

Analysis of the Autonomous System Network Topology

Damien Magoni
LSIIT-ULP

Bld Sébastien Brant
Illkirch, France

magoni@dpt-info.u-strasbg.fr

Jean Jacques Pansiot
LSIIT-ULP

Bld Sébastien Brant
Illkirch, France

pansiot@dpt-info.u-strasbg.fr

ABSTRACT

Mapping the Internet is a major challenge for network researchers. It is the key to building a successful modeling tool able to generate realistic graphs for use in networking simulations. In this paper we provide a detailed analysis of the inter-domain topology of the Internet. The collected data and the resulting analysis began in November 1997 and cover a period of two and a half years. We give results concerning major topology properties (nodes and edges number, average degree and distance, routing policy, etc.) and main distributions (degree, distance, etc.). We also present many results about the trees of this network. The evolution of these properties is reviewed and major trends are highlighted. We propose some empirical laws that match this current evolution. Four new power-laws concerning the number of shortest paths between node pairs and the tree size distribution are provided with their detailed validation.

Categories and Subject Descriptors

C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—*Internet*

General Terms

Measurement

Keywords

Autonomous System, AS network topology analysis, power laws

1. INTRODUCTION

For network researchers and engineers, the study of the Internet itself is fascinating. Internet, unlike the networks it is composed of, has no authority defining its topology evolution. That is why nobody can give us today a fully detailed map of the Internet. This is a main challenge because many network protocol developers would appreciate such information. Hopefully one aspect of the Internet topology is easier to capture than the others. It is the topology made by the Autonomous Systems of the Internet. These network entities are used at the inter-domain routing level of the Internet. An important repository of inter-domain routing data is available at the National Laboratory for Applied Network Research [10]. Several

studies are based on these data and we provide a follow-on analysis to existing work.

Based on six instances of BGP data (from November 1997 to May 2000) we have calculated many average properties of the AS network from distributions concerning degree, distance, number of shortest paths, trees, etc. We have found that some of them can be concisely described by power-laws. We have also examined the evolution of these average properties and we have derived some empirical laws.

Our work will provide useful information for inter-domain routing protocol developers and particularly for those involved in multicast inter-domain routing research. Indeed we study the degree distribution (important for inter-domain multicast), the distance (i.e. path length) distribution, the routing policy (to check for the efficiency of inter-domain routing), the number of distinct shortest path distribution (to measure the amount of redundancy), the bi-connectivity (a measure of reliability vs connection failure) and many properties concerning the trees (interesting for routing management).

The rest of this paper is organized as follows. Section 3 of this paper defines a few notions and terms that will be used in the rest of the paper. Section 4 is devoted to the problems of the analysis while section 5 is dedicated to the analysis of the last sample of our data set (i.e. May 2000). In section 6 we analyze the evolution of the Autonomous System (AS) network since the inter-domain observer began collecting information in November 1997. Finally, in section 7, we present four new power-laws that we found while analyzing the various distributions shown in section 5.

2. RELATED WORK

This paper deals with AS network topology, a topic already widely studied. Thus its content is not completely new and the topological properties that we study are part of a framework set by several previous studies. One of the earliest studies of the AS network topology was carried out by Govindan et al. [3]. They recovered the traces of the BGP updates of one backbone BGP router from June 1994 to June 1995 and from another route server from August 1995 to November 1995. They inferred from the traces many topological results as well as route stability results. We study some topological properties already examined in their work and we compare our measures with theirs. Another study by Faloutsos et al. [1] used BGP data recovered from a special BGP router connected to several peers, from November 1997 to December 1998. They defined three power-laws that hold for the three AS network topology instances (one every six months) built by using this BGP data. In our paper we use the *same* source of information up to May 2000.

Thus we study six instances of the AS network topology. We define five new power-laws exactly in the way Faloutsos et al. did and thus our work is an extension of their work. At the router-level, an early study was carried out by Pansiot et al. [11] by using source routing. We use roughly the same taxonomy applied to the AS network topology and we study a problem that they tackled in their work, namely routing policy overhead. Another similar study whose aim was to discover the Internet routers map was recently undertaken by Govindan et al. using an heuristic called hop-limited probes [4]. They also noticed that some of the Faloutsos et al. power-laws hold for their router-level Internet instance of 1999.

3. AS NETWORK PRESENTATION

Before we present the analysis of the collected data on the network of Autonomous Systems, we provide below a few definitions of terms that we will use throughout the paper.

BGP routing tables give potentially multiple AS paths to a set of IP prefixes [16] and all of these represent AS level connectivity. Hence we can build the graph of ASs by analyzing these AS paths: two adjacent ASs in a path have a BGP connection that we model by an edge. We can also calculate the *Routing Policy Ratio (RPR)* by using these AS paths. For a given couple of nodes, the *RPR* is equal to the advertised path length between the two divided by the shortest path length computed in the graph of ASs.

The ASs are usually classified depending on the way they manage transit traffic [14] (i.e. traffic that does not originate or terminate in the AS):

- Stub AS: has only one connection to another AS.
- Multi-homed AS: has two or more connections to other ASs but refuses to carry transit traffic.
- Transit AS: has two or more connections to other ASs and carries both local and transit traffic.

We keep these definitions and add the following ones, considering the AS network as an undirected graph:

- Cycle AS: an AS that belongs to a cycle (i.e. it is on a closed path of disjoint ASs).
- Bridge AS: an AS which is not a cycle AS and is on a path connecting 2 cycle ASs.

We then divide the ASs into two exclusive broad categories:

- In-mesh AS: an AS which is a cycle AS or a bridge AS.
- In-tree AS: an AS which is not an in-mesh AS (i.e. it belongs to a tree).

We then define the **mesh** as the set of in-mesh ASs and the **forest** as the set of in-tree ASs. All ASs in the forest can also be put into one of the next two exclusive categories:

- Branch AS: an in-tree AS of degree at least 2.




Figure 1: Different kinds of AS

- Leaf AS: an in-tree AS of degree 1 (synonym of a stub).

The cycle and bridge AS definitions are useful to define the mesh and are needed by the study of the bi-connectivity. The in-mesh AS definition is necessary to classify ASs in the mesh or in the forest. This is useful prior to the analysis of the trees (number of roots, size and depth distributions, etc.) Finally an AS can also have the following qualification(s):

- Root AS: An in-mesh AS which is the root of a tree (i.e. it is adjacent to two or more in-mesh ASs and to one or more in-tree ASs).
- Relay AS: an AS having exactly 2 connections.
- Border AS: an AS located on the diameter of the network.
- Center AS: an AS located on the radius of the network (i.e. belonging to the center of the network).

Note that some qualifications are cumulative while others are exclusive. Each in-tree AS is connected to one or more in-tree ASs or to a root AS and it can belong to only one tree. Each tree is connected to the mesh via its root AS and there can be only one root AS per tree (otherwise we would have a cycle). A root AS belongs to the mesh. A root AS is considered in-tree only in a few cases (e.g. when counting the size of a tree). Figure 1 shows the different kinds of ASs in an inter-domain level network. As a last remark we shall point out that a non-root AS in the mesh could be a root AS with a hidden tree (not seen due to BGP aggregation mechanism).

4. BGP DATA ANALYSIS

In our paper we use the same source of information as Faloutsos et al. [1]. This source is a BGP router that collects routes from the 23 peers currently active as of May 2000. The BGP data of this router is stored on a server managed by the Network Measurement and Analysis team of the National Laboratory for Applied Network Research (NLNAR). The team has set up a Network Analysis Infrastructure that provides, among other data, BGP raw routing table information [10]. The set of data now extends from November 1997 to May 2000 in our study.

Date	11/97	5/98	11/98	5/99	11/99	5/00
D. S.	10	15	18	20	21	27
E. S.	10	15	17	19	19	23
PR's	55957	57597	55867	60013	70511	84295
ASs	3025	3653	4351	5043	6214	7624
PR/AS	18.5	15.8	12.8	11.9	11.3	11.1

Table 1: BGP speakers information

4.1 Origin of the data

This BGP router, called *route-views.oregon-ix.net*, has established BGP connections with a number of BGP peers. Some of them are located in very big ASs. Its purpose is to be an observer of the AS paths currently advertised in the AS network. This observer dumps its Adj-RIBs-In routing table (thus before any processing) so we get the BGP routing information of all the sources for all the destination prefixes at the most. We say "at the most" since sometimes some IP prefixes are not advertised by some peers. This router has been given AS number 65534. This number has been chosen from among the block of AS numbers reserved for private use and not advertised on the global AS network [6]. Therefore *route-views* does not propagate any information, it only recovers all the routes given by its peers. It has been dumping the BGP routing tables every day since November 8, 1997. The dump of the day is in a file called *ASmap.date.time-stamp.txt*. We recovered these files up to May 20, 2000. Their sizes range from roughly 40 MBytes to 120 MBytes. In the subsequent analysis we only keep data files sampled every six months to avoid overfilled charts and graphs.

4.2 Relevance of the data

The first thing we checked is the validity of the data source over a chosen period of time. When the BGP router *route-views* was started, it had 10 connections with BGP peers (we also call them sources) but this quickly evolved as new BGP peers were connected to it. As of May 2000, *route-views* has 30 declared neighbors located in 27 ASs (but only 23 emitting, not counting multiple sources per AS) ! The evolution of the number of declared (D.) and emitting (E.) BGP peers of *route-views* are shown in table 1. If we look at the BGP peers list (not shown) from the beginning we can see that not only the total number increased, but also some peers canceled their connections and were replaced by others, and some peers were even disappearing and reappearing through time (probably due to connectivity problems). In fact, we have found only 8 sources that have been connected to *route-views* ever since the beginning. We call them the 8 origin sources. Despite the increase in the number of BGP peers (by a 2.3 factor), the increase in the number of ASs and prefixes is mainly due to the sole growth of the AS network during this two and a half year period. An explanation for this statement is given later on in this section.

We can see in table 1 the evolution of the number of ASs and the number of prefixes (noted PR) seen by *route-views*. A prefix is a partially masked IP network address. It is the fundamental mechanism of the classless inter-domain routing (CIDR) which enables aggregation of IP network addresses. CIDR is fully explained in [2] and has been deployed since late 1992 in the Public Internet to contain the growth rate of the routing information by enforcing hierarchical routing [17]. This enforcement needs an address assignment strategy to avoid the dissemination of routing information in the AS network (e.g. domains having non-contiguous network numbers) and to reduce the accumulation of routing information in backbone

routers [15]. An Internet Registry (IR) system has been set up to control IP allocation to help CIDR deployment [7]. Notice that some prefixes can be subsets or supersets of others. But if they are in the routing file, it means that they are both advertised and each of them counts in the total number of prefixes. We see that the number of prefixes in November 1998 is lower than the number of prefixes of the two previous AS network instances. Maybe this is due to the policy of prefix aggregation (i.e. the replacement of a group of prefixes by a common *shorter* prefix). Govindan et al. [3] reports 531 ASs and 21524 prefixes in November 1994 and 909 ASs and 31470 prefixes one year later. This gives, roughly, a 70% and 46% increase per year for ASs and prefixes respectively. By using the table 1 we infer an average increase of 44% and 18% per year since November 1997. We can clearly see the effects of the AS number allocation strategy [6] and the CIDR strategy in the dramatic reduction of the regular increase of AS numbers and prefixes since 1994. In two and half years, the number of ASs was multiplied by 2.5 while the number of prefixes was multiplied by only 1.5. This also shows the effect of the strong policy of aggregation in the prefix management. Notice that the average number of prefixes per AS has fallen from 18.5 to 11.1. This is not necessarily a good thing as explained in [6]: as many prefixes as possible should be placed within a given AS, provided all of them conform to the same routing policy.

To see the impact of the variation in the number of sources on the results since November 1997, we analyzed the data given by all sources at each chosen time instance on the one hand, and the data given by the 8 origin sources on the other hand. The deviation is given in % for some main topological properties of the AS network in figure 2. We can see that the deviation is usually very low with a few percentile points. For example, in May 2000, the number of ASs seen by the origin sources only differ by 0.6% from the number seen by the 23 sources. This means that the increase of the BGP peers (sources) number has not influenced the AS number increase. Among the 45 newly discovered ASs (0.6%), 37 are leaf ASs and 8 are non-root cycle ASs (4 of them have degree 2, 3 have degree 3 and 1 has degree 7). The biggest gap can be observed for the number of connections in the AS network (i.e. the number of edges) which is much underestimated when we consider only the 8 origin sources (8.5% difference in May 2000). This means that part of the increase in the number of connections is due to the BGP peer number increase and not only to the growth of the AS network itself. However we will see later the numbers related to the growth of the AS network and notice that this problem is alleviated by the sole increase of the AS network size. Among the 1317 new connections found (8.5%), 61 were brought by the 45 new ASs and 1256 were linking existing *in-mesh* ASs. Among the 61 edges connecting the new ASs, 37 were connecting an in-tree AS to an in-mesh AS. This means that all 37 leaf ASs were directly connected under root ASs. The other 24 edges were all connecting a new AS and an existing in-mesh AS. It is clear that these edges were connecting the 8 new cycle ASs.

To further investigate the influence of the number of sources, we have measured the % of destinations seen by any n number of sources taken from among the 8 origin sources in May 2000. (We did not take the 23 sources set because the number of combinations would have exploded.) In theory, every source or combination of sources must see all the destinations. Figure 3 shows the % of ASs seen by any combination of n sources among 8. The *max* curve is the combination that saw the highest % of destinations while the *min* curve shows the lowest. With 2 or more sources, more




Figure 2: Deviation % of origin vs all sources




Figure 3: % of destinations seen by a n -combination of sources

than 95% of the destinations seen by 8 sources are detected in the lowest combination case. The poor score of the *min* curve for 1 source is probably due to connectivity problems for this specific source. Furthermore figure 2 shows that there is only a 0.6% difference between the number of destinations seen by 8 sources and by 23 sources. The asymptote is quickly reached, and the number of sources has no impact on the number of ASs existing in reality. This doesn't mean that all real ASs are seen but probably the biggest part of them. In fact, some ASs can be hidden due to BGP AS path aggregation mechanism.

We proceeded in the same way as above for the % of connections seen. In theory, every source should almost always be the root of a shortest path tree. This means that a source does not see all connections in the network. It is by cumulating the views of many sources that we will be able to unveil most of the connections in the network. Figure 4 shows the % of connections seen by any combination of n sources from among 8. We can see that if we take any one source out we will still see at least 95.9% of the connections seen by all 8 sources. We also notice that with at least 5 sources we can observe at least 90% of the connections that we would see with 8 sources. Although the *max* curve seems asymptotic, this is clearly not the case for the *min* curve. This means that if we add more than 8 sources, the total number of connections will probably increase. In fact this is the case, since we see 8.5% more connections with 23 sources than with 8 sources. The main problem we want to solve with these figures is: how many sources shall we add to get as close as possible to the true total number of connections existing in the AS network (not taking into account dynamic problems such as link failure). Figure 4 shows us that although 8 sources are not




Figure 4: % of connections seen by a n -combination of sources




Figure 5: Influence of source position

enough, an asymptote can exist (i.e. both curves are not linear and tend to flatten a lot when the number of sources increases).

The fact that it is very difficult to see all real connections can be partly explained by what we call *the viewing problem*. Each emitting source sees a tree of AS and the merging of all these trees enables us to build the graph of AS. But as shown in figure 5, this graph can be skewed: some zones are seen as trees but they could be heavily meshed. This mostly depends on the location of the sources. The farther they are from each other and the more they are scattered, the better the view will be.

To conclude this section, we can say that the number of sources (i.e. BGP peers) connected to the BGP observer router has nearly no influence on the number of ASs detected and a moderate influence on the number of connections detected (up to 8.5% as we saw earlier). This is why, in the remaining part of this paper, we will only present results obtained by processing the data given by *all* sources at any given date (e.g. the 23 sources of May 2000, the 19 sources of November 1999, etc.).

5. AS NETWORK TODAY

In this section some major properties and distributions of the AS network are presented.

5.1 Global & average properties

Table 2 shows the main characteristics of the AS network as of May 2000. The network roughly contains two times more edges (undi-

Nb of AS	7624
Nb of connections	15234
Mean distance	3.65
Mean eccentricity	7.02
Diameter	10
Radius	5
Mean degree	4.0
Max degree	1704
Mean mesh degree	5.15
Mesh size	4825
Center size	3
Border size	8
Nb of trees	591
Mean tree size (w/ root)	5.74
Max tree size	312
Mean tree depth	1.1
Max tree depth	3
Nb of cutpoints	663
Nb of bicomponents	2810

Table 2: AS network properties in may 2000

rected) than nodes. The average path length (distance) between any two nodes is between 3 and 4. 63% of the ASs are in the mesh while the remaining 37% are in trees (roots excluded). Root ASs are only 12% of the mesh ASs but they lead to 37% of the ASs. The ten biggest ASs, with respect to the degree, are root ASs and they are nearly all interconnected (i.e. at one hop from one another). We note that six of them are sources and three are origin sources. There are few border and center ASs.

5.2 Connections

Figure 6 shows the AS distribution by degree. It is highly skewed and complies with power-law 2 (see section 7). Notice that the number of nodes of degree 2 is higher than the number of nodes of degree 1.

On one hand this can be a true anomaly which means that the power-law does not apply either to the degree 1 nodes or to the degree 2 nodes. On the other hand this can be a measurement problem which means that we under-estimate the number of leaf nodes. This under-estimation can be the consequence of the use of AS aggregation or the fact that there is no need for having leaf ASs.

97.3% of the in-tree nodes have degree 1. This implies that a vast majority of nodes of degree 2 are in the mesh (97.1%). Only 26% of the ASs have a degree of 3 or above. It is nearly the same result as the one found between 1994 and 1995 by Govindan et al. [3]. This means that an inter-domain multicast routing protocol using reduced trees (i.e. a technique abstracting relay nodes) would leverage a high efficiency from the inter-domain topology. It is also true at the router level where some architectures deploying reduced trees have already been studied [12, 5].

Figure 7 shows the ASs degree distribution by rank. It complies with power-law 1 (see section 7). As said above, the ten biggest AS are roots and 38% of all the edges end up on one of these ten nodes. As a last note we recall that, at the router-level, Faloutsos et al. [1] found that power-laws 1 and 2 held in 1995 [1] and Govindan et al. found that power-law 2 held in 1999 [4].




Figure 6: AS distribution by degree




Figure 7: AS distribution by rank

5.3 Distances between AS

This section gives some results concerning the inter-domain distances. The diameter is 10, it is the same value as the one found in 1995 by Govindan et al. [3]. This enforces empirical law 3 (see section 6). The eccentricity distribution seems to be Gaussian with a mean eccentricity of 7. The radius is 5 and only three nodes are in the center. We notice that the eccentricity distribution of the routers in 1995 [11] has the same shape (i.e. Gaussian) as the one in figure 8.

The distribution of the ASs mean distance seems more erratic. Maybe this is due to our sampling (i.e. we rounded the values to the tenth). Nodes having a mean distance around 3 are dominant. It is worth remembering that at the router-level, Pansiot et al. [11] found an average distance of 21.8 and a diameter of 31 (and a radius of 16) in 1995, while Paxson [13] found 16 and above 30 respectively in 1997.

5.4 Distinct shortest paths

We begin this section with the definition of distinct shortest paths:

Definition 1. Let u and v be two distinct vertices of a connected graph G . Two paths joining u and v are *disjoint* if they have no vertices other than u and v in common. These two paths are *distinct* if they have at least one vertex not in common.

The distribution of the *Number of distinct Shortest Paths (NSP)* is useful for evaluating the amount of redundancy edges involved in shortest paths. In a tree graph or in a complete graph, for example, any pair of nodes has one and only one shortest path, hence the *NSP* distribution is limited to one value (i.e. 100% of pairs).




Figure 8: AS distribution by eccentricity




Figure 9: AS distribution by mean distance

have 1 shortest path). In the case of a broader distribution, higher values mean that if one edge of a shortest path of a pair of node is removed, there is still a probability for another shortest path of the same length to exist for this pair. Figure 10 shows how many pairs of nodes (in %) have the same number of distinct shortest paths. We only plot the beginning of the distribution, but examining all the results show that it exhibits power-law number 5 described in section 7. Figure 11 shows the % of pairs having the same shortest path length. The given distribution looks Gaussian with an average equal to the mean distance (i.e. 3.65). This distribution is closely related to the mean distance distribution, which means that although the latter is erratic, a better scale may show it as being Gaussian. It is worth noticing that the average number of distinct shortest paths (not necessarily disjoint) is 5.21 (it was 3.59 in November 1997). This increase is a consequence of the mesh growth (see section 6. A study of the AS shortest paths has recently been carried out by Tangmunarunkit et al. [18].

5.5 Trees

We present here two distributions concerning trees partially shown in figures 12 and 13. The first is the frequency of the tree sizes. As above, we only show the beginning of the distribution, but if we examine all the results we see that it can be concisely described by the power-laws 6 and 7 detailed in section 7. The second distribution concerns the frequency of the tree depths. Notice that no tree has a real arborescence given the distribution of figure 13. 90% of trees is simply composed of leaves directly connected to their corresponding root. Less than 10% of trees has depth 2 and only a few trees have depth 3 which is the maximum depth.




Figure 10: AS pair distribution by the number of distinct shortest paths




Figure 11: AS pair distribution by shortest path length

5.6 Connectivity

The May 2000 instance of the AS network has 663 cutpoints. 39% of them have a degree comprised between 2 and 5, and 35% have a degree between 6 and 15. The biggest cutpoint is the biggest AS (with respect to degree). The AS network mesh contains only 10 cutpoints which means that most cutpoints are induced by in-tree ASs (apart from leaves). The mesh also contains 11 bicomponents (8 of size 3, 2 of size 4 and 1 of size 4803). Hence it has no bridges. As the mesh contains 4825 ASs, it is clear that the biggest part of the mesh is biconnected (it is the bicomponent with 4803 ASs). This property ensures that we have at least one backup (disjoint) path for any given path in the biggest part of the mesh.

5.7 Routing policy

As we have the AS paths in the BGP routing files, we can use them to compare the routing distance of a given (source, destination) pair to the distance given by the corresponding shortest path calculated after having built the graph. We did this for each of the 23 sources.

We noticed that the number of destination ASs is not equal to the number of ASs. We verified this in each of the six instances of AS network we studied. There is, on average for all sources, roughly 0.5 to 1.5% of the ASs that are never a destination AS in all the advertised AS paths. We looked for ASs that are never a destination for *all* the sources and we did not find any. This means that some ASs are not destinations when we consider a given source, but are destinations when we consider another source. The average % of non-destination ASs for all sources masks the fact that it is never




Figure 12: Tree distribution by size




Figure 13: Tree distribution by depth

the same group of ASs that is concerned given any one source. To conclude we can say that **no** AS is never a destination AS for all sources.

For a given source, we define any non-destination AS as an *unreachable* AS. The remaining ASs constitute valid destinations and for a given AS pair, we can compare the distance given by the best AS path (remember that we have paths from all sources: see section 4.1) with the distance given by the shortest path. The former divided by the latter gives a ratio that we call *Routing Policy Ratio* (*RPR*).

We calculated the ratio for all pairs (source, destination). These pairs represent a sample of 0.3% of the total number of pairs. The sources have a mean distance of 2.75 while all ASs have a mean distance of 3.65. This suggests that the sources are closer to the center and it means that we are probably underestimating the *RPR* by a few % points (the true *RPR* is unlikely to be inferior to the calculated one). The ratio values range from 1 (the trivial minimum) to 4.

The *RPR* distribution is calculated by taking the average of the policy routing ratio distributions of each of the 23 emitting sources. However, we excluded 5 sources which had more unreachable destination ASs than reachable destination ASs. As these sources had probably connectivity problems, taking them into account would have strongly biased the average distribution. The distribution is shown in figure 14. We can notice two important facts. First, 73.7% of the routes are equal to the shortest paths (i.e. there is no distance overhead). Second, 4.8% of the destinations are unreachable, which is high even if this value is an average. If we calculate the




Figure 14: Routing policy ratio distribution

mean of all the ratios for all the reachable destinations we obtain:

$$RPR_{5/00} = 1.087$$

This means that due to policy decisions, paths are on average 8.7% longer than their corresponding shortest paths. This ratio is quite good, and we can say that inter-domain routing is very efficient and cooperative. We remind the reader that at the router level, Pansiot et al. found a *RPR* of 32% in 1995 [11] and Tangmunarunkit et al. found in 2001 [18] that some 20% of Internet paths are inflated by more than 50%.

6. AS NETWORK EVOLUTION

This section shows the evolution of the properties of the AS network over a 30 month period, from November 1997 to May 2000.

6.1 Evolution since late 1997

Figures 15 and 16 show the values of the properties for each instance, compared with the values of the properties of the first instance having a base of 100 (e.g. if a property has value 170 in May 1999, it means that it has increased by 70% since November 1997). We can clearly distinguish two trends in the curves. A group of properties is quickly increasing at a regular rate (although the rate has not the same value for all properties), while another group seems to stay constant.

In the evolving group, for instance, the number of ASs and the number of connections grow both at near regular rates of 20% and 24% increase respectively every six months. In the stable group, we can find the diameter, the radius, the average distance and so on. We can see that the mesh slowly takes over the forest part because the size of the mesh (i.e. the number of nodes in the mesh) is increasing at a faster rate than the size of the whole graph. This means that the network gets more connected. This result was already found by Faloutsos et al. [1] at the end of 1998.

6.2 Biannual average evolution

If we calculate the % of deviation of some properties of the AS network, for each instance compared with its previous instance, we get the figure 17. We can see that some curves can be roughly averaged by a horizontal line. This suggests that the evolution rates of these properties are fairly constant.

Table 3 gives the average of the % of variations calculated for the five instances (after the first) of some major properties of the AS network. We can see that the number of degree 2 ASs is increasing by 26% every six months, nearly twice as much as the increase of the number of degree 1 ASs. We also see that the % of root ASs




Figure 15: Properties evolution since 11/97

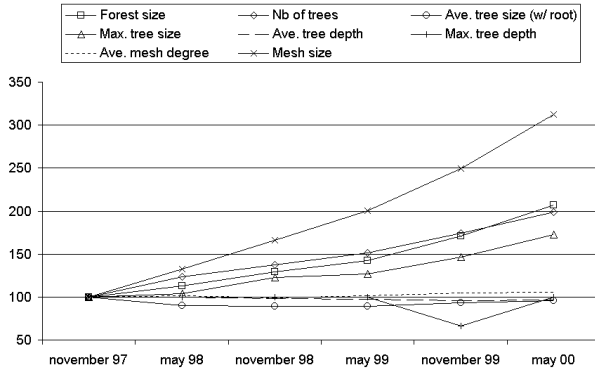


Figure 16: Properties evolution cont'd

versus the mesh ASs is decreasing every six months by 9%. This is because the size of the mesh is growing faster than the number of trees. Among properties that **regularly increase** (i.e. that have an average variation above 5%, such as the number of ASs), it is striking to see that the standard deviation of their average variation is quite low. This means that the increase rate is very regular (we can see this on the curves of figure 17).

We deduce from the previous paragraph that the growing properties are ruled by empirical laws that we infer from the table 3. In the laws description, the term *currently* means that the law has been holding for 30 months (2.5 years) and that it will probably hold in the very near future (a few years).

LAW 1. Empirical Law 1 (ASs growth) Currently, the number of ASs in the AS network increases by 45% each year.

LAW 2. Empirical Law 2 (Connections growth) Currently, the number of BGP connections in the AS network increases by 53% each year.

Notice that there is only an 8% difference between the variation ratio and the ratio of empirical law 1. This means that the majority of the ASs are added with one edge. It enforces the incremental theory of the power-laws' origins [9]. The additional edges are mostly used to increase connections of degree 1 ASs as discussed earlier.




Figure 17: Evolution of the biannual % variation of several measures

	Mean	Std deviation
Nb of AS	20.3	2.6
Nb of connections	23.8	3.3
Mean distance	-0.4	1.1
Diameter	2.8	12.3
Radius	0.7	11.6
Mean degree	2.9	1.8
Max degree	23.8	3.7
Mean mesh degree	1.1	2.1
Mesh size	25.6	3.9
Nb of trees	14.8	4.8
Mean tree size (w/ root)	-0.7	5.1
Mean tree depth	-0.7	1.0
% of roots vs mesh	-8.6	1.6
Nb of in-tree AS	13.7	5.2
Nb of in-mesh AS	25.6	3.9
Nb of AS of degree 1	14.1	5.4
Nb of AS of degree 2	26.0	4.2
Nb of AS of degree 3+	23.7	6.8

Table 3: Biannual variation % since November 1997

LAW 3. Empirical Law 3 (Distance invariants) Currently, the average distance, the diameter and the radius of the inter-domain graph of the AS network stay constant.

LAW 4. Empirical Law 4 (Trees growth) Currently, the number of trees in the AS network increases by 32% each year.

LAW 5. Empirical Law 5 (Tree invariants) Currently, the average tree size and the average tree depth of the inter-domain trees of the AS network stay constant.

Note that the average tree depth seems to be slowly decreasing with time (-3.6% in 2 years). Of course, these laws, like all empirical laws, are approximate and their life-cycle is very short. An upper bound can be easily given on the number of ASs for example. AS numbers are coded with a 16 bit integer and the IANA has reserved the block 64512-65535. A quick calculation tells us, if the empirical law 1 holds, how many years it can last as from November 1997:

$$n = \frac{\ln \frac{64511}{3025}}{\ln 1.45} = 8.23 \quad (1)$$

Date	11/97	5/98	11/98	5/99	11/99	5/00
<i>RPR</i>	1.082	1.087	1.070	1.076	1.082	1.087
<i>%NR</i>	4.53	5.16	5.80	7.98	6.98	4.83

Table 4: *RPR* evolution

The law, if valid, tells us that the AS number exhaustion will arise in approximately 8 years starting from 1997. So we can roughly announce that empirical law 1 (and most probably the other empirical laws) can be used *at most* until circa January 2006, if the AS count keeps growing at the same pace. These laws can be seen as reflecting the teenage growth of the AS network.

6.3 Routing policy evolution

Table 4 shows the evolution of the *RPR* described in section 5. The *RPR* has been quite stable since November 1997, imposing a small average overhead of 8% on the distances to travel for two and a half years.

The *%NR* line in table 4 shows the % of non reachable (NR) destination ASs. For any instance, sources that have seen more non reachable ASs than reachable ASs have been excluded from the calculation (as they most probably had connectivity problems at this time). The NR ASs % seems to remain at approximately 5% over time. The highest score was reached in May 1999 where nearly 8% of the ASs were seen as unreachable.

6.4 Connectivity

Table 5 shows the evolution of the number of bi-connectivity elements (cutpoints, bicomponents and bridges). They are given for the whole graph and for the mesh only (denoted by a small *m*), although the latter values are more significant.

In the whole graph, all branch ASs and root ASs are automatically cutpoints. Furthermore, all in-tree edges are automatically bridges. Note that for the mesh, the roots are included but all edges from the roots to in-tree nodes are removed (i.e. most roots are not cut-points anymore).

We point out that cutpoints are sensitive ASs that should not fail in order to maintain connectivity. Although the mesh contains roughly 1.5% of all the cutpoint ASs, these ASs do not really separate the mesh into equal bicomponents. The bicomponent size distribution (not shown) proves that, for each instance, there is always one bicomponent nearly as big as the mesh and a few other very small bicomponents (i.e. containing around half a dozen ASs).

Notice that two bridges appeared in the mesh of the instance of November 1999. This is a very sensitive area and the connectivity will be lost if one of these two AS connections fails (note that more than one IP-level link can be underlying an AS interconnection). A deeper investigation showed that these two bridges were linked by one on-bridge AS and that it had no real effect on the connectivity since the biggest bicomponent of the mesh had 3834 of the 3852 in-mesh ASs. This means that these bridges did not divide the mesh in two big bicomponents, but only cut off a small bicomponent from the big part of the mesh.

7. POWER-LAWS

In this section we present four power-laws that we found in our analysis of the AS network. We follow exactly the same presentation as the one used by Faloutsos et al. in [1]. Their paper

Date	11/97	5/98	11/98	5/99	11/99	5/00
<i>Cutpt</i>	359	441	475	516	582	663
<i>Bicomp</i>	1489	1610	1787	1957	2372	2810
<i>Bridge</i>	1479	1604	1779	1948	2364	2799
<i>Cutpt_m</i>	6	4	7	7	9	10
<i>Bicomp_m</i>	10	6	8	9	10	11
<i>Bridge_m</i>	0	0	0	0	2	0

Table 5: Biconnected elements

presents the three first power-laws. We continue their enumeration and therefore number our power-laws from number 4 up to number 7. After discussing how to get a power-law given a certain data set, we detail each of the four new power-laws.

7.1 Validation of a power-law

During our analysis, when we came across a distribution that could be governed by a power-law, we calculated the log of the two series of values, and we made a linear regression on this data by applying the *least squares fitting* method [19]. We determined the quality of the fit of the data by calculating the *product-moment correlation coefficient* also called Pearson's correlation [8]. The correlation coefficient varies from -1 to 1 but one usually takes its absolute value (ACC). To qualify for a line fitting, the ACC value should be at least equal to 0.95. We can consider that we have a very good fit if the ACC value is 0.98 or higher.

7.2 The pair rank exponent

We study the *Number of distinct Shortest Paths (NSP)* of each pair of vertices. The number of distinct shortest paths between two vertices is the number of shortest paths such as any of these paths have at least one vertex not in common (see the definition in section 5.4). It is not only the count of disjoint shortest paths but the count of all paths, excluding identical paths (i.e. having the same list of vertices) of course.

We sort the pairs in decreasing number of shortest paths (*NSP*), n_p , and define the pair rank r_p as being the index of the pair in the sequence. We plot the (n_p, r_p) pairs in log-log scale up to the rank of the last unique pair for a given *NSP*. The measures for the May 2000 instance are shown in figure 18. The data values are plotted with diamonds. The solid line slope and *y*-axis crossing value are given by linear regression (explained in the previous subsection). We do not show the plot of each power-law for each instance of the AS network because there would have been 24 plots to display. Instead we give in tables the exponent and ACC values of each AS network instance for each power-law.

The diamonds of figure 18 are well approximated by the linear regression with an ACC of 0.997. We infer the following power-law and associated definition.

LAW 6. Power-Law 4 (pair rank exponent) *The number of distinct shortest paths, n_p , between a pair of nodes p , is proportional to the rank of the pair, r_p , to the power of a constant, \mathcal{P} :*

$$n_p \propto r_p^{\mathcal{P}} \quad (2)$$

Definition 2. Let us sort the pairs of nodes of a graph in decreasing order of number of distinct shortest paths. We define the pair rank exponent, \mathcal{P} , as being the slope of the plot of the number of




Figure 18: Number of shortest paths *vs* pair rank (May 2000)

Date	11/97	5/98	11/98	5/99	11/99	5/00
\mathcal{P}	-0.26	-0.26	-0.27	-0.29	-0.33	-0.23
ACC	0.973	0.962	0.998	0.996	0.991	0.997

Table 6: Pair rank exponent

distinct shortest paths of the pairs versus the rank of the pairs in log-log scale.

7.3 The number of shortest paths exponent

We study the distribution of the *NSP*. We define the frequency of a *NSP*, f_n , being the number of pairs having a value of *NSP* of n (i.e. the pairs have exactly n shortest paths). We plot the (f_n, n) pairs in log-log scale up to a value of n whose total number of pairs is inferior to the number of self-pairs, being of course equal to the number of vertices (e.g. 3024 in November 1997). As above, the measures for the May 2000 instance are shown in figure 19.

The plots of figure 19 fit the linear regression line with accuracy and we infer the following power-law and associated definition.

LAW 7. Power-Law 5 (number of shortest paths exponent) The frequency, f_n , of a number of distinct shortest paths between a pair of nodes, n , is proportional to the number of distinct shortest paths to the power of a constant, \mathcal{N} :

$$f_n \propto n^{\mathcal{N}} \quad (3)$$

Definition 3. We define the number of shortest paths exponent, \mathcal{N} , as being the slope of the plot of the frequency of the number of distinct shortest paths versus the number of distinct shortest paths in log-log scale.

Notice that this power-law is strongly linked to the previous one (the pair rank exponent). Each of them represents a facet of the relationship between a pair and its *NSP*. The same remark applies for power-laws 1 and 2 found by Faloutsos et al. Both come from the relationship between a node and its outdegree. A power-law accurately represents one of the tails of a distribution, but the other tail is usually much less accurately defined. Hence the use of




Figure 19: Frequency of pairs *vs* number of distinct shortest paths

Date	11/97	5/98	11/98	5/99	11/99	5/00
\mathcal{N}	-1.95	-2.01	-1.97	-1.98	-1.96	-1.97
ACC	0.986	0.988	0.987	0.988	0.991	0.986

Table 7: Number of shortest paths exponent

another power-law, defined rather differently in order to model the behavior of the other tail.

The power laws 4 and 5 derived from the number of distinct shortest path distribution are tied to the amount of redundancy edges involved in shortest paths. We have been able to generate graphs having redundancy edges (i.e. not trees) and complying with power-laws 1 and 2 (from the degree distribution) that do not comply with power-laws 4 and 5. This means that laws 4 and 5 do not stem from laws 1 and 2 and thus are interesting indicators for characterizing an AS network map.

7.4 The tree rank exponent

We study the size of each tree, defined by the sum of the vertices composing the tree and including the root. We sort the trees in decreasing tree size, s_t , and define the tree rank r_t as being the index of the tree in the sequence. We plot the (s_t, r_t) pairs in log-log scale up to the rank of the last unique tree for a given tree size. The measures for the May 2000 instance are shown in figure 20. The solid line slope and y -axis crossing value are given by linear regression.

The plots of figure 20 match the linear regression line and, consequently, we infer the following power-law and definition.

LAW 8. Power-Law 6 (tree rank exponent) The size, s_t , of a tree t , is proportional to the rank of the tree, r_t , to the power of a constant, \mathcal{T} :

$$s_t \propto r_t^{\mathcal{T}} \quad (4)$$

Definition 4. Let us sort the trees of a graph in decreasing order of size. We define the tree rank exponent, \mathcal{T} , as being the slope of the plot of the sizes of the trees versus the rank of the trees in log-log scale.




Figure 20: Tree size *vs* tree rank

Date	11/97	5/98	11/98	5/99	11/99	5/00
\mathcal{T}	-0.77	-0.75	-0.73	-0.72	-0.74	-0.75
ACC	0.983	0.983	0.985	0.989	0.990	0.990

Table 8: Tree rank exponent

The same remark as above can be made. This power-law is closely related to the next one (the tree size exponent). Each of them represent a facet of the relationship between a tree and its size.

7.5 The tree size exponent

We study the distribution of the size of the trees. We define the frequency of a tree size, f_s , being the number of trees having a tree size of s . We plot the (f_s, s) pairs in log-log scale up to a value of s owned by only one tree. The measures for the May 2000 instance are shown in figure 21.

The plots of figure 21 fit the line quite well but the first two ACCs are not good. They are around 0.94 which is somewhat below what should be acceptable. The reasons can be multiple: the lack of information due to a reduced number of sources at the time (10 in Nov. 1997 and 15 in May 1998), the graphs not yet being big enough for this law, etc. The ACC values after May 1998 are all above 0.95 and we feel confident that a power-law governs the tree size.

LAW 9. Power-Law 7 (tree size exponent) The frequency, f_s , of a tree size, s , (including the root), is proportional to the tree size to the power of a constant, \mathcal{S} :

$$f_s \propto s^{\mathcal{S}} \quad (5)$$

Definition 5. We define the tree size exponent, \mathcal{S} , as being the slope of the plot of the frequency of the tree sizes versus the tree sizes in log-log scale.

8. CONCLUSIONS

Although we seem to model the AS level topology and its evolution with precision, we can not ensure that average values, empirical laws and power-laws, will hold with the same parameters or even at all in the middle to long-term future. A technological breakthrough could probably change the shape of Internet in a dramatic way. The




Figure 21: Frequency of trees *vs* tree size

Date	11/97	5/98	11/98	5/99	11/99	5/00
\mathcal{S}	-2.01	-2.35	-1.98	-2.46	-2.38	-2.25
ACC	0.940	0.934	0.951	0.975	0.963	0.992

Table 9: Tree size exponent

advent of switched-circuit protocols such as ATM could completely mask the physical topology to the network layer. Such changes could strongly affect the AS level topology of Internet. To resume, the goals of our study were two-fold:

- to give the network researcher a detailed view of the current AS network topology as well as a view of its on-going evolution.
- to provide a lot of information such as additional power-laws to model the AS network as accurately as possible.

We proposed new empirical laws that have seemed to be valid since November 1997. We also added a stone to the pioneering work of Faloutsos et al. by discovering four new power-laws characterizing quantities of the AS network not yet studied, as far as we know. These power-laws have been validated over a thirty month period.

9. ACKNOWLEDGMENTS

We would first like to thank the National Laboratory for Applied Network Research [10] for making available their huge collection of BGP data. The NLANR project is based on work sponsored by the National Science Foundation under Cooperative Agreement No. ANI-9807479. We also thank the anonymous reviewers for their helpful comments.

10. REFERENCES

- [1] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of ACM SIGCOMM'99*, Cambridge, Massachusetts, USA, September 1999.
- [2] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless inter-domain routing (cidr) : an address assignment and aggregation strategy. Request For Comments 1519, Internet Engineering Task Force, September 1993.

- [3] R. Govidan and A. Reddy. An analysis of internet inter-domain topology and route stability. In *Proceedings of IEEE Infocom'97*, Kobe, Japan, April 1997.
- [4] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *Proceedings of IEEE Infocom'00*, Tel Aviv, Israël, 2000.
- [5] D. Grad, J.-J. Pansiot, and S. Marc-Zwecker. Distributed computation of reduced multicast trees. In *Proceedings of TDP'96*, pages 91–107, La Londe les Maures, June 1996.
- [6] J. Hawkinson and T. Bates. Guidelines for creation, selection and registration of an autonomous system (as). Request For Comments 1930, Internet Engineering Task Force, March 1996.
- [7] K. Hubbard, M. Koster, D. Conrad, D. Karrenberg, and J. Postel. Internet registry ip allocation guidelines. Request For Comments 2050, Internet Engineering Task Force, November 1996.
- [8] J. Kenney and E. Keeping. *Mathematics of Statistics*, chapter 15, Linear Regression and Correlation, pages 252–285. Van Nostrand, 3rd edition, 1962.
- [9] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *ACM Computer Communication Review*, 2000.
- [10] National Laboratory for Applied Network Research, <http://moat.nlanr.net/Routing/rawdata/>. *BGP data from route-views*. Supported by NSF.
- [11] J.-J. Pansiot and D. Grad. On routes and multicast trees in the internet. *ACM Computer Communication Review*, 28(1):41–50, January 1998.
- [12] J.-J. Pansiot, D. Grad, and S. Marc-Zwecker. Towards a logical addressing and routing sublayer for internet multicasting. In *Proceedings of PROMS'95*, pages 521–535, Salzburg, Austria, October 1995.
- [13] V. Paxson. End-to-end routing behavior in the internet. *IEEE / ACM Transactions on Networking*, 5(5):601–615, October 1997.
- [14] Y. Rekhter and P. Gross. Application of the border gateway protocol in the internet. Request For Comments 1772, Internet Engineering Task Force, March 1995.
- [15] Y. Rekhter and T. Li. An architecture for ip address allocation with cidr. Request For Comments 1518, Internet Engineering Task Force, September 1993.
- [16] Y. Rekhter and T. Li. A border gateway protocol 4 (bgp-4). Request For Comments 1771, Internet Engineering Task Force, March 1995.
- [17] Y. Rekhter and T. Li. Implications of various address allocation policies for internet routing. Request For Comments 2008, Internet Engineering Task Force, October 1996.
- [18] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin. The impact of routing policy on internet paths. In *Proceedings of IEEE Infocom'01*, Anchorage, Alaska, USA, 2001.
- [19] E. Whittaker and G. Robinson. *The Calculus of Observations: A Treatise on Numerical Mathematics*, chapter 9, The Method of Least Squares. Dover, 4th edition, 1967.