

Quantifying the Pitfalls of Traceroute in AS Connectivity Inference

Yu Zhang¹, Ricardo Oliveira², Hongli Zhang¹, and Lixia Zhang² *

¹ Harbin Institute of Technology, Harbin, 150001, China,
{yuzhang,zhanghongli}@hit.edu.cn

² University of California, Los Angeles, CA 90024, USA,
{rveloso,lixia}@cs.ucla.edu

Abstract. Although traceroute has the potential to discover AS links that are invisible to existing BGP monitors, it is well known that the common approach for mapping router IP address to AS number (IP2AS) based on the *longest prefix matching* is highly error-prone. In this paper we conduct a systematic investigation into the potential errors of the IP2AS mapping for AS topology inference. In comparing traceroute-derived AS paths and BGP AS paths, we take a novel approach of identifying mismatch fragments between each path pair. We then identify the origin and cause of each mismatch with a systematic set of tests based on publicly available data sets. Our results show that about 60% of mismatches are due to IP address sharing between peering BGP routers in neighboring ASes, and only about 14% of the mismatches are caused by the presence of IXPs, siblings, or prefixes with multiple origin ASes. This result helps clarify an argument that comes from previous work regarding the major cause of errors in converting traceroute paths to AS paths. Our results also show that between 16% and 47% of AS adjacencies in two public repositories for traceroute-derived topology are false.

Key words: AS topology measurement, traceroute, BGP

1 Introduction

The Internet is a vast distributed system formed by a myriad of networks called Autonomous Systems (ASes) that exchange routing information using the Border Gateway Protocol (BGP). There have been two basic approaches to measuring AS-level connectivity: (1) passive measurement through collecting BGP routing updates, and (2) active measurement using traceroute. In the BGP-based measurement, AS adjacencies can be directly extracted from the ASPATH attribute in BGP updates collected from the monitors/routers by Routeviews [4] and RIPE-RIS [3]. But because of policy filters and best path selection, each BGP monitor only provides a limited partial view of the topology. Most monitors in

* This work was partially supported by the National Basic Research Program of China (973 Program) under grant No.2005CB321806 and by the US National Science Foundation under Contract No CNS-0551736.

traceroute measurement projects, such as CAIDA’s Ark [1] and DIMES [16], are placed in different ASes than BGP monitors, thus ideally they can complement the topology inferred from existing BGP sets. It is also easier to deploy a traceroute monitor than to obtain a new BGP feed [16].

However converting the router IP addresses on traceroute paths to AS numbers, termed *IP2AS mapping*, is a difficult problem. Typically this conversion is done by finding the origin AS of each IP address in the traceroute path from the BGP routing table using *longest prefix matching* (LPM). Unfortunately this approach is known to generate potentially false AS links, and the following question emerges: *what’s the impact of inference errors of traceroute-derived AS paths on the AS topology map when using LPM?*

Several previous efforts have studied the problem of traceroute-derived measurement and articulated possible causes for the mismatch between the traceroute-derived path and the BGP path [7, 13, 9, 10, 12]. However, these previous efforts did not provide answers to our question because of the following reasons: (1) They quantified mismatch causes in the unit of *path*, *e.g.* either there was a match in the converted path or not, which does not pin down all individual points on the topology that the two paths differ; and (2) They did not investigate the accuracy of traceroute-derived topology.

In this paper we conduct a systematic and exhaustive investigation into the impact of pitfalls of LPM-based traceroute measurement on topology inference. Our contributions can be summarized as follows. (1) We identify differences in pairs of traceroute and BGP paths systemically. This allows us to pinpoint multiple mismatches in the same AS path pair and to identify each mismatch point shared by multiple path pairs. (2) We collect a comprehensive set of publicly available data and develop a set of tests to infer the cause of each mismatch more systematically than before. (3) Our results show that about 60% of mismatches occur because of IP address sharing between neighbor routers. This result is a departure from previous work [13, 10] that attributed the causes of errors mainly to Internet eXchange Points (IXPs), sibling ASes under the same ownership, and prefixes originated from multiple origin ASes. (4) We find that between 16% and 47% of the traceroute-derived adjacencies in public data sets widely used by the community may be bogus.

2 BGP vs. traceroute

Generally speaking, the data path inferred from traceroute and BGP control path should match. There are however some scenarios where the two paths differ, either because the data path is not completely aligned with the control path, or because of IP2AS shortfalls in converting IP addresses to AS numbers. We describe different reasons why the BGP AS paths may differ from the AS paths measured by LPM-based traceroute method.

(1) There may be divergence between data path and control path due to BGP aggregation, multi-hop sessions, tunneling, layer-2 switching, and abnormal routing. (2) The traceroute path may be incomplete because of non-responsive

hops. In addition, the BGP routing tables may not tell exactly the original ASN of a given prefix, *e.g.* (3) an unannounced prefix or (4) a Multiple Origin ASes (MOAS) prefix.

The IP addresses announced by a given AS X may be used by another AS Y . We call those addresses the *foreign addresses* of Y . (5) A typical case is that one prefix is shared by multiple participants in the IXP, which is a shared infrastructure where multiple networks peer with each other publicly. (6) ASes under the same ownership, *i.e.* siblings, may also share the same IP address space. (7) Another typical case is *IP address sharing between neighbor ASes*, where a border router owned by AS Y replies to a traceroute using one of its interfaces whose IP address is borrowed from the neighboring AS X to enable the point-to-point connection. For example, when Y has a private peering with X , two incident routers' interfaces are typically numbered from a /30. If the /30 is coming from X , then routers at Y may reply with X 's address range.

According to our measurements, 63~88% of path pairs had a match (no extra links in traceroute path). In the remaining cases, at most 3.7% of mismatch path pairs are originated by divergence of control path and traceroute path, and for the rest we provide evidence for their occurrence due to errors in IP2AS mapping. Therefore, we believe it makes sense to use BGP paths as the reference by default as [13, 10, 12]. If in the vast majority of cases the data path would not be align with the control path (or BGP), there would be a significant number of mismatch cases we could not explain, which is not the case.

3 Related Work

Measuring Internet AS-level connectivity from traceroute data has attracted many research efforts over recent years. One of the first such studies was done by Chang *et al.* [7], which alerted for possible errors in AS topology inferred from traceroute data using the LPM approach. They presented a technique to identify the ownership of border routers based on IP alias resolution, and presented some heuristics to fill the holes of unmapped hops in traceroute paths. This work was probably the first that pointed out potential errors in traceroute-derived AS paths because of IP address sharing between neighbor ASes.

In a later work, aiming at an accurate AS-level traceroute tool, Mao *et al.* [13] compared BGP paths with traceroute paths launched from the same AS where the BGP table was extracted. They investigated a comprehensive set of possible causes of mismatch and developed heuristics to correct the IP2AS mapping. In a following work [12], they presented a dynamic programming algorithm to reassign /24 prefixes to ASes to minimize the number of mismatched path pairs. The main outcome of this work was a method to correct the mismatches due to unmapped hops, MOAS prefixes, IXPs and siblings.

At the same time, Hyun *et al.* [10] presented a path pair comparison and quantified the mismatched pairs due to IXPs and siblings. They adopted the algorithm for the longest common subsequence (LCS) problem to describe the pattern of unexplained mismatches. Later [9], they presented the concept of *third-*

party addresses, but its definition does not clearly address the issue of IP address sharing between BGP neighbors.

As far as we can tell, our paper is the first to propose a systematic method of identifying mismatches between each traceroute-derived path and BGP path pair. This method allows us to pinpoint multiple mismatch points in the same AS path pair and align mismatched portions of a pair of paths, giving local context to the comparison and explaining the cause of the mismatches. Our result asserts that the main cause of mismatch is the IP address sharing between neighbor ASes in accordance with [7] and departing from [13, 10] that attributed the mismatches to the presence of IXPs, siblings and MOAS. Note that we do not use the LCS algorithm to describe the unexplained mismatches (as [10]), but we enhanced it to identify the mismatches.

4 Data Sets

4.1 AS Path Pair Data

We collect traceroute raw data and the corresponding BGP routing updates from 4 ASes. Table 1 lists the number of destination IP addresses and prefixes probed by the traceroute vantage points, as well the corresponding BGP information.

UCLA: From a host located at UCLA, we performed probes targeting all /24 blocks in the BGP routing table, using the traceroute tool *scamper* (<http://www.wand.net.nz/scamper/>) with ICMP-paris [6]. At the same time, we collected BGP updates and tables from a backbone router at UCLA.

CAIDA Ark: There are 3 CAIDA Ark monitors that happen to be located in ASes which provide a BGP feed to either RouteViews or RIPE-RIS collectors. For each /24 block, the latest traceroute result is picked.

Table 1. Information of AS path pair data sources.

Monitor	ASN	#pair	#prefix	Collector	Organization	Date
ucla	52	7.6M	272K	ucla	UCLA	2009-02-22~03-10
ams-nl	1103	5.2M	218K	ris-rrc03	SURFnet	2009-02-01~03-12
nrl-jp	7660	4.9M	212K	rv2-oix	APAN	2009-02-01~03-12
she-cn	4538	5.2M	218K	rv-wide	CERNET	2009-02-01~03-12

The traceroute AS paths are generated by the LPM-based IP2AS mapping on the BGP routing table of the AS where the traceroute is launched from. Since there is no guarantee that BGP routers will have consistent tables inside a large AS (*e.g.* different routers in same AS can have different tables), we only collect the path pairs where the next-hop AS is the same in order to reduce ambiguities. A traceroute path is paired with its corresponding BGP path to the same prefix, only if there is no change observed in the local BGP route to the destination prefix during the traceroute probe, otherwise the paths are discarded.

4.2 AS Adjacencies

To evaluate the accuracy of traceroute-derived AS adjacencies, we collect data from CAIDA Ark, DIMES, and UCLA IRL [5]. The data from IRL is also used to explain the mismatch with the assistance of data from Internet Routing Registry (IRR) [2] and iPlane [11].

CAIDA Ark: Two traceroute-derived AS topologies are obtained by merging all snapshots in Feb. 2009 [1]: (1) The topology with only direct links, in which every consecutive pair of ASes have a pair of contiguous hops in the traceroute path; and (2) The topology with both direct links and indirect links, in which two IP addresses in different ASes may be separated by one or more unmapped or non-responsive hops.

DIMES: We also collect DIMES' monthly traceroute-derived AS topology in Feb. 2009 [16]. This graph include the AS links which are observed at least once in the given month and at least twice considering all period.

BGP: We use the BGP-derived AS adjacencies available at UCLA IRL [5], which is extracted from RouteViews and RIPE-RIS. For the sake of completeness, the data is accumulated over a period of 5 months ending at March 2009, following the methodology in [15].

IRR: The Internet Routing Registry (IRR) [2] is a central repository where ISPs explicitly insert information such as routing policies and BGP adjacencies. We are able to extract 28,700 total AS numbers and 156,094 total AS adjacencies from all available IRR databases as of 2009-03-05.

iPlane: iPlane [11] project provides a list of routers' alias, *i.e.* a set of interface IP addresses belonging to the same router. This information can be used to explain mismatches due to IP address sharing between BGP neighbors, since we can look up each interface alias in BGP tables and estimate which ASes have BGP sessions in a same router. We extracted a total of 286,043 IP interface addresses on 67,430 routers on 2009-03-05.

4.3 IXP and Sibling Lists

To help identify ASNs used by IXPs and ASes with sibling relationship, we extracted the name/description of each AS from all WHOIS databases.

IXPs: We compiled a list of 404 /24 prefixes belonging to IXPs by crawling three websites, peeringDB.com, PCH.net and euro-ix.net, on 2009-03-09. Additionally, we search a list of ASNs associated with IXP names (from the previous websites) and the common words "internet exchange", "exchange point", "access point" and "gigapop", carefully filtering the false IXP records, *e.g.* a description "peering at an IXP". We end up with a total of 323 ASNs belonging to IXPs.

Siblings: We look for similarities in AS names/descriptions of a given pair of ASes using approximate string matching except in the cases where the name is a word appearing in an English dictionary. The acquisition history of all Tier-1 ISPs from wikipedia is also used to group ASes. After computing the transitive closure of sibling relationships and cleaning up the candidate sibling groups with size greater than 20 manually, we get 3,490 sibling groups with 13,639 ASes.

5 Mismatch Analysis: Breaking Paths into Fragments

In this section we develop a technique for comparing BGP paths and traceroute-derived AS paths obtained in Sec. 4.1. We use the classic file comparison command `diff`-like method to find the *longest common subsequence* (LCS) that is present in both traceroute path and BGP path [8]. The LCS solution is described as a minimum array of binary operations needed to transform the BGP AS path into the traceroute AS path: insertion ‘+’, deletion ‘-’, or unmodified ‘=’. The consecutive ‘=’ operations represent the common segments, while the ‘+’ and ‘-’ operations indicate the difference.

To pinpoint multiple mismatches in the same AS path pair and describe a mismatch in its local context, we define a *mismatch fragment* between two AS paths as a sequence of ‘-’ and/or ‘+’ wrapped around by two ‘=’. For example, Figure 1(a) shows the solution (F1) and mismatch fragments (F2) of a one-to-one substitution case. Note that the same mismatch fragment at the AS level may have the different IP-level fragment. We develop five additional steps to detect the mismatch fragment systematically:

	(a) substitute	(b) end-extra	(c) non-responsive	(d) tie-break	(e) loop	(f) end-substitute
BGP:	A B C D	A B	A B C	A B C D	A B	A B C
Traceroute:	A E C D	A B C	A * C	A C B D	A C A B	A D
F1:	=A -B +E =C =D	=B +C =	=A -B ** =C	=A +C =B, =B -C =D	=A +C +A =B	=A -B -C +D =
F2:	=A -B +E =C			=A -B =C, =C +B =D	=A +C =A	=A -B +D =

Fig. 1. Examples of AS path pairs and their mismatch fragments

1. We add 4 special tokens to the traceroute AS path: ‘*’ representing consecutive non-responsive hops, ‘?’ representing consecutive unmapped hops; ‘^’ and ‘\$’ represent the beginning and end of a path, respectively. See examples in Figure 1(b) and (c).
2. When there are multiple alternative solutions with the same number of operations, the one whose ‘=’ operator appears earlier in the BGP path is picked. The goal of this tie-break is to concentrate the errors in the least number of original hops as possible. In Figure 1(d), F1 is picked from two solutions.
3. The mismatch fragment is replaced with its inside loop, since loops describe differences more properly as shown in Figure 1(e).
4. Among more than one ‘-’ operations in the substitution at the end of path, only the first is kept, *e.g.* F1 is replaced with F2 in Figure 1(f).
5. Mismatch fragments whose modifying operations include only ‘+*’ or ‘+?’ or only deletions (‘-’) at the end of path are discarded, because our interest is in extra links brought by traceroute.

We obtained a total of 39K unique mismatch fragments (15~20K per monitor) including 44% *extra* (‘+’-only), 20% *missing* (‘-’-only), and 36% *substitute* (both ‘-’ and ‘+’) patterns. Among the *extra* mismatch fragments, 39% are

loops. Overall, there are 12~37% of path pairs containing one or more mismatch fragments. And we also observe that the appearance frequency of mismatch fragments follows a heavy-tailed distribution, which means that there are a small number of mismatch fragments shared by a large number of the path pairs.

6 Inferring the Causes of Mismatch

In this section we look into causes of mismatch between BGP paths and traceroute derived paths, and classify them into 7 types as described in Sec. 2. Our classification algorithm follows an if-then-else process, *i.e.* it starts by checking whether the mismatch is of type 1, if yes then the classification stops, otherwise it continues and checks for type 2, and so on until it's put in *Unknown* bin.

1. **Divergence:** These are cases where the control plane is not aligned with the traceroute path. We detect these cases whenever the edit distance of the mismatch fragment (*i.e.* the number of modifying operations), has a high value, *i.e.* greater than 3 for *substitute* and *loop* patterns, or greater than 2 for *missing* and *extra* patterns. Below are two examples we find in our data:
 - (a) **Tunneling:** The Amateur Packet Radio Network (AMPR.org) uses the prefix 44/8 announced by AS7377 (UCSD) and an overlay network on the Internet to tunnel traffic (including ICMP) between different parts of the network. So the BGP path to 44/8 ends at AS7377, while the traceroute path comes in AS7377, travels cross the overlay network on other ASes, and then ends in AS7377.
 - (b) **Routing Dynamics:** We observed some *substitute* and *extra* mismatch fragments at the end of paths, where AS2512 (CalREN), the provider of UCLA, is appended. An example is ‘= AS12969, -AS43571, +AS2512, =\$', where AS12969 is more than 2 AS hops away from AS2512. This happened because some of our traceroute probes were actually falling in a routing loop within AS2512 immediately after reaching AS12969.
2. **Unannounced Prefixes:** In a *substitute* fragment, only ‘?’ is inserted.
3. **Non-responsive Hops:** In a *substitute* fragment, only ‘*’ is inserted.

For each following cause, given a mismatch fragment M , we conduct the specific tests on one or more AS pairs which are adjacent operands in M . Once one AS pair pass the test, the corresponding cause for M is determined. Let the AS pair be X and Y . For *extra* pattern, the pair is ‘+ X ’ and ‘= Y ’. For *missing* pattern, the pair is ‘= X ’ and ‘- Y ’. For *substitute* pattern, the pair is ‘+ X ’ and ‘- Y ’, or ‘+ X ’ and ‘= Y ’.

4. **MOAS Prefixes:** The matching prefix is announced by both X and Y .
5. **IXPs:** X is an IXP ASN, or the IP addresses mapped to X are used in IXPs.
6. **Siblings:** X and Y belong to the same sibling group.

7. **Neighbors:** Three types of tests are conducted: 1) According to the iPlane’s alias list, the IP address mapped to X belongs to a router that has another interface mapped to Y . 2) X and Y are neighbors in BGP topology. 3) X and Y are neighbors in the topology from the IRR. The contributions of these three tests are 18%, 77% and 5%, respectively.

Table 2. Taxonomy of causes of mismatch as measured in units of paths and fragments.

%	Paths				Fragments			
	ucla	ams-nl	nrt-jp	she-cn	ucla	ams-nl	nrt-jp	she-cn
1 Divergence	0.58	3.34	0.34	2.00	6.25	4.09	6.11	4.46
2 Unannounced	1.49	0.83	7.22	3.49	2.02	1.81	2.31	2.97
3 Non-responsive	14.88	2.77	22.65	24.60	7.28	4.35	3.66	5.08
4 MOAS	9.22	0.50	0.90	1.12	2.42	2.61	2.09	1.97
5 IXP	32.56	1.77	10.86	34.66	6.45	3.21	3.19	4.22
6 Siblings	8.38	3.43	5.25	7.20	5.61	6.99	6.96	6.83
7 Neighbors	37.85	91.72	63.53	29.64	60.52	62.91	62.84	61.84
8 Unknown	0.63	0.36	0.66	1.12	9.45	14.05	12.84	12.65

Table 2 shows the fraction of cases in each class, measured in the percentages of paths and fragments. The path values are relevant for comparison with previous work [13, 10]. We can see that the majority of mismatch cases are the result of foreign addresses including IXPs, siblings and BGP neighbors. And over half of mismatch fragments are due to IP addresses sharing between BGP neighbors, that supports the view of [7]. The contribution of IXPs, siblings and MOAS only sum up to nearly 14%, although the previous work [13, 10, 12] considered them as the major causes. In addition, 6~9% of mismatch cases are due to holes, *i.e.* unannounced or non-responsive hops, in traceroute paths.

Comparing the results in units of path and fragment side by side, we note that the fragment-based results is more robust to the monitor location and the possible flaws in cause inference than the path-based. And there are two types of bias in path-based counting: (1) Overestimating the influence of some causes, such as *IXPs* in *ucla* and *she-cn*. This is mainly because multiple paths may often share a single point of mismatch close to the monitor. (2) Underestimating the difficulty to infer causes. In *Unknown* bin, only about 1% of paths contain 9~14% of fragments. About 93% of fragments in *Unknown* bin are at the end of path. Most of these cases may either correspond to BGP sessions not visible in the current BGP topology or due to misclassified *Divergence* cases.

7 Accuracy of Traceroute-derived AS Connectivity

In this section we assess the accuracy of traceroute-derived AS adjacencies. According to our previous work [14], the BGP table of a monitor should reveal

almost all its AS neighbors over time. The BGP AS graph from UCLA IRL is denoted by G_{bgp} . There are about 180 monitors providing full tables residing in 112 different ASes connected to 13.6K unique ASes through a total of 43K links. Let G_{truth} denote this set of AS adjacencies. A traceroute-derived AS link $X - Y$ is bogus, if either X or Y is in our set of 112 ASes but the link $X - Y$ does not exist in G_{truth} .

To evaluate the inaccuracy of traceroute-derived AS links, we inspect two values: $L_{\overline{bgp}}/L$ and L_{bogus}/L , where L is the number of links discovered by traceroute; $L_{\overline{bgp}}$ is the number of extra links not in G_{bgp} ; L_{bogus} is the number of bogus links. The value $L_{\overline{bgp}}/L$ can be considered as an upper bound of the error rate, while the value L_{bogus}/L should be seen as a lower bound of inaccuracy.

To understand how the extra links and bogus links were created, we search for these links in our mismatch fragments and group them in the causes described in the previous section. Table 3 shows the inaccuracy of traceroute-derived AS topology by causes. The results from different monitors are similar. We see that the cause *Neighbors* are responsible for most of the links not seen in BGP, and contribute to the highest chunk of bogus links.

Table 3. Inaccuracy of traceroute-derived AS topology by causes.

%	$L_{\overline{bgp}}/L$				L_{bogus}/L			
	ucla	ams-nl	nrt-jp	she-cn	ucla	ams-nl	nrt-jp	she-cn
1 Divergence	1.85	2.19	2.61	2.26	0.68	0.65	0.82	0.78
2 Unannounced	0.76	0.78	1.09	1.31	0.42	0.37	0.55	0.29
3 Non-responsive	2.67	1.88	1.67	2.27	1.01	0.84	0.58	1.01
4 MOAS	0.62	0.68	0.50	0.54	0.10	0.09	0.09	0.08
5 IXP	2.29	1.19	1.32	1.51	0.35	0.35	0.27	0.40
6 Siblings	1.08	1.42	1.61	1.47	0.22	0.34	0.39	0.28
7 Neighbors	14.64	13.28	13.92	12.89	5.24	3.95	4.36	4.07
8 Unknown	3.22	5.68	5.63	5.50	0.65	1.05	1.01	0.97
Total	24.96	24.24	25.39	24.94	6.19	4.30	4.22	4.61

We also verify the accuracy of the AS adjacency sets provided by Ark and DIMES in Table 4. About 47% of AS adjacencies in the traceroute-derived topologies are not seen in BGP. In addition, we verify that about 16% of the traceroute AS adjacencies are false. Discarding the indirect links, that is caused by non-responsive hops or unannounced prefixes, in CAIDA’s data can reduce the fraction of bogus links from 13% to 8%. However still 31% of the extra links not seen in BGP are actually bogus ($L_{bogus}/L_{\overline{bgp}}$).

8 Conclusion

In this paper we develop a systematic approach to identify and classify errors in AS paths inferred from traceroute using the LPM method. Our results shed

Table 4. Inaccuracy of public traceroute AS topology data sets.

	L	L_{bgp}	L_{bogus}	L_{bgp}/L	L_{bogus}/L	L_{bogus}/L_{bgp}
DIMES	77358	32159	11204	41.6%	14.5%	34.8%
Ark _{direct}	56014	14515	4510	25.9%	8.0%	31.1%
Ark _{indirect}	69962	25215	9059	36.0%	12.9%	35.9%
Total	104844	49731	17062	47.4%	16.3%	34.3%

light into the major pitfalls of traceroute-based AS topology measurement and show the limitations of publicly available AS topologies derived from traceroute. Since most of the inconsistencies originate from IP address sharing between BGP neighbors, we believe that building an accurate database of router interface aliases can bring significant improvement to the accuracy of the router path to AS path conversion process, and this is part of our future work.

References

1. Archipelago Measurement Infrastructure, <http://www.caida.org/projects/ark/>
2. Internet Routing Registry, <http://www.irr.net/>
3. RIPE routing information service project, <http://www.ripe.net/>
4. RouteViews routing table archive, <http://www.routeviews.org/>
5. UCLA IRL Internet topology collection, <http://irl.cs.ucla.edu/topology/>
6. Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Avoiding traceroute anomalies with paris traceroute. In: IMC '06 (2006)
7. Chang, H., Jamin, S., Willinger, W.: Inferring AS-level Internet topology from router-level path traces. In: SPIE ITCOM (2001)
8. Hunt, J.W., McIlroy, M.D.: An algorithm for differential file comparison. Tech. rep., Bell Laboratories (1976)
9. Hyun, Y., Broido, A., kc claffy: On third-party addresses in traceroute paths. In: Proc. of Passive and Active Measurement Workshop (PAM) (2003)
10. Hyun, Y., Broido, A., kc claffy: Traceroute and BGP AS path incongruities. Tech. rep., CAIDA (2003)
11. Madhyastha, H., Isdal, T., Piatek, M., Dixon, C., Anderson, T., Krishnamurthy, A., Venkataramani, A.: iPlane: an information plane for distributed services. In: Proc. of OSDI (2006)
12. Mao, Z.M., Johnson, D., Rexford, J., Wang, J., Katz, R.H.: Scalable and accurate identification of AS-level forwarding paths. In: INFOCOM 2004 (2004)
13. Mao, Z.M., Rexford, J., Wang, J., Katz, R.H.: Towards an accurate AS-level traceroute tool. In: Proc. of ACM SIGCOMM (2003)
14. Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.: In search of the elusive ground truth: The Internet's AS-level connectivity structure. In: Proc. ACM SIGMETRICS (2008)
15. Oliveira, R., Zhang, B., Zhang, L.: Observing the evolution of Internet AS topology. In: ACM SIGCOMM (2007)
16. Shavitt, Y., Shir, E.: DIMES: Let the Internet measure itself. ACM SIGCOMM Computer Comm. Review (CCR) (2005)