Quantifying Interference between measurements on the RIPE Atlas platform

Thomas Holterbach, Cristel Pelsser, Randy Bush, Laurent Vanbever

► To cite this version:

Thomas Holterbach, Cristel Pelsser, Randy Bush, Laurent Vanbever. Quantifying Interference between measurements on the RIPE Atlas platform. Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication (CoRes 2016) CoRes 2016, May 2016, Bayonne, France. hal-01306027

HAL Id: hal-01306027

<https://hal.science/hal-01306027>

Submitted on 25 Apr 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantifying interference between measurements on the RIPE Atlas platform

Thomas Holterbach¹ Cristel Pelsser^{2 †} Randy Bush³ Laurent Vanbever^{1 ‡}

¹ETH Zürich

²University of Strasbourg & CNRS

³Internet Initiative Japan & Dragon Research Labs

Public measurement platforms composed of low-end hardware devices such as RIPE Atlas have gained significant traction in the research community. Such platforms are indeed particularly interesting as they provide Internet-wide measurement capabilities together with an ever growing set of measurement tools. To be scalable though, they allow for concurrent measurements between users. This paper answers a fundamental question for any platform user : Do measurements launched by others impact my results ? If so, what can I do about it ?

We measured the impact of multiple users running experiments in parallel on the RIPE Atlas platform. We found that overlapping measurements do interfere with each other. We found that increasing hardware CPU greatly helped in limiting interference on the measured delays.

Keywords: Measurement, Delay, RIPE Atlas

1 Introduction

Public measurement platforms composed of many low-end devices or *probes*, such as RIPE Atlas [RIPc], are increasingly used by researchers and network operators. To scale and be practical, measurement platforms schedule measurements in parallel, without providing feedback to the user. When put together with the limited hardware and software capabilities, this raises the question of measurement interferences. What is the impact of an increased load on the precision of measurements performed ? Do the measurements performed by one participant impact the results obtained by others ? If so, by how much ? Can this have an impact on previous research results ? [HPBV15] answers these questions empirically for the RIPE Atlas platform. Here we present of subset of these results.

As of April 2015, RIPE Atlas is composed of over 6,700 public probes scattered in 197 countries. Three versions of the probes exist, differing only by their hardware. Version 1 and version 2 are identical except for the amount of RAM they have. Both are Lantronix XPort Pro with a 167MHz CPU, 8MB or 16MB of RAM, respectively and a 16MB flash. The version 3 probe is a revamped TP-Link TL-MR3020 router with a 400MHz CPU, 32MB of RAM and a 4MB NAND. The v3 probes are therefore more powerful.

The number of v3 probes increased rapidly after they started to be distributed in 2013. While the number and proportion of v1 and v2 probes is decreasing, they remain non-negligible, accounting for 28.2% of the probes in April 2015.

RIPE Atlas provides different types of measurements. As of 2015, RIPE Atlas offers four[§] types of measures to its users : `ping`, `traceroute`, `DNS` and `SSL` [RIPb]. In RIPE Atlas, a measurement is defined by a type, a frequency and set of probes. It can therefore refer to an arbitrary number of individual measurements performed from multiple probes. Users can also provide a start date and an end date. If none

[†]The credits go to IIJ for supporting Thomas and Cristel's work.

[‡]We thank both the RIPE Atlas and the NLNOG RING support teams for accommodating our measurements and promptly replying to our questions.

[§]. HTTP measurements are possible but are restricted to researchers and other interested users on a case-by-case basis.

is provided, the measurement will start as soon as possible and has to be stopped manually. Measurements can be repeated or run only once (*one-off*). One-off measurements are near real-time if no start time is defined : users should expect results within 10 seconds [RIPb].

Atlas probes are popular sources of measurements and are increasingly used in research. Since its inception, Atlas performed almost 30 million individual measurements[¶]. V3 probes hosted 2/3 of the measurements, while v1 and v2 probes hosted the rest. In March 2015, the user who used the most credits spent 83.3 million credits [RIPa]. This is enough to perform more than 2,700,000 traceroutes. During the same month, the most used Atlas probe (a v1) provided 608,824 results [RIPa], one every 4 seconds. Finally, the number of concurrent measurements is important. As an illustration, the platform was executing 592,000 concurrent individual measurements when we collected these statistics. An increasing number of research papers use RIPE Atlas [MTS15, FPA15, FSC15, CJR⁺15].

Some Atlas probes are more used than others. Due to their geographical position, IPv6 capability or a NAT gateway, some probes are more attractive than others. As a result some probes are highly loaded.

2 Quantifying interference

We describe how we quantify interference between measurements performed on a RIPE Atlas probe. We take the perspective of one user λ and one probe ρ and measure the effects on the results reported by ρ to λ when ρ originates concurrent measurements. In particular, we look at changes in the delay reported by ρ when concurrent one-off traceroutes are originated on ρ . We launch an increasing number of one-off traceroutes on the probe. We use NL Ring nodes [NLN] as destinations of the pings sourced on ρ .

FIGURE 1: We load the probe with traceroutes and pings to distant targets while we measure the delay toward local destinations.

We want to focus on the behavior of the probe and avoid network interferences. For each experiment, we measure the delay between the tested Atlas probe and a colocated Ring node in the same LAN (*i.e.* there is no IP hop between them). Because packets between the Atlas probe and its colocated Ring node always stay in the same LAN, we prevent our measurements from being polluted by Internet variations (Figure 1). We obtained these pairs of colocated Atlas probe and Ring node by a traceroute campaign between each Ring node and Atlas probes in the same AS. The results depicted in Table 2 all come from measurements done between an Atlas probe and its colocated NL Ring node.

3 Decreased precision

We now use our methodology (§2) to measure the decrease in precision of delay-based measurements. We performed our measurements on multiple probes (at least two per version) to ensure conformity. As

[¶]. As a measurement may involve a large number of probes, the number of individual measurements is more representative of the load of the platform.

	50 th	95 th	stdev
(on : 100 traceroutes + 1.4 ping/s)			
v1	1.10 ms	7.30 ms	16.3 ms
v2	1.20 ms	7.70 ms	7.40 ms
v3	0.06 ms	0.10 ms	0.00 ms

FIGURE 2: Quantification of interferences for v1, v2 and v3 probes. On the left, the probe is loaded by sourcing 100 one-off traceroutes. We look at the impact of a load on the ping delay reported by the probe. The probe sends 1.4 ping/s. With more powerful hardware, v3 probes are less sensitive to load than v1 and v2. On the right, we see that delays measured *from* a v2 probe systematically increase when concurrent one-off traceroutes are launched on this probe.

their number is not negligible and their decrease in precision is serious, the next figure only focusses on v2 probes.

Delays measured *from* the probe increase when concurrent measurements are launched *on* it. We launched ping measurements from the Atlas probe and towards eight random Ring nodes plus the colocated Ring node. The ping rate towards each destination is 9 ping/min, averaging 1.4 ping/s over all destinations. We increase the load on the probe by launching successively 10, 25, 50, 100, 250, and 500 one-off traceroutes.

Figure 2 shows the impact of the concurrent one-off traceroutes on the delay measured from a v2 probe. The blue points are RTTs between the Atlas probe and its colocated Ring node, while the red points are RTTs between the Atlas probe and another Ring node. The gray areas are the periods when one-off traceroutes are running. The number above each gray area is the number of one-off traceroutes executed. To quantify the impact, we compare the median, 95th percentile, and standard deviation of the ping measurements before the one-off traceroutes (the white area preceding the gray area) and during the one-off traceroutes execution time (gray area). The difference is reported in Table 2.

Delays measured from the probe systematically increase when one-off traceroutes are performed. Starting 100 one-off traceroutes increases the median delay of the concurrent pings by more than 1 ms. For v1 and v2 Atlas probes, the standard deviation is seriously impacted : +16.3 ms (v1) and +7.4 ms (v2). Atlas probes v3 show less effect, the median is only increased by 0.06 ms while the standard deviation is not impacted ; this is due to v3 probes having more power.

Key points. We observed significant interferences on delay measurements for v1 and v2 probes. These probes compose 28% percent of the platform. An important portion (34%) of the public experiments available result from experiments on v1 and v2 probes.

Impact for researchers. RIPE Atlas makes publicly available all the results collected with the platform since its inception in 2010 [RIPc]. Researchers using these data should consider the impact of interferences. Especially for data collected before 2013—prior to v3 probes. We suggest researchers to be very careful when using publicly available delay measurements.

Provide feedback to users with a measurement confidence index. A fundamental problem with Atlas is that the user has no visibility on the concurrent load of the platform. For that, we argue that RIPE can return a “confidence index” along with each result. The index would be function of the platform concurrent load. High (resp. low) load would lead to low (resp. high) confidence.

4 Related Work

Other researchers have observed measurement interference and its impact on RTT. As an example, the effects virtualization can produce on measured delays have already been pointed out [WST09, SPBP06]. As a number of large-scale platforms use VMs [Pla, ML, NLN], tools such as [PP06] for PlanetLab, provide users with information about the state of these platforms and their nodes. In contrast, RIPE Atlas does not use virtualization and relies on a scheduler to share resources among users.

In [BES15], Bajpai *et al.* showed that RTTs from v1 and v2 probes to the first hop router are consistently higher than for v3 probes. They do not however study the relation between the measured delays and the load of the probes.

5 Conclusion

We presented the first measurement study of user-induced interferences on the RIPE Atlas platform. We found that measurements do interfere with each other. Delays reported from the probe increase and vary more when they compete with concurrent measurements.

Our findings also bring up new, non-trivial research questions : how can we design measurement platforms that provide more isolation between users, while still being efficient (*i.e.*, not requiring a global lock). We plan to explore this direction in the future.

Références

- [BES15] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons learned from using the ripe atlas platform for measurement research. *SIGCOMM Comput. Commun. Rev.*, 45(3):35–42, July 2015.
- [CJR⁺15] Danilo Cicalese, Diana Joumblatt, Dario Rossi, Marc-Olivier Buob, Jordan Auge, and Timur Friedman. A fistful of pings : Accurate and lightweight anycast enumeration and geolocation. In *IEEE INFOCOM*, 04/2015 2015.
- [FPA15] Roderick Fanou, Francois Pierre, and Emile Aben. On the Diversity of Interdomain Routing in Africa. In *PAM*, 2015.
- [FSC15] Pierdomenico Fiadino, Mirko Schiavone, and Pedro Casas. Vivisecting whatsapp in cellular networks : Servers, flows, and quality of experience. In *TMA*, 2015.
- [HPBV15] Thomas Holterbach, Cristel Pelsser, Randy Bush, and Laurent Vanbever. Quantifying interference between measurements on the ripe atlas platform. In *IMC*, 2015.
- [ML] M-Lab. Measurement lab. <http://www.measurementlab.net>.
- [MTS15] Guilherme Machado, Christos Tsiraras, and Burkhard Stiller. Schengen Routing : A Compliance Analysis. In *AIMS*, 2015.
- [NLN] NLNOG Ring. <https://ring.nlnog.net>.
- [Pla] Planetlab. Planetlab : An open platform for developing, deploying and accessing planetary-scale services. <http://planet-lab.org>.
- [PP06] KyoungSoo Park and Vivek S. Pai. Comon : A mostly-scalable monitoring system for planetlab. *SIGOPS Oper. Syst. Rev.*, 40(1):65–74, January 2006.
- [RIPa] Community Information, contributions, and hosts that stand out. <https://atlas.ripe.net/get-involved/community/>.
- [RIPb] RIPE Atlas - User-Defined Measurements. <https://atlas.ripe.net/docs/udm/>.
- [RIPc] RIPE NCC. RIPE Atlas. <https://atlas.ripe.net>.
- [SPBP06] Neil Spring, Larry Peterson, Andy Bavier, and Vivek Pai. Using planetlab for network research : Myths, realities, and best practices. *SIGOPS Oper. Syst. Rev.*, 2006.
- [WST09] Jon Whiteaker, Fabian Schneider, and Renata Teixeira. Explaining packet delays under virtualization. *SIGCOMM Comput. Commun. Rev.*, 41, 2009.