

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

Kai Chen David R. Choffnes Rahul Potharaju Yan Chen
Fabian E. Bustamante Dan Pei† Yao Zhao

Northwestern University

†AT&T Labs – Research

{kchen,drchoffnes,rpo219,jingo}@u.northwestern.edu, {ychen,fabianb}@northwestern.edu
†peidan@research.att.com

ABSTRACT

An accurate Internet topology graph is important in many areas of networking, from deciding 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 limited number of either BGP- or traceroute-based data sources. While techniques for charting the topology continue to improve, the growth of the number of vantage points is significantly outpaced by the rapid growth of the Internet.

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. Based on traceroute measurements from more than 992,000 IPs in over 3,700 ASes distributed across the Internet hierarchy, our proposed heuristics 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 correctly filter out all false links introduced by public IP-to-AS mapping. We have made the identified set of links and their inferred relationships publicly available.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations - Network monitoring

General Terms

Measurement, Management

Keywords

Internet measurement, AS-level topology, Traceroute, BGP

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

CoNEXT'09, December 1–4, 2009, Rome, Italy.

Copyright 2009 ACM 978-1-60558-636-6/09/12 ...\$10.00.

1. INTRODUCTION

An accurate Internet topology graph is important in many areas of networking, from deciding ISP business relationships to diagnosing network anomalies. Appropriately, several research efforts have investigated techniques for measuring and generating such graphs [1–6].

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 designed to potentially cover every corner of the Internet given sufficient numbers of vantage points (VPs).¹ However, all existing traceroute-based projects are restricted by their limited number of VPs. Furthermore, the traceroute measurements provide an IP-level map while our interest is the AS-level map. Converting an IP-level topology to an accurate AS-level one remains an open area of research [7].

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 the diversity of paths covered by traceroutes gathered from hundreds of thousands of peer-to-peer (P2P) users worldwide (Section 2). Specifically, the probes are issued from over 992,000 P2P user 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 3). 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 mappings 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 as ground-truth and show that they filter out all of the false links introduced by public IP-to-AS mapping for this ISP.

¹Vantage 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
Skitter	24	≤ 24
iPlane	192	≤ 192
DIMES	8,059	200
Ono	600,000	6,000

Table 1: Approximate numbers of VPs for topology-gathering projects at the time of publication.

Third, we characterize the new links discovered by our P2P measurements (Section 4). 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.

Fourth, we derive a number of root causes behind the identified missing links, presenting a detailed analysis of their occurrences, and quantify the number of missing links due to each of those reasons (Section 5). Interestingly, many of the missing links (75.02% in our dataset) are missing due to multiple, concurrent reasons.

We discuss limitations of this work in Section 6, review closely related research in Section 7 and conclude in Section 8.

2. P2P FOR TOPOLOGY MONITORING

Understanding and characterizing the salient features of the ever-changing Internet topology requires a system of observation 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, 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 [8], an extension to the Vuze BitTorrent client. The software has been installed more than 600,000 times by hosts located in over 40,000 routable prefixes, spanning more than 6,000 ASes and 192 countries. 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 1 compares the number of unique machines and VPs in our study and in a set of related efforts including Routeviews [9], RIPE/RIS [10], iPlane [11], DIMES [12] and Skitter [13]. For Ono it is difficult to determine the number of unique machines, so we use the number of times the software was installed.

As we show in Section 4, about 23,914 new links are discovered

²Users 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.

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 unobtainable through other approaches such as BGP table dumps and strategic active probing from dedicated infrastruc-




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

Figure 1 shows the distribution of VPs across hierarchical tiers for the publicly-available BGP data and the actually used P2P traceroutes. Note that P2P traceroutes have significantly more VPs compared to the publicly-available BGP data, especially in lower-tier networks. This unique perspective allows us to view previously hidden regions of the network and determine their impact on properties of the Internet topology. The following sections present our methodology for AS-level topology inference and report on our study of missing links.

3. METHODOLOGY

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

3.1 Data Collected

3.1.1 P2P traceroutes

The traceroutes in our dataset are collected by P2P users recording the result of the `traceroute` command provided by their operating system. 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.

The measurement is performed using default settings except that the timeout for router responses is 3 seconds and no reverse DNS lookups are performed. Each peer running our software performs at most one measurement at a time; after each traceroute completes, the peer issues another to a randomly selected destination from the set of connections it has established through BitTorrent.³

³Note that Ono biases BitTorrent connections towards nearby

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. The data was collected from 992,197 distinct peer IPs⁴ in 3,723 unique ASes. Together, these peers probe more than 84 million distinct destination IPs.

3.1.2 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 [9] 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 [10]. We use 10 months of data gathered between Dec 1, 2007 and Sep 30, 2008, the same time period for our P2P traceroute data. Furthermore, we download AS links from the UCLA IRL lab [14] which also contain those links collected from route servers, looking glasses, and IRR [15]. Because the 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” [2, 3]. According to Oliveira et al. [2–4], 10 months of the public view data should be enough to cover “all” the hidden links⁵

3.1.3 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 [2] to process these files and extract the ground-truth AS links that can be used as baseline for our validation.

3.2 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. For one, 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 2, we address this issue through measurements from P2P users. Another limitation is that 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.

peers, so there is a slightly higher probability that traceroutes will be issued to them. We posit that this assists the discovery of new AS links because these nearby peers are often located in access networks at lower tiers of the AS topology.

⁴The number of unique installs and the number of distinct IPs are not equal because each user is often assigned dynamic IP addresses and some users disable traceroute probes.

⁵Hidden links are those policy-allowed links which do not always show up in public view. For example, links only on sub-optimal paths do not show up in public view unless the primary paths fail.




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

Figure 2 illustrates the steps we take to convert traceroute IP paths into their corresponding AS-level paths. In the next subsection, we discuss the processing we perform on IP-level paths (Steps 1–3). Then, after obtaining traceroute AS-level paths based on public IP-to-AS mappings, we adjust the paths to correct for inconsistencies with the corresponding BGP AS paths (Steps 4–6).

3.2.1 IP-level Adjustments

Step 1. Before performing IP-to-AS mappings, 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. Such routers are effectively hidden from our measurement and could lead to falsely inferred AS links. There are other known problems such as load balancing, zero-TTL forwarding and address rewriting of gateway routers that would cause routing loops [16]. To avoid the potential problems of these issues, we conservatively remove the entire path from our analysis.

Step 2. As explained by Mao et al. [7], paths that traverse Internet eXchange Points (IXPs) can lead to falsely inferred AS links. Using a list of known IXP prefixes, such as PCH [17], PeeringDB [18] and Euro-IX [19], we remove from each path any hop that belongs to an IXP. 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 kind of false links because they are known to be incomplete. Our heuristics in the next subsection will address the remaining problems of IXPs at the AS level.

Step 3. After the first two steps, we simply convert IP-level paths to AS-level ones by directly using the AS mappings provided by Team Cymru [20], which incorporates both publicly available and private BGP information.

3.2.2 IP-to-AS Mapping

The next phase in our analysis is to address the issues contained in the conversion from IP-level paths into AS-level ones. While previous work has investigated the problem of accurate IP-to-AS mappings in networks where BGP data is available [7], our study is the first to address the problem for an arbitrary (and large) set of networks. The proposed techniques will consider the scenario where the traceroute VPs are not the same as the BGP VPs, which makes them generically applicable. Furthermore, unlike previous work using traceroutes, we expect to see a significant number of new links compared to the public view because our software monitors a larger portion of the Internet. The key challenge that we address in this section is how to distinguish the real new links from those that are falsely inferred due to incorrect IP-to-AS mappings. To evaluate the quality of our heuristics, we compare our results with ground-truth from a tier-1 ISP.

In the first phase of our analysis, we simply convert IP-level paths to AS-level ones (*i.e.*, Step 3). The authors in [7] identify

	Problem	Symptom				Filtering heuristic(s)
		Loop	Missing hop	Substitute hop	Extra hop	
Incomplete paths	Unresolved hops within an AS Unmapped hops between ASes MOAS hops at the end	Problem addressed in [7]				Steps 1, 4 Step 4 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 2: Problems within traceroute-inferred AS-level paths, symptoms for these problems, and the step(s) we take to solve them. Note that we do not consider the symptoms for “incomplete paths” because they are addressed in [7]. Reading the last row of the table, “private peering interface IPs” will cause missing hop problem, and we address this problem in Step 6 of our techniques.

several patterns of discrepancies between traceroute and BGP paths (as shown in Table 2), each of which entails a difference of at most one AS hop (e.g., an AS is missing from the path, an extra AS appears in the hop, or a substitute AS appears in the path). To account for these discrepancies while still preserving true new AS links discovered by traceroute measurements, we mark a new link to be *pending* if it could be corrected by techniques used by Mao et al. [7]; otherwise, we assume that the new link is real. In our implementation, we conservatively modify all the *pending* links such that they are consistent with the corresponding BGP paths. We emphasize that this approach prevents false positives, but may filter out real links not present in BGP. Note that unlike the work in [7], we only correct the AS-level paths generated by traceroutes so that we can confidently infer new links. Correcting the IP-to-AS mappings is beyond the scope of this paper.

We show in Table 2 that our implementation for converting IP paths to AS paths can address most of the well-known problems identified by Mao et al. [7]. Since their work addressed the problem of incomplete paths, we directly apply their techniques to our dataset (Steps 1 & 4). However, identifying and modifying falsely mapped AS links is a significant challenge that we address in this work.

Step 4. Besides dealing with incomplete paths in this step, we further filter additional IXPs. While we have used the available IXP prefixes to delete the AS hops belonging to any IXP, our list of IXP prefixes is not complete. For those IXPs that do not make their prefixes publicly available, we still can identify them by using the IXP participating AS list we have. We pick out the AS hops in the middle of traceroute AS paths that are publicly mapped to multiple ASes and check if these multiple ASes are collocated in an IXP. This occurs when the shared infrastructure address is originated into BGP by multiple participating ASes. However, we cannot use this approach to identify IXP that use their own AS numbers – a limitation that we address in Step 6.

Step 5. A single organization may own and manage multiple sibling ASes. Among two sibling ASes, one AS may use some address blocks from another to number its equipment or during route propagation only one AS includes its AS number in the BGP AS path while the other does not. This would cause problems within the traceroute AS paths. To mitigate such problems, we download the known sibling ASes from CAIDA [21]. 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 measurement, 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 that mitigate the remaining problems when sibling ASes cause discrepancies between traceroute AS paths and BGP AS paths.

PROCEDURE Address issues within traceroute AS paths
1 Initialization: set the DISTANCE of each AS link on the traceroute AS paths;
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 $DISTANCE(B, C) = 1$ then
4 AS link $B-C$ is considered <i>true</i> ;
5 if $DISTANCE(B, C) = 2$ 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 DISTANCE of each of these two links as 1 (For multiple X s, choose the longest matched one; For instance, both [$A B X_1 C D$] and [$A' B X_2 C D'$] exist, the first one matches the traceroute AS path [$A B C D$] better, hence X_1 is preferred;
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 DISTANCE of each link as 1 (longest match for multiple X s);
12 else
13 if $DISTANCE(A, C)=1$ (or $DISTANCE(B, D)=1$) then
14 Delete B (or C) and set the DISTANCE of link $A-C$ (or $B-D$) as 1 ;
15 if $DISTANCE(A, C) \neq 1$ and $DISTANCE(B, D) \neq 1$ then
16 Mark $B-C$ as a real link and set DISTANCE(B, C) as 1 ;
17 end
18 if $DISTANCE(B, C) \geq 3$ then
19 if $DISTANCE(A, C)=1$ (or $DISTANCE(B, D)=1$) then
20 Delete B (or C) and set the DISTANCE of link $A-C$ (or $B-D$) as 1 ;
21 if $DISTANCE(A, C) \neq 1$ and $DISTANCE(B, D) \neq 1$ then
22 Mark $B-C$ as a real link and set DISTANCE(B, C) as 1 ;
23 end
24 end
25 Return the traceroute AS path when DISTANCES for all links are 1;

Algorithm 1: Heuristics in Step 6 of Figure 2.

Step 6. Algorithm 1 addresses a variety of issues with traceroute AS paths that remain after the previous five steps. Below, we discuss how these heuristics apply each symptom, i.e., loops and missing/extra/substitute hops as shown in Table 2.

- **Loop.** Loops in traceroute AS paths can happen due to unan-

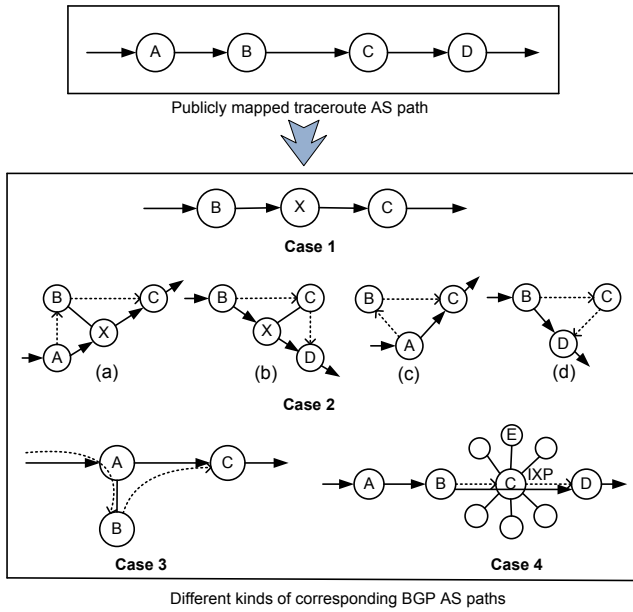


Figure 3: Relationship between traceroute AS paths and BGP AS paths for several cases. Dotted arrows are traceroute AS paths and solid arrows are the corresponding BGP AS paths.

nounced IP addresses, sibling ASes, or route anomalies on the forwarding paths. In our dataset, the loops in the traceroute AS paths are rare. While not all AS-level loops are invalid, we conservatively discard these paths.

- **Missing hop.** We check the public view connectivity graph to calculate the *DISTANCE* (*DISTANCE* is defined as the number of hops between the two end point ASes of the traceroute AS link according to the BGP public view graph) of each link on each traceroute AS path. If the *DISTANCE* is 2 and we find in the BGP AS paths the corresponding route(s) [...*BXC*...] (case 1 in Figure 3), we conservatively add one hop 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 reasons:

- **Private peering interface IPs.** AS links *B-X* and *X-C* are both private peerings using the IP addresses from *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;
- **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;
- **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, the *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;
- **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 *DISTANCE* 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:

- **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*...]) – this is a substitute hop problem.

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 diverse addresses can falsely generate an 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 AS path is [...*AC*...] (or [...*BD*...]) – this is an extra hop problem;

- **IXPs or Sibling ASes.** As explained in previous subsections, IXPs can lead to extra hops and sibling ASes can lead to substitute/extra hops.
- **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 traceroute AS paths consistent with BGP AS paths by replacing the middle hop with *X* or deleting it (line 11 ~ 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.

- **Special case of Extra hop.** Though rare, we found cases where traceroute AS links have a *DISTANCE* ≥ 3 . We posit one plausible scenario in case 4 of Figure 3. Here *C* is an IXP with who has its own AS number but only announces its addresses via a particular participant, say AS *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 case. 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.

3.2.3 Validation

After applying all the heuristics in the previous section, we are left with 100,000 AS links discovered through P2P traceroutes. 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 3: Statistics of the identified missing links (PV stands for public view; New# is the number of missing links not in PV).

now validate a significant portion of these links with the ground-truth information from a tier-1 AS (the number is on the order of thousands⁶). 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 percentage 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 indicated thousands of links to 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 4.

Line # in Algorithm 1	False links left
-	48.80%
8	10.47%
11	5.13%
13	0.47%
18	0

Table 4: Percent of false links remaining after each filtering step.

Distance(B, C)=2 and [...BXC...] exists in BGP (line 8): We see several hundreds of unique 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.

Distance(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 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 in 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.

Distance(B, C) \geq 3 (line 18): We have no specific ground-truth files 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

⁶Because this information is proprietary, we cannot disclose the precise number of AS links so we use percentages in this section.

0.47% of the links to the tier-1 AS had DISTANCE \geq 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 have not seen at least 21.3% of total links in the tier-1 AS’s 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.

3.3 Policy Inference

After extracting the AS links, we infer the business relationships between ASes based on the PTE algorithm proposed by Xia [22]. After improving the seminal work by Gao [23], the PTE approach is considered to outperform most other approaches [6]. Most AS links are classified as one of three kinds of relationships: *customer-provider* links, *peering* links, and *sibling* links. In our study, we also decompose *customer-provider* links into *customer-to-provider* links and *provider-to-customer* links directionally. Further, we assume that the AS relationships did not change significantly within our ten-month measurement period. To justify this, we sample the AS relationships from CAIDA [21] for the past five years. We check the relationships at ten-month intervals and find that more than 98.5% of AS pairs do not change their relationships.

We also 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. [2, 3], which relies on the number of downstream customer ASes to classify each AS.

4. THE MISSING LINKS

After generating an AS topology from P2P traceroutes, we found a significant number of new AS links (including *customer-provider*, *peering* and *sibling*), as shown in Table 3. In this section, we use our set of missing links to determine the public view’s coverage of each class of AS links and where these links are missed by public views.

4.1 Coverage of tier-1 AS links

We begin by focusing on the tier-1 AS connectivities, listed in Table 5. 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 [2] 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 explanations for this to occur. 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)

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 5: 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).




Figure 4: The missing provider links.

current public view for each AS may capture an incomplete view of the AS. In addition, some tier-1 ISPs intentionally do not announce all of their prefixes (*e.g.*, those longer than /24), which prevents the public view from seeing the corresponding links.

4.2 Coverage of customer-provider links

We now turn our attention to the set of *customer-provider* links discovered by P2P traceroutes. Table 3 shows that P2P traceroutes discover 12.86% additional *customer-provider* links missing from public views. To put this in context, recent work [4] 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 views do 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, due to factors such as route aggregation (explained later in Section 5.2.2, the assumption is often violated thus these public views are 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 3.3 to classify each AS into a tier, then group all of the fractions for each tier. Figure 4 and Figure 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 ev-




Figure 5: The missing customer links.




Figure 6: The missing peering links.

ery tier. More importantly, we observe that the fraction of missing provider links of an AS somewhat correlates 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.

4.3 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 [2, 6]. 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 5.

4.4 Missing sibling AS links

We revealed 216 additional *sibling* links which are missing from public view. We think that one reason behind these missing *sibling* links could be due to route announcement in BGP. To illustrate this case, consider two *sibling* ASes: AS1 and AS2. During BGP route

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 6: Numbers of missing/visible links in each pattern of Figure 7. 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 7: 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.

announcement, the AS path announced by these ASes might contain the AS number of either AS1 or AS2 but not both. The result is that the public view cannot see this *sibling* link; however, when probes between P2P users in these two ASes traverse this link, they reveal both AS numbers and thus the *sibling* link.

5. IN SEARCH OF ROOT CAUSES

In the previous section, we characterized links found through 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 Internet-wide AS topology. In our analysis, we observe that the missing of an AS link *cannot* be explained by one or more root causes. Thus, we determine a *set* of root causes that could be responsible for a missing link.

5.1 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. Our method is to check the BGP AS paths to trace the routes from VPs to missing links. In other words, for a missing link AS1-AS2 and any AS path containing AS1 or AS2, we record the route pattern from the VP to the associated AS. All the found route patterns are shown in Figure 7. 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.

5.1.1 Observations

Table 6 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 6 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:

- It is generally believed that a monitor with full BGP table⁷ can discover all the connections of its upstream providers [2, 3]. 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 well known that many *peering* links are missed in the low-tiers of the Internet hierarchy [2, 6], and our result for pattern (h) in Table 6 confirms this fact. However, we also find many instances of upstream *peering* links being invisible to downstream full table monitors (for instance, pattern (b)). This indicates that ASes located low in the hierarchy are not solely responsible for missing peering links.

5.2 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 4.4). While we cannot prove that

⁷A full-table VP means the VP covers a complete prefix space.

our list of root causes is exhaustive, we believe it accounts for most missing links.




Figure 8: Illustrations of (a) sub-optimal path to a VP and (b) route aggregation.

5.2.1 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 [24, 25]. 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 8(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 .

5.2.2 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 8(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 5.3. 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 [20] 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 prefixes 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. We confirmed this fact through a thread of discussion on the NANOG mailing list [26]. Thus, default routing (and any other type of non-BGP routing) may prevent links from appearing in the topologies inferred from the public view.

5.2.3 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 [23]: while exporting to a *provider* or *peer*, an AS will export the routes from its *customers* and itself, but not its *providers* or *peers*; 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.

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

Table 7: 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).

Based on these policies, all missing links in Figure 7 can fall into two categories as in Table 7: 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 [2, 3, 5, 6]. 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 6) 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 7. 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 5.3, we will quantify the impact of this reason on missing links; below, we introduce a special case.

Partially cooperative VPs: It seems counterintuitive that VPs cannot see the direct *peering* links and *customer-provider* links for their ASes. We conjecture that one possible 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

peers and providers. 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. To validate this result, we contacted a Routeviews administrator who confirmed our findings [27]. While Routeviews [9] 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.

		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	c2p						•
	p2p						•
	p2c		•		•	•	
d	c2p						•
	p2p						•
	p2c		•		•	•	
e	c2p						•
	p2p						•
	p2c		•		•	•	
f	c2p						•
	p2p						•
	p2c		•		•	•	
g	c2p						•
	p2p						•
	p2c		•		•	•	
h	c2p						•
	p2p						•
	p2c		•		•	•	

Table 8: The potential root causes for each kind of missing link (c2p, p2p, and p2c) under each kind of missing pattern (from pattern (a)–(h)) in Figure 7. The first three reasons are special cases, while the last three reasons are the main root causes in our analysis. Reading the first row of the table, pattern (a) may miss c2p links due to partially cooperative VPs and default routing, miss p2p links due to partially cooperative VPs, and miss c2p link due to completely hidden ASes and route aggregation.

5.3 Categorizing the Missing Links

The previous subsection broadly categorized missing links according to their location relative to VPs and Table 8 summarized the possible root causes under each pattern of Figure 7; here, we provide a fine-grained link classification. When categorizing missing links in this way, there could be more than one plausible reason for them to be absent from the public view. For instance, when manually investigating a set of missing links, we found that they were on a valley-containing path with respect to one VP and a valley-free path with respect to a different VP. In this section, we focus on the three main root causes: (α) valley-free policy, (β) route aggregation, (γ) sub-optimal paths to VPs. Though this is not an exhaustive list, we believe that a combination of these three root causes explains most of the missing links.

Notation	Description
\mathbb{M}	the missing links set $\mathbb{M} = \{m_i, i = 1, 2, \dots\}$
\mathbb{V}	the VPs set $\mathbb{V} = \{v_j, j = 1, 2, \dots\}$
\mathbb{P}	the missing patterns set $\mathbb{P} = \{p_k, k = 1, 2, \dots\}$
$valley(m_i, v_j, p_k)$	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 9: Table of notations.

PROCEDURE Finding Reasons for Missing Links	
1	See notations in Table 9; 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	$f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup \text{“}(\alpha)\text{”}$;
8	else
9	$f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup \text{“}(\beta)\text{”}$;
10	end
11	else
12	if both ASes attached to m_i are visible to v_j then
13	foreach node attached to missing link m_i do
14	if $valley(m_i, v_j, p_k) = 1$ then
15	$f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup \text{“}(\alpha)\text{”}$;
16	else
17	$f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup \text{“}(\gamma)\text{”}$;
18	end
19	end
20	end
21	end
22	$f_2(m_i, v_j) := \bigcup_{p_k \in \mathbb{P}} f_3(m_i, v_j, p_k)$;
23	end
24	$f_1(m_i) := \bigcup_{v_j \in \mathbb{V}} f_2(m_i, v_j)$;
25	end
26	Return $f_1(m_i)$: reasons for missing link m_i ;

Algorithm 2: Assigning reasons to missing links.

Our heuristic for determining the root causes for missing links is shown in Algorithm 2. A high level, 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 because this policy prevents it from being seen by the VP.

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 10: 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. Reading the table, $\{\alpha\}$ means 1.38% of the missing links are solely due to valley-free policy; $\{\alpha, \beta\}$ means 0.26% are exactly due to both *valley-free policy* and *route aggregation*; $\{\alpha, \beta, \gamma\}$ means 75.02% are due to all these three reasons simultaneously.

- When at least one of the ASes of a missing link is hidden from a VP, this link is classified as missing due to aggregation. Sometimes, default routing is the reason for a missing link; we regard it 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 10. 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.

6. 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 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 all work using traceroutes to extend the AS topology. Using a tier-1 AS’s ground-truth as baseline, we have validated our result related to this AS and demonstrated that our proposed heuristics can filter *all* of the false links. It should be admitted that we cannot determine the extent to which this result applies to other ASes. Especially, 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. However, validating with the known tier-1 AS’s ground-truth increases our confidence about our result. In addition, our publically available uncovered links also enable researchers to collectively validate

our results utilizing ground truth accessible to them and 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 our conservative approach to selecting traceroutes to include for inferring the AS topology mitigates this issue, it is possible that other unidentified issues affect our measurements.

7. RELATED WORK

The Internet’s connectivity structure is defined by ISP interactions via the BGP, which generates and advertises AS paths for routing messages. Chang et al. [1] were among the first to study the completeness of commonly used BGP-derived topology maps. Several projects (*e.g.*, [2, 3]) focused on evaluating and quantifying the public view’s coverage of different components of Internet topology. In [4], 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. [6] 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 Skitter [13], DIMES [12], and iPlane [11] are providing views of the Internet structure from active measurements. The reach of these platforms have been limited by scalability and/or coverage of active probes from relatively few vantage points. In addition, Lo et al. [28] used active measurements to expose hidden prepending policies and hidden ASes but their work concentrated more on BGP routing dynamics than the AS topology. Recently, Shavitt et al. [29] studied the importance of vantage points distribution in Internet topology measurements, however they did not investigate the accuracy of their inferred AS links.

In contrast to all previous work, our paper is the first to use a P2P approach to discover AS-level paths through traceroute probes. As part of our work, we develop comprehensive heuristics to accurately convert our IP-level paths to AS-level. Using the largest and most widely distributed set of vantage points to date, we are able to expose previously hidden regions of the Internet, identifying links missing from the public view and investigating the causes.

8. 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 ac-

curately 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. 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 (including missing ones) and the inferred relationships publicly available at:

- <http://aqualab.cs.northwestern.edu/projects/SidewalkEnds.html>

Acknowledgment

We would like to thank our shepherd Renata Teixeira and the anonymous reviewers for their valuable comments. We are also grateful for Ricardo Oliveira's help with data processing. This work is supported by a grant from NSF NeTS Award 0917233. Bustamante and Choffnes are supported in part by US NSF CAREER Award CNS-0644062. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources.

9. REFERENCES

- [1] 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.
- [2] 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*, Annapolis, MD, June 2008.
- [3] R. Oliveira, D. Pei, W. Willinger, B. Zhan, and L. Zhang, "Quantifying the completeness of the observed Internet AS-level structure," UCLA, Tech. Rep. 080026, September 2008.
- [4] R. Oliveira, B. Zhang, and L. Zhang, "Observing the evolution of Internet AS topology," in *Proc. of ACM SIGCOMM*, Kyoto, Japan, August 2007.
- [5] R. Cohen and D. Raz, "The Internet Dark Matter: on the Missing Links in the AS Connectivity Map," in *Proc. of IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [6] 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*, Cambridge, MA, April 2007.
- [7] Z. M. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an accurate AS-Level traceroute tool," in *Proc. of ACM SIGCOMM*, Karlsruhe, Germany, August 2003.
- [8] D. Choffnes and F. Bustamante, "Taming the torrent: A practical approach to reducing cross-ISP traffic in P2P systems," in *Proc. of ACM SIGCOMM*, Seattle, WA, August 2008.
- [9] ROUTEVIEWS, <http://www.routeviews.org/>.
- [10] RIPE, <http://www.ripe.net/projects/ris/>.
- [11] 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*, Seattle, WA, November 2006.
- [12] Y. Shavitt and E. Shir, "DIMES: Let the Internet measure itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, 2005.
- [13] CAIDA, "Skitter," http://www.caida.org/tools/measurement/skitter/as_adjacencies.xml.
- [14] "Internet topology collection," <http://irl.cs.ucla.edu/topology/>.
- [15] IRR, "Internet Routing Registry," <http://www.irr.net>.
- [16] B. Augustin, X. Cuvelier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute," in *Proc. of IMC*, Rio de Janeiro, Brazil, October 2006.
- [17] Packet Clearing House, "PCB," <http://www.pch.net/resources/data.php?dir=/exchange-points>.
- [18] PeeringDB, "PeeringDB," <http://www.peeringdb.com/>.
- [19] European Internet Exchange Association, "Euro-IX," <http://www.euro-ix.net>.
- [20] Team Cymru, "Ip-to-asn service," <http://www.team-cymru.org/Services/ip-to-asn.html>.
- [21] CAIDA, "AS Relationships," <http://www.caida.org/data/active/as-relationships/>.
- [22] J. Xia, "On the Evaluation of AS Relationship Inferences," in *Proc. of IEEE GLOBECOM*, Dallas, USA, November 2004.
- [23] L. Gao, "On inferring Autonomous System relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, April 2001.
- [24] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," in *Proc. of ACM SIGCOMM*, Portland, Oregon, August 2004.
- [25] R. Teixeira and J. Rexford, "A measurement framework for pin-pointing routing changes," in *ACM SIGCOMM Workshop*, Portland, Oregon, August 2004.
- [26] North American Network Operators Group, "NANOG email list archives," <http://www.merit.edu/mail.archives/nanog/>.
- [27] Authors and Admin, "Personal Communication with Routeview Administrators," 2008.
- [28] S. Lo, R. K. Chang, and L. Colitti, "An Active Approach to Measuring Routing Dynamics Induced by Autonomous Systems," in *Proc. of ExpCS*, San Diego, CA, June 2007.
- [29] Y. Shavitt and U. Weinsberg, "Quantifying the Importance of Vantage Points Distribution in Internet Topology Measurements," in *Proc. of IEEE INFOCOM*, Rio de Janeiro, Brazil, March 2009.