Recursives in the Wild: Engineering Authoritative DNS Servers

Moritz Müller[1,2], **Giovane C. M. Moura**[1],
Ricardo de O. Schmidt[1,2], John Heidemann[3]

[1]SIDN Labs
The Netherlands

[2]University of Twente
The Netherlands

[3]USC/Information Sciences Institute
U.S.

IETF99 - IEPG
Prague, CZ, July 16th, 2017

Introduction

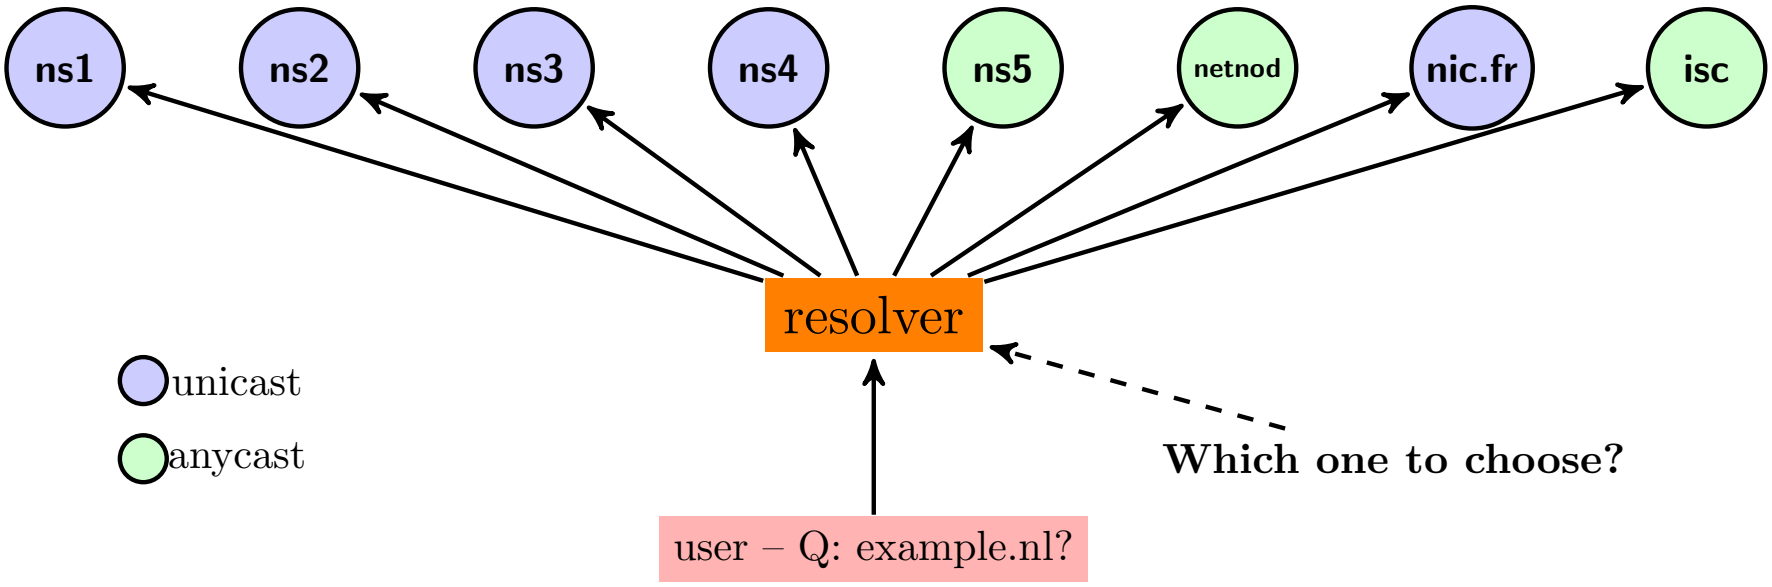

Figure: Resolving a Name under `.nl` TLD

Introduction

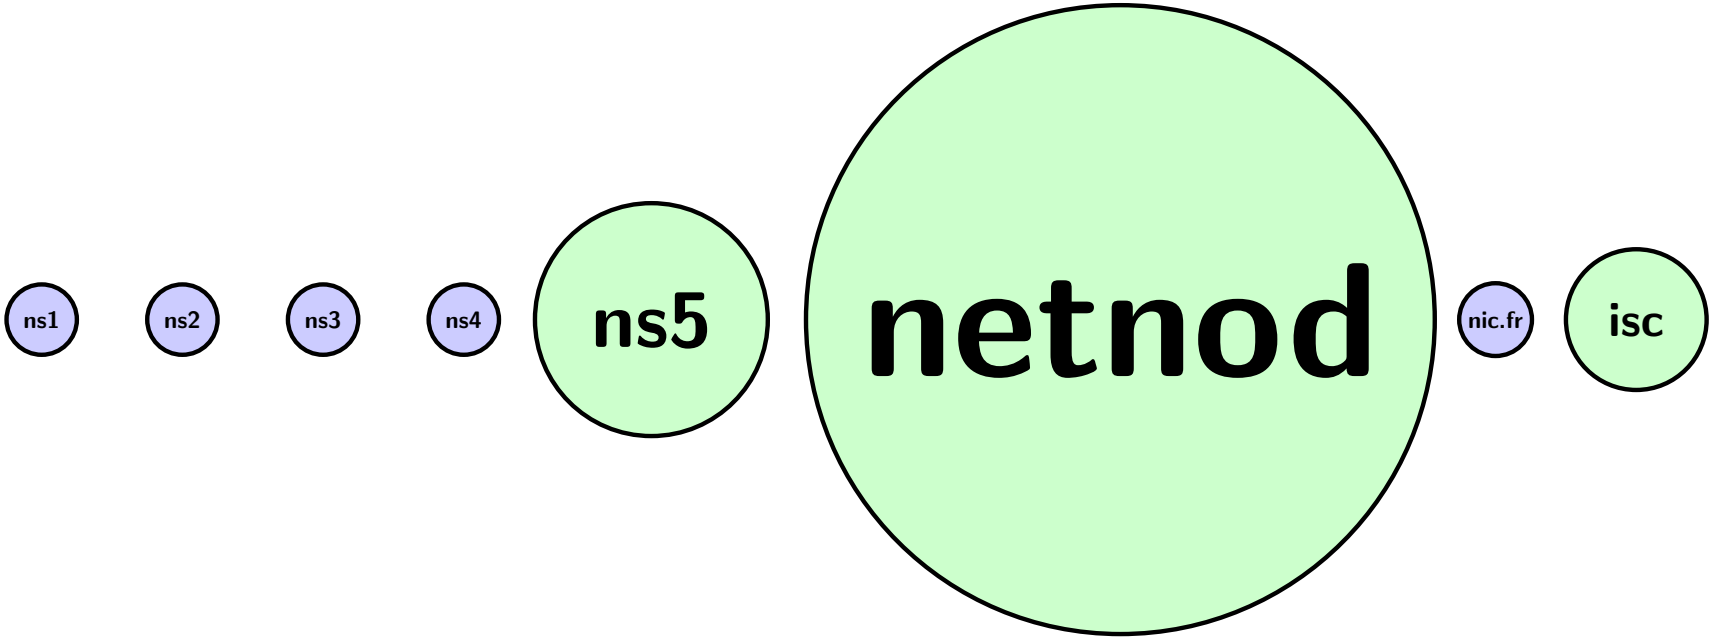

Figure: Authoritative Servers by Size (sites) - area proportional to number of sites on `.nl` (June 2016)

Introduction

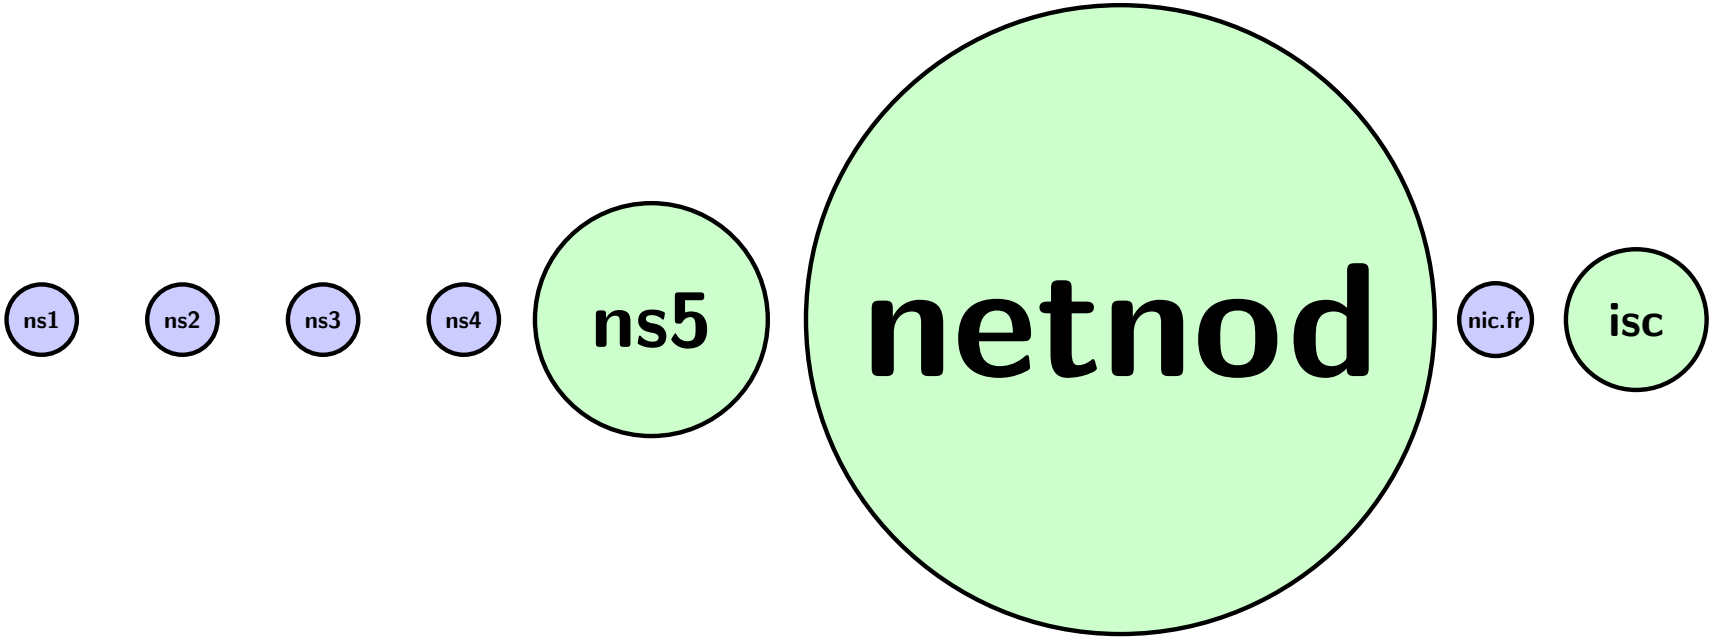

Figure: Authoritative Servers by Size (sites) - area proportional to number of sites on `.nl` (June 2016)

- *The larger, the more queries it gets, right?*

Introduction

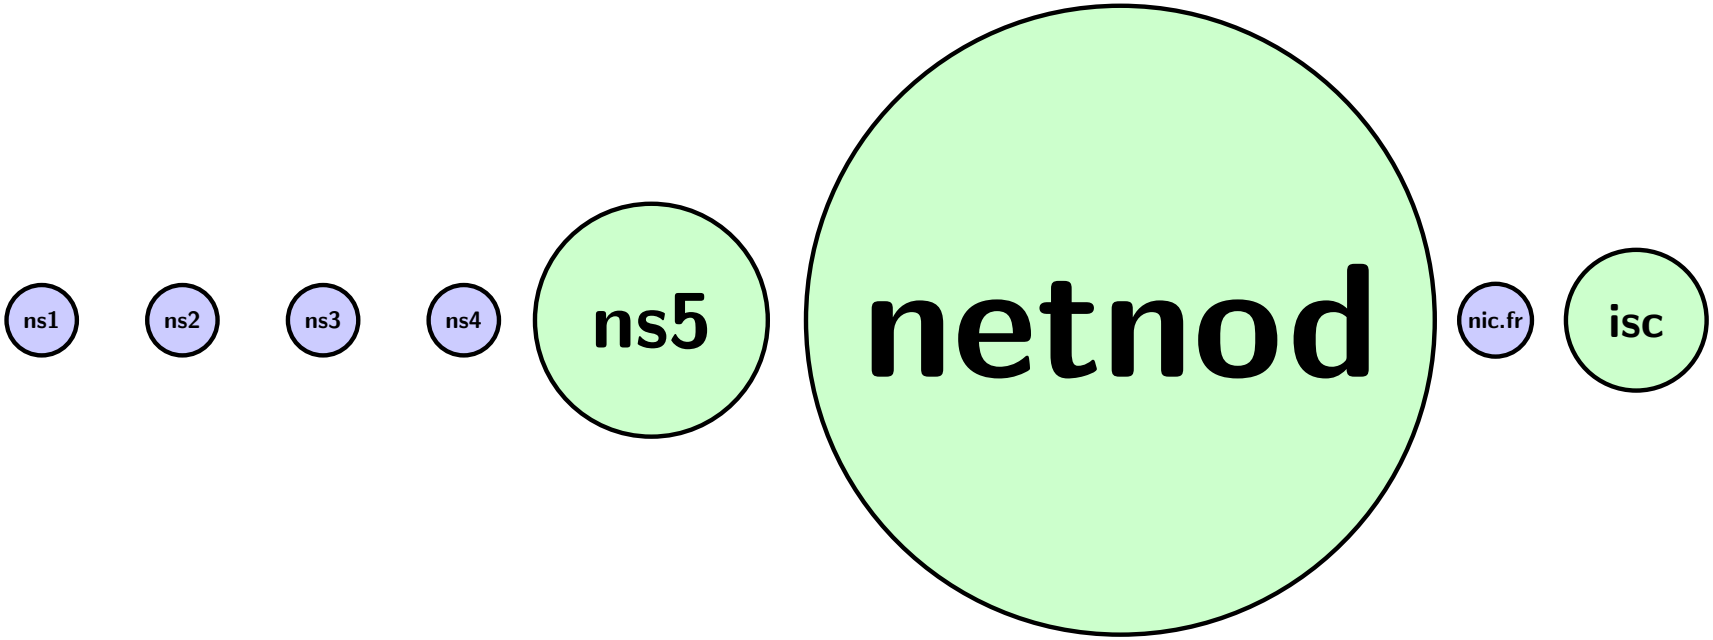

Figure: Authoritative Servers by Size (sites)

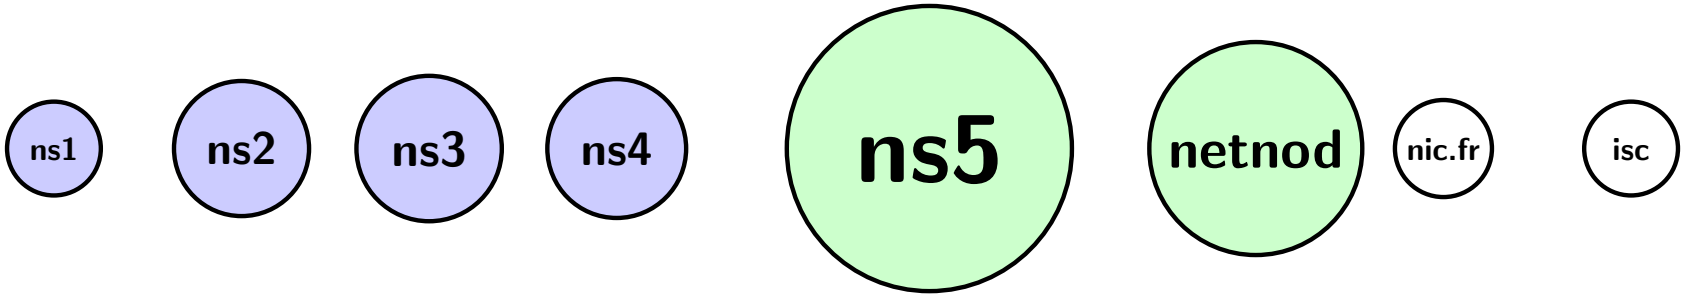

Figure: Authoritative Servers by Query Volume (June 2016)

Introduction

- Why is this hapenning?
- Meaning: why recursives choose this way → how do they behave in the wild?
- Study goal: analyze how recursives behaves in the wild with the goal with **better enginnering authoritative servers**
  - Previous work (Yu et al., [1]) was done in 2012, controlled environment
  - Recursives typically prefered low latency authoritatives

Approach

1. Set up an authoritative server infrastructure at 2LD (`ourtestdomain.nl`), using 7 Amazon AWS datacenters, IPv4
2. VPs: 9000+ Ripe Atlas probes
   - VP = probe_id + IP of local recursive
3. We vary the number/location of servers (NS records) and measure how VPs choose authoritatives
4. We use TXT records to determine which server responded to each probe/recursive
   - e.g: similar to chaos queries
   - Every 2min, for 1 hour
   - NS record TTL of 5 seconds (to ensure cold cache)
5. We also look at the root and `.nl` auth data

Setup

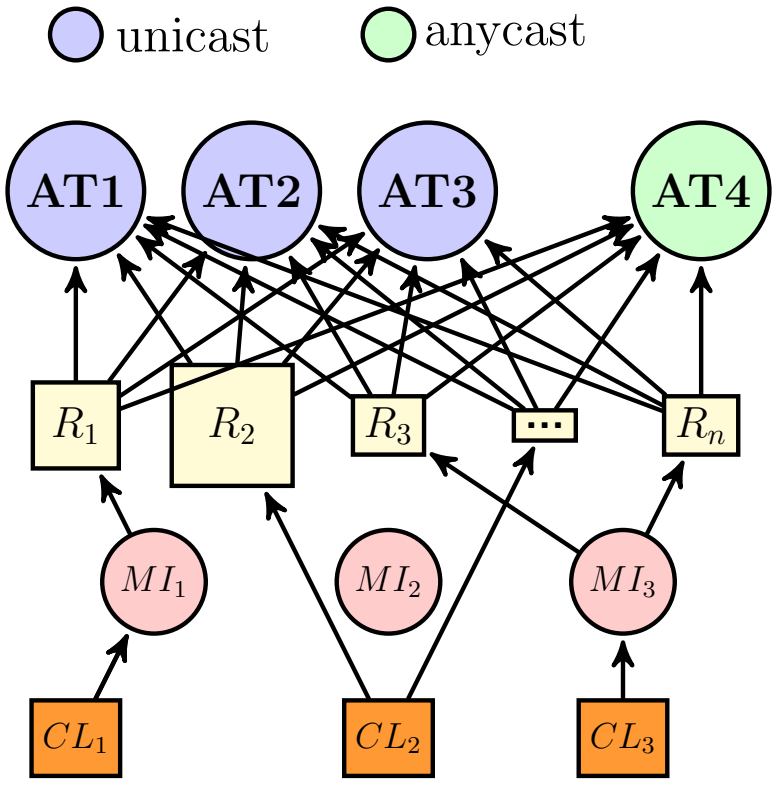

Figure: TLD Setup, Recursives, Middleboxes and Clients.

Setup

ID	locations (airport code)	VPs
2A	GRU (São Paulo, BR), NRT (Tokyo, JP)	8,702
2B	DUB (Dublin, IE), FRA (Frankfurt, DE)	8,685
2C	FRA, SYD (Sydney, AU)	8,658
3A	GRU, NRT, SYD	8,684
3B	DUB, FRA, IAD (Washington, US)	8,693
4A	GRU, NRT, SYD, DUB	8,702
4B	DUB, FRA, IAD, SFO (San Francisco, US)	8,689

Table: Combinations of authoritatives we deploy and the number of VPs they see.

Do recursives query all authoritatives?

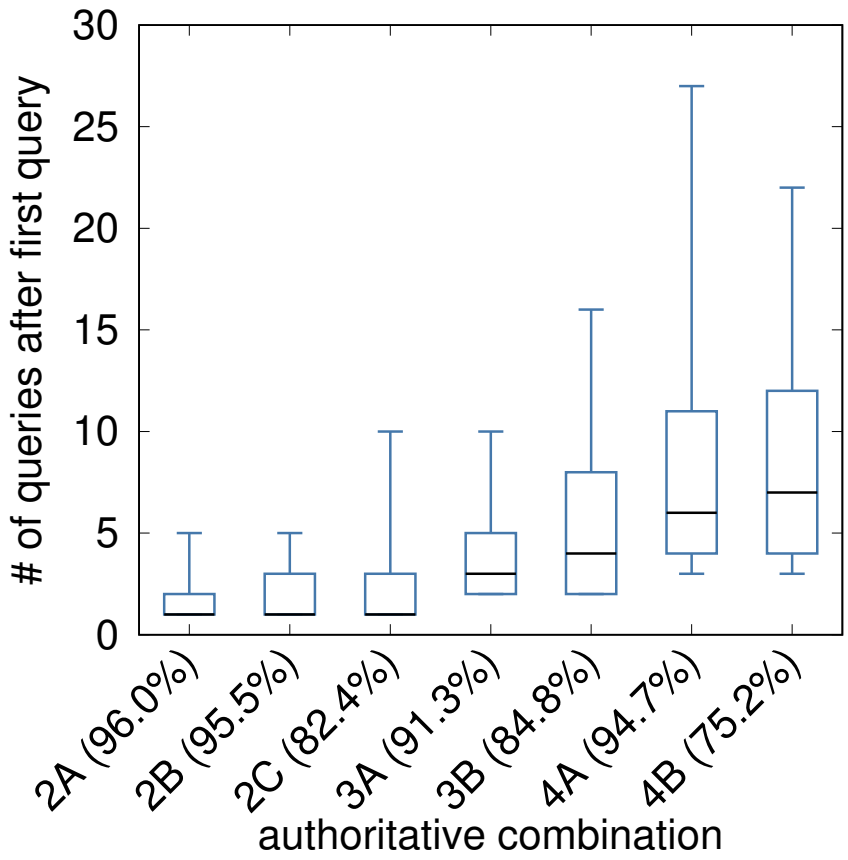

Figure: Queries to probe all authoritatives, after the first query.

- Yes! Most query all!

How are queries distributed over time?

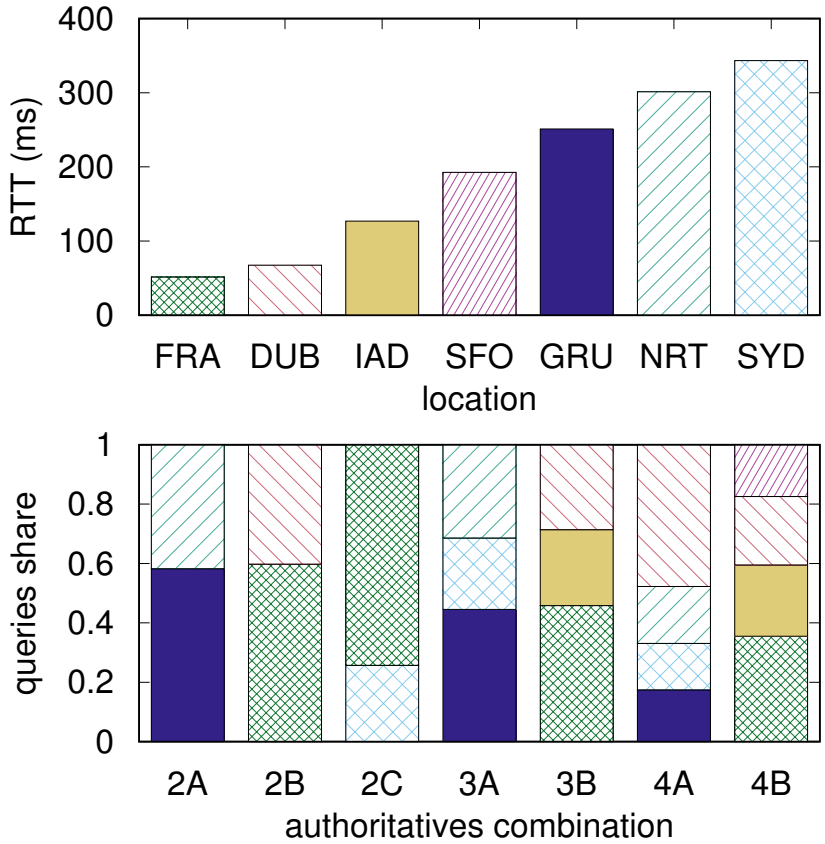

Figure: Median RTT (top) and query distribution (bottom) for combinations of authoritatives.

Confirming [1], but now in the wild

How do recursives distribute queries?

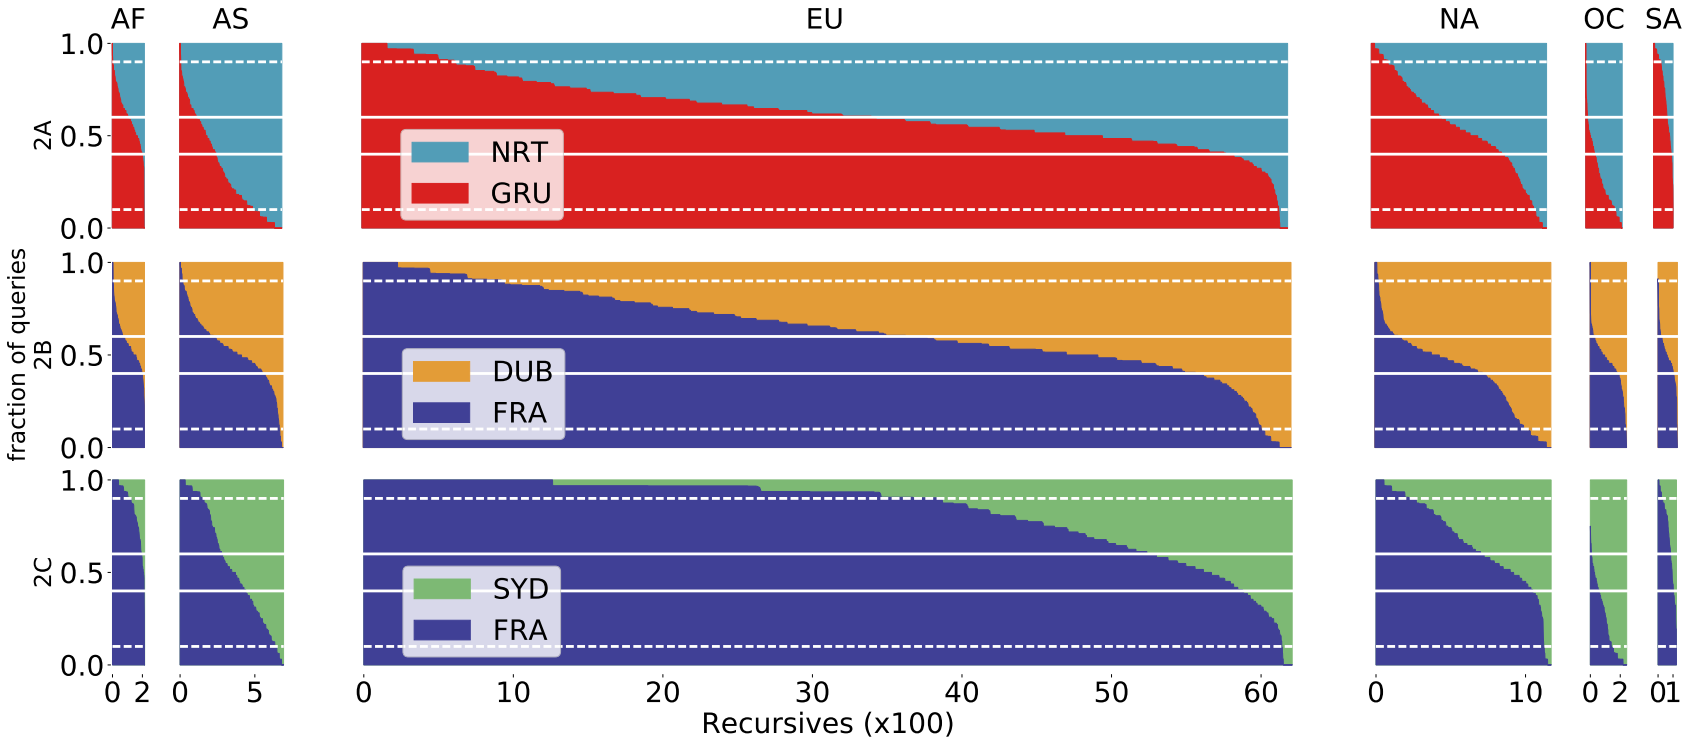

Figure: Recursive queries distribution for authoritative combinations 2A (top), 2B (center) and 2C (bottom). Solid and dotted horizontal lines mark VPs with weak and strong preference towards an authoritative.

How do recursives distribute queries?

- 59-69% of resolvers have a a week preference for an auth (60% of queries to one auth)
- 10-37% have strong pref to one auth (90% of queries to one auth)
- **Distribution is inversily proportional with median RTT**

How do recursives distribute queries?

- What happens when Auth are more less the same?

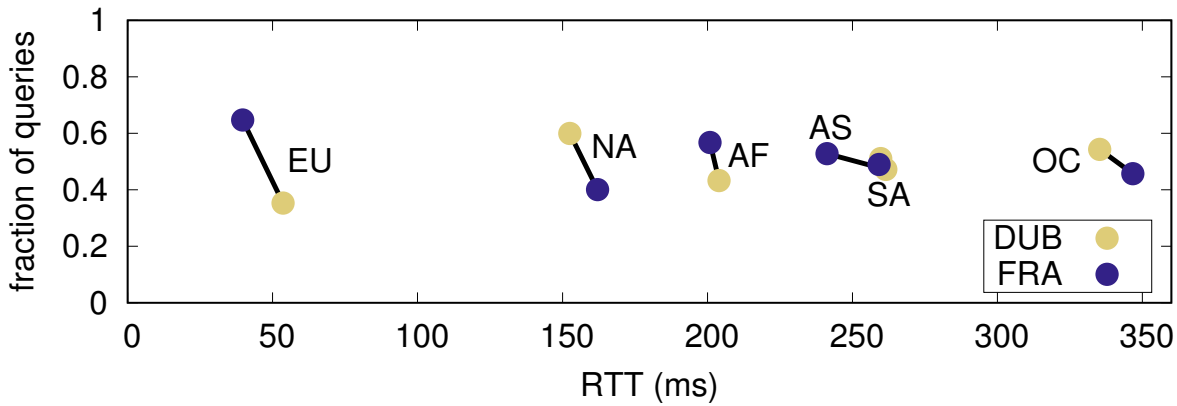

Figure: RTT sensitivity of 2B

- EU VPs get to FRA faster (13ms)
- Thus they prefer FRA slightly over DUB
- Asian VPs divide more equaly (despite 20.3ms diff!!)
- **RTT influence decreases for far away resolvers**

How query frequency affects the results?

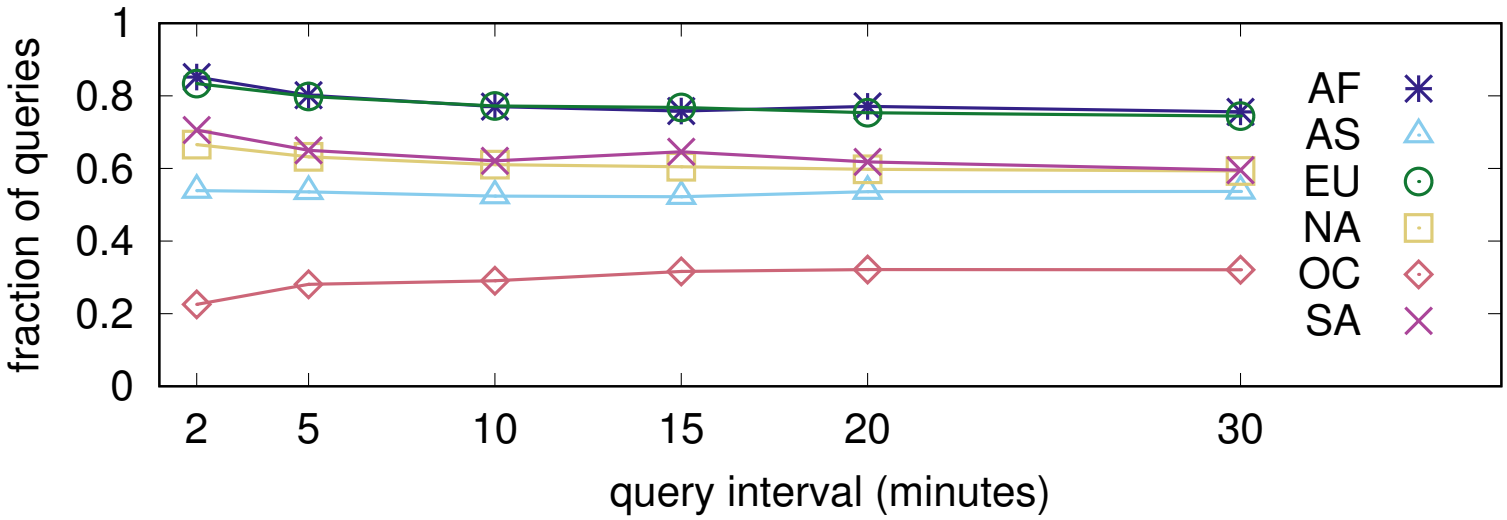

Figure: Fraction of queries to FRA (remainder go to SYD, configuration 2C), as query interval varies from 2 to 30 minutes.

Higher frequency, higher preference (infra-cache)

What about production zones? (`root` and `.nl`)?

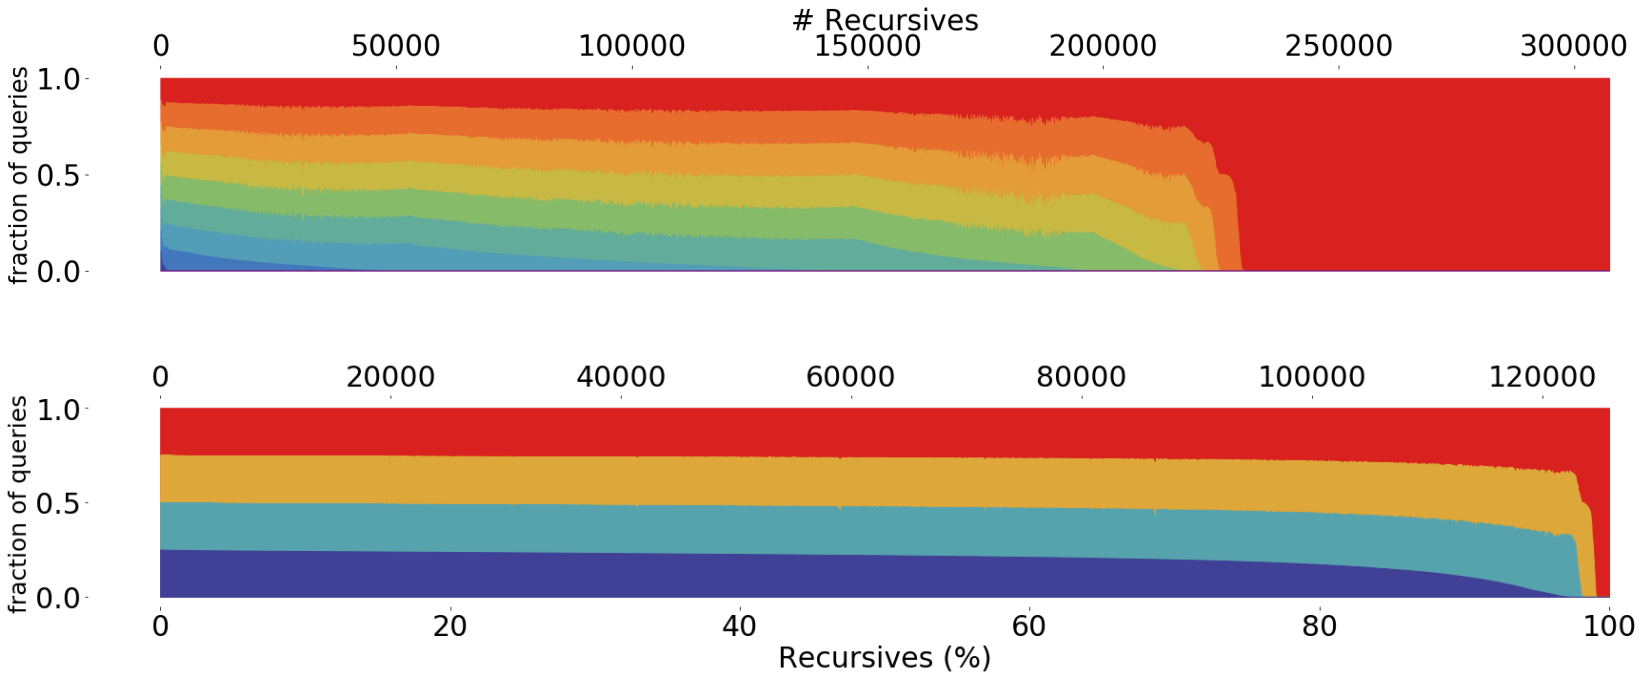

Figure: Distribution of queries of recursives with at least 250 queries across 10 out of 13 Root letters (top) and across 4 out of 8 name servers of `.nl` (bottom).

Conclusions and Recommendations

Main conclusion:

- Worst-case latency limited by the least anycasted authoritative
  - recursives use all authoritatives, query more often the better performing one (but diversity is important for them)
  - We (`.nl`) see 23% of incoming traffic in NL-based auth servers from the US, because of this (we're moving to anycast on all NSes)

Recommendation:

- Use Anycast on all your NS, and peer them very well, with multiple sites[2]
  - also important for DDoS[3]

Questions?

Contact details

Giovane C. M. Moura

giovane.moura@sidn.nl

Download our paper and data at:
https://tinyurl.com/y7exc5ts

Bibliography I

Y. Yu, D. Wessels, M. Larson, and L. Zhang, "Authority server selection in dns caching resolvers," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 2, pp. 80–86, Mar. 2012. [Online]. Available: http://doi.acm.org/10.1145/2185376.2185387

R. d. O. Schmidt, J. Heidemann, and J. H. Kuipers, "Anycast latency: How many sites are enough?" in *Proceedings of the Passive and Active Measurement Workshop*. Sydney, Australia: Springer, Mar. 2017, pp. 188–200. [Online]. Available: http://www.isi.edu/%7ejohnh/PAPERS/Schmidt17a.html

Bibliography II


G. C. M. Moura, R. de O. Schmidt, J. Heidemann, W. B. de Vries, M. Müller, L. Wei, and C. Hesselman, "Anycast vs. DDoS: Evaluating the November 2015 root DNS event," in *Proceedings of the ACM Internet Measurement Conference*, Nov. 2016. [Online]. Available: http://www.isi.edu/%7ejohnh/PAPERS/Moura16b.html