

Remote Peering: More Peering without Internet Flattening

Ignacio Castro^{§†*} Juan Camilo Cardona^{§†*} Sergey Gorinsky[§] Pierre Francois[§]

[§] IMDEA Networks Institute [†] Open University of Catalonia [‡] Carlos III University of Madrid

ABSTRACT

The trend toward more peering between networks is commonly conflated with the trend of Internet flattening, i.e., reduction in the number of intermediary organizations on Internet paths. Indeed, direct peering interconnections bypass layer-3 transit providers and make the Internet flatter. This paper studies an emerging phenomenon that separates the two trends: we present the first systematic study of remote peering, an interconnection where remote networks peer via a layer-2 provider. Our measurements reveal significant presence of remote peering at IXPs (Internet eXchange Points) worldwide. Based on ground truth traffic, we also show that remote peering has a substantial potential to offload transit traffic. Generalizing the empirical results, we analytically derive conditions for economic viability of remote peering versus transit and direct peering. Because remote-peering services are provided on layer 2, our results challenge the traditional reliance on layer-3 topologies in modeling the Internet economic structure. We also discuss broader implications of remote peering for reliability, security, accountability, and other aspects of Internet research.

Categories and Subject Descriptors

C.2.1 [Computer-Communications Networks]: Network Architecture and Design—*Network topology*

General Terms

Measurements

Keywords

Internet; interconnection economics; evolution; transit; peering; remote peering

1. INTRODUCTION

The Internet economic structure is important for reliability, security, and other aspects of Internet design and operation. In spite of its importance, the economic structure remains poorly understood. It is typically modeled on layer 3 of Internet protocols because economic relationships can be inferred from BGP (Border Gateway Protocol) [55] and IP (Internet Protocol) [53] measurements. In particular, BGP identifies ASes (Autonomous Systems) on announced paths, enabling inference of layer-3 structures where ASes act as economic entities interconnected by transit or peering relationships [30]. ASes are imperfect proxies of organizations, e.g., multiple ASes can be owned by a single organization and act as a single unit. Nevertheless, AS-level topologies [19,64] have proved themselves useful for reasoning about Internet connectivity, routing, and traffic delivery. While being useful, layer-3 models struggle to detect and correctly classify a significant portion of all economic relationships in the dynamic Internet.

Validated evolutionary changes in the economic structure include a trend toward more peering. The proliferation of peering relationships is caused partly by their cost advantages over transit. IXPs (Internet eXchange Points) are layer-2 switching facilities where peering commonly takes place [1,8]. IXPs keep growing in the number of their members and amount of peering traffic [15].

Internet flattening refers to a reduction in the number of intermediary organizations on Internet paths [13,24,31]. For example, the Internet becomes flatter when a major content provider expands its own network to bypass transit providers and connect directly with eyeball networks, which primarily serve residential users. Internet flattening is routinely conflated with the trend towards more peering. Indeed, peering relationships are commonly established to bypass transit providers and thus reduce the number of organizations on end-to-end paths.

This paper presents the first empirical and analytical study on an emerging phenomenon of remote peering. Remote peering is an interconnection where a remote network reaches and peers with other networks via a layer-2 intermediary called a remote-peering provider. Remote-peering providers include not only new companies, such as IX Reach [37] and Atrato IP Networks [6], but also traditional transit providers that leverage their traffic-delivery expertise to act as remote-peering intermediaries. By buying a remote-peering service, networks can peer without extending their own infrastructures to a shared location.

* Both authors contributed equally to this paper.

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 ACM 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.

CoNEXT'14, December 2–5, 2014, Sydney, Australia.

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

ACM 978-1-4503-3279-8/14/12 ...\$15.00.

<http://dx.doi.org/10.1145/2674005.2675013>.

Remote peering separates the trends of increasing peering and Internet flattening. On layer 3, remote peering is not distinguishable from direct peering. When a network buys a remote-peering service to establish new paths around a transit provider, the Internet becomes flatter on layer 3 because the new paths bypass the layer-3 transit provider. However, the layer-3 perspective is misleading because the new paths replace the transit provider with the layer-2 remote-peering provider. When one broadens the layer-3 perspective to include the organizations that provide layer-2 services, remote peering does not necessarily reduce the number of intermediary organizations on Internet paths in spite of the enabled additional peering relationships. Hence, remote peering means more peering without Internet flattening.

Our paper reports two measurement studies and a mathematical model that generalizes the empirical findings. First, we develop a ping-based method that conservatively estimates the spread of remote peering. We apply the method in 22 IXPs worldwide and detect remote peers in more than 90% of the studied IXPs, with remote peering by up to 20% of the members at an IXP. Our second study evaluates how remote peering can affect traffic patterns. Based on ground truth from a research and education network, we estimate the amount of transit-provider traffic that this network might offload via remote peering when the number of reached IXPs varies from 1 to 65. The results show a significant offload potential around 25% of the traffic in some scenarios. While the measurements reveal diminishing marginal utility of reaching an extra IXP, we generalize this property in the mathematical model and derive conditions for economic viability of remote peering.

By demonstrating the wide spread and significant traffic offload potential of remote peering, our results challenge the research community’s reliance on layer-3 topologies in representing the Internet economic structure. Due to the failure to include the organizations that provide layer-2 remote-peering services, layer-3 models substantially distort the economic structure and can lead to incorrect conclusions about its properties, e.g., by conflating the trends of increasing peering and Internet flattening. Our findings call for new topological models to represent the prominent role of layer-2 organizations in the Internet economic structure.

The wide spread of remote peering also has broader implications for Internet research. When a provider offers transit and remote peering, buying both might not yield reliable multihoming. The presence of intermediaries that are invisible to layer-3 protocols adds to existing security concerns, e.g., the invisible layer-2 intermediaries can monitor traffic or deliver it through undesired geographies. For Internet accountability, it is a challenge to associate an action with the responsible invisible entity. As a new economic option, remote peering opens a whole new ballgame for connectivity, routing, and traffic distribution, e.g., via newly enabled IXPs [44]. To sum up, our paper makes the following main contributions:

- The paper reports the first systematic study of remote peering. The work illuminates the emerging phenomenon that many in the research community are unaware of. Even those who already know about remote peering benefit from our quantification of its wide spread and significant traffic offload potential.
- Our work reveals separation between the trends of increasing peering and Internet flattening. While en-

abling additional peering relationships, remote peering does not necessarily decrease the number of intermediary organizations on Internet paths.

- The results call for a rethink of modeling the Internet economic structure as layer-3 topologies. There is a need for new models to represent the increasing role of layer-2 entities in the Internet structure.
- The demonstrated prominence of remote peering also has broader implications for reliability, security, accountability, economics, and other aspects of Internet research.

The rest of the paper is organized as follows. Section 2 provides background on Internet economic interconnections. Section 3 empirically studies the spread of remote peering. Section 4 estimates a network’s potential to offload transit traffic to remote peering. Section 5 analyzes economic viability of remote peering versus transit and direct peering. Section 6 discusses broader implications of our findings. Section 7 presents related work. Finally, section 8 sums up the paper.

2. INTERCONNECTION LANDSCAPE

We start by providing relevant background on economic relationships between networks in the Internet.

2.1 Transit

Transit refers to a bilateral interconnection where the customer pays the provider for connectivity to the global Internet. In a common setting, transit traffic is metered at 5-minute intervals and billed on a monthly basis, with the charge computed by multiplying a per-Mbps price and the 95th percentile of the 5-minute traffic rates [25, 62]. In the early commercial Internet, traffic flowed mostly through a hierarchy of transit relationships, with a handful of tier-1 networks at the top of the hierarchy.

2.2 Peering

Peering is an arrangement where two networks exchange traffic directly, rather than through a transit provider, and thereby reduce their transit costs. The exchange is commonly limited to the traffic belonging to the peering networks and their customer cones, i.e., their direct and indirect transit customers. To reduce costs further, peering is typically done at IXPs.

Networks differ in their policies for recognizing another network as a potential peer. The peering policies are typically classified as open, selective, and restrictive [45, 52]. An open policy allows the network to peer with every network. A network with a selective policy peers only if certain conditions are met. A restrictive policy has stringent terms that are difficult to satisfy.

Costs of peering and transit have different structures. Peering involves a number of traffic-independent costs, e.g., IXP membership fees and equipment maintenance expenses at the IXP. Peering also has traffic-dependent costs, e.g., IXP ports for higher traffic rates are more expensive. Over the years, peering relationships have proven themselves as cost-effective alternatives to transit.

Partly due to the lower costs, peering has spread widely, with the IXPs growing into major hubs for Internet traffic. Peering relationships bypass layer-3 transit providers and thus make the Internet flatter, at least on layer 3.

In this paper, *direct peering* at an IXP refers to peering by a network that has IP presence in the IXP location. If a network is not co-located with the IXP already, the network can establish its IP presence at the IXP by contracting an IP transport service or extending its own IP infrastructure to reach the IXP location.

2.3 Remote peering

Remote peering constitutes an emerging type of interconnection where an IP network reaches and peers at a distant IXP via a layer-2 provider [12]. The remote-peering provider delivers traffic between the layer-2 switching infrastructure of the IXP and remote interface of the customer. On the customer's behalf, the remote-peering provider also maintains networking equipment at the IXP to enable the remote network to peer with other IXP members. Figure 1 depicts a typical setting for the remote-peering relationship.

Remote peering provides a smaller connectivity scope than transit. Instead of global Internet access, this service limits the connectivity to the reached IXP members and their customer cones. Technologically, remote peering can be implemented with standard methods, such as those used in layer-2 MPLS (MultiProtocol Label Switching) VPNs (Virtual Private Networks). The main innovation of remote peering lies in its economics.

Remote peering has both traffic-dependent and traffic-independent costs. In comparison to direct peering, the traffic-independent cost is lower, and the traffic-dependent cost is higher: the remote-peering provider has multiple customers and reduces its per-unit costs due to traffic aggregation and acquisition of IXP resources in bulk. Compared to transit, remote peering has lower traffic-dependent costs. Thus, from the cost perspective, remote peering represents a trade-off between direct peering and transit.

IXPs and remote peering are highly symbiotic. IXPs benefit from remote peering because the latter brings extra traffic to IXPs, enriches geographical diversity of IXP memberships, and strengthens the position of IXPs in the Internet economic structure. To promote remote peering, AMS-IX (Amsterdam Internet Exchange), DE-CIX (German Commercial Internet Exchange), LINX (London Internet Exchange), and many other IXPs establish partnership programs that incentivize distant networks to peer remotely at the IXP. For example, some IXPs reduce membership fees for remotely peering networks. AMS-IX started its partnership program around year 2003. According to our personal communications with AMS-IX staff, about one fifth of the AMS-IX members were remote peers at the time of our study.

Implications of remote peering for transit providers are mixed. On the one hand, remote peering gives transit customers alternative means for reaching distant networks. On the other hand, remote peering is a new business niche where transit providers can leverage their traffic-delivery expertise.

According to anecdotal evidence, remote peering successfully gains ground and satisfies diverse needs in the Internet ecosystem. In this paper, we focus on usages where remote peering at IXPs is purchased by distant networks or other IXPs. For example, AMS-IX Hong Kong and AMS-IX interconnect their infrastructures via remote peering to create additional peering opportunities for their members [60]. We do not consider an alternative usage where remote peering at an IXP is bought by a local network to benefit from cost




Figure 1: Directly and remotely peering networks, and probing of their IP interfaces from an LG server

reductions that remote peering provides even over short distances [21].

3. SPREAD OF REMOTE PEERING

In this section, we report measurements that conservatively estimate the spread of remote peering in the Internet.

3.1 Measurement methodology

Because remote peering is provided on layer 2, conventional layer-3 methods for Internet topology inference are unsuitable for the detection of remote peering. For instance, traceroute and BGP data do not reveal IP addresses or ASNs (AS Numbers) of remote-peering providers.

The basic idea of our methodology for detecting a remotely peering network at an IXP is to measure propagation delay between the network and IXP. Specifically, we use the ping utility to estimate the minimum RTT (Round-Trip Time) between the IXP location and the IP interface of the network in the IXP subnet. If the minimum RTT estimate exceeds a threshold, we classify the network as remotely peering at the IXP.

While our ping-based method is intuitive, the main challenges lie in its careful implementation and include: identification of probed interfaces, selection of vantage points, adherence to straight routes, sensitivity to traffic conditions, identification of networks, choice of IXPs, threshold for remoteness, IXPs with multiple locations, impact of blackholing, and measurement overhead. We discuss these challenges below.

Identification of probed interfaces: The targets of our ping probes are the IP interfaces of the IXP members in the IXP subnet. IXP members do not typically announce the IP addresses of these interfaces via BGP. To determine the IP addresses of the targeted interfaces, we look up the addresses on the websites of PeeringDB [52], PCH (Packet Clearing House) [51], and IXP itself.

Selection of vantage points: The ping requests need be launched into the IXP subnet from within the IXP location so that the requests take the direct route from the IXP location to the probed interface. We send the ping requests from LG servers that PCH and RIPE NCC (Réseaux IP Européens Network Coordination Centre) [57] maintain at IXP locations. Figure 1 depicts our probing of IP network interfaces from an LG server at an IXP.

Adherence to straight routes: With our choice of the vantage points, the ping requests and ping replies are expected to stay within the IXP subnet. It is important to keep the probe routes straight because otherwise the RTT mea-

IXP acronym	IXP name	Location		Peak traffic (Tbps)	Number of members	Number of analyzed interfaces
		City	Country			
AMS-IX	Amsterdam Internet Exchange	Amsterdam	Netherlands	5.48	638	665
DE-CIX	German Commercial Internet Exchange	Frankfurt	Germany	3.21	463	535
LINX	London Internet Exchange	London	UK	2.60	497	521
HKIX	Hong Kong Internet Exchange	Hong Kong	China	0.48	213	278
NYIIX	New York International Internet Exchange	New York	USA	0.46	132	239
MSK-IX	Moscow Internet eXchange	Moscow	Russia	1.32	367	218
PLIX	Polish Internet Exchange	Warsaw	Poland	0.63	235	207
France-IX	France-IX	Paris	France	0.23	230	201
PTT	PTTMetro São Paulo	São Paulo	Brazil	0.30	482	180
SIX	Seattle Internet Exchange	Seattle	USA	0.53	177	175
LoNAP	London Network Access Point	London	UK	0.10	142	166
JPIX	Japan Internet Exchange	Tokyo	Japan	0.43	131	163
TorIX	Toronto Internet Exchange	Toronto	Canada	0.28	177	161
VIX	Vienna Internet Exchange	Vienna	Austria	0.19	121	134
MIX	Milan Internet Exchange	Milan	Italy	0.16	133	131
TOP-IX	Torino Piemonte Internet Exchange	Turin	Italy	0.05	80	91
Netnod	Netnod Internet Exchange	Stockholm	Sweden	1.34	89	71
KINX	Korea Internet Neutral Exchange	Seoul	South Korea	0.15	46	71
CABASE	Argentine Chamber of Internet	Buenos Aires	Argentina	0.02	101	68
INEX	Internet Neutral Exchange	Dublin	Ireland	0.13	63	66
DIX-IE	Distributed Internet Exchange in Edo	Tokyo	Japan	N/A	36	56
TIE	Telx Internet Exchange	New York	USA	0.02	149	54

Table 1: Properties of the 22 IXPs in our measurement study on the spread of remote peering

measurements might be high even for a directly peering network. Potential dangers include an unexpected situation where the device of a probed IP interface replies from one of its other IP interfaces and thereby sends the ping reply through an indirect route with multiple IP hops. A more realistic danger is that some of our targeted IP addresses are actually not in the IXP subnet because the respective website information is incorrect. To protect our method from such dangers, we examine the TTL (Time To Live) field in the received ping replies. When ping replies stay within the layer-2 subnet, their TTL values stay at the maximum set by the replying interface [66]. When the path of a ping reply includes an extra IP hop, the TTL value in the reply decreases. Therefore, we discard the ping replies with different TTL values than an expected maximum. We refer to this discard rule as a *TTL-match filter*. For the expected maximum TTL, our experiments accept two typical values of 64 and 255 hops. Although ping software might set the maximum TTL to other values (e.g., 32 or 128 hops), these alternative settings are relatively infrequent, and ignoring them does not significantly increase the number of discarded ping replies in our experiments. Also, different ping replies from the same interface might arrive with different TTL values, e.g., because the replying interface changes its maximum TTL. Whereas we are interested in a conservative estimate for the extent of remote peering, we discard all replies from an IP interface if their TTL value changes during the measurement period. We call this rule a *TTL-switch filter*.

Sensitivity to traffic conditions: Even if a probe stays within the IXP subnet, RTT might be high due to congestion. To deal with transient congestion, we repeat the measurements at different times of the day and different days of the week for each probed IP interface, and record the minimum RTT observed for the interface during the measurement period. This minimum RTT serves as a basis for

deciding whether the interface is remote. Again to be on the conservative side, we exclude an IP interface from further consideration if we do not get at least 8 TTL-accepted ping replies from this interface for each probing LG server. We call this rule a *sample-size filter*. The limit of 8 replies and other parameter values in our study are empirically chosen to obtain reliable results while keeping the measurement overhead low. If less than 4 of the collected ping replies have RTT values within the maximum of 5 ms and 10% of the minimum RTT, i.e., below $RTT_{min} + \max\{5 \text{ ms}, 0.1 \cdot RTT_{min}\}$, we apply an *RTT-consistent filter* to disregard the interface. For an IXP that has both PCH and RIPE NCC servers, we probe each IP interface from both LG servers and exclude the interface from further consideration if the larger of the two respective minimum RTTs is not within the maximum of 5 ms and 10% of the smaller one. We refer to this rule as an *LG-consistent filter*.

Identification of networks: To identify the network that owns a probed IP interface, we use the network’s ASN. We map the IP addresses to ASNs through a combination of looking up PeeringDB, using the IXPs’ websites and LG servers, and issuing reverse DNS (Domain Name System) queries. If the ASN of an IP interface changes during the measurement period, we exclude the IP interface from further consideration. This exclusion rule is called an *ASN-change filter*.

Choice of IXPs: In choosing IXPs, we strive for a global scope surpassing the regional focuses of prior IXP studies. Our choice is constrained to those IXPs that have at least one LG server. Under the above constraints, we select and experiment at 22 IXPs in the following 4 continents: Asia, Europe, North America, and South America. After manually crawling the websites of the IXPs in January 2014, we collect data on their location, peak traffic, and number of members. Table 1 sums up these data. While informa-




Figure 2: Cumulative distribution of the minimum RTTs for all the analyzed interfaces

tion at IXP websites is often incomplete, out of date, or inconsistent in presenting a property (e.g., peak traffic), our measurement method does not rely on these data. We report this information just to give the reader a rough idea about the geography and size of the studied IXPs. For each studied IXP, table 1 also includes the *number of analyzed interfaces*, i.e., interfaces that stay in our analyzed dataset after applying all 6 aforementioned filters. Across all the 22 IXPs, we apply the filters in the following order: sample-size, TTL-switch, TTL-match, RTT-consistent, LG-consistent, and ASN-change. After the filters discard 20, 82, 20, 100, 28, and 5 interfaces respectively, we have a total of 4,451 analyzed interfaces. The high count of TTL-switch discards is likely due to operating system changes during our measurements.

Threshold for remoteness: We classify a network as remotely peering at an IXP if the minimum RTT observed for its IP interface at the IXP exceeds a threshold. Despite the redundancy of our RTT measurements, the minimum RTT might still include non-propagation delays, e.g., due to persistent congestion of the IXP subnet or probe processing in the network devices. To minimize the possibility that such extra delays trigger an erroneous classification of a directly peering network as remote, the threshold should be sufficiently high. Figure 2 plots the cumulative distribution of the minimum RTTs for all the 4,451 analyzed interfaces. A majority of the analyzed interfaces have minimum RTTs distributed almost uniformly between 0.3 and 2 ms. This is a pattern expected for directly peering networks. The likelihood of a network being a direct peer declines as the minimum RTT increases. Our manual checks do not detect any directly peering network with the minimum RTT exceeding 10 ms. Thus, we set the remoteness threshold in our study to 10 ms. While this relatively high threshold value comes with a failure to recognize some remotely peering networks as remote peers, the false negatives do not constitute a significant concern because we mostly strive to avoid false positives in estimating the spread of remote peering conservatively.

IXPs with multiple locations: If an IXP operates interconnected switches in multiple locations, probes from an LG server at one location to an IP interface at another




Figure 3: Classification of the analyzed interfaces with respect to 4 ranges of minimum RTTs

location might have a large RTT. The chosen remoteness threshold of 10 ms is sufficiently high to avoid false positives in cases where all locations of the IXP are in the same metropolitan area. False positives are possible if the geographic footprint of the IXP is significantly larger, e.g., spans multiple countries. We do not observe such situations in our experiments. In a more common scenario, two partner IXPs from different regions, e.g., AMS-IX Hong Kong and AMS-IX, interconnect by buying layer-2 connectivity from a third party. Our methodology correctly classifies such scenarios as remote peering.


Impact of blackholing: If a probed interface intentionally blackholes or accidentally fails to respond to ping requests, the IP interface might be excluded from our analyzed data due to a low number of ping replies for the interface, as discussed above. In a hypothetical (not observed in our experiments) scenario where the probed interface forwards the probe to another machine that sends a ping reply on the interface's behalf, the ping reply is discarded by our TTL-match filter and does not affect accuracy of our RTT measurements.

Measurement overhead: While our method relies on probing from public LG servers, it is important to keep the measurement overhead low. The probes are launched through HTML (HyperText Markup Language) queries to the servers. The LG servers belonging to RIPE NCC and PCH react to an HTML query by issuing respectively 3 and 5 ping requests. For any LG server, we submit at most one HTML query per minute and generally spread the measurements over 4 months. The maximum number of ping replies received from any probed IP interface is 21 and 54 for respectively RIPE NCC and PCH servers.


We conducted the measurements during the 4 months from October 2013 to January 2014. The measurement data are available at [56].

3.2 Experimental results

Figure 3 classifies all the 4,451 analyzed interfaces across the 22 IXPs according to the minimum RTT measured for



(a) Distributions of the IXP counts



(b) Interfaces of all the 285 remotely peering networks

Figure 4: IXP-count distributions and interface classifications for identified networks

each interface. Our conservative estimate finds remote peering in 91% of the studied IXPs. While the numbers of remote interfaces are large in the 3 biggest IXPs (AMS-IX, DE-CIX, and LINX), these numbers are also large at smaller IXPs such as France-IX in France, PTT in Brazil, JPIX in Japan, and TOP-IX in Italy. Despite using the high remoteness threshold of 10 ms, the classification does not reveal remote interfaces in only two IXPs (DIX-IE and CABASE). Hence, our method independently confirms wide presence of remote peering in the Internet economic structure.

The classification in figure 3 looks at the remote interfaces in greater detail by considering the following 3 ranges for the minimum RTT: [10 ms; 20 ms), [20 ms; 50 ms), and [50 ms; ∞) which roughly correspond to intercity, intercountry, and intercontinental distances respectively. We detect the intercontinental-range peering at 12 IXPs, i.e., a majority of the studied IXPs. For example, Italian network E4A remotely peers at both TIE and TorIX, based in the USA and Canada respectively. Brazilian networks comprise most of the remote peers at PTT, the largest among the 21 IXPs of the PTTMetro project in Brazil. The high fraction of remote interfaces at the Turin-based TOP-IX likely results from the IXP’s interconnections with VSIX and LyonIX, two other Southern European IXPs located in Padua and Lyon respectively.

Switching the perspective from the interfaces to the networks that own them, we apply our network identification method (described in section 3.1) to determine ASNs for 3,242 out of the 4,451 analyzed interfaces. While a network might have interfaces at multiple IXPs, we identify a total of 1,904 networks. We refer to the number of the studied IXPs where a network peers as an *IXP count* of the network. Figure 4a presents the distribution of the IXP counts for all the 1,904 identified networks. While a majority of the networks connect to only one IXP, some networks peer at as many as eighteen IXPs.

285 of the identified networks have a remote interface at a studied IXP. Business services offered by the remotely peering networks are diverse and include transit (e.g., Türk Telekom), access (e.g., E4A and Invitel), and hosting (e.g., Trunk Networks). Figure 4a also plots the distribution of the IXP counts for all the 285 remotely peering networks. Both distributions in figure 4a are qualitatively similar, suggesting

that the choice of IXPs for a network to peer is relatively independent of whether the network peers directly or remotely.

We also examine the remotely peering networks with respect to the minimum RTTs of their analyzed interfaces. For each IXP count, we consider all the analyzed interfaces of the remotely peering networks with this IXP count and classify the interfaces in regard to the following 4 ranges of minimum RTTs: [0 ms; 10 ms), [10 ms; 20 ms), [20 ms; 50 ms), and [50 ms; ∞). Figure 4b depicts the fractions of these 4 categories. While our study sets the remoteness threshold to 10 ms, the remotely peering networks with the IXP count of 1 have no interfaces with the minimum RTT below 10 ms. As the IXP count increases, the fraction of the remote interfaces tends to decline because some interfaces of the remotely peering networks are used for direct peering. E4A exemplifies networks with a large number of remote interfaces: 6 of its 9 analyzed interfaces are classified as remote.

3.3 Method validation

While our methodology employs a series of filters and high remoteness threshold of 10 ms to avoid false positives, this section reports how we validate the method and its conservative estimates of remote peering.

First, we use ground truth from TorIX, an IXP located in Toronto. TorIX staff confirmed that their members classified as remotely peering networks in our study are indeed remote peers. In one case, the TorIX staff initially thought that a network identified as a remote peer by our method was rather a local member with a direct peering connection. Nevertheless, a closer examination showed that throughout our measurement period this local member conducted maintenance of its Toronto PoP (Point of Presence) and connected to TorIX from its remote PoP via a contracted layer-2 facility.

Then, we take a network-centric perspective and focus on E4A and Invitel. Both networks specialize in providing Internet access. Based on the measurements, our method classifies the E4A interfaces at DE-CIX, France-IX, LoNAP, TorIX, and TIE as remote. Using public information on IXP websites [3, 43] and insights from private conversations, we confirm that E4A indeed peers remotely at these 6 IXPs.




Figure 5: Network contributions to the transit-provider traffic and offload potential with peer group 4

Our method also identifies Invitel as a remote peer at AMS-IX and DE-CIX, with the minimum RTTs of 22 and 18 ms respectively. Our private inquiries indicate that Invitel uses remote-peering services of Atrato IP Networks to reach and peer at AMS-IX and DE-CIX.

Finally, we receive an independent confirmation that our RTT measurement methodology is accurate. On our request, the TorIX staff measured minimum RTTs between the TorIX route server and member interfaces. Their results for our analyzed interfaces closely match our RTT measurements from the local PCH LG server. The mean and variance of the differences are respectively 0.3 and 1.6 ms.

4. TRAFFIC OFFLOAD POTENTIAL

While section 3 demonstrates that remote peering is common, we now estimate the *traffic offload potential* of a network, i.e., the transit-provider traffic that the network might be able to offload via remote peering. Based on ground truth from a research and education network, we also examine sensitivity of the offload potential to the number of reached IXPs and choice of peers at the reached IXPs.

4.1 Traffic data

We collect and use traffic data from RedIRIS, the NREN (National Research and Education Network) in Spain. This network interconnects with GÉANT (backbone for European NRENs), buys transit from two tier-1 providers, peers with major CDNs (Content Delivery Networks), and has memberships in two IXPs: CATNIX in Barcelona and ES-panix in Madrid. In February 2013, we used NetFlow to collect one month of traffic data at the 5-minute granularity in the ASBRs (Autonomous System Border Routers) of RedIRIS.

Utilizing the BGP routing tables in the ASBRs, we determine the AS-level path and traffic rate for each of the traffic flows. While a network can be associated with a traffic flow in the role of a traffic origin, destination, or intermediary, we classify the traffic flows associated with a network as its:

- (a) *origin traffic*, i.e., originated in the network;
- (b) *destination traffic*, i.e., terminated in the network;
- (c) *transient traffic*, i.e., passing through the network.

Among all the inter-domain traffic, only the traffic between RedIRIS and its transit providers might contribute to the offload potential. Depending on whether RedIRIS receives the traffic from its transit providers or sends the traffic to them, we respectively classify the transit-provider traffic as *inbound* or *outbound*. The collected dataset identifies networks by their ASNs and contains records for 29,570 networks that are origins of the inbound traffic or destinations of the outbound traffic.

To illustrate the contributions of the 29,570 networks to the transit-provider traffic of RedIRIS, we report how much traffic each individual network contributes as an origin of the inbound traffic and destination of the outbound traffic. Figure 5a plots the average traffic rates for the respective inbound and outbound contributions by the individual networks during the measurement period. The figure ranks the networks in the decreasing order of the contributions. While a few networks make huge contributions close to the Gbps mark, most networks contribute little. In the range where the networks are ranked about 20,000 and contribute average traffic rates around 100 bps, the distributions of the inbound and outbound traffic exhibit a similar change in the qualitative profile of the decreasing individual contributions: a bend toward a faster decline. While the raw data exhibit




Figure 6: Origin and destination traffic vs. transient traffic for top contributors to the offload potential

the bend as well, reasons for the bend constitute an interesting topic for future work. Figure 5b reveals daily and weekly fluctuations in the transit-provider traffic of RedIRIS, with the periodic fluctuations being clearly pronounced for the inbound traffic.

4.2 Offload scenarios

RedIRIS cannot offload all of its transit-provider traffic. The offload potential depends on the set of IXPs that the network is able to reach via remote peering. Also, the memberships of the reached IXPs do not include all the networks that contribute to the transit-provider traffic of RedIRIS. Finally, not all the members of the reached IXPs are likely to peer with RedIRIS.

For the set of IXPs that RedIRIS might be able to reach, we consider the Euro-IX association formed, as of February 2013, by 65 IXPs from all the continents [28]. The considered 65 IXPs are a superset of the 22 IXPs studied in section 3, with the set enlargement made feasible by removing the constraint of having LG servers in the IXPs. Based on Euro-IX data from February 2013, we limit potential peers of RedIRIS to the members of these 65 IXPs.

We further trim the group of potential peers by excluding the networks that are highly unlikely to peer with RedIRIS. First, we do not consider the transit providers of RedIRIS as its potential peers because transit providers typically do not peer with their customers. It is worth noting that no network sells transit to these two tier-1 providers, and thus no such network needs to be excluded due to its transitive transit relation with RedIRIS. Second, since RedIRIS already has memberships in CATNIX and ESpanix, the other members of these two IXPs are disregarded as candidates for remote peering with RedIRIS. In particular, we exclude all the other tier-1 networks because they have memberships in ESpanix. Third, due to the cost-effective interconnectivity that comes with the GÉANT membership, we do not consider the other GÉANT members as potential peers of RedIRIS. After applying the above three rules, the group of potential remote peers of RedIRIS reduces to 2,192 networks. Even after eliminating the highly unlikely peers, there remains a significant uncertainty as to which of the 2,192 networks might actually peer with RedIRIS.

To deal with the remaining uncertainty about potential peers, we examine a range of *peer groups*, i.e., groups of networks that might peer with RedIRIS. Using PeeringDB which reports peering policies of IXP members [45, 52], we compose the following 4 peer groups so that the peering policies of their members comprise:

- [peer group 1] *all open policies*;
- [peer group 2] *all open and top 10 selective policies*, which adds to peer group 1 the 10 networks that have the largest offload potentials among the networks with selective policies;
- [peer group 3] *all open and selective policies*;
- [peer group 4] *all policies*, i.e., all open, selective, and restrictive policies.

Peer group 4 constitutes our upper bound on the likely peers of RedIRIS. When RedIRIS reaches all the 65 IXPs, this peer group 4 includes all the aforementioned 2,192 networks. Peer group 1 represents a lower bound on the networks that might actually peer with RedIRIS. It is common for such open-policy networks to automatically peer with any interested IXP member via the IXP route server [58].

For each peer group, we determine the offload potential of RedIRIS by fully shifting to remote peering the traffic that the networks of this peer group and their customer cones contribute to the transit-provider traffic of RedIRIS. While RedIRIS is in control of its outbound transit traffic, we assume that the networks of the peer group shift the inbound transit traffic of RedIRIS to remote peering as well.

In addition to studying sensitivity of the offload potential to the peer groups, we also evaluate its sensitivity to the choice of reached IXPs. Specifically, our evaluation varies the set of reached IXPs from a single IXP to all the 65 IXPs in the Euro-IX data.

4.3 Offload evaluation results

We start by estimating the maximal offload potential with peer group 4 (all policies) when RedIRIS reaches all the 65 IXPs. In this scenario, RedIRIS offloads traffic of 12,238 networks including the 2,192 members of the peer group. Figure 5a shows how much traffic these 12,238 networks contribute to the offload potential in the inbound and outbound directions. The plot ranks the networks in the decreasing order of their traffic contributions. The results suggest that the maximal offload potential is substantial: RedIRIS offloads around 27% and 33% of its transit-provider traffic in the inbound and outbound directions respectively. While the inbound traffic dominates the outbound traffic, figure 5b reveals that the peaks of the transit-provider traffic and offload potential of RedIRIS consistently coincide, implying that the traffic offload can reduce transit bills, which are typically determined by traffic peaks.

Figure 6 zooms in on the top 30 contributors to the maximal offload potential. These 30 networks make the largest traffic contributions to the combined inbound and outbound




Figure 7: Offload potential at a single IXP

offload potential. The top contributors include Microsoft, Yahoo, and CDNs (Content Delivery Networks), suggesting that content-eyeball traffic features heavily in the offload potential. For a majority of the top contributors, the origin and destination traffic dominates the transient traffic.

Switching to the sensitivity analyses, we first evaluate the offload potential for the 4 peer groups when RedIRIS reaches a single IXP. This single IXP is chosen among the 10 IXPs where RedIRIS has the largest offload potential. Figure 7 reports the offload potential of RedIRIS at each of the 10 IXPs. The top 4 of the IXPs include the big European trio (AMS-IX, LINX, and DE-CIX) and Terremark from Miami, USA. For any of the 4 peer groups, the offload potential is similar across the 3 largest European IXPs because these IXPs have many common members. On the other hand, the offload potential at Terremark is significantly different due to its different membership: numerous members of Terremark from South and Central America [49] contribute significantly to the transit-provider traffic of RedIRIS and are not present in Europe.

We now assess the additional value of reaching a second IXP after RedIRIS fully realizes its offload potential at a single IXP. When the two reached IXPs have common members that contribute to the transit-provider traffic of RedIRIS, realizing the offload potential at the first IXP reduces the amount of traffic that RedIRIS can offload at the second IXP. For peer group 4 (all policies), figure 8 illustrates this effect when AMS-IX, LINX, DE-CIX, and Terremark act as either first or second IXP. When LINX and AMS-IX act as the first and second IXPs respectively, the offload potential remaining at AMS-IX after fully realizing the offload potential at LINX is 0.2 Gbps, which is much lower than the full potential of 1.6 Gbps at AMS-IX. When Terremark acts as the second IXP, the decrease in its offload potential is less pronounced because Terremark shares only about 50 of its 267 members with either of the 3 largest European IXPs.

Generalizing the above, we examine the additional value for RedIRIS to reach an extra IXP. We iteratively expand the set of reached IXPs by adding the IXP with the largest remaining offload potential. For peer group 4, the first 4 reached IXPs are added in the following order: AMS-IX, Terremark, DE-CIX, and CoreSite. For all the 4 peer groups, figure 9 plots the remaining transit-provider traffic




Figure 8: Additional value of reaching a second IXP after realizing the offload potential at a single IXP

of RedIRIS as the number of reached IXPs increases. The overall reduction in transit-provider traffic of RedIRIS varies from 8% for peer group 1 (all open policies) to 25% for peer group 4 (all policies). Figure 9 shows that the marginal utility of reaching an extra IXP diminishes exponentially and that reaching only 5 IXPs enables RedIRIS to realize most of its overall offload potential.

While the previous results are specific to RedIRIS, we now present evidence that the property of diminishing marginal utility of reaching an extra IXP holds in general. These additional experiments change the metric from transit-provider traffic to the number of IP interfaces reachable only through transit providers. By adding the IXP that reduces this metric the most, we iteratively expand the set of reached IXPs. Figure 10 plots the number of IP interfaces reachable only through transit providers as the number of reached IXPs increases. Without reaching any IXP, around 2.6 billion IP interfaces are reachable through the transit hierarchy. For peer group 4 (all policies), the number of IP interfaces reachable only through transit providers decreases to about 1 billion after reaching the first IXP. The marginal utility of reaching an extra IXP declines with any of the 4 peer groups. Figures 9 and 10 show that this decline is qualitatively consistent with the exponentially diminishing pattern in the above RedIRIS traffic study. Note that the generalized result for the metric of reachable IP interfaces does not depend on particulars of RedIRIS or another network.

5. ECONOMIC VIABILITY

While section 4 exposes the significant offload potential of remote peering and diminishing marginal utility of reaching an extra IXP, we now generalize the empirical results in the mathematical model and derive conditions for economic viability of remote peering versus transit and direct peering.

5.1 Model

In our model, a network delivers its global traffic via 3 options: (1) transit, (2) expansion of its own infrastructure for direct peering at n IXPs, and (3) remote peering at m IXPs. The respective traffic fractions are denoted as t , d , and r :

$$t + d + r = 1. \quad (1)$$




Figure 9: Additional value for RedIRIS to reach an extra IXP

Total cost C of the traffic delivery consists of transit, direct-peering, and remote-peering components C_t , C_d , and C_r :

$$C = C_t + C_d + C_r. \quad (2)$$

While section 4.3 shows diminishing marginal utility of reaching one more IXP, we fit the RedIRIS data to exponential decay and model the transit traffic fraction as the following function of the number of IXPs where the network peers either directly or remotely:

$$t = e^{-b \cdot (n+m)}. \quad (3)$$

Equality 3 generalizes our empirical results via parameter b that controls how quickly the transit traffic fraction declines. While $b = 0$ represents networks that cannot reduce its transit traffic by peering at distant IXPs, $b = \infty$ enables offload of all transit traffic by reaching a single IXP. Low values of b are characteristic for networks with mostly global traffic, e.g., Google and other networks with highly distributed traffic. The results in figure 4a suggest that networks with high b values are more common. With parameter p denoting the normalized transit price, we model the transit cost as

$$C_t = p \cdot t = p \cdot e^{-b \cdot (n+m)}. \quad (4)$$

The direct-peering cost depends on both the number of reached IXPs and traffic delivered through them:

$$C_d = g \cdot n + u \cdot d. \quad (5)$$

While parameter g accounts for membership fees and other traffic-independent costs of the network in the distant IXPs, parameter u reflects traffic-dependent costs.

The remote-peering cost has a similar structure with traffic-independent parameter h and traffic-dependent parameter v :

$$C_r = h \cdot m + v \cdot r. \quad (6)$$

According to section 2, the per-IXP traffic-independent cost for remote peering is lower than for direct peering:

$$h < g, \quad (7)$$

and the per-unit traffic-dependent cost for remote peering is larger than for direct peering but smaller than for transit:

$$u < v < p. \quad (8)$$




Figure 10: Generalized additional value of reaching an extra IXP

Combining equalities 2, 4, 5, and 6, we express the total traffic-delivery cost of the network as

$$C = p \cdot e^{-b \cdot (n+m)} + g \cdot n + u \cdot d + h \cdot m + v \cdot r. \quad (9)$$

5.2 Analysis

Seeking to minimize its total cost, the network might first consider only transit and direct peering at distant IXPs without purchase of remote peering, i.e., $m = 0$ and $r = 0$. Under this strategy, the total cost is

$$C = (p - u) \cdot e^{-b \cdot n} + u + g \cdot n, \quad (10)$$

and the network minimizes the cost by reaching $\tilde{n}$ IXPs to offload traffic fraction $\tilde{d}$ via direct peering:

$$\tilde{n} = \frac{\log\left(\frac{b \cdot (p-u)}{g}\right)}{b} \text{ and } \tilde{d} = 1 - e^{-b \cdot \tilde{n}}. \quad (11)$$

Continuing from the above solution, the network might widen its strategy to include remote peering, with the total cost becoming

$$C = (p - v) \cdot e^{-b \cdot (\tilde{n}+m)} + (v - u) \cdot e^{-b \cdot \tilde{n}} + g \cdot \tilde{n} + u + h \cdot m. \quad (12)$$

The network minimizes the cost in equality 12 by remote peering at $\tilde{m}$ extra IXPs:

$$\tilde{m} = \frac{\log\left(\frac{g \cdot (p-v)}{h \cdot (p-u)}\right)}{b}, \quad (13)$$

Inequality $\tilde{m} \geq 1$ means that remote peering at one or more IXPs reduces the total cost. Thus, we establish the following condition for economic viability of remote peering:

$$\frac{g \cdot (p - v)}{h \cdot (p - u)} \geq e^b. \quad (14)$$

The economic viability condition in inequality 14 implies that remote peering is more viable for networks with lower b values, i.e., networks with global traffic. Google, major content providers, and other networks with large volumes of global traffic can afford extending their own infrastructures to peer directly at distant IXPs. By realizing this potential for direct peering, the networks with large-volume global traffic contribute to Internet flattening. For such networks, remote peering is just an extra option for bypassing their

transit providers. On the other hand, there are also networks with global traffic that is too small in volume to justify the costs of extending their own infrastructures. Invitel and E4A are examples of such networks. For networks with small-volume global traffic, remote peering constitutes the only economically viable opportunity to reach and peer at distant IXPs. By taking the opportunity to buy the services of remote-peering providers, the networks with small volumes of global traffic increase the extent of peering in the Internet without necessarily making the Internet flatter.

The economic viability condition contains g/h , i.e., the ratio of the per-IXP traffic-independent costs for direct and remote peering. In regions such as Africa, h tends to be much smaller than g because local IXPs offer little opportunities to offload traffic, and transit is expensive [33]. Thus, our analytical model explains why remote peering is economically attractive for African networks. A comprehensive validation of the economic viability model and its implications represents an interesting topic for future work.

6. DISCUSSION

The trends toward more peering and Internet flattening are typically conflated because direct peering interconnections bypass transit providers and thereby reduce the number of intermediary organizations on Internet paths. Complementing direct peering, remote peering enables additional peering as well. However, this increase in peering involves a remote-peering provider that acts as a middleman. Furthermore, the intermediary that sells the remote-peering service can be the same company that provided the bypassed transit services. Hence, remote peering increases peering without necessarily flattening the Internet economic structure.

The observed separation of the two trends questions the usage of AS-level topologies for representing the Internet economic structure. With remote-peering services provided on layer 2, layer-3 modeling of the Internet structure fails to distinguish remote peering from direct peering and ignores the intermediary presence of remote-peering providers. Below, we elaborate on various dangers posed by this omission of the intermediary economic entities.

Layer-3 topologies can make the Internet structure look more reliable than it is. When a company employs the same physical infrastructure to provide transit and remote-peering services, buying both might not translate the redundancy into higher reliability for the multihomed customers.

The emergence of remote peering makes AS-level paths even less representative of the underlying physical paths. The hidden presence of layer-2 remote-peering infrastructures in layer-3 paths creates an additional reason why a path with the smallest number of ASes does not necessarily provide the shortest delay of data delivery.

While it is common to use AS-level models for reasoning about Internet security, the hidden presence of layer-2 intermediaries adds to existing security concerns. The invisible intermediaries might be unwanted entities, e.g., those associated with problematic governments. The risks include monitoring or modification of traffic by the intermediaries and exposure of traffic to other parties, e.g., by delivering it through undesired geographies.

The reliance on layer-3 models also compromises accountability. Whereas a layer-2 intermediary might delay or discard traffic, attribution of responsibility for such perfor-

mance disruptions is complicated because the middleman is invisible on layer 3.

Because remote peering has different economics than transit and direct peering, the omission of the layer-2 intermediaries from layer-3 models weakens economic understanding of the Internet. In developing markets such as Africa, remote peering becomes a cost-effective alternative for reaching well-connected areas in Europe and North America [33]. Since remote peering has a smaller connectivity scope than transit, adoption of remote peering necessitates new strategies for traffic distribution. IXPs greatly benefit from remote peering: existing IXPs gain members, and new IXPs are enabled by bringing together a critical mass of traffic [44]. Ignoring the remote-peering providers distorts substantially the Internet economic landscape.

Thus, our results call for alternative models of the Internet structure that explicitly represent layer-2 entities. The relevant additions include not only remote-peering providers but also other layer-2 economic entities such as IXPs. With the growing prominence of IXPs and remote-peering connectivity to them, integrated modeling of the Internet structure on layers 2 and 3 becomes increasingly important for understanding the Internet. The refined mapping of the Internet economic structure will likely require novel methods for inference of economic entities and their relationships [27].

7. RELATED WORK

The Internet structure is highly important for network accountability [4], multihoming [2], routing security [32], traffic delivery economics [10, 11], and various problems in content delivery via overlay systems [14, 26, 36, 38, 40, 47, 50, 63, 67]. By clarifying the Internet structure, our study of remote peering enables further advances in these and other significant practical domains.

Because network operators do not publicly disclose connectivity of their networks, the research community relies on measurements and inference to characterize the Internet structure [34]. A prominent means for the topology discovery is the traceroute tool that exposes routers on IP delivery paths [19, 22]. For example, iPlane [48] and Hubble [39] use traceroute to generate and maintain annotated Internet maps. Paris traceroute enhances traceroute with the ability to discover multiple paths [7]. A complementary approach is to utilize BGP traces [5, 30, 35, 61, 64]. Our work employs active probing in the data plane to understand the role of remote peering in the Internet structure.

Delay measurements are common in Internet studies, e.g., to understand evolution of Internet delay properties [42] or Internet penetration into developing regions [33]. Our paper uses delay measurements to investigate how geography affects peering of networks.

While previous research studies only mention remote peering [1, 21, 33, 54, 60], our paper is the first to closely examine this emerging type of network interconnection. Our results show wide spread, significant traffic offload potential, and conditions for economic viability of remote peering.

The Internet structural evolution [46] changes the dominant sources of traffic [23] and diversifies network types and their interconnection arrangements [17, 18, 29]. For example, partial transit [65] and paid peering [24] complement the dominant relationships of transit and peering. Our study of remote peering confirms the trend toward diversification of interconnection types.

The arguments that the Internet structure becomes flatter are multifaceted [24, 31, 41], with the continued growth of IXPs [1, 16, 60] cited in support of this trend. Our work reveals separation between the trends of increasing peering and Internet flattening. Also, while analyses of interconnection options are typically restricted to networks that share a location [20, 59], our work exhibits remote peering as a cost-effective solution that enables distant networks to peer over a layer-2 intermediary.

With our results showing the increasing opaqueness of the Internet structure from layer-3 perspectives, the opaqueness is likely to increase further with adoption of software-defined networking [9]. Our paper calls for new approaches to mapping the Internet economic structure on both layers 2 and 3.

8. CONCLUSION

The paper presented the first empirical and analytical study on remote peering. Using careful measurements of RTTs at 22 IXPs worldwide, our ping-based method exposed wide spread of remote peering, with remote peering in more than 90% of the examined IXPs and peering on the intercontinental scale in a majority of them. Based on ground truth from RedIRIS, we also estimated how much transit-provider traffic the network can offload via remote peering when the number of reached IXPs varies from 1 to 65. The assessment showed a significant traffic offload potential around 25% in some cases. While the results exhibited diminishing marginal utility of reaching an extra IXP, reaching only 5 IXPs realized most of the overall offload potential. After generalizing the diminishing marginal utility in the mathematical model, we derived conditions for economic viability of remote peering versus transit and direct peering.

While important in itself as an emerging factor in the Internet ecosystem, remote peering was argued to have broader implications for Internet research. Remote peering revealed separation between the commonly conflated trends of increasing peering and Internet flattening. With remote peering provided on layer 2, we discussed why the omission of remote-peering providers from traditional layer-3 representations of the Internet topology compromised research on Internet reliability, security, accountability, and economics. Finally, we called for refined modeling of the Internet economic structure on both layers 2 and 3.

9. ACKNOWLEDGMENTS

We are grateful to RedIRIS for the access to its traffic data. We also thank TorIX and Jon Nistor for their help in validating our RTT measurements. This research was financially supported in part by the European Commission (FP7-ICT 288021, EINS), Spanish Ministry of Science and Innovation (RYC-2009-04660), and Cisco Systems (CG 573362).

10. REFERENCES

- [1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *Proc. SIGCOMM*, 2012.
- [2] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman. A Measurement-based Analysis of Multihoming. In *Proc. SIGCOMM*, 2003.
- [3] Amsterdam Internet Exchange (AMS-IX). <https://www.ams-ix.net>.
- [4] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker. Accountable Internet Protocol (AIP). In *Proc. SIGCOMM*, 2008.
- [5] D. G. Andersen, N. Feamster, S. Bauer, and H. Balakrishnan. Topology Inference from BGP Routing Dynamics. In *Proc. IMC*, 2002.
- [6] Atrato IP Networks. <https://www.atrato.com>.
- [7] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proc. SIGCOMM*, 2006.
- [8] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *Proc. SIGCOMM*, 2009.
- [9] J. Bailey, R. Clark, N. Feamster, D. Levin, J. Rexford, and S. Shenker. SDX: A Software Defined Internet Exchange. In *Proc. SIGCOMM*, 2014.
- [10] P. Banger and S. Gorinsky. Impact of Prefix Hijacking on Payments of Providers. In *Proc. COMSNETS*, 2011.
- [11] P. Banger and S. Gorinsky. Economics of Traffic Attraction by Transit Providers. In *Proc. Networking*, 2014.
- [12] S. Biernacki. Remote Peering at IXPs. EURONOG 1 presentation, 2011.
- [13] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *Proc. IMC*, 2013.
- [14] F. Cantin, B. Gueye, D. Kaafar, and G. Leduc. Overlay Routing Using Coordinate Systems. In *Proc. CoNEXT*, 2008.
- [15] J. C. Cardona and R. Stanojevic. A History of an Internet eXchange Point. *CCR*, 2012.
- [16] J. C. Cardona and R. Stanojevic. IXP Traffic: A Macroscopic View. In *Proc. LANC*, 2012.
- [17] I. Castro and S. Gorinsky. T4P: Hybrid Interconnection for Cost Reduction. In *Proc. NetEcon*, 2012.
- [18] I. Castro, R. Stanojevic, and S. Gorinsky. Using Tuangou to Reduce IP Transit Costs. *ToN*, in press, 2014.
- [19] H. Chang, S. Jamin, and W. Willinger. Inferring AS-level Internet Topology from Router-Level Path Traces. In *Proc. ITCOM*, 2001.
- [20] H. Chang, S. Jamin, and W. Willinger. To Peer or Not to Peer: Modeling the Evolution of the Internet AS-level Topology. In *Proc. INFOCOM*, 2006.
- [21] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There Is More to IXPs than Meets the Eye. *CCR*, 2013.
- [22] 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 *Proc. CoNEXT*, 2009.
- [23] K. Cho, K. Fukuda, H. Esaki, and A. Kato. Observing Slow Crustal Movement in Residential User Traffic. In *Proc. CoNEXT*, 2008.
- [24] A. Dhamdhere and C. Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *Proc. CoNEXT*, 2010.

- [25] X. Dimitropoulos, P. Hurley, A. Kind, and M. Stoeklin. On the 95-percentile Billing Method. In *Proc. PAM*, 2009.
- [26] Z. Duan, Z.-L. Zhang, and Y. T. Hou. Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning. *ToN*, 2003.
- [27] R. Durairajan, J. Sommers, and P. Barford. Layer 1-Informed Internet Topology Measurement. In *Proc. IMC*, 2014.
- [28] Euro-IX. <https://www.euro-ix.net>.
- [29] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy. In *Proc. TPRC*, 2007.
- [30] L. Gao. On Inferring Autonomous System Relationships in the Internet. *ToN*, 2001.
- [31] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proc. PAM*, 2008.
- [32] P. Gill, M. Schapira, and S. Goldberg. Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security. In *Proc. SIGCOMM*, 2011.
- [33] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. In *Proc. PAM*, 2014.
- [34] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier. Network Topologies: Inference, Modeling, and Generation. *Communications Surveys & Tutorials*, 2008.
- [35] S. Hasan and S. Gorinsky. Obscure Giants: Detecting the Provider-free ASes. In *Proc. Networking*, 2012.
- [36] S. Hasan, S. Gorinsky, C. Dovrolis, and R. Sitaraman. Trade-offs in Optimizing the Cache Deployments of CDNs. In *Proc. INFOCOM*, 2014.
- [37] IX Reach. <http://www.ixreach.com>.
- [38] W. Jiang, S. Ioannidis, L. Massoulié, and F. Picconi. Orchestrating Massively Distributed CDNs. In *Proc. CoNEXT*, 2012.
- [39] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. E. Anderson. Studying Black Holes in the Internet with Hubble. In *Proc. NSDI*, 2008.
- [40] R. Keralapura, N. Taft, C.-N. Chuah, and G. Iannaccone. Can ISPs Take the Heat from Overlay Networks. In *Proc. HotNets*, 2004.
- [41] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-domain Traffic. In *Proc. SIGCOMM*, 2010.
- [42] D. Lee, K. Cho, G. Iannaccone, and S. Moon. Has Internet Delay Gotten Better or Worse? In *Proc. CFI*, 2010.
- [43] London Internet Exchange (LINX). <https://www.linx.net>.
- [44] LINX NoVA. <https://www.linx.net/service/publicpeering/nova>.
- [45] A. Lodhi, N. Larson, A. Dhamdhare, C. Dovrolis, and K. Claffy. Using PeeringDB to Understand the Peering Ecosystem. *CCR*, 2014.
- [46] R. Ma, J. Lui, and V. Misra. On the Evolution of the Internet Economic Ecosystem. In *Proc. WWW*, 2013.
- [47] H. V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani. A Structural Approach to Latency Prediction. In *Proc. IMC*, 2006.
- [48] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *Proc. OSDI*, 2006.
- [49] D. Madory. "Crecimiento" in Latin America. Renesys Blog, 2013.
- [50] V. Misra, S. Ioannidis, A. Chaintreau, and L. Massoulié. Incentivizing Peer-Assisted Services: a Fluid Shapley Value Approach. In *Proc. SIGMETRICS*, 2010.
- [51] Packet Clearing House. <https://www.pch.net>.
- [52] Peering Database. <https://www.peeringdb.com>.
- [53] J. Postel. Internet Protocol DARPA Internet Program Protocol Specification. RFC 791, 1981.
- [54] A. H. Rasti, N. Magharei, R. Rejaie, and W. Willinger. Eyeball ASes: from Geography to Connectivity. In *Proc. IMC*, 2010.
- [55] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771, 1995.
- [56] Remote Peering Data. <https://svnnext.networks.imdea.org/repos/RemotePeering>.
- [57] Réseaux IP Européens Network Coordination Centre (RIPE NCC). <http://www.ris.ripe.net/cgi-bin/lg/index.cgi>.
- [58] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *Proc. IMC*, 2014.
- [59] S. Shakkottai and R. Srikant. Economics of Network Pricing with Multiple ISPs. *ToN*, 2006.
- [60] J. H. Sowell. Empirical Studies of Bottom-Up Internet Governance. In *Proc. TPRC*, 2012.
- [61] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP Topologies with Rocketfuel. *ToN*, 2004.
- [62] R. Stanojevic, N. Laoutaris, and P. Rodriguez. On Economic Heavy Hitters: Shapley Value Analysis of 95th-percentile Pricing. In *Proc. IMC*, 2010.
- [63] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante. Drafting Behind Akamai: Inferring Network Conditions Based on CDN Redirections. *ToN*, 2009.
- [64] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *Proc. INFOCOM*, 2002.
- [65] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. Vazirani. How Many Tiers? Pricing in the Internet Transit Market. In *Proc. SIGCOMM*, 2011.
- [66] H. Wang, C. Jin, and K. G. Shin. Defense Against Spoofed IP Traffic Using Hop-count Filtering. *ToN*, 2007.
- [67] M. Yu, W. Jiang, H. Li, and I. Stoica. Tradeoffs in CDN Designs for Throughput Oriented Traffic. In *Proc. of CoNEXT*, 2012.