

Recursives in the Wild: Engineering Authoritative DNS Servers

IETF 100 – IRTF MAPRG | 2017-11-13 | Singapore

Moritz Müller^{1,2}, **Giovane C. M. Moura**¹,

Ricardo de O. Schmidt^{1,2}, John Heidemann³

¹SIDN Labs, ²University of Twente, ³USC/Information
Sciences Institute



Before we start...

- How many DNS operators in the room?
- Role: point of view of a DNS operator that wants to reduce latency to its services
- Why? : **Time (latency) is money**

Many reports (google, amazon): report : high latency, less searches

- Paper presented at ACM IMC2017 in London



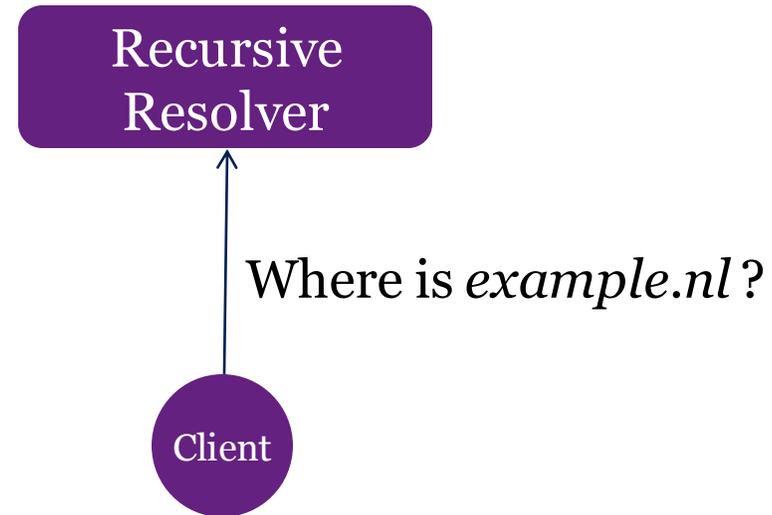
Introduction



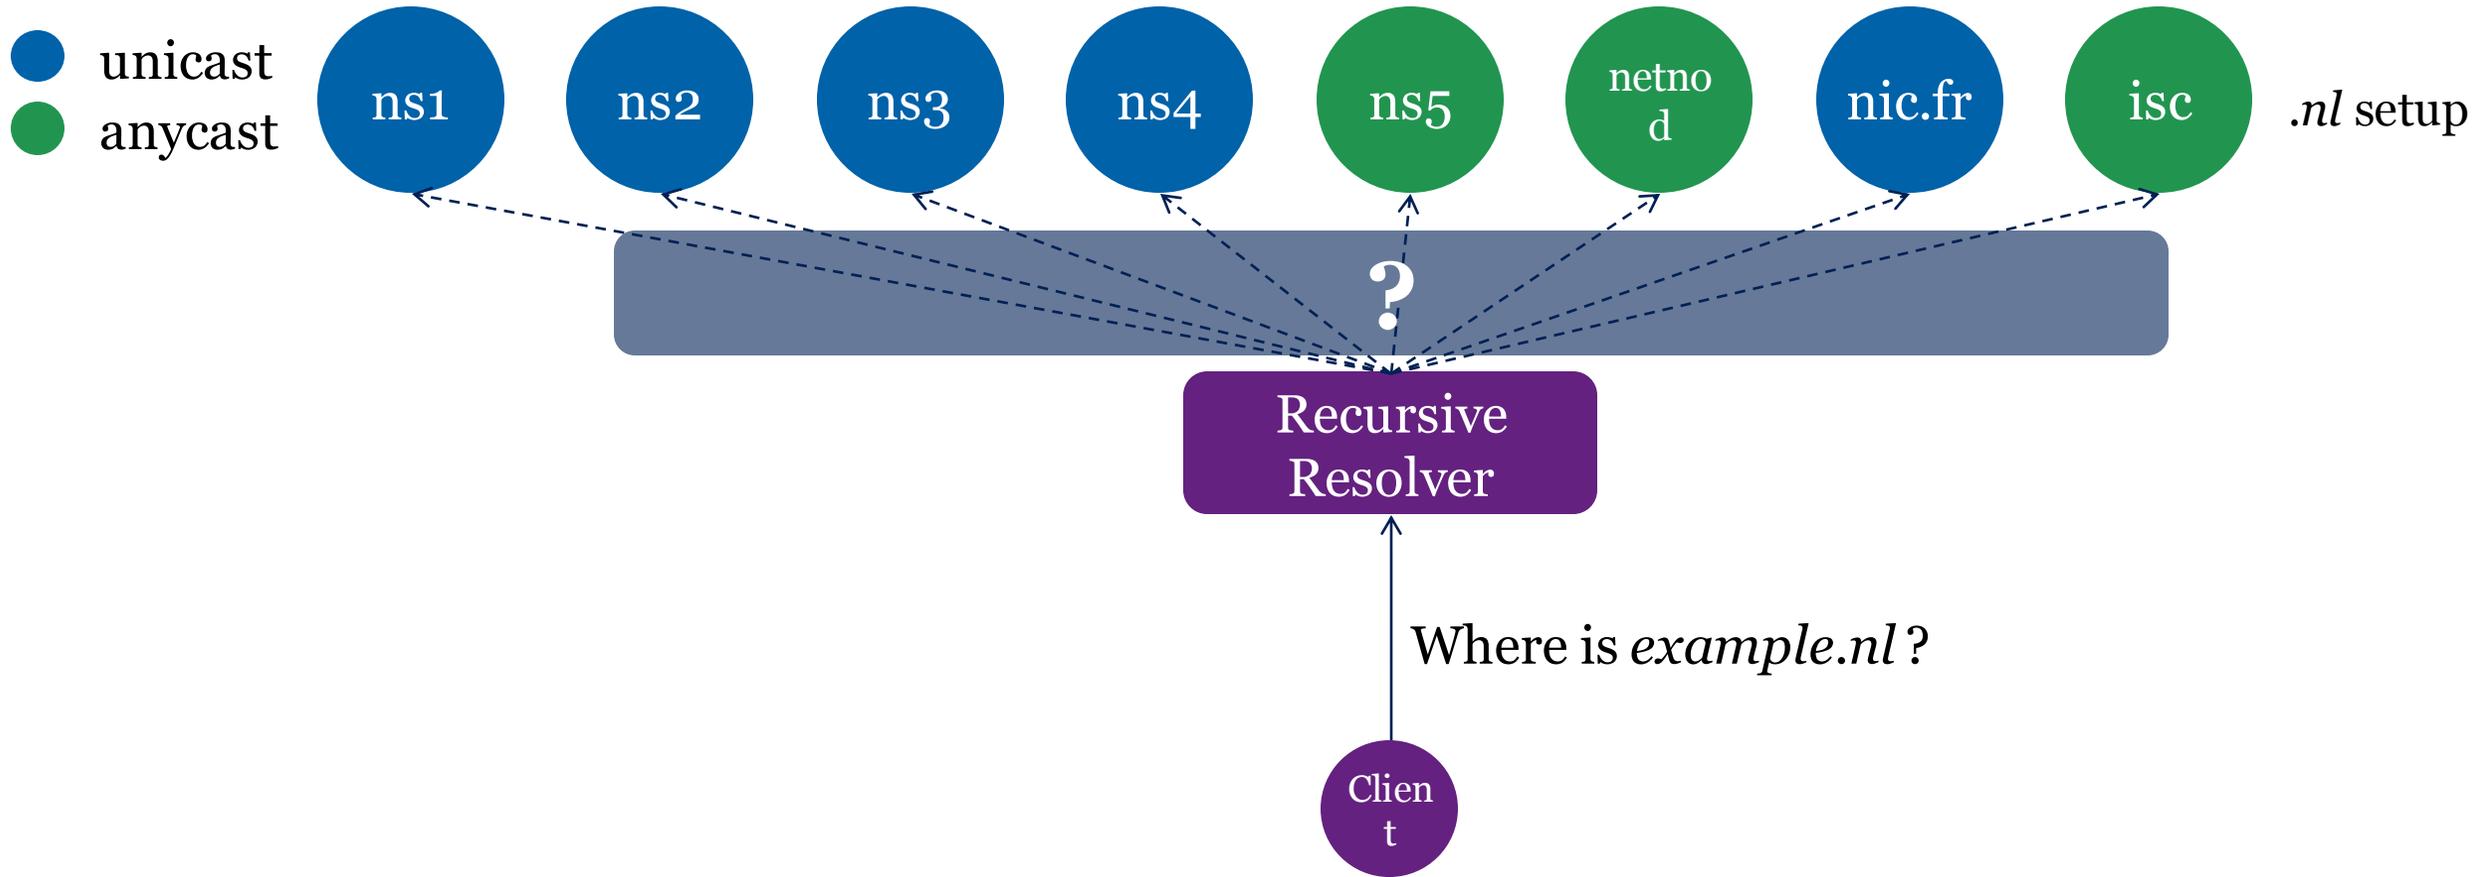
- Example of a DNS authoritative setup:
8 authoritative name servers for the same zone (.nl)
Could be the same for any second-level domain
- Redundancy for high availability



Introduction

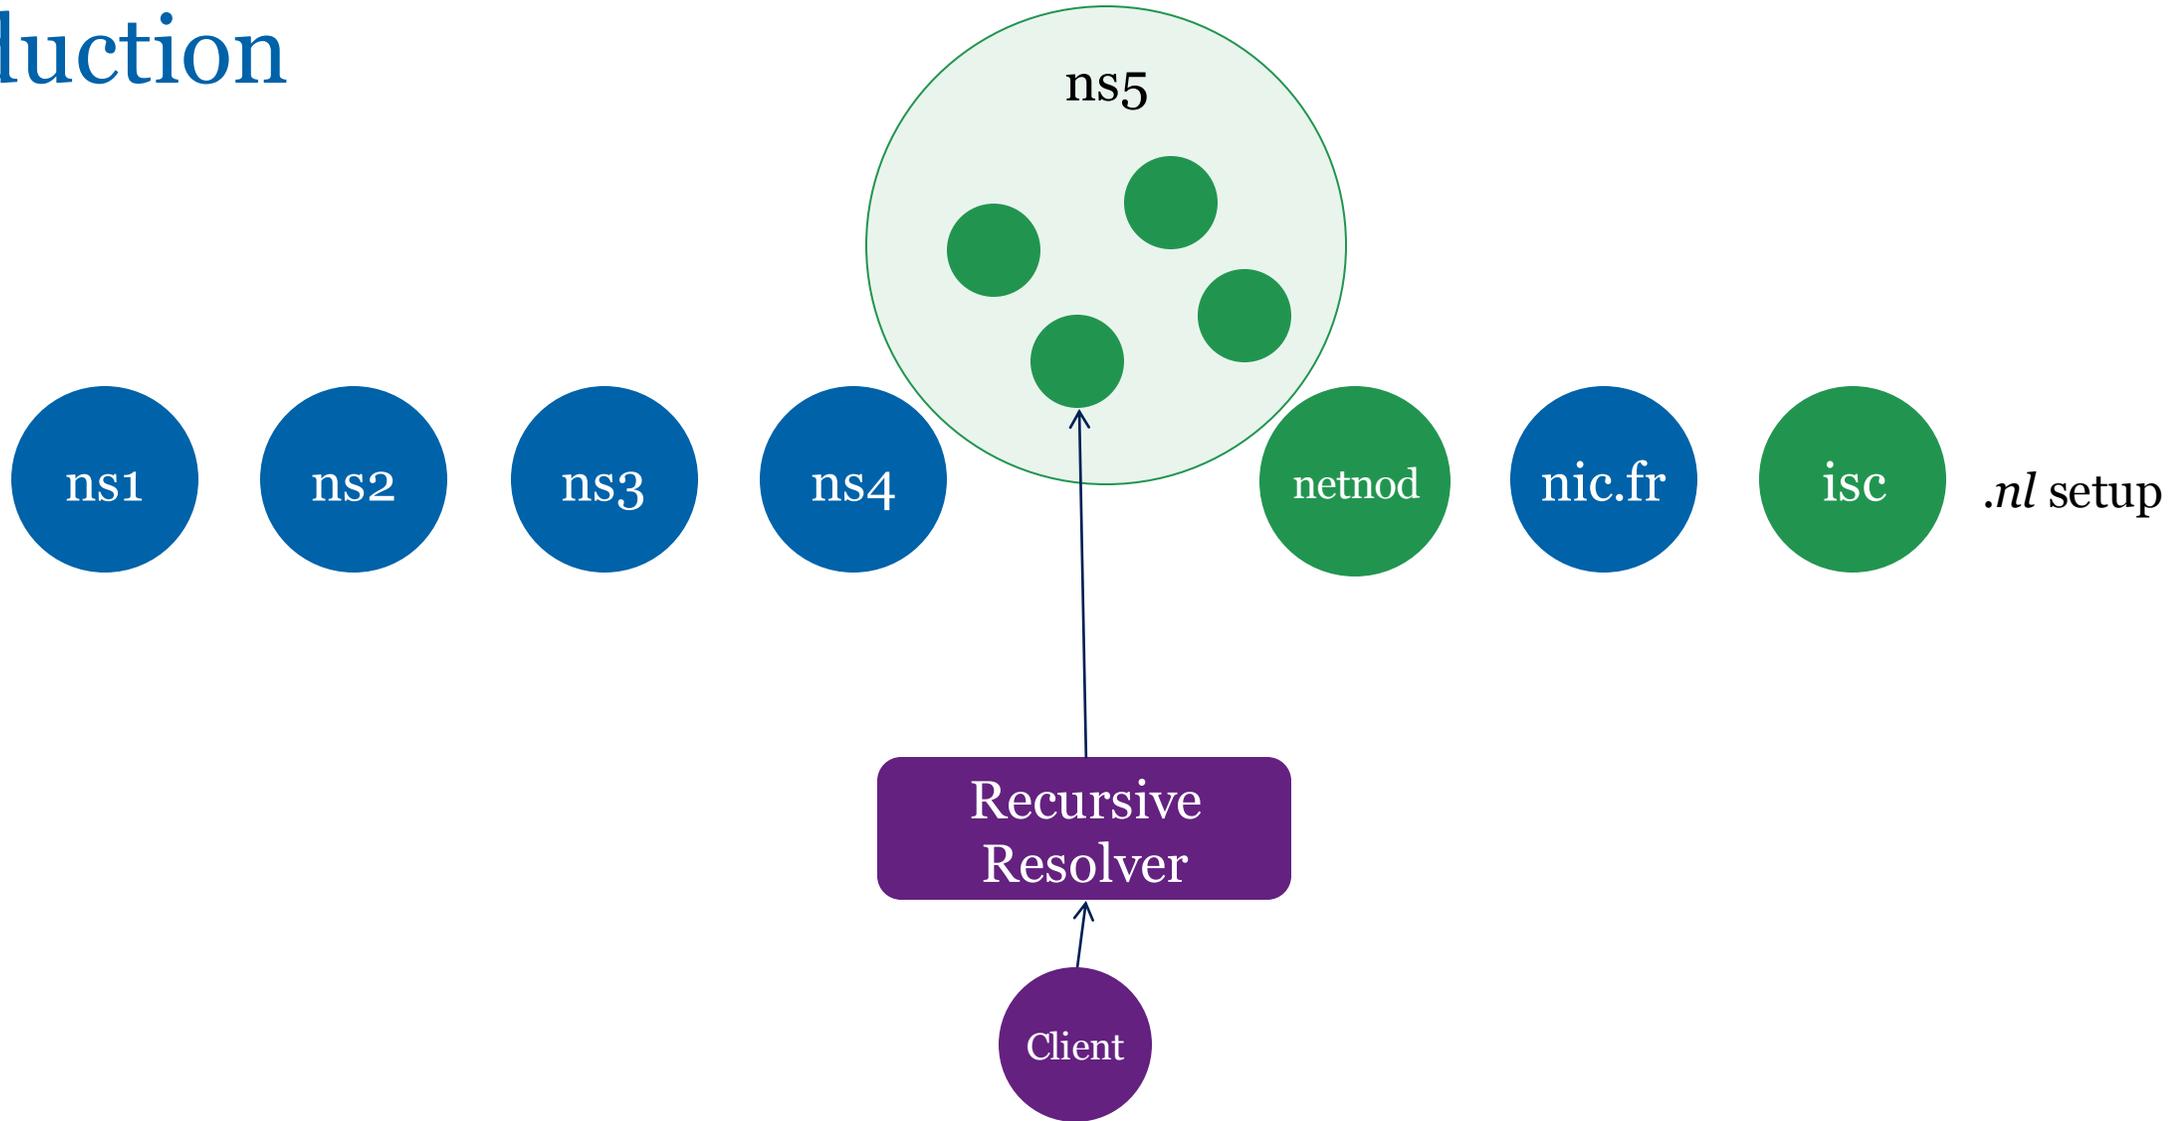


Introduction



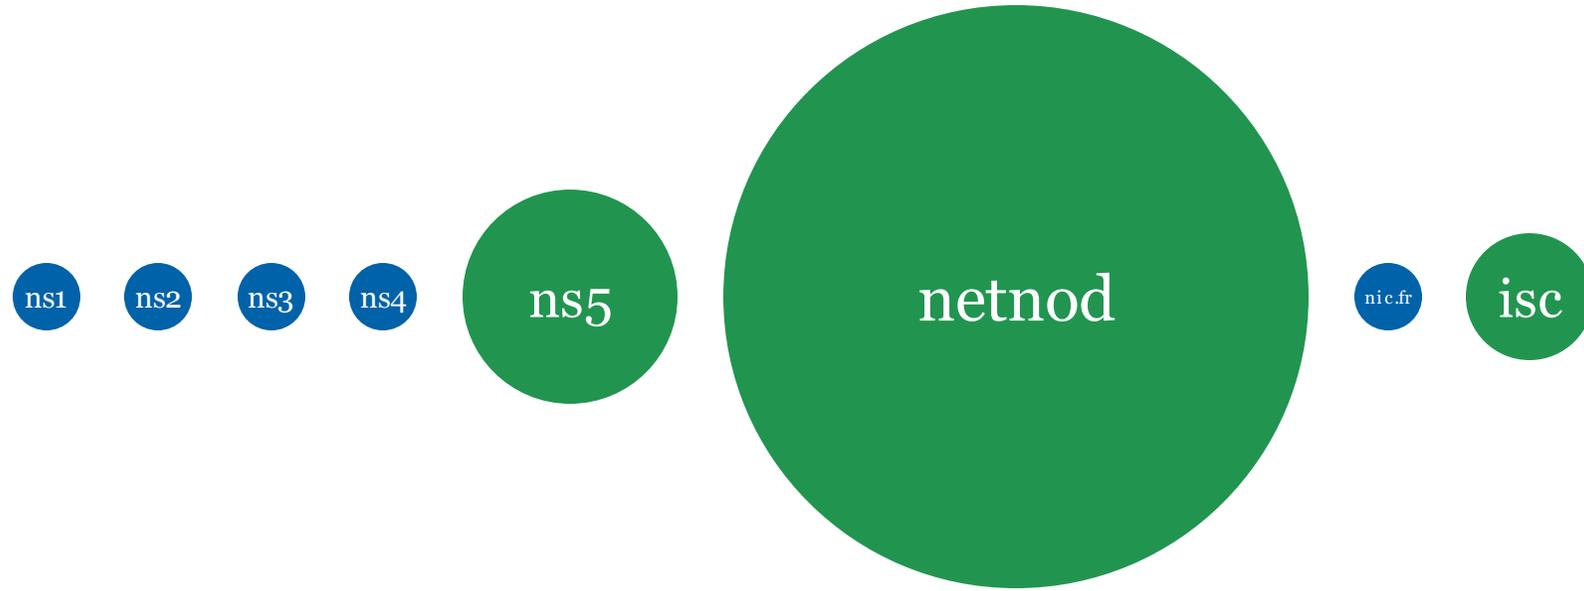
Introduction

- unicast
- anycast



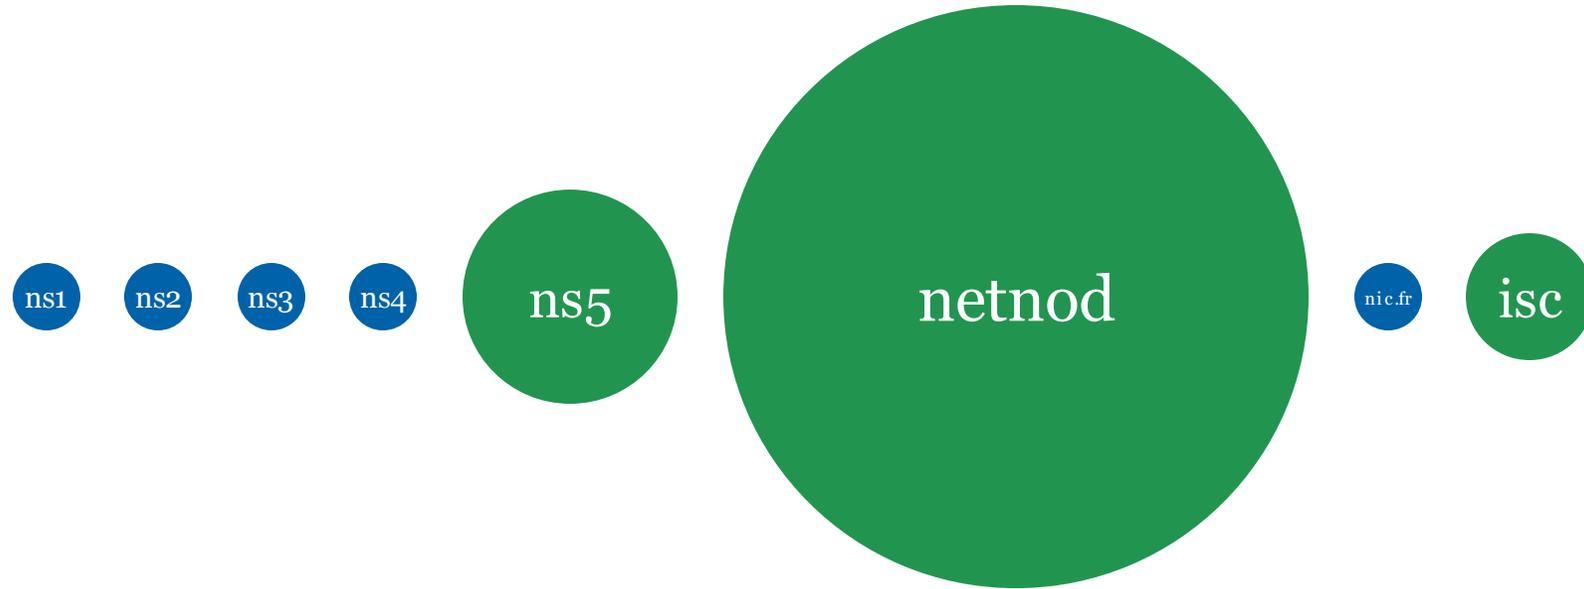
Introduction

area
relative to
the number
of *sites*



Introduction

area relative
to the
number of
sites

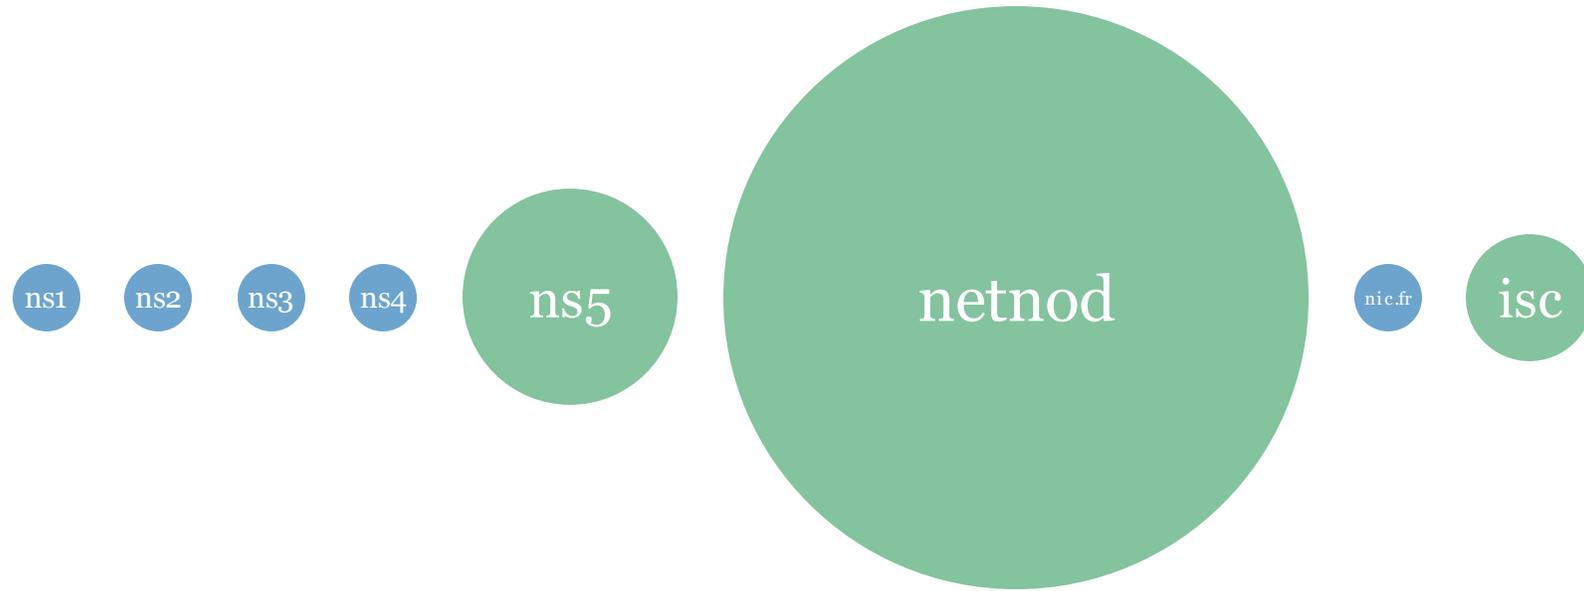


area relative
to the
number of
queries

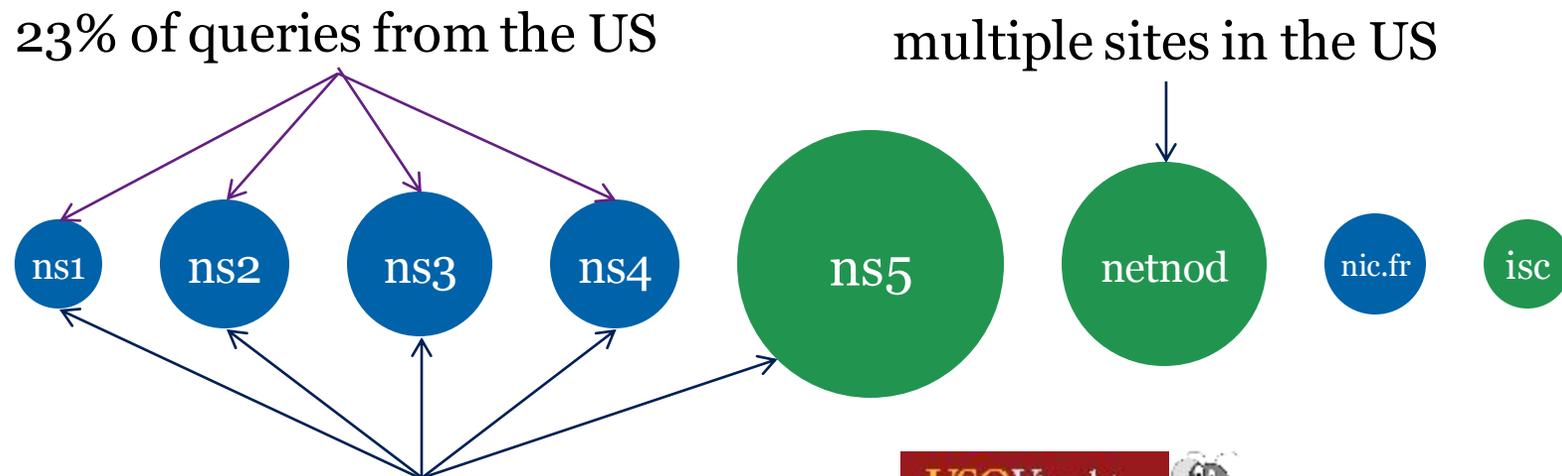


Introduction

area relative to
the number of
sites

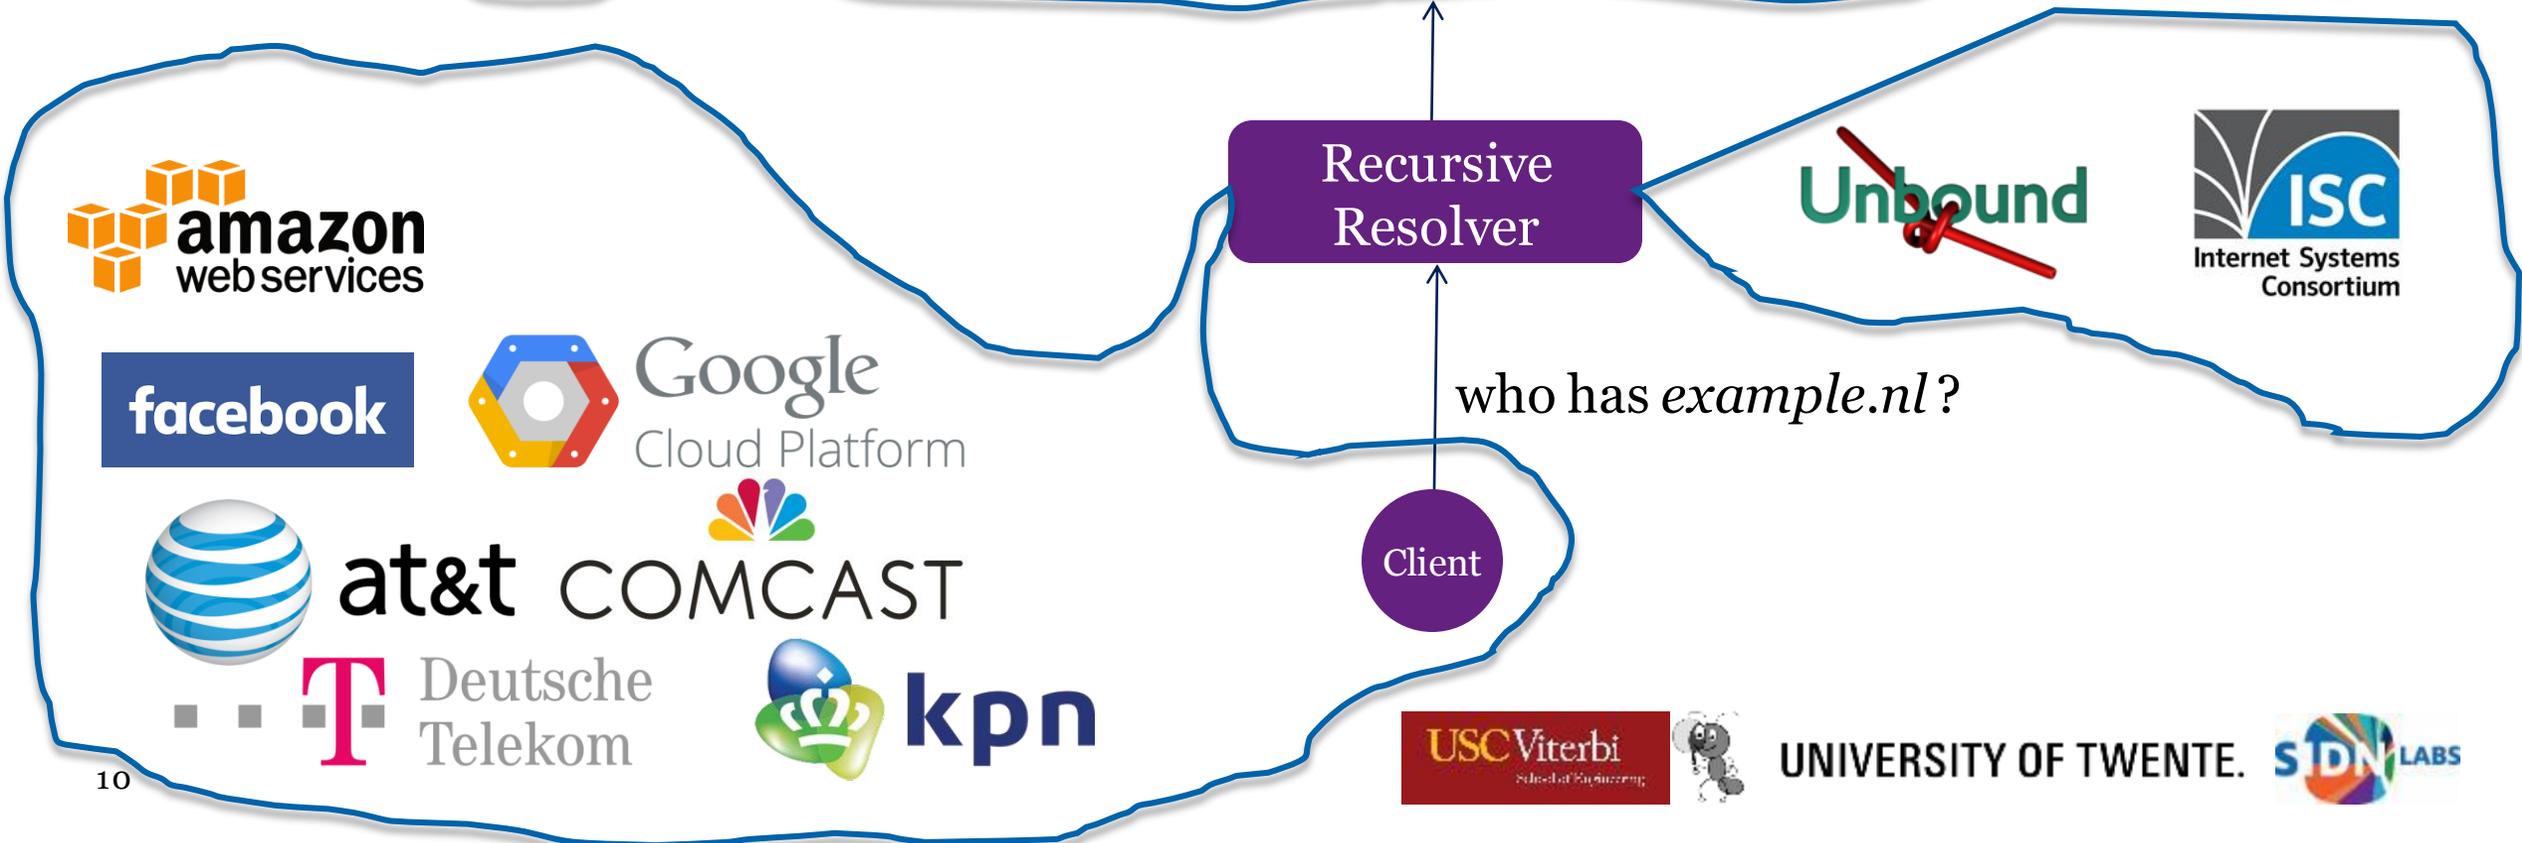


area relative to
the number of
queries



Why is that?

- unicast
- anycast



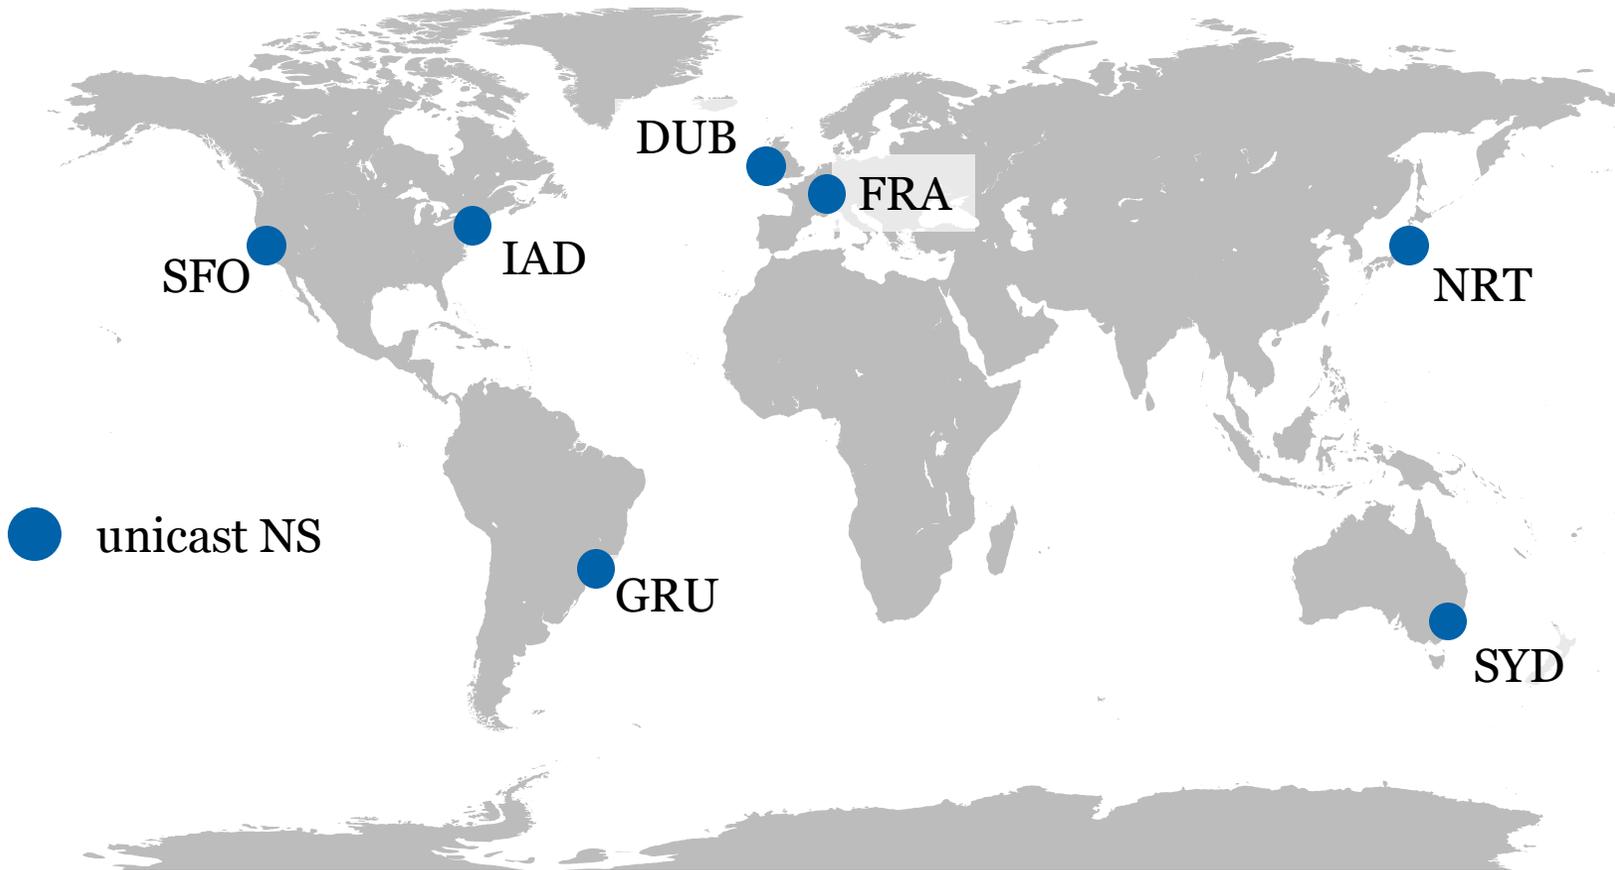
Research Questions

- How do recursive resolvers select authoritative name servers?
- [1] says, most *implementations* prefer faster responding authoritatives
- but what is the overall behaviour *in the wild*?
- To improve performance, how should DNS operators design their authoritatives?

[1] Yu, Y., Wessels, D., Larson, M., and Zhang, L.
Authority Server Selection in DNS Caching Resolvers.
SIGCOMM Computer Communication Review 42,
2 (Mar. 2012)



Measurement Design



7 measurement setups:

GRU+NRT

DUB+FRA

FRA+SYD

GRU+NRT+SYD

DUB+FRA+IAD

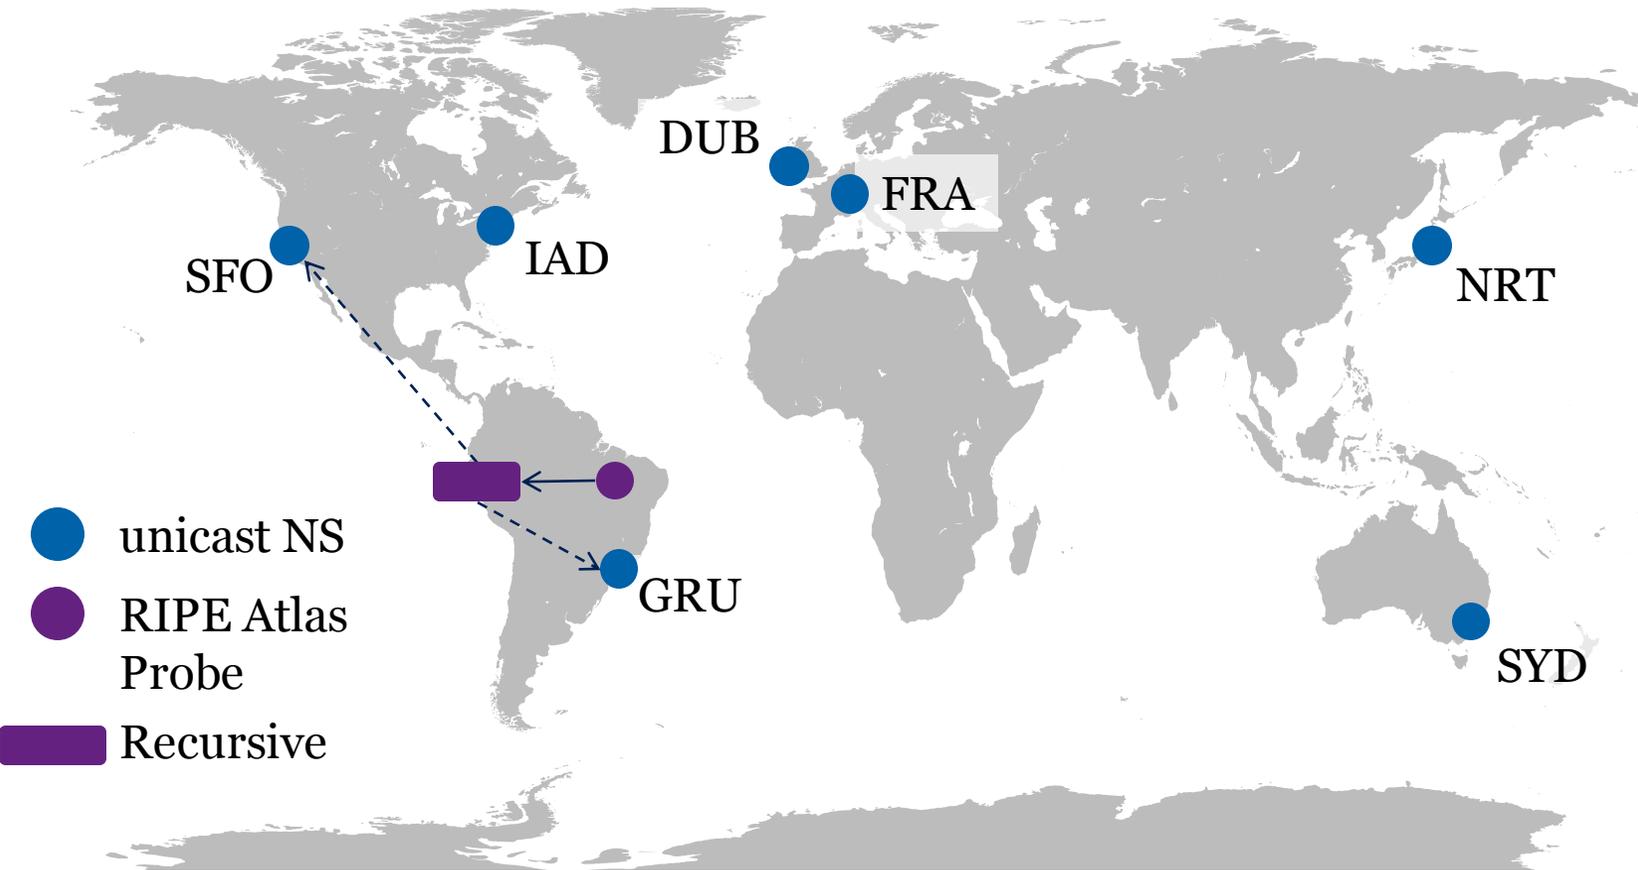
DUB+GRU+NRT+SYD

DUB+FRA+IAD+SFO

IPv4 only (for now)



Measurement Design



7 measurement setups:

GRU+NRT

DUB+FRA

FRA+SYD

GRU+NRT+SYD

DUB+FRA+IAD

DUB+GRU+NRT+SYD

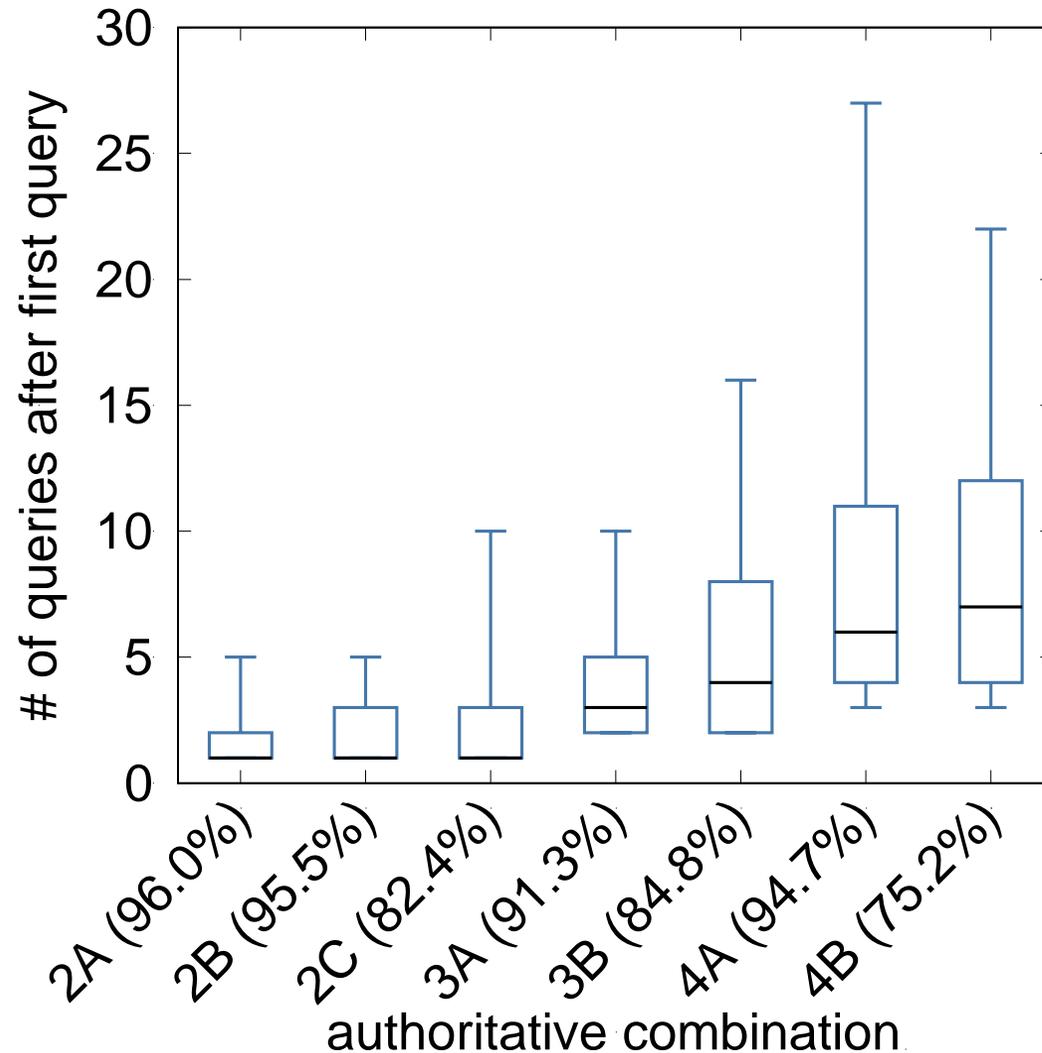
DUB+FRA+IAD+SFO

1 hour each measurement, ever 2 min

~9000 Ripe Atlas probes from ~2500 ASes

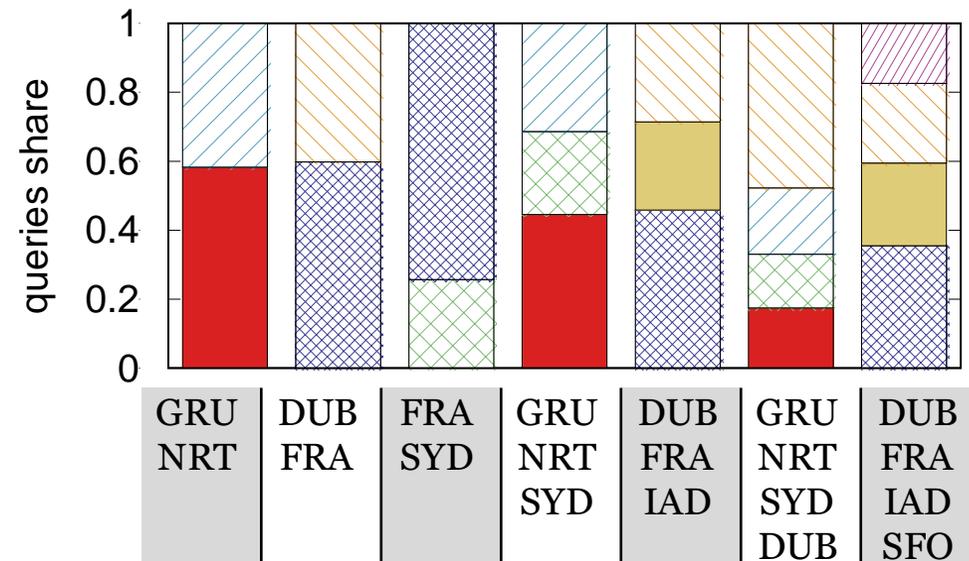
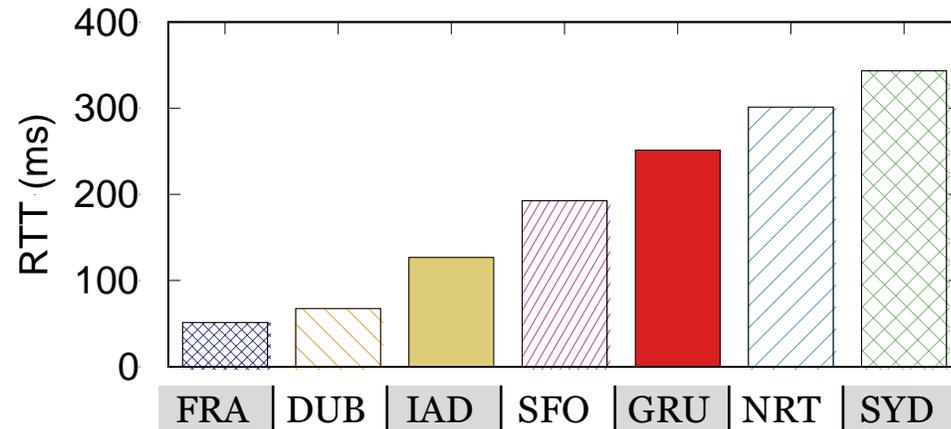


Do recursives query all authoritatives?

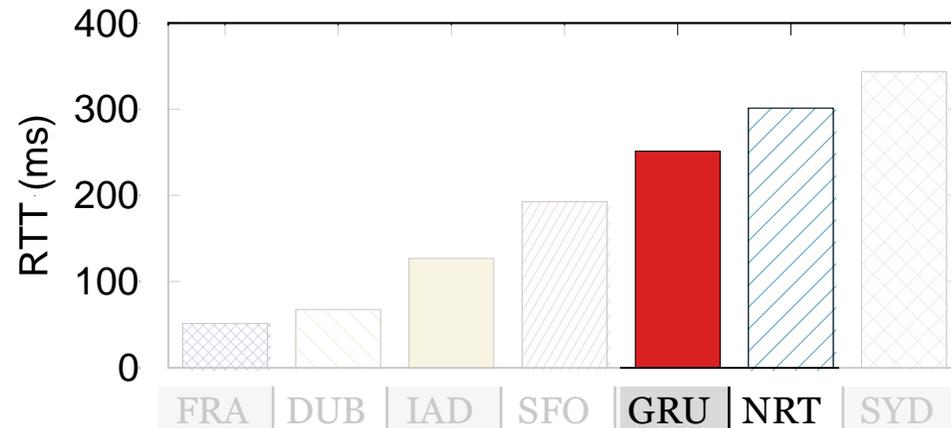


- Majority will quickly query ALL authoritatives
- Meaning “one bad apple may spoil the whole bunch” ?

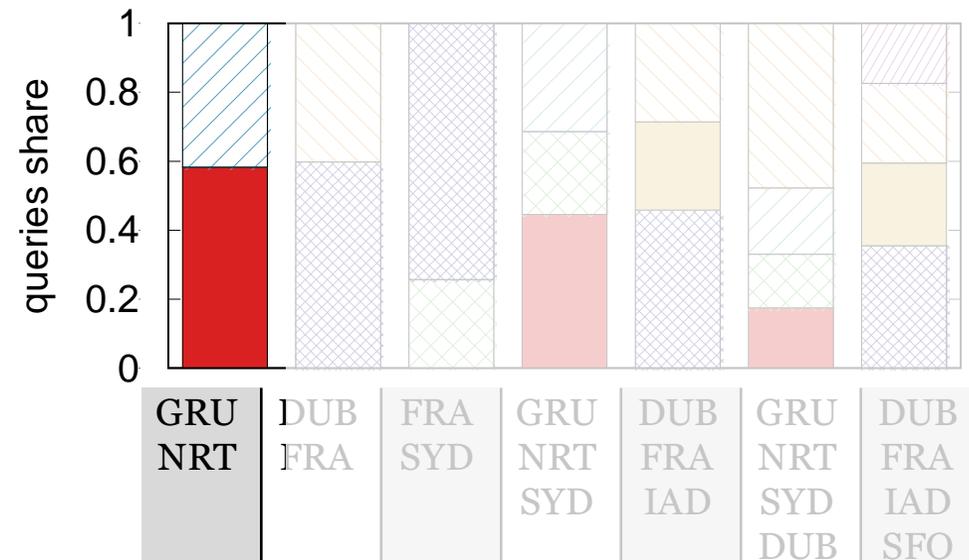
How do recursives distribute their queries over time?



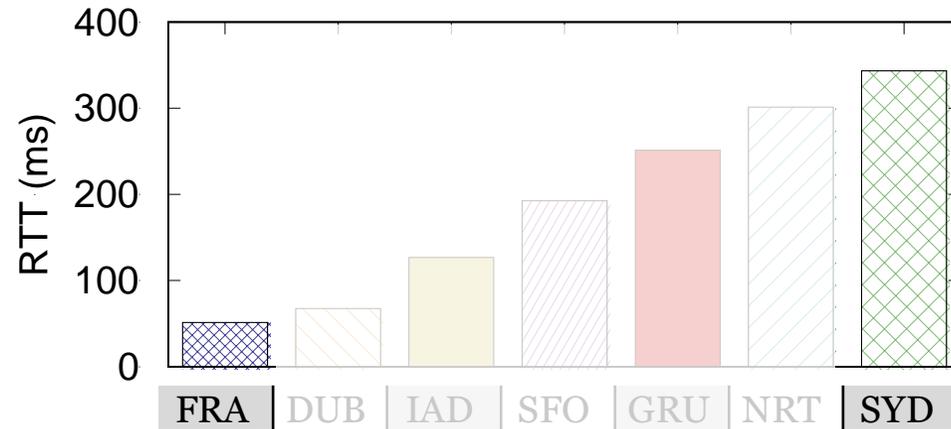
How do recursives distribute their queries over time?



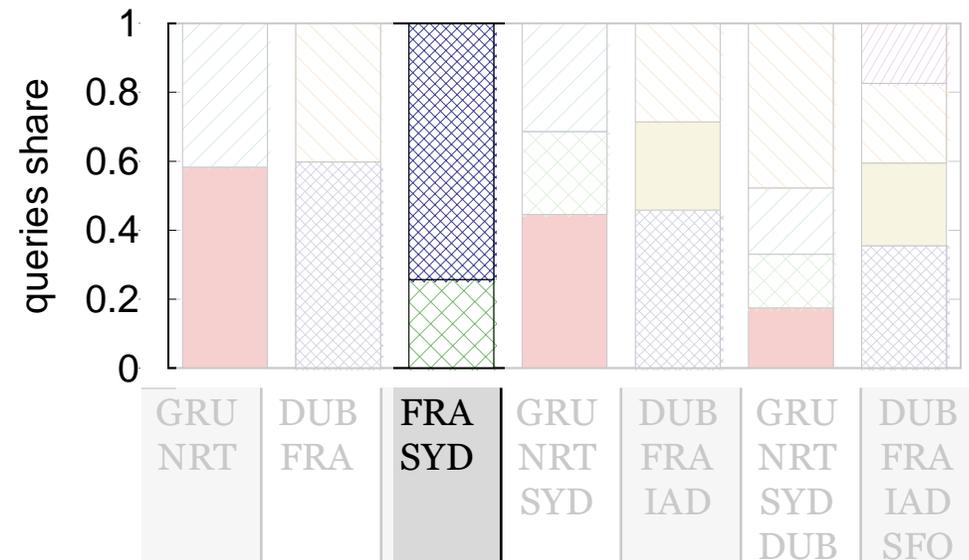
- Authoritatives with similar latency get similar number of queries



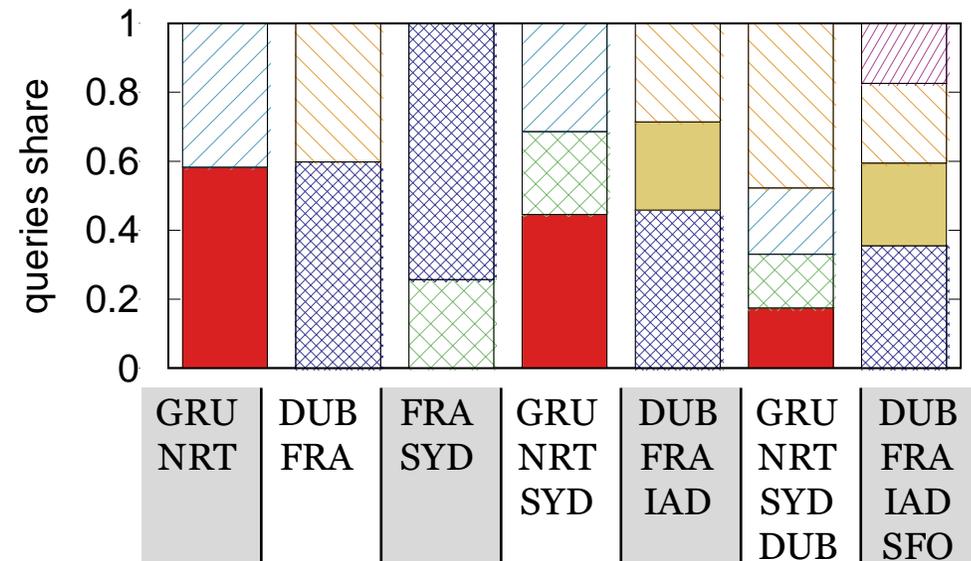
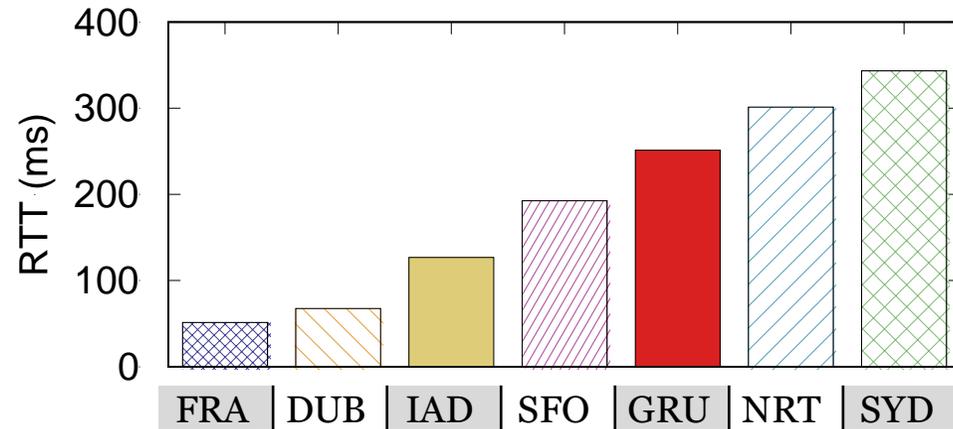
How do recursives distribute their queries over time?



- Authoritatives with similar latency get similar number of queries
- Larger difference leads to larger preference

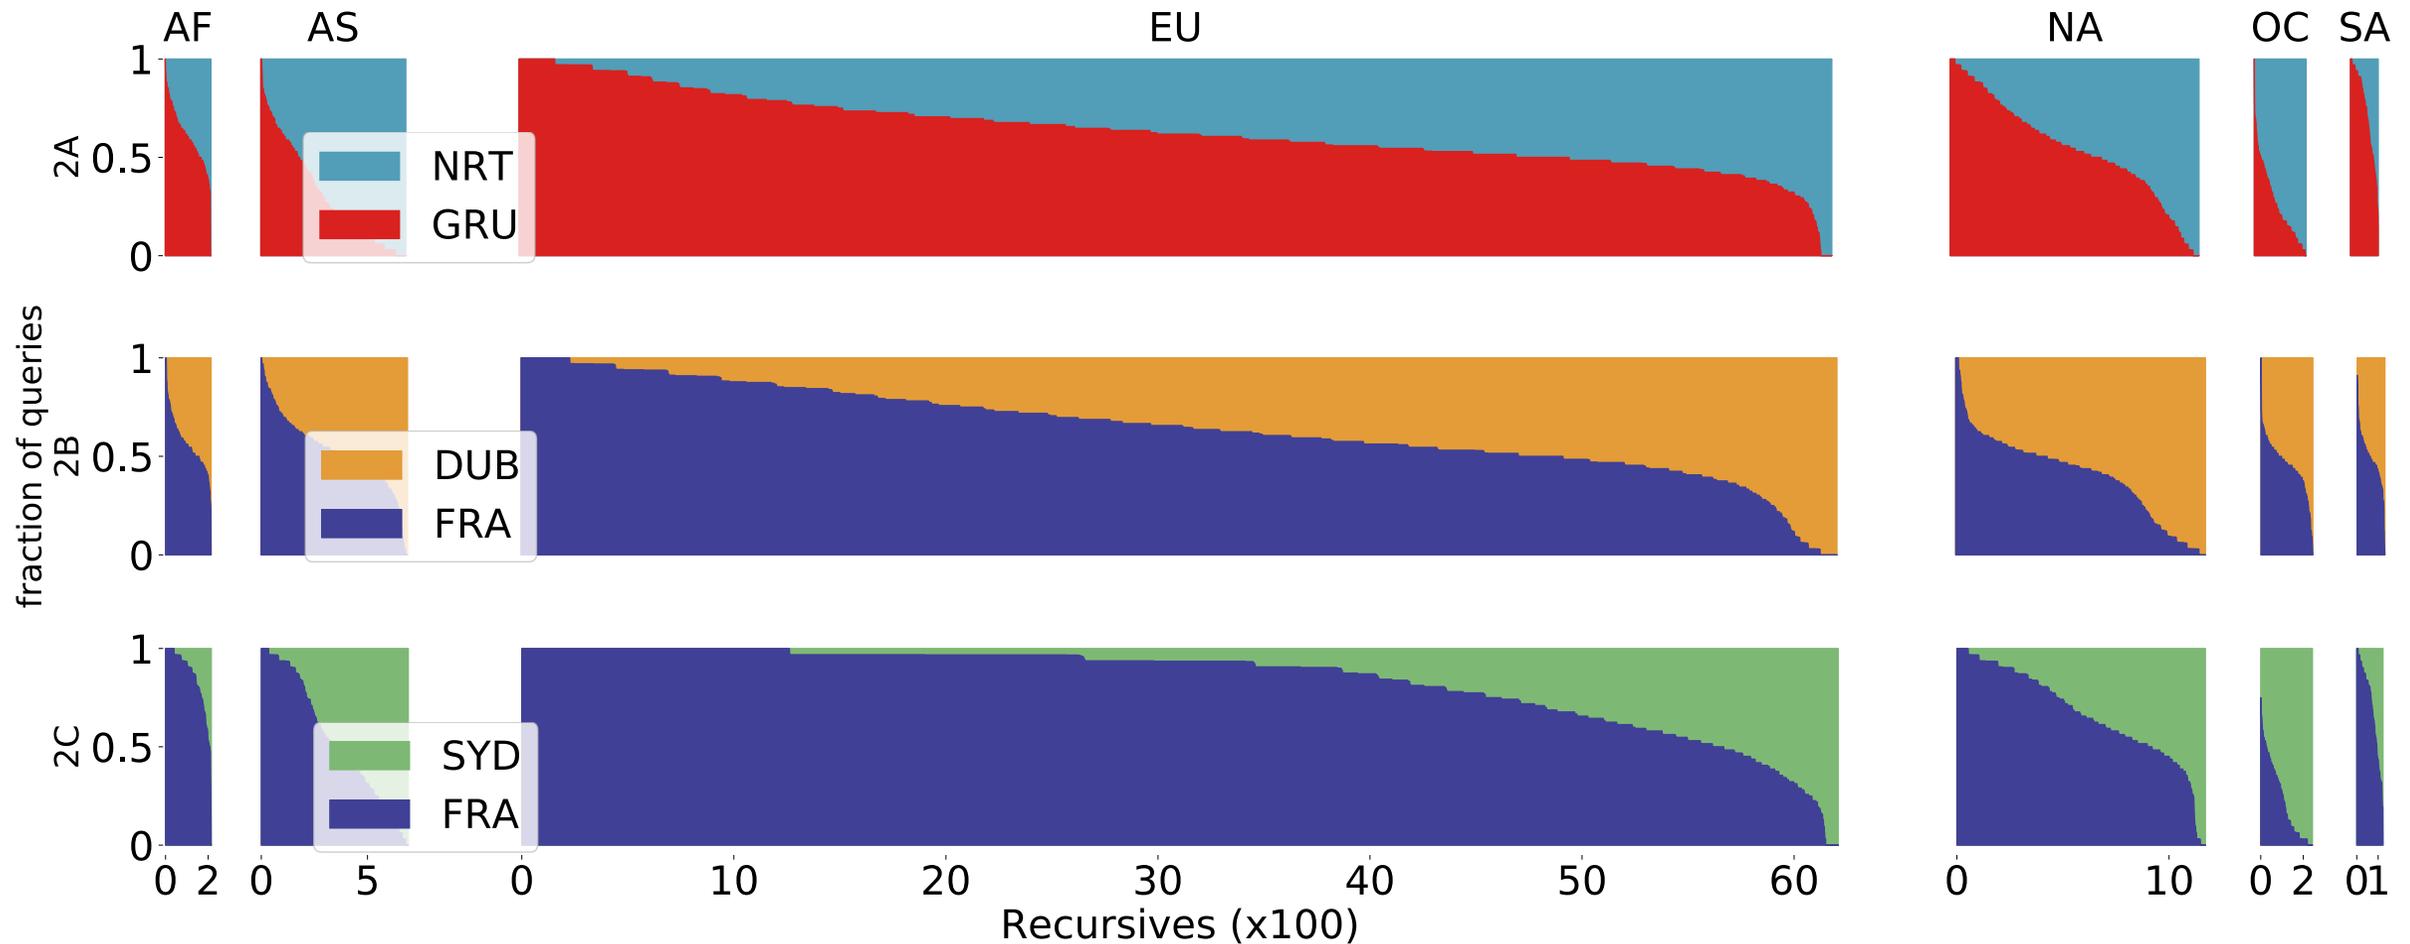


How do recursives distribute their queries over time?

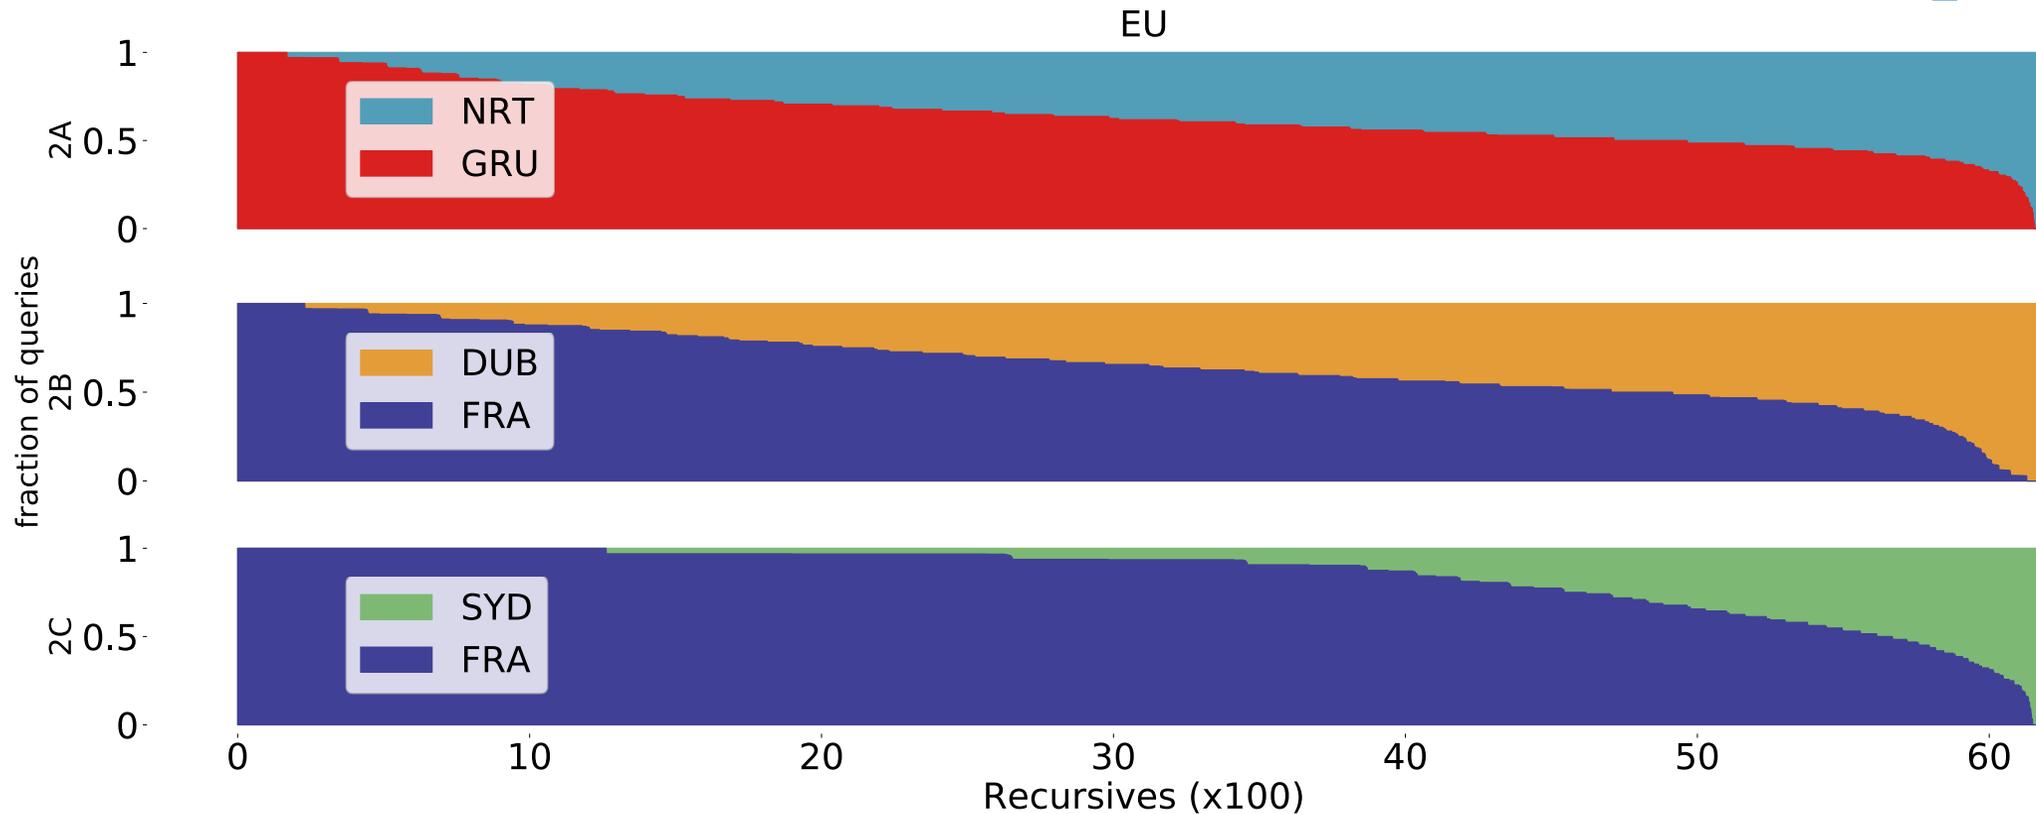


- Authoritatives with similar latency get similar number of queries
- Larger difference leads to larger preference
- Authoritatives that respond faster are in general preferred
- Confirms previous work, but now in the wild

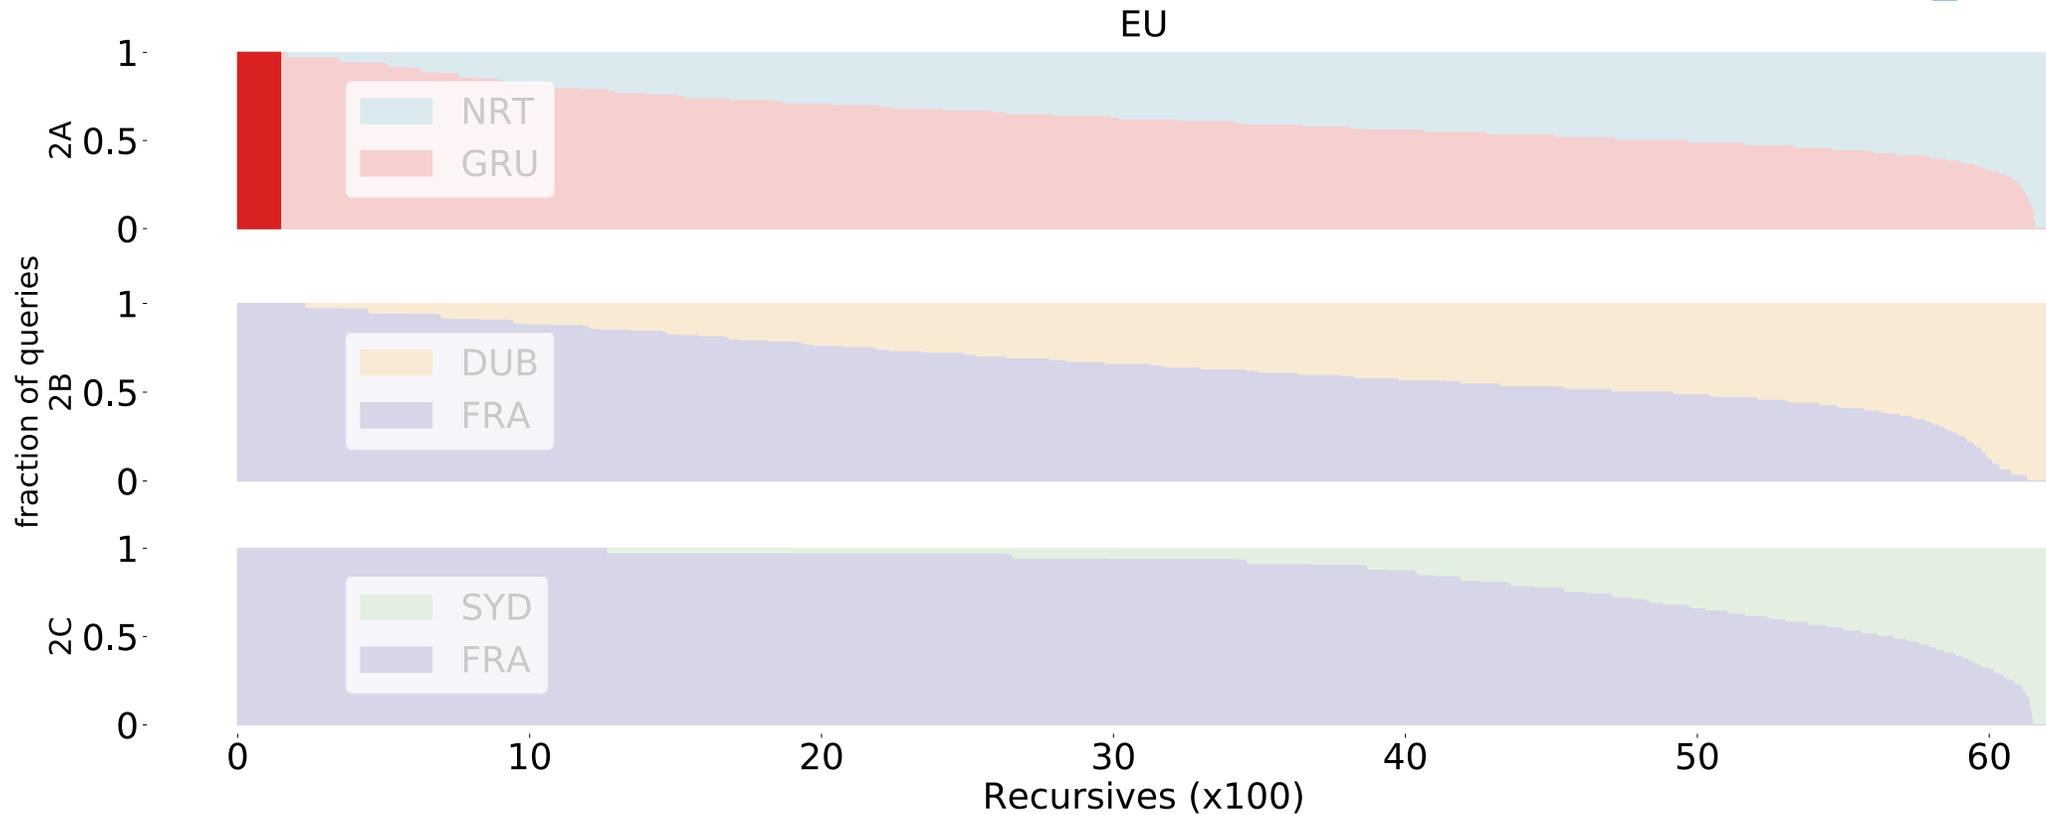
How do *individual* recursives distribute their queries?



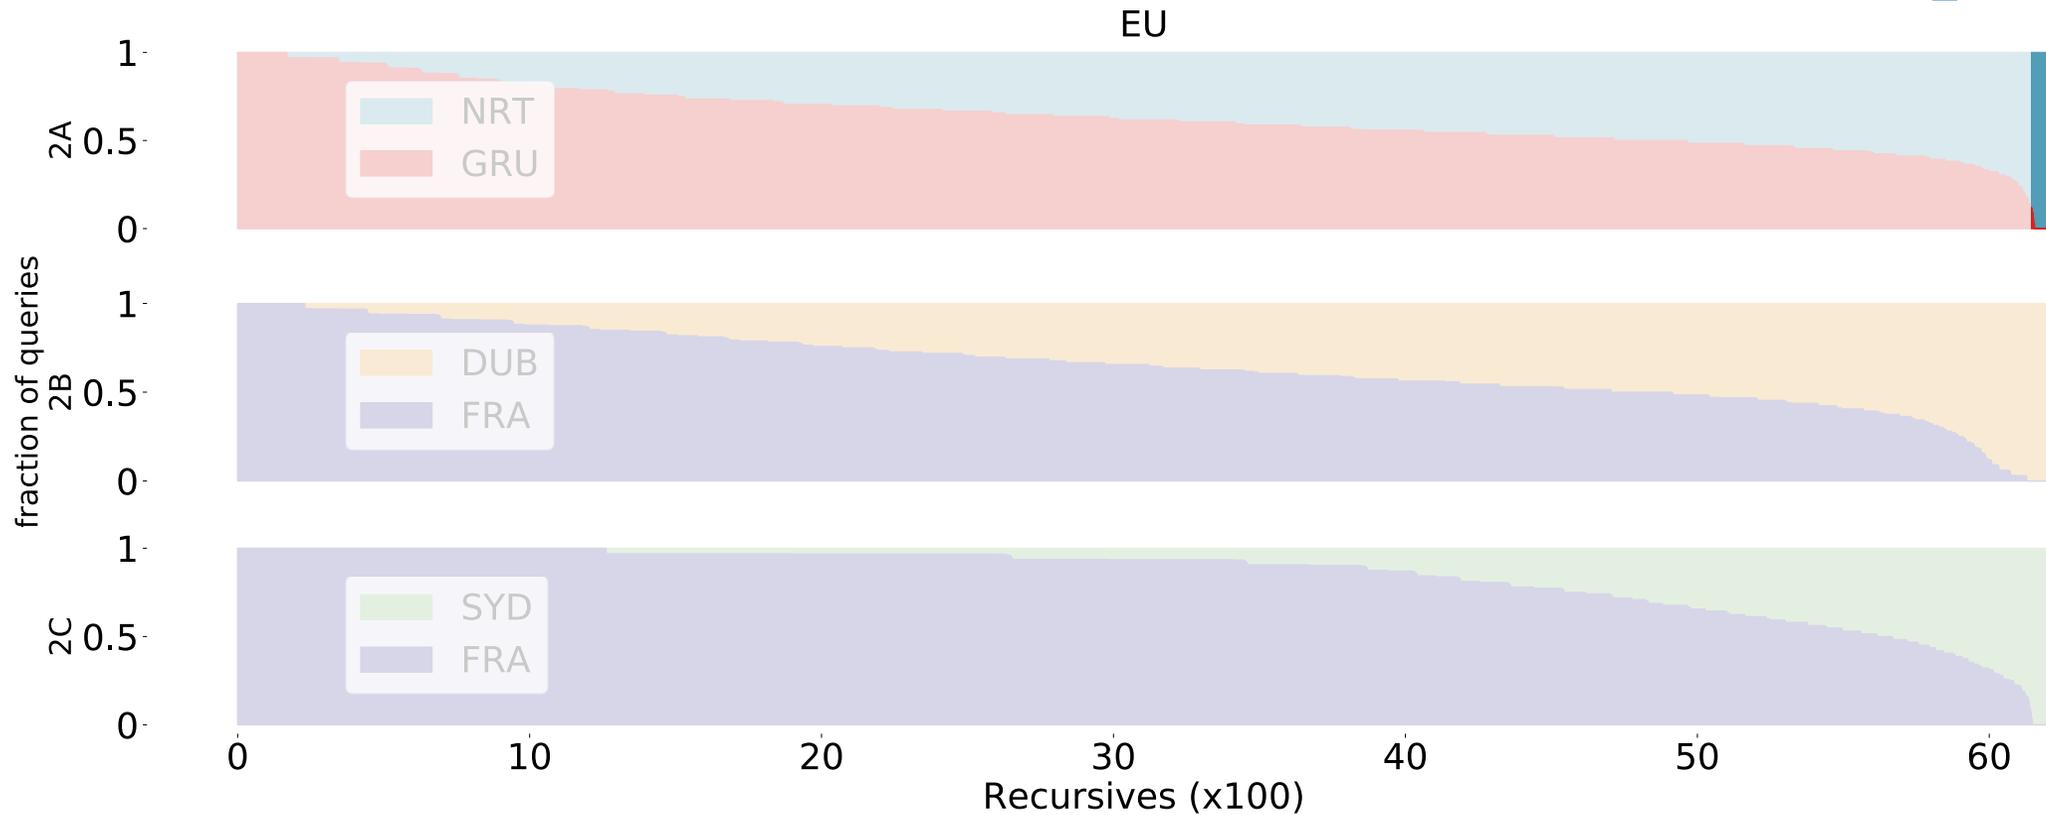
How do *individual* recursives distribute their queries?



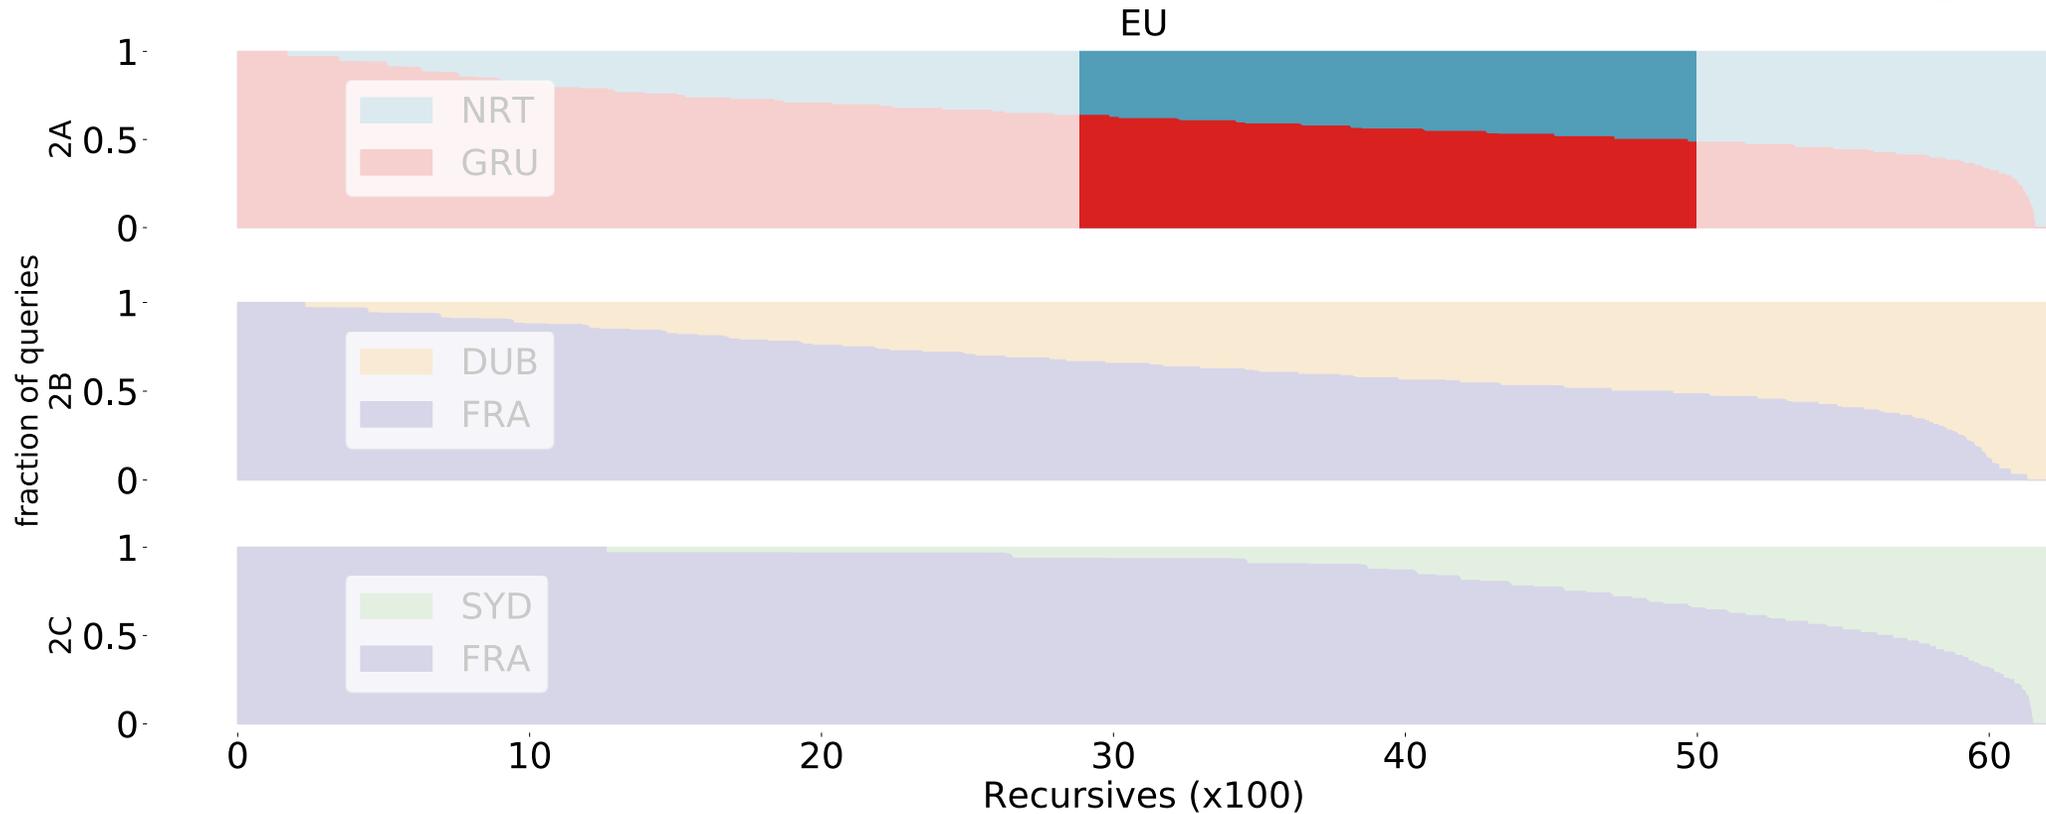
How do *individual* recursives distribute their queries?



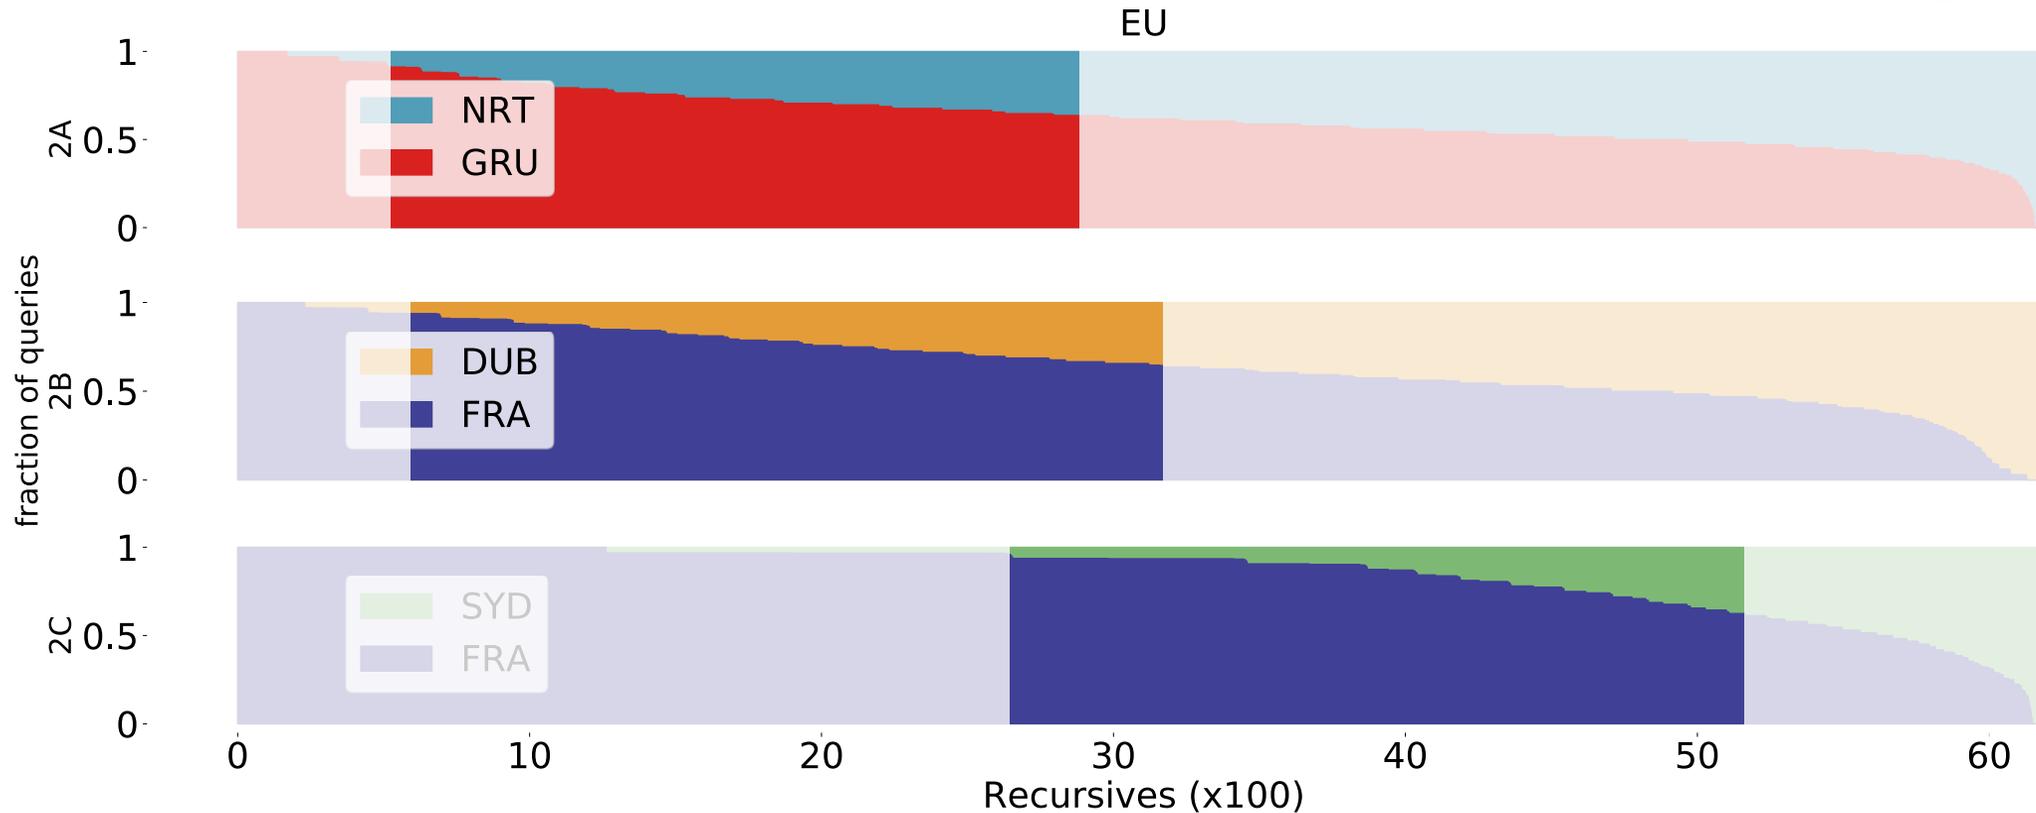
How do *individual* recursives distribute their queries?



How do *individual* recursives distribute their queries?



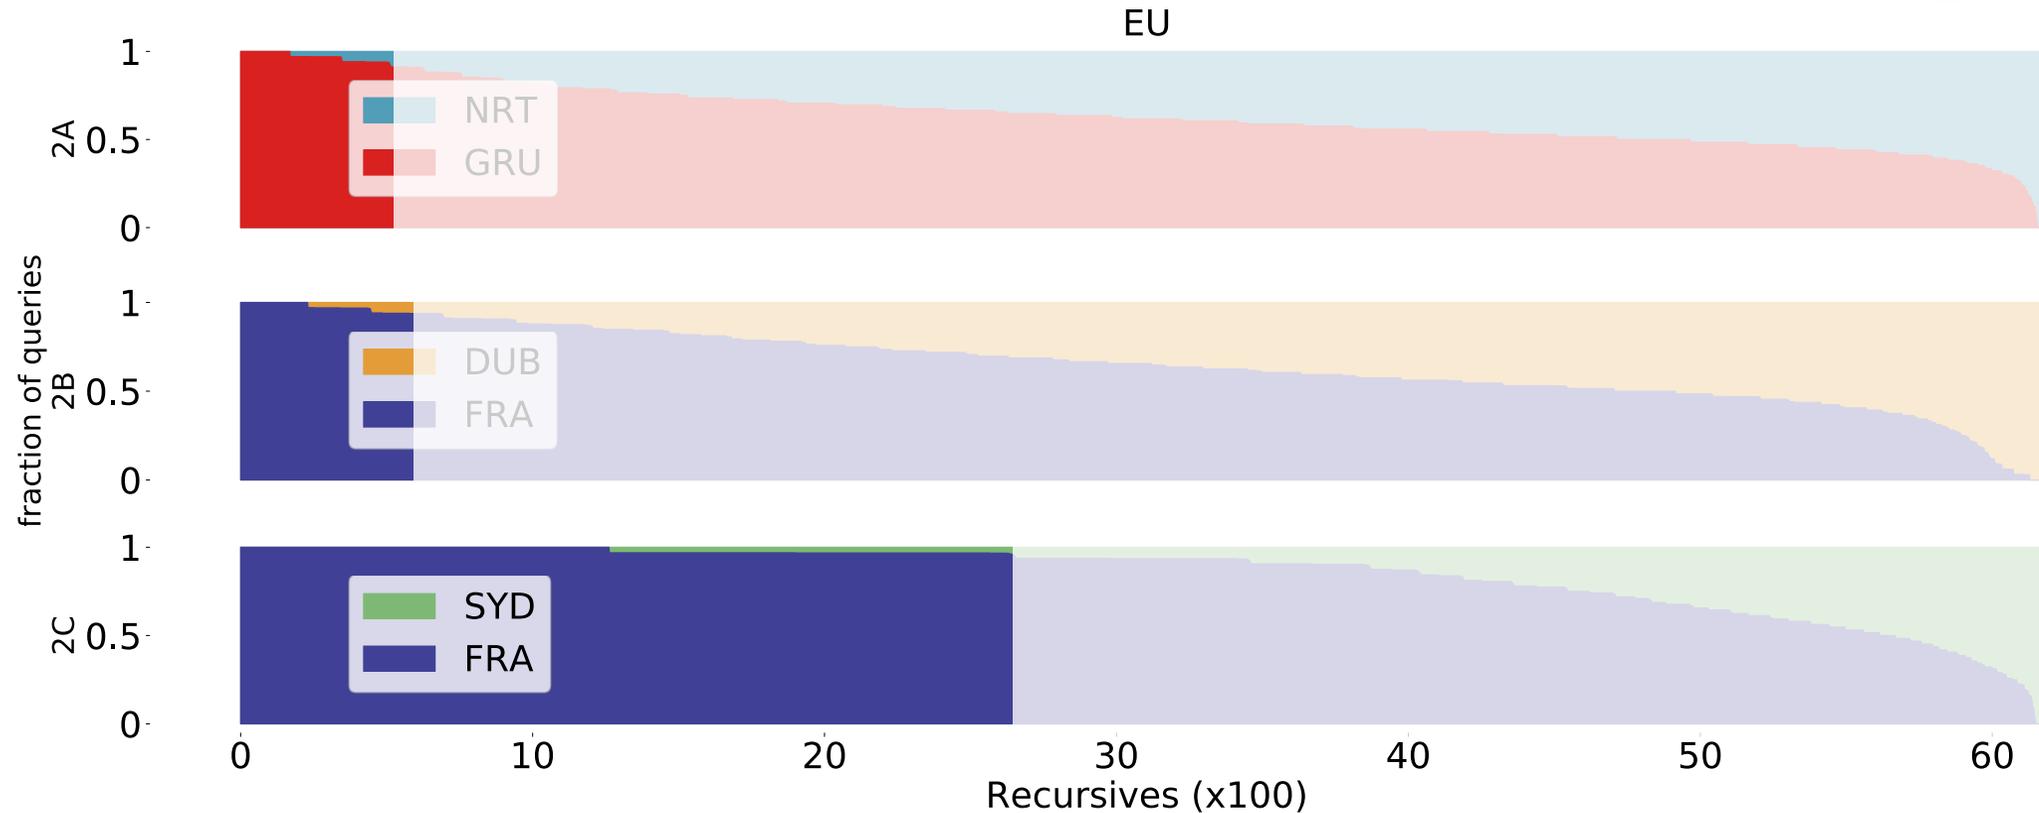
How do *individual* recursives distribute their queries?



**Up to 69% of resolvers have a weak preference
(60% to 90% of their queries to one NS)**

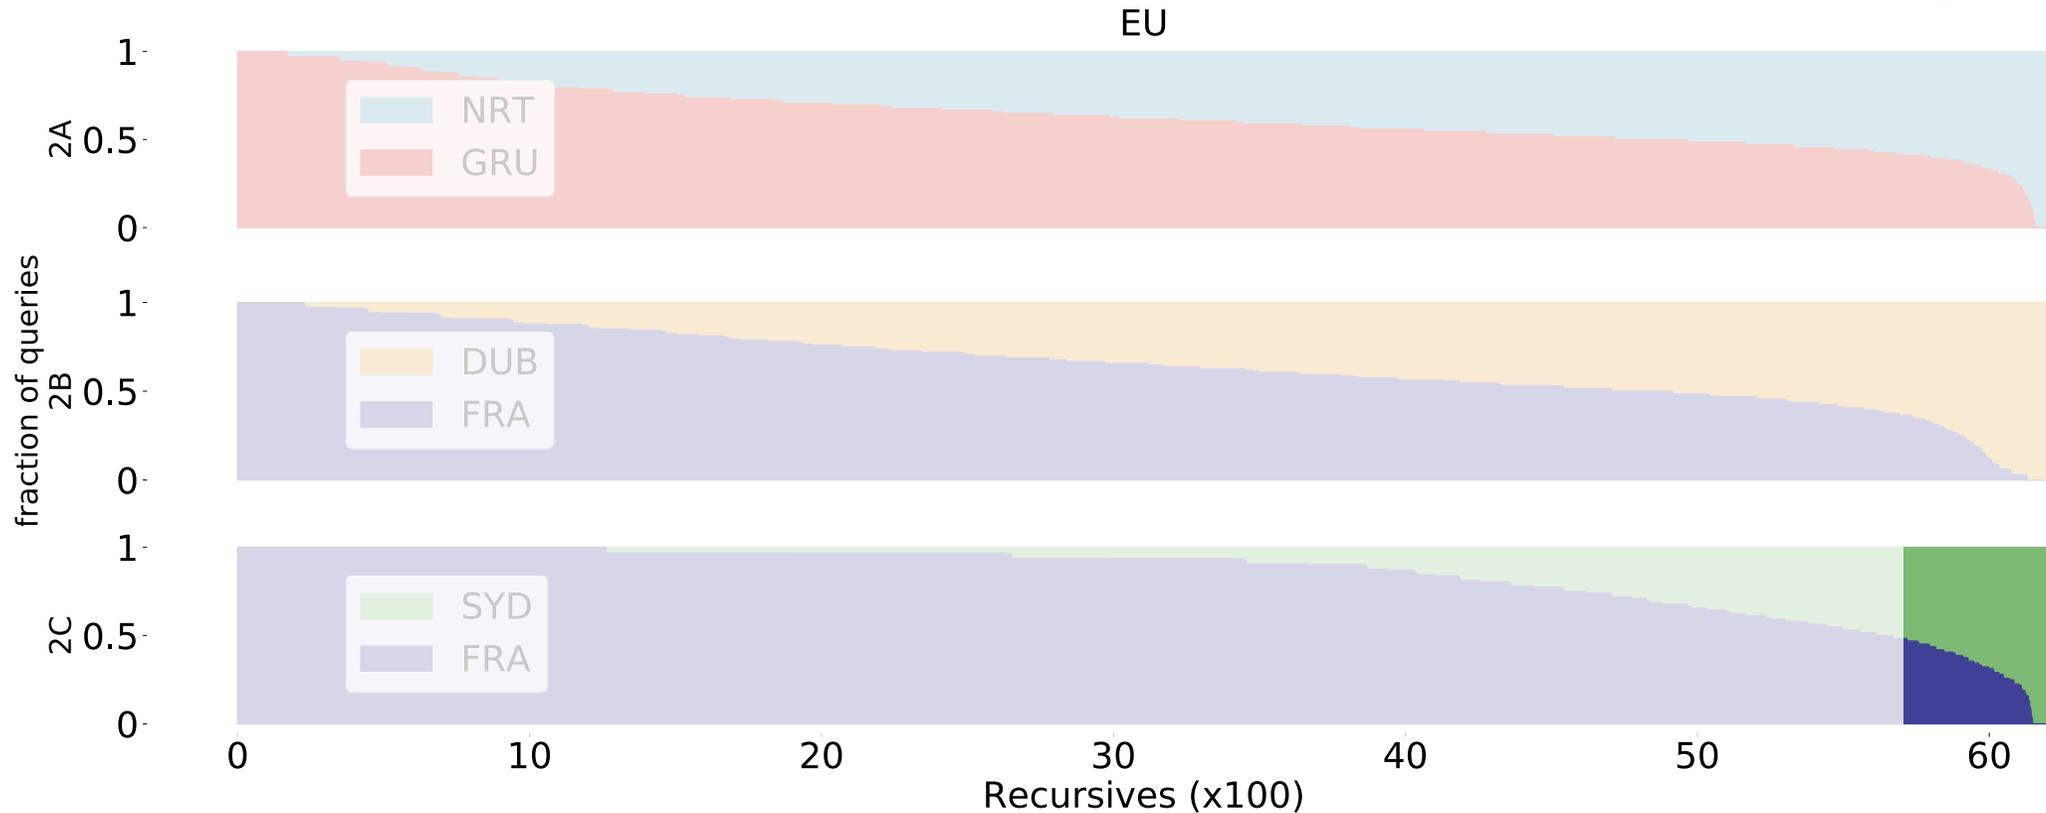


How do *individual* recursives distribute their queries?



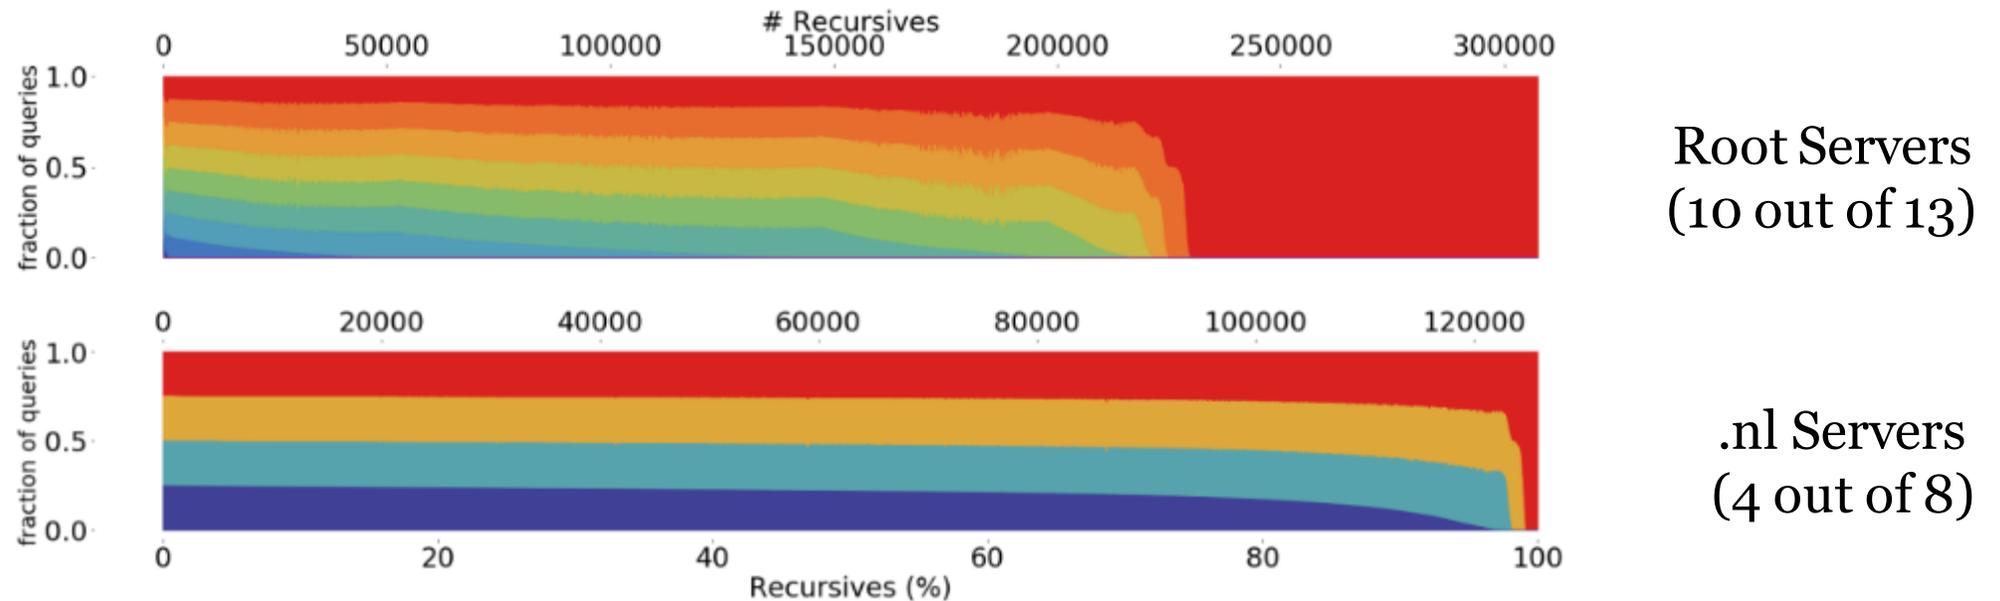
**Up to 37% of resolvers have a strong preference
(more than 90% of their queries to one NS)**

How do *individual* recursives distribute their queries?



Some resolvers always prefer the slower NS

Validation: Authoritatives in Production



- **Root: +60% query at least 6 servers**
- **.nl: +90% query at least 4 servers**
- **Overall confirms the observations from our test bed**

Measurement Summary

- Resolvers will query ALL your authoritative servers
 - Distribution is inversely proportional with the median RTT
 - Recursives prefer faster responding authoritatives
 - But they also query slower authoritatives from time to time
-
- Additional findings:
 - Lower RTT becomes more relevant if competing NSes are closer (<150 ms)
 - Stronger preference when querying more frequent (< 10min interval)

Recommendations for DNS Operators

- The slowest authoritative limits the response time of a DNS service
- **Recommendation:**
 - Use anycast on *all* your name servers
 - Anycast sites need to be well connected with good peering

→ Based on this work .nl is replacing unicast NSes with anycast



Data Sets

All data sets (but one) available:

<https://ant.isi.edu/datasets/dns/index.html#recursives>



Data Sets

All data sets (but one) available:

<https://ant.isi.edu/datasets/dns/index.html#recursives>

Yes, the majority of resolvers query every authoritative

Paper available at: <https://conferences.sigcomm.org/imc/2017/papers/imc17-final12.pdf>

Questions?

Giovane C. M. Moura

Email: giovane.moura@sidn.nl

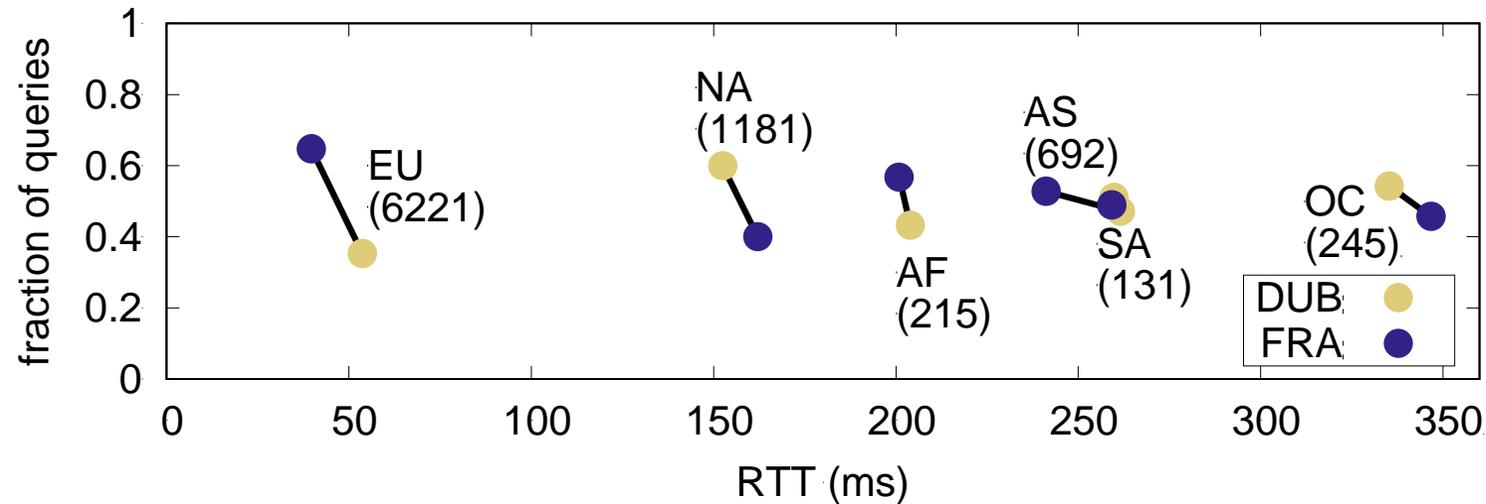
twitter: @giomourasec



Additional Slides

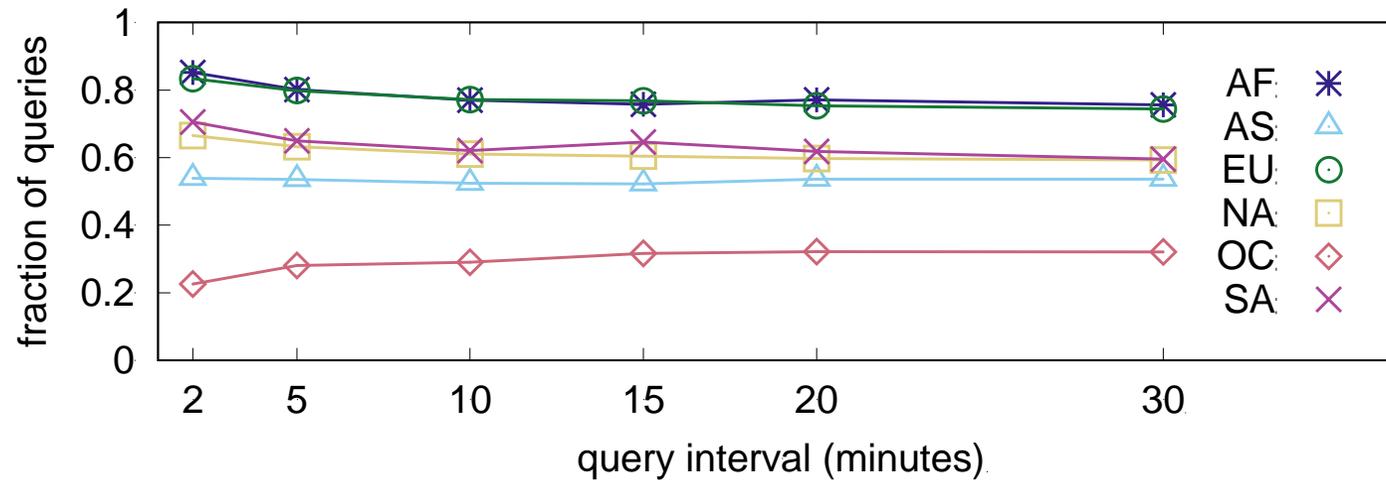


Does preference change for distant recursives?



- VPs in EU reach Frankfurt 13 ms faster than Dublin
 - Thus, they clearly prefer Frankfurt
 - VPs in Asia reach Frankfurt 20 ms faster, but distribute their queries almost equally
- Lower RTT becomes more relevant if competing authoritatives are closer to the recursive

How does query frequency affect the results?



- A higher query frequency leads to a stronger preference
- However, preference persists even after the default timeout of resolvers like Bind and Unbound