Realtime Detection of Degradation in WiFi Network's Goodput Due to Probe Traffic

Dheryta Jaisinghani
IIIT-Delhi
New Delhi, India
dherytaj@iiitd.ac.in

Vinayak Naik
IIIT-Delhi
New Delhi, India
naik@iiitd.ac.in

Sanjit K. Kaul
IIIT-Delhi
New Delhi, India
skkaul@iiitd.ac.in

Sumit Roy
University of Washington
Seattle, WA
sroy@u.washington.edu

***Abstract*—IEEE 802.11 WLAN (Wireless or WiFi LAN) clients discover neighboring APs (Access Points) by active or passive scanning. Such an active scan of WLAN injects probe frames in the network. Network conditions like packet losses, roaming, etc. result in increased active scanning and hence, an excessive increase of the probe traffic. Of the several causes inherent to WLANs like interference, we find an excessive probe traffic also has a potential of hampering goodput of a WiFi network. We confirm this behavior in a controlled home environment as well as in an uncontrolled enterprise environment. Our analysis of 36 hours of wireless traffic collected over a period of 5 months with approximately 45 million wireless frames reveals that goodput of a WLAN drops exponentially with increase in the probe traffic. Therefore, realtime detection of increase in probe traffic and knowledge of a threshold for acceptable probing is crucial for WLAN's performance. In this paper, we formulate a metric to measure the increase in probe traffic in realtime and evaluate its functioning empirically. The metric not only reflects increase in probe traffic correctly, it is even simple enough to allow its realtime measurement.**

I. Introduction

WiFi clients perform active scanning to look for neighborhood APs. During the active scanning, they inject probe traffic in to the network that consists of probe requests, which further result in probe responses from AP(s). It is an essential network management traffic. Its purposes are (1) to let a new client to associate in a 802.11 network, (2) achieve roaming in a 802.11 network, and (3) to allow stationary clients to switch APs [1]. Raghavendra et al. in [2] found that even stationary clients at times wrongly initiate handoffs with an aim to find better AP and improve throughput. However, such handoffs only inject unwanted probe traffic and hardly improve network throughput.

We find that excessive probe traffic disrupts network's performance in terms of lowered goodput. It is known that frames at low rate result in lower throughput for high data rate clients. However, unlike other low rate data frames, probe requests are not only transmitted at 1 Mbps, but they are also mostly sent on each channel in the spectrum after allowed timeout. Further, all APs present in the network send probe response frames for every probe request they listen. Probe responses are also sent at low data rates, typically 1 Mbps. A probing cycle is considered complete when probing device responds with ACK (acknowledgement) frame to the AP, else AP retries the probe response until it gets an ACK or maximum retries timer expires.

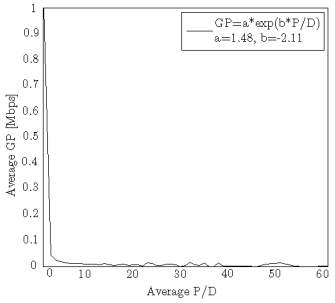

Fig. 1: Exponential loss in MAC goodput as a function of increase in probe traffic as compared to data traffic. Increase in probe traffic is measured by metric P/D.

This traffic degrades network's goodput exponentially, as shown in Figure 1. The plot shows drop in MAC (Medium Access Control) GP (goodput) with relative increase in probe traffic, which is measured by the metric P/D. This metric is the ratio of number of P (probe) frames to number of fresh D (data) frames *i.e.* data frames which were successfully transmitted in first attempt. The derived relationship is based on 36 hours of traffic captured over 5 months in an uncontrolled enterprise environment.

The solutions found in the literature is to disable probing in the network – (1) by not allowing devices to actively probe, (2) by implementing modified scanning algorithms in devices [3]–[5], or (3) by not allowing APs to respond to probe requests. After disabling probing, the other option is passive scanning, which is time consuming and hence not preferred [6]–[9]. However, probing may not be a reason behind network's degraded goodput. Interference [10]–[12], hidden terminals [13], [14], exposed terminals [15], rate adaptation [16] etc. are amongst few causes of a network's goodput drop. Large scale network deployments have large number of WiFi devices

978-3-9018-8273-9/15/ ©2015 IFIP

[17] - clients and APs which exaggerate these causes.

Therefore, we suggest to detect in realtime if probe traffic is reducing network performance in terms of lowered GP before taking action against probing. To eliminate other causes of goodput drop, first we study increased probe traffic and its effects in a controlled network environment of home, followed by uncontrolled network environment of an enterprise building. We introduce the metric P/D to monitor the realtime increase in probe traffic and empirically evaluate its functioning. Our metric can be used to find out whether it is necessary to disable probing, by any of three aforementioned methods, or not. We also present an empirically derived mathematical relation to demonstrate the exponential drop in MAC GP due to excessive probing.

Organization of the paper: We begin with formulation of the metric P/D in Section II. In the section, we present a need of considering only probe frames to study the impact on network's GP, then we show the effect of probe frames on data frames, and discuss the metric P/D. In Section III, we present network deployment, experiment setup, data collection, results and analysis of both controlled and uncontrolled experiments. We also present evaluation of metric P/D in that section. We conclude the paper in Section IV.

II. Formulation of the Metric P/D

A. Why consider probe frames?

IEEE 802.11 [18] standard specifies more than 35 types of frames under 3 categories of Management, Control, and Data. These frames are used for establishing, managing, and controlling data transfers using WiFi connections. Data frames carry information to be transmitted from source to destination. Control frames help in efficient transfer of data frames. Management frames help in establishing and maintaining WiFi connections. We want to study which of the control/management frames may have an adverse effect on network GP. We measure network GP as MAC GP, which is defined as the amount of data acknowledged per unit time. For a given Time Slot T_s, GP is calculated as per Equation 1,

$$GP = \frac{\sum_{i=1}^{N} D_s}{T_s} \text{bits/second} \quad (1)$$

where, N is the number of acknowledged data frames and D_s is the size of acknowledged data frames in bits. Initiation and transmission of control frames in the network is dependent on number of data frames, which is not the case for management frames. Therefore, we suspect increase in number of management frames may have adverse effect on transfer of data frames. Except beacons and probe frames, other management frames are very few, typically less than 1% in a 30 minute capture of WiFi traffic. Thus, we consider only beacons and probes to analyze their effect on GP.

Beacons help already associated clients to know the current status of BSS (Basic Service Set), such as current channel, BSS load, and new clients to discover the APs by passive scanning. They are transmitted by AP at regular scheduled intervals as determined by TBTT (Target Beacon Transmission Time). Since, the number of APs as well as TBTT do not change often in the network, the number of these frames in a given interval of time does not change much. Probe frames also serve similar purpose as beacons, with a difference that they are initiated by clients as part of active scanning procedure. Clients send probe requests and APs respond with probe responses. Number of probe request frames are dependent on factors such as client roaming, packet losses, and vendors. Number of probe response frames are dependent on number of APs, channel overlap [2], etc . Since these factors are dynamic in nature, the number of probe frames is a dynamic entity, hence the growth of these frames change often with time. These frames have low data rate and they are broadcast in nature. An excessive number of these frames will result in lower GP. Therefore, we study the effect of increased probe frames in the network.

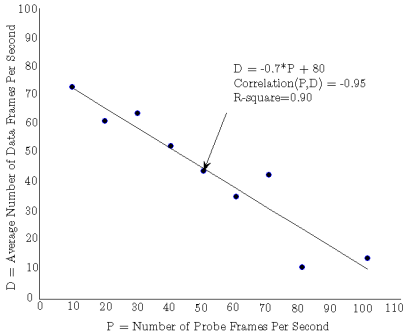

Fig. 2: Average number of D frames versus P frames transmitted in a second. D decreases by slope of -0.7 with increase in P, also P and D are inversely proportional with negative correlation coefficient -0.95.

B. Effect of probes frames on data frames

Figure 2 shows a relationship between number of P and D frames. Instead of client specific P and D frames, we consider network wide frames. Our network setting always had data to be sent. The data presented is averaged over 36 hours of captured traffic in an uncontrolled enterprise environment. For every 10 P frames/second injected in the network, each point on the plot is an average of number of D frames transmitted in 1 second versus number of P frames in that time slot. D frames decrease with a slope of 0.7 with increase in P frames, as shown by Equation 2,

$$D = -0.7 * P + 80 \quad (2)$$

P and D frames are negatively correlated, with correlation coefficient equal to -0.95, therefore an increase in P implies decrease in D. Since our network always had data to be sent, when P increases, channel contention also increases resulting

in (1) Decrease in number of D frames, (2) Loss of ACKs , (3) Increase in data frame retries and (4) Triggering of rate adaptation. All these, further result in reduction of network GP.

C. Understanding the metric P/D

To measure the relation of P and D in a time slot, T_s, we take the ratio of P and D, where considered data frames are fresh data frames. Equation 3 shows a calculation for this metric,

$$P/D = \frac{N_p}{N_d} \tag{3}$$

where, N_p is the number of P frames including probe requests and probe responses and N_d is the number of fresh data frames in a T_s.

Assuming there is always data to be sent, when $P/D < 1$, most of the D frames are able to get channel access and network GP is not affected due to P frames. When $P/D = 1$, both P frames and D frames are able to win the channel contention in equal numbers. In this case, lesser number of D frames are able to win the contention, thus network GP starts reducing. Severe decrease in network GP can be seen when the instances of $P/D > 1$ increase. This growth essentially means P frames are increased to an extent that D frames are unable to get the channel access. Rate of increase of $P/D > 1$ allows us to detect in realtime if GP is reducing due to increase in probe traffic. We take a conservative approach for early detection of growth in probe traffic by considering only fresh data frames for the calculation of P/D. Since, increase in P frames, increase frame retransmissions, considering retransmitted data frames as well, will give a better P/D, thereby hiding actual network condition.

Empirically, we find relation between GP and metric P/D, as given in Equation 4,

$$GP = a * e^{-b*P/D} \tag{4}$$

where, a, b are the constants specific to network under study. This equation is based on curve fitting, with the statistical coefficients for the goodness of fit as $SSE = 0.003338$, $R\text{-}square = 0.9966$, $Adjusted\ R\text{-}square = 0.9965$, and $RMSE = 0.008428$.

D. Using the metric P/D

There is a counter for each of the 3 cases (1) $P/D < 1$, (2) $P/D = 1$, and (3) $P/D > 1$. For each time slot T_s, the counter is incremented in a cumulative manner as per the case. We use $T_s = 1$ second in our analysis. The cumulative increase in the case of $P/D > 1$ for every T_s is plotted against time for a duration of 1 minute, before the counters are reset. Data suggests this duration is long enough to diagnose the growth in probe traffic. Our analysis shows that with 90% accuracy when the slope of this plot is equal to or greater than 0.10, GP starts reducing. Thus, the metric allows us to measure the growth of probe frames in realtime. Our hypothesis is that as the duration of cumulation increases, the accuracy of detection would increase. We will evaluate this hypothesis in our future work.

III. Experiments and Results

Large scale network deployments have varying and large number of clients as well as APs which exaggerate the factors, discussed in Section I, responsible for GP drop of a WLAN. Therefore, to show an increase in probe traffic reduces network GP we first perform a controlled experiment with fixed number of clients and APs as well as controlled network traffic. Later, we demonstrate the drop in network GP due to excessive increase in probe traffic in an uncontrolled enterprise network. All the devices work in g mode. We measure the growth in probe traffic in realtime with the metric P/D in both controlled and uncontrolled network setups.

Experiment Methodology: We inject UDP traffic over a WiFi link in the network and capture wireless traffic to measure GP in presence and absence of excessive probe traffic. UDP traffic is injected to prevent the chances of transport layer re-transmissions. `Iperf` client and server setup [19] is used for traffic injection and bandwidth monitoring of the network. `Iperf` client is executing on a laptop with wireless link, sending data to an `Iperf` server executing on an ethernet client.

Data Collection and Processing: We collect WiFi traffic using a WiFi interface card in monitor mode in controlled network setup and a WIPS (Wireless Intrusion Prevention System) device in an uncontrolled network setup. Both capturing devices scan 2.4 Ghz band passively. They capture the traffic across all channels (1-13) in a round-robin fashion, which is saved in the form of PCAPs (Packet Captures). These PCAPs are further analyzed for probe requests, probe responses, data rates, channels, frame retries, and loss in ACKs.

A. Controlled Network

1) Network Deployment: Figure 3 shows this setup with 1 WiFi client, 1 AP, 1 ethernet client and 5 probing clients. Table I lists the specifications of these devices. The WiFi client is associated with the AP and the ethernet device connected to the ethernet port of the AP. `Iperf` client is executed on WiFi client and `Iperf` server is executed on ethernet client.

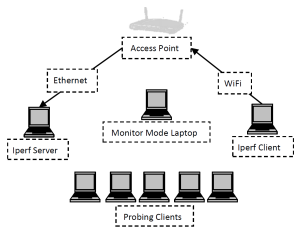

Fig. 3: Experiment setup for controlled network environment.

TABLE I: Specifications of devices used in controlled network

Device Type	Model	Qty
AP	Wireless N-150, WGR614v10	1
Probing Clients	Intel Centrino Wireless-N 1030	1
	TP-Link TL WN721N	2
	Intel Corp. PRO/Wireless 5100	1
	Broadcom Corp. BCM4312	1
	Broadcom Corp. BCM4313	1
`Iperf` Client	Atheros Comm. AR9485	1
`Iperf` Server	Broadcom Corp. BCM5784	1

We initiate UDP data transfer from `Iperf` client to `Iperf` server for an hour. While this data transfer is going on, we enable active probing on 5 laptops, sending 1 broadcast probe request per second. These laptops are not associated to the AP, hence do not transmit any data frames. We used 5 laptops for probing because lesser than 5 probing laptops had merely any effect on bandwidth. These laptops probe the network for 30 minutes before we turn off active probing. Along with our AP, their requests are also responded by neighborhood APs in the residential area. There were 2-4 APs responding to probe requests. Other APs are also far enough to cause any drop in our network's GP. In order to eliminate other causes of GP drop we disable rate adaptation on the `Iperf` client.

2) Analysis and Results: Figure 4 shows network bandwidth as monitored by `Iperf` and Figure 5 shows the realtime detection of growth in probe traffic. This experiment demonstrates (1) an increase in probe traffic reduces network GP and (2) metric P/D detects the growth of probe traffic in realtime. For the first 1000 seconds, active probing is disabled on laptops, it is enabled for 1000 - 3000 seconds, and disabled again for last 1000 seconds. The rate of growth in P/D as seen from the plot is faster for 1000 - 3000 seconds. For this period, drop in bandwidth to as low as 10 Mbps from 20 Mbps is observed. ACK losses increase from 0.04% to 0.3% and frame retries increase from 0.8% to 4%. These factors ultimately reduce GP by upto 6%. The average reduces from 1.6Mbps to 1.5Mbps.

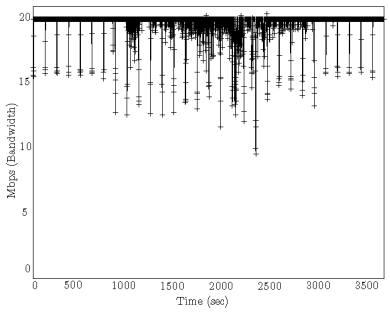

Fig. 4: Bandwidth monitoring of WiFi link as reported at the `Iperf` client in absence and presence of excessive probe traffic. The drop in bandwidth can be noticed from 1000 to 3000 seconds, when network experiences increase in probe traffic.

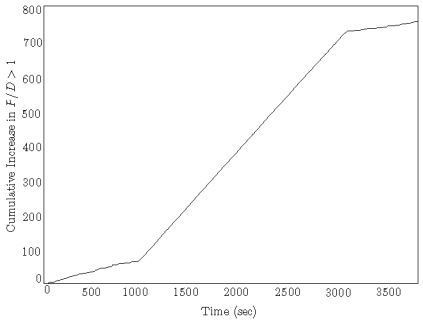

Fig. 5: Realtime detection of growth in probe traffic using metric P/D. Notice the increase in slope from 1000 to 3000 seconds, when clients begin active scanning and decrease from 2700 seconds onwards, when active scanning is turned off at the clients.

TABLE II: Specifications of devices used in uncontrolled network setup

Device Type	Model	802.11
AP	Cisco Aironet 1140 Series	a/b/g/n
`Iperf` Client	TP-Link TL WN721N	a/b/g/n
`Iperf` Server	Ethernet	NA
WIPS device	Airtight Sensor SS-300-AT-C50/60	a/b/g/n

B. Uncontrolled Network

Controlled network setup had fixed number of clients and APs, controlled data traffic, and rate adaptation disabled. The number of devices were lesser as compared to an uncontrolled enterprise network. Our enterprise network had 20 to 40 WiFi clients, upto 15 APs including non-enterprise APs, overlapping channels, rate adaptation enabled and uncontrollable WiFi traffic. In this section, we will demonstrate the effect of increased probe traffic in such a live network.

1) Network Deployment: Figure 6 shows the floor map of the enterprise WLAN, where we collected network traffic of uncontrolled environment. 3 floors of the enterprise building are shown here. We collected traffic on 3^{rd} floor B-wing section of the map shown. The WLAN is deployed with a WLAN controller, 1 WIPS device per floor, and 3 APs per floor. WLAN controller aids in centralized management of APs and clients. It manages all APs by taking care of functions like RRM (Radio Resource Management), QoS (Quality-of-Service), Roaming [20], etc. Interference manage-

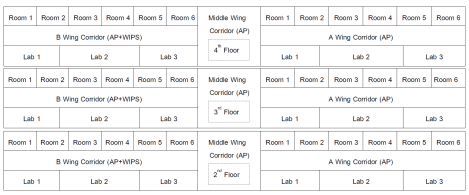

Fig. 6: Building floor map with placement of APs and WIPS.

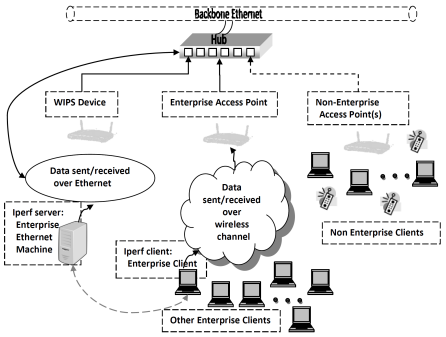

Fig. 7: Experiment setup for uncontrolled network environment.

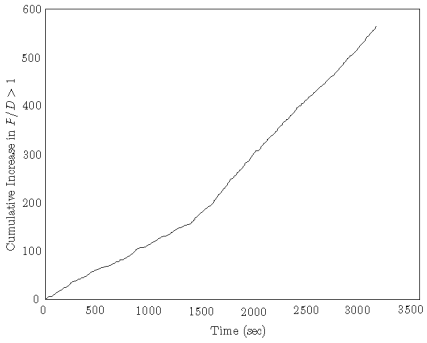

Fig. 8: The plot shows realtime growth in the number of cases corresponding to $P/D > 1$. Notice the increase in slope of metric from 1500 seconds onwards, when probe traffic started to increase.

ment algorithms as part of RRM at the controller takes care of assigning channels dynamically to the APs. Every AP here is broadcasting up to 7 SSIDs. Along with enterprise APs, we notice presence of non-enterprise or external APs in the network. These APs are either mobile hotspots or APs brought by students for their own experiments. They operate in conjunction with enterprise APs, however they are not under the control of the controller. Hence, their operating channels may or may not overlap with those of enterprise AP channels. Setup and specifications of these devices is shown in Figure 7 and Table II, respectively. We analyzed an hour of traffic collected per day for 36 days over a period of 5 months. We tried to collect the data in peak hours of operation, usually 10 AM to 3 PM in working days.

2) Analysis and Results: We present here results of one of the 36 experiment days. However, all days who saw growth in probe traffic experience similar results. Figure 8 shows the change in slope of plot $P/D > 1$ with increase in probe traffic. Network saw 162% growth in probe traffic. Growth in P frames result in decreased network GP as shown in Figure 9. It drops from 2.5 Mbps to as low as 0 Mbps due to increase in unacknowledged data frames, which eventually result in frame retries and drop in frame data rates. Figure 10 shows, the growth in unacknowledged and retried data frames, which eventually drop data rates from 54 Mbps to

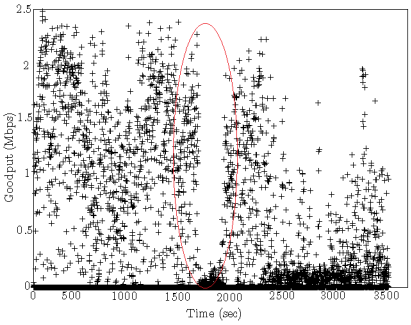

Fig. 9: The plot highlights realtime drop in network MAC GP to as low as 0 Mbps from 1500 seconds with increase in probe traffic.

TABLE III: Details of Probe Responses

Parameter	Value
Probe Responses per probe request	1-7
Responding enterprise APs	2-6
Responding non-enterprise APs	1-5
Locations of responding enterprise APs	3^{rd} Floor: B,M,A-Wing, 2^{nd} Floor: B,M-Wing, 4^{th} Floor: M,A-Wing
Number of enterprise APs whose channels overlap with other enterprise APs	0-5
Number of non-enterprise APs whose channels overlap with other enterprise APs	0-2

as low as 1-2 Mbps. We also analyze the APs, who respond to the probe requests. Table III summarizes the details of all responding APs, with averages calculated over all experiment days. As seen from WIPS device, enterprise APs on the floors above and below across all 3 wings as well as non-enterprise APs also respond to probe requests. The channels of enterprise APs also overlap with other enterprise APs as well as non-enterprise APs. Even the RRM algorithms on controller, do not always assign non-overlapping channels to all the APs. Overlapping channels complicate the situation by introducing RF interference and increased probe responses due to overhearing [3].

IV. Conclusion and Future Work

We presented an empirical study to understand the effect of increased probe traffic on the goodput of a WiFi network and found that increased probe traffic can exponentially reduce network goodput. We introduced a metric of Probes to Fresh Data Frames Ratio (P/D) to measure the increase in probe traffic in realtime and deduce if it is affecting network goodput. We showed that the metric works by testing it in both controlled and uncontrolled network setups. We are working towards a competent solution to reduce excessive probe traffic, such that its implementation does not affect network adversely. Our current experiments do not reveal cases where the metric P/D might fail, it should be tested for false positives and negatives. Lastly, the network conditions, which trigger the increase in probe requests need detailed investigation.

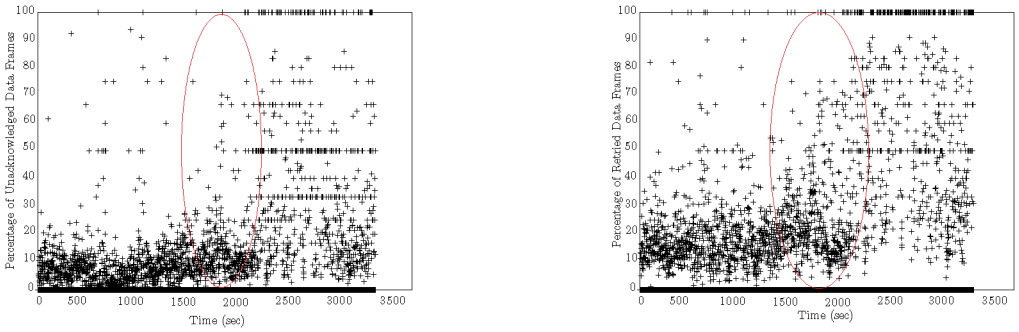

(a) (b)

Fig. 10: Corresponding to Figure 8, as highlighted in the plot as probe traffic increases from 1500 seconds, unacknowledged data frames increase from approximately 10% to almost 100% and retried data frames increase from approximately 20% to 100%.

ACKNOWLEDGMENT

The authors acknowledge support provided by the ITRA project funded by DEITy Government of India under a grant with Ref. No. ITRA/15(57)/Mobile/HumanSense/01 and AirTight Networks.

REFERENCES

[1] my80211, "802.11 client active and passive scanning," [Online; accessed 5-January-2015]. [Online]. Available: http://www.my80211.com/home/2010/1/11/80211-client-active-and-passive-scanning.html

[2] R. Raghavendra, E. Belding, K. Papagiannaki, and K. Almeroth, "Unwanted link layer traffic in large ieee 802.11 wireless networks," *Mobile Computing, IEEE Transactions on*, vol. 9, no. 9, pp. 1212–1225, Sept 2010.

[3] M. Youssef, L. Shahamat, M. Kleene, and A. Agrawala, "The ieee 802.11 active probing analysis and enhancements," in *Wireless Networks, Communications and Mobile Computing, 2005 International Conference on*, vol. 2, June 2005, pp. 1539–1544 vol.2.

[4] S. Lee, M. Kim, S. Kang, K. Lee, and I. Jung, "Smart scanning for mobile devices in wlans," in *Communications (ICC), 2012 IEEE International Conference on*, June 2012, pp. 4960–4964.

[5] I. Ramani and S. Savage, "Syncscan: practical fast handoff for 802.11 infrastructure networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 1, March 2005, pp. 675–684 vol. 1.

[6] Microsoft, "Scanning 802.11 networks," [Online; accessed 5-January-2015]. [Online]. Available: http://msdn.microsoft.com/en-us/library/windows/hardware/ff564113%28v=vs.85%29.aspx

[7] Cisco, "Voice over wireless lan 4.1 design guide - voice over wlan roaming [design zone for mobility] - cisco," [Online; accessed 13-October-2014]. [Online]. Available: http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book/vowlan_ch5.html

[8] I. Purushothaman and S. Roy, "Fastscan: A handoff scheme for voice over ieee 802.11 wlans," *Wirel. Netw.*, vol. 16, no. 7, pp. 2049–2063, Oct. 2010. [Online]. Available: http://dx.doi.org/10.1007/s11276-010-0243-5

[9] M. Emmelmann, S. Wiethoelter, and H.-T. Lim, "Opportunistic scanning: Interruption-free network topology discovery for wireless mesh networks," in *World of Wireless, Mobile and Multimedia Networks Workshops, 2009. WoWMoM 2009. IEEE International Symposium on a*, June 2009, pp. 1–6.

[10] W. L. Tan, M. Portmann, and P. Hu, "A systematic evaluation of interference characteristics in 802.11-based wireless networks," in *Advanced Information Networking and Applications (AINA), 2011 IEEE International Conference on*, 2011, pp. 646–652.

[11] W. Kim, J. Lee, T. Kwon, S.-J. Lee, and Y. Choi, "Quantifying the interference gray zone in wireless networks: A measurement study," in *Communications, 2007. ICC '07. IEEE International Conference on*, 2007, pp. 3758–3763.

[12] U. Das and C. Hood, "Using a shielded room to characterize udp performance in the presence of interference in ieee 802.11 wireless networks," in *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, 2008, pp. 1–5.

[13] S. Khurana, A. Kahol, and A. Jayasumana, "Effect of hidden terminals on the performance of ieee 802.11 mac protocol," in *Local Computer Networks, 1998. LCN '98. Proceedings., 23rd Annual Conference on*, 1998, pp. 12–20.

[14] Y. Zhou and S. Nettles, "Balancing the hidden and exposed node problems with power control in csma/ca-based wireless networks," in *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 2, 2005, pp. 683–688 Vol. 2.

[15] L. Wang, K. Wu, and M. Hamdi, "Combating hidden and exposed terminal problems in wireless networks," *Wireless Communications, IEEE Transactions on*, vol. 11, no. 11, pp. 4204–4213, 2012.

[16] M. Loiacono, J. Rosca, and W. Trappe, "The snowball effect: Detailing performance anomalies of 802.11 rate adaptation," in *Global Telecommunications Conference, 2007. GLOBECOM '07. IEEE*, Nov 2007, pp. 5117–5122.

[17] IEEE, "Ieee-sa - registration authority ma-l public listing," 2014, [Online; accessed 29-September-2014]. [Online]. Available: http://standards.ieee.org/develop/regauth/oui/public.html

[18] "Ieee standard for information technology–telecommunications and information exchange between systems–local and metropolitan area networks–specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications am," *IEEE Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Am*, pp. –.

[19] U. of Illinois, "Iperf - the tcp/udp bandwidth measurement tool," [Online; accessed 5-January-2015]. [Online]. Available: https://iperf.fr/

[20] L. Phifer, "Buyer's guide to enterprise wlan controllers," 2011, [Online; accessed 9-October-2014]. [Online]. Available: http://www.enterprisenetworkingplanet.com/netsysm/article.php/3924291/Buyers-Guide-to-Enterprise-WLAN-Controllers.html