Mixing Biases: Structural Changes in the AS Topology Evolution

Hamed Haddadi[1], Damien Fay[2], Steve Uhlig[3], Andrew Moore[4], Richard Mortier[5], Almerima Jamakovic[6]

[1] Royal Veterinary College, University of London**
[2] National University of Ireland, Galway
[3] Deutsche Telekom Laboratories/TU Berlin
[4] University of Cambridge
[5] University of Nottingham
[6] TNO ICT, The Netherlands

Abstract. In this paper we study the structural evolution of the AS topology as inferred from two different datasets over a period of seven years. We use a variety of topological metrics to analyze the structural differences revealed in the AS topologies inferred from the two different datasets. In particular, to focus on the evolution of the relationship between the core and the periphery, we make use of a recently introduced topological metric, the weighted spectral distribution.

We find that the traceroute dataset has increasing difficulty in sampling the periphery of the AS topology, largely due to limitations inherent to active probing. Such a dataset has too limited a view to properly observe topological changes at the AS-level compared to a dataset largely based on BGP data. We also highlight limitations in current measurements that require a better sampling of particular topological properties of the Internet. Our results indicate that the Internet is changing from a core-centered, strongly customer-provider oriented, disassortative network, to a soft-hierarchical, peering-oriented, assortative network.

1 Introduction

The Internet continuously evolves: new networks are created, old ones disappear, and existing ones grow or merge. At the same time, business dynamics cause interconnections between networks to change. Both these effects cause the underlying topology of the Internet to be in a constant state of flux. Studying the evolution of this topology is important as it impacts a variety of factors relevant to network users and application designers, such as scalability and performance. For example, different network structures affect the speed of propagation of both legitimate (e.g., routing) and illegitimate (e.g., hijacked prefixes) information.

Most efforts to understand the structure of the Internet have focused on the Autonomous System (AS) topology. There are over 30,000 ASes today, each

** This work was done while the author was at Max Planck Institute for Software Systems, Germany.

representing a single administrative authority with its own network and peering policies. Thus, the AS topology is a graph reflecting the interconnections between the networks that compose the Internet. Relationships between ASes are typically classified as either customer-provider, sibling-sibling or peer-peer. Note that as the Internet has grown, many larger networks have come to be represented as more than one AS (i.e., to advertise more than one AS number). As a result, the AS topology may contain edges that do not directly represent a business relationship between two distinct networks. However, the AS topology serves as an available, albeit approximate, measure of the complexity of the Internet's structure at the network level.

Characterizing the structure of the AS topology has proven to be difficult, but it is usually simplified to identifying a richly connected core, including the fully meshed tier-1 Internet Service Providers (ISPs), providing connectivity for the large number of smaller ISPs and customer networks at the periphery of the network. These edge ISPs may connect to only a single upstream provider, or may connect to many for resilience, performance and cost reasons. Recent work has shown that the trend is for networks to try to connect directly in the periphery of the Internet, rather than to the core, bypassing the largest providers [8]. However, no direct evidence of a corresponding large-scale change in the topological structure had been shown.

In this paper we analyze the evolution of the AS topology using two significant datasets, each generated by a different measurement technique: the Skitter dataset using traceroute, and the UCLA dataset using BGP. We are aware that there are problems with biased measurements in both data sets and one of the aims of this paper is to highlight such differences and biases, which could potentially affect many simulations, protocol designs and publications based on these datasets. However, we still aim to draw conclusions mindful of these drawbacks in this paper. We focus on the overall structure of the topology, rather than local features such as node degree, using a recently introduced metric called the *weighted spectral distribution* (WSD) [7]. This allows us to distinguish topologies with different mixing properties, i.e., how much the core can be differentiated from the periphery of the topology. A clear distinction between the core and the periphery is believed to be one of the strongest features of the Internet topology [19, 20].

This paper makes three contributions. First, we explain how the WSD depicts the mixing between core and periphery in the AS topology (Section 2). Second, we find that the AS topology has evolved from a highly hierarchical graph with a clearly distinct core towards a "softer" hierarchy where the core and non-core parts of the topology are less distinct (Section 3). Third, we show how the two different measurement techniques, traceroute and BGP, both provide limited but complementary coverage of the AS topology: the traceroute dataset has increasing difficulty sampling the periphery, while the BGP dataset can improve its sampling of the transit part of the Internet (Section 4). Section 3.1 studies the evolution of the AS topology seen in the Skitter dataset, and Section 3.2 then studies the evolution of the AS topology seen in the UCLA dataset. We compare these views of the AS topology in Section 4, where we also discuss the likely evolution of the "real" AS topology.

We are aware of the problems associated with traceroute sampling and we are also aware of the efforts in DIMES project to remedy these issues[7], however this data is currently only available since January 2007 and hence not long enough for a thorough comparison of Internet topology evolution.

2 Theoretical background

The *weighted spectral distribution* (WSD) is a graph theoretic metric based on the random walk cycles in a graph. A random walk starts at a node, say u, with degree d_u, and transitions to a connected node with probability $1/d_u$. After several such steps, say N, if the random walk returns to the starting node, then this is called a random walk cycle of length N. The WSD takes the struture of the graph to be all such random walk cycles as expressed via the normalised Laplacian (roughly speaking, how the graph appears over short walks taken from every node). The normalised Laplacian matrix of a graph, G, defined as:

$$L(G)(u,v) = \begin{cases} 1, & \text{if } u = v \text{ and } d_v \neq 0 \\ -\dfrac{1}{\sqrt{d_u d_v}}, & \text{if } u \text{ and } v \text{ are adjacent} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

Expressing L using the eigenvalue decomposition,

$$L(G) = \sum_i \lambda_i e_i e_i^T \tag{2}$$

where e_i and λ_i are the eigenvalues and eigenvectors of L respectively[8]. The WSD is based on the following theorem from [7]:

Theorem 1. *The eigenvalues, λ_i, of the nomalised Lapacian matrix for an undirected network are related to the random walk cycle probabilities as:*

$$\sum_i (1-\lambda_i)^N = \sum_C \frac{1}{d_{u_1} d_{u_2} \ldots d_{u_N}} \tag{3}$$

where d_{u_i} is the degree of node u_1 and $u_1 \ldots u_n$ denotes a path from node u_1 of length n ending at node n, i.e. an n-cycle. The number of N-cycles is related to various graph properties. The number of 2-cycles is just (twice) the number of edges and the number of 3-cycles is (six times) the number of triangles. Hence $\sum_i (1-\lambda)^3$ is related to the clustering coefficient[9]. An important graph property

[7] `http://www.netdimes.org`

[8] These are in general different from the eigenpairs of the walk Laplacian.

[9] The clustering coefficient, $\gamma(G)$, is defined as the average number of triangles divided by the total number of possible triangles

$$\gamma(G) = 1/n \sum_i \frac{T_i}{d_i(d_i-1)/2}, d_i \geq 2 \tag{4}$$

is the number of 4-cycles. A graph which has the minimum number of 4-cycles, for a graph of its density, is quasi-random, i.e., it shares many of the properties of random graphs, including, typically, high connectivity, low diameter, having edges distributed uniformly through the graph, and so on. For a proof see [7]. Theorem 1 states that the probability of taking a random walk of length N that returns to the original node, is directly related to the weighted eigenvalues of L. This probability is the 'local structure' of a node, i.e. its local connecivity. Noting that the λ_i are unique [10] to a graph it can be seen that the WSD gives a "thumbprint" of the graph structure. In [7] the distribution of eigenvalues, $f(\lambda = k)$, rather than the eigenvalues themselves is used to form a graph metric (we refer the reader to [7] for details). Specifically the *weighted spectral distribution* is then defined as:

$$WSD : G \to \Re^{|K|}\{k \in K : ((1-k)^N f(\lambda = k))\} \tag{5}$$

Where K is the set of bins used to estimate the distribution. Of interest in this paper is the spectral clustering coefficient, $\omega(G, 3)$ defined as:

$$\omega(G, 3) = \sum_K ((1-k)^3 f(\lambda = k)) \tag{6}$$

which gives a measure of the proportion of paths length 3 in the network which form triangles. As shown in [7] the WSD and $\omega(G, 3)$ can be used for estimating the parameters of a topology generator that produce graphs which are close (in the WSD sense) to the target graph. It is also shown in [7](Section V.A) how the WSD represents the core and periphery of a graph in terms of easily identifiable peaks. However, in this paper we apply the technique for tracking the evolution of the AS level graph. The WSD enables us to view the distinct features of the core and periphery more clearly than in the past.

3 Evolution of the Internet

In this section we look at the evolution of the Internet seen through the two datasets, over a total period of more than 7 years and 3 joint years of the two datasets. We rely on a number of topological metrics presented in [10].

3.1 Skitter topology

The first dataset we study consists of 7 years of traceroute measurements, starting in January 2001, collected by the CAIDA Skitter project [12]. Traceroutes are initiated from several locations in the world toward a large range of destination IP addresses. The IP addresses reported in the traceroutes are mapped to

where T_i is the number of triangles for node i and d_i is the degree of node i.

[10] This is not strictly true but the proportion of co-spectral graphs is thought to be insignificant.

AS numbers using RouteViews BGP data. We use a monthly union of the set of all unambiguous links collected on a daily basis by the project.[11]

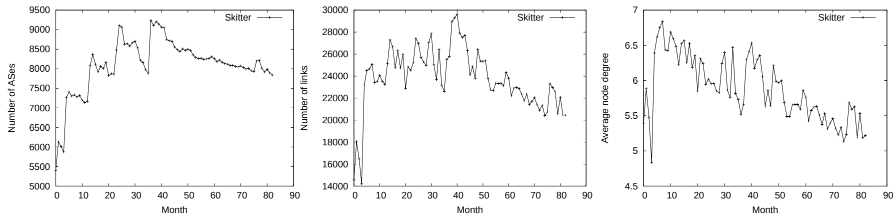

(a) Number of nodes (b) Number of links (c) Average degree

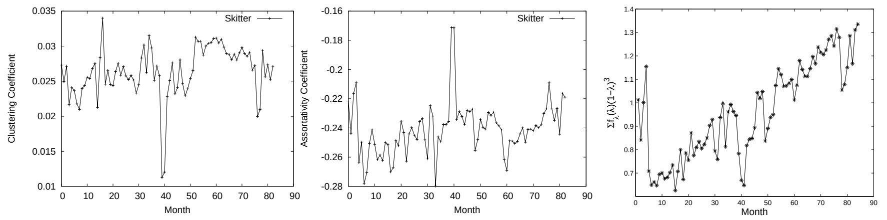

(d) Clustering coefficient (e) Assortativity coefficient (f) $\omega(G, 3)$

Fig. 1. Topological metrics for Skitter AS topology.

Figure 1 presents the evolution over the 7 years of a set of topological metrics computed on the AS topology of the Skitter dataset.

The number of ASes seen by Skitter exhibits abrupt changes during the first 40 months. At the end of those 40 months, changes were made in the way probing was performed.[12] The large increases in the number of ASes, observed during the first 40 months, are due to new monitors being added to the system. After each increase in the number of ASes a smooth decrease follows, corresponding to a subset of the IP addresses of the Skitter list that no longer respond to probes, e.g., because a firewall starts blocking the probes. The variations in the number

[11] A link may be ambiguous for a variety of reasons, principally due to problems resolving an IP address to an AS number. The Skitter IP address list includes some IP addresses which matched a prefix with two or more origin ASes. This can happen for a number of reasons such as a provider stripping the customer AS from the AS path. Since it is not known which AS is the true origin, the dataset lists both ASes. We filter out such instances as it is not possible to identify the authenticity of such links.

[12] These changes were subject to caveats and bugs affecting measurements, and, thus, the values of the resulting metrics at month 40. For more information we refer to `http://www.caida.org/data/active/skitter_aslinks_dataset.xml/`.

of ASes seen by Skitter are not caused by changes in the AS topology itself, but are artifacts of the probing. Such artifacts should be reported and accounted for in topological studies.

The number of AS edges and the average node degree both follow the behavior of the number of ASes seen. We only observe a large increase in the number of links during the first few months, during which new monitors are added resulting in new regions of the Internet being covered by Skitter measurements. Given the difficulty of building a list of destination IP addresses that will answer probes and cover most of the ASes, especially at the edge [2], a new monitor will typically discover new ASes close to its location.

The AS edges that Skitter no longer observes probably still exist but can no longer be seen by Skitter due to its shrinking probing scope. To be effective in observing topology dynamics, traceroute data collection must update destination lists constantly to give optimal AS coverage. This limitation of Skitter is visible in the decreasing average node degree. We would expect to see a net increase in the average node degree as ASes tend to add rather than remove peering links, and the results of the BGP data support this view. If the sample of the AS topology of the Skitter measurements was not worsening, we should see an increasing average node degree.

The lower three graphs of Figure 1 present the evolution of the clustering coefficient, the assortativity coefficient, [13] and the weighted spectrum with $N = 3$, $\omega(G, 3)$ (related to the topology's clustering)[14]. We observe that changes were made to the way Skitter probes the Internet around month 40: the metrics take an unusual value, very small for the clustering and very high for assortativity. The values of the clustering and the assortativity coefficients fluctuate wildly over the 7 years, as if the sampling of the AS topology by Skitter at the AS-level is not stable. Neither the clustering nor the assortativity seem to decrease or increase over the 7 years. The value of $\omega(G, 3)$ shows a long-term increasing trend, similar to the decreasing trend in the average node degree. Although related to the clustering, $\omega(G, 3)$ gives different weights to different parts of the topology. The subset of the topology that corresponds to duplicated topological structures, e.g. different ASes at the periphery that connect to the same set of upstream providers, receives a smaller weight than the rest. The increasing $\omega(G, 3)$ is likely to be caused by the shrinking network sampled by Skitter, that contains more 3-cycles on average.

Figure 2(a) presents four WSDs sampling the entire duration of the Skitter dataset. Notice the eigenvalues at zero, indicating the presence of several disconnected components. The WSD in January 2002 shows a single peak at $\lambda = 0.4$. As time passes, a second peak appears around $\lambda = 0.3$. The WSD computed

[13] Assortativity is a measure of the likelihood of connection of nodes of similar degrees [14]. This is usually expressed by means of the *assortativity coefficient* r: assortative networks have $r > 0$ (disassortative have $r < 0$ resp.) and tend to have nodes that are connected to nodes with similar (dissimilar resp.) degree.

[14] See [7] and [9] for a detailed explanation on the mathematical measures and different datasets.

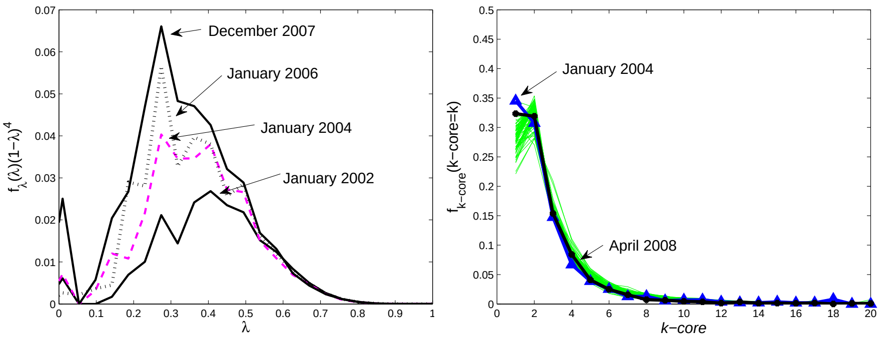

(a) WSD, Skitter topology. (b) k-core proportions, Skitter topology.

Fig. 2. Clustering and spectral features of Skitter topology

from the Skitter data suggests an Internet moving from a less hierarchical to more hierarchical topology, as if the core was becoming more dominant. This contradicts current observations that the AS topology is becoming less hierarchical, with increasing numbers of ASes peering at public Internet Exchange Points (IXPs) to bypass the core of the Internet [8].

To understand the unexpectedly dominant core seen in the Skitter dataset, we rely on the k-core metric. A k-core is defined as the maximum connected subgraph, H, of a graph, G, with the property that $d_v \geq k \ \forall v \in H$. As pointed out by [1] and [3] the k-core exposes the structure of a graph by pruning nodes with successively higher degrees, k, and examining the maximum remaining subgraph; note this is not the same as pruning all nodes with degree k or less. Figure 2(b) shows the proportion of nodes in each k-core as a function of k. There are 84 plots shown, but as can be seen there is little difference between each of them, demonstrating that the proportion of nodes in each k-core is not fundamentally changing over time. The WSD on the Skitter data is therefore not really observing a more dominant core, but a less well-sampled edge of the AS topology. We provide explicit evidence in Section 4 that Skitter has increasing problems over time to sample the non-core part of the topology.

There is a practical explanation for the sampling bias of Skitter: the Skitter dataset is composed of traceroutes rooted at a limited set of locations, so the k-core is expected to be similar to *peeling the layers from an onion* [1]. From a topology evolution point of view, Skitter's view of the AS evolution is inconclusive, due to its sampling bias. Skitter is not sampling the periphery of the Internet and so cannot see evolutionary changes in the whole AS topology. Based on our evidence, we cannot make claims about the relative change of the core compared to the edge, as we can with the UCLA dataset.

We insist on the fact that the purpose of this paper is not to blame the Skitter dataset for its limited coverage of the AS topology, as it aims at sampling the router-level topology. Datasets like Skitter that rely on active probing do provide some topological information not visible from BGP data, as will be shown in Section 4.

3.2 UCLA

We now examine the evolution of the AS topology using 52 snapshots, one per month, from January 2004 to April 2008. This dataset, referred to in this paper as the UCLA dataset, comes from the Internet topology collection[15] maintained by Oliviera *et al.* [16]. These topologies are updated daily using data sources such as BGP routing tables and updates from RouteViews, RIPE,[16] Abilene[17] and LookingGlass servers. Each node and link is annotated with the times it was first and last observed. Note that due to the multiple sources of data used by the UCLA dataset, there is a risk of pollution and bias when combining such differing data sources, which may contain inconsistencies or outdated information.

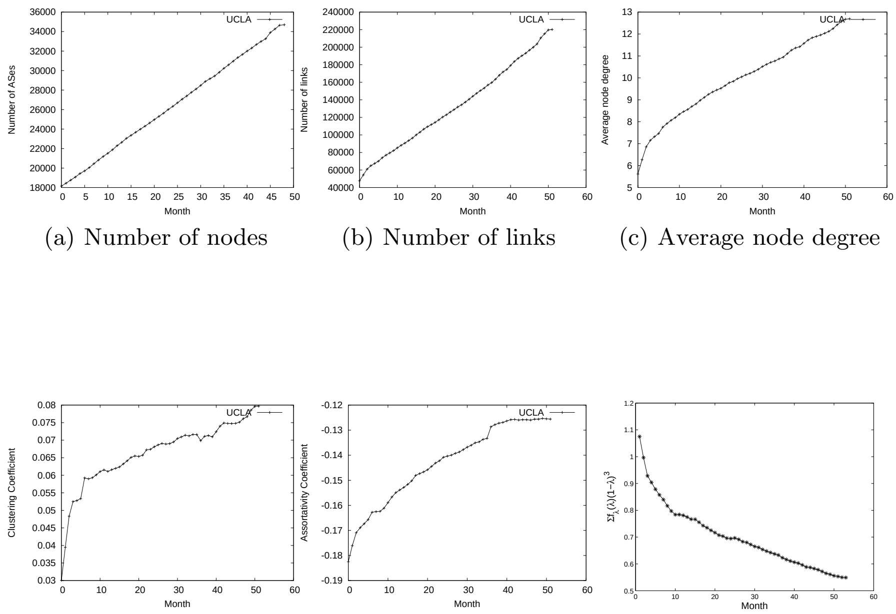

(a) Number of nodes (b) Number of links (c) Average node degree

(d) Clustering coefficient (e) Assortativity coefficient (f) $\omega(G, 3)$

Fig. 3. Topological metrics for UCLA AS topology.

Figure 3 presents the evolution of the same set of topological metrics as Figure 1, over 4 years of AS topologies in the UCLA dataset.

[15] `http://irl.cs.ucla.edu/topology/`
[16] `http://www.ripe.net/db/irr.html`
[17] `http://abilene.internet2.edu/`

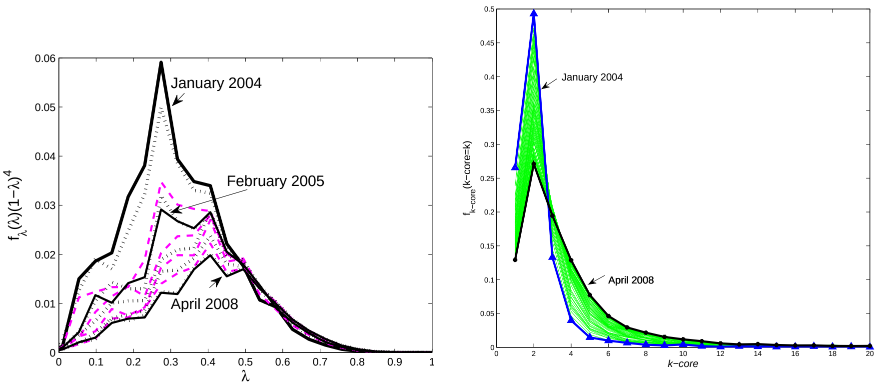

(a) WSD, UCLA topology. (b) k-core proportions, UCLA topology.

Fig. 4. Clustering and spectral features of UCLA topology

The UCLA AS topologies display a completely different evolution compared to the Skitter dataset, more consistent with expectations. As the three upper graphs of Figure 3 show, the number of ASes, AS edges, and the average node degree are all increasing, as expected in a growing Internet.

The increasing assortativity coefficient indicates that ASes increasingly peer with ASes of similar degree. The preferential attachment model seem to be less dominant over time. This trend towards a less disassortative network is consistent with more ASes bypassing the tier-1 providers through public IXPs [8], hence connecting with nodes of similar degree. Another explanation for the increasing assortativity is an improvement in the visibility of non-core edges in BGP data. We will see in Section 4 that the sampling of core and non-core edges by UCLA and Skitter biases the observed AS topology structure. Contrary to the case of Skitter, $\omega(G, 3)$ for UCLA decreases over time. As a weighted clustering metric, $\omega(G, 3)$ indicates that the transit part of the AS topology is actually becoming relatively sparser over time compared to the periphery. Increasing local peering with small ASes in order to reduce the traffic sent to providers decreases both the hierarchy induced by strict customer-provider relationships, and in turn decreases the number of 3-cycles on which $\omega(G, 3)$ is based.

If we look closely at Figure 4(a), we see a spectrum with a large peak at $\lambda = 0.3$ in January 2004, suggesting a strongly hierarchical topology. As time passes, the WSD becomes flatter with a peak at $\lambda = 0.4$, consistent with a mixed topology where core and non-core are not so easily distinguished.

Figure 4(b) shows the proportion of nodes in each k-core as a function of k. There are 52 plots shown as a smooth transition between the first and last plots, emphasized with bold curves. The distribution of k-cores moves to the right over time, indicating that the proportion of nodes with higher connectivity is increasing over time. This adds further weight to the conclusion that the

UCLA dataset shows a weakening hierarchy in the Internet, with more peering connections between nodes on average.

4 Reconciling the Datasets

The respective evolutions of the AS topology visible in the Skitter and UCLA datasets differ, as seen from topological metrics. Skitter shows an AS topology that is becoming sparser and more hierarchical, while UCLA shows one that is becoming denser and less hierarchical. Why do these two datasets show such differences? The explanation lies in the way Skitter and UCLA sample different parts of the AS topology: Skitter sees a far smaller fraction of the complete AS topology than UCLA, and even UCLA does not see the whole AS topology [15]. A far larger number of vantage points than those currently available are likely to be necessary in order to reach almost complete visibility of the AS topology [17].

To check how similar the AS topologies of Skitter and UCLA are, we computed the intersection and the difference between the two datasets in terms of AS edges and ASes. We used a two-year period from January 2006 until December 2007. In Table 1 we show the number of AS edges and ASes that Skitter and UCLA have in common during some of these monthly periods (labeled "intersection"), as well as the number of AS edges and ASes contributed to the total and coming from one of the two datasets only (labeled "Skit-only" or "UCLA-only"). We observe a steady increase in number of total ASes and AS edges seen by the union of the two datasets. At the same time, the intersection between the two datasets decreases. In late 2007, Skitter had visibility of less than 25% of the ASes and less than 10% of the AS edges seen by both datasets. As Skitter aims at sampling the Internet at the router-level, we should not expect that it has a wide coverage of the AS topology. Such a limited coverage is however surprising, given the popularity of this dataset. Note that Skitter sees a small fraction of all AS edges, which is not seen by the UCLA dataset. This indicates that there is potential in active topology discovery to complement BGP data.

From Table 1, we may conclude that the Skitter dataset is uninteresting. To the contrary, the relatively constant, albeit decreasing, sampling of the Internet core by Skitter gives us a clue about which part of the Internet is responsible for its structural evolution.

In Table 2 we show the number of AS edges belonging to the tier-1[18] mesh (labeled "T1 mesh") as well as other AS edges where a tier-1 appears. More than 30% of the AS edges sampled by Skitter cross at least a tier-1 AS, against about 15% for UCLA. Both dataset see almost all AS edges from the tier-1 mesh. Note that the decrease in the number of AS edges in which a tier-1 appears in Skitter is partly related to IP to AS mapping issues for multi-origin ASes [8].

The evolutions of the AS topology observed by the Skitter and UCLA datasets are not inconsistent. Rather, the two datasets sample differently, the AS topol-

[18] We rely on the currently accepted list of 12 tier-1 ASes that provide transit-only service: AS174, AS209, AS701, AS1239, AS1668, AS2914, AS3356, AS3549, AS3561, AS5511, AS6461, and AS7018.

	Autonomous Systems				AS Edges			
Time	**Total**	**Intersect.**	**Skit-only**	**UCLA-only**	**Total**	**Intersect.**	**Skit-only**	**UCLA-only**
Jan. 2006	25,301	32.6%	0%	67.4%	114,847	15.4%	5.3%	79.3%
Mar. 2006	26,007	31.6%	0%	68.4%	118,786	14.9%	4.4%	80.7%
May. 2006	26,694	30.5%	0%	69.5%	124,052	13.8%	4.6%	81.5%
Jul. 2006	27,396	29.5%	0%	70.5%	128,624	13.2%	3.7%	83.1%
Sep. 2006	28,108	28.7%	0%	71.3%	133,813	12.6%	3.4%	84.0%
Nov. 2006	28,885	27.9%	0%	72.1%	139,447	12.4%	3.4%	84.2%
Jan. 2007	29,444	27.2%	0%	72.8%	144,721	11.6%	3.1%	85.3%
Mar. 2007	30,236	26.5%	0%	73.5%	151,380	11.2%	3.0%	85.8%
May. 2007	30,978	25.6%	0%	74.4%	157,392	10.5%	2.7%	86.8%
Jul. 2007	31,668	25.9%	0%	86.1%	166,057	10.0%	3.8%	86.2%
Sep. 2007	32,326	24.5%	0%	75.5%	168,876	9 .7%	2.5%	87.8%
Nov. 2007	33,001	23.9%	0%	76.1%	174,318	9 .5%	2.2%	88.3%

Table 1. Statistics on AS and AS edge counts in the intersection of both Skitter and UCLA datasets, and for each dataset alone.

	Skitter			UCLA		
Time	**Total**	**T1 mesh**	**Other T1**	**Total**	**T1 mesh**	**Other T1**
Jan. 2006	23,805	66	7,498	108,720	64	19,149
Mar. 2006	22,917	66	7,289	113,555	64	19,674
May. 2006	22,888	64	7,504	118,331	64	20,143
Jul. 2006	21,740	65	7,192	123,842	64	20,580
Sep. 2006	21,400	65	6,974	129,228	64	21,059
Nov. 2006	22,034	66	7,159	134,636	65	21,581
Jan. 2007	21,345	65	6,898	140,216	65	22,531
Mar. 2007	21,366	65	6,774	147,000	65	23,194
May. 2007	20,738	65	6,694	153,156	65	23,769
Jul. 2007	22,972	65	6,838	159,792	65	24,310
Sep. 2007	20,570	64	6,510	164,770	65	24,888
Nov. 2007	20,466	64	6,430	170,431	65	25,480

Table 2. Coverage of tier-1 edges by Skitter and UCLA.

ogy, leading to different bias. A large fraction of the AS topology sampled by Skitter relates to the core, i.e., edges containing at least a tier-1 AS. With its wider coverage, UCLA observes a different evolution of the AS topology, with a non-core part that grows more than the core. The evolution seen from the UCLA dataset seems more likely to reflect the evolution of the periphery of the AS topology. The non-core part of the Internet is growing and is becoming less and less hierarchical. We wish to point out that, despite a common trend towards making a *union* of datasets in our networking community, such simple addition is not appropriate for the UCLA and Skitter datasets. Each dataset has its own biases and measurement artifacts. Combining them blindly will only add these biases together, potentially leading to poorer quality data. Further research is required in order to devise a correct methodology that takes advantage of different datasets obtained from different sampling processes.

The above observations suggests that the Internet, once seen as a tree-like, disassortative network with strict power-law properties [6], is moving towards an assortative and highly inter-connected network. Tier-1 providers have always been well connected, but the biggest shift is seen at the Internet's periphery where content providers and small ISPs are aggressively adding peering links among themselves using IXPs to avoid paying transit charges to tier-1 providers. Content distribution networks are partly the reason behind such changes [13].

A different view of the Internet evolution can be obtained using the WSD, shown in Figures 2(a) and 4(a). One possible cause for this behavior is increased mixing of the core and periphery of the network, i.e. the strict tiered hierarchy is becoming less important in the network structure. This is given further weight by studies such as [15] which show that the level of peering between ASes in the Internet has greatly increased during this period, leading to a less core-dominated network. Given that a fraction of AS edges are not visible from current datasets and that visibility is biased towards a better visibility of customer-provider peering relationships, we believe that our observations actually underestimate the changes in the structure of the AS topology. Using a hierarchical and preferential attachment-based model to generate synthetic AS topologies is likely to be less and less justified than ever. The AS topology structure is becoming more complex than in the past.

5 Related work

In this section we outline related work, classified into three groups: evolution of the AS topology, spectral graph analysis of the AS topology, and analysis of the clustering features of the AS topology.

Dhamdhere and Dovrolis [4] rely on available estimation methods for type of relationships between ASes in order to analyze the evolution of the Internet ecosystem in last decade. They believe the available historic datasets from RouteViews and RIPE are not sufficient to infer the evolution of peering links, and so they restrict their focus to customer-provider links. They find that 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 multi-homing degree of transit and content/access providers. Relying on geo-location tools, they find that the AS ecosystem is now larger and more dynamic in Europe than in North America. In our paper we have relied on two datasets, covering a more extensive set of links and nodes, in order to focus on structural growth and evolution of the Internet. We use a large set of graph-theoretic measures in order the focus on the behavior of the topology. Due to inherent issues involved with inference of node locations and types of relationships [11], we treat the AS topology as an undirected graph.

Shyu *et al.* [18] study the evolution of a set of topological metrics computed on a set of observed AS topologies. The authors rely on monthly snapshots extracted from BGP RouteViews from 1999 to 2006. The topological metrics they study are the average degree, average path length, node degree, expansion, resilience, distortion, link value, and the Normalized Laplacian Spectrum. They find that the metrics are not stable over time, except for the Normalized Laplacian Spectrum. We explore this metric further by using WSD.

Oliveira et al. [16] look at the evolution of the AS topology as observed from BGP data. Note that they do not study the evolution of the AS topology structure, only the nodes and links. They propose a model aimed at distinguishing real changes in ASes and AS edges from BGP routing observation artifacts. We use the extended dataset made available by the authors, in addition to 7 years of AS topology data from an alternative measurement method.

6 Conclusions

In this paper we presented a study of two views of the evolving Internet AS topology, one inferred from traceroute data and the other from BGP data. We exposed discrepancies between these two inferred AS topologies and their evolution. We reconciled these discrepancies by showing that the topologies are not directly comparable as *neither* method sees the entire Internet topology: BGP data misses some peering links in the core which traceroute observes; traceroute misses many more peering links than BGP in the periphery. However, traceroute and BGP data do provide complementary views of the AS topology.

To remedy the problems of decreasing coverage by the Skitter traceroute infrastructure and the lack of visibility of the core by UCLA BGP data, significant improvements in fidelity could be achieved with changes to the existing measurement systems. The quality of data then collected by the traceroute infrastructure would benefit from greater AS coverage, while the BGP data would benefit from data showing intra-core connectivity it misses today. We acknowledge the challenges inherent in these improvements but emphasize that, without such changes, the study of the AS topology will forever be subject to the vagaries of imperfect and flawed data. Availability of traceroute data from a larger number of vantage points, as attempted by the Dimes project, will hopefully help remedy these issues. However, even such measurements have to be done on a very large scale, and ideally performed both from the core of the network (like Skitter), as well as the edge (like Dimes). Efforts in better assessment of the biases inherent to the measurements are also necessary.

In an effort to provide a better perspective on the changing structure of the AS topology, we used a wide range of topological metrics, including the newly introduced weighted spectral distribution. Our analysis suggests that the core of the Internet is becoming less dominant over time, and that edges at the periphery are growing more relative to the core. The practice of content providers and content distribution networks seeking connectivity to greater numbers of ISPs at the periphery, and the rise of multi-homing, both support these observations. Further, we observe a move away from a preferential attachment, tree-like disassortative network, toward a network that is flatter, highly-interconnected, and assortative. These findings are also indicative of the need for more detailed and timely measurements of the Internet topology, in order to build up on works such as [5], focusing on the economics of the structural changes such as institutional mergers, multi-homing and increasing peering relationships.

Acknowledgements

We thank Mickael Meulle for his help with the BGP datasets.

References

1. J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani. k-core decomposition of Internet graphs: hierarchies, self-similarity and measurement biases. *Networks and Heterogeneous Media*, 3:371, 2008.
2. R. Bush, J. Hiebert, O. Maennel, M. Roughan, and S. Uhlig. Testing the reachability of (new) address space. In *Proceedings of the 2007 SIGCOMM workshop on Internet network management (INM'07)*, 2007.
3. S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of Internet topology using k-shell decomposition. *PNAS*, 2007.
4. A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the Internet ecosystem. In *Proceedings of ACM/Usenix Internet Measurement Conference (IMC) 2008*, 2008.
5. N. Economides. The economics of the Internet backbone. *NYU, Law and Research Paper No. 04-033; and NET Institute Working Paper No. 04-23*, June 2005.
6. M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of ACM SIGCOMM 1999*, pages 251–262, Cambridge, Massachusetts, United States, 1999.
7. D. Fay, H. Haddadi, S. Uhlig, A. W. Moore, R. Mortier, and A. Jamakovic. Weighted Spectral Distribution. *To appear in IEEE/ACM Transactions on Networking (TON).*
8. P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *Proceedings of Passive and Active Measurement Conference (PAM)*, April 2008.
9. H. Haddadi, D. Fay, A. Jamakovic, O. Maennel, A. W. Moore, R. Mortier, M. Rio, and S. Uhlig. Beyond node degree: evaluating AS topology models. Technical Report UCAM-CL-TR-725, University of Cambridge, Computer Laboratory, July 2008.
10. H. Haddadi, D. Fay, A. Jamakovic, O. Maennel, A. W. Moore, R. Mortier, M. Rio, and S. Uhlig. On the Importance of Local Connectivity for Internet Topology Models. In *21st International Teletraffic Congress (ITC 21)*, 2009.
11. H. Haddadi, G. Iannaccone, A. Moore, R. Mortier, and M. Rio. Network topologies: Inference, modelling and generation. *IEEE Communications Surveys and Tutorials*, 10(2), 2008.
12. B. Huffaker, D. Andersen, E. Aben, M. Luckie, k Claffy, and C. Shannon. The skitter as links dataset, 2001-2007.
13. C. Labovitz, S. Iekel-Johnson, D. McPherson, F. J. J. Oberheide, and M. Karir. ATLAS Internet Observatory 2009 Annual Report. *NANOG47, http://tinyurl.com/yz7xwvv*, June 2009.
14. M. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20):871–898, 2002.

15. 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 *ACM SIGMETRICS*, Annapolis, USA, June 2008.
16. R. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of Internet AS topology. In *Proceedings of ACM SIGCOMM 2007*, Kyoto, Japan, Aug. 2007.
17. M. Roughan, S. J. Tuke, and O. Maennel. Bigfoot, sasquatch, the yeti and other missing links: what we don't know about the as graph. In *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 325–330, New York, NY, USA, 2008. ACM.
18. L. Shyu, S.-Y. Lau, and P. Huang. On the search of Internet AS-level topology invariants. In *Proceedings of IEEE Global Telecommunications Conference (GLOBECOM) 2006*, pages 1–5, San Francisco, CA, USA, 2006.
19. L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proceedings of IEEE Infocom 2002*, June 2002.
20. S. Zhou. Characterising and modelling the Internet topology, the rich-club phenomenon and the PFP model. *BT Technology Journal*, 24, 2006.