

Characterization of the Internet at the AS Level

Inas Khalifa
School of Engineering Science
Simon Fraser University, Burnaby, BC V5A 1S6
ikhalifa@sfu.ca

Abstract

Modeling the Internet topology is an important open problem that has attracted researchers attention in the past few years. This report presents an overview of the different research activities related to Internet characterization at the AS level. It describes the different methods used to measure the Internet AS topology, as well as some of the recent visualization and modeling techniques. It then outlines the different topological properties inferred from AS graphs, and their significance. In particular, it focuses on the AS degree distribution demonstrating its high variability using snapshots of AS topology, and attempts to explain its causes. Finally, it presents a software tool specially developed to facilitate manipulation and analysis of AS topology data.

1 Introduction

There has been a significant increase in research activities related to studying the topology of the Internet in the past few years. These activities include collecting and visualizing topology information, analyzing and modeling these measurements to infer the Internet's connectivity graph and describe its properties, and explaining the origins and causes of some of the observed surprising features [1].

Studying the Internet's topology can be useful in a multitude of ways. It allows designing more efficient protocols that take advantage of topological properties [2]. It also facilitates creating realistic network models for simulation purposes, by building topology generators that produce graph structures that closely match the measured Internet connectivity graphs, e.g. [3]. Moreover, estimating topological parameters of Internet graphs is useful for protocol analysis and speculation of the Internet topology in the future.

The Internet connects thousands of *Autonomous Systems* (ASes) operated by many different administrative domains such as ISPs, companies, and uni-

versities [4]. An AS is defined as a connected group of one or more IP prefixes run by one or more network operators, which has a *single* and *clearly defined* routing policy [5]. In other words, an AS appears to other ASes as having a single coherent interior routing plan and presents a consistent picture of what networks are reachable through it. An AS is identified by a globally unique 16-bit number. In North America, AS numbers are assigned by the American Registry for Internet Numbers (ARIN).

The infrastructure of the Internet thus can be studied at two different granularities, both using graph representation, as shown in Figure 1. At the **router level**, each graph vertex represents a router, and each edge represents a link between two routers. Alternatively, the Internet can be studied at the **AS level**, where vertices represent ASes and edges represent inter-AS interconnections [6].

This report presents an overview of the ongoing research activities related to Internet characterization at the AS level. It describes the different methods used to measure the Internet AS topology, and outlines some of the recent visualization and modeling techniques. It then describes the different topological properties inferred from AS graphs, and their significance. In particular, it focuses on the AS degree distribution demonstrating its high variability using snapshots of AS topology, and attempts to explain its causes. Fi-

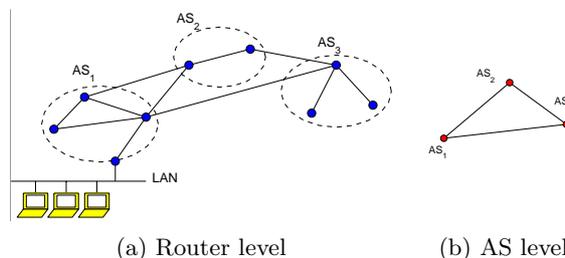


Figure 1: Representing the Internet infrastructure.

nally, it presents a software tool specially developed to manipulate and analyze AS topology data.

The rest of this report is organized as follows. AS topology measurement techniques are discussed in Section 2. Visualization and modeling methods are described in Section 3. Section 4 examines the high variability of the AS degree distribution and its causes. Section 5 describes the implementation and applications of the topology analysis tool. Finally, the conclusion is given in Section 6.

2 Mapping Techniques

The Internet, unlike the networks it is composed of, has no authority defining its topology evolution. That is why a fully detailed map of the Internet is not available [7]. There are two main approaches used to collect Internet topology data. The first technique collects topology information using routing tables obtained from routers in different ASes. The other technique uses probing to construct a router-level map and then constructs an *AS overlay*, i.e. groups routers into ASes. These two techniques are described in this section.

2.1 BGP-Based Techniques

Routing within an AS is accomplished using *interior routers* which implement algorithms such as OSPF, IS-IS, and RIP [4]. Routing between ASes uses the Border Gateway protocol (BGP) and is accomplished by *exterior routers*, or border gateways. A pair of ASes interconnect via one or more dedicated links and/or public network access points. Each such link connects a border gateway from each AS. Each AS contains a router called *border gateway speaker*, which implements BGP [8]. Each BGP speaker keeps a *BGP routing table*, which stores a set of candidate routes for that router.

BGP is a path-vector protocol that constructs paths by successively propagating *updates* between pairs of BGP speakers that establish peering sessions [9]. Each update concerns a particular prefix (a partially masked IP network address used for CIDR) and includes the list of ASes along the *AS path*. Each BGP speaker creates updates for one or more prefixes, and can send these updates to its immediate neighbors, or *peers*, via BGP sessions. The simplest path-vector protocol would employ a shortest AS path routing policy. BGP, however, allows a variety of routing policies in order to honor contractual agreements that control traffic exchange between ASes. Each AS sends *only*

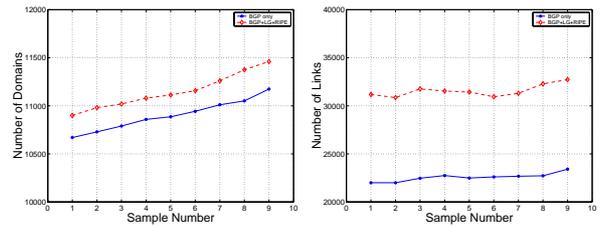


Figure 2: The incompleteness of BGP data.

its best route for a prefix to a neighbor AS. Upon receiving an update, a router uses its *import policies* to decide whether or not to use that update's AS path. Moreover, if the path is chosen, the router decides whether or not to propagate the update to its neighboring ASes according to its *export policies* [4].

By construction, each BGP routing table contains two kinds of connectivity information. It contains the AS connectivity map, i.e. its immediate neighbors, in addition to portions of the connectivities of its neighbors—since an AS path may traverse more than two ASes. As such, BGP routing tables from different ASes contain overlapping subsets of the Internet AS connectivity. Thus, by carefully aggregating the information in such tables, an AS interconnection map can be constructed. Since November 1997, The University of Oregon Route Views Project [10] has been collecting BGP routing tables once daily from a route server with BGP connections to multiple geographically distributed target operational routers. It processes this information (more than 40 Mbytes of data daily) and constructs AS interconnection maps labeled by AS numbers.

On the plus side, BGP tables are available publicly and are relatively easy to parse, process and comprehend. This is why it has been used by many researchers, who do not collect their own data, to study the Internet topology, e.g. [6, 7, 11]. However, since BGP tables only include the selected (best) routes, rather than all possible routes stored in the router, topology maps obtained by merging BGP tables tend to be sparse [12]. Change *et al.* performed an extensive study to examine the incompleteness of topology information deduced from BGP tables [1]. They used two additional data sources to augment BGP data. First, they used BGP summaries available from *Looking Glass* sites, which are BGP routers owned by ISPs. ISPs made available public but limited access to those routers, in order to troubleshoot Internet-wide routing problems. Second, they used the Internet Routing Registry's (IRR) RIPE database, which contains

relatively up-to-date information about ASes in Europe. Figure 2 illustrates the incompleteness of the data obtained using BGP tables only. The figures illustrate the number of ASes and the number of AS inter-connections obtained using BGP tables only, versus those obtained using BGP only, in addition to data from looking glass sites (LG) and the RIPE database. The results were generated using nine samples collected weekly starting in mid March 2001 (shown on the x -axis). The figure indicates that BGP only data, while missing only about 2% of the ASes, it misses more than 20% of the interconnections.

2.2 Probing-Based Techniques

The process of constructing an AS map using probing involves two stages. First, a router-level map of the Internet is constructed. Then, ASes and their connectivity relationships are inferred from the router-level map, thus creating an AS overlay. Techniques for constructing router-level maps are typically based on the *traceroute* utility. Traceroute is a utility that records the route that a packet follows through the Internet between a source and a specified destination host. First, the utility sends a packet with a *time to live* (TTL) value of 1, designed to be exceeded by the first router that receives it. The router drops the packet and returns an error message to the source host. This enables traceroute to determine the time required for the hop to the first router, as well as its IP address. Then, the source increments the packet's TTL by 1 and sends it so that it reaches the second router in the path to the destination. The second router will drop the packet and reply with an error message. This process is repeated until the packet reaches its destination.

Mapping techniques use the traceroute utility in different ways in order to achieve different levels of efficiency and coverage (number of destinations discovered). Instead of selecting destination addresses randomly, it is customary to incorporate some intelligence, or heuristics, to decide how to choose targets for probing [13]. It should be noted that the graphs constructed using trace routes are in essence *directed graphs*, since they capture forward paths, thus they reflect the Internet connectivity more accurately.

The following are examples of router-level mapping techniques. Cheswick and Burch [14] started a large-scale Internet mapping project in 1997. They have constructed a map of the Internet containing about 90,000 routers and 100,000 links. However, since they use a single source (Lucent) for probing, their data set renders a connectivity coverage bias towards their

transit provider. Govindan and Tangmunarunkit [15] have developed *Mercator*, a program that uses informed random address probing to carefully explore the IP address space, uses source-route capable routers to enhance the fidelity of the resulting map, and employs mechanisms for resolving aliases (interfaces belonging to the same router). However, *Mercator* is relatively slow at processing probes and its use of source routing tends to generate user and ISP complaints and is less practical for large-scale studies. Broido and Claffy [12] run the *Skitter* measurement tool developed at the Cooperative Association for Internet Data Analysis (CAIDA) [16], on more than 20 monitors around the globe, collecting forward path and round trip time to about 400,000 hosts. They constructed the largest IP graph available so far, consisting of about 665,000 nodes, by processing about 220 Million traceroutes.

An AS overlay is constructed from the router-level map by assigning an AS to each router in the map using the information available from BGP routing tables. More precisely, for a given router interface address, the matching route entry in the BGP table is found, and the router is assigned to the origin AS in the AS path associated with that route entry. Then, a collapsing algorithm is used to generate the AS map by grouping sets of routers into their ASes. Tangmunarunkit *et al* [17] as well as Chang *et al* [18] used this technique to construct AS maps. Although probing-based techniques provide richer connectivity than that produced by BGP techniques, they usually don't capture all existing ASes as indicated in [18].

2.3 Comparison

After describing the two mapping approaches, a natural question that comes to mind is which one is better. Table 1 compares the different aspects of the two approaches. As the table indicates, each technique has its own strengths and weaknesses. Thus, it can be concluded that data acquired using either technique is not sufficient to construct a representative AS-level topology, and hence both techniques must be used in order to correctly infer the properties of AS topology. It is also worth noting that even if both data sets are combined, it is still difficult to construct a complete map of the Internet due to security restrictions, since not every AS administrator is willing to share their peering agreements, or allow probing packets to go through.

Table 1: Comparison of the two AS mapping approaches.

	BGP-based techniques	Probing-based techniques
Algorithm complexity	Straight forward	Complex
Output data	High-level connectivity graph	Low-level directed reachability graph
Availability	Publicly available (WWW)	Available upon request
Completeness	Less interconnections (edges)	Less ASes (vertices)
Graph edges	Represent peering relationships	Represent physical connectivity

3 Visualization and Modeling

Once the AS adjacency graph is acquired, the next step is to plot, or *layout*, the data in a meaningful way. There are two types of graph layout techniques, namely, geographical and logical layout [19].

Geographical layout techniques place each node on a map at its precise physical location, e.g. longitude and latitude, and then draw lines between connected nodes. In 1997, CAIDA developed *MapNet*, a public Java applet for visualizing the infrastructure of multiple international backbone providers simultaneously. It divides each backbone into a group of nodes (POPs) and pipes between these nodes, and draws them based on their geographical location on a map of the world [20]. Since then, CAIDA developed other visualization tools such as GeoPlot and MantaRay. The main complexity in designing geographical layout visualization tools is in rendering very large data sets and allowing users to interactively manipulate them. However, positioning the nodes and edges of a graph is straight forward, since the positions of the graph elements are available. Unfortunately, determining the exact geographical locations of network elements is not an easy task due to a number of reasons. First, ISPs and network administrators are rarely willing to disclose exactly where their routers are. Moreover, using a machine's name to infer its location is still difficult, since some host names do not explicitly include information about their location, and even when they do, it generally only indicates the location at a coarse resolution, i.e. city or state name. Additionally, many hosts on the Internet do not have host names, they are only identified by an IP address.

Logical layout techniques position graph elements based on their connectivity rather than physical location. Such techniques reveal information that may be more valuable to a researcher than mere distance, e.g. the number of hops between two routers. The design of logical layout techniques is more difficult than geographical techniques. Since, in addition to having to render and manipulate huge data sets, they also have to determine where to position graph nodes such that

the graph remains *readable*. Cheswick *et al.* modeled router-level maps as physical systems where nodes are connected by springs. The optimal node positions are then calculated as those that minimize the total energy of the system [14]. They also suggested that computing and plotting the minimum distance spanning tree of a graph, though incomplete, is relatively faster and achieves better graph readability.

Huffaker *et al.* of CAIDA developed *Otter*, a general-purpose visualization tool that allows both layout techniques [21]. The main strength of *Otter* is that it is capable of displaying large data sets that originate from different environments such as topology discovery data, Internet core routing tables and web site directory structures. It allows both geographical placement and logical placement as well as *quasi-geographical* placement when only partial geographic information is available. A sample map of the Internet AS core network is available at CAIDA's web site [22]. It contains 7,563 ASes and 25,005 peering sessions.

3.1 High-Level Models

Visualization tools convert the huge data sets representing the Internet AS topology into user-friendly, and often fascinating, graphical models. However, since they attempt to show all the available information, it is difficult to draw any quantitative conclusions, just by examining the maps visually. Moreover, although the Internet is assumed to be hierarchical by construction, the maps are usually too densely connected to any sort of hierarchy. Thus, the need arises for higher level models that describe the Internet while hiding the overwhelming details.

In the very early days of Internet, the AS topology was approximately tree-structured. A single national *backbone* provided transit to several *regional providers*, each serving one or more metropolitan areas. Campus networks, or *stubs*, obtained global connectivity by connecting to their respective regions [23]. The Internet has evolved tremendously since then. Now several backbones are interconnected, regional service providers may interconnect with each other as well as

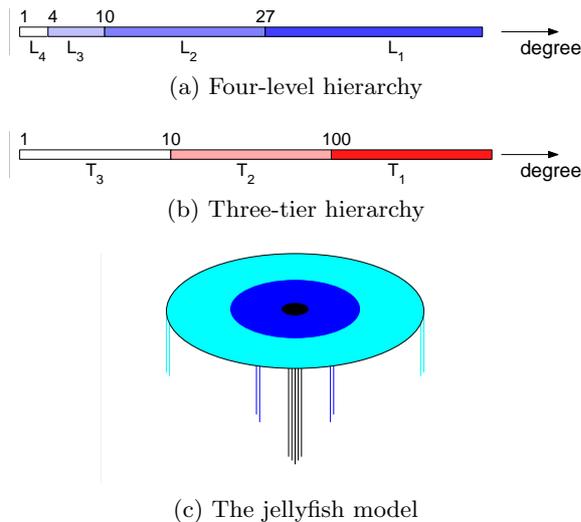


Figure 3: Hierarchical models of the Internet.

with more than one backbones, and stubs may connect to one or more regionals or backbones.

Govindan and Reddy used the distribution of the AS degrees to decompose the Internet’s AS topology into a hierarchy with *four* levels, as in Figure 3(a). They used three snapshots of the AS topology collected every six months starting in November 1994. They plotted the cumulative distribution function, *CDF*, of AS degrees in each snapshot. They found that the *CDF* is well approximated by a piece-wise linear function with four components, where each component defines a degree range. They suggested that the four ranges represent a *natural* hierarchy of ASes classified by function, i.e. level L_1 corresponds to national or international backbones, L_2 corresponds to large North American regional or European national providers, L_3 corresponds to smaller regional providers, and L_4 corresponds to multi-campus corporate or academic networks [23]. It is not likely that such classification will hold for the current AS topology. Chang *et al.* also used the AS degree distribution to categorize ASes. However, they divided ASes into *three* hierarchical levels (Figure 3(b), *Tiers 1, 2,* and *3*, where the degrees of Tier-1 (Tier-2) ASes are one order of magnitude larger than those of Tier-2 (Tier-3) ASes [1].

Tauro *et al.* proposed yet another conceptual model of the Internet AS topology. They first identify the *core*, L_0 , as the maximal clique that contains the highest-degree nodes. Then, they define the first layer, L_1 , to contain all the nodes that are neighbors of the core. Similarly, they defined L_2 to contain neighbors of L_1 except for L_0 . By repeating the same procedure,

they identified *five* layers, in addition to the core [11]. They showed that the *importance* of the nodes decreased as we move away from the core, which validates their selection of the core. They suggested that using this model, the Internet AS topology can be visualized as a *jellyfish*, Figure 3(c), where the core is the center of the cap of the jellyfish surrounded by layers of nodes, called *shells*, and the legs of the jellyfish are those degree-one nodes that attach to the different shells. Although this model seems simple enough, it must be noted that they have used BGP data to construct it. As pointed out in Section 2, BGP data only presents a sketchy view of the Internet AS topology. Thus, it remains to be seen if such a model will hold when applied to data acquired using other techniques.

4 AS Degree Distribution

To gain more insight into the structure of the Internet’s AS graphs, quantitative analysis of the topological properties of such graphs is essential. The properties of an AS graph can be broadly classified into the following types:

- **Global properties:** such as the total number of ASes and the total number of AS interconnections, e.g. Figure 2, the maximum and minimum AS degree, and the graph diameter. Such properties provide a global picture of the data and can be used to determine, for example, how much the Internet’s AS structure has changed over time.
- **Averages:** such as the average AS degree and the average number of ASes within h interconnections, or hops, from each other. Since average values imply a uniform data distribution, they do not provide an accurate view of the data, and usually miss the interesting features that one would want to capture [6].
- **Distributions:** such as the probability density function (*pdf*)¹, cumulative distribution function (*CDF*), and the complementary *CDF* (*CCDF*) of the different graph properties. Unlike averages, examining the distributions provides a simple means for identifying the existing *trends* in property values.

The rest of this section is devoted to examining the distribution of the AS outdegree. The outdegree, or simply degree, of an AS is defined as the number of

¹A more accurate term to use is the frequency distribution or *histogram*

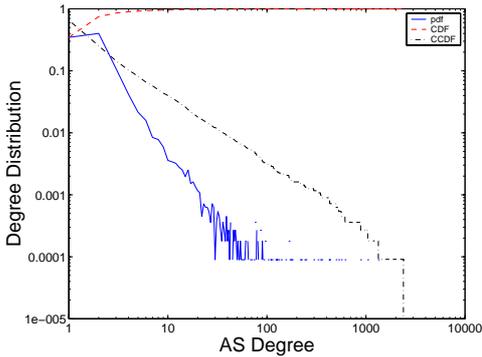


Figure 4: *pdf*, *CDF* and *CCDF* of AS degree.

interconnections that this AS has. Alternatively, using graph representation, the degree, $d(v)$, of a graph vertex v is defined as the number of graph edges with which it is incident. Analysis of the AS degree distribution has been gaining increased interest from researchers over the past few years. The reason for that is, in most cases, it is among very few pieces of information that one can use to classify ASes, as discussed in Section 3. Recall that AS topology data is usually *anonymized* to protect the privacy of AS owners.

Figure 4 depicts the *pdf*, *CDF* and *CCDF* of the AS degree in the AS topology acquired from BGP routing tables during the week of May 26th, 2001, using log-log scale. The figure indicates that the observed AS degrees range over three or four orders of magnitude. In other words, the local connectivity of ASes is free of scale [24]. Such a distribution is referred to as a *heavy tailed* or *highly variable*, meaning that it doesn't decay exponentially as the AS degree increases. Also, note that the *CCDF* illustrates this high variability more clearly than both the *pdf* and *CDF*. Other topology samples obtained over different time periods starting as early as 1997 produce similar AS degree distribution.

Faloutsos *et al* used similar topology samples to propose that the frequency, f_d , of occurrence of an outdegree d is proportional to the outdegree to the power of a constant [6], i.e.,

$$f_d \propto d^o, \quad o \text{ is a constant.}$$

In other words, the frequency of AS degrees (which is equivalent to the *pdf*) follow a *power-law*. However, the topology data acquired using BGP-based technique only misses a large number of the actual AS interconnections, as shown in Section 2. Figure 5 illustrates the *CCDF* of AS degrees for the same nine samples used in Section 2. The figure indicates that

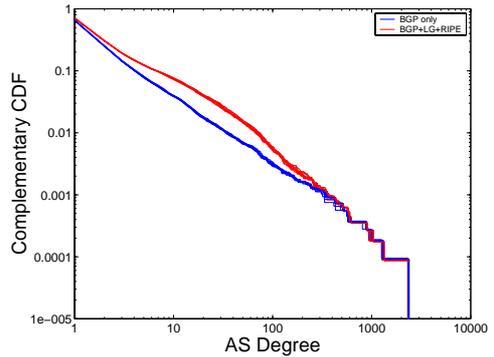


Figure 5: *CCDF* of AS degree for BGP data and BGP+LG+RIPE data.

the two distributions are highly variable. However, there is a noticeable difference in the AS degree range from about 5 to 200. This observation stresses the conclusion drawn in Section 2, that one should not use a single mapping technique exclusively, to determine the properties of the Internet AS topology. The next section elaborates on the causes of the high variability of AS degrees.

4.1 Explaining the High Variability of AS Degrees

In 1999, Barabasi and Albert proposed a simple model for random networks such as the actor collaboration graph, WWW graph, and the citation patterns of the scientific publications [24]. This model, the *BA model*, produces a power-law vertex degree distribution, by using two dynamics:

1. **Incremental Growth:** The model starts with n_0 nodes. A single new vertex is added at every time step. The new vertex connects to m existing vertices.
2. **Preferential Connectivity:** The probability, $P(k)$, that a new vertex connects to an existing vertex i depends on the connectivity or degree, k_i , of that vertex, i.e. $P(k) = \frac{k_i}{\sum_j k_j}$, i.e. a *rich-get-richer* mechanism.

Chen *et al* showed that the BA model can not be used to explain the high-variability of AS degrees in the Internet AS topology. They showed that although the AS topology grows incrementally, the degrees of the new vertices follow a distribution that heavily favors degrees 1 and 2, instead of the fixed number m . Also, the new ASes have a much stronger preference

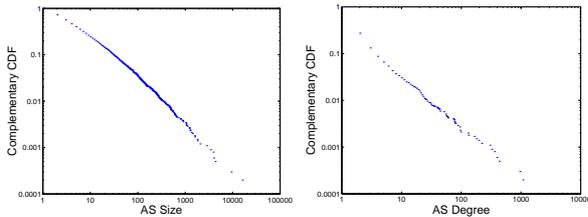


Figure 6: High variability of AS sizes and degrees.

to connect to high-degree ASes than predicted by the BA simple linear preferential model. Moreover, as shown in Figure 5, the AS degree distribution does not strictly follow a power-law.

Tangmunarunkit *et al* proposed an alternative explanation for the high variability in AS degrees. They showed that AS sizes are highly variable, where the size of an AS is defined as the number of routers in that AS. They found that the AS degrees are strongly correlated with AS sizes, having correlation coefficients around 94%. Thus, they concluded that the high variability of AS sizes is likely the cause of the high variability AS degrees. Figure 6 depicts the *CCDF* of AS degrees and sizes computed using AS topology acquired in May of 2001². The figure clearly shows the high variability of both AS sizes and degrees and the correlation coefficient was found to be 95.9%.

5 The Topology Analysis Tool

An AS topology analysis tool has been specially developed for the purpose of this project. The main goal of this tool is to provide a simple means to analyze, manipulate, and aggregate multiple topology data instances. It also facilitates comparing the topological properties of different data samples, e.g. samples collected at different points in time using the same data source or samples collected at the same time period but using different mapping technique. A screen shot of the tool’s graphical user interface (GUI) is shown in Figure 7. More specifically, the following are examples of possible uses of this tool:

- By loading a single topology data sample, the tool automatically calculates and prints the total number of ASes and their interconnections, and plots the *CCDF* of the AS degree. For example, the first seven samples shown in Figure 7 represent daily measurements of the AS topology. By

²The data used to generate this figure was supplied by H. Tangmunarunkit of USC-ISI

inspecting the results, it is clear that the AS degree distribution is constant, and that there are very slight differences between the properties of each sample.

- By loading multiple data samples, the tool aggregates the information from all samples, reports the contribution of each sample, and again computes the topological properties. For example, by simultaneously loading the seven samples mentioned above, a snapshot of the AS topology during one week is constructed. This snapshot contains a more accurate picture of the topology by avoiding possible short-term AS router failures.
- The tool allows saving the currently displayed data, possibly consisting of multiple data samples, into a single topology graph. In other words, one can combine data samples measured over a period of time into a single graph.
- By viewing samples collected using different methods, one can easily determine how the methods differ. For example, the first two samples in Figure 8 depict two topology data sets collected during the same week, the first obtained using BGP routing tables only, while the second is augmented by LG and IRR’s RIPE data. It is easy to see that the second set contains more interconnections, and that the degree distributions of the two data sets are quite different.
- The last two samples in Figure 8 are *synthetic* samples used to illustrate the two extreme cases in the size-degree relationship. In the first sample, AS sizes are chose so that they are perfectly correlated with the AS degrees, while in the other sample, the AS sizes are selected randomly for each AS.

The topology analysis tool was developed in Java, and thus can be used on multiple platforms.

6 Summary and Conclusion

This report presented an overview of the different research efforts related to Internet characterization at the AS level. It described the two approaches typically used in measuring the Internet AS topology. It also outlined some of the recent visualization methods as well as high-level hierarchical models of the Internet. It then focused on the AS degree distribution demonstrating its high variability and explained its possible causes. Finally, it described the implementation of an

AS topology analysis tool that was developed to facilitate the analysis, manipulation, and comparison of the topological properties of different AS graphs.

The main conclusions that can be drawn from the material presented in this report can be summarized in the following points:

- Analyzing the topology of the Internet is very important as it can be used to improve protocol design and network simulation.
- In order to obtain a more complete picture of the Internet AS topology, data acquired using both BGP-based techniques as well as probing-based techniques must be integrated.
- High-level models provide a more compact way to visually illustrate the hierarchy of the Internet AS topology.
- AS degrees are highly variable, quite likely due to the high variability of AS sizes.

Future work will involve obtaining AS topology samples from probing-based techniques and comparing them with BGP-data samples acquired at the same period.

References

- [1] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "On inferring AS-level connectivity from BGP routing tables," <http://citeseer.nj.nec.com/chang01inferring.html>.
- [2] C. Labovitz, A. Ahuja, R. Wattenhofer, and S. Venkatachary, "The impact of Internet policy and topology on delayed routing convergence," in *Proceedings of the IEEE INFOCOM '01*, Anchorage, AL, April, 22-26 2001, vol. 1, pp. 537–546.
- [3] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *Proceedings of the 9th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, Cincinnati, OH, August 15-18 2001, pp. 346–353.
- [4] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, December 2001.
- [5] J. Hawkinson and T. Bates, "RFC 1930: Guidelines for creation, selection, and registration of an autonomous system (AS)," <http://www.ietf.org/rfc/rfc1930.txt>, March 1996.
- [6] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proceedings of the ACM SigComm '99*, Cambridge, MA, August 31-September 3 1999, pp. 251–261.
- [7] D. Magoni and J. Pansiot, "Analysis of the autonomous system network topology," *ACM Computer Communication Review*, vol. 31, no. 3, pp. 26–37, July 2001.
- [8] J. Walrand and P. Varaiya, *High-Performance Communication Networks*, Morgan Kaufmann, 2000.
- [9] J. W. Stewart, *BGP4: Inter-Domain Routing in the Internet*, Addison-Wesley, 1999.
- [10] University of Oregon, "Route views project," <http://www.caida.org/analysis/routing/>.
- [11] S. Tauro, C. Palmer, G. Siganos, and M. Faloutsos, "A simple conceptual model for the Internet topology," in *Proceedings of IEEE Global Telecommunications Conference, GLOBECOM '01*, San Antonio, TX, November, 25-29 2001, vol. 3, pp. 1667–1671.
- [12] A. Broido and K. Claffy, "Internet topology: Connectivity of IP graphs," in *Proceedings of SPIE, ITCOM 2001*, Sonia Fahmy and Kihong Park, Eds., vol. 4526, pp. 172–187. Denver, CO, August, 19-24 2001.
- [13] A. Danesh, S. H. Rubin, M. H. Smith, and Lj. Trajkovic, "Mapping the Internet," in *Proceedings of IFSA/NAFIPS 2001*, Vancouver, BC, July 2001, pp. 687–692.
- [14] B. Cheswick, H. Burch, and S. Branigan, "Mapping and visualizing the Internet," in *Proceedings of USENIX Annual Technical Conference*, San Diego, CA, June, 18-23 2000.
- [15] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *Proceedings of the IEEE INFOCOM '00*, Tel-Aviv, Israel, March, 26-30 2000, vol. 3, pp. 1371–1380.
- [16] The Cooperative Association for Internet Data Analysis (CAIDA), "Skitter," <http://www.caida.org/tools/measurement/skitter/>.
- [17] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The impact of routing policy on Internet paths," in *Proceedings of the IEEE INFOCOM '01*, Anchorage, AL, April, 22-26 2001, vol. 2, pp. 736–742.
- [18] H. Chang, S. Jamin, and W. Willinger, "Inferring AS-level Internet topology from router-level path traces," in *Proceedings of SPIE, ITCOM 2001*, Sonia Fahmy and Kihong Park, Eds., vol. 4526, pp. 196–207. Denver, CO, August, 19-24 2001.
- [19] H. Burch and B. Cheswick, "Mapping the Internet," *Computer*, vol. 32, no. 4, pp. 97–98, April 1999.
- [20] The Cooperative Association for Internet Data Analysis (CAIDA), "Mapnet," <http://www.caida.org/tools/visualization/mapnet/>.
- [21] B. Huffaker, E. Nemeth, and K. Claffy, "Otter: A general-purpose network visualization tool," in *Proceedings of The Internet Society's Internet Global Summit, INET'99*, San Jose, CA, June, 22-25 1999.
- [22] The Cooperative Association for Internet Data Analysis (CAIDA), "Visualizing Internet topology at a macroscopic scale," http://www.caida.org/analysis/topology/as_core_network/.

- [23] R. Govindan and A. Reddy, “An analysis of Internet inter-domain topology and route stability,” in *Proceedings of the IEEE INFOCOM '97*, Kobe, Japan, April, 7-12 1997, vol. 2, pp. 850–857.
- [24] A.-L. Barabasi and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, pp. 509–512, October 1999.

A Acronyms

AS	Autonomous System.
BGP	Border Gateway Protocol.
CDF	Cumulative Distribution Function.
CCDF	Complementary CDF.
CIDR	Classless Inter-Domain Routing.
ICMP	Internet Control Message Protocol.
IS-IS	Intermediate System-to-Intermediate System protocol.
ISP	Internet Service Provider.
OSPF	Open Shortest Path First.
pdf	Probability density function.
POP	Point of Presence.
RIP	Routing Information Protocols.

Figure 7: A screen shot of the topology analysis tool.

Figure 8: A screen shot of the topology analysis tool.