



UNIVERSITY CARLOS III OF MADRID

Department of Telematics Engineering

Master of Science Thesis

Empirical characterization of Internet Exchange Points

Author: **Juan Camilo Cardona Restrepo**

Supervisor: **Rubén Cuevas, Ph.D.**
External Supervisor: **Rade Stanojević, Ph.D.**

Leganés, September 2011

Abstract

Today's public Internet eXchange Points (IXP) are an important factor in the Internet ecosystem, carrying around 15-20% of the Internet inter-domain traffic. In spite of their importance, the community still lacks empirical data on the nature of the traffic exchanged through IXPs. We strive in obtaining a deeper understanding of exchange points by closely examining two medium-size Europeans IXPs: the Slovakian-IX and the Milan-IX.

By crawling the web archive `waybackmachine.org` we obtained several snapshots of the Slovak-IXP website from the last 14 years. This historical data allows us to study the dynamics of the IXP in terms of: the (low-tier) AS-level topology, the traffic dynamics, the port capacity and the traffic matrix. Our datasets show that, different from Customer-Provider links, peering links are very stable as we observe that once they are created they are very unlikely to disappear. After the proliferation of content service providers in 2006, the traffic profiles of the ISPs peering at the IXP experienced significant changes reflected in the extreme growth of the content ISPs traffic and heavily imbalanced traffic ratios. An analysis on the distribution of traffic of the IXP hints that since the beginning of SIX a small fraction of heaviest peering pairs carry the majority of the SIX traffic while most of the peering pairs exchange very low amount of traffic and hence enjoy close-to-zero monetary gain.

In order to provide an analysis from a different perspective, we use another set of data from the Milan-IX to compare the two IXP and draw several important conclusions. We demonstrate that the peering (AS-level) topology within the IXP is impacted by the pricing model of the IXP. We also quantify the relationship between the different types (access/content/transit) of ISPs present in the two studied IXPs, both in terms of traffic volumes and peering intensity. Furthermore, the distribution of traffic per peering pairs for both IXPs is similar and dominated by a few heavy peering pairs.

We strongly believe that in contrast with confidential datasets typically used in studying the Internet traffic characteristics, IXP's data provide rich and publicly available resources crucial for understanding various aspects of the Internet.

Table of Contents

1	Introduction	1
2	State Of the Art	2
2.1	Background on Internet Exchange Points	2
2.2	Related Work	3
3	The evolution of Slovak-IX	5
3.1	Datasets	5
3.2	Peering matrix	6
3.3	Traffic dynamics	11
3.4	Capacity dynamics	15
3.5	Summary of Contributions	18
4	Comparison of SIX and MIX	19
4.1	Datasets	19
4.2	Pricing model matters!	20
4.3	Attraction and repulsion	21
4.4	Traffic distribution	21
4.5	ISP-type	22
4.6	Peering intensity vs. member size	23
4.7	Summary of Contributions	23
5	Conclusions	25
	Appendix	26
	References	29

List of Figures

3.1	A sample of the peering matrix of SIX as seen on its webpage.	7
3.2	Snapshot of mrtg data.	7
3.3	Membership duration for all of the 84 ISPs in SIX.	8
3.4	Number of members of SIX along the last 14 years.	8
3.5	Evolution of peering density of SIX (aggregate) and per ISP type.	9
3.6	Histogram of peering density for all 84 (historic) members of SIX. For each ISP, the last peering density value is used.	10
3.7	Histogram of link creation times for 1711 peering pairs in the history of SIX.	10
3.8	SIX's aggregated traffic and the traffic of top-5 current members.	12
3.9	The percentage of total inbound/outbound traffic per ISP type.	13
3.10	The evolution of ISP traffic imbalance (IB), median, 10-th, 90-th percentile and the weighted average.	14
3.11	The fraction of the total traffic from the top-1%, top-10% and bottom 50% of peering pairs.	14
3.12	Median savings of peering for SIX in the last 14 years.	16
3.13	Per member port(s) utilization of SIX since 1998: median, 10th, 90th-percentile and weighted average.	17
3.14	The estimated utilization at the time of port upgrade.	17
4.1	Exchanged traffic between pairs at SIX and MIX. Top 100 pairs shown.	20
4.2	Attraction vs Repulsion for members of SIX and MIX.	21
4.3	Fraction of total traffic generated by heaviest peering pairs at SIX and MIX.	22
4.4	Traffic distribution for SIX and MIX in 2008.	22
4.5	Peering intensity vs. traffic of ISPs peering at MIX and SIX.	24
5.1	The pseudocode for solving gravity.	27
5.2	Real throughput compared to throughput estimated using the gravity model for MIX.	28
5.3	Cumulative distribution of exchanged traffic for peering pairs in both the estimated TM and the real TM of MIX.	28

Chapter 1

Introduction

Several recent research studies have detected that stub autonomous systems are more frequently establishing peering directly among each other, thereby producing a flatter Internet structure[8, 12]. This trend could be expected since for most ISPs a direct interconnection can reduce expenditures and provide different technical benefits. For most ASes it is not economically viable to deploy dedicated infrastructure to support all their peerings, as the number of potential peers for an AS can be large. Consequently, ISPs attempt to reduce cost by using Internet eXchange Points (IXP), which are physical “meeting points” for ISPs interested in peering over specific locations. Globally there are few hundred IXPs varying in size from small regional exchanges serving single-digit number of members to large international IXPs serving hundreds of ISPs. Due to the large number of interconnections supported by IXPs, they are key elements of the Internet ecosystem.

Many IXP release statistics of their aggregated traffic and a few of them provide per-member traffic stats and data on the interconnections established inside of them. The information released by IXP has been largely underutilized, yet it can be very useful for analyzing the Internet. Furthermore, there have been different studies that attempted to estimate the number of peering links located in IXPs[2, 14, 32], however, the analysis and modeling of specific IXPs is still limited.

In this work, we attempt to understand more about IXPs by studying two European exchanges: the Milan IXP (MIX) and the Slovakian IXP (SIX). From the former, data containing the traffic exchange between members for one single day in 2008 was obtained. For the Slovakian IXP (SIX), a large amount of data, including its peering matrix and per-member traffic information, was procured using an Internet Archive[18]. We undertook two independent analyses using these two datasets. First of all, the data of the last 14 years of SIX is used to show important characteristics of the evolution of the IXP in terms of peering, traffic and port utilization (Chapter 3). On the other hand, we compare the two IXPs, showing interesting similarities and differences between them and present some initial results in the influence of pricing in IXPs (Chapter 4).

Chapter 2

State Of the Art

2.1 Background on Internet Exchange Points

Internet Exchange Points are facilities that provide ASes with the infrastructure required to establish interconnection agreements. Basically, an IXP is a layer 2 domain that can span one or multiple locations. An IXP can be implemented in different ways, from the installation of a few simple switches to more elaborated architectures that can include resiliency mechanism or complex protocols like MPLS [21]. IXPs offer ISPs the opportunity to use the same backhaul, equipment and operation personnel to support more than one interconnection, thus reducing the overall costs per individual peering.

Nowadays there are more than 200 IXP around the world [2], but detailed information on each one of them is scarce. Hence it is hard to measure how much of the Internet topology depends on IXPs, which is an important factor for different research projects (e.g. evaluating the resiliency of the Internet). Xu et al. [32] presented one of the first estimations of the amount of links held by IXPs, this work was later extended by He et al. [14] and Augustin et al. [2]. Table 2.1 shows a summary of the data found in [2]. From the 57.6K links detected by [2], around 30K were not found in other large AS-link datasets (CAIDA, DIMEA or PlanetLab), which shows how undervalued this type of links were in past Internet measure projects. The data gathered in [2] is the latest and richest source of IXP's peering links, however, as it is often the case, the methodology is efficient only for interconnections between ASes that possess some kind of "probe" (e.g. Looking Glass or a loose-source record-route router) that can be openly accessed. Since most ISPs do not have this kind of resource available, the data in [2] is still considered to be fairly incomplete, which is certified by comparing the dataset with the peering matrices of some IXPs that publish this information [20, 16, 30, 22, 28]. To give an example, the estimated peering density¹ of MIX and SIX according to [2] is 14% and 10% respectively, while their data shows a peering density of around 58% and 69%.

Other important parameter to quantify is the percentage of Intra-domain traffic routed through Exchange Points. Labovitz et al. calculate in [17] the total Internet inter-domain traffic (sum of all entries in the global AS-AS traffic matrix) in July 2009 to be around 39Tbps. Using the yearly growth of 45% from [17, 15], we can calculate the current inter-

¹The peering density is calculated as the ratio between the actual number of peerings and the number of possible peerings inside the IXP

Region	Number of IXPs	Number of Links
Europe	86	44605
Australia and Ocenia	14	2694
Asia	25	2761
USA	48	8233
Rest of America	17	609
Africa	4	22
Total	194	57577

Table 2.1: Summary of discovered IXP's AS-links included in [2].

domain traffic to around 81Tbps. On the other hand, by summing the total traffic published by the most important IXPs [25], we can estimate the fraction of total Internet inter-domain traffic exchanged via IXPs, to be in the range of 15-20%. For some ISPs the fraction of IXP traffic is even higher; e.g. in some European academic networks that openly publish their traffic stats, IXP traffic corresponds to 40-50% of their totals [26, 4]; in some other heavily localized Internet markets such as Japan this fraction can be as high as 70% [5].

The last analysis reinforces the general belief that IXPs are an important part of the Internet. Even when this statement is already understood by the research community, there is still a lot to be learned from IXPs. In this work, we broaden our knowledge of IXPs by closely studying the MILAN-IX and the Slovak-IX.

2.2 Related Work

Despite its critical importance in the Internet ecosystem, the empirical studies of existing IXPs are only starting to appear. In recent years different groups have attempted to measure the number of peering relationships at the world's IXPs [32, 14, 2]. Although the precise number of links inside IXPs is still unknown, the data from these projects has been used to study different implications of IXPs on the AS-level topology of the Internet [1, 13]. In a more economical perspective, D'Ignazio et al. use the peering matrices from different IXPs to analyze competitiveness in the Internet market [10, 9]. We are not aware of any other empiric research on IXPs that goes beyond the AS-level topology analysis.

Dhamdhere and Dovrolis [7] study the evolution of the AS-level graph, and observe, that the CP links are fairly unstable as many ISPs regularly change the providers as a part of the internal cost optimization process. Olivier et al. [24] propose a model that characterizes the changes of the AS-topology and observe from their data that CP links are more stable than peering links. Our dataset, however, allows for studying the evolution of AS-level dynamics at the peering-link level inside IXPs, and we see that peering links are very stable: once they are created they almost never disappear.

Cho et al. [5, 6] and Borgnat et al. [3] analyze the traffic trends of Japanese residential and transpacific traffic respectively, by looking at the relevant packet traces over multiple-year periods (4 years in [6], 7 years in [3]). Cho et al. [6] show that residential traffic has exponential growth and that per-subscriber traffic usage distribution remains lognormal with

mean (and variance) growing in time.

Chapter 3

The evolution of Slovak-IX

In this chapter we examine various dynamic and invariant properties of the Slovak IXP.

3.1 Datasets

The analysis presented in this chapter is based on historical data from the Slovak IX (SIX). SIX publishes in its webpage the peering matrix with info on which member pairs do peer and which do not. Additionally, SIX provides traffic statistics for each member, including maximum traffic capacity (port capacity) and traffic volumes in both directions (average and maximum)[28]. In order to obtain the data of past years (SIX has been publishing these data since 1997), we use the Internet Archive [18]. We provide hereafter more details on the data obtained for each of these characteristics:

Peering matrix is a matrix that indicates whether two members of the IXP peer between each other or not. Since peering relationship is symmetric, the peering matrix is symmetric at all times. There could be many reasons for peering or not-peering between two ISPs, and typically two ISPs would peer if and only if such peering provides financial or technical benefits to *both* of the ISPs [2, 13]. Figure 3.1 shows a recent peering matrix of SIX published in their webpage. The Internet Archive [18] stores 155 historic snapshots for the peering matrix which we extended with 3 snapshots from the months of March, May and July of 2011. One important property of the data is that snapshots are spread-out throughout the time in a non-uniform fashion: in 2007 there were 49 snapshots, while each of the years '97,'98,'99 and '09 has less than 5 snapshots.

Each participating ISP in SIX is identified by the name and AS number. In several cases, either name or AS number of an ISP change at some point, and for the purpose of our study we consider that ISP to be the same as the one before the change. However, sometimes ISPs stop peering at SIX and since 1997, we found 32 ISPs that used to peer at SIX that do not peer any more. With 52 ISPs that peer at SIX today, it brings the total number of ISPs that have participated in SIX to 84. In Figure 3.3 we sketch the lifetime graph of all of the 84 ISPs.

In addition to mapping SIX members to the ISPs (that may change name and AS number), we classified each ISP based on the type of business they are involved with: access,

content and network service providers (NSP). We used *PeeringDB*¹ service [31] to map ISPs to their type, and in few cases in which no type info was found in *peeringDB*, we manually inspected the ISP type. In Figure 3.4 we depict the evolution of the number of SIX participants: per type and total.

Per-member traffic demands and port capacity are extracted from the *mrtg* data available for *each* member at the SIX webpage. Figure 3.2 shows a (partial) snapshot of a webpage generated using the *mrtg* tool. It contains the data from UPC (large access provider), time-stamped on 12/4/2006, with the port capacity² of 1Gbps (value *Max Speed* in the graph), listing the average, max and current demand for both the inbound and outbound traffic, on daily (shown), weekly, monthly and yearly (not shown in Figure 3.2) basis. The *mrtg* tool also produces visual images depicting the daily/weekly/monthly/yearly traffic trends, and we designed a script that transforms these visual images into the numeric data. However, for the purpose of this paper, we exclusively work with the numeric data provided directly by *mrtg*: max and average.

The Internet Archive does not store all the pages in a domain simultaneously, in fact, it is usual that some web pages in a site are more frequently stored than others. The traffic statistics of each member of SIX are stored in different html files and for a few of them we do not have very frequent snapshots. However, most ISPs do have at least one *mrtg* sample data point per year, which is enough for (relatively accurate) capturing multiple-year dynamics. In the intervals between the available information, we use simple linear interpolation, to estimate the traffic at any point in time.

The data formats for both the *mrtg* and peering matrix data have varied across the years of existence of SIX website. This required careful manual tuning of our crawler, to create a consistent dataset across multiple years of operation.

3.2 Peering matrix

Due to its importance, the AS-level topology of the Internet has been one of the most comprehensively studied objects in the networks research community. In spite of a tremendous amount of work on this topic, the existing measurement tools have very low accuracy in measuring (inferring) the AS-level topology in the lower tiers of the AS ecosystem [2, 14, 13, 27]. The IXP data we use here, offer unique opportunity to examine not only the current state AS-level topology among a subset of low-tier ISPs, but also to evaluate its dynamics.

Peering density dynamics. A major difference between the AS-level topology at the IXP level and the global AS-level graph is in the density of interconnections. Namely an ‘average’ ISP typically peers with a large fraction of other ISPs from the IXP. To quantify how likely the two IXP members are to peer, we use the *peering density*, a common metric defined as the ratio between the number of peering links and the number of all possible pairs of ISPs participating at the IXP. In Figure 3.5 we plot the density of SIX across the 14 years of operation and observe that this quantity has been fairly stable over the time and is in the

¹Peering DB is an online data-base that stores the basic info on most ASes involved in peering. For the few ASes without an entry in PeeringDB, we manually inspected their type.

²Some members lease more than one port, in which case the sum of all port capacities is shown as the port capacity.





Figure 3.1: A sample of the peering matrix of SIX as seen on its webpage.

System: SIX-UPC-gw
 Interface: GigabitEthernet
 IP: (192.108.148.205)
 Max Speed: 1 Gbit/s (GigabitEthernet)
 The statistics were last updated **Wednesday, 12 April 2006 at 5:20**,
 at which time 'SIX-GigaSwitch' had been up for **332 days, 9:38:14**.

'Daily' Graph (5 Minute Average)



	Max	Average	Current
In	170.0 Mb/s (17.0%)	100.9 Mb/s (10.1%)	54.0 Mb/s (5.4%)
Out	191.6 Mb/s (19.2%)	98.8 Mb/s (9.9%)	35.6 Mb/s (3.6%)

Figure 3.2: Snapshot of mrtg data.

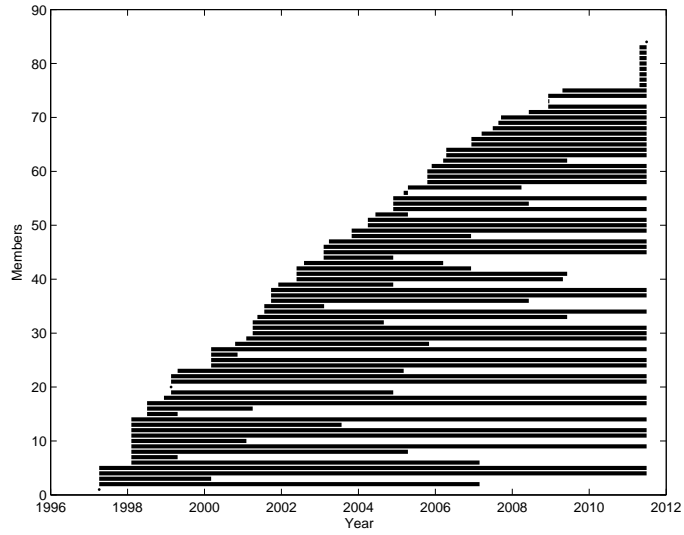


Figure 3.3: Membership duration for all of the 84 ISPs in SIX.




Figure 3.4: Number of members of SIX along the last 14 years.




Figure 3.5: Evolution of peering density of SIX (aggregate) and per ISP type.

range around 70%. This value is ‘normal’ for European standards, as the peering density in several of the IXPs that publish their peering matrix, including Czech NIX, Viena VIX, Irish INEX, Milano MIX is in the range of 50%-85%. The peering densities in the US-based IXPs are reportedly lower than the European, while the IXPs in Australia, New Zealand and far-east are slightly denser in terms of peering [2]. At the same figure we also plot, per-type peering density of content, access and network service providers, and, similarly to the results in Section 4.5, we observe no significant dependence between the type of an ISP and the peering density.

Per-member peering density. Another important property of the IXP is that different members have different peering policies. That means that those ISPs with restrictive peering policies are expected to peer with significantly lower number of ISPs than the members with an open peering policy. This property is reflected in our data. In Figure 3.6 we plot the histogram of per-member peering density (fraction of other members a member peers with) for all 84 ISPs that have participated in SIX. For the ISP density, we used the density of the last data sample in which the ISP was a SIX member. We can observe a bi-modal distribution of per-member peering density, indicating the groups with open and restrictive peering policies.

Peering link creation. Throughout the history of SIX, there have been 1711 pairs of ISPs peering between each other. For those pairs of ISPs that eventually start peering, how long does it take from the moment the newer of them appears at SIX until they engage into the peering (we call this value the *link creation time* - LCT)? In Figure 3.7 we depict the distribution of LCT for the 1711 peering pairs. We can observe that when the peering happens, it is usually created within a year from the appearance of the newer ISP, yet there are a dozen of pairs, that required more than 5 years of simultaneous existence to start peering. We stress that frequency of our sample points in time is one sample point in average per several months, implying that we are not capable to accurately identify the appearance of an ISP or a peering on sub-yearly time scales.




Figure 3.6: Histogram of peering density for all 84 (historic) members of SIX. For each ISP, the last peering density value is used.




Figure 3.7: Histogram of link creation times for 1711 peering pairs in the history of SIX.

ISP Name (ASN)	Type	AGR
aggregated	N/A	103.2%
t-com (AS6855)	Access	137.2%
gts (AS5578)	NSP	77.88%
vnet (AS29405)	Content	189.8%
orange (AS15962)	Access	143.4%
datacamp (AS39392)	Content	99.26%

Table 3.1: Type and Annualized Growth Rate (AGR) for aggregated traffic and current Top-5 ISPs.

Peering link removal. Dhamdhare and Dovrolis [7], analyzed AS-level customer-provider (CP) links over a 10-year period, and showed that these links appear and disappear quite often, citing as a major reason the cost optimization process customers perform when searching for the most cost-effective provider. The peering links, on the other hand, have a different business objective, and the reasons for CP link disappearance (other than the death of one of the ASes) are not present for the IXP peering. Therefore, one may guess that peering links are unlikely to be broken by one of the ISPs once they are created. Our data shows, that is indeed the case. Out of 1711 links, existing in SIX, only 20 link pairs de-peered. The other source of the peering link removal is the disappearance of one of the ISPs. The presented results question, at least for the particular situation of SIX, the relatively instability of peering links found in past projects [24]. As we mentioned overall 32 ISPs that participated in SIX do not participate anymore and indeed peering links they were engaged with, are not present anymore.

3.3 Traffic dynamics

Traffic growth. The growth of the Internet traffic, has been shown to follow an exponential pattern in the residential broadband networks [6], transpacific traffic [3] as well as in the inter-domain [17]. We observe a similar pattern in SIX, which can be seen from Figure 3.8 that depicts the traffic growth for SIX and top-5 (in terms of traffic volume) ISPs that are currently active in the exchange point. In contrast to residential and global inter-domain traffic that grow with *annual growth rate* (AGR) of around 40-50% [6, 15, 17], the AGR of the IXP traffic both in aggregate and individual ISP terms is much higher: around 100% or even more for some ISPs. This hints the growth of the relative fraction of the inter-domain traffic that is exchanged via peering at the IXP, and therefore a decay of the relative fraction of the inter-domain traffic that reaches end-customers via transit, a phenomenon consistent with the widely observed flattening-of-the-Internet trend [17].

Another interesting property of the growth is that it is not uniform among involved ISPs: the traffic of some ISPs grows quicker than for the others. Table 3.1 contains the AGR for top-5 active ISPs and the total SIX aggregate. Relatively wide range of AGR reveals significant differences in the growth of different ISPs. The exact reasons for such variability of AGR are out of scope of this work and will be part of our future work.

Remark: The annualized growth rates are obtained using the linear least-square fitting of the traffic growth curves in log scale.

Traffic per ISP type. As the Internet ecosystem matures, we are likely to expect the




Figure 3.8: SIX's aggregated traffic and the traffic of top-5 current members.

emergence of the specialized ISPs that consolidate around particular type of business. By looking in our data, we can see the emergence of one such class: content ISPs. Namely, until 2006, the content ISPs carried a very small fraction of the SIX traffic: under 10% in both directions. Since mid-2006, the number of specialized content ISPs doubled from 5 to today's 10; see Figure 3.4. The relative outbound traffic of those members, however, grew for an order of magnitude, from under 5% to over 40% as can be seen from Figure 3.9.

Traffic imbalance (and ISP consolidation). From Figure 3.9 one can notice that inbound to outbound traffic ratios vary significantly for different ISP types. We obtain the ISP type from the PeeringDB and manually for several ISPs that do not have an entry in their database. We stress that there is no sharp line between these three categories, since often each network may provide multiple services (e.g. a NSP offering access services). An ISP with a very uniform customer population is expected to have a very unbalanced in/out traffic; e.g. a fully residential ISP traffic is likely to be heavily inbound, while an ISP serving only content is likely to have heavily outbound traffic. Here we look at the evolution over time of the traffic imbalance which we define for ISP X at time t as:

$$IB(X, t) = \max \left(\frac{T_{inbound}(X, t)}{T_{outbound}(X, t)}, \frac{T_{outbound}(X, t)}{T_{inbound}(X, t)} \right)$$

Where $T_{inbound}(X, t)$ and $T_{outbound}(X, t)$ is the inbound and outbound traffic of the ISP X at time t , respectively. In Figure 3.10 we depict the evolution of the median ISP IB , as well as the 10th and the 90th percentile. We also plot the (traffic) weighted average of the imbalance across all ISPs present in SIX at time t :

$$WA_IB(t) = \frac{\sum_{X \in Active_t} T(X, t) \cdot IB(X, t)}{\sum_{X \in Active_t} T(X, t)}, \quad (3.1)$$

where $T(X, t) = T_{inbound}(X, t) + T_{outbound}(X, t)$ is the total traffic of ISP X at time t . From Figure 3.10 we can see a growing trend in the traffic imbalance, indicating the

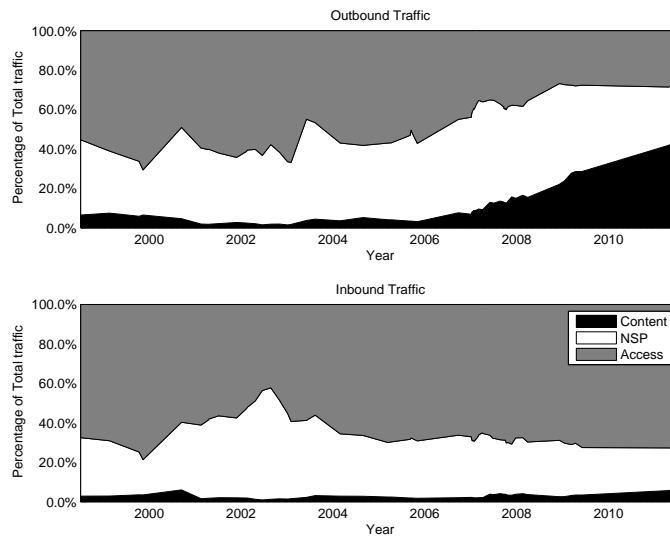


Figure 3.9: The percentage of total inbound/outbound traffic per ISP type.

increasing focus of the ISPs in particular end-customer groups. The fact that the weighted average is close to the 90th-percentile indicates that the large ISPs are more pronounced in such traffic imbalance, compared to the small and medium ISPs.

Traffic matrix dynamics. We can utilize an algorithm based on the gravity model (described in the Appendix) to estimate TM entries with reasonable accuracy. The expected errors of any TM estimation tool can be relatively large for a single origin-destination (O-D) pair. However, we do not analyze specific O-D pairs, but rather focus on aggregate statistics of the TM entry distribution, to draw the relevant conclusions. An invariant property of the SIX traffic is the skewness of the per-pair traffic volume distributions across time: few heaviest pairs, generate most of the SIX traffic. In Figure 3.11 we show the evolution of the fraction of the total traffic generated by the top-1%, top-10% and bottom-50% of the peering pairs. We can see that after initial transients, these quantities have remained relatively stable since 2003, indicating mentioned skewness of the per-pair traffic volume distributions.

Reasons for peering: cost-reduction or performance?. The most widely cited reason for creating a peering relationship between two ISPs is cost reduction: if the traffic between ISP *A* and ISP *B* is exchanged directly via peering, there is no need to be delivered via transit provider(s), hence the transit cost for both ISPs is expected to reduce. While such reasoning is indeed valid when the traffic volume between ISPs *A* and *B* is large, it becomes questionable when the traffic is small, since engaging into a peering relationship has a non-zero monetary cost (for legal agreements, staff, maintenance, etc.) associated with it. As shown in Figure 3.11, the sum of the traffic generated by half of the peering pair in SIX is relatively low and one may wonder whether these peerings really represent any economical benefits to ISPs. We use the history of the price of transit found in [23] (Figure 3.12b) to calculate the maximum potential savings obtained by the 50% lower peering pairs in time. In order to do this, we implement the gravity model algorithm described in the Appendix to estimate the median exchanged traffic between peering pairs of SIX (Figure 3.12a) and multiply it by the cost of transit. The results of this operation (depicted in Figure 3.12c)




Figure 3.10: The evolution of ISP traffic imbalance (IB), median, 10-th, 90-th percentile and the weighted average.

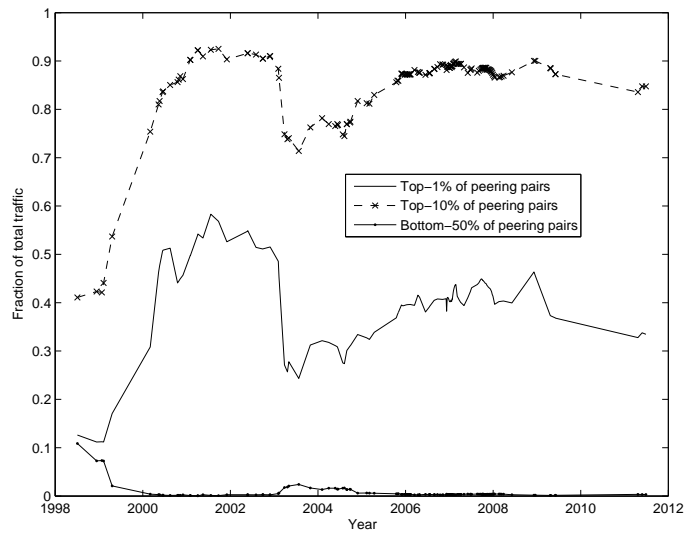


Figure 3.11: The fraction of the total traffic from the top-1%, top-10% and bottom 50% of peering pairs.

show that since 1998 the monetary gain for these peerings would represent savings of less than 10 USD, which is not enough for ISPs to justify these peering based on economical benefits.


3.4 Capacity dynamics

In order to exchange traffic at SIX and most other IXPs, each member needs to pay a monthly fee for each port it uses for traffic exchange. The prices of ports increase with the port capacity, but however are sub-additive, with price per *Mbps* going down as the higher port capacity is purchased. The reasons for upgrading from lower port capacity to a higher one can be different, but roughly speaking most of the upgrades happen because the member's traffic reaches port-utilization that is above some threshold (indicating the congestion). There are also several cases in which the port capacity becomes obsolete, and the member is moved to the new minimum capacity.


Information on network utilization and upgrades are almost always considered confidential by the ISPs. Literature often cites 50% rule-of-thumb for net upgrades [19, 29], but we are not aware on any empiric study across a set of diverse ISPs that evaluate such statements. The data we study here offers an opportunity to shed light on the upgrade practice and the port utilization in dozens of operational ISPs.

Port utilization. In Figure 3.13 we depict the median, the 10th and the 90th-percentile of peak utilization among all members of SIX at each time instance for which we have a SIX snapshot. We used the peak utilization as the maximum monthly utilization averaged in 2-hour slots, available in *mrtg* data. In the first couple of years, the SIX ports were very highly utilized with median members' peak utilization being greater than 50%. In the early 2000's, a significant increase in the port capacities occurred and brought the utilization of many members down. From 2009 to today we observe the increase in the utilization levels, yet still 90% of the members currently have the peak utilization of their port(s) under 60%. We also depict the evolution of the weighted average of the utilization, weighted with the traffic of the ISP, similarly to Eq. (3.1). We can see that, since 2004, most of the heavy ISPs are running their ports at a higher utilization than the median.


Port upgrades. In any ISP, increasing capacity of its networks is the most important mechanism for ensuring the high level of quality of service. Typically, the capacity upgrades in operational networks require a lot of planning and serious implementation efforts. In contrast, upgrading the port capacity at the IXP level is rather straightforward and can be done almost instantaneously. Hence, the data from SIX can offer insights on the operational practices of the involved ISPs. Figure 3.14 contains two histograms on the peak utilization at the moments of capacity-upgrades, one for the period from 1997-2004, and another for the period since 2004. In the earlier period, majority of upgrades happened because of the port overload, however since 2004, most ISPs upgrade their capacity early to avoid congestion of their SIX port. While most of the upgrades happen when the utilization reaches the range [40%-60%], there is still a non-negligible fraction of upgrades that happen outside of this range.



(a) Median exchanged traffic by peering pairs in SIX.



(b) Cost of transit in time [23]



(c) Savings for the median exchanged traffic in SIX

Figure 3.12: Median savings of peering for SIX in the last 14 years.




Figure 3.13: Per member port(s) utilization of SIX since 1998: median, 10th, 90th-percentile and weighted average.




Figure 3.14: The estimated utilization at the time of port upgrade.

3.5 Summary of Contributions

We highlight here some of the contributions presented in this Chapter:

- Internet Exchange Points often publish basic information that can be used to analyze the global network. We bring attention to the community of these valuable data that has not been sufficiently exploited for research purposes.
- We observe a bimodal distribution for the per-member peering density among the participating ISPs, a majority of them are very open for peering and the rest are very selective for peering. In contrast to customer-provider links that have been shown to change relatively frequently, AS links at the IXP are stable: once they are created they are very unlikely to be broken unless one of the ISPs disappears.
- We observe the recent explosion of the traffic generated by Content Distribution Networks (CDN) and verify with SIX the exponential growth of the traffic observed in several studies [5, 17, 3].
- Traffic matrix analysis reveals that the skewness of the per peering-pair traffic distribution: a majority of peering-pairs exchange very little traffic while a minority dominates the aggregate traffic. This finding implies that for most of the peering pairs the monetary gain of peering is minor which questions the common belief that the direct cost-reduction is the main reason for peering between the ISPs.

Chapter 4

Comparison of SIX and MIX

In this chapter we compare different characteristics of SIX and MIX and their members.

4.1 Datasets

For this chapter we use data from MILAN-IXP (MIX) and the Slovak IX (SIX). The dataset from MIX consists of traffic exchanged among all its members (i.e. the traffic matrix). For *each* pair of members, the total traffic in both directions is reported for forty eight 30-minute time slots during one weekday in 2008 (February 26). From SIX's dataset¹ we use the peering matrix and per-member traffic obtained for the end of February 2008. Table 4.1 summarizes the basic stats of these two IXPs.

In order to better compare the two IXPs, we needed an estimation of the traffic matrix for SIX. We used the data from MIX to develop and test an algorithm based on the gravity law to calculate a traffic matrix using per-member traffic stats and the peering matrix. The gravity law states that the amount of traffic T_{ij} flowing from network i to network j is proportional to a repulsion factor R_i of i and an attraction factor A_j of j . The description of the gravity model algorithm and its evaluation is included in the Appendix. We use the developed algorithm to calculate the attraction and repulsion and estimate the traffic matrix of SIX for this particular day.

¹The dataset from SIX is completely described in Section 3.1

	SIX	MIX
Members	46	61
avg throughput (daily)	5Gbps	10Gbps
peering density	69%	58%

Table 4.1: Basic stats on MIX and SIX in February 2008.




Figure 4.1: Exchanged traffic between pairs at SIX and MIX. Top 100 pairs shown.

4.2 Pricing model matters!

The two studied IXPs are similar in many ways: they host similar number of members, with similar traffic levels, and similar peering density, see Table 4.1. Additionally, both of them are the dominant national IXP, virtually all of their members are operating business in Italy or Slovakia respectively. However, there is one striking difference between them in terms of peering between ‘heavy’ pairs. In order to define the weight of the pair of members, we use the repulsion and attraction parameters to obtain a first-order approximation of the traffic flowing between the pairs of members in an IXP (for those that peer and also those that do not peer). We define the weight as the total traffic exchanged for a pair of members, that is:

$$ET_{ij} = (A_i \cdot R_j + R_i \cdot A_j)[Mbps]. \quad (4.1)$$

As we show in the Appendix, our gravity model algorithm provides a fairly accurate approximation for the ‘heavy’ pairs of the IXP. We extrapolate such reasoning for the pairs that do not peer, and use (4.1) as an estimate of traffic flowing between any pair of peers. In Figure 4.1 we depict the values of the exchanged traffic for top-100 pairs (ranked by the exchanged traffic). As we can observe from this figure, most of the pairs with high exchanged traffic at MIX do not peer, while this is not the case at SIX. There can be many factors that impact the decision on whether a member should exchange traffic with another member at the IXP or not. However, we believe that the critical factor that influences the peering decision between the heavy pairs is the financial one. Namely, the pricing model of the two IXPs are fundamentally different: SIX employs flat fee [28] (traffic-volume independent pricing) while MIX used usage based, 90th-percentile pricing [20]². Additionally, the price

²After 2008 the pricing model in MIX changed from usage-based to flat-fee. However, we did not see significant increase in the peering density between heavy pairs, which is partly due to the fact that once a private peering is created, there is virtually zero-cost for maintaining such peering.




Figure 4.2: Attraction vs Repulsion for members of SIX and MIX.

per *Mbps* in MIX case was prohibitively expensive, calling for non-IXP peering between heavy pairs (willing to peer).

4.3 Attraction and repulsion

Even though closely correlated, the repulsion/attraction parameters of an ISP are not directly derivable from the upstream/downstream traffic measurements. For example a heavy attractor/repulser may peer with a low number of peers, which can result in low amount of traffic exchanged. Therefore we use the repulsion/attraction parameters as indicators of the ISP demand (size) rather than the observed traffic itself.

In Figure 4.2 we plot attraction against repulsion for each member of MIX and SIX. As mentioned in Section 3.3, we used the *type* classification from PeeringDB [31] to distinguish between access, content and network service providers (NSP). From Figure 4.2 we can deduce two important observations. First, not surprisingly, the access networks typically have larger attraction/repulsion ratio while the content providers have the opposite property. And second, the distributions of repulsion and attraction parameters are very skewed, with a few members dominating the others in terms of size.

4.4 Traffic distribution

Another important property of the IXP traffic is the heterogeneity of the traffic volumes among the peering pairs. In other words, the IXP traffic is dominated by the several heaviest peering pairs. In Figure 4.3 we depict the fraction of the IXP traffic generated by the top- $x\%$ peering pairs. For MIX both the direct measurement and the gravity estimates are shown, while for SIX only the estimated values (from gravity model) are shown. We can observe that top-10% of peering pairs generate 85-90% of traffic and top-1% of peering pairs generate 30-40% of the IXP traffic. A different perspective of the distribution of traffic among peering pairs is depicted in Figure 4.4, which shows a histogram of the traffic exchanged by peering pairs in both MIX and SIX. The traffic distribution of both IXPs is quite similar and in both

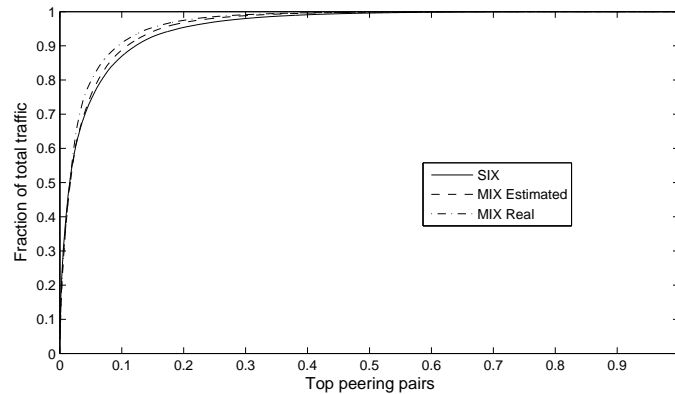


Figure 4.3: Fraction of total traffic generated by heaviest peering pairs at SIX and MIX.




Figure 4.4: Traffic distribution for SIX and MIX in 2008.

cases half of the peering pairs exchanged less than $300Kbps$, an amount of traffic that, based on the cost of transit in 2008[23], would only represent ISPs with savings of a few USD.

4.5 ISP-type

Different types of networks peer through IXPs. We use the information obtained from Peering DB to classify the members of SIX and MIX into three categories: access, content and network service providers (NSP). We stress again that this classification is relatively crude as some members (particularly at MIX) offer multiple services and cannot be clearly classified in one of these three types. Table 4.2 lists the fraction of the total IXP traffic flowing between different member types. Not surprisingly, we observe that access networks are sinks to between 50% of 60% of the total traffic. Content providers were for both IXPs in 2008 the source of less than 16% of the traffic, however, in recent years the inter-domain

		To		
		Access	Content	NSP
SIX	From			
	Access	22.3%	1.4%	12.5%
	Content	10.7%	0.2%	4.4%
	NSP	35.1%	2%	11%
MIX	From			
	Access	24.6%	1.6%	12.9%
	Content	7.7%	0.7%	6.3%
	NSP	18.1%	2.3%	25.4%

Table 4.2: Relative traffic exchanged between the different types of members for SIX and MIX.

		To		
		Access	Content	NSP
SIX	From			
	Access	76.3%	78%	63%
	Content	78%	70%	60%
	NSP	63%	60%	51.4%
MIX	From			
	Access	52.4%	67.4%	44%
	Content	67.4%	80.9%	52%
	NSP	44%	52%	43.3%

Table 4.3: Density of peering between types of members for SIX and MIX.

traffic flowing from Content Providers has increased[17], this particular point is further developed for SIX in Section 3.3. The similarities between SIX and MIX in Table 4.2 are evident and the few inconsistencies could be explained by the fact that several large Italian access and content providers are represented at MIX by another company which is an NSP.

One may wonder whether peering density varies between different types of members (ie. whether access ISP is more likely to peer with content ISP than NSP). We summarize our findings in Table 4.3, from which we can observe lack of evidence for significant difference of peering density between different member types.

4.6 Peering intensity vs. member size

Another question that arises in the analysis of IXP is whether the member's traffic volume affects the number of peerings it engages at the IXP, or not. We do not observe such dependence at neither MIX nor SIX. In Figure 4.5 we plot the number of peerings for each member of the IXPs, against its size. Apparently, the peering intensity (the number of peering relationship the member engages) of a member do not appear to depend on its size, in spite of large differences in the member traffic volumes (multiple orders of magnitude).

4.7 Summary of Contributions

We highlight here some of the contributions presented in this chapter:

- We study basic characteristics of the traffic exchanged at IXPs. In particular, we show

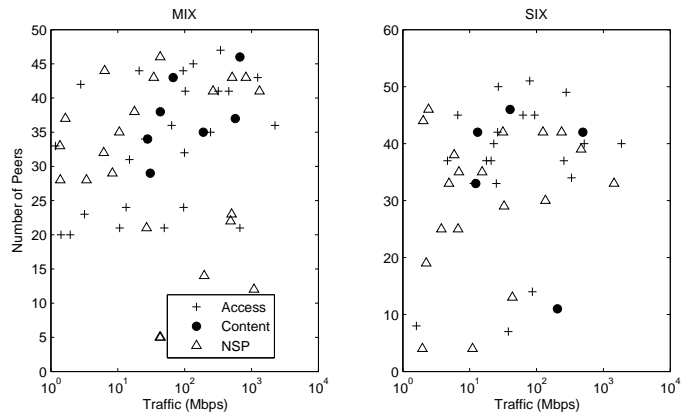


Figure 4.5: Peering intensity vs. traffic of ISPs peering at MIX and SIX.

that the traffic distribution per peering pairs in both IXPs is very similar and dominated by a few heavy peering pairs (Sec. 4.4).

- We demonstrate strong correlation between the pricing model employed by the IXP and the peering density among the heavy peers: an IXP with a flat-fee pricing is significantly more likely to see heavy pairs peer compared to an IXP with usage-based pricing (Sec. 4.2).
- We do not observe significant correlation between the peering intensity (the number of peerings) of a member and its type nor its traffic volume (Sec 4.5 and 4.6).

Chapter 5

Conclusions

Although located in the periphery of the Internet, IXPs are an important part of the Internet ecosystem, currently mediating the exchange of 15-20% of the inter-domain traffic. With the trend of flattening-the-Internet [8] (i.e. bypassing the transit providers with direct peering), the relative importance of IXPs, and peering in general, is likely to increase. In this paper we study detailed peering and traffic data from two European IXPs (Slovakian and Milan) and draw several important conclusions. We compare both IXP and demonstrate that the peering (AS-level) topology within the IXP is impacted by the pricing model of the exchange point. By examining the evolution of SIX, we found interesting facts in regard to the peering and traffic trends at the IXP. We discovered that once it has been created, a peering between two ISPs is very unlikely to disappear unless one of the members leaves the IXP. By quantifying the peering density (fraction of other ISPs a member peers with) of SIX members we observe a bimodal distribution: majority of members have high peering density, a minority have low peering density with almost no members in the middle. The emergence of content ISPs and heavily-unbalanced ISPs (indicating the ISP consolidation) are two major factors characterizing the last 5 years of SIX operation. We believe that the data available from IXPs can be of great use when understanding the topology of the Internet.

Appendix: Gravity model and traffic matrix estimation

In this Appendix we quantify to which extent the so called gravity law applies to IXP traffic. The gravity law says that the amount of traffic T_{ij} flowing from network i to network j can be approximated as:

$$T_{ij} = \frac{R_i \cdot A_j}{F_{ij}}, \quad (5.1)$$

where R_i is the repulsion factor of i , A_j is the attraction factor of j and F_{ij} is the friction parameter typically inversely proportional to the “distance” (in some metric space) between i and j . Since the two studied IXPs cover relatively small geographic area (a country), the friction parameter is unlikely to depend on the geographic distance, and in this paper we take it to be constant: $F_{ij} = (Mbps)^{-1}$.

In case the peering graph G is the full mesh (i.e. every member peers with every other member) computing the repulsion and attraction factors from the upstream/downstream traffic stats is trivial: they are equal (up to the multiplicative constant) to the corresponding upstream and downstream traffic volume, respectively. However, normally the peering graph is not a full mesh and the computation of repulsion and attraction is not straightforward. Namely in that case

$$A_j \sum_{(i,j)\text{peer}} R_i = d_j, \text{ for all } j, \quad (5.2)$$

$$R_i \sum_{(i,j)\text{peer}} A_j = u_i, \text{ for all } i, \quad (5.3)$$

where u_k and d_k is the upstream and downstream traffic volume of network k , respectively. It is not hard to see that the system (5.2)-(5.3) is under-determined: if $(A_1, \dots, A_N, R_1, \dots, R_N)$ is a solution to (5.2)-(5.3), then for any positive scalar α the vector $(\alpha A_1, \dots, \alpha A_N, R_1/\alpha, \dots, R_N/\alpha)$ solves (5.2)-(5.3) as well. To enforce the uniqueness of repulsion and attraction parameters we require the following normalization constraint to be met:

$$\sum_{k=1}^N A_k = \sum_{k=1}^N R_k, \quad (5.4)$$

Thus, the attraction and repulsion factors are the solution of the nonlinear system: (5.2)-(5.4). Pseudocode in Figure 5.1 provides a simple iterative method for solving it. We

```

1 SolveGravity(u, d, Peering_Matrix)
2 for  $k = 1 : N$ 
3    $A_k(0) = \sqrt{d_k}$ 
4    $R_k(0) = \sqrt{u_k}$ 
5 endfor
6  $t = 0$ 
7 until convergence
8   for  $k = 1 : N$ 
9      $\hat{A}_k(t+1) = \frac{d_k}{\sum_{s \text{ peer with } k} R_s(t)}$ 
10     $\hat{R}_k(t+1) = \frac{u_k}{\sum_{s \text{ peer with } k} A_s(t)}$ 
11  endfor
12   $\eta = \sqrt{\frac{\sum_{k=1}^N \hat{A}_k(t+1)}{\sum_{k=1}^N \hat{R}_k(t+1)}}$ 
13   $\mathbf{A}(t+1) = \frac{\hat{\mathbf{A}}(t+1)}{\eta}$ 
14   $\mathbf{R}(t+1) = \eta \hat{\mathbf{R}}(t+1)$ 
15   $t = t + 1$ 
16 repeat

```

Figure 5.1: The pseudocode for solving gravity.

compare the estimated TM, obtained by solving the system (5.1)-(5.4), with the real TM in Figures 5.2 and 5.3. In the former, we plot the estimated average throughput against the real throughput for all pairs of members that peer at MIX. The estimation is reasonably accurate for ‘heavy’ pairs, but does not deliver precise results for the pairs that exchange lows amounts of traffic. Figure 5.3 compares the cumulative distribution of both TM for the number of pairs with less than certain amount of traffic. From both Figures, it can be stated that even if the system does not provide a highly accurate estimation on the individual member pairs, it approximates well the general distribution of traffic among them. It is likely that more precise TM estimation methods could be developed, probably by optimizing existing techniques[11, 33, 34], however, accurate estimation of TM is out of scope of our work. We rather seek for a simple method that would allow us to explicitly estimate the IXP TM with reasonable accuracy, which is indeed achieved by the gravity method tailored for the IXP case.




Figure 5.2: Real throughput compared to throughput estimated using the gravity model for MIX.

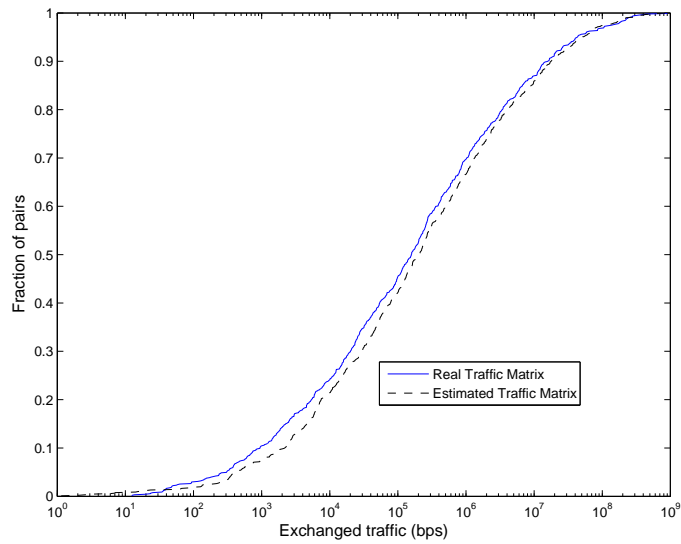


Figure 5.3: Cumulative distribution of exchanged traffic for peering pairs in both the estimated TM and the real TM of MIX.

References

- [1] M. Ahmad and R. Guha. Understanding the impact of internet exchange points on internet topology and routing performance. In *Proceedings of the ACM CoNEXT Student Workshop*, page 18. ACM, 2010.
- [2] B. Augustin, B. Krishnamurthy, and W. Willinger. Ixps: mapped? In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 336–349. ACM, 2009.
- [3] P. Borgnat, G. Dewaele, K. Fukuda, P. Abry, and K. Cho. Seven years and one day: Sketching the evolution of internet traffic. In *INFOCOM 2009, IEEE*, pages 711–719. IEEE, 2009.
- [4] CESNET. Traffic stats., 2010.
- [5] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 207–218. ACM, 2006.
- [6] K. Cho, K. Fukuda, H. Esaki, and A. Kato. Observing slow crustal movement in residential user traffic. In *Proceedings of the 2008 ACM CoNEXT Conference*, page 12. ACM, 2008.
- [7] A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the internet ecosystem. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement (IMC)*, pages 183–196. ACM, 2008.
- [8] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of the 6th International Conference (Co-NEXT)*, page 21. ACM, 2010.
- [9] A. D’Ignazio and E. Giovannetti. Antitrust analysis for the internet upstream market: A border gateway protocol approach. *Journal of Competition Law and Economics*, 2(1):43, 2006.
- [10] A. D’Ignazio and E. Giovannetti. Asymmetry and discrimination in internet peering: evidence from the linx. *International Journal of Industrial Organization*, 27(3):441–448, 2009.

- [11] V. Erramill, M. Crovella, and N. Taft. An independent-connection model for traffic matrices. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement (IMC)*, pages 251–256. ACM, 2006.
- [12] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *Proceedings of the 9th international conference on Passive and active network measurement (PAM)*, pages 1–10. Springer-Verlag, 2008.
- [13] E. Gregori, A. Improta, L. Lenzini, and C. Orsini. The impact of ixps on the as-level topology structure of the internet. *Computer Communications*, 2010.
- [14] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: a framework for discovering missing links in the internet topology. *IEEE/ACM Transactions on Networking (TON)*, 17(2):391–404, 2009.
- [15] C. Index. Forecast and methodology, 2009–2014. *White paper, CISCO, June, 2*, 2010.
- [16] INEX. Peering matrix, 2010.
- [17] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. *ACM SIGCOMM Computer Communication Review*, 40(4):75–86, 2010.
- [18] W. Machine. The internet archive. <http://wayback.archive.org/web/>.
- [19] J. Martin and A. Nilsson. On service level agreements for ip networks. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 855–863. IEEE, 2002.
- [20] MIX-IT, 2010.
- [21] I. Nakagawa, H. Esaki, and K. Nagami. A design of a next generation ix using mpls technology. In *saint*, page 238. Published by the IEEE Computer Society, 2002.
- [22] NIX, 2010.
- [23] B. Norton. Internet transit prices historical and projected., August 2011.
- [24] R. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of internet as topology. *ACM SIGCOMM Computer Communication Review*, 37(4):313–324, 2007.
- [25] S. Radovic. European internet exchange association. www.euro-ix.net, 2011.
- [26] SANET. Traffic stats, 2010.
- [27] Y. Shavitt and E. Shir. Dimes: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5):71–74, 2005.
- [28] SIX. Peering matrix, 2010.

-
- [29] T. Telkamp. Traffic characteristics and network planning. In *Proc. Internet Statistics and Metrics Analysis Workshop*, 2002.
 - [30] VIX. Peering matrix, 2010.
 - [31] P. Website.
 - [32] K. Xu, Z. Duan, Z. Zhang, and J. Chandrashekar. On properties of internet exchange points and their impact on as topology and relationship. *NETWORKING 2004, Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications*, pages 284–295, 2004.
 - [33] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg. Fast accurate computation of large-scale ip traffic matrices from link loads. In *ACM SIGMETRICS Performance Evaluation Review*, volume 31, pages 206–217. ACM, 2003.
 - [34] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. An information-theoretic approach to traffic matrix estimation. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 301–312. ACM, 2003.