A Measurement Study of the Origins of End-to-End Delay Variations

Yuval Shavitt
School of Electrical Engineering
TEL AVIV UNIVERSITY אוניברסיטת תל-אביב
shavitt@eng.tau.ac.il
http://www.netDIMES.org
http://www.eng.tau.ac.il/~shavitt

DIMES Status Report

- Now also use PlanetLab (currently mostly PE)
- New agent:
  - New Traceroute
    - Stop using MTR code
  - Paris Traceroute (ICMP & UDP)
  - Bidirectional packet train module
  - Higher measurement rate (5 or 6 per minute)

A Measurement Study of the Origins of End-to-End Delay Variations

Yuval Shavitt and Udi Weinsberg
Tel-Aviv University

Problem Setting

- The Internet exhibits non-stable routes
  - Failures
  - Load balancing
  - Changes incommercial agreements
- This often affect delay, which affects many applications
  - Inconsistent delays (Jitter)
  - Asymmetric delays

Work Goal

- Understand the origins of e2e delay variations
  - Result from existence of multiple routes
    - designed load balancing or transient failures
  - Result of problems within each route
    - intra-route issues (congestion, failures)

Related Work

- [Wang *et al*., Pucha *et al*.] studied the impact that specific routing events have on the overall delay
  - Routing changes result in significant RTT delay increase
  - However, variability is small
- [Augustin *et al*.] examined the delay between different parallel routes in short time epoch
  - Only 12% have a delay difference which is larger than 1ms
- [Pathak *et al*.] studied the delay asymmetry
  - There is a strong correlation between one-way delay changes and route changes

Key Differences

- We study the RTT delay along longer time periods
- Examine the difference of the delay *distribution* between parallel routes
- Focus on the origin of delay variability
  - Within each route (e.g., congestion)
  - Due to multiple routes (e.g., load-balance)

How do we measure?

- Use DIMES for conducting two experiments
  - 2006 and 2009
  - 100 agents measures to each other
  - Broad set of ASes and geo locations
  - Active traceroute (ICMP and UDP)
  - Each agent probes each IP address twice every two hours
  - 4 days of probing
  - Collect the route IPs and e2e delay

Agent Statistics (1)

- 2006
  - 102 agents
  - Million traceroutes
  - 6861 e2e pairs
  - VPs in North America (70), Western Europe (14), Australia (10), Russia (6), Israel (2)
- 2009
  - 105 agents
  - Million traceroutes
  - 10950 e2e pairs
  - VPs in Western Europe (41), North America (38), Russia (14), Australia (4), South America (2), Israel (2), Asia (4)

Agents Statistics (2)

- 2006
  - 18% tier-1
  - 78% tier-2
  - 3% small companies
  - 1% educational
- 2009
  - 14% tier-1
  - 58% tier-2
  - 28% educational

Only 7 agents participated in both

Identifying Routes and Pairs

- Using community-based infrastructure:
  - Routes can start and end in private IP space
  - Users can measure from different locations
- Only the routable section of each path is considered
  - The source (*S*) is the first routable IP
  - The destination (*D*) is the last routable IP

Some Accounting

- The e2e pair $P_i=(S,D)$ contains all the routes that were measured between *S* and *D*
- For pair P_i, each route *j* was seen in $|E^i_j|$ different paths
- For pair P_i, the dominant route E^i_r is the route that was seen the most times
  - There can be several dominant routes with equal prevalence
  - For brevity we assume there is one at index *r*

What do we measure?

- Stability of e2e routes
  - Use Edit Distance (ED) as a measure for difference between two routes
    - Counting insert, delete, and substitute operations
  - Normalize ED by the maximal route length
    - Can compare between ED of routes with different length
    - $\widehat{ED}^{i}_{jr}$ marks normalized ED for pair *i* between routes *j* and *r*

What do we measure?

- Stability of e2e routes
  - The stability is the weighted average of ED of all non-dominant routes to the dominant route of nearest length:

$$RouteISM_i = \sum_{j \neq r} \left(|E^i_j| \cdot \widehat{ED}^{i}_{jr} \right) / \sum_{j \neq r} |E^i_j|$$

  - A second stability measure is the prevalence of the dominant route

What do we measure?

- Stability of RTT delays
  - Each route E^i_j has a set of RTT delays, corresponding to each measured path
  - Treat each delay value as a *sample*, consider the 95% confidence interval surrounding the mean delay – *CI(*E^i_j*)*
  - High variance samples result in long CI

What do we measure?

- Stability of RTT
  - RTT stability of a two routes is the intersection between their CI's, normalized by the minimal CI

$$\widehat{O}^i_{jk} = \frac{CI(E^i_j) \bigcap CI(E^i_j)}{\min\{|CI(E^i_j)|, |CI(E^i_j)|\}}, \forall j \neq k$$

Key Concept

- Overlapping CI's (left)- intra-route delay variance
- Non-overlapping (right) - inter-route delay variance

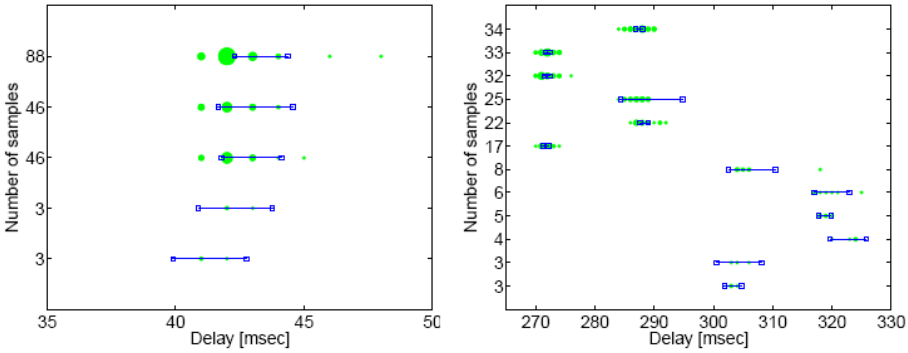

(a) 217.0.116.82 → 80.91.184.206 (b) 69.134.208.1 → 134.159.160.58

Take Home Message

- For 70% of the pairs and for over 95% of the academic pairs, the delay variations are mostly <u>within the routes</u>
- Internet e2e routes are mostly stable, however these intra-route delay variations still affect application!

Things to Note

- We measure RTT values
  - Capture forward and reverse path delay
  - Stability is only on the forward path
    - However, 90% of our routes have very similar forward and reverse paths
    - Indicating that stability of one-way is a good estimation
- Using UDP and ICMP
  - Capture all possible routes, not flows
  - Upper bound for instability

Results – Route Statistics

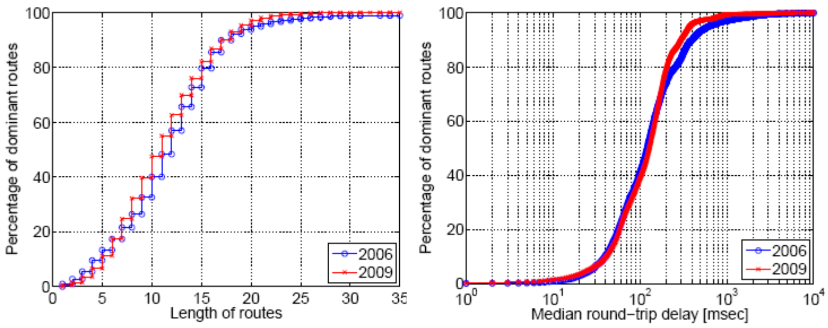

(a) Route Length (b) Median Delay

- Both have roughly the same route length and median delay
- Shorter routes than Paxson's (11-12 hops)

Results – Route Stability

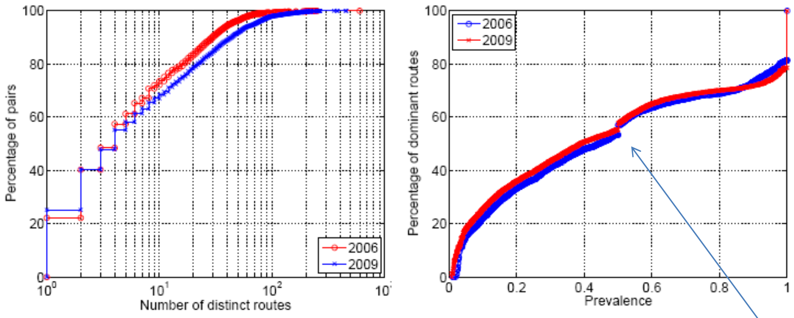

(a) Distinct routes per pair (b) Dominant route prevalence

- Overall stable ie2e routing
- Stability slightly increased over time
- Academy pairs have higher stability
- USA pairs have slightly higher stability
- RouteISM < 0.2 for over 90% of the pairs

Load balancers
Not visible in academy pairs

Results – Origin of Delay Instability

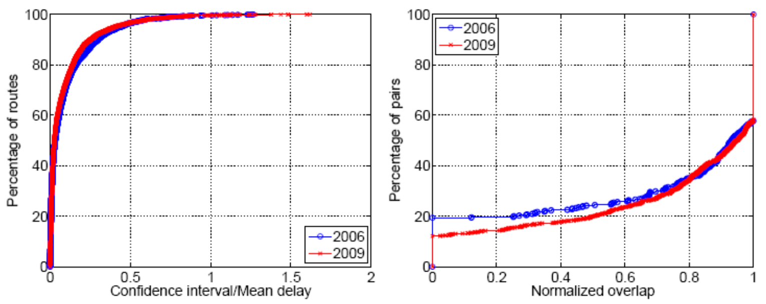

(a) Confidence interval/mean delay (b) Normalized overlap

- The delay confidence interval are not “too long”, and extend only for routes with high variance
- 70% of the cases, changes in route delay cannot be attributed to multiple path routing, but rather to changes between the routes
- In 15% of the cases (20% in the 2006 data sets) the change in delay is mainly due to route changes

Results – Route and Delay Instability

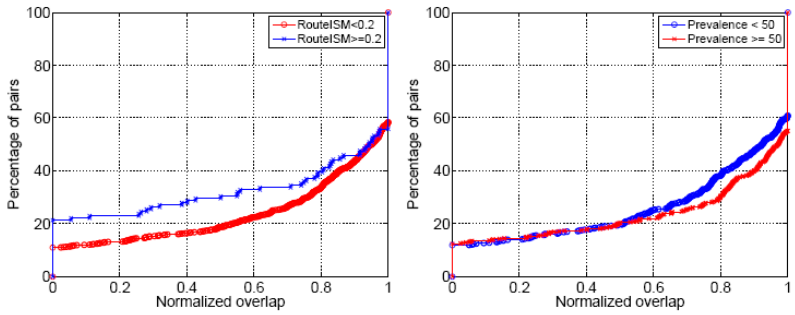

(a) RouteISM (b) Prevalence

- When the difference between routes is high, higher chances of different delay distribution
- Prevalence does not significantly indicate level of overlap!

Results – Additional Findings

- Over 95% of the pairs that have academic source and destination ASes have an overlap of over 0.7
  - Academic networks having small route difference induced by local load-balancing and little usage of "spill-over" backup routes
- Only 5% of the pairs that have both source and destination in the USA witnessed overlap of 0

Conclusions

- A measurement study of the e2e delay variance and its origins using overlap of confidence intervals
- Techniques for quantifying route stability
- For roughly 70% of the pairs and for over 95% of the academic pairs, the delay variations are mostly within the routes and not between different routes

DIMES

Thank You!