UNIVERSITY OF TWENTE.

DACS
Design and Analysis of Communication Systems

USC Viterbi
School of Engineering

ant.isi.edu

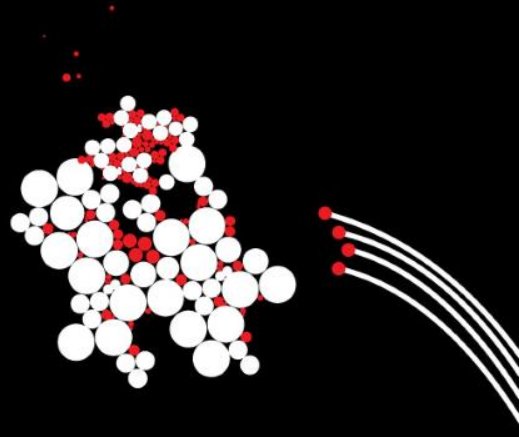

BROAD AND LOAD-AWARE ANYCAST MAPPING WITH VERFPLOETER

WOUTER B. DE VRIES, RICARDO DE O. SCHMIDT, WES HARDAKER, JOHN HEIDEMANN, PIETER-TJERK DE BOER AND AIKO PRAS

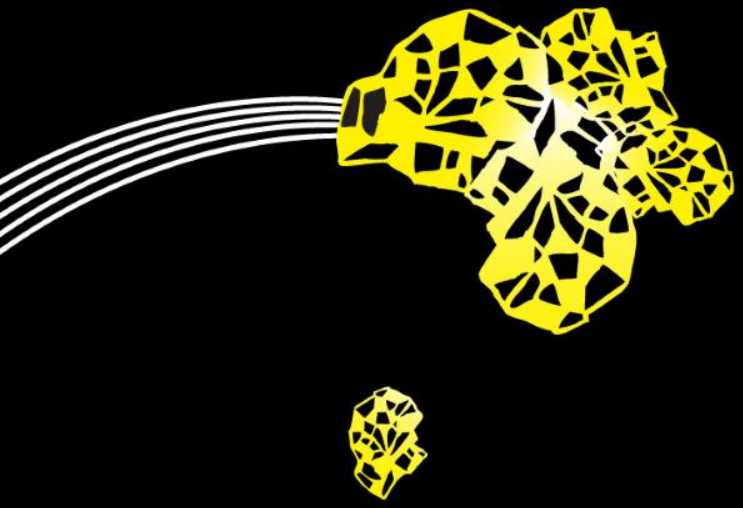

London - November 3, 2017

INTRODUCTION

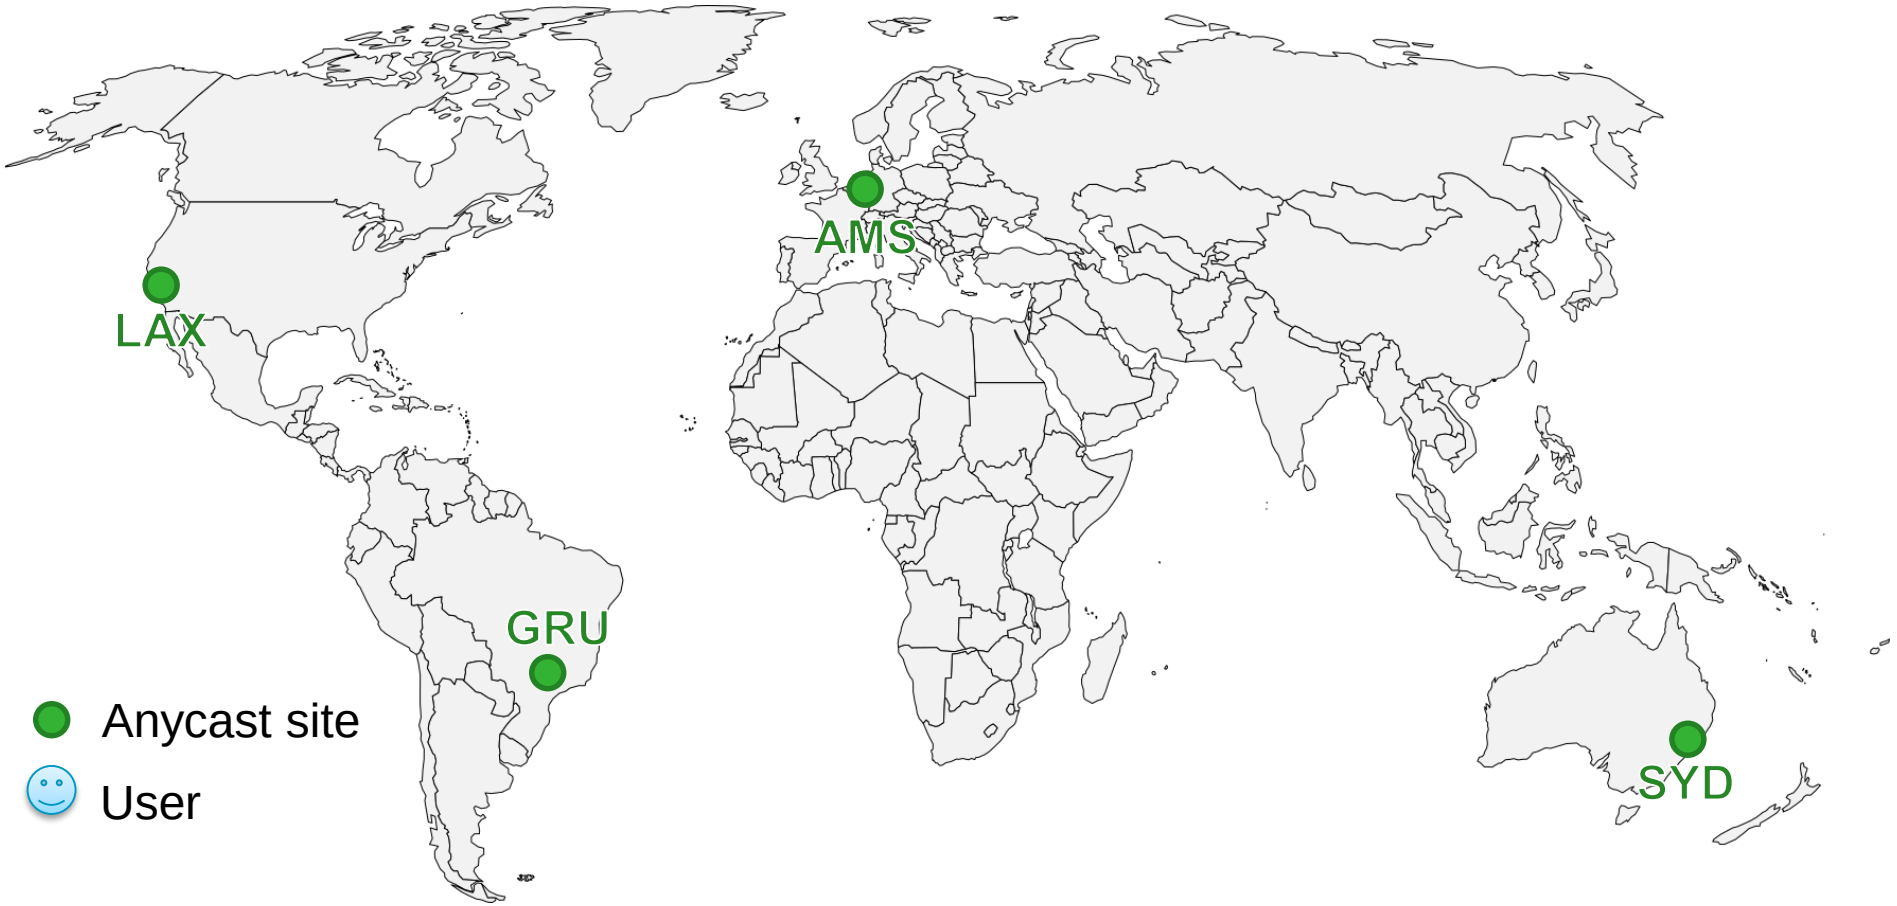

INTRODUCTION

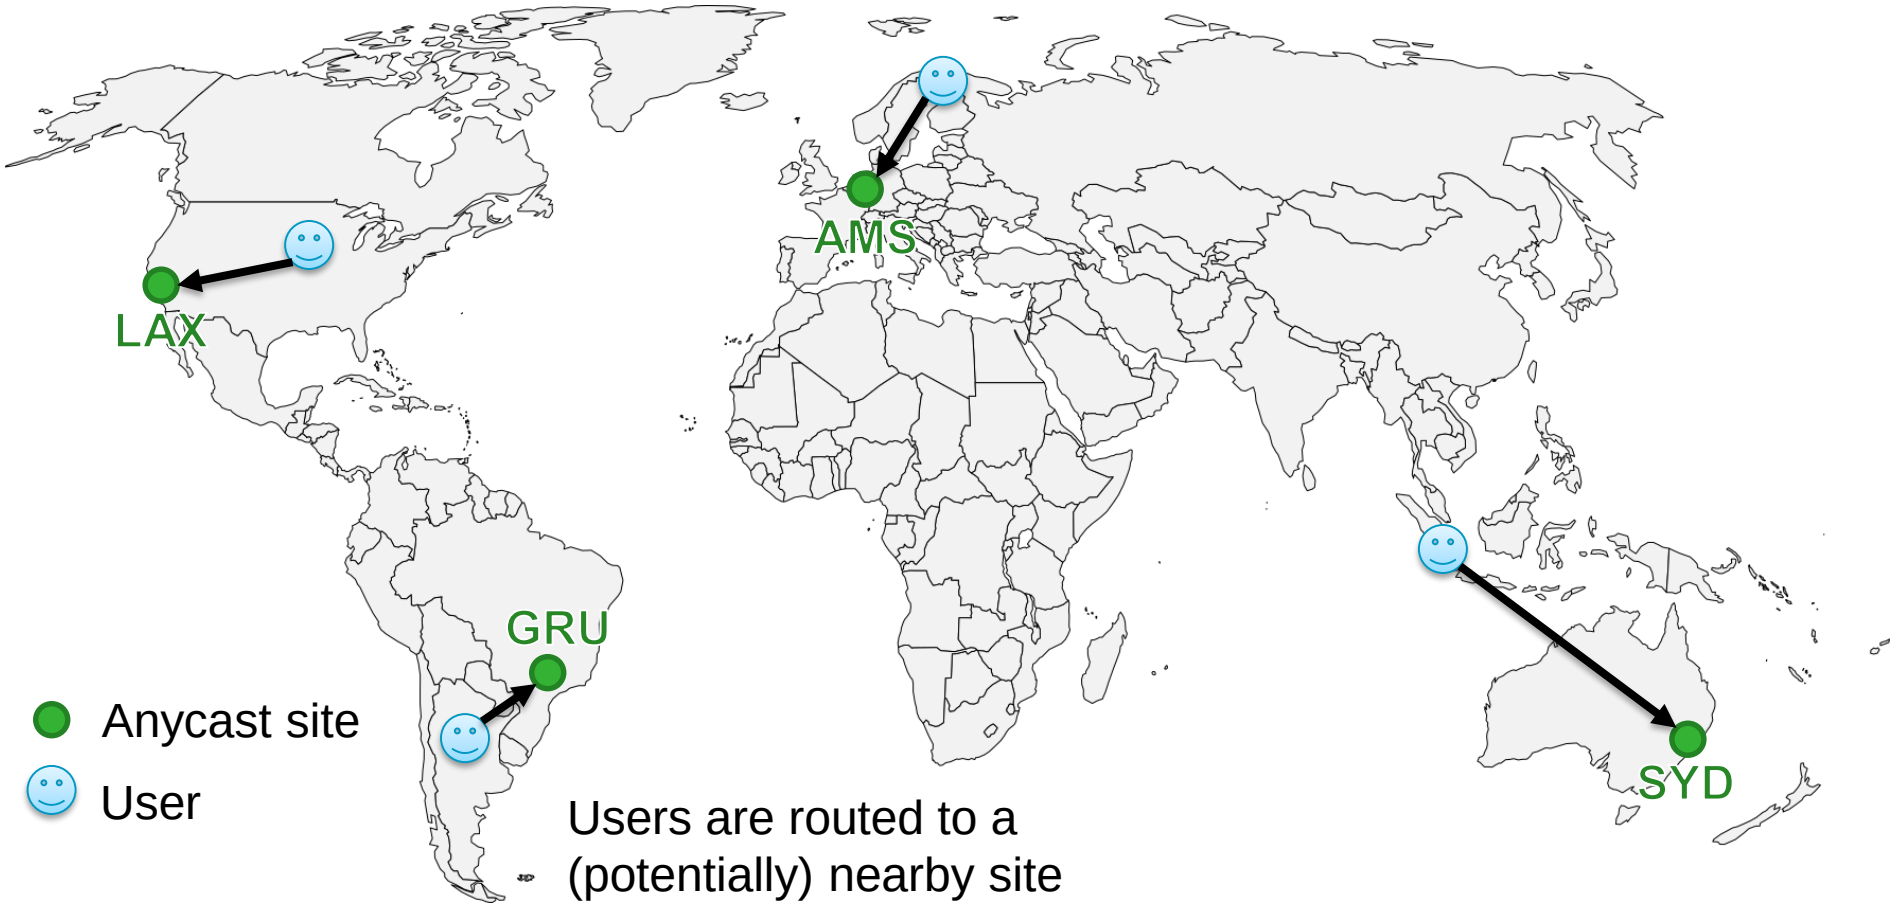

INTRODUCTION

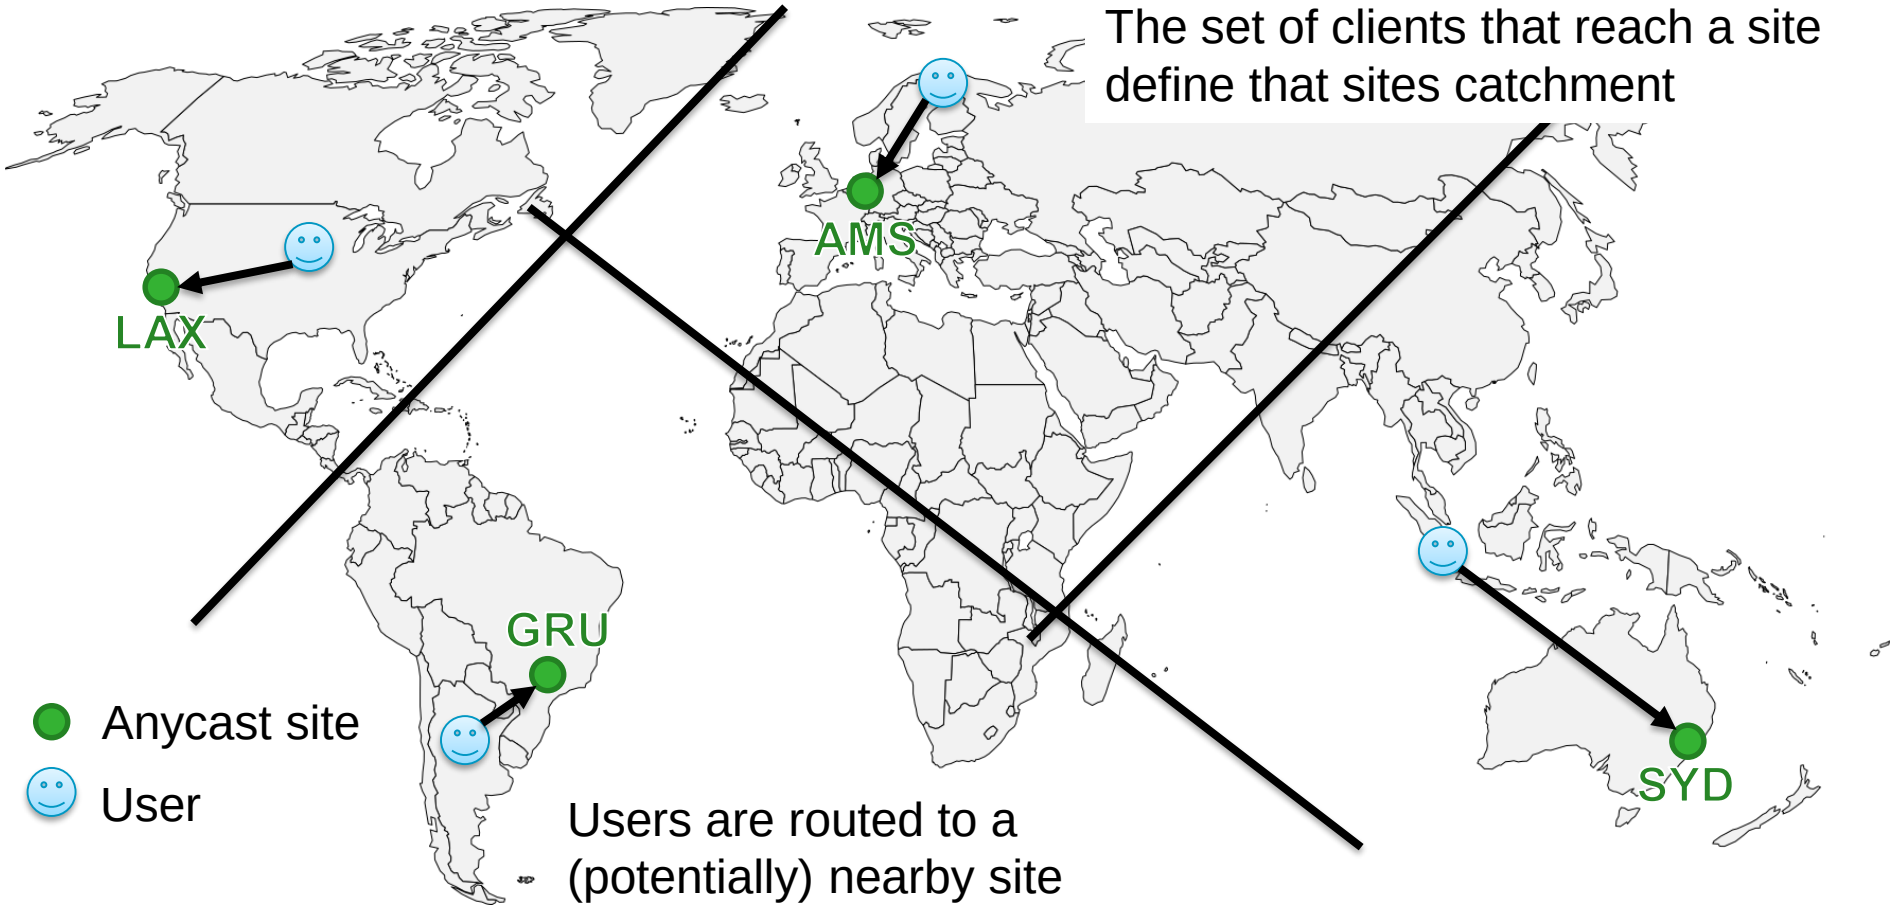

INTRODUCTION

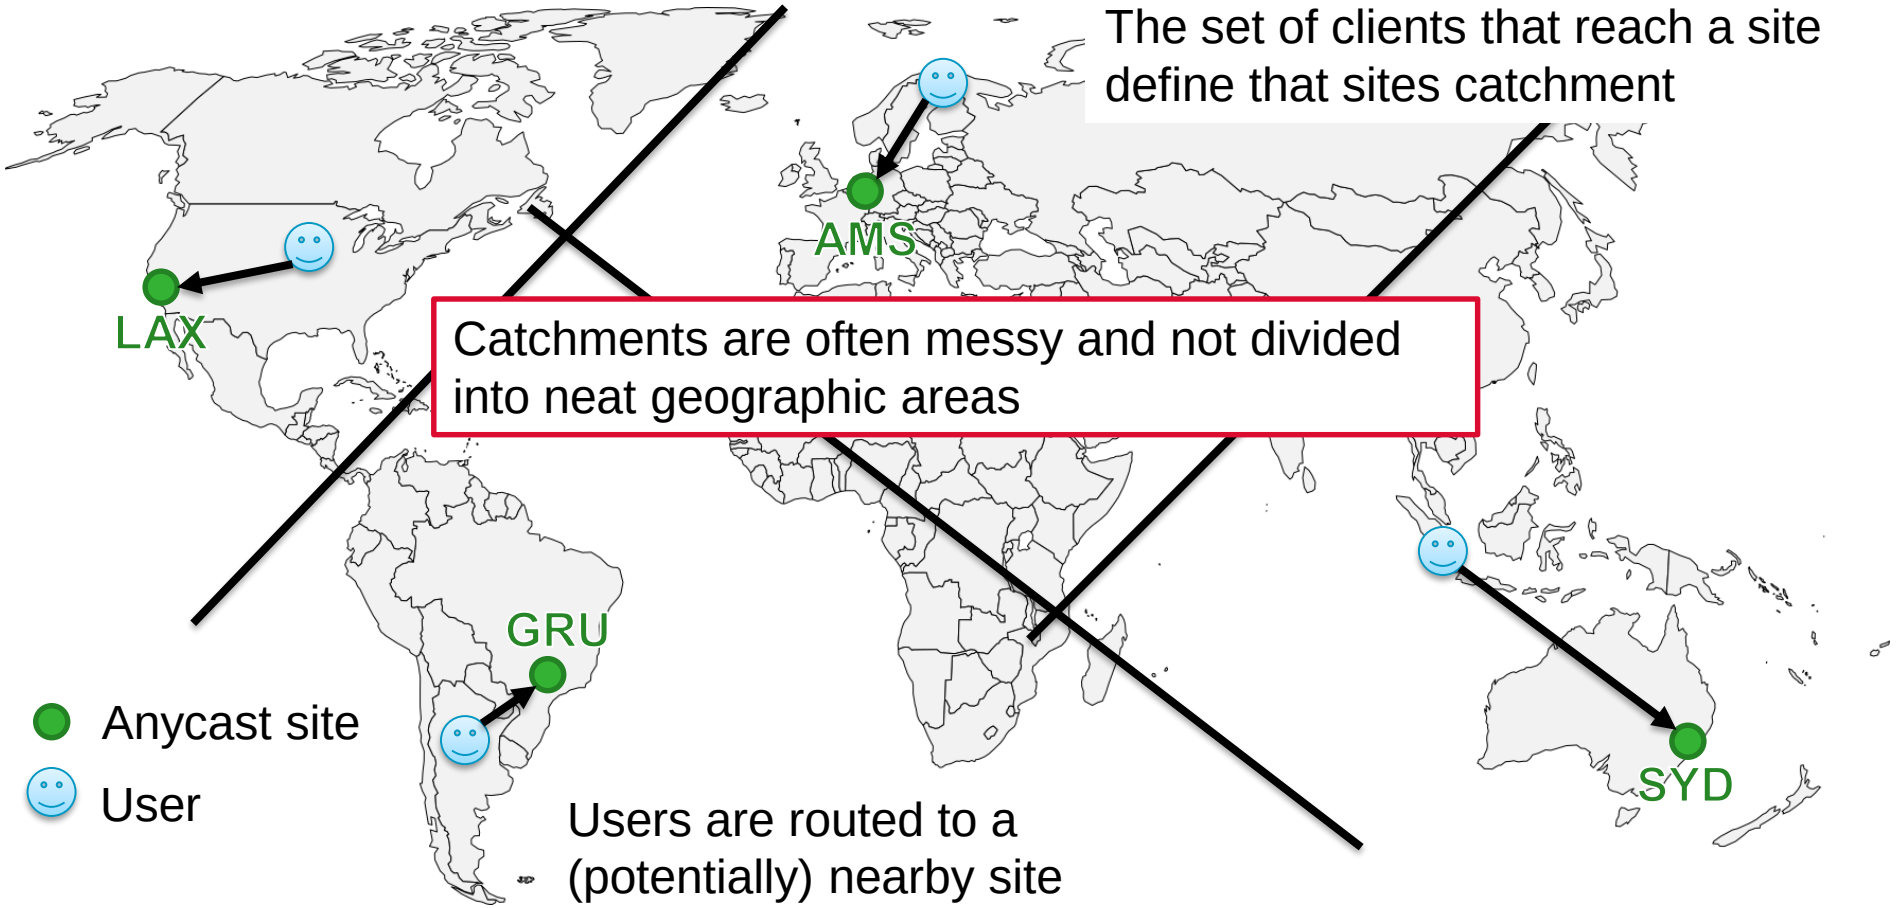

INTRODUCTION

Peering agreements, Transits and BGP policies influence the routing process.

You don’t know where a client from a certain prefix will be routed *until* he is routed there.

INTRODUCTION

Peering agreements, Transits and BGP policies influence the routing process.

You don't kn prefix will be routed *uni*

Suboptimal routing can lead to severely degraded performance, in terms of RTT, load balancing and resilience

DETERMINING CATCHMENT - TODAY

- Analyze log files
  - Only if you have them
- Use active probing (e.g. RIPE Atlas)
  - Requires VPs at many locations

DETERMINING CATCHMENT - TODAY

- Analyze log files
  - Only if you have them
- Use active probing (e.g. RIPE Atlas)
  - Requires VPs at many locations

NL: 563 VPs

DETERMINING CATCHMENT - TODAY

- Analyze log files
  - Only if you have them
- Use active probing (e.g. RIPE Atlas)
  - Requires VPs at many locations

NL: 563 VPs

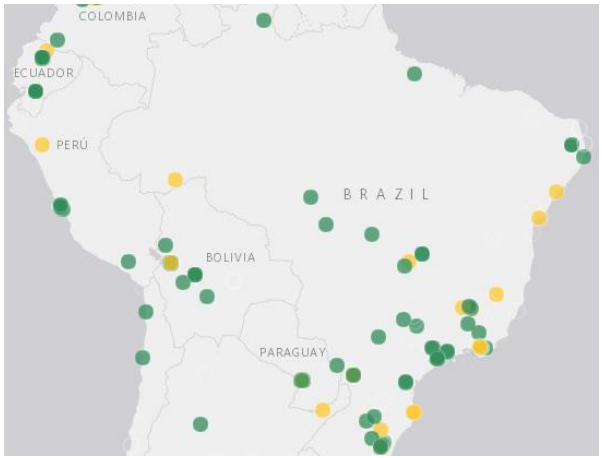

Brazil: 63 VPs

China: 19 VPs

GOALS

- Develop a method to accurately map anycast catchments
  - Without requiring input from users
- “Calibrate” the anycast catchment to the actual client base of the service

VERFPLOETER: METHODOLOGY

ACTIVE MEASUREMENT WITH “PASSIVE” VANTAGE POINTS

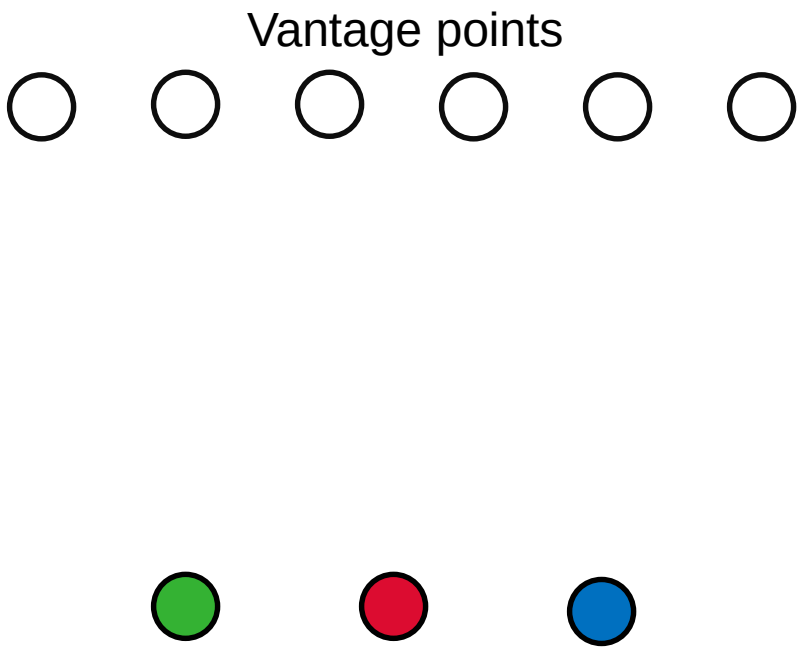

VERFPLOETER: METHODOLOGY

ACTIVE MEASUREMENT WITH “PASSIVE” VANTAGE POINTS

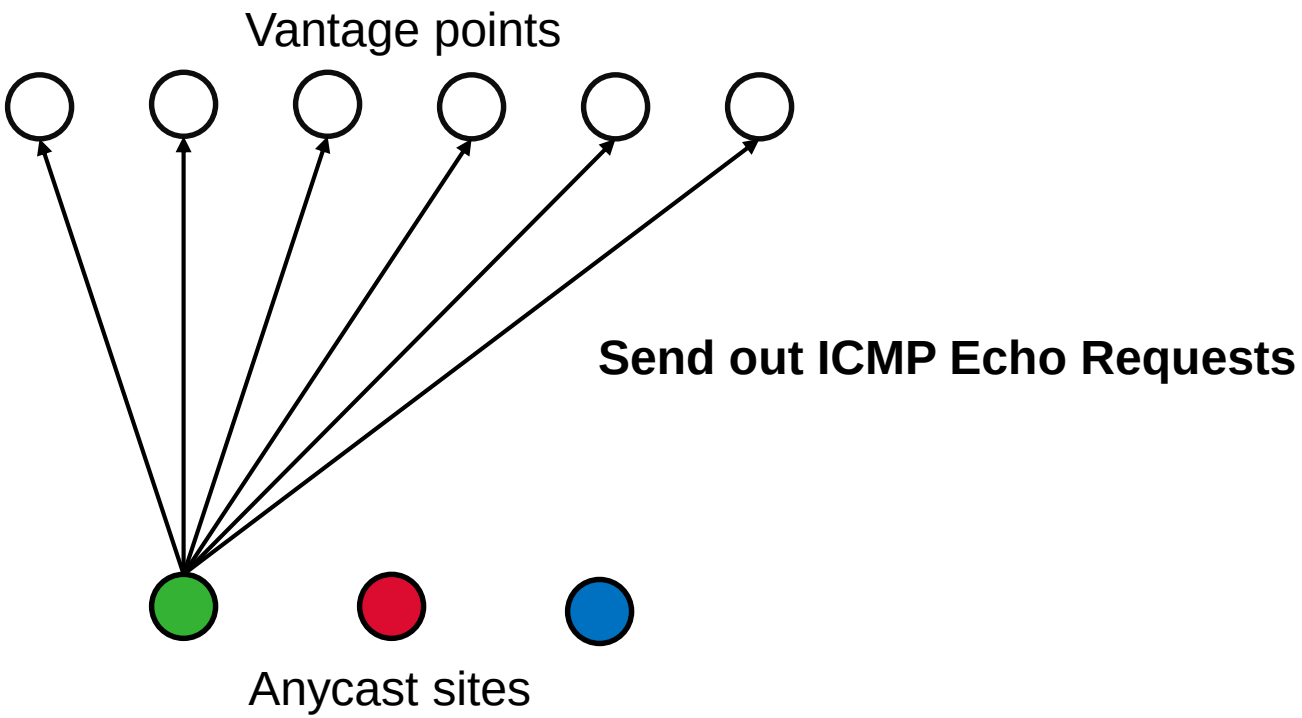

VERFPLOETER: METHODOLOGY

ACTIVE MEASUREMENT WITH “PASSIVE” VANTAGE POINTS

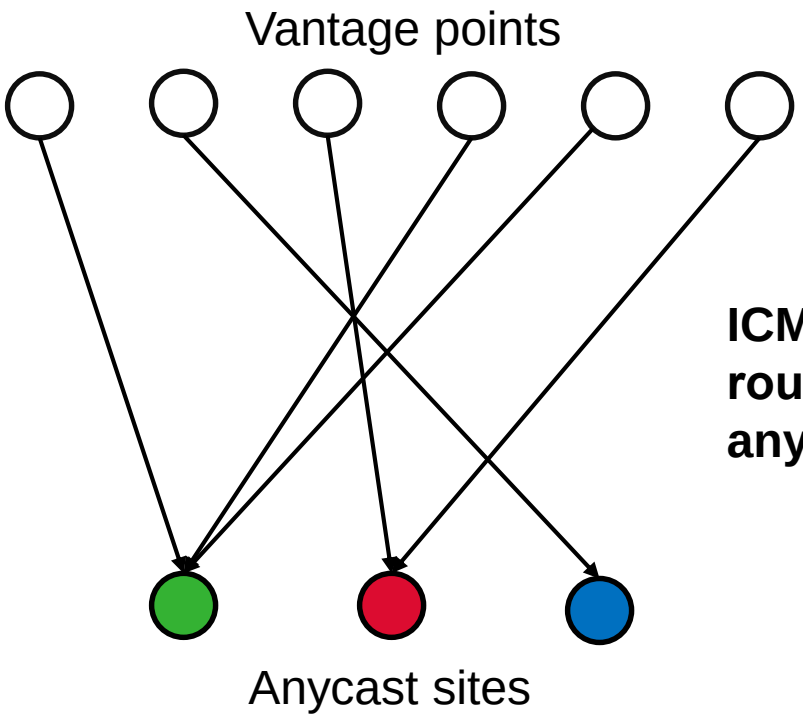

ICMP Echo Replies are routed towards “nearest” anycast site

VERFPLOETER: METHODOLOGY

ACTIVE MEASUREMENT WITH “PASSIVE” VANTAGE POINTS

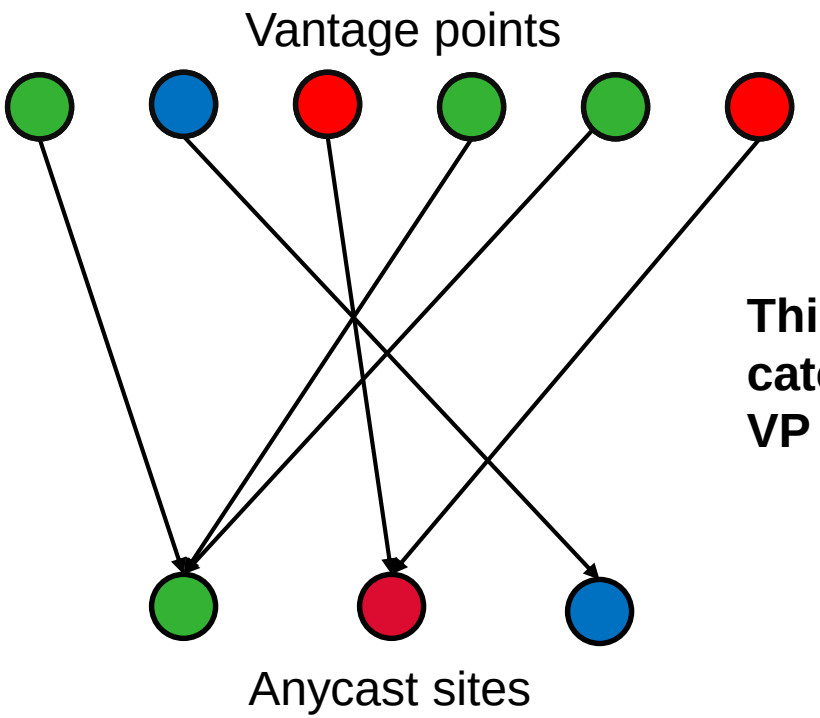

This defines the catchment that each VP belongs to

WHAT ARE THE PASSIVE VPS

ICMP-responding IPv4 or IPv6 addresses

We use a recent IPv4 hitlist from USC/ISI:
https://ant.isi.edu/datasets/ip_hitlists/

1 IPv4 address per /24 *block:* 14.7 million addresses

6.3 million after removing those in which there was never a responding host

```
1.0.0.0/24 - 1.0.0.57
1.0.1.0/24 - 1.0.1.85
1.0.2.0/24 - 1.0.2.4
1.0.3.0/24 - 1.0.3.204
```

USING VERFPLOETER: TWO CASE STUDIES

- The B Root DNS
  - Anycast since 1st of May 2017
  - 2 anycast sites
- Our anycast testbed
  - 9 anycast sites

Measured catchment using both Atlas and Verfploeter

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

Geolocated using MaxMind GeoLite2

	Atlas	Verfploeter
/24 blocks seen	8,677 (of 9083)	3,786,907 (of 6,877,175)
Geolocatable	8,677	3,786,229
Unique	2,079	3,606,300

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

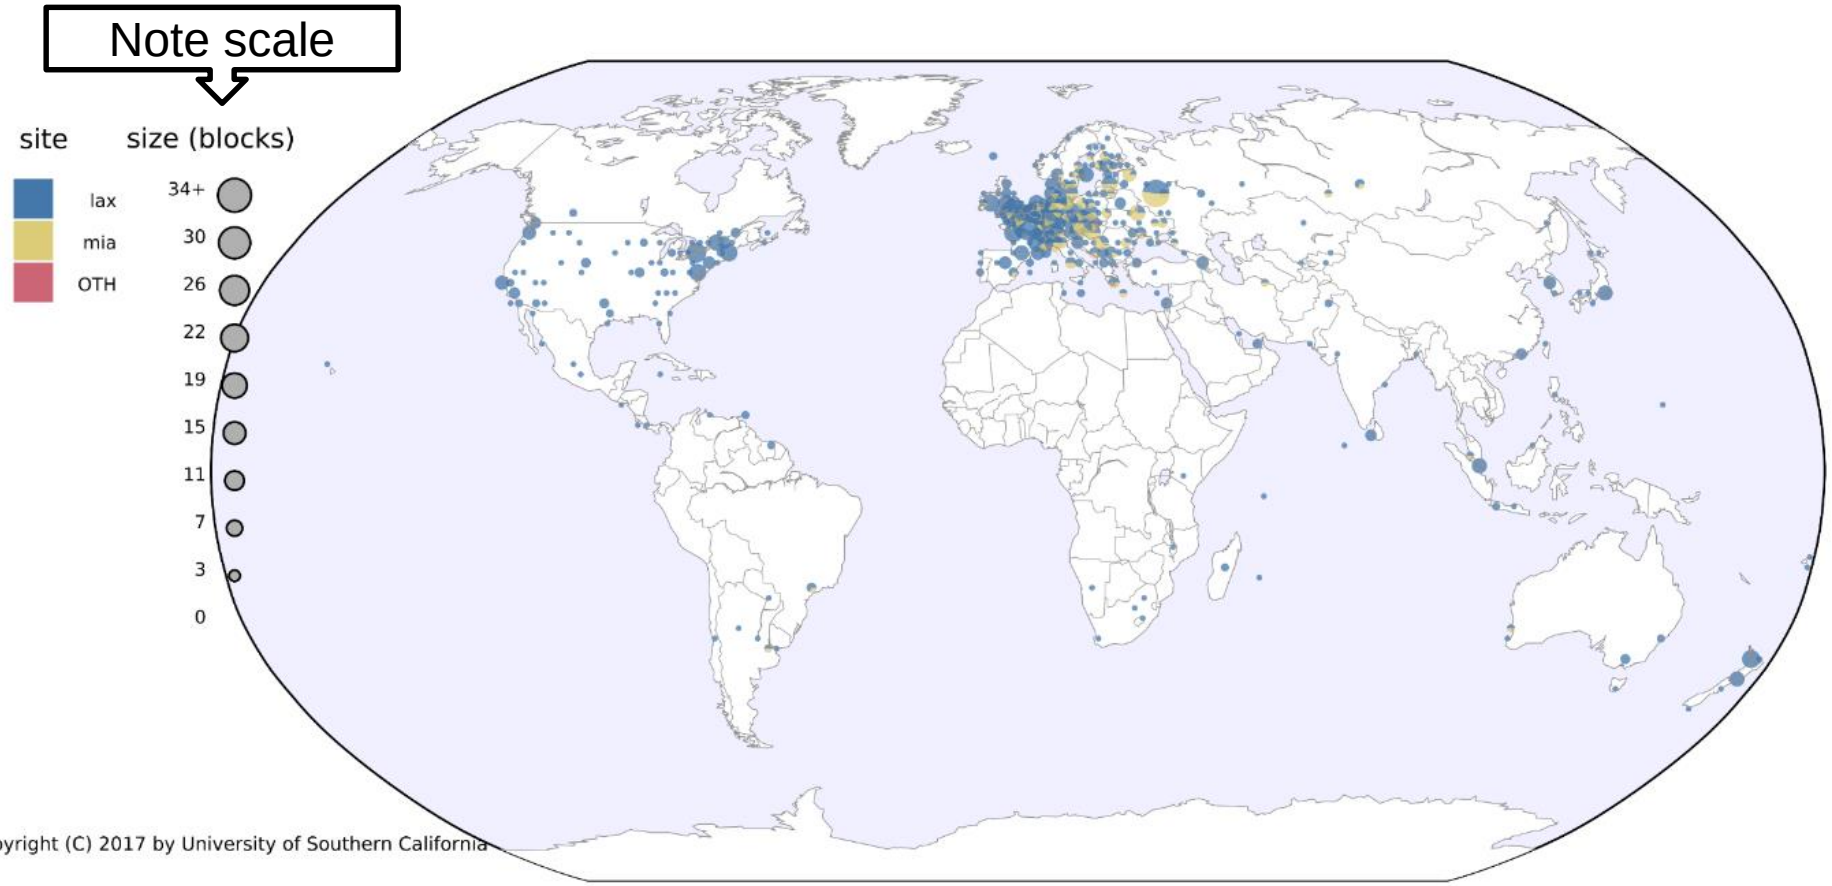

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

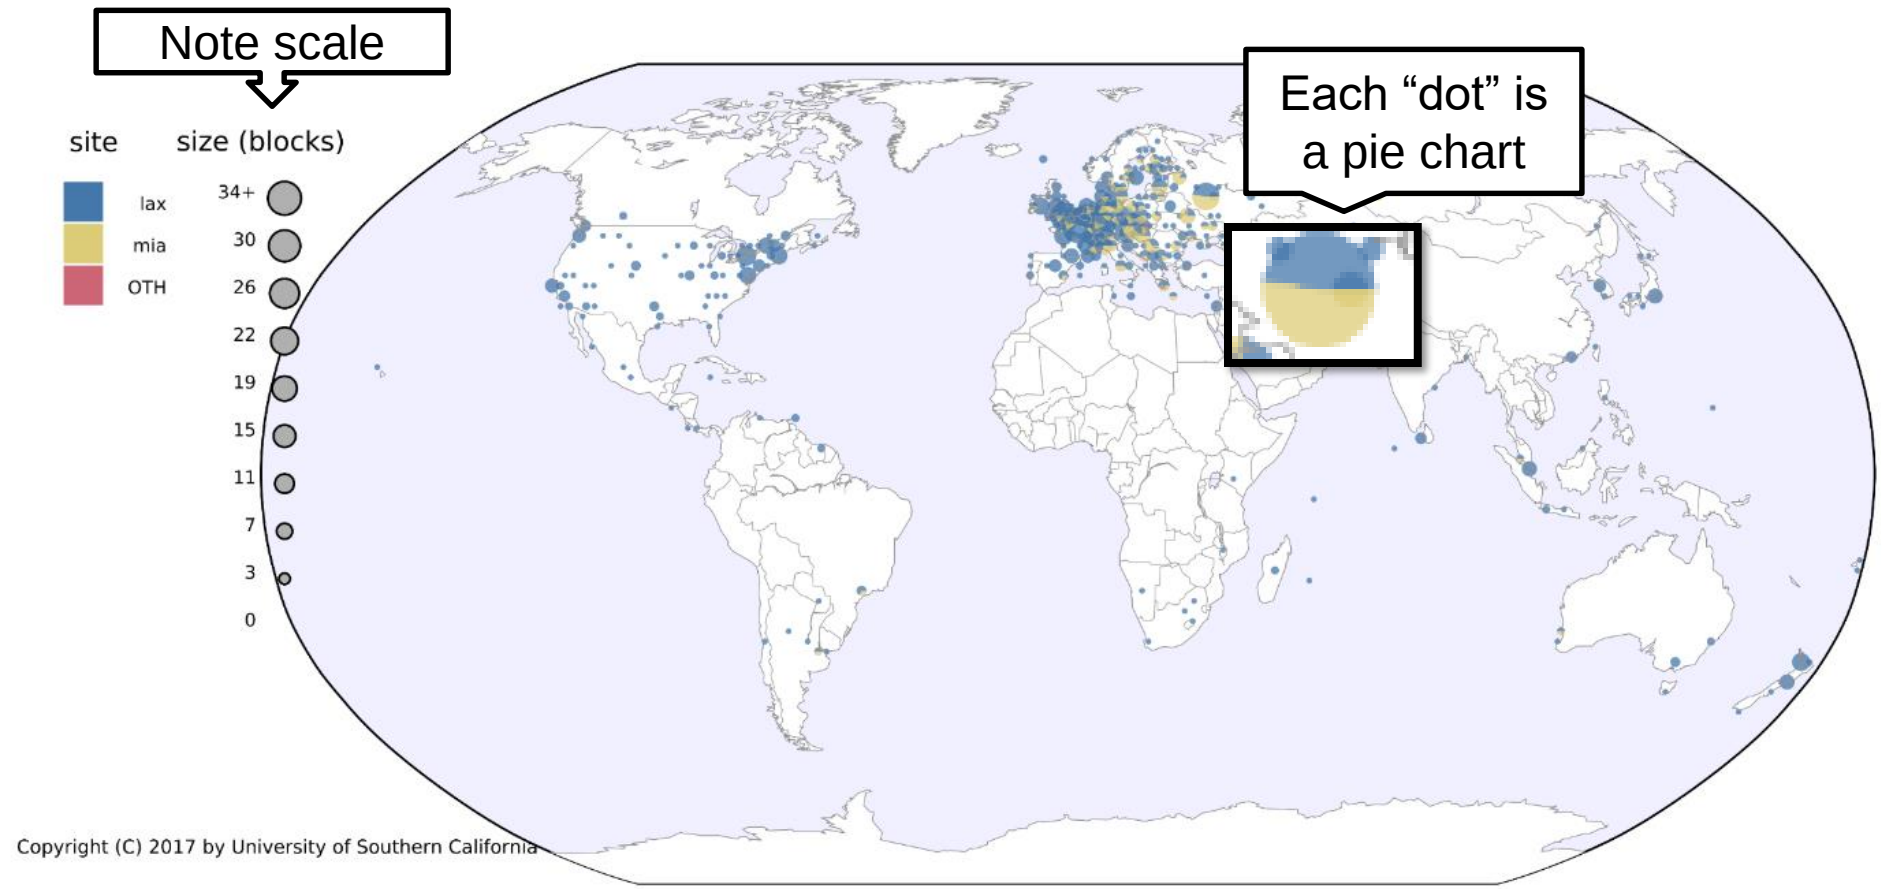

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

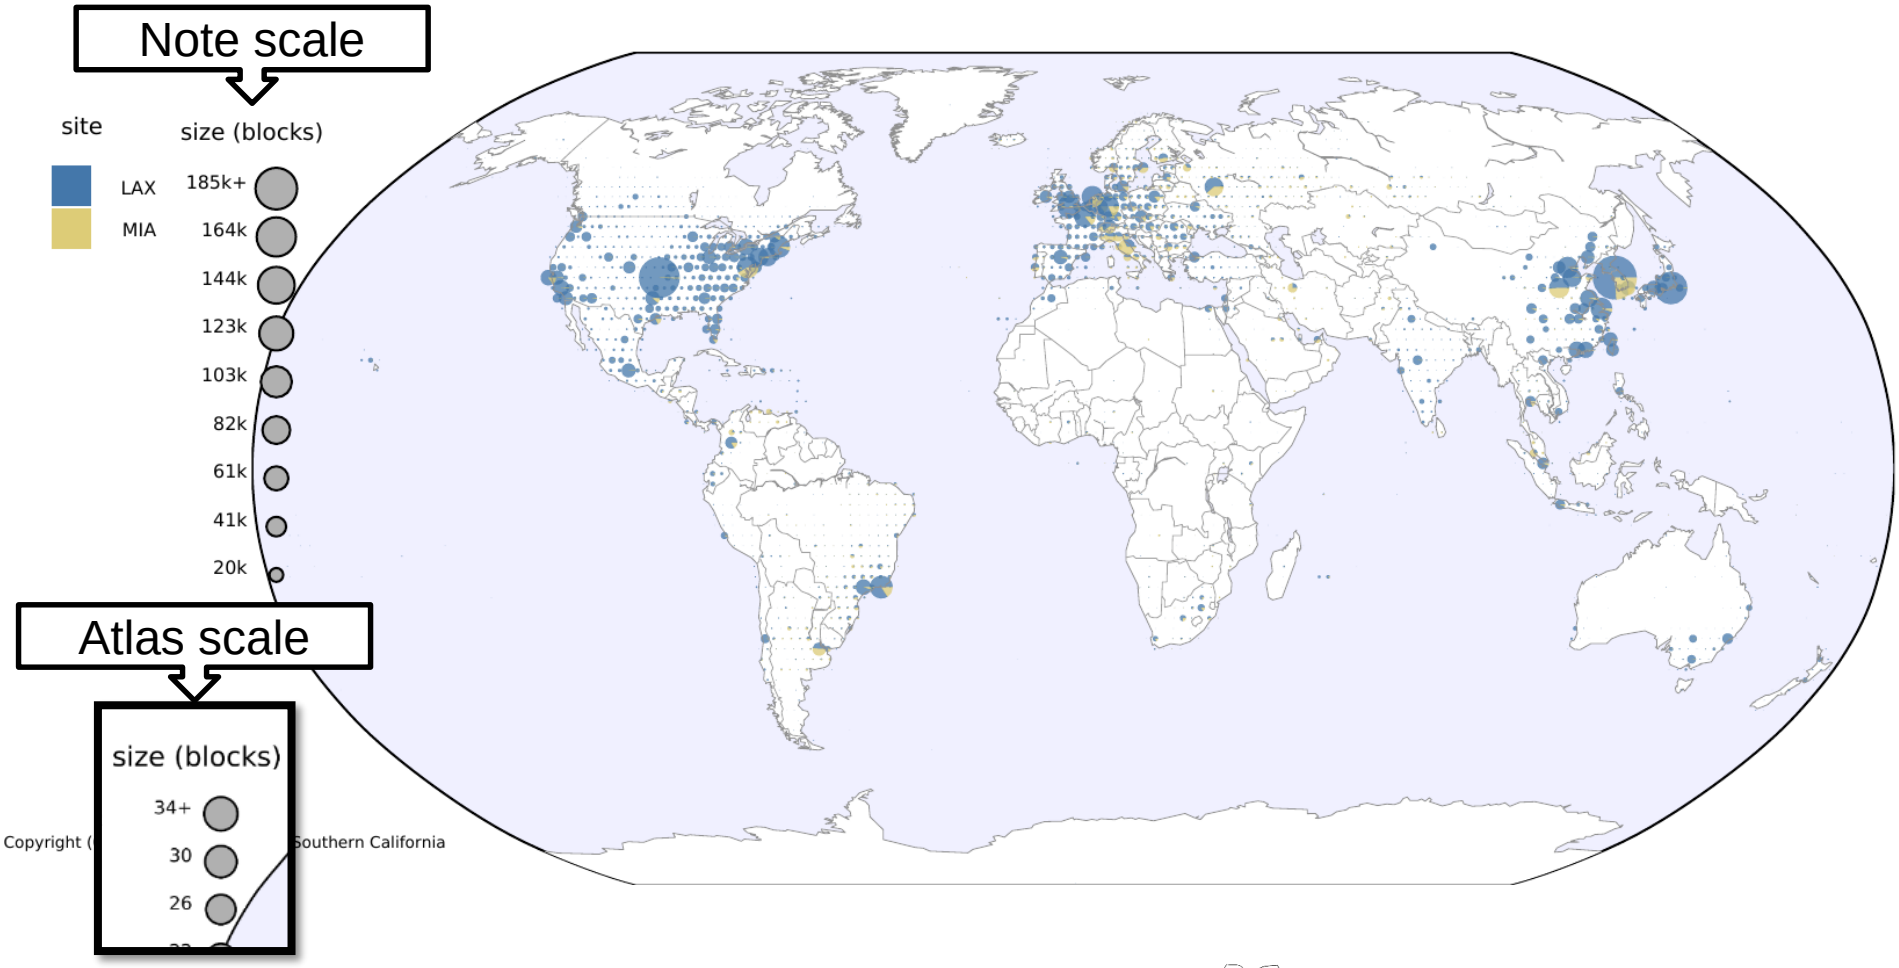

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

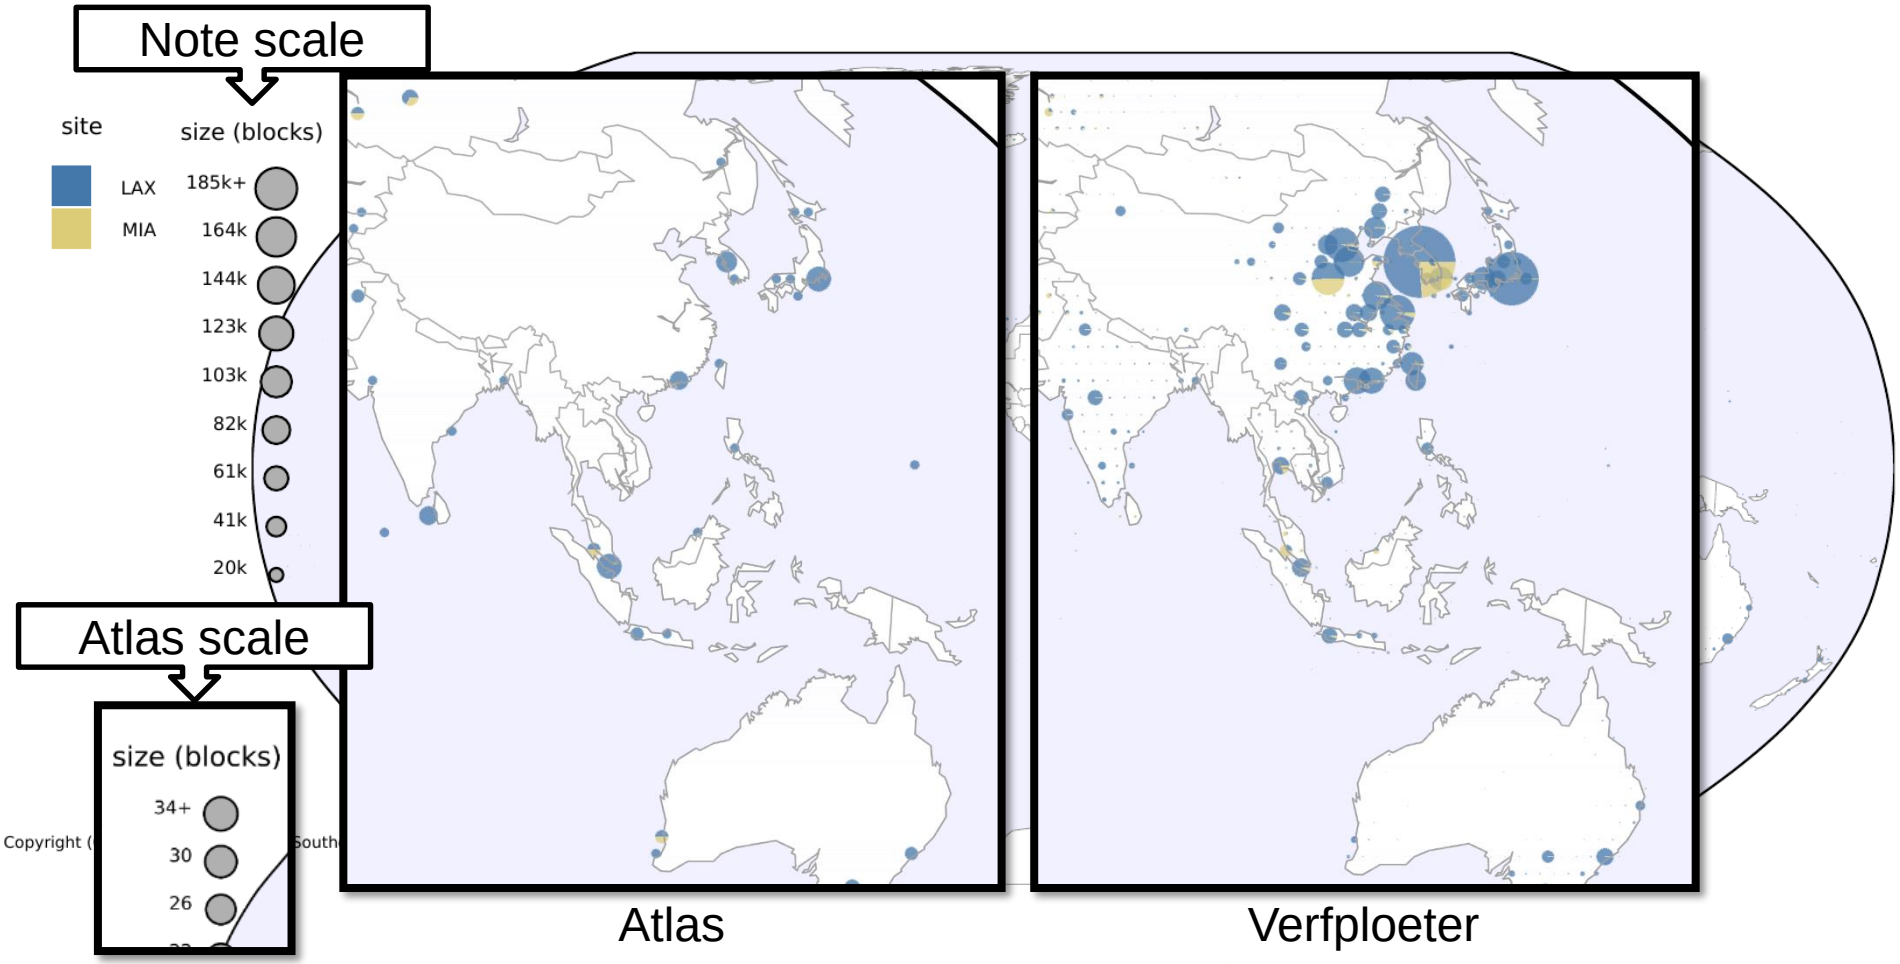

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

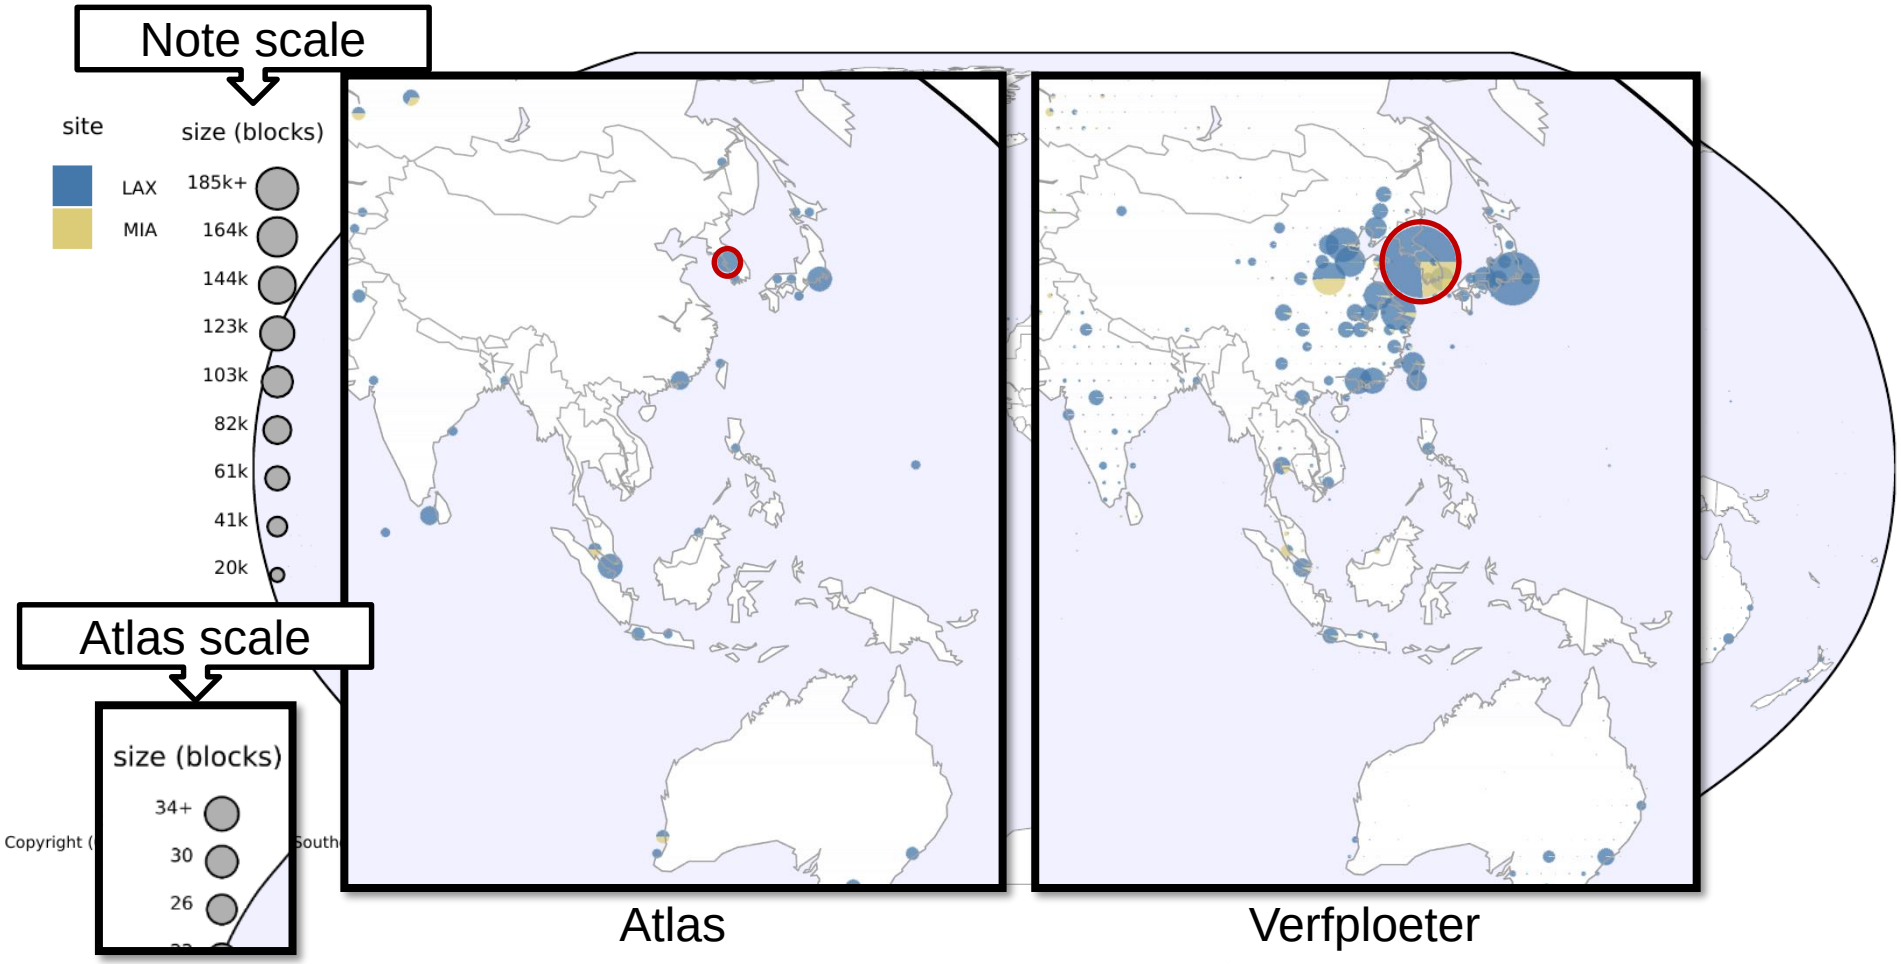

FROM ATLAS TO VERFPLOETER

B-ROOT CASE STUDY

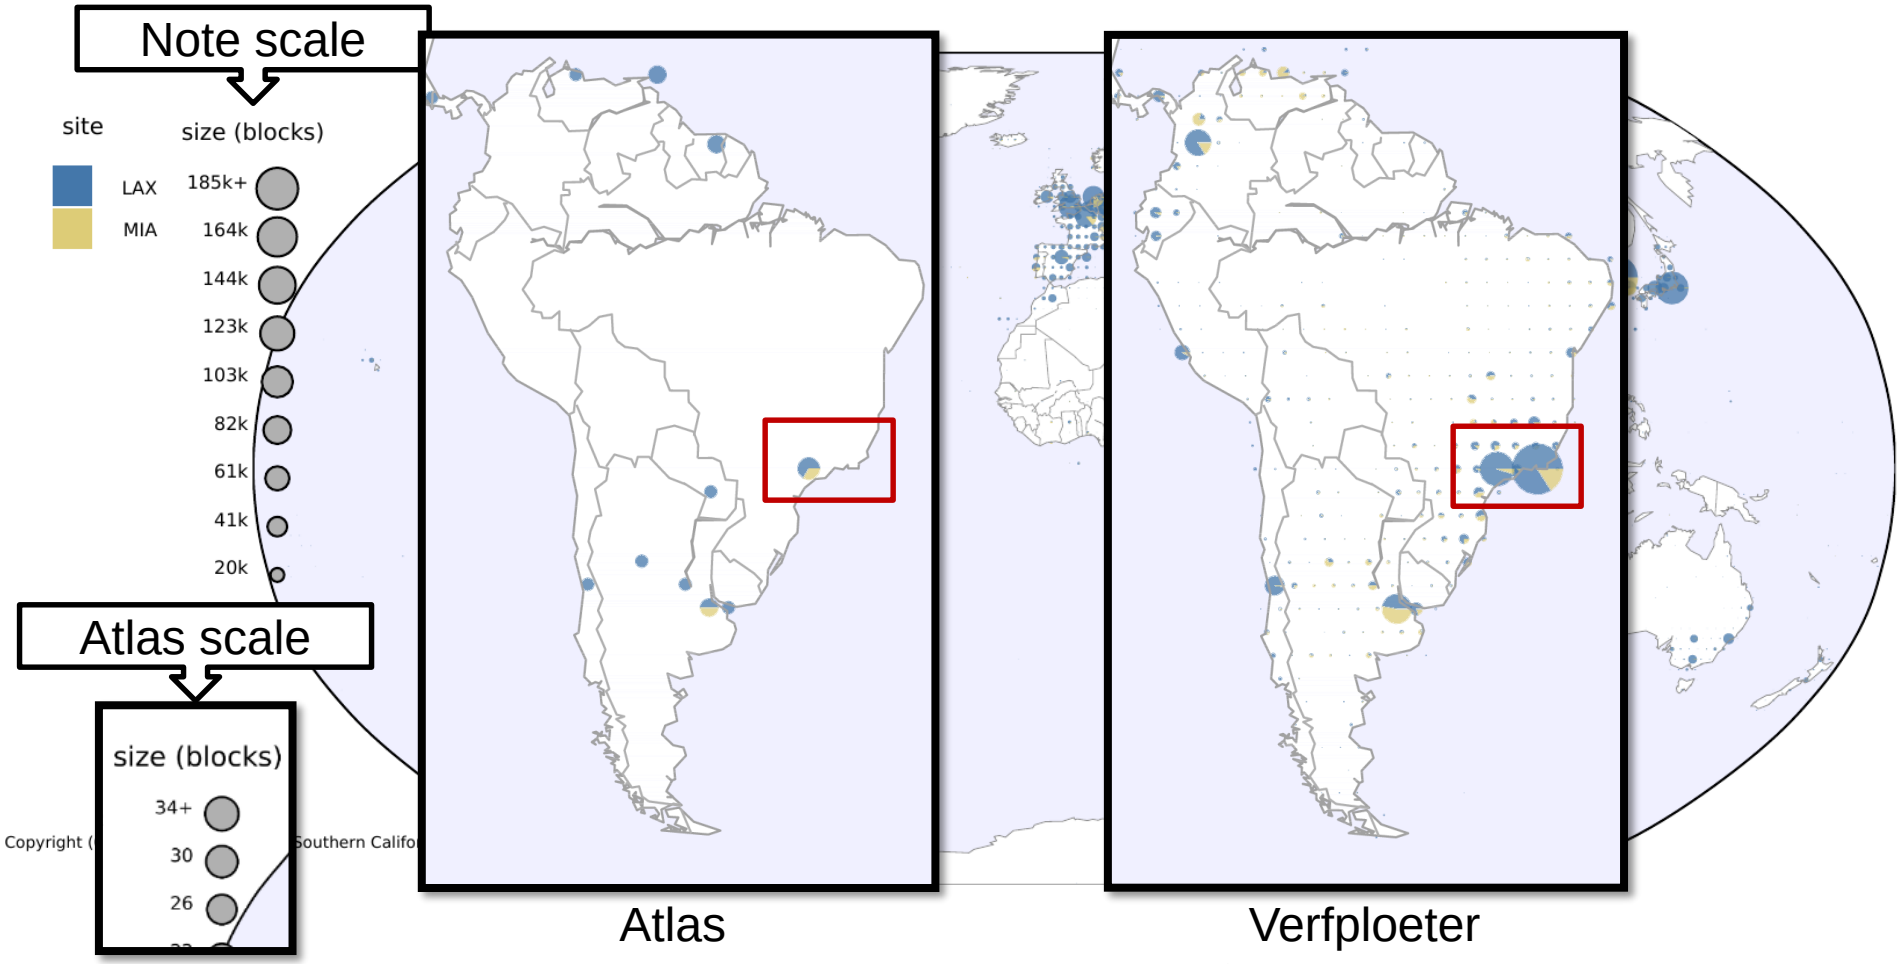

MEASURING AN ANYCAST TESTBED: TANGLED

A more diverse ground for verfploeter measurements:

Our 9-site anycast testbed

Location		Host	Upstream
AU	Sydney	Vultr	AS20473
FR	Paris	Vultr	AS20473
JP	Tokyo	WIDE	AS2500
NL	Enschede	Univ. of Twente	AS1103
UK	London	Vultr	AS20473
US	Washington	USC/ISI	AS1972
US	Miami	Florida Int. Univ.	AS20080
BR	Sao Paulo	Florida Int. Univ.	AS1251
DK	Copenhagen	DK-Hostmaster	AS39389

MEASURING AN ANYCAST TESTBED: TANGLED

A more diverse ground for verfploeter measurements:

Our 9-site anycast testbed

Location		Host	Upstream
AU	Sydney	Vultr	AS20473
FR	Paris	Vultr	AS20473
JP	Tokyo	WIDE	
NL	Enschede	Univ. of Twente	AS1103
UK	London	Vultr	AS20473
US	Washington	USC/ISI	AS1972
US	Miami	Florida Int. Univ.	AS20080
BR	Sao Paulo	Florida Int. Univ.	AS1251
DK	Copenhagen	DK-Hostmaster	AS39389

Shared upstream

MEASURING AN ANYCAST TESTBED: TANGLED

ATLAS VS VERFPLOETER

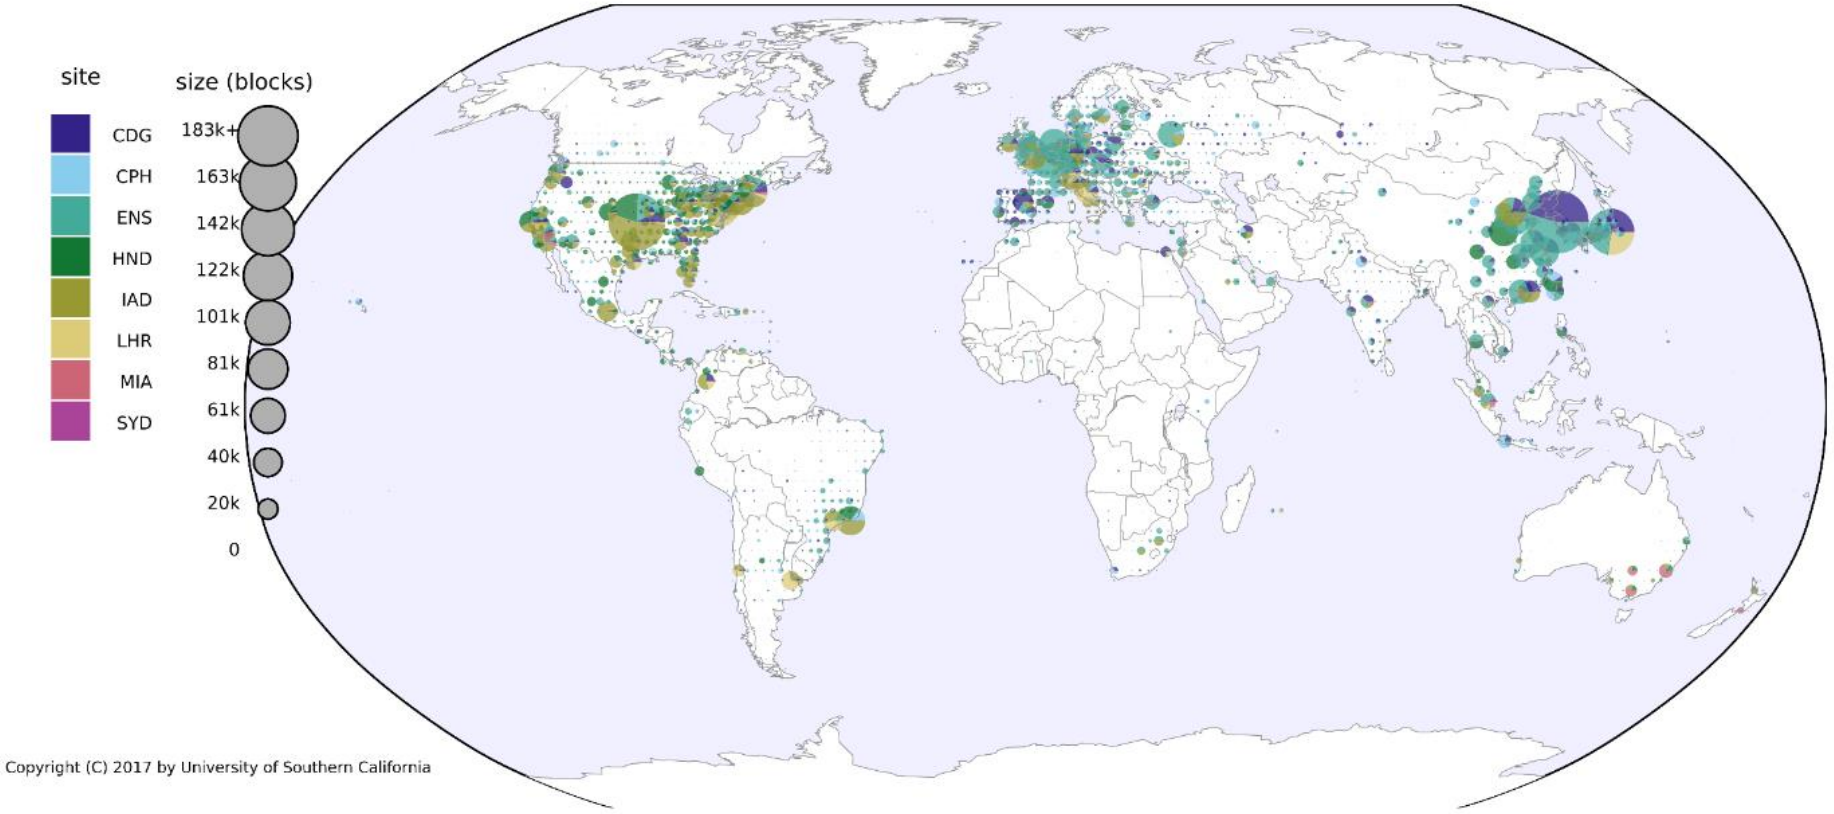

MEASURING AN ANYCAST TESTBED: TANGLED

ATLAS VS VERFPLOETER

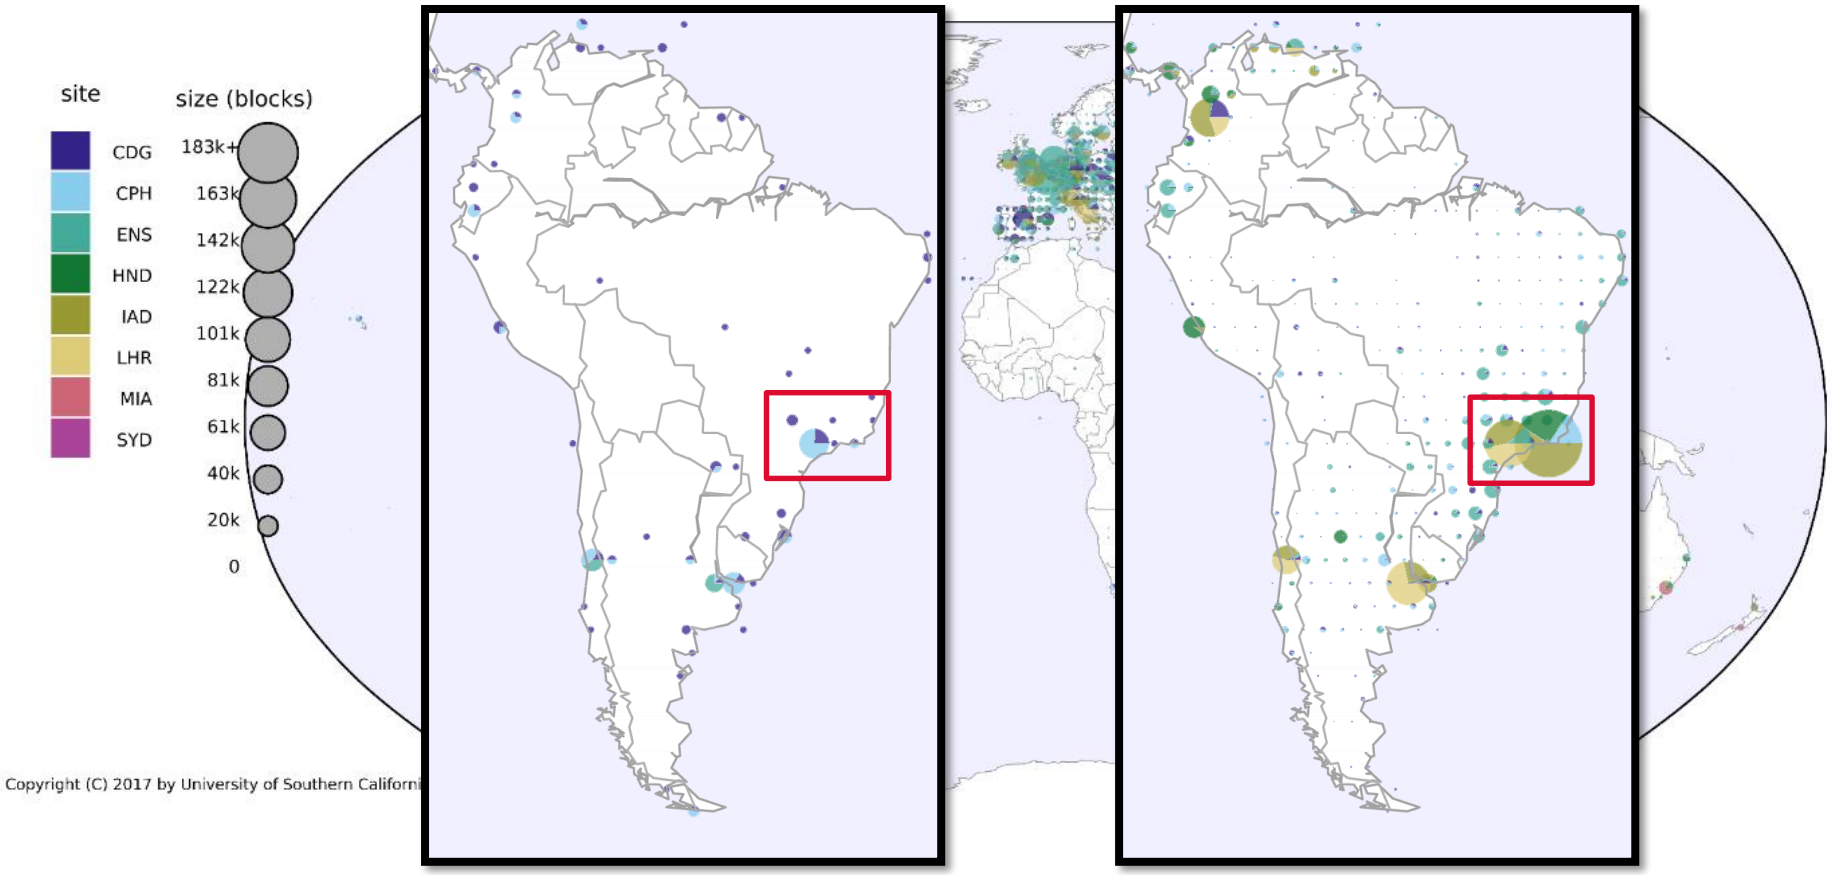

Atlas

Verfploeter

UNIVERSITY OF TWENTE. DACS Design and Analysis of Communication Systems

USC Viterbi School of Engineering ant.isi.edu

MEASURING AN ANYCAST TESTBED: TANGLED

ATLAS VS VERFPLOETER

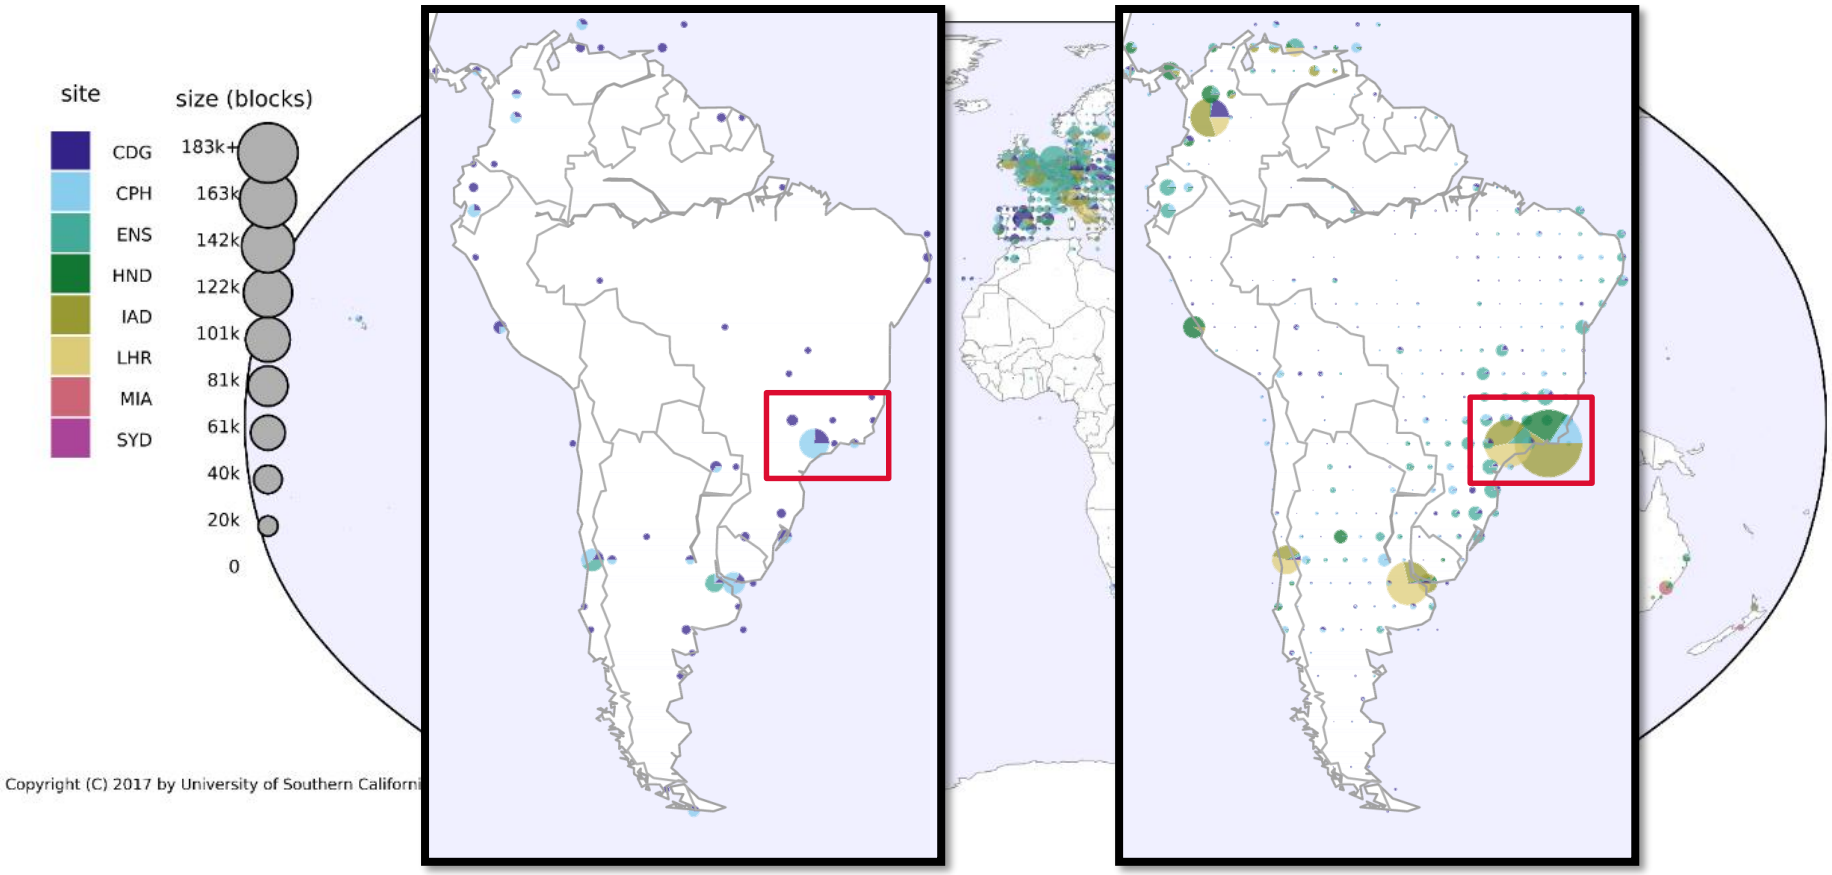

Atlas

Verfploeter

UNIVERSITY OF TWENTE. DACS Design and Analysis of Communication Systems

USC Viterbi School of Engineering ant.isi.edu

CALIBRATING CATCHMENT

Not all measured prefixes are equal
Some will send no traffic, others a lot

e.g. due to NAT or ISP DNS Forwarders

CALIBRATING CATCHMENT

Obtained with Verfploeter

Prefix	Site
1.0.0.0/24	LAX
1.0.1.0/24	LAX
1.0.2.0/24	LAX
1.0.3.0/24	MIA

RIGHT JOIN

Obtained from (historic) log files or from a service with a similar user base

Prefix	Load
1.0.0.0/24	10 Q/s
1.0.1.0/24	0 Q/s
1.0.2.0/24	0 Q/s
1.0.3.0/24	15 Q/s

CALIBRATING CATCHMENT

Prefix	Site
1.0.0.0/24	LAX
1.0.1.0/24	LAX
1.0.2.0/24	LAX
1.0.3.0/24	MIA

RIGHT JOIN

Prefix	Load
1.0.0.0/24	10 Q/s
1.0.1.0/24	0 Q/s
1.0.2.0/24	0 Q/s
1.0.3.0/24	15 Q/s

Prefix	Site	Load
1.0.0.0/24	LAX	10 Q/s
1.0.1.0/24	LAX	0 Q/s
1.0.2.0/24	LAX	0 Q/s
1.0.3.0/24	MIA	15 Q/s

CALIBRATING CATCHMENT

Prefix	Site
1.0.0.0/24	LAX
1.0.1.0/	
1.0.2.0/	
1.0.3.0/	

Prefix	Load
1.0.0.0/24	10 Q/s

SITE	Fraction of traffic
LAX	0.40
MIA	0.60

Prefix	Site	Load
1.0.0.0/24	LAX	10 Q/s
1.0.1.0/24	LAX	0 Q/s
1.0.2.0/24	LAX	0 Q/s
1.0.3.0/24	MIA	15 Q/s

CALIBRATING CATCHMENT AT B-ROOT

1. Record traffic from unicast B-root (prior to anycast deployment)

Date	Site	Queries/day	Queries/sec
2017-04-12	LAX (Unicast)	2.34G	27.1k
2017-05-15	Both	2.20G	25.4k
	LAX	1.78G	20.6k
	MIA	0.407G	4.71k

2. Run Verfploeter on anycast B-root

CALIBRATING CATCHMENT AT B-ROOT

RECORDED TRAFFIC COMBINED WITH VERFPLOETER

Before calibration

Date	Method	Measurement	% to LAX
May 15th	RIPE Atlas	9682 VPs	82.4%
May 15th	Verfploeter	3.923M /24s	87.8%
May 15th	Verfploeter x Load		81.6%
May 15th	Actual load	2.188G q/day	81.4%

After calibration

CALIBRATING CATCHMENT AT B-ROOT

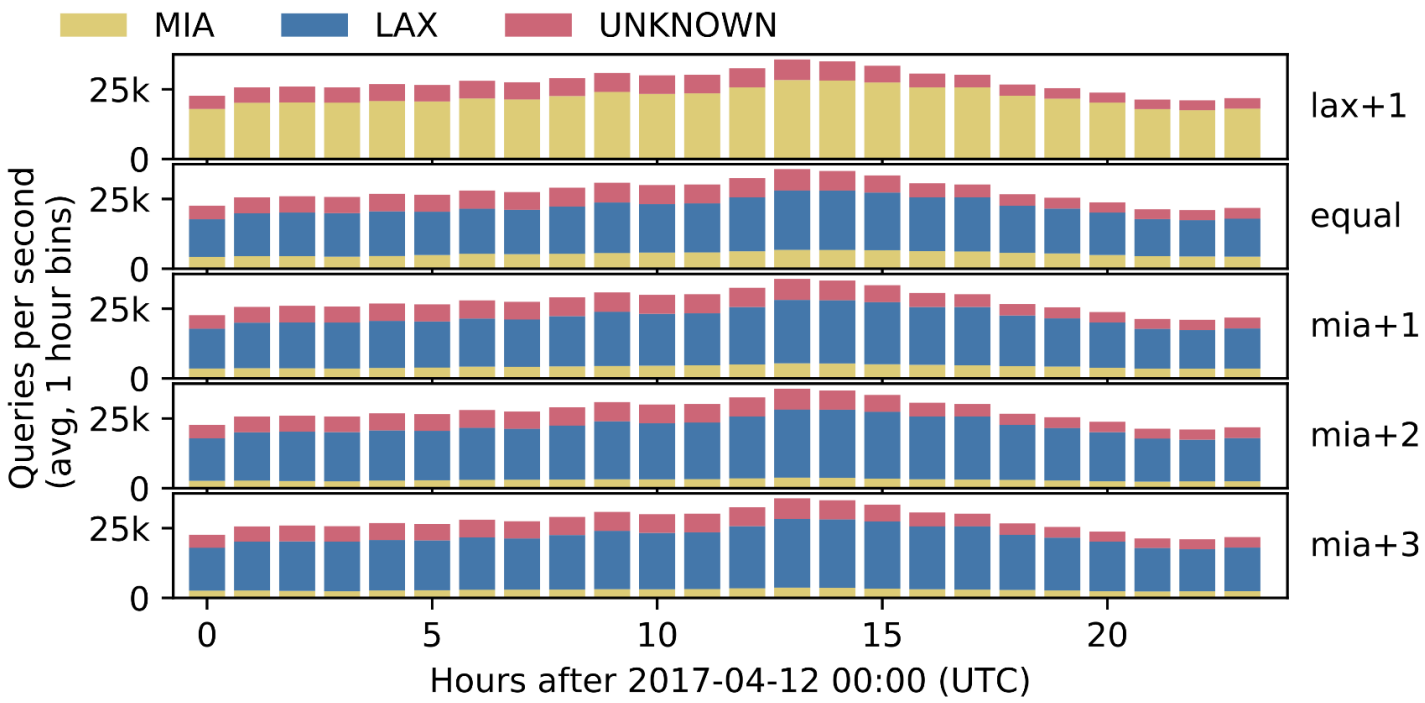

Verfploeter allows you to experiment with various configurations and how that affects the load on the sites of your service

INSIDE AUTONOMOUS SYSTEMS

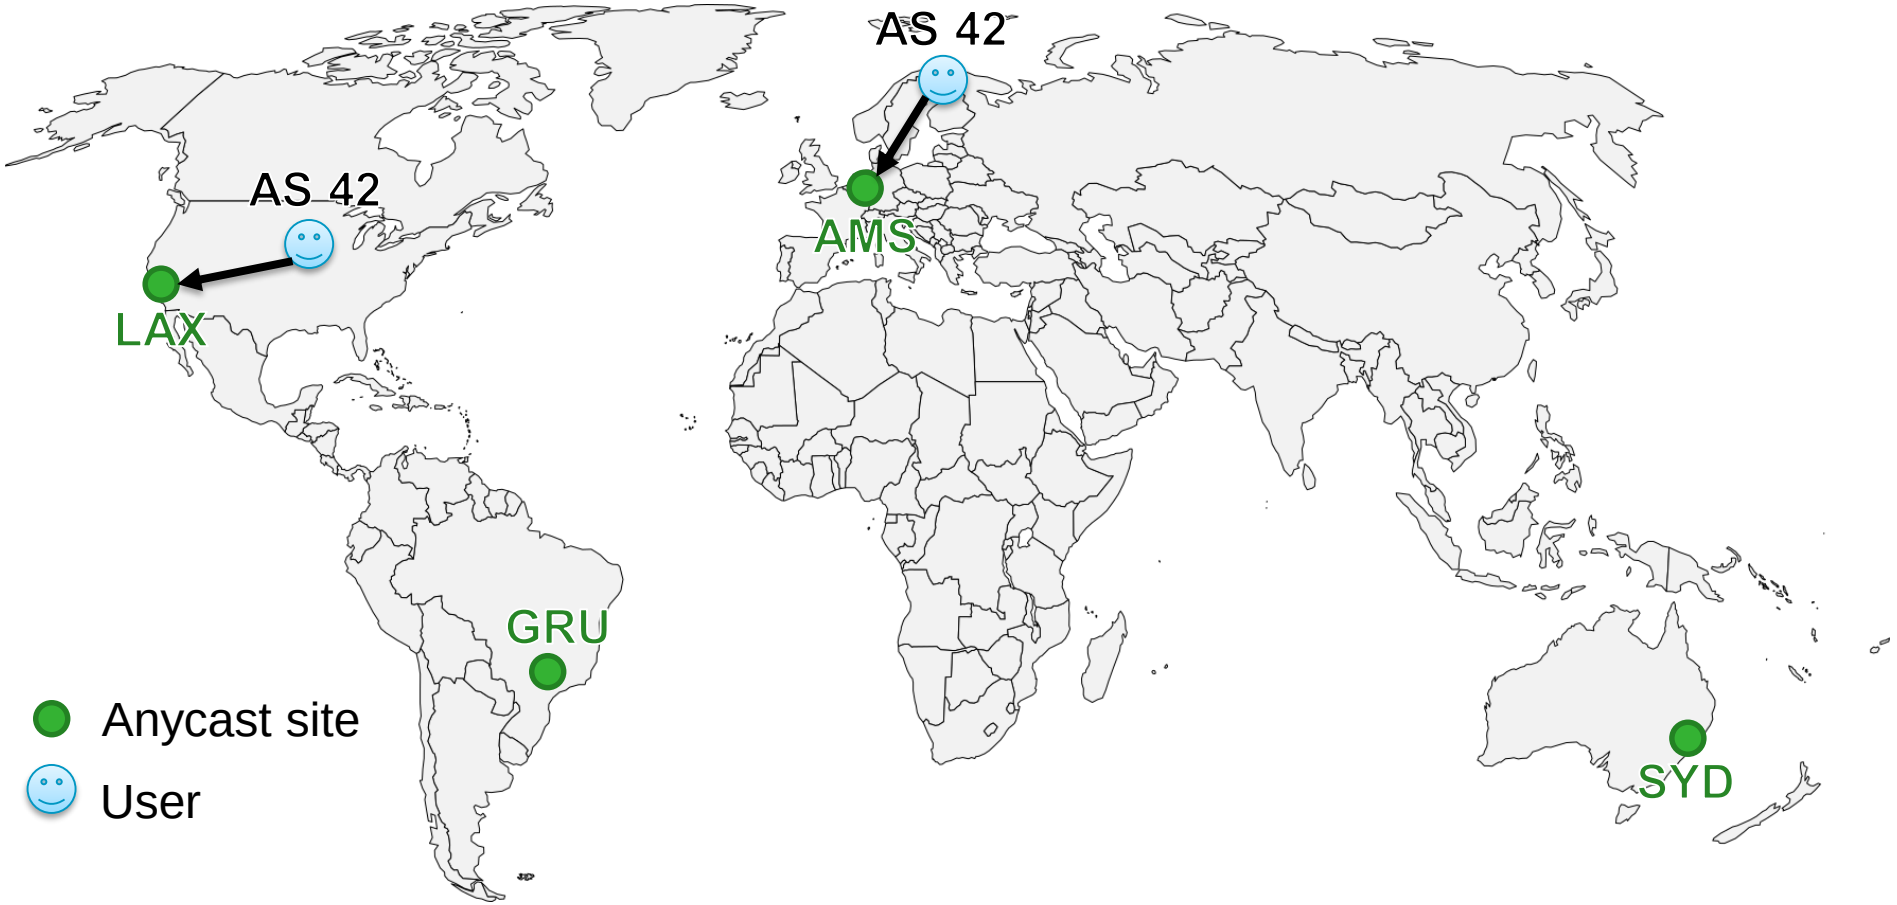

INSIDE AUTONOMOUS SYSTEMS

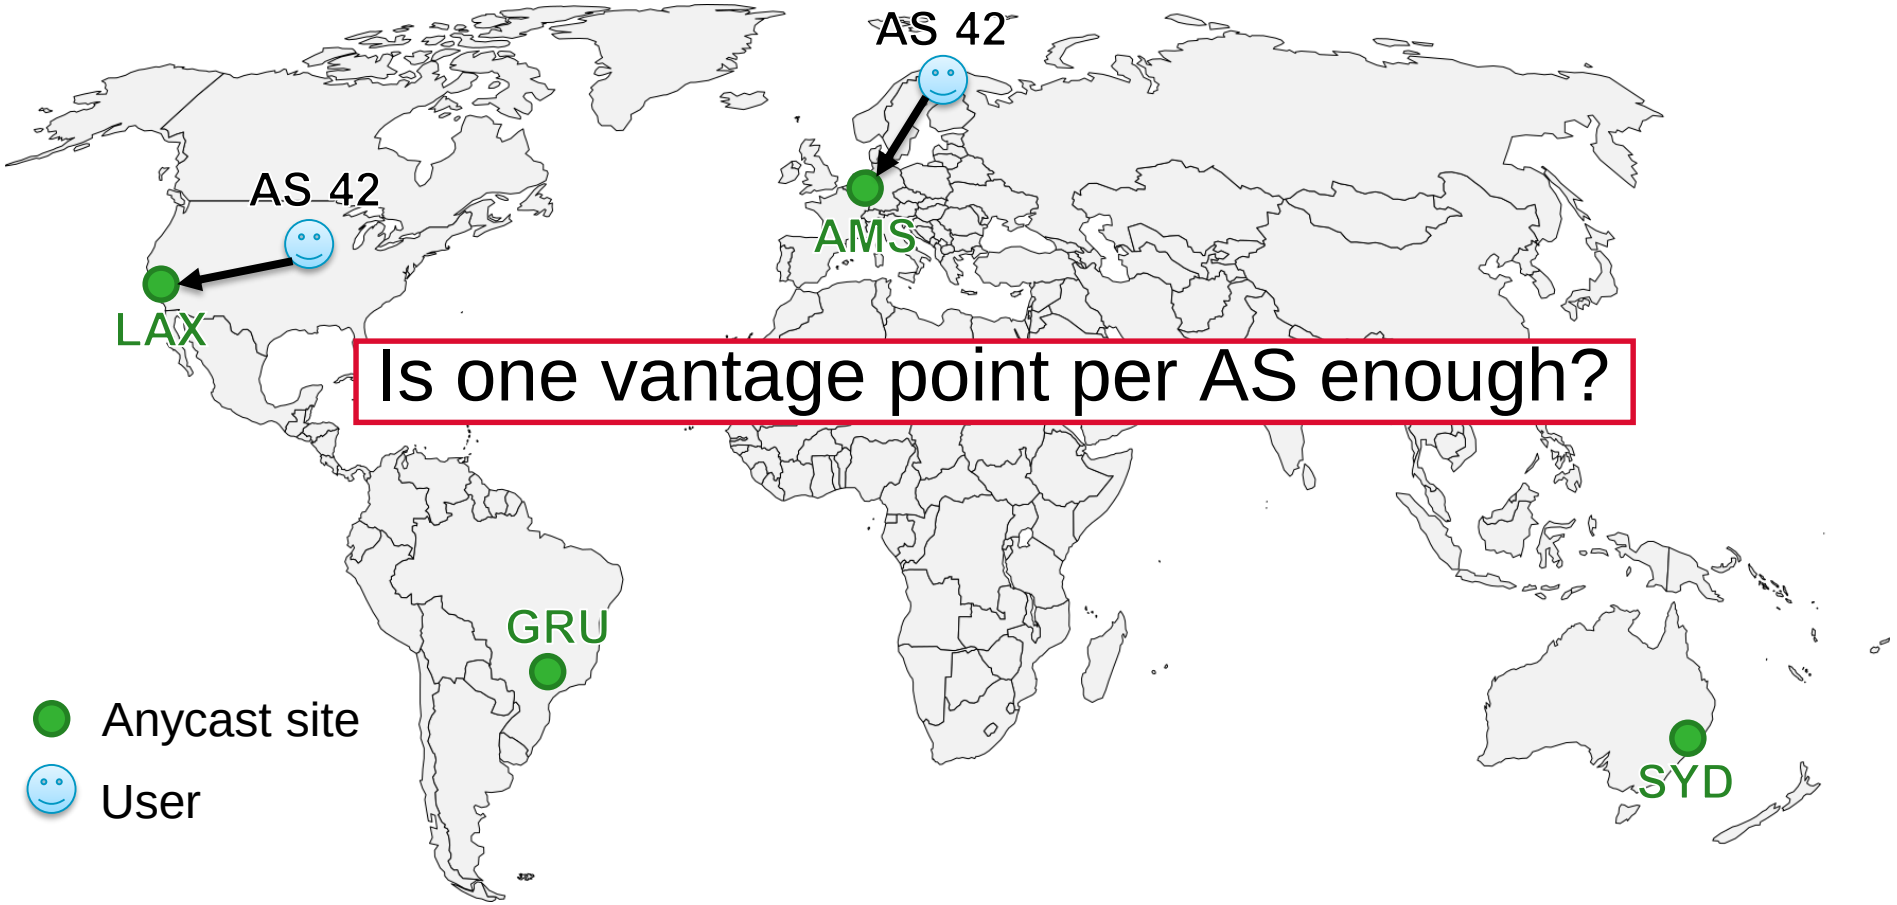

INSIDE AUTONOMOUS SYSTEMS

NUMBER OF SITES VERSUS NUMBER OF ANNOUNCED PREFIXES

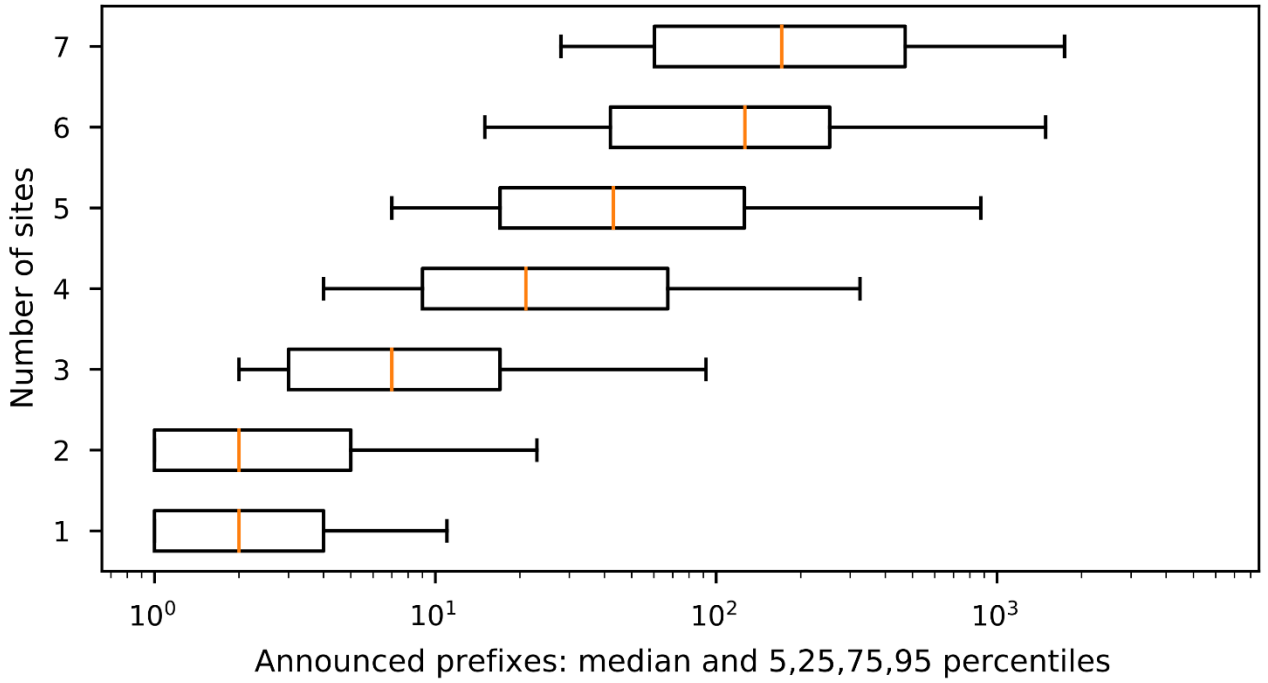

We see more than 1 site from around 12.7% of all ASes

CONCLUSIONS

- Verfploeter has ~430x more “vantage points” than Atlas
  - Atlas has some blocks that we don’t have
- Calibration of measured catchment to actual user base is required

Improve your anycasted service using Verfploeter

https://github.com/woutifier/

Do your own analysis with our public datasets

https://ant.isi.edu/datasets/anycast/index.html#verfploeter

UNIVERSITY OF TWENTE.

USC Viterbi School of Engineering

DACS Design and Analysis of Communication Systems

ant.isi.edu

QUESTIONS AND COMMENTS