Inferring Multilateral Peering

Vasileios Giotsas
University College London
v.giotsas@cs.ucl.ac.uk

Shi Zhou
University College London
s.zhou@ucl.ac.uk

Matthew Luckie
CAIDA / UC San Diego
mjl@caida.org

kc claffy
CAIDA / UC San Diego
kc@caida.org

ABSTRACT

The AS topology incompleteness problem is derived from difficulties in the discovery of p2p links, and is amplified by the increasing popularity of Internet eXchange Points (IXPs) to support peering interconnection. We describe, implement, and validate a method for discovering currently invisible IXP peering links by mining BGP communities used by IXP route servers to implement multilateral peering (MLP), including communities that signal the intent to restrict announcements to a subset of participants at a given IXP. Using route server data juxtaposed with a mapping of BGP community values, we can infer 206K p2p links from 13 large European IXPs, four times more p2p links than what is directly observable in public BGP data. The advantages of the proposed technique are threefold. First, it utilizes existing BGP data sources and does not require the deployment of additional vantage points nor the acquisition of private data. Second, it requires only a few active queries, facilitating repeatability of the measurements. Finally, it offers a new source of data regarding the dense establishment of MLP at IXPs.

Categories and Subject Descriptors

C.2.5 [**Local and Wide-Area Networks**]: Internet; C.2.1 [**Network Architecture and Design**]: Network topology

Keywords

AS relationships; routing policies; multilateral peering; IXP; BGP; topology

1. INTRODUCTION

The Internet inter-domain routing infrastructure is composed of self-organized Autonomous Systems (ASes), which connect to each other using the Border Gateway Protocol (BGP). The BGP-induced AS topology has important implications on the performance, security, and quality-of-service of protocols and applications, and has therefore attracted significant research interest from a variety of disciplines. Many findings have been characterized as controversial due to the widely documented incompleteness of the existing AS topology datasets [22, 31, 23, 35, 21, 33]. The most widely-used data sources for compiling the AS graph use BGP and/or traceroute data. The incompleteness problem derives from the inability of these data sources to observe most p2p links due to policy restrictions imposed by operators that limit propagation of p2p links. The proliferation of Internet eXchange Points (IXPs) to support cost-effective dense peering amplifies the gap between what exists and what can be observed. Researchers have proposed a number of methods to find missing p2p links, such as more effective placement of vantage points [44, 39, 27], aggressive deployment of traceroute monitors at edges of the network [21, 38, 37], or combining different data sources including Internet Routing Registries (IRRs) and looking glass servers [19, 42, 28, 15]. In 2012, Ager *et al.* used sampled traffic (sFlow) data from a large European IXP to discover more peering links at this single IXP than were previously believed to exist in the whole Internet [14].

In this paper we describe, implement, and validate a new method to reveal currently invisible AS peerings using publicly available BGP data sources. Our approach infers peerings established over IXP-provided route servers. Most ASes at an IXP use a route server to implement multilateral peering (MLP) and maximize their peering. MLP is the prevalent peering paradigm in terms of number of p2p links [41, 20], therefore unearthing those links is a critical step towards more complete AS topologies. The default behavior of route servers is to advertise all the routes they learn to all connected peers, though many IXPs allow their members to control how their prefixes are advertised by using a set of special-purpose BGP community values. We implement an algorithm to mine these community values, extract the route server participants, and infer their export policies. By combining BGP data from multiple sources we infer more than 206K p2p links from 13 IXPs for May 2013, 88% of which are not visible in publicly available BGP AS paths.

This paper significantly extends our preliminary work, in which we used IXP route servers and BGP community strings to infer 36K peering links [26]. In this work we develop additional more efficient methods using public looking glasses, which enable us to infer 206K peering links – 5.7 times more than in [26] – at mostly the same set of IXPs. We validate 26K of these links using 70 publicly available looking glasses, finding at least 98.4% exist. In particular, our new methods find most of the peering at DE-CIX: 54K

links up from 7K in our previous work. This IXP had more than 50K links in 2011 [14]. The peering we do not find is established bilaterally across the IXP peering fabric; we claim this is a small fraction of the peering at our IXPs because our result for DE-CIX is comparable to that in [14], and the ASes that engage in bilateral peering are more likely to be selective or restrictive in their peering decisions, i.e. peer with a small fraction of members at an IXP. We also analyze revealing topological characteristics of the discovered links. Notably, 25K (12.4%) of the links are between two stub ASes, making them impossible to observe via BGP unless a vantage point is present in one. Similarly, 114K in our set (55.6%) involve at least one stub, reinforcing Ager's reports of dense peering connectivity at the edge [14]. We show that while some ASes publicly advertise a restrictive or selective peering policy, they engage in dense multilateral peering in selected geographical regions.

The rest of this paper is organized as follows. We begin with an overview of AS topology data sources and their incompleteness. In section 3 we describe how MLP is implemented using route servers and how individual ASes can control which other ASes receive their routes. Section 4 presents our inference algorithm and how we apply it to multiple BGP data sources: BGP data archived by Route Views and RIPE RIS, as well as active queries of looking glasses located at the IXP or clustered close to it. In Section 5 we present our results, demonstrating that 88% of the links we infer were previously invisible, and that at least 98.1% are valid. Section 6 suggests some auspicious next steps for improving our understanding of the Internet topology.

2. BACKGROUND

2.1 Inter-domain Routing

A network of routers under the same administrative entity is called an Autonomous System (AS) and comprises a routing domain. Each AS is identified by a unique 32-bit number (ASN). ASes are autonomous in the sense that they can independently decide which Interior Gateway Protocol (IGP) will be used for routing inside their own domains. To achieve global reachability ASes interconnect and exchange prefix reachability information using the inter-domain routing protocol known as Border Gateway Protocol (BGP). In inter-domain routing, physical connectivity does not imply reachability, which is fundamentally determined by routing policies specified by AS operators, which depend largely on business relationships between ASes.

Business relationships define the economics of routing between ASes and are coarsely divided into three categories. A customer-to-provider (c2p) relationship is established when an AS (customer) pays a better-connected AS (provider) to transit traffic to the rest of the Internet. A peer-to-peer (p2p) relationship is when two ASes agree to exchange traffic only between themselves and their customers to minimize the costs of sending traffic through their providers. Finally, a sibling relationship expresses a connection between ASes owned by the same organization, which freely exchanges traffic without cost or routing restrictions.

BGP routes are usually exported following the *valley-free* rule [24], i.e. a customer route can be exported to any neighbor, but a route from a peer or a provider can only be exported to customers. Hence, a path (a series of adjacent AS links) is valley-free if it complies with one of the following patterns (the sibling links can be found in any position of the path without changing the valley-free property): (1) $n \times$c2p + $m \times$p2c; or (2) $n \times$c2p + p2p + $m \times$p2c; where n and m ≥ 0. Most routing paths are valley-free because they follow from the business interests of ASes, i.e. to minimize operation cost and maximize revenue, although exceptions exist [25].

2.2 AS Topology Data Sources

BGP data is the most widely used source of AS topology data. Route collectors operated by Route Views [12] and RIPE RIS [10] passively collect BGP messages and provide public archives of routing tables and update messages. Routing tables and update messages include an AS Path attribute which identifies the sequence of ASes visited before the route was received, also known as the *control path*. The AS path is the primary source of AS links and is generally considered a reliable source. However, misconfiguration, route hijacks, and path poisoning may induce false links [33, 36]. Other sources of BGP data include Looking Glass (LG) servers that allow the remote execution of non-privileged BGP commands through a web interface or remote login. In the general case LG servers do not allow full BGP table dumps, and typically they are used in one-off queries and not periodic data collection due to the high cost of performing prefix-specific queries.

A second popular source of topology information is IP-level paths collected through globally distributed traceroute monitors. AS links can be inferred by mapping the collected IP addresses to ASNs. However, such mapping is a heuristic and can produce considerable artifacts arising from IP addresses returned by routers in paths that map to a third-party ASes [43]. A third source is the Internet Routing Registry (IRR), a publicly accessible database where AS administrators voluntarily and manually register adjacency and policy information. IRR data are frequently inaccurate, incomplete or intentionally false, although certain databases - notably RIPE - are more reliable. It has been shown that with the proper filtering techniques the IRR can provide a useful source of topology data [40, 16]. Other topology data sources exist, e.g. syslogs, but they are usually proprietary and not available to the research community.

2.3 The Topology Incompleteness Problem

The most significant limitation of the existing BGP data collection projects is the large number of missing links. Missing links are of two types: *hidden* and *invisible* [34]. Hidden links are usually backup c2p links that can be observed if the preferred path towards a prefix changes. Other (typically p2p) links are inherently invisible links due to the limited number and placement of vantage points, and the route propagation restrictions on p2p relationships associated with valley-free routing policies of most ISPs. Invisible p2p links constitute the majority of missing links, and are mostly located in the periphery of the AS graph [35, 33]. BGP feeds are mostly provided by high-tier ASes and some geographic areas are poorly covered. Furthermore, two-thirds of all contributing ASes configure their connection with the BGP collector as a p2p link, which means they advertise only routes learned from customers. AS paths collected from LG servers suffer from similar limitations [30]. Theoretically optimal placement of BGP monitors might mitigate this incompleteness [44, 39, 27], but in practice ASes participate voluntar-

Figure 1: Bilateral (a) vs Multilateral (b) peering for a full-mesh connection between six ASes. Bilateral peering requires $n \cdot (n-1)/2$ BGP sessions, while multilateral peering requires only $c \cdot n$ sessions with c route servers.
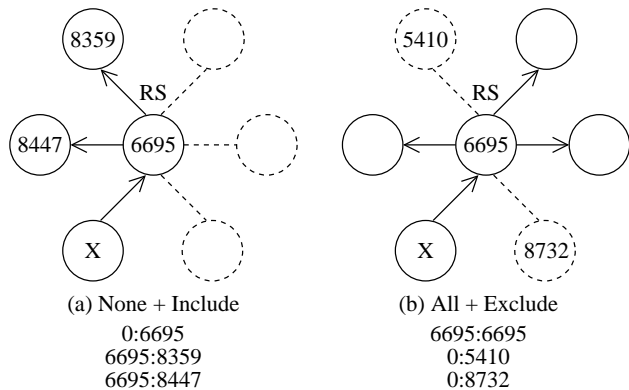

Figure 2: Controlling route advertisements in a route server using BGP communities. In (a), X advertises a route to two selected peers, while in (b) X advertises to all peers except two.

	DE-CIX	MSK-IX	ECIX
RS-ASN	6695	8631	9033
ALL	6695:6695	8631:8631	9033:9033
EXCLUDE	0:peer-asn	0:peer-asn	64960:peer-asn
NONE	0:6695	0:8631	65000:0
INCLUDE	6695:peer-asn	8631:peer-asn	65000:peer-asn

Table 1: Examples of patterns of community values for controlling announcements by a route server. Typically, members use ALL+EXCLUDE or NONE+INCLUDE to control announcements.

ily in such data collection projects so optimal placement is not possible. Some researchers suggest highly distributed traceroute monitoring infrastructures [38, 21] are a promising approach to discover invisible AS links, yet the visibility improvement so far is limited compared with the links discovered at just a single IXP by Ager *et al.* [14].

Another increasingly prominent aspect of the Internet interconnection ecosystem is the proliferation of IXP infrastructure to facilitate cost-effective dense peering. Recent work [28, 15] has demonstrated a vast number of p2p links at IXPs, most of which were not visible in any public data set. This hypothesis has been confirmed by an Internet-wide traceroute study that targeted the discovery of IXP peerings [15]. In total, 58K IXP links were inferred, of which 44K were not visible in any public dataset. Despite the novelty of these techniques [15, 28] they have not been used to provide periodic data due to the complexities of data acquisition. Conducting large-scale targeted traceroute measurements is computationally expensive as well as time-consuming; the methodology in [15] required 16 million traceroute and LG queries (14 days to complete) to discover 44K links.

Recently, Ager *et al.* [14] used sFlow traffic data to discover a rich peering fabric at a large European IXP. Based on actual traffic exchanged they inferred more than 50K p2p links at this IXP alone, about 10K more than was visible in public BGP data for the same period, and more than that discovered across all IXPs using traceroute in [15].

3. MULTILATERAL PEERING

An increasing number of IXPs offer two interconnection paradigms, bilateral and multilateral peering, both of which are illustrated in figure 1. Bilateral agreements require establishing a new BGP session for every peering, which scales poorly at IXPs that mostly have participants with open peering policies who wish to maximize their peering. For IXPs with more than 50K reported peerings [14], managing a separate BGP session for each peer would involve considerable overhead. MLP offers a scalable way to support such dense peering; participants connect with one or more route servers which reflect routes learned from one participant to other participants. Further, some ASes will not enter into bilateral peering unless traffic requirements are met, but will advertise routes to a route server so that smaller ASes can reach them directly. The most notable example is Google,

whose peering policy requires at least 100Mbps peak traffic to establish bilateral peering; they invite networks with less than 100Mbps traffic to peer with them via route servers [6]. Although connection to route servers is optional, a large percentage of IXPs' members opt in. For example, about 77% of the members of the two largest IXPs (DE-CIX and AMS-IX) are connected to the route servers. For the rest of this paper we will use the term *RS members* to refer to route server members.

By default, routes sent to a route server are advertised to all RS members. However, members can control which networks receive their routes through the route server. Filtering mechanisms are essential for IXP participants because even ASes with very open routing policy may not wish to peer with everybody at a given IXP. There are several techniques to implement policy filters, but the most popular practice is through the use of BGP communities, an optional 32-bit BGP attribute used to encode additional information on a BGP route [29]. The values of BGP communities are not standardized, but IXPs clearly document the usage of their community values in IRR records or support pages. There are four common community types among all the IXPs we studied that define the following actions:

- ALL: routes are announce to all RS members. This is the default behavior.

- EXCLUDE: block an announcement toward a specific member. This action can be used in combination with the ALL community to exclude specific RS members from receiving a route.

- **NONE:** Block an announcement toward all RS members. When a community type signals this action, no member receives the route unless they are listed with an INCLUDE community.

- **INCLUDE:** Allow an announcement toward a specific member. This action can be used in combination with the NONE community to allow only specific RS members to receive a route.

Table 1 shows some examples of community values for different IXPs. The `peer-asn` corresponds to the ASN of the RS member that will be included or excluded from receiving an advertisement. In this paper we label route server community values *RS communities*. Because the `peer-asn` part of a community value is 16 bits wide, it is not possible to directly encode 32-bit ASNs. Many IXP operators map the 32-bit ASNs of their members to 16-bit ASNs in the private ASN range to enable filtering of 32-bit ASNs.

Figure 2 illustrates how operators use RS communities to control the announcement of routes by a route server. Figure 2a shows an example of the NONE+INCLUDE scenario listed in table 1. A route tagged with communities `0:6695` `6695:8359` `6695:8447` is advertised by the route server to ASes 8359 and 8447 only. Figure 2b shows an example of the ALL+EXCLUDE scenario. A route tagged with communities `6695:6695` `0:5410` `0:8732` is advertised to all members except ASes 5410 and 8732. Therefore, two ASes can peer via a route server if two requirements are satisfied: *connectivity* and *reachability*. Connectivity is enabled by establishing a session with the route server. Reachability is enabled by configuring outbound filters using RS communities and inbound AS-PATH filters.

4. LINK INFERENCE ALGORITHM

The discovery of IXP links is key to obtaining complete AS topologies. In this section, we explore how one could discover IXP links on a regular basis, at low cost, with public or reproducible measurements. This section presents a framework for the inference of invisible IXP peering links through public BGP data. The key idea behind our methodology is that by obtaining both connectivity and reachability data via IXP infrastructure, it is possible to infer the peering links established with a route server without having to observe them in a BGP or traceroute path.

Connectivity data, namely which ASes are connected to a given route server, can be obtained from three sources of data: (i) looking glass (LG) servers that provide an interface to route servers, (ii) RPSL AS-SETs registered in IRRs by AS operators, and (iii) IXP websites that list connected networks. Information obtained from LGs is the most reliable as it explicitly reports the status of the route server routing table, although previous studies (and our own analysis) have found the other two sources to be accurate and current [28, 15]. Reachability data (RS community strings in route advertisements) can be extracted from public BGP data sources; these include passive BGP measurements (e.g. Route Views and RIS repositories) and active BGP queries of available IXP LG servers.

Before we describe the link inference methodology in detail we first illustrate an example in figure 3 to explain the logic behind our algorithm. In figure 3 ASes A, B, C and D are connected to an IXP route server operated by ASN
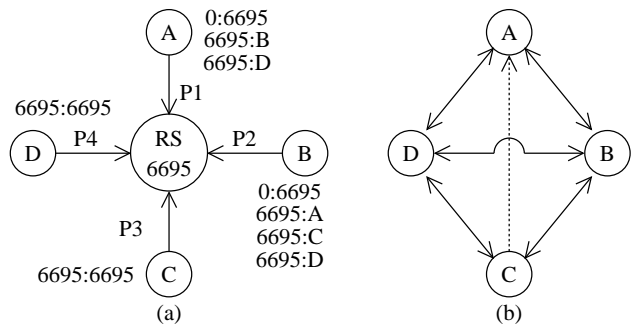

Figure 3: Inferring peering links over a route server using RS communities. The communities sent to the route server are shown in (a) while (b) shows the links that result. C's routes are received by A, but C blocks A from receiving its routes, so we do not infer a p2p link between A and C.

6695. Each of these ASes advertises routes tagged with a set of RS community values. According to these RS communities, all ASes will receive each others' routes except C which is excluded by A. To infer peering links over the route server we also need to know the import filters because an AS may filter some routes received. However, we do not have import filters of all RS members; we overcome this limitation by making the following *reciprocity* assumption: *If an RS member i does not exclude another RS member j from receiving its prefix advertisements, i will also not block the incoming advertisements from j.* Hence, we infer a p2p link between two RS members if they have a reciprocal ALLOW export policy. In figure 3 (b) only A and C do not have a p2p link. In section 4.4 we validate the correctness of this reciprocity assumption against a 230 IRR-based import and export filters set by AMS-IX RS members.

4.1 Inference based on Active BGP Queries

Many IXPs provide public LG interfaces to their route servers which allow the use of non-privileged BGP commands to query the status of the route server routing table. The following steps describe the basic version of our algorithm for using the LG commands to infer the peering links over a route server:

Step 1: Obtain the ASNs and IXP IP addresses of the networks connected to the route server using the `show ip bgp` command. Let A_{RS} be the set of all connected networks on the route server.

Step 2: For each ASN $a \in A_{RS}$ collect the set of prefixes advertised to the route server using the `show ip bgp neighbor [address] routes` command. Let P_a be the set of advertised prefixes for an ASN $a \in A_{RS}$.

Step 3: For each ASN $a \in A_{RS}$ query the prefix information for a subset of its prefixes $P'_a \in P_a$ using the `show ip bgp [prefix]` command. The prefix information will give us the set of RS community values $C_{a,p}$ applied by ASN a when it advertises a prefix $p \in P'_a$ to the route server.

Step 4: From steps 1-3 we have obtained both the connectivity data (A_{RS}) and the reachability data ($C_{a,p}$), so we can infer the peering links established via the route server. For each ASN $a \in A_{RS}$ we construct a set $N_a = \bigcap_p N_{a,p}$ with $N_{a,p} \subset A_{RS}$, which contains the route server partici-

pants toward which all of its routes are advertised. We have two cases depending on the type of RS communities used:

1. *ALL + EXCLUDE*: $N_{a,p} = A_{RS} - E_p$, where E_p is the list of the route server participants excluded by the RS communities, with $E_p \subset A_{RS}$.

2. *NONE + INCLUDE*: $N_{a,p} = I_p$, where I is the list of the route server participants included by the RS communities, with $I_p \subset A_{RS}$.

Step 5: For every pair of ASNs $(a, a') \in A_{RS}$ we infer a p2p link between them if $a \in N_{a'}$ and $a' \in N_a$, i.e. only if both ASes a and a' allow each other to receive their routes.

If the IXP does not provide a LG, we can obtain RS communities from third-party LGs of networks connected to the IXP, i.e., RS members. However, these third-party LGs cannot provide the full view of the RS communities for all the RS members, but only for those members that allow their routes to be advertised to the network that operates the LG.

4.2 Inference based on Passive BGP Data

In addition to active BGP querying via LG servers, our algorithm also works with passive BGP collections from Route Views and RIPE RIS archives. Using passive BGP collections offers three benefits. First, it allows us to extend our inference of p2p links, since not all IXPs provide LG interfaces to their route servers. Second, it can also reduce the number of LG queries necessary, which we explain in this section. Third, it allows us to infer historical peering trends, in conjunction with historical connectivity data through archived IRR records or website archiving services such as the Wayback Machine [13].

Although passive collections do not show the complete route server BGP table, it is still possible to obtain some RS community values. BGP communities are optional transitive attributes; they can be propagated by BGP speakers in the control plane. If at least one route server participant (or one of its customers) provides a vantage point to a Route Views or RIS route collector, it is possible to obtain the RS communities for a large number of IXP participants, depending on the number of RS peerings. For example, we are able to collect the RS communities for 101 of LINX's RS members from AS11666 which participates in LINX RS and contributes a BGP view to the Route Views EQIX collector. Figure 4 extends the example from figure 3. Suppose that AS E is a customer of AS D and contributes a BGP view to a collector. The IXP links that involve D (D–A, D–B, D–C) will be visible in AS paths archived by the route collector, although links (A–B, B–C) will not be seen in AS paths observed by E. However, these paths are also accompanied by the RS communities that B, C and D have applied if RS community values are propagated from the route server to AS D and then to E. Since D is the AS that provides a view of the IXP route server to the BGP collector, we will call it the *RS feeder*. Even if we have only one RS feeder from an IXP, if this feeder is densely connected we can obtain the RS communities for a large number of RS members.

One challenge with the inference of route server peerings using archived BGP data is to determine which AS applied the RS communities and at which IXP. We are able to determine the IXP based either on the upper or the lower 16 bits of the RS community values which typically encode the ASN of the route server. For example, we infer the RS community


Figure 4: Inference of Route Server peering links through passive BGP measurements. Even though we don't observe AS paths with A–B and B–C links, we infer they exist through the community values attached to routes received by E.

values in figure 4 were set at DE-CIX route servers because they contain DE-CIX's ASN (6695). However, sometimes there may be no RS community that encodes the ASN of the IXP. For example, consider the RS community values defined by MSK-IX in table 1. In the ALL+EXCLUDE scenario only the ALL community encodes 8631 which is the ASN of MSK-IX route servers. Since the ALL community is unnecessary because it is the default behavior it may be omitted. Instead, the RS communities may only contain an array of EXCLUDE values of type `0:peer-asn` which makes it difficult to determine the IXP route server as we described above. In such cases we infer the IXP by examining the excluded ASes; each of the excluded ASes may connect to route servers at different IXPs, but often the combination of ASes is only found at a single IXP.

After the RS communities have been extracted and the route server has been identified, we need to pin-point the AS that applied these communities, which we call the *RS setter*. We check every AS in the path against the list of the IXP's participants obtained either through RPSL objects or the IXP's website. We distinguish the following cases:

1. If the AS path contains less than two IXP participants we cannot pin-point the setter.

2. If the AS path contains two IXP participants, we identify as the RS setter the AS closest to the origin. For the first route in figure 4 we know that D and A are RS members; if AS E is not also a RS member then we infer A is the RS setter.

3. If the AS path contains more than two IXP participants, we need to determine which two have a p2p relationship (normally only one p2p relationship should be observed in an AS path, as explained in 2.1). For this purpose we use the AS relationships from [32] which have been shown to have over 99% accuracy for c2p relationships inferred. After we find the IXP participants with the p2p relationship, we identify the RS setter as the AS whose position in the path is closer to the IP prefix. For example, in figure 4, if all E, D and A are members of the route server we check the relationships of the links E – D, and D – A. Since E – D is of c2p type we infer A as the RS setter.

Having identified the RS setters S_{RS} and their community values $C_{s,p}$ we can then infer the route server peerings fol-

Figure 5: CCDF of the number of RS members advertising a given prefix to the DE-CIX route server. 48.4% of prefixes were announced by more than one member.

lowing steps 4 and 5 in section 4.1. When both active and passive measurement data are available for the same IXP route server we combine their data before executing steps 4 and 5.

4.3 Querying Cost

The cost of active measurements is expressed in terms of the number of queries that must be issued and processed [15]. Minimizing this querying cost facilitates more frequent collection. The cost c of our algorithm is given by the following equation:

$$c = 1 + |A_{RS}| + \bigcup_a P'_a \qquad (1)$$

The total number of queries depends on the number of RS members ($|A_{RS}|$) and the number of prefix information commands issued for each prefix queried at step 3. During our five month analysis (January - May 2013) we found the community values applied by each RS member were remarkably consistent among their different prefix announcements toward a specific route server. In fact we found less than 0.5% of cases when a RS member had prefix announcements with different RS communities, and these differences were only found in less than 2% of their prefixes.[1] Therefore by randomly selecting 10% of the prefixes advertised by each RS member, with a maximum of 100 prefixes, we can obtain a consistent view of the RS communities.

We further reduce the querying cost by carefully selecting which IP prefixes to query. Figure 5 shows that 48.4% of prefixes received by the DE-CIX route server were advertised by at least two RS members. By strategically querying BGP information from an LG server for a prefix advertised by multiple RS members, we can obtain the RS communities attached by those members with a single prefix query. Consequently, we sort in decreasing order the prefixes in P_a according to the number of RS members m_p that advertise each prefix to the RS, and we start querying prefix information from the prefix with the highest m_p. In the case of the DE-CIX LG this optimization reduces the total number of queries to 8,400, which is the maximum cost we observed among all the route servers. Without these optimizations we would require 18x more queries. To even further reduce the number of LG queries, we exclude from the active queries

---

[1] Here we refer to prefix advertisements toward the same route server. When an AS is member of more than one IXP route servers its prefix announcements can differ significantly across IXPs, but our algorithm is applied on each route server separately.

the RS members for which the communities are collected through passive BGP measurements. Thus our optimized querying cost can be now expressed as:

$$c = 1 + |A_{RS} - A_{RS}^{passive}| + \bigcup_a (P'_a - P_a^{passive}) \qquad (2)$$

where $A_{RS}^{passive}$ is the set of RS members whose RS communities can be obtained through passive measurements, and $P_a^{passive} = \emptyset$ for each $a \notin A_{RS}^{passive}$. Excluding those prefixes from the active queries reduces the total number of queries to 5,922. By conducting active measurement queries for different IXPs in parallel we can complete all measurements in less than 17 hours even with a rate limit of 1 query per 10 seconds.

4.4 Import and Export Filters

At the beginning of this section, we stated a *reciprocity assumption* according to which we infer the import filters: if a RS member is (not) blocked by the export filter it will also be (not) blocked by the import filter. In other words, if an AS is willing to send traffic to a RS member it will also be willing to accept traffic from it. To validate the correctness of this assumption we use import and export filter data from the members of the AMS-IX route servers. AMS-IX uses an IRR-based filtering mechanism from which the BGP configurations are automatically generated [1], so both import and export filters can be obtained for the AMS-IX members that utilize IRR filtering.

We extracted the filters of 230 AMS-IX RS members by parsing RIPE, ARIN and RADB databases from April 2013. For all of the AMS-IX members the import filters were at most as restrictive as the export filters, and often more permissive. None of the IRR import filters blocked an AS that was not blocked by the corresponding export filter, confirming our assumption. However, about half of the import filters blocked fewer ASes than the export filters, meaning that many RS members are more open at receiving traffic, even from ASes which they do not wish to send traffic. Hence, our assumption is conservative, it does not introduce false-positives but it will miss asymmetric peering links where traffic flows in a single direction.

5. RESULTS

We gathered the RS members and supported RS communities for 13 large European IXPs listed in table 2. We first collected RS communities through passive BGP data to minimize the querying load on LG servers. We accumulated daily BGP table dumps and update messages from Route Views and RIPE RIS repositories for 1-7 May 2013. We filtered out paths that contain (1) reserved, unassigned, and private ASNs (i.e. 23456 and 63488–131071) which should not have appeared in BGP advertisements and (2) path cycles that resulted from misconfiguration and poisoning. We also filtered out transient AS paths to avoid inferring short-lived links that may result from misconfigurations in setting community values. The next step is to query the available LGs for the RS communities of the remaining RS members. We wrote a script to automate this (HTTP) querying of LGs and parsing of responses. Nine of the IXPs provided an LG interface to their route servers[2]; for the remaining

---

[2] France-IX's LG does not output RS communities, so we used Renater's LG which has a feed from France-IX RS.

IXP	LG	ASes	RS	Pasv	Active	Links
AMS-IX	N	574	444	296	55	49249
DE-CIX	Y	483	369	113	256	54082
LINX	N	457	177*	137	39	14759
MSK-IX	Y	374	348	23	325	58501
PLIX	Y	222	211	37	174	21911
France-IX	Y	193	169	103	25	8117
LONAP	N	120	109	30	65	4458
ECIX	Y	102	83	33	50	2751
SPB-IX	Y	89	78	0	78	2828
DTEL-IX	Y	74	71	0	66	1725
TOP-IX	Y	71	52	19	33	1272
STHIX	N	69	42	4	23	340
BIX.BG	Y	53	52	0	52	950

Table 2: Results for the inference of MLP links per IXP. The *ASes* column shows the number of ASes at each IXP, and the *RS* column shows how many of these ASes are connected to the route server; LINX is marked with an asterisk because it does not provide a list of RS members either from its website or an AS-SET. We could only obtain partial data by searching the IRR records of LINX's members for AS8714, the ASN of LINX's route server. The *pasv* column shows the number of RS members whose community strings we obtain from passive BGP data, and the *active* column shows the number of RS members whose community strings we obtain via querying LGs, either directly from the IXP's LG (Y in the *LG* column) or from the LG of a member of that IXP. Finally, the *links* column shows the number of MLP links inferred for each IXP.

IXPs we use 11 LGs provided by their RS members. We can only obtain partial connectivity using only passive BGP data and third-party LGs to collect an IXP's RS communities; in the future we plan to integrate more data sources to expand our view of these IXPs. By combining connectivity information and the RS communities collected by passive and active measurements, we inferred 206,667 multilateral peering links between 1,363 different ASNs. The summation of links in table 2 is larger than 206,667 because 11,821 of the inferred links appear between ASes that co-locate in multiple IXPs. AMS-IX and DE-CIX have the largest link overlap with 7,502 common route server peerings, which is expected given that 123 of the ASes in our study are connected in both IXPs. However, we were only able to collect community data for a part of the AMS-IX and LINX members and therefore the actual link overlap may be larger.

Figure 6 compares the visibility of the MLP links in our dataset against topological data obtained from passive (BGP AS paths) and active (AS paths inferred from traceroute [2, 3]) measurements during the same period. We ordered the RS members by their inferred MLP links, so the line breaks correspond to clustering of IXPs; that is, the BGP-based visibility of each IXP is relatively consistent for all members of that IXP. In the Route Views and RIPE RIS data there were 58,952 visible peering links out of 153,837 total AS links. Only 24,511 (11.9%) of the p2p links are common between our dataset and the public BGP view. Hence, our measurements reveal 209% more peering links, and 18% more AS links than the public BGP view. The inferred links
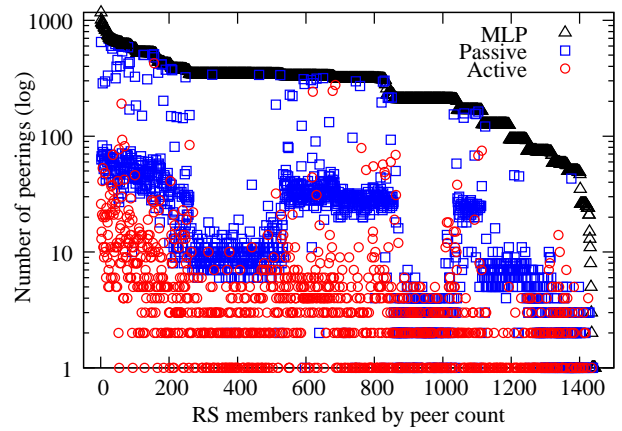

Figure 6: Comparison of the number of MLP links found through our algorithm against passive BGP data (Route Views, RIPE RIS, PCH) and active traceroute data (Ark, DIMES). The inferred MLP links have little overlap with links observed from current active and passive data sources.
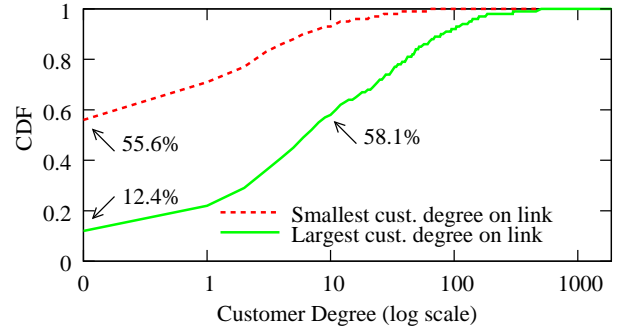

Figure 7: For each p2p link inferred from RS communities, the number of customers of the ASes involved. We plot separate lines representing smallest and largest degrees involved in each peering link. 55.6% of links involve one stub, 12.4% are between two stubs, and 58.1% involve ASes with fewer than 10 customers.

have very little overlap with the links visible to the existing publicly available traceroute topology data: we found only 3,927 links that also appeared in Ark-based and DIMES-based (traceroute-inferred) AS links for the same period [2, 3]. This minimal overlap can be explained by the fact that both Ark and DIMES do not infer links across IXP Route Servers, but report them as links between the RS members and the Route Servers.

To further explore the low visibility of these links in BGP and traceroute-derived AS paths, we examined the customer degrees of the ASes we infer established a p2p link via a route server. The customer degree expresses the number of customers to which an AS provides transit. The customer degree is not affected by the number of invisible links which are almost exclusively p2p links (see section 2.3). Figure 7 shows the degree distributions for the AS with the smallest degree on the link and for the AS with the largest degree. In contrast to what is observable in public BGP data, 12.4% of

Figure 8: The fraction of successfully validated MLP links per AS, classified by the LG type used for validation. Circles correspond to ASes whose LG displays all paths, while triangles correspond to ASes whose LG displays only the best path. LGs that display only the best path may result in lower successful validation ratio when less preferred links are hidden.

IXP	Links	Validated	Links	Confirmed
DE-CIX	6250	11.6%	6152	98.4%
AMS-IX	6190	12.6%	6134	99.1%
MSK-IX	4171	7.1%	4122	98.8%
LINX	3597	24.4%	3492	97.1%
PLIX	2565	11.7%	2515	98.1%
France-IX	1162	14.3%	1126	96.9%
DTEL-IX	732	42.4%	726	99.1%
ECIX	553	20.1%	548	99.1%
LONAP	460	10.3%	455	98.9%
SPB-IX	425	15.0%	422	99.3%
TOP-IX	288	22.6%	288	100%
STHIX	77	22.6%	76	98.7%

Table 3: Validation of the inferred MLP links per IXP. We validate between 7.1% and 42.4% of the links inferred per IXP, and we confirm between 96.9% to 100% of the links tested.

the p2p links in our MLP-inferred dataset are between two stub ASes at the edge of the network. Because these links occur at the edge they are visible in a path only from the ASes involved; unsurprisingly only 1.4% of these links are present in BGP AS paths at Route Views and RIPE RIS. Reinforcing the dense peering at the edge which is enabled by IXPs, 55.6% of all 206K links involve a stub, and 58.1% of the links involve ASes with at most 10 customers.

5.1 Validation of Link Inference Algorithm

To validate the correctness of our link inference framework we test the agreement of the inferred links against connectivity information extracted from other public LG servers. By querying the PeeringDB database we collected the addresses of 70 LGs that are relevant to the inferred links, meaning the LG offers an interface to the collectors of an RS member or one of its customers. For every inferred link relevant to a particular LG, we try to confirm its existence by examining the AS paths returned from the command `show ip bgp [prefix]`. We use up to six different prefixes to ensure that path diversity due to traffic engineering policies will not cause our validation to miss existing links. We select the prefixes to be as geographically distant from each other as possible, based on Maxmind's geolocation database [5].

We repeated our validation at two different time periods to ensure that the correctness of our algorithm is stable over

time. In May 2013 we tested 18,100 links and we successfully confirmed the existence of 98% of those links. In October 2013 we tested 14,513 links and we were able to confirm 98.9% of them. In total we tested 26,392 different RS peerings, and we succeeded in confirming the existence of 98.4% of them. As shown in table 3 the validation results are consistently above 97% for all of the IXPs under study. In the intervening period between link inference and link validation some RS members were disconnected from the route servers or became idle. These cases were filtered out from the October 2013 results explaining the higher validation rate.

Observing a link in the BGP paths of an LG's output confirms the existence of this link, but the reverse does not necessarily hold (not observing a link does not necessarily mean that it does not exist). A path can be hidden from a LG if another path with higher local preference or lower hop count is available, and the LG displays only the active (best) path and not all the available paths. As a result the existence of links that are part of less preferred paths cannot be confirmed by querying LGs that only show active paths. Typically, paths learned from customers are assigned with higher local preference and may hide paths learned from peers. We also found that 14 ASes (out of the 70 used for validation) assigned a higher local preference value to bilateral peers than route server peers. Moreover, in 3 cases the ASN of the route server was not removed from the path making the path appear artificially longer. In all the October 2013 cases where a link failed validation a more preferred path existed. Figure 8 compares the validation results between the looking glass that display all the available paths against the looking glasses that display only the active path. We can see that LGs that only show the active paths restrict our validation effort.

5.2 Peering Policies of RS Members

As shown in table 2, on average 73% of an IXP's members chose to participate in MLP via an available route server. To explain this high participation rate, we collect self-reported peering policy information of IXP members from PeeringDB [9] or on the IXPs' websites where the information is available, e.g., PLIX. We were able to collect policy data for 904 out of the 1,667 IXP members, 72% of which reported an open peering policy, 24% reported a selective peering policy, and 4% claimed a restrictive peering policy.
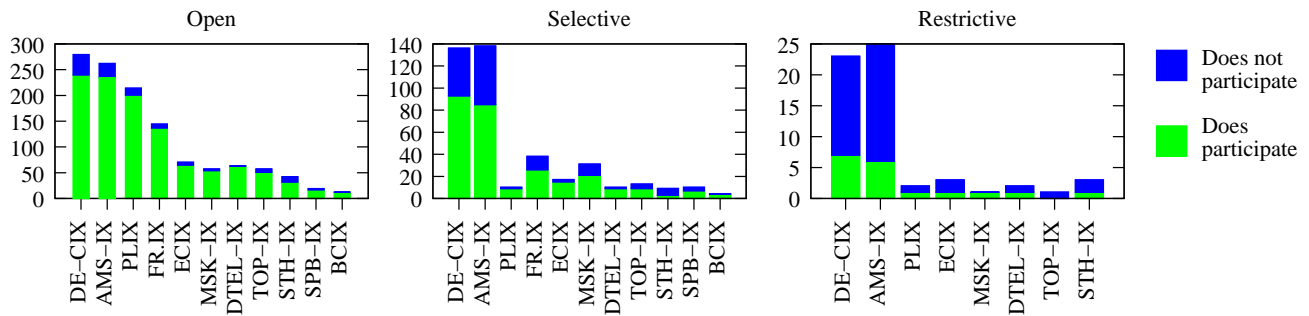
Figure 9: Participation in route servers compared to the self-reported peering policy. Most participants have an open peering policy and use the route servers. Route server participation is common among selective ASes, and rare with restrictive ASes.
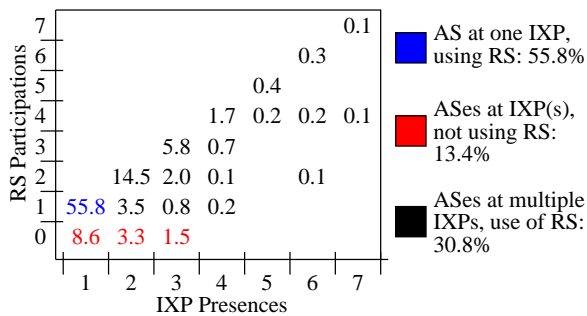

Figure 10: Number of IXPs an AS is connected to against route server participation at those IXPs. 55.8% of ASes in our set are at a single IXP and use its route server. 13.4% of ASes in our set do not use any route server at any of these IXPs (bottom row).

Figure 11: Fraction of RS members allowed to receive prefix announcements from an AS via the RS as a function of that AS's self-reported peering policy in PeeringDB. Nearly all ASes allow either (1) nearly all, or (2) only a small fraction of RS members to receive their routes; the use of RS communities does not scale well for implementing finer-grained filtering. Dots are scattered within each bin to allow visibility.

Figure 9 shows for each of the three self-reported peering policies of IXP members, the distribution of their participation in route servers at 11 IXPs (or only 8 IXPs that involved participants with restrictive peering policies). 92% of the ASes with open peering policies are connected to at least one route server, consistent with their self-reported interest in peering. Interestingly, 75% of the ASes with selective policy and 43% of the ASs with restrictive policy are also connected to at least one route server. These ASes decide to connect to a route server depending not only on their peering policy but also on their business strategy at a given IXP location and their desired relationships with that IXP's candidate peers.

Figure 10 compares the number of route servers where an AS is connected against the total number of the IXPs where the same AS is present. Most (55.8%) ASes in our set are at a single IXP and use its route server. 7.9% of the ASes with presence at multiple IXPs do not have consistent RS participation. For instance, AS9002, a large Ukrainian ISP with a selective peering policy, opts out of connecting to the route servers of Eastern-European IXPs (DTEL-IX and MSK-IX) where many of its customers co-locate, but appears to have an open peering policy in the route servers of Western-European IXPs (DE-CIX and AMS-IX). This suggests peering policies often can have local scope.
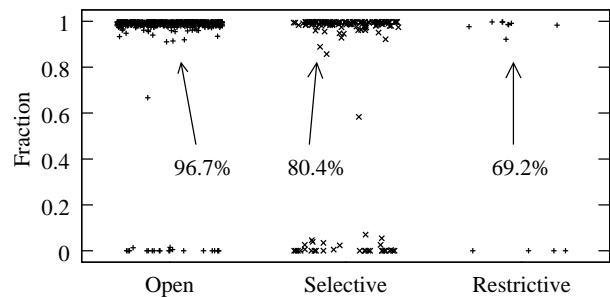
5.3 Route Filtering Patterns of RS Members

Export filters determine the set of RS members with which another RS member intends to peer. Figure 11 shows that there is a binary pattern for most ASes: either very few ASes receive routes, or the vast majority do. Specifically, almost all RS members block fewer than 10% or allow fewer than 10% of other RS participants from receiving their paths. This pattern is congruent with the nature of the most common RS community filters (ALL + EXCLUDE and NONE + INCLUDE); the use of these RS communities does not scale well for implementing finer-grained filtering over route servers, especially for IXPs with hundreds of members.

While the fraction of ASes that use an IXP route server is smaller for ASes that self-report their routing policies as selective and restrictive than for those self-reporting as open (figure 9), it is still the case that most ASes using route servers have an open peering policy. Also, a network's observable MLP behavior is not always consistent with its reported peering policy. Indeed, 69.2% of ASes self-reporting as restrictive and connecting to route servers are also open in establishing RS peerings, though perhaps in only certain regions. A more meaningful classification of RS participants' policies relies on observing their export filters, rather than

Figure 12: Density of peering links per RS member per IXP. For each member at each IXP, we plot the fraction of links the member established that it could have possibly formed using the IXP. The red crosses show the mean density of RS peering observed at each IXP. Dots are scattered within each IXP for visibility.
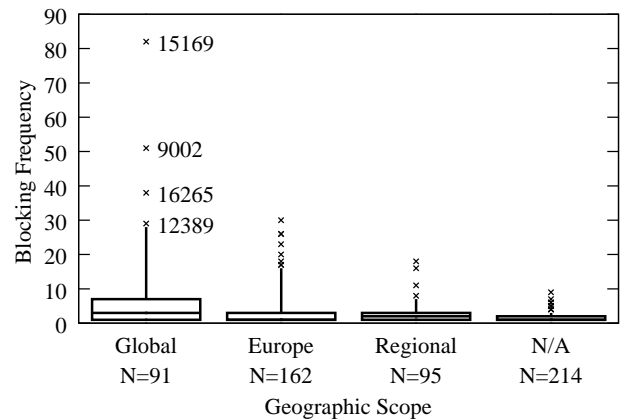

Figure 13: The distribution of blocking frequency using exclude community values, by geographic scope. ASes with an N/A geographic scope did not register their scope in PeeringDB. Google's AS (15169) is blocked 82 times by 75 different ASes that have another peering with Google that they prefer.

the peering policy they self-report in PeeringDB. Studying individual IXPs reveals region-specific (differences in) peering policies. For example, AS12779, an Italian ISP, allowed only 3 RS members of TOP-IX (an IXP in Italy) to receive its prefixes, but blocked only 5 RS members in DEC-IX (an IXP in Germany).

5.4 Peering Density

Peering density can be expressed as the fraction of peering links that a RS member has established over the number of all possible peering links they could establish via the RS. Figure 12 plots this peering density metric for members of the route servers for which we have full connectivity data through the corresponding LG interfaces.[3] The density of RS peering links is between 80%-95%, depending on the number of RS members with open peering policies. Previous work has found the overall peering density of European IXPs to be around 70%, including multilateral and bilateral agreements [15, 14, 17], suggesting that the density of RS-based peering environments is higher than the density of bilateral peering environments.

5.5 Repellers

RS members that are blocked by the EXCLUDE communities can be described as *repellers* by the ASes that set those community values. Of the 1,363 RS members in all the IXPs, 570 are blocked by at least one AS. Figure 13 shows the number of times an AS is blocked compared to the geographic scope of its operations. Although it may seem surprising that global networks are the top repellers of routes by other ASes, these networks are connected to multiple IXPs and there is a significantly larger number of potential RS members they may repel. Of the 1,795 applications of EXCLUDE RS communities, only 12% are set by a provider to block a customer that co-locates in the same route server. To explore the rest of the exclusions, we calculate the customer cones for each RS member using the

algorithm in [32]. The customer cone includes the set of ASes in the downstream path of a provider. We find that 77% of the EXCLUDE community values are used to block an AS that is part of the customer cone. Interestingly the most widely blocked network is AS15169 (Google), which is blocked 82 times by 75 different ASes. This filtering behavior is counter-intuitive for networks characterized by open peering, as AS15169 is an attractive peering partner due to the large volume of traffic that it carries. But in these cases, the AS that blocks AS15169 has a private peering with it in the same IXP or another PoP and prefers the use of the direct peering over the multilateral peering, which we confirmed through the available looking glasses and the IRR records of the RS members. The same behavior is also observed for other large content providers like AS20940 (Akamai), which is blocked 14 times. Hence, the repulsiveness of an RS member is relative to the route server and not necessarily a global characteristic.

5.6 Hybrid Relationships

We observed that 1,230 of the RS links visible in passive BGP data are inferred as provider-customer by CAIDA's relationship inference algorithm [32]. We attempted to clarify whether these relationships are indeed hybrid p2p/p2c relationships, if they have been mistakenly inferred p2c, or if they are actual transit relationships over IXPs which are known to exist although IXP operators discourage such relationships. We collected 422 relationship-tagging and ingress-point tagging BGP community values, defined by 85 different ASes that are involved in 440 hybrid relationships. Combining the two community types, we can learn the relationship type at the different points-of-presence. We were able to verify 202 of these relationships as location-specific hybrid relationships. However, it is difficult to generalize these findings for all the links given the lack of additional data.

---

[3]DTEL-IX is excluded because its LG server restricts queries for 5 RS members who do not wish to disclose their connectivity.

5.7 The Full Picture

Our results reveal that the incompleteness problem is much larger than previously believed. Not only the publicly available data miss the majority of peering links but also past works on link discovery underestimated the total number of peering links (e.g. [15, 21]). To put into perspective the topology incompleteness problem and the contribution of our work, in this section we attempt to estimate the number of IXP peerings globally.

Past works have found that the peering density of Europeans IXPs ranges between 60% – 70% [14, 18]. These findings are supported by data obtained through public IXP peering matrices [8, 7, 11]. According to Cardona *et al.* [18] the peering density depends heavily on the pricing model. Flat-fee pricing encourages the establishment of more peerings and results in peering density close to 70%, while usage-based pricing leads to peering density close to 60%. The availability of route servers is a second important factor that leads to increased peering density by enabling multilateral peering [41, 20]. We utilize these findings to calculate the number of peering links within large European IXPs. For the 37 largest European IXPs with at least 50 members, we collected data from peering registries [4, 9] and individual IXP websites on the members, the pricing model, and the availability of route servers. For IXPs with a flat-fee pricing and available route servers we assume a peering density of 70%, for IXPs with usage-based pricing and available route servers we assume a peering density of 60%, while for IXPs with no route servers we assume a peering density of 50%. Based on these assumptions we estimate the number of European IXP peerings to 558,291. In the case of the highest possible link overlap among those IXPs we estimate 399,732 unique AS peering links.

To expand our estimation we gathered data for all the IXPs globally with at least 50 members. In total we compiled data for 61 IXPs (37 in Europe, 14 in North America, 11 in Asia/Pacific, 1 in Latin America, and 1 in Africa), in which 8,577 different ASes are connected. We maintain the same assumptions on peering density except for the North American IXPs that have a primarily for-profit business model that results in lower peering density [15, 20]. Assuming a peering density of 40% for North American IXPs we estimate the global number of IXP peering links to 686,104, and the number of unique AS peering links to 510,870. For a more conservative estimation we assume that no IXP has peering density over 60%; in this case the global number of IXP peerings would be 596,011, while the total number of unique AS peerings would be 422,423.

5.8 Limitations of our methodology

Despite the substantial number of invisible links discovered, a large fraction of the AS topology is still missing. First, our methodology is limited to the discovery of multilateral peering agreements, and does not capture bilateral peering links. MLP is more prevalent in the European Internet, where we focused this work. IXPs in North America support mostly bilateral peerings, although there are notable exceptions such as Equinix, Any2 and Telx. Moreover, our algorithm requires that a route server utilize BGP communities to control path advertisements, and that these communities are not filtered out. Although this configuration is popular, alternative techniques exist. For example, the Vienna Internet Exchange (VIX) and the Hong Kong Internet

Exchange (HKIX) provide a web portal to configure export filters, while the Netnod route servers strip out all community values before propagating paths to its RS members.

6. CONCLUSIONS

Using new techniques to mine IXP route server data with a mapping of BGP community values, we inferred 206K p2p links from 13 large European IXPs, four times more p2p links than are directly observable in public BGP data. Our approach uses only existing BGP data sources, and requires only few active queries of LG servers, facilitating reproducibility of our results.

As discussed in section 2.3, a variety of challenges in sustainable measurement instrumentation, the nature of the Internet's routing architecture, and the complexity of the ecosystem render topology incompleteness an inherent part of Internet science. Since no single source of topological data can provide a complete topology, advancing our understanding of Internet topology will require extracting all insights we can from all data sources available, including combining data sources where possible and appropriate. Although we implement and validate a new capability to analyze the dense establishment of peering at IXPs, we emphasize that a significant open question relates to how much traffic such IXP peering links carry, which would be an enlightening measure of their relative importance in the global topology.

In the future we plan to apply our methods to additional IXPs and BGP data sources. We will also investigate how our algorithm can reduce the cost of traceroute measurements that target the discovery of IXP peerings, to enable sustainable operational measurement infrastructure using such methods.

7. REFERENCES

[1] AMS-IX Route Servers. `https://www.ams-ix.net/technical/specifications-descriptions/ams-ix-route-servers#IRRdb`.

[2] CAIDA: Archipelago Measurement Infrastructure. `http://data.caida.org/datasets/topology/ipv4.allpref24-aslinks/`.

[3] DIMES AS Edges. `http://www.netdimes.org/PublicData/csv/ASEdges4_2012.csv.gz`.

[4] Euro-IX. `https://www.euro-ix.net/`.

[5] GeoLite IP Geolocaton Database. `http://dev.maxmind.com/geoip/legacy/geolite`.

[6] Google Peering Policy. `https://peering.google.com/about/peering_policy.html`.

[7] GR-IX Peering Matrix. `http://www.gr-ix.gr/services/peering-matrix.shtml`.

[8] INEX Peering Matrix. `https://www.inex.ie/ixp/peering-matrix`.

[9] PeeringDB. `http://www.peeringdb.com`.

[10] RIPE Routing Information Service. `http://www.ripe.net/ris`.

[11] RONIX Connection Matrix. `http://www.ronix.ro/matrice.php`.

[12] RouteViews. `http://www.routeviews.org/`.

[13] The Wayback Machine. `http://archive.org/web/web.php`.

[14] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large European IXP. In *SIGCOMM 2012*, pages 163–174, 2012.

[15] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped? In *IMC '09*, pages 336–349, 2009.

[16] G. D. Battista, T. Refice, and M. Rimondini. How to extract BGP peering information from the Internet routing registry. In *MineNet '06*, pages 317–322, New York, NY, USA, 2006. ACM.

[17] J. C. Cardona Restrepo and R. Stanojevic. A history of an Internet exchange point. *SIGCOMM Comput. Commun. Rev.*, 42(2):58–64, Mar. 2012.

[18] J. C. Cardona Restrepo and R. Stanojevic. IXP traffic: a macroscopic view. In *LANC '12*, pages 1–8. ACM, 2012.

[19] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. Towards capturing representative AS-level internet topologies. *Comput. Netw.*, 44(6):737–755, 2004.

[20] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is more to IXPs than meets the eye. *ACM SIGCOMM CCR*, 43(5), October 2013.

[21] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the sidewalk ends: extending the Internet AS graph using traceroutes from P2P users. In *CoNEXT '09*, pages 217–228, 2009.

[22] Q. Chen, H. Chang, R. Govindan, and S. Jamin. The origin of power laws in Internet topologies revisited. In *INFOCOM 2002*, volume 2, pages 608–617, 2002.

[23] R. Cohen and D. Raz. The Internet dark matter - on the missing links in the AS connectivity map. In *INFOCOM 2006*, pages 1–12, April 2006.

[24] L. Gao and J. Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9:681–692, December 2001.

[25] V. Giotsas and S. Zhou. Valley-free violation in Internet routing: Analysis based on BGP Community data. In *IEEE ICC 2012*, pages 1193–1197, June 2012.

[26] V. Giotsas and S. Zhou. Improving the discovery of IXP peering links through passive BGP measurements. In *16th IEEE Global Internet Symposium*, pages 3365–3370, Turin, Italy, Apr. 2013.

[27] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. On the incompleteness of the AS-level graph: a novel methodology for BGP route collector placement. In *IMC '12*, pages 253–264, 2012.

[28] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: a framework for discovering missing links in the Internet topology. *IEEE/ACM Transactions on Networking*, 17(2):391–404, 2009.

[29] N. Hilliard, E.Jasinska, R. Raszuk, and N. Bakker. Internet Exchange Route Server Operations. `http://tools.ietf.org/html/draft-ietf-grow-ix-bgp-route-server-operations-01`, Aug. 2013.

[30] A. Khan, T. Kwon, H.-c. Kim, and Y. Choi. As-level topology collection through looking glass servers. In *IMC '13*, pages 235–242, 2013.

[31] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *INFOCOM 2003*, volume 1, pages 332–341, 2003.

[32] M. Luckie, B. Huffaker, k. claffy, A. Dhamdhere, and V. Giotsas. AS Relationships, Customer Cones, and Validation. In *IMC '13*, pages 243–256, 2013.

[33] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)completeness of the observed Internet AS-level structure. *IEEE/ACM Transactions on Networking*, 18:109–122, February 2010.

[34] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. In search of the elusive ground truth: the Internet's AS-level connectivity structure. In *SIGMETRICS 2008*, pages 217–228, 2008.

[35] M. Roughan, S. J. Tuke, and O. Maennel. Bigfoot, sasquatch, the yeti and other missing links: what we don't know about the AS graph. In *IMC '08*, pages 325–330, 2008.

[36] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the Internet's autonomous systems. *JSAC*, 29(9):1810–1821, 2011.

[37] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing experiments to the internet's edge. In *USENIX NSDI*, April 2013.

[38] Y. Shavitt and E. Shir. DIMES: let the Internet measure itself. *CCR*, 35(5):71–74, 2005.

[39] Y. Shavitt and U. Weinsberg. Quantifying the importance of vantage points distribution in Internet topology measurements. In *INFOCOM 2009*, pages 792–800, Apr. 2009.

[40] G. Siganos and M. Faloutsos. Analyzing BGP policies: methodology and tool. In *INFOCOM 2004*, volume 3, pages 1640–1651, 2004.

[41] B. Woodcock and V. Adhikari. Survey of characteristics of Internet carrier interconnection agreements. Technical report, Packet Clearing House, May 2011.

[42] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level topology. *CCR*, 35(1):53–61, Jan. 2005.

[43] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, J. Bi, H. Zhang, and L. Zhang. A framework to quantify the pitfalls of using traceroute in AS-level topology measurement. *JSAC*, 29(9):1822–1836, Oct. 2011.

[44] Y. Zhang, Z. Zhang, Z. M. Mao, C. Hu, and B. Maggs. On the impact of route monitor selection. In *IMC '07*, pages 215–220, 2007.