Stability and Symmetry of Internet Routing

Udi Weinsberg
School of Electrical Engineering
Tel-Aviv University, Israel
Email: udiw@eng.tau.ac.il

Yuval Shavitt
School of Electrical Engineering
Tel-Aviv University, Israel
Email: shavitt@eng.tau.ac.il

Yaron Schwartz
School of Electrical Engineering
Tel-Aviv University, Israel
Email: yaronsch@eng.tau.ac.il

I. Introduction

The end-to-end (e2e) stability of Internet routing has been studied for over a decade. In his seminal work from 1995, Paxson [1] used over 40,000 e2e routes from 37 vantage points (vps) for identifying common route pathologies and symmetries. Both Paxson's work and more recent work [2] use data originating from vps that are mostly located in academic networks, possibly contributing to biased results [3]. Moreover, the metrics used by these works for measuring stability compare distinct elements between paths and ignore the order in which the elements appear.

In this work, we analyze the e2e route stability of more than 100,000 routes, measured from over 100 broadly distributed vps. We study the path stability and symmetry in 6 levels of granularity: router, point of presence (PoP), address prefix (AP), autonomous systems (AS), city and country, using a novel metric for stability which is based on the edit distance algorithm. Additionally, we study the spatial properties of Internet stability and show that different geographical regions have different stability properties.

The data used in this paper is obtained from DIMES [4], a community-based Internet measurement system that performs active probing using software agents installed on users' PCs. This data provides a broad diversity in the types and geographical locations of observed ASes, enabling us to better capture the real Internet e2e routing behavior.

II. Experimental Setup

Using the DIMES experiments planner, we selected 120 globally distributed agents and designed 96-hours experiments in which each agent executes UDP and ICMP traceroute measurements to all other agents. Out of the possible 120 agents, 104 agents returned over 1 million traceroutes (IP-level paths), providing us with 6,942 source-destination directed pairs (an average of 144 measurements per source-destination pair). For the purpose of this analysis we took only the first 96 measurements which account for 4 days.

Raw traceroute data was filtered by removing traceroutes that have only non-routable IP addresses. In addition we removed traces that exhibit known traceroute problems [5], namely routing loops and destination impersonation. Resolving IP paths to higher levels of granularity is achieved using several data sources: PoP indices are inferred using iPlane [6], address prefixes and AS numbers are resolved using RouteViews [7] and geographic locations are extracted from RIR databases (mainly WhoIs [8]).

Agents were selected for the experiments according to the following criteria – recently active, distributed in a large set of ASes and distributed in a large set of geographical regions. These criteria were selected to achieve e2e routes with diverse lengths that traverse through various ASes spread across different countries and continents. The resulting agents are located in a variety of geographical locations and types of ASes – 78% in USA and Canada, 11% in Western Europe, 5% in Australia and New Zealand, 3% in Russia and 3% in Israel. 29% of the agents are in tier-1 ASes, 55% are in tier-2, 3% are in small commercial organizations and 1% are located in academic networks (the remaining 12% were not listed in the datasource we used [9]).

III. Analysis Methodology

The method used to quantify the path stability of an e2e route is the Edit Distance (ED) measure. The ED algorithm calculates the number of add, delete and modify operations required in order to convert between two character strings. We modified the algorithm to consider a single-hop IP address, PoP index, prefix, AS number, city and country as a single character and thus the ED value represents the "distance" between two paths.

Previous works used various metrics for attempting to quantify the stability of a path. Paxson [1] simply counted how many different elements are there between two paths of exact length. Pucha *et al.* [2] defined the similarity coefficient as the number of similar elements divided by the total number of distinct elements in the two paths $\frac{|A \cap B|}{|A \cup B|}$. We argue that ED better captures stability since it takes into account the order of elements in each path, resulting in large ED values for paths that have the same items with different order.

Hops that were not resolved properly (due to an unresponsive router or failure in mapping) are used in the ED algorithm as "wild-cards", meaning that they always match a hop in the compared path. This approach was taken in order to get a lower bound estimation of the stability, therefore preventing measurements and analysis difficulties from increasing the estimated instabilities. Since the number of unresolved hops is less than 5%, this technique does not significantly bias the instability analysis.

Measuring the stability of an e2e path (between a source-destination pair) is performed by first finding the dominant route of a source-destination pair, i.e., the route that appears the most times for the given pair. We then measure the ED of all non-dominant routes of the pair to the corresponding dominant route. In cases where there is more than one dominant path, each non-dominant path is compared to the dominant path that has the most similar number of hops. Since different granularity levels results in different path lengths, we normalize the resulting ED value by dividing it with the maximal length of the two compared paths. The result of this normalization is referred to as the *stability measure* (SM) of a path in a source-destination pair for a given level of granularity. The stability measure is in the range of 0-1, where higher values stand for lower stability.

IV. Preliminary Results

Figure 1 shows the distribution of SM values in the various granularity levels. The figure shows that over 50% of paths are stable (SM value equals zero) at the various levels, except for IP-level paths. The result that only 10% of the paths are stable at the IP-level is the basis for further analysis. The mutual influence of instabilities at different levels will also be evaluated in future analysis since we expect that instabilities flow in a cascading fashion from coarse to finer paths' granularity.

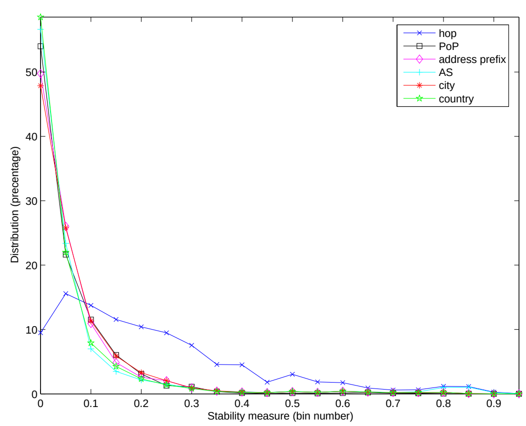

Fig. 1. Distribution of stability measure values

We evaluate the spatial stability of paths in different geographical regions. On average, the IP hop and AS SM values of routes that started and ended in North America are smaller than those of the routes that started and ended in Europe. Routes from Europe to north America and routes from north America to Europe have similar SM characteristics.

Two aspects of routing symmetry are analyzed – checking whether paths in opposite directions are symmetric and analyzing whether instability in one direction of a path indicates instability in the opposite direction. The first is evaluated by calculating the SM value between the dominant path of each pair and the reversed dominant path of the opposite pair. We expect that this analysis will result in highly asymmetric paths mainly due to hot-potato routing [10]. The latter was calculated by finding the difference between SM values in opposite directions of each pair. Figure 2 shows that stability properties are kept for opposing pairs, meaning that the routes in both directions are either both stable or not-stable.

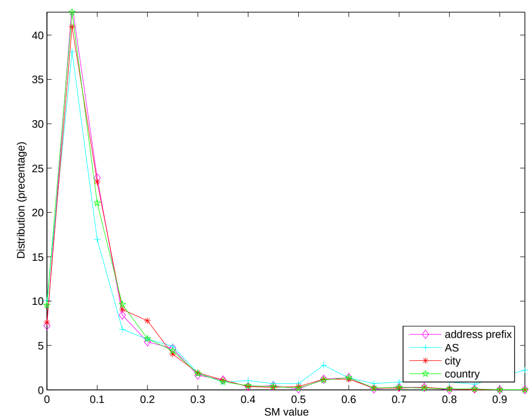

Fig. 2. Distribution of opposite pairs differential SM values

V. Conclusion

This paper presents a large-scale analysis of the stability and symmetry of routing in the modern Internet using adapted edit distance algorithm as a measure for stability. This improves on existing work by capturing changes in the order of hops in paths as well as changes in the hops themselves. Preliminary results show that this measure alongside with diverse and broad active measurements have potential to help understand better the complex routing of Internet packets.

References

[1] V. Paxson, "End-to-End Routing Behavior in the Internet," in *IEEE/ACM Transactions on Networking*, 1996, pp. 601–615.

[2] H. Pucha, Y. Zhang, Z. M. Mao, and Y. C. Hu, "Understanding network delay changes caused by routing events," *SIGMETRICS*, vol. 35, no. 1, pp. 73–84, 2007.

[3] Y. Shavitt and U. Weinsberg, "Quantifying the importance of vantage points distribution in internet topology measurements," in *IEEE Infocom (to appear)*, 2009. [Online]. Available: http://www.eng.tau.ac.il/ udiw/papers/vp.pdf

[4] Y. Shavitt and E. Shir, "Dimes: Let the internet measure itself," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 5, pp. 71–74, 2005.

[5] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute," in *IMC'06*, 2006, pp. 153–158.

[6] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An information plane for distributed services," in *OSDI*, 2006, pp. 367–380.

[7] "University of Oregon RouteViews Project (http://www.routeviews.org/)." [Online]. Available: http://www.routeviews.org/

[8] "WhoIs.Net DOMAIN-BASED RESEARCH SERVICES. http://www.whois.net." [Online]. Available: http://www.whois.net.

[9] X. Dimitropoulos, D. Krioukov, G. Riley, and kc claffy, "Revealing the autonomous system taxonomy: The machine learning approach," 2006. [Online]. Available: http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0604015

[10] T. Renata, S. Aman, G. Tim, and R. Jennifer, "Dynamics of hot-potato routing in ip networks," in *SIGMETRICS*, New York, NY, USA, 2004.