

Observing Home Wireless Experience through WiFi APs

Ashish Patro, Srinivas Govindan, Suman Banerjee
Department of Computer Sciences, University of Wisconsin Madison
{patro, govindan, suman}@cs.wisc.edu

ABSTRACT

We present a measurement study of wireless experience in a diverse set of home environments by deploying an infrastructure, we call WiSe. Our infrastructure consists of OpenWrt-based Access Points (APs) that have been given away to residents for free to be installed as their primary wireless access mechanism. These APs are configured with our specialized measurement and monitoring software that communicates with our measurement controller through an open API. We have collected wireless performance traces from 30 homes for a period in excess of 6 months. To analyze the characteristics of these home wireless environments, we have also developed a simple metric that estimates the likely TCP throughput different clients can expect based on current channel and environmental conditions. With this infrastructure, we provide multiple quantitative observations, some of which are anecdotally understood in our community. For example, while a majority of links performed well most of the time, we observed cases of poor client experience about 2.1% of the total time.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network Monitoring; C.2.1 [Network Architecture and Design]: Wireless communication

Keywords

Home WiFi networks; Deployment; Measurement; Interference; Characterization; WiSe

1. INTRODUCTION

Residential wireless networks have continued to become more complex and diverse over time. A typical home network today has different kinds of WiFi enabled devices such as laptops, handhelds (smartphones and tablets), printers, entertainment devices including game controllers (XBox, Wii), interactive television (e.g., Apple TV, Google TV), wireless HDTV transmitters [8], and wireless-based security cameras and other security systems. In addition, there is a plethora of other diverse wireless appliances in our homes — various cordless handsets, Bluetooth headsets, various types of sensors and actuators, and even microwave ovens that operate in the same unlicensed bands. Over the last decade or so, there has been a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom'13, September 30-October 4, Miami, FL, USA.
Copyright 2013 ACM 978-1-4503-1999-7/13/09 ...\$15.00.

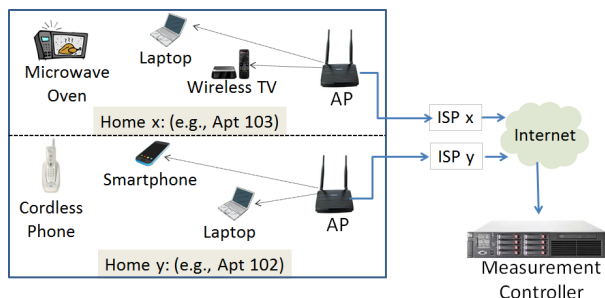


Figure 1: WiSe measurement framework.

few individual studies that have tried to quantify the experience and performance of wireless networks. Most famous among them include Jigsaw for WiFi activities in a university campus [4], and a few in home environments [14, 5]. In such prior work, researchers have attempted to understand and evaluate specific wireless characteristics through passive observations by carefully deployed sniffers.

Due to the above body of work and through various anecdotal experiences, our community has a broad understanding of WiFi performance in different settings. In general, we know that users are able to connect and get reasonably “good” performance most of the time, although there are instances where the networks appear to be frustratingly slow or unavailable. In this paper, we have attempted to perform a more systematic study of WiFi experiences in home environments and present a detailed characterization.

A view through home WiFi APs: Unlike prior wireless measurement studies that have deployed passive sniffers, in our work we have taken a significantly different approach — we have given away 30 of our own OpenWrt-based WiFi APs, enhanced with our network measurement and analysis software, to various home residents for free. We collect wireless-specific measurements *inline* using these WiFi APs as vantage points. Together, they form a unique type of wireless measurement infrastructure that we call WiSe (**W**ireless **i**nfrastructure for **i**nline **S**ensing). Unlike out-of-band wireless sniffers that observe a subset of the traffic, *our WiFi APs can exactly observe all traffic to and from its clients*, correlate them with wireless performance, and can hence conduct types of measurements not possible otherwise (Figure 1).

The WiSe WiFi APs are equipped with dual WiFi NICs — one provides the AP functionality to clients, and the other serves some unique and additional measurement functions when necessary. Deployed over the last 6 months, the average WiFi AP observes more than 2 GB of data traffic per day.

In choosing residents to give these WiFi APs to, we attempted to maximize the mutual proximity of these APs in the bulk of our deployment. In fact, two-thirds of our APs are concentrated in two dense apartment buildings in the downtown area of an urban city.



Figure 2: WiSe deployment in the downtown area and some suburban locations in Madison, Wisconsin. The stars indicate the two apartment complexes and the other points indicate locations with deployment of single APs.

Such a deployment allows multiple of our WiSe APs to often observe the same traffic from different vantage points enabling some specific types of wireless measurements. Some APs have been distributed to residents in homes further away from the city downtown that are much more suburban in nature. This allows for some comparisons between different types of wireless environments (Figure 2).

In a way, this measurement effort is related to the recent BISmark project [20] and the SamKnows effort [2] which focused on measurements of broadband ISP behavior across different communities. The key different between BISmark, SamKnows, and WiSe is that the first two projects primarily focus on characterization of the wired Internet path from the ISP’s network into the home and not on the wireless network properties, while WiSe focuses solely on the wireless network properties and not on the rest of the ISP path. Our measurement infrastructure is quite stylized both in software and hardware to our specific wireless measurement goals.

Objectives of measurement study and a wireless performance metric: In this paper, our goal has been to evaluate our community’s collective intuition of WiFi networks performance. In particular, we wanted to answer a number of questions such as: (i) how often do home WiFi networks provide us with good, mediocre, or bad performances? (ii) when the wireless performance is bad, what are likely causes of such experience and how persistent are these experiences? (iii) how much interference do we see in these environments from a) other WiFi sources and b) various non-WiFi sources?, and (iv) how do users configure their WiFi networks?

To evaluate and answer all of these questions, one of the first requirements is to define a wireless performance metric that can capture the overall “goodness” of the network. While a performance metric is necessary for our measurement study purposes, we believe that a good metric may have broader applicability.

In general, many different performance metrics are possible. In defining a metric, we posited four desirable properties for our purposes: (i) it should be easy to observe this metric from any passive vantage point and especially so for a WiFi AP; (ii) the computation of the metric should not require us to impose additional load on the wireless network; (iii) the metric should be relevant to only the wireless part of any user’s end-to-end path and ignores potential latencies and bottlenecks that exist uplink from a WiFi AP at home; and (iv) it is application-agnostic.

As a result of this, we cannot use certain simplistic measures as our metric, e.g., end-to-end throughputs of all flows of different users — which is a function of upstream wired bottlenecks and also the individual applications that generate these flows. Instead, we propose a specific metric that can provide the following estimate — given the current wireless environmental conditions, what is the likely TCP throughput of a new flow between the AP and the different WiFi clients connected to the AP? Given that TCP elastically adjusts its throughputs to current conditions and bottlenecks, this metric would require to consider the busy time of the channel, signal quality to and from the client, consequent loss properties, and re-transmissions,

and also overall airtime expended by the AP to and from its different clients. We call this metric, *WiFi-based TCP throughput* or *Witt*. Note that in computing our metric we do not conduct active TCP measurements, but instead estimate it *through a model based on various passive observations*. While it is possible to define many other metrics that capture the wireless experience, given the constraints of our measurement infrastructure and the goals of our study, we find this metric simple to compute, useful in understanding the wireless performance of WiFi environments, and fairly easy to explain to lay users of wireless networks. We provide a more precise definition of Witt in §3.

Key contributions: In this paper, we make the following main contributions.

- Present a unique measurement study of WiFi experience in home environments through the lens of WiSe APs across 30 homes (§2.3). Through our measurements we observed that while most of the WiFi clients experience moderate to good performance, poor performance plagues these environments about 2.1% of the time. The major cause of poor network performance (airtime, signal strengths) was dependent on the environment. Some APs experienced short periods of high impact interference (81% degradation) from external sources (e.g., microwaves). Also, majority of APs at homes tend to have static channel configurations over time, indicating that these APs do not adapt to interference or contention experienced by APs due to external sources.
- Describe the unique measurement and monitoring infrastructure deployed for this purpose that is being continually enhanced in scope and volume. The WiFi APs as part of this infrastructure support an open API through which wireless performance properties (WiFi and non-WiFi) are gathered and analyzed.
- Define a simple wireless performance metric, called Witt, that can be used to get a quick and easy measure of the wireless experience of a WiFi AP and its clients (§3).

2. WiSe INFRASTRUCTURE AND MEASUREMENT FRAMEWORK

Passive measurements of wireless systems, usually through carefully placed sniffers [4, 12] have been popular mechanisms to understand various properties of these systems. Our proposed approach in collecting measurements and understanding WiFi performance is passive as well. Unlike such prior wireless measurement efforts, our work (inspired in part by the BISMark project) places specialized measurement capabilities embedded into regular WiFi APs, and observes various wireless properties inline. This approach has various unique advantages as described next

Inlined measurements from home WiFi APs: Measurement capability installed in WiFi APs provide a unique vantage point. This is a useful approach for collecting wireless measurements and understanding wireless properties for the following reasons. First, our infrastructure is able to observe all traffic to and from its clients without missing any data, unlike external passive sniffers. Second, APs can sometimes be re-configured by users to operate on different channels or with different transmit power levels. Since our measurement infrastructure sits inside these APs, we are both aware of these changes and can ensure that such changes do not affect our measurement process. Third, we can use different built-in functionalities in the APs to capture various types of information, e.g., CRC errors, CCA threshold settings, busy time of channel as reported by the AP’s WiFi NICs, etc. The availability of this information directly from the AP itself proves to be particularly useful.

Perhaps, the most important benefit on inlined measurements is that we can provide better incentives to residents to deploy the

Type	Description
AP statistics	Record aggregate statistics local to the AP such as airtime utilization, overheard beacons from external APs, total received packets, packet counts with CRC errors.
Client statistics	Record aggregate downlink statistics per associated client (e.g., Total packets sent, received, retried, client's signal strength at AP).
Non-WiFi devices	Report non-WiFi devices detected by the AP (type, start time, duration, RSSI), e.g., using Airshark [16].
Per-packet summaries	Record packet summaries for all links overheard by the AP. Each packet summary contains: received timestamp, packet length, PHY rates, retry bit and RSSI (average overhead < 1%).
Flow statistics	Report aggregate flow level statistics (e.g., sent, rcv, packet retries per domain).

Table 1: The data gathered by the WiSe framework through an open API about the wireless network by using each connected WiSe AP as a vantage point.

measurement infrastructure. In our deployment so far, we have given away the APs for free to residents. The requirement to receive our AP is for the residents to deploy them *as their primary WiFi AP in their apartments and homes*, and a willingness to participate in our measurement framework. If for some reason, this AP has any problems, the users are incentivized to fix them (and inform us if necessary) almost immediately. In our prior experience, users are not always as meticulous in ensuring that passive sniffers stay online and functional for long running measurements.

User privacy in our measurement study: An infrastructure can potentially be highly intrusive to users. However, in this work, we have only needed to observe IEEE 802.11 frame header information for the most part, that do not carry any private information of users. As described to our Institutional Review Board, we have informed our participating users (and in some cases the relevant ISPs) that we do not capture user identifiable information in the study. In a few aspects when information such as flow types are analyzed (e.g., Netflix video), we have taken care to anonymize user and home identities (using an identity decoupling technique) in the measurements that this information cannot be mapped back to any individual user or home.

2.1 Infrastructure description

WiSe Access Points: For our study, we deployed OpenWrt [1] based APs using the ALIX 2d2 platform [6] (having a 500 MHz AMD Geocode CPU, 256 DDR RAM and slots for flash disk, Mini PCI and USB accessories). Each AP is equipped with two Atheros 9220 Mini-PCI WiFi NICs. The "primary" WiFi NIC is setup in the Access Point mode to allow users to connect their WiFi based devices to the AP. The "secondary" wireless NIC is used as back-up wireless NIC to perform various wireless measurements for the purpose of this study. The nodes were set to use a default transmit power of 17 dBm (50 mW). We also benchmarked our routers against commercial routers with similar configurations and achieved comparable performance under different workloads.

Measurement controller: A measurement controlling server collects periodic wireless measurements from the APs as well remotely configures the WiSe APs. Currently, our controller runs on a standard Linux server (3.00 GHz dual core CPU, 4 GB DRAM) with a public hostname. In our current setup, the controller is deployed in our lab's server cluster with public hostname/IP.

2.2 Measurements and Management API

In order for us to perform and collect measurements from WiSe APs, we have defined and implemented an open API that expose a set of simple capabilities. Our measurement controller uses this open API to specify all types of measurements necessary and also to remotely manage and configure the various APs. Table 1 provides a high level description of this open API. Over the last 6 months, we have incrementally added all of these capabilities into

the WiSe APs to collect various statistics. We now briefly describe these different measurements performed over the course of this study.

Basic Statistics. Each WiSe AP periodically collects aggregate statistics (in 10 second intervals) such as overheard WiFi packets and beacons from neighboring APs present in the vicinity. These aggregate statistics include the total number of CRC errors for received packets and the per link packet transmission summaries for the overheard links. Local packet statistics collected the WiSe APs include aggregate statistics per local WiFi link (using kernel statistics), the nature of wired traffic by finding the TLDs (Top Level Domains) corresponding to the flows.

Airtime utilization. To measure the airtime utilization, we use the aggregate "busy time" statistics maintained by the Atheros NICs that we used with our APs. Whenever the energy level is detected to be higher than the CCA (Clear Channel Assessment) threshold (due to WiFi or non-WiFi sources), the channel is marked as busy, otherwise its marked as free. The driver records cumulative statistics about the total channel busy period observed over a period of time. In addition, the driver also records airtime utilization statistics due to local packet transmissions as well as utilization due to overheard WiFi packets only.

Per-packet Summaries. To perform WiFi interference analysis (§5.2), each WiSe AP records a short 10 byte per-packet summary for the WiFi data packet headers of its own links as well as the overheard data packets from other WiFi links on the same channel. Each packet's summary (grouped by link) contains the MAC timestamp (from driver) for the packet's transmission start time (32-bit timestamp with microsecond granularity), packet length, PHY rate, retry bit and RSSI. These statistics are periodically reported back to the controller to provide fine grained information about the WiFi activity on the channel. Across our deployment, the average WiFi packet size was between 900 bytes to 1000 bytes. So, it translates to only 1% average overhead per data packet.

Non-WiFi device detection: In our prior work, we designed a tool called Airshark [16] which demonstrates the feasibility of Atheros 9280-based WiFi NICs to detect non-WiFi devices (e.g., microwave oven, cordless phones) by using commodity WiFi NICs available in the market (e.g., Atheros 9280 chipsets). Airshark uses subcarrier-level energy samples received from the WiFi NIC as input to algorithms to detect the presence of different types of non-WiFi devices. For our work, we ported Airshark to run on our WiSe Access Points to detect the presence of non-WiFi devices in residential deployments. Since Airshark requires the WiFi NIC to be placed in monitor mode, we implement this functionality in the secondary WiFi NIC of the AP and have it scan all WiFi channels periodically. We also use the information to study the impact of non-WiFi interference in our deployments.

Remote configuration of WiSe APs. Our measurement controller can also allow automated management and remote configuration

Location (AP IDs)	Count	Deployed Since
Bldg 1 (APs 1 - 14)	14	Sep 2012
Bldg 2 (APs 25 - 30)	6	Dec 2012
Others (APs 15 - 24)	10	Oct 2012

Table 2: Summary of the deployment of WiSe Access Points




Figure 3: Distribution of the daily data download over the WiFi network per WiSe AP (Sep 1, 2012 - Jan 31, 2013). 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure.

of the WiSe APs. The configuration capability includes picking channels of operation, choice of transmit power, and remote install of new software components if needed. Configurations are typically left to residents (using a local configuration utility), but this capability also allows us to remotely make any modifications when necessary (or requested by the residents). Configuration changes are installed only when the AP has no clients associated to it. This ensures that users are not affected by these choices.

2.3 Deployment

To understand the wireless characteristics in real home environments, we distributed our Access Points to volunteers who used them as the primary wireless APs inside their homes (Table 2) to connect all their WiFi devices. During this period, all APs were configured to use a randomly selected channel on different days to capture the impact of different channels on our APs. Channel changes were done sometime late at night when the APs had no associated clients. We now discuss the details of our deployment.

Multiple APs within two residential apartments. We collaborated with the property managers of two apartment buildings (Bldg 1 and Bldg 2) in the downtown area of our city to sign up volunteers, who wanted to be a part of our study in exchange for getting free wireless APs. Both these apartments were multi-storied with many 1-2 bedroom apartment units on each floor. In the case of Bldg 1, we focused our deployment on a specific portion of the building (5 floors) having 12 units per floor. We have currently given away 14 APs to Bldg 1 residents such that each of these floors have multiple APs (usually between 2 and 4). Bldg 2 is a 10-story building with dormitory style housing (1-2 bedrooms per apartment unit) in which the property manager provides its own wireless service to residents by deploying its own off-the-shelf APs. Hence, working with the property manager and the local ISP, in this building we deployed 6 APs on 2 consecutive floors (3 per floor) in the designated areas, that replaced their existing APs. We were told that these building APs were still the primary form of Internet access in this residential building. Thus, the two buildings had slightly different approaches to wireless AP deployment and provided us with some diversity in our measurements.

Across different types of user homes. We also distributed our APs to volunteers and colleagues staying at different low to high density residential units spread across different locations in our city (Others). This was done to capture the wireless network properties across different types of home environments.

Causes	Indicator
Non-WiFi devices, Rate Anomaly, Heavy external traffic	Airtime utilization
WiFi and non-WiFi Interference, AP and Client Transmit Power + Signal Environmental/Location effects	Losses + Data Rates
Local WiFi clients sharing an AP	Local contention

Table 3: The list of potential factors can cause the degradation of a WiFi link’s performance and potential indicators of these issues.

Volume of observed traffic. Over the course of our deployment, we observed wide-ranging daily WiFi usage characteristics across our users (Figure 3). The median WiFi usage across users varied between a low 30 MB per day to over 5.6 GB per day. The 90th percentile usage for some users was as high as 8-9 GB per day. The usage of these wireless networks can be driven by a number of factors such as the kind of wireless devices used, wired access link capacity, the nature of the traffic/service used and user behavior.

3. QUANTIFYING WIRELESS EXPERIENCE THROUGH A METRIC

In order to analyze and understand the large volume of wireless performance data that we have observed, we wanted to categorize them through the use of a metric. In defining this metric, we set out some modest goals as outlined in Section 1 with our intent being two-fold: (a) to quickly identify the good and bad performances of our network, and (b) to be able to explain to residents or ISP personnel, how often there are performance problems of some type.

As discussed, there can be many different ways to approximate the “goodness” of a wireless network. In this paper, we wanted to pick something that is passively observable in real-time, captures the wireless experience alone, and is application-agnostic. We call our defined metric, *Witt*, and explain its construction next.

3.1 WiFi-based TCP Throughput metric

The idea of our proposed metric is fairly simple — it measures the *likely* TCP throughput between a client device and its AP (or a server located on the same uncongested LAN as the AP), given the current wireless conditions that exist. The metric is clearly, a property of the client and AP combined and the protocol used (802.11g vs. 802.11n). In this paper, we also consider the average value of the metric for the entire AP as a single aggregate. In such cases, we take the average of *Witt* for all its *active* clients. An active client (for this paper) is defined as one which has sent at least a minimum level of traffic in recent time window. We use a threshold of at least 500 packets in the last 10 second window — we want to bias the metric towards client devices that are imposing a higher load at the current time, than the ones that are less active.

Given that TCP-based flows are a dominant fraction of Internet traffic, this metric likely captures most of the user’s experience when the wireless link is the bottleneck (short of analyzing performance on an application-by-application basis). We believe that such a metric would actually capture a lot of wireless properties on the network. Factors that imply poor wireless conditions, e.g., low signal strength, high degree of interference from various sources leading to losses, increased latency on the WiFi path due to reduced PHY rates or multiple re-transmission attempts, high airtime utilization leading to a reduced ability to send traffic, all will reduce the likely TCP throughput estimate, and vice versa. Thus, we focus on estimated TCP throughput as a direct measure of the link’s performance.

Feature	Correlation Coefficient
Airtime	0.321
CRC errors	0.345
Local contention	0.463
Signal strength	0.536
Effective rate	0.882
Effective rate + Airtime	0.915
Preferred "Link exp" model (Eqn. 2)	0.958

Table 4: Correlation of metrics with the observed TCP throughput (802.11g). The best individual and overall metrics are highlighted.

	Coefficients (β_0, β_1)	95% Confidence interval for β_1
802.11g	(0.167, 0.422)	(0.403, 0.441)
802.11n	(-0.493, 0.733)	(0.720, 0.746)

Table 5: Parameters of the linear model for predicting Witt.

Causes and indicators related to wireless performance. Table 3 presents a short summary of such causes for poor performance experience by wireless links. The table also shows that the impact of multiple such causes can be captured by using a set of key indicators observed locally by the APs.

For example, the impact of non-WiFi activity as well as presence of WiFi traffic using low PHY rates (rate anomaly) can be observed by APs through an increase in *airtime utilization*, i.e., the fraction of airtime occupied by only external WiFi and non-WiFi transmissions. Similarly, an increase in packet losses for a link indicates the presence of factors such as WiFi or non-WiFi interference at the receiver, poor signal quality at the client, location dependent performance problems etc. High local contention at an AP caused by the presence of multiple clients with high traffic demand can reduce the available throughput capacity per client. In the next section, we use these indicators to build the Witt metric.

3.2 How to measure Witt?

A key objective of our metric is that it should be easy to obtain through passive observations at the AP and should not impose additional traffic. Hence, we built a simple model of Witt based on likely factors that will impact the metric. To do this, we first collected targeted ground truth data, built and tested our model of separate parts of the data, and then used it for analyzing wireless network performance. We first describe our ground truth measurements and then our model, leading to the metric.

Ground truth measurements. To collect ground-truth measurements of WiFi-based TCP throughput under different conditions, we placed 4 of our own clients (laptops) at 8 different deployment locations within the apartment buildings (excluding lab experiments). These laptops co-existed with the actual users of the WiSe APs and performed TCP downlink throughput runs (using iperf between APs and clients) that lasted 20 seconds each. This setup allowed us to create a scenario with the WiFi link being the bottleneck.

In the case of the actual home deployments, the clients ran throughput measurements in intervals of 5 to 10 minutes over the course of a week through which we collected hundreds of these measurements. Further, these clients connected to different APs within the apartment buildings to emulate different types of link conditions. These experiments were automatically conducted at different times of the day and hence, covered a diverse set of link characteristics, channel and operating conditions. Table 6 shows the distribution of different parameter values from our ground truth measurements.

Creating the Witt metric. To understand the impact of different factors on the observed throughput, we correlated the TCP through-

Parameter	Airtime	CRC errors	MAC retries	Signal (at AP)
Minimum	17.6%	1.2%	6.2%	-71 dBm
10th percentile	20.5%	6.4%	13.3%	-70 dBm
50th percentile	38.3%	17.2%	24.9%	-61 dBm
90th percentile	47.5%	41.6%	56.7%	-52 dBm
Maximum	72.5%	98.8%	86.4%	-49 dBm

Table 6: Parameter value ranges in our ground truth measurements.

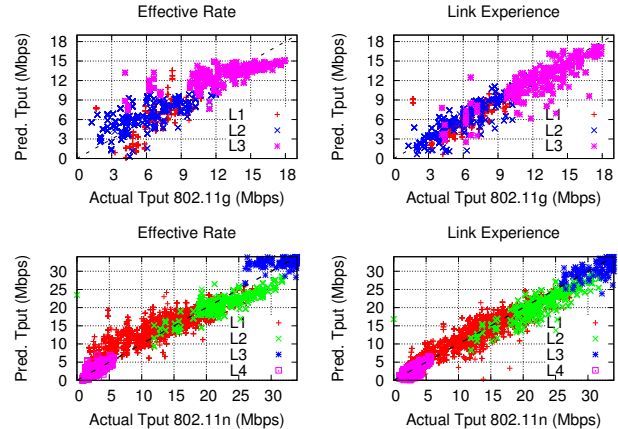


Figure 4: Scatter plot showing the actual vs. predicted TCP throughput values for the 802.11g (top) and 802.11n (bottom) ground truth experiments by using "Effective Rate" and "Link Experience". Points near the "x=y" line indicate instances with accurate prediction of TCP throughput. Each set of points corresponds to a different WiFi link.

puts from our clients with various wireless statistics recorded by the WiSe APs in 10 seconds intervals (Table 1). The following are some of the example statistical features that we tested to predict Witt: (i) airtime utilization, (ii) CRC error rate, (iii) client signal strength, (iv) local contention, and (v) effective rate — we define the last two precisely next.

- *Local contention (c)*: the relative amount of other client traffic transiting through this AP as a fraction of total traffic passing through this AP. It captures the fraction of time the AP spends on transmitting (and receiving) traffic of its other clients, and hence, reduces the ability of this client to receive (and send) traffic.

- *Effective rate (r)*: captures the net effect of packet losses and choice of PHY rate used on an AP-client link. It uses aggregate kernel statistics about the number of successful (s_i) and total packet (p_i) transmissions at each PHY rate (r_1, \dots, r_n) used by an AP-client pair. It is defined as:

$$r = \frac{1}{\sum_i p_i} \sum_i s_i r_i, \quad 1 \leq i \leq n \quad (1)$$

$$\text{link_exp} = (1 - a) * (1 - c) * r, \quad 0 \leq a \leq 1, \quad 0 \leq c \leq 1 \quad (2)$$

Table 4 shows the correlation between these features available at the AP (individually and also in various combinations) and the observed TCP throughput from our ground truth measurements. Among all possible combinations tested (not all are shown in Table 4), the best combination of features that has the highest correlation is given by Equation 2. We call it the "link experience" model — which depends on a combination of airtime utilization from external sources (a), local contention (c), and effective rate (r). Together this model intuitively captures various aspects that govern the experience likely to be expected for a wireless link. Note that all these features

Witt	≥ 16 Mbps	8 - 16 Mbps	4 - 8 Mbps	1 - 4 Mbps	< 1 Mbps
Rating	(V. Good)	(Good)	(Moderate)	(Poor)	(V. Poor)

Table 7: Using Witt to rate link quality.

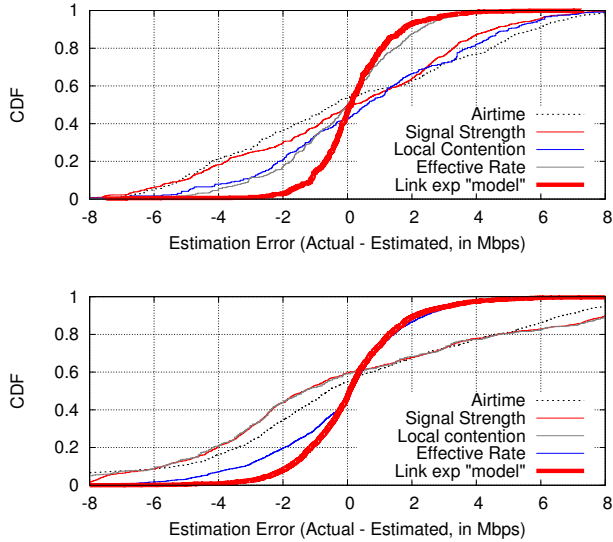


Figure 5: CDF of errors between the actual and predicted TCP throughput values obtained by using different metrics for 802.11g (top) and 802.11n (bottom).

are based on aggregate statistic per link (10 second intervals in our current system) and can be easily obtained using kernel statistics reported by wireless drivers.

Finally, to map “link experience” to Witt, we create a simple linear model between the two (Equation 3). We tested the fidelity of Witt by dividing our ground truth data into training and testing data sets, each corresponding to different days in our entire dataset. The parameters for our linear model as well as the error estimates for the slope, β_1 obtained using the training data, is indicated in Table 5. The small range of the the 95% confidence interval shows that the linear model is a reasonable estimate for predicting the throughput.

$$Witt = \beta_1 * link_exp + \beta_0 \quad (3)$$

A subjective rating. Based on the Witt value, we also use a subjective rating to classify links and APs into a range from *Very Good* to *Very Poor* based on the range of throughput observed (Table 7). Given that the highest Witt values estimated in our infrastructure approached 19 Mbps for 802.11g networks and 34 Mbps in 802.11n networks under best conditions, anything very low (1-4 Mbps) or lower was deemed poor or very poor. Further, our rating moved from across categories mostly with doubling of Witt value to indicate the ability of the APs to support higher TCP throughput.

3.3 Benchmarking the Witt metric

To evaluate the performance of the Witt metric, we present a couple of benchmarks in this section to compare the performance of the “link experience” model with other alternatives.

Effective Rate vs. Link Experience. As discussed earlier, effective rate exhibited the highest correlation amongst the individual features in estimating the TCP throughput. We used our ground truth TCP throughput measurements and compared them with the predicted TCP throughput values using the linear regression model (Equation 3) for effective rate and the link experience model. Figure 4

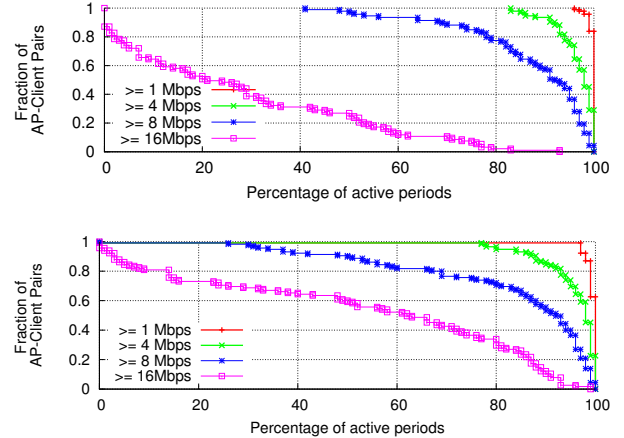


Figure 6: Distribution of Witt across AP-client pairs using 802.11g (top) and 802.11n (bottom) in our deployment that were active for at least 20 days.

compares the actual and predicted TCP throughputs for a few representative WiFi links from our ground truth experiments. It shows that the link experience model resulted in a better fit compared to effective rate. Only using effective rate as a feature can result in an overestimation of the link’s TCP throughput (e.g., in the presence of high channel utilization) which leads to higher prediction errors. The link experience model also accounts for the potential reduction in TCP throughput due to factors such as high airtime and/or local contention which results in a better prediction.

CDF of prediction errors using different features. Figure 5 shows the CDF of the errors of our proposed link experience model to create the Witt metric in estimating the WiFi TCP throughput and compares that to the other alternatives. The figure shows that our proposed formulation performs effectively and better than the baseline metrics with an estimation error under 1.5 Mbps and 2 Mbps for more than 80% of 802.11g and 802.11n instances respectively. In summary, Witt provides a quick estimate of likely TCP throughput that an average WiFi link can achieve based on observed traffic characteristics at the AP and without requiring any active measurements.

4. USING Witt TO CLASSIFY WIRELESS EXPERIENCE IN THE WILD

The Witt metric allows us to classify links into different categories based on their performance from the large volume of data collected through our measurements (over 100 GB). In this section, we analyze on the impact of various factors (local and external) that affected the performance of WiSe APs. To study periods with WiFi activity at the APs, we focus on the periods when a WiSe AP has at least one active client (§3.3).

How did the link performance vary across APs over time?

During the course of our study, a diverse set of clients associated with the WiSe APs (e.g., laptops, handhelds, entertainment devices etc.). We measured their Witt values (§3) during active periods grouped their results based on their Witt values. Figure 6 shows the distribution of the performance experienced by AP-client pairs who were active for at least 20 days in our deployment. While a majority of our deployment consisted of 802.11g based WiSe APs, others used 802.11n based APs. For the different Witt threshold values used to characterize link performance in Table 7. For example, a line corresponding to ≥ 16 Mbps shows the fraction of these links

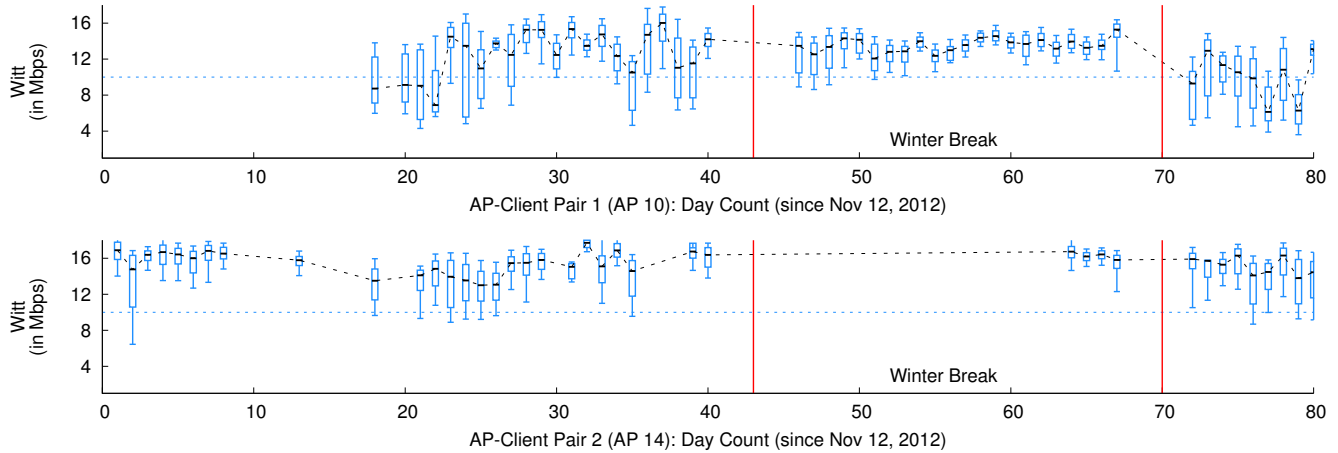


Figure 8: Distribution of Witt for 2 different AP-client pairs in our deployment from 12 Nov 2012 to Jan 31, 2013 and shows their variation in performance over the span of this period. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure.




Figure 7: Comparison of average Witt for AP - client pairs over 802.11g and 802.11n (3 days each).

(an AP-client pair) whose Witt value was above 16 Mbps for a given percentage of their active periods.

The graph shows that a majority of links performed well for most of the time. For example, around 80% AP-client pairs experienced Witt \geq 8 Mbps during more than 80% of their active periods. But, there were occasional periods for many AP-client pairs where the link performance was "Poor" ($<$ 4 Mbps). About 8 % of these pairs experienced Poor performance ($<$ 4 Mbps) during more than 10% of their active periods, indicating some periods of poor performance which may be observed by users. In the later sections, we discuss the causes of poor performance across our different deployments.

802.11g vs. 802.11n. To study the performance gains of using 802.11n vs. 802.11g for a given link, we configured the 802.11n based APs to use 802.11g for a few days. Figure 7 compares the average Witt values on 802.11g vs. 802.11n for clients that support both protocols. The Witt values for most AP-client pairs increased between 12% to 100% due to usage of higher PHY rates (e.g., 65 Mbps) and use of 802.11n specific MAC layer optimizations such as frame aggregation. Interestingly, some clients experienced similar performance on both protocols (e.g., AP-Client pairs 2 and 3), indicating that using 802.11n did not ensure higher throughput for all clients. This can happen due to packet losses based on the link quality or the prevalence of interference from co-existing 802.11b/g WiFi transmitters on the same channel [19].

Causes for "Poor" wireless experience. During the course of 80 days from November 12, 2012 to January 31, 2013, we detected a total of 186 and 2031 minutes of the "Very Poor" and "Poor" instances respectively, across all 30 WiSe APs (average of 2.3 and 25 minutes,

Indicators				Bldg 1		Bldg 2	
$A \uparrow$	$S \downarrow$	$L \uparrow$	$R \downarrow$	V. Poor	Poor	V. Poor	Poor
✓	×	×	×	0%	18.4%	0%	1%
×	×	✓	×	24.2%	49.5%	25.2%	78.1%
✓	×	✓	×	61.8%	26.7%	2.1%	1.4%
×	✓	✓	×	2.3%	1.1%	20%	15.8%
×	✓	✓	✓	9.4%	0%	51.6%	3.4%
Others				2.3%	4.3%	1.1%	1.3%

Table 8: Distribution of causes responsible for "Poor" and "Very Poor" periods in Bldg 1 and Bldg 2. Each "cause" is composed of 1 or more of following indicators and corresponding threshold values: High Airtime ($A \uparrow$, $>$ 60%), High MAC Loss Rates ($L \uparrow$, $>$ 50%), Low signal strengths ($S \downarrow$, $<$ $-70dBm$) and Low PHY rates ($R \downarrow$, \leq 12Mbps). Others correspond to remaining combination of factors such as high contention, signal strengths etc.

respectively per day across all APs). Thus, the "Very Poor" periods are rare but the "Poor" periods can occur intermittently depending on the link and the location. Overall, these cases accounted for 2.1% of the active periods during the 80 days.

In section 2.3, we had discussed that our deployments consisted of two apartments with multiple WiSe APs deployed within the same building. We aggregated the instances of poor wireless performance across WiSe APs in each of these apartments (Bldg 1 and Bldg 2). Table 8 breaks down the periods of poor wireless performance at these APs. Each row in Table 8 is a combination of indicators of poor performance: high airtime ($>$ 60%) [9], high wireless MAC layer packet losses ($>$ 50%) indicating the fraction of wireless transmissions that were retried, low signal strengths ($<$ -70dBm) and low PHY rates (\leq 12Mbps). We use the signal strengths from clients observed at the APs as estimate of link quality at clients since we do not have access to the clients.

The table shows that the presence of both high airtime and wireless losses ($A \uparrow + L \uparrow$) were the main cause (61.8%) of "Very Poor" performance in Bldg 1. This can happen due to the network congestion at the WiSe APs and clients leading to high packet losses. In Bldg 2, the major cause of "Very Poor" performance (51.6%) was poor signal strengths ($<$ -70dBm) which lead to high packet losses as well as usage of low PHY rates by the rate adaptation algorithm ($S \downarrow + L \uparrow + R \downarrow$). The impact of other factors ("Others"), such as

period of high local contention (>0.5) from other clients associated to an AP was quite low at both locations ($\leq 4.3\%$). The prevalence of low local contention at the wireless hop is due to the fact that it was not common for multiple clients associated to the same AP to generate high traffic demand during the same time interval (10 second intervals). In some cases where there were multiple *active clients* at the same AP, bottlenecks on the wired link or low traffic demand led to lower contention on wireless hop.

The most likely cause for the above observations about poor wireless performance is the *nature of the wireless deployments in the two buildings*. Bldg 1 has private APs per apartment unit resulting in a dense wireless deployment. Thus, some APs experience occasional high airtime utilization ($> 60\%$) due to neighboring sources of traffic. Bldg 2 provides centralized wireless service to its residents and thus, some users can occasionally experience poor signal quality based on the client device and location. The impact of losses due to low signal strengths was much lower in Bldg 1 ($<12\%$) due to good signal coverage within each apartment. But, performance issues can still arise due to the dense nature of these deployments. Finally, high WiFi loss rates ($L \uparrow$) was the major cause for "Poor" performance in both Bldg 1 and Bldg 2 (49.5% and 78.1% respectively). Some potential causes of these losses are packet reception issues at some clients or packet collisions at the receivers due to external WiFi/non-WiFi sources. Through our interference analysis from WiFi sources (§5.2) in Bldg 1, we were able to attribute at least 19% of these lossy cases during "Poor" periods to strong interference from external WiFi transmitters.

Variability in wireless experience over time. Figure 8 shows the Witt for 2 AP-client pairs (Bldg 1) across a period of 80 days since 12th November 2012. These were amongst the most actively used clients in our deployment and were chosen to show the diversity in link performance over time across different locations and clients. Amongst the 2 clients, the one at AP 14 experienced consistently good performance over this period due to low neighboring wireless activity (median Witt around 14 - 16 Mbps). On the other hand, the client at AP 10 experienced a higher fluctuation in performance as shown by the wide range of the observed median Witt across days. One of the reasons for this behavior was high airtime utilization caused by some neighboring WiFi transmitters at AP 10 (§5.1). The "Winter Break" corresponds to the period between Dec. 25, 2012 and Jan. 21, 2013 during which the wireless activity in Bldg 1 was lower compared to the other days¹. During this period, the median Witt value for AP 10 was consistently high and stable (around 14 Mbps) compared to other periods of time due to less utilized wireless channel conditions. This observation indicates the high impact of neighboring wireless traffic at this AP.

Summary. The majority of the WiFi clients ($\geq 80\%$) in our deployment of WiSe APs experienced "Good" wireless performance during most of their active periods while some clients (8%) experienced poor performance for greater than 10% of their active periods. Overall, the client's experience was poor during 2.1% of the total active periods across APs. The location and type of deployment (e.g., Bldg 1 vs. Bldg 2) influenced the causes of "Poor" instances. In a dense deployment of private APs (Bldg 1), high airtime utilization from neighboring APs was the major cause of performance degradation while low performance due to weak signal strengths were more prevalent in a centralized home deployment (Bldg 2). We also observed high variability in link performance over the day as well as

¹A large fraction of residents in the building are university students, and many of them were out of town during the break.




Figure 9: Time-series airtime utilization at AP 9 with and without the presence of traffic from an external AP that used low PHY rates. AP 9 was inactive during this period.

during across a period of time (e.g., high Witt values during periods of low wireless traffic in the building).

5. A MORE DETAILED VIEW

We now analyze the impact of external factors on the wireless clients in the wild: contention from transmitters using low PHY rates, hidden terminals and non-WiFi interference.

5.1 Contention from low data rate senders

The presence of transmitters using low PHY rates during the active periods of APs can cause their Witt to suffer. This is due to the rate-anomaly [21] problem caused by the loss in channel availability. Figure 9 shows an example of airtime utilization at WiSe AP 9 over a period of 2 hours in the presence and absence of traffic from an external AP (AP x). During AP x's activity, AP 9's airtime utilization (10 second average) increased from an average of 30% to around 65 - 70% due to the usage of low PHY rates by AP x. This was one of the major cause of "Poor" performance in Bldg 1 (§4).

Prevalence and impact of contention from low rate senders. We analyzed the presence of such transmitters and their impact on the WiSe APs from 12th Nov. to 21st Dec (40 days). Figure 10 (top and center) shows the duration of *active periods* per day (in minutes) and the number of days respectively, during the 40 day period, over which the WiSe APs experienced contention from external WiFi senders that transmitted at least 500 packets while using low average PHY rates (≤ 15 Mbps). Thus, it only shows the impact on WiSe APs while they were sending actual traffic to their clients. Figure 10 (top) shows that some APs (e.g. APs 1, 6, 10, 28) experienced contention from such senders lasting over multiple minutes (median between 20 - 48 minutes). Figure 10 (center) shows that 4 WiSe APs faced at least 5 minutes of contention from such low PHY rate traffic during 10 or more days while they were transmitting data to their clients. This happened due to the presence of nearby external APs at these locations (both Bldg 1 and Bldg 2) that consistently used low PHY rates for some clients across multiple days. Figure 10 (bottom) compares the distribution of Witt at the impacted APs during active periods with and without (5 minutes before and after) such contention. Some APs experienced consistent reduction in Witt during the periods. For example the 75th percentile value for AP 10 reduced from 12.5 Mbps to 6 Mbps during such periods. These external APs was the main cause for the low values of Witt at AP 10 as observed in §4.

In addition to activity from external low PHY rate transmitters, the period of channel contention experienced (Figure 10, top) is also partially dependent on the activity at the AP itself. For example, AP 6 (Figure 3) was the most active AP in our deployment which resulted in the higher periods of contention (median of 40 minutes)

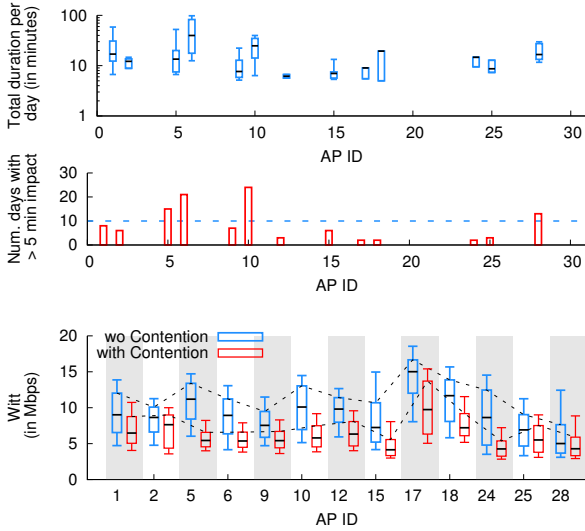


Figure 10: (top) Candlestick graph showing the distribution of active periods per day (Y-axis in logscale) during which the WiSe APs experienced contention from external transmitters (using average PHY rate ≤ 15 Mbps), (center) Bar graph showing the total number of days during which WiSe APs experienced this contention for a minimum duration of 5 minutes (from 12th Nov. to 21st Dec, 40 days), (bottom) Witt during active periods for clients with and without contention from such transmitters. 10th, 25th, 50th, 75th and 90th percentiles are shown here.

experienced by it. While, it is not possible for us to determine the signal strength properties of these external links that used low PHY rates, we analyzed our traces and found that many of these transmitters sent large data packets while using conservative rate adaptation. These links switched to low bitrates after a single failure resulting in high usage of low data rates and thus, airtime usage.

5.2 Packet losses due to WiFi sources

High packet losses were a major cause for the "Poor" cases in our deployment. Amongst external factors, hidden terminal (HT) style interference at a wireless receiver from nearby links can reduce a link's Witt by increasing packet losses. We leveraged our deployment of 14 WiSe APs in Bldg 1, to collect timestamped (microsecond level) WiFi packet summaries (§2.2) for all observed WiFi traffic from multiple vantage points. These packet summaries from neighboring APs are time-synchronized and merged at the controller using prior techniques from [4, 18] using common data packet summaries present in both traces.

We use these synchronized and merged packets summaries from multiple APs in Bldg 1 to compute "hidden terminal events" in "epochs" of 15 seconds. We mark an epoch as a hidden terminal event for a WiSe AP when the loss rates for one of its links is 40% higher for packets *overlapped in time* by the interferer compared to packets *not overlapped by any other transmitter* [18]. When an epoch is marked as a hidden terminal event, the main cause is packet losses at the receiver caused by the overlapping packet transmissions from the interferer. To minimize false positives, we used a constraint requiring a minimum of 1000 WiFi packets for a link and a minimum of 100 packet overlaps from a potential interferer per epoch to check the presence of a conflict. Thus, our results are a conservative estimate of the interference experienced by the APs in Bldg 1.

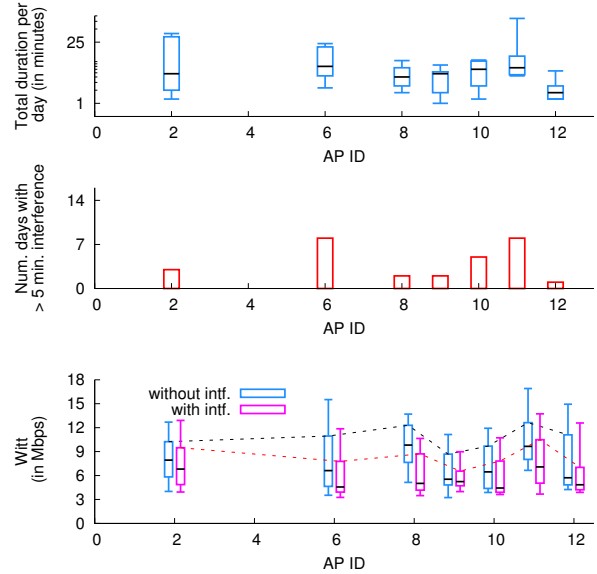


Figure 11: (top) Distribution of time per day during which the WiSe APs in Bldg 1 experienced hidden terminal interference from external links. (center) Bar graph showing the total number of days during which WiSe APs experienced hidden terminal interference for a minimum duration of 5 minutes per day over a period of 2 weeks. (bottom) Distribution of Witt with and without the presence of HT interference (within 5 minutes of the interference event). Min, max, 25th, 50th and 75th percentiles are shown.

What was the impact of HT interference on APs in Bldg 1?

We ran the interference detection measurements across APs in Bldg 1 for a period of two weeks (Figure 11). Across the 14 WiSe APs in Bldg 1, our analysis detected the occurrence of HT interference for 7 APs. Figure 11 (center) shows the number of days during which we detected at least 5 minutes of HT interference across different APs. During the two weeks, APs 6, 10 and 11 experienced HT interference between 5 - 8 days, indicating that some APs were repeatedly impacted by nearby WiFi interferers across multiple days. The median duration of interference per day (Figure 11, top) varied between 3 to 7 minutes but some APs (2, 6 and 11) experienced higher periods of interference (maximum of 23 - 87 minutes) during this two week period. Figure 11 compares the Witt of the impacted APs during periods with and without HT interference (within 5 minutes of HT event). The 75th percentile values dropped between 0.5 Mbps (AP 2) to 4.3 Mbps (AP 8) indicating high variation in impact of interference across some APs. AP 2 experienced minor reduction in throughput due to interference because, even though the loss rates for packets overlapped by the interferer increased by 40%, the proportion of these packets were low ($< 20\%$) compared to the total transmitted packets.

Our analysis shows the HT interference can occur intermittently, mostly for short periods of time. The occurrence of such HT interference is a property of both the receivers' and interferers' traffic. We observed high burstiness of WiFi links in home environments. For example, in our deployment only about 10% of total periods of continuous activity at the WiSe APs lasted more than 3 minutes at most APs (Figure 12). This is one of the reasons for the small periods of interference in homes. The length of these bursty periods are dictated by the nature of the underlying traffic and can be highly variable depending on the type of traffic. For example, we observed highest periods of continuous activity during video stream-




Figure 12: Distribution of "on-periods" across different WiSe APs. We represent a consecutive period of 10 second intervals with more than 100 data packets each as the *on-period* of the APs.

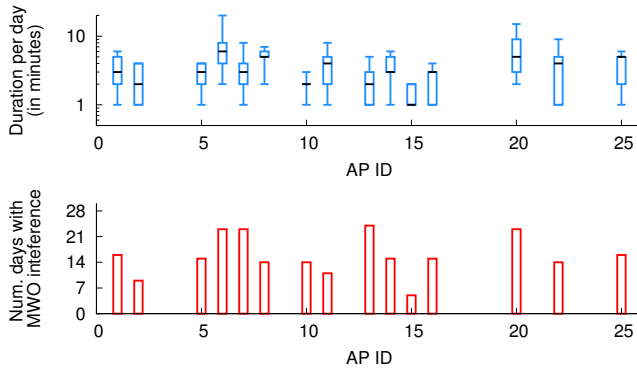


Figure 13: (top) Distribution of duration per day during which microwave interference reduced Witt by 20%. (bottom) Number of days from a 30 day period, during which the APs experienced at least 1 minute of such interference.

ing sessions (e.g., Netflix) which can download data in the order of 50 MB while buffering [15]. We found that, at APs 6 and 11, the periods of highest HT interference (39 and 87 minutes respectively) coincided with the usage of Netflix. The APs are more sensitive to such issues during these continuous periods of high activity, while short periods with bursty traffic are less likely to experience these performance issues from external sources due to the lower volume of the transmitted traffic.

5.3 Non WiFi interference activity

Another factor that can degrade the performance of WiFi links in homes is interference caused by commonly available non-WiFi devices, such as microwave ovens, Cordless Phones etc. Unlike WiFi transmissions, these devices do not sense the medium before transmitting energy into the spectrum. Different non-WiFi devices can impact WiFi links differently based on their transmit power, transmission protocol as well as their distance from the WiFi transmitters and receivers. By running Airshark [16], WiSe APs detect the presence of non-WiFi devices operating in their vicinity and report them to the controller (§2.2). In this section, we report the properties of non-WiFi interference in homes and quantify their impact on the Witt across the APs.

How did non-WiFi interference impact nearby WiFi links?

In this section, we focus on the impact of microwave oven devices on nearby links, since they were the most ubiquitous non-WiFi interferers in our deployment. Microwave ovens have a duty cycle of 50% with alternative periods of approximately 8ms of active and quiet periods (60Hz cycle). We compared the Witt for the active WiSe AP-client pairs with and without microwave oven activity

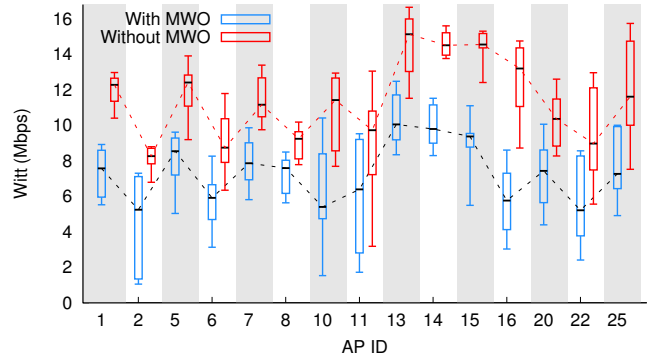


Figure 14: Candlestick plot comparing the estimated average Witt for WiSe APs with active WiFi links during and after the completion of microwave oven activity. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure. The dotted line compares the 50th percentile values for both cases.

(within 5 minutes of the activity). This allowed us to quantify their impact on these links by comparing their performance during these two periods. Since microwave ovens mainly impact channels 8-11, we analyzed this data for a period of 30 days during which the APs were configured to use channel 11.

Figure 13 shows the duration per day and the number of days (from the 30 day period) during which different APs experienced at least 20% degradation of Witt in the presence of microwave oven interference during active periods. While these periods are short, they can cause significant reduction in Witt at some APs (e.g., AP 2, Figure 14).

Figure 14 compares the estimated average Witt for different WiSe APs during active periods with microwave oven activity and active periods without microwave oven activity (within 5 minutes). The dotted line compares the 50th percentile values of the Witt for both periods to focus on the performance during periods of high interference. Some APs (e.g., 2, 5 and 16) experienced high performance losses during periods of microwave activity. For example, the 25th percentile Witt dropped from 7.8 Mbps to 1.5 Mbps for AP 2 (in Bldg 1), indicating a performance drop of 81% due to microwave oven activity.

Impact on airtime and effective rate. To provide greater insight into the causes of the low Witt measured at some APs during periods of interference, Figure 15 shows the impact of microwave ovens on factors such as airtime and effective rate (§3). Figure 15 (left) shows the average airtime utilization (10 second average) across APs during the two periods. Some APs reported significant increase in airtime utilization in the presence of microwave ovens. For example, the 75th percentile value of airtime utilization for AP 16 increased from 0.33 to 0.70, indicating an absolute increase in airtime utilization of 37 percentage points due to the microwave oven. Other APs, such as AP 10, 11 and 22 also experienced high airtime utilization in the presence of microwave ovens (over 60%).

Figure 15 (right) shows the impact of microwave ovens interference on the effective rates. At some APs (e.g., APs 13-15), there was little reduction during most instances of interference. Thus, these APs were mostly impacted by higher airtime utilization at the sender. Other APs, such as APs 2 and 16 experienced high reduction for the effective rates due to microwave oven interference. For these APs, both high airtime utilization and packet losses contributed to a lower Witt. This analysis shows the diversity of impact of microwave

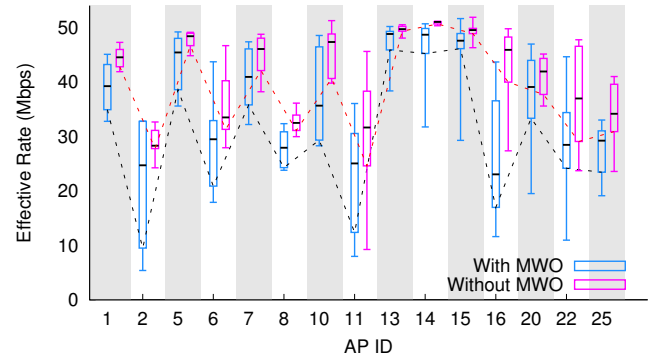
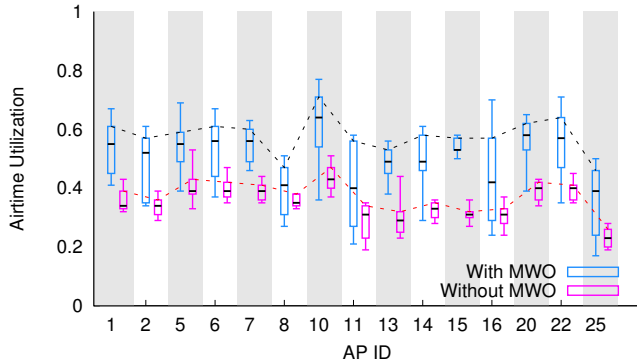


Figure 15: Candlestick plot comparing the airtime utilization (left) and effective rates (right) for WiSe APs with active WiFi links with and without microwave oven interference. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure. The dotted line compares the 25th and 75th percentile values for airtime utilization and effective rate respectively.




Figure 16: Average number of microwave instances detected per 5 minute bin (on channel 11) at 2 WiSe APs experiencing low (AP 15) and high (AP 22) microwave oven activity (30 day period).

ovens on WiFi links. Using different mitigation techniques based on the interferer type (e.g. channel hopping vs. fixed frequency) can be helpful for APs that experience high degradation in the link’s performance. As discussed next, leveraging context information about non-WiFi device activity (e.g., time of day) can also be helpful in avoiding interference from these devices when they are more likely to be used.

How does microwave oven activity vary across locations?

Figure 16 provides an example of microwave oven activity at two different APs (channel 11). For this analysis, we computed the average number of microwave oven instances seen per 5-minute bin at each WiSe AP during the one month period. A value of 1 on the y-axis indicates that an average of 1 microwave oven instance was reported during the specific 5-minute bin of day over the one month period. AP 22 observed the highest activity during 11am - 1:30 pm and 5pm - 8pm. On the other hand, AP 15 observed low frequency of microwave oven activity on the same channel during most of the day except 10:30 am - noon. Learning about this “time of day” context about non-WiFi device activity can enable home APs to make better decisions in order to avoid the impact of such devices during periods of highest non-WiFi device activity. We plan to further explore this aspect in our future work in building mitigation mechanisms.

5.4 Channel Usage Patterns

Our deployment of WiSe APs across homes showed that WiFi links in dense deployments occasionally experience poor performance due to external sources. Changing channels can partially help in dealing with repeated performance issues on a channel (e.g., non-WiFi interference on a single channel).

To study the usage of channel selection algorithms by other home APs, we configured the WiSe APs to periodically scan all WiFi channels to overhear beacons from neighboring APs, especially to

Unique channels	Num. APs (Overall)	Percentage (Overall)	Num. APs (Bldg 1 only)	Percentage (Bldg 1 only)
1	171	56.1%	99	61.1%
2	61	20%	33	20.4%
>=3	73	23.9%	30	18.3%

Table 9: Number of unique channels used by neighboring external APs as observed by the WiSe APs over a 1 month period. The table shows the overall values as well as the APs specifically observed in Bldg 1.

include external APs, on different channel. Table 9 shows the number of distinct WiFi channels used by these external neighboring APs observed by all WiSe APs as well for Bldg 1 only over a period of one month. It shows that around 56% of the overall 305 APs observed, used a single WiFi channel for the entire period, indicating the fact the majority of home APs used a static WiFi configuration, and are never re-assigned by residents after they are deployed. These channel configuration patterns were prevalent across different locations (as shown for Bldg 1), indicating that this property was not biased towards a particular location.

Using a static channel assignment may not be an issue in low density deployments or if the link’s experience is good for most of the time. But, as discussed in §4, some APs observed lower Witt values for various reasons. Such APs can benefit from channel re-assignments based on the performance experienced on the current channel.

5.5 Summary

1. The impact of interference (WiFi and non-WiFi) is dependent on both the link’s and interferer’s traffic. Majority of the interference durations were short (1 - 7 minutes) due to the bursty nature of traffic at homes. But, some links experienced extended periods of interference (tens of minutes) during periods of continuous activity due to either high airtime utilization (at sender) or packet losses (at receiver).
2. Even though most interference periods were short, some had a high impact on the APs. For example, microwave ovens caused high degradation of Witt at some APs (81%).
3. Learning context (e.g., time of day) about interference activity (e.g., periods of non-WiFi device usage at homes) can enable APs to avoid interference from such devices.
4. Majority of APs (56%) observed at homes used static channel configurations over time (30 days) indicating that they rarely or never get configured once they get deployed.

6. RELATED WORK

Characterizing network performance. The BISmark project [20] aims at understanding the performance of wired access networks in homes through the long-term deployment of gateways in homes serviced by a diverse set of ISPs. We use a similar approach but focus on building a framework to study the properties of dense residential wireless networks and deploy multiple APs within the apartment building. Previous studies [11, 4, 12, 3] have evaluated wireless network deployments in enterprises and homes by collecting passive traces or user traffic statistics. For example, Jigsaw [4] used a large number of passive sniffers in an enterprise WLAN to collect wireless packet traces for debugging wireless problems. Instead of using passive sniffers, we perform measurements directly through the wireless APs. This allows us to obtain more comprehensive information about the AP's view of the network in addition to packet level statistics (e.g., airtime utilization, non-WiFi devices etc.). By performing controlled experiments in three houses, authors in [14] observed high asymmetry in links and variability in link performance at different locations within homes. We measure the link performance in homes over time and identify the causes of occasional poor performance experienced by links. Prior work such as [5, 7] has focused on end-host techniques for measuring end to end performance in home networks. We use AP-centric techniques, since it allows us to monitor home WLANs by using a single instrumented vantage point.

Wireless network debugging and diagnosis. Prior work on model based techniques [10] require active network measurements to predict link capacity and build conflict graphs. Our focus is to passively estimate a wireless link's performance by only using coarse-grained local observations, since any additional traffic from the WiSe APs may impact the existing traffic. To avoid measurement overhead, passive techniques have been proposed [17, 18, 13] for detecting the presence of interference in WLANs. We build upon prior work from PIE [18] that uses microsecond level timing information to passively detect the presence of wireless interference for a link by comparing the link's loss rates in presence and absence of overlapping packets from a nearby WiFi sender. We ported Airshark [16] to the WiSe platform to detect the presence of non-WiFi devices near WiSe APs and measure their impact on user traffic.

7. CONCLUSION

We describe a unique measurement and management infrastructure, WiSe, to perform inline measurements of wireless properties in homes that uses APs as vantage points. We have been operating this infrastructure for more than 6 months in 30 homes. The deployed homes have diverse characteristics — some are in dense apartments while others are in sparser suburban neighborhoods, some residents directly own the APs while others use a shared infrastructure deployed in the apartment. We also present a simple metric to estimate the wireless TCP throughput in these deployments. Our infrastructure operates through an open API and so in the future other types of APs can be added to the infrastructure. We believe studies such as these can provide our community with valuable ground truth data. Hence, we also plan to release our anonymized wireless traces broadly to the community in the near future.

8. ACKNOWLEDGMENTS

We thank the participants in this study for their support. We also thank Nathan Moeller for his assistance in the initial stages of this deployment. We would like to acknowledge the anonymous reviewers and our shepherd, Alex Snoeren, whose comments helped bring the paper into its final form. Srinivas Govindan and Suman Banerjee are supported in part by the US National Science Founda-

tion through awards CNS-1040648, CNS-0916955, CNS-0855201, CNS-0747177, CNS-1064944, and CNS-1059306.

9. REFERENCES

- [1] "OpenWrt". <https://openwrt.org/>.
- [2] "SamKnows". <http://www.samknows.com/broadband/index.php>.
- [3] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in Chaotic Wireless Deployments. *MobiCom '05*.
- [4] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis. *SIGCOMM '06*.
- [5] L. DiCioccio, R. Teixeira, and C. Rosenberg. Impact of home networks on end-to-end performance: controlled experiments. *HomeNets '10*.
- [6] P. Engines. alix2d2. <http://www.pceengines.ch/alix2d2.htm>.
- [7] C. Gkantsidis, T. Karagiannis, P. Key, B. Radunovic, E. Raftopoulos, and D. Manjunath. Traffic management and resource allocation in small wired/wireless networks. *CoNEXT '09*.
- [8] H. Guru. "Do Wireless HDTV Systems Do The Job?". <http://hdguru.com/do-wireless-hdtv-systems-do-the-job-latest-models-reviewed/6178/>.
- [9] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding congestion in IEEE 802.11b wireless networks. *IMC '05*.
- [10] A. Kashyap, S. Ganguly, and S. R. Das. A measurement-based approach to modeling link capacity in 802.11-based wireless networks. *MobiCom '07*.
- [11] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. *MobiCom '02*.
- [12] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the mac-level behavior of wireless networks in the wild. *SIGCOMM '06*.
- [13] J. Manweiler, P. Franklin, and R. Choudhury. RxIP: Monitoring the health of home wireless networks. In *INFOCOM*, 2012.
- [14] K. Papagiannaki, M. Yarvis, and W. S. Conner. Experimental Characterization of Home Wireless Networks and Design Implications. In *INFOCOM* 2006.
- [15] A. Rao, A. Legout, Y.-s. Lim, D. Towsley, C. Barakat, and W. Dabbous. Network characteristics of video streaming traffic. *CoNEXT '11*.
- [16] S. Rayanchu, A. Patro, and S. Banerjee. Airshark: Detecting non-WiFi RF devices using commodity WiFi hardware. *IMC '11*.
- [17] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. *SIGCOMM '06*.
- [18] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki. PIE in the sky: online passive interference estimation for enterprise WLANs. *NSDI '11*.
- [19] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee. 802.11n under the microscope. *IMC '08*.
- [20] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet Performance: A View From The Gateway. *SIGCOMM '11*.
- [21] Tan, Godfrey and Gutttag, John. Time-based fairness improves performance in multi-rate WLANs. In *Proceedings of the annual conference on USENIX ATC, 2004*.