

Verfploeter: Broad and Load-Aware Anycast Mapping

Wouter B. de Vries, Ricardo de O. Schmidt, Wes Hardaker,
John Heidemann, Pieter-Tjerk de Boer, Aiko Pras foo

2017-08-30




Our Starting Goals:

- ▶ Develop a technique to:
 - ▶ Accurately map anycast catchments
 - ▶ Accurately study B's anycast IPv4 catchments
 - ▶ Predict load in advance of changes


Traditional "Active" Anycast Probing

- ▶ Traditional techniques require lots of Vantage Points
- ▶ e.g. RIPE Atlas



"Verfploeter" uses "Passive" Vantage Points

- ▶ Verfploeter uses the Internet as its vantage points
- ▶ We collect response traffic to ICMP "pings"



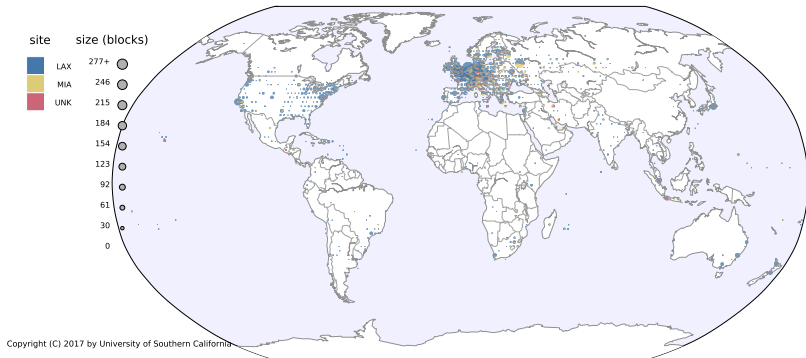
Input: An IPv4 Hitlist

- ▶ Previous work supplied us with ICMP-responsive IP-addresses in each IPv4 /24 block:

*Xun Fan and John Heidemann,
"Selecting Representative IP Addresses for Internet
Topology Studies"*

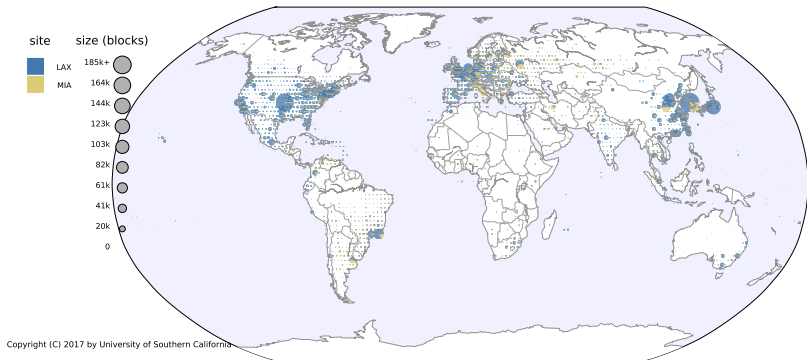
- ▶ https://ant.isi.edu/datasets/ip_hitlists/
- ▶ Example addresses:
 - 1.0.100.37
 - 1.0.101.100
 - 1.0.102.123
 - 1.0.103.1
 - 1.0.104.1
 - 1.0.105.106
 - 1.0.106.123

RIPE Atlas Coverage of B-Root



(measured 2017/05/15)

Verfploeter Coverage of B-Root

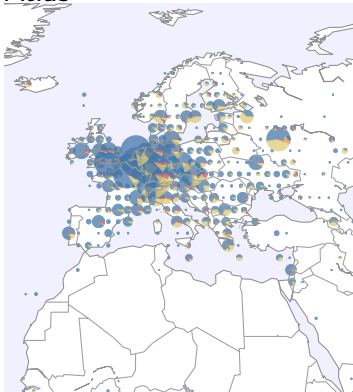


Note: huge scale difference


B-Root Coverage Differences: Europe

- Atlas has good coverage of Europe

Atlas




Verfploeter




B-Root Coverage Differences: North America

- ▶ Atlas has reasonable coverage in North America
- ▶ Verfploeter has significantly better coverage

Atlas



Verfploeter



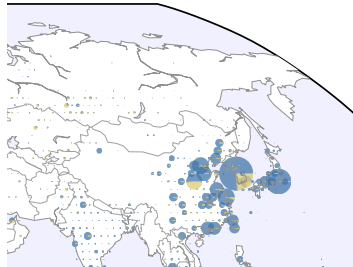
B-Root Coverage Differences: China

- ▶ Atlas has sparse coverage elsewhere (e.g. China)
- ▶ Verfloploeter fills in the holes

Atlas



Verfploeter



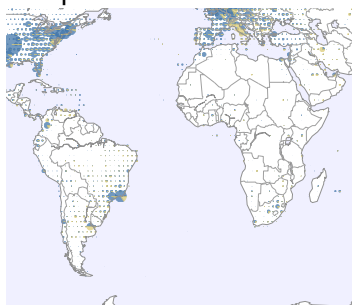
B-Root Coverage Differences: Southern Continents

- Verfploeter has more South America and Africa coverage

Atlas



Verfploeter



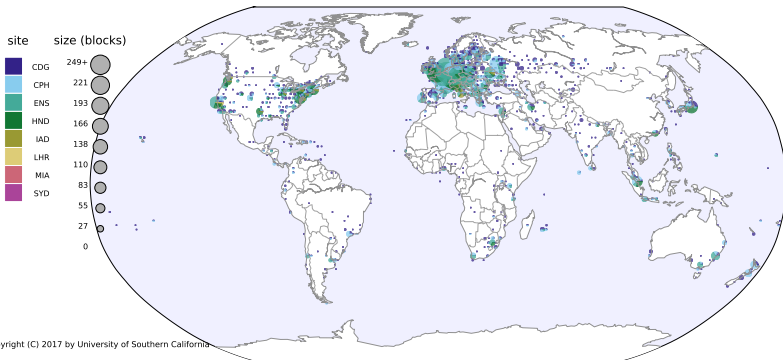
Tangled: Studying A Larger Anycast Testbed

- ▶ Tangled: A 9-site anycast testbed University of Twente
- ▶ Tangled provides a more complex anycast test framework

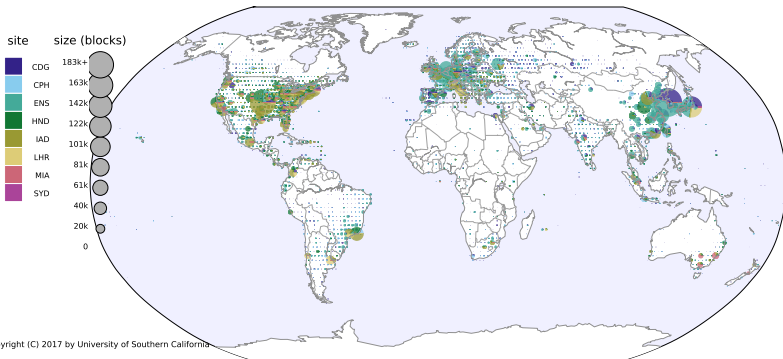
Location		Host	Upstream
AU	Sydney	Vultr	AS20473
FR	Paris	Vultr	AS20473
JP	Tokyo	WIDE	AS2500
NL	Enschede	Univ. of Twente	AS1103
UK	London	Vultr	AS20473
US	Miami	Florida Int. Univ.	AS20080
US	Washington	USC/ISI	AS1972
BR	Sao Paulo	Florida Int. Univ.	AS1251
DK	Copenhagen	DK-Hostmaster	AS39389

- ▶ Note: some locations have common upstreams

RIPE Atlas Coverage of Tangled



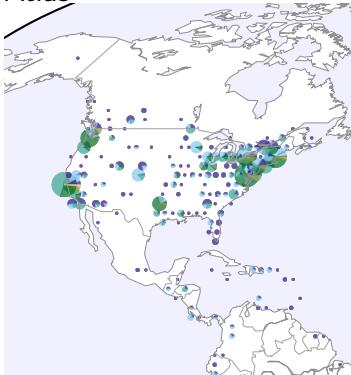
Verfploeter Coverage of Tangled



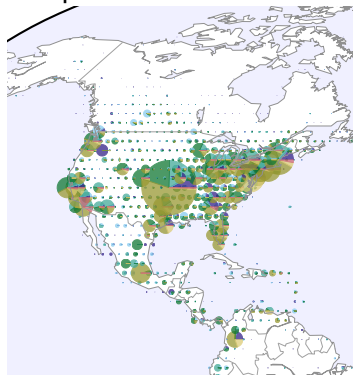
Tangled Coverage Differences: North America

- ▶ IAD (olive colored):
 - ▶ Has little coverage according to Atlas
 - ▶ Significant NA coverage according to Verfploeter

Atlas




Verfploeter




Tangled Coverage Differences: Australia

- ▶ Australia also shows significantly different coverage

Atlas



Verfploeter



Estimating Load Changes

Question:

- ▶ Can we accurately estimate the future load of changes to an Anycast network?
- ▶ Even before adding a new anycast instance?
- ▶ To test this theory, we conducted a study at B-Root

Estimating Load Changes at B-Root

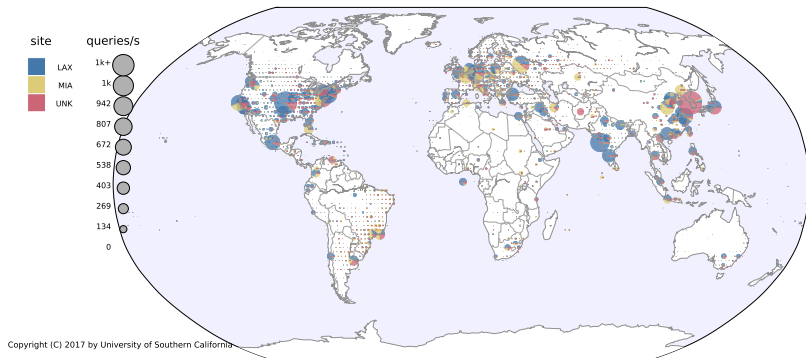
1. We recorded traffic measurements for our single LAX site
 - ▶ (before deploying our second site at MIA)
 - ▶ Some /24s send us more traffic than others
 - ▶ (e.g. India, which has heavy NAT deployment)
2. We ran Verfploeter against a deployed test prefix
 - ▶ (the test prefix was served by MIA ahead of time)
3. We combined these to calculate an estimated load

Measured Load at B-Root

Date	Site	Queries / Day	Queries / Sec
2017-04-12	LAX	2.34G	27.1k
2017-05-15	both	2.20G	25.4k
	LAX	1.78G	20.6k
	MIA	0.407G	4.71k

- Averaged over a 24h UTC period

Verfploeter Load Prediction of B-Root



- ▶ Red “unknown” blocks show holes in Verfploeter coverage
 - ▶ 12.9% of /24s did not respond to our probes (e.g. KR)

Estimated vs Measured

Date	Method	Measurement		% LAX
May 15	RIPE Est.	9682	VPs	82.4%
May 15	Verfploeter /24s	3.923M	/24s	87.8%
May 15	Verfploeter Est.			81.6%
May 15	Actual Load	2.188G	q/day	81.4%

Notes:


- ▶ **Unknown** /24s (12.9%) were given an average weight
- ▶ Weighting shifted Verfploeter results 87.8% \Rightarrow 81.6%
- ▶ \therefore **Weighting is required to get accurate predictions**
- ▶ We couldn't use RIPE to estimate loads – too few /24s

Do We Need To Look Inside ASes?

Given our significant number of new vantage points:


- ▶ Can we study traffic catchments **within an AS**?
- ▶ Do the number of sites seen change:
 - ▶ with prefix size?
 - ▶ with AS size?
- ▶ Can we measure routing stability using Verfploeter?
 - ▶ (**Yes!** – see result in paper)

Sites Seen vs Number of Prefixes Announced



- ASes announcing more prefixes are likely to see more anycast sites

Sites Seen vs Prefix Size



- ▶ 80% of prefixes smaller than /16 reach a single site
- ▶ Larger prefixes are more likely to see more sites

Verfploeter Sizing Summary

- ▶ Verfploeter sees $\sim 430\times$ more network blocks than Atlas
 - ▶ Atlas has some unique blocks though
- ▶ Difference in /24s seen by Verfploeter and Atlas:

	Atlas	Verfploeter
/24 blocks seen	8677	3,786,907
geolocatable	8677	3,786,229
unique	2079	3,606,300

Verfploeter: Ready For Use

Verfploeter provides a **novel mechanism for studying anycast**

- ▶ Paper:

- ▶ <https://www.isi.edu/~johnh/PAPERS/Vries17a.pdf>

- ▶ Datasets:

- ▶ <https://ant.isi.edu/datasets/anycast/index.html#verfploeter>

- ▶ Software:

- ▶ <https://ant.isi.edu/software/verfploeter/index.html>

Please feel free to contact us for help or to share your results