



**RIPE
NCC**

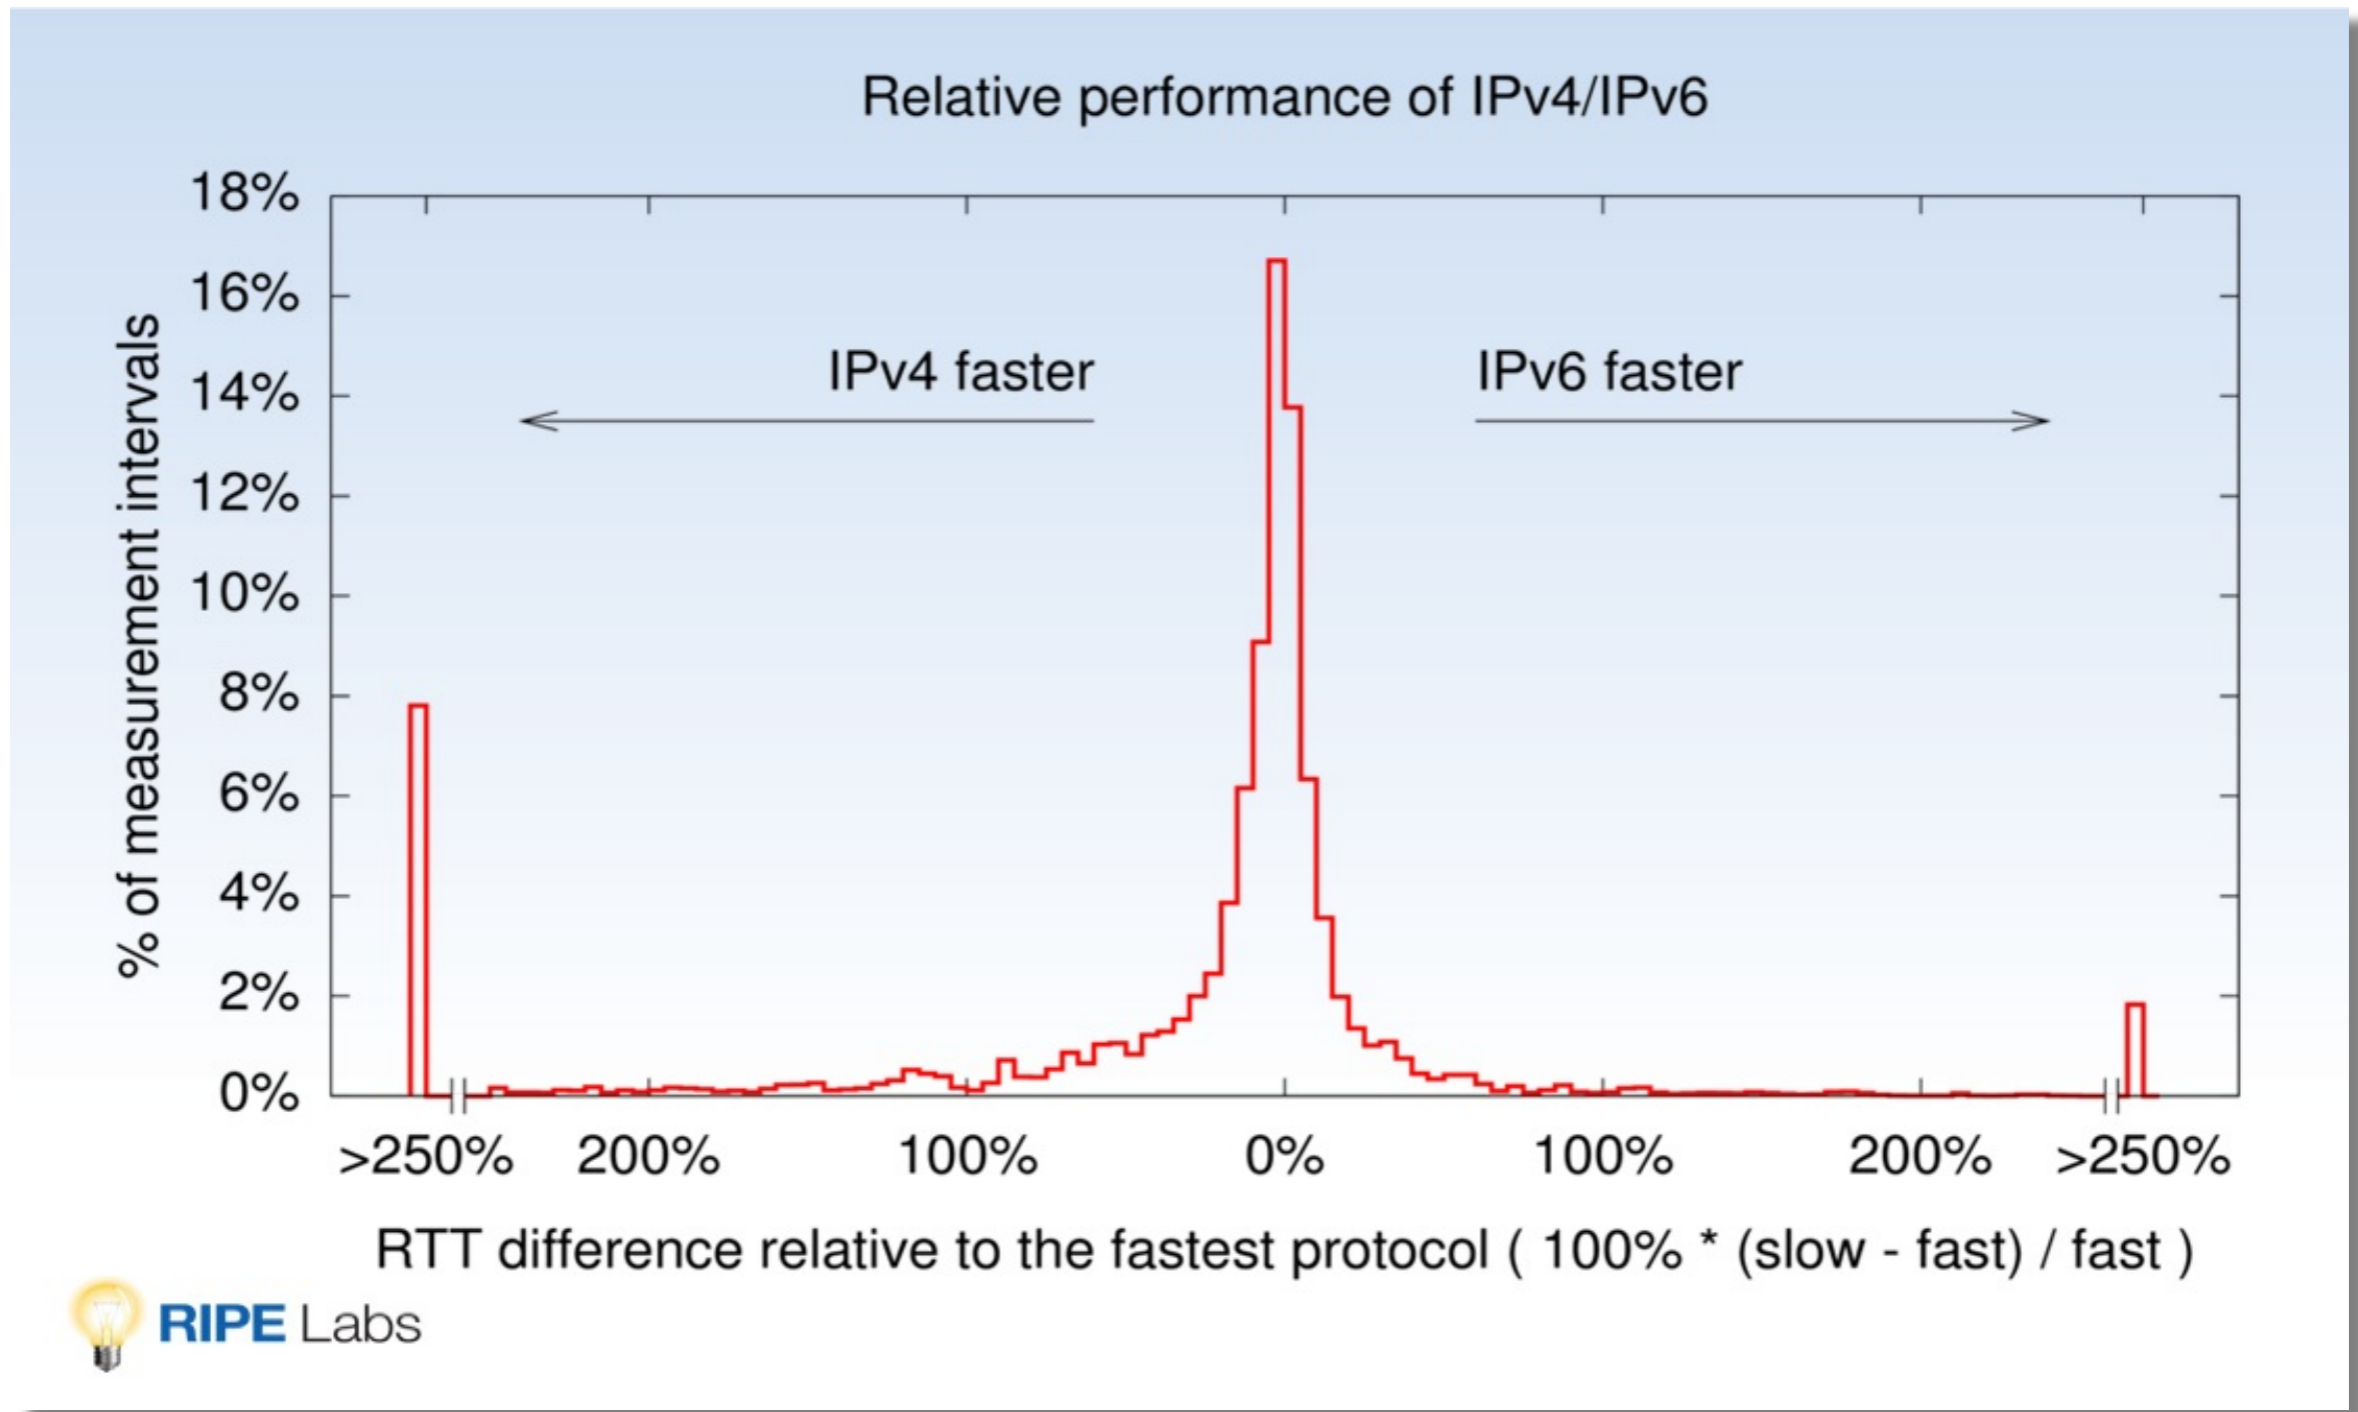
YA-IPv4/IPv6-C (yet another IPv4/ IPv6 comparison)

emile.aben@ripe.net



Latency





from <https://labs.ripe.net/Members/emileaben/measuring-world-ipv6-launch-comparing-ipv4-and-ipv6-performance>

- Facebook @Scale 2015
 - IPv6 performs better than IPv4 between US Mobile LTE networks and Facebook
 - Lots of measurements, from single (distributed) source network to a homogenous set of destination networks
 - Proves there are (very large) environments where IPv6 is consistently outperforming IPv4

... and now we have RIPE Atlas Anchors



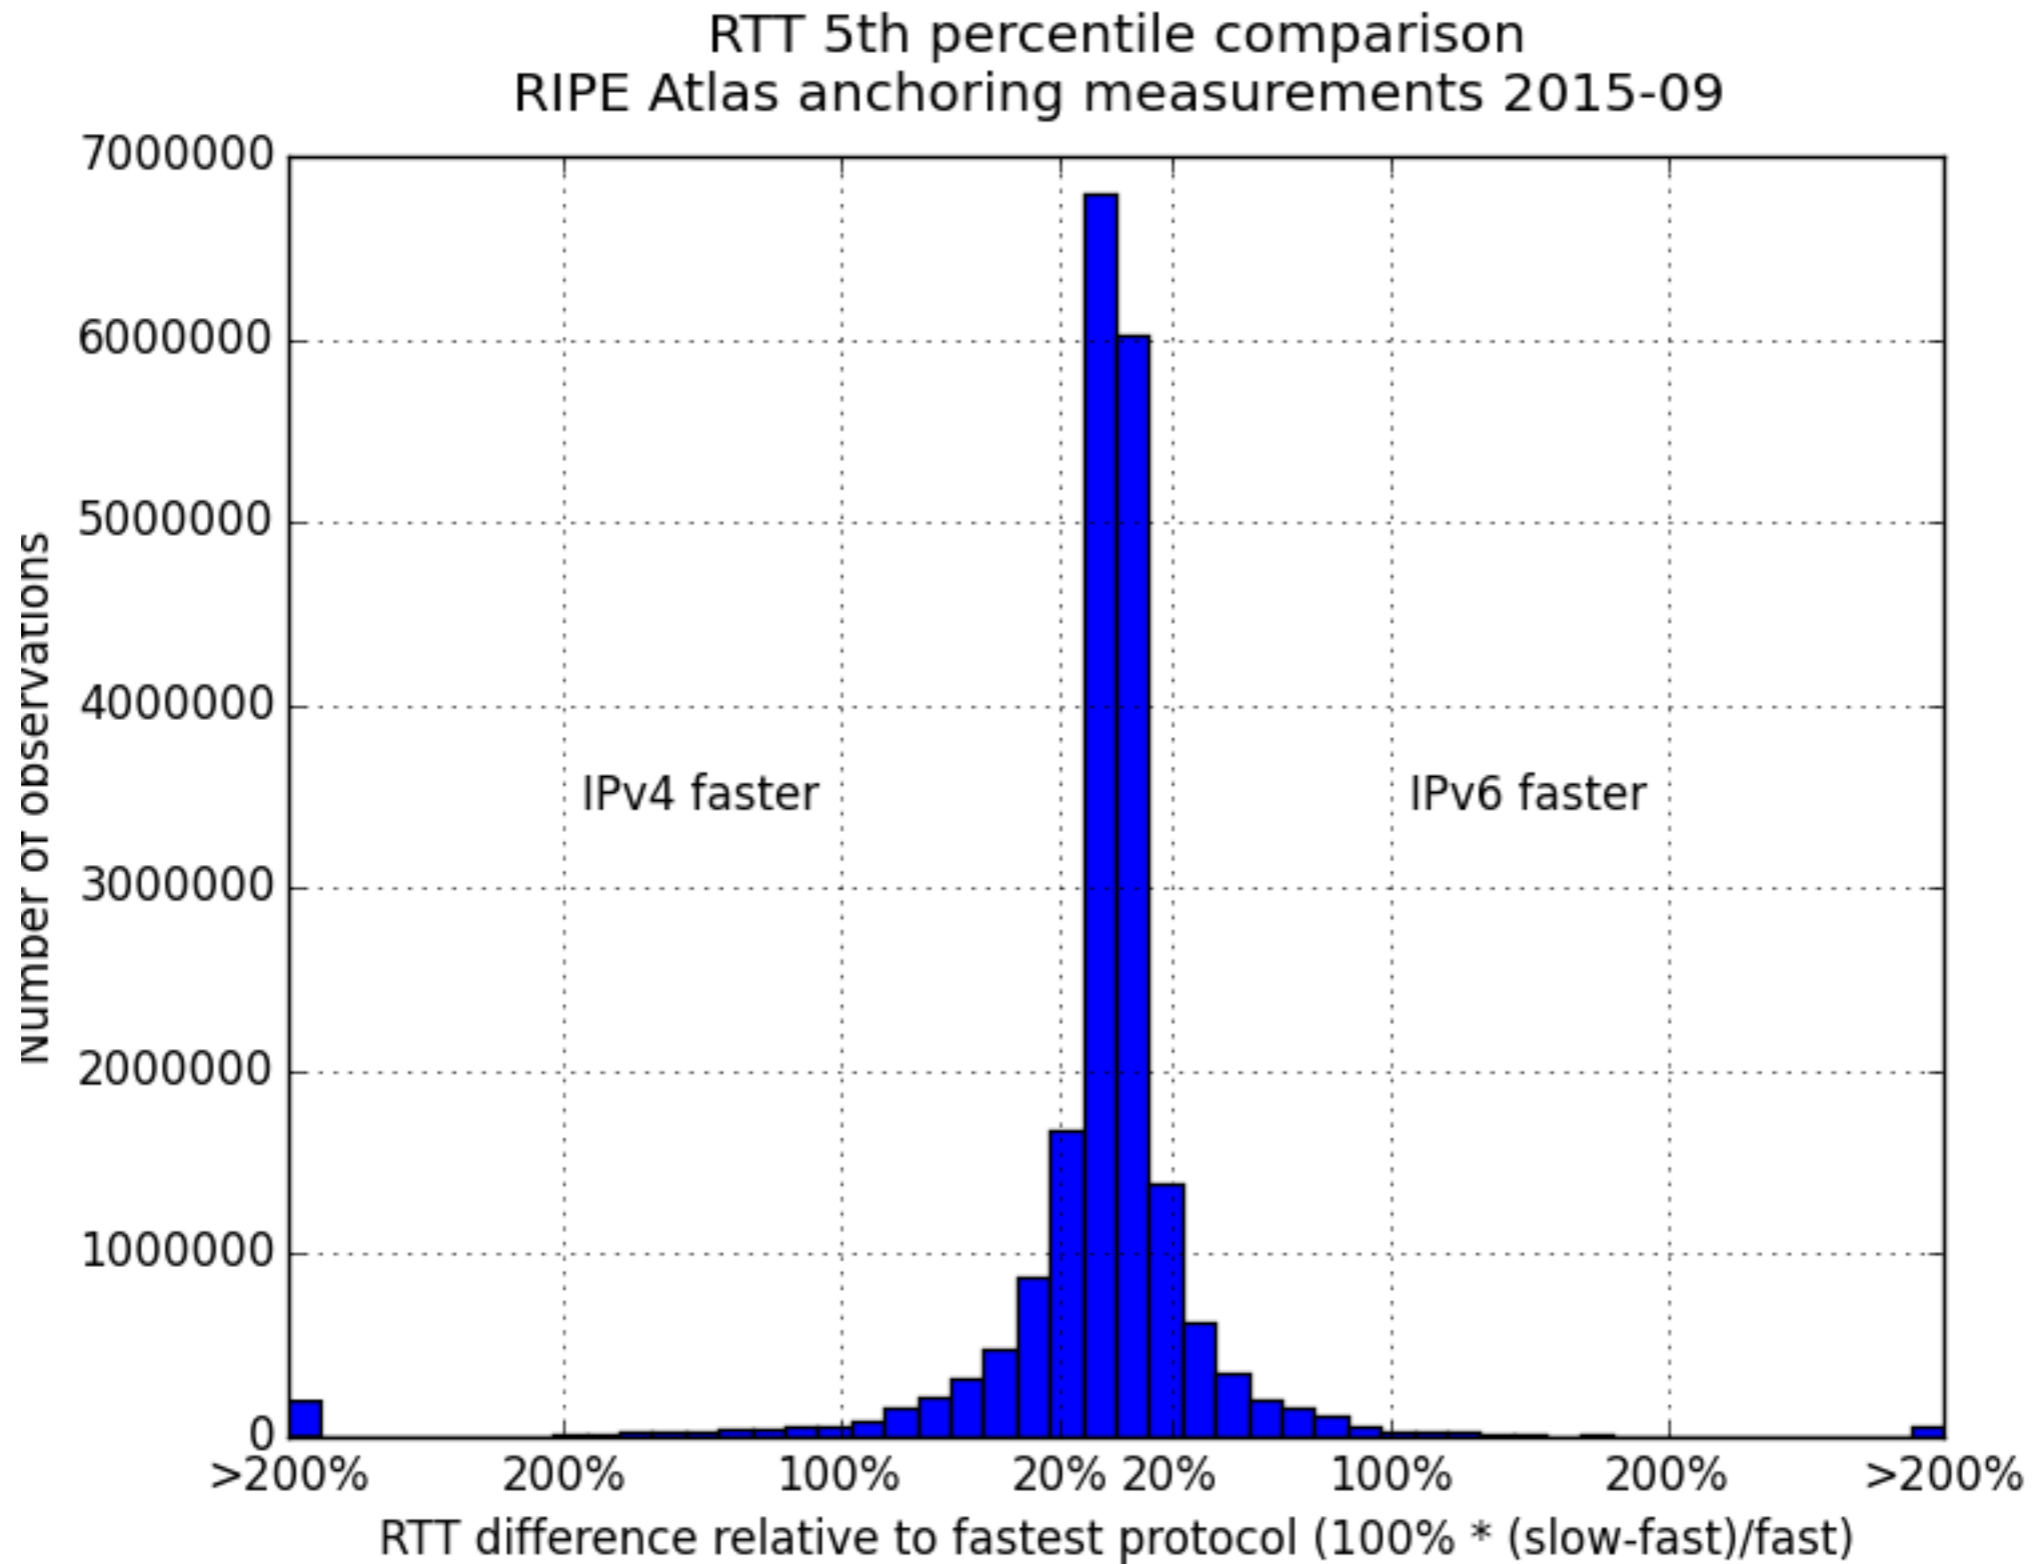
Total number of RIPE Atlas anchors: 152

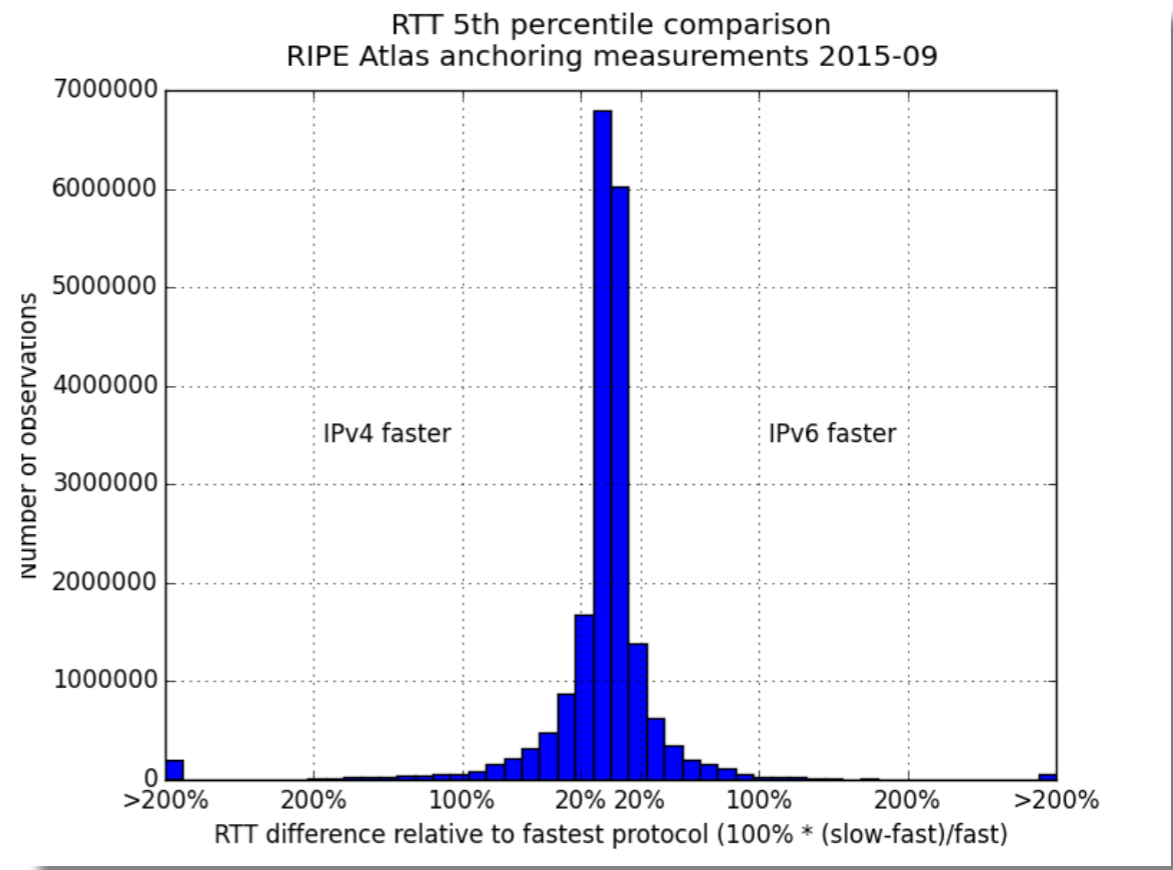
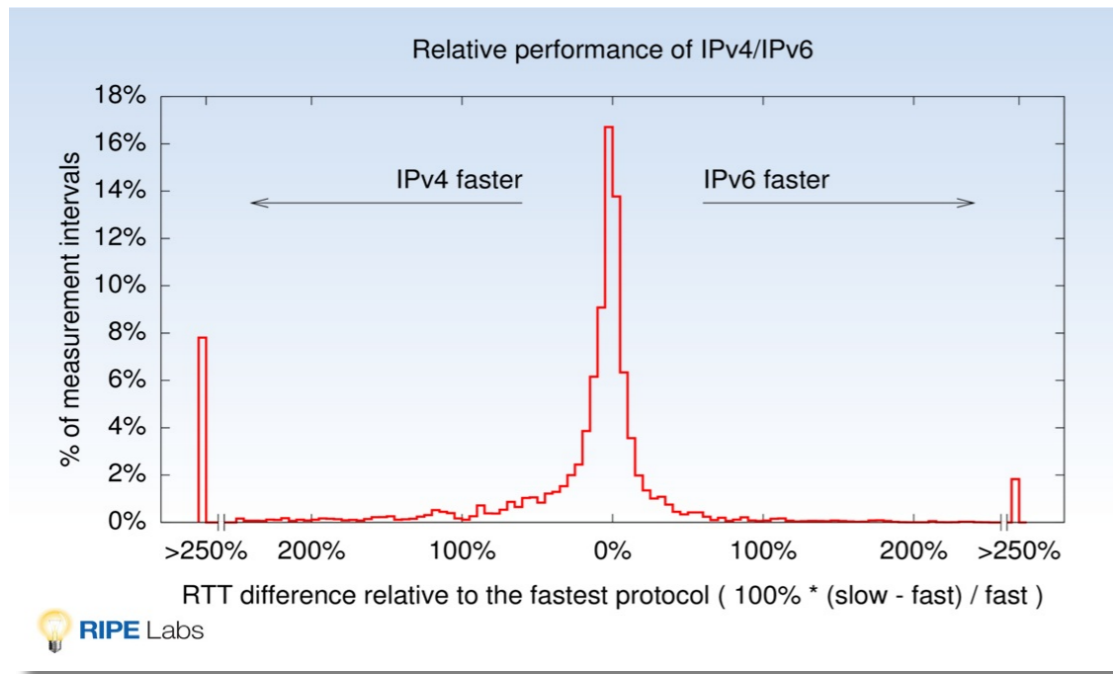
<https://atlas.ripe.net/anchors/map/>

<https://atlas.ripe.net/anchors/list/>

- RIPE Atlas Anchors
 - Biased set (“clue core”)
 - IPv4 and IPv6 are prerequisites for hosts
- Anchoring measurements, since 2013-09
 - Inbound ping/traceroute to each anchor
 - From all other anchors
 - + ~100 “randomly” selected probes
 - Both address families (if available)
- A lot of v4/v6 comparison points!
 - Same source hardware, same destination hardware

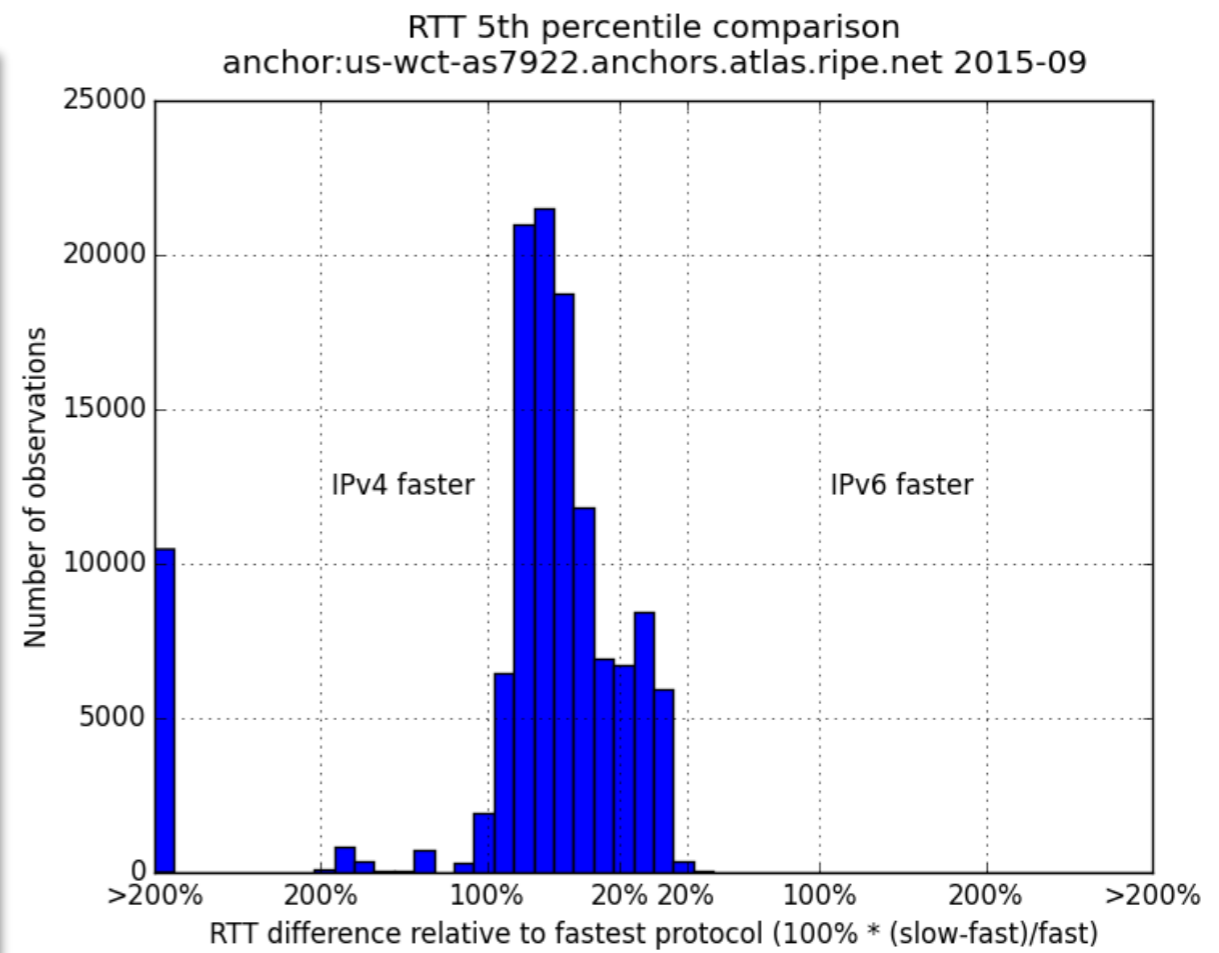

Then	Now
2012	2015
sources: 52	sources: 2941 sources (RIPE Atlas probes and anchors)
dests: 33 webservers (by hostname)	dests: 149 RIPE Atlas anchors





- Similar, distribution slightly fatter on IPv4 side
- Extremes have mostly disappeared
 - Likely cause: Difference in measurement setup

- Same network with 2 anchors: radically different
 - Contact with operator: routing config problem, being fixed
 - ie. to measure is to improve, to measure and report even more



- If your network hosts an anchor:
- <http://sg-pub.ripe.net/emile/v4v6-anchors/>

- Probes are to low number of destinations, so very low diversity of networks tested.

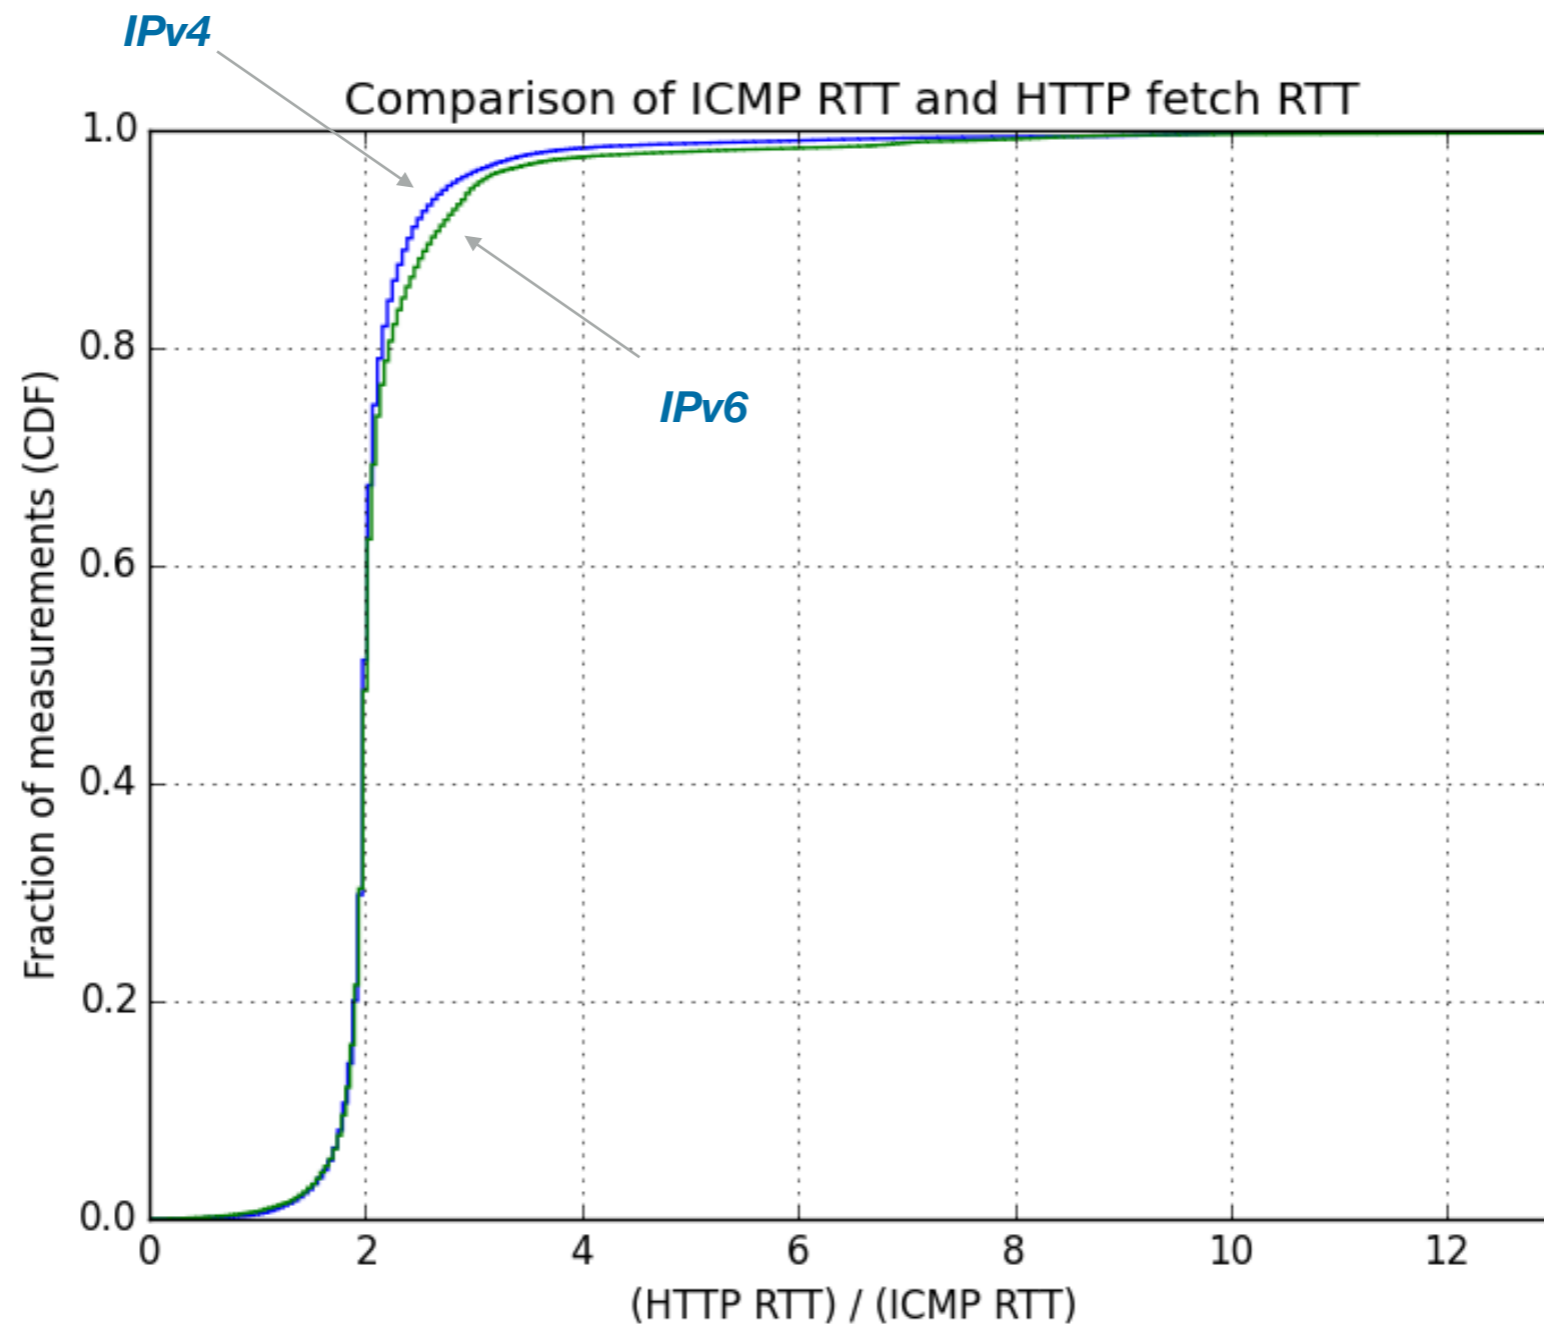


ICMP vs. HTTP

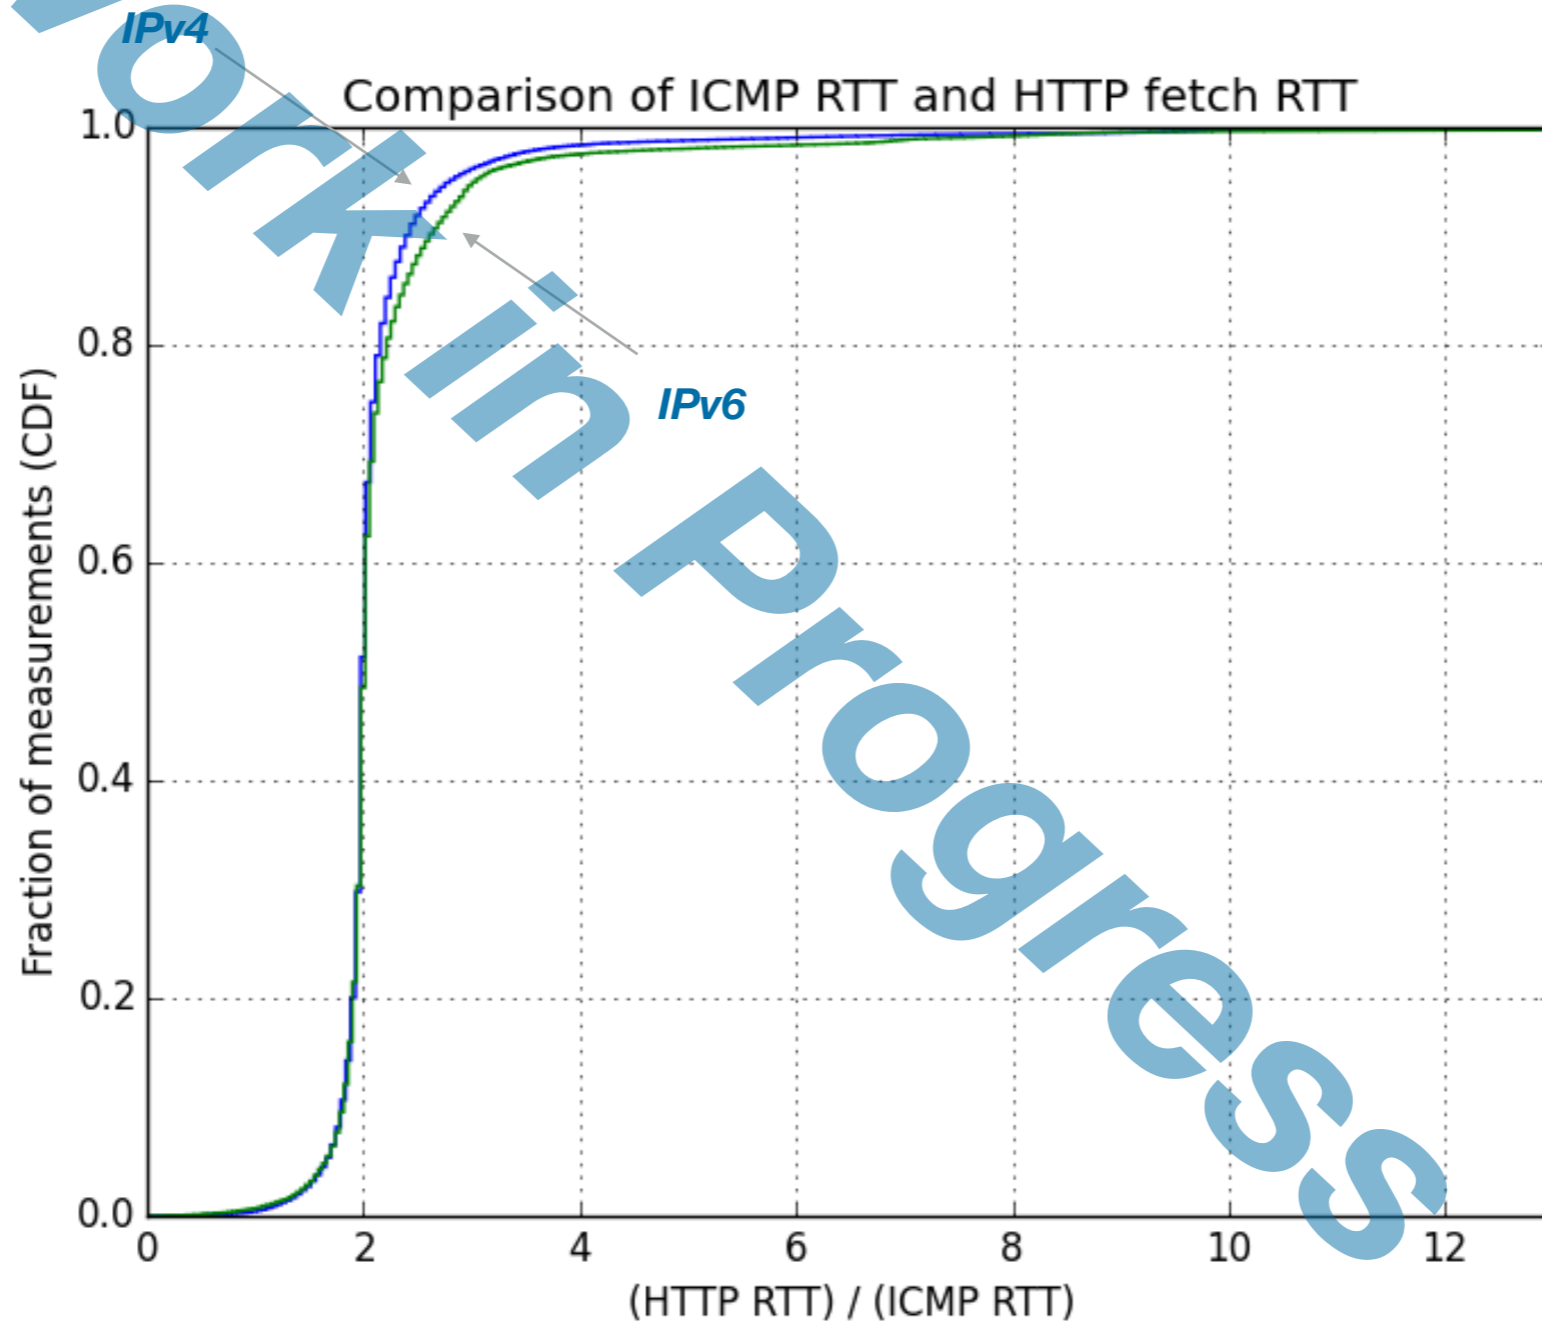


- HTTP fetch for same src-dst combinations as ping/traceroute anchoring measurements
- Experimental, let us know what you think
- Current settings:
- 4096 byte HTTP fetch, no options tweaking
 - ie. 2x round-trip (TCP handshake) + ϵ for data
- Measurement IDs for all anchoring measurements:
 - <https://atlas.ripe.net/api/v2/anchor-measurements/?format=json>
- Only few days of data

- Comparing only dual-stack pairs



- Comparing only dual-stack pairs





Availability



Leveraging the IPv4/IPv6 Identity Duality using Multipath Transport

Ioana Livadariu, Simone Ferlin, Ozgu Alay, Thomas Dreibholz, Amogh Dhamdhere, Ahmed Elmokashfi

IEEE Global Internet Symposium (GI), Hong Kong, Apr 2015.

In this paper, we first investigate the AS level congruency of IPv4 and IPv6 paths in the Internet. We find that more than 60% of the current IPv4 and IPv6 AS-paths are non-congruent at the AS-level, which motivates us to explore how MPTCP can utilize the IPv4/IPv6 identity duality to improve data transfer performance.

Leveraging the IPv4/IPv6 Identity Duality using Multipath Transport

Ioana Livadariu, Simone Ferlin, Ozgu Alay, Thomas Dreibholz, Amogh Dhamdhere, Ahmed Elmokashfi

IEEE Global Internet Symposium (GI), Hong Kong, Apr 2015.

In this paper, we first investigate the AS level congruency of IPv4 and IPv6 paths in the Internet. We find that more than 60% of the current IPv4 and IPv6 AS-paths are non-congruent at the AS-level, which motivates us to explore how MPTCP can utilize the IPv4/IPv6 identity duality to improve data transfer performance.

- Can we see leverage-able effects in Anchoring measurements too?

Leveraging the IPv4/IPv6 Identity Duality using Multipath Transport

Ioana Livadariu, Simone Ferlin, Ozgu Alay, Thomas Dreibholz, Amogh Dhamdhere, Ahmed Elmokashfi

IEEE Global Internet Symposium (GI), Hong Kong, Apr 2015.

In this paper, we first investigate the AS level congruency of IPv4 and IPv6 paths in the Internet. We find that more than 60% of the current IPv4 and IPv6 AS-paths are non-congruent at the AS-level, which motivates us to explore how MPTCP can utilize the IPv4/IPv6 identity duality to improve data transfer performance.

- Can we see leverage-able effects in Anchoring measurements too?
 - Maybe, <numbers left out>

- Results multi-interpretable
- IPv4 slightly “better” than IPv6
 - IPv4/IPv6 agility (Happy Eyeballs) covers up problems
 - With good agility only single stack users will notice/
complain
- IPv6 can cover up when IPv4 is bad and vice versa

- Results multi-interpretable
- IPv4 slightly “better” than IPv6
 - IPv4/IPv6 agility (Happy Eyeballs) covers up problems
 - With good agility only single stack users will notice/complain
- IPv6 can cover up when IPv4 is bad and vice versa

*if ipv4_ipv6_protocol_agility and dual_stack:
2 * chance_for_best_performance*

