Internet Topology, Geography and other Random things

Emile Aben
System Architect
RIPE NCC

We have lots of interesting data for viz

- Examples:
  - RIRstats: Resource allocations (with countrycode)
    - ftp://ftp.ripe.net/pub/stats/
  - IPv6 AS stats: % IPv6 enabled ASes per country
    - http://v6asns.ripe.net/
  - IPv6 RIPEness: IPv6 readiness per LIR/country
    - http://ipv6ripeness.ripe.net/

- Interesting to combine with other per country data?

We have lots of interesting data for viz

- Examples: RIRstats: Address use growth in 2011

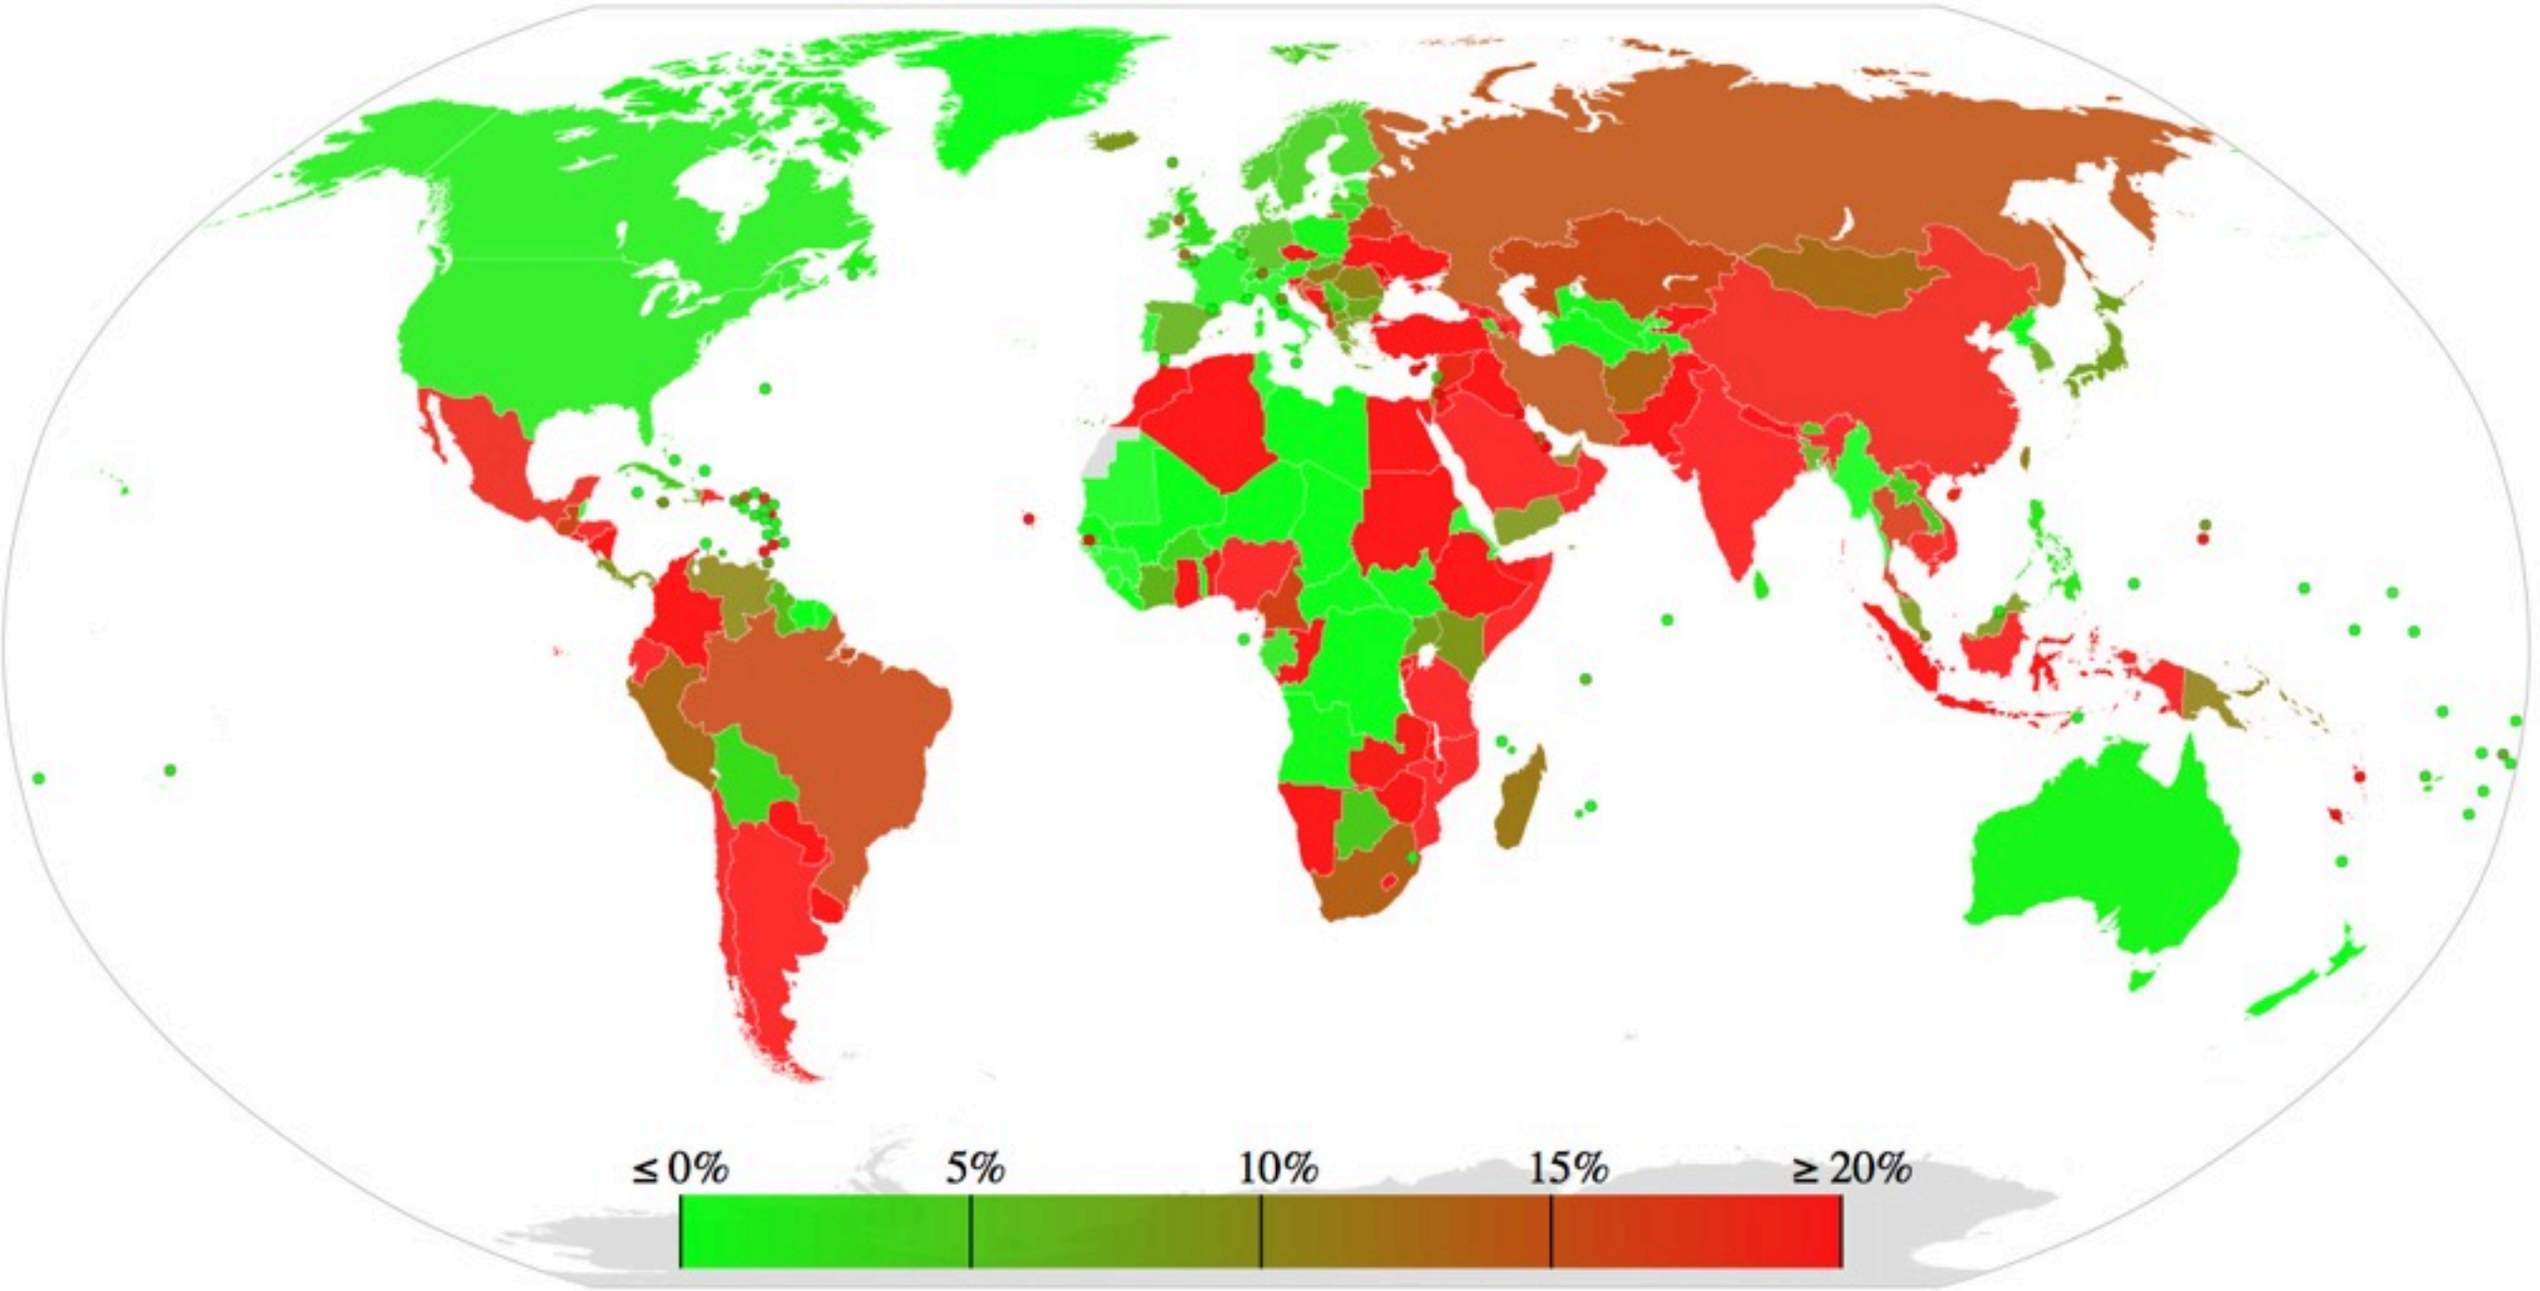

Where is d-root?

RIPE Atlas RTTs Map produced at: 2012-01-23 10:32:46 UTC.
Measurements taken between 2012-01-23 10:02 UTC and 2012-01-23 10:32 UTC.

Round Trip Time to d.root-servers.net

http://atlas.ripe.net/atlas/maps_index.html

Where is f-root?

RIPE Atlas RTTs Map produced at: 2012-01-23 10:37:02 UTC.
Measurements taken between 2012-01-23 10:07 UTC and 2012-01-23 10:37 UTC.

Round Trip Time to f.root-servers.net

http://atlas.ripe.net/atlas/maps_index.html

Massive Internet measurements and maps

- Problem: Too much data to plot and visually make sense of
- What is the Internet?
  - Collection of networks (ASes) that interconnect
  - 1 AS ≈ 1 ISP
  - Current Internet > 30 000 ASes
- Obvious solution: Aggregate info per AS

CAIDA's IPv4 & IPv6 AS Core
AS-level INTERNET GRAPH

Archipelago August 2010

IPv4

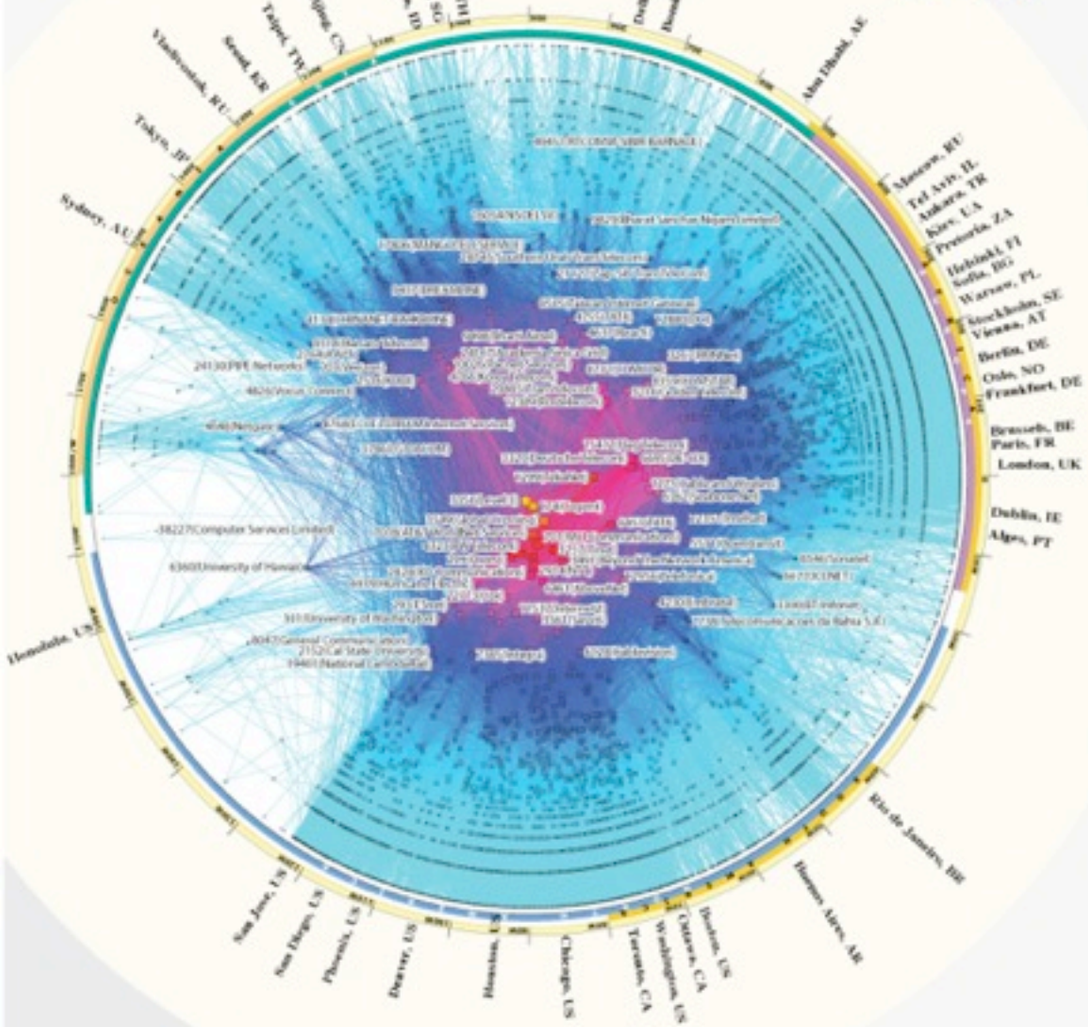

IPv6

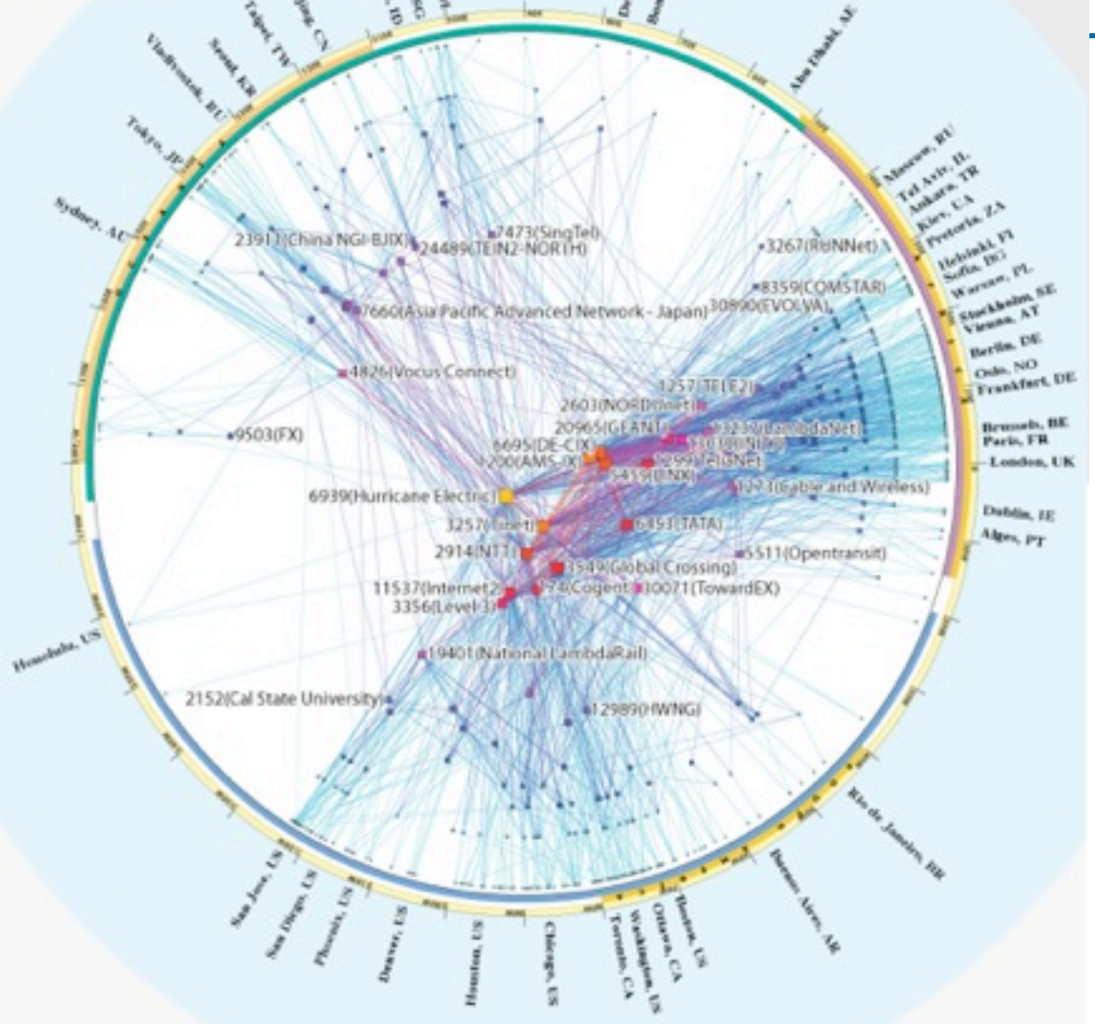

ANALYSIS TEAM Bradley Huffaker, bc claffy
SOFTWARE DEVELOPMENT Young Hyun, Matthew Luckie
POSTER DESIGN Colette Ijya, Will Michaelsen

	Number of IP addresses	Number of IP links	Number of ASes	Number of AS links
IPv4	16,802,061	18,796,744	26,7021	85,104
IPv6	8,551	21,852	715	1,672

ARK HOSTS AARNet, Acreo, AMS-IX, APAN, ARIN, ASTI, CAIDA, Canarie, CENIC, CNRST, CYMRU, Exolxa Telecom, FORTH, FunkFeuer, HEANet, Hurricane Electric, Indonesian IPv6 Task Force, Internet Systems Consortium, Iowa State Univ., KREONet2, Level 3 Communications, Men and Mice, National Research Council Canada, NCAR, NIC Chile, NIC Mexico, Northeastern Univ., Public Univ. of Navarra, Purdue Univ., RNP, Southern Methodist Univ., SURFnet, TKK, TWAREN, UCAD, Univ. Leipzig, Univ. Politecnica de Catalunya, Univ. of Cambridge, Univ. of Hawaii, Univ. of Napoli, Univ. of Nevada at Reno, Univ. of Oregon, Univ. of Waikato, Univ. of Washington, Univ. of Zurich, US Army Research Lab

COOPERATIVE ASSOCIATION FOR INTERNET DATA ANALYSIS
San Diego Supercomputer Center . University of California, San Diego
9500 Gilman Drive, mc0505 . La Jolla, CA 92093-0505 . 858-534-5000
http://www.caida.org/research/topology/as_core_network/

This visualization represents macroscopic snapshots of IPv4 and IPv6 Internet topology samples captured in 2010. The plotting method illustrates both the extensive geographical scope as well as rich interconnectivity of nodes participating in the global Internet routing sytem.

For the IPv4 map, CAIDA collected data from 45 monitors located in 24 countries on 6 continents. Coordinated by our active measurement infrastructure, Archipelago (Ark[1]), the monitors probed paths toward 124 million /24 networks that cover 96% of the routable prefixes seen in the Route Views[2] Border Gateway Protocol (BGP) routing tables on 1 August 2010.

For the IPv6 map, CAIDA collected data from 12 Ark monitors located in 6 countries on 3 continents. This subset of monitors probed paths toward 307 thousand IPv6 prefixes which represent 99.6% of the globally routed IPv6 prefixes seen in Route View's BGP tables on 1 August 2010. We aggregate this IP-level data to construct IPv4 and IPv6 Internet connectivity graphs at the Autonomous System (AS) level. Each AS approximately corresponds to an Internet Service Provider (ISP). We map each observed IP address to the AS responsible for routing traffic to it, i.e., to the origin (end-of-path) AS for the IP prefix representing the best match for this address in BGP routing tables collected from Route Views.

The position of each AS node is plotted in polar coordinates (radius, angle) calculated as follows:

$$\text{radius} = 1 - \log\left(\frac{\text{outdegree(AS)} + 1}{\text{maximum outdegree} + 1}\right)$$

$$\text{angle} = \left(\frac{\text{longitude of the AS's BGP prefixes in netacq}}{}\right)$$

Our IPv6 graph grew from 515 AS nodes in January 2009 to 948 nodes in August 2010 (84% growth). Over the same period, the number of ASes in our IPv4 graph grew 22%, from 23K to 28K.

Most ASes grew their observed peering degree in both our IPv4 and IPv6 graphs, although at different rates, which alters their relative degree-based rank over time. In the IPv4 graph, AS 3356 remained dominant, with the largest observed degree in both 2009 and 2010. The second and third largest ASes, AS 174 and AS 3549, also maintained the same observed degree relative to the largest degree AS (3356) for the last two years. In contrast, ASes 7018, 701, and 1239 saw observed peering degree declines relative to the largest degree AS 3356, slipping to 4th, 5th, and 7th place. Note that we rank each AS independently; some network providers have topology spread across multiple ASes. A more accurate topology-based ranking of providers would require a validated list of AS ownership -- data not currently available.

The observed IPv6 AS ranking experienced greater change. AS 6939 moved up from 2nd place in 2009 to 1st place in 2010. AS 1200 dropped from 1st to 3rd place. AS 12008 and 6175 fell out of the top ten, allowing AS 1299 and 174 to rise to 9th and 10th place. The third and fourth ranked ASes -- AS 1200 and AS6695 -- are both exchange points rather than transit providers, reflecting the less mature state of the IPv6 topology, i.e., characterized by relatively fewer private peering relationships.

In neither 2009 nor 2010 are the top degree-ranked ASes the same across IPv4 and IPv6. The IPv4 core is centered primarily in the United States, while the IPv6 core includes Europe as well as the United States. We observed no high-degree "hub" IPv6 ASes in Asia, surprising given the reportedly large IPv6 deployment in Asia. This gap may reflect the geographic bias of our IPv6-capable monitor deployment: five in the US, four in Europe, and only one in Asia.

Top Ranked ASes

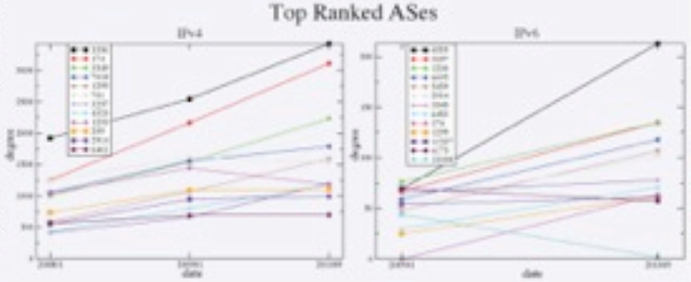

[1] Ark: http://www.caida.org/projects/ark/
[2] Route Views: http://www.routeviews.org/

ASes have very different sizes

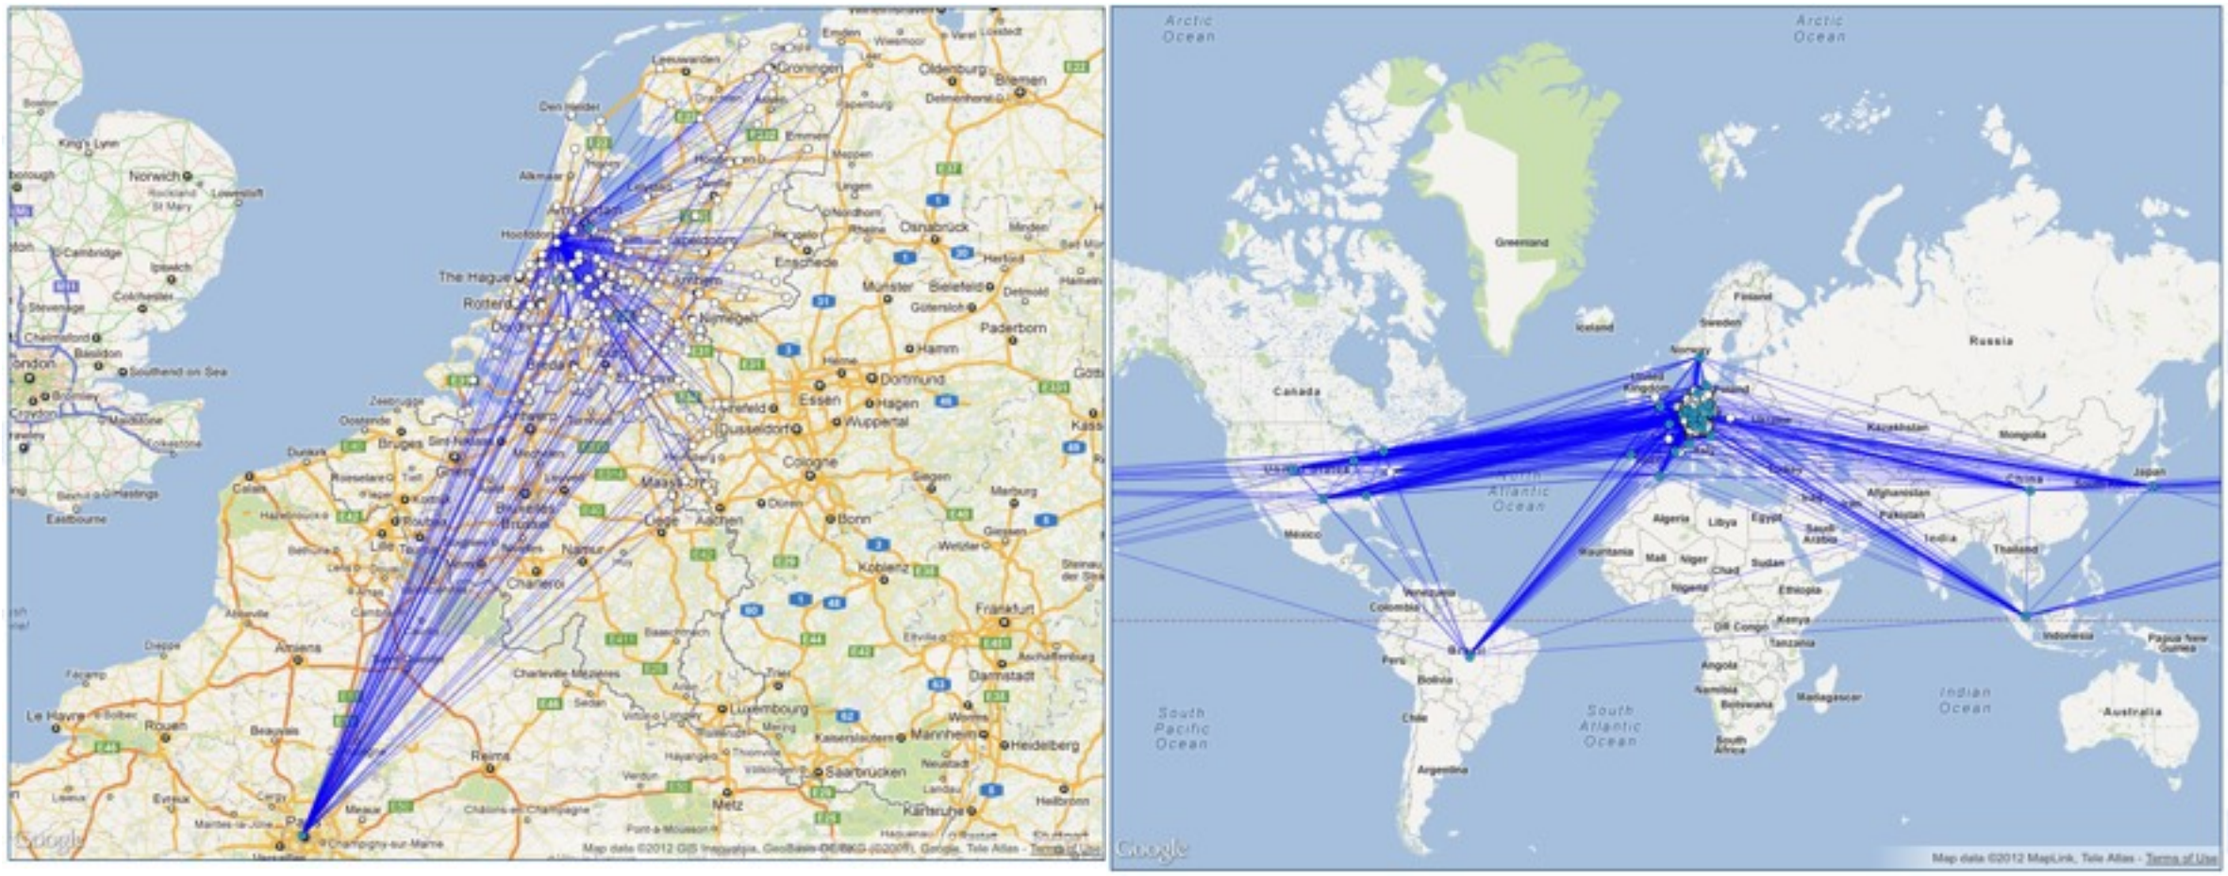

Source: http://as-rank.caida.org/

Problem: IP geolocation DBs inaccurate, especially for infrastructure

ASes overlap

Source: http://as-rank.caida.org/

A bit more Difficult? Traceroute

```
traceroute to www.ucla.edu (169.232.55.224), 64 hops max, 52 byte packets
 1  gw.telrtr.guestnet.ripe.net (193.0.10.3)  1.265 ms  0.624 ms  0.991 ms
 2  amsix1.eun.ams.as8218.eu (195.69.145.47)  1.087 ms  0.859 ms  0.825 ms
 3  xe0-1-0.tcr1.thn.lon.as8218.eu (83.167.63.235)  6.761 ms  6.801 ms  6.818 ms
 4  xe0-0-0.tcr1.swad.nyc.as8218.eu (83.167.56.145)  82.622 ms  82.842 ms  82.673 ms
 5  83.167.55.86.static.not.updated.as8218.eu (83.167.55.86)  83.003 ms  83.055 ms  82.620 ms
 6  rtr.loss.net.internet2.edu (198.32.118.161)  82.907 ms  83.007 ms  82.906 ms
 7  xe-2-1-0.0.chic0.tr-cps.internet2.edu (64.57.20.251)  107.871 ms  108.142 ms  107.777 ms
 8  137.164.129.2 (137.164.129.2)  171.237 ms  171.094 ms  171.048 ms
 9  xe-1-0-0.0.paix0.tr-cps.internet2.edu (64.57.20.222)  171.141 ms  171.012 ms  171.051 ms
10  137.164.131.94 (137.164.131.94)  172.468 ms  172.525 ms  172.156 ms
11  dc-svl-core1--svl-px1-10ge-1.cenic.net (137.164.46.204)  172.951 ms  172.313 ms  172.870
ms
12  dc-lax-core2--svl-core1-10ge-1.cenic.net (137.164.46.96)  173.324 ms  174.151 ms  173.039
ms
13  dc-lax-agg1--lax-core2-ge.cenic.net (137.164.46.106)  172.875 ms  173.237 ms  172.271 ms
14  * * *
15  border-2--core-1-ge.backbone.ucla.net (169.232.4.104)  173.546 ms  174.226 ms  173.386 ms
16  core-1--anderson-1-ge.backbone.ucla.net (169.232.8.27)  173.849 ms  173.194 ms  173.171 ms
17  * * *
18  * * *
19  * * *
```

Multi-origin traceroute

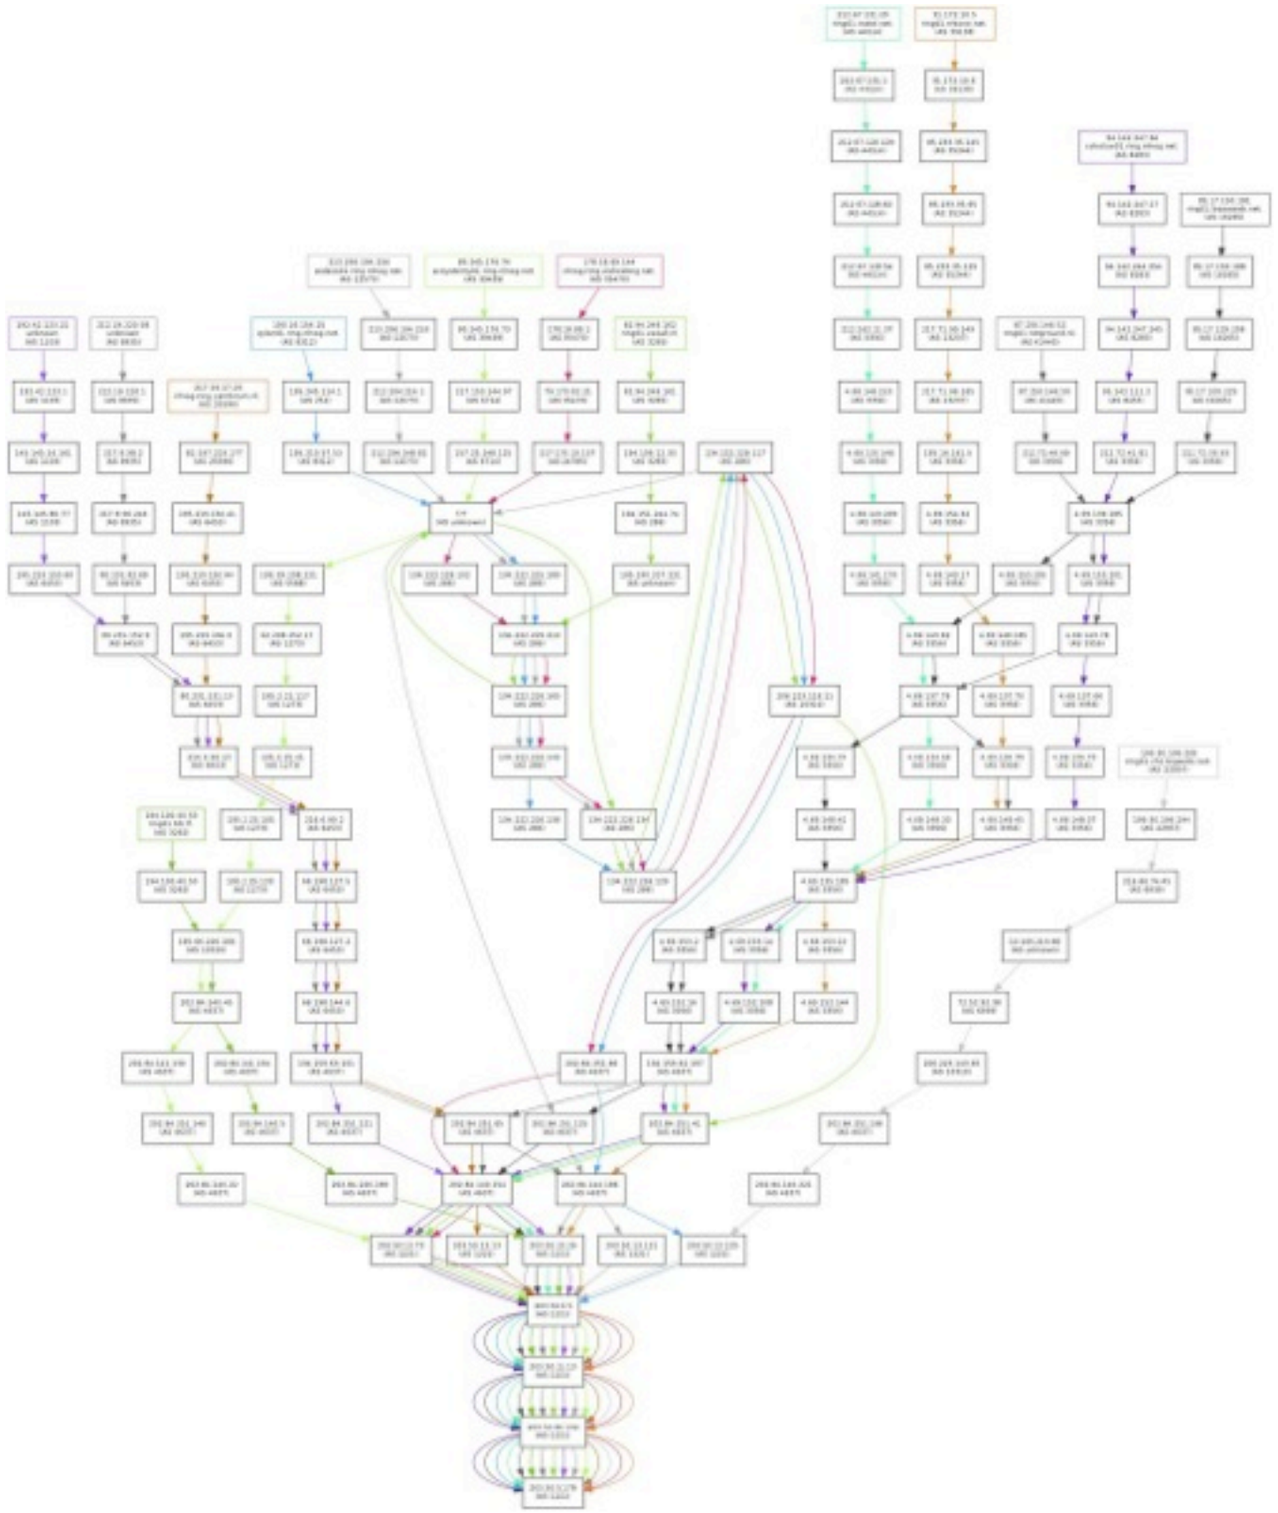

- Source: Teun Vink
- (NLNOG The Ring)

Multi-origin traceroute

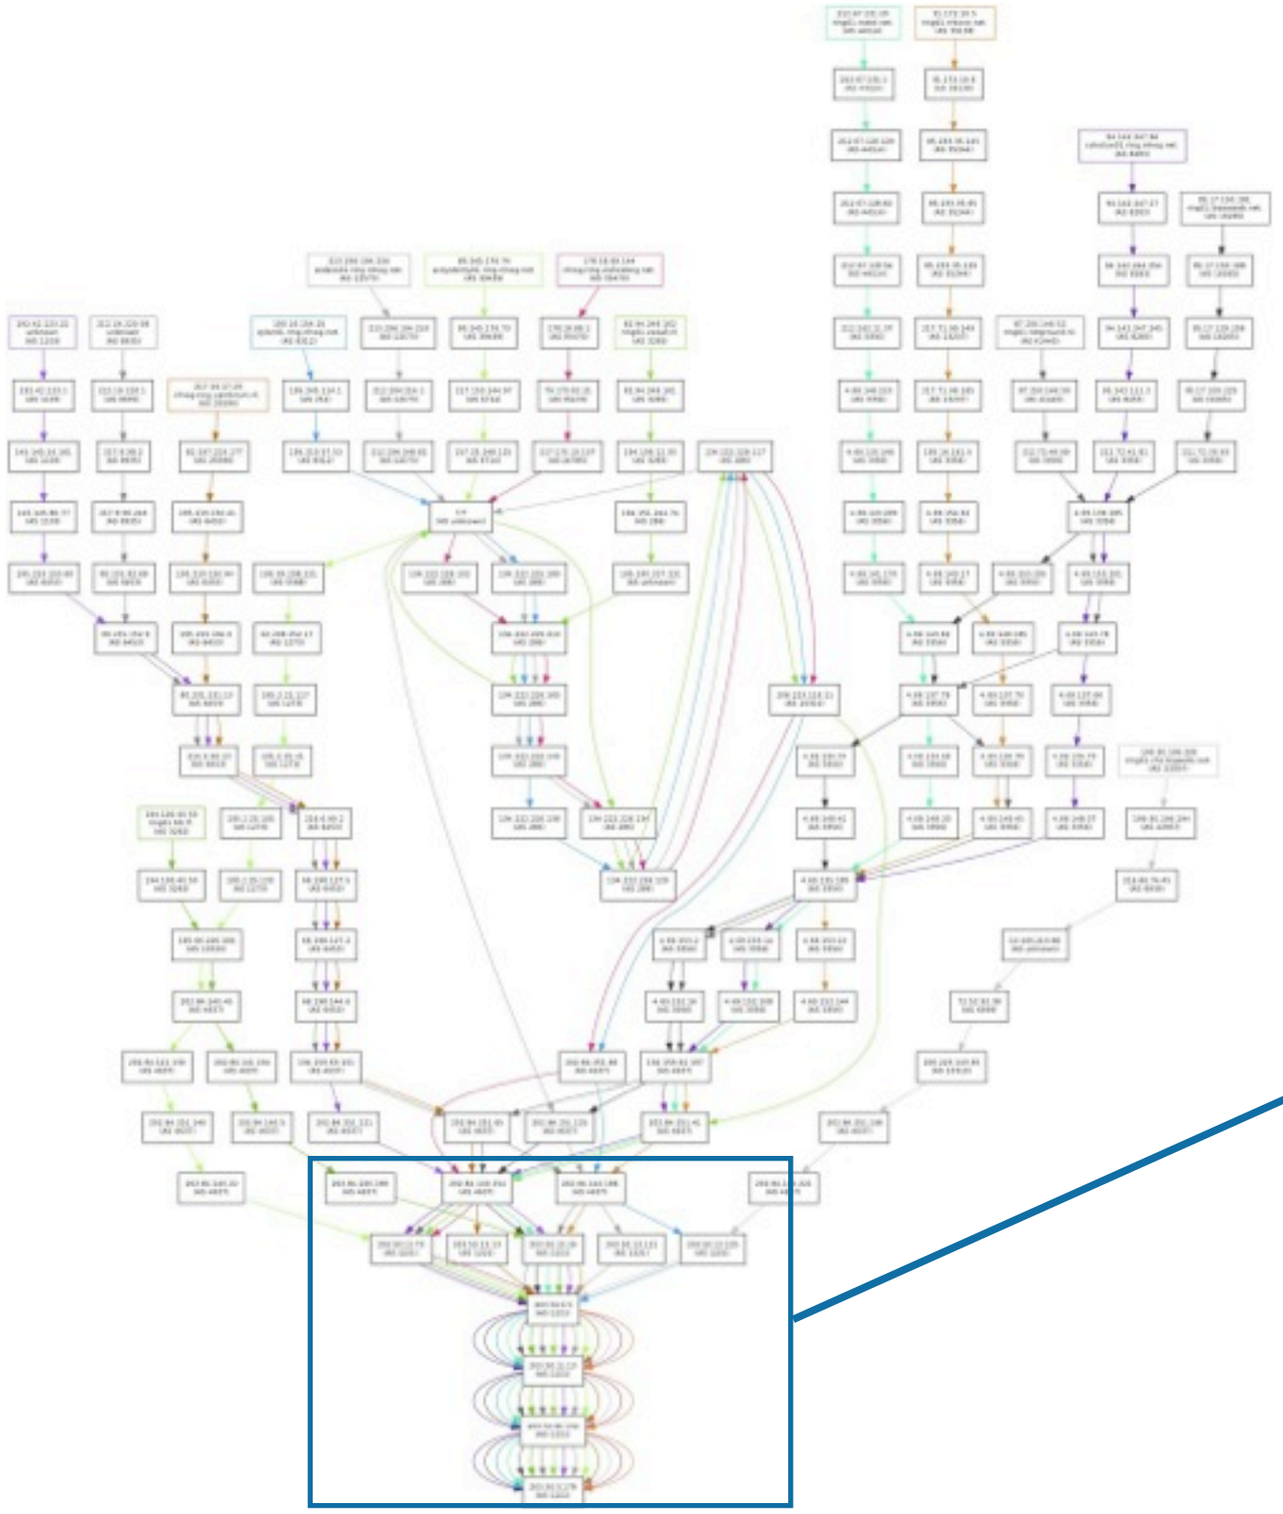

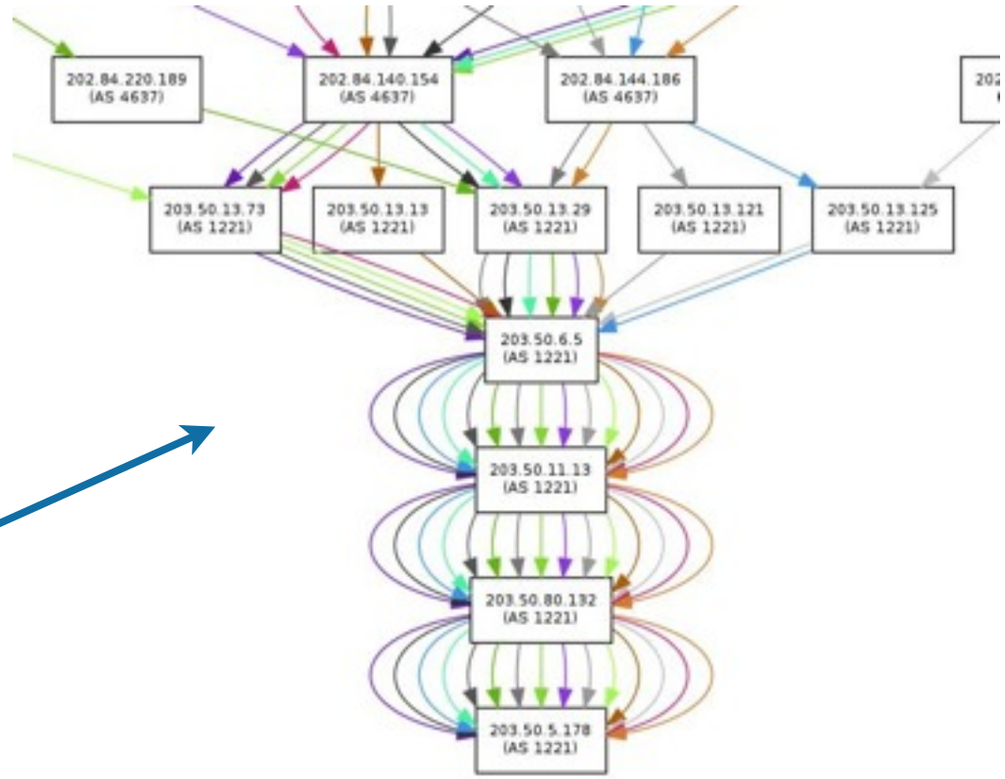

- Source: Teun Vink
- (NLNOG The Ring)

Multiple-origin traceroute “subway map”

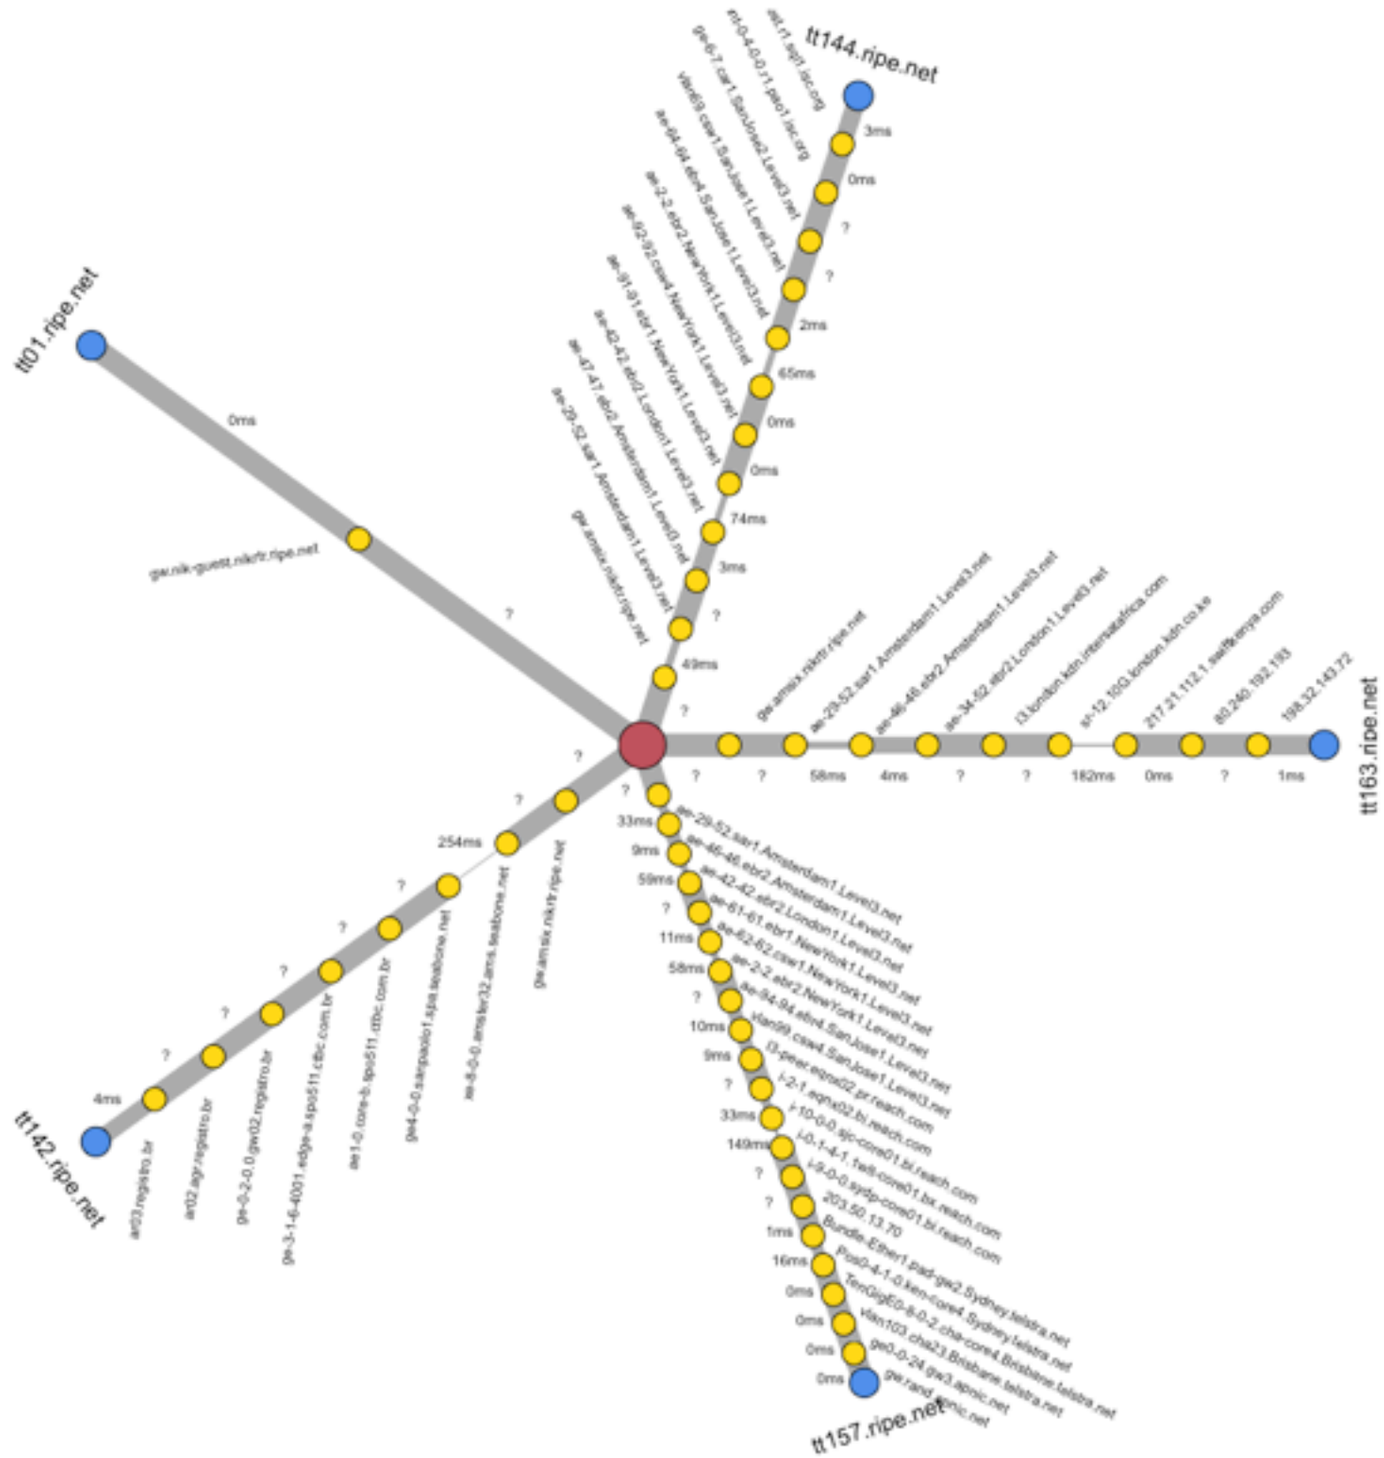

Traceroute per AS

```
traceroute to www.ucla.edu (169.232.55.224), 64 hops max, 52 byte packets
 1  gw.telrtr.guestnet.ripe.net (193.0.10.3)  1.265 ms  0.624 ms  0.991 ms

 2  amsix1.eun.ams.as8218.eu (195.69.145.47)  1.087 ms  0.859 ms  0.825 ms
 3  xe0-1-0.tcr1.thn.lon.as8218.eu (83.167.63.235)  6.761 ms  6.801 ms  6.818 ms
 4  xe0-0-0.tcr1.swad.nyc.as8218.eu (83.167.56.145)  82.622 ms  82.842 ms  82.673 ms
 5  83.167.55.86.static.not.updated.as8218.eu (83.167.55.86)  83.003 ms  83.055 ms  82.620 ms

 6  rtr.loss.net.internet2.edu (198.32.118.161)  82.907 ms  83.007 ms  82.906 ms
 7  xe-2-1-0.0.chic0.tr-cps.internet2.edu (64.57.20.251)  107.871 ms  108.142 ms  107.777 ms
 8  137.164.129.2 (137.164.129.2)  171.237 ms  171.094 ms  171.048 ms
 9  xe-1-0-0.0.paix0.tr-cps.internet2.edu (64.57.20.222)  171.141 ms  171.012 ms  171.051 ms
10  137.164.131.94 (137.164.131.94)  172.468 ms  172.525 ms  172.156 ms

11  dc-svl-core1--svl-px1-10ge-1.cenic.net (137.164.46.204)  172.951 ms  172.313 ms  172.870
ms
12  dc-lax-core2--svl-core1-10ge-1.cenic.net (137.164.46.96)  173.324 ms  174.151 ms  173.039
ms
13  dc-lax-agg1--lax-core2-ge.cenic.net (137.164.46.106)  172.875 ms  173.237 ms  172.271 ms
14  * * *

15  border-2--core-1-ge.backbone.ucla.net (169.232.4.104)  173.546 ms  174.226 ms  173.386 ms
16  core-1--anderson-1-ge.backbone.ucla.net (169.232.8.27)  173.849 ms  173.194 ms  173.171 ms
17  * * *
18  * * *
19  * * *
```

Split up traceroute in components

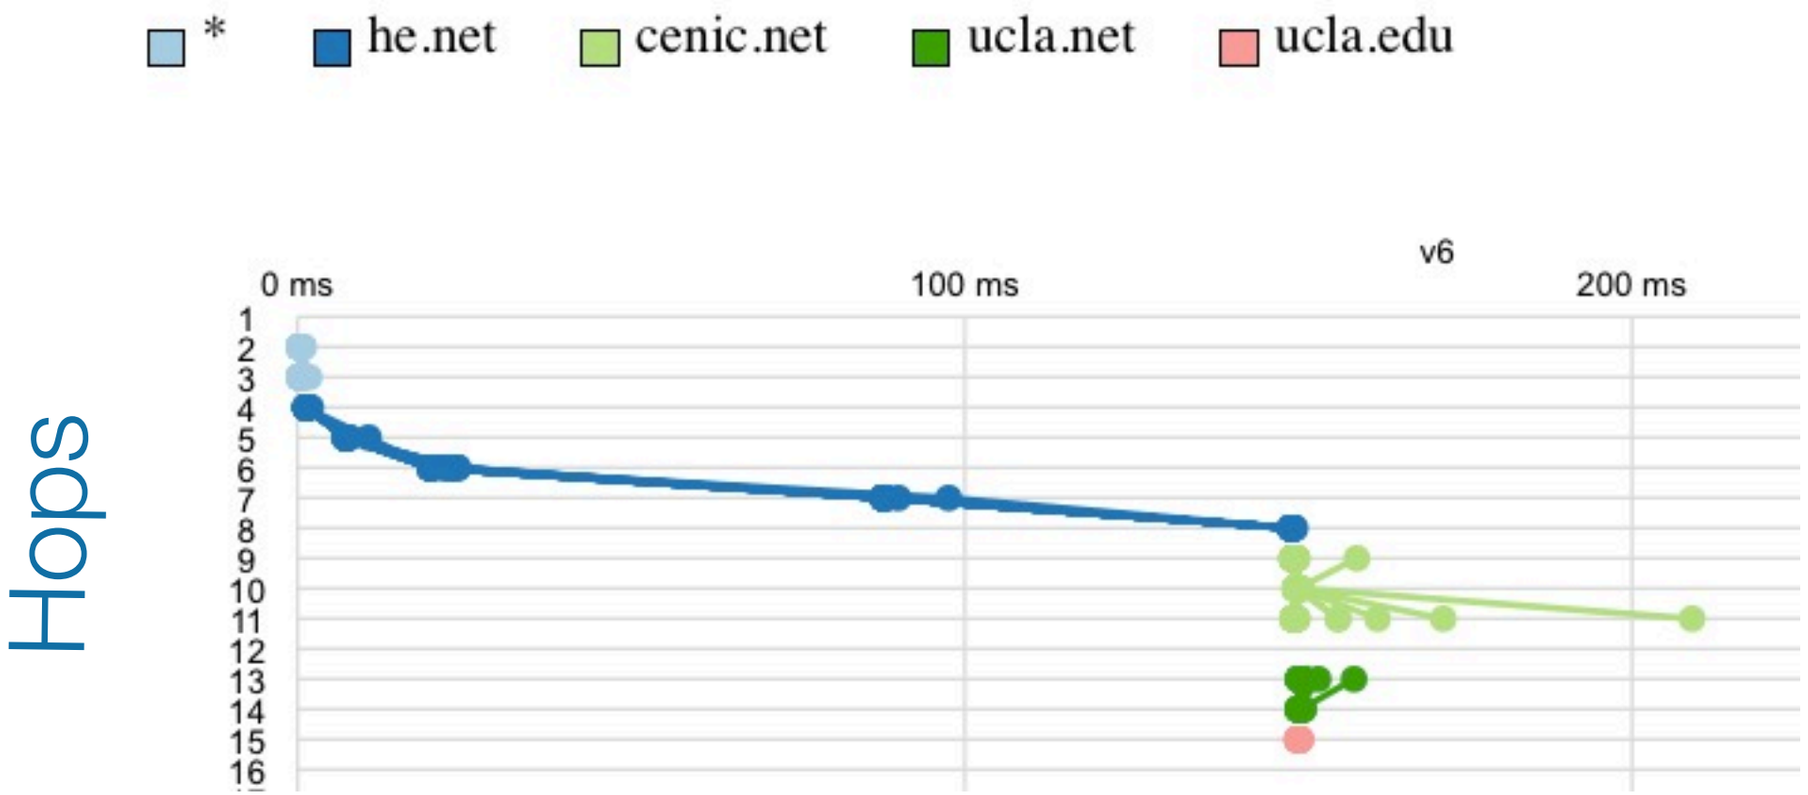

Round Trip Time

Traceroute from Switzerland to US West Coast

Split up traceroute in components

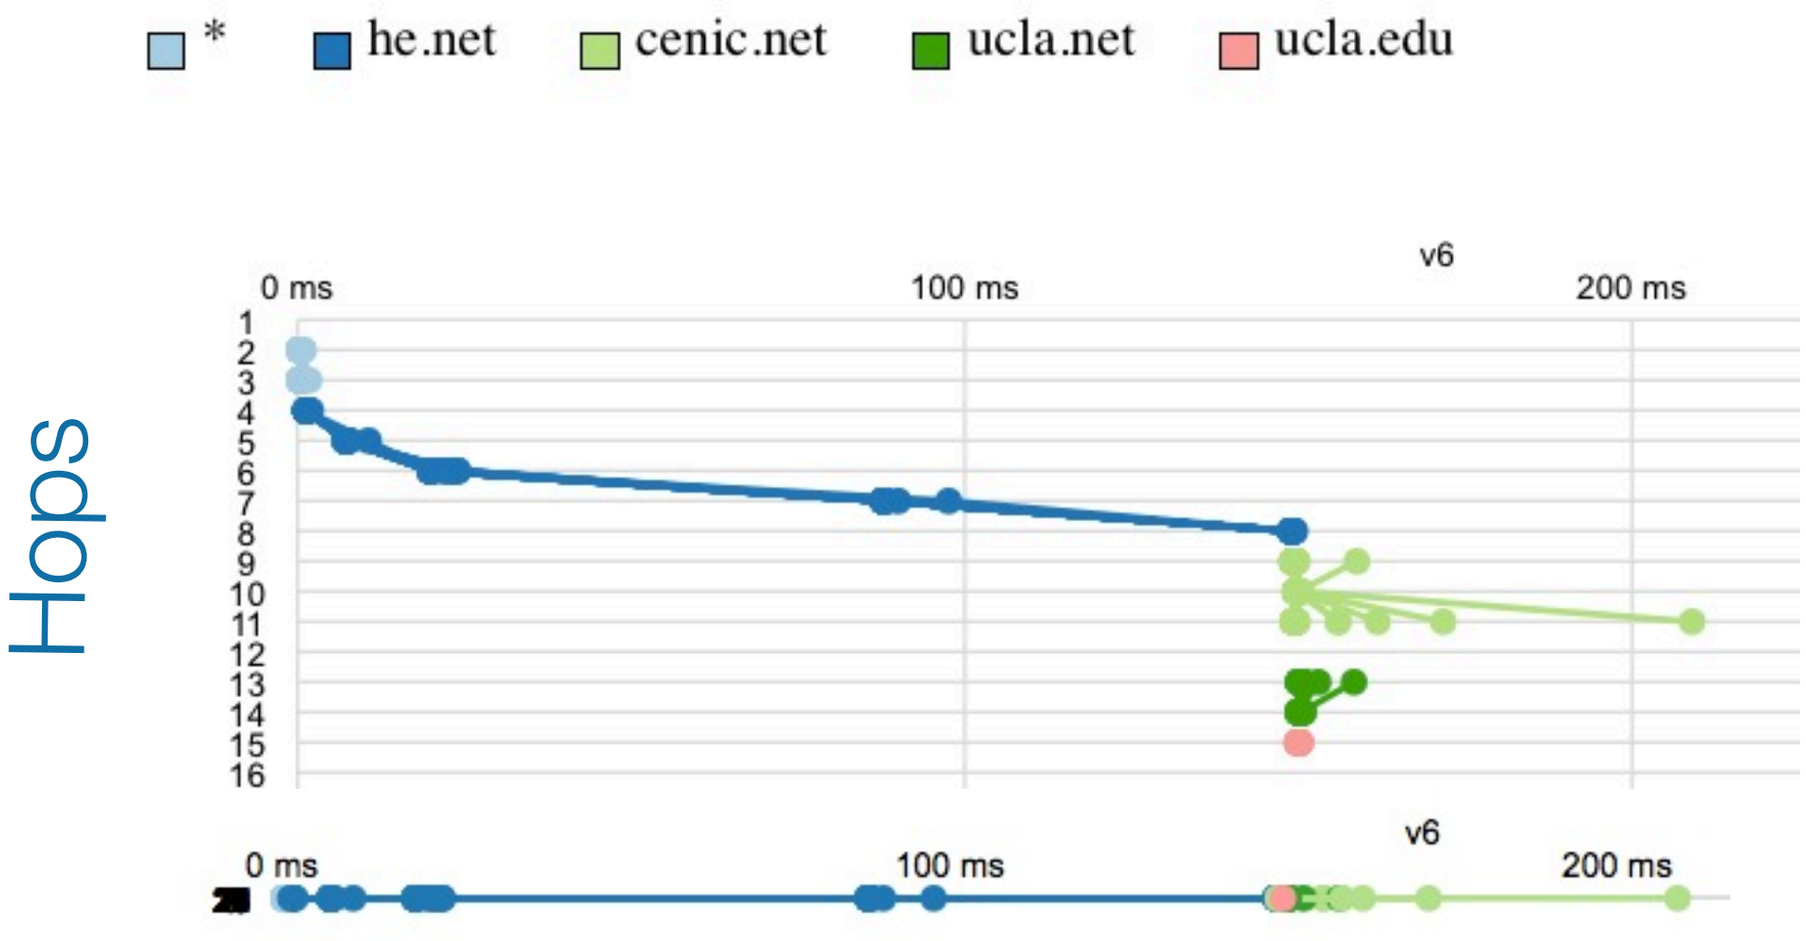

Traceroute from Switzerland to US West Coast

Next step : combine multi-origin + AS

- Only take AS entry and exit points

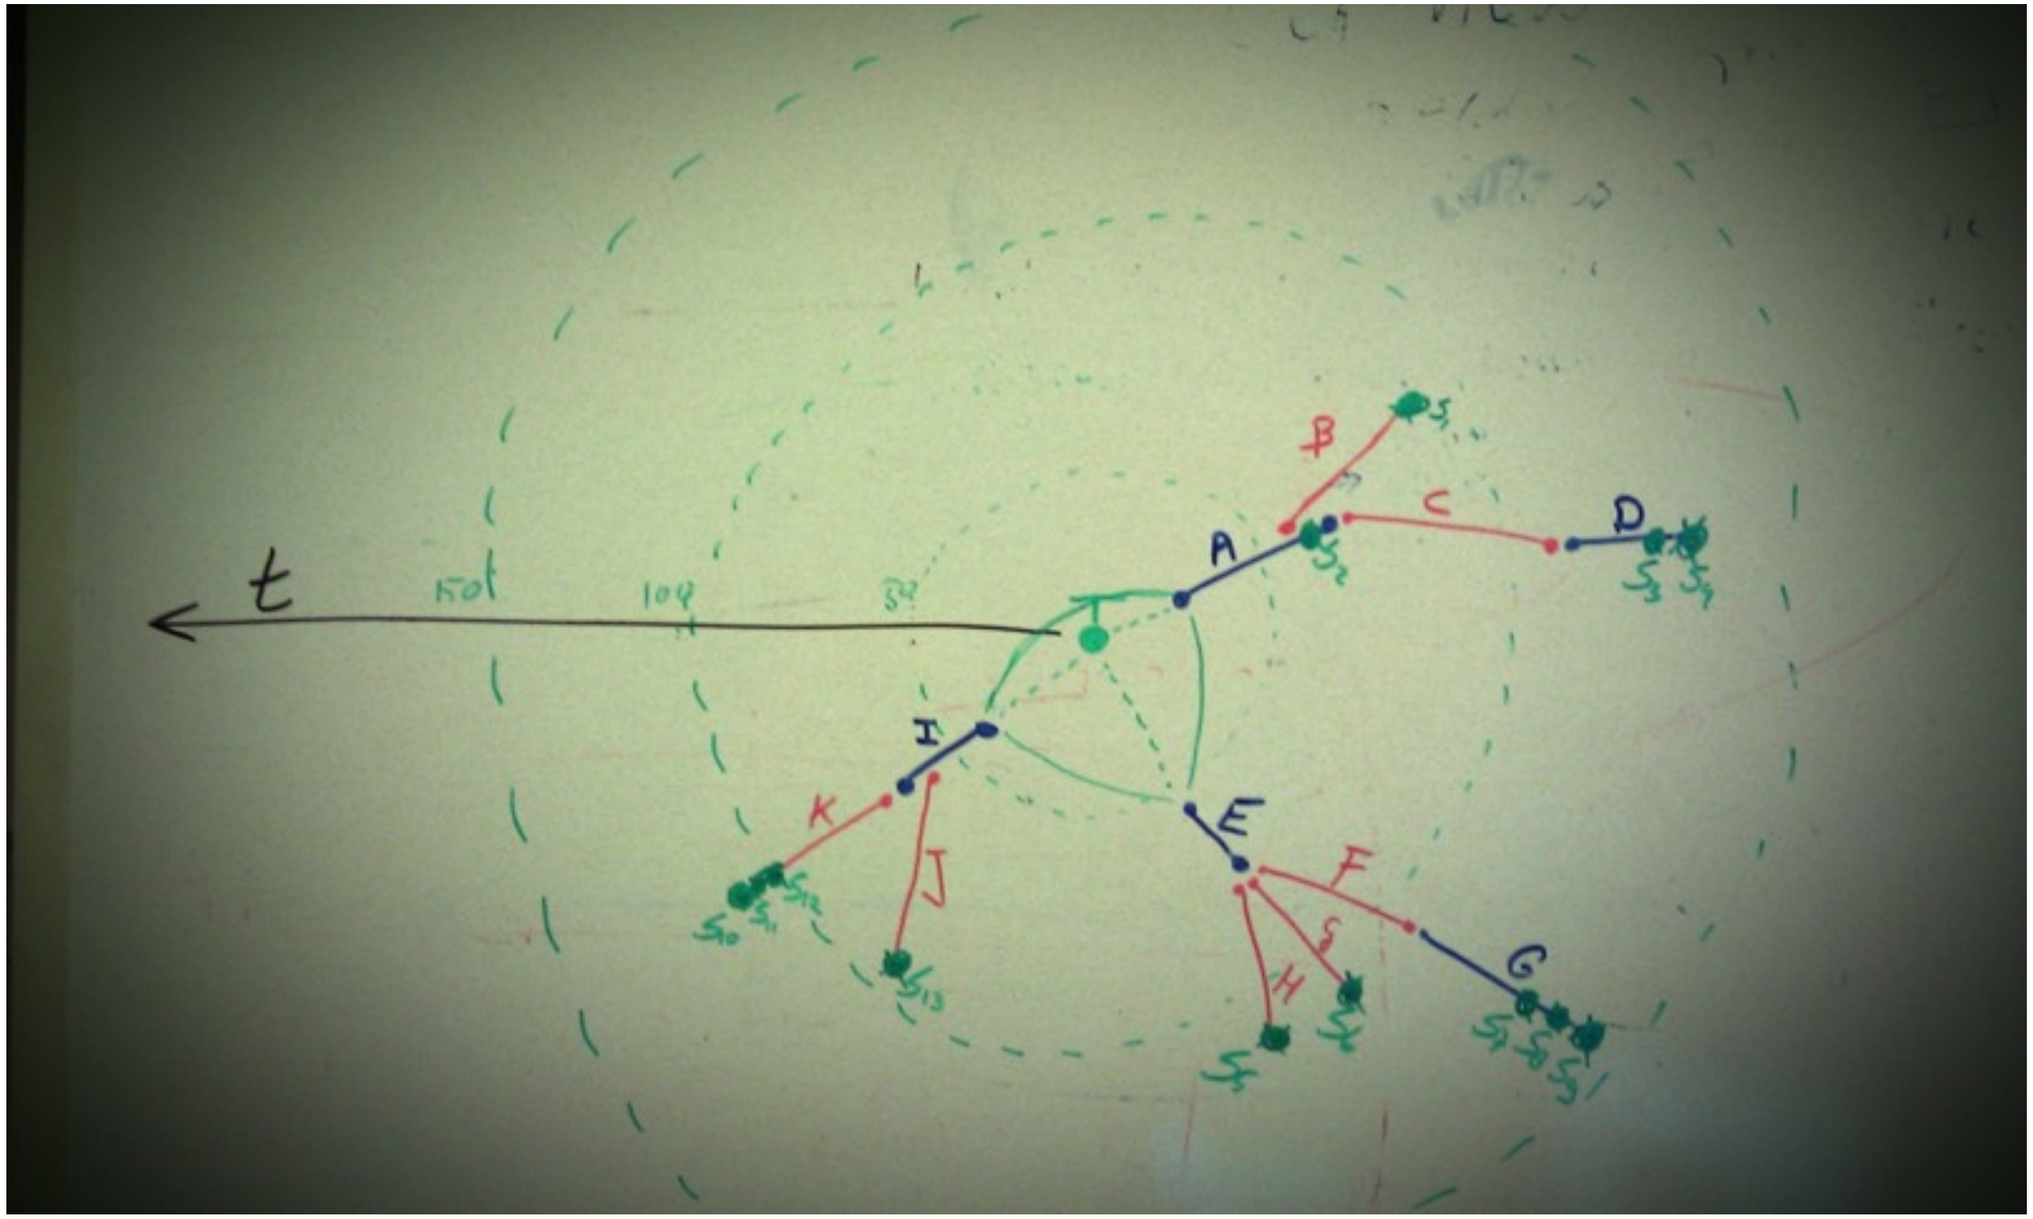

Questions?