On inferring regional AS topologies

Yohei Kuga
KEIO University
sora@sfc.wide.ad.jp

Kenjiro Cho
IIJ
kjc@iijlab.net

Osamu Nakamura
KEIO University
osamu@sfc.wide.ad.jp

ABSTRACT

Analysis of AS-level topologies is an effective way to understand the Internet structure and its development. An AS topology represents logical connection relationships of networks, and thus, is independent of geographic locations of connections. However, the actual Internet topologies are constructed with geographical and physical limitations. In the meantime the global Internet topology is shifting from a U.S. centered star topology to a more distributed topology interconnecting regional hub ASes. Therefore, regional topology views would provide new insights to understanding changes in the Internet structures.

Our goal is to understand differences and similarities amongst Internet structures in regional views such as from Europe and from Asia. We examine inference method for regional AS topologies, and compare the resulting regional AS topologies. Our inference technique extracts AS boundaries in traceroute data collected from multiple vantage points. Then, we illustrate the resulting regional AS topologies using AS Core Maps in order to identify the degree structure and AS locations in each regional view. We show that the distributions of AS out-degrees similarly follow power-law but Tier-1 ASes play different roles in regions. We also identify limitations such as the lack of vantage points in Africa and Oceania in the current data sets.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network topology*

General Terms

Measurement

Keywords

Internet topology

1. INTRODUCTION

Analysis of AS topologies is an effective way to understand the Internet structure and its development. We can analyze the AS topology by two techniques: traceroute data[1] collected from multiple vantage points and BGP (Border Gateway Protocol) data[2]. Meanwhile, with the Internet being globalized, the global traffic has become increasingly common between two countries. In order to achieve lower latency and reduce the cost, the global Internet topology is shifting from a U.S. centered star topology to a more distributed topology interconnecting regional hub ASes with the advent of regional IXes and core ISPs in each region.

Originally, a topology represents logical connection relationships of networks, and thus, is independent of geographic locations. Therefore, the target of the existing methods is only logical topologies such as AS-level or Router-level topologies. However, the actual Internet topologies are constructed with geographical and physical limitations. To understand the Internet structure within a geographic region such as city-level, country-level or continent-level, we need to analyze the Internet topology taking geographic information into consideration. To analyze the regionality of the Internet, it is important to observe not only AS relationships but also the locations of AS boundaries.

Our goal is to understand differences and similarities among Internet structures in regional views such as from Europe and from Asia. Such regional topologies will allow us to identify specific features of the Internet structures in the region as well as the status of the Internet deployment in the region. Also, it would be useful for policy making and investment planning. For example, an enterprise customer may want to know an ISP which well covers geographic areas of interest. Or, an ISP or ASP may want to prioritize investment in facilities at regional hubs.

As the first step of the research, we examine methods of continent-level AS topology analysis using traceroute

data measured from multiple vantage points. The target continents are Africa, Asia, Europe, Oceania, and the United States. The United States is further divided into four regions by time zones: US-EST, US-CST, US-MST, and US-PST. One of the challenges is to find an effective method to divide geographically adjacent regions from the traceroute data.

Our analysis uses the topology data measured by the Skitter project[3] of CAIDA. Our method extracts AS boundaries in the same region from the Skitter data. And this method infers the geographic information between AS boundaries by using RTT between the IP address pair at an AS boundary. Finally, we extract the continent-level AS topology structures by clustering the AS boundary addresses in each region.

This method allows us to analyze and compare the structures of the AS topology of each region in the United States, Africa, Europe, Asia, and Oceania. We compare continent-level AS topology structures by out-degree distributions and AS Core Maps[4]. The results of the out-degree distribution analysis show that the distribution of AS out-degree in each continent follows power-law, which is common among the regions. Moreover, the AS Core Maps of regional views showing the AS relationships among ASes in each continent reveal differences in the role of Tier-1 ASes in each region.

We also identify that the current vantage points of the Skitter project are not sufficient in some regions for our purposes, especially in Africa and Oceania. In the future, we are planning to provide more vantage points in order to analyze AS topologies in such regions as well as in more details; e.g., in the country-level or in the city-level.

Our work makes a contribution to the field of Internet topology research and macroscopic capacity planning. The regional Internet topology analysis allows a geopolitical point of view to determine which AS is better connected. When the relationship between the Internet topology and geographical location is shown, we can identify regional hubs with the presence of major ISPs in the area. And these information can be used for macroscopic capacity planning for the Internet. For instance, there is an important challenge to identify the presence of a single point of failure in a region or between regions.

Section 2 presents the analysis methods of the existing AS topology research related to our research. In Section 3, we propose a method for inferring continent-level AS topologies. Section 4 shows the analysis results of our approach. In Section 5, we compare the continent-level AS topologies, and discuss the results. In Section 6, we describe the present activities and problems for future research. Section 7 concludes the paper.

2. RELATED WORK
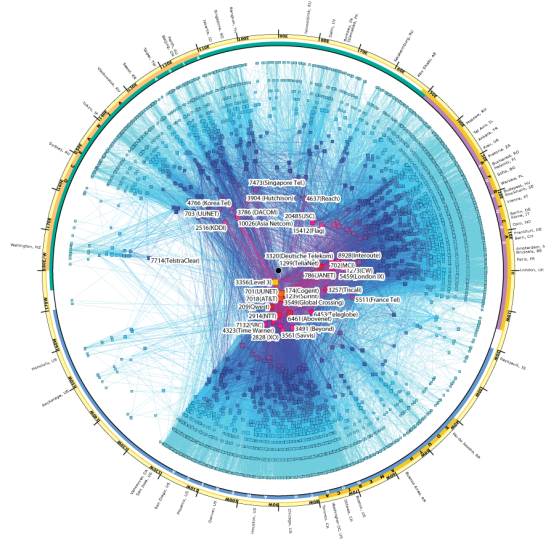
Figure 1: CAIDA's IPv4 AS Core Map, January 2008

There are a large number of the existing research projects that analyze the Internet topologies[5, 6, 7, 8]. The analysis methods of AS topologies can be classified into three by data types: traceroute data from multiple vantage points, the AS path information in BGP tables, and the AS relationship information in the WHOIS database[9]. The Skitter project is collecting the Internet topology data with the Skitter and the Scamper[10] tools derived from the traceroute tool for analyzing AS topologies as well as for their visualization. The Skitter data also includes the AS number of each hop by means of mapping hop's IP address to the origin AS of the best matching prefix found in the BGP table. The data measured from multiple vantage points is used to infer global AS topologies.

The AS Core Map is a visualization method of an AS topology used by the Skitter project of CAIDA. Figure 1 shows the AS Core Map 2008 created by CAIDA. The map plots ASes by mapping the registered geographic location of the AS in the WHOIS database to the corresponding longitude. A link between two ASes shows that the two ASes are directly connected. The radius of an AS, the distance from the center, is computed based on the out-degree of the AS that is a number of its outgoing links (shown in Equation 1) so that ASes with higher out-degree are placed closer to the center. Thus, the AS Core Map allows to identify AS structures by the geographic locations of ASes and their coreness based on the out-degree.

The original AS Core Map shows the logical topology structure from the global view, and does not provide regional views. Our work tries to provide regional views

by extracting a topology observed in a region, and then, by visualizing it by means of the AS Core Map.

$$radius = 1 - log\left(\frac{outdegree\,(AS) + 1}{max.outdegree + 1}\right)$$
$$angle = \left(\begin{array}{c} \text{longitude of the AS's} \\ \text{BGP prefixes} \end{array}\right) \quad (1)$$

3. TECHNIQUES FOR INFERRING REGIONAL AS-LEVEL TOPOLOGIES

In this section, we discuss the methods for inferring continent-level AS topologies. Our method 1) uses traceroute data from multiple vantage points, 2) extracts the AS boundary's IP address pairs, and 3) infers continent-level AS topologies by clustering IP addresses using geographic information.

The clustering method for continent-level AS topologies are as follows:

1. We extract pairs of IP addresses that are both ends of an AS border link. Then, we cluster the IP addresses by combining two links if they have the same IP address at one end assuming the two links are located in the same place.

2. However, there are cases in which Step (1) combines clusters at distant locations due to Layer-2 wide-area VPNs and MPLS where both ends of a border link are not located in the same place. To exclude these long-haul links, we use the per-hop delay for the border link, and exclude links whose per-hop delay is larger than a threshold.

3. Step (1) and (2) create many small clusters, and it is difficult to merge these small clusters only by their IP addresses. In order to combine clusters in a region, we make use of landmarks for regions. Here, we use known prefixes for IXes and hints in reverse DNS names of the AS border IP addresses. The landmarks are also used to infer the geographic locations of clusters.

The details of each Step are described in the following subsections.

3.1 Clustering AS boundaries using border routers' IP addresses

Our method infers AS topologies in a region at the continent-level by clustering inter-AS links based on IP addresses. This section describes the traceroute data used for the research, and the inference technique to cluster inter-AS links at the continent-level to obtain AS topologies in a region.

AS topologies can be obtained by BGP tables, the WHOIS database, or traceroute data. However, BGP tables and the WHOIS database have few clues for inferring geographic information of AS boundaries. Thus,
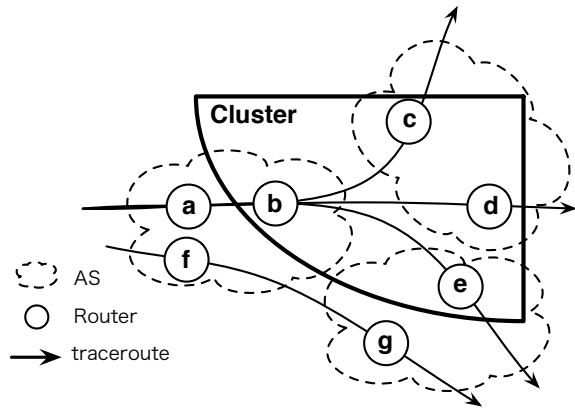
**Figure 2: The clustering method of AS boundaries using common IP addresses.**

our method uses traceroute data from multiple vantage points. The advantages of traceroute data are that we can infer the geographic location from the IP address of an AS border as well as we can make use of RTT information to improve the inference accuracy.

In order to extract AS boundary links from traceroute data, the IP address of each hop in the traceroute data is associated with the corresponding AS number. We create a mapping table from a prefix to the corresponding AS number by extracting the announced prefix and its origin AS number in full-route BGP tables. By longest prefix matching in this mapping table, IP addresses in the traceroute data are mapped to the AS numbers. Then, inter-AS links which have different AS numbers at both ends are extracted.

The next step is to cluster inter-AS links based on common IP addresses. We assume that, if two inter-AS links have a common IP address at either end, they are connected to the same router and the links are likely to be located in the same location. Figure 2 shows the clustering method of AS boundary links based on common IP addresses. *bc*, *bd*, *be*, and *fg* in the figure show AS boundaries. Some AS boundaries include the same IP address of the *b* router. The AS boundaries, *bc*, *bd*, and *be*, are inferred to belong to the same region, while *fg* is not clustered by this method.

3.2 Identifying remote inter-AS links

Both ends of AS boundaries are usually located closely as it is a common practice that an AS peers with another AS at the same location (e.g., at an Internet Exchange Point). However, some inter-AS links have both ends in remote locations, often connected through Layer-2 VLAN or MPLS networks. Therefore, we need to identify such remote inter-AS links and exclude them for the clustering. Otherwise, distant routers connected by a remote inter-AS link would be clustered into the
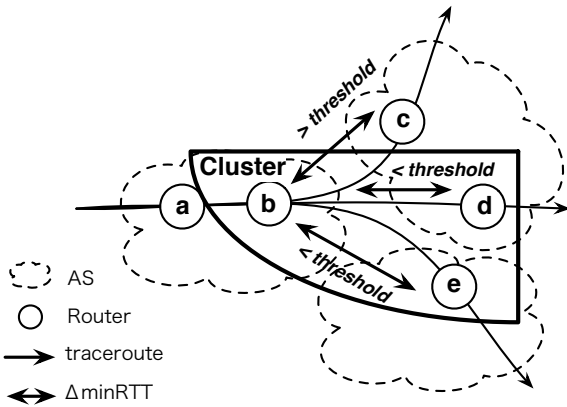
Figure 3: The clustering method with a delay threshold

same location.

In order to identify remote inter-AS links, we make use of RTT information in the traceroute data. The traceroute data includes RTT information from the observer to each hop so that it is possible to infer the delay of a link connecting two AS boundaries. When the delay of an inter-AS link exceeds a threshold value, we assume it be a remote link connecting distant locations, and exclude it from the clustering.

Figure 3 shows the clustering method that considers the delay threshold. *bc*, *bd*, and *be* in the figure show the AS boundaries. These AS boundaries contain the common IP address of the *b* router but the delay for *bc* exceeds the threshold. AS boundaries exceeding the delay threshold are excluded from the clustering. As a result, we infer that only *bd* and *be* (but not *bc*) are AS boundaries in the same region.

3.3 Clustering AS boundaries by DNS names

The clustering method with the delay threshold is not enough and has two issues. One is that it cannot distinguish two regions which are geographically close. For example, it cannot correctly distinguish countries in Europe as some of the European countries are closely located geographically and the delay for such inter-country links could be lower than the threshold. The other issue is that the clustering method based on common IP addresses cannot merge clusters if they do not have a common IP address. Thus, we need another method to improve the accuracy of clustering.

To this end, we use landmarks. Landmarks are IP addresses associated with geographic information. The landmarks are used both to divide and to merge clusters. The landmarks are also used to label the resulting clusters as the geographic information of the landmarks in a cluster is likely to represent the geographic information of the whole cluster.

We use two types of landmarks. One is known prefixes of IXes, and the other is reverse DNS names of routers. Many of inter-AS links are located at major IXes in a region, and the prefixes used at IXes along with their locations are often publicly available (e.g., Packet Clearing House[11] and PeeringDB[12]). We have created a mapping table from the known prefixes used at IXes to the corresponding geographic information (e.g., city and country).

Moreover, it is common amongst major ISPs that routers have reverse DNS names associated with their POP locations. For instance, when the reverse DNS name of a router of NTT Communications (AS2914) is *xe-1-3-0.r21.tokyjp01.jp.bb.gin.ntt.net*, we can infer its location to be in Tokyo because *tokyjp01* in the DNS name indicates *Tokyo, JP*.

The Scriptroute project [13] of Washington University maintains an open-source tool, Undns, to infer the geographic information from DNS names [14]. Undns includes the geographic information database for DNS names of 271 ASes, as of April, 2008. We apply Undns to the IP addresses of the extracted AS boundaries.

4. TOPOLOGY DATA

4.1 traceroute data

The traceroute data measured from multiple vantage points is important for analyzing a global AS topology. We used the traceroute data measured by Skitter of CAIDA. Skitter widely measures path information in the Internet from 16 probes in 11 countries as of March, 2008. The probes exist in the United States, Britain, Spain, South Korea, Ireland, Finland, Philippines, Japan, Australia, Canada, and Morocco.

Table 1 describes the overview of the data measured by Skitter. In this paper, we used path data measured from March 14, 2008 to March 19, 2008. The data format is the scamper format (warts). There were 14,032,566 entries in the measured path data.

Table 1: Skitter data

data format	scamper (warts)
Num of tr-data	14,032,566
Num of probes	16
measured length	2008/03/14 - 2008/03/19

4.2 Clustering and landmarks

We used BGP tables from the Routeviews project[15] to associate IP addresses with the corresponding AS numbers as described in Section 3.1. The BGP table had 275,636 entries.

For landmarks of IX prefixes described in Section 3.3, we used the database published by Packet Clearing House[11].
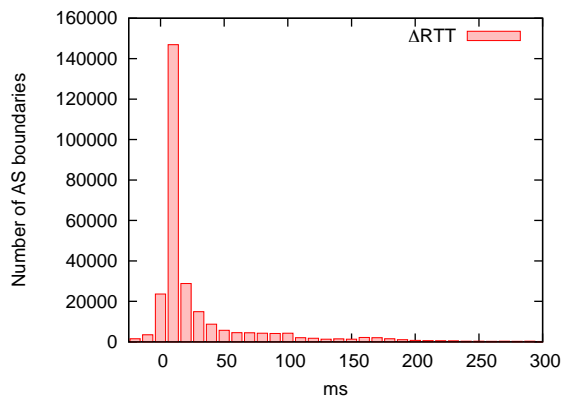
Figure 4: Histogram of difference of RTT in traceroute data

Figure 5: CCDF of AS out-degree

Table 2 shows the number of landmarks used for clustering. The number of AS boundary's IP addresses matching to IX prefixes is 228. The number of AS boundary's IP addresses matching by Undns is 78,192. These landmarks were also used to map clusters to continent names(e.g., Asia, Europe).

Table 2: Number of landmarks

IXes	228
Undns	86,612

5. RESULTS

Figure 4 shows the difference in the delay of all the AS boundaries included in the traceroute data. This figure is a histogram with the bin size of 10ms. The figure shows that 60% of the AS boundaries are connected with the delay equal to or smaller than 10ms. Moreover, there are AS boundaries with the delay more than 200ms, which is likely to be remote inter-AS links. Therefore, it is necessary to exclude these links by a delay threshold.

Table 3: Size of Cluster of AS boundary

Area	AS Boundaries	Uniq ASN
Europe	6,112	3,182
US-EST	2,208	1,287
US-PST	1,816	892
US-CST	1,286	777
Asia	923	583
Oceania	395	248
US-MST	270	214
Africa	40	38

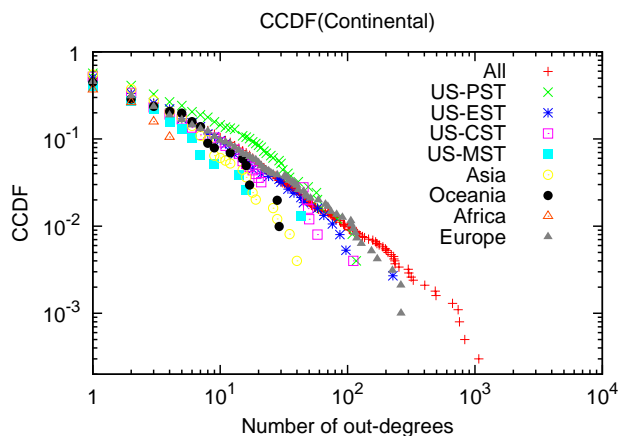The threshold for the delay should be set not to exclude the AS boundaries between countries in the same continent. Thus, the delay threshold is set to 200ms. This threshold was selected to include the longest fiber link within Asia but exclude links between Asia and Europe or the United States. Selecting the optimal value for the threshold is still for future work.

Table 3 shows the size of the clusters for each continent extracted with the delay threshold of 200ms. For the United States, there are four regions with different time zones.

6. COMPARISON OF REGIONAL AS TOPOLOGIES

This paper proposed, and analyzed the method for inferring continent-level AS topologies. This section compares the resulting AS topologies for each continent, and evaluates the results.

6.1 Distribution of AS out-degree

Figure 5 shows Complementary Cumulative Distribution Functions (CCDF) of AS out-degree for each continent. *All* in the figure shows AS out-degree of all the continents. The figure shows that the distributions of AS out-degrees follow power-law, and are similar among the continents, although the number of ASes in Africa and Oceania is insufficient.

6.2 Visualization of AS topologies by AS Core Map

We visualized the continent-level AS topologies by the AS Core Map. This method plots a macroscopic snapshot of the Internet by means of the out-degree and the longitude of ASes. The AS Core Map clearly shows the core ASes in the global Internet and the quantity of connectivity between locations represented by ASes. Our goal is to understand differences and similarities among Internet structures in regional views. So, the AS Core Map provides information about who are hubs
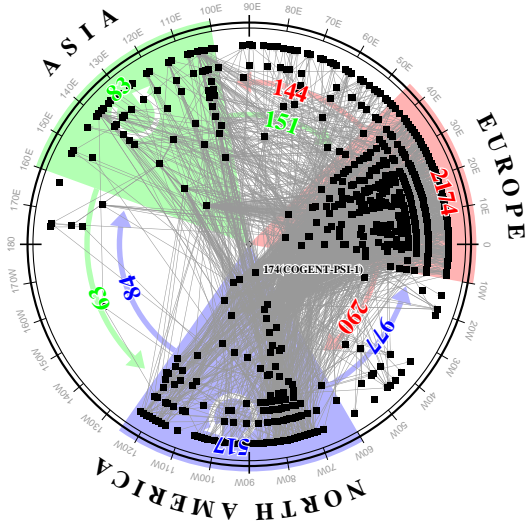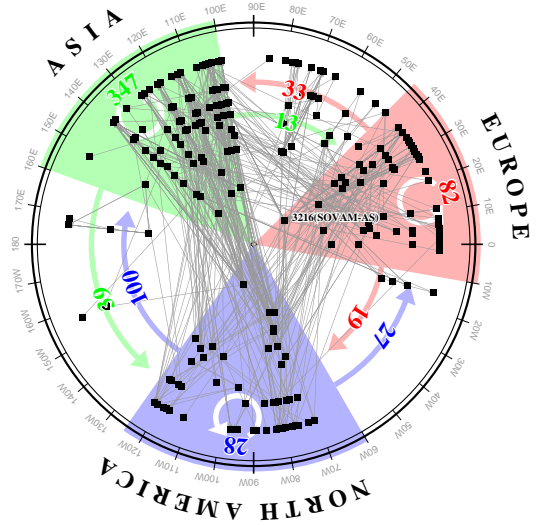
Figure 6: AS Core Map in Europe

Figure 7: AS Core Map in Asia

in each region and differences of connectivity between locations represented by ASes in each region.

We mapped the ASes to geographic locations by using the set of announced prefixes from the BGP table and the WHOIS database. We converted the border router's IP address to AS number using prefixes from Routeviews, and identified the location where an AS registered in the WHOIS database. And then, we converted the AS number to longitude using the Google Maps API[16].

Figure 6, 7, 8, 9 plot the continent-level AS topologies using the AS Core map. Each figure represents a AS topology within Europe, Asia, U.S.(EST area), or U.S.(PST area). The dots in figure show ASes observed in the region, and links show direct connections between ASes in the region. The position of an AS is determined by the number of links and the longitude of the AS. ASes with higher out-degree are plotted closer to the center, and the radius is determined by the longitude of ASes. The arrows in the figure show the connections between the United States, Asia and Europe, and represent the presence of ASes from other continents.

Figure 6 and Figure 7 show the AS topologies in Europe and Asia. It can be observed that, compared with the European AS topology, the Asian AS topology has fewer Tier-1 ASes from the United States, suggesting that there are a limited number of U.S. ASes which have presence in Asia.

In the Asian AS topology, there are a lot of ASes from Moscow (at longitude 40 degrees) and from New Zealand (at longitude 175 degrees). It could suggest the increasing activities of Russian ISPs in the Asian market, although it could be caused by some errors in our clustering.

In the European AS topology, it is observed that there are not only ASes in West Europe but also a large number of ASes in East Europe. It is also observed that there are ASes from Singapore (at longitude 105 degrees) and from Beijing (at longitude 117 degrees).

Figure 8 and Figure 9 show the AS topologies in the East Coast (EST timezone) and the West Coast (PST timezone) of the United States. The EST AS topology has more ASes from Europe and less ASes from Asia. On the other hand, the PST AS topology has more ASes from Asia and less ASes from Europe.

By comparing the continent-level AS topologies, we can observe characteristics of the AS structures in the region. By observing the regional AS topologies, we can also identify hub ASes connecting many ASes in the region. There are differences even within the United States, between the East Coast and the West Coast. On the other hand, we can observe only 40 AS boundaries in the African AS topology, which suggests the lack of vantage points in the current traceroute data sets for this type of analysis.

7. FUTURE WORKS

This section discusses future research topics.

7.1 more traceroute data in some continents

We have compared continent-level AS topologies, and the results reveal the lack of vantage points to cover Africa and Oceania in the current data sets. However, it is challenging to obtain reliable measurements from
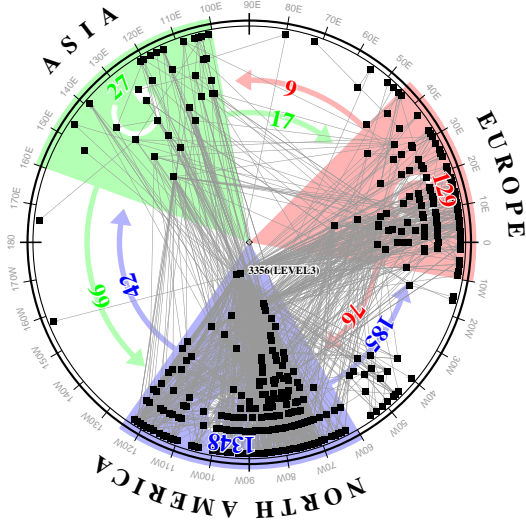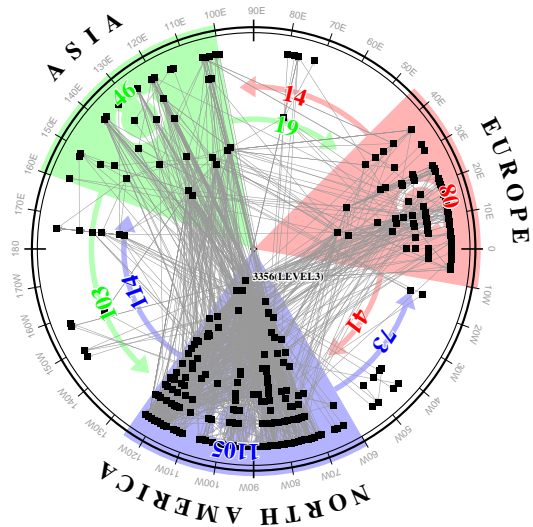
Figure 8: AS Core Map in U.S.(EST)

Figure 9: AS Core Map in U.S.(PST)

developing countries because the measurement environment such as computing facilities and networks are limited and less reliable.

A possible way to improve the coverage is to make use of additional traceroute data collected by different projects, although different projects have different measurement methods and goals so that we cannot blindly merge different data sets.

7.2 Improving the inference method

Our method extracts AS boundaries based on IP addresses associated with AS numbers. However, the actual boundary could be off by one hop because a router could have an IP address borrowed from the AS at the other end of the inter-AS link. Thus, it is difficult to pinpoint the inter-AS link only from a single traceroute output, but we can use heuristics similar to solving interface aliases of a router in order to improve the accuracy.

In this paper, we have focused on continent-level AS topologies. For the future, we intend to extend the method for country-level and city-level AS topologies. However, it will require considerable improvement in the delay threshold and the landmarks. We have used the fixed delay threshold of 200ms, but the threshold should be adaptively selected considering the location of a link in question. For example, the threshold can be raised in Japan as being islands but needs to be lowered to distinguish European countries.

7.3 Inferring the AS topology between regions

We have compared AS topologies only within regions,

and did not look into AS boundaries crossing regions. However, in order to understand how the global Internet looks from a region, it is important to include inter-AS links to the outside of the region.

8. CONCLUSIONS

The existing Internet topology research has examined mainly logical connectivities at the AS-level and the router-level. In this paper, we have proposed the methods for inferring regional AS topologies, and shown the continent-level AS topologies using traceroute data measured from multiple vantage points. From the results, we have identified, to some extent, similarities and differences in the regional AS topology structures. However, there are still a lot of rooms to improve in the analysis and evaluation methods. Currently, our approach is targeted for continent-level AS topologies but we intend to extend it to the country-level and city-level in the future. We believe that it is essential to understand regional characteristics in the global Internet, and we hope our research will be of help to make the Internet truly global.

9. REFERENCES

[1] traceroute.
    http://www.traceroute.org/#source%20code.
[2] Y. Rekhter and T. Li. A Border Gateway Protocol
    4 (BGP-4). RFC 1771, IETF, March 1995.
[3] K. Claffy, Tracie E. Monk, and Daniel McRobb.
    Internet tomography. *Nature*, January 1999.

[4] Bradley Huffaker, Daniel Plummer, David Moore, and k claffy. Topology discovery by active probing. *Symposium on Applications and the Internet (SAINT)*, January 2002.

[5] Michalis Faloutsos, Petros Faloutsos, and Christos 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*, pages 251–262, New York, NY, USA, 1999. ACM.

[6] Sharad Jaiswal, Arnold L. Rosenberg, and Don Towsley. Comparing the Structure of Power-Law Graphs and the Internet AS Graph. In *ICNP '04: Proceedings of the Network Protocols, 12th IEEE International Conference*, pages 294–303, Washington, DC, USA, 2004. IEEE Computer Society.

[7] Shi Zhou and Raul J. Mondragon. Accurately modeling the Internet topology, 2004.

[8] Lun Li, David Alderson, Walter Willinger, and John Doyle. A first-principles approach to understanding the internet's router-level topology. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–14, New York, NY, USA, 2004. ACM.

[9] Priya Mahadevan, Dmitri Krioukov, Marina Fomenkov, Xenofontas Dimitropoulos, k c claffy, and Amin Vahdat. The Internet AS-level topology: three data sources and one definitive metric. *SIGCOMM Comput. Commun. Rev.*, 36(1):17–26, 2006.

[10] Scamper. *http://www.wand.net.nz/scamper/*.

[11] Packet Clearing House - Internet Exchange Directory. *http://www.pch.net/ixpdir/Main.pl*.

[12] PeeringDB. *https://www.peeringdb.com/*.

[13] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A public internet measurement facility. *4th USENIX Symposium on Internet Technologies and Systems*, 2002.

[14] Undns. *http://www.scriptroute.org/source/*.

[15] University of Oregon RouteViews Project. *http://www.routeviews.org/*.

[16] Google Maps API - Google Code. *http://code.google.com/apis/maps/*.