On the Stability and Diversity of Internet Routes in the MPLS Era

Zakaria Al-Qudah[†*1] , Ibrahim Jomhawy [†], Mohammad Alsarayreh[†] and Michael Rabinovich[‡]

[†] University of West Florida, Pensacola, FL.

[*] Yarmouk University, Irbid, Jordan.

[‡] Case Western Reserve University, Cleveland, OH.

Abstract

Stability and diversity of end-to-end routes are key properties of Internet that have a large effect on the design of networks and networked systems. As the Internet evolves and deploys new technologies, it is important to re-assess these properties in the face of the new realities. This paper evaluates the stability and diversity of Internet routes with emphasis on the impact of the widely deployed Multi-Protocol Label Switching (MPLS). In particular, we analyze two datasets: (1) high frequency (once per minute) traceroutes originating from a number of PlanetLab hosts to a set of destinations distributed around the Internet (random hosts extracted from the Alexa list of top-1M websites), and (2) traceroutes originating from a large number of Ripe Atlas probes to a number of hosts (mostly root DNS servers). We find that Internet routes are significantly less stable than previously reported, at least according to one common metric (the route persistence), and much more diverse. At the same time, these more frequent route changes do not translate to high variability of route-trip time (RTT) delays between hosts, as RTTs tend to stay similar across route changes. Notably, with regard to diversity, we show that higher diversity is likely due not to a real change in the Internet but rather to underestimation by previous studies. Finally, our assessment of MPLS role in increased path instability produced inconclusive results as MPLS appears to be a significant contributor in the Ripe Atlas data set but not in the Plant Lab dataset.

Keywords: Internet path stability, MPLS, persistence, prevalence, dominant route, diversity and RTT.

1. Introduction

The end-to-end characteristics of Internet routing have been an active field of study for decades. Several Internet path properties have been measured

[1]Corresponding author. Email: zalqudah@uwf.edu

including end-to-end route stability and diversity [1–3]. Previous studies have established that Internet paths are generally stable with routes between two end-points lasting for several hours or even days [1].

These findings have been very influential as they have been relied upon in designing experimental apparatus for measurement studies (e.g., [4]) as well as in designing systems (e.g., [5]). In particular, researchers and system designers rely on these findings to probe paths with low frequency to reduce measurement overhead assuming that the stability of these paths would still result in a valid approximation of the Internet paths at any given time. Furthermore, the stability of Internet paths (or the lack thereof) may affect the end-to-end performance of applications. For example, switching paths in the middle of a Transmission Control Protocol (TCP) session may result in a significant change in Round Trip Time (RTT). TCP may need time to learn the new path characteristics, leading to sub-optimal communication performance.

Updating our understanding of the current end-to-end Internet path properties is crucial since the Internet continues to evolve. Several new technologies that could potentially affect these properties have been deployed including multi-homing and Multi-Protocol Label Switching (MPLS). In particular, MPLS is widely deployed today for traffic engineering, as a tool for service providers to better utilize their resources and to offer better services to their customers [6, 7]. While there are studies that explicitly consider the impact of load balancing on Internet route stability [2], we are unaware of prior studies that have considered the impact of traffic engineering based on MPLS on the stability of these routes. Our present study addresses these gaps. In particular, our study, has the following goals:

- Re-assessing the stability and diversity of the Internet routes regardless of the cause of route changes (routing event, load balancing, traffic engineering, etc.)
- Investigating the variability in Round Trip Time (RTT) that is associated with a path change between a source-destination pair.
- Evaluating the impact of MPLS deployments on the stability, diversity, and RTT variability of end-to-end routes in the Internet.

To achieve goals, we analyze two datasets collected from two sources. The first dataset is collected by conducting a large-scale active measurement via PlanetLab [8]. We use several PlanetLab nodes as vantage points to monitor the paths to a large number of hosts on the Internet through traceroute measurements. We conduct our traceroute measurements on a short time scale (one minute) for a full day per source-destination pair. The second dataset has been downloaded from the Ripe Atlas platform. It involves traceroutes conducted as part of the built-in measurements that the platform conducts continuously, originating from all probes to mostly root DNS servers. The time spacing between these traceroutes is 30 minutes.

Our contributions include the following key findings.

- We find that Internet routes are significantly less stable than previously reported, at least according to the route persistence metric, which is a measure of how often a path between a host-pair changes. To this end, our findings could indicate a real change in the Internet or that previous studies underestimated path persistence. Particularly, previous studies probed Internet's paths at low frequency which could cause them to miss some path changes. Our results show less clear trends with respect to route prevalence, another commonly used metric of route stability reflecting how often the most common path between a host-pair is taken.
- We also find significantly higher Internet route diversity than previously reported, which is a measure of the number of distinct paths used by host-to-host communication over time. However, unlike route persistence, we show that the lower diversity we observed is likely not an indication of a true change in the Internet characteristics but due to previous studies probing at inadequate frequency, and thus missing many existing paths.
- While we find Internet routes to become less stable, we find that route changes usually don't lead to significant changes in round-trip time delays between hosts. While existing studies evaluate the general variability of RTT between a host-pair or within a TCP connection, we measure the variability of RTT when the path between a host-pair changes, which allows us to understand the implications of path changes on the stability of RTT.
- Contrary to expectation, we did not find conclusive evidence that MPLS contributes significantly to the decrease in path stability. While we do find significant impact of MPLS on path instability in the Ripe Atlas data set, no such contribution is found in the PlanetLab data set.

Our results further emphasize the importance of sound measurement practices. In particular, our seemingly large data sets disagreed on some properties. Thus drawing a general conclusion is usually not possible from a single data set. Additionally, parametrization of measurement experiments can significantly impact the conclusions drawn from these experiments. Probing at a high frequency in our case revealed a much higher route diversity than possible with low frequency probing. Therefore, we recommend that measurement studies should explore the sensitivity of their results to the choice of experiment parameters whenever possible to improve applicability of the drawn conclusions. This paper extends our preliminary investigation in [9]. The present paper adds key new insights by analyzing a second independent dataset and by examining the effect of probing frequency on the outcome of the analyses.

2. Related Work

2.1. Path Characteristics

The characteristics of the Internet have been studied extensively. In [1], the stability and symmetry of end-to-end routing were analyzed using 40,000 end-to-end routes between 37 Internet sites. The study found that Internet paths are relatively stable.

The results of [1] were re-assessed by Schwartz et al. [3] and Cunha et al. [2]. In [3] a systematic evaluation of the diversity, stability and symmetry of Internet routes was performed. The authors employed more than 100 distributed vantage points in their evaluation. The Internet stability was found to be less than what is found in Paxson's study [1]. However, the conclusion was that the overall results of [1] remain valid. The authors in [2] measured and characterized the end-to-end route dynamics in the presence of load balancing. They found that when removing route changes that are attributed to load balancing, the Internet stability has not changed significantly when compared to Paxson's results.

Our experiments differ from these studies in various ways: (*i*) We assess the stability of Internet paths on short time-scale by probing at 1-minute intervals in one data set and 30-minutes intervals in the other data set whereas other studies probed at intervals of hours and days. Probing at short-time scale can reveal high frequency route flaps which would be missed by probes at coarser time-scales. In fact, we do find much higher path diversity than reported in [3], and provide evidence that this difference is due to the low frequency probing of the previous study, not to a real increase in route diversity in the Internet.(*ii*) We employ broader, more realistic definition of path equality as explained in Section 3.2. (*iii*) We investigate the impact of MPLS on path stability and diversity.

The stability of the Internet and the wide area backbone failures were studied experimentally in [10] using observation of five of the largest U.S. Internet Exchange Points (IXPs) over three years. The authors observed a slightly higher chances of route failure and fail-over (and thus lower stability) than [1]. The path diversity of Internet topology, both inside and across autonomous systems, was analyzed in [11]. Our study focuses on the actual routes experienced by probing packets, rather than on the set of all potential paths that constitute the Internet topology.

2.2. MPLS

The prevalence and characteristics of MPLS in the open Internet were evaluated in [6]. The results indicate that the total number of MPLS tunnels observed varied with time, and 25% of paths observed had at least one MPLS tunnel. In addition to that, MPLS deployment increased over the time.

The authors in [12] shed the light on the availability of obscured tunnels. They develop methods for revealing the deployment and characteristics of obscured MPLS tunnels from traceroute data to assess the prevalence of MPLS tunnels and study their characteristics. The results estimate that 30% of paths had at least one MPLS tunnel, and 90% of the MPLS tunnels were short (had less than five hops).

A number of solutions have leveraged MPLS for dynamic traffic engineering [7, 13–15]. These solutions have the potential of decreasing the stability of Internet routes due to MPLS tunnels. We are investigating this potential impact in the present paper. We refer the reader to [16] for a tutorial on the use of MPLS in traffic engineering.

2.3. RTT

Previous studies that considered the general variability of RTT between a host-pair arrived at conflicting conclusions. The authors in [17] measured the HTTP latency as an approximation of RTT and found it to be generally stable. Another study that considered RTT in the context of the "buffer bloat" problem [18] also found RTT to be generally stable. On the other hand, the authors in [19] found significant RTT variations when they monitored the paths between a client and 5K+ servers. Similarly, several studies analyzed RTT within a TCP connection or across TCP connections between a host-pair found a significant variability in RTT [20–22]. In this study, we are not interested in the general stability of RTT. Instead, we are looking to understand the implications of a path change between a host-pair on the stability of RTT.

3. Methodology

3.1. Data Sets

3.1.1. PlanetLab

We collected one of the datasets for this study using 88 PlanetLab [8] nodes as vantage points. We assign for each PlanetLab node a set of 1000 IP addresses for top web servers from the QuantCast list of top websites[2] and monitor the path between the PlanetLab node and the corresponding 1000 IP addresses. Therefore, we monitor 88K paths in total.

For each IP address, the corresponding PlanetLab node issues a traceroute once every 60 seconds for 24 hours, a total of 1440 traceroutes. We used Scamper tool version scamper-cvs-20141211a [23] with the traceroute option that implements ICMP extensions for MPLS. There are two features that routers can support in relation to these extensions. First, they can implement the *ttl-propagate* feature, which involves the first router in the MPLS tunnel copying the IP-level time-to-live (TTL) to the MPLS-level TTL. Second, they can adopt RFC 4950 and embed the MPLS information in the returned ICMP time-exceeded message. Depending on which of these two features a router implements, there are four ways the MPLS tunnel is reflected in traceroute. When both features are implemented, the tunnel and its internal hops are revealed and marked as MPLS hops (explicit tunnel). If the first router in the tunnel enables ttl-propagate but RFC 4950 is not enabled, the internal hops of the tunnel will be revealed. However, they will not be marked as MPLS hops (implicit tunnel). If the MPLS routers enable RFC 4950, but the first router in the tunnel does not enable ttl-propagate, this tunnel will show as a single MPLS hop (the last hop

[2]The current URL for the QuantCast list is https://ak.quantcast.com/quantcast-top-sites.zip, retrieved on 12/30/2019. The present study used this list as obtained on November 29, 2014.

Item	PlanetLab	Ripe Atlas
Source-destination pairs	88K	$\sim$ 95K
Traceroutes performed	$\sim$ 126M	$\sim$ 4M
Traceroutes reached their destinations	$\sim$ 80M	$\sim$ 3.5M
Distinct IPs	$\sim$ 141K	$\sim$ 70K
Distinct ASes	5686	3628

Table 1: Summary of the data sets

in the tunnel) in trace route (opaque tunnel). Finally, when routers implement neither ttl-propagate nor RFC 4950, the tunnel will be completely invisible in traceroute (invisible tunnel). See [12] for a full discussion of these cases.

The traceroute timeout is set to one second (i.e., traceroute waits for the ICMP echo response for one second before announcing the hop unresponsive). The choice of one second timeout simplifies the management of the experiment. At the same time, previous measurement studies show that the RTT of the vast majority of Internet paths is less than one second [24], and therefore we believe our choice of this parameter does not affect our results. We configure Scamper to issue one traceroute attempt per hop instead of the three-attempts default. Using three attempts would effectively mean probing at an interval that is shorter than our intended 60-second interval. The Scamper tool is configured to abort probing if it receives no response from five consecutive hops to avoid looking like an attack traffic. This experiment was conducted on August 24 2015.

When manually examining the results, we noticed that 8 vantage points (PlanetLab nodes) had their traceroutes always dropped only few hops away from them which makes them useless for our analysis. We eliminate the results from these 8 vantage points from further analysis. These were the only vantage points with this behavior.

3.1.2. Ripe Atlas

The Ripe Atlas data set contains traceroutes from all Atlas probes (nearly 9K probes) to 17 destinations (DNS root servers and ripe.net hosts). This experiment issues a UDP traceroute from each of the active probes once every 30 minutes. Each hop is sent three probes. We analyze the results of the first probe to be consistent with the PlanetLab data set where we issue one probe per hop. The timeout for these probes is 1 second. While this experiment has been running since 2010, we only analyze the results of August 24, 2015 to be consistent with the PlanetLab experiment.

3.1.3. Data Sets Summary

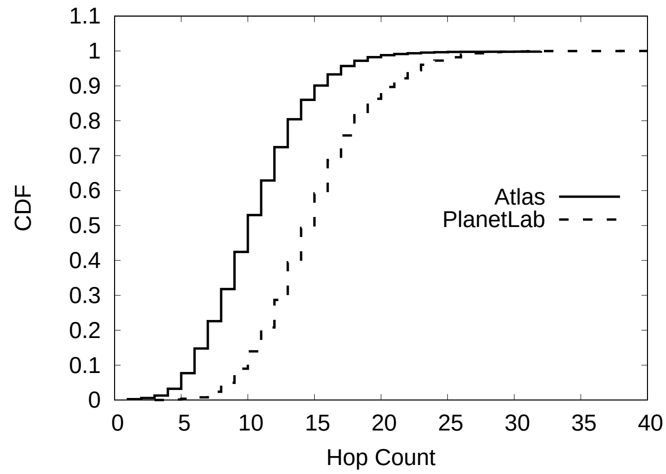

Figure 1: Path lengths in PlanetLab and Ripe Atlas data sets

Table 1 summarizes our data sets. The PlanetLab data set contains paths traversing the total of 5686 distinct autonomous systems and 141K distinct IP addresses. The Ripe Atlas data set involves 3628 distinct autonomous system with 70K distinct IP addresses. Therefore, we believe our data sets represent a significant slice of the Internet. In terms of temporal coverage, our data sets represent a full day which is enough to cover any potenial diurnal effects. Obviously, our study is not designed to offer insights on the longitudinal behaviour of Internet routes in terms of stability and diversity.

Figure 1 shows the CDF of the path length (hop count) for the two data sets. As shown, the PlanetLab data set exhibits significantly longer paths than Ripe Atlas. We believe that this might be due to the wide spread anycasting of the root DNS servers.

We note that the two data sets reflect very different environments. The PlanetLab vantage points are mainly academic networks and the targets are popular websites. On the other hands, the Ripe Atlas vantage points involves various types of networks including residential ones. The targets in the Ripe Atlas data set are widely anycasted. Furthermore, the probing frequency is once per minute for the PlanetLab data set and once per 30 minutes for the Ripe Atlas data set.

3.2. Path Equality

To measure the stability and diversity of Internet paths, we need to clearly define the equality of the paths between a source-destination pair. This is necessary since some nodes do not respond to traceroute probes (which results in traceroute displaying '*' for such probes). Previous studies have handled this issue in various ways. In [1] for example, paths that include a non-responding node have been removed from the analysis. While such an action might not have affected the results in [1] since the percentage of these cases was small at the time of that study, we cannot do the same thing in our study because a significant number of nodes in the Internet today do not respond to traceroute probes. Indeed, in contrast to [1], where eliminating paths with non-responsive hops resulted in discarding less than 10% of all path measurements (reducing the number of measurements used for the analysis from 35,109 to 31,709), we find that we would have to remove 66% and 32% of the paths in our PlanetLab and Ripe Atlas data sets, respectively, if we were to discard paths with non-responding hops. Obviously, discarding such paths would skew our results to reflect only fully responsive paths, and this skew would affect our study to a much greater degree than [1] because the results would represent much smaller fraction of the Internet paths. We note that our finding agrees with the common knowledge that today's hosts are less likely than in the past to return ICMP messages as prescribed by the ICMP protocol. As other evidence, [25] found 53.2% of live hosts they sampled did not respond to pings, and neither did 21-27% of top-50K webservers in [26].

In [3], the authors regarded a non-responding host as a wild-card when comparing the paths between a source-destination pair. That is, the path "A * C" has been considered to be equal to "A B C". In this study, we introduce three definitions of equality for the paths between a source-destination pair.

- Strict: according to this definition, two paths are considered equal only if the hops along the two paths are exactly the same. That is, the path "A * C" is considered to be not equal to the path "A B C" nor equal to the path "A * C". In other words, this definition considered an unknown hop unequal to any other hop whether known or unknown.

- Mid: according to this definition, unknown hops in the same position are regarded as the same hop. However, a known hop is considered to be unequal to an unknown hop. For example, when we encounter two paths in the form "A * C" and "A * C" we consider them to be equal. However "A * C" is considered unequal to "A B C".

- Loose: This definition is similar to the one used in [3]. In other words, an unknown hop is considered to be equal to another unknown hop in the same position in the path, and is considered to be equal to a known hop in the same position. That is, "A * C" is equal to "A * C" and is equal to "A B C".

With the strict definition, our estimation on path equality will be the most conservative. This is because we regard two unknown hops in the same position as unequal despite the possibility of them being equal (e.g., the same hop did not respond in both cases). Therefore, stability estimation based on this definition represents a lower bound. On the other hand, the loose definition regards unknown hops in the same position as equal despite the fact that the traceroute probe might have passed through a different unknown hop in the second probe than in the first probe. Furthermore, the loose definition regards an unknown hop to be equal to a known hop in the same position which might not be true. We believe the mid definition to be the most realistic but we consider the strict and loose definitions to provide bounds to our results and, in the case of loose definition, to match the methodology of [3]. Because as we mentioned, PlanetLab paths are much more likely to contain unknown hops than Ripe Atlas paths (66% vs. 32%), the bounds provided by strict and loose definitions are typically extremely wide for the PlanetLab results. We mainly use the mid definition for drawing conclusions about the path characteristics in those cases.

3.3. Metrics

We use metrics from previous studies for assessing stability and diversity of Internet paths. We borrow the two metrics used by [1] and [3] to evaluate the stability of Internet paths: persistence and prevalence. We further employ the number of observed paths between a host pair, first proposed in [27] and used to characterize Internet routes in [3], as a metric to assess the path diversity on the Internet. We specify these metrics more precisely below.

3.3.1. Persistence

This metric reflects the amount of time it takes for an observed path between a source-destination pair to change. In our measurements, the persistence of the path between two end-points is calculated as the ratio of "no path change" events to the maximum possible "no path change" events. For instance, in our PlanetLab data set, we perform 1440 measurements per source-destination pair. Therefore, we have a total of 1439 possible "no path change" events. We measure with this metric the actual "no path change" events that occur out of these 1439 possible events. We calculate this metric over all traceroutes whether the traceroute reached the final destination or not. This is because a change in the subpath indicates a change in the total path. However, even when we regard two subpaths to be equal, the paths could be different in reality since a change could occur in the unknown part of the path. Therefore, our persistence estimation is conservative and represents only a lower bound.

3.3.2. Prevalence

This metric evaluates how often a certain path between two end-points is taken, as the fraction of times this path has been observed over the total number of measurements. In particular, we measure the prevalence of the path that was taken the most (called the dominant path). For this metric, we only consider probes that reached their final destination.

To show the relationship between persistence and prevalence, consider the following examples. Suppose the sequence of paths observed in four traceroute measurements between a source destinations pair was P1, P1, P1, P1. In other words, the probing packets traveled through the same route in all 4 observations. The persistence in this example equals to 1. The prevalence of the dominant route also equals to 1 meaning that the dominant route was taken in all observations. Consider, however, if the sequence of observed paths was P1, P2, P1, P2. The persistence in this case equals to 0 meaning that the route has changed for every probe. The prevalence however equals to 0.5 meaning that the dominant route (either P1 or P2 could be used as the dominant in this case since they were used equally often) was taken in half of the observations. Therefore, prevalence and persistence capture two different (although not completely independent since low prevalence entails low persistence) aspects of path stability.

3.3.3. Route diversity

As mentioned, we use the number of distinct paths observed between a host-pair over a given period of time, as the metric to characterize the diversity of Internet paths. Route diversity reflects different aspects of Internet connectivity from the two route stability metrics. Indeed, the same number of changes (persistence) may involve different number of distinct routes (diversity), and the prevalence – unlike diversity – says nothing about the variability of non-dominant paths. Still, this metric is not completely independent from the path stability metrics since persistence or prevalence equal to 1 entails the diversity also equal to 1.

3.4. MPLS Impact

To assess the impact of MPLS on the stability and diversity of Internet paths, we need to classify source-destination pairs into those that communicate over paths involving MPLS tunnels ("MPLS pairs") and those that don't ("No-MPLS pairs"). We examine several definitions for the MPLS source-destination pairs: the pair is an MPLS pair (*i*) if at least one traceroute from the source to the destination includes an MPLS tunnel and (*ii*) if X% of all traceroutes from the source to destination (including those that produce the same route) contain an MPLS tunnel where x = 30, 50, 70, 90, and 100%. A traceroute is considered to include an MPLS tunnel if at least one hop reported an MPLS label. According to the classification in [12], this covers only explicit and opaque tunnels, and therefore is only a conservative estimate on MPLS tunnels. In other words, the No-MPLS pairs may in fact include routes with unobserved MPLS tunnels (involving invisible or implicit tunnels). Table 2 shows the percentage of MPLS pairs for various values of the threshold X. As shown in the table, the choice of the parameter X does not seem to affect the classification significantly (unless we require every single traceroute to include an MPLS tunnel before we classify the pair as an MPLS pair). Therefore, we choose 50% of the traceroutes as the classification threshold. As mentioned, our detection of MPLS pairs is conservative, in that some non-MPLS pairs may communicate over undetected MPLS tunnels.

We note that Ripe Atlas data sets involves significantly lower percentage of MPLS pairs. One possible explanation is the shorter path length: shorter paths are less likely to encounter an MPLS tunnel.

Probes %	1 Probe	30%	50%	70%	90%	100%
PlanetLab- MPLS host pairs	45.0%	42.5%	42.0%	41.0%	40.0%	27.0%
Ripe Atlas- MPLS host pairs	17.2%	15.7%	15.0%	14.4%	13.9%	12.8%

Table 2: Classification of MPLS source-destination pairs for various values of the threshold X.

4. Results

4.1. Route Persistence

We first evaluate the persistence of Internet routes over the two data sets. Figure 2 shows cumulative distribution of route persistence for both data sets under our definitions of path equality.

Both data sets demonstrate a sharp decrease in the persistence of Internet routes from the results reported in Paxson's study [1]. Indeed, Paxson found that two-thirds of the measured routes persisted for days or weeks, which means a persistence value of one would have been obtained for those host-pairs when monitored for one day as our study does. However, for the most realistic mid definition of path equality, close to 46% of the source-destination pairs in the PlanetLab data set (Figure 2a) and 63% in the Ripe Atlas data set (Figure 2b) had 0.6 or less persistence[3]. In fact, under strict definition of path equality, nearly 50% and 65% of source-destination pairs in the PlanetLab and Ripe Atlas data sets respectively, have their paths changing for every traceroute we perform (i.e., show zero persistence in the figures). Such significant departure in path persistency from the results in [1] may be due to a real change in the Internet behavior as the Internet has changed dramatically since the time of that study. Another potential explanation is that the analysis in [1] overestimated persistence due to its low probing frequency. Indeed, probing the Internet at low frequency might result in missing some path changes.

Comparing the path persistence in the two datasets, they show very different results. For instance, under the loose definition, the PlanetLab data set shows that nearly 18% of source-destination pairs have 0.6 or less persistence, whereas nearly 58% of the host pairs have these levels of persistence in the Ripe Atlas data set. Overall, the Ripe Atlas data set exhibits a significantly lower persistence than the PlanetLab data set, as visualized clearly in Figure 2c, which repeats the mid definition curves for both datasets. A potential source for the discrepancy is different probing frequency – 1 minute in the Planet Lab dataset vs. 30 minutes in the Ripe Atlas dataset. However, when we remove this factor by recomputing the PlanetLab persistence for the probing frequency of once every 30 minutes (by taking every 30^{th} data point in our PlanetLab data set for

[3] We follow [3] in choosing 0.6 as the benchmark value when calling out our results. In fact, while [3] used this value for route prevalence only, since that was the focus of their study, we utilize this value throughout.

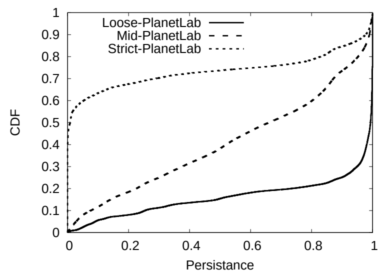

(a) PlanetLab

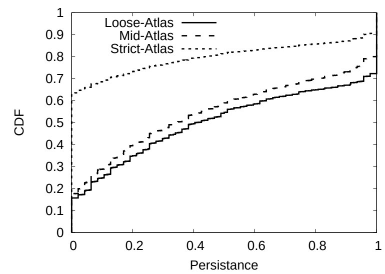

(b) Ripe Atlas

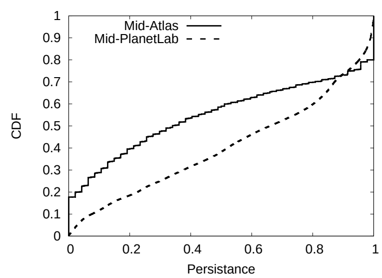

(c) PlanetLab vs. Ripe Atlas (mid definition)

Figure 2: Route persistence

analysis), route persistence in the Planet Lab dataset does not change appreciably (we omit the corresponding graphs for brevity). Therefore, we conclude that the difference in the probing frequency between the data sets is not the cause of discrepancy.

We don't speculate about the cause beyond noting again that both datasets reflect very different environments. We observe discrepancies in other metrics between the two datasets as well throughout this study (as presented below). These discrepancies highlight the dangers of generalizing point observations to the entire Internet.

4.2. MPLS Impact on Route Persistence

We turn to assessing the impact of MPLS on route persistence, by comparing this metric for MPLS and no-MPLS host-pairs. Figure 3 compares CDFs of route persistence of these two classes of host-pairs in the PlantLab data set for our definitions of path equality. Figure 4 shows the same results for the Ripe Atlas data set. While the MPLS and No-MPLS host-pairs exhibit similar route persistence in the PlanetLab dataset, the Ripe Atlas figures suggest that MPLS pairs experience lower persistence for all three definitions of path equality. For

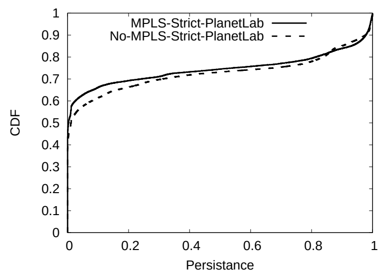

(a) Strict definition of path equality

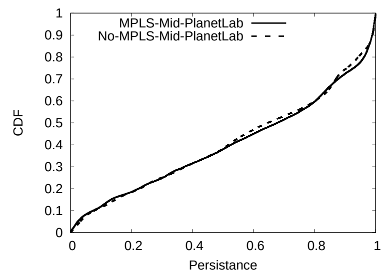

(b) Mid definition of path equality

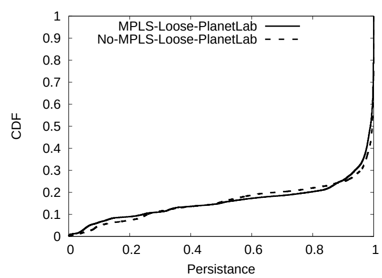

(c) Loose definition of path equality

Figure 3: MPLS impact on route persistence: PlanetLab data set

example, 80% of MPLS pairs experience 0.6 or less persistence compared to only 60% for the No-MPLS pairs according to the mid definition. Therefore, MPLS seems to be decreasing the stability of Internet routes according to one data set. Since the two data sets reflect very different environments as explained in Section 3.1.3, a separate study is needed to confirm this finding, perhaps via additional data sets.

4.3. Route Prevalence

We now assess the Internet path stability via the prevalence of the dominant route. Figure 5 shows the CDF of the dominant route prevalence under the three definitions of path equality for both datasets.

We find significantly lower route prevalence in both our data sets than the results in the Paxson study [1]. For example, only 52.5% and 40% (mid-definition) of the source-destination pairs in, respectively, PlanetLab and Ripe Atlas datasets have prevalence of the dominant route over 0.6 whereas the corresponding value in [1] is about 70%. Unless route changes are somehow synchronized with the path measurements (which seems unlikely), the fraction of

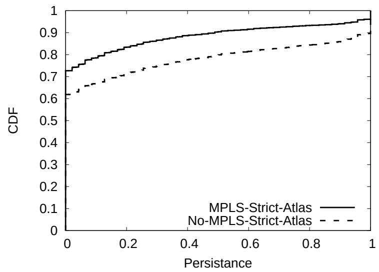

(a) Strict definition of path equality

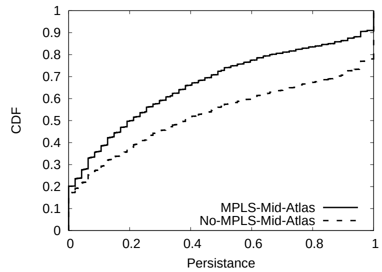

(b) Mid definition of path equality

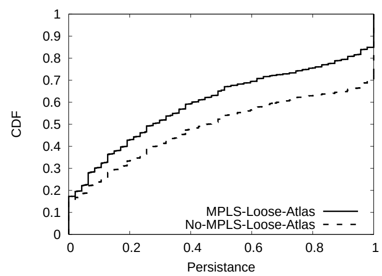

(c) Loose definition of path equality

Figure 4: MPLS impact on route persistence: Ripe Atlas data set

time measurements observe the dominant route should be independent of the measurement frequency as long as large enough total number of measurements per each host-pair is collected. We thus believe these discrepancies represent a change in the Internet behavior from the time of the Paxson study. We can further compare our results to [3] according to the loose path equality definition, since [3] used only the loose definition as mentioned earlier. The Ripe Atlas dataset exhibits route prevalence similar to [3] under the loose definition, while PlanetLab dataset shows – again – much lower route prevalence.

Comparing route prevalence exhibited by both datasets to each other is inconclusive as the results depend on the definition of path equality and the value of the prevalence chosen for comparison. As an example, Figure 5c contrasts the cumulative distributions of route prevalence for host-pairs in both datasets for the most realistic mid-definition.One case with clear distinction between the two datasets is for the loose definition, where the PlanetLab route prevalence is drastically higher, with twice as many host-pairs (80.7% vs. only 42% in the Ripe Atlas dataset) having higher than 0.6 prevalence.

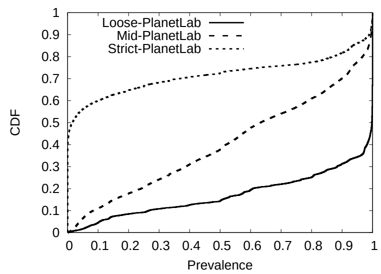

(a) PlanetLab

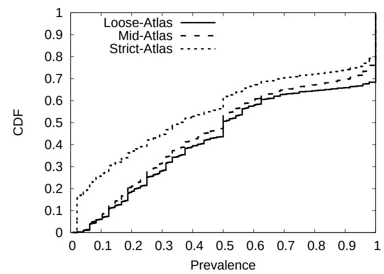

(b) Ripe Atlas

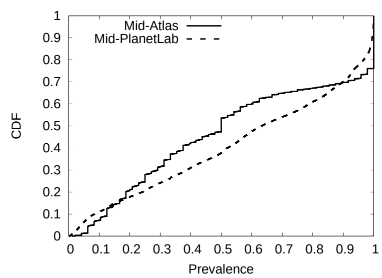

(c) PlanetLab vs. Ripe Atlas (mid definition)

Figure 5: Dominant route prevalence

4.4. MPLS Impact on Route Prevalence

Figures 6 and 7 show the impact of MPLS on route prevalence, by comparing the probability distribution of route prevalence for the MPLS pairs versus the No-MPLS pairs. Similar to path persistence, MPLS does not appreciably affect the path prevalence in the PlanetLab data set but has noticeable impact in the Ripe Atlas data set.

4.5. Route Diversity

The next property that we assess is the diversity of the routes between a source-destination pairs. Figure 8 shows the CDF of the number of distinct paths that were taken between source-destination pairs for the duration of our measurements in both datasets according to the three definitions of path equality. Note that as we investigate this property, we remove host-pairs with no probes reaching their destination. The Ripe Atlas data set exhibits significantly lower diversity than PlanetLab data set. According to the mid definition in the PlanetLab data set, over 60% of source-destination pairs experienced more than 10 distinct routes throughout the 24 hours of our experiment. On the

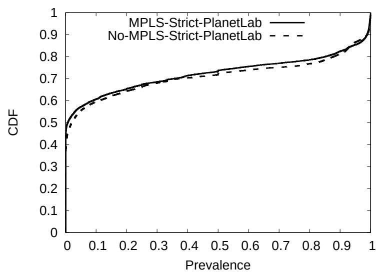

(a) Strict definition of path equality

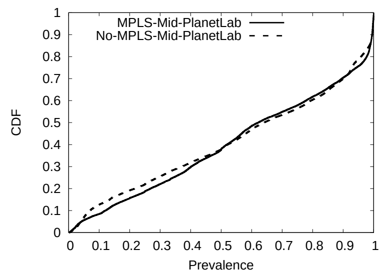

(b) Mid definition of path equality

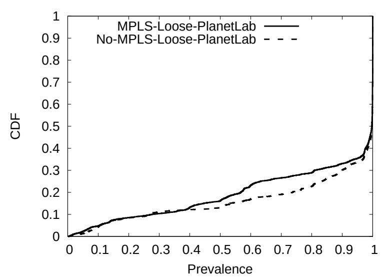

(c) Loose definition of path equality

Figure 6: MPLS impact on dominant route prevalence: PlanetLab dataset

other hand, the Ripe Atlas data set shows about 50% of pairs experienced less than three distinct routes.

Similar to route persistence, the discrepancy between the two data sets could be due to the difference in probing frequency, since probing at a low rate may miss some paths. To factor out this aspect, we coarsen our PlanetLab data set to match the probing frequency of the Ripe Atlas data set (once per 30 minutes). Figure 9a depicts the resulting CDF for route diversity among host-pairs. As shown, the two data sets become consistent when the same probing frequency is used.

When compared to [3], our data sets appear to suggest an increase in diversity of Internet routes where the authors reported that only around 30% of source-destination pairs experienced more than 10 distinct routes over the course of four days (i.e., 2.5 or more distinct routes per day on average), while as mentioned our host-pairs utilize substantially more routes over the day of our experiment. However, it should be noted that the authors in [3] probed the Internet once per hour for about four days which could also mean that they might have missed some routes. Therefore, we further coarsen the probing fre-

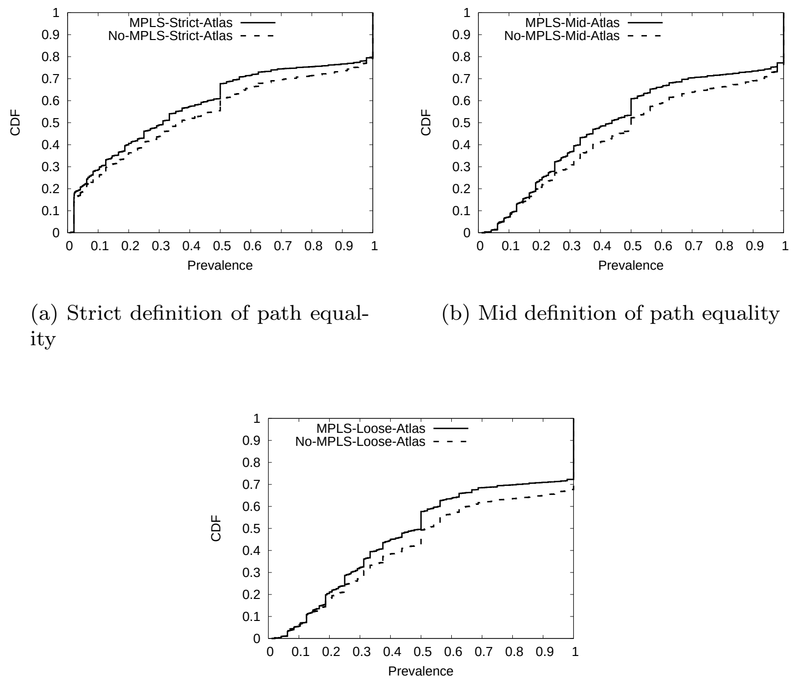

(a) Strict definition of path equality

(b) Mid definition of path equality

(c) Loose definition of path equality

Figure 7: MPLS impact on dominant route prevalence: Ripe Atlas dataset

quency in our data sets to once per hour to match the probing frequency of [3]. Figures 9b and 9c show the resulting CDF of route diversities. Now, about 35% and 28% of host pairs experience three or more distinct routes according to the mid definition of PlanetLab and Ripe Atlas data sets respectively, which is roughly consistent with the findings in [3].

Therefore, we conclude that end hosts enjoy a significant route diversity, and this diversity has been under-reported in the literature due to low probing frequency. Furthermore, the findings of this experiment illustrate the importance of properly parameterizing the measurements experiments (such as choosing appropriate probing frequency) for obtaining sound results.

4.6. MPLS Impact on Route Diversity

Figures 10 and 11 show the impact of MPLS on route diversity in our two data sets, by comparing the route diversity of MPLS and no-MPLS host-pairs. One can notice some tendency towards higher diversity for MPLS pairs; however with the exception of Ripe Atlas data set under strict path equality definition, the difference is very small.

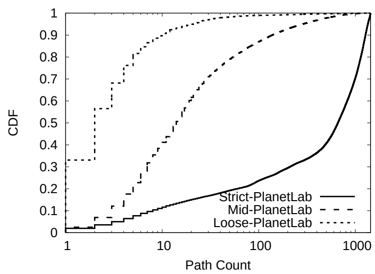

(a) PlanetLab

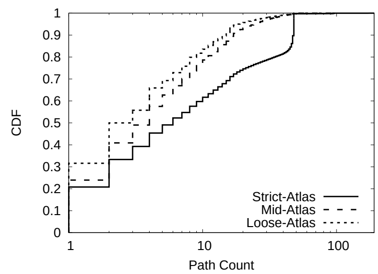

(b) Ripe Atlas

Figure 8: Route diversity

4.7. RTT Variability

Up to now we have considered path stability in terms of the sequence of routers traversed by packets. However, the impact of path change on applications and user experience often depends on how different the RTT becomes after the path changes. To understand this aspect of path (in)stability, we examine the cumulative distribution of the amount of change in RTT when a path changes for the PlanetLab and Ripe Atlas datasets (The figures are dropped to maintain the reader's focus). The key observation from these distributions is that, in both data sets, path changes have no appreciable effect on RTT. In particular, in the PlanetLab data set, about 90% of the path changes result in less than 10 ms change in RTT. Furthermore, about 70% of path changes entail an insignificant RTT change of under 1 ms. Similarly, in relative terms, about 90% of changes in RTT were smaller than 10%. The different definitions of path equality produce similar RTT results. We also examined the separate distributions of the amount of increase in RTT and decrease in RTT and again found similar results. We therefore drop these figures for brevity. Similarly, the Ripe Atlas dataset suggests an insignificant difference in the RTT before and after the path change. For instance, around 80% of hosts have RTT change under 10 ms, and for over 80% of host pairs the change is within 10%. MPLS and No-MPLS groups also exhibit similar type of change in RTT, and we omit these figures for brevity. Overall, our results indicate that RTT between host-pairs tends to stay stable despite path changes.

5. Conclusions and Future Work

This study evaluates the stability and diversity of Internet paths and the impact of MPLS on these properties. The study uses two separate data sets. One is collected from an active measurement study using PlanetLab and the other is collected using the Ripe Atlas platform. We find that Internet paths

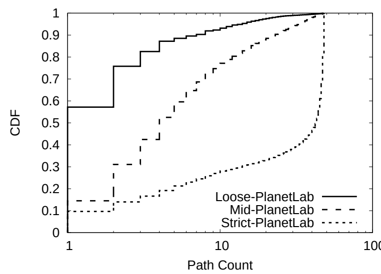

(a) PlanetLab, 30 min probing period

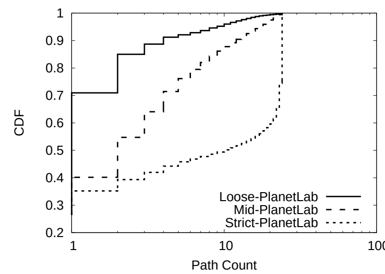

(b) PlanetLab, 60 min probing period

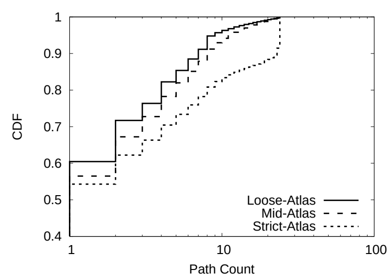

(c) Ripe Atlas, 60 min probing period

Figure 9: Route diversity at coarser probing rates

are significantly less stable than previously reported with respect to route persistence. While the PlanetLab data set suggests MPLS is not contributing to this decrease in stability, the Ripe Atlas data set suggests a significant contribution to this decrease. Moreover, we find that path changes are usually associated with an insignificant change in RTT.

Our experiments demonstrate that Internet routes between a given host-pair exhibit a significant diversity, While previous studies indicated that this diversity is much less than what we report here, we show that this is likely due to previous studies not probing Internet paths frequent enough, and therefore, missing some paths.

We believe that these results represent a significant departure from the current view that the research community has about the stability and diversity of Internet paths. Due to the importance of these properties in designing systems and experiments, we believe that our findings may be practice changing. In particular, we provide an evidence that packets between a host-pair switches between a larger number of paths (high path diversity) and very quickly (low path persistence) compared to what is currently assumed about these properties.

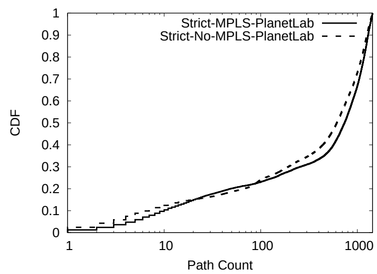

(a) Strict path equality definition

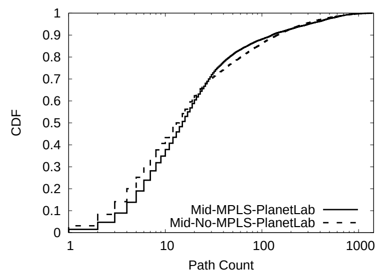

(b) Mid path equality definition

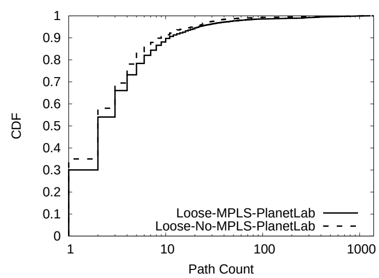

(c) Loose path equality definition

Figure 10: MPLS impact on route diversity: PlanetLab dataset

This difference calls for a greater care when designing systems and experiments that rely on assumptions about these path properties. For example, the authors in [5] probe routes once a day relying on the results of previous studies to reduce probing load. Our findings in this study show that this assumption is not valid anymore. Therefore, re-evaluation of their system is needed in the light of our findings. Moreover, we expect further studies to evaluate the impact of this change on protocols such as TCP and to optimize them to handle the greater path diversity and lower path persistence that exist in the Internet today.

The findings of this study provide yet another example on the importance of following best measurement practices. First, due to the continuity of changes of the Internet properties, studies that re-evaluate previously established properties are essential. Second, measurement studies need to parametrize their experiments carefully in oder to draw useful conclusions. Third, drawing general conclusions from point measurements might lead to incorrect conclusions. Indeed, our study involves two relatively large data sets, and they disagreed on some properties.

While we measured the variability of RTT when a path between a source-destination pair changes, the variability in the available bandwidth remains

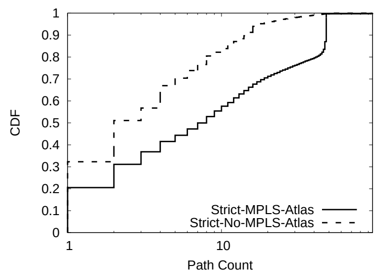

(a) Strict path equality definition

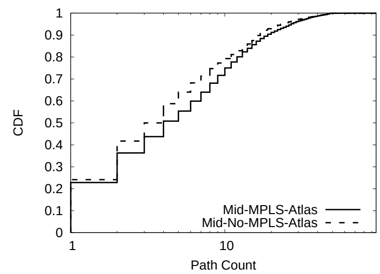

(b) Mid path equality definition

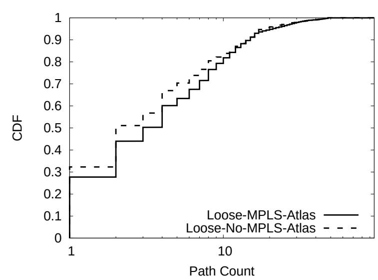

(c) Loose path equality definition

Figure 11: MPLS impact on route diversity: Ripe Atlas dataset

unexplored. Measuring the available bandwidth is challenging since most of the "accurate" bandwidth estimation tools require control of both ends of the path whereas in our experimental setup, we do not have access to both ends. There are two different ways to tackle this problem. First, we can monitor the variability of available bandwidth between PlanetLab nodes themselves. Second, we can rely on tools that estimate the available bandwidth with the control of only one side of the path such as [28].

References

[1] Vern Paxson. End-to-end Routing Behavior in the Internet. In *SIGCOMM*, pages 25–38, New York, NY, USA, 1996. ACM.

[2] Ítalo Cunha, Renata Teixeira, and Christophe Diot. Measuring and Characterizing End-to-end Route Dynamics in the Presence of Load Balancing. In *PAM*, pages 235–244, Berlin, Heidelberg, 2011. Springer-Verlag.

[3] Yaron Schwartz, Yuval Shavitt, and Udi Weinsberg. On the Diversity, Stability and Symmetry of End-to-End Internet Routes. In *IEEE INFOCOM Workshops*, pages 1–6, March 2010.

[4] Italo Cunha, Renata Teixeira, Darryl Veitch, and Christophe Diot. Predicting and Tracking Internet Path Changes. *ACM SIGCOMM Computer Communication Review*, 41(4):122–133, 2011.

[5] Yan Chen, David Bindel, Hanhee Song, and Randy. H. Katz. An Algebraic Approach to Practical and Scalable Overlay Network Monitoring. In *ACM SIGCOMM*, pages 55–66, New York, NY, USA, 2004. ACM.

[6] Joel Sommers, Paul Barford, and Brian Eriksson. On the Prevalence and Characteristics of MPLS Deployments in the Open Internet. In *Internet Measurement Conference (IMC'11)*, pages 445–462, New York, NY, USA, 2011. ACM.

[7] Xipeng Xiao, A. Hannan, B. Bailey, and L. M. Ni. Traffic engineering with MPLS in the Internet. *IEEE Network*, 14(2):28–33, Mar 2000.

[8] Planet Lab. http://www.planet-lab.org.

[9] Zakaria Al-Qudah, Mohammad Alsarayreh, Ibrahim Jomhawy, and Michael Rabinovich. Internet path stability: Exploring the impact of mpls deployment. *2016 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7, 2016.

[10] Craig Labovitz, Abha Ahuja, and Farnam Jahanian. Experimental Study of Internet Stability and Backbone Failures. In *FTCS*, pages 278–285, Washington, DC, USA, 1999. IEEE Computer Society.

[11] Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker. Characterizing and Measuring Path Diversity of Internet Topologies. In *SIGMETRICS*, pages 304–305, San Diego, CA, USA, 2003. ACM.

[12] Benoit Donnet, Matthew Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. Revealing MPLS Tunnels Obscured from Traceroute. *SIGCOMM Comput. Commun. Rev.*, 42(2):87–93, March 2012.

[13] Anwar Elwalid, Cheng Jin, Steven Low, and Indra Widjaja. MATE: MPLS Adaptive Traffic Engineering. In *INFOCOM'11*, pages 1300–1309, 2001.

[14] Karl Hendling, Brikena Statovci-Halimi, Gerald Franzl, and Artan Halimi. A New Bandwidth Guaranteed Routing Approach for Online Calculation of LSPs for MPLS Traffic Engineering. In *IFIP/IEEE MMNS '03*, pages 220–232, Sep 2003.

[15] J. Boyle, V. Gill, A. Hannan, D. Cooper, D. Awduche, B. Christian, and W. S. Lai. Applicability Statement for Traffic Engineering with MPLS (RFC 3346). 2002.

[16] MPLS TE Basics. https://sites.google.com/site/amitsciscozone/home/important-tips/traffic-engineering/mpls-te-basics; Retrieved on 12/30/2019.

[17] Radek Vingralek, Yuri Breitbart, Mehmet Sayal, and Peter Scheuermann. Web++: A System for Fast and Reliable Web Service. In *USENIX Annual Technical Conference, General Track*, pages 171–184, 1999.

[18] Mark Allman. Comments on bufferbloat. *ACM SIGCOMM Computer Communication Review*, 43(1):30–37, 2013.

[19] Robert L. Carter and Mark E. Crovella. On the Network Impact of Dynamic Server Selection. *Computer Networks*, 31(23-24):2529–2558, 1999.

[20] Jay Aikat, Jasleen Kaur, F Donelson Smith, and Kevin Jeffay. Variability in tcp round-trip times. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 279–284. ACM, 2003.

[21] Toke Høiland-Jørgensen, Bengt Ahlgren, Per Hurtig, and Anna Brunstrom. Measuring latency variation in the internet. In *Proceedings of the 12th International on Conference on emerging Networking EXperiments and Technologies*, pages 473–480. ACM, 2016.

[22] Srinivas Shakkottai, Nevil Brownlee, Andre Broido, et al. The rtt distribution of tcp flows on the internet and its impact on tcp based flow control. Technical report, Cooperative Association for Internet Data Analysis (CAIDA), 2004.

[23] Scamper. http://www.caida.org/tools/measurement/scamper/.

[24] Hsiao keng Jerry Chu. Tuning TCP Parameters for the 21st Century. http://www.ietf.org/proceedings/75/slides/tcpm-1.pdf, 2009.

[25] Hal Burch and Dawn Song. A Security Study of the Internet: An Analysis of Firewall Behavior and Anonymous DNS. *Technical Report CMU-CS-04-141, Carnegie Mellon University,School of Computer Science*, 2004.

[26] Matthew Luckie and Ben Stasiewicz. Measuring Path MTU Discovery Behaviour. In *IMC'10*, pages 102–108, 2010.

[27] Himabindu Pucha, Ying Zhang, Z Morley Mao, and Y Charlie Hu. Understanding network delay changes caused by routing events. *ACM SIGMETRICS performance evaluation review*, 35(1):73–84, 2007.

[28] Demetres Antoniades, Manos Athanatos, Antonis Papadogiannakis, Evangelos P Markatos, and Constantine Dovrolis. Available Bandwidth Measurement As Simple As Running Wget. In *PAM'06*, 2006.