

# Anycast Latency

## How many sites are enough?

Presented by

Ricardo de Oliveira Schmidt



October 25, 2016  
Madrid, Spain

Reference:

## **Anycast Latency: How Many Sites Are Enough?**

Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers

Technical Report ISI-TR-2016-708, USC/Information Sciences Institute, May 2016

- <http://www.isi.edu/~johnh/PAPERS/Schmidt16a.pdf>

# Anycast

## *Definition*

Multiple copies of a service at *different locations* configured to answer with the *same IP address*

## *Purpose*

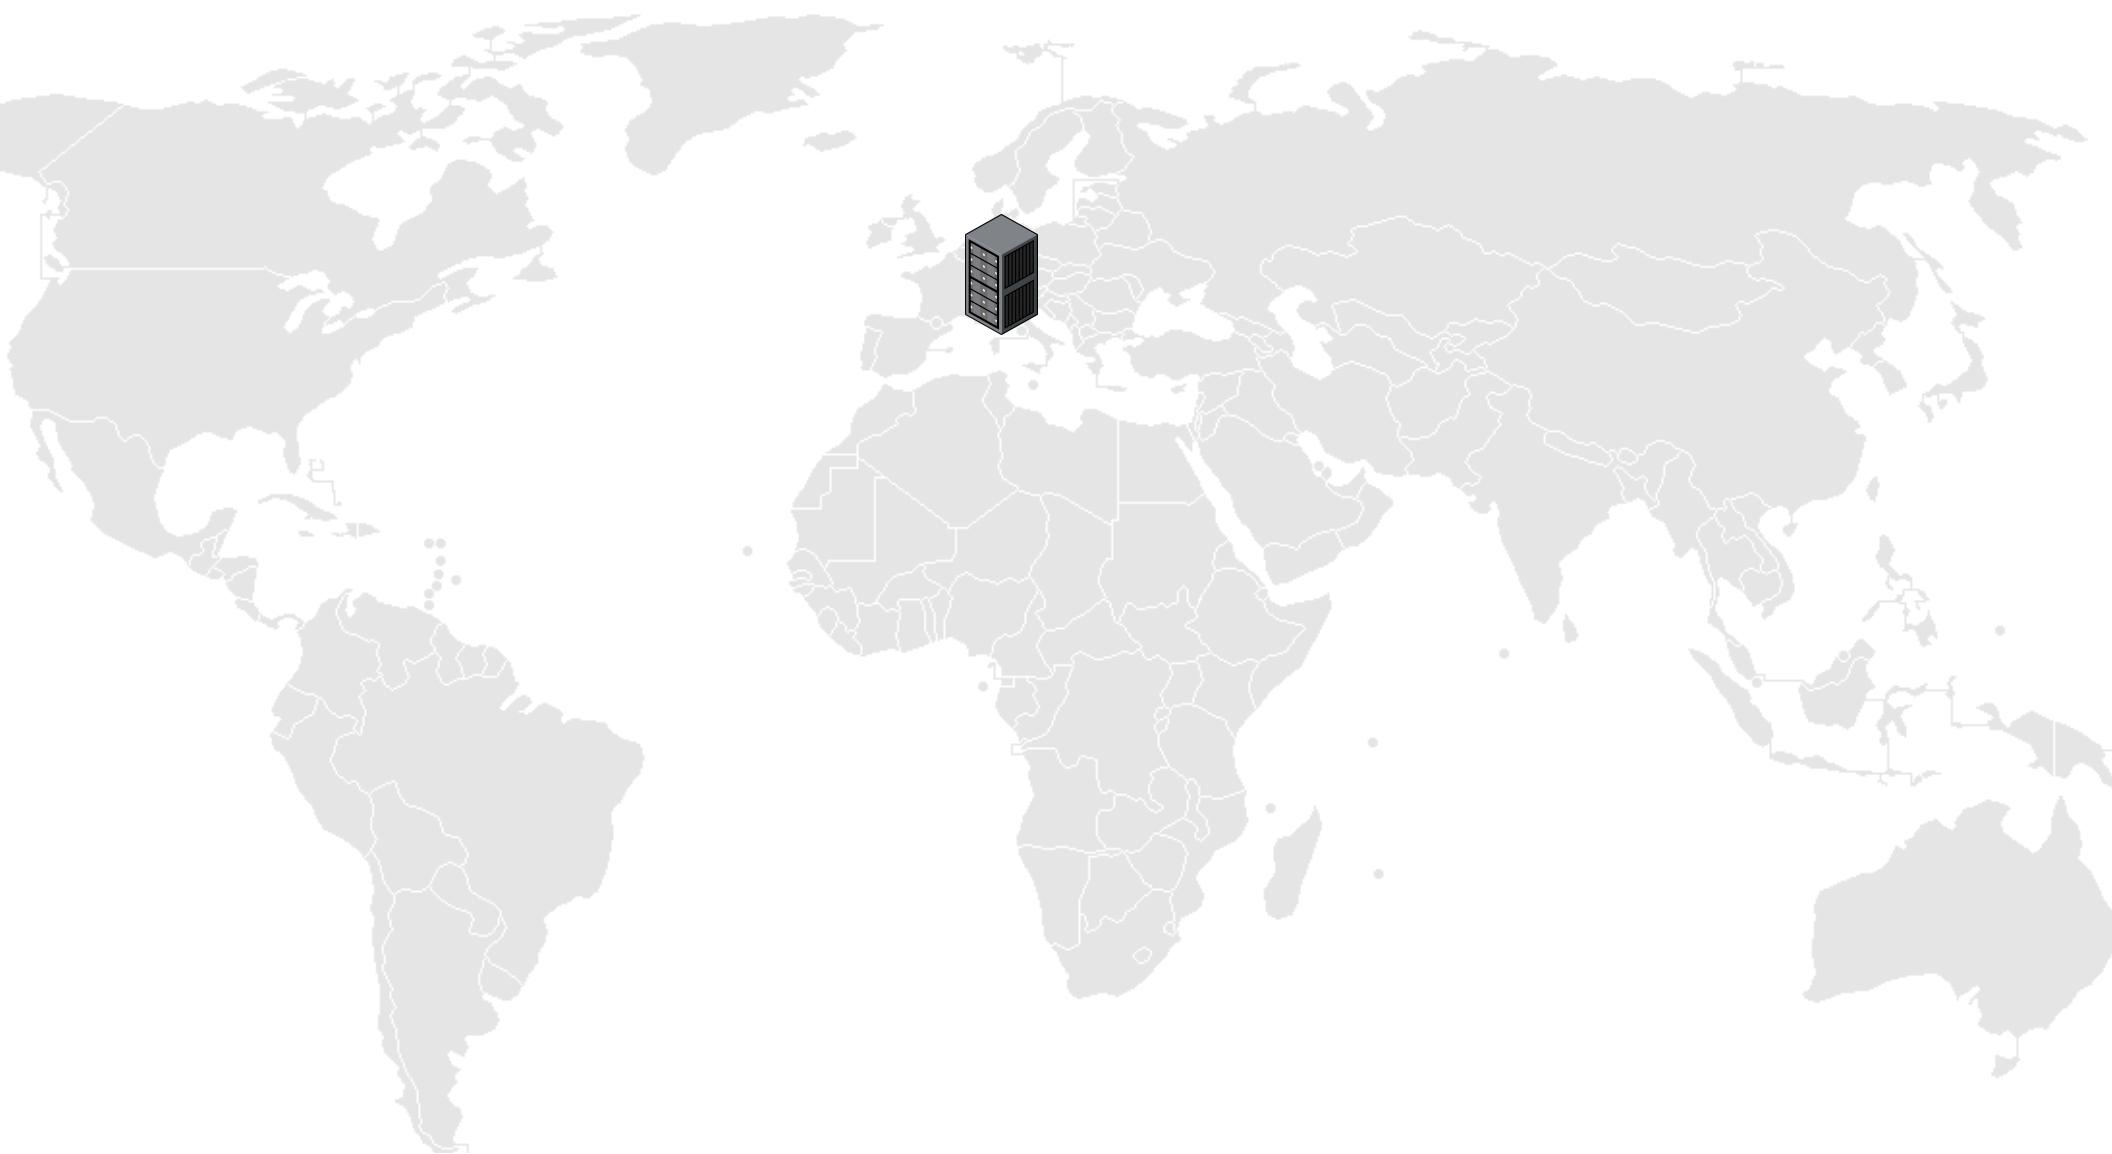
Distribution, redundancy, resilience, **performance**

## *Where*

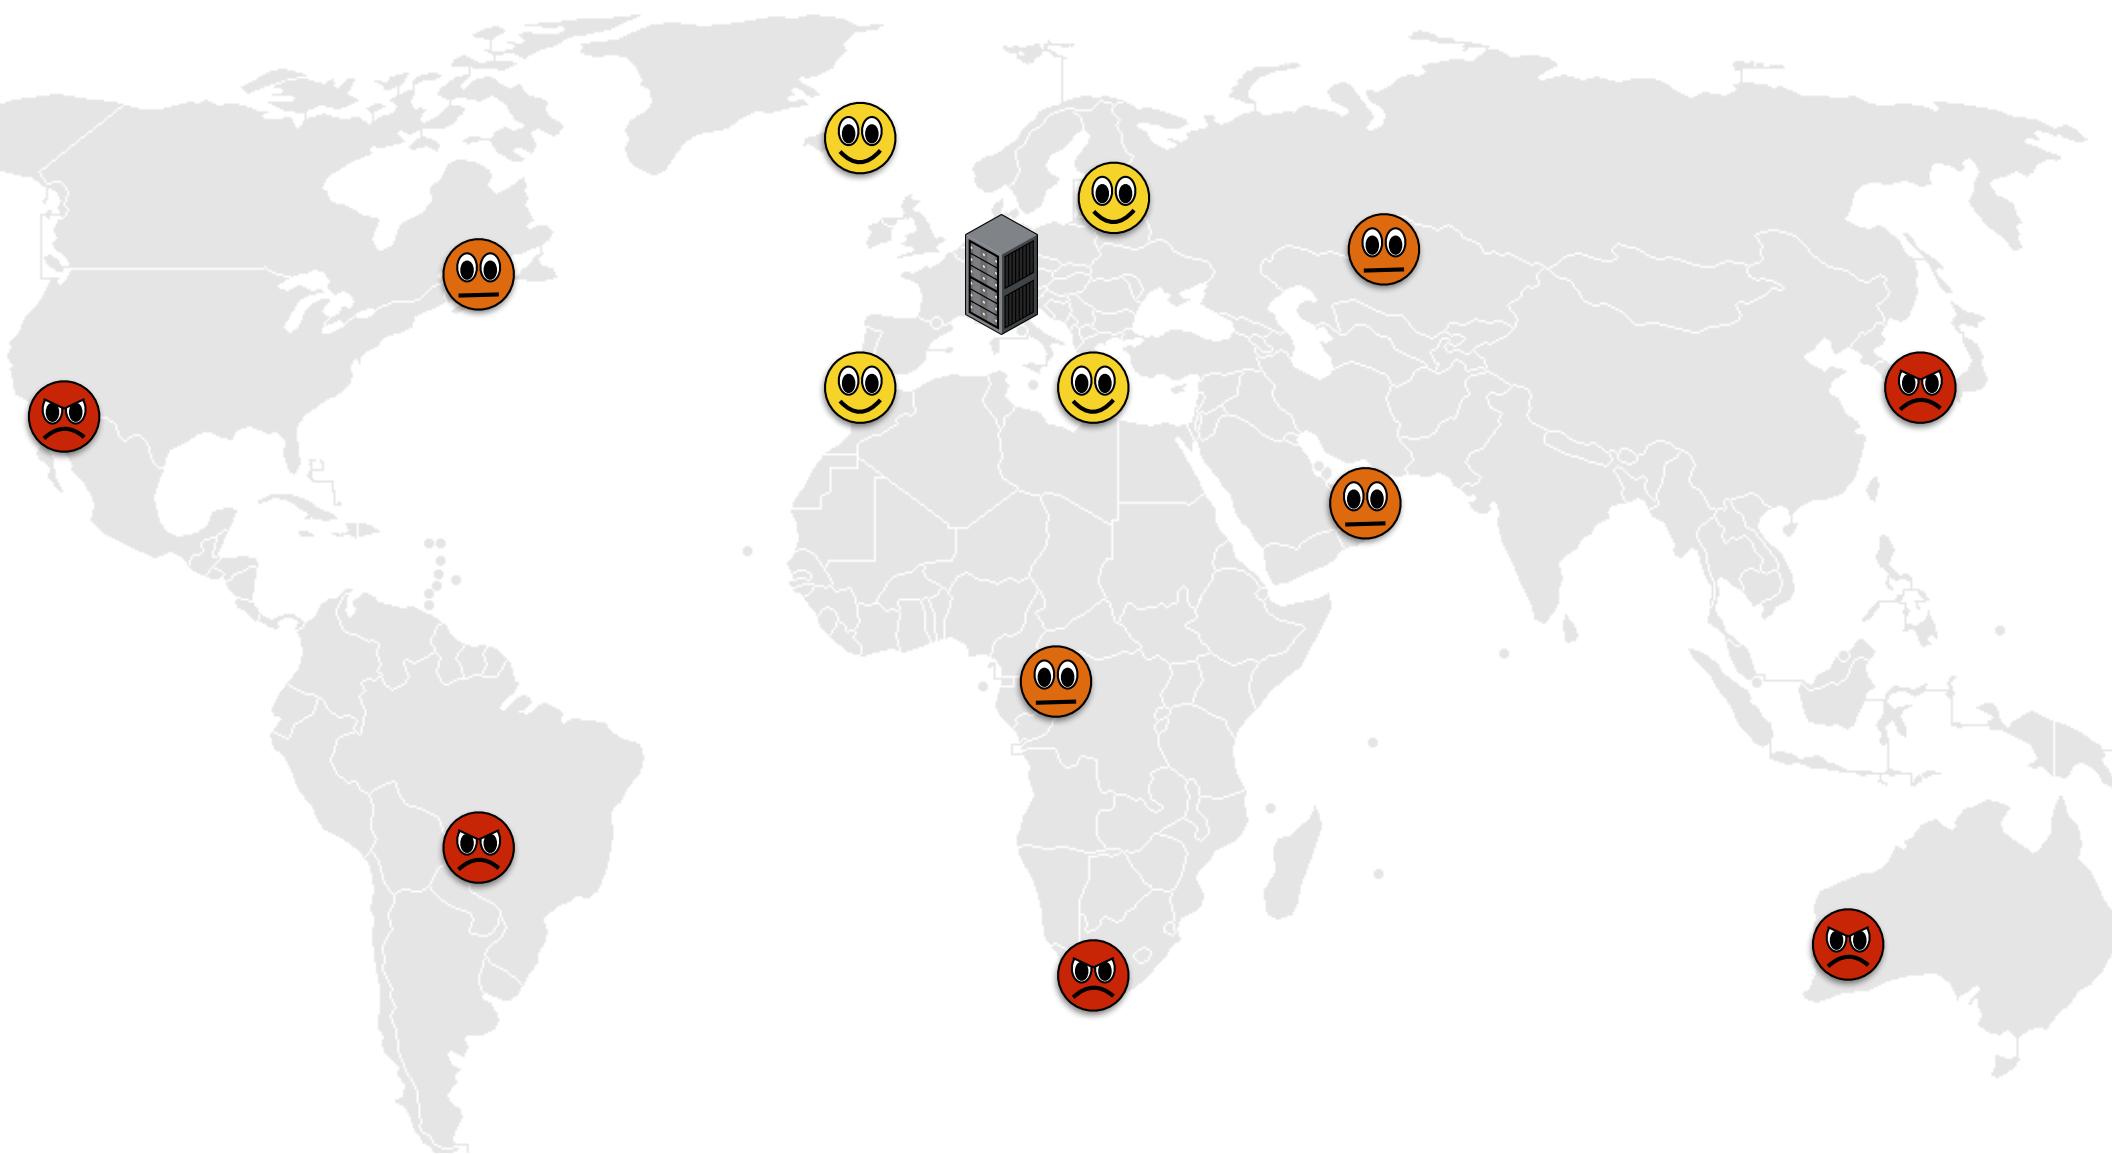
Content Delivery Networks (CDNs)

DNS at the root and TLD levels

# Anycast



# Anycast



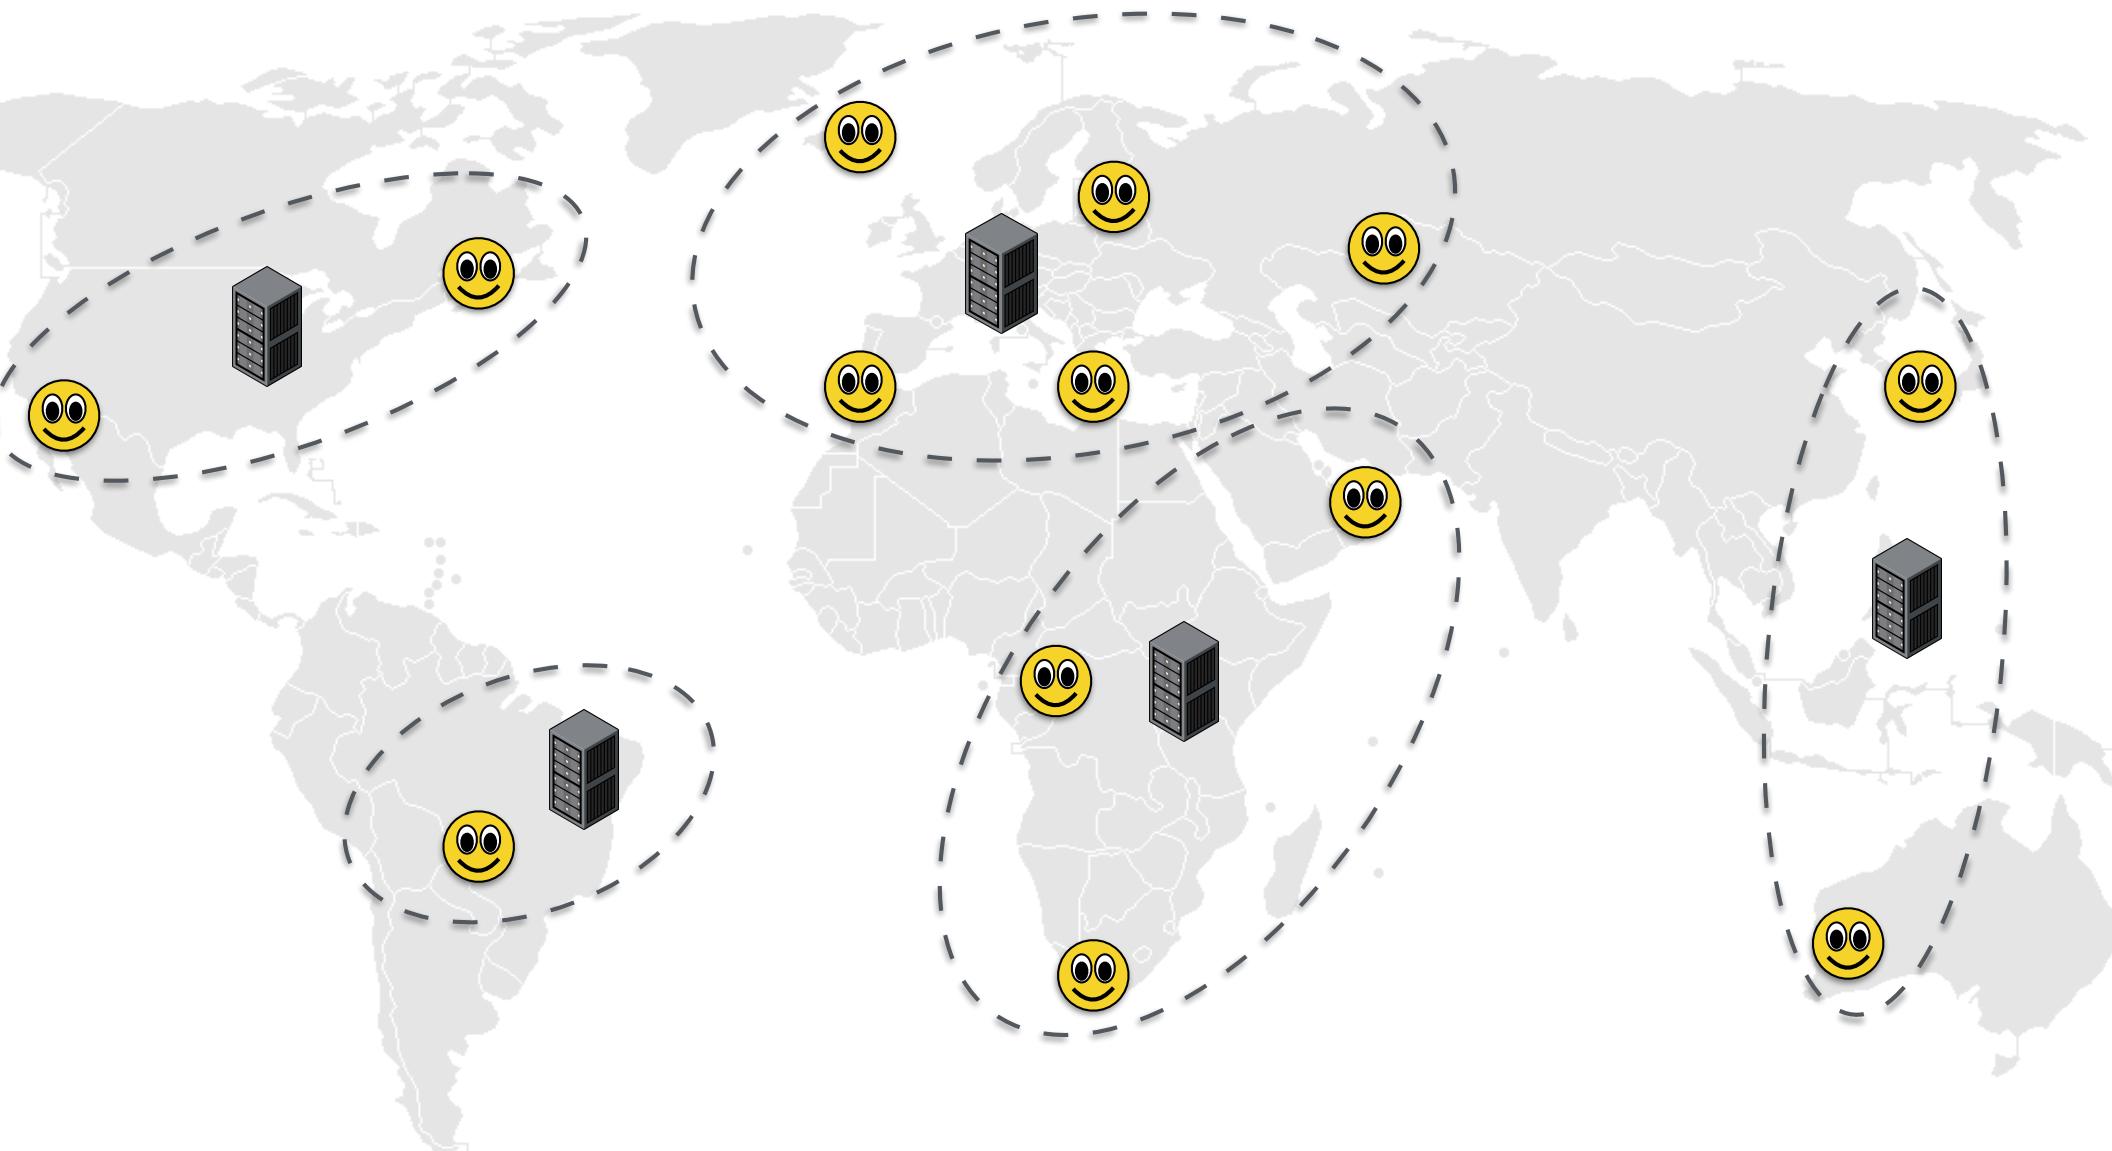
# Anycast



# Anycast

Anycast uses **BGP** to map users to sites

But BGP only **approximates closest match**



# RQ's and Approach

If BGP only approximates closest in anycast:

1. How good (or bad) is this approximation?
2. What is the impact of location of anycast sites?
3. What is the impact of routing policy?

**Can we determine a minimum number of anycast sites as optimal?**

To answer these questions we...

- ... study the relationship between latency and number of anycast sites
- ... look at the Root DNS as case of study

# Root DNS (or a part of it)

Our case of study were **four letters** of the Root DNS:

**C**: 8 sites

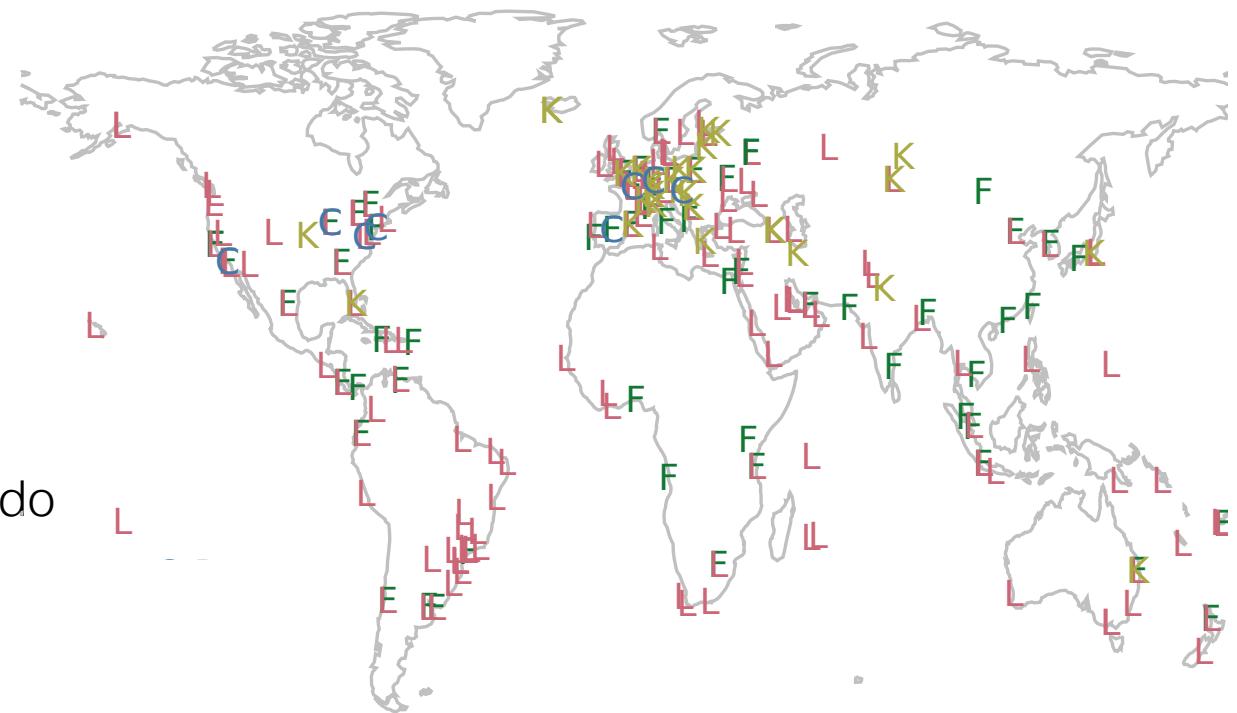
**F**: 58 sites

**K**: 33 sites

**L**: 144 sites

We **DO NOT** focus on:

What a letter should or not do



# RIPE Atlas

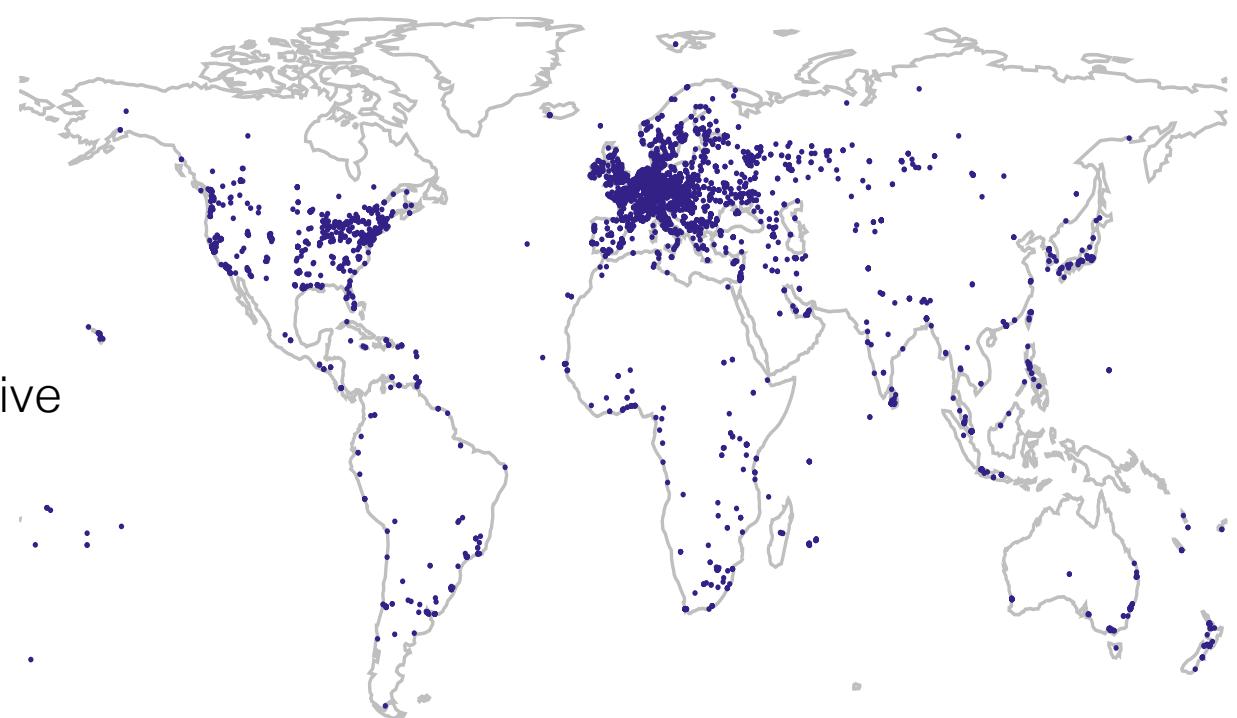
Our vantage points (VPs) are RIPE Atlas probes

7,900 of them

174 countries

2927 ASes

The bias towards Europe  
**does not** impact our qualitative  
results and conclusions

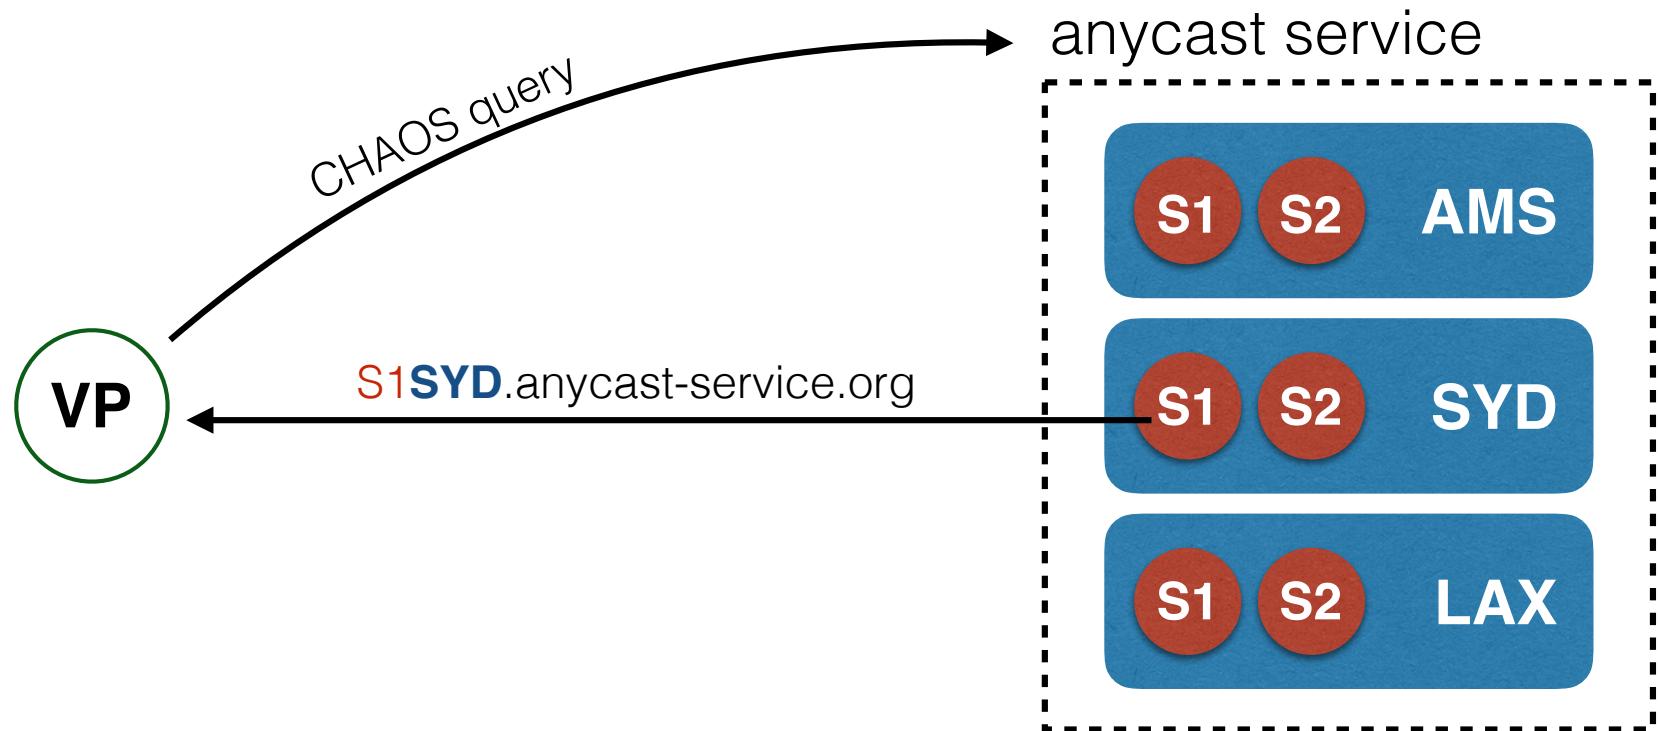


# Measurements

## Step 1: determine catchment

CHAOS queries to `hostname.bind` to anycast prefix...

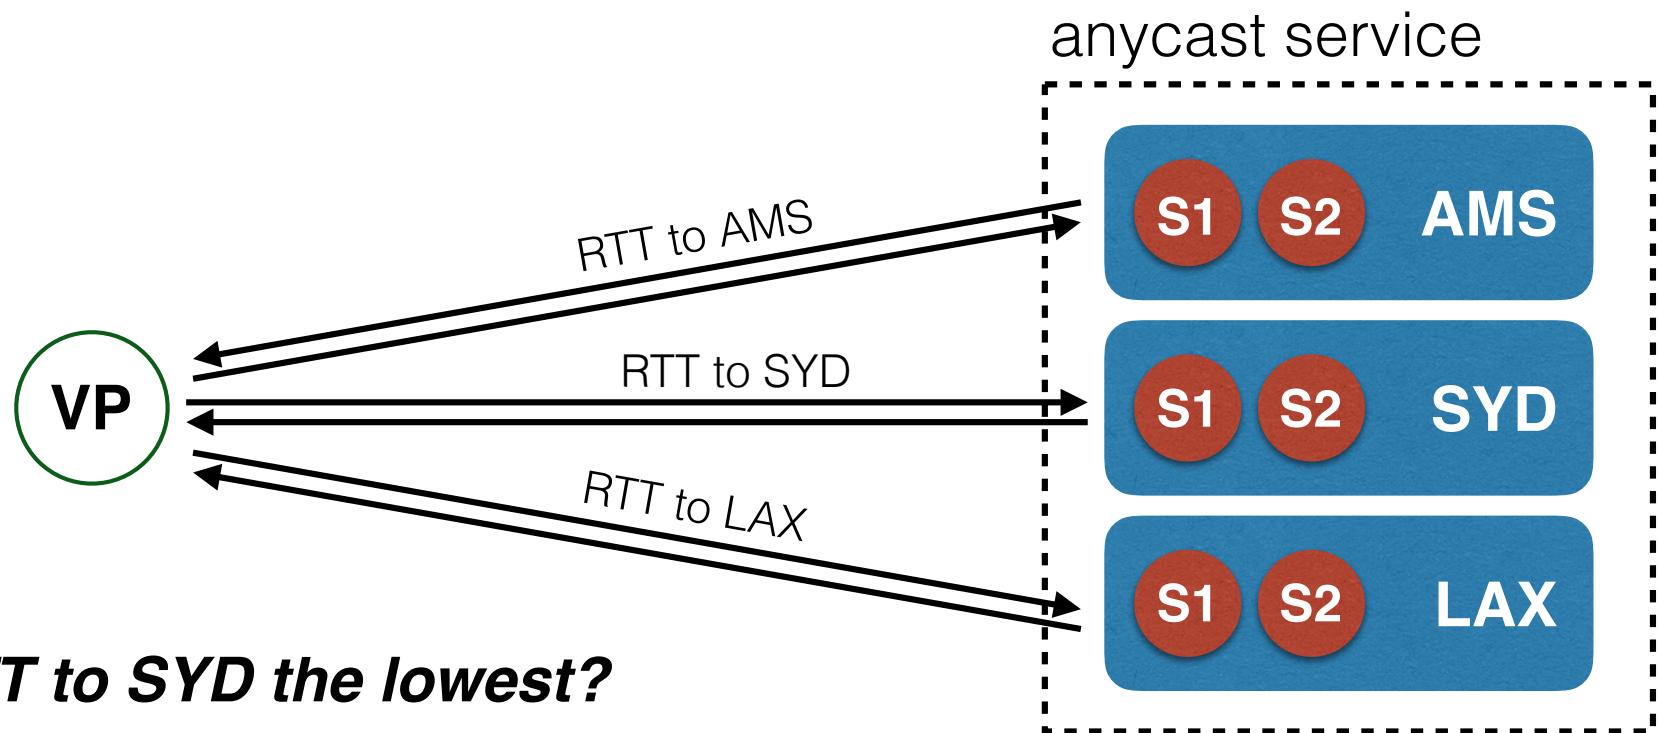
... returns a string that identifies a server



# Measurements

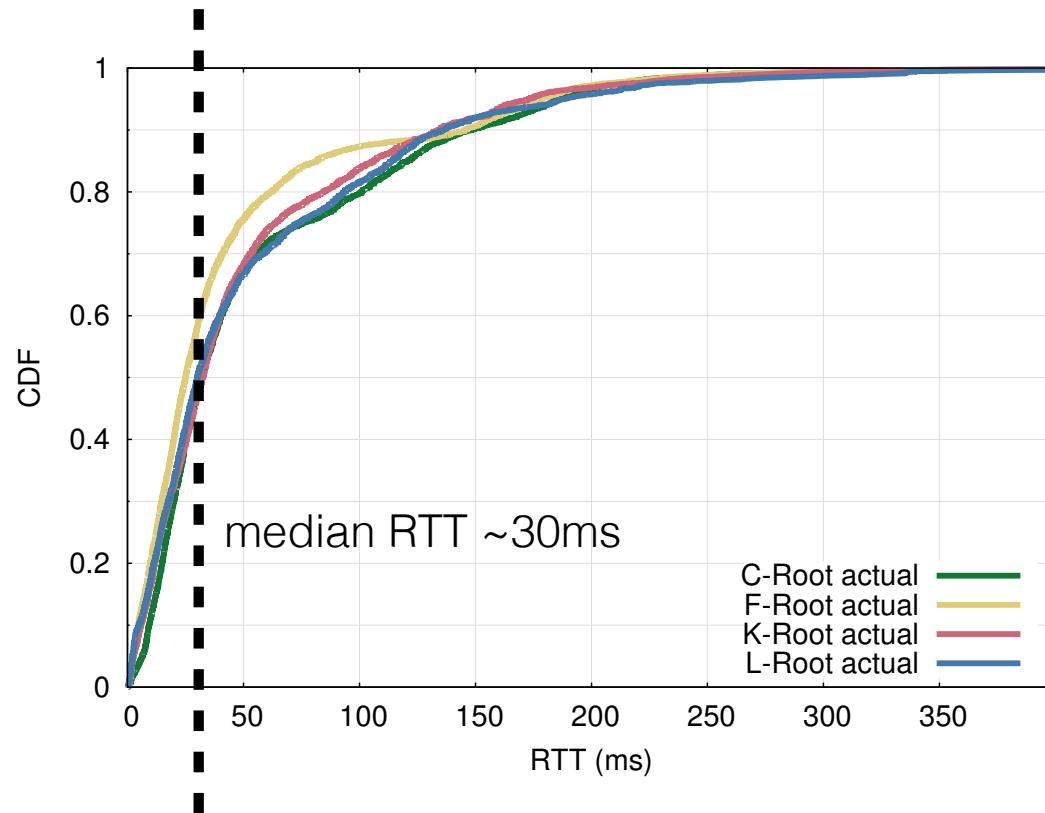
## Step 2: determine (optimal) latency

ICMP (ping) to all **sites** of the anycast...  
... as given by the identifier from CHAOS



# What **Performance** do we see?

Distribution of RTT for all letters:



We do not determine a threshold for good or bad performance...  
... but we do believe that 30ms is *very good*

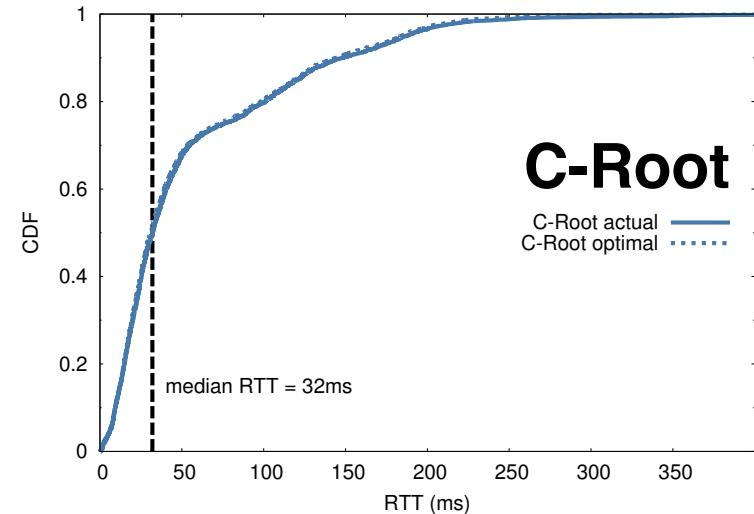
# What RTT is possible?

Comparing actual to optimal possible:

C-Root with 8 sites is at its optimal

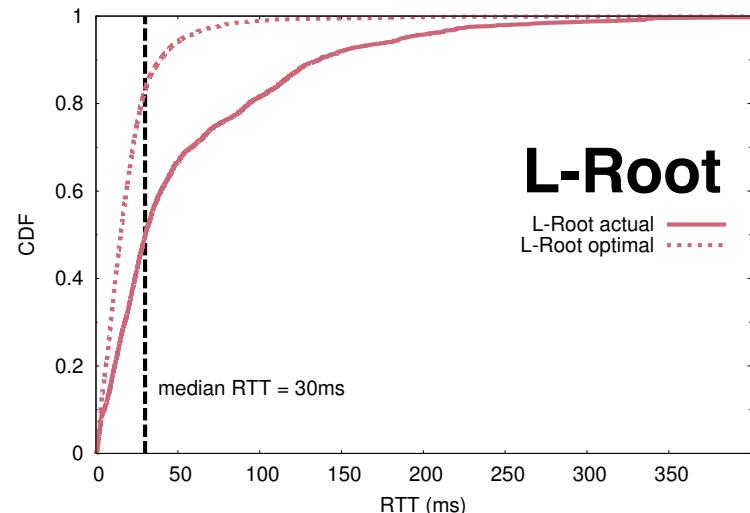
L-Root with 144 sites has plenty of room

But their median is roughly the same



Note:

Comparing anycast and unicast routing



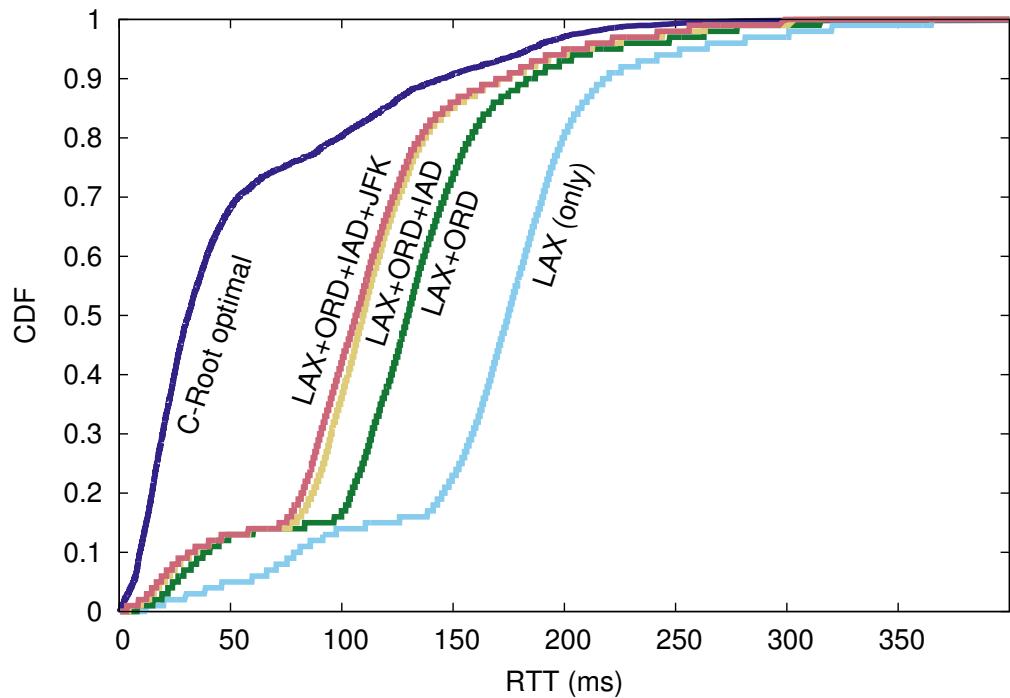
# What about **location?**

**Location matters!**

Simulation using C-Root sites  
from west to east coast

Note:

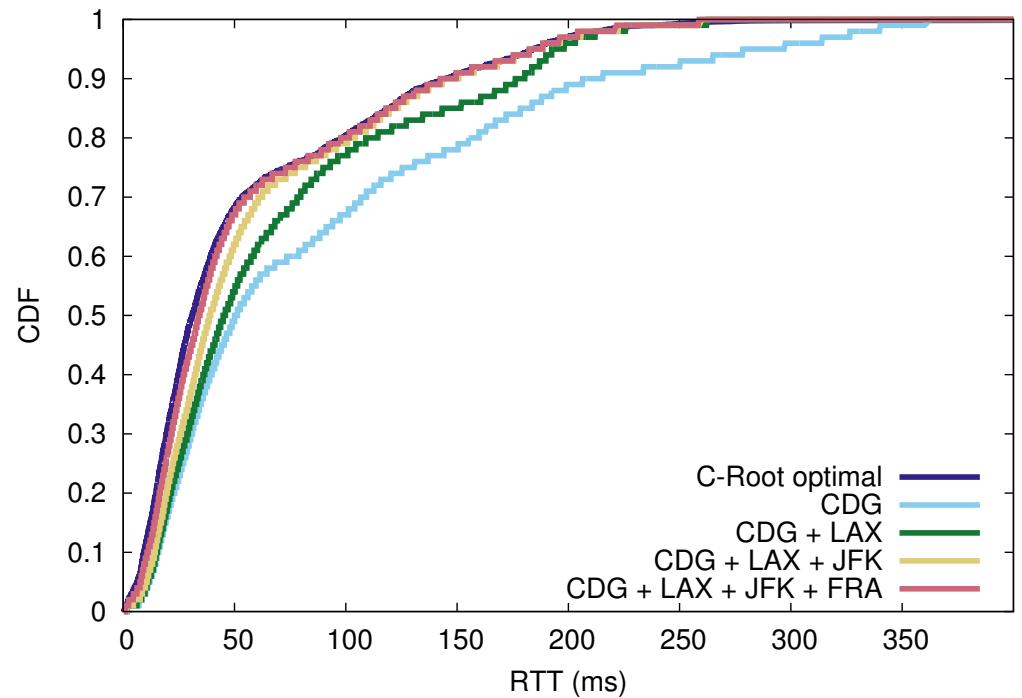
C-Root has sites in Europe  
Bias from RIPE Atlas is visible



# What about **location?**

**Location matters!**

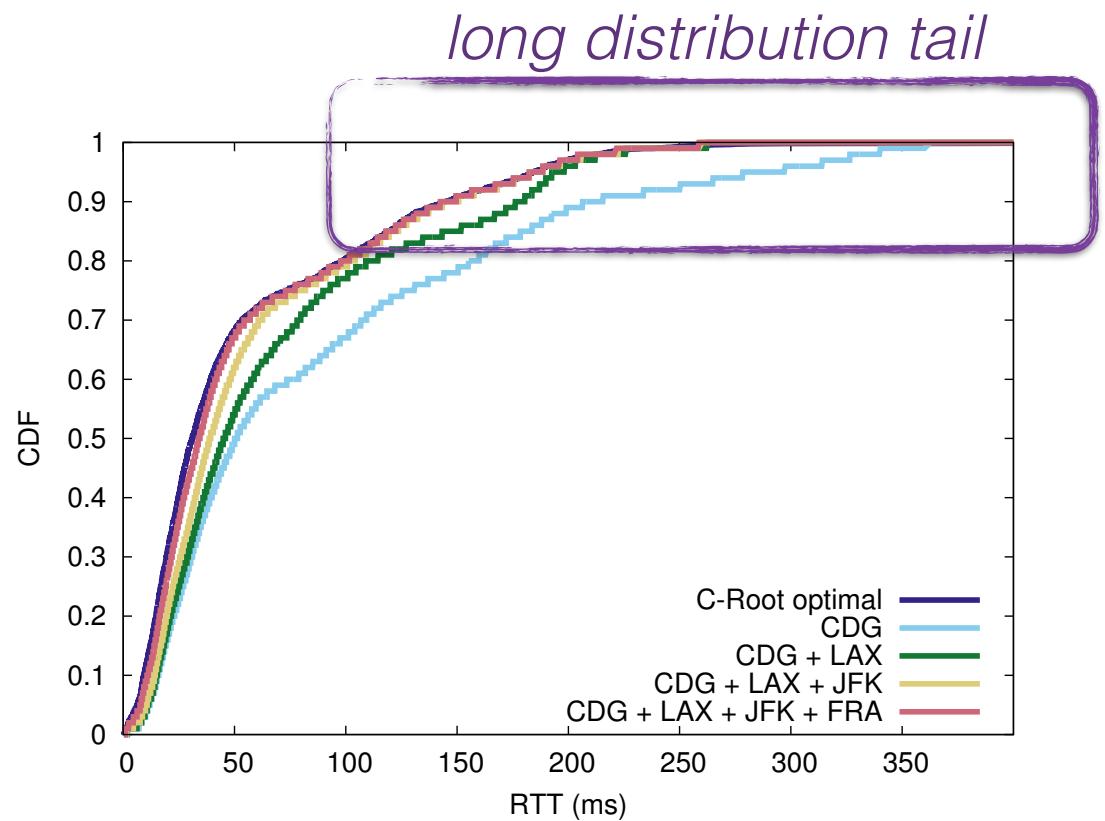
Simulation using C-Root sites  
from US and Europe



# What about **location?**

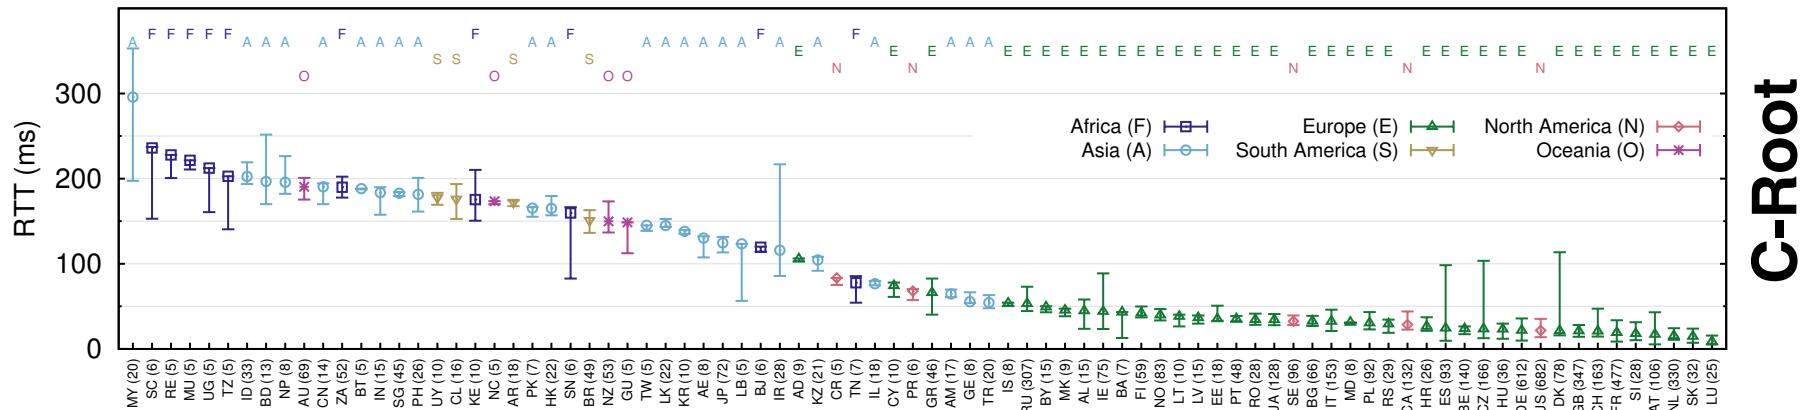
**Location matters!**

Simulation using C-Root sites  
from US and Europe



# Sites vs Location

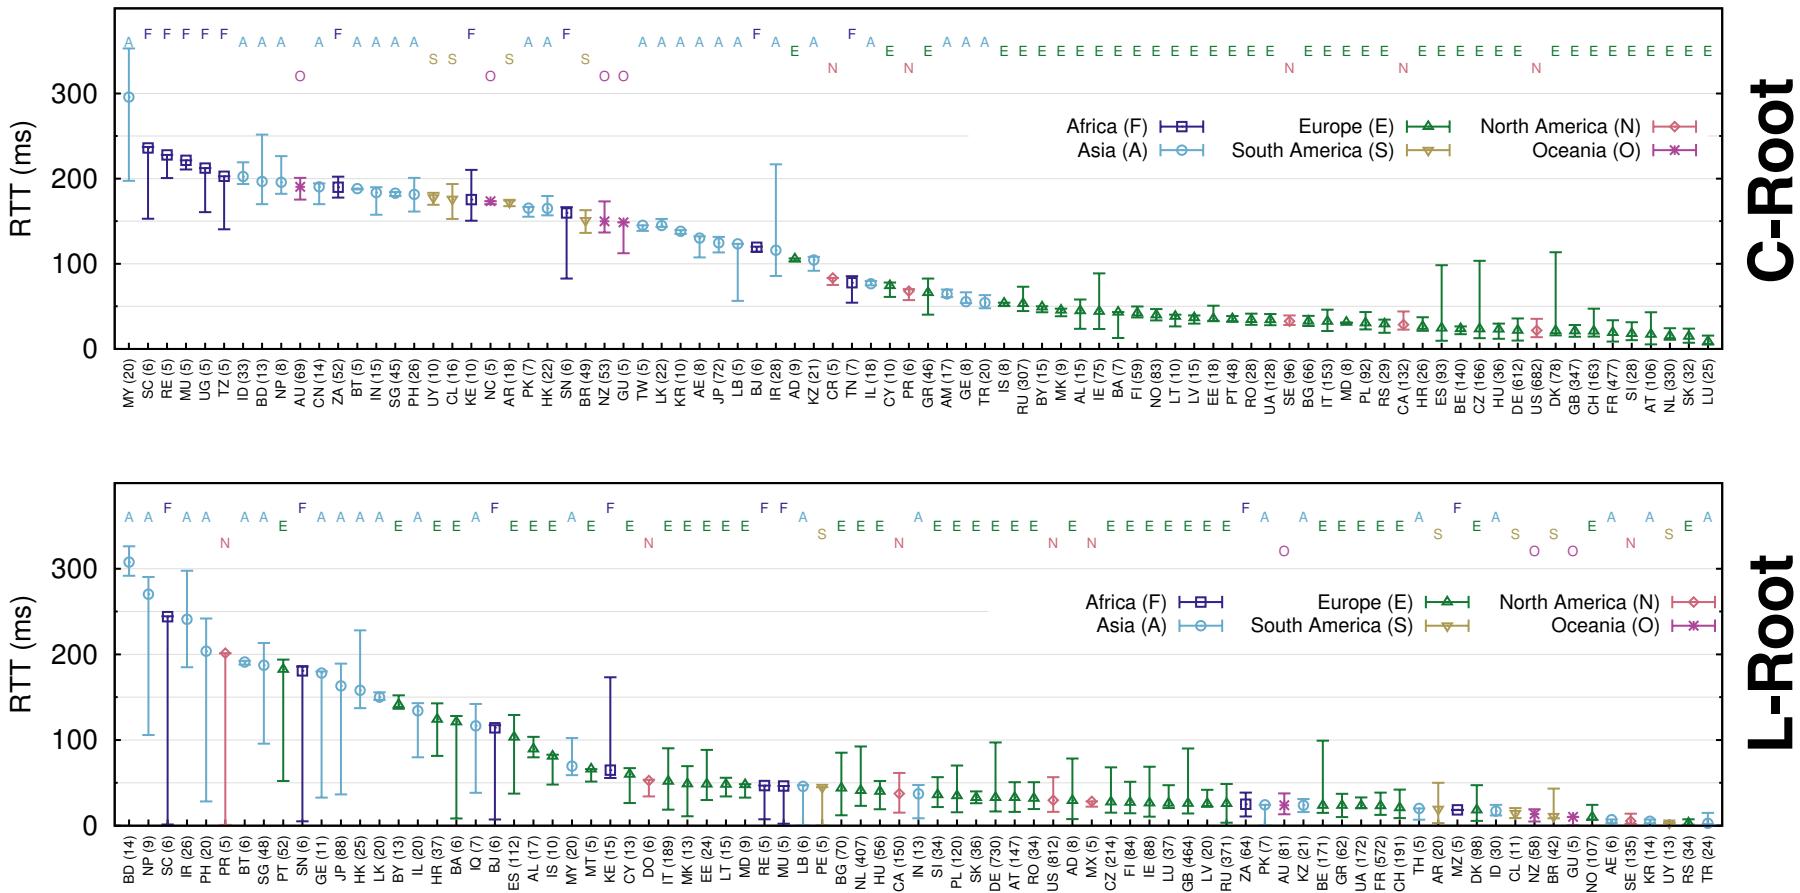
So, at the end many sites do help those at the end of the distribution tail...  
 ... but it depends where and how connected



C-Root

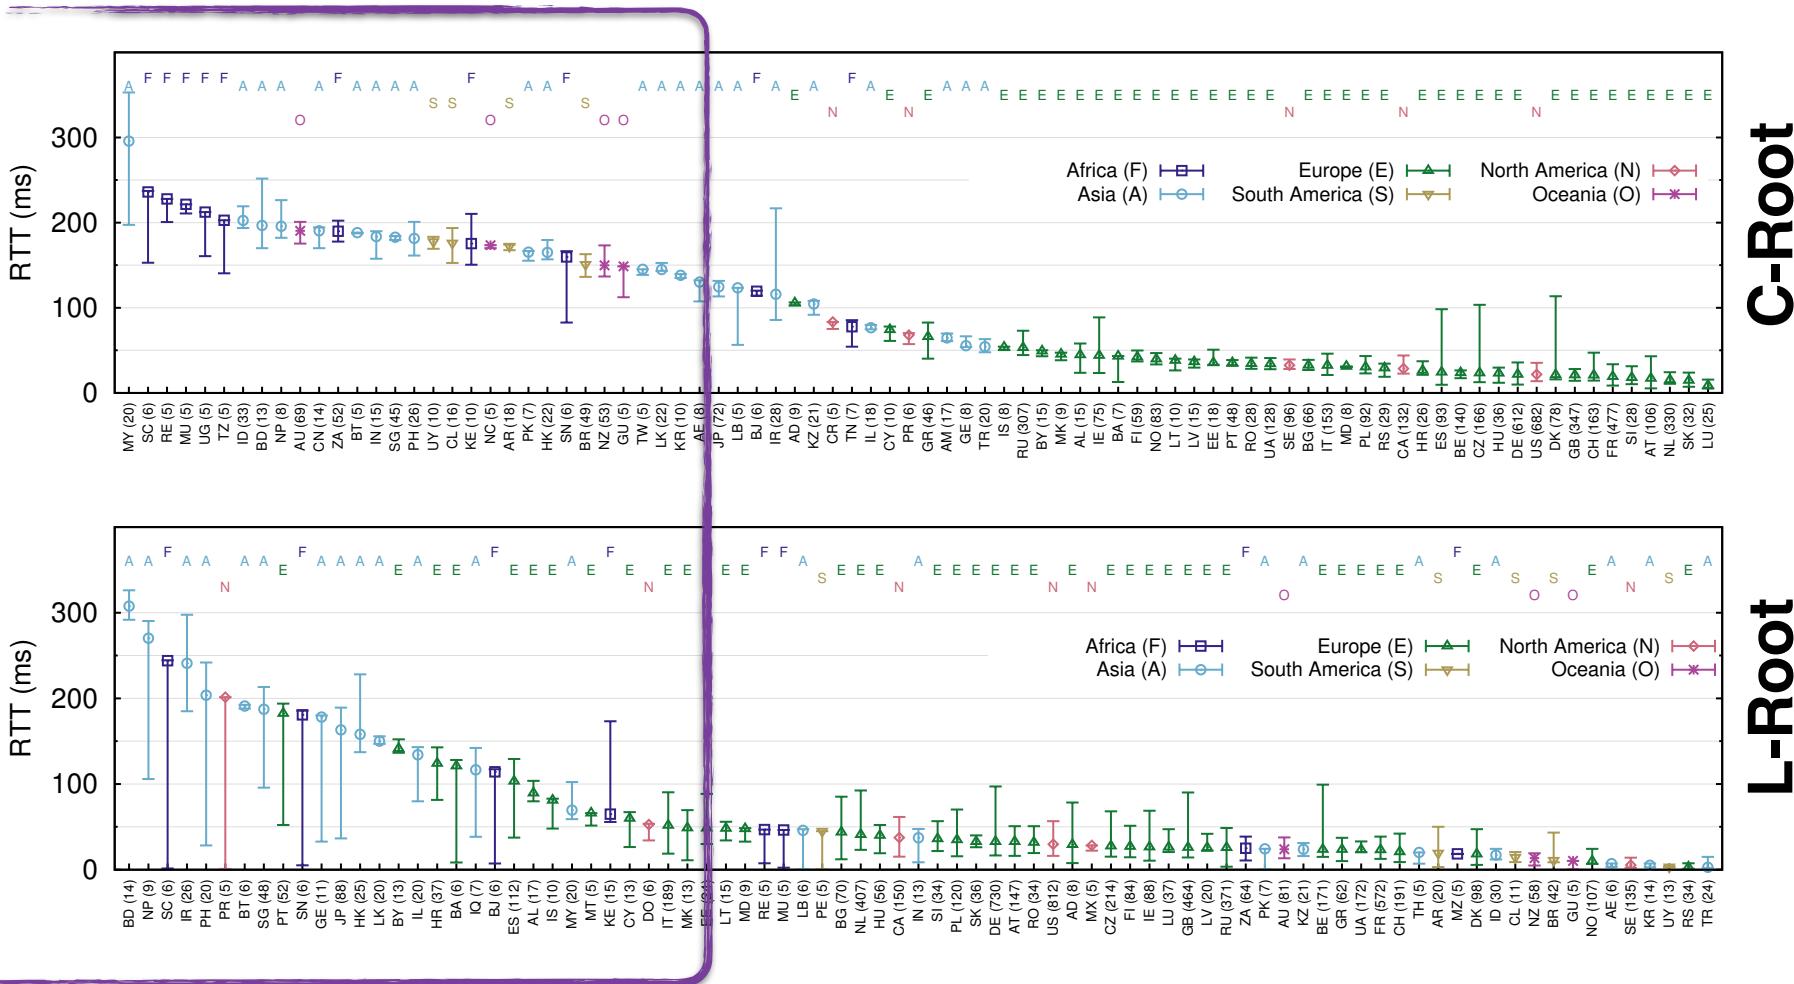
# Sites vs Location

So, at the end many sites do help those at the end of the distribution tail...  
 ... but it depends where and how connected



# Sites vs Location

So, at the end many sites do help those at the end of the distribution tail...  
... but it depends where and how connected



# Final Considerations

Looking at real-world anycast deployments we learned that...

... number of sites can make a difference on performance, but...

**... *location and connectivity of sites have higher impact***

... we believe that ~12 well-connected sites is "enough" for *performance*

For the future, we will focus on other purposes of anycast:

Resilience to Denial of Service attacks

Load balancing

# Final Considerations

Looking at real-world anycast deployments we learned that...

... number of sites can make a difference on performance, but...

**... *location and connectivity of sites have higher impact***

... we believe that ~12 well-connected sites is "enough" for *performance*

For the future, we will focus on other purposes of anycast:

Resilience to Denial of Service attacks

Load balancing

*in about 20 minutes  
(depending on Wouter)*

[r.schmidt@utwente.nl](mailto:r.schmidt@utwente.nl)  
<http://www.ricardoschmidt.com/>

## Acknowledgements:

Geoff Huston, George Michaelson, Ray Bellis, Cristian Hesselman, Benno Overeinder, Jaap Akkerhuis, Duane Wessels, Paul Vixie, Romeo Zwart, Anand Buddhdev, and C-Root operators.

SIDN Labs, NLnet Labs and SURFnet

*Self-managing Anycast Networks for the DNS (SAND) project | <http://www.sand-project.nl/>*  
*DNS Anycast Security (DAS) project | <http://www.das-project.nl/>*

UNIVERSITY OF TWENTE.

