

On the Geographical Properties of BGP Routing Tables

Juan I. Nieto-Hipólito José M. Barceló

Computer Architecture Department, Technical University of Catalunya, Spain

{jnieto, joseb}@ac.upc.es.

Abstract: Recently, Faloutsos et al, published an unexpected result [1]: Autonomous Systems (AS's) interconnectivity exhibit a power-law degree distribution. This result has led to the construction of scale-free network models to characterize Internet topologies. Most of these works use the tables published by Oregon Route Views [2], and frequently the distributions measured do not exactly follow a power-law. This work is different of previous ones in the sense that we use routing tables from several geographical sites, and the union of all of them having a broader vision of AS's connectivity. Besides, we compare BGP tables of different AS's sizes and make use of the CCDF (Complementary Cumulative Density Function) to better fit the power-law. This comparison is based on the AS degree connectivity, clustering coefficient and path length. Our results show that the topology of Internet at the AS level could be modeled both as a scale-free network and as a small-world network. We also show that the model of Barabasi-Albert does not fit very well for small and medium AS degrees.

Index Terms—Internet topology, BGP routing tables, Power-laws, Small-world networks.

I. INTRODUCTION

An Autonomous System (AS) is a set of routers under a single technical administration, using one or various IGP (Interior Gateway Protocol, i.e., RIP, OSPF, etc) and common metrics to route packets inside the AS and using an EGP (Exterior Gateway Protocol, i.e., BGPv4) to route packets to other AS. BGP comes from Border Gateway Protocol and is an inter-autonomous system routing protocol, [3]. BGPv4 is extensively used to connect ISP (Internet Service Providers) and to interconnect enterprises to ISP. ISP's usually are providers (provide connectivity) of other ISP that at the same time are providers of smaller ISP. In the periphery of Internet there are small ISP that usually give services to enterprises that are not ISP. ISP can be classified as transit ISP when they offer transit of traffic, multihomed ISP when they are connected to more than one ISP and do not offer transit of traffic and stub ISP when they are connected to only other ISP. An ISP can have more than one AS number assigned and give services to other ISP on large geographical areas. We will

Juan Ivan Nieto Hipólito is with UABC, Autonomous University of Baja California, México. He is persecuting his Ph. Degree in the Department of Architecture and Technology of Computers of the Technical University of Catalunya, Spain, with a grant of the Mexican government.

This work was supported by the Ministry of Education of Spain under grant CYCIT TIC-2001-SGR-00226 and by the Department of Universities, Research and Information Society (Generalitat of Catalonia) under grant CIRIT 2001-SGR-00226.

consider for simplicity that an ISP is only one Autonomous System and then study topologies based on BGP interconnectivity. This assumption is not far from reality since Internet topology is studied based on inter-domain interconnectivity and routing policy handle by BGP.

The primary function of BGP is to exchange network reachability information with other BGP peers on neighbors AS. The AS's apply different policies to the way they import/export routes when using BGP. *Routing policies* define how routing decisions are taken in Internet. If we have two networks separated by two routers belonging to different AS, policy comes into play when one AS decides to announce networks (prefix routes) to the other AS. The import policies allow an AS to transform incoming route updates. For example, denying or allowing an update, assigning a local preference attribute to an incoming route depending on AS origin, AS path, etc. The AS's only send their best route to their AS's neighbors. Export policies allow an AS to determine whether to send this route to a neighbor and, if it does, send the route with or without hints such as the community attribute or the MED (Multi Exit Discriminator). A route is expressed as a prefix, it is to say, a group of one or more networks. Routing policies are not applied to each prefix separately but to a group of prefixes defined at the AS. In fact and without exception, an AS must have only one routing policy, even if the AS is connected to more than one exit point.

BGP exchanges reachability information with other BGP peers of the same AS or of different AS's. This information will allow us to construct a graph of AS's connectivity, this is to say, the Internet topology at Autonomous Systems level. A BGP peer builds a routing table database consisting of the set of all feasible paths and the list of networks (prefix's) reachable through each feasible path or AS path. An example of a BGP table is shown in Table I, where the symbol ">" expresses the "best path". As BGP routers only send their best path to BGP peers, a BGP router will have a particular view of the Internet topology depending on where they are placed. For this reason to have a wider perspective of the Internet topology it is necessary the union of several BGP Tables and taken from different places, as we do in this work.

Based on the work of [1] and making use of the Oregon RouteViews BGP table [2], many researchers have recently investigated the Internet topology, [4], [5], [6], [8], [9], [10]. The importance of these studies resides in the development of Internet topology generators, the deployment of content

networks, the placing of web servers or in the development of inter-domain traffic engineering models.

Table 1: Example of a BGP table.

Network	Next Hop	Metric	Weigh	Path
*> 0.0.0.0	195.66.225.254		0	5459 i
* 3.0.0.0	204.42.253.253		0	267 1225 701 80 i
*	206.157.77.11	70	0	1673 701 80 i
*	12.127.0.249		0	7018 701 80 i
*>	204.70.4.89		0	3561 701 80 i
*	205.158.2.126		0	2828 701 80 i
*	158.43.206.96		0	1849 702 701 80 i

Other authors have centred their study in the novel definition of complex networks. A complex network displays certain organization principles that are encoded in their topology. These works study three main topology models: the classical Erdős-Renyi model for random graphs, the small-world model motivated by small paths between two nodes and high clustering coefficients and the scale-free model that presents power-law degree distributions. Some of these works are [11], [12] or [13].

Authors of [10] and [15] argue that tables taken from only one point such the Oregon Route Views give a poor vision of the Internet and they propose to use more than one point of vantage. Authors from CAIDA, [14], propose to generate active probes to complete the AS interconnectivity. T. Bu et al, [6], propose using the CCDF to better fit the power-law behavior described by Faloutsos et al. Finally, A. Lakhina et al, [16], describe sampling biases in IP topology measurements using traceroutes.

Here in this work we compare several BGP tables from different geographical sites and their union, using classical measures of network's topology: AS path distribution, clustering coefficient and degree distribution together with information about number of AS's and number of edges seen in the different perspectives. Instead of using the frequency degree distribution we will use the CCDF, as has recently been proposed in [6]. We will compare tables of different sizes and from AS's with different interconnectivity: core AS's (connected to a high number of AS) and leaf AS (connected to a few number of AS).

In section II we define the metrics selected to compare the BGP tables, section III is devoted to explain the methodology used, section IV shows the results gotten and section V finalizes with some conclusions.

II. METRICS

In this section we define the metrics selected to compare our sources of data, for that propose we consider the AS level topology as the graph $G=(N,E)$, where N is the number of vertices or AS's and E the number of edges or links that connects the vertices.

At each of the different BGP tables we will investigate the AS degree rank, the AS degree CCDF (Complementary Cumulative Density Function), the clustering coefficient and the AS path length. We define the adjacency matrix $\{a_{ij}\}$ where

a_{ij} is 1 if node i has an edge joining node j and 0 elsewhere.

A. Degree Rank

This metric ranges AS's in decreasing order of degree d and plots the pair degree d versus rank r in log-log scale. The degree, d_i , of node i is the number of neighbors directly connected with it, so $d_i=\sum a_{ij}$ for all j . This metric was introduced by [1] with the name of Power-Law 1 (rank exponent) and establishes that the degree, d_i , of an AS is proportional to the rank of the AS, r_i , to the power of a constant η : $d_i \propto r_i^\eta$, where the symbol \propto means proportional.

B. CCDF

Faloutsos et al, [1], introduced the metric f_d as the fraction of nodes with degree d and demonstrated that it follows a power-law of the type: $f_d \propto d^{-\xi}$, the exponent ξ is obtained using a linear regression. T. Bu et al, [6], define the empirical complementary distribution (*ecd*) as the CCDF of the degree distribution. This metric plots the fraction of nodes with degree greater than or equal to d versus the degree of the AS. In a probabilistic sense the *ecd* is defined as $F_d = Prob\{D \geq d\} = \sum f_i$ for $d \leq i < \infty$, where D is a random variable that indicates the number of incident links upon an AS, i.e., its degree d . F_d also follows a power-law with exponent α , $F_d \propto d^{-\alpha}$.

C. Clustering Coefficient

The clustering coefficient, C , is a concept used in the theory of small-world networks [17] and is a metric that indicates the grade of connectivity of any node, by definition $0 \leq C \leq 1$. The clustering coefficient C_i for any node i in the network (graph) is defined as the ratio between the number of connections among the d_i neighbors of a given node i and its maximum possible value, $d_i(d_i-1)/2$, where d_i is the degree of node i . C_i is the fraction of the edges that actually exists, and $C(G)$ is the average value of C_i .

$$C(G) = \frac{1}{N^{\geq 2}} \sum_{i \in G} C_i \quad C_i = \frac{\text{\#of edges in } G_i}{d_i(d_i-1)/2}$$

In the above equations G_i is the subgraph of node i , and is defined taking only into account the neighbors of node i . C_i is only defined for AS's with degree $d_i \geq 2$. This is because for $d_i=1$, C_i is undetermined. $N^{\geq 2}$ is the set of AS's with degree $d_i \geq 2$ and $N^=1$ is the set of AS's with $d_i=1$ and $N=N^{\geq 2} + N^=1$. The clustering coefficient provides a measure of how well are interconnected the neighbors of a node. Fully connected networks have a clustering coefficient of 1. Isolated nodes have a C_i of 0.

D. AS path-length

To measure the AS path-length, l , we will obtain the CCDF of the AS path length, F_l . This metric plots the cumulative fraction of path-lengths greater than or equal to l versus l . In a probabilistic sense F_l could be seen as $F_l = Prob\{L \geq l\} = \sum f_i$ for $l \leq i < \infty$. L is a random variable, limited by $1 \leq L \leq L_m$, where L_m is the longest AS path found in the BGP table, f_i is the fraction of paths with length l and l is the number of AS's traversed to get the target.

III. METODOLOGY

In this work, we use besides the data of [2] six public available BGP Tables [7]. In order to get the adjacency matrix, $\{a_{ij}\}$, we need to analyze BGP routing tables, which provide us with AS paths and links contained in them. It is important to note that BGP is a protocol of peering relationships and not of physical connections. For that reason the local view of AS A located in Europe could be (and is) different of AS B located in Asia and so on, that is the motivation of this paper, investigate the difference and similarities of the distinct tables analyzed. We capture 6 BGP tables, 2 located in USA, 3 in Europe and 1 in Asia.

Table 2 shows these data sources and their parameters. All BGP tables were collected over the same period of time (October/2002), and the difference of collection times between the first and the last one was of 1 week. Of these tables Oregon and Ripe-rc are Remote Route Collectors, which have a repository where the complete data can be gotten via an anonymous ftp. Ripe-rc is a set of 9 remote route collectors, 7 deployed in Europe, 1 in Japan and 1 in USA, [18]. The rest of the data sources were gotten using the CISCO or Zebra “sh ip bgp” command via telnet to the site. Swinog is a medium size AS with 41 neighbors and Exodus Comm. Europe, Exodus Comm. Asia and Opentransit are leaf AS’s with only one neighbour. See Table 2.

From Table 2, we can see also that Swinog, Ripe-rc and Oregon are numerically (around 14000 AS’s seen) very similar in spite of the difference of neighbors. The three leaf AS’s capture less AS connectivity, being Opentransit the weakest with only 9413 AS’s seen. Another interesting fact is the column AVG Degree (average degree) that is the ratio between the number of links (doubled) and nodes in the graph. For small AS’s the average degree is around 2.8 while for Swinog and the repository tables, the average degree is 4.08. We can observe that a repository table has almost the same average degree like a medium AS such as Swinog.

Table 2: BGP Tables analyzed

Source of data	Total # of AS’s	Total # of links	Max. Degree	AVG Degree	Neighbors
Exodus Comm. Eur.	13801	17311	1616	2.5086	1
Opentransit	9413	12245	1075	2.6017	1
Swinog	14018	28653	2609	4.0880	41
Exodus Comm. Asia	11264	15819	2041	2.8087	1
Oregon route view	14154	28918	2606	4.0861	61
Ripe-rc	14053	27122	2601	3.8599	170

IV. RESULTS

In this section we will investigate the AS connectivity. We analyze each of the BGP Tables and their union.

A. POWER-LAW

Figure 1 shows the degree rank in a log-log plot for the BGP Tables of Opentransit and Oregon, that is the strongest and the weakest in a numerically sense, respectively.

The parameters of the Power Law-1 for every one of the data sources are shown in Table 3. Where η is the pendent of the curve, c is a constant calculated as $d_i = c r^\eta$ when $\eta=0$ and R^2 is the correlation coefficient.

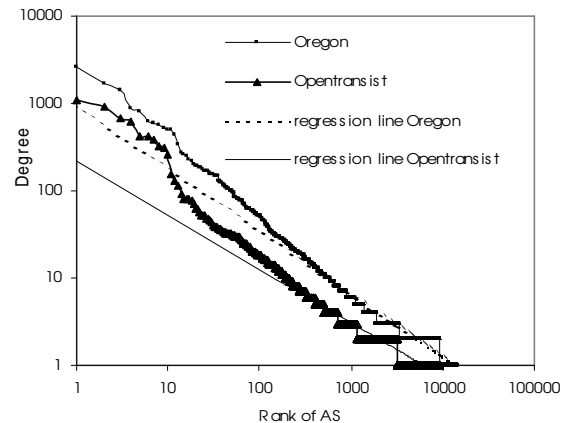


Figure 1: Power Law -1 (Rank Exponent).

We observe from figure 1, that the curve shape of Opentransit is similar to Oregon, and this behavior is followed by the other four data sources (for the sake of clarity, the other curves are not shown in the figure), that is, the most complete is the data the curve is shifted upward.

Table 3: Power Law -1.

Source of data	c	η	R^2
Oregon	965.62	-0.7226	0.9318
Opentransit	216.68	-0.616	0.9098
Swinog	1062.8	-0.7365	0.9375
Exodus Comm. Europe	233.1	-0.6023	0.8867
Exodus Comm. Asia	324.77	-0.6443	0.9251
Ripe-rc	829.46	-0.7083	0.9315

This tendency is due to the fact that the adjacency matrix, $\{a_{ij}\}$, is denser for BGP tables that contain more paths. This observation allow us to conclude that the union of all the BGP tables may offer better samples than individual ones. With respect to the correlation coefficient, R^2 , we can say that this is a low value for each table and that the regression line does not fit very well with the data except for high degree ranks. These high ranks correspond with the long tail of the data, that is, AS’s with degrees lower than 10 for the Opentransit table and lower than 40 for the Oregon table. Low degrees are more than 80% of the AS’s. Furthermore, the denser is the adjacency matrix $\{a_{ij}\}$, better fits the power-law with higher degrees. As a conclusion, we can say that the rank (and therefore f_d) is not good for fitting with a power-law due to the deviation of the low ranks (high degrees) with respect the regression line.

B. CCDF

The CCDF, F_d , seems to better fit a power-law than the rank degree or the density f_d . Figure 2 plots the CCDF of Oregon and Opentransit. Like in figure 1, the curve shape of Opentransit is similar to Oregon. This behavior is followed by the other four data sources (not shown in the figure), and the most complete is the data the curve is shifted upward.

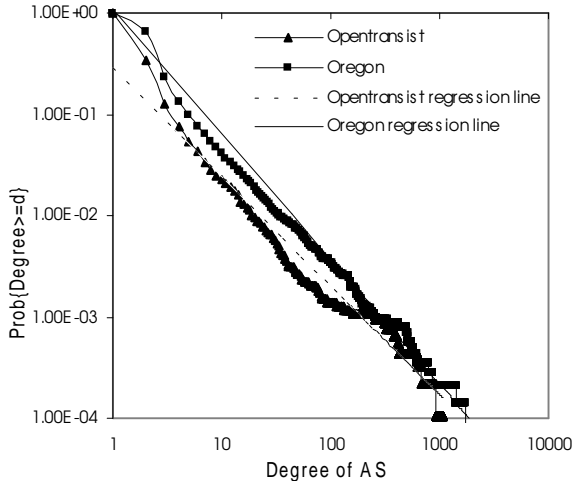


Figure 2: CCDF. $\text{Prob}\{\text{Degree} \geq d\}$.

The parameters of the CCDF are shown in Table 4. The CCDF is calculated as $F_d = c d^\alpha$ with α and c being constants and R^2 is the correlation coefficient. We notice that the correlation coefficients, R^2 , are higher than the ones gotten with power law-1 indicating a better fitting. We observe that Ripe-rc has the higher coefficient of correlation producing an exponent $\alpha = -1.197$. Again AS's with denser adjacency matrix $\{a_{ij}\}$ have better correlation coefficients, however the repository ones have no much better parameters than a medium size AS such as Swinog.

Table 4: CCDF parameters.

Source of data	c	α	R^2
Oregon	1.0988	-1.2368	0.9735
Opentransit	0.2874	-1.0759	0.9273
Swinog	1.2773	-1.261	0.971
Exodus Comm. Europe	0.2156	-1.0502	0.9583
Exodus Comm. Asia	0.3094	-1.104	0.955
Ripe-rc	0.7968	-1.1975	0.9763

C. Power-law of the Union

Figure 3 shows the Power Law-1 of the union of the six data sources. The parameters of its regression line are: $c=1735.2$, $\eta = -0.7838$ and $R^2 = 0.9402$. The most complete is the adjacency matrix $\{a_{ij}\}$ the power law fits better. That is the correlation coefficient R^2 of the union is higher than any single source of data.

D. CCDF of the Union

The last statement about the adjacency matrix is confirmed with the CCDF, aggregating all the adjacency matrices from all the data sources. Figure 4 shows the CCDF of the union of

the six data sources. The parameters of its regression line are: $c = 2.1929$, $\alpha = -1.3287$ and $R^2 = 0.9769$. The evolution of the union of the data sources is shown in Table 5. We observe an increment respect to Oregon of 16,11% on the number of links (edges) with only an increment of 0,71% of AS's. This is an important observation, because due to the nature of BGP some AS's remain hidden from others and some links do not appear due to the fact that only best paths are exported to other AS's. If we had more sources of data located in different sites we would had a very much complete picture of Internet at the AS level topology. Unfortunately the web sites where a complete BGP Table is freely available are few.

Table 5. Evolution of the union of data sources.

Source of data	Total # of AS's	Total # of links	Max. Degree	AVG Degree
Asia+swinog+opentransit=A	14056	28989	2619	4.12479
A + Exodus Europe = B	14073	29060	2620	4.12989
B + Oregon = C	14233	32441	2635	4.55856
C + Ripe-rc	14256	33578	2639	4.7107

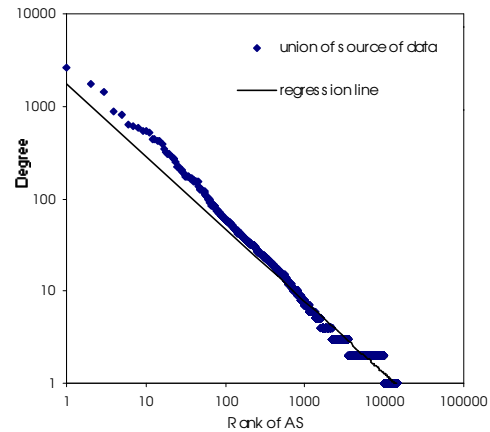


Figure 3: Power Law -1 of the six data sources

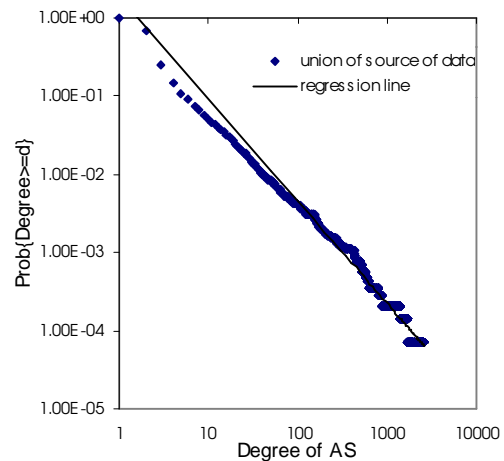


Figure 4: CCDF of the union of the six data sources

E. Clustering coefficient

Table 6 shows the clustering coefficient for each data source and the union of all ones. In this table we note that leaf AS's have the highest number of AS's with degree 1. Oregon has the lower number, 4812, of AS's with degree equal to 1. If we aggregate the tables, the adjacency matrix $\{a_{ij}\}$ is denser and the number of AS's with degree 1 decreases to 4506. That means a 6,35% less with respect to Oregon, a 11,43% less with respect to Swinog and 27,84% less with respect to Opentransit. The most diverse is the source of data the most complete is the picture of the Internet. This result coincides with the conclusions recently obtained in [16].

74,88% of the AS's have degree 1 or degree 2. These are the leaf AS's that provide services to enterprises. Normally, leaf AS's are connected to 2 ISP in order to obtain redundancy in case one ISP connection fails. Intuitively, one can imagine that the average AS path length has to be small, since one leaf AS will traverse a few number of medium or big AS to reach other leaf AS.

Table 6: Clustering coefficient and AS's with degree 1 and 2.

Source of data	N^1	N^2	N^1+N^2	C
Oregon	4812	6054	10866	0.4235
Opentransit	6245	1992	8237	0.1979
Swinog	5088	5775	10863	0.4222
Exodus Comm. Europe	9945	2285	12230	0.1792
Exodus Comm. Asia	6766	2968	9734	0.3050
Ripe-rtc	4754	6022	10776	0.3912
Union of source of data	4506	6170	10676	0.4624

If we compare Opentransit with Exodus Comm Europe, two leaf AS's, we can note that Exodus sees a higher number of AS's, 13801 (see Table 2) than Opentransit, 9413 AS's. However, the clustering coefficient is higher for Opentransit. This observation is accentuated if we compare Exodus Comm Europe (13801 AS's and $C=0.179$) with Exodus Comm Asia (11264 AS's and $C=0.305$). That means that even a leaf AS that has higher number of AS's in its BGP table sees less AS interconnectivity than a leaf AS with lower number of AS's in its BGP table. Furthermore, a medium AS such as Swinog has a clustering coefficient similar than Oregon, a repository of 61 AS's. Joining tables from different geographical locations, even if they are of different sizes, improves the interconnectivity picture.

F. AS path length.

AS path length measures the separation between two AS's. The AS path length is the number of nodes traversed to get the target AS. In practice BGP only considers the "best path", which is signaled with the symbol ">" in the BGP table, like we observe in Table 1.

In Figure 5, we show the CCDF, F_i , of the path length for the Opentransit BGP table and the Union of all the tables. We notice that the probability of having small paths is greater than the probability of having long paths. All individual tables have similar AS path length distribution. We also show in the figure the regression line of the Union of the tables. The curves fit an exponential-law: $y=3.5359e^{-0.7335x}$ and a correlation coefficient $R^2=0.9911$. This high value of R^2 confirms the validity of the

exponential-law to model the AS path length.

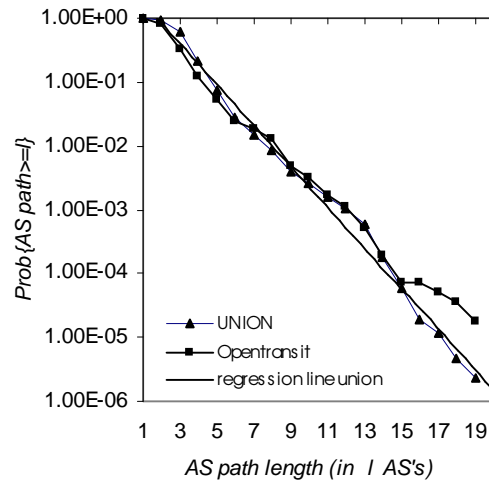


Figure 5: CCDF of the AS path length, F_i .

G. Small-world networks

Finally, following the results we can talk about the duality of the topology of Internet at the AS level, in which the Internet can be seen as a small-world network and a scale-free network [11]. A small-world network is that one in which the AS path length is small and the clustering coefficient is high (compared with a random graph). We can observe from Table 6 and Figure 5 that the AS connectivity behaves as a small world network with high clustering coefficients and small AS path lengths. However, a small world network is not an indication of some type of organization principle (scale-free networks). This property is given for the presence of degree power-laws. From Figure 2, we can observe that the CCDF behaves a power-law, indicating scale-free behaviors.

Barabasi and Albert (AB model), [11], propose a topological model that exhibits "preferential attachments" following degree distributions. This model predicts a clustering coefficient following approximately a power-law $C_i \sim d_i^{-0.75}$. We have plotted, in figure 6, the clustering coefficient for each node and the AB model. We can observe that for nodes with high degrees, larger than a five hundred neighbors, the AB model behaves well. However, for medium and small degrees, the curve does not fit very well.

In the last year, a lot of works discussing preferential models have appeared. We believe that more work in this area is needed and better models than that of Barabasi and Albert may be proposed.




Figure 6: Clustering coefficient of node i versus degree

V. CONCLUSION

In this paper, we compare BGP tables from different sites and of different sizes. We have chosen six complete BGP Tables. We have obtained the adjacency matrix $\{a_{ij}\}$ of these tables and of the union of all the tables. Since the degree rank and the degree density function do not fit very well a power-law 1, we have chosen the CCDF. We have shown that the CCDF follows a power-law $F_d \propto d^{-\alpha}$ that fits better than the degree density function and that at most complete is the data sources the regression line fits better the model. We have taken into account the union of more than a hundred AS's.

We have also compared BGP tables of different sizes. The results show that repository collectors capture more of the AS connectivity. However, medium AS's also capture most of the AS connectivity. Taking the union of the BGP tables means to see a denser adjacency matrix and therefore improve the observed AS interconnectivity. Also mention that the number of leaf AS's practically remain invariant, around a 75% of the AS's.

Our result shows that to get a most complete picture of the AS connectivity it is necessary to have access at more BGP Tables. Having more tables means an increment in the number of links although the number of AS's seen remains practically invariant.

Since the CCDF confirms a power-law degree distribution, this means that the AS's topology can be seen as a *scale-free network*. Any network to be considered as a *small-world network* must comply two conditions: a high value of clustering coefficient and a small path length. Our results show that the AS connectivity of Internet covers these considerations. Specially with respect to the clustering coefficient we can say that at most complete the data, the AS's with degree 1 tend to decrease and with this the coefficient tends to augment, that means an improvement in the connectivity of the Internet at the AS level. Furthermore, models such as the AB model do not model very well the AS connectivity. We believe that more work in this area is needed and better models than that of Barabasi and Albert may be

proposed. These models would have to take into account not only classical graph metrics but also BGP connectivity.

REFERENCES

- [1] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in Proceeding of ACM SIGCOMM'99, August 1999.
- [2] Route Views "University of Oregon Route Views Project". <http://www.antc.uoregon.edu/route-views/>.
- [3] Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)". Request For Comment – RFC 1771.
- [4] A. Broido, K. C. Claffy, "Internet expansion, refinement and churn", ETT January 2002.
- [5] G. Huston, "Analyzing the Internet's BGP Routing Table", The Internet Journal Protocol, Vol. 4, Mar 2001.
- [6] T. Bu and D. Towsley, "On Distinguishing between Internet Power Law Topology Generators". In Proceedings of INFOCOM, 2002.
- [7] "Public route server and looking glass list", <http://www.traceroute.org>.
- [8] Zihui Ge, D. R. Figueiredo, S. Jaiswal, L. Gao, "On the Hierarchical Structure of the Logical Internet Graph". In Proceeding of SPIE ITCOM 2001.
- [9] H. Chang, R. Govindan, S. Jamin, S. J. Shrenker and W. Willinger, "On Inferring AS-Level Connectivity from BGP Routing Tables". In Proceeding of INFOCOM, 2002.
- [10] L. Subramanian, S. Argawal, J. Redxford, R. H. Katz, "Charectirizing the Internet hierarchy from multiple vantage points", In Proceeding of INFOCOM 2002.
- [11] R. Albert, A.-L. Barabási, "Statistical Mechanics of Complex Networks". Reviews of Modern Physics, 74, 47 (2002).
- [12] A. Vázquez, R. Pastor-Satorras, A. Vespignani, "Large-scale topological and dynamical properties of the Internet", Physical Review E., Vol 65, June 2002.
- [13] W. Willinger, R. Govindan, S. Jamin, V. Paxson, S. Shenker, "Scaling phenomena in the Internet: critically examining critically", Proceedings of the National Academy of Sciences, USA, Vol 99, supp 1, pp 2573-2580, February 2002.
- [14] B. Huffaker, D. Plummer, D. Moore and K.C. Claffy, "Topology discovery by active probing," Symposium on Applications and the Internet (SAINT), 2002.
- [15] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker and W. Willinger, "The Origin of Power Laws in Internet Topologies Revisited", INFOCOM 2002.
- [16] A. Lakhina, J. W. Byers, M. Crovella, P. Xie, "Sampling biases in IP topology measurements", INFOCOM 2003, San Francisco
- [17] D.J. Watts, S.H. Strogatz "Collective dynamics of "small-world" networks." Nature 393, 440-442, 1998.
- [18] RIPE Routing Information Service, <http://data.ris.ripe.net>