Anycast Discovery: Daily Mapping of Anycast

RIPE 89, PRAGUE, CZECH REPUBLIC

28 OCT – 1 NOV, 2024

Remi Hendriks

UNIVERSITY OF TWENTE.

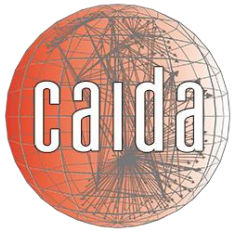

Our Team @ DACS

Remi Hendriks
PhD Candidate
University of Twente

Raffaele Sommese
Assistant Professor
University of Twente

Mattijs Jonker
Assistant Professor
University of Twente

Roland Van Rijswijk-Deij
Full Professor
University of Twente

Content

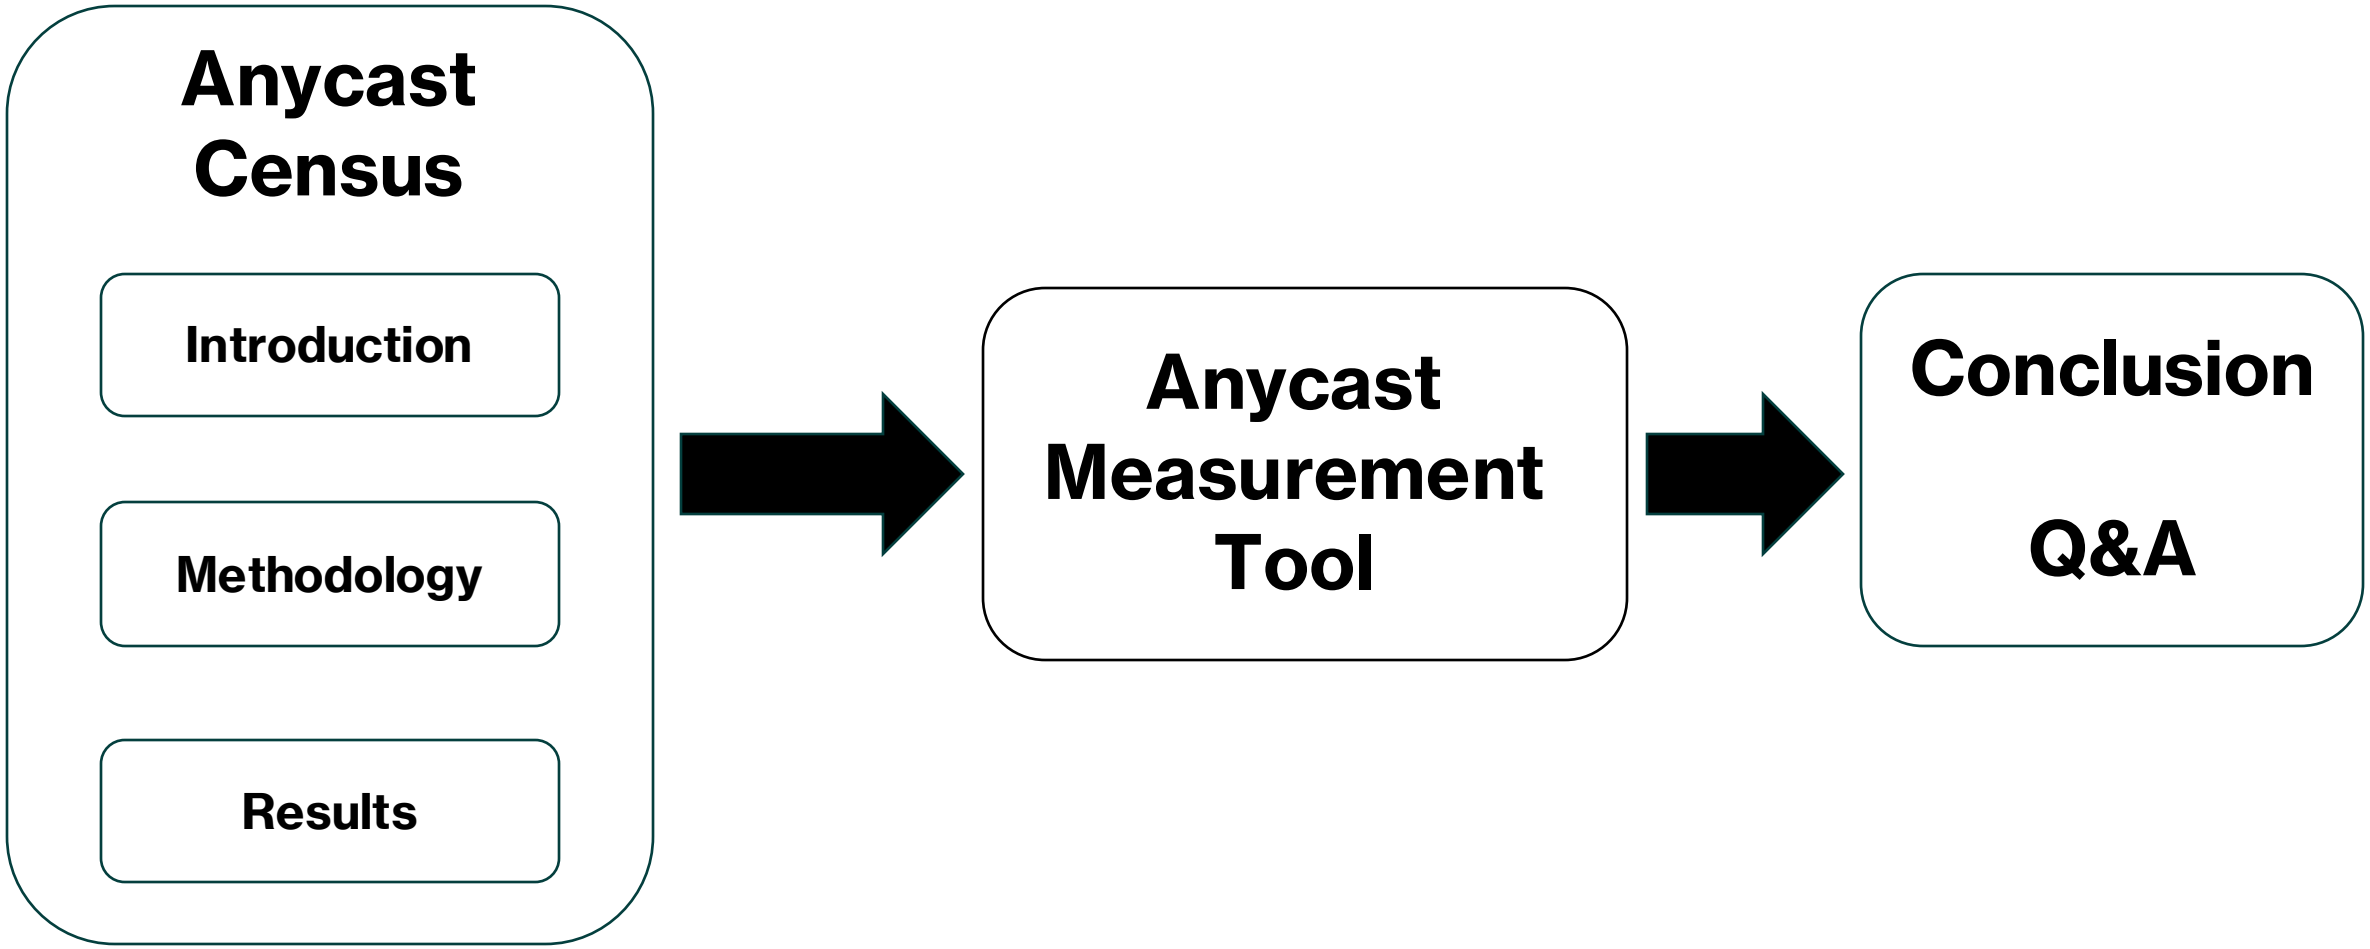
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

Content

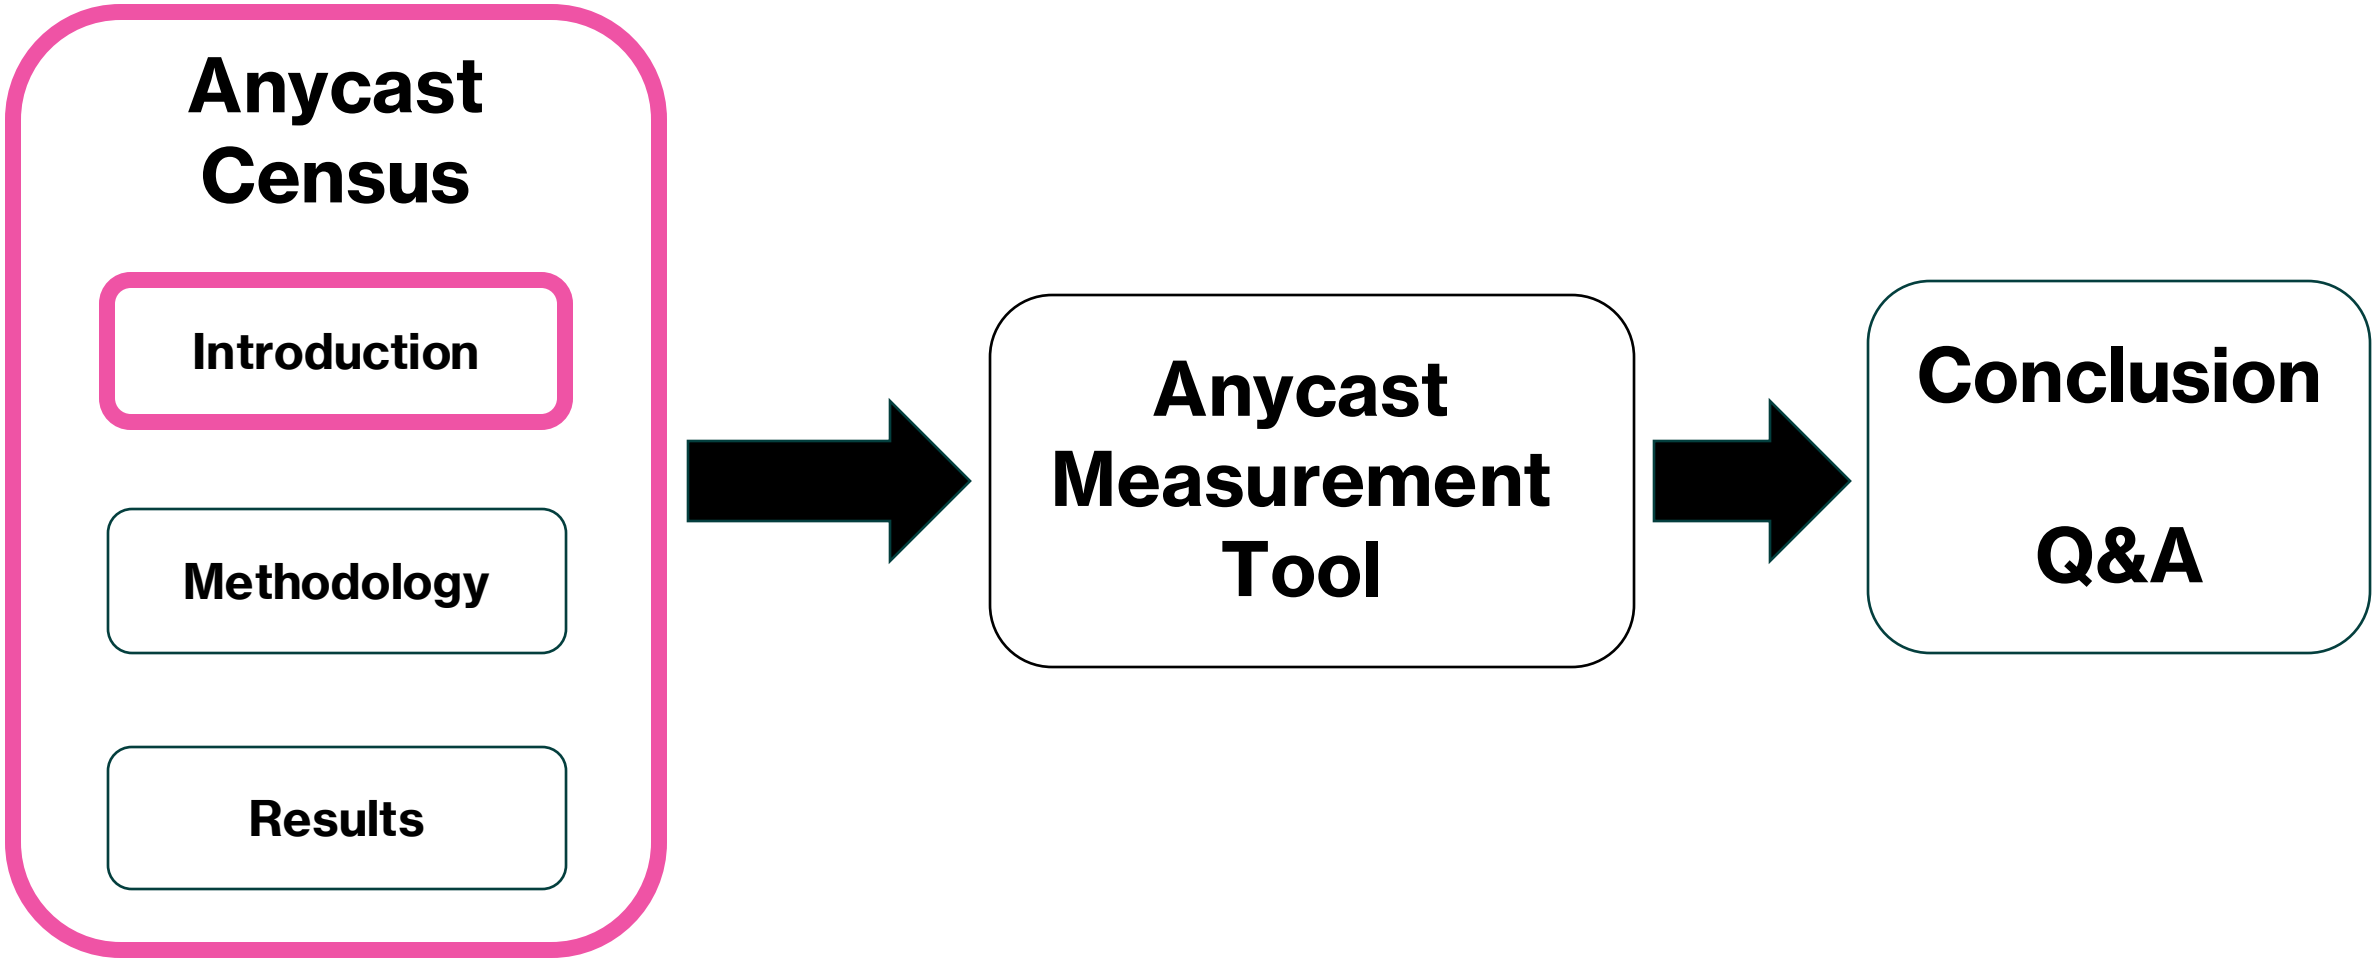
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

What is anycast?

Anycast is the geographical distribution of Internet services

- Achieved by announcing a single IP address at multiple locations
- Widely used as it provides resilience, low-latency, and load distribution

What is anycast?

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

Providing a daily anycast census

- Funded by **RIPE NCC Community Project Fund** (50 K€)
  - Deployment, infrastructure, and other research costs
- Establishing a measurement pipeline
- **A reliable daily census of anycast**
  - **Publicly available** for operators and researchers

RIPE NCC

Anycast census; why?

- 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)

Anycast census; why?

- 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)
- Anycast is **opaque**
  - Unknown if an address is anycast
  - Unknown if a service is provided using anycast
  - Unknown where anycast sites are located

Why do operators care?

Knowing what is, and what is not, anycast is useful for:

- Making **better Traffic Engineering** decisions
  - Troubleshoot network problems
  - Anycast to anycast routing problems
- **Resilience assessment** of third parties

Content

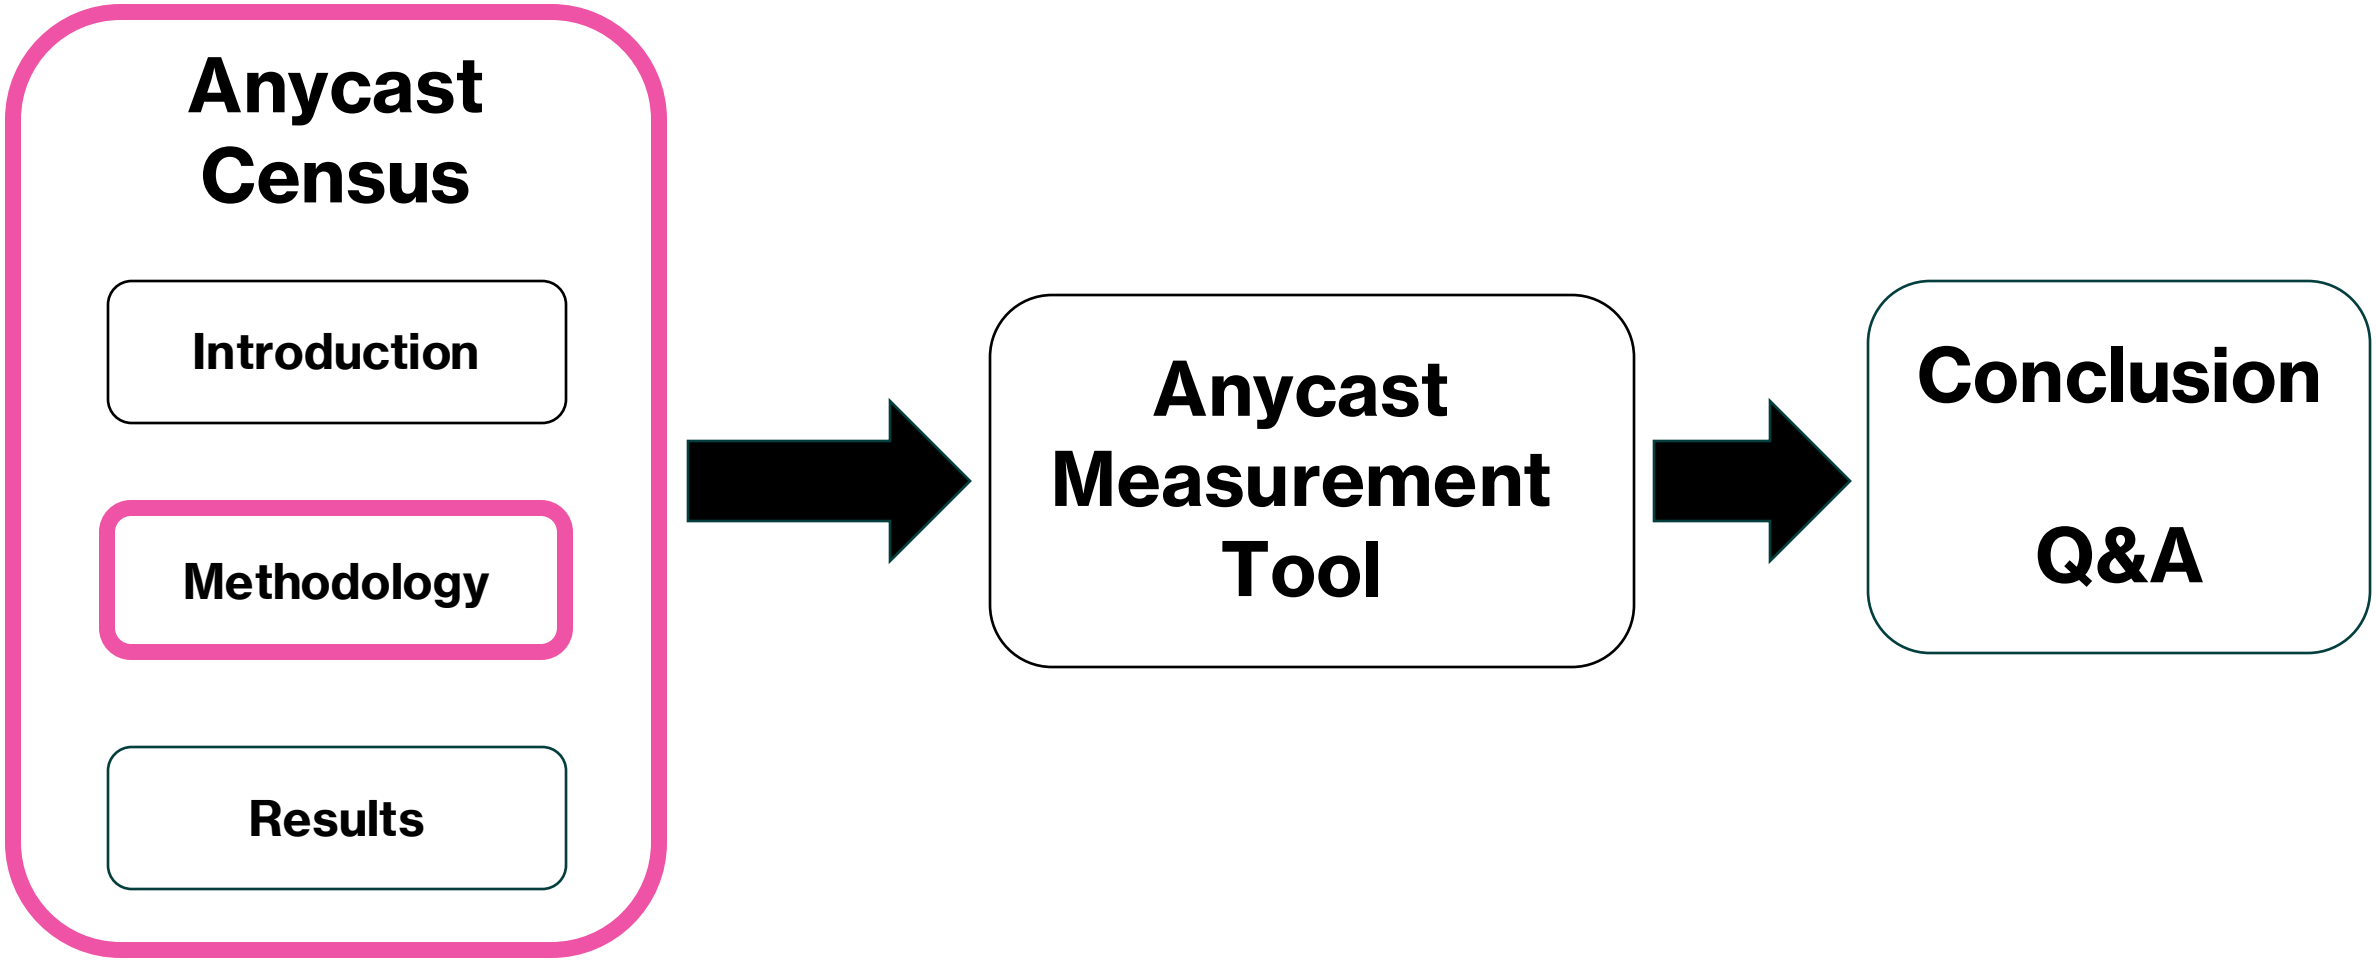
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

Realizing the census

- Daily measurement pipeline utilizes **two methodologies**
  - **Anycast-based (ICMP, TCP)**
  - **Latency-based (ICMP, TCP, UDP/DNS)**

- Anycast-based measurement uses **MAnycast²** approach
  - Developed in an IMC2020 submission [1]
  - Leverages the concept of using anycast to measure anycast

- Latency-based measurement based on iGreedy [2]

[1] Sommese et al. "MAnycast²: Using Anycast to Measure Anycast" ACM IMC '20

[2] Cicalese et al. "Latency-Based Anycast Geolocation: Algorithms, Software, and Data Sets," in IEEE Journal on Selected Areas in Communications '16

Anycast-based measurement Set-up

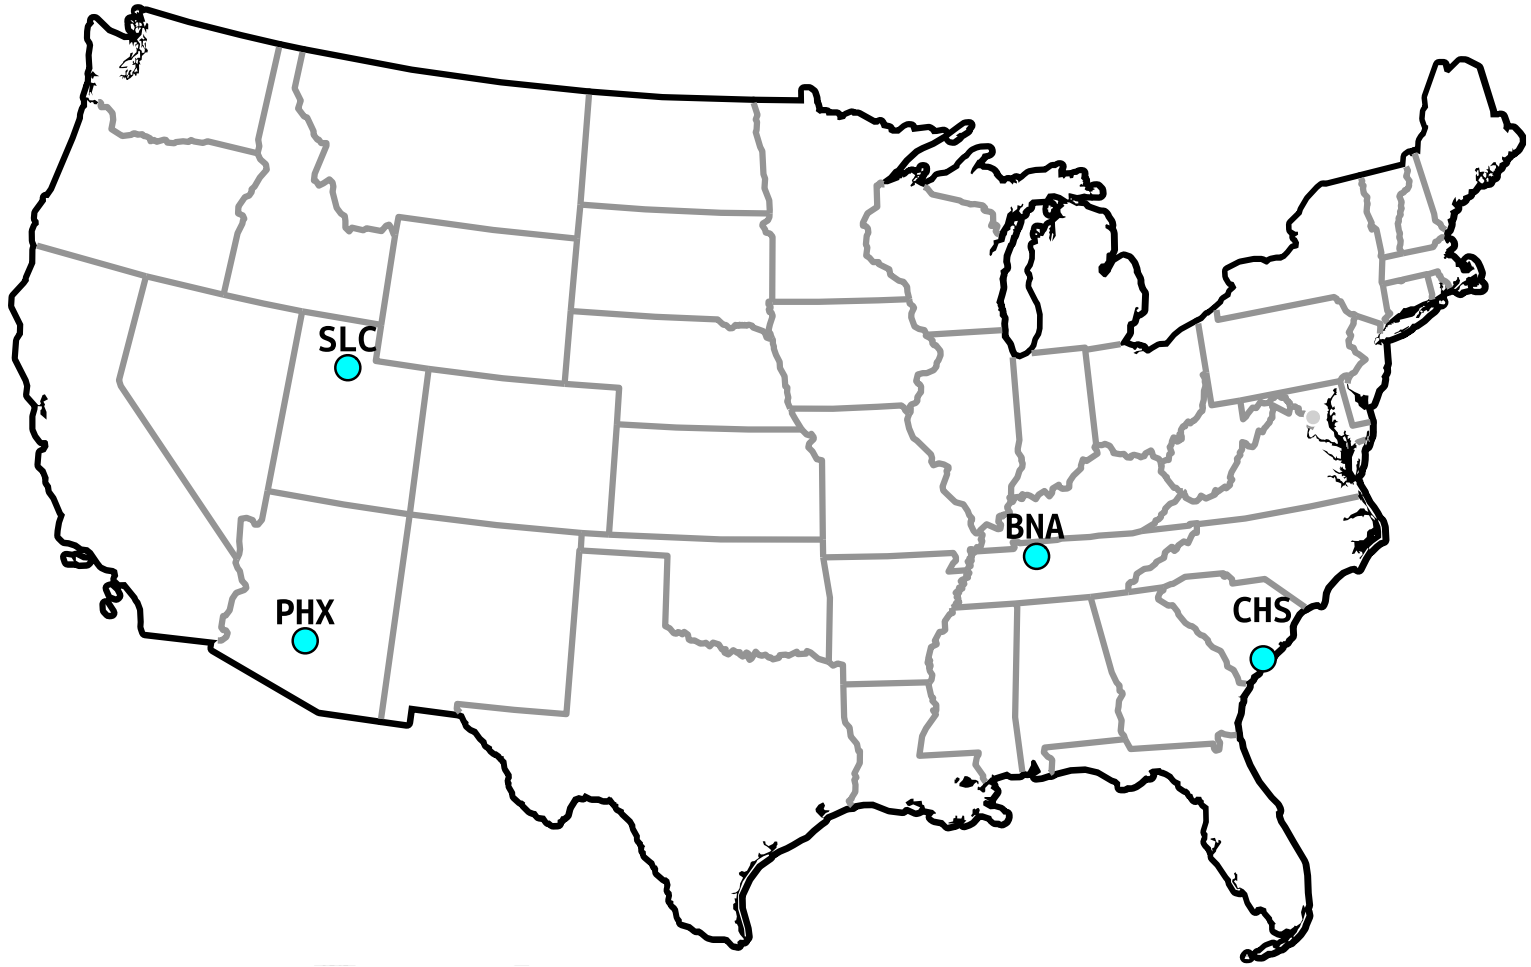

Anycast-based measurement
Unicast

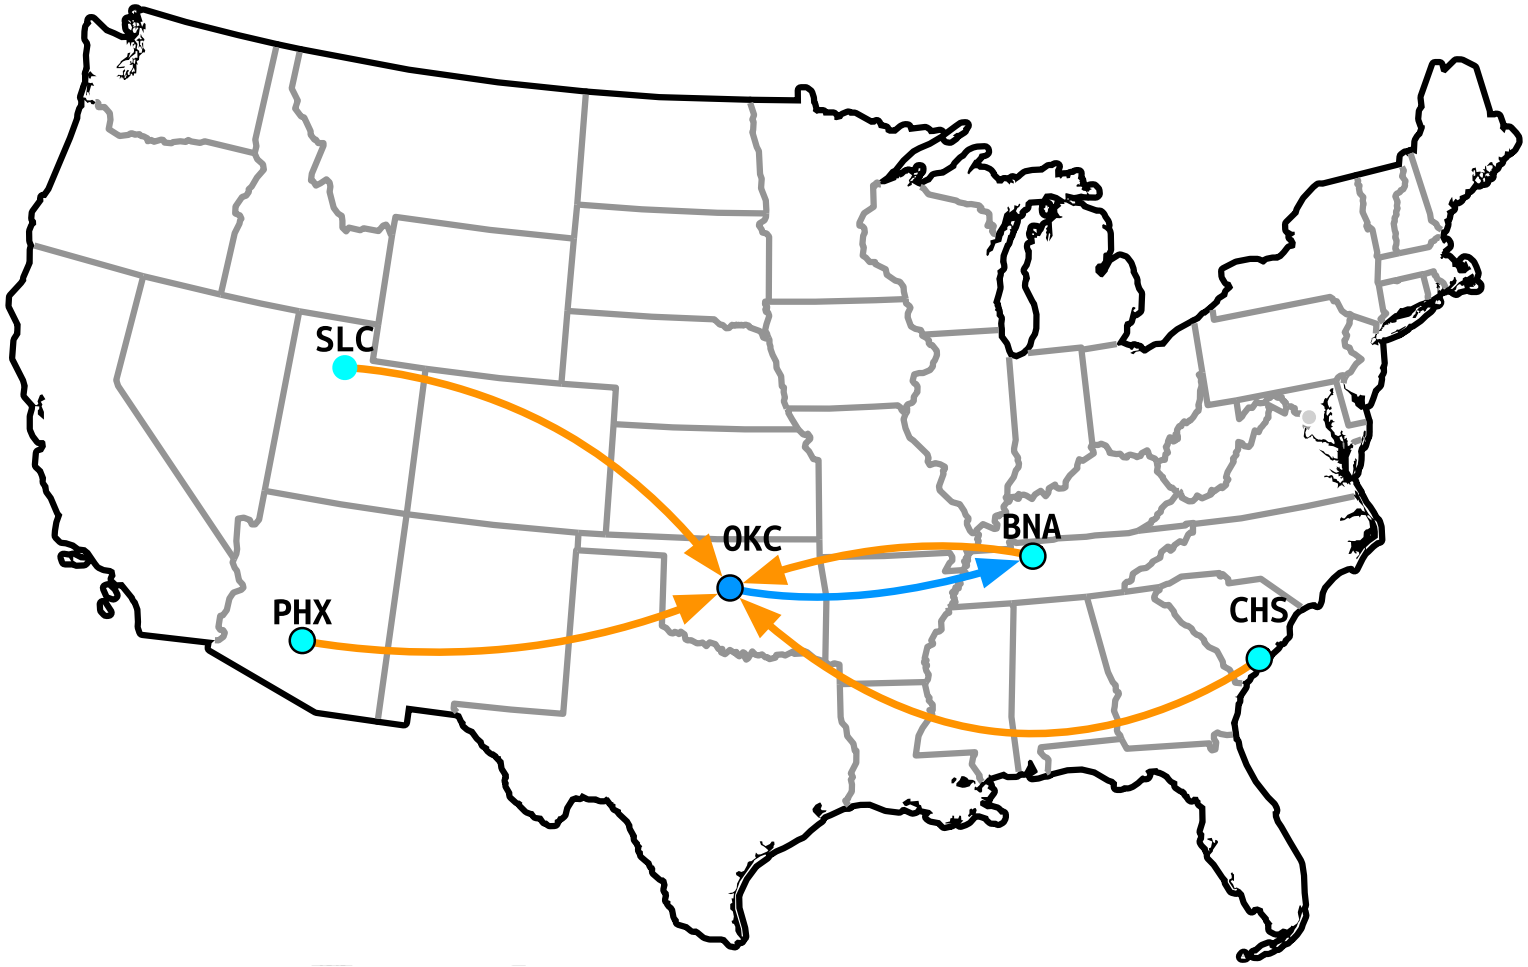

Anycast-based measurement
Anycast

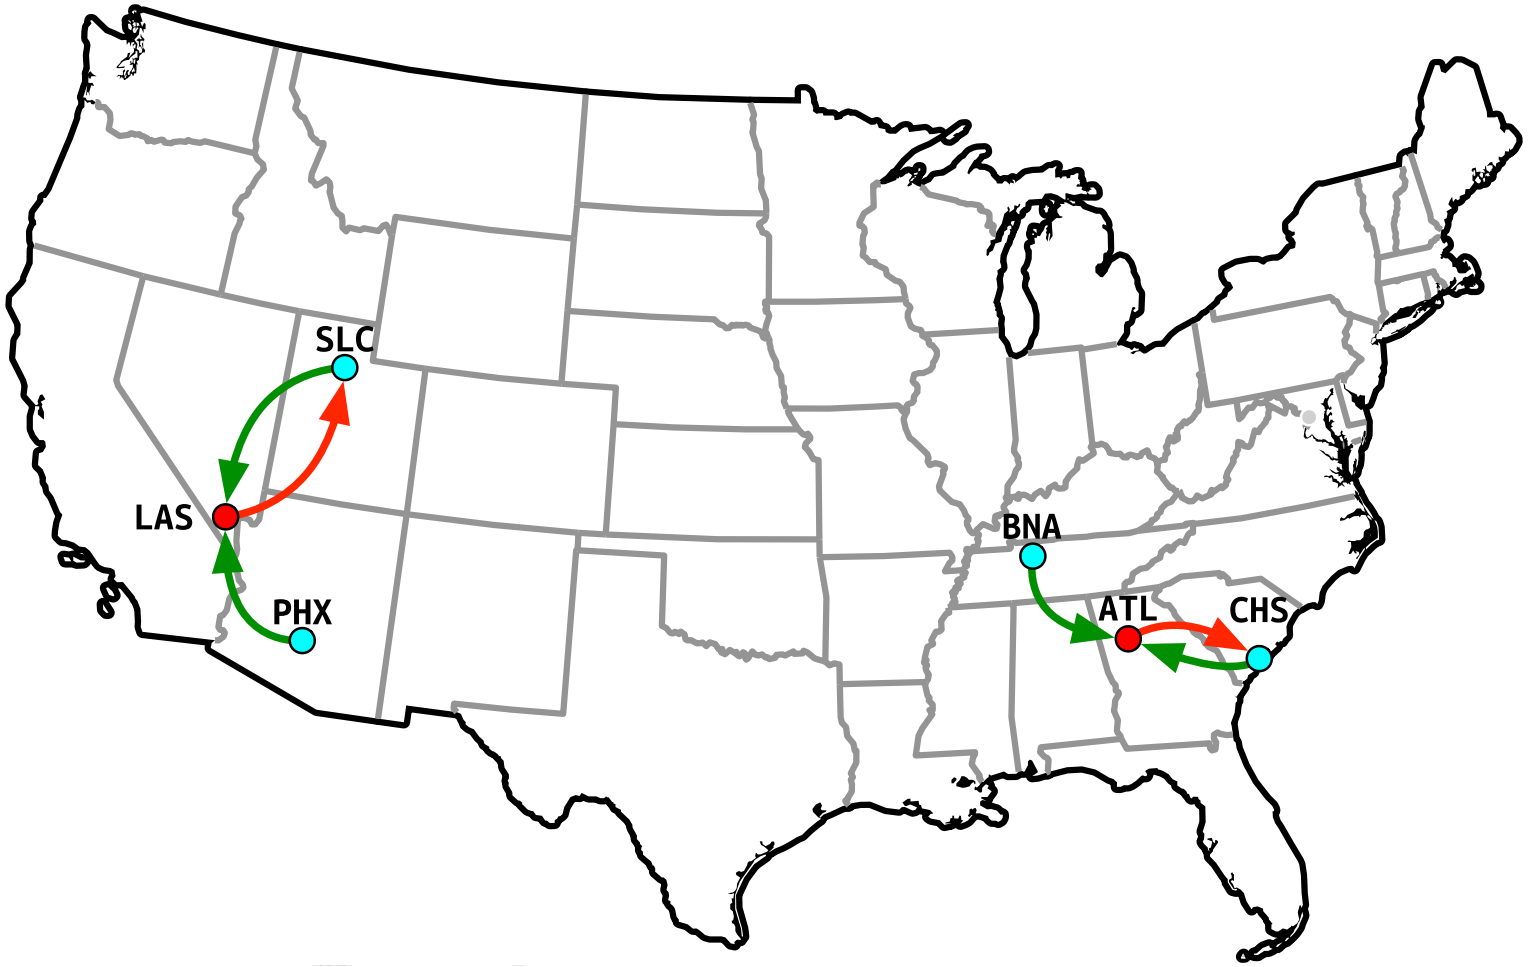

Anycast-based measurement
Pros and cons

- Pros:
  - **Low probing-cost** (suitable for Internet scale measurement)
  - **Low FN rate** (rarely misclassifies anycast as unicast)
- Cons:
  - **Considerable FP rate** (falsely classifying unicast as anycast)
  - **No geolocation** of sites (only detection & enumeration)

Latency-based measurement (GCD)

- GCD (**Great Circle Distance**)
- Latency-based measurement using **speed-of-light violations**
- Current state-of-the-art

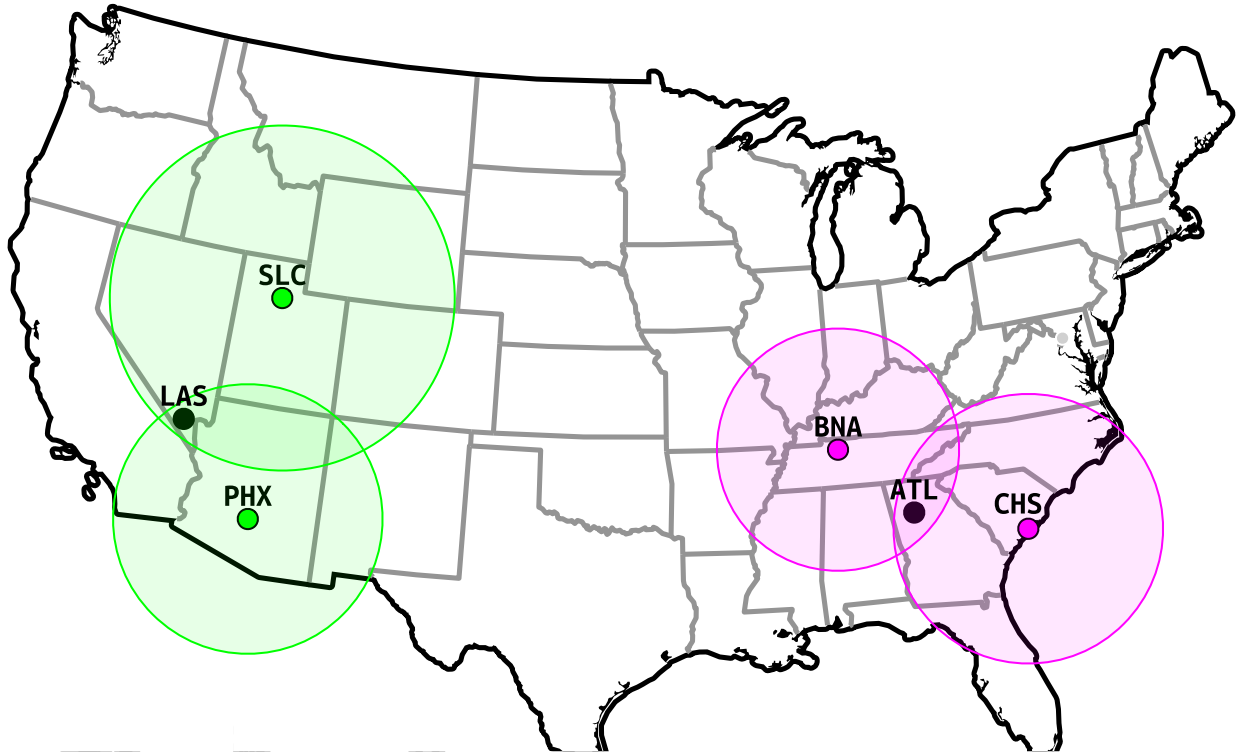

Latency-based measurement (GCD)

- GCD (**Great Circle Distance**)
- Latency-based measurement using **speed-of-light violations**
- Current state-of-the-art
- Pros:
  - Low FP/FN rate (**highly accurate**)
  - **Geolocation possible**
- Cons:
  - Requires large measurement platform (e.g., RIPE Atlas, CAIDA Ark)
  - High probing cost

  (unsuitable for Internet scale)

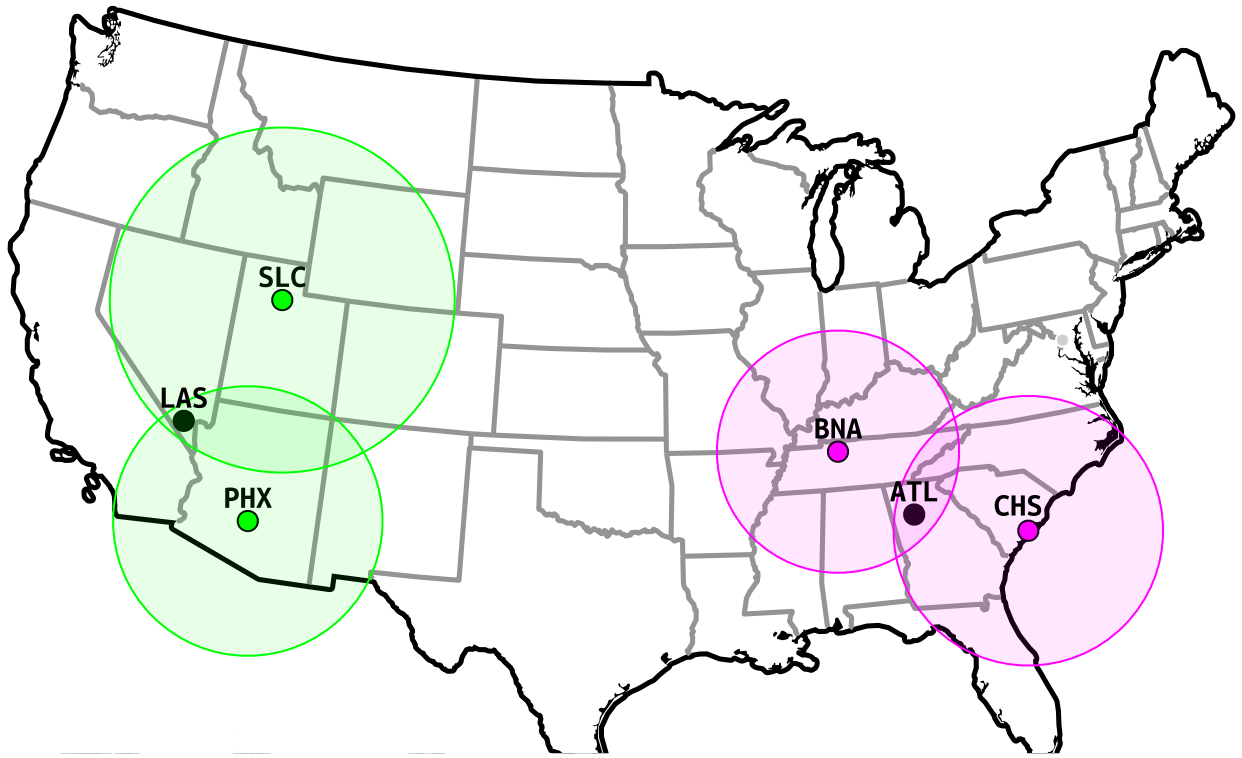

GCD measurement
RIPE Atlas example

Unicast

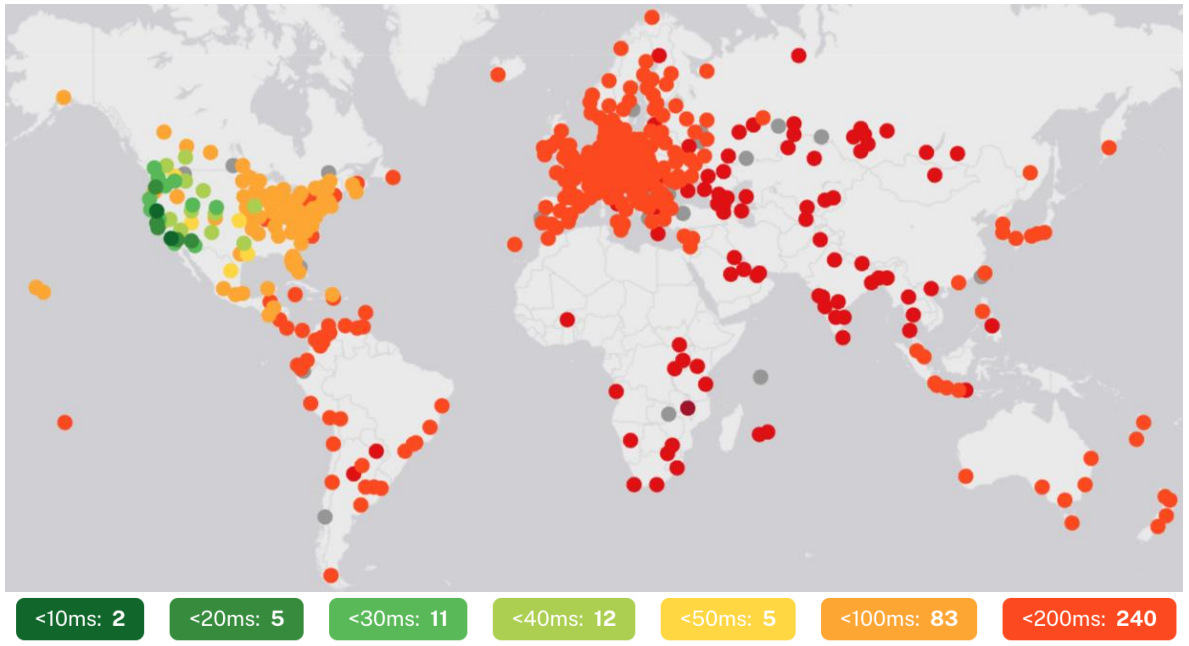

https://atlas.ripe.net/measurements/79147535/

Anycast

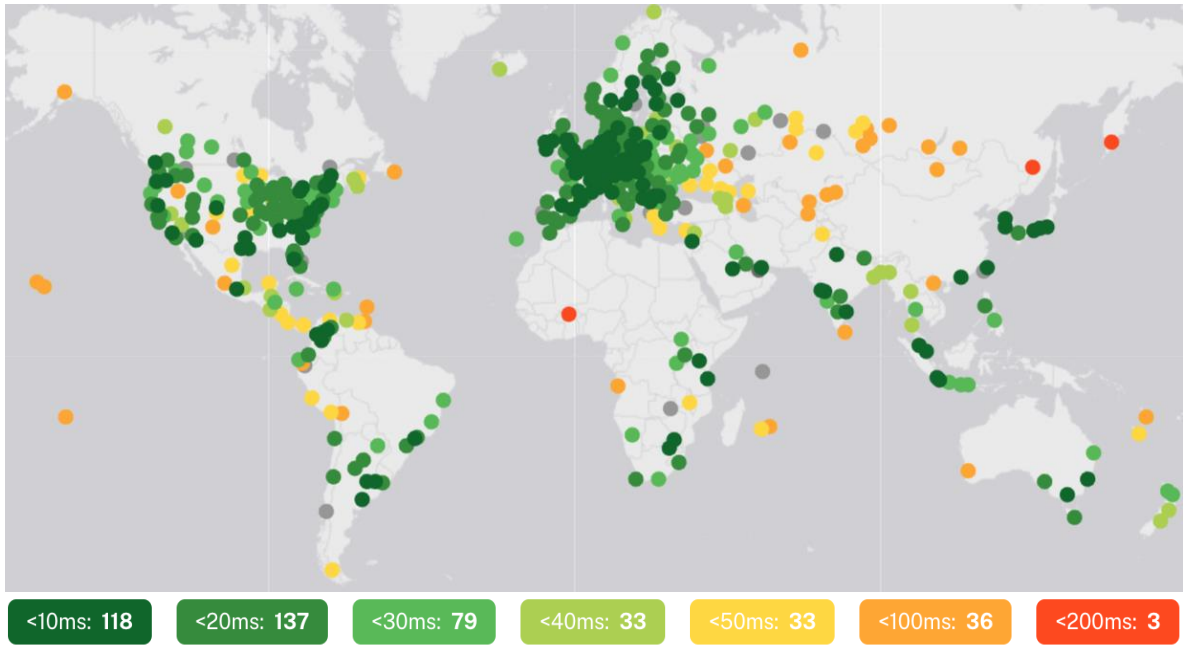

https://atlas.ripe.net/measurements/79137270/

RIPE NCC
RIPE Atlas

We combine the two

- Perform anycast-based census
  - Input: Internet wide hitlist

    (10^6 prefixes)
  - Output: set of "*anycast targets*" (***AT***) (includes TPs and FPs)
- Perform GCD-based measurement
  - Input: ***AT***

    (10^4 prefixes)
  - *Output: Anycast prefixes + enumeration + locations individual sites*

Our hitlists

- Hitlist: a set of responsive Internet hosts
- /24 granularity for IPv4, /48 for IPv6
  - Smallest routable prefix size
- IPv4
  - USC/ISI hitlist
- IPv6
  - TUM public IPv6 hitlist
  - AAAA record addresses from OpenINTEL

Pipeline

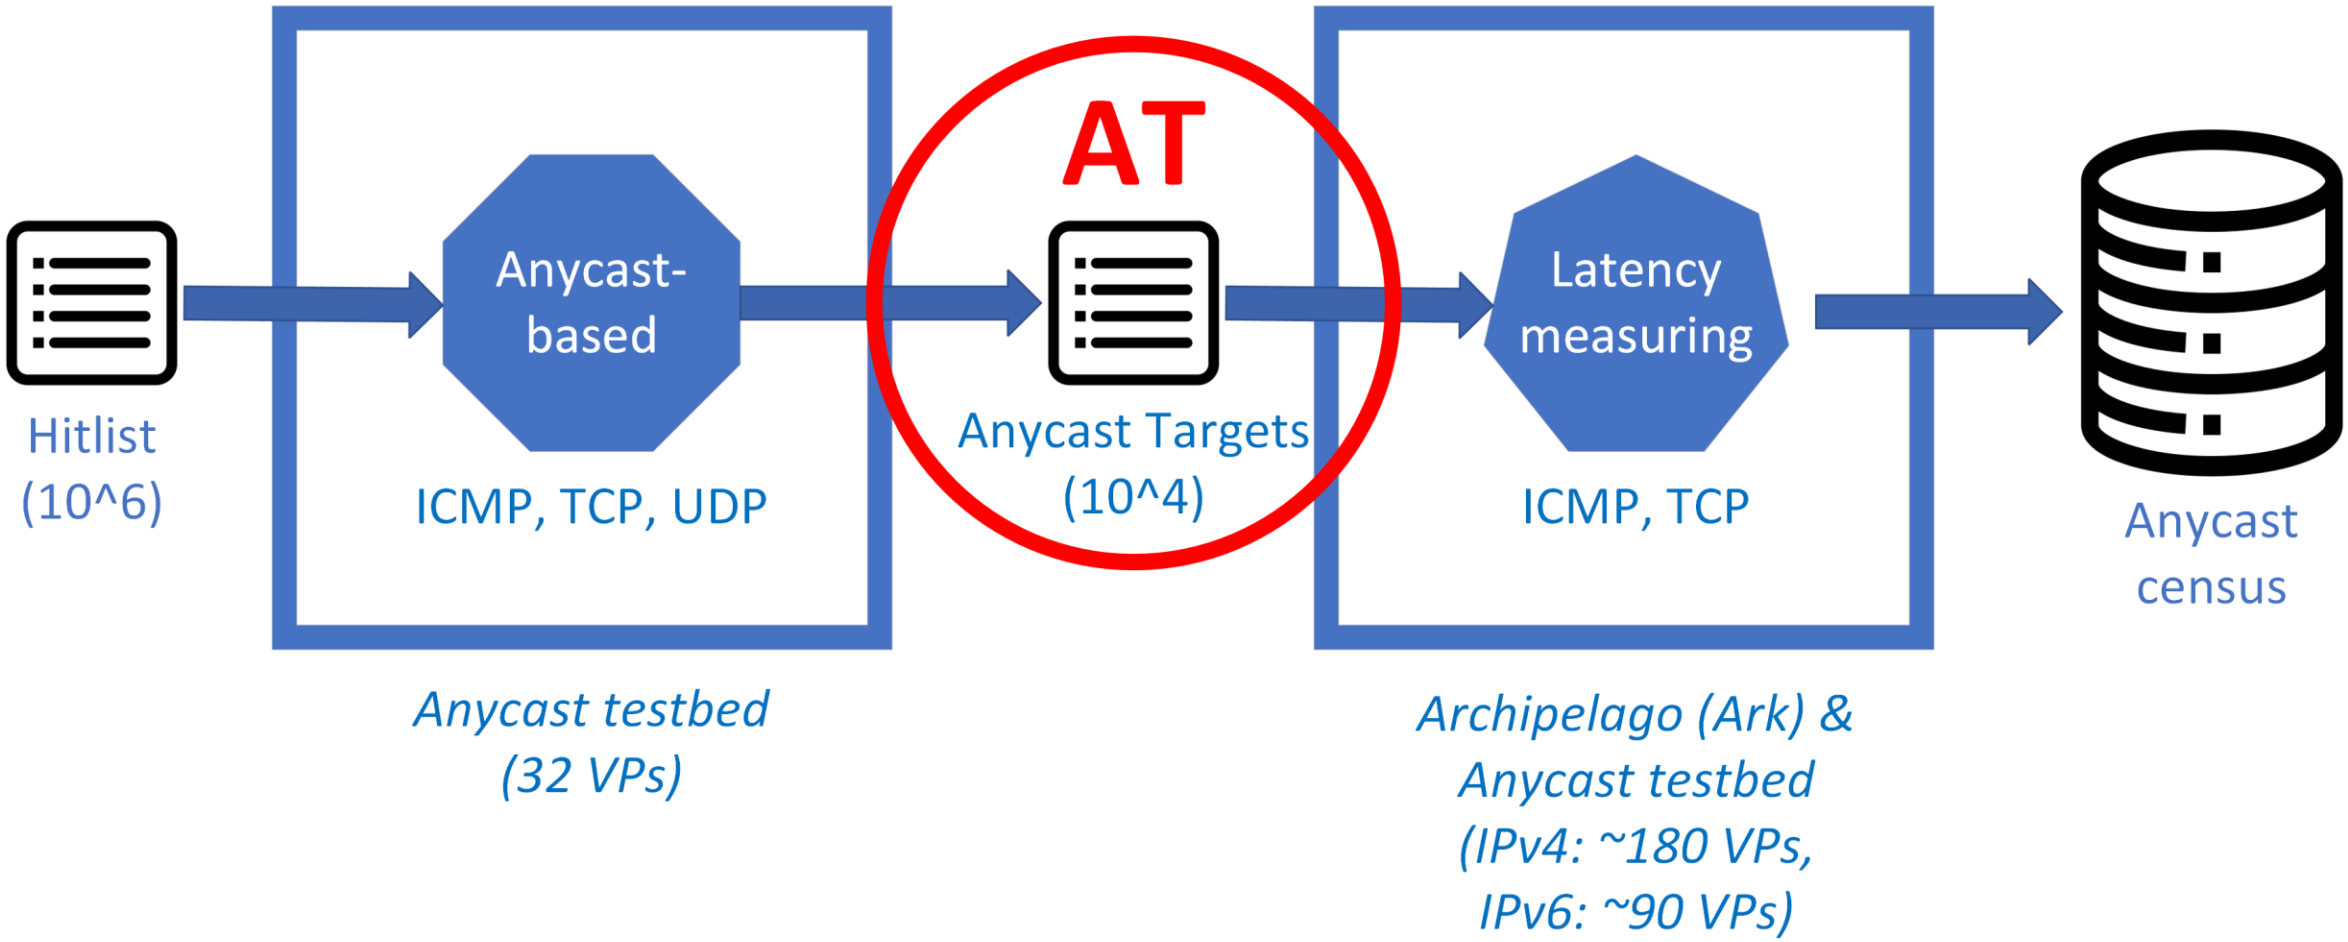
Hitlist
(10^6)
Anycast-
based
ICMP, TCP, UDP
Anycast testbed
(32 VPs)
AT
Anycast Targets
(10^4)
Latency
measuring
ICMP, TCP
Archipelago (Ark) &
Anycast testbed
(IPv4: ~180 VPs,
IPv6: ~90 VPs)
Anycast
census

Anycast-based measurement tool

- Developed anycast measurement tool
  - More on this later …

- Deployed using Vultr
  - 19 countries, 6 continents

Our anycast deployment (32 locations)
[https://www.vultr.com/features/datacenter-locations/]

Latency-based measurement system

- GCD measurements with CAIDA's Ark [1] and Vultr VPs
  - ~ 180 vantage points
- Implemented using Scamper [2] tool
- Accurate geolocation and enumeration of 'small' anycast deployments
  - ** Fails to differentiate between sites with near geographic proximity*
- Geolocation and enumeration up to 60 sites for 'large' deployments

[1] https://www.caida.org/projects/ark

[2] https://www.caida.org/catalog/software/scamper/

A combined view

- Neither methodology perfect
- Census contains both anycast-based and latency-based results
  - Latency-based approach has rare cases of FNs

    Example: anycast in Belgium + Netherlands
  - Anycast-based approach has FPs
- Criteria up to you…
  - Filter on both -> accept some FNs
  - Filter on either -> accept some FPs

Content

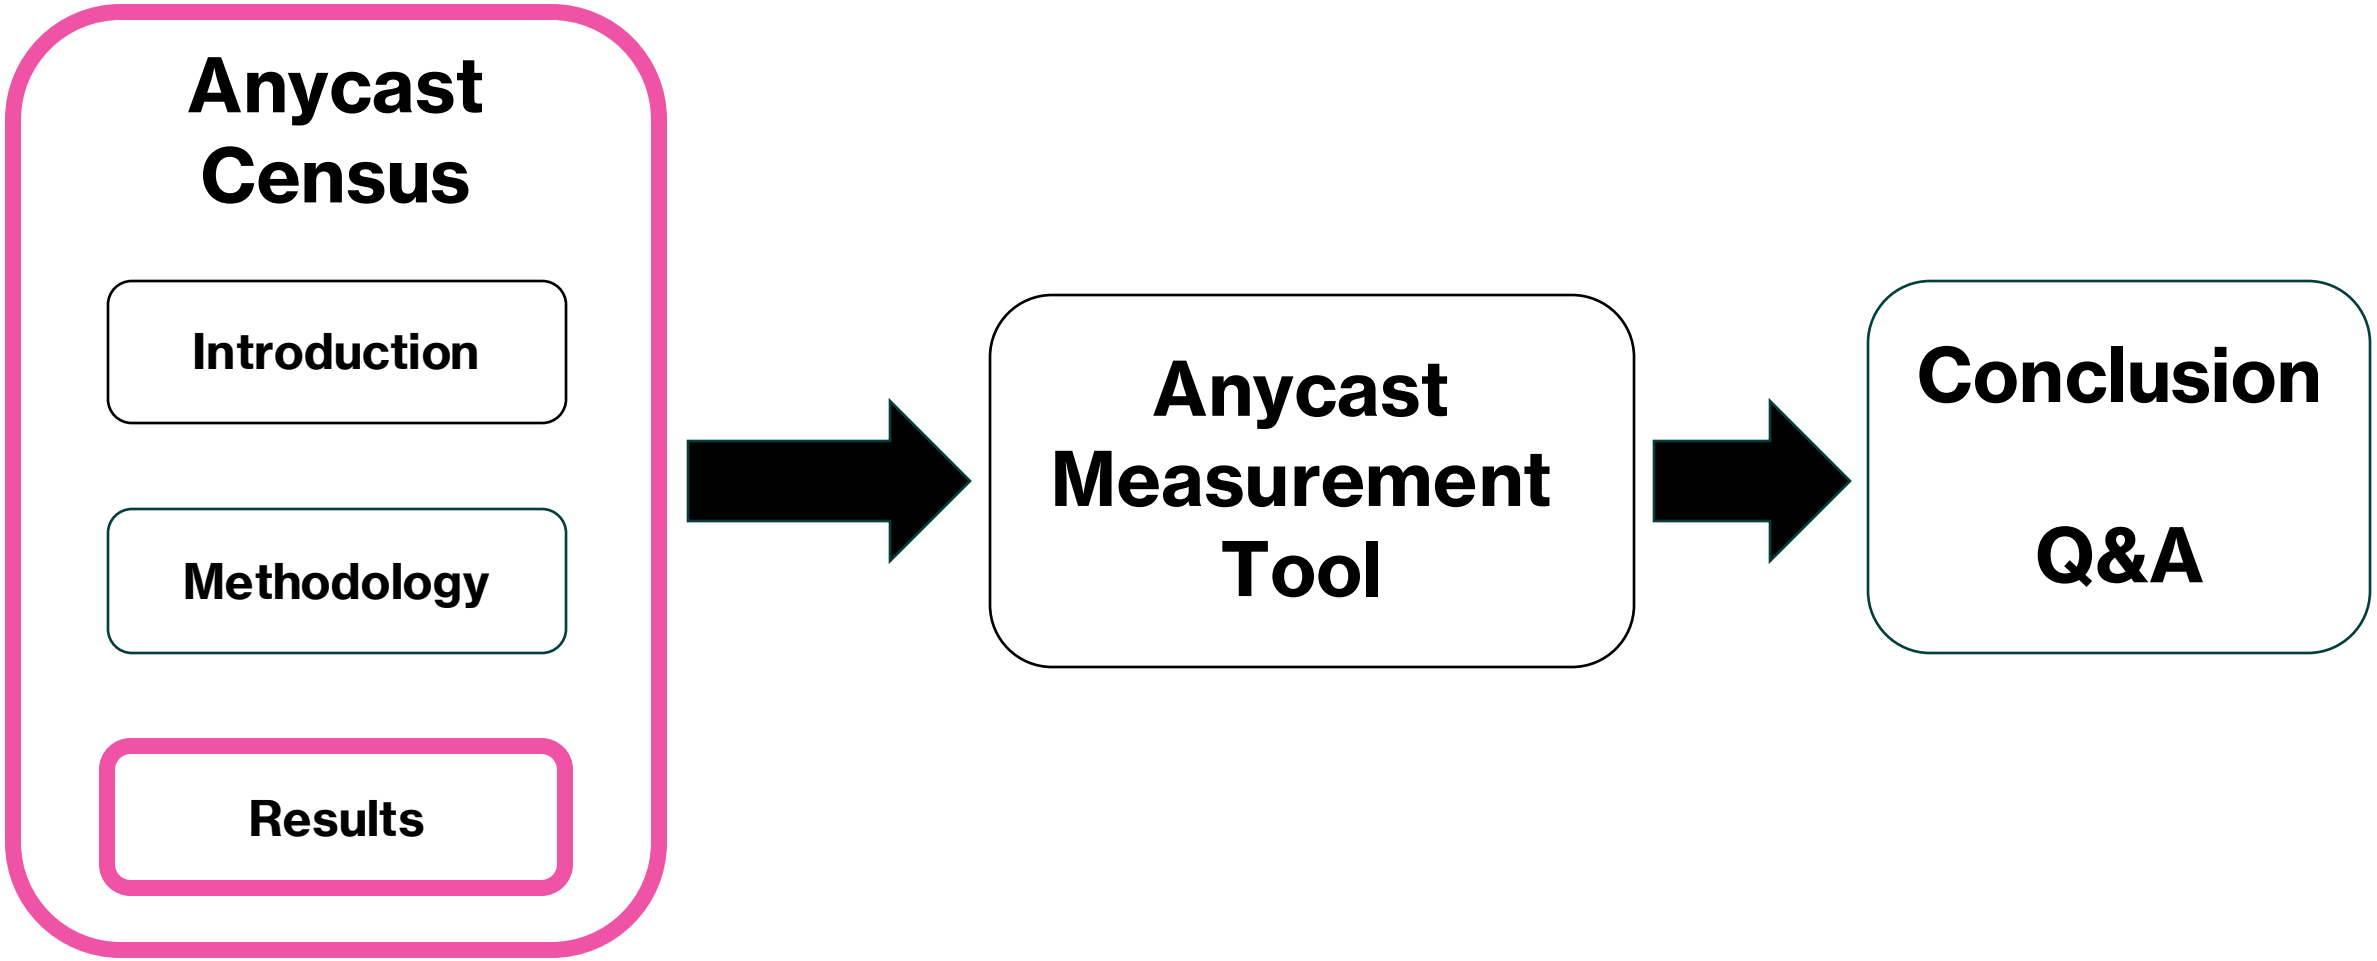
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

Results
GCD-confirmed prefixes

- ~12.3k anycasted /24s (769 ASes)
- ~6.0k anycasted /48s (462 ASes)
- 299 ASes found to anycast both IPv4 and IPv6

Results
GCD-confirmed prefixes

AS	Organization	IPv4	IPv6
396982	Google Cloud	3,345	3
13335	Cloudflare	3,131	162
16509	Amazon	1,235	86
54113	Fastly	438	56
15169	Google	282	6
209242	Cloudflare Spectrum	234	2,836
19551	Incapsula	2	292
12041	Afilias	222	207
44273	GoDaddy	31	122

Table 6: Largest ASes originating anycast prefixes for IPv4 and IPv6.

Results
GCD-confirmed prefixes

AS	Organization	IPv4	IPv6
396982	Google Cloud	3,345	3
13335	Cloudflare	3,131	162
16509	Amazon	1,235	86
54113	Fastly	438	56
15169	Google	282	6
209242	Cloudflare Spectrum	234	2,836
19551	Incapsula	2	292
12041	Afilias	222	207
44273	GoDaddy	31	122

Top 5 IPv4 anycast ASes

Table 6: Largest ASes originating anycast prefixes for IPv4 and IPv6.

Results
GCD-confirmed prefixes

AS	Organization	IPv4	IPv6
396982	Google Cloud	3,345	3
13335	Cloudflare	3,131	162
16509	Amazon	1,235	86
54113	Fastly	438	56
15169	Google	282	6
209242	Cloudflare Spectrum	234	2,836
19551	Incapsula	2	292
12041	Afilias	222	207
44273	GoDaddy	31	122

Top 5 IPv6 anycast ASes

Table 6: Largest ASes originating anycast prefixes for IPv4 and IPv6.

Results
Anycast-based by protocol

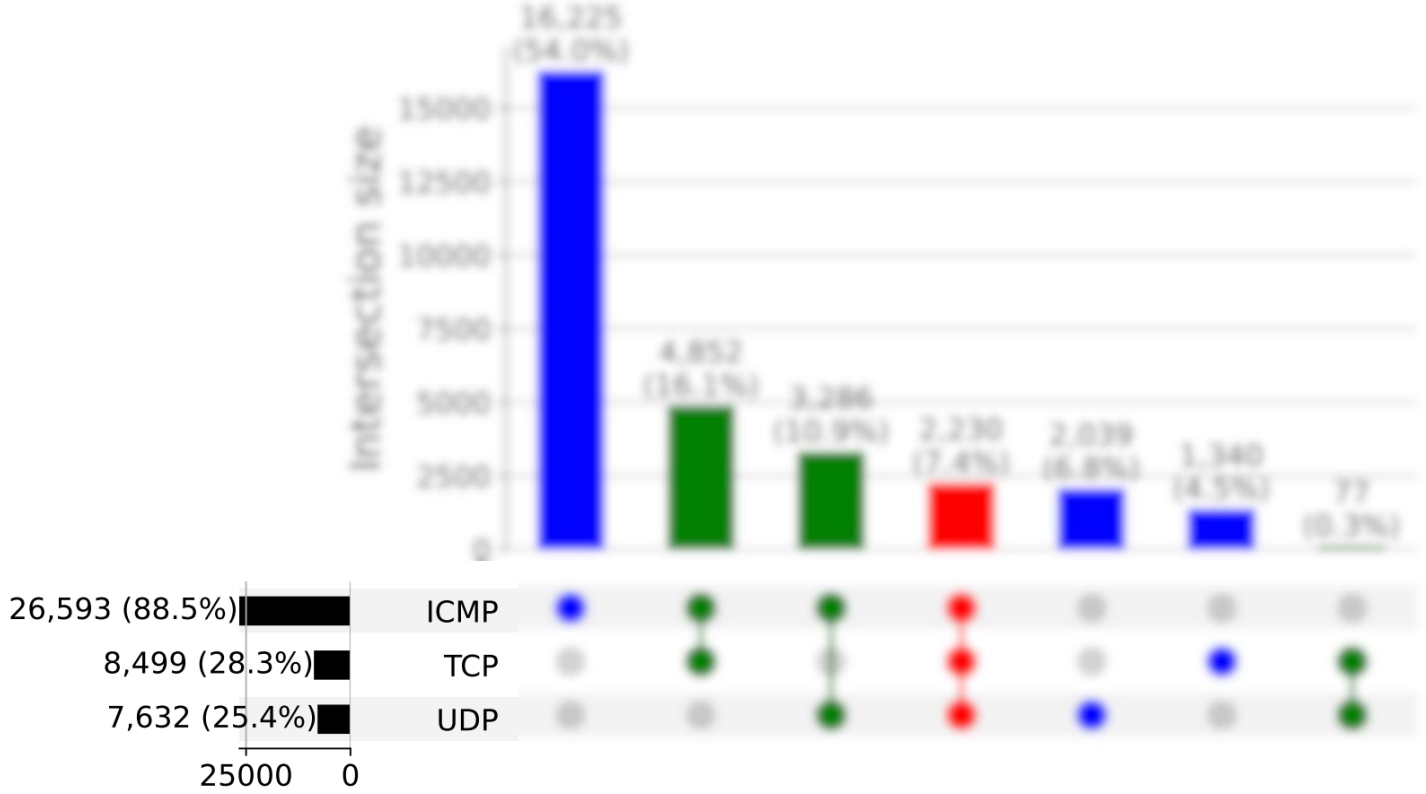

Figure 6: MAnycastR detection of anycast candidates for ICMPv4, TCPv4, and UDPv4.

Results
Anycast-based by protocol

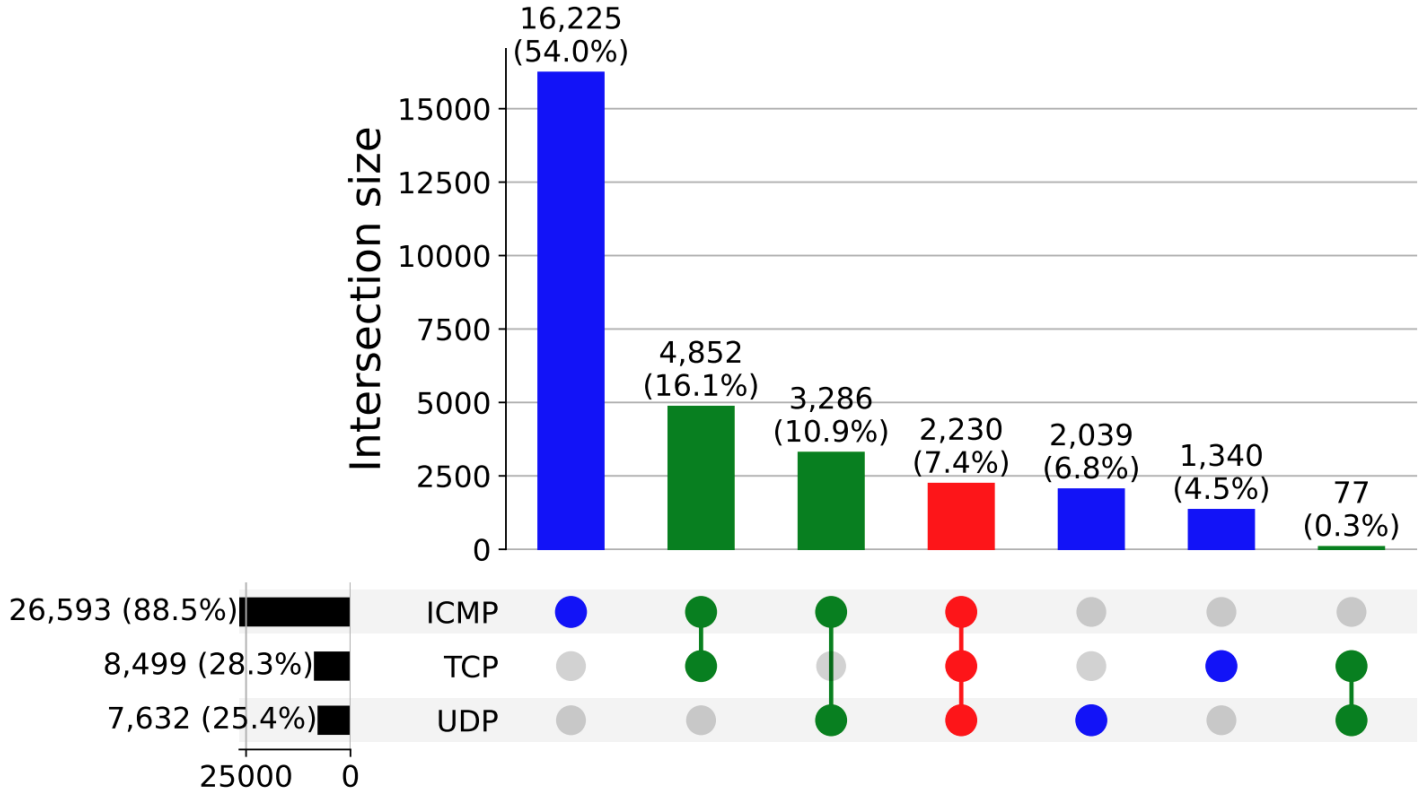

Figure 6: MAnycastR detection of anycast candidates for ICMPv4, TCPv4, and UDPv4.

Results
Anycast-based by protocol

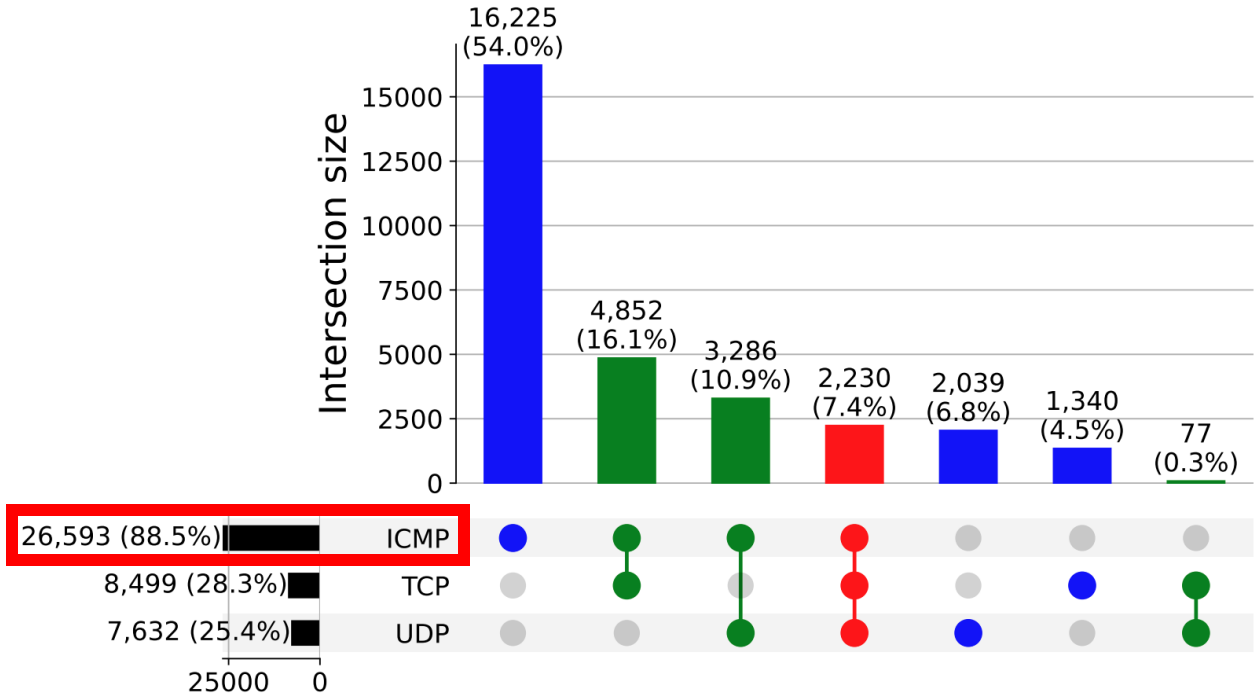

Figure 6: MAnycastR detection of anycast candidates for ICMPv4, TCPv4, and UDPv4.

Main takeaways:

- Majority detectable using ICMP (89%)

Results
Anycast-based by protocol

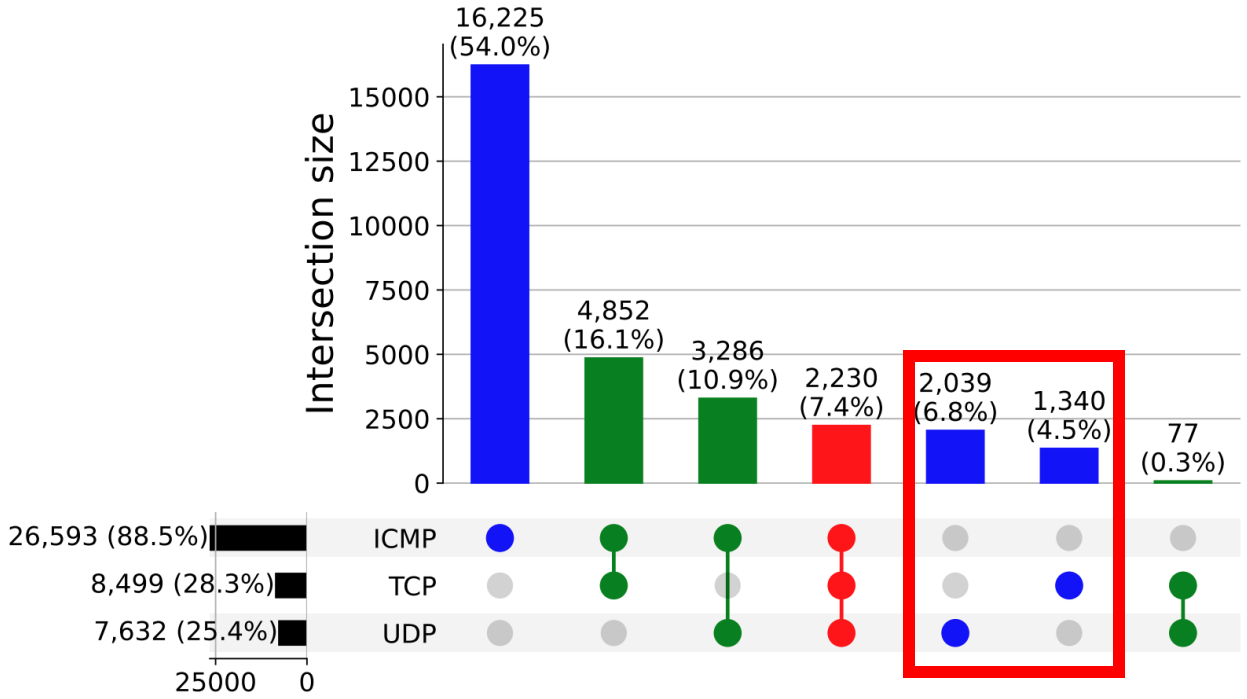

Figure 6: MAnycastR detection of anycast candidates for ICMPv4, TCPv4, and UDPv4.

Main takeaways:

- Majority detectable using ICMP (89%)
- Additional prefixes found (not detected with traditional approaches)
  - 2.0k UDP
  - 1.3k TCP

Results
Enumerating deployments

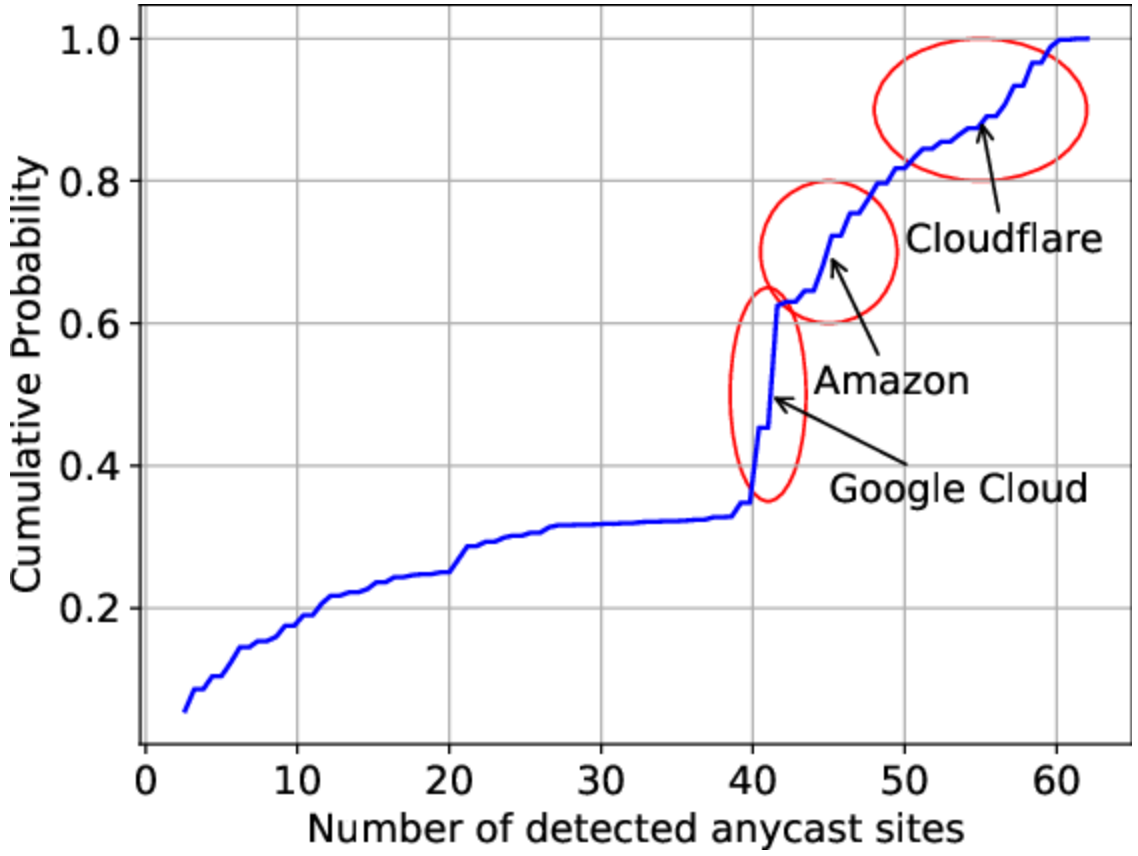

Results
Enumerating deployments

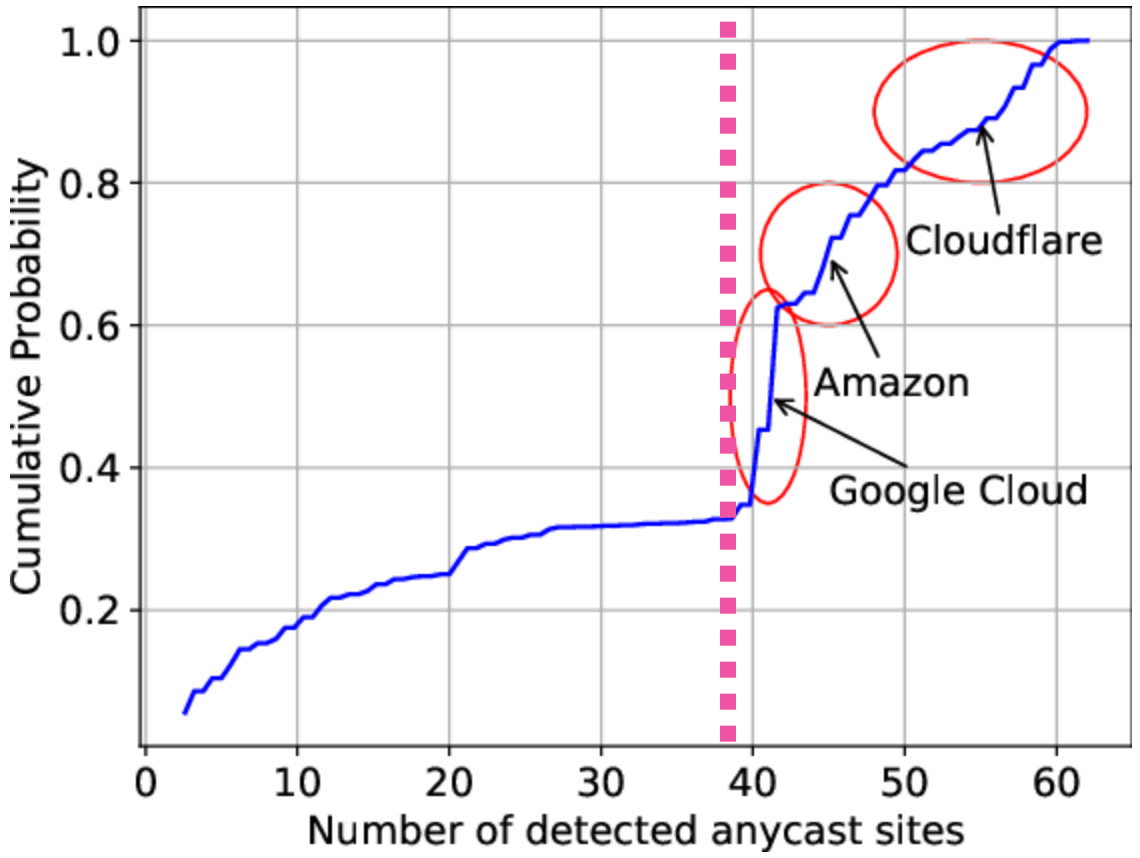

1/3 of anycast prefixes we detect **< 40 sites**

2/3 of anycast prefixes we detect **>= 40 sites**

Results
Enumerating deployments

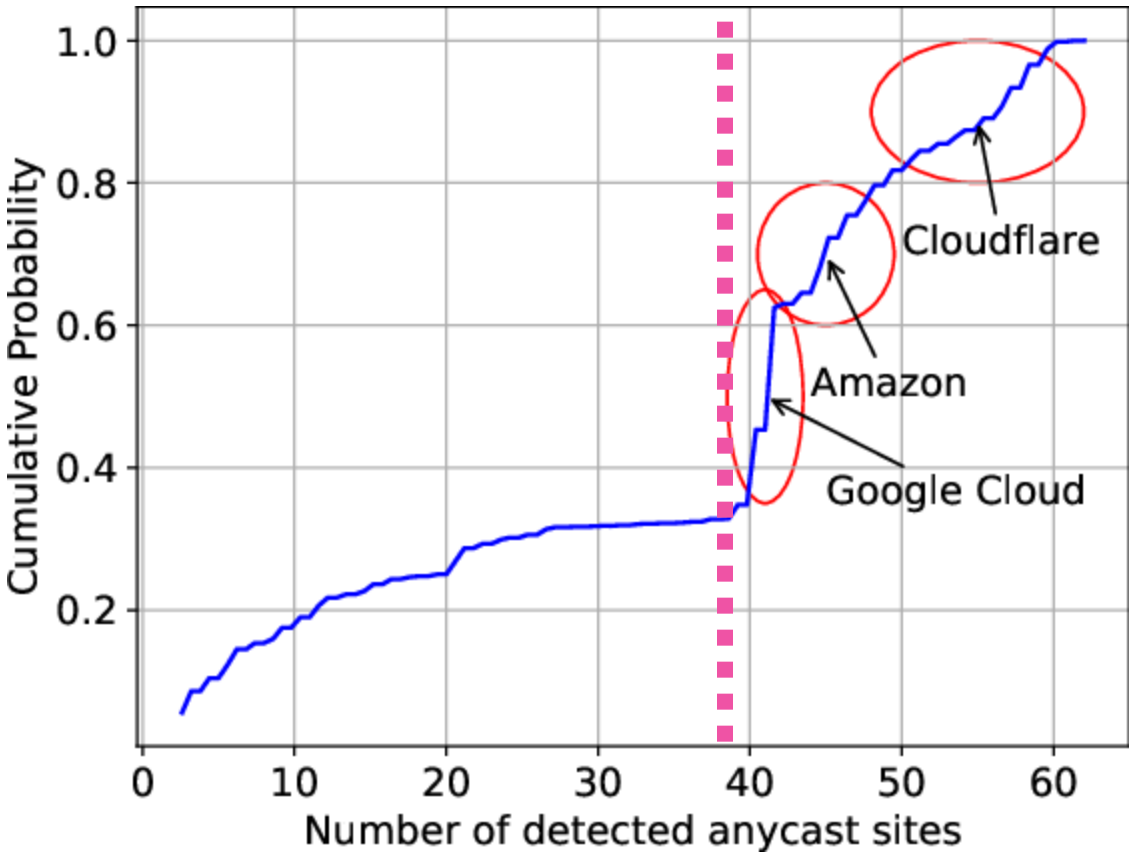

1/3 of anycast prefixes we detect < 40 sites

2/3 of anycast prefixes we detect >= 40 sites

Enumeration:

- **Lower bound** of real deployment
- Fairly **accurate** for **small deployments**
- **Indication of size** for **large deployments**

Longitudinal observations (preliminary)

- Longitudinal measurement lets us map the development of Anycast

- Daily measurement allows us to observe:
  - Anycast deployments regularly changing in size
  - Prefixes switching between unicast and anycast (on-demand anycast)
  - Cases of BGP prefix hijacking
  - Cases of temporary anycast (anti-DDoS?)
  - Anycast outages (entirely or at particular sites)

Census summary

- Created **responsible**, **scalable**, **accurate** anycast measurement pipeline

- A **daily** census of anycast
  - Detection, enumeration, and geolocation
  - Providing results from two methodologies

- We hope our census is useful to the community

Content

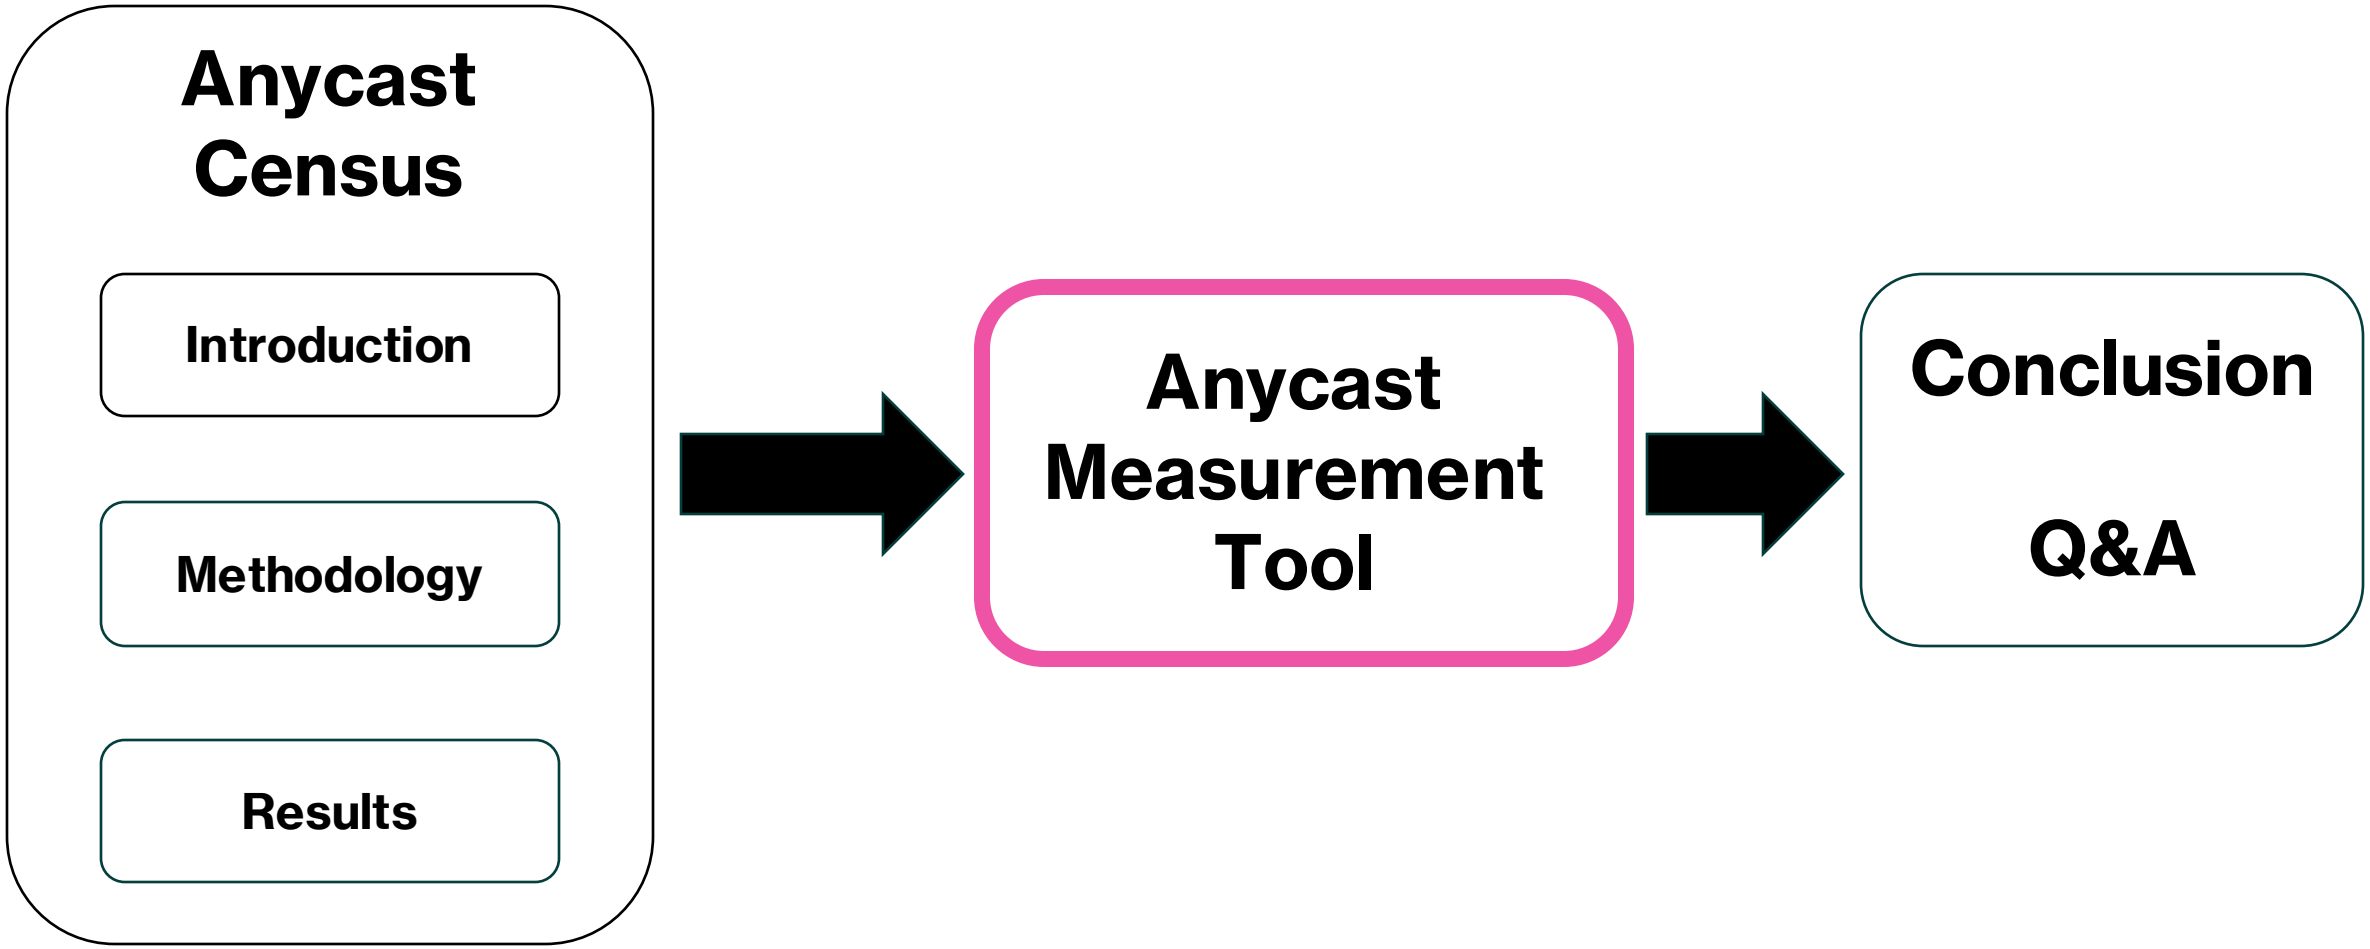
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

Anycast measurement tool

- So far, we looked at measuring anycast deployments externally
- Tool also built for measuring anycast deployments themselves

- Allows anycast operators to assess the performance of their deployment

- Example: Catchment mappings (like done with Verfploeter [1])
  - *I.e., mapping for each address, the anycast site that it routes to*
  - Tool can map the IPv4 space in a few minutes with ease

[1] De Vries et al. "Verfploeter: Broad and Load-Aware Anycast Mapping" ACM IMC '17

Catchment analysis
Anycast RTT data

- Measurement tool can synchronously probe the hitlist from all VPs
  - *i.e.,* each hitlist target receives a probe from all anycast sites (with 1-second intervals)
  - Technique used to perform anycast-based measurement
- Filter on *sender == receiver,*

  for all probed IPs
  - Example: filter on probe from BNA (discard all others)
  - Result: RTT data (like regular unicast ping) for entire hitlist

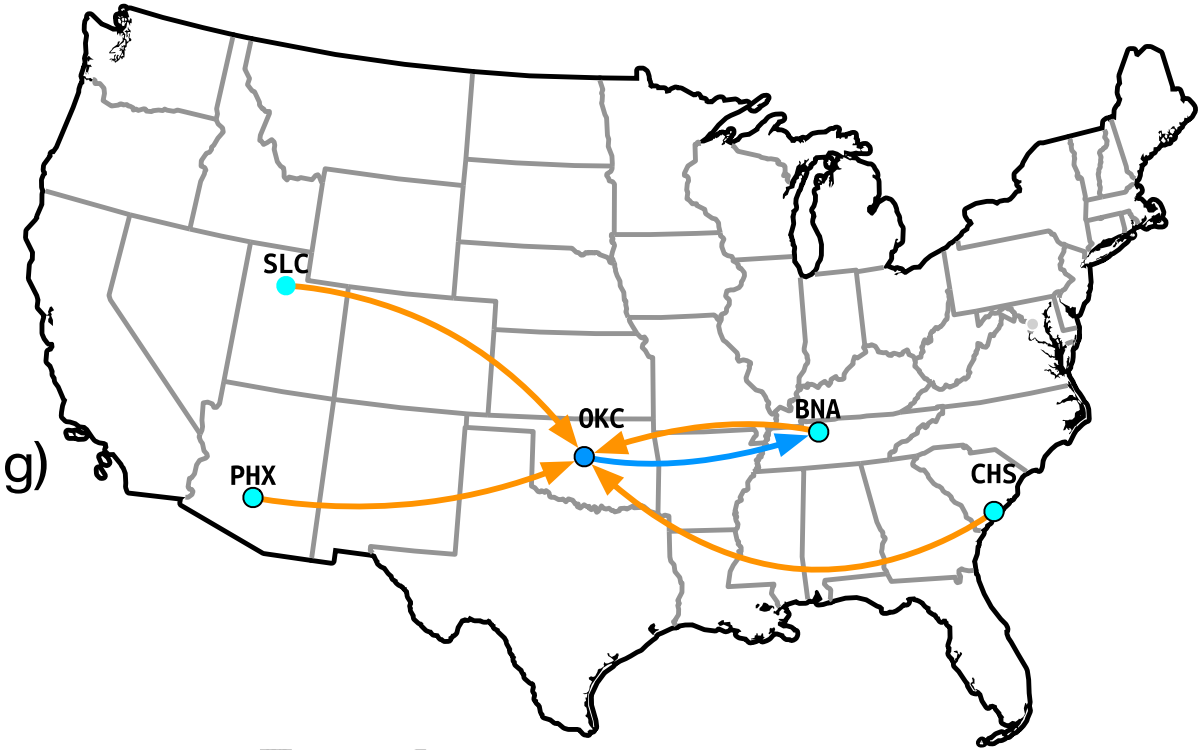

Catchment analysis
Anycast RTT data (example)

Latency towards
USC/ISI IPv4 hitlist

Average	65.7 ms
25%	9.8 ms
50%	23.5 ms
75%	74.1 ms

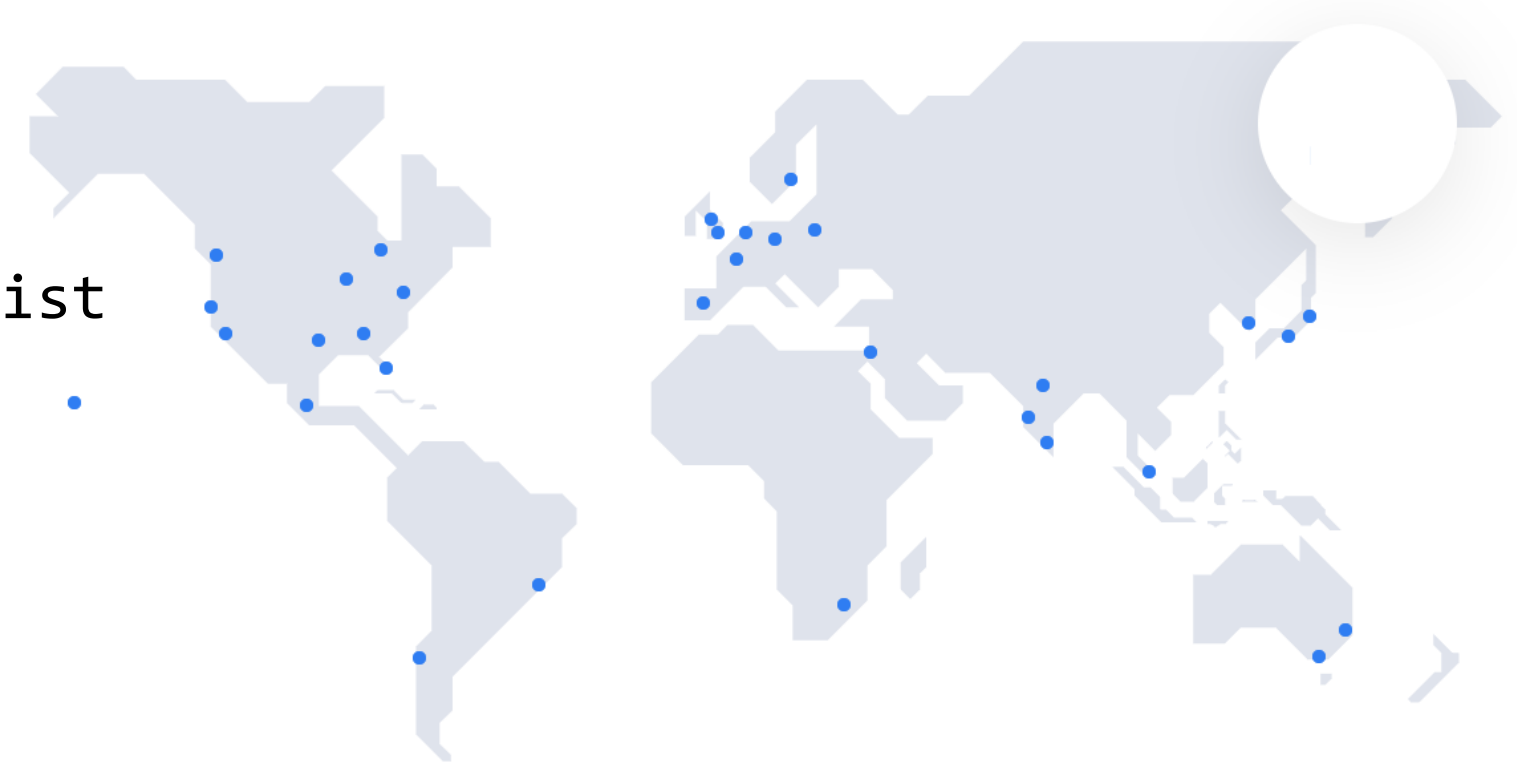

Our anycast deployment (32 locations)

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- Measurement tool allows for probing hitlist from all VPs with their unicast IPs
  - Result: RTT data from all anycast sites to the probed IP

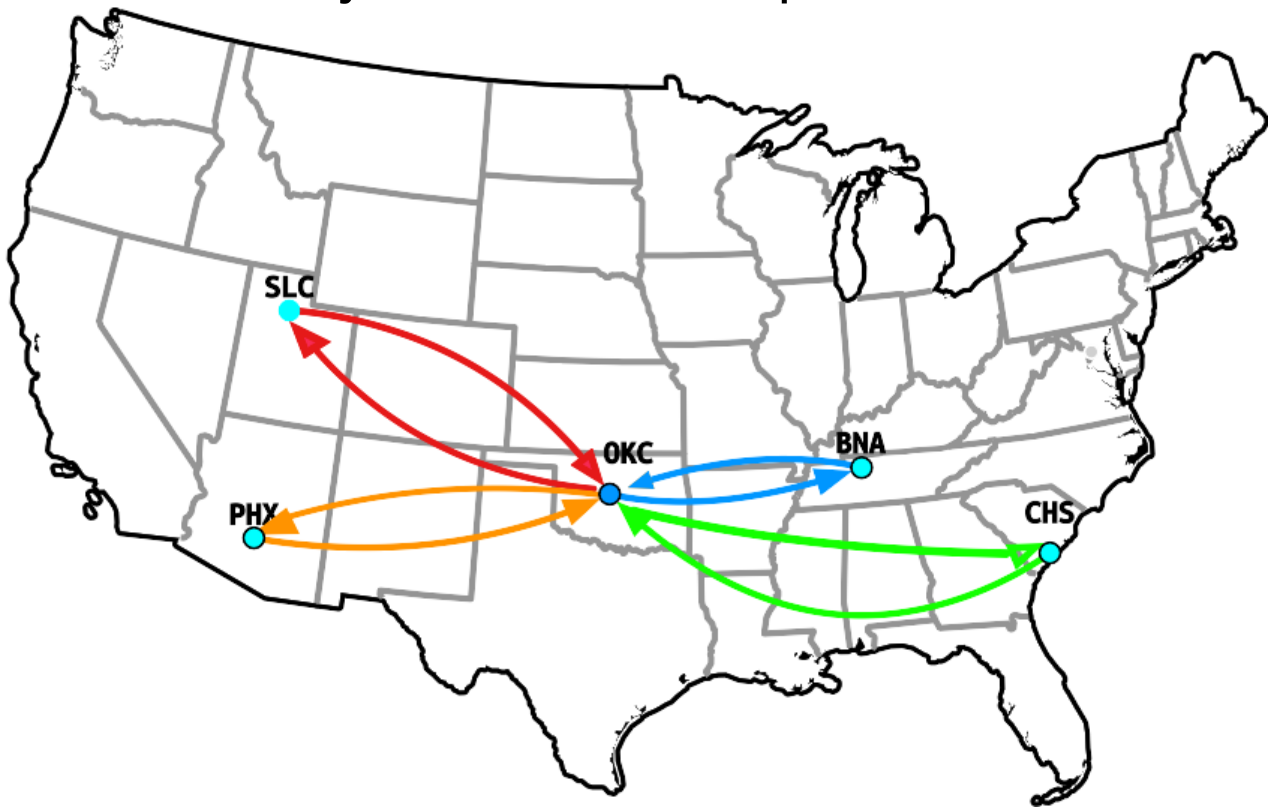

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- Filter on min($RTT_{unicast}$) -> best-case unicast RTT
  - I.e., the site that has the lowest RTT towards a target
  - The highest achievable performance (in terms of latency) for the anycast deployment

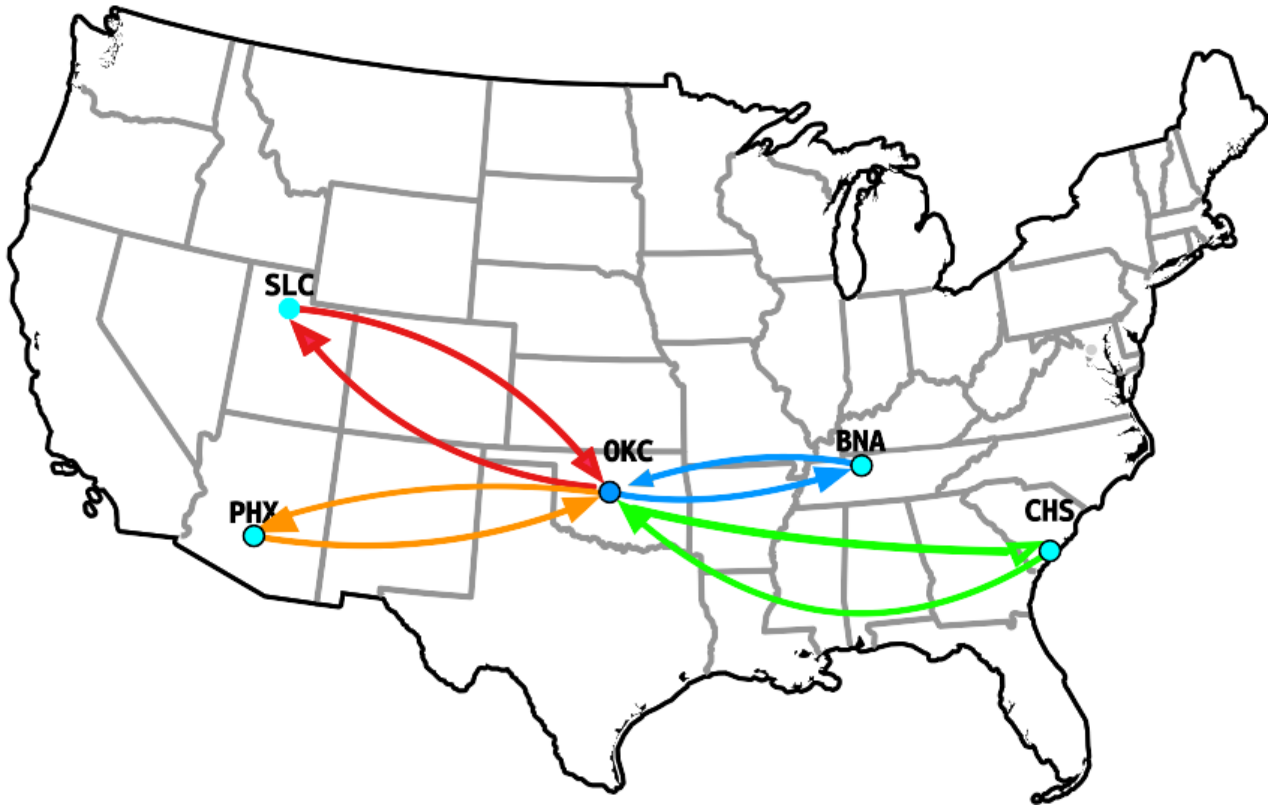

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- min($RTT_{unicast}$) mean is 39.4 ms

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- $\min(RTT_{unicast})$ mean is 39.4 ms
- $RTT_{anycast}$ mean is 65.7 ms

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- $\min(RTT_{unicast})$ mean is 39.4 ms
- $RTT_{anycast}$ mean is 65.7 ms
- $\Delta = RTT_{anycast} - \min(RTT_{unicast}) =$ 26.3 ms

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- $\min(RTT_{unicast})$ mean is 39.4 ms
- $RTT_{anycast}$ mean is 65.7 ms
- $\Delta = RTT_{anycast} - \min(RTT_{unicast}) =$ 26.3 ms
  - Substantial latency inflation due to sub-optimal anycast routing
  - Large gain in performance (+40%) possible using *e.g.,* BGP prepending, selective announcements, ...

Catchment analysis
Anycast RTT data vs. best-case unicast RTT

- $\min(RTT_{unicast})$ mean is 39.4 ms
- $RTT_{anycast}$ mean is 65.7 ms
- $\Delta = RTT_{anycast} - \min(RTT_{unicast}) =$ 26.3 ms
  - Substantial latency inflation due to sub-optimal anycast routing
  - Large gain in performance (+40%) possible using *e.g.,* BGP prepending, selective announcements, ...

- Future work:
  - Automated detection and solving of sub-optimal anycast routing

Measurement tool
Further use cases

- Identifying load distribution (as done in the Verfploeter paper)
- Simultaneous probing with multiple prefixes
  - Side-to-side comparison of catchment mapping for two different prefixes
    - e.g., assessing effectiveness of changed BGP announcements
      - e.g., what if site *X* were unavailable, where would its traffic go?
- IPv6, TCP, UDP/DNS and CHAOS probing
- Detect network regions experiencing 'site flipping' due to Load-Balancers
- And much more...

Content

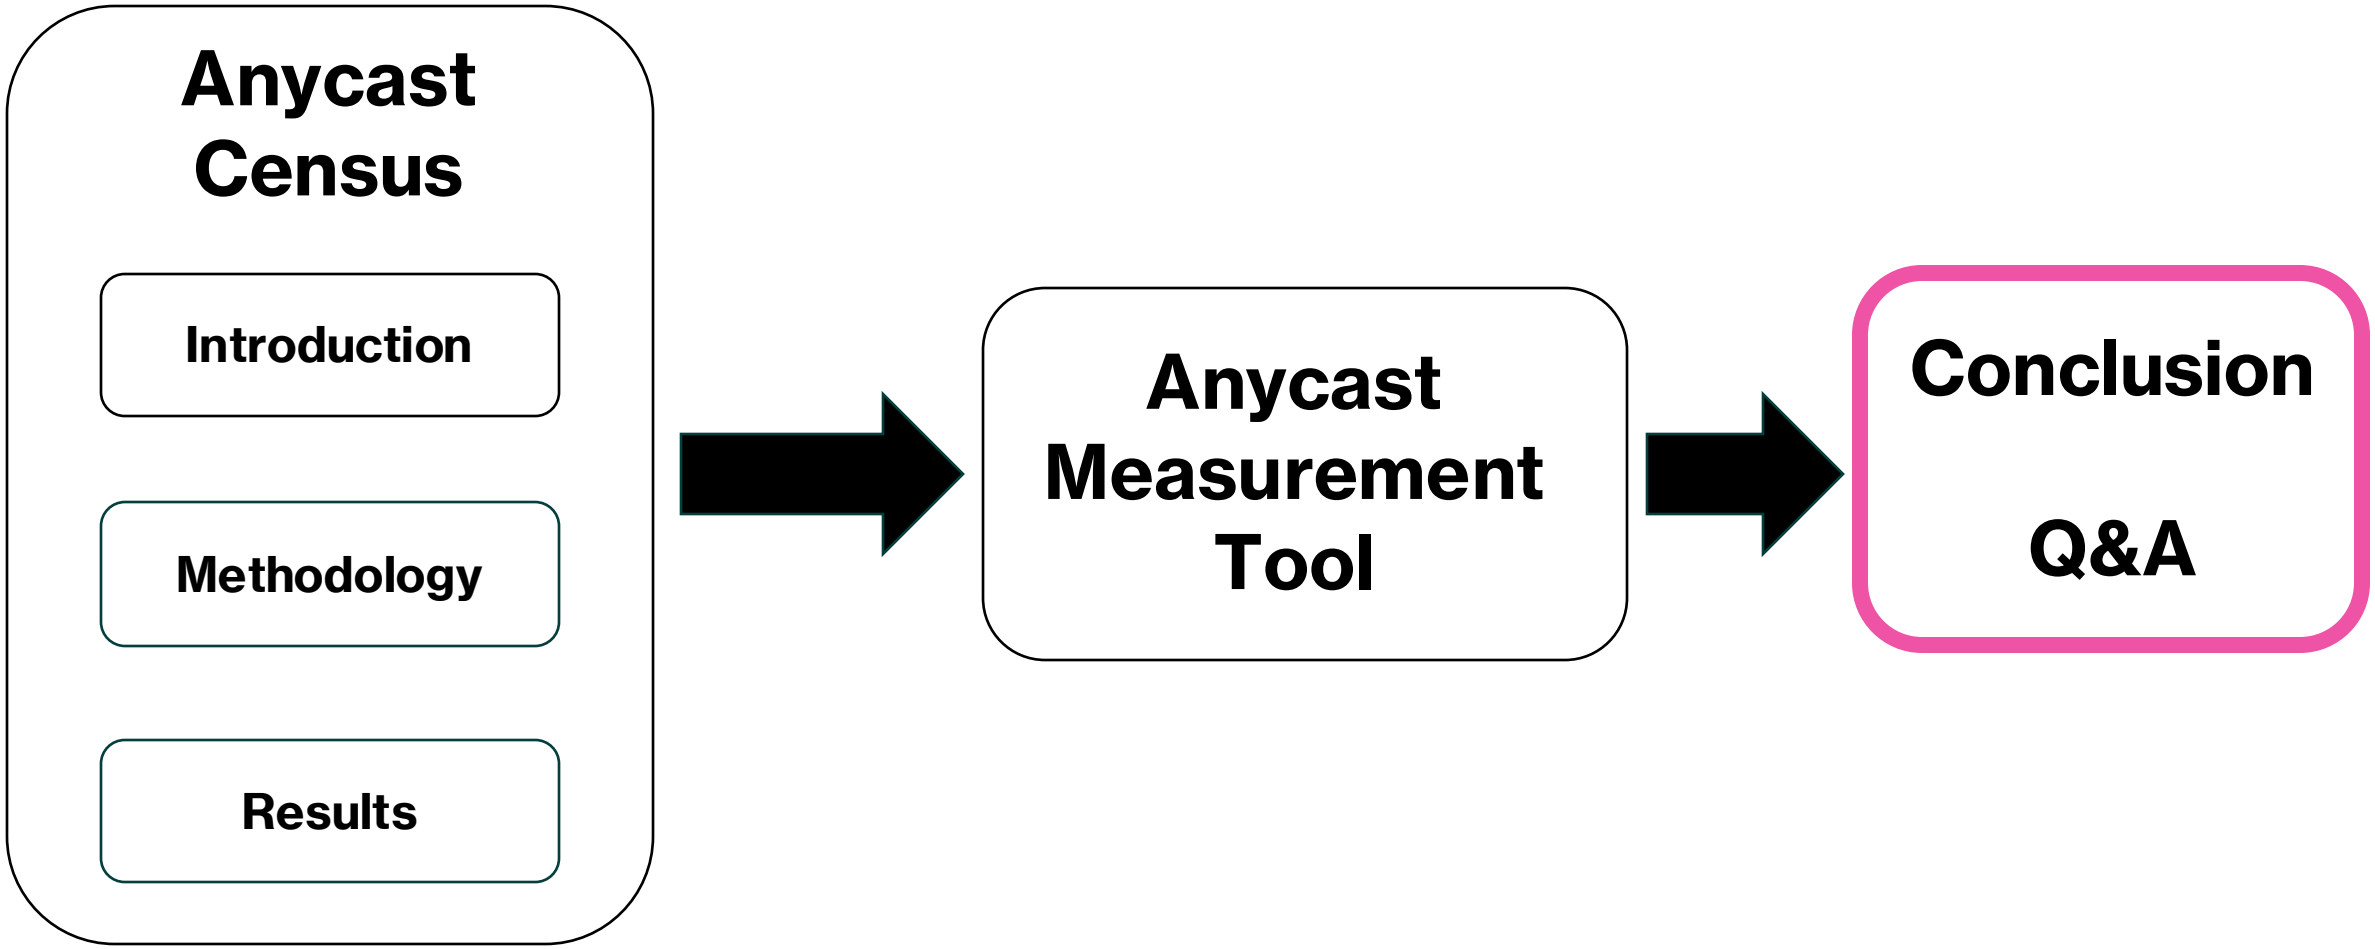
Anycast Census
Introduction
Methodology
Results
Anycast Measurement Tool
Conclusion
Q&A

Conclusion

- A daily census of anycast
  - Using two methodologies
    - Anycast-based approach
    - Latency-based approach
  - Publicly available

- Developed measurement tool
  - Public release soon

Public census repo
github.com/anycast-census/anycast-census

Future

- Refining/improving pipeline
  - E.g., canary outage detection
- Web-interface/API for live measurements
- User-friendly dashboard to visualize data
- Signal-based measurement

  of short-lived anycast
  - Anti-DDoS ASes announcements
  - BGP route collector
- Longitudinal analysis of anycast

Public census repo

github.com/anycast-census/anycast-census

Call for contribution

Please let us know if your prefix is covered by our

We need ground truth validation.

We are looking to **collaborate**

- Expand testbed infrastructure (unicast & anycast)
  - New providers/upstreams
  - More geographical coverage
  - Interested in economically developing regions
- Measurement tooling
  - Feel free to contact us

- **Contact**:
  - remi.hendriks@utwente.nl

Public census repo
github.com/anycast-census/anycast-census

Appendix slides

Catchment analysis
Anycast RTT data (example)

Site	Mean RTT (in ms)	Catchment size
de-fra	111.3	385,264
kr-icn	60.8	379,341
us-ewr	29.3	268,475
in-bom	146.1	264,222
jp-nrt	88.3	233,544
us-ord	29.3	194,599
us-mia	109.2	150,184
us-lax	51.3	146,907
br-sao	101.6	146,285
gb-lhr	33.1	140,632

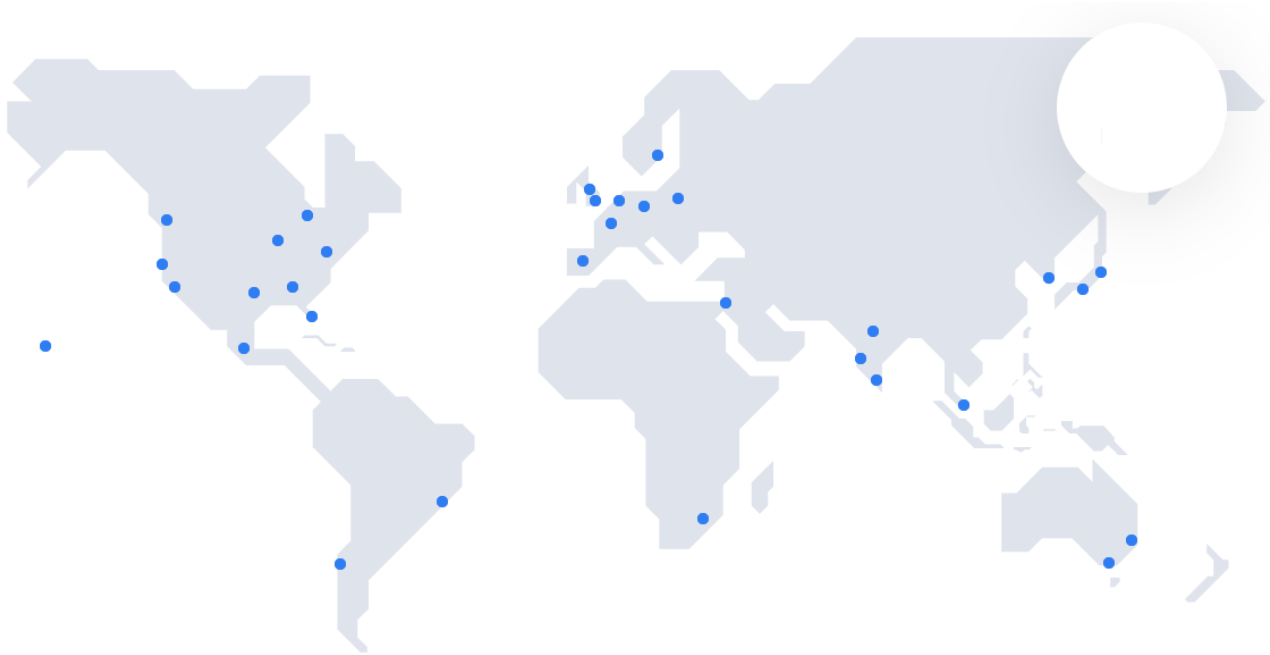

Our anycast deployment (32 locations)

Catchment analysis
Anycast RTT data (example)

Site	Mean RTT (in ms)	Catchment size
de-fra	111.3	385,264
kr-icn	60.8	379,341
us-ewr	29.3	268,475
in-bom	146.1	264,222
jp-nrt	88.3	233,544
us-ord	29.3	194,599
us-mia	109.2	150,184
us-lax	51.3	146,907
br-sao	101.6	146,285
gb-lhr	33.1	140,632

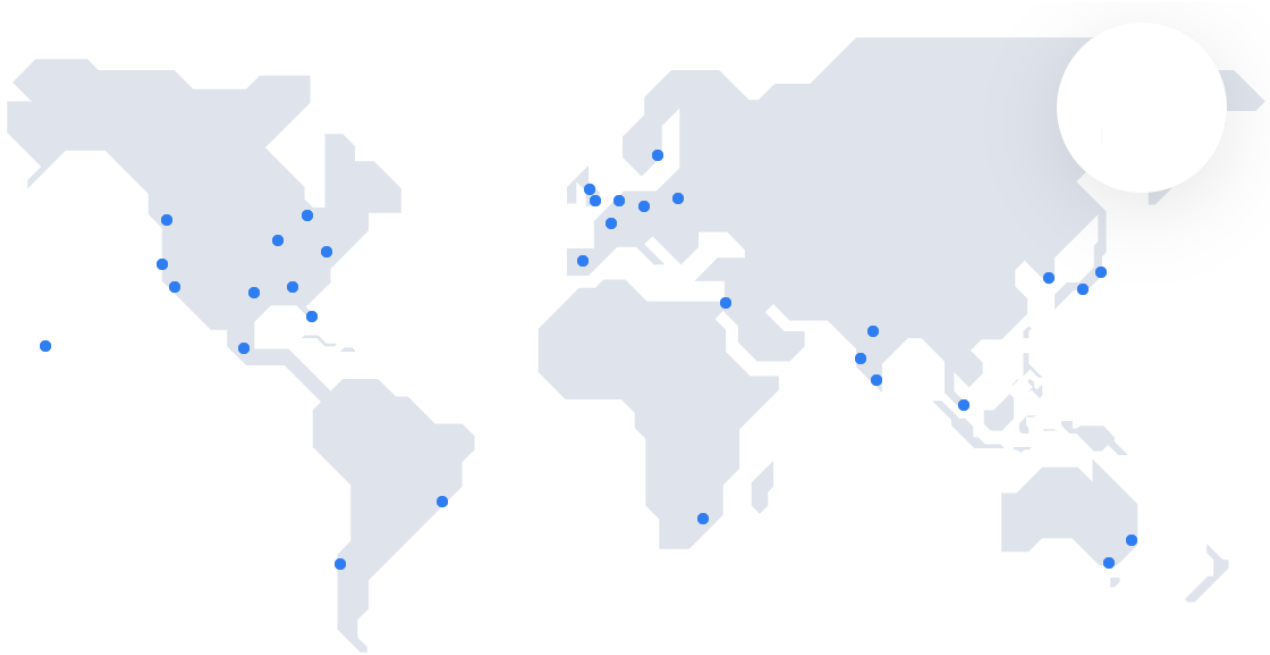

Our anycast deployment (32 locations)

Catchment analysis
Anycast RTT data (example)

- Zooming in on de-fra
- Mapping target locations using *ip2location*
- Large number of Chinese prefixes (99k)
  - Most from a few ASes

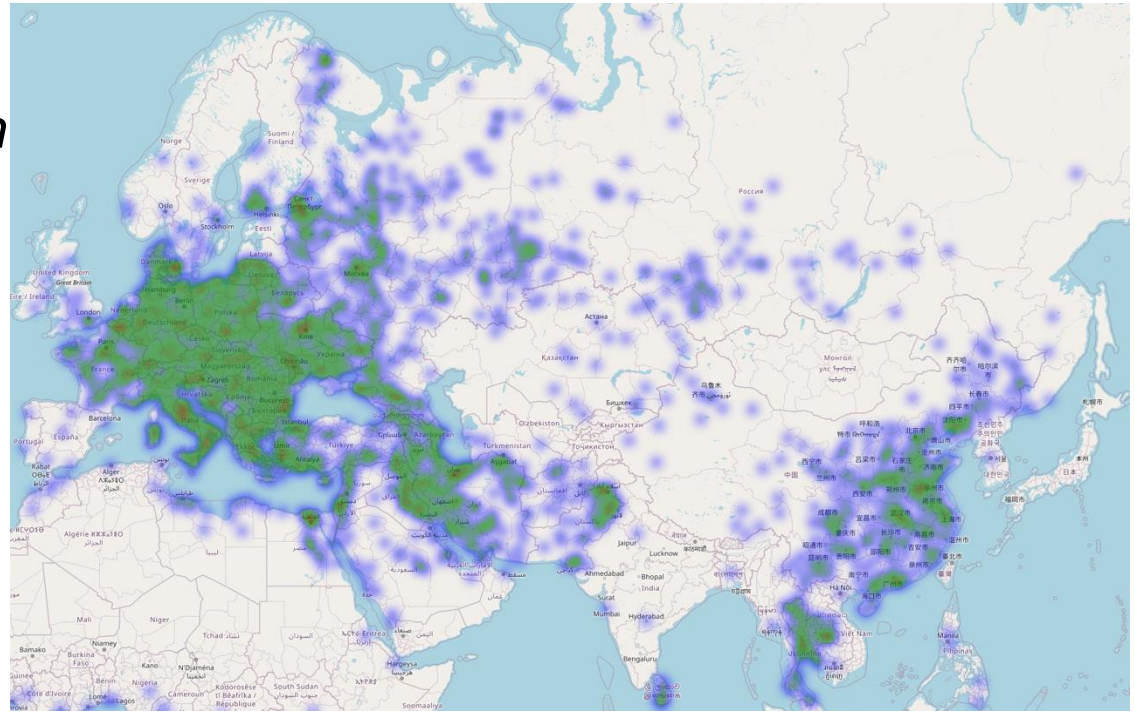

Anycast measurement tool

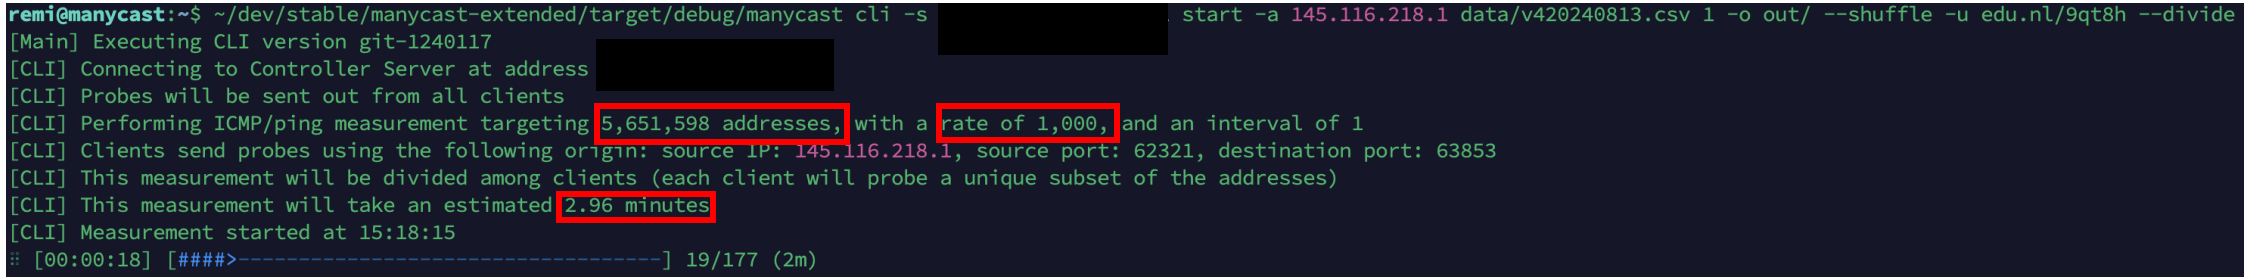

Catchment mapping of 5.7 million addresses in 3 minutes (using a probing rate of 1,000 at each anycast site)

- Measurement tool allows for mapping the catchment of the IPv4 space in a few minutes with ease
  - **Catchment mapping**: *the set of anycast targets that route to each anycast site*

- Speed and 'ease' makes daily (or even hourly) catchment mappings trivial

Verfploeter

- Methodology to perform catchment mapping

  i.e., mapping the set of addresses that route to each anycast site

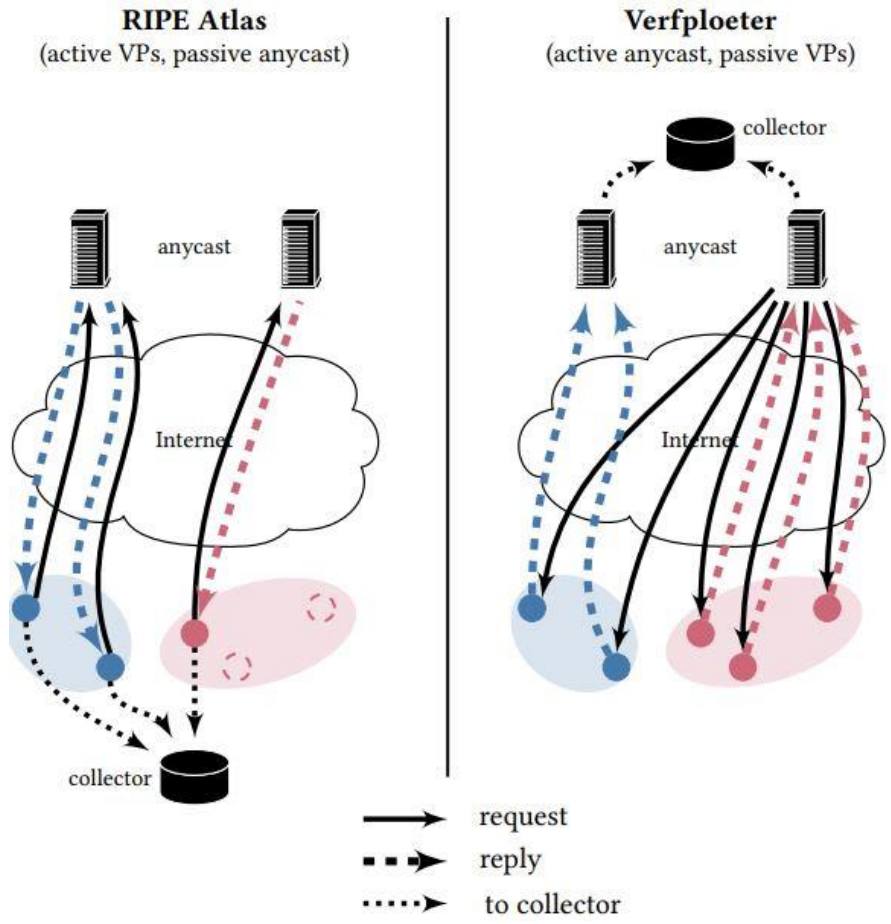

Verfploeter

- Methodology to perform catchment mapping

  i.e., mapping the set of addresses that route to each anycast site

- Verfploeter allows for catchment mapping at Internet scale
  - Uses responsive hosts on the Internet
  - Not restricted by deployment of e.g., RIPE Atlas
  - Allows for pro-active analysis

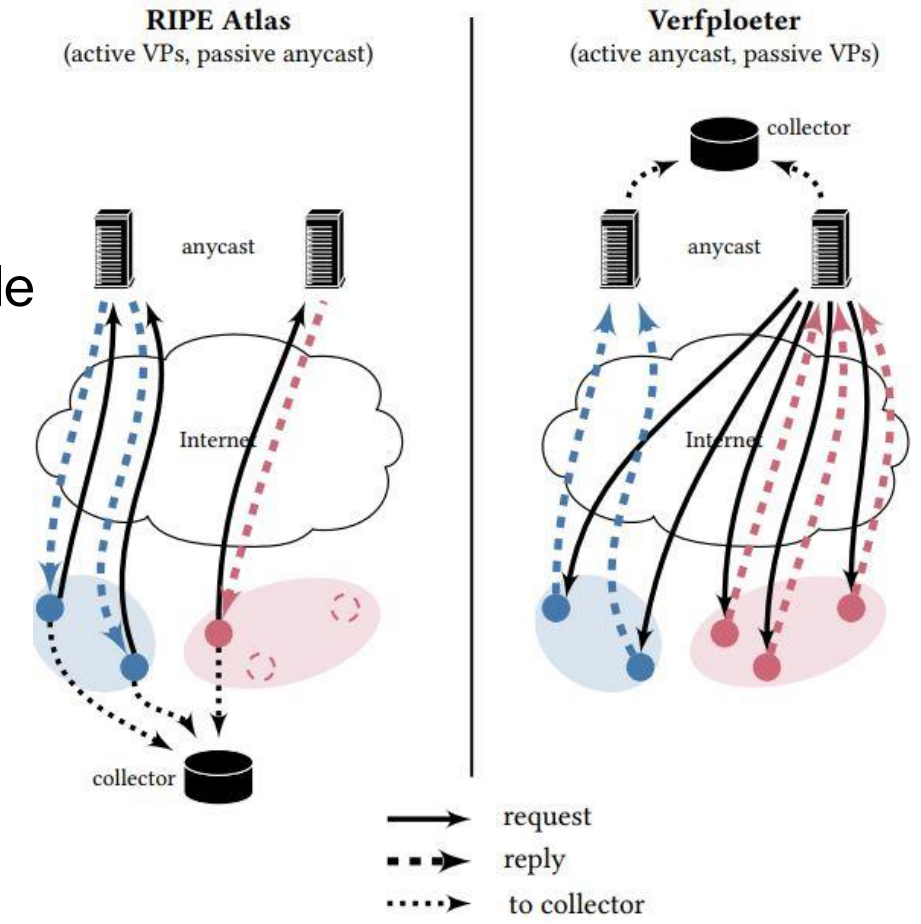

Verfploeter

- Methodology to perform catchment mapping

  i.e., mapping the set of addresses that route to each anycast site

- Verfploeter allows for catchment mapping at Internet scale
  - Uses responsive hosts on the Internet
  - Not restricted by deployment of e.g., RIPE Atlas
  - Allows for pro-active analysis
- Catchment mappings help operators to
  - Assess anycast performance
  - Plan infrastructure expansions

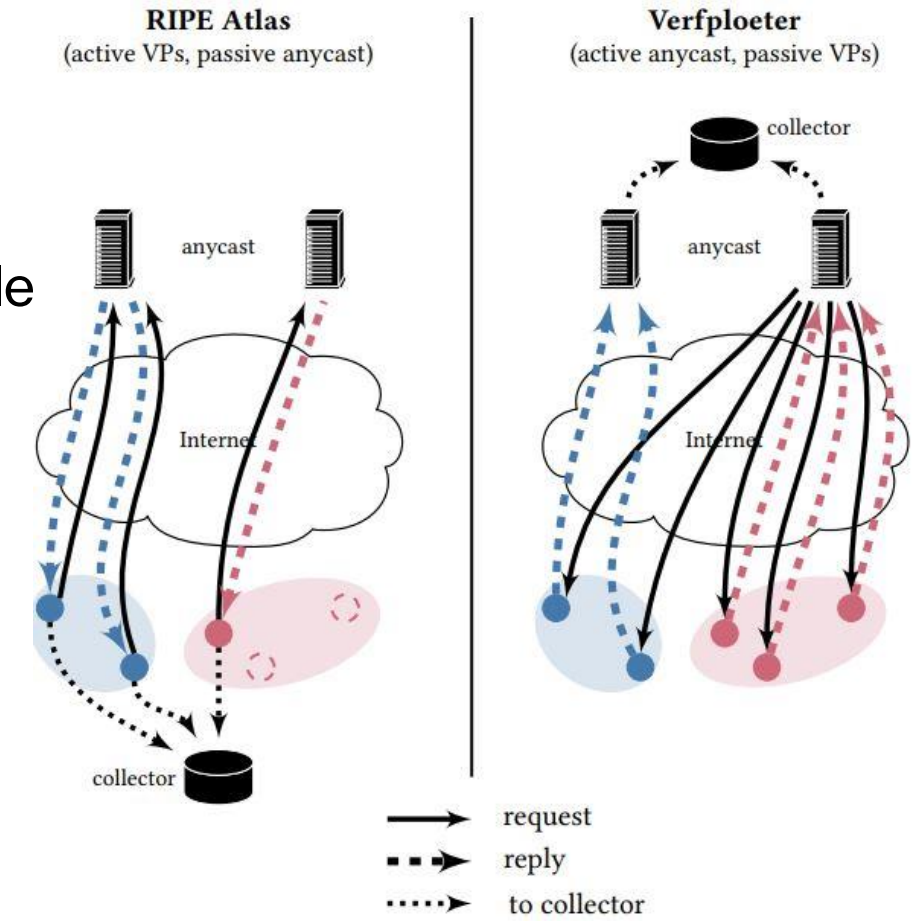

Appendix slides
Daily census output snippet classification + enumeration

File: YYYY/MM/DD/YYYY-MM-DD_v4.json


```
{"prefix": "1.0.0.0/24", "characterization":
    {"MAnycastICMPv4": {"anycast": true, "instances": 26},
    "MAnycastTCPv4": {"anycast": true, "instances": 26},
    "MAnycastUDPv4": {"anycast": null, "instances": 0},
    "iGreedyICMPv4": {"anycast": true, "instances": 60},
    "iGreedyTCPv4": {"anycast": true, "instances": 26}
    }
},
...
```

Appendix slides
Daily census output snippet locations

File: YYYY/MM/DD/YYYY-MM-DD_v4_locations.json


```
 {
   "prefix": "1.0.0.0/24",
   "count": 60,
   "instances": [
    {
      "marker": {
        "city": "Honolulu",
        "code_country": "US",
        "id": "HNL",
        "latitude": 21.3187007904,
        "longitude": -157.9219970703
      }
    },
...
```

Appendix slides
MAnycastR FNs breakdown

# of sites receiving	Candidate anycast	TPs	FPs	TP rate
2	14,966	620	14,346	4.14%
3	421	347	74	82.42%
4	307	284	23	92.51%
5	400	364	36	91.00%
5-10	918	901	17	98.15%
10-15	831	825	6	99.28%
15-20	3,890	3,886	4	99.90%
20-25	2,288	2,288	0	100.00%
25-32	2,572	2,572	0	100.00%
Total	26,593	12,087	14,506	45.45%

Table 3: Anycast-based ICMPv4 TP rate per number of sites receiving replies, taking GCD_{Ark} as ground-truth.

Appendix slides
MAnycastR IPv6 ATs by protocol

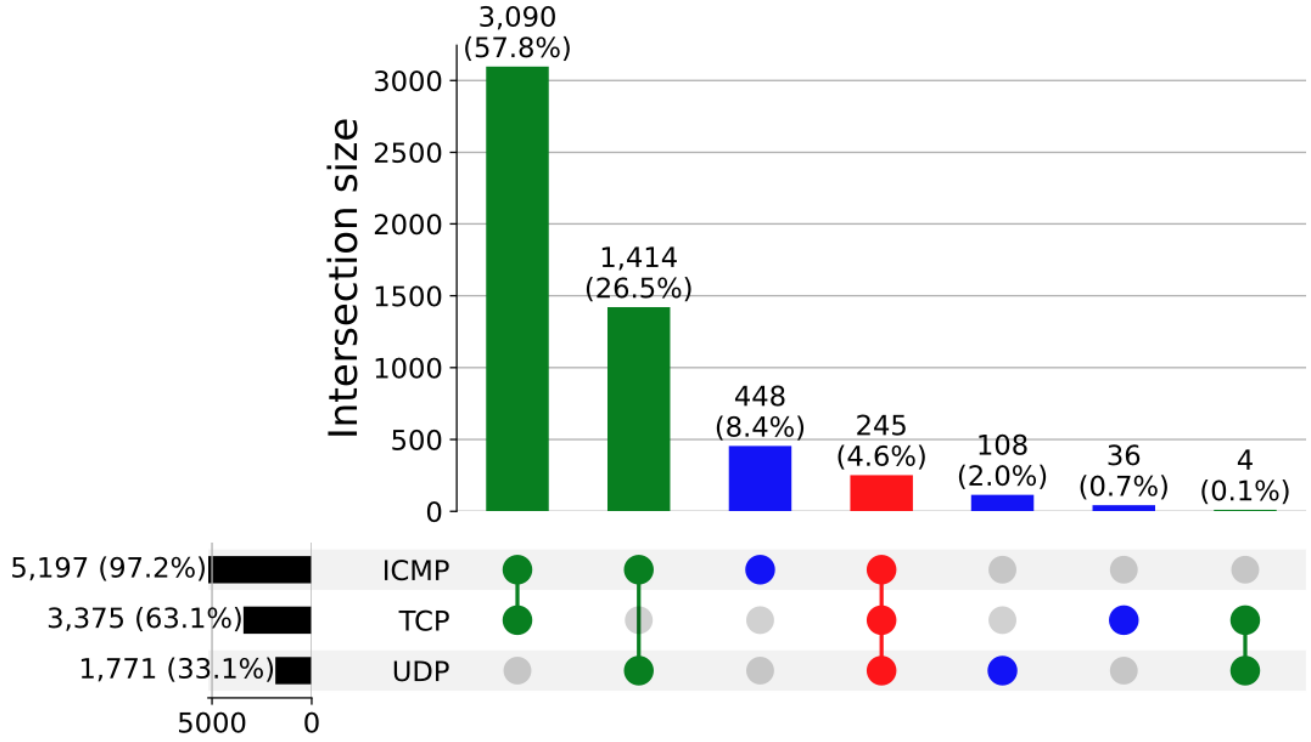

Figure 7: MAnycastR detection of anycast candidates for ICMPv6, TCPv6, and UDPv6.