



Measuring Home Networks with HomeNet Profiler

Lucas Di Cioccio, Renata Teixeira, Catherine Rosenberg

► To cite this version:

Lucas Di Cioccio, Renata Teixeira, Catherine Rosenberg. Measuring Home Networks with HomeNet Profiler. Passive and Active Measurement Conference, Mar 2013, Hong Kong, Hong Kong SAR China. pp.176-186, 10.1007/978-3-642-36516-4_18 . hal-00835033

HAL Id: hal-00835033

<https://hal.sorbonne-universite.fr/hal-00835033>

Submitted on 17 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 Home Networks with HomeNet Profiler

Lucas DiCioccio^{1,2}, Renata Teixeira², and Catherine Rosenberg³

¹ Technicolor, ² CNRS and UPMC Sorbonne Universités, ³ University of Waterloo

Abstract This paper designs HomeNet Profiler, a software that runs on any computer connected inside a home network, to collect a wide range of measurements about home networks including the set of devices, the set of services (with UPnP and Zeroconf), and the characteristics of the WiFi environment. To attract a larger number of users, HomeNet Profiler runs one-shot measurements upon user demand. We validate this design choice against periodic measurements taken from six home networks. Data collected from these six homes and with HomeNet Profiler in more than 1,600 homes in France shed light on the diversity of devices that connect to home networks and of the WiFi neighborhood across home networks.

1 Introduction

The availability of cheap broadband Internet is popularizing Internet access from homes. A household today can have a variety of networked devices ranging from personal devices like laptops and smartphones to printers and media centers. These devices connect among themselves and to the Internet via a local-area network—the *home network*. Although there is increasing interest in home networking [1–3, 12, 19], there is yet little data on current home networks. Most prior work has focused on measuring and characterizing residential Internet access [4, 5, 9, 11, 13, 14, 17, 18]. The lack of data on home networks is partially due to the challenges of measuring home networks at large scales. The vast majority of home networks are behind network-address translators, so a device outside the home often cannot measure the characteristics of the home network itself. Some prior studies have deployed measurement points inside the homes of a few volunteers [10, 12, 15], but it is hard to get representative results from a few homes.

This paper designs HomeNet Profiler, a tool to measure home network configuration and performance (§2). Users run HomeNet Profiler on-demand from an end-system directly connected to their home network. HomeNet Profiler scans the local network for active devices and services advertised via protocols such as Universal Plug and Play (UPnP). It also measures the wireless environment per home. HomeNet Profiler incorporates features to help recruit a large number of volunteers. For example, it performs on-demand, one-time measurements, because many users feel uncomfortable downloading software that will run continuously in their machines. We validate this design choice with a testbed from which we periodically measure six homes in France (§3).

Between April 2011 and May 2012, users from 46 different countries ran HomeNet Profiler. This paper presents our analysis of home networks in France, where we have data from over 1,600 homes. We analyze devices present in home networks (§4) and the WiFi environment (§5). Our results show that in four out of five homes, users connect less than a dozen of devices to their home network. In addition, only a small number

of these devices, mainly home gateways, are active at any given time. We also observe that the density of the WiFi neighborhood varies considerably across homes.

The main contribution of this paper is the design and validation of HomeNet Profiler. Our initial experience shows that HomeNet Profiler was able to reach a large number of users (over 2,432 homes worldwide) and can hence get more representative results on home networks. Our testbed measurements validate our design choices and put HomeNet Profiler results into perspective by analyzing the dynamics of both the set of devices and the WiFi neighborhood, which we cannot study with HomeNet Profiler’s one-shot experiments. As more users run HomeNet Profiler, we plan to conduct a larger characterization study to shed light on home networks worldwide and help obtain better models to improve home networks and their WiFi environment.

2 Design

This section discusses the requirements of HomeNet Profiler as well as our design and implementation decisions.

2.1 Requirements

The primary requirement for a home network data collection tool is that it *runs from inside the home*. Measurements from outside the home cannot have visibility into the home network configuration and its devices. The goal of measuring a large diversity of home networks and the fact that it is not possible to collect data inside a user’s home without explicit user participation impose additional requirements such as:

Ease of use. The tool should be simple to run, even for users who are not tech-savvy.

Portability. The tool should run on all home networks and end-systems.

Respect users’ privacy. Users are unlikely to run a measurement tool inside their homes if the tool collects information that they consider private or any personally identifiable information. Our data collection effort complies with the rules of the French National Commission of Informatics and Freedom.¹

Light user commitment. We ask home users to do us a favor by allowing us to collect data inside their homes. We cannot ask users to commit too much time or resources without running the risk of reducing the number of users willing to participate.

Incentive for participation. Some users will run research tools altruistically. However, if users can get something out of the experiment, then we are more likely to get a larger number of participants.

2.2 Design and implementation decisions

The design requirements outlined in the previous section lead to some high-level design and implementation decisions.

First, HomeNet Profiler runs on *end-systems*. We considered deploying measurements on the home router/gateway or on one of the end-systems connected to the home

¹ Commission nationale de l’informatique et des libertés (CNIL): <http://www.cnil.fr/english/>.

network. Although some home users deploy routers with measurement capabilities [18], a hardware deployment has higher cost and more complicated logistics.

Second, inspired by the success of Netalyzer [13], HomeNet Profiler runs *one-shot* measurements on user demand. On one hand, long term, periodic measurements would give us a complete picture of home networks. On the other hand, users may be uncomfortable with installing a permanent software on their machine for privacy concerns and because of the possible impact on machine performance. We evaluate one-shot measurements against periodic measurements in § 4 and § 5.

Third, HomeNet Profiler is a *Java executable JAR*. We considered implementing HomeNet Profiler as a signed Java applet similar to prior work [13, 16], but it is hard to load system libraries such as the Windows Native WiFi interface from an applet and we need sudo rights on some Linux distributions (e.g., Archlinux), which is not possible from an applet. Instead, a Java executable JAR can collect the datasets we want and yet it is portable and simple for users to run.

Finally, HomeNet Profiler users are offered to complete a survey. This survey allows us to obtain information that would be hard to infer automatically from the measurements (such as finding devices which are turned off). To see the survey questions, we invite the reader to run HomeNet Profiler. As an incentive for users to run HomeNet Profiler, we output a detailed report of their home network.² Before the measurements begin, HomeNet Profiler lets the user select which measurements to execute. Hence, users who are uncomfortable with some measurements can still run HomeNet Profiler with a subset of the measurements.

System overview. We design HomeNet Profiler as a client-server application. The server hosts the HomeNet Profiler website, which users visit to fetch and run the HomeNet Profiler client. HomeNet Profiler starts in a separate window. Users then complete the survey while the measurement modules run in the background. Upon completion, the client sends all collected data to the server and redirects the web browser to the report page. When HomeNet Profiler exits, it leaves a random identifier on the user’s machine to track multiple runs from the same end-system. A *run* refers to one execution of HomeNet Profiler.

2.3 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 experiment. Our main goal is to discover the devices connected to the home network, the protocols they support (for instance, home devices are often expected to support UPnP and Zeroconf), and the services they provide as well as the network technologies connecting the home to the Internet and inside the home. We also characterize the WiFi neighborhood by measuring the quality of all visible WiFi networks. In addition to these direct measurements, we collect the configuration of the machine running HomeNet Profiler as well as the list of applications running on the machine. This extra information helps us interpret the results in case some configuration affects some of our measurements (e.g., a firewall or a VPN).

² For an example report refer to: <http://cmon.lip6.fr/hnp/example>

The HomeNet Profiler client has the following measurement modules.³

Device scan: searches the home network for active network devices. This module first populates the ARP cache by sending UDP packets on Port 9 (i.e., the discard port) to all IP addresses in the sub-network of the end-system. We force a 10 seconds timeout on the scan to avoid long delays when sub-networks are too large. Our data confirms that the vast majority of scans finishes before the timeout. This module then reads the ARP cache to collect the vendor ID (OUI) and the SHA1 hash of the MAC of each network interface on the LAN. If the associated IP address is private we also collect it, otherwise we just record the presence of a public IP.

WiFi scan: collects a list of access points found with one WiFi scan. For each access point we collect ESSID (the network name), the BSSID (the MAC address of the access point), the channel number, and the Received Signal Strength Indicator (RSSI). We anonymize the ESSIDs and BSSIDs. We distinguish between the *home WiFi*, which is the one the end-system is connected to, and *neighbor Wifis*. On MacOS, the airport command-line tool provides all this information. On Linux, we use *iwconfig* and *iwlist*. On Windows, we use the Win32 Native WiFi API. This library is not available on windows XP prior to SP3, so we can only collect WiFi information for newer Windows machines. We also observe that some Linux WiFi drivers only report information for the network the end-system is associated to.

Service scan: queries two protocols commonly-used to advertise services in home electronics: Zeroconf and UPnP. We opt for querying these protocols instead of a port scan per device because a port scan is intrusive and may take too much time.

Netalyzr [13]: performs a number of tests related to the access network configuration, security, and performance. At each execution, HomeNet Profiler downloads and runs the latest version of Netalyzr’s command-line client.

Configuration of the UPnP gateway: in cases where the home gateway supports UPnP, HomeNet Profiler collects the model of the gateway, the connection type, and the connection speed. It also tests traffic counters using UPnP queries.

Aside from these measurements taken from the client, when HomeNet Profiler’s server receives the collected measurements, it maps the client’s public IP address to its geographical location and AS number using the Maxmind database. We then discard the public IP address. HomeNet Profiler also sends meta-data such as the time taken by each module and whether HomeNet Profiler was running with sudo privileges.

This paper reports preliminary results on devices (§ 4) and WiFi (§ 5). We report on the Netalyzr and UPnP gateway configuration measurements in prior work [6].

3 Measurements

Testbed. In most cases, users run HomeNet Profiler once, but both the WiFi neighborhood and the devices connected to a home network vary over time. We thus complement HomeNet Profiler by instrumenting six different home networks in Paris. We installed laptops to colleagues from Technicolor and UPMC Sorbonne Universités. The households have between one and three members. Each laptop runs the WiFi scan module

³ To address privacy concerns and comply with French laws, we anonymize all personally-identifiable information using SHA1 hash.

every ten seconds using an Intel WiFi card. Every ten minutes, laptops also run the device scan module on an Ethernet adapter. We collect data from March 19, 2012 to July 31, 2012. These six homes are not representative of the population, but instrumenting a larger number of homes would represent a practical challenge. Nevertheless this testbed allows us to validate HomeNet Profiler and put collected data into perspective.

HomeNet Profiler data. We announced HomeNet Profiler by email to family, friends, colleagues, and mailing lists of networking researchers as well as through grenouille.com, a French website often accessed by people who want to monitor their ISP performance. Between April 2011 and May 2012, a total of 2,721 distinct end-systems ran HomeNet Profiler 3,634 times. Some users run HomeNet Profiler multiple times on the same end-system or from multiple end-systems in the same home. Users may also run HomeNet Profiler when they are not at home. For our analysis, we select a single representative run per home using two heuristics described in our technical report [7]. After applying these heuristics, we infer that our data comes from a total of 2,432 distinct homes. Users ran HomeNet Profiler from home networks in 46 countries and 210 different ASes. We detail the per-country and per-measurement modules breakdown of our dataset in a previous work [8]. This paper focuses on the 1,682 homes in France.

4 Set of devices in home networks

This section studies the set of devices that connects to home networks. We first use our testbed to analyze the dynamics of devices over time. Then, we analyze differences in number of devices across homes in the HomeNet Profiler data.

4.1 Evaluation of the completeness of device scans

Some devices may be disconnected from the home network at the time when users run HomeNet Profiler. Hence device scans in HomeNet Profiler are likely incomplete. We evaluate whether HomeNet Profiler benefits from extra device scans.

Repeated device scans in our testbed observe different sets of devices. Fig. 1 shows the presence of a given device during the four months of data collection. The x-axis is the time of each device scan. The y-axis represents individual devices measured in each home network (identified by their MAC address). We label the y-axis with the home-id and below each home-id, the number of devices observed in that home network during our measurements. We order devices per home based on their occurrence. The most prevalent device of all six homes is the home gateway. Note that there are gaps in the data collection because of maintenance or other measurement campaigns running on the same testbed. These gaps are easily identified by the vertical bars with no points per home. We ignore these gaps in the following discussion.

The number of devices measured per home in four months varies between 6 and 19 depending on the home. We ask the volunteers of the testbed to manually label each device observed over the whole data collection period. We divide devices into types: *home devices*, which are those that belong to members of the household; and *visitor devices*, which belong to friends who are just using the home network for a short




Figure 1: Observed devices per home network




Figure 2: Number of active devices versus the total number of devices per home. The number in the legend represents the number of homes included in the respective curve.

stay. The topmost devices of each home, those we only observed in a small fraction of the scans, correspond to visitor devices. We observe two types of home devices: always-on devices and on-off devices. *Always-on devices* are the ones users leave on all the time after the device first connects to the home network and until the device is decommissioned. These typically include home gateways and access points/routers. In Home-6 we also observe an IP printer and an IP security camera that were always on, and in Home-1 a network disk appears just before Week 5. *On-off devices* have prevalence between always-on home devices and visitor devices. We observe two types of on-off devices: personal mobile devices (such as laptops and smartphones) that leave the house with their owners; and devices that people turn on when needed (for example, a weighing scale and a gaming console).

We compute the fraction of the home devices observed in a single scan over the number of home devices. Given that HomeNet Profiler requires at least one user in the home to run the tool, we only count device scans when at least one laptop or desktop is on. Overall, we find that a single device scan only observes a small fraction of the home devices. For example, 92% of the scans with at least one laptop/desktop observe at most half of the home devices. Nevertheless, one single device scan captures all always-on devices more than 99.5% of the time. Hence, one-shot measurements are well-suited for studies that measure always-on devices such as home gateways [6].

We do not observe many more devices by aggregating the results of two consecutive scans (85% of pairs of scans would still observe at most half of the home devices). Only periodic measurements of the home network can observe all the home devices. We find that it takes approximately eight days on average (and a median of four days) to discover all home devices in the six homes we measured. To alleviate the lack of periodic measurements, HomeNet Profiler’s survey explicitly asks users to list the devices they typically connect to their home network.

4.2 Homes networks in France

We use the HomeNet Profiler data to study the devices that connect to home networks in France. We infer the *number of active devices* in a home network by counting MAC addresses present in the device scan. We remove devices with a MAC address belonging

to a virtual device.⁴ Given we only have one-shot measurements, we take the answers to the survey as ground truth for the *total number of devices*. Although users may misrepresent the number of devices in their home, we expect most users to answer this question correctly.

Fig. 2 shows the cumulative distribution of the number of active devices and the total number of devices across measured homes as well as the difference for homes where users selected both measurements (i.e., the number of total minus active devices for each home). The total number of devices per home ranges between 2 to 29, presenting a much wider spread than what we observe in our testbed. HomeNet Profiler’s larger scale measurements are essential to capture this wider spread of home network sizes. The range of the number of active devices, however, is smaller than that of the total number of devices. Approximately 75% of homes have at most four active devices during our measurements. This result is in agreement with our validation that shows that just a small fraction of home devices are on at any given time. The ‘difference’ curve confirms that many home devices are not connected when HomeNet Profiler runs.

The size of each household (i.e., the number of members living in a household) may explain the number of devices in a home network. However, some devices such as printers serve all members of a household. For the 400 homes for which users reported the size of their household, we find that the number of active devices and the size of the household have a Pearson correlation coefficient of only 0.18. The coefficient increases to 0.33 when considering the total number of devices and to 0.37 when considering only laptops and desktops. These results imply that the size of a household does have a moderate positive correlation with the total number of devices and hence the size of the household should be considered to model the total number of devices in the home.

5 WiFi neighborhood

This section characterizes the WiFi neighborhood as seen by end-systems at home. We first study the dynamics of the results of WiFi scans in our testbed to justify our choice of doing a single WiFi scan in HomeNet Profiler. We then study the WiFi neighborhood of French homes.

5.1 Accuracy of neighborhood characterization in one-shot measurements

It is important that HomeNet Profiler accurately measures WiFi neighbors with strong signals because their use could enhance home networks’ connectivity. The set of neighbor WiFis can vary considerably even in short time windows (of seconds), because lost WiFi beacons prevent us from inferring the presence of an ESSID-BSSID pair. We study the short-term dynamics of the WiFi neighborhood of each of the six homes in two-minute intervals; during each two-minute interval we perform 12 consecutive WiFi scans. We assume that the aggregate set of measured ESSID-BSSID pairs in the 12 scans represents the complete WiFi neighborhood during the two-minute interval.⁵ Then, we compute the *fraction of the WiFi neighborhood observed*, which is the number of ESSID-BSSID pairs observed in the first scan of a two-minute interval divided by the number of ESSID-BSSID pairs of the WiFi neighborhood in this interval.

⁴ In our dataset, the OUI for virtual machines are VMWare, Hyper-V, and Parallels.

⁵ It is practically impossible to get ground truth on the WiFi neighborhood.

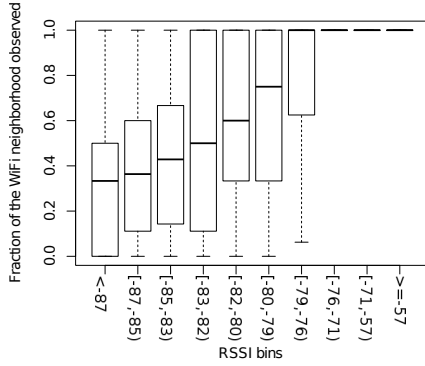


Figure 3: Fraction of the WiFi neighborhood observed with one scan for different RSSI bins




Figure 4: CDF of the number of neighbor ESSID-BSSID pairs

Intuitively, the probability of a WiFi scan to observe an ESSID-BSSID pair will be lower if the pair has low RSSI. To better understand this effect, we group the ESSID-BSSID pairs into ten *RSSI bins* based on the mean RSSI of each pair during a two-minute interval. We pick bin boundaries at every 10th-percentile of the distribution of mean RSSI per two-minute interval for all ESSID-BSSID pairs to ensure that every RSSI bin has 10% of the points. Fig. 3 shows the boxplot of the fraction of the WiFi neighborhood observed. The x-axis presents the RSSI bins (note that the x-axis is not linear). Boxes represent the inter-quartile range of the distribution of the fraction of the WiFi neighborhood observed for ESSID-BSSID pairs in a given RSSI bin; the solid line inside the box is the median, the whiskers represent the minimum and maximum values. The 802.11 standards do not specify units for RSSI and each vendor may use a different scale. All machines in our testbed have the same hardware and software. Hence we can aggregate RSSIs from the six machines in our testbed.

Fig. 3 confirms our intuition that ESSID-BSSID pairs with stronger signals are easier to observe. For example, the leftmost bin shows that half the time, a single WiFi scan observes no more than 34% of the ESSID-BSSID pairs with RSSI lower than -87, whereas a single scan is sufficient to observe all ESSID-BSSID pairs with RSSI higher than -71. The lower the RSSI, the more scans we need to observe all the ESSID-BSSID pairs. In summary, one scan is enough to collect all the ESSID-BSSID pairs with strong RSSI and to frequently get a large fraction of those with lower RSSI. This result validates HomeNet Profiler’s approach of performing a single WiFi scan to speed-up the data collection. ESSID-BSSID pairs with strong RSSI are more likely to interfere with the home WiFi and are also the ones that home users could potentially use for backup connectivity, for instance.

5.2 WiFi neighborhood in France

HomeNet Profiler successfully collects WiFi results in 1,131 homes in France. Some end-systems do not have a WiFi interface or lack support from the OS to run the WiFi scan. Because some WiFi access points may broadcast ESSID-BSSID pairs for more than one network (e.g., a guest network), we cannot tell which ESSID-BSSID pairs originate from a given WiFi access point. We consider that all ESSID-BSSID pairs other than the one the end-system is associated to are neighbor WiFi. In total, aggregating home and neighbor WiFi, we study 7,154 distinct ESSID-BSSIDs.

We focus on the 2.4 GHz band, which is the most used (96% of homes we measured). When two neighbor WiFis operate on the same or close channels, they might interfere. We say that two neighbor ESSID-BSSID pairs are *overlapping* if they are on channels where numbers differ by 4 or less. Channels 1, 6, and 11 are the non-overlapping channels in the 2.4 GHz band and hence are recommended for use. In our measurements, 18% of the ESSID-BSSID pairs operate on non recommended channels. We also notice that 39% of ESSID-BSSID pairs operate on Channel 11. We believe that some ISPs ship home gateways with hardcoded WiFi configuration.

WiFi neighborhoods are generally crowded in France. Fig. 4 plots the cumulative distribution of the number of neighbor ESSID-BSSID pairs across all measured homes. We present three distributions: for all neighbor WiFis; for ESSID-BSSID pairs that overlap with the home WiFi; and for ESSID-BSSID pairs on the same channel as the home WiFi. Overall, the number of ESSID-BSSID pairs of the WiFi neighborhood varies considerably across homes (from 1 to 52 neighbor WiFis) and more than 75% of homes have an overlapping WiFi neighbor. The actual number of WiFi neighbors is likely larger because HomeNet Profiler misses some WiFi neighbors with low RSSI.

The quality of the home WiFi also depends on the strength of the received signal. Since end-systems have different WiFi adapters, their RSSI measurements are not directly comparable. Thus we only compare RSSIs of different ESSID-BSSID pairs measured on the same end-system. Further, if the home access point broadcasts ESSID-BSSID pairs for a guest network, then their RSSI will be similar to the RSSI of the home WiFi. French ISPs offer country-wide community networks with well-known ESSIDs. After removing these ESSIDs, we find that in 13% of homes, the end-system has stronger RSSI to a neighbor WiFi that overlaps with the home WiFi. We have high confidence on this result because our testbed validation shows that we always observe WiFis with strong RSSI and here we are only studying the two strongest WiFis.

6 Conclusion

This paper designs HomeNet Profiler, a tool that home users run on an end-system to measure home networks. HomeNet Profiler scans the local network for active devices and services, observes the WiFi neighborhood, and complements measurements with a user survey. We design HomeNet Profiler as a one-shot measurement tool. Our testbed results show that one-shot measurements capture practically all always-on devices, but only a small fraction of on-off devices. As a result, HomeNet Profiler’s survey is an important complement to understand the full set of home devices at a large number of homes. In addition, the testbed results show that one-shot measurements are sufficient to capture all WiFi neighbors with strong signal and a significant fraction of neighbors with lower signal. WiFi neighbors with strong signal are more likely to interfere with the home WiFi or to be useful as backup links. Hence, HomeNet Profiler captures an essential part of the WiFi neighborhood. The biggest advantage of this one-shot approach is that it requires little effort/commitment from users and hence allow us to reach a large number of users. So far, users have run HomeNet Profiler from over 2,400 homes. Our analysis of 1,600 homes in France shows that the number of home devices vary considerably across homes and that only a small fraction of home devices are active at any given time. We also find that WiFi neighborhoods are crowded in France. We hope to

attract more users in other countries in the near future to perform a larger scale characterization. We also plan develop a service that will provide aggregate statistics based on HomeNet Profiler data to give an up-to-date view on home networks to the community.

References

1. K. L. Calvert, W. K. Edwards, N. Feamster, R. E. Grinter, Y. Deng, and X. Zhou. Instrumenting Home Networks. In *ACM SIGCOMM HomeNets Workshop*, 2010.
2. M. Chetty, R. Banks, R. Harper, T. Regan, A. Sellen, C. Gkantsidis, T. Karagiannis, and P. Key. Who's Hogging The Bandwidth?: The Consequences Of Revealing The Invisible In The Home. In *Proc. ACM CHI*, 2010.
3. M. Chetty, D. Halsem, A. Baird, U. Ofaha, B. Summer, and R. E. Grinter. Why Is My Internet Slow?: Making Network Speeds Visible. In *Proc. ACM CHI*, 2011.
4. D. R. Choffnes, F. E. Bustamante, and Z. Ge. Crowdsourcing Service-Level Network Event Monitoring. In *Proc. ACM SIGCOMM*, 2010.
5. D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. Biersack. Capacity Estimation of ADSL links. In *Proc. CoNEXT*, 2008.
6. L. DiCioccio, R. Teixeira, M. May, and C. Kreibich. Probe and Pray: Using UPnP for Home Network Measurements. In *Proc. PAM*, 2012.
7. L. DiCioccio, R. Teixeira, and C. Rosenberg. Characterizing Home Networks With HomeNet Profiler. Technical Report CP-PRL-2011-09-0001, Technicolor, 2011.
8. L. DiCioccio, R. Teixeira, and C. Rosenberg. Measuring and Characterizing Home Networks (Poster). In *Proc. ACM SIGMETRICS*, 2012.
9. M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu. Characterizing Residential Broadband Networks. In *Proc. IMC*, 2007.
10. C. Dixon, R. Mahajan, S. Agarwal, A. Brush, B. Lee, S. Saroiu, and V. Bahl. An Operating System for the Home. In *Proc. NSDI*, 2012.
11. D. Han, A. Agarwala, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan. Mark-and-Sweep: Getting the Inside Scoop on Neighborhood Networks. In *Proc. IMC*, 2008.
12. T. Karagiannis, E. Athanasopoulos, C. Gkantsidis, and P. Key. HomeMaestro: Order from Chaos in Home Networks. Technical Report MSR-TR-2008-84, MSR, 2008.
13. C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzer: Illuminating the Edge Network. In *Proc. IMC*, 2010.
14. G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *Proc. IMC*, 2009.
15. K. Papagiannaki, M. Yarvis, and W. S. Conner. Experimental Characterization of Home Wireless Networks and Design Implications. In *Proc. IEEE INFOCOM*, 2006.
16. A. Ritacco, C. Wills, and M. Claypool. How's my Network? A Java Approach to Home Network Measurement. In *ICCCN*, 2009.
17. M. Siekkinen, D. Collange, G. Urvoy-Keller, and E. Biersack. Performance Limitations of ADSL Users: A Case Study. In *Proc. PAM*, 2007.
18. S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet Performance: A View From the Gateway. In *Proc. ACM SIGCOMM*, 2011.
19. J. Yang and W. K. Edwards. A Study on Network Management Tools of Householders. In *ACM SIGCOMM HomeNets Workshop*, 2010.