

The Unexpected Responsiveness of Internet Hosts

Neil Spring



Me


- Measure the Internet to evaluate and justify protocols that increase network reliability.
- Thesis work - measuring how routers are connected in practice to evaluate and enhance routing protocols in terms of how they exploit common network designs in routing around failures.
- Recent work - measuring when residential links fail to determine how people and protocols should respond to faults.

Residential Link Reliability

- Residential links are:
 - Important: VoIP/911, Security cameras, Thermostats
 - Vulnerable: Exposed to weather, loss of power, singly-connected



It's Personal



What I mean by “how ... to respond to faults”

- Small-scale individual questions:
 - Should I get more than one provider? Or change?
 - Is it just me?
- System builder questions:
 - Would it help to coordinate with neighbors for mutual backup?
 - What fraction of errors can “Network Diagnostics” diagnose?
- Policy questions:
 - Do cities with more buried wiring fare better or worse?
 - How does Maryland compare to Virginia, North America to Europe?

How to detect network failures

- “ping” is the fundamental tool.
 - Innocuous packets that have only one purpose (excuse me, are you alive?)
 - A response shows that the recipient is reachable and alive.



No response \Rightarrow failure






- IP service allows four bad things to happen to your packets: **delay**, duplication, corruption, and **loss**.
- A lost echo request (are you there?) or reply (I sure am!) should happen 1-3% of the time without major failure.

ThunderPing

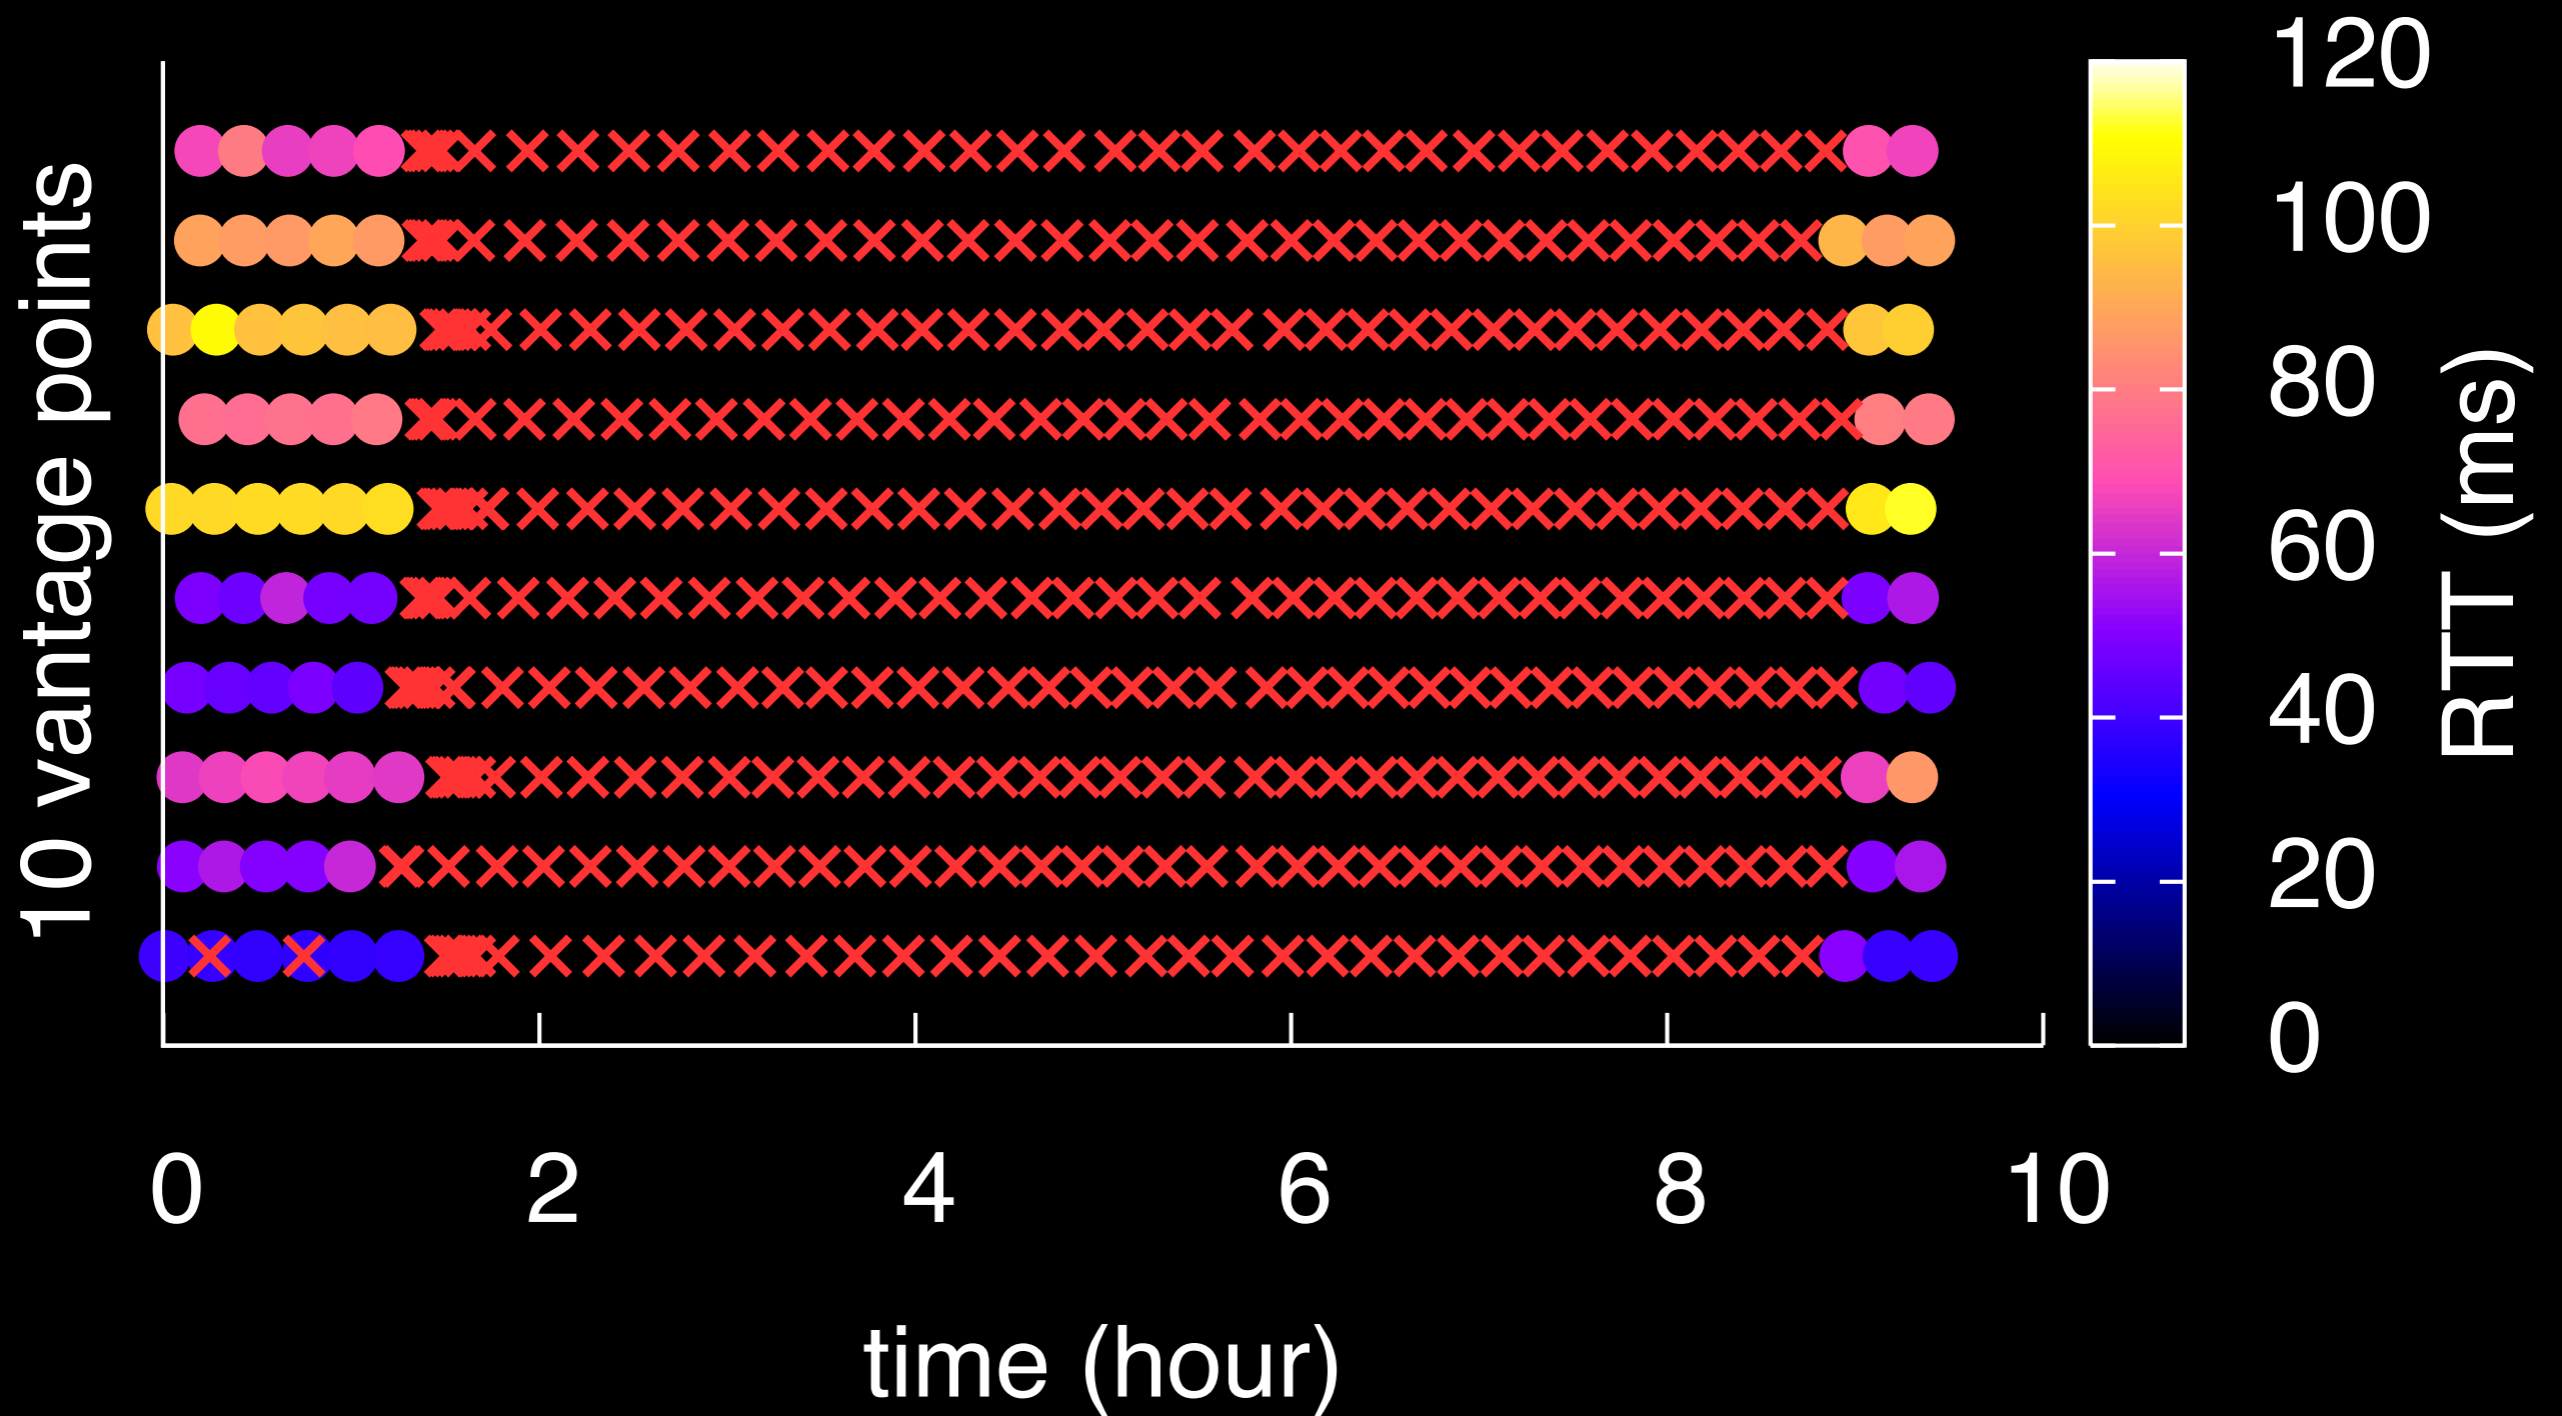


- 1. Watch for severe weather alert forecasts
- 2. Ping addresses thought to be in that region before, during, and after the alert
- 3. Figure out if there actually was weather, correlate failures with conditions




	.3%
	.4%
	.3%
	2.0%
	3.0%

Lost pings \Rightarrow outages



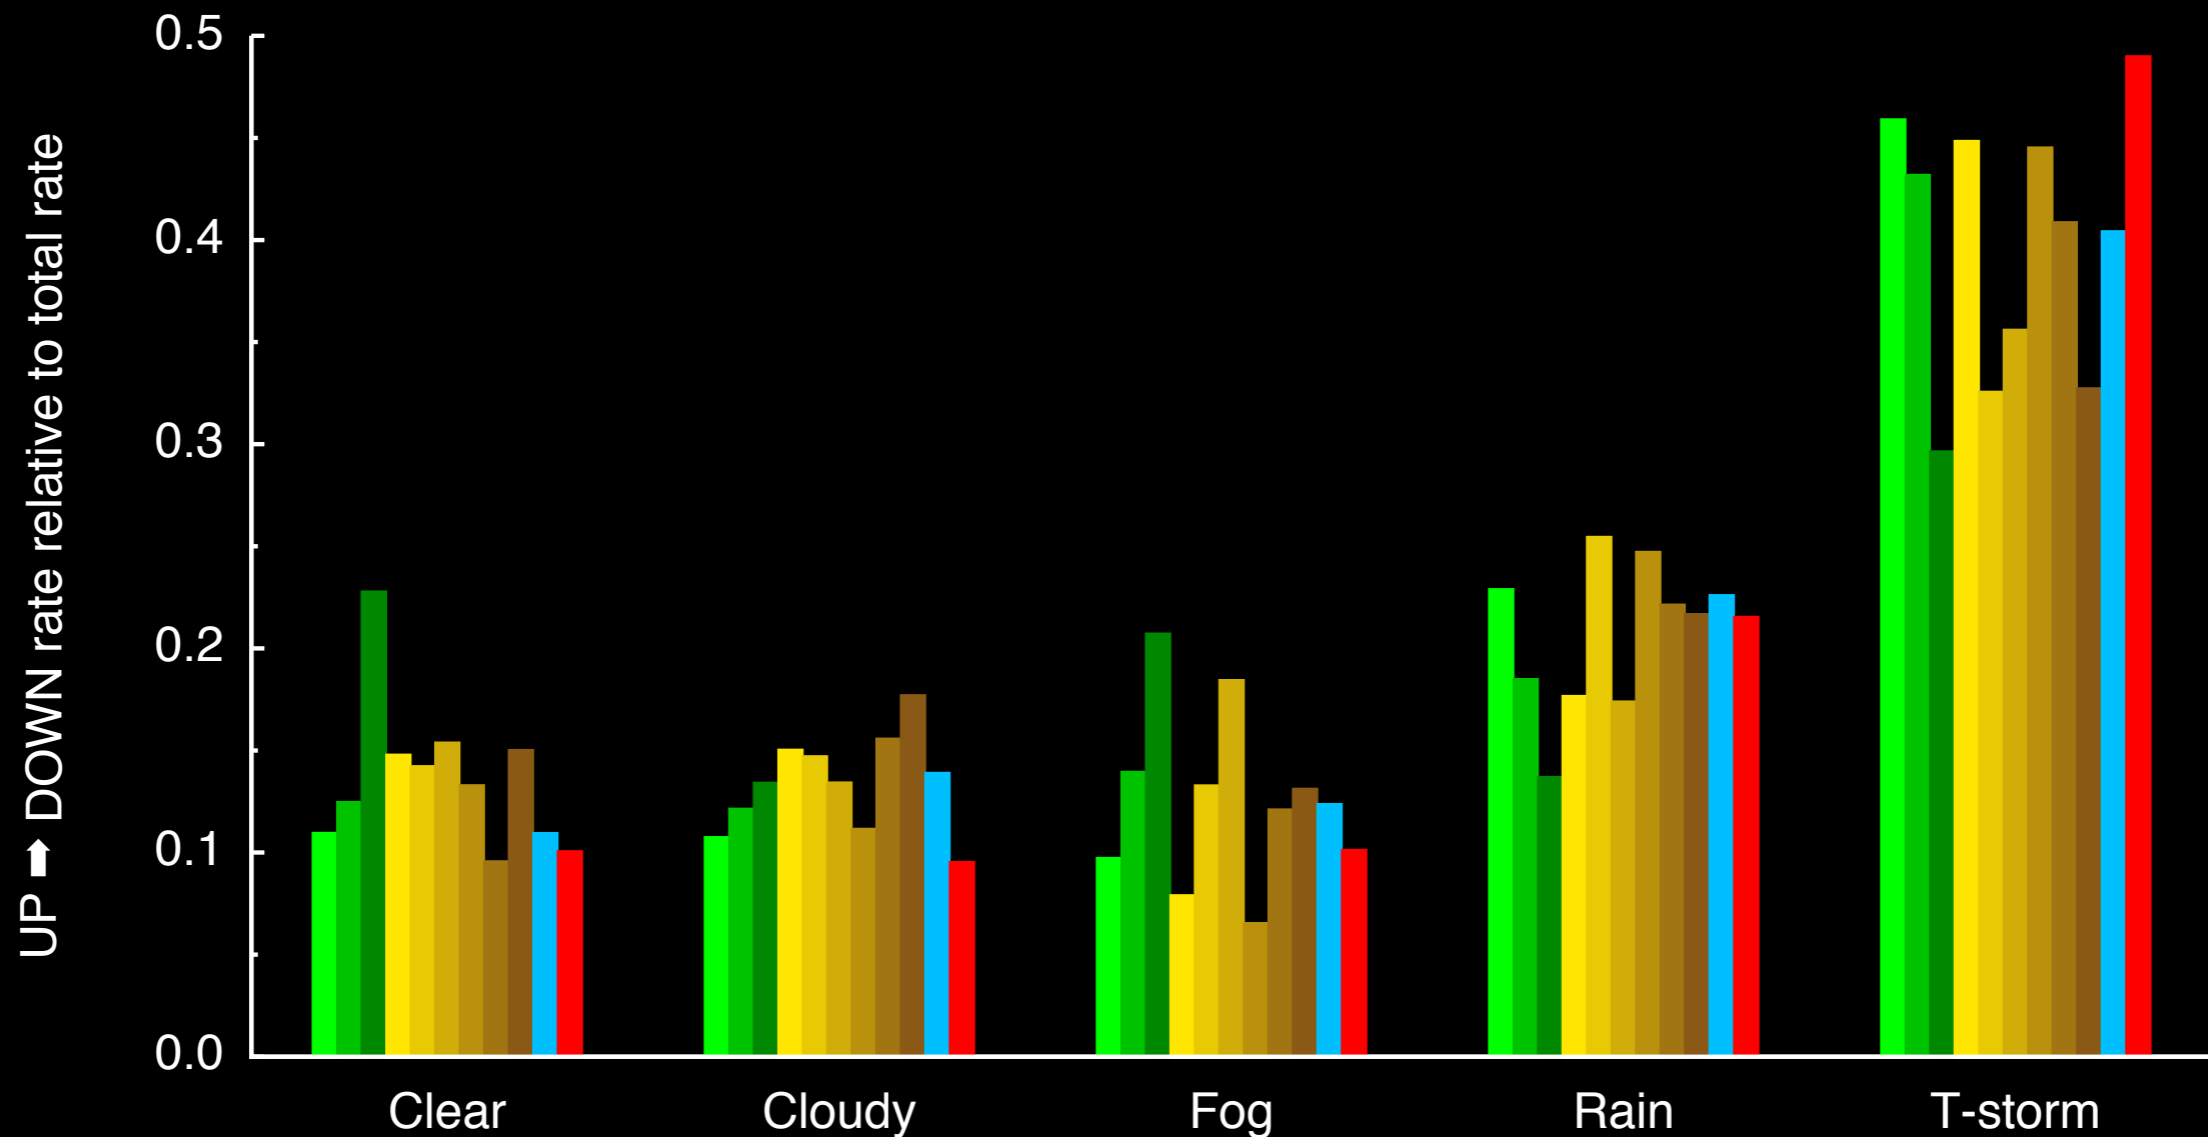
October 29, 2012 05:53 EDT



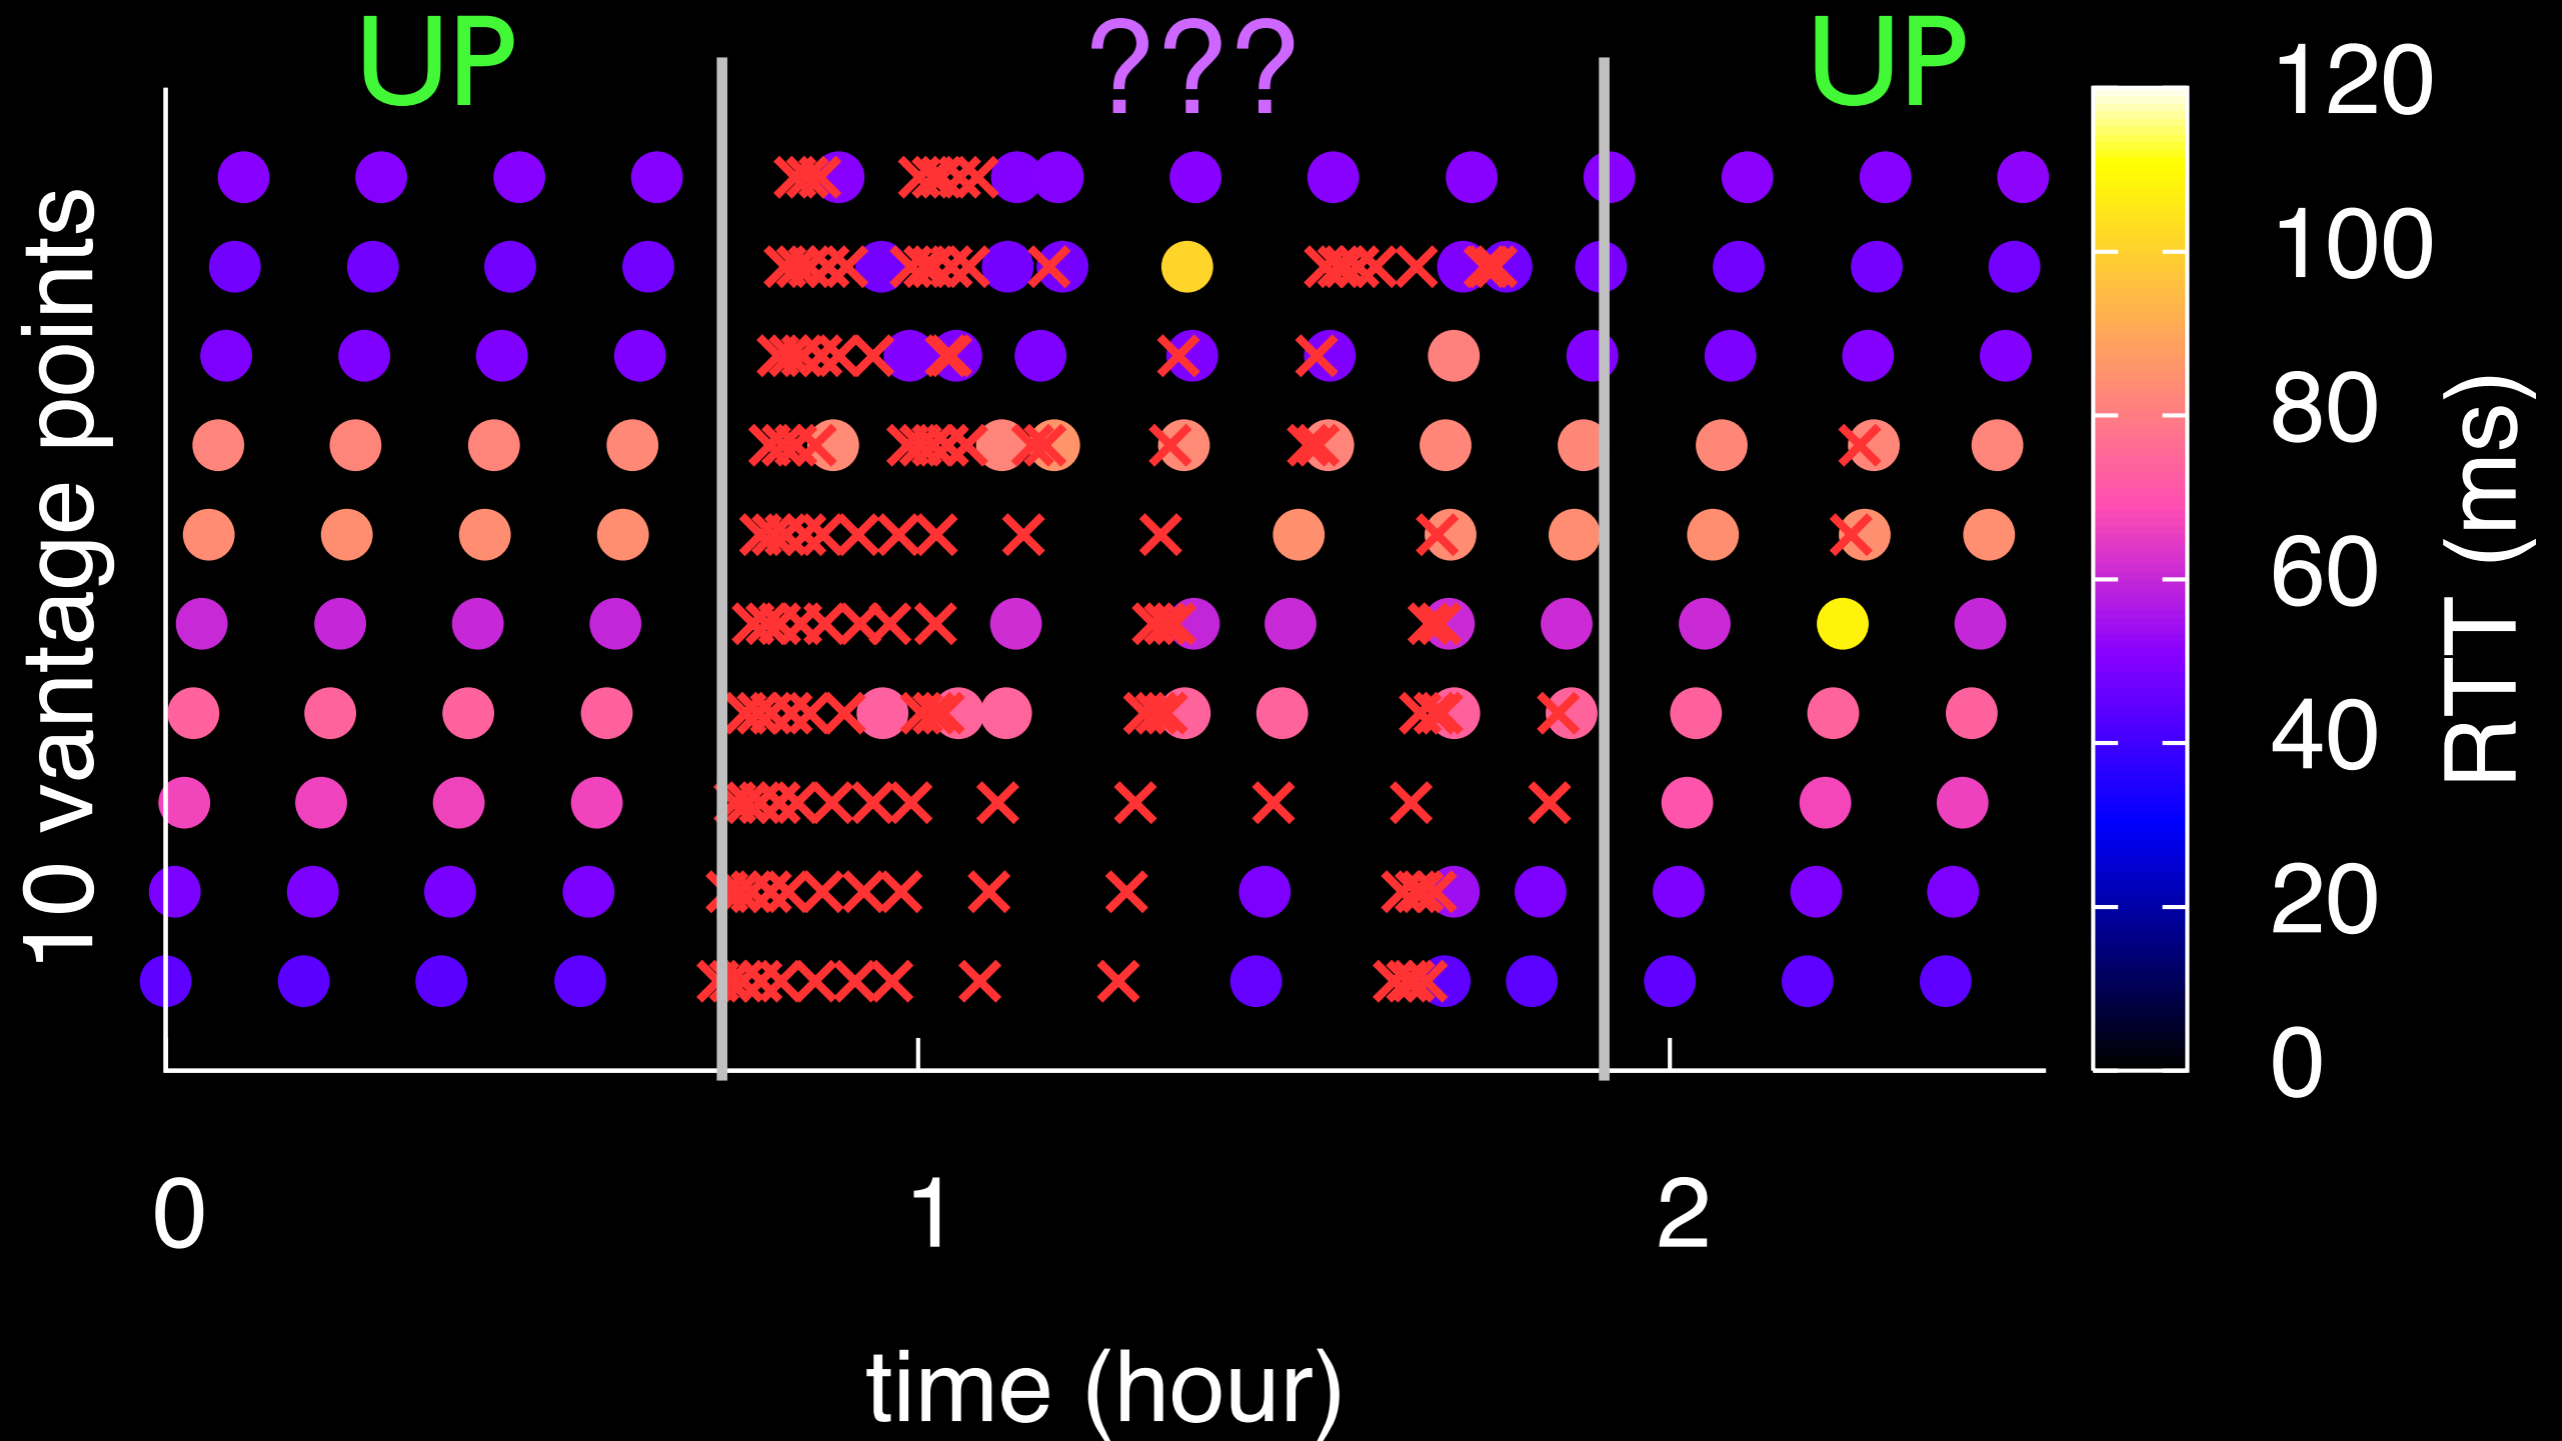
	%UP	#HOSTS
CT	99	1560
DC	99	823
DE	100	966
GA	99	1409
IN	99	123
KY	99	3349
MA	99	4234
MD	99	4955
ME	99	1016
MI	99	3801
NC	99	5712
NH	99	990
NJ	99	6066
NY	98	8438
OH	99	3483
PA	99	11165
RI	99	985
SC	98	884
TN	99	3475
VA	99	14063
VT	99	443
WV	99	1861



Failures in weather



Some lost pings \Rightarrow ??



Two Questions

- Could high delay create false outages?
- Could renumbering cause false outages and alter their duration?

When should pings time out?

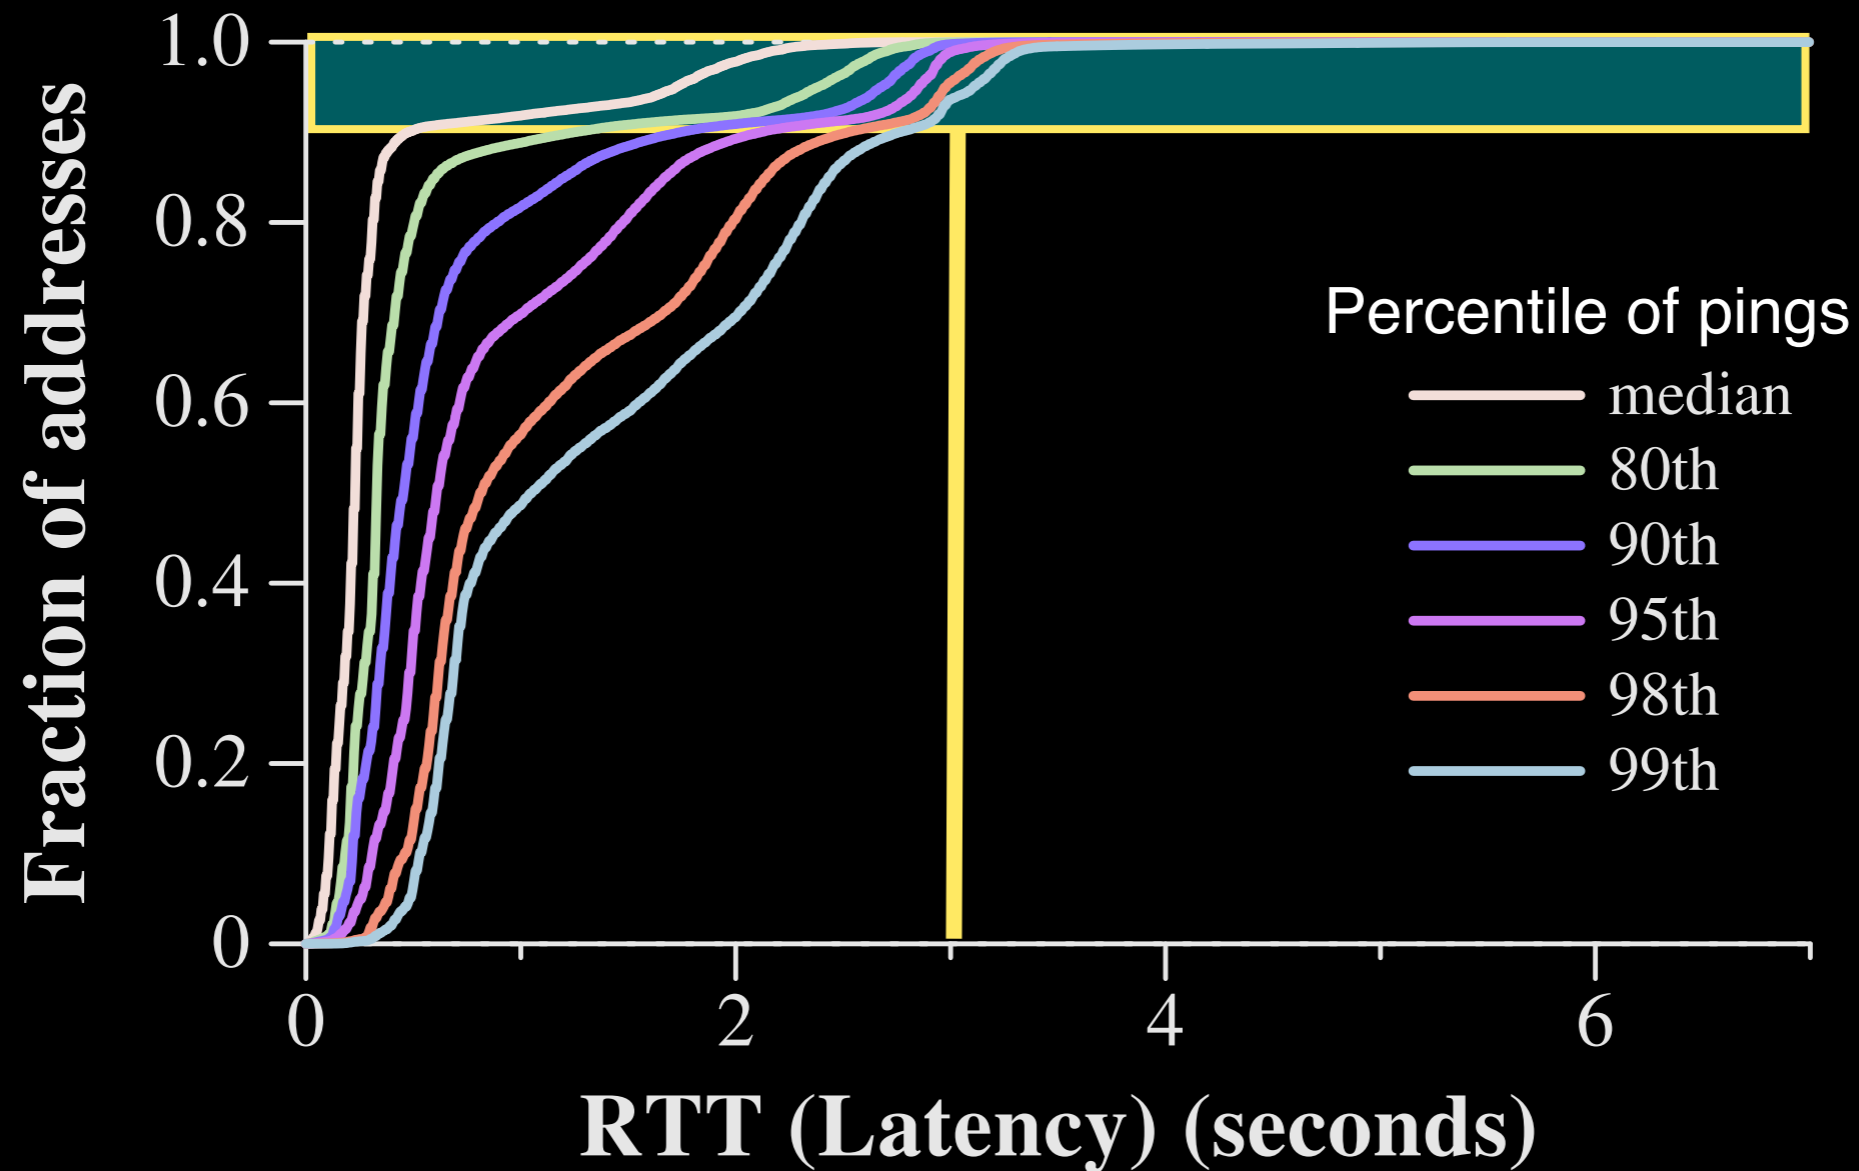
When should pings time out?

Measurement platform	Timeout (seconds)
RIPE Atlas	1
Scamper	2 (configurable)
Hubble / iPlane	2 (one retry)
SamKnows	3
Scriptroute / Thunderping	3 (configurable)
ISI survey	3 (collects all)

Let's confirm ~3s!

- Dataset: ISI survey data: 1% of routed /24's, pinged every 11 minutes.
 - Precise timing below 3s timeout.
 - Imprecise timing above 3s timeout. **Any received echo reply** is logged with time and source.
- Approach: Look at all response times, including those longer than the timeout.

Survey-detected RTTs



About 10% of addresses routinely respond after one second.
The distribution appears clipped by the 3s limit.

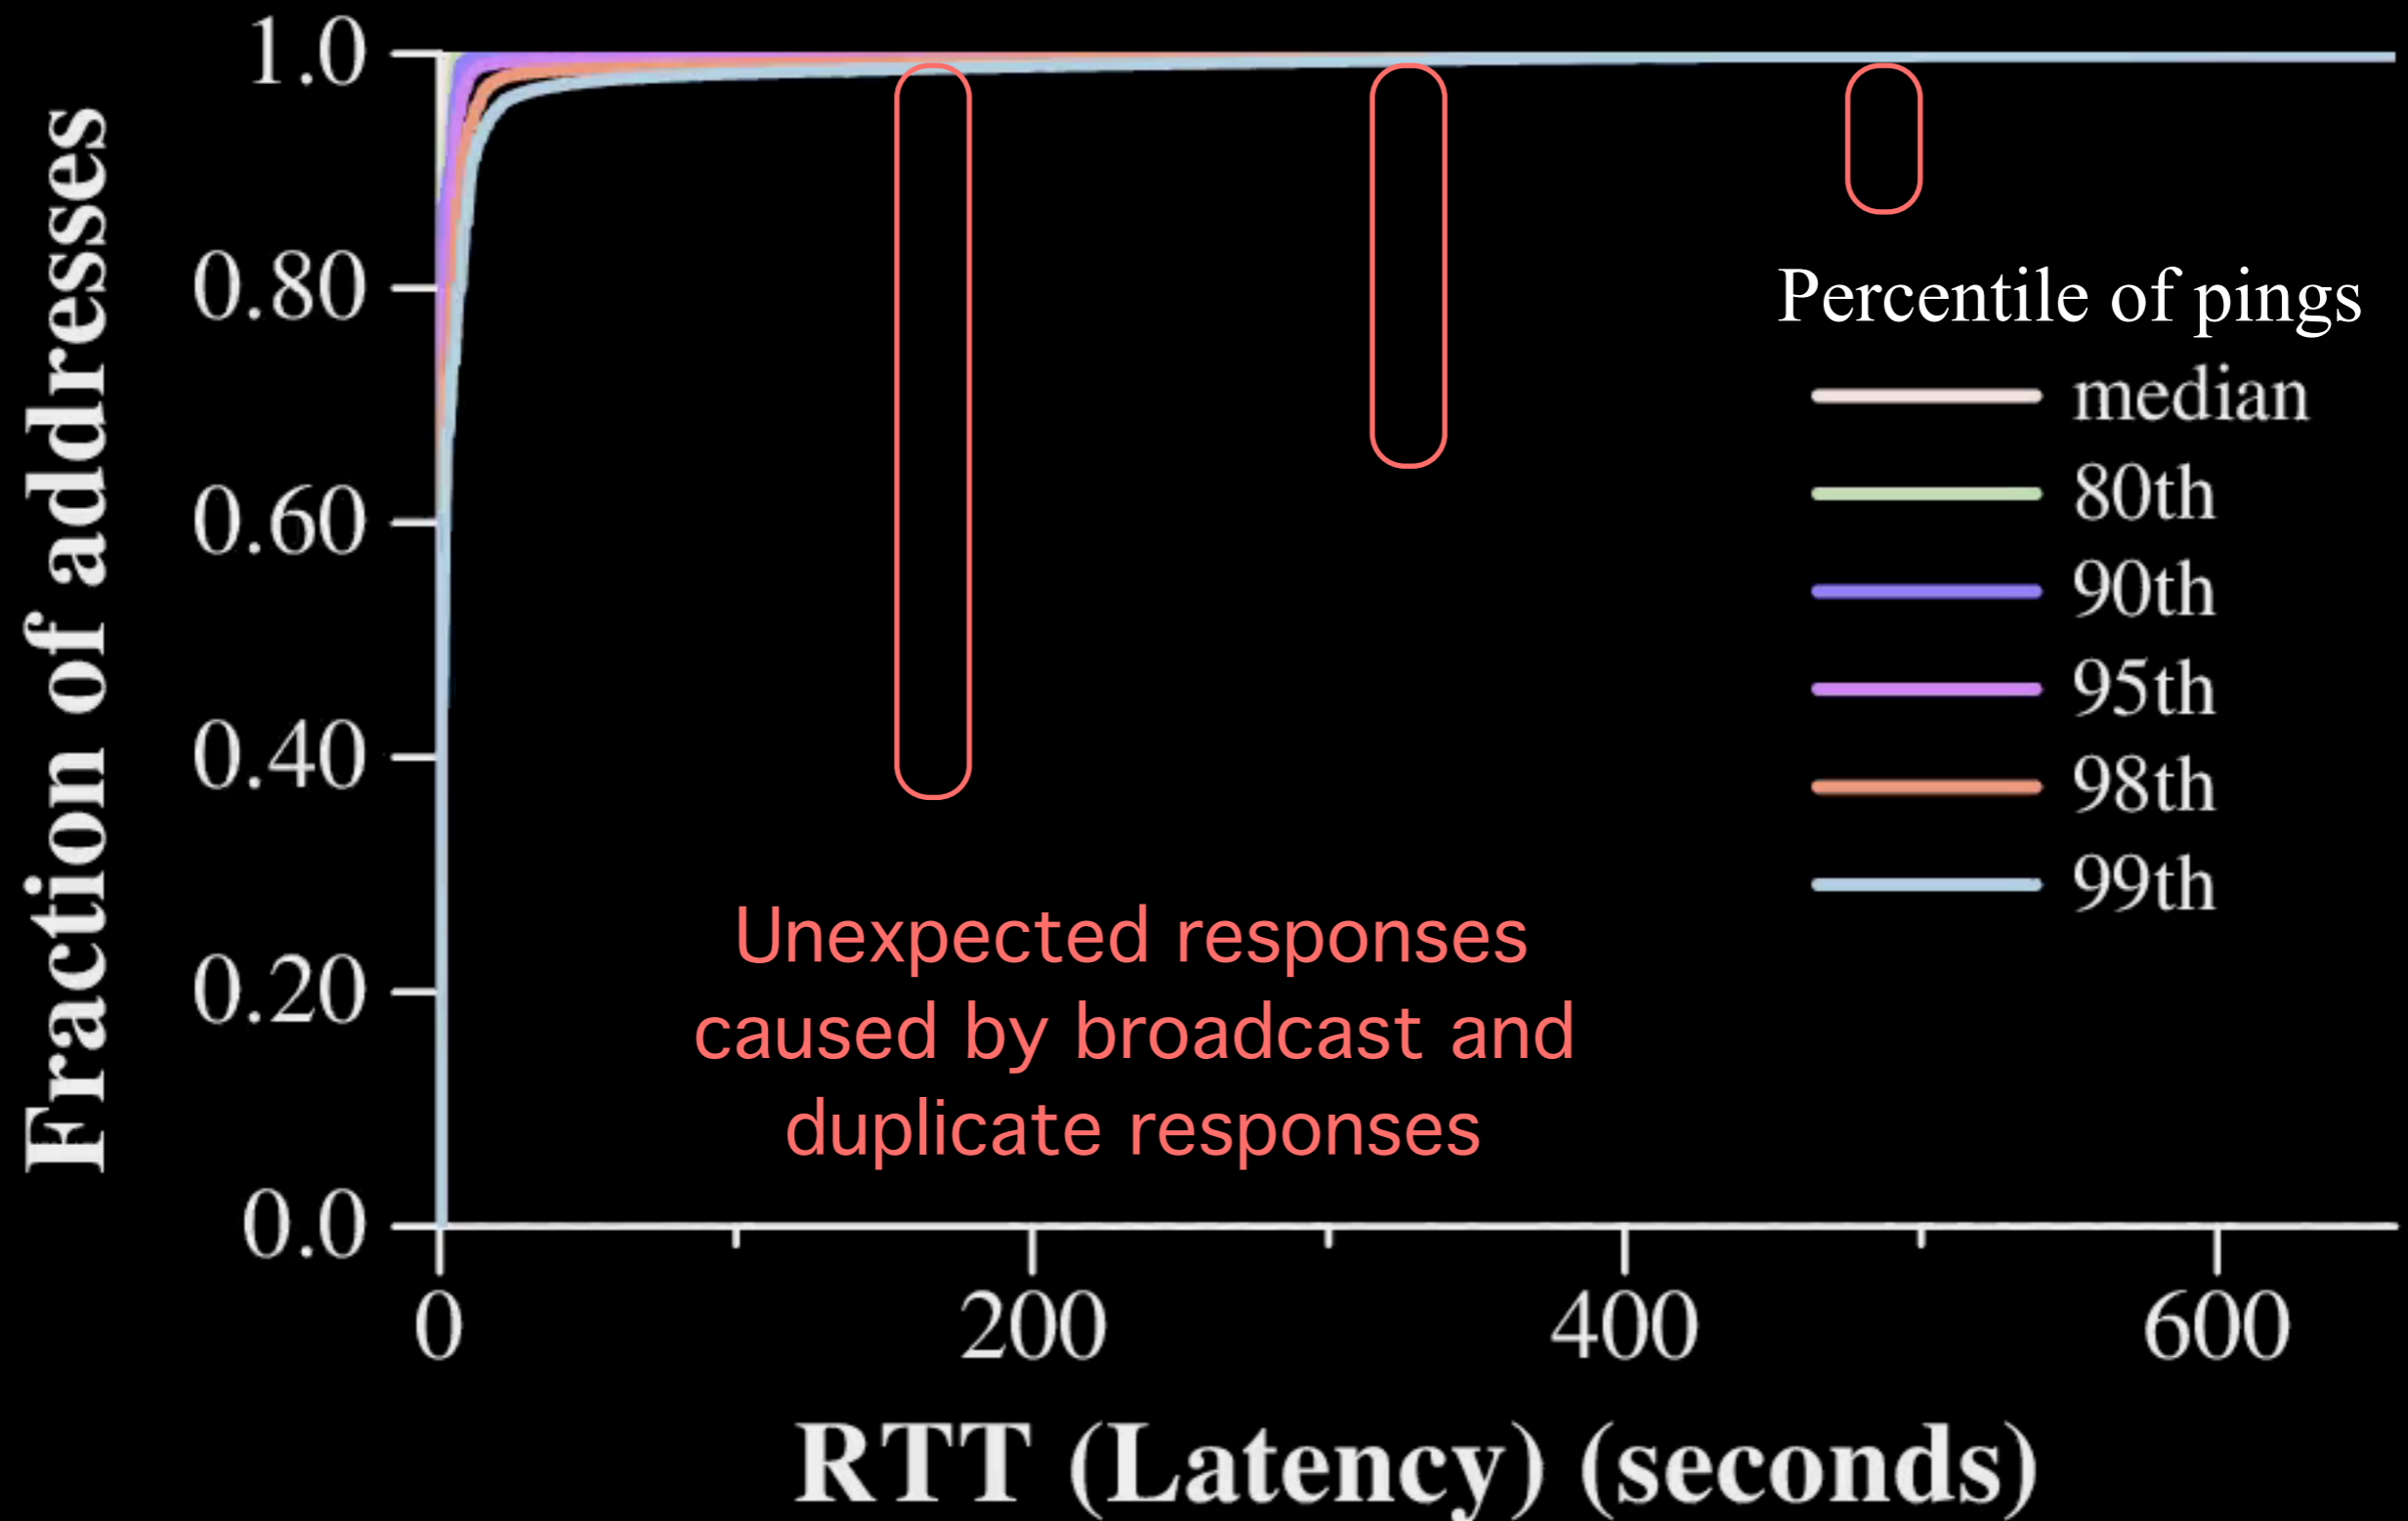
Transform Survey Data

Time		Probe Destination	Reply Source	RTT	
[1320291701.0]	P	v119 1.99.16.242	1.99.16.242	2960.995	45
[1320292364.0]	P	v119 1.99.16.242	1.99.16.242	2767.092	45
[1320293027.0]	P	v119 1.99.16.242	error_time_out		
[1320293031.0]	P	v119 no_probe_ip	1.99.16.242	0.000	45 [d004]
[1320293691.0]	P	v119 1.99.16.242	error_time_out		
[1320293696.0]	P	v119 no_probe_ip	1.99.16.242	0.000	45 [d005]
[1320294354.0]	P	v119 1.99.16.242	error_time_out		
[1320294358.0]	P	v119 no_probe_ip	1.99.16.242	0.000	45 [d004]
[1320295017.0]	P	v119 1.99.16.242	error_time_out		
[1320295030.0]	P	v119 no_probe_ip	1.99.16.242	0.000	45 [d013]

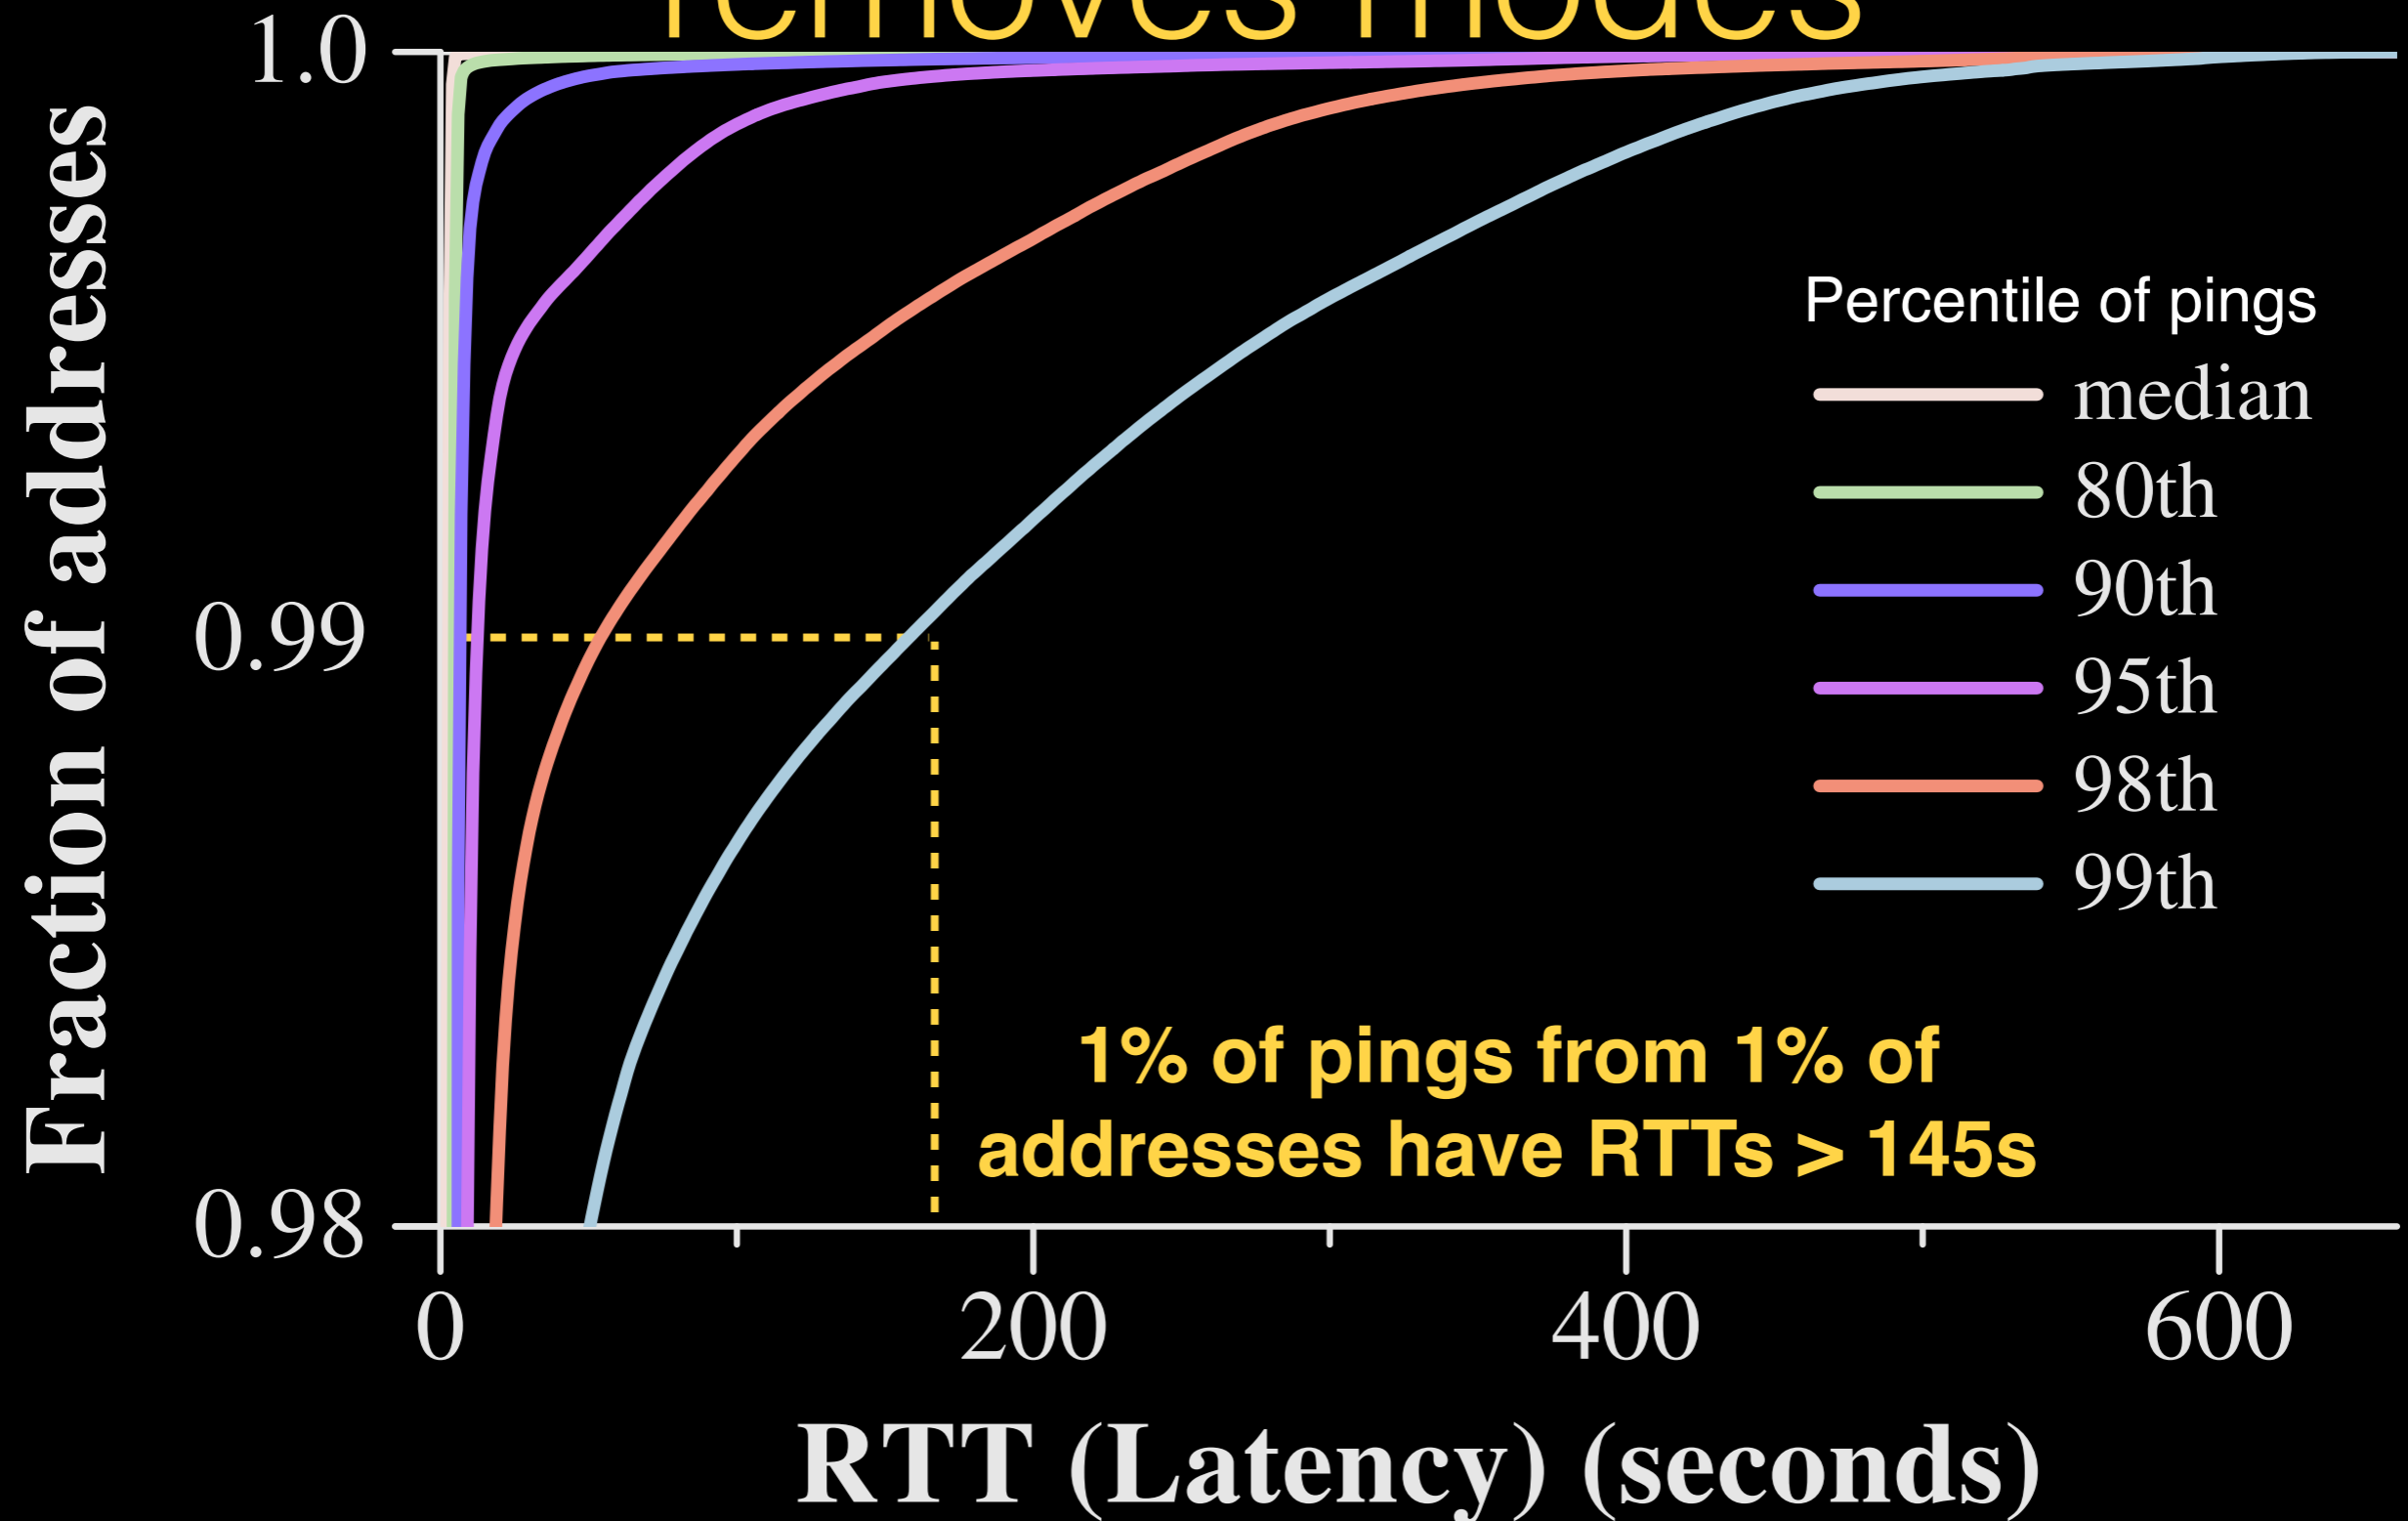


[1320291701.0]	P	v119 1.99.16.242	1.99.16.242	2960.995	45
[1320292364.0]	P	v119 1.99.16.242	1.99.16.242	2767.092	45
[1320293027.0]	P	v119 1.99.16.242	1.99.16.242	4000.0000	45
[1320293691.0]	P	v119 1.99.16.242	1.99.16.242	5000.0000	45
[1320294354.0]	P	v119 1.99.16.242	1.99.16.242	4000.0000	45
[1320295017.0]	P	v119 1.99.16.242	1.99.16.242	13000.0000	45

Absurdly long RTTs



Filtering broadcast responses removes modes



When should probes time out?

		% of pings						
		1%	50%	80%	90%	95%	98%	99%
% of addresses	1%	0.01	0.03	0.04	0.07	0.10	0.13	0.18
	50%	0.16	0.19	0.21	0.26	0.42	0.53	0.64
	80%	0.19	0.26	0.33	0.43	0.54	0.74	1.21
	90%	0.22	0.31	0.42	0.57	0.84	1.31	3
	95%	0.25	1.42	2.38	3	5	9	15
	98%	0.30	1.94	4	6	12	41	78
	99%	0.33	2.31	4	8	22	76	145

350ms timeout misses 50% of pings from 5% of addrs

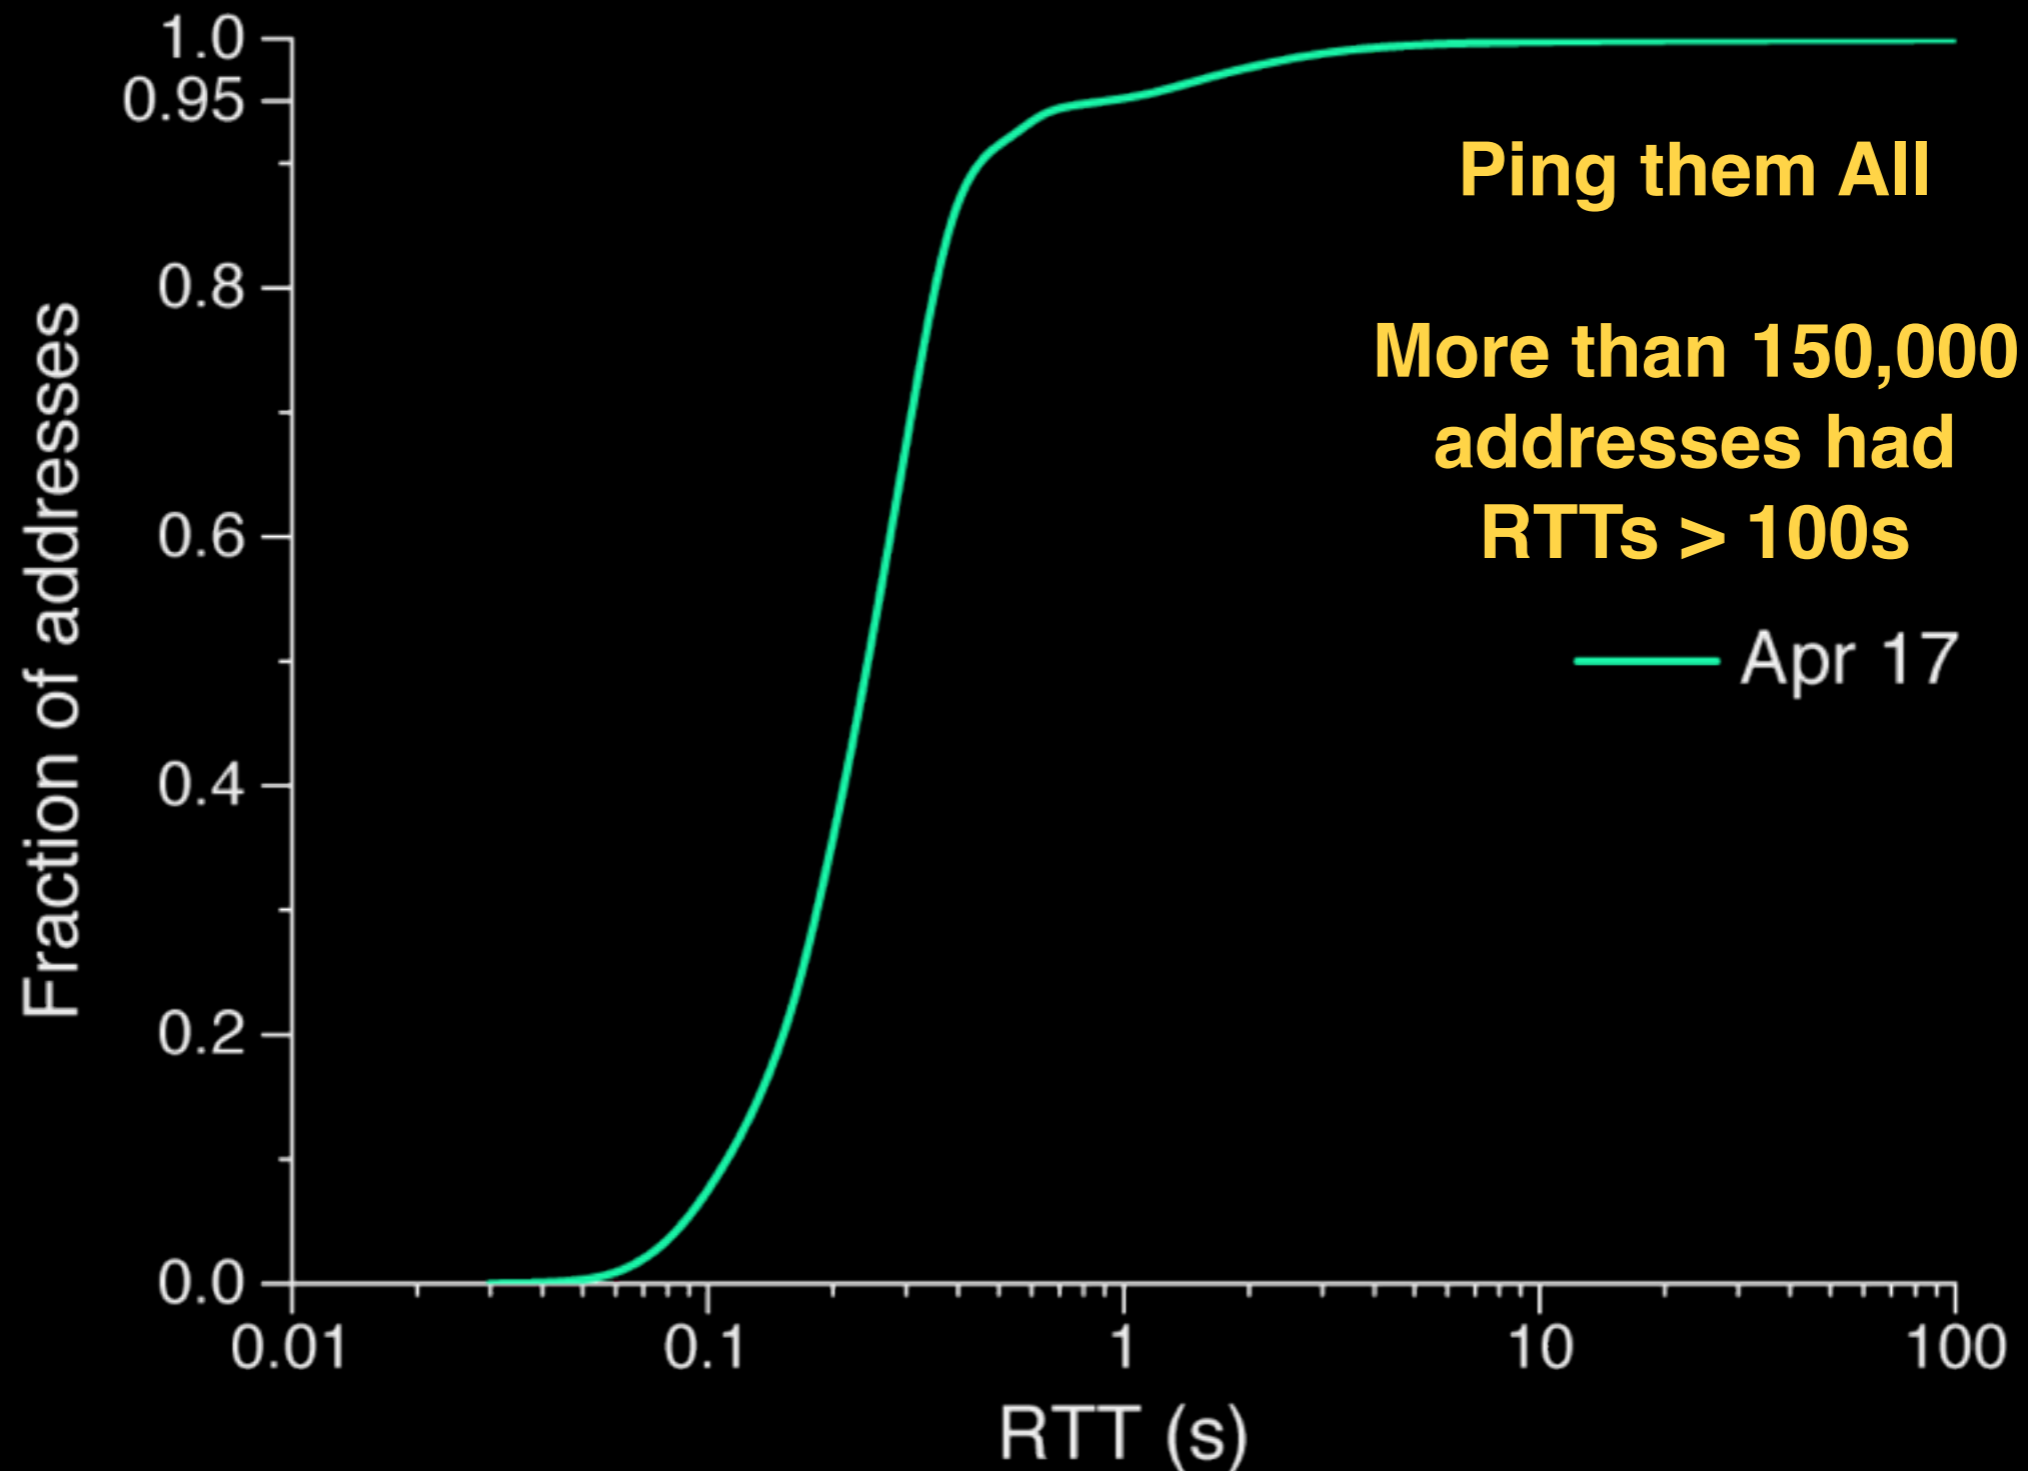
Most addresses can respond within 350ms

99% of pings from 99% of addrs have RTTs < 145s

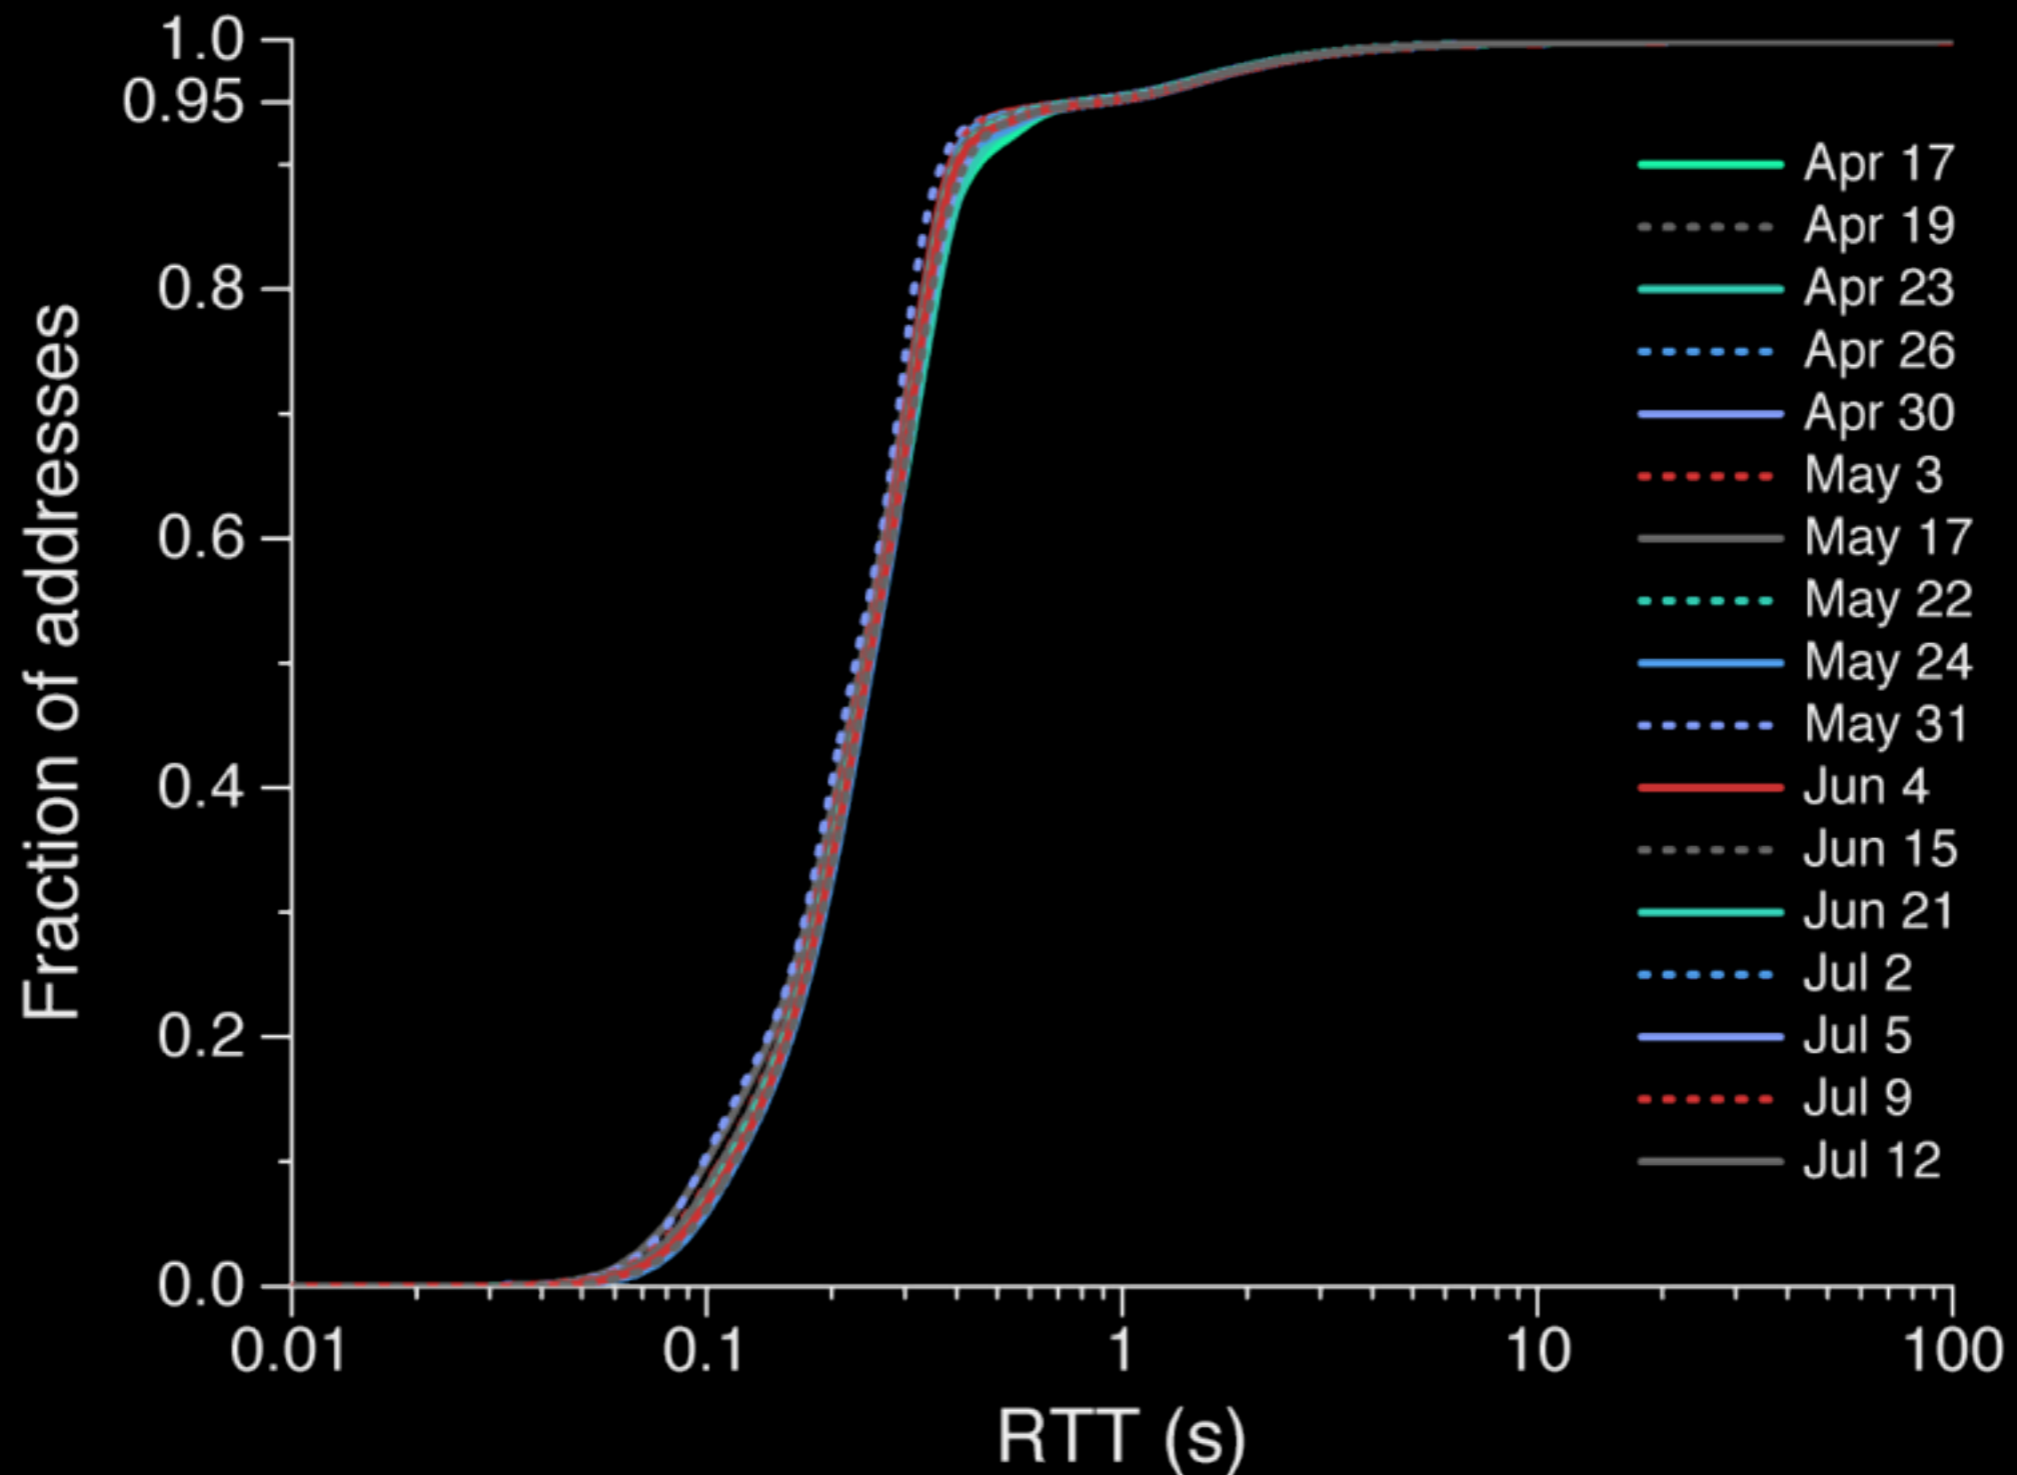
Inconceivable!

- Is it an unrepresentative sample?
- Is it temporary?
- Is it just ICMP (the protocol used by ping)?
- Is this new?
- What addresses take so long to respond?

Did we sample bad addresses?



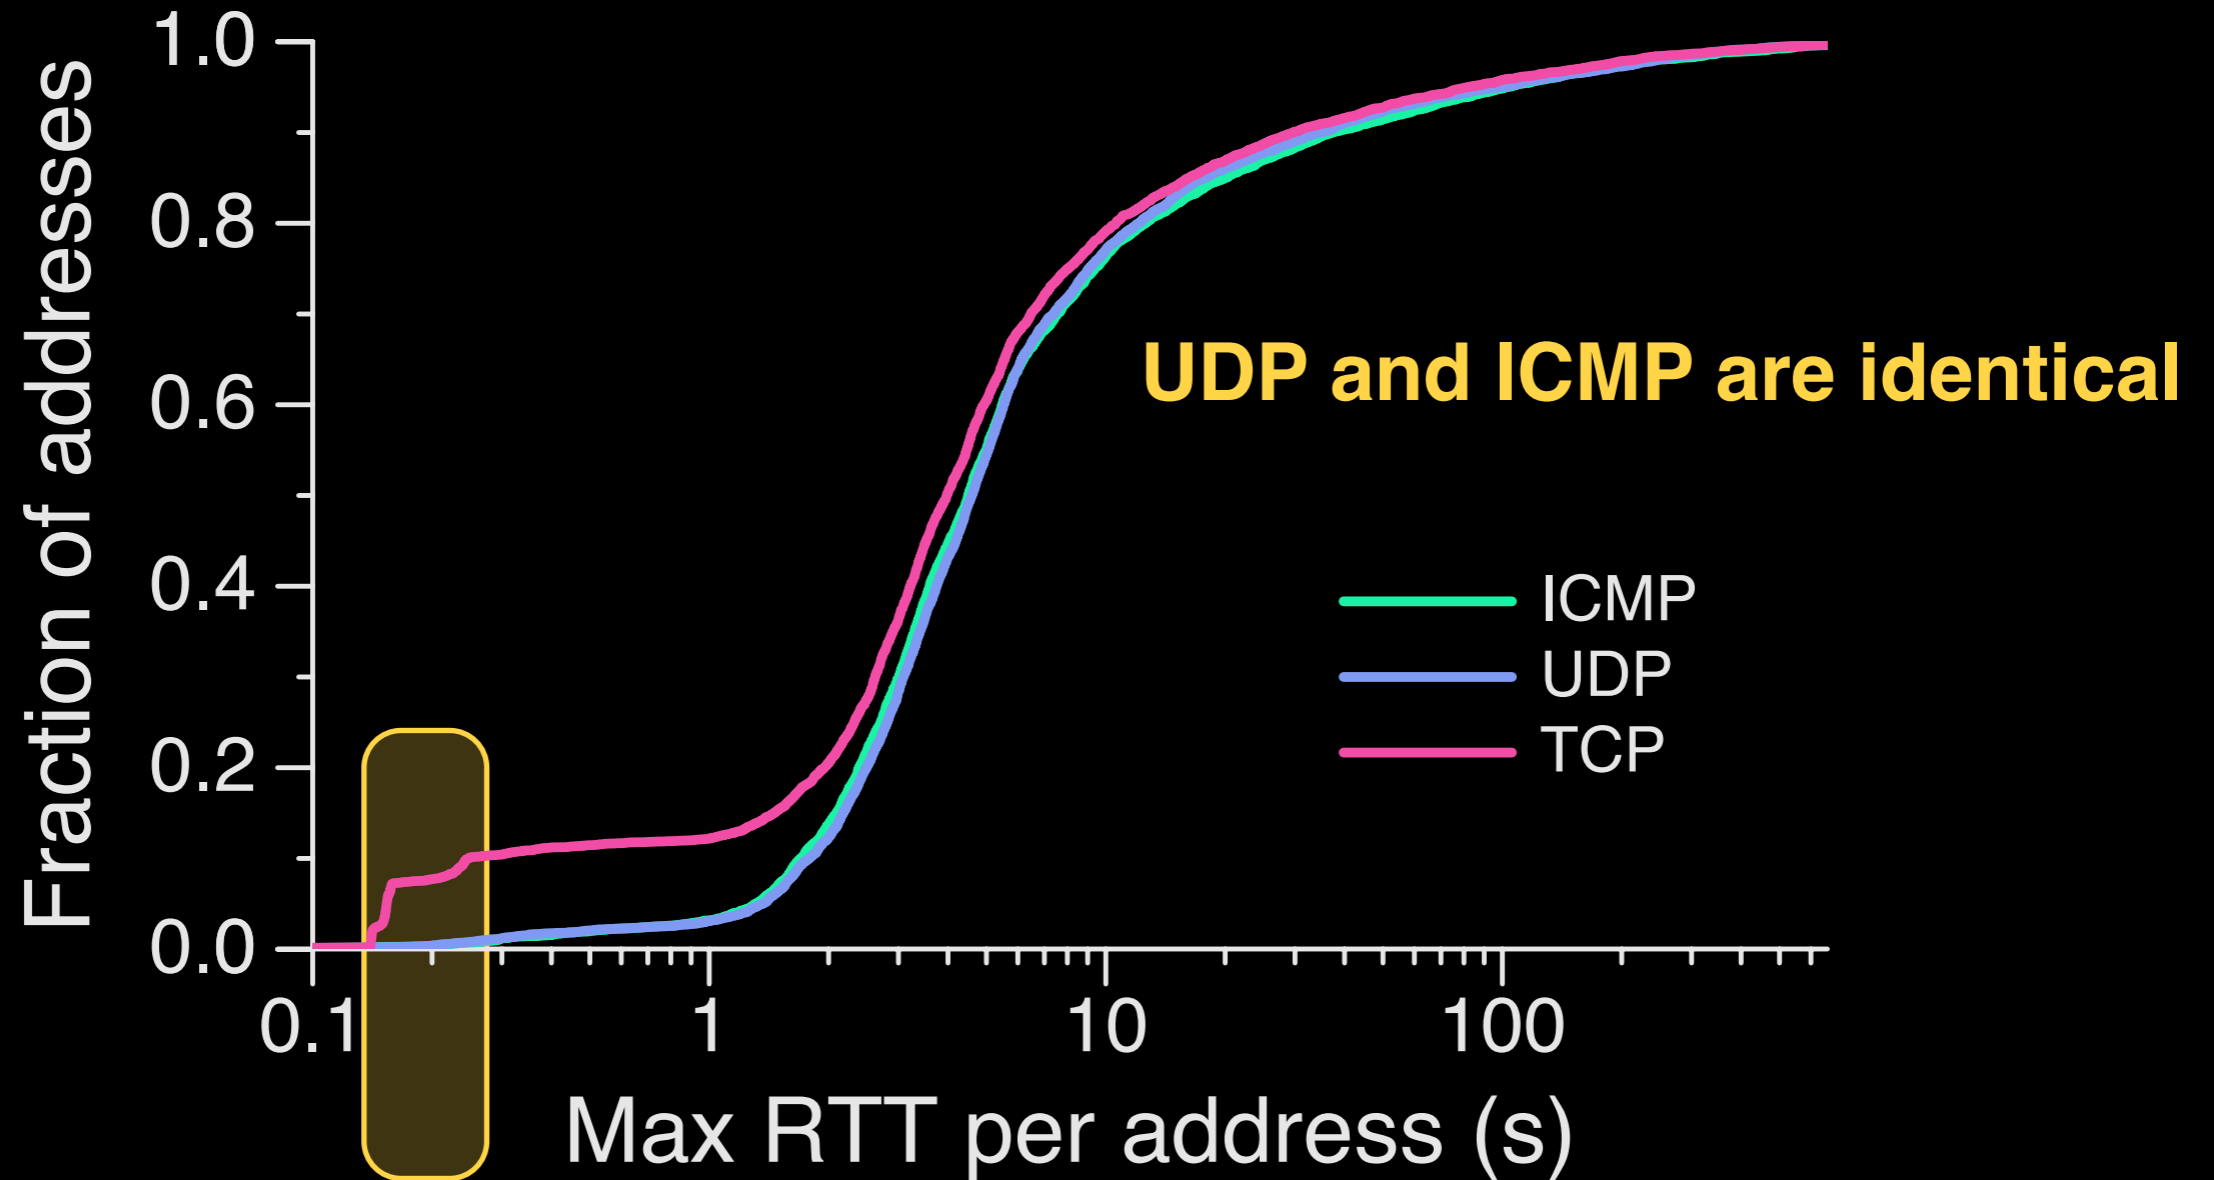
Is it temporary?



Is it just ICMP?

- Use Scamper to send TCP, UDP, ICMP probes to **high-latency** addresses
 - “**high-latency**”: ~5K addresses from ISI 2015 surveys whose 50th, 80th, 90th or 95th percentile RTTs are in the top 5%
- Sent ICMP, UDP, TCP packets 20 mins apart, for 36 hours

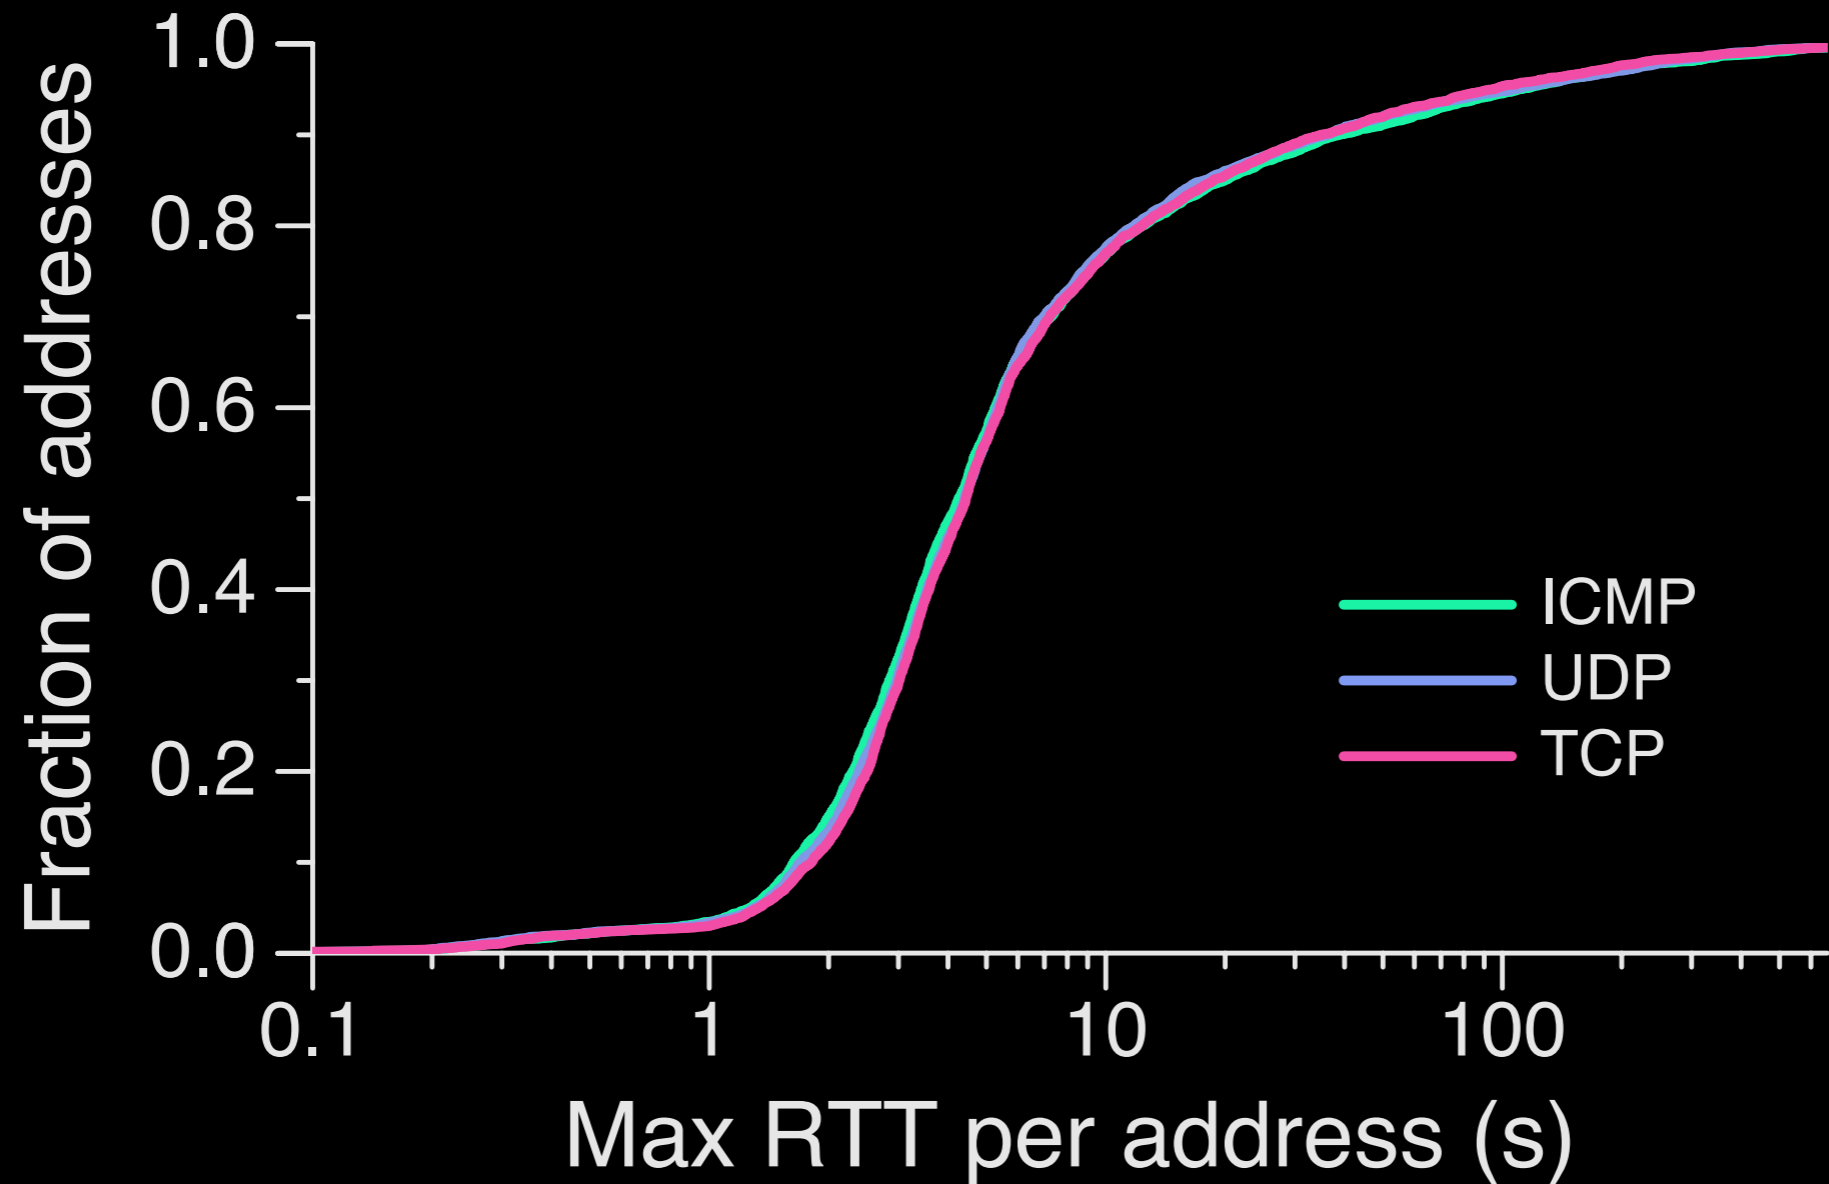
Is it just ICMP?



Mode for TCP


Likely caused by firewall

Is it just ICMP?

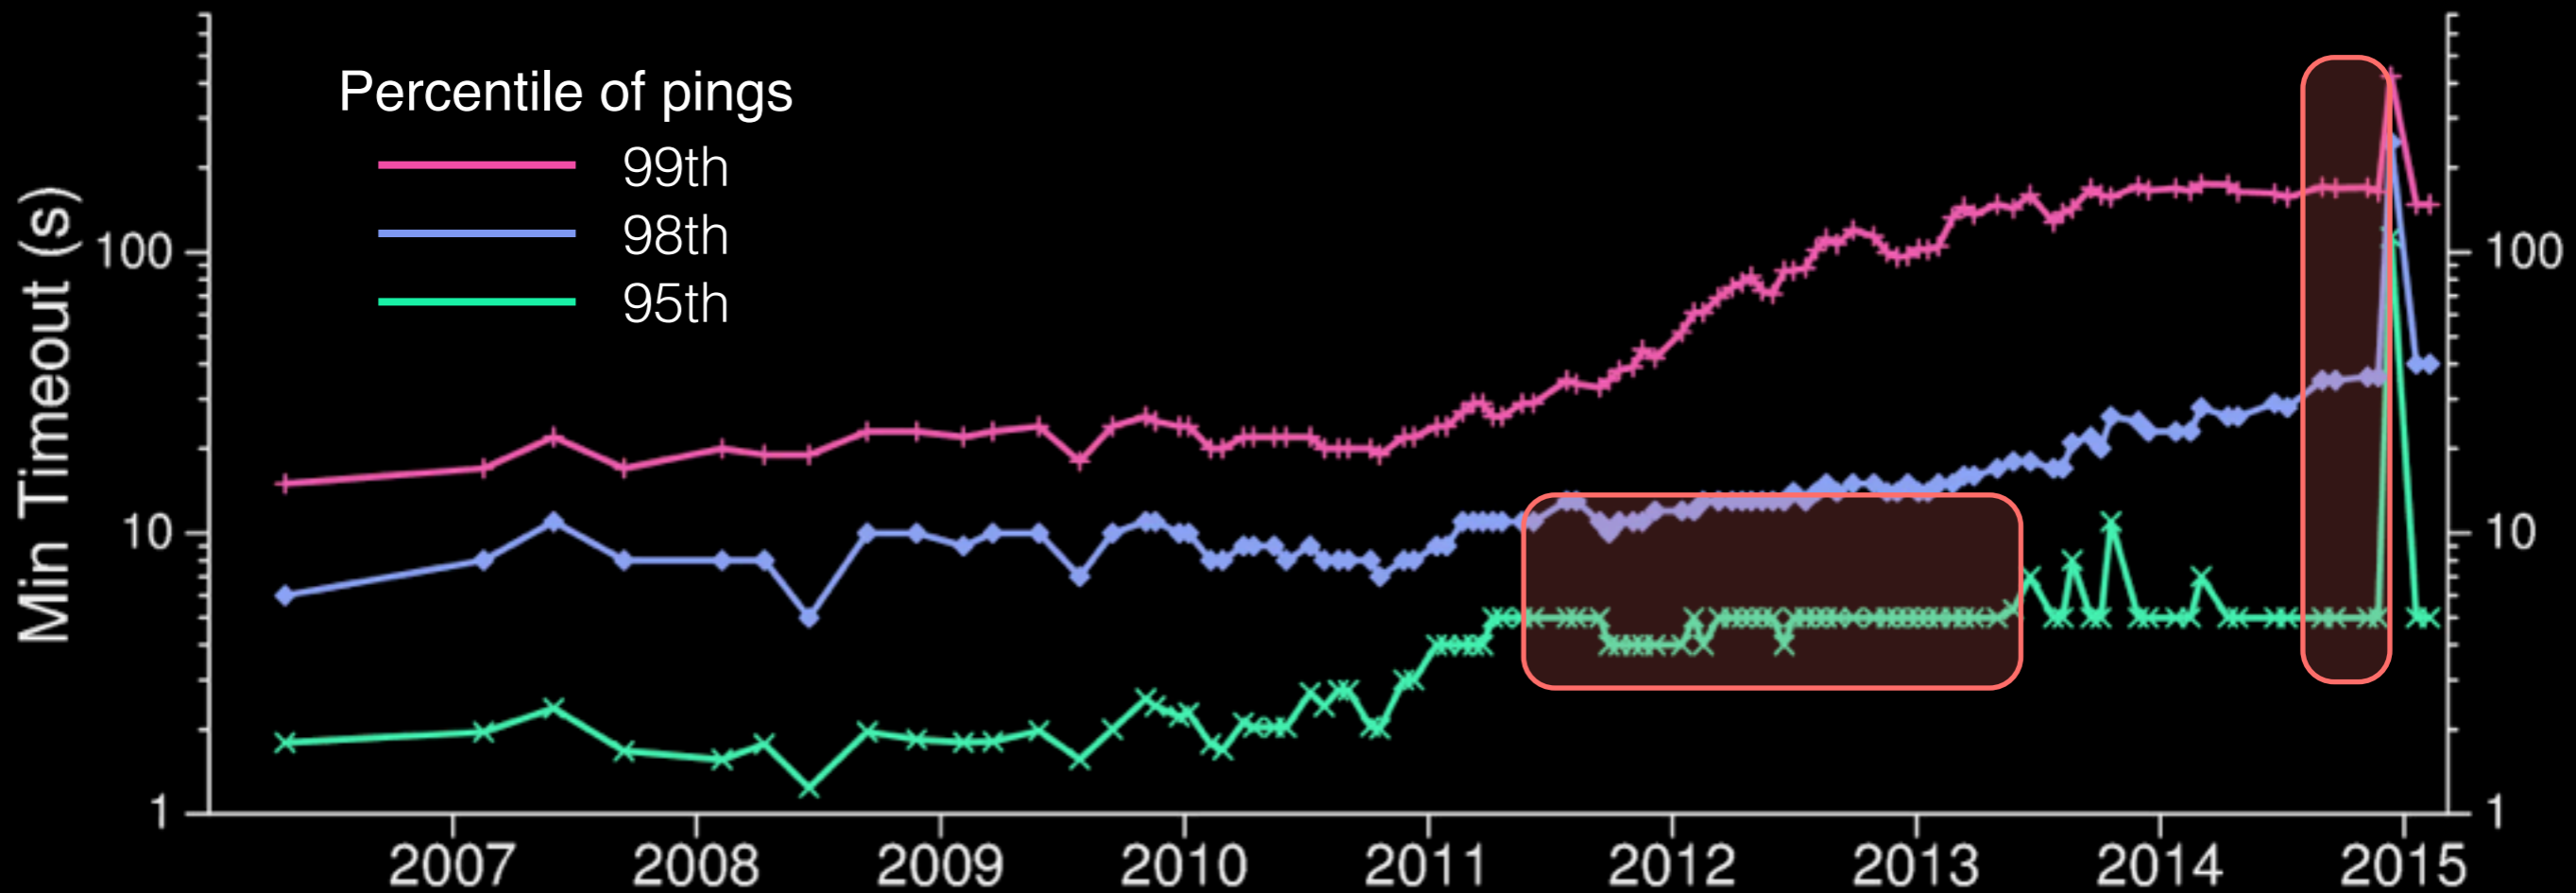


Removed ~500 addresses belonging to firewalling AS

Is this new?



Is this new?



What addresses take so long?

1/2: Where?

Continent	July 2015 high RTT addresses	
	Number	% per continent
South America	8.05M	26.9
Asia	4.56M	3.2
Europe	2.32M	2.4
Africa	1.30M	31.7
North America	1.14M	1.2
Oceania	0.08M	3.7

What addresses take so long?

2/2: Which providers?

Autonomous System	July 2015 high RTT addresses	
	Number	% per AS
Telefonica Brasil	4.20M	77
Tim Celular S.A.	1.72M	71.6
Bharti Airtel Ltd.	1.03M	79.2
Cellco Partnership	0.63M	72.7
Tele2	0.58M	67.4

**All
cellular**

**Majority of
responsive
addresses**

Lessons

- Pings reach cell phones; may use power, expose activity.
- Duration of buffering across disconnection is extraordinary, violates TTL and MSL.
- Long timeouts necessary to disambiguate outages from disconnection.

Two Questions

- Could high delay create false outages?
- Could renumbering cause false outages and alter their duration?

What's Renumbering


- “Dynamic” addresses may change because:
 - The administrator needs to reassign devices to networks
 - A long outage allows the network to forget
 - A rebooted machine gets a new address
 - The provider limits the lifetime of addresses

Data: RIPE Atlas Probes

- Logs show when these devices:
 - Get a new address
 - Reboot
 - Lose connectivity

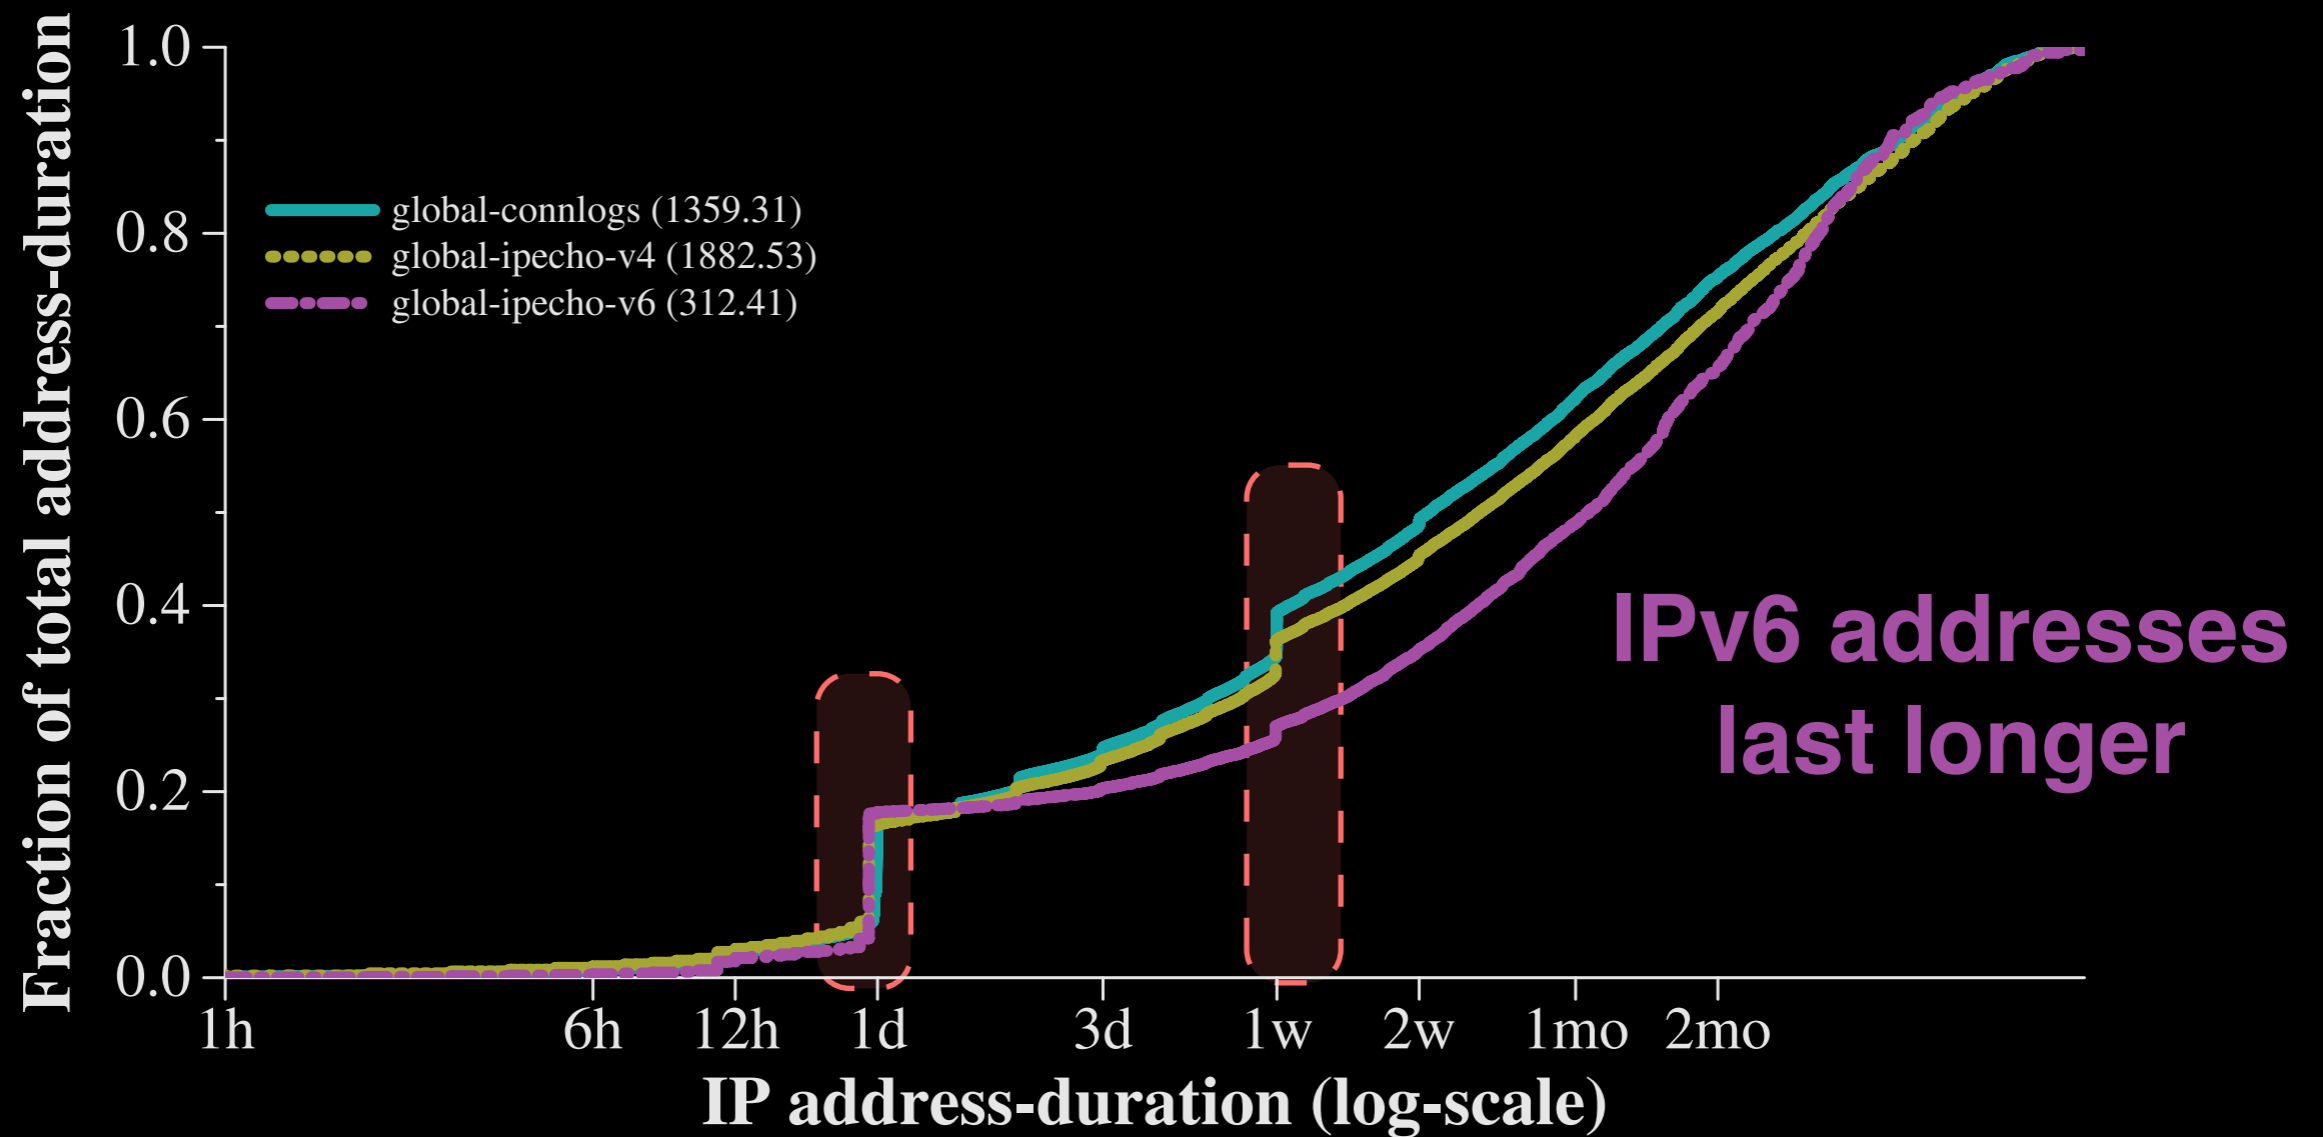


Weight address durations

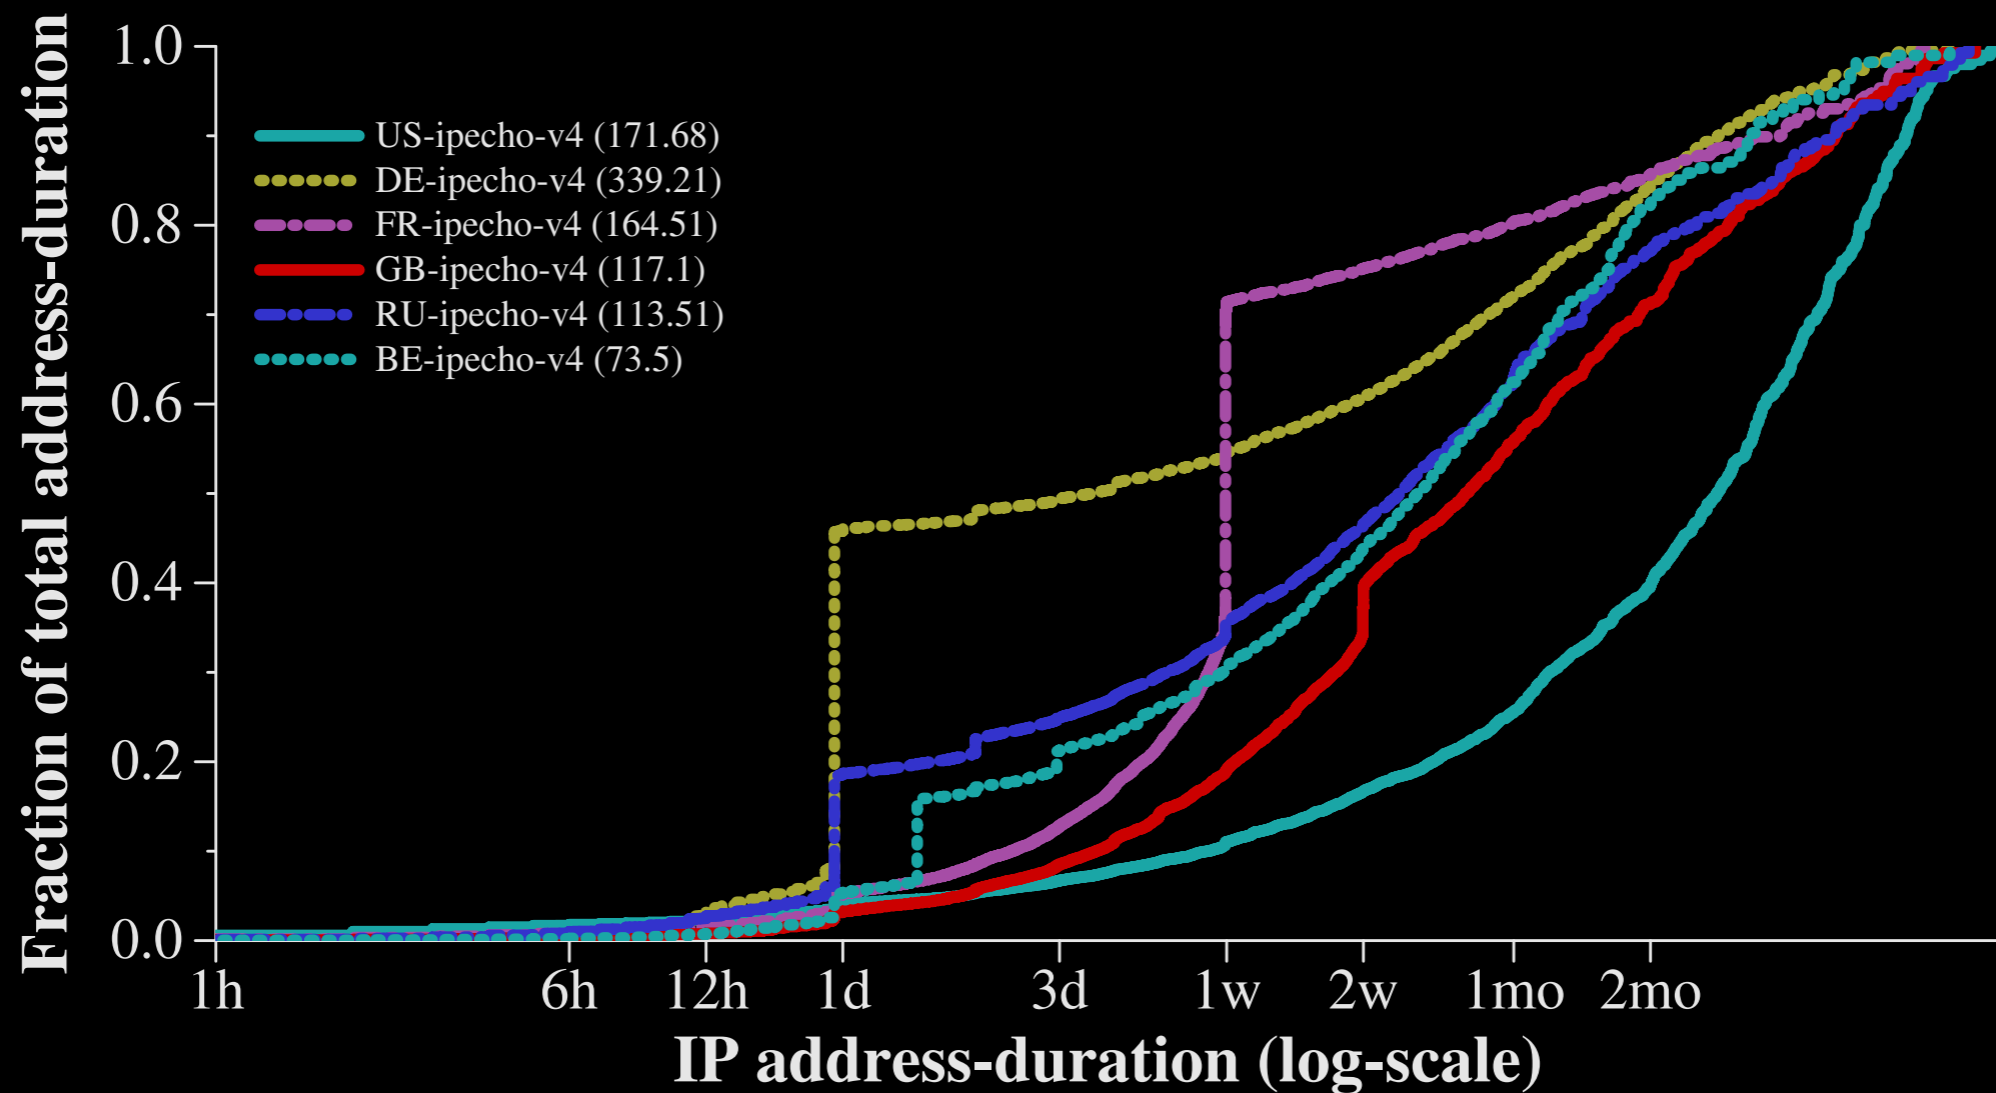


	Address	Duration
IP ₁	79.194.205.144	NA
IP ₂	79.194.192.169	24
IP ₃	79.194.196.241	24
IP ₄	79.194.194.4	12
IP ₅	91.9.219.235	NA
		Sum: 60


Addresses often last days



Periodic address durations are common in Germany and France



Cable seems stable.




Could renumbering cause false outages?

- We don't see periodic renumbering in the US, so, unlikely here.
- Where there is periodic renumbering, can account for it.

Two Questions

- Could high delay create false outages?
- Could renumbering cause false outages and alter their duration?

Renumbering by outage duration from Atlas probes for one ISP





Now

- Building tools to identify hosts after address changes and outages
- Studying how a sample of address space can be representative
- Providing information to users about their own and adjacent networks

Remember

- When sending a packet into the Internet, you might see a response after minutes.
- When blacklisting an IP address for misbehavior, you might see the same machine at a different address in a few hours.

Great Students



Questions?