

Timeouts: Beware Surprisingly High Delay

Ramakrishna Padmanabhan, Patrick Owen,
Aaron Schulman, Neil Spring



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)
------------	------------------

**Can we use observed RTTs to tell
what the timeouts should be?**

When should probes 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**)

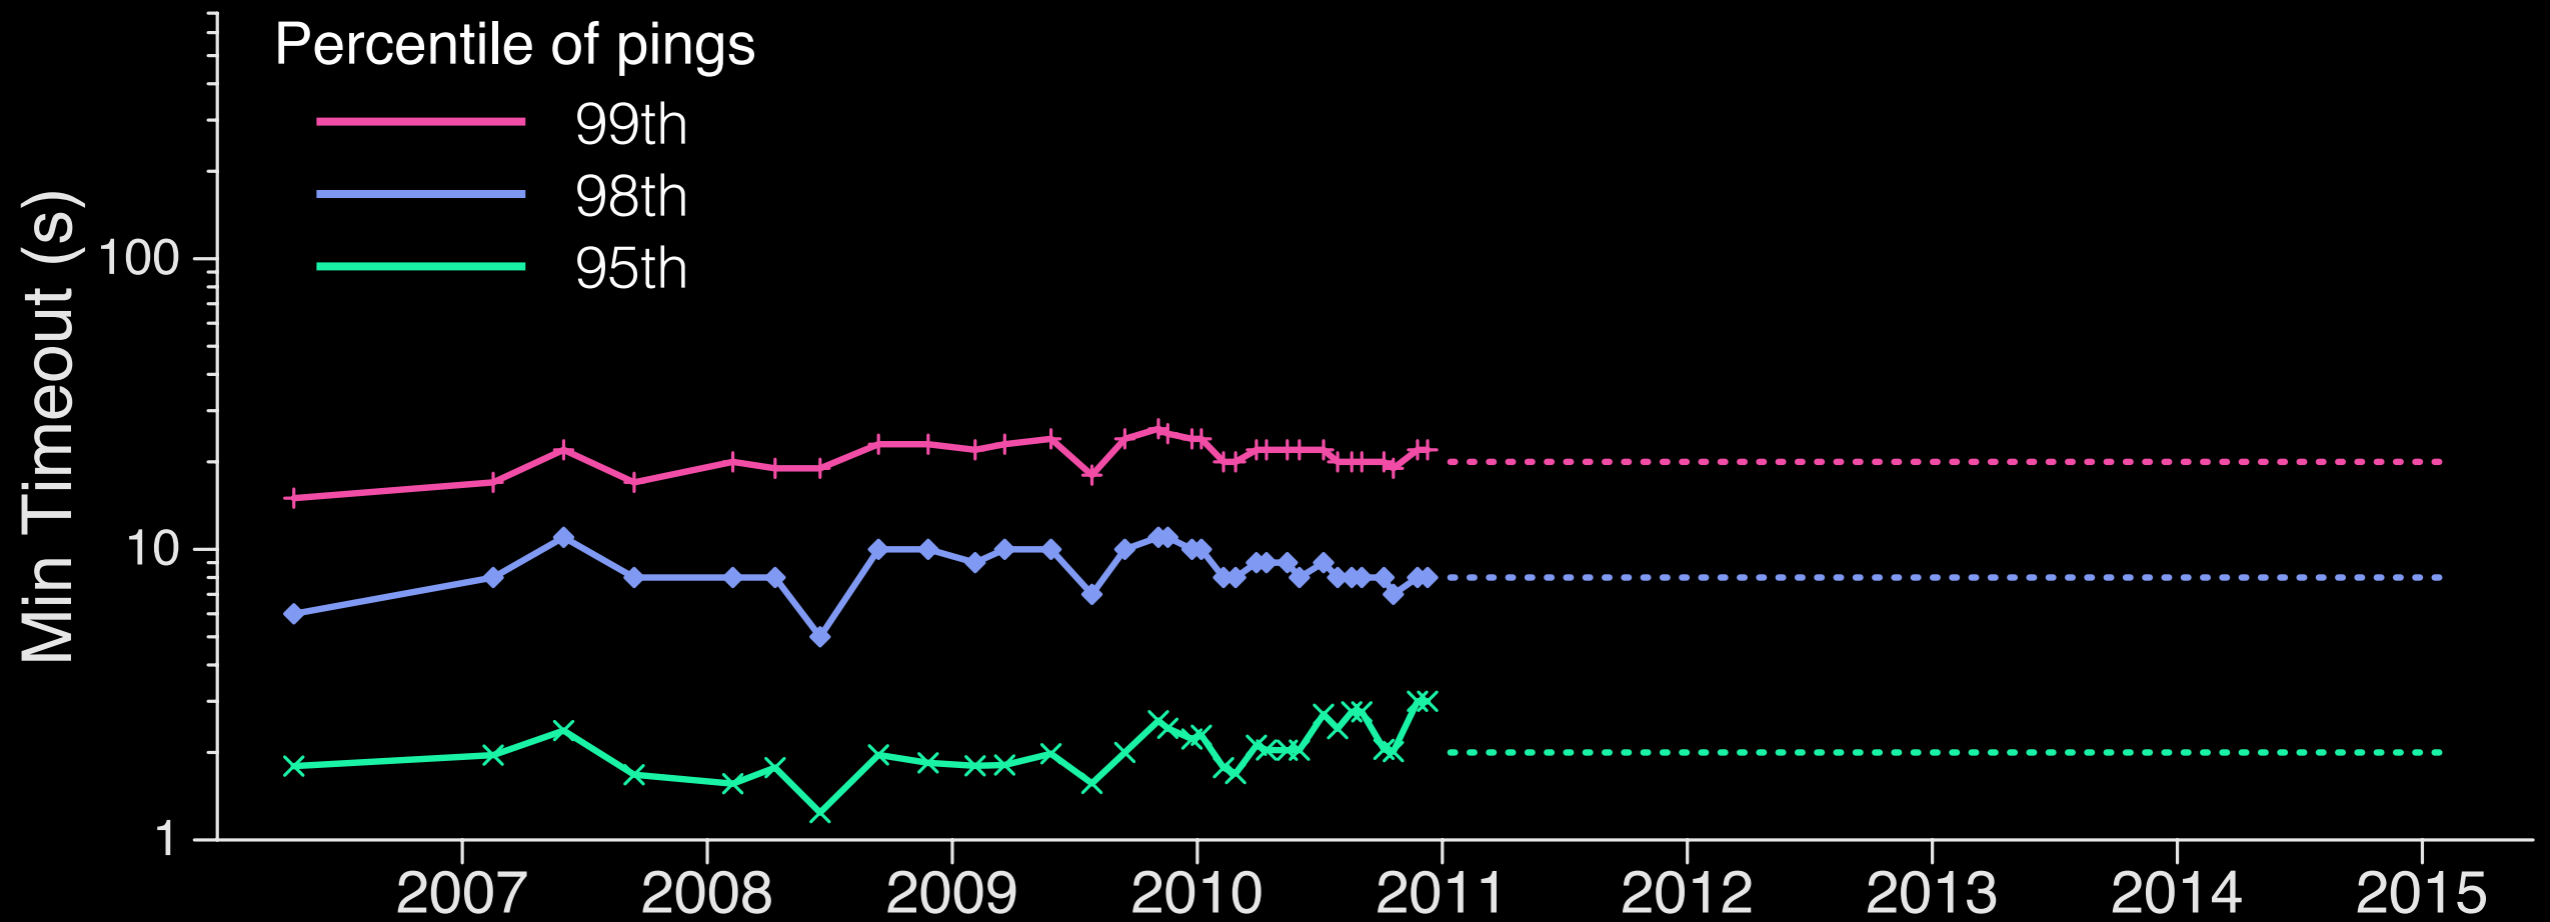
Why is this not widely known?

		% 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	1.61	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


RTTs until 2011 were stable



Expected RTT trends



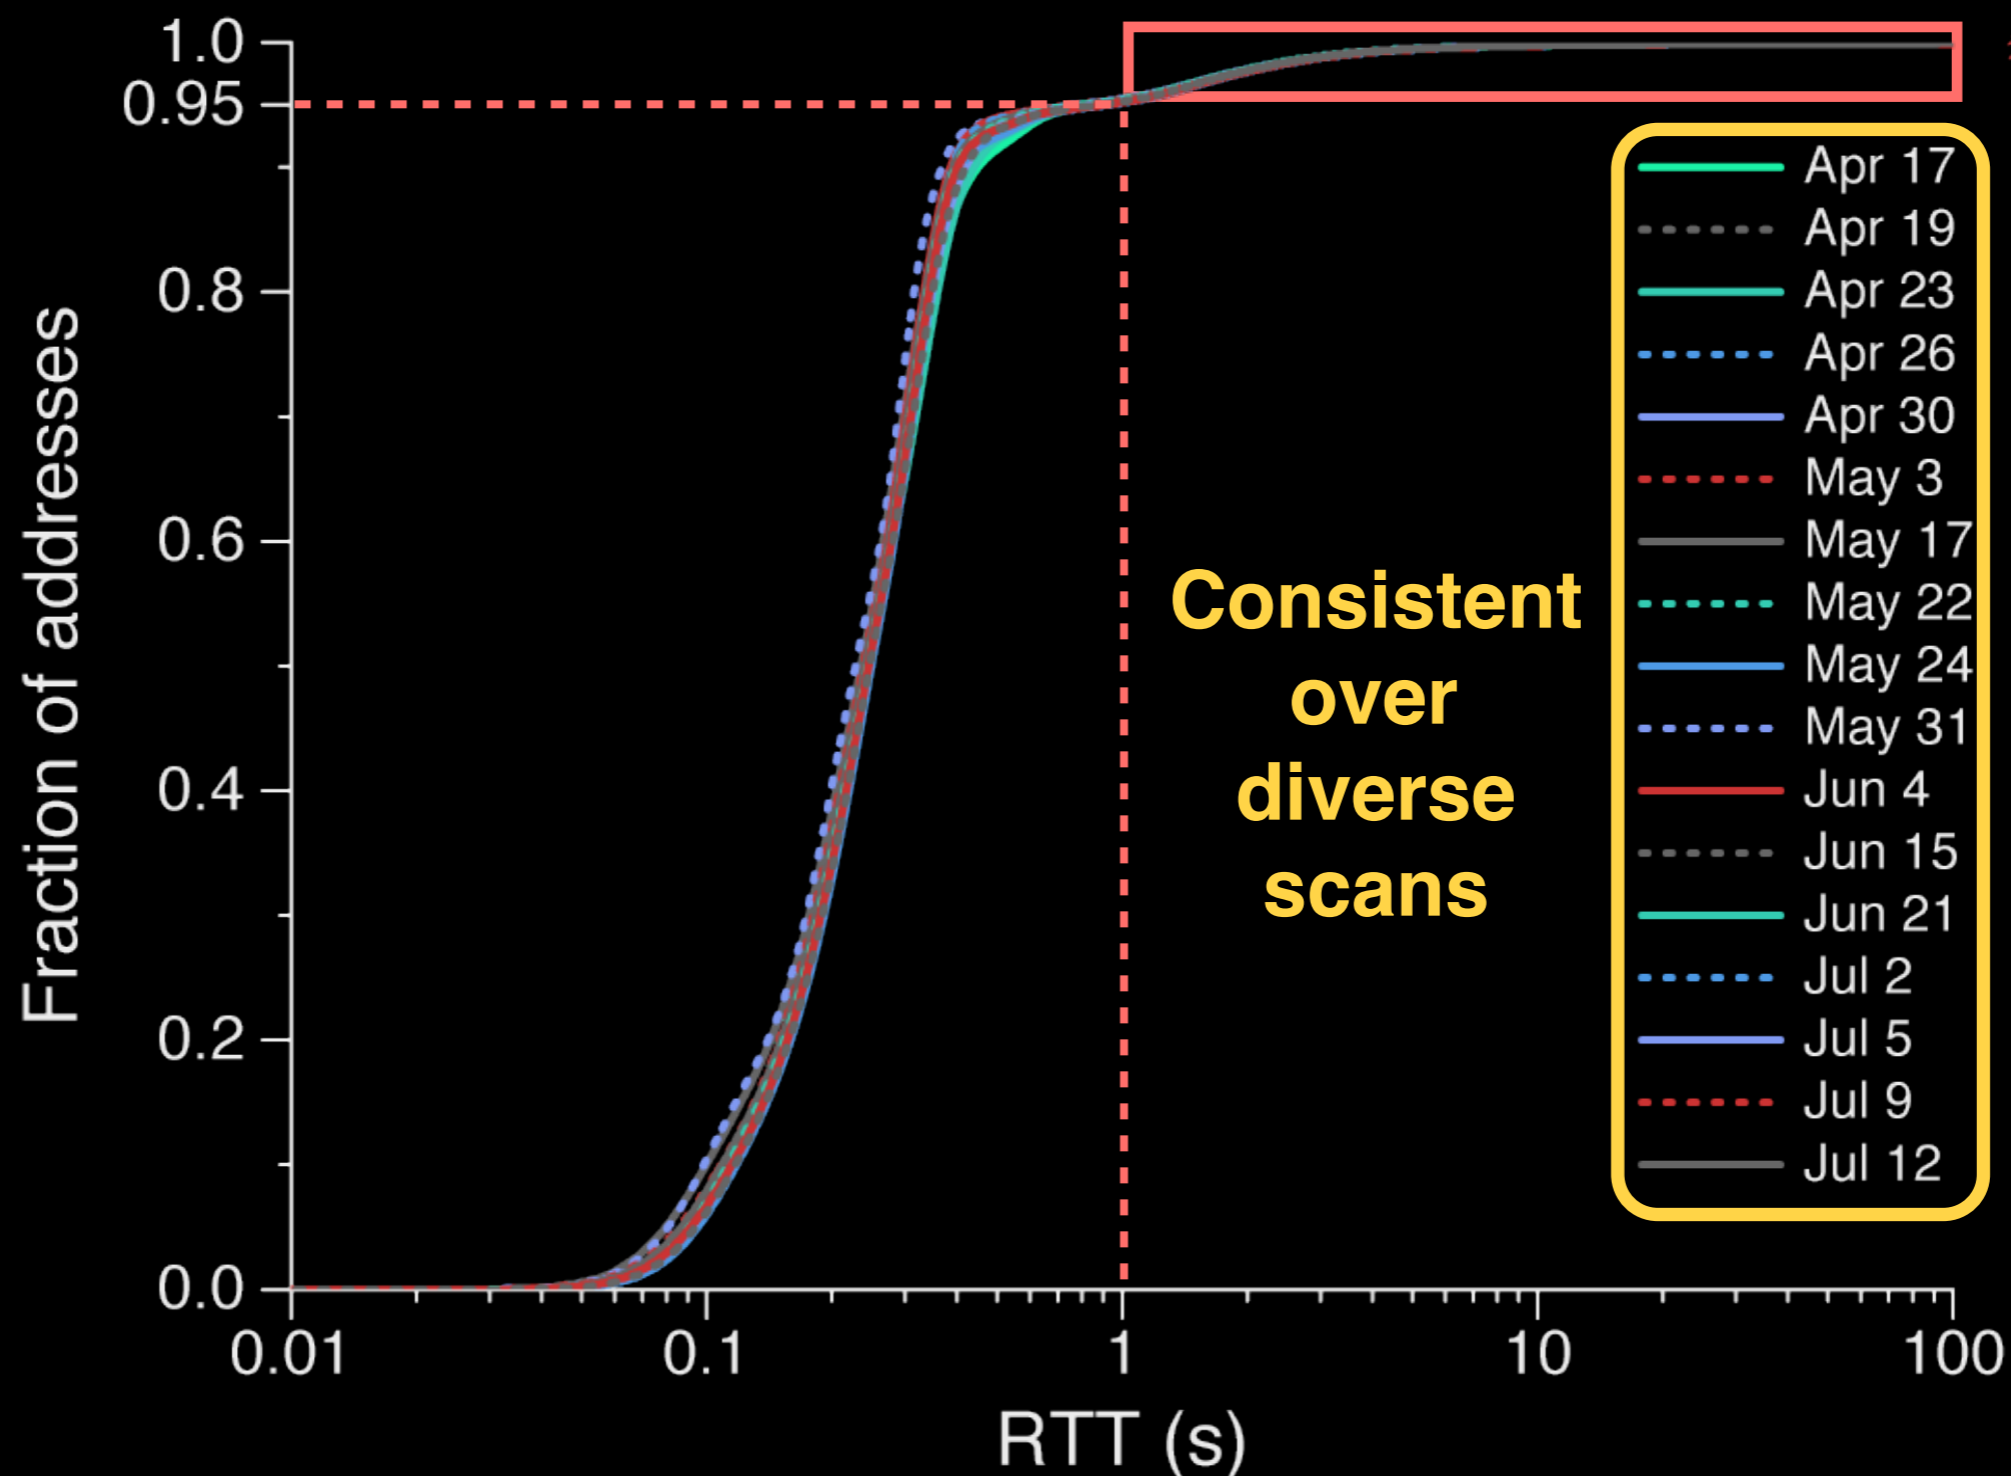
RTTs increasing since 2011



High RTTs from Zmap scans are consistent in time in 2015



Turtles



When should probes time out?

When should probes time out?

- Perhaps never? (depending upon your needs, e.g.: outage detection)
- If they must, beware high latencies
 - 5% of pings from 5% of addresses take longer than 5s
 - High latencies are persistent and consistent
 - High latencies affect ICMP, UDP, TCP