Towards Stable and Hybrid UDP-TCP Relay Routing for Streaming and VoIP Services

Salim Mohamed[1], Ossama Mohammed[2]

[1]Electrical and Computer Engineering, Michigan State University, East Lansing, USA
[2]Service Delivery and Management, Innovaway, Napoli, Italy
[1]Email: mohame26@msu.edu
[2]Email: osama.mohammed@it.ibm.com

September 27, 2020

Authors

[1]Salim Mohamed, Ph.D. candidate at the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, USA. Areas of interest: Routing optimization for wired/wireless networks, machine learning and algorithmic complexity reduction.

[2]Ossama Mohammed, data Communication engineer at the Service Delivery and Management, Innovaway, Napoli, Italy. Areas of interest: Wireless communication and machine learning.

Agenda

Introduction

Background

- Relay or overlay routing for IP networks has been well-documented in past years. However, the implementation cost of relay solutions has not yet been conclusively identified.
- Relay routing introduces new implementation-concerns, such as probing overhead (cost) and its processing latency, stability, availability, and sensitivity to underlay routing changes.
- Dynamic-relay routing relies on periodic probing for enhanced performance while static-relay routing uses less and non-periodic probes to measure latency and packet loss.
- There exists considerable research focused on understanding routing dynamics. However, the literature has insufficient exploration of relay attributes, such as stability and mechanisms for reducing the relay probing burden.
- HTL is introduced in to assist in predicting minimum and stable relay paths while minimizing probing overhead.
- This work has two parts.
  - 1^{st} is an analysis of the stability characteristics of relay (overlay) routing.
  - 2^{nd} is a preliminary, performance evaluation of a new streaming scheme called hybrid UDP-TCP streaming.

Methodology

Platform and Experiments

Platform

Our measurements were conducted using 140 Planetlab nodes distributed as: North America 63.57%, South America 4.29%, Australia 3%, Asia 17.86% and Europe 12.86%.

Experiments

- Ping and IPerf were used to conduct measurements over $19,460$ and $14,762$ end-to-end paths, respectively in a network of $n = 140$ nodes.
- Ping sends its bulk of packets in four distinct sizes: 0.05, 0.1, 0.25 and 0.5 MBytes. Ping packets were also scheduled in the same order 4 times in 16 experiments.
- IPerf datagrams were sent at 12 distinct demand-rates. Our diverse measurements were used to examine the HTL characteristics, and design a stable HTL-based path estimation.
- Having a diverse measurement interval as suggested in [1], provides more confidence in capturing possible routing changes.

Probing

Probing Daemon

Throughout the measurements period, all experiments followed the exact probing abstraction illustrated.

Probing Success

The probability of no node to be measured simultaneously by more than one of the m probers to not influence actual delay. We leveraged our simultaneous active probers m to be as small as possible so that the probability of success is still high:

$$\Pr(\text{success}) \approx \exp\left\{-\frac{m^2}{2(n-m)}\right\}$$

For success demand equals 70% and $n = 140$, $m \approx \frac{1}{2}\sqrt{(2.8534n + 0.50887)} - 0.35667 \approx 10.$

Probing Abstraction

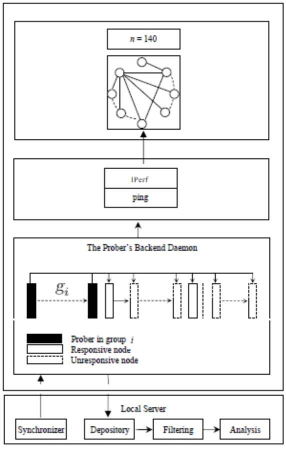

Relay Path Design and Structure

Single Session Relay

- Our design is a network-layer overlay, in which we force **intermediate** Planetlab nodes to not pass traffic to its application-layer.
- Overlay communication can either at application-layer such as application-multicast or network-layer such as IP-multicast.
- Application-layer overlay is a multi-session since each intermediate hop (overlay node) is re-sending the same packet to the immediate-next overlay node (source-destination IP addresses change at each hop).
- On contrast, network-layer overlay allows source to efficiently send its traffic in a single transmission.

Single Path Relay

- Our design is a single path as we use only a unique minimum delay overlay path for any given source-destination pair.
- TCP can force its traffic to traverse along multiple paths instead of one (each with private session).

Relay over Planetlab IP Layer

Vsys

Programable sudo command for privilege allocation tool that:

- Defines privileges at an arbitrary granularity by filtering data between the host and guest domains.

Relay Setup

Encapsulation Examples

GRE and IP-in-IP (IPIP) are two simple and similar tunneling approaches.

Generic Routing Encapsulation (GRE) Encapsulation

- Adds an additional header from 4 to 16 bytes between the original and tunnel IP headers depending on enabled options (default four bytes).
- GRE Features:
  - Encapsulate any network-layer packet not just IP.
  - GRE checksum is not useful only at the final destination in case of TCP, i.e., double check occurs.
  - Specify a tunnel key.
  - Enforce packet sequencing.

Used GRE Packet

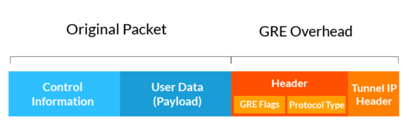

Contribution

Contribution

- The major contribution of our work is in identifying short and long-term analyses of the minimum RTT relay paths, such as long-term stability.
- Introduce HTL alternation sequences.
- Further probabilistic relay attributes, such as prevalence are studied.
- For analyzing HTL characteristics, we performed $311,360$ RTT measurements for all paths in a network of 140 nodes.
- Through extensive analysis, we found that the HTL prevalence shares a similar behavior to the TTL prevalence studied in [1]. Therefore, we conclude that both HTL and RTT are sufficient routing metrics for predicting relay changes and thus, reducing the relay cost.
- This paper proposes two different schemes for using hybrid UDP-TCP as an alternative for TCP-based streaming and VoIP services and shows the difference in performance between the two schemes.

Motivation

Motivation

- The motivation behind using hybrid UDP-TCP instead of TCP is slightly similar to the Quick UDP Internet Connections (QUIC) protocol in [2].
- Relay routing introduces new implementation-concerns, such as probing overhead (cost) and its processing latency, stability, availability, and sensitivity to underlay routing changes.
- Finding how stable is a measurement-less model of relay paths over $24-$hours so that a single instance of underlay measurement could be reused for estimating new stable paths.
- Look into the scale of achieved benefit when implementing our estimation model into real-time scenarios, such as a path serving for the hybrid UDP-TCP protocol.

Results

HTL Stability

Discussion

- Since paths with HTL ≤ 4 hops are dominant in logical routing, we found that they either prefer to remain at constant HTL or switch to shorter HTL counts.
- Their tendency to reduce the number of hops is uniform, in particular for paths of HTL ≤ 3 hops. Therefore, the focus should be on HTL $\in [2 \rightarrow 4]$ hops when designing stable relay routing. Beyond 4 hops, paths are less stable in maintaining constant HTL.

Relay HTL Stability

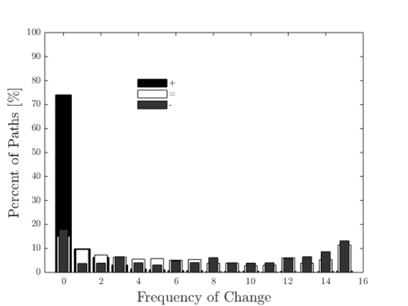

HTL Frequency Sequence

Discussion

- This sequence is called frequency sequence $\mathsf{F}_s(h)$. For a given physical path r, if $\mathsf{M}_s(h) = 3, 3, 5, 2, 4, 5$, then $\mathsf{F}_s(h) = 3, 5, 2, 4$.
- There were $8,863$ paths whose $\mathsf{F}_s(h)$ sequences demonstrated multi-hop paths (relay was always better), and suffered no HTL miss.
- Additional $5,824$ paths with physical candidates still included in $\mathsf{F}_s(h)$ also suffered no HTL miss.
- This indicates about 75.4% of paths change their HTL within a stable $\mathsf{F}_s(h)$ during the measurement period.
- Excluding fixed physical paths, we found only 949 paths suffered HTL miss(es). Such a number is quite small compared to the total of $19,640$ paths.The result details the fraction of paths per each HTL miss.

The Overall HTL Miss

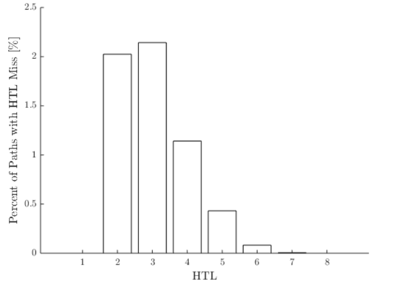

HTL Transition Sequence

Discussion

- The frequency sequence neither investigates the nature of HTL switching, gradual or random, nor how often HTL miss(es) occur.
- This sequence $\mathsf{T}_s(h)$ takes time into consideration for identifying HTL miss(es) that occur between consecutive measurements. Generally, $\mathsf{M}_s(h)$ is a sub-sequence of $\mathsf{T}_s(h)$, and therefore, $\mathsf{T}_s(h)$ can be generated by placing all missing HTLs.
- For HTL $=$ 2 hops, about 90% of our relay paths had no miss(es) during path switching. From the result, we can conclude that nature of HTL switching in relay follows a gradual transition rather than a random one.

HTL Transition Miss Breakdown

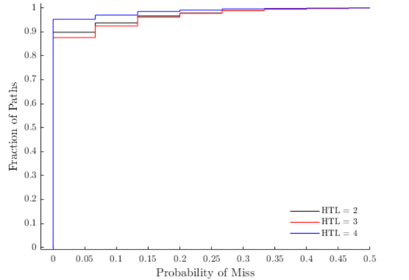

Dominant Path Prevalence

Discussion

- The prevalence is defined by the steady-state likelihood of the most frequent relay observation of a path during measurement.
- We examined prevalence at each routing granularity of the following: Node, Autonomous System (AS), city, HTL, RTT in order to determine how stable is.
- The result shows path fluctuations at g_a, and implies changes at g_c and vice-verse for both dominant sets, since prevalence at g_c is always strictly below the one at g_a.
- Since g_n curve is strictly above g_a, g_c and g_h, with 100% any changes at g_n are as well reflected at any other granularity.
- On average nodes have 70% of their paths as considerably stable.

Dominant Path Prevalence

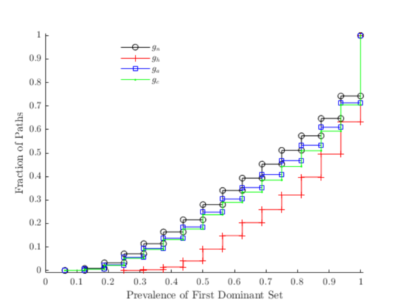

Detection of HTL Changes

Discussion

- The proposed HTL mechanism for detecting relay changes requires including the HTL field in the encapsulated relay header. As a result, relay nodes can determine if a change has occurred without extra probing.
- In the result all associated False-Positive (FP), False-Negative (FN), and total errors when relying on relay routing at g_h for predicting actual relay changes. The FN rate can be reported at every granularity.
- The HTL does not result in FP at g_n since any no-change in the nodes involved in a relay path will not be reported as a change by HTL.

TTL vs. RTT Reduction

Table: HTL CHANGE DETECTION

Granularity	FN %	FP %	Error %
g_n	41	0.00	15
g_a	35	0.00	12
g_c	33	0.00	10

HTL-Based Relay Estimation

Discussion

The figure represents our approach to handling UDP-TCP streaming requests. The controller passes the requester ID to a relay path selection, and receives back all possible stable bandwidth relay paths. The controller then compares the performance of the selector's returned sub-topology in order to choose a relay scheme for the request.

HTL-Based UDP-TCP Request Handling

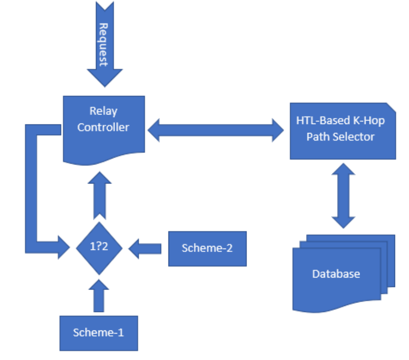

Bandwidth Evaluation

Expensive Probing

- For many applications, an acceptable packet loss in UDP streams is $\%1$.
- We evaluated our bandwidth under different drop-rates, and examined the change in our measurements as τ varies.
- For a path p, we defined $b_p = \arg \underset{i}{max} \quad \{r_i | d_p <= \tau\}$ where r_i is the rate-demand, d_p is the drop-rate of p and τ is the drop threshold.
- However, evaluating b_p is based-on IPerf accuracy but expensive in probing.
- The result shows the cumulative gain of the expensive b_p as τ varies for all $147, 62$ paths. The increase in bandwidth refers to the difference between the physical and the relay b_p.
- Slow CDF convergence indicates more paths gaining more bandwidth this scheme.
- We noticed that few paths gained higher relay bandwidths as τ increases.
- The range of $[0 \rightarrow 100]$ Mbps seems to be the dominant bandwidth gain.

Bandwidth Evaluation - Expensive Probing

Expensive Bandwidth

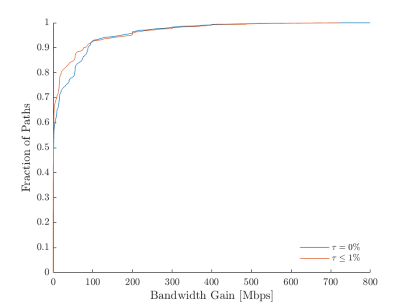

Bandwidth Evaluation

Inexpensive Probing

- This paper provides a less expensive probing relaying scheme by:
  - Choosing a smaller set of demand-rates to evaluate bandwidth.
  - For a Youtube rate-demand, a centralized controller will receive a request from a client with a specific demand, and then probe stable HTL relay paths within the network at the requested demand-rate.
  - This scheme is less expensive in terms of the probing overhead than the previous one.
- This scheme is less expensive in terms of the probing overhead than the previous one. This result compares the performance from the previous expensive scheme with only a single set of measurements at a user rate-demand of 800 Mbps.
- As we reduce our drop-rate demand, or equivalently, increasing τ, we noticed within $[0 \rightarrow 100]$ Mbps, as expected the performance of the inexpensive surpasses its counterpart as the CDF of the later converses earlier.
- Using a less-probing overhead, quickly the inexpensive relay scheme is able to serve UDP-TCP streams at higher rates without exactly determining the available bandwidth for each path.

Bandwidth Evaluation

Expensive vs. Inexpensive Probing (The Upper for $\tau = 0$, and Lower for $\tau \leq 1$)

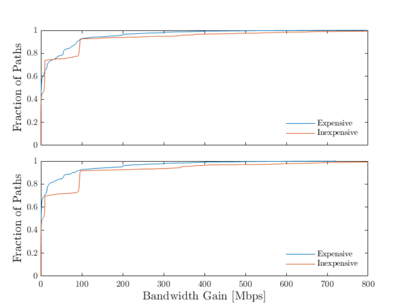

Hybrid UDP-TCP

Discussion

- For physically compromised client-server connections, we have been trying to determine the likelihood of obtaining a relay path with a minimum drop-rate given a particular rate-demand. Through this path, a UDP client will receive the entire stream without involving TCP to resend incorrect or lost datagrams.
- The increase in the likelihood of occurs as we raise the rate-demand.
- This because we find that the current underlay routing is prepared to send UDP streams with a small drop-rates only at low demand-rates.
- Therefore, as the demand-rate increases, packets start to experience more loss, and consequently, hybrid relay routing is a suitable candidate for such an issue.

Hybrid UDP-TCP Success

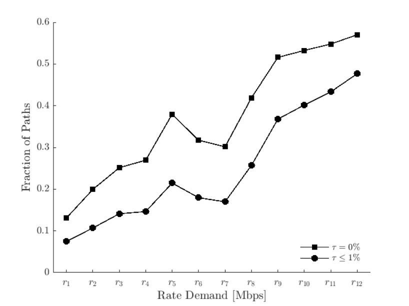

Conclusion

Conclusion

- This study demonstrates that relay paths are more stable in terms of HTL, and results show that an HTL-based relay path selection assists in reducing the search overhead for better paths as services demand.
- Focused reduce the delay within magnitude of $[1 \rightarrow 100]$ ms.
- some services are short-time lived ones, and looking for better paths in large networks might be disadvantageous, especially when the searching time is longer than the service life-time.
- The paper is an attempt for minimizing the propping overhead in overlay schemes and detailing many statistical boundaries for the relay forwarding by analyzing a wide-set of real-time measurements.
- Our work recommends that an HTL-based relay path prediction is able to determine future paths that reduce drop-rates in streaming services when high transmission-rates are demanded.
- This reduction in probing is a result of the self-similarity model of Internet data.
- The study shows that relay paths with HTL ≤ 4 hops are dominant, and they either tend to remain at constant HTL or switch to shorter HTL counts.

References

References

[1] V. Paxson, "End-to-End Internet Packet Dynamics," In Proceedings of SIGCOMM, pp. 139-152, 1997.
[2] J. Iyengar and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," Internet-Draft, `https://tools.ietf.org/html/draft-ietf-quic-transport-29`, retrieved: 09-21-2020.