MAnycast^2 Using Anycast to Measure Anycast

R. Sommese, L. Bertholdo,
G.Akiwate, M. Jonker,
R. van Rijswijk-Deij,A. Dainotti,
K. Claffy,A. Sperotto

IMC2020 – Virtual Conference

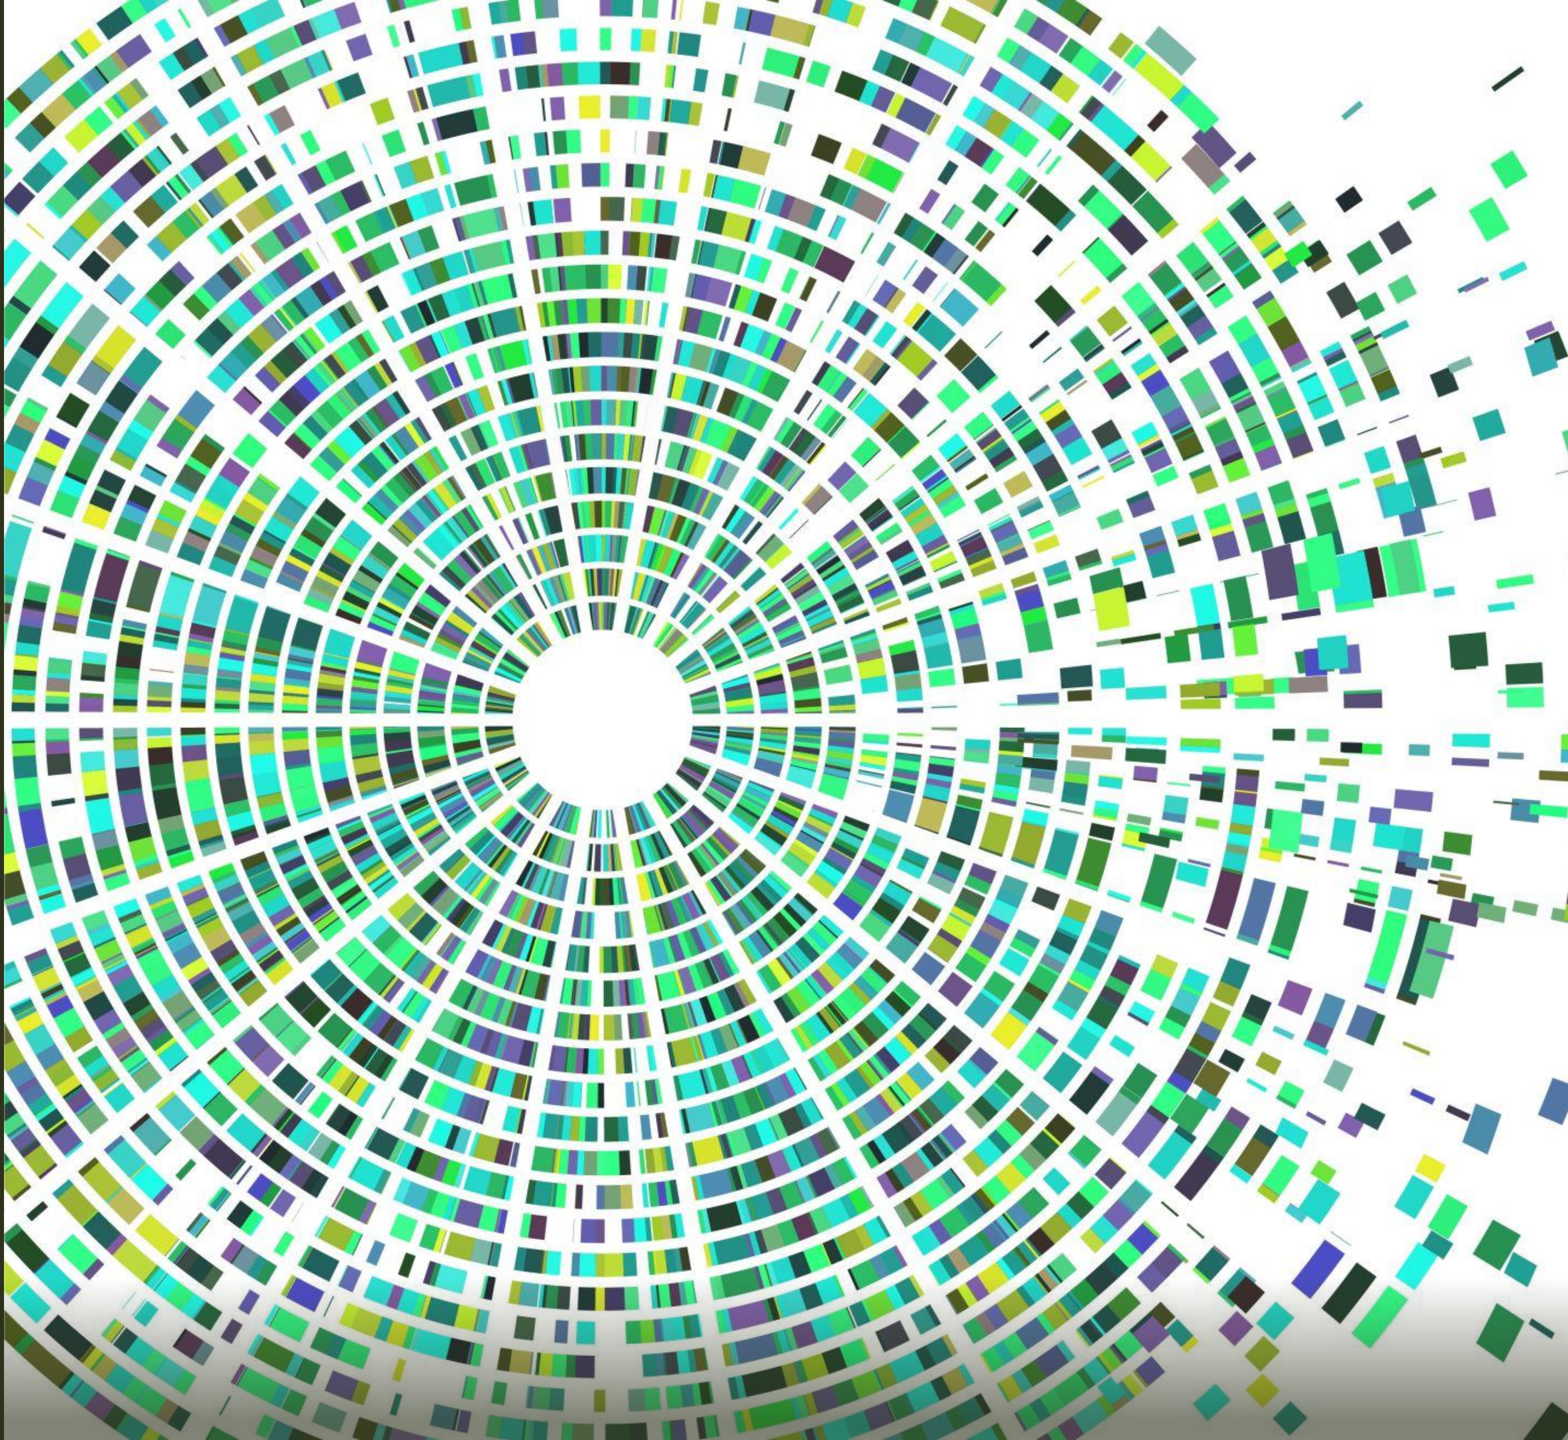

Homeland Security

UNIVERSITY OF TWENTE.

Netherlands Organisation for Scientific Research

MAnycast^2 Using Anycast to Measure Anycast

R. Sommese, L. Bertholdo, G. Akiwate,
M. Jonker, R. van Rijswijk-Deij,
A. Dainotti, K. Claffy, A. Sperotto

IMC2020 – Virtual Conference

Introduction

Anycast became a common way to improve resiliency of Internet services.

Identifying address prefixes that are anycast would enable more accurate assessment of resilience properties.

IPv6 introduced a special format for anycast addresses, whereas the IPv4 approach relies on assigning the same unicast IP to multiple hosts and leverage on routing.

The routing opacity creates a measurement challenge.

Approaches in detecting anycast

- First anycast DNS implementations use CHAOS-class record to provide information regarding the anycast server instances.
- IETF proposed two Best Current Practice: NSID and Unique ASN per instance to distinguish different anycast instances.
- In 2013, Xun et al. inferred the use of anycast for DNS top-level domain by using CHAOS query and traceroute.
- In 2015, Cicalese et al. introduced iGreedy, a method based on the Great-Circle Distance (GCD), performing later a census of anycast deployment on the Internet.
- In 2019, Bian et al. proposed a passive approach to detect anycast prefixes that did not rely on any active measurements, but rather used public BGP data from route collectors

Challenges in detecting anycast

- Xun et al. approach relies on the "collaboration" of the server, which had to provide information regarding the anycast instances.
- iGreedy requires a location-aware large-scale platform (such as PlanetLab or RIPE Atlas) in order to perform the measurements and has a significant footprint in terms of traffic Moreover, iGreedy is sensitive to latency.
- Passive approach of Bian et al. requires reliable ground-truth data.
- How can we responsibly and efficiently perform a regular census of anycast deployment?

MAnycast^2: Overview

- In this paper, we propose *MAnycast^2 - Measuring anycast using anycast*: a new measurement and inference technique to efficiently detect anycast prefixes.
- MAnycast^2 is inspired by De Vries et al. NOMS 2020 study.
- We use anycast vantage points (VPs) as sources to infer whether a set of target destination IP addresses are themselves anycast.

Working principle

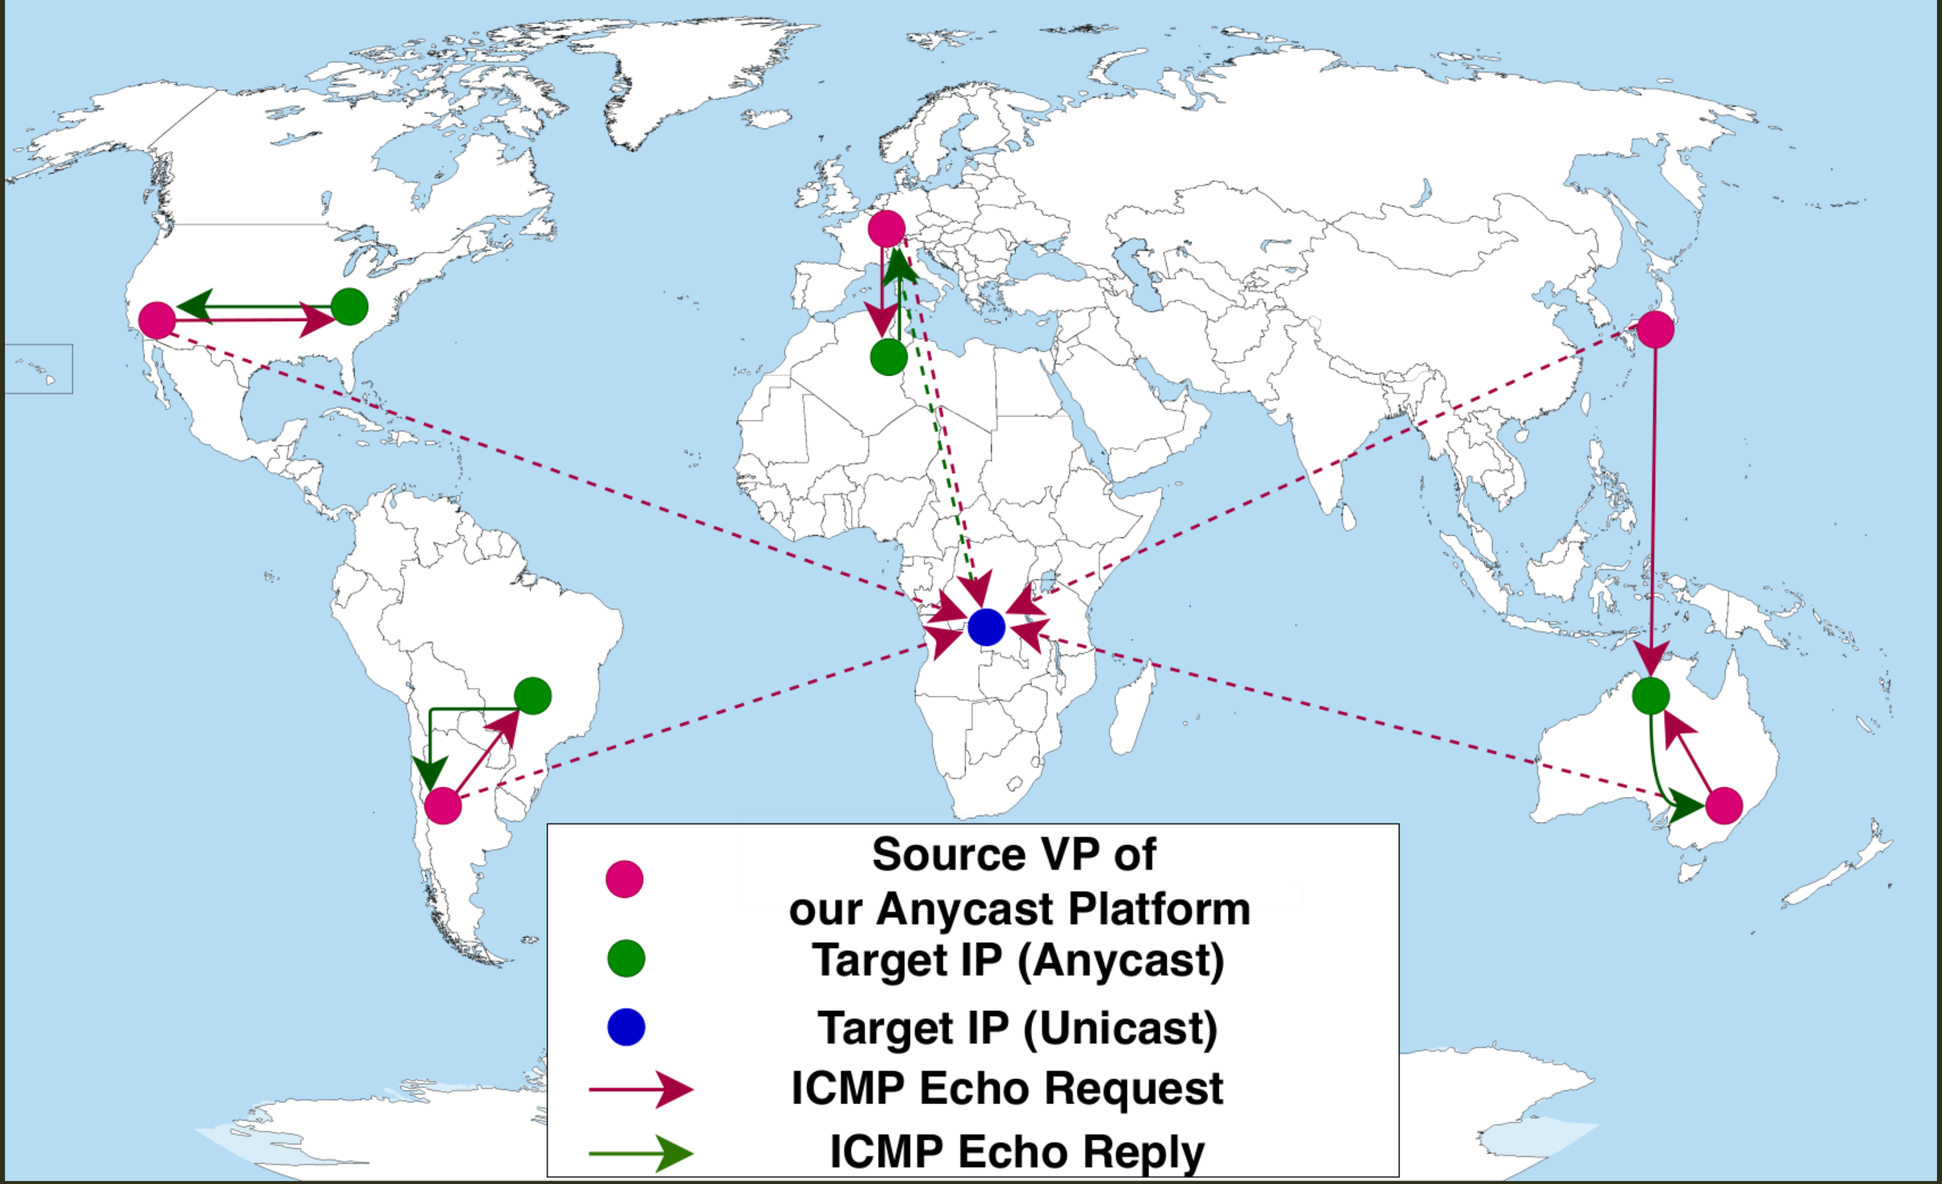

Unicast Detection

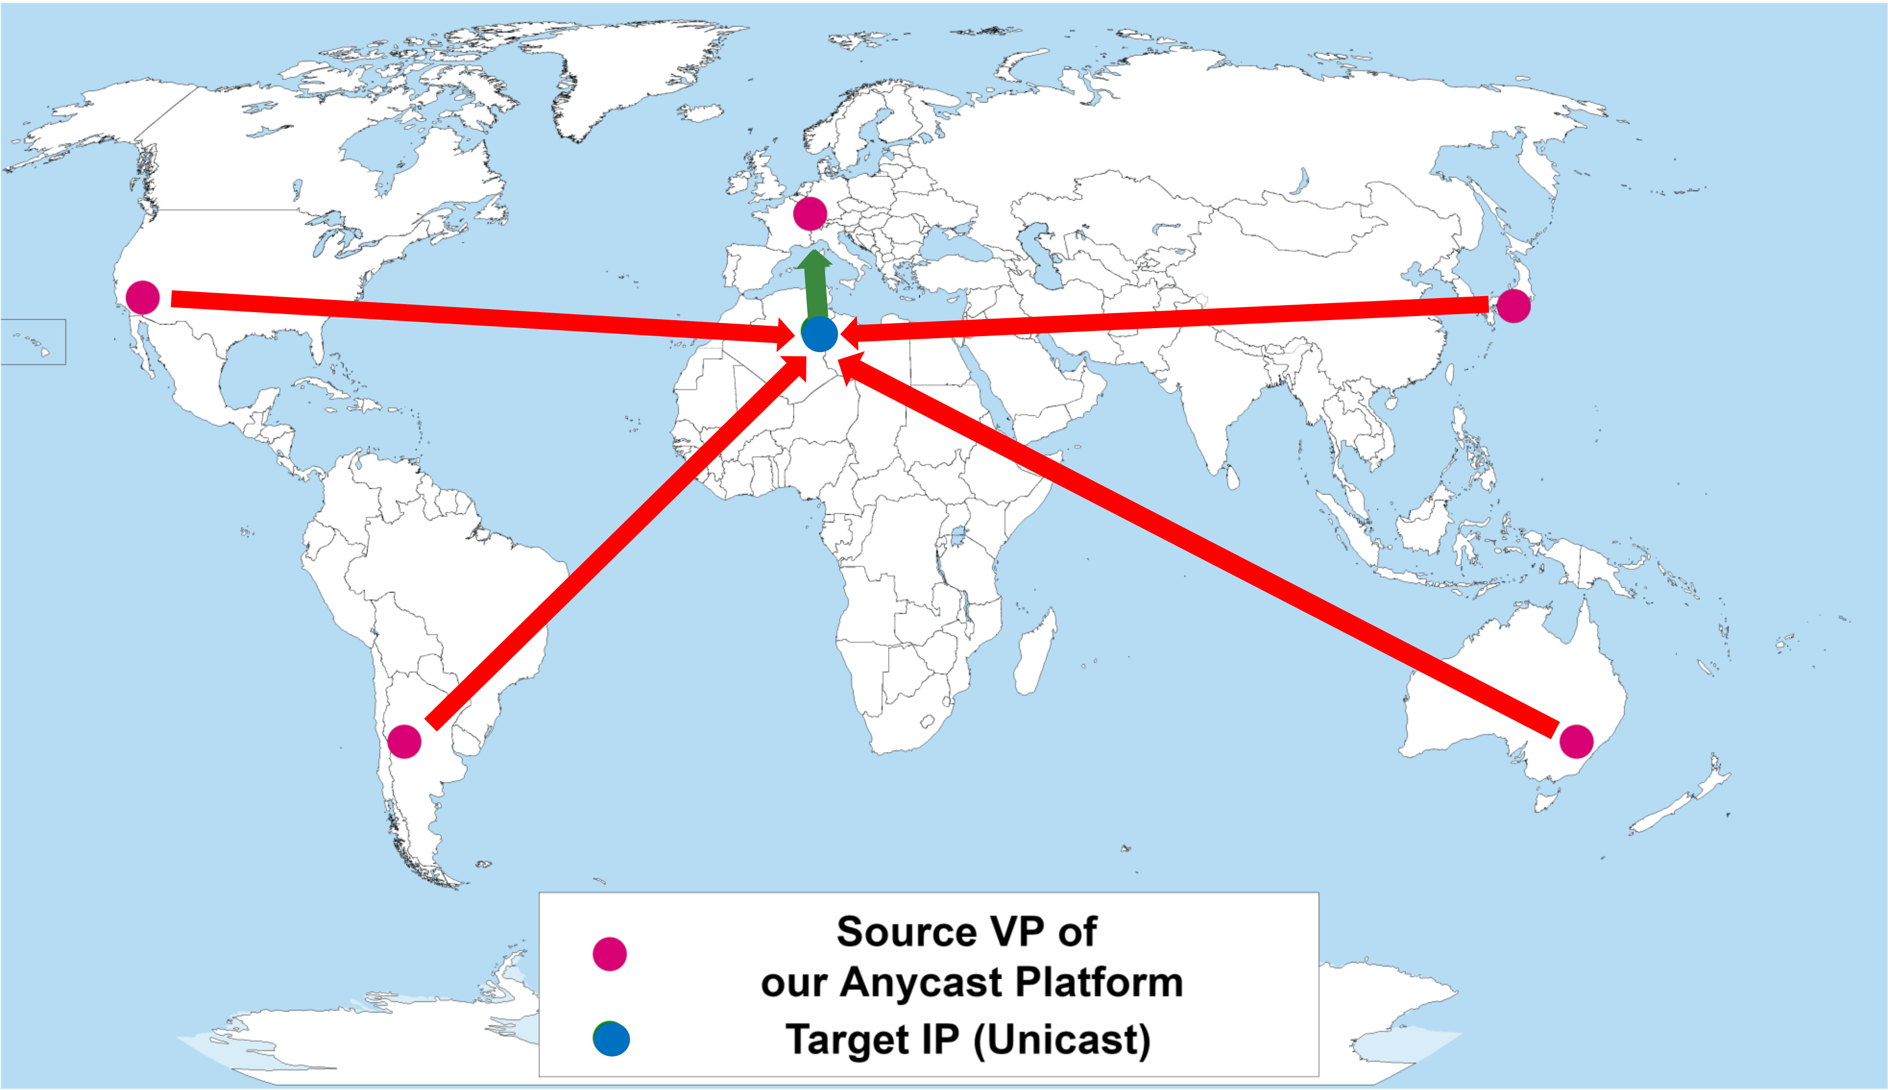

Anycast Detection

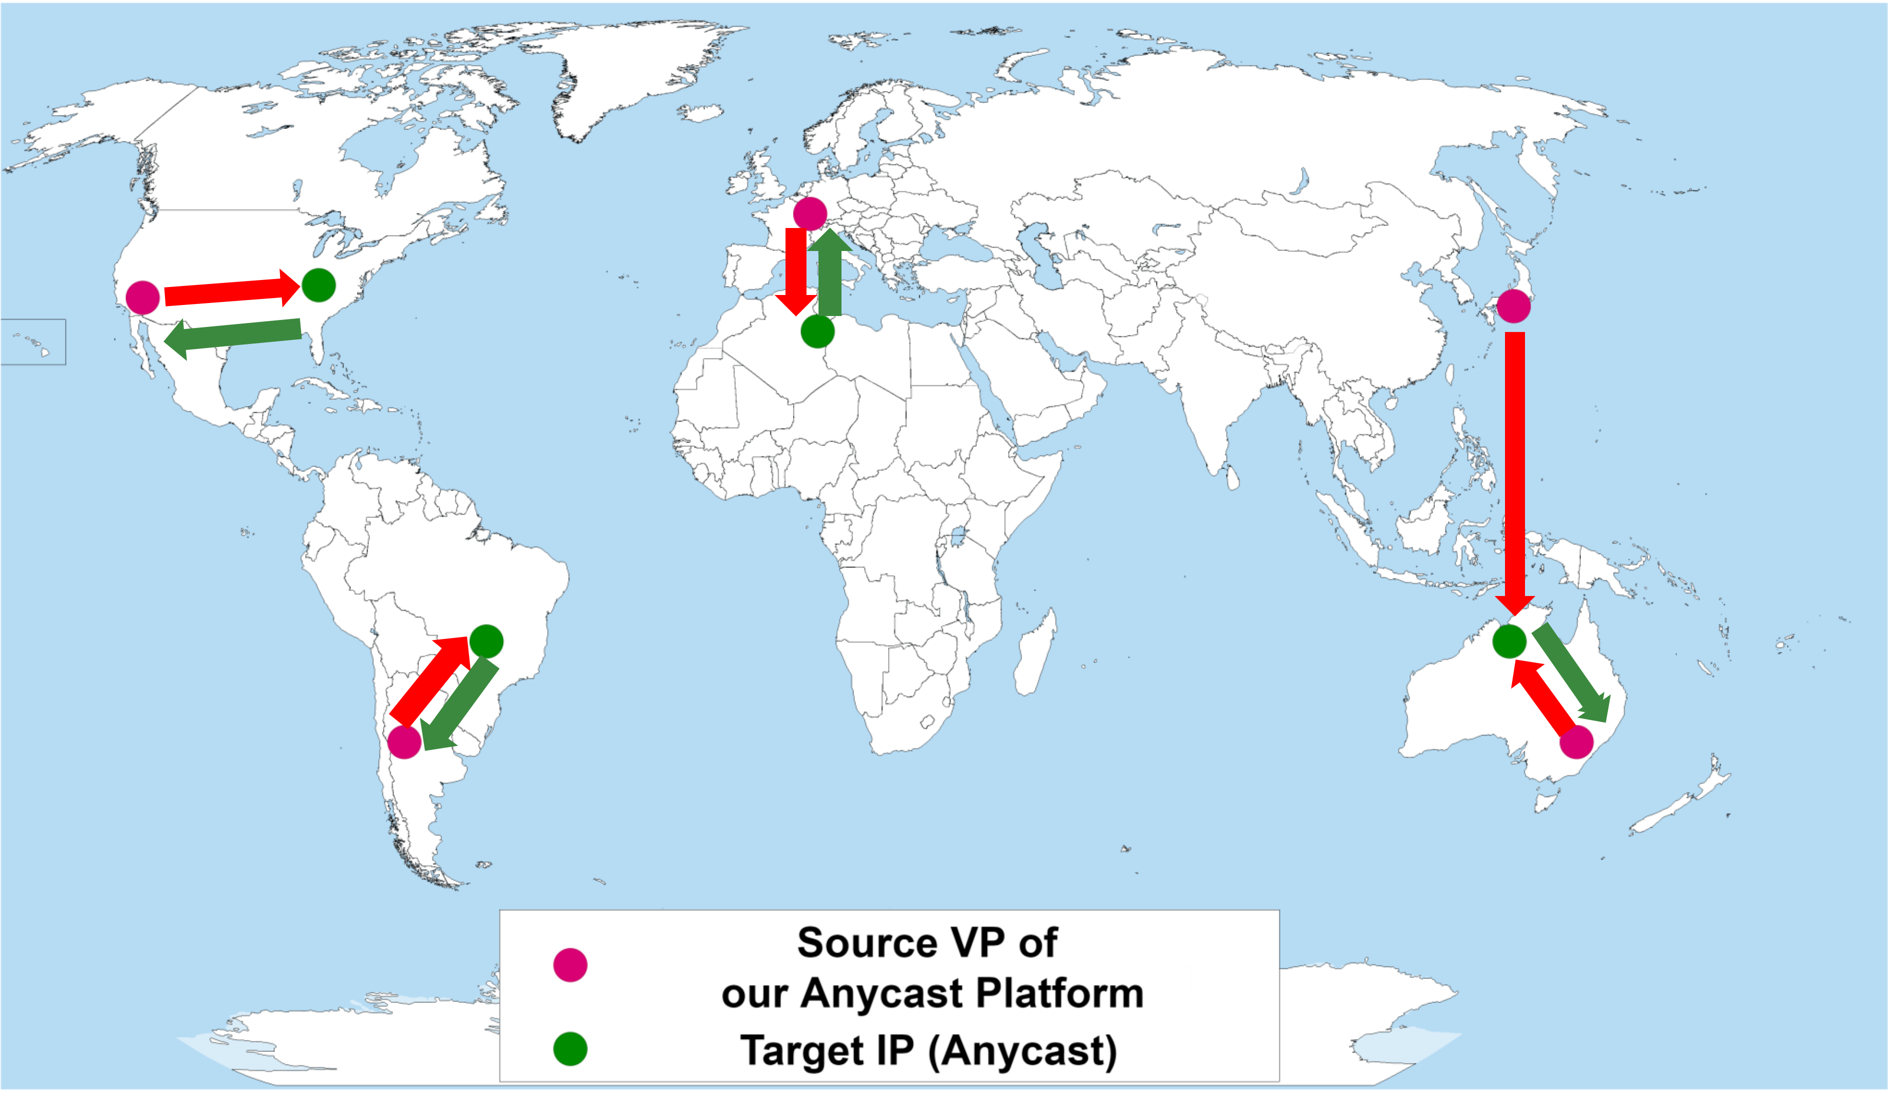

Methodology

We use the Tangled framework to implement our anycast measurement infrastructure.

Tangled has ten anycast instances (VPs), receiving transit from a combination of ISPs, commercial data-centers, academic networks, and IXPs.

For probing we use Verfploeter, a global probing system running on the Tangled testbed that was developed to monitor anycast catchment distributions.

Our implementation infers whether an IP address is unicast or anycast with one ping from each of ten VPs.

Preliminary Results

- On May 5, 2020, we tested all 6,125,756 target /24 prefixes from the ISI IPv4 hitlist in ~2.5 hours, choosing 1 IP for each /24 prefix.

- 3.47 million (56.5%) responded to at least one probe.

- For 3.45 million /24 prefixes (99.55% of the responding prefixes), only one VP received ICMP Echo Replies, so we classified these prefixes as unicast.

- For 15,540 (0.45%) /24 prefixes, between 2 and 10 VPs received ICMP Echo Replies. We considered them candidate anycast prefixes.

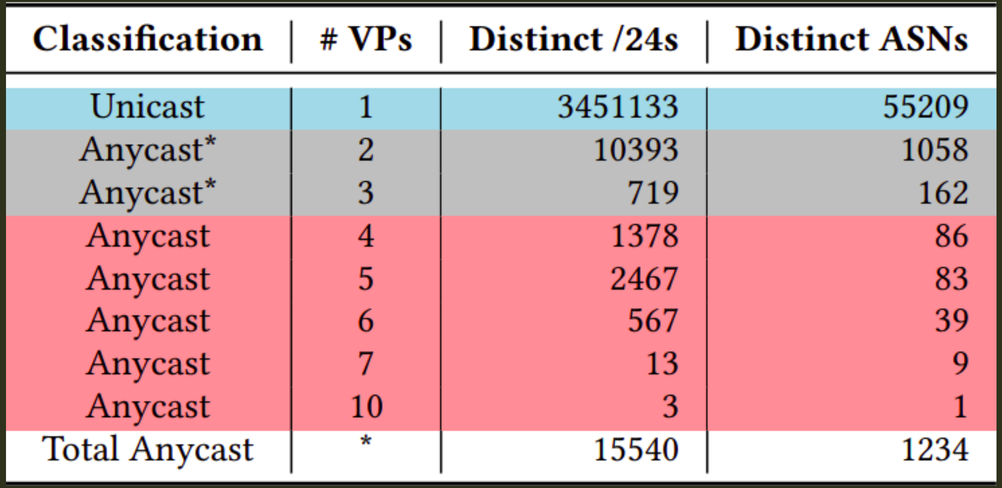

Classification	**# VPs**	**Distinct /24s**	**Distinct ASNs**
Unicast	1	3451133	55209
Anycast*	2	10393	1058
Anycast*	3	719	162
Anycast	4	1378	86
Anycast	5	2467	83
Anycast	6	567	39
Anycast	7	13	9
Anycast	10	3	1
Total Anycast	*	15540	1234

Ground Truth Validation

- Our first approach to validation used well known anycast services such as DNS root servers and public DNS resolvers.
- We correctly classifed all root servers as anycast except C-Root.
- C-Root's misclassification was a false negative (i.e., we failed to detect an anycast IP).
- We also correctly classified the anycast prefixes serving the public DNS resolvers of CloudFlare, OpenDNS, and Quad9.
- We incorrectly classified as unicast the prefix for the Google Public DNS Resolver.

Validation from AS Operators

Org Name	MAnycast^2	Operator
Cloudflare	3127	Confirmed
PCH	134	134
Amazon	4870	524
Akamai	212	90
Microsoft	75	51

Comparison against iGreedy

- We compared our results with the iGreedy technique.
- We used 200 random RIPE Atlas probes, as geographically diverse unicast measurement nodes.
- We sampled ~2% of prefixes we identified as unicast and take all the prefixes we identified as anycast to run iGreedy against them.
- In total, we ran iGreedy on 82,270 /24 prefixes.

Comparison against iGreedy

- We observed low false negative rate (0.1%) for the sample unicast prefixes.
- For answers received on 4 or more VPs, our results almost agrees always with iGreedy.
- The disparity was extremely high when we receive answer on 2 or 3 VPs

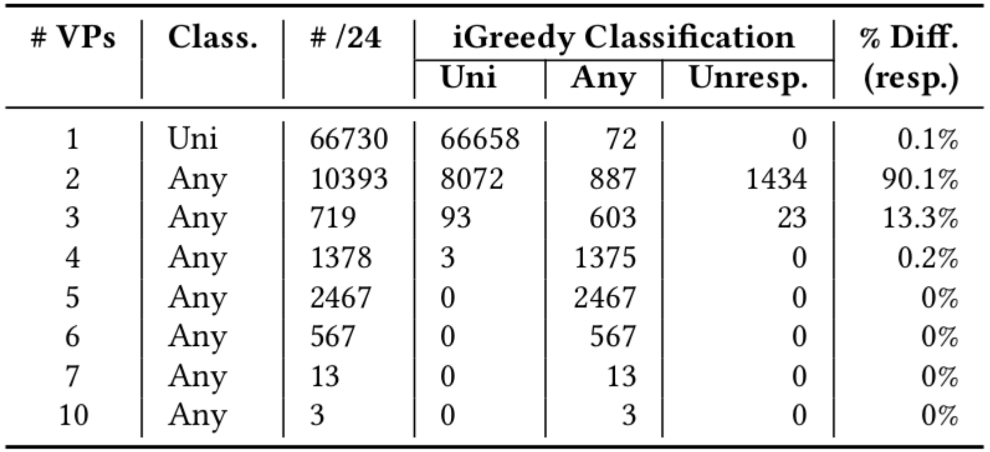

# VPs	Class.	# /24	iGreedy Classification			% Diff.
			Uni	Any	Unresp.	(resp.)
1	Uni	66730	66658	72	0	0.1%
2	Any	10393	8072	887	1434	90.1%
3	Any	719	93	603	23	13.3%
4	Any	1378	3	1375	0	0.2%
5	Any	2467	0	2467	0	0%
6	Any	567	0	567	0	0%
7	Any	13	0	13	0	0%
10	Any	3	0	3	0	0%

Considerations on Results

We have a low false negative rate and a low or zero false positive rate for answers received on 4 or more VPs.

However, MAnycast² misclassifies two prominent anycast services (C-Root and Google Public DNS).

We also find ambiguous results when only 2-3 VPs receive responses.

These issues draws open challenges for our methodology.

Conditions for Success

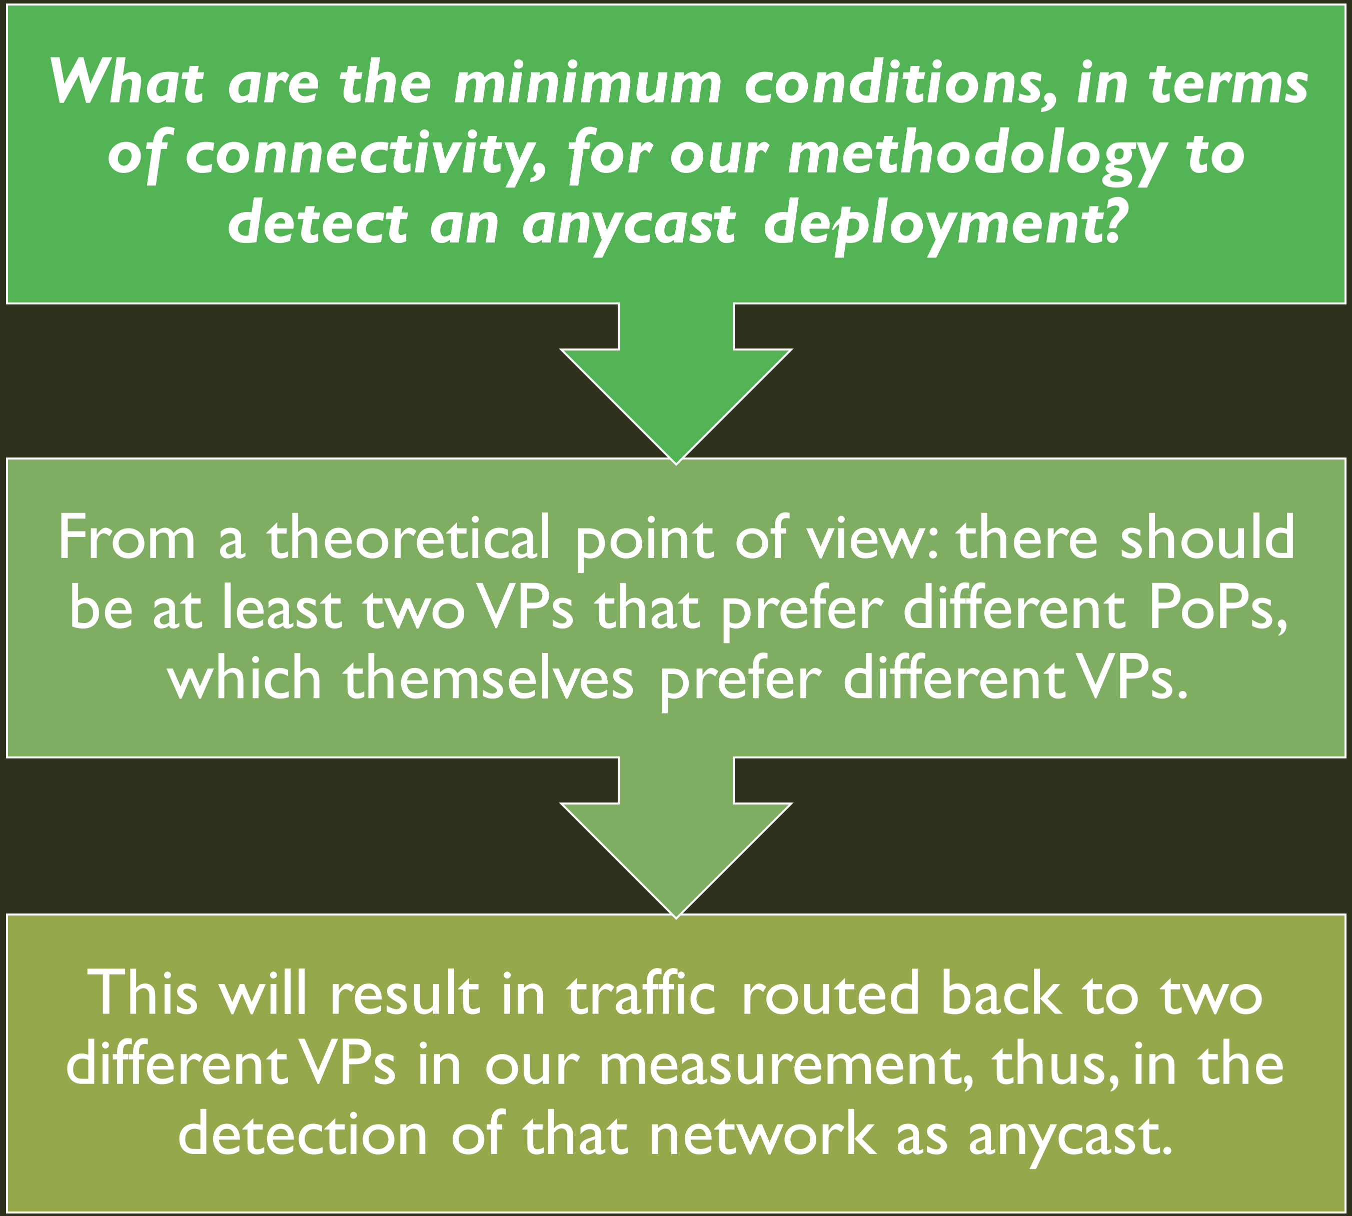

What are the minimum conditions, in terms of connectivity, for our methodology to detect an anycast deployment?

From a theoretical point of view: there should be at least two VPs that prefer different PoPs, which themselves prefer different VPs.

This will result in traffic routed back to two different VPs in our measurement, thus, in the detection of that network as anycast.

What happen if conditions are not met?

- MAnycast 2 misclassified C-Root as unicast, because all the answers were received on the London VP node.
- This behaviour is due to the fact that Cogent consider our London upstream provider (Vultr) as preferred-route.
- In the same way, Google prefer to route all the packet to the São Paulo IXP, where our testbed is directly interconnected with them.
- These examples establish an open challenge for our methodology: understanding and accommodating preferred routing strategies from large network operators.

Routing Flaps and Load Balancing

- Another routing phenomenon, which bring us to misclassify unicast prefixes as anycast, are routing flaps and load balancing.
- This is mostly likely to happen when we receive responses at only two (or occasionally three) VPs.
- A key factor seems to be the time that elapses between probing a target IP address from distinct VPs.
- Reduced this time and repeating measurements can help to resolve some incorrect classifications.

Regional and Topological Blindspots

- Our method's accuracy appears to vary by region, due to variation in density of connectivity relative to different VPs in our testbed.
- These may prevent detection of regional anycast services.
- Latency-based approaches face similar challenges in detecting small anycast deployments.
- Regional anycast services are challenging to detect and require a widely distributed geographical infrastructure with many nodes.

Considerations on Applicability

- At the actual stage, a possible use of our methodology is to filter out, efficiently and at scale, unicast addresses.
- Then, one can apply the heavier-weight latency-based method on a smaller remaining set of prefixes for which we are uncertain (2 or 3 VPs).
- The combined approach provides classification results close to iGreedy with a substantially reduced measurement overhead.
- Another improvement could be, when VPs are in IXP, considering each incoming upstream connection as a separate VPs.
- We repeated the measurment with the PEERING Testbed showing an overlap of 90% with the measurement performed on the Tangled Testbed.

Conclusions

- Our contribution is MAnycast^2, a new measurement technique based on the idea of using anycast to measure anycast.
- Major characteristic of MAnycast^2 is that it is lightweight and scalable, with a low false-negative rate and low false-positive rate, under certain condition.
- Our results, compared to the state-of-the-art latency-based methodology shows promising results expecially in combined approaches.
- Future improvements to our methodology will focus on reducing the false-negative classification rate.

Questions?

r.sommese@utwente.nl