

The Evolution of the Internet: Topology and Routing

Georgos Siganos Michalis Faloutsos Christos Faloutsos

Abstract—In this paper, we study the evolution of the Internet topology over the last three years. We study the evolution at three different levels: a) each node individually, b) the network as a whole, c) the path level. First, we find that the degrees of the nodes increase in a “rich get richer” fashion: the increase is proportional to the degree of the node. Second, we identify that new edges prefer nodes that have high degree of connectivity. Third, we observe a “small world” phenomenon: the network grows exponentially, but distances remain the same. Fourth, we find that the routing inflation increases over time.

Keywords—internet topology, power-laws, topology evolution, path inflation

I. INTRODUCTION

In this paper, we study the evolution of the Internet topology at the Autonomous System or inter-domain level over the last three years. We try to understand the dynamics of the growth of the network and its routing. Furthermore, we identify causal relationships in these properties. We relate microscopic phenomena at the node level with the macroscopic properties of the network as a whole. Understanding the Internet topology and its dynamics can have significant impact in interpreting its behavior and improving its performance.

The motivation for this work is the limited understanding of the Internet topology and its evolution. This lack of understanding is a significant factor of “Why We Don’t Know How To Simulate The Internet” [21], [9]. In practice, researchers need the topology in their efforts to a) design efficient protocols [27], [6], b) interpret measured and simulated data [8], c) detect and resolve distributed Denial of Service attacks [20]. d) calculate the routing inflation of the paths. Therefore, network analysis and management without topological knowledge is similar to trying to solve the traffic problem of a city without looking at the street layout.

There are a lot of aspects of the Internet topology evolution that we do not understand. The size and the constant change are factors that make such a study very challenging. In the last few years, some aspects of topology have been studied, but there does not seem to be a com-

prehensive evolution study. By comprehensive, we mean a study of the topology that examines the evolution at multiple levels (network, nodes, and routing) and with an effort to understand the co-evolution and co-dependencies of these levels. Over the last few years, several efforts have studied the topology in a static manner [19], [12], [5], [7]. Some recent efforts provide some analysis of the evolution, but typically focus on properties at the network level [14] [15] [11]. In the most recent effort, the authors analyze the evolution of major topological properties [16].

The most popular model for the growth mechanism of the Internet is an elegant theoretical approach based on preferential attachment of new nodes to nodes with high degrees [2]. Recent work raises doubts whether the model is verified by the empirical data [4]. Using a different and much simpler approach we conclude that even though the model doesn’t capture exactly the evolution of the growth, it can be a good approximation. Finally, there are very few studies of the routing paths and topological paths. They used one instance and examined the topology at the router level [25] and at the router level and AS level [26].

The purpose of this work is to highlight the major trends in the Internet growth and to provide novel insight on the relationships of growth-phenomena at different scales. We analyze the topology daily from November 1997 to March 2001 for a total of 916 instances. We study the evolution of three different levels: a) each node individually, b) the network as a whole, c) the routing and topological paths. Furthermore, we attempt to explain the macroscopic properties of the network through the microscopic phenomena at the node level. A secondary goal is to examine the validity of the theoretical model [2] for the Internet growth. Our main findings can be summarized in the following points:

- *The “rich-get-richer” phenomenon:* Nodes obtain new edges with rate proportional to their existing degree.
- *New nodes attachment is preferential.* We study how new nodes attach to the network. We observe that the “popularity” of an existing node among new nodes is a function of its degree, but this relation isn’t linear as the theoretical model proposes [2]. Note though that it could be a good approximation.
- *Internal edges generation is preferential.* We find that a significant part of the edges appears between nodes that existed in the topology. Again a preferentiality exists but it is not linear [1].

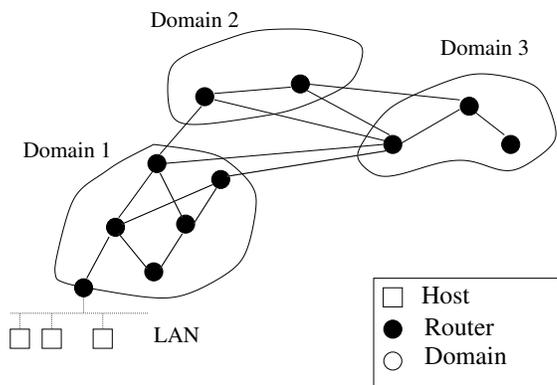


Fig. 1. The structure of Internet

- *The degree distribution remains the same.* The exponent of the power-law of the degree distribution, remains practically constant. Furthermore, the topological power-laws observed in [7] hold for all our instances.
- *Exponential growth but distances remain the same: a “small world” phenomenon.* The topology is compact and it becomes more compact over time. More than 99% of nodes are *within 6 hops* over all our instances. This means that with the same hops we reach more and more nodes given the growth.
- *Routing inflation increases with time.* The “inflation” of the routing paths compared to the topological distances at the Autonomous system level seems to increase over time. The percentage of routing paths that were inflated started from 25% for November 1997, and by March 2001 has increased to 32%.

Our work in perspective. Our goal is to identify general growth trends, find relationships between growth mechanisms, and capture the “expected” behavior of the network and its node. Naturally, the behavior of a large complex system shows some deviations. Furthermore, any real data are bound to suffer from inaccuracies and internal inconsistencies. We explain the methods we use to filter artifacts, and present the consistency of our results with that of related research efforts.

The rest of this paper is structured as follows. In section II, we present some definitions and previous work. In section III, we describe the Internet instances that we use. In section IV, we analyze the time evolution at the node level. In section V, we study the evolution of the network as a whole. In section VI, we analyze the paths and their evolution in time. In section VII we identify relationships between our observations and relate micro and macro phenomena. In section VIII, we conclude our work.

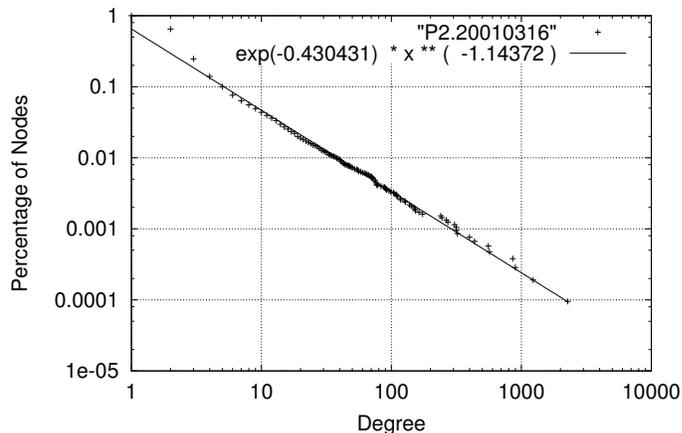


Fig. 2. RCDF: Plot of the percentage of nodes which have degree higher than a degree versus that degree.

II. BACKGROUND AND PREVIOUS WORK

We study the topology of the Autonomous System or inter-domain level. Autonomous Systems or domains are connected subnetworks that are under separate administrative authorities, as shown in Figure 1. This way, the topology of the Internet can be studied at two different levels. At the router level, each router is represented by a node. At the inter-domain level, each domain is represented by a single node and each edge is an inter-domain interconnection. The study of either level is equally important, since different protocols are employed inside a domain and between domains.

Given a graph, the **degree** of a node is defined as the number of edges incident to the node. The **rank** of a node is defined as the index in the order of decreasing degree. The node with the highest degree has rank 1, the node with the second highest has rank 2 and so on. The **distance** between two nodes is the number of edges of the shortest path between the two nodes. For the path level study, we define as **topological path** between two nodes the shortest path based on the topology. The **routing path** between two nodes is the actual path that a packet will follow. This is the path that is being advertised in the BGP routing table. Due to policies at the BGP level the routing paths may not be the topological shortest paths. We define as **inflation** of a routing path to be the difference between its length and the topological distance in hops, and **relative inflation** the ratio of its length and the topological distance.

For data fitting, we use linear regression based on the least-square errors method [22]. The accuracy of the approximation is indicated by the absolute value of the correlation coefficient, which is a number between 0 and 1. An ACC value of 1 indicates perfect linear correlation, i.e.,

the data points are exactly on a line. Typically, values over 0.97 are considered very good fits.

Related work. Several interesting static studies of the topology exist [19], [12], [5], [7]. Govindan and Reddy [11] study the growth of the inter-domain topology of the Internet between 1994 and 1995. The authors observe an increase in the connectivity over time. In recent studies, Huston [14] and Jin et al. [15] study the evolution of the network as a whole during overlapping time intervals, and identify exponential growth. For completeness and consistency, we conduct our own study.

A very elegant growth model has been proposed by Barabasi and Reka [2]. Their model grows a graph by adding nodes. The probability of a new node connecting with node i of degree d_i is proportional to its degree: $\frac{d_i}{\sum d_j}$, where $\sum d_j$ is the sum of the degrees of all current nodes. We call this model **linear preferentiality**. In a more recent work, the same authors propose a more general model that includes generation of edges between existing nodes [1]. This growth mechanism, which we call **internal edge generation**, is also suggested to follow the linear preferentiality.

Regarding paths and routing, Tangmunarunkit et al. [25], [26] examine how BGP policy makes the routing paths longer than the shortest paths. They examine 2 different data instances and report that 80% of the paths at the router level are inflated by at least one router hop. They also check the inflation of one instance at the AS level and they find that 5% of the paths are inflated by at least one AS hop.

Power-Laws of the Internet topology. Faloutsos et al. [7] introduced the use of power-laws to describe the Internet topology¹. Power-laws seem to describe several topological properties such as the degree distribution. The exponents of these power-laws can characterize concisely the topology, and they have already been used in assessing the realism of graph generators [2] [17] [24] [15] [28]. Here we study the **degree exponent**². We use the Reverse Cumulative Distribution Function or RCDF of a degree. RCDF is the percentage of nodes that have degree greater or equal to a given degree. We plot the RCDF versus the degree in log-log scale in Figure 2. The fit is spectacular with a correlation coefficient of 0.996.

¹Power-laws are expressions of the form $y \propto x^a$, where a is a constant, x and y are the measures of interest, and \propto stands for “proportional to”.

²The law we present here is slightly different than the one presented in [7], which uses the probability distribution function (PDF) of the same distribution. The two power-laws are equivalent, but the exponent differs by one, given the integral-derivative relationship of PDF and RCDF.

III. DATA INSTANCES

In our study, we wanted to use a data set that has an extensive spatial coverage and long time span. To the best of our knowledge, these criteria are best met by the data repository of the National Laboratory for Applied Network Research [10]. Furthermore, the main thrust of our analysis is the evolution, therefore, the minimal requirements for the data are the following:

- the data should be a representative subset of the total
- the measurements should monitor the evolution in a consistent way

Therefore, we do not argue that our data is the complete AS map topology [3]. But, we argue that the data meets the above two requirements: it is representative and consistent.

We present the origin of the data we use in our experiment [10]. The data is the union of a number of real routing tables used in cooperating routers. The data is collected by a route server at Oregon Route Views Project [18], from BGP³ routing tables of multiple geographically distributed routers with BGP connections to the server. For 97% of the instances, there are 15 or more routers contributing to the graph.

We examine the inter-domain topology of the Internet from 8th of November 1997 till 16th of March 2001. We represent the topology of the inter-domain by an undirected graph. We filter the initial data to remove incomplete data files that they do not represent correctly the topology. For example, on the 29th of August 1999, the size changed from 5627 to 103 nodes and became 5633 nodes the next day.

We highlight our arguments for the representativeness and the consistency of the data in our study. Naturally, such a task is ill-defined and open-ended and any such time evolution study is bound to suffer. Whenever possible we compare our work with similar research efforts.

Data representativeness. Our graph instances seem to have the same qualitative properties as graphs obtained by other sources and tools. Tangmunarunkit et al. [25] use a graph that they obtain from translating a router-level graph to the corresponding AS graph. They find that their AS graph and the NLNR graphs are qualitatively the same.

Data consistency. We want to verify that the collected data is a consistent representation of the topology. For this reason, we conducted two studies with different numbers of contributing routers: one using all the reported routers each time, and one using only nine routers that appeared consistently over a two-year period. *The results of both*

³BGP stands for the Border Gateway Protocol [23], and is the inter-domain routing protocol.

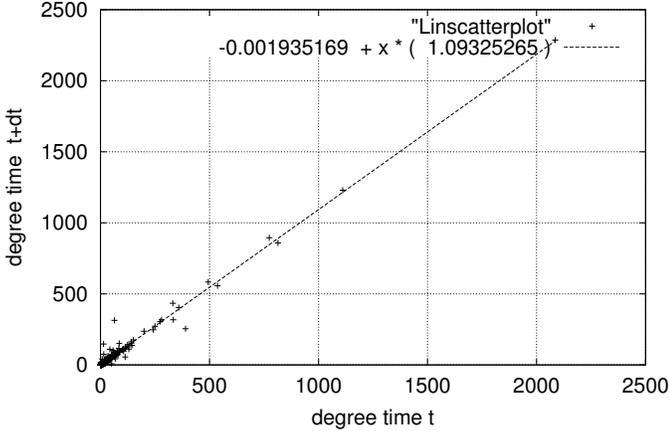


Fig. 3. New degree versus old degree, in linear linear plot

analysis are consistent. Given the agreement of the two studies, here we present the study using all the routers, which spans a larger interval. This strengthens greatly our confidence in the consistency of the collected data. Furthermore, the authors in [16] conducted some similar experiments using only a subset of routers for their analysis and their observations are in agreement with ours.

As we already mentioned, we do not argue that the data is complete. The graph includes most of the AS, but we can not quantify the percentage of the edges that are accounted for. More specifically, our graph captures the vast majority of the Autonomous Systems. Intuitively, this is to be expected to ensure the connectivity of the network. However, for the edges, things are not as straightforward. In appendix A, we study the increase in the number of observed edges as we add more routers in the observation group.

IV. THE EVOLUTION AT THE NODE LEVEL

In this section, we focus on the evolution at the node level. We examine how the degree of a node changes over time, and observe a rich-get-richer phenomenon. Then, we try to find the causes of this phenomenon. We examine how new nodes attach to the network, and how new internal edges appear between existing nodes. We find that our empirical data does not agree with the theoretical growth model. According to the model, both these processes obey the linear preferentiality. We find that they do not, although linear preferentiality could be considered as a rough approximation.

A. Degree Evolution

Observation 1: The degree increase of a node is proportional to its degree (rich-get-richer phenomenon).

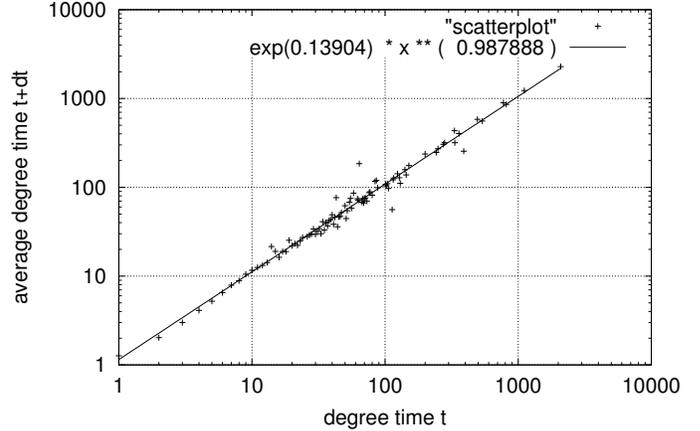


Fig. 4. Average new degree versus old degree, in log-log plot

We examine the degree evolution of the nodes in time. Given a time interval, we plot the final degree of the node versus its initial degree. In Figure 3, we show the degree evolution as above for a three month period, from August 3rd till September 1st of 2000. We approximate the plot using linear regression and the correlation coefficient is 0.99. The slope of the plot $c_{90} = 1.093$ characterizes how nodes change degree. More specifically, the degree of a node, d , in two instances in time t , and $t + \Delta t$ is given by:

$$d(t + \Delta t) = c_{\Delta t} d(t) \quad (1)$$

In other words, the increase of the degree is proportional to the existing degree of a node. High-degree nodes increase their connectivity faster than low-degree nodes. We find that this equation holds for all time intervals we tested, all the slopes are within 1.07 – 1.10 with an average of 1.085. We used also one and six month intervals and found consistent results.

Intrigued by the simplicity of this rule, we wanted to verify this property in an alternative way. We find the average final degree of all nodes with the same initial degree. In Figure 4, we plot this average final degree versus the original degree in log-log scale⁴. The plot has an excellent linear fit with a correlation coefficient of 0.99. Note that the slope is one in log-log space, which indicates that the x and y quantities have a linear relationship. This verifies the first observation.

We can go even further and validate the accuracy of the parameter $c_{\Delta t}$. We find that the estimation of parameter $c_{\Delta t}$ with both methods is sufficiently consistent. In

⁴Note that we plotted both this graph and the previous in both log-log and linear-linear scales with similar results. In general, in log-log scale, points with high values affect less the linear approximation. Due to space limitations, we decided to show one graph for each case.

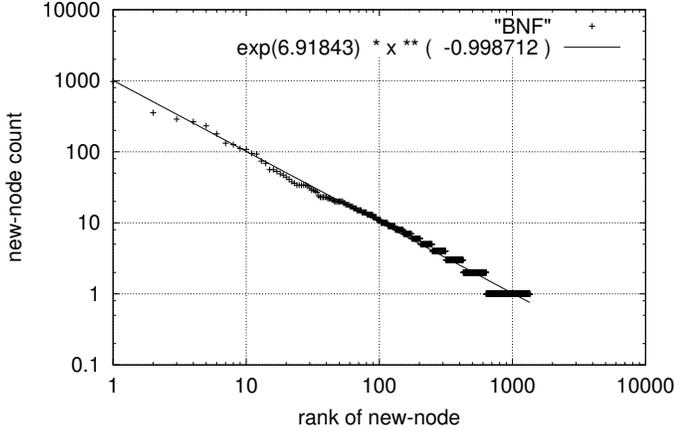


Fig. 5. New-node count of every node in order of non-increasing count. New-node count is the number of new nodes connecting to a node

more detail, we can relate the two approaches starting from equation 1. We can take the average degree \bar{d}_t over each set of nodes with equal degree, and then take the logarithm of the resulting equation:

$$\bar{d}(t+\Delta t) = c_{\Delta t} \bar{d}(t), \Rightarrow \log d(t+\Delta t) = \log c_{\Delta t} + \log d(t)$$

The slope in the linear plot must be equal with the intercept ($x=0$) in the logarithmic scale. The slope in the linear plot is equal to $c_{90} = 1.093$, and the estimate using the intercept is equal to 1.14, which gives us less than 5% error.

In conclusion, the degree increase can be characterized by a “rich-get-richer” phenomenon. Naturally, there are nodes that deviate from this rule, but this rule captures the average and expected behavior of the nodes with very high accuracy.

B. New Edges and Preferentiality

We want to study how new edges appear in the graph. The edges of the network are altered with one of the following processes:

- added edges between new nodes and existing nodes
- removed edges between dead nodes and existing nodes
- added edges between existing nodes
- removed edges between existing nodes

Studying the evolution of edges is more challenging than the evolution of nodes. Intuitively, this is because a routing table is more likely to be aware of all possible ASs than it is to know all possible ways to reach an AS. This makes the task of finding all possible edges between existing ASs complicated. A positive note here is that finding the edges between one-degree nodes is not as difficult,

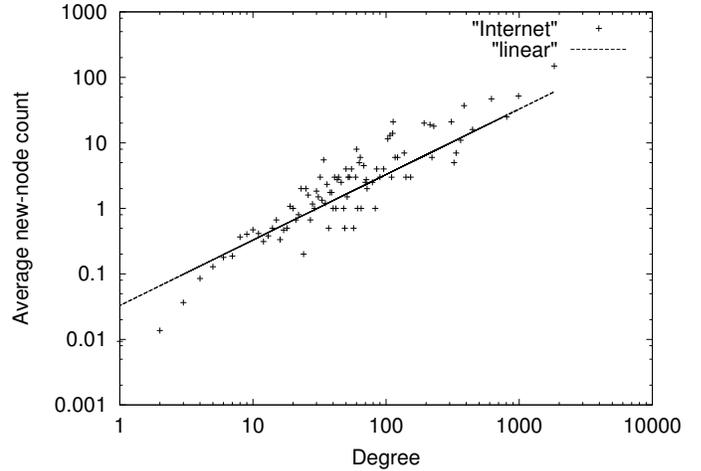


Fig. 6. Average new-node count of nodes with the same initial degree versus this degree

since a routing table will have to have at least one edge for each destination. New nodes or dying nodes have typically low number of edges and therefore we believe that we can capture their edges with more accuracy than the edges between existing nodes. Note that our study with nine only routers gave qualitatively similar results as the study with all the routers, and therefore we present only the latter.

Observation 2: New nodes prefer higher degree nodes but not with linear preferentiality. First, we want to establish the existence of preferentiality. For every node in the graph, we count how many new nodes have connected to it, which we call **new-node count**. We do this for the duration of the three years. In Figure 5 we plot the new-node count for each node in the order of non-increasing new-node count. For example, point (1,1000) means that the first most popular node connected with 1000 new nodes. It is interesting to observe that we can approximate the plot using linear regression with a correlation coefficient of 0.983. The distribution of the frequency is skewed. This preferentiality is one of the factors that cause the “rich get richer” phenomenon we saw in the previous section. Note that in this plot, each node changes many degrees, so we can not easily relate its popularity with its degree. We describe our effort to do this in the next paragraph.

Node attachment does not follow the linear preferentiality. It is not easy to relate the preferentiality of the new nodes, and the degree of the existing nodes, and for this we resort to an approximate method. In a nutshell, the difficulty lies in that we can not measure something without altering its initial state. More specifically, as new nodes join, they change the degree of the existing nodes. Thus, on the one hand, we need to get enough new nodes for sta-

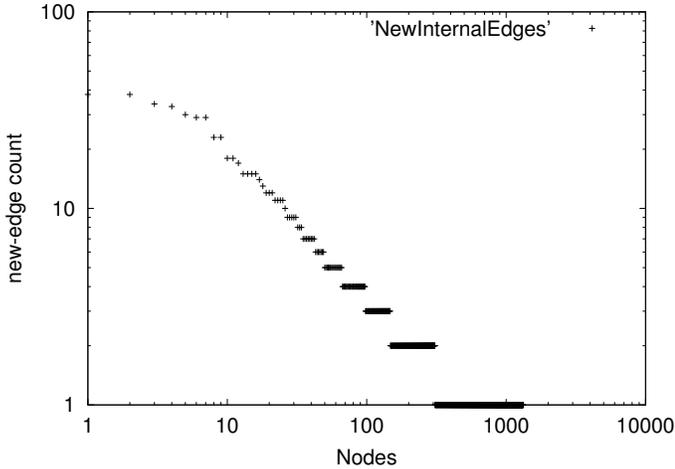


Fig. 7. New-edge count of every node in order of non-increasing count. New-edge count is the number of new edges between existing nodes.

tistical purposes, on the other hand, we do not want a large degree change. We pick a time instance and we record the degree of every node in the graph, and we assume that the degree stays constant during this interval. Then, we find the new-node count for each node, and relate the new-node count with the initial degree of the node. We choose to use an interval of three months, in which we have an increase of approximately 10%. We also tried intervals of 1 month and 6 months and we found similar results.

In Figure 6, we plot the average new-node count of nodes of the same initial degree versus that degree. We plot both the result from the real data and the expected result from the linear preferential model [2]. Although there exists a preferentiality based on the degree, the data does not agree with the linear model. It seems like nodes with higher degree are preferred more, and nodes with lower degree are preferred less, than what the model suggests.

Note that nodes that disappear from the network also follow a similar preferentiality. More dead nodes are adjacent to high degree nodes than that of lower degree nodes. The relationship again is not linear. The plots are omitted for brevity.

Observation 3: New internal edges prefer higher degree nodes but not with linear preferentiality.

We study how new edges appear between existing nodes in the graph. We use a similar approach to the one we used to assess the preferential connectivity of the new nodes. We pick two instances of the domain topology. We remove from these two instances the nodes that are not common in both of them. In this way we have removed the new and dead nodes from the graphs and now both of the graphs have the same nodes. We compare the neighbors of a node in both times. We define as new edges the edges that exist

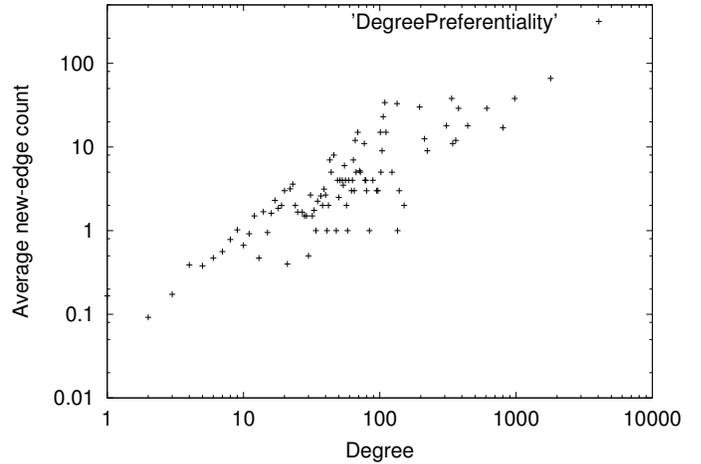


Fig. 8. Average new-edge count of nodes with the same initial degree versus this degree.

only in the second instance, and we refer to their number by **new-edge count**. We also define as dead edges, the edges that exist only in the first instance, and we denote their number as **dead-edge count**. For our analysis we used a time interval of 40 days, from August 3rd till September 12th of 2000.

First, we find that there exist some preferentiality in the way new edges are added between nodes. In figure 7, we plot the new-edge count of every node that obtained at least one edge in order of non-increasing count. The distribution is not uniform and some nodes obtain more edges. We want to investigate whether the new-edge follow a preferentiality according to the degree. We use the same method we employed for the node attachment preferentiality. For the given interval, we consider that the degree of a node does not change significantly. In figure 8, we plot the average new-edge count of all nodes with a given initial degree. We can see that the preferentiality is a function of the degree, but it is not a linear relationship.

We observed similar plots also for the edges that disappear between existing nodes. There exists preferentiality but the preferentiality is not linear.

C. Degree of New and Dead Nodes

In this section we study the dynamics of the nodes that appear or disappear from the network. We are interested in the degree with which they appear and disappear. We make the following definitions. We define as **birth** the first time we see a node. We define as **death** the last time we see a node. Naturally, the first and last day are handled appropriately.

We are interested on the number of births and deaths we have per day. In order to consider one day in our study, the

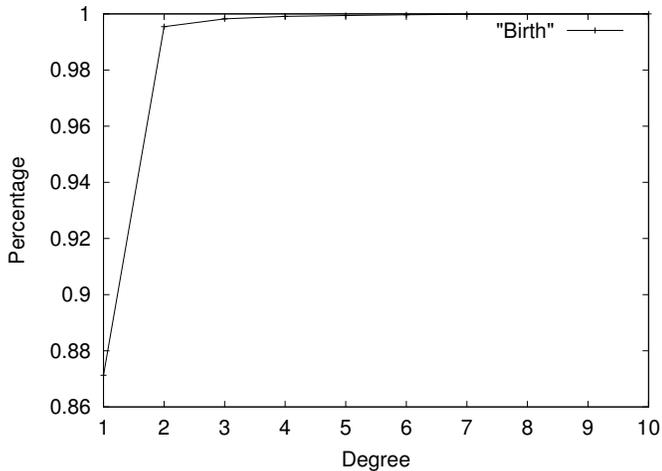


Fig. 9. Percentage of new nodes that have degree higher than a degree versus that degree.

exact previous day must exist. If we didn't force that rule there would be days where instead of measuring one day we would measure multiple days since it's not uncommon to have gaps between measurement days.

Observation 4: The distribution of the degree of new and dead nodes is skewed. In Figure 9, we plot the Cumulative distribution of the degrees of the nodes when they first appear in the Internet. The distribution is skewed and we can observe that the majority of the nodes connect to the network with degree 1. The percentage of nodes that connect with degree 1 is 87% and for degree 2 the percentage is 12%, which leaves only 1% for all the other degrees. The distribution of the final degree of the nodes when they die is similar to the one of the new nodes. The majority of the nodes that die have degree 1, with a percentage of 81% and 15% for degree 2, and 4% for the rest of the degrees. From these results we can see that most of the changes happens on the edge of the network.

V. EVOLUTION OF THE NETWORK AS A WHOLE

In this section, we study the time evolution of the properties of the inter-domain topology as a whole. First, we observe that the number of nodes and edges grow super-linearly. We find that the growth could be approximated best by an exponential function, but we also find that cubic or quadratic approximation can also describe the growth very well. Then, we observe that the degree power-law [7] holds for every instance, and examine the evolution of its exponent.

Observation 5: The number of nodes grows exponentially. Specifically, we find that the number of nodes doubles every two years approximately. In Figure 10, we plot the number of nodes versus the time in log-linear scale.

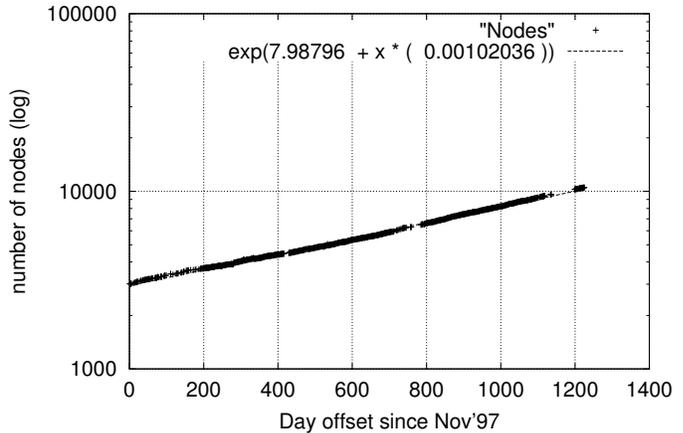


Fig. 10. The time evolution of Number of Nodes

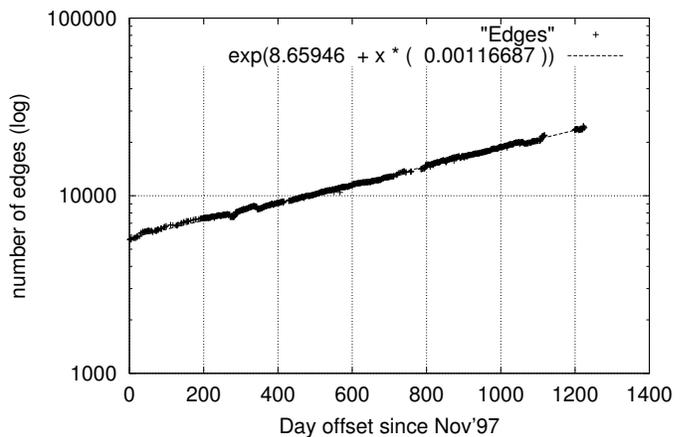


Fig. 11. The time evolution of Number of Edges

The number of nodes $N(t)$ seem to be exponential with time t :

$$N(t) = K_n e^{at}, \quad K_n, a \in \mathfrak{R} \quad (2)$$

The values of the constants are $K_n = 2874$ and $a = 10.2 \cdot 10^{-4}$. For the calculation, we use linear regression and the correlation coefficient is 0.998. We can approximate the growth with a cubic or a quadratic function with excellent correlation coefficient, 99.5% and 99.6%, respectively. We prefer to go with the exponential model, since the correlation coefficient is slightly better.

The number of nodes doubles every two years approximately. It is easy to calculate the amount of days x in which the number of nodes doubles: $N(t+x) = 2N(t)$. Using equation 2, we can solve for x and we get: $x = 680$ days, which is approximately two years.

Observation 6: The number of Edges grows exponentially. In Figure 11, we plot the number of edges ver-

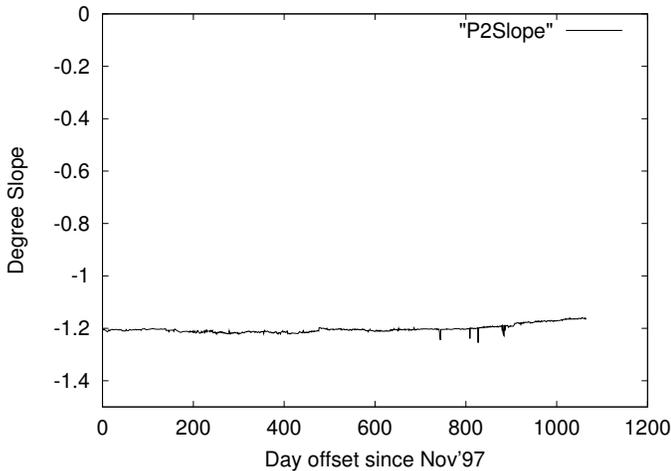


Fig. 12. Time evolution of the degree exponent.

sus the time in log-linear scale. The number of Edges $E(t)$ seem to be exponential with time t :

$$E(t) = K_e e^{bt}, \quad K_e, b \in \mathfrak{R} \quad (3)$$

The values of the constants are $K_e = 5614$ and $b = 11.494 \cdot 10^{-4}$. For the calculation, we use linear regression in the log-linear plot, and the correlation coefficient is 0.998. Again here we find that the growth can be described with a cubic or quadratic function, but with a slightly smaller correlation coefficient.

Note here that the exponent of edges is larger than the exponent of the nodes. This suggests that the number of edges grows faster than the number of nodes, which agrees with the observed increase of the average degree.

A. Persistence of Power-Law exponents in Time

The power-laws reported in [7] hold for all the instances from November 1997 till March 2001. Due to space limitations, we examine the evolution of the slope of one of the power-laws, the degree power-law. First, the striking observation is that the power-law holds for every instance. Second, the slope of the power-law seems to remain practically constant.

Observation 7: The degree exponent remains practically constant. We study the slope of the degree exponent and its evolution in time. In Figure 12, we plot the degree exponent versus time. The correlation coefficient for every instance is always higher than 0.99. We observe from the graph that the slope is always between $-1.2 \pm 3\%$. Therefore, we claim that the slope has remained approximately constant during these three years⁵.

⁵Actually, the last few months the slope seems to have an increasing

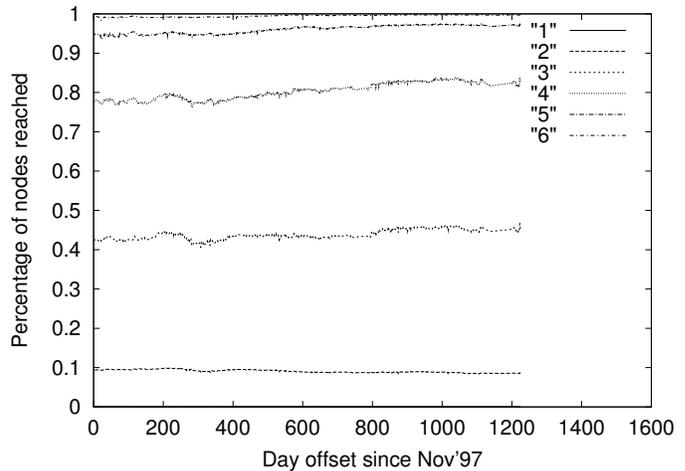


Fig. 13. Percentage of nodes reached versus time for the topological paths. Each line represents the percentage for different number of hops.

VI. THE EVOLUTION OF PATHS

We study the evolution of the paths and distances of the network. For the topological distances, we use the topology as we described before. For the routing paths, we use the unprocessed BGP tables of the routers from Oregon. Our data is over 107 Gbytes and contain over one billion paths which correspond to 411,766 unique source-destination pairs over a period of more than three years. These paths are the actual paths that a packet at the reporting router would follow towards a given destination, at some point in time. Each BGP table has entries which specify a path for a specific IP range of the destination address. This means that there are several paths for the same end-points and a packet follows a path according to the destination IP. This is a manifestation of traffic engineering and multihoming [13].

In a nutshell, we find that the topological distances remain the same, while the routing paths become longer.

A. Topological Paths

Observation 8: The distribution of the topological distances has remained practically the same. We find that the distances in the network do not change in the time period we examine. This is somewhat surprising, if we think that the network size increases exponentially.

In Figure 13, we plot the percentage of nodes we can reach for a given number of hops versus the day that each instance was collected. Each line corresponds to a different number of hops. We see that the neighborhood of a

tendency, but the increase and the duration of this increase is too small to distinguish between an actual long term trend from a temporary phenomenon or measurement noise.

node is roughly a constant percentage of the total nodes. For example, we find that 99% of the pairs are within six hops, and 45% of the pairs are within three hops. Given that the network increases, it is clear that the size of the neighborhood in absolute size (not as a percentage of the total nodes) is increasing. In other words, the network is dense and it becomes denser over time.

We also examine the evolution of distances between specific pairs of nodes. In the previous analysis, we compare the distribution of distances from different snapshots with possible different nodes. We examine the evolution of the distances of specific node pairs to factor out the effect of new and dead nodes. We conclude that the distances still remain the same. Among all possible pairs of nodes⁶, we choose the 411,766 pairs of nodes that appear as end points of routing paths, and report the change between the first and last time they appear in a BGP table. We found that on average the lengths remained the same. In more detail we found that 63.4% of the paths have the same length, 16.8% are longer and 19.6% are shorter. The change in size, both for the increase and the decrease, follow similar distributions and were comparable.

The significance of the observation is that the Internet follows a “small world” structure: the distances do not increase with size. In contrast, if the topology was a two dimensional grid, as it was often modeled before, the distances would scale roughly according to the square root of the size. This is a qualitative shift in the way we view the Internet topology and its evolution.

B. Routing Paths

Observation 9: Routing inflation becomes worse with time. First, we find that the evolution of inflation is fairly smooth and stable. In Figure 14, we plot the percentage of the routing paths for a given inflation versus time. Each line corresponds to a different inflation. We show only the inflation from 0 to 4 hops. The percentage of paths that have inflation more than 4 hops is marginal with a percentage of less than 0.04%.

Given the relative stability of the inflation, we can focus on the first and last instance to quantify the evolution of inflation. In Figure 15, we plot the cumulative distribution of the relative inflation for the first and last instance. We see that the line of the last instance is always below the line of the first instance. The percentage of paths that are inflated increased from 25% to 32%. Also the percentage of paths with relative inflation more than 1.5, increased from 2.5% to 5%.

⁶The possible pair of nodes is close to 50 million, so it is apparent that we need to sample that.

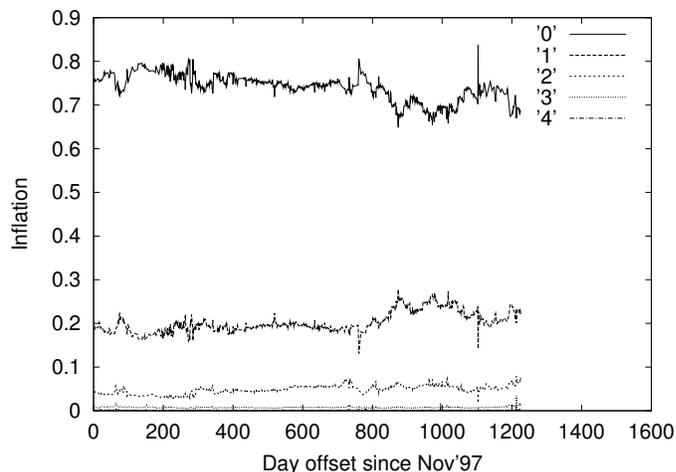


Fig. 14. Time evolution of the inflation. Each line corresponds to a different inflation, for example line 0 corresponds to zero inflation etc.

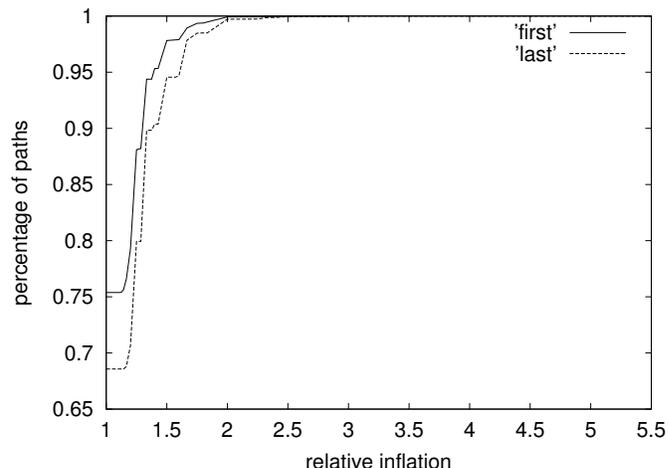


Fig. 15. Cumulative Distribution Function of the relative Inflation for the first and last instance.

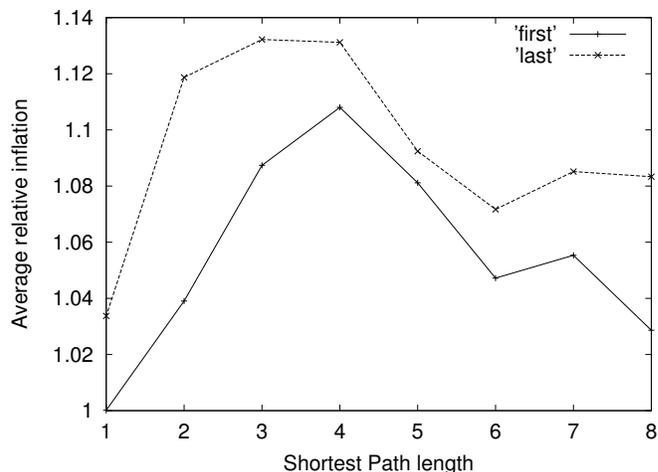


Fig. 16. Average relative inflation versus length of the topological path.

We want to examine which paths are more inflated. We group routing paths according to the distance of their end-points and calculate the average relative inflation. In Figure 16, we plot the average relative inflation versus the length of the equivalent shortest path. We can see that the worse inflation happens for the medium length paths of lengths 3 and 4. In this Figure we see again that inflation over time became worse for all categories of paths.

The inflation at the AS level seems less than at the router level. We want to compare the inflation at the AS level with the routing inflation at the router level as reported in [25]. They find that 80% of the paths are inflated on the router level, while for the same time period we found that only 30% are inflated on the AS level. In a nutshell, the router level seems to suffer from more inflation than the AS level. However, this comparison should be considered only as an indication, since the graphs and the sets of paths we compare are different.

VII. DISCUSSION

Our observations provide novel insight into the dynamics of the Internet growth. A first step in this direction is to identify dependencies and causal relationships between these observations. Here, we show how the phenomena at the node level can explain partly the evolution properties at the network level.

Rich-get-richer phenomenon leads to exponential edge evolution. The rich-get-richer phenomenon is a fundamental property of the network growth. We can show that this property “leads” to the exponential growth of the edges. In other words, we want to link the two behaviors one at the node level and one at the network level: if node degrees increase in a rich-get-richer fashion, then the total number of edges increases super-linearly and in our case exponentially. More specifically, we can prove the following lemma.

Lemma 1: If the degree increase of the nodes of a graph is proportional to their degree $d(t)$: $d(t + \Delta t) = c_{\Delta t} d(t)$, with $c_{\Delta t} > 1$ then we can prove that we have

1. exponential edge growth: $E(t) = K_e e^{bt}$
2. and $c_{\Delta t} = e^{b \Delta t}$

PROOF. *First clause.* Let us consider a time interval $t, t + \Delta t$. We take the sum over all nodes.

$$\sum_{i \in \text{nodes}(t)} d_i(t + \Delta t) = c_{\Delta t} \sum_{i \in \text{nodes}(t)} d_i(t) \quad (4)$$

We have that:

$$\sum_{i \in \text{nodes}(t)} d_i(t) = 2E(t)$$

Observe that the left side of equation 4 corresponds to the number of Edges in time $t + \Delta t$, except the edges from nodes that attach after time t , which we define as $E_{NewNodes}$.

$$\sum_{i \in \text{nodes}(t)} d_i(t + \Delta t) = 2E(t + \Delta t) - 2E_{NewNodes}$$

If time period Δt is small, we can assume that $E_{NewNodes} \simeq 0$. This way, we get:

$$c_{\Delta t} = \frac{E(t + \Delta t)}{E(t)} \quad (5)$$

Let us assume that $\Delta t = 1$, and $E(t_0)$ edges at time t_0 . By using induction, we can prove that the edges grow exponentially in time as follows

$$E(t + \Delta t) = c_1^t E(t_0) \quad (6)$$

Thus, we have exponential growth with $K_e = E(t_0)$, and $b = \log_e c_1$.

Second clause. Let us consider an interval $(t, t + \Delta t)$. We have already shown that the edges grow exponentially. Therefore, for the given interval we get:

$$\frac{E(t + \Delta t)}{E(t)} = e^{b \Delta t} \quad (7)$$

Comparing equations 5 and 7 we get the following relationship for b and $c_{\Delta t}$: $c_{\Delta t} = e^{b * \Delta t}$.

Experimental Verification. The measured data is in accordance to the theoretical expectations of lemma 1. Let us consider an interval of approximately three months: $\Delta t = 90$. From section IV-A, we substitute the measured values in the left hand side of the equation: $c_{90} = 1.09$ and $b = 0.00116687$. We find that the estimated value of $c_{90} = e^{b * 90} = 1.11$ which is within 2% of the measured c_{90} .

Rich-get-richer phenomenon helps preserves the degree distribution. The rich-get-richer phenomenon explains partly the persistence of the degree distribution and its power-law. If we take the logarithm in equation 1, we can get:

$$\log d(t + \Delta t) - \log d(t) = \log c_{\Delta t}$$

This means that for every node, the difference of the logarithms of the degree change is the same, since $c_{\Delta t}$ is independent of the node.

We can observe this graphically. In Figure 17, we plot the degree of each node in the order of non-increasing degree⁷. We plot the degrees of two instances in time: the

⁷This corresponds to the rank power-law of [7], and it is an alternative way to see the degree distribution. Actually, the degree power-law and this power-law are equivalent. We prefer this power-law because it depicts more clearly the point.

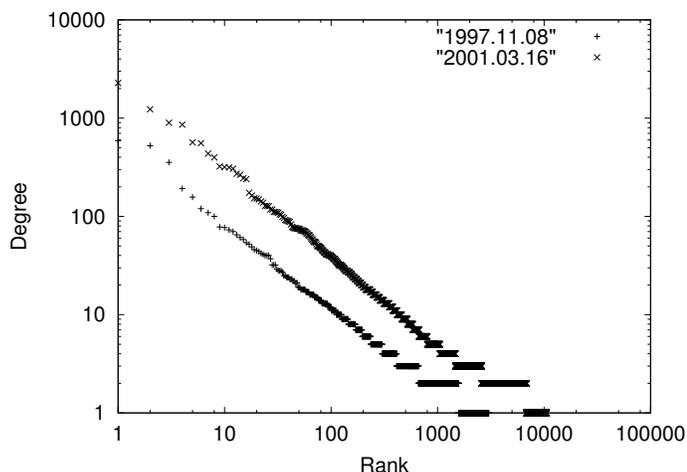


Fig. 17. Degree of nodes versus rank

first and last day. The degrees appear as line in the log-log plot [7]. The lower dataset corresponds to the first day and the higher dataset corresponds to the last day. It looks as if the two lines has shifted as it was expected.

VIII. CONCLUSIONS

We present a large scale evolution study of the Internet topology and routing at the Autonomous Systems level. We analyze 916 daily instances from November 1997 until March 2001 and study more than 1 billion routing paths. We study the AS dynamics at three different levels: a) the node level, b) the network as a whole, and c) the paths both topological and routing ones. We identify a number of phenomena that give us a clearer understanding of the dynamics of the Internet evolution.

Our work leads to the following observations:

1. The degree increase of a node is proportional to its degree (rich-get-richer phenomenon).
2. New nodes prefer higher degree nodes but not with linear preferentiality.
3. New internal edges prefer higher degree nodes but not with linear preferentiality.
4. The distribution of the degree of new and dead nodes is skewed, with very small degree.
5. The number of nodes grows exponentially, doubling every year approximately.
6. The number of Edges grows exponentially.
7. The degree exponent remains practically constant.
8. The distribution of the topological distances has remained practically the same.
9. Routing inflation becomes worse with time.

We note a small paradox: on the one hand we have exponential growth, and on the other hand, we have some invariant topological properties. Intrigued by this, we

study how phenomena at the node (micro) level can explain phenomena at the network (macro) level. We show that the rich-get-richer phenomenon leads to the exponential edge increase. We also show how the rich-get-richer phenomenon contributes to the preservice of the degree distribution and the related power-law.

Future work. Our observations bring us closer to understanding the dynamics of the Internet growth. We want to incorporate our observations into a realistic graph generation model. Such a model would give us two main benefits: a) we would be able to generate realistic graphs for simulations purposes, b) we could estimate and predict the evolution of the Internet. Overall, understanding the laws that the topology and routing obey can help us design better protocols and utilize the Internet more effectively.

ACKNOWLEDGMENTS. This work was partly supported by NSF CAREER proposal 9985195 and DARPA grant N66001-00-1-8936.

REFERENCES

- [1] R. Albert and A. Barabasi. Topology of complex networks: local events and universality. *Phys.Review*, 85, 2000.
- [2] A. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 8, October 1999.
- [3] Hyunseok Chang, Ramesh Govindan, Sugih Jamin, Scott Shenker, and Walter Willinger. On inferring as-level connectivity from bgp routing tables. <http://topology.eecs.umich.edu> submitted for publication, 2001.
- [4] Qian Chen, Hyunseok Chang, Ramesh Govindan, Sugih Jamin, Scott J. Shenker, and Walter Willinger. The origin of power laws in internet topologies revisited. <http://topology.eecs.umich.edu> submitted for publication, 2001.
- [5] Bill Cheswick and Hal Burch. Internet mapping project. *Wired Magazine*, December 1998. See <http://cm.bell-labs.com/cm/cs/who/ches/map/index.html>.
- [6] M. Faloutsos, A. Banerjee, and R. Pankaj. QoS MIC: a QoS Multicast Internet protocol. *ACM SIGCOMM*, Sep 2-4, Vancouver BC 1998.
- [7] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *ACM SIGCOMM*, pages 251-262, Sep 1-3, Cambridge MA, 1999.
- [8] A. Feldmann, A.C. Gilbert, P. Huang, and W. Willinger. Dynamics of ip traffic: A study of the role of variability and the impact of control,. *ACM SIGCOMM*, Sep 1999.
- [9] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM Transactions on Networking*, 9(4):392-403, Aug 2001.
- [10] National Laboratory for Applied Network Research. Routing data. Supported by NSF, <http://moat.nlanr.net/Routing/rawdata/>, 1998.
- [11] R. Govindan and A. Reddy. An analysis of internet inter-domain topology and route stability. *Proc. IEEE INFOCOM*, Kobe, Japan, April 7-11 1997.
- [12] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. *Proc. IEEE INFOCOM*, Tel Aviv, Israel, March 2000.
- [13] Bassam Halabi and Danny McPherson. *Internet Routing Architectures*. CISCO Press, 2000.

- [14] Geoff Huston. Analyzing the internet's bgp routing table. <http://www.telstra.net/gih>, 2001.
- [15] Cheng Jin, Qian Chen, and Sugih Jamin. Inet: Internet topology generator. *Technical Report UM CSE-TR-433-00*, 2000.
- [16] Damien Magoni and Jean Jacques Pansiot. Analysis of the autonomous system network topology. *ACM Computer Communication Review*, July 2001.
- [17] A. Medina, I. Matta, and J. Byers. On the origin of powerlaws in Internet topologies. *ACM Computer Communication Review*, 30(2):18–34, April 2000.
- [18] University of Oregon Route Views Project. Online data and reports. <http://www.routeviews.org/>.
- [19] J.-J. Pansiot and D Grad. On routes and multicast trees in the Internet. *ACM Computer Communication Review*, 28(1):41–50, January 1998.
- [20] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets. *ACM SIGCOMM*, San Diego, Aug 2001.
- [21] V. Paxson and S. Floyd. Why we don't know how to simulate the internet. *Proceedings of the 1997 Winter Simulation Conference*, December 1997.
- [22] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 2nd edition, 1992.
- [23] Y. Rekhter and T. Li (Eds). A Border Gateway Protocol 4 (BGP-4). RFC 1771, 1995.
- [24] Junhong Sun and M. Faloutsos. Generating small realistic internet topologies. *Network Group Communication (poster)*, 2000.
- [25] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin. The impact of routing policy on internet paths. *Proc. IEEE INFOCOM*, Anchorage, Alaska, March 2001.
- [26] Hongsuda Tangmunarunkit, Ramesh Govindan, and Scott Shenker. Internet path inflation due to policy routing. *SPIE IT-Com*, 2001.
- [27] D. Zappala. Alternate path routing for multicast. *Proc. IEEE INFOCOM*, Tel-Aviv, Israel 2000.
- [28] E. W. Zegura, K. L. Calvert, and M. J. Donahoo. A quantitative comparison of graph-based models for internetworks. *IEEE/ACM Transactions on Networking*, 5(6):770–783, December 1997. <http://www.cc.gatech.edu/projects/gtitm/>.

APPENDIX

I. THE SIGNIFICANCE OF ROUTING TABLES

Here we study the effect of using more routing tables on the coverage of the data. We find that the graph does not change significantly, if we add more similar in nature routing tables. We start with a graph based on the routing table of only one router. As we add the routing tables of other routers, we measure what is the combined number of nodes and edges. The results are illustrated in fig. 18, for the nodes, and fig. 19, for the edges. In fig. 18, we plot the number of nodes in the graph versus the number of router tables that are used to create the graph. Note that we have ordered the routing tables in order of decreasing number of nodes. We can see that the first router captures the majority of the nodes in the network. The consequent routing tables add less than 2% new nodes. In fig. 19, we plot the number of edges in the graph versus the number of routers that are used to create the graph. We observe that the importance of having more routers decreases fairly quickly, though slower than for the nodes. We see that 15 routers capture approximately 95%

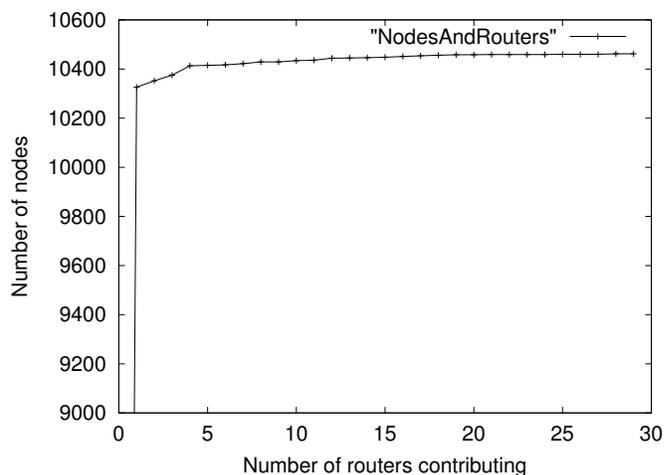


Fig. 18. Number of nodes in the graph versus the number of routers used to create the graph.

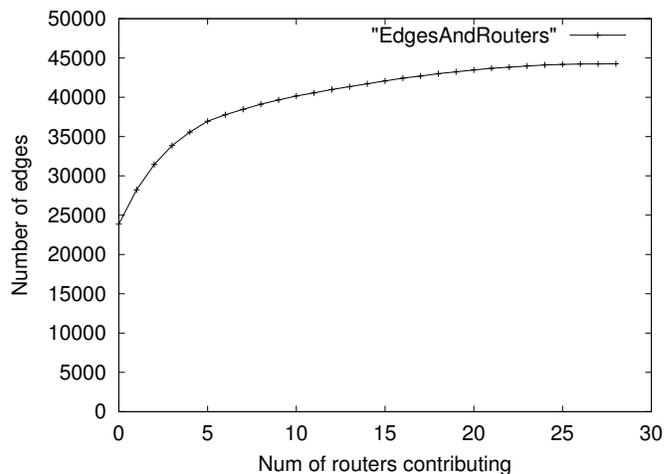


Fig. 19. Number of edges in the graph versus the number of routers used to create the graph.

of all the edges. Note that the last router adds only 0.01% of new edges. Note that this result does not imply that with 15 routers we can capture the whole of the AS map topology. It could very well be that some links are very difficult to observe. The routing tables that are used are more likely routing tables at the backbone, and it is possible to miss links at the periphery. A consolation is that these links are most likely not major backbone links.