

# Locating and Enumerating Anycast a Comparison of Two Approaches

Remi Hendriks ; Tim Betzer ; Ben Du

Raffaele Sommese ; Mattijs Jonker ; Roland van Rijswijk-Deij

**UNIVERSITY  
OF TWENTE.**

Technische  
Universität  
München

**TUM**



# What is anycast?

Anycast is the geographical distribution of Internet services

- Achieved by announcing IP addresses at multiple locations
- Widely used to provide resilience, reduce latency, ...
- Examples include
  - Content Delivery Networks (CDNs)
  - Critical infrastructure (e.g., DNS)

# What is anycast?



*Cloudflare's global network*  
<https://www.cloudflare.com/en-gb/network/>

# Motivation

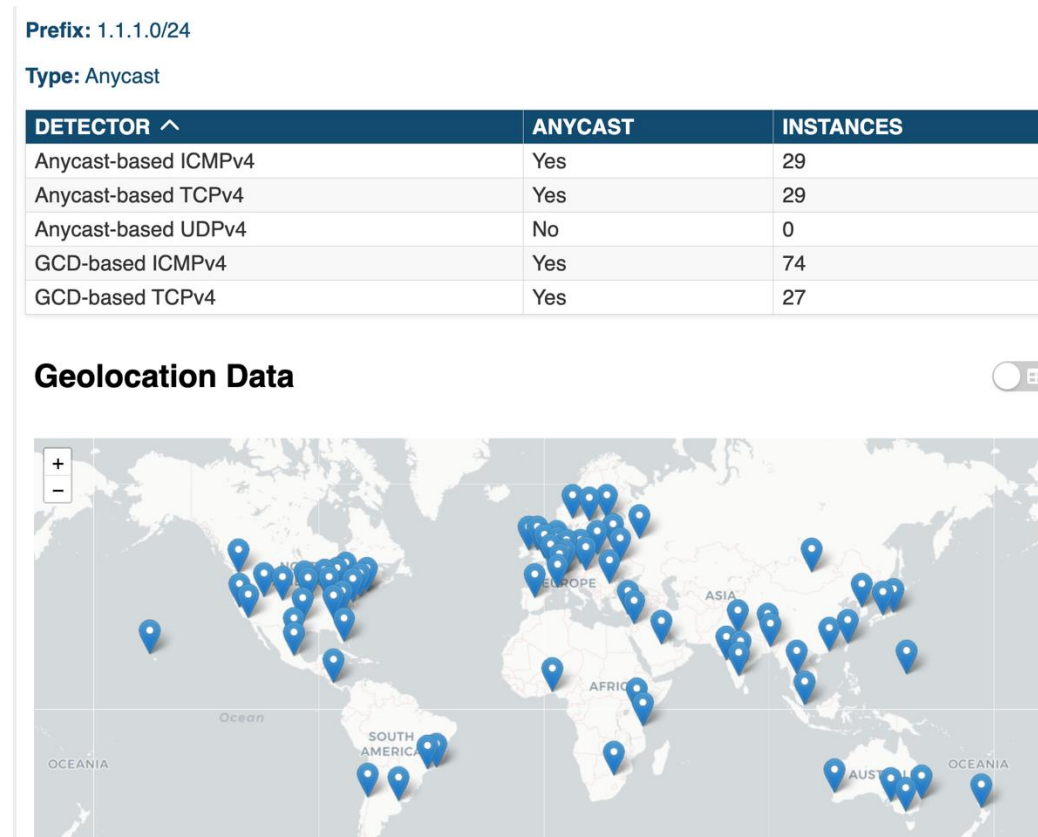
- Anycast is one of the most effective distribution and **resilience** techniques
  - Used for critical Internet services (e.g., DNS)
  - Deployed by CDNs for low-latency, reliability, ..
  - DDoS mitigation (used to provide DDoS protection services)

# Motivation

- Anycast is one of the most effective distribution and **resilience** techniques
  - Used for critical Internet services (e.g., DNS)
  - Deployed by CDNs for low-latency, reliability, ..
  - DDoS mitigation (used to provide DDoS protection services)
- Due to anycast opaqueness, a census is needed
- Knowing what is anycast (and where it is located) is useful for
  - Making better Traffic engineering decisions
  - Troubleshooting network problems
  - Resilience assessment of third parties
  - Assessing location of services (e.g., GDPR compliance)

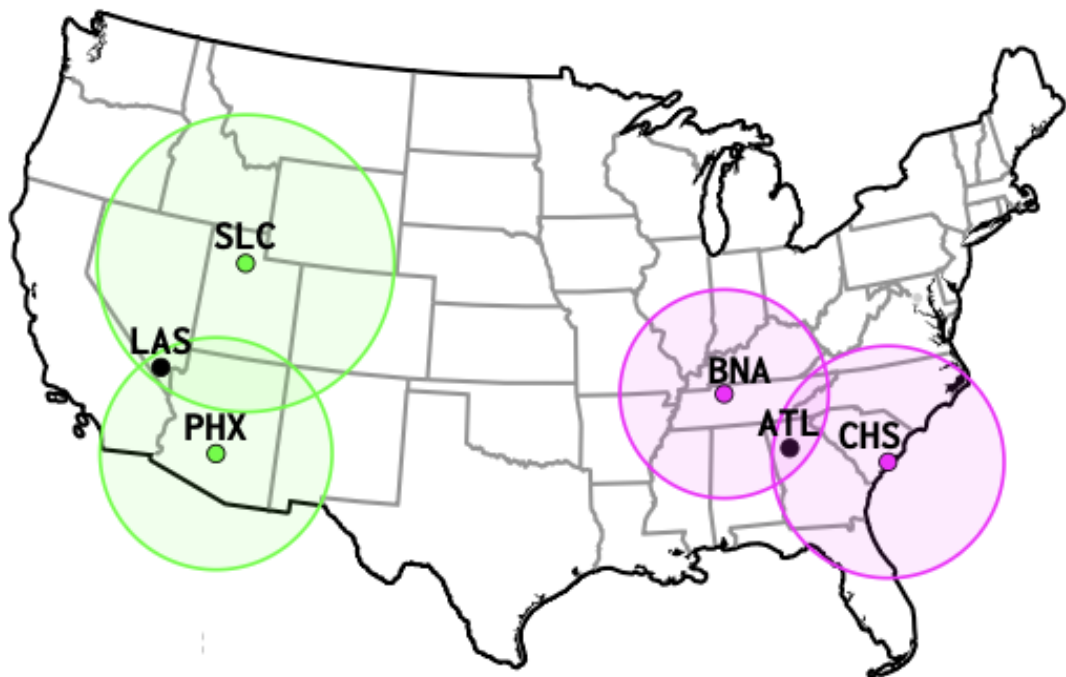
# State of the art

- Current anycast geolocation dataset [manycast.net](https://manycast.net) [1]

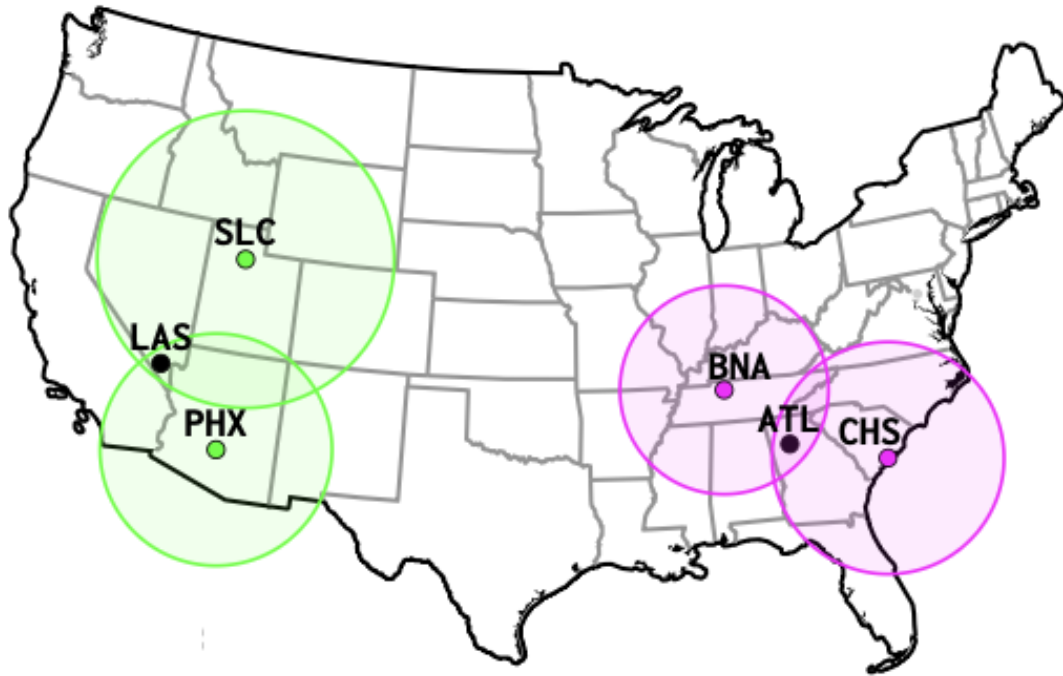


# State of the art Methodology

- Uses iGreedy [2] to geolocate anycast
- Speed-of-light in fibre optics and latency measurements
- Geolocates using minimum set of overlapping circles
- Infers the anycast site to be at the largest metropolitan city



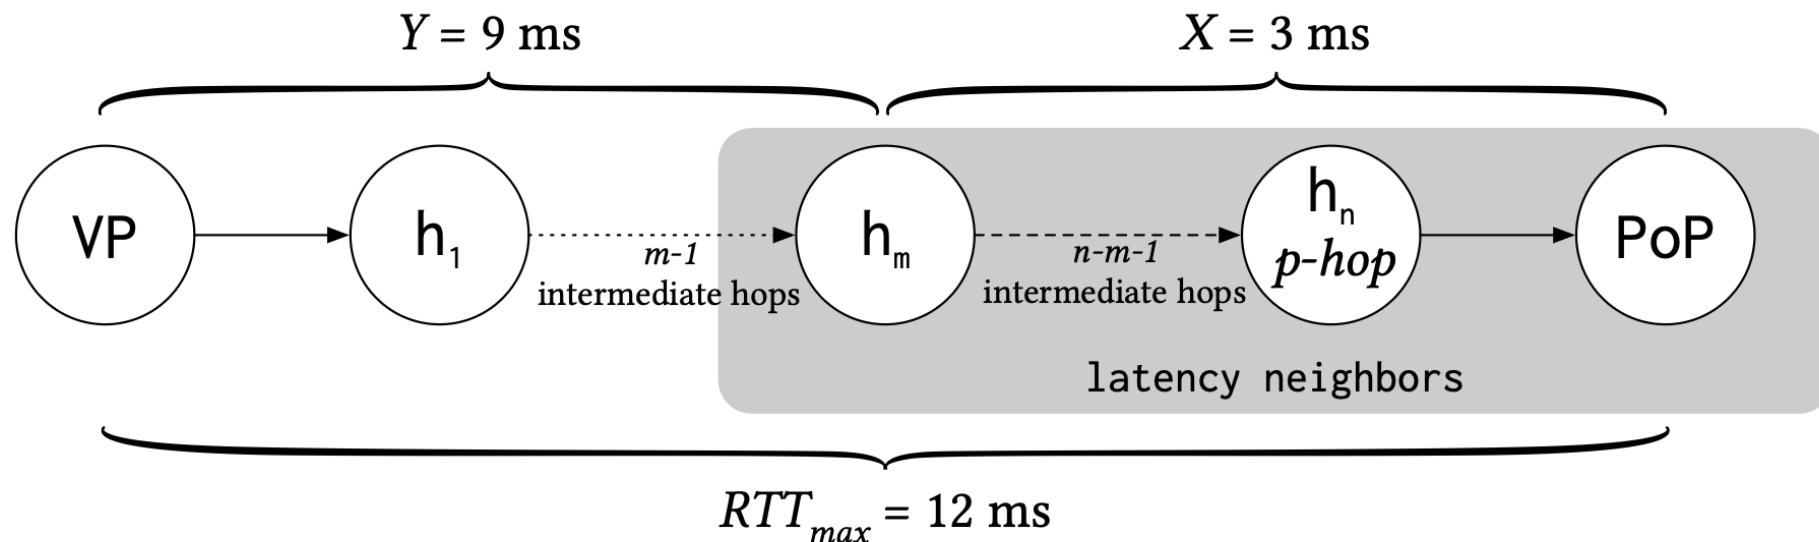
# State of the art Methodology



- Uses iGreedy [2] to geolocate anycast
- Speed-of-light in fibre optics and latency measurements
- Geolocates using minimum set of overlapping circles
- Infers the anycast site to be at the largest metropolitan city
- However...
  - Latency measurements are noisy
  - Known to suffer with anycast sites in near proximity (e.g., regional anycast)

# Improving the state-of-the-art methodology

- Traceroute is used to improve unicast geolocation
  - Can it improve anycast geolocation?
- Our approach uses 'latency neighbors' [3]
  - Latency neighbors are near the anycast site reached
  - Geolocate latency neighbor using Hoiho/IPInfo -> infer anycast site location



# Measurement campaign

- Measurements performed using CAIDA's Archipelago
  - 274 VPs inside 66 distinct countries covering 178 unique ASes
  - Targets 13,765 anycast /24s from census
  - Totaling 3.4M traceroutes
- Side-to-side comparison with state-of-the-art (iGreedy) using the same VPs



# Results

## Is it more accurate?

- Accuracy assessment using root-letter deployments
  - Large variety in deployment sizes and distribution
  - Ground-truth available (root-servers.org)

### Root Servers

The screenshot shows the 'Root Servers' page for Verisign, Inc. The operator is listed as 'Verisign, Inc.'. Below the operator name, there are navigation buttons for 'Homepage', 'Statistics', 'Peering Policy', and 'Email'. The 'Locations' section shows a total of 57 sites, with a list of locations including Amsterdam, NL; Ashburn, US; El Segundo, US; Frankfurt, DE; Guangzhou, CN; London, GB; LOS ANGELES, US; Manassas, US; Marseille, FR; Miami, US; New York, US; Paris, FR; Plano, US; San Jose, US; Seattle, US; Singapore, SG; Stockholm, SE; Tokyo, JP; and Washington DC, US.

Operator: Verisign, Inc.

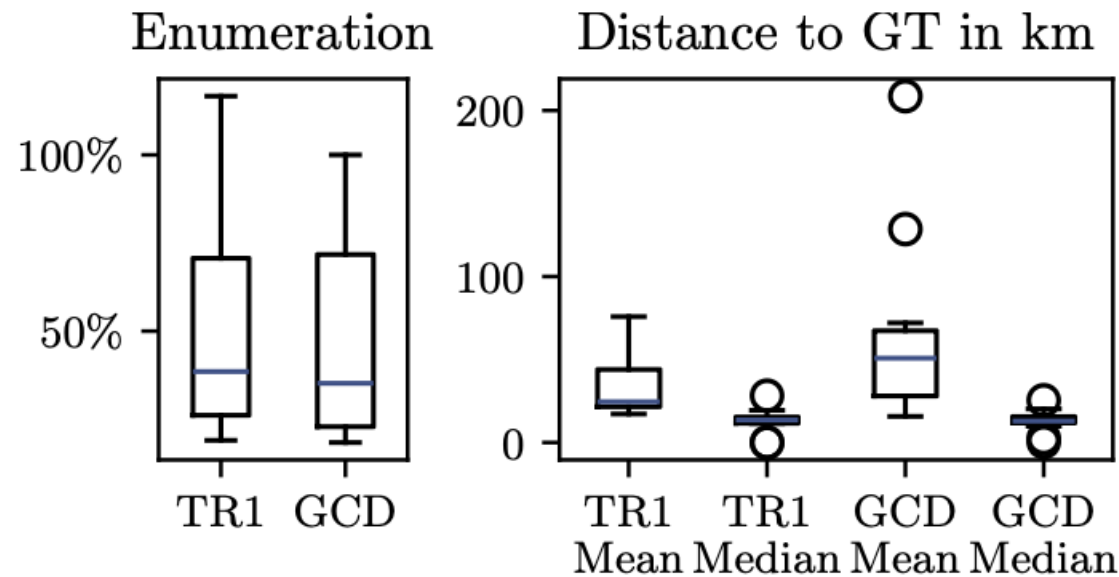
Locations: Sites: 57

- Amsterdam, NL
- Ashburn, US
- El Segundo, US
- Frankfurt, DE
- Guangzhou, CN
- London, GB
- LOS ANGELES, US
- Manassas, US
- Marseille, FR
- Miami, US
- New York, US
- Paris, FR
- Plano, US
- San Jose, US
- Seattle, US
- Singapore, SG
- Stockholm, SE
- Tokyo, JP
- Washington DC, US

# Results

## Is it more accurate?

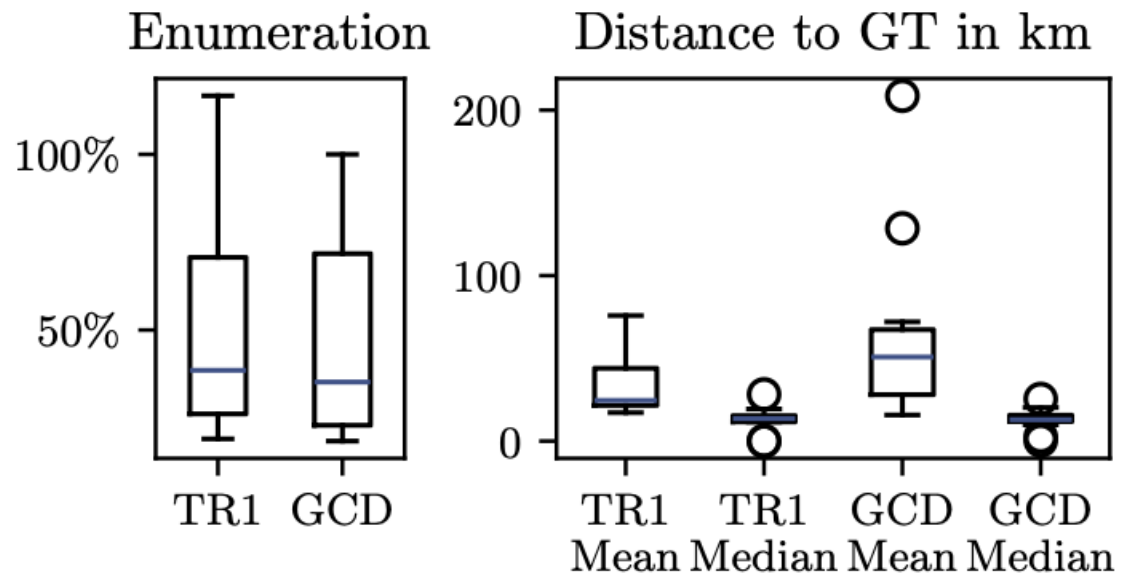
- Accuracy assessment using root-letter deployments
  - Large variety in deployment sizes and distribution
  - Ground-truth available (root-servers.org)
- TR1 = Traceroute methodology, GCD = iGreedy



# Results

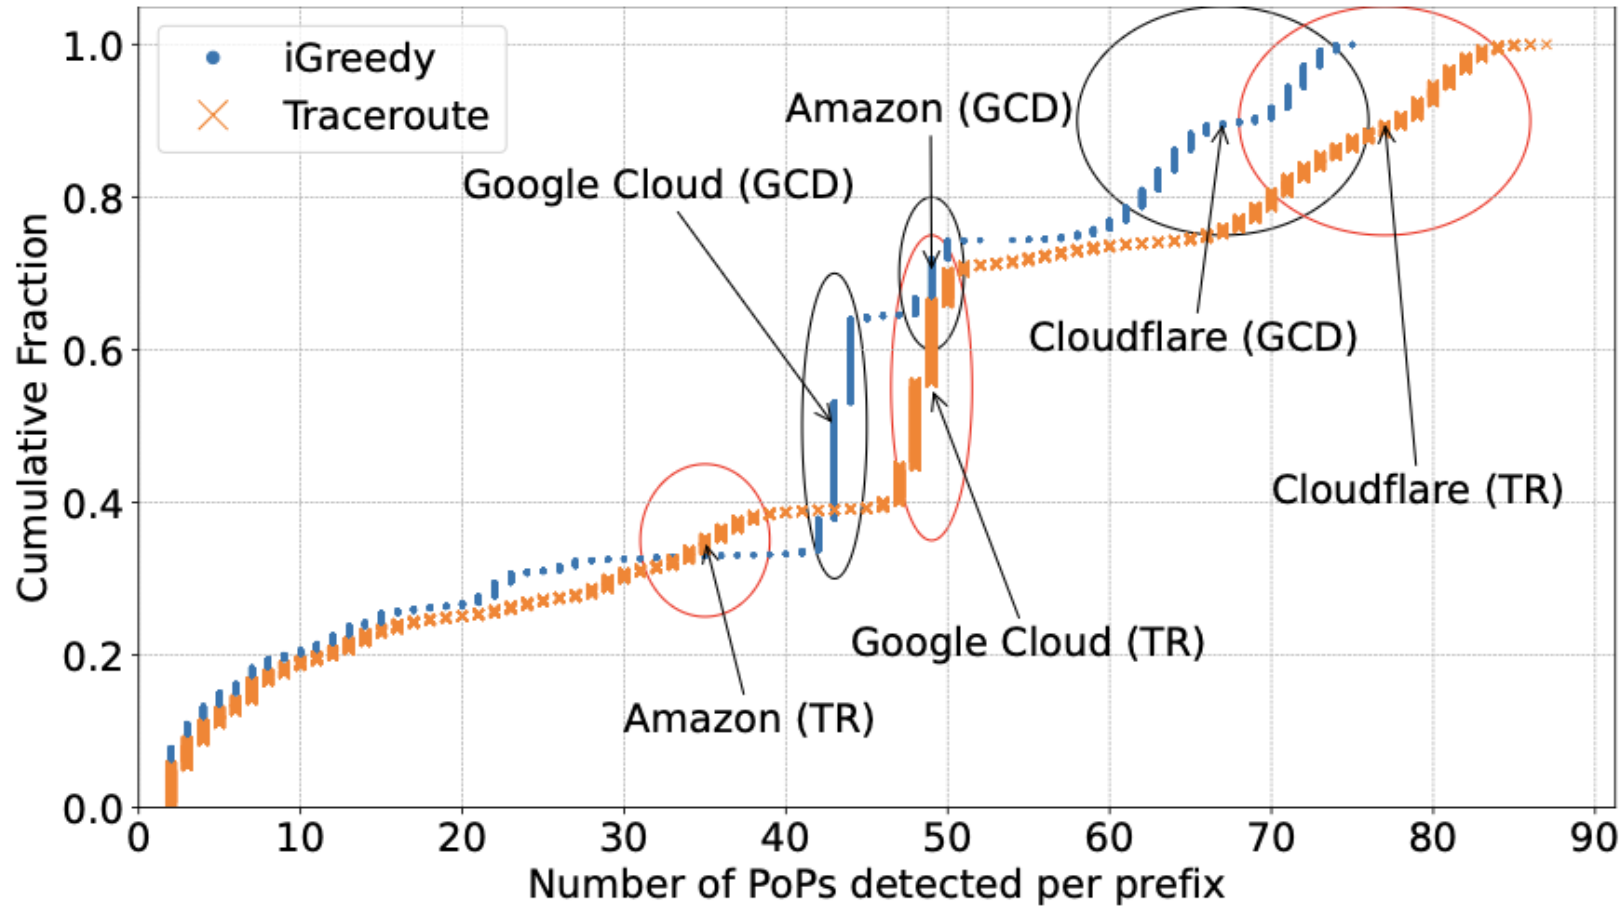
## Is it more accurate?

- Accuracy assessment using root-letter deployments
  - Large variety in deployment sizes and distribution
  - Ground-truth available (root-servers.org)
- TR1 = Traceroute methodology, GCD = iGreedy
- **Takeaways:**
  - **Traceroute achieves higher enumeration**
  - **\* but suffers from overcounting**
  - **Traceroute achieves higher accuracy**  
(median error 21km against 51km)



# Results

## When does it find more?

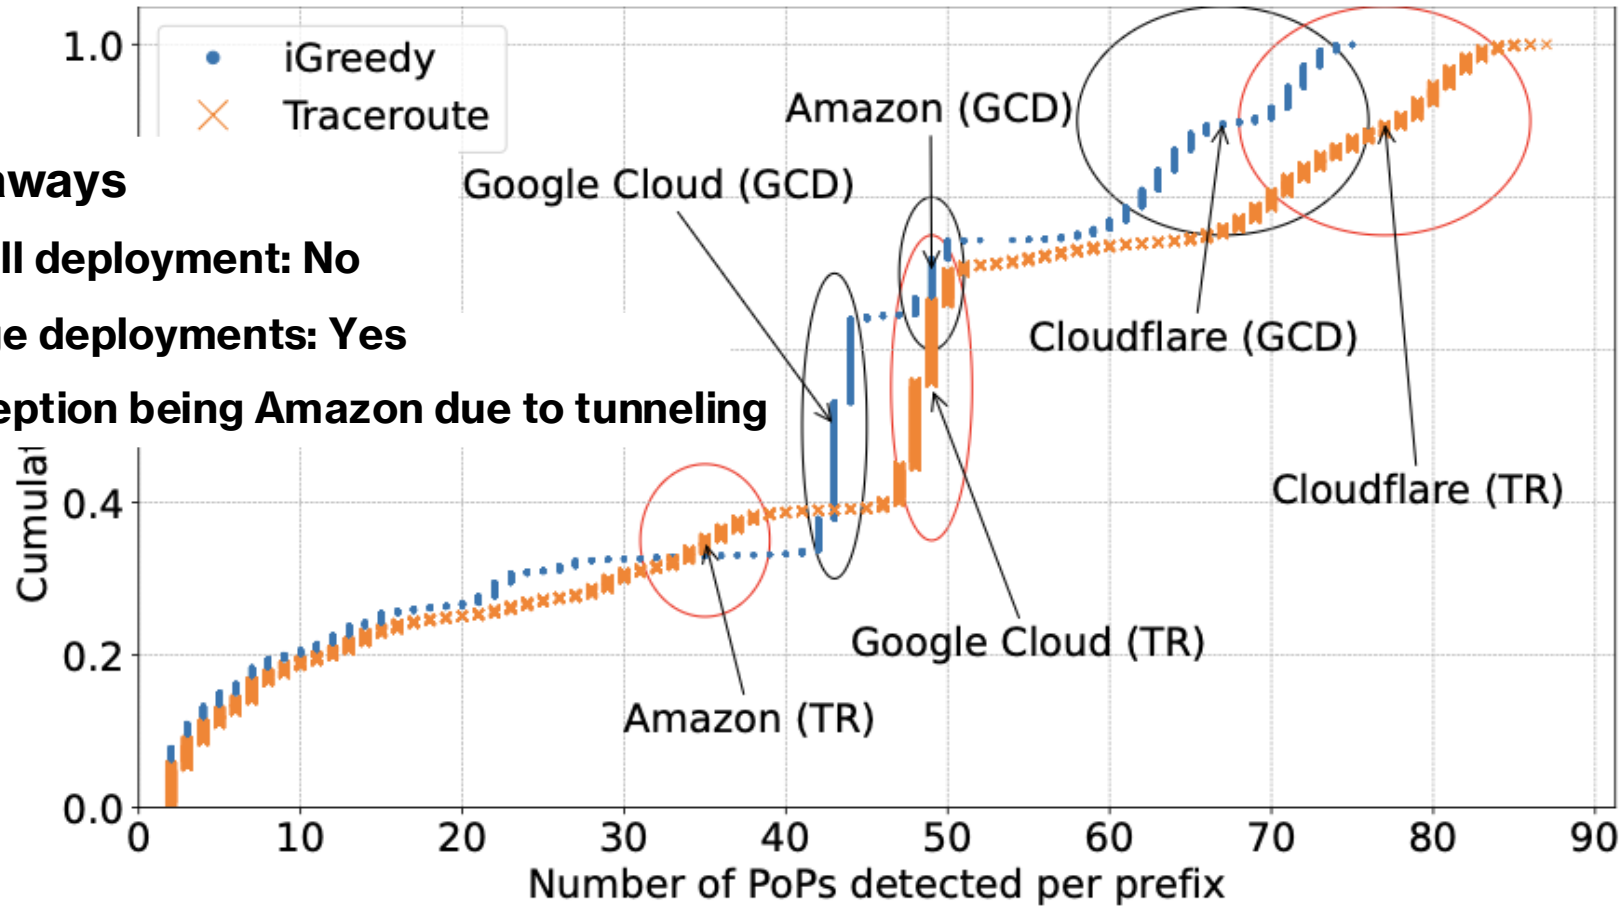


# Results

## When does it find more?

- **Takeaways**

- **Small deployment: No**
- **Large deployments: Yes**
- **Exception being Amazon due to tunneling**



# Results

## What is the probing cost?

Technique	Probes	# of VPs	Total
TR (full)	9.83	274	2,693
TR (skipping same-AS hops)	6.96	274	1,907
TR (latency constrained)	8.40*	124 <sup>†</sup>	1,316
TR (reduced)	5.53*	124 <sup>†</sup>	960
iGreedy	1	274	274

\*Initial measurement required from all 274 VPs.

<sup>†</sup> Average number of VPs within 12 ms RTT.

**Table 3: Probing cost per target breakdown.**

# Results

## What is the probing cost?

- **Takeaways**
  - Probing cost can be reduced by limiting traceroutes
  - Final cost 4X that of iGreedy

<b>Technique</b>	<b>Probes</b>	<b># of VPs</b>	<b>Total</b>
<b>TR (full)</b>	9.83	274	2,693
<b>TR (skipping same-AS hops)</b>	6.96	274	1,907
<b>TR (latency constrained)</b>	8.40*	124 <sup>†</sup>	1,316
<b>TR (reduced)</b>	5.53*	124 <sup>†</sup>	960
<b>iGreedy</b>	1	274	274

\*Initial measurement required from all 274 VPs.

<sup>†</sup> Average number of VPs within 12 ms RTT.

**Table 3: Probing cost per target breakdown.**

# Discussion and future work

- Traceroute methodology outperforms iGreedy
  - Both in enumeration and geolocation accuracy (mean error 26km compared to 51km)
  - However, 4x probing cost makes it unattractive as a census technique
  - Suffers from overcounting in rare cases
- Future work
  - Run methodology using historical traceroute datasets
  - Detect causes of sub-optimal anycast routing
- Traceroute dataset, code, and more details available in the paper

## Locating and Enumerating Anycast: a Comparison of Two Approaches

Remi Hendriks  
University of Twente  
Enschede, the Netherlands  
remi.hendriks@utwente.nl

Tim Betzer  
Technical University of Munich  
Munich, Germany  
betzer@net.in.tum.de

Ben Du  
CAIDA/UC San Diego  
La Jolla, CA, USA  
bendu@ucsd.edu

Raffaele Sommese  
University of Twente  
Enschede, the Netherlands  
r.sommese@utwente.nl

Mattijs Jonker  
University of Twente  
Enschede, the Netherlands  
m.jonker@utwente.nl

Roland van Rijswijk-Deij  
University of Twente  
Enschede, the Netherlands  
r.m.vanrijswijk@utwente.nl

### Abstract

Anycast allows for providing services from multiple, geographically distant Points of Presence (PoPs), using a single IP address, to, e.g., improve resilience. Due to its opaqueness, it is often unknown which addresses are provisioned using anycast and, if so, where the PoPs are located. As anycast is widely used for critical Internet infrastructures (e.g., the DNS) efforts have been made to map anycast deployments. The current state-of-the-art mapping technique, iGreedy, relies on latency-based measurements, and is adversely affected by noise caused by, e.g., network processing delays. Previous work has shown that traceroute can alternatively be used to detect anycast. As traceroute reveals the hops a packet traverses, it may also be used to locate sites using geolocation data for hops near the anycast PoPs.

This paper is the first to assess the performance of the traceroute-based approach at scale, by targeting 14k prefixes from an anycast census. Using ground truth we show traceroute achieves a slight increase in enumeration and geolocation precision over iGreedy. However, it suffers from overestimating the number of PoPs and incurs a 4x increase in probing cost, making it unattractive for anycast censuses.

### CCS Concepts

• Networks → Network measurement.

### Keywords

Anycast; Geolocation; Traceroute



This work is licensed under a Creative Commons Attribution 4.0 International License.  
ANRW '25, Madrid, Spain  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2009-3/2025/07  
<https://doi.org/10.1145/3744200.3744783>

### ACM Reference Format:

Remi Hendriks, Tim Betzer, Ben Du, Raffaele Sommese, Mattijs Jonker, and Roland van Rijswijk-Deij. 2025. Locating and Enumerating Anycast: a Comparison of Two Approaches. In *Applied Networking Research Workshop (ANRW '25)*, July 22, 2025, Madrid, Spain. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3744200.3744783>

### 1 Introduction

Anycast is the practice of making an IP address available in multiple discrete locations [20]. This allows operators to offer replicated services nearer to clients and increase resilience through redundancy. Examples include DNS resolvers and nameservers, and CDNs providing web caching. While the number of anycast addresses on the Internet is small (<0.1% [18]) they serve a large share of Internet traffic [6].

For this reason, efforts have been made to map anycast deployments on the Internet to e.g., measure Internet resilience. The most notable are iGreedy [7] and MAnycast<sup>2</sup> [27], where the former, iGreedy, also enumerates and geolocates PoPs behind anycast addresses. iGreedy geolocates using latency measurements by probing an anycast address from multiple geographically distributed Vantage Points (VPs), similar to conventional IP unicast geolocation [32]. However, such latency measurements are known to be noisy due to e.g., network processing and queuing delays.

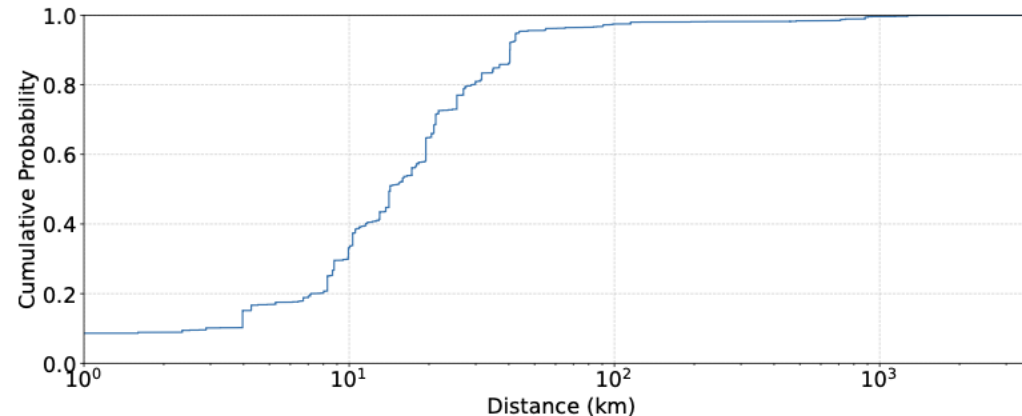
Unicast IP geolocation often involves the use of tools such as traceroute to improve precision [32]. Traceroute allows for measuring the path that an Internet packet takes to reach its target. When determining the location of a target, in case no location information is available for the target itself, traceroute can be used to find nearby hops for which location information is available. It has also been used to geolocate anycast using the location of hops near anycast PoPs, to infer the location of PoPs themselves [31]. One example is verifying GDPR compliance when sending personal data to anycast addresses [24]. However, traceroute has limitations: operators may block the ICMP Time Exceeded messages

# Appendix slides

# Inferring site reached

## Is the p-hop method valid?

- Previous work has used p-hop as a proxy of the anycast destination reached
- Results show that this is a valid approach
  - 95% of distant p-hops within 100km of a nearby anycast site





**Figure 4: Distances (log scale) between invalid latency neighbor *p*-hops and the nearest airport found using the latency neighbor method.**

# Opaqueness

## Anycast is opaque

- Unknown whether an address is unicast or anycast
- Unknown where anycast addresses are announced
- Commercial IP to location datasets return a single location

# Opaqueness ip2location

<b>Permalink</b>	<a href="https://www.ip2location.com/1.1.1.1">https://www.ip2location.com/1.1.1.1</a> 
<input checked="" type="checkbox"/> <b>IP Address</b>	1.1.1.1
<input type="checkbox"/> <b>Country</b>	 <a href="#">Australia</a> [AU]
<input type="checkbox"/> <b>Region</b>	Queensland
<input type="checkbox"/> <b>City</b>	Brisbane

# Opaqueness ipinfo.io

## IP Geolocation

City

Brisbane

State

Queensland

Country



Australia

# Opaqueness maxmind

IP Address	<u>Location</u>	<u>Network</u>	<u>Postal Code</u>	<u>Approximate Latitude / Longitude*, and Accuracy Radius</u>
1.1.1.1	-	1.1.1.0/24	-	- (-)