



Measuring the Evolution of Internet Peering Agreements

Amogh Dhamdhere, Himalatha Cherukuru, Constantine Dovrolis, Kc Claffy

► To cite this version:

Amogh Dhamdhere, Himalatha Cherukuru, Constantine Dovrolis, Kc Claffy. Measuring the Evolution of Internet Peering Agreements. 11th International Networking Conference (NETWORKING), May 2012, Prague, Czech Republic. pp.136-148, 10.1007/978-3-642-30054-7_11 . hal-01531958

HAL Id: hal-01531958

<https://inria.hal.science/hal-01531958>

Submitted on 2 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Measuring the Evolution of Internet Peering Agreements

Amogh Dhamdhere¹, Himalatha Cherukuru², Constantine Dovrolis², and Kc Claffy¹

CAIDA¹ Georgia Tech²
{amogh,kc}@caida.org dovrolis@cc.gatech.edu

Abstract. There is much interest in studying the structure and evolution of the Internet at the Autonomous System (AS) level. However, limitations of public data sources in detecting settlement-free peering links meant that prior work focused almost exclusively on transit links. In this work, we explore the possibility of studying the full connectivity of a small set of ASes, which we call *usable monitors*. Usable monitors, while a subset of the ASes that provide BGP feeds to Routeviews/RIPE collectors, are better suited to an evolutionary study than other ASes. We propose *CMON*, an algorithm to classify the links of usable monitors as transit or non-transit. We classify usable monitors as transit providers (large and small), content producers, content consumers and education/research networks. We highlight key differences in the evolution of connectivity of usable monitors, and measure transitions between different relationships for the same pair of ASes.

Keywords: Internet topology, Autonomous Systems, Peering, Economics

1 Introduction

The Internet consists of thousands of autonomous systems (ASes) connected together to provide end-to-end reachability. Connections between ASes are typically bilateral in nature, with an underlying business relationship. At the two ends of the spectrum of AS relationships, we have transit and settlement-free peering¹. There has been great interest recently in studying the evolution and dynamics of the Internet topology at the AS-level. Unfortunately, existing publicly available data can reliably capture only transit links, while most settlement-free peering links are invisible, especially those that are topologically lower in the Internet hierarchy than the route monitors [17]. As a result, evolutionary studies have focused on transit links [9, 18].

In this work, rather than studying only a subset of the connectivity of all ASes in the Internet, we take the approach of focusing on *the complete connectivity of a subset of ASes*. We use a subset of the ASes that provide routing feeds to Routeviews/RIPE collectors, which we call usable BGP route monitors (or *usable monitors* for short); we believe usable monitors are good candidates for an evolutionary study, as more of their AS links are visible from the local view than from remote ASes. We propose a new

¹ In a transit relationship, the customer pays the provider for carrying traffic, while no money is exchanged in a settlement-free peering relationship.

heuristic *CMON* to classify the links of usable monitors as transit or settlement-free. We study the evolution of the connectivity of usable monitors from 2006-2010, focusing on differences in the connectivity and link dynamics based on their business types.

Our study, though focused on a relatively small number of usable monitors, yields important insights into the evolution of AS connectivity. First, we confirm previous results [17] about the visibility of links of different types of ASes. While 99% of links of tier-1 ASes are visible from remote BGP route monitors, up to 75% of the links of Content Producer and Content Consumer ASes are observable only if a BGP route monitor is present at those ASes. We find that customer networks treat smaller transit providers as backup providers, while choosing larger providers as their primary transit providers. This trend has economic implications for transit providers, who get paid based on the amount of traffic they carry; providers that are used only as backup will see lower revenues. Finally, we measure the probability with which a link between two ASes changes relationship type, and produce a state transition diagram for each AS type. These state transition diagrams can be input to a model of macroscopic AS-level topology dynamics and evolution. Given the economic implications of topology dynamics, such a model can be used to predict the economic health of the transit providers ecosystem.

2 Datasets and methodology

2.1 Topology data

We collected historical BGP data from the two major public repositories at RouteViews [4] and RIPE [3]. We rely only on these two data sources because no other source of topological/routing data (routing registries, traceroutes, looking glass servers, etc.) provides historical information. The use of Routeviews/RIPE repositories of BGP data has been shown to be inadequate to expose the *complete Internet topology* [7, 8, 13]. In particular, even though most ASes are visible, a significant fraction of peering and backup links at the edges of the Internet are missed [6, 13, 20]. However, we are only interested in primary links used most of the time; missing backup and transient links are not an issue, and in fact we should remove them altogether as they are not a sound reflection of long-term evolutionary dynamics. To remove backup and transient links, we apply the method of “majority filtering” described in our previous work [9] on the set of BGP AS paths² obtained from Routeviews and RIPE collectors. We construct a topology snapshot by collecting 5 sets of AS paths over a duration of 3 weeks, only using AS paths that were seen in a majority of those five samples. Our previous work [9] presents a detailed description of the data collection and pre-processing.

2.2 Identifying usable monitors

We define a BGP route monitor as an AS that provides a routing feed to Routeviews/RIPE collectors. Previous work [9, 17] has shown that many AS links (particularly settlement-free peering links) are not visible in Routeviews/RIPE data unless at least one endpoint of the link is a BGP route monitor. Even such route monitors, however, may not export

² We do not use BGP updates, as these reveal backup and transient links which we want to filter.

their complete connectivity to Routeviews/RIPE collectors. We define a usable monitor as a monitor AS, such that *a larger number of AS links of this monitor AS are visible using the monitor's local view as compared to those visible from remote monitors*. We believe that an evolutionary study is more accurate if restricted to usable monitors, as we are able to see more of their connectivity as compared to other ASes. Note that a usable monitor does not necessarily provide a default-free table to Routeviews/RIPE collectors, so we cannot simply look for ASes that provide a complete routing table to Routeviews/RIPE collectors to identify usable monitors. We use the following procedure to identify usable monitors.

Let X be an route monitor AS as defined previously, and $\mathcal{L}_l(X)$ be the set of X 's AS links obtained from AS paths that AS X exports (we refer to $\mathcal{L}_l(X)$ as the “local view” of AS X). Similarly, route monitors other than X export AS paths that may contain AS X . Using AS paths from such *remote monitors*, we obtain the set $\mathcal{L}_r(X)$ of X 's links that are visible from remote monitors. We use the number of links in $\mathcal{L}_l(X)$ and $\mathcal{L}_r(X)$ to determine if X is a usable monitor. If $\mathcal{L}_l(X) \geq \mathcal{L}_r(X)$, then more links of AS X are visible in X 's local view than are visible at remote monitors. In such a case, we classify X as a usable monitor. In practice, we allow a small slack factor s for detecting usable monitors. If $\frac{|\mathcal{L}_r(X) - \mathcal{L}_l(X)|}{|\mathcal{L}_r(X) \cup \mathcal{L}_l(X)|} < s$, where s is the *slack factor*, then the local view from AS X misses at most a fraction s of all of X 's links. We detect usable monitors by setting s to a small value.




Fig. 1: The number of usable monitors as a function of the slack factor.

To choose an appropriate value for the slack factor, we measure the number of Routeviews/RIPE monitors that we classify as usable monitors for different values of the slack factor. Figure 1 shows that when the slack factor exceeds 0.1, there is a “plateau” effect, where the number of usable monitors does not increase sharply until the slack factor reaches around 0.7, a trend that is seen in all snapshots (The figure shows three snapshots from 2001, 2005 and 2010). For the purposes of this study, we choose to be conservative in identifying usable monitors, and set the slack factor to 0.1.

This yields fewer usable monitors, but increases the confidence that we observe their complete connectivity. For an evolutionary study, we need to balance the tradeoff between a long enough duration, and the number of ASes that are usable monitors for that entire duration. We use 17 continuous snapshots from 2006-2010 for our study, which gives us 58 ASes which were usable monitors for that entire duration³. We use peeringDB [2] and organization webpages to classify the 58 usable monitor ASes according to their business type. The 58 usable monitors consist of 11 transit providers that advertise global presence and large traffic volumes (“Large Transit Providers” or LTPs), 14 transit providers that have regional presence (“Small Transit Providers” or STPs), 12 Content Consumers (CC), 6 Content Providers (CP), 2 Enterprise Customers (EC), and 11 Education/Research networks (ER).⁴

2.3 Visibility of links of monitor ASes

Figure 2 shows, for the topology snapshot in July 2009⁵, the number of links of each of the 58 usable monitors observed from the monitor itself, and the number of links seen from remote route monitors. We find that all LTPs lie close to the diagonal – most of their links are visible from remote monitors. On the other hand, we find that a large number of links of CPs and CCs are visible only from the local monitor. We compute the *invisibility fraction* for each type of monitor AS, i.e., the fraction of links of that type of monitor AS that are invisible from remote monitors. LTPs have the smallest invisibility fraction (0.5%) – nearly all of their links are visible from remote monitors. The invisibility fraction for small transit providers (STPs) is 40%. On the other hand, 75% of the links of CP monitors are invisible from remote monitors, while the invisibility fractions for CCs and ERs are 55% and 60%, respectively. Our analysis thus confirms the findings of Oliveira et al. [17], who relied on case studies to show that most tier-1 network links are visible from remote monitors, but many – possibly most – links of content providers are not. This partial visibility of the complete AS topology is mostly due to the low visibility of the connectivity of CPs and CCs, and further confirms the limitations of existing public BGP snapshots, each of which is only a partial view of global interdomain connectivity.

3 Classifying AS links of monitor ASes

We describe *CMON*, a two-step algorithm for classifying links of monitor ASes. First, *CMON* determines if a link X-Y of monitor AS X is a transit or non-transit (settlement-free, paid-peering or other) link. *CMON* then classifies transit links as provider-customer or customer-provider, from the perspective of the monitor AS. *CMON* relies on the notion of a “hierarchy” in the AS topology, where large-degree tier-1 providers are at the top, followed by regional tier-2 providers, and so on. Though there is recent evidence that this hierarchy is flattening [10, 12], the approximate placement of ASes in this hierarchy is sufficient for *CMON*.

³ The set of ASes that are usable monitors changes over time, hence we identify 58 ASes that were full monitors throughout the study duration

⁴ Our datasets are available at www.caida.org/~amogh/monitors/datasets.html

⁵ We observed qualitatively similar trends in other snapshots.




Fig. 2: Link visibility of monitor ASes.

3.1 Step 1: Transit vs non-transit links

CMON classifies a link X - Y of a usable monitor X as a transit or a non-transit link (settlement-free or paid-peering link) based on the visibility of that link from remote monitors. Due to the “no-valley, prefer-customer, prefer-peer” rule of thumb that is generally followed for interdomain routing, a link X - Y of monitor X is visible from a remote monitor M only if M is *lower in the hierarchy than* X . In Figure 3(a), the remote monitor L , which is higher in the hierarchy than X , does not see the non-transit link X - Y . This is because X does not export the route X - Y - Z to L . On the other hand, the transit link X - W is visible from the monitor L , because X exports a path X - W to L . AS W can see the link X - Y , because X exports the route X - Y - Z to W . The crux of our heuristic is to find remote usable monitors that are higher in the hierarchy than X , and then test the visibility of links of X from this set of remote monitors. To find remote monitors that are higher in the hierarchy than X , we obtain the set of 15 largest ASes according to their average degree over the last 10 years, which also provide full tables to Routeviews/RIPE collectors. These ASes (most of which are well known “tier-1” providers) are likely to be higher in the hierarchy than monitor ASes whose links we try to classify. We refer to this set as $\mathcal{H}$. If X is itself high in the hierarchy (or in $\mathcal{H}$), then X would not advertise settlement-free peering links to other tier-1 networks, and our heuristic still works.

Let $n(l)$ be the number of usable monitors in $\mathcal{H}$ that see a link l of monitor AS X . Let $n_{max}(X)$ be the maximum number of remote monitors that see any link of monitor X . We compute $f(l) = \frac{n(l)}{n_{max}(X)}$, the normalized fraction of remote LTP monitors that see link X - Y . If $f(l) \geq C$, then link l is likely to be a transit link, where C is the threshold used for distinguishing between transit and non-transit links. Using data from a set of usable monitors for which we have complete ground truth information (described in more detail in Section 3.3), we found that a value of $C=0.25$ gives the best accuracy for classifying transit vs. non-transit links. To avoid misclassifying links for which $f(l)$ is

close to C , we classify l as “unknown-ANY” (UNK-ANY) if $0.75C < f(l) < 1.25C$, i.e., we cannot be certain if the link is a transit or non-transit link.

As the classification of a link of a monitor X as transit vs. non-transit relies on the visibility of that link from remote monitors, we discuss the factors that affect this visibility, and how to interpret a link that *CMON* classifies as non-transit. As already described, settlement-free (and paid-peering) links are not visible from remote monitors high in the hierarchy. On the other hand, a provider-customer link X - R of monitor X may also not be visible from remote monitors, because this link is either a backup transit link or is used for only one direction of traffic in load balancing. If R prepends route announcements over the link X - R , then remote ASes may prefer other routes to reach R . In such cases, *CMON* classifies the link X - R as a non-transit link. As this link is either not used at all (backup link), or is only used for one direction of traffic, we believe that *CMON* correctly distinguishes that link from a transit link used for all traffic. We further discuss the issue of non-transit links when we validate *CMON* in Section 3.3.




Fig. 3: Intuition behind the heuristics in *CMON*.

3.2 Step 2: Provider vs customer links

In Step 2, *CMON* classifies transit links of a monitor AS X (after filtering out non-transit links in Step 1) as either customers or providers of X , relying on the position of monitor AS X in the Internet hierarchy. If a remote monitor L is higher in the hierarchy than monitor X , then an AS path of the form “ $L..X Y$ ” indicates that Y should be a customer of X , i.e., L sees the link X - Y in the *customer form*. On the other hand, a path from L of the form “ $L..Z X$ ” indicates that Z should be a provider of X , i.e., remote LTP monitor L sees the link X - Z in the *provider form*. Figure 3(b) shows an example where remote LTP monitors L_1 and L_2 are higher in the hierarchy than monitor X . Both L_1 and L_2 see link X - Y in the customer form and link X - Z in the provider form. On the other hand, the remote monitor L_3 is lower than X in the hierarchy, and sees the customer form for link X - Z . If the remote monitors we use are *all higher in the hierarchy than X* , then they should all see the customer form for link X - Y , and the provider form for link X - Z . The key is to choose remote monitors that are higher in the hierarchy than monitor X . We use the set of high-degree usable monitors $\mathcal{H}$ that we determined in Section 3.1. For each link l of monitor AS X , let $n_c(l)$ and $n_p(l)$ be the number of monitors in $\mathcal{H}$ that see the customer and provider form for l , respectively. Instead of choosing the form seen by a simple majority of monitors in $\mathcal{H}$, we use the following (conservative) test. If

$n_p(l) \geq 2n_c(l)$, then we classify l as a provider link; if $n_c(l) \geq 2n_p(l)$, we classify l as a customer link. If neither is true, then we classify l as 'UNK-TRANSIT', i.e., we are certain that the link is a transit link, but cannot determine if it is a provider or customer. Again, if X is itself a tier-1 network (or is in $\mathcal{H}$), then networks in $\mathcal{H}$ would see the customer form for customers of X (X would not have transit providers in this case).

3.3 Validation

To validate *CMON*, we used ground truth information from a set of 6 networks⁶. As *CMON* is designed specifically to classify the links of monitor ASes, our ground truth must also be from monitor ASes. We had access to the full routing tables for 3 ground truth networks, while the remaining networks are a usable monitor at Routeviews/RIPE. We had access to partial ground truth information for two Routeviews/RIPE usable monitors (ESNET and IJ). We find that the accuracy of *CMON* is 90% for the set of 4 ground truth networks (26 errors out of 260 total links). In the case of 18 out of 26 errors, the ground truth was a customer, while *CMON* classified it as a non-transit or UNK link. Evaluation results for *CMON* on the two partial ground truth networks are promising. One of these networks indicated that they had one provider, which *CMON* identified correctly. The second partial ground truth network indicated that they have a single customer which again *CMON* identified correctly. We also performed some sanity checks on the relationship inferences from *CMON*. We tested *CMON* on Routeviews/RIPE usable monitors that are well-known tier-1 networks: AT&T (AS7018), Cogent (AS174), Level3 (AS3356), and Hurricane Electric (AS6939). *CMON* produces only a few (or no) providers, and a large number of customers, which is the result we expect for tier-1 networks. We plan to extend the validation of *CMON* using ground truth we are collecting as part of CAIDA' AS-rank project [1].

CMON does classify a number of links of transit providers as non-transit, while we expect large tier-1 networks to have only a few settlement-free peering links. As discussed in Section 3, the non-transit category includes backup transit links, or transit links used for only one direction of traffic. For each non-transit link $M-X$ of a monitor M , we determine if X advertises at least one prepended route over this link. Prepending indicates that X does not prefer to receive traffic from M over this link, and X uses the link $M-X$ either only for outbound traffic, or as a backup transit link. We find that for 6 out of 11 LTP monitors, more than 50% of non-transit links are prepended. These fractions are quite high for some networks – 85% for AT&T (AS7018), 64% for Savvis (AS3561), 62% for Level 3 (AS3356), and 51% for Cogent (AS174). These fractions are lower for STPs, with Telstra (AS1221) the largest at 50%. For all CPs and CCs, this fraction is less than 12%, indicating that most non-transit links of these ASes are settlement-free or paid-peering links.

Finally, we compared *CMON* with the algorithm by Gao [11] (*GAO*) using the same ground truth data. We find that *GAO* has an accuracy of 80% on the ground truth, compared to 90% for *CMON*. In 45 out of 50 errors by *GAO*, the ground truth is a peering link, while *GAO* classifies it as either a provider or customer link. The agreement between *CMON* and *GAO* is 72%. Note that *GAO* classifies certain links as siblings, a

⁶ Georgia Tech, SOX, SWITCH, Media Network Services, ESNET, and IJ

relationship type that *CMON* would classify as non-transit. We treat such cases as disagreements between *GAO* and *CMON*. In 5589 out of 6750 disagreements between *CMON* and *GAO* (out of a total of 24758 compared links), *CMON* classified the link as non-transit, while *GAO* classified it as either provider or customer.

4 Evolution of Connectivity of Monitor ASes

We first make some general observations about trends in the number of customers, providers and non-transit links for monitors of each AS type from 2006-2010. We find that customer links account for most (74%) of the connectivity of LTP monitors in 2010, while non-transit links (due to load balancing/backup relationships, as described earlier) account for 24%. The average number of customers of LTPs shows an increasing trend (1149 in 2006 to 2111 in 2010), while the average number of non-transit links is almost constant (168 in 2006 to 207 in 2010). On the other hand, 80% of the connectivity of STPs in 2010 was due to non-transit links, and only 18% due to transit links. The average number of non-transit links for STPs increased (167 to 253), while the number of customer links declined (35 to 24). This indicates that *customers increasingly prefer LTPs as their primary transit providers*. Non-transit links accounted for 90% of the connectivity of CPs and CCs in 2010. As only a small fraction ($\leq 12\%$) of routes advertised over non-transit links of CPs and CCs are prepended (see Section 3.3), these non-transit links are likely settlement-free or paid-peering links. The average number of non-transit links of CPs and CCs shows an increasing trend from 2006-2010 (124 to 165); this is expected, as they tend to peer aggressively to reduce upstream transit costs.

We compare the growth rates of transit and non-transit links for usable monitors with the growth rate of transit and non-transit links in the overall graph from 2006 to 2010. To obtain the growth rates of AS links in the overall graph, we use data from our prior work [9], which used Gao's AS relationship algorithm to classify each link in the AS topology over the same time span. The number of non-transit links in the AS topology grew by a factor of 1.77 from 2006 to 2010, while transit links grew by a factor of 1.6. The top graph in Figure 4 shows the growth rate in the number of transit links of each monitor AS from 2006 to 2010, and the trend line shows the growth rate of non-transit links in the overall graph. We see that monitor ASes mostly follow the trend seen in the overall graph. Most LTP monitors show a larger growth in transit links than the overall trend, while most STP monitors show a smaller growth than the overall trend. The bottom graph in Figure 4 shows the growth rate of non-transit links of each monitor AS from 2006 to 2010. Interestingly, most of the monitor ASes we study lie below the trend line – non-transit links of monitor ASes grow more slowly than non-transit links in the overall graph. The ASes for which non-transit links grow faster than the overall rate of non-transit link growth are mostly Content Consumers (CC) and STPs. The number of non-transit links grows more slowly than the overall rate for almost all LTPs.

5 Dynamics of Connectivity of Monitor ASes

In this section, we analyze dynamics in the links of monitor ASes over time. We define the *span* of a link as the time between the first and last occurrence of the link.




Fig. 4: Change in the number of customer and non-transit links of monitor ASes from 2006-2010. The trend line shows the growth rate in the complete graph for the same duration.

5.1 Link durations

We first study the duration of various link types. We consider each link of our set of usable monitor ASes, and count the number of snapshots in which the link persists *with the same relationship type*. A change of relationship, e.g., from customer to non-transit, is counted as the death of a customer link, and the birth of a non-transit link. We exclude links that *CMON* classifies as UNK at any point in their lifetime. Including links that were UNK at any point during their lifetime would underestimate the duration of the known relationships associated with those links.

First, we classify all links of usable monitor ASes into two broad types: transit and non-transit, and compute the durations separately for links of each type. We find that non-transit links have slightly shorter lifetimes than transit links – a median duration of 5 snapshots (15 months) for non-transit links, 6 snapshots (18 months) for provider links, and 7 snapshots (21 months) for customer links.

Next, we study how the duration of customer, provider and non-transit links differs based on the type of monitor AS. Figure 5 shows CDFs of the durations of different link types, for each monitor type. Provider links of LTPs are the shortest lived (median 3 months), while those of CP, CC and STP monitors are longest-lived (median 24 to 33 months). This difference is consistent with the fact that LTPs aim to be “transit-free”, i.e., have no transit providers at all. The customer links of all monitor types show similar durations (median 15 to 21 months, graph not shown). We find that non-transit links of LTPs have the shortest duration (median 9 months), while those of STPs, CPs and CCs are the longest (median 21 months). We conjecture that this is due to customers of LTPs that frequently reconfigure their primary/backup providers. The link between such customers and the LTP monitor often changes relationship type between provider and non-transit.

Figure 6 compares the durations of transit and non-transit links for STP (top) and LTP (bottom) monitors. For LTPs, non-transit links have a median lifetime of 9 months, half that of the median for transit links. For STPs, non-transit links are longer-lived




Fig. 5: Link duration of transit and non-transit links for monitor ASes.

than transit links (median of 24 and 18 months, respectively). As discussed earlier, a significant fraction of non-transit links of STPs and LTPs are due to load-balancing or backup configurations. Interestingly, such non-transit links of STPs persist longer than those of LTPs, suggesting that customer networks tend to stay with LTPs as their primary providers, and use STPs as backup providers (or as transit providers for only one direction of traffic).




Fig. 6: Comparing link durations of link types for STP and LTP monitors

5.2 Link state transition probabilities

Finally, we study transitions between relationship types for the same pair of ASes. For each link l of monitor AS M , let $s_l(t)$ represent the relationship at time t : either NONE




Fig. 7: Transition probabilities for links of STP, LTP, and CP/CC/ER monitors

(no link), provider, customer or non-transit. We estimate the conditional probability $P(s_l(t+1) = s_2 | s_l(t) = s_1)$ that the link is in relationship state s_2 in the snapshot after it is in relationship state s_1 (s_1 can be the same as s_2). Figures 7(a) to 7(c) show the transition probabilities of the links of STP, LTP, and CP/CC/ER monitors (CPs, CCs and ERs had similar state transition probabilities, so we merged these three classes for this analysis). We only show transition probabilities that are at least 0.02. The bold arrow shows the transition with the largest probability *out* of any state.

The major difference between these transition diagrams is in the transition out of the “NONE” state, which occurs when the link first appears, or disappears and reappears at some point during its span. For LTP monitors, a transition from NONE goes to the customer state with the highest probability (0.66), while it goes to non-transit with probability 0.23. On the other hand, for STPs, the largest transition out of NONE is to the non-transit state with probability 0.70, while it goes to customer with probability 0.18. These transition probabilities indicate that customers are more likely to choose LTPs as their primary providers, while customers tend to select STPs as backup providers, or they are used for only one direction of traffic. The transition probability from non-transit to customer is negligible for STPs, and 0.03 for LTPs, suggesting that networks which start their relationship with STPs in the non-transit state are unlikely to become full transit customers in the future. For LTPs, this transition occurs with a low, but non-negligible, probability. For CP, CC and ER monitors, the largest transition probability out of NONE is to non-transit. This is intuitive, as these networks tend to create a large number of non-transit (mostly settlement-free or paid-peering) links. In each of these transition diagrams, the probability of staying in the same relationship state is quite high. Also, the transition from customer to provider or vice-versa is rare for all monitor types – less than 0.01 probability and not shown in the graphs.

We envision a *top-down* model that uses these state transition probabilities to predict evolution dynamics in the AS topology at a macroscopic level. While *bottom-up* models have been favored in the literature (see [5, 10] and references therein), those models are highly complex, and must be parameterized with precise data about interdomain traffic, economics, and geography. An interesting question is whether an evolution model using only the state transition probabilities for different AS types, *agnostic to underlying factors*, can still accurately model topology dynamics and evolution.

6 Related Work

Several measurement studies highlighted the incompleteness of AS topologies derived from publicly available BGP data, particularly the limited visibility of settlement-free peering links [6, 8, 13, 16, 20]. Given that the inferred topologies are incomplete, previous work proposed methods to capture as much of the Internet topology as possible [13, 20]. Due to the incompleteness problem, prior measurement studies of topology evolution had to either focus on transit links, or on macroscopic properties of the AS graph. A recent study measured the average degree and effective diameter of the Internet AS graph and concluded that the AS graph is *densifying* [14]. Siganos *et al.* [19] observed exponential growth and preferential attachment in the Internet from 1997-2001. Magoni *et al.* [15] found exponential growth in the number of ASes and links during that same time period. Oliveira *et al.* [18] tackled the problems of topology liveness and completeness, i.e., how to differentiate genuine link births and deaths from routing transients. Dhamdhere *et al.* [9] studied the evolution of the Internet ecosystem (focusing mostly on transit links) over the last decade. A previous study by Zhang *et al.* [21] investigated the effect of route monitor placement on topology inference and AS path prediction, without using these route monitors to study topology evolution and dynamics. Our work differs from previous work in two significant ways. First, our study is the first to use BGP route monitors to study the evolution of the AS topology, including settlement-free peering links. Second, we focus on the dynamics of the connectivity of individual ASes, and not on macroscopic topological properties.

7 Summary and Future Work

A great deal of research over the last decade has focused on measuring the Internet topology at the AS level. The major results from those measurement studies, however, were negative, emphasizing the fact that interdomain connectivity, particularly of settlement-free peering links, is notoriously hard to measure. In this paper, we presented an approach to study the complete connectivity of a subset of ASes in the Internet – BGP route monitors that provide a feed to public BGP data repositories. We proposed *CMON*, a heuristic to classify the links of monitor ASes. As future work, we plan to extensively validate *CMON* using data from CAIDA's AS rank [1]. We studied the evolution and link dynamics of monitor ASes, highlighting trends and differences depending on the type of monitor. We continue to collect topology snapshots from Routeviews/RIPE, and tracking the evolution of full monitors over time is part of our ongoing work.

8 Acknowledgements

A. Dhamdhere, C. Dovrolis and K. Claffy were supported in this work by the National Science Foundation (grant CNS-1017139).

References

1. CAIDA AS Rank. <http://http://as-rank.caida.org/>
2. Peering database. <http://www.peeringdb.com>
3. RIPE Network Coordination Centre. <http://www.ripe.net>
4. University of Oregon Route Views Project. <http://www.routeviews.org>
5. Chang, H., Jamin, S., Willinger, W.: To Peer or Not to Peer: Modeling the Evolution of the Internet's AS-Level Topology. In: Proc. IEEE Infocom (2006)
6. Chang, H., Willinger, W.: Difficulties Measuring the Internet's AS-Level Ecosystem. In: Proc. Annual Conference on Information Sciences and Systems (2006)
7. Chen, Q., Chang, H., Govindan, R., Jamin, S., Shenker, S., Willinger, W.: The Origin of Power-Laws in Internet Topologies Revisited. In: Proc. IEEE Infocom (2002)
8. Cohen, R., Raz, D.: The Internet Dark Matter - On the Missing Links in the AS Connectivity Map. In: Proc. IEEE Infocom (2006)
9. Dhamdhere, A., Dovrolis, C.: Ten Years in the Evolution of the Internet Ecosystem. In: Proceedings of ACM Internet Measurement Conference (2008)
10. Dhamdhere, A., Dovrolis, C.: The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In: Proc. ACM CoNEXT (2010)
11. Gao, L.: On Inferring Autonomous System Relationships in the Internet. IEEE/ACM Transactions on Networking 9(6) (2001)
12. Gill, P., Arlit, M., Li, Z., Mahanti, A.: The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In: Proc. PAM (2008)
13. He, Y., Siganos, G., Faloutsos, M., Krishnamurthy, S.V.: A Systematic Framework for Unearthing the Missing Links: Measurements and Impact. In: Proc. USENIX/SIGCOMM NSDI (2007)
14. Leskovec, J., Kleinberg, J., Faloutsos, C.: Graph Evolution: Densification and Shrinking Diameters. ACM Transactions on Knowledge Discovery from Data (2007)
15. Magoni, D., Pansiot, J.J.: Analysis of the Autonomous System Network Topology. ACM SIGCOMM CCR (2001)
16. Mahadevan, P., Krioukov, D., Fomenkov, M., Huffaker, B., Dimitropoulos, X., kc claffy, Vahdat, A.: The Internet AS-Level Topology: Three Data Sources and One Definitive Metric. ACM SIGCOMM CCR (2005)
17. Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.: In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure. In: Proc. ACM SIGMETRICS (2008)
18. Oliveira, R.V., Zhang, B., Zhang, L.: Observing the Evolution of Internet AS Topology. In: Proc. ACM SIGCOMM (2007)
19. Siganos, G., Faloutsos, M., Faloutsos, C.: The Evolution of the Internet: Topology and Routing. University of California, Riverside technical report
20. Zhang, B., Liu, R., Massey, D., Zhang, L.: Collecting the Internet AS-level Topology. ACM SIGCOMM CCR (2005)
21. Zhang, Y., Zhang, Z., Mao, Z.M., Hu, C., Maggs, B.M.: On the Impact of Route Monitor Selection. In: Proc. ACM SIGCOMM IMC (2007)