On the analysis of Internet Paths with DisNETPerf,
a Distributed Internet Paths Performance Analyzer

Sarah Wassermann, Pedro Casas, Benoit Donnet

Université de Liège, AIT Austria

sarah.wassermann@student.ulg.ac.be - http://wassermann.lu

Why DisNETPerf?

- Load-balancing: user-requests redirected to servers at hundreds of milliseconds
- Performance-monitoring of reverse paths (uncontrolled server→customer) challenging
- Reverse traceroute[2] relies on IP spoofing and IP Record Option: security concerns

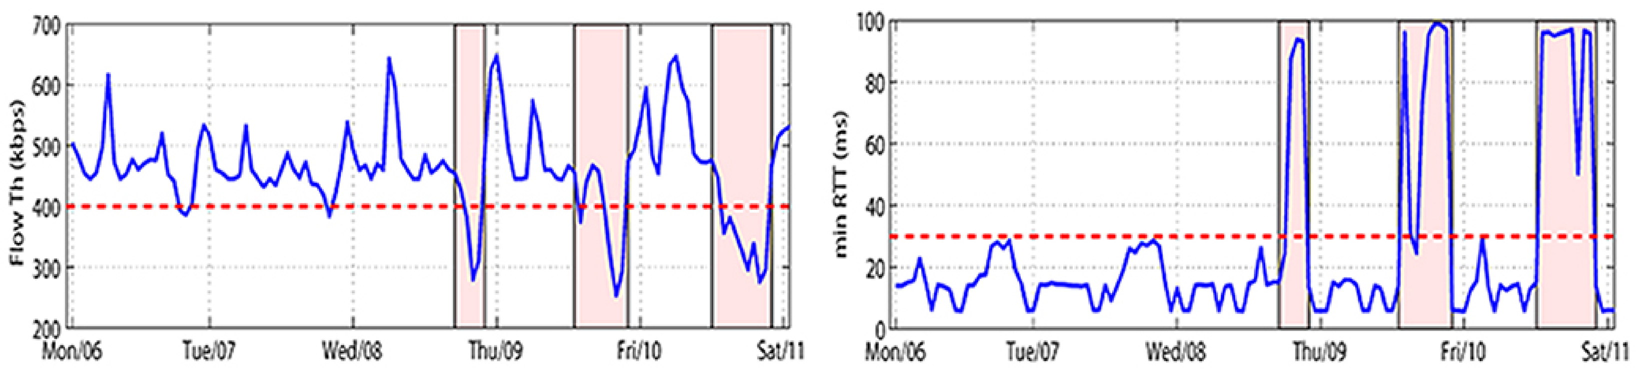

Real case study: Google load balancing policy impacts the QoE of a large number of customers watching YouTube[1]

Path asymmetry study

- Path asymmetry study between RIPE Atlas boxes in different regions of the world
- Asymmetric path: A→B different from B→A
- Metric: *Route Similarity* index (RSIM) - ratio between #shared links and #total links
- Path asymmetry indicated by RSIM(A→B, inverse(B→A))
- More asymmetric paths in larger zones
- Path asymmetry non negligible phenomenon!

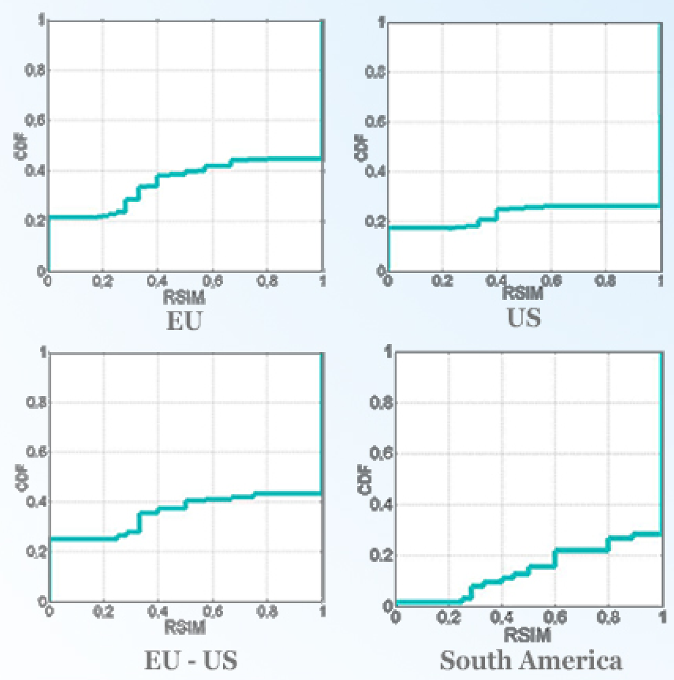

Overview

- DisNETPerf relies on RIPE Atlas (RA) distributed measurements
- Given: content server IP_s, destination IP_d, locate IP_c - closest RA probe to IP_s
- Combined topology- and delay-based distance: first locate candidate RA probes by AS and then elect IP_c by propagation delay
- Periodically run `traceroute` measurements from IP_c to IP_d
- Data can be used to troubleshoot paths from the content server (mimicked by IP_c) to the target customer

Distribution of RA probes

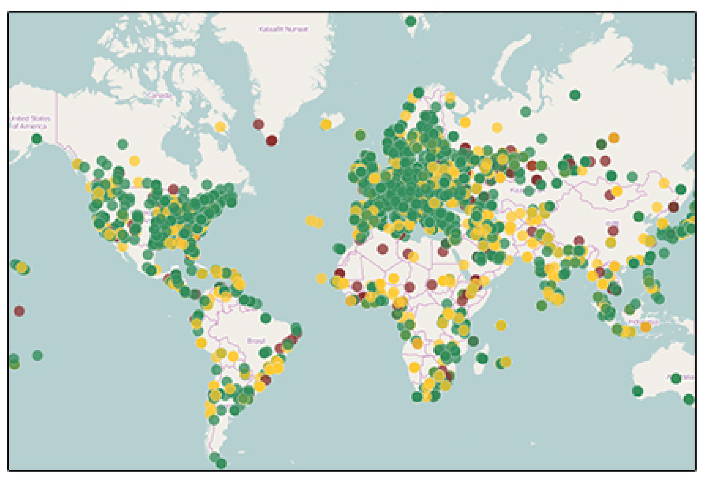

Probe selection: 2 approaches [3]

Smallest latency (SL) approach

1. Launch `ping` measurements from candidate RA probes to IP_s
2. Probe with smallest minimum RTT elected as IP_c

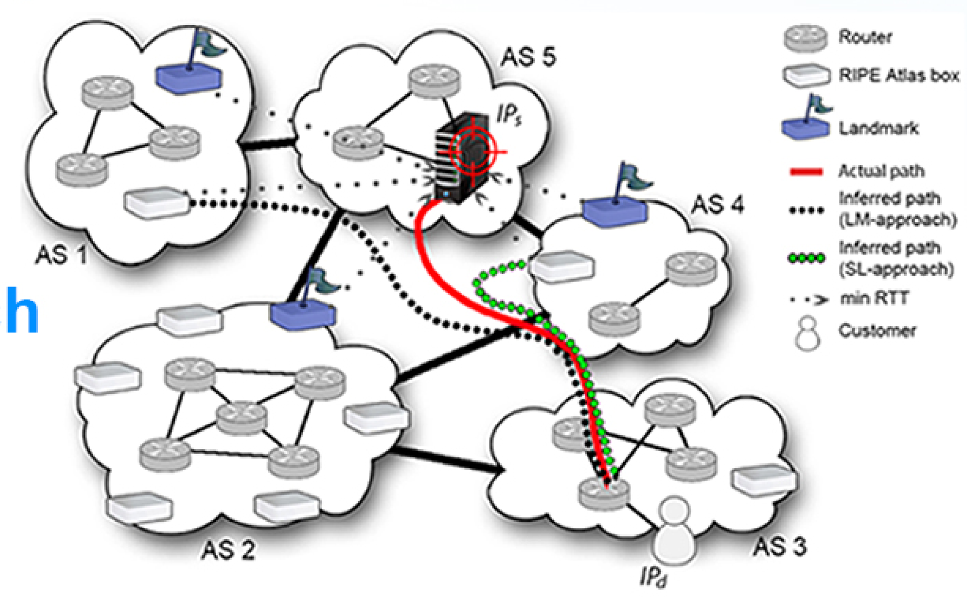

Landmark (LM) approach

1. Use some candidate RA probes as landmarks for distance reference
2. Launch `ping` measurements from landmarks to IP_s and all other candidate RA probes
3. Elect as IP_c probe with smallest normalized distance to IP_s:

$$D_{ij} = \frac{1}{K}\sum_{l=1}^{K}|d_{il} - d_{jl}|$$

where K = #landmarks providing a RTT for both IP_i and IP_j, d_{il} = min RTT between IP_i and landmark l

Evaluation

- IP_c is a *good* probe w.r.t. IP_s and IP_d if the path from IP_c to IP_d is highly similar to the path from IP_s to IP_d
- *Route Similarity* index - $RSIM(IP_c, IP_s, IP_d) = \frac{2 \times C_{links}(IP_c, IP_s, IP_d)}{T_{links}(IP_c, IP_s, IP_d)}$ where C_{links} = #common links, T_{links} = #total links
- Ground-truth: RA probes as IP_s and IP_d to compute real server→customer paths (300 tests)
- RSIM computed for different link definitions (AS, PoP, and IP level)
- AS-level measurements more relevant (troubleshooting faulty ASes)→high RSIM when RA probe located in same AS
- 2 approaches very accurate: select the best RA probe (highest RSIM) for more than 80% of tests

Observed RSIMs for elected RA boxes

Numerical difference between highest observed RSIMs and RSIM of elected RA boxes

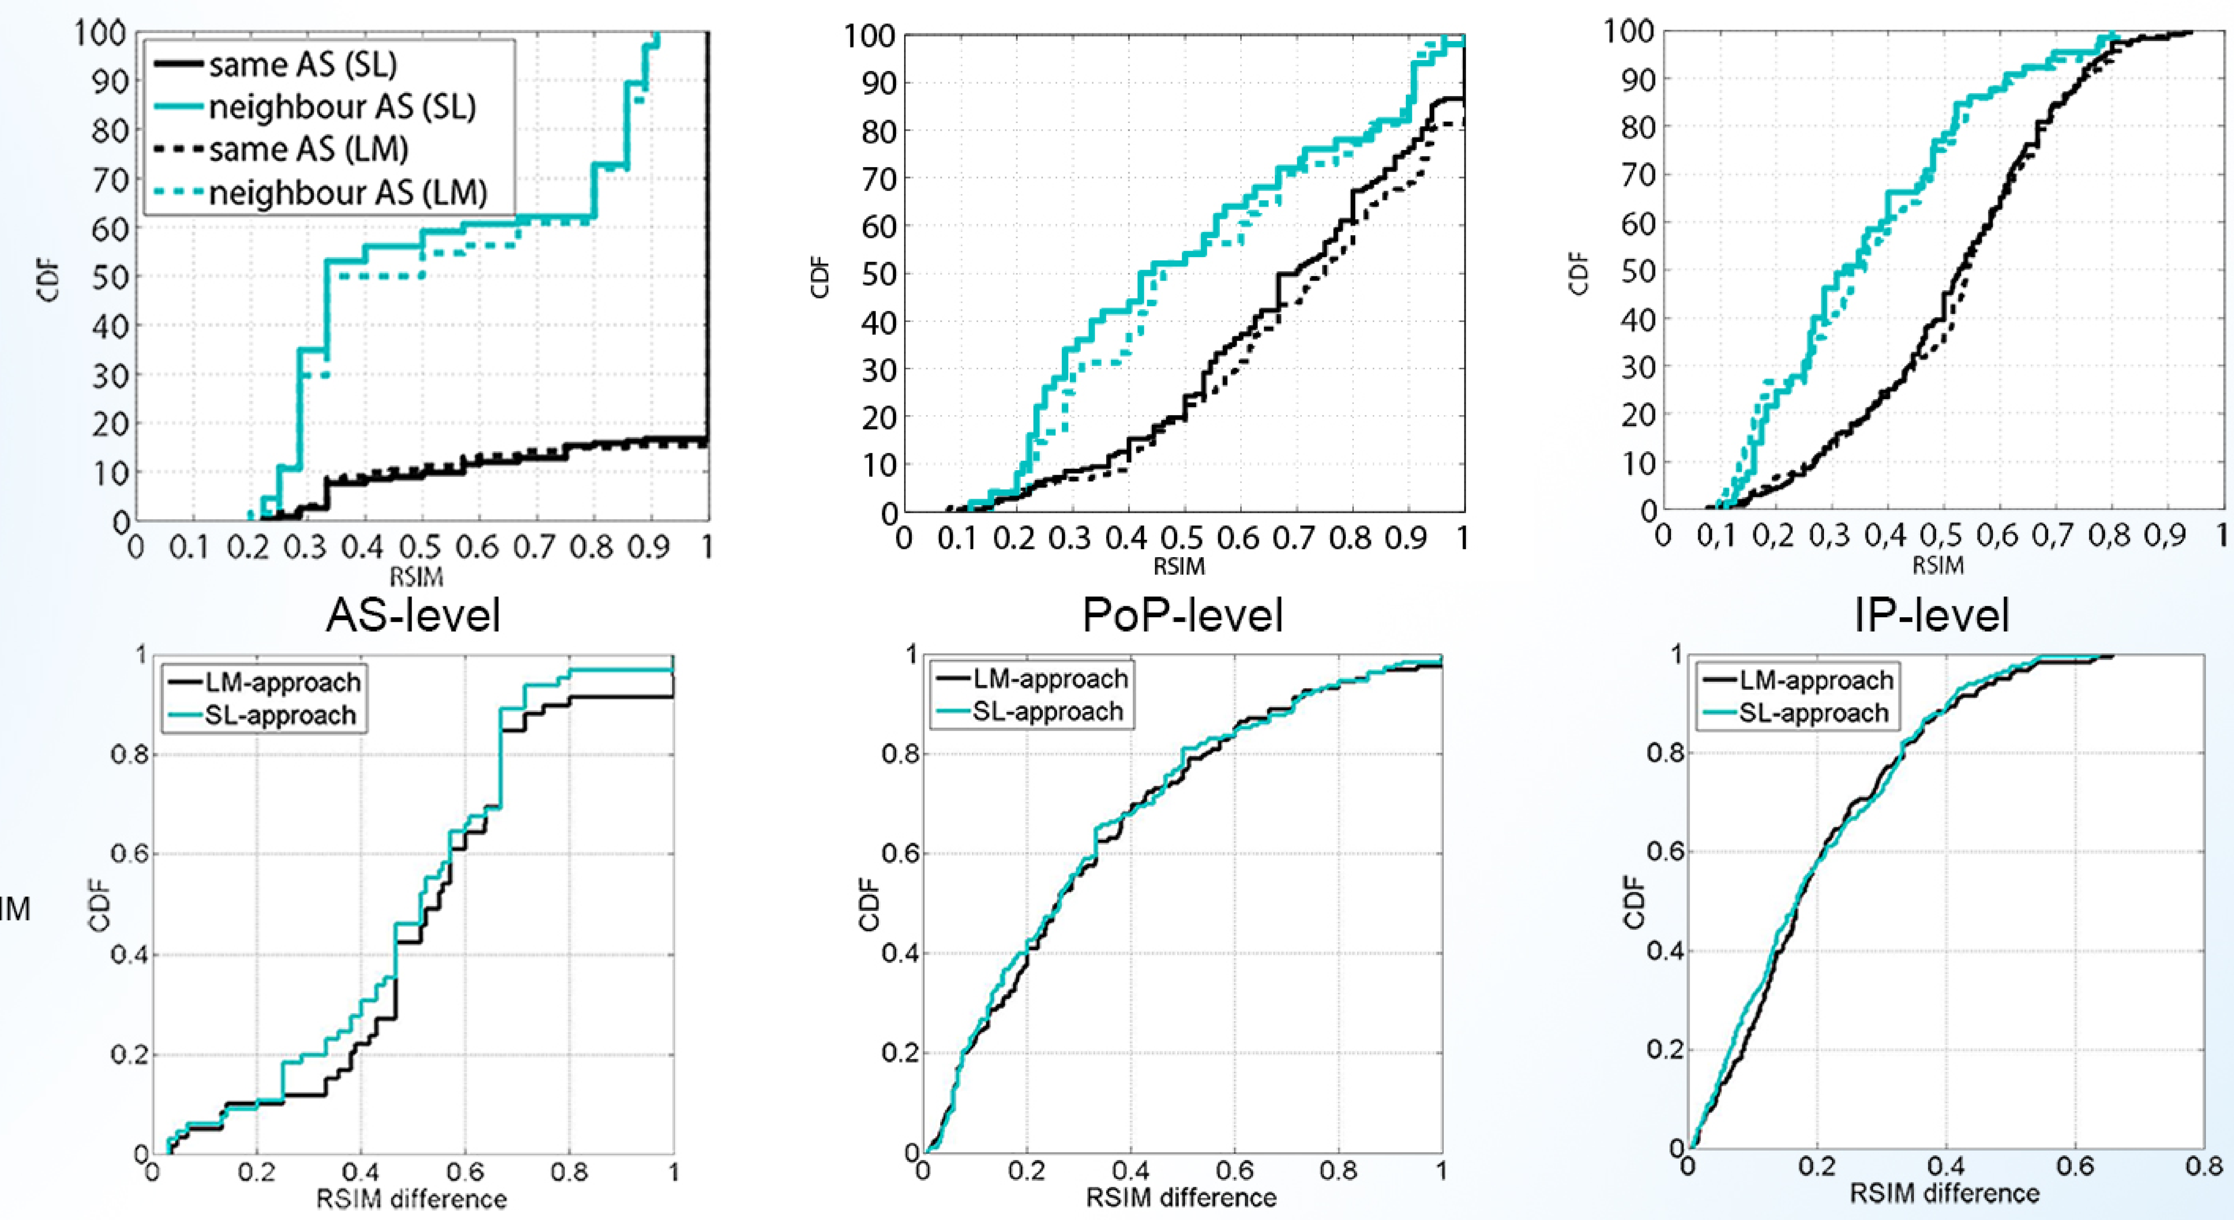

- **Ongoing**: combined analysis of passive (@access ISP) and active measurements (DisNETPerf) for Youtube analysis

References

[1] P. Casas, A. D'Alconzo, P. Fiadino, A. Bär, A. Finamore, and T. Zseby, "When YouTube does not work. analysis of QoE-relevant degradation in Google CDN traffic," *IEEE Transactions on Network and Service Management, vol. 11, no. 4, pp. 441–457*, December 2014.

[2] E. Katz-Bassett, H. Madhyastha, V. Adhikari, C. Scott, J. Sherry, P. van Wesep, A. Krishnamurthy, and T. Anderson, "Reverse traceroute," in *Proc. USENIX Symposium on Networked Systems Design and Implementations (NSDI)*, June 2010.

[3] S. Wassermann, P. Casas, B.Donnet, "Towards DisNETPerf: a Distributed Internet Paths Performance Analyzer", in *Proc. ACM CoNEXT Student Workshop, December 2015*

ACKNOWLEDGEMENTS: This work has been supported by the European Commission through the Seventh Framework Grant Agreements mPlane (FP7-318627).