

On the Suitability of ping to Measure Latency

Cristel Pelsser <cristel@ij.ad.jp>

Luca Cittadini <luca.cittadini@gmail.com>

Stefano Vissicchio <stefano.vissicchio@uclouvain.be>

Randy Bush <randy@psg.com>

We Use Ping

- We were using Atlas and found ping variance to be unexpectedly high
- What is the actual distribution of the RTTs?
- Could there be a rich distribution?
- So we decided to calibrate our tools
- We ran a paris traceroute series from Roma Tre to Ashburn

Paris Traceroute Uses
Flow-ID
To Explore Hashed
ECMP / LAGged
Alternative Paths

Source port	Destination port
Length	Checksum

Figure 1: UDP header [16]. Fields in bold are part of the flow-id.


Type	Code	Checksum
Identifier		Sequence Number

Figure 2: ICMP echo message [17]. Echo request messages have type=8 and code=0. Echo reply messages have type=0 and code=0.

Type	Code	Checksum
unused (zero)		
IP Header + 64 bits of payload		

Figure 3: ICMP port unreachable message [17]. Type and code fields are both set to 3.

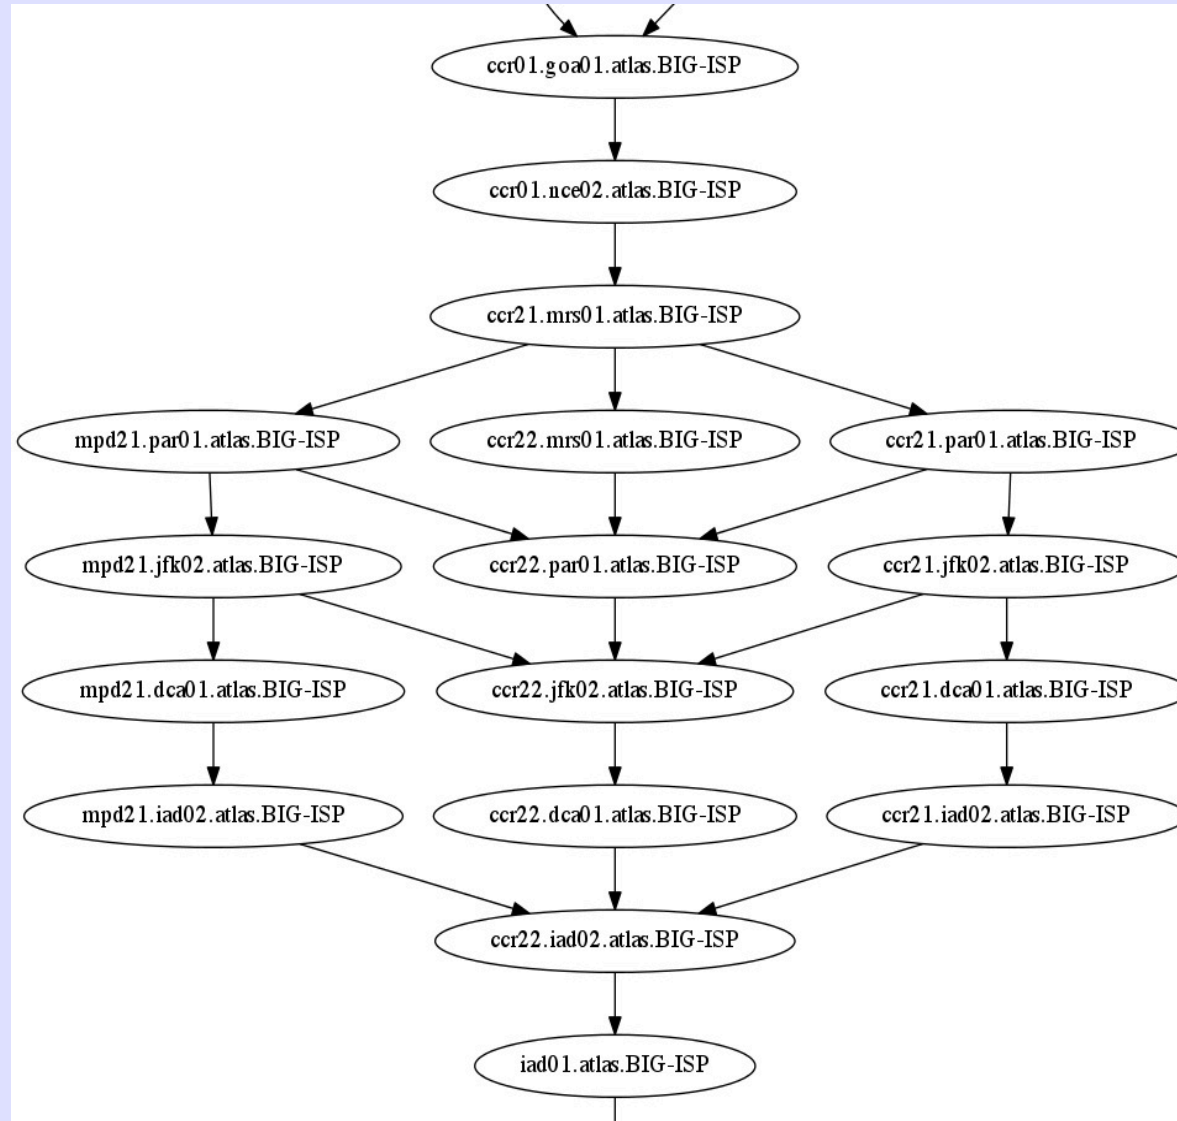
Roma Tre - Ashburn




Notice Linux Rounding

ECMP

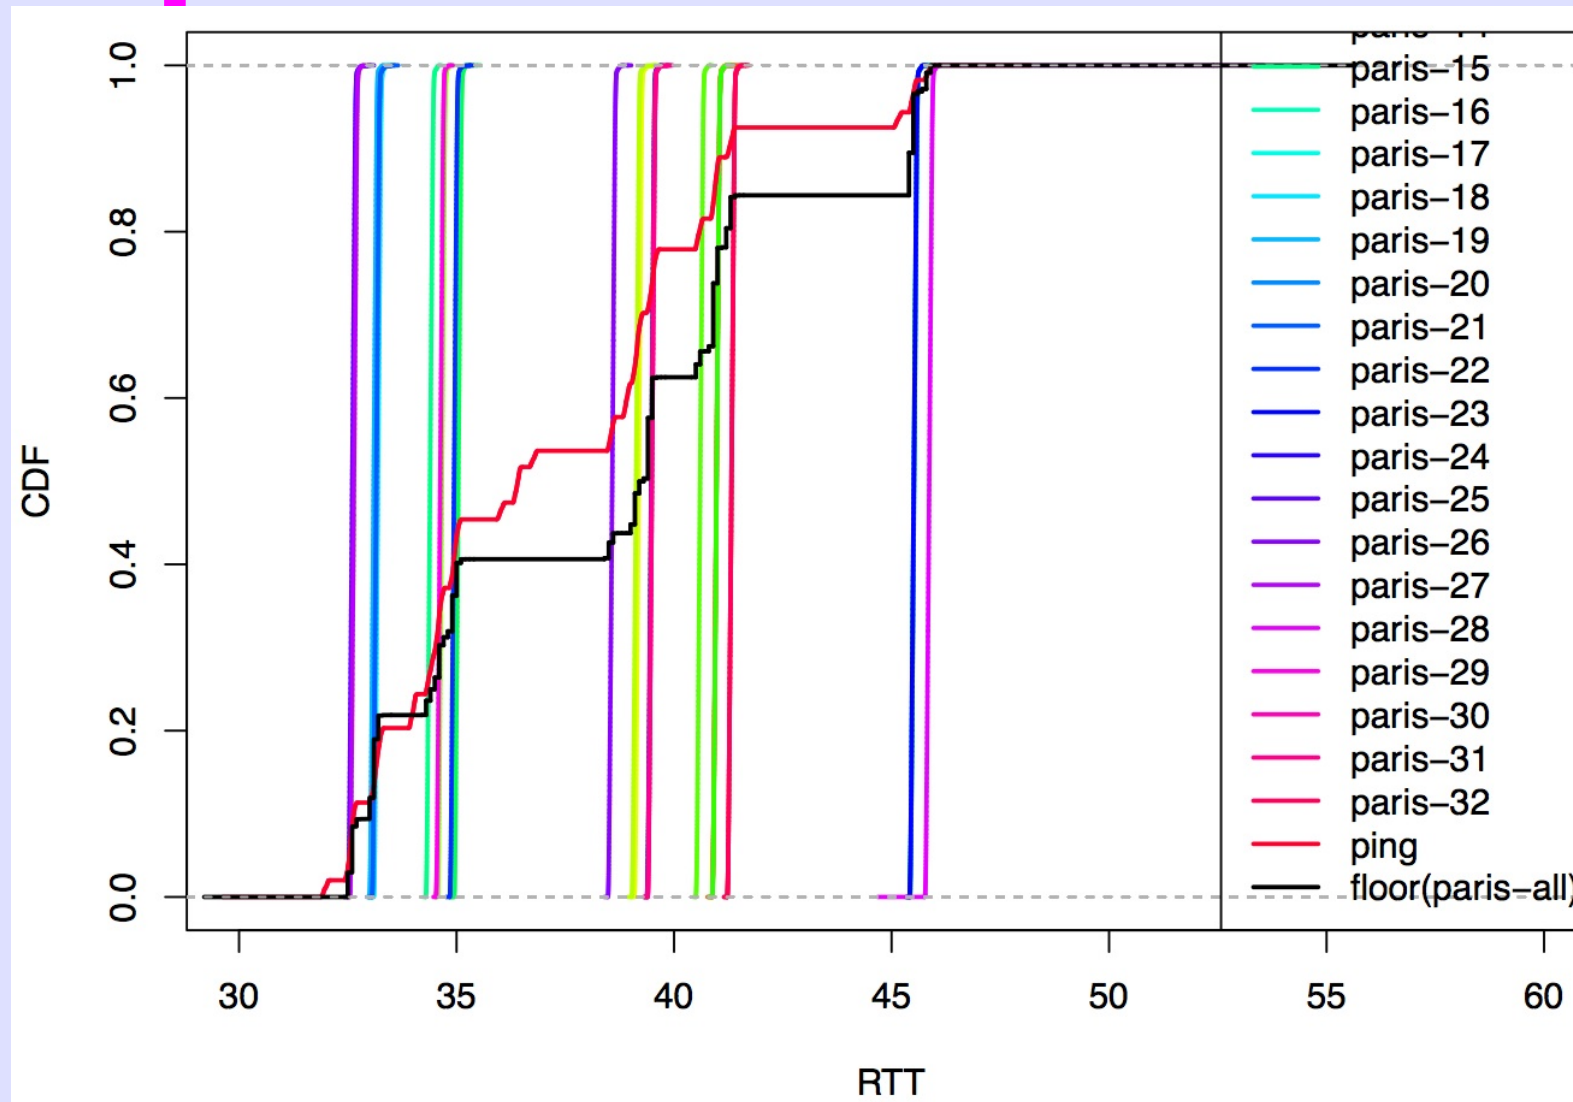
!= Equal Latency Multi-Path



So, A Simple Path





Simple Path Used LAG



FreeBSD Does Not Round (yay!)

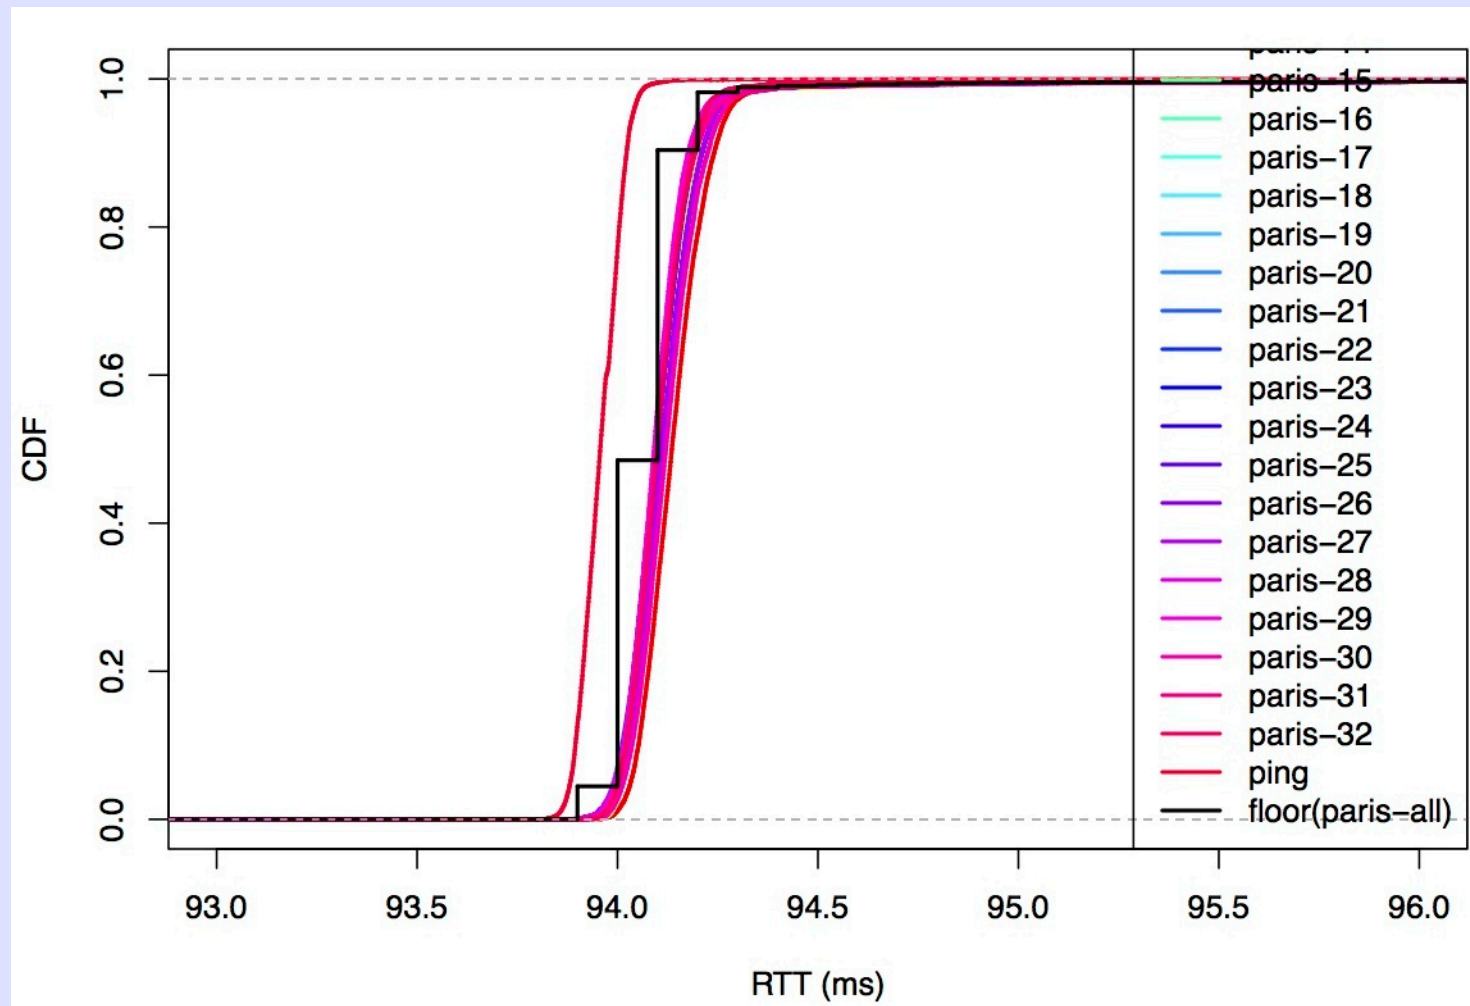
It's Not The Hosts



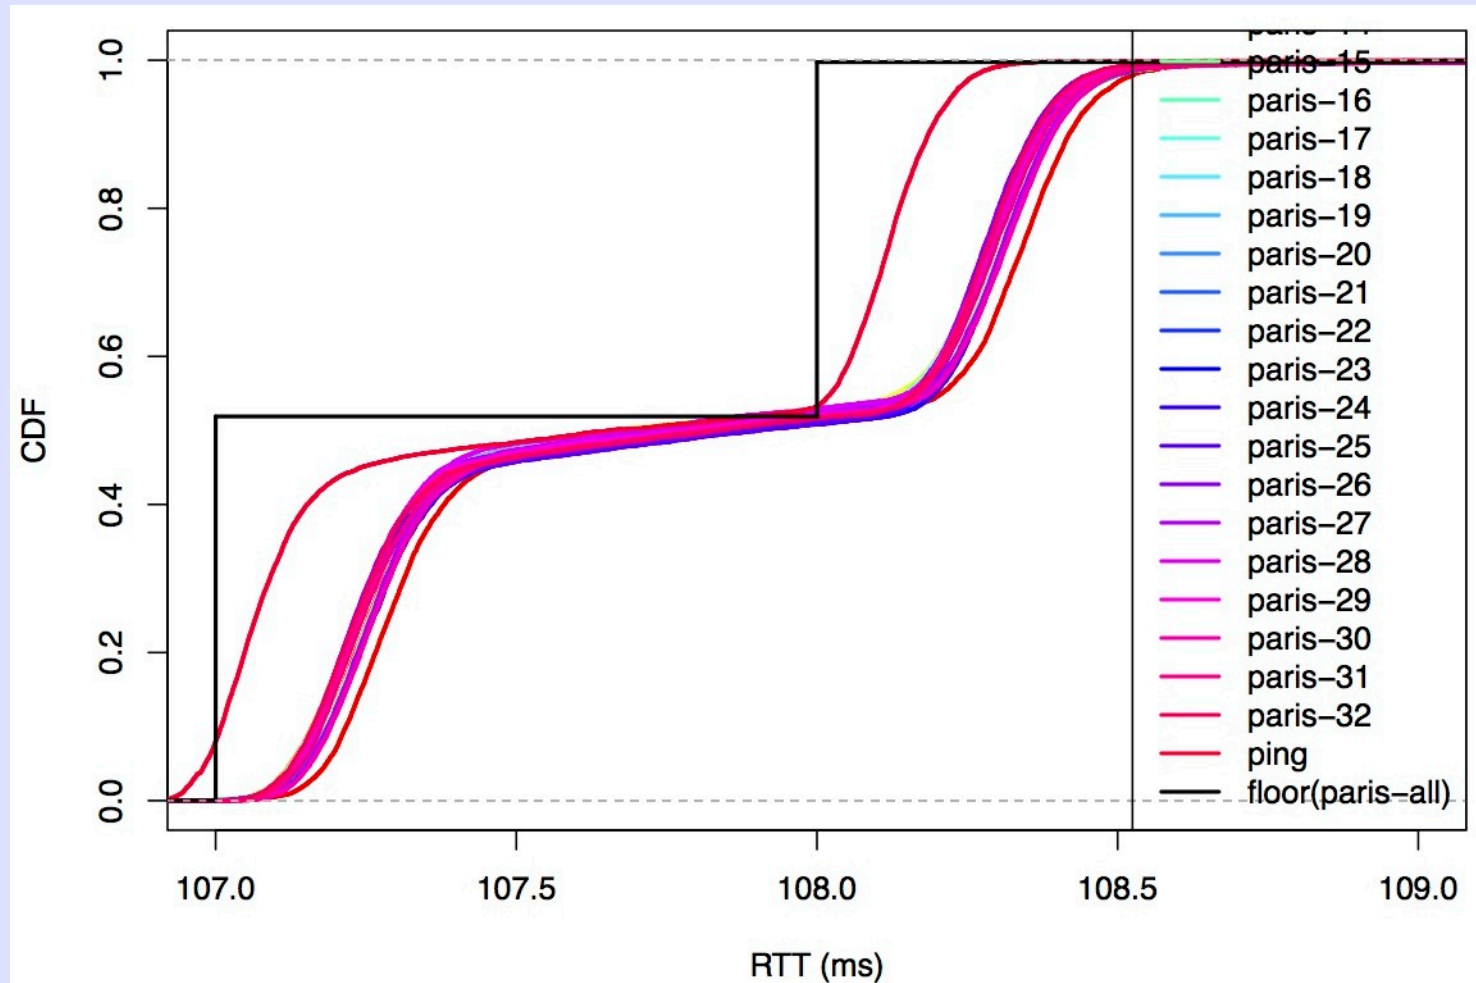
Looking for Causes

- Very Reproducible
- Circuit Loading $\leq 50\%$
- Same different times of day
- Same different probe timing/spacing
- Same UDP or ICMP
- MPLS seems not to affect
- Layer > 3 Hashing \Rightarrow Large Effect

Layer-3 Only Hash



Event During Run



Selected Summary

Src	Dst	ISPs	ECMP	LAG	Dispersion
ROM	ASH	> 1	yes	yes	4.21%
ROM	BXL	> 1	yes	yes	5.02%
ROM	LON	> 1	yes	yes	5.20%
TYO	SJO	1	yes	yes	23.59%
DAL	LON	1	no	yes	24.32%
PAR	ASH	1	no	yes	0%
DAL	ASH	1	no	yes	42.23%
DAL	SEA	1	yes	yes	15.46%
DAL	DAL	1	no	no	0.57%
DAL	DAL	1	no	yes	1.39%

We'd Like a Large
Number of
Measurements

But Atlas is Imprecise