

On the Latency Impact of Remote Peering

Fabricio Mazzola¹, Pedro Marcos², Ignacio Castro³, Matthew Luckie⁴, and
Marinho Barcellos⁴

¹ UFRGS fmmazzola@inf.ufrgs.br

² FURG pbmarcos@furg.br

³ QMUL i.castro@qmul.ac.uk

⁴ University of Waikato mjl@wand.net.nz, marinho.barcellos@waikato.ac.nz

Abstract. Internet Exchange Points (IXPs) play an essential role in the Internet, providing a fabric for thousands of Autonomous Systems (ASes) to interconnect. Initially designed to keep local traffic local, IXPs now interconnect ASes all over the world, and the premise that IXP routes should be shorter and faster than routes through a transit provider may not be valid anymore. Using BGP views from eight IXPs (three in Brazil, two in the U.S., and one each in London, Amsterdam, and Johannesburg), a transit connection at each of these locations, and latency measurements we collected in May 2021, we compare the latency to reach the same addresses using routes from remote peers, local peers, and transit providers. For four of these IXPs, at least 71.4% of prefixes advertised by remote peers also had a local peering route, BGP generally preferred the remote route due to its shorter AS path, but the local route had lower latency than the remote route in the majority of cases. When a remote route was the only peering route available at an IXP, it had slightly lower latency than a corresponding transit route available outside the IXP for >57.6% of the prefixes for seven of the eight IXPs.

1 Introduction

How to deliver traffic is an increasingly complex aspect of the Internet today as many applications generate large volumes of traffic and have strict service requirements. As a consequence, Autonomous Systems (ASes) are constantly increasing their interconnection capacities and expanding their footprint. Internet Exchange Points (IXPs) are key elements of this process, as they can shorten Internet paths and reduce interconnection cost [31, 17, 4, 10]. As of May 2021, there were more than 800 IXPs deployed worldwide [46, 22, 29]. The largest IXPs have surpassed 1000 members [36, 35, 32] and 10 Tbps of peak traffic [16, 19, 36, 7].

An original motivation of IXPs was to keep local traffic local by having ASes physically present at an IXP facility. However, IXPs no longer only interconnect members physically present at IXP facilities. Remote peering – where an AS is not physically present at an IXP facility and reaches the IXP through a layer-2 provider – allows ASes to widen their peering footprint with a quicker setup, no additional hardware, and lower installation costs compared to local peering [20, 9,

15]. For example, ASes from 85 different countries connect to LINX remotely [36] as of May 2021. To cope with the demand for peering, IXPs and remote peering resellers have expanded their offerings [12, 48, 25] with some IXPs having up to 55 official partners selling remote peering services [32, 8, 37]. Network operators prefer to steer traffic through IXPs instead of transit providers because of the reduced transit and operational costs [18, 21]. However, the ability to interconnect with remote members at IXPs adds complexity to traffic engineering choices.

Given the public debate about remote peering performance [34, 6, 41, 2, 1, 5, 43], which is currently data-poor, and to understand the latency properties of BGP routes at IXPs, we analyze latency and latency variability when using different interconnection methods (remote peering, local peering, and transit) to reach addresses in prefixes announced by remotely connected members in eight IXPs identified in Table 1. These eight IXPs include six of the world’s ten largest IXPs by membership, and are deployed in five countries (three in Brazil, two in the U.S., and one each in London, Amsterdam, and Johannesburg). Our contributions are as follows.

First, we find that inferring remote ASes using the state-of-the-art methodology [42] based on latency and colocation data is insufficient for some IXPs (§3). Incomplete and/or inaccurate colocation data in regions, such as Latin America, yields a high number of unknown inferences (more than 68.6% for three IXPs). Because we need to infer which ASes are remotely connected to a given IXP in order to identify prefixes announced by remote ASes, we infer geographically distant remote ASes [15] and complement these inferences with ground-truth data (§4). We found that at least 26.2% of all ASes connected to major IXPs, such as PTT-SP and AMS-IX, were remotely connected members in May 2021. These remotely connected members announced fewer than $\approx 15\%$ of the prefixes visible at the IXP, for most IXPs.

Next, we classify prefixes announced by remote ASes in BGP data collected from PCH and IXP looking glass servers. We focus our analysis on prefixes that had routes available through both remote and local ASes (§5). We found that for 82.5% of these prefixes, on average, the AS path for the route from the remote peer was shorter or had the same length as the route from the local peer in the four IXPs with most of these prefixes (LINX, AMS-IX, Eq-Ash, and Eq-Chi). Using BGP views from RouteViews peers, we confirmed that remote routes tended to be preferred by BGP. However, our latency measurements indicate that the local route had a lower latency in most cases.

Finally, we examine the prefixes announced exclusively by remote members at IXPs (§6). Our findings suggest that remote routes can have lower latency to reach addresses in prefixes announced by remote ASes when compared with a transit route, though not by a considerable margin for six out of eight IXPs: using the remote route or the transit had a latency difference no higher than 5ms for 78.1% of the measured prefixes. However, for NAPAfrica in South Africa, remote peering routes had a lower latency than transit routes, with a latency benefit of more than 40ms for 81.4% of the measured prefixes.

IXP	Location	Observed Interfaces	BGP VPs		Reseller
			LG	PCH	Ground Truth
PTT-SP	Sao Paulo, BR	2,169	✓	✗	✓
LINX	London, UK	911	✓	✓	✓
AMS-IX	Amsterdam, NL	907	✓	✓	✗
NAPAfrica	Johannesburg, ZA	542	✗	✓	✗
PTT-RJ	Rio de Janeiro, BR	462	✓	✗	✓
PTT-CE	Fortaleza, BR	395	✓	✗	✓
Eq-Ash	Ashburn, VA, US	365	✗	✓	✗
Eq-Chi	Chicago, IL, US	259	✗	✓	✗

Table 1: The eight IXPs analyzed in our study, along with the availability of BGP VPs and ground truth data on remote peering.

2 Measurement Architecture

In this section, we discuss the measurement architecture we used. First, we present the IXPs we measured (§2.1). Next, we describe the datasets we used in our work, including the IXP ground truth and BGP routing data (§2.2). Finally, we characterize the vantage points (VPs) along with the active measurements we performed (§2.3).

2.1 Peering Infrastructure Selection

To identify networks connected via remote peering, and prefixes and routes announced via remote peering, we need peering infrastructures that have (1) publicly available BGP routing data, and (2) an active measurement VP attached to the IXP switching fabric. Table 1 presents the eight selected IXPs where we had both BGP routing data and active measurement capability. These IXPs include six of the world’s ten largest IXPs by membership [22, 29] 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) [3, 14, 13]. The eight IXPs together comprise 3466 unique ASes.

2.2 Datasets

Remote Peering Reseller Ground Truth Data. We obtained ground truth information for the ASes remotely connected via resellers for four of the analyzed IXPs: LINX, PTT-SP, PTT-RJ, and PTT-CE. The data set contains information about the ASN and IP interface of remote ASes reaching the IXPs through shared ports or VLANs associated with resellers. For the PTT IXPs, we obtained the ground truth data from their operators on the 20th April 2021. The set of ASes reaching LINX through resellers or locally connected to the IXP is publicly available at their member portal [37] (collected on 5th May 2021). LINX

representatives confirmed that ASes with Port Type labeled as ConneXions correspond to ASes using resellers. The ground truth for the four IXPs comprise a list of 1634 unique ASes using remote peering through resellers.

Membership and Interface Addresses. To identify the peering router’s IP and ASN of all members at each IXP, we combine multiple public data sources for all IXPs except for LINX, which publishes this information through their member portal [37]. We collected membership data and subnet information from Euro-IX [22] and the publicly available databases of Hurricane Electric (HE) [29], PeeringDB (PDB) [46], and Packet Clearing House (PCH) IXP Directory [44]. In cases of conflicts, we followed the preference ordering described in [42]: *Euro-IX* > *HE* > *PDB* > *PCH*.

BGP Datasets and Sanitization. We used two sources of 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 [45]. 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 AMS-IX and LINX, we observed 3.4–3.9x more routes and 1.9–2.0x more prefixes from PCH than from LG views. As our goal is to understand the latency difference between routes announced at IXPs by different peering types (i.e., remote and local peering), we prefer the dataset from PCH whenever it is available, as it provides us with better visibility of the IXP routes (*PCH* > *LG*). On IXPs with only LG views (PTT sites), we have observed that LGs are configured to output only the best routes at the time of our BGP routing data collection, lowering the number of cases with multiple routes for the same prefix. Additionally, we collected BGP data from RouteViews collectors at each IXP to understand the types of routes that RouteViews peers actually chose. For each IXP, we obtained a BGP snapshot corresponding to the same period our measurements were performed (5-6 May 2021). We discarded: (i) routes with artifacts, such as reserved/unassigned ASes [30] and loops; (ii) prefixes shorter than /8 or longer than /24.

2.3 Data plane measurements

Vantage Points. At each IXP listed in Table 1, we used RouteViews collectors which were directly connected to the IXP LAN to conduct active measurements using scamper [38]. Figure 1 illustrates the measurement architecture of each RouteViews collector and how we used them to conduct active measurements.

Measurement Types. We conducted two types of measurements. In the first, we measured the latency to each IXP member’s peering router. These measurements use the IP address that the collector has in the IXP LAN (X.2), so that probes and responses cross the IXP LAN, as in when we probe X.3 in Figure 1. In the second, we measured the path and latency to IP addresses within prefixes announced by each IXP member. Note that these prefixes are peering routes, and not transit routes. These measurements go out via a selected IXP member




Fig. 1: Architecture of our data plane measurements. We used RouteViews collectors with an interface connected to a transit provider and 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.3) used the collector’s IP address in the IXP LAN (X.2), so the probes and responses crossed the IXP LAN. Other measurements used the Transit IP address T.1 as the source address, and were delivered to each IXP member using the layer-2 address corresponding to their IXP LAN IP address (e.g., X.4).

(e.g. AS B, using the layer-2 address of X.4 in Figure 1) but used the collector’s Transit IP address T.1 as the source address, so that we could receive a response. This strategy allowed us to maintain the same return path from the probed address back to the RouteViews collector, while varying the forward path as we selected different IXP members. We provide further details about the measurement methodology in the sections describing our results (§4,5,6).

3 Challenges in Inferring Remote Peering

Our method needs to know which networks connect via remote connections at IXPs. However, there are two different notions of remote peering.

Notions of Remote Peering. Conversations with IXP and reseller representatives revealed that notions of remote peering varied. Some considered remote peering based on the AS connection type (e.g., using shared ports via resellers), regardless of location (even those in the same city as the IXP). Other representatives viewed remote peering based on the geographical distance to the IXP.

Figure 2 shows different ways that ASes can connect to IXPs. *Local ASes* connect directly to an IXP switch using a router deployed in the same facility as the switch (ASes A, B, C). ASes can also connect via resellers. *Resellers* provide ports and transport to the IXP, usually connecting the routers of the remote ASes to the IXP switches via layer-2 transport. ASes located close to the IXP (ASes D, E) use resellers to lower peering equipment and installation costs. Resellers can also bridge large geographical distances by connecting members located far from the IXP (ASes F, G). Finally, an AS may also connect remotely without




Fig. 2: ASes connect to IXPs via local (ASes A, B, C) and remote connections, either via a reseller (ASes D, E, F, G) or by purchasing transport from the remote location to the IXP switch (AS H). Remote networks can be physically located near the IXP (ASes D, E) or be geographically distant (ASes F, G, H).

any reseller ports, using its own port at the IXP and purchasing transport to the port from the remote location (AS H).

Available Data Limits Accuracy of Remote Peering Inferences. The current state-of-the-art methodology for inferring remote peering proposed by Giotsas *et al.* [27] infers remote peering (1) through a reseller and/or (2) geographically distant from the IXP. The method combines delay measurements with additional features, such as port capacity and AS presence at colocation facilities; if an AS is not present in one of the feasible IXP facilities, their method infers the AS is remotely connected. We used available ground truth (§2.2) for four IXPs (LINX, PTT-SP, PTT-RJ, and PTT-CE) and applied their method to all interfaces connected to these IXPs.

We implemented the four steps from the Giotsas *et al.* [27] method. The first step (*ping measurement campaign*) measures the latency to IXP member interfaces from a VP within the IXP. Using the scamper probers on the RouteViews collectors (§2.3), we performed delay measurements to the peering interfaces of IXP members every two hours for two days, and discarded measurements where the replies might have come from outside the peering infrastructure because they had an IP-TTL value that appeared to have been decremented (i.e., the received IP-TTL was not 64 or 255). The second step (*colocation-informed RTT interpretation*) computes a geographical area where the IXP member router could be located using an AS to colocation facility mapping obtained from PeeringDB and IXP websites. Then, we obtained publicly available RIPE Atlas IPv4 traceroute measurements collected on the same days as our ping campaign and applied step 3 (*multi-IXP router inference*) and step 4 (*finding remote peers via port capacities and lack of private connectivity*) to complete the methodology.

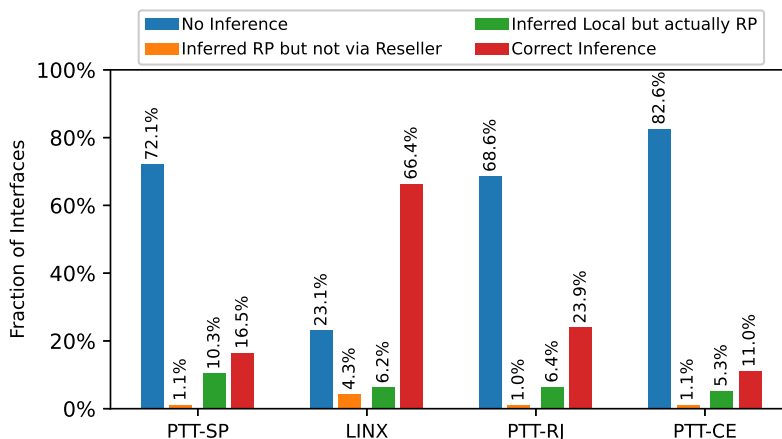


Fig. 3: Classification of interfaces we obtained when we applied our implementation of the current state-of-the-art methodology for inferring remote peering [27]. The high percentage of no inference for the three Brazilian IXPs was a consequence of the method’s high reliance on public information (PeeringDB) which was not widely available for members of Brazilian IXPs.

Figure 3 presents the results we obtained. In [27], public information about AS presence at colocation facilities was missing for $\approx 25\%$ of remote peers and $\approx 18\%$ of local peers. When we reproduced the study, the number of unknown inferences for LINX was low and the fraction of remote and local interfaces inferred was similar with the published work [27], which we hypothesize was because the PeeringDB coverage for LINX members that had valid information about presence in IXP facilities was high (83.0%). The case for Brazilian IXPs was different. For PTT-SP and PTT-CE, only 27.0% of the members had PeeringDB entries that reported both the IXP and facilities where they were present, leading the current state-of-the-art method to only classify 17.1%, on average, of the interfaces at the Brazilian IXPs. This low classification was because few ASes connected to the Brazilian IXPs shared their information in PeeringDB. Openly publishing peering data has only recently been encouraged by IXP operators in Brazil as best practice [39].

In addition, 5.3-10.3% of the interfaces inferred as local peerings were actually remote, according to ground truth. We believe the misclassification was related to incorrect information about the presence of ASes in colocation facilities. In many cases, an AS using a reseller recorded the facility their reseller connected to in their PeeringDB record, leading the method [27] to infer the AS was locally connected. The other 1.0-4.3% of interfaces inferred as remote were correct, but they did *not* observably connect to the IXP via a reseller. In summary, the

IXP	Interfaces (I)	Reseller Remote Peering		P also Local
		Routes (R)	Prefixes (P)	
PTT-SP	1,265 of 2,169 (58.3%)	28,385 of 154,509 (18.4%)	27,148 of 158,880 (17.1%)	577 of 27,148 (2.1%)
LINX	189 of 911 (20.7%)	107,533 of 1,018,593 (10.6%)	90,633 of 486,171 (18.6%)	71,357 of 90,633 (78.7%)
PTT-RJ	172 of 462 (37.2%)	5,525 of 128,961 (4.3%)	5,502 of 128,478 (4.3%)	25 of 5,502 (0.5%)
PTT-CE	214 of 395 (54.2%)	7,098 of 26,025 (27.3%)	7,095 of 26,012 (27.3%)	10 of 7,095 (0.1%)

Table 2: Number and percentage of routes and prefixes announced by members using a shared port via resellers. Members connecting to an IXP via a reseller announced fewer routes than members connecting locally. LINX had a considerable percentage (78.7%) of the same prefixes being announced by both remote and local peers.

methodology of [27] may not be suitable for accurately inferring remote peering for IXPs that have incomplete or inaccurate publicly available data.

4 Remote Peering at IXPs

Inferring remote peering (RP) based solely on reseller connections is imprecise, as it ignores geographically distant ASes not using reseller ports which also incur a latency penalty. However, examining only remote peers that are geographically distant overlooks RP through resellers. This diversity in the notion of RP led us to evaluate RP both by (1) connection type (*Reseller RP*), and (2) geographical distance to the IXP (*Geographical RP*).

To identify members using Reseller RP, we used ground truth that identified members connected to an IXP using a reseller for four IXPs (§2.2). To infer members using Geographical RP at all eight IXPs, we used the method in [15], 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 §2.2. We performed latency measurements to these addresses on 5-6 May 2021. From each RouteViews scamper instance, we probed each interface every two hours for two days, 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 [33, 49]. We adopted [15]’s method because its latency threshold alone yielded ac-

IXP	Geographical Remote Peering			
	Interfaces (I)	Routes (R)	Prefixes (P)	P also Local
PTT-SP	681 of 2,169 (31.4%)	20,289 of 158,932 (12.8%)	19,612 of 154,561 (12.7%)	1,118 of 19,612 (5.7%)
LINX	121 of 911 (13.3%)	92,975 of 1,015,040 (9.2%)	71,452 of 482,643 (14.8%)	65,060 of 71,452 (91.1%)
AMS-IX	238 of 907 (26.2%)	67,397 of 978,225 (6.9%)	63,323 of 485,933 (13.0%)	56,503 of 63,323 (89.2%)
NAPAfrica	40 of 542 (7.4%)	7,256 of 159,100 (4.6%)	7,252 of 144,513 (5.0%)	88 of 7,252 (1.2%)
PTT-RJ	61 of 462 (13.2%)	3,861 of 129,135 (3.0%)	3,850 of 128,652 (3.0%)	355 of 3,850 (9.2%)
PTT-CE	139 of 395 (35.2%)	6,870 of 26,610 (25.8%)	6,869 of 26,597 (25.8%)	8 of 6,869 (0.1%)
Eq-Ash	35 of 365 (9.6%)	49,157 of 967,133 (5.1%)	46,752 of 525,688 (8.9%)	43,455 of 46,752 (92.9%)
Eq-Chi	17 of 259 (6.6%)	8,382 of 347,788 (2.4%)	8,120 of 271,855 (3.0%)	5,795 of 8,120 (71.4%)

Table 3: Number and percentage of routes and prefixes announced by inferred geographically remote members. Members we infer to connect to an IXP from some geographical distance announced fewer routes than members connecting locally. LINX, AMS-IX, Eq-Ash, and Eq-Chi all had a considerable percentage (71.4%) of the same prefixes announced by both remote and local peers.

curate results for single metropolitan area peering infrastructures [27], which is the case of the analyzed IXPs in our work (see §2.1).

To further assess the correctness of our inferences – and similar to step 2 in [27] (*colocation-informed RTT interpretation*) – we obtained the colocation facilities of each of the eight 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.

4.1 Remotely Connected Members

Tables 2 and 3 summarize the number and percentage of interfaces connected via remote peering at each IXP.

Reseller RP. We observed a large percentage of Reseller RP at the three Brazilian IXPs, representing more than 37.2% of their member base (Table 2). According to network operators at these IXPs, the IXPs’ members are spread across Brazil, which has a large land mass, and members connect to the IXP to reach

large content and cloud providers. We encountered a substantially smaller fraction of Reseller RP at LINX (20.7%).

Geographical RP. We inferred that at least a quarter of the ASes connected to PTT-CE, AMS-IX, and PTT-SP were Geographical RP (Table 3). The remaining IXPs had less than 13.3% Geographical RP members inferred. This indicates that even though remote peering is widely used at IXPs (as shown by [42, 27]), a considerable fraction of member ASes are physically connected to the IXPs or closely located to them.

4.2 Remotely Announced Prefixes and Routes

For each IXP, we examined the proportion of BGP routes in the IXP routing data, and the percentage of prefixes that could be reached via both local and remote peers (i.e. local and remote routes). To identify whether routes were local or remote, we compared routes observed in the BGP data with inferred remote networks. We labeled routes as remote when the next-hop IP interface belonged to the IXP subnet and belonged to the list of networks we classified as remote.

We show the percentage of remote interfaces, routes, and prefixes we inferred at each IXP, along with absolute numbers, in Table 2 for Reseller RP and in Table 3 for Geographical RP. In all IXPs, remote peers announced proportionally fewer routes than local peers, both for Reseller RP (Table 2) and Geographical RP (Table 3). For example, in PTT-SP and PTT-RJ, the fraction of peers using Reseller RP was 3.2x and 8.7x higher than the fraction of routes they announced, respectively. For LINX, the 189 remote peers (20.7% of all interfaces) announced just 10.6% of the routes (107k/1M). For the Geographical RP inferences, PTT-RJ shows the highest difference between the fraction of remote interfaces and remote routes (4.4x), with 61 (13.2%) remote interfaces announcing just 3.0% of all routes (67k/981k). The results suggest that remotely connected ASes tend to announce fewer prefixes than local networks into the IXP. Conversations with IXP network operators revealed that remote peers mainly use their connections to obtain specific content not available at their local IXPs.

Interestingly, we observed a sizeable percentage of prefixes announced by both remote and local peers in some IXPs. At LINX, AMS-IX, Eq-Ash, and Eq-Chi, at least 71.4% of remotely announced prefixes also had a route announced by a local peer in May 2021. These cases can be a problem for traffic engineering since remote peering is invisible to Layer-3 protocols, and there is no guarantee that BGP will choose the lowest latency route.

5 Choosing Between Remote and Local Peering

Sending traffic via an IXP rather than a transit provider can potentially offer lower latency by keeping local traffic local. However, it is currently unknown whether remote peering might hinder that benefit. The geographical distance of an AS or its connection type can introduce undesired latency implications to peering. In this section, we first investigate whether remote routes have shorter

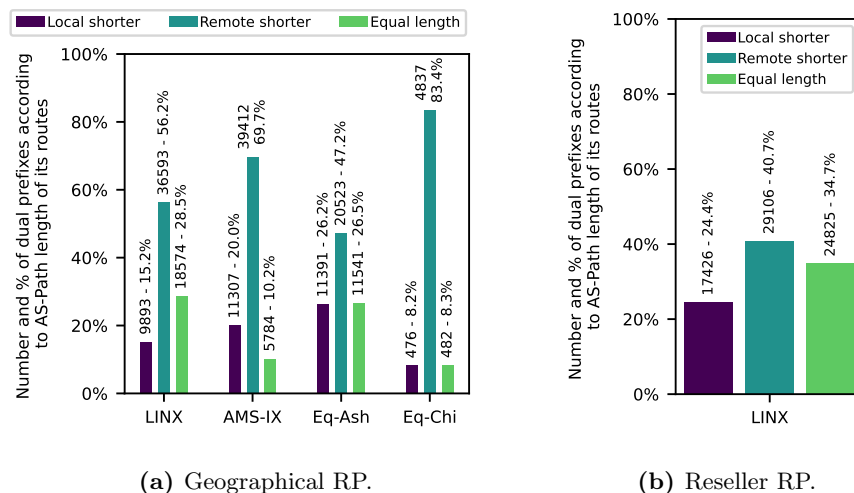


Fig. 4: AS path lengths of prefixes reachable via both remote and local peers. Regardless of the method to infer RP, the majority of prefixes with both local and remote routes had remote routes with an AS path length shorter or the same length as the local route, and therefore likely chosen by BGP, a hypothesis we have confirmed using data from RouteViews peers (§5.2).

AS paths than local routes (§5.1). Next, we analyze routing data from RouteViews collectors at each IXP and find that remote routes are chosen by BGP in the majority of cases (§5.2). Then, we measure latency, and compare the latency of remote routes with the latency of local routes (§5.3). Finally, we measure the latency variation of each route and evaluate if remote peering introduces higher latency variability compared to the local route (§5.4).

5.1 Which Route had the Shortest AS Path?

Prefixes with both local and remote routes can be problematic for traffic engineering because an AS might choose a higher-latency route with a shorter AS path, since AS path length is the BGP second tie-breaker (after local preference) [47]. To examine whether this was the case, we compared the AS path length of routes for every prefix announced via remote and local peerings seen in IXP routing data, reporting the analysis for the IXPs that had a considerable number of these cases, namely LINX, AMS-IX, Eq-Ash, and Eq-Chi (§4.2). To compare routes, we selected the shortest AS path route of each type, local and remote. In order to observe the path lengths as they appear in the routing data, we do not reduce paths with AS path prepending.

Remote Routes had Shorter AS Paths than Local Routes. Figure 4 shows the percentage of prefixes with a shorter AS path length per peering type.

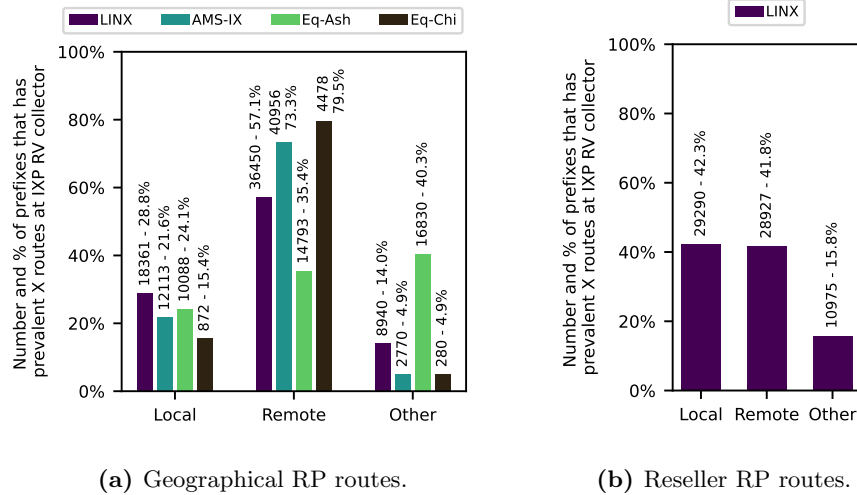


Fig. 5: The type of selected route by peers of RouteViews collectors at each IXP for prefixes with both local and remote routes. The remote route was more likely to be selected for Geographical RP. For Reseller RP, preference between remote and local routes was the same – $\approx 42\%$.

In Figure 4a, most Geographical RP routes (an average of 82.5%) had shorter (or equal) AS path lengths, with the remaining 17.4% having a shorter AS path for the local route. Thus, BGP may choose a remote route over a local route if BGP uses AS path length as a tie breaker. The difference in AS path lengths for most prefixes with different length routes was a single ASN (82.1%, 79.0%, 73.9%, 89.9% for LINX, AMS-IX, Eq-Ash, and Eq-Chi). This happened because the local route was usually announced by large transit providers connected to the IXPs, which include the transit provider’s ASN in the path.

Figure 4b, shows the distribution when looking at the Reseller RP inferences for LINX. We only show LINX because the PTT-SP, PTT-RJ, and PTT-CE results are similar but from a much smaller number of prefixes associated with resellers (fewer than 600 prefixes each). Again, we find that the remote routes tend to have shorter AS paths – 40.7% of remote prefixes had the shortest AS path, whereas only 24.4% of local prefixes had the shortest AS path. The difference in path length for most prefixes with different length routes was also a single ASN – 62.5% of the prefixes with different AS path lengths for LINX.

5.2 Are Shorter AS Path Remote Routes Chosen?

Next, we want to understand the extent to which remote routes are preferred over local routes. We analyze how frequently the remote routes appear in routes shared by RouteViews peers in the IXPs (§2.2). For each prefix with both local

IXP	Reseller RP		Geographical RP	
	Remote lower	Local lower	Remote lower	Local lower
PTT-SP	131 (51.1%)	125 (48.9%)	112 (20.9%)	423 (79.1%)
LINX	21,001 (45.5%)	25,155 (54.5%)	13,721 (33.0%)	27,903 (67.0%)
AMS-IX	-	-	6,644 (38.8%)	10,477 (61.2%)
NAPAfrica	-	-	14 (28.0%)	36 (72.0%)
PTT-RJ	10 (76.9%)	3 (23.1%)	53 (26.1%)	150(73.9%)
PTT-CE	4 (57.1%)	3 (42.9%)	4 (66.7%)	2 (33.3%)
Eq-Ash	-	-	2,230 (9.4%)	21,561 (90.6%)
Eq-Chi	-	-	830 (25.0%)	2,486 (75.0%)

Table 4: Number of prefixes that had lower latency via remote or local peers. Generally, a route from a local peer had lower latency than a route from a remote peer to reach addresses in the same prefix.

and remote routes announced, we find all the routes the RouteViews peers see and compare them with the routes in the dataset used in the previous section. A remote (or local) route is prevalent among RouteViews peers when most peers see the route. It was also possible that most peers reported a different route, neither local nor remote, which we did not observe in the IXP routing data that we used.

Figure 5 shows how often each kind of route was preferred according to RouteViews peers: the local, the remote, or a different route which was not in our data set (*other* in Fig 5). We find that the remote route was more commonly chosen. For Geographical RP routes at LINX, AMS-IX, and Eq-Chi, these remote routes were chosen for at least 57.1% of the prefixes, compared to 28.8% or fewer local routes, and 14.0% or fewer other routes. When a remote route was prevalent among RouteViews peers, the remote route had the shortest AS path among the routes (local, other) for most prefixes (83.5%, 90.0%, 81.3%, and 98.5% of these prefixes, respectively, for LINX, AMS-IX, Eq-Ash, and Eq-Chi). When local routes were prevalent, they were not always the shortest AS path routes available, and the IXP had a remote route with shorter or equal AS path length (64.5%, 39.7%, 76.6%, and 61.0%, respectively, for LINX, AMS-IX, Eq-Ash, and Eq-Chi). This suggests that operators might have been using local policy to prefer local routes so that the remote routes with shorter AS paths were not selected by BGP.

For Reseller RP routes (Figure 5b) the situation was different: preference between remote and local routes was similar ($\approx 42\%$), with other paths accounting for the remaining 15.8%. For 75.2% of the prefixes with remote routes prevalent, the remote paths had shorter AS paths. When local routes were prevalent, 58.4% of prefixes had a remote alternative with shorter or equal AS path length available at the IXP.

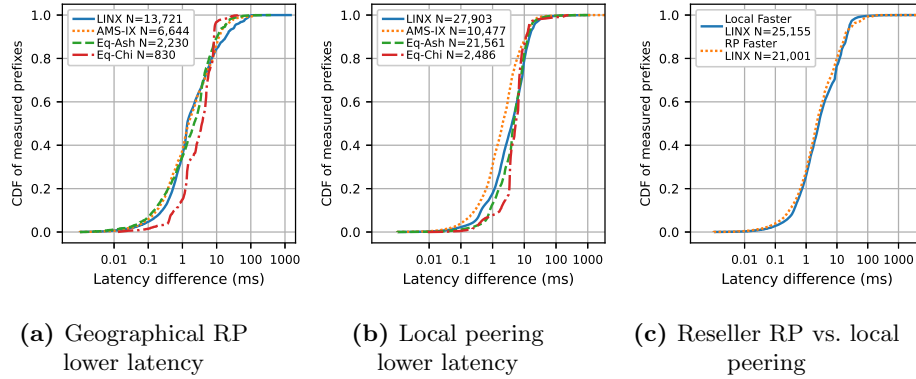


Fig. 6: Latency difference between remote and local routes measured by end-to-end latency to reach an address in a remote prefix. For Geographical RP, when local routes had lower latency, the advantage compared to the remote route was more than 5ms for at least 44.7% of prefixes in three IXPs

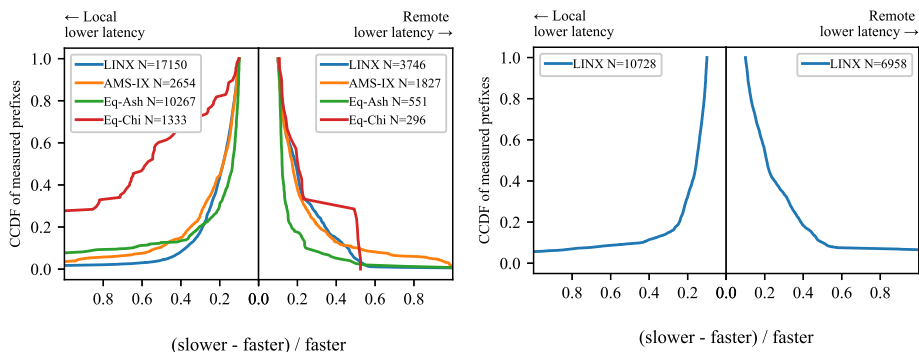
5.3 Is There a Latency Penalty Using a Remote Route?

Considering the current preference for peers to select remote routes, we wanted to understand whether they were also the best route latency-wise. We performed active measurements, using traceroutes toward IP addresses within the prefixes set seen in IXP routing data. Since we did not have a pre-selected list of responding servers, we initially probed the first ten addresses in the IP block of the prefix, followed by thirty IP addresses randomly selected, from a system external to the IXP. Because not every prefix had a responsive address, the set of measured prefixes is smaller than the original set of prefixes. We then ran ICMP-Paris traceroute measurements to these IP addresses from RouteViews VPs in the IXPs over two days and compared the latency of the remote and local routes, provided we had obtained at least five responses from addresses in each type of route. Because a prefix can have multiple remote or local routes, we used the lowest latency measured when comparing each route type – i.e., we compared the lowest latency local and remote routes.

Local Routes had Predominantly Lower Latency than Remote Routes.

Table 4 shows the number (percentage) of prefixes where a remote route had lower latency than the local routes. Looking at Geographical RP first, local routes had lower latency than remote routes for nearly all analyzed IXPs. When focusing on the IXPs with a higher prevalence of prefixes with both local and remote routes (e.g., LINX, AMS-IX, Eq-Ash, and Eq-Chi), up to 90.6% of the measured prefixes had lower latency using a local route. Similarly, for the Reseller RP inferences in LINX, the majority of prefixes also had a lower latency local route.

The previous analysis was binary – which route had the lowest latency. We now analyze the differences in latency. Figure 6 shows the latency difference



(a) Geographical RP perspective.

(b) Reseller RP perspective.

Fig. 7: Relative comparison of end-to-end latencies. For Geographical RP, when either the local or remote route had lower latency, the route had up to 30.7% lower latency than when compared with the other route type for 75.1% of prefixes in three IXPs. For Reseller RP, when a remote route had lower latency, its advantage over the local route tended to be higher than vice-versa.

between remote and local routes. The figures have a different number of points, as the number of prefixes with lower latency for remote or local routes shown in Table 4 are different. Figure 6a shows that when a Geographical RP provided a route with lower latency than the local route, the advantage was small: for at least 72.9% of the prefixes, the latency benefit of the remote route was restricted to 5ms or less for three IXPs. In contrast, when the local route was faster, as shown in Figure 6b, the latency advantage was more pronounced. For at least 44.7% of prefixes in three IXPs, the latency benefit for the local route was more than 5ms when compared to the corresponding remote route. When looking at Reseller RP for LINX in Figure 6c, we observe that the distribution of latency differences was similar for both remote and local routes, with nearly 20% of the prefixes having a latency difference above 10ms.

Figure 7 shows a CCDF of the *relative* latency difference between remote and local routes when the latency differed by more than 5ms. The left side of the figure shows the prefixes where the local route had lower end-to-end latency than the remote route, while the right side shows when the remote route had a lower latency than the local route. The x-axis represents how much faster one route was when compared with the other. For example, an x equals 0.2 shows that for some fraction of prefixes (in the y-axis), one type of route was 20% faster than the other type of route. We see on the left side of Figure 7a that local routes are up to 30% faster (better) for 75.1% of prefixes observed in three IXPs. For Eq-Chi, 50% of prefixes are at least 57.8% faster (better) via a local route than using the remote one. On the right side, we see a similar pattern, where remote routes have RTTs less than 30.7% lower (better) for 75.1% of prefixes observed in three IXPs. The situation was different for Reseller RP inferences for LINX. As

IXP	Total prefixes	Remote lower latency, <i>longer</i> AS path length	Remote lower latency, <i>equal</i> AS path length	Local lower latency, <i>longer</i> AS path length	Local lower latency, <i>equal</i> AS path length
LINX	41,624	1,177 (2.8%)	2,185 (5.2%)	12,950 (31.1%)	9,636 (23.2%)
AMS-IX	17,121	1,397 (8.2%)	657 (3.8%)	4,798 (28.0%)	1,828 (10.7%)
Eq-Ash	23,791	270 (1.1%)	674 (2.8%)	9,547 (40.1%)	5,579 (23.5%)
Eq-Chi	3,316	57 (1.7%)	161 (4.9%)	2,149 (64.8%)	111 (3.3%)

Table 5: Breakdown per IXP when comparing remote and local routes for each prefix in terms of latency and AS path length – Geographical RP only. A large number of local routes had lower latency but had a longer AS path than the remote route.

shown in Figure 7b, when the remote routes via reseller had lower latency, they were at least 20% faster for 54.6% of prefixes, while when the local route had lower latency, they were at least 20% faster for only 32.5% of measured prefixes. In summary, the results suggest that with proper configuration and knowledge about these cases, ASes can decide which route to select and steer their traffic, potentially enabling better performance according to their specific goals.

The Path with Lowest Latency was Not Always Preferred by BGP.

Table 5 shows the percentage of prefixes where the route with lowest latency would not match the route specified in a BGP tie-breaker. We observed a small percentage of prefixes where the remote route had lower latency but also had a longer AS path when compared to the local route (no more than 8.2%). In contrast, there were proportionally more cases of prefixes for which the local route had lower latency but a longer AS path than the remote route, varying from 28% (AMS-IX) up to 64.8% (Eq-Ash). When both the remote and local routes had the same path length, the local peering predominantly had a latency advantage over the remote routes despite the latency benefit not being higher than 5ms for most routes. The results for Reseller RP, obtained from LINX, follow a similar pattern (as in Table 5) and are omitted. In summary, the results indicate that the shortest AS path route may often not match the route with the lowest latency.

5.4 Do Remote Routes Have More Latency Variability than Local Routes?

In discussion with network operators, there was a concern about potential latency variability that could be introduced by a layer-2 connection or the geographic distance separating the AS’s router to the IXP. To compare the relative latency variability of remote routes over local routes, we performed active measurements by sending at least 120 ping packets from the scamper prober at the IXP RouteViews node to an address in each of the prefixes with both local and remote routes seen in Table 4 over ≈ 4 days (depending on the size of the IXP):




Fig. 8: Latency variability to remotely announced prefixes via remote and local routes. The latency variability to reach remote destinations was similar for both local and remote routes, suggesting that reseller connections and geographical distance had limited impact on latency variability.

at least 60 packets via the local route and at least 60 via the remote route. We computed the latency standard deviation for the best remote and local routes for the prefixes we used in the latency comparison in the previous section.

Remote and Local Routes had Similar Latency Variability. Figures 8a and 8b show the latency variability was similar between remote and local routes. Regardless of peering type or remote peering perspective, 75% of the prefixes had less than 10ms of latency variability. More specifically, for three of the four analyzed IXPs, the same fraction of prefixes had latency variability below 5ms. The results indicate that variability was *not* a distinguishing feature at least for the IXPs we considered.

6 Does Remote Peering have Lower Latency than Transit?

When remotely announced prefixes do not also have routes from a local peer at the IXP, ASes must decide between delivering their traffic via the remote peer at the IXP or using a transit provider. Which connection type presents the lower latency to reach these prefixes? Discussions in the network operator community concern whether remote peering is an inferior alternative to transit in both latency and connection stability [40, 34].

To assess whether remote peering or transit had lower latency to reach addresses in prefixes exclusively announced at an IXP via remote peers, we performed traceroute measurements through the remote peers at eight IXPs, as well as a transit provider from the same location (§6.1). We compared the latency

IXP	Reseller RP latency		Geographical RP latency	
	Remote lower	Transit lower	Remote lower	Transit lower
PTT-SP	8,886 (74.2%)	3,085 (25.8%)	5,657 (72.0%)	2,205 (28.0%)
LINX	10,342 (77.7%)	2,973 (22.3%)	2,724 (71.0%)	1,108 (29.0%)
AMS-IX	-	-	2,651 (57.6%)	1,950 (42.4%)
NAPAfrica	-	-	1,787 (98.1%)	35 (1.9%)
PTT-RJ	1,929 (64.9%)	1,045 (35.1%)	1,113 (59.6%)	754 (40.4%)
PTT-CE	3,014 (71.7%)	1,190 (28.3%)	2,648 (71.3%)	1,065 (28.7%)
Eq-Ash	-	-	708 (28.9%)	1,740 (71.1%)
Eq-Chi	-	-	1,204 (94.6%)	69 (5.4%)

Table 6: Latency comparison between remote peering or transit, showing the number of prefixes with lower latency. For Reseller RP, in four IXPs, at least 64.9% of the prefixes had lower latency via Reseller RP routes than via transit. For Geographical RP, seven of eight IXPs had at least 57.6% of prefixes with lower latency via remote peering routes than via transit.

variability of both RP and transit to reach these remotely announced prefixes (§6.2).

6.1 Does Transit Offer Lower Latency than Remote Peering?

We collected latency measurements to addresses in prefixes announced by remote peers both using the remote routes and a transit route using a similar approach to §5.3 – we first identified remote prefixes without a local route and responsive IP addresses in each prefix. We collected at least five latency samples for each remote prefix using a remote peer and the transit provider.

Table 6 shows the number of probed prefixes per IXP, along with the connection type (remote or transit) with lowest latency. Note that the number of prefixes with a measurement is lower than the number of prefixes observed in the routing table (§5.3), as in some cases we failed to identify a responsive address for the prefix. The remote route had lower latency for most prefixes: 57.6% of the prefixes had lower latency with Geographical RP routes for seven out of eight IXPs, and 64.9% for Reseller RP.

Remote Routes can have a Substantial Latency Advantage. Figure 9a and 9b show the absolute latency difference for Geographical RP. Figure 9a shows that some remote routes had latencies substantially lower than the transit alternative in some IXPs. In NAPAfrica, 81.4% of remote routes with lower latency than transit had at least 40ms lower latency. When we discussed our results with resellers, they suggested that high IP transit prices, along with poor ISP interconnectivity and performance in Africa, made remote peering a lower latency and cheaper option, in line with the published literature [28, 23, 24]. For the remaining IXPs, the latency difference between remote routes and

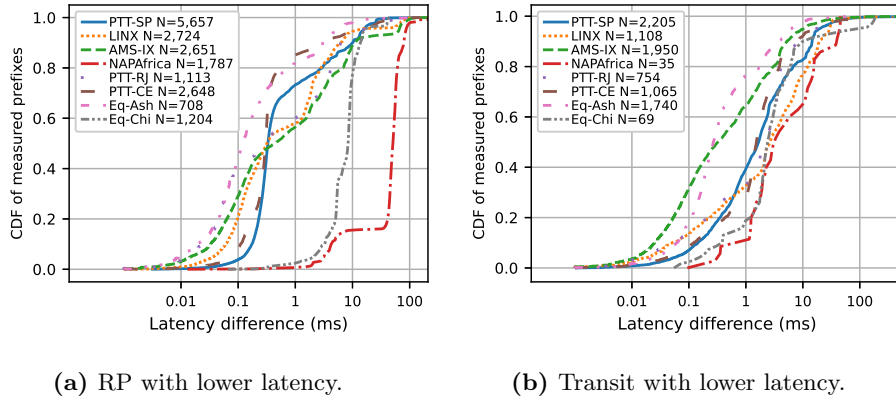


Fig. 9: Latency difference between Geographical RP and transit provider routes measured by latency to addresses in remote prefixes. Remote peering had a substantial advantage for a few IXPs (NAPAfrica, Eq-Chi), but not as a substantial advantage for others (less than 5ms for 78.1% of measured prefixes).

transit was not substantial. Regardless of which route had lower latency, in six IXPs, we observed that the latency difference was below 5ms for at least 78.1% of the measured prefixes.

Figure 10a and 10b show the results for Reseller RP. Figure 10a suggests that any latency advantage of remote peering was not substantial. For more than 67.2% of remote routes with lower latency, the latency advantage was within 1ms. In comparison, Figure 10b suggests that when transit was faster for three out of four IXPs, the latency advantage was a bit higher: in at least 53.1% of transit routes with lower latency, the advantage was more than 1ms.

6.2 RTT Variability of Remote Prefixes

In §5.4, the latency variability to reach addresses using either remote or local routes was similar. To understand if using a transit provider introduces more latency variability, we performed ping measurements to exclusively announced prefixes seen at Table 6. Similar to the previous measurements, we sent at least 120 ping probes from to each prefix over up to 4 days (depending on the size of the IXP): 60 (at least) via the transit provider and 60 (at least) via the remote route. We then computed the latency standard deviation among the ping probes for the measurements via remote peering and transit.

Transit and Remote Peering had Similar Latency Variability. Figures 11a and 11b show the latency variability for remote peering. The latency variability to reach prefixes exclusively announced at an IXP via a remote peer was equivalent for both remote and transit. PTT-SP and Eq-Ash were the only IXPs where a fraction of the prefixes had higher latency variability (see Figure 11a).

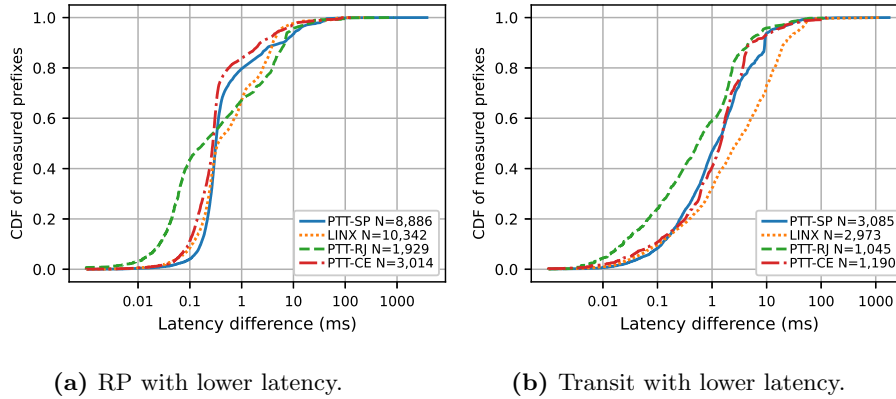


Fig. 10: Latency difference between Reseller RP and transit measured by the latency to reach remote prefixes. When Reseller RP had lower latency, the latency advantage was not substantial (below 1ms for over 67.2% of the measured prefixes). When transit routes had lower latency, the latency advantage was a bit higher (more than 1ms for 53.1% of the measured prefixes in three IXPs).

Still, for all the IXPs, the standard deviation for 75% of the prefixes was below 10ms. We observed a similar trend for Reseller RP inferences, where resellers and transit had comparable latency variability.

7 Related Work

With the growing deployment of remote peering, there have been several efforts to investigate this interconnection practice. We divide related work into two categories: (1) methods to identify remote peering at IXPs, and (2) studies to explore implications of remote peering on the Internet.

Inferring Remote Peering. Two main related methodologies have been proposed in the literature. In 2014, Castro *et al.* [15] introduced a conservative inference method based on measuring propagation delay to IXP interfaces connected to it via pings. Responses to ping probes sent to IXP interfaces that presented latency more than 10ms and whose IP-TTL had not been decremented were classified as remote. The authors reported that 91% of the 22 studied IXPs showed networks connecting via remote peering. Further, using ground-truth traffic from a National Research and Education Network, the paper demonstrated that a network could offload up to 25% of its transit-provider traffic via remote peering.

In 2018, Nomikos *et al.* [42] also proposed a methodology to infer remote peering. Using ground-truth data from seven IXPs, the authors showed that latency alone was not sufficient to make accurate inferences in some cases, such as IXPs with switching fabrics distributed across different countries. The paper proposed combining latency measurements with additional remote peering fea-

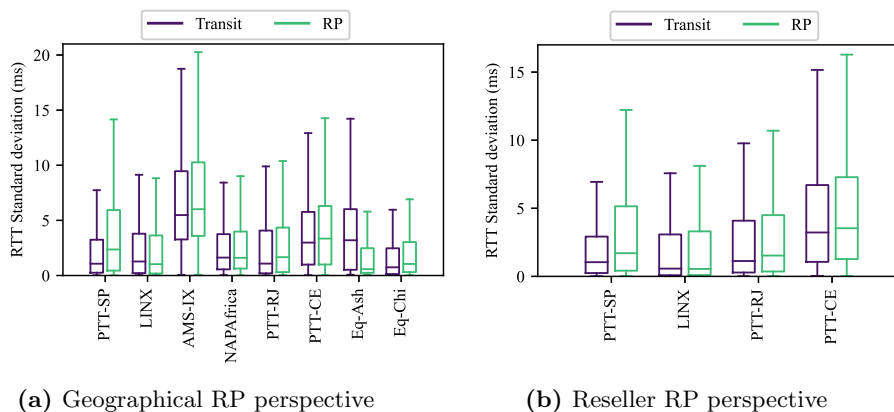


Fig. 11: Latency variability to remotely announced prefixes via remote peers and transit providers. The latency variability to reach addresses in remote prefixes was similar between transit and remote peering in all IXPs (latency standard deviation less than 10ms for 75% of measured prefixes), suggesting that neither transit or the remote peering had a substantial effect on latency variability.

tures, such as port capacity and AS presence at colocation facilities, to obtain a more trustworthy inference methodology. Their method computes the geographical area where an IXP member’s router could be located and associates the router with the feasible facilities that a local peering could use. They used this method to infer RP in 30 IXPs worldwide, and reported that 90% of the analyzed IXPs had more than 10% of their members using remote peering, with two of the largest IXPs in terms of members (DE-CIX and AMS-IX) having up to 40% of remote members. In 2021, the authors extended the previous work [27], with changes in the methodology and additional analysis on Wide-Area IXPs.

Implications of Remote Peering. In 2017, Giotsas *et al.* [26] proposed a methodology for detecting peering infrastructure outages, such as colocation facilities and IXPs. The authors reported that the rise of remote peering made it easier for localized failures in IXP and colocation facilities to become widespread. For two outages observed in London (2016), they showed that more than 45% of the interfaces related to the affected links were from outside England, with more than 20% of them being located outside Europe.

In 2019, Bian *et al.* [11] proposed a methodology to characterize anycast based on archived BGP routing information collected globally. While trying to infer anycast prefixes, the authors found that remote peering caused a significant element of inaccuracy in their method. They reported that RP can cause unintended consequences on anycast performance and potentially affect 19.2% of the anycast prefixes. Active measurements found that 38% of such prefixes were indeed impacted with an average latency increase of 35.1ms.

8 Limitations and Future Work

Route Selection. Route selection is a complex problem faced by network operators, as there are many metrics that could affect traffic delivery performance. In this paper we focused on investigating AS-Path length and latency (§5 and §6). Analyzing routing by other metrics is challenging, because of the lack of reliable information in publicly available datasets regarding transit costs, economic decisions, and local preference.

Path Relevance. Despite analyzing a considerable number of remote routes, one question that stands is the relevance of such paths, both in terms of destination popularity and traffic carried. Investigating this problem requires data protected by confidentiality terms and not publicly available (e.g., IXP traffic data) for all IXPs. Additionally, many IXPs do not have an implemented and automated way to measure traffic flowing through each announced route, and are able to only share aggregated traffic per AS.

IPv6. We focused on IPv4 IXP interfaces and IPv4 announced prefixes. Six out of eight RouteViews collectors used in our work did not have IPv6 transit that would enable us to study IPv6. We hope to investigate IPv6 routes in the future.

Distributed IXPs. Our analysis considered only IXP facilities within a single metropolitan area, avoiding wide-area peering infrastructures. Our method would not work for distributed IXPs because we used a delay-based methodology and ground truth data to infer remote peering [15]. In distributed IXPs, local members connected at facilities far from the IXP region could present very high latencies and, consequently, be inferred as remote.

Future Work. Our findings help to characterize the latency impact of remote peering. Beyond the analysis we performed, we believe that considering additional IXPs, and analyzing IPv6 prefixes would improve the community’s understanding of remote peering in the context of other available route types. Improving current methodologies is also crucial to promote further research on RP implications to performance and security. Our methodology used a 10ms latency threshold to infer geographical remote peering. While the threshold is conservative, it was adequate to identify networks connected far from IXPs. However, a deeper analysis of the impact of using different latency thresholds (e.g., 2ms and 5ms) is needed. We also plan to leverage our ground truth data about networks connected via resellers to investigate better approaches to infer remote peering connections.

9 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 shed light on the latency impact

of reaching addresses in remotely announced prefixes at IXPs via remote routes, local peering routes, and regular transit, and had the following key findings.

Inferring remote peering is still challenging. Using IXP ground truth and delay measurements, we showed that current state-of-the-art methodologies have limitations. We show that relying on public network data can result in a sizable fraction of unknown inferences for some IXPs, caused by public data being unavailable for some classes of networks. Compared to the European, American, and Asian IXPs evaluated in [42], reduced data availability in some regions, such as Latin America, limits the accuracy of remote peering inferences.

The route preferred by BGP is not always the lowest latency route. When investigating the use of remote routes in the BGP routing, we detected a high prevalence of prefixes announced both by remote and local peerings in four IXPs (LINX, AMS-IX, Eq-Ash, and Eq-Chi). We found that most remote routes for these prefixes had a shorter or equal AS path length compared to the available local routes and tended to be preferred by the peers of RouteViews collectors. Despite being shorter and indeed preferred, they were not necessarily the lowest latency route. For at least 61.2% of these prefixes in seven IXPs, the local route had lower latency compared to the geographically distant remote peering routes.

Remote routes are a reliable option to deliver traffic at IXPs. Some prefixes have only remote routes at IXPs, and ASes must choose between delivering their traffic via remote peering or a transit provider. Our measurements suggest that relying on remote routes can be an advantageous option for end-to-end latencies. In some scenarios (NAPAfrica and Eq-Chi), remote routes at the IXPs had considerably better latency results when compared to transit, showing latency improvements of at least 40ms for 81.4% of the measured prefixes, when the remote route was faster than transit. For the other six IXPs, we observed that the latency difference of using the remote route or the transit was no higher than 5ms for 78.1% of the measured prefixes.

The connection type or geographical distance does not directly impact latency variability for remote routes. A concern about remote peering growth at IXPs is that networks using a reseller or being geographically distant limits the original performance benefits of peering. Our measurements suggest that remote peering does not introduce additional latency variability to reach addresses in these prefixes. For 75% of the remote prefixes, we observed less than 10ms of latency variability for remote connections.

Acknowledgments

We thank the anonymous reviewers and our shepherd, Shuai Hao, for their valuable feedback on our paper. We are also very thankful to all network operators and NIC.br for their valuable insights regarding the deployment and different aspects of remote peering. This study was sponsored (in part) by the Coordenacao de Aperfeicoamento de Pessoal de Nivel Superior – Brasil (CAPES) –

Finance Code 001, by the EPSRC (EP/W032473/1, EP/S033564/1), and by National Science Foundation (NSF) grant CNS-2120399. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the sponsors.

References

1. 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)
2. Panel: What are the challenges of remote peering?, European peering forum 2016, Sofia, Bulgaria. <https://www.peering-forum.eu/agenda?year=2016> (2016)
3. IX.br - Brazilian public IXP project. <https://ix.br/intro> (2020)
4. Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., Willinger, W.: Anatomy of a large European IXP. In: SIGCOMM (2012)
5. Ali, Z.: Panel: Remote peering, ripe 65, Amsterdam, Netherlands. https://ripe65.ripe.net/presentations/196-RIPE65_EIXWG_ZA.pdf (2012)
6. Almeida, L.: Good practices for ISP connection to international IXPs. In IX Forum 13 Brazil (2019), <https://youtu.be/bIfwuFVm0iI>
7. AMS-IX: AMS-IX infrastructure (2018), <https://ams-ix.net/technical/ams-ix-infrastructure>
8. AMS-IX: The AMS-IX partner program (2018), <https://www.ams-ix.net/ams/partners>
9. AMS-IX: Pricing. <https://www.ams-ix.net/ams/pricing> (2021)
10. Augustin, B., Krishnamurthy, B., Willinger, W.: IXPs: Mapped? In: IMC. pp. 336–349 (2009)
11. Bian, R., Hao, S., Wang, H., Dhamdhere, A., Dainotti, A., Cotton, C.: Towards passive analysis of anycast in global routing: Unintended impact of remote peering. SIGCOMM CCR **49**(3), 18–25 (Nov 2019)
12. BICS: BICS launches remote peering to offer communications providers an affordable solution for network expansion. <https://www.prnewswire.com/news-releases/bics-launches-remote-peering-to-offer-communications-providers-an-affordable-solution-for-network-expansion-248317651.html> (2014)
13. Brito, S.H.B., Santos, M., dos Reis Fontes, R., Perez, D.A.L., Rothenberg, C.E.: Dissecting the largest national ecosystem of public Internet eXchange points in Brazil. In: PAM (2016)
14. Carisimo, E., Fiore, J.M.D., Dujovne, D., Pelsser, C., Alvarez-Hamelin, J.I.: A first look at the Latin American IXPs. SIGCOMM CCR (2020)
15. Castro, I., Cardona, J.C., Gorinsky, S., Francois, P.: Remote peering: More peering without Internet flattening. In: CoNEXT. p. 185–198 (2014)
16. CGL.br: IX.br reaches mark of 10 tb/s of peak Internet traffic (2021), <https://cgl.br/noticia/releases/ix-br-reaches-mark-of-10-tb-s-of-peak-internet-traffic/>
17. Chatzis, N., Smaragdakis, G., Feldmann, A., Willinger, W.: There is more to IXPs than meets the eye. SIGCOMM CCR (2013)

18. Cloudflare: Bandwidth costs around the world. The Cloudflare Blog (2016), <https://blog.cloudflare.com/bandwidth-costs-around-the-world/>
19. DE-CIX: DE-CIX Frankfurt statistics (2018), <https://www.de-cix.net/en/locations/germany/frankfurt/statistics>
20. DE-CIX: Connect to DE-CIX via a reseller (2021), <https://www.de-cix.net/en/access/how-to-connect/connect-via-reseller>
21. Dr Peering: What are the economics of Internet peering? (2012), <https://drpeering.net/FAQ/What-are-the-economics-of-peering.php>
22. Euro-IX: Euro-IX. the IXP database. <https://ixpdb.euro-ix.net/en/> (2021)
23. Fanou, R., Valera, F., Dhamdhere, A.: Investigating the causes of congestion on the African IXP substrate. In: IMC. pp. 57–63 (2017)
24. Formoso, A., Chavula, J., Phokeer, A., Sathiaselan, A., Tyson, G.: Deep diving into Africa’s inter-country latencies. In: INFOCOM. pp. 2231–2239 (2018)
25. 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)
26. Giotsas, V., Dietzel, C., Smaragdakis, G., Feldmann, A., Berger, A., Aben, E.: Detecting peering infrastructure outages in the wild. In: SIGCOMM. pp. 446–459 (2017)
27. Giotsas, V., Nomikos, G., Kotronis, V., Sermpezis, P., Gigis, P., Manassakis, L., Dietzel, C., Konstantaras, S., Dimitropoulos, X.: O peer, where art thou? uncovering remote peering interconnections at IXPs. *IEEE/ACM Transactions on Networking* **29**(1), 1–16 (2021)
28. Gupta, A., Calder, M., Feamster, N., Chetty, M., Calandro, E., Katz-Bassett, E.: Peering at the Internet’s frontier: A first look at ISP interconnectivity in Africa. In: PAM. pp. 204–213 (2014)
29. Hurricane Electric: Internet exchange report. <https://bgp.he.net/report/exchanges> (2021)
30. IANA: IANA autonomous system (AS) numbers. <https://www.iana.org/assignments/as-numbers/as-numbers.xml> (2020)
31. Internet Society: IXPs make the Internet faster and more affordable (2021), <https://www.internetsociety.org/issues/ixps/>
32. IX.br: IX.br members (2018), <http://ix.br/particip/sp>
33. Katz-Bassett, E., John, J.P., Krishnamurthy, A., Wetherall, D., Anderson, T., Chawathe, Y.: Towards IP geolocation using delay and topology measurements. In: IMC. pp. 71–84 (2006)
34. Levy, M.J.: 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)
35. LINX: 100GE member facing port milestone reached (2021), <https://www.linx.net/newsletter/linx-newsletter-august-2021/>
36. LINX: LINX quick facts (2021), <https://quickfacts.linx.net/>
37. LINX: Members by IP/ASN (2021), <https://portal.linx.net/members/members-ip-asn>
38. Luckie, M.: Scamper: A scalable and extensible packet prober for active measurement of the Internet. In: IMC. pp. 239–245 (2010)
39. Mendes, J.L.: IX forum 14: PeeringDB (2020), https://forum.ix.br/files/apresentacao/arquivo/1020/IX_Forum_14_Sao_Paulo-peeringdb-20201204-V3.pdf

40. NANOG: Virtual or remote peering. <https://mailman.nanog.org/pipermail/nanog/2017-August/091963.html> (2017)
41. Nipper, A., Hedges, B., Verhoef, E., Sirota, J.: Remote peering panel - providers. In IX Forum 12 Brazil (2018), https://www.youtube.com/watch?v=HQZY-DJN1xU&ab_channel=NICbrvideos
42. Nomikos, G., Kotronis, V., Sermpezis, P., Gigis, P., Manassakis, L., Dietzel, C., Konstantaras, S., Dimitropoulos, X., Giotsas, V.: O peer, where art thou? uncovering remote peering interconnections at IXPs. In: IMC. p. 265–278. New York, NY, USA (2018)
43. Norton, W.B.: The great remote peering debate. http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2012/9/18_The_Great_Remote_Peering_Debate.html (2012)
44. PCH: Packet clearing house, Internet exchange directory. <https://www.pch.net/ixp/dir> (2020)
45. PCH: PCH. daily routing snapshots. <https://www.pch.net/resources/> (2021)
46. PeeringDB: PeeringDB. IXPs and colocation database. <https://www.peeringdb.com> (2020)
47. Rekhter, Y., Hares, S., Li, T.: A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Jan 2006)
48. Telecomdrive Bureau: Epsilon grows remote peering presence at LINX. <https://telecomdrive.com/epsilon-grows-remote-peering-presence-at-linx/> (2020)
49. Trammell, B., Kühlewind, M.: Revisiting the privacy implications of two-way Internet latency data. In: PAM. pp. 73–84 (2018)