

PREDICTING ANYCAST CATCHMENTS USING UNICAST TRACEROUTES

Wouter B. de Vries
w.b.devries@utwente.nl

Pieter-Tjerk de Boer
p.t.deboer@utwente.nl

Aiko Pras
a.pras@utwente.nl

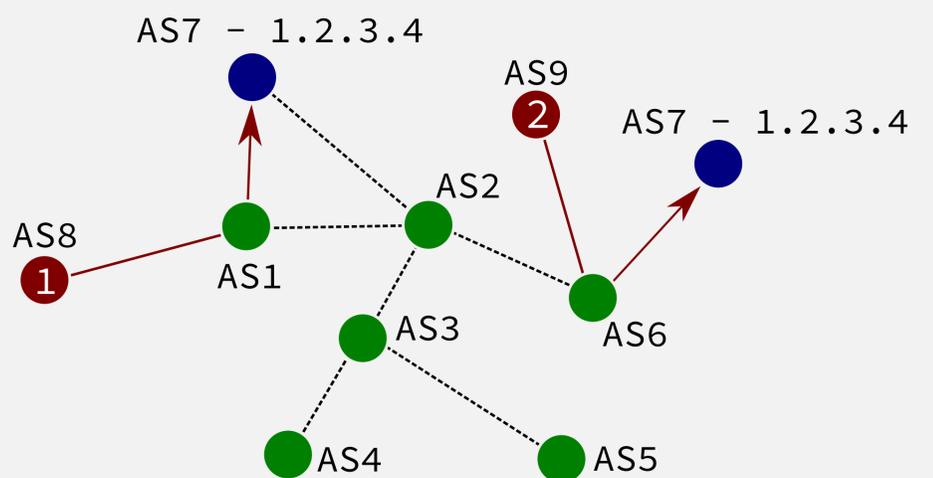
IP anycast is widely being used to distribute essential Internet services, such as DNS, across the globe. The distribution of clients across anycast sites (i.e., the catchment) depends on the routing interaction between each of the sites of the same service. The question is, how to predict if a new anycast site is going to interact well (i.e., result in a desirable catchment distribution) with an existing site?

Can anycast catchments be sufficiently accurately predicted using unicast traceroutes?

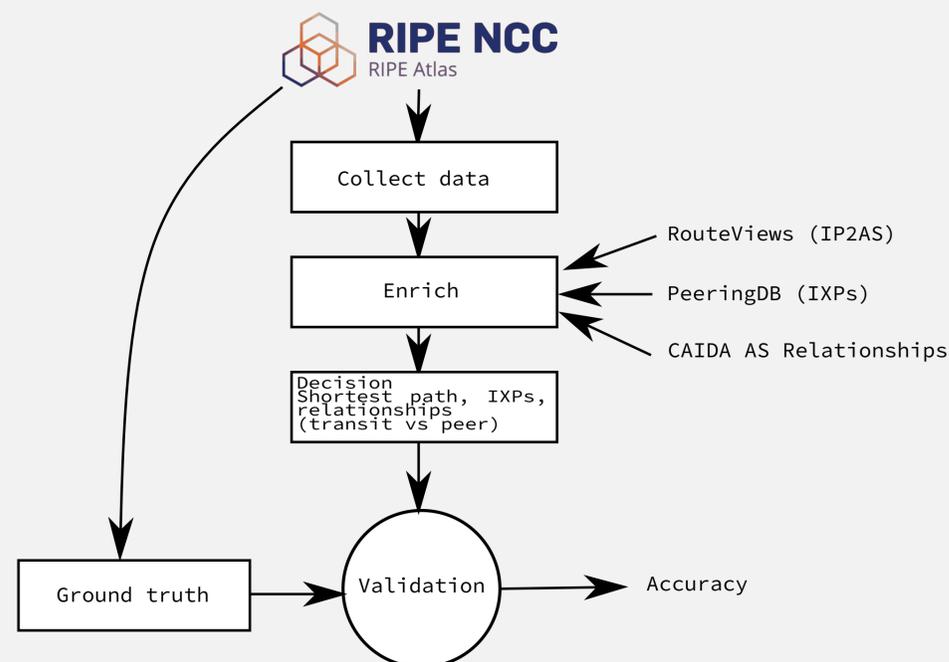
- How do routers make their decisions?
- What are the common preferences amongst operators?
- What is the best way to use this knowledge to predict routing decisions? (Machine learning?)

Goal: to provide operators with a tool to predict the anycast catchment given two network locations, prior to actual deployment

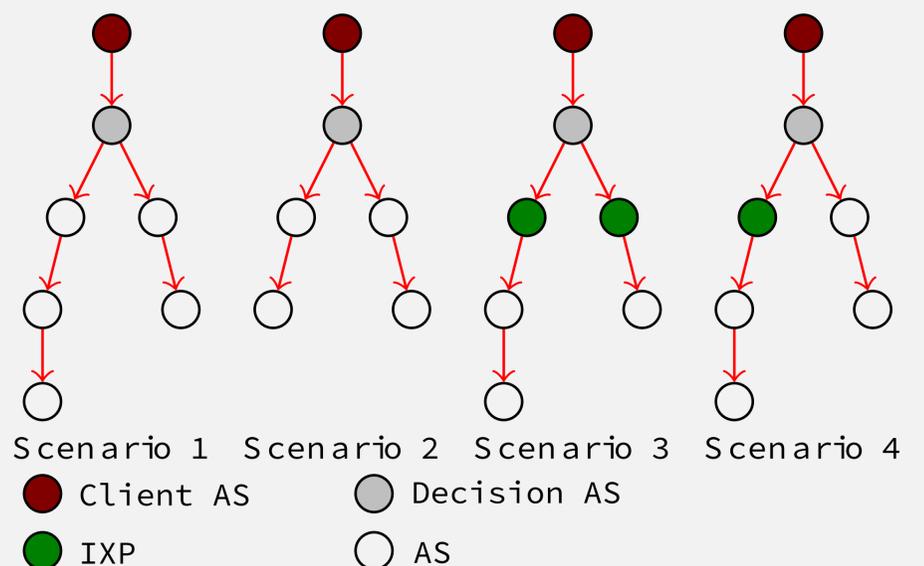
"Anycast is a network addressing and routing methodology in which datagrams from a single sender are routed to the topologically nearest node in a group of potential receivers"



Methodology



Example scenarios



A basic decision process, based on AS path length, AS Relationships and IXPs achieves reasonable results. A case study on B-root (2 anycast sites) shows that from a total of 6632 measurements, in 3322 cases the right site could be predicted, in 1018 cases the wrong site, and in a further 2283 cases no prediction could be made.