Characterizing Home Wireless Performance: The Gateway View

Ioannis Pefkianakis[†], Henrik Lundgren[‡], Augustin Soule[‡], Jaideep Chandrashekar[‡], Pascal Le Guyadec[‡],
Christophe Diot[‡], Martin May[‡], Karel Van Doorselaer[*], Koen Van Oost[*]
HP Labs[†], Technicolor Research[‡], Technicolor Connected Home[*]
Email: *ioannis.pefkianakis@hp.com*[†], *firstname.lastname@technicolor.com*[‡*]

***Abstract*—In this paper, we analyze a large dataset of passive wireless measurements and obtain insights about wireless performance. We monitor 167 homes continuously for 4 months from the vantage point of the gateway, which allows us to capture all the activity on the home wireless network. We report on the makeup of the home wireless network, traffic activity, and performance characteristics. We find that in most homes, a small number of devices account for most of the observed traffic volume and the bulk of this traffic activity occurs in the evenings. Studying link performance, we find that overall, the vast majority of transmissions are carried out at high data rates and the wireless networks have good coverage. We find a small number of episodes where performance is poor; a few homes have a disproportionate number of poor performance reports. Investigating further, we observe that most of these are not caused by poor coverage (pointing to network interference). Our results significantly add to the understanding of home wireless networks and will help ISPs to understand their subscriber networks.**

I. Introduction

Home wireless networking problems are very hard to diagnose and troubleshoot for end-users. Unsurprisingly, this is one of the three most common issues reported at ISP helpdesks[1]. Thus, having a baseline understanding of home wireless performance and characterizing the common problems will greatly benefit both ISPs and their subscribers. However, a multitude of factors pose significant challenges to such an effort. Active WiFi devices from the subscribers' networks, as well as those on neighboring networks, compete for wireless capacity and interfere with each other. Furthermore, non-WiFi devices, such as microwave ovens, operate on the unlicensed frequency bands and cause RF interference. While previous studies have investigated home networks, the findings are based on small-scale deployments [2], [3], or else they do not mainly focus on the wireless network performance [4]–[6].

In this paper, our goal is to characterize wireless home networks in general and to extract a baseline understanding of the related performance characteristics. To this end, we collect and analyze a large dataset of wireless measurements taken from 167 subscribers of a large European ISP for a period of 4 months. Our resulting dataset with over 16 million measurements contains metrics extracted from commodity gateways under normal operation in subscriber homes, without requiring specialized hardware. In addition, our dataset includes both wireless link performance metrics and wireless traffic information, which allows us to relate WiFi link performance to WiFi device data usage. To our knowledge, this is the largest such dataset relating to wireless home performance to date.

Analyzing this dataset, we first perform a broad characterization of the wireless home networks; we study the penetration of particular WiFi technologies, high level device demographics, and the wireless neighborhood for each subscriber home, as well as each home's spatial and temporal wireless traffic characteristics. Subsequently, we qualitatively characterize link performance (good or bad) based on the reported PHY rate metric and quantify one class of WiFi performance problems. We correlate the performance metrics with various factors (WiFi technology, network density, etc.) to obtain insights on the impact of each on wireless link performance. Finally, to understand the causes of poor performance episodes, we examine whether poor wireless coverage is to blame (for these episodes). We summarize our key findings below.

We uncover significant diversity across the homes in our dataset; this in terms of 802.11 technology, number of devices, and usage. Overall, we find the WiFi interfaces of both gateways and devices to be under-utilized, accommodating very low traffic volumes. We also observe that when there is traffic, most of it is accounted for by a few stations (typically 2-3), and the bulk of this activity occurs in the evening.

We find that, across all the homes in our dataset and over the entire 4 month period, poor wireless performance is relatively uncommon. Roughly 7.6% out of over 6 million transmissions actually use the lowest PHY rates. However, a few homes have a disproportionate number of poor performance reports. By further analyzing coverage metrics during these transmissions, we rule out poor coverage as the main cause for the poor performance and we point to wireless interference. Surprisingly, and contrary to what is often conjectured, we find that the wireless link performance is *largely uncorrelated* with the wireless density in the neighborhood. Specifically, we observe almost no correlation between wireless link performance in the home and either (i) the number of neighboring wireless networks in range, or the (ii) cumulative effects of their associated RSSIs. Interestingly, commodity APs incorporate such metrics into their channel selection algorithms. We believe that our findings provide valuable insights to ISPs that could lead to better handling of customer complaints about wireless performance, and also suggest improvements in how gateway access points are configured and operated.

This paper is organized as follows. Section II discusses related work and Section III describes our data collection. Section IV broadly characterizes homes' WiFi networks. Sec-

Ioannis Pefkianakis carried out this work while at Technicolor.

[1]Based on informal conversations with two large ISPs.

978-1-4799-8381-0/15/$31.00 ©2015 IEEE

tions V and VI analyze the wireless link performance and study the root cause of poor performance for the homes of our dataset. Section VII concludes the paper.

II. Related Work

The performance of 802.11 wireless networks in campus, enterprise, urban, and rural environments has been extensively studied [11]–[14]. However, there are limited studies on 802.11 wireless home network performance, which collect feedback either from the *end-host*, or from *customized gateways*. End-host based measurement tools run as applications on the client side [6], [7], and typically suffer from a number of limitations. They cover only a single device's perspective and they provide only one-shot measurements, when the specific client device is connected to the network. Finally, they collect limited network feedback available at the application level.

A gateway-based approach on the other hand, enables continuous measurements at a fine time scale, and collects 802.11 MAC-layer feedback and higher layer traffic characteristics from all the WiFi devices in the home. So researchers have deployed customized APs in volunteers' homes [2]–[5]. While this approach allows to customize and instrument certain measurements at fine time scale, such efforts typically suffer from poor scalability and lack of generality. They mainly rely on recruited, technically-inclined volunteers [4], who need to obtain and install the customized gateway. For example in [3], authors study the wireless TCP throughput of 30 homes, where two-thirds of their APs are concentrated in two dense apartment buildings. Although recent work [4], [5] which leverages BISmark infrastructure [1] monitors a larger number of homes (64), it only seeks to identify if the wireless or the access network is the performance bottleneck. In contrast to existing work, we collect data directly from subscribers of a large ISP under normal service operation. This enables a larger scale deployment with very diverse WiFi environments, includes a more representative population sample, and still allows for fine feedback collection from all the devices connected to the gateway. Different from existing work, we study wireless performance in concert with data usage, which allows us to identify performance bottlenecks.

Finally, recent studies [8]–[10] design systems to identify 802.11 pathologies. These systems either require active measurements [8] which are typically disabled by the ISPs for performance reasons, or a dedicated WiFi radio to collect signal samples [9], or collaboration among APs [10], which are not available in commodity gateways. Although these systems can admittedly provide richer wireless feedback at finer time scales, our study uses metrics available in any commodity 802.11 driver, which are still provided at a sufficient time scale to broadly characterize wireless home network performance.

III. Dataset Description

In this paper, we analyze the wireless data collected from the Internet home gateways of 167 residential broadband subscribers, of a large European ISP. The most gateways of our deployment (71.3%) use fiber; 91% of the fiber plans provide 100/10 Mbps downstream/upstream speed, and for the rest it is 30/3 Mbps. The remaining gateways use ADSL technology with 24/1 Mbps downstream/upstream speeds. The subscribers are distributed over a large geographic area and span 10 cities. These 167 subscribers volunteered to be part of a project and were aware of the data collection. The results presented in this paper span a 4 month period (June - September, 2013) and contain over 16 million measurement samples. For ease of exposition, we use the terms "home" and "gateway" interchangeably, as also "user" and "subscriber". We next elaborate on our monitoring infrastructure and metrics.

A. Data Collection

Each subscriber in our deployment has almost identical (the differences are mainly in CPU speeds) home gateway platform with the following specifications: (i) ADSL2+ modem or fiber WAN access link, (ii) 4 ethernet ports, (iii) a WiFi access point enabled by a Broadcom 802.11b/g/n 2x2 radio with MIMO support. The WiFi radio supports 2 transmit and receive (2x2) RF chains (antennas) and both spatial diversity and spatial multiplexing MIMO modes. The 802.11 interface operates at 2.4GHz band and supports both 20MHz and 40MHz channels. Importantly, none of our gateways operates at the 5GHz band.

The data collection on the gateway is done by a lightweight OSGI module that periodically queries some data models exposed by the gateway software API and reports these metrics to a backend server. A particular set of design choices – polling frequency, data format, etc. – were dictated by very strict operational constraints; the monitoring should not have any negative impact on the subscribers, cause any instability on the gateways (generally resource constrained), or affect any of the services on the gateway (e.g., VoIP) adversely. A considerable amount of trial and error resulted in the final set of choices. The monitoring module generates reports every 30 seconds and sends them to the backend as JSON objects of 15-20KB size (including more metrics than what we analyze in this study).

B. Metrics

Our metrics are based on generic feedback provided by commodity 802.11 drivers. We use PHY rate as an indicator of the wireless link speed [16], [17] and RSSI to identify the root cause of low PHY rate [13]. Our gateways also report the achieved (actual) throughput generated by the user's running applications. We next elaborate on our metrics.

PHY rate: The PHY rate R, is the speed at which a station communicates with the gateway. The gateway reports the PHY rate of the last transmitted and received data frame, for each station associated with it, every 30 seconds. As we collect millions of records, this time granularity allows us to get statistically meaningful results. In commodity devices, the typical reasons for a rate change are: (i) a shift in the Received Signal Strength (RSS), (ii) frame losses caused by interference from hidden terminals or non-WiFi devices [13], [16], [17], (iii) rate sampling (rate adaptation (RA) algorithm evaluates wireless channel quality [18]). Typically, a high PHY rate is an indicator of good coverage and negligible interference. In our study, *we equate "good performance" with high PHY rates being used by the gateway*. As PHY rate is an upper bound of the wireless throughput (e.g., it does not account for losses, contention, 802.11 overheads), we also calculate an effective throughput metric, as we discuss later.

RSSI (dBm)	Expected PHY rate (Mbps)
[min, -88]	6.5
[-87, -86]	13
[-85, -83]	19.5
[-82, -81]	26
[-80, -75]	39
[-74, -73]	52
[-72, -71]	58.5
[-70, max]	65

TABLE I. RSSI - R_E MAPPING (802.11N 1X1, 20MHZ SETTING).

RSSI and Expected PHY rate (R_E): RSSI captures the signal strength between the station to the gateway. Per station RSSI values are measured at the gateway based on received data frames. The value exported is an average of the RSSIs of the frames received during the reporting interval (30 secs). The gateway does not report RSSIs for corrupted (e.g., due to interference) data frames, so RSSI measurements are unlikely to be affected by wireless interference. To identify the impact of RSSI on link performance, we calculate an *expected PHY rate* R_E, which is the best throughput transmission PHY rate for a given RSSI value. In general, the relationship between RSSI and R_E depends on 802.11 radio characteristics, and varies between vendors. We carry out detailed, controlled measurements in an RF shielded environment, using a wireless channel emulator from Azimuth Systems to obtain the RSSI-R_E mappings for both 802.11b/g and all the MIMO 802.11n configurations of our gateway platform. For example, we present the RSSI-R_E mappings for the 802.11n 1x1 setting in Table I. Notice that lower RSSIs correspond to lower R_E, indicating a performance degradation due to poor coverage. A caveat in converting RSSI to a rate is that, for the MIMO 802.11n setting, the effective rate also depends on the multi-path properties of the environment. In this case, we consider R_E as a coarse indicator.

RateGap: RSSI can remain stable, but the PHY rate can still vary due to interference losses [16], [17][2], or when the RA samples different rates in order to converge to a particular rate [18]. To identify rate variations which are not triggered by RSSI variations, we calculate the difference between the expected PHY rate R_E and the reported PHY rate R:

$$RateGap = Rate\_index\{R_E\} - Rate\_index\{R\} \tag{1}$$

We measure the rate gap in rate options (indexes) and not in Mbps, because the difference between two adjacent PHY rates in Mbps may highly vary among 802.11 technologies. Table II presents a high-level interpretation of relationship between PHY rate R and expected PHY rate R_E. If both R and R_E are high (upper-left), we expect good performance. If both are low (bottom-right), we expect poor performance attributed to poor coverage. On the other hand, if the expected PHY rate R_E is high, but the used PHY rate R is low (bottom-left) the result is a positive $RateGap$ and poor performance attributed to interference or RA dynamics. Conversely, if the expected PHY rate R_E is low, but the used PHY rate R is high, the result is a negative $RateGap$ that can be attributed to RA dynamics; typically the RA algorithm samples PHY rates higher than the available channel capacity to estimate channel quality, which often leads to poor performance. Note, it is out of the scope of this work to distinguish between interference and RA dynamics as the cause of non-zero $RateGap$. That would require feedback such as channel contention, RA parameters, collected at fine time scales, which are not available to us.

[2]We verify this phenomenon with controlled hidden-terminal experiments.

		Expected PHY Rate (R_E)	
		High	Low
PHY Rate (R)	High	good performance	poor performance (RA dynamics)
	Low	poor performance (interference/RA dynamics)	poor performance (poor coverage)

TABLE II. INTERPRETING PHY RATES R AND R_E.

We note two caveats in our approach that relate to RSSI. First, the reported RSSI is measured at the gateway (not the station) and consequently, the R_E and $RateGap$ metrics are computed only for uplink (from station to gateway). However, the PHY rate R is recorded at both directions. Second, the reported RSSI is an average for the reporting interval, while the PHY rate is reported only for the last frame; thus, the interpretation for the $RateGap$ is likely to be less accurate under dynamic wireless channels. However using the allan deviation methodology described in [12], we identify stable wireless channels (RSSI and PHY rate) in our home networks.

Effective throughput (T_E): is the maximum amount of bits that can be transmitted per time unit over wireless, given the 802.11 protocol overhead and the instantaneous channel quality (losses, contention) [16]. Our gateways do not expose 802.11 frame loss statistics (apart from excessive retries as we discuss next), and contention time with other devices on the same frequency. Thus we calculate an upper-bound to T_E, based on the PHY rate reported by our gateway and the 802.11 protocol overhead, which can be considered an 802.11 technology dependent constant.

Achieved throughput (T_A): The gateway periodically (30 secs) reports the actual throughput T_A (sent and received) for each station connected to it. T_A is the actual total number of bits over time, and it captures the actual demand on the wireless network. This metric depends on the application's data rate and is bounded by the wireless effective throughput or the access link throughput (for traffic exiting/entering the home). The throughput T_A is computed as the transmitted (or received) bits over a 30 second period, and it is reported by the gateway in Kbps. As gateway's driver reports only integer values (most commodity wireless drivers do not have floating-point support), we cannot capture T_A smaller than 1 Kbps. Consequently zero T_A values imply $T_A < 1Kbps$.

Additional metrics: The gateway reports each station's capability (e.g., 802.11b/g, 802.11n 1x1 or 2x2) and transmit/receive traffic counters in bytes. It also reports the 802.11 *excessive retries* (i.e., the frame transmissions where no ACK was received after a maximum number of retransmissions) for each station. Excessive retries are reported only in downlink (gateway to station), and we use them as an additional performance indicator. Finally, the gateway's automatic channel selection (ACS) reports a list of detected neighboring SSID's and the associated RSSIs, at a longer timescale (every 2 hours) than the above metrics.

Summary: In a nutshell, we leverage PHY rate (and effective throughput) metrics to identify performance bottlenecks, and achieved throughput metric to capture traffic demand. The limited feedback reported by our gateways prevent us from estimating a more realistic application-level throughput. However, PHY rate is still an important performance proxy,

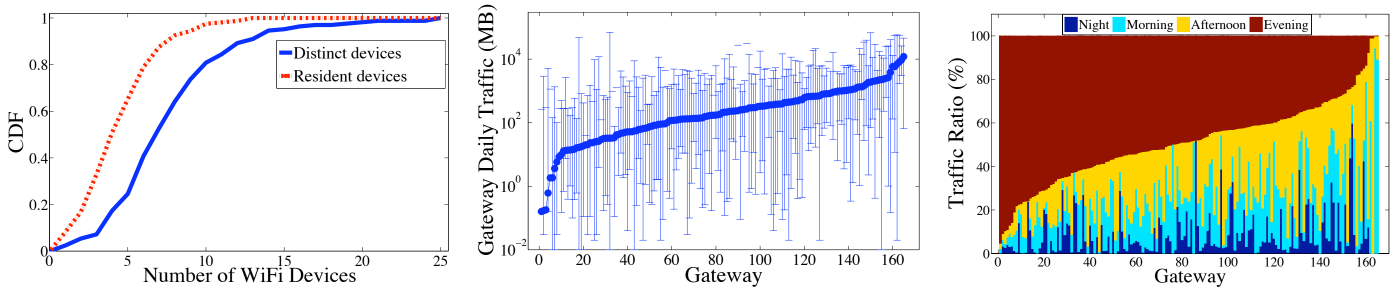

(a) Total and resident WiFi devices.
(b) Daily wireless traffic across homes.
(c) Traffic at different times of day.

Fig. 1. Device population and traffic dynamics.

802.11 device distribution	.11b 0.45%	.11g 42.55%	.11n 1x1 45.48%	.11n 2x2 11.52%
Homes with only 802.11x devices	.11b/g 28.1%	.11n 8.4%	.11b/g/n 63.5%	

TABLE III. HOME WIFI TECHNOLOGY OVERVIEW.

which can capture poor performance instances [13], [16], [17]. A limiting factor to identify the root cause of performance bottlenecks is the lack of contention and interference feedback from external sources (i.e., WiFi devices connected to neighbouring APs and non-WiFi wireless devices). However, we still leverage RSSI to detect wireless coverage problems.

IV. HOME WIRELESS ENVIRONMENT AND TRAFFIC CHARACTERIZATION

In this section, we report on wireless devices in each home, their 802.11 technologies and their WiFi neighbourhood (Section IV-A). We also study the wireless traffic activity of the homes of our deployment (Section IV-B).

A. Network Characterization

We identify a total of 1328 distinct devices (by MAC address) associated to the gateways of our deployment. Some devices are only observed for a few days over the entire dataset; others are seen every day. To differentiate these *transient devices* from those which are owned and operated (regularly) by people in the homes in our dataset, we use a threshold of 7 days. In other words, devices which are observed on fewer than 7 days are labeled as *transient*, while devices appearing on more than 7 days are termed *resident*. Transient devices may be associated with visitors or they may be turned on very infrequently (e.g., wireless printers). Figure 1(a) plots the distribution of the total devices recorded in each home, and also the distribution of resident devices. The total device count for a home varies from 1 to 25. The median number of devices found in a home is 7, including transient devices. If we consider only the resident population (dotted curve), we see that the median home has 4 resident devices, 12% of the homes have more than 8 devices, with 13 being the largest value. In the rest of our analysis, we include both resident and transient devices unless mentioned otherwise.

Home WiFi technology: Wireless throughput is tied to 802.11 technology; newer technologies (e.g., 802.11n) support higher transmission rates. As can be seen in Table III, in our dataset we find a high penetration of newer high-throughput 802.11n technology across the homes. Specifically, 57% of the 1328 wireless devices are 802.11n, and roughly 3 out of 4 homes (72%) have at least one 802.11n device. However, only 11.5% of the total devices support the high speed 802.11n MIMO feature. We also see plenty of legacy technology; 42.6% of the devices use 802.11g and a tiny fraction (0.45%) use the older 802.11b technology. To allow a high level comparison of performance, we further classify homes into *high* and *low* speed (homes) based on the device technologies. In the former, we exclusively see 802.11n, while in the latter, we only observe legacy technologies (802.11b/g). Across our dataset, we find 8.4% (28.1%) of the homes to be high (low) speed respectively; the rest have a combination of devices that span these technologies. In general, we expect the high speed homes to have better network performance characteristics than the low speed homes. However, wireless dynamics can significantly compromise the network performance gains, and may inverse this relationship (cf. Section V-B).

Home WiFi neighborhood: We next leverage the feedback from our gateways' ACS algorithm to analyze their WiFi neighbourhoods. Overall, we uncover diverse and dense wireless environments around each of our gateways. The number of neighboring SSIDs around each gateway varies between 1 (perhaps a non-urban dweller) and 78 (possibly city apartment dweller), with the median value being 17. This density is potentially significant since a large number of neighboring networks is believed to adversely affect performance due to interference. We study this in detail in Section VI.

We also find that the RSSIs from each of the SSIDs vary significantly. Specifically, the average RSSI from a neighboring AP to a home gateway of our deployment varies from -10.7 dBm to -91 dBm, implying different proximity of neighboring APs to the home gateway. -91 dBm corresponds to a signal that can barely be sensed, while -10.7 dBm is a very strong signal and likely originates from a device in direct proximity of our gateway, such as a WiFi extender inside the home. Finally, the neighboring SSIDs of our gateways are likely to implement ACS algorithms, as we observe them to appear at different channels over time. They mostly use channels 1, 6, and 11, which are the non-overlapping at 2.4GHz.

B. Traffic Characterization

We next examine the wireless traffic activity in space and time, for the homes and devices in our dataset.

Spatial-temporal traffic characteristics: We first present the daily traffic, sum of uplink and downlink, for each home in Figure 1(b). Each point on the x-axis represents a single gateway and the y-axis indicates the range (min-max) over all the days in our dataset. Each vertical line is also annotated with

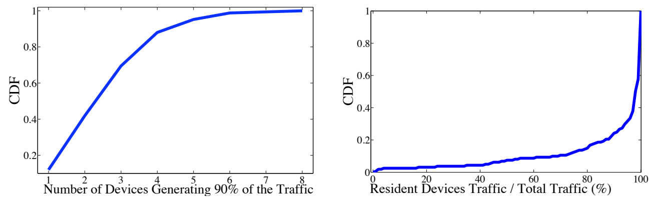

(a) Devices contribution to traffic. (b) Traffic from resident devices.
Fig. 2. Wireless traffic correlation with devices.

a marker corresponding to the median daily traffic volume. We observe that there are significant variations across different homes as well as within a single home. The median daily traffic volume can vary up to four orders of magnitude across homes (the y-axis is in log-scale); from less than 1 MB to more than 10 GB. For a single home, the daily (wireless) traffic volume can vary between a few MB up to to 67 GB.

Next, we study how time of day impacts the observed traffic volume. To this end, we partition the 24 hours of a day into four periods: morning [6am-12pm), afternoon [12pm-6pm), evening [6pm-12am) and night [12am-6am) times, and compute the daily fraction of traffic seen in each period (for each day and for each gateway). Figure 1(c) shows the fraction of traffic generated at each time period, across the homes in our dataset. We find that for all but a few gateways, the most traffic is generated during the evening. Specifically, for 73.3% of the gateways, the evening period records the largest traffic fraction. For the remaining gateways, 12.7%, 10.3% and 3.7% generate most traffic during afternoon, morning and night, respectively. These differences can be attributed to variations in device usage patterns and applications used in each home.

Number of devices and traffic: We next seek to understand which devices contribute the most to a home's traffic volume. We therefore plot in Figure 2(a), the minimum number of devices that account for at least 90% of the observed traffic (on the x-axis) and the corresponding fraction of gateways on the y-axis. We see that only a small number of devices generate 90% of the traffic; in roughly 70% of homes, only 3 devices account for this traffic. In less than 15% of the homes, more than 4 devices contribute to 90% of the total wireless traffic. We further distinguish between transient and resident devices. Figure 2(b) plots the fraction of a gateway's wireless traffic that is attributed to resident devices, for all gateways. We observe that, for the vast majority of the homes (80%), at least 87% of the wireless traffic is generated by resident devices. While we do not have information about the device types or their usage, we speculate that this traffic concentration (over a small number of devices) hints at individuals in the households owning multiple devices, some of which are used infrequently (and do not generate significant traffic).

Traffic symmetry: It is widely presumed that downlink traffic (AP to stations) is considerably larger than uplink traffic in 802.11 networks [11] because of the "pull" nature of client-server applications (web surfing, email, etc.). We indeed see in our dataset that, the overall downlink traffic dominates (over uplink), with up to two orders of magnitude. However, for 7.2% of the stations and 4.8% of the homes we see this condition to be reversed; gateways and stations can generate up to one order of magnitude higher uplink than downlink traffic. We speculate that services such as cloud syncing (e.g., Dropbox), or uploading media files to social networks could contribute to the cases where we observed larger upload traffic.

Takeaways: Our results show a high penetration of high-speed 802.11n devices and very diverse wireless home networks, both in terms of device population and WiFi neighborhood. Although there is noteworthy diversity in the number of home devices, most of the traffic is typically generated by few devices, during the evening. Traffic activity also varies in time, and gateways can remain often highly under-utilized.

V. Link Performance Characterization

In this section, we study the wireless link performance across individual homes and 802.11 technologies. We further compare effective and achieved throughput metrics.

A. *Overall Performance Distribution*

We start by analyzing the PHY rates reported in individual reports, over all the gateways and devices. Recall that we equate high PHY rates with good wireless link performance. In Figure 3(a) we show the PHY rate values for 802.11n 1x1 stations (and for 802.11g stations in 3(b)); this is over all the reports in the dataset. Note that 802.11n 1x1 and 802.11g stations account for 88% of all the devices in our dataset. In Fig. 3(a), we see that at least 64% (and 53%) of the gateways' transmissions (and receptions) are carried out using the *highest PHY rate*, i.e., 65 Mbps. This observation holds also in the 802.11g case; 70.7% and 58.7% of the activity is associated with the 2 highest rates (54 and 48 Mbps). We also see the same pattern for 802.11b and 802.11n 2x2 devices[3].

Thus, we see the vast majority of frame transmissions and receptions at high PHY rates. The average PHY rates are close to the technology limit for each 802.11 technology, as shown in the first three rows of Table IV. This leads to the conclusion that in our dataset, good wireless link performance is the norm and there are few limiting factors for link performance. However, as can be seen in Figures 3(a)-3(b), there is a small number of samples at lower PHY rates. Overall, we find 7.6% of the total transmissions at PHY rates of 6.5 Mbps or lower, from all the records in our dataset. In the next section, we analyze these in detail and elaborate on their causes.

While PHY rate governs the number of bits that can be transmitted over the link, this includes overhead from the 802.11 protocols. The effective throughput (T_E), derived from PHY rate by accounting for protocol overheads (§III-B), more closely approximates what an application on a wireless device is likely to experience. We compute this metric by assuming 1470-byte packets and specific to 802.11n MPDU aggregation along with A-MPDUs of 16 (27) frames for 802.11n 1x1 (and 2x2) stations.[4] Table IV (bottom) presents the T_E statistics for each 802.11 technology. We observe the average T_E to be as high as 83.5 Mbps, 48.8 Mbps or 26.7 Mbps (802.11n 2x2, 1x1 and 802.11g, respectively). These throughputs are higher than the capacity of currently deployed access link technologies (ADSL2+ and DOCSIS 2.0 [1]) and popular application bandwidth requirements (e.g., Netflix). Thus, in

[3]The 802.11n 2x2 rate distribution is slightly lower than the other technologies, mainly since the samples are spread over its larger set of rate options.

[4]It is the average framing observed in UDP packet train experiments on Atheros NICs. We leave the evaluation of various packet sizes as future work.

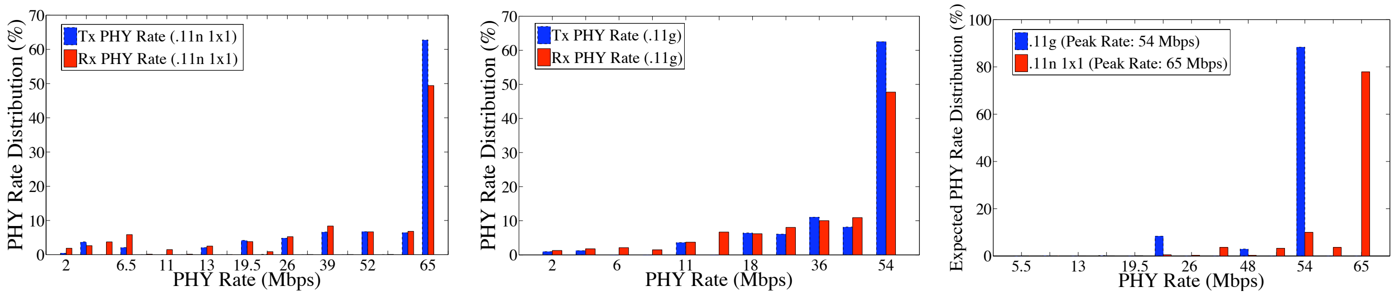

(a) Transmit/Receive PHY rates (.11n 1x1). (b) Transmit/Receive PHY rates (.11g). (c) Expected Rx PHY rates (.11g/n).

Fig. 3. PHY rate and Expected PHY rate distributions aggregated over all stations and gateways.

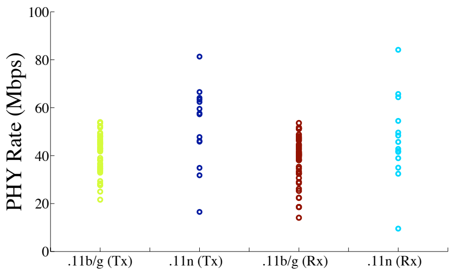

Fig. 4. Performance variations among homes of different 802.11 technologies.

these cases it is likely that applications will be performance limited by the access link rather than the wireless channel.

Our analysis of the excessive retries metric which is an indicator of severe performance problems, further corroborates that performance is overall good. We observe that the average number of excessive retries in a 30-second time bin across all the reports of our dataset, is small (2.9). The median retries is zero. The average loss ratio calculated from the excessive retries is only 1.9% (note this does not include losses from frames which were not excessively retried).

PHY Rate and Exp.Rate R_E (Mbps)	Station technology .11b	.11g	.11n 1x1	.11n 2x2
Max Supported PHY Rate	11	54	65	130
Mean Tx PHY Rate	11	44.75	53.35	91.75
Mean Rx PHY Rate	8.75	39.47	46.59	71.87
Mean Exp. Rx Rate R_E	11.0	51.2	61.76	111.99
Effective throughput T_E (Mbps)	**.11b**	**.11g**	**.11n 1x1**	**.11n 2x2**
Mean Tx T_E	6	26.70	48.76	83.54
Mean Rx T_E	4.75	23.98	42.59	65.86

TABLE IV. AVERAGE WIRELESS LINK PERFORMANCE.

Link asymmetry: Our results show that the average downlink (Tx) PHY rate is higher that the uplink (Rx) for 78.5% of the stations in our dataset. Overall, the average transmission PHY rate calculated over all the stations and gateways is from 13.4% to 27.7% higher than the average reception PHY rate (Table IV). This asymmetry is observed for 802.11n 1x1 and 802.11g stations, in Figures 3(a)-3(b). These differences can manifest due to various factors – antenna diversity, transmit power, interference and receiver noise floor variations [15]. As we consistently observe such differences, we speculate this is due to transmit power asymmetry. The gateways in our study are configured at 19 dBm Tx power, larger than what we expect on small mobile devices. We confirmed this through controlled experiments on a few smartphones (Samsung Galaxy S3 and S4) and Atheros 802.11n NICs, which showed average Tx powers of 14 dBm and 16 dBm, respectively. Yet another reason for link asymmetry, for 802.11n 2x2 stations, is asymmetric Rx/Tx capabilities. These stations may have 2 receive RF chains (antennas), but transmit using only one transmit RF chain. This applies to 33.8% of the 802.11n 2x2 stations in our dataset. As a consequence, the downlink PHY rates (gateway to station) will be higher than the uplink, for these stations. The 7.2% stations in our dataset for which uplink traffic dominates (cf. Section IV-B) may be more sensitive to these PHY rate asymmetries depending on their applications' traffic demands.

B. Variation Across Homes

We now investigate wireless link performance variations among the homes of our deployment. Our results show that the average PHY rate across homes can vary up to an order of magnitude. To understand the impact of technology on performance, in Figure 4, we plot the PHY rate performance (Tx and Rx) across different technology clusters. Each point in the scatter plot represents a single home, and the y-axis represents the average home PHY rate. Given the fact that 802.11n technology offers higher performance than legacy technology, we would expect that the points corresponding to these would be higher than for the two legacy technologies. Surprisingly, this is not the case. Specifically, the highest performance 802.11b/g home achieves 3.3 and 5.6 times higher transmit and receive PHY rate performance than the lowest performance 802.11n home. Furthermore, even for a single technology, we see a wide gamut of performance. The best performing 802.11n home has an average PHY rate that is 9 times the worst performing one (the latter perform worse than all homes with legacy technology). We attribute these differences among technologies to artifacts such as interference and coverage. We elaborate on these causes in Section VI.

C. Achieved Throughput

Having so far looked at the wireless link and effective throughput T_E, we now analyze reports of the actual throughput T_A. This metric represents the *actual* number of bits that were sent/received on the link (in a short period of time). By relating this to T_E (which is a capacity measure), we can comment as to whether the wireless link is a bottleneck.

Figures 5(a) (and 5(b)) plot the distribution of downlink (and uplink) T_A, over all recorded samples, by particular 802.11 technologies. From the figures, we find that $T_A = 0$ in a very large number of samples. Recall that the gateway wireless driver only reports integer values; thus, $T_A = 0$ corresponds to something in the range $[0, 1)\text{Kbps}$. Over all the various technologies, from Figures 5(a) and 5(b), roughly

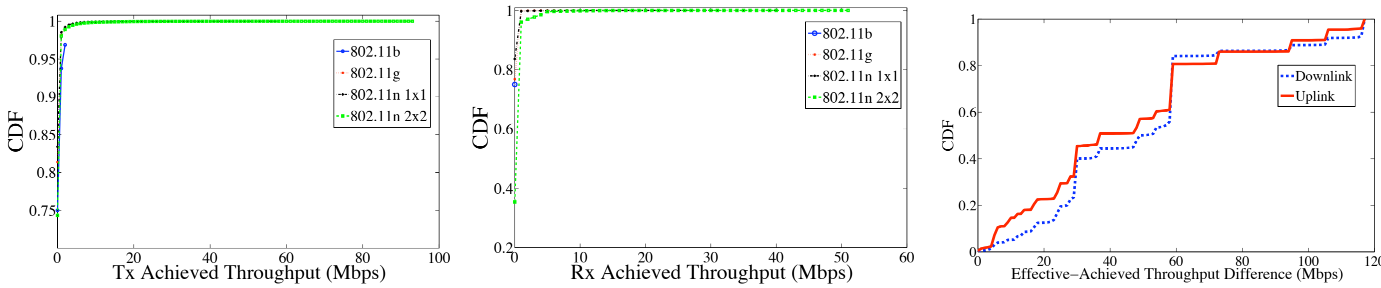

(a) Downlink T_A. (b) Uplink T_A. (c) $T_E - T_A$ difference.

Fig. 5. Achieved throughput distribution and comparison with the effective throughput.

81.2% and 73.3% of the samples, for downlink and uplink, record an actual throughput $< 1\text{Kbps}$. This implies that the wireless stations (when connected) have low traffic activity, and their WiFi interfaces are typically highly under-utilized. However, we do see sporadic reports of high actual throughput: the peak value of T_A is 93.4 Mbps and 51.9 Mbps (downlink and uplink) – recorded for 802.11n 2x2 stations. This indicates that there are instances where high throughput WiFi is required to accommodate the traffic rate. To consolidate the above observations, we compute for each episode, the *aggregate* throughput of the gateway, i.e., the sum of downlink and uplink T_A across all the associated stations. We also observe for the vast majority of the episodes the aggregate gateway T_A to be zero. Interestingly, the average gateway aggregate T_A over all episodes is at most 0.5 Mbps, for the majority of the gateways ($\sim$90%) of our deployment.

Effective vs. achieved throughput: We point out that from these results, we cannot conclude if the wireless link (or the application) is the bottleneck. Conceivably, T_A is low because channel conditions are bad and cannot support high throughput. To understand if this is the case, we analyze the difference ($T_E - T_A$). This represents the fraction of throughput that is not being used. Note that T_E can be estimated *only when there is traffic*, so the samples we analyze are restricted to those with non-zero T_A values. Figure 5(c) plots the distribution of this throughput difference for both uplink and downlink directions. From the figure we see a difference of at least 20 Mbps in 87.5% (downlink) and 77.4% (uplink) of the samples, which is significantly larger than the traffic needs of applications in common use today. It should be pointed out that even here, we cannot conclusively state if, the application or the wireless channel is the limiting factor. It is conceivable that T_E is high even when there are a large number of stations contending for the channel, and this may reduce the share to each station (reducing T_A). However, considering the very large gap between effective and the actual throughput, and the overall low values for T_A (not changing over time) we consider it unlikely that the wireless link is the bottleneck. Finally, as seen in Figure 5(c), we do find isolated instances where the throughput gap is small: at most 2 Mbps in 0.7% of downlink samples (1.7% in uplink). In these cases, the wireless link can a be a performance bottleneck. As Figure 5(c) indicates, the $T_E - T_A$ difference is overall smaller in the uplink case, which makes performance bottlenecks more likely to be in the uplink. This can be explained by our previous discussion – PHY rates are often smaller in the uplink direction.

Takeaways: Our results show that the wireless link performance is overall good with the effective wireless throughput to be significantly higher than the actual generated throughput. However, there are still instances of poor performance. In the next section, we elaborate on these instances.

VI. Understanding Poor Performance

In this section, we identify poor performance episodes and seek their root cause(s). We use a subjective rating to classify poor performance; we consider transmissions at 6.5 Mbps PHY rates or lower as indicators of poor performance. This is motivated by the fact that, it is the lowest 802.11n PHY rate and that our controlled experiments (under various WiFi conditions) have shown that these PHY rates are likely to yield performance below popular applications' bandwidth requirements (e.g., Youtube). Returning to Figures 3(a) and 3(b), we see that 6.3%/14.2%, and 2.2%/ 5.2% of the frames are transmitted/received at PHY rates less or equal to 6.5 Mbps, for 802.11n 1x1 and 802.11g stations, respectively[5]. Similarly, we identify poor performance instances for 802.11n 2x2 and 802.11b stations. Overall, 7.6% of the total instances in our dataset are poor performance episodes.

The fraction of poor performance episodes per home over the total performance instances varies across homes, as shown in Figure 6(a). For the majority of the homes the performance is good most of the time. Specifically, for 60.1% of the homes, the fraction of poor performance episodes is at most 6%. We observe 9% of the homes to have no poor performance episodes. However, the fraction of poor performance episodes can be up to 45% and 66%, for two homes of our deployment (Figure 6(a)). We next study their root cause.

Coverage: We gauge coverage problems by calculating the expected PHY rate R_E from RSSI for the uplink transmissions (station to gateway) as discussed in Section III-B. High R_E implies that the channel capacity can accommodate high transmission rates. Figure 3(c) shows that 78% and 88% of the total receptions are performed at the highest R_E option, for 802.11n 1x1 and 802.11g stations, respectively. On the other hand, only 0.03% of the R_E samples are less or equal to 6.5 Mbps. Table IV further shows that the average R_E is only up to 11.2% lower than the the maximum supported PHY rate. We further illustrate coverage performance in Figure 6(b), by presenting the RSSI distribution of all the stations of our deployment, including 802.11b and 802.11n 2x2 technologies. We observe that approximately 84% of the RSSI samples are

[5]802.11n stations can still use legacy, smaller than 6.5 Mbps PHY rates (e.g., for low RSSIs).

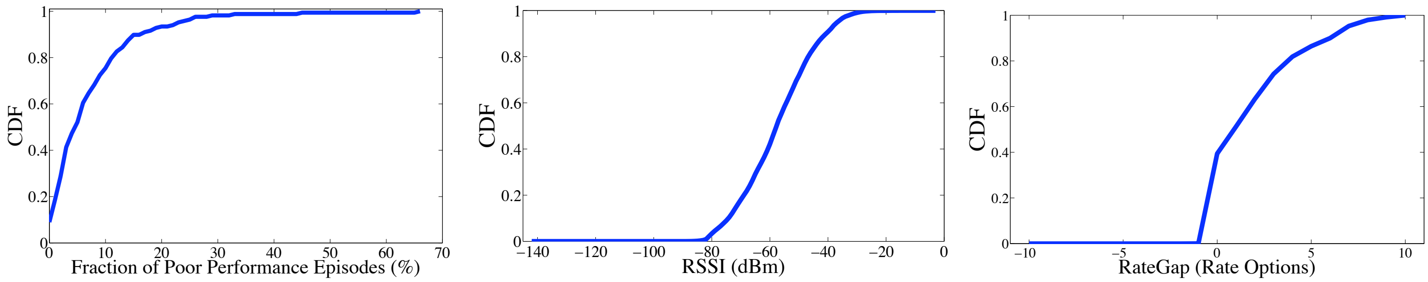

(a) Poor performance across homes. (b) RSSI distributions. (c) *RateGap* in rate options.

Fig. 6. Poor performance across homes and root causes.

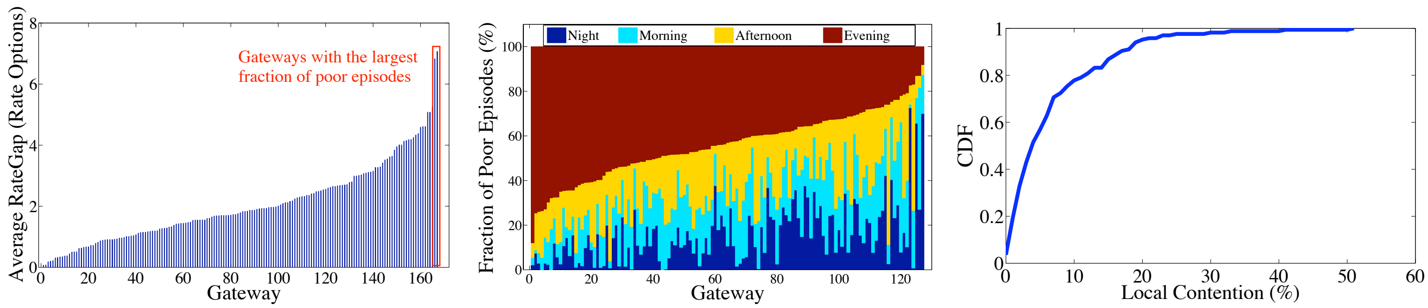

(a) Average *RateGap* across homes. (b) Poor performance in time. (c) Local contention (%).

Fig. 7. Poor performance in space and time and local contention.

at least -70 dBm, which corresponds to the highest R_E option, for all the station technologies. We see only 0.02% of the RSSI values to be less than -87 dBm, which corresponds to rates R_E less or equal to 6.5 Mbps. We conclude that *poor coverage is not the main cause of poor wireless performance.*

Other causes: We next seek to identify and quantify performance degradation which is not triggered by low RSSI, by leveraging the $RateGap$ metric from equation 1. In Figure 6(c), we present the $RateGap$ for all the records in our dataset. Recall that we compute $RateGap$ only for the uplink (cf. Section III-B). We can observe instances where $RateGap$ is high; 18% of the frame transmissions have a positive $RateGap$ greater than 4 PHY rate options, while the maximum $RateGap$ is 10 (observed for 802.11n 2x2 stations). This $RateGap$ can explain the poor performance variation among homes (cf. Figure 6(a)). In Figure 7(a), we show the average $RateGap$ performance for each gateway of our deployment, which includes all the instances from all gateway's stations. We observe the average $RateGap$ to vary from almost zero to 7 rate options. Interestingly, the two homes with the highest average $RateGap$ (as indicated in the plot), are the homes with the largest fraction of poor performance episodes.

A positive $RateGap$ can be attributed to interference, which can be high in the gateways of our deployment, since they operate at the congested 2.4GHz band. We expect interference to be high mainly during evening times, where people are at home, using their devices (cf. Section IV-B). We verify our intuition by plotting the distribution of poor performance episodes during different times of day for each gateway in Figure 7(b). We consider only the gateways which have relatively high number of poor episodes (at least 200). We observe that for 77.2% of the gateways, the majority of poor performance instances are during evening times.

Interference has been tied to the density of the WiFi environment [6], since neighboring 802.11 devices can interfere with home devices. We therefore investigate this, by correlating the $RateGap$ (which can be caused by interference) with the WiFi environment density. Specifically, we correlate $RateGap$ with home's number of neighboring wireless networks N (SSIDs), a metric which is used by our gateway's channel selection algorithm. Apart from the number of neighboring SSIDs N, we devise another WiFi environment density metric (named RSSID - RSSI Density), which considers both the number of SSIDs and their RSSIs to the home gateway. It is $RSSID = \sum_{i=1}^{N} \frac{1}{|RSSI_i|}$ and it increases with the number of SSIDs and their signal strength to the gateway.

Interestingly, we find that *there is no strong correlation between WiFi performance and the density of the neighboring WiFi environment.* Figures 8(a) and 8(b) correlate our gateways' average $RateGap$ with the number of SSIDs and the RSSID metric, respectively. In the scatterplot, each point corresponds to a home. We observe that, higher WiFi density (higher number of SSIDs, or RSSID) does not necessarily result in higher $RateGap$. To better understand our finding, we compute Kendall, which is a rank correlation that does not make any assumption about the underlying distributions, noise, or the nature of the relationships. Kendall correlations of $RateGap$ with the number of SSIDs and the RSSID metric are very low (0.09 and 0.08 respectively). We make similar observations with Spearman and Pearson correlations. Likewise, we did not observe any correlations for other metrics such as the max/min RSSIs from the neighboring SSIDs, with WiFi performance. Note that we need to interpret our findings with cautiousness. First, our metrics do not consider the channel used by the neighboring SSIDs. This is because our gateways scan for SSIDs very infrequently (at most once or twice per day), and consequently, we do not have fine-grained feedback

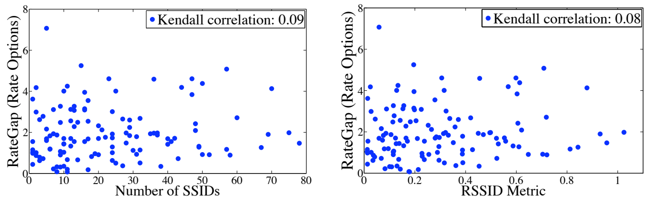

(a) RateGap vs. SSIDs. (b) RateGap vs. RSSID metric.

Fig. 8. WiFi RateGap and environment correlation.

of neighboring SSIDs' channels. Second, our gateways do not report the traffic from neighbouring networks, which would allow us to gain deeper insights for the correlation between interference and neighboring APs.

Local contention (i.e., stations associated to the home gateway contending for the wireless medium) is also one possible source of interference, as it reduces the ability of the stations to send traffic and can result in hidden terminals. To estimate local contention, we first split the time in 30-second bins. We then define as local contention, the fraction of the time bins where we observe traffic from at least two stations. These stations are likely to contend for the wireless medium. Figure 7(c) shows the local contention aggregated over all the gateways of our dataset. We observe that for 77.8% of the gateways, the local contention is less than 10% and the median is 6.6%. We further observe that, for the majority (78%) of the poor performance episodes, there is no local contention. Therefore, local contention is unlikely to cause interference in our gateways. We also exclude interference from overlapping 40MHz channels [19] as a cause, since almost all the frame transmissions in our dataset use 20MHz bandwidth PHY rates.

Finally, our results show a tiny fraction (0.5%) of the samples to have negative $RateGap$ (Figure 6(c)). This is because, RA algorithms may sample high rates to estimate channel quality (cf. Table II), which can cause frame losses.

Takeaways: Overall, we eliminate poor wireless coverage as the root cause of poor performance. We speculate that there can be interference from external WiFi and non-WiFi devices operating on the same frequencies as our gateways.

VII. Conclusion

This paper presents a study of 802.11 wireless home network performance, using data from a 4 month measurement campaign of 167 homes, subscribing to a large ISP. Our results show that, wireless home networks have adopted the high-throughput 802.11n technology, and they are overall high performing in terms of their PHY rates. This result corroborates previous studies that have been carried out at a smaller scale [3]. We still identify instances of poor performance, where we eliminate poor coverage to be their root cause. We believe the overall good coverage can be explained by the fact that the subscribers in our dataset are mostly urban residents in the three largest cities of a European country where apartment dwelling is the norm (rather than houses). In the case of these poor performance episodes, we point to interference as their likely cause. Since we use basic metrics available in commodity gateways that can be collected at a large scale, we cannot accurately quantify interference effects. On the other hand, previous efforts have used metrics that can uncover interference, but sacrifice wide deployability [2], [3]. Ideally, a combination of the two approaches is likely to be fruitful in understanding wireless performance.

While we identify rare instances of "poor wireless", anecdotally these are a common cause for helpdesk calls for ISPs. From discussions with the ISP helpdesk operators and from analyzing a small sample of helpdesk logs from our ISP, we see that WiFi related calls are mostly unrelated to link performance and instead relate to issues such as gateway misconfigurations, authentication problems, etc. We plan to study these problems in future work. Finally, our analysis of home wireless traffic activity shows very low resource utilization across gateways and WiFi devices. Overall, we consider our efforts an important step to better understand wireless home networks.

Acknowledgment

This work was partially funded by the European ICT FP7 User Centric Networking project (grant no. 611001).

References

[1] S. Sundaresan, W. Donato, N. Feamster, R. Teixeira, S. Crawford, A. Pescape. Broadband Internet Performance: A View From the Gateway. *ACM SIGCOMM'11.*

[2] K. Papagiannaki, M. Yarvis, W. S. Conner. Experimental Characterization of Home Wireless Networks and Design Implications. *IEEE INFOCOM'06.*

[3] A. Patro, S. Govindan, S. Banerjee. Observing Home Wireless Experience through WiFi APs. *ACM MOBICOM'13*

[4] S. Grover, M. S. Park, S. Sundaresan, S. Burnett, H. Kim, N. Feamster. Peeking Behind the NAT: An Empirical Study of Home Networks. *ACM IMC'13.*

[5] S. Sundaresan, N. Feamster, R. Teixeira, Y. Grunenberger, D. Papagiannaki, D. Levin. Characterizing Home Network Performance Problems. *UPMC Technical Report, 2013.*

[6] L. Di Cioccio, R. Teixeira, C. Rosenberg. Measuring Home Networks with HomeNet Profiler. *PAM'13.*

[7] C. Kreibich, N. Weaver, B. Nechaev, V. Paxson. Netalyzr: Illuminating the Edge Network. *ACM IMC'10.*

[8] P. Kanuparthy, C. Dovrolis, K. Papagiannaki, S. Seshan, P. Steenkiste. Can User-Level Probing Detect and Diagnose Common Home-WLAN Pathologies. *ACM CCR'12.*

[9] S. Rayanchu, A. Patro, S. Banerjee. Airshark: Detecting Non-Wifi RF Devices Using Commodity Wifi Hardware. *ACM IMC'11.*

[10] J. Manweiler, P. Franklin, and R. Choudhury. RxIP: Monitoring the health of home wireless networks. *IEEE INFOCOM'12.*

[11] A. Gupta, J. Min, I. Rhee. WiFox: Scaling WiFi Performance for Large Audience Environments. *ACM CoNEXT'12.*

[12] D. Aguayo, J. Bicket, S. Biswas, G. Judd, R. Morris. Link-level Measurements from an 802.11b Mesh Network. *ACM SIGCOMM'04.*

[13] B. Raman, K. Chebrolu, D. Gokhale, S. Sen. On the Feasibility of the Link Abstraction in Wireless Mesh Networks. *IEEE/ACM Trans. Netw. 17(2): 528-541, 2009.*

[14] D. Kotz, K. Essien. Analysis of a Campus-wide Wireless Network. *ACM MOBICOM'02.*

[15] G. Judd, P. Steenkiste. Understanding Link-level 802.11 Behavior: Replacing Convention With Measurement. *ACM WICON'07.*

[16] I. Pefkianakis, Y. Hu, S. H.Y. Wong, H. Yang, S. Lu. MIMO Rate Adaptation in 802.11n Wireless Networks. *ACM MOBICOM'10.*

[17] S. H. Wong, H. Yang, S. Lu, V. Bharghavan. Robust Rate Adaptation for 802.11 Wireless Networks. *ACM MOBICOM'06.*

[18] J. Bicket. Bit-rate Selection in Wireless Networks. *MIT MSc Thesis'05.*

[19] V. Shrivastava, S. Rayanchu, J. Yoon, S. Banerjee. 802.11n under the Microscope. *ACM IMC'08.*