

O Peer, Where Art Thou?

Uncovering Remote Peering Interconnections at IXPs

George Nomikos
FORTH, Greece
gnomikos@ics.forth.gr

Petros Gigis
FORTH / University of Crete, Greece
gkigkis@ics.forth.gr

Stavros Konstantaras
AMS-IX, Netherlands
stavros.konstantaras@ams-ix.net

Vasileios Kotronis
FORTH, Greece
vkotronis@ics.forth.gr

Lefteris Manassakis
FORTH, Greece
leftman@ics.forth.gr

Xenofontas Dimitropoulos
FORTH / University of Crete, Greece
fontas@ics.forth.gr

Pavlos Sermpezis
FORTH, Greece
sermpezis@ics.forth.gr

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

Vasileios Giotsas
Lancaster University, England
v.giotsas@lancaster.ac.uk

ABSTRACT

Internet eXchange Points (IXPs) are Internet hubs that mainly provide the switching infrastructure to interconnect networks and exchange traffic. While the initial goal of IXPs was to bring together networks residing in the same city or country, and thus keep *local traffic local*, this model is gradually shifting. Many networks connect to IXPs without having physical presence at their switching infrastructure. This practice, called *Remote Peering*, is changing the Internet topology and economy, and has become the subject of a contentious debate within the network operators' community. However, despite the increasing attention it attracts, the understanding of the characteristics and impact of remote peering is limited. In this work, we introduce and validate a heuristic methodology for discovering remote peers at IXPs. We (i) identify critical remote peering inference challenges, (ii) infer remote peers with high accuracy (>95%) and coverage (93%) per IXP, and (iii) characterize different aspects of the remote peering ecosystem by applying our methodology to 30 large IXPs. We observe that remote peering is a significantly common practice in all the studied IXPs; for the largest IXPs, remote peers account for 40% of their member base. We also show that today, IXP growth is mainly driven by remote peering, which contributes two times more than local peering.

CCS CONCEPTS

• **Networks** → **Network measurement; Network architectures; Network properties;**

ACM Reference Format:

George Nomikos, Vasileios Kotronis, Pavlos Sermpezis, Petros Gigis, Lefteris Manassakis, Christoph Dietzel, Stavros Konstantaras, Xenofontas Dimitropoulos, and Vasileios Giotsas. 2018. O Peer, Where Art Thou? Uncovering

Remote Peering Interconnections at IXPs. In *2018 Internet Measurement Conference (IMC '18)*, October 31–November 2, 2018, Boston, MA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3278532.3278556>

1 INTRODUCTION

Internet eXchange Points (IXPs) are crucial components of today's Internet ecosystem [25, 29, 37, 38], that provide infrastructure for the direct interconnection (*peering*) of Autonomous Systems (ASes). Currently, there exist more than 700 IXPs around the world, with more than 11K member networks (i.e., *peers*); these correspond to approximately 20% of the total number of ASes [11, 15, 16]. The largest IXPs host more than 800 networks each [1, 7], and handle aggregate traffic that peaks at or exceeds 6 Tbps [3, 8].

IXPs were originally created to locally interconnect ASes at layer-2 (L2), and *keep local traffic local* [39]. Under this model, networks peer at IXPs to *directly connect* with each other and avoid connections through third parties, and thus reduce costs, improve performance (e.g., lower latency), and better control the exchanged traffic [26, 67]. However, the ever-increasing traffic flowing at the edge of the Internet, creates pressure for denser and more diverse peering that challenges the traditional IXP model. As a result, the IXP ecosystem is undergoing a fundamental shift in peering practices to respond to these requirements: networks may establish peering connections at IXPs from *remote* locations, to broaden the set of networks they reach within one AS-hop [41, 69], either over a (owned or rented) "long cable" or over *resellers* that provide ports on the IXP and L2 access through their own network [13, 18]. This practice contradicts the traditional view of IXPs as local hubs of direct peering and is commonly referred to as *Remote Peering* [67] by IXP operators, where "remote" denotes a distant and/or indirect IXP connection:

DEFINITION 1. *Remote Peering (RP) is when a network peers at an IXP without having physical presence in the IXP's infrastructure and/or through a reseller.*

While RP has been actively advertised by IXPs, it has also fired up a heated debate within the operators' community [23, 67]. The *proponents* of RP highlight the benefits in connectivity and cost reduction for the IXP members, whereas the *opponents* emphasize on the risks and implications for network performance and resilience.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

IMC '18, October 31–November 2, 2018, Boston, MA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5619-0/18/10...\$15.00

<https://doi.org/10.1145/3278532.3278556>

Irrespective of which side in this debate one stands for, the reality is that RP is fundamentally changing the IXP peering landscape, with unclear effects on Internet economics and performance. Today we lack the tools and techniques to answer even simple questions, such as “Which peers of an IXP are remote and which are local?”. The answer to this question could significantly benefit Internet operations and drive routing policies and peering decisions (e.g., eyeballs or content providers that seem local at an IXP may not be local). Such knowledge is therefore important both for IXP operators to understand the characteristics of their member base, and IXP members to perform e.g., traffic engineering (TE) based on peering policies. Moreover, it enables researchers to explore different facets of RP ecosystems.

In this paper, we propose a methodology to infer RP, and analyze its main characteristics. Our primary objective is to enable transparency, a property which is desired by all stakeholders, regardless of which side they pick in the RP debate. We first provide the necessary background on this debate as well as related work in Section 2. After presenting our measurement datasets (Section 3), we make the following contributions:

Identify inference challenges (Section 4). We first identify the difficulties of inferring RP, by collecting and analyzing a best-effort validation dataset of remote/local peers in 15 large IXPs. We show that inference based exclusively on latency measurements, as proposed by Castro *et al.* [36], is not capable of accurately inferring RP at scale.

Infer remote peers (Section 5). We design a novel methodology to infer whether a peer is remote or local to an IXP. Due to the involved complexity and challenges, we take into account multiple dimensions of peering, such as latency, colocation and IXP facility information, IXP port capacity and router connectivity, and combine them to achieve an accurate inference. Comparing our inferences against validation data shows that our approach achieves a 95% accuracy and 93% coverage, while the corresponding percentages of the state-of-the-art [36] are 77% and 84%, respectively.

Characterize remote peering (Section 6). We apply our methodology to 30 large IXPs, and analyze characteristics of RP. While an extensive evaluation of RP characteristics and implications is outside the scope of this work, we consider use cases that exhibit the applicability of our inference approach. We find that RP is prevalent today, with 28% of the peers being remote. Our results also show that today, IXP growth is mainly driven by remote peering, which contributes *two times more* than local peering with respect to the number of new IXP members.

We further discuss relevant insights which arise from our study, including potential implications of RP (Section 7). Finally, we describe follow-up research directions, such as traffic analysis and a large-scale longitudinal study (Section 8).

2 BACKGROUND & RELATED WORK

Peering at IXPs. ASes connect and exchange traffic (*i.e.*, *peer*) with each other via bi- or multi-lateral setups at IXPs, which operate L2 switching platforms. Typically, ASes become *members* of an IXP by connecting to its infrastructure through their own router(s),

colocated at the facility where the IXP has presence. This enables them to peer with other IXP members.

Remote peering at IXPs. Remote peering does not require physical presence of networks’ routing equipment in the IXP fabric [63]. The connection is performed through: (i) *resellers* [71] of IXP ports that connect the remote peer’s router(s) to the IXP switches, (ii) *L2 connections* (“long cables”) to the IXP facility (with ports bought by the peer itself), either with privately owned cables or by using a carrier, and (iii) *IXP federations* [4, 9], *i.e.*, IXPs belonging to the same organization (like DE-CIX Frankfurt and DE-CIX New York), which are interconnected so that local peers of one IXP are remote to the other and vice versa¹.

Wide-area IXPs and Remote Peering. Some IXPs are geographically distributed entities, possessing switching infrastructure in multiple facilities in different metropolitan areas²/countries. We call such cases, where the IXP’s L2 network spans large geographical areas, *wide-area* IXPs. An example of a wide-area IXP is NL-IX [14], spanning the European continent³. The members of a wide-area IXP are local peers, as long as they are directly patched to the switching infrastructure of at least one facility of the IXP (see Definition 1); otherwise they are remote. Note that such IXP setups can heavily complicate remote peering inferences (see Section 4.2).

The remote peering debate. The increasing attention that remote peering is drawing has also given rise to a recent debate within the networking community [23], placing emphasis on the impact of remote peering on Internet routing and economics.

Remote peering is good! There are several advantages and new possibilities for *networks* peering remotely:

- *Monetary savings.* CAPEX is reduced since there is no need for additional routing equipment, or colocation and installation fees [67, 68]. Remote peering can also be an option for offloading transit traffic [36].
- *Increased connectivity.* Networks can easily establish direct connections with more peers (e.g., content providers present at remote IXPs), and have better control over traffic routed from / towards them.

For the *IXP*, remote peering leads to:

- *More members/customers.* IXPs can attract members which are present in different cities or countries, and thus, increase their market share. IXPs with many members are more visible and appealing to potential customers.
- *Reseller ecosystem.* The IXP can benefit from reseller organizations, which handle new IXP memberships at scale, and therefore the setup and billing of new members is simplified.

Remote peering is bad! On the other hand, some network and IXP operators claim that remote peering is a disservice to the Internet [23, 67]. IXPs have been originally created as peering hubs to keep “*local traffic local*” [39]. Changing this trend might lead to:

¹The involved IXPs still use their own route servers and BGP communities and serve their own member base.

²We consider as metropolitan area a disk with diameter 100 km.

³While IXPs such as DE-CIX may have presence in multiple cities (e.g., Frankfurt, New York), they are not considered as wide-area IXPs, since they operate an independent/separate IXP at each city. In contrast, NL-IX is a sole IXP entity with a network distributed among multiple countries/cities.

- *Degradation of performance.* Links over IXPs involving peers at distant locations from IXPs are expected to have larger latency (RTTs) than links between local peers. Hence, direct peering connections on IXPs might not necessarily lead to improved quality in communication. Additionally, resellers usually offer low capacity IXP ports (e.g., 100Mbps; see Section 5.1.1), which can cause congestion [43].
- *Loss of resilience.* While a network might have separate L3 connections with its peers on an IXP, in the case of remote peering some of these connections might share a common port (e.g., resellers sell fractions of the same physical IXP port to multiple remote peers). A single outage on this port can thus affect (a) multiple connections, and (b) networks hundreds or thousands of kilometers away from the IXP. As a result, neither traffic nor outages “stay local”.

Need for transparency. While there is no consensus on whether remote peering is a good or bad practice, both its proponents and opponents acknowledge the necessity for understanding the characteristics of remote peering. Network operators want to know which peers are local or remote, where they are located, and the implications on the communication (e.g., latency, bandwidth, resilience) among peers. This knowledge is critical since it can guide traffic engineering and peering policies.

Related work. Prior works on IXPs explore various aspects of the IXP ecosystem and show its impact on the Internet’s hierarchical topology [25, 29], traffic exchange economics [40, 57], and content delivery [30, 39, 73]. Others discuss multilateral peering over IXPs at scale [49] and show that interconnection strategies, such as RP, and extensive colocation practices [48], create unexpected interdependencies among peering infrastructures [46]. Other work investigates the impact of RP [50] on the topology or the performance of continental peering ecosystems, such as Africa [43]. Castro *et al.* [36] aimed to explore the traffic offloading capabilities of RP and provided a simple RTT-based approach for inferring RP.

However, in our work we show that RTT alone [36] is not sufficient to achieve accurate inference (see Section 4). Instead, we combine RTT measurements with several other domain-specific design aspects of remote peering and achieve significantly larger accuracy and coverage levels, calculated using a substantial validation dataset. Our goal is to establish a general, thoroughly validated RP inference methodology and yield valuable insights on the global RP ecosystem.

3 DATASETS & MEASUREMENTS

3.1 Active Measurement Sources

We employ ping measurements to estimate the latency (RTT) between an IXP and its member ASes, and traceroute measurements to extract the IP-level paths traversing peering links.

Pings. We conduct ping measurements from a number of Vantage Points (VPs), namely *Looking Glasses (LGs)* and *RIPE Atlas probes (RA)*; the exact location of these VPs is known. Castro *et al.* [36] used the PCH LGs [15] that provided access to PCH border routers deployed in 22 IXPs. Unfortunately, PCH does not allow ping queries through their LGs anymore. Instead, using IXP websites, we compiled a list of 23 publicly accessible LGs, that provide

Table 1: Overview of the IXP (IPv4) dataset and contribution of each data source.

Source	IXP Prefixes			IXP Interfaces		
	Total	Unique	Conflicts	Total	Unique	Conflicts
Websites	42	4		12409	24	
HE	429	51	1 (.010 %)	29866	7659	80 (.27 %)
PDB	638	187	1 (.005 %)	22146	1162	62 (.28 %)
PCH	467	129	1 (.007 %)	5922	256	22 (.37 %)
Total	731			31690		

direct interfaces inside the IXP networks, e.g., to an IXP route server. To automate the querying of these LGs we use the Periscope platform [45].

We augment the set of the ping-enabled VPs through RA [19], a well-established global Internet measurement platform with more than 25,000 probes. To identify RA probes colocated with IXP infrastructure, we search for probes with source IPs in the address space of an IXP’s peering LAN, and for probes which resolve to an ASN assigned to an IXP NOC⁴. We discovered 66 such RA probes.

Merging the available LG and RA VPs provides good coverage in the RIPE (29 IXPs) and APNIC (11 IXPs) regions. Only 6 IXPs are covered in the ARIN and LACNIC regions, and none under AFRINIC.

Traceroutes. We collect all the publicly available RA IPv4 traceroute measurements (*i.e.*, built-in and user-defined) [19]. In total, we study 3.15 billion traceroute paths towards 600K IPs, probed between Jan. 2017 and Mar. 2018. We use the collected traceroute paths to extract IP-level IXP crossings (see Section 3.3 and steps 3, 4 of Section 5.2), as well as private connections between ASes over facilities (see step 5 of Section 5.2).

3.2 IXP Peering LANs and Ports

Our methodology combines multiple sources of IXP-related information with the measurements of Section 3.1.

IXPs, members, and interfaces. To identify traceroute hops that traverse IXPs, and feed our methodology with IXP-related information, we combine multiple sources to build an up-to-date list of IXPs, their *members*, and the *associated IXP interfaces* (*i.e.*, IP addresses belonging to IXP prefixes that are assigned to IXP member ASes). We retrieve the related IXP information directly from IXP websites by parsing the provided Euro-IX [52] json and/or csv machine-readable formats, and the publicly available databases of Hurricane Electric (HE) [11], PeeringDB (PDB) [16], and Packet Clearing House (PCH) [15].

To address cases of conflicting data, we consider IXP websites as the most reliable source of information since the data are directly provided by the IXP operators; in fact, while websites may share peering policy information with e.g., PeeringDB, they maintain their own IXP-related information, such as membership lists. We then rank the other IXP sources based on their fraction of conflicting entries compared to the website data (Table 1). Consequently, we apply the following preference ordering to resolve conflicts: *IXP websites* > *HE* > *PDB* > *PCH*.

⁴ Note that probes connected to the IXP members themselves are not useful for our methodology, since these members can be also remote to the IXP, and thus may affect the RTT-based inference step biasing the ping measurements.

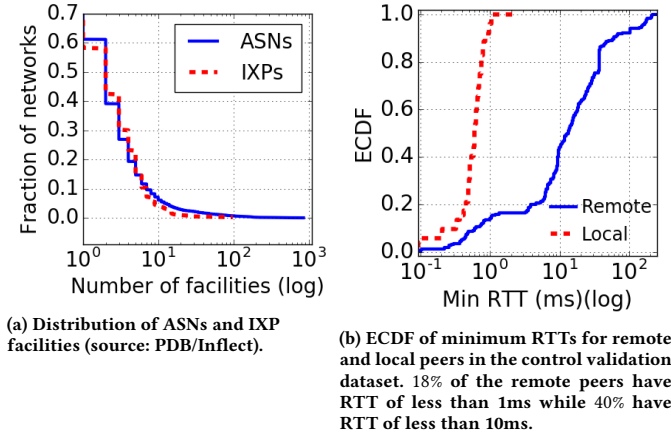


Figure 1: Overview of facilities and VP-to-IXP interface RTT.

The final dataset includes 31,690 IXP IP-to-AS mappings (*IXP interfaces*) and 729 IXP prefixes from 703 IXPs (Table 1). Interestingly enough, the IXP prefixes and interfaces that are unique in the websites are quite few (4 and 24 respectively), since the other databases are usually populated with up-to-date entries. To the best of our knowledge, the collected dataset comprises the most complete list of IXPs, IXP prefixes, and IXP interfaces to-date.

IXP port capacity. We record the capacity of the peering ports allocated to each IXP member, using the json/csv datasets directly provided through the IXP websites, and the PDB records. For each IXP, we also compile the available port capacity options through the pricing section of its website [70]. As we explain in Section 5.1.1, knowing the port capacities allows us to distinguish IXP peers that obtain virtual ports through port resellers from peers that obtain physical ports directly from the IXP.

3.3 Detecting IXP Crossings in Traceroutes

We process traceroute measurements (Section 3.1) and IXP information (Section 3.2) with traIXroute [21, 65] to identify paths that cross IXPs. We configure traIXroute to identify IXP crossings in a path, when (i) there exists a sub-path of three IPs (*i.e.*, *IP triplet*) that contains an IXP IP in the middle of the triplet and this IXP IP belongs to the same AS as the 3rd IP, (ii) the AS of the 1st IP in the triplet is different, and (iii) these two ASes are members of the IXP (whose prefix the IXP IP of the triplet belongs to).

3.4 Colocation Facilities

To infer the remoteness or locality of peers, we also use the location of the facilities where IXPs and their members are present. We first collect the facility list from PDB and *Inflect* [12], a database for Internet infrastructure services (whose data comes either directly from service providers or trusted third-party sources). For each facility we keep the geographical coordinates provided by PDB, which are independently verified through *Inflect* to filter-out spurious information [64]. Our dataset includes 656 IXPs which are associated with 1,078 facilities. The *Inflect* dataset allows us to correct the geographical information for 308 of these facilities. Moreover,

Table 2: Validation data retrieved from IXP operators (top, 6 rows) and websites (bottom, 9 rows). IXPs with superscript ‘C’ (‘T’) are part of the “control” (“test”) subset.

	IXP	#Facilities	#Total Peers	#Validated Peers	#Local	#Remote
OPERATORS	AMS-IX ^T	14	878	463	258	205
	DE-CIX FRA ^T	28	795	323	103	220
	LINX LON ^T	15	770	170	71	99
	DE-CIX NYC ^C	25	162	80	59	21
	LINX MAN ^T	3	99	37	17	20
	LINX NoV ^T	4	48	21	12	9
WEBSITES	EPIX KAT ^C	3	465	233	135	98
	EPIX WAR ^C	6	308	170	93	77
	France-IX PAR ^T	9	402	292	127	165
	Seattle IX ^T	11	296	246	180	66
	Any2 LA ^T	2	299	212	147	65
	D. Realty ATL ^C	3	142	85	42	43
	France IX MRS ^C	2	77	31	19	12
	AMS-IX HK ^C	2	46	24	14	10
	AMS-IX SF ^C	4	36	23	16	7
Total	131	4823	2410	1293	1117	

we extract information related to which facility each AS (*i.e.*, IXP member) is present. As shown in Fig. 1a, around 60% of IXPs and ASes are present in a single facility, with only 5% in more than 10 facilities. To alleviate possible incompleteness in PDB/Inflect data, we extend the colocation dataset by *manually* extracting the facility list from the websites of the 50 IXPs with most AS members. IXP websites provide additional facility data for 48% of the IXPs, allowing us to compile an as complete as possible dataset for the most prominent IXPs.

PDB vs. Websites. We have encountered some discrepancies between PDB and IXP/facility websites. For example, the NL-IX website provides additional information on 17 (~15%) of its data centers not present in PDB (incompleteness). On the other hand, for the CoreSite LA1 facility, PDB reports 108 ASes (~43%) that are not listed in Coresite’s list of locally deployed networks [6], indicating possible inaccuracies in PDB. Even in the face of such artifacts, the combination of the heuristics we apply in Section 5 results to high accuracy/coverage.

3.5 IXP Local/Remote Members for Validation

Inferring remote peering accurately, requires thorough investigation of the challenges related to interconnectivity between IXPs and their members, as well as information to validate the peering inference itself. To this end, we contacted IXP operators and requested lists specifying which of their members are local and/or remote. We received validation data⁵ for 6 IXPs. However, the provided lists do not cover the entire list of the members of these IXPs. This is due to the fact that IXP operators usually know whether their members are connected through resellers, but not where they are located, or if they use a L2 carrier to access their colocation facilities. In essence, they do not/cannot know “what goes on beyond that cable” [23]; a gap that is the primary motivation of this work.

We further augmented the validation dataset by manually extracting lists of remote and local members from websites of IXPs that publish the port type of their members (physical or virtual through a reseller). In total, we collected validation data for 6 IXPs directly from their operators, and for 9 more IXPs from their websites. In addition, we enriched the total IXP list in the validation

⁵The validation dataset we use is a best-effort collection of relevant trusted data.

dataset with the facilities at which the IXPs are present based on data from Section 3.4. All relevant statistics are shown in Table 2.

We split the validation dataset in two subsets, “*control*” and “*test*”, depending on whether they include IXPs with publicly accessible colocated VPs from which ping measurements can be triggered. The reason for this discrimination is that we need to (i) re-evaluate existing inference approaches [36] and investigate further challenges in order to fine-tune our approach, and (ii) properly validate the full methodology using active measurements. Since only the *test* subset contains IXPs with accessible ping-enabled VPs, we used the *control* subset to evaluate latency-wise inference challenges (see Section 4), and the *test* subset to ping local and/or remote target interfaces in order to compare our inference results with the reported ones (see Section 5.3).

4 RTT-BASED INFERENCE CHALLENGES

Here, we use the *control* subset of our validation dataset to investigate the challenges and limitations of inferring RP based exclusively on latency measurements (Section 4.1), placing emphasis on the fairly common case of wide-area IXPs (Section 4.2).

4.1 RTT is not enough

For each IXP in our control dataset, there is no publicly available VP to execute RTT measurements, but we obtained one-time access to results from pings executed within the IXP infrastructure targeting the peering interfaces of all the remote and local members of the IXP. We apply the *TTL match* and *TTL switch* filters proposed in [36] to discard replies with TTL values less than the expected maximum (64 and 255 hops) that may indicate ping replies outside of the IXP subnet. We repeat the measurements every 20 minutes for two days, and we calculate the minimum RTT per IXP interface. As shown in Fig. 1b, RTT values above $2ms$ are a very strong indication of remote peers, with 99% of the local peers having RTT values less than $1ms$. This result is consistent with previous works that exhibited that a delay of $1ms$ corresponds roughly to a distance of 100 km [54, 75], approximating the coverage (i.e., disk diameter) of a single metropolitan area. However, low RTT does not necessarily mean that a peer is local. Surprisingly, **18% of the remote peers in our control dataset are within $1ms$ from the IXP**, while 40% are within $10ms$, which is the “*remoteness threshold*” used in [36].

4.2 Wide-area IXP challenges

Conservative latency thresholds do not ensure the elimination of peers which are falsely identified as remote for *wide-area IXPs*. In fact, IXP members which are present in any of the facilities of such IXPs are local to the IXP but can be remote to the measurement VP, even if the VP is also hosted in one of the IXP’s facilities. An indicative example is NET-IX, which has distributed its switching fabric in facilities across 18 different countries [24]. To understand the RTT characteristics among the different facilities of such a geographically distributed IXP, we obtained pairwise delay measurements between 16 of NET-IX’s international sites. NET-IX measures the delay between its different facilities based on the Y.1731 Performance Monitoring standard [22], by sending precisely timestamped test packets across its MetroNID network demarcation points. The results are shown in Fig. 2a. For 87% of the facility pairs the median

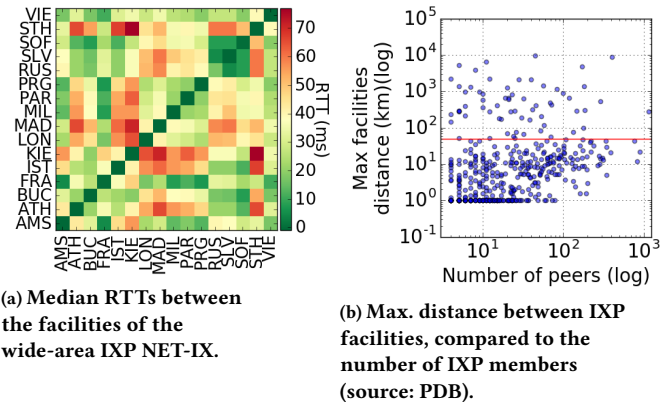


Figure 2: Features of wide-area IXPs.

RTT is above $10ms$. Note that we also observe facilities in different countries with less than $10ms$ delay between them; for instance, Frankfurt (FRA) and Prague (PRA) have a $7ms$ delay. **Therefore, a remoteness RTT-threshold is not meaningful for wide-area IXPs.**

Next, we quantify the popularity of the model of the wide-area IXPs. We use our colocation dataset compiled in Section 3.4, and we classify an IXP as wide-area if its switching fabric is deployed among multiple facilities, and at least two of them are in different metropolitan areas. Since there can be different naming conventions used for the same city/metro area, we calculate the geodesic distance between each pair of IXP facilities, by applying Karney’s method [53] on their geographical coordinates. We consider facilities more than $50km$ apart as located in different metropolitan areas. For April 2018, we found that 64 of the 446 (14.4%) IXPs in PDB with at least two IXP members are wide-area, including 10 of the 50 (20%) largest IXPs in terms of the size of their IXP member list (Fig. 2b). **Therefore, wide-area IXPs are fairly common and not just some exceptional cases.** Note that the infrastructure of some IXPs can be thousands of *kms* apart. For instance, NL-IX has facilities in London and Bucharest that are over $1,300km$ away from each other.

The results of this section highlight that although RTT measurements have the potential to provide useful insights w.r.t. the peering approach employed by an IXP member, alone they are not adequate to accurately infer remote peers. A $10ms$ -threshold is very conservative in the case of IXPs concentrated in a single metropolitan area, while it yields a large number of false positives in the case of wide-area IXPs.

5 INFERENCE METHODOLOGY

To address the limitations of remote peering inference based exclusively on latency measurements, we introduce a “first-principles” [56] approach. We rely on domain-specific knowledge to identify technological (beyond latency) and economic aspects of peering connectivity (Section 5.1), and build upon these aspects to design a methodology for inferring remote and local peers (Section 5.2). We validate the proposed methodology in Section 5.3.

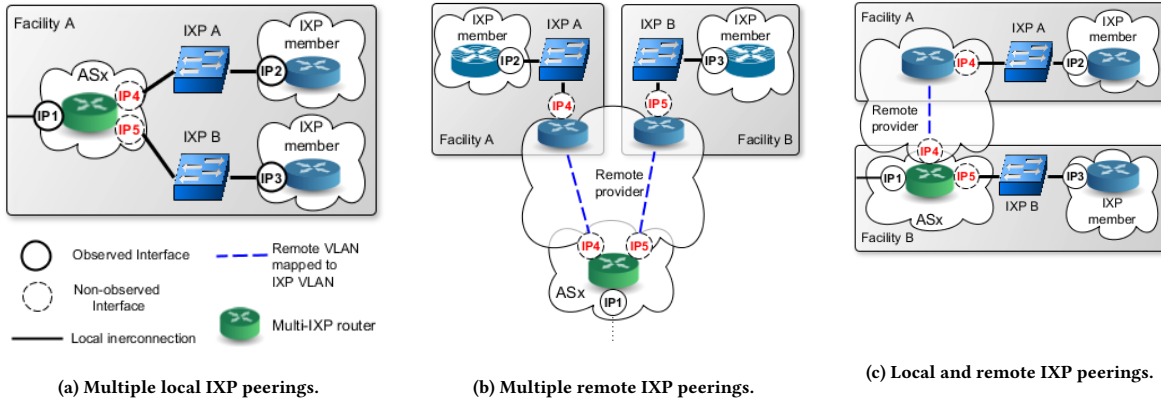


Figure 3: Different scenarios of *multi-IXP* routers, for which we may observe different traceroute paths where IP_1 precedes both IXP interfaces IP_2 and IP_3 , indicating the presence of a multi-IXP router in AS_x .

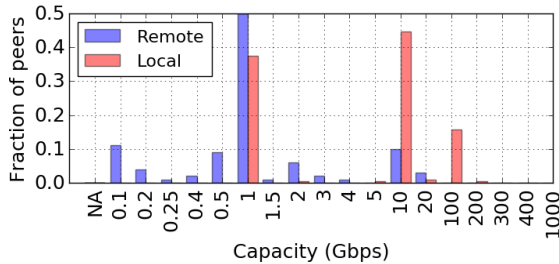


Figure 4: Capacity of IXP ports for remote and local peers in our control validation dataset. Fast Ethernet (FE) carries traffic at the rate of 100 Mbit/s and Gigabit Ethernet (GE) at 1 Gbit/s.

5.1 Design Aspects

5.1.1 Port Capacity. IXPs offer to ASes connectivity to switch ports, whose capacity is typically between 1GE and 100GE [2]. To make remote peering an attractive service, resellers split their physical ports to multiple virtual ports (e.g., via sub-interfaces/VLANs) of lower capacity (rate-limiting), and offer them to remote peers at lower prices. *Fractional port capacities can be purchased only through resellers today*⁶. Thus, this information can indicate a network that peers remotely, via a reseller, at an IXP. Figure 4 shows the port capacity for remote and local peers in our control validation dataset. No local peer has port capacities below 1GE (which is the minimum capacity for physical ports offered by the corresponding IXPs), while 27% of remote peers access the IXP through ports of 1FE – 5FE capacity; on the other hand, ports of 100+GE are allocated only to local peers.

5.1.2 Presence at Colocation Facilities. To establish a direct connection to an IXP, an AS needs to deploy routing equipment in at least one colocation facility where the IXP has deployed switching equipment. Therefore, *it is not possible for an AS to be a local peer of an IXP if they are not colocated in a facility*. As Fig. 5 shows, all local peers of an IXP in our control validation dataset are present in at least one IXP facility, while 95% of the remote peers do not have any common facility with the IXP. Hence, assuming perfect knowledge

⁶In rare cases, some old IXP members are connected to physical ports of capacity less than the minimum offered today. This can be also due to stale entries in PDB.

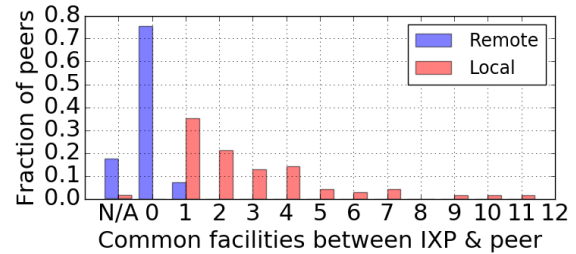


Figure 5: Number of IXP facilities where local and remote peers in our control validation dataset are present.

of the facilities where IXP members are present, identifying RP would be a straightforward lookup process. However, the available colocation data for IXP members are incomplete and noisy. For example, in Fig. 5, there are no available data for 18% of the remote peers, while 5% of them appear to have presence in one IXP facility.

To further investigate the latter 5% of RP cases, we contacted the IXP operators. Their feedback suggested that such cases are either an artifact of remote peers (not colocated with the IXP) adding the facility of their port reseller in their PDB record, or a consequence of the fact that peers (colocated with the IXP) prefer to connect through a port reseller in order to buy virtual ports of lower capacity at a discount price (see Section 5.1.1).

5.1.3 Multi-IXP Routers. An AS may connect to multiple IXPs through the same border router to reduce operational costs; we call such routers *multi-IXP routers*. The IP interfaces of a multi-IXP router might appear in different traceroute paths to be interconnected with different IXPs. We distinguish three cases where this is possible:

- (1) When multiple IXPs are present in the same facility, a colocated AS may connect directly to all of them using a single router (Figure 3a).
- (2) Remote peers may connect through the same provider (port reseller) to multiple remote IXPs where this provider has presence (Figure 3b).
- (3) An AS may connect with the same router to both local and remote IXPs, if it is e.g., colocated with one IXP and uses a reseller for another (Figure 3c).

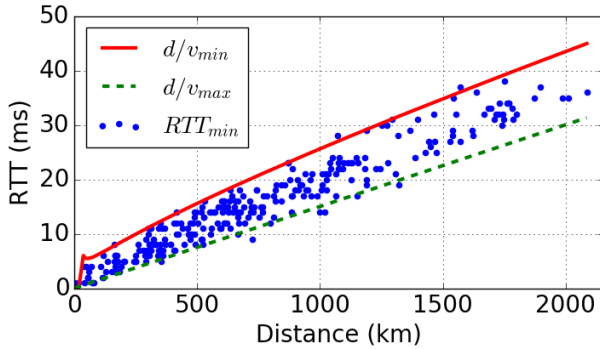


Figure 6: Inter-facility RTT as a function of distance, based on Y.1731 Perf. Monitoring measurements from NL-IX and NET-IX.

5.1.4 *Private Connectivity.* Two networks collocated at the same IXP-hosting facility can interconnect with each other (*private peering*) without using the IXP infrastructure, e.g., by directly connecting their routers. This might be a more economical solution in case they exchange large volumes of traffic [66]. Therefore, when an IXP member appears to be privately connected with several ASes which are collocated at the facility of the same IXP, this is a strong indication that this member is local to the IXP.

5.2 Algorithm

We next describe our methodology for inferring remote peering, by combining RTT measurements with the four peering aspects discussed in Section 5.1. While the steps of the methodology can be validated independently (see Section 5.3), the order in which they are applied matters and was selected as follows. Step 1 (Port Capacities) is first since it reliably infers RP, albeit with small coverage. Step 2 (RTT measurement) generates data used for step 3. Step 3 (RTT+colocation) is required as input by Step 4 (multi-IXP routers) and 5 (private connectivity). Step 4 comes before step 5 due to its higher accuracy; step 5 is the last resort for missing inferences. Note also that while an individual step may miss some cases for different reasons (e.g., incomplete collocation data or RTT outliers in Step 2), these cases can be captured by a following step.

Step 1: Finding reseller customers via port capacities. IXP members that reach the IXP through a reseller are identified as remote peers (see Definition 1). As discussed in Section 5.1.1, members can be connected to IXP ports of capacity lower than the minimum physical port capacity C_{min} offered by the IXP, only if they reach the IXP through a reseller⁶. Hence, as a first step, for each IXP member AS_x we compare the port capacity C_x , reported either in the IXP website or the Inflect and PDB databases, to the C_{min} value reported in the pricing section of the IXP’s website. If $C_x < C_{min}$, we infer that AS_x is a remote peer using a virtual port obtained through a reseller.

Step 2: Ping RTT Measurements. From every VP in an IXP (see Section 3.1), we execute ping measurements to every IXP IP interface of the IXP’s members (see Section 3.2). To reduce the sensitivity of the results to network conditions, we repeat the measurements every two hours for two days, which results in 24 measurements in total for each {VP, IP interface} pair. Similarly to Section 4, we apply

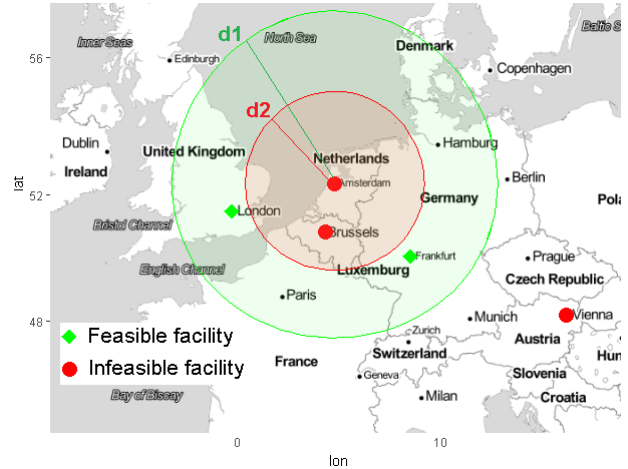


Figure 7: Example of combining RTT measurements with IXP collocation data to infer local peers at geographically distributed IXPs.

the *TTL match* and *TTL switch* filters to discard measurements without consistent TTL values. Finally, for each responsive IP interface we store the minimum RTT value, RTT_{min} , to counter transient latency inflation artifacts [51].

Step 3: Colocation-informed RTT interpretation. To infer local and remote peers, we analyze the collected RTT_{min} values. Besides the collocation information of the IXPs and its members (see Section 3.4), the exact locations of the VPs are also known in all ping measurements. From the value of the RTT_{min} we calculate a geographical area (circle or ring) around the VP location where the IP interface (and thus the router) of the IXP member can be located. The presence (or not) of a facility of the IXP in this area, denotes a local (or remote) peering, respectively.

More precisely, we first calculate the distance between the involved VPs and each of the IXP’s facilities, as described in Section 4. Then, from the observed RTT_{min} , we calculate the potential distance between the VP and the ping target (IP interface at a member’s router). Katz-Bassett *et. al* [54] found that the end-to-end probe packet speed is at most $v_{max} = \frac{4}{9} \times c$, where c is the speed of light. As shown in Fig. 6 (green/dashed curve), our dataset of facility-to-facility delays based on Y.1731 measurements obtained from NL-IX and NET-IX confirms this. Through data fitting, we also find an approximate lower bound (red/continuous curve in Fig. 6) for the speed $v_{min}(d) = 10^7 \cdot (\ln(d) - 3)$, where d is the distance. Based on these bounds⁷, we estimate that the ping target is within a distance range $D_{feasible} = [d_{min}, d_{max}]$ (green area in Fig. 7) from the VP, where $d_{min} = v_{min} \times RTT_{min}$ and $d_{max} = v_{max} \times RTT_{min}$. We call the facility that is located in $D_{feasible}$, a *feasible facility*.

Based on the estimated area defined by $D_{feasible}$ (see e.g., Fig. 7), and the distances between the IXP facilities and the VP, we infer that the IXP member that owns the queried IP interface (ping target) is local or remote to the IXP, as follows:

- (1) **Remote peer:** if (i) the IXP has no available feasible facility, or (ii) the IXP has at least one feasible facility, but the peer is present in another feasible facility where the IXP is *not* present.

⁷ Out-of-bounds outliers do not impact the high accuracy of this step (see Table 4).

- (2) **Local peer**: if the IXP has at least one feasible facility, and the IXP member is also colocated in one of the feasible IXP facilities.
- (3) **No inference**: if the IXP has at least one feasible facility, but the IXP member is *not* present at any feasible facility.

In the latter case, it is likely that our colocation dataset is incomplete w.r.t. the given peer. In this case we do not make an inference yet, but instead we leverage *multi-IXP router* and *private connectivity* information (see Section 5.1) as described in the following steps.

Combining RTT values with colocation information allows us to alleviate false positives caused by wide-area IXPs. Figure 7 shows an example of such a case, based on the topology of the NL-IX IXP. The IXP has distributed its peering fabric across multiple cities, including Amsterdam, Brussels, London, Frankfurt and Vienna. Our measurement VP is in an IXP facility in Amsterdam, from which we ping the IXP peering interfaces. Assume that for an interface IP_x we measure an RTT_{min} of 4ms. Without taking into consideration the geographical footprint of the IXP’s infrastructure we would infer the corresponding peer as remote assuming a “reasonable” (see Fig. 1b) 2ms-threshold. Instead, we find that the IXP has two feasible facilities (London and Frankfurt) in the ring between $d_1 = 532km$ and $d_2 = 299km$ from the VP, as defined by our v_{max} (green area) and v_{min} (red area) bounds respectively, allowing us to infer as local the IXP members colocated at these facilities.

Similarly, we can avoid false negatives due to remote peers that are in close proximity to the IXP. For instance, for a peer located in Rotterdam connected remotely to the IXP’s facility in Amsterdam (57km distance) we will typically measure $RTT_{min} < 2ms$. By using the peer’s colocation data we can correctly determine that, despite the low RTT, the peer is not local.

Step 4: Multi-IXP router inference. The previous steps may not be able to infer the peering type due to missing facility data or missing RTT values from unresponsive IXP interfaces. In such cases, we proceed to use the multi-IXP router feature (see Section 5.1.3), for inferring remoteness (or locality).

To identify multi-IXP routers we first collect traceroute paths from public RIPE Atlas measurements in the same period as our ping campaign (two days). We then extract the IP-level IXP crossings, as explained in Section 3.3, and we collect all sequences of hops $\{IP_x, IP_{x+1}^{IXP}\}$, where the interface IP_{x+1}^{IXP} belongs to the address space of an IXP, and the interface IP_x belongs to an AS that is a member of this IXP. For each AS that appears to peer at more than one IXP in different IXP crossings, we perform alias resolution on all its IP interfaces using MIDAR [55] to map these interfaces to routers⁸. For interfaces on the same router, we find the set of IXPs that appear as next hops in traceroute paths. If a router appears to have connections to more than one IXPs, we characterize it as a multi-IXP router.

For example, assume two sequences of IP hops, $\{IP_a, IP_{IXP1}\}$ and $\{IP_b, IP_{IXP2}\}$, where both IP_a and IP_b are owned by the same AS and are mapped to the same router R , and IP_{IXP1} and IP_{IXP2} belong to the peering LANs of $IXP1$ and $IXP2$, respectively. In this

case, R has layer-3 connectivity with both IXPs, and therefore we characterize R as a multi-IXP router.

We then classify the multi-IXP routers in one of the categories described in Fig. 3, and infer each one based on geolocation data from Section 3.4 as follows:

- (1) **Local multi-IXP router**: A multi-IXP router is local to all involved IXPs (Fig. 3a), if (i) the involved AS has been inferred as local peer –from previous steps– in at least one of the IXPs, and (ii) the involved IXPs have at least one common facility. Then *the AS is inferred as a local peer to all involved IXPs.*
- (2) **Remote multi-IXP router**: A multi-IXP router is remote to all involved IXPs (Fig. 3b), if (i) the involved AS has been inferred as remote peer –from previous steps– in at least one of the IXPs (e.g., IXP_R), and (ii) at least one of the following holds:
 - (a) all the involved IXPs have at least one common facility.
 - (b) the maximum distance between the facilities of any involved IXP and IXP_R , is smaller than the minimum possible distance d_{min} between all the facilities of the involved AS and all the facilities where IXP_R is present.

Then *the AS is inferred as a remote peer to all involved IXPs.*

- (3) **Hybrid multi-IXP router**: A multi-IXP router is local to a subset of the involved IXPs (Fig. 3c) and remote to another IXP subset, if (i) the involved AS has been inferred as local peer –from previous steps– in at least one of the IXPs (e.g., IXP_L) of the local subset, and (ii) at least one of the following conditions is true for the remote subset:
 - (a) IXP_L does not have any common facility with the other involved IXPs.
 - (b) the minimum distance between the facilities of IXP_L and any other involved IXP, is larger than the maximum possible distance d_{max} between all the –common– facilities where both the involved AS and IXP_L are present.

Then *the AS is inferred as a local peer to IXP_L and remote peer to all other involved IXPs in the remote subset.*

To understand the intuition behind conditions 2(b) and 3(b), assume that $R_x \in AS_x$ is a multi-IXP router peering with two IXPs, IXP_{ams} in Amsterdam, and IXP_{lon} in London. The minimum distance between the facilities of the two IXPs is 300km, while the maximum distance is 360km. If from the first two steps we inferred that AS_x is remote to IXP_{ams} , with $d_{min} = 500km$, then R_x cannot be local to any facility of IXP_{lon} (condition 2(b) holds). Similarly, if we inferred that AS_x is local to IXP_{ams} with $d_{max} = 50km$, then R_x cannot be local to any facility of IXP_{lon} (condition 3(b) holds).

Step 5: Localization of private connectivity. If Steps 1-4 fail to infer whether a peer is local or remote, we use the private connectivity of an IXP member and apply a “voting” scheme similar to the Constrained Facility Search (CFS) approach [48].

Let \mathcal{F}_{IXP} be the set of feasible facilities for the IXP, AS_x an IXP member identified based on the dataset of Section 3.2, and \mathcal{I}_{IXP} the set of all IP interfaces of the multi-IXP routers identified in Step 4.

- (1) We parse all the collected traceroute paths, perform IP-to-AS mapping [34] and extract all the AS sequences over *private interconnections* (not over an IXP), i.e., from a sequence $\{IP_i, IP_j\}$, where IP_i belongs to AS_i and IP_j to AS_j ($\neq AS_i$), we extract the sequence $\{AS_i, AS_j\}$. Let \mathcal{I}_{priv} be the set of all interfaces involved in such private AS-level interconnections.

⁸ There are two available datasets based on MIDAR: (i) one based on aliases resolved with MIDAR and `iffinder` [32], yielding the highest confidence aliases with very low false positives, and (ii) one also including aliases resolved with `kapar` [33], which significantly increases coverage at the cost of accuracy. We selected the first dataset to favor accuracy over completeness.

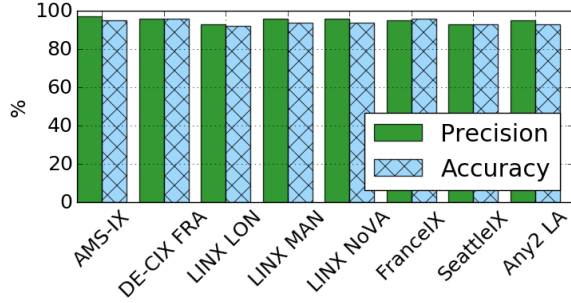


Figure 8: Validation results per IXP in our test validation dataset.

- (2) We run alias resolution on the interfaces in $I_{IXP} \cup I_{priv}$, that belong to IXP members for which we have not made an inference yet. For each router R_x (belonging to an AS_x) with at least one interface $i \in I_{IXP}$, we compile the set N_x of the (private) AS neighbors of AS_x .
- (3) Based on our AS-to-facility mapping from Section 3.4, we find the most common facilities \mathcal{F}_{common} among the majority of the ASes in N_x .

If $|\mathcal{F}_{IXP} \cap \mathcal{F}_{common}| = 1$, i.e., only one facility of the IXP belongs to both sets, then we infer AS_x as a local peer to the IXP. Otherwise we infer the peer as remote to the IXP. The intuition behind this heuristic is that private interconnections are typically established within the same facility, as explained in 5.1.4. Nonetheless, we do not require all the private AS neighbors to be present in \mathcal{F}_{common} because tethered private interconnections across facilities –although less common– are still possible [48].

It should be noted that our aim is not to pinpoint the exact AS boundaries, nor to derive the AS-level topology from IP hops, both of which have been shown to be non-trivial processes [59, 61]. Instead, we aim to infer a router’s colocation facility based on its adjacent ASes. For example, a reply from a third-party interface may result in a spurious AS-level link; however, the interface (no matter to which AS it is mapped) belongs to the same router, and thus the facility inference is not affected.

Table 3: Validation Sets and Metrics for RP Inference.

	Name	Definition
Sets	$\mathcal{V}\mathcal{D}_R$	Remote Peers in Validation Dataset
	$\mathcal{V}\mathcal{D}_L$	Local Peers in Validation Dataset
	$\mathcal{V}\mathcal{D}$	$\mathcal{V}\mathcal{D} = \mathcal{V}\mathcal{D}_R \cup \mathcal{V}\mathcal{D}_L$
	$\mathcal{I}\mathcal{N}\mathcal{F}_R$	Inferred Remote Peers
	$\mathcal{I}\mathcal{N}\mathcal{F}_L$	Inferred Local Peers
	$\mathcal{I}\mathcal{N}\mathcal{F}$	$\mathcal{I}\mathcal{N}\mathcal{F} = \mathcal{I}\mathcal{N}\mathcal{F}_R \cup \mathcal{I}\mathcal{N}\mathcal{F}_L$
Metrics	COV	$\frac{ \mathcal{I}\mathcal{N}\mathcal{F} \cap \mathcal{V}\mathcal{D} }{ \mathcal{V}\mathcal{D} }$ (Coverage)
	FPR	$\frac{ \mathcal{I}\mathcal{N}\mathcal{F}_R \cap \mathcal{V}\mathcal{D}_L }{ \mathcal{I}\mathcal{N}\mathcal{F} \cap \mathcal{V}\mathcal{D}_L }$ (False Positives rate)
	FNR	$\frac{ \mathcal{I}\mathcal{N}\mathcal{F}_L \cap \mathcal{V}\mathcal{D}_R }{ \mathcal{I}\mathcal{N}\mathcal{F} \cap \mathcal{V}\mathcal{D}_R }$ (False Negatives rate)
	PRE	$\frac{ \mathcal{I}\mathcal{N}\mathcal{F}_R \cap \mathcal{V}\mathcal{D}_R }{ \mathcal{I}\mathcal{N}\mathcal{F}_R }$ (Precision)
	ACC	$\frac{ \mathcal{I}\mathcal{N}\mathcal{F}_R \cap \mathcal{V}\mathcal{D}_R + \mathcal{I}\mathcal{N}\mathcal{F}_L \cap \mathcal{V}\mathcal{D}_L }{ \mathcal{I}\mathcal{N}\mathcal{F} }$ (Accuracy)

Table 4: Validation of each step of the algorithm.

Methodology	Feature	FPR	FNR	PRE	ACC	COV
Steps	RTT_{min} [36]	17.5%	25.7%	85%	77%	84%
Step 1:	Port Capacity	-	-	96%	-	11%
Step 2+3:	$RTT_{min} + \text{Colo}$	1.1%	7%	98.5%	95.6%	76%
Step 4:	Multi-IXP	7%	7%	93%	93%	53%
Step 5:	Private Links	10%	16%	90%	86.5%	49%
	Combined	4%	7.2%	95%	94.5%	93%

5.3 Validation

We validate each step of our methodology independently by comparing inference results (see Section 6.1) against the *test* subset of the validation dataset (see Section 3.5). The validation metrics we use and the sets that we consider are defined in Table 3. Note that concerning validation data it holds that $\mathcal{V}\mathcal{D}_R \cap \mathcal{V}\mathcal{D}_L = \emptyset$ (on the interface level), and in the metrics we do not take into account inferences for peers with no validation data (i.e., $\mathcal{I}\mathcal{N}\mathcal{F} - \mathcal{V}\mathcal{D} = \emptyset$). Table 4 shows the validation results for all IXPs in the test dataset, for each step separately, as well as the entire algorithm.

State-of-the-art. As a baseline, we first validate the remote inference when using only RTT_{min} (step 2), assuming a remoteness threshold of 10ms [36], to quantify the improvement versus the state of the art [36] achieved by our algorithm. RTT_{min} yields a high FPR due to mis-inferring local peers at wide-area IXPs as remote. We calculated that when excluding wide-area IXPs the FPR of the RTT_{min} approach drops to 2%. At the same time, the FNR is also high since many of the remote peers have $RTT_{min} < 10ms$.

Proposed methodology. When combining RTT_{min} with colocation data from Section 3.4 (step 3) we improve significantly all validation metrics; only the coverage metric has a small decrease, due to the fact that both latency and facility data are required. The false-negative inferences of $RTT_{min} + \text{Colo}$ are either due to spurious colocation data, or reseller customers collocated at the IXP. The latter false negatives are alleviated by taking into account *Port Capacity* data as described in Section 3.2 (step 1). For port capacity we validate only the precision metric, since we use it to infer only remote peers. For the next two steps we utilize traceroute data from Section 3.1. Specifically, the *Multi-IXP* step (step 4) also exhibits very high PRE and ACC, but can be used only for half of the interfaces. Finally, the *Private Links* step (step 5) has the lowest ACC and PRE compared to the other steps, but still outperforms vanilla RTT-based inference and is used only as a “last-resort” heuristic. When all the five steps are combined, they yield $\sim 95\%$ ACC and PRE, and cover 93% of the tested IXP interfaces. Fig. 8 shows the precision and accuracy metrics per IXP in our test validation dataset, ordered by the size of IXP. The results are consistent across all IXPs. For SeattleIX we obtain the lowest precision (92%), due to incomplete colocation data. Our inferences for LINX LON have the lowest accuracy (91%), because of a higher –than the other IXPs– number of collocated members connected through remote providers using non-fractional ports. These inaccuracies may indicate potential errors in the port capacities dataset.

6 INFERRING RP IN THE WILD

Here, we apply our inference methodology on the 30 largest IXPs in our dataset, step by step (Section 6.1). Having inferred RP at IXPs,

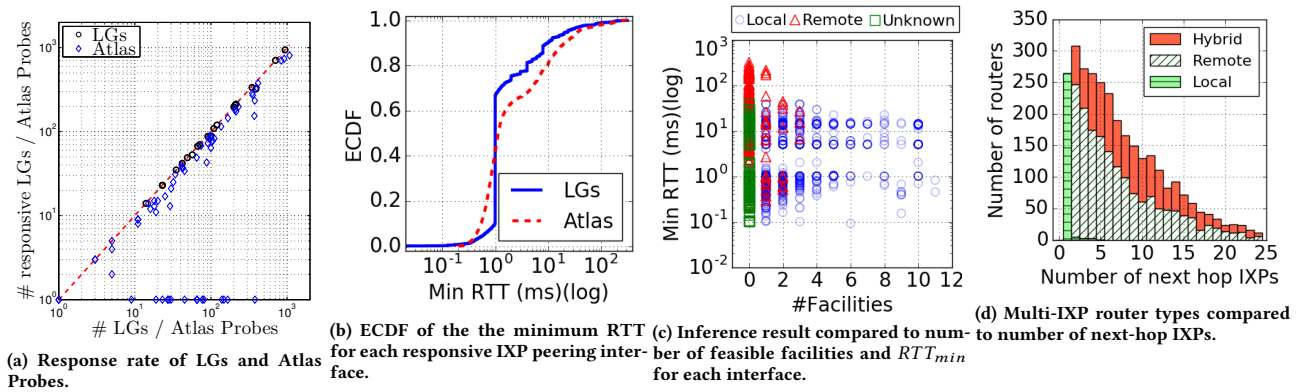


Figure 9: Measurement results for RTTs, feasible facilities and multi-IXP routers.

we investigate some relevant use cases. Indicatively, we focus on RP features in Section 6.2. We further study aspects of the evolution of the RP ecosystem over time (Section 6.3), as well as routing implications involving a large IXP (Section 6.4).

6.1 Application of Step-wise Inference

Step 1. We first infer the IXP members that reach the IXP through resellers, by comparing the port capacities of each member against the minimum physical port capacity offered by that IXP. For some IXPs (see Fig. 10a), such as France-IX, which cooperates with more than 20 resellers [10], 40% of the inferences can be made by using only port capacity information. However, for other IXPs that do not allow port reselling (e.g., HKIX), this step fails to make any inference. On average, this step contributes for approx. 10% of the total inferences (see column COV in Table 4).

Step 2. We then execute a ping measurement campaign between 7-9 Apr. 2018 from each LG and Atlas VP, to the peering interfaces of the IXP that hosts the VP. LGs achieve high response rates (Fig. 9a) due to being directly attached to the IXP peering LAN. In contrast, 50 of the 66 Atlas probes are colocated within an IXP facility, but are not inside the IXP’s LAN. Therefore, pings from them to IXP LAN IP addresses are more likely to fail for various reasons [62]. 14 of the Atlas probes do not receive any ping response.

Figure 9b shows the RTT_{min} distributions between VPs (LGs and Atlas probes) and IXP interfaces. 75% of the IXP interfaces are within 2ms from the respective VP. **More than 20% of the interfaces have $RTT_{min} > 10ms$, a 2-fold increase since 2014** [35, 36].

However, we found Atlas probes with consistently inflated RTT values⁹. Such probes may be deployed in the IXP’s management LAN which may not be in the IXP’s facilities, but still abide to the *TTL match* filter (see Section 4) which is set to $TTL_{max} - 1$ for Atlas probes. Thus, we discard probes that have $RTT_{min} \geq 1ms$ between the probe and the IXP’s route server. This filter removes another 21 Atlas probes from the set of usable VPs. Also, note that a large number of minimum RTTs obtained from LGs are exactly 1ms, which happens because many LGs round up the RTT value to the nearest integer. For such LGs we calculate the d_{min} distance between the IXP interface and the VP assuming $RTT'_{min} =$

⁹Atlas probes can yield measurement errors [51]; in our campaign, we account for non-persistent inflation by considering minimum RTTs over time.

$RTT_{min} - 1$, and we use the rounded-up RTT_{min} to calculate the corresponding d_{max} distance.

Table 5 provides the statistics of the queried interfaces that were used for our inferences after filtering out the unusable VPs.

Step 3. We calculate the feasible IXP and AS facilities for each peering interface, based on the measured RTT_{min} , and infer the interfaces as local, remote or unknown, based on the combined latency and colocation information. Figure 9c shows the RTT_{min} for each IXP interface versus the number of feasible facilities. Each $(RTT_{min}, \#facilities)$ data point is tagged with its inferred peering type. **94% of the remote interfaces have no feasible common facility with the IXP** (which further validates the colocation “principle”), while for 6% we have at least one feasible facility. Drilling down on this 6%, 40% of the involved interfaces exhibit $RTT_{min} > 2ms$, indicating spurious colocation information. Moreover, 5% of them are in a facility within the same metro area as the IXP VP but not affiliated with the IXP, while the rest are cases of IXP members colocated with the IXP but connecting through a reseller via a low-capacity virtual port (inferred at Step 1).

Step 4. For the *unknown* interfaces of Step 3, we investigate if they are part of multi-IXP routers. Figure 9d shows the number (per inferred type) of IXP routers compared to the number of IXPs with which they are connected (next-hop IXPs). Surprisingly, we find that 20% of the *unknown* interfaces and **~80% of the corresponding routers have multiple IXP connections, with 25% of them connecting to more than 10 IXPs**. This result highlights that the AS-level and IXP-level peering diversity of such IXP peers are misleading indicators of their resilience, since **all of their interconnections depend on the same physical equipment** (i.e., the multi-IXP router). We further observe that cases of remote multi-IXP routers are more prevalent than hybrid ones.

Table 5: Statistics of interfaces involved in the ping campaign. For our measurements we used the 30 largest IXPs with usable VPs.

VP Type	# VPs	# Interfaces		# Members	# IXPs
		Queried	Resp. (Fig. 9a)		
LG	23	3,806	3,617 (95%)	2,347	18
Atlas	22	6,457	4,861 (75%)	4,097	22
Total	45	10,578	7,738 (73%)	6,444	30

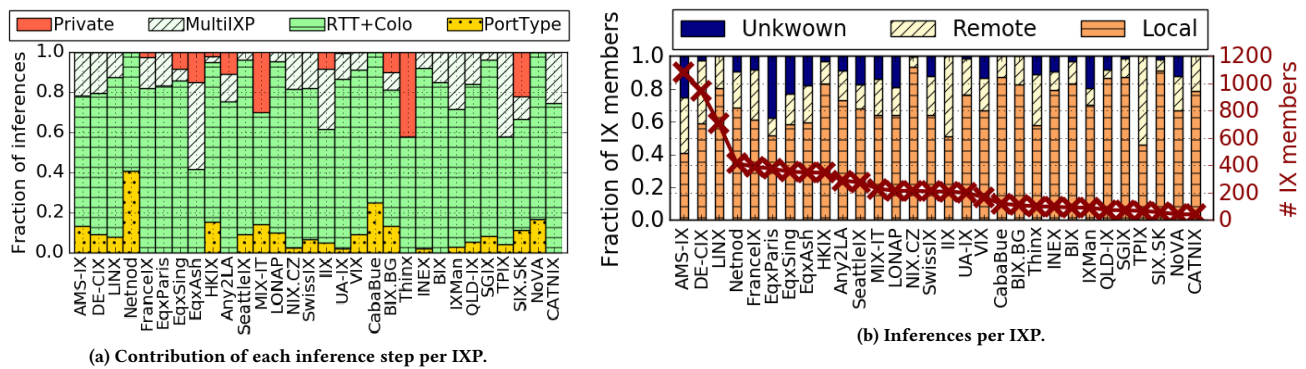


Figure 10: Inference results for the 30 largest IXPs with LG/Atlas VPs.

Step 5. Finally, for the remaining unknown interfaces we infer locality or remoteness based on their private connectivity. As shown in Fig. 10a, we had to apply this heuristic only for 11 of the top 30 IXPs, because previous steps did not manage to successfully infer some of the IXP interfaces of these IXPs as remote or local.

Overall. In total, the contribution (in terms of fraction of inferences) of each step of the methodology is shown in Fig. 10a. Steps 2 (*RTT + colo*) and 3 (multi-IXP routers) account for the majority of the inferences. Moreover, Fig. 10b shows the final inference results for the top 30 IXPs. Overall, **we find 28% of all the IXP interfaces for which we made an inference to be remote. Also, for 90% of the IXPs, it holds that more than 10% of their members are remote peers.** Finally, we find that for the two largest IXPs (DE-CIX and AMS-IX) almost 40% of their members are remote.

6.2 Features of Remote Peers

Having inferred remote and local peers per IXP, we proceed to investigate what are the features of remote peers and if/how they differentiate from local peers. We examine 3 features for each IXP member: (i) the size of its customer cone, as reported by CAIDA [5], (ii) its traffic levels and countries of presence as reported by PDB [16], and (iii) the user population it serves, as reported by APNIC [28]. We classify an IXP member network as follows: “remote” if it has only remote connections; “local” if it has only local connections; “hybrid” if it has both types (in the same or multiple IXPs). Out of 2959 total inferred AS-peers in 30 IXPs, we find that 63.7% are local, 23.4% are remote and 12.9% are hybrid.

In Fig. 11a we show the fractions of remote, local and hybrid IXP members with respect to the size of their customer cone. We observe that remote peers (red line) have quite similar patterns with the local ones (blue line). In fact, whether a network chooses to engage in local or remote peering (which is a matter of network design) at an IXP is not reflected on the size of its customer cone. This is probably due to the fact that both practices achieve similar Internet reachability to/from the local/remote peer’s customers. Interestingly enough, member ASes that are local peers in some IXPs and remote in others tend to have one order of magnitude larger customer cones than the other cases. This is because hybrid IXP members are usually large ISPs that have diverse peering policies over large geographical areas, engaging both in local and remote

peering depending on their business needs per market segment. Note that the insights pertaining to the customer cones of local, remote and hybrid peers are also reflected in the estimated user populations by APNIC, as expected (results omitted for brevity).

Regarding the country distribution of the IXP members, we found that most local (13.86%) and hybrid (11.04%) peers are headquartered in GB, while PL seems to host the most remote peers (12.88%).

With respect to the traffic levels associated with each network¹⁰, as shown in Fig. 11b, the observed pattern seems to comply with the insights related to the cones and user populations of the IXP members. The distributions of the traffic levels for remote and local peers are similar (albeit with the fraction of local peers per traffic level being larger as expected), while hybrid peers seem to be present also at very high traffic levels, together with locals. It is also interesting that networks with vastly different traffic levels (ranging from 100s of Mbits to 100s of Gbits) engage in RP practices.

6.3 RP Evolution

To understand aspects of the evolution of RP over time, we collect (i) daily RTT measurements (pings) from available LG VPs in 5 IXPs (LINX, HKIX, LONAP, THINX and UAIX), (ii) PDB dumps, and (iii) Atlas traceroutes between 2017/07/04 - 2018/09/10, and we use them to infer remote and local peers across time. Based on this information, we can calculate aggregate growth (*i.e.*, a new member joins an IXP) and departure (*i.e.*, an old member leaves an IXP) rates *per peering type*. We observe that **the number of remote peers grows twice as fast as the number of local peers**, indicating that today, remote peers are the primary drivers of IXP growth (Fig. 12a). These results are confirmed by IXP annual reports from some of the largest IXPs (AMS-IX, DE-CIX, France-IX) [27, 42, 44], indicating that IXPs that already service the majority of local networks in their respective country-level peering ecosystems, seek to expand their market pool by attracting remote peers. However, remote peers also exhibit higher (+25%) departure rates than local ones; reseller customers do not commit substantial resources to establish their IXP connectivity (*e.g.*, routing equipment at the IXP), therefore it might be easier for them to terminate it. For the same time period we also found 18 cases of peers that switched from remote to local interconnections.

¹⁰In Figure 11b, we refer to the aggregate –self-reported via PeeringDB– traffic levels exchanged by the network themselves and not their peering connections.

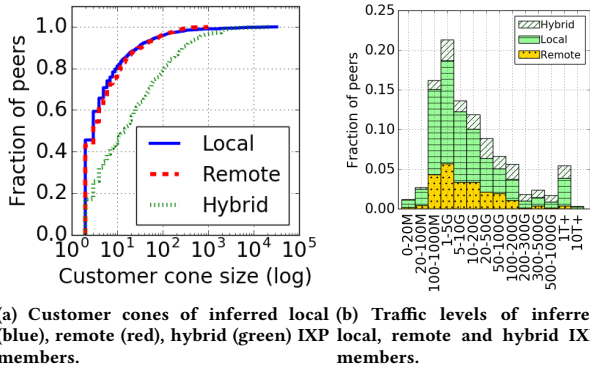


Figure 11: Features of all inferred IXP members.

6.4 RP Routing Implications

Here, as another use case demonstrating the applicability of our inference methodology, we investigate the interplay between remote peering and Internet routing. Specifically, we consider the DE-CIX Frankfurt (FRA) IXP, and examine the routing behavior between its 314 remote members (as inferred by our methodology) and any other of its 781 (local or remote) members (available at the time of measurement). Let AS_R be a remote member of DE-CIX FRA, and AS_x another DE-CIX FRA member (remote or local; $AS_R \neq AS_x$), which peers in at least one more common IXP with AS_R . We are interested in *circuitous paths* that start at AS_R and end at AS_x , which we find with the following process. (i) We randomly choose maximum 5 available (up and running) RIPE Atlas probes within AS_R . (ii) We extract the routed prefixes that AS_x advertises via BGP, using the RIPEstat service [20]. (iii) We select the first IP address (.1) of a randomly chosen prefix among these prefixes. (iv) We run traceroutes from the chosen probes in AS_R towards the selected IP address of AS_x . (v) We extract all traceroute paths involving an IXP crossing (see Section 3.3), either over DE-CIX FRA or another common IXP.

We analyze the results for all possible $\{AS_R, AS_x\}$ pairs ($\sim 245k$ in total). We identify 5941 IXP crossings involving AS_R and AS_x as the two peering IXP members. As described above, these crossings involve either DE-CIX FRA or another IXP where both ASes peer. AS_R and AS_x are also the source and destination of the traceroute(s), respectively. In the majority of the cases (66%), we observe that the routing decision of AS_R seems to comply with an expected hotpotato exit strategy [31, 74], *i.e.*, the IXP involved in the crossing is the closest one to AS_R among the IXPs where both AS_R and AS_x are present. Interestingly enough, on the one hand, we identify cases (18%) where traffic is exchanged via the RP interconnection of AS_R at DE-CIX FRA, while there exists another common IXP that is closer to AS_R . By using this closer IXP, instead of the RP in DE-CIX FRA, AS_R could offload traffic 100s of km closer to its network. On the other hand, there are cases (16%) where the two peers use another (local or remote) peering link (*i.e.*, not over DE-CIX FRA) to exchange traffic, while the facilities of DE-CIX FRA are closer to the AS_R . In the latter cases, AS_R could use the RP over DE-CIX FRA to offload traffic hundreds of km closer to its network.

The reason why in some cases these networks do not make a “seemingly better” (latency-wise) routing decision has to do

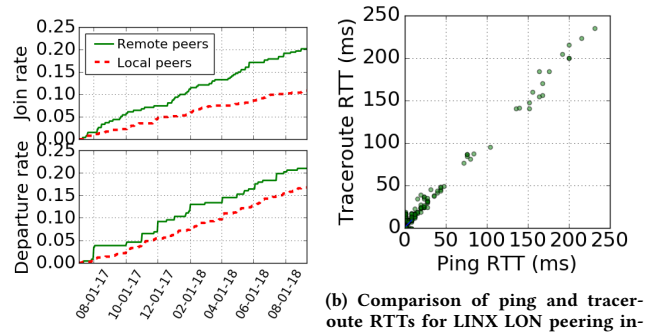


Figure 12: Analysis of archived RTT data.

with their own policies, which are not typically known. Note also that both routing options (over remote peering in DE-CIX or local/remote peering in another IXP), are over peer-to-peer links via IXPs. Therefore, we cannot distinguish routing preferences based on coarse AS-relationships [60] (customer vs. peer vs. provider); this would require taking into account also additional features, such as BGP communities [47], which is the subject of future work.

7 DISCUSSION & INSIGHTS

Ubiquity and growth. We found that RP becomes an increasingly popular practice and is almost ubiquitous in the global IXP ecosystem (Fig. 10b). For instance, in AMS-IX, 90% of new customers come through reseller programs [23]. Exceptions are IXPs that do not support port reselling, but even they facilitate RP for physically distant members, *e.g.*, over L2 carriers. It is worth mentioning that in the past, one of the reasons why local IXP traffic remained local, was because distance meant cost. In light of significantly reduced transport costs, that is no longer the case. Indeed, the largest IXPs have more distant members compared to the average. This is an example of a network effect; the more members an IXP has, the more valuable that IXP is to networks [68]. We observe that most new members at the largest IXPs are remote (Fig. 12a). Interestingly, even smaller IXPs exhibit growing tendencies in terms of RP.

Implications. There are RP cases that have a clear impact on routing paths, and thus, on performance and resiliency. We find evidence that RPs support suboptimal routing choices and introduce latency penalties (Section 6.4). In fact, in many cases, exchanging traffic at an IXP where both traffic source and destination are colocated as members, would be more beneficial for performance (*e.g.*, lower latency; Fig. 1b). For ASes with a global footprint the lack of visibility in whether a peer is local or remote makes their traffic engineering considerably harder. In particular, anycast routing employed by CDNs is affected by RP practices that drive traffic away from the intended load-balancing center, *i.e.*, the IXP itself. In contrast, we also find cases where RP can improve performance by offering better routing choices to a broader set of networks (Section 6.4).

In terms of resiliency, there are potential issues with RP setups. While an extensive investigation of these issues is the subject of future work, here we reason about some obvious resilience implications. Multiple peers connect via the same reseller’s physical port to

one IXP (see Fig. 4); one remote peer may connect (through different resellers) to multiple IXPs with just one router (see Fig. 9d). In these cases, an outage in a single IXP or IXP switch port: (i) propagates much further than the metropolitan area of the IXP over the RP link, and (ii) may affect several IXP members at once. Even if these members have backup paths/links that are activated upon such a cascading outage (e.g., over their transit providers), there is some unavoidable delay and packet loss until BGP routing converges on the backup paths [46]; even after this happens, routing may be sub-optimal in terms of packet-level end-to-end delay.

The IXP's point of view. IXPs themselves do not discriminate members as local or remote; they are simply interested in (i) attracting as many ASes as possible (potentially expanding in multiple geographical regions and colocation facilities or via RP), and offering (ii) short paths and low latencies, (iii) high throughput, and (iv) additional services. However, they are aware of which customers are “virtual” (Section 3.5), *i.e.*, connected on a virtual port offered by a port reseller, since they need to terminate the inner VLAN and configure the IXP-end of the virtual port. For IXPs, RP inference (uncovering also distant, non-virtual IXP peers; see Steps 2-5 in Section 5.2) is interesting for two main reasons: (a) to overcome the local saturation and find new attractive markets/locations for expansion, and (b) to offer to their customers transparency as to who is local and who is not. Finally, we note emerging RP flavors. For instance, NL-IX is a reseller for AMS-IX; DE-CIX offers the *GlobePeer* [9] product that allows an IXP member to acquire access to all DE-CIX peers irrespective of their location.

8 FUTURE RESEARCH DIRECTIONS

Being equipped with a methodology that reliably infers RP, we identify the following directions for future work in this field.

Traffic Analysis. Knowing which IXP peering interconnections are remote/local, a natural follow-up step is to investigate the importance of RP in terms of the actual traffic flowing over the switching infrastructure of an IXP. To achieve this, we need datasets containing the traffic levels of remote and local IXP peering connections.

Beyond Pings. Measuring RTTs using pings from VPs within the IXP suffers from limited and unstable VPs. However, traceroutes from VPs anywhere in the world, can provide an additional source of useful RTT measurements that cover much more in space and time than pings. In fact, the difference of the RTTs between the VP and the consecutive IP interfaces involved in a –potential– IXP crossing, as observed in a traceroute path, can provide an indication of the delay between the associated IXP peers. Fig. 12b shows the RTTs between the LINX LON LG and the interfaces of the members of LINX, using traceroutes (see Section 3.1) and pings; we observe that the RTT patterns are close, supporting such an approach. However, traceroute-based approaches come with their own set of challenges such as asymmetric paths, load-balancing artifacts, ICMP rate-limits and heterogeneous opaque layer-2 connectivity mechanisms [58, 72]. We plan to investigate inference approaches that are robust against such artifacts and enable us to scale up our methodology, decoupling it from ping-based measurements.

Longitudinal Study. In Section 6.3 we analyzed the RP growth of 5 IXPs over a period spanning more than a year. Understanding

whether our observations represent an actual trend, and not just recent developments in the IXP ecosystem, requires digging deeper into history. Since daily RTT measurements (pings) from LGs in IXPs are not available for all IXPs (e.g., during the time-frame of Section 6.3), we aim to apply a traceroute-based methodology to perform an extensive analysis in space (more IXPs) and time (years).

Other Implications/Trade-offs. Evaluating the impact of RP and routing policies (Section 6.4) on the performance of CDNs and anycast services might be of interest to the community. RP is associated to implications for performance, resilience/reliability, and security, and it comes with certain trade-offs (e.g., debugging is more complex when third-party layer-2 infrastructures are involved). Follow-up work could focus on assessing such trade-offs and comparing RP to more traditional connectivity practices, such as classic transit.

9 CONCLUSION

In this work, we introduce, validate, and apply a methodology that can infer remote and local peers at IXPs with high accuracy and coverage. Our methodology is built upon the observation that RP is not driven only by technical factors, but is actually a business decision guided by economic considerations. In particular, taking into account port capacities, colocation strategies, multi-IXP peerings as well as latencies and private connectivity practices, we achieve very high accuracy (95%) and coverage (93%) levels, outperforming the state-of-the-art by +18% and +9%, respectively. At the same time we reduce almost 4 times the false positive and negative rates. The primary objective of this approach is to enable IXPs and existing or new potential IXP members to understand which peers of an IXP are *physically* local, allowing for better-informed peering and routing decisions. Moreover, by equipping researchers with a reliable inference methodology, we enable the in-depth investigation of multiple facets of the peering ecosystem, such as the detection of routing inefficiencies that undermine the resilience and performance of traffic exchange. In our measurement-based study of 30 of the largest IXPs worldwide, we found that more than 90% of them have more than 10% of their members as remote peers. Strikingly, for large IXPs, this share may exceed 40%. The number of remote peers grows twice as fast compared to local peers, driving the IXP growth. The remote peers show similar patterns with local peers in terms of customer cones, user populations and aggregated traffic levels, indicating that remote peering is widely adopted practice across networks. Moreover, we observe that several remote peer routers are connected to more than 10 IXPs, while we also find evidence of hybrid (remote & local) IXP peering interconnections on the same router, with profound implications for routing resilience.

Prototype and Portal. To automate our remote peering inference methodology and make our results publicly accessible to the community, we have implemented a web portal at [17], through which we publish monthly snapshots of our inferences, and visualize the geographical footprint of IXPs and their connected members.

Acknowledgements. We thank the anonymous reviewers and Walter Willinger for their constructive feedback, and Eleni Fragkiadaki for her technical insights. This work has been funded by the European Research Council grant agreement no. 338402.

REFERENCES

- [1] AMS-IX Amsterdam: connected networks. https://ams-ix.net/connected_parties. Accessed: 13.05.2018.
- [2] AMS-IX Amsterdam: services and pricing. <https://ams-ix.net/services-pricing/pricing>. Accessed: 13.05.2018.
- [3] AMS-IX Amsterdam: traffic statistics. <https://ams-ix.net/technical/statistics>. Accessed: 13.05.2018.
- [4] AMS-IX EasyAccess Service. <https://ams-ix.net/services-pricing/easyaccess>. Accessed: 13.05.2018.
- [5] CAIDA AS Relationships. <http://www.caida.org/data/as-relationships>.
- [6] Coresite Carrier List. <https://www.coresite.com/resources/resource-library/additional/carrier-list>. Accessed: 24.09.2018.
- [7] DE-CIX Frankfurt connected networks. <https://www.de-cix.net/en/locations/germany/frankfurt/connected-networks>. Accessed: 13.05.2018.
- [8] DE-CIX Frankfurt traffic statistics. <https://www.de-cix.net/en/locations/germany/frankfurt/statistics>. Accessed: 13.05.2018.
- [9] DE-CIX GlobePEER Remote Service. <https://www.de-cix.net/en/de-cix-service-world/globepeer-remote>. Accessed: 13.05.2018.
- [10] France-IX Resellers List. <https://www.franceix.net/en/members-resellers/resellers>. Accessed: 13.05.2018.
- [11] Hurricane Electric, Internet Exchange Report. <https://bgp.he.net/report/exchanges>. Accessed: 27.03.2018.
- [12] Inflect: Find the right data center. <https://inflect.com>. Accessed: 13.05.2018.
- [13] IX Reach Remote Peering Service. <http://ixreach.com/services/remote-peering>. Accessed: 13.05.2018.
- [14] NL-IX: The Interconnect Exchange. <https://www.nl-ix.net>. Accessed: 13.05.2018.
- [15] Packet Clearing House, Internet Exchange Directory. https://prefix.pch.net/applications/ixpdir/menu_download.php. Accessed: 30.04.2018.
- [16] PeeringDB. <https://www.peeringdb.com>. Accessed: 30.04.2018.
- [17] Remote IXP Peering Portal. <http://remote-ixp-peering.net>.
- [18] RETN Remote Peering Service. <http://retn.net/services/remote-ix>. Accessed: 13.05.2018.
- [19] RIPE Atlas - Measurements. <https://atlas.ripe.net/about/measurements>. Accessed: 13.05.2018.
- [20] RIPEStat Service. <https://stat.ripe.net>. Accessed: 13.05.2018.
- [21] traIXroute - Source Code. <https://github.com/gnomikos/traIXroute>.
- [22] ITU-T Y.1731 Performance Monitoring In a Service Provider Network. https://www.cisco.com/c/en/us/td/docs/ios/cether/configuration/guide/ce_y1731-perfmon.html, Mar 2011. Accessed: 13.05.2018.
- [23] Remote Peering Panel Discussion, 29th Euro-IX Forum, Krakow, Poland. <https://www.euro-ix.net/en/events/fora/29th-euro-ix-forum>, Nov 2016.
- [24] NET-IX Network Map. https://www.netix.net/network_map, Apr 2018.
- [25] AGER, B., CHATZIS, N., FELDMANN, A., SARRAR, N., UHLIG, S., AND WILLINGER, W. Anatomy of a large European IXP. *ACM SIGCOMM CCR* 42, 4 (2012), 163–174.
- [26] AHMED, A., SHAFIQ, Z., BEDI, H., AND KHAKPOUR, A. Peering vs. transit: Performance comparison of peering and transit interconnections. In *Proc. of IEEE ICNP* (2017).
- [27] AMS-IX. AMS-IX 2016 Annual report. https://ams-ix.net/annual_report/2016. Accessed: 05.09.2018.
- [28] APNIC. IPv6 Measurement Campaign. <https://stats.labs.apnic.net/v6pop>. Measurement Methodology: <https://labs.apnic.net/measureipv6>, Dataset collected on: 22.05.2017.
- [29] AUGUSTIN, B., KRISHNAMURTHY, B., AND WILLINGER, W. IXPs: mapped? In *Proc. of ACM IMC* (2009).
- [30] BÖTTGER, T., CUADRADO, F., TYSON, G., CASTRO, I., AND UHLIG, S. A Hypergiant's View of the Internet. *ACM SIGCOMM CCR* 47, 1 (2017).
- [31] CAESAR, M., AND REXFORD, J. BGP routing policies in ISP networks. *IEEE network* 19, 6 (2005), 5–11.
- [32] CAIDA. iffinder. <http://www.caida.org/tools/measurement/iffinder>.
- [33] CAIDA. kapar. <http://www.caida.org/tools/measurement/kapar>.
- [34] CAIDA. Routeviews prefix2as Dataset. <http://data.caida.org/datasets/routing/routeviews-prefix2as>.
- [35] CASTRO, I., CARDONA, J. C., GORINSKY, S., AND FRANCOIS, P. Remote Peering Data. <https://svnnext.networks.imdea.org/repos/RemotePeering>. Accessed: 13.05.2018.
- [36] CASTRO, I., CARDONA, J. C., GORINSKY, S., AND FRANCOIS, P. Remote peering: More peering without internet flattening. In *Proc. of ACM CoNEXT* (2014).
- [37] CHATZIS, N., SMARAGDAKIS, G., BÖTTGER, J., KRENC, T., AND FELDMANN, A. On the benefits of using a large IXP as an Internet vantage point. In *Proc. of ACM IMC* (2013).
- [38] CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. On the Importance of Internet eXchange Points for Today's Internet Ecosystem. *ACM SIGCOMM CCR* (2013).
- [39] CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. There is more to IXPs than meets the eye. *ACM SIGCOMM CCR* 43, 5 (2013), 19–28.
- [40] CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. Quo vadis Open-IX? *ACM SIGCOMM CCR* 45, 1 (2015), 12–18.
- [41] CHIU, Y.-C., SCHLINKER, B., RADHAKRISHNAN, A. B., KATZ-BASSETT, E., AND GOVINDAN, R. Are we one hop away from a better internet? In *Proc. of ACM IMC* (2015).
- [42] DEC-IX. DE-CIX Annual Report 2016. <https://goo.gl/qwCM23>. Accessed: 05.09.2018.
- [43] FANOU, R., VALERA, F., AND DHAMDHARE, A. Investigating the Causes of Congestion on the African IXP substrate. In *Proc. of ACM IMC* (2017).
- [44] FRANCE-IX. France-IX Annual Report 2017. <https://www.franceix.net/annual-report-2017>. Accessed: 05.09.2018.
- [45] GIOTSAS, V., DHAMDHARE, A., AND CLAFFY, K. C. Periscope: Unifying looking glass querying. In *Proc. of PAM* (2016).
- [46] GIOTSAS, V., DIETZEL, C., SMARAGDAKIS, G., FELDMANN, A., BERGER, A., AND ABEN, E. Detecting Peering Infrastructure Outages in the Wild. In *Proc. of ACM SIGCOMM* (2017).
- [47] GIOTSAS, V., LUCKIE, M., HUFFAKER, B., AND CLAFFY, K. C. Inferring complex AS relationships. In *Proc. of IMC* (2014).
- [48] GIOTSAS, V., SMARAGDAKIS, G., HUFFAKER, B., LUCKIE, M., AND CLAFFY, K. C. Mapping Peering Interconnections to a Facility. In *Proc. of ACM CoNEXT* (2015).
- [49] GIOTSAS, V., ZHOU, S., LUCKIE, M., ET AL. Inferring Multilateral Peering. In *Proc. of ACM CoNEXT* (2013).
- [50] GUPTA, A., CALDER, M., FEAMSTER, N., CHETTY, M., CALANDRO, E., AND KATZ-BASSETT, E. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. In *Proc. of PAM* (2014).
- [51] HOLTERRBACH, T., PELSSER, C., BUSH, R., AND VANBEVER, L. Quantifying interference between measurements on the RIPE Atlas platform. In *Proc. of ACM IMC* (2015).
- [52] IXPDB. IXP Database. <https://www.ixpdb.net/en/ix-f/ixp-database>. Accessed: 13.05.2018.
- [53] KARNEY, C. F. Algorithms for geodesics. *Journal of Geodesy* 87, 1 (2013), 43–55.
- [54] KATZ-BASSETT, E., JOHN, J. P., KRISHNAMURTHY, A., WETHERALL, D., ANDERSON, T., AND CHAWATHE, Y. Towards IP geolocation using delay and topology measurements. In *Proc. of ACM IMC* (2006).
- [55] KEYS, K., HYUN, Y., LUCKIE, M., AND CLAFFY, K. Internet-scale IPv4 alias resolution with MIDAR. *IEEE/ACM ToN* 21, 2 (2013), 383–399.
- [56] LI, L., ALDERSON, D., WILLINGER, W., AND DOYLE, J. A first-principles approach to understanding the Internet's router-level topology. In *ACM SIGCOMM CCR* (2004), vol. 34, pp. 3–14.
- [57] LODHI, A., DHAMDHARE, A., AND DOVROLIS, C. Open peering by Internet transit providers: Peer preference or peer pressure? In *Proc. of IEEE INFOCOM* (2014).
- [58] LUCKIE, M., DHAMDHARE, A., CLARK, D., HUFFAKER, B., AND CLAFFY, K. C. Challenges in Inferring Internet Interdomain Congestion. In *Proc. of ACM IMC* (2014).
- [59] LUCKIE, M., DHAMDHARE, A., HUFFAKER, B., CLARK, D., AND CLAFFY, K. C. Bdrmap: Inference of borders between IP networks. In *Proc. of ACM IMC* (2016).
- [60] LUCKIE, M., HUFFAKER, B., DHAMDHARE, A., GIOTSAS, V., AND CLAFFY, K. C. AS relationships, customer cones, and validation. In *Proc. of ACM IMC* (2013).
- [61] MARDER, A., AND SMITH, J. M. MAP-IT: Multipass accurate passive inferences from traceroute. In *Proc. of ACM IMC* (2016).
- [62] MASON, A. G., AND NEWCOMB, M. J. *Cisco secure Internet security solutions*. Cisco press, 2001.
- [63] NIPPER, A. Remote Peering (with A look at Resellers as well), 29th Euro-IX Forum, Krakow, Poland. <https://goo.gl/1ynm26>, Nov 2016. Accessed: 13.05.2018.
- [64] NIPPER, A. PeeringDB Update, DENOG 9. <https://docs.peeringdb.com/presentation/20171123-DENOG9-nipper.pdf>, Nov 2017. Accessed: 13.05.2018.
- [65] NOMIKOS, G., AND DIMITROPOULOS, X. traIXroute: Detecting IXPs in Traceroute Paths. In *Proc. of PAM* (2016).
- [66] NORTON, W. B. *The 2014 Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2014.
- [67] NORTON, WILLIAM B. The Great Remote Peering Debate. <https://goo.gl/yMrELB>, 2012. Accessed: 13.05.2018.
- [68] NORTON, WILLIAM B. Understanding Remote Peering (presentation). <https://goo.gl/3WruyV>, 2013. Accessed: 13.05.2018.
- [69] ROUGHAN, M., WILLINGER, W., MAENNEL, O., PEROULI, D., AND BUSH, R. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *IEEE JSAC* 29, 9 (2011), 1.
- [70] SNIJDERS, J., ABDEL-HAFEZ, S., AND STRONG, M. IXP megabit per second cost & comparison. <https://goo.gl/3nx1FJ>. Accessed: 13.05.2018.
- [71] SOUQUET, S. France-IX reseller programme - Challenges and evolution after 4 years, 29th Euro-IX Forum, Krakow, Poland. <https://goo.gl/jRXS4b>, Nov 2016. Accessed: 13.05.2018.
- [72] STEENBERGEN, R. A. A practical guide to (correctly) troubleshooting with traceroute. *NANOG* (2009), 1–49.
- [73] STOCKER, V., SMARAGDAKIS, G., LEHR, W., AND BAUER, S. Content may be King, but (Peering) Location matters: A Progress Report on the Evolution of Content Delivery in the Internet. In *Proc. of ITS Europe* (2016).
- [74] TEIXEIRA, R., SHAIKH, A., GRIFFIN, T., AND REXFORD, J. Dynamics of hot-potato routing in IP networks. In *ACM SIGMETRICS PER* (2004), vol. 32, pp. 307–319.
- [75] TRAMMEL, B., AND KÜHLEWIND, M. Revisiting the Privacy Implications of Two-Way Internet Latency Data. In *Proc. of PAM* (2018).