

Quantifying the Importance of Vantage Points Distribution in Internet Topology Measurements

Yuval Shavitt
School of Electrical Engineering
Tel-Aviv University, Israel
Email: shavitt@eng.tau.ac.il

Udi Weinsberg
School of Electrical Engineering
Tel-Aviv University, Israel
Email: udiw@eng.tau.ac.il

Abstract—The topology of the Internet has been extensively studied in recent years, driving a need for increasingly complex measurement infrastructures. These measurements have produced detailed topologies with steadily increasing temporal resolution, but concerns exist about the ability of active measurement to measure the true Internet topology. Difficulties in ensuring the accuracy of every individual measurement when millions of measurements are made daily, and concerns about the bias that might result from measurement along the tree of routes from each vantage point to the wider reaches of the Internet must be addressed. However, early discussions of these concerns were based mostly on synthetic data, oversimplified models or data with limited or biased observer distributions.

In this paper, we show the importance that extensive sampling from a broad distribution of vantage points has on the resulting topology and bias. We present two methods for designing and analyzing the topology coverage by vantage points: one, when system-wide knowledge exists, provides a near-optimal assignment of measurements to vantage points; while the second one is suitable for an oblivious system and is purely probabilistic.

The majority of the paper is devoted to a first look at the importance of the distribution’s quality. We show that diversity in the locations and types of vantage points is required for obtaining an unbiased topology. We analyze the effect that broad distribution has over the convergence of various autonomous systems topology characteristics. We show that although diverse and broad distribution is not required for all inspected properties, it is required for some. Finally, some recent bias claims that were made against active traceroute sampling are revisited, and we empirically show that diverse and broad distribution can question their conclusions.

I. INTRODUCTION

The study of the topological structure of the Internet, ranging from the finest IP-level to the coarsest Autonomous-Systems (AS) level, is the driving force of several measurements effort in recent years.

There are two common techniques for Internet topology measurements, passive and active. RouteViews [1] is the major passive measurement project; it relies on collection of BGP announcements and updates from a few tens of vantage points (vps). Most other topology measurement projects rely on active probing, mostly using dedicated instrumentation boxes (e.g., Archipelago [2]) or utilize PlanetLab nodes [3] (e.g.,

iPlane [4] and RocketFuel [5]). A third approach is to use community-based software agents. DIMES [6] is currently the only project that deploys a large number of software agents and maintains an active community of participants.

Distribution based on dedicated hardware boxes is often limited in quantity. Using PlanetLab allows increasing the number of vps, but there are mainly sited in academic networks. However, both approaches create a relatively stable and consistent output that is more easily analyzed. Community-based projects benefit from the contribution of a large and widespread community, but often produce intermittent results that are more challenging to analyze. Currently, there are three operational distributed active probing infrastructures, namely Archipelago, iPlane and DIMES. The data used in this paper is obtained from DIMES and iPlane. Both are highly distributed active measurement infrastructures with hundreds of measurements points.

It is already accepted that attempting to infer the Internet topology from a single or even a few vps leads to an incomplete [7], [6] and, even more important, biased topologies [8], [9]. A problem with previous work that discovered bias and attempted to quantify it, is the usage of either synthetic networks or real data obtained from infrastructures that suffers from the problems mentioned above.

This paper aims to bridge this gap by studying the effects of the measurement distribution on the quality of the observed topology. We begin with a presentation of two methods for assignment of measurement activities to vantage points, one is suitable for stable infrastructures like iPlane or Ark, where there is full knowledge of the vp capability; the other is suitable for community-based efforts where measurement points may be active for short periods. The first (omniscient) method is useful for optimizing the assignment of measurements, while the second (oblivious) method is random and useful for analysis of the coverage abilities of the measurement effort given a time frame and number of participating vantage points. Of course, in both cases increasing the number of vps and the number of measurements per vp significantly improves the topology coverage.

However, this alone is later shown to be insufficient since diversity in the location and the type of vps also play a crucial role in topology discovery. While some question the possibility of cleaning up data from a community-based infrastructure

[†]This work was partially funded by the Evergrow, OneLab II and the MOMENT consortia, which are partly financed by the European Commission; the Israeli Science Foundation, center of knowledge on communication networks; and the Israel Internet association. Special thanks to Scott Kirkpatrick and the anonymous reviewers for the insightful comments and suggestions.

[8], we describe various simple filtering techniques and show that given a sufficiently diverse and broad distribution (in sense of geography, type and quantity), it is possible to obtain data with comparable quality to other controlled-deployed infrastructure. The filtered data is then used to explore the benefits of having a broad distribution in order to reevaluate some recent bias claims. Moreover, various properties of the Autonomous Systems (AS) graph are analyzed, showing that broad distribution can further assist in reducing the bias of the results. Techniques recently presented by Latapy and Magnien [10] are employed to show that some graph properties require more than 40 different vps in order to converge to a value that represents the measured topology but surprisingly, some require only few vps to converge.

There are many efforts that attempt to measure the Internet and abundant of papers analyzing different aspects of the resulting data. However, only few papers perform in-depth analysis of the measurement infrastructures themselves in general, and highly distributed projects in particular. A notable work of Mahadevan *et al.* [11] performed an analysis of the AS-level topology using three different data sources– Skitter, RouteViews and WHOIS. The authors showed that topologies created from active traceroutes and passively collected BGP are similar to one another but differ substantially from the user-maintained WHOIS topology.

A recent work of Oliveira *et al.* [12] focuses on creating an evolutionary model of the AS topology and provides an evaluation of the proposed model using different data sources. The authors use BGP as the basis for their analysis and extend it with IRR data and active probing including Skitter, DIMES and iPlane, and mention that although active traceroute probing is an important source for topology information, it has a problem of broadness (covering all sampled topology) and freshness (updating the destinations list).

The remaining of this paper is organized as follows. Section II provides two methods for assigning measurements to vps, showing the importance of distribution to Internet topology coverage. Section III examines the qualitative aspects of diverse and broad distribution and revisits some conclusions of recent bias claims. We conclude the paper in section IV.

II. TOPOLOGY K -COVERAGE

A fundamental objective of all measurement projects is the ability to probe as much of the Internet topology as possible. Usually, active measurement infrastructures attempt to maximize their coverage of some topology elements (prefixes, ASes, PoPs, cities, etc.). We present two methods for selecting which set of elements each vp needs to probe in order to gain maximal topology coverage. We define an Internet topology element as being K -Covered when it is probed from at least K different vps, for a given $K \geq 1$.

The omniscient method requires knowledge about the ability of vps to probe the inspected elements. This method results in an optimal assignment of elements to vps but is computationally hard. The oblivious method uses a randomized assignment

and therefore does not requires knowledge about the ability of the vps to probe elements.

Note that neither model considers the path traversed from the vp to the destination, which obviously contributes additional probing of other elements. Therefore, these models provide a lower bound to the number of times each element is probed.

This omniscient method is appropriate to infrastructures that statically deploy measuring devices, such as iPlane and Archipelago. Let P be the set of elements we want to measure during a time frame T and N be the set of actively measuring vps. We describe an integer program that aims to find the optimal assignment of elements to vps given that any given vp i can perform at most $M_i(T)$ measurements and can probe a pre-known partial set of elements $P_i(T) \subseteq P$.

Maximizing the number of K -covered Internet topology elements requires solving the following optimization problem:

$$\begin{aligned} \text{Maximize} \quad & \sum_{j=1}^{|P|} Y_j \\ \text{Subject to:} \quad & \sum_{j=1}^{|P|} X_{ij} \leq M_i(T), \quad 1 \leq i \leq N \\ & X_{ij} \leq Z_{ij}, \quad 1 \leq i \leq N, 1 \leq j \leq |P| \\ & Y_j = \begin{cases} 1, & \sum_{i=1}^N X_{ij} \geq K \\ 0, & \text{otherwise} \end{cases} \end{aligned}$$

Where $Z_{ij} = 1$ when element $j \in P_i(T)$ and 0 otherwise, and $X_{ij} = 1$ when element j is assigned to be measured by vp i and 0 otherwise. When $K=1$, this is a standard integer program which is in NP . Therefore, the general problem is at least as hard.

The oblivious method is appropriate to community-based projects, such as DIMES, where there is churn in the set of active vps, and vps cannot be assumed continuously active. Pre-assigning prefixes to vps might result in poor coverage as some become inactive, while new vps will not be taken into account. Instead, each vp that is active is assigned a set of random targets (topology elements) $P_i(T)$, each is uniformly selected from the complete set P . In the DIMES implementation, a vp is known to be active only when it connects to the server. Since it is not known for how long it will remain active, the set $P_i(T)$ is split into small subsets of random elements. However, without harming the results we simplify the analysis and assume that $P_i(T) = P_{i_1}(T) \cup P_{i_2} \cup \dots \cup P_{i_n}(T)$. Given a total of $M = \sum_{i=1}^N M_i(T)$ measurements, the expected number of elements that are K -covered is given by [13]:

$$E(K, M) = |P| \cdot \sum_{i=K}^{\infty} \binom{M}{i} \left(1 - \frac{1}{|P|}\right)^{M-i} \left(\frac{1}{|P|}\right)^i$$

Notice that although this model does not enforce at least K different probing vps for each element, given that N is large and $|P_i(T)| \ll |P|$, the probability that the same element

will be probed twice by the same vp is small and equal to $1 - E(2, |P_i(T)|) / |P_i(T)| = 1 - (1 - 1/|P|)^{|P_i(T)|} \simeq 0$.

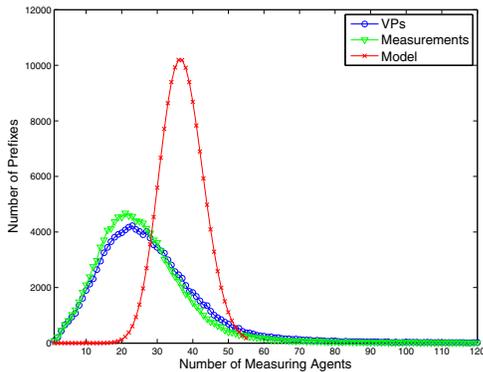


Fig. 1. Comparison of the oblivious model to the assignment of prefixes to agents using DIMES

We compare the oblivious model to the measurements results of an experiment conducted using DIMES agents (vps) (Section. III-A provides a brief overview of DIMES). The experiment utilizes 906 agents that together performed over 4.5 million traceroutes. Each agent performed traceroute measurements to a random set of scripts, each containing 60 destination IP addresses that were randomly selected from a list that contains three IP addresses from each of the approximately 155 thousand prefixes that exist in DIMES.

Figure 1 shows the distribution of the number of measurements per prefix as expected by the model, the actual number of measurements per prefix and the distribution of different vps per prefix as performed by DIMES. As shown above, the distribution of vps and measurements is similar, showing that indeed the probability for a vp to measure the same prefix more than once is low.

However, although the experiment’s distribution exhibits the same normal distribution as the model, it is shifted to the left, scaled down and has a long tail (x-axis reaches over 700 but was truncated for brevity). This is mainly a result of the three IP addresses per prefix and double random selection, causing some prefixes to be measured from hundreds of agents (an event with extremely low theoretical probability according to the model) and some not to appear at the list at all. This caused the very long tail, resulting in the scaled down and left shift of the distribution. Note that in November 2008, DIMES changed the way it assigns IP addresses to scripts and scripts to agents. A list of one IP address per prefix is created, randomly shuffled and split into files, each containing 60 IP addresses. These files are assigned to agents in a round-robin fashion, resulting in a near-equal probing of each IP address.

Since $|P|$ is relatively constant (roughly 170,000 prefixes, 25,000 ASes, etc) the ability to increase $E(K, M)$ is possible only by increasing M . This can be achieved by increasing the number of measurements per vp $M_i(T)$, increasing the analyzed time-frame T or increasing the number of vps N . Increasing $M_i(T)$ is difficult since it can cause high load on

vps and might cause the generated traffic to be flagged as hostile. Increasing T might result in capturing transient effects that will cause the observed topology to be incorrect. Increasing the number of vps is feasible, especially in community-based infrastructures where it is relatively easy (although not trivial) to convince additional users to participate and act as new vps.

Both models show that having a large number of active vps can help increasing the number of K -covered topology elements. Although having large number of vps satisfies this model, the remainder of this paper examines additional aspects of vps distribution and shows that for some, this alone is not enough, but rather a high diversity in the locations and types of vps is also required to obtain an unbiased topology.

III. VANTAGE POINTS DISTRIBUTION ANALYSIS

As noted before, increasing the number of vps is a challenging task for all Internet measurement projects, either due to the need of purchasing and deploying new specialized machines, or convincing users to install an agent on their PCs. It was a natural assumption that increasing the number of vps is important for improving the observed network topology until Barford *et al.* [14] showed that the utility of adding vps beyond the second one, quickly diminishes in terms of ASes (nodes) and AS-links (edges). However, more recently Shavitt and Shir [6] argued that although the utility indeed diminishes, the data from adding hundreds and thousands of vps have a substantial effect on the resulting topology. We wish to study empirically how using a large number of vps affects the observed topology.

For this end, we use traceroute data from DIMES collected during the week 12 of 2008. Due to its community-based design, DIMES may exhibit unpredicted behaviors, depending on the activity of its community. The following section briefly highlights some of the important aspects of DIMES previously discussed in [6] and provides the understanding of how to detect and filter out unwanted behavior, leading towards a more accurate data analysis.

A. Infrastructure Overview

DIMES performs measurements using hundreds of software agents installed on users’ PCs. Agents perform measurements by following a script that is sent to them from a central server. An agent can perform Traceroute and Ping measurements using either ICMP or UDP packets, with a default of two measurements per minute, inducing minimal bandwidth overhead. Upon completion of a script, an agent submits the results and requests a new script to perform.

The measurement scripts aim to cover the entire IP prefix space. DIMES collects the list of prefixes from the RouteViews BGP tables, and using each prefix a set of IP destinations is constructed. A single script includes traceroute and ping commands to 60 destination IP addresses.

Raw measurement data that is reported back to the server is filtered in order to remove measurement artifacts that can later cause analysis mistakes. Traceroutes that exhibit some known problems [15], namely routing loops and appearance

of the destination address in the middle of the traceroute, are discarded for analysis. These measurements account for less than 0.1% from the total number of traceroutes performed.

On average, an agent performs approximately 10,000 weekly traceroute measurements. Agents that performed less than 100 measurements were filtered out (approximately 200 such agents). These are either agents that run for less than two hours during the entire week, or agents that send the server results that get filtered out due to some of the known traceroutes problems described above. Although the former type may provide some interesting measurements we prefer not to include them in this study as they require special care in order to avoid “noise”.

B. Data Filtering

When using DIMES data, a vp is simply an AS that homes one or more agents. We ignore cases where there are many agents within a single large AS. Since vps often change, creating the set of vps (denoted by VP) requires the discovery of the ASes from which agents perform measurements. This discovery is done by following each of the traceroutes, until reaching a hop with a routable IP address that can be resolved into a valid AS. We limit the search to the first four hops in the path to minimize incorrect vp identification. However, it is not completely error proof since a loss of measurements from the routers in the hosting AS will result in mis-identification.

This analysis reveals that there are approximately 28% of the agents that appear to be homed in more than one vp; 15% of the agents are homed in more than 5. However, most of the measurements of an agent are performed from one vp and the rest of the vps appear to have a few (usually less than ten) measurements which is a clear sign of mis-identification. Since DIMES is an at-home project, some users install an agent on their laptops, and really do measure from different ASes.

Let $tr(a_i, vp_j)$ be the number of traceroutes performed by agent a_i from vp $vp_j \in VP$. We filter out all the data for which $tr(a_i, vp_j) < T$. Applying this filter for $T = 100$ reduces the number of agents that have more than 5 different vps to less than 4% and removes all the agents that have more than 10 different vps. Filtering with $T = 400$ keeps only agents that have at most 5 vps. There is a clear trade-off between the accuracy of vp identification and loss of data due to over-filtering. Therefore, in this paper we selected $T = 100$ and refer to this as the *filtered* data.

For each vp_i we extract the number of agents, $agents(vp_i)$ and the number of traceroute measurements, $tr(vp_i)$. We then find the degree k_i of its hosting AS in the undirected AS graph, i.e., number of links to neighboring ASes, and calculate the average number of traceroute measurements per degree:

$$tr(k) = \sum_{i \in \{k_i = k\}} tr(vp_i) / |\{k_i = k\}|$$

and average number of agents per degree:

$$agents(k) = \sum_{i \in \{k_i = k\}} agents(vp_i) / |\{k_i = k\}|$$

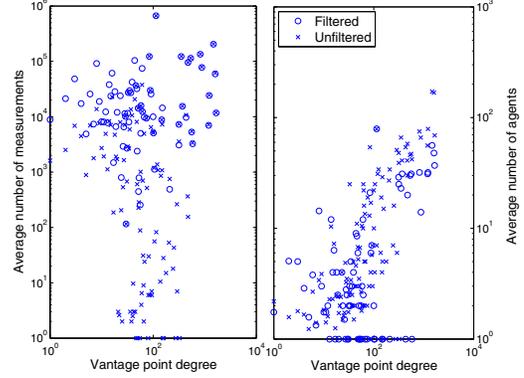


Fig. 2. Average number of agents and measurements per vp degree. Filtered data includes only agents from vps where $tr(a_i, vp_j) \geq 100$

We repeat this calculation for the filtered set, and find $tr^f(k)$ and $agents^f(k)$.

Figure 2 shows that vps with higher degree tend to have more agents (using degrees extracted from RouteViews produced the same results). This indicates that this is not a sampling bias caused by measurement artifact [16] but rather a hint that large ASes tend to “host” more agents. Additionally, vps with a single agent in them are mostly in the mid-level degrees section, and vps with more than one agent are scattered in a wide variety of degrees. Figure 2 shows that vps with mid-level degrees and single agent perform relatively few traceroute measurements, resulting in some of them being filtered out. Other than this, there is no direct relationship between the number of measurements to the vp degree, showing that there is a good distribution of measuring agents across the different vps.

C. Diminishing Returns

We now have the data needed to revisit the diminishing returns claim by examining how the observed topology changes as data from more vps is added. Evaluating the effect of adding vps is done by first building the *local* AS topology as observed from each of the vps, i.e., the AS topology seen by agents that are hosted in each vp.

Creating the AS-level topology from IP-level traceroutes provided from DIMES is achieved by performing AS resolution for each hop in all paths. DIMES performs AS resolution by first applying longest-prefix-matching on BGP tables obtained from the RouteViews archive (using the same week being studied). This resolves approximately 98% of the IP addresses. The remaining 2% are queried against two WhoIs databases, namely RIPE and RADB. This resolves additional 1.5% of the IP addresses. The remaining 0.5% unresolved IP addresses are discarded and do not contribute nodes and edges to the topology. Since this method of IP-to-AS mapping can result in minor distortions of the topology [17], [18], when constructing the complete DIMES AS topology, we only include ASes and links that exist in the local topologies of the

filtered set of vps, VP^f (we drop the f notation since all further analysis uses only the filtered data).

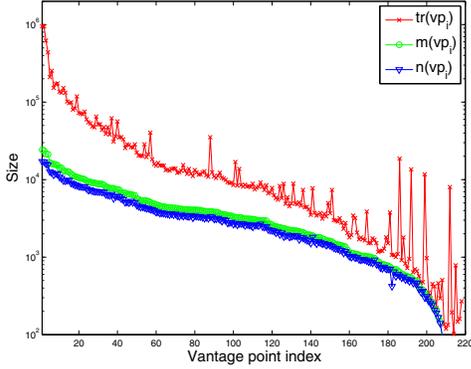


Fig. 3. Vantage points details sorted by number of discovered AS-links

We denote by $tr(vp_i)$, $n(vp_i)$ and $m(vp_i)$ the number of traceroutes performed, nodes and edges in the local topology of vp_i , respectively. We then sort the set VP in a non-increasing order of $m(vp_i)$. Figure 3 shows, for each vp_i , the values of $tr(vp_i)$, $n(vp_i)$ and $m(vp_i)$ in the sort order. There is a high correlation between the number of measurements and the size of the topology. However, this correlation breaks at the tail of the ordered list.

Although this sort order seems quite intuitive, and we use it for further analysis, we point out that it holds a certain bias. VPs that discover the largest local topologies are mostly located in the United States and Europe. Out of the top ten vps, six are in the US and 4 are in Europe¹. The 16th vp is the first in a different geographical region (Israel) and the 35th is in South-Africa. This means that regions outside the areas mostly used by other active measurement projects are considered only later in the sorted vp list. Although the analysis in this paper would benefit from having large local topologies for all vps, from a distributed infrastructure point of view, each vp is responsible only for mapping its vicinity.

D. Aggregation of Local Topologies

Once we have the sorted set VP we build a set of aggregated AS topologies, Agg . An aggregated topology agg_i includes the ASes and links from the set of local topologies, $agg_i = \cup_{j=1}^i vp_j$. For each agg_j we find the number of ASes, n_j , AS links, m_j , and difference in ASes, $\Delta n_j = n_j - n_{j-1}$ and AS links, $\Delta m_j = m_j - m_{j-1}$.

This data is plotted in Figure 4. The figure shows that the number of ASes remains relatively constant after a few vps. The number of AS-links on the other hand gradually increases even when going over 80 vps. Note that there are vps, which do not contribute new ASes or links and are not shown due to the semi-log scale.

¹The first five ASes in the ordered list are Teleinform Russia (AS42286), Verizon (AS19262), UKRTelnet Ukraine (AS6849), Comcast (AS7015) and Qwest (AS209).

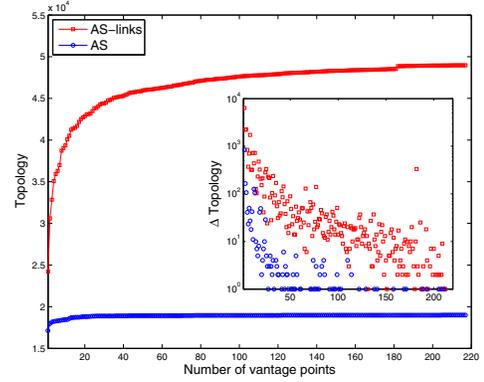


Fig. 4. Number of ASes and AS-links in the aggregated topology (inset shows the change in topology size)

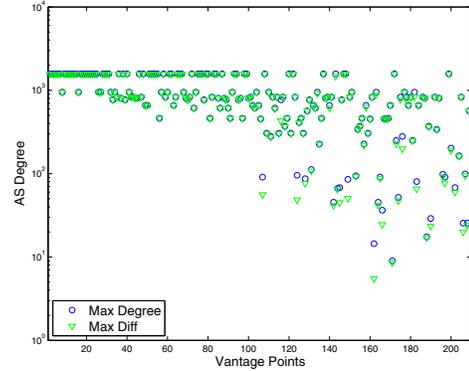


Fig. 5. Maximal AS degree and maximal neighboring ASes degree difference of additional AS links

The figure questions the diminishing returns argument in two ways. First, for AS-links, the return diminishes much slower than was observed in the past. Second, the tail of the distribution is indeed thick. This is especially interesting given the fact that in general as we add more vps we get those with fewer measurements (see Figure 3) due to the sort order. Therefore, the tail of the distribution in Figure 4, where the number of measurements is in the few thousands, is actually a lower bound on the vp contribution.

To further understand the contribution of vps in the tail, we collect all AS links that were added between two consecutive aggregated topologies, agg_j and agg_{j+1} , and the set of ASes adjacent to these links. In this set of ASes, we find the maximal AS degree, and the maximal degree difference between adjacent ASes. The maximal degree shows whether the vp detects links in the core of the Internet. High degree difference indicates that the vp manages to detect “radial” links (meaning customer-to-provider [19] links towards the core) and low degree difference indicates that the vp manages to detect “tangential” links. Figure 5 shows that even vps in the tail manage to discover new links towards the core. However, starting from roughly the 150th vp, we see vps that only

contribute “tangential” links or links connecting low degree ASes to regional providers.

We further study how merging observed topologies from distributed vps affects several graph characteristics that are commonly used in graphs analysis. Following Latapy and Magnien [10], we examine the convergence trend of each analyzed property to the value of the overall AS graph by analyzing the graph characteristics of each AS graph agg_i , $1 \leq i \leq |VP|$. Each graph’s property is plotted to show how the values extracted from the AS graphs agg_i converge to the overall value of the property which is calculated on the complete AS graph using VP . We show that while some properties converge to the overall value using only a few vps, others converge slowly and require many vps to reach the value of the complete topology.

We start by analyzing the **node degree distribution** which is the probability that a randomly selected node is of degree k . Let $n(k)$ be the number of nodes with degree k . The node degree distribution is: $P(k) = n(k)/k$. The degree distribution has become one of the most frequently used Internet topology characteristic [11] since the work of Faloutsos *et al.* [20] that showed that the degree distribution of Internet topologies follows a power-law, meaning $P(k) \sim k^{-\gamma}$, where γ is a positive exponent. We use the closely related Zipf [21] distribution, $n(k) \sim k^{-\alpha}$, and calculate $\gamma = 1/\alpha + 1$. Evaluating how close a sampled degree distribution is to a power-law degree distribution is obtained by using the fact that power-law degree distributions have a theoretical maximum power-law degree property, $k_{max}^{PL} = n^{1/(\gamma-1)}$ [22].

Figure 6(a) and Figure 6(b) show that the average degree and PDF exponent monotonically converge, reaching the vicinity (within 10%) of the overall value after roughly 40 vps, and about 180 vps to come within 1% of the eventual value. Interestingly, the exponent value we get using all vps ($\gamma = 2.197$) is very similar to the one reported by Faloutsos [20] who used a single vp ten years ago ($\gamma \simeq 2.20$).

Figure 6(c) shows that the maximum degree converges even faster, indicating that the first few vps accurately map the highest degree AS, namely Level3 (AS-3356). However, the theoretical maximum degree exceeds the sampled maximum degree after roughly 60 vps, and reaching over 3500. This indicates that the maximal power-law degree measure is not accurate for the AS topology.

Density of a graph is used for measuring how close a graph is to a clique, and is equal to $2m/(n(n-1))$. Figure 6(d) shows that the density monotonically converges. This is expected since Figure 4 shows that the number of links increases while the number of ASes remains relatively stable, resulting in a slowly increasing density.

We decomposed each AS graph into *shells* using the k -pruning algorithm (the higher the shell index of an AS the better is its connectivity [23]). The *nucleus* is the shell with the highest index, k_{max} . Unlike the density, the nucleus index (not shown in the figure) converges to its overall value ($k_{max} = 36$) after less than 80 vps, showing that the density does not change evenly in the graph. Discovering links in the nucleus requires

only a few vps since these core links are traversed by many traceroute paths due to valley-free [19] routing.

Betweenness centrality (bc) is commonly used for measuring the centrality of a node or a link. Node betweenness measures the number of shortest paths passing through a node as an estimate to the potential traffic load on this node assuming uniformly distributed traffic which follows shortest paths [11]. We calculate the maximal betweenness over all ASes in each Agg_j as a measure for possible congested nodes in the graph. In order to compare topologies with different sizes, we normalize the average betweenness by the maximal possible betweenness value, $n_j(n_j - 1)$.

Clustering Coefficient (cc) of a graph measures the local cliquishness of a node’s neighborhood [24]. Simply put, clustering coefficient estimates how close a given node and its immediate neighbors are from being a clique. A graph’s average clustering coefficient is used to estimate how close a graph is to a small-world network, such that graphs with higher average clustering coefficient can be better modeled by small-world network.

Figure 6(e) and Figure 6(f) show that bc and cc converge within 10% of the overall value after roughly 20 vps. Since bc is a measure of load, adding links almost always decreases the average load on ASes (except special cases like the Braess paradox), thus the monotonic descent. Since high-degree ASes are more “central” than low-degree ASes, we expect that tier-1 ASes will have the maximal betweenness values, indicating possibly congested nodes. Indeed, Level3 (AS3356), a tier-1 AS, is the AS with maximal bc for all resulting graphs. Considering that adding radial links reduces the cc while adding tangential links mostly increases it, explains the smooth and non-monotonic cc graph and the gradual increase after roughly 140 vps (these vps contribute mostly tangential links).

E. Sampling Bias

Lakhina *et al.* [16] showed that AS degrees inferred using traceroute-like sampling technique, are highly affected by the location of the vps. The authors claimed that this bias could cause the observed power-law distribution, and suggested as a criteria for detecting bias in traceroute studies to check whether the highest-degree ASes tend to be near the measuring sources. Using this criteria, the authors showed that some commonly used datasets exhibit bias. However, Dall’asta *et al.* [25] explored the origins of the bias and showed that various traceroute exploration strategies produce sampled topologies with minimal bias. More recently, Cohen *et al.* [26] showed that when the underlying graph degree distribution obeys a power-law with an exponent larger than 2, such as the Internet AS-level graph, a tree-like sampling process produces a negligible bias in the sampled degree distribution.

It is expected that achieving a broad distribution of vps will result in a good sampling of the underlying topology so that it will exhibit less bias resulting from the closeness of vps and observed ASes. We calculate the distance (number of hops) in

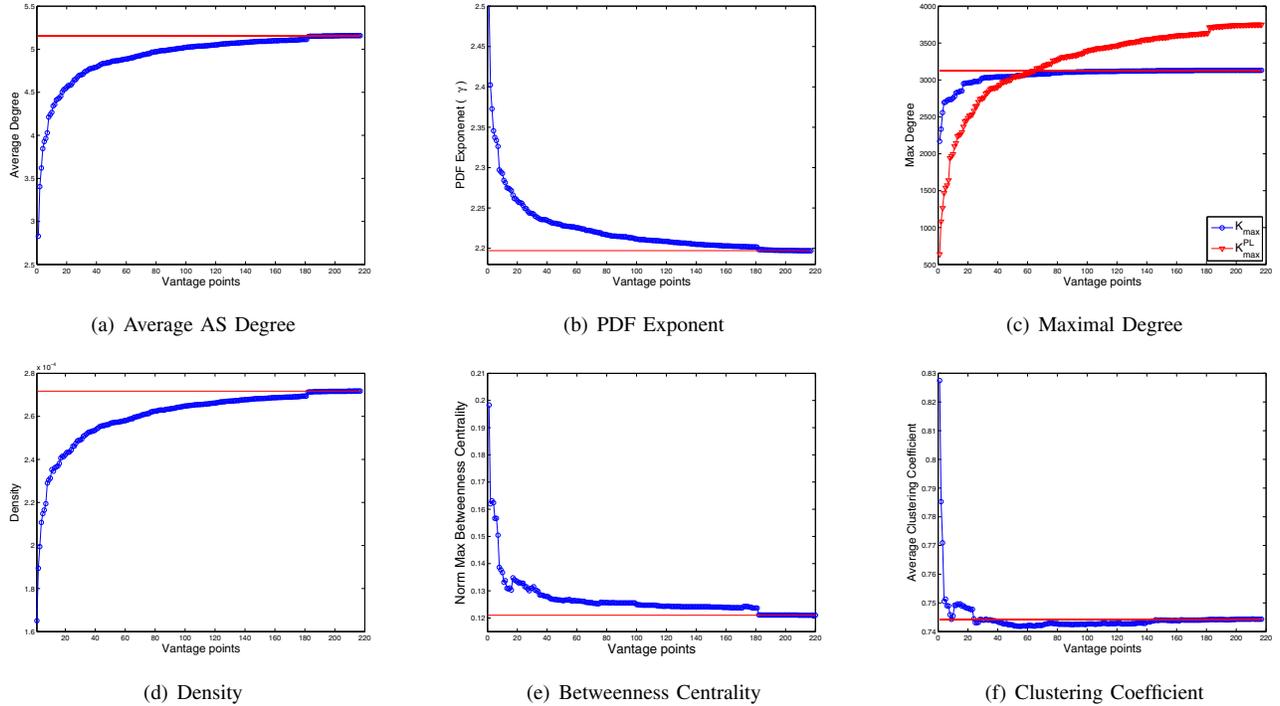


Fig. 6. Topology characteristics analysis of AS graphs agg_i , the horizontal line shows value calculated on the complete AS graph

the valley-free² path between an AS and all vps that observed it using a customized valley-free dijkstra. Calculating a valley-free path requires the inference of the type-of-relationship between adjacent ASes (customer-provider, peers or siblings) in the AS graph. The relationships are evaluated using the near-deterministic type-of-relationship algorithm [27].

First, for each AS we find all the vps that include it in their observed topology (referred to as “observing vps”), and calculate the number of hops from the AS to each of them. Figure 7 shows the average number of observing vps calculated over all ASes with a given degree obtained from DIMES and RouteViews AS-level topologies. As expected, the figure shows that low-degree ASes are observed from much fewer vps than high-degree ASes. This is attributed to the fact that small degree ASes are harder to detect and probe. However, there are a few ASes that have high degree (mostly in RouteViews) and are observed by only few vps, but these are quite rare.

Table I shows the distribution of the number of ASes per distance to the nearest observing vp. The table shows that most ASes are 1 to 3 hops away from the nearest vp, some ASes have a vp that exists in them, and a small fraction of ASes have vps that are 4 or 5 hops away. The resulting average distance is 1.82 hops with a standard deviation of 0.74. This shows that having a well-spread set of vps can reduce sampling bias.

Finally, Figure 8 shows the average distance over all ASes

²A valley-free path follows a strict hierarchical structure – an uphill segment of zero or more customer-to-provider or sibling links, followed by zero or one peer-to-peer links, followed by a downhill segment of zero or more provider-to-customer or sibling links.

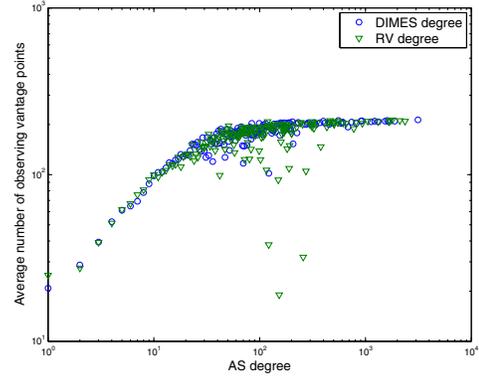


Fig. 7. Number of observing vps per AS degree

TABLE I
DISTANCES FROM AS TO NEAREST VANTAGE POINT

Distance to nearest vp	0	1	2	3	4	5
Number of ASes	213	6365	9329	2869	204	6

with a given degree (using DIMES and RouteViews) to the nearest vp. The resulting standard deviation surrounding the average is less than 1 hop. While it is possible to see a correlation between the distance and AS degrees, the reason for this is not sampling bias but rather probabilistic distribution of agents. As seen in Figure 2 the probability that an agent is located in a given AS is roughly proportional to the degree of the AS. When using degrees from the RouteViews AS graph, the bias is even less noticed, since there are high-degree ASes

that are relatively far away (2 hops) from the nearest vp.

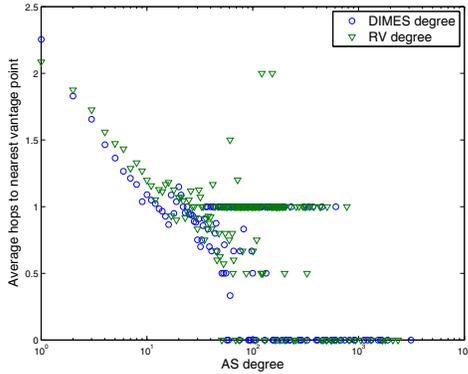


Fig. 8. Number of hops to nearest vp per AS degree

F. Diversity Bias

We further wish to study the effect of vps distribution by examining how the diversity of types of ASes from which measurements are performed affect the observed topology.

We first study the geographical distribution of DIMES and iPlane’s vps. Finding the country of a vp was done by querying the commercial MaxMind database with the IP addresses of the DIMES agents and iPlane’s sources. VPs that are resolved to exist in more than one geographic location are removed from this analysis.

In DIMES, the largest number of vps are in the United States (40.8%), followed by United Kingdom (8%), Russian Federation (7.31%), Germany (6.5%) and Canada (5.8%). Other vps are spread over the entire globe, including most of the European countries, some countries in the Middle East, the Far East, Australia and South Africa. On the other hand, there was only one vp in South America (Argentina) and no vps in Central Africa, China and the Arab countries in the middle east. iPlane’s vps are also mostly located in the US (43.3%). However, iPlane has a wider variety of geographic locations, such as Japan (4.3%), Taiwan (2.8%), South Korea (2.5%) and China (1%) in Asia, Brazil (1 vp) and Uruguay (1 vp) in South America. However, like DIMES, there are no vps in Central Africa and the Arab countries. This analysis shows that both projects have a good geographical distribution of vps that should assist in locating vps in remote location.

We compare the vps of iPlane and DIMES by examining their AS types. The AS type is determined using the work of Dimitropoulos *et al.*[28] that uses a machine learning approach to classify an AS as a large ISP (t1), a small ISP (t2), an academic network (edu), an Internet exchange point (ix), a network information center (nic) which holds important network infrastructure, or a customer (comp) of either the small or large ISPs. We manually classified additional ASes, such as those having a description that contains the word “university” that were missing from the inferred set. In total there are 18,639 classified ASes which does not cover all the

ASes in the datasets but is sufficient for this analysis. ASes with unknown types are ignored since both infrastructures have a similar number of unclassified ASes.

TABLE II
VANTAGE POINTS’ TYPES OF IPLANE AND DIMES

	t1	t2	edu	comp	ix	nic	unknown
iPlane	17	104	117	22	3	5	46
DIMES	23	129	6	15	0	0	47

Table II provides the number of vps for each AS type. It shows that iPlane uses much more academic vps than DIMES. This is mainly the contribution of the PlantLab nodes that iPlane utilizes. DIMES has about 20% more vps in tier-2 ISPs which aligns with its at-home distribution.

TABLE III
ASES’ TYPES FOR WHICH DIMES DEGREES ARE LARGER THAN IPLANE’S (RATIO<1), SMALLER (RATIO>1) AND EQUAL (RATIO=1)

	t1	t2	edu	comp	ix	nic	unknown
ratio<1	20	1487	108	1398	7	34	1286
ratio>1	24	1295	137	1219	9	36	1123
ratio=1	0	1401	248	3304	11	61	2445

We found the degree of each AS in the AS-level graphs of the two infrastructures and calculated the ratio between the degree in the AS-graph of iPlane and DIMES using all the ASes that existed on both graphs (15,886 ASes out of roughly 19,000 on each graph). Table III shows the number of ASes for each type when DIMES degrees are larger than iPlane’s (ratio<1), smaller (ratio>1) and equal (ratio=1).

It clearly shows that DIMES has higher degrees for tier-2 ASes and lower for Universities which exactly correlates to the location of the vps. In addition, DIMES also have more customer ASes (comp) which can also be contributed to the fact that it has more vps that reside in tier-1 and tier-2 ASes that provide connectivity to these customer ASes, therefore its vps can probe them better. Surprisingly, the two infrastructures do not agree on the degree of even a single tier-1 AS. However, recalling that the average degree of t1 ASes is very high and noting that the average ratio is 1.074 (with standard deviation of 0.33), shows that the degree difference is small.

This analysis shows that although increasing the number of vps can help reduce sampling bias, it still does not guarantee unbiased results. Although both iPlane and DIMES have a very broad distribution, the types of ASes in which their vps are located still result an observed topology that is slightly biased towards these types. Overcoming this bias cannot be achieved by simply increasing the number of vps but rather a broad diversity in types is required.

TABLE IV
AVERAGE AS DEGREES USING ROUTEViews, DIMES AND IPLANE FOR PER AS TYPES

	t1	t2	edu	comp	ix	nic	unknown
RouteViews	575.84	9.1	3.4	2.29	14.85	5.42	3.2
DIMES	528.31	8.36	3.34	2.3	7.96	4.06	3.86
iPlane	565.31	8.47	5.83	2.45	22.82	4.12	4.5

Since neither iPlane nor DIMES can be considered as the “ground truth” of the Internet topology, we compare their average AS degrees per AS type to the RouteViews AS graph. RouteViews passively collects BGP messages and is considered as a relatively reliable source for the AS topology. Table IV provides the average AS degree per type and shows that for most AS types both iPlane and DIMES observed a smaller average degree than RouteViews. However, iPlane has a significant higher average degree in academic ASes and exchange points. This can be attributed to iPlane having a high number of vps in these AS types, allowing it to probe them better, further stressing the need of diversity in vps types.

TABLE V
AVERAGE VP DEGREE PER TYPE, CALCULATED ORIGIN ASes OF IPLANE, DIMES AND BOTH

VPs	Avg. Degree	t1	t2	edu	comp	ix	nic
iPlane	iPlane	951.5	36.0	9.4	8.1	172.0	3.5
	DIMES	961.3	27.0	5.1	5.5	28.6	3.5
DIMES	iPlane	904.7	38.0	2.6	9.1	-	-
	DIMES	903.1	50.1	18.6	25.8	-	-
Both	iPlane	1157.0	89.5	3.0	2.0	-	-
	DIMES	1161.0	81.3	14.0	3.0	-	-

Finally, we evaluate how a presence of a vp within an AS affects its observed connectivity to other ASes. Table V provides the average vp degree per type, calculated over origin ASes of iPlane, DIMES and both. As expected, ASes that are a vp mostly result in better probing, therefore exhibit higher inferred average degree. An exception are t1 vps, however, their average degrees are very close to one another, showing that both infrastructures share similar t1 coverage (indeed, out of the 23 t1 vps of DIMES and 17 of iPlane, 13 vps are shared by both). Interestingly, the average degrees of t1 and t2 ASes that host both iPlane and DIMES vps are relatively high, showing that its easier for both projects to locate vps in high degree ASes. This analysis quantifies the assumption that measuring from within a network can help discover more of its links with other networks.

IV. CONCLUSION

This paper presents an analysis of the significance of the distribution of vps in an active Internet measurements infrastructure. We showed that diverse and broad distribution can help overcome sampling bias and uncover hidden parts of the Internet topology. However, even broadly distributed infrastructures still exhibit some bias towards the type of ASes from which measurements are performed, further stressing the need for a diversity of vps types and geographic locations. For this end, community based infrastructures are best posed as their growth potential is practically unlimited. Looking at various commonly analyzed graph properties, we showed that some require more than 40 vps to converge, but surprisingly several others converge with only a few vps.

REFERENCES

[1] University of Oregon Advanced Network Technology Center, “Route Views Project, <http://www.routeviews.org/>.”

[2] CAIDA Macroscopic Topology Project Team, “Caida archipelago, next generation active measurement infrastructure (<http://imdc.datacat.org/collection/>).”

[3] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, “Planetlab: An overlay testbed for broad-coverage services,” *ACM SIGCOMM Computer Communications Review*, vol. 33, no. 3, July 2003.

[4] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, “iPlane: An information plane for distributed services,” in *OSDI*, 2006, pp. 367–380.

[5] N. T. Spring, R. Mahajan, and D. Wetherall, “Measuring ISP topologies with rocketfuel,” in *SIGCOMM*, 2002, pp. 133–145.

[6] Y. Shavitt and E. Shir, “Dimes: Let the Internet measure itself,” *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 5, 2005.

[7] R. Cohen and D. Raz, “The Internet dark matter - on the missing links in the AS connectivity map,” in *IEEE Infocom 2006*, Barcelona, Spain, Apr. 2006.

[8] kc claffy, M. Crovella, T. Friedman, C. Shannon, and N. Spring, “Community-oriented network measurement infrastructure (conmi) workshop report,” *SIGCOMM Computer Communication Review*, vol. 36, no. 2, pp. 41–48, 2006.

[9] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore, “On the bias of traceroute sampling: or, power-law degree distributions in regular graphs,” in *STOC*. Baltimore, MD, USA: ACM, 2005, pp. 694–703.

[10] M. Latapy and C. Magnien, “Complex network measurements: Estimating the relevance of observed properties,” in *IEEE Infocom*, 2008.

[11] 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 Computer Communications Review*, vol. 36, p. 17.

[12] R. V. Oliveira, B. Zhang, and L. Zhang, “Observing the evolution of internet as topology,” *ACM SIGCOMM Computer Communications Review*, vol. 37, no. 4, pp. 313–324, 2007.

[13] M. Raab and A. Steger, “Balls into bins - a simple and tight analysis,” in *RANDOM '98: The Second International Workshop on Randomization and Approximation Techniques in Computer Science*, 1998, pp. 159–170.

[14] P. Barford, A. Bestavros, J. Byers, and M. Crovella, “On the marginal utility of network topology measurements,” in *IMW'01: the 1st ACM SIGCOMM Workshop on Internet Measurement*, 2001, pp. 5–17.

[15] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, “Avoiding traceroute anomalies with paris traceroute,” in *IMC'06*, 2006, pp. 153–158.

[16] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie, “Sampling biases in IP topology measurements,” in *INFOCOM*, Apr. 2003.

[17] Z. Mao, J. Rexford, J. Wang, and R. H. Katz, “Towards an accurate AS-level traceroute tool,” in *SIGCOMM*, 2003.

[18] Z. Mao, D. Johnson, J. Rexford, J. Wang, and R. H. Katz, “Scalable and accurate identification of AS-level forwarding paths,” in *INFOCOM*, 2004.

[19] L. Gao, “On inferring autonomous system relationships in the Internet,” *IEEE/ACM Trans. on Networking*, vol. 9, no. 6, pp. 733–745, 2001.

[20] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the Internet topology,” *ACM SIGCOMM Computer Communications Review*, vol. 29, no. 4, pp. 251–262, 1999.

[21] G. K. Zipf, *Human Behavior and the Principle of Least-Effort*. Cambridge, MA: Addison-Wesley, 1949.

[22] S. Dorogovtsev and J. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, Jan. 2003.

[23] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir, “A model of Internet topology using k -shell decomposition,” *Proceedings of the National Academy of Sciences USA (PNAS)*, vol. 104, no. 27, pp. 11 150–11 154, Jul. 3 2007.

[24] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, June 1998.

[25] L. Dall’asta, I. Alvarez-Hamelin, A. Barrat, A. Vazquez, and A. Vespignani, “A statistical approach to the traceroute-like exploration of networks: theory and simulations,” Jun. 2004.

[26] R. Cohen, M. Gonen, and A. Wool, “Bounding the bias of tree-like sampling in IP topologies,” Tech. Rep. cs.0611157, 2006.

[27] Y. Shavitt, E. Shir, and U. Weinsberg, “Near-deterministic inference of AS relationships,” Tech. Rep. cs.0711.4562, 2007.

[28] X. Dimitropoulos, D. Krioukov, G. Riley, and kc claffy, “Revealing the autonomous system taxonomy: The machine learning approach,” Tech. Rep. cs.0604015, 2006.