MEASURING AND VISUALIZING DNS WATERSHEDS

Jim Cowie

Internet History Initiative

29 October 2024

INTERNET HISTORY INITIATIVE

- Collect and preserve the network operators' community legacy of Internet measurement datasets
- Extract time series data that reflect key aspects of regional Internet growth and diversification
- Study similarities and differences in Internet development across world regions
- Make these time series available to researchers studying different (potentially non-technical) aspects of international development

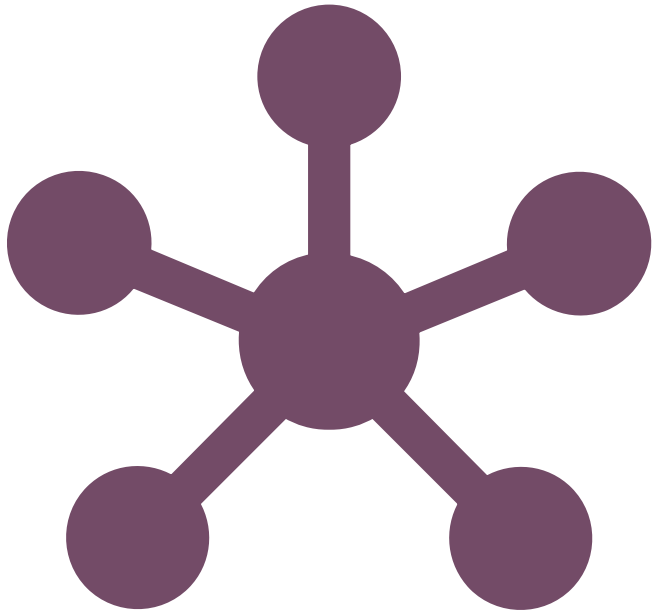

ATLAS PROBE DENSITY, OCT 2024 (LOG10)

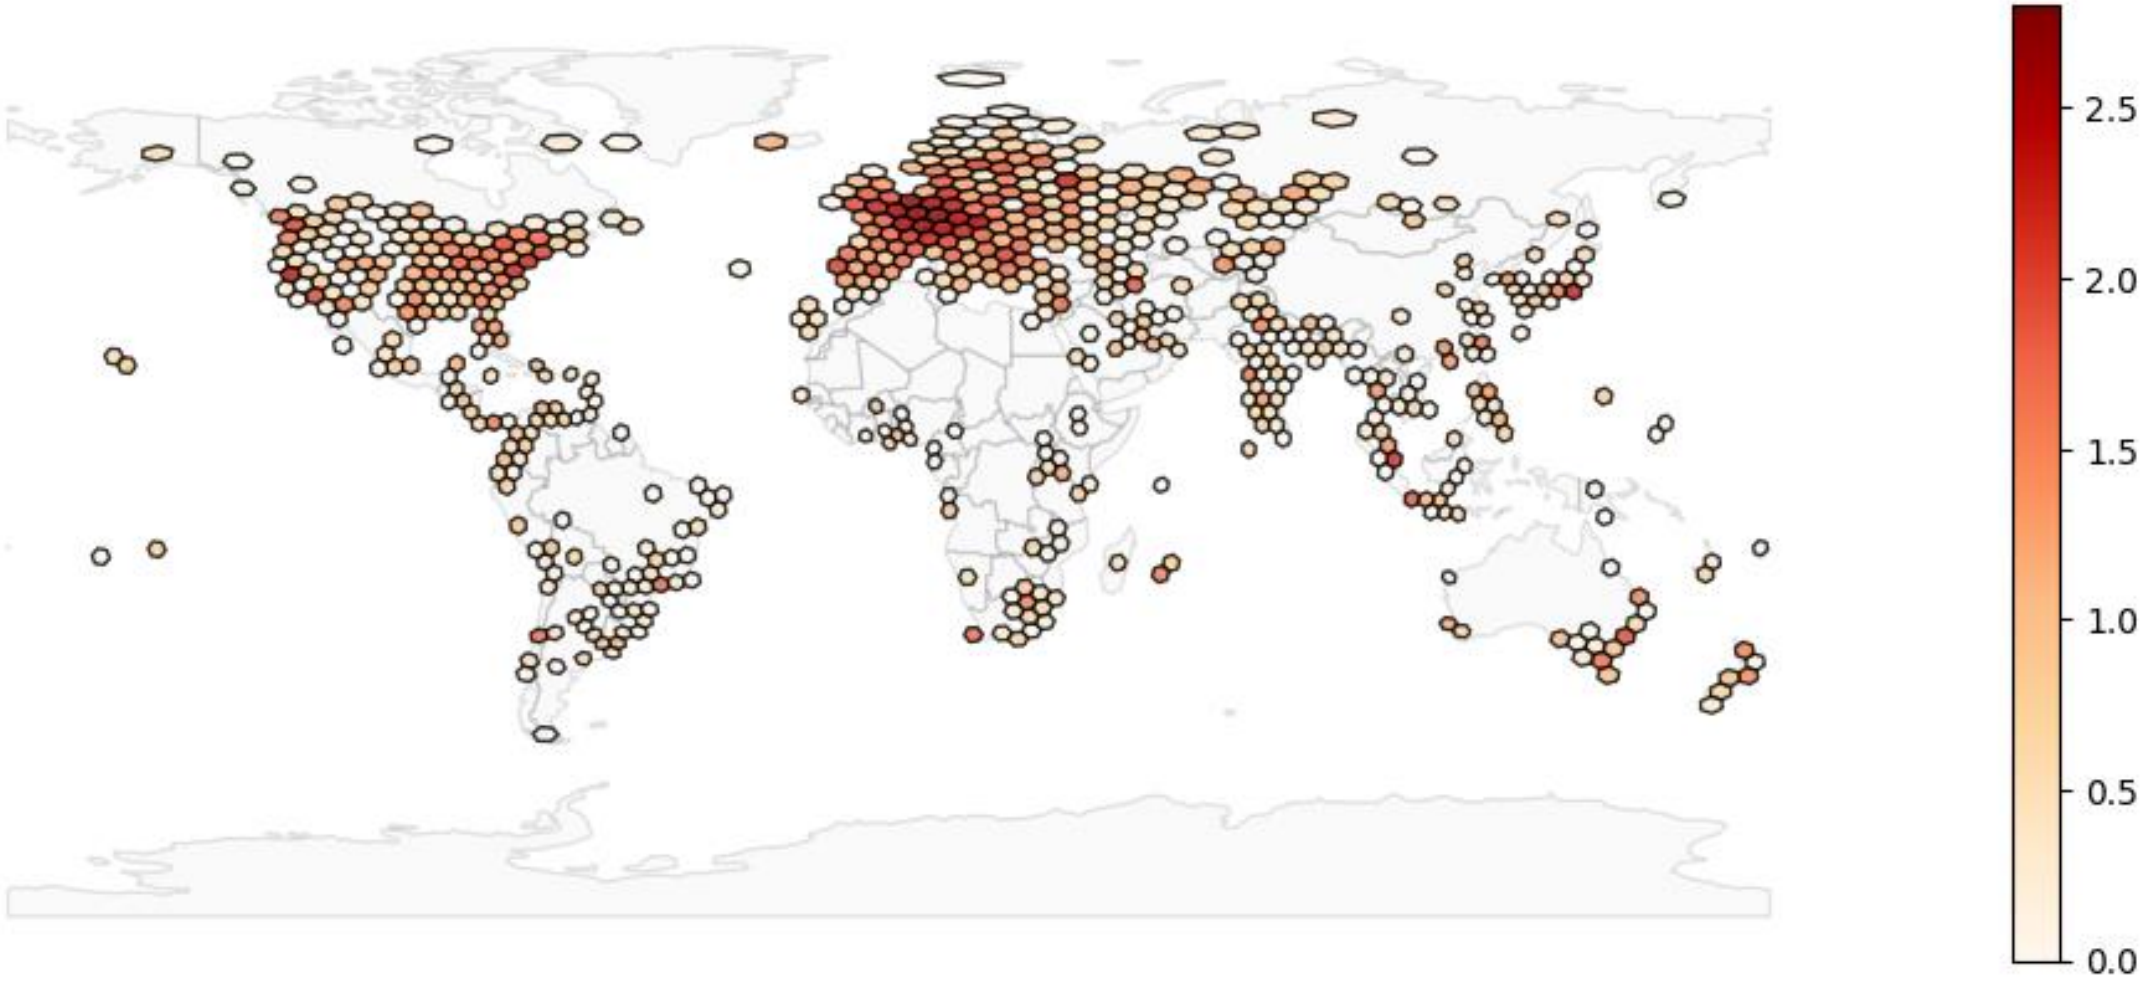

INTERPRETATIVE CAVEATS

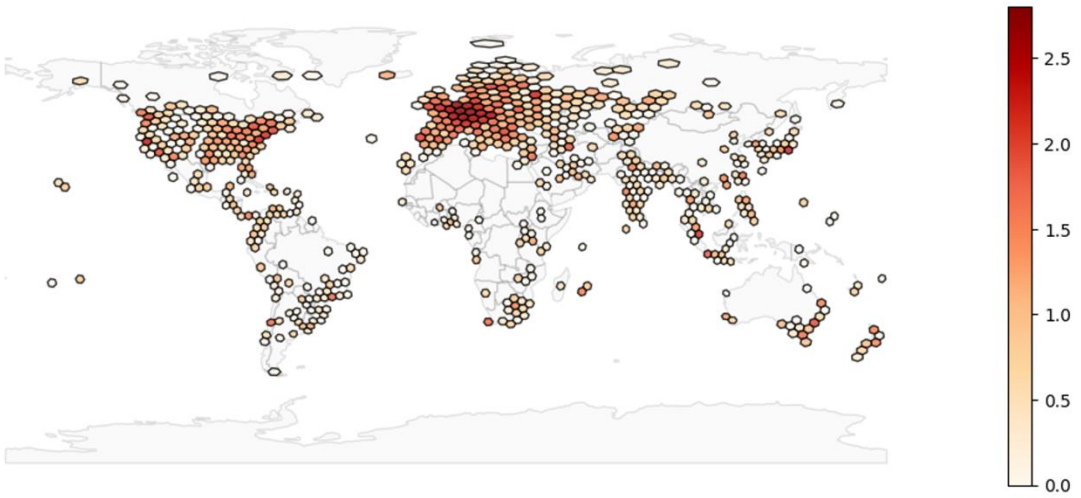

- ATLAS probes are wonderful but their placement has the potential to introduce significant biases for global research
- Probes' "power user" distribution is not ideally matched to any consumer population
- Undercounts are particularly significant in Africa and South Asia; no mobile probes
- Nothing here should be interpreted as a market share calculation

ATLAS LONG-RUNNING “WHO AM I” DNS QUERIES

- https://atlas.ripe.net/api/v2/measurements/8310237
  - dig +short -t txt o-o.myaddr.l.google.com.

- https://atlas.ripe.net/api/v2/measurements/8310245/
  - dig +short whoami.akamai.net.

- Returns “my address” --- i.e., the IP address performing the query
- Reveals the ‘ultimate resolver’ hitting the authoritative server
- All probes, once an hour, since August 2017

DNS RESOLVER SELECTION

- These long-running daily ATLAS experiments allow us to see the **final recursive resolver** that makes queries to authoritative resolvers on behalf of each ATLAS probe, every hour, for more than 7 years
- This IP address can be classified as local (same ASN) or (sometimes**) as part of the unicast set identified with the back end of an anycast global provider (e.g., Google 8.8.8.8, Cloudflare 1.1.1.1, Quad9 9.9.9.9)
- This approach looks past forwarders – if your local DNS forwards to 8.8.8.8, you're counted as 8.8.8.8
- ** Herein lies the complication….

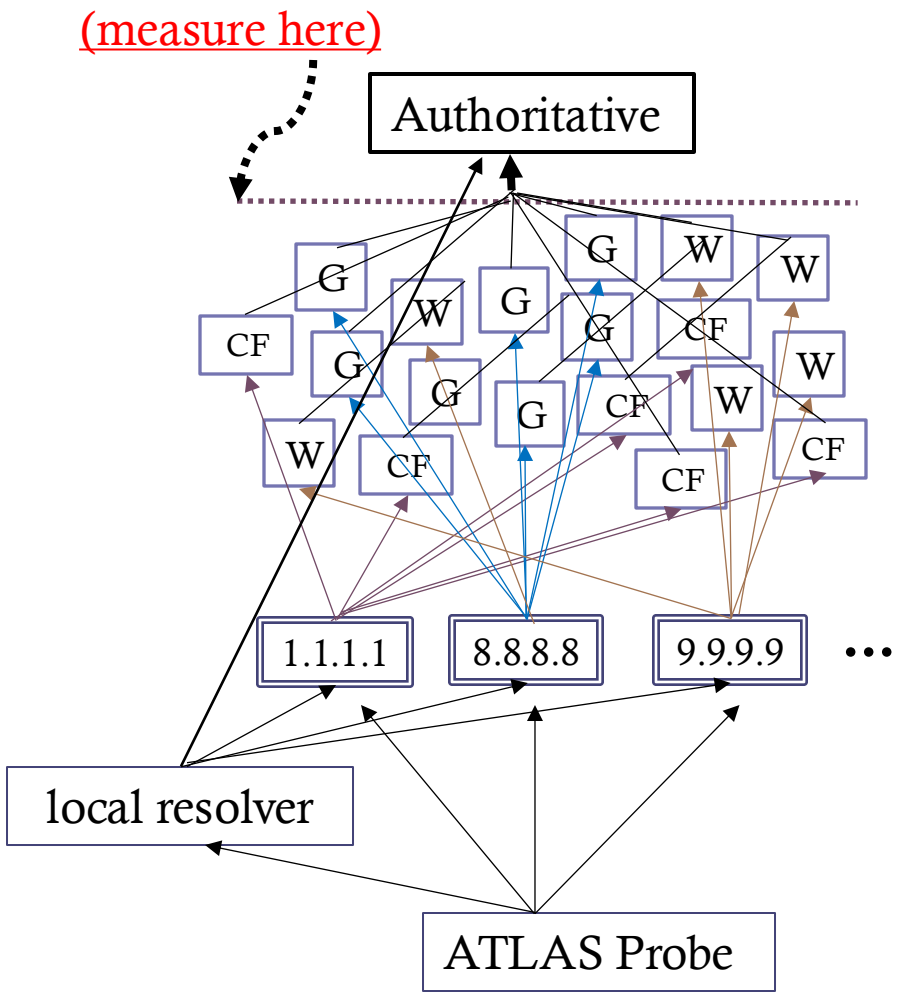

HOW MANY RESOLVERS REPRESENT A PROBE?

- MODE (most common situation): 2
- MEDIAN (50% more, 50% less): 94
- MAXIMUM: 8922

At the high end, many probes appear to be using resolver space across all major providers, changing every hour.

Hundreds of unique IPs each from Google, Cloudflare, Turk Telecom, Telecom Italia, Deutsche Telekom, Orange, Vodafone….

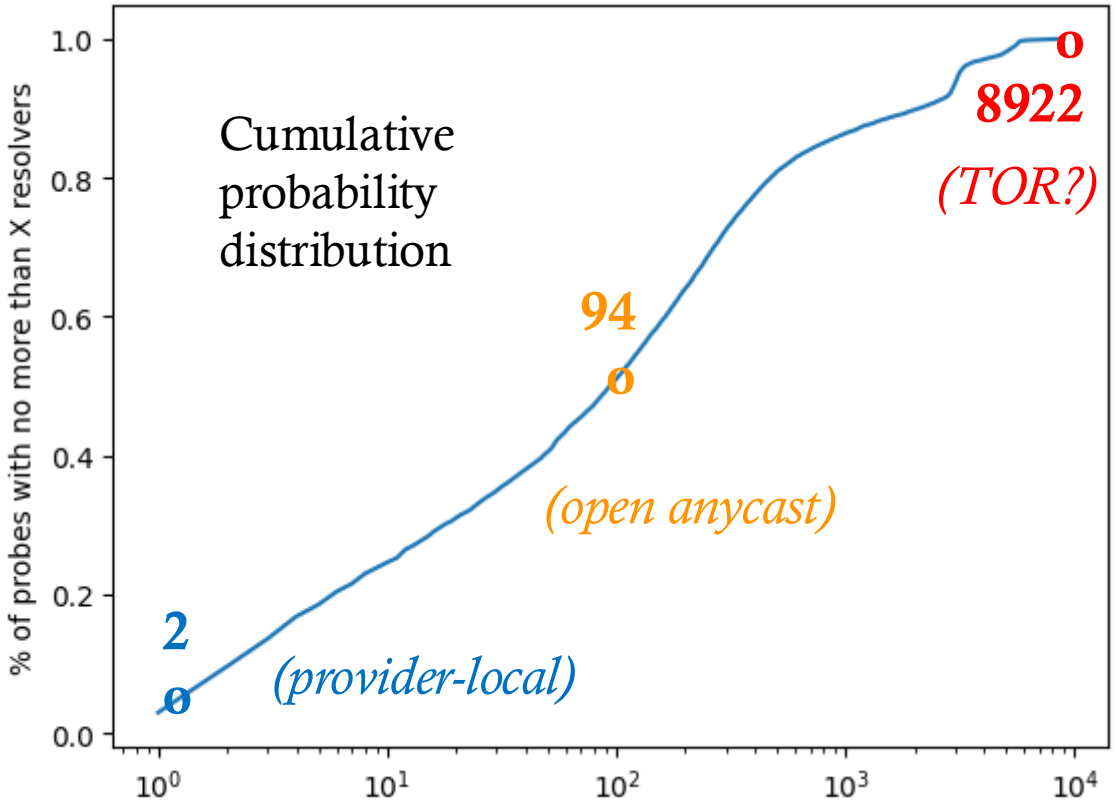

WORLDWIDE (ALL PROBES), LAST 5 YEARS

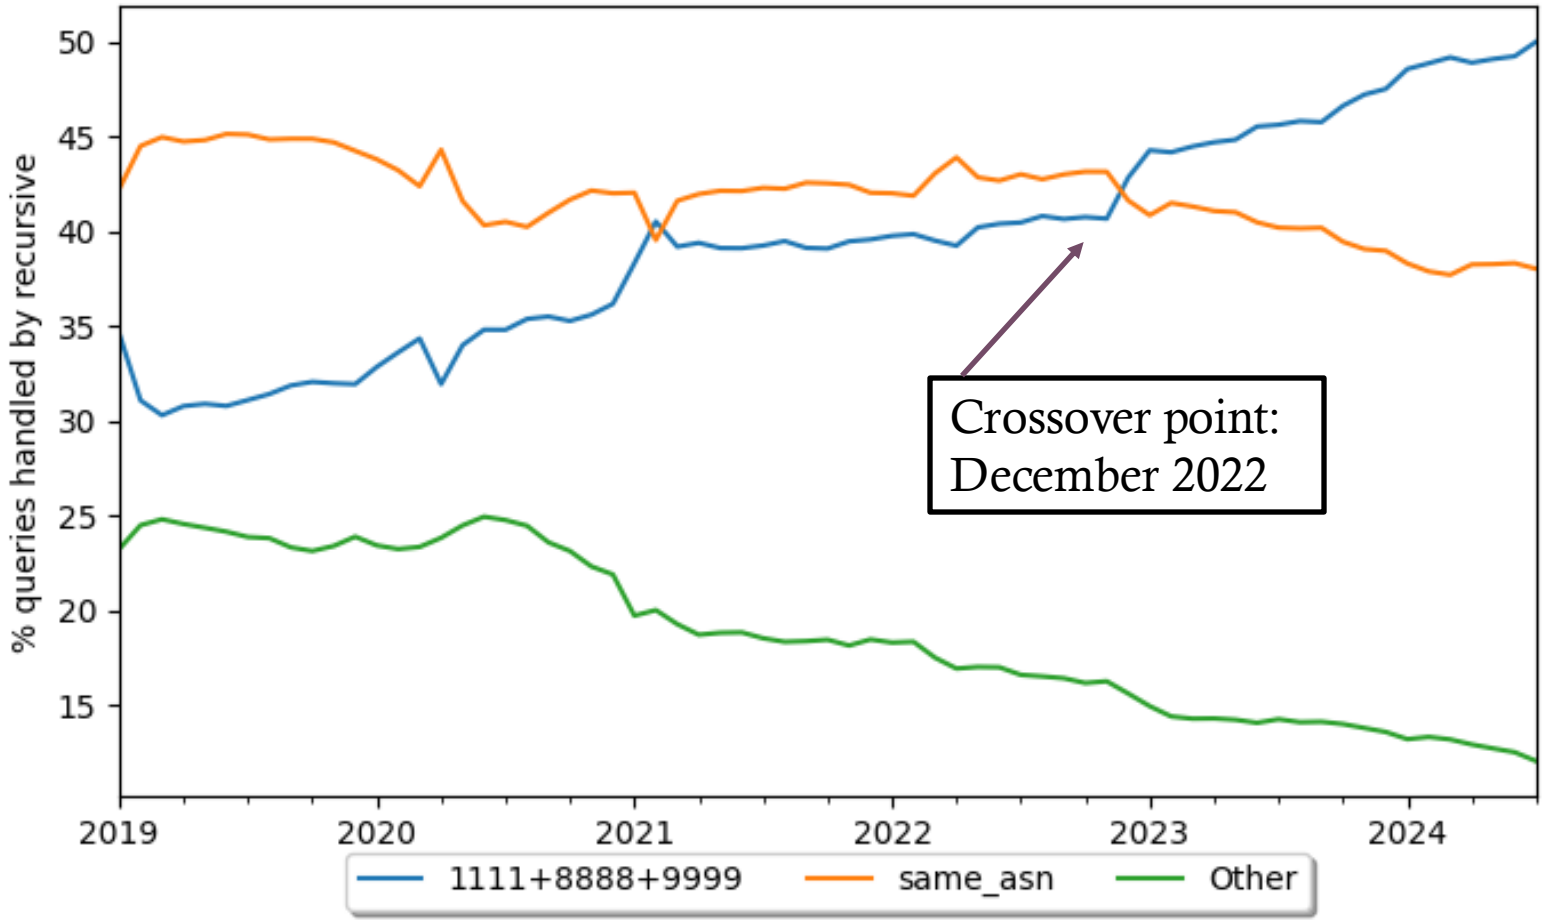

- **Big anycast open resolvers** have grown steadily, from ~30% to nearly 50%
- **Same ASN** resolver share has declined, but only modestly, by about 5%, to ~37%
- **Other resolver services** have declined.. one interpretation would be big-DNS consolidation

ASIA (ALL REGIONS), LAST 5 YEARS

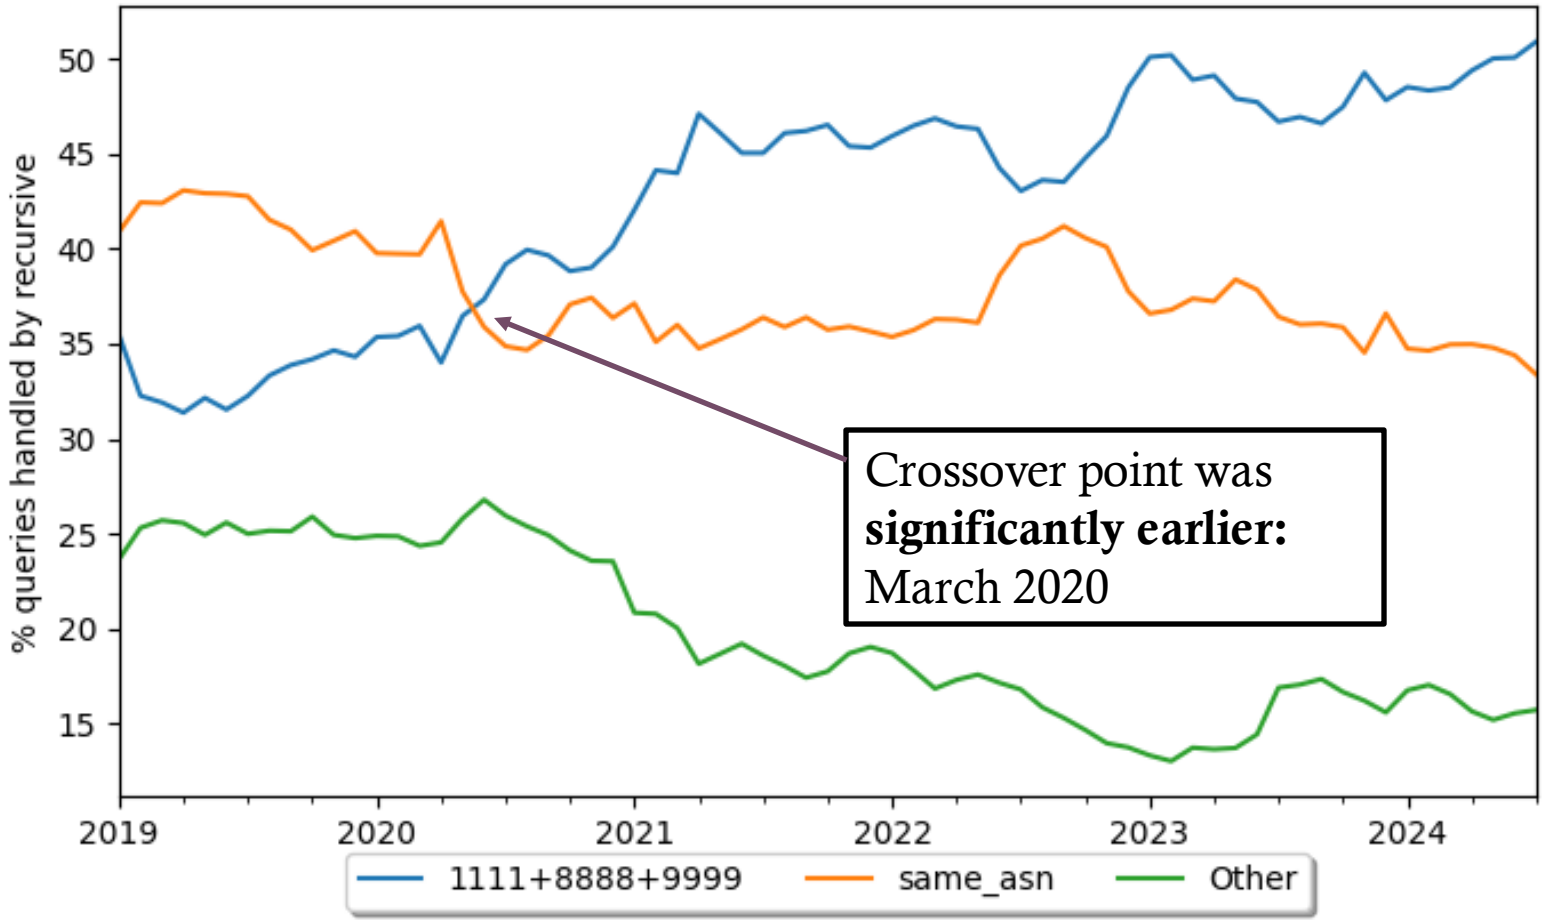

- **Big anycast open resolvers**
- **Same-ASN Resolver**
- **Other resolver services**

COMPLICATIONS: UNICAST ROSTERS FOR ANYCAST

- Seeing only the ultimate resolver unicast IPs, how can we subset and interpret the data?
- Some anycast services put together their backends out of whatever IPv4 resources are available to host their nodes, making it easy to undercount their attribution
- Some anycast services publish official rosters of IP resources and even geolocate their backend services (thank you!)
- The final resolver before authoritative may not be the initial resolver node you are mapped to locally

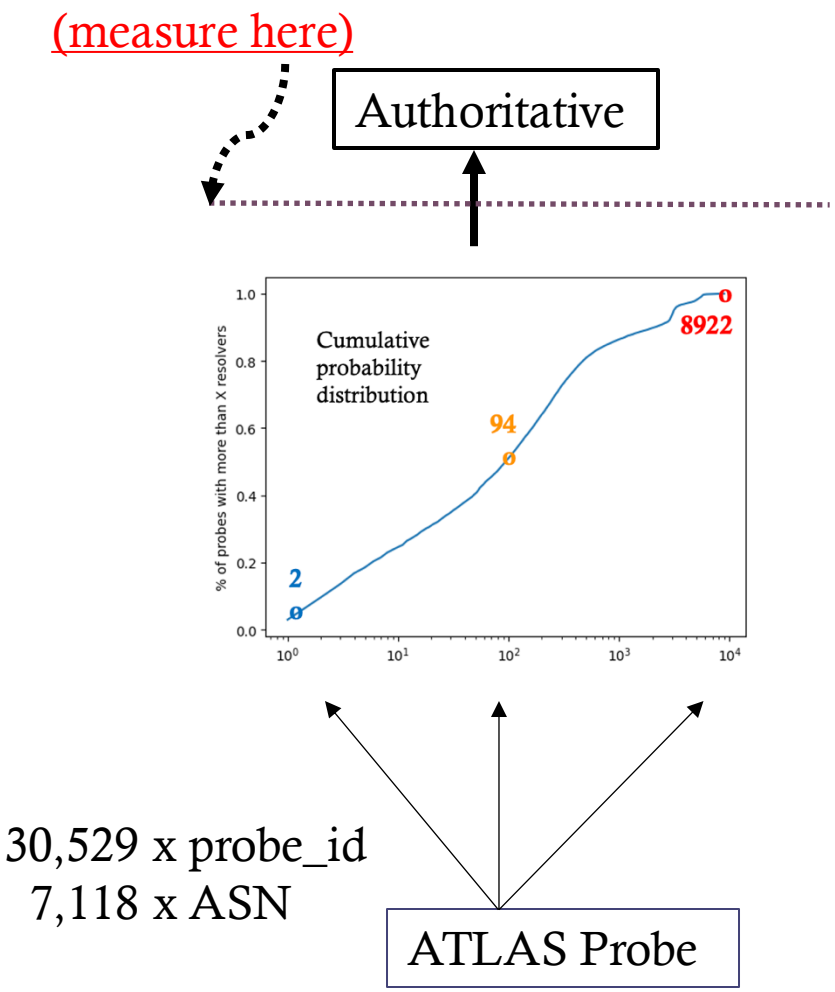

8.8.8.8 BACKEND

- https://developers.google.com/speed/public-dns/faq#locations
- https://www.gstatic.com/ipranges/publicdns.json

```
34.64.0.0/24 icn
34.64.1.0/24 icn
34.64.2.0/24 icn
34.101.0.0/24 cgk
34.101.1.0/24 cgk
34.101.2.0/24 cgk
34.153.64.0/24 dia
34.153.65.0/25 dia
34.153.65.128/26 dia
34.153.65.192/26 dmm
34.153.66.0/24 dmm
74.125.16.128/26 bom
172.253.11.0/25 zrh
172.253.11.128/26 cmh
172.253.11.192/26 grq
172.253.12.0/25 zrh
172.253.12.128/25 mil
172.253.13.0/25 kix
172.253.13.128/26 mil
172.253.13.192/26 waw
172.253.14.0/25 zrh
172.253.14.128/26 cmh
172.253.14.192/26 cgk
172.253.15.0/25 kix
172.253.252.0/24 icn
172.253.253.0/24 icn
172.253.254.0/24 dls
172.253.255.0/24 waw
173.194.90.0/24 cbf
173.194.91.0/24 scl
173.194.93.0/24 tpe
173.194.94.0/24 cbf
173.194.95.0/24 tul
173.194.96.0/25 dub
173.194.96.128/25 fra
173.194.97.0/24 chs
```

CLOUDFLARE BACKEND

AS13335

https://www.cloudflare.com/ips/

No geolocation volunteered?

QUAD9 BACKEND

- Reconstruct by having probes query **`CH TXT id.server. @9.9.9.9`**
- Replies are of the form
  - res111.pao.rrdns.pch.net
  - res200.syd.rrdns.pch.net
  - res231.qsin1.rrdns.pch.net
  - res320.iad.rrdns.pch.net
- Resolve these to unicast addresses, determine BGP origin ASNs, note airport geolocation hints
- Use the resulting end-of-time ASN/prefix set as an **admittedly incomplete** map of the historical quad9 backend
- AS{42,715,49544,7195,54285,51095,49115}

SIDEBAR: H3 VISUALIZATION

H3 is Uber's hierarchical hexagonal geospatial indexing system, originally designed to support the ridesharing dispatch system.

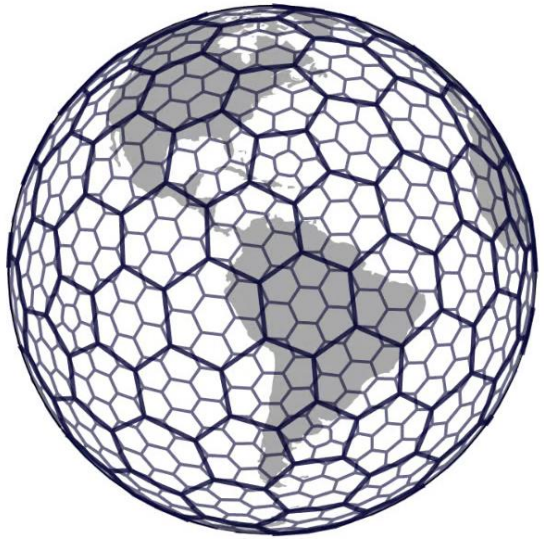

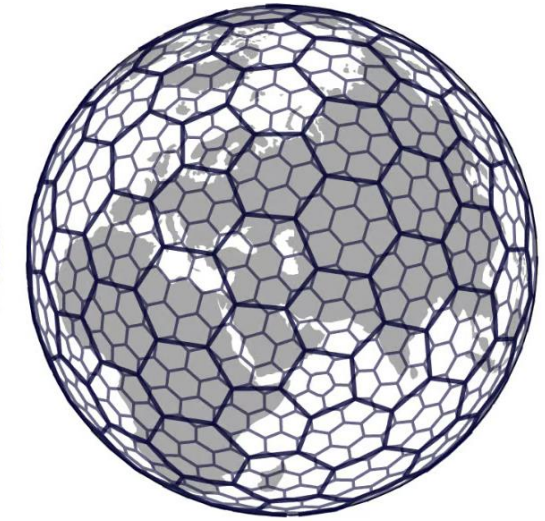

- Packages: h3py, geopandas
- hex = h3.latlng_to_cell(lat, lon, level)
- shapely.Polygon([pt for pt in h3.h3_to_geo_boundary(hex,True)])
- gpd.GeoDataFrame(df, geometry=df.hex.apply(…), crs='EPSG:4326')

GOOGLE (OCT ‘24)

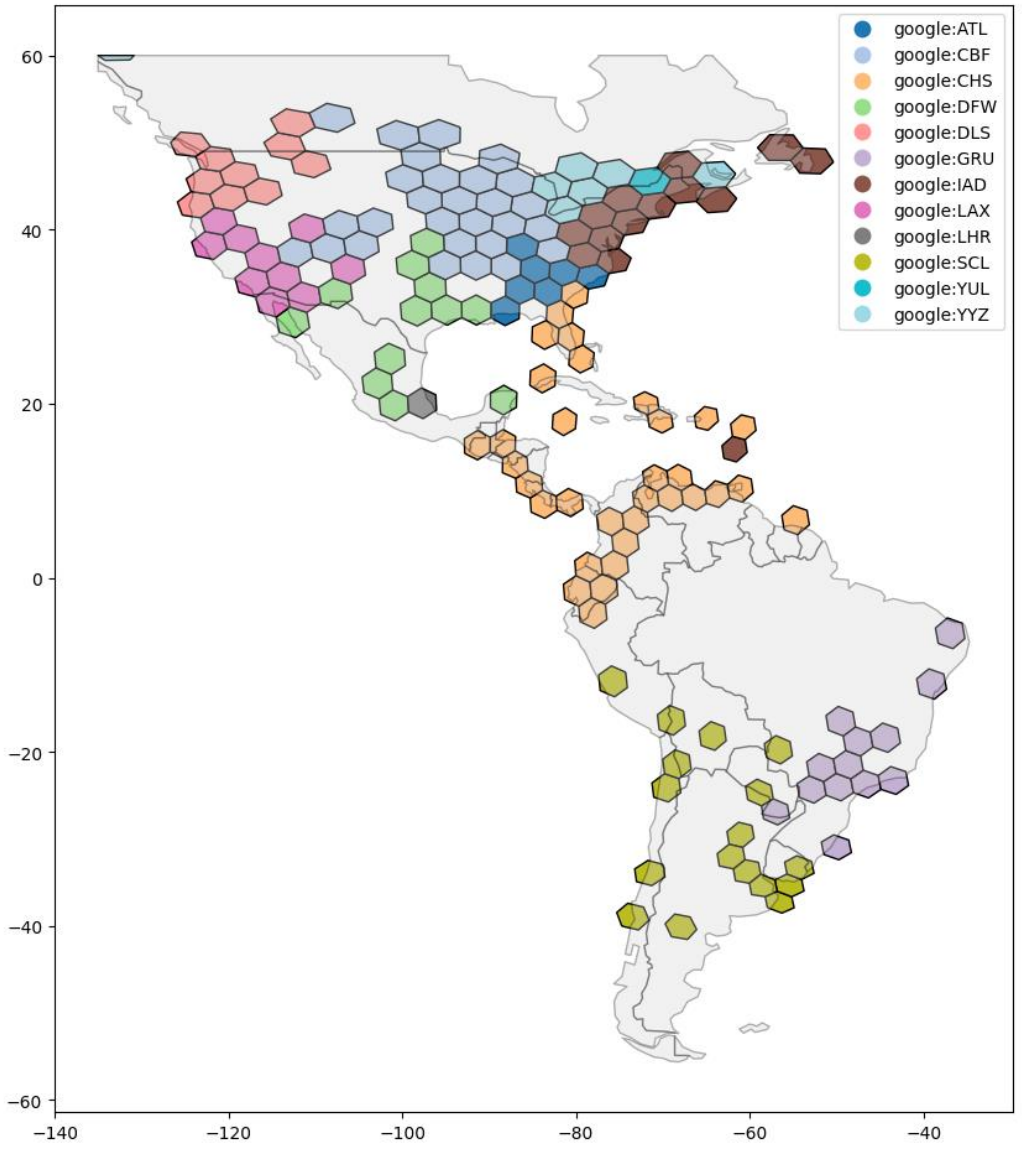

QUAD9 (OCT ‘24)

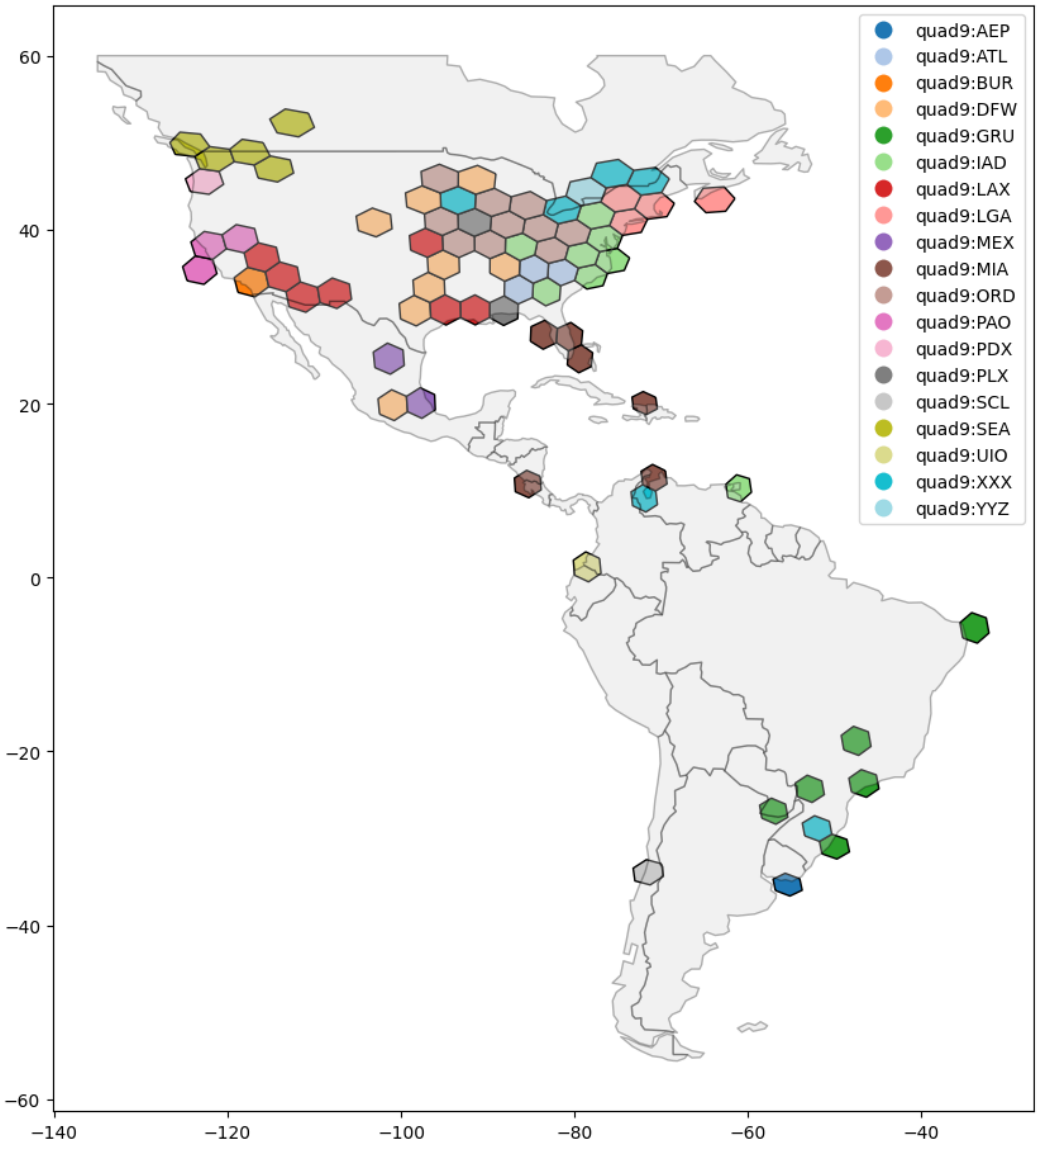

QUAD9 (OCT ‘24)

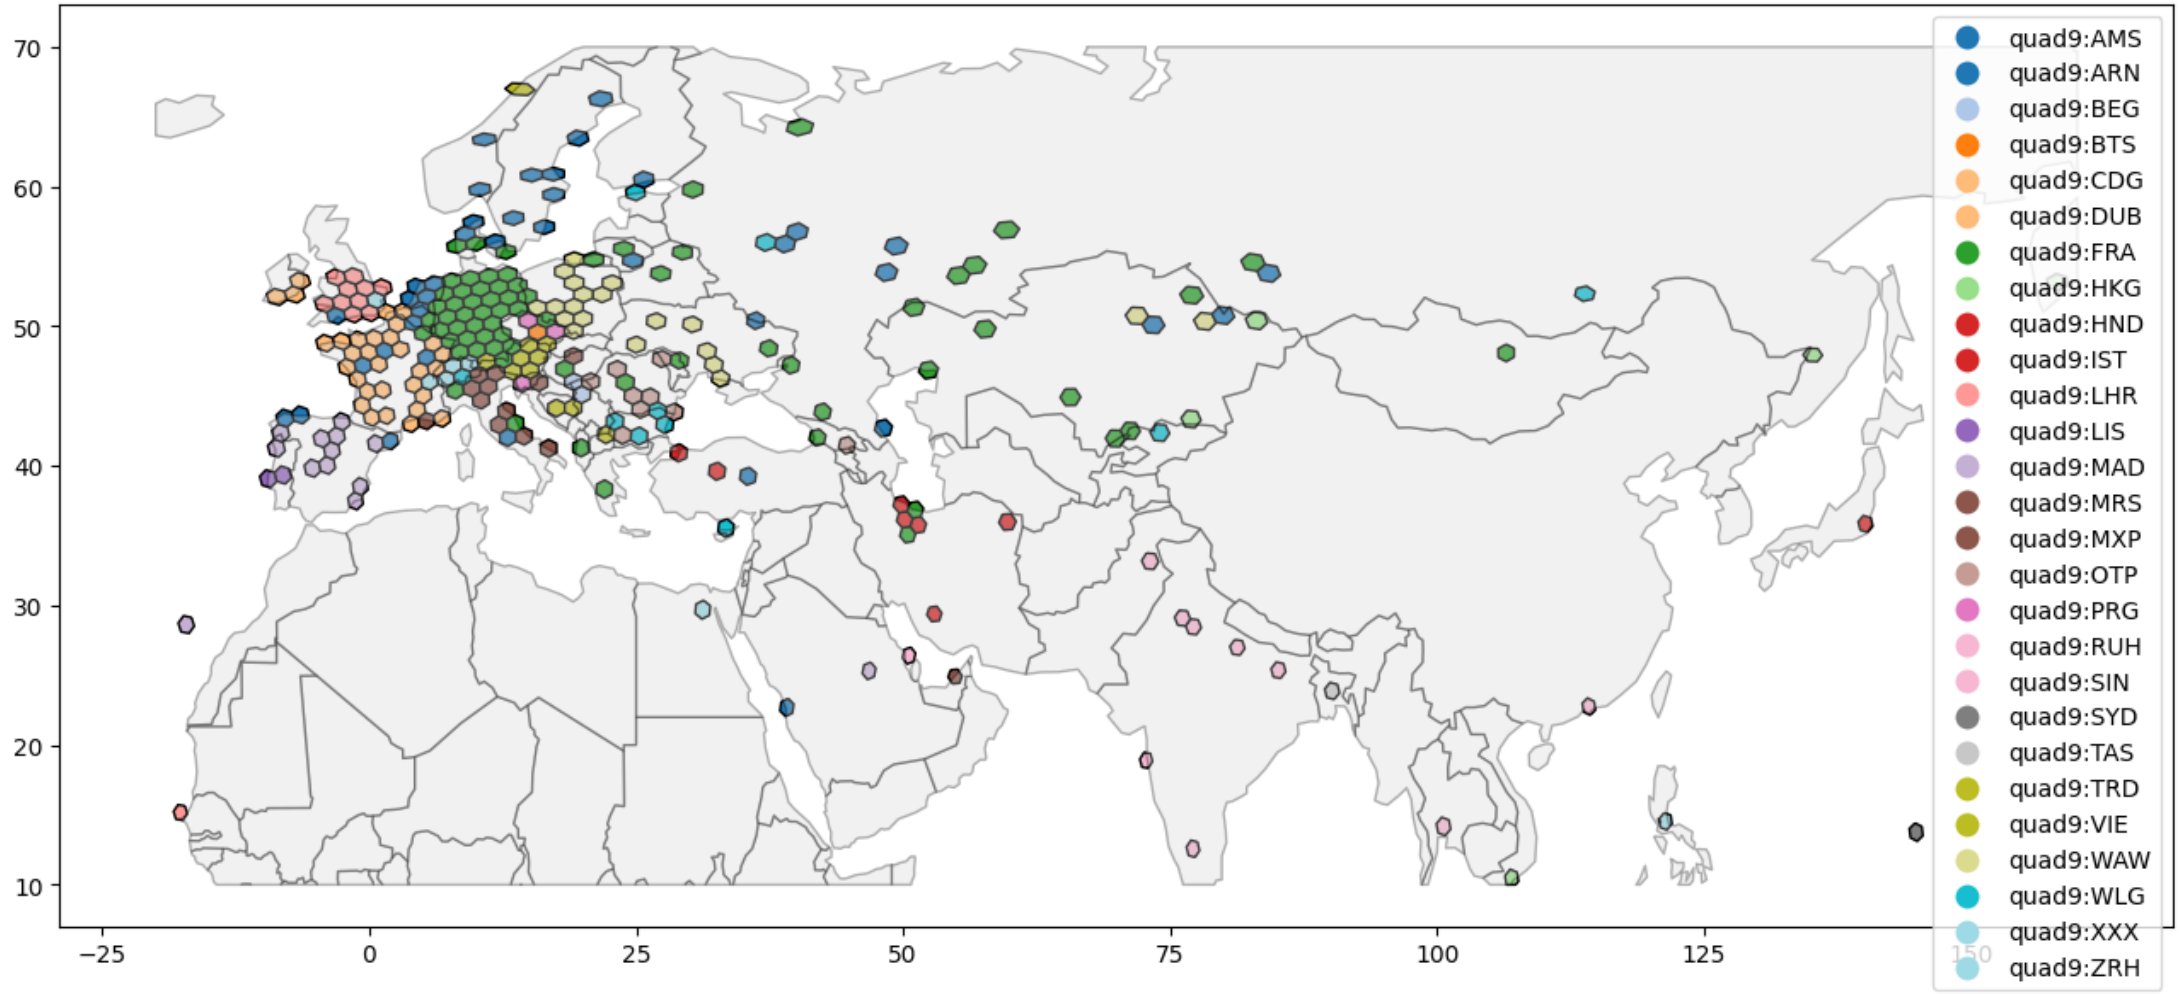
70
60
50
40
30
20
10
−25
0
25
50
75
100
125
150
quad9:AMS
quad9:ARN
quad9:BEG
quad9:BTS
quad9:CDG
quad9:DUB
quad9:FRA
quad9:HKG
quad9:HND
quad9:IST
quad9:LHR
quad9:LIS
quad9:MAD
quad9:MRS
quad9:MXP
quad9:OTP
quad9:PRG
quad9:RUH
quad9:SIN
quad9:SYD
quad9:TAS
quad9:TRD
quad9:VIE
quad9:WAW
quad9:WLG
quad9:XXX
quad9:ZRH

GOOGLE (OCT ‘24)

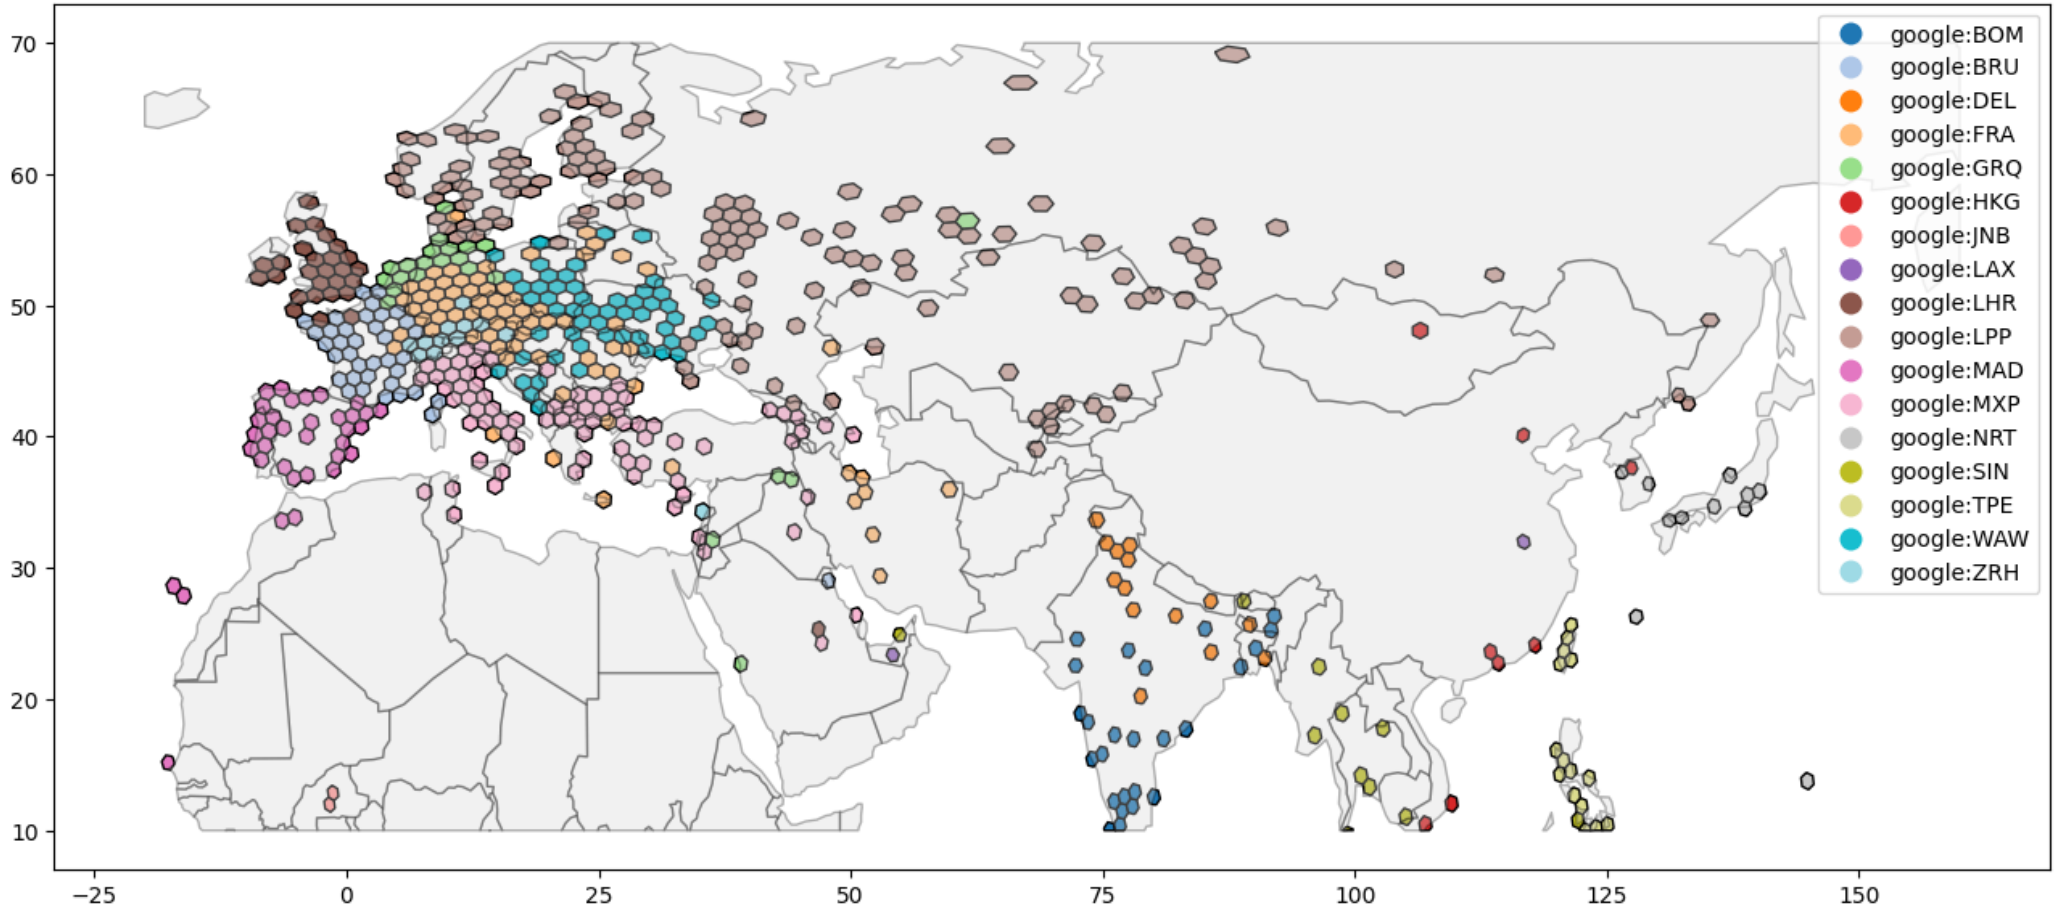

CONCLUSIONS

- Wonderful things are lurking in petabytes of ATLAS historical data!
- This is one of many data legacies that our community needs to preserve for future historians
- ATLAS probe placement doesn't always make it easy to deduce global and regional trends, but with our eyes open, we can still draw some conclusions
- Among ATLAS probe hosters, we see some evidence that use of "Big Recursive DNS" is growing and (perhaps) consolidating over the last 5 years, while local resolution has declined
- These big global anycast DNS providers do an amazing job of bringing resolution and caching close to clients, creating visible 'watersheds' of local service.

THANK YOU!

https://www.internethistoryinitiative.org

Mastodon: @IHI@cooperate.social

Email: jacowie@cyber.harvard.edu