

Studying the Effects of Internet Exchange Points on Internet Topology

Mohammad Zubair Ahmad* and Ratan Guha

Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, Florida, USA

Abstract

The recent uncovering of a high number of peering links at Internet Exchange Point (IXP) locations across the world has made these exchange switches a critical component of the Internet Autonomous System (AS) level ecosystem. Studies concentrating on the internet topology evolution have surmised that numerous links hidden at these exchange points hold the key towards solving the missing links problem in studying the evolution of the AS-level topology of the Internet. In this work, we study the effect of this set of hitherto unseen peering links on the visible Internet topology. Starting with a set of measurements determining the growth of IXPs in the inter-domain routing architecture of the Internet and continuing with a more advanced graph based metric analysis of available Internet topology data, we conclude that IXP links follow power law increase characteristics while exhibiting definitive clustering characteristics. Moreover, these additional links affect the joint degree distributions of nodes with higher degrees while leaving most other types of nodes unchanged. We conclude that the currently inferred AS-level maps of the Internet demonstrate considerable variations with the incorporation of these new links and could eventually lead to a remodeling of our understanding of Internet topology evolution.

Keywords: Internet topology; Internet exchange points; Autonomous systems; Topology evolution; Route views

Introduction

The explosive growth of the internet as a collection of Autonomous Systems (AS) has led to a plethora of efforts in trying to understand the current internet as topology and its evolution. Analysis of the Internet topology is needed for better network planning and designing optimal routing strategies [1]. The Border Gateway Protocol (BGP), which is used to route all Internet traffic causes packet loss and transient disconnectivity during convergence [2]. Thus, creating a more robust routing architecture requires a greater understanding of the underlying Internet AS topology evolution.

The setting up of Internet Exchange Points (IXPs) has been beneficial primarily from an economic perspective for ASes to peer directly with other member ASes at these locations [3]. Increased peering at these IXP switches has led to more recent research [4-6] showing a significant number of new links being uncovered at these locations impacting our understanding of the Internet topology at the AS level. Augustin et al. [6], present a framework to uncover these hidden links and report the presence of almost 18K more links than previously known, the majority of which are of the peer-to-peer type.

It has been suggested [4] that the extra peering links at these IXPs may hold the key to solving the missing links problem for the AS-level Internet and [6] shows that this hypothesis is probably true. However, the task ahead of us does not stop at uncovering these peering links. These additional links obtained need to be studied and analyzed in detail with respect to the existing Internet topology and their effects measured before a final conclusion can be arrived at. Any number of questions arises: Do the extra IXP links uncovered have a significant effect on the growing topology dynamics of the Internet? If the effects of these links are significant then how do we change our outlook in conducting topology research to accommodate these newer changes? Does solving the hidden links problem with these newer IXP links actually mean that we can accurately predict the growth of the Internet and verify previous evolution models as correct or not?

In this paper we study AS visibility at IXPs with the primary aim of establishing the role of these IXPs in determining the evolving Internet topology. We try to find out if IXP data presents significant connectivity information not present in the more conventional data

sources such as RouteViews BGP data [7] or Skitter data from CAIDA [8] among others.

The primary contribution of this paper is to carry out graph based studies aimed at finding an answer to one primary question: do the recently uncovered peering links significantly alter the state of the Internet topology as we know it? The constant evolution of the internet topology is undergoing a sea change with the advent of increased peering (leading to a widely inferred 'flattening' of the internet [3,9] and we carry out a graph based study into the macroscopic properties of these IXP peering links with respect to the rest of the visible Internet. We choose a set of metrics discussed in [10] to study and analyze the topological properties of the Internet from various data sources in addition to the extra peering links obtained at the Preprint submitted to Computer Communications December 2, 2012 IXPs. Using numerous available data sources enables us to create a representative 'graph' of the AS-level Internet which we then analyze. We observe that while the extra links affect the topology for specific metrics, the core power law growth behavior is not drastically altered. Our studies point out for a need to keep a definite track of links being created and destroyed at IXP locations specially with a significant percentage of Internet routes passing through IXP routers.

Our paper presents results pointing to the effect IXP links are having on the visible Internet topology and serves as a precursor to more work needed to come up with a concrete view about the net effect these links have on the Internet topology.

We organize the rest of the paper as follows: Section 3 talks about the architecture of the IXPs and is followed by a brief description of some related work in section 3.3. The growth of IXPs is quantified in

*Corresponding author: Mohammad Zubair Ahmad, Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, Florida, USA, E-mail: zubair@eecs.ucf.edu

Received October 20, 2012; Accepted December 11, 2012; Published December 13, 2012

Citation: Ahmad MZ, Guha R (2012) Studying the Effects of Internet Exchange Points on Internet Topology. J Inform Tech Softw Eng 2:114. doi:10.4172/2165-7866.1000114

Copyright: © 2012 Ahmad MZ, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

section 4 while section 5 explains the graph analysis procedures we used. Section 6 presents the results observed with relevant discussions and is followed by section 7 with the analysis and related discussions. Finally, we conclude in section 8 discussing the overall summary of our results, limitations and course for future work.

IXPs and Topology Evolution

This section describes the role of IXPs and their effects in the growth of the Internet ecosystem. We present a brief introduction to the IXP architecture which leads us to the actual reasoning behind why they are an important component in the study of Internet topology evolution.

IXP architecture and growth

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 while decreasing network performance. Only if ASes have a local connection (AS C and D) are these problems mitigated. IXPs enable public peering between member ASes by providing physical connectivity infrastructure and the decision to initiate BGP sessions between AS pairs is left to the individual AS networks themselves. Most IXPs connect members through a common layer-2 switching fabric [5]. The public peering at the IXP then becomes simpler due to the availability of physical infrastructure, with member ASes A and B (Figure 2) initiating a BGP session to exchange packets through the IXP switch. On the other hand if E needs to send data to F, it requires the set up of BGP sessions between routers in the Internet cloud for it

to be able to successfully transfer data to F. Figure 2 shows a 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.

The question arises as to when should an AS subscribe to an IXP? It is dependent on a variety of factors, primarily economic in nature. In the scenario shown in figure 2, if there is a significant volume of daily traffic between AS E and F, then it would probably be better off for F to peer at the IXP. Assuming both are stub ASes, the amount both would have to pay their respective transit providers would be far greater than the cost of setting up a peering link at the IXP. Data transfer costs, which in turn is dependent on traffic volumes are generally the determining factors behind AS peering at IXPs.

The advantage of peering at IXPs has led to a significant growth in the number of ASes peering at these switching points worldwide. As more and more ASes start peering there are a greater percentage of data packets being routed in the Internet through these switches. In the following section we conduct some measurements and show that almost thirty percent of all routes in the Internet traverse an IXP. This leads to a greater number of peering links being formed at the IXPs thereby affecting the various characteristics of the Internet topology.

Data sources and identifying IXP peering links from traceroutes

Internet topology evolution is typically studied by using various established datasets made available to the research community. BGP routing table dumps from the University of Oregon's RouteViews project [7] is the most extensively used resource. AS links appearing in the BGP tables represent existing links with a high probability of being alive and is thus a more reliable source of information. However, if a link breaks or a node is down, the information takes some time to be updated through the network through BGP updates thereby leading to higher routing table convergence times. These updates have also been used as topology snapshots since they show a greater number of AS links over time [11].

Another widely available source of data is the data released by CAIDA under the Archipelago (Ark) infrastructure for research use [12]. From various vantage points across the internet, ICMP probe packets are sent to a set of destination IP addresses using the traceroute tool. iPlanes [13] and Dimes [14] are other important and widely used sources of data publicly available for use in the study of Internet topology evolution.

There is a limited availability of data with respect to IXPs. PCH [15] maintains and makes available a set of BGP tables collected from a set of IXP routers worldwide while Peering DB [16] is another project where IXP information is manually updated by individual providers. The recent IXP mapping effort by Augustin et al. [6] present IXP specific datasets including

IXP IDs and network prefixes. Using a variety of tools developed, the authors come up with a list of IXP members and a set of peering links at these IXPs. They successfully discover and validate the existence of 44K IXP peering links which is roughly 75% more than reported in previous studies [4,5]. This additional dataset of peering links at IXPs is used in this paper to create a more complete Internet topology graph.

IXP peering links have been mentioned as the hidden links which may be the key to solving [4,5] the well known missing link problem in

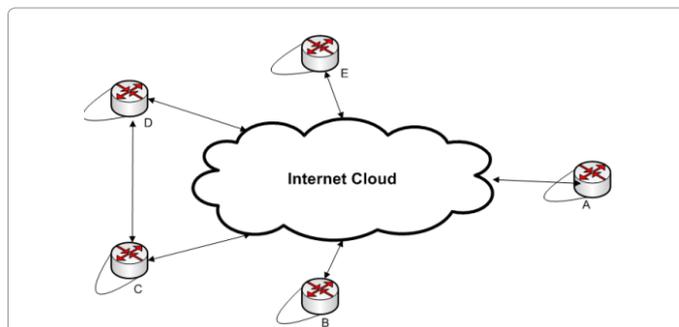


Figure 1: A set of ASes transmitting data to each other through the Internet. AS C and D share data through a direct peering link.

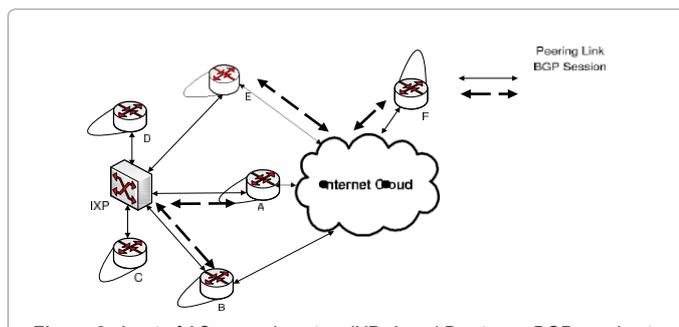


Figure 2: A set of ASes peering at an IXP. A and B set up a BGP session to exchange data while E and F use the Internet cloud to transmit data to each other. Any AS peering at the IXP may initiate BGP session with a peering AS.

Dataset source	Name
RouteViews BGP [7]	<i>RVIEWS</i>
CAIDA (Ark/Skitter) [12]	<i>CAIDA</i>
Packet Clearing House [15]	<i>PCH</i>
DIMES [14]	<i>DIMES</i>
IXP Mapping [6]	<i>IXPMAP</i>
<i>RVIEWS + CAIDA + DIMES + IXPMAP</i>	<i>IXPALL</i>

Table 1: 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

Table 2: Datasets analysed and nomenclature.

the study of Internet topology evolution. Table 2 presents a summary of the various data sets used and the nomenclature used throughout this paper.

Identifying IXPs in a traceroute has been described extensively in [4] and [17]. 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 lists of IXP address blocks are available at PCH [15] and Peering DB [16]. 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 [4] 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. We use these techniques to identify paths traversing an IXP in a later section.

Related work

Internet topology evolution research is traditionally carried out with active measurements with [18] being one of the earliest works constructing topology snapshots from BGP routing tables and updates. This led to the general technique of constructing AS or router-level graphs of the Internet topology using both traceroute and BGP data. The authors in [10,19,20] analyzed these graphs based on various graph theoretical metrics. The focus has mostly been on designing measurements to maximize the number of links uncovered and solve the incompleteness problem [5,21]. Researchers have all along concentrated on finding new links [22] and removing the expired links [23] formed due to the constantly changing Internet dynamics.

Topology evolution needs to be studied in detail to help in the design and implementation of better topology generators and evolution models. These topology generators play a major role as newer and more efficient routing architectures can only be designed when effective topology maps can be created. Models proposed in [24,25] aim to generate graphs which exhibit desired graph characteristics of the Internet.

IXPs were recently identified as an integral component of the Internet architecture and were made a focal point of the study in [17] and [6]. He et al. [4,22] carry out significant studies on uncovering IXP peering links and suggest that these locations hold the key of solving the hidden links problem in Internet topology research. By using the very comprehensive study carried out by Augustin et al. [6], we aim to measure the impact these IXP peering links are having on

the evolving Internet topology today. Gregori et al. [26], presented an initial work discussing the impact IXP links are having on the AS-level Internet topology while we provide a more in-depth analysis and characterization of various graph based topology metrics in our work. Our aim is to interpret and analyze the effects these IXP peering links are having on the Internet topology.

Growth of IXPs

An increasing number of IXPs are being deployed across the world to enable more efficient traffic delivery over the 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 numbers 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 Internet 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 (now renamed Ark) 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 Peering DB, we search for routes consisting of hops within these prefixes. An IXP route is thus defined as a route which contains at least 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 at least 30 percent of observed 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 1 presents the total number of routes observed along

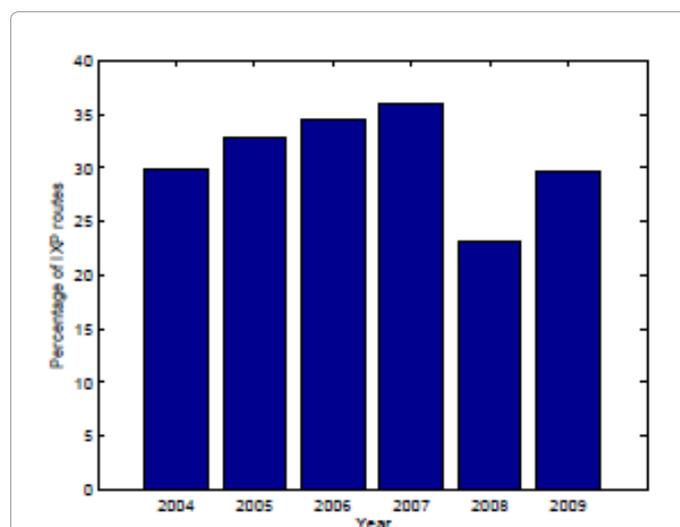


Figure 3: Percentage of IXP routes visible in one cycle of Skitter traceroute data every year for the month of September.

with the total number of IXP routes obtained. Oliveira et al. [23] 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 has led to a decrease in the number of IXP routes too, but it still shows a significant percentage of routes being taken going through an IXP thereby underlying the importance of IXPs in the evolution of the Internet ecosystem.

AS Graph Analysis

In this section, we present our methodology to obtain AS information from the different datasets we choose to consider.

Our main aim is to identify the set of ASes visible, the number of AS links visible and other important network metrics representing important properties of the resultant graph. We look at topology metrics considered by Mahadevan et al. [10], as they appear to fundamentally characterize Internet AS topologies and have been widely used. As this study is primarily meant for comparison purposes, we decided to obtain a snapshot of Internet topology data from the data sources for a period of 31 days in October 2009. A month's worth of data provides a reasonable snapshot of the evolving Internet topology with enough time for different ASes and links to either show up or go down. We obtain AS-level graphs from each data source as mentioned next and merge the 31 daily graphs into one graph per dataset.

Graph construction

RouteViews [7] collects and archives static snapshots of BGP routing tables from a set of monitors which can be accessed from the RouteViews data archives. Deriving the graphs from October 2009 we obtain a set of AS paths which we then convert to a set of AS links. The unique AS links obtained are set aside from which every individual AS visible is then recorded. The final combined monthly graph we refer to as the RVIEWWS graph in the rest of the paper.

CAIDA's IPv4 Routed/24 topology dataset [12] uses team probing to distribute the work of probing the destinations among the available monitors using the scamper tool and forms a part of the Archipelago (Ark) topology infrastructure (which was formerly known as Skitter). Scamper probes are currently sent to a random destination prefix from a set of 7.4 million prefixes. As specified in [10] private ASes generate indirect links which we filter out during creation of the AS-level graphs and are then combined to form the final CAIDA graph.

PCH [15] releases the BGP routing tables at various IXP routers (currently 63) from various locations around the world. These routing table formats are the same as the RouteViews tables and hence are analyzed using a similar technique. We construct the PCH graph from these daily graphs.

The DIMES Internet mapping project is a distributed technique carrying out traceroute measurements from individual users located worldwide. Millions of traceroute/ping measurements are carried out by the low footprint DIMES agents installed on volunteer local hosts to present a detailed view of the Internet with a significant percentage of new links compared to those found in RVIEWWS and CAIDA.

The IXP Mapping project [6] releases data specific to IXPs across the Internet with only peering links unearthed at these IXPs. We term this dataset IXPMAP. This is the most comprehensive set of peering links present at IXPs currently available to the research community and we make it the primary source of study in this paper.

The peering links in IXPMAP are however not useful by themselves as they do not in any way give a complete picture of the

Internet. As in other similar topology related studies, we combine these peering links with the other views of the Internet we obtain from the different datasets available to us. As stated earlier, we have the CAIDA traceroute based dataset (representing the data plane) and the RVIEWWS BGP based dataset (representing the control plane). We compare the links obtained from the PCH data with the other BGP based dataset (RVIEWWS) and present the result in table 3. It is observed from the table that PCH contains only 370 unique links in comparison to RVIEWWS and the other IXP-specific dataset with a high number of links (almost 71k) being common among the BGP based datasets. The reasoning behind such similarity between these datasets is the fact that both are derived from BGP tables at a set of routers some of which are actually common to both sources. Due to such a characteristic of the PCH data we simply combine the unique links obtained from this dataset to the RVIEWWS graph to simplify our analysis and reduce the number of graphs generated to three.

We complete the entire picture of the Internet by combining CAIDA, RVIEWWS, DIMES and IXPMAP to one entire IXPALL graph. This graph is characterized by the data plane (CAIDA), the control plane (RVIEWWS), extensive peer to peer links (DIMES) and the peering links (IXPMAP) and built over a one month period, is relatively representative of the Internet during that period of time.

Validity of chosen datasets

As detailed in the subsection above, we carry out a careful consideration of each of the available datasets before combining them to create the final combined graph of the Internet. While each of the links made available are validated by the sources before release, it can be considered that over time some of the links may simply expire and new ones created. This is especially true for the IXPMAP dataset which is not maintained by the original developers any more. However we do not consider the dataset to have become corrupt and rendered useless. By using historical data (from CAIDA and RVIEWWS) for that particular month we obtain a relatively clear and correct snapshot of the Internet for that particular period and study the graphs. The question of the current validity of the peering links could be raised when the topology evolution is being studied over an extended period of time, something which is not the goal in this work. The IXP peering links would have a high probability of remaining valid for the period considered and thus enable an accurate study of their effects on the AS-level topology of the Internet.

We carry out graph based comparison studies in the next section between CAIDA, RVIEWWS and the IXPALL datasets and do not report the results of the DIMES dataset individually. This is because both CAIDA and RVIEWWS present distinctly different views of the Internet as mentioned earlier (the data and control planes respectively) while DIMES presents an overall view based on the locations of the user agents. However, the unique links from DIMES are used in creating our view of the complete Internet in IXPALL.

	Links
PCH only (GP)	370
RVIEWWS only (GB)	1408
IXPMAP only (GM)	47507
PCH + RVIEWWS (GP ∩ GB)	71284
PCH + IXPMAP (GP ∩ GM)	57
RVIEWWS + IXPMAP (GB ∩ GM)	159
PCH + RVIEWWS + IXPMAP (GP ∩ GB ∩ GM)	4250

Table 3: Comparing the number of observed links in the PCH, RVIEWWS and IXPMAP graphs.

Topology characteristics

Degree distribution

The node degree distribution is the probability distribution of the node degrees in a graph. In other words, it is the probability that a node selected randomly is of k -degree and this probability is calculated by:

$$P(K) = \frac{n(k)}{n}$$

where $n(k)$ is the number of k -degree nodes in a graph with total number of nodes n . Scale-free networks such as the Internet have been shown to exhibit power law degree distributions [27] and hence the power law exponent is computed for this metric. This power law model has had a significant effect on Internet topology research and topology generators [1,25] are designed primarily adhering to this characteristic.

From figures 4 and 5 we observe distinct power law characteristics being followed by all three topology datasets for a wide range of node degrees. The average node degrees (Table 4) are in k -order with $RVIEWS \leq CAIDA \leq IXPALL$ and the average node degree in I X PALL exhibiting a significantly higher value than the others. This is largely due to popular IXP nodes exhibiting high degrees due to multiple peering ASes at one location. The power law exponents computed are not affected significantly by these additional high degree nodes with the γ value for the combined IXPALL graph being slightly higher than the others (refer to 5 for complete details). The authors in [10] point out that a natural cut off at power-law maximum degree is obtained at: $k_{PL} = n(\gamma - 1)$.

From table 4 we observe that the maximum node degree k_{max} for the IXPALL is closest to the power law thereby meaning that the power law approximation for this set is relatively accurate.

This result shows that the degree distribution of the IXPALL graph still does follow a power law but with different parameters. By uncovering of these new peering links at IXPs the basic topology evolution characteristic of the Internet does not deviate from the existing power law characteristic and its behavior remains the same. The CCDFs of these graphs also reiterate this conclusion. The addition of an extremely high number of unique peering links does not break the power law characteristics of the graph. Figure 5 shows that the IXPALL graph has a greater number of nodes for corresponding node degrees in comparison with the CAIDA and RVIEWS graphs. This is simply due to the fact that a high number of low to medium degree ASes (degrees of 10 to 1000) peer at the IXP switches with each other. The newer links uncovered are between these peering ASes increasing the total number of ASes with these degree characteristics. However it is evident from the figure that the net characteristic of the Internet's degree distribution still remains the same even with the addition of the IXP peering links

Power law degree distributions

The now famous paper by Faloutsos et al. [27] exhibiting a power-law degree distribution of the Internet graph at the router level led to a plethora of research in this evolution characteristic of the Internet. Suggested scale-free network models based on preferential attachment [28] describe the power law degree distributions with an exponent α between 2 and 3. However there has been a large amount of follow up work where the degree distribution characteristic has been shown to be a result of an inherent bias of traceroute based measurement mechanisms. Lakhina et al. in [29] show that traceroutes from a small set of sources to a larger set of destinations measure edges in a highly biased manner with the degree distribution results differing sharply from that of the actual underlying graph. Achlioptas et al. in [30]

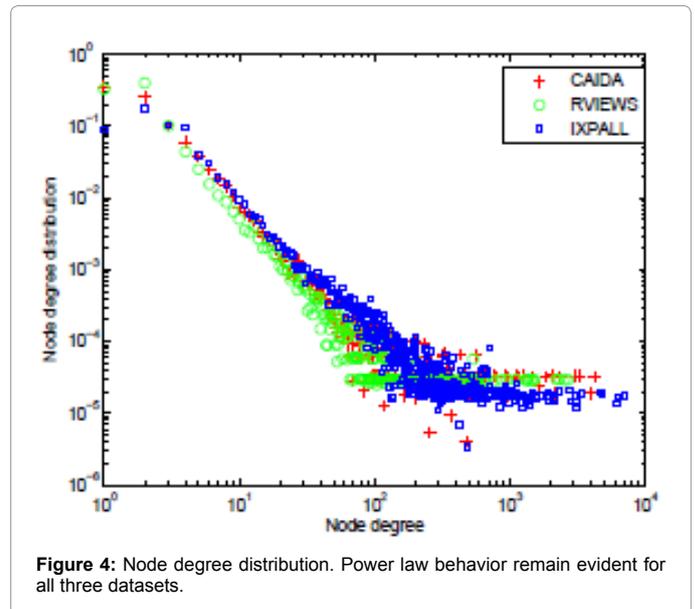


Figure 4: Node degree distribution. Power law behavior remain evident for all three datasets.

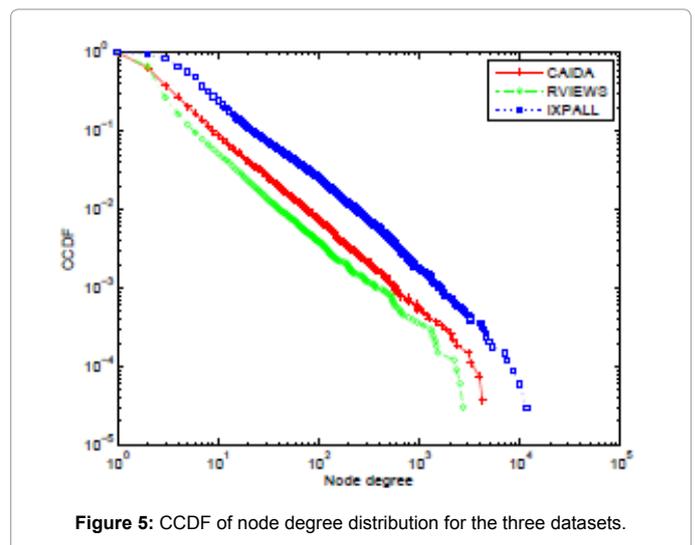


Figure 5: CCDF of node degree distribution for the three datasets.

provide a mathematical proof of the results obtained in [29] while a recent work by Willinger et al. [31] discuss the origin and reasons behind the scale-free Internet myth. We discuss this particular issue in this paper as in our first result we do show that the combined Internet graph exhibits the power-law distribution with an exponent of 2.18 (table 4). It has to be noted however that the basis for not supporting this power law characteristic is for traceroute based studies from a very small set of source monitors to thousands of destination IP addresses across the globe. The authors of [6] carefully select a large number of traceroute enabled looking glass (LG) servers (about 2300) from which they send out targeted traceroute probes to responding target hosts within (or a neighbor of) an AS peering at a known IXP prefix. We believe this technique will not be subject to the traceroute sampling biases as discussed in [29,30] and the IXP peering links obtained also do not show such a property when analyzed in isolation. When combined with the other datasets to represent the entire Internet, these links end up affecting the graph properties but nearly not enough when node degree distributions are studied.

The IXPMAP dataset is inherently free of the traceroute bias in

Metric	Property description	CAIDA	RVIEWS	IXPALL
Average Degree	Number of nodes (n)	26957	33199	33606
	Number of edges (m)	94161	77101	320728
	Average node degree (k^-)	6.98	4.64	19.08
Degree Distr	Max node degree ($kmax$)	4249	2717	11623
	Power law max degree (k^{PL}) max	7301	9690	7809
	Exponent of $P(k)(-\gamma)$	2.14	2.13	2.16
	Maximum degree ratio	0.16	0.08	0.35
Joint degree distr	Avg neighbor degree ($k^-nn/(n-1)$)	0.028	0.015	0.019
	Assortative coefficient (r)	-0.16	-0.20	-0.07
Clustering	Mean clustering (C^-)	0.39	0.25	0.29
	Clustering coefficient (C)	0.02	0.01	0.04
Coreness	Average node coreness (κ^-)	2.05	1.33	4.34
	Max node coreness (κmax)	38	25	87
	Core size ratio (n_{core}/n)	$3 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
	Minimum deg in core ($kmin$) $core$	75	38	119
	Fringe size ratio (n_{fringe}/n)	0.37	0.33	0.25
	Max degree in fringe	3	8	6
Distance	Average distance (d^-)	3.364	3.844	3.333
	Std deviation of distance (σ)	0.661	0.848	0.655
Eccentricity	Graph radius	4	6	5
	Average eccentricity (ϵ^-)	5.522	7.443	6.291
	Graph diameter	8	11	9
Betweenness	Avg node betweenness	$8.78 \cdot 10^{-5}$	$8.57 \cdot 10^{-5}$	$6.96 \cdot 10^{-5}$
	Avg edge betweenness	$3.57 \cdot 10^{-5}$	$4.99 \cdot 10^{-5}$	$1.45 \cdot 10^{-5}$

Table 4: Table detailing graph summary statistics.

our opinion thus making it beneficial for us to study its effects on the Internet topology. Moreover, the objective of this work is a complete understanding of the IXP link effects (and not only degree distributions), which we carry out for other important topology metrics.

Joint degree distribution

The joint degree distribution gives us an idea of the general neighborhood of a randomly chosen node with an average degree. The immediate one hop neighborhood of the node gives significant information not only about the interconnections between nodes but also the structure of the area around the node. Mahadevan et al. in [10] define the joint degree distribution (JDD) as the probability

$$p(k_1, k_2) = \frac{m(k_1, k_2)}{m}$$

that a randomly selected edge connects k_1 and

k_2 -degree nodes, where $m(k_1, k_2)$ is the total number of edges connecting nodes of degree k_1 and k_2 . Figure 6 shows the JDD for the different graphs. Since CAIDA has the highest number of radial links connecting low degree customer AS nodes to high-degree provider AS nodes, it is at the top for lower node degrees. Since IXPALL contains all these nodes and links from CAIDA its behavior is very similar initially. However the effect of IXP peering is evident for medium to high degree nodes (10 to 1000). Numerous peerings between ASes at different locations worldwide result in tangential links between ASes of similar higher degrees resulting in the I X PALL graph showing consistently high values throughout the middle and latter sections of the graph. Figure 7 presents the ccdf of the average neighbor connections against average node degrees. A higher percentage of CAIDA nodes have an average neighbor degree greater than RVIEWS but the effect of the extra peering links added in IXPALL is not extensive when combined with the graphs. This is because only a small number of extra nodes with higher number of links are included, thereby not affecting the actual number of nodes. Thus we can accurately conclude that the

peering links at IXPs again significantly affect the JDD of the Internet topology graphs obtained from the traditional sources.

A summary statistic of the JDD is the average neighbor connectivity, the average neighbor degree of the average k -degree node. The average neighbor degree for the different graphs is listed in table 4. As seen in the degree distribution plots, CAIDA exhibits values greater than the BGP based graphs but the IXP peering nodes have high average neighbor degrees, which has an overall effect in increasing the average degree of the neighbor nodes in IXPALL.

Another scalar value summarizing the JDD is the assortative coefficient [32] which measures mixing patterns between nodes. The coefficient r , which lies between -1 and 1 denotes the correlation between a pair of nodes, with negative values of r indicating relationships between nodes of different degrees and positive values of r showing that nodes have correlations between nodes of the same degree. With the scale free nature of Internet, it is not surprising to see all our graphs being disassortative in nature with a high number of radial links connecting nodes of different degrees [10]. Since the traceroute based studies are unable to find a high number of tangential links, all the graphs show higher disassortative trends. However the peering links in IXPALL are the source of the tangential links between high degree nodes thereby resulting in a relatively higher assortative coefficient value.

Clustering coefficient

The value for the local clustering coefficient of a node denotes how close its neighbors are to forming a clique. This metric serves as a supplement to the JDD by providing more information about how the neighbors interconnect. If the average number of links between k -degree nodes is $m_{mn}(k)$, then the local clustering coefficient $C(k)$ is (from [10]): $C(k) = \frac{2m_{mn}(k)}{k(k-1)}$

If two neighbors of a node are also connected, then it forms one triangle while a triplet of nodes is formed when out of three nodes either two or three nodes are connected to each other. An open triplet is formed with two connections while a closed triplet is created when all the nodes are connected to each other.

The global clustering coefficient is a percentage of the number of closed triangles (made up of three closed triplets) in the entire graph over the total number of triplets in the graph.

From a high local clustering value of a node it can be inferred that its neighbors have greater interconnections which in turn leads to greater path variance. Such a characteristic would provide interesting ramifications for ASes peering at individual IXP locations. A pair of ASes would be more eager to peer if there is a potential to peer with other ASes already present at that location. With a high local clustering value, all ASes at the IXP would be able to transmit traffic to each other more efficiently through a subset of peering ASes. These highly

clustered networks would also help in the routing performance under different conditions. From table 4 we observe CAIDA to have a higher mean clustering value but IXPALL exhibits a clustering coefficient double that of the former. As mentioned in [10], this is due to greater differences in disassortativity and JDD values. In figure 8 we observe IXPALL exhibits high clustering values for lower degree nodes. These are due to the CAIDA nodes which are highly disassortative, meaning that lower degree nodes have a higher probability of being connected to high degree nodes. For higher degree nodes, the local clustering values are significantly higher. This is because the average node degree k for IXPALL nodes is much greater in comparison. The ccdf of local clustering values (figure 9) obtained reinforce the above conclusions whereby there is always a higher probability of nodes exhibiting a particular local clustering value.

Rich club connectivity

The Rich club connectivity (RCC) metric, introduced by Zhou and Mondragon in [25,33] provides an insight into the properties of power law networks. Rich nodes are a small number of nodes with large numbers of links forming a core club of nodes which are very well connected to each other. As defined in [10], if $\rho=1\dots n$ is the first ρ nodes ranked in decreasing order of node degrees, then the RCC $\phi(\rho/n)$ is the ratio of the number of links in the subgraph induced by these ρ nodes to the maximum possible links $\rho(\rho - 1)/2$. It is pointed out in [25] that the RCC is a key component in characterising Internet AS-level topologies.

Figure 10 presents the RCC for the various graphs and it can be seen that CAIDA exhibits the highest RCC values. Even though IXPALL has a greater number of links its lower RCC means that the higher degree nodes are not connected extensively with each other. The subgraphs induced from these high degree nodes do not come close to forming cliques which can be explained from the location based nature of IXPs. IXPs in general are not connected to each other and the peering links created at these locations remain localized. These peering links denote a cooperation only between a pair of nodes which are independent of other peering links. The IXPALL graph would exhibit higher RCC values if more ASes at the IXP peer with a greater number of ASes already peering there. The potential for a greater IXP utilization is evident from this result as there is an opportunity for more ASes to come up with peering agreements and ensure even better connectivity.

Node coreness

The authors in [10,34] define the k -core of a graph as the subgraph obtained from the original graph by the iterative removal of all nodes of degree less than or equal to k . The node coreness (κ) can be defined as the highest k for which the node is present in the k -core but removed in the $(k + 1)$ -core. Thus all one degree nodes have coreness equal to 0 while the maximum node coreness κ_{max} is termed the graph coreness. In this case the κ_{max} -core of the graph is not empty but the $(\kappa_{max} + 1)$ -core is. The graph fringe is defined as the set of nodes in the graph displaying minimum coreness κ_{min} .

The node coreness is a more advanced version of node connectivity than the node degree as it tells us how well the node is connected to the entire graph. A node may have a high degree but its connectivity to other parts of the graph is dependent largely on its neighbors. The best example to describe this is a high degree hub of a star which has a coreness of 0 with its neighbors only having a very low degree (one), which when removed leaves the hub disconnected.

From table 4 we observe IXPALL exhibits significantly higher

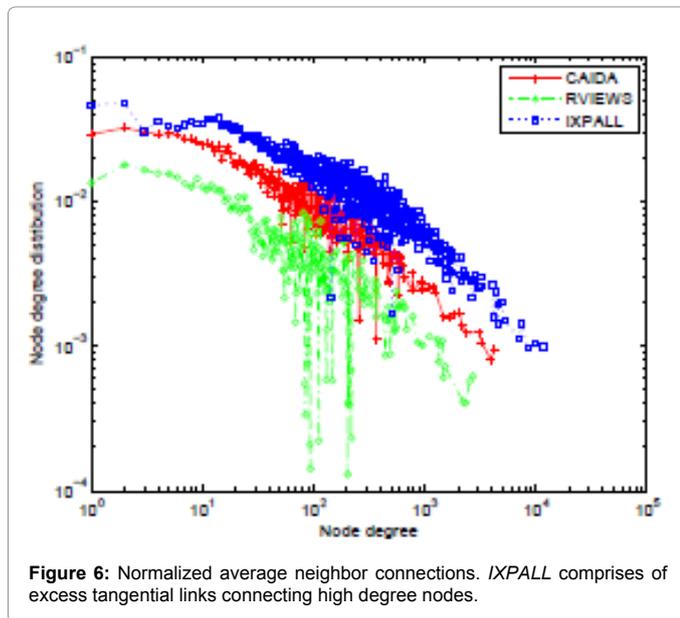


Figure 6: Normalized average neighbor connections. IXPALL comprises of excess tangential links connecting high degree nodes.

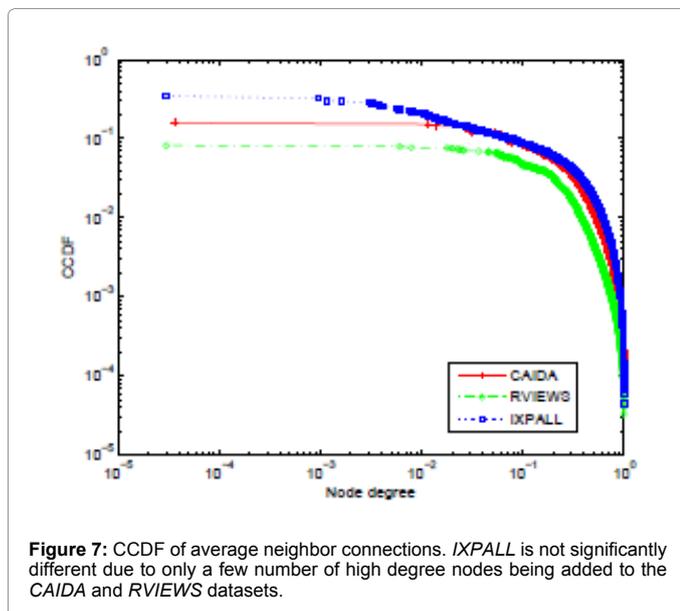


Figure 7: CCDF of average neighbor connections. IXPALL is not significantly different due to only a few number of high degree nodes being added to the CAIDA and RVIEWWS datasets.

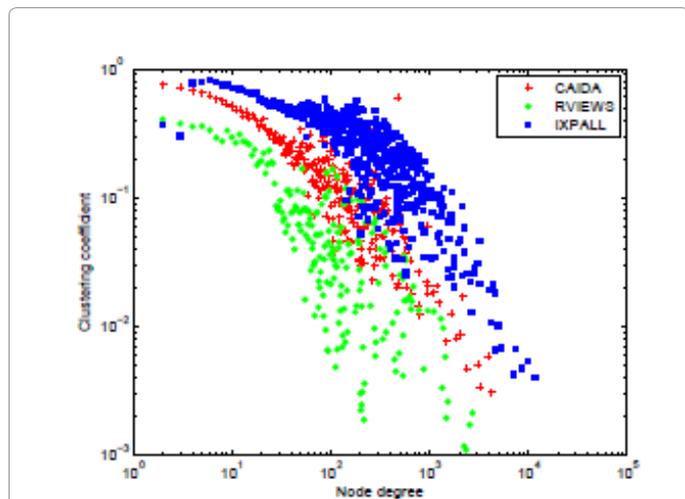


Figure 8: Local clustering with increasing node degrees. *IXPALL* exhibits constant high clustering values due to a high number of links being clustered at the IXP nodes.

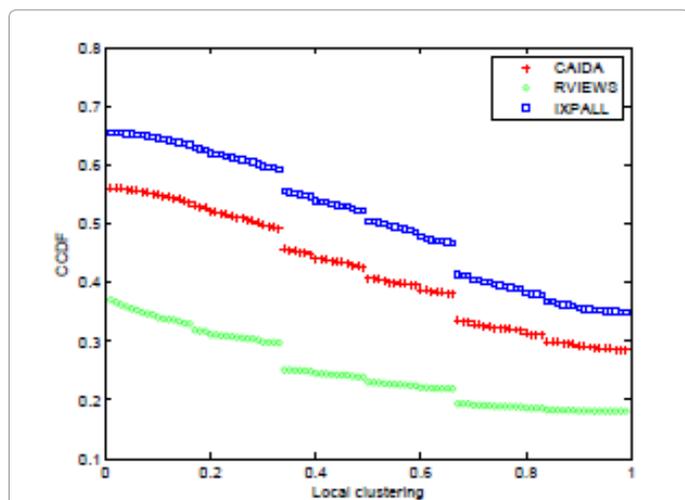


Figure 9: CCDF of local clustering values. *IXPALL* shows a consistently high probability for all clustering values considered.

average node coreness (k) and maximum coreness (k_{max}) values. The core size ratio is also higher indicating the general higher general connectivity due to IXP links induced in the graph. Figure 11 displays this result showing the effect of the IXP peering links increasing the overall coreness for nodes with all low, medium and high degrees. It is also evident from the figure that the increase in node coreness follows a power law increase for nodes upto degrees of 100 before remaining stable for higher degree nodes. Likewise the fringe size ratio is also the lowest in *IXPALL* which means fewer nodes with minimum coreness thereby leading to a better connected graph than the two others. The coreness result presents an important characteristic: the fact that the greater number of links also leads to better connectivity. These new links are not all only tangential links between low degree nodes but contain a generous amount of radial links leading to better node connectivity.

Distance and eccentricity

The distance distribution $d(x)$ is the probability for a pair of

random nodes to be at a distance of x hops within each other whereas eccentricity is the maximum distance between the pair of nodes. Thus the maximum eccentricity in a graph is also the maximum distance and is termed the graph diameter. This metric is important while designing efficient routing policies to enable paths with lesser hops to be chosen. The authors in [10] also point out that the distance distribution plays a major role in helping the network recover from virus attacks. Figure 12 presents the distance distribution values of the three graphs studied. We observe that about 55 percent of nodes in *IXPALL* are separated by a distance of 5 hops while it is lower for the other graphs. Even though *IXPALL* has a greater number of links (which means that average distances should decrease), the average distance value is greater suggesting that deployment of IXPs do not decrease the path lengths between end-hosts on the Internet. There could be routing performance efficiencies through IXP deployment but the number of hops traversed largely remain the same. Figure 13 shows that maximum distances for a majority of the nodes are similar across all graphs with almost 70

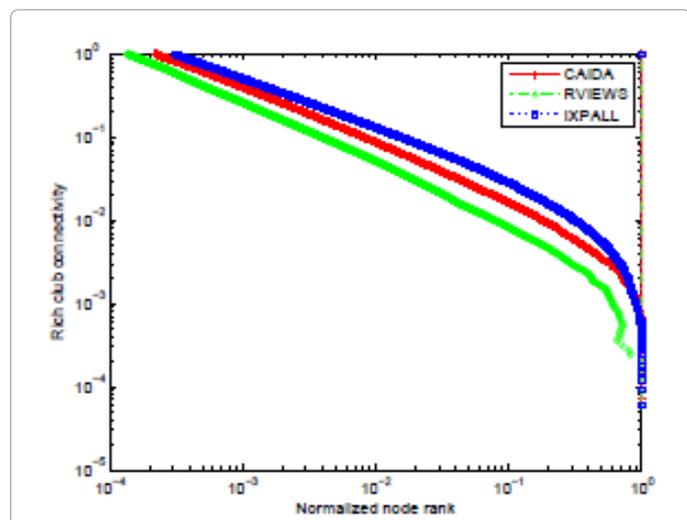


Figure 10: CCDF of the Rich club connectivity (RCC) for the three graphs. The highest connectivity among high degree nodes is in *CAIDA* while *IXPALL* high degree nodes are not connected between themselves.

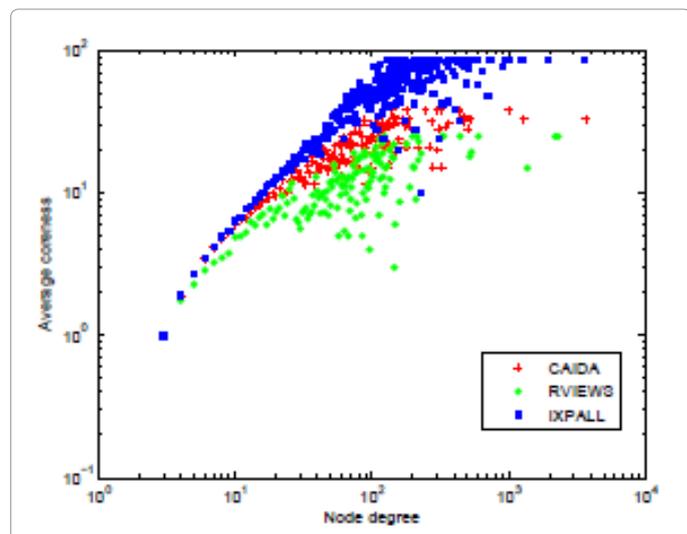


Figure 11: Average node coreness with increasing node degrees. The increase in coreness roughly follows a power law for all graphs for low and medium and degree nodes before becoming stable.

percent of IXPALL nodes separated by a maximum of 6 hops from each other.

Betweenness

The most common and effective means of measuring node centrality is betweenness. Nodes which appear on a greater number of shortest paths between any pair of nodes in the graph exhibit a higher betweenness value. Such nodes are considered to be more central than others since it is assumed that majority of the traffic on a network is sent along the shortest path from source to destination. Potential traffic load on nodes/links may be estimated from betweenness values of certain critical nodes which would also point to locations for potential congestion. Using a relatively quick algorithm [35] to calculate the betweenness centrality of the nodes, we obtain the normalized betweenness distribution with increasing node degrees. Since the maximum number of paths possible in a graph is $n(n-1)$, all the graphs are normalized by this value and the results shown in

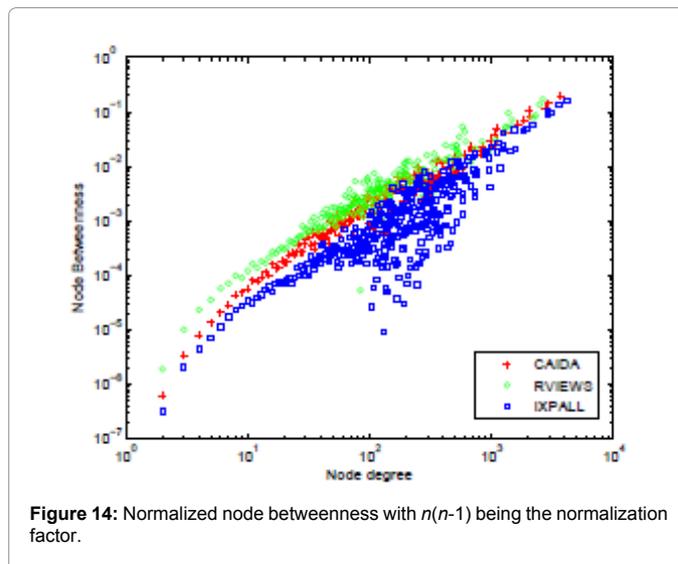


Figure 14: Normalized node betweenness with $n(n-1)$ being the normalization factor.

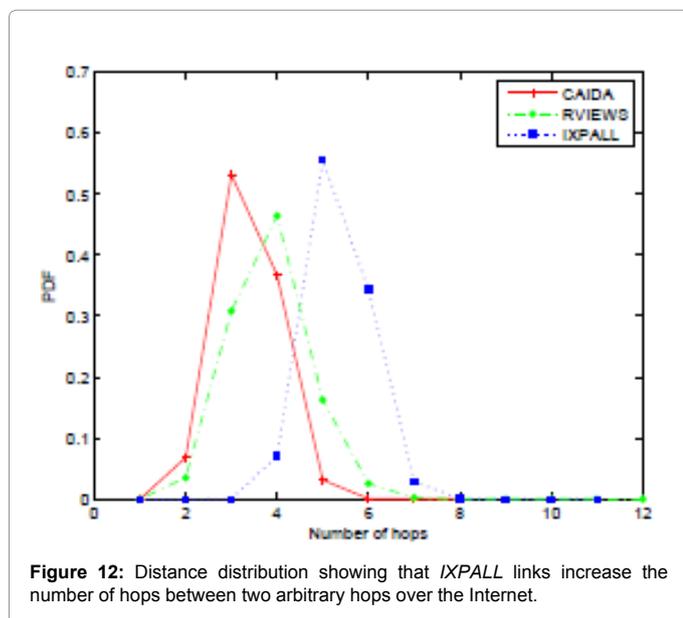


Figure 12: Distance distribution showing that IXPALL links increase the number of hops between two arbitrary hops over the Internet.

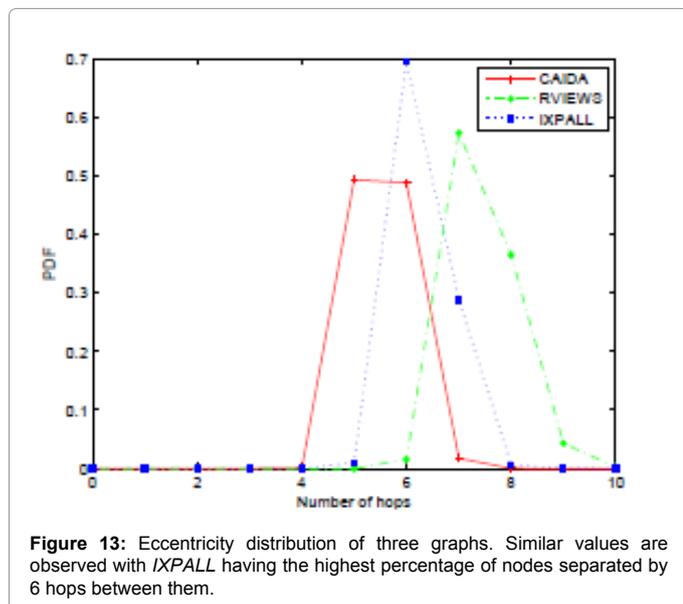


Figure 13: Eccentricity distribution of three graphs. Similar values are observed with IXPALL having the highest percentage of nodes separated by 6 hops between them.

figure 14. It can be observed that all three graphs exhibit a power-law function of node betweenness with increasing node degrees with IXPALL exhibiting lower values overall. Higher numbers of nodes of all degrees (mainly medium degrees) in IXPALL leads to greater path diversity. This means there is a presence of a greater number of nodes for paths of equal distance between all pairs of nodes leading to the lower betweenness values observed.

Continuing from the node betweenness values exhibited in figure 14 we compute the normalized edge betweenness for the graphs and present the results in figures 15. The figure shows the CDF of the log of betweenness values for all edges in the three graphs. It can be seen that IXPALL has the highest percentage of edges with the lowest edge betweenness values (as is evident in the scatter plot in figure 16 of edge centrality, with a high concentration of points with very low betweenness values). This means that a high percentage of IXP peering edges (along with nodes) do not fall on the available shortest paths between nodes in the entire graph. It has to be noted here that inter-domain routing in the Internet does not follow conventional shortest path approaches and is actually determined by inter-ISP routing policies and hot-potato routing in BGP. Betweenness can thus not be considered as an entirely accurate indicator of Internet path performance except to give an idea of the relative importance of the nodes/edges along a shortest path. We may conclude from this result that IXPs do not necessarily decrease the hop count of paths between ASes peering at those locations as path lengths essentially remain similar to other established paths from source to destination AS.

Analysis and Discussions

Combining the extra peering links visible at the IXPs with the general structure of the Internet has given us a varied set of characteristics of the completed picture of the Internet (the data plane combined with the control plane and the peering links). Comparing the derived topologies based on the available graph metrics gives us an insight into the effects the peering links uncovered at the IXPs are having on the topology evolution of the Internet.

The most widely studied node degree distribution behavior of the Internet remains essentially unchanged even after the addition of all the peering links. The scale-free nature of the Internet graph, based on the different views considered, does remain the same. Numerous instances of related work have noted that the IXP peering links hold

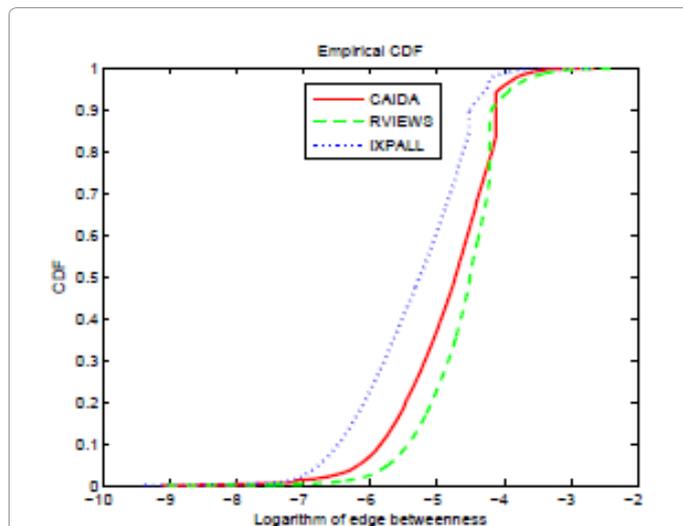


Figure 15: CDF of log of edge betweenness for the three graphs. IXPALL has the highest percentage of edges with the lowest edge betweenness values.

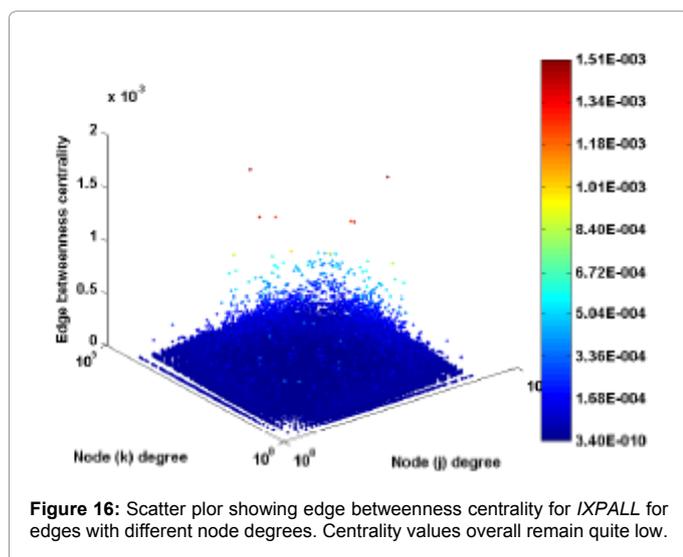


Figure 16: Scatter plot showing edge betweenness centrality for IXPALL for edges with different node degrees. Centrality values overall remain quite low.

the key to solving the missing links problem and our findings suggest peering links provide a part of the solution to the problem. However it has to be mentioned that there has been work following the famous paper by Faloutsos et al. [27] which have discounted the scale-free nature of the Internet [29-31] due to inherent biases in the traceroute mechanisms. Observing the effects of IXP peering links on other important metrics leads to some interesting insights. Higher JDD values for medium to high degree nodes means that well connected ASes (and providers, preferably the higher tier ISPs) are setting up peering relationships at exchange points. Such peering links lead to higher average neighbor degrees. A generous mix of both tangential and radial links are evident in the IXPALL graph unlike in CAIDA where there is a high number of radial links connecting nodes of vastly different degrees. The high JDD also comes with high levels of local clustering due to IXP peering links. This characteristic should and does directly serve to provide an incentive to ASes to peer at an IXP. A high number of links inevitably leads to greater local clustering but the RCC on the other hand displays the fact that there is little interconnection between the IXPs between themselves. Such connections between

IXPs are however not needed since they are constructed to provide a platform for local interconnectivity amongst coordinating ASes.

The node coreness metric which points out how "deep in the core" the node is situated [10], shows that the nodes in the IXPALL graph are mostly well connected with well connected neighbors. The IXP substrate has thus become an important component of the Internet's infrastructure leading to a 'flatter' Internet from a hierarchical one. Gill et al. in [9] reported the changing characteristic of the Internet to a more "flat" architecture which can be inferred by the results obtained by us with the coreness metric. The greater number of peering links between ASes at IXPs lead to those ASes getting deeper into the core of the Internet with decreasing emphasis on connections with upper tier ASes. The authors in [36,37] have all pointed towards this evolution characteristic of the Internet and our coreness metric based result presents a theoretical confirmation of these observations.

Node and edge betweenness are two measures of centrality from which further inferences can be made about the effects of IXPs peering links. Both these metrics point towards lower values for IXPALL which means not many AS-AS peering links are a part of the shortest paths between ASes. Zheng et al. [38] show that routing policies and the layer 2 technology used on peering links may lead to cases of Triangle Inequality Violations (TIVs) [39,40] in the Internet and not necessarily provide significant savings on RTT measurements between ASes. With most detour paths [41] forming TIVs, peering links do not necessarily lead to shorter paths along the Internet. The results we obtain again largely confirms this Internet path characteristic from a theoretical perspective.

Coming back to the questions we posed at the beginning of the paper we observe that IXP links indeed play a major role on topology characteristics of the Internet. Their effects on various important topology metrics should make the Internet topology research community stand up and take notice of this integral component and give due attention to uncovering more peering links at IXP locations worldwide. While Augustin et al. in [6] present a first step in carrying out a comprehensive study to uncover peering links, there is no sustained effort in the community to continue such studies at the moment. On the other hand, the flattening of the Internet topology structure [9] shows the growing trend among the ASes to move away from higher tier transit ISPs towards creating inter-AS peering links. These characteristics and the incredibly high number of IXP peering links point towards the fact that IXPs are indeed the key towards solving the missing links problem and with their addition to the visible Internet topology we will go a long way to verifying the validity of topology generators and evolution models.

Conclusions

We have discussed the addition of IXPs as an important component of the Internet's AS level ecosystem and analyzed various graph based properties of the Internet after incorporating a set of peering links unearthed at various IXPs around the world. Our studies confirm that IXPs are an integral part of the Internet and the peering links being created at these locations do in fact have a relatively large impact on the growth and evolution of the Internet.

Our work is bound by a set of limitations which we identify and aim to redress in future work. Firstly, the validity and importance of the IXP peering links dataset examined [6] in this work was released around October 2009 and is hence subject to changes over a period of time. While the characteristic is certainly true and newer datasets would be more helpful, we generate graphs of the Internet from

historical CAIDA and RouteViews data for that same time period. Our focus lies on analysis for that period of time for which we believe our dataset to be valid and consistent. We also limit our data collection to a single month which in our opinion provides us with a representative snapshot of the Internet topology for that period. However this would not be enough to carry out evolution studies over a period of time. We also use a set of graph based metrics which we believe is important and has been used in various other topology studies, but it is certainly not an exhaustive list.

There are also variations of these metrics which provide different views of the results obtained, something which we leave as a course for future work.

Overall through this work we emphasize the importance of IXPs while studying the Internet's growth. With topology evolution playing a major role in the development and implementation of future Internet architectures, IXPs and peering links within these locations play a pivotal role in understanding and solving the missing links problem. This paper presents and analyzes some of the important effects IXP peering links are having on the Internet topology and we hope it encourages the implementation and validation of newer and more accurate topology generators with the ultimate goal of a more efficient Internet routing architecture for the future.

References

- Haddadi H, Uhlig S, Moore A, Mortier R, Rio M (2008) Modeling internet topology dynamics. SIGCOMM Comput Commun Rev 38: 65–68.
- Kushman N, Kandula S, Katabi D, Maggs B (2007) R-BGP: Staying Connected in a ConnectedWorld. 4th USENIX Symposium on Networked Systems Design & Implementation, Cambridge, UK.
- Dhamdhere A, Dovrolis C (2010) The internet is flat: modeling the transition from a transit hierarchy to a peering mesh. Proceedings of the 6th International Conference, Co-NEXT '10, ACM, New York, NY, USA.
- He Y, Siganos G, Faloutsos M, Krishnamurthy S (2009) Lord of the links: a framework for discovering missing links in the internet topology. IEEE/ACM Trans Netw 17: 391–404.
- Oliveira RV, Pei D, Willinger W, Zhang B, Zhang L (2008) In search of the elusive ground truth: the internet's as-level connectivity structure. SIGMETRICS Perform Eval Rev 36: 217–228.
- Augustin B, Krishnamurthy B, Willinger W (2009) IXPs: mapped? Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference, ACM, USA.
- Routeviews routing table archive.
- Huffaker B, Hyun Y, Andersen D, Claffy K (2009) Skitter as links dataset.
- Gill P, Arlitt M, Li Z, Mahanti A (2008) The flattening internet topology: natural evolution, unsightly barnacles or contrived collapse? Proceedings of the 9th international conference on Passive and active network measurement, Springer-Verlag, Berlin, Germany.
- Mahadevan P, Krioukov D, Fomenkov M, Huffaker B, Dimitropoulos X, et al. (2005) Lessons from three views of the internet topology.
- Zhang B, Liu R, Massey D, Zhang L (2005) Collecting the internet AS level topology. SIGCOMM Comput Commun Rev 35: 53–61.
- Young Hyun EA, Huffaker B, Luckie M (2009) The ipv4 routed /24 topology dataset.
- Madhyastha HV, Isdal T, Piatek M, Dixon C, Anderson T, et al. (2006) iplane: An information plane for distributed services. Proceedings of the 7th symposium on Operating systems design and implementation USA.
- Shavitt Y, Shir E (2005) Dimes: let the internet measure itself. SIGCOMM Comput Commun Rev 35: 71–74.
- Packet clearing house. PCH.
- Public exchange points search/list.
- Xu K, Duan Z, Zhang ZL, Chandrashekar J On properties of internet exchange points and their impact on as topology and relationship.
- Govindan R, Reddy A (1997) An analysis of internet inter-domain topology and route stability. Proceedings of Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies.
- Mahadevan P, Krioukov D, Fall K, Vahdat A (2006) Systematic topology analysis and generation using degree correlations. Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications.
- Chang H, Govindan R, Jamin S, Shenker SJ, Willinger W (2004) Towards capturing representative AS-level internet topologies. Computer Networks: The International Journal of Computer and Telecommunications Networking, USA.
- Oliveira R, Pei D, Willinger W, Zhang B, Zhang L (2010) The (in)completeness of the observed internet AS-level structure. IEEE/ACM Transactions on Networking 18: 109–122.
- He Y, Siganos G, Faloutsos M, Krishnamurthy SV (2007) A systematic framework for unearthing the missing links: Measurements and impact. NSDI'07 Proceedings of the 4th USENIX conference on Networked systems design & implementation, USA.
- Oliveira RV, Zhang B, Zhang L (2007) Observing the evolution of internet as topology. Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, USA.
- Bu T, Towsley D (2002) On distinguishing between internet power law topology generators. Proceedings of Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies.
- Zhou S, Mondragon RJ (2004) Accurately modeling the internet topology. Phys Rev E 70, 066108.
- Gregori E, Improta A, Lenzini L, Orsini C (2011) The impact of ixps on the as-level topology structure of the internet. Computer Communications 34: 68–82.
- Faloutsos M, Faloutsos P, Faloutsos C (1999) On power-law relationships of the internet topology. Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication, USA.
- Barabasi AL, Albert R (1999) Emergence of Scaling in Random Networks. Science 286: 509–512.
- Lakhina A, Byers J, Crovella M, Xie P (2003) Sampling biases in IP topology measurements. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications.
- Achlioptas D, Clauset A, Kempe D, Moore C (2005) On the bias of traceroute sampling: or, power-law degree distributions in regular graphs. STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing, ACM, New York, NY, USA 694–703.
- Willinger W, Alderson D, Doyle JC (2009) Mathematics and the Internet: A Source of Enormous Confusion and Great Potential. Notices of the American Mathematical Society 56: 586–599.
- Newman MEJ (2002) Assortative mixing in networks. Phys Rev Lett 89: 208701.
- Zhou S, Mondragon R (2004) The rich-club phenomenon in the internet topology. Communications Letters, IEEE 8: 180–182.
- Gaertler MD, Patrignani M (2004) Dynamic analysis of the autonomous system graph. International Workshop on Inter-domain Performance and Simulation.
- Brandes U (2001) A faster algorithm for betweenness centrality 25: 163–177.
- Haddadi H, Fay D, Uhlig S, Moore A, Mortier R, et al. (2010) Mixing biases: Structural changes in the AS topology evolution. Proceedings of the Second international conference on Traffic Monitoring and Analysis, Germany.
- Fay D, Haddadi H, Thomason A, Moore AW, Mortier R, et al. (2010) Weighted spectral distribution for internet topology analysis: theory and applications. IEEE/ACM Trans Netw 18: 164–176.
- Zheng H, Lua EK, Pias M, Griffin TG (2005) Internet routing policies and round-trip-times. PAM 3431: 236–250.
- Wang G, Zhang B, Ng TSE (2007) Towards network triangle inequality violation aware distributed systems. Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, USA.

-
40. Lumezanu C, Baden R, Spring N, Bhattacharjee B (2009) Triangle inequality and routing policy violations in the internet. Proceedings of the 10th International Conference on Passive and Active Network Measurement, USA.
41. Gummadi KP, Madhyastha HV, Gribble SD, Levy HM, Wetherall D (2004) Improving the reliability of internet paths with one-hop source routing. Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation, USA.