NORTHWESTERN
UNIVERSITY

Electrical Engineering and Computer Science Department

Technical Report
NWU-EECS-09-16
July 28, 2009

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

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 number of vantage points continues to shrink relative to the fast-paced 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 software, we show that this approach indeed reveals significant new topological information. Based on traceroute measurements from more than 580,000 hosts in over 6,000 ASes distributed across the Internet hierarchy, our proposed heuristics identify 23,914 new AS links not visible to the public view – 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. In addition, for the benefit of the community, we will make all our identified missing links publicly available.

Keywords: AS topology, P2P, traceroute

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
†Department of Electrical Engineering and Computer Science, Northwestern University
‡AT&T Labs – Research

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 number of vantage points continues to shrink relative to the fast-paced 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 software, we show that this approach indeed reveals significant new topological information. Based on traceroute measurements from more than $580,000$ hosts in over $6,000$ ASes distributed across the Internet hierarchy, our proposed heuristics identify $23,914$ new AS links not visible to the public view – 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. In addition, for the benefit of the community, we will make all our identified missing links publicly available.

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 either BGP- or traceroute-based data sources. 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 available BGP paths do not cover the entire Internet due to issues such as route aggregation, hidden sub-optimal paths and policy filtering.

While BGP paths represent the "control plane" for the Internet, they do not necessary reflect the true paths data packets travel. Traceroute measurements provide the ability to infer the actual paths that data packets take when traversing the Internet. Because they are active measurements, traceroute probes can be designed to cover every corner of the Internet given sufficient numbers of vantage points (VPs).[1] However, all current traceroute-based projects are restricted by their limited number of VPs. Further, these measurements provide an IP-level rather than the most commonly used 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 P2P users worldwide (Section 2). Specifically, the probes are issued from over 580,000 P2P users in 6,000 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 each of these problems. Our proposed techniques for correcting IP-to-AS mappings are generic and work for the scenarios that 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 it correctly filters out all the false links introduced by public IP-to-AS mapping.

Third, we characterize the new links discovered by our P2P measurements (Section 4) – Our study reveals 12.86% more *customer-provider* links than what can be found in the public view. We also find that some common assumptions about the visibility of paths according to ISP relationships are routinely violated. For example, we have successfully found 40.99% more missing *peering* links. However, these links do not necessarily fall below the level of VPs. In other words, a VP could even miss its upstream *peering* links.

Fourth, we derive a number of root causes behind the uncovered missing links, presenting a detailed analysis of their

[1] By *vantage point*, we mean a unique AS.

Project	# unique machines	# unique ASes
Routeviews/RIPE	790	438
Skitter	24	≤ 24
iPlane	192	≤ 192
DIMES	8,059	200
Ours	580,000	6,000

Table 1: The VPs for each project, approximately.

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.

In the remainder of this paper we discuss the value of paper as well as its limitations in Section 6, review closely related work 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 VPs involved in the measurements.

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, we eliminate the need to place BGP monitors in each ISP; rather, each participating host in our system can contribute to the AS topology measurement study simply by performing traceroute measurements.

Through an extension to a popular BitTorrent client currently installed by 580,000 peers[2] located in over 40,000 routable prefixes, spanning more than 6,000 ASes and 192 countries, our software collects traceroute measurements between connected hosts. 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 contrasts the number of unique machines and VPs in our study and in a set of related efforts including Routeviews [8], RIPE/ RIS [9], iPlane [10], DIMES [11] and Skitter [12].

As we show in Section 4, 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 public view 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 infrastructure.

Figure 1 shows the layer-wise distribution of VPs for the public view and our existing P2P traceroutes. It is remarkable that our P2P traceroutes have an overwhelming advantage over the public view, especially in these low tier networks. The better coverage of P2P VPs could conceive a different perspective of the Internet graph and its potentially missing links. The following sections present our methodology for AS-level topology inference and report on our study of missing links.

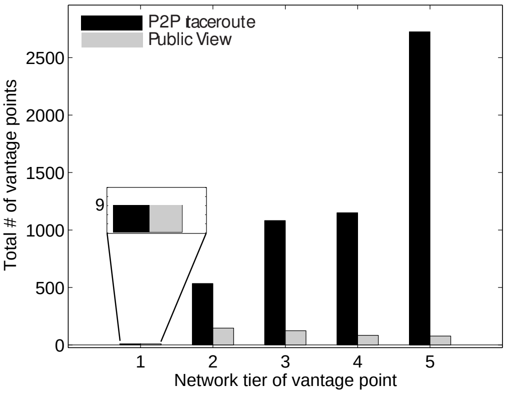

Figure 1: Distribution of VPs with respect to their network tiers.

3. METHODOLOGY

In this section, we present our methodologies. After we describe our datasets, we 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 detail 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. Not surprisingly, 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.

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 more than 580,000 distinct peers in 6,600 unique ASes.

3.1.2 BGP feeds

[2] We use "peer" for P2P user, and italic "*peer*" for peering ISP.

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 [8] collected at route-views.oregon-ix.net, which is the most widely used BGP archive so far, 6 other Oregon route servers and 16 route collectors of RIPE/RIS [9]. 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 UCLA IRL lab [13] which also contain those links collected from route servers, looking glasses, and IRR [14]. However, they do not provide BGP AS paths and their corpus has not included information from some newly added VPs. So combining all these sources of AS links together, we are able to get the most complete AS links dataset. Throughout 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 links[3] in the Internet graph.

3.1.3 Ground-truth data

We have the data of proprietary router configurations and syslogs from a tier-1 ISP. This is a major source for deriving the ground-truth connectivities of this tier-1 network. The data include historical configuration files of more than one thousand routers in these two networks, historical syslog files from all routers in the tier-1 network. We also have access to iBGP feeds of several routers in this network. We directly use the heuristics in [2] to process these files and extract desired data.

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 (e.g., due to aggregation and unannounced shared infrastructure addresses). 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 each of the remaining limitations.

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

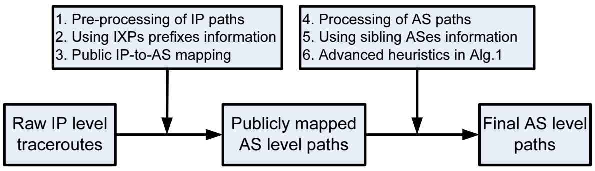
1. Pre-processing of IP paths
2. Using IXPs prefixes information
3. Public IP-to-AS mapping
4. Processing of AS paths
5. Using sibling ASes information
6. Advanced heuristics in Alg.1
Raw IP level traceroutes
Publicly mapped AS level paths
Final AS level paths

Figure 2: The flowchart from IP paths to AS paths.

3.2.1 IP-level Adjustments

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 problem such as routing loops and appearance of the destination address in the middle of the traceroutes [15]. To prevent these problems from happening, we conservatively remove the entire path from our analysis.

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 [16], 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.

3.2.2 IP-to-AS Mapping

The next step in our analysis of traceroute data is to convert 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. Unlike previous work, we expect to see a significant number of new links compared to the public view simply 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 by directly using the AS mappings provided by Team Cymru [17], which incorporates both publicly available and private BGP information. The authors in [7] identify several patterns of discrepancies between traceroute and BGP paths, 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 assume

[3] By hidden links, they mean *customer-provider* links and policy-allowed *peering* links, *i.e.*, upstream *peering* links, on backup paths.

	Problem	Symptom				Filtering Heuristic
		Loop	Missing hop	Substitute hop	Extra hop	
Incomplete paths	Unresolved hops within an AS	Problem addressed in [7]				Steps 1, 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 2: Causes for incorrect traceroute-inferred AS-level paths, symptoms for these causes, and the step(s) we take to eliminate them. Note that we do not consider the symptoms for "incomplete paths" because they are addressed in [7]. To understand the table, for instance, "private peering interface IPs" will cause missing hop, and we address this problem in our step 6 of our heuristics.

that a link is false if it could be corrected by techniques used by Mao et al. [7] (Table 2); otherwise, we assume that the new link is real. 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 of converting IP paths to AS paths can address most of the well-known problems identified within IP-to-AS mapping [7]. The incomplete paths have been addressed in [7] and we use their techniques to filter in ***Steps 1*** and ***4***; while the challenges are within identifying and modifying falsely mapped AS links.

Step 4: Besides addressing incomplete paths in this step, we further filter additional IXPs. Although we have used the available IXP prefixes to inspect the traceroute IP paths and directly delete the hops belonging to any IXP, our list of IXP prefixes is not complete. Fortunately, an IXP may either have its own ASN to originate routes or announce its infrastructure addresses by one or more of the participating ASes. So in the traceroute AS paths, if we see the cases where an intermediate hop is mapped to multiple ASes, we check and delete it when it is from an IXP. Note that we cannot find the IXPs who have their own ASNs here. But they are addressed in our advanced heuristics.

Step 5: To mitigate the problems of sibling ASes, we download the known sibling ASes from CAIDA [18]. 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...]; In our measurement, we also see the cases where the traceroute AS path is [...WY...] while a corresponding BGP AS path is [...WXZ...]. For all these cases, we use the BGP AS path to modify the traceroute AS paths. Again, publicly available information is limited. Our advanced heuristics will mitigate problem for the rest when they cause mismatch between traceroute AS paths and BGP AS paths.

Step 6: When analyzing how the advanced heuristics address the potential problems, we use the symptoms to help understand our algorithm:

Loops: Loops on the traceroute AS paths can happen due to unannounced IP addresses, sibling ASes, or route anomalies on the forwarding paths. Fortunately, in our dataset, the looping traceroute AS paths only account for a very small portion. To be conservative, we discard these looping AS paths and only focus on normal AS paths.

```
   PROCEDURE Convert traceroute IP paths to AS level paths
1  Initialization: set the DISTANCE of each AS link on the traceroute AS
   paths according to the connectivity graph of public view;
2  foreach AS link on the traceroute AS paths, e.g., link $B$-$C$ in top Figure 3 do
3      if DISTANCE($B, C$) = 1 then
4          AS link $B$-$C$ is considered true and move on;
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 and move on (If there are multiple
               $X$s match, using the longest match to choose one);
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 and move on (longest match for
                   multiple $X$s);
               else
12                 if DISTANCE($A, C$)=1 (or DISTANCE($B, D$)=1)
                   then
13                     Delete $B$ (or $C$) and set the DISTANCE of link
                       $A$-$C$ (or $B$-$D$) as 1 and move on;
14                 if DISTANCE($A, C$)$\neq$ 1 and DISTANCE($B, D$)$\neq$ 1
                   then
15                     Regard $B$-$C$ as a truly new link and set
                       DISTANCE($B, C$) as 1 and move on;
               end
16     if DISTANCE($B, C$) $\geq$ 3 then
17         if DISTANCE($A, C$)=1 (or DISTANCE($B, D$)=1) then
18             Delete $B$ (or $C$) and set the DISTANCE of link $A$-$C$ (or
               $B$-$D$) as 1 and move on;
19         if DISTANCE($A, C$)$\neq$ 1 and DISTANCE($B, D$)$\neq$ 1 then
20             Regard $B$-$C$ as a truly new link and set
               DISTANCE($B, C$) as 1 and move on;
   end
21 Return the traceroute AS path when DISTANCEs for all links are 1;
```

Algorithm 1: Advanced heuristics in Figure 2.

Missing hop: Our advanced heuristics check the public view connectivities to calculate the DISTANCE 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 (in ***Step-6.8***, *i.e.*, line 8). This mismatch could result from the following cases: 1) B-X and X-C are both private peerings using the IP addresses from A and C respectively. When traceroute probings go from A to X and then immediately exit X to enter B, it would cause [...BC...] in traceroute AS path

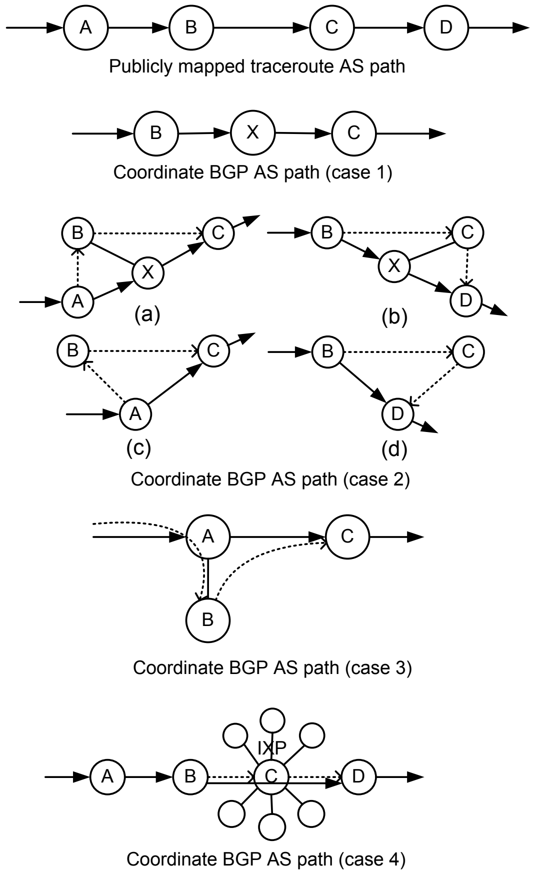
A
B
C
D
Publicly mapped traceroute AS path
B
X
C
Coordinate BGP AS path (case 1)
B
C
X
A
(a)
B
C
X
D
(b)
B
C
A
(c)
B
C
D
(d)
Coordinate BGP AS path (case 2)
A
C
B
Coordinate BGP AS path (case 3)
IXP
A
B
C
D
Coordinate BGP AS path (case 4)

Figure 3: Correspondence between traceroute AS path and BGP AS path, dotted arrows are traceroute AS paths and solid arrows are the corresponding BGP AS paths

and [...BXC...] in BGP AS path; 2) X is a customer or sibling of B who uses the IP addresses from B but does not announce them publicly, so the traceroute probing responses from X are still falsely mapped into B which causes [...BC...] in traceroute AS path and [...BXC...] in BGP AS path; 3) BGP border router in X uses outgoing interface for ICMP causing X's missing; 4) B-C is a true link somewhere that has never been observed by the BGP monitors, and this produces false negatives in our results.

Substitute hop and Extra hop: If the DISTANCE is 2 and we could not find any corresponding route like [...BXC...] in the BGP AS paths, it may either be due to the incompleteness of BGP AS paths we collected from publicly available VPs or because AS path [...BXC...] is invalid and not existing in BGP. The latter one could result from the following scenarios: 1) X is multihomed to its providers B and C and used IP addresses from its provider (B or C) to configure its machines but does not announce them publicly (case 2a/2b in Figure 3), this would produce a traceroute AS path of [...ABC...] (or [...BCD...]) while its corresponding BGP AS path as [...AXC...] (or [...BXD...]). This causes the substitute hop problem; 2) Sometimes, X could be A (case 2c/2d in Figure 3). This is because in addition to the unannounced addresses from its provider, AS A owns and announces some other addresses. This causes a potential problem when traceroutes go through one part to another part within an AS, its publicly mapped AS path would have an inter AS link. For example, in case 2c/2d of Figure 3, while traceroute AS path is [...ABC...] (or [...BCD...]) its underlying BGP AS path is [...AC...] (or [...BD...]) which is the extra hop problem; 3) It is known that sibling ASes would also cause substitute/extra hop problem. 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 (***Step-6.11*** $\sim$ ***Step-6.13***). Similarly, our conservative method risks in discarding true links here.

Another case that would cause one substitute or extra hop in traceroute AS path is due to the use of outgoing interface to reply to ICMP message. In case 3 of Figure 3, AS A's last hop router uses its outgoing interface towards C to reply to an ICMP message (connection between A and B is using addresses from B), which cause [...ABC...] in traceroute AS path and a corresponding [...AC...] in BGP AS path. Again, ***Step-6.11*** and ***Step-6.13*** can address the false substitute/extra middle hop.

Special case of extra hop: One scenario for the traceroute AS link that has a DISTANCE $\geq$ 3 between B and C is case 4 of Figure 3. AS C is an IXP having its own AS number, IXP C announce its addresses through some particular participant, say AS E, which is not B or D. If E is not a neighbor of B, this shows up B and C as being at least 3 hops away. Our heuristics take care of this scenario in ***Step-6.17***.

3.2.3 Validation

After applying all the heuristics, we have harvested around 100,000 AS links through our P2P traceroutes. We extract all the links with a tier-1 AS (the number is on the order of thousands[4]) and validate these links with the ground-truth of this AS. We are excited to find that all our links are in the ground-truth. Our measurements have not covered all the links with this tier-1 AS. However, we are not claiming to get all the links, instead we pay more attention to the quality of our found links and extending the public view.

Validating with the tier-1 network, we inspect how our heuristics help to filter the false links caused by the IP-to-AS mapping procedure. We focus on our proposed advanced heuristics. Before applying these advanced heuristics, our P2P traceroutes find thousands of AS links of this tier-1 AS when we are done with the first five steps. Compared with the AS's ground-truth connectivities, we are surprised to find that 48.8% of links we extracted from the publicly mapped AS paths are false. Using the tier-1 ground-truth as baseline, in Table 3, we calibrate the number of false links filtered and modified by each of our advanced heuristics.

Confirmation on line 8: In our measurements, we see several hundreds of unique cases where [...T_1, C...] is in our traceroute AS paths while [...T_1, X, C...] are in BGP AS paths. Checking with the router configuration files of the tier-1 net-

[4] Note that the exact number of AS links have been concealed for non-disclosure reasons; Furthermore, we use percentages across the section.

Line No. in Algorithm 1	False links left
-	48.80%
8	10.47%
11	5.13%
13	0.47%
18	0

Table 3: Filtering false links with each of our advanced heuristics.

work, we found that, in 94% of the cases, the last IP hop that publicly mapped to T_1 is actually belonging to a third AS X. These false links may happen due to the private peering issue or unannounced IP addresses. This evidence increases our confidence of modification on line 8 of our algorithm which is adding an extra hop. Another interesting observation is that we did not find any of these T_1-C links to be true according to the ground-truth. Promisingly, this particular heuristic successfully filtered 38.33% false links.

Confirmation on line 11 and 13: We have hundreds of cases with this tier-1 AS where [...A, T_1, C...] (or [...B, T_1, D...]) is in our traceroute AS paths. Using the ground-truth data we did the following validation: we first returned back to the corresponding IP level paths and extracted the IPs that mapped T_1, then we searched these IPs in the router configuration files to see if they are used to configure real routers of the tier-1 network or not. In 93% of the cases, we found that these IPs are not used in by this tier-1 network. This tells us that these IPs should be allocated to the tier 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, finding such real cases indicates that our heuristics accurately analyze the root causes for wrong mappings and so we can take proper way to modify the false links. According to Table 3, the two lines help us rectify 5.34% and 4.66% false links.

Observation on line 18: We have no specific ground-truth files that can help us validate our heuristic here. However, the tier-1 network connectivities can give us estimation on how successful this line would be to delete false links. Our 0.47% false links in this scenario is completely filtered.

3.3 Policy Inference

After extracting the AS links, we infer the business relationships between ASes based on the PTE algorithm proposed by Xia [19]. After improving the seminal work by Gao [20], 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 [18] 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 used 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 4. In this section, we use our set of missing links to determine the public view's coverage of each class of Internet AS links and where these links are missed.

4.1 Coverage of tier-1 AS links

We begin by focusing on the tier-1 AS connectivities, listed in Table 5. It is worthwhile to note that, although we have uncovered 23,914 missing links, we have not been able to see many new additional tier-1 AS links: 1) no new link is found for three tier-1 networks AT&T, Sprint, and SAVVIS; 2) a very small percentage (up to 3.14%) of new links are found for the rest tier-1 networks. Note that our finding is confirming the observation in [2] that, in general, tier-1 AS links are covered "fairly complete" by public view over time. On the other hand, our results also disclose that public view still miss some tier-1 links although they have monitors in these tier-1 networks. This is because: 1) the current public view has very few feeds (i.e., peered routers) in each AS, and a tier-1 AS could contain thousands of routers and each router potentially has its constrained view; 2) some tier-1 ISPs have their rules not to announce long prefixes (*e.g.*, longer than /24) which cause public view lose to see such tier-1 links.

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 public view (second column), number of new links in P2P traceroutes (third column), and the percentage (fourth column).

4.2 Coverage of customer-provider links

Table 4 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. Our results indicate that these public views are not as complete as previously suggested.

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 4: Newly found links under conservative extraction. (PV stands for public view; New # is the number of links missing from PV)

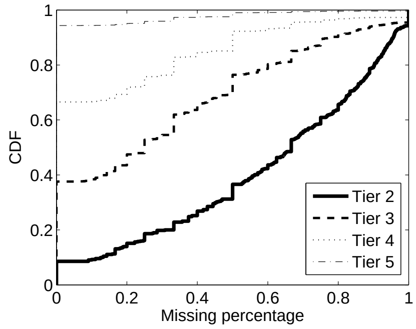

Figure 4: Missing percentage for the missing provider links of each AS on each tier.

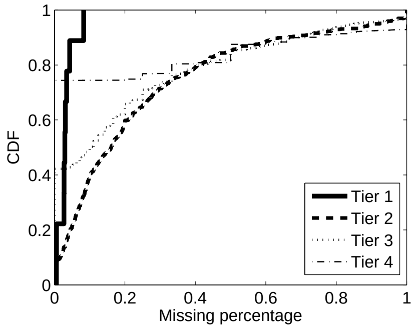

Figure 5: Missing percentage for the missing customer links of each AS on each tier.

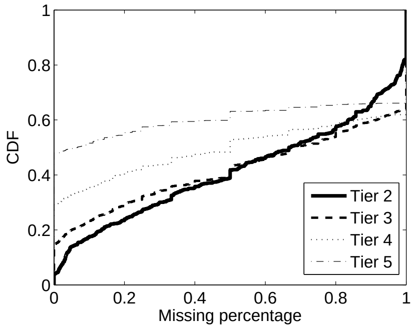

Figure 6: Missing percentage for the missing peering links of each AS on each tier.

We categorize the missing links according to their business relationships: the percentage of missing provider links and the percentage of missing customer links[5]. We use our classification from Section 3.3 to categorize each AS into a tier, then group all of the missing rates for each tier. Figure 4 and Figure 5 show the CDF for missing provider links and missing customer links, where each point (x, y) means that y percent of ASes in the tier have at most x percent of their links missing from the public view. Note that tier-1 ASes have no providers and tier-5 ASes have no customers and thus are omitted from the corresponding figures. The figure clearly shows that *customer-provider* links can be missed in every tier. Another interesting observation from our dataset is that the missing percentages for provider links somewhat correlate with tiers: the higher the tier of an AS, the higher possibility to 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 tier 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 links are expected to appear at lower tiers of the Internet hierarchy, where there is less coverage from BGP feeds. However, we also find that a significant number of *peering* links are missing from public view at higher tiers. Similar to the previous section, we calculate missing rate for *peering* links and plot the CDF in Figure 6. The graph shows that high tier networks have relatively higher missing rates of missing links than lower tier networks. We will investigate the reasons behind these unexpected missing links in Section 5.

4.4 Missing sibling AS links

We revealed 216 additional *sibling* links which are missing from public view. We believe that one reason behind these missing *sibling* links could be due to route propagation in BGP. To illustrate this fact, consider two *sibling* ASes: AS1 and AS2. During BGP route propagation, 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. To address this challenge, we observe that a missing 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 for a missing link, it is useful to first determine where it occurs with respect to the VPs of the public view. Tracing from VPs to each missing link, we get the eight corresponding patterns in Figure 7. For instance, if a link AS_x-AS_y is missing, then we check the public view for any path containing AS_x or AS_y. For each of these AS paths, we then record the pattern from the VP to the associated AS. For simplicity, we degenerate a continuous series of *customer-to-provider* (or *provider-to-customer*) links into one logical *customer-to-provider* (or *provider-to-customer*) link and this

[5] The next section contains a more complete analysis of the root causes for these missing links.

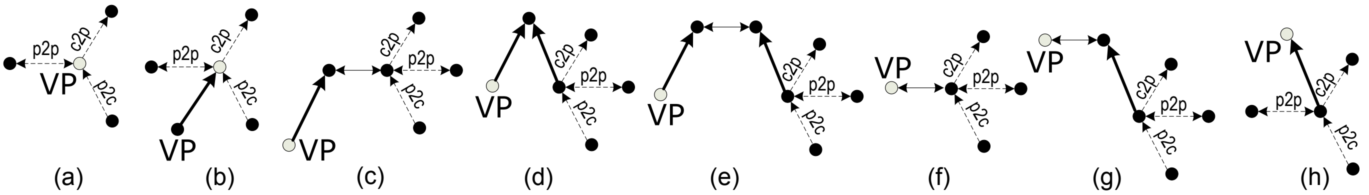

Figure 7: Eight patterns for all the links (visible/missing) relative to the VPs. A bold arrow represents a *customer-to-provider* link or a combination of *customer-to-provider* links; a bidirectional (thin) arrow represents only one *peering* link; A dotted arrow represents a missing link; The sibling links have been degenerated. To understand the graph, for example, pattern (b) means there are c2p, p2p, and p2c missing links at the places when going from a VP and traveling along one (or multiple) *customer-to-provider* links.

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
Missing percentage	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 missing pattern.

makes no difference on our analysis. 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. As we prove next, this list is exhaustive.

5.1.1 *Completeness Proof*

We now prove that there are eight patterns of relative position for a link in terms of a VP under valley-free routing.

Theorem: The eight relative position patterns illustrated in Figure 7 are exhaustive.

Proof: Based on export policy settings, a *provider-to-customer* ($p2c$) or a *peering* ($p2p$) edge can only be followed by *provider-to-customer* or *sibling* ($s2s$) edges [20]. This can be used to define a valid AS path 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.

5.1.2 *Observations*

Table 6 shows the number of *customer-provider* and *peering* 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 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 table can discover all the connections of its upstream providers [2, 3]. However, from our measurement, we found that even 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. For example, in our measurement, we found the first 100 full table VPs have missed 1096 adjacencies of themselves.
- 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 such as pattern (h) in Table 6 confirms this fact. However, we also find that there are a lot of instances of upstream *peering* links being invisible to downstream full table monitors (for instance, pattern (b)). This tells us that low hierarchy is not the only reason to blame for the missing peering links.

5.2 Identifying root causes

We now further categorize the eight missing patterns into two categories: missing under valley and missing under non-valley (shown in Table 7). This section discusses our findings and focuses on the latter category. While the former reason has been used to explain the missing links, the latter one has received little attention in the previous literatures. In particular, we determine 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 our list of root causes is exhaustive, we believe it accounts for most missing links.

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

Relationship	Missing under non-valley	Missing under valley
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 7: Patterns for links missing under valley-free and valley patterns.

an AS could contain hundreds of routers while different routers may potentially have different routes even for a same prefix [21, 22]. This leads to the problem that the public view data could miss a lot of AS links and even the directly connected links of vantage point ASes. What's more, 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.

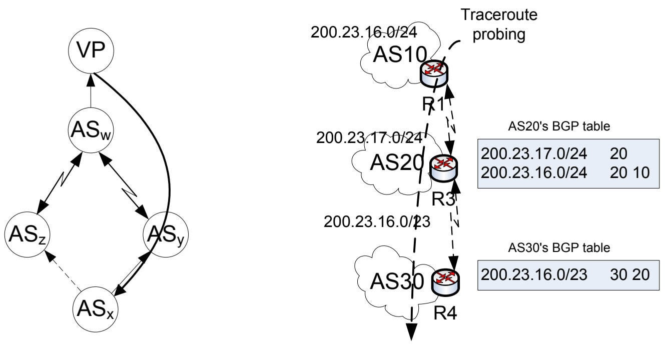

200.23.17.0/24	20
200.23.16.0/24	20 10

200.23.16.0/23	30 20

(a) Sub-optimal path

(b) Route aggregation

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

5.2.2 *Route aggregation*

BGP uses prefix aggregation to reduce the size of routing tables. This is a process that combines the characteristics of several different routes in such a way that advertisement of a single route is possible. 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, aggregation may omit AS_PATH information for individual prefixes, causing 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 public view. However, Cymru [17] has access to private BGP feeds that may contain ASNs not in public view, which helps us successfully uncover 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 them 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 [23]. 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 routing policy*

In general, 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 policy settings [6, 20]: while exporting to a *provider* or *peer*, an AS will export the routes from its *customer* network and from itself, but usually will not export its *provider* or *peer* routes; while exporting to a *customer* or *sibling*, an AS will export its routes and its *customer* routes, as well as its *provider* or *peer* routes, *i.e.*, its full table. These exporting rules indicate that AS path should be valley-free, i.e., after a *provider-to-customer* link or a *peering* link, the AS path cannot traverse *customer-to-provider* links or another *peering* link.

The valley-free routing policy is well known and widely used to explain the missing links, especially the missing *peering* links [2, 3, 5, 6]. We have also used valley-free routing to explain many instances of the missing links; however, in contrast with most previous work, while we find that valley-free routing policy can cause missing *peering* links, we observe a large number of missing *customer-provider* links on the valley-free AS path to VPs, as shown in Table 6. This is because the existing public view only has one or two BGP monitors in each vantage point AS, hence greatly restricts its visibility about the global view. With the help of large-scale P2P traceroutes, we can reveal some of these missing *customer-provider* links (Recall that they are classified as on the sub-optimal paths to VPs in our analysis). In Section 5.3, we will further quantify the impact of valley-free routing policy on missing links. Below we will see a special case of this.

Partially cooperative VPs: It would seem counterintuitive that the VPs cannot see the direct *peering* links and *customer-provider* links for their ASes; we now suggest four possible reasons why this happens. 1. Contrary to common belief, 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 because they export only a selective portion of their routing tables; 2. Before exporting to the central route collectors, the VPs may have already aggregated some neighboring ASes, which causes the public view to miss these directly connected links; 3. The default routing configuration of the peering routers of some VPs prevent them from reporting the default connectivity to the central collector, and hence cause directly connected missing links; 4. The restricted view from one or two monitors placed in the VPs cannot uncover all the connectivity.

We focus on the first reason; the last three reasons are due to route aggregation and sub-optimal paths to VPs, which was covered in the previous section. Our heuristic for testing this hypothesis is that vantage points in this category should not export any other *peering* link or *customer-to-provider* link to route collectors. Our experimental results show that, for the 344 vantage points that miss at least one *peering* link or *customer-to-provider* link, the public view has not seen any direct *peering* or *customer-to-provider* link from 148 ($148/344 = 43\%$) VPs, yielding a number of 2116 links missing due to this reason. To validate this result, we contact a Routeviews administrator [8] who confirmed our findings [24]. While the Routeviews administrators request all of their peered VPs to treat them as a "*customer*" and thus export their entire routing tables, not all the participating ISPs comply for policy reasons. Instead, some neighbors could treat each route collector as a "*peer*" and selectively export partial information from their routing tables.

5.3 Categorizing the Missing Links

The previous section 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 fine-grained link classification. When categorizing missing links in this way, there could be more than one plausible reason for them to be missing. For instance, when we manually checked a set of links, we observed that they were explained as a valley-free path with respect to one VP and a non valley-free path with respect to a different VP. In this section, we focus on the identification and quantification of the three main root causes: (a) valley-free routing policy, (b) route aggregation, (c) sub-optimal paths. Though this is not an exhaustive list, we believe that a combination of these three root causes explains most of the missing links.

		Partially cooperative VPs	**Completely hidden ASes**	**Default routing**	**Route aggregation**	**Sub-optimal path to VP**	**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 links (c2p, p2p, and p2c) under each kind of missing patterns (from pattern (a) to pattern (h)) in Figure 7. The first three reasons are special cases, while the last three reasons are the main root causes in our analysis. To understand the table, for instance, the first row means that 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.

Our heuristic for determining the root causes for missing links is shown in Algorithm 2. The algorithm essentially does the following:

- When the link is found to be on a valley path to a VP, it is classified as missing under valley-free routing policy.
- 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 attribute 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 61 for completely hidden ASes.

Root cause	$\{a\}$	$\{b\}$	$\{c\}$	$\{d\}$	$\{a,b\}$	$\{a,c\}$	$\{b,c\}$	$\{a,b,c\}$	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: a - valley-free routing policy, b - route aggregation, c - sub-optimal paths, d - missing sibling links, "unknown" is because relationships of these link have not yet been determined. To understand the table, for instance, $\{a\}$ means 1.38% of the missing links are solely due to valley-free routing policy; $\{a,b\}$ means 0.26% are exactly due to both valley-free routing policy and route aggregation; $\{a,b,c\}$ means 75.02% are due to all these three reasons simultaneously.

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 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
See notations in Table 9; Initialization: $f_3(m_i, v_j, p_k) = \Phi$, $f_2(m_i, v_j) = \Phi$, and $f_1(m_i) = \Phi$;
foreach missing link $m_i \in \mathbb{M}$ do
    foreach VP of public view $v_j \in \mathbb{V}$ do
        foreach missing pattern $p_k \in \mathbb{P}$ do
            if $\exists$ one AS attached to $m_i$ that is not visible to $v_j$
            then
                if $valley(m_i, v_j, p_k) = 1$ then
                    $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "$(a)$";
                else
                    $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "$(b)$";
                end
            else
                if both ASes attached to $m_i$ are visible to $v_j$ then
                    foreach node attached to missing link $m_i$ do
                        if $valley(m_i, v_j, p_k) = 1$ then
                            $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "$(a)$";
                        else
                            $f_3(m_i, v_j, p_k) := f_3(m_i, v_j, p_k) \cup$ "$(c)$";
                        end
                    end
                end
            end
        $f_2(m_i, v_j) := \bigcup_{p_k \in \mathbb{P}} f_3(m_i, v_j, p_k)$;
    end
    $f_1(m_i) := \bigcup_{v_j \in \mathbb{V}} f_2(m_i, v_j)$;
end
Return $f_1(m_i)$: reasons for missing link $m_i$;
```

Algorithm 2: Assigning reasons to missing links.

- *BGP policies have a significant effect*: A significant number of links are missing due to valley-free routing policy and sub-optimal paths. 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 under hundreds of the public view VPs. For instance, 1.38% of the missing links are classified as valley-free routing policy, 0.33% as route aggregation, and 0.27% as sub-optimal paths. However, there are 75.02% of the links are missed because all the three causes occur simultaneously.

6. DISCUSSION

In this paper we showed that using P2P traceroutes reveals a significant number of missing AS links; namely, 12.86% more *customer-provider* links and 40.99% *peering* links are missing from 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 IP-to-AS mappings gathered from publicly available information are neither 100% complete nor accurate. This is a limitation of all work using traceroutes to extend the AS topology. However, using a tier-1 AS's ground-truth as baseline, we have validated our results related to this AS and demonstrated that our proposed heuristics can filter all the false links efficiently. This greatly increases our confidence about our measurements results for other ASes, though we admit that this is not a sufficient translation condition. Furthermore, the AS relationship inference algorithm is not 100% accurate either, and this can potentially influence the accuracy of classification of newly discovered links and root causes. In addition, 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. Our approach mitigates this issue through a comprehensive set of heuristics. It is possible, however, that other unidentified issues affect our measurements.

7. RELATED WORK

The Internet's connectivity structure is defined by ISP interactions via the BGP, through which AS paths are advertised for the purpose of 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 the missing links. However, all of them are related to passive data from BGP public VPs. The topologies are inherently limited to the information made available by these VPs.

Measurement platforms, such as Skitter [12], DIMES [11], and iPlane [10] are providing views about the Internet structure from active measurements. Unfortunately, they are either limited to the scale or do not provide serious and accurate AS level topology. In addition, Samantha et al. [25] used active measurements to expose hidden prepending policies and hidden ASes but their work concentrated more on BGP routing dynamics than the AS topology. More recently, Shavitt et al. [26] studied the importance of vantage points distribution in Internet topology measurements, however they had no guarantee or estimation on 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. We have the largest scale of probing infrastructure and detailed heuristics to convert IP level paths to AS level graph. We identify where these links are missing and why they occur.

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 over 580,000 machines in 6,000 ASes distributed across the whole Internet to probe its topology and a series of heuristics to modify the false links, we have uncovered 23,914 new links hidden from the public view. While we confirmed that tier-1 ASes connectivity is covered fairly well by the public view, our results also indicated that: 1) the public view can miss a substantial number of *customer-provider* links; 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 proposed an algorithm for classifying them into a number of root causes. We 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 the more complete AS topology affects other commonly held beliefs about Internet properties such as caching and resiliency. Finally, to facility other research in this area, we will make our missing links publicly available to the community.

9. REFERENCES

[1] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," in *Computer Netw.*, 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 *ACM SIGMETRICS*, 2008.

[3] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "Quantifying the Completeness of the Observed Internet AS-level Structure, Tech. Rep. 080026, 2008.

[4] R. Oliveira, B. Zhang, and L. Zhang, "Observing the Evolution of Internet AS Topology," in *ACM SIGCOMM*, 2007.

[5] R. Cohen and D. Raz, "The Internet Dark Matter: on the Missing Links in the AS Connectivity Map," in *IEEE Infocom*, 2006.

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

[7] Z. M. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an Accurate AS-Level Traceroute Tool," in *ACM SIGCOMM*, 2003.

[8] ROUTEVIEWS, "Routeviews project," http://www.routeviews.org/.

[9] RIPE, "Routing Information Service," http://www.ripe.net/projects/ris/.

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

[11] Y. Shavitt and E. Shir, "DIMES: Let the Internet measure itself," in *ACM SIGCOMM CCR*, 2005.

[12] CAIDA, "Skitter AS adjacency list," http://www.caida.org/tools/measurement/skitter/as_adjacencies.xml.

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

[14] "Internet Routing Register," http://www.irr.net.

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

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

[17] Cymru, "IP-to-AS mapping," http://www.team-cymru.org/Services/ip-to-asn.html.

[18] "The CAIDA AS Relationships Dataset, 2004-2008," http://www.caida.org/data/active/as-relationships/.

[19] J. Xia, "On the Evaluation of AS Relationship Inferences," in *In IEEE GLOBECOM*, 2004.

[20] L. Gao, "On inferring Autonomous System relationships in the Internet," *IEEE/ACM Trans. Netw.*, 2001.

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

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

[23] "NANOG mailing list," http://www.merit.edu/mail.archives/nanog/.

[24] Authors and Administrator, "Question about Routeviews," in *Email communication between administrator of Routeview project*, 2008.

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

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