A Framework for Modeling Internet Topology Dynamics and Discovering Missing Links in the Internet Topology

M.Saiprasad[1], A.Srinivas[2], T.Mahesh[3], R.Satya Teja[4]

[1]M.Tech (CSE), Sree Chaitanya College of Engineering, Karimnagar, AP,India
[2]Assistant professor in CSE Department, Sree Chaitanya College of Engineering , Karimnagar, A.P, India
[3] Assistant professor in CSE Department, Vivekananda Institute of Technology and Science SET, Karimnagar, AP, India
[4]M.Tech (CS), Vivekananda Institute of Technology and Science SET, Karimnagar, AP, India

Abstract— The lack of an accurate representation of the Internet topology at the Autonomous System (AS) level is a limiting factor in the design, simulation, and modeling efforts in inter-domain routing protocols. In this paper, we design and implement a framework for identifying AS links that are missing from the commonly-used Internet topology snapshots. We apply our framework and show that the new links that we find change the current Internet topology model in a non-trivial way. An accurate and complete model of the topology would be important for protocol design, performance evaluation and analyses. The goal of our work is to develop methodologies and tools to identify and validate such missing links between ASes. In this work, we develop several methods and identify a significant number of missing links, particularly of the peer-to-peer type. Interestingly, most of the missing AS links that we find exist as peer-to-peer links at the Internet Exchange Points (IXPs). First, in more detail, we provide a large-scale comprehensive synthesis of the available sources of information. We cross-validate and compare BGP routing tables, Internet Routing Registries, and traceroute data, while we extract significant new information from the less-studied Internet Exchange Points (IXPs). We identify 40% more edges and approximately 300% more peer-to-peer edges compared to commonly used data sets. All of these edges have been verified by either BGP tables or traceroute. Second, we identify properties of the new edges and quantify their effects on important topological properties. Given the new peer-to-peer edges, we find that for some ASes more than 50% of their paths stop going through their ISPs assuming policy-aware routing. A surprising observation is that the degree of an AS may be a poor indicator of which ASes it will peer with.

Index Terms— Internet, inter-domain, measurement, BGP missing links, routing, topology.

I INTRODUCTION

An accurate and complete model of the Internet topology is critical for future protocol design, performance evaluation, simulation and analysis [1]. The current initiatives of rethinking and redesigning the Internet and its operation from scratch would also benefit from such a model. However, it remains as a challenge to develop an accurate representation of the Internet topology at the AS level, despite the recent flurry of studies [2]–[9]. Currently, there is a list of sources that contain such topological information. The list includes archives of BGP routing tables, archives of BGP routing updates, Internet Routing Registries, and archives of trace route data. Each of these sources has its own advantages, but each of them also provides an incomplete, sometimes inaccurate view of the Internet AS topology, while these sources are often complementary. Furthermore, as far as we know, IXPs (Internet Exchange Points) have not received attention in terms of Internet topology discovery, although they play a major role in the Internet connectivity.

There are two major contributions in this work. First, we design and implement a systematic framework for discovering missing links in our current Internet topology snapshot, and provide two novelties compared to previous studies—the comprehensive synthesis of different data sources and the extraction of topological information from IXPs. Second, we apply our framework and conduct an in-depth study of the importance of these new links, and improve our understanding of the Internet topology at the AS level.

In more detail, our framework first identifies and validates a significant number of AS links by a careful cross-reference and synthesis of most known sources of information: BGP tables, traceroute, and IRR [10].1 Second, our framework extracts significant new topological information from Internet Exchange Points (IXPs); such information is typically not used in topological studies. While prior work [11] has proposed methods to identify participating ASes at IXPs, our study greatly extends their work and overcomes certain limitations.

Note that we set a highly selective standard in our framework: we only accept edges which are verified by BGP tables or from traceroute data. In other words, we do not provide a union of the existing sources of information, but a critical synthesis. To achieve this goal, we develop a large scale traceroute-based tool, RETRO, to confirm the existence of edges, which we suspect exist.

We arrive at several interesting observations. First, we find a significant number of new edges, including 40% more edges (15%) and approximately 300% more peer-to-peer edges (65%) as compared to the widely used Oregon Routeviews data set (all available BGP routing tables, respectively). Second, most of the newly discovered edges are peer-to-peer edges: the current topological models have a bias by under-representing peer-to peer edges. Third, most missing peer-to-peer AS links that we find are at the IXPs: Our results show that nearly 95% of the peer-to-peer links missed from the BGP tables are incident at IXPs. This suggests that exploring the connectivity at IXPs may help us identify hidden edges between ASes that participate at IXPs. Fourth, IRR is a good source of hints for finding new edges, especially after it is filtered using a state of the art tool [12] for this purpose.

We find that the new edges significantly change our view of the Internet AS topology, and we also identify interesting patterns of the new edges. First, the new edges change the models of Internet routing and financial implications that previous research studies may have arrived at by using the incomplete topology models.

We quantify the routing decision changes in the routing model due to the peer-to-peer edges not considered previously. We find that for some ASes (mostly of degrees 10 to 300), more than 50% of their paths stop going through a provider, compared to a less complete topology. The financial implication is that these ASes maynot pay their providers to the extent that was earlier expected. Clearly, business-oriented studies should consider all peer-to peer edges for accurate results. Second, we find that provider-customer and peer-to-peer edges have significantly different properties and they should be modeled separately: We find that the degree distribution of the provider-customer only edges can be accurately described by a power-law (with correlation coefficient higher than99%) in all the topological instances that we examine. In contrast, degree distribution of the peer-to-peer only edges is better described by a Weibull distribution with correlation coefficient higher than 99%, which corroborates previous studies [9], [7]. Third, the degrees of the nodes of a peer-to-peer link can vary significantly: 50% of the peer-to-peer edges are between nodes whose degrees, d1, d2 differ a lot either in absolute (|d1- d2>144|) or relative value (max (d1/d2, d2/d1) >4). This has direct implications on how we think about and model peer-to-peer edges. For instance, this observation suggests that researchers need to use caution when using the degree as an indication of whether two ASes could have a peer-to-peer relationship. Our results can provide guidelines to AS policy inference algorithms, which partly rely on the node degree. Fourth, we provide an educated guess on how many edges we may still be missing. We estimate the edges to be roughly 35% compared to the peer-to-peer edges we know at the end of this study.

The rest of this paper is organized as follows. We review the data sources and previous work in Section II. In Section III, we present our framework and the motivation behind its design. In section IV, we discuss the impact of the new edges. In Section IV, we summarize our work.

II. BACKGROUND

A. Data Sources and Their Limitations

In this section, we describe the most popular data sources and their two main limitations: incompleteness and a bias in the nature of the discovered links.

BGP routing table dumps are probably the most widely used resource that provides information on the AS Internet topology. Each table entry contains an AS path, which corresponds to a set of AS edges. Several sites collect tables from multiple BGP routers, such as Routeview [13] and RIPE/RIS [14]. An advantage of the BGP routing tables is that their link information is considered reliable. If an AS link appears in a BGP routing table dump, it is almost certain that the link exists. However, limited number of vantage points makes it hard to discover a more complete view of the AS-level topology. A single BGP routing table has the union of “shortest” or, more accurately, preferred paths with respect to this point of observation.Asa result, such a collection will not see edges that are not on any preferred path for this point of observation. Several theoretical and experimental efforts explore the limitations of such measurements [15], [16].Worse, such incompleteness may be statistically biased based on the type3 of the links. Some types of AS links are more likely to be missing from BGP routing table dumps than other types. Specifically, peer-to-peer links are likely to be missing due to the selective exporting rules of BGP. Typically, a peer-to-peer link can only be seen in a BGP routing table of these two peering ASes or their customers. A recent work [9] discusses in depth this limitation.

BGP updates are used in previous studies [3], [5] as a source of topological information and they showthat by collecting BGP updates over a period of time, more AS links are visible. This is because as the topology changes, BGP updates provide transient and ephemeral route information. However, if the window of observation is long, an advertised link may cease to exist [3] by the time that we construct a topology snapshot. In other words, BGP updates may provide a superimposition of a number of different snapshots that existed at some point in time. Recently, Oliveira et al. [17] explicitly distinguished this commonly overlooked “liveness problem” from the “completeness problem", which is the central topic of this paper. Note that BGP updates are collected at the same vantage points as the BGP tables in most collection sites. Naturally, topologies derived from BGP updates share the same statistical bias per link type as from BGP routing tables: peer-to-peer links are only to be advertised to the peering ASes and their customers. This further limits the additional information that BGP updates can provide currently. On the other hand, BGP updates could be useful in revealing ephemeral backup links over long period of observation, along with erroneous BGP updates.

By using traceroute, one can explore IP paths and then translate the IP addresses to AS numbers, thus obtaining AS paths. Similar to BGP tables, the traceroute path information is considered reliable, since it represents the path that the packets actually traverse. On the other hand, a traceroute server explores the routing paths from its location towards the rest of the world, and thus, the collected data has the same limitations as BGP data in terms of completeness and link bias. One additional challenge with the traceroute data is the mapping of an IP path to an AS path. The problem is far from trivial, and it has been the focus of several recent efforts [18]–[20].

Internet Routing Registry (IRR) [10] is the union of a growing number of world-wide routing policy databases that use the Routing Policy Specification Language (RPSL). In principle, each AS should register routes to all its neighbors (that reflect the AS links between the AS and its neighbors) with this registry. IRR information is manually maintained and there is no stringent requirement for updating it. Therefore, without any processing, AS links derived from IRR are prone to human errors, could be outdated or incomplete. However, the up-to-date IRR entries provide a wealth of information that could not be obtained from any other source. A recent effort [12] shows that, with careful processing of the data, one can extract a nontrivial amount of correct and useful information.

B. Related Work and Comparison

There has been a large number of measurements studies related to topology discovery, with different goals, at different times, and using different sources of information.

Our work has the following characteristics that distinguish it from most previous other efforts, such as [9], [2]: 1) We make extensive use of topological information from the Internet Exchange Points to identify more edges. It turns out that IXPs "conceal" many links which did not appear in most previous topology studies. 2) We use a more sophisticated, comprehensive and thorough tool [12] to filter the less accurate IRR data, which was not used by previous studies. 3) We employ a "guess-and-verify" approach for finding more edges by identifying potential edges and validating them through targeted traceroutes. This greatly reduced the number of traceroutes that were needed. 4) We accept new edges conservatively and only when they are confirmed by a BGP table or a traceroute. In contrast, some of the previous studies included edges from IRR without confirming them with traceroute.

The most relevant previous work is done by Chang et al. [2] with data collected in 2001. They identify new edges by looking at several sources of topological information including BGP tables and IRR. They estimate that 25%-50% AS links were missing from Oregon Routeview BGP table, the most commonly used data set for AS topology studies. Their work was an excellent first step towards a more complete topology.

In a parallel effort, Cohen and Raz [9] identify missing links in the Internet topology. Our studies corroborate some of the observations there. Note that, their work does not include an exhaustive measurement, data collection and comparison effort as our work. For example, IXP information was not used in their work.

Several other interesting measurement studies exist. Net-Dimes [4] is an effort to collect large volumes of host-based traceroute information. The key here is to increase the number of traceroute points by turning cooperative end hosts into observation points. The challenge now becomes the measurement noise removal, the collection, and processing of the information [21]. Our approach and NetDimes could complement and leverage each other towards a more complete and accurate topology. Donnet et al. [22] propose efficient algorithms for large-scale topology discovery by traceroute probes. Rocketfuel [23] explores ISP topologies using traceroutes. In [5], the authors examine the information contained in BGP updates.

Most of these studies and our work seek a complete snapshot of the Internet topology. In otherwords, short-lived backup links are most likely not included in most such studies. Some ASes have such links, which normally are not "visible" unless the primary links are down. Recently, active BGP probing [8] has been proposed as a method for identifying backup AS links, and this could complement our work and the efforts mentioned above.

There are several efforts that study the topology and they would benefit from an accurate and complete topology. A plethora of efforts attempts to model the topology and to generate realistic topologies (e.g., [24]). Some studies [15], [25] document the limitations of the sources of topological information, but without necessarily attempting to identify a more complete topology. A recent study [7] models the evolution of the Internet topology by investigating the process of AS peerings. Another recent work [17] models the evolution using a constant rate birth–death process. Our work can be seen as a basis that can provide more complete and accurate information for such studies.

The exhaustive identification of IXP participants has received limited attention. Most previous work focuses on identifying the existence of IXPs. Xu et al. [11] develop what appears to be the first systematic method for identifying IXP participants. Inspired by their work, our approach subsumes their method, and thus, it provides more complete and accurate results.

III. FRAMEWORK FOR FINDING MISSING LINKS

In this section, we present a systematic framework for extracting and synthesizing the AS level topology information from different sources. The different sources have complementary information of variable accuracy. Thus, we cannot just simply take the union of all the edges. A careful synthesis and cross-validation is required. At the same time, we are interested in identifying the properties of the missing AS links.

In a nutshell, our study arrives at three major observations regarding the properties of the missing AS links: 1) most of the missing AS edges are of the peer-to-peer type; 2) most of the missing AS edges from BGP tables appear in IRR; and 3) most new found AS edges are incident at IXPs. At different stages of the research, these three observations direct us to discover even more edges, some of which do not appear in any other source of information currently.

We present an overview of our work in order to provide the motivation for the different steps that we take. We start with the data set from Oregon routeviews BGP table Dump (OBD) [13], the BGP table dumps collected at route-views.oregon-ix. net, which is by far the most widely used data archive. Our work consists of four main steps.

A. BGP routing tables: We consider the AS edges derived from multiple BGP routing table dumps [3], and compare them to the Routeview data (OBD). The question we try to answer is what the information that the new BGP tables bring is. We use the term BD to refer to the union data from all available BGP table Dumps. Table I lists the acronyms for our data sets.

B. IRR data: We systematically analyze the IRR [10] data and identify topological information that seems trustworthy by

TABLE I
TOPOLOGICAL DATA SETS USED IN OUR STUDY

OBD	The Oregon routeviews BGP table dump
BD	OBD and other additional BGP table dumps
IRRnc	IRR edges processed by Nemecis with non-conflicting policy declarations
IRRdual	IRRnc edges correctly declared by both adjacent ASes
BD+IRR	BD and the edges of IRRdual confirmed by RETRO
IXPall	Union of cliques of IXP participants
ALL	BD+IRR and the potential IXP edges that are confirmed by RETRO

TABLE II
STATISTICS OF THE TOPOLOGIES

Name	Nodes	Edges	p-c	p-p
OBD	19.8k	42.6k	36.7k	5.5k
BD	19.9k	51.3k	38.2k	12.7k
BD+IRR	19.9k	56.9k	38.2k	18.3k
ALL	19.9k	59.5k	38.2k	20.9k

Nemecis [12]. We follow a conservative approach, given that IRR may contain some outdated and/or erroneous information. We do not accept new edges from IRR, even after our first processing, unless they are confirmed by traceroutes (using our RETRO tool). Overall, we find that IRR is a good source of hints for missing links. For example, we discover that more than 80% of the new edges found in the new tables (i.e., the AS edges in BD but not in OBD) already exist in IRR. Even compared to BD, IRR has significantly more edges, which are validated by RETRO as we explain below.

C. IXPs and potential edges: We identify a set of potential IXP edges by applying our methodology on inferring IXP participants from Section V. We find that many of the peer-to-peer edges missing from the different data sets could be IXP edges.

D. Validation using RETRO: We use our traceroute tool, RETRO, to verify potential edges from IRR and IXPs. First, we confirm the existence of many potential edges we identified in the previous steps. We find that more than 94% of the RETRO-verified AS edges in IRR indeed go through IXPs. We also discover edges that were not previously seen in either the BGP table dumps or IRR. In total, we have validated 300% more peer-to-peer links than those in the OBD data set. The statistics of the topologies generated from the different data sets in our study are listed in Table II.

IV. THE IMPACT OF THE NEW EDGES

We follow a commercial characterization: an AS pays for sending traffic through its provider. How much do the new edges affect these decisions? We count for each AS how many of its paths stop going through one of its providers once the new edges are added. We refer to these paths as ex-provider paths. We plot the number of the ex-provider paths for each node in Fig. 1. We see that the effect on the routing on individual ASes is dramatic: there are many ASes, for each of which, several thousand out of the total 20K paths (to all other ASes) stop going through a provider. Another interesting observation is that the nodes which seem to bene_t the most from these changes have degrees in the range from 10 to 100 (right y-axis).

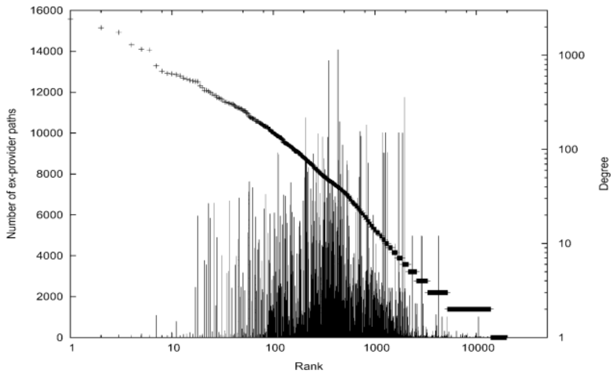

Fig. 1. The number of ex-provider paths (shown as impulses on the left y-axis) of each node in order decreasing node degree (shown as a semi diagonal line corresponding to the right y-axis). The x-axis shows the rank of the nodes in the order of descending degree.

V. CONCLUSION

Accuracy of simulations on Internet protocols and applications is strongly dependent on the correct use of synthetic topologies at large scales. The Internet should not simply be seen as a power law network of transit and stub ASes. The core is a dense mesh of tier-1 ISPs interconnected by peering relationships. At the edge on the other hand, costumer ISPs aim at increasing their peering relationships with other established ASes, through Internet eXchange Points (IXP). Content distribution networks are also increasingly adding peering links with as many as they can afford. The availability of rich topology information from numerous vantage points (at least at the AS-level) calls for an end to blind use of random networks or power law models for simulations. In a nutshell, our work develops a systematic framework for the cross-validation and the synthesis of most available sources of topological information. We are able to find and confirm approximately 300% additional edges. Furthermore, we recognize that Internet Exchange Points (IXPs) hide significant topology information and most of those new discovered peer-to-peer AS links are incident at IXPs. The reason for such a phenomenon is probably because, most missing peer-to-peer links are likely to be at the middle or lower level of the Internet hierarchy, and peering at some IXP is a cost-efficient way for the ASes to setup peering relationships with other ASes. We show that by adding these new AS links, some research results based on previous incomplete topology, such as routing decision and ISP profit/cost, change dramatically. Our study suggest that business-oriented studies of the Internet should make a point of taking into consideration as many peer-to-peer edges as possible. So, how many AS links are still missing from our new snapshot of the Internet topology? Our findings suggest that if we know the peering matrix of all the IXPs, we might be able to discover most of the missing peer-to-peer AS links. Unfortunately, very few IXPs publish their peering matrices. Futhermore, the published peering matrices are not necessarily accurate, complete or up-to-date. In our conservative estimates, there might be still 35% hiding peer-to-peer edges, in addition to what we already have in current Internet AS graph. Our future plans have two distinct directions. First, we want to continue the effort towards a more complete Internet topology instance. Using the framework we developed here, we are in a good position to quickly and accurately incorporate new information, such as new BGP routing tables, or new traceroute servers. Second, given our more complete AS topology, we are in a better position to understand the structure of the Internet and the socioeconomic and operational factors that guide its growth. This in turn could help us interpret and anticipate the Internet evolution and, indirectly, give us guidelines for designing better networks in the future.

REFERENCES

[1] S. Floyd and V. Paxson, "Difficulties in simulating the Internet,"
IEEE/ACM Trans. Networking, vol. 9, no. 4, pp. 392–403, Aug. 2001.
[2] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger,
"Towards capturing representative AS-level Internet topologies,"
Computer Networks, vol. 44, no. 6, pp. 737–755, 2004.
[3] B. Zhang, R. A. Liu, D. Massey, and L. Zhang, "Collecting the Internet
as-level topology," ACM SIGCOMM Comput. Commun. Rev. (CCR),
vol. 35, no. 1, pp. 53–61, Jan. 2005.
[4] Y. Shavitt and E. Shir, "DIMES: let the Internet measure itself," ACM
SIGCOMM Comput. Commun. Rev. (CCR), vol. 35, no. 5, pp. 71–74,
Oct. 2005.
[5] X. Dimitropoulos, D. Krioukov, and G. Riley, "Revisiting Internet
AS-level topology discovery," in Proc. Passive and Active Measurement
(PAM) Workshop, Boston, MA, 2005.
[6] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos,
K. claffy, and A. Vahdat, "The Internet AS-level topology:

Three data sources and one definitive metric," ACM SIGCOMM
Comput. Commun. Rev. (CCR), vol. 36, no. 1, pp. 17–26, Jan. 2006.
[7] H. Chang, S. Jamin, and W. Willinger, "To peer or not to peer: Modeling
the evolution of the Internet's AS-level topology," in Proc. IEEE
INFOCOM, 2006, 12 pp.
[8] L. Colitti, G. DiBattista, M. Patrignani, M. Pizzonia, and M. Rimondini,
"Investigating prefix propagation through active BGP probing,"
Microprocessors Microsyst., vol. 31, no. 7, pp. 460–474, 2007.

[9] R. Cohen and D. Raz, "The Internet dark matter—on the missing links
in the AS connectivity map," in Proc. IEEE INFOCOM, 2006, 12 pp.
[10] Internet Routing Registry. [Online]. Available: http://www.irr.net
[11] K. Xu, Z. Duan, Z. Zhang, and J. Chandrashekar, "On properties of
Internet exchange points and their impact on AS tolology and relationship,"
in Networking, 2004, pp. 284–295.
[12] G. Siganos and M. Faloutsos, "Analyzing BGP policies: Methodology
and tool," in Proc. IEEE INFOCOM, 2004, vol. 3, pp. 1640–1651.
[13] Oregon Routeview Project. [Online]. Available: http://www.routeviews.
org
[14] Ripe Route Information Service. [Online]. Available: http://www.ripe.
net/ris
[15] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie, "Sampling biases in
IP topology measurements," in Proc. IEEE INFOCOM, 2003, vol. 1,
pp. 332–341.
[16] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore, "On the bias
of traceroute sampling, or power-law degree distributions in regular
graphs," in Proc. STOC'05, Baltimore, MD, May 2005, pp. 694–703.
[17] R. V. Oliveira, B. Zhang, and L. Zhang, "Observing the evolution of
Internet AS topology," in Proc. ACM SIGCOMM, 2007, pp. 313–324.
[18] Z. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an accurate
AS-level traceroute tool," in Proc. ACM SIGCOMM, 2003, pp. 365–378.
[19] Z. M. Mao, D. Johnson, J. Rexford, J.Wang, and R. Katz, "Scalable and
accurate identification of AS-level forwarding paths," in Proc. IEEE
INFOCOM, 2004, vol. 3, pp. 1605–1615.
[20] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M.
Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies
with Paris traceroute," in Proc. ACM IMC'06, Rio de Janeiro, Brazil,
Oct. 2006, pp. 153–158.
[21] E. Shir, Dec. 2005, Personal communication via e-mail.
[22] B. Donnet, P. Raoult, T. Friedman, and M. Crovella, "Efficient algorithms
for large-scale topology discovery," in Proc. ACM SIGMETRICS,
Jun. 2005, pp. 327–338.
[23] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring
ISP topologies with Rocketfuel," IEEE/ACM Trans. Netw., vol. 12, no.
1, pp. 2–16, Feb. 2004.
[24] S. Jaiswal, A. Rosenberg, and D. Towsley, "Comparing the structure of
power law graphs and the Internet AS graph," in Proc. 12th IEEE Int.
Conf. Network Protocols (ICNP'04), 2004, pp. 294–303.
[25] S. Jin and A. Bestavros, "An empirical study of inherent routing bias
in variable-degree networks," Boston Univ., Boston, MA, Tech. Rep.,
2003.