

# Is It Me? Understanding the Impact of the Home Network on End-to-End Measurements

Lucas Di Cioccio  
Thomson / UPMC Paris  
Universitas  
lucas.dicioccio@thomson.net

Renata Teixeira  
CNRS / UPMC Paris  
Universitas  
renata.teixeira@lip6.fr

Catherine Rosenberg  
University of Waterloo  
cath@ecemail.uwaterloo.ca

## ABSTRACT

End users are often curious to verify the performance of their network connection. Mostly, they accomplish this task with end-to-end active measurements. Unfortunately, the home network can affect end-to-end measurement and can lead to erroneous conclusions. This paper performs measurements in a controlled environment to understand the effect of cross traffic from home on the measurement. We show that although the impact from cross-traffic can be large, it is hard to detect it with active measurements in the home network.

## Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement techniques

## General Terms

Measurement, Experimentation

## Keywords

Cross-traffic detection, End-host monitoring, Home network

## 1. INTRODUCTION

Home users of Internet services have many reasons to monitor the performance of their Internet Service Providers (ISPs). First, the customer of an ISP needs a way to verify whether the quality of his connection corresponds to the one promised in his SLA (service level agreement, i.e., contract) with the ISP. Second, home users can share their network experiences to rank different ISPs. The need for monitoring ISPs is clearly expressed by a number of efforts that try to organize home users to monitor ISPs collectively (for example, the Grenouille project in France [4]).

All these efforts require deploying monitoring tools at the end-hosts. Monitoring tools can test several characteristics of the performance of an ISP. Grenouille measures transfer speeds and round-trip times (RTT) towards specific servers. Cross-traffic from the end-host could bias end-to-end measurements, i.e., lead to the wrong inference that the provider's network is performing poorly. To discard biased reports, Grenouille client verifies that the network card is idle before performing measurements.

Simply checking the end host's activity suffices if the computer is directly connected to the Internet. Unfortunately,




Figure 1: Scenario of Reference

it is rarely the case. Instead, users might own several computers and connect to a home gateway, which also provides TV or Voice over IP (VoIP) services, yielding a relatively complex home network. All these devices compete for the bandwidth provided by the ISPs. As a result, end-host measurements will share the access link with user traffic that cannot be seen directly from the end-host. Moreover, the home network has often a wireless (WiFi) component. The quality of WiFi connectivity varies widely according to the end-host position, concurrent sessions from other nodes and other external interferences. All these factors may bias the inference of ISP performances.

This paper addresses the problem of removing any measurement bias that might come from the home. First, we need to detect if there exists a local bias and then localize the source of the “measurement noise”. Second we need to understand its impact on the end-to-end (e2e) metric of interest. Indeed, even if the home environment seems very noisy locally, its impact on the e2e metric could be negligible and hence the metric could be representative of the performance of the e2e path.

## 2. APPROACH

By comparing the traffic on the WAN interface of the gateway with the traffic on its LAN interfaces, we can assess the presence of competing flows during a measurement cycle. There are two main strategies for retrieving traffic/flow information: by querying it passively if available, otherwise we must infer it from active measurements. A passive technique to query information from the gateway, would be to use a dedicated signalling protocol like UPNP [1]. UPNP can give indication about cross-traffic by looking at the quantity of data transferred by an interface, but it does not give information about other types of interferences (e.g, WiFi reception quality). So, we start by studying active measurements techniques.

There are many factors that can impact e2e measurements. Here we focus on events that happen inside the home. Let's illustrate with the example in Figure 1. Client *C*, connected to the gateway *G* via a home network performs a measurement to a server *S*. There is also a potential inter-

ferer  $I$  on the same LAN. We assume that  $G$  is the property of the ISP and that end-users cannot install software on it. Thus we cannot perform one way measurements like Pathload [3], for instance. Still, we assume some cooperation in the sense that the  $G$  and  $S$  will reply to a set of basic active measurement requests like ICMP echo requests (ping).

We have set up a testbed on which we reproduce some scenarios experienced by the monitoring end-hosts. Our testbed allows us to vary different parameters independently and study the impact of the home environment on e2e measurements. We next present a list of candidate parameters from Netprints [2] that we envisage to vary.

- In terms of network usage: We want to investigate if the presence of another active computer, or IPTV flows  $I$  on the same LAN, will impact the e2e measurements.
- Impact of a wireless LAN: We want to investigate if the fact that  $C$  is connected via WiFi has a major impact on the e2e measurement.
- Hardware/software versions: We want to investigate the importance of the hardware/software configurations of  $C$  or  $G$ .
- CPU usage: The CPU usage of  $C$ ,  $G$ , and  $S$  can have an impact on the measured values, especially when the operating systems of the devices are not real-time.

### 3. PRELIMINARY RESULTS

We report early results corresponding to measurements performed in a controlled environment. We setup a testbed to reproduce the scenario of reference of Figure 1. Our DSLAM (Zhone AuD8000) connects two gateways ( $G_A$ , Livebox and  $G_B$ , TelefonicaBox) and a PPP server (Linux 1.7GHz, 512Mb). The client  $C$  and the interferer  $I$  are two commodity laptops and are changed from one home to another. In all measurements, the gateways are synchronized with the DSLAM at 10/1Mbps, the DSLAM is connected to  $S$  with full-duplex Ethernet 100Mbps, and the laptops are both directly connected to the same gateway with Ethernet 100Mbps.

We want to check if we can detect precisely if the home network is busy by pinging  $G$ . We measure the RTT from  $C$  to  $G$  and  $S$  with 500 pings varying packet sizes, one ping every 200ms. To simulate an activity,  $I$  transfers a file to  $S$ . The upload consists of up to ten parallel UDP (constant bit rate 1Mbps) or TCP streams. We compare  $G_A$ , which has a regular Linux IP stack, to  $G_B$ , which implements real-time scheduling.

Our results (not shown for space constraints) show that the interfering traffic (be it UDP or TCP) has a big impact on the ping measurements to the server. Without interfering traffic, the RTTs to  $S$  are around 30ms for both types of gateways, whereas, with one or more interfering flows the RTTs to  $S$  is around 450ms when using  $G_A$  and around 100ms with  $G_B$ .

In Figure 2, we compare the cumulative distribution of the RTTs to the gateway for  $G_A$  and  $G_B$ . In absence of cross-traffic, the gateway takes less than 1ms to answer pings in 90% of the cases. In presence of 10 competing UDP flows, 70% of the RTTs are under 1ms with  $G_A$  and only 15% for  $G_B$ . Meaning that  $G_B$  was much slower to answer ping. There is almost no deviation under 5 parallel UDP flows for

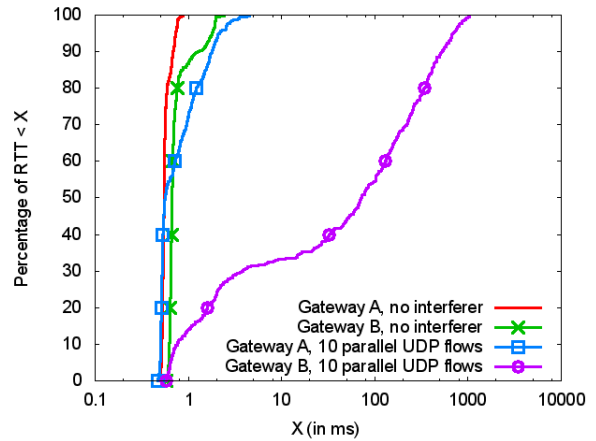


Figure 2: Comparison of RTTs from  $C$  to  $G$

both gateways. No deviation is observed on the RTT to  $G_A$  or  $G_B$  when the competing traffic is made of TCP streams. Varying the size of the ping payload does not change the shape of the distributions.

Our understanding is that the real-time gateway answers to pings with a small priority under load. We can leverage on such a deviation of RTT, to infer the presence of cross-traffic in the home network. A conservative approach would be to throw away every measurements where the RTT to the gateway is larger than a threshold that reflects the best case. We will also need to investigate if other conditions, that do not affect the e2e measurements, can lead to an increase of the RTT to the gateway.

To sum up, the home environment can have a significant impact on e2e measurements. However, pinging the gateway only detects large UDP transfers from another device in the home. We are currently investigating other techniques that can work for all types of cross-traffic.

### 4. FUTURE WORK

Our preliminary results illustrate the difficulty to detect the presence of cross traffic inside the home network by pinging the gateway. Thus, there is a real need for protocols like UPNP to query statistics on the network usage. These statistics need to be collected passively, because an ISP might enhance results of direct measurements. Our final goal is to propose methods to “clean” the data reported by the end-hosts by taking into account the activity of the home network at the time of the measurements.

### 5. ACKNOWLEDGMENTS

This work was supported by the ANR project C'Mon.

### 6. REFERENCES

- [1] Universal plug & play forum. <http://www.upnp.org/>.
- [2] B. Aggarwal et al. Netprints: Diagnosing home network misconfigurations using shared knowledge. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, 2007.
- [3] M. Jain and C. Dovrolis. Pathload: A measurement tool for end-to-end available bandwidth. In *Proceedings of Passive and Active Measurements (PAM) Workshop*, pages 14–25, 2002.
- [4] Projet Grenouille. <http://www.grenouille.com/>.