

Ten Years in the Evolution of the Internet Ecosystem ^{*†}

Amogh Dhamdhere, Constantine Dovrolis
School of Computer Science
Georgia Tech
{amogh,dovrolis}@cc.gatech.edu

ABSTRACT

Our goal is to understand the evolution of the Autonomous System (AS) ecosystem over the last decade. Instead of focusing on abstract topological properties, we classify ASes into a number of “species” depending on their function and business type. Further, we consider the semantics of inter-AS links, in terms of customer-provider versus peering relations. We find that the available historic datasets from RouteViews and RIPE are not sufficient to infer the evolution of peering links, and so we restrict our focus to customer-provider links. Our findings highlight some important trends in the evolution of the Internet over the last decade, and hint at what the Internet is heading towards. After an exponential increase phase until 2001, the Internet now grows linearly in terms of both ASes and inter-AS links. The growth is mostly due to enterprise networks and content/access providers at the periphery of the Internet. The average path length remains almost constant mostly due to the increasing multihoming degree of transit and content/access providers. In recent years, enterprise networks prefer to connect to small transit providers, while content/access providers connect equally to both large and small transit providers. The AS species differ significantly from each other with respect to their rewiring activity; content/access providers are the most active. A few large transit providers act as “attractors” or “repellers” of customers. For many providers, strong attractiveness precedes strong repulsiveness by 3-9 months. Finally, in terms of regional growth, we find that the AS ecosystem is now larger and more dynamic in Europe than in North America.

Categories and Subject Descriptors

C.2.1 [COMPUTER-COMMUNICATION NETWORKS]: Network Architecture and Design

*The AS topology data used in this paper will be made available at www.cc.gatech.edu/~amogh/topology.html

†This work was supported by NSF awards CNS-0347374 (CA-REER) and CNS-0831848.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC’08, October 20–22, 2008, Vouliagmeni, Greece.

Copyright 2008 ACM 978-1-60558-334-1/08/10 ...\$5.00.

General Terms

Measurement, Economics

Keywords

Internet Topology, Measurement, Evolution, Economics

1. INTRODUCTION

The Internet, as a network of Autonomous Systems (ASes), resembles in several ways a natural ecosystem. ASes of different sizes, functions, and business objectives form a number of *AS species* that interact to jointly form what we know as the global Internet. ASes engage in competitive transit (or customer-provider) relations, and also in symbiotic peering relations¹. These relations, which are represented as inter-AS logical links, transfer not only traffic but also economic value between ASes. The Internet AS ecosystem is highly dynamic, experiencing *growth* (birth of new ASes), *rewiring* (changes in the connectivity of existing ASes), as well as *deaths* (of existing ASes). The dynamics of the AS ecosystem are determined both by external “environmental” factors (such as the state of the global economy or the popularity of new Internet applications) and by complex incentives and objectives of each AS. Specifically, ASes attempt to optimize their utility or financial gains by dynamically changing, directly or indirectly, the ASes they interact with. For instance, the objective of a transit provider may be to maximize its profit, and it may approach this goal through competitive pricing and selective peering. The objective of a content provider, on the other hand, may be to have highly reliable Internet access and minimal transit expenses, and it may pursue these goals through aggressive multihoming and an open peering policy.

Our study is motivated by the desire to better understand this complex ecosystem, the behavior of entities that constitute it (ASes), and the nature of interactions between those entities (AS links). How has the Internet ecosystem been growing? Is growth more important than rewiring in terms of the formation of new links? Is the population of transit providers increasing (implying diversification of the transit market) or decreasing (consolidation of the transit market)? Given that the Internet grows in size, does the average AS-path length also increase? Which ASes engage in aggressive multihoming? What is the preferred type of transit provider for AS customers? Which ASes tend to constantly adjust their set of providers? Are there regional differences in how the Internet evolves? These are some of the questions we ask in this paper.

Understanding the evolution of the Internet ecosystem is important for several reasons. First, we believe that there is a need to

¹We refer to “settlement free interconnection” as a “peering relation” and “paid transit” as a “customer-provider” relation.

develop *bottom-up* models of Internet topology evolution that capture the complex interactions between the constituent entities. As such, it is necessary to study the differences between the types of ASes that form this ecosystem in terms of business function and incentives. Second, understanding the evolution of the Internet is critical for studying the performance of protocols and applications over time. For instance, to answer the question “How will BGP perform 10 years from now?” we first need to answer the question “How will the Internet look 10 years from now?”. Third, there is much recent interest in generating synthetic AS graphs for simulation and analysis. A study of the evolution of the Internet can provide valuable inputs to such topology generators, such as the types of ASes in the Internet and their topological and behavioral properties over time. Finally, in light of the recent interest in re-designing the Internet with “clean-slate” approaches, it is crucial to understand how the existing Internet has evolved. Doing so could help us identify new architectures and mechanisms that have an intrinsic capability to evolve towards desirable economic, reliability and performance conditions.

There is an extensive literature on AS-level topology measurement and modeling (reviewed in detail in Section 8). A large portion of that literature, however, takes a graph-theoretic perspective, viewing all ASes as nodes in a graph and all inter-AS relations as edges, without considering the type of relation (customer-provider versus peering) or the role of the participating ASes (customer versus provider). Viewing all ASes as the same type of node ignores the major differences in the function and objectives of different ASes. Further, even though most of the previous work on AS-level topology modeling mentions the terms ‘evolution’ or ‘dynamics’, the main focus has been on measurements and modeling of growth, ignoring rewiring. The latter is very important, however, as it represents the attempt of individual ASes to optimize their connectivity. Finally, most of the earlier work on AS-level topologies has focused on macroscopic properties and metrics, such as the degree distribution, the clustering coefficient or the graph diameter, without considering the local policy and semantics of inter-AS relations. The latter are very important as they control the flow of traffic and value in the AS ecosystem.

In this paper, we attempt to measure and understand the evolution of the Internet ecosystem during the last decade (1998-2007). We propose a method to classify ASes into a number of types depending on their function and business type, using observable topological properties of those ASes. The AS types we consider are large transit providers, small transit providers, content/access/hosting providers, and enterprise networks. We are able to classify ASes into these AS types with an accuracy of 80-85%. We focus on *primary* inter-AS links, meaning links that are used under “normal operating conditions”, to distinguish with backup links that appear under failure conditions or routing convergence. We also consider the semantics of inter-AS links, in terms of customer-provider (CP) versus peering (PP) relations, and distinguish between the customer, provider and peering role of an AS in each relation. Unfortunately, we find that the available historical datasets from RouteViews and RIPE are *not sufficient to infer the population and evolution of peering links*. So we restrict the focus of this study to the evolution of the population of AS types and of customer-provider links.

The rest of this paper is structured as follows. In Section 2, we describe the data collection and filtering methodology. In Section 3, we focus on the evolution of the global Internet. In Section 4, we present a classification scheme of ASes into four AS types based on their expected business function. Then, we examine the evolution of each AS type at a global scale as well as regionally. In Sections 5

and 6, we investigate the evolution of customer-provider relations in the Internet, from the perspective of the customer and provider, respectively. In Section 7, we present some results on the evolution of the Internet peering ecosystem. These results should be viewed as “conjectures” because of the limitations in detecting the complete set of peering links. We discuss related work in Section 8 and conclude with a summary of our main findings in Section 9.

2. DATASETS AND METHODOLOGY

A study of the evolution of the Internet ecosystem needs frequent snapshots of the AS-level Internet topology, annotated with policy information for each link. Given that such historical information is not available, we have to rely on measurement and inference, collecting data from multiple sources and considering the limitations of each dataset. This section describes the datasets we use and the subsequent filtering and validation processes.

We collected BGP AS-paths from BGP table dumps obtained from the two major publicly available repositories at RouteViews [30] and RIPE [29]. The RouteViews collection process started in November 1997, providing an invaluable resource in the past ten years. The first RIPE collector became active in October 1999. We rely only on these two repositories because no other source of topological/routing data (routing registries, traceroutes, looking glass servers, etc.) provides historical information. Note that the use of AS-paths has been shown to be inadequate to expose the *complete Internet topology* [9, 10, 18]. In particular, even though most ASes are detected, a significant fraction of peering and backup links at the edges of the Internet are missed [8, 18, 37]. In fact, it has been estimated that there are at least 40% more peering links in the Internet than those obtained from AS-paths [8, 10]. We are well informed of these limitations, which are further exposed later in this section. There are, however, three important points to consider. First, *we do not aim to detect backup links*; instead, we are only interested in primary Internet links, used most of the time (as opposed to backup links that are only used upon failures or overload conditions). We describe later how to avoid backup links in the data filtering process. Second, *the main focus of this evolutionary study is customer-provider links*. As we show later in this section, the available monitors from RouteViews and RIPE are not enough to detect all peering links or the births and deaths of those links. Third, even though missing links can be detrimental for complex inference applications (such as AS path prediction or BGP root-cause analysis), it has been shown recently that they are *less critical in topology inference* [38].

Filtering of backup and transient links: Next, we describe how to only detect primary links, avoiding backup links and false AS-paths that often appear during BGP convergence. First, note that short-term failures and routing transient events can “confuse” an evolutionary study, misinterpreting link disappearances and appearances due to transient failures as link deaths and births respectively. For instance, suppose that the primary link l_p between AS-x and AS-y fails at time t_1 , causing the activation of a backup link l_b between AS-x and AS-z. l_p is repaired at t_2 and the connectivity returns to its original state. Since we focus on primary links, our goal is to ignore the transient event during (t_1, t_2) and to *not* detect l_b . On the other hand, a change of routing policy that exchanges the role of links l_p and l_b (so that l_b becomes the primary link) should be detected as the death of l_p and the simultaneous birth of l_b .

To achieve the previous objective we follow the “majority filtering” approach described next. Note that a *snapshot*, in the following discussion, does not refer to a time instant but to a period of 21 days. During a certain snapshot, we collect at N different times the unique AS-paths that are exported from all active RouteViews

and RIPE monitors. The period between these successive *samples* is T_s , with $N T_s=21$ days. Then, we keep only those AS-paths that appear in the majority of the samples and ignore the rest. This process is designed to filter out links that appear due to routing transient events, as well as due to “hard” failures of interdomain links (e.g. due to router crashes or fiber cuts). Routing transients typically persist for less than a few hours, while it is reasonable to expect that hard failures are repaired within 10 days. In each of these cases, the majority filtering rule successfully filters out the transient links.² Note that if a certain link X-Y is used as primary in one AS path but as backup in another path, it will be included in our snapshot.

To select an appropriate value of N , we do the following. We collect all visible AS-paths for each day of January 1998. Next, we divide the month into N blocks of the same duration, and collect the set of visible AS-paths from a randomly selected instant in each of the N blocks. Then, we perform majority filtering, considering only AS-paths that appear in the majority of the N samples. Finally, we measure the number of visible AS links. We vary N from 1 to 10, and repeat the previous process multiple times for each value of N . As N increases, the average number of visible links decreases (from about 5850 to 5725 during that month) because fewer backup links become visible. Additionally, the variability in the number of visible links decreases. We observe that $N=5$ results in about the same average as higher values of N , and reasonably low variance (standard deviation of 12 links). In the rest of this study, $N=5$ samples.

The trade-off behind the selection of the snapshot duration (21 days in our study) is explained next. If the snapshot duration is too long (say more than a month), then we may miss several birth-death (or death-birth) transitions of the same link. On the other hand, if the snapshot duration is too small (say a few days), then the majority filtering mechanism may not be able to filter out backup links that appear during long-lasting failures such as fiber cuts. Finally, a new snapshot is collected every three months, providing us with 40 snapshots (10 years) from January 1998 to October 2007.

Variations in the number of active monitors: Another issue we need to consider is that the number of BGP monitors in both RouteViews and RIPE has been increasing significantly over the last ten years, from about 10 in 1997 to almost 400 at the end of 2007. The increase in the number of monitors has been less than 20% in 35 out of the 39 pairs of successive snapshots. As the number of monitors increases, some previously existing links may become visible for the first time at a certain snapshot. How do we distinguish those first appearances of existing links from genuine link births? Similarly, sometimes monitors are removed. How do we distinguish between the disappearance of existing links from genuine link deaths? Also, can we bound the estimation error in the number of link births and deaths between each pair of successive snapshots?

To answer the last question we perform the following analysis. Let the set of monitors at snapshots T_1 and T_2 be \mathcal{M}_1 and \mathcal{M}_2 respectively. Let \mathcal{L}_1 and \mathcal{L}_2 be the set of links observed at T_1 and T_2 , respectively. $\mathcal{L}_2 - \mathcal{L}_1$ is the set of new link appearances between T_1 and T_2 . This is our estimate for the set of new link births. This set includes the links that were genuinely born between T_1 and T_2 , but it may also include an error term that is the set of links that were present at T_1 but became visible at T_2 due to the monitor set increase. To derive an upper bound for the latter, we do the following.

First, determine the set of links \mathcal{L}'_2 that would be observed at T_2

²A similar process was used by Dimitropoulos et al. [11], but considering an AS-path only if it appears in *all* N samples.

using the set of monitors that were common between \mathcal{M}_1 and \mathcal{M}_2 , i.e., $\mathcal{M}_1 \cap \mathcal{M}_2$. The set $\mathcal{L}'_2 - \mathcal{L}_1$ (where $\mathcal{L}'_2 - \mathcal{L}_1 \subseteq \mathcal{L}_2 - \mathcal{L}_1$) includes links that were definitely born between T_1 and T_2 , and hence it gives a *lower bound* on the number of actual link births. On the other hand, the number of links in the set $(\mathcal{L}_2 - \mathcal{L}_1) - (\mathcal{L}'_2 - \mathcal{L}_1)$ is an upper bound for the error between the estimated and actual number of link births. So, the *worst case relative error* (WCRE) in the number of link births between T_1 and T_2 is:

$$\text{WCRE} = \frac{|(\mathcal{L}_2 - \mathcal{L}_1)| - |(\mathcal{L}'_2 - \mathcal{L}_1)|}{|(\mathcal{L}'_2 - \mathcal{L}_1)|} \quad (1)$$

We measured the WCRE for every pair of snapshots. In 30 out of the 39 snapshots pairs, the WCRE is less than 10%. For all but one pair, the WCRE is less than 20%. In the remainder of this paper, we omit the pair of snapshots for which the WCRE was larger than 20% (Jan-Apr 2000). We also measured the WCRE separately for customer-provider (CP) links and peering (PP) links. Unfortunately, the WCRE is very high for peering links and in 9 out of 39 snapshots it is greater than 100%. On the other hand, the WCRE for CP links is quite low, and for all except one pair of snapshots (Jan-Apr 2000), it is less than 10%.³

The previous analysis considers the effect of an increased set of monitors on the measurement of link births. A similar problem occurs while measuring link deaths, as some monitors are occasionally disconnected temporarily or permanently from the RouteViews and RIPE collectors. We performed a similar analysis to determine the effect of monitor deaths on the estimated number of link deaths. We find that the WCRE in the estimated number of link deaths is less than 10% for 37 out of the 39 snapshot pairs.

The previous WCRE analysis showed that, even though we can estimate well (within 10%) the link births/deaths of CP links, we do not get a reasonable accuracy for the link births/deaths of PP links. This is a negative but significant result, which should be considered by future studies that rely on RouteViews and RIPE topological data. It also implies that the conclusions of several previous topological studies should be re-examined.

Sensitivity of population counts to number of monitors: We next examine the visibility of CP and PP links, as well as of ASes, when we vary the number of used monitors. Consider first the population of ASes. Let n_{AS} be the set of visible ASes if we use all available monitors at a given snapshot. We then randomly select a fraction f of the available monitors, and determine the population of ASes that is visible using that subset of monitors. We repeat this experiment 100 times, and determine for each run the number of ASes visible with a fraction f of the available monitors $n_{AS}(f)$. Figure 1 shows the median, 10th and 90th percentile values of the ratio $n_{AS}(f)/n_{AS}$ for the snapshot Jan 2007, together with the corresponding ratios for the populations of CP links and PP links. We repeated this analysis for all snapshots, and the results are quantitatively similar across time, without any noticeable trends.

Note that the number of visible ASes is strongly insensitive to the number of available monitors. Even with 10% of the monitors we practically see the same set of ASes that is visible with all monitors. The fraction of CP links is also insensitive to the number of available monitors, as long as we use more than 60-70% of the available monitors in the given snapshot. So, we expect that a 10-20% increase in the number of available monitors across successive snapshots will not cause a significant variation in the number of visible CP links. The situation is very different with PP links however. The fraction of visible PP links increases roughly linearly with the fraction of used monitors. This means that if we had

³Note that the WCRE is calculated for every pair of snapshots, and so it does not accumulate over time.

more monitors we would probably see significantly more PP links. So, *the estimated population size of PP links should be viewed as lower bound on the actual population size*. Similar observations were recently reported by Oliveira et al. [26].

The previous observations have two major implications. First, on the positive side, *it appears that the RouteViews and RIPE historical datasets contain enough monitors to detect the ASes and CP links in a robust manner*. Even though we cannot be certain that we see *all* ASes or CP links, we at least have evidence that these populations would not differ by a large number if we had more monitors. Second, on the negative side, *it is clear that the RouteViews and RIPE datasets are not sufficient to detect the population or the birth/death rates of PP links*. Consequently, in the rest of the paper we focus on the evolution of CP links. When we present some results for PP links, the reader should recall that those figures are lower bounds on the actual number of PP links.




Figure 1: Visibility of ASes, CP and PP links as a function of the number of monitors used in a snapshot.

Policy inference: After collecting and filtering the data as described earlier, the final data processing step is to use the AS-paths in each snapshot (those that passed the majority filtering process) to infer the underlying AS topology and the relationships between adjacent ASes. For this purpose we use the well-tested algorithm described by Gao in [16]. Despite the significant follow-up work on AS relationship inference [11, 33], we prefer Gao’s algorithm because of its ability to infer relationships using only observed AS paths, without any additional information such as data from routing registries or active probes. Comparison studies for the accuracy of related algorithms in [33] and [16] have shown that Gao’s algorithm is more accurate in identifying peering relationships. Further, a recent study [38] showed that the AS relationship inferences made by this algorithm are quite stable with respect to variations in the observed AS-paths. Gao’s algorithm results in four types of AS relationships: Customer-Provider (CP), Peering (PP), Sibling, and Unknown. We ignore the last two categories, as they account for less than 2% of the visible links in any snapshot.

Finally, the AS topology and relationship matrix provide an annotated graph for each snapshot. The differences between successive snapshots show the evolutionary events of link and node births and deaths, which form the core of the analysis in the following sections. Note that if a certain link has changed role at some snapshot (say from CP to PP), we view that event as the death of a CP link and the simultaneous birth of a PP link between the corresponding

ASes. The reader may be wondering about the frequency of link type changes, from CP to PP or the opposite. Even though we cannot answer this question in a definite manner (due to the visibility problem with PP links), we measured that 9% of the PP links in a snapshot become CP links in the next snapshot (This number is the average over all pairs of snapshots). The fraction of CP links that become PP links appears to be much less (1%) but that is probably due to the poor visibility of PP links. Also, these changes are not cumulative, as we run the relationship inference algorithm separately for each snapshot.

3. GROWTH AND REWIRING TRENDS

We first examine the evolution of some major characteristics of the global Internet.

Growth of ASes and inter-AS links: Figure 2 shows the number of ASes and inter-AS links in each snapshot. Due to the previously discussed issues with measuring PP links, we only count the number of CP links in each snapshot. A first observation is that, despite the economic recession of 2001-03 and the well documented turmoil in the telecom market, *the Internet AS-level topology has been increasing in size over the last ten years*. Second, it appears that the Internet has gone through two distinct growth phases so far: *an initial phase, up to mid-2001, in which the Internet grew exponentially in terms of the number of ASes and links* (of the form $y = a * e^{bx}$). Then, the growth process switched to linear for both the number of ASes and links (of the form $y = ax + b$). We find that the number of ASes from 1998 to mid-2001 can be modeled as $y = 3150 * e^{0.094x}$, where x is the snapshot number ($x = 0, 1, \dots$). In the last six years, the number of ASes can be modeled as $y = 2537 + 604x$. Regarding the number of CP links, the corresponding functions are $y = 5462 * e^{0.102x}$ and $y = 1499x - 35$. Each of the previous regression formulae gives a correlation coefficient that is at least 99%. To eliminate the possibility that this trend shift is an artifact of the measurement infrastructure (e.g. the changing set of monitors), we measured the number of visible ASes and CP links with a set of monitors that remained the same in the last ten years. The results, even though revealing a lower number of links, still show a trajectory change from exponential to linear in mid-2001. Huston [19] observed a similar trend shift in the number of ASes (but not CP links) around mid-2001.

To determine the boundary at which the trajectory shifted from exponential to linear, we perform the following test. We assume that the number of CP links and ASes can be modeled as $y = a e^{bx}$ when $x \leq z$ and $y = ax + b$ when $x > z$. We then compute the value z_{min} that minimizes the total sum-of-squares error (SSE) for the above regression formula. z_{min} is our estimate for the snapshot where the growth trajectory changed from exponential to linear. It appears that the exponential phase lasted for the first 15 snapshots for ASes and 16 snapshots for CP links, *ending in mid/late 2001*. Figure 2 also shows the exponential and linear regression curves for the number of ASes and CP links.

Evolution of CP link count (and lower bound estimates of PP link count): Next, we distinguish between CP and PP links, and examine the growth trends separately for these two link types. We emphasize again that the number of PP links we report here should be viewed as a *lower bound* on the actual number of peering links. Figure 3 shows the number of CP and PP links, as well as their fractions, over time. Both link types have been increasing in absolute numbers. As shown earlier, the number of CP links shows an initial exponential growth followed by a linear growth after 2001. Modeling the growth of PP links is difficult with the given measurements. It appears, however, that that growth process has followed a different trajectory than that of CP links.




Figure 2: Evolution of the number of ASes and CP links. The regression curves are also shown.

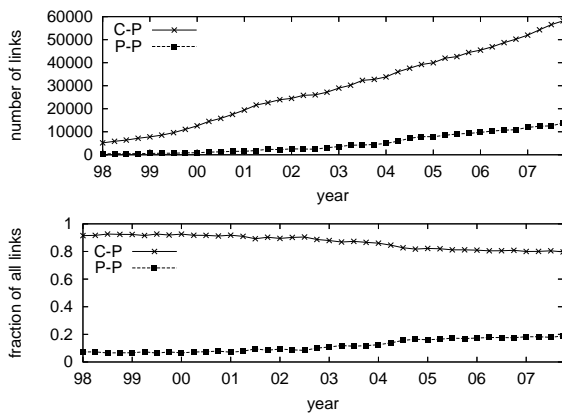


Figure 3: Evolution of CP and PP links in absolute numbers and as a fraction of the total number of links.

The bottom panel in Figure 3 shows the fraction of CP and PP links. *The fraction of PP links has been increasing steadily after 2001*, even though the growth rate of CP links is larger than that of PP links. The reason is that the relative increase rate of PP links is larger than that of CP links. Given that we probably underestimate the number of PP links, *the fraction of PP links at the end of 2007 is at least 20%*.

Evolution of AS-path length and multihoming trends: Next, we investigate the evolution of the average AS-path length (after removing AS-path prepending). We do so by calculating the average length of the AS-paths observed in each snapshot. The upper panel in Figure 4 shows that *the average path length has remained practically constant (at 4.2 AS hops) over the last 10 years*. This is interesting, given the significant growth of the underlying network. The fact that the average AS-path length has remained constant points to a *densification process* that increases the average degree of ASes at a sufficiently high rate to keep the average AS-path length constant. Indeed, the upper panel of Figure 4 shows that the average AS degree, counting only CP links, has increased consistently over time, from 3.2 links to 4.3 links per AS. The median degree (not shown) is dominated by small networks that have just 1 or 2 providers, and hence it does not show an increasing trend. This densification process has also been studied by Leskovec et al. [22],

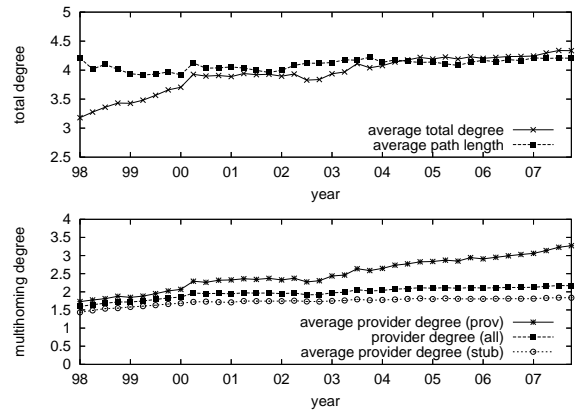


Figure 4: Evolution of average AS degree, AS-path length, and multihoming degree.

who observed that the effective diameter⁴ of the AS graph *slowly decreases* with time. Earlier modeling work, such as the preferential attachment growth model of Albert and Barabasi [1], predicted an average path length that grows slowly with the size of the network ($O(\ln \ln n)$), when a newly attached node has at least two edges. Such a growth model would result in an increase in the average path length from 4.2 to 4.7 over the last 10 years, contrary to the constant average path length of 4.2 that we observed. An interesting possibility is that ASes choose their providers in a “distance-aware” manner, so that the AS-path length from/to their major Internet sources/destinations remains practically constant with time.

A plausible explanation for the densification of the Internet is the increasing popularity of *multihoming* for economic, reliability and performance reasons. The bottom panel of Figure 4 shows the average *multihoming degree*,⁵ defined as the number of providers of a given AS, for two broad classes of ASes: *stubs* (i.e., ASes that never had customers during their observed lifetime), and *non-stubs* (the rest of ASes). We find that *the average multihoming degree has been increasing in both classes*. Non-stubs, however, have been increasing their average multihoming degree much faster than stubs (from 1.5 to about 3.5), in particular after 2003. This may be because non-stubs, which are typically content/access/hosting/transit providers, attempt to optimize their connectivity, and at the same time improve their reliability, by multihoming to several upstream transit providers. For many stubs, on the other hand, one or two (primary) transit providers is often enough.

Growth versus rewiring: Next, we seek to understand the relative significance of *growth versus rewiring*. Growth refers to the addition of new ASes in the network (together with their corresponding links), while *rewiring refers to changes in the connectivity of existing ASes*. Specifically, we focus on the number of CP link births due to AS births (growth) versus CP link births due to rewiring. We also look at the number of CP link deaths due to AS deaths versus CP link deaths due to rewiring. The top panel of figure 5 shows, for each pair of snapshots, the number of CP link births due to AS births and due to rewiring. Initially, the CP link births due to AS births and rewiring were comparable in number. Since 2001, however, we find that the number of CP link births due to internal

⁴The effective diameter of a graph is the minimum value of d such that at least 90% of the connected node-pairs are at distance at most d . A smoothed version of this metric is used in [22].

⁵Multiple physical links between two ASes are counted as a single inter-AS link.




Figure 5: Evolution of the number of CP link births (and deaths) due to node births (and deaths) versus rewiring.

rewiring has increased much faster than that due to AS birth. Currently, *around 75% of link births are associated with existing ASes (rewiring)*. A similar analysis, shown in the bottom panel, shows that the number of CP link deaths due to rewiring is significantly higher than that due to AS deaths. About 80% of the link deaths are due to rewiring and this fraction is increasing. These observations are important for two reasons. First, most of the literature on AS topology modeling has focused on growth, ignoring rewiring. Second, rewiring represents the effort of individual ASes to optimize their performance, reliability, profitability or other objectives. An intriguing possibility is that rewiring implies that the Internet, as a multi-agent and self-organized system, attempts to optimize a certain, still unknown, global objective in a distributed manner. This possibility has also been discussed by Chang et al. [6].

Given the increasing significance of rewiring, we next focus on the births and deaths of links between existing nodes in two successive snapshots. Let G_1 and G_2 be the graphs representing the primary AS topology in two consecutive snapshots. We construct G'_1 from G_1 by removing all nodes that are not present in G_2 ; similarly construct G'_2 from G_2 . Note that G'_1 and G'_2 have the same set of nodes. Let E'_1 and E'_2 be the set of links in G'_1 and G'_2 respectively. We use the following graph-level metric, referred to as **Jaccard Distance**, to quantify the rewiring between G'_1 and G'_2 .

$$s(E'_1, E'_2) = \frac{|(E'_1 - E'_2) \cup (E'_2 - E'_1)|}{|E'_1 \cup E'_2|} \quad (2)$$

Note that $s(E'_1, E'_2)$ captures both link births and deaths between the two snapshots. The Jaccard distance thus quantifies the difference between the sets of links in two consecutive snapshots. For example, a Jaccard distance of 0.5 indicates that 50% of the links seen in the two snapshots were either born before the second snapshot or died after the first.

We calculate the Jaccard distance separately, first, on the CP graph where the customer is a stub, and second, on the CP graph where the customer is a non-stub. Figure 6 shows these metrics for each pair of snapshots over the last 10 years. We find that the Jaccard distance is much smaller for the CP graph where the customer is a stub, as compared to the CP graph where the customer is a non-stub. This indicates that *non-stubs have consistently been more aggressive than stubs in changing their upstream connectivity*. We further investigate this effect after proposing a finer classification of AS types in the Internet in the next section.




Figure 6: The Jaccard distance for CP links where the customer is stub versus non-stub.

4. EVOLUTION OF AS TYPES

When we think of the Internet as a graph, it is important to recognize that not all nodes are the same. ASes connect to the Internet with different requirements and business interests, and hence optimize their connectivity in different ways [13]. The topology changes that we observe represent the outcome of a complex multi-constraint optimization process that individual ASes conduct.

AS classification scheme: We propose a simple classification scheme for ASes according to their business type. The initial classification consists of the following five AS types.

Enterprise Customers (EC) represent various organizations, universities and companies at the network edge that are mostly users, rather than providers of Internet access, transit or content. Typically, ECs do not have AS customers.

Small Transit Providers (STP) are often regional ISPs that provide Internet access and transit services. STPs aim to maximize their customer base in their geographical area and to reduce their upstream transit costs through *selective peering* with other regional ISPs. STPs often peer selectively rather than openly to avoid peering with ASes already in their customer tree, or ASes that are likely to become customers at a future time. We count national and academic/research transit networks also as STPs.

Large Transit Providers (LTP) are international ISPs with a large footprint, both in terms of number of AS customers and geographical presence. LTPs aim to maximize their customer base, peering with other ASes only when it is necessary to maintain reachability (*restrictive peering*).

Access/Hosting Providers (AHP) are ISPs that offer Internet access (e.g., DSL, cable modem, dial-up, leased lines) and/or server hosting. Their access customers can be residential users or enterprises that do not have AS numbers, while their server hosting customers are content/service providers that also do not have AS numbers⁶. AHPs often engage in selective peering to minimize the transit costs paid to their upstream providers.

Content Providers (CP) are not in the business of offering Internet transit or access. Instead, their revenues result from providing content that users pay for. CPs aim to minimize transit costs, and so often have *open peering* policies.

Similar classifications have been proposed in previous work. Chang et al. [5] classified ASes (for the purposes of determining interdo-

⁶A limitation of AS topologies derived from BGP tables is that they include only the organizations that have AS numbers.

main traffic matrices) into “web hosting”, “residential access” and “business access”. Dimitropoulos et al. [12] classified ASes into large and small ISPs, customer networks, universities, Internet exchange points and network information centers. We chose the previous five AS types based on the terminology used in discussions on the NANOG mailing list and in W. Norton’s white papers [25].

Note that the difference between LTPs and STPs is quantitative, as both AS types have the same business function. LTPs are basically the major ISPs that are often referred to, rather informally, as “tier-1” transit providers. The “tier-1” label is often associated with 10-20 ASes. We choose to be more inclusive, defining as LTPs the top-30 ASes in terms of the average number of customers during the time period in which an AS was seen in the last decade. That average is larger than 140 AS customers for the LTPs in our datasets.

This leaves us with around 27,000 ASes (in the latest snapshot) that cannot be classified manually. Instead, we first pick a training set of 50 ASes for each of the remaining four AS types (EC, STP, AHP and CP) that are definitely of the corresponding type (based on information obtained from their webpages). For ECs, we pick well-known universities and corporations. For STPs, we choose transit providers that are mostly regional in terms of their coverage and customer size. For CPs and AHPs, we pick well-known content providers, hosting sites, and large broadband/dial-up residential/business access ISPs. Next, we observe the topological properties of the ASes in each training set, in terms of the average number of customers C , providers P , and peers R for that AS in the last decade. We found significant overlap in the number of providers among the four AS types, and so we do not rely on that metric. On the other hand, the number of customers and peers (C , R) allows us to distinguish between ECs, STPs and CPs. Unfortunately, we are unable to distinguish CPs from AHPs. These two AS types, even though have different business roles, largely overlap in terms of both C and R . So, in the rest of the paper we merge these two AS types in what will be referred to as **Content/Access/Hosting Providers (CAHPs)**. Figure 7 shows the average number of customers and peers for ASes in the four training sets. Most ECs have zero customers and peers, and they are not shown in this graph.

The next step is to determine a set of boundaries in the two-dimensional (C , R) space that separate the training sets of the four AS types with the minimum number of misclassifications. We apply the well known machine learning technique of *decision trees* on the training samples to obtain the following C and R coordinate boundaries for each AS type:

$$\text{EC: } C < 2.1, R \leq 1$$

$$\text{STP: } 2.1 \leq C < 140, R < 3.5 \text{ and } 33.1 \leq C < 140, R \geq 3.5$$

$$\text{LTP: } C \geq 140$$

$$\text{CAHP: } C < 2.1, R > 1 \text{ and } 2.1 \leq C < 33.1, R \geq 3.5$$

Based on the previous boundaries, we next use the average C and R values of each AS (measured over the snapshots in which that AS was present in the ten-year dataset) to classify it into one of the four AS types. Note that the AS types we consider are quite distinct from each other in terms of their function and business goals. It is thus reasonable to expect that ASes do not change from one AS type to another during their lifetime. To examine this hypothesis, we performed the following test. We rerun the decision tree algorithm to classify each AS using a two-year dataset from 2006 and 2007. We then compared this more recent classification with that based on the ten-year dataset. We found that only 3% of the ASes that appear in both datasets were classified differently. In most of these cases, it appears that the classification change was due to a large shift in the customer and peer degrees of that AS. For

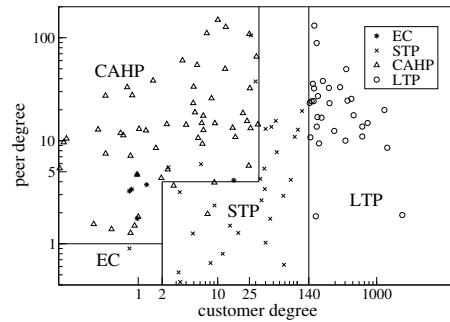


Figure 7: Coordinate boundaries for the four AS types we consider.

example, AS-1 has a large average customer degree over the ten-year dataset and is classified as an LTP. However, in the two-year dataset it has a customer degree of 0, and is classified as an EC. AS1 was originally owned by Genuity Inc., a large global ISP. In 2004, Genuity sold AS-1 to Level3 Communications, also a global ISP. Level3 does not use that AS number for its transit services, and this is why that AS has no customers in the last couple of years.

To evaluate the accuracy of the previous classification scheme, we perform the following. We select a random sample of 150 ASes (50 ECs, 50 STPs and 50 CAHPs), and mix these samples to remove any information about the classification of these ASes (to avoid any subjective bias in the validation process). Then, we use information from WHOIS servers and the webpages of those ASes to infer their main business function. If the actual business function does not match the classification produced by our algorithm, we count that AS as a misclassification. We find that the classification accuracy for ECs is 78%. The errors in this category are due to some residential access providers that are classified as ECs because they have no AS customers and no peers. The accuracy for STPs is 86%. The errors here are due to ASes that mainly offer content hosting services. These providers have few AS customers and a small number (or none) of peers and hence they get classified as STPs. The classification accuracy for CAHPs is 86%. The errors in this case are mostly due to some academic/research backbones that get classified as CAHPs due to their large number of peers. Dimitropoulos et al. [12] reported a similar accuracy figure (78%) for their AS classification scheme.

Population trends for each AS type: Figure 8 shows the population of each AS type over the last ten years. These curves show two distinct phases, similar to the global growth trends observed in Section 3, with a change of slope around 2001. The STP population shows a small growth rate (increase by factor of 1.23 over the last six years). The LTP population remains almost 30 by definition. The EC population shows a strong growth trend (increase by factor of 2.33 in the last six years), contributing most of the growth in the number of ASes. The CAHP population, even though much smaller in absolute numbers than ECs, has also been growing significantly (increase by factor of 1.6 in the last six years). ECs and CAHPs represent the periphery of the network, where the users and content reside. If we judge by the population of this AS type, *the Internet edge grows at a significant and stable pace*. On the other hand, LTPs and STPs represent the core of the Internet. Even though the STP population was growing significantly before 2001, their growth rate in the last few years has decreased. This may be an indication that the number of transit providers is stabilizing.




Figure 8: Evolution of the population of AS types.




Figure 9: Regional distribution of AS types over time.

Geographical trends for each AS type: To classify ASes into broad geographical regions, we use the “registry” field of the corresponding WHOIS entries. Figure 9 shows the fraction of ASes of each AS type that were registered to ARIN (North America) and RIPE (mostly Europe). The other registries (APNIC, LACNIC and AFRINIC) account for the remaining small fraction, and are not shown here. Interestingly, we see that the population of ECs in the two continents (NA and Europe) converges. It is likely that in the next few years there will be more ECs registered in Europe than in North America. This has already happened in the case of STPs, and the number of STPs is now slightly higher in Europe. LTPs, though, are mostly still based in North America. On the other hand, the fraction of CAHPs in Europe has always been higher than in North America, probably because of the many regional access providers (several per country) in Europe. These trends imply that *the Internet market, in terms of the number of access/hosting, transit and content providers will soon be larger in Europe than in North America*, if this is not happening already.

Rewiring activity for each AS type: The differences in the business function and incentives of the four AS types could also appear in their rewiring activity. To measure this quantity between a pair of snapshots, we calculate the Jaccard distance for the set of CP links of each AS. We then compute the average Jaccard distance for all ASes of the same AS type. The top panel of Figure 10 shows these averages over time. We see that, clearly, *ECs show the lowest*

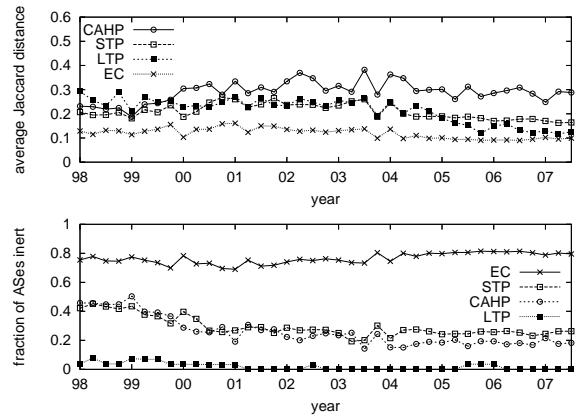


Figure 10: Rewiring activity and fraction of inert ASes for each AS type.

rewiring activity throughout the last ten years. *STPs and LTPs have similar rewiring activity, while CAHPs exhibit the highest rewiring activity especially since 2001*. CAHPs rewire their CP links frequently, as they attempt to minimize their transit costs and provide good performance/reliability to their customers.

A related metric is the fraction of nodes in each AS type that are *inert*, meaning that they do not undergo any change in their set of CP links between two successive snapshots. The bottom panel of Figure 10 shows the fraction of inert nodes for different AS types over time. We find that the fraction of inert ECs increased slightly with time, from 74% in 2001 to 80% currently. This implies that ECs at the network edge are becoming increasingly stable with respect to the connectivity to their providers. The fraction of inert STPs has stayed almost constant since 2001 (between 25% and 30%). We examined the set of STPs that are inert in every pair of snapshots since 2001, and found that several of the inert STPs are national monopoly providers or research and educational backbone networks. Such STPs have a fairly stable customer base, and do not have the incentive to constantly optimize their connectivity. As expected, the fraction of inert LTPs is very low and it approaches zero, because large transit providers have a constant churn in their customers. The most interesting trend is that the fraction of inert CAHPs has decreased significantly, from 46% to 18%. This again suggests that such access/hosting/content providers have an incentive to constantly optimize their connectivity.

5. EVOLUTION OF CP RELATIONS: CUSTOMER-SIDE PROPERTIES

Number of providers per AS type: Figure 11 shows the average number of providers per customer (or the average multihoming degree) for each AS type. The median number of providers (not shown) shows similar trends. *The multihoming degree for ECs has increased very slowly over the last decade (from 1.5 to 1.9), and is almost constant since 2001*. On the other hand, *the multihoming degree for STPs has increased significantly (from 1.9 to 3.9), LTPs⁷ (from 2 to 5.5), and CAHPs (from 2.1 to 8.2)*. The dramatic increase in the multihoming degree of content/hosting/access providers and

⁷Tier-1 ASes are commonly attributed as not having any providers. Recall, however, that we define LTPs as the top-30 providers in terms of average number of AS customers. This set includes ASes that have providers.




Figure 11: Evolution of average number of providers for each AS type.




Figure 12: Evolution of the distribution of the number of providers of CAHPs.

transit providers is probably the main reason behind the densification of the Internet, discussed earlier.

We further study the *distribution* of the number of providers of different AS types. We find that the distribution of the number of providers for ECs has not changed significantly in the last 10 years. On the other hand, the largest change is for CAHPs. Figure 12 shows the distribution of the number of providers for CAHPs in 5 snapshots over the last 10 years. We see that the distribution has been shifting consistently towards the right, indicating an increase in the number of providers for CAHPs. Further, we find that the median number of providers for CAHPs has been quite close to the average, and 50% of CAHPs in the latest snapshot (2007.10) have more than 7 providers. This means that the average number of providers for CAHPs seen in Figure 11 is not biased by a small number of CAHPs that have many providers.

STPs versus LTPs: We are also interested in differences in the *type of provider* that each AS type connects to when acting as the customer in a customer-provider relation. Figure 13 shows the number of links in each transit category over time. Interestingly, we find that both EC-LTP links (meaning, the customer is an EC and the provider is an LTP) and EC-STP links show an exponential increase up to 2001, which matches the trend of the total number of CP links. Thereafter the growth rate slowed down, following a lin-

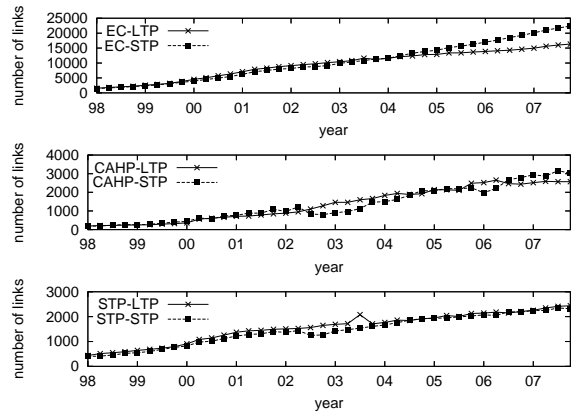


Figure 13: Evolution of CP links between different pairs of AS types.

ear increase. We find that until 2004 the number of EC-STP links was almost the same as the number of EC-LTP links. After 2004, the growth rate of EC-STP links has been higher than that of EC-LTP links (240 links/month vs 106 links/month), meaning that ECs increasingly prefer to connect to smaller, regional providers. There are several possible reasons why ECs may prefer STPs over LTPs. One possibility is that STPs are cheaper than LTPs. Another possibility is that ECs connect to STPs due to regional factors such as national monopolies and regulations, or region-specific marketing by STPs.

The middle panel of Figure 13 shows the evolution of provider links for CAHP customers, while the bottom panel shows the number of provider links for STP customers. The numbers of CAHP-LTP and CAHP-STP (also STP-STP and STP-LTP) links have been increasing at roughly the same rate. Unlike ECs, CAHP and STP customers do not prefer one type of provider over the other.

Rewiring activity of AS customers: Next, we investigate the rewiring activity of AS customers according to the broad geographical region in which they belong. Specifically, we first find the set of active customers (customers that made some change to their set of providers) between pairs of successive snapshots. Then, we calculate the fraction of those active customers that belong to each geographical region. Figure 14 shows these trends. The fractions for Asia-Pacific (APNIC), Latin America (LACNIC) and Africa (AFRINIC) are practically constant and significantly lower than for Europe (RIPE) and North America (ARIN). Interestingly, we find that after 2004-2005, there are more active customers based in Europe than in North America. In Section 4, we showed that Europe is catching up with North America in terms of the population of ECs, and the population of STPs is already larger in Europe. We conjecture that this has created a more competitive market in Europe than in North America, with European customer ASes being more active in adjusting their upstream connectivity.

6. EVOLUTION OF CP RELATIONS: PROVIDER-SIDE PROPERTIES

Preferential attachment and preferential detachment: First, we measure the total number of CP links that were born and died between two consecutive snapshots. We define the *attractiveness* A_p of a provider p as the fraction of CP links born in the second snapshot that connected to provider p . Similarly, the *repulsiveness* R_p of a provider p is the fraction of CP links that died in the second

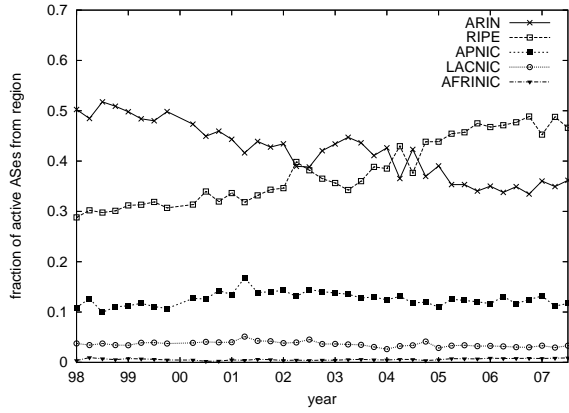


Figure 14: Fraction of active customer ASes in each geographical region.

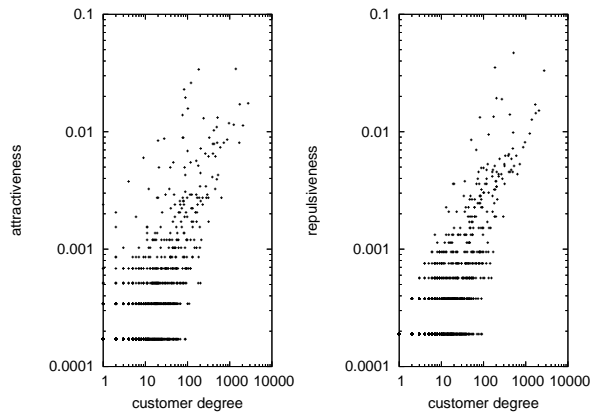


Figure 15: Attractiveness and repulsiveness versus customer degree.

snapshot and that belonged to provider p . These two metrics, attractiveness and repulsiveness, associate a business property (the ability to attract and retain customers) with a topological property (number of customer links of a provider AS).

Figure 15 shows the scatter plots of attractiveness and repulsiveness versus the number of customers, for a recent pair of snapshots in 2007. The left plot shows that *the likelihood with which a provider gains a CP link shows positive correlation with the customer degree of that provider*, as one would probably expect from the “rich get richer” principle. However, there are several outliers, and the correlation coefficient is only 64%. The low correlation indicates that a simple model in which the attractiveness of a node is proportional to its customer degree would not be very accurate. The graph at the right is also interesting because it shows *an equally strong positive correlation between the repulsiveness of a provider and its customer degree*. Thus, when we consider the rewiring of CP links, we observe not only a “preferential attachment” behavior, but also, an equally strong *preferential detachment* behavior. Preferential detachment has been largely ignored in the earlier literature, with the exception of a brief mention in [32].

Attractors and repellers: Figure 15 also shows that *there are a few providers that have very large attractiveness and repulsiveness*. We are interested in the properties of these *attractors* and *repellers* of

AS customers, and use the following approach to identify them. For each pair of snapshots, we calculate A_p and R_p for each provider p . We find that in all snapshot pairs, around 50-100 providers account for more than 60% of the total number of CP link births in the Internet. Henceforth, we identify the *attractors* of a snapshot pair as the set of providers with the highest attractiveness that account for at least 60% of the total CP link births. Similarly, we identify the *repellers*, based on the set of maximum repulsiveness providers that account for at least 60% of the total CP link deaths.⁸

Next, we examine the number of attractors and repellers between each pair of snapshots over time. Figure 16 shows the evolution of the total number of attractors and repellers, distributed among AS types. A decreasing trend in the number of attractors would imply that the customer gains are shared by a decreasing set of providers, indicating a shift towards an oligopoly or even monopoly. What we see, however, is that *the number of attractors and repellers shows an increasing trend*. This is significant because it implies that the gains and losses of customers are increasingly shared by a larger set of providers. In other words, the Internet is not heading towards an oligopoly or consolidation of providers; instead, the market of competing providers is increasing in size. We find that since 2001, the number of LTPs in the set of attractors and repellers has stayed almost constant. This is because around 25 out of the 30 LTPs appear in these sets in any given snapshot pair. *The increase in the number of attractors and repellers is mainly due to an increasing number of STPs in these sets*.

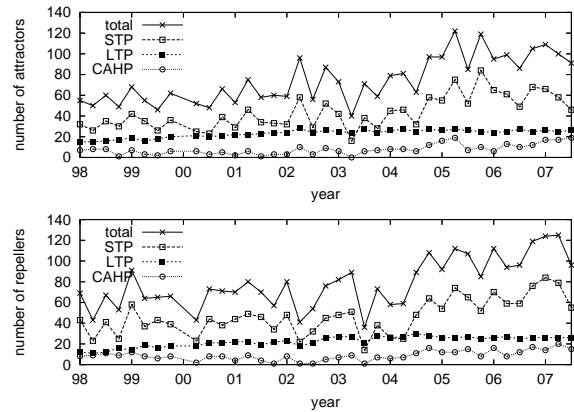


Figure 16: Evolution of the number of attractors and repellers (total and among AS types).

Figure 17 shows the number of attractors and repellers in different geographical regions. Initially, it was the case that most attractors and repellers were registered in North America. Since 2003-04, however, providers from Europe have outnumbered those from North America in the attractor and repeller sets.

In addition to the number of attractors and repellers in each geographical region, we examine the total attractiveness and repulsiveness in different regions. The total attractiveness (repulsiveness) of a set of providers is the fraction of CP link births (deaths) that are contributed by providers in that set. The top (bottom) panel of Figure 18 shows the total attractiveness (repulsiveness) of the attractors (repellers) in each geographical region. From 1998 until 2003-04, the attractors in North America had a greater total attractiveness than those in Europe (coinciding with the period in which

⁸Choosing different values for this threshold yields qualitatively similar results.

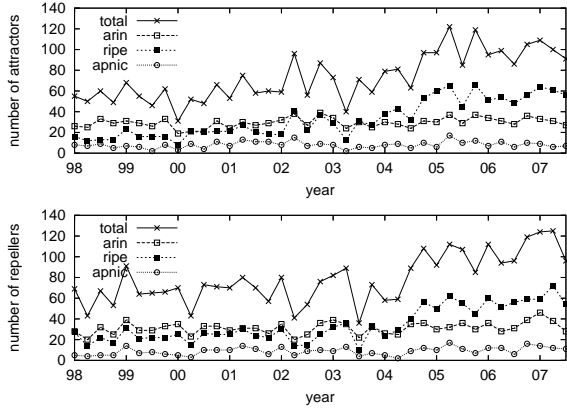


Figure 17: Evolution of number of attractors and repellers in each geographical region.

the number of attractors in North America was larger than that in Europe). It is interesting, however, that after 2003-04 the attractors in Europe and North America have similar total attractiveness. This means that even though the number of attractors is larger in Europe, they account for a similar fraction of the total CP link births than the attractors in North America. Similar trends are seen for the total repulsiveness in Europe and North America.




Figure 18: Evolution of total attractiveness of attractors and repellers in each geographical region.

Correlation of attractiveness and repulsiveness for the same AS: We have seen that providers can act as attractors or repellers of AS customers. Here, we examine whether a correlation exists between these two properties of the same provider. If so, how do these correlations vary at different time lags? To answer these questions, we calculate the crosscorrelation of the attractiveness $A_p(t)$ and repulsiveness $R_p(t)$ timeseries of the same provider at different lags. Instead of examining all providers, we restrict this analysis only to those providers that were classified as either attractors or repellers (according to the 60% rule described earlier) at some point in their lifetime. We refer to this set of providers as AR , where $|AR|=638$. For each provider in AR , we compute the cross-correlation at different lags, and also the confidence bounds at 99% significance level. The confidence bounds are used to determine whether there is a significant correlation between the attractive-




Figure 19: Lag of maximum absolute correlation for each AS provider in AR .

ness and repulsiveness time series at a particular lag. We find 317 providers for which a significant correlation exists at some lag. For each of those providers, we then determine the lag that shows the maximum absolute correlation.

Figure 19 shows, for the previous 317 providers, the lag at which the maximum (in absolute value) correlation occurred. Interestingly, we find that *in almost all cases the correlation is positive*. Further, we find that in 85% of the cases, the maximum correlation occurs at positive lags. In particular, most of the mass is at lags 1, 2 and 3 snapshots (44.7%, 13.5% and 9.1% of the providers, respectively). Note that a positive lag l means that we correlate the attractiveness at time t with the repulsiveness at time $t + l$, and each lag corresponds to 3 months. So, for a large number of providers, *strong attractiveness precedes strong repulsiveness by a period of 3-9 months*. There are several possible reasons for this effect. We conjecture that some providers attract many new customers due to advertising and promotions. These providers are not always able to keep their new customers, leading to significant repulsiveness a few months later. This may be due to customers that change providers frequently (such as CAHPs), or due to follow-up advertising/promotions from competitors.

7. CONJECTURES ON THE EVOLUTION OF PEERING

Given that a large fraction of peering links may not be visible through RouteViews and RIPE routing tables, we do not study in detail the evolution of peering relations in this paper. In this section, we only present some tentative results, which should be viewed as unproven “conjectures” about the evolution of peering. The following observations need to be re-examined in a future study, when the research community obtains sufficient visibility of the peering links in the Internet.

Figure 20 shows the median peering degree for each of the four AS types. We prefer to use the median degree in this case because the average peering degree is heavily influenced by a single LTP provider (AS13237) that appears to have over 200 peers. ECs and STPs have median peering degrees of zero. It is interesting that *the median peering degree of CAHPs has increased significantly since 2003, from 2 to 10*. It is not surprising that LTPs establish many peering links; those links are needed for global reachability when it is not possible to directly reach some destinations through customers. CAHPs, on the other hand, have the incentive to create

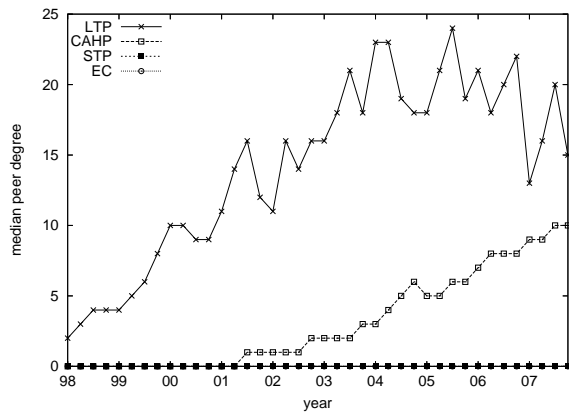


Figure 20: Median number of peers for each AS type over time.

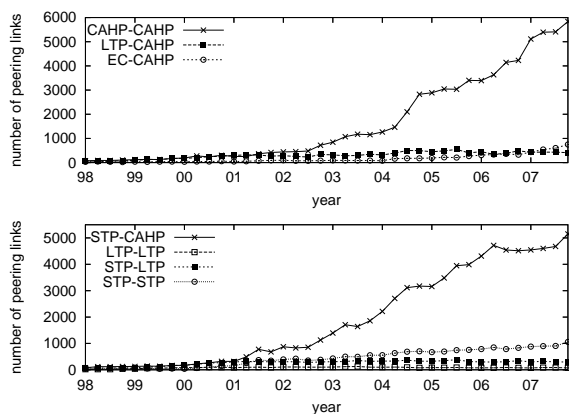


Figure 21: Number of PP links of the most common types.

many peering links to reduce their transit costs paid to upstream providers. Their revenues result from the content they offer or from access/hosting fees.

Figure 21 shows the number of peering links in each category over time. We see several interesting trends. First, the number of peering links that involve CAHPs (CAHP-CAHP, EC-CAHP, STP-CAHP) increased significantly between 2001-2005, and it shows a persistent growth rate thereafter. The exception is for the links of type LTP-CAHP, which are almost constant in number since 2003. The largest number, as well as the highest growth rate, is for links of the type CAHP-CAHP and CAHP-STP. This could be because content/hosting/access providers have the incentive to get as close as possible to the destinations/sources of their traffic. These destinations/sources of traffic are other CAHPs or they are networks that are reachable through STPs. Another interesting observation is that the number of STP-LTP peering links has remained almost constant over the last 5-6 years. We conjecture that this is due to the “restrictive” peering policy of most large transit networks. The previous observations confirm the anecdotal evidence, mentioned in various white papers (see [25] and related references), that content/access providers are rising in the peering ecosystem as the dominant players. The underlying reason is that such ASes mostly have an *open peering policy*, while transit providers have *selective or restrictive* policies, peering by necessity rather than by choice.

8. RELATED WORK

A major research effort during the last decade aimed to characterize the AS-level topology. One of the most well cited papers, by Faloutsos et al. [15], argued that the Internet AS-level topology is “scale-free”. This observation was questioned by Chen et al. [9], who showed that the degree distribution in the Internet, though heavy-tailed, does not obey a strict power-law distribution. Tangmunarunkit et al. [34] attempted to explain the heavy-tailed degree distribution, and conjectured that this could simply be due to the heavy-tailed AS size distribution. Most previous measurement studies focused on static topological properties of the Internet, such as degree distribution or clustering, and did not examine the evolution of the topology over time.

A recent study measured the average degree and effective diameter of the Internet AS graph and concluded that the AS graph is *densifying* [22]. Two other measurement studies [23, 32] studied the evolution of the Internet topology in the period 1997-2001 with respect to several microscopic and macroscopic properties. Siganos et al. [32] observed the exponential growth of the Internet during that time period, and showed that a rich-get-richer form of preferential attachment leads to exponential growth in the number of edges. Magoni et al. [23] examined the evolution of some global Internet characteristics and found exponential growth in the number of ASes and links during that time period.

The observation that the degree distribution follows a power-law led to several topology generation models that could produce such distributions. These models focused on “growing” a topology that could match the Internet topology with respect to certain measurable graph metrics. The most well known work in this area is the preferential attachment model of Barabasi et al. [2]. An extension of that model [1] incorporated the random rewiring of a fixed number of existing links. Several variants of preferential attachment models were later proposed [3, 36, 39]. Park et al. [28] compared different growth models for Internet topology with respect to several metrics such as the average diameter and clustering coefficient. The models in this research thread have been mostly descriptive. More recent work has attempted to incorporate the effect of economic factors in the evolution of the Internet topology, most notably [31, 35].

The previous descriptive models received considerable criticism (for instance, see [20, 21]) because they mostly focus on the degree distribution and clustering, ignoring important features of the Internet topology such as hierarchy or the presence of links of different types (transit versus peering). Further, the previous models do not explain how the Internet topology is evolving. This led to new models that view the Internet topology as the outcome of optimization-driven activity of individual ASes. These concepts were first introduced by Carlson and Doyle [4], and later applied in the context of the Internet in [14] and elsewhere. Chang et al. [6] used domain-specific information about the Internet to model AS interconnection practices. A recent work by Chang et al. [7] models the behavior of an AS in two distinct economic roles (customer and peer), and examines the topological effects of actions of individual ASes when acting in different roles. A recent editorial [17] stresses the need to further understand the dynamics of the AS topology. Norton [25] discusses, mainly using anecdotal evidence, how economic and competitive interests influence peering and transit connectivity in the Internet. Economides [13] discusses the economics of the Internet backbone (without looking at topology dynamics).

Several measurement studies have highlighted the incompleteness of the topologies inferred from publicly available routing data [8, 10, 18, 24, 37]. Given that the inferred topologies are incomplete, much attention has been devoted to methods that capture as much of

the Internet topology as possible, most notably the work by Zhang et al. [37] and He et al. [18]. Zhang et al. [38] investigated the effect of the selection of route monitors on different applications such as topology inference and AS path prediction. Their main observation was that for applications such as topology and relationship inference, the publicly available BGP data are reasonably accurate, and data from an increasing set of monitors is only marginally useful. Oliviera et al. [27] tackle the problems of topology liveness and completeness, i.e., how to distinguish between topology changes that are genuine link births and deaths, versus those that are caused due to link appearance and disappearance during routing transients.

9. CONCLUSIONS

We measured the evolution of the AS-level topology over the last 10 years in terms of growth as well as rewiring, four distinct economic/business classes of ASes, and customer-provider links. Our findings highlight some important trends, trend shifts, and sketch what the Internet may be heading towards. The main findings are summarized next.

The Internet has gone through two growth phases in terms of ASes and transit links: an initial exponential phase up to mid/late-2001, followed by a linear phase. Even as the network grows, the average path length remains practically constant, meaning that the network *densifies*. We find that currently, around 75% of link births are associated with existing ASes rather than new ASes (rewiring versus growth); similarly, about 80% of the link deaths are due to rewiring.

We proposed a simple AS classification scheme according to economic considerations and business types. In terms of the population of these AS types, we find that the ASes at the network edge (ECs) contribute most of the overall growth. The average multihoming degree has remained roughly constant for ECs, has increased significantly for STPs, LTPs and CAHPs. The aforementioned densification process is thus driven by transit providers and access/hosting/content providers. In terms of rewiring activity, CAHPs are the most active, while ECs are the least active.

We introduced two provider metrics, attractiveness and repulsiveness, to measure the ability of a provider to attract and retain customers. We see positive correlations between the attractiveness and repulsiveness of a provider and its customer degree. Also, for many providers, strong attractiveness precedes strong repulsiveness by a period of 3-9 months. There are a few providers that have very large attractiveness and repulsiveness (attractors and repellers). The total number of attractors and repellers between successive snapshots shows an increasing trend.

In terms of regional growth, we find that the Internet market, in terms of the number of access/hosting/content and transit providers will soon be larger in Europe than in North America. This is also reflected by the fact that since 2004-2005, a larger fraction of active customers are based in Europe than in North America, and providers from Europe increasingly feature in the set of attractors and repellers. This is important, because much of the regulatory and policy debate about the Internet has been focused on North America. Our measurements hint at an increasing European influence on the Internet ecosystem.

We have explained the previous measurement results with conjectures about the causes of the observed densification, the high activity of CAHPs, and the incentives that lead certain AS types to connect to other AS types. Unfortunately it is hard to validate these conjectures, mainly due to the lack of pricing and other economic data about various AS types. Obtaining such data and further explaining the previous observations in an economic or optimization

basis is a valuable direction for future work. An even more ambitious goal is to model the provider and peer selection process for each AS type. We would then be able to simulate the AS ecosystem and better understand both the dynamics of the AS-level topology as well as the impact of various, existing or proposed, provider selection and peering strategies.

Acknowledgements

We would like to thank our shepherd Lixin Gao and the anonymous reviewers for their helpful feedback. We would also like to thank Pierre Francois, Manish Jain, Amund Kvalbein, Aman Shaikh, Georgos Siganos, Mukarram Bin Tariq, Renata Teixeira and Walter Willinger for their valuable comments.

10. REFERENCES

- [1] R. Albert and A. L. Barabasi. Topology of Evolving Networks: Local Events and Universality. *Physical Review Letters* 85, 5234, 2000.
- [2] A. L. Barabasi and R. Albert. Emergence of Scaling in Random Networks. *Science* 286 509512, 1999.
- [3] T. Bu and D. Towsley. On Distinguishing Between Internet Power Law Topology Generators. In *Proceedings of IEEE Infocom*, 2002.
- [4] J. M. Carlson and J. Doyle. Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems. *Physical Review E* 60, 1999.
- [5] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger. An Empirical Approach to Modeling Inter-AS Traffic Matrices. In *Proceedings of the Internet Measurement Conference (IMC)*, 2005.
- [6] H. Chang, S. Jamin, and W. Willinger. Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach. In *Proceedings of ACM SIGCOMM Workshop on Models, Methods and Tools for Reproducible Network Research (MoMeTools)*, 2003.
- [7] H. Chang, S. Jamin, and W. Willinger. To Peer or Not to Peer: Modeling the Evolution of the Internet's AS-Level Topology. In *Proceedings of IEEE Infocom*, 2006.
- [8] H. Chang and W. Willinger. Difficulties Measuring the Internet's AS-Level Ecosystem. In *Proceedings of the 40th Annual Conference on Information Sciences and Systems*, 2006.
- [9] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of Power-Laws in Internet Topologies Revisited. In *Proceedings of IEEE Infocom*, 2002.
- [10] R. Cohen and D. Raz. The Internet Dark Matter - On the Missing Links in the AS Connectivity Map. In *Proceedings of IEEE Infocom*, 2006.
- [11] X. Dimitropoulos, D. Krioukov, M. Fomenkov, Y. Hyun, K. Claffy, and G. Riley. AS Relationships: Inference and Validation. *ACM SIGCOMM Computer Communication Review*, 2007.
- [12] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy. Revealing the Autonomous System Taxonomy: The Machine Learning Approach. In *Proceedings of Passive and Active Measurement Conference (PAM)*, 2006.
- [13] N. Economides. The Economics of the Internet Backbone. *Handbook of Telecommunications Economics Ed. S. Majumdar, I. Vogelsang, M. Cave. Amsterdam: Elsevier Publishers*, 2006.

- [14] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet. In *ICALP '02: Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, 2002.
- [15] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-law Relationships of the Internet Topology. In *Proceedings of ACM SIGCOMM*, 1999.
- [16] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6), 2001.
- [17] H. Haddadi, S. Uhlig, A. Moore, R. Mortier, and M. Rio. Modeling internet topology dynamics. *SIGCOMM Computer Communication Review*, 2008.
- [18] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy. A Systematic Framework for Unearthing the Missing Links: Measurements and Impact. In *Proceedings of 4th USENIX/SIGCOMM NSDI*, 2007.
- [19] G. Huston. The 32-bit AS Number Report. <http://www.potaroo.net/tools/asn32>.
- [20] E. F. Keller. Revisiting "Scale-free" Networks. *BioEssays* 27, Wiley Periodicals Inc., 2005.
- [21] D. Krioukov, kc claffy, M. Fomenkov, F. Chung, A. Vespignani, and W. Willinger. The Workshop on Internet Topology (WIT) Report. *ACM SIGCOMM Computer Communication Review*, 2007.
- [22] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph Evolution: Densification and Shrinking Diameters. *ACM Transactions on Knowledge Discovery from Data (ACM TKDD)*, 2007.
- [23] D. Magoni and J. J. Pansiot. Analysis of the Autonomous System Network Topology. *ACM SIGCOMM Computer Communication Review*, 2001.
- [24] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, kc claffy, and A. Vahdat. The Internet AS-Level Topology: Three Data Sources and One Definitive Metric. *ACM SIGCOMM CCR*, 2005.
- [25] W. B. Norton. The Evolution of the U.S. Internet Peering Ecosystem. *Equinix white papers*, 2004.
- [26] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure. In *Proc. ACM SIGMETRICS*, 2008.
- [27] R. V. Oliveira, B. Zhang, and L. Zhang. Observing the Evolution of Internet AS Topology. In *Proceedings of ACM SIGCOMM*, 2007.
- [28] S. Park, D. M. Pennock, and C. L. Giles. Comparing Static and Dynamic Measurements and Models of the Internet's AS Topology. In *Proceedings of IEEE Infocom*, 2004.
- [29] RIPE. RIPE Network Coordination Centre. <http://www.ripe.net>.
- [30] Routeviews. University of Oregon Route Views Project. <http://www.routeviews.org>.
- [31] M. A. Serrano, M. Boguna, and A. D. Guilera. Modeling the Internet. *The European Physics Journal B*, 2006.
- [32] G. Siganos, M. Faloutsos, and C. Faloutsos. The Evolution of the Internet: Topology and Routing. *University of California, Riverside technical report*.
- [33] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet hierarchy from Multiple Vantage Points. In *Proceedings of IEEE Infocom*, 2002.
- [34] H. Tangmunarunkit, J. Doyle, R. Govindan, W. Willinger, S. Jamin, and S. Shenker. Does AS size Determine Degree in AS Topology? *ACM SIGCOMM Computer Communication Review*, 2001.
- [35] X. Wang and D. Loguinov. Wealth-Based Evolution Model for the Internet AS-Level Topology. In *Proceedings of IEEE Infocom*, 2006.
- [36] S. H. Yook, H. Jeong, and A. L. Barabasi. Modeling the Internet's Large-scale Topology. *Proceedings of the National Academy of Sciences*, 2002.
- [37] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level Topology. *ACM SIGCOMM Computer Communication Review*, 2005.
- [38] Y. Zhang, Z. Zhang, Z. M. Mao, C. Hu, and B. M. Maggs. On the Impact of Route Monitor Selection. In *Proceedings ACM SIGCOMM Internet Measurement Conference (IMC)*, 2007.
- [39] S. Zhou. Understanding the Evolution Dynamics of Internet Topology. *Physical Review E*, vol. 74, 2006.