Recursives in the Wild:
Engineering Authoritative DNS Servers

RON++ Meeting | 2017-12-15 | Utrecht

Moritz Müller[1,2], Giovane C. M. Moura[1],

Ricardo de O. Schmidt[1,2], John Heidemann[3]

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

UNIVERSITY OF TWENTE.

Introduction

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Introduction

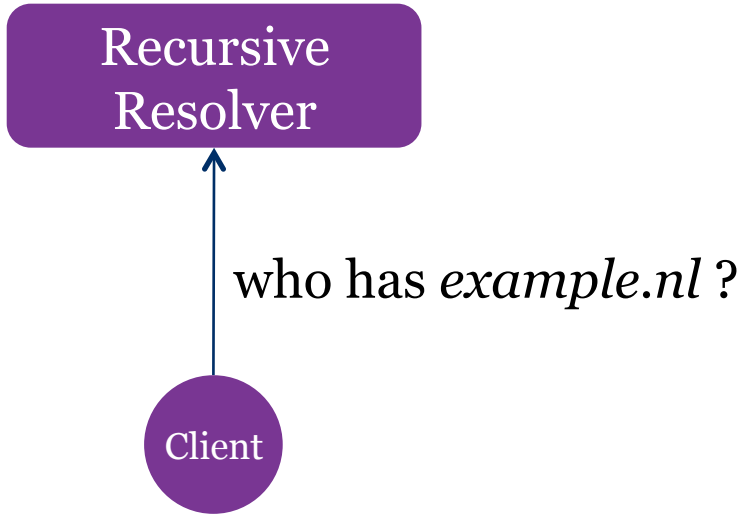

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Introduction

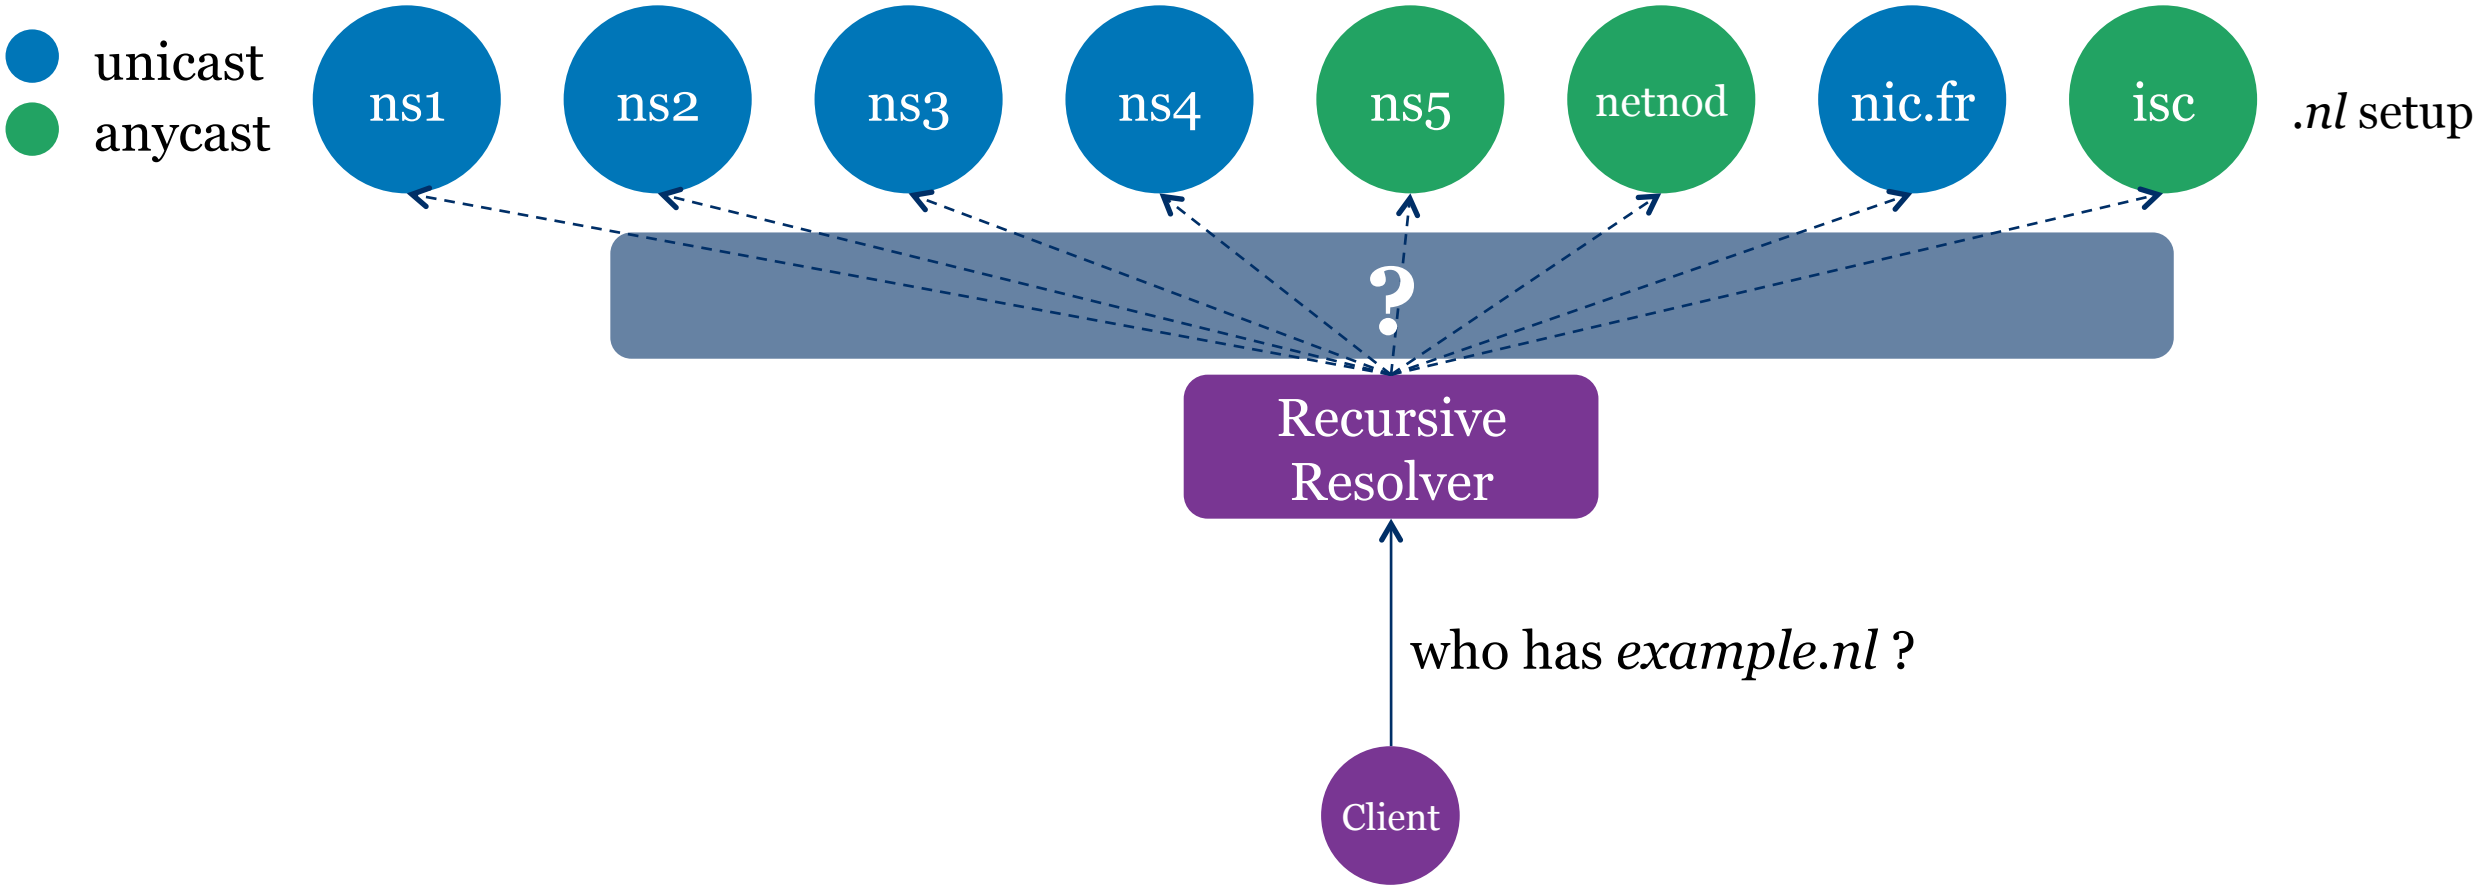

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Introduction

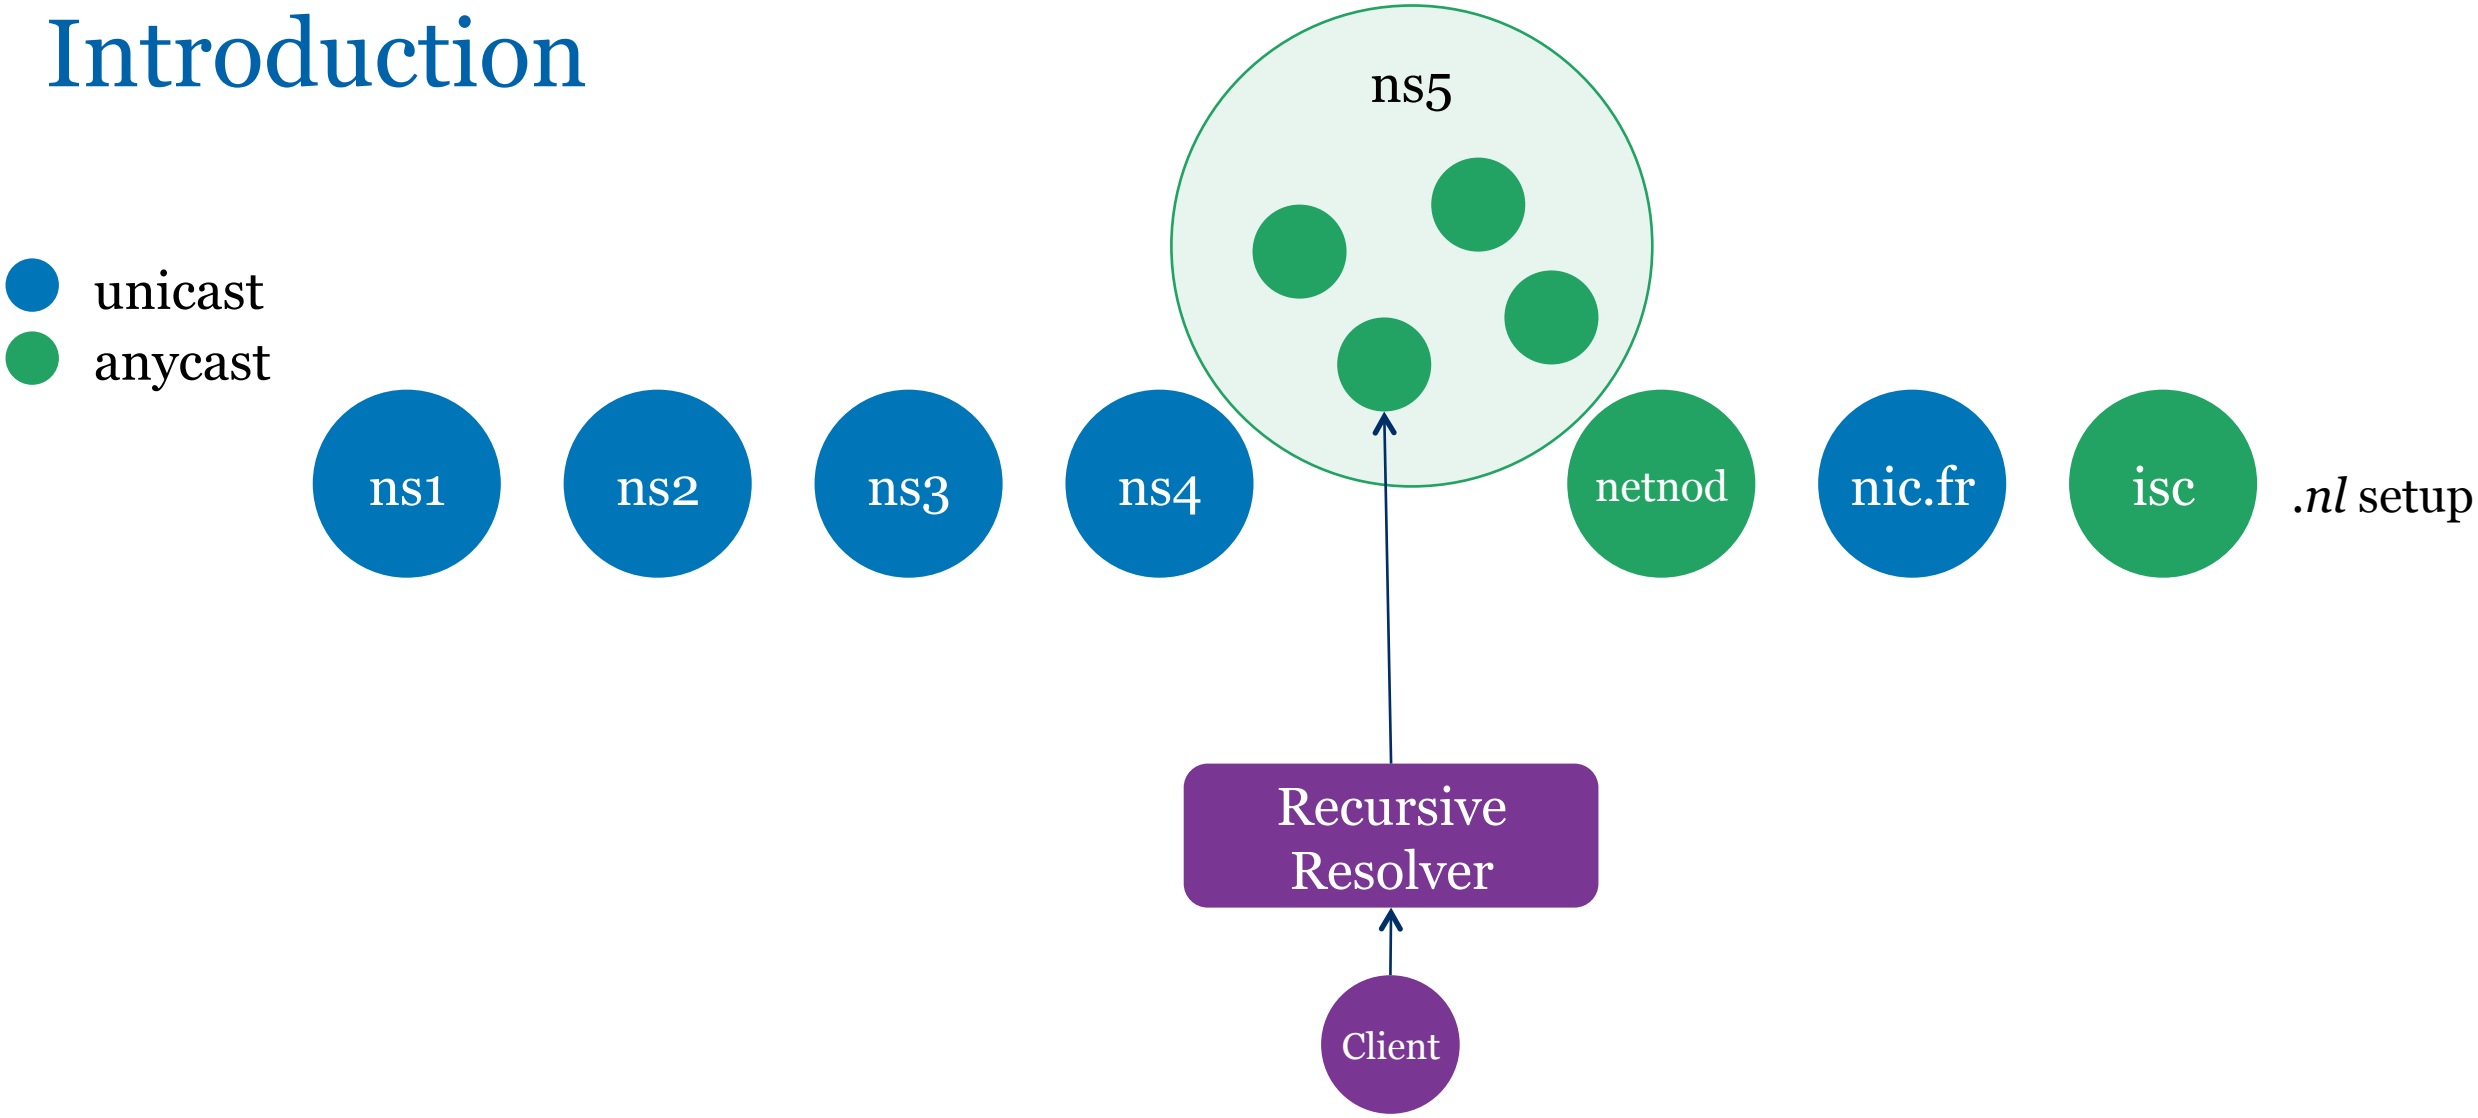

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Introduction

area relative to the number of *sites*

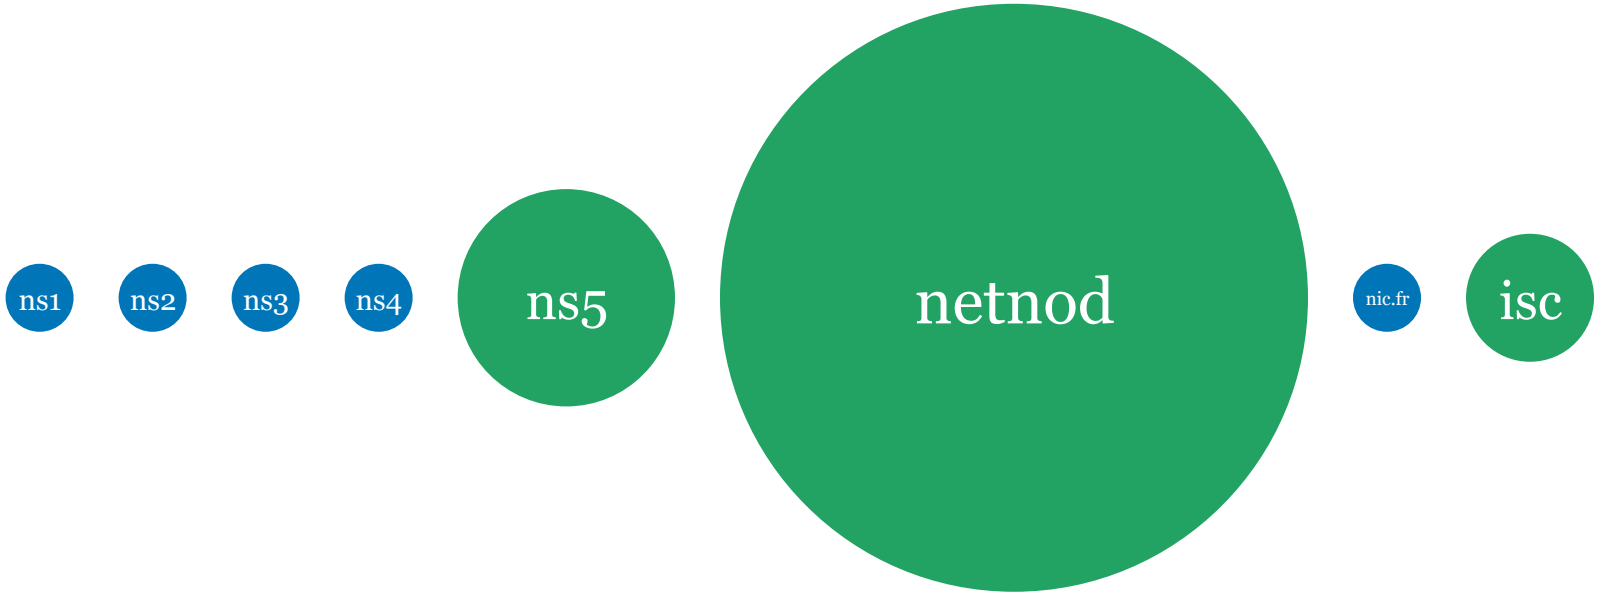

Introduction

area relative to the number of *sites*

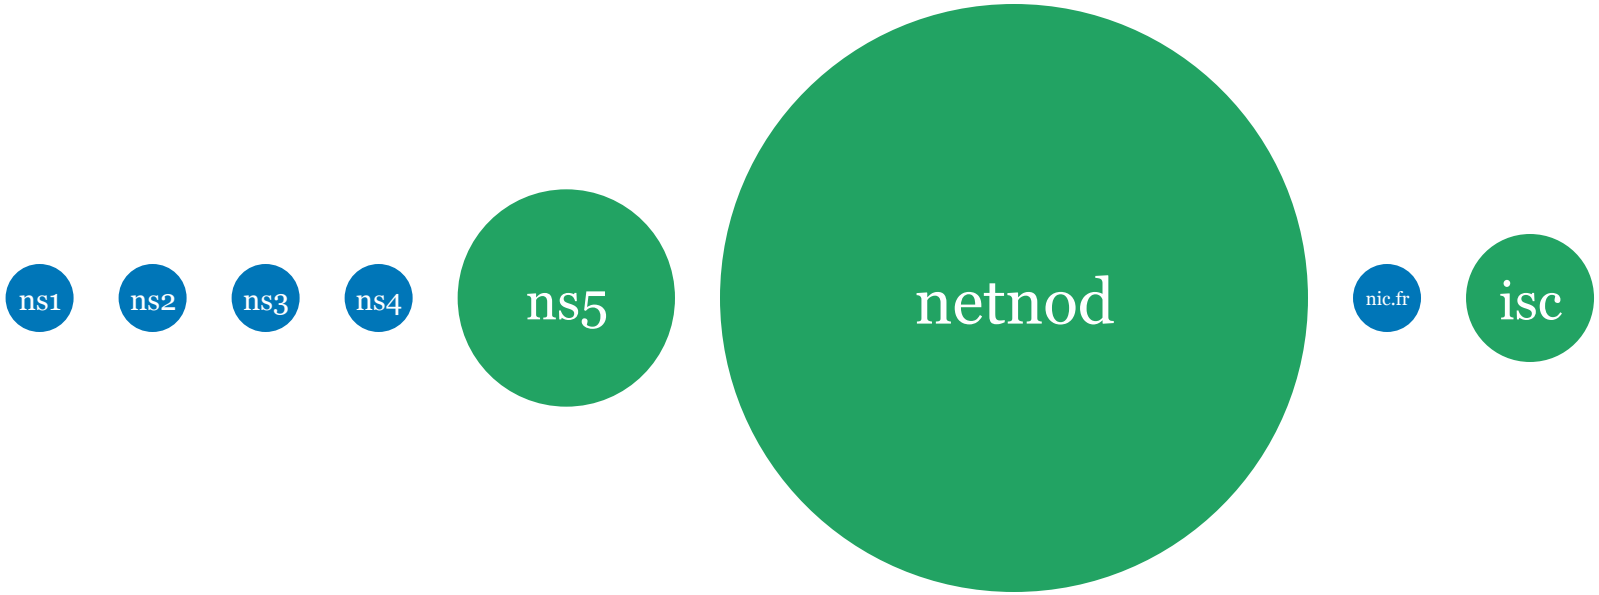

area relative to the number of *queries*

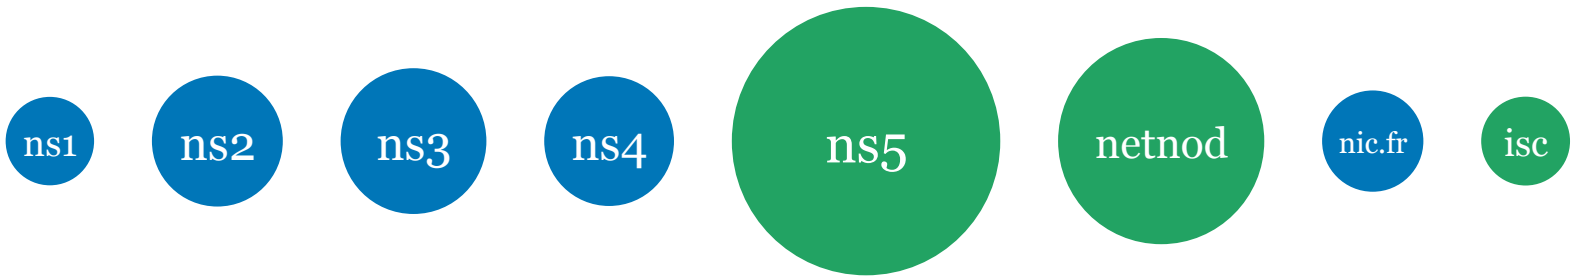

Introduction

area relative to the number of *sites*

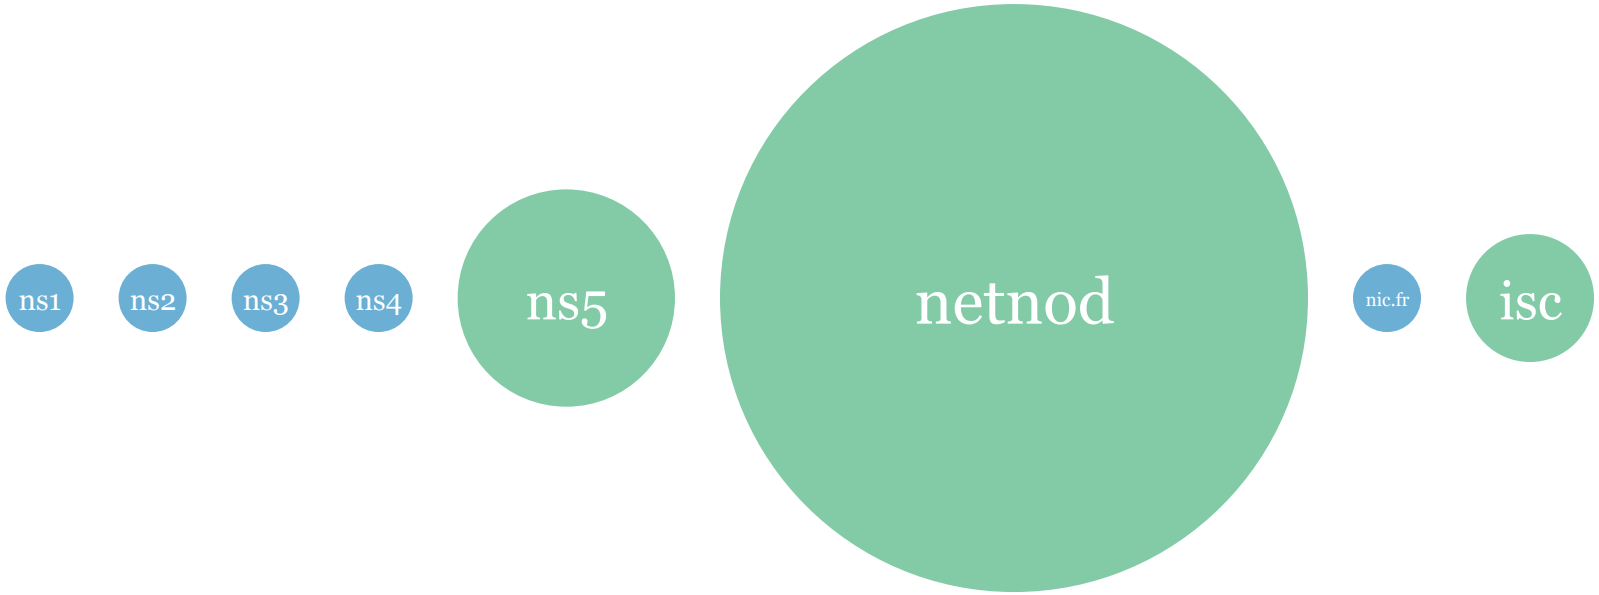

area relative to the number of *queries*

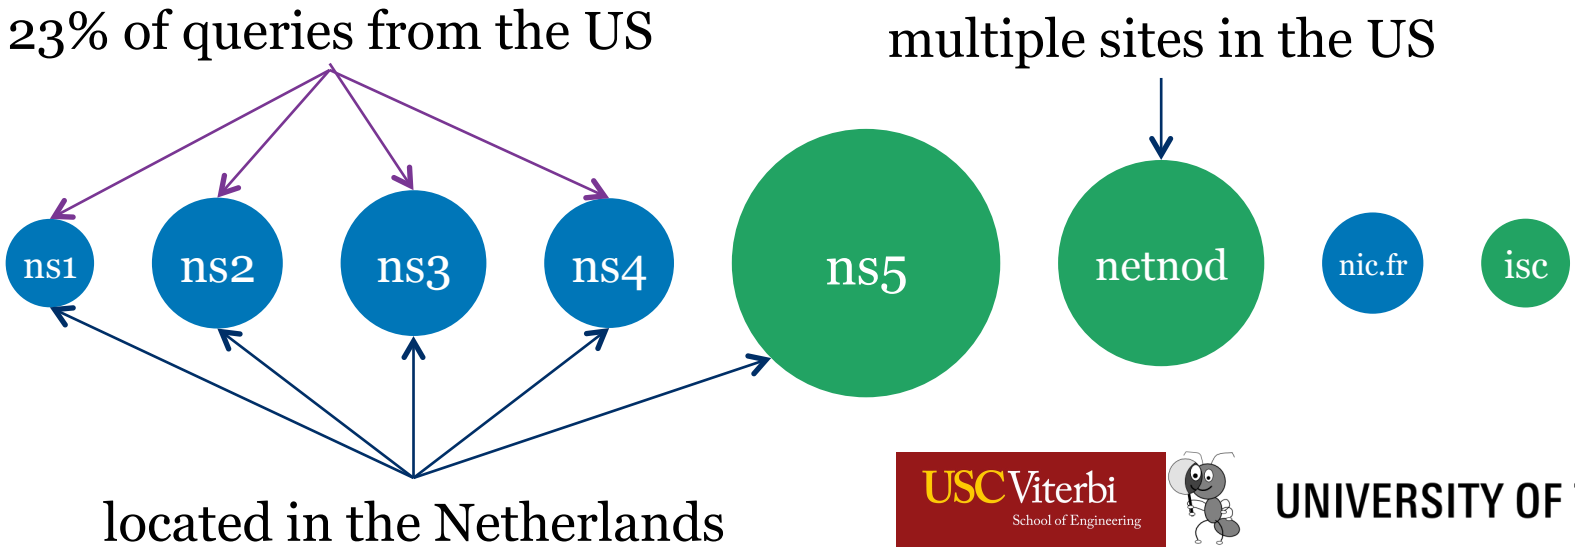

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

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 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), 80–86.

Measurement Design

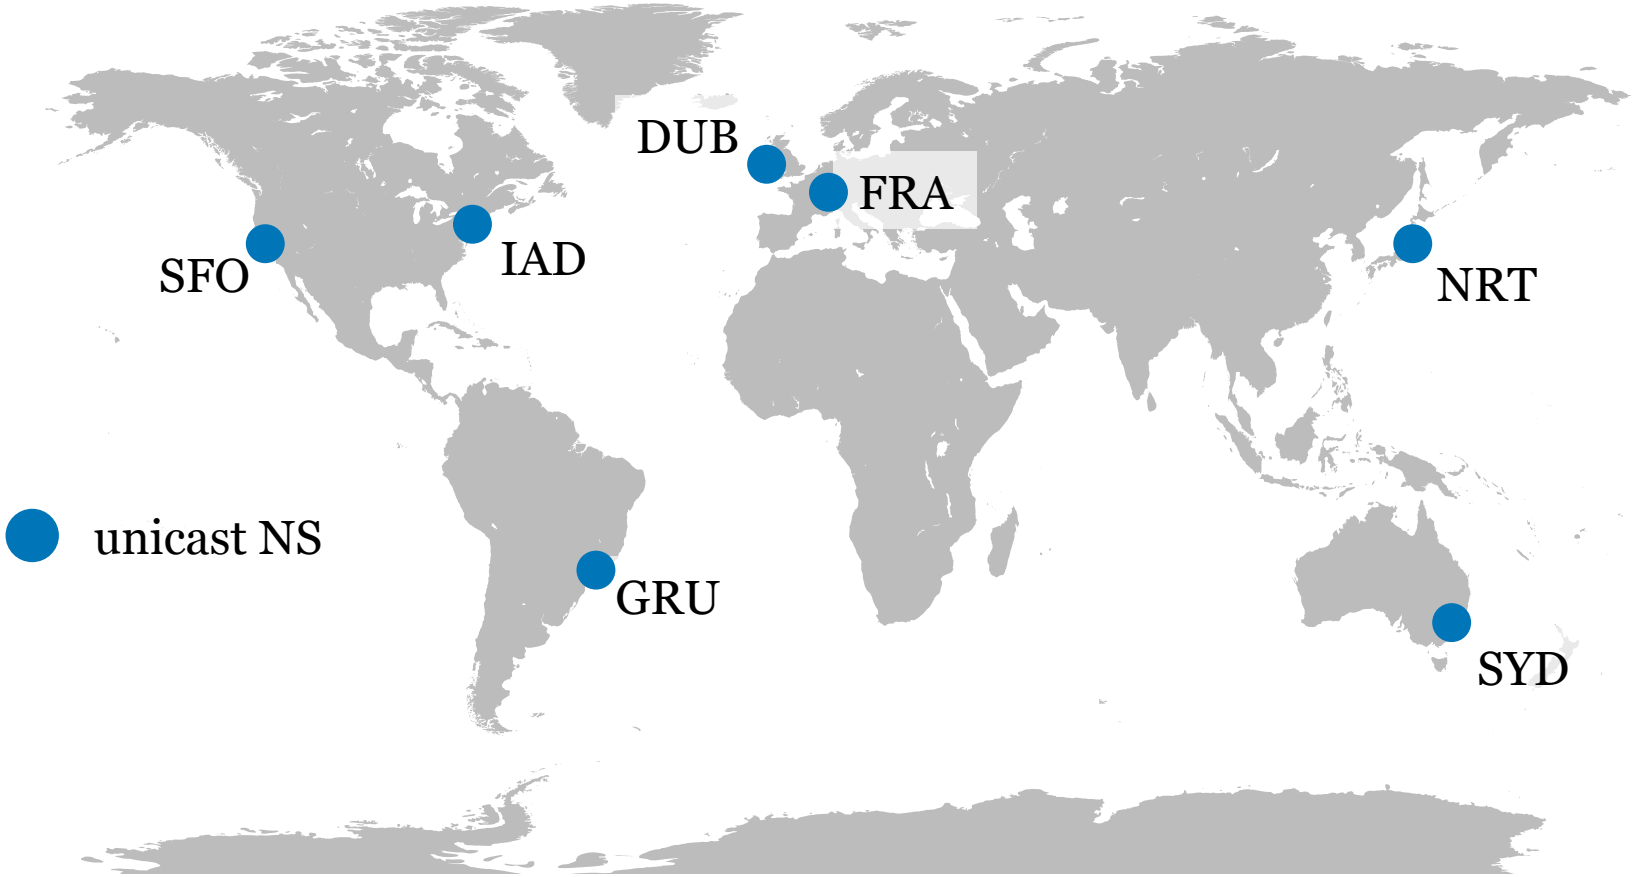

Setups:

GRU+NRT

DUB+FRA

FRA+SYD

GRU+NRT+SYD

DUB+FRA+IAD

DUB+GRU+NRT+SYD

DUB+FRA+IAD+SFO

IPv4 only (for now)

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Measurement Design

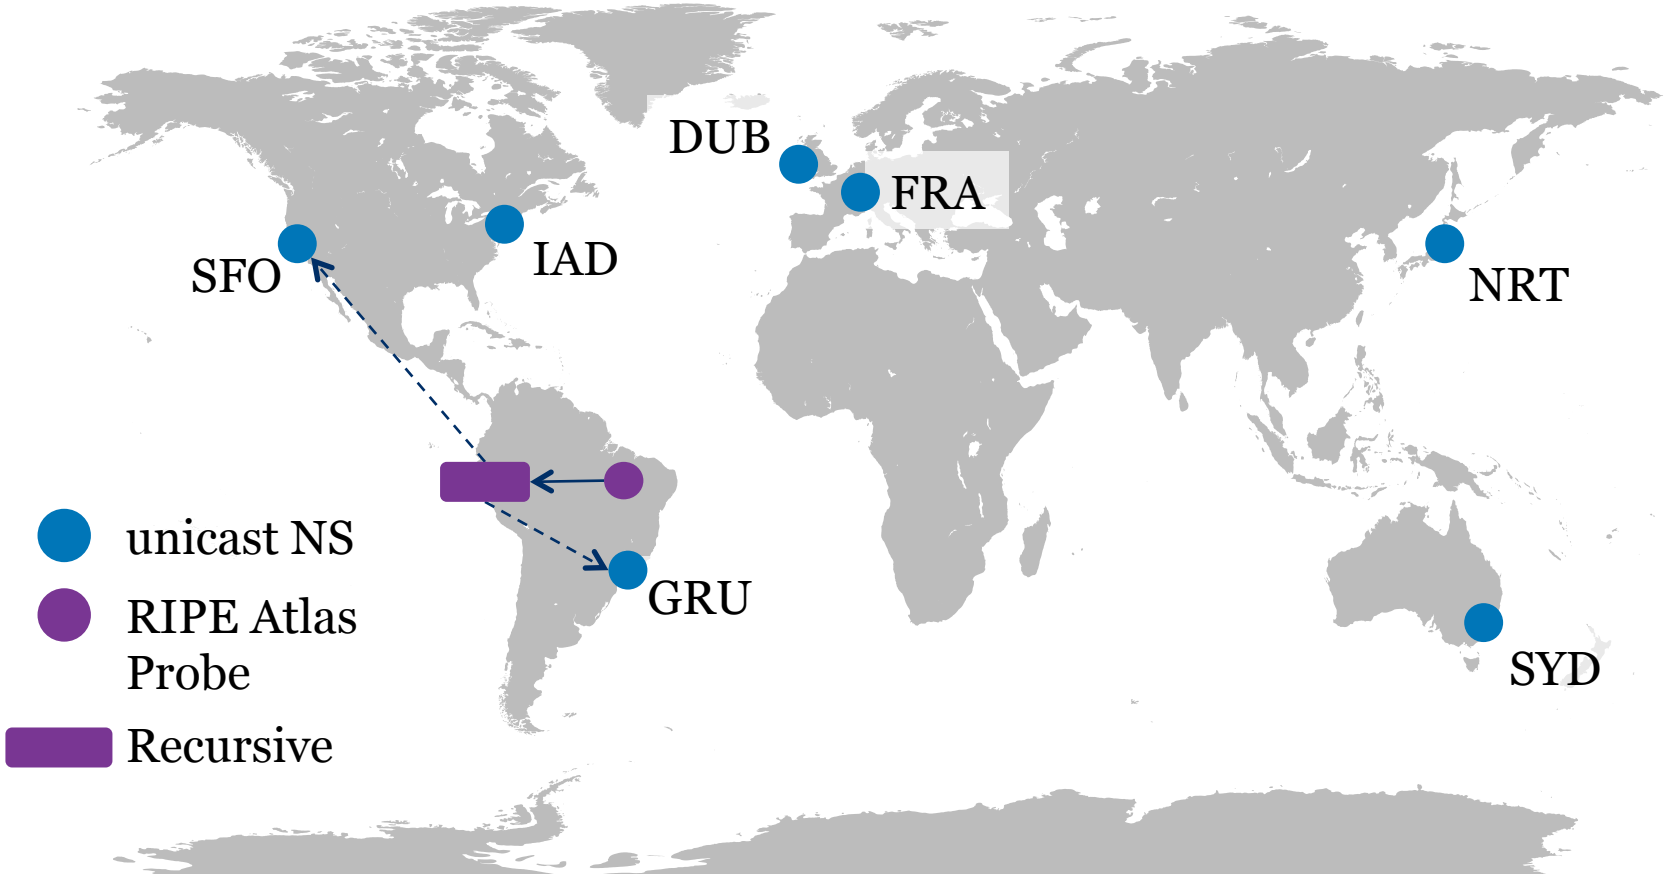

Setups:

GRU+NRT

DUB+FRA

FRA+SYD

GRU+NRT+SYD

DUB+FRA+IAD

DUB+GRU+NRT+SYD

DUB+FRA+IAD+SFO

IPv4 only (for now)

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Do recursives query all authoritatives?

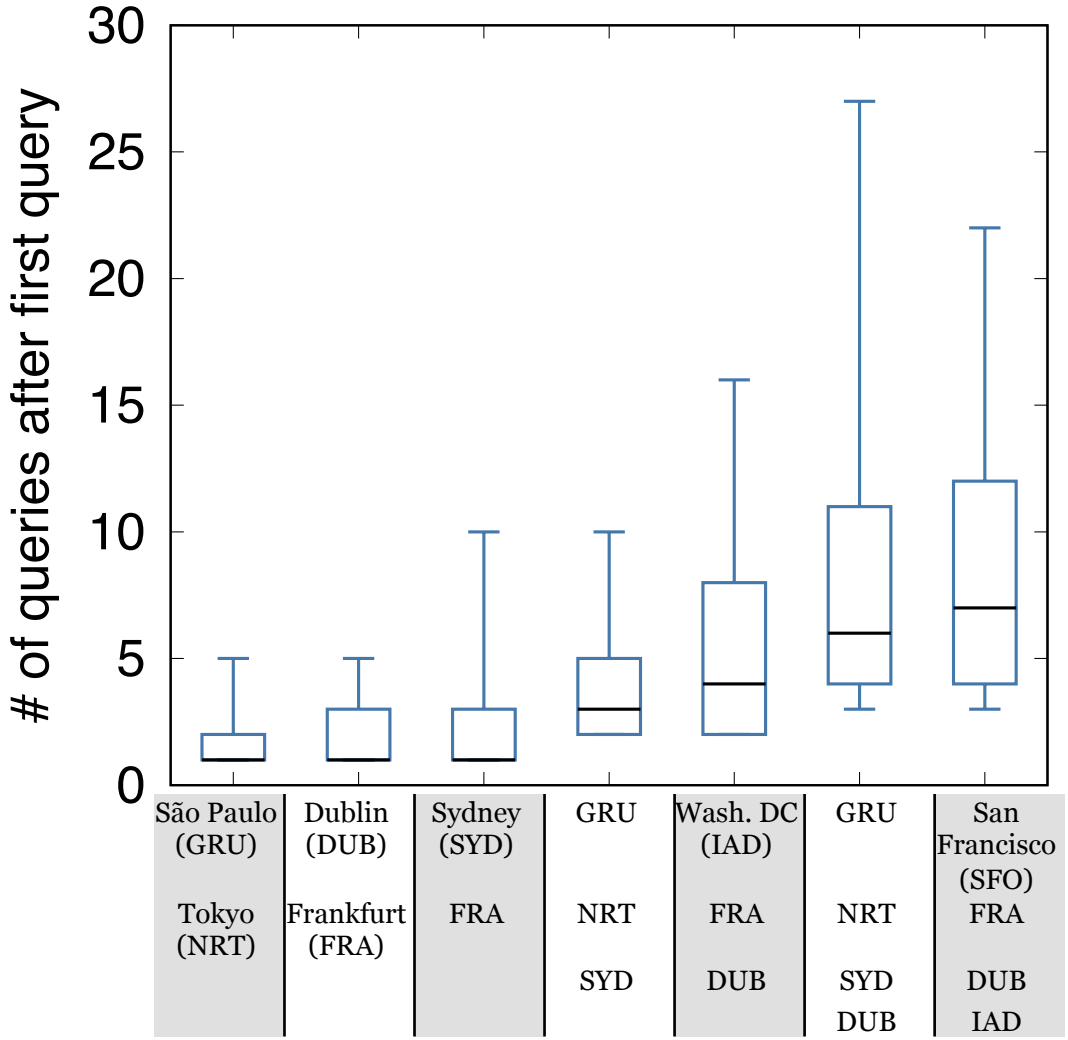

Yes, the majority of resolvers query every authoritative

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do recursives distribute their queries over time?

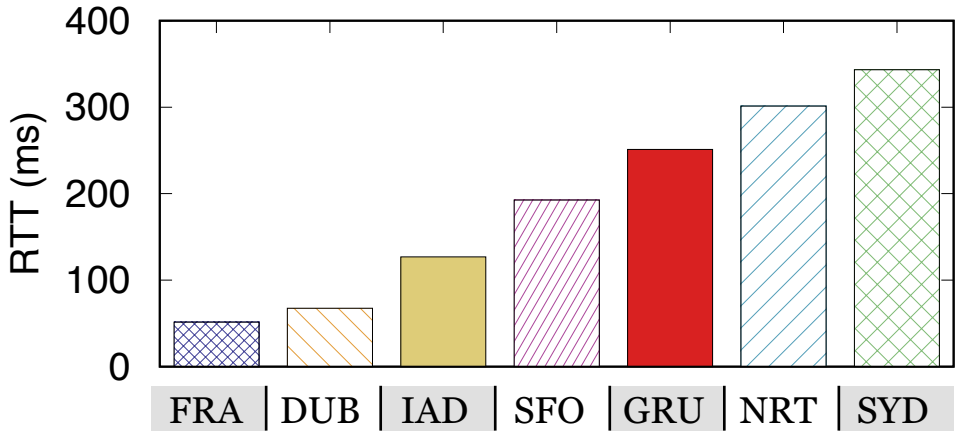

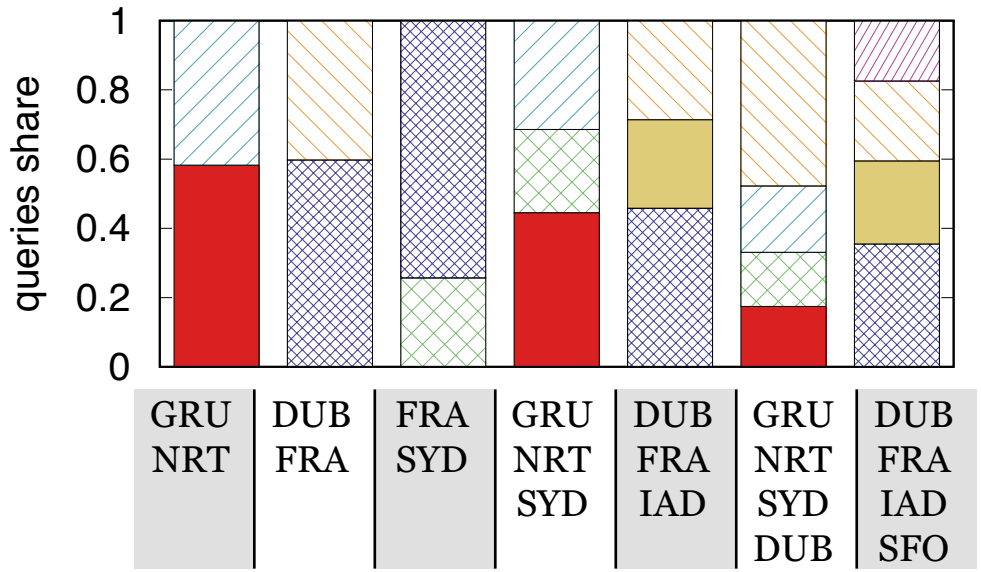

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do recursives distribute their queries over time?

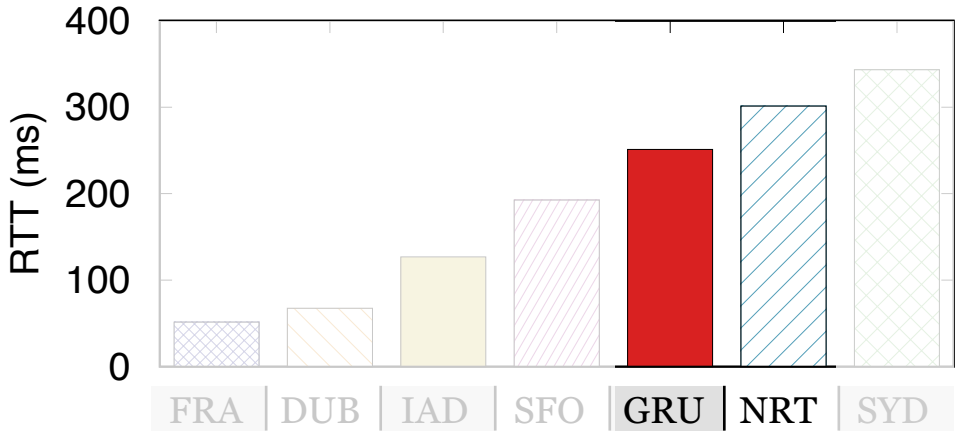

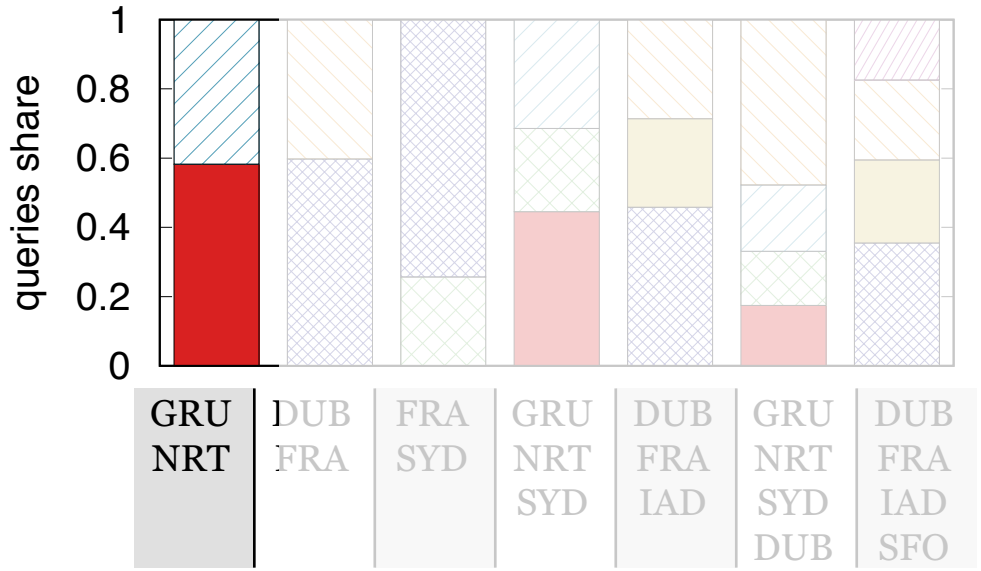

- Authoritatives with similar latency get similar number of queries

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do recursives distribute their queries over time?

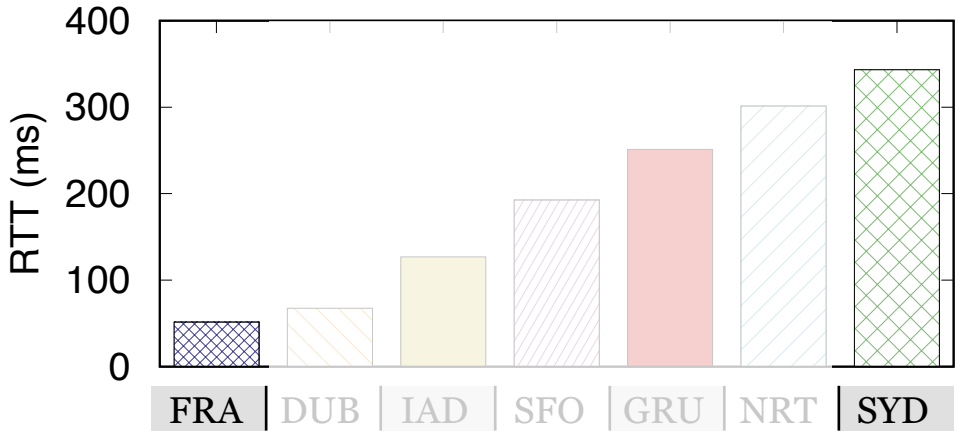

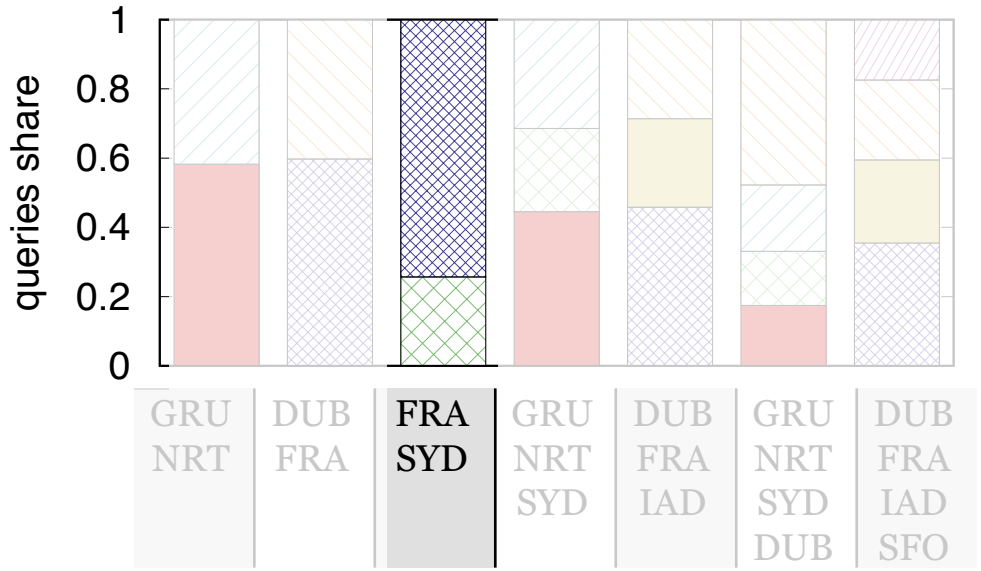

- Authoritatives with similar latency get similar number of queries
- Larger difference leads to larger preference

USC Viterbi School of Engineering
UNIVERSITY OF TWENTE.
SIDN LABS

How do recursives distribute their queries over time?

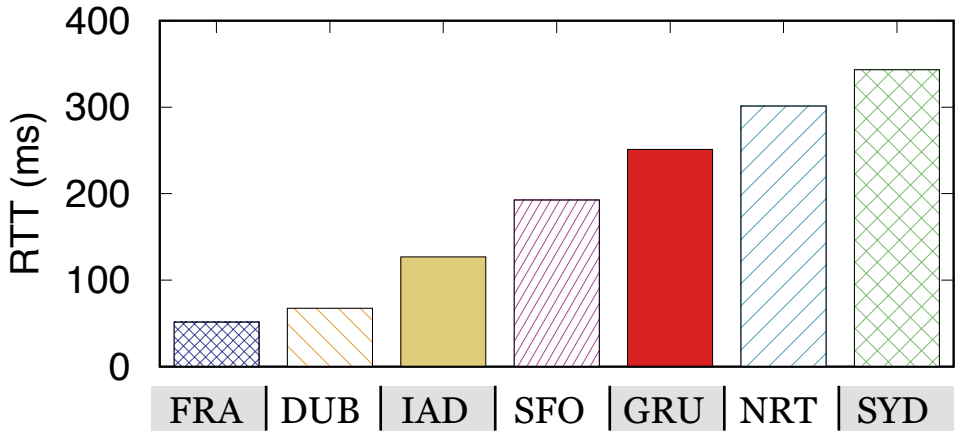

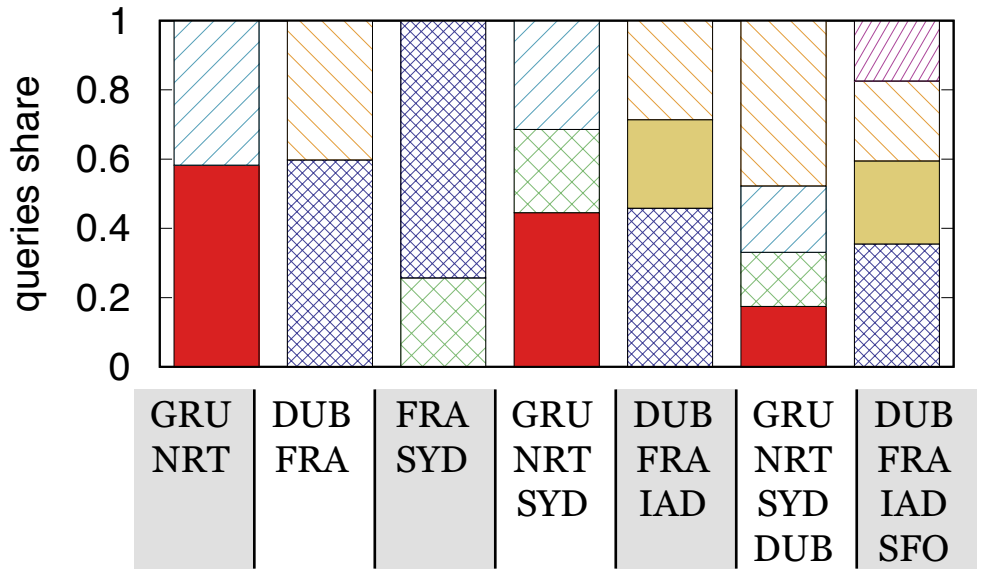

- Authoritatives with similar latency get similar number of queries
- Larger difference leads to larger preference
- Authoritatives that respond faster are in general preferred
- Confirms previous work, but now in the wild

USC Viterbi School of Engineering UNIVERSITY OF TWENTE. SIDN LABS

How do *individual* recursives distribute their queries?

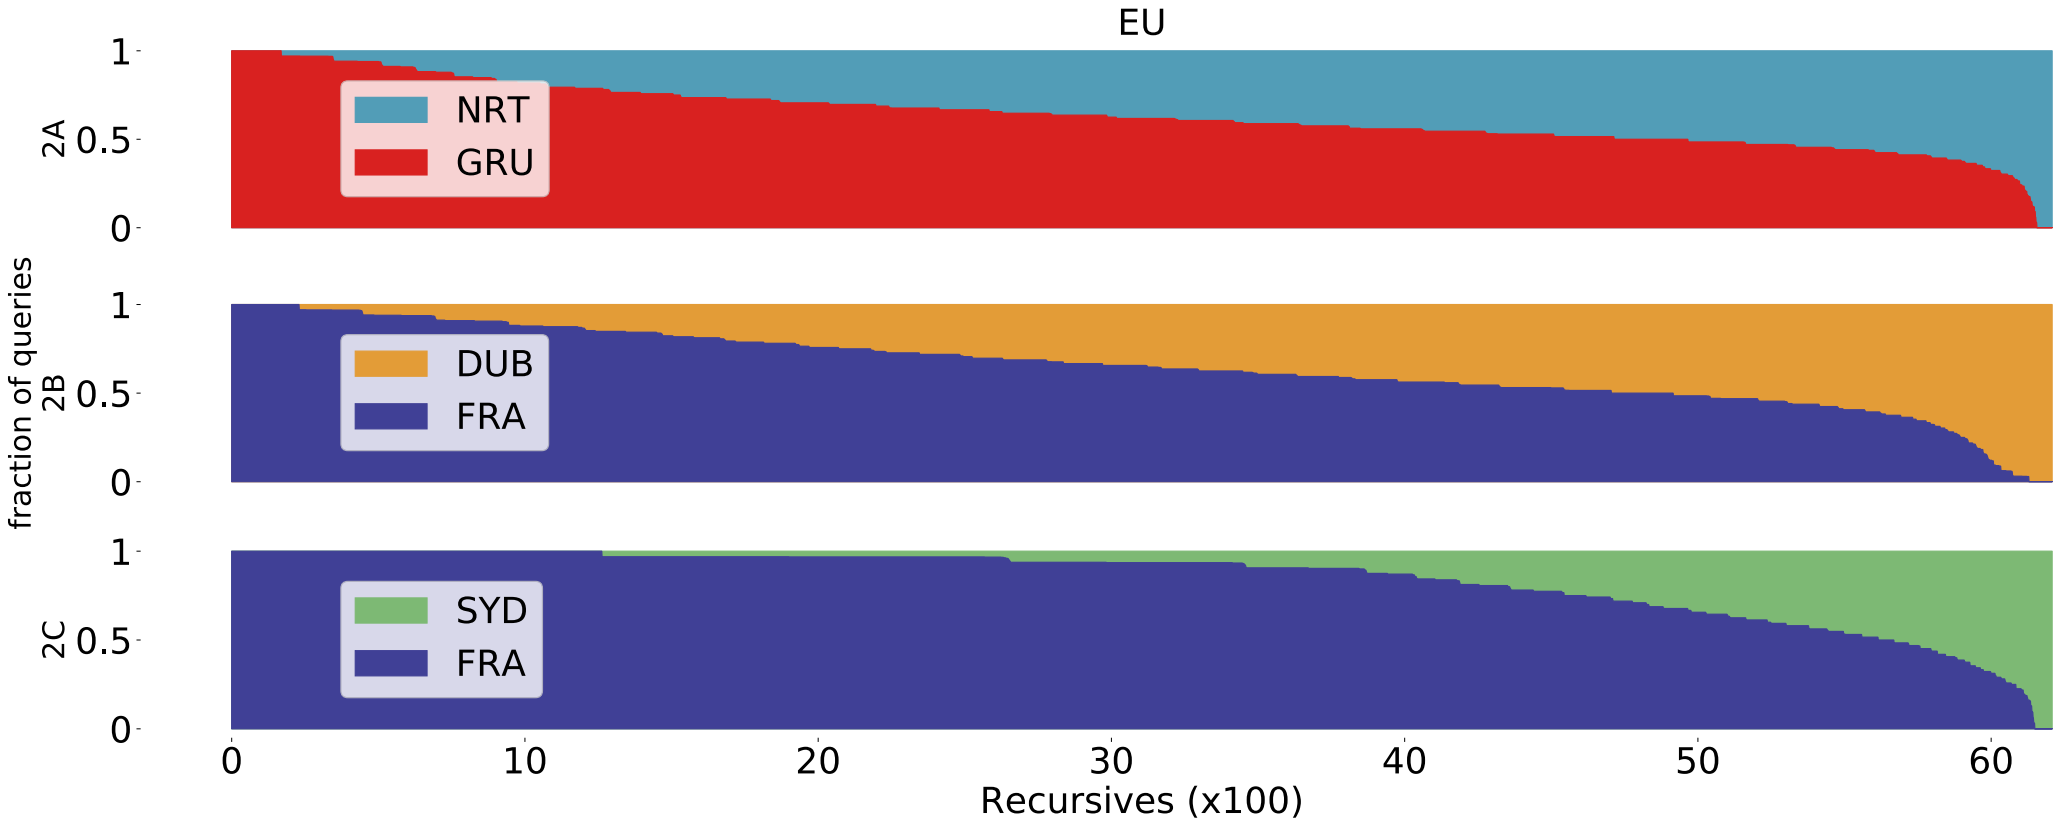

How do *individual* recursives distribute their queries?

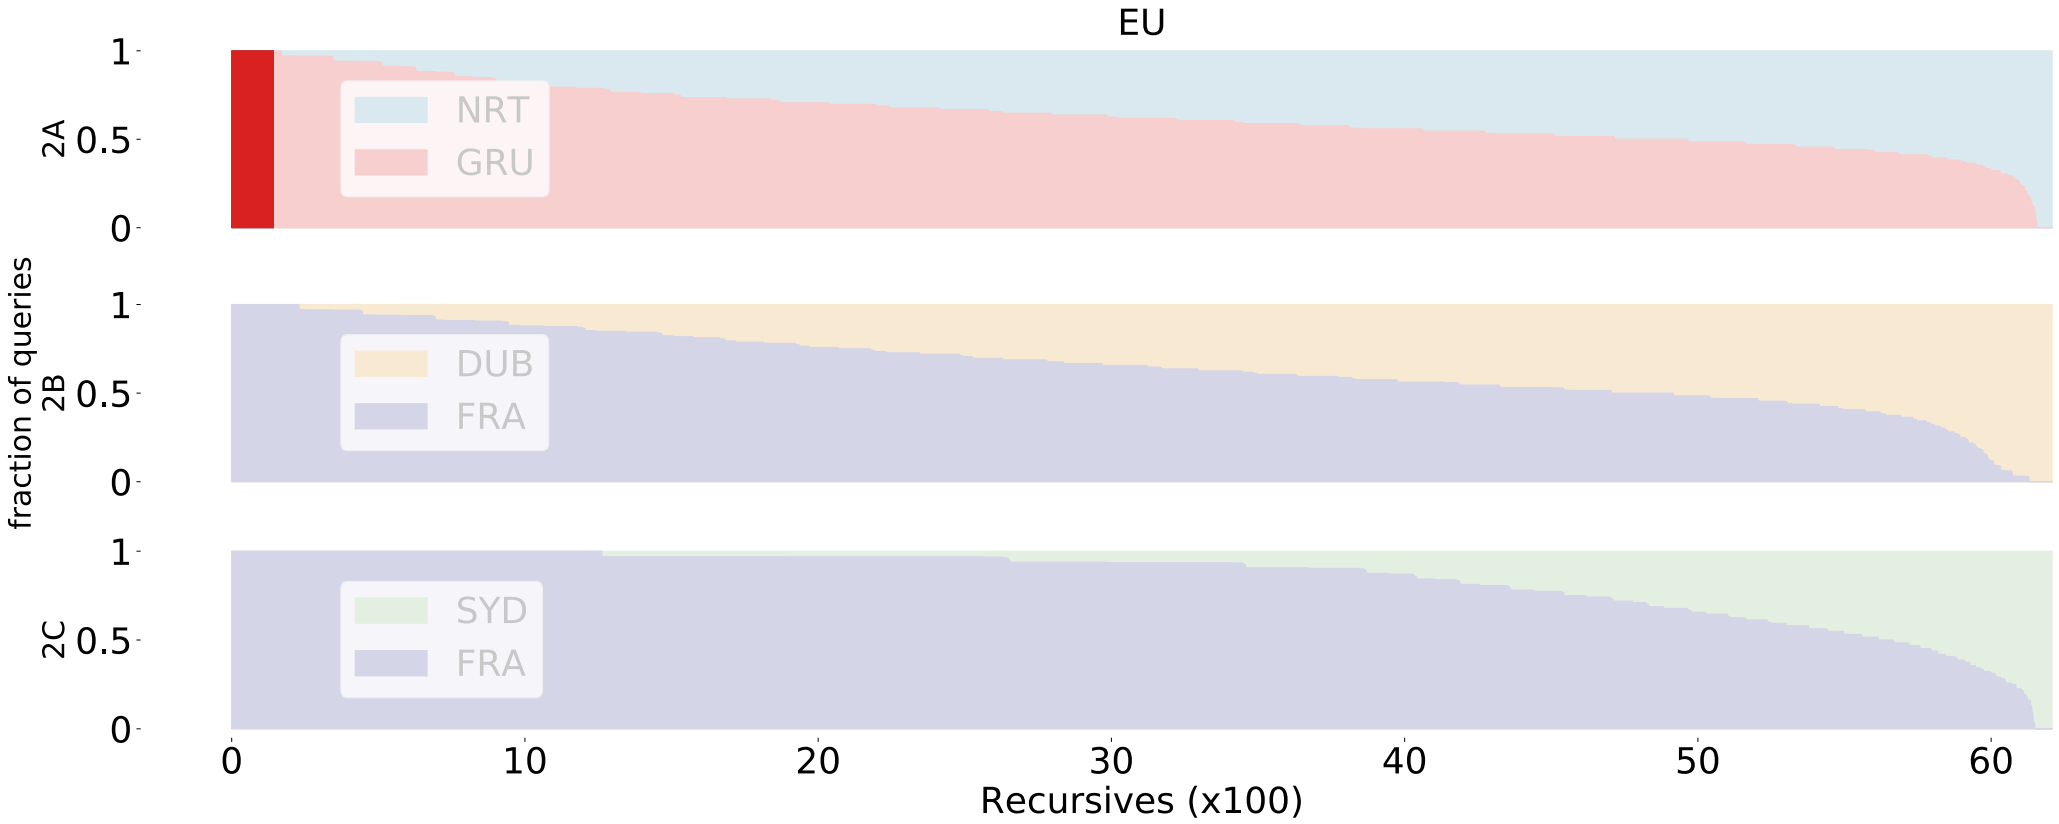

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do *individual* recursives distribute their queries?

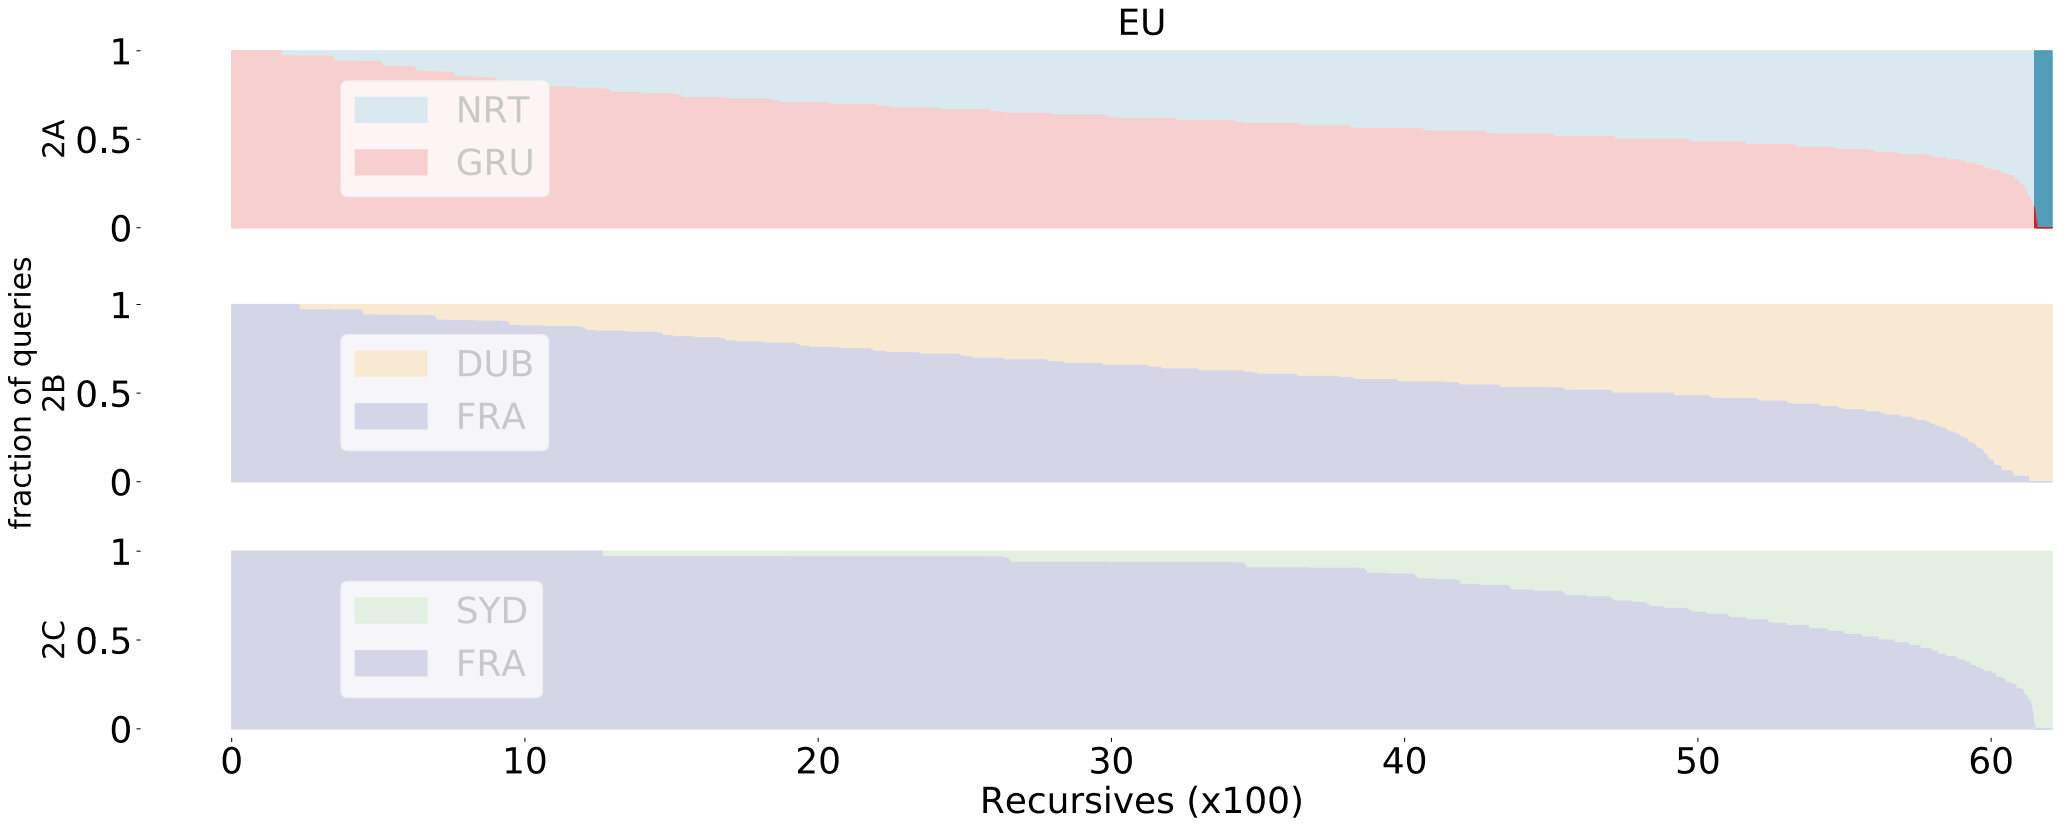

How do *individual* recursives distribute their queries?

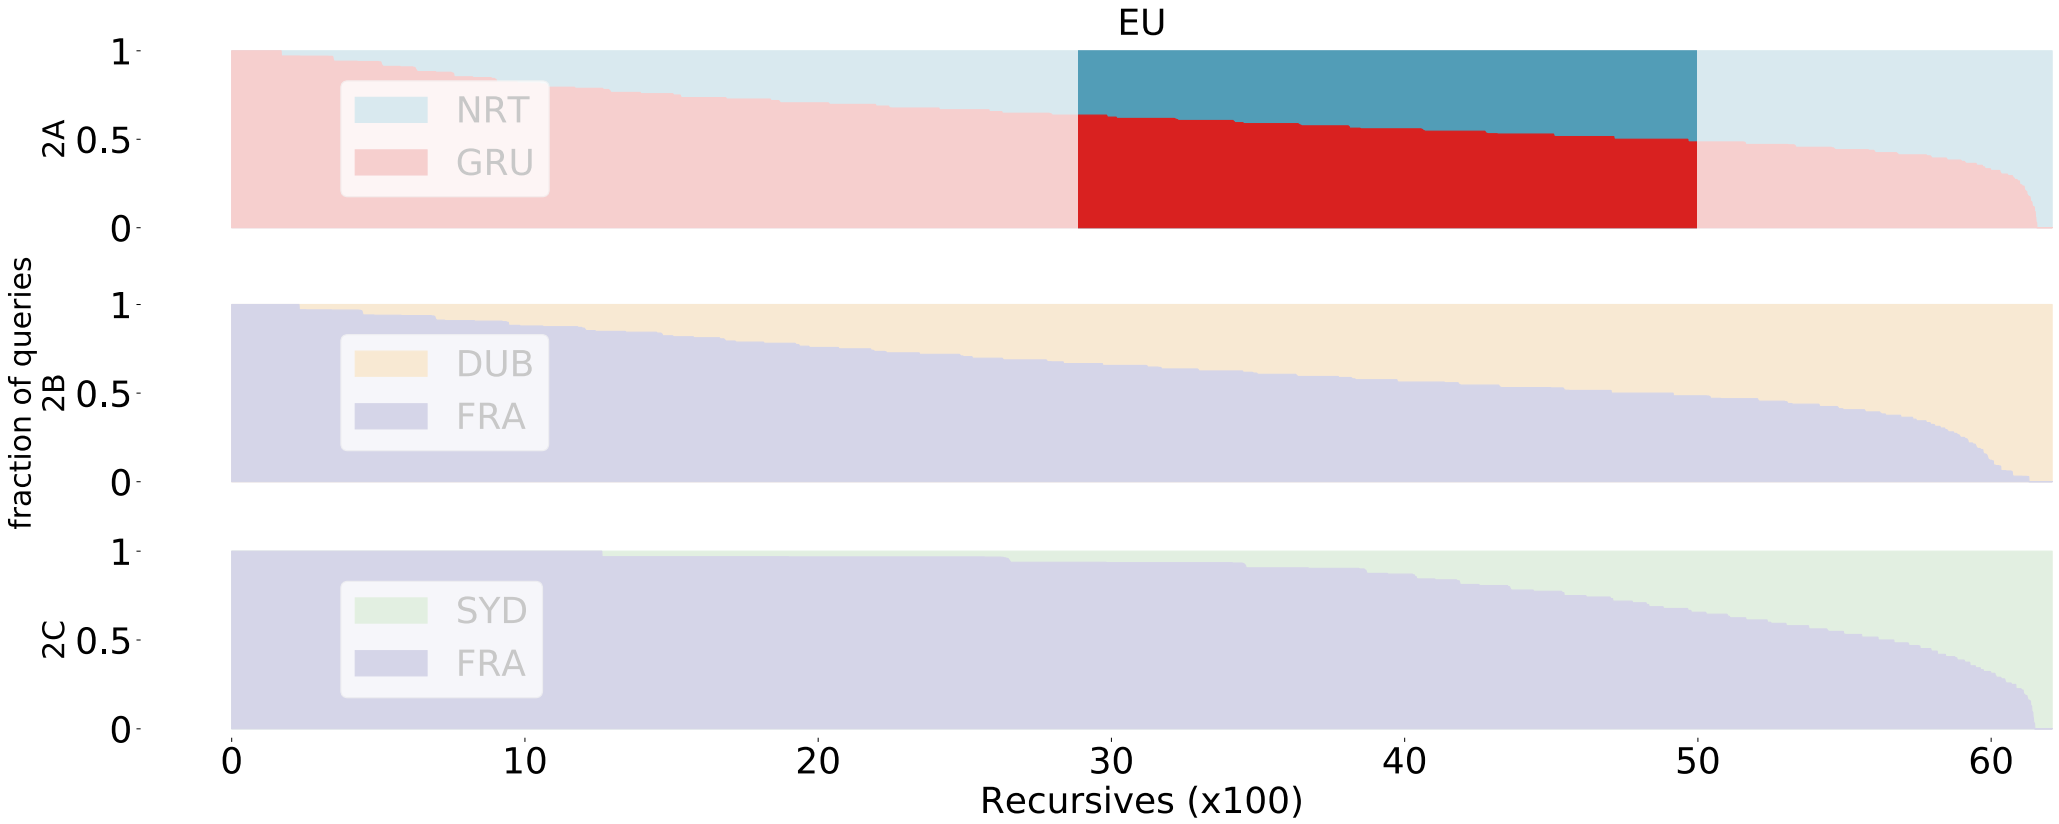

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do *individual* recursives distribute their queries?

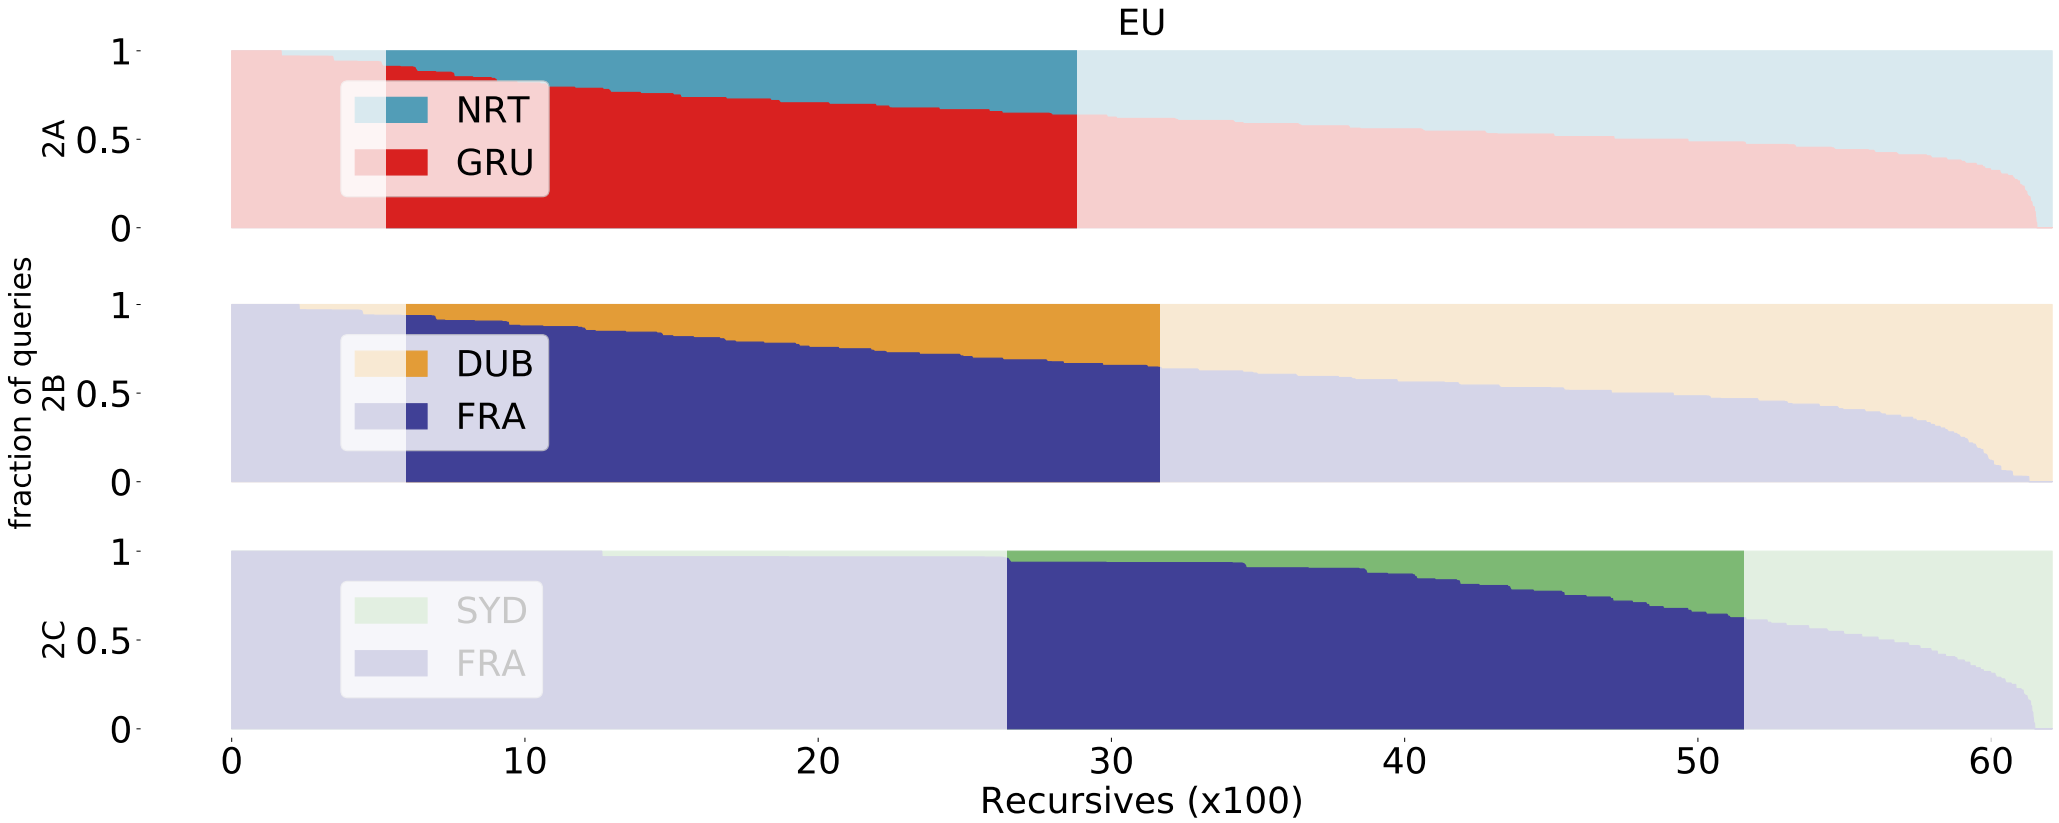

Up to 69% of resolvers have a weak preference
(60% to 90% of their queries to one NS)

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do *individual* recursives distribute their queries?

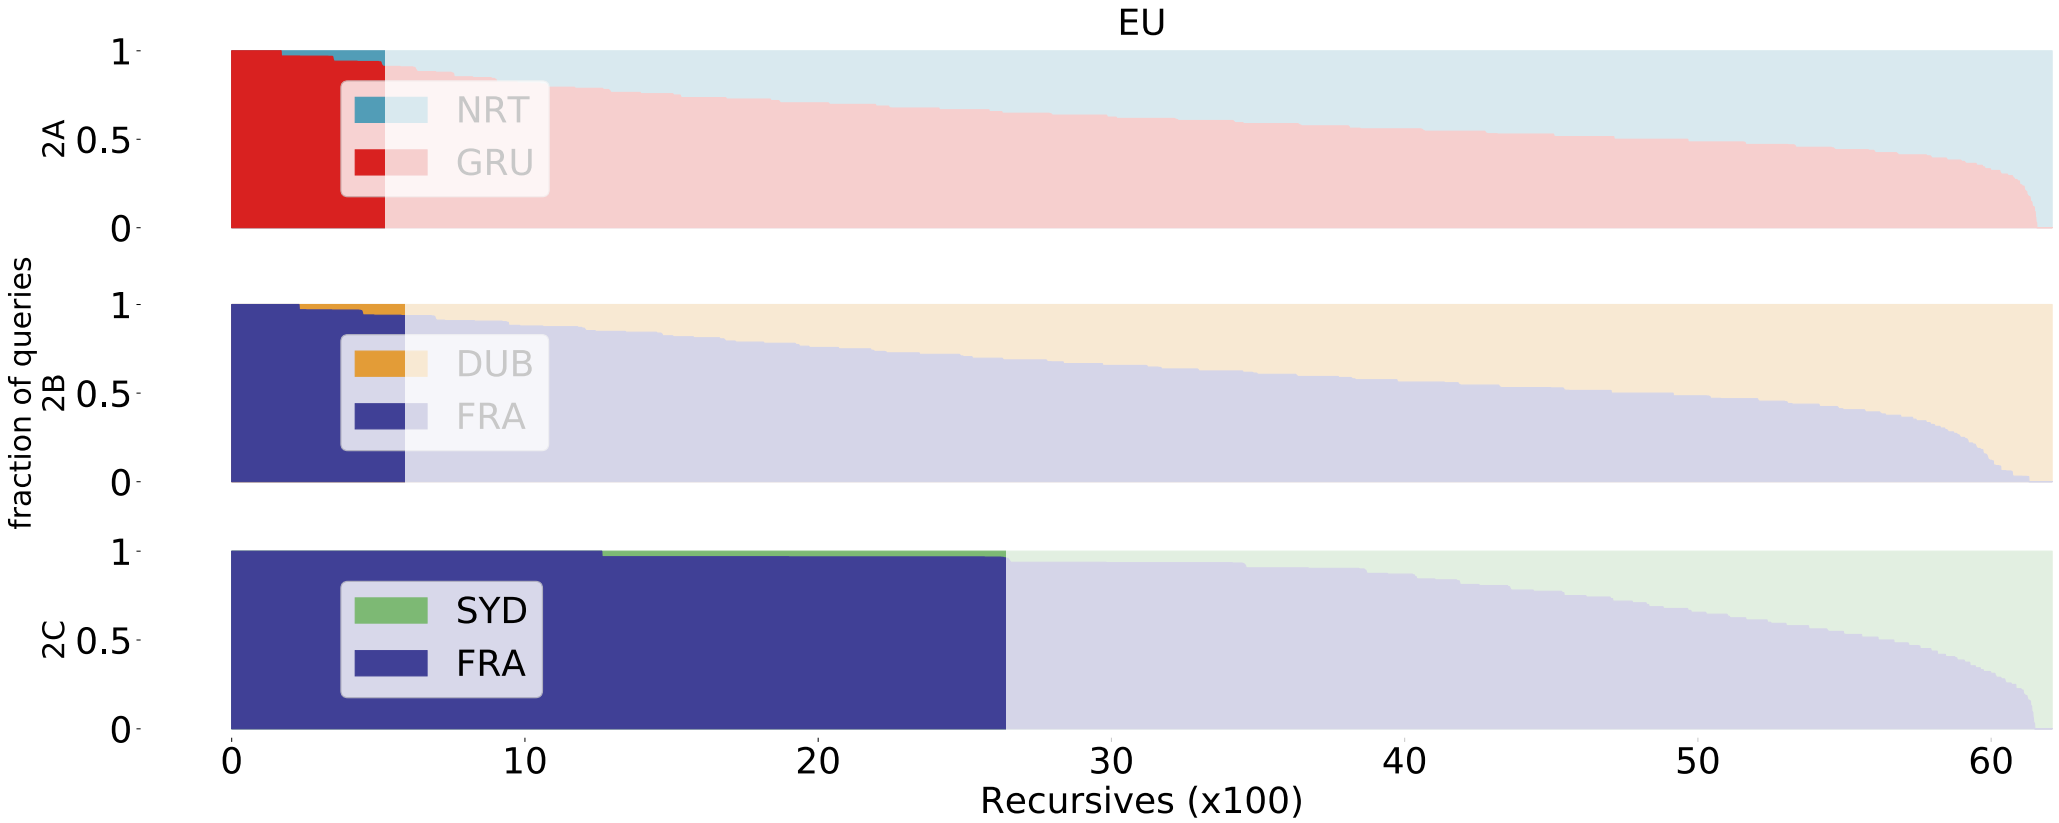

Up to 37% of resolvers have a strong preference
(more than 90% of their queries to one NS)

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How do *individual* recursives distribute their queries?

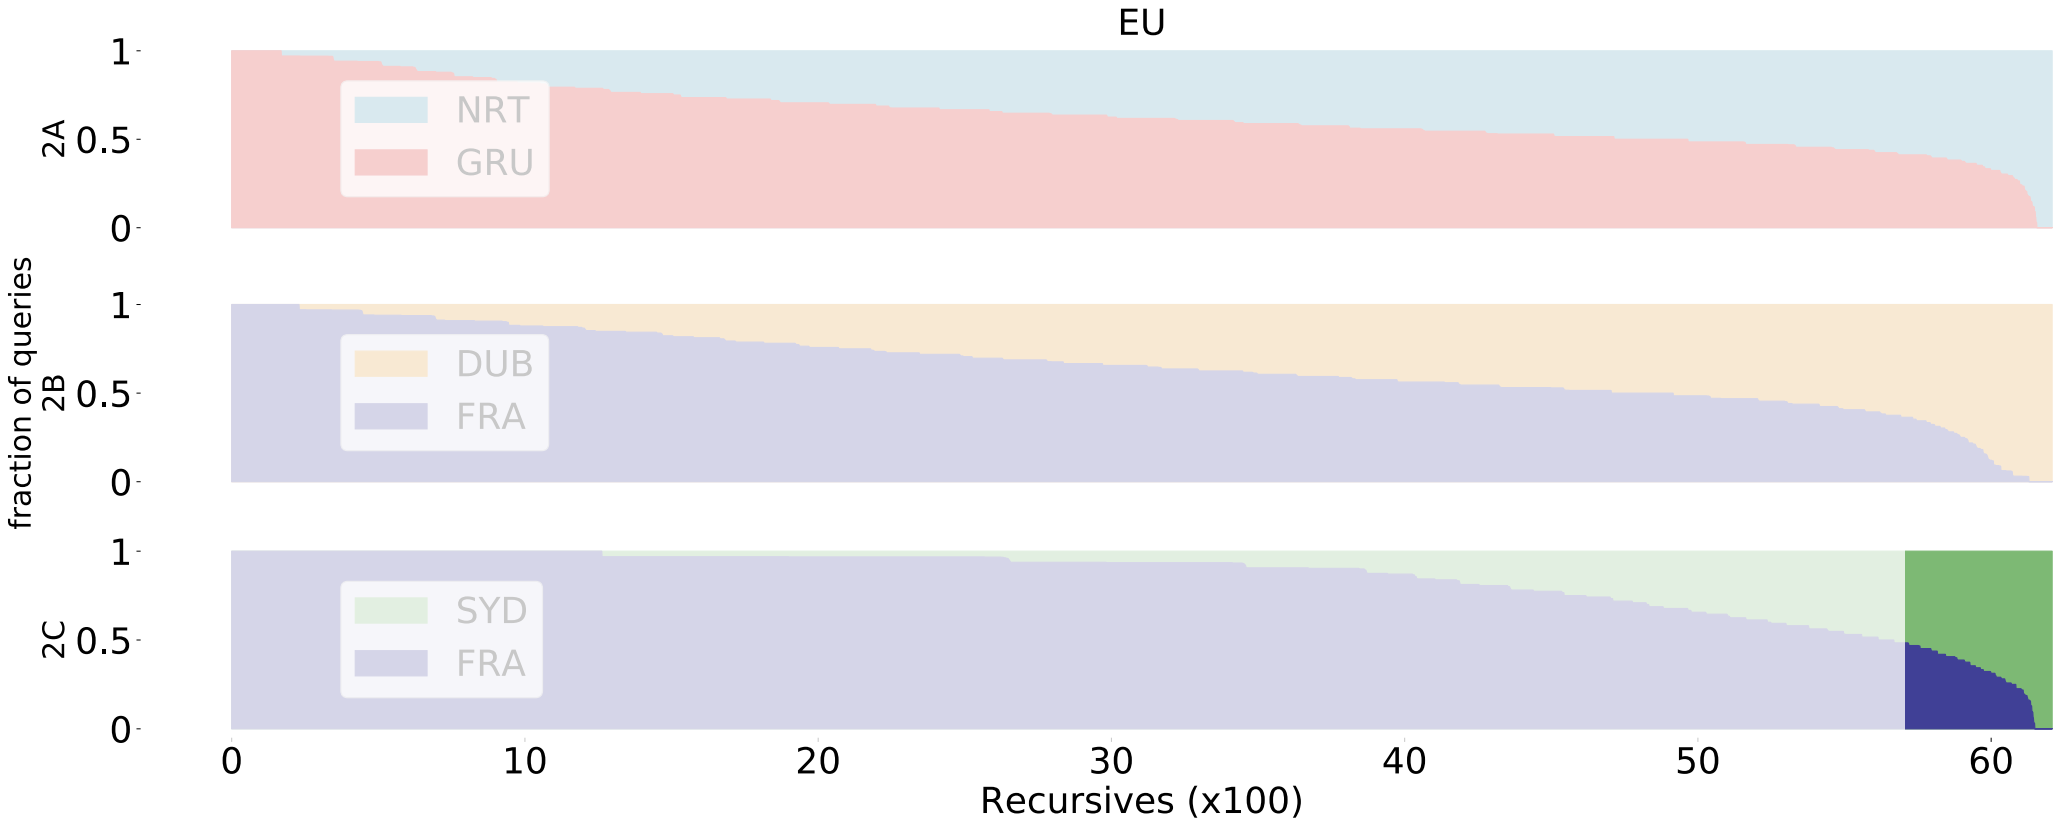

Some resolvers always prefer the slower NS

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

Validation: Authoritatives in Production

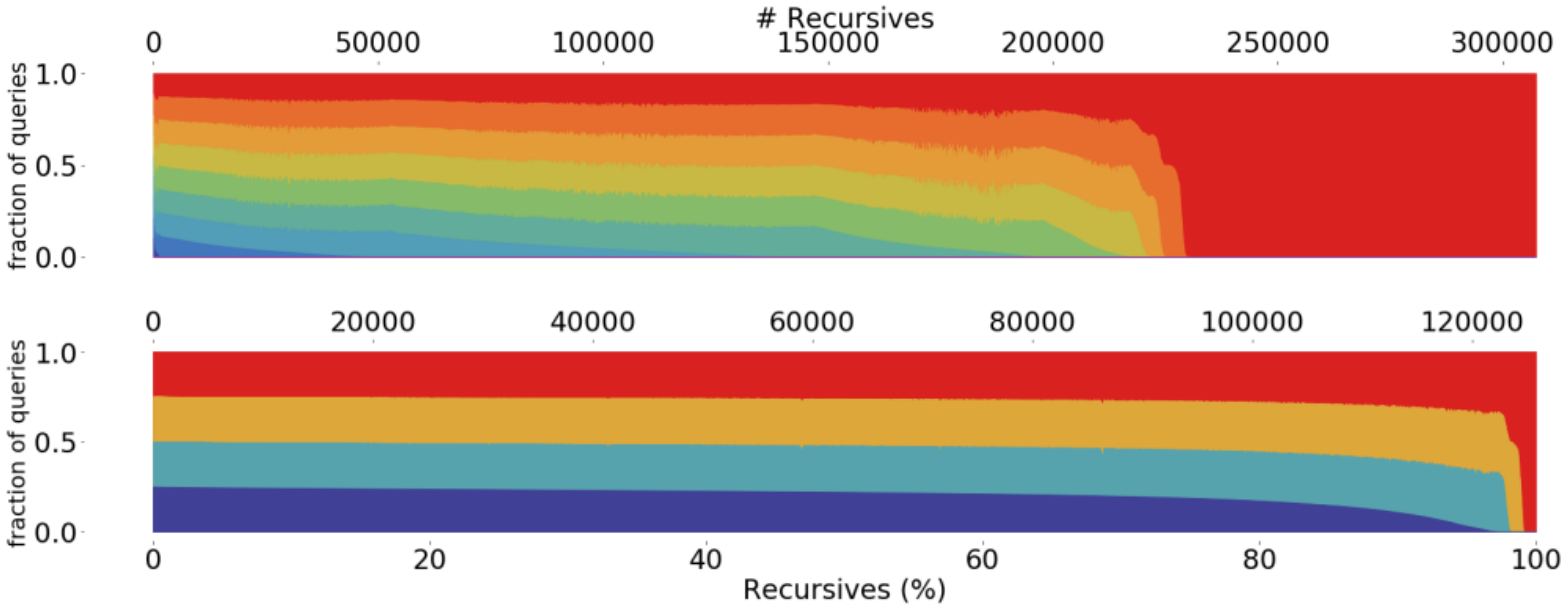

Root Servers
(10 out of 13)

.nl Servers
(4 out of 8)

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

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS

How does query frequency affect the results?

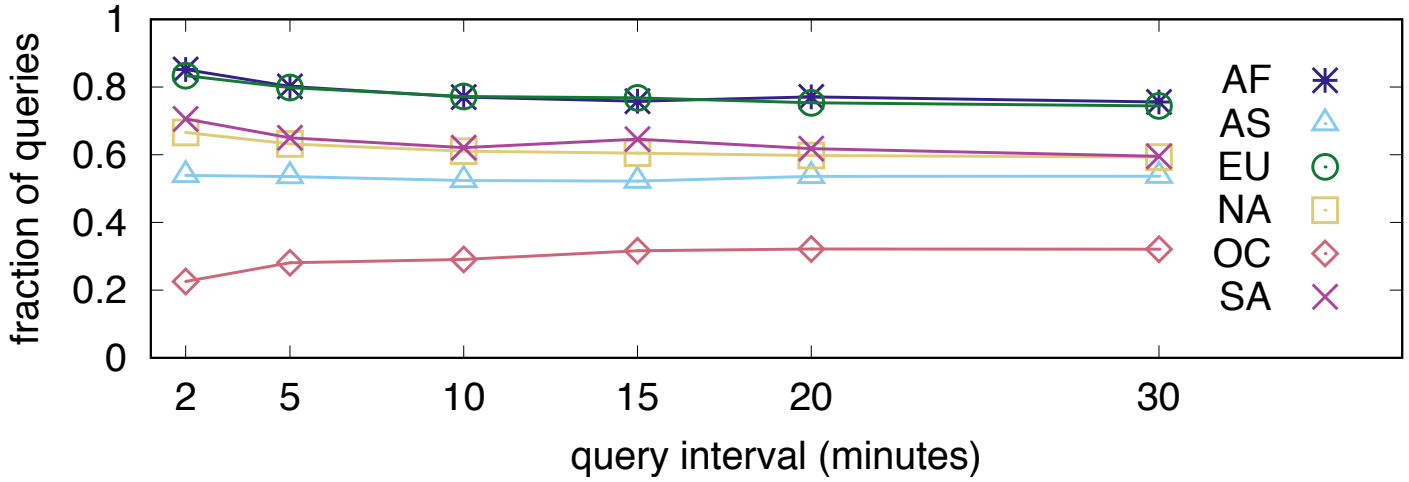

- A higher query frequency leads to a stronger preference
- Preference persists even after the default timeout

Measurement Summary

- 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:
  - Stronger preference when querying more frequent (< 10min interval)
  - Lower RTT becomes more relevant if competing NSes are closer (<150 ms)

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

Questions?

Moritz Müller
email: moritz.muller@sidn.nl
twitter: @moritzcm_

Additional Slides

Does preference change for distant recursives?

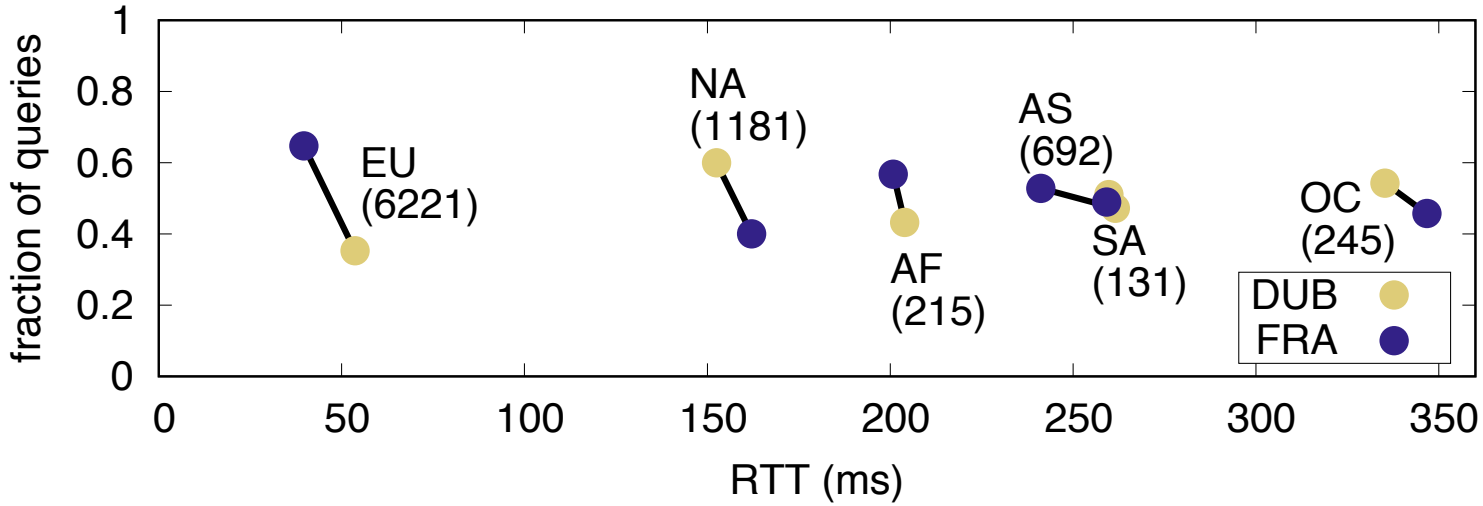

- VPs in EU reach Frankfurt 13 ms faster than Dublin, they clearly prefer Frankfurt
- VPs in Asia reach Frankfurt 20 ms faster, query distribution almost equal

→ Lower RTT becomes more relevant if competing authoritatives are closer

How do *individual* recursives distribute their queries?

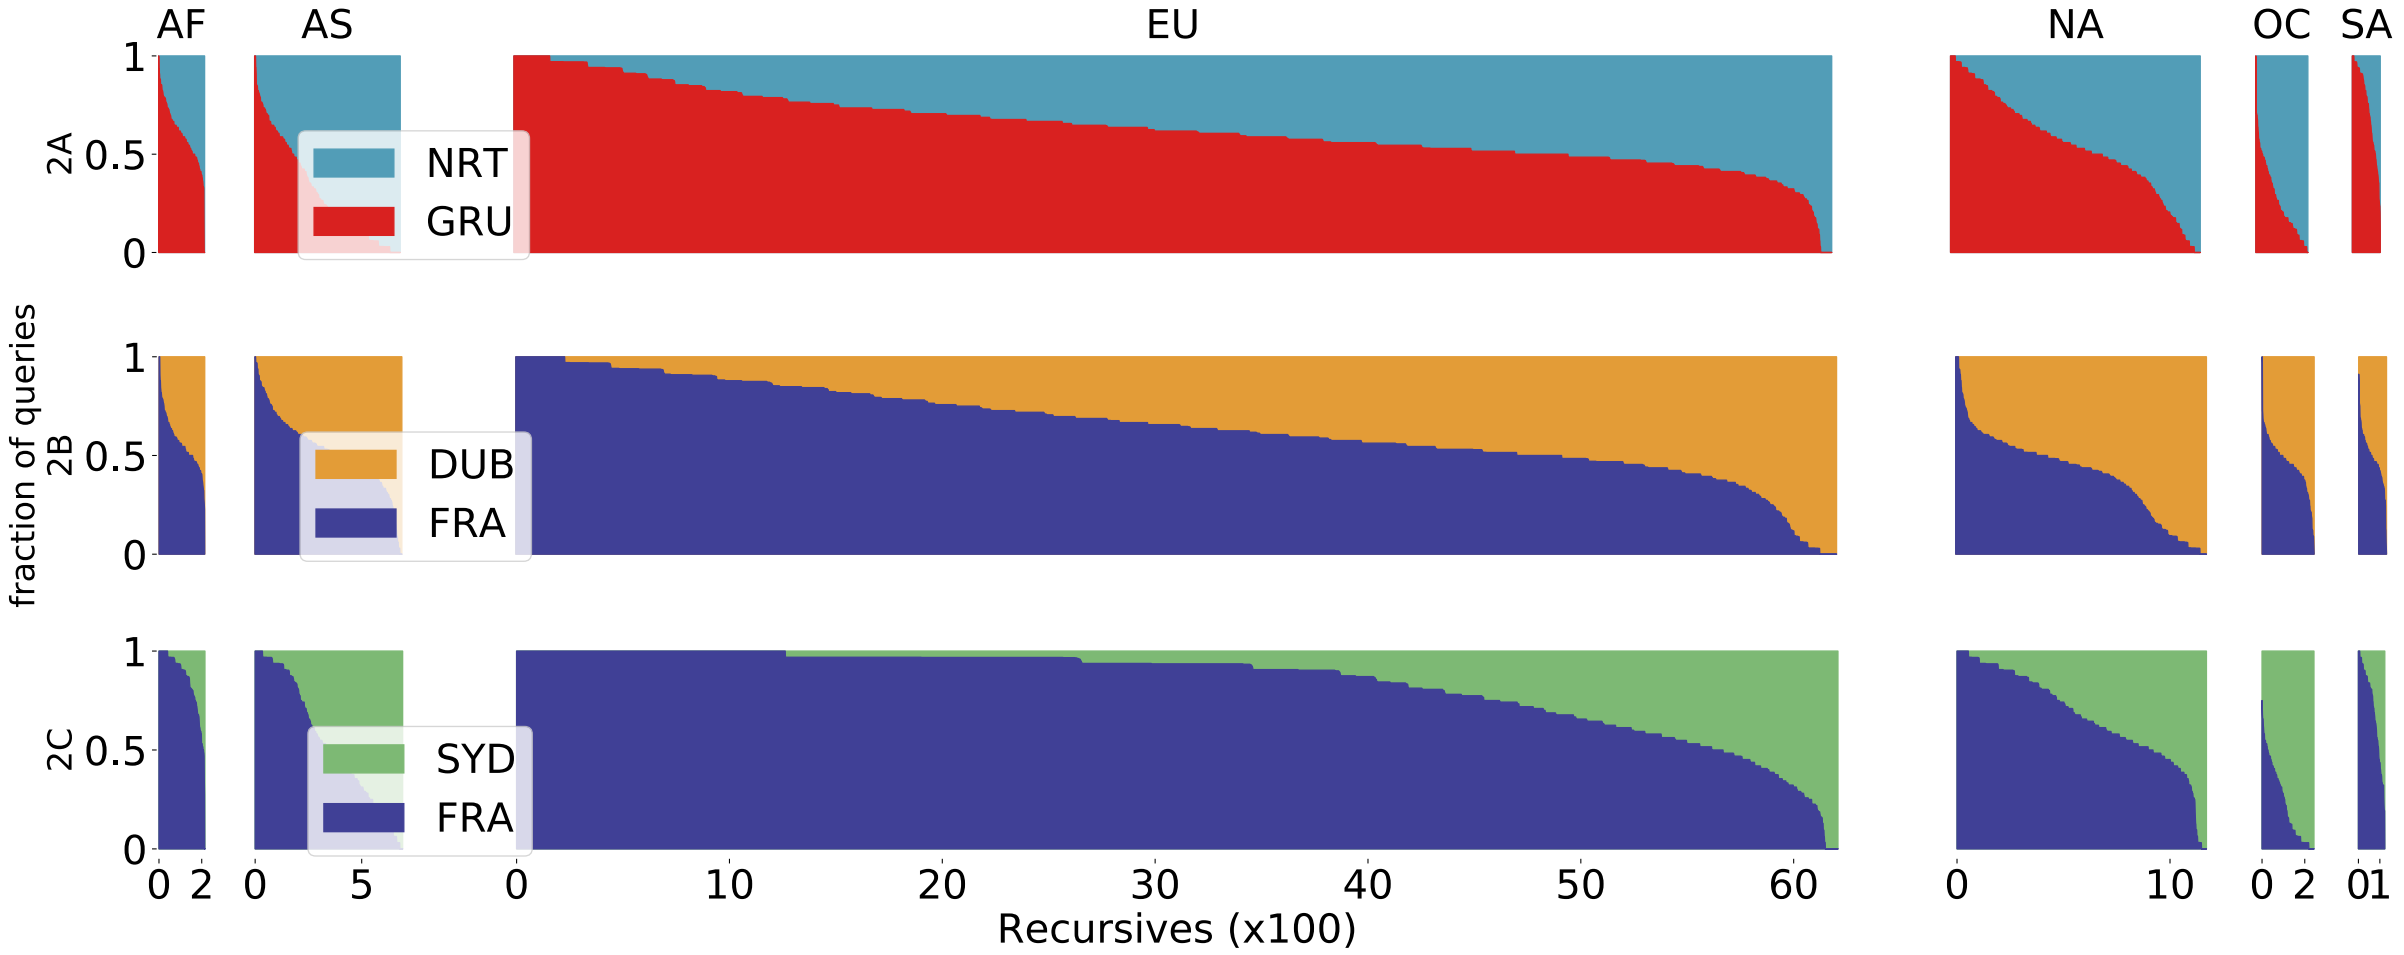

USC Viterbi School of Engineering

UNIVERSITY OF TWENTE.

SIDN LABS