



HAL
open science

Measuring and characterizing home networks

Lucas Di Cioccio, Renata Teixeira, Catherine Rosenberg

► **To cite this version:**

Lucas Di Cioccio, Renata Teixeira, Catherine Rosenberg. Measuring and characterizing home networks. ACM SIGMETRICS, Jun 2012, Londres, United Kingdom. pp.383-384, 10.1145/2254756.2254804 . hal-00835387

HAL Id: hal-00835387

<https://hal.sorbonne-universite.fr/hal-00835387v1>

Submitted on 18 Jun 2013

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.

Measuring and Characterizing Home Networks

Lucas DiCioccio
Technicolor
UPMC Sorbonne Universites
lucas.dicioccio@technicolor.com

Renata Teixeira
CNRS
UPMC Sorbonne Universites
renata.teixeira@lip6.fr

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

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement techniques

Keywords

End-host monitoring, Home network

1. INTRODUCTION

A household today can have a variety of devices ranging from personal devices like laptops and smartphones to printers, game consoles, and media centers. These devices connect among themselves and to the Internet via a local-area network—a *home network*. A better understanding of home network performance and configuration should enable research in a number of areas, for example: home network troubleshooting and management techniques.

Unfortunately, there is little data on home networks. Most previous measurement efforts have focused on residential access links [3, 4]. The lack of data on home networks is partially due to the challenges of collecting home network data on a large scale. The vast majority of home networks are behind network-address translators, so a measurement point outside the home cannot measure the characteristics of the home network itself. Although it is feasible to deploy measurement points inside the homes of a few volunteers (as HomeMaestro did in the UK [2]), a small set of homes cannot offer a characterization of home networks at large. The effort necessary to recruit a large number of volunteers is a hurdle in itself.

We address this challenge by designing HomeNet Profiler.¹ Users run HomeNet Profiler from an end-system directly connected to their home network. HomeNet Profiler collects home network configuration and performance data. Because it runs from inside the home, HomeNet Profiler can perform a number of measurements to learn about the home network, for instance scan the local network for active devices. This paper describes the design and implementation of HomeNet Profiler. We also evaluate the performance of HomeNet Profiler – in particular how long HomeNet Profiler takes to run in real homes – and describe our initial data collection effort.

2. HomeNet Profiler DESIGN

Requirements.

The primary requirement for a home network data collection tool is that it *runs from inside the home*. The goal of measuring a large

¹HomeNet Profiler is available at <http://cmon.lip6.fr/hnp>.

diversity of home networks and the fact that it is not possible to collect data inside a user’s home without explicit user agreement and participation imposes additional requirements. HomeNet Profiler should be easy to run, even for users who are not tech-savvy. In particular, the tool should be portable and work for most home networks and end systems. Also, HomeNet Profiler should respect users’ privacy: users are unlikely to run a measurement tool inside their homes if the tool collects information that they consider private. Further, we are asking home users to do us a favor by allowing us to collect data inside their homes. Hence, HomeNet Profiler must not be a burden for users, especially in terms of time and resources. Finally, HomeNet Profiler should offer an incentive for participation. Indeed, if users can get something out of the tool, then we are more likely to get a larger number of participants.

System overview.

We implement HomeNet Profiler as a client/server application. The server hosts the HomeNet Profiler website, which users visit to download the HomeNet Profiler client. Once the client finishes loading in the user’s end-system, HomeNet Profiler starts in a separate window. Upon completion, the client sends all collected data to the server and redirects the web browser to a report page. Both the HomeNet Profiler client and the HomeNet Profiler website are available in French and English. A *run* refers to one execution of HomeNet Profiler in an end-system.

Measurement modules.

We select a broad range of measurements to learn as much as possible about the home network. At the same time, measurements should not take too long to execute, otherwise users might give up in the middle of the run. The HomeNet Profiler client has the following measurement modules.

Device scan: searches for active network interfaces. HomeNet Profiler sends UDP packets to all IP addresses in the end-system’s LAN and then reads the ARP cache of the end-system.

Service scan: queries two commonly-used protocols to advertise services in home electronics: Zeroconf and UPnP. We opt for querying these protocols instead of a port scan per device, which is intrusive and may take several minutes to complete.

Configuration of the UPnP gateway: collects (in cases where the home gateway has UPnP) the home gateway model, the upstream connectivity type and synchronization speeds (e.g., Cable 4Mbps upload, 1Mbps download), as well as traffic counters, which report the number of bytes and packets transferred.

WiFi networks: collects the list of access points found with a WiFi scan. For each access point we collect the network name, the base station identifier, the channel number, and the received signal strength indicator.

Netalyzr [3]: performs a number of tests related to the access network configuration, security, and performance. HomeNet Profiler invokes Netalyzr’s command-line client.

Computer configuration: collects the name and version of the operating system; the end-system’s network configuration, including the list of DNS servers and TCP parameters; and the list of network interfaces with the corresponding IP addresses.

Running applications: captures the list of processes running on the end-system as well as the list of TCP ports listening for incoming connections and open UDP ports. We also collect system services in MacOS and Windows.

Installed applications: lists the applications we find on the PATH environment variable. It may not reflect the full list of installed applications.

User survey: HomeNet Profiler complements and validates measurements with a user survey which runs in parallel to the measurement modules. We design the survey to be fast for users to complete and easy for us to interpret the results. Hence, all questions but one are multiple choice.

Finally, when HomeNet Profiler’s server receives the client reports, it maps the client’s public IP address to its geographical location and AS number using the Maxmind database.

User incentives and privacy.

As incentive to attract users, HomeNet Profiler presents a report after a measurement run.² The report provides some simple advice on improving home network performance. We also address privacy questions by anonymizing all personally-identifiable information using SHA1 and comply with French laws. At the end of a run, HomeNet Profiler only leaves a random identifier on the user’s machine to track multiple runs from the same end-system. Moreover, all measurements and the user survey are optional. As a result, users can adjust HomeNet Profiler to match their time and privacy requirements.

Pilot study.

We ran a pilot study with a small group of students, colleagues, and friends from France, Brazil, Canada, and the United States. The purpose of this study was to test measurement modules in different homes and operating systems and to adjust the survey questions. During the pilot, testers ran HomeNet Profiler 152 times from 47 different end-systems.

3. HomeNet Profiler DEPLOYMENT

Starting on April 4, 2011, we sent emails advertising HomeNet Profiler to family, friends, and colleagues as well as mailing lists of networking researchers. On April 18, 2011, we posted an announcement on Grenouille.com, a French website often accessed by people who want to monitor their ISP performance. We now evaluate how HomeNet Profiler performs in this deployment and which measurement modules are most popular.

Table 1 shows the total number of runs and end-systems that ran HomeNet Profiler between April 4, 2011 and March 21, 2012 and the number of runs per measurement module. Overall, users ran HomeNet Profiler 3634 times. Individual users ran HomeNet Profiler from 1 to 87 times, hence there are fewer end-systems than runs. Also, users skipped different measurement modules (users only select to run all measurement modules in 51% of the runs). Table 1 also shows the median and the 95th percentile duration of each module. The median time to complete the survey (3 minutes)

²For an example of report please visit <http://cmon.lip6.fr/hnp/example.html>.

Module	Runs	End systems	Homes	Duration	
				Median	95%
UPnP Gateway	3541	2692	2374	29 s.	39 s.
Netalyzr	3491	2668	2371	5 min.	7 min.
Computer conf.	3519	2687	2367	< 1 s.	3 s.
WiFi	3521	2671	2360	10 s.	11 s.
Service scan	3512	2678	2356	10 s.	21 s.
Device Scan	3533	2686	2355	1 s.	11 s.
Running apps.	3268	2485	2178	< 1 s.	4 s.
Installed apps.	3167	2402	2098	2 s.	11 s.
Survey	2142	2005	1736	3 min.	9 min.
Total	3634	2721	2432	6 min.	10 min.

Table 1: Popularity and duration of measurements.

is relatively short and reflects our desire to make the survey easy and quick to answer. The fact that different users select different measurement modules suggest that to attract a larger number of users, it is important to let users customize data collection.

Users may run HomeNet Profiler from enterprise or academic networks. Hence, we filter all runs from machine running outside home networks. We also ensure that we select only one run per home. In our technical report [1], we devise heuristics to perform these tasks. After applying these filtering heuristics, we infer that our data comes from a total of 2432 distinct homes. Users ran HomeNet Profiler from home networks in 46 countries and 210 different ASes. Table 2 shows the number of home networks and ASes we observe overall and for the top-five countries in our data. France dominates our data, mainly because of the announcement on the Grenouille.com website. This table shows that we obtain measurements from a fair number of different ASes in the top five countries.

Country	Homes	ASes
France	1699	33
USA	272	62
Canada	62	12
Brazil	59	10
Italy	57	6
Overall	2432	210

Table 2: Homes and ASes per country.

4. FUTURE WORK

Our future work will characterize home networks with HomeNet Profiler data. We will compare the diversity of home networks and their WiFi environment. HomeNet Profiler data will help us understand typical problems that may arise in home networks. We will complement HomeNet Profiler one-shot measurements with continuous measurements in a small number of homes.

5. ACKNOWLEDGEMENT

This work was supported by the Agence Nationale de la Recherche grant C’MON and by the European Community’s Seventh Framework Programme (FP7/2007-2013) no. 258378 (FIGARO).

6. REFERENCES

- [1] L. DiCioccio, R. Teixeira, and C. Rosenberg. Characterizing Home Networks With HomeNet Profiler. Technical Report CP-PRL-2011-09-0001, Technicolor, Sep 2011.
- [2] T. Karagiannis, E. Athanasopoulos, C. Gkantsidis, and P. Key. HomeMaestro: Order from Chaos in Home Networks. Technical Report MSR-TR-2008-84, MSR, Redmond, WA, USA, May 2008.
- [3] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the Edge Network. In *Proc. IMC*, Oct 2010.
- [4] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *Proc. IMC*, Oct 2009.