

On the Analysis of Internet Paths with DisNETPerf, a Distributed Paths Performance Analyzer

Original

On the Analysis of Internet Paths with DisNETPerf, a Distributed Paths Performance Analyzer / Wassermann, Sarah; Casas, Pedro; Donnet, Benoit; Leduc, Guy; Mellia, Marco. - ELETTRONICO. - (2016), pp. 72-79. (IEEE 41st Conference on Local Computer Networks Workshops (LCN Workshops), 2016 Dubai, UAE November 2016) [10.1109/LCN.2016.031].

Availability:

This version is available at: 11583/2666872 since: 2017-03-11T14:34:48Z

Publisher:

IEEE

Published

DOI:10.1109/LCN.2016.031

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

On the Analysis of Internet Paths with DisNETPerf, a Distributed Paths Performance Analyzer

Sarah Wassermann*, Pedro Casas†, Benoit Donnet*, Guy Leduc*, Marco Mellia‡

*Université de Liège, Montefiore Institute

sarah.wassermann@student.ulg.ac.be, {forename.surname}@ulg.ac.be

†AIT Austrian Institute of Technology

pedro.casas@ait.ac.at

‡Politecnico di Torino

mellia@polito.it

Abstract—`traceroute` is the most widely used Internet path analysis tool today to study the topology of the Internet and to diagnose routing failures as well as poor performance events. A major limitation of `traceroute` when the destination is not controllable by the user is its inability to measure *reverse paths*, i.e., the path from any given destination back to the source. This is a major drawback for ISPs, who need to understand the performance of the Internet paths connecting popular services (e.g., YouTube and Facebook) to their customers. Even if public servers and distributed measurement platforms can provide partial reverse path visibility through ad-hoc measurements, there is still a need for a structured approach capable of analyzing the performance of Internet paths connecting any pair of nodes (servers, routers, hosts, etc.). While the problem of *reverse traceroute* has been addressed in the past, proposed techniques rely on IP address spoofing – which might lead to security concerns, and assume the availability of certain route-tracking options –, which might not be available. In this paper, we introduce and evaluate DisNETPerf, a new tool which provides exactly the same type of information as `traceroute`, but for paths connecting arbitrarily selected nodes. DisNETPerf works by firstly locating probes (i.e., measurement points) that are the closest to a given target node, using them to perform `traceroute` measurements from the target point-of-view to a given destination for path performance monitoring and troubleshooting purposes. We propose two techniques for probe location, and demonstrate that the reverse path (from server to users) can be measured with very high accuracy in certain scenarios. We also analyze relevant characteristics of Internet paths and distributed measurement platforms, which reinforce the applicability and relevance of DisNETPerf in current Internet.

Index Terms—Distributed Active Measurements; Reverse Traceroute; Internet Paths’ Symmetry; RIPE Atlas Measurement Framework

I. INTRODUCTION

Which are the Internet paths connecting a certain server to my customers? How many hops have these paths? What is the performance of these paths? Which parts of the paths are congested? These are standard questions a network operator poses herself when she needs to diagnose performance issues impacting her customers, potentially originating outside the boundaries of her own network. A normal approach the

operator would follow to answer these questions is to run `traceroute` measurements from some controlled node at the edge of her network (connecting her customers to the Internet) toward some servers which are getting impacted (e.g., YouTube and Facebook servers). Indeed, `traceroute` is still today the de-facto standard tool used by operators to investigate routing failures and performance problems [1]. By assuming path symmetry between her controlled node and the targeted servers, she would be able to get some initial hints on the performance of the end-to-end paths. However, assuming path symmetry at the Internet level is a major mistake [2], [3], [4], as Internet paths often become asymmetric, especially at network boundaries, due to administration policies changes among others. This has been flagged as the number one “plague” of `traceroute` [1], and one can only assume that `traceroute` shows relevant information for the forward path. The reverse path itself is therefore completely invisible. However, `traceroute` has a major limitation when the destination is not accessible: it cannot measure the reverse path (i.e., from destination back to the source) as one cannot run measurements from the inaccessible destination. This is exactly the problem we tackle in this paper: performing *reverse traceroute* measurements.

Let us present as motivation example a real performance degradation event impacting a large number of customers during several days for a specific bandwidth-demanding service: YouTube.

Internet-scale services such as YouTube are provisioned from geo-distributed servers, using large Content Delivery Networks (CDNs). While user-requests are normally redirected to the closest servers (in terms of latency), internal CDN load-balancing policies may select servers which lie at hundreds of milliseconds from customers, potentially impacting their Quality of Experience (QoE), especially under congestion of the downlink paths connecting YouTube servers to customers. In such a context, it is very difficult for an ISP to find the root cause of the problem, as reported in previous work [5]. The authors of the latter show that a drop in the throughput has a major impact on the customers’ QoE.

At the same time of this anomaly, the minimum RTT to YouTube servers increases by more than 300%, corresponding to a different load balancing policy selecting servers at much farther locations [6]. The root causes behind the anomaly could be either that the selected YouTube servers were not correctly dimensioned, or that there is some heavy network congestion at peak time in the paths from the selected servers to the customers. Having a tool that can measure the performance of the paths from these YouTube servers toward the monitored customers becomes paramount for the ISP, to diagnose and troubleshoot the detected performance degradation.

Previous work proposed a tool to perform *reverse* traceroute [7], i.e., from the YouTube servers to the customers. A major drawback of the proposed approach is that it heavily relies on IP-spoofing and IP Record Route Option, both being not necessarily allowed everywhere [8], [9] and potentially leading to security concerns for the case of IP-spoofing. Therefore, in this paper, we introduce DisNETPerf, a Distributed Internet Paths Performance Analyzer, which can monitor any Internet path using distributed active measurements. While DisNETPerf is not strictly tied to any particular distributed measurement platform, in this paper – as well as in current DisNETPerf implementation –, we rely on the RIPE Atlas framework [10]¹ for running and evaluating the proposed solution. While recent studies have shown that RIPE Atlas might incur into biased results due to measurement interference issues [11], its large scale in terms of deployed nodes and geographical diversity provides a better framework for DisNETPerf, as compared to other popular distributed measurement frameworks such as CAIDA’s Ark [12] or PlanetLab [13].

DisNETPerf uses a combined topology- and delay-based distance notion to locate a probe that is as close as possible to a desired target destination, from which reverse traceroute measurements should be run. By doing so, DisNETPerf aims at locating probes which offer a very high path similarity to the real reverse path. In our study, we evaluate two different probe selection techniques to locate these probes, and show that reverse Internet paths can be measured with relatively high accuracy at the AS level, PoP level, and even router-interface level. This paper represents a continuation of our first research efforts around the DisNETPerf system [14].

The remainder of this paper is organized as follows: Sec. II describes the basics behind DisNETPerf, focusing on the different probe selection techniques. Sec. III presents an evaluation on the performance achieved by the probe selection techniques used by DisNETPerf, using RIPE Atlas as distributed measurement platform. In this section, we additionally provide characterization results on the topology and diversity of RIPE Atlas, which is critical to the proper operation of DisNETPerf. Finally, we also perform an Internet paths symmetry analysis using RIPE Atlas, to further support the paramount contributions of DisNETPerf. In Sec. IV, we provide additional




Figure 1. DisNETPerf overview. The first step of DisNETPerf consists of selecting a monitoring point or probe located as close as possible to a target server, to later on perform traceroute measurements towards specific destinations.

discussion on the obtained results, presenting some highlights on the application of DisNETPerf in the practice. Sec. V briefly reviews the related work. Finally, Sec. VI concludes this work.

II. REVERSE TRACEROUTE WITH DISNETPERF

To compute and monitor the path from a given content server to a specific user, we developed DisNETPerf, a distributed Internet paths performance analyzer. Current implementation of DisNETPerf locates the closest RIPE Atlas probe to this content server, and gathers information about the path leading from the selected probe to the customer. DisNETPerf is open source and freely available (<https://github.com/SAWassermann/DisNETPerf>).

Fig. 1 describes the overall idea behind the DisNETPerf approach. In a nutshell, given a certain content server with IP address IP_s , and a destination customer with IP address IP_d , DisNETPerf pinpoints the closest box, namely IP_c , using a combined topology- and delay-based distance: probes are located first by AS – using BGP routing proximity to select probes in the same AS as IP_s – and then by propagation delay – for electing the closest probe to IP_s . DisNETPerf then periodically runs traceroute measurements from IP_c to IP_d , collecting different path performance metrics such as RTT per hop, end-to-end RTT, etc. This data might then be used to troubleshoot paths from the content server (mimicked by IP_c) to the target customer.

Current DisNETPerf implementation uses two different probe selection approaches for locating IP_c , partially proposed in the literature for IP geolocation [15], [16], [17]. We called these selection approaches the *smallest latency* (SL) approach (Sec. II-A) and the *landmark* (LM) approach (Sec. II-B), which we describe next.

A. Probe Selection by Smallest Latency

The SL approach starts by determining whether RIPE Atlas probes are located in the same AS as the targeted content server IP_s . If this is not the case, the SL approach tries to

¹This research is supported by measurements obtained from RIPE Atlas, an open measurements platform operated by the RIPE NCC.

locate probes in the neighbor ASes of IP_s . Neighborhood information is obtained through AS relationships. In this paper, we use CAIDA’s AS relationships dataset [18]. If no probes are found in the neighbor ASes neither, then the SL approach randomly selects a large (and configurable) set of boxes among all the available ones. Once RIPE Atlas probes have been identified or pre-selected, the selection of IP_c can start. We call these pre-selected probes the “candidate probes”.

The SL approach then selects as IP_c the candidate probe with the smallest latency to the target IP_s . Latency is computed on the basis of standard ping measurements; in particular, the SL approach issues ten ping measurements from each of the candidate probes toward IP_s . The candidate probe with the smallest minimum RTT to IP_s is finally elected as the representative probe of the content server, i.e., IP_c . We consider the minimum RTT, as it provides a rough estimation of the propagation delay between two IP addresses.

B. Probe Selection using Landmarks

The first step of the LM approach is exactly the same as the one followed by the SL approach, i.e., candidate probes are firstly selected based on their AS. However, the continuation is slightly different. The next step consists of grouping the candidate probes in two different sets: the *landmarks* and the probes that can be elected as IP_c . Landmarks are chosen randomly among all the candidate probes. Then, ten ping measurements are issued from each of the landmarks toward IP_s and toward all the candidate probes belonging to the other set. For each pinged IP address, a feature vector d is computed, containing the minimum RTT from each landmark to this IP address. Finally, IP_c is selected as the probe with the most similar feature vector to the one of IP_s , according to the following normalized distance:

$$D_{ij} = \frac{1}{K} \sum_{l=1}^K |d_{il} - d_{jl}|,$$

where K is the number of landmarks providing a RTT for both IP_i and IP_j , and d_{il} is the minimum RTT between IP_i and landmark l . When D_{ij} is small, we assume that IP_i and IP_j are close to each other. In the evaluations section, we select 20 landmarks for each IP_s .

III. EVALUATION

In this section, we evaluate the performance achieved by the two probe selection approaches used by DisNETPerf, namely SL and LM approaches. Given that, in this paper, we rely on RIPE Atlas, we firstly provide a brief study on the characteristics of the RIPE Atlas measurement platform (Sec. III-A). As described in Sec. II, both the SL and the LM probe selection approaches depend on the distribution of the available RIPE Atlas probes. We thus shed some light on how distributed the probes in current RIPE Atlas deployment are.

Also, given that DisNETPerf is mainly needed in regions where Internet paths are not symmetric – even if this is not exclusive, as explained in [1] –, we perform an analysis on

Table I
TOP 5 ASes W.R.T. NUMBER OF HOSTED PROBES.

# probes	% probes	AS	Org.
348	4%	AS7922	Comcast
299	3%	AS6830	Liberty Global
232	2.5%	AS3320	Deutsche Telekom
228	2.5%	AS12322	Free SAS
148	1.6%	AS3215	Orange




Figure 2. Number of probes hosted per AS. The distribution of probes per AS is highly concentrated on small values, and for more than 96% of the ASes covered by RIPE Atlas, there are at most 10 active probes. Still, some ASes host a large share of probes, with more than 300 boxes in some cases.

the path asymmetry in specific regions of the world, using the RIPE Atlas framework for conducting the measurements (Sec. III-B).

Finally, we evaluate DisNETPerf, and show that the tool yields very accurate results under certain evaluation scenarios (Sec. III-C).

A. Study of the RIPE Atlas Framework

To gain a better understanding of the RIPE Atlas framework, we study the distribution of their probes across the network. Approximately 9,200 connected boxes are hosted by about 3,300 different ASes all over the world, while a significant part of the probes are located in Europe and the United States. AS7922 (Comcast Cable Communications, Inc.) is the one including the highest number of probes. Indeed, we found 348 probes connected through this AS, which corresponds to about 4% of all the active boxes. Table I indicates the five ASes anchoring most of the connected RIPE Atlas probes. RIPE Atlas boxes are scattered across a large number of ASes, and the majority of ASes only host a limited number of probes: in 96% of the ASes, less than 10 probes are actively contributing to the RIPE Atlas network. It is also worth mentioning that the top five ASes – controlled by large ISPs – host 13.6% of all the available probes. Fig. 2 underlines this unequal repartition among ASes, depicting the distribution of probes per AS. As a first conclusion, we expect that the performance of DisNETPerf when using RIPE Atlas would be much better in EU countries and major ISPs, where a large number of probes are available.




Figure 3. Path symmetry analysis in different regions at the AS level. Even at this level, the prevalence of path asymmetry is notorious at all the analyzed scales. As expected, highly connected and dense regions such as the EU and the US show the highest degree of asymmetric paths. The intercontinental asymmetry suggests a highly dense network connecting the US and the EU.

B. Analysis of Path Asymmetry

We study next the path asymmetry between a set of RIPE Atlas probes in several areas of the world. A path is considered as being asymmetric if the path taken from probe A toward probe B is distinct from the one from B to A . Similarly to Hu and Steenkiste [19], we define path similarity as the ratio between the number of shared links among both paths and the total number of links, using the *Route Similarity Index* (RSIM), a value in the interval $[0, 1]$, where 0 means completely different paths and 1 means exactly the same path. More precisely, for each pair of probes A and B , we compute the RSIM of the path leading from A to B , and the opposite path from B to A to assess path asymmetry. Given a generic Internet path from IP_A to IP_C , and another path from IP_B to IP_C , we define the RSIM index as:

$$RSIM(IP_A, IP_B, IP_C) = \frac{2 \times C_{links}(IP_A, IP_B, IP_C)}{T_{links}(IP_A, IP_B, IP_C)},$$

where C_{links} refers to the links shared in common by both $IP_A \rightarrow IP_C$ and $IP_B \rightarrow IP_C$ paths, and T_{links} to the total number of links. A high RSIM indicates a high similarity between the considered paths. To evaluate the similarity between a forward path $IP_A \rightarrow IP_B$ and a reverse path $IP_B \rightarrow IP_A$, we simply reverse the latter and compare it to the former, using RSIM. A path pair is thus considered as being asymmetric if and only if the computed RSIM is strictly lower than 1.

In the study, we evaluate path asymmetry country-wise within France and the US, continent-wise within Europe and South America, and intercontinental-wise between Europe and the US. We randomly select 100 RIPE Atlas boxes in the country- and continent-wise scenarios (50 for the intercontinental-wise analysis) and run `traceroute` measurements between them, analyzing the path symmetry at the AS level. We focus on the analysis at the AS level for practical reasons: AS routing represents the most relevant level for our purposes of detecting which AS on the Internet path from servers to customers might be responsible for certain performance degradation, and thus do not take into account the intra-AS routing. In addition, analyzing paths at the IP

level is subject to inaccuracies when comparing a forward to its corresponding reverse path. For instance, `traceroute` might overestimate load balancers on the paths, and thus provide measurement results indicating multiple IP paths whereas they are actually the same on the router level [20]. Another source of inaccuracy is the problem of zero-TTL forwarding [21], which implies that mis-configured routers potentially forward packets whose TTL is equal to zero, leading to a wrong IP path inference.

In this study, as well as in the evaluation presented in Sec. III-C, IP2AS mapping is done through the database provided by MaxMind [22] – note that, even if MaxMind is less accurate than other IP2AS mapping solutions, we have verified that, in our measurements, MaxMind provides the same results as other solutions such as Team Cymru (<http://www.team-cymru.org/>).

Fig. 3 depicts the obtained results at the AS level, for the three proposed scenarios. The number of asymmetric path pairs is higher in larger areas. Indeed, only about 10% of the evaluated AS paths in France are asymmetric (fully or partially) whereas this number increases to approximately 45% when considering the whole European continent. However, the prevalence of path asymmetry is notorious at all the analyzed scales. As expected, highly connected and dense regions such as the EU and the US show the highest degree of asymmetric paths. The intercontinental asymmetry suggests a highly dense network connecting the US and the EU. This brief study demonstrates that the phenomenon of path asymmetry is far from negligible at the Internet scale, and that a dedicated tool which can accurately measure the reverse path such as DisNETPerf is highly needed to track path performance issues on server-to-customer paths.

C. DisNETPerf Evaluation

Let us now focus on the evaluation of the accuracy provided by the probe selection approaches used by DisNETPerf. We say that IP_c is a *good* probe with respect to IP_s and IP_d if the path from IP_c to IP_d is highly similar to the path from IP_s to IP_d . As in Sec. III-B, we use RSIM as path similarity metric, computing $RSIM(IP_c, IP_s, IP_d)$. In this




Figure 4. Probe selection evaluation, based on RSIM. There is a significant difference between groups at the AS level, and the case of same AS co-location results in near optimal results.




Figure 5. Difference between the RSIM of selected probe and highest observed RSIM. The most significant differences are observed at the AS level.




Figure 6. Number of probes presenting the highest RSIM. In most of the cases, we observe less than 50 such probes, which is influenced by the characteristics of the RIPE Atlas framework.

evaluation, we consider links at multiple granularities; in particular, we consider links at the AS level, at the PoP level, and the IP interface level. IP2AS mapping is achieved through MaxMind’s database already referred to in Sec. III-B; IP2PoP mapping is done through the datasets made available by iPlane [23].

Fig. 4 presents the results showing the applicability of DisNETPerf in terms of path similarity. The goal is to assess whether the probe selection approaches select probes with the highest path similarity to the one we want to actually monitor. We use RIPE Atlas probes as source and destination (i.e., IP_s and IP_d) so as to compute the real path (i.e., the ground truth) between servers and customers. We randomly select 300 RIPE Atlas source probes IP_{s_i} , and consider a single fixed destination probe IP_d . For each source IP_{s_i} , we run

DisNETPerf to locate the closest probe IP_{c_i} , obtain both the ground-truth path $IP_{s_i} \rightarrow IP_d$ and the DisNETPerf path $IP_{c_i} \rightarrow IP_d$, and compute $RSIM(IP_{c_i}, IP_{s_i}, IP_d)$.

We compute the RSIMs at the AS level, PoP level, and IP level, and plot the resulting distributions. Results are reported for the two probe selection approaches and for two different groups, the first one in which IP_{c_i} and IP_{s_i} are located in the same AS (black lines), and the second one in which IP_{c_i} is located in a neighbor AS (gray lines). SL/LM-approach results are depicted by plain/dotted lines. There is a significant difference between groups at the AS level, and the case of same AS co-location results in near optimal results, with more than 80% of the tests resulting in a perfect path matching (i.e., RSIM equals to one). Nevertheless, we observe that about 40% of the tests carried out for the SL approach and 45%




Figure 7. Minimum RTTs from candidate probes towards the target. A non-negligible fraction of probes are located near the target, at less than 20 ms for 20% of the candidate probes.

of the tests performed for the LM approach yield a RSIM index ≥ 0.5 . Note that the most relevant segment of the path to monitor for troubleshooting purposes is the one closer to the customer (where problems generally occur), thus a RSIM ≥ 0.5 is actually very good.

Differences between both groups are less relevant for paths at the PoP and IP level. At the PoP level, the RSIM index computed using the SL and LM approaches is ≥ 0.5 for about 50% of the tests when IP_{c_i} and IP_{s_i} are in different ASes, and for more than 80% of the tests when they are in the same AS. We observe the highest rate of low RSIM indexes at the IP level: only about 20% of the tests yield a RSIM index ≥ 0.5 when IP_{c_i} and IP_{s_i} are not co-located in the same AS. However, IP-level paths are generally less relevant, and we plan to evolve to router-based paths, clustering interfaces belonging to the same router (IP aliases) [24]. In general, we note that the results for both SL and LM approaches are comparable and highly similar. Moreover, we observed that probes selected by DisNETPerf using both approaches generally correspond to paths with the highest similarity to the ground-truth ones: in more than 80% of the performed tests for the two techniques, $RSIM(IP_{c_i}, IP_{s_i}, IP_d)$ results in the highest RSIM among all the selected candidates when considering the AS level, which, as we said, is the most relevant level for our purposes.

In the case where DisNETPerf does not choose a probe among the ones presenting the highest RSIM, we compute the difference between the RSIM of IP_{c_i} and the best one. Fig. 5 summarizes the obtained results for the SL and LM approaches at the AS, PoP, and IP level. We notice that the most significant differences are observed at the AS level: considering the SL approach, the difference is above 0.5 for 54% of the measured paths whereas this is the case for approximately 57% of the selected probes in the context of the LM approach. The RSIMs at the PoP and IP level are much more similar. Indeed, more than 75% and 95% of the selected IP_{c_i} yield a RSIM with a difference below 0.5 with respect to the highest RSIM respectively. This can be explained by the paths length. They are actually much shorter at the AS level, and thus one single link difference becomes more relevant when calculating and

comparing RSIMs at this level.

Another point we consider is the number of candidate probes with the best path. Fig. 6 shows that this number is in most of the cases lower than 50. The reason for this can be two-fold. Firstly, as exposed in Sec. III-A, RIPE Atlas boxes are quite well distributed over the world, but most of the ASes only host a fairly restricted amount of probes. This implies that, in some cases, not many probes can be selected as candidate probes. Secondly, identical paths are rarer when the analysis is done at more fine-grained levels (the PoP and IP level, for instance).

Finally, we gain some insight on the distance in terms of delay of the candidate probes with respect to IP_{s_i} for the SL approach. For this purpose, we record, for each candidate probe, its minimum RTT towards IP_{s_i} . Fig. 7 depicts the distribution of the measured minimum RTTs. We note that a non negligible fraction of probes are located near the target: about 20% of the candidate probes are less than 20 ms away from their target.

IV. DISCUSSION

The results obtained by the presented probe selection approaches are encouraging, especially as the best results are obtained at the AS level, which is the most relevant level for our monitoring and performance diagnosis purposes. However, we have observed less promising results when moving to more fine-grained levels, such as the PoP and IP level. We intend to consider more complex probe-selection approaches to improve results. Candidate techniques are the so-called *Routing State Distance* by Gürsun et al. [25], and the combined topology and routing techniques presented in [26]. Even though these techniques would probably provide more accurate results, we foresee a much heavier computational and instrumental complexity than in the considered SL and LM approaches. We are therefore investigating whether the gain in precision is worth switching to a computationally intensive algorithm.

To conclude with the study, we provide some discussion on the applicability of DisNETPerf in practice for continuous reverse path performance monitoring and analysis. Traditionally, either active or passive measurements are considered for monitoring large-scale networks such as the Internet. While active measurements allow to know exactly the workload injected by the application into the network, the passive measurements can offer a more detailed view of transport and network layer impacts. The combination and correlation of active and passive measurements for troubleshooting purposes at the Internet scale has been almost neglected so far. DisNETPerf offers the possibility of running continuous reverse path performance measurements, for example from servers provisioning relevant services to the customers of a certain ISP. As a preliminary example analysis, we have been running DisNETPerf for several days, monitoring the reverse paths from a set of 50 Google servers provisioning YouTube traffic in Italy, to a set of ISP customers. More precisely, DisNETPerf firstly elected the closest probes to the corresponding YouTube servers, and issued `traceroute` measurements towards an entry point at




Figure 8. Results obtained by monitoring 50 Google servers in Italy with DisNETPerf; each point in both plots corresponds to one measurement. While the observed Internet paths tend to remain stable in terms of number of traversed ASes, the RTTs measured by DisNETPerf show high variability for some specific reverse paths, suggesting both transient as well as periodical congestion events, in some cases during several hours.

the edge of an ISP where customers connect to the Internet every 30 minutes, during one week.

Fig. 8 depicts the collected results, both in terms of number of traversed ASes (top) as well as RTT (bottom). While the observed Internet paths tend to remain stable in terms of number of traversed ASes, the RTTs measured by DisNETPerf show high variability for some specific reverse paths, suggesting both transient as well as periodical congestion events, in some cases lasting several hours. For example, some reverse paths show a latency increase higher than 300%, potentially resulting in serious performance degradation during such events. In some cases, major RTT variations span several hours, leading to overall RTTs as high as one second, thus clearly suggesting underlying performance issues. Even more, in some of the monitored paths, there is a noticeable correlation between changes in the number of traversed ASes and a strong variation on the measured RTTs, suggesting that path inflation/deflation and routing issues might also be involved in those cases. Combining the monitoring results provided by DisNETPerf with the ones of passive measurements is a promising approach to locate and redress these issues.

V. RELATED WORK

There is a very rich literature in the problem of using active measurements such as *traceroute* to understand the Internet topology [27] as well as diagnosing performance issues. In this paper, we focus on three specific aspects of this broad literature: (i) similar systems to DisNETPerf [23], [7], [28], relying on *traceroute* measurements to evaluate the performance of Internet paths; (ii) techniques for (geo)location of probes at the Internet scale [16], [26], [15], [17]; and (iii) metrics and approaches for assessing the similarity [25], [19], [29] and Internet paths (a)symmetry [2], [3], [4]. The starting point of DisNETPerf is the reverse *traceroute* system proposed by Katz-Bassett et al. [7]. They proposed a combined active measurement technique to incrementally piece together the path from a certain destination

back to the source, and test its performance using PlanetLab. While the proposed solution shows accurate results in many different evaluation scenarios, the – already mentioned – main limitation of this approach is the usage of IP-spoofing and IP Record Route Option, both being not necessarily allowed everywhere. Reverse *traceroute* is built upon a more complete and exhaustive Internet path measurement system called *iPlane* [23]; *iPlane* offers a service for predicting the performance of Internet paths by building a structural model of the Internet using *traceroute* and opportunistic bandwidth-measurements. In a nutshell, *iPlane* constructs a periodically updated sort of annotated map of the Internet to predict latency, bandwidth, capacity, and loss rates between arbitrary Internet hosts. Authors from these papers took a further step in this direction very recently, by proposing *Sibyl* [28], a query-based system capable of providing *traceroute* measurements on arbitrarily selected Internet paths. Similar to *iPlane*, *Sibyl* combines multiple historical *traceroute* measurements with prediction techniques to reconstruct measurements on segments not necessarily measured before. These papers show the relevance and (still) novelty of the addressed problem.

When it comes to (geo)location of probes in the Internet, there is also an assorted list of papers worth mentioning. Works such as [15] and [16] study the problem using a combination of techniques, including delay-based measurements, network topology measurements, and geo-constrained measurements. Madhyastha et al. [26] also tackle the same problem, for the specific case of predicting path latencies. To do so, they introduce a complex yet effective approach to locate probes at relevant locations, using a combination of Internet’s routing topology, PoP connectivity, and routing policies. One of the probe location techniques used in this paper is inspired by our previous work on network probe proximity and neighborhood models [17], where we infer proximity among nodes without measuring the latency between them, but rather relying on a set of stable landmark probes to define latency-based location

coordinates, similar to [16].

When it comes to the evaluation of path similarity at the Internet level, different papers propose multiple metrics based on routing [25], path-segments matching [19], and Bloom filters [29]. In the evaluations done in this paper, we use the notions of path similarity defined in [19], due to their simplicity and very intuitive interpretation. Finally, multiple papers have studied the problem of path asymmetry in the Internet at the routing level, performing analysis on both AS-level and link-level paths [2], [3], [4]. The overall understanding is that assuming symmetric paths at the Internet level does not hold in many cases, especially in highly dense and connected regions.

VI. CONCLUSION

Tracking the performance of Internet paths is a very important task for ISPs to diagnose the Quality of Experience of their customers when problems occur. This task is challenging, as paths are in general not symmetric, and thus standard tools such as `traceroute` do not suffice to troubleshoot paths performance degradations. In this paper, we introduced DisNETPerf, a Distributed Internet Paths Performance Analyzer, capable to measure reverse paths. Our evaluation shows that DisNETPerf computes the reverse path with very high accuracy in some scenarios. The best results are obtained at the AS level, which is the most relevant level for our purposes. Indeed, our primary goal is to pinpoint the potential guilty AS for the performance issues on the path connecting servers to customers.

We are investigating other probe location techniques previously proposed in the literature to improve the accuracy of DisNETPerf, but still, we acknowledge that the simplicity behind the currently employed techniques provides relatively high accuracy with a very fast computation, which is another element to take into account, more particularly when thinking about using DisNETPerf in near real time monitoring scenarios.

Finally, we leave for future work a deep comparison between DisNETPerf and the reverse `traceroute` system.

REFERENCES

- [1] R. A. Steenbergen, "A practical guide to (correctly) troubleshooting with `traceroute`," January 2009, nanog 45. See http://www.nanog.org/meetings/nanog45/presentations/Sunday/RAS_traceroute_N45.pdf.
- [2] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker, "On routing asymmetry in the Internet," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, November/December 2005.
- [3] V. Paxson, "End-to-end routing behavior in the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 5, pp. 41–56, October 2006.
- [4] W. Vries, J. J. Santanna, A. Sperotto, and A. Pras, "How asymmetric is the Internet," in *Proc. IFIP International Conference on Autonomous Infrastructure, Management, and Security (AIMS)*, June 2015.
- [5] P. Casas, A. D'Alconzo, P. Fiadino, A. Bär, A. Finamore, and T. Zseby, "When YouTube does not work. analysis of QoE-relevant degradation in Google CDN traffic," *IEEE Transactions on Network and Service Management*, vol. 11, no. 4, pp. 441–457, December 2014.
- [6] P. Fiadino, A. D'Alconzo, A. Bär, A. Finamore, and P. Casas, "In the detection of network traffic anomalies in content delivery network services," in *Proc. International Teletraffic Congress (ITC)*, September 2014.

- [7] E. Katz-Bassett, H. Madhyastha, V. Adhikari, C. Scott, J. Sherry, P. van Wesep, A. Krishnamurthy, and T. Anderson, "Reverse `traceroute`," in *Proc. USENIX Symposium on Networked Systems Design and Implementations (NSDI)*, June 2010.
- [8] R. Beverly, A. Berger, Y. Hyun, and k. claffy, "Understanding the efficacy of deployed Internet source address validation filtering," in *Proc. ACM Internet Measurement Conference (IMC)*, November 2010.
- [9] R. Fonseca, G. M. Porter, R. H. Katz, S. Shenker, and I. Stoica, "IP options are not an option," University of California at Berkeley, Technical Report UCB/EECS-2005-24, December 2005.
- [10] RIPE NCC Staff, "RIPE Atlas: a global Internet measurement network," *Internet Protocol Journal*, vol. 18, no. 3, pp. 2–26, September 2015.
- [11] T. Holterbach, C. Pelsser, R. Bush, and L. Vanbever, "Quantifying interference between measurements on the RIPE atlas platform," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2015.
- [12] k. claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: from art to science," in *Proc. IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*, March 2009.
- [13] PlanetLab, "PlanetLab, an open platform for developing, deploying, and accessing planetary-scale services," August 2015, see <https://www.planet-lab.org/>.
- [14] S. Wassermann, P. Casas, and B. Donnet, "Towards DisNETPerf: a Distributed Internet Paths Performance Analyzer," in *Proc. of the ACM CoNEXT Student Workshop*, December 2015.
- [15] E. Katz-Bassett, J. P. John, T. Anderson, A. Krishnamurthy, Y. Chawathe, and D. Wetherall, "Towards IP geolocation using delay and topology measurements," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2006.
- [16] V. Padmanabhan and L. Subramanian, "An investigation of geographic mapping techniques for internet hosts," in *Proc. ACM SIGCOMM*, August 2001.
- [17] Y. Liao, W. Du, and G. Leduc, "A lightweight network proximity service based on neighborhood models," in *Proc. IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*, November 2015.
- [18] The CAIDA UCSD, "AS relationships," June 2015, <http://data.caida.org/datasets/as-relationships/serial-1/20150601.as-rel.txt.bz2>.
- [19] N. Hu and P. Steenkiste, "Quantifying Internet end-to-end route similarity," in *Proc. Passive and Active Measurement Workshop (PAM)*, April 2006.
- [20] P. Marchetta, V. Persico, A. Pescapé, and E. Katz-Bassett, "Don't Trust `Traceroute` (Completely)," in *Proc. of the ACM CoNEXT Student Workshop*, December 2013.
- [21] 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 Internet Measurement Conference (IMC)*, October 2006.
- [22] MaxMind, "Geolocation and online fraud prevention from MaxMind," June 2015, see <http://download.maxmind.com/download/geoip/database/asnum/GeoIPASNum.dat.gz>.
- [23] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An information plane for distributed services," in *Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 2006.
- [24] K. Keys, "Internet-scale IP alias resolution techniques," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 1, pp. 50–55, January 2010.
- [25] G. Gürsun, N. Ruchansky, E. Terzi, and M. Crovella, "Routing state distance: A path-based metric for network analysis," in *Proc. ACM Internet Measurement Conference (IMC)*, November 2012.
- [26] H. V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani, "A structural approach to latency prediction," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2006.
- [27] B. Donnet and T. Friedman, "Internet topology discovery: a survey," *IEEE Communications Surveys and Tutorials*, vol. 9, no. 4, pp. 2–15, December 2007.
- [28] I. Cunha, P. Marchetta, M. Calder, Y.-C. Chiu, B. Machado, A. Pescapé, V. Giotsas, H. V. Madhyastha, and E. Katz-Bassett, "Sibyl: a practical Internet route oracle," in *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, March 2016.
- [29] B. Donnet, B. Gueye, and M. A. Kaafar, "Path similarity evaluation using Bloom filters," *Computer Networks (COMNET)*, vol. 56, no. 2, pp. 1389–1286, February 2012.