On the Hierarchical Structure of the Logical Internet Graph

Zihui Ge, Daniel R. Figueiredo, Sharad Jaiswal*, Lixin Gao+
*Department of Computer Science
+Dept. of Electrical and Computer Engineering
University of Massachusetts
Amherst - MA 01002

ABSTRACT

The study of the Internet topology has recently received much attention from the research community. In particular, the observation that the network graph has interesting properties, such as power laws, that might be explored in a myriad of ways. Most of the work in characterizing the Internet graph is based on the physical network graph, i.e., the connectivity graph. In this paper we investigate how logical relationships between nodes of the AS graph can be used to gain insight to its structure. We characterize the logical graph using various metrics and identify the presence of power laws in the number of customers that a provider has. Using these logical relationships we define a structural model of the AS graph. The model highlights the hierarchical nature of logical relationships and the preferential connection to larger providers. We also investigate the consistency of this model over time and observe interesting properties of the hierarchical structure.

1. INTRODUCTION

The Internet has been growing at an exponential rate in the past few years. The growth has been physical (number of new nodes and links), functional (an explosion in the number of applications, protocols, etc.) and perhaps less importantly geographical (it now extends into almost any place on earth). Though the functional nature of the Internet (the applications that run on it) and its control mechanisms (the protocols) have been studied extensively, a fascinating area of study which is not as well understood is the physical structure of the Internet graph, in particular its topology. The topology of the Internet and its functional aspects are interrelated, and the manner in which the nodes are arranged in the network graph has a direct bearing on the performance of protocols.

In its higher level of abstraction, the Internet topology is characterized by the Autonomous Systems (AS) graph. An Autonomous System is a collection of routers and links in the Internet which are under one administrative domain and which present a unified face to the rest of the world in terms of accessibility and routing policies. There are currently about 11,000 ASes in the Internet. BGP (Border Gateway Protocol)[1] is the routing protocol used to exchange path information between ASes and enables them to construct routing tables. The exchange of such information between any two ASes can be disallowed or limited due to commercial agreements. These agreements can be viewed as logical relationships between ASes and have a direct influence in the paths taken over the graph. In this work, we introduce a hierarchical structure of the AS graph that is based on the inference of logical relationships between ASes.

A few studies on the nature of the Internet topology have been carried out recently with preliminary results aiming at both physical and logical characterization of different metrics of the Internet graph. The work by Faloutsos[2] explores the physical AS graph of the Internet and observes that certain metrics are described by *power law* distributions. Another direction of research has dealt with the problem of automatically constructing a realistic Internet topology.[3] Barabási et al[4] have also studied and developed an interesting model for the incremental growth of the Internet, which captures the "small world" phenomenon observed in complex networks. Both of these Internet topologies studies have primarily focused on the physical graph. A different approach emphasizes the importance of logical relationships that are inherently present in the AS graph due to routing policies. In,[5] the author introduces a classification of the possible logical relationships between ASes and uses these classifications to simplify router

L. Gao was supported in part by NSF grant NCR-9977555, ANI-0085848, and NSF CAREER Award grant ANI-9875513. The other authors are supported in part by NSF grants ANI-9809332, ANI-9805185, CDA-9502639, EIA-0080119 and by DARPA (TASC subcontract 2000-012). D. Figueiredo is also supported in part by CAPES (Brazil). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

configuration. In a recent work, Gao,[6] has introduced a heuristic algorithm to *infer* the logical relationships of the AS graph from BGP routing tables and used it to quantify the different relationships of the AS graph.

As noted above, most of the work on understanding the Internet topology has not taken into consideration the logical relationships between ASes. Moreover, even though there has been a vague notion of the existence of clusters (i.e., tiers) of ASes and their hierarchical arrangement, a precise methodology for classifying ASes according to specific criteria does not exist. An important observation is that the logical relationships between ASes can provide helpful insights into the structure of the AS graph. In particular, the logical edges of the graph can be used to precisely classify ASes into tiers and construct a hierarchy of clusters. In this work, we show that such logical information is indeed useful and provides a natural way to classify nodes into tiers and form a hierarchical structure of the graph. Such a simple hierarchical classification can be used to outline an approach to generate realistic Internet topologies.

In this paper, we analyze the Internet AS graph generated from different views and discuss properties of both the physical and logical graph. We provide evidence that a few different metrics of the AS graph follow a *power law*, and in particular, we identify the existence of a power law in the logical AS graph, namely, the number of customers that a provider has. We perform a comparison between the routing table data collected at two different and independent places and illustrate their difference. Hence, we use the union of these two data sets to attain a more complete view of the AS graph. Based on our observations and using the logical relationships between nodes, we formalize the concept of tiers and provide an algorithm to precisely classify AS nodes into these tiers. We also present a method for creating a hierarchical arrangement of the tiers, which follows naturally from the dominant relationships of the graph. Finally, we present different metrics over the hierarchical structure and observe interesting properties, such as the peering relationships between ASes in different tiers and the prevalence of *multi-homing* amongst the providers and top tiers.

The remainder of this paper is organized as follows. In Section 2 we provide a background and discuss the data collection methodology. In Section 3 we present the characteristics of the physical and logical AS graph. The proposed hierarchical structure of the logical AS graph is introduced in Section 4. We present results of different metrics and discuss some interesting properties of hierarchical model in Section 5. Finally, we conclude the paper and present ideas for future work in Section 6.

2. DATA COLLECTION AND BACKGROUND

An important step in analyzing the Internet topology is the construction of the autonomous system graph. The connectivity graph can be inferred by inspecting the routing table entries, which provide the nodes and edges of the network. The logical inference mechanism also uses routing table information to infer the relationship among connected ASes. The availability of more information (i.e., more routing table entries, different collection points) can be used by the algorithm to provide a more complete inference. In this work, we rely on the data collected by other institutions that are publicly available for general use.

The University of Oregon runs a project called Route Views that has as its main goal, to obtain a global view of the routing system from the perspective of several different ASes in the Internet. [*] Figure 1 illustrates how the collection of this global information is done. Basically, a router with its own autonomous system (AS number 65534) connects to many different ASes around the world using a "peering" session. This connection is done using the BGP routing algorithm with the multi-hop option. Thus, neighbors of Route View router advertise all their routes currently being used, and any possible changes using BGP. The Route Views router never advertises anything to any of its neighbors, so it has no impact in the Internet routing. The archive of the data is maintained by NLANR and more recently by the University of Oregon and is publicly available. [†]

The RIPE organization also runs a project called RIS (Routing Information Service) whose goal is to collect routing information between autonomous systems (AS). [‡] RIPE uses the same collecting strategy as the Route View project, but it peers with different ASes. Moreover, just recently RIPE has introduced different collection points that peer with different ASes, but the major collection site is the RRC00, located in Amsterdam. Note that while collecting the same information, that is routing table information, none of the snapshots can obtain the complete picture of the AS graph.

[*] Details of the project can be found in http://www.antc.uoregon.edu/route-views/

[†] The data is located on http://moat.nlanr.net/AS/Data

[‡] Details of the project can be found in http://www.ripe.net/ripencc/pub-services/np/

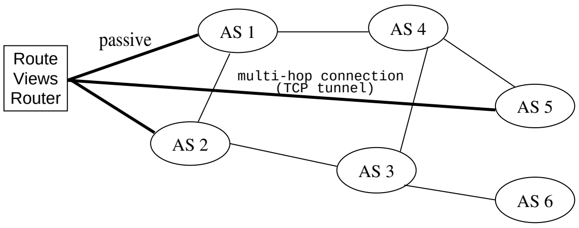

Figure 1. Route Views router and peer connection to other ASes

2.1. Inferring Logical Relationships in the Graph

The AS graph constructed from the routing table entries as described in the previous section provides the AS connectivity graph. However, routing policies among neighboring ASes often constrain the type of traffic that will be carried in that edge, as discussed in.[5] For example, an AS might have two providers and set its policy such that it will not act as a transit node (e.g., carry traffic) between the providers. This motivates the need to explore the logical AS graph, where relationship between nodes are considered. This logical relationship graph provides a more precise view of the actual Internet graph.

The idea of annotating the edges of the physical AS graph into logical relationships was recently introduced in.[6] In that work, the author defines three possible logical relationships for a link: customer-provider, peering, sibling. The customer-provider relationship is directed, thus one node is the provider and the other its customer. The other two relationships are undirected and both incident nodes play the same role. The paper goes beyond this categorization and provides a heuristic algorithm to infer these relationships from the information provided in the snapshot of the routing tables.

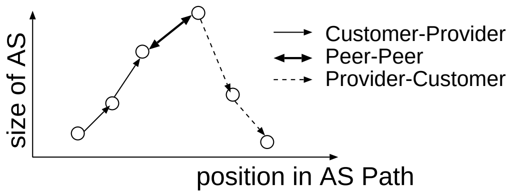

Figure 2. Inferring logical relationships from the routing table entries

The heuristic inference algorithm uses the physical degree of an AS as a measure of its "size". This ranks the importance of an AS according to its degree, assuming more important ASes have larger degree. This assumption was also used in.[7] The input to the inference algorithm is the routing table data file collected by the projects mentioned above. The algorithm basically uses the AS path information to construct and infer the possible relationships.

The main idea of the algorithm is illustrated in Figure 2, which depicts a single AS path entry of the routing table. Given an AS path entry, the algorithm infers that this path must traverse customer-provider edges until it reaches the highest AS. At the highest AS, the path either traverses a peering, sibling or provider-customer link. After that the path follows provider-customer edges until it reaches its final destination. Note that the AS with the highest degree is the top provider in a given path. Deciding if a neighbor of the highest AS has either a peering or sibling relationship requires examining other AS path entries. By elimination, we can discard the possibilities as the algorithm progresses, and at the end we are left with a single possible classification. Note that the algorithm is an heuristic and can incorrectly classify relationships. The detailed algorithm and the proofs of existence of such relationships can be found in.[6]

We implemented a robust and efficient version of the algorithm, which constructs both the physical AS graph and the annotated logical AS graph from the routing table data. [§] Note that our purpose here is not to improve neither explore the heuristics of the algorithm, but rather obtain the annotated AS graph for further analysis.

2.2. Metrics

A set of measures are introduced in order to characterize and analyze the AS graph. These metrics are useful in understanding the properties of the AS graph and will give insights to general observations that can be drawn from

[§] The inference program is publicly available at http://www-net.cs.umass.edu/~ ratton/AS

	1999/12	2000/06	2000/11	2001/05
nodes	6312	7891	9545	11038
edges	12801	16782	20950	23791
avg degree	4.05608	4.25345	4.38973	4.31074
max degree	1418	1744	2076	2376
avg distance	3.676	3.613	3.616	3.616
diameter	9	10	10	11

Table 1. Characteristics of the physical AS graph on different dates.

the graph.

In particular we investigate the total number of nodes and edges in the graph, which provide a measure of the graph size; the node degree, which is defined as the number of edges connected to a given node; the shortest edge hop distance between a given pair of nodes; and the diameter of the graph, which is the maximum shortest path distance between any two nodes.

Besides these traditional graph metrics, we will investigate the degree rank of the nodes and the CCDF (Complementary Cumulative Density Function) of the degree of the nodes. The former metric establishes an ordering in the nodes of the graph, where the node with the highest degree has rank 1, and so on. This metric was introduced in,[2] where the authors measure the AS graph at three instances of time. The latter is a probability metric, defined as Prob. $[D \geq d]$, where D is a random variable that indicates the number of incident edges upon a node (i.e., its degree) and was introduced in.[8] It is interesting to note that both these metrics exhibit the important property of being a *power law*, which we depict in the following section.

2.3. Power Laws

Power laws are an important characteristic of functions that are generally associated with probability theory. In a general sense, a power law is any expression of the form $y \propto x^{-\alpha}$ where α is a constant. In probability theory, power laws have a more strict definition. Let X be a random variable. We say X exhibits a power law if $\mathrm{P}[X > x] \approx x^{-\alpha}$ as $x \to \infty$ and $0 < \alpha < 2$.[9] The exponent of the power law has important implications on the moments of X. In particular, if $0 < \alpha < 1$, $\mathrm{E}[X] \to \infty$ and if $1 \leq \alpha < 2$, $\mathrm{Var}[X] \to \infty$. Note that a condition for X to be considered a power law is to let $x \to \infty$, which is systematically relaxed when analyzing real data. As it turns out, all measured data are finite, and usually the value of x cannot be made arbitrarily large. Moreover, noise in the data is usually present for very large values of x, where the number of samples is small.

Power laws have been observed in many different scenarios.[9] Intuitively, a power law relation assumes small values with high probability and large values with low probability. However, this low probability does not decay exponentially, which leads to the possibility of events with very large values. A power law can be identified by plotting the function in a log-log graph and fitting the curve to a straight line. If the function is well approximated by a line, then we say that it obeys a power law. Moreover, the exponent α of the power law is given by the negative value of the slope of the fitted line.

3. PHYSICAL AND LOGICAL CHARACTERISTICS OF THE AS GRAPH

In this section we will investigate the physical and logical characteristics of the Internet AS graph. By analyzing different properties of the graph one can gain a better understanding of its global structure. These properties are defined by the set of measures introduced in the previous section that can be applied both on the physical and logical graph.

The analysis is based on four different dates, that are around 6 months apart from each other, spanning a one and a half year period. The dates were carefully chosen to avoid any corrupted snapshots from any of the collection points. The most recent data available at the time of this write up was used. In particular the dates chose are: 2001/05/07, 2000/11/29, 2000/06/13 and 1999/12/06. All results presented in this paper are based on the union of the data from the snapshots taken by the Route View and RIPE collection points.

3.1. Physical AS Graph

In Table 1 the results for the metrics defined are presented for the four different points in time. A few interesting observations can be drawn from these results. We start by noting that the number of nodes and edges in the graph has almost doubled over the 18 month period. More surprising is the fact that a network with over 11000 nodes has such a small diameter and low average pairwise distance, and that despite the doubling in nodes and edges, both of these metrics have remained practically the same with a small decrease in the average distance. Another observation is the increase in the maximum degree, and the fact that, currently, there exists an AS with over 2376 neighbors. Finally, we note that the average degree, which is the ratio between the number of edges (doubled) and nodes in the graph, is roughly stable and presents a small increase over time.

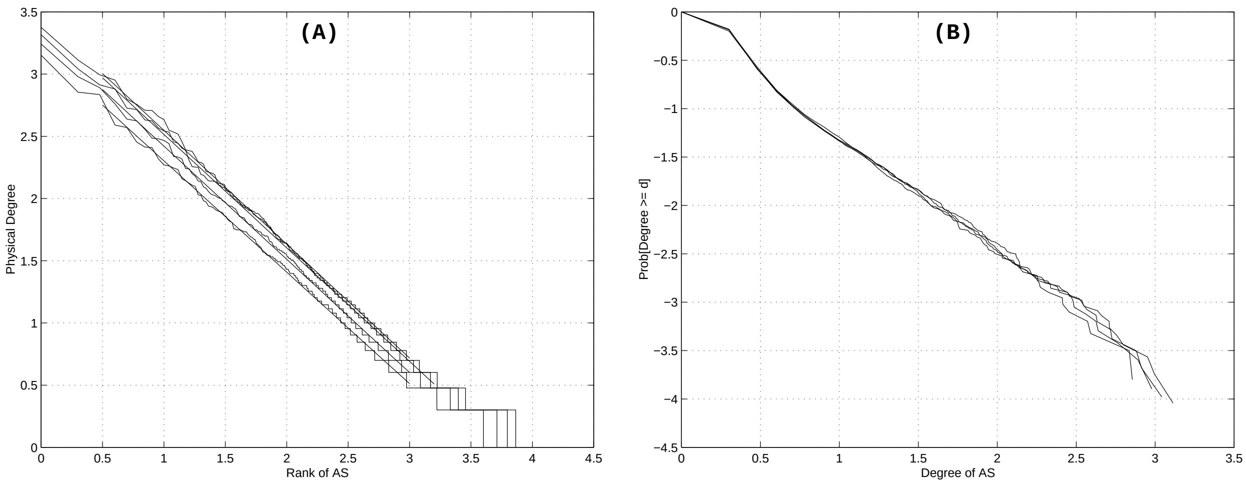

Figure 3. (A) Log-log plot of AS degree rank for different dates. (B) CCDF of AS degree for different dates.

The values for the other two metrics are now presented. Figure 3.A plots the degree rank of all ASes in a log-log plot (base 10) for different instances of time. We observe that the curves for more recent data are always above older dates. This is expected, since the degree of an AS of rank i is likely to increase over time. Thus as time passes, the curves will be shifted upwards. This behavior is due to the definition of the degree rank metric, which is an absolute measure.

Figure 3.A also shows that the degree rank for each time instance can be fitted to a straight line, indicating that this metric exhibit a power law. The straight line fit was performed aiming at the best possible match considering the points in the middle of each curve and discarding possible outliers in both low value and high value ends. The slope of the line at each time instance characterizes the parameter of the power law. The values obtained for α are 0.896, 0.910, 0.912 and 0.915 for the respective dates in chronological order. This indicates that the trend for the exponent of this power law is to remain constant over time. We note that this observation is not consistent with the results in,[2] where different dates showed different values for the power law exponent. The main reason for this discrepancy is the use of different dates and different approaches in fitting the curves to a straight line.

The last metric is illustrated in Figure 3.B, which plots in log-log scale the CCDF of the node degree for different dates in the same graph. The curves for different dates are very similar to one another and differences arise only in the tail, where the degree is high and only few data points are available. This similarity is confirmed by the very small differences in the α parameter of the power law, which were 1.19, 1.16, 1.13 and 1.11, respectively, in chronological order. While fitting the curves with a straight line, we focused on the middle values of the curves and ignored possible outliers at both low and high values.

From Figure 3 we conclude that the CCDF is a better measure to characterize the AS graph, since it is a probabilistic measure and does not depend on the actual graph size. Thus, we expect the slope of the linear fit of the curves to remain the same over time. Another interesting property of this metric is that as time passes, more data points will become available in the high degree region, since ASes have the tendency to increase their degrees. This will allow a better line fit for the current high degree region of the plot.

	1999/12	2000/06	2000/11	2001/05
P-C	90.66% (11606)	91.49% (15354)	90.15% (18887)	91.01% (21651)
PEER	7.90% (1011)	7.34% (1231)	8.23% (1724)	7.36% (1752)
SIB	1.44% (184)	1.17% (197)	1.62% (339)	1.63% (388)
avg distance	3.707	3.641	3.643	3.641
diameter	16	17	17	14

Table 2. Characteristics of the logical AS graph on different dates.

3.2. Logical AS Graph

So far we have only investigated the characteristics of the physical AS graph. However, as introduced in section 2, policy issues among neighboring ASes can have impact on connectivity. Thus, we now focus on the logical AS graph, which is the physical AS graph with its edges annotated with a logical relationship. We will use the same metrics as the ones presented for the physical AS graph, but will distinguish among different edges according to their classification. An important consideration is the distance between a given pair of nodes in the logical graph.

3.2.1. Computing Distances on the Logical AS Graph

⊕	SIB	C2P	P2C	PEER	CPC	NOT VALID
SIB	SIB	C2P	P2C	PEER	CPC	NOT VALID
C2P	C2P	C2P	CPC	PEER	CPC	NOT VALID
P2C	P2C	NOT VALID	P2C	NOT VALID	NOT VALID	NOT VALID
PEER	PEER	NOT VALID	PEER	NOT VALID	NOT VALID	NOT VALID
CPC	CPC	NOT VALID	CPC	NOT VALID	NOT VALID	NOT VALID
NOT VALID	NOT VALID	NOT VALID	NOT VALID	NOT VALID	NOT VALID	NOT VALID

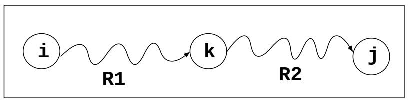

(B)

Figure 4. Algebra rules for logical relationships. **(B)** Relationships using the algebra rules.

Note that in the physical graph all edges were of the same type and the computation of the distance between a given pair was simply based on hop count. However, in the logical annotated AS graph the distance between a pair of nodes should be computed only over legitimate paths. To better understand how links with different relationships can be put together to form a valid path, we developed a small algebra of link relations, which is illustrated in Figure 4. This algebra determines the relation an "edge" (i.e., path) obtains when joining two "edges" (i.e., sub-paths) with different relationships.

Figure 4.B provides an example of the algebra. Suppose node i has a path to node k of relation type R1; and node k has a path to node j of relation type R2. Thus, a path between i and j exists if the operation R1 + R2 is valid, and its relation is given by the corresponding table entry.

This table can be constructed by analyzing case by case and considering all possible paths between any two nodes of the annotated AS graph. An example of how the entries in the table were obtained is by considering the case of joining C2P + C2P sub-paths. Using Figure 4, it is easy to see that if two nodes are separated by two C2P relationships then their resulting relationship is also C2P. A similar argument holds true for PEER and SIB relationships. We also introduced a new relationship called CPC (Customer-Provide-Customer) which typically denotes the logical relationship between any two nodes that are separated by the same provider.

Using the algebra and the table illustrated in Figure 4, it is simple to to compute the shortest valid path between all pairs of nodes. This computation can be carried out using the classical Floyd-Warshall[10] algorithm, which runs in $O(N^3)$, where N is the number of nodes in the graph.

3.2.2. Analysis of the Logical AS Graph

The annotated AS graph generated by the logical relationship inference algorithm provides a logical picture of the AS graph. Using the metrics defined in the previous section, we obtain the characteristics of this logical graph. Table 2 summarizes the results for the different dates. A striking observation is that 90% of all logical relationships between AS pairs is of the type customer-provider. Of the remaining, around 7% percent have a peer-to-peer relationship, and a very small fraction, sibling-sibling. This observation is consistent with the results in,[6] but we show that this trend holds over a longer period of time and over the unified AS graph snapshots (Route View plus RIPE).

The average distance between any two AS nodes in the logical graph is almost the same as in the physical graph, while the diameter is much longer (17, compared to 11 in the physical graph). We owe this to the fact that though

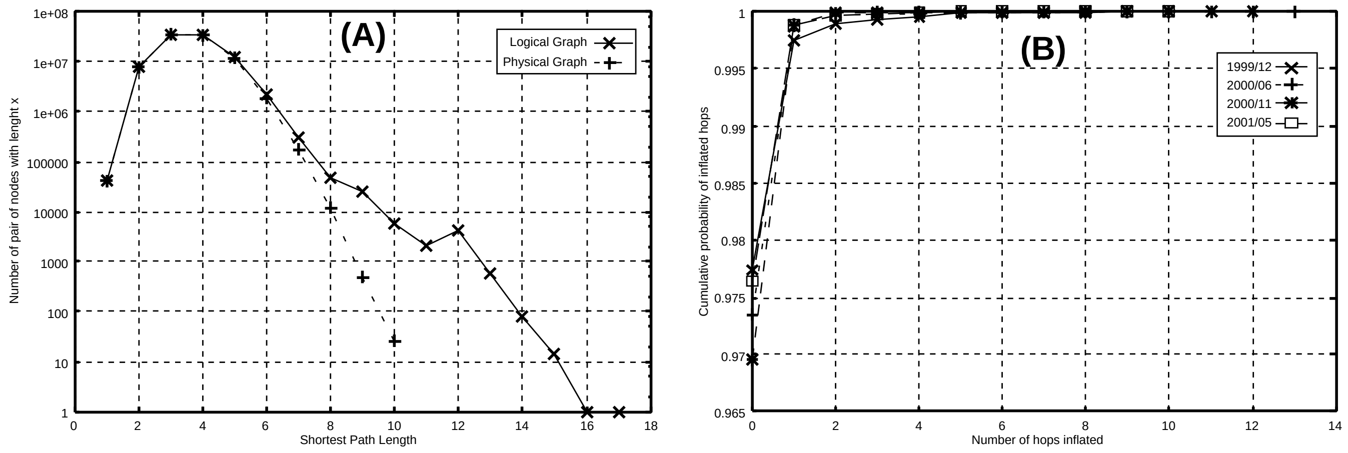

Figure 5. (A) Histogram of path lengths for all node pairs (2000/11). (B) Cumulative distribution of path inflation in the logical AS graph.

some AS nodes might be physically separated from each other by a small number of links, these links may not be traversable due to the nature of their logical relationships. This translates to some nodes being separated by much longer distances than their counterparts in the physical graph yielding the difference in diameters.

This fact is also apparent in Figure 5.A, where a histogram of the number of ASes that have distance k is plotted for the both physical and logical graphs of the 2000/11 data. The histogram for the logical graph shows a much longer tail than the physical graph, corroborating our hypothesis that some pairwise distances are inflated. The same phenomena is also present for the other dates.

However, the path inflation observed in Figure 5.A is not significant to alter the main structure of paths the graph in the physical graph. Another metric that provides a different view of the path inflation is the inflation difference, which measures the absolute difference between the physical and logical path. This metric was introduced in,[11] where the authors analyze the impact of the hierarchical structure of ASes on router-level path inflation. Figure 5.B shows the cumulative distribution of the inflation difference for the different dates. We observe that nearly 97% of the paths between two nodes are *not* inflated. Moreover, only around 0.5% of the paths are inflated by 2 or more hops. This result does not agree with the observations in,[11] however, in that work the path inflation is considered in the router level, while we analyze the path inflation in the AS level (e.g., a hop in their work is a link between routers while here is between ASes).

A final observation consists on characterizing the logical relationship degree of each node. We have already observed that the vast majority of edges have a provider-customer relationship, and we now further investigate the structure of these relationships in the overall AS graph. In particular, we are interested in the distribution of the number of nodes that are customers of a given provider. Figure 6 shows in log-log scale the CCDF of the number of customer a provider has for all four dates. As illustrated before for the physical graph, the curves plotted can be fitted to a straight line, which indicates the presence of a power law. This phenomena only occurs with the provider-customer relationship and does not hold true for the other logical relationships. We note that this plot is comparable with the power law in the CCDF of the physical degree illustrated in Figure 3.B, which leads us to conclude that the provider-customer relationship also dominates the structure of the edges in the graph. This is important when presenting a model for the structure of the AS graph.

3.3. Differences Between Independent Snapshots

The analysis and observations presented so far have relied on the unified snapshots (Route View and RIPE) of the AS graph taken the same day and almost at the same time. A natural question that might rise is how similar are these two snapshots? It should be clear that these independent snapshots can be different from each other due to the different peering points. The peering nodes provide different levels of detail over different portions of the Internet AS graph. In this section we give a brief comparison between the two snapshots and conclude that there are potential differences in the logical graph.

	2001/04				2001/05			
	Route View	Ripe RRC00	Ripe RRC01	Union	Route View	Ripe RRC00	Ripe RRC01	Union
nodes	10856	10861	10591	10920	10966	10994	10737	11038
edges	22753	21743	14383	23912	22536	21533	14512	23791
avg degree	4.19178	4.00387	2.71608	4.37949	4.11016	3.91723	2.70318	4.31074
max degree	2354	2312	1825	2371	2353	2323	1855	2376
avg distance	3.626	3.662	4.254	3.607	3.638	3.675	4.239	3.616
diameter	10	10	11	10	11	11	12	11

Table 3. Characteristics of the physical AS graph from different collection points.

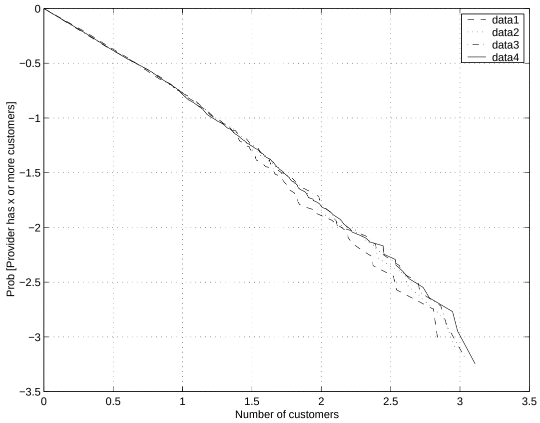

Figure 6. CCDF of AS degree of only P2C relationship for different dates

We characterize three different and independent snapshots: Route View, RIPE RRC00 and RIPE RRC01. The physical and logical properties of each individual snapshot is presented together with their union for two different dates. Table 3 illustrates the results obtained. We first note that all three snapshots agree well on the number of nodes in the AS graph, with RIPE RRC00 having the largest value among the three. However, the picture is a bit different for the number of edges. Route View has the largest number of edges with 4.5% more than RIPE RRC00 and 37% more than RIPE RRC01. Clearly, RIPE RRC01 does not capture the essence of the graph, since the graph obtained solely from this snapshot has a much lower number of edges. This incomplete information leads to significant discrepancies on the degree metrics. We also note that neither Route View nor RIPE capture the maximum degree of the graph, since the union has the node with the largest degree. The union also shows a substantial increase in the average node degree, due to the larger overall number of edges. However, the average distance and the diameter of the union, when compared to the Route View or RIPE RRC00 data is practically the same. The edges added by each snapshot enhance the graph connectivity, but have small impact on physical distances.

A comparison between the snapshots can also be drawn in the logical AS graph. Table 4 shows the percentage of

	2001/04				2001/05			
	Route View	Ripe RRC00	Ripe RRC01	Union	Route View	Ripe RRC00	Ripe RRC01	Union
P-C	91.35% (20786)	94.31% (20505)	97.87% (14077)	90.49% (21637)	92.02% (20737)	94.76% (20404)	97.94% (14213)	91.01% (21651)
PEER	7.26% (1651)	4.54% (987)	1.80% (259)	7.78% (1860)	6.76% (1523)	4.18% (900)	1.75% (254)	7.36% (1752)
SIB	1.39% (316)	1.15% (251)	0.33% (47)	1.74% (415)	1.22% (276)	1.06% (229)	0.31% (45)	1.63% (388)
avg distance	3.659	3.689	4.589	3.633	3.673	3.704	4.486	3.641
diameter	15	15	12	21	16	15	14	14

Table 4. Characteristics of the logical AS graph from different collection points.

	2001/04		2001/05	
	Route View	Ripe RRC00	Route View	Ripe RRC00
nodes	59 (0.54%)	63 (0.58%)	44 (0.40%)	71 (0.65%)
degree 1	46	58	33	64
degree 2	10	5	9	6
degree ≥ 3	3	0	2	1
edges	2169 (4.99%)	1159 (2.55%)	2258 (5.24%)	1255 (2.78%)
P-C	5.61% (1283)	3.57% (772)	6.45% (1397)	4.09% (886)
PEER	46.67% (868)	20.32% (378)	47.95% (840)	20.43% (358)
SIB	4.34% (18)	2.17% (9)	5.41% (21)	2.84% (11)

Table 5. Differences between Route View and RRC collection points.

each type of relationship for each particular snapshot together with the absolute value. From this table we can observe that the Route View percentage breakdown has a better match with the union, indicating that it contains the essence of the logical AS graph. Note the increased number of customer-provider relationships in both RIPE snapshots and the reduced number of peering relationships. From this we conclude that the edges that are not captured in these snapshots have a significant impact in the logical AS graph, and mostly affect the peering relationships of the graph.

Table 5 illustrates the uniqueness of each individual snapshot. By uniqueness we mean the nodes, edges and relationships that appear in one snapshot but not the other. We also characterize the unique nodes and investigate their degree. From the table we note that the number of unique nodes relative to the total number of nodes is relatively small (less than 1%). Moreover, we observe that the degree of these unique nodes is also very small. For example, around 90% of the unique nodes in the RIPE data have degree 1, thus are end-systems in the AS graph. We conclude that these unique nodes have minor impact in the essence of the graph and are almost insignificant.

However, the same does not hold true for the number of unique edges and unique relationships in the graph. We note that around 5% of the edges are unique to the Route View data and that around 2.5% are unique to the RIPE snapshot. These edges have a significant impact on the logical structure of the graph. Note that around 53% of the peering relationships are unique to the Route View snapshot and that 16% are unique to RIPE. Thus, the union of both snapshots provide a better picture of the physical AS graph, in particular considering the edges, and also provide more information about the paths to the inference mechanism.

4. LOGICAL HIERARCHY OF AS GRAPH

We shall now present a hierarchical model of the logical AS graph. To begin with, we present the intuition behind the construction of this hierarchy. Network traffic originating or terminating at an AS node, in most cases, has to be routed through its provider. For example, if an AS wants to communicate with any other destination node, it will start by forwarding traffic to its provider. The provider, depending on its connectivity, might in turn have to route through its own provider or through a peer node. This immediately presents the picture of messages being routed up the ladder from one customer to its provider, until they reach either a "top-level" provider or a peering point, and then routed down the ladder to the destination AS. This is a rough picture of the hierarchy of AS nodes defined by customer-provider relationships. Based on the inference of logical relationships between ASes, we have observed that customer-provider relationships dominate in the AS graph. This reaffirms our intuition that the hierarchy imposed by these relationships is the inherent functional structure of the whole graph.

Using this intuition, we now present a precise methodology for classifying nodes into this hierarchy. Of all the provider ASes in the Internet, a small fraction are responsible for a very large number of customers and have a very high connectivity in the backbone. This set of providers are usually denominated tier 1 providers (represented as the set $Tier_1provider$). More formally,

1. A node belongs to $Tier_1provider$ if it does not have a provider and forms a clique with others in the set.
2. Define the remaining levels of the hierarchy recursively as follows: a node belongs to $Tier_{i+1}$ if it has a provider that belongs to $Tier_iprovider$. Furthermore, if it has a customer, then it belongs to the set $Tier_{i+1}provider$ else it is placed in the set $Tier_{i+1}customer$.

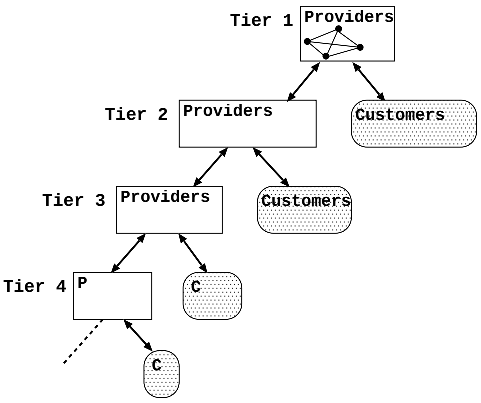

Figure 7. The Hierarchical model for the AS graph

Figure 7 illustrates this classification succinctly. A precise definition of **Provider** and **Customer** nodes are presented below in Phase 1 of the classification algorithm.

We now describe a detailed methodology consisting of three phases that classify each AS node into its position in the hierarchy. In the first phase a node will be classified as either a **Customer** or a **Provider**. This information is used to determine in which box (customer or provider) a node will be placed once its tier level has been determined. In the second phase, we identify all nodes belonging to the $Tier_1provider$ set. In the third and final phase we classify the remaining nodes in their respective tiers. A step-by-step description and accompanying pseudo-code is provided below.

Phase I: Classification of AS nodes as Customer or Provider
Input: Annotated AS graph ($\mathbf{V}$, $\mathbf{E}$, $\mathbf{R}$) where $\mathbf{R}$ is the mapping of edges to relations
Output: Set **Customer** and **Provider**

$\mathbf{Provider} \leftarrow \emptyset$
for each AS node $i \in \mathbf{V}$
    for each edge $\langle i, j \rangle \in \mathbf{E}$
        if $\mathbf{R}(\langle i, j \rangle)$ = Provider-Customer
            $\mathbf{Provider} \leftarrow \mathbf{Provider} \cup \{i\}$
$\mathbf{Customer} \leftarrow \mathbf{V} - \mathbf{Provider}$

Figure 8. *Initial classification of nodes*

Phase II: Classification of $Tier_1$ providers
Input: Annotated AS graph ($\mathbf{V}$, $\mathbf{E}$, $\mathbf{R}$) where $\mathbf{R}$ is the mapping of edges to relations
Set **Provider**
Output: Set $\mathbf{Tier_1provider}$

Sort **Provider** by decreasing degree, store the result in vector $Sorted\_Provider$
$\mathbf{Tier_1provider} \leftarrow \{ \, Sorted\_Provider[0] \, \}$
for $i \leftarrow 1$ **to** the size of $Sorted\_Provider$
    for each AS nodes $j \in \mathbf{Tier_1provider}$
        if $(\langle Sorted\_Provider[i], j \rangle \notin \mathbf{E})$
            next i
    $\mathbf{Tier_1provider} \leftarrow \mathbf{Tier_1provider} \cup \{ \, Sorted\_Provider[i] \, \}$

Figure 9. *Classification of $Tier_1$ nodes*

Phase 1: In Phase 1 each AS will be categorized as either a **Provider** or a **Customer**. This will determine, once the tier level of the node has been determined, whether it would go into the left hand provider box or to the customer box in the right. In order to do this, we follow the following rule of thumb. If from any node, a provider-to-customer relationship exists to any other node, this node is classified a **Provider**. All remaining nodes are classified as a **Customer**. The pseudo-code for this classification is illustrated in Figure 8.

```
Phase III: Assigning each AS node to its tier in the hierarchy
Input: Annotated AS graph (V, E, R) where R is the mapping of edges to relations
       Set Provider, Tier1provider
Output: Set Tier2provider, Tier2customer, Tier3provider, Tier3customer, ...
level ← 1
Unclassified ← V − Tier1provider
repeat
    no_new_AS_classified ← True
    Tier_{level+1}provider ← ∅
    Tier_{level+1}customer ← ∅
    for each AS nodes i ∈ Unclassified
        for each AS nodes j ∈ Tier_{level}provider
            if ((⟨j, i⟩ ∈ E) and (R(⟨j, i⟩) = Provider-Customer))
                no_new_AS_classified ← False
                if (i ∈ Provider)
                    Tier_{level+1}provider ← Tier_{level+1}provider ∪ { i }
                else
                    Tier_{level+1}customer ← Tier_{level+1}customer ∪ { i }
                Unclassified ← Unclassified − { i }
    level ← level + 1
until (no_new_AS_classified = True)
```

Figure 10. *Classifying each node in its particular tier*

Phase 2: In Phase 2, we shall identify all the $Tier_1$ nodes. Previously we had stipulated that a node belongs to $Tier_1$ if it has no provider of its own and forms a clique with all other $Tier_1$ nodes. Validating that a node does not have a provider depends on our inference heuristic, which could be incomplete. This could result in leaving out some very big providers which provide connectivity in the backbone. Hence, we adopt an inference-independent notion of Tier 1 which is the set of all ASes that form a clique. Since finding the maximal clique in a graph is a *NP-hard* problem, we have introduced a heuristic which allows us to build a clique of $Tier_1$ providers. We sort all AS nodes in the order of decreasing degree. Parsing down this sorted list of nodes, we identify all nodes which compose a clique in the physical graph. We assume that the node with the largest degree would definitely be in $Tier_1$ and use it to initialize the clique. The pseudo-code for this process is illustrated in Figure 9. The nodes in the clique are classified as $Tier_{1provider}$. The heuristic works well in identifying the correct set of $Tier_1$ nodes since it is reasonable to assume that for a node to be in $Tier_1$ it would almost certainly have a large number of customers and hence a large degree.

Phase 3: Now we classify the remaining nodes. We carry out an iterative operation illustrated in the pseudo-code presented in Figure 10. Beginning with $Tier_1$ (the nodes of which have already been classified) we identify all unclassified ASes such that a provider-to-customer (P2C) relationship exists between a $Tier_1$ node and the unclassified AS. Of all such nodes, we classify as *Tier*$_2$*provider* all nodes currently labeled as **Provider**, and the others as *Tier*$_2$*customer* In a similar fashion we go down the hierarchy, to identify $Tier_3$, $Tier_4$ ... etc.

Unfortunately, these three phases do not completely classify all nodes of the AS graph. A small fraction of nodes (less than 1%) are left out of the hierarchy. This could be due to erroneous routing information which result in incorrect logical inference or mistakes in the heuristic itself. In order to classify these nodes we add the following rules:

An unclassified node which is either a sibling, peer or provider of a node in *Tier*$_1$*provider* is also placed in *Tier*$_1$*provider*. If this unclassified node has either a sibling or peering relationship to a node in $Tier_i$, it is placed in *Tier*$_i$*provider* or *Tier*$_i$*customer*, according to its Phase 1 classification. If not, and the unclassified node is a provider of a $Tier_i$ node, then it belongs to the set *Tier*$_{i-1}$*provider*.

Level	Type	1999/12	2000/06	2000/11	2001/05
1	provider	0.19% (12)	0.24% (19)	0.27% (26)	0.20% (22)
2	provider	7.97% (503)	8.90% (702)	8.50% (811)	8.23% (908)
	customer	34.92% (2204)	39.10% (3085)	39.21% (3743)	39.14% (4320)
3	provider	7.62% (481)	6.56% (518)	6.30% (601)	6.16% (680)
	customer	32.67% (2062)	31.80% (2509)	31.71% (3027)	31.83% (3513)
4	provider	1.70% (107)	1.10% (87)	1.19% (114)	1.21% (134)
	customer	12.91% (815)	10.59%	11.64% (1111)	11.43% (1262)
5	provider	0.17% (11)	0.13% (10)	0.13% (12)	0.15% (17)
	customer	1.57% (99)	1.42% (112)	0.94% (90)	1.42% (157)
6	provider	0.03% (2)	0.01% (1)	-	0.03% (3)
	customer	0.13% (8)	0.14% (11)	0.10% (10)	0.14% (16)
7	provider	-	-	-	-
	customer	0.13% (8)	0.01% (1)	-	0.05% (6)
Total	provider	17.7% (1116)	16.9% (1337)	16.4% (1564)	16.0% (1764)
	customer	82.3% (5196)	83.1% (6554)	83.6% (7981)	84.0% (9274)

Table 6. Characteristics of the size of each box in the hierarchical model.

5. PROPERTIES OF HIERARCHY AND AS GRAPH OVER TIME

In this section we present a few properties of the proposed hierarchical structure and show that this logical arrangement of the AS graph is precise and consistent over time. The analysis of the characteristics of the hierarchical structure will also provide useful insights on how the AS graph is formed and how it grows over time.

The first property investigated is the size of each box in the hierarchy, that is, the number of ASes in each cluster. As described, each level of the AS graph has two clusters, one for provider nodes and the other for customers. The relative and absolute size of each box is presented Table 6. We start by noting that the number of nodes that are customers is around 83%, which can be seen in the last line of the table. This indicates that the vast majority of ASes do not provide any service, but act solely as customers, and form the "edge" of the Internet. However, a trend can be observed in the percentage number of providers, which is slowly decreasing over time. A plausible explanation is the increasingly popularity of the Internet and the necessity of large businesses to become interconnected, which can create ASes for private purposes that do not offer service to others. We also note that the number of levels in the hierarchy is 7 for three of the dates observed and 6 for the other, which are fairly low values. This indicates that the large growth of the Internet is not reflected in its depth but rather occurs in its breadth.

An surprising observation is that the distribution of nodes in the levels of the hierarchy is very skewed. From Table 6 we observe that around 39% of all AS nodes are customers of tier 1, and that around 32% of the node are customers of tier 2 providers. This indicates that the vast majority of the nodes are located in the first couple of levels of the hierarchical structure. Also note that the number of customers in each level decrease as the level increases. In particular, it has a significant jump from level 2 to level 3, and from level 3 to level 4. This phenomena is also present in the providers' nodes, but starts with level 2. The number of nodes in tier 1 is quite small, and most providers are located in tier 2. From this level below, this size also decreases, having a significant jump from level 3 to level 4. Economical reasons might explain why the number of tier 2 providers is larger than tier 1, since the cost-benefit of offering service as a tier 2 provider might be higher. Moreover, tier 1 providers usually are represented by international or national carriers which have very large networks and may form a private backbone, while tier 2 providers are formed by regional and smaller carriers.

An important observation is that the relative size of all boxes tend to be somewhat constant over time, despite the significant increase in their absolute number. Note that the number of ASes that are customers of tier 1 nearly doubled over the time period, but the relative number remained nearly the same. This same trend is also observed at other levels of the hierarchy and is also present for the providers side. This trend supports the hierarchical structure proposed showing its robustness over time.

The logical hierarchical structure proposed relies on the customer-provider relationships which dominate the percentage of edges in the AS graph. To investigate how the proposed hierarchy divides these relationships we analyze the distribution of the customer-provider relationships across different levels. This is illustrated in Figure 11 which shows the percentage of provider-customer edges between a provider in level i and a customer in level $i + 1$; between providers and customers in the same level i; between providers and providers that are different levels (a provider node in level i has a customer which is also a provider in a higher level); and between providers and customers that are more than one level apart. We note that the vast majority of these edges follow the hierarchical

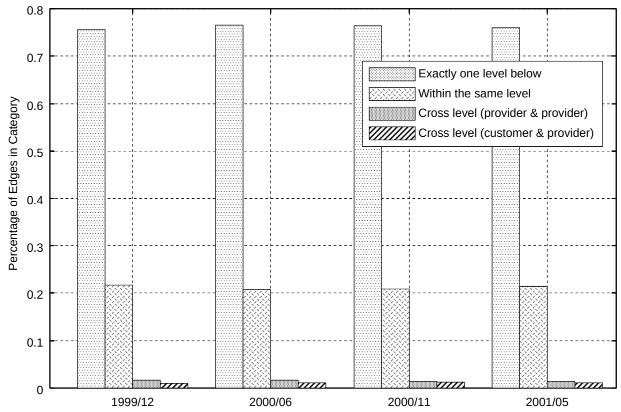

Figure 11. Customer-Provider relationship distribution over levels

structure, that is, are either exactly one level below or are within the same level. Around 2% of all the provider-customer relationships violate the hierarchical structure crossing from levels not directly connected. However, this value is small enough not to perturb the logical structure proposed, and are potentially due to relationship inference errors. Another key observation is that the distributions are nearly constant over time, as illustrated in Figure 11.

The peering relationship is the second most important type of edge since it is responsible for around 7% of all edges in the AS graph, as presented in section 3. Thus, we investigate how the peering edges are distributed on the hierarchical structure. Figure 12.A plots the distribution of peering relationships between the two types of clusters. As one might expect, the largest number of peering occurs between provider nodes, a much smaller number occurs between provider and customer nodes, and just a small fraction occurs between customer nodes. This is another characteristic that distinguishes provider from customer nodes, since providers can benefit the most from peering relationships.

To understand where the peering relationships are occurring in the hierarchy, we observe the distribution of the peering edges over different levels. Figure 12.B illustrates the percentage of peering edges that are incident to a particular level. We first note that the majority of the peering relationships are incident to level 2, which can be observed by the relative high bar in the graph. Peering with level 1 comes second followed by peering with level 3, while the remaining peering relationships is very small.

These results lead to an interesting observation which states that level 2 has most of the peering relationships of the AS graph, and that other nodes peer with level 2 more than with any other level of the hierarchy. We should note that level 2 has the largest provider cluster in the hierarchy. However, other reasons besides its size can contribute for the high peering relations with this level. Tier 2 providers are strategically placed in the hierarchy, attracting other ASes to peer with them. It is also known, that peering relationships are usually established between ASes of the same size or importance, increasing the chances of other ASes to peer with tier 2 instead of tier 1. Moreover, we can also notice the increase in the relative percentage of peering relationships with level 2, while peering with other levels are decreasing.

Another correlated observation is that out of the peering relationships that originate from a level, the vast majority of them are incident to level 2. In Figure 12.B, each bar that indicates the percentage of peering with that level is further divided into three categories that show the peering between that level and level 1, level 2 and other levels, respectively from top to bottom. Note that peering relationships, as with any other relationship, can occur within the same level or cluster. The middle category is the percentage of peering between that level and level 2. We note that this percentage is largest in all levels, showing indeed that most peering from any level involve a node from level 2. Moreover, this percentage is also increasing over time as can be observed from the figure.

Another property of this logical hierarchy is the connectivity of each cluster across different levels. We use the degree and pairwise distance to characterize this property. The results are presented in Table 7, which shows the degree and distance of each cluster for the different dates. We start by noting that tier 1 providers have an average degree of around 400, which is not surprising given the degree distribution presented in section 3. Since the nodes with the highest degree tend to form the tier 1 set, their average degree is expected to be very high. An important distinction appears in the degree of the customer and provider clusters. The customer cluster has a very low degree,

		1999/12		2000/06		2000/11		2001/05	
Level	Type	degree	distance	degree	distance	degree	distance	degree	distance
1	provider	388.083	1.30	368.632	1.43	331.269	1.70	419.727	1.75
2	provider	15.342	2.43	14.872	2.55	16.501	2.55	16.518	2.55
	customer	2.078	2.60	2.086	2.68	2.137	2.69	2.139	2.67
3	provider	7.412	3.59	7.012	3.71	7.606	3.72	8.056	3.66
	customer	1.634	4.07	1.740	4.07	1.722	4.06	1.736	4.03
4	provider	3.879	4.59	4.471	4.65	4.070	4.67	4.187	4.76
	customer	1.372	5.15	1.349	5.23	1.361	5.16	1.371	5.18
5	provider	3.273	5.16	3.400	5.29	2.583	5.71	2.588	5.21
	customer	1.121	6.35	1.080	5.46	1.167	6.19	1.045	5.74
6	provider	8.500	6.00	2.000	-	-	-	4.000	6.67
	customer	1.000	6.82	1.000	5.45	1.000	6.80	1.062	6.08
7	provider	-		-	-	-	-	-	-
	customer	1.000	5.43	1.000	-	-	-	1.000	5.73
Total	provider	14.702	3.25	16.082	3.21	17.302	3.23	17.192	3.24
	customer	1.770	3.76	1.841	3.69	1.859	3.69	1.861	3.68

Table 7. Characteristics of the degree and pairwise distance within each box in the hierarchical model.

which is slightly above 2 for tier 2 customers and below 2 for all other levels. The high number of multi-homing nodes in the tier 2 customers help explains the relative higher degree for that level. The provider degree is comparatively higher, indicating that providers have much better connectivity than customers. We can also observe that the connectivity degrades as the level of the hierarchy increases. This occurs for both customer and provider clusters.

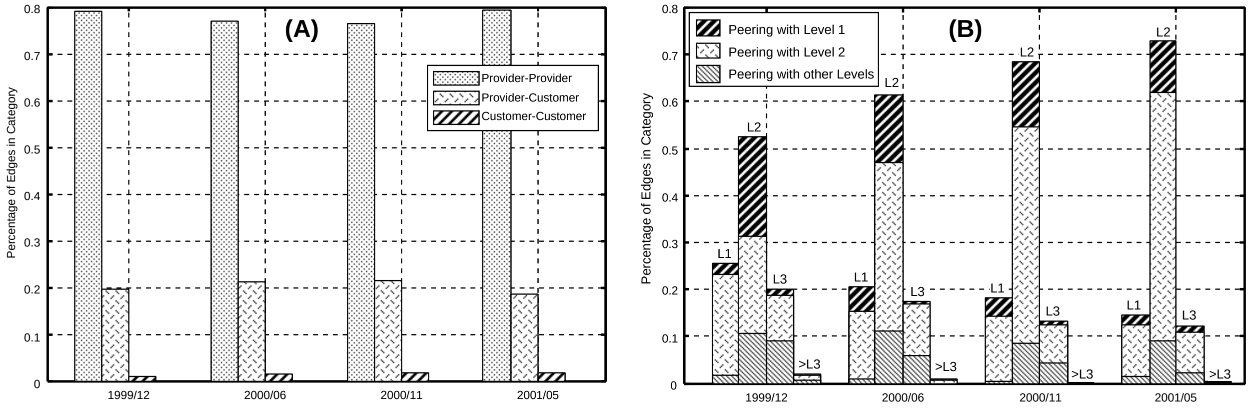

Figure 12. (A) Peering relationships distribution over clusters. (B) Peering relationships between different tiers, and between a tier and tier 2

We now characterize an important characteristic of ASes, namely the *multi-homing* effect, and how is this related to the hierarchical structure. In this paper, we define an AS to be multi-homed if it has more than one customer-provider relationship, thus having more than one provider. Table 8 shows the results of the multi-homing property for different levels of the logical hierarchy. We note that the number of multi-homing is particularly large for the second level for both customers and providers. For tier 2 customers, nearly 73% of the nodes are multi-homed, while tier 2 providers have around 83% nodes multi-homing. This higher value is expected since a provider node is more likely to have more than one provider, since it is important to maintain connectivity in the presence of failures. Note that the percentage of nodes multi-homing decreases with the level of the hierarchy, indicating that nodes located near tier 1 are more likely to be multi-homed. The Multi-homing effect seen in tier 1 is anomalous and is probably an artifact of the classification of left over nodes from Phase 3 of the algorithm. By construction, the core of the tier 1 set is a clique, and multi-homing has no practical significance.

Another observation is that the percentage of nodes multi-homing in a particular level is increasing slowly over time. We can observe this property on both customer and provider boxes, and on different levels. This indicates that AS nodes are willing to increase their connectivity by adding a second provider, which can lead to better service and greater reliability. Multi-homing has a direct impact on route advertisement, and consequently on routing table sizes, and has been observed in.[12] When an AS has multi-homing, route aggregation is either not possible or diminished, leading to multiple path advertisements. Finally, we note that the values for multi-homing in level 5 are

Level	Type	1999/12	2000/06	2000/11	2001/05
1	provider	50.00% (6)	42.11% (8)	57.69% (15)	50.00% (11)
2	provider	79.92% (402)	82.19% (577)	81.63% (662)	83.26% (756)
	customer	72.78% (1604)	72.45% (2235)	73.04% (2734)	73.75% (3186)
3	provider	64.24% (309)	61.58% (319)	67.72% (407)	68.68% (467)
	customer	46.94% (968)	51.06% (1281)	50.45% (1527)	53.40% (1876)
4	provider	23.36% (25)	39.08% (34)	38.60% (44)	41.79% (56)
	customer	29.69% (242)	29.67% (248)	31.86% (354)	33.60% (424)
5	provider	9.09% (1)	10.00% (1)	8.33% (1)	5.88% (1)
	customer	11.11% (11)	6.25% (7)	15.56% (14)	3.82% (6)

Table 8. Multi-homing characteristics of each box in the hierarchical model.

not significant due to the relative small number of nodes in that level and potential logical inference errors.

6. CONCLUSION

In this paper we have investigated the logical structure of the AS graph in the Internet. In particular, we propose a hierarchical organization to place AS nodes into specific clusters forming a tree-like hierarchy. This hierarchy naturally follows from the dominant logical relationships present in the AS graph, which are customer-providers edges. A specific algorithm to classify and place each AS node into the proper level in the hierarchy is described. We present substantial quantitative evidence through different measures that indicates that the classification proposed captures the underlying functional structure by aggregating nodes with similar logical properties.

From the hierarchical model, some interesting and consistent trends of the AS graph can be observed. More specifically, the nature of peering relationships and the multi-homing phenomena. We also show that the relationships that do not fit directly into the framework of model can either be logically explained or have insignificant effect.

We base our observations on data publicly available from Route View and RIPE repositories and use an inference heuristic to extract logical relationships between ASes. Our studies also explores the differences between snapshots of the AS graph taken by these different and independent institutions. We conclude that different snapshots are quite similar in a broader sense (i.e., number of nodes) but differ on specific details (i.e., number of edges). For that reason the union of the snapshots is used throughout this paper.

This simple and precise representation of the AS graph provides a natural starting point for further studies concerning the Internet graph. For example, we are currently working on determining a growth model that is based on the properties of this hierarchical structure. We are also conducing further studies that concentrate on characterizing the precise structural details of each cluster in our hierarchical model. A combination of these results can lead to an algorithm that is able to generate a realistic topology.

REFERENCES

1. Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)." Request For Comment - RFC 1771, Mar. 1995.
2. M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proc. of Sigcomm'99*, (California, USA), Aug. 1999.
3. E. Zegura, K. L. Calvert, and M. J. Donohoo, "A quantitative comparison of graph-based models for internetworks," *IEEE/ACM Transactions on Networking* **5**, pp. 770 – 783, Dec. 1997.
4. A. L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science* (286), pp. 509 – 512, 1999.
5. C. Alaettinoglu, "Scalable router configuration for the internet," in *Proc. IEEE IC3N*, Oct. 1996.
6. L. Gao, "On inferring autonomous system relationships in the internet," in *Proc. of Globecom'00*, (USA), Dec. 2000.
7. R. Govindan and A. Reddy, "An analysis of internet inter-domain toplogy and route stability," in *Proc. of IEEE INFCOM*, (California, USA), Apr. 1997.
8. T. Bu and D. Towsley, "On distinguishing the internet power law topology generators," tech. rep., University of Massachusetts - Computer Science, May 2001.
9. R. J. Adler, R. E. Feldman, and M. S. Taqqu, *A practical guide to heavy tails*, Birkhaüser, 1998.
10. T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms*, ch. 26, pp. 560 – 562. Prentice-Hall, 2 ed., 1998.
11. H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The impact of routing policy on internet paths," in *Proc. of IEEE INFCOM*, (Alaska, USA), Apr. 2001.
12. T. Bu, L. Gao, and D. Towsley, "On the growth of bgp routing table," tech. rep., University of Massachusetts, May 2001.