How Bad Are The Rogues' Impact on Enterprise 802.11 Network Performance?

Kaixin Sui, Youjian Zhao, Dan Pei *, Li Zimu
Tsinghua National Laboratory for Information Science and Technology
Tsinghua University, Beijing 100084, China
Email:{suijx11@mails., peidan@, zhaoyoujian@}tsinghua.edu.cn, lzm@cernet.edu.cn

Abstract—**Enterprise 802.11 Network (EWLAN) is an important infrastructure to the Mobile Internet, but its performance is being significantly impacted by the ever-increasing Rogue access points (RAPs). For example, in the university EWLAN we studied, the number of RAPs is more than seven times that of the enterprise APs. In this paper, we propose a generic methodology to measure RAP's carrier sense interference and hidden terminal interference, and it only uses readily available SNMP metrics, without any additional measurement hardware. Our results show that, on average, the carrier sense interference due to RAPs causes only 5% access delay increase at the MAC layer, because of careful engineering and software optimization. However, hidden terminal interference due to RAPs causes (a much more severe) up to 30% MAC layer loss rate increase on average, because no existing approach has explicitly dealt with the hidden terminal impact from rogue APs. Overall, the RAP interference would increase the IP layer delay at the WiFi hop by up to 50%.**

I. Introduction

As people rely more and more on the Mobile Internet in their daily life and work, 802.11 (WiFi) networks have become increasingly important. WiFi is often more preferred over cellular connection, especially in developing countries, because WiFi connection is typically free to the users. This is evidenced by the fast growing of various type of WiFi networks. First, universities, companies, cities and towns deploy enterprise WiFi networks and offer it to employees and residents free. Second, commercial providers deploy WiFi networks at airports, hotels, public transportations, shopping centers. Third, there are a lot of personal access points, including smart phone turned "personal hot-spot", USB dongle AP, software AP on laptops, residential broadband routers connected to Ethernet ports at work place or small business. Different 802.11 networks often occupy the same spatial area and compete for the same wireless spectrum, potentially causing interference and performance degradation.

This paper focuses on studying how *one specific* enterprise 802.11 network (Enterprise Wireless LAN, or **EWLAN** for short)'s performance is impacted by other 802.11 networks. In such a context, the access points (APs) of the EWLAN are often called enterprise APs (**EAP** for short in the rest of the paper), those of other networks (including neighbor EWLAN) are called rogue APs (**RAP** for short). While the EAP placement, channel assignment, and transmission power levels are carefully designed by engineers and optimized by the vendor software for the EWLAN, the RAPs are placed and assigned channel totally at the will of the RAP owner, without any consideration of the EWLANs. For example, there are more than 15000 unique RAPs surrounding 2000 EAPs in the China university T network we studied.

Chaotic RAP deployment potentially can cause performance degradation of EWLAN. Operators cannot clearly tell whether the performance problem is led by RAPs or not in a scalable way because of the following challenges. First, there lacks a measurement methodology for measuring RAP impact, without dedicated measurement hardware (*e.g.* sniffers used in SHAMAN [1], JIGSAW [2]). Second, the measurements results need to be actionable – we need to not only measure the overall impact, but also find the responsible RAPs, so that the EWLAN operators can take actions. In other words, we need to isolate the offending RAPs and quantify their impact on the EWLAN. Third, it is nontrivial to find hidden terminal nodes given that they should be "hidden" and to distinguish loss caused by hidden terminal and that caused by low SNR.

To address the above challenges, we propose an approach that only utilizes the SNMP data readily available at the access point controller and make the following contributions.

- To the best of our knowledge, our measurement is one of the largest-scale WiFi measurements so far, and quantifies the prevalence of Rogue 802.11 networks in a large scale for the first time in the literature.
- We propose a generic methodology to measure RAP's carrier sense interference and hidden terminal interference, and it only uses widely available SNMP metrics, without any additional measurement hardware. Further, we develop an metric for roughly quantifying RAP's service quality impact. This methodology can be used on any EWLAN.
- We observe that the studied EWLAN, despite the large number of surrounding RAPs, is not severely impacted by carrier sense interference. On average, carrier sense interference only causes 5% access delay increase at the MAC layer. This surprising result is explained by the fact that the studied EWLAN has automatic EAP power adjustment and automatic EAP channel switching to switch away from the channel most interfered by carrier sense interference, and that the EWLAN topology is carefully engineered by the operators.
- However, we found that RAPs' hidden terminal interference causes up to 30% MAC layer loss rate increase at EAP on average, much more severe than

* Dan Pei is the corresponding author.

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

carrier sense interference. This is because, even though EAP can hear the very weak signal from those hidden terminal RAPs, the current EAP software does not do anything about and simply ignore the hidden terminal RAPs, when optimize the EAP's channel and powers.

- We report a large-scale phenomenon that human traffic (and accompanying mobile devices) has significant impact on signal strength (RSSI) of the WiFi devices (RAPs in our case). Also, we show that RAPs are basically stationary.

The rest of the paper is organized as follows. Section II introduces our measurement dataset. Section III presents measurement results that show the prevalence of the Rogue APs, and the results that show enterprise network's performance is sometimes less than ideal. Section IV presents our overall measurement methodology, and details for identifying hidden terminal nodes. Section V presents the impact measurement results. Section VI reviews the related work, and finally a conclusion is presented in VII.

II. Data Set

This section introduces our measurement data set, user behavior and basic classification of EAPs. We study the EWLAN in a University T in China. T campus covers an area about 4 km^2, with $\sim$42,000 students, $\sim$11,000 faculties and staff. Our five-day (Monday 07/14/2014 to Friday 07/18/2014) data set covers the 11 wireless controllers (called **AC**) and 2,002 EAPs serve over 50,000 802.11 devices in 79 buildings. The EAPs' models are all Cisco Aironet series, which must be controlled and configured by the ACs. The AC dynamically controls the EAP transmission power by Transmit Power Control (**TPC**) [3] algorithm to avoid interference between EAPs. Because Dynamic Channel Assignment (**DCA**) [3] option is turned on, each EAP switches to a channel with the least carrier sense interference between one of the three orthogonal 2.4G channels $\{1,6,11\}$, when the detected carrier sense interference in the current channel is above a certain threshold (10% in T EWLAN).

A. SNMP Data Set

SNMP objects are readily available at the AC/AP vendor hardware (common industry practice [4]). Some of them meet the 802.11 standards, thus further facilitate our data retrieving. We set up a SNMP walker program to poll these SNMP objects at regular intervals (every 10 minutes) from 11 ACs. Each timestamped SNMP walk returns pairs of structured key and value. The information contained in the key is shown in Table I: Object name, which is one-to-one mapped from OID (a sequence of digits and dots to identify an SNMP metric uniquely), EAP MAC address, EAP radio (2.4 or 5GHz), channel number, RAP MAC address, and client MAC address.

Table I summarizes all the SNMP objects used in this paper. The objects descriptions drawn from our testbed experiments and Cisco MIB 7.6 have been confirmed by Cisco engineers. There are basically two types of objects: counter objects and sampled objects, shown in the second column of the table with the reporting interval at the ACs. For a *sampled* object, the object value is sampled and updated at regular intervals. For *counter* object, the counter keeps incrementing in the background at the EAP, but the SNMP value is updated at each AC only once regularly. Our SNMP polling interval is set to be 10 minutes, larger than the maximum of various object update intervals (180 seconds) at the AC in order to balance the frequency of information and CPU burden to the AC. The data acquisition system should have no negative influence on the AC, since our measurement shows that the total CPU usage remains below 10% during our study period.

B. User Behavior: Strong Diurnal Pattern

As shown in [5], network characteristics are heavily affected by the behavior of the users, and building types in a university can provide useful information about the aggregated user behavior. Therefore, similar to the building classification method in [5], we divide the 79 buildings in our data set into 5 categories: administrative, classroom (including the study rooms in the libraries), cafeteria, department, and dorm, before we study the user behavior. In Table II, we list the information for each building type. Manually collected information is shown in rows 1-5, and some basic measurement results are shown in rows 6-12. Information in Table II is used throughout the paper to help explain our measurement results.

Fig. 1 shows the user behavior in T's EWLAN: the total number of EAP clients and RAP clients (right Y-axis), EAP traffic volume (normalized by the maximum value) and channel utilization (left Y-axis) in all channels per SNMP poll as time and building type vary. These curves are based on or derived from objects 2, 17, 11, and 3 in Table I respectively. Two-day period out of our five-day study period is used for the clarity of presentation. The diurnal pattern in all the curves in Fig. 1 is consistent with the occupants' schedules (rows 1 and 2) in Table II.

All four curves in each figure peak around the same time: 2 peaks a day in administrative; three peaks a day in classroom, department, and cafeteria; and one peak in dorm. Furthermore, the curves of EAP client count, RAP client count, and channel utilization all peak as the number of the occupants peaks, and these curves are always proportional to each other. However, in the EAP traffic volume curves, the afternoon peaks at 16:00 in department, administrative, and classroom buildings(where people "work") always have much larger value than the other peaks of the same building types in the same day. This interesting observation can be explained by the fact that people tend to concentrate on work in the morning, and surf the Internet or even watch video in the afternoon.

Likewise, we derive from Fig. 1 the rush hour for each building type by considering the peaks of all four curves. The resulting rush hour is 16:00-17:00 for administrative, department and classroom ("Internet surfing at work" time), 12:00-13:00 for cafeteria (dining time), and 23:00-24:00 for dorm (student entertainment time). The user-perceivable performance at the rush hour is what we concerned mostly when we study the RAP's impact later in Section V.

III. Prevalence of Rogue APs

In this section, we first show the prevalence of the RAPs in T campus through measurement results. Then we show that T's EWLAN performance is much less than ideal, especially during the rush hour when the number of users and traffic

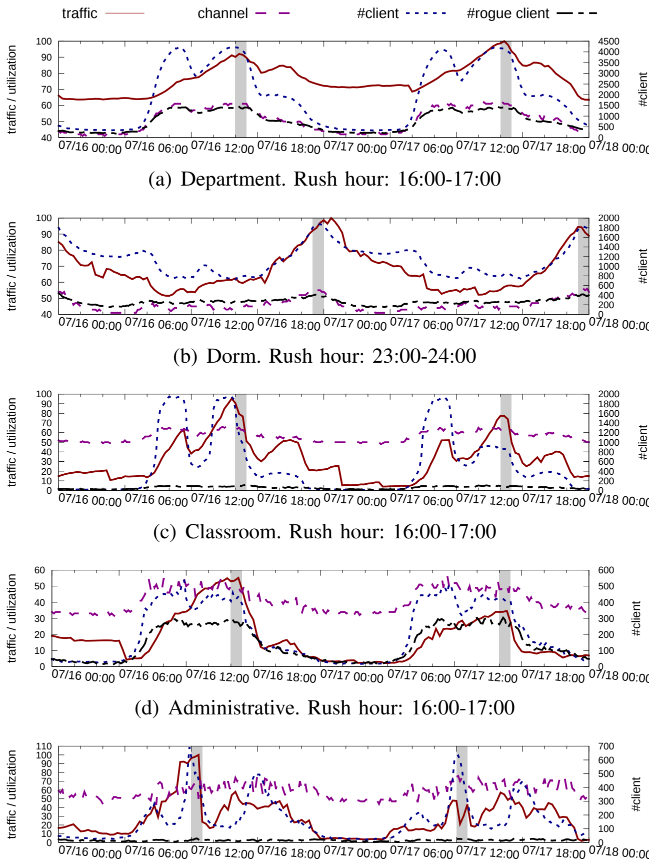

(a) Department. Rush hour: 16:00-17:00

(b) Dorm. Rush hour: 23:00-24:00

(c) Classroom. Rush hour: 16:00-17:00

(d) Administrative. Rush hour: 16:00-17:00

(e) Cafeteria. Rush hour: 12:00-13:00

Fig. 1. The total number of EAP clients and RAP clients (right Y-axis), EAP traffic volume (normalized) and channel utilization (left Y-axis) in all channel per SNMP poll as time and building type vary. The rush hours are indicated by shading.

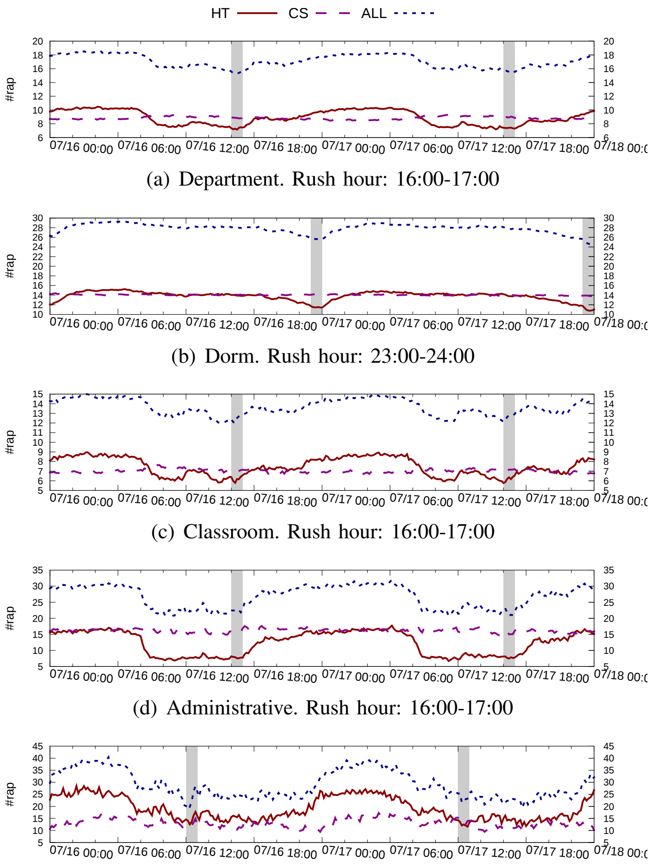

(a) Department. Rush hour: 16:00-17:00

(b) Dorm. Rush hour: 23:00-24:00

(c) Classroom. Rush hour: 16:00-17:00

(d) Administrative. Rush hour: 16:00-17:00

(e) Cafeteria. Rush hour: 12:00-13:00

Fig. 2. The average number of **Detected** RAPs, Carrier Sense RAPs, and Hidden Terminal RAPs in all channel per EAP per SNMP poll as time and building type vary. The rush hours are indicated by shading.

TABLE I. SNMP data summary: 11 ACs; 2002 EAPs; EAP reports object every 90 seconds. polling interval is 10 minutes.

ID	Object name	Type/Reporting interval	Description	Location Key
1	bsnAPIfPhyChannelNumber	sampled/∼180s	Current channel number of the AP radio.	EAP
2	bsnApIfNoOfUsers	sampled/∼90s	Number of clients associated with this radio.	EAP
3	bsnAPIfLoadChannelUtilization	sampled/∼180s	Time percentage used by all non WiFi and WiFi traffic of current channel.	EAP
4	bsnAPIfInterferenceUtilization	sampled/∼180s	Time percentage used by interference from other 802.11 networks on this channel.	EAP
5	bsnAPIfDot11TransmittedFrameCount	counter/∼180s	This counter shall increment for each successfully transmitted MSDU.	EAP
6	bsnAPIfDot11RetryCount	counter/∼180s	The number of attempts made by the EAP before transmitting the MSDU successfully.	EAP
7	bsnAPIfDot11FailedCount	counter/∼180s	This counter shall increment when an MSDU is not transmitted successfully due to the number of transmit attempts exceeding either the bsnAPIfDot11ShortRetryLimit or dot11LongRetryLimit, (7 and 4 respectively in T)	EAP
8	bsnMobileStationAPMacAddr	sampled/∼90s	802.11 MAC address of the AP to which the client is associated.	EAP, EAP client
9	bsnMobileStationAPIfSlotId	sampled/∼90s	Radio of the AP to which the client is associated.	EAP, EAP client
10	bsnMobileStationSnr	sampled/∼90s	The difference between signal strength of the client and noise.	EAP, EAP client
11	bsnMobileStationBytesSent(Received)	sampled/∼180s	Bytes sent to (received from) Mobile Station	EAP, EAP client
12	bsnMobileStationPacketsSent(Received)	sampled/∼180s	Packets sent to (received from) Mobile Station	EAP, EAP client
13	cldcClientDataRetries	sampled/∼180s	The number of attempts made by the client before transmitting the MSDU successfully.	EAP, EAP client
14	bsnRogueAPChannelNumber	sampled/∼90s	The advertised channel number of the rogue picked up from the AP. It is different from bsnAPIfPhyChannelNumber.	EAP, RAP
15	bsnRogueAPAirespaceAPRSSI	sampled/∼90s	RSSI of the rogue AP as seen by EAP.	EAP, RAP
16	bsnRogueAPTotalClients	sampled/∼90s	Total number of clients detected on this rogue.	EAP, RAP
17	bsnRogueClientAirespaceAPRSSI	sampled/∼90s	RSSI seen by AP from the rogue client.	EAP, RAP client

TABLE II. INFORMATION ABOUT 79 BUILDINGS. ROWS 1-5: MANUALLY COLLECTED INFO; ROWS 6-12 ARE MEASUREMENT RESULTS.

	Info /Building type	Administrative	Cafeteria	Classroom	Department	Dorm
1	Permanent occupants and schedule	Staff 08:00-18:00	Staff 06:00-23:00	Staff 08:00-22:00	Faculties & Students 08:00-00:00 Staff 08:00-18:00	Students 24 hour
2	Temporary occupants and schedule	Visitors 08:00-18:00	Students 07:00-09:00 11:00-13:00 17:00-20:00	Students 08:00-22:00	Visitors 08:00-00:00	Visitors 08:00-23:00
3	Ethernet ports available	Yes	No	No	Yes	Yes
4	Neighbor building types	Department	Dorms	Classroom Department	Administrative Classroom	Dorm Cafeteria
5	Major potential sources of RAPs	Residential in own & Neighbor buildings	Residential in Neighbor buildings	Residential in Neighbor buildings 3G gateway	Residential USB 3G gateway	Residential Neighbor Enterprise
6	**Rush hour**	16:00-17:00	12:00-13:00	16:00-17:00	16:00-17:00	23:00-24:00
7	**#Building [79]**	20	6	8	35	10
8	**#EAPs [2002]**	103	25	346	913	615
9	**#EAP Clients in 5 days [51269]**	16376	13514	16121	33760	10280
10	**#RAPs in 5 days [15110]**	2699	2144	2268	7788	4379
11	**#Concurrent RAPs [5374]**	673	428	555	2486	1838
12	**#RAP Clients in 5 days [44996]**	7349	1677	5269	27270	10114

volume both peak, implying that the large number of RAPs might be the suspect.

Note that the number of RAPs in 2.4 GHz (15110 in 5 days) is much more than 5 GHz (292 in 5 days). It is because the technology for 5 GHz is newer than 2.4 GHz, and wireless card work with 5 GHz is more expensive and less common. Hence, our study **focus on 2.4 GHz** in the rest of the paper.

A. RAPs' Diurnal Patterns Are Opposite to Human Traffic

A huge number of RAPs exist in T campus. According to rows 8, 10, 11 in Table II, for just around 2,000 EAPs, there are more than 15,000 unique RAPs, and the median number of unique concurrent RAPs per SNMP polls also reaches 5374. The RAP (concurrent RAP) to EAP ratio is more than 7:1 (2:1). On the other hand, the number of EAP clients (about 51000, row 9) is similar to the number of RAP clients (about 45000, row 12 in Table II). This shows that almost every client of T campus EAPs has access to some RAPs. And the average number of RAP clients per RAP is much smaller than the number clients per EAP because a RAP is most of time protected by private passwords known to only a small number of persons, while EAPs are public infrastructure open to every campus member with centralized authentication servers.

Based on the data of object 13 in Table I, we find that the RAPs are mainly allocated at channel 1, 6, 11, while channel 11 is used less than channels 1 and 6. The median number of concurrent RAP on the channel 1, 6, 11 is 1673, 1665, 1258 respectively, and other channels altogether is 840.

Fig. 2 shows the average RAP count in all channel per EAP per SNMP poll as time and building type vary. The meaning of Carrier Sense RAPs and Hidden Terminal RAPs will be discussed later in the paper. The difference between different buildings can be explained by the information in Table II, such as occupants, Ethernet availability, building count, and surrounding building type.

There is one interesting observation about all these figures: the number of RAPs is usually lower when there are more human traffic (daytime except dorm) than when there are less human traffic (night time except dorm). This is counterintuitive at the first glance, because we suppose more RAPs (e.g. mobile RAPs) are turned on during the day. However, according to [6],"multipath fading occurs when the reflected signal paths refract off people, furniture, windows, and scatter the transmitted signal, and sometimes could produce an additional loss of signal power on the order of 20 dB or more". Therefore, more human traffic (bodies and mobile devices) causes more multipath fading. The RAPs with very weak signal "appear" when the multipath fading (less human traffic) is weak and "disappear" when the multipath fading (more human traffic) is strong. As a result, the number of concurrent RAPs (row 11) is always nearly a third of total detected RAP (row 10) in Table II. We will see consistent evidence supporting this explanation later in the paper. The large number of RAPs and large RAP to EAP ratio pose a significant concern about RAPs' potentially large impact on T campus EWLAN.

B. RAP Classification and Mobility

Different RAPs might behave differently. In order to better understand RAPs impact on EAPs, we classify RAP into 6 types in Table III first using RAP SSID keywords (highlighted using bold fonts), and if SSID cannot tell, we use the OUI (Organizationally Unique Identifier) [7] in RAP MAC address to map the RAP to its manufacturer. For example, *3g* and *pocket* in SSID indicate a type of mobile 3G/WiFi gateway, while an OUI of *tp-link* indicates the popular *residential AP* manufacturer TP-LINK. Based on this table, we can see that the potentially mobile RAPs (3G gateway, smart phone, USB WiFi dongle, and software AP on laptops) only contribute to 2451 RAPs (less than 17%) out of about 15000 RAPs. The most dominant types are residential AP (installed on campus) which contributes to more than 70% of all RAPs, and APs from Neighbor EWLAN (e.g. installed by telecom companies) which contributes to about 13% of all RAPs.

For those most frequent SSID and OUI, we also use the number in the bracket to show its total RAP count, which shows that the popular TP-LINK residential APs alone contributes to more than 30% (5387) of all RAPs. The same manufacturer typically has the same default configuration parameters (e.g. default channel), which can affect its interaction with the EAPs.

TABLE III. RAP CLASSIFICATION BASED ON SSID FIRST AND THEN OUI.

Type	Count.	Mobile?	SSID or OUI Keyword
3G Gateway	456 [3.02%]	Yes	**[ssid] 3g, hame, incar, mobile, pocket, portable, u+net**
Smart Phone	857 [5.67%]	Yes	[oui] meizu, mobile, nokia, oppo, samsung, xiaomi **[ssid] android, coolpad, htc, ipad, iphone, mylgnet, ruitel**
USB WiFi Dongle	532 [3.52%]	Yes	**[ssid] 360wifi(504), baidu(28)**
Software AP on Laptops	606 [4.01%]	Yes	[oui] apple, lenovo, toshiba **[ssid] asus, dell, hp, lenovo, thinkpad, vaio [software] connectify, dubawifi, liebaofree**
Neighbor Enterprise WiFi	1981 [13.11%]	No	[oui] aruba, cisco(383), h3c(829), juniper **[ssid] cernet, chinacomm, chinanet(623), chinaunicom(308), cloudwifi, cmcc(208), ctt-T (259), cu_campus(69), cwic, gehua(51), ivi, nuctech(183), videophone**
Residential AP	10678 [70.67%]	No	[oui] d-link(482), hiwifi(103), huawei(1144), netgear(360), tenda(603), tp-link(5387), zte(312) **[ssid] b-link, buffalo, dd-wrt, d-link, dorm, fast, feixun, hiwifi, huawei, iptime, jcg, lab, mercury, netcore, netgear, openwrt, phicomm, print, room, tenda, toto-link, tp-link, volans, zte ...**

TABLE IV. RAP CLASSIFICATION AND DAILY MOVEMENT

Type	Daily Movements and Percentage			
	0	<= 5	<= 10	MAX
3G gateway	0.8469	1	1	4.8
Smart phone	0.8789	0.9983	1	11.6
Software on laptops	0.8731	0.9907	0.9963	22.8
USB dongle	0.8971	0.9980	1	6.4
Neighbor enterprise	0.9589	0.9895	0.9948	17.6
Residential	0.9374	0.9965	0.9997	20.4
All	0.9226	0.9953	0.9984	22.8

We also observe that there are a few neighbor EWLAN (identified mostly using SSID) with hundreds of APs of their own, which (e.g. telecom company) is in some sense more difficult to coordinate with than the owners of other RAPs (e.g. a graduate student). This further highlights the challenges of managing EWLAN.

Furthermore, we study the mobility of the RAPs in Table IV. We first define a RAP's *coordinate* as the set of buildings of EAP that can detect it with a relatively strong signal. When its coordinate changes, we consider the RAP moves. For theoretically mobile RAPs (3G gateway, smart phone, USB dongle, and Software), their power is typically less than those stationary ones. Thus, we generate the building set of RSSI (object 15 in Table I) $>$ -85 detected by EAP belongs to that building for mobile RAPs, and use RSSI $>$ -70 for theoretically stationary RAPs.

Table IV shows the percentage of averaged movement times of the RAPs per day. Only under 8% of the RAPs are actually mobile. One MAC address can be shared by multiple APs and signal extender in the neighbor EWLAN, and thus signal strength fluctuation can cause an "AP" (identified by the MAC address) appear to move from one coordinate to another, but indeed this is not a real movement. So overall Table IV is an overestimate of the RAP mobility, and even so, the RAP mobility is not much in our data set.

C. EAP Performance Is Less than Ideal

The channel utilization of the EAP is heavy. According to rows 9, 12 in Table II, there are more than 50,000 EAP clients and 40,000 rogue clients in T campus. The great amount of dense 802.11 users fill all the available channels with their traffic. As shown in Fig. 1, the channel utilization reaches its peak at the rush hour everyday. Fig. 3(a) shows the CDF of EAP channel utilization (object 3 in Table I) using all the sample points with one or more clients associated at the rush hour of our five day study period. It shows that the median busy time percentage of the EAP channel is around 60%, while the busiest reaches 80% or even higher in nearly 20% cases. The channel is much more vulnerable and sensitive to additional interference under the heavy traffic.

According to [3], "interference is any 802.11 traffic that is not part of your wireless LAN, including rogue access points and neighboring wireless networks. Lightweight access points can simultaneously scan all valid 802.11a/b/g channels looking for sources of interference. The access points go "off-channel" for a period not greater than 60 ms to monitor these channels for noise and interference. Packets collected during this time are analyzed to detect rogue access points, rogue clients, ad-hoc clients, and interfering access points." The interference values measured by EAP are reported to the controller at regular intervals and stored in the controller SNMP MIB.

Fig. 3(b) shows the interference utilization (object 4 in Table I) as a CDF at the rush hour. It shows that on average the interferer of other 802.11 network occupies only around 3% of the channel usage time, and lower than 10% in most cases. The low interference value is explained by the fact that the switch *bsnAPIfDot11MacParamsConfigType* is set to be automatic at all the EAPs in T campus, which means the "controller may dynamically rearrange EAPs' channel assignments to increase system performance in the presence of the interference by Dynamic Channel Assignment (DCA) algorithm if the amount of 802.11 interference exceeds a predefined configurable threshold (the default is 10 percent)" according to [3].

The loss rate of the EAP IP layer and MAC layer shown in Fig. 4 is derived from counter objects 5, 6, 7 in Table I. The calculation method of loss rate will be explained later in Section V-B. Fig. 4(a) shows that the EAP IP layer loss rate is more than 1% on average, and Fig. 4(b) shows that the MAC layer loss rate is more than 40% on average. We consider the 50% to be the **high MAC loss rate threshold** referring to [8]. In Fig. 4(b), the MAC layer loss rate reaches a considerably high level over 50% in nearly 20% cases at the rush hour.

The overall EWLAN performance is sometimes much less than ideal, especially during the rush hour when the number of users and traffic volume both peak. Fig. 3 and Fig. 4 indicate that the EAPs in T campus are suffering the heavy channel utilization, continual interference, and high packet loss. However, the EAP can somehow avoid the interference by DCA and TPC. Unfortunately, the EAP takes no extra measures to avert the packet loss, therefore, the packet loss is significantly severe than any other impacts on EAP. While the EAPs are

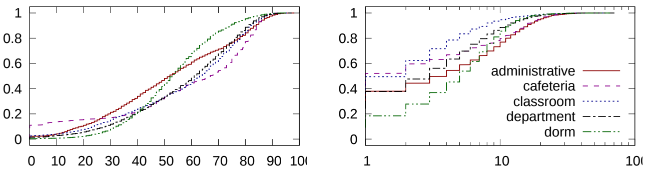

(a) Channel utilization (b) Interference utilization

Fig. 3. The CDF of **Channel Utilization** and **Interference Utilization** of EAPs using all the sample points with one or more clients associated at the rush hour of our five day study period

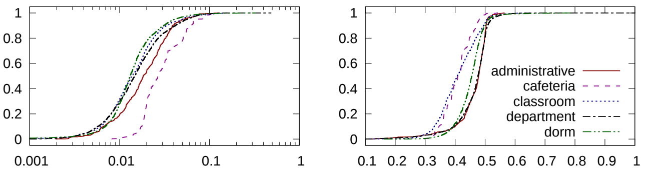

(a) EAP IP loss rate (b) EAP MAC loss rate

Fig. 4. The CDF of **MAC layer loss rate** and **IP layer loss rate** of EAPs using all the sample points with one or more clients associated at the rush hour of our five day study period

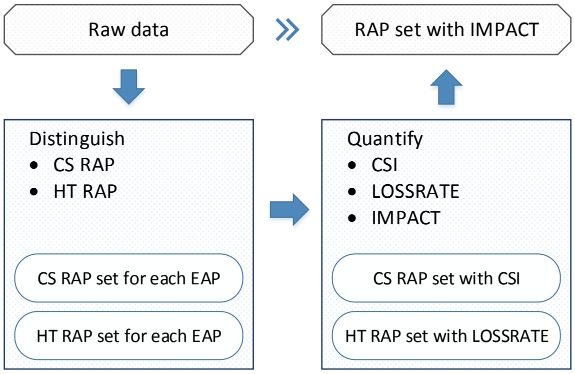

Fig. 5. Overall measurement methodology.

carefully designed with reasonable power, appropriate position, staggered channel, and centralized controlled, the chaotic RAP deployment without any consideration of the EAPs is the obvious suspect of the EWLAN performance degradation. The impact of the RAP should be quantified and alleviated.

IV. Distinguishing CS vs HT

The overall measurement methodology (Section IV, V) of the paper is organized as follows. First, we isolate the offending RAPs and propose an approach to identify hidden terminal nodes from all RAPs detected by EAP in Section IV. Secondly, we measure the impact of Carrier Sense RAP (**CS RAP** for short) and Hidden Terminal RAP (**HT RAP** for short) in Section V. Processing procedures from the raw data to the actionable results are summarized in Fig. 5.

A. Validating The Value of Carrier Sense Threshold (CST)

RAPs detected by the EAP generally fall into two broad categories, CS RAP and HT RAP. The CS RAP mainly competes for the same wireless spectrum with the EAP, which results in the increasing EAP access delay. The HT RAP leads to severe packet loss of EAP. The impact of CS RAP and HT RAP needs to be measured respectively due to the totally different influences on EAP.

We isolate the offending RAPs based on object 15 in Table I, and use Carrier Sense Threshold *CST* as the boundary to distinguish CS RAP and HT RAP within the availability of the data.

According to [9], “the *CST* indicates the minimumpower/energy that an RF receiver must receive to detect the transmission of a wireless signal.” The RAP with RSSI below *CST* seen by EAP is considered to be the HT RAP given that they should be “hidden”, or else belongs to the CS RAP. We choose -85 dBm as the value of *CST*, which means the RAP with RSSI below -85 dBm is HT RAP, while the RAP with RSSI greater than or equal to -85 dBm is CS RAP. However, even though EAP can hear the very weak signal of those HT RAPs, the current EAP software does not do anything about and simply ignore the HT RAPs, when optimizing the EAP’s channel and powers.

Why we choose -85 dBm as the value of *CST*? According to [9], “most wireless card manufacturers conservatively set this threshold to a low value -85 dBm.” For further validation, we conduct an experiment as shown in Fig. 6. EAP and RAP are in the same clean channel with no background traffic. During the experiment, RAP PC sends UDP packets to RAP client via RAP at the rate of 4 Mbps. Meanwhile, EAP PC sends UDP packets to EAP client via EAP at the rate of 400 Kbps. Similar to the method of validating hidden terminal impact in [10], we move the RAP further away from the EAP. While the EAP gradually gets out of the RAP communication range and the EAP client is still in, the RAP completes the transition from CS RAP to HT RAP of the EAP.

Fig. 7 shows when RAP RSSI value (object 15 in Table I) crosses from above -80 dBm to below -90 dBm, the variation of channel utilization, interference utilization, and retry count (object 3, 4, 6 in Table I) of EAP significantly change. Note that we turn off the DCA of the EAP during the experiment. When RAP RSSI $>=$ *CST* (the white zone in Fig. 7), the RAP is more like a CS RAP. Carrier sense impact is mainly represented on EAP such as low retry count, high interference utilization caused by RAP, and high channel utilization used by both EAP and RAP. When RSSI $<$ *CST* (the shadow zone in Fig. 7), the RAP is more like a HT RAP. Hidden terminal impact is mainly represented on EAP such as high retry count, low interference utilization, and partly channel utilization used by EAP. As the RSSI of RAP decreases, the channel utilization, interference utilization, and retry count of EAP show distinctive differences between two zones separated by *CST* $=$ -85 dBm.

The channel utilization and interference utilization suddenly fall down when RAP RSSI $<$ *CST* (-85 dBm), because the packets from RAP with RSSI below *CST* may not be received by the EAP wireless card. We call these packets hidden packets. Without detecting the hidden packets in the channel based on CSMA/CA mechanism, the EAP considers the channel to be idle and sends its packet to EAP client. There is a high probability that the packets from EAP loss

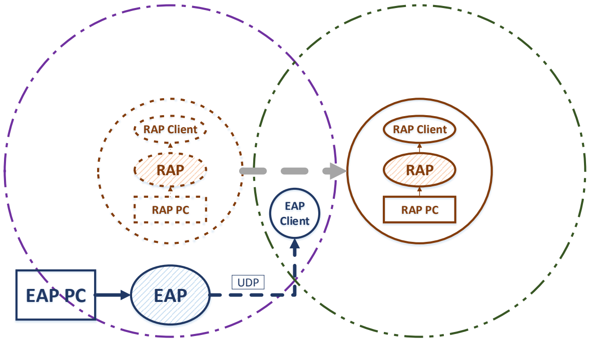

Fig. 6. *CST* experiment method. The RAP moves further away from the EAP. The EAP gradually gets out of the RAP communication range and the EAP client is still in. The RAP changes from CS RAP to HT RAP during the experiment.

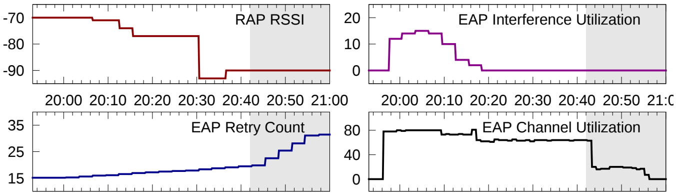

Fig. 7. *CST* experiment results. The white zone is when the RAP is more like a CS RAP, and the shadow zone is when the RAP is more like a HT RAP. As the **RSSI** of RAP decreases, **Channel Utilization**, **Interference Utilization** and **Retry Count** of EAP show distinctive differences between two zones separated by *CST* = -85 dBm. DCA is not enable during the experiment.

on a collision with the hidden packets from RAP at the EAP client within the communication range of both EAP and RAP. As such, the retry count of EAP increases rapidly due to severe losses when RAP RSSI $<$ *CST* (-85 dBm).

The experiment validates -85 dBm as the value of *CST*. For each RAP, we use ***CST*** to distinguish whether it is a CS RAP (RSSI $>=$ *CST*) or a HT RAP (RSSI $<$ *CST*).

B. More HT RAPs Than CS RAPs

Fig. 2 shows the average number of detected RAPs, CS RAPs and HT RAPs per EAP per SNMP poll as time and building type vary.

The difference between different buildings can be explained by the information in Table II such as occupants, Ethernet availability, building count, and surrounding building type. The HT RAP count is less stable than the CS RAP count, because the HT RAP with low RSSI is easily faded by multipath fading. The HT RAP with very weak signal “appear” when the multipath fading is weak and “disappear” when the multipath fading is strong. So the number of HT RAPs is usually lower when there are more human traffic than when there are less human traffic, while the number of CS RAPs holds comparatively firm along with time.

C. #RAPs in Overlapping Channels with An EAP

There are only 3 orthogonal channels $\{1, 6, 11\}$, and each channel overlaps with neighbor channels. Due to the EAP is only interfered by the RAPs in overlapping channels with it, we only concern about the RAPs **in overlapping channels** with each EAP when measuring the impact of RAPs in the rest of the paper.

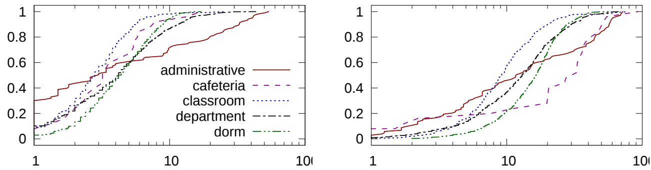

(a) CS RAP at an EAP

(b) HT RAPs at an EAP

Fig. 8. The CDF of CS RAPs and HT RAPs count **in overlapping channels** with an EAP per day for each type of buildings during five days

Fig. 8 is derived form object 1, 14, 15, 16 in Table I. Fig. 8(a) shows the number of CS RAPs in overlapping channels with an EAP per day for each type of buildings as a CDF, when one or more RAP clients associated. Fig. 8(b) shows the number of HT RAPs in overlapping channels with an EAP per day for each type of buildings as a CDF. If a CS RAP has no client associated, it means that the RAP have no traffic thus no carrier sense influence on the EAP. However, it does not validate to the HT RAP, because their traffic is too weak for EAP to unpack and record in MIB.

The EAPs in classroom have both small amount of CS RAPs and HT RAPs, because the classroom building do not have the Ethernet port (row 3 in Table II), and its neighbor buildings of other types are far away from the teaching areas in T campus. The EAPs in cafeteria have large amount of HT RAPs but small amount of CS RAPs, because RAPs of cafeteria building are mainly from the neighbor buildings of dorms. The RAPs RSSI seen by EAP in cafeteria are mostly low due to the fading channel such as building walls, space distance.

On average, around 5 CS RAPs per day in the overlapping channels with the EAP, might compete channel with the EAP and lead to the access delay. And around 15 HT RAPs, might be potential hidden terminal nodes of the EAP and lead to packet loss. The number of HT RAPs in overlapping channels with the EAP is sometimes triple the number of CS RAPs. The large RAP in overlapping channels to EAP ratio reinforces the significant concern about RAPs’ potentially large impact on T campus EWLAN, especially the HT RAPs.

V. Rogue APs’ Impact

The EWLAN performance is less than ideal at the rush hour as mentioned in Section III-C. Although the RAPs count is more at the off peak hours, the RAPs impact on EAP is less, because the actual traffic of RAP and active users of EAP are smaller than at the rush hour. Therefore, the user-perceivable performance at the rush hour is what we mostly concerned. We quantify the impact of CS RAP and HT RAP respectively by measuring *CSI*, *LOSSRATE*, and finally conclude the overall *IMPACT* at the rush hour in this section.

A. Carrier Sense RAPs’ Impact: CSI

According to 802.11 CSMA/CA mechanism, EAP cannot send traffic but have to wait when hearing the channel is occupied. A CS RAP mainly competes for the same wireless spectrum with the EAP, which results in the increasing EAP access delay. We measure the impact of CS RAP by the metric *CSI* to roughly estimate the additional back-off time of EAP caused by CS RAP.

CSI = IU / (CU - IU)

CU is short for channel utilization based on object 3 in Table I, and *IU* is short for interference utilization based on object 4 in Table I.

When the channel utilization is high, the *CSI* is approximate the increased delay caused by the carrier sense interference. The larger the value for *CSI*, the more severe the EAP impacted by CS RAPs. Note that with no access to the reliable non-WiFi data to filter the non-WiFi utilization from denominator, the *CSI* might be underestimated.

Fig. 9(a) shows the CDF of CSI when one or more EAP clients associated during the channel heavy load time with *CU* $>$ 60. On average carrier sense interference only cause 5% access delay increase at the MAC layer, and below 10% in most cases. As mentioned in Section II, the EAP deployment in T campus is carefully designed and optimized by DCA and TPC, so the carrier sense interference somehow can be avoided. Despite the large number of surrounding RAPs, the impact of CS RAP is not severely.

B. Hidden Terminal RAPs' Impact: LOSSRATE

Unlike CS RAP, a HT RAP does not postpone the EAP sending time, because EAP can barely hear the HT RAP traffic. However, more seriously, without detecting the hidden packets in the channel based on CSMA/CA mechanism, the EAP considers the channel to be idle and sends its packet to EAP client. There is a high probability that the packets from EAP loss on a collision with the hidden packets from RAP at the EAP client within the communication range of both EAP and RAP. The HT RAP leads to severe packet loss of EAP.

In T EWLAN, *RTSThreshold* is set to be 0 for all EAPs, which means the RTS/CTS handshake is turning on for all frames. The *LongRetryLimit* and *ShortRetryLimit* are set to be 7, 4 respectively in T EWLAN, which means all the failure frame with size larger than or equal to 0 (*RTSThreshold*) have already been transmitted 4 (*LongRetryLimit*) times.

Based on objects 5, 6, 7, 12, 13 in Table I, the *SUCCESS* counter adds up if the MSDU transmits successfully, the *RETRY* counter will increase by each retransmission, and the *FAILED* counter will increase only if the MSDU is finally failed after several transmit attempts. The packet *LOSSRATE* from EAP to EAP client of MAC and IP layer can be handled conveniently by equations as follows.

IP LOSSRATE = F / (F + S)

MAC LOSSRATE = (L*F + R) / (L*F + R + S)

L is short for *LongRetryLimit*, *F* for *FAILED* counter, *R* for *RETRY* counter, and *S* for *SUCCESS* counter.

There are several reasons for packet loss: fading channel, non-WiFi interference and hidden terminal RAP. However, packet loss due to fading channel and non-WiFi interference represents the obvious similarity in low SNR of the associated clients. Therefore, packet loss from EAP to high SNR EAP client is mainly caused by HT RAP.

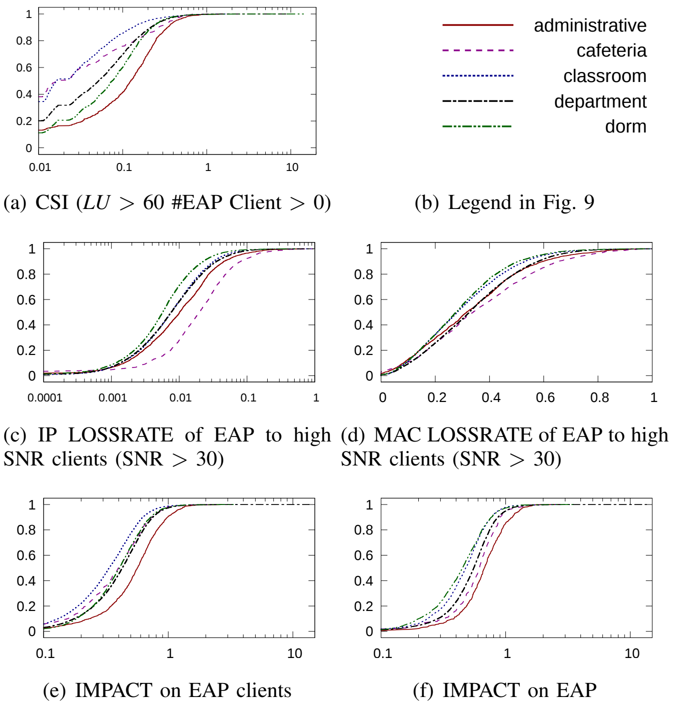

(a) CSI (LU $>$ 60 #EAP Client $>$ 0)

(b) Legend in Fig. 9

(c) IP LOSSRATE of EAP to high SNR clients (SNR $>$ 30)

(d) MAC LOSSRATE of EAP to high SNR clients (SNR $>$ 30)

(e) IMPACT on EAP clients

(f) IMPACT on EAP

Fig. 9. Impact measurement results. The CDF of ***CSI***, ***IP LOSSRATE***, ***MAC LOSSRATE***, and overall ***IMPACT***, at the rush hour during five days.

In order to get rid of the loss caused by low SNR, we measure the impact of HT RAP by ***LOSSRATE*** of pakets from EAP to **high SNR EAP client**. We select 30 dBm as the high SNR threshold, because the client with SNR above 30 dBm always associated successfully and very fast according to [11]. Fig. 9(c) shows the CDF of high SNR clients' IP loss rate. Fig. 9(d) shows the CDF of high SNR clients' MAC loss rate. The IP loss rate of high SNR clients is nearly 1% on average. The MAC loss rate is around 30% on average. Critically, the T EWLAN users suffer an extremely high MAC loss rate more than 50% in almost 20% cases.

The observation that the impact of HT RAPs is much more severe than CS RAPs is consistent with their quantitative relations in Fig. 2 and Fig. 8. This is because the current EAP software does not do anything about and simply ignore the HT RAPs as mentioned in Section IV.

C. Overall Impact

Although *CSI* and *LOSSRATE* can measure the RAP impact from two different points of view, there is still lack of a directly quantified value to measure the overall impact. We define a new metric *IMPACT* to represent the overall impact derived from *CSI* and *MAC LOSSRATE.*

IMPACT = (1 + CSI) * (1 + MAC LOSSRATE) - 1

Overall, Fig. 9(e) and Fig. 9(f) show that a fairly high increasing delay (over 50%) of the WiFi hop is caused by RAP. In extreme cases, the *IMPACT* even exceeds 10, it seems that the EAP wireless network performance is terrible, and the

EAP client could not access the Internet through it due to the abnormal high delay caused by RAPs. Because the queueing delay is not taken into consideration, the *IMPACT* is always an underestimate of delay, thus it could be even worse in the practical environment.

In Fig. 9(f), buildings with greater *IMPACT* always have more RAPs, RAP clients and traffic volume per EAP. The impact of CS RAP and HT RAP are naturally greater than the other buildings of the same type. For example, the F building of department always suffers the high impact above 2 due to RAPs. Because most of the occupants in F building are students from the computer science department. They have more personal APs and network devices than the other students. There are even several wireless experimental environments in F building. Buildings with impact greater than 2 should be optimized firstly.

Based on the measurement results of the overall impact and the responsible RAPs, lots of approaches could be adopted to alleviate the impact of the RAPs and optimize the EWLAN performance. The operators in T EWLAN should pay more attention to the hidden terminal problem, due to the observation that hidden terminal interference is more severe than the carrier sense interference.

VI. Related Work

[12], [13], [14], [15] use active measurements to gain the interference map or the conflict graph. [1], [2] monitor the wireless environment with additional measurement hardwares. [8] measures the interference of neighbor APs by capturing packets based on OpenWrt [16]. We calculate the interference of Carrier Sense RAPs and the Hidden Terminal RAPs just by using the SNMP data.

Some previous works also use the SNMP data. [17] collects SNMP and tcpdump [18] data of 12 APs and 74 users during 12 weeks in Stanford. [5] collects syslog, SNMP, and tcpdump data of 476 APs and over 1700 users in Dartmouth. We collect SNMP data of 2002 APs and over 50,000 users in T campus as of July 1, 2014.

VII. Conclusion

This paper presents the first large-scale measurement study on the rogue access points' impact on the Enterprise WiFi Network (EWLAN) using the readily available SNMP data without any additional measurement hardware. We show that, for the studied EWLAN, the rogue APs outnumbers EWLAN APs by about seven times. We observe that carrier sense interference due to RAPs are not severe, because of the careful topology engineering and EAP automatic channel and power adjustment. However, RAPs' hidden terminal interference causes (a much more severe) up to 50% MAC layer loss rate increase because EAPs currently do nothing about them.

We believe that our work is an important step towards in-depth understanding and mitigating service quality problems in operational EWLAN, and the dynamic interaction between different 802.11 LANs. Future directions include a more thorough study on the mitigation strategies, and measuring the Non-WiFi interferer's impact and mitigation solutions.

Acknowledgement

This work has been supported in part by the National High Technology Research and Development Program of China (863 Program) under Grant No.2013AA013302, the National Key Basic Research Program of China (973 program) under Grant No.2013CB329105, and the National Natural Science Foundation of China (NSFC) under Grant No.61233007, No.61472210, No.61472214.

We would like to thank the reviewers for their valuable comments. We sincerely appreciate Fenghua Li, Qingguo Zhang, and Xudong Li for their useful experience on analyzing the Cisco SNMP MIB. Also, we express thanks to Dapeng Liu for his assistance on verifying the Cisco SNMP OID. At last, we thank Juexing Liao, Yaodong Zhang, Changhua Pei, Xiaohui Nie, and Guo Chen for their proofreading of this paper.

References

[1] Yu-Chung Cheng, Mikhail Afanasyev, Patrick Verkaik, Peter Benko, Jennifer Chiang, Alex C Snoeren, Stefan Savage, and Geoff M Voelker. Shaman automatic 802.11 wireless diagnosis system. 2010.

[2] Yu-Chung Cheng, John Bellardo, Péter Benkö, Alex C Snoeren, Geoffrey M Voelker, and Stefan Savage. *Jigsaw: solving the puzzle of enterprise 802.11 analysis*, volume 36. ACM, 2006.

[3] Cisco wireless lan controller configuration guide. http://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/71113-rrm-new.html.

[4] Snmp object navigator. http://tools.cisco.com/Support/SNMP/do/BrowseOID.do.

[5] David Kotz and Kobby Essien. Analysis of a campus-wide wireless network. *Wireless Networks*, 11(1-2):115–133, 2005.

[6] Mustafa Ergen. Ieee 802.11 tutorial. *University of California Berkeley*, 70, 2002.

[7] Snmp object navigator. http://standards.ieee.org/develop/regauth/oui/oui.txt.

[8] Ashish Patro, Srinivas Govindan, and Suman Banerjee. Observing home wireless experience through wifi aps. In *Proceedings of the 19th annual international conference on Mobile computing & networking*, pages 339–350. ACM, 2013.

[9] Nabeel Ahmed. *Interference Management in Dense 802.11 Networks*. PhD thesis.

[10] Justin Manweiler, Peter Franklin, and Romit Roy Choudhury. Rxip: Monitoring the health of home wireless networks. In *INFOCOM, 2012 Proceedings IEEE*.

[11] How to: Define minimum snr values for signal coverage. http://www.wireless-nets.com/resources/tutorials/define_SNR_values.html.

[12] Nabeel Ahmed and Srinivasan Keshav. Smarta: a self-managing architecture for thin access points. In *Proceedings of the 2006 ACM CoNEXT conference*.

[13] Nabeel Ahmed, Usman Ismail, Srinivasan Keshav, and Konstantina Papagiannaki. Online estimation of rf interference. In *Proceedings of the 2008 ACM CoNEXT Conference*, page 4. ACM, 2008.

[14] Dragoş Niculescu. Interference map for 802.11 networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 2007.

[15] Jitendra Padhye, Sharad Agarwal, Venkata N Padmanabhan, Lili Qiu, Ananth Rao, and Brian Zill. Estimation of link interference in static multi-hop wireless networks. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, pages 28–28. USENIX Association, 2005.

[16] Openwrt. https://openwrt.org/.

[17] Diane Tang and Mary Baker. Analysis of a local-area wireless network. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 1–10. ACM, 2000.

[18] Tcpdump. http://www.tcpdump.org/.