On the suitability of two large-scale Internet measurement platforms

Thomas Holterbach, Cristel Pelsser and Randy Bush

November 4, 2014

Problematic

RIPE Atlas probes are **small devices** with **low capacity** and **anyone** in the world is able to start measurements from and toward this probe.

Question

Are the probes powerful enough to perform these measurements ?

How to measure the suitability of the RIPE Atlas probes ?

Idea

We combine two measurement platforms : RIPE Atlas and NLNOG Ring

- ∼300 Ubuntu VMs in different ASs
- Provides powerful measurement tools such as **Scamper**

How to measure the suitability of the RIPE Atlas probes ?

12 Atlas probes and NL Ring nodes are on the same local network

- Each of these 12 NL Ring nodes performs pings toward
  - the 11 other NL Ring nodes and their associated Atlas probes
  - and two BSD servers
- We perform ping for 16 different flow-IDs

How to measure the suitability of the RIPE Atlas probes ?

12 Atlas probes and NL Ring nodes are on the same local network

- Each of these 12 NL Ring nodes performs pings toward
  - the 11 other NL Ring nodes and their associated Atlas probes
  - and two BSD servers
- We perform ping for 16 different flow-IDs

Consequence

We are now able to compare RIPE Atlas probes with

- NL Ring nodes
- and the two BSD servers

How to measure the suitability of the RIPE Atlas probes ?

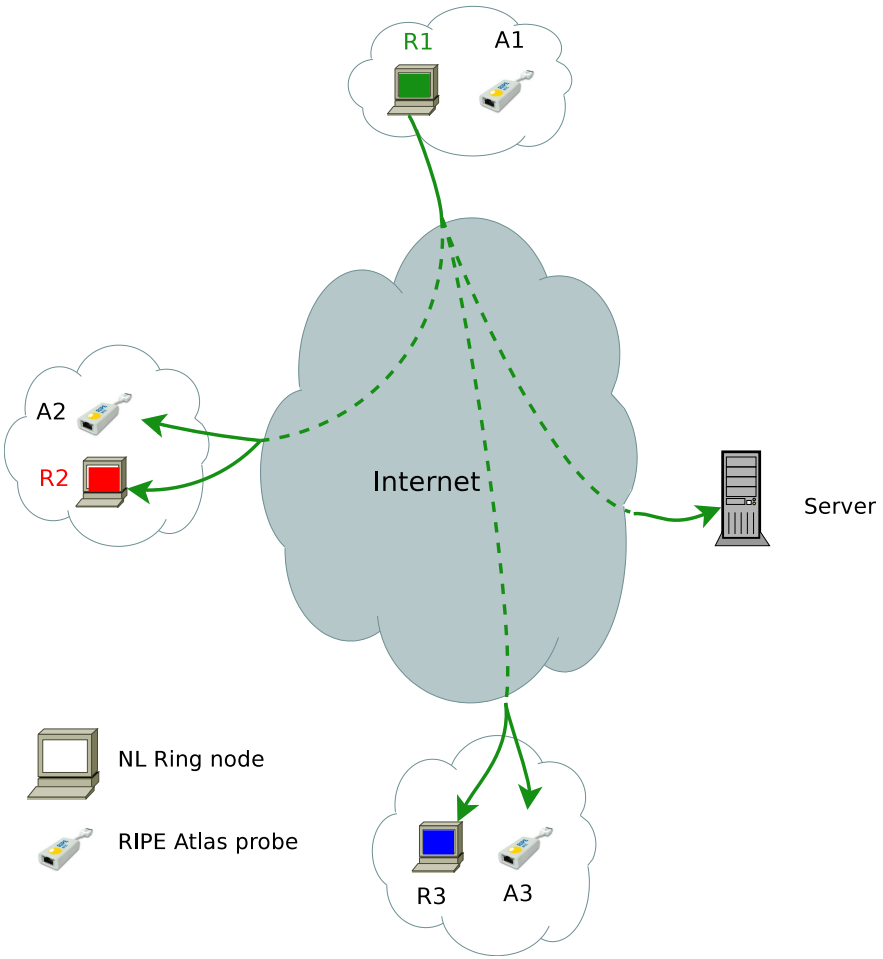

How to measure the suitability of the RIPE Atlas probes ?

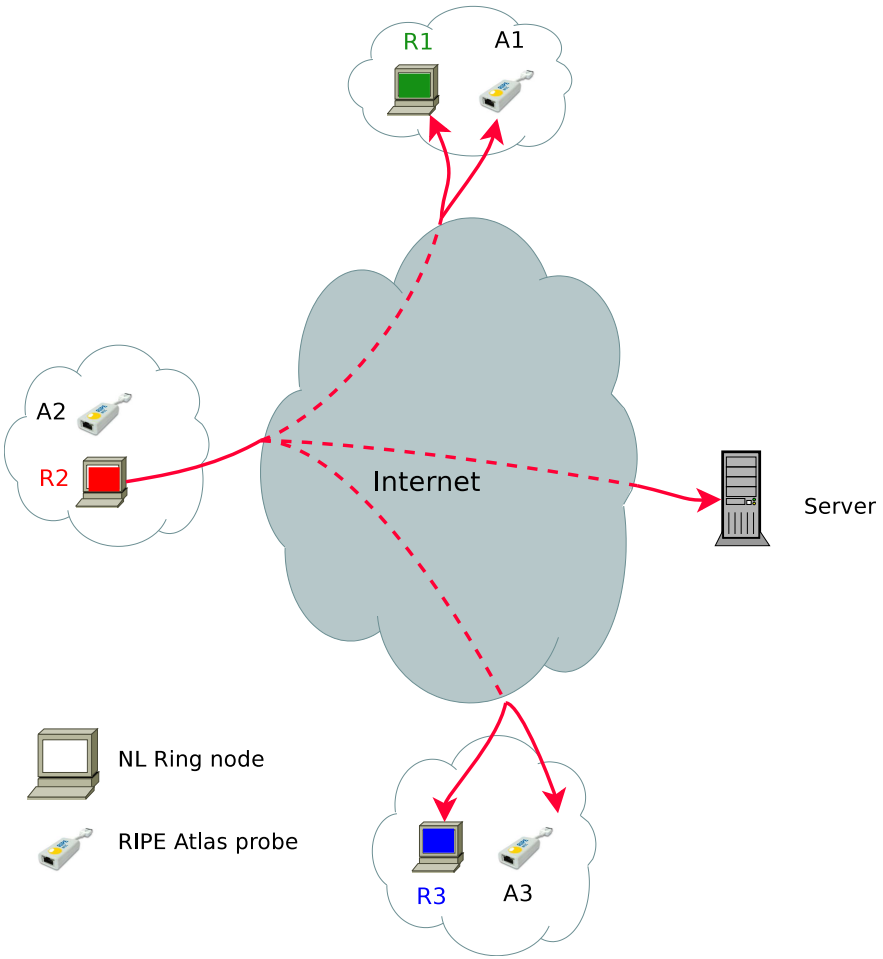

How to measure the suitability of the RIPE Atlas probes ?

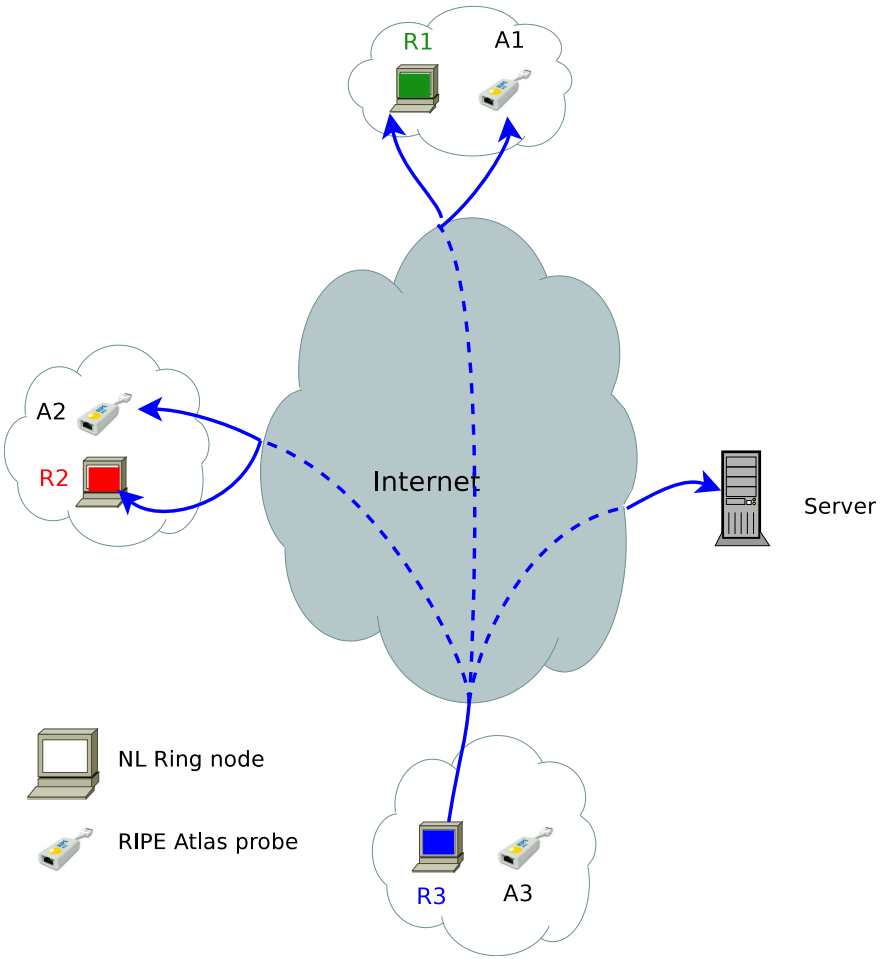

How to measure the suitability of the RIPE Atlas probes ?

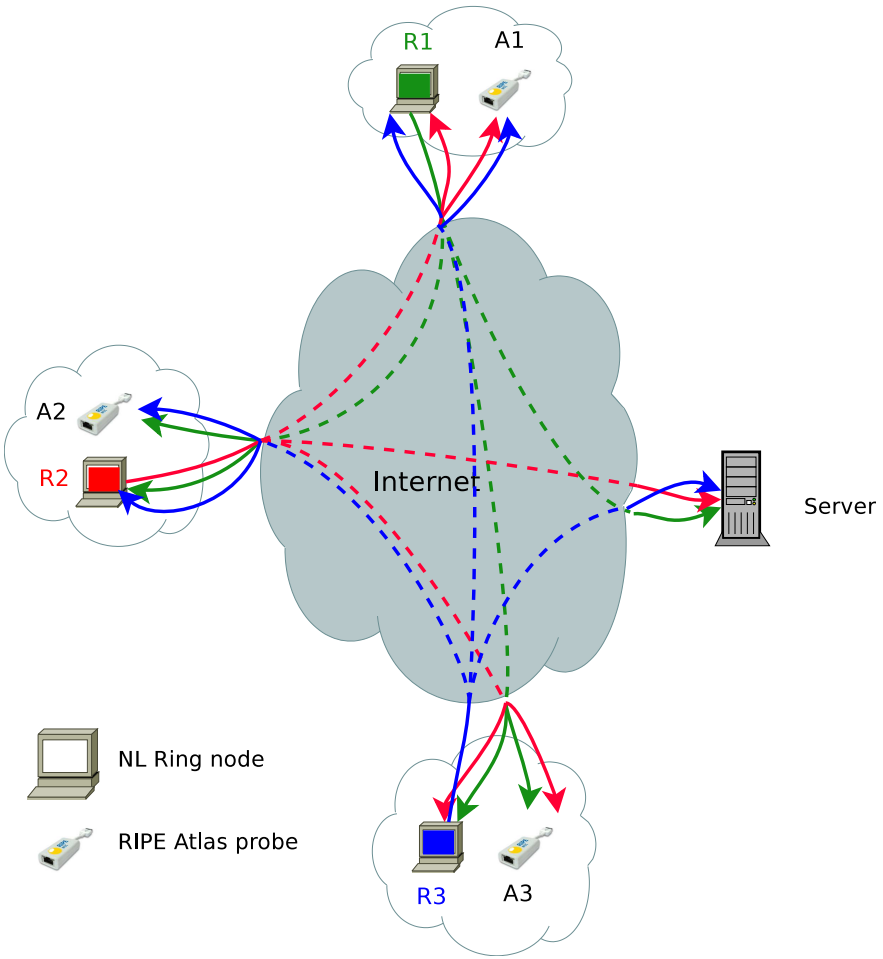

Actually, we have **12 Ring nodes**, **12 Atlas probes** and **2 servers**
$\Rightarrow$ 288 couples source-destination in total

Measurement orchestration

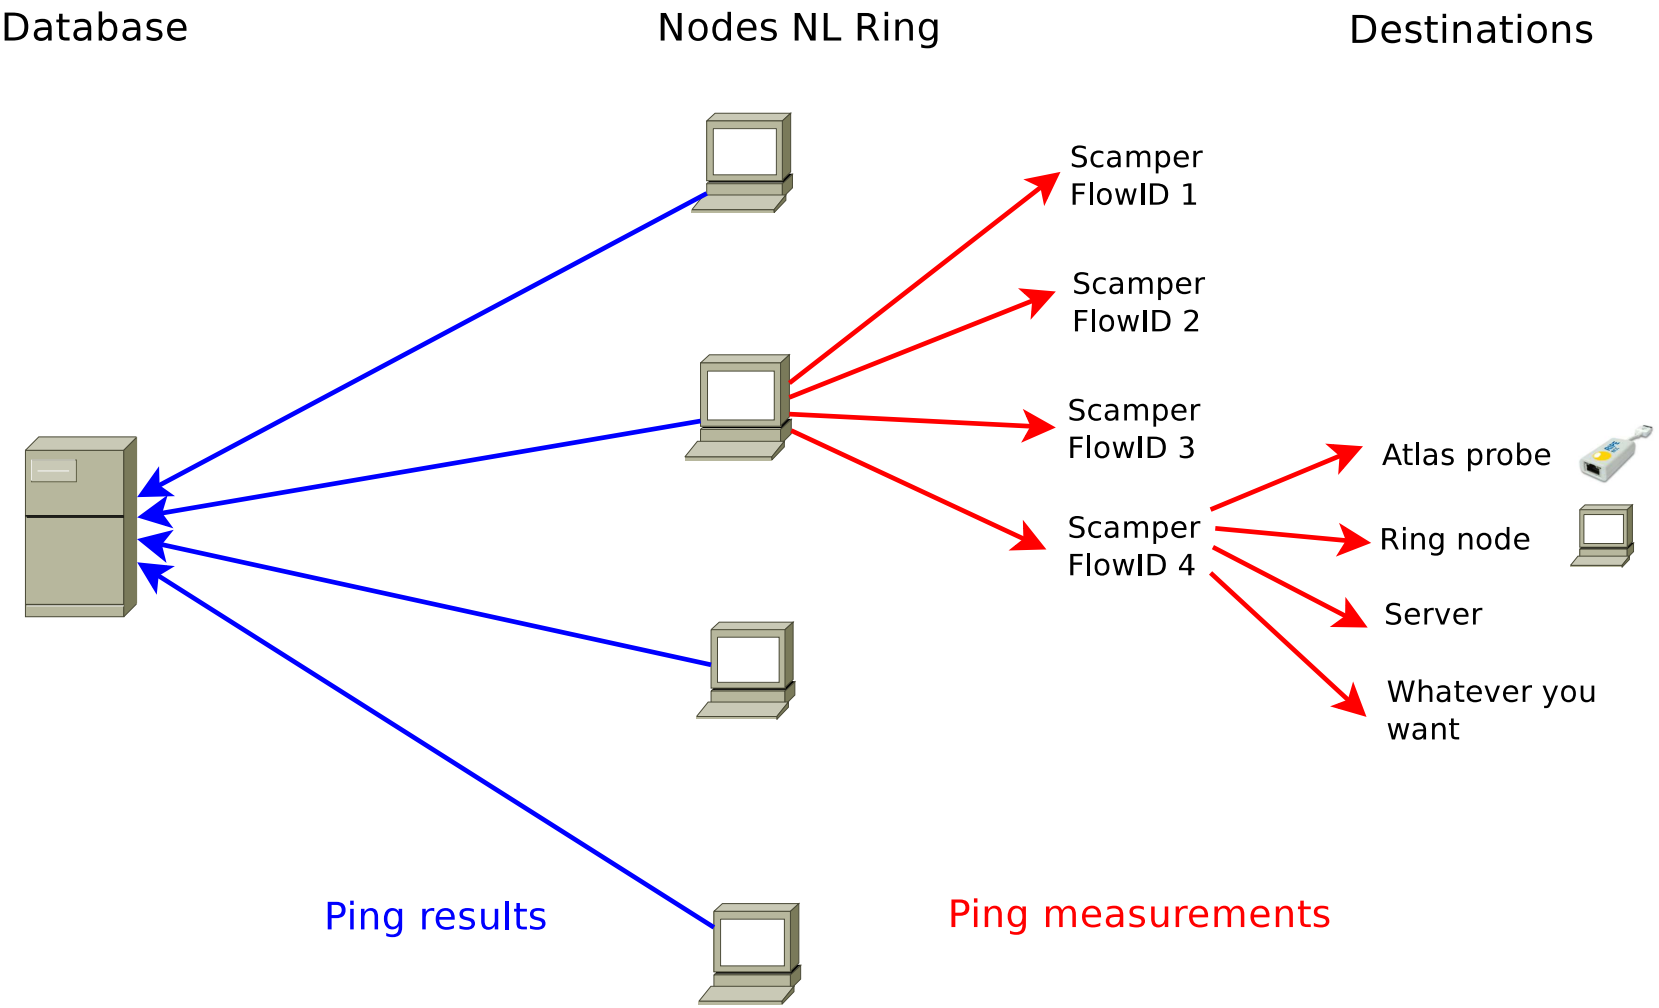

Measurement calibration

Ping frequency

Thanks to the NL Ring nodes, we are able to perform ping with a high frequency $\Rightarrow$ 1 ping every two seconds for each flow-ID

		Theoretical				Actual	
		Every two seconds		One week		One week	
Platform	Number	Ping sent	Ping received	Sent	Received	Sent	Received
Ring node	12	384	176	116M	53M	107M	49M
Atlas probe		0	176	0	53M	0	49M
Server	2	0	192	0	58M	0	53M

Is an event coming from the Internet or the end-points ?

- Events occur both on the NL Ring node and the RIPE Atlas probe
  $\Rightarrow$ These events are coming from the Internet

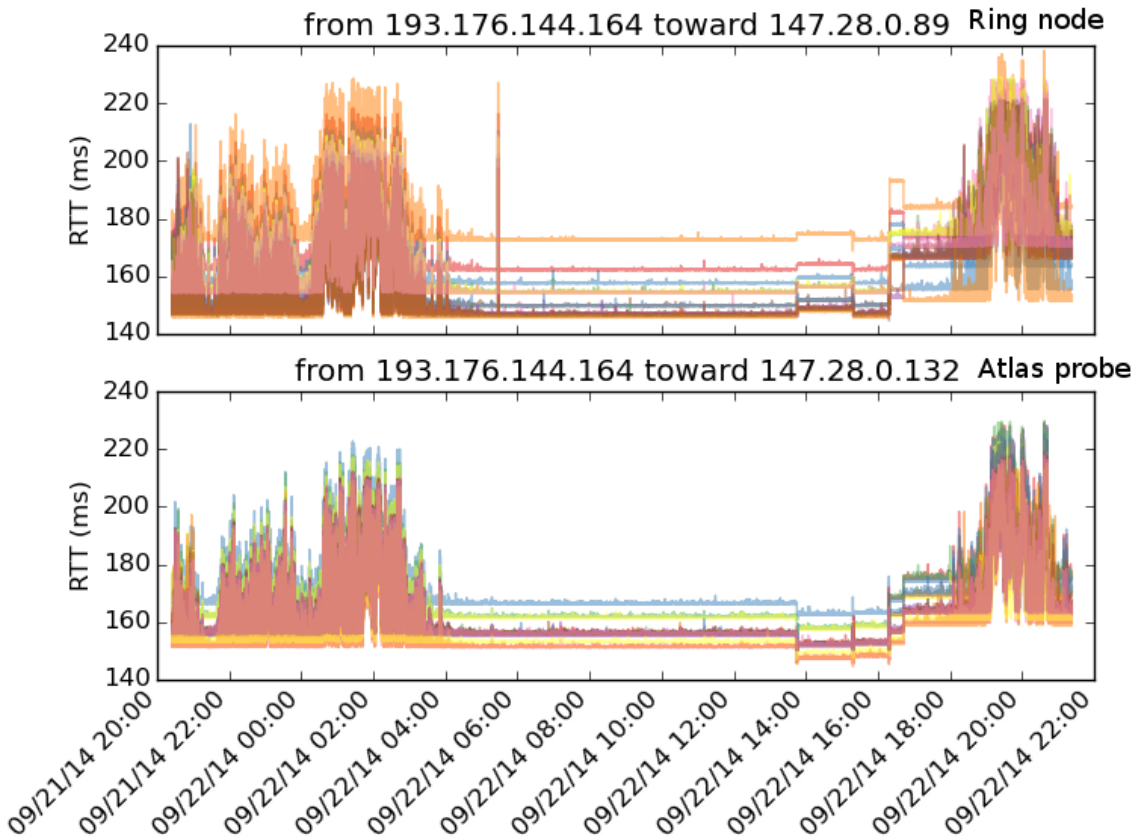

Is an event coming from the Internet or the end-points ?

- An event between 14:00 and 20:00 only occurs with the Atlas probe
- This event might be due to this RIPE Atlas probe

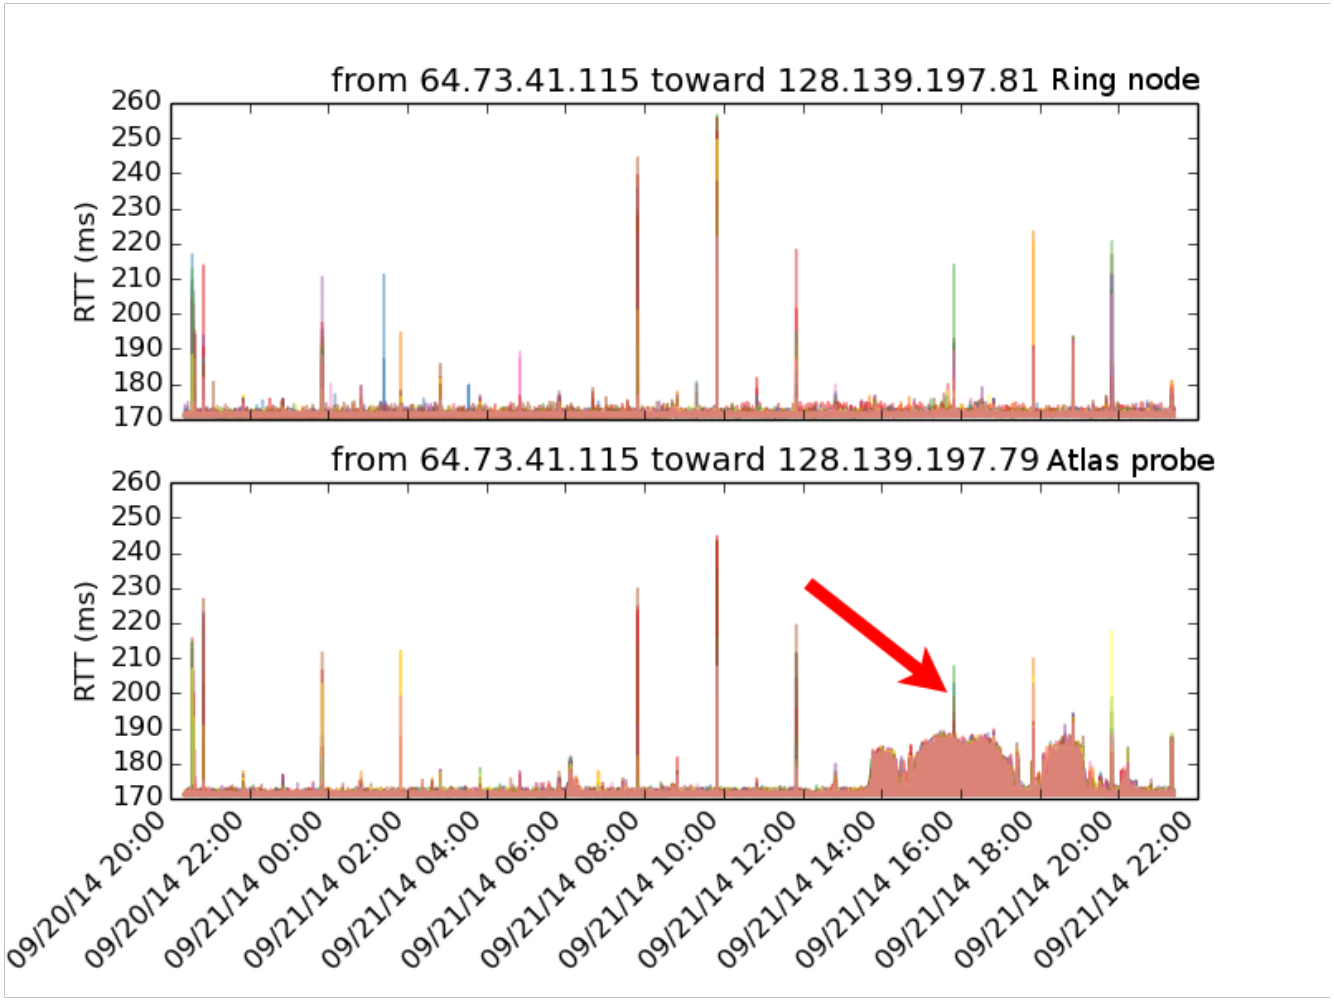

Is an event coming from the Internet or the end-points ?

- An event between 14:00 and 18:00 only occurs with the Ring node
- This event might be due the NL Ring node

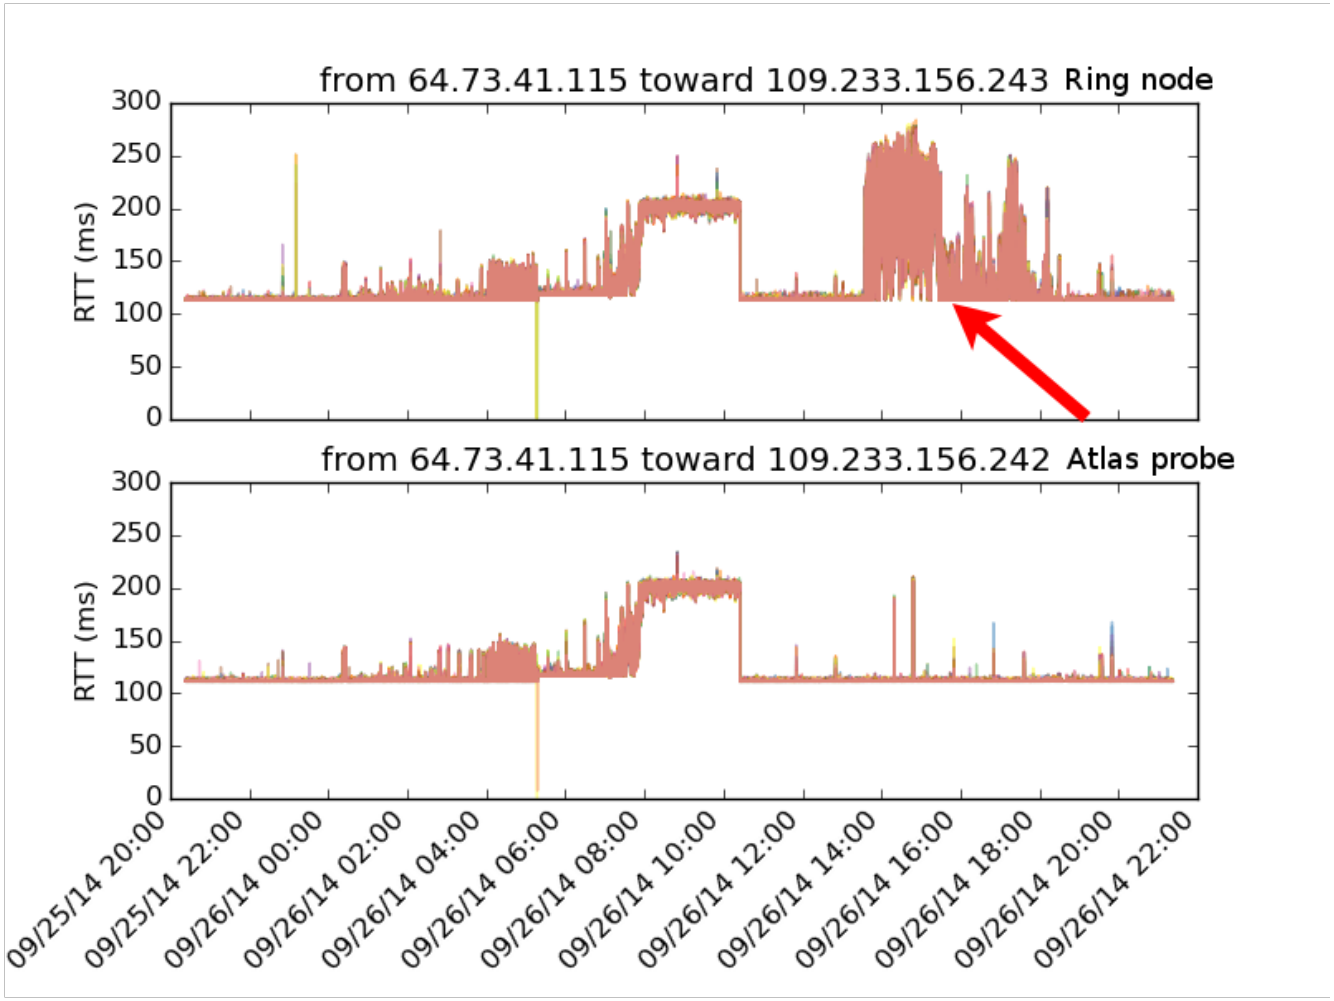

Is an event coming from the Internet or the end-points ?

- Same destination, same time, another source
- The previous event on the NL Ring node does not occur this time
  $\Rightarrow$ the previous event is not coming from the NL Ring node

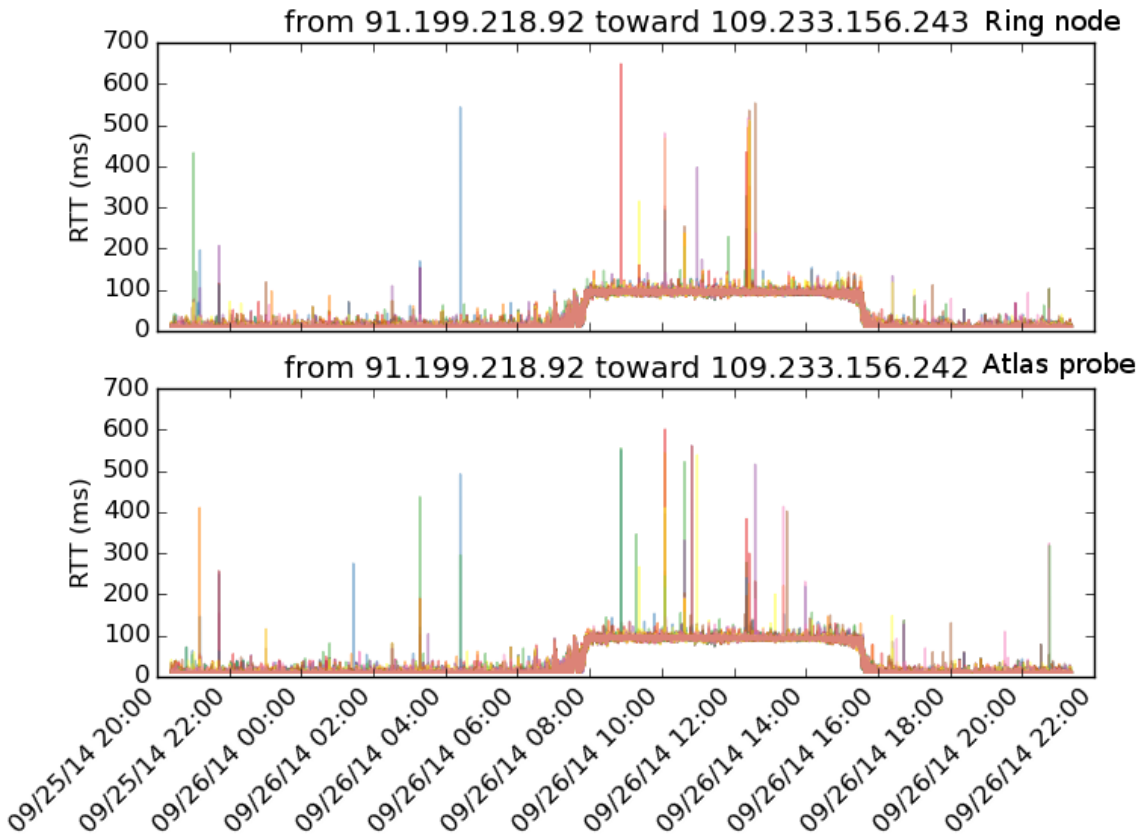

Some statistics over one week

Mean RTT

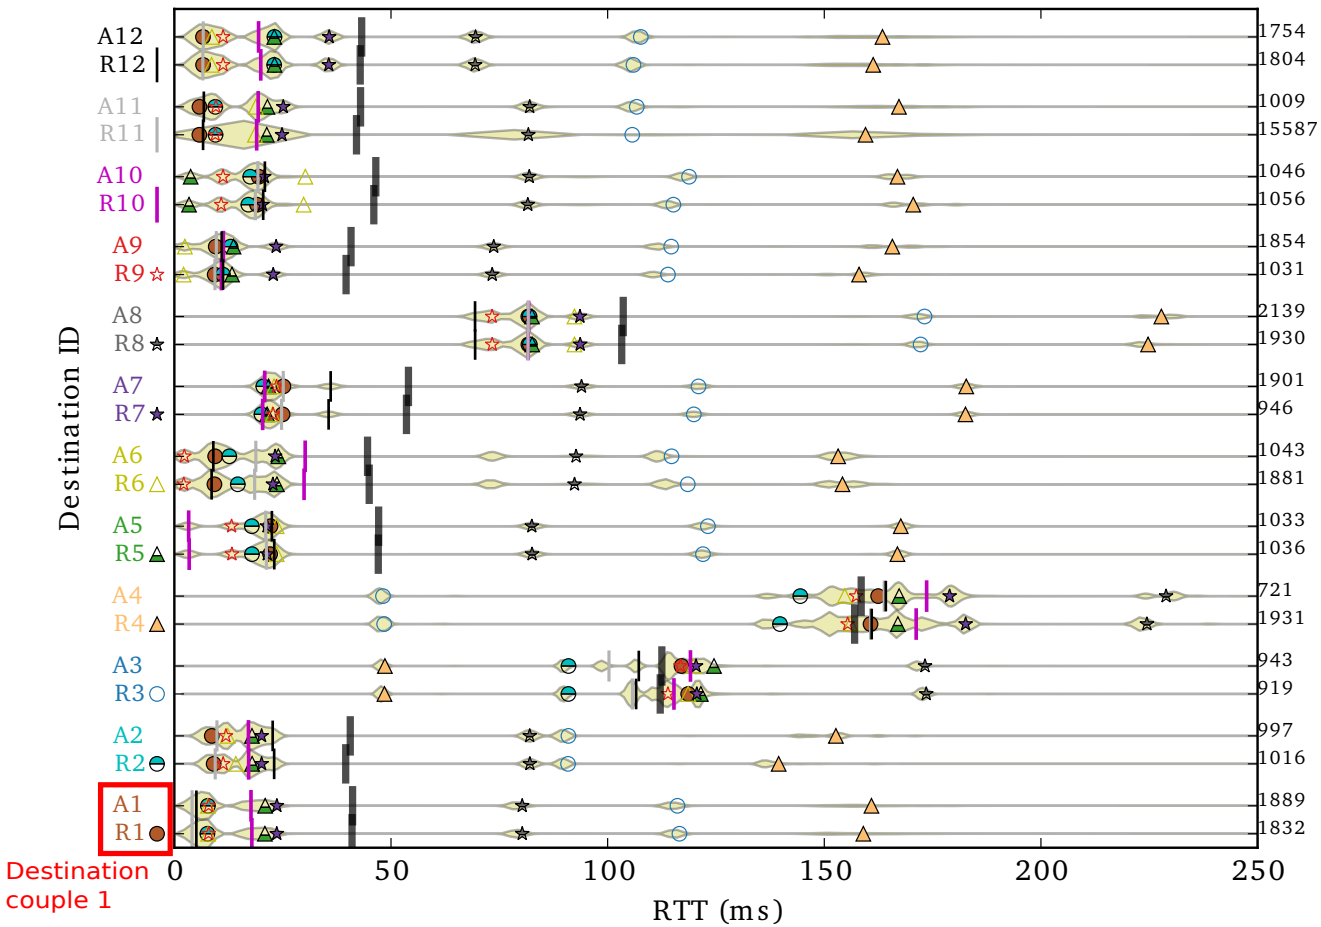
A12
R12
A11
R11
A10
R10
A9
R9
A8
R8
A7
R7
A6
R6
A5
R5
A4
R4
A3
R3
A2
R2
A1
R1
1754
1804
1009
15587
1046
1056
1854
1031
2139
1930
1901
946
1043
1881
1033
1036
721
1931
943
919
997
1016
1889
1832
Destination ID
Destination couple 1
0
50
100
150
200
250
RTT (ms)

Mean RTT

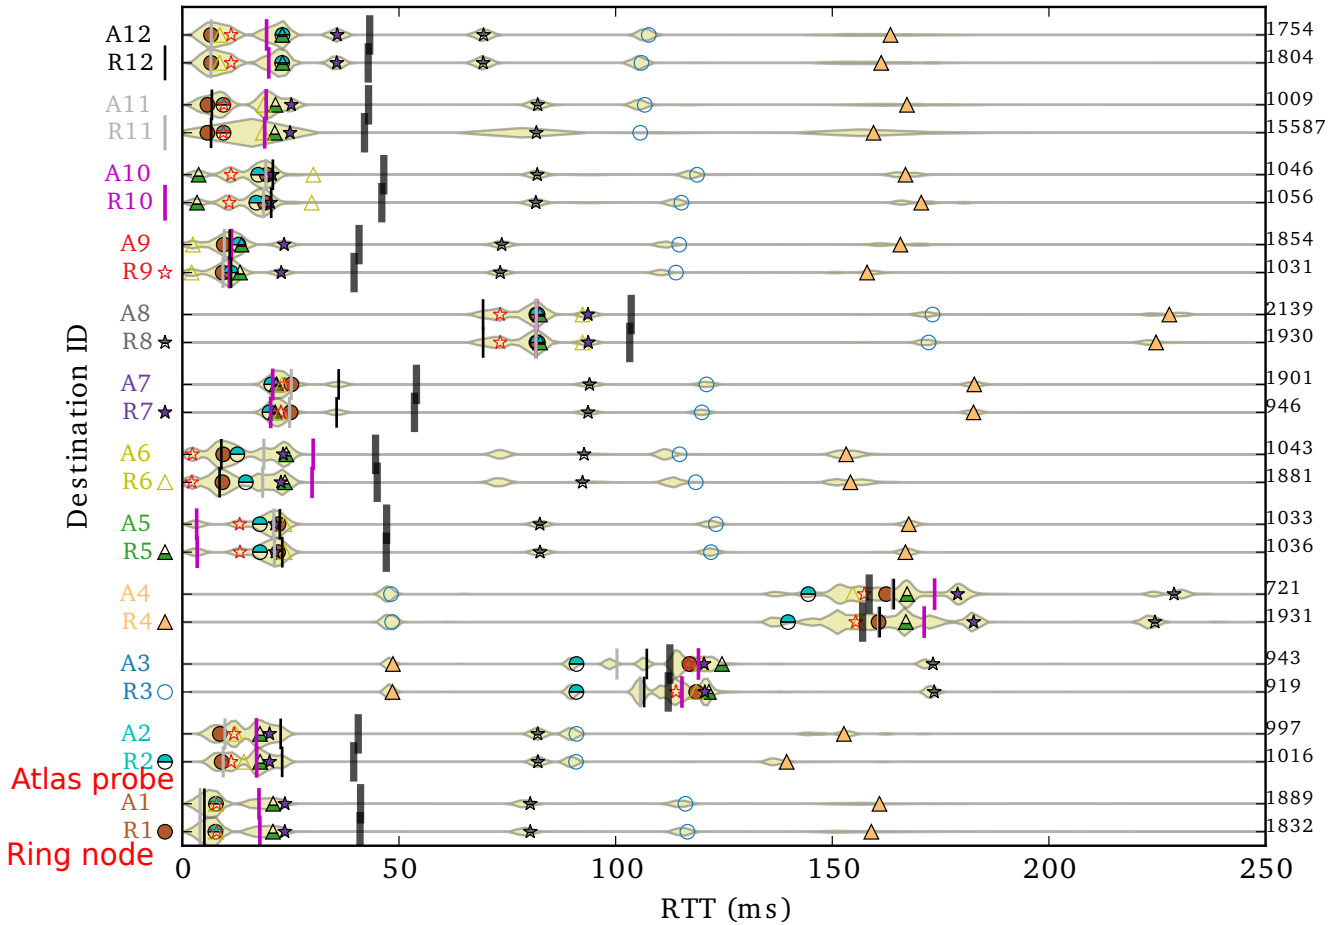
A12
R12
A11
R11
A10
R10
A9
R9
A8
R8
A7
R7
A6
R6
A5
R5
A4
R4
A3
R3
A2
R2
A1
R1
Atlas probe
Ring node
Destination ID
1754
1804
1009
15587
1046
1056
1854
1031
2139
1930
1901
946
1043
1881
1033
1036
721
1931
943
919
997
1016
1889
1832
0
50
100
150
200
250
RTT (ms)

Mean RTT

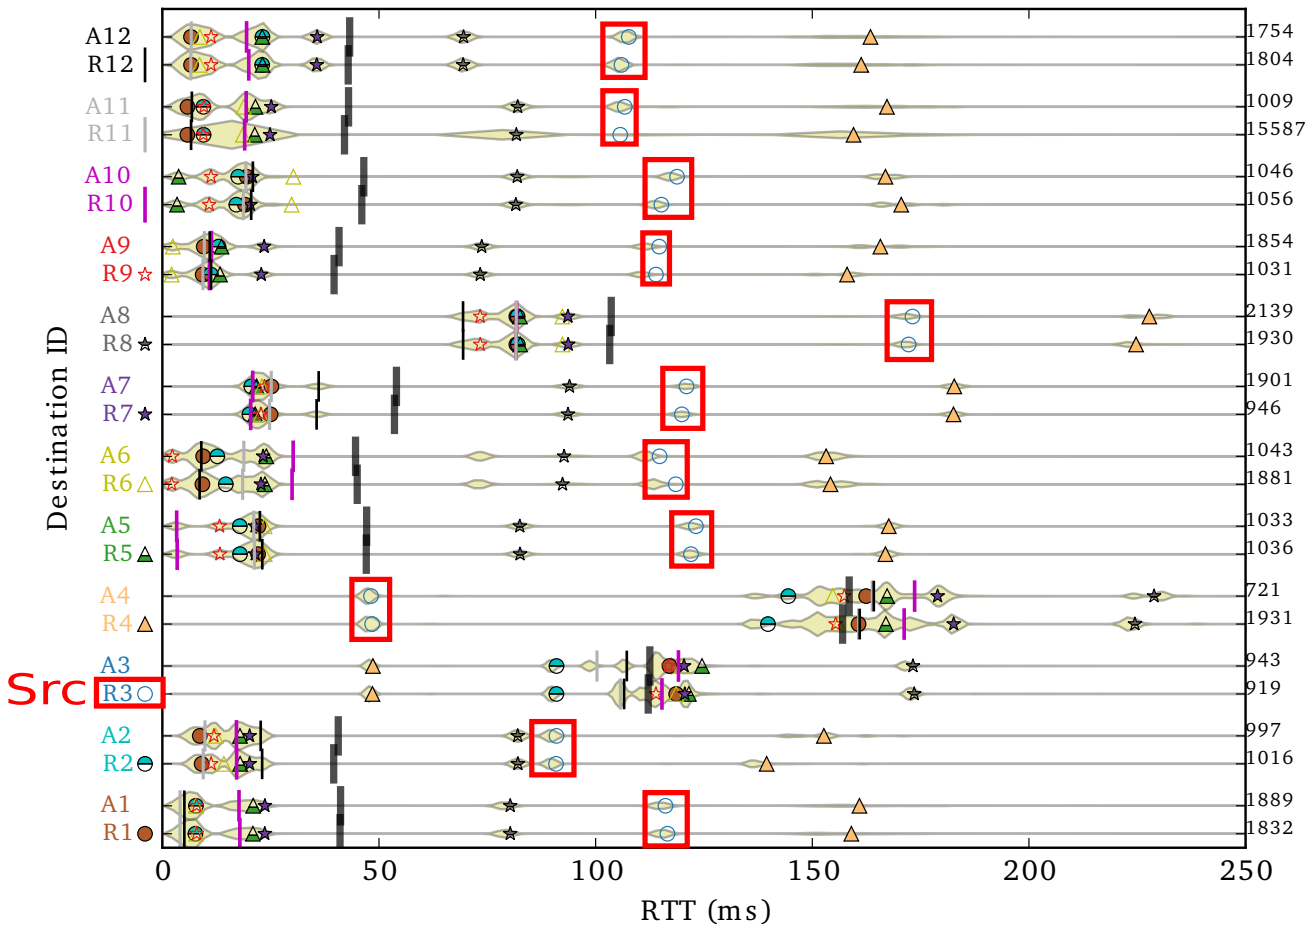
A12
R12
A11
R11
A10
R10
A9
R9
A8
R8
A7
R7
A6
R6
A5
R5
A4
R4
A3
R3
A2
R2
A1
R1
Src
Destination ID
1754
1804
1009
15587
1046
1056
1854
1031
2139
1930
1901
946
1043
1881
1033
1036
721
1931
943
919
997
1016
1889
1832
0
50
100
150
200
250
RTT (ms)

Mean RTT

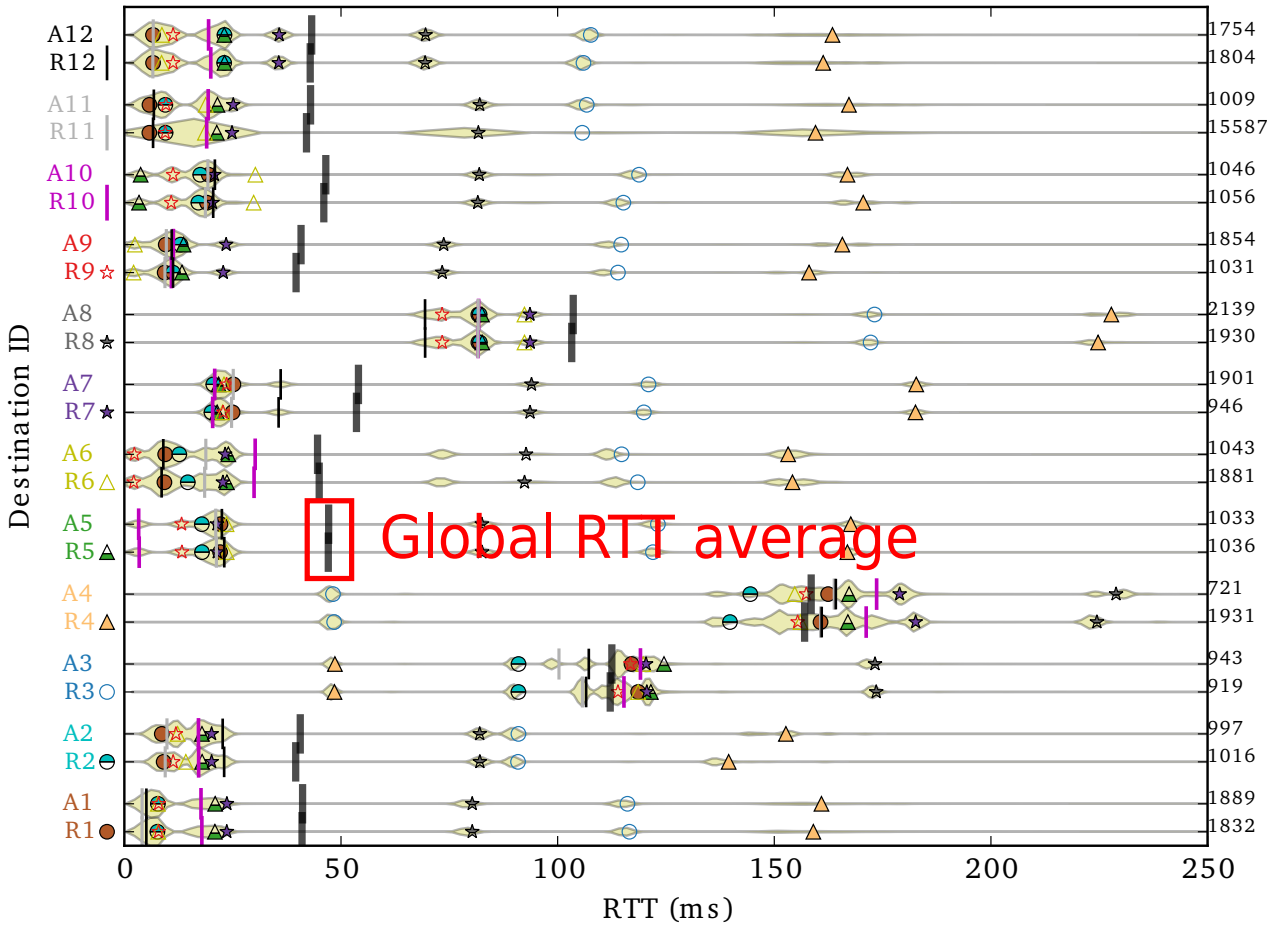
A12
R12
A11
R11
A10
R10
A9
R9
A8
R8
A7
R7
A6
R6
A5
R5
A4
R4
A3
R3
A2
R2
A1
R1
1754
1804
1009
15587
1046
1056
1854
1031
2139
1930
1901
946
1043
1881
1033
1036
721
1931
943
919
997
1016
1889
1832
Global RTT average
Destination ID
0
50
100
150
200
250
RTT (ms)

Mean RTT

Between these two platforms

- the global average RTT difference is 0.5ms
- there is at most 1.5ms difference in global RTT average

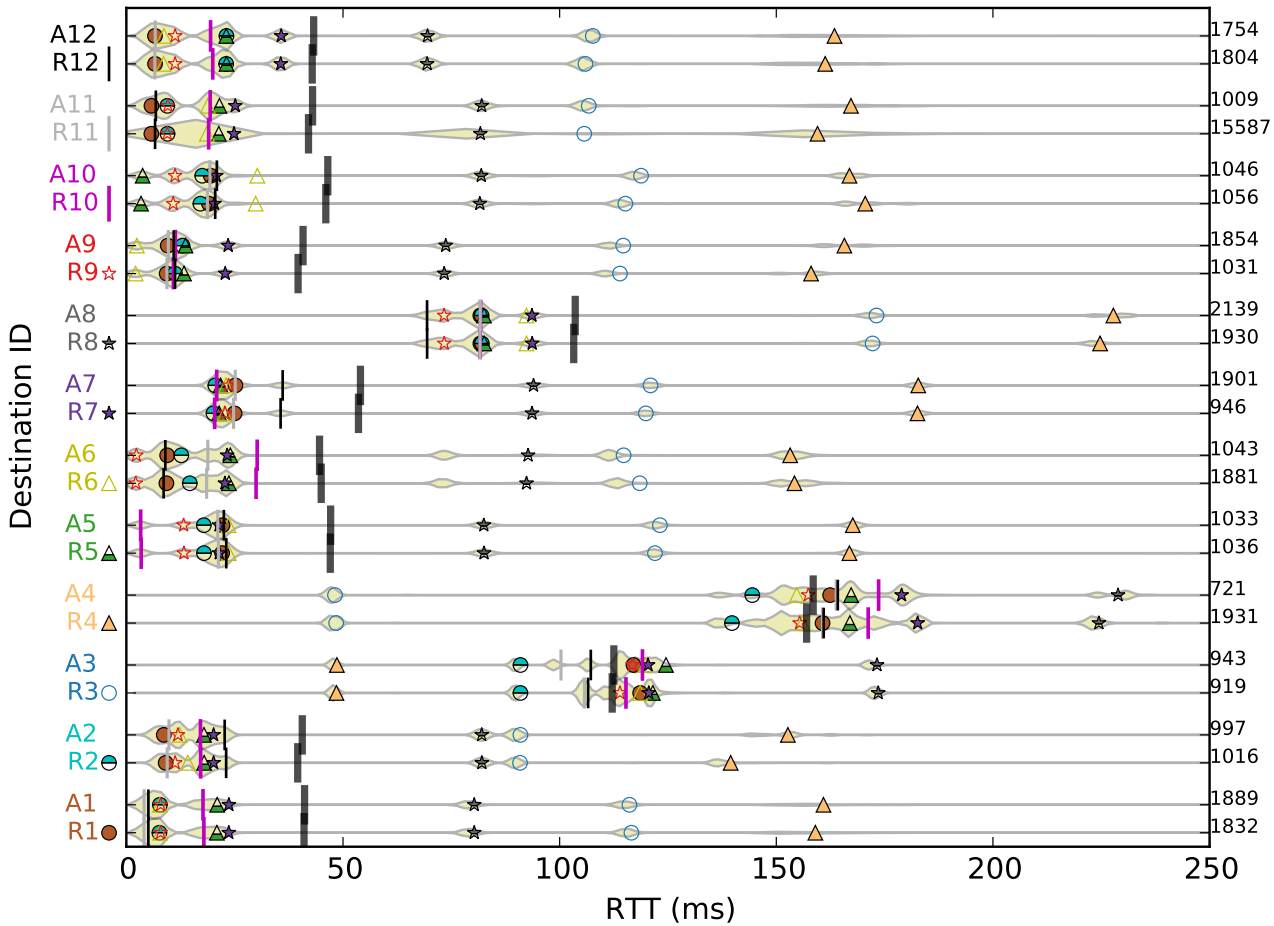

Mean RTT

Between these two platforms

- the global average RTT difference is 0.5ms
- there is at most 1.5ms difference in global RTT average

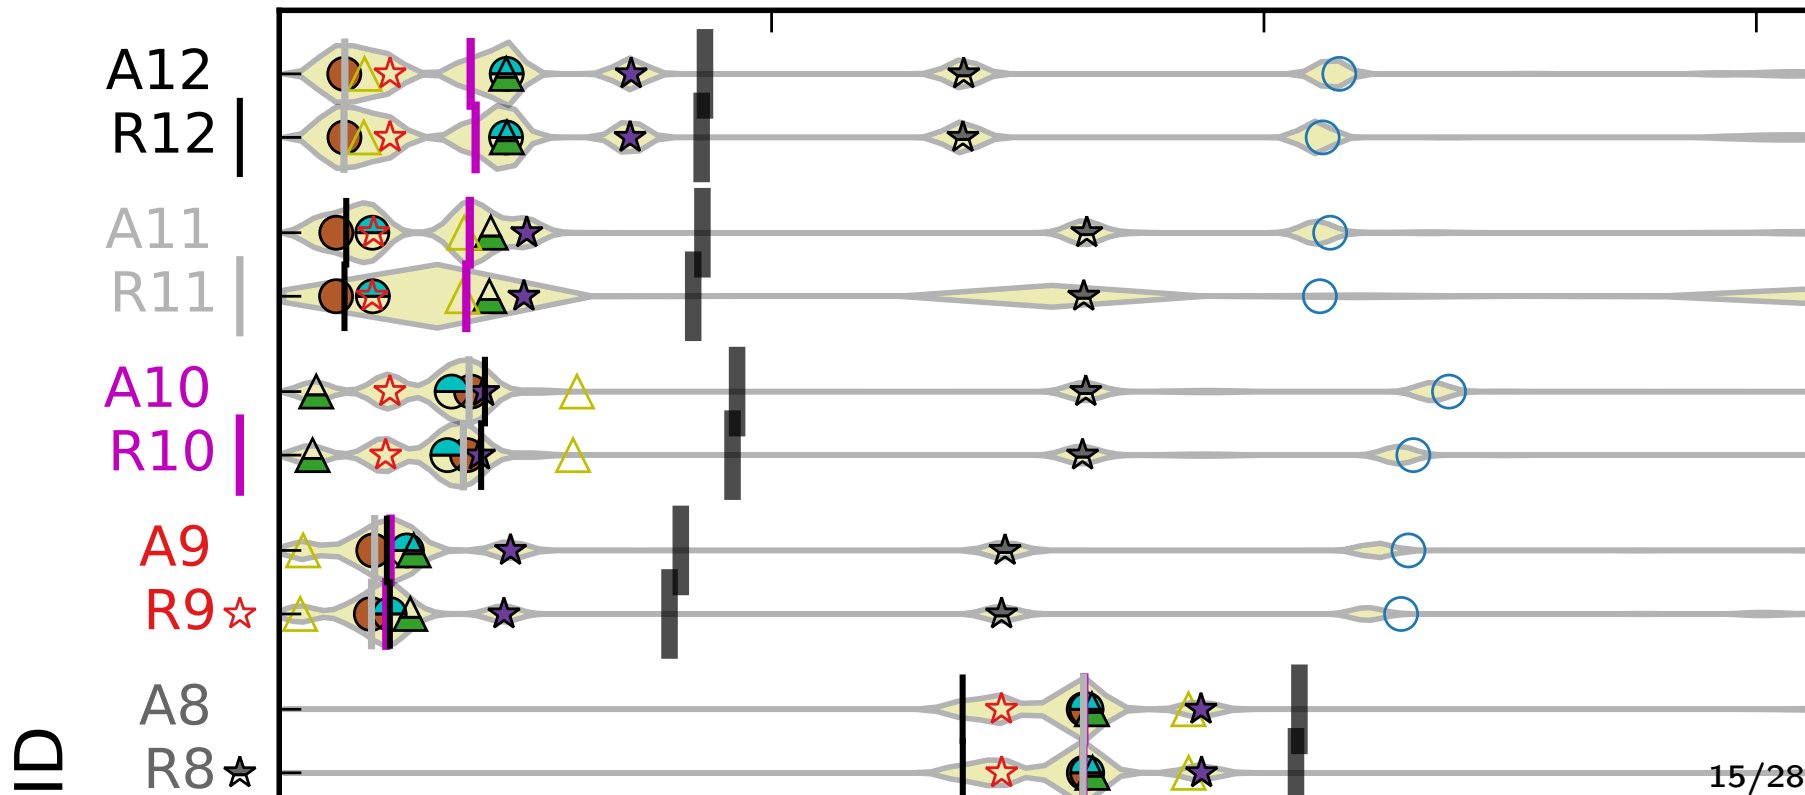

Standard deviation

To take into account the flow-IDs effect, we define a new RTT variability metric

RTT variability metric

Δ_s^d is the average RTT standard deviation over the 16 flow-IDs between the source s and the destination d

$$\Delta_s^d = \frac{\sum_{f=0}^{15} SD_s^d(f)}{16}$$

$SD_s^d(f)$ is the RTT standard deviation between the source s and the destination d for flow-ID f

Mean of the Standard Deviation Across Flow-IDs Versus Global Standard deviation

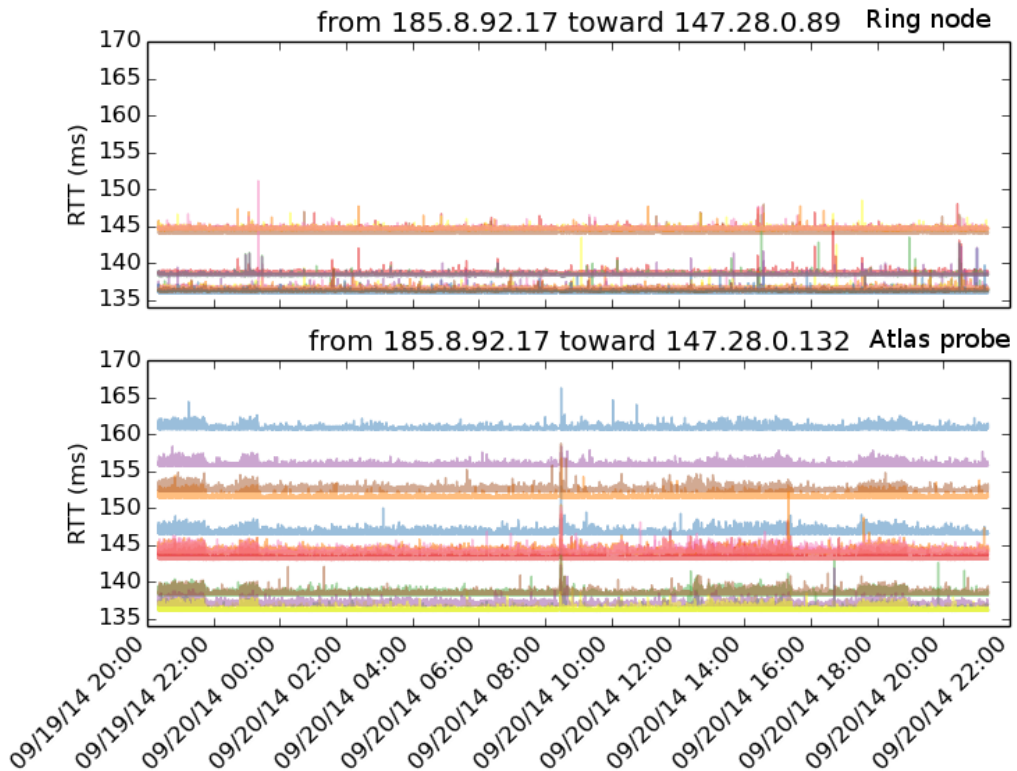

Pair	$\frac{\sum_{f=0}^{15} SD_s^d(f)}{16}$ (ms)	Global standard deviation (ms)
R2 - R4	2.7	4.4
R2 - A4	1.4	7.3

Average standard deviation δ

Because there is several sources S (11 Ring nodes) for one destination :
$\delta(d)$ is the Δ_s^d mean over all the sources for the destination d

$$\delta(d) = \frac{\sum^{\forall s \in S} \Delta_s^d}{|S|} = \frac{\sum^{\forall s \in S} \sum_{f=0}^{15} SD_s^d(f)}{16 * |S|}$$

Average Standard deviation δ

Half of the destinations perform better than the two servers regarding the RTT variability.

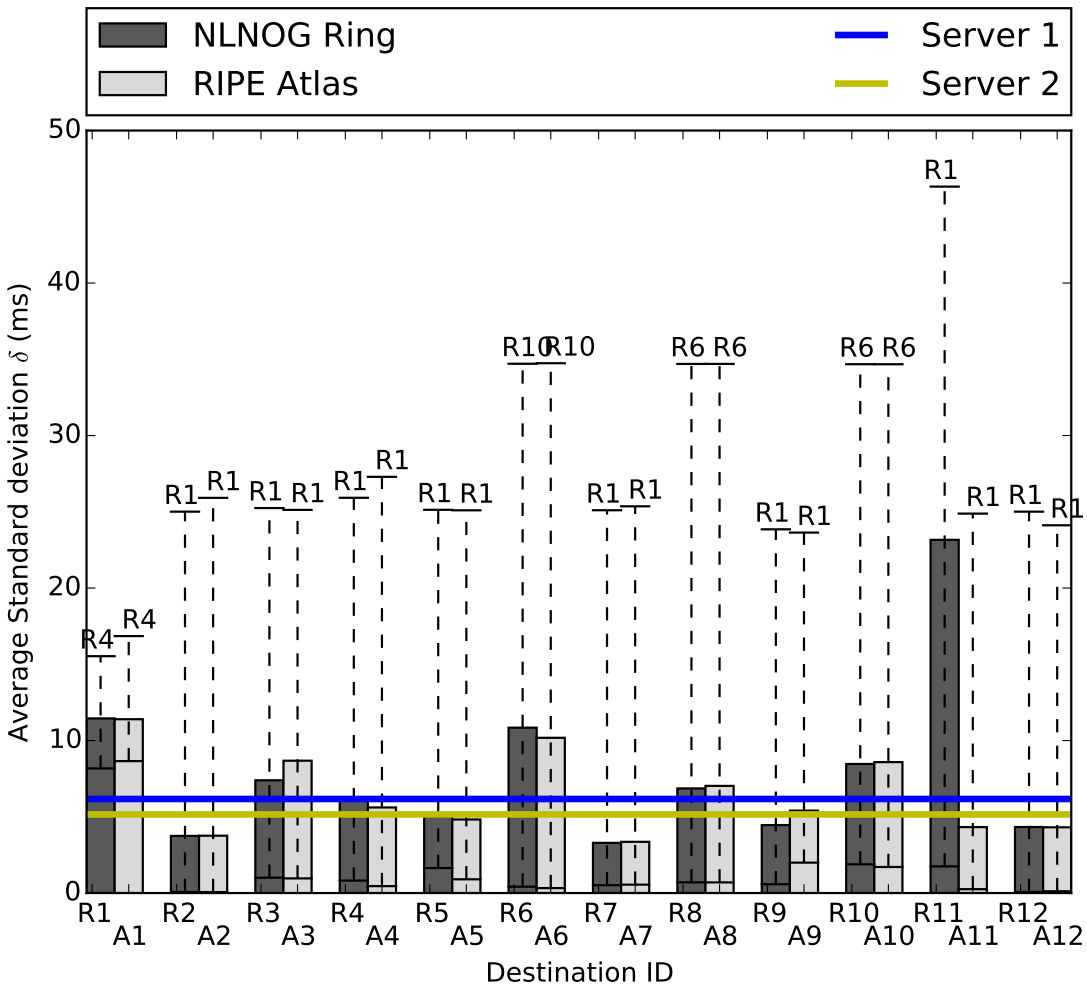

Packet loss percentage

RIPE Atlas probes don't lose more packets than the NL Ring nodes and the two servers

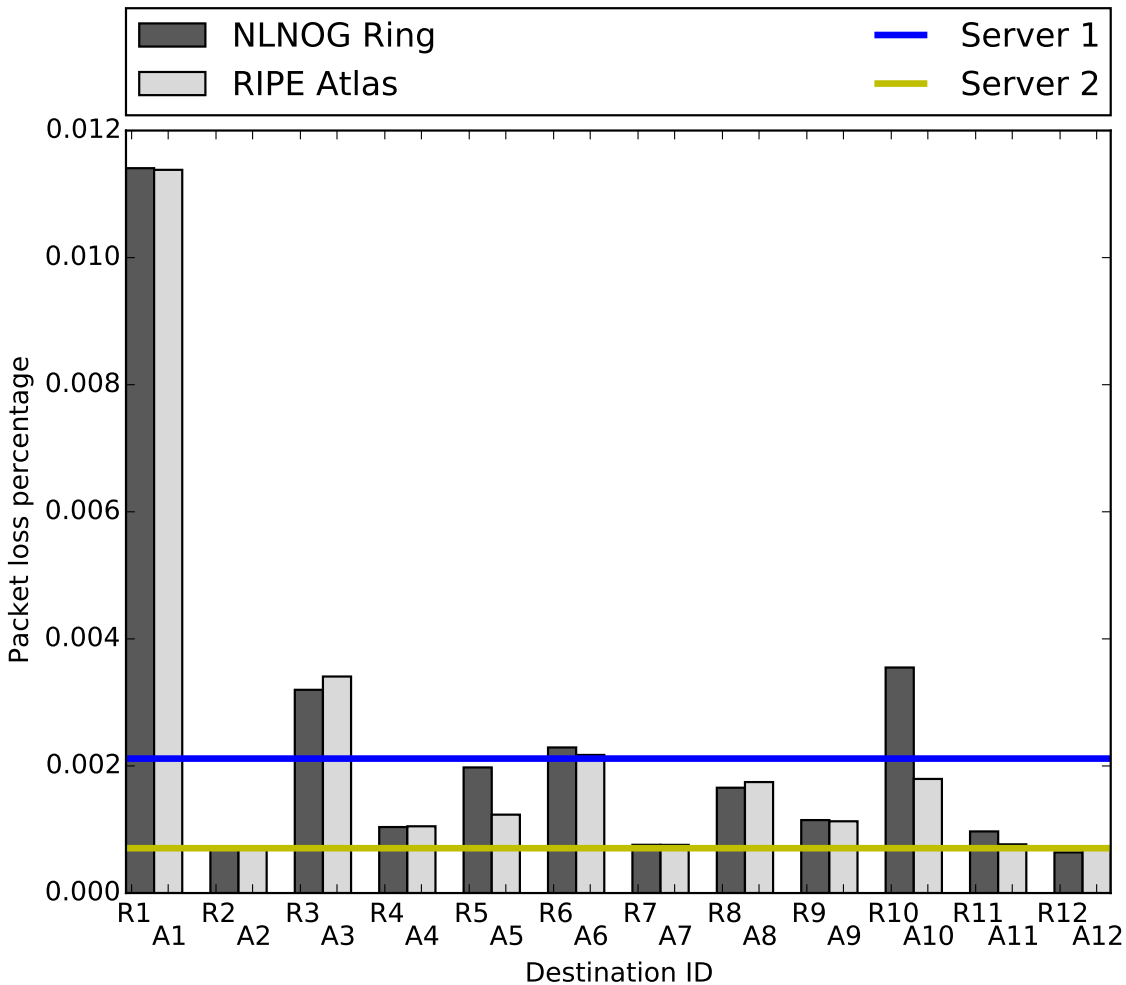

What if we increase the set of Atlas probes and Ring nodes up to 250 ?

New experiment with more Atlas probes and Ring nodes

Ring node X performs pings toward

- the 12 selected Ring nodes
- ∼250 random Ring nodes

New experiment with more Atlas probes and Ring nodes

Ring node X performs pings toward

- the 12 selected Ring nodes
- ∼250 random Ring nodes

Ring node Y performs pings toward

- the 12 selected Atlas probes
- ∼250 random Atlas probes

New experiment with more Atlas probes and Ring nodes

Ring node X performs pings toward

- the 12 selected Ring nodes
- ∼250 random Ring nodes

Ring node Y performs pings toward

- the 12 selected Atlas probes
- ∼250 random Atlas probes

Experiment Duration	1 week
Number of flow-ID	16
Theoretical frequency	1 ping every 20 seconds for each flow-ID
Actual frequency	1 ping every 21.5 seconds for each flow-ID

Average Standard deviation δ

- In both cases, the 12 selected have similar behavior than the 250
- 5 Atlas probes have 100% losses compared to only 2 Ring nodes

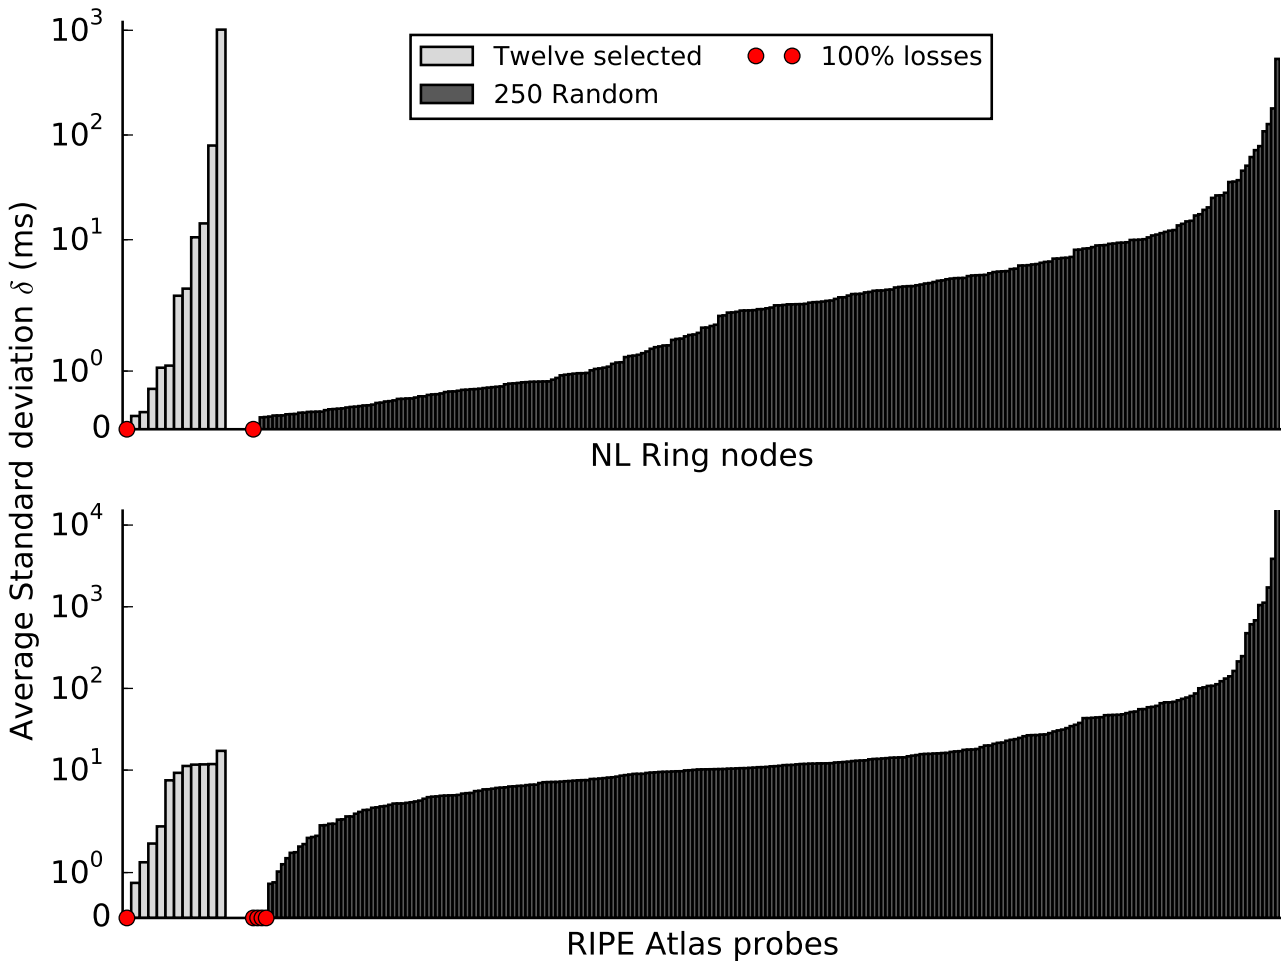

Packet loss percentage

- In both cases, the 12 selected have similar behavior than the 250
- 12 Ring nodes have no losses compared to only 1 Atlas probe
- 4 Ring nodes have 50% or more losses compared to 25 Atlas probes

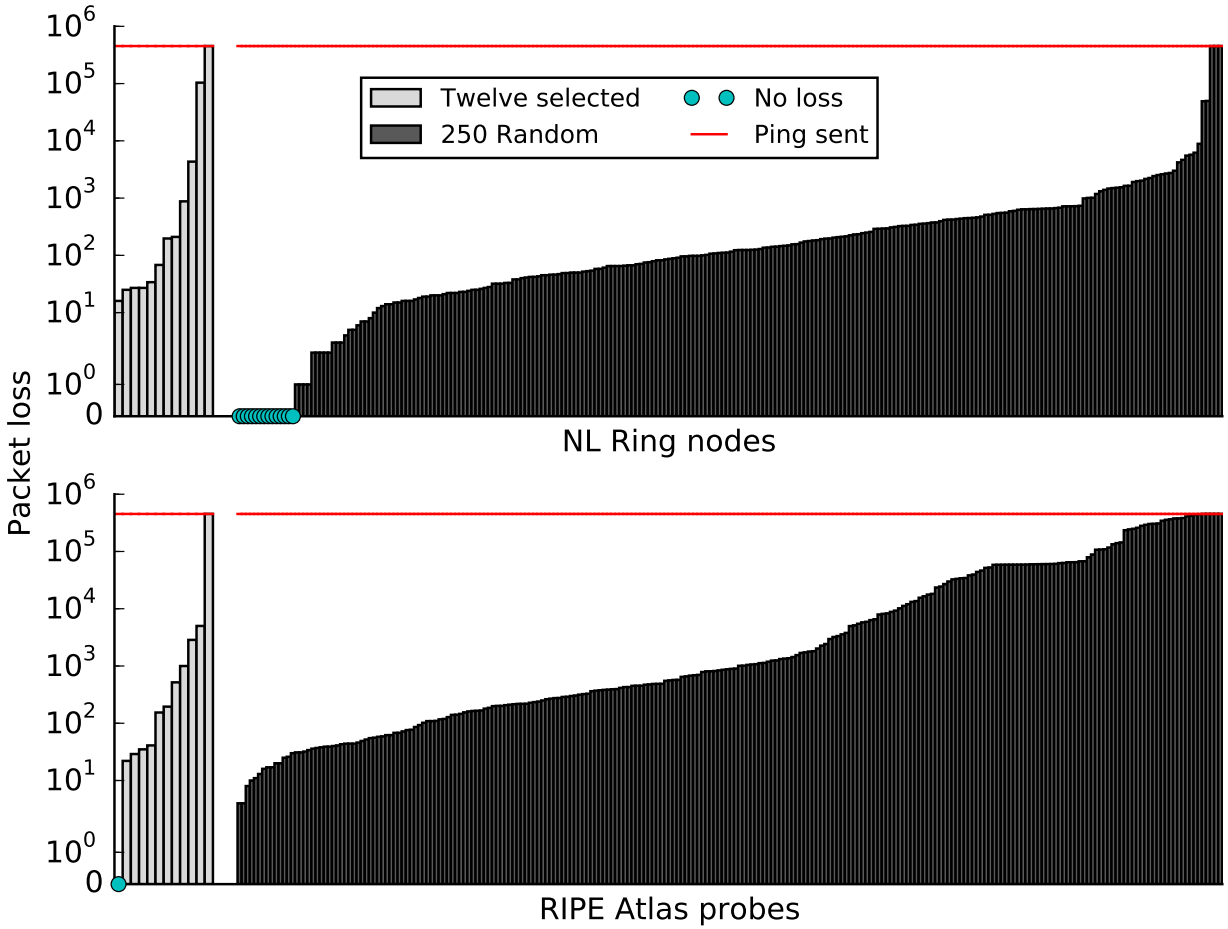

Further observations

Our method is able to efficiently catch internet events

Some events only affect a subset of the paths used

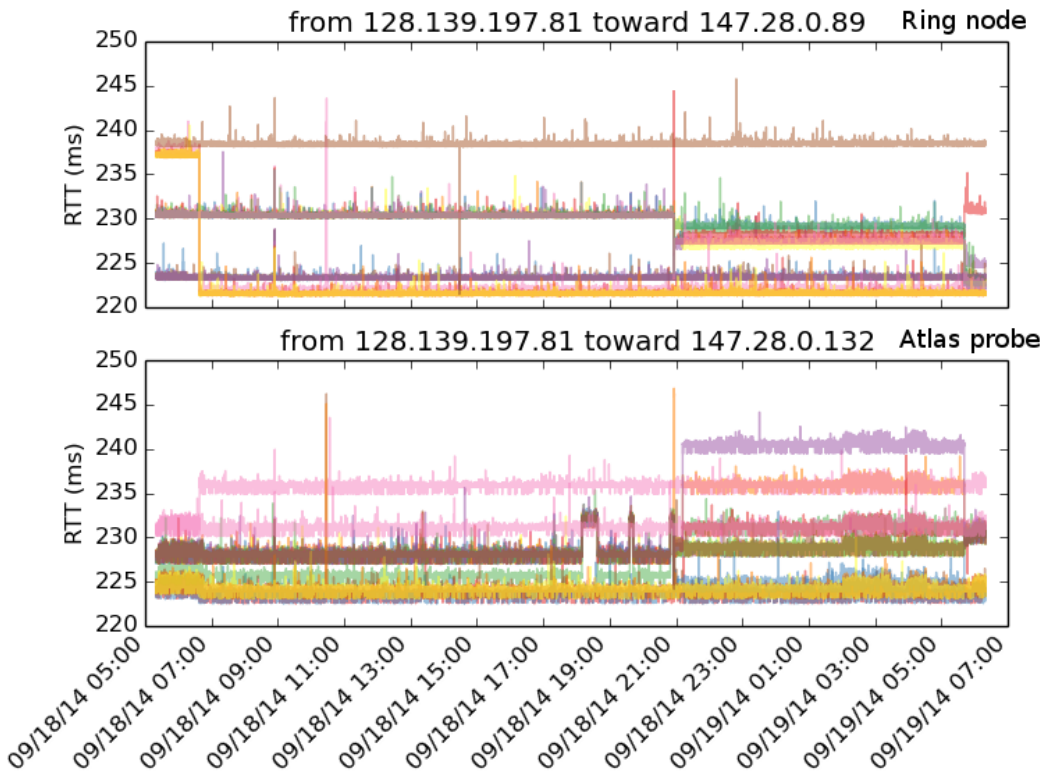

Good news

RIPE Atlas probes are recently able to perform ping measurements with different flow-ID

2014-08-25

New firmware release 4660

- Bugfix: fixed a bug in paris traceroute where the ICMP version would have the wrong paris id in outgoing packets

Conclusion

As destination, the RIPE Atlas probes we used are able to provide comparable results than the NL Ring nodes and the two BSD servers

Further work :

- Study the RIPE Atlas probes suitability when they are the sources
- Find the maximum ping frequency Atlas probes are able to cope as destination