Mapping an Anycast CDN Using RIPE Atlas

verizon digital media services

University of Glasgow — School of Computing Science

Stephen McQuistin, University of Glasgow
Marcel Flores, Verizon Digital Media Services

Introduction

Anycast CDNs announce the same IP address blocks from different points-of-presence (PoPs), relying upon BGP routing to map clients to these PoPs. This defines *catchments*[2]: the set of clients served by a given PoP. In this poster, we outline a methodology for mapping anycast catchments, and evaluating changes in anycast announcements at a large CDN. Understanding and optimising these catchments is important, given their impact client performance and PoP load.

Data sources

- Traceroute data from RIPE Atlas probes[1].
- Probe-PoP mappings: Use traceroutes and BGP session information from each PoP.
- Test IP blocks: a control block (consistent with current announcements), and an experimental block (with proposed configuration).

Catchments

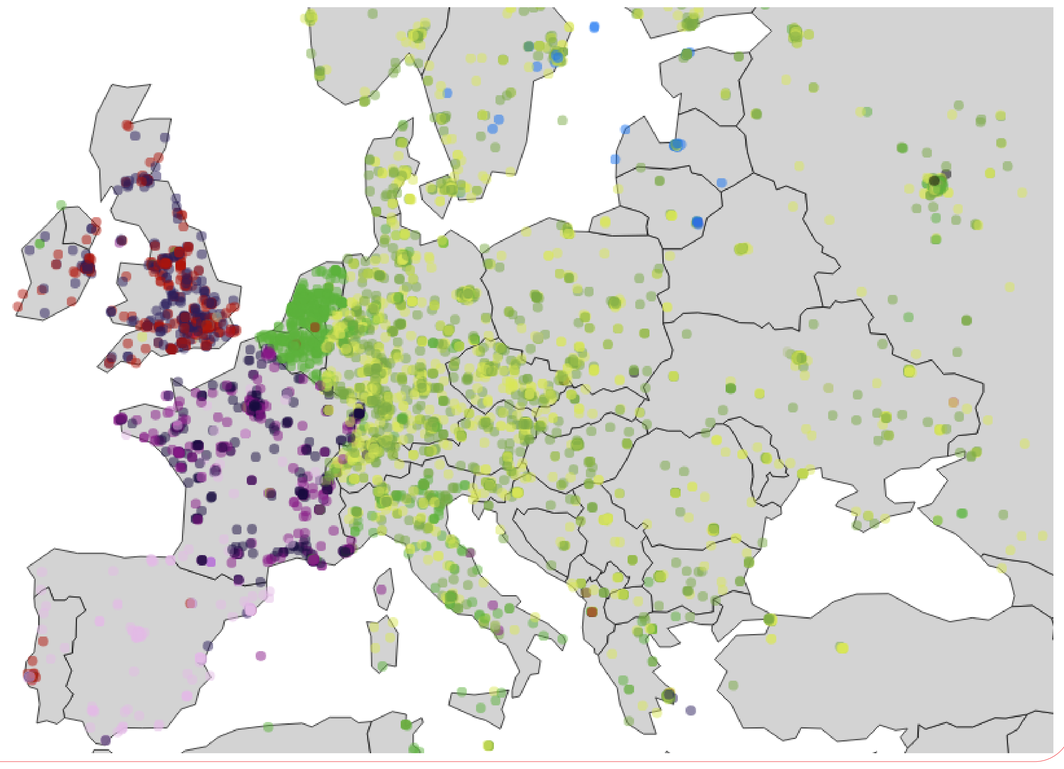

Grouping probes

- Group probes together for consistent measurement
- Ideally, groups would exhibit uniform path[3] and RTT behaviour
- Find that ASN combined with state (country outside of US) provides best trade off

For All US Networks:

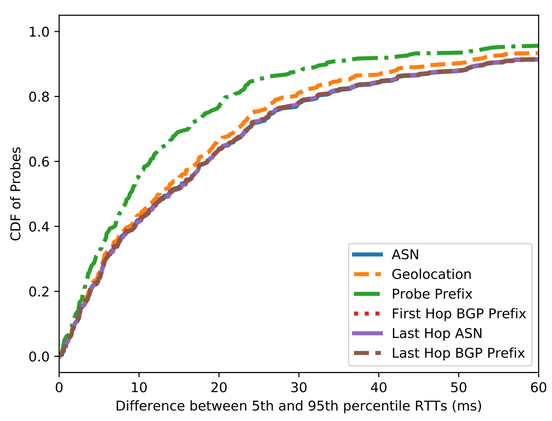

For *Large* US Networks:

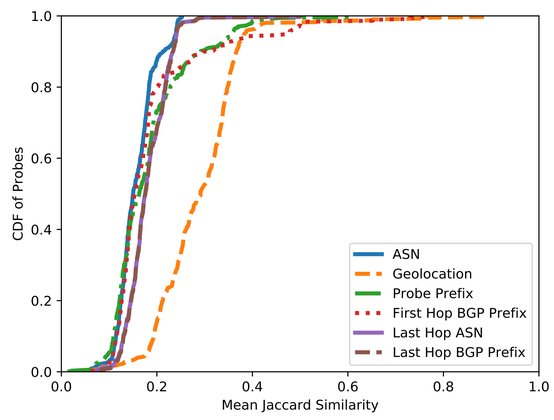

For All US Networks:

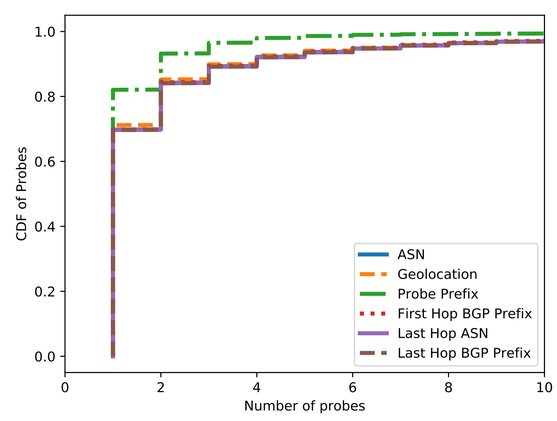

Scoring

In order to quantify the value of a configuration change, for each group g_i we consider $\Delta_{g_i} = \frac{\text{rtt}_{ctl}^{g_i} - \text{rtt}_{exp}^{g_i}}{\text{rtt}_{ctl}^{g_i}}$. To normalize these values, we apply a logistic function:

$$s(\Delta_{g_i}) = \frac{2}{1 + e^{-2\Delta_{g_i}}} - 1 \in [-1, 1].$$

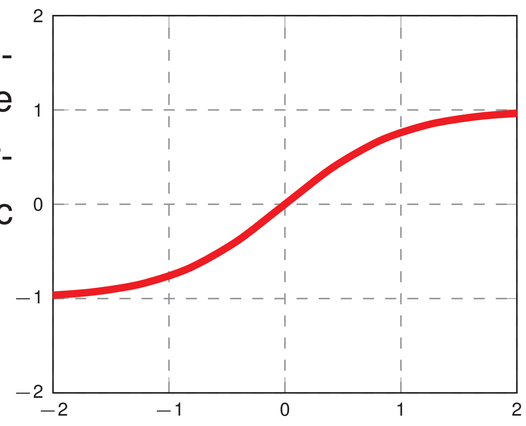

Methodology

1. Perform traceroute from each probe to control and test blocks
2. Group probes together, based on AS number and geolocation
3. Score the change within each group, based on CDN client popularity of ASN.
4. Rank groups: positive scores indicate that performance has improved; negative scores show that it has degraded

Case study

The best and worst peformers with a tested configuration. Larger, more important networks are weighted higher.

Group	# Probes	rtt_{ctl}	rtt_{exp}	Score
A	83	50.37	13.27	.039
B	13	55.62	16.92	.026
C	12	19.70	20.77	−.002
D	4	13.32	15.01	−.003

Several large networks saw improvements, reflected in higher scores and richer catchments.

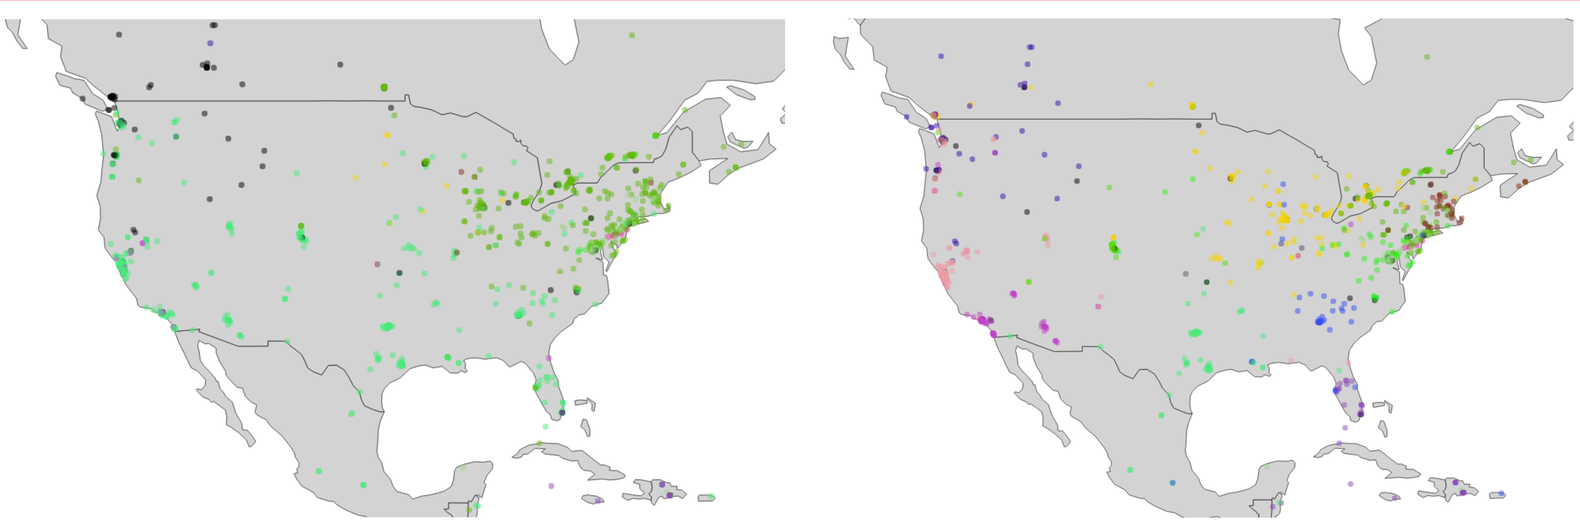

References

[1] RIPE Atlas - RIPE Network Coordination Centre, 2017. `https://atlas.ripe.net/`.

[2] W. B. de Vries, R. d. O. Schmidt, W. Hardaker, J. Heidemann, P.-T. de Boer, and A. Pras. Verfploeter: Broad and load-aware anycast mapping. Technical Report ISI-TR-719, USC/Information Sciences Institute, 2017.

[3] T. Holterbach, E. Aben, C. Pelsser, R. Bush, and L. Vanbever. Measurement Vantage Point Selection Using A Similarity Metric. In *Proceedings of the 2017 Applied Networking Research Workshop*. ACM, 2017.