Measuring the Internet Topology with Smartphones

Adriano Faggiani
IIT-CNR
Pisa, Italy
adriano.faggiani@iit.cnr.it

Enrico Gregori
IIT-CNR
Pisa, Italy
enrico.gregori@iit.cnr.it

Luciano Lenzini
Dip. Ing. dell'Informazione,
University of Pisa
Pisa, Italy
luciano.lenzini@iet.unipi.it

Valerio Luconi
Dip. Ing. dell'Informazione
University of Pisa
Pisa, Italy
valerio.luconi@iet.unipi.it

ABSTRACT

Despite the very well known smartphone issues such as on-off behaviour and battery/bandwidth limitations, in this paper we show that smartphones can be successfully employed in a crowdsourcing system to perform Internet AS-level topology discovery. We propose and illustrate a measurement methodology that takes these issues into account. We implemented such methodology in Portolan, our smartphone-based crowdsourcing system, and ran six months of measurements. We show that smartphones mobility allows to obtain measurements from 706 different ASes with just 200 active devices. Moreover, we show that our methodology manages to bring novelty with relatively few measurements. On average 27.75% of the AS links found by Portolan are not found by BGP measurements.

CCS Concepts

•General and reference → Measurement; •Networks → Network measurement; Public Internet;

Keywords

Crowdsourcing, smartphones, Internet measurements

1. INTRODUCTION

In the last decade the crowdsourcing paradigm has been intensively adopted for running large-scale tasks with the help of the masses. Even in the networking community the crowdsourcing approach has gained consensus as a viable way to gather various kinds of measurements on several aspects of the Internet (e.g. structure, performance, traffic, neutrality). Several crowdsourcing systems have been deployed. Some of them are based on (almost) fixed devices. For example, Dasu [19], that uses a plugin for a BitTorrent client to perform broadband characterization, or RIPE Atlas [18], that provides an infrastructure of fixed probes that everyone can use to run its own network measurements, or DIMES [20] that focuses on discovering the topology of the Internet at the Autonomous System (AS) level of abstraction, by distributing a software agent for PCs. Recently a whole range of systems have been relying on mobile devices, such as smartphones. Among them the most relevant are NetAlyzr [16], that provides a smartphone app that allows a user to measure the performance of its network and detect possible anomalies, and MobiPerf [15], that provides to the research community a measurement facility based on hundreds of mobile devices. Smartphones offer advantages, like their mobility and pervasiveness, that make them appealing to build a crowdsourcing system. However smartphones pose also several challenges, as they are low-resource devices. Battery and bandwidth are relevant issues, and their unlimited consumption could bother the user. Thus the number of measurements that smartphones can carry out is restricted. Moreover, smartphones show an on-off behaviour, thus the assignment of measurement tasks is not always trivial.

Despite all these limitations, in this paper we show that smartphones can be successfully exploited to build a system for discovering the Internet AS-level topology. We briefly show how we coped with these limitations to develop Portolan [4, 7, 3], that, to the best of our knowledge, is the first and only smartphone-based crowdsourcing system explicitly designed for performing Internet AS-level topology discovery with active measurements (i.e. traceroute). We then propose a measure-

C2B(I)D '15, August 17, 2015, London, United Kingdom
 ISBN 978-1-4503-3539-3/15/08…$15.00
DOI: http://dx.doi.org/10.1145/2787394.2787398

ment methodology that takes into account these issues and focuses on reducing the number of measurements, in order to keep measurement campaigns sustainable for smartphones, without losing the ability to discover new AS links. We implemented our methodology in Portolan and we present data gathered by Portolan in six months, from July 1st 2014 to December 31st 2014, divided in snapshots of two months each, showing pros, cons and peculiarities of conducting Internet AS-level topology discovery with smartphones. We show that the mobility of smartphones can bring several advantages in terms of vantage points, in fact we were able to carry out measurements from 706 different ASes with an average of just 200 active devices. We also show that our methodology is able to find a significant number of AS links with a rather limited number of traceroutes, and that on average 27.75% of the links found by Portolan were not found in the BGP-inferred topology. On the other hand, we found that our methodology can still be improved, since a large percentage of our traceroutes still provide redundant information.

The rest of the paper is structured as follows. Section 2 briefly describes the Portolan architecture. In Sect. 3 we describe the measurement methodology. Section 4 shows the results of the measurement campaign conducted by Portolan in July-December 2014 and Sect. 5 concludes the paper.

2. PORTOLAN

Portolan is a smartphone-based crowdsourcing system that we built for measuring the Internet AS-level topology[1]. Portolan was used to conduct the measurements discussed in this paper. The architecture of Portolan has been widely discussed in [4, 7, 3]. However, to make the paper self contained, we provide an essential description of Portolan.

Portolan is composed by a client side represented by a smartphone application for Android[2], and a server side. The smartphone app implements a UDP version of Paris traceroute [1]. UDP traceroute is the only one that can be run on Android without root privileges [4]. To mitigate the effect of smartphones on-off behaviour, the assignement of measurement tasks to smartphones is done via a polling mechanism. The measurement engine of the smartphone app starts at boot and runs in the background. Periodically it polls the server to receive a short list of maximum 100 target IP addresses that are to be probed. This way, every time the smartphone is on, it performs measurements for Portolan. Once a smartphone has finished probing the target list, it returns the results back to the server. The server then aggregates data from all smartphones to produce AS-level topologies. Since smartphones have limited resources, we restricted the amount of traffic produced by the Portolan app to maximum 2MB/day when the smartphone is connected to a cellular (2G, 3G or 4G) network. When the device is connected to a WiFi network, traceroutes are unlimited. However, to not consume too much battery, Portolan stops running if the battery level goes below 40%.

[1]Since January 2015 Portolan also provides a software agent for PCs, however their contribution has not been considered in this paper

[2]https://play.google.com/store/apps/details?id=it.unipi.iet.portolan.traceroute

3. METHODOLOGY

Our methodology consists of three phases. The first is a target selection phase that is the core of our methodology. Starting from the assumption that the Internet AS-level topology inferred with BGP data is reliable although incomplete, we focus on exploring just the portion of the Internet that BGP fails to reveal, i.e. the Internet periphery [8, 6, 17]. The selected targets must thus be limited in number, to keep the measurement campaigns sustainable, and must be located in the periphery of the Internet, to avoid as possible to discover portions of the Internet already inferred via BGP data. This phase is performed offline by our server. The second phase is the actual data collection phase performed by Portolan smartphones. In this phase the target list produced by the first phase is probed cyclically. Finally the third phase is a data analysis and filtering phase that aims at reducing biases due to known issues related to traceroute [5]. In must be noted that the proposed methodology is general and could be adopted by any system that aims at Internet AS-level topology discovery with traceroute measurements, especially systems that rely on low-resource devices such as smartphones.

3.1 Target Selection

This phase has two goals: *(i)* generate a short list of targets that allows low resource devices such as smartphones to finish all the measurements in a reasonable amount of time, given the limitations discussed above; *(ii)* keep the measurements confined in portions of the Internet not already discovered with BGP data, i.e. the Internet periphery.

To keep the measurements confined in the Internet periphery, both measurement monitors and target IP addresses should be located in that portion of the Internet. Since we have no control over the position of Portolan smartphones, as they are under the control of volunteer users, we can only choose target IP addresses to ensure that at least the measurements are directed to the periphery. We thus select our target IP addresses from the address space of stub ASes. Stub ASes are ASes that do not have any customers, thus they are in the lowest layer of the Internet hierarchy, i.e. they are the leaves of the Internet AS graph. To identify stub ASes and their announced address space we use BGP data collected by RouteViews[3] and RIPE RIS[4].

[3]http://www.routeviews.org/

[4]http://www.ripe.net/

Choosing stub ASes as targets has also the positive effect to reduce the number of measurements. Since stub ASes are usually small ASes confined to a single geographical region with a small announced address space, we can select just one target IP address per stub. However, having such a small number of targets, we cannot rely on randomly chosen targets. This because only a small portion of all the IP addresses that are announced with BGP are allocated and respond to UDP probes [12]. Therefore, to obtain effective measurements we need to choose IP addresses that are guaranteed to respond with an ICMP *Port Unreachable* message to Portolan UDP traceroute probes. If the target IP address is a non-responding or non-allocated address, the traceroute could stop at an address of the upstream provider of the target AS, thus missing the last AS link. To identify the IP addresses responding to UDP probes within the network announced by stub ASes, we use the ISI ANT Lab Internet Address History dataset[5]. This dataset provides all the IP addresses responding to ICMP Ping, together with a responsivity index. Since Portolan uses the UDP traceroute, we probe all the IP addresses belonging to stub ASes with UDP Ping. From those that respond to UDP Ping, for each stub AS we choose the one that is tagged as the most responsive in the Internet Address History dataset.

3.2 Data Collection

Once the target list has been defined, it is probed cyclically and separately from each AS that hosts at least one Portolan smartphone (from now on *source AS*). A cycle for a given source AS is completed when all the addresses of the target list are probed with measurements starting from that AS. The amount of time needed to complete a cycle depends on how many Portolan smartphones are located in one AS, since the target list is divided among all the smartphones in that AS. However, since we have no control over the devices running the Portolan app, we cannot know a priori the duration of a cycle of measurements. The results of the measurements are analysed at intervals of two months to produce snapshots of the discovered topology.

The reason to run cyclic measurements is that it allows to observe the evolution of the topology over time. Cyclic measurements could also increase the probability of discovering temporary backup configurations set up after network faults. Moreover, certain targets may also not respond at certain times, thus probing them repeatedly could reduce these misses.

3.3 Data Analysis and Filtering

In this section we describe the hygiene phase run on raw traceroute data to limit the biases that can result when inferring the Internet AS-level topology from traceroute data. Issues can occur at the IP level or when converting the IP paths discovered by traceroute into AS paths.

At the IP level, most of these issues are solved using Paris traceroute [1]. However, other issues such as loops due to router misconfigurations remain open. The UDP Paris traceroute implementation of Portolan performs a pre-filtering phase aimed at detecting and correcting these anomalies before sending measurement results back to the server. If a traceroute probe stops at IP interfaces already encountered in previous hops, the measurement is immediately stopped.

Once data is stored in the Portolan server, IP paths must be converted into AS paths with a process named as IP-to-AS mapping, in order to extract the links between ASes. This process is subject to several known issues that could lead to the inference of false AS links [13, 9, 2]. Thus, after the IP-to-AS mapping we run a data hygiene process to filter out those links that are potentially wrong. The entire procedure involves five steps.

Step 1 – IP-to-AS mapping. We map IP interfaces to the ASes they belong to by performing the longest prefix match in BGP routing tables extracted from RouteViews and RIPE RIS BGP data, enhanced with prefixes from the peering LANs of IXPs publicly available at PeeringDB[6]. We thus obtain an AS path for each traceroute. In this step we filter out those AS paths that contain AS loops. We are aware that these AS paths may also contain valid links, but, in order not to introduce false links, we prefer a conservative approach.

Step 2 – Link extraction. In the AS paths we look for pairs of consecutive hops that belong to different ASes and pairs of hops belonging to different ASes separated by one single hop belonging to the peering LAN of an IXP. Thus, we define two kinds of links, *direct links* and *IXP links*. *Direct links* are direct connections between two ASes, i.e. without any hop belonging to other entities in the middle. *IXP links* instead, are connections established via an IXP. This distinction will be used in the next steps of the filtering phase, where we apply two separate filters to the two link types (Step 4 to *IXP links* and Step 5 to *direct links*).

Step 3 – MOAS (Multi-Origin ASes) prefixes filtering. Some IP addresses could be mapped to MOAS prefixes, i.e. prefixes that could belong to multiple ASes. Thus, we filter out all those links where one of the IP addresses is mapped to more than one AS.

Step 4 – IXP links filtering. This data filtering involves only *IXP links*. We collect the lists of participants of all known IXPs monthly provided by PeeringDB. In our *IXP links* set every link is tagged with the IXP used to establish that link. We filter out from our *IXP links* set all the links where at least one AS is not present in the participant list of the IXP used to establish that link.

Step 5 – Direct links filtering. This last step

[5]http://www.isi.edu/ant/

[6]https://www.peeringdb.com/

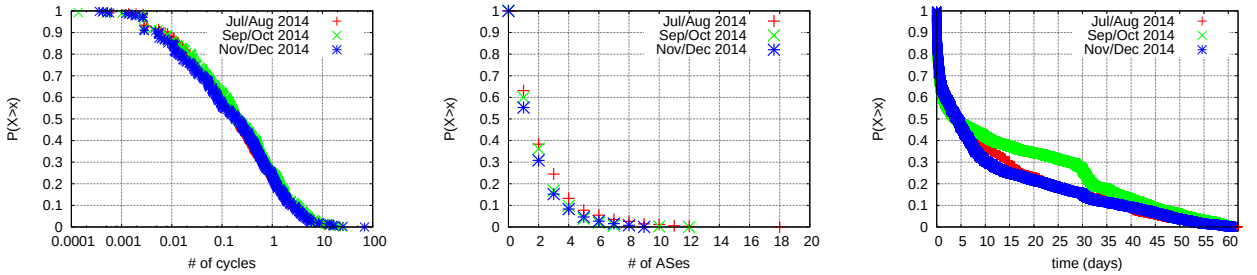

(a) Cycles completed by each source AS.
(b) Number of ASes visited by each device.
(c) Link discovery in time (all links).

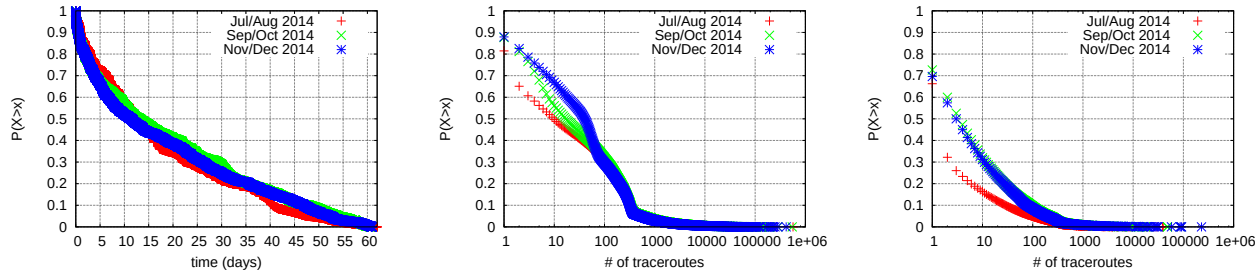

(d) Link discovery in time (new links).
(e) Number of traceroutes per link (all links).
(f) Number of traceroutes per link (new links).

Figure 1: CCDFs.

consists in filtering out possible false AS links from the *direct links* set using the BGP AS paths collected during the same period of traceroute measurements. Some of these links could be wrong due to wrong mapping of AS Border Routers or occurrence of *third-party addresses* [9, 14, 11]. To partially avoid these errors, for each AS connection A–B found after the IP-to-AS mapping, we check whether a BGP AS path ...–A–X_1–...–X_n–B–... exists. If it does, we assume an incorrect inference, due to missing AS hops or third-party addresses, and we discard the AS link because it may be a false positive. Note that the AS link A–B may also be correct but not observed by the public BGP monitors due to their placement. However, in order to not introduce false positives we prefer to adopt a conservative approach.

4. RESULTS

In this section we show the results of the traceroute measurements carried out by Portolan between July 1st 2014 and December 31st 2014 divided in three periods of two months each[7] (from now on *period*). In that interval the Portolan target list was counting 36,183 IP addresses. We chose a period length of two months because of the following considerations. Since Portolan has a daily average of 200 active devices, the period should be long to obtain an adequate number of measurements (given the limited resources of smartphones). However, a too long period could lead to a low reactivity to changes in the topology (due to activation or termination of AS links). Thus the two months choice seems to us a good trade-off between these two aspects.

We start our analysis by showing the advantages of using smartphones as measurement monitors. Fig. 1a depicts the CCDF of the number of cycles completed by each source AS in the three periods. It can be noted that the source ASes completing at least one cycle are just 25% of the total. This can be due to two main reasons. First, some source ASes could host just few smartphones, thus the time to complete a cycle increases. Second, some source ASes could have been visited by few smartphones for a limited amount of time, thus only a few measurements are carried out from these source ASes. However, this can be seen as a positive aspect, because even if a smartphone stays in a source AS for a short time, it manages to carry out some measurements and to provide a new point of view. By observing Fig. 1b it can further be noted the positive aspect of using smartphones as measurement agents. The figure depicts the CCDF of the number of source ASes visited by each device. It can be observed that on average on the three periods 60% of the devices visit more than one source AS, with a maximum that goes from 8 to 18 ASes visited by a single device. This shows how smartphones are useful for Internet topology discovery, as a single device acts as multiple vantage points. This can be further enphasized if observing Table 1. The number of source ASes for each period is on average higher than

[7]All raw data is available at http://portolan.iet.unipi.it/data/traceroute/

Table 1: Results summary.

	Portolan				BGP
	# links	**# new links**	**# traceroutes**	**# source ASes (# exclusive)**	**# links**
Jul–Aug	73,278	29,605 (40.40%)	15,898,119	461 (156)	228,365
Sep–Oct	67,122	13,946 (20.78%)	17,008,252	404 (80)	234,678
Nov–Dec	72,242	15,952 (22.08%)	16,521,242	397 (110)	235,248
Union	109,404	42,828 (39.15%)	–	706	–
Intersection	39,891	5,324 (13.35%)	–	196	–

400, and the union is 706. This is a very interesting result if we consider that the number of active Portolan smartphones is on average 200, and can be explained by the high mobility of smartphones.

We continue our analysis showing that our measurement methodology implemented in Portolan is able to bring novelty with respect to BGP topology. Table 1 shows the number of links found by Portolan and BGP in the three periods, and the number (and percentage) of links found by Portolan that are not found in BGP-inferred topologies (from now on *new links*). On average on the three periods 27.75% of the links found by Portolan are not found with BGP. We notice that in Jul-Aug 2014 the number of new links is significantly higher if compared to that of the other two periods. This can be explained by observing the number of links in BGP topology, which is lower than that of the other two periods, and the number of source ASes, which is significanlty higher. This means that in the Jul-Aug perdiod Portolan smartphones have shown a higher mobility than the other periods. The reason could be that Jul-Aug is vacation time, and usually people visit new places and thus new networks.

To deepen our analysis on the topologies produced by Portolan we consider the union and the intersection of the three sets of links and new links (Table 1). We can observe that these sets are not highly overlapped. This means that in each period there is a significant percentage of links that are not found in the other periods. This can again be explained with the smartphones mobility and intermittence, by observing that this behaviour is also found in the union and intersection of the source ASes sets. Moreover, if we consider the exclusive source ASes for each period (i.e. the source ASes that were visited only in that period) we can notice that they are always a substantial percentage of the total set. Thus we can conclude that our measurements are very useful to enhance the BGP-inferred topologies, but they cannot be used as the only source of Internet topology data, to extract meaningful properties of the Internet topology. Another consequence of the mobility and intermittence of smartphones can be observed in Fig. 1c-1d. The figures show the CCDFs of the link discovery in time. We notice that the discovery of both links and new links is distributed among the whole period, with an average of 25-30% of links discovered in the second month. This means that the whole period is used to bring novelty.

Finally, we provide some efficiency considerations. If we consider Table 1 we can notice that Portolan is able to discover an average of 70K links per two months period with an average of 16M traceroutes. This is a high efficiency if compared to results achieved by systems with similar aims. For example, CAIDA Archipelago[8] discovers on average 90K links with 500M traceroutes in two months[9]. However, the efficiency of the measurement methodology can be further enhanced. We consider Fig. 1e-1f, that show the CCDFs of the number of traceroutes per link. We notice that in 95% of cases (i.e. for 95% of links) the number of traceroutes is less than 1,000. We also notice that 5% of links is found in a huge amount of traceroutes, that thus could provide redundant information. To quantify this redundancy we calculated for each period the total number of traceroutes that discover the links seen by up to 1,000 traceroutes (Table 2), to have an insight on how many traceroutes provide information with low redundancy. If we consider the 95% of links that are found in less than 1,000 traceroutes each, the total number of traceroutes is approximately 6M, (approximately 40% of the total). This means that even if the total number of measurements is not so high, there is still margin of improvement in eliminating redundant information. However, it must be noticed that a certain degree of redundancy is typical of crowdsourcing systems and of traceroute measurements in general, as the paths close to the same sources or the same destinations are always similar [10].

5. CONCLUSIONS

In this paper we showed that the smartphone platform can be successfully adopted to carry out Internet AS-level topology discovery within a crowdsourcing system. We showed the challenges that the smartphone platform poses (i.e. low battery and bandwidth resources, on-off behaviour), and how we faced them within Portolan. We introduced a measurement methodology that takes into account the limited number of measurements that each smartphone can carry out. Our methodology focuses on campaigns with a limited num-

[8]http://www.caida.org/projects/ark/

[9]We are aware that the purpose of CAIDA Archipelago is slightly different, however this is just to give a rough comparison

Table 2: Non redundant traceroutes.

	traceroute/link			
	≤ 1 (avg. 15% links)	≤ 10 (avg. 40% links)	≤ 100 (avg. 68% links)	≤ 1000 (avg. 95% links)
Jul–Aug	13,591	98,772	704,853	6,281,824
Sep–Oct	8,297	112,265	768,057	6,032,234
Nov–Dec	8,725	79,006	1,285,591	6,536,042

ber of targets specifically chosen to be responsive and to be residing in the Internet periphery. Finally we prensented results from six months of data collection with Portolan, focusing of the pros and cons of using smartphones. We showed that smartphones are particularly suited to act as vantage points, due to their mobility. In fact they have been able to conduct measurements from over 700 source ASes with an average of just 200 active devices. We also showed that with our methodology Portolan has been able to bring novelty with relatively few measurements, discovering an average of 27.75% of links not found by BGP measurements. On the other hand, we showed that the links found by Portolan in one period show low overlap with the links of other periods, thus they can only be used to enhance other topologies (e.g. BGP ones). Moreover, the number of measurements can still be lowered, given that the percentage of useful traceroutes is approximately 40%.

6. REFERENCES

[1] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proc. ACM SIGCOMM IMC '06*, pages 153–158, 2006.

[2] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes From P2P users. In *Proc. ACM SIGCOMM CoNEXT '09*, pages 217–228, 2009.

[3] A. Faggiani, E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio. Smartphone-based crowdsourcing for network monitoring: Opportunities, challenges, and a case study. *IEEE Comm. Mag.*, 52(1):106–113, 2014.

[4] A. Faggiani, E. Gregori, L. Lenzini, S. Mainardi, and A. Vecchio. On the feasibility of measuring the Internet through smartphone-based crowdsourcing. In *Proc. WiOpt '12*, pages 318–323, 2012.

[5] A. D. Flaxman and J. Vera. Bias Reduction in Traceroute Sampling – Towards a More Accurate Map of the Internet. In A. Bonato and F. Chung, editors, *Algorithms and Models for the Web-Graph*, volume 4863 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin Heidelberg, 2007.

[6] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. A Novel Methodology to Address the Internet AS-Level Data Incompleteness. *IEEE/ACM Trans. Netw.*, PP(99):1–1, 2014.

[7] E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio. Sensing the Internet through crowdsourcing. In *Proc. PerMoby '13*, pages 248–254, 2013.

[8] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy. Lord of the Links: A Framework for Discovering Missing Links in the Internet Topology. *IEEE/ACM Trans. Netw.*, 17(2):391–404, 2009.

[9] Y. Hyun, A. Broido, and k. claffy. On Third-party Addresses in Traceroute Paths. In *Proc. PAM Workshop '03*, 2003.

[10] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proc. INFOCOM '03.*, pages 332–341 vol.1, 2003.

[11] M. Luckie and k. claffy. A Second Look at Detecting Third-Party Addresses in Traceroute Traces with the IP Timestamp Option. In *Proc. PAM '14*, pages 46–55, 2014.

[12] M. Luckie, Y. Hyun, and B. Huffaker. Traceroute Probe Method and Forward IP Path Inference. In *Proc. ACM SIGCOMM IMC '08*, pages 311–324, 2008.

[13] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an Accurate AS-level Traceroute Tool. In *Proc. ACM SIGCOMM '03*, pages 365–378, 2003.

[14] P. Marchetta, W. de Donato, and A. Pescapé. Detecting Third-Party Addresses in Traceroute Traces with IP Timestamp Option. In *Proc. PAM '13*, pages 21–30. 2013.

[15] MobiPerf. http://www.mobiperf.com/.

[16] Netalyzr. http://netalyzr.icsi.berkeley.edu/.

[17] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (In)Completeness of the Observed Internet AS-level Structure. *IEEE/ACM Trans. Netw.*, 18(1):109–122, Feb. 2010.

[18] RIPE Atlas. https://atlas.ripe.net/.

[19] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing Experiments to the Internet's Edge. In *Proc. USENIX NSDI '13*, pages 487–499, 2013.

[20] Y. Shavitt and E. Shir. DIMES: Let the Internet Measure Itself. *ACM SIGCOMM Comput. Commun. Rev.*, 35(5):71–74, 2005.