

Impact of Internet Exchange Points on Internet topology evolution

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

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 from a simple visual analysis and continuing with a more advanced graph based metric analysis, 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 provide us a better understanding of Internet topology evolution.

I. INTRODUCTION

The rapidly evolving Internet topology has led to a plethora of efforts in trying to characterize the dynamics and trends of this growth and evolution. This is due to an underlying need to better understand the constraints of the current routing protocols and come up with more efficient designs to push forward the development of the next generation Internet. Analysis of the Internet topology is needed for better network planning and designing optimal routing strategies [1]. 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 Autonomous Systems (ASes) to peer directly with other member ASes at these locations. Increased peering at these IXP switches have led to more research ([2]–[4]) showing a significant number of new links being uncovered at these locations which in turn impacts our understanding of the Internet topology at the AS level. Augustin et al. [4] 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 [2] that the extra peering links at these IXPs may hold the key to solving the *missing links* problem for the AS-level Internet and [4] 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 analyzed in detail with respect to the existing Internet topology and their effects measured before a final conclusion can be drawn. Any number of questions arise: (1) Do the extra IXP links uncovered have a significant effect on the growing topology dynamics of the Internet? (2) If the effects of these links are significant then how do we change our outlook in conducting topology research to accommodate these newer changes? (3) 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?

The primary contribution of this paper is to carry out graph based studies aimed at finding an answer to these questions and open up interesting challenges for future directions. We choose a set of metrics discussed in [5] to study and analyze the topological properties of the Internet from standard data sources in addition to the extra peering links obtained at the IXPs. Topological graphs of the Internet from a period of 31 days in October 2009 is first created from CAIDA’s skitter dataset [6], the RouteViews BGP tables [7], the BGP tables released by PCH [8] and the peering links observed by the IXP mapping project [4]. We carry out a complete analysis of the Internet topology based on specific graph-based metrics. 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.

II. IXPs AND TOPOLOGY EVOLUTION

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 [3]. The public peering at the IXP then becomes simpler due to the availability of physical infrastructure with member ASes A and B (as shown in figure 1) 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




Fig. 1. 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.

routers in the Internet cloud for it to be able to successfully transfer data to F.

III. 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. in [5] 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.

A. 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 *BGP* graph in the rest of the paper.

CAIDA's IPv4 Routed /24 topology dataset [9] 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 [5] private ASes generate indirect links which we filter out during creation of the AS-level graphs which are then combined to form the final *Skitter* graph.

PCH [8] 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 obtain the *PCH* graph from these daily graphs.

The IXP Mapping project [4] releases data specific to IXPs across the Internet with only peering links unearthed at these IXPs. We term this dataset *IXPMAP*. To study their effect on the observed topology we combine this with the *PCH* graph to form the *IXPALL* dataset.

IV. RESULTS

A. Number of nodes and edges

From table I we observe that the BGP tables find a higher number of ASes (nodes) than *Skitter* while the *IXPMAP* dataset only has just a tenth of the number of nodes visible in the others. This is due to the fact that only a certain number of ASes peer at the IXPs and ASes not having any peering relationship are not visible in this dataset. The total number of links visible in *IXPMAP* is almost 52K while only 3756 nodes are present. Combining these links with the *PCH* for 31 days, there more than 45K links visible only in the combined *IXPALL* graph. With 366 unique nodes in the combined graph, there could be a maximum of $\frac{n(n-1)}{2}$ links, which roughly equals 66K links in a complete graph assuming every AS peers with all other ASes at the IXPs. We find about 45K links denoting a high level of peering prevalent at the IXPs.

B. Node degree distribution

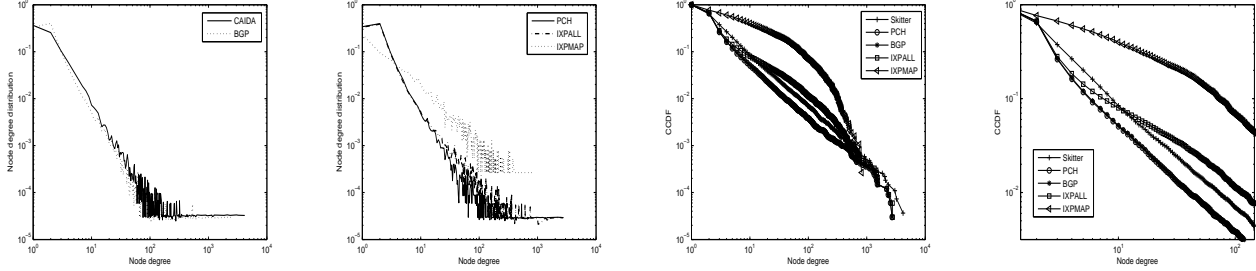
The node degree distribution is the probability distribution of the node degrees in a graph. In other words, it is the probability that a node selected randomly is of k -degree and this probability is calculated by: $P(k) = \frac{n(k)}{n}$, where $n(k)$ is the number of k -degree nodes in a graph with total number of nodes n . Scale-free networks such as the Internet have been shown to exhibit power law degree distributions [10] and hence the power law exponent is computed for this metric. Due to the closely spaced nature of the plots, we divide the plots into separate figures (2(a) and 2(b)) for better understanding and readability.

The results show that the different topologies all show a power law behavior with the average node degree in the BGP based graphs to be lesser than the *Skitter* graph. With the addition of the higher number of IXP links, *IXPALL* ends up with a greater average node degree than *Skitter*. The power law exponent γ (where $P(k) = k^{-\gamma}$) also remains high for the *skitter* graph. A novel observation is that the *IXPMAP* dataset also exhibits a significant power law degree distribution greater than the others with a higher power law exponent. The authors in [5] point out that a natural cut-off at power-law maximum degree is obtained at: $k_{max}^{PL} = n^{\frac{1}{\gamma-1}}$. From table I we observe that the maximum node degree k_{max} for the combined graph 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

TABLE I
COMPARING BASIC GRAPH PROPERTIES BETWEEN DIFFERENT DATA SOURCES

Property description	<i>Skitter</i>	<i>BGP</i>	<i>PCH</i>	<i>IXPMAP</i>	<i>IXPALL</i>
Number of nodes	27357	33199	33266	3756	33527
Number of edges	86918	77101	75961	51973	123627
Average node degree	6.59	4.64	4.56	27.67	7.37
Power law exponent (γ)	2.47	2.25	2.18	2.91	2.28
Maximum node degree	4185	2717	2706	821	2772
Power law maximum	1839	1032	6795	74	3431
Maximum degree ratio	0.152	0.081	0.081	0.218	0.082
Average neighbor degree	0.024	0.0154	0.0152	0.049	0.0162
Top clique size	13	7	7	10	7
Clustering coefficient	0.0193	0.0105	0.010	0.267	0.085
Assortativity coefficient	-0.1714	-0.2015	-0.2007	-0.1686	-0.1805



(a) Degree distribution for *Skitter* and *BGP*. Power law behavior is apparent from the nature of the graph. (b) Degree distribution for *PCH*, *IXPALL* and *IXPMAP*. *IXPALL* displays power law characteristics. (c) Complete CCDF of node degree distribution for all datasets (d) CCDF of node degree distribution for all datasets with node degree less than 100

Fig. 2. Node degree distributions vary according to power law characteristics for all graphs constructed. From Table 1 we observe the power law exponent for *IXPALL* remains similar to the *Skitter* and *BGP* graphs, thereby denoting IXP links to conform to Internet growth characteristics.

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. Figure 2(c) show that each of the graphs are very similar with the *BGP* based graphs showing almost identical characteristics. A closer look at the number of nodes with degree lesser than 10 (fig. 2(d)) shows that about 8 percent of nodes in *IXPALL* have degree less than or equal to 10 while the corresponding percentage for the *IXPMAP* set is almost 40 percent. In comparison to the *Skitter* graph there is initially a lesser percentage of lower degree nodes but it gradually increases as the effect of the additional IXP peering links come into play for node degrees between 10 and 100 with a skew towards higher degree nodes.

C. 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 [5] define the joint degree distribution (JDD) as the probability that a randomly selected edge connect k_1 and k_2 -degree nodes: $P(k_1, k_2) = \frac{m(k_1, k_2)}{m}$, where $m(k_1, k_2)$ is the total number of edges connecting nodes of degree k_1 and k_2 . Figure 3(a) shows the JDD for the *Skitter* and *BGP* graphs while figure 3(b) presents the results for

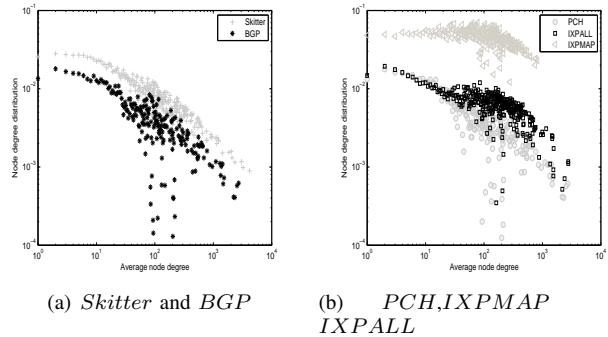


Fig. 3. CCDF with average neighbor node degree. *IXPALL* shows a higher JDD due to high degree of nodes in *IXPMAP*

the others. We observe that *Skitter* has a high number of links which connect lower degree nodes to high degree nodes. These links are mainly radial links connecting customer ASes to provider ASes which have a high degree and so it enables the *skitter* graph to remain higher than the *BGP* graph. Figure 3(b) shows that almost all nodes peering at the IXPs have high degree neighbors which thereby result in the combined *IXPALL* graph to show a higher joint degree distribution as well. Infact the highest number of nodes with larger neighbor node degrees at IXPs is for those nodes with degrees greater than 100.

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 datasets are listed in table I. As seen in the degree distribution plots, *Skitter* exhibits values greater than the BGP based graphs but the *IXP MAP* nodes have much higher average neighbor degrees, which has an overall effect in increasing the average degree of the neighbor nodes in *IXP ALL*.

Another scalar value summarizing the JDD is the assortative coefficient [11] 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 [5]. Since the traceroute based studies are unable to find a high number of tangential links all the graphs show a higher disassortative trends including the links obtained in *IXP MAP*.

V. DISCUSSION AND CONCLUSIONS

We analyze the different AS-level topology graphs of the Internet obtained: the *Skitter* graph based on measurements from CAIDA, the BGP tables obtained from RouteViews and IXP routers from PCH, the extra links uncovered by the IXP mapping project by [4] and finally the topology snapshot created by combining the PCH and IXP mapping graphs. The numerous additional links affect the existing graph topologies but the overall power law characteristic of node degree distributions remain similar to existing results.

Basic topology measurement metrics such as average node degree and node degree distribution show that these values are affected significantly with the addition of new, uncovered links at the IXPs. This means that these links may not be ignored during the construction of topology evolution models or topology generators. However, the addition of these links does not affect the power law behavior of the degree distribution. The new links uncovered exhibit an increase based on a power law with a higher exponent while the combined graph also shows power law characteristics with minor modifications to the previous growth parameters. Joint degree distributions denoting the effect of the new links on average degree neighbors are also significant on a particular set of nodes with higher degrees but the effect on an overall percentage of the nodes is not considerable. However the high levels of clustering associated with these links tend to effect the clustering metrics of the observed Internet topologies. The comparatively lower number of nodes at these IXP locations coupled with high number of extra peering links present an intriguing scenario to figure out if the effects are indeed consequential or not in the larger view of the complex network of the Internet.

A study such as this is bound by some limitations as well. Firstly, we analyze data for a period of a single month only and assume that it would hold up as a representative snapshot of the different AS level topologies of the Internet. While historical

data is applicable, we believe that data from 31 consecutive days would present a stable view of the Internet from the different sources. The newer links uncovered and used in this paper is made available by the authors of [4] but which is not updated since the completion of their work. There would inevitably be numerous more links and nodes being created or removed in this dataset. As the authors state, their work leaves little scope for finding any more hidden links at IXPs worldwide, but what is needed is to keep the dataset updated.

A study comprising of these metrics can be categorized only as initial study into the effects that these IXP links are having on the overall study of Internet topology evolution and more work is needed to come to definite conclusions about how the growth of the Internet is being affected by this increased peering between ASes at IXPs worldwide. Ongoing work being carried out by us is a more extensive graph analysis using advanced metrics such as distance, coreness, spectrum and betweenness associated with these graphs. These are additional advanced metrics defined by the authors in [5] and will go a long way towards helping us understand the nature of the combined graph. We plan to report these additional results in a future work.

REFERENCES

- [1] H. Haddadi, S. Uhlig, A. Moore, R. Mortier, and M. Rio, "Modeling internet topology dynamics," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 65–68, 2008.
- [2] 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.
- [3] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "In search of the elusive ground truth: the internet's as-level connectivity structure," *SIGMETRICS Perform. Eval. Rev.*, vol. 36, no. 1, pp. 217–228, 2008.
- [4] 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.
- [5] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. Claffy, and A. Vahdat, "Lessons from three views of the internet topology," Aug 2005.
- [6] B. Huffaker, Y. Hyun, D. Andersen, and K. Claffy, "Skitter as links dataset - jan 2000 to may 2009," available at: caida.org/data/active/skitter_aslinks_dataset.xml.
- [7] "Routeviews routing table archive," available at: routeviews.org.
- [8] "Packet clearing house," available at: pch.net.
- [9] E. A. Young Hyun, Bradley Huffaker and M. Luckie, "The caida ipv4 routed /24 topology dataset - jan 2009," available at: caida.org/data/active/ipv4_routed_24_topology_dataset.xml.
- [10] 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.
- [11] M. E. J. Newman, "Assortative mixing in networks," *Phys. Rev. Lett.*, vol. 89, no. 20, p. 208701, Oct 2002.