

On the Diversity, Stability and Symmetry of End-to-End Internet Routes

Yaron Schwartz, Yuval Shavitt, Udi Weinsberg
Tel-Aviv University

Problem Setting

- The Internet is an extremely complex communication network
 - Evolving structure
 - Failures
 - Load balancing
- Directly affects on e2e data flow
 - Diversity – multiple routes between end points
 - Stability – consistency of routes
 - Symmetry – routes in opposite directions

Work Goal

- A measurement study aiming to quantify various aspects of e2e routes
 - Diversity, Stability and Symmetry of e2e routes
 - What has changed since Paxson's work (1995)?
 - Understand the bias in existing work due to VP distribution

Take Home Message

- Internet e2e routes are diverse but stable
- VP types highly affect the results
- Routes are mostly asymmetric but are either both stable or unstable in opposite directions
- Longitudinal analysis shows that diversity and stability are consistent, indicating trade-offs between the Internet growth and changing trends in its connectivity

Related Work

- [Paxson] was the first to study the stability of e2e flows
 - Conducted active probes from 37 VPs, mostly academic, back in 1994 and 1995
 - Found a relatively stable Internet.
- [He *et al.*] and studied asymmetry
 - Low levels of routing asymmetry in AS level
 - Few end-points are “to blame”
- [Rexford *et al.*] studied stability of popular prefixes
 - Popular prefixes have outstanding stability

Key Differences

- Leverage a broad and diverse set of VPs
 - Over 100 VPs in various AS types
- Attempt to discover all possible paths
 - Inducing per-packet and per-flow load-balancing
 - Use both ICMP and UDP probing
- Use fix time intervals in two experiments

Key Concept

- Repeat active measurements between source and destination hosts
- Resolving IP-level paths to higher levels of granularity (prefix, AS, City and Country)
- Quantify the diversity, stability and symmetry of routes as observed from measured paths

How do we measure?

- Use DIMES for conducting two experiments
 - 2006 and 2009
 - Over 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 hops and e2e delay

Vantage Point Statistics (1)

■ 2006

- 113 VPs
- Million traceroutes
- 7040 e2e pairs
- VPs in North America (79), Western Europe (16), Australia (10), Russia (6), Israel (2)

■ 2009

- 107 VPs
- Million traceroutes
- 10408 e2e pairs
- VPs in Western Europe (49), North America (35), Israel (9), Russia (5), Australia (3), South America (3), Asia (3)

Vantage Point Statistics (2)

■ 2006

- 16% tier-1
- 80% tier-2
- 1% educational
- 3% small companies

■ 2009

- 7% tier-1
- 69% tier-2
- 21% educational
- 3% small companies



Only 15 agents participated both

Filtering and Processing

- Removed traceroutes
 - Only non-routable IP addresses
 - Traceroutes with repeating IP addresses
 - Traceroutes with loops (IP and AS level)
 - 400k traceroutes in 2006 and 800k in 2009 remain
- Resolution to higher level
 - Prefix and AS-level using RouteViews and WhoIS
 - Geographic resolution using MaxMind

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 $|E_j^i|$ different paths
- For pair P_i , the dominant route E_r^i 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
 - Prevalence is the overall appearance ratio of a route j of pair P_i

$$Prevalence_j^i = |E_j^i| / \sum_{j=1}^{k_i} |E_j^i|$$

- As a stability measure, we use the prevalence of the dominant route 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}_{jr}^i 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:

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

What do we measure?

- Symmetry


- Differential RouteISM is the difference between the RouteISM in opposite directions
- SymRouteISM is the normalized ED between the dominant route of one direction and the inverted dominant route of the opposite direction

$$SymRouteISM_i = \widehat{ED}(R(E_r^i), \overleftarrow{R}(E_r^j))$$

Things to Note

- Using UDP and ICMP
 - Capture all possible routes, not flows
 - Upper bound for instability


Results – Route Statistics



(a) Dominant route length

- 2006 and 2009 have roughly the same route length, similar to those found by Paxson's
- Probably a tradeoff between the increase in topology size with richer connectivity and layer-2 tunnels


Results – Route Diversity



(b) Distinct routes per pair

- Roughly 25% of the pairs in both years have exactly one route
- Roughly 30% of the pairs witnessed more than 10 different routes
- A higher level of stability than reported by [Pucha *et al.*] (using a 20 days study, only 6% had one dominant route)
 - The longer study reduces the chances to see the exact same route


Results – Route Stability




(c) Dominant route prevalence

- 100% prevalence is attributed to single route
- 8% increase in the 50% prevalence is attributed to load-balancing or prolonged flaps

Results – Route Stability




(a) 2006




(b) 2009

- IP-level is least stable
- 2009 is more stable for all levels
- Although the Internet is increasing in topology size, the overall stability slightly improves over time
 - Improving knowledge of operators?
 - Easier management of devices?

Results – VP Bias




(a) Prevalence




(b) Stability

- Routing in academic networks is much more stable
 - 20% in commercial and 30% in academic have single route
 - Slightly observed using RouteISM of the AS-level
- No usage of load-balancers in academic pairs
- Cross-continent pairs are slightly more stable
 - Since there are not many alternative routes
- Important to have diverse set of VPs

Results – Symmetry



(a) RouteISM symmetry



(b) Differential RouteISM

- IP level is obviously not symmetric
- AS-level routes are more symmetric than city-level routes
 - Indicates the existence of points-of-presence (PoPs) that belong to the same AS but reside in different cities
- Differential stability
 - Approximately 90% of the pairs have differential stability of less than 0.3
 - When instability exists in one direction, it is likely to appear in the opposite direction as well

Conclusions

- A measurement study of e2e routes
- The Internet today is less stable than 1995
 - Still exhibits different behavior depending on network type
- Longitudinal analysis shows that e2e route properties did not significantly change in recent years
- We attribute this to a trade-off
 - Increasing topology size of the Internet and usage of load balancers
 - The adoption of tunneling technologies that result in more stable IP-level routes

Thank You!