# On Properties of Internet Exchange Points and Their Impact on AS Topology and Relationship*


Kuai Xu[1], Zhenhai Duan[2], Zhi-Li Zhang[1], and Jaideep Chandrashekar[1]

[1] Department of Computer Science, the University of Minnesota,
200 Union St SE, Minneapolis, MN 55455, USA
{kxu, zhzhang, jaideepc}@cs.umn.edu

[2] Department of Computer Science, the Florida State University,
Tallahassee, FL 32306, USA
{duan}@cs.fsu.edu



**Abstract.** Internet eXchange Points (IXPs) are one of two primary methods for Autonomous Systems (ASes) to interconnect with each other for exchanging traffic and for global Internet reachability. This paper explores the properties of IXPs and their impact on the AS topology and AS business relations using Scriptroute and Skitter traceroute probes, BGP routing archives and other data. With these datasets we develop an algorithm to discover IXPs and infer ASes that participate at these IXPs. Using the discovered IXPs and their inferred AS participants, we analyze and characterize the properties of IXPs and their participants such as size, geographical locations. We also investigate the impact of IXPs on the global AS topology and business relations between ASes. Our study sheds light on the Internet interconnection practices and the evolution of the Internet, in particular, the potential role IXPs play in such evolution.


## 1 Introduction

The Internet is a collection of interconnected Autonomous Systems (ASes), which are independently administrated by various Internet Service Providers (ISPs), universities and other institutions. In general, two ASes interconnect with each other either via a dedicated circuit or through an IXP (See Fig. 1 for an illustration). In the former case, depending on their business agreement, either one party pays the full cost of the circuit for interconnection with the other party, or both parties pay half of the circuit cost to exchange traffic. Clearly via a dedicated circuit, only one pair of ASes can interconnect with each other.

In contrast, an IXP is a *shared interconnection infrastructure*, where multiple ASes can interconnect with one another through switches and routers at the IXP. Each AS only bears the cost of a circuit from its own premises to the IXP, in


* This work was supported in part by the National Science Foundation under the grants ANI-0073819 and ITR-0085824. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

addition to facility fees paid for the IXP. More importantly, through the same infrastructure, an AS can interconnect with some or all of the other participants at the IXP, subject to mutual business agreements. Therefore, IXPs provide a cost-effective model for different ASes to interconnect. Indeed, hundreds of IXPs have been deployed around the world, and many ASes participate at such IXPs for Internet interconnection.


(a) A dedicated circuit (b) An example of IXP

**Fig. 1.** AS interconnections

Although IXPs play an important role in the development of the Internet, many of their properties and their impact on the evolution of the Internet such as its topology and AS relationship, have not been systematically studied. For example, how are IXPs distributed on the Internet and how are they used in the different parts of the world? How do different ASes value and consequently utilize IXPs? Why do some ASes participate at an IXP, while others do not? How does an AS participating at an IXP decide which other AS to interconnect and what business relation to enter into? We believe that a better understanding of IXP properties and their impact on the global Internet will shed light on the Internet interconnection practices and the evolution of the Internet.

Towards this goal, we study the properties of IXPs such as geographical locations and size distributions, and investigate their impact on AS topology and relationship. To identify participants at an IXP, we design an inference algorithm based on traceroute probes and reverse DNS lookups. Using this algorithm, we find that IXPs appear to be more prevalent in Europe than in other continents. We also find that *transit* ASes (i.e., ASes with customer ASes) participating at IXPs have a higher percentage of peering relationship among themselves than, say, the average AS on the Internet. This reflects the fact that transit ASes generally peer more aggressively with other transit ASes to reduce transit cost, while stub ASes typically participate at an IXP primarily for transit service.

To investigate the impact of IXPs on AS topology and relationship, we decompose the AS topology into a *ring structure* based on the "distance" of an AS to the *Internet core* (or I-core in short). Using this ring classification, the (current) Internet AS topology can be decomposed into six rings. The I-core ASes normally only have a small percentage of IXP neighbors relative to their degree. This is not surprising, as the I-core ASes are known to prefer interconnection with their customer ASes via dedicated circuits for obvious economic reasons. However, a large percentage of transit ASes in Ring 2 and Ring 3, especially in Europe and Asia, have a significant number of IXP connections. These Ring 2 transit ASes more aggressively peer with each other to reduce transit costs and their reliance on the I-core, as is evidenced by the much higher percentage of peering relations among themselves.

The remainder of the paper is structured as follows. In Section 2, we describe the data sources and methodology used in the paper. We study the properties of IXPs in Section 3, investigate the impact of IXPs on AS topology and relationship in Section 4. The paper is concluded in Section 5.

## 2 Data Sources and Methodology

### 2.1 Data Sources

In order to have a more comprehensive set of *BGP route views*, we collected BGP routing tables from three vantage points, namely, the Route-views project [11], RIPE NCC [13], and a border router at the University of Minnesota and merged them to get an aggregated BGP routing table.

We use the *public IXP database* available at Packet Clearing House (PCH) [12] for some basic information about IXPs. For each IXP, the database records geographical location (e.g., city, country), name, status (e.g., active, planning, defunct), IP address blocks, IP version (e.g., IPv4, IPv6), and some other attributes. In this paper, we mainly study the properties of *active layer-3 IXPs* supporting IPv4, where there are 148 such IXPs.

We use *traceroute* to discover IXPs and identify their AS participants. The data is collected from two sources, Scriptroute project [15] and Skitter project [1]. Scriptroute provides a measurement infrastructure for us to launch traceroute probes from more than two hundred nodes to a large set of sample destinations on the Internet, while Skitter project deploys about 20 monitors periodically sending out traceroute to end hosts on the Internet. The traceroute data were collected between 11/02/2003 and 01/16/2004.

### 2.2 Inferring IXP Participants

As observed in [9, 7], if the forward path of a traceroute probe traverses an IXP, the IP addresses of the routers (or more precisely their interfaces) connecting to the IXP's switching facility are also recorded. Based on this intuition, we develop an IXP inference algorithm to infer the participants at an IXP. First of all, we filter out anomalous probes, such as those containing AS path loops. Then, we check if a forward IP path contains the IP address associated with an IXP by searching the IP address blocks in the IXP database. If a traceroute hop maps to an IXP, we consider the ASes before and after the hop as potential participants at the IXP. If the mapping fails, or if an IP address is mapped to multiple origin ASes, we do not consider them as IXP participants. Although [14] also describes a similar approach, it does not discuss several important issues such as multiple origin AS conflicts and AS loops [9] in traceroute results.

To obtain a larger set of IXP participants, we use another inference mechanism based on reverse DNS name lookups. The border router of an AS present at an IXP has an interface directly connected to the IXP switch facility. If the interface is associated with a DNS name, we will have a high probability to discover the participating AS from the domain name of the interface. Based on this, we design a simple scheme to identify more participants at an IXP. Given the IP address block of an IXP, we conduct brute-force reverse DNS lookups for the IP addresses within this block. If a reverse DNS lookup is successful, we extract the domain name from the host name. By looking up the domain name via DNS query, we obtain an IP address from the domain. We then map this IP address to an AS number, which is considered as a *potential* participant at the IXP. An AS inferred from this approach is labelled as a "true participant" at the IXP if the corresponding interface is alive by active probing using ping or traceroute.

### 2.3 Identifying IXP Edges and Inferring AS Relationship

An IXP provides a shared physical infrastructure for participating ISPs to interconnect. However, it is not necessary that every participant at an IXP will interconnect with every other. In other words, how a participant connects to others at the IXP depends on its mutual agreements, e.g., provider-customer or peering, with the others. We now briefly describe how to determine whether there is an interconnection between two ASes and how we infer their business relation. First we need to define some notation that will be used throughout the paper. We refer to the topology graph constructed *directly* from the aggregate BGP routing table as the AS topology graph, where each node represents an AS, and there is an edge between two ASes if they are neighbors on any AS path in the routing table. From the aggregate BGP routing table we constructed an AS topology graph with 16,014 nodes and 32,748 AS edges. Note that from traceroute measurements we may discover some ASes interconnections ("edges") at the IXPs that are not observed in the AS topology graph. We refer to them as *hidden (IXP) edges*. To infer the AS relation between two ASes, we apply the inference algorithm proposed by Gao [5] on the AS topology graph. In the remainder of the paper, we classify an AS as either a *transit* or *stub* AS based on its relations with its neighboring ASes. If an AS does not have any customer, it is a stub AS; otherwise, it is a transit AS.

## 3 Properties of IXPs and Their AS Participants

In this section we characterize the properties of IXPs in terms of their size, geographic location and hidden IXP edges. We further use three IXPs located in Asia, Europe and North America as representative examples to illustrate the relations between the IXPs and their AS participants and among the participants.

### 3.1 IXP Characteristics

Using the methodology described in the previous section, we discover 82 IXPs (referred to as *probed* IXPs) in the traceroute probes out of the 148 listed in the PCH database. We define the *size of an IXP* as the number of ASes participating at the IXP. Fig. 2(a) shows the *inferred* size distribution of the 82 IXPs, ranked by their size. Note that because traceroute probes may not traverse all participants of every IXP, the actual size of IXPs can be larger than what is reported here. In general, for IXPs close to traceroute monitors, their participants tend to be well discovered, as a large number of random probes generated by the traceroute monitors are more likely to traverse every single physical connection between the hosting ASes and their IXP neighbors. Hence, there is *bias* towards the IXPs close to traceroute monitors in terms of their inferred size [8]. To validate the coverage of the probed IXP participants, we use public participant lists released by six IXPs[3] on their Web sites. On average, 77.8% participants at these IXPs are discovered by our inference algorithm. Note that it is possible that we may not obtain such high coverage for the other IXPs due to limited traceroute probes.

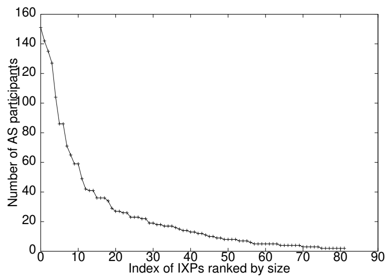


(a) IXPs size distribution

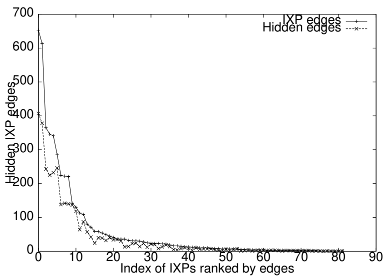


(b) Hidden IXP edges

**Fig. 2.** The size distribution of probed IXPs and hidden IXP edges

The 82 probed IXPs are located in five continents on the world: 7 in Asia (AS), 34 in Europe (EU), 35 in North America (NA), 3 in Oceania (OC), and 3 in South America (SA). It is not surprising to observe that North America has the largest number of IXPs as it contains more than 50% ASes on the Internet. However, 34 of 82 (41.5%) probed IXPs are from Europe, although it contains only 25% ASes on the Internet. These observations confirm the common belief that IXPs are a more prevalent method of Internet interconnection in Europe than other continents. Moreover, these statistics suggest that our inference algorithm is adequate to capture the characteristics of IXP distributions on the world.

[3] They are AMS-IX, DE-CIX, HKIX, JPIX, LINX-Extreme and LINX-Foundry.

Recall that we may discover some IXP edges that are not observed in the AS topology graph, and such an edge is called a *hidden* (or *missing*) edge. This may happen, for example, if ASes $A$ and $B$ have a *peering* relation at IXP $X$, and there is no BGP view data (i.e., BGP routing tables) from $A$, $B$ and their respective customers. Fig. 2(b) shows the hidden IXP edges as well as the total number of IXP edges for the probed IXPs found in the traceroute dataset. We see that at 61 IXPs, 50% or more of the IXP connections among the AS participants are missing from the AS topology constructed from the BGP routing data alone. These results quantify the observation in [2] that a rich set of IXP edges are not shown in BGP-derived AS topology graph.

### 3.2 IXPs and Their AS Participants: A Detailed Look

In this section we study the role IXPs play in Internet interconnections and their influence on the relations among the AS participants by examining three representative IXPs in more detail. The three IXPs, PAIX, LINX-Foundry, and JPIX, are located in North America, Europe, and Asia, respectively. We classify AS participants at each IXP into *transit* and *stub* ASes. Table 1 lists the number of transit and stub ASes in these IXPs. We see that more than 70% of ASes participating at PAIX and LINX-Foundry are transit ASes, while JPIX has significantly more stub ASes (38%). These results suggest that these IXPs could potentially play different roles in AS interconnections among their AS participants by region, as the economics and interconnection structure are different and market conditions vary.

**Table 1.** AS types of participating ASes of selected IXPs

| IXP Name | Size | # Transit ASes(percentage) | # Stub ASes(percentage) |
|---|---|---|---|
| PAIX | 135 | 94 (70.0%) | 41 (30.0%) |
| LINX-Foundry | 151 | 119(78.8%) | 32 (21.2%) |
| JPIX | 71 | 44 (62.0%) | 27 (38.0%) |

We also study the business relations among IXP participants and observe that most transit ASes participate at IXPs to exchange traffic with other transit ASes, thus entering into a peering agreement with each other. Some are also there to provide transit service to stub (or other smaller transit) ASes. All stub ASes participate at IXPs to buy (less expensive) transit service from transit ASes at the IXPs. There are also slightly different roles for the IXPs: PAIX and LINX-Foundry are examples of a "peer traffic exchange" where mostly transit ASes participate to exchange traffic among themselves, with a relatively small number of ASes present to buy transit service. Whereas, JPIX seems to be an example of "transit service exchange" where many transit ASes not only participate to exchange traffic among themselves, but also provide transit service for stub (or smaller transit) ASes.

## 4 Impact of IXPs on AS Topology and Relationship

In this section, we first propose a top-down approach to decompose the Internet into a ring structure and classify the ASes into the "rings". For each ring, we further discuss the impact of IXPs on their AS interconnections and relationship.

### 4.1 IXPs and AS Topology

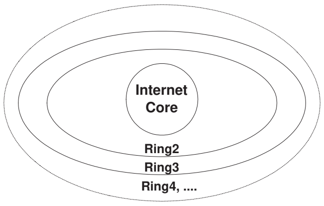


(a) The ring structure

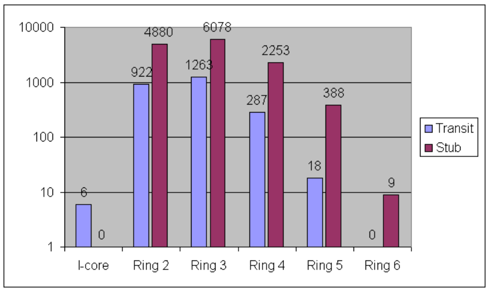


(b) The size of Internet core and rings

**Fig. 3.** The ring structure of AS topology and their breakdown

**The Ring Structure of AS Topology** It is well-known that the Internet *AS topology* has an inherent *hierarchical* structure, arising from the interconnection relations among the ASes [6, 17, 16]. At the top, there is a collection of biggest ASes that *peer* with each other to exchange traffic with no monetary settlement. In other words, none of them carry transit traffic for another AS in the collection. Hence by definition, they form a full *peering* mesh. We will refer to this collection of ASes at the top of the AS hierarchy as the *Internet core*, or *I-core* in short. The rest of the ASes in the Internet rely on one or more ASes in the I-core to reach part of the Internet directly or indirectly.

In this study, we use the I-core identified in [3, 10, 18]. The I-core consists of the six largest ISPs that form a peering mesh, AS701(UUNET), AS1239(Sprint), AS7018(AT&T), AS3356(Level 3), AS209(Qwest) and AS3561(Cable & Wireless)[4]. We classify the rest of ASes based on their "distance" – the number of AS hops – from the I-core. We call the collection of ASes that are directly connected (i.e., one hop away) to one or more ASes in the I-core the *Ring 2*; the collection

[4] The membership of the I-core has been changing as the Internet evolves. For instance, Genuity was used to be a member of the I-core. After being acquired by Level 3, it is no longer part of the I-core. The ASes in the I-core in our study are those identified during our data collection period.

of ASes that are two hops away from the I-core (i.e., they can only reach an AS in the I-core via an AS in Ring 2) as *Ring 3*; and so forth. Using this ring classification, the Internet AS topology can be decomposed into six rings, as schematically depicted in Fig. 3(a).

To understand the role that IXPs play in the interconnections and business relations among ASes, we further divide each Ring $i$, $i = 2, 3, 4, 5$, into stub ASes (referred to as *Ring i stub*) or transit ASes (referred to as *Ring i transit*). Fig. 3(b), with y-axis in log-scale, shows the number of transit ASes and stub ASes in each ring. The majority of ASes are in the Ring 2 (922 transit, 4880 stub) and Ring 3 (1263 transit, 6078 stub), comprising more than 80% of the total number of ASes in the current Internet. The I-core (transit ASes only) and Ring 6 (with 9 stub ASes) are the two smallest. Although this ring structure is very similar to various Internet topology models proposed in the literature, such as [6, 17, 16], there are two important difference between the ring structure and other models. First, the ring structure decomposes the Internet from two dimensions, namely, topology and business relations. Second, we choose the I-core based on the Internet peering practice [3, 10, 18], instead of simply picking a set of fully-connected "tier-1" ISPs[5].

**IXPs and the I-core ASes** If an AS connects to an IXP, we refer to such a connection as an *IXP connection*, and the neighboring ASes thus connected via the IXP as *IXP neighbors*. By examining the IXP neighbors of the ASes in the I-core, we see that all the ASes are connected to a few of ASes via some IXPs. However, with the exception of AS3356, the number of ASes connected via IXPs relative to their degrees are fairly small, and the percentage of IXP neighbors are less than 2.5%. Hence the majority of ASes are connected to the I-core via "private peering", i.e., *dedicated circuits*. Although AS3356 has a relative high percentage (6.8%) IXP neighbors worldwide, nearly 75% of such IXP neighbors are located in Europe. In other words, Europe is the main location where AS3356 exchanges traffic with these neighbors through IXPs.

**IXPs and Ring 2 ASes** Using the traceroute probe data, we have found that a total of 1083 ASes are participating at at least one IXP. Among these 1083 ASes, 394 are from Ring 2 and 593 are from Ring 3. Among these 394 Ring 2 ASes, over 75% (300) are in Ring 2 transit, the rest (94) are in Ring 2 stub. Fig. 4 shows the number of IXP connections for the 300 ASes in Ring 2 transit and the 94 ASes in Ring 2 stub respectively, ranked by the decreasing number of IXP connections. We see that more than half of the ASes with IXP connections in Ring 2 transit are present at at least two IXPs, about 68 ASes are present at five or more IXPs. In contrast, only one out of 94 ASes with IXP connections in Ring 2 stub is present at five IXPs, while most of them are participating at only one IXP. This is consistent with the nature of transit and stub ASes, as transit ASes are more likely to participate at one or more IXPs to peer and exchange traffic with other transit ASes, while stub ASes participate at an IXP mostly for gaining transit service from one of the transit ASes at the IXP.

[5] Typically, the definition of "tier-1" is an ISP who does not buy any transit from others. Due to marketing purposes, some non tier-1 ISPs claim to be "tier-1".

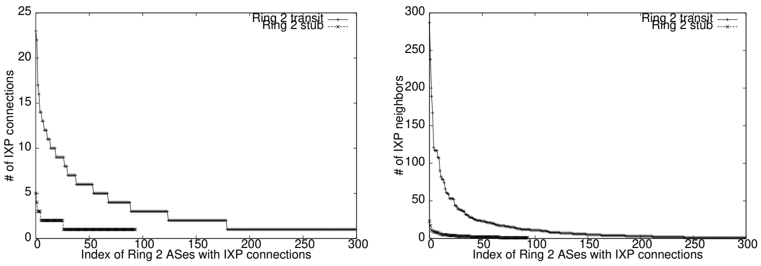


(a) # of IXP connections (b) # of IXP neighbors

**Fig. 4.** IXP connections and neighbors of Ring 2 ASes

Table 2 shows the *geographic* break-down[6] of the ASes with IXP connections in Ring 2 as compared to that of all ASes in Ring 2. The large majority of ASes in Ring 2 are in North America, as most I-core ASes are in the United States. However, among these ASes, only a small percentage of the Ring 2 ASes are found to have IXP neighbors, i.e., connected to them via IXPs. In contrast, Europe, Asia as well as other regions have far fewer Ring 2 ASes, but they have a significantly higher percentage of Ring 2 ASes with IXP neighbors. Given that we tend to *underestimate* the number of IXPs and IXP neighbors in regions such as Asia where there are only a few traceroute monitors, due to the "biased" and partial data we have, the percentage of ASes with IXP neighbors can be considerably higher. These results illustrate that connection via IXPs is far more prevalent in Asia, Europe and other regions than in North America.

We also study the *degree distribution* of Ring 2 ASes and find out that Ring 2 ASes (especially transit ASes) *with* IXP connections tend to have higher degree than those *without* IXP connections. This finding indicates the strong influence of IXPs on the degree of ASes *with* and *without* IXP connections for ASes in Ring 2 transit and stub.

**Table 2.** Locations of Ring 2 and Ring 3 ASes with IXP neighbors

| Ring | AF | AS | EU | NA | OC | SA | Total | Ring | AF | AS | EU | NA | OC | SA | Total |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Ring2(R2)<br>(R2 Transit) | 5<br>(4) | 118<br>(83) | 406<br>(217) | 5242<br>(595) | 7<br>(4) | 24<br>(19) | 5802<br>922 | Ring3(R3)<br>(R3 Transit) | 184<br>(14) | 1385<br>(299) | 2196<br>(458) | 3098<br>(387) | 227<br>(54) | 251<br>(51) | 7341<br>1263 |
| R2 w/ IXPs<br>(R2 Transit) | 0<br>(0) | 43<br>(40) | 193<br>(148) | 153<br>(107) | 2<br>(2) | 3<br>(3) | 394<br>300 | R3 w/ IXPs<br>(R3 Transit) | 0<br>(0) | 158<br>(93) | 261<br>(148) | 136<br>(64) | 17<br>(10) | 21<br>(15) | 593<br>330 |
| Percentage (%)<br>(R2 Transit) | 0.0<br>(0.0) | 36.4<br>(48.2) | 47.5<br>(68.2) | 2.92<br>(18.0) | 28.6<br>(50.0) | 12.5<br>(15.8) | 6.8<br>32.5 | Percentage (%)<br>(R3 Transit) | 0.0<br>(0.0) | 11.4<br>(31.1) | 11.9<br>(32.3) | 4.4<br>(16.5) | 7.5<br>(18.5) | 8.4<br>(29.4) | 8.1<br>26.1 |

**IXPs and Ring 3 ASes** Among these 1083 ASes with IXP connections, 593 are from Ring 3. Table 2 also shows the geographic breakdown of ASes in Ring 3 and those with IXP connections. Relatively speaking, we see that a

[6] The AS to country and continent mapping is derived from CAIDA NetGeo database [1], which is updated every week.

large portion of ASes in Asia, Europe and other continents are in Ring 3, most of them in Ring 3 stub, as compared to North America. This is most likely due to the fact that most ASes, especially stub networks, in Asia, Europe and other continents typically connect to one or more regional transit ASes for global Internet reachability. Fig. 5(a) shows the number of IXP connections and IXP neighbors for Ring 3 *transit* vs. *stub* ASes that are present at IXPs. Similar to the case of Ring 2 ASes, Ring 3 transit ASes have relatively larger number of IXP neighbors than Ring 3 stub ASes, as transit ASes usually peer with others for economic reasons. Fig. 5(b) illustrates the degree distributions of Ring 3 transit and stub ASes with IXP connections. Furthermore, we also find that Ring 3 ASes (especially transit ASes) with IXP connections tend to have higher degree than those without. This again demonstrates the significant advantages and benefits of participating IXPs for Internet connections.

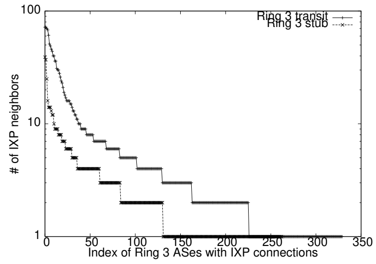


(a) # of IXP neighbors

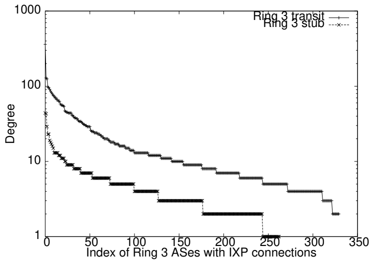


(b) Degree distribution

**Fig. 5.** IXP neighbors and degree of Ring3 ASes w/ IXP connections

**IXPs and Ring 4, 5, 6 ASes** Among the 2,955 ASes in Ring 4, Ring 5 and Ring 6, only 90 have IXP connections. Among them, 83 are in Ring 4, 7 in Ring 5 and none in Ring 6. Most of these ASes are stub ASes located in Asia and Europe, participating only at one IXP to buy transit service.

### 4.2 IXPs and AS Relationship

We now investigate the role of IXPs in influencing the business relations among ASes. To this end, we apply the AS relation inference algorithm to the aggregated BGP routing table. The results suggest that overall most dominant relations among ASes are of *provider-customer* type (93.9%), with only 6.1% inferred to be of *peering* relation type. However, among ASes that are participating at IXPs, the percentage of inferred peering relations are significantly higher (about 42.5%). This confirms the common belief that many (transit) ASes peer with each other at IXPs to exchange traffic was to reduce transit costs.

To gain a deeper understanding of how ASes interconnect at IXPs, we study the relations among ASes with IXP connections in the different rings of the AS topology. In general, we see that the I-core ASes participate at IXPs to provide transit service for other (Ring 2) ASes, *not* to peer with other I-core ASes. This is in contrast to Ring 2 transit ASes.

In Table 3, we examine how Ring 2 ASes form relations with their IXP neighbors. We present the number of provider-customer (PC) and peering (PP) relations and the percentage between Ring 2 and the following four groups: 1) I-core, 2) Ring 2 transit, 3) Ring 2 stub, and 4) Ring 3 ASes. The results suggest that most I-core ASes are providers of Ring 2 ASes, while most Ring 2 transit peer with each other. Most Ring 3 ASes are customers of Ring 2 transit.

**Table 3.** Relations: Ring 2 ASes and their IXP neighbors

| Ring 2 | I-core | | Ring 2 Transit | | Ring 2 Stub | | Ring 3 | |
|---|---|---|---|---|---|---|---|---|
| | PC (%) | PP (%) | PC (%) | PP (%) | PC (%) | PP (%) | PC (%) | PP (%) |
| transit | 35 (72.9) | 13 (27.1) | 236 (41.0) | 339 (59.0) | 39 (67.2) | 19 (32.8) | 262 (62.5) | 157 (37.5) |
| stub | 8 (100) | 0 (0.0) | 39 (67.2) | 19 (32.8) | 0 (0.0) | 6 (100) | 4 (57.1) | 3 (42.9) |

We conduct a similar analysis for Ring 3 ASes and study the relations between Ring 3 ASes and Ring 2 transit/stub and Ring 4 ASes. The results confirm the previous finding that most Ring 3 ASes are customers and they buy transit from Ring 2 transit ASes. On the other hand, some Ring 3 transit ASes also provide Internet connectivity for Ring 4 ASes at some IXPs.

## 5 Conclusions

In this paper we have studied the properties of IXPs and investigated their impact on the AS topology and AS relationship, using traceroute results, BGP archives and other data sources. We developed an algorithm to discover IXPs and infer ASes that participate at these IXPs. Using the discovered IXPs and their inferred AS participants, we analyzed and characterized the properties of IXPs and their participants such as size, geographical location. We found that IXPs are a more prevalent method of Internet connections in Europe. Most participants at IXPs are transit ASes that more aggressively peer with each other, most likely to reduce transit costs. Furthermore, the percentage of peering relationship among transit ASes at IXPs are significantly higher than the average.

We believe that our study is one of the first to perform a detailed analysis of the IXP properties and their impact on AS topology and AS relationship. Previous studies have focused on the properties of Internet topology at either AS-level or router-level [4], or on AS relationship [5, 16]. The work [14] is most similar to ours, which also studies the properties of IXPs and their participants with only Skitter traces, but does not provide as detailed an analysis as we have done here, in particular, in terms of their impact on global AS topology and AS relationship. Our results indicate that it is interesting and important to take methods of Internet interconnections such as IXPs into account when analyzing the Internet topology and how it evolves. Furthermore, IXPs can also be used to provide additional information in inferring AS relationship. We believe that gaining a better understanding of IXPs and their role in Internet interconnections can shed light on the Internet interconnection practices and the evolution of the Internet, in particular, the potential role economics plays in such evolution. Due to limited data, our study constitutes only an initial step towards this goal.

**Acknowledgments:** We thank CAIDA, Route-views, RIPE, PCH for making the data available. We also thank Neil Spring for providing Scriptroute facility. We are grateful to David Farmer and University of Minnesota NTS for helping us set up measurement facilities. We are also grateful to Bill Woodcock for many helpful comments and discussions.

## References

1. CAIDA. http://www.caida.org.
2. H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. Towards capturing representative AS-level Internet topologies. In *Proc. ACM SIGMETRICS*, 2002.
3. G. Cook. The Cook report on Internet protocol, technology, economics, policy: Economics of IP network interconnection peering transit and IXs, Nov. 2002.
4. M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-law relationships of the Internet topology. In *Proc. ACM SIGCOMM*, Aug. 1999.
5. L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Tran. on Networking*, Dec. 2001.
6. Z. Ge, D. Figueiredo, S. Jaiswal, and L. Gao. On the hierarchical structure of the logical Internet graph. In *SPIE ITCom*, 2001.
7. Y. Hyun, A. Broido, and k claffy. Traceroute and BGP AS path incongruities. In *Internetworking International Conference*, June 2003.
8. A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proc. IEEE INFOCOM*, Apr. 2003.
9. Z. M. Mao, J. Rexford, J. Wang, and R. Katz. Towards an accurate AS-level traceroute tool. In *Proc. ACM SIGCOMM*, Aug. 2003.
10. W. B. Norton. The evolution of the U.S. Internet peering ecosystem. Nov. 2003. http://www.equinix.com/pdf/whitepapers/PeeringEcosystem.pdf.
11. U. of Oregon. Routeviews archive project. http://archive.routeviews.org/.
12. PCH. IXPs database. http://www.pch.net/resources/data/exchange-points/.
13. RIPE. Routing information service raw data. http://data.ris.ripe.net/.
14. E. Silenok. Peering relationships through IX vs. direct peering, 2003. http://www.caida.org/~elena/projects.html.
15. N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A public Internet measurement facility. In *USITS*, 2003.
16. L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proc. IEEE INFOCOM*, 2002.
17. S. L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos. A simple conceptual model for the Internet topology. In *Global Internet*, 2001.
18. B. Woodcock. Internet topology and economics: How supply and demand influence the changing shape of the global network. Jan. 2003.