Observing Home Wireless Experience through WiFi APs

MobiCom '13 September 2013

A. Patro, S. Govindan, S. Banerjee
University of Wisconsin Madison

Presented by Bob Kinicki

Advanced Computer Networks
Fall 2013

Motivation

- Generally, while home WiFi users get reasonably good performance most of the time, there remain instances when home network performance remains frustratingly slow.
- Most researchers over the last decade have deployed passive sniffers to understand and evaluate specific wireless characteristics.

Research Goals

1. To perform a more systematic study of WiFi experience in home environments and provide a detailed characterization.

2. To evaluate our community's collective intuition of WiFi network performance.

Objectives

To answer these questions:

- How often does home WiFi provide good, mediocre or bad performance?
- When performance is bad –what are the causes and how long does it persist?
- How much interference do we see and what sources provide the interference?
- How do users configure their WiFi networks?

Research Approach

- Define a wireless performance metric that captures overall network goodness.
- This metric should consider ONLY wireless part of user's end-to-end path.
- Metric is "application-agnostic" while focusing on TCP elasticity.
- Witt :: WiFi-based TCP throughput
- Evaluate and use Witt as a key metric in wireless measurement study.

Outline

- Introduction
- WiSe Infrastructure and Framework
- How was Witt constructed?
- Use Witt to Classify WiFi Experience
- Analyze detailed Results from Measurement Study
  - To answer the posed questions.
- Summary and Critique

WiSe Measurement Framework

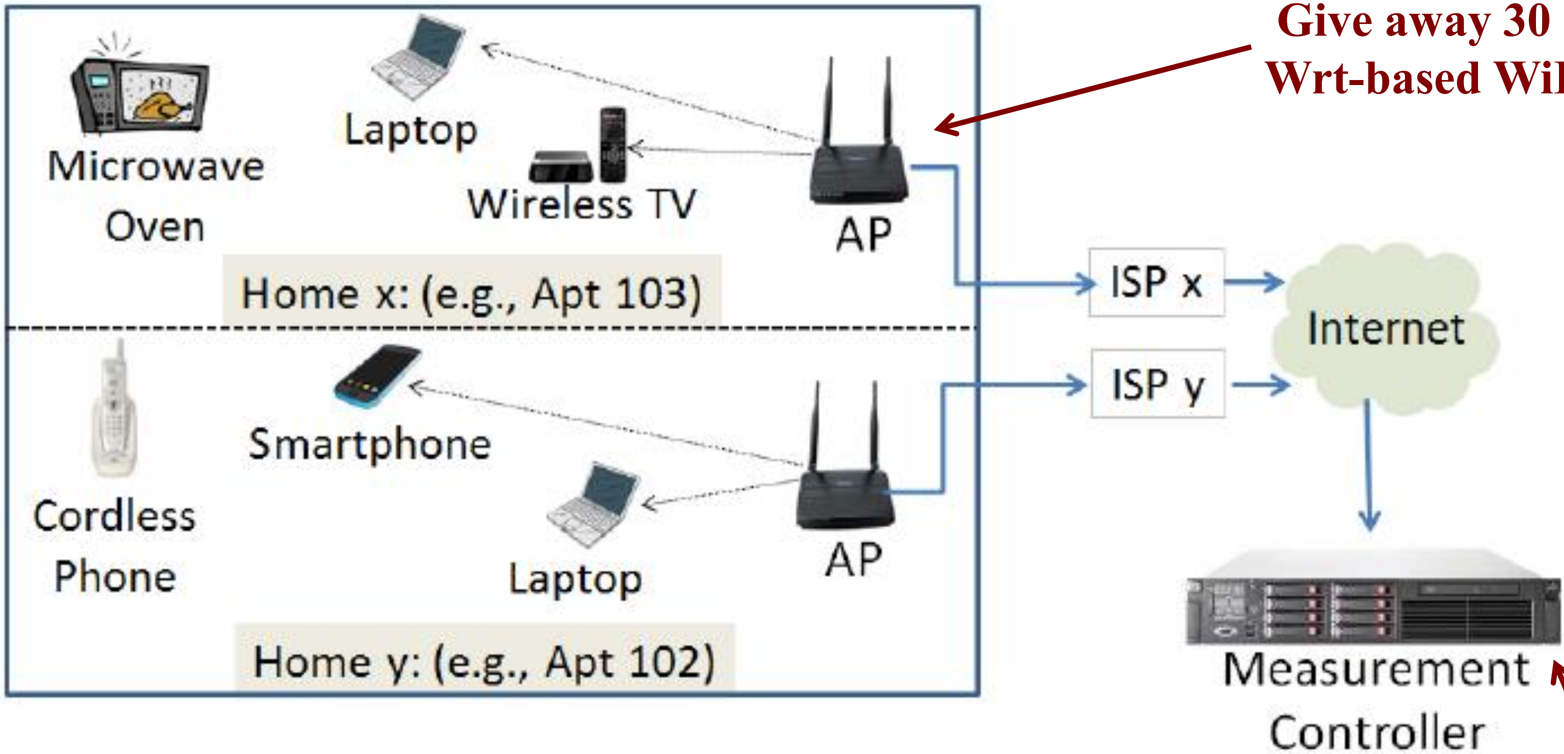

Figure 1: WiSe measurement framework.

Give away 30 Open Wrt-based WiFi APs

Uses open API to remotely manage and configure APs

Open API High Level Description

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. **(Every 10 secs.)**
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, recv, 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.

10-byte per packet summaries including AP's own links and overheard data packets on the same channel

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.

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

Wide-Ranging Daily WiFi Usage

8-9 GB per day
(90th percentile usage)

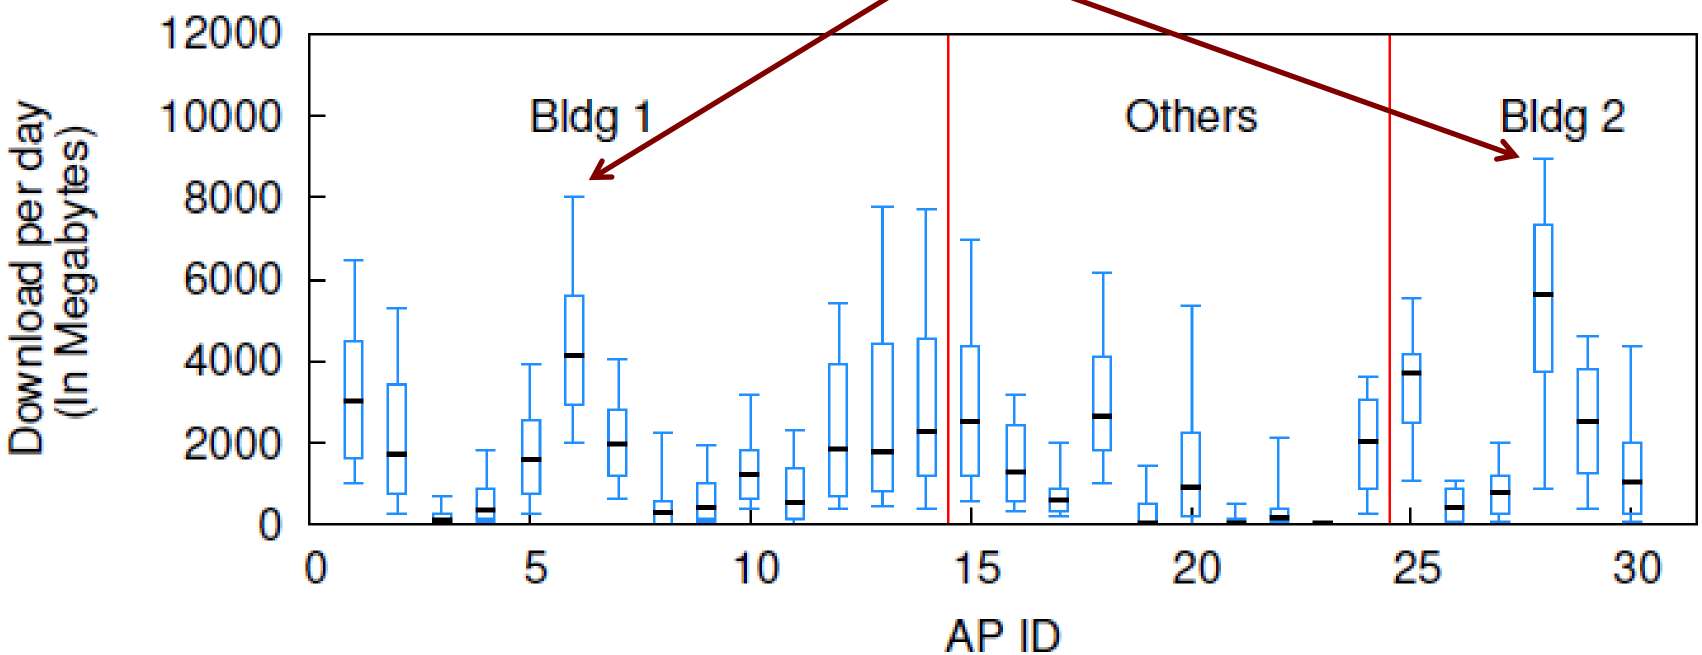

Figure 3: Distribution of the daily data download over the WiFi network per WiSe AP (Sep 1, 2012 - Jan 31, 2013). 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure.

Outline

- Introduction
- WiSe Infrastructure and Framework
- How was Witt constructed?
- Use Witt to Classify WiFi Experience
- Analyze detailed Results from Measurement Study
  - To answer the posed questions.
- Summary and Critique

Witt: WiFi-based TCP Throughput

- Metric idea – measure (passively at the AP) the likely TCP throughput between a client and its AP given the existing wireless conditions.
- Consider also the average value for all active clients* as a single aggregate for the entire AP.

*To be considered active, a client has to send at least 500 packets in the last 10-second window.

Degradation Factors and Indicators

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.

More details wrt factors: low signal strength, increased delay due to reduced PHYrates or multiple retransmissions, high airtime reduces ability to send

How to measure Witt?

Collect 'ground truth' measurements under a variety of conditions.

- Four clients (laptops) co-existed at eight different deployment locations in the apartment buildings.
- Iperf TCP download run between APs and clients for 20 seconds.
- Clients ran throughput measurements in intervals of 5 to 10 minutes over the course of a week.

How to measure Witt?

- – Clients were connected to different APs to emulate different link conditions.
- – Experiments automatically conducted at different times of the day.
- – Collected hundreds of measurements (see Table 6).

- Based on the key factors from the measurements build a model of Witt.
- Use benchmarks to evaluate Witt.

Ground Truth Measurements

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.

Feature Definitions

- Airtime utilization :: aggregate busy statistic that seems to include time when transmitting and receiving and overhearing.
- Local contention (c) :: the relative amount of other client traffic through an AP as a fraction of the total traffic passing through this AP.
- Effective rate (r) :: captures the net effect of packet losses and choice of PHY rate used on an AP-client link. (see equation 1)
- Link experience (link_exp) :: (see equation 2)

Effective Rate and Link Experience

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

$$\text{link\_exp} = (1 - a) * (1 - c) * r, \ 0 \leq a \leq 1, \ 0 \leq c \leq 1 \qquad (2)$$

where*

s_i is the number of successful packet transmissions and p_i is the total number of packet transmissions at each PHY rate (r_1, …, r_n) used by an AP-client pair.

a is the airtime utliization from external sources.

*Note – all features are based on aggregate stats per link (collected over 10 second intervals).

How to Create Witt ?

To create Witt, wireless statistics recorded by WiSe APs in 10-sec intervals were evaluated via correlations as potential important features to used to predict Witt.

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.

Build a linear model of Witt

Link experience is mapped to Witt using a linear model (equation 3).

$$Witt = \beta_1 * \text{link\_exp} + \beta_0 \qquad (3)$$

By dividing ground truth data into training and testing data sets, authors test fidelity of linear model and develop 95% confidence intervals that show model is a reasonable estimate for predicting throughput.

	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.

Benchmarks to Evaluate Witt

- Ground truth TCP throughput measurements are compared against predicted TCP throughputs using linear regression of effective rate and link experience (see Figure 4).
- CDF of errors between actual vs predicted TCP throughputs using different metrics (see Figure 5).

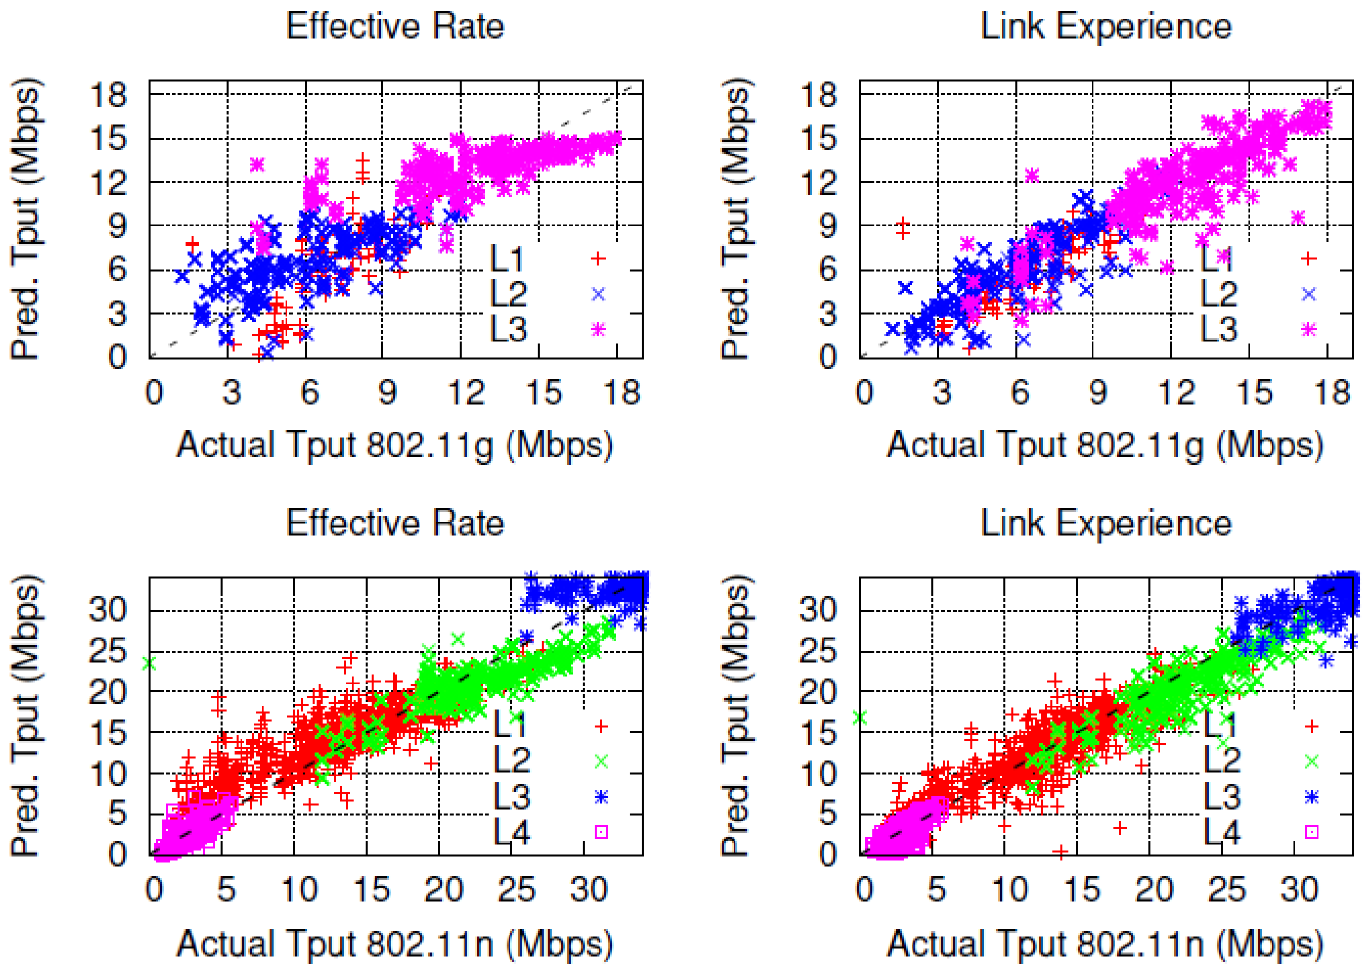

Figure 4: Scatter plot showing the actual vs. predicted TCP throughput values for the 802.11g (top) and 802.11n (bottom) ground truth experiments by using "Effective Rate" and "Link Experience". Points near the "x=y" line indicate instances with accurate prediction of TCP throughput. Each set of points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

CDF of Errors for Various Metrics

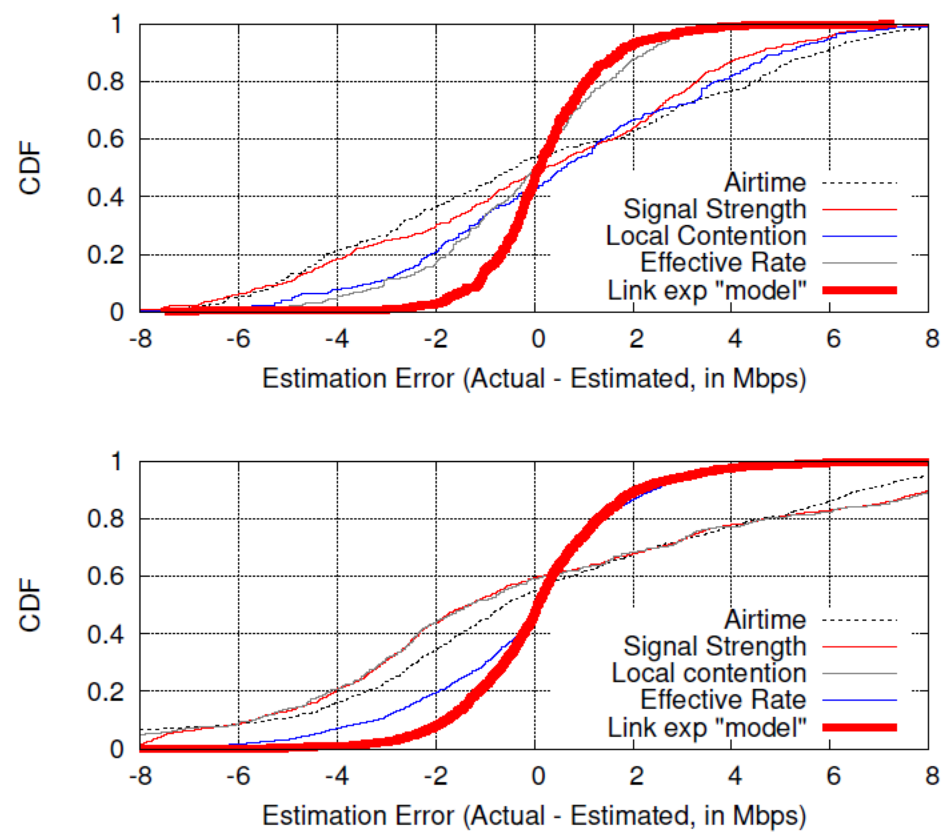

Figure 5: CDF of errors between the actual and predicted TCP throughput values obtained by using different metrics for 802.11g (top) and 802.11n (bottom).

Outline

- Introduction
- WiSe Infrastructure and Framework
- How was Witt constructed?
- Use Witt to Classify WiFi Experience
- Analyze detailed Results from Measurement Study
  – To answer the posed questions.
- Summary and Critique

Use Witt to Classify Wireless Experience

- Focus on periods when WiSe AP has at least one active client.

How did link performance vary across APs over time?

- A diverse set of clients associated with WiSe APs.
- Measured Witt values during active periods and group results bases on Witt values.
- In Figure 6 – AP clients are active for at least 20 days.

Figure 6

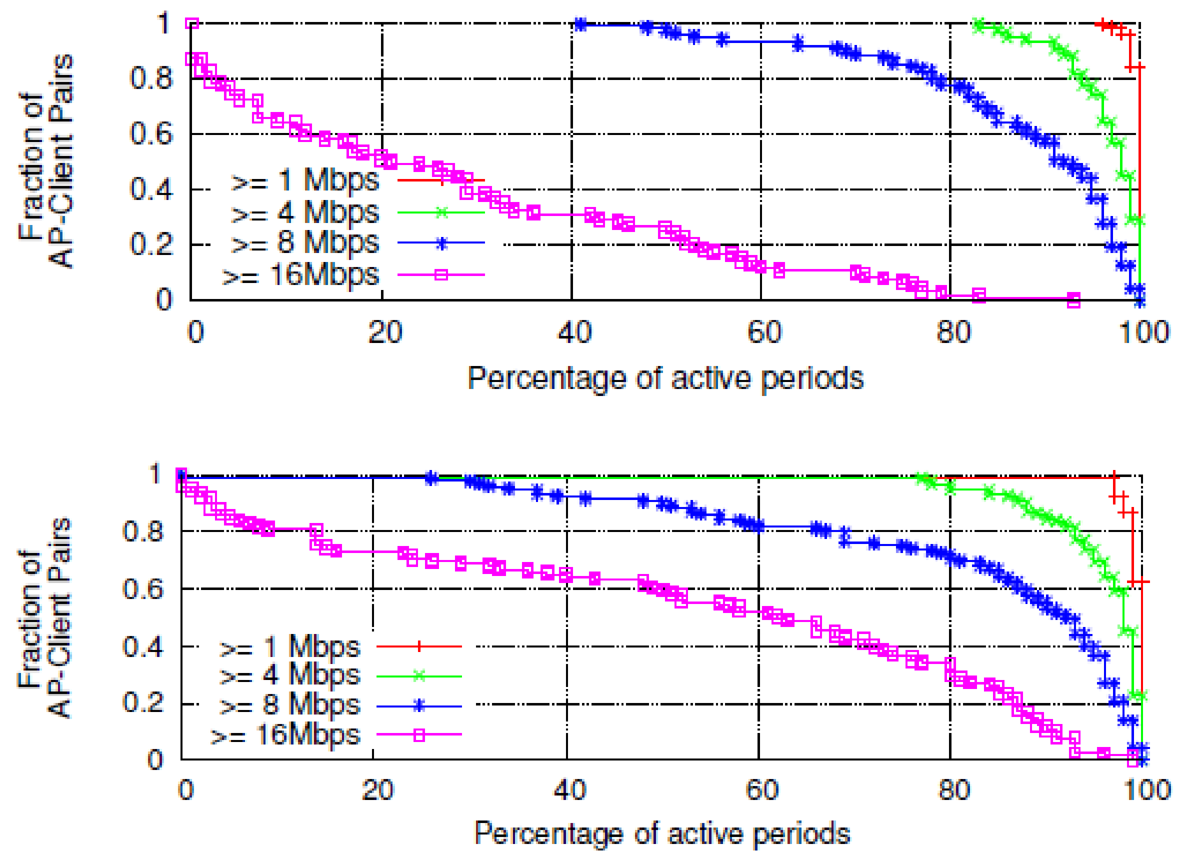

Figure 6: Distribution of Witt across AP-client pairs using 802.11g (top) and 802.11n (bottom) in our deployment that were active for at least 20 days.

Table 7 & Figure 7

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.

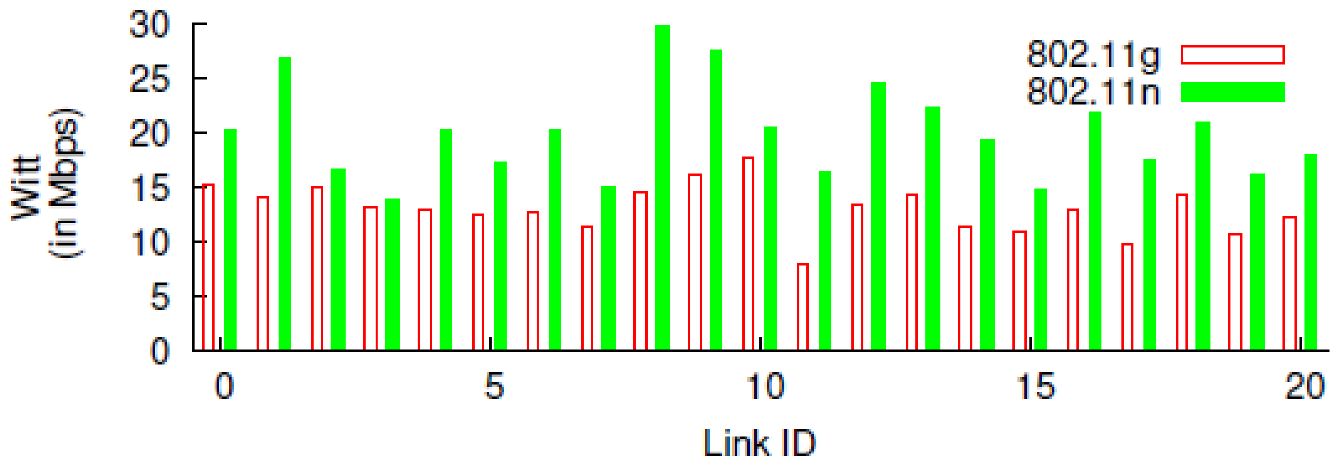

Figure 7: Comparison of average Witt for AP - client pairs over 802.11g and 802.11n (3 days each).

Causes for Poor Wireless Experience

- Over 80 days (Nov 2012 – Jan 2013) detected 186 and 2031 minutes of "Very Poor" and "Poor" instances across all 30 WiSe APs (2.1% of the active periods).
  - Very poor periods rare; Poor periods occur intermittently depending on link and location.
- Aggregate instances of poor performance across WiSe APs in each apartment.

Causes of Poor Experience

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 ($R\downarrow, <= 12Mbps$). Others correspond to remaining combination of factors such as high contention, signal strengths etc.

High density produces higher airtime and losses.

High frame losses cause poor results.

Low signal strength yields low performance.

Impact of Other Factors

- Impact of other factors (including local contention from other clients) was low (<= 4.3%).
- Prevalence of low local contention at wireless hop is due to it is uncommon for multiple clients to generate high traffic during the same interval.
- In cases where there were multiple active clients at AP, bottleneck at the wired link led to lower contention.

Variability in Wireless Experience

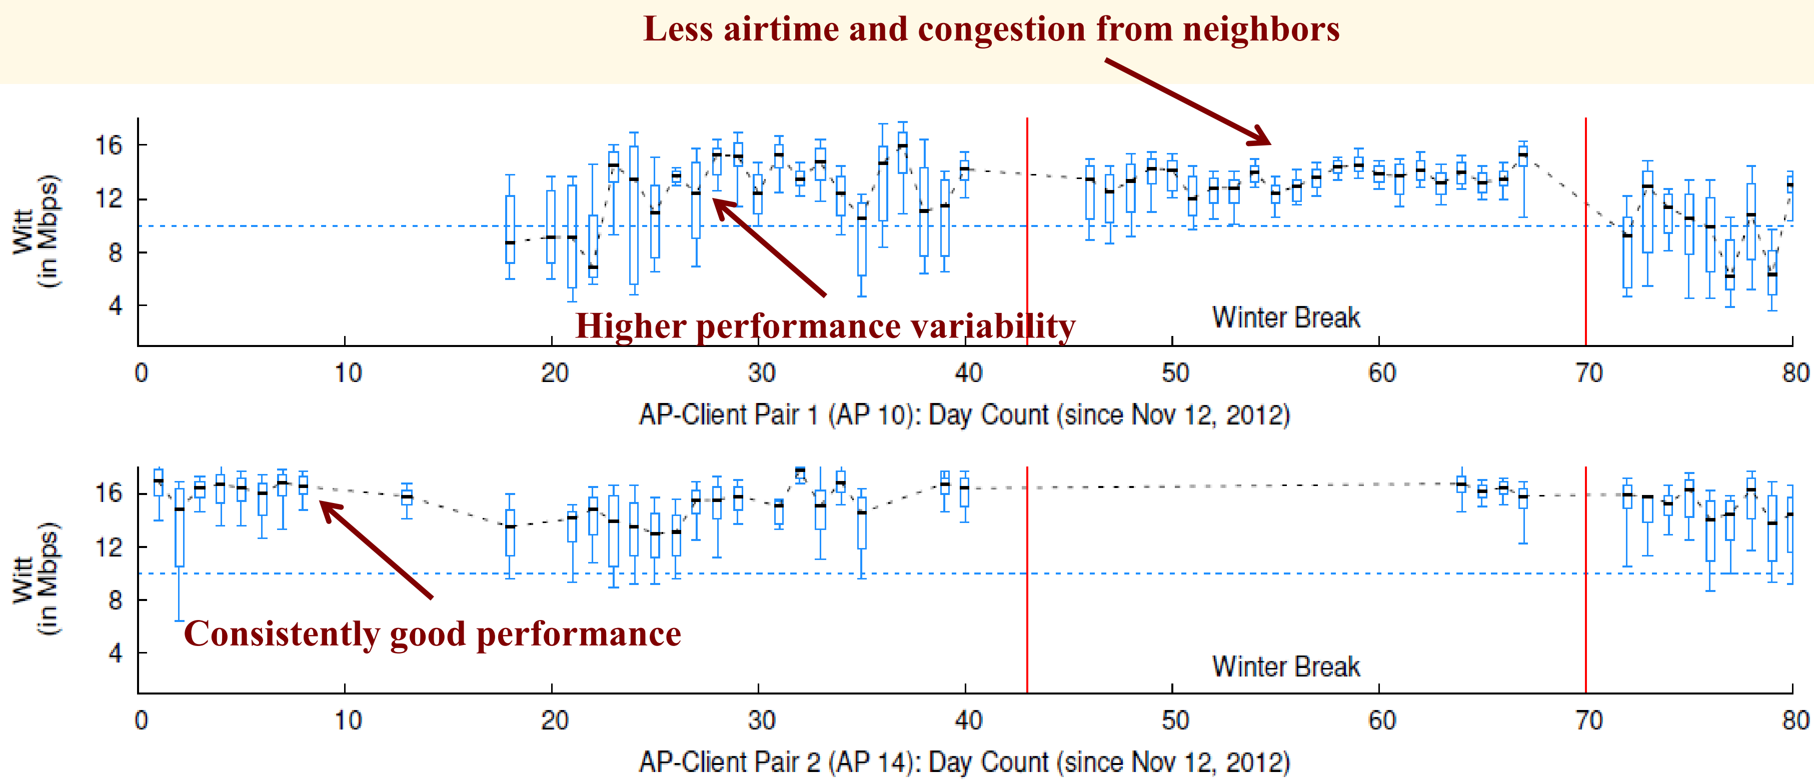

Figure 8: Distribution of Witt for 2 different AP-client pairs in our deployment from 12 Nov 2012 to Jan 31, 2013 and shows their variation in performance over the span of this period. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure.

Outline

- Introduction
- WiSe Infrastructure and Framework
- How was Witt constructed?
- Use Witt to Classify WiFi Experience
- Analyze detailed Results from Measurement Study
  - To answer the posed questions.
- Summary and Critique

Detailed View

Analyze impact of external factors on wireless clients in the wild:

- Contention from low data rate senders
- Packet loss due to hidden terminals
- Non-WiFi interference activity

Contention from low data rate senders

Due to performance anomaly (aka rate anomaly) transmitters using low PHY rates can cause their Witt to suffer.

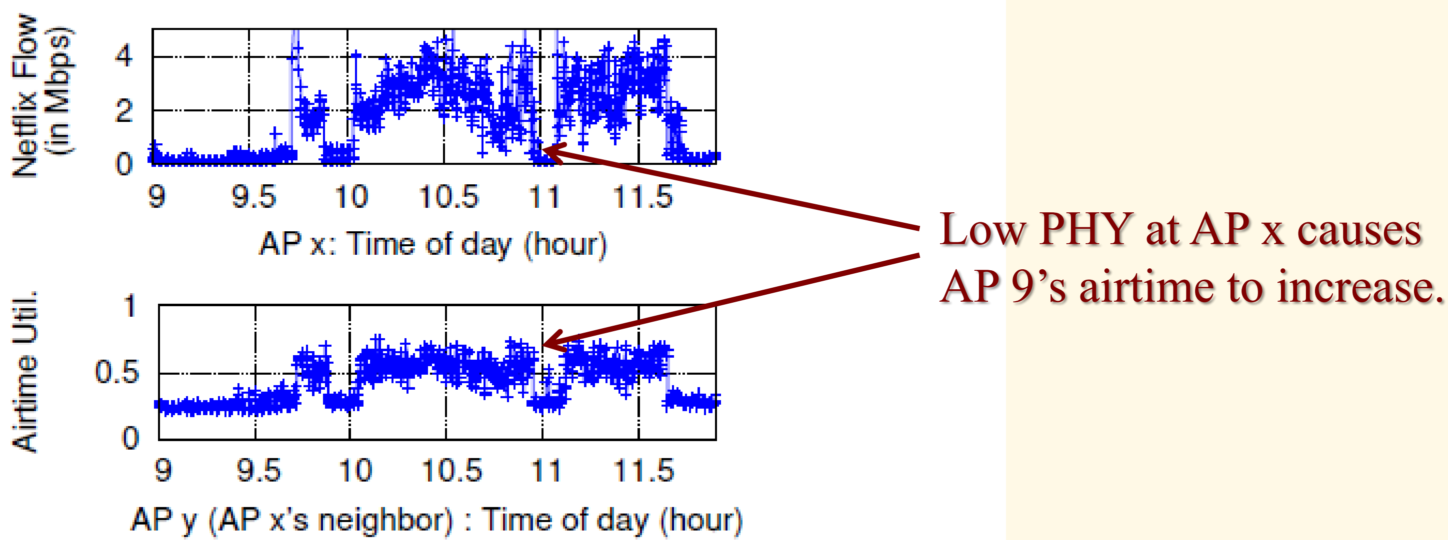

Figure 9: Time-series airtime utilization at AP 9 with and without the presence of traffic from an external AP that used low PHY rates. AP 9 was inactive during this period.

AP 6 had highest activity in Figure 3

Impact of low PHY from external APs

Impact of contention

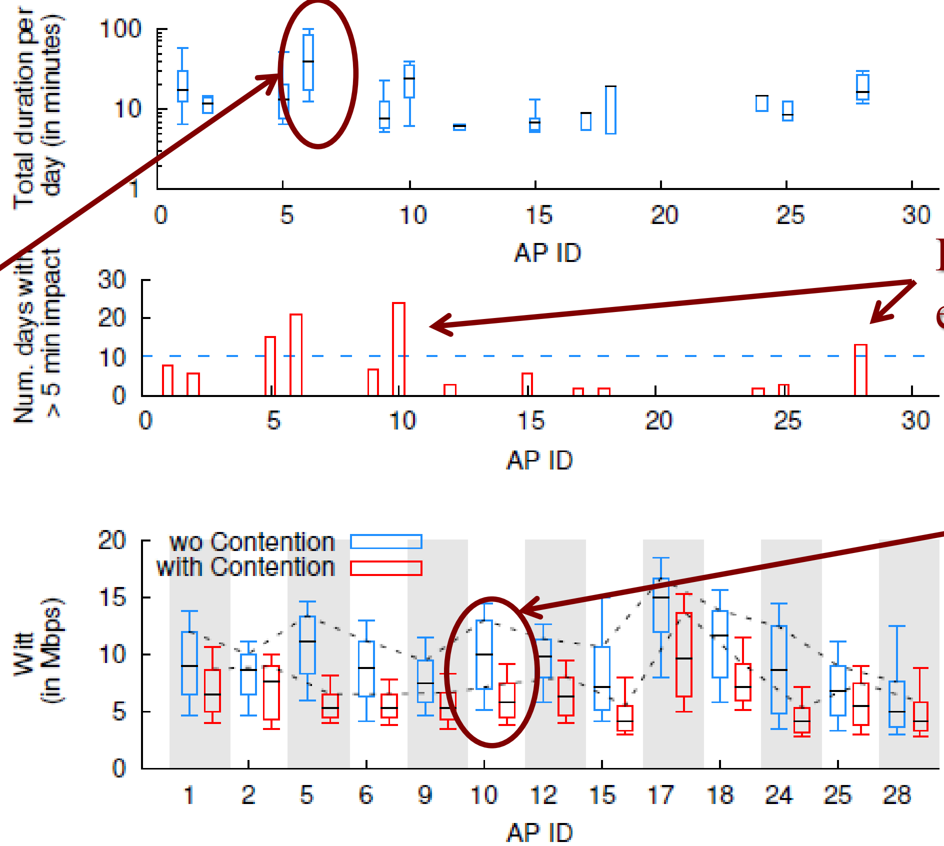

Figure 10: (top) Candlestick graph showing the distribution of active periods per day (Y-axis in logscale) during which the WiSe APs experienced contention from external transmitters (using average PHY rate <= 15 Mbps), (center) Bar graph showing the total number of days during which WiSe APs experienced 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.

Packet Loss due to Hidden Terminals

- High packet loss was a major cause for “Poor” cases.
- Hidden Terminals (HT) are an external factor that can reduce link’s Witt by increasing packet loss.
- Used synchronized and merged packet summaries from multiple APs in Bldg 1 to compute HT events in 15-second epochs.
- Packet loss at a receiver due to overlapping packet transmissions from the interferer is the main cause for a hidden terminal event.

Packet Loss due to Hidden Terminals

- For 15 second "epochs", epoch is marked as HT event for WiSe AP when one of its link's loss rates is 40% higher for packets overlapped in time by the interferer compared to packets not overlapped by any other transmitter.
- Required constraint of 1000 packet minimum for a link and minimum of 100 packet overlaps from potential interferer per epoch to check for HT conflict (makes this a conservative estimate of interference experienced).

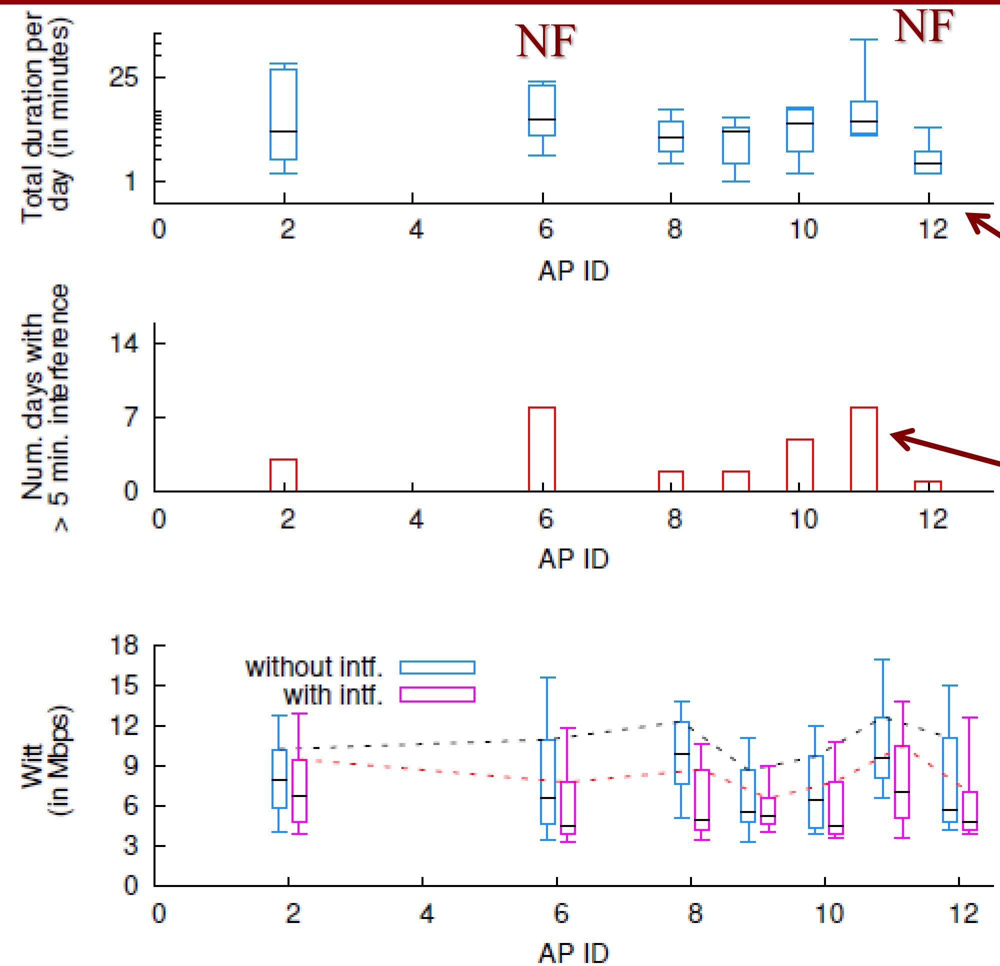

HT interference for seven APs

6, 10, 11 repeatedly impacted by HTs

High variability in HT impact due high burstiness of WiFi links

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.

High Burstiness of Traffic

- Only about 10% of total periods of continuous activity at the WiSe APs lasted more than three minutes.

{explains small periods of interference in homes}

Example – Netflix video streaming

APs 6 and 11 (NF in Figure 11) periods of highest interference coincided with usage of Netflix. APs are more sensitive to HT interference issues during periods of high activity.

Non-WiFi Interference Activity

- Interference by commonly available non-WiFi devices can degrade WiFi link performance (e.g,. Microwaves).
- These devices do NOT have carrier sense before transmitting.
- Authors use **Airshark** to detect presense of non-WiFi devices.
- Since microwaves impact channels 8-11, conducted 30 day experiment with APs using channel 11.

Figure 13 Microwave Interference

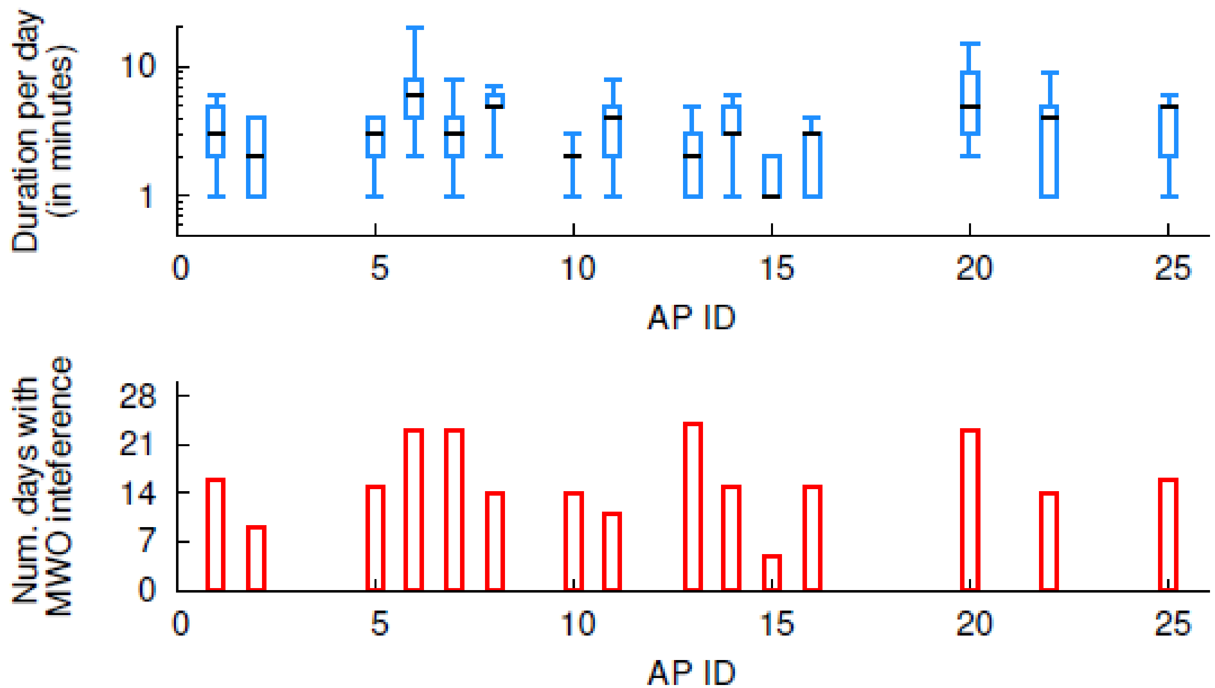

Figure 13: (top) Distribution of duration per day during which microwave interference reduced Witt by 20%. (bottom) Number of days from a 30 day period, during which the APs experienced at least 1 minute of such interference.

Figure 14 Impact of Microwave Activity on Witt

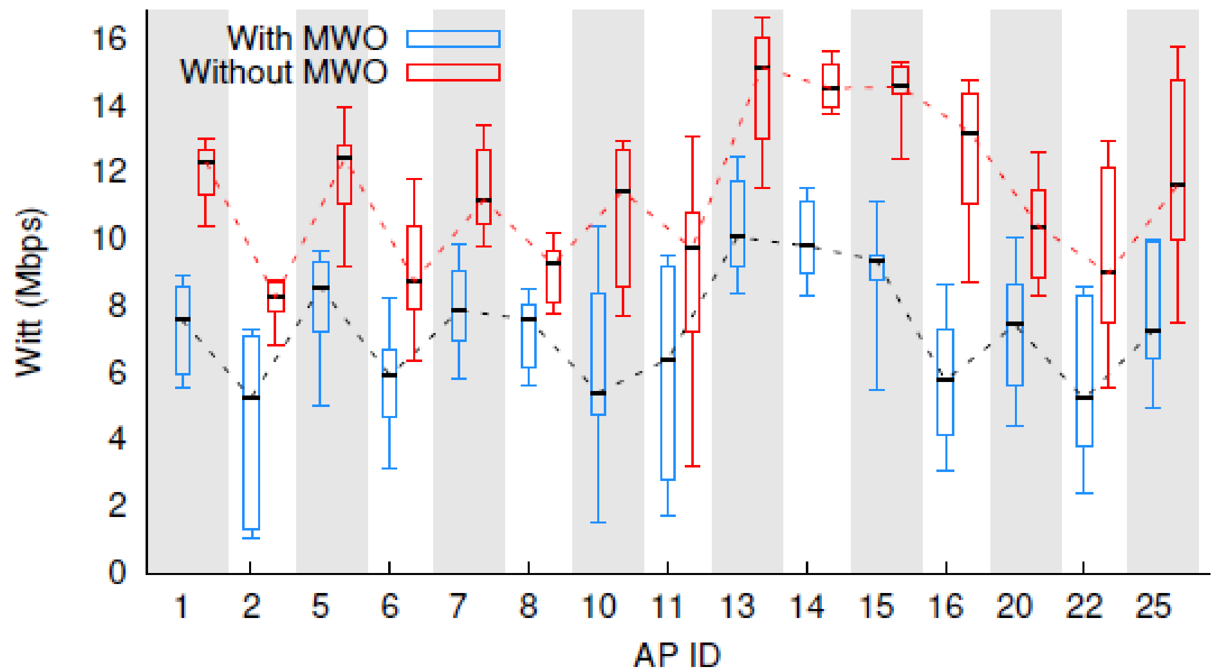

Figure 14: Candlestick plot comparing the estimated average Witt for WiSe APs with active WiFi links during and after the completion of microwave oven activity. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure. The dotted line compares the 50th percentile values for both cases.

Impact of Microwave Activity on Airtime and Effective Rate

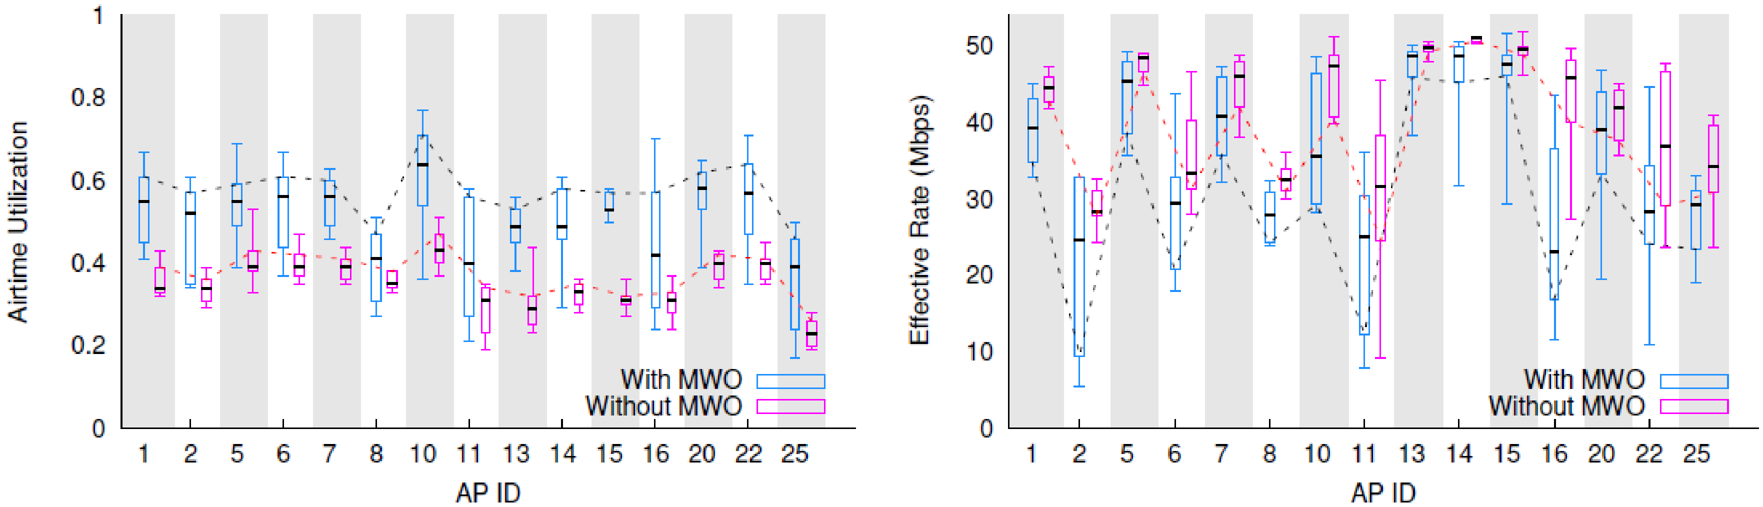

Figure 15: Candlestick plot comparing the airtime utilization (left) and effective rates (right) for WiSe APs with active WiFi links with and without microwave oven interference. 10th, 25th, 50th, 75th and 90th percentiles are shown in this figure. The dotted line compares the 25th and 75th percentile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

Figure 16 Microwave Instances

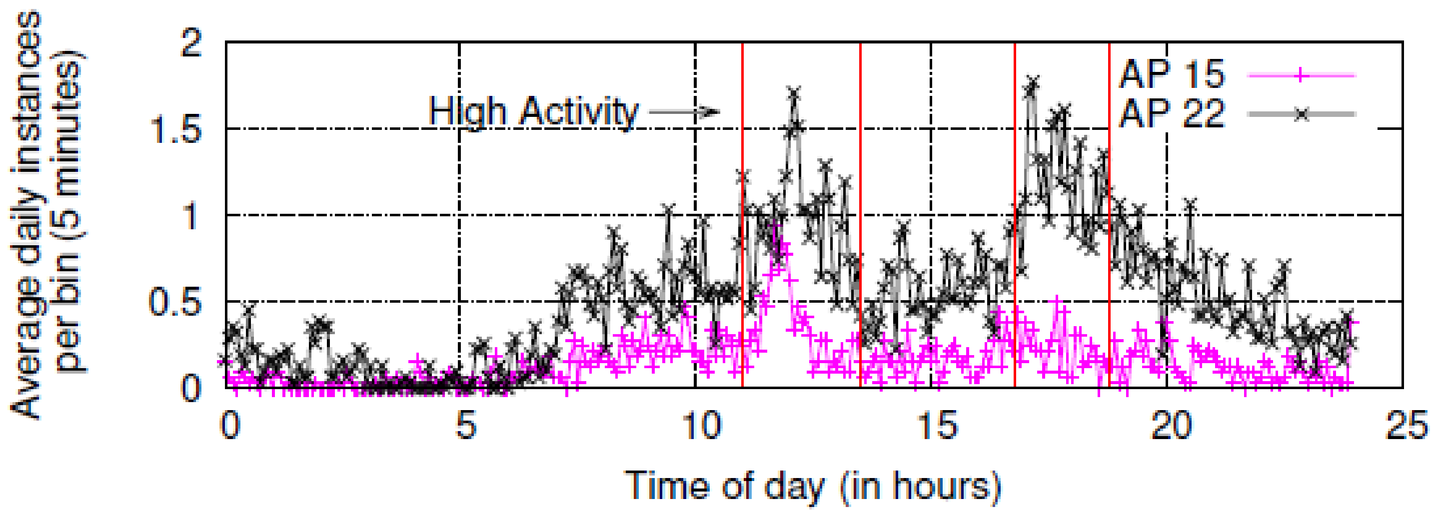

Figure 16: Average number of microwave instances detected per 5 minute bin (on channel 11) at 2 WiSe APs experiencing low (AP 15) and high (AP 22) microwave oven activity (30 day period).

Figure 16 demonstrates differences over APs and over time of day.

Channel Usage Patterns

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.

Done by periodically scanning all channels to overhear beacons from neighboring APs (including external APs).

Section 5 Summary

- Impact of interference (WiFi and non-WiFi) depends on the traffic of both link and interferer. Majority of interference durations are short.
- Some interferers had high impact on the APs (e.g., microwave ovens severely degraded performance of some APs).

Section 5 Summary (cont)

- Learning the context about interference activity (e.g., time of day) can enable APs to avoid interference.
- Majority of APs observed use static channel configurations.

Conclusions

- WiSe APs are used to measurement wireless properties in homes.
- Simple metric, Witt , is developed, tested and used in this investigation.
- Paper provides detailed results about causes of poor performance, contention from low data rate senders, packet loss caused by hidden terminals, and interference due to non-WiFi devices.

Critique

Did authors

answer these questions:

- How often does home WiFi provide good, mediocre or bad performance? Y
- When performance is bad –what are the causes and how long does it persist? Y
- How much interference do we see and what sources provide the interference? Y
- How do users configure their WiFi networks? N

Critique/Questions

Top Level comments:

- Paper is structured well.
- Scientific methodology (factors and features) was strong.
- Used very thorough experimentation with unusual set up.
- Several graphs/experiments were not well-explained
- Is Witt the only important metric?

Critique/Questions

More detailed comments:

- Provide little analysis of non-apartment performance.
- There were a number of small grammar mistakes.
- Figures/Tables and descriptive prose not always close together.
- While many good, detailed results are given, this does not inform well my intuition on wireless behavior.

Observing Home Wireless Experience through WiFi APs

Questions?

Thank you!