Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für
Technische Informatik und
Kommunikationsnetze

Demian Jäger

Analyzing the time dynamics in IXP datasets

Semester Thesis SA-2014-43
October 20, 2014 to February 8, 2015

Advisor: Dr. Bernhard Ager
Co-Advisor: Vasileios Kotronis
Supervisor: Prof. Dr. Bernhard Plattner

Abstract

Internet exchange points (IXPs) interconnect different Autonomous Systems directly with each other, reducing latency and costs by not using an intermediate network provider. Despite existing for a long time, IXPs became only recently a topic of interest to researchers. This thesis contributes to this research by providing an analysis of the time dynamics in an IXP dataset. We analyze daily snapshots of PeeringDB [14], a database which contains information on the IXP ecosystem, from March 2014 to December 2014. It turns out, that most of the traffic exchange happens at a few large IXPs, which grow fast in terms of their capacity. Furthermore, an analysis of the available capacity at IXPs shows, that the growth approximately follows an exponential function. While all types of connected networks grow in terms of their capacity, content providers and content distributors grow a lot faster than others.

Contents

1 Introduction **9**

2 Dataset **11**
2.1 PeeringDB 11
2.1.1 Internet Exchange Point (IXP) 11
2.1.2 Participant 11
2.1.3 Facility 12
2.1.4 Summary 12
2.2 Collection of the Dataset 13
2.3 Artifacts in Dataset 13

3 Analyzing the Time Dynamics of PeeringDB **15**
3.1 Difference Algorithm 15
3.1.1 Difference Data Format and Examples 15
3.1.2 Algorithmic Steps 16
3.1.3 Uniqueness of IDs 16
3.2 Extract Information 17
3.3 Growth Analysis 18

4 Results **19**
4.1 Dynamics of IXPs 19
4.2 Dynamics of Participants 19
4.3 Evolution of the Overall Capacity 23
4.4 Update Frequency of Database 23

5 Conclusion and Future Work **27**
5.1 Conclusion 27
5.2 Future Work 27

List of Figures

3.1 First, difference files are built. Extracting the information of potential interest forms the next step. Eventually, the information is concatenated and shown in a time series analysis. 16
3.2 The CDF of a similarity function shows that elements with identical IDs are very akin. Elements with different IDs have very little similarity. We thus assume that IDs are unique and consistent over multiple snapshots. Here, the snapshot from June 26, 2014 is compared to the one from October 21, 2014. 17

4.1 The biggest outlier which grew by 54% is PTT Sao Paulo. The three outliers with the largest member count at start are AMS-IX, DE-CIX and LINX. 20
4.2 The capacity of IXPs grew in the observation period. At larger IXPs, the growth of capacity is related to their size. The two rightmost points represent AMS-IX and DE-CIX, two of the largest Internet exchange points in the world. 20
4.3 Comparison of the member size history of AMS-IX and PTT Sao Paulo. While AMS-IX grew by 6%, PTT Sao Paulo increased by 54%. At some days up to 5 new members joined the exchange point in Brazil during one single day. 21
4.4 The average capacity of an IXP increased over the last 6 years. The distribution did not change in general: most of the overall capacity is still shared among a few large IXPs. 21
4.5 Most of the peering participants did not change their network capacity between March 13, 2014 and December 3, 2014. However, a few of them, mostly content providers or distributors, grew a lot in terms of capacity (i.e., regarding their peering interconnections). 22
4.6 While Apple Inc. started with almost no connection capacity, and grew very fast during this time, Netflix already had many connections and prepared for their market launch in September, which is marked as a dotted line. 23
4.7 The overall capacity of all connections between IXPs and publicly peering networks grew during the observation period. While IPv4 and IPv6 addresses grew approximately by the same value, undetectable IP versions grew less. 24
4.8 In the last few months networks distributing content grew a lot faster than other network types. While network service providers grew by 19%, Cable/DSL/ISP by 25%, content distributors or content providers increased their capacity in total by 38% from March to December 2014. As discussed in Section 3.3 only IPv4 connections are considered. 24
4.9 The CDF of the ages of peering participants are shown. Here, the snapshot from March 13, 2014 was compared to the one from December 3, 2014. That corresponds to 265 days, which is depicted with a vertical line. 25

List of Tables

4.1 IXPs which grew by more than 30 members are either large ones in Europe or located in South America. 20
4.2 Participants with a capacity growth above 400 Gbps during the 9 month observation period. Most of them are content providers. 22

Chapter 1

Introduction

Traditionally, traffic between two small Autonomous Systems (ASes) connected to the Internet is exchanged with the help of an upstream Internet service provider (ISP). Such connections and their dynamics have been studied extensively in the past. Internet exchange points (IXPs) provide infrastructure to establish a direct peering connection between two ASes without using an external ISP, also enabling the capability of multilateral peering [7]. They are often built as a high capacity layer 2 switch. Potential participants who are willing to exchange traffic with each other can rent ports for a monthly fee. Since IXPs are, at least in Europe [4], often operated in a non-profit fashion, this fee is rather small compared to monthly costs of an upstream connection. Peering relations themselves are often established for free. To find potential peering partners, participants often rely on a peering database called PeeringDB. Its entries are self-contributed and are publicly available. The database is maintained by the exchange points and the participants themselves. While the traditional way to establish peering links between two ASes has been studied throughly, IXPs became only recently a topic of interest to researchers [6, 8, 1, 4, 12].

Those recent studies show that IXP connections play a much more important role in the Internet ecosystem than anticipated [1]. The number of peering links which were found at one single IXP exceeded the assumed total of all peer-to-peer links in the whole Internet. Ahmad et al. [2] discovered that routing delays on routes over IXPs are smaller than paths taking the traditional approach of hierarchical routing (i.e., first client AS, upstream ISP hierarchy, downstream ISP hierarchy, second client AS), even though they are just slightly shorter than non-IXP paths in terms of hop count. Chatzis et al.[4] evaluated the IXP ecosystem in general and encouraged the research community to put more effort in analyzing this still underestimated topic. Lodhi et al. [12] used, as done in this thesis, PeeringDB to investigate the peering ecosystem and verified the quality of parts of the data. They concluded that the data in PeeringDB may be a reliable source of information. On the other hand, they did not investigate the time dynamics and the temporal evolution of the PeeringDB dataset. In 2012, Cardona et al. [3] examined the history of one single IXP, namely the Slovak Internet eXchange (SIX). Their research is based on data gathered from the official website of SIX.

The goal of this thesis is, not to examine historical details of one single IXP, but rather look at the changes of the IXP ecosystem as a whole using a comprehensive dataset. The aim is to examine the time dynamics of the IXP ecosystem as this particular field has not been studied in depth yet. Therefore, we investigate whether larger IXPs grow faster or slower than smaller ones and identify where and how much the bandwidth capacity of connected networks changes. For this purpose, and with the help of an already existing framework, PeeringDB is crawled on a daily basis and its content is stored in an appropriate python data structure. We use PeeringDB as it provides a broad view on the peering ecosystem. This thesis extended the framework with the following capabilities:

- build a file which shows all differences between two snapshots of the database
- analyze the changes using a time-series analysis
- import further database snapshots given as historical SQL dumps

First, the format and the content of PeeringDB are described in Chapter 2, as well as the dataset we used to perform our analysis. Chapter 3 describes our approach, the functions we used to analyze the changes, and the properties of the dataset we relied on. We will see in Chapter 4 that the capacity grew almost exponentially during a period of 6 years, that IXPs grow related to their size with respect to their capacity, and that participants that increase their bandwidth capacity at IXPs are mostly content distribution networks or content providers. Finally, Chapter 5 summarizes the key findings and provides an outlook on future work that could be done by further analyzing the huge dataset.

Chapter 2

Dataset

In this chapter, we describe the dataset used in this thesis. Section 2.1 explains PeeringDB and in which format it stores the data. In Section 2.2 we present which data was available and how we retrieved the data. A description of data artifacts is given in Section 2.3.

2.1 PeeringDB

Using an IXP instead of an upstream provider in order to connect to other parts of the Internet can lead to certain advantages such as potential latency reduction, packet loss decrease, and lower costs. While connecting to an IXP is simple and cheap, establishing peering relationships or finding the most convenient peering location is more difficult. PeeringDB assists potential participants in finding peering partners. It gives IXPs and potential participants a way to share information like their points of presence, their peering policy, or their traffic levels. PeeringDB contains information about three types of entities, which we describe in the following sections.

2.1.1 Internet Exchange Point (IXP)

Exchange points are the central elements of the database. They provide information such as their IP ranges, supported protocols, or contact details. A list shows all participants (see 2.1.2) which want to peer at this location. Another list presents all facilities (see 2.1.3) connected to this IXP. Large exchange points, such as AMS-IX or DE-CIX, have hundreds of participants and are present at multiple facilities.

An entry of an IXP includes the following properties:

- *ID*: a unique identifier (see Subsection 3.1.3)
- *general information*: common name, long name, city, country, continental region, media type, supported protocols (IPv4, IPv6, multicast)
- *contact information*: website, traffic statistics website, email, phone number
- *CIDR address blocks*: address blocks of IXP for IPv4 and IPv6
- *local facilities*: a list of all local facilities this IXP is present at, including the ID, their participant count and location
- *peering participants*: a list of all participants peering at this IXP, including the ID, name, ASN, one of their ip addresses, number of ip addresses, and their peering policy

2.1.2 Participant

Networks which are either connected to public peering exchange points or private peering facilities are called participants. PeeringDB distinguishes between several types of participants:

Internet service providers (36%), network service providers (30%), content providers / distributors (22%), enterprises (4%), research networks (4%), or others (4%).[1] An entry of a participant includes the following properties:

- *ID*: a unique identifier (see Subsection 3.1.3)
- *general information*: company name, company website, geographic scope
- *contact information*: role, contact name, phone number, email
- *network information*: primary ASN, traffic levels, traffic ratio (inbound / outbound), network type (ISP, CDN, etc.), approximate amount of prefixes, route server url, supported protocols (IPv4, IPv6, multicast)
- *last updated*: indicates when the entry was updated last
- *peering information*: general peering policy (open, selective, or restrictive), peering policy URL, ratio requirements, location requirements, contract requirements
- *notes*: often used to clarify their peering policy, or explain their peering strategy in more detail
- *public peering points*: a list of all IXPs the participant is connected to, including the ID of the IXP, the name of the IXP, the capacity of the link, the IP address, and the ASN
- *private peering facilities*: a list of all IDs of the facilities the participant is connected to including the ID, name, ASN, location, and type of connection

2.1.3 Facility

Facilities are entities, such as data centers, which make private peering possible. Some facilities (28%) are connected to one or more public exchange points. These connections create the possibility for their participants to use the rich public peering environment of an IXP. An entry of a facility includes the following properties:

- *ID*: a unique identifier (see Subsection 3.1.3)
- *general information*: name, facility management, address, contact information, website
- *local exchanges*: a list of all IXPs connected to this facility, including the ID, name and participant count
- *private peering participants*: a list of all participants present at this facility, including the ID, name, ASN, and type of connection

2.1.4 Summary

As shown, there are three basic elements in PeeringDB of which each contains meta-information: IXPs, participants, and facilities. These entities are interconnected with each other, with the connections being represented with links to the ID of connected entities. Some information, such as IP addresses are listed on both sides of the connections, while others, such as link capacities, are only listed on one side. This thesis focuses on IXPs and their participants rather than the facilities.

[1] This information was compiled using the December 3, 2014 dataset snapshot.

2.2 Collection of the Dataset

In order to perform the time analysis we rely on a series of daily snapshots and some historic data. During the time from March 13, 2014 to December 3, 2014 we rely on 247 snapshots. These were taken by crawling and parsing the PeeringDB [14] website on a daily basis. Since PeeringDB is providing snapshots of its MySQL database, we were able to find 7 additional snapshots on the Internet Archive [9]. The historic data dates back to February 2008, and includes October 2009, October 2011, December 2011, January 2012, April 2012, and June 2012.

The used data structure, as presented in Section 2.1, is closely oriented to the website's data representation rather than the one of the internal database. Since the historical snapshots found at the Internet Archive are MySQL dumps of the internal PeeringDB database, converting them into the data structure used by the analysis tools was necessary. The dumps were imported into a local MySQL server and queries were executed to gather the data. Artifacts, such as connections between non-existent IXPs and participants, were found in the dumps. The challenge was to adjust the queries accordingly, in order to ignore these artifacts the same way as it is done when viewing the PeeringDB website. We also encountered coding issues, since the database mixes the MySQL latin with Unicode encoding.[2]

2.3 Artifacts in Dataset

Most of the data was gathered by crawling the PeeringDB website. All the information from each single snapshot is distributed across 8000 single web pages. Since the crawling process is not an atomic operation, it sometimes happened that entries in the database changed during the crawling process. As a consequence, a connection from a participant to an IXP sometimes disappeared for one day and reappeared the next day, which led to bandwidth changes for one single day. Such problems often occurred at large entities. These entries were removed during our sanity checks.

Due to the self-contributing nature of PeeringDB, some data is unparsable, wrong or outdated. Sometimes it is impossible to distinguish between IPv4 and IPv6 addresses. This is the case if participants do not report correctly formated addresses. Snijders [17] lists some of the most challenging IP address entries to parse such as "2001:7F8:20:101::(245/247):61/64", or "Soon". During the analyses of the capacity (e.g., in Section 4.3), ultra high capacity links between IXPs and participants were detected; these connections were unrealistic since there is no technology available supporting such large capacity and they were therefore corrected. The following data needed to be adjusted in order to obtain realistic results.

- The connection of Virgin Media (ID 1412) to the exchange points LINX Extreme LAN (ID 321) and LINX Juniper LAN (ID 18) was changed from 1100 Gbps to 110 Gbps. This corresponds to the website of LINX [13].
- IX Australia (ID 7021) connects 5 IXPs in Australia with each other. These connections were set as unrealistic 100 Tbps and were thus removed.

Overall, we removed connections with a capacity of approximately 502 Tbps. The capacity of connections with undetectable IP version decreased from 508 to 6 Tbps.

Lodhi et al. [12] already mentioned that some data is not correct or outdated. However, we do not believe that our analysis suffers too much because of these artifacts since we focus on the general dynamics and trends of the IXP ecosystem. The update frequency of the participants' entries is discussed in Section 4.4.

[2] `https://github.com/wbolster/mysql-latin1-codec`, and `http://www.whitesmith.co/blog/latin1-to-utf8/` explain coding issues in MySQL database.

Chapter 3

Analyzing the Time Dynamics of PeeringDB

In the course of analyzing how the PeeringDB [14] dataset changes over time, it is essential to know which properties change more frequently than others. Therefore, we conduct two types of analyses in this thesis. On one hand, we examine the differences between two single points in time in PeeringDB in order to detect frequent changes and outliers. On the other hand, we study the dynamics of the database with a time series analysis. As long as daily snapshots are available, this day-by-day analysis shows when the outliers' growth took place (e.g., the evolution of the bandwidth of one ISP) and enables us to further investigate the cause of the changes.

Section 3.1 explains the algorithm used to get all the properties which change between two points in time. This information is used to detect frequent changing attributes. Analyzing the change of IXPs' member size or participants' connection capacity is explained in Section 3.3.

The analysis is, depending on the type, divided into two or three steps. Figure 3.1 illustrates these steps. First, all differences between each of the involved snapshots are calculated and stored on disk. Second, the interesting information is extracted. In case the analysis considers only two points in time, this information is used to produce a plot. Otherwise, the information is concatenated in a third and last step for a further time series analysis. With such an analysis we investigate for example the capacity of an IXP over time. To get the initial bandwidth, we compare the start snapshot with an empty snapshot. This results in the addition of all connections. Now, time slot by time slot, we add capacity changes, calculated by the difference algorithm.

3.1 Difference Algorithm

In order to examine the time dynamics of PeeringDB we develop a program that compiles the differences between two snapshots. We describe in this section how the algorithm works, and on which properties it relies. The goal of the algorithm is to compare two snapshots of PeeringDB and write all changes into one file, called difference file. There, all context information that is needed to identify the changing element is stored in one line per change. Doing so makes it easy to parse the changes afterwards with command line tools like `grep` or `awk`.

3.1.1 Difference Data Format and Examples

A line in the difference file is tab separated and contains the following three parts: *context information*, *changed property*, and the *change* itself. The *context information* describes at which element something changed, for example at an IXP with ID 17. The *changed property* describes the attribute that changed, e.g. Common Name. *Change* can be either `added`, `deleted`, or `old_value changed to new_value`. The lines below show an example of a change. The participant with the ID 1470 added a connection to the IXP with the ID 73, and the IPv4 connection capacity between these two entities changed from 0 Mbps to 100 Mbps.

```
Participant  1470  IXP  73  added
Participant  1470  IXP  73  ipv4_bandwidth   0  changed to  100
```

As seen in this example, we identify entities by their respective ID. Subsection 3.1.3 shows that these IDs are unique in one single snapshot and do not change across multiple snapshots.

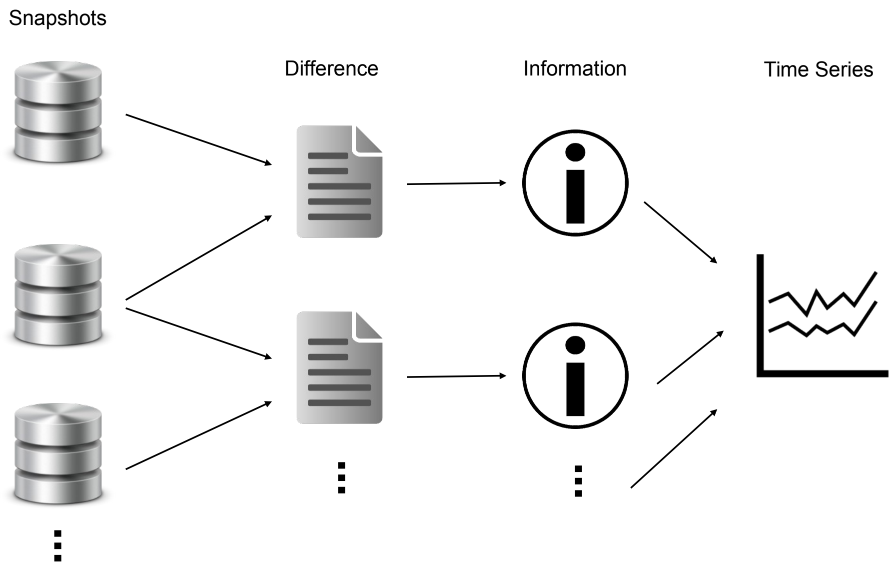

Figure 3.1: First, difference files are built. Extracting the information of potential interest forms the next step. Eventually, the information is concatenated and shown in a time series analysis.

3.1.2 Algorithmic Steps

The following steps are performed to write all the changes of one IXP into a file: the first step is to check which IDs of the IXPs only exist in the first snapshot, and which only exist in the second snapshot. We list these IXPs, including all their properties (listed in 2.1.1), as deleted and added, respectively. Now, there are only IXPs left which exist in the old and the new snapshot. Each property of the old IXP is compared with the same attribute of new IXP. If they differ, a line is written in the difference file.

Participants and facilities are handled analogously.

3.1.3 Uniqueness of IDs

We expected that all IDs of IXPs, participants, and facilities are unique. The confirmation that IDs are unique in one single snapshot is found by checking the database scheme provided in the SQL dumps of PeeringDB. The IDs are defined as unique primary keys. We therefore assumed that IDs do not change from one snapshot to another, which would make it easier to compare two snapshots. To verify this assumption, entities with different IDs and entities with identical IDs were compared to each other using a similarity function. Given two IXPs, participants, or facilities, the function returns a value between 0 and 1. If the similarity is 1, the entities are identical. If the function returns 0.5, half of the attributes are identical. If 0 is returned, they have no property in common at all. The algorithm ignores fields which are empty in both entities; otherwise two IXPs only reporting their names and leaving every other property empty would return a similarity close to 1. We show the exact approach in Algorithm 1.

similarity = 0;
max_similarity = 0;
foreach *property of element* **do**
  if *property is the same in old and new element and non-empty* **then**
    similarity +=1;
    max_similarity+=1;
  else if *property is non-empty in one element* **then**
    max_similarity+=1;
end
return similarity / max_similarity

Algorithm 1: Similarity Function

The CDFs of similarities between elements with matching IDs and with different IDs are compared in Figure 3.2.

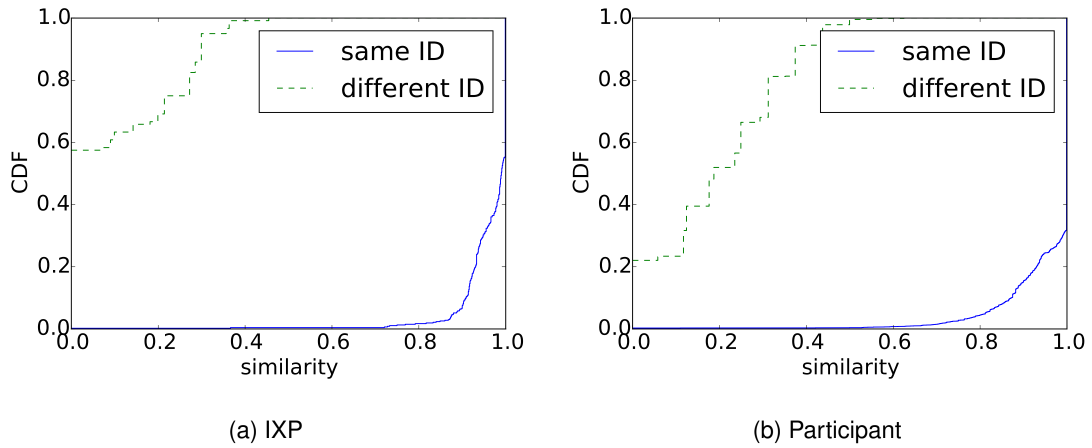

Figure 3.2: The CDF of a similarity function shows that elements with identical IDs are very akin. Elements with different IDs have very little similarity. We thus assume that IDs are unique and consistent over multiple snapshots. Here, the snapshot from June 26, 2014 is compared to the one from October 21, 2014.

3.2 Extract Information

The difference files produced in the first step of the difference algorithm are text files which contain one change per line. In this step, information is extracted out of these lines, is summarized and is then put in a python data structure. This information can afterwards be used to perform the actual time series analysis.

One possible piece of information gained from a difference file is the overall change of bandwidth across all participants. For example, the difference file may show the following lines:

```
Participant  1470  IXP  73  ipv4_bandwidth   0  changed to  100
Participant  2432  IXP  73  ipv4_bandwidth  30  changed to  100
IXP  1  Member  4224  ip_address  2001:::1  changed to  2001::2
```

The script then aggregates interesting information (here line 1 and 2) and ignores the facts that are not relevant for our analysis (here line 3), such as the change of the IP address. The following information is extracted from the example lines above:

- Participant 1470: increased bandwidth by 100 Mbps
- Participant 2432: increased bandwidth by 70 Mbps
- IXP 73: increased its total bandwidth by 170 Mbps

- Overall bandwidth of all entities increased by 170 Mbps

IXPs do not report the capacity of their connections to participants. However, participants list this capacity. As seen in the example, this information is added to the IXP. The connection capacity of IXPs can only be derived from looking at the information listed at its participants.

3.3 Growth Analysis

Chapter 4 shows that IXPs and participants grew in general regarding their member count and their bandwidth capacity during our observation period. In order to explain this fact we compared their growth with the size they had at the start of the period. First, the absolute member increase of the IXP versus the initial member count is visualized in a scatter plot. Since many data points are very close to each other, we do not see the complete situation in this plot. The cumulative distribution function (CDF) of the relative growth will help to see the general dynamics. To define the relative growth simply as $\frac{\texttt{new\_size-old\_size}}{\texttt{old\_size}}$ would result in values $\pm\infty$ if everything was added or deleted. We thus define the relative growth as: $\texttt{relative\_growth} := \frac{\texttt{new\_size-old\_size}}{\texttt{old\_size+new\_size}}$. This results in values from -1 to 1, where -1 is the complete deletion, 0 means stagnation, and 1 implies an initial addition. Halving the bandwidth results in a relative growth of -1/3, doubling in +1/3.

When comparing the bandwidth capacity of links, only IPv4 connections were considered. We assumed that all connections supporting IPv6 also support IPv4 and thus do not need to be considered separately. This assumption is justified by the fact that only a very little fraction of the traffic on IXPs is IPv6 traffic [10]. Renting an interface for only running an IPv6 stack would not be cost-effective. Furthermore, PeeringDB Beta [15] confirms this assumption by only listing one connection and assign both IP addresses to this single connection. Thirdly, at AMS-IX, taken as an example, all connection capacities summed up result in a capacity of 10.8 Tbps. This is only 7% less than what their website reports: 11.6 Tbps. Since IP addresses are stored as strings in PeeringDB, we used the python module *ipaddress* to match the strings to IPv4, IPv6, or unknown IP versions.

Chapter 4

Results

While examining the differences between two single points in time in PeeringDB [14], March 13 and December 3, 2014, outliers were detected: IXPs which grow faster than others, or participants which have extensively increased their capacity. We see that connections between IXPs and peering participants changed the most frequently: 3565 connections were established, while 899 were deleted. Since there are many changes in connections between IXPs and participants, the member count and bandwidth capacity of IXPs is examined in Section 4.1. Most of the fast growing participants are content distribution networks (CDN) or content providers, which can be seen in Section 4.2. By looking at the evolution of the overall bandwidth capacity of all IXPs together, Section 4.3 shows that most of the recently added connections support IPv4 and IPv6.

4.1 Dynamics of IXPs

During the 9 month observation period 86 IXPs were added to the database. During the same period, only 6 IXPs disappeared. It is hard to say whether this small number is accurate or not. It may also be possible that IXPs do not remove their PeeringDB entries when they disappear. Figure 4.1 depicts a comparison of the absolute and the relative change of the member count. The member count change correlates vaguely with a coefficient of 0.59 to the member count at start. The correlation of the absolute capacity to the capacity at start is, on the other hand, significantly higher with a correlation coefficient of 0.92 (see Figure 4.2). All the IXPs with more than 150 members grew, as Subfigure 4.1a shows, and we clearly see one outlier, PTT Sao Paulo, on which we will have a closer look later on. While most IXP gained or lost members, the amount of members of 45% of the IXPs remained stable. Subfigure 4.1b indicates that most of the changing IXPs started either with size zero, or grew between 0 and 100% (recall from Section 3.3 that the line at +1/3 means a doubling of the member count). PTT Sao Paulo, the outlier seen in Subfigure 4.1a, grew by 127 members from 235 to 362. The member size history of PTT Sao Paulo is compared to the one of AMS-IX in Figure 4.3. This fast growing IXP, which grew by up to 5 members a day, is part of the PTTMetro project [16]. This project, organized by the Brazil Ministry of Communications, maintains over 20 IXPs and interconnects them all. Fast growing networks of IXPs exist also in other parts of South America, namely Argentina (Cámara Argentina de Internet). Table 4.1 lists all IXPs which grew by more than 30 members during the entire observation period.

Most of the overall capacity is located at a few large IXPs as seen in Figure 4.4. While this did not change over the last years, the size of the average IXP, however, increased. The 10 with respect to the connection capacity largest IXPs (2%) claim almost half (49%) of the overall capacity.

4.2 Dynamics of Participants

While the capacity increase at IXPs is closely related to their size, we see in Figure 4.5 that this growth is more scattered for participants. A correlation coefficient of 0.60 confirms this im-

IXP Name	Start	End	Growth	relative
PTT Sao Paulo	235	362	127	+54%
Cámara Argentina de Internet	10	56	46	+460%
DE-CIX, the global Internet Exchange	470	525	55	+12%
London Internet Exchange Ltd. (LINX)	474	513	39	+8%
Amsterdam Internet Exchange (AMS-IX)	601	636	35	+6%

Table 4.1: IXPs which grew by more than 30 members are either large ones in Europe or located in South America.

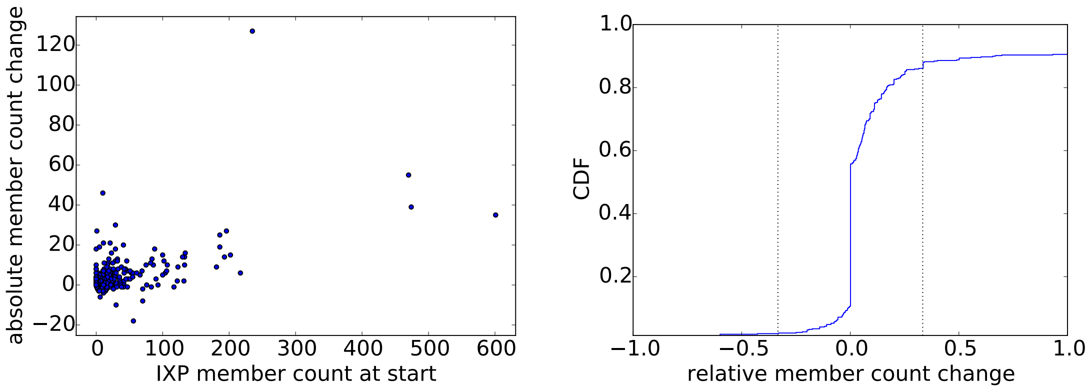

(a) change of absolute member size

(b) CDF of relative change of IXP's member count

Figure 4.1: The biggest outlier which grew by 54% is PTT Sao Paulo. The three outliers with the largest member count at start are AMS-IX, DE-CIX and LINX.

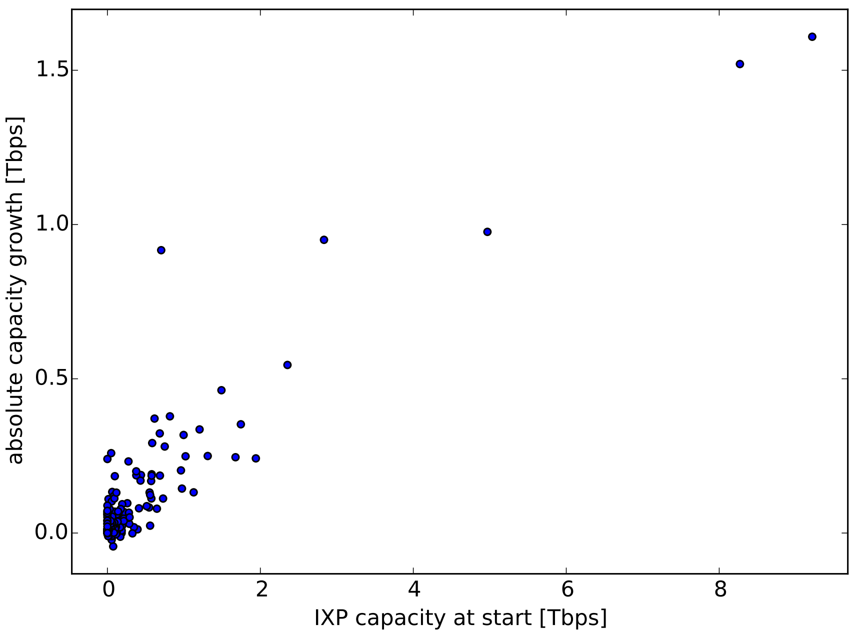

Figure 4.2: The capacity of IXPs grew in the observation period. At larger IXPs, the growth of capacity is related to their size. The two rightmost points represent AMS-IX and DE-CIX, two of the largest Internet exchange points in the world.

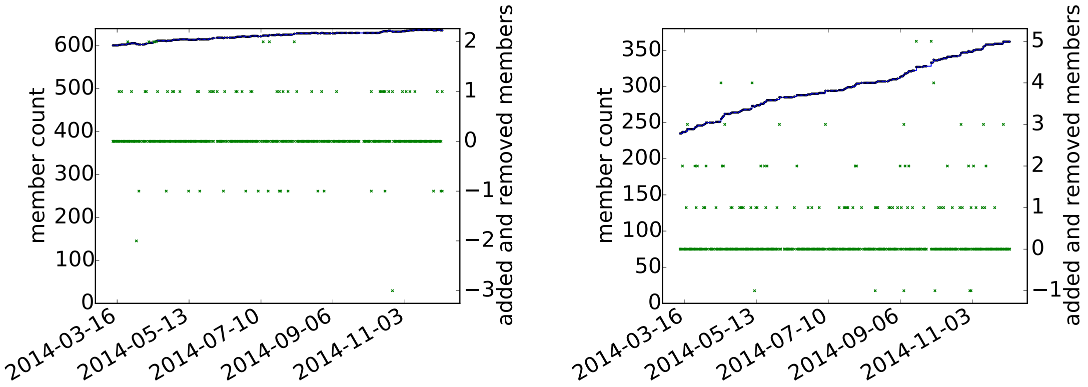

(a) AMS-IX

(b) PTT Sao Paulo

Figure 4.3: Comparison of the member size history of AMS-IX and PTT Sao Paulo. While AMS-IX grew by 6%, PTT Sao Paulo increased by 54%. At some days up to 5 new members joined the exchange point in Brazil during one single day.

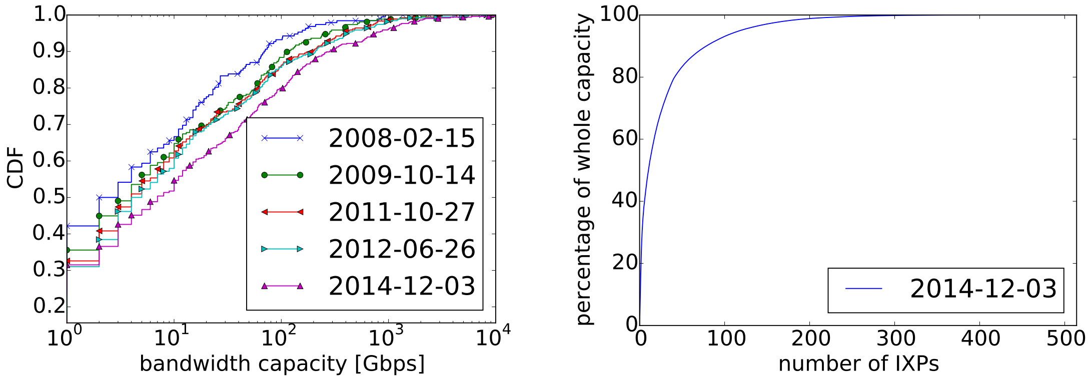

(a) CDF of the capacity distribution of IXPs

(b) IXPs' capacity distribution

Figure 4.4: The average capacity of an IXP increased over the last 6 years. The distribution did not change in general: most of the overall capacity is still shared among a few large IXPs.

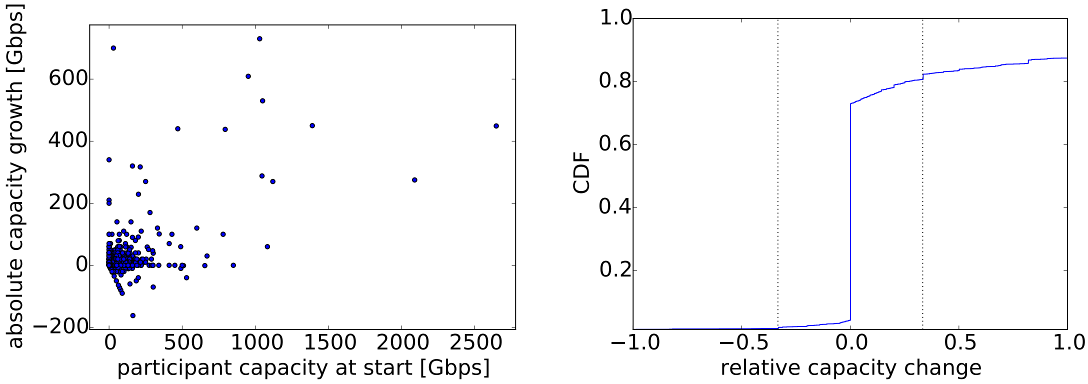

(a) absolute growth of networks

(b) relative growth of networks' capacity

Figure 4.5: Most of the peering participants did not change their network capacity between March 13, 2014 and December 3, 2014. However, a few of them, mostly content providers or distributors, grew a lot in terms of capacity (i.e., regarding their peering interconnections).

pression. 45% of the IXPs did not report a change in their member count and as Subfigure 4.5b indicates, there are even more participants which do not report a change in their connection capacity (69%). Subfigure 4.5a shows that there are a few networks which grew intensively. 8 participants increased their peering capacities by more than 400 Gbps. Most of them announce themselves as content providers or distributors and are listed in Table 4.2.

Company Name	Type	Start	End	Growth	relative
Facebook	Content	1030 Gbps	1760 Gbps	730 Gbps	+70%
Apple Inc	Content	30 Gbps	730 Gbps	700 Gbps	+2333%
Amazon.com	Enterprise	952 Gbps	1561 Gbps	609 Gbps	+64%
EdgeCast Networks, Inc.	Content	1050 Gbps	1580 Gbps	530 Gbps	+50%
Netflix	Content	1390 Gbps	1840 Gbps	450 Gbps	+32%
Akamai Technologies	Content	2648 Gbps	3097 Gbps	449 Gbps	+17%
Microsoft	NSP	794 Gbps	1232 Gbps	438 Gbps	+55%
Twitch Interactive[a]	Content	470 Gbps	910 Gbps	440 Gbps	+94%

[a] acquired by Amazon.com

Table 4.2: Participants with a capacity growth above 400 Gbps during the 9 month observation period. Most of them are content providers.

Two examples of fast growing content providers are compared in Figure 4.6. The two cases are different, since Apple's capacity started at a very low value compared to other content providers and grew by 700 Gbps. The reason for Apple's growth - or its appearance at all - is that the company no longer wants to completely rely on networks of third parties while distributing content to their users. Netflix, on the other hand, already had a lot of connection capacity to IXPs and increased it selectively. Netflix did their market launch in Europe in mid of September 2014. In order to manage their content distribution, the company set up a connection to France IX on July 29th, ECIX Frankfurt on August 9th and to ECIX Düsseldorf on August 15th, each having a capacity of 100 Tbps. This rollout process is reflected in Subfigure 4.6b. Both of the companies, Apple and Netflix, have announced hardly any invalid IP addresses and all their connections are IPv6 enabled.

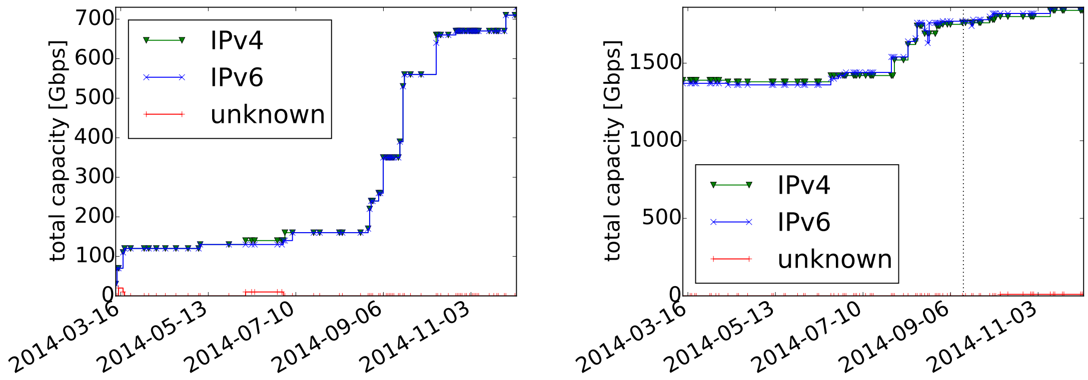

(a) Apple Inc. (b) Netflix, Inc.

Figure 4.6: While Apple Inc. started with almost no connection capacity, and grew very fast during this time, Netflix already had many connections and prepared for their market launch in September, which is marked as a dotted line.

4.3 Evolution of the Overall Capacity

Networks that peer at exchange points announce their capacities of their peering connections (i.e., IXP-facing interfaces). Figure 4.7 shows all these announced capacities added up. The IPv4 capacity grew from 6.3 Tbps to 87.2 Tbps from February 2008 to December 2014. The growth looks similar to an exponential behavior, just as Labovitz et al. [11] described. We see that not only SIX, as shown by Cardona et al. [3], but also the IXP environment in general shows a tendency to grow in an exponential manner. We thus conducted an exponential fit with the method of least squares by using all available data points. Doing so leads to a yearly growth of 37.2%, while Cisco estimated in their VNI in 2010 [5] a yearly growth of 34%.

A comparison between different IP versions shows that most of the added connections support IPv6. The gap between the two versions does not widen. We thus conclude, that the connections between IXPs and its participants are mostly IPv6 ready. While IPv4 and IPv6 connections grow fast, the amount of incorrect entered IP addresses, on the other hand, grows slower.

In Section 4.2 we saw that most of the fast growing participants are content providers. Figure 4.8 shows all the IPv4 connections separated by their respective network type. While all types grow, content providers grow by far the fastest. This is not surprising as an IXP is an excellent point to feed content into the network without having to pay huge connection fees.

4.4 Update Frequency of Database

If a peering participant changes some of its properties, its "Last Updated" field is updated. This provides some information on how up-to-date the entries in PeeringDB are. Out of 3858 participants which existed at the start and the end of the observation period of 265 days only 1504, which corresponds to 39%, updated any of its properties. Figure 4.9 shows the CDF of the entries' age. In general the participants' entries became more up to date. This is mainly due to recently added participants or very old entries which were removed. Nevertheless, only 55% of all participants changed any of its entries during the last year.

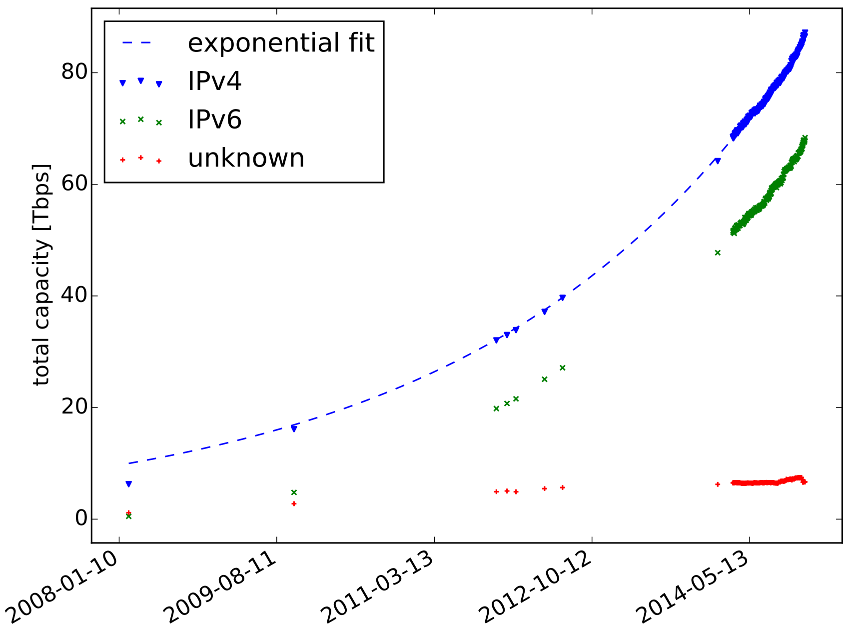

Figure 4.7: The overall capacity of all connections between IXPs and publicly peering networks grew during the observation period. While IPv4 and IPv6 addresses grew approximately by the same value, undetectable IP versions grew less.

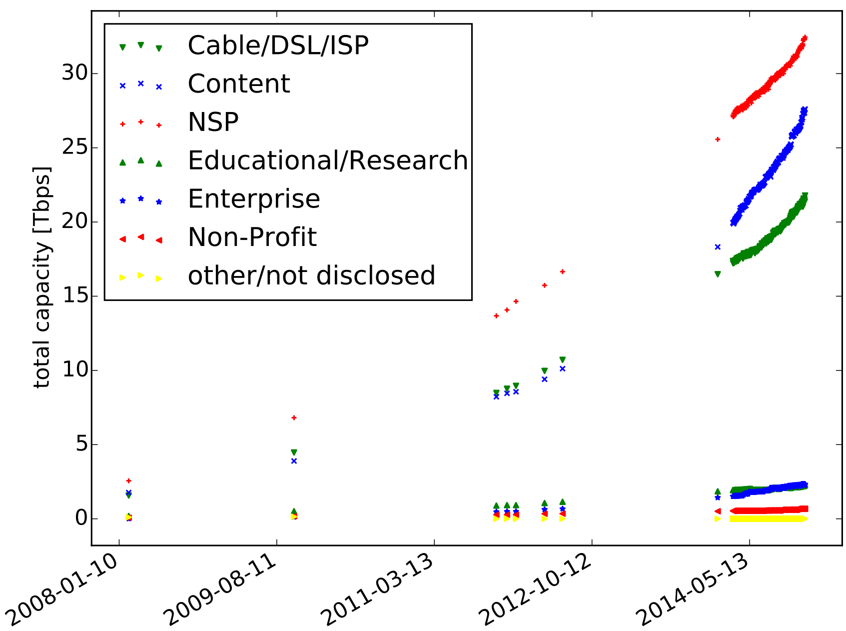

Figure 4.8: In the last few months networks distributing content grew a lot faster than other network types. While network service providers grew by 19%, Cable/DSL/ISP by 25%, content distributors or content providers increased their capacity in total by 38% from March to December 2014. As discussed in Section 3.3 only IPv4 connections are considered.

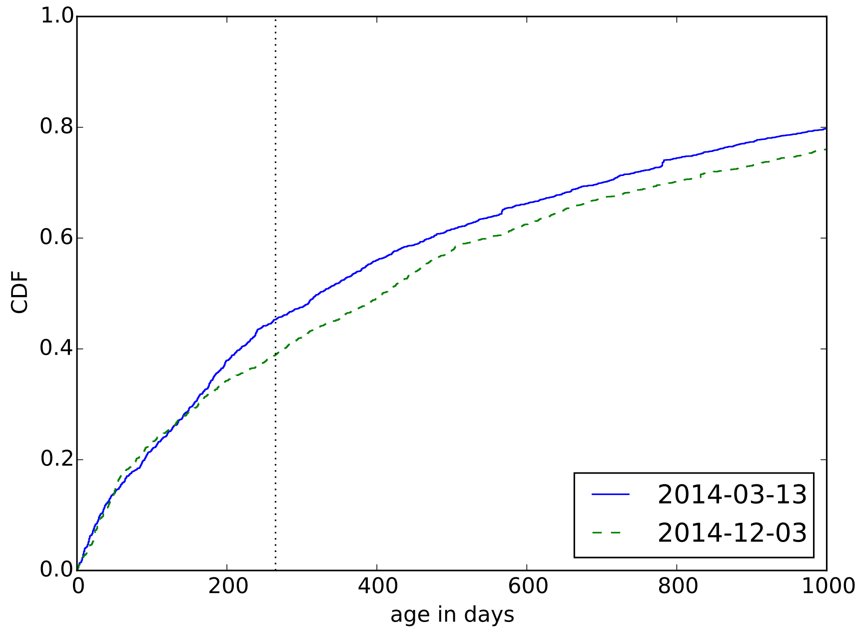

Figure 4.9: The CDF of the ages of peering participants are shown. Here, the snapshot from March 13, 2014 was compared to the one from December 3, 2014. That corresponds to 265 days, which is depicted with a vertical line.

Chapter 5

Conclusion and Future Work

5.1 Conclusion

Internet exchange points recently became a topic of interest to researchers. This thesis examines the time dynamics of the PeeringDB dataset. Given a framework which crawls the PeeringDB website and fetches its content into a python data structure, we extended the framework to make it capable of analyzing the time dynamics. Additionally, it is now possible to include the database content if it is given as a SQL dump.

We discovered that, with respect to their bandwidth, IXPs grow related to their capacity. Many IXPs increased in terms of member count. Outliers were found in Argentina and Brazil, where some IXPs grow very fast. The largest part of all the available capacity at IXPs is located at very few IXPs: 50% of the bandwidth is distributed across only 2% of IXPs. While all network types grow, the fastest growing group of peering networks are content providers or distributors. These are either operated by large companies such as Apple, Facebook, or Netflix for their own use, or companies which provide this service for clients, such as EdgeCast Networks, or Akamai Technologies. While Cardona et al. [3] saw a very fast traffic growth in some connections at SIX (Slovak Internet eXchange) (sometimes more than 100% per year), the overall capacity across all exchange points grows fast, but slower than the traffic at SIX. As available IPv4 addresses are becoming scarce, it is a good sign that IXPs and their peering participants are in general ready for the next generation of the Internet protocol, namely IPv6.

5.2 Future Work

The dataset of PeeringDB contains an immense amount of information. This thesis focused on interesting pieces of information which were found by looking at often changing properties. It would be interesting to find automatically properties which may be worth to have a closer look on. One could find correlated properties by applying PCA (principal component analysis) or ANOVA (analysis of variance). Correlating changes with events such as company takeovers, or stock prices may lead to interesting results. We compared some findings of the work of Cardona et al. [3] with our own analysis. Examining if more of their conclusions can be extrapolated from one IXP to the whole ecosystem would be an enriching extension. Information from other data sources such as *The European Internet Exchange Association* (Euro-IX) as well as *Packet Clearing House* (PCH) may be considered and it should be checked whether the dynamics of these datasets are consistent with our findings.

Bibliography

[1] AGER, B., CHATZIS, N., FELDMANN, A., SARRAR, N., UHLIG, S., AND WILLINGER, W. Anatomy of a large european IXP. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication* (2012), ACM, pp. 163–174.

[2] AHMAD, M. Z., AND GUHA, R. Studying the effect of internet exchange points on internet link delays. In *Proceedings of the 2010 Spring Simulation Multiconference* (2010), Society for Computer Simulation International, p. 103.

[3] CARDONA RESTREPO, J. C., AND STANOJEVIC, R. A history of an internet exchange point. *ACM SIGCOMM Computer Communication Review 42*, 2 (2012), 58–64.

[4] CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. There is more to IXPs than meets the eye. *ACM SIGCOMM Computer Communication Review 43*, 5 (2013), 19–28.

[5] CISCO SYSTEMS, INC. Cisco visual networking index: Forecast and methodology, 2009–2014. `http://large.stanford.edu/courses/2010/ph240/abdul-kafi1/docs/white_paper_c11-481360.pdf`, June 2010. Last accessed on: 2015-01-02.

[6] DHAMDHERE, A., AND DOVROLIS, C. The internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of the 6th International COnference* (2010), ACM, p. 21.

[7] GIOTSAS, V., ZHOU, S., LUCKIE, M., ET AL. Inferring multilateral peering. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies* (2013), ACM, pp. 247–258.

[8] GREGORI, E., IMPROTA, A., LENZINI, L., AND ORSINI, C. The impact of IXPs on the as-level topology structure of the internet. *Computer Communications 34*, 1 (2011), 68–82.

[9] INTERNET ARCHIVE. `https://archive.org`. Last accessed on: 2015-01-12.

[10] KIM, J., SARRAR, N., AND FELDMANN, A. Watching the ipv6 takeoff from an ixp's viewpoint. *arXiv preprint arXiv:1402.3982* (2014).

[11] LABOVITZ, C., IEKEL-JOHNSON, S., MCPHERSON, D., OBERHEIDE, J., AND JAHANIAN, F. Internet inter-domain traffic. *SIGCOMM Comput. Commun. Rev. 41*, 4 (Aug. 2010), –.

[12] LODHI, A., LARSON, N., DHAMDHERE, A., DOVROLIS, C., ET AL. Using peeringdb to understand the peering ecosystem. *ACM SIGCOMM Computer Communication Review 44*, 2 (2014), 20–27.

[13] LONDON INTERNET EXCHANGE. Linx member info. `https://www.linx.net/pubtools/member-techinfo/member_id/100050`. Last accessed on: 2015-01-12.

[14] PEERINGDB. PeeringDB. `http://www.peeringdb.com`. Last accessed on: 2015-01-08.

[15] PEERINGDB. PeeringDB Beta. `https://beta.peeringdb.com/`. Last accessed on: 2015-02-08.

[16] PTTMETRO. PTTMetro. http://www.ptt.br. Last accessed on: 2015-01-08.

[17] SNIJDERS, J. PeeringDB accuracy: Is blind faith reasonable?, 2013.