Understanding Hop Characteristics for Enhanced Relay Services

Salim Mohamed
Electrical and Computer Engineering
Michigan State University
East Lansing, USA
Email: mohame26@msu.edu

Osama Mohammed
Service Delivery and Management
Innovaway
Napoli, Italy
Email: osama.mohammed@it.ibm.com

***Abstract*—Relay or overlay routing for IP networks has been well-documented. However, the implementation cost of relay solutions has not yet been conclusively identified. There is also a shortage of studies that present and analyze relay characteristics for reducing the relay cost or the probing burden. This paper, in particular, introduces new relay characteristics, such as the number of relay hops in a minimum delay path or Hop-To-Live (HTL). The HTL and Round-Trip-Time (RTT) are examined for reducing the relay cost and estimating stable relay paths. This paper considers a wide-set of $311,360$ Ping measurements on a network of 140 Planetlab hosts. The main results emphasize the redundant probing burden over a $24-$hours period and concluded that relay hosts have favorable dominant relay paths with high (not maximum) prevalence. Unfortunately, such a prevalence causes relay paths to suffering from a weak persistence. The RTT demonstrates the ability to detect such dominant changes with a 25% average error. In contrast to recent studies, our work shows that relay paths are asymmetrical on each examined granularity.**

Keywords-Hop and RTT characteristics; relay routing; Internet measurements.

I. Introduction

Overview: Recent Software Defined Networks (SDNs) leverage the control on relay paths, but the probing cost remains high. The study in [6] characterizes the Time-To-Live (TTL) changes of the physical (underlay) routing. On the contrary, we introduce and examine the influence of a new relay (overlay) routing metric that defines the number of relay hops in a minimum delay path. For the remainder of this paper, we refer to this metric by the Hop-To-Live (HTL). The main point of this paper is in questioning the benefit of $24-$hours measurements for estimating stable relay paths and detecting their changes.

Experiments: Ping was used to perform a set of $311,360$ delay measurements on $19,460$ end-to-end paths in a network of $n = 140$ hosts. Ping sends bulk of packets of four distinct sizes: 0.05, 0.1, 0.25 and 0.5 MBytes. These distinct load measurements were used to examine the HTL characteristics. Pings were scheduled in load order of 4 times in 16 experiments. These measurements were designed on diverse intervals as suggested in [6] to achieve more confidence when capturing routing changes.

Probing Daemon: The network hosts were divided into a number of groups. We performed a single measurement in each group g_i where $i \in [1 \rightarrow m]$. Each prober utilizes two loops: The first, is to probe $n-1$ hosts, and the second one is to probe unresponsive hosts again. The actual group time is: $t_i = \sum_{k=1}^{|g_i|} \lambda_i(k) + \beta_i(k)$ as $\lambda_i(k)$ and $\beta_i(k)$ are probing loop-times. These times can be determined as: $\lambda_i(k) = \sum_{j=1}^{n-1} \bar{\epsilon}$ and similarly $\beta_i(k) = \eta \sum_{j=1}^{\theta_i(k)} \bar{\epsilon}$, where $\theta_i(k)$ is the number of unresponsive hosts in the first loop. The $\bar{\epsilon}$ is the average network probing time while $\eta = 1$ is the average re-probing count. The probing scheme is a server-based synchronization to minimize the effect of the probing conflict occurs when two daemons or more probe a host simultaneously. The groups-sizes are optimized to reduce imperfect measurements, and daemons randomly select hosts as in [2].

Contribution: The major contribution of this paper is in identifying short and long-term characterization for the minimum Round-Trip-Time (RTT) relay paths, such as the dominant prevalence in Section III. For analyzing HTL characteristics, we performed $311,360$ RTT measurements for all paths in a network of 140 hosts. Through analysis, we found that the HTL prevalence has similar behavior to the TTL prevalence studied in [6]. The study proposed new metrics to govern inexpensive relay routing. This study concluded that both HTL and RTT are sufficient metrics to predict relay changes and reduce the probing cost.

The rest of this paper is structured as follows: Section II summarizes related studies and relay schemes. HTL characteristics such as prevalence, persistence, and symmetry are outlined in Section III. The proposed RTT detection scheme for detecting path changes is briefed Section IV. Section IV also describes the RTT statistical prevalence and RTT symmetry. In Section V, we conclude our paper.

II. Related Research

The study in Jiang et al. [1] refers to an estimation-based relay scheme for Skype users. The lack of measurements for end-pairs in [1] was compensated by a network tomography-based delay estimation. This approach is not a perfect replacement for the probing-based relay routing because our study shows that both the physical and the relay paths are asymmetric in general. Researchers in [2] used stable HTL relay paths to construct a Layer-3 forwarding scheme. The difficulty to perform a passive analysis as in [3] on content-based networks like [4] is convincing that it is possible to infer the network

978-3-903176-31-7 © 2020 IFIP

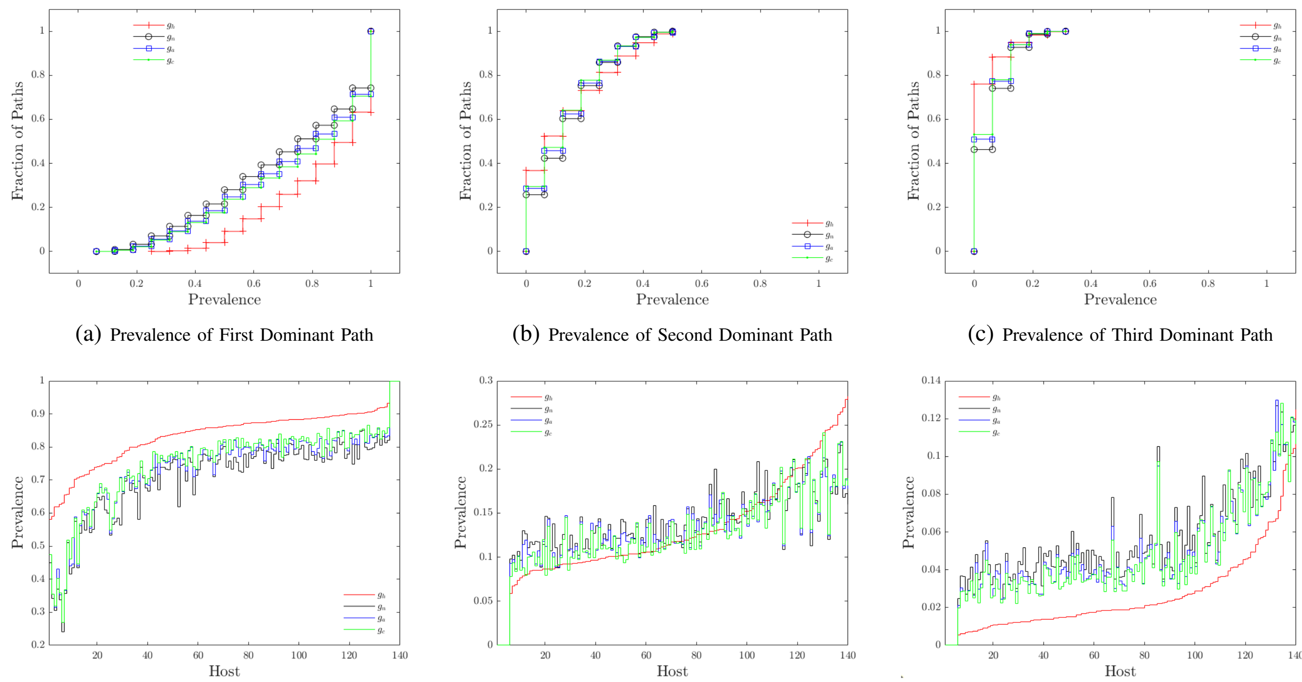

(a) Prevalence of First Dominant Path

(b) Prevalence of Second Dominant Path

(c) Prevalence of Third Dominant Path

(d) Host Prevalence of First Dominant Path

(e) Host Prevalence of Second Dominant Path

(f) Host Prevalence of Third Dominant Path

Fig. 1. Path and Host Prevalence.

conditions from passive content-distributions. The work in [5] is a cost-effective probing that relies on measurement reusing. This study is a data-reuse and analysis of relay characteristics for stable relay services. The work in [6] examines physical TTL characteristics. This work is motivated by [7] to design stability metrics for the IPv4 relay routing. The research in [8] is an instance of using an HTL-based relay for hybrid UDP-TCP streaming. In [9], the authors argued the need for new estimation-based routing schemes as this paper does. The study in [10] uses the structure-based stability of the relay paths to overcome the routing inefficiency when not considering certain physical links. This study, in contrast, examines both the HTL-based and the path structure-based relay stability. The work in [11] has concluded that the relay performance is limited by the routing resources and host location, which is a motivation for characterizing many granularities.

III. HTL Characteristics

A. Dominant Prevalence

The prevalence of an relay path r is defined by a set of steady-state likelihoods for each relay instance of r. For each dominant set $\mathcal{D}_i(r)$, $p(\mathcal{D}_i(r)|r) = \frac{\omega_i(r)}{m}$ is the corresponding prevalence of such a set. Here, $\omega_i(r)$ is the size of $\mathcal{D}_i(r)$. In our analysis, we examined such a characteristic at each path granularity, such as node, Autonomous System (AS), city, HTL, RTT. We abbreviate each granularity as g_n, g_a, g_c, g_h, g_d respectively. The study in [8] presents the behaviour of the first dominant set's prevalence. This study, however, illustrates the complete view of the decay in prevalence and the estimated host prevalence for each host u at each granularity as i increases. The behaviour of g_d is also briefed for a variety of delay statistics in Section IV-B. Depicting such characteristics is a main contribution of this work. For a dominant set $\mathcal{D}_i(r)$ consisted of a set of relay paths each is with m measurements and originated at u, the host prevalence likelihood is estimated by:

$$p(s|\mathcal{R}_s) = \frac{\frac{1}{m}\sum_{r\in\mathcal{R}_s}\omega_i(r)}{|\mathcal{R}_s|} \tag{1}$$

Both the path and host prevalence represent long-term stability measures. Fig. 1(a) implies that for any path change at g_h, a change at each remaining granularity is possible. By comparing Figures 1(b) and 1(c), the distribution of the first dominant set converges at higher prevalence values since the chance of observing the next dominant sets diminishes the measurement-size. Fig. 1(a) classifies 30% of the examined relay paths as long-lived of prevalence equals one, and at g_n, approximately 51% of paths are dominated by a prevalence $\leq$ 0.75. Clearly, the host prevalence also declines as we examine the next dominant sets since equation 2 is an average of the path prevalence explained earlier. Nearly about 70% of hosts have an average prevalence of 0.8 at each granularity. Unfortunately, our dataset is not spanned over long periods of time in order to confirm the g_h behavior for detecting path changes for a given host.

B. HTL Persistence

For physical routing, persistence is more difficult task to be analyzed correctly than prevalence [6]. This because a

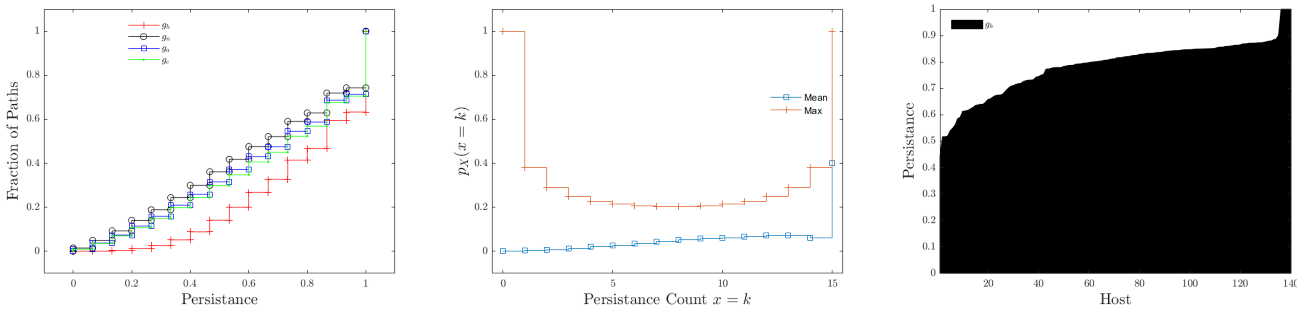

(a) Path Cumulative Persistence

(b) Path HTL Binomial Persistence

(c) Host HTL Persistence

Fig. 2. Path and Host Persistence.

serious of measurements for a particular route do not indicate a lack of change [6] over a wide range of time scales from seconds to hours. This section focuses on the relay persistence of the minimum delay paths to understand the implications of the changes that occur within short-time scales. For relay paths, we argue that capturing short-lived paths is useless because of the challenge to implement them at short scales. Therefore, our analysis only considered changes within the range $[10 \to 20]$ minutes. For a relay route r with a set of observations, r_1, r_2, r_1, r_1 $\ldots$, r_n, we refer to the estimated persistence or no-change probability by $\hat{p}_r$. Fig. 2(a) describes the the cumulative distribution of $\hat{p}_r$. Clearly, there is a similarity between Figures 1(a) and 2(a). This similarity and small persistence below 0.1 is because of the high prevalence (but not maximum) of the first dominant that is responsible for all path changes. The larger persistence when $\hat{p}_r = 1$ is for relay paths with a single dominant path. Fig. 2(a) explains that about 50% of relay paths have $p_r \leq 0.4$. Having an estimate for $\hat{p}_r$ leads to model the likelihood $p_X(x)$ of noticing k no-change transitions within a set $\mathcal{O}(r)$ of relay observations. This likelihood is modeled as Binomial distribution: $X \sim Bin(|O(r)|, \hat{p}_r)$ as in 2, where $|O(r)|$ is 15.

$$p_X(x=k) = \binom{|O(r)|}{k} \hat{p}_r^k (1-\hat{p}_r)^{|O(r)|-k} \qquad (2)$$

For our measurements, the overall network mean and maximum view of $p_X(x)$ is in Fig. 2(b). The $\simeq 0.04$ average likelihood again proofs the lack of persistence in relay routing with high dominant prevalence. The $p_X(x)$ is high when $k = 15$ for paths with a single dominant prevalence, with which path osculations are rare. The maximum curve of $p_X(x)$ is concave with only 296 paths have zero likelihood for noticing a persistent transition while 5013 single dominant paths with $p_X(x) = 1$. However, for a host u, the average persistence at g_h of its relay paths within $[0.5 \to 1]$. This indicates that all 140 hosts do exhibit persistent relay paths.

C. Path's Node and HTL Symmetry

Routing symmetry is an important characteristic for Internet. The existence of asymmetric paths highlights difficulties of providing a consistent topology for Internet [6]. The asymmetries increase as we examine paths at fine-grained granularities. For instance, overlay paths are expected to be more asymmetric at g_n than they are at g_c. Here, we determine our path symmetry measure $N_s(a,b)$ for two paths between hosts u and v as follows: For $r_1 : u, n_1, \ldots, n_n, v$ and $r_2 : v, n'_1, \ldots, n'_m, u$, we define $\mathcal{R}_{r_1} = \{n_i \mid i \in [1 \to n]\}$ as a set of overlay nodes involved in r_1 and similarly $\mathcal{R}_{r_2}$. Therefore, our path symmetry is evaluated as:

$$P_s(u,v) = \frac{|\mathcal{R}_{r_1} \cap \mathcal{R}_{r_2}|}{|\mathcal{R}_{r_1} \cup \mathcal{R}_{r_2}|} \qquad (3)$$

The symbol $|\cdot|$ refers to the set cardinality. However, measuring such a symmetry requires simultaneous measurements on both directions. Instantaneously, using a centralized

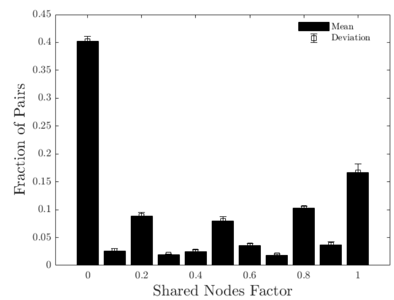

(a) Node Symmetry

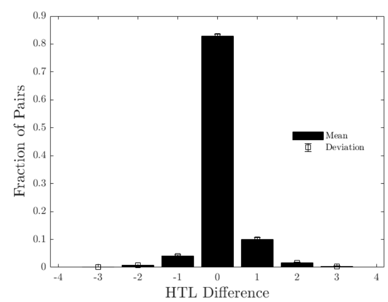

(b) HTL Symmetry

Fig. 3. Path's Node and HTL Symmetry.

probing control, we were able to synchronize Pings of four distinct loads between end-hosts. The study in [6] shows that 49% of measured physical paths were asymmetric by at least one different city. Our symmetry analysis at g_n shows that 40% of our 9,730 observed paired relay paths are completely asymmetric, and only 13% are fully symmetric as Fig. 3(a) demonstrates. Despite the increased asymmetry of the relay routing at g_n, relay paths are highly symmetric at g_h as Fig. 3(b) illustrates. For a particular pair of relay-ends, (u, v), the HTL symmetry factor is simply determined by absolute difference in HTL between both directions of $u \rightleftharpoons v$. Fig. 3(b) shows 82% of our relay measurements are symmetric at g_h. This symmetry factor is approximated by a Laplace distribution with parameters, $\mu = 0$ and diversity $\beta = 0.6$. Thus, $p(x|\mu, \beta) = \frac{1}{2\beta}\exp\{-\frac{x-\mu}{\beta}\}$ approximates the probability mass-function of x that depicts the HTL symmetry difference.

IV. RTT Characteristics

The authors in [8] discussed the use of HTL for detecting path changes. In contrast, this section discusses the use of RTT for detecting relay changes, and new relay RTT characteristics, such as the relationship between RTT and the path symmetry at g_n. Further, we discuss the behavior of physical prevalence when using distinct RTT statistics.

A. RTT and Change Detection

This subsection discusses the use of RTT for detecting relay changes. This analysis shows a symmetrical prevalence by having the g_n curve prevalence in-between the g_d and g_h curves. This causes g_n changes detectable by either one. However, using g_d introduces a False Positive (FP) error comparing to HTL. Table I summarizes the used RTT classification policies that reduce the False Negative (FN) and the Total Error (TE) rates by more than a half as in Table II. The accuracy of such a classifier could be enhanced by imposing extra physical RTT constraints and forming a hybrid classification of HTL and RTT.

TABLE I. RTT CLASSIFICATION POLICIES

	Granularity		
	Physical RTT	Relay RTT	Node
FP	No change	Change	No change
FN	Change	No change	Change

TABLE II. RTT DETECTION ERRORS

Granularity	FN %	FP %	TE %
g_n	16	31	25
g_a	13	32	26
g_c	12	33	26

B. Physical RTT Prevalence

Many protocols consider RTT as an essential metric. For instance, the TCP uses the minimum RTT estimate to determine the Bandwidth-Delay-Product (BDP) that represents the congestion window size or *cwnd*. The maximum RTT, however, is used to decide the TCP's retransmission time. Moreover, streaming applications rely on average RTT as a performance measure. For relay routing, we need to understand which RTT statistic should be considered for constructing relay paths. The minimum RTT curve has the highest prevalence by maintaining 56% of the physical paths frequently on minimum RTTs as Fig. 4(a) proofs. Consequently, a higher host prevalence as in Fig. 4(b) is inevitable since 75% of the hosts have an average prevalence within $[0.5 \rightarrow 1]$. This scenario occurs when the path queue status is not busy so that at least one packet of each measurement is served at a minimum RTT. However, the challenge for physical routing is in maintaining constant RTTs. The maximum RTT prevalence, as a result, could be smaller because of cross-traffics or path changes when balancing loads. Therefore, the TCP could experience additional re-transmission times than the *cwnd* when using the maximum-based RTT mechanism for unacknowledged packets. The average RTT prevalence has a near curve to minimum RTT, and thus, more recommended than the maximum-based RTT re-transmission. Notably, the clear separation between all three prevalence curves is useful when using one RTT statistic to estimate the other statistics.

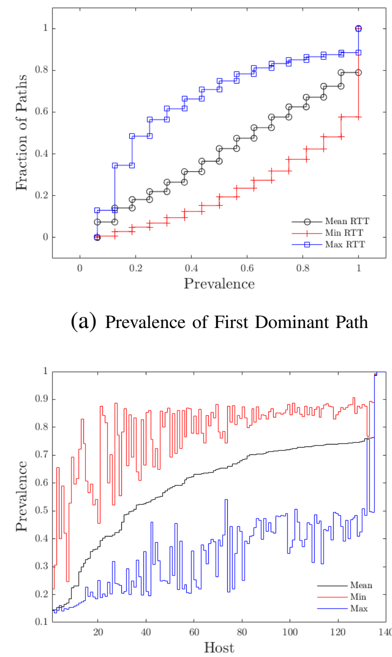

(a) Prevalence of First Dominant Path

(b) Host Prevalence of First Dominant Path

Fig. 4. Path and Host Statistical Prevalence.

C. Path's RTT Symmetry

In Section III-C, we have examined relay path symmetries at g_n and g_h. In contrast, here we establish a definition for whether two paths are symmetric at g_d or not in physical and relay routing. This view is strongly related to the discussion in Section III-C as we have concluded that relay paths are highly symmetric at g_h but weakly symmetric at g_n. This section examines for a particular host u connected with the remaining $n-1$ hosts, the number of paths of $n-1$, in which node u experiences similar RTT values on both directions. The definition of such symmetry factor for node u is:

$$H_s(u) = \frac{\sum^{n-1}\{\mathbb{1}|d(r:u \rightarrow v - d(r:v \rightarrow u) \leq \tau\}}{n-1} \quad (4)$$

The $\mathbb{1}$ is a numeric indicator equals one for any connection $u \rightleftharpoons v$ whose absolute RTT difference between both directions is smaller than a given threshold τ. The larger the τ, the more symmetric a host at g_d is being. Fig. 5 shows a compassion in distribution between the physical and relay RTT symmetries using a $\tau = 0$, 2 and 5 milliseconds. The figure indicates that overlay is more symmetric as τ is relaxed. Fig. 5(a) shows that nearly 14% and $10\%$of the physical and relay hosts respectively behave in full delay-asymmetric manner when $\tau = 0$ millisecond while 18% physical and 31% relay hosts have 0.5 symmetry factor. This factor raises gradually as expected when relaxing τ towards 10 milliseconds. This is because of the high symmetry in relay paths at g_h as discussed in Section III-C. Relay paths, therefore, converge faster than their corresponding physical paths in RTT symmetry to larger symmetry factors due to the highly symmetric HTL behaviour.

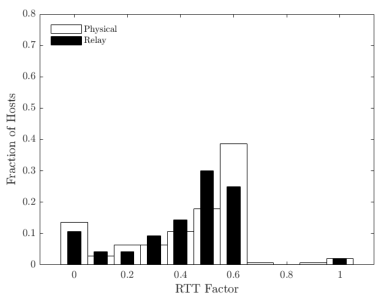

(a) $\tau =$ zero Millisecond

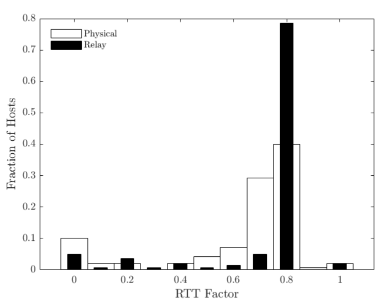

(b) $\tau =$ Five Milliseconds

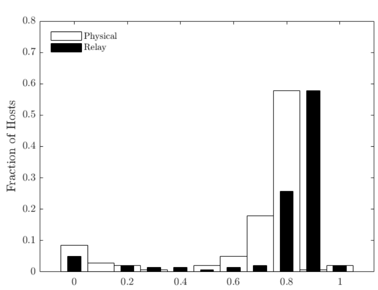

(c) $\tau =$ Ten Milliseconds

Fig. 5. Host's Physical and Relay RTT Symmetry.

V. 9 Conclusion and Future Works

This study shows that relay paths are more stable in terms of HTL, and recommends that a collaboration between HTL and RTT-based routing could lead to inexpensive relay schemes. Nearly, 25% of relay paths have their first dominant prevalence equals 0.5. Further, 70% of hosts have an average prevalence of 0.8 at each granularity. The study shows that the likelihood of observing change-free transitions in relay paths is 0.04. Relay paths have been realized to be highly symmetric at g_h while experiencing a lack of symmetry at every other granularity. HTL as a stability metric detects relay changes by 15% total error while RTT has a 25% total error. The study concludes that the minimum RTT prevalence is more dominant to be noticed in physical paths. In close, this work is planned to include an analysis of the characteristics of the next $k-$shortest RTT paths.

References

[1] J. Jiang et al., "VIA: Improving Internet Telephony Call Quality Using Predictive Relay Selection," SIGCOMM, pp. 286–299, 2016.

[2] S. Mohamed, S. Das, S. Biswas, and O. Mohammed, "On The Significance of Layer-3 Traffic Forwarding," WWIC, Bologna, Italy, pp. 170-181, 2019

[3] J. Kim, A. Chandra and Weissman, "OPEN: Passive Network Performance Estimation for Data-intensive Applications," Technical Report, pp.8-41, 2008.

[4] http://www.akamai.com, retrieved: 09-21-2020.

[5] D. Choffnes and F. Bustamante, "On the Effectiveness of Measurement Reuse for Performance-Based Detouring," INFOCOM, pp. 693-701, 2009.

[6] V. Paxson, "End-to-End Internet Packet Dynamics," In Proceedings of SIGCOMM, pp. 139-152, 1997.

[7] F. Golkar, T. Dreibholz, and A. Kvalbein, "Measuring and Comparing Internet Path Stability in IPv4 and IPv6," In Proceedings of the 5th IEEE International Conference on the Network of the Future (NoF), pp. 1-5, 2014.

[8] S. Mohamed and O. Mohammed, "Towards Stable and Hybrid UDP-TCP Relay Routing for Streaming and VoIP Services," ICNS, pp. 55-65, 2020.

[9] R. Fontugne, J. Mazel, and K. Fukuda, "An Empirical Mixture Model for Large-Scale RTT Measurements," In Proceedings. IEEE INFOCOM, pp. 2470-2478, 2015.

[10] A. Lareida, D. Meier, T. Bocek, and B. Stiller, "Towards Path Quality Metrics for Overlay Networks," IEEE 41st Conference on Local Computer Networks, pp. 156–159, 2016.

[11] J. Han, D. Watson, and F. Jahanian, "An Experimental Study of Internet Path Diversity," IEEE Transactions on Dependable and Secure Computing, pp. 273-288, 2006.