

# An Analysis of the Largest National Ecosystem of Public Internet eXchange Points: The Case of Brazil

Samuel Henrique Bucke Brito, Mateus Augusto Silva Santos, Ramon dos Reis Fontes, Danny Alex Lachos Perez, Hirley Dayan Lourenco da Silva and Christian Esteve Rothenberg

**Abstract**—Internet eXchange Points (IXP) have become an increasing research target when aiming at understanding the complex and evolving Internet ecosystem. IXPs are shared infrastructures where Autonomous Systems (AS) implement peering agreements for their traffic exchange and thus represent an interesting microcosm of the Internet diversity and a strategic vantage point to deliver end-user services.

In this article, we provide an in-depth analysis of the largest set of public IXPs in a single country, namely the case of Brazil. The Brazilian public peering ecosystem counts with over 25 IXPs maintained by an overarching project called IX.br following a non-profit business model that facilitates multilateral agreements. The nation-wide peering initiative provides an appealing environment for innovation and fostering IP connectivity market practices. Without IX.br, access providers are limited in terms of coverage, performance, cost, and dependence on transit providers. The open and incentive-rich IXP approach can be regarded as an interesting development that may inspire other development countries as well as more established regional markets. Based on BGP data from all looking glass servers in IX.br, we provide insights into the peering ecosystem per IXP and from a nation-wide perspective by inspecting properties of the connectivity graphs and the IPv4 and IPv6 prefix distribution. We propose peering affinity as a metric well-suited to measure the connectivity between different types of ASes and overall found lower peering density in IX.br when compared to more mature ecosystems, such as AMS-IX, DE-CIX, LINX, and MSK-IX. When dissecting AS-level graphs we also observe the formation of IXP-enabled k-clique communities. As a final contribution, we have shared our 16 GB dataset along all supporting code to allow for new research studies by the community.

**Index Terms**—IXP, BGP, Autonomous System, Routing.

## I. INTRODUCTION

Internet eXchange Points (IXP) are a relevant approach to promote Internet development in terms of connectivity and performance. IXP facilities, located at strategic places throughout nations, allow dozens or hundreds of Autonomous Systems (AS) to connect and agree on their traffic exchange. The increased participation of ASes at IXPs is contributing to the critical role of these tactical IXPs infrastructures in the overall Internet ecosystem [2, 8].

Many ASes justify their interest on peering at IXPs because of performance benefits [28]. With video traffic representing

50% and growing of the total Internet traffic, peering at IXPs allows a better distribution of content closer to end users and reducing transit costs. During the 2014 Soccer World Cup, Brazilian IXPs played a critical role in delivering the traffic<sup>1</sup> and are expected to be again an important infrastructural piece during the 2016 Olympic Games.

Recently, remote peering [6] has emerged as an IXP connectivity trend that allows ASes becoming members of regional IXPs without the costs of placing routers, renting physical space (collocation), in addition to human resources OPEX. Remote peering is contributing to the increase of IXP participants, leading to larger amounts of traffic exchange and more routes available within the IXP ecosystem, altogether increasing the strategic value of IXPs and “flattening” the Internet (cf. [6, 11]).

The OpenIX<sup>2</sup> and Euro-IX<sup>3</sup> are efforts in the US and Europe, respectively, to promote the development of IXPs. One remarkable example is AMS-IX<sup>4</sup>, operating in Netherlands with more than 700 members, which is currently the world’s largest IXP. The Brazilian Internet Steering Committee, through Brazil’s National Internet Registry (NIC.br), leads the Latin American region by operating more than half of the existing IXPs. Since 2006, Brazil has grown from 4 IXPs to the current 25 in operation –with 16 new locations under evaluation. Under the code name IX.br, the project manages all public IXPs in Brazil, including PTT-SP (Sao Paulo), the largest in Latin America averaging around 600 Gbps traffic with peaks of 1 Tbps. As shown in Table I, IX.br is among the world’s top ten IXPs in terms of traffic and PTT-SP alone is the fifth largest one in the number of members.

Many efforts have been devoted towards a better understanding of the complex Internet ecosystem [21] through studies based on diverse data, such as (i) information obtained in interviews or centralized databases, (ii) data plane tools (e.g. traceroute), and (iii) BGP tables available in route servers [1, 5, 15, 16, 25, 26, 34]. IXPs have become attractive research targets because they represent a microcosm of Internet diversity [7], serving a variety of members such as ISPs of different nature, content providers, as well as public and private organizations.

In this article, we bring out an extensive data collection and in-depth analysis work considering all public IXPs operating in Brazil. We first classify all AS participants and generate

The authors are members of Information & Networking Technologies Research & Innovation Group (INTRIG) at the School of Electrical and Computer Engineering (FEEC) at University of Campinas (UNICAMP), in Campinas, São Paulo, Brazil. Feel free to contact shbbrito@dca.fee.unicamp.br.

A preliminary version of this paper was presented in International Conference on Passive and Active Network Measurement (PAM). Heraklion, Crete, Greece. March 31 and April 01, 2016 [3].

Manuscript received November 17, 2015; revised August 23, 2016.

Digital Object Identifier: 10.14209/jcis.2016.23

<sup>1</sup><https://labs.ripe.net/Members/emileaben/internet-traffic-during-the-world-cup-2014>

<sup>2</sup><http://www.open-ix.org>

<sup>3</sup><https://www.euro-ix.net>

<sup>4</sup><https://ams-ix.net>

TABLE I  
TRAFFIC OF SOME OF THE WORLD’S LARGEST PUBLIC IXPs (AUGUST 03, 2016).

IXP	Country	Members	Maximum Throughput (Gbps)			Average Throughput (Gbps)		
			Daily	Monthly	Yearly	Daily	Monthly	Yearly
(01) DE-CIX	Germany	600+	4,637.10	4,735.50	5,178.40	2,966.50	2,878.00	2,912.70
(02) AMS-IX	Netherlands	797	8,517.00	-	9,419.00	5,663.00	-	5,699.00
(03) LINX	United Kingdom	630	2,610.00	2,740.00	3,340.00	1,910.00	1,900.00	2,460.00
(04) MSK-IX	Russia	384	1,696.30	1,601.50	2,135.50	1,130.10	1,023.80	1,061.50
(05) NL-ix	Netherlands	576	2,440.00	-	-	1,582.18	-	-
<b>(06) IX.br</b>	<b>Brazil</b>	<b>715</b>	<b>1,760.00</b>	<b>1,800.00</b>	<b>1,570.00</b>	<b>1,160.00</b>	<b>1,140.00</b>	<b>1,060.00</b>
(07) HKIX	Hong Kong	241	1,101.85	1,260.31	1,260.31	737.90	785.35	637.39
(08) SIX	USA, Canada	224	1,122.74	1,122.74	1,136.42	825.33	781.12	718.26
(09) JPIX	Japan	138	315.54	-	-	200.00	-	-
(10) JINX	South Africa	24	15.50	29.30	20.80	11.30	13.90	16.80

(01) <http://www.de-cix.net/about/statistics/>  
 (03) <https://www.linx.net/pubtools/trafficstats.html>  
 (05) <https://www.nl-ix.net/network/traffic/>  
 (07) <http://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html>  
 (09) <http://www.jpix.ad.jp/en/technical/traffic.html>

(02) <https://ams-ix.net/technical/statistics>  
 (04) <http://www.msk-ix.ru/network/traffic.html>  
 (06) <http://ix.br/cgi-bin/all>  
 (08) <http://www.seattleix.net/agg.htm>  
 (10) <http://stats.jinx.net.za/showtotal.php>

AS-level connectivity graphs (per IXP and nation-wide) for the analytic studies. The observed topologies shed light on the existing peering density (individually by AS and grouped by category) evidencing the peering potential at IXPs. We also look at advertised routes (AS-PATH), average vertices’ degree, path depth, traffic engineering practices based on AS-Prepend BGP knobs, and IPv4/v6 prefix announcements. In addition, we study emerging k-clique communities in the IXP ecosystem. Different from related work on AS-level topologies (e.g., [22] [27] [33] [13]), our empirical analysis is the first with focus on the Brazilian public IXP ecosystem, which corresponds to the largest set of public IXPs worldwide. When simultaneously looking at all national IXPs, we use new metrics and methodologies to compute the peering density between different types of ASes.<sup>5</sup>

Our publicly available dataset –currently the largest in the context of the Latin America IXPs– has more than 16 GB of data. It provides information from all 25 Brazilian public IXPs, including the classification of their members, (IPv4 and IPv6) BGP tables, spreadsheets, connectivity graphs, as well as all coding and supporting tools (e.g. gnuplot) used in this paper. In order to facilitate the work of fellow researchers, the full dataset is available through the research group public repository<sup>6</sup>.

The remainder of this article is organized as follows. Section II brings background information on IXP history, architecture, and business and operational models. Section III describes our methodology based on the pro workflow and developed framework to gather and process the datasets, including the reconnaissance of limitations in our research. Section IV is the largest section of the paper and details the analytic studies providing a discussion of the results. Section V discusses related work. Finally, Section VI concludes the paper and points to avenues for future work.

<sup>5</sup>The peering affinity metric and findings from our first analyses were introduced in [3].

<sup>6</sup><https://github.com/intrig-unicamp/ixp-ptt-br>

II. PRIMER ON INTERNET EXCHANGE POINTS

A. Historical Overview

From 1987 until 1994 the NSFNET backbone was under the management of National Science Foundation (NSF) and it was in 1992 that started a plan to transfer the Internet core operations to the private sector. It was in the context of this new commercial ecosystem that emerged three key elements: (i) Network Service Providers (NSP), responsible for the backbone operation; (ii) Network Access Points (NAP), to carry traffic between NSPs from distributed locations inside USA; and (iii) Routing Arbiter (RA), to collect and advertise routing information at NAPs (similar to modern Route Servers). The inception of IXPs can be rooted back to this point in time when NAPs were created to be a physical point to connect several NSPs. During the following years, NAPs became weak because they were maintained by telecommunications operators that had their own interests, a situation that motivated the disconnection of competitor operators.

The big telecommunications operators had interest to directly connect (peering) with other same level telcos, showing no interest in open peering with smaller telcos. Around 1999, the need for point-to-point circuits between big telcos was scaling linearly and was expensive, besides the fact that for some telephone companies it took more than a year to deliver such connectivity [28]. In the face of this situation, the big telcos perceived that the establishment of private peering through a neutral-IXP was faster and cheaper, which contributed to the adoption of the IXP peering model.

B. Fundamentals of an IXP

An IXP is a shared infrastructure at a certain location installed to facilitate the interconnection of ASes through peering agreements [7]. IXPs contribute to the Internet performance by keeping the traffic exchange as local as possible between different networks in the same geographical region, reducing the number of hops between ASes. The IXP network core is a high-performance switching arrangement capable of sustaining high traffic rates between dozens or even hundreds




Fig. 1. IXP General Architecture

of IXP members’ BGP-speaking devices (routers). The IXP network architecture is arguably simple and resembles those of modern data centers (e.g., non-blocking, fat-tree or spine-leaf designs) with the main goal being a high-available and high-performance Ethernet switch-fabric for AS router L2 connectivity in addition to a number of supporting services (e.g., VLAN isolation, security). An example architecture of an IXP can be observed in Figure 1.

After physical connectivity is up, IXP members agree to establish multilateral (open) peering with all other members or bilateral (private) peering of selective or restrictive nature, implemented through BGP policy configuration to IP filtering and prefix advertisement. To avoid the hazards of establishing peering (BGP sessions) in a full-mesh topology, Route Servers (RS) are used as centralized control plane points so that an AS can reach all members by establishing a single BGP session with the RS (multilateral peering). Looking Glass (LG) servers are commonly used to mirror all existent routes in the BGP table of RS, allowing IXP’s members to verify errors and validate new configurations via telnet access.

C. Business & Operation Models

The nature and services of an IXP largely depend on its business and operational model, i.e. the entity running and operating the IXP infrastructure and its vision, incentives, and commercial considerations. Following the approximate classification of [8], we can distinguish “for-profit” and “non-profit” IXPs, which can be further divided into “cooperative” and “managed” non-profit IXPs (e.g., DE-CIX, AMS-IX, LINX). The latter, mainly found in Europe, are considered among the most vibrant and innovative IXPs [8]. In the US, the predominant business model of IXPs is private, for profit. In most cases, the common goal of IXPs is to provide an open and indiscriminate environment of shared nature to motivate collaboration and win-win situations that improve the traffic exchange and leverage the overall quality of the Internet traffic while simplifying operations.

D. The case of Brazil: IX.br

The case of Brazil follows an interesting approach that may inspire other countries, especially in development regions. Brazilian IXPs are part of an overarching project called IX.br

and adopt a non-profit business model managed and fully funded by NIC.br, the Brazilian Internet Steering Committee that takes care of (and financial income from) DNS registry services, IP allocation, in addition to Internet development activities funded by the government. The attractive cost proposition for ASes added to the open policies of the IX.br business and operational model are the main factors in the leadership of Brazil in terms of the amount of IXPs, including the largest one in the city of Sao Paulo. <sup>7</sup>

One relevant observation is that ASes peered at IX.br are not allowed to rely on the public IXP as their only Internet link. For this reason, to join the free IX.br IXP services, ASes need to prove they already reach the global Internet through some transit provider.

There are national plans to install new IXPs all over the country, especially in the north, west and central regions where there is a concerning deficit of Internet connectivity compared to the south, southeast, and northeast. The goal behind the IX.br expansion plan is to attract ISPs (access providers) to those isolated areas lacking of connectivity by offering the IXP incentives (free co-location, peering opportunities, etc.). According to NIC.br, there are currently 45 candidates interested in hosting the new IXPs. Interested entities, be it commercial or not, are only requested to operate neutrally and free of fees to IXP participants.

III. METHODOLOGY

We now present the workflow and framework used for our in-depth analysis of the Brazilian ecosystem of public IXPs. We designed and implemented a data collection and analytics framework to extract insights from all public IXPs of IX.br, which turns out to be the largest set of public IXPs worldwide.

A. Information Sources

Two initial sources of information were (i) the web page of IX.br project, under responsibility of Brazil’s National Internet Registry (NIR) called NIC.br, and (ii) the online tool PeeringDB<sup>8</sup>. During our analysis we found that PeeringDB data regarding Brazilian IXPs and their members were outdated and incomplete, an impasse confirmed when we crossed the information from PeeringDB with those provided by NIC.br.

The most relevant information source was compiling BGP raw data through telnet access to all IX.br LGs, totaling an input dataset of over 11 GB.

B. Workflow: Data In, Knowledge Out

Our workflow and processing framework to gather data and generate outputs (knowledge) is as follows. As shown in Figure 2, the first step is to access every Brazilian IXP via telnet through the address lg.<code>.ptt.br, where <code> is replaced by the two or three letters in column “Code” of Table II. Table II also contains the list of each

<sup>7</sup>Similar to most important IXPs worldwide, NIC.br maintains a website (<http://ix.br>) with detailed technical specification of the infrastructure, up-to-date membership lists, traffic statistics, and additional information.

<sup>8</sup><http://www.peeringdb.org>

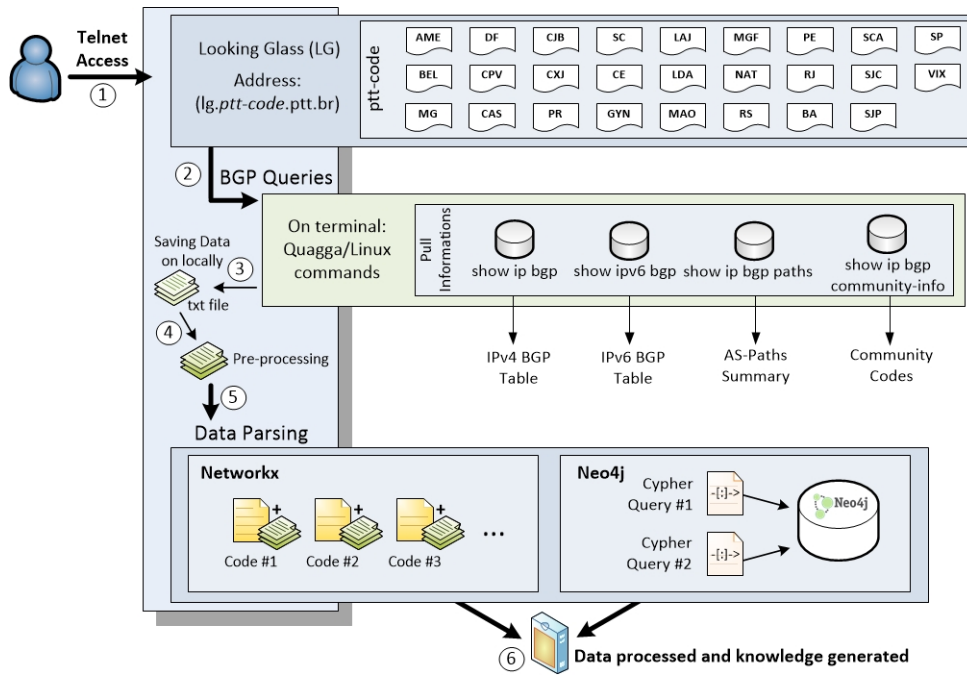


Fig. 2. Workflow of data gathering and processing.

Brazilian IXP and its operation region, the average exchange traffic (Gbps) and the quantity of members (column “M”). Once connected by telnet at each IXP LG, the second step is to query BGP<sup>9</sup> to collect the following data: (i) control plane BGP table (both IPv4 and IPv6), (ii) list of BGP AS-PATH, and (iii) community codes. The raw dataset with the output of these BGP queries is first locally stored as simple text files (step 3), and then parsed/pre-processed (step 4). Finally, the datasets go through a set of analytic functions (steps 5 and 6) implemented with different graph-oriented tools, as explained next in Section III-C.

The manual and time-consuming task described in steps 1, 2 and 3 were automated through the developed framework consisting of a set of scripts to automatically access every Brazilian IXP by telnet and save the outputs from the different BGP queries in the corresponding text files, as described in Algorithm 1. The script contains as input the list of codes of each Brazilian IXP, allowing it to automatically open a telnet session to each one of the IXPs based on their standardized address format `lg.<code>.ptt.br`. After the session is established, command queries to the routing stack (be it based in Quagga, BIRD, IOS, JunOS, etc.) are used to retrieve the relevant BGP data and produce the output text files.

<sup>9</sup>By means of the issuing the corresponding BGP routing stack commands of the Route Server implementation, typically Quagga in the Brazilian IX.br IXPs but also BIRD is very popular in IXPs around the world.

TABLE II  
PUBLIC IXPs OPERATING IN BRAZIL. (OCTOBER 09, 2015)

#	City	State	Code	Avg Gbps	Members
01	Belem	PA	BEL	0.41	14
02	Belo Horizonte	MG	MG	2.33	33
03	Brasilia	DF	DF	6.18	31
04	Campina Grande	PB	CPV	0.62	10
05	Campinas	SP	CAS	4.98	36
06	Cuiaba	MT	CGB	0.78	9
07	Caxias do Sul	RS	CXJ	0.93	5
08	Curitiba	PR	PR	16.66	67
09	Florianopolis	SC	SC	1.32	35
10	Fortaleza	CE	CE	3.21	30
11	Goiania	GO	GYN	1.06	24
12	Lajeado	RS	LAJ	0.14	8
13	Londrina	PR	LDA	2.85	31
14	Manaus	AM	MAO	0.03	9
15	Maringa	PR	MGF	0.42	21
16	Natal	RN	NAT	0.32	13
17	Porto Alegre	RS	RS	36.39	122
18	Recife	PE	PE	1.26	17
19	Rio de Janeiro	RJ	RJ	61.70	73
20	Salvador	BA	BA	2.39	48
21	Sao Carlos	SP	SCA	0.00	3
22	Sao Jose dos Campos	SP	SJC	0.06	13
23	Sao Jose do Rio Preto	SP	SJP	0.03	11
24	Sao Paulo	SP	SP	550.68	686
25	Vitoria	ES	VIX	0.70	21

There are filters in PTT-SP and PTT-PR compromising the BGP table.

**Algorithm 1** Script to collect raw data from IXPs.

**Input:** `code[ixp1, ixp2, ixp3, ... , ixp25]` in

*LOOP Process*

- 1: **for**  $i = code[1]$  **to**  $code[25]$  **do**
- 2:   telnet `lg.code[i].ptt.br`
- 3:    $code[i]$ -IPv4.txt  $\leftarrow$  sh ip bgp
- 4:    $code[i]$ -IPv6.txt  $\leftarrow$  sh ipv6 bgp
- 5:    $code[i]$ -Paths.txt  $\leftarrow$  sh ip bgp paths
- 6:    $code[i]$ -Comm.txt  $\leftarrow$  sh ip bgp community
- 7:   quit telnet session
- 8: **end for**

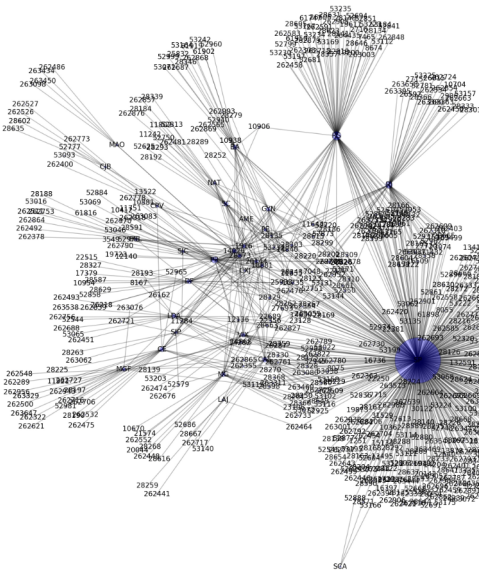


Fig. 3. Graph of the Brazilian IXP ecosystem (IX.br), where each central dot refers to one of the 25 Brazilian IXPs. Notice the existence of ASes that are simultaneously peered at multiple IXPs.

C. Generation of AS-level Graphs

Our framework includes two different tools for the job of generating and analyzing the AS-level graphs based on the input BGP data: (1) NetworkX<sup>10</sup> software for complex networks, and (2) Neo4j<sup>11</sup> graph-oriented database. Both tools are used to build the AS-level connectivity graphs of each IXP and a nation-wide graph based on interconnecting all IXPs through their common AS members. The inputs to both tools are the adjacency matrices generated from the files extracted from each IXP LG. In the resulting graphs, nodes (vertices) are ASes of BGP AS-PATH attributes and edges represent the inter-domain BGP sessions. When observing the nation-wide IXP landscape in the unified graph (see Fig. 3), a second type of node is introduced to represent the IXPs themselves.

Python-based algorithms for NetworkX and Java-based programs for Neo4j were developed to generate the graph and compute a number of properties detailed in the upcoming section. As an informative note, the process of all data import and graph generation takes around 2 hours for each IXP using an Intel Core I7-4790 3.6 GHz with 16 GB RAM. We reckon that the developed software is not optimized for performance and this time could be highly improved. However, once the input data is processed and the generated graphs are loaded in memory, both NetworkX and Neo4j allow running graph queries in a simple fashion, yielding results in very short times.

Our extensive analysis upfront is mainly based on the above-mentioned graphs. In addition to these datasets, all complementary data and code (e.g., algorithms, spreadsheets, gnuplot graphs, etc.) for all 25 IXPs and the unified, Brazil-wide graph (see Fig. 3), are part of the public repository.

D. Limitations

Equally important to the analysis effort presented in this paper is to recognize some limitations of research work when building AS-level topologies [21]. Although we are confident about our methodology and results over the dataset we are sharing, it is important to highlight that information collected from public servers do not represent the totality of traffic exchange between several Brazilian ASes, but only a fraction of everything that can be publicly observed. In terms of peering at IXPs, **only multilateral agreements are mirrored on public servers** and can be observed through LGs, while bilateral agreements cannot be observed as they are directly established between two ASes in a private fashion.

Another relevant observation is that public LGs should mirror all routes inside IXPs RS(s), but it is impossible to assure that observations from public access do not suffer **additional filtering by IXPs operators**. For example, during our analyses it became clear that BGP tables extracted from both PTT-SP and PTT-PR were incomplete by comparing them with the remaining IXPs –an information later confirmed by NIC.br representatives on the grounds of scalability issues. This limitation explains a higher concentration of routes with lower depths and the low percentage of density found on our analysis of PTT-SP, while the remaining IXPs present higher concentration of routes with higher depths and higher percentage of density.

One piece of information included in our shared dataset is the summary of **BGP communities** used at Brazilian IXPs. However, we could not use this information to infer hidden links following the methodology proposed in [16] because NIC.br does not apply any standardized codification system for the use of BGP communities in Brazilian PTTMetro IXPs.

Last but not least, we are aware that our studies are centered around the control plane information (BGP) and should be complemented with **data plane measurements** to enrich the insights on the Brazilian IXP ecosystem.

IV. ANALYSES & RESULTS DISCUSSION

In this section, we analyze the results from the datasets and discuss the main findings from the following studies: (a) classification of IXP members; (b) density of peering; (c) connectivity degree; (d) path depth; (e) traffic engineering practices with AS-Prepend; (f) IPv4 vs IPv6 prefix announcements, and (g) AS-level k-clique communities. We present individual results from four representative IXPs: medium (PTT-MG), medium-to-large (PTT-RJ), large (PTT-RS), and small IXP (PTT-DF), located in the capital, Brasilia. However, results for every IXP are available in the public repository.

The tables ahead include a column called “Brazil” with the average and confidence intervals of the results considering all IX.br IXPs. The high values of the Brazil-wide standard deviation confirm the diversity of ASes and IXPs in the national landscape. The size of an IXP ends up being a relevant factor for many of the observed metrics. For example, for smaller IXPs with fewer members, it is natural to find a higher connectivity degree (density) because of the reduced amount of possible combinations. However, this characteristic cannot be assumed as an exception-free rule.

<sup>10</sup><https://networkx.github.io/>

<sup>11</sup><http://neo4j.com/>

One noteworthy observation is the validity of the results of two IXPs, namely PTT-SP and PTT-PR. BGP data collected from both IXP LGs revealed that filters are being applied to the exported routing tables, a fact confirmed by IX.br representatives due to performance and scalability issues of the LG servers in operation. For this reason, most of the analyses do not include these IXPs.

#### A. Members Classification: Who is who?

A first effort to organize our dissection was manually classifying all 1,142 ASes registered at IX.br. Actually there are 715 unique members registered at IX.br, but there are 1,142 ASes considering the overlap of members peered at multiple IXPs<sup>12</sup>. Our “ground truth” attempt to classify the type of ASes present at IXPs is relevant for an accurate view on the current profile of the members interested in peering in every region of Brazil. The classification task was executed following a manual, “divided and conquer” approach by members of our research group and included individual cross-validation actions. In addition to *whois* services of both NIC.br (Brazil) and LACNIC (Latin America), content from the AS Web sites was used to sort each AS into the categories presented in Table III. Again, the complete dataset with the individual classification of all ASes can be found in our public repository.

**Access Providers dominate.** Looking at the profile results in Table III, despite some variations in the percentile numbers, it is clear that the majority of members peering at IXPs are access providers of local coverage. This is an expected result given the economic incentives of access providers to exchange the maximum amount of traffic as possible through multilateral agreements at IXPs, reducing thereby the transit costs of upstream customer-provider links sold by large operators. The increasing presence of smaller access providers at IXPs directly impacts the ISP prices applied to their downstream customers in a local scenario of competition between multiple access providers. Like in most of development countries, the average quality of Internet connectivity in Brazil is still low compared to developed nations. The public IXP initiative however is contributing to revert this situation by keeping traffic regionalized and reducing the distance between endpoints, and may be more importantly providing an incentive-rich environment for innovation and healthy IP connectivity market practices that are motivating ISPs, mostly access providers, to extend their reach to far locations with poor connectivity options. Without the cost-attractive, shared infrastructure of the IXPs, access providers would need to rely on transit providers, resulting in higher costs and fewer competition in the access provider arena – arguably the most interested parties in open peering.

**In the capital things are different.** The only exception to the dominance of access providers happens at PTT-DF IXP where the presence of government and public organizations is high, a regional particularity at the federal capital of Brazil. Among the PTT-DF members we can highlight the Federal Senate,

Federal Police, Federal Service of Data Processing (Serpro), the Information and Technology Company of Social Security (Dataprev), Telebras (a state telecommunications company), and others.

**Few but heavy Content Providers.** A relative low participation of content providers was observed at IXPs of different Brazilian regions, such as newspapers, magazines, radio and television stations, etc. The majority of content providers are companies that operate Content Delivery Networks (CDN) and are known to be responsible for a large fraction of the traffic. This result highlights the fact that few Brazilian content providers are exploring the benefits of IXP peering due to cost savings and reduced hop distance to eyeball ISPs. We recognize as a plausible reason the common practice of content providers relying on CDN providers to deliver their content closer to the users in the context of a wider geographical span (including internationally), as opposed to IXPs that bring more localized benefits.

**Low presence of private companies.** This fact can be explained by the main motivation of private companies to increase their redundancy through multi-homed connections with IXP members including larger ASes (telcos). These telcos can reach the whole Internet in contrast to IXPs with more restricted reachability towards their local region. While the amount of private companies at IXPs is low, the observed peering density (amount of open peering with all types of ASes) is relatively high, according to the results to be presented in Sec. IV-B.

**Majority incentives lead to the predominance of open peering.** Based on the IX.br records, currently, 97.72% of ASes opt for open peering through multilateral agreement – a high amount in harmony with the spirit and efforts on public, open policies conducted by NIC.br. Only 2.28%, mainly transit providers, choose private peering based on bilateral agreements. The observed ratio is coherent with the fact that large telcos sell transit to local access providers and hence lack economic incentives to openly exchange traffic except those ASes of similar size and nature. Small, regional ISPs are mostly customers that already buy transit somewhere else, since one requirement of IX.br is not relying on their IXP as their only Internet access. Considering the experiences from more mature IXPs in Europe, we can conjecture that the current high fraction of open peering is due to IX.br ASes still being in a “learning” phase seeking peerings with a very open approach.

**Low AS presence at multiple IXPs.** Table IV reports our findings on ASes simultaneously peered at multiple IXPs of IX.br. The broad majority (83.68%) of ASes (759 of 907) are peered to only one IXP. This result can be expected because ASes choose to peer at IXPs mostly to benefit only from local traffic as usually they are not physically present at more than one location. Another relevant detail is that ASes are commonly multi-homed through at least one transit or access ISP to reach the entire Internet, once this is a required policy to become member of IX.br. The set of ASes peering at more than one IXP is mostly composed of access providers that sell services in more than one region. Again, the access providers motivation in exploiting open peering as much as possible

<sup>12</sup>Although IX.br is a national project, we highlight that a member peered at one IXP of IX.br is not connected to the members of other IXPs.

TABLE III  
 PROFILE OF ASES AT BRAZILIAN IXPs AS OF MARCH 25, 2015. THE BRAZIL COLUMN INCLUDES ABSOLUTE AND PERCENTILE NUMBERS WHEREAS INDIVIDUAL IXPs REFLECT ONLY PERCENTILES (ABSOLUTE NUMBERS CAN BE OBSERVED IN TABLE II).

Classification	Brazil (*)	DF	MG	RJ	RS	SP	VIX
<b>1. Internet Provider</b>	743 (65.1% ± 20%)	37.5%	55.9%	51.9%	68.0%	73.1%	75.0%
1.1 Transit Provider	98 (8.6% ± 09%)	20.8%	14.7%	19.2%	5.0%	5.6%	10.0%
1.2 Access Provider	645 (56.5% ± 21%)	16.7%	41.2%	32.7%	63.0%	67.5%	65.0%
<b>2. Services Provider</b>	115 (10.1% ± 07%)	8.3%	8.8%	17.3%	5.0%	12.5%	5.0%
2.1 Content Provider	37 (3.2% ± 06%)	0.0%	2.9%	5.8%	3.0%	4.7%	0.0%
2.2 Hosting Provider	78 (6.8% ± 05%)	8.3%	5.9%	11.5%	2.0%	7.8%	5.0%
<b>3. Public Organization</b>	140 (12.3% ± 21%)	37.5%	20.6%	15.4%	11.0%	4.4%	15.0%
3.1 Public University	20 (1.8% ± 19%)	0.0%	0.0%	0.0%	2.0%	1.1%	0.0%
3.2 Government	100 (8.8% ± 13%)	33.3%	17.6%	13.5%	8.0%	2.2%	15.0%
3.3 Other	20 (1.8% ± 03%)	4.2%	2.9%	1.9%	1.0%	1.1%	0.0%
<b>4. Private Organization</b>	144 (12.6% ± 09%)	16.7%	14.7%	15.4%	16.0%	10.0%	5.0%
4.1 Private University	8 (0.7% ± 03%)	0.0%	2.9%	0.0%	4.0%	0.0%	0.0%
4.2 Private Company	119 (10.4% ± 09%)	16.7%	8.8%	15.4%	10.0%	8.9%	5.0%
4.3 Other	17 (1.5% ± 09%)	0.0%	2.9%	0.0%	2.0%	1.1%	0.0%

(\*) Average of ALL 26 Brazilian IXPs.

is clear: cost savings by avoiding transit links sold by big telecommunication operators. Large ASES can be identified by their simultaneously peering practices in over half of all IXPs (14 to 23) and are predominantly two big telecom operators that we consider transit providers in the national landscape, namely NET and GVT, in addition to public organizations managing the DNS root servers (ICANN), a national Internet performance measurement service (NIC.br), and the National Education and Research Network (RNP).

TABLE IV  
 #ASES PEERED AT MULTIPLE IXPs. (MARCH 25, 2015)

# Peering IXPs	#ASEs	% of ASES	Major Category (*)
01	759	83.68%	All
02	112	12.35%	1.2
03	20	2.21%	1.2
04	2	0.22%	1.2 / 2.2
05	6	0.66%	1.2 / 3.2
06	2	0.22%	1.2
07	1	0.11%	1.1
14	1	0.11%	3.3
15	2	0.22%	1.1 / 3.2
19	1	0.11%	1.1
23	1	0.11%	3.2
<b>Total</b>	907	100%	-

(\*) AS classification presented in Table III

**B. Peering Density: How much peering?**

We consider density of peering as the ratio between the quantity of active BGP connections (peering links) of the  $n$  ASES at an IXP and the sum of all possible peerings ( $n * (n - 1) / 2$ ). The observed peering density (Tab. V) shows wide dispersion in different regions and peering density below

50% points to the potential to expand direct traffic exchange between current IXP members.

We find lower values of peering density in IX.br compared to those presented in previous works regarding other more mature ecosystem of IXPs, such as AMS-IX, DE-CIX, LINX and MSK-IX. While the average percentage of peering density in IX.br is around 40%, more mature IXPs exhibit an average peering density between 79%-95% [16]. Furthermore, when dissecting the AS-level graphs we observe the formation of IXP-enabled k-clique communities (Sec. IV-G), but also lower communities compared to other more mature ecosystem of IXPs.

These two complementary studies confirms our observation that there is a relevant empty space for peering between ASES exchanging traffic at IX.br. One possible explanation to the lower peering density is that IX.br is still young compared to more mature IXPs such as the above-mentioned European IXPs. While IX.br started its first IXP (PTT-SP) in 2004, the peering initiative in Europe started in the early 90s –the oldest IXP of IX.br has only half the lifetime of the largest European IXPs. Another fact could be the relatively limited market that national ASES get through IX.br given that the traffic patterns in Brazil show strong international components according to Alexa’s Ranking for Brazil<sup>13</sup>.

**Who peers with whom? Tell me your AS type...**

To analyze the inter-AS connectivity, we generated a peering matrix for every IXP in the spirit of an adjacency matrix, where  $x$  and  $y$  axis contain all IXP members (ASEs) in a symmetrical fashion. We also considered a unified matrix with all the IXPs to provide a wider view on the nature of peering in the national landscape, i.e., integrating all individual IXPs.

Figure 4 depicts the individual peering matrices of some IXPs. A gray pixel (bit 1) indicates the existence of peering between two ASES while a white pixel (bit 0) indicates the absence of peering. The last column illustrates a scale of

<sup>13</sup><http://www.alexa.com/topsites/countries/BR>

TABLE V  
PEERING DENSITY AT IX.BR (BRAZIL) AND EUROPEAN MORE MATURE IXPs.

Description	Brazil (*)	DF	MG	RJ	RS	DE-CIX	MSK-IX
Peering Links	126	57	79	271	1,952	-	-
Density (%)	44.2% ± 23%	20.7%	34.2%	21.3%	63.4%	79%	95%

(\*) Average of 23 Brazilian IXPs without filters, that is, excluding PTT-PR and PTT-SP.

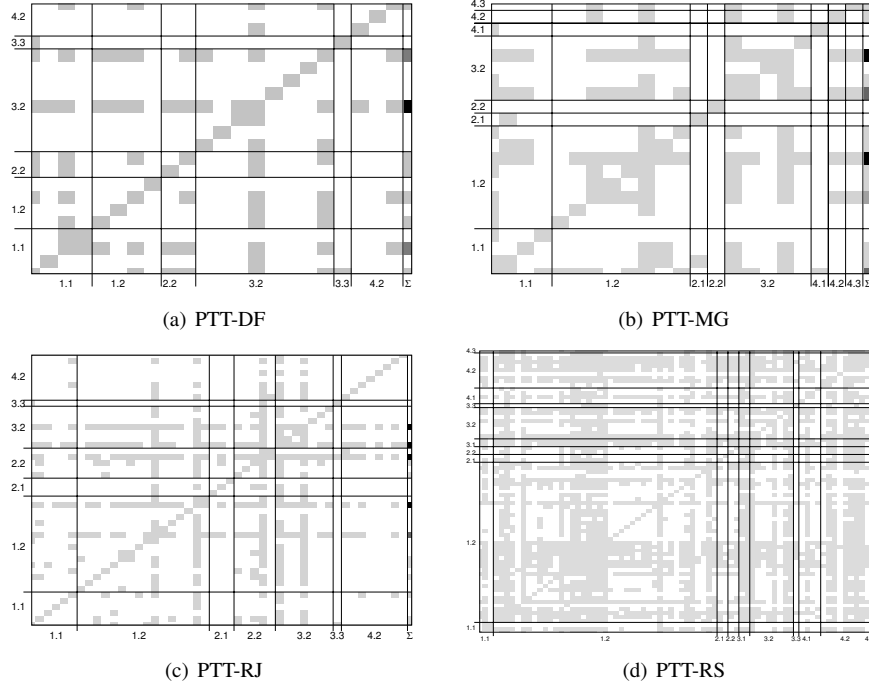


Fig. 4. Peering Matrices indicating the level of connectivity between ASes sorted by category, (a) PTT-DF is a small size IXP, (b) PTT-MG is a medium size IXP, (c) PTT-RJ is a medium-to-large size IXP, and (d) PTT-RS is a large size IXP.




Fig. 5. Peering affinity: matrix of all Brazilian IXPs, where the amount of peering follows a color scale and both axes are grouped in a symmetrical fashion by ASes profiles according to Table III.

the amount of connections between an individual AS with other ASes of each respective category previously presented in Table III, where darker shades mean more connections. The horizontal and vertical lines traversing the graphics are the boundaries between AS categories. We can visually identify

through the long vertical and horizontal lines that some ASes (mostly from access ISP 1.2 cat.) tend to peer more with all types of ASes.

The result of the nation-wide analysis is presented in Figure 5 on a scale from 0 to 2. The color scale is a function of the ratio between the sum of connections (peering between ASes) and the number of vertices of both crossed categories.

In order to quantify (and not just visualize) the amount of peering between different types of ASes, we propose *Peering Affinity (PA)* as a cross-AS-type peering metric defined as follows. Let  $P$  and  $Q$  be sets of ASes such that each set represents a single profile, including the case in which both sets are the same. Let  $c(AS_i, AS_j)$ , such that  $AS_i \in P$  and  $AS_j \in Q$ , be the *connection function*:

$$c(AS_i, AS_j) = \begin{cases} 1 & \text{if } AS_i \text{ and } AS_j \text{ are peers} \\ 0 & \text{otherwise} \end{cases}$$

Then, the *peering affinity* function in respect to  $P$  and  $Q$ ,  $\mathbb{P}A(P, Q)$ , is defined as:

$$\mathbb{P}A(P, Q) = \frac{\sum_{AS_i \in P} \sum_{AS_j \in Q} c(AS_i, AS_j)}{|P| + |Q|}$$

We opt to divide by the sum of vertices resulting in a scale from 0 to 2 instead of dividing by their product because it



returns a more convenient scale to highlight the differences in peering degree. It is possible to have results higher than 2 for PA, but we considered 2 as the highest value of our scale to normalize the results. The majority of PA results are equal or below 2 and very few results of PA are above 2. The numerical results of PA are publicly available in our repository <sup>14</sup>.

Taking as example the peering affinity between members of categories 1.2 (Access Providers) and 2.1 (Content Providers), the amount of connections between all peered ASes of categories 1.2 and 2.1 totals 98, divided by the number of vertices of both categories (236)<sup>15</sup>, returns 0.42 as the cross-AS-type peering affinity metric.

Figure 5 presents the result of the nation-wide analysis regarding peering affinity with the color scale being a function of the ratio between the sum of connections (peering between ASes) and the number of vertices of both crossed categories.

We can observe a relatively high density of peering between ISPs, either transit or access providers. We also observe high density between public organizations, more specifically from the government. The availability of PTT-DF in the federal capital is certainly an enabler to the increased connectivity between many government agencies. Despite such spikes of high density, there is relatively low peering affinity between among the remaining AS categories, which points to the room for increasing the peering density in the national public IXP landscape.

To the best of our knowledge, the peering affinity analysis is the first one that considers a set of peering matrices where ASes are grouped together by their type. When crossing the peering matrices in Fig. 4 with the numbers of Table V, we find a coherent results in support of our methodology. The PA matrix is useful to indicate which specific category of ASes is low connected compared to others, thus helping the adoption of management policies by IX.br aiming to increase the peering density between its members.

It is relevant to mention that choosing multilateral agreement in IX.br means an AS is willing to peer to all other ASes in the same spirit of an open peering policy, but this does not mean that ASes are automatically peered between themselves and explains why the matrices are not full meshes. Most ASes have inbound filters to deal with a selective peering policy with a set of criteria (eg. minimum traffic volume) that peers must meet.

### C. Vertice Degree: How many peers?

Another effort of our dissection was finding the vertices' degree (both distribution and average values) in each IXP graph. By doing so, we aim at revealing and understanding the behavior of the ASes in terms of the amount of neighbouring peers. In sought of connectivity patterns, we discovered that nodes with higher degrees correspond to the older ASes judging by their AS numbers (ASN).

Figures 6 and 7 show the vertices' degree of the following IXPs graphs: PTT-DF (small size), PTT-MG (medium size)

and PTT-RS (large size). Figure 7 plots the degree distribution for all ASes. Since the amount of vertices is large and diverse on an individual AS granularity, in Figure 6, we sorted the AS numbers in the  $x$  axis in a growing fashion and computed the degree for a group of ASes in bins of size around 8,000 ASN. Each bar representing the average degree of a bin of ASes accompanied by an indicator of its confidence interval at a 95%-confidence level. This approach allows showing that ASes registered for longer time (i.e. with smaller ASN) exhibit a higher average degree, a coherent result considering that vertices with higher degrees commonly correspond to telecommunications operators with more adjacencies because of the nature of their transit business and their longer time in operation.

### D. Depth / Diameter: How far are you?

As advertised prefixes traverse BGP domains, the previous AS number is added to the list of ASes (AS-PATH) already visited with the goal to discard any content that include its own AS number on AS-PATH to avoid loops. By observing the AS numbers that exist on AS-PATH it is possible to find the quantity of hops and which ASes must be traversed by a packet before the source of an advertised prefix can be reached by a member of the IXP.

Figure 8(a) illustrates the study of depth from routes advertised by IXPs members based on the AS-PATH attribute present in BGP routing tables. Depth values equal to 1 mean that the AS-PATH is composed of only one AS, a situation that shall be interpreted as ASes directly connected on IXPs that are effectively advertising their own prefixes (low percentage). The remaining routes with depths higher than 1 are being learned by IXP's members from other ASes, which means these routes are not directly advertised by adjacent members of an IXP. The amount of instances of depth equal to 1 and 2 is very low compared to the remaining route advertisements and hence they are omitted in Figure 8(a). For the sake of a more accurate view on the real AS-level distances, duplicated information was removed. A source of duplicated entries in the BGP table are the AS-Prepend practices, commonly used to achieve some sort of inbound traffic engineering as discussed next in Section IV-E.

Taking PTT-RS as an example, in both Figure 8(a) and Figure 8(b), we observe that approximately 35% (y-axis) of the total routes found on BGP table have 5 ASes and another 35% have 6 ASes in their AS-PATH attribute, that is, around 70% of all routes have a depth in order of 5 or 6. By observing all IXPs together, we found the average depth of all routes advertised at IXPs varying from 4 to 6, remarking that the higher concentration of routes has a depth of 5. Another relevant detail is the value of 8 as the highest depth, that is, the most distant sources of AS announcements reaching IXPs are 8 hops far away, recalling that each hop means an entire AS (not a single router as `traceroute` would reveal). However, it is important to note that, inside BGP tables of an IXP, depths higher than 8 (from 9 to 14) do exist but are inexpressive in the face of the total number of advertised routes. This fact is reinforced by Figure 8(a), where the accumulated depth on the  $x$ -axis shall be read as *until  $n$  hops*.

<sup>14</sup><https://github.com/intrig-unicamp/ixp-ptt-br/blob/master/analysis/Adjacencies/peeringMatrixByProfile/heatmap>

<sup>15</sup>PTT-SP and PTT-PR are out of this sum due to filters in the BGP table.

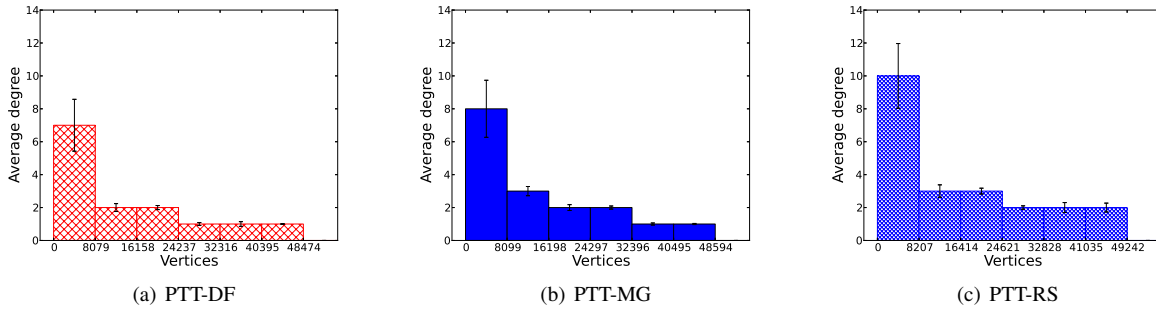


Fig. 6. Average Degree

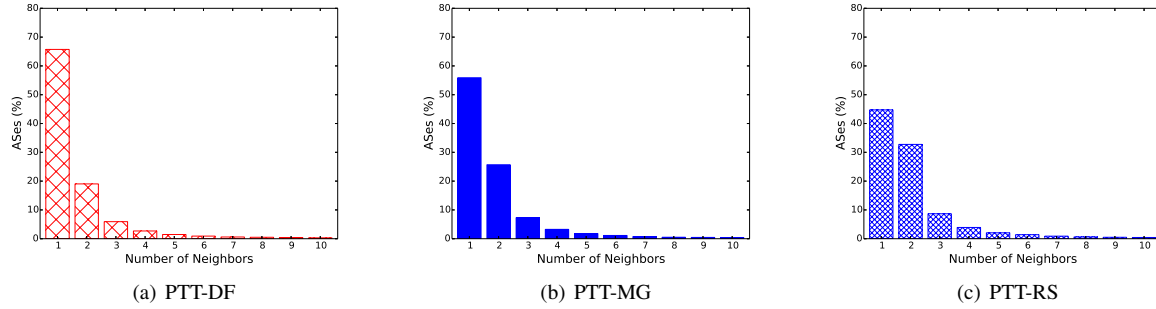


Fig. 7. Degree Distribution

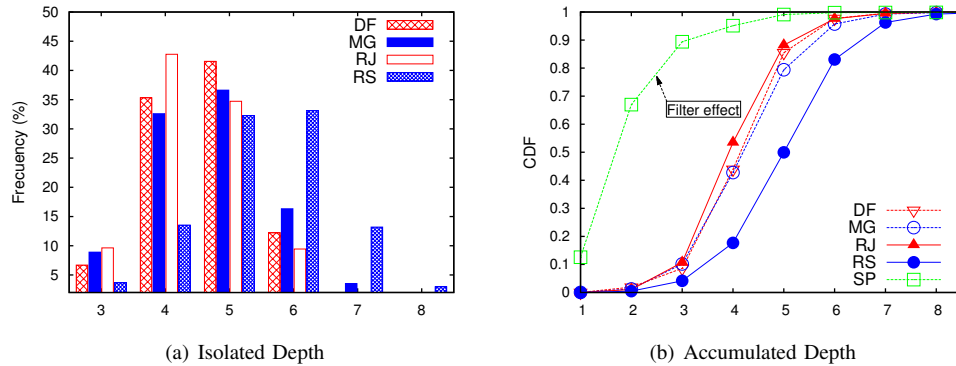


Fig. 8. Depth of AS-PATH

As mentioned earlier, the results of PTT-SP and PTT-PR are *out of the curve* and shall be carefully considered due to the LG BGP table filtering practices. We include PTT-SP in Figure 8(b) precisely to illustrate the effect of route filtering when attempting to carry this kind of Internet measurement research.

To put the results into international perspective, we can compare these numbers with those found in Italy [2], where around 70% of MIX IXP routes have a cumulated depth of 8 and around 40% of them have depth of 4 to 6.

*E. Traffic Engineering with AS-Prepend*

The default behavior of BGP is preferring routes to prefixes with a smaller list of AS-PATH elements, that is, those paths with smaller depth/diameter. Hence, the AS-PATH attribute is commonly used as an indirect mechanism to influence BGP routing policies to select between multiple routes to

the same destination (IP prefix). The “traffic engineering” practice consists of an AS prepending its own AS number more than once to the AS-PATH attribute to turn less attractive the reachability to a given prefix through itself [4]. AS-Prepend is a BGP “knob” for inbound traffic engineering considered harmful because it compromises the integrity of routing information [35]. Other BGP attributes suitable to implement routing policies for traffic engineering purposes include the Community and Multi-Exit Discriminator (MED) attributes. However, we were not able to find meaningful results from the available data.

The results of Table VI reinforce that AS-Prepend is in fact commonly used, as observed from the total number of routes with AS-Prepend. Note that while the second set of rows in the table refers to all ASes of the Internet (be them IXP members or not) observed from IXPs’ graphs, the last set of rows reflects only the prepend practices from AS members directly connected at the IXPs, remarking that not all members

may be advertising routes. We observe a higher amount of AS-Prepend per AS at IXPs, a fact suggesting the needs for traffic engineering in IXP peering links.

When looking at the numbers in Table VI it is worth to recall the difference between two distinct concepts present in the BGP table information: (i) number of routes, and (ii) number of prefixes. While the full BGP table currently features around 512,000 prefixes, it is usual to find BGP tables at IXPs with millions of entries due to the advertisement of multiple routes towards the same prefix.

#### F. IPv4 vs IPv6 Prefixes and Routes

We now turn our attention to the IPv4 and IPv6 prefix advertisement practices in the IX.br ecosystem.

Best practices toward conserving the IP address space and limiting the growth rate of global Internet routing tables encourage ASes the advertisement of aggregated prefixes [14], for example of sizes /15 and /21 from large ISPs, and small ISPs, respectively. Although the Best Current Practices (BCP) [14] are recommended by most of the Regional Internet Registries (RIR), there is no actual hard requirement regarding the size of advertisements. A number of situations are known to negatively impact the effectiveness of aggregation (e.g., multi-homing, renumbering, traffic engineering through advertisements of more specific routes), contributing to sub-optimal levels of aggregation observed in practice.

**Similar practices for IPv4 prefix announcements.** When looking into the size of prefixes for each route inside the BGP tables of all Brazilian IXPs (Figures 9(a) and 9(b)), we observe an homogeneous behavior in most of the Brazilian IXPs. This fact is reinforced by the low standard deviation values. Around 50% of all advertised routes are of the size /24, with near 20% of both /22 and /23. The remaining 30% are mostly between /16 and /21. All other prefixes together, from /8 to /15 and from /25 to /32, represent less than 2.5%. Even though it is well known that prefixes higher than /24 (/25 to /32) are generally filtered by providers, actually we found a small percentage of these prefixes, probably advertised by ASes with less technical BGP staff.

**As well as for IPv6 (still well aggregated) prefixes.** According to RFC 5375 [10] on IPv6 addressing considerations, larger IPv6 addresses can lead to larger routing tables unless network designers actively pursue aggregation. The national policy suggested by NIC.br is that from the /16 block allocated by LACNIC (prefix 2801::/16), the following assignment plan should be followed: (i) /32 for ISPs, (ii) /48 for enterprises, and (iii) /56 for residential users. As in IPv4, similar practices and homogeneous behavior can be observed in most IXPs when looking into IPv6 address family of BGP tables.

As shown in Figures 9(c) and 9(d), the majority of the advertised prefixes are /32 (approx. 35% in average) and /48 (approx. 42% in average). We observe some /29 (approx. 2.5% in average), /40 (approx. 5% in average), and /44 (approx. 4% in average). Through these results we can confirm that the presence of unaggregated IPv6 prefixes is still inexpressive—most ISPs are adequately advertising their /32 while companies are advertising their /48. However, we recognize that the

adoption of IPv6 is still too small compared to the amount of IPv4 prefixes.

**Low adoption of IPv6.** Comparing between IPv4 and IPv6 prefixes and routes, only 2.95% of the prefixes (6.95% of the routes) are IPv6 against the remaining 97.05% (93.05% of the routes) of IPv4. These numbers clarify that the adoption of IPv6 in Brazil is still very slow despite all the efforts of NIC.br, specifically by its working group called IPv6.br. This result was expected given the early stage of IPv6 adoption worldwide, i.e. the result is in line with global adoption rates [12].<sup>16</sup>

#### G. AS-level Network Communities

The concept of network community has been used as a building block for a number of methods to analyze the Internet AS-level topology. In recent years, related work has highlighted the relevance of detecting and interpreting Internet communities (e.g., [17, 18, 29]). Recent studies [29] present evidence that correlate Internet structural properties and IXPs. For instance, the proliferation of IXPs has triggered the formation of denser sub-graphs. In addition, IXPs have been shown to influence the formation of communities [18]. For example, bigger communities can be found composed mainly by ASes connected to a common IXP.

Motivated by the aforementioned importance of network communities, we enrich our analysis of the Brazilian IXP ecosystem by building a subgraph induced by all IXP members. We then use the resulting subgraph to detect communities and provide results on their fundamental properties, namely size and density.

**Detecting Communities.** Recognized methods in the literature on extracting Internet communities include  $k$ -clique,  $k$ -dense and  $k$ -core [29]. The common feature of these three methods is the capacity to detect dense sets of ASes connected to each other. We focus on the  $k$ -clique communities method because of its particular feature of exposing overlapping communities [30], allowing the separation of those communities with a common amount of nodes.

A  $k$ -clique is a complete sub-graph with  $k$  nodes. Two  $k$ -cliques are adjacent if they have  $k - 1$  nodes in common. A  $k$ -clique community, or  $community(k)$ , is the maximal set formed by the union of all adjacent  $k$ -cliques. Another property is that each  $k$ -clique community has one and only one community ( $k-1$ )-clique such that  $community(k)$  is a sub-graph of  $community(k - 1)$  [18], that is, smaller values of  $k$  allow to identify more communities than would be possible to identify with higher values of  $k$ .

**Characterizing the observed Communities.** As mentioned earlier, we extracted communities from a subgraph induced by all the members of the Brazilian IXPs. For each value of  $k$ , we considered the following metrics:

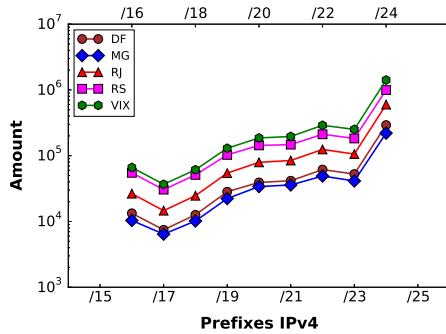
- (i) *Largest Community*: the community with a maximum number of ASes;
- (ii) *Density*: ratio between existing connections and possible connections. This metric only considers connections within the community;

<sup>16</sup><https://www.google.com/intl/en/ipv6/statistics.html>

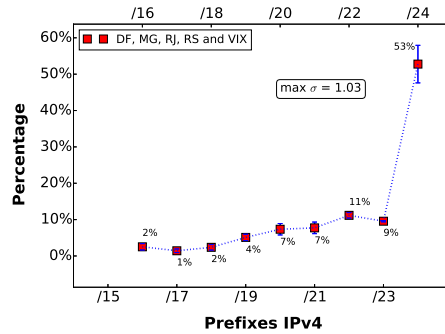
TABLE VI  
STATISTICS ON AS-PREPEND VIEWED THROUGH IXPS.

Metric Description	Brazil (*)	DF	MG	RJ	RS	VIX
Routes	832,989	559,159	434,264	1,150,905	1,947,453	2,663,751
Routes with AS-Prepend	295,909	127,184	245,129	294,663	1,710,070	623,965
<b>AS-Prepend X Routes (%)</b>	<b>30.8% ± 22%</b>	<b>22.7%</b>	<b>56.4%</b>	<b>25.6%</b>	<b>87.8%</b>	<b>23.4%</b>
ASes at Graph	43,333	47,176	46,939	47,632	48,351	47,474
ASes with AS-Prepend	7,305	6,206	8,629	8,890	10,803	9,124
<b>AS-Prepend X ASes (%)</b>	<b>16.1% ± 04%</b>	<b>13.2%</b>	<b>18.4%</b>	<b>18.7%</b>	<b>22.3%</b>	<b>19.2%</b>
Members Advertising	18	24	22	51	79	18
Members Advertising with AS-Prepend	6	7	6	19	36	5
<b>AS-Prepend X Members (%)</b>	<b>22.5% ± 19%</b>	<b>29.2%</b>	<b>27.3%</b>	<b>37.3%</b>	<b>45.6%</b>	<b>27.8%</b>

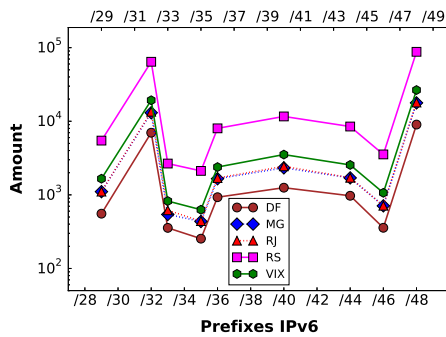
(\*) Average of 23 Brazilian IXPs without filters, that is, excluding PTT-PR and PTT-SP.



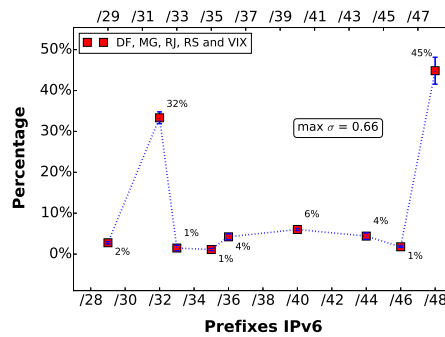
(a) Amount of IPv4 Prefixes



(b) Average Percentage of IPv4 Prefixes



(c) Amount of IPv6 Prefixes



(d) Average Percentage of IPv6 Prefixes

Fig. 9. Amount and distribution of IPv4 and IPv6 prefixes advertised at IX.br. Figures (a) and (b) show that more than 50% of IPv4 prefixes are /24, followed by /23 and /22. In (c) and (d), we observe more than 70% of IPv6 prefixes being /32 (ISPs) and /48 (enterprises).

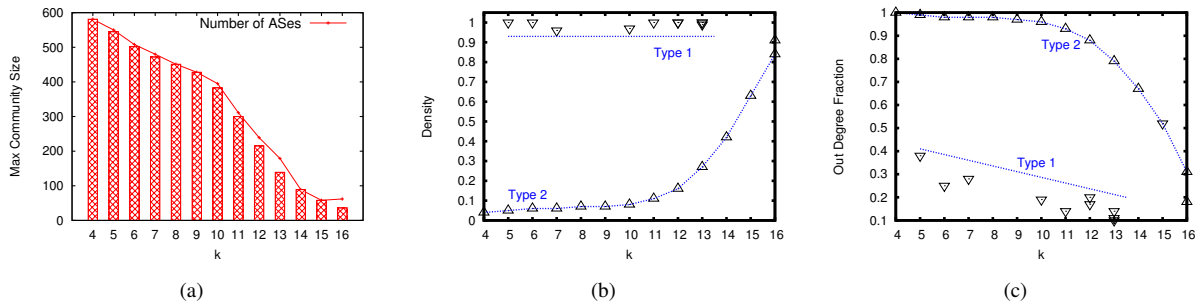


Fig. 10.  $k$ -clique communities in the scope of the IXP members: (a) largest communities, where upper bound line represents the total number of ASes for all the communities found for a given  $k$ ; (b) density; and (c) Out Degree Fraction (ODF).

(iii) *Out Degree Fraction (ODF)*: a node's ODF [18] is defined as the ratio between its degree in the community graph and the overall degree. The ODF value of a community is the average of its nodes' ODF.

Figure 10 presents our findings on the  $k$ -clique communities detected in the IX.br ecosystem. The number of ASes of the largest community (Fig. 10(a)) fluctuates between 581 and 36 with respect of  $k$ -values set to 4 and 16, respectively. The

latter is the biggest  $k$  that results in at least one community.

The upper bound line in Figure 10(a) represents the total number of ASes for all the communities found. There is a clear influence from the size of the largest community since only a few values of  $k$  show a gap between the bars and the line. Note that the gaps are relatively small. This result is a consequence of both the small number of communities (mostly 1 or 2) and the absence of more than one community with a large number of ASes.

In order to differentiate classes of communities, Figures 10(b) and 10(c) have blue dashed lines that point out Type 1 and Type 2 communities. The former represent groups of either high density communities or low ODF communities. A set of communities that contains either an exponential growth or an exponential decay is a set of Type 2 communities. Our results indicate that smaller communities are of Type 1, while Type 2 mostly represents groups of larger communities.

Figure 10(b) shows the density for each community detected. Considering the cases in which there is only a single community, the density is mostly less than 10% ( $k = \{4, 8, 9\}$ ) or limited ( $k = 14$ ), except for the case of  $k = 15$ . This result indicates, on one hand, the existence of small groups of strongly connected ASes that, on the other hand, are loosely connected between them. More connectivity between groups of ASes could be created. For the cases with more than one community, it can be observed that only one of them has low density, which is precisely the community with the largest size in number of ASes (not shown in the figure). An interesting exception of this result is noted for  $k = 16$ , since there are 2 communities with high density.

Figure 10(c) shows the average ODF of the communities. Lower values indicate sets of nodes mostly connected to outside their community. This is the case of smaller communities while bigger communities always have high ODF values, and, thus, are mostly connected to nodes inside the community. Figures 10(b) and 10(c) show opposite results, since high density communities have low ODF values.

**Anatomy of IX.br Communities: Wrap Up.** We summarize our main findings as follows:

- there is always a small number of communities;
- there is either a large or a small number of ASes when comparing multiple communities of a given  $k$ -value;
- smaller communities are usually dense and larger communities are usually sparse;
- the above item has an important exception: the largest value of  $k$  has 2 dense communities;
- smaller communities have small ODF values while larger communities have high ODF values.

The small number of communities highlights tightly connected groups of IXPs' members. Note that many isolated communities would suggest a low level of connectivity between sets of nodes. On the other hand, the induced subgraph of the Brazilian IXP ecosystem does not have a  $k$ -value as high as other IXPs' induced subgraphs such as the London Internet Exchange (LINX) and the German Internet Exchange (DE-CIX) [18]. In particular, LINX had a  $k$ -value greater than 28 (data computed 4 years ago) while the Brazilian ecosystem is bounded today by a value of only 16 according to our results.

Thus, there is still room to improve connectivity in order to draw more ASes into the existing communities.

## V. RELATED WORK

Khan *et al.* [22] highlight the less-known capabilities of LGs servers to construct the Internet AS topology by collecting raw data from 245 LGs servers in 110 countries worldwide. The authors found 11,000 AS links and 686 ASes previously not found in other BGP, traceroute, and IRR based AS topologies. The authors highlight that LG-based methods to collect raw data is less error prone than traditional traceroute-based approaches.

A work based on 1.5 years of active layer-3 probing is presented by Durairajan *et al.* [13], where one of the targets were IXPs because of their fixed geographic location. The goal was to identify the Internet infrastructure and the authors demonstrate the capability of their method using LG vantage points distributed throughout the Internet.

Richter *et al.* [33] explain that the use of route servers at IXPs has greatly simplified inter-domain routing facilitating members to peer with many other ASes present at the same IXP. Similar to our work, they report an empirical analysis based on a collection of IXP-provided datasets from two European IXPs. The main difference of their dataset and ours is that the European IXPs provided them with both control plane (BGP data from RS) and data plane (sFlow measurements).

A temporal analysis by Ager *et al.* [1], based on nine months of registries collected from one of the biggest world's IXP, includes a member classification pointing to their diversity. One discovery of this work was that the amount of peering of a single IXP in 2012 exceeded the total quantity of peering between ASes of the entire Internet in 2010.

The authors of a 2012 study focused in the Slovak Internet eXchange (SIX) [32] argue that there are many efforts and measurements intended to comprehend the Internet as a whole, but little is known about the Internet in the local context of a specific region. The authors classified the evolution of the ecosystem of providers connected at SIX and also the traffic profile distributed by its members.

A method to discover hidden peering on IXPs through the mining of BGP community attribute can be found in [16], the same problem discussed in [9]. The method is particularly interesting because an AS can restrict its advertised routes, making incomplete the list of AS-PATH (BGP paths) of IXP's route servers. By applying this technique, the authors could infer 206,000 point-to-point links from 13 IXPs in Europe, four times more links than the amount of links directly observed inside BGP tables.

According to Giotsas *et al.* [15], the relationships between ASes are traditionally classified as: (i) transit (provider-to-customer), (ii) peering (peer-to-peer) and (iii) sibling (similar domains), however there are advanced settings that imply in hybrid or complex relations. To better understand these complex relationships, the CAIDA [5] algorithm was improved to be capable of automatically parses BGP tables, outputs from traceroute and geolocation data. By running this algorithm it

was possible to observe that 4.5% of a universe of 90,272 relationships first classified as transit (provider-to-customer) were hybrid or complex relationships.

Lodhi *et al.* [26] recent work on data mining PeeringDB, one of the few public sources of information about peering that consists of an online tool where Internet members contribute by inserting information regarding its policies, traffic amount, and geographic location. In this study the authors found consistent correlations between PeeringDB data compared to measurements from BGP prefixes advertised in the Internet.

Another recent in-depth analysis of PeeringDB can be found in [36], which contains relevant information regarding the public and private peering ecosystems in the five regions of the world, including how IXP facilities are distributed across the five regions. We previously mentioned in our paper that we found PeeringDB data was outdated regarding Brazil, an information coherent considering that Kloti *et al.* [23] bring a first cross-comparison of three well-known publicly available IXP databases and found 40.2% more IXPs and 66.3% more ASes than PeeringDB.

By looking at graphs of ASes connected on IXPs, [18] defines as major communities those with the maximum value of  $k$  and conclude that these communities are typically big in size and have low density of connection, trending to connect to external nodes outside the community. A similar study using  $k$ -dense communities is done in [29], which conclusion shows that communities with the maximum  $k$ -value trend to be composed of tier-2 providers and content providers.

A recent study closest to ours characterized the nature of interdomain Internet connectivity in Africa [19], specifically focused on JINX (in Johannesburg) and KIXP (in Nairobi), two major IXPs in Africa. The authors measured the presence of local ISPs at various African IXPs and which of them chose to interconnect at these exchanges. An interesting result of this work was finding that 66.8% of the paths between residential users and Google leave the continent, mainly because local ISPs are not present at these IXPs or because they are not peered between each other. The study presented individual peering matrices for JINX and KIXP that inspired our similar efforts in Section IV-B leading to the proposed peering affinity metric applied to the national-wide peering matrix grouping ASes by their type.

## VI. CONCLUSION AND FUTURE WORK

This paper presents an in-depth analysis towards comprehending the ecosystem of the largest set of public IXPs which happen to be in Brazil. By compiling information collected from public servers (and also missing data provided by the IXP operators), AS-level connectivity graphs of all public Brazilian IXPs were used to dissect the underlying anatomy and identify several features of the peering ecosystem.

The developed framework allows scalable and rich analytic studies around the Brazilian IXP connectivity graphs, including studies moving beyond traditional sets of individual peering matrices to include a single national-wide AS graph. Sorting ASes by their category, we leverage a novel metric called *peering affinity* to measure the degree of peering

between different types of ASes. We found lower values of peering density in the IX.br ecosystem compared to those observed in more mature IXPs in Europe. Furthermore, when dissecting the AS-level graphs we observe the formation of IXP-enabled  $k$ -clique communities, which also exhibit lower community sizes compared (again) to more mature ecosystems. These complementary studies point to the potential to increase peering between ASes exchanging traffic at IX.br.

Our ongoing work includes a temporal analysis based on collecting datasets over a longer period, providing different snapshots that will allow us to comprehend the dynamic aspects and evolution of the Brazilian IX.br ecosystem, besides enabling the creation of a predictive model. Regarding the study of the AS-level network communities, an interesting research direction would be to characterize the level of *dependency* of adjacent  $k$ -cliques on possibly small sets of nodes. It is likely that such sets exist because there is a small number of communities. This line could motivate the creation of new regional IXPs to distribute the interconnection of communities. Finally, the collected dataset is subsidizing our research group work on Software-Defined Networking (SDN) [24] peering towards allowing new inter-domain exchange applications currently not feasible with BGP [20], such as content-based peering or application-based peering, as well as new approaches to inter-domain policies towards more secure and flexible traffic engineering techniques, including the delivery of Application-Layer Traffic Optimization (ALTO) information services [31].

## ACKNOWLEDGMENTS

This research was supported by the Innovation Center of Ericsson, Brazil. The authors would like to thank Antonio Moreiras (NIC.br) and Antonio Galvao de Rezende Filho (IX.br) for their help in accessing all LGs. We also thank Giovanni Comarella and Steve Uhlig for their insightful comments on earlier versions of the paper.

## REFERENCES

- [1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. *Anatomy of a Large European IXP*. 2012. SIGCOMM'12, August 13-17, Helsinki, Finland.
- [2] M. D. Bartolomeo, G. D. Battista, R. di Lallo, and C. Squarcella. *Is It Really Worth to Peer at IXPs? A Comparative Study*. 2015. 20th IEEE ISCC. July 06-09, 2015. Larnaca, Cyprus.
- [3] S. H. B. Brito, M. A. Santos, R. dos Reis Fontes, D. A. L. Perez, and C. E. Rothenberg. Dissecting the largest national ecosystem of public internet exchange points in brazil. In *International Conference on Passive and Active Network Measurement*, pages 333–345. Springer International Publishing, 2016. doi: 10.1007/978-3-319-30505-9\_25.
- [4] M. Caesar and J. Rexford. *BGP Routing Policies in ISP Networks*. 2005. Network. IEEE. Volume 19, Issue 6. Pages 5-11. ISSN 0890-8044. doi: 10.1109/mnet.2005.1541715.
- [5] CAIDA. *CAIDA - Center for Applied Internet Data Analysis*. 2014. www.caida.org.

- [6] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. *Remote Peering: More Peering without Internet Flattening*. 2014. CoNEXT'14. Dec 2-5, 2014. Sidney. Australia.
- [7] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. *There is More to IXPs than Meets the Eye*. 2013. CCR ACM SIGCOMM. Volume 43, Number 5. October. doi: 10.1145/2541468.2541473.
- [8] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. *On the Importance of Internet eXchange Points for Today's Internet Ecosystem*. 2013. Computer Communications Review (CCR). ACM SIGCOMM.
- [9] 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*. 2009. CoNEXT'09, December 01-04, Rome, Italy.
- [10] G. V. de Veldi, C. Popoviciu, T. Chown, O. Bonness, and C. Hahn. *IPv6 Unicast Address Assignment Considerations*. 2008. IETF. RFC 5375.
- [11] A. Dhamdhare and C. Dovrolis. *The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh*. 2010. CoNEXT'10, Nov 30 - Dec 3, Philadelphia, USA.
- [12] A. Dhamdhare, M. Luckie, B. Huffaker, K. Claffy, A. Elmokashifi, and E. Aben. *Measuring the Deployment of IPv6: Topology, Routing and Performance*. 2012. IMC'12, November 14-16, Boston, Massachusetts. USA.
- [13] R. Durairajan, J. Sommers, and P. Barford. *Layer1-Informed Internet Topology Measurement*. 2014. IMC'14. November 05-07, 2014. Vancouver, BC, Canada.
- [14] V. Fuller and T. Li. *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*. 2006. IETF. RFC 4632.
- [15] V. Giotsas, M. Luckie, B. Huffaker, and K. Claffy. *Inferring Complex AS Relationships*. 2014. IMC'14. November 6-7. Vancouver, BC, Canada.
- [16] V. Giotsas, S. Zhou, M. Luckie, and K. Claffy. *Inferring Multilateral Peering*. 2013. CoNEXT'13, December 9-12. Santa Barbara, California, USA.
- [17] E. Gregori, L. Lenzini, and S. Mainardi. Parallel k-clique community detection on large-scale networks. *Parallel and Distributed Systems, IEEE Transactions on*, 24(8):1651-1660, Aug 2013. doi: 10.1109/tpds.2012.229.
- [18] E. Gregori, L. Lenzini, and C. Orsini. k-clique communities in the internet as-level topology graph. In *Distributed Computing Systems Workshops (ICDCSW), 2011 31st International Conference on*, pages 134-139, June 2011.
- [19] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Candalro, and E. Katz-Bassett. Peering at the internet's frontier: A first look at isp interconnectivity in africa. In *Passive and Active Measurement, PAM 2014, 15th International Conference on*, pages 204-213, March 2014.
- [20] A. Gupta, L. Vanbever, M. Shabaz, S. P. Donovan, B. Schlinker, N. Feamster, et al. *SDX - A Software Defined Internet Exchange*. 2014. SIGCOMM'14. Chicago, USA.
- [21] H. Haddadi and O. Bonaventure. *Recent Advances in Networking. Chapter 1: Internet Topology Research Redux*. 2013. ACM SIGCOMM eBook. Volume 1. August.
- [22] A. Khan, T. T. Kwon, H. chul Kim, and Y. Choi. *AS-level Topology Collection Through Looking Glass Servers*. 2013. IMC'13, October 23-25, 2013. Barcelona, Spain.
- [23] R. Kloti, B. Ager, V. Kotronis, G. Nomikos, and X. Dimitropoulos. *A Comparative Look into Public IXP Datasets*. 2016. ACM SIGCOMM CCR. Volume 46, Number 1, January 2016. doi: 10.1145/2875951.2875955.
- [24] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1):14-76, 2015. doi: 10.1109/jproc.2014.2371999.
- [25] C. Labovitz, S. Iekel-Johson, D. McPherson, J. Oberheide, and F. Jahanian. *Internet Inter-Domain Traffic*. 2010. SIGCOMM'10, Aug 30 - Sep 3, New Delhi, India.
- [26] A. Lodhi, N. Larson, A. Dhamdhare, C. Dovrolis, and K. Claffy. *Using PeeringDB to Understand the Peering Ecosystem*. 2014. CCR ACM SIGCOMM. April.
- [27] M. Luckie, B. Huffaker, A. Dhamdhare, V. Giotsas, and K. Claffy. *AS Relationships, Customer Cones, and Validation*. 2013. IMC'13, October 23-25, 2013. Barcelona, Spain.
- [28] W. B. Norton. *The Internet Peering Playbook: Connecting to the Core of the Internet*. 2014. DrPeering Press. 2014 Edition. USA.
- [29] C. Orsini, E. Gregori, L. Lenzini, and D. Krioukov. Evolution of the internet k-dense structure. *IEEE/ACM Transactions on Networking*, 22(6):1769-1780, Dec 2014. doi: 10.1109/tnet.2013.2282756.
- [30] G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814-818, 2005. doi: 10.1038/nature03607.
- [31] D. A. L. Perez, S. H. B. Brito, R. dos Reis Fontes, and C. E. Rothenberg. Delivering application-layer traffic optimization services based on public routing data at internet exchange points. 2016. 34rd Brazilian Symposium on Computer Networks and Distributed Systems.
- [32] J. C. C. Restrepo and R. Stanojevic. *A History of an Internet eXchange Point*. 2012. CCR ACM SIGCOMM. Volume 42, Number 2, April. doi: 10.1145/2185376.2185384.
- [33] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. *Peering at Peerings: On the Role of IXP Route Servers*. 2014. IMC'14. November 05-07, 2014. Vancouver, BC, Canada.
- [34] B. Schlinker, K. Zarifis, I. Cunha, N. Feamster, and E. Katz-Bassett. *PEERING: An AS for US*. 2014. HotNets'14. October 27-28. Los Angeles, CA, USA.
- [35] P. Smith. *BGP Best Practices*. 2006. RIPE NCC Regional Meeting. 14-15 November. Manama, Bahrain. <https://www.ripe.net/ripe/meetings/regional-meetings/manama-2006/BGPBCP.pdf>.

- [36] R. P. Venkata and L. Ruan. *A Region-Centric Analysis of the Internet Peering*. 2016. Computer Science Technical Reports. Paper 381. June 2016.



**Samuel Henrique Bucke Brito** is a PhD student at School of Electrical and Computer Engineering (FEEC) at University of Campinas (UNICAMP). He holds a Master Degree in Electrical Engineering (Telecommunications) and a BSc in Computer Science. Recognized as IPv6 Forum Certified Trainer and CCAI/ITQ (Cisco Certified Networking Academy Instructor Trainer). Author of two best-sellers in Brazil, including one recognized as IPv6 Forum Officially Certified Guide. Currently holds the position of Assistant Professor and Academic


Coordinator at Methodist University of Piracicaba (UNIMEP).



**Hirley Dayan Lourenco da Silva** is a telecommunication engineer undergraduate at INATEL, Brazil, with 15 years of experience on delivering telecommunication solutions to many operators in many countries around the world. He works as the Innovation Technical Coordinator for Ericsson and also as a team member of Computing and Industrial Automation Department at UNICAMP, in Brazil, working on his graduation research on Cloud and SDN-related topics.



**Mateus Augusto Silva Santos** received his Ph.D. from Universidade de So Paulo (USP) in 2014. From 2013 to 2014 he was a research scholar with the Inter-Networking Research Group at UC Santa Cruz. His research interests are in software-defined networking, network security and wireless networks. He is currently a postdoctoral researcher with University of Campinas (UNICAMP) and a researcher at Ericsson.



**Ramon dos Reis Fontes** is Professor at Federal Institute of Education, Science and Technology of Bahia (IFBA) in Brazil and phd Student in Electrical and Computer Engineering from State University of Campinas, SP, Brazil



**Christian Esteve Rothenberg** is an Assistant Professor in the Faculty of Electrical and Computer Engineering at University of Campinas (UNICAMP), where he received his Ph.D. (2010) and currently leads the the INTRIG (Information & Networking Technologies Research & Innovation Group). From 2010 to 2013, he worked as Senior Research Scientist in the areas of IP systems and networking at CPqD R&D Center in Telecommunications (Campinas, Brazil). Christian holds the Telecommunication Engineering degree from the Technical University of Madrid (ETSIT - UPM), Spain, and the M.Sc. (Dipl. Ing.) degree in Electrical Engineering and Information Technology from the Darmstadt University of Technology (TUD), Germany, 2006. Christian holds two international patents and has over 80 publications including scientific journals and top-tier networking conferences such as SIGCOMM and INFOCOM.



**Danny Alex Lachos Perez** is a PhD student in the Faculty of Electrical and Computer Engineering at University of Campinas (UNICAMP). He is involved as a researcher in the Advanced Peering with Software-Defined Knowledge Plane project supported by the Innovation Center of Ericsson Telecomunicacoes S.A (Brazil). He research interests are: Software-Defined Networking, Application-Layer Traffic Optimization, Software Engineering and Graph Database. He holds the Computer Engineering degree from Pedro Ruiz Gallo University,

Per, 2004 and also a Master degree in Computer Engineering from University of Campinas, Brazil, 2016.