Internet Path Stability: Exploring the Impact of MPLS Deployment

Zakaria Al-Qudah
CSE Dept.
American University of Ras Al Khaimah (AURAK)
Ras Al Khaimah, UAE

Mohammad Alsarayreh, Ibrahim Jomhawy
Computer Engineering Dept.
Yarmouk University
Ibird, Jordan

Michael Rabinovich
EECS Dept.
Case Western Reserve Univeristy
Cleveland, OH

Abstract

With constant evolution of the Internet, many of its even well-established properties continue to change. This study re-evaluates the stability and diversity of Internet paths and whether it is affected by the widely deployed Multi-Protocol Label Switching (MPLS). In particular, using traceroutes between large number of source-destination pairs with the Internet Control Message Protocol (ICMP) extensions for the MPLS protocol, we study the stability and diversity of Internet paths as well as the amount of change in Round Trip Time (RTT) associated with a path change between a source-destination pair. Our results indicate that Internet routes are significantly less stable than previously reported. However, MPLS does not contribute significantly to this instability, and most path changes are associated with an insignificant change in RTT. While our route diversity results are not directly comparable with previous results, we find a great diversity of Internet routes: 60% of source-destination pairs in our experiments experienced 10 or more distinct routes in the course of 24 hours.

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

I. Introduction

The end-to-end characteristics of Internet routing has been an active field of study for decades. Several Internet path properties have been measured 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 several measurement studies (e.g., [4]) as well as in designing systems (e.g., [5]). In these studies, the authors probe paths with low frequency to reduce measurement overhead while relying on the stability of these paths for maintaining recent enough view of the Internet paths landscape. Furthermore, the stability (or the lack of) of Internet paths may affect the end-to-end performance of applications. For example, switching paths in the middle of a Transmission Control Protocol (TCP) session may require TCP some time to learn the new path characteristics which may lead to a sub-optimal performance in data delivery.

Updating our understanding of the current end-to-end Internet path properties is crucial since the Internet continues to evolve. For example, 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, providing another tool for service providers to better utilize their resources and to provide better services [6], [7].

While there are studies that explicitly consider the impact of load balancing on Internet route stability [2], we are unaware of any study that has considered the impact of traffic engineering based on MPLS on the stability of these routes. Therefore, the goals of this study are as follows.

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

To achieve goals, we conduct a large-scale active measurements study to probe Internet paths. We use several PlanetLab [8] nodes as vantage points to monitor the paths to a large number of hosts on the Internet. We conduct our traceroute measurements on a short time scale (one minute) for a full day per source-destination pair.

Analyzing the results, we find a decrease in the stability of Internet paths when compared to previously published results. However, MPLS does not seem to contribute significantly to this drop in stability. Despite the significant decrease in stability, we find that the variability in RTT when a path changes is generally insignificant.

II. Related Work

A. Path Characteristics

The Internet behavior has 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 Scheartz 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 where other studies probed at intervals of hours and days. Probing at short-time scale has the potential of revealing high frequency route flaps which would otherwise be missed if we probe at larger time-scales.(*ii*) We employ broader, more realistic definition of path equality as explained in Section III-B. (*iii*) We investigate the impact of MPLS on path stability and diversity. Overall, we find much higher path variability than the above studies.

A quantitative evaluation of asymmetry in the Internet for forward and reverse paths of end-to-end points was performed at the Autonomous System (AS) level in [9] for academic and commercial networks. The findings indicate that the level of asymmetry in commercial networks is higher than that in educational and research networks. The delay asymmetry in the Internet was studied in [10]. They found that the delay asymmetry depends on routing changes, and therefore, it is considered to be a dynamic property.

The stability of the Internet and the wide area backbone failures were studied experimentally in [11]. The Internet stability is found to depends on the telecommunication switching system and the software and hardware components of routing systems. The study in [11] used five of the largest U.S. Internet Exchange Points (IXPs) over three years. Therefore, several trends of Internet stability and failure were described. The path diversity of Internet topology was analyzed in [12] using different Internet topologies to describe the path diversity inside an AS and across multiple ASes. The authors in [13] showed that the Internet policy affects the length of Internet paths. In [14], the impact of routing events on the performance of the end-to-end Internet paths was evaluated using geographically and topologically diverse vantage points.

B. 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 [15] shed the light on the availability of obscured tunnels. The study develops methods for revealing the deployment and characteristics of obscured MPLS tunnels from traceroute data. 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 general description of MPLS system design for traffic engineering in an ISP network and a procedure for deploying MPLS system was presented in [7]. The authors indicate that the MPLS system is feasible and useful in large ISP network. In [16], a Multipath Adaptive Traffic Engineering (MATE) was described which can be used for switched networks such as MPLS. An analytical model was proposed to improve the stability and optimality of MATE, which can remove any traffic imbalances and can avoid network congestion.

In [17] a new online algorithm for dynamically routing bandwidth guaranteed Label Switched Paths (LSPs) was built. The results gave better bandwidth performance than other existing algorithms. The applicability and limitations of MPLS in traffic engineering were discussed in [18].

III. Methodology

A. Data Set

To collect the data set for this study, 88 PlanetLab [8] nodes were used as vantage points. We collect 88K distinct IP addresses for hosts on the Internet by collecting a number of top web hosts from Alexa [19] and resolving these host names. We then assign for each PlanetLab node a set of 1000 IP addresses and monitor the path between the PlanetLab node and the corresponding 1000 IP addresses. Therefore, the total number of paths monitored is 88K paths.

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 [20] with the traceroute option that implements ICMP extensions for MPLS. There are two features that routers can implement 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 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 [15] 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

Source-destination pairs	88K
Traceroutes performed	$\sim$ 126M
Traceroutes reached their destinations	$\sim$ 80M
Touched IPs	$\sim$ 1.45B
Distinct IPs	$\sim$ 141K
Distinct ASes	5686

Table I
SUMMARY OF THE DATA SET

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 [21], and therefore we believe our choice of this parameter will not affect our results. We configure Scamper to issue one traceroute attempt per hop instead of the three-attempts default. Issuing three attempts basically means probing at a time-scale that is shorter than our probing frequency of 60 seconds. 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. The paths observed by our traceroutes touched upon a total of 5686 distinct autonomous systems and 1.45B IP address hops, of which 141K addresses were distinct. Table I summarizes the data set obtained from this experiment.

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.

B. 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 ICMP echo requests (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 affect 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 since a significant number of nodes in the Internet today do not to respond to ICMP echo requests; indeed, [22] found 53.2% of live hosts they sampled did not respond to pings, and neither did 21-27% of top-50K webservers in [23]. 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 differentiate among 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 ICMP packet 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 is might not be true. Though, we consider this definition to match the methodology of [3]. Therefore, our estimation on path stability based on this definition represents an upper bound. The mid definition lies in between.

C. Metrics

We borrow the two metrics used by [1] to evaluate the stability of Internet paths. These are:

1) Persistence: This metric evaluates 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 total possible "no path change" events. Since we perform 1440 measurements per source-destination pair, 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 cloud 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.

2) Prevalence: This metric evaluates how often a certain path between two end-points is taken. In our measurements, we measure the number of times each distinct path was taken out of the 1440 probes. We measure the prevalence of the path that was taken the most (called the dominant path). For

Probes %	1 Probe	30%	50%	70%	90%	100%
MPLS Src/Dst pairs	45%	42.5%	42%	41%	40%	27%

Table II
CLASSIFICATION OF MPLS SOURCE-DESTINATION PAIRS FOR VARIOUS VALUES OF THE THRESHOLD X.

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 probes for the route 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 probes 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 "P1 or P2" 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.

D. *MPLS vs. No-MPLS*

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 [15], 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 II 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 does 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.

IV. RESULTS

As mentioned, our detection of MPLS pairs is conservative, in that some non-MPLS pairs may communicate over undetected MPLS tunnels. Further, while our results below indicate high path variability, we in fact underestimate path variability, as path changes involving obscure and invisible MPLS tunnels are undetected.

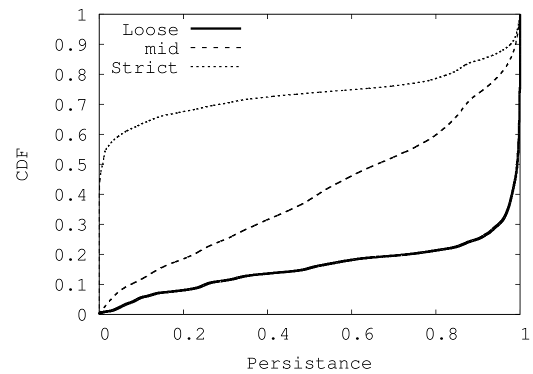

Figure 1. Route Persistence.

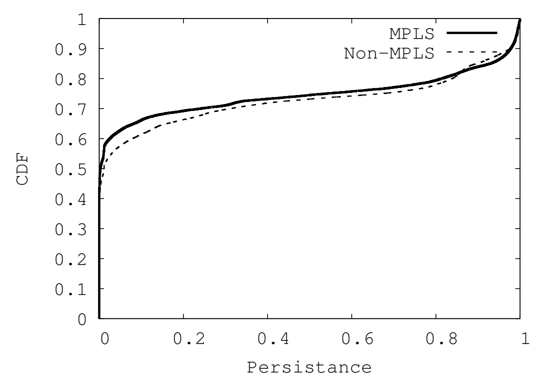

Figure 2. Route Persistence. Comparing MPLS pairs versus No-MPLS pairs with strict definition of path equality.

A. *Persistence*

We first evaluate the persistence of Internet routes over the entire data set. Figure 1 shows the results. As expected, the strict definition of path equality results into the least route persistence. In particular, for nearly 50% of source-destination pairs, the path kept changing for every traceroute we perform (i.e., showed zero persistence in the figure). As mentioned above, this represents a lower bound on our estimation of route persistence. According to the mid definition, close to 46% of the source-destination pairs had 0.6 or less persistence whereas the number is 18% for the loose definition. While the way the results in [1] are presented does not allow us to make an apple-to-apple comparison with our results, the author found that two-thirds of measured routes persisted for days or weeks.

Figures 2, 3, and 4 show the results for strict, mid, and loose definitions of path equality respectively featuring MPLS vs. No-MPLS pairs. As the figures show, MPLS and No-MPLS pairs experience similar persistence , and therefore, MPLS does not seem to affect the stability of Internet routes in terms of persistence.

B. *Prevalence*

We now assess the prevalence of Internet routes via reporting on the prevalence of the dominant route. Figure 5 shows the results for the three definitions of path equality. As shown, 74.5%, 47.5% and 19.3% of source-destination

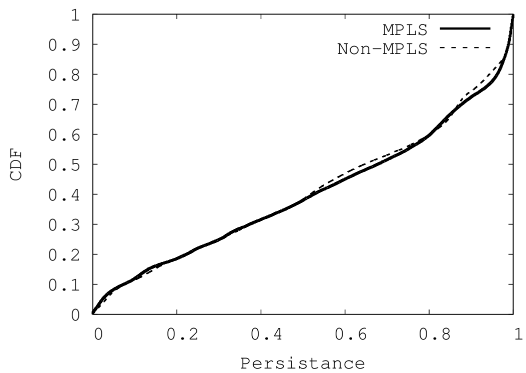

Figure 3. Route Persistence. Comparing MPLS pairs versus No-MPLS pairs with mid definition of path equality.

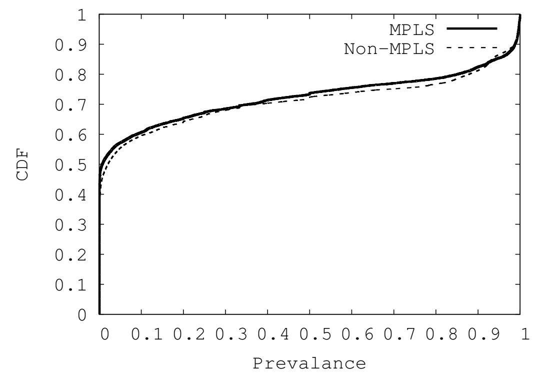

Figure 6. Dominant route prevalence. Comparing MPLS pairs versus No-MPLS pairs with strict definition of path equality.

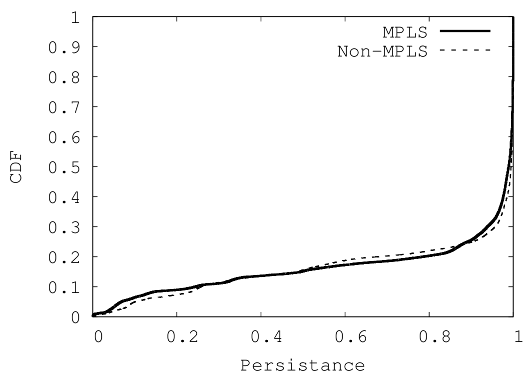

Figure 4. Route Persistence. Comparing MPLS pairs versus No-MPLS pairs with loose definition of path equality.

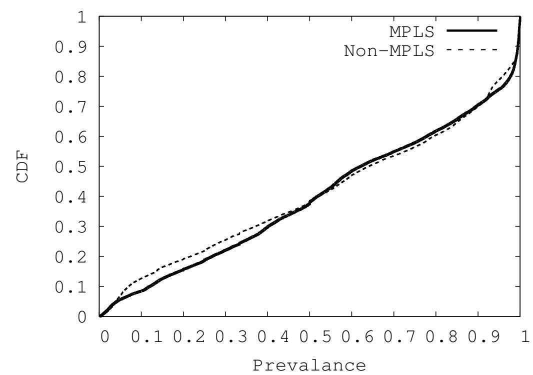

Figure 7. Dominant route prevalence. Comparing MPLS pairs versus No-MPLS pairs with mid definition of path equality.

pairs had 0.6 or less prevalence for the strict, mid, and loose definitions respectively. The prevalence of the dominant route is significantly lower than reported in [1]. For example, only 52.5% (mid-definition) of the source-destination pairs examined have prevalence of the dominant route over 0.6 whereas the value was reported in [1] to be about 70%. Furthermore, our prevalence results are consistent with the findings in [3].

Figures 6, 7, and 8 show the dominant route prevalence according to the strict, mid, and loose definitions of path equality, for the MPLS pairs versus the No-MPLS pairs. The results seem similar for both groups (MPLS vs. No-MPLS) with only marginal decrease in stability in the No-MPLS group according to the loose definition.

C. Route Diversity

The next property that we assess is the diversity of the routes between a source-destination pairs. Figure 9 shows the CDF of the number of distinct paths that were taken between

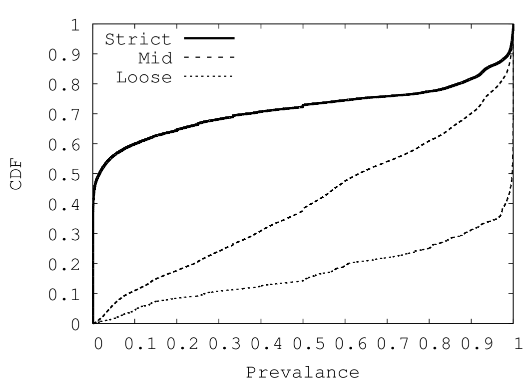

Figure 5. Dominant route prevalence.

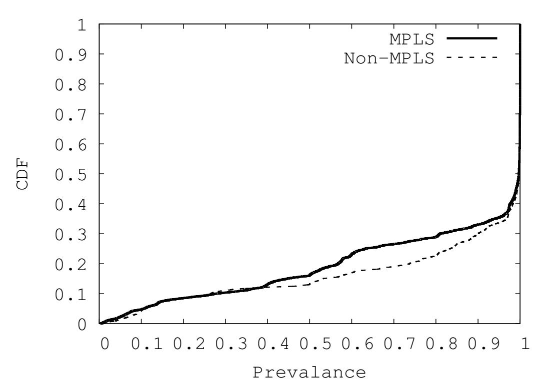

Figure 8. Dominant route prevalence. Comparing MPLS pairs versus No-MPLS pairs with loose definition of path equality.

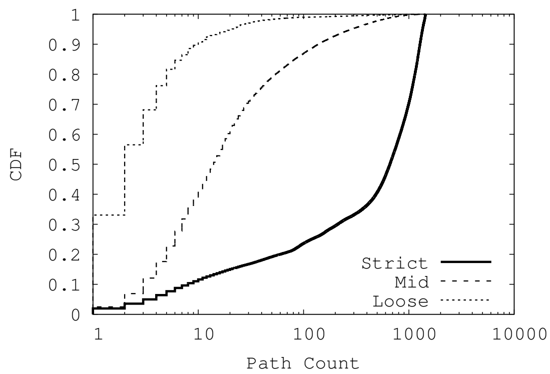

Figure 9. Route diversity.

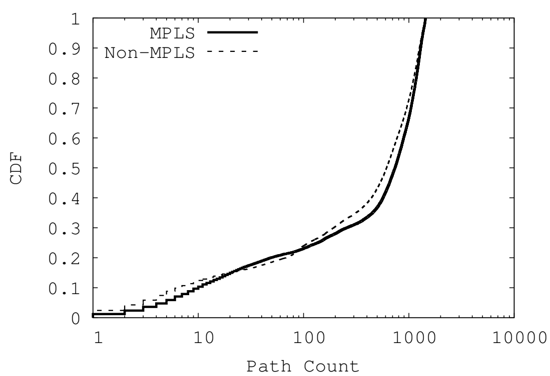

Figure 10. Route diversity under strict definition of path equality, MPLS pairs versus No-MPLS pairs.

source-destination pairs for the duration of our measurements according to the three definitions of path equality. As expected, the strict definition exhibits the largest diversity. According to the mid definition, over 60% of source-destination pairs experienced more than 10 distinct routes throughout the life of our experiment. According to [3], 30% of source-destination pairs experienced more than 2.5 distinct routes on overage.

Figures 10, 11, and 12 show the diversity of the routes between a source destination pair according to the strict, mid, and loose definitions of path equality, for the MPLS pairs versus the No-MPLS pairs. The results indicate the diversity of paths is slightly larger for the MPLS group for all definitions of path equality.

D. RTT Variability

We next plot the amount of change in RTT in Figure 13 and the percentage of change in RTT in Figure 14 for all three definitions. As shown, about 90% of the path changes result in a drop or increase in RTT that is less than 10 ms. Furthermore, about 70% of path changes result in an insignificant change in RTT (less than 1 ms). Similarly, about 90% of changes in RTT were smaller than 10%. The different definitions of path equality result into similar RTT results. Similar results also hold for paths that connect MPLS and No-MPLS pairs. We also examined the cases where we observe an increase in RTT and those where we observe a decrease in RTT separately

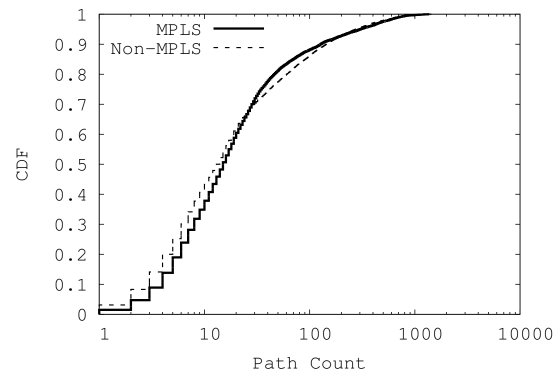

Figure 11. Route diversity under mid definition of path equality, MPLS pairs versus No-MPLS pairs.

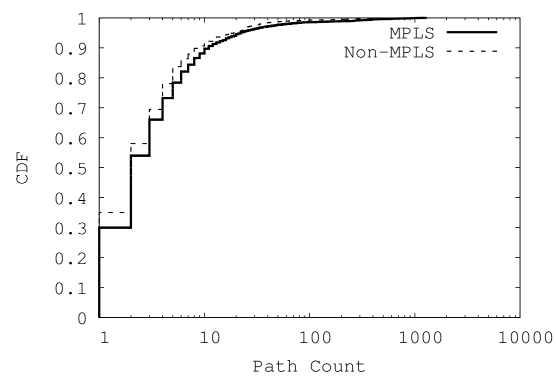

Figure 12. Route diversity under loose definition of path equality, MPLS pairs versus No-MPLS pairs.

and found similar results. Therefore, we drop these figures. Previous studies reported contradictory results on the stability of RTT. For example, the authors in [24] found extreme RTT variations when they monitored single source-destination pair for two days, once per 30 seconds. The authors in [25] found that HTTP latency correlates with RTT, and that it is generally stable. Our results are consistent with the results of a study that considered RTT in the context of the “buffer bloat“ problem [26].

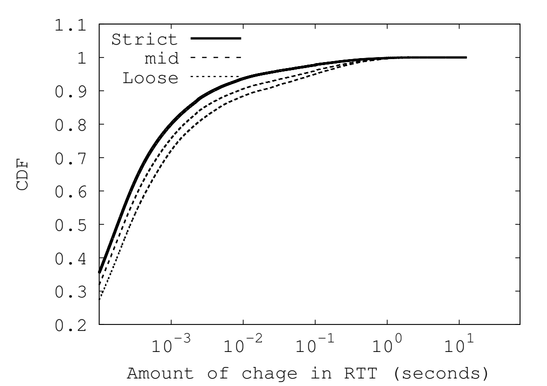

Figure 13. Distribution of amount of change in RTT for the three definitions of path equality.

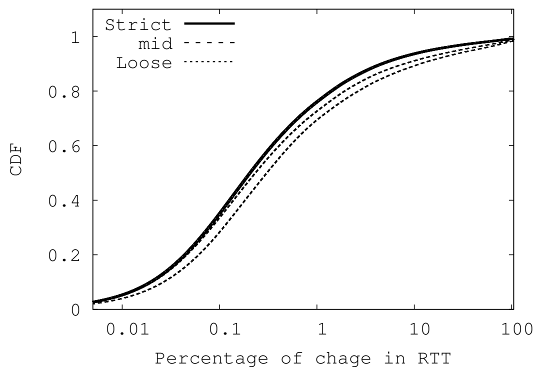

Figure 14. Distribution of percentages of change in RTT for the three definitions of path equality.

V. 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 a number of PlanetLab nodes as vantage points to monitor a large number of paths from these vantage points to hosts on the Internet using traceroute. We find that Internet paths are significantly less stable than previously reported. However, we find that the RTT does not vary significantly when a path changes. Furthermore, our results also indicate that the role of MPLS in the decease of route stability on the Internet is marginal. Our experiments demonstrate that Internet routes between a given host-pair enjoy a significant diversity, and host-pairs that predominantly communicate over paths that include MPLS tunnels show more path diversity than other host-pairs.

The findings of this study provide yet another example on the continuity of changes of the Internet properties. Therefore, studies that re-evaluate previously established properties are essential. Furthermore, this study recommends that such dynamic properties of the Internet are taken into consideration when designing systems or protocols. For example, many of the systems and measurement studies that assumed the stability of Internet routes need to be re-evaluated in the light of the new findings.

While we measured the variability of RTT when a path between a source-destination pair changes, the variability in the available bandwidth remains 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 control only the PlanetLab nodes. 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 [27].

References

[1] Vern Paxson. End-to-end Routing Behavior in the Internet. SIGCOMM ’96, 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. PAM ’11, 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] Yihua He, Michalis Faloutsos, Srikanth Krishnamurthy, and Bradley Huffaker. On Routing Symmetry in the Internet. In *GLOBECOM ’05*, Nov 2005.

[10] Abhinav Pathak, Himabindu Pucha, Ying Zhang, Y. Charlie Hu, and Z. Morley Mao. A Measurement Study of Internet Delay Asymmetry. PAM’08, pages 182–191, 2008.

[11] Craig Labovitz, Abha Ahuja, and Farnam Jahanian. Experimental Study of Internet Stability and Backbone Failures. FTCS ’99, pages 278–285, Washington, DC, USA, 1999. IEEE Computer Society.

[12] Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker. Characterizing and Measuring Path Diversity of Internet Topologies. SIGMETRICS ’03, pages 304–305, San Diego, CA, USA, 2003. ACM.

[13] Hongsuda Tangmunarunkit, Ramesh Govindan, Scott Shenker, and Deborah Estrin. The Impact of Routing Policy on Internet Paths. In *INFOCOM’01*, pages 736–742, 2001.

[14] Feng Wang, Zhuoqing Morley Mao, Jia Wang, Lixin Gao, and Randy Bush. A Measurement Study on the Impact of Routing Events on End-to-end Internet Path Performance. SIGCOMM ’06, pages 375–386, New York, NY, USA, 2006. ACM.

[15] 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.

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

[17] 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.

[18] 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.

[19] Alexa. http://www.alexa.com.

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

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

[22] 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.

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

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

[25] 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.

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

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