying them to their routes. We also identify the ASes with the highest use of RP BGP communities and highlight their usage patterns.
- 5) We contribute a web portal to the community, making our results publicly accessible (§VII). The portal depicts the data we collected between August and October 2022, along with interactive analysis and plots.

II. MEASUREMENT ARCHITECTURE

In this section, we explain the measurement architecture used to infer remote interfaces and prefixes for our intended analysis in four parts: we justify the selection of IXPs considered (§II-A), the control plane datasets (with BGP information) we collected for the study (§II-B), how we used VPs to perform necessary dataplane measurements (§II-D), and the methodology used to infer RP (§II-C). In the subsequent sections (§III, §IV, and §V), we provide further details about the specific methodology and metrics used in each corresponding analysis.

A. Peering Infrastructure Selection

To identify networks connected via remote peering, and prefixes and routes announced via remote peering, we required peering infrastructures that had (i) publicly available BGP routing data, and (ii) an active measurement vantage point (VP) attached to the IXP switching fabric. Table I presents the seven selected IXPs where we had both BGP routing data and active measurement capability. These IXPs include five of the world’s ten largest IXPs by membership [6], [7] and are deployed in five different countries. The three Brazilian IXPs (i.e., PTT sites) are part of the largest ecosystem of public IXPs in the world (IX.br) and are the leading Latin American IXPs in terms of average traffic volumes (≈ 12.9 , 9.2 , and 1.4 Tbps, respectively) [33]–[35].

B. Datasets

To identify the peering router’s IP and ASN of all members at each IXP, we combine multiple public data sources for

IXP	Location	Observed Interfaces	BGP VPs	
			LG	PCH
PTT-SP	Sao Paulo, BR	2,163	✓	✗
LINX	Londres, UK	876	✓	✓
NAPAfrica	Johannesburg, ZA	611	✗	✓
PTT-RJ	Rio de Janeiro, BR	418	✓	✗
PTT-CE	Fortaleza, BR	419	✓	✗
Eq-Ash	Ashburn, VA, US	375	✗	✓
Eq-Chi	Chicago, IL, US	265	✗	✓

TABLE I: Seven analyzed IXPs in our work, along with the number of interfaces observed on 2022-10-14 (2022-10-17 for PTT-RJ) and the BGP data availability in each peering infrastructure. These dates represent the days with the maximum number of interfaces at each IXP.

all IXPs except for LINX, which publishes this information through their member portal [21]. We collected membership data and subnet information from Euro-IX [6] and the publicly available databases of Hurricane Electric (HE) [7], PeeringDB (PDB) [5], and Packet Clearing House (PCH) IXP Directory [36]. In cases of conflicts, we followed the preference ordering described in [29]: *Euro-IX* > *HE* > *PDB* > *PCH*.

We used two sources of BGP routing data: (i) Looking Glass (LG) of the IXP which observes routes from the IXP’s Route Server and (ii) routes from the archive collected by PCH [37]. For IXPs with both PCH and LG views, we used data archived by PCH because it has greater visibility of routes advertised by IXP members. For example, when comparing both datasets for LINX, we observed 3.9x more routes and 2.0x more prefixes from PCH than from LG views. We discarded: (i) routes with artifacts, such as reserved/unassigned ASes [38] and loops; (ii) prefixes shorter than /8 or longer than /24.

At each IXP listed in Table I, we used RouteViews collectors which were directly connected to the IXP LAN to conduct active measurements using scamper [39]. Figure 1 illustrates the measurement architecture of each RouteViews collector and how we used them to conduct active measurements.

To identify which networks are connected remotely at IXPs, we conduct active latency measurements to each IXP member’s peering router. These measurements use the IP address that the collector has in the IXP LAN (X.1), so that probes and responses cross the IXP LAN, as in when we probe X.2 in Figure 1. The *ping* responses help us to understand the geographic distance between the IXP and the IXP member’s peering router and, consequently, infer how far they are from the analyzed peering infrastructures.

C. Distinguishing remote from local peering

To infer members using RP at all seven IXPs, we used the method in [16], which uses latency measurements and empirically obtained thresholds as a proxy of physical distance, with the following approach. For each IXP, we associated IXP member ASes and their assigned IXP IP addresses using the datasets mentioned in §II-B. We performed latency measurements to these addresses between 2022-08-01 and 2022-10-31. From each RouteViews scamper instance, we probed




Fig. 1: Architecture of our data plane measurements (similar to [32]). We used RouteViews collectors with an interface in the IXP LAN as VPs for data plane measurements. Delay measurements to the peering router of each IXP member (e.g., X.2) used the collector’s IP address in the IXP LAN (X.1), so the probes and responses crossed the IXP LAN.

each interface with 30 *pings* per day and used the minimum latency for each address to account for cases of transient congestion. To ensure that the ping replies returned directly over the peering infrastructure, we discarded measurements where the replies had an IP-TTL value that appeared to have been decremented (i.e., not 64 or 255). If the minimum latency from a given interface was 10ms or higher, we classified the member’s router as remotely connected to the IXP; a latency of 10ms would roughly correspond to a distance of up to 1000km from the IXP [40], [41]. We adopted the method in [16] because its latency threshold alone yielded accurate results for single metropolitan area peering infrastructures [42], which is the case of the analyzed IXPs in our work (see §II-A).

D. Data plane measurements

To further assess the correctness of our inferences – and similar to step 2 in [42] (*colocation-informed RTT interpretation*) – we obtained the colocation facilities of each of the seven analyzed IXPs in public data sources (IXP websites and PeeringDB) and computed the distance between them. We observed that Equinix Ashburn has the largest distance between facilities (i.e., 80km), which corresponds to a latency of ≈ 1 ms. Therefore, any IXP peer interface with latency consistently higher than 10ms is unlikely to be a local peer at the IXPs we examined.

III. LONGITUDINAL ANALYSIS OF RP

RP is widely deployed in IXPs around the world nowadays, representing up to 40% of the member base in the largest IXPs [32], [42]. According to [42], as of mid-2018, the deployment of new remote peering connections had been major factor in the recent IXP growth.

Methodology. We pose two questions: (i) are RP connections still a major contributor to the growth of IXPs? (ii) How does the prevalence of remote peering changes over time in IXPs? We answer these questions in three steps, as follows. To infer RP at IXPs, we use a methodology similar to [32]: during three months, we continuously collect IXP membership data,

	Remote Interfaces		
	05-05-2021 [32]	16-08-2022	14-10-2022 (17-10-2022 for PTT-RJ)
PTT-SP	681 of 2,169 (31.4%)	720 of 2,156 (33.4%)	728 of 2175 (33.5%)
LINX	121 of 911 (13.3%)	129 of 878 (14.7%)	130 of 888 (14.6%)
NAPAfrica	40 of 542 (7.4%)	52 of 612 (8.5%)	49 of 619 (7.9%)
PTT-RJ	61 of 462 (13.2%)	67 of 427 (15.7%)	67 of 444 (15.1%)
PTT-CE	139 of 395 (35.2%)	169 of 423 (39.9%)	161 of 421 (38.2%)
Eq-Ash	35 of 365 (9.6%)	31 of 375 (8.3%)	30 of 378 (7.9%)
Eq-Chi	17 of 259 (6.6%)	18 of 265 (6.8%)	19 of 265 (7.2%)

TABLE II: Number of interfaces connected via RP at IXPs and the percentage they represent on the total IXP members. The prevalence of RP grew slightly in most IXPs (6/7) between 05-2021 and 10-2022 (matched by an increase in the absolute number of remote interfaces). Eq-Ash was the only IXP with a decrease in remote interfaces, both absolutely and in prevalence.

announced routes at IXPs, and latency measurements from a VP inside the IXP to each member interface (§II-C). Then, we compare our results regarding remotely inferred interfaces (§III-A) and prefixes (§III-B) with the most recent results from the literature.

A. How does remote membership vary with time?

First, we investigate how the prevalence of remote peering (in terms of member interfaces) changes with time, as shown in Table II. We select two days from our 2022 data collection (one in the first month and one in the third) and compare them with the numbers obtained in May 2021, from earlier work, extending the analysis period from 3 months to 1.5 years. We selected the two days with the highest numbers of prefixes and connected interfaces at IXPs for each month, to increase the visibility on the IXP routing data.

We found that remote member interfaces’ deployment has grown since 2021 in almost all IXPs (6/7) and that the amount of (aggregated) growth varied according to the characteristics of the IXP. The numbers indicate that well-established IXPs, such as LINX and PTT-SP, have grown less, which conforms intuition: these massive IXPs already have a widespread member base and less space to grow in local and remote networks compared to emergent infrastructures, such as NAPAfrica and PTT-CE. In three months, the growth ranged between 6.9% (PTT-SP) and 22.5% (NAPAfrica). The odd case, Equinix Ashburn, actually had a decrease in the number of remote member interfaces (-14.29%), showing that the RP growth cannot be simply assumed.

We contrast our findings with earlier work [42], whose authors looked at the RP evolution at five IXPs between 2017

	Remote Prefixes		
	05-05-2021 [32]	16-08-2022	14-10-2022 (17-10-2022 - PTT-RJ)
PTT-SP	19,612 of 154,561 (12.7%)	13,764 of 122,237 (11.3%)	12,453 of 130,707 (9.5%)
LINX	71,452 of 482,643 (14.8%)	67,794 of 472,746 (14.3%)	70,163 of 463,036 (15.2%)
NAPAfrica	7,252 of 144,513 (5.0%)	8,301 of 174,449 (4.8%)	8,435 of 175,797 (4.8%)
PTT-RJ	3,850 of 128,652 (3.0%)	2,654 of 98,176 (2.7%)	1,640 of 93,589 (1.7%)
PTT-CE	6,869 of 26,597 (25.8%)	5,356 of 28,017 (19.1%)	5,367 of 27,137 (19.8%)
Eq-Ash	46,752 of 525,688 (8.9%)	21,698 of 600,632 (3.6%)	21,679 of 597,586 (3.6%)
Eq-Chi	8,120 of 271,855 (3.0%)	8,686 of 344,233 (2.5%)	8,745 of 344,312 (2.5%)

TABLE III: Number and percentages of remote and prefixes at IXPs. Remote prefixes decreased in four IXPs even though they showed a growth in remote members over time.

and 2018. They concluded that remote member interfaces on the five IXPs grew 20% over one year, but this analysis was for all IXPs combined *aggregated*. We look in more depth and find that this growth is not equally distributed, being influenced by basic properties of the IXP (e.g., size, traffic, location).

We can also observe (from Table II) that the fraction of remote interfaces remained relatively stable, with changes under 3% (in PTT-CE, it went from 35.2% to 38.2%). This is so because the number of local interfaces fell in some IXPs *despite* the growth in RP. For PTT-SP and PTT-RJ, the number of local interfaces has decreased -2.76% and -5.99%, respectively. In virtually all IXPs (6/7), the number of remote interfaces increased, reaching up to 10 $\times$ when compared than local interfaces. NAPAfrica, the younger/smaller IXP in our set, had a subtle difference (1.6 $\times$), hinting that for emergent IXPs the accelerated grow may happen similarly for both remote and local interfaces.

B. How do remotely announced prefixes vary with time?

To answer the question, first we combine the information about remote interfaces with IXP BGP routing data to identify the prefixes announced by these members. To classify a *prefix* as local or remote, we tag it according to the *nextHop* interface of the BGP route, which refers to the IXP member announcing it at the IXP. A prefix that has routes being announced by *both* remote *and* local members is denoted as *hybrid*.

Table III shows the results obtained for the three-month data collection compared with data from 2021 [32]. For most IXPs (5/7), the number of remote prefixes decreased from 2021 to 2022. The reasons for this, according to our private talks with network operators, is that networks use RP more often to fetch/download content (faster) from remote locations (via more established IXPs) than to deliver/upload content to distant networks. We did not observe a change pattern in terms of remote prefix prevalence, ranging broadly between -57.4% (2,210 fewer) and 16.3% (1,183 more). One might expect that an increase in the number of remote members would lead to

an increase in remote prefixes, but we found no correlation: the four IXPs with a reduction in remote prefixes *actually increased* the number of remote members in the same period (as shown in §III-A). In nearly half of the IXPs (4/7), we saw a shift in the prevalence of prefixes, from remote to local ones.

IV. ROUTING MISPRACTICES CAUSED BY RP

If remote peering is not correctly configured by operators, ASes connected to multiple IXPs with different types of peering (i.e., a combination of local and remote connections) may introduce *trombone paths*. In this section, we investigate likely routing mispractice cases associated with RP misconfiguration. Trombone paths happen when two networks connected in a geographically-close location (e.g., London) exchange traffic via a distant peering facility (e.g., Chicago), causing poor routing performance and poor connectivity.

The Internet routing best practices indicate that more specific prefixes have a preference (as the intended destination for packets) over less specific prefixes. For example, a prefix 1.0.0.0/24 will be preferred over 1.0.0.0/16. In a peering scenario, the more specific prefixes (routes) announced by one AS in one IXP will be preferred for traffic exchange over less specific prefixes that can be announced at one transit provider. As reported by [23], an AS (in São Paulo) that has peering connections to a local IXP and a remote IXP can damage its peering connection performance if the operator does not correctly configure its prefix announcements in both IXPs. We exemplify this scenario in Figure 2, where an AS A connects to an IXP nearby (i.e. local peering on IXP X) and to a distant IXP (i.e. RP on IXP Y). By wrongly announcing more specific prefixes to the remote IXP instead of the local IXP, it can inadvertently steer its traffic exchange thousands of kilometres away (green path) from the optimal peering infrastructure (IXP X). For example, if an AS B also connected at IXP X via local peering and IXP Y via RP, the more specific prefix wrongly announced would lead the traffic exchange between them to go via the distant IXP Y, even though the better option (IXP X) is where both ASes are locally connected. In the rest of this section, we denote these likely problematic cases as *trombone prefixes*.

Methodology. We use the BGP data collected between 2022-08-01 and 2022-10-31 along with ping measurements performed from VPs in IXPs to corresponding IXP member interfaces (§ III). The data allows us to identify the ASes that are present in multiple IXPs (in our set), by definition with one local connection and two or more, remote. For these ASes, we check the set of prefixes they announce in more than one IXP. The last step is to verify the of trombone prefix cases, that is, where a more specific prefix is announced at the remote IXP while a less specific prefix is announced at the local IXP.

We quantify the daily number of trombone prefixes and identify the ASes that announce them. We observed more than 480 trombone prefixes/day, with a peak of 1069/day. The daily number of ASes varied between 4 and 31, with an average of 17 ASes daily over the entire period; the 31 ASes were responsible for 1408 prefixes. We discover that both the number of trombone prefixes and ASes responsible




Fig. 2: Routing mispractice behavior when announcing at remote and local IXPs. AS A announces a more specific prefix to the IXP it connects via RP and a less specific prefix to the IXP it connects via Local Peering. When AS B receives both routes, it will give preference to the one via IXP R (red path) considering Internet routing best practices.

for them can vary over time, indicating that they usually represent a *transient* routing problem that ASes fix in a matter of days. Looking further at the ASes that own such prefixes, we see just two ASes being accountable for 56.8% of all the overall trombone prefixes on average in the two-month period (AS52320 and AS262589, with 12616 and 8417 occurrences each). Not surprisingly, perhaps, we find that these two ASes were also the ones responsible for the peak occurrences (representing 78.6% of the prefixes on this day). Both ASes are continental networks offering IP transit, cloud and remote peering services for Latin American customers. We believe these prefixes represent routes from these ASes' customers that were improperly announced by the transit provider and subsequently fixed.

IX.br and LINX are involved in most cases identified. Here we shift the analysis to the *IXPs* most commonly involved in the trombone prefixes. Figure 3 shows the most relevant combinations of IXPs for two months; we only display the IXP pairs with more than 1000 occurrences, which represent 93.0% of the total. The order is relevant, such that the first and second IXPs of the pair are, respectively, announcing the most specific prefix and the less specific prefix. We can observe in the figure that 50% of the cases are due to less specific prefixes being announced locally by networks connected at one of the IX.br IXPs (i.e., PTT-SP, PTT-RJ, and PTT-CE) and more specific versions of it announced remotely at LINX. Five ASes were responsible for more than 86.8% of the cases on the three IXPs: the previously mentioned AS52320 and AS262589, and AS28329, AS14840, and AS267613. The latter three are Brazilian networks with continental coverage, also offering IP transit and cloud services to Latin American clients.

Even though looking only at control plane data cannot guarantee that these cases represent actual routing problems, anedoctally, they often do. Our study reveals that there is a considerable number of occurrences happening over time. If these cases are not dealt with properly, they may affect the performance and connectivity of these networks and their




Fig. 3: IXPs involved in the identified prefixes with route announcements mispractices.

customers and impose financial losses.

V. CONNECTION STABILITY OF IXP MEMBERS

In this section, we investigate to which extent RP can affect the stability of IXP members' interface connections to the IXP route server. Understanding how much a peering connection type is less stable than others can be helpful since downtime can cause significant financial, performance and reputational loss for ASes [43]. Based on earlier work [44], we expect to see remote interfaces being less stable than local ones because they are not physically connected at the IXPs and rely on a third-party infrastructure in which they do not have control to reach the peering facilities. Besides, since many RP resellers usually connect multiple remote networks through shared logical ports, in which higher instability to one physical interface can affect many clients.

Methodology. To analyze the connection stability of remote and local interfaces at IXPs, we collected membership data from IXP LGs every 15 minutes over two months (2022-11-19 to 2023-01-23). This analysis is only feasible in IXPs with LGs allowing each of the IXP interfaces to be classified according to their remoteness: LINX, PTT-SP, PTT-RJ, and PTT-CE. In AMS-IX, the LG presents all the IXP member interfaces under the Route Server interface address, while in NAPAfrica, Equinix-Ash, and Equinix-Chi, no LGs offer open access to BGP data (see §I). We exclude from the analysis a small fraction of IXP interfaces (1.6%) which either changed from remote to local and vice-versa or were not present during the entire collection period. We collected 7008 samples of the connected interfaces to RS on four IXPs (LINX, PTT-SP, PTT-RJ e PTT-CE) between 2022-11-19 and 2023-01-23 (see § II-B).

Interfaces switch between two states, up or down, which we use to compute the following metrics:




Fig. 4: Median uptime between failures per IXP interface. Remote interfaces stay less time continuously without failures.

- **Median Uptime Between Failures.** It represents the stability of a connection and is computed by the median of all the hours one interface stayed in the up state uninterrupted during the collection period. If the interface remains up between data collection X and $X + 1$, we add 15 minutes of uptime to the interface.
- **State changes.** The metric indicates how common errors are on IXP's interfaces. We compute it as the sum of state changes (up to down, and vice-versa), between all data collections.
- **Mean Time Between Failures.** MTBF is (a well-known metric) which we use to indicate the level of availability of interfaces. We compute it by dividing the total uptime by the number of failures.
- **Reliability.** It measures the ability of a system or component to perform its intended function for a specified period. The system's reliability can be calculated using various methods, including reliability models, failure data analysis, and reliability testing. We computed it using the MTBF metric and two different intervals, day and month:

$$Reliability = e^{-\frac{1}{MTBF} \times Time}$$

RP interfaces stay less time continuously functioning than local ones. We first evaluate whether the connection type affects the interfaces' stability, by comparing the median uptime between failures for remote and local interfaces. Figure 4 shows that the uptime for local interfaces is considerably higher than remote ones. In other words, remote interfaces tend to stay less time continuously functioning before a change of state (e.g., failure, BGP session issue). This difference is striking in the case of LINX and PTT-CE. For LINX, the disparity reaches two times higher uptime for local members (1749 hours to local vs 871 hours to remotes) when considering at least 50% of the interfaces. For PPT-CE, at least 50% of the local interfaces have uptime higher than 874 hours, while the uptime of at least 50% of the remote ones does not reach 236 hours (3.7 times lower).

Local interfaces suffer fewer state changes. Next, we look




Fig. 5: Number of state changes per IXP interface; the y axis is in log scale. Remote interfaces are subject to many more state changes.


at the state change number for each connection type. Too many state changes indicate lower stability, since the BGP session of the members is constantly shutdown and re-established. Figure 5 shows the results using a log-scale on the Y-axis. We observe that remote interfaces consistently present more state transitions than local ones, reaching up to $3.5\times$ more changes. While 75% of the local members in each IXP have no more than between 10 and 2 changes, depending on the IXP, the corresponding values for remote connections were between 16 and 5. Combined with the uptime hours previously presented, these numbers strongly indicate less connection stability for RP.

Local peers are greatly more reliable than RP. Finally, we compare the metric Reliability for the IXP interfaces, for periods of different durations, day and month. Members with higher reliability also present higher availability and less time between intermittent failures. Figures 6a and 6b show the results 1-day and 1-month, respectively. Even though reliability is directly related to the time interval, remote interfaces have lower reliability than local ones in both scenarios for all IXPs. However, the results vary considerably depending on the duration of the period considered. While in the 1-day analysis, there was little difference between local and remote (the median differences varied between 1.37% and 3.98%), in the 1-month, there are massive differences (reaching 47.05% and 33.85% for PTT-CE and LINX, respectively).


All results combined showed that, in addition to having less uptime and more state changes, remote interfaces also show less reliability over time and are less stable when compared to local ones in the analyzed IXPs. This confirms our expectations and shows there can be a substantial difference considering in some large IXPs.

VI. USAGE OF RP BGP COMMUNITIES FOR TE

BGP communities have been used in IXP to tag routes and help with traffic engineering [45]. Communities can also be used to perform TE in remote peering connections, such as



(a) Interface's reliability considering 1-day period.



(b) Interface's reliability considering 1-month period.

Fig. 6: Level of reliability using the Reliability metric. Results show that remote interfaces are less reliable than local ones, reaching a difference up to 47.05%.

tagging connections whose latency exceeds some threshold. IX.br started supporting BGP communities targeting remote networks in March 2022. This was accomplished with two changes: the RS in IXPs started tagging remote routes with *informational communities* to assist ASes in making traffic engineering decisions; and IXPs started supporting *action communities* added by the IXP members to impose their routing policies on their routes. We look at which communities networks add when sharing routes with the IXP.

Methodology. We analyze the usage of RP-related BGP communities for the three IXPs among our set that have support for such communities (PTT-SP, PTT-RJ, and PTT-CE). RP-related communities can be used to avoid route exports and to add different levels of prepending for ASes when matching certain conditions, such as RTT thresholds, packet loss thresholds, or announcement within some RIR region. The IX.br is responsible for measuring both latency and packet loss and marking up the routes with such characteristics.

Using the LG API in each IXP, we collected a routing data snapshot on 2023-01-17 from their primary IPv4 RS. The information captured for every route includes a prefix, next-hop address, AS-Path, and lists of BGP communities. We also fetched the RS configuration file containing the semantics of available informational and action BGP communities. To identify the communities defined by the IXPs, we build a *dictionary* using the LG API. We identified 96 communities related to RP on IX.br, out of the total 649 available. Their actions allow ASes to avoid route export and path-prepend once, twice or thrice to peers with RTT higher than 10, 50, 100, 150, 200 and 250 ms.

We process every route observed in the IXP BGP data for our analyses and compare the BGP communities in it with our community dictionary. We then counted the number of RP communities used in each route by each AS (i.e., if a route have two RP communities, we add two to the AS count).

Which type of BGP communities is more popular? First, we identify which specific RP-related communities are prevalent

in the IXP route server member's routes. On PTT-SP, the most common actions are *adding a single prepend* to members with RTT higher than 10ms and 50ms (4911 and 3860 instances, respectively). In contrast, the most common action in PTT-CE and PTT-RJ was to *avoid route export* for ASes with RTT higher than 10ms (2024 and 188 occurrences, respectively).

Are RP-related BGP Communities used? These fine-grained traffic engineering BGP communities for RP were introduced in March 2022 to all the 35 Brazilian IXPs. According to network operators from IX.br, this was prompted by a broad request from the network operation community. The general understanding has been that the communities would be widely used, and to confirm whether this was true or not, and to which extent, we examined the usage of RP BGP communities. Surprisingly, we discover that they are still not so widespread, at least when compared to the rest of BGP communities. For example, in PTT-SP, the fraction of RP-related action communities was under 0.27%.

We believe the low usage is because of two main reasons: the lack of knowledge about the potential impacts that remote peers may cause on performance; and the lack familiarity with the communities available at the IXPs. According to network operators, measurements to evaluate the peering performance are not proactively and regularly done to every individual peer. Instead, networks generally perform more detailed traffic engineering analysis when clients complain of reduced connectivity or performance. Unfortunately, the lack of periodic analysis and the opportunity to use RP-related BGP communities may hinder optimal peering decisions overall. As shown in [32], using remote connections to deliver traffic can impose latency penalties compared to using a local route alternative for many cases. With respect to the lack of familiarity, given that different IXPs provide different services and deploy distinct BGP communities, network operators may still need to learn about the possibility of fine-grained traffic engineering. We discussed this with the operational community, and learned that often they deploy a single configuration model: a configuration is

performed considering communities that are available at every IXP (e.g., avoid export to ASN, export only to ASN) and then replicated to all other IXPs.

Do ASes apply different RP communities in distinct IXPs to steer traffic? We now examine the numbers and characteristics of the ASes that are using RP BGP communities to their routes. We found that only a tiny number of ASes use these communities, representing 0.5% (11 ASes’ interfaces out of 2175) in PTT-SP and 1.1% (5 ASes’ interfaces out of 471) in PTT-CE. When looking further at these ASes, we can highlight that one AS (AS7195) uses RP communities for its routes on the three IXPs examined, and three ASes (AS61832, AS53202, and AS61568) use them in two IXPs (PTT-SP and PTT-CE). AS7195 (i.e., EdgeUno) is a continental company focusing on IP transit, remote peering, and cloud hosting services. It adds the community to avoid route export to ASes with RTT > 10 to, on average, 2.5% of their announced prefixes. They could represent prefixes needing better performance, which would not benefit from having long-distance networks redistributing them to their customers. The other three ASes, AS61832, AS53202, and AS61568, are ISPs connected to multiple IXPs worldwide. AS53202 and AS61568 add communities to avoid route export for ASes with RTT > 50ms to 7.4% and 0.4%, respectively, of their announced prefixes in both IXPs. AS61832, in contrast, makes a distinguished use of these actions against remote peers: in PTT-CE, AS61832 uses the “No export to RTT > 10ms” action on 74.4% of its prefixes, but on PTT-SP, it applies the community “Add one prepend to RTT > 10ms” to all its announced routes. This approach inflates the AS-Path of the routes distributed in PTT-SP and forces traffic for some of its routes to go preferably via its connection in PTT-CE.

VII. WEB PORTAL

We created a web portal, located at [46], to allow users to interactively dissect the results. The portal allows users to view the longitudinal analyses we conducted on RP deployment between August and October 2022. For each IXP, we provide results over time for the number and type of IXP interfaces, announced prefixes, routes, and the minimum/average RTT for each IXP member (see Figure 7). Additionally, users can obtain a detailed view of the announced BGP routes at IXPs, including prefix, next-hop, and AS Path information per day. Furthermore, we present results that demonstrate the potential routing impacts caused by prefixes announced by ASes connected via RP or local peering at IXPs. Users can interact with and query all of the results using a variety of selectable options.

We plan to enable data download for the community to encourage reproducibility. We invite interested IXPs to contact us if they want to have their peering infrastructure on the portal.

VIII. RELATED WORK

With the growing deployment of remote peering, there have been several efforts to investigate this interconnection practice. We divide related work into three categories: (1) methods




Fig. 7: Example of analysis in the portal’s GUI.

to identify remote peering at IXPs, (2) studies to explore implications of remote peering on the Internet, and (3) work investigating the characteristics and usage patterns of BGP communities at IXPs.

Inferring Remote Peering. Two main related methodologies have been proposed in the literature. In 2014, Castro *et al.* [16] introduced a conservative inference method based on measuring propagation delay to IXP interfaces connected to it via pings and identifying RP via a empirically obtained 10ms latency threshold. Results showed that 91% of the 22 studied IXPs showed networks connecting via RP, with potential to offload up to 25% of its transit-provider traffic via RP. In 2018, Nomikos *et al.* [29] also proposed a methodology to infer RP, combining latency measurements with additional RP features, such as port capacity and AS presence at colocation facilities. They inferred that 90% of the analyzed IXPs had more than 10% of their members using RP, with a few having up to 40% of their members as remote. In 2021, the authors extended the previous work [42], with changes in the methodology and additional analysis on Wide-Area IXPs.

Implications of Remote Peering. In 2017, Giotsas *et al.* [31] showed that the rise of remote peering made it easier for localized failures in IXP and colocation facilities to become widespread. Bian *et al.* [30] proposed a methodology to characterize anycast based on archived BGP routing information collected globally. While trying to infer anycast prefixes, the authors found that RP peering could potentially affect 38% of anycast prefixes with an average latency increase of 35.1ms. Bertholdo *et al.* [44] analyzed the stability of participants’ connections to IXP route-servers. Their results show that unstable interfaces were mainly caused by large regional ASes connected in just one IXP or ASes connected via Remote Peering at IXPs. Finally, Mazzola *et al.* [32] compared the latency impact to reach prefixes announced by remote peers at eight different IXPs using routes from remote peers, local peers, and transit providers. Results showed that, despite remote routes were generally preferred by BGP given its shorter AS path, the local routes had lower latency than the remote route in the majority of cases.

Usage of BGP Communities. In [47], Giotsas *et al.* collected communities from router server with semantics defined by the IXP in order to infer p2p links. In [48], Philipp *et al.* examined the role of route servers in IXPs, using communities for some inferences. Krenc *et al.* [49] observed announcements

at BGP collectors (e.g. RIPE and RouteViews) aiming to understand better community usage, but limited to when/how ASes *add* communities to announcements and when they *remove*. Mazzola et al [45], on the other hand, evaluated how action communities used for traffic engineering are used by ASes in IXPs and performed a characterization about BGP communities usage patterns.

IX. FINAL REMARKS

IXPs are critical infrastructures that support ever-increasing data volumes and service requirements of modern Internet services. However, the recent growth of remote peering introduces new challenges for traffic engineering because peering may no longer keep local traffic local. Our paper sheds light on the RP growth over the last year in seven IXPs and explores some of the implications on the use of RP on the Internet. We highlight some of the following key findings.

RP growth varies according to IXP characteristics. Using results from a three-month data collection and comparisons with state-of-the-art data, we found that remote interfaces have grown since 2021 in almost all IXPs evaluated. The growth was directly related to how developed and prevalent the peering infrastructure was. While the growth was lower in well-established IXPs, RP development was predominant in more emergent peering infrastructures.

Undesirable trombone prefixes may be common, due to route announcements mispractices. ASes connected to multiple IXPs with a combination of local and remote connections may unintentionally incur in route announcements mispractices. Our analysis showed that more than 60 distinct ASes announced most specific prefixes on the remote IXP instead of prioritizing their local IXP connection. We found over 37k prefixes with highly likely trombone paths, impairing peering performance.

Remote interfaces are consistently more unstable than local ones. A concern about remote peering growth at IXPs is that networks using a shared port or being geographically distant would impose higher instability on their connection to the IXP RS, affecting other members. Our results indicate that this is indeed the case in all analyzed IXPs. Remote interfaces seemed to be less reliable, presenting differences to local peers that reach up to 47.05% on a monthly analysis. Besides, the remote interfaces remain less time in up state between failures, with local interfaces staying up to $3.7\times$ more active. Lastly, remote interfaces showed up to $3.5\times$ more state changes (from up to down) than local peers.

BGP communities' usage to perform TE on remote peers is still not widespread. Many modern IXPs started to offer specific BGP communities to filter route export or perform some action (e.g., prepend) to networks connected remotely. Our analysis of the usage of these communities at three IXPs (PTT-SP, PTT-RJ, and PTT-CE) revealed that they are still negligible, representing less than 0.27% of all the action communities seen at the IXP routes.

Our findings helped characterize some of the RP implications to Internet routing and its growth in different IXPs. We

believe that considering additional IXPs, and analyzing IPv6 prefixes would improve the community's understanding of the remote peering evolution and deployment in other contexts. Besides, even though our analysis was comprehensive in including some of the largest IXPs on the planet, in Internet measurements we say it is virtually always possible to grow the dataset to include more VPs (IXPs) and for longer periods, so "it is never enough". Improving current methodologies is also crucial to promote further research on RP implications to performance and security.


REFERENCES

- [1] Internet Society, "IXPs make the Internet faster and more affordable," 2021. [Online]. Available: <https://www.internetsociety.org/issues/ixps/>
- [2] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is more to IXPs than meets the eye," *SIGCOMM CCR*, 2013.
- [3] B. Ager, N. Chatzis, A. Feldmann, N. Sarraf, S. Uhlig, and W. Willinger, "Anatomy of a large European IXP," in *SIGCOMM*, 2012.
- [4] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *IMC*, 2009, pp. 336–349.
- [5] PeeringDB, "PeeringDB. IXPs and colocation database." <https://www.peeringdb.com>, 2020.
- [6] Euro-IX, "Euro-IX. the IXP database." <https://ixpdb.euro-ix.net/en/>, 2021.
- [7] Hurricane Electric, "Internet exchange report," <https://bgp.he.net/report/exchanges>, 2021.
- [8] LINX, "LINX quick facts," 2023. [Online]. Available: <https://quickfacts.linx.net/>
- [9] —, "100GE member facing port milestone reached," 2021. [Online]. Available: <https://www.linx.net/newsletter/linx-newsletter-august-2021/>
- [10] IX.br, "IX.br members," 2018, <http://ix.br/particip/sp>.
- [11] CGI.br, "IX.br reaches mark of 10 tb/s of peak Internet traffic," 2021. [Online]. Available: <https://cgi.br/noticia/releases/ix-br-reaches-mark-of-10-tb-s-of-peak-internet-traffic/>
- [12] DE-CIX, "DE-CIX Frankfurt statistics," 2018, <https://www.de-cix.net/en/locations/germany/frankfurt/statistics>.
- [13] AMS-IX, "AMS-IX infrastructure," 2018, <https://ams-ix.net/technical/ams-ix-infrastructure>.
- [14] DE-CIX, "Connect to DE-CIX via a reseller," 2021. [Online]. Available: <https://www.de-cix.net/en/access/how-to-connect/connect-via-reseller>
- [15] AMS-IX, "Pricing," <https://www.ams-ix.net/ams/pricing>, 2021.
- [16] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois, "Remote peering: More peering without Internet flattening," in *CoNEXT*, 2014, p. 185–198.
- [17] BICS, "BICS launches remote peering to offer communications providers an affordable solution for network expansion," <https://prn.to/3WE4tJe>, 2014.
- [18] Telecomdrive Bureau, "Epsilon grows remote peering presence at LINX," <https://telecomdrive.com/epsilon-grows-remote-peering-presence-at-linx/>, 2020.
- [19] France-IX, "Angola cables launches remote peering service at France-IX Marseille," <https://www.franceix.net/en/events-and-news/presse/angola-cables-lance-un-service-de-peering-distance-france-ix-mar/>, 2017.
- [20] AMS-IX, "The AMS-IX partner program," 2018, <https://www.ams-ix.net/ams/partners>.
- [21] LINX, "Members by IP/ASN," 2021. [Online]. Available: <https://portal.linx.net/members/members-ip-asn>
- [22] M. J. Levy, "Peer locally, with local network on local IXs or not! (alternate title: Remote peering is bad for content)," https://www.globalpeeringforum.org/pastEvents/gpf14/presentations/Wed_2_MartinLevy_remote_peering_is_bad_for.pdf, 2019.
- [23] L. Almeida, "Good practices for ISP connection to international IXPs," in *IX Forum 13 Brazil*, 2019. [Online]. Available: <https://youtu.be/bIfwuFVmoil>
- [24] A. Nipper, B. Hedges, E. Verhoef, and J. Sirota, "Remote peering panel - providers," in *IX Forum 12 Brazil*, 2018. [Online]. Available: https://www.youtube.com/watch?v=HQZY-DJNixU\&ab_channel=NICbrvideos
- [25] "Panel: What are the challenges of remote peering?, European peering forum 2016, Sofia, Bulgaria." <https://www.peering-forum.eu/agenda?year=2016>, 2016.

- [26] “Panel: Remote peering, with a look at resellers as well, 29th Euro-IX Forum, Krakow, Poland,” https://www.euro-ix.net/media/filer_public/ba/61/ba61bc32-f506-4ac7-b037-4dab8549c8a2/e-an-20161108-remote_peering.pdf, 2016.
- [27] Z. Ali, “Panel: Remote peering, ripe 65, Amsterdam, Netherlands,” https://ripe65.ripe.net/presentations/196-RIPE65_EIXWG_ZA.pdf, 2012.
- [28] W. B. Norton, “The great remote peering debate,” http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2012/9/18_The_Great_Remote_Peering_Debate.html, 2012.
- [29] 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 *IMC*, New York, NY, USA, 2018, p. 265–278.
- [30] R. Bian, S. Hao, H. Wang, A. Dhamdere, A. Dainotti, and C. Cotton, “Towards passive analysis of anycast in global routing: Unintended impact of remote peering,” *SIGCOMM CCR*, vol. 49, no. 3, p. 18–25, Nov. 2019.
- [31] V. Giotsas, C. Dietzel, G. Smaragdakis, A. Feldmann, A. Berger, and E. Aben, “Detecting peering infrastructure outages in the wild,” in *SIGCOMM*, 2017, pp. 446–459.
- [32] F. Mazzola, P. Marcos, I. Castro, M. Luckie, and M. Barcellos, “On the latency impact of remote peering,” in *Passive and Active Measurement*, O. Hohlfeld, G. Moura, and C. Pelsner, Eds. Cham: Springer International Publishing, 2022, pp. 367–392.
- [33] “IX.br - Brazilian public IXP project,” <https://ix.br/intro>, 2020.
- [34] E. Carisimo, J. M. D. Fiore, D. Dujovne, C. Pelsner, and J. I. Alvarez-Hamelin, “A first look at the Latin American IXPs,” *SIGCOMM CCR*, 2020.
- [35] S. H. B. Brito, M. Santos, R. dos Reis Fontes, D. A. L. Perez, and C. E. Rothenberg, “Dissecting the largest national ecosystem of public Internet eXchange points in Brazil,” in *PAM*, 2016.
- [36] PCH, “Packet clearing house, Internet exchange directory,” <https://www.pch.net/ixp/dir>, 2020.
- [37] —, “PCH. daily routing snapshots,” <https://www.pch.net/resources/>, 2021.
- [38] IANA, “IANA autonomous system (AS) numbers,” <https://www.iana.org/assignments/as-numbers/as-numbers.xml>, 2020.
- [39] M. Luckie, “Scamper: A scalable and extensible packet prober for active measurement of the Internet,” in *IMC*, 2010, pp. 239–245.
- [40] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe, “Towards IP geolocation using delay and topology measurements,” in *IMC*, 2006, pp. 71–84.
- [41] B. Trammell and M. Kühlewind, “Revisiting the privacy implications of two-way Internet latency data,” in *PAM*, 2018, pp. 73–84.
- [42] V. Giotsas, G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, and X. Dimitropoulos, “O peer, where art thou? uncovering remote peering interconnections at IXPs,” *IEEE/ACM Transactions on Networking*, vol. 29, no. 1, pp. 1–16, 2021.
- [43] T. Holterbach, S. Vissicchio, A. Dainotti, and L. Vanbever, “Swift: Predictive fast reroute,” in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 460–473. [Online]. Available: <https://doi.org/10.1145/3098822.3098856>
- [44] L. M. Bertholdo, J. M. Ceron, L. Z. Granville, and R. van Rijswijk-Deij, “Forecasting the impact of ixp outages using anycast,” in *TMA*, 2021.
- [45] F. Mazzola, P. Marcos, and M. Barcellos, “Light, camera, actions: Characterizing the usage of ixps’ action bgp communities,” in *Proceedings of the 18th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 196–203. [Online]. Available: <https://doi.org/10.1145/3555050.3569143>
- [46] “Remote peering portal,” <https://bit.ly/3LC4RG4>, 2023, [Online].
- [47] V. Giotsas, S. Zhou, M. Luckie, and k. claffy, “Inferring multilateral peering,” in *ACM CoNEXT*, 2013, pp. 247–258.
- [48] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger, “Peering at peerings: On the role of IXP route servers,” in *ACM IMC*, 2014, pp. 31–44.
- [49] T. Krenc, R. Beverly, and G. Smaragdakis, “As-level bgp community usage classification,” in *ACM IMC*, 2021, p. 577–592.



Fabrício Mazzola received the B.S. degree in computer science from UFRGS, in 2019, and the Ph.D. degree from UFRGS in 2023. His research interests include Internet measurements, Internet routing, networking, and distributed systems.



Augusto Setti received the B.S. degree in computer science from FURG in 2023. His research interests include Internet measurements, Internet routing, networking, and distributed systems.



Pedro Marcos Pedro Marcos is an assistant professor at the Federal University of Rio Grande (FURG), Brazil. He has a PhD in Computer Science from the Federal University of Rio Grande do Sul (UFRGS), also in Brazil. Pedro’s research interests are in the Internet interconnection and routing aspects, focusing on measuring the interactions between the different entities that compose the Internet and extracting valuable information to improve Internet operations. Pedro’s research has been funded by LAC-NIC/FRIDA and by FAPESP/MCTI/MCom/CGI.br.



Marinho Barcellos has been at the University of Waikato, New Zealand, since 2019. Before, he was an Associate Professor at UFRGS, Brazil. Barcellos has contributed with research in topics including network security, Internet measurements, and programmable data planes. He was Program Chair of ACM CoNEXT 2022 and General Chair of ACM SIGCOMM 2016, a member of the SIGCOMM Executive Committee 2017–2021, and currently Chair of the Advisory Board of PACMNET and of the SIGCOMM CARES Committee.