

Towards Capturing Representative AS-Level Internet Topologies

[Extended Abstract]*

Hyunseok Chang[†] Ramesh Govindan^{††} Sugih Jamin[†] Scott J. Shenker^{††} Walter Willinger[§]

[†]University of Michigan

^{††}ICSI

[§]AT&T Labs-Research

{hschang,jamin}@eecs.umich.edu

{ramesh,shenker}@icsi.berkeley.edu

walter@research.att.com

1. INTRODUCTION

For the past two years, there has been a significant increase in research activities related to studying and modeling the Internet's topology, especially at the level of *autonomous systems* (ASs). A closer look at the measurements that form the basis for all these studies reveals that the data sets used consist of the BGP routing tables collected by the Oregon route server (henceforth, *the Oregon route-views*) [1]. So far, there has been anecdotal evidence and an intuitive understanding among researchers in the field that BGP-derived AS connectivity is not complete. However, as far as we know, there has been no systematic study on *quantifying* the completeness of currently known AS-level Internet topologies. Our main objective in this paper is to quantify the completeness of Internet AS maps constructed from the Oregon route-views and to attempt to capture *more representative* AS-level Internet topology. One of the main contributions of this paper is in developing a methodology that enables quantitative investigations into issues related to the (in)completeness of BGP-derived AS maps.

2. ON THE COMPLETENESS OF BGP-DERIVED AS-LEVEL TOPOLOGY

While we do not have a complete map of the Internet at the AS level, we can discover the AS neighbors of each AS if we have access to its BGP routing table. We call the neighbors of an AS_{*i*} so discovered the *local* view of AS_{*i*}. From AS_{*j*}'s

*A full version of this paper is available as [2].

[†]This project is funded in part by NSF grant number ANI-0082287 and by ONR grant number N000140110617. Sugih Jamin is further supported by the NSF CAREER Award ANI-9734145, the Presidential Early Career Award for Scientists and Engineers (PECASE) 1998, and the Alfred P. Sloan Foundation Research Fellowship 2001. Additional funding is provided by AT&T Research, and by equipment grants from Sun Microsystems Inc. and Compaq Corp.

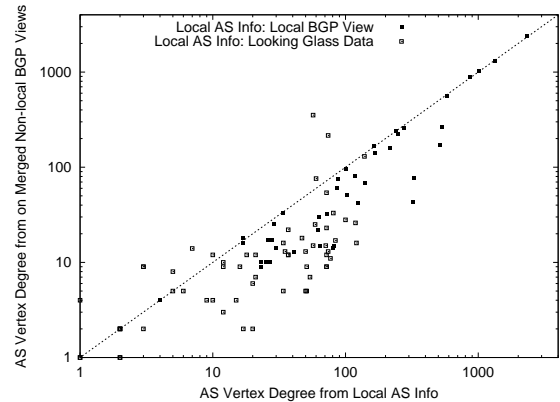


Figure 1: The Completeness of Non-local Views

BGP routing table, we can also infer some portion of AS_{*i*}'s neighbors. We call the neighbors of AS_{*i*} so inferred by AS_{*j*} the *non-local* view of AS_{*i*}. Next we ask, if we try to infer the neighbors of a given AS (say, AS_{*i*}) from the BGP routing tables of *other* ASs, i.e., by merging the non-local views of AS_{*i*}, how complete would the inferred neighbors of AS_{*i*} be compared to its own local view? To answer this question, we collected (1) BGP routing tables of 41 distinct ASs from 12 existing public route servers, and (2) BGP summary data of about 70 distinct ASs from Looking Glass sites (see [2] for further details). In Fig. 1, we compare the vertex degree of a given AS inferred by its merged non-local views (*y*-axis) and its actual vertex degree (*x*-axis). The actual vertex degree of a given AS is based on either its local BGP view or its Looking Glass data. Aside from a few exceptional cases, it is clear that the vertex degrees predicted by non-local views fall short of actual vertex degrees.

3. AUGMENTING CONNECTIVITY USING THE INTERNET ROUTING REGISTRY

To further quantify the difference between BGP-derived AS connectivities and actual inter-AS connections, we next consulted the IRR databases, which are public repositories that maintain individual ISP's routing policy information. The IRR database of our choice is the RIPE database, which is believed to contain reasonably up-to-date entries. After carefully eliminating invalid records from this RIPE database,

Source	# of nodes (%inc)	# of edges (%inc)
Oregon route-views	11,174	23,409
+ RSs	11,268 (0.84%)	26,324 (12.5%)
+ RSs + LG	11,320 (1.3%)	27,899 (19.2%)
+ RSs + LG + RIPE	11,456 (2.5%)	32,759 (40.0%)

Table 1: AS Graph Statistics

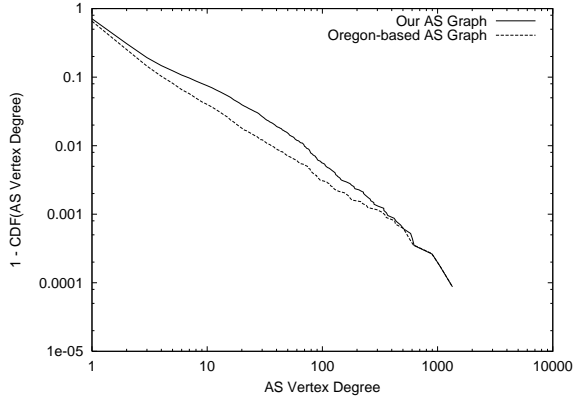


Figure 2: Frequency Distribution of AS Degree.

we extract AS connectivity information from the remaining records. We find that AS graphs reconstructed from the Oregon route-views, the Looking Glass data, as well as RIPE information have typically about 40% more edges (and about 2% more nodes) than their counterparts that rely solely on Oregon route-views (see Table 1).

Next, we check how increasingly denser AS graphs affect the power-law characteristics that have been identified by Faloutsos et al. in the Oregon-based AS graphs [3]. We compare the Oregon-based AS graph (Table 1, first row) against our more complete AS graph (Table 1, last row). In Figure 2 we plot the complementary cumulative distribution function of the AS degree. The vertex degree distribution of the Oregon-based AS graph appears to be consistent with the strict power-law result reported in [3]. However, our AS graph, though not necessarily complete, shows more ASs with vertex degrees ranging from 4 to 300, resulting in a curved line that deviates significantly from the Oregon-based counterpart.

4. ON VALIDATING AND GENERALIZING OUR AS GRAPH

In this section, we address two questions regarding our much denser AS graph. First, could our AS graph be *artificially* inflated by our use of non BGP-derived connectivity information such as IRR-based data? Second, given that we use only the European RIPE registry database, can our observations on the renewed AS vertex-degree distribution be expected to hold for more *global* version of our AS graph?

Noticing the rich AS connectivity of existing public exchange points (EPs), we validate the newly-found edges by identifying those ASs that are *physically co-located* at EPs. We retrieved AS co-location information of 16 European EPs from Looking Glass sites. We find that as much as 50% of the newly-found edges in our AS graph occur at one of

these 16 EPs. We also find that a majority of large maximal cliques newly formed in our AS graph are induced by all physically co-located ASs. These results strongly suggest that the added connectivity of our AS graph reflects the actual existing AS-level connectivity.

The generalization of our findings to the global AS-level topology is based on the following observations: (1) our AS graph (as well as the existing BGP-derived AS graph) captures the connectivity of both tier-1 ASs and end-customer ASs reasonably well, (2) there are more than 150 EPs *world-wide*, and we typically see some 10–100 ASs co-located at each of them. Referring back to Figure 2, the implication of these observations is as follows. Adding more AS connectivity data of the global Internet, we would see the head and the tail of degree frequency distribution curve remain essentially unchanged, but the middle portion of the curve would be expected to experience a further shift upwards and to the right. The resulting distribution will still remain *heavy-tailed* or *highly-variable*, but will very likely no longer conform to the strict power-law behavior that characterizes the Oregon-based AS maps.

5. CONCLUSION

In this paper, we show that, as far as past investigations into the Internet topology at the AS level are concerned, many of the findings that have been reported in the literature have used the publicly available BGP measurements without realizing the possible pitfalls associated with taking the data at face value, or without examining whether or not the use of the data is justified for inferring the Internet AS connectivity. Our results confirm that while the actual connectivity of the Internet at the AS level is quite high, BGP measurements typically see only a portion of all existing AS connections. This observation comes as no surprise as BGP is a path-vector protocol, and the propagation of AS connectivity is determined by routing policy. In this sense, the main lesson learned from the study presented in this paper is that since network-related measurements often reflect network protocol-specific features, arguing for the general validity of an empirical finding about the Internet should typically include a careful investigation into the sensitivity of the findings to known deficiencies and inaccuracies of the measurements at hand.

Acknowledgements

We thank Tim Griffin for valuable discussions and for offering us the UUNET’s BGP routing table.

6. REFERENCES

- [1] “University of Oregon Route Views Project,” <http://www.antc.uoregon.edu/route-views/>.
- [2] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, “Towards Capturing Representative AS-Level Internet Topologies,” Tech. Rep. CSE-TR-454-02, EECS Department, University of Michigan, 2002.
- [3] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the Internet topology,” in *Proceedings of ACM SIGCOMM ’99*, August 1999.