Where the Sidewalk Ends:
Extending the Internet AS Graph Using Traceroutes From P2P Users

Kai Chen*, David Choffnes†, Rahul Potharaju‡, Yan Chen§, Fabian Bustamante§, Dan Pei♯, Yao Zhao♮
*HKUST, †University of Washington, ‡Purdue, §Northwestern University, ♯AT&T Labs–Research,♮Bell Labs

***Abstract*—An accurate Internet topology graph is important in many areas of networking, from understanding ISP business relationships to diagnosing network anomalies. Most Internet mapping efforts have derived the network structure, at the level of interconnected autonomous systems (ASes), from a rather limited set of vantage points. In this paper, we argue that a promising approach to revealing the hidden areas of the Internet topology is through active measurement from an observation platform that scales with the growing Internet. By leveraging measurements performed by an extension to a popular P2P system, we show that this approach indeed exposes significant new topological information. Our study is based on traceroute measurements from more than 992,000 IPs in over 3,700 ASes distributed across the Internet hierarchy, many in regions of the Internet not covered by publicly available path information. To address this issue we develop heuristics that identify 23,914 new AS links not visible in the publicly-available BGP data – 12.86% more *customer-provider* links and 40.99% more *peering links*, than previously reported. We validate our heuristics using data from a tier-1 ISP, and show that they successfully filter out all false links introduced by public IP-to-AS mapping. We analyze properties of the Internet graph that includes these new links and characterize why they are missing. Finally, we have made the identified set of links and their inferred relationships publicly available.**

***Index Terms*—Internet measurement, AS topology, Traceroute**

I. Introduction

The Internet AS topology graph is important for many applications such as inferring ISP business relationships, designing new Internet routing protocols and diagnosing network anomalies [1]–[4]. As a result, many research efforts have investigated techniques for measuring and generating such graphs [5]–[10].

Most Internet mapping efforts have derived the network structure, at the AS level, from a limited number of data sources for either BGP paths or traceroute traces. The advantage of using BGP paths is that they can be gathered passively from BGP route collectors and thus require minimal measurement effort for obtaining a large number of Internet paths. Unfortunately, the publicly available BGP paths do not cover the entire Internet due to issues such as visibility constraints, route aggregation, hidden sub-optimal paths and policy filtering. In contrast, traceroute measurements provide the ability to infer the data paths that packets take when traversing the Internet. Because they are active measurements, traceroutes can be potentially issued from every corner of the Internet given sufficient numbers of vantage points (VPs),[1] facilitating the discovery of new network links. However, most existing traceroute-based projects are restricted by their currently limited number of VPs. Furthermore, the traceroute measurements provide an IP-level map that is too fine-grained for many applications. Converting an IP-level topology to an accurate AS-level one remains an open area of research [11].

In this paper, we argue that a promising approach to revealing the hidden areas of the Internet topology is through active measurement from an observation platform that scales with the growing Internet. Our work makes the following key contributions. First, we collect and analyze paths measured by traceroutes from hundreds of thousands of peer-to-peer (P2P) users worldwide (Section II). Specifically, the probes are issued from over 992,000 IPs in 3,700 ASes, making our measurement study the largest-ever in terms of the number of VPs and network coverage.

Second, we provide a thorough set of heuristics for inferring AS-level paths from traceroute data (Section III). To this end, we present a detailed analysis of issues that affect the accuracy of traceroute measurements and how our heuristics address these problems. Our proposed techniques for correcting IP-to-AS mapping are generic and work for the scenarios where traceroute VPs are poorly correlated with public BGP VPs. Furthermore, we validate our heuristics using data from a tier-1 ISP and show that they filter out all of the false links introduced by public IP-to-AS mappings for this ISP.

Third, we characterize the new links discovered by our P2P measurements (Section IV). We find that some common assumptions about the visibility of paths according to AS relationships are routinely violated. For example, while we have found 40.99% more *peering* links, we further observe that a VP can even miss some of its upstream *peering* links. More importantly, we reveal 12.86% more *customer-provider* links than what can be found in the publicly-available BGP data. With these new links, we analyze several properties of the Internet AS graph (Section IV-E).

Fourth, we derive a number of root causes behind the identified missing links, present a detailed analysis of their occurrences, and quantify the number of missing links due to each of those reasons (Section V). Interestingly, many of the missing links (75.02% in our dataset) are caused by multiple concurrent reasons.

We discuss limitations of this work in Section VI, review related research in Section VII and conclude in Section VIII.

II. P2P for Topology Monitoring

Understanding and characterizing the salient features of the ever-changing Internet topology requires a system of observa-

[1]Vantage points can be defined as locations with distinct network views. Because this paper focuses on AS topologies, we use *vantage point* to refer to a unique AS.

Project	# unique machines	# unique ASes
Routeviews/RIPE	790	438
Archipelago	41	≤ 41
iPlane	192	≤ 192
DIMES	8,059	200
Ono	**600,000**	**6,000**

TABLE I
APPROXIMATE NUMBERS OF VPS FOR TOPOLOGY-GATHERING PROJECTS AS OF DECEMBER 2009.

tion points that grows organically with the network. Because ISP interconnectivity is driven by business arrangements often protected by nondisclosure agreements, one must infer AS links from publicly available information such as BGP and traceroute measurements. The success of either approach ultimately depends on the number of measurement VPs.

To achieve broad coverage, we believe that it is essential to use a platform built upon large-scale emergent systems, such as P2P, that grow with the Internet itself. By piggybacking on an existing P2P system, one can eliminate the need to place BGP monitors in each ISP; rather, each participating host in the system can contribute to the AS topology measurement study simply by performing traceroute measurements.

As a first step toward this goal, we use data gathered from Ono [12], an extension to the Vuze BitTorrent client. The software has been installed more than 600,000 times by hosts assigned 992,000 IPs located in over 40,000 routable prefixes, spanning more than 6,000 ASes and 192 countries as of December 2009. Ono collects traceroute measurements between connected hosts to ensure that the software meets its goal of improving download performance while reducing cross-ISP traffic. Volunteers report this data to our central servers for offline analysis.[2] This platform constitutes the most diverse set of measurement VPs and is the largest set of traceroute measurements collected from end hosts to date. Table I compares the number of unique machines and VPs in our study and in a set of related efforts including Routeviews [13], RIPE/RIS [14], iPlane [15], DIMES [16] and Archipelago [17] as of December 2009. For Ono it is difficult to accurately determine the number of unique machines, so we use the number of times the software was installed.

As we show in Section IV, about 23,914 new links are discovered through these traceroute measurements. These new links include 26 ASNs (AS numbers) that do not appear in the publicly-available BGP data and thus are truly "dark networks" when viewed through the lens of the public BGP servers. Thus the view of the network from P2P users contributes a vast amount of information about network topology that is not seen by other existing approaches such as BGP table dumps and strategic active probing from dedicated infrastructure.

Figure 1 shows the distribution of VPs across hierarchical tiers (using the technique in [7]) for the publicly-available BGP data and the P2P traceroutes used in this study. Note that each bar represents the number of ASNs with VPs. The P2P traceroutes have significantly more VPs compared to the publicly-available BGP data, especially in lower-tier

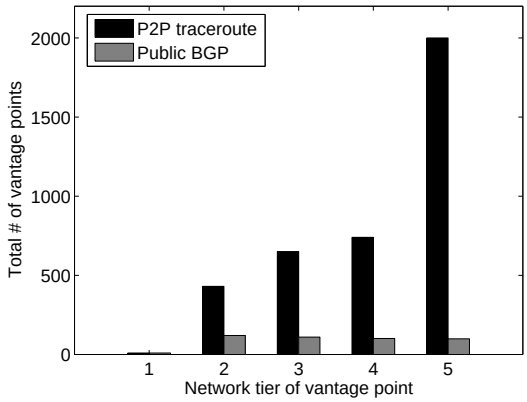

Figure 1. Distribution of VPs with respect to their network tiers.

networks. This unique perspective allows us to view previously hidden regions of the network and determine their impact on properties of the Internet topology.

While our dataset enables more complete coverage of the Internet topology than any previous study, it is important to note that we cannot make claims about the portion of the Internet not covered by this data (given the lack of ground truth) or about the impact of this missing data on key topological properties. Instead, an important goal of our work is to show the promise of using P2P for topology discovery, and we show that even a simple technique for path probing reveals significant amounts of new topology information.

III. METHODOLOGY

In this section, we first describe the datasets we use in this study. Second, we present a systematic approach to addressing the challenges associated with inferring AS-level paths from traceroute data. Third, we discuss how we validate our resulting topologies. Finally, we explain the algorithms used for inferring properties of the AS topology.

A. Data Collected

1) P2P traceroutes: The traceroutes in our dataset are collected by BitTorrent users recording the result of the `traceroute` command provided by their operating system. The measurements in this study were issued for the purpose of evaluating the effectiveness of reducing cross-ISP traffic in BitTorrent [12] and thus our destinations are restricted to BitTorrent peers connected to each of our measurement sources.

The data was collected from 992,197 distinct peer IPs[3] in 3,723 unique ASes. The destinations of these traceroutes are BitTorrent peers participating in the same swarms as our VPs. Specifically, the measurements are issued to destinations that established BitTorrent connections with a VP, generally in the order that the connections appeared. Together, the peers in our dataset probe more than 84 million distinct destination IPs.

Each peer performs traceroute measurements continuously while running our software, with at most one measurement

[2]Users are informed of the diagnostic information gathered by the plugin and are given the chance to opt out. In any case, no personally identifiable information is ever published.

[3]The number of unique installs in Table I and the number of distinct IPs are not equal because each user is often assigned dynamic IP addresses and some users disable traceroute probes.

active at a time; after each traceroute completes, the peer issues another to the next destination from the list of BitTorrent connections. Note that Ono biases connections towards nearby peers, so there is a slightly higher probability that traceroutes will be issued to them. While traceroutes between endpoints in the same AS do not reveal new AS-level topology information, biased connections between endpoints in nearby, but different, ASNs may cross private peering links or other portions of the topology unobservable from relatively distant vantage points (e.g., BGP monitors and PlanetLab), as discussed in Sec. V.

Measurements to the same destination are performed no more often than once every five minutes to ensure no destination is overwhelmed with probes. The measurements are performed using default settings except that the timeout for router responses is 3 seconds and no reverse DNS lookups are performed. Because the software performing the measurements is cross-platform, there are multiple traceroute implementations that generate data for our study. The vast majority of the data that we gather comes from the Windows traceroute implementation.

There are three measurements for each router hop; the ordered set of hops is sent to our central data-collection servers along with the time at which the measurement was performed. We use the data collected between Dec 1, 2007 and Sep 30, 2008, which consists of 541,023,742 measurements containing over 6.2 billion hops. We are making this data available to researchers upon request via the EdgeScope project[4].

2) BGP feeds: The BGP data used in this study includes a collection of BGP routing tables from 790 BGP speaking routers in 438 unique ASes. Specifically, we combine several BGP feeds: Routeviews [13] collected at route-views.oregon-ix.net, which is the most widely used BGP archive so far, six other Oregon route servers and 16 route collectors of RIPE/RIS [14]. We use 10 months of data gathered between Dec 1, 2007 and Sep 30, 2008, the same time period for our P2P traceroute data. We also download AS links from the UCLA IRL lab [18], which contains links collected from route servers, looking glasses and IRR [19]. Because UCLA data does not include BGP AS paths, nor information from new VPs added near the time of publication, we combine all of these sources of AS links to obtain the most complete set of AS links. For the rest of this paper, we will refer to this dataset as the "public view" [6, 7]. According to Oliveira et al. [6]–[8], ten months of public view data should be enough to cover all the *hidden links*, *i.e.*, policy-allowed links that do not always show up in the public view. For example, links only on sub-optimal paths do not appear in the public view unless the primary paths fail.

3) Ground-truth data: To validate our inferred AS links, we use router configurations and syslogs from a tier-1 ISP as ground-truth connectivity information. The data includes historical configuration and syslog files for more than 800 routers in this network. We simply leverage the heuristics in [6] to process these files and extract the ground-truth AS links that can be used as the baseline for our validation.

[4]http://www.aqualab.cs.northwestern.edu/projects/EdgeScope.html

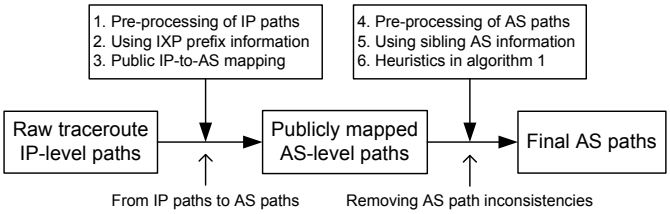

Figure 2. High-level architecture for converting IP paths to AS paths.

B. Inferring AS Links Using Traceroutes

While traceroute probes can provide detailed network topology information, there are a number of issues that prevent their widespread use in AS topology generation. First, the number of probe sources and targets required to reveal new topological information grows with the size of the Internet. As we discussed in Section II, we address this issue through measurements from P2P users. Second, traceroutes provide IP-level views of the topology and the IP-to-AS mappings gathered from publicly available information are incomplete and potentially incorrect. Finally, traceroute measurements are subject to the constraints of the routers they visit, which can drop probes, silently forward them without altering the TTL or even erroneously modify the TTL in ways that affect the inferred path. When using traceroutes as a telescope for viewing the AS topology, one must expect a blurry lens with many artifacts. In this section, we discuss a systematic approach for sharpening and clarifying this view by addressing these limitations.

Figure 2 illustrates the steps we take to convert traceroute IP paths into their corresponding AS paths. In the following, we first address the issues with traceroute IP-level paths (**steps 1–2**). Then, we obtain AS-level paths based on public IP-to-AS mappings (**step 3**). Finally, we address the issues with the directly mapped AS-level paths (**steps 4–6**).

Step 1: Pre-processing IP paths. Before performing IP-to-AS mapping, we inspect each IP-level path. First, we search for those measurements that contain repeated, consecutive IP addresses in the path. When this occurs, the repeated IP is likely to be upstream from a router that is not decrementing the traceroute probe's TTL. There are other problems such as load balancing, zero-TTL forwarding and address rewriting of gateway routers that would cause routing loops [20]. All these cases could lead to falsely inferred AS links. To avoid any potential problems related to these issues, we conservatively remove the measurement from our analysis.

Step 2: Removing IPs within IXPs. Paths that traverse Internet eXchange Points (IXPs) can lead to falsely inferred AS links. This is because IXPs usually create an extra AS hop along the AS path [11]. Using a list of known IXP IP prefixes, such as PCH [21], PeeringDB [22] and Euro-IX [23], we remove from each IP path any hop that belongs to an IXP prefix. This allows us to correctly infer direct links between the ASes that connect to each other at an IXP. However, we cannot rely on the publicly available information to completely eliminate such false links because they are known to be incomplete. Our heuristics in the next subsection will address this issue for AS link inference.

Step 3: Public IP-to-AS mapping. After the first two steps,

	Problem	Symptom				Filtering heuristic(s)
		Loop	Missing hop	Substitute hop	Extra hop	
Incomplete paths	Unresolved hops within an AS	Problem addressed in [11]				Step 4
	Unmapped hops between ASes					Step 4
	MOAS hops at the end					Step 4
False AS links	Internet exchange points(IXPs)				✓	Steps 2, 4, 6
	Sibling ASes	✓	✓	✓	✓	Steps 5, 6
	Unannounced IP addresses	✓	✓	✓	✓	Step 6
	Using outgoing interface IPs		✓	✓	✓	Step 6
	Private peering interface IPs		✓			Step 6

TABLE II
PROBLEMS WITHIN TRACEROUTE-INFERRED AS-LEVEL PATHS, SYMPTOMS FOR THESE PROBLEMS, AND THE STEPS WE TAKE TO SOLVE THEM.

we convert IP-level paths to AS-level ones by directly using the IP-to-AS mapping provided by Team Cymru [24], which incorporates both publicly available and private BGP information. For private addresses, we use a placeholder instead of an ASN for the hop. As we discuss in the following paragraphs, the placeholder is filled with a valid link if it matches a pattern in the public view; otherwise, no links on either end of the placeholder are inferred.

We next address the issues with converting IP-level paths into AS-level ones. While previous work has investigated the problem of accurate traceroute AS paths where BGP paths at the same VP are available [11], our study is the first to address the problem for an arbitrary (and large) set of traceroute paths. Note that our proposed techniques are general in that they consider the scenario where traceroute VPs are not the same as BGP VPs. The key challenge we address here is to distinguish falsely inferred AS links caused by incorrect IP-to-AS mappings. To evaluate the quality of our heuristics, we compare our results with ground-truth data from a tier-1 ISP.

Recall that the above step directly converts IP paths to AS ones using public IP-to-AS mapping. However, Mao et al. [11] identify several patterns of discrepancies between traceroute and BGP paths (as in Table II) using such mapping, each of which entails a difference of at most one AS hop, *e.g.*, an AS is missing on the path, and an extra or substitute AS appears on the path. To account for these discrepancies while still preserving true new AS links discovered by traceroute measurements, we mark any new link as *pending* if it can be corrected by techniques used in [11]; otherwise, we assume that the new link is true. In our approach, we conservatively modify all the pending links such that they are consistent with the corresponding BGP paths. We show in Table II that our method for converting IP paths to AS paths can address all the problems identified by [11]. We emphasize that this method reduces false positives (and in our ground truth data, eliminates them), but may filter out real AS links not present in available BGP paths.

Step 4: Pre-processing AS paths. As shown in Table II, to address the incomplete AS path problem, we directly apply the techniques developped in [11] to our dataset with the following modifications. To avoid inferring false links when multiple origin AS (MOAS) hops appear at the end of a path, we conservatively drop the last hop in our analysis. Further, we remove IXPs from the traceroute-generated AS paths based on published IXP prefixes. In addition, we explore the fact that IP addresses within an IXP prefix are mapped to multiple ASes when the shared infrastructure address is announced into BGP by multiple participant ASes. To capture these cases, we identify the hops along the traceroute paths that are publicly mapped to multiple ASes, and check if these ASes are collocated in an IXP. If so, we remove these hops from the path. However, we cannot use this approach to identify IXPs that use their own AS numbers, a limitation we address in Step 6.

Step 5: Addressing issues with sibling ASes. A single organization may own and manage multiple sibling ASes. One AS may use some address blocks from its sibling to number its equipment or during route propagation only one of two sibling ASes includes its ASN in the BGP AS path. To address this issue, we download the known sibling ASes from CAIDA [25]. For a sibling AS pair (X, Y), we may see the cases where traceroute AS path is $[...WXYZ...]$ while a corresponding BGP AS path is $[...WXZ...]$ or $[...WYZ...]$. For this case, we modify the traceroute AS path to be $[...W\{X, Y\}Z...]$; In our dataset, we also find instances where the traceroute AS path is $[...WYZ...]$ while a corresponding BGP AS path is $[...WXZ...]$. In those cases we use the BGP AS path to modify the traceroute AS paths. Again, publicly available sibling AS information is limited. In the next step, we use heuristics to mitigate the remaining problems when sibling ASes cause discrepancies between traceroute AS paths and BGP AS paths.

Step 6: Addressing remaining issues. Algorithm 1 addresses the inconsistencies between traceroute AS paths and BGP AS paths that remain after the previous 5 steps. Below, we discuss how the algorithm deals with each of the symptoms, *i.e.*, loops and missing/extra/substitute hops, as shown in Table II. In this section, we refer to the *DIST* for each traceroute-based AS link, defined to be the number of hops between these two ASes with respect to the public view.

- **Loop.** Loops in traceroute AS paths can happen due to unannounced IP addresses, sibling ASes, or route anomalies on the forwarding paths. In our dataset, loops in the traceroute AS paths are rare. While not all AS-level loops are invalid, we conservatively discard these paths.
- **Missing hop.** We use the public view connectivity graph to calculate the *DIST* of each link on each traceroute AS path. If the *DIST* is 2 and in the BGP AS paths the corresponding route is $[...BXC...]$ (case 1 in Figure 3), we conservatively add one hop X in the middle to make traceroute AS paths consistent with BGP AS paths (as in line 8 of Algorithm 1). This mismatch could result from the following:
  1) ***Private peering interface IPs.*** AS links B-X and X-C are both private peerings using IP addresses from

PROCEDURE Addressing issues within traceroute AS paths
1 Initialization: set the *DIST* of each traceroute-based AS link;
2 **foreach** *AS link in the traceroute AS paths (e.g., use B-C at the top of Figure 3 as illustration)* **do**
3 **if** *DIST(B,C)=1* **then**
4 AS link B-C is considered *true*;
5 **if** *DIST(B,C)=2* **then**
6 Check the public view BGP AS paths;
7 **if** *There exists an AS path ...B X C...* **then**
8 Fix B-C using B-X-C and set *DIST* of each link as 1; /*For multiple Xs, choose the one on the longest matched path, *e.g.*, suppose both [A B X_1 C D] and [A' B X_2 C D'] exist, the first path matches [A B C D] better, so X_1 is preferred than X_2*/;
9 **if** *There does not exist an AS path ...B X C...* **then**
10 **if** *...A X C... (or ...B X D...) appears in BGP AS paths* **then**
11 Replace B (or C) with X and set the DIST of each link as 1 (longest match for multiple Xs);
  else
12 **if** *DIST(A,C)=1 (or DIST(B,D)=1)* **then**
13 Delete B (or C) and set the DIST of link A-C (or B-D) as 1 ;
14 **if** *DIST(A,C)≠1 and DIST(B,D)≠1* **then**
15 Mark B-C as a real link and set DIST(B,C) as 1 ;
  end
16 **if** *DIST(B,C)≥3* **then**
17 **if** *DIST(A,C)=1 (or DIST(B,D)=1)* **then**
18 Delete B (or C) and set the DIST of link A-C (or B-D) as 1 ;
19 **if** *DIST(A,C)≠1 and DIST(B,D)≠1* **then**
20 Mark B-C as a real link and set DIST(B,C) as 1;
end
21 Return the traceroute AS path if *DIST* of each link on it is 1;

Algorithm 1: Heuristics in **Step 6** of Figure 2.

B and C respectively. When traceroute probes travel from B to X and then immediately exit X to enter C, the resulting traceroute-based AS path would be [...BC...] while [...BXC...] is the (true) BGP AS path;

2) ***Sibling ASes.*** AS X is a sibling of B (or C) and uses its sibling's address blocks for equipment numbering. As we discussed in step 5, this would cause a traceroute AS path to miss one hop;
3) ***Unannounced IP addresses.*** AS X is a customer of B (or C), and uses the IP addresses from B (or C) but does not announce them publicly. In this case, X responds to traceroute probes with IPs that are falsely mapped to B, which causes [...BC...] to incorrectly appear in the traceroute AS path while [...BXC...] is the correct BGP AS path;
4) ***Using outgoing interface IPs.*** A border router in AS X uses its outgoing interface for ICMP, so the hop is not mapped to X.

Note that it is possible for traceroute-based link B-C in the above cases to be real – but not observed by the public BGP monitors. Because we conservatively filter out such links, we may introduce false negatives in our results.

- **Substitute hop and extra hop.** If the *DIST* is 2 and the intermediate node connecting B and C is X, but we could not find any corresponding route [...BXC...] in the BGP AS paths, it may either be due to insufficient coverage of BGP AS paths from publicly available VPs or because AS path [...BXC...] is invalid. The substitute/extra hop problem could result from the following scenarios:

1) ***Unannounced IP addresses.*** Consider an AS X that is multihomed to its providers B and C and uses IP addresses from one of them (B or C) to set up its equipment but does not announce them publicly (case 2a/2b in Figure 3). This would produce a traceroute AS path of [...ABC...] (or [...BCD...]) while the corresponding BGP AS path is [...AXC...] (or [...BXD...]) – a substitute hop.
   Another issue can arise if an AS A not only uses unannounced addresses from its provider but also owns and announces some other addresses. When traceroutes traverse this AS, these different addresses can generate a false inter-AS link based on public IP-to-AS mappings. For example, in case 2c/2d of Figure 3, while traceroute AS path is [...ABC...] (or [...BCD...]) its BGP path is [...AC...] (or [...BD...]) – an extra hop;
2) ***IXPs or sibling ASes.*** As explained in previous subsections, IXPs can lead to extra hops and sibling ASes can lead to substitute/extra hops.
3) ***Using outgoing interface IPs.*** In case 3 of Figure 3, for example, AS A's last-hop router uses its outgoing interface (facing C) to reply to an ICMP message (the connection between A and B uses addresses from B). This causes one extra or substitute hop in traceroute AS path: [...ABC...] appears in the traceroute AS path and [...AC...] appears in the BGP AS path. Further, if the traceroute traverses only one hop in A, then it would cause A to be falsely substituted with B.

For these scenarios, if we can find the corresponding routes in BGP, we make the traceroute AS paths consistent with BGP AS paths by replacing the middle hop with X or deleting it (line 11 $\sim$ line 13 in Algorithm 1). Similar to the missing hop cases, our conservative approach could discard true links. For instance, we may omit true sibling AS links.

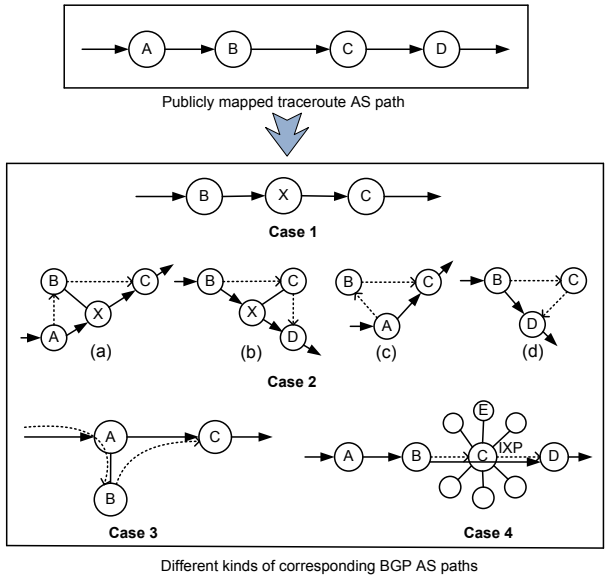

Figure 3. Relationship between traceroute AS paths and BGP AS paths for several cases. Dotted arrows are traceroute AS paths and solid arrows are the corresponding BGP AS paths.

- **Other special cases.** Though rare, we found cases where traceroute AS links have a *DIST* ≥ 3. We provide one plausible scenario in case 4 of Figure 3. Here, C is an IXP with its own ASN that is announced only via a particular participant, say E. If E is not a neighbor of B (*i.e.*, ≥ 2 hops), this would cause B and C to be at least 3 hops away in BGP. Our algorithm addresses the special case in lines 17 and 18. Otherwise, we assume the link to be *true* if it could not be explained by this. While it is possible for other unaccounted scenarios to exist, we believe the impact of these scenarios is sufficiently limited by the scarcity of the examples in our dataset.

C. Validation

After applying all the heuristics in the previous section, we are left with 100,000 AS links discovered through P2P traceroutes. We now validate a significant portion of these links with the ground-truth information from a tier-1 AS (the number of AS links is on the order of thousands[5]). Most importantly, we find that *all* the P2P-based links are in the ground-truth information.

Using the tier-1 network (denote T_1), we calculate the percent of false links filtered out by each of our heuristics, focusing on those in Algorithm 1. After applying **Steps 1–5** (and before applying these heuristics), our P2P traceroutes identified thousands of links with this tier-1 AS. Compared with the ground-truth connectivity, 48.8% of these traceroute-based AS links were false. We now discuss how each aspect of Algorithm 1 reduces the percent of false links; the list of values is presented in Table III.

DIST(B, C)**=2 and [...BXC...] exists in BGP (line 8):** We see several hundred cases where $[...T_1, C...]$ is in our traceroute AS paths while $[...T_1, X, C...]$ is in BGP AS paths. Checking with the router configuration files of the tier-1 network, we found that, in 94% of the cases, the last IP hop that publicly mapped to T_1 actually belongs to a third AS X. These false links may happen due to private peering or unannounced IP addresses. This lends strong evidence that line 8 of the algorithm, which adds an extra hop to a traceroute-based AS path, is valid. We further note that we did not find a single T_1-C link to be valid according to the ground-truth. After this step, slightly more than 10% of the links are false.

DIST(B, C)**=2 but [...BXC...] does not exist in BGP (lines 11 and 13):** Our traceroute dataset contains hundreds of cases where $[...A, T_1, C...]$ (or $[...B, T_1, D...]$) appears for this tier-1 AS. To validate this, we first used IP-level paths and extracted those IPs that were mapped to T_1. Then we searched for these IPs in the router configuration files to see if they are indeed used to configure real routers of the tier-1 network. In 93% of the cases, we found that these IPs are not used by this tier-1 network. This indicates that the IPs are probably allocated to the AS's customers (or siblings), say X. Given the data available to us, we have no way to determine which AS this X is. However, this result indicates that our heuristics accurately identify the corresponding cases for incorrect mappings, allowing us to filter out (or correct) the false links. After accounting for these issues, only 0.47% of the links are false.

[5]Because this information is proprietary, we cannot disclose the precise number of AS links so we use percentages in this section.

Line # in Algorithm 1	False links left
-	48.80%
8 (address missing hop)	10.47%
11 (address substitute hop)	5.13%
13 (address extra hop)	0.47%
18 (address special case)	0

TABLE III
PERCENT OF FALSE LINKS REMAINING AFTER EACH FILTERING STEP.

DIST$(B, C) \geq$ **3 (line 18):** We have no specific ground-truth data that can help us validate our heuristic here. However, the tier-1 network connectivity information allows us to estimate whether this line removes any false links. In this study, we found only 0.47% of the links to the tier-1 AS had *DIST* ≥ 3. After applying the rule (lines 17 to 18), all of these 0.47% false links are properly removed.

Finally, we note that the goal of this work is to increase the accuracy of AS path inference from P2P traceroutes so that we can extend the AS topology, but we do not claim that P2P traceroutes alone can cover the entire AS topology. For instance, we miss at least 21.3% of the tier-1 AS's links according to the ground-truth. As such, our P2P-based dataset does not introduce any false links in this tier-1 AS, nor does it discover all the links in the AS.

D. Policy Inference

After extracting the AS links, we infer the business relationships between ASes based on the PTE algorithm proposed by Xia [26]. After improving the seminal work by Gao [27], the PTE approach is believed to outperform most other approaches [10]. Accordingly, most AS links are classified as one of three kinds of relationships: *customer-provider* links, *peering* links, and *sibling* links. We note that these are simplistic assumptions of real business relationships. And we also note that the PTE inference algorithm can incorrectly infer relationships, and this would potentially influence the accuracy of link classification and root cause analysis.

We use our topology to classify ASes into hierarchical tiers. There are many techniques for hierarchical classification, including use of the degrees of individual ASes, the number of prefixes originated by the ASes and the number of distinct AS paths seen from a particular AS. However, without accounting for the ASes' contractual relationships, these heuristics may be misleading. Thus, we use the technique proposed by Oliveira et al. [6, 7], which relies on the number of downstream customer ASes to classify each AS. In this paper, we classify ASes into 5 tiers, with tier-1 as the Internet core and tier-5 as the edge.

IV. THE MISSING LINKS

After generating a more complete AS topology from P2P traceroutes, we found a significant number of new AS links (including *customer-provider*, *peering* and *sibling*), as shown in Table IV. In this section, we use our set of newly identified links to determine the public view's coverage of each class of AS links and where these links are missed by public view. We

General AS links			Customer-provider links			Peering links			Sibling links		
PV #	New #	Fraction %	PV #	New #	Fraction %	PV #	New #	Fraction %	PV #	New #	Fraction %
119470	23914	20.02%	83783	10775	12.86%	31054	12729	40.99%	4545	216	5.75%

TABLE IV
STATISTICS FOR NEWLY IDENTIFIED LINKS (*PV* STANDS FOR PUBLIC VIEW; *New#* IS THE NUMBER OF LINKS NOT IN PV).

Tier-1 network	In PV	New in P2P	Percentage
AT&T (AS7018)	2668	0	-
Sprint (AS1239)	2293	0	-
Level3 (AS3356)	2774	53	1.91%
Qwest (AS209)	1656	34	2.05%
Verio (AS2914)	1116	35	3.14%
UUNET (AS701)	3692	17	0.46%
SAVVIS (AS3561)	713	0	-
Cogent (AS174)	2451	44	1.80%
GBLX (AS3549)	1721	49	2.85%

TABLE V
NUMBER OF AS LINKS FOR TIER-1 NETWORKS IN THE PUBLIC VIEW (2ND COLUMN), NUMBER OF NEW LINKS FROM P2P TRACEROUTES (3RD COLUMN), AND THE CORRESPONDING PERCENTAGE (4TH COLUMN).

then evaluate several topological properties of the AS graph with these newly identified links.

A. Coverage of tier-1 AS links

We begin by focusing on the tier-1 AS connectivity, listed in Table V. Note that although we have uncovered 23,914 new links, we discovered few new tier-1 AS links: 1) we did not find any new links for three of the tier-1 ASes, and 2) we found a small percentage (up to 3.14%) of new links for the remaining tier-1 networks. This result is consistent with previous work [6] indicating that tier-1 AS links are covered fairly completely by the public view over time. On the other hand, our results also indicate that the public view still misses some tier-1 links, even though there are monitors in these networks. We offer the following possible explanations. First, a tier-1 AS could contain thousands of routers, each potentially with a constrained view of the AS. In this case, the relatively small number of feeds (*i.e.*, peered routers) of the current public view for each AS may capture an incomplete view of the AS. In addition, some tier-1 ISPs do not announce all of their prefixes (*e.g.*, those longer than /24), which prevents the public view from seeing the corresponding links.

B. Coverage of customer-provider links

We now turn our attention to the set of *customer-provider* links discovered by P2P traceroutes. Table IV shows that P2P traceroutes discover 12.86% additional *customer-provider* links missing from the public view. To put this in context, recent work [8] investigating the AS graph based on BGP data suggests that a time window of ten months captures all non-optimal paths and that the public view does not miss *customer-provider* links in general if valley-free policy is strictly followed – thus each link should be on some paths of at least one prefix. Our results indicate that, due to factors such as route aggregation (explained later in Section V-B2), the assumption is often violated – thus this public view is not as complete as previously suggested.

We categorize the missing links according to their relationships: the fraction of missing provider links and the fraction of missing customer links. We use the method from Section III-D to classify each AS into a tier, then group all of the fractions for each tier. Figures 4 and 5 show the CDF of the fractions of missing links, where the fraction for one AS is calculated as the number of missing provider (or customer) links divided by the total number of provider (or customer) links. Note that tier-1 ASes have no providers and tier-5 ASes have no customers. The figures clearly show that *customer-provider* links can be missed in every tier. More importantly, we observe that the fraction of missing provider links of an AS somewhat correlates to its tier in the Internet hierarchy: the higher the tier number of an AS, the more likely that the public view will miss its provider links.

C. Coverage of peering links

Previous work has shown that the public view misses a large number of *peering* links, especially in the lower tiers of the Internet routing hierarchy [6, 10]. Our study finds that P2P traceroutes reveal an additional 40.99% *peering* links, which confirms these prior results. Such missing *peering* links are expected to appear at lower tiers of the Internet hierarchy, where there is less coverage from BGP feeds. However, we find that a significant number of *peering* links are missing from the public view at higher tiers. Similarly, we calculate the fraction for missing *peering* links and plot the CDF in Figure 6. The graph shows that high tier networks have relatively higher fractions of missing links than low tier networks except that tier-1 ASes do not miss *peering* links. We will investigate the reasons behind these missing *peering* links in Section V.

D. New sibling AS links

We revealed 216 additional *sibling* links which are missing from the public view. We believe that one reason behind these new *sibling* links could be route announcement in BGP – *e.g.*, the paths announced by a pair of sibling ASes, AS1 and AS2, contains the ASN for only one of the siblings. Such announcements hide the *sibling* link from the public view; however, traceroute probes traversing this link can reveal IPs for both sibling ASes and thus the *sibling* AS link.

E. AS Graph Properties with New Links

Most previous work on the Internet AS graph properties relies on BGP information; we now use our P2P traceroute measurements to generate a more complete AS graph and evaluate its topological properties. We focus only on several popular aspects of the AS graph properties; an exhaustive analysis is beyond the scope of this paper.

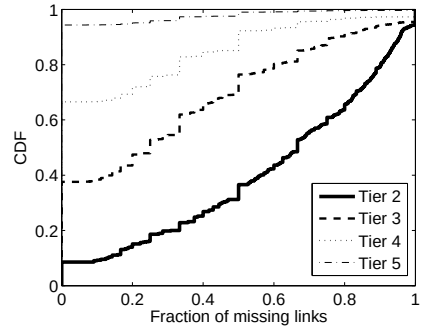

Figure 4. Revealed provider links.

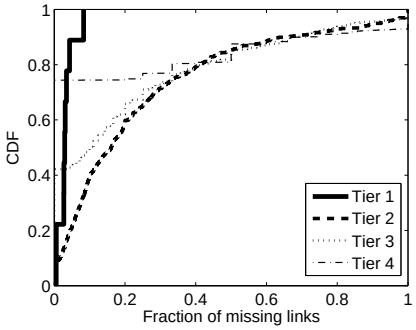

Figure 5. Revealed customer links.

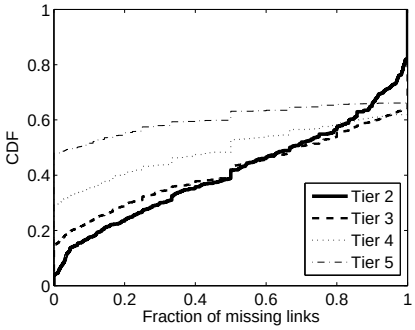

Figure 6. Revealed peering links.

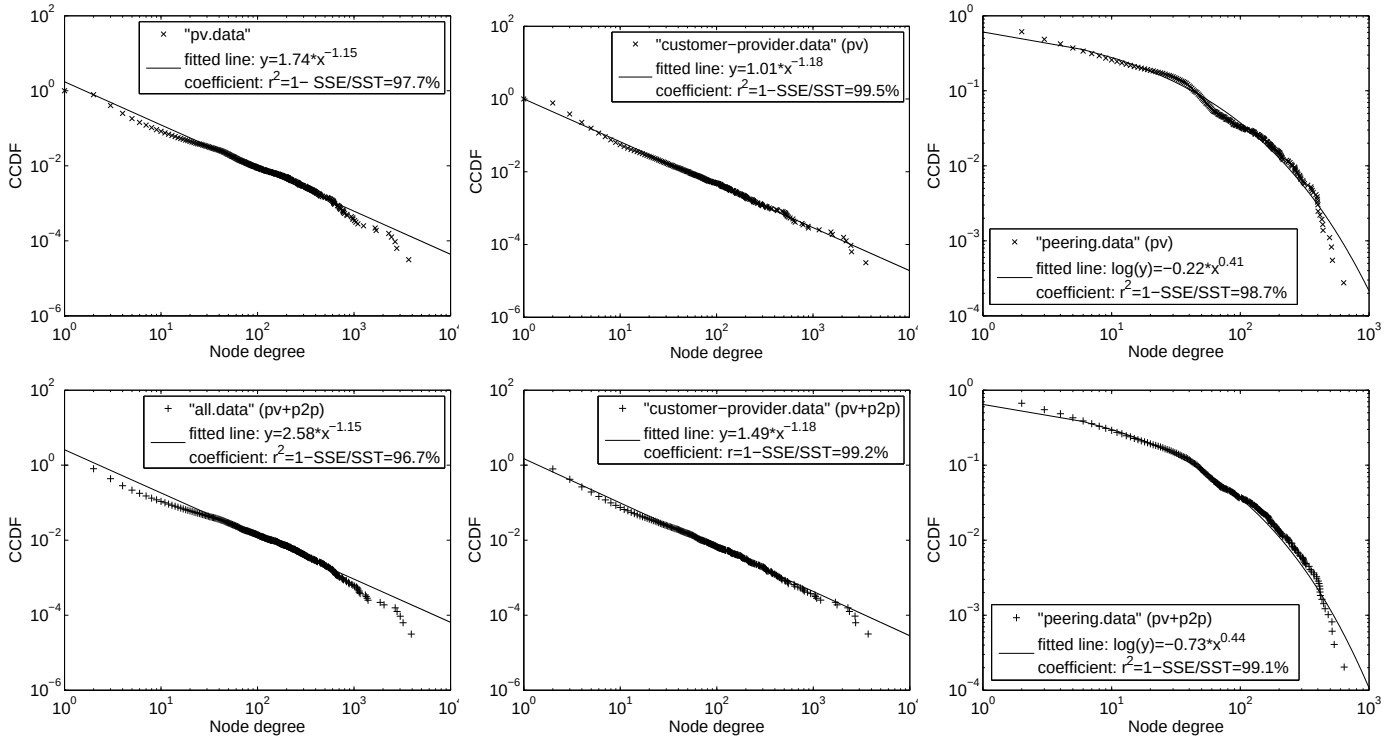

Figure 7. Degree distribution of the entire AS graph (left two), *customer-provider* subgraph (middle two), and *peering* subgraph (right two). Each pair is derived from the public view (pv) and from the public view plus p2p traceroute (pv+p2p). SSE: sum of squares due to error; SST: total sum of squares.

1) Degree Distribution: Since the seminal power-law distribution of Internet topology study [28], there has been debate [5, 9, 10, 29] as to whether this observation is valid when more links are identified. With the most complete AS graph to date, we now re-examine the distributions. We model the degree distributions for different types of links in this graph. We first check the distribution for the entire AS graph, then the *customer-provider* subgraph and at last the *peering* subgraph. The *customer-provider* subgraph is composed of all *customer-provider* links, and the *peering* subgraph is composed of all *peering* links. We analyze the *customer-provider* subgraph to evaluate the "rich-become-richer" phenomenon, which states that a new node prefers to connect to richly connected nodes to get better service. Based on this assumption, a power-law distribution should hold for the *customer-provider* subgraph. On the other hand, the power-law may not hold for the *peering* subgraph because there is less incentive for a new node to connect to a node having more *peers*.

- Figure 7 (left two) validates that the power-law property can still be used to describe the more complete AS graph. With around 20.02% newly identified links, the degree distribution still follows a power-law function with a coefficient of determination of 96.7%. Note that there is a slight decrease (*i.e.*, 1%) compared with the coefficient of the public view. This is because we have identified 40.99% more *peering* links which impact the original power-law distribution.

- The *customer-provider* degree distribution, depicted in the middle graphs of Figure 7, is precisely modeled by a power-law curve. Such properties are essentially unchanged with respect to the public view alone although we have additional 12.86% new links included. This strongly confirms the "rich-become-richer" rule behind the *customer-provider* subgraph.

- Figure 7 (right graphs) show that pure *peering* subgraphs are explained by the Weibull [30] distribution, which is not heavy-tailed. Our newly identified *peering* links improves the fit to this model.

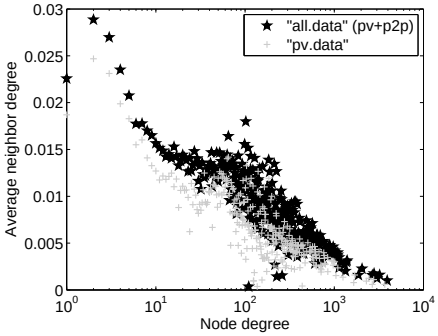

Figure 8. Average neighbor connectivity(normalized).

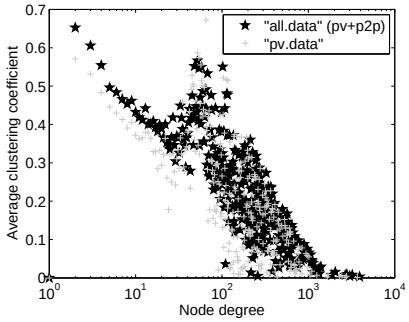

Figure 9. Average clustering coefficient.

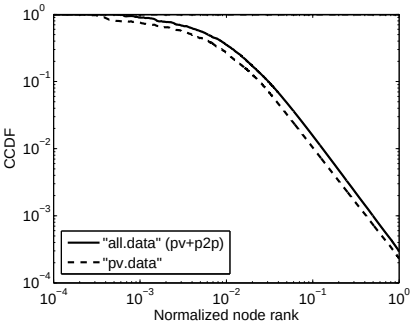

Figure 10. Rich-club connectivity.

Note that our results corroborate observations in previous work [10] that are based on a one-day snapshot and the IRR data. Nonetheless, we have demonstrated that these distributions persist even with a long-term view of Internet topology from the public view and P2P traceroutes.

2) Average Neighbor Connectivity: The node degree distribution tells how many nodes of a given degree are in the network, but it fails to provide information on the interconnection between these nodes. Instead, the average neighbor connectivity is used to show average neighbor degree of a k-degree node. This metric reflects whether ASes of a given degree preferentially connect to high or low-degree ASes. Figure 8 shows the average neighbor connectivity in terms of node degree in two topologies. The number is normalized by its maximal value n (n is achieved when the graph is full mesh). The figure indicates that the AS graph derived from the public view underestimates the average neighbor connectivity, especially for low and middle-degree nodes.

3) Clustering: The clustering coefficient describes the local connections among a node with its neighbors, and thus demonstrates the local robustness in the topology. Let d_v be the number of neighbors for node v and m_v be the number of links between these d_v neighbors. Average clustering coefficient C_{ave}, is defined as $\frac{1}{|V|}\sum_{v\in V} C_v$ where $C_v = m_v/\binom{d_v}{2}$ is the local clustering coefficient for v. We observe that the C_{ave} of the Internet AS graph increased from 0.40 to 0.46 after we add the newly identified links, meaning that the AS graph from the public view underestimates the local path diversity.

In addition, we find that the clustering density did not increase uniformly. In Figure 9, we show the average clustering coefficient in terms of node degree between two topologies. The graph shows that the neighborhoods of low-degree ASes have become more clustered, likely because the public view fails to see many links near the edge of the network. While most of the coefficients for middle-degree ASes increase, we find that the coefficients for some of them decrease. This occurs because these ASes are connecting with nodes that have lower connectivity compared to those in the public view.

4) Rich-club Connectivity: Rich-club connectivity (RCC) shows the extent to which nodes with high degree are connected to each other. Specifically, let $\sigma = 1...n$ be the first σ nodes with a non-increasing order in a graph with n nodes. RCC is the ratio of the number of links in a subgraph induced by the first σ highest-degree nodes with the maximum possible number links in this subgraph, which is $\binom{\sigma}{2}$. This essentially measures how tightly this σ-induced subgraph is connected. Figure 10 shows the CCDFs of RCC with normalized node rank (σ/n) of two AS graphs. We can observe a gap between these two curves, indicating that the public view underestimates this property. Another observation is that both curves follow straight lines when $\sigma/n \geq 0.01$, which shows a power-law distribution of RCC for both graphs.

V. In Search of Root Causes

The above section characterized links found via P2P traceroutes that were absent from the public view. By determining why these links are missing, we can better understand how to extend our results to build models for generating AS graphs.

An analysis of root causes for missing links is particularly difficult because we lack the ground-truth information required to validate our conclusions. This is a limitation of any work on an Internet-wide AS topology. In our analysis, we observe that a single missing AS link may have one or more possible root causes. Thus, we determine a set of root causes that could be responsible for a missing link.

A. Exploring Missing Patterns

To identify the cause(s) for a missing link, we first determine where it occurs with respect to the VPs of the public view. Specifically, we use BGP AS paths to identify routes from VPs to missing links and classify them according to route patterns depicted in Figure 11. For simplicity, and without loss of generality, we condense a continuous series of *customer-to-provider* (or *provider-to-customer*) links into one logical *customer-to-provider* (or *provider-to-customer*) link. Note that in some rare cases, the public view does not contain information about either AS in a link found through P2P traceroutes; we omit these links in the following analysis.

1) Observations: Table VI presents both visible and missing links for each pattern. Note that the sum of *peering*, *customer-to-provider*, and *provider-to-customer* links can be different from the sum of links in each pattern in Table VI because we omit *sibling* links and links for which the relationship cannot be inferred. Also, one link could appear in a pattern both as a *customer-to-provider* link and as a *provider-to-customer* link. After classifying missing links in this way, we make the following key observations:

Patterns	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)
# of unique links observed	75817	78746	54869	55731	40518	54262	40666	52331
# of *peering*	19474	16492	N/A	N/A	N/A	N/A	N/A	N/A
# of *customer-to-provider*	5036	4550	N/A	N/A	N/A	N/A	N/A	N/A
# of *provider-to-customer*	49194	55948	52092	53830	39024	51681	39290	50604
# of unique links missed	5185	22535	23094	23909	23889	22676	23691	23884
# of *peering*	3330	12395	12576	12726	12706	12473	12579	12709
# of *customer-to-provider*	1521	7220	7973	10563	10410	7484	9722	10274
# of *provider-to-customer*	1343	6852	7692	10444	10583	7077	9914	10469
Percentage of missing links	6.83%	28.62%	42.09%	42.90%	58.96%	41.79%	58.26%	45.64%

TABLE VI

NUMBERS OF MISSING/VISIBLE LINKS IN EACH PATTERN OF FIGURE 11. READING COLUMN 2, 75,817 VISIBLE LINKS FIT PATTERN (A) WHILE 5,185 MISSING LINKS FIT PATTERN (A). "N/A" MEANS NO LINK HAS BEEN OBSERVED VIA THE CORRESPONDING PATTERNS (DUE TO VALLEY-FREE POLICY).

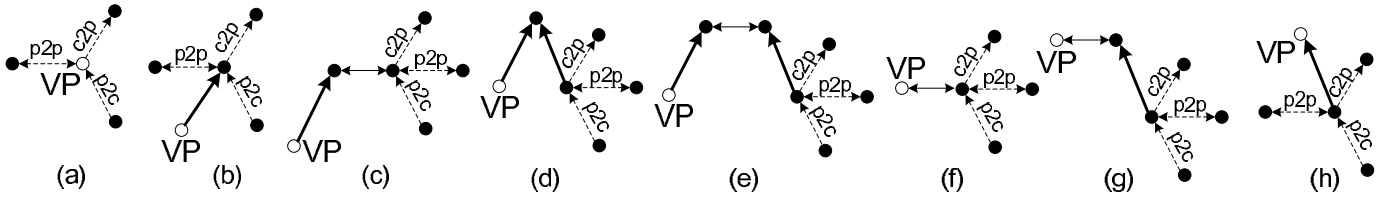

Figure 11. Eight patterns for the locations of missing links relative to the VPs. A bold arrow represents a *customer-to-provider* link or a combination of *customer-to-provider* links; a bidirectional (thin) arrow represents only one *peering* link; a dotted arrow represents an identified missing link. Reading the figure, pattern (b) means there are missing c2p, p2p, and p2c links when starting at a VP and traversing one (or multiple) *customer-to-provider* links.

- It is reported that a monitor with full BGP table[6] can discover all the connections of its upstream providers [6, 7]. However, we found that a full-table VP may not cover all of the links belonging to its AS, nor all those belonging to the AS's upstream providers (such as pattern (a) and (b)). In our measurements, we found the first 100 full table VPs missed 1096 links adjacent to the VP's AS.
- While *peering* links are expected to be missing from the public view, we note that we found a significant number of missing *customer-provider* links.
- It is known that many *peering* links are missed in the low-tiers of Internet hierarchy [6, 10], and our result for pattern (h) in Table VI confirms this fact. However, we also find many instances of upstream peering links being invisible to downstream full table monitors (for example, pattern (b)). This means that ASes located low in the hierarchy are not solely responsible for missing *peering* links.

2) Completeness: We find that there are at most 8 patterns as shown in Figure 11 for a link to be missing.

Proof: ASes follow certain guidelines in their export policy settings [27]. Put in simple words, a *provider-to-customer* ($p2c$) or a *peering* ($p2p$) edge can only be followed by *provider-to-customer* or *sibling* ($s2s$) edges. This can be used to define a valid AS path p as:

$$valid\_path(p) = x(c2p|s2s) + y(p2p) + z(p2c|s2s) \quad (1)$$

where $x, z = \{0, 1, 2, 3...\}$ and $y = \{0, 1\}$. When we do not consider the $s2s$ links and abstract a continuous series of $c2p$ and $p2c$ links into one logical $C2P$ and $P2C$ link respectively, Eq. 1 can be reduced to:

$$valid\_path(p) = x'(C2P) + y(p2p) + z'(P2C) \quad (2)$$

where $x', y, z' = \{0, 1\}$. Thus, due to the presence of the three binary variables, we can get at most $2^3 = 8$ patterns. ■

[6]In our experiment, a VP with a full BGP table contains routing table entries that cover nearly the entire Internet prefix space.

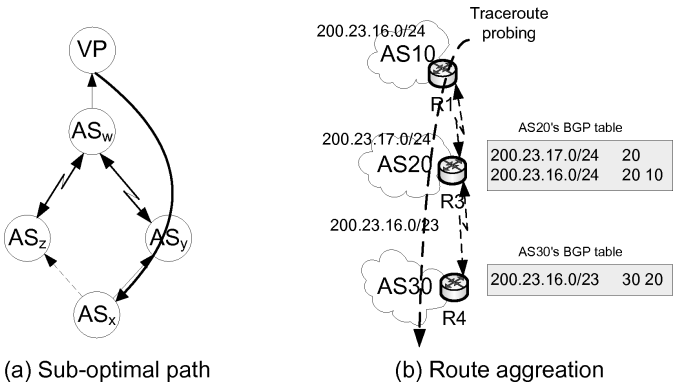

Figure 12. Examples of sub-optimal path to a VP and route aggregation.

B. Identifying Root Causes

In this section, we exploit the reasons why a *customer-provider* or a *peering* link would not appear in the public view and provide examples to explain these cases (the reason for the missing sibling links was discussed in Section IV-D). While we cannot prove that our list of root causes is exhaustive, we believe it accounts for most missing links.

1) Sub-optimal Paths to VPs: The current BGP public view monitoring system has only one or two feeds (*i.e.*, peered routers) in each peered AS, and an AS could contain hundreds of routers while different routers may potentially have different routes even for the same prefix [31, 32]. From this is clear that the public view data could miss many AS links, even those directly connected to vantage-point ASes. Further, according to the BGP specification, if a router receives multiple routes to a prefix, it usually selects one best path according to its policies and exports only that path to its neighbors. For example, consider Figure 12(a), where AS_x is multi-homing to its upstream providers AS_y and AS_z. During the propagation to the VP, some arbitrary AS_w or the VP itself might choose the path between AS_x and AS_y instead of AS_z. The result is that the VP will have no knowledge of the link AS_x-AS_z.

2) Route Aggregation: BGP uses prefix aggregation to reduce the size of routing tables by combining several different

routes into a single one. For instance, in Figure 12(b), AS-20 aggregates two prefixes 200.23.16.0/24 and 200.23.17.0/24 from AS-10 and itself by announcing 200.23.16.0/23 instead. During this process, the previous prefix with the previous AS_PATH is no longer propagated and there is a new route with a new AS_PATH, say $200.23.16.0/23$ 20, which causes the corresponding AS link AS10-AS20 to be hidden.

Without an alternative source for AS path information, BGP paths from the public view are insufficient for determining the effects of route aggregation on inferred AS topologies. By combining AS paths derived from P2P traceroutes with paths from BGP routing tables, however, we are the first to extensively quantify the problem in Section V-C. In the rest of this section, we introduce two special cases of route aggregation: completely hidden ASes and default routing.

Completely hidden ASes: We found 61 of our 23,914 missing links are absent because one of their associated ASes is completely hidden from all the public view VPs. We believe this occurs because all prefixes that are exported via these particular ASes are aggregated between the origin and every VP, making them invisible to all of the monitors. Of these missing links, there are 26 distinct AS numbers absent from the public view. However, Cymru [24] has access to private BGP feeds that may contain ASNs not in the public view, which allows us to discover these new AS numbers. Most of the new ASes ($21/26 = 81\%$) are stub ASes, i.e., they appear at the end of P2P AS paths. Intuitively, such ASes at the edge of the network are relatively far from the public view VPs and thus more likely to be aggregated by their upstream providers before reaching the VPs.

Default routing: We found that over 50% of the public view VPs see only hundreds of prefixes or fewer. We analyzed these VPs and found that they miss significant parts of the active IP address space. For example, the VP of AS8487 observes only the following six prefixes {78.41.184.0/21, 91.103 .239.0/24, 91.103.232.0/22, 82.138.64.0/23, 91.103.232.0/21, 77.95.71.0/24}, and the combination of these prefixes is a small subset of the full IP address space. For such routers, it is likely that a (non-BGP) default forwarding policy is being used to forward traffic for prefixes that are not in the routing table. Thus, default routing (and any other type of non-BGP routing) may prevent links from appearing in the topologies inferred from the public view.

3) Valley-free Policy: Internet routing consists of import and export policies. Import policies specify whether to accept or deny a received route and assign a local preference indicating how favorable the route is, while export policies allow ASes to determine whether to propagate their best routes to the neighbors. Most ASes use the following guidelines in their export settings [27]: while exporting to a *provider* or *peer*, an AS will export the routes from its *customer*s and itself, but not its *provider*s or *peer*s; while exporting to a *customer* or *sibling*, an AS will export its routes and its *customer* routes, as well as its *provider* and *peer* routes. This implies that an AS path should be *valley-free* – after a *provider-to-customer* link or a *peering* link, the AS path cannot traverse another *customer-to-provider* or *peering* link.

Based on these policies, all missing links in Figure 11 can fall into two categories as in Table VII: on the valley-containing path(s) to VPs and on the valley-free path(s) to VPs. The valley-free policy is well known and often explains the missing links, especially the low-tier missing *peering* links [6, 7, 9, 10]. In addition to the missing *peering* links, we observe a substantial number of missing *customer-provider* links with the large-scale P2P traceroutes (as shown in Table VI) for which the valley-free policy is one contributing root cause, for instance, the missing *customer-to-provider* links in patterns (c)-(h) of Figure 11. All these links allow us to evaluate the extent to which the valley-free policy prevents the public view from seeing the AS links. In Section V-C, we will quantify the impact of this reason on missing links; below, we introduce a special case.

Relationship	Valley-free	Valley-containing
peering	(a)(b)	(c)(d)(e)(f)(g)(h)
customer-to-provider	(a)(b)	(c)(d)(e)(f)(g)(h)
provider-to-customer	(a)(b)(c)(d)(e)(f)(g)(h)	N/A

TABLE VII
CATEGORIES FOR MISSING LINKS RELATIVE TO VPS (*Valley-free* MEANS THE LINKS IN RELATED PATTERNS ARE ON THE VALLEY-FREE PATHS TO VPS; *Valley-containing* MEANS THE LINKS IN RELATED PATTERNS ARE ON THE VALLEY-CONTAINING PATHS TO VPS).

Partially cooperative VPs: It seems counterintuitive that VPs cannot see the direct *peering* links and *customer-provider* links for their ASes. While aggregation is a possible reason, we conjecture that another important reason is that some ASes do not treat their route collectors as a "*customer*;" rather, they treat the collector as a "*peer*" and thus do not export their peer and provider routes. We refer to such cases as partially cooperative VPs. Our heuristic for testing this hypothesis is that VPs in this category should not export any other *peering* link or *customer-to-provider* link to route collectors. In our dataset, we found 344 vantage points that miss at least one *peering* link or *customer-to-provider* link. Of these, the public view does not contain *any* direct *peering* or *customer-to-provider* link from 148 ($148/344 = 43\%$) VPs, corresponding to 2116 missing links. While Routeviews [13] asks all of its peered VPs to treat it as a "*customer*" and export their entire routing tables, not all the participating VPs comply for policy reasons. Instead, some VPs treat Routeviews as a "*peer*" and selectively export partial information from their routing tables.

C. Categorizing the Missing Links

The above section broadly categorized missing links according to their location relative to VPs and Table VIII summarized the possible root causes[7] under each pattern of Figure 11; here, we provide a fine-grained classification of missing links. Note that there could be multiple possible explanations for each missing link – for instance, manual inspection revealed that a set of missing links were on a valley-containing path with respect to one VP and a valley-free path with respect to a different VP. The following analysis focuses on the three main root causes: (α) valley-free policy, (β) route aggregation, (γ) sub-optimal paths to VPs. Though this may not be an

[7]The first three reasons are special cases, while the last three reasons are main root causes in our analysis.

		Partially cooperative VPs	Completely hidden ASes	Default routing	Route aggregation	Sub-optimal paths to VPs	Valley-free policy
a	c2p	•		•			
	p2p	•					
	p2c		•		•		
b	c2p			•		•	
	p2p					•	
	p2c		•		•	•	
c/d/e/f/g/h	c2p						•
	p2p						•
	p2c		•		•	•	

TABLE VIII
THE ROOT CAUSES FOR EACH MISSING LINK (C2P, P2P, AND P2C) UNDER EACH MISSING PATTERN (FROM PATTERN (A)-(H)) IN FIGURE 11.

Notation	Description
$\mathbb{M}$	the missing links set $\mathbb{M} = \{m_i, i = 1, 2, ...\}$
$\mathbb{V}$	the VPs set $\mathbb{V} = \{v_j, j = 1, 2, ...\}$
$\mathbb{P}$	the missing patterns set $\mathbb{P} = \{p_k, k = 1, 2, ...\}$
$valley(m_i, v_j, p_k)$	under pattern p_k, if the link m_i is on the valley-containing path to VP v_j
$f_1(m_i)$	the reasons for missing link m_i
$f_2(m_i, v_j)$	the reasons for VP v_j to miss link m_i
$f_3(m_i, v_j, p_k)$	under pattern p_k, the reasons for VP v_j to miss link m_i

TABLE IX
TABLE OF NOTATIONS.

exhaustive list, we believe that a combination of these root causes explains most of the missing links.

Our heuristics for determining the root causes for missing links are shown in Algorithm 2. The main notations used in the algorithm are explained in Table IX. At a high level, the algorithm does the following:

- When a link is found to be on a valley-containing path to a VP, it is classified as missing under valley-free policy since the policy prevents it from being seen by the VP.
- When at least one of the ASes of a missing link is hidden from a VP, this link is classified as missing due to route aggregation. Furthermore, we regard default routing as a special case of route aggregation.
- When both the ASes of a missing link are seen by the VP, the link is classified as missing because it is on a sub-optimal path. Note that this link could also be affected by aggregation, but to be conservative, we do not assign aggregation as one of the causes.

The result of applying the algorithm to our dataset is shown in Table X. To understand the table, $\{\alpha\}$=1.38% means 1.38% of the missing links are solely due to valley-free policy; $\{\alpha, \beta\}$=0.26% means 0.26% are exactly due to both *valley-free policy* and route aggregation; $\{\alpha, \beta, \gamma\}$=75.02% means 75.02% are due to all these three reasons simultaneously. The following can be observed from the table:

- *Route aggregation is a dominant factor*: Though our approach to revealing route aggregation is conservative, we found that about $(\frac{80+61+116+17941}{23914}) = 76.10\%$ of the missing links are related to route aggregation. These missing instances include the 26 completely hidden ASes.
- *BGP policies have a significant effect*: A significant number of links are missing due to valley-free policy and sub-optimal paths to VPs. This confirms previous observations; however, we are the first to quantify their effect on the inferred topology.
- *Missing links have multiple reasons*: Most of missing links are explained by multiple root causes when they are missed by hundreds of the public view VPs. For instance, 1.38% of the missing links are due to valley-free policy, 0.33% due to route aggregation, and 0.27% due to sub-optimal paths to VPs. However, there are 75.02% of the links are missed because all the three causes occur simultaneously.

PROCEDURE Finding Reasons for Missing Links
1 See notations in Table IX; Initialization: $f_3(m_i, v_j, p_k) = \Phi$, $f_2(m_i, v_j) = \Phi$, and $f_1(m_i) = \Phi$;
2 **foreach** *missing link* $m_i \in \mathbb{M}$ **do**
3 **foreach** *VP of public view* $v_j \in \mathbb{V}$ **do**
4 **foreach** *missing pattern* $p_k \in \mathbb{P}$ **do**
5 **if** $\exists$ *one AS attached to* m_i *that is not visible to* v_j **then**
6 **if** $valley(m_i, v_j, p_k) = 1$ **then**
7 $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "(α)";
8 **else** $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "(β)";
9 **if** *both ASes attached to* m_i *are visible to* v_j **then**
10 **foreach** *node attached to missing link* m_i **do**
11 **if** $valley(m_i, v_j, p_k) = 1$ **then**
12 $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "(α)";
13 **else** $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "(γ)";
14 $f_2(m_i, v_j) := \bigcup_{p_k \in \mathbb{P}} f_3(m_i, v_j, p_k)$;
15 $f_1(m_i) := \bigcup_{v_j \in \mathbb{V}} f_2(m_i, v_j)$;
16 Return $f_1(m_i)$: reasons for missing link m_i;

Algorithm 2: Assigning reasons to missing links.

VI. LIMITATIONS

In this paper we showed that using P2P traceroutes reveals a significant number of missing AS links; namely, our dataset adds 12.86% more *customer-provider* links and 40.99% *peering* links to the public view. Thus, publicly available information alone is insufficient for generating more accurate and complete topologies. Note, however, that our approach to extending the AS topology is not meant to replace existing approaches for generating those topologies; rather, it is complementary to existing systems that gather AS topological information.

There are limitations, however, to using traceroutes to extend the AS topology. For one, traceroutes provide IP-level views of the topology, and the public IP-to-AS mapping is neither 100% complete nor accurate. This is a limitation of

Root cause	$\{\alpha\}$	$\{\beta\}$	$\{\gamma\}$	$\{\delta\}$	$\{\alpha,\beta\}$	$\{\alpha,\gamma\}$	$\{\beta,\gamma\}$	$\{\alpha,\beta,\gamma\}$	Unknown
# of links	330	80	65	216	61	4911	116	17941	194
Percentage	1.38%	0.33%	0.27%	0.90%	0.26%	20.54%	0.49%	75.02%	0.81%

TABLE X
CATEGORIZING MISSING LINKS: α - VALLEY-FREE POLICY, β - ROUTE AGGREGATION, γ - SUB-OPTIMAL PATHS, δ - MISSING SIBLING LINKS, "UNKNOWN" IS BECAUSE WE COULD NOT DETERMINE THE RELATIONSHIPS OF THESE LINKS.

all work using traceroutes to extend the AS topology. Using a tier-1 AS's ground-truth as baseline, we have validated our results for this AS and demonstrated that our heuristics can filter *all* of the false links. We cannot, however, determine the extent to which this result applies to other ASes. Specifically, our dataset contained some additional tier-1 links for some other tier-1 ASes but we lack access to their ground-truth to validate these links. By making our uncovered links publicly available we hope to enable researchers with other sources of ground-truth data to collectively validate our results and/or improve the heuristics.

We also note that the AS relationship inference algorithm can incorrectly infer relationships, and this can potentially influence the accuracy of classification of newly discovered links and root causes. Finally, we point out that traceroute measurements are also subject to the constraints of the routers they visit, which can drop probes, silently forward them without altering the TTL or even erroneously modify the TTL in ways that affect the inferred path. While we conservatively select traceroutes to be included in AS topology inference to mitigate this issue, it is possible that other unidentified issues affect our measurements.

Our traceroute measurements discovered large numbers of AS links not visible from public views even though they were restricted to connected BitTorrent hosts as destinations, selected essentially at random. It is likely that P2P users could reveal even greater portions of the AS topology with more sophisticated and controlled measurements that include probes to non-P2P hosts. The design and implementation of techniques to efficiently maximize topology discovery safely (e.g., ensuring probes do not cause a DDoS) is an important area of future work.

VII. RELATED WORK

The Internet connectivity structure is defined by ISP interactions via BGP, which generates and advertises AS paths for routing messages. Chang et al. [5] were among the first to study the completeness of commonly used BGP-derived topology maps. Several projects (*e.g.*, [6, 7]) focused on evaluating and quantifying the public view's coverage of different components of Internet topology. In [8], the authors observed the tradeoff between topology liveness and completeness, and proposed an empirical liveness model to differentiate link birth and death during routing dynamics. He et al. [10] presented a framework to find missing AS links from the commonly-used Internet topology snapshots based on other sources such as additional BGP routing tables, IRR and IXPs.

Measurement platforms, such as DIMES [16], iPlane [15] and Archipelago [17] are providing views of the Internet structure from active measurements. The reach of these platforms has been limited by scalability and/or coverage of active probes from relatively few vantage points. In addition, Lo et al. [33] used active measurements to expose hidden prepending policies and hidden ASes but their work concentrated more on BGP routing dynamics than AS topology. Shavitt et al. [34, 35] studied the importance of vantage points distribution in Internet topology measurements, however they did not investigate the accuracy of their inferred AS links. Augustin et al. [36] leverage active measurements to infer links inside IXPs. More recently, Huffaker et al. [37] take steps toward merging the router-level and AS-level views of the Internet by using a collection of traces from different traceroute measurements.

VIII. CONCLUSION

This paper demonstrates that an approach to measuring the network that leverages P2P systems can significantly improve our understanding of the AS topology. By leveraging measurements from more than 992,000 IPs in 3,700 ASes broadly distributed throughout the Internet, we use a comprehensive set of heuristics to identify 23,914 new links hidden from the public view. While we confirmed that tier-1 AS connectivity is well covered by the public view, our results also indicated that: 1) the public view can miss a substantial number of *customer-provider* links, and 2) missing *peering* links can occur at tiers higher than the VPs in the Internet hierarchy. We characterized the Internet graph properties with these new links. To further understand the reasons behind the missing links, we classified them into a number of root causes and presented the first detailed empirical study that demonstrates the effects of these different root causes on the missing links.

As part of our future work, we intend to investigate how this more complete AS topology affects other commonly held beliefs about Internet properties such as caching and resiliency. To facilitate other research in this area, we have made the set of links used in our study and the inferred relationships publicly available at:
- *http://aqualab.cs.northwestern.edu/projects/SidewalkEnds.html*

REFERENCES

[1] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," in *IEEE Trans. Netw.*, 2001.

[2] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica, "HLP: A Next Generation Interdomain Routing Protocol," in *ACM SIGCOMM*, 2005.

[3] O. Maennel and A. Feldmann, "Realistic BGP Traffic for Test Labs," in *ACM SIGCOMM*, 2002.

[4] K. Park and H. Lee, "On the effectiveness of Route-based packet filtering for distributed DoS attack prevention in power-law Internets," in *ACM SIGCOMM*, 2001.

[5] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," *Computer Networks: The International Journal of Computer and Telecommunications Networking*, vol. 44, no. 6, April 2004.

[6] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "In search of the elusive ground truth: The Internet's AS-level connectivity structure," in *Proc. of ACM SIGMETRICS*, June 2008.

[7] R. Oliveira, D. Pei, W. Willinger, B. Zhan, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Transactions on Networking*, Feburary 2010.

[8] R. Oliveira, B. Zhang, and L. Zhang, "Observing the evolution of Internet AS topology," in *Proc. of ACM SIGCOMM*, August 2007.

[9] R. Cohen and D. Raz, "The Internet Dark Matter: on the Missing Links in the AS Connectivity Map," in *Proc. of IEEE INFOCOM*, April 2006.

[10] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A Systematic Framework for Unearthing the Missing Links: Measurements and Impact," in *Proc. of USENIX NSDI*, April 2007.

[11] Z. M. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an accurate AS-Level traceroute tool," in *Proc. of ACM SIGCOMM*, August 2003.

[12] D. Choffnes and F. Bustamante, "Taming the torrent: A practical approach to reducing cross-ISP traffic in P2P systems," in *Proc. of ACM SIGCOMM*, August 2008.

[13] ROUTEVIEWS, http://www.routeviews.org/.

[14] RIPE, http://www.ripe.net/projects/ris/.

[15] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An Information plane for Distributed Services," in *Proc. of USENIX OSDI*, November 2006.

[16] Y. Shavitt and E. Shir, "DIMES: Let the Internet measure itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, 2005.

[17] CAIDA, "Archipelago," http://www.caida.org/projects/ark/.

[18] "Internet topology collection," http://irl.cs.ucla.edu/topology/.

[19] IRR, "Internet Routing Registry," http://www.irr.net.

[20] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute," in *Proc. of IMC*, October 2006.

[21] Packet Clearing House, "PCB," http://www.pch.net/resources/data.php?dir=/exchange-points.

[22] PeeringDB, "PeeringDB," http://www.peeringdb.com/.

[23] European Internet Exchange Association, "Euro-IX," http://www.euro-ix.net.

[24] Team Cymru, "Ip-to-asn service," http://www.team-cymru.org/Services/ip-to-asn.html.

[25] CAIDA, "AS Relationships," http://www.caida.org/data/active/as-relationships/.

[26] J. Xia, "On the Evaluation of AS Relationship Inferences," in *Proc. of IEEE GLOBECOM*, November 2004.

[27] L. Gao, "On inferring Autonomous System relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, April 2001.

[28] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-law Relationships of the Internet Topology," in *ACM SIGCOMM*, 1999.

[29] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The Origin of Power Laws in Internet Topologies Revisited," in *IEEE Infocom*, 2002.

[30] W. Weibull, "A statistical distribution function of wide applicability," in *J. Appl. Mech. 18:293-7*, 1951.

[31] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," in *Proc. of ACM SIGCOMM*, August 2004.

[32] R. Teixeira and J. Rexford, "A measurement framework for pin-pointing routing changes," in *ACM SIGCOMM Workshop*, August 2004.

[33] S. Lo, R. K. Chang, and L. Colitti, "An Active Approach to Measuring Routing Dynamics Induced by Autonomous Systems," in *Proc. of ExpCS*, June 2007.

[34] Y. Shavitt and U. Weinsberg, "Quantifying the Importance of Vantage Points Distribution in Internet Topology Measurements," in *Proc. of IEEE INFOCOM*, March 2009.

[35] ——, "Quantifying the Importance of Vantage Point Distribution in Internet Topology Mapping," *IEEE Journal on Selected Areas in Communication*, October 2011.

[36] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *IMC*, 2009.

[37] B. Huffaker, A. Dhamdhere, M. Fomenkov, and K. Claffy, "Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers," in *PAM*, 2010.

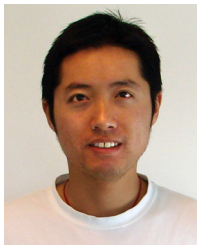

Kai Chen is an Assistant Professor with Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong. He received a Ph.D. degree in Computer Science from Northwestern University, Evanston, IL in 2012. His research interest includes networked systems design and analysis, data center networks, cloud computing, and Internet measurement.

David R. Choffnes received the M.S. and Ph.D. degrees from Northwestern University in 2006 and 2010, respectively, and B.A. degrees in Physics and French from Amherst College in 2002.

He is currently a postdoctoral research associate at the University of Washington, where he was named a Computing Innovation Fellow. His current work focuses on practical solutions to real problems in Internet-scale distributed systems and networks.

Rahul Potharaju is currently pursuing his Ph.D. at Purdue University. Prior to that, in 2009, he earned his M.S. in Computer Science from Northwestern University and worked as a Research Engineer at Motorola Applied Research Center. His current work focuses on large-scale Internet measurements and security aspects of distributed systems.

Yan Chen is an Associate Professor in the Department of Electrical Engineering and Computer Science at Northwestern University, Evanston, IL. He got his Ph.D. in Computer Science at University of California at Berkeley in 2003. His research interests include network security, measurement and diagnosis for large scale networks and distributed systems. He won the Department of Energy (DoE) Early CAREER award in 2005, the Department of Defense (DoD) Young Investigator Award in 2007, and the Microsoft Trustworthy Computing Awards in 2004 and 2005 with his colleagues. Based on Google Scholar, his papers have been cited for over 5,200 times and his h-index is 29.

Fabian E. Bustamante received the M.S. and Ph.D. degrees from the Georgia Institute of Technology in 1997 and 2001, respectively, and the Licenciatura in Computer Science (5-year degree) from the Universidad Nacional de La Patagonia San Juan Bosco, Argentina in 1993.

He is currently an Associate Professor in the Department of Electrical Engineering and Computer Science at Northwestern University. His research focuses on large-scale distributed systems and networking. Bustamante is the recipient of a National Science Foundation CAREER Award (2007) and the Science Foundation of Ireland E.T.S. Walton Fellow Award (2009). He is a Senior Member of the ACM and a member of the IEEE and USENIX.

Dan Pei is a researcher at AT&T Research. He received his PhD degree from UCLA in 2005, and his Bachelor's and Master's degrees from Tsinghua University in 1997 and 2000. His current research interests are network measurement and security.

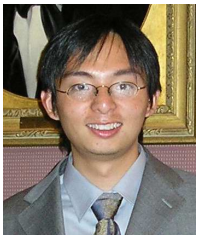

Yao Zhao is a research developer in Bell Labs, Alcatel-Lucent Venture. He got his Bachelor and Master degree in Computer Science at Tsinghua Univ at 2001. And he received his Ph.D. degree of Electrical Engineering and Computer Science at Northwestern University in Jun 2009. His research interests include network measurement, monitoring and security, wireless ad-hoc and sensor networks.