Timelapse of AS-Level Topology Graphs Using BGP Advertisements

Sourabh Gawande*, TSRK Prasad†, Amrita Suresh‡, Neena Goveas§ and Bharat Deshpande¶
Department of Computer Science & Information Systems
BITS Pilani - K K Birla Goa Campus, Goa, India - 403726
{*f2011013, †tsrkp, ‡f2012054, §neena, ¶bmd} @goa.bits-pilani.ac.in

***Abstract*—Internet is a rich fabric of interconnections between autonomous systems (ASes). These interconnection / peering agreements are subjected to constant change and are governed by mostly nontechnical considerations such as geopolitics, law enforcement, business and personal relationships. The impact of changes in peering agreements are felt in network performance and routing stability of Internet.**

Here we represent the interconnections between autonomous systems (ASes) using AS-level topology graphs. We plot the temporal variations of AS-level topology graphs to illustrate dynamism inherent in Internet topology. We use IPv4 BGP data for the year 2013 from Route Views project to represent temporal variations using a timelapse of all AS-level topology graphs. These temporal variations can be used to understand the underlying changes in interconnection fabric of the Internet. We generate AS-level geolocation databases to solve the graph layout problems. The AS-level geolocation databases are also helpful in geography specific analysis of AS topology. Visualizations produced as part of this work are very useful in debugging network performance.

I. Introduction

Internet is defined as the interconnection of networks. The interconnected networks are mostly independent networks operated by competing organizations. These independent networks are referred to as autonomous systems (ASes). Internet Service Providers (ISPs) are the most prominent ASes connected to the Internet. With the growth of content networking, even the content hosting and delivery organizations have become full fledged ASes. These ASes use Border Gateway Protocol (BGP) for exchanging routing and reachability information.

BGP is a well-established path-vector protocol; When an AS wants to advertise its' network prefixes, AS sends BGP advertisements to its' neighbouring ASes. A BGP advertisement comprises of a number of attributes, such as ASPATH, AS_SET and PREFIX. A sample BGP path advertisement is shown in Figure 1. The path advertisement shown in Figure 1 is valid only for routers that maintain BGP peering session with router using IP interface address of 208.51.134.246. In Figure 1, PREFIX specifies the destination network, 1.0.6.0/24; ASPATH lists the ordered sequence of ASes that must be traversed in order to reach the specified destination network. With respect to Figure 1, an IP packet destined for 1.0.6.0/24 network shall traverse 3549 6939 6939 4826 38803 56203 ASes. Thus ASPATH attribute implicitly lists AS connectivity information between adjacent ASes mentioned in ASPATH. We use ASPATH attribute to generate the AS-level topology graph.

```
TIME: 12/12/13 12:00:00
TYPE: TABLE_DUMP_V2/IPV4_UNICAST
PREFIX: 1.0.6.0/24
SEQUENCE: 3
FROM: 208.51.134.246 AS3549
ORIGINATED: 12/05/13 06:31:50
ORIGIN: IGP
ASPATH: 3549 6939 6939 4826 38803 56203
NEXT_HOP: 208.51.134.246
MULTI_EXIT_DISC: 13899
```

Fig. 1: A typical BGP path advertisement; ASPATH attribute is highlighted. The sequence of numbers shown are AS numbers.

The entire Internet can be viewed as an AS-level topology graph where each AS is a node, and the BGP peering between two ASes is a link. This AS-level topology has important implications for both day to day Internet operations and Internet research. Estimates of the AS-level topology have been used in a variety of research activities, including analyzing Internet topological properties, network debugging [1], inferring AS relationship hierarchy [2], building network topology generators for simulations, and evaluating the effectiveness of new protocols and improvements [3].

Center for Applied Internet Data Analysis (CAIDA) has been a pioneer in the generation of Internet topologies. CAIDA runs Ark, short for Archipelago, measurement infrastructure for performing distributed traceroutes. These traceroutes are supplemented with BGP tables from Route Views project to generate macroscopic network topology graphs. CAIDA's efforts in internet mapping and network measurements are summarized in [4].

Visualization of AS-level topology poses its' own challenges and few attempts have been made to come up with meaningful visualizations. RIPE NCC provides BGPlay [5], [6] which visualizes a network prefix propagation across ASes. Cyclops [1] is another attempt at visualizing changes in the interconnections of an AS. Few projects like NetViews, Internet Atlas, and VisTracer also attempt to integrate geolocation information into the visualization.

The Route Views Project of University of Oregon [7] collects real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. The fourteen Route Servers of Route Views Project peer with various ISPs around the world. Route Views project collects the Routing Information Bases (RIBs) every 2 hours and archive them.

There were some efforts to use route views data for BGP topology inference [8]. A comprehensive survey of network topology inference methods is provided in [9]. The trade offs involved in performing traceroute-based topology inference is contrasted with BGP routing information based topology inference are clearly explained in [10]. The AS-level topology graphs generated by using distributed network measurement platforms such DIMES, Ark, Ripe Atlas and Route Views are compared in [9].

Here we use one collection session of RIBs to generate one snapshot of AS-level topology graph. This graph is a visualization of Internet topology for the chosen BGP data sample. We create many such snapshots over a period of time and combine them to generate a timelapse. Individual snapshots of BGP data as well as timelapses (using these snapshots) can be created using the procedure detailed in section II. We use the generated AS-level graphs of IPv4 BGP data to visualize the geographic distribution of ASes and the relationship between these ASes. A timelapse of AS-level topology graphs is a good way to visualize and understand the dynamic nature of the Internet topology.

II. GRAPH GENERATION

A BGP path advertisement contains the AS path for a network prefix. One RIB snapshot of the Route Views Project contains millions of path advertisements. We generate the AS-level topology graph by the following BGP path transformations.

1) Each AS number is treated as a node.
2) Any two adjacent nodes are connected by an undirected edge.

A few problems arise because of route leakages and path padding. Route leakages occur when private AS numbers (range from 64512 to 65535; these are not to be seen on Internet) are seen in some path advertisements. These advertisements are invalid as per the BGP standard. We ignore nodes coming out of private AS numbers and edges leading to these eliminated nodes. Any autonomous system number (ASN) repetition due to path padding is ignored for graph generation purpose.

Another issue is with AS_SET BGP attribute. Recent network security guidelines advise against the use of AS_SET attribute in the interest of secure inter-domain routing. But AS_SET aggregation is still used by many AS operators in BGP path advertisements. If advertised AS_PATH is 300 $\{200\ 100\}$, then undirected edges (300,200) and (300,100) are added to the AS-level topology graph.

A. Graph Layout

Since the graph contains thousands of nodes and millions of edges, graph layout becomes a non-trivial issue. We follow the layout conventions set by CAIDA Internet engineers. According to AS-topology layout algorithm of CAIDA [11], the nodes are all arranged on concentric circles. The exact position of a node is described by the polar coordinates of the node (r, θ). The radius r decreases as the outdegree of the nodes increases. The angle θ is the longitude from the geocoding or geolocation information, which is represented on the circle as geo-specific sectors.

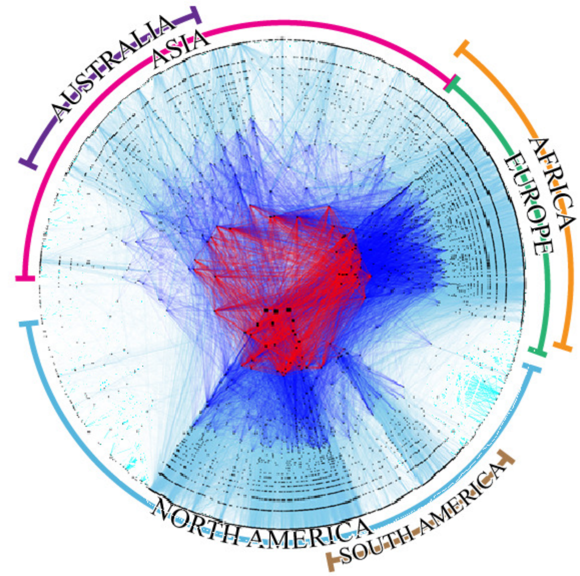

Fig. 2: AS-level topology generated using Route Views IPv4 BGP data snapshot taken at 01-January-2013 2:00AM.

The placement of the nodes in the graph is done based on the global degree of the AS representing the node and the longitude of the AS.

1) The radius of an AS is determined by

$$r = 1 - log\left(\frac{outdegree + 1}{max\_outdegree + 1}\right)$$

2) The angle is obtained in one of two ways. If WHOIS databases of Regional Internet Registry (RIRs) contain the physical address of AS, we map the address to the corresponding longitude coordinates. This is the most preferred approach. Else we obtain the longitude coordinates of IP prefixes owned by the AS using Maxmind geolocation database and use their mean as representative location for AS.

A sample AS-topology graph is shown in Figure 2. The size of a node is a function of its outdegree. The thickness of an edge is a function of degree of two nodes connected by that edge. At the core, nodes have very high degree and have large number of interconnections amongst themselves. This is the typical interconnection pattern amongst largest ISPs of the world; they are very densely connected. The small, access ISPs of the world have very low degree and are represented on the periphery. The small ISPs are connected in a tree-like topology with the large ISPs. Edges connecting to the access ISP nodes are thin as an indication of their relatively low node degree.

Individual snapshots of BGP data as well as a timelapse (using these snapshots) can be created using the procedure described above. After a single snapshot has been generated,

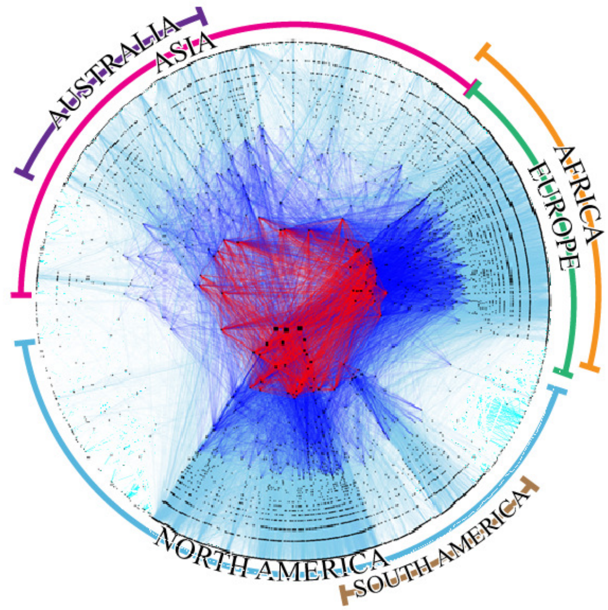

(a) January, 2013

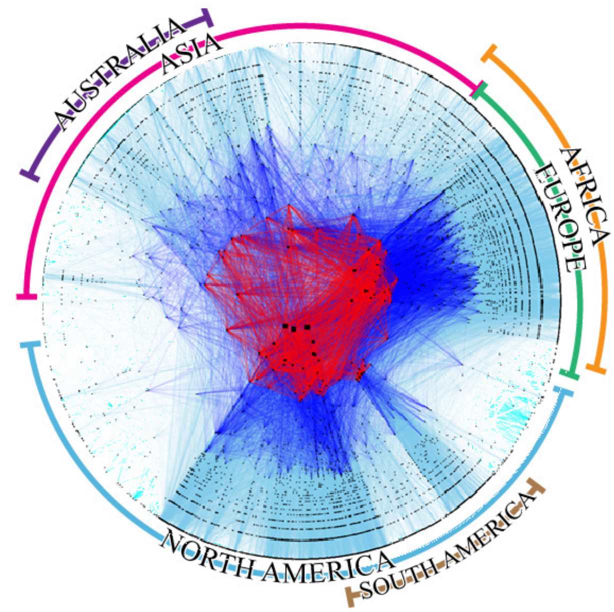

(b) December, 2013

Fig. 3: AS-level topology graph generated using Route Views IPv4 BGP data dump taken at the beginning of two months in 2013. The connectivity among core ASes increased in December, 2013 when compared with January, 2013. This change is reflected as more dense graph in core of December, 2013 graph.

the same steps can be used to get snapshots for BGP data collected at different times. Putting these snapshots together gives a timelapse.

III. CONCLUSIONS

We wanted to generate CAIDA IPv4 AS core like plot and succeeded in that effort. CAIDA uses Route Views BGP data, Netacuity database and traceroute information from Ark project to generate their IPv4 AS core plot. We were able to generate an equivalent plot using only Route Views BGP data and freely available Maxmind geoip database.

We are able to generate a timelapse for the year 2013 and the time lapse shows quite a bit of dynamism. Because of the work done as part of this project, we are able to build a database of ASN to geolocation mapping. This database can be used by other research projects as well. We made the AS geolocation database, AS-topology layouts and timelapse for the year 2013 available at: *https://github.com/prasadtalasila/ASTopology*

Figure 3 shows a complete view of AS-level topology graphs for two months, namely January and December of the year 2013. Densification of core is evident in Figure 3b when compared with Figure 3a. Figure 4 highlights a specific portion of AS core. We can clearly see densification of edges in the core. Moreover the circled region in Figure 4a highlights a set of ASes from January, 2013 graph. A corresponding region from December, 2013 graph is shown in Figure 4b.

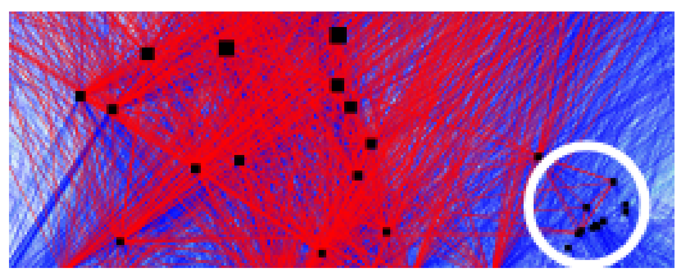

(a) January, 2013

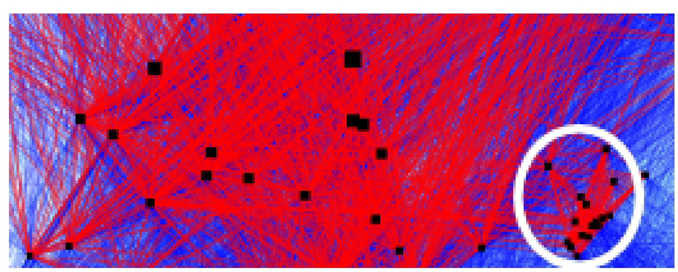

(b) December, 2013

Fig. 4: A slice of AS-level topology graphs as shown in Figure 3. ASes in the encircled region show better connectivity in December, 2013 when compared with January, 2013.

IV. FUTURE WORK

The process of creating a single AS-level topology graph takes significant amount of time, approximately 15 minutes. Such a large amount of time is required because of initial BGP data set size and the resulting graph size. Further work can be done to reduce the graph generation time. Improvements in this processing time will make real-time interaction with Route Views data set a distinct possibility. The current implementation only uses the RIBs that are collected by Route Views route collectors every 2 hrs. Apart from this, Route Views router collectors also archive updates done to the RIBs of their peers (i.e. addition and withdrawal of routes) every 15 min. These updates can be used to create timelapses of finer granularity.

We plan to integrate the AS-level topology graph on a geographic map to show geospecific variations on the Internet topology more clearly. We can also show AS-level path trace from a source IP address to a destination IP address. Such a path trace will not require sending out ICMP PING messages at all. Suggested improvements will bring in necessary transparency to the process of managing interconnections among autonomous systems.

REFERENCES

[1] Y.-J. Chi, R. Oliveira, and L. Zhang, "Cyclops: The AS-level connectivity observatory," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 5, pp. 5–16, Sep. 2008.

[2] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001. [Online]. Available: http://dx.doi.org/10.1109/90.974527

[3] K. Park and H. Lee, "On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets," *SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 4, pp. 15–26, Aug. 2001.

[4] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: From art to science," in *Conference For Homeland Security, 2009. CATCH '09. Cybersecurity Applications Technology*, March 2009, pp. 205–211.

[5] G. Di Battista, F. Mariani, M. Patrignani, and M. Pizzonia, "BGPlay: A system for visualizing the interdomain routing evolution," in *Graph Drawing*, ser. Lecture Notes in Computer Science, G. Liotta, Ed. Springer Berlin Heidelberg, 2004, vol. 2912, pp. 295–306.

[6] P. Angelini, L. A. Clarucci, M. Candela, M. Patrignani, M. Rimondini, and R. Sepe, "BGPlay3D: Exploiting the ribbon representation to show the evolution of interdomain routing," in *Graph Drawing*. Springer, 2013, p. 526.

[7] "University of Oregon Route Views Project," 2014. [Online]. Available: http://www.routeviews.org/ [Last accessed: 03.11.2014]

[8] G. Siganos, S. L. Tauro, and M. Faloutsos, "Jellyfish: A conceptual model for the AS Internet topology," *Communications and Networks, Journal of*, vol. 8, no. 3, pp. 339–350, Sept 2006.

[9] B. Huffaker, M. Fomenkov, and K. Claffy, "Internet topology data comparison," 2012.

[10] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: inference, modeling, and generation," *Communications Surveys Tutorials, IEEE*, vol. 10, no. 2, pp. 48–69, Second 2008.

[11] CAIDA, "IPv4 and IPv6 AS core: Visualizing IPv4 and IPv6 internet topology at a macroscopic scale in 2013," 2014. [Online]. Available: http://www.caida.org/research/topology/as_core_network/2014/ [Last accessed: 03.11.2014]