Studying the effect of Internet eXchange Points on Internet link delays

Mohammad Zubair Ahmad and Ratan Guha
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando, Florida
Email: {zubair,guha}@eecs.ucf.edu

***Abstract*—Internet exchange points (IXPs) are an increasingly important constituent of the Internet at the Autonomous System (AS) level. IXPs are set up with the goal of enabling greater efficiency in traffic exchange between ASes both from economical and technical perspectives. Little is however known about the effect that IXPs have in packet transmission between peering ASes, with a significant focus of the research community being in determining their effect on the topology evolution of the Internet. In this paper, we report on the increasing deployment of IXPs around the world over the past few years and carry out a set of experiments to try and establish the effectiveness of Internet routes traversing an IXP. We find that IXP links entail lesser delays than normal links on the Internet even though their presence does not decrease the length of a path. Our results present pointers towards developing more extensive experiments to verify the effectiveness of deploying IXPs worldwide.**

I. Introduction

The Internet comprises of an immense network of smaller networks termed Autonomous Systems (ASes), which are linked together both physically (through hardware such as cables, routers, satellite and so on) and logically (through software). These ASes are characterized by varying sizes and functions and are constantly interacting with each other thereby helping maintain global connectivity. The connectivity is ensured through careful design considerations which enables an AS pair to establish a logical link between themselves. AS links on the other hand are either customer-provider in nature or a cooperative peering link where a mutual traffic sharing agreement is agreed upon, leading to efficient and cost-effective transmission of data between ASes. Internet Exchange Points (IXPs) play an important part in this system by enabling ASes to form peering links and exchange traffic directly.

The complex growth of the AS topology has led to the increased use of IXPs throughout the globe. By enabling the creation of higher number of peering links between ASes, substantial cost savings can be made by direct exchange of local traffic instead of using expensive long-distance links across the world. IXPs also reduce delivery latencies by enabling fewer hops to the destination, free up bandwidth and reduce losses due to congestion. Overall, IXPs are shaping out to be an important component in the evolution of Internet topology and thus need to be studied in more detail.

The key contribution of this paper is to study the effect of IXPs on packet routing throughout the Internet. We aim to find out the effect an IXP has on a packet traversing it on a route through that IXP. A key metric for route efficiency is link delay and in this work we measure and compare delays in IXP and non-IXP routes. We carry out an extensive set of experiments in PlanetLab [1] selecting vantage points situated all across the globe to enable a wider reach across the Internet. We also motivate the importance of IXPs in the growth of the Internet ecosystem by monitoring the increasing percentage of routes involving IXPs over the last five years using available traceroute data from CAIDA's Skitter project [2].

This paper is orqanized as follows: we follow up this section with a brief overview of previous research on IXPs in Section 2 and discuss the IXP architecture and the increasing popularity of the IXPs in Section 3. Section 4 presents a detailed explanation of the experiments we carried out in PlanetLab and our analytical process. This is followed by the results observed in Section 5 followed by our conclusions and the course of future work.

II. Related Work

There has not been any prior work in evaluating the effect of deploying IXPs in the Internet. IXPs have been looked into with regard to studying the dynamics of Internet topology evolution [3] as a high number of peering links at IXPs are mostly hidden from traceroute based studies. He at al. [4], [5] concentrate on searching for a greater number of peering links at these IXPs and is followed by another extensive study by Augustin et al. [6] where a huge portion of the IXP substrate is mapped. However there has not been any significant work which deals with evaluating the technical implications of Internet routes passing through IXPs.

III. The IXP architecture

IXPs are independently maintained physical infrastructures enabling public peering of member ASes. An IXP provides physical connectivity between the different member networks while the decision to initiate BGP sessions between AS pairs is left to the individual networks themselves. Figure 1 represents a regular scenario where a set of ASes (A to E) transmit data to each other using the Internet. Here local ASes end up using international links to transmit data which increases costs alongwith decreasing network performance. Only if ASes

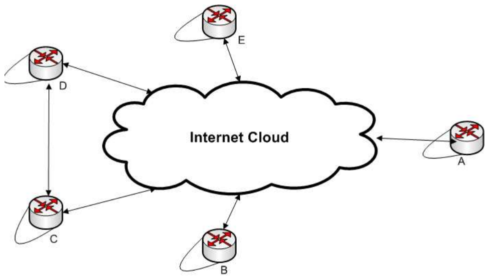

Fig. 1. A set of ASes transmitting data to each other through the Internet. AS C and D share data through a direct peering link.

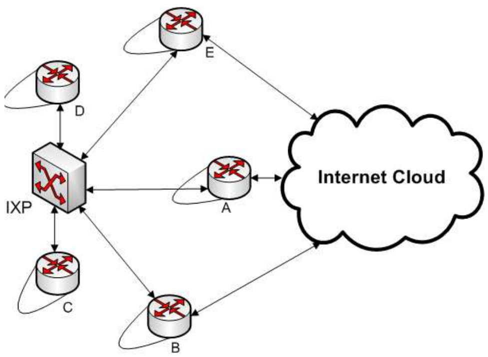

Fig. 2. A set of ASes transmitting data to each other through an IXP.

have a local connections (AS C and D) are these problems mitigated. Figure 2 shows a similar scenario with the ASes peering at the IXP switch. In this case, data sent between these ASes need not traverse the entire Internet and can be directly shared through the IXP. These peering links reduce transmission delays, use lesser international bandwidth and thus reduce overall costs of exchanging data for every IXP member AS.

A. IXPs and member ASes

Identifying IXPs in a traceroute has been described extensively in [5] and [7]. IXPs are assigned an IP address block and each AS peers at the IXP with a definite IP address for the interface within the given block. The list of IXP address blocks are available at PCH [8] and PeeringDB [9]. With the known list of IXP address prefixes we can search for every prefix from traceroute data and identify routes which include an IXP hop. As stated in [5] AS participants may then be identified by following the sequence of IP addresses before and after the known IXP address. By mapping the IP address of the participants to their AS numbers we can obtain the participants at that particular IXP. Not all IXP participants may be identified using this technique which however is not a focus of this paper. Our primary aim is to study the effect of peering at the IXPs for the participant ASes and so are not focused on finding all missing AS-IXP links.

B. Motivation: Growth of IXPs

An increasing number of IXPs are being deployed across the world to enable more efficient traffic delivery over the

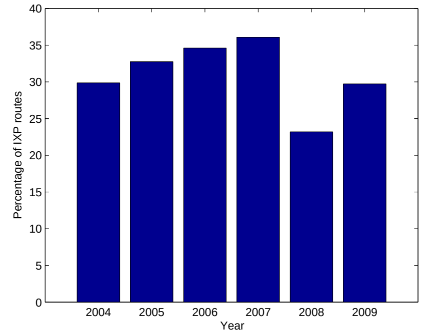

Fig. 3. Percentage of IXP routes visible in one cycle of Skitter traceroute data every year for the month of September.

Internet. This growth in the number of IXPs has been skewed with regard to the geographical location of these new IXPs being set up. There are numerically higher number of IXPs in Europe and North America than those in Asia or Africa for example. However, there is no denying the fact that with an increasing number of IXPs coming up and with more ASes peering at these IXPs, the net traffic going through these IXPs has increased over the years.

To study the impact of IXP routes we first need to quantify the percentage of routes going through any IXP in the Internet. To do this, we obtain one complete cycle of Skitter traceroute data from the year 2004 to 2009 for the month of September. A complete cycle of data represents different skitter vantage points across the world sending out traceroute probes to the standard CAIDA destination list and records the paths taken. Based on the available list of IXP prefixes obtained from PCH and PeeringDB, we search for routes consisting of hops within these prefixes. An IXP route is thus defined as a route which contains atleast one hop through the network with a known IXP prefix. We count the number of IXP routes obtained within one cycle and calculate its percentage based on the total number of routes obtained for the same cycle period. Figure 3 presents the percentage of IXP routes obtained every year and we observe that for most years we have atleast 30 percent of all routes traversing an IXP. This means that almost one in every three routes goes through an IXP. The drop in percentage in 2008 and 2009 can be attributed to the fact that CAIDA's skitter architecture underwent a major change that year transferring to the Ark architecture. This resulted in a fewer traceroute probes being sent out and thus there were lesser routes recorded during this time. Table I presents the total number of routes observed along with the total number of IXP routes obtained. The total number of routes decrease significantly in 2008 and 2009 also due to the fact that CAIDA's destination list has not been updated and this leads to a decrease in the total number of reachable IP addresses. Oliveira et al. [10] point out that a high number of links and routes are not visible in the Skitter data due to its shrinking probing scope. The number of routes visible have decreased which is has led to a decrease in the number of

TABLE I

IXP GROWTH OBTAINED FROM SEARCHING KNOWN IXP PREFIXES FROM ONE CYCLE OF SKITTER DATA FOR THE MONTH OF SEPTEMBER

Year	IXP Routes found	Total routes visible	Percentage
2004	6963592	23312823	29.87
2005	6999045	21370051	32.75
2006	6387175	18455760	34.60
2007	5606309	15541716	36.07
2008	1629327	7020300	23.20
2009	1906532	7407891	25.73

IXP routes too, but it still shows a significant percentage of routes being taken going through an IXP thereby underlying the importance of an IXP in the evolution of the Internet ecosystem.

IV. EXPERIMENTS

This section presents our data collection and analysis methodology in detail. Our data collection is broadly divided into two phases: (i) Phase-1: perform active tracing from various vantage points to a destination list and (ii) Phase-2: ping IXP participants and record delay data. From the ping data collected we calculate individual link delays between IXP participants which gives us a measure of actual IXP performance.

A. Phase-1: Traceroute from various vantage points

We carry out traceroute measurements from 35 different PlanetLab [1] vantage points over 6 days to come up with a dataset containing more than 20 million traces from which we search for the set of routes incorporating an IXP prefix. The main reason for carrying out our own specific traceroute based study instead of using the standard data sources such as CAIDA [2], DIMES [11] or iPlane [12] are because we need to follow it up with specific delay based studies from these vantage points ourselves. Using PlanetLab to carry out such a task is very suitable for our experiment. The vantage points chosen represent a global coverage from all the continents with an increased number of vantages selected from North America and Europe. This is because a significant amount of Internet traffic is generated and delivered within these geographical locations and hence the job of the IXP becomes very critical. The traceroutes are all conducted to same destination list and the different routes taken are recorded along with the round trip times.

B. Phase-2: Ping IXP participants and record delay data

Determining the IXP participants has been active area of research in the area of Internet topology evolution [13]. Xu et al. [7] proposed a simple technique by which one hop neighbors of IXP ASes could be identified as participants at the IXPs. He et al. [5] significantly extend the procedure proposed in [7] and come with an increased number of IXP links. For our studies, the former technique of determining one hop IXP participants is good enough as we are presented with a large number of ASes participating at each IXP. Our goal is not to find all the links but choose a significant number of these links to carry out our delay studies. Hence, we use the simpler method of finding a good number of links to start our ping experiments on.

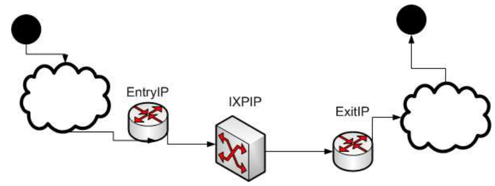

Fig. 4. IXP triplets with one hop participants.

On a given IXP route, we record the IP address of the hop before and the IP address of the hop after the IXP prefix and name these IPs *EntryIP* and *ExitIP* respectively. These IP adresses are a part of the AS participating at this IXP as shown in figure 4. After listing out a set of (EntryIP,IXPIP,ExitIP) triplets, we now start the process of pinging each of these IP addresses from the vantage point from where the initial traceroute was conducted. However, it has to be noted that not all the different IP addresses selected actually respond to ping probes. Some of the participating ASes may block ping traffic or the machine with the particular IP address may be down at the time of probing. We write a set of scripts to first parse out the reachable IPs from these triplets. A triplet is accepted as a candidate for the ping experiment only if all the IPs in the triplet are reachable. We select 300 such working triplets and ping them for a continuous period of 24 hours. The ping results are stored at each vantage point for further analysis as explained next.

C. Calculate link delays

For every (EntryIP,IXPIP,ExitIP) triplet which we individually ping we take every corresponding delay value and carry out the following computation:

$$D_i = D_{IXPIP} - D_{EntryIP}$$

$$D_j = D_{ExitIP} - D_{IXPIP}$$

$$\forall i, j \leq n$$

where D_{IXPIP}, $D_{EntryIP}$ and D_{ExitIP} are sequentially corresponding ping delay values and n is the number of ping packets sent during the 24 hour period.

The values for D_i and D_j hence represent the delays on the links between the participating ASes and the IXPs. The purpose of calculating these link delays is to determine how the presence of an IXP on a route affects the delivery time of a packet. Does the presence of an IXP actually help in faster packet delivery betweeen source and destination or is there no significant performance improvement? To carry out a comparitive study, we need to carry out similar ping experiments with links in regular routes across the Internet.

TABLE II
PARTIAL LIST OF ISPS FROM WHICH LINKS WERE SELECTED GROUPED AS TIER-1 (TOP), TIER-2(MIDDLE) AND TIER-3(BOTTOM).

ISP Name	AS Number	Location	Degree
ATT	7018	US	1490
Verizon	701	US	2569
Sprint	1239	US	1735
Colt	8220	Europe	161
Eqip	3300	Europe	67
Telia	1299	Europe	256
Hong Kong Telecom	4637	Asia-Pacific	201
AOL	1668	US	156
Cogent	16631	US	187
Broadwing	6395	US	231
UUNet-Europe	702	Europe	587
Easynet	4589	Europe	86
Soul	9942	Australia	114
VSNL	4755	Asia-Pacific	49
Espire	6467	US	30
Exodus	3967	US	43
MFN	6461	US	498
Telstra	1221	Australia	66
ATT-AP	2687	Asia-Pacific	24
Singapore Telecom	7543	Asia-Pacific	8

D. *Link delays in regular routes*

Using the same traceroute data obtained we select a set of paths which do not contain any hop through an IXP AS. Our main aim is to choose a set of links similar to the ones chosen for IXPs and carry out another set of ping experiments to determine an average value for link delays. Choosing any set of three simultaneous hops does not however prove to be a representative set of links on a given traceroute. We devise a chosen set of criteria to choose all types of links on a regular route to obtain link delay values which would be close to most standard link delays obtained over the Internet. We start off with choosing an initial hop randomly within the first ten hops of a route and then select the remaining two hops based on one of the following criteria:

- The second hop is selected from within the same AS as the first hop. In such a case, the third hop is selected from the first different AS following this IP address.
- The second hop is selected from the next nearest Tier 1 ISP. In this case, the third hop is selected from the next Tier 2 ISP.
- If the first hop is a Tier 1 ISP the second hop is selected from the nearest Tier 2 ISP and the third from the next unique ISP (whether Tier 2 or Tier 3).
- The second and the third hops both belong to the same AS as the first.

Table II lists some of the major ISPs from where we selected IPs to identify representative links.

We select an equal number of IP triplets based on each of the criteria above and we believe it forms a basis for a set of representative links in the Internet. We also do not choose IP addresses more than five hops away from the previous one as this would denote a greater distance and thereby a much larger delay than what we would essentially like to compare the IXP link delays with. We name the IP address triplets as *(RegEntryIP,RegMiddleIP,RegExitIP)* and carry out

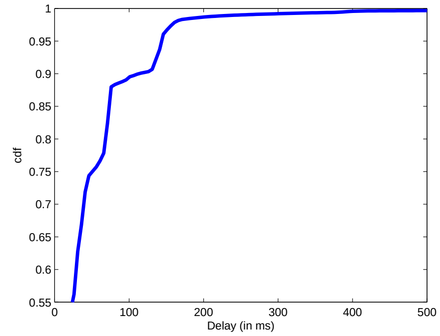

Fig. 5. CDF of delays obtained from IXPs.

individual pings to each of these IPs simultaneously over a 24 hour period. Similar to the IXP links we again compute the difference in delays between these IPs:

$$RD_i = D_{RegMiddleIP} - D_{RegEntryIP}$$

$$RD_j = D_{RegExitIP} - D_{RegMiddleIP}$$

$$\forall i, j \leq n$$

where $D_{RegMiddleIP}$, $D_{RegEntryIP}$ and $D_{RegExitIP}$ for each IP triplet. The RD_i and RD_j values are recorded.

E. *Route lengths*

We also aim to study the effect of IXPs on the average path length (based on the number of hops) taken by a packet from a source to the destination. Our aim is to observe if the packet takes a greater or lesser number of hops on its path and if the presence of IXPs significantly alter the path characteristics. From the traceroute data obtained we measure the path lengths for the paths including one or more IXP prefix and the other non-IXP paths. Even though the number of non-IXP paths are greater we aim to measure the average and CDF of the path lengths of both types of routes from every vantage point.

V. RESULTS

We present the results of the experiments carried out in this section.

A. *Link delays*

Figures 5 and 6 present the CDF values for the link delays computed as explained in the previous section. Monitoring the individual link delays between IXP participants we observe that almost 95 percent of the links monitored have a delay less than 175ms. This means that only 5 percent of links from all vantage points measured have delays greater than 175ms. On the other hand for 95 percent of the links monitored on non-IXP routes, the delays are significantly greater at more than 250ms as shown in figure 6. This means that IXPs actually help in reducing delays in packet delivery for ASes participating at them. This would enable better and quicker packet delivery and thus a higher network throughput on these links along with a higher degree of user satisfaction. This would have a greater

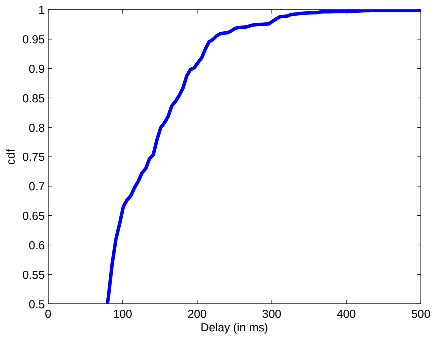

Fig. 6. CDF of delays obtained from regular ISPs.

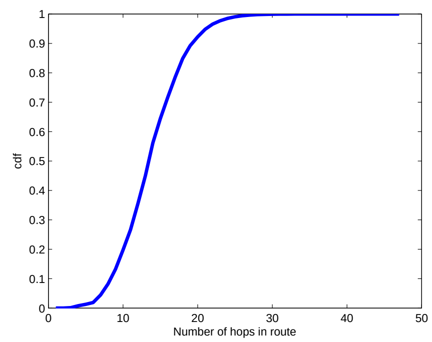

Fig. 7. CDF of average number of hops per route obtained from IXP routes.

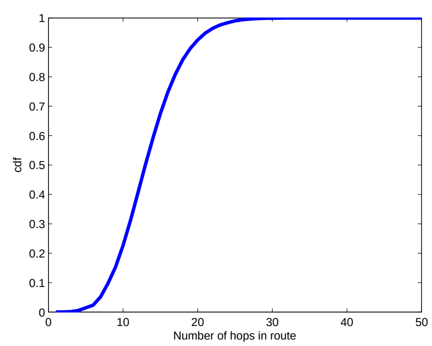

Fig. 8. CDF of average number of hops per route obtained from regular ISP routes.

effect on streaming or voice/video data being transmitted as these applications are sensitive to even small delay and jitter values. By reducing the delivery latencies on a route, the IXPs thus enable more efficient data transfer across the Internet.

B. Path lengths

Figures 7 and 8 present the CDF of the average number of hops per route for an IXP route and a non-IXP route respectively. Both figures show an almost identical distribution with regard to the hop count with almost 90 percent of IXP routes have a slightly lower hop count than for a similar percentage of non-IXP routes. This shows that even with the presence of IXPs the number of hops a packet has to take to reach the destination still remains largely the same with no significant difference.

C. Discussion

Our main aim in this paper has been to quantify the impact of IXPs on Internet routing performance which leads us to carry out the delay experiment for a representative set of IXP and non-IXP links on the Internet. Delay is one of the most important metrics in determining the effectiveness of routes on the Intenet. From the results obtained we observe that IXPs do in fact lead to lower packet delivery latencies. Keeping the economical advantages of deploying IXPs aside, we can see that these traffic switches enable a faster and more efficient means of data delivery from any source to destination. Even though the average number of hops in an IXP route is not significantly different than those routes which do not traverse an IXP, the delivery delays are lower. More studies need to be carried out to determine the effect of packet switching and queueing at the IXP in case of it being heavily loaded. If hundreds of ASes start peering at an IXP, there will be a point where the hardware would not be able to handle all the excess traffic efficiently. This could lead to greater delays in packet delivery than if a non-IXP route is selected as the packets may end up being queued for a greater period of time. Most IXP websites publish daily traffic volume data but more information with regard to IXP hardware availability and peering matrices need to be made available before we can have a better understanding of the effects of peering traffic at the IXPs.

VI. Conclusions and future work

We report the results of an initial study on the effect of IXPs on routing delays in the Internet and observe that IXPs decrease packet delivery latencies for routes in comparison to routes not traversing an IXP. Our experiments also point out that the number of hops in an IXP route is generally the same as in a regular route. This characteristic builds on the effectiveness of an IXP as it means that the best policy for ASes exchanging significant amounts of traffic, peering at IXPs would be a good idea. While the effect of IXPs in Internet topology evolution has been studied extensively, there has not been any study towards evaluating the performance of these IXPs in the Internet ecosystem. Our work presents for the first time such a delay performance evaluation with scope for a much more detailed and exhaustive study. Deploying our experimental code in a greater number of vantage points will give us a better reachability and a better picture of link availability and delays. Also, more experiments need to be carried out to determine route effectiveness, flow throughput and other important metrics used to measure routes on the Internet. Carrying out probe-train experiments (as a case-study) between a fixed source and destination over a period of time will also help determine how the IXP links fluctuate and effects of BGP convergence on link failure. We are currently carrying out these experiments and look forward to gaining a better understanding of the working of the IXP links in the near future.

REFERENCES

[1] "Planet lab website," http://www.planet-lab.org.

[2] D. A. Bradley Huffaker, Young Hyun and K. Claffy, "Skitter as links dataset - jan 2000 to may 2009," http://www.caida.org/data/active/skitter_aslinks_dataset.xml.

[3] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*. New York, NY, USA: ACM, 1999, pp. 251–262.

[4] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A systematic framework for unearthing the missing links: Measurements and impact," in *NSDI*, 2007.

[5] 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 Trans. Netw.*, vol. 17, no. 2, pp. 391–404, 2009.

[6] B. Augustin, B. Krishnamurthy, and W. Willinger, "Ixps: mapped?" in *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. New York, NY, USA: ACM, 2009, pp. 336–349.

[7] K. Xu, Z. Duan, Z.-L. Zhang, and J. Chandrashekar, "On properties of internet exchange points and their impact on as topology and relationship," *Networking*, vol. 3042, 2002.

[8] "Packet clearing house," http://www.pch.net.

[9] "Public exchange points search/list," http://www.peeringdb.com/private/exchange list.php.

[10] R. V. Oliveira, B. Zhang, and L. Zhang, "Observing the evolution of internet as topology," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 313–324, 2007.

[11] Y. Shavitt and E. Shir, "Dimes: let the internet measure itself," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 71–74, 2005.

[12] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iplane: an information plane for distributed services," in *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*. Berkeley, CA, USA: USENIX Association, 2006, pp. 367–380.

[13] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the internet as-level topology," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 53–61, 2005.