IEEE 802.11 in the Large: Observations at an IETF Meeting

Andrea G. Forte, Sangho Shin and Henning Schulzrinne
Department of Computer Science
Columbia University
Email: {andreaf, sangho, hgs}@cs.columbia.edu

Abstract— **We observed wireless network traffic at the 65th IETF Meeting in Dallas, Texas in March of 2006, attended by approximately 1200 engineers. The event was supported by a very large number of 802.11a and 802.11b access points, often seeing hundreds of simultaneous users. We were particularly interested in the stability of wireless connectivity, load balancing and loss behavior, rather than just traffic. We observed distinct differences among client implementations and saw a number of factors that made the overall system less than optimal, pointing to the need for better design tools and automated adaptation mechanisms.**

I. Introduction

In recent years, IEEE 802.11 networks have experienced rapid growth. One of the main problems is how to deploy an 802.11 network in very crowded environments so that each user has a minimum amount of guaranteed bandwidth. Few studies have been conducted on large scale wireless networks and all of them revealed many limitations of the current 802.11b/g standards in highly congested environments.

We analyze the data collected by monitoring the IEEE 802.11a/b wireless network deployed at the 65th IETF meeting held in Dallas, TX between March 19th and March 24th, 2006. The meeting was attended by roughly 1200 engineers and the IEEE 802.11 network comprised more than 90 Access Points (APs), making this one of the largest indoor IEEE 802.11 networks analyzed to date.

We took our measurements on three of the six days of meetings, from March 21st to March 23rd and collected 25GB of data. For the analysis of this data, we focused on one of the three days, March 22nd and on one of the many rooms in which sessions were held, the room named Chantilly. This choice was made because March 22nd was one of the busiest days at the IETF meeting and room Chantilly was used throughout the day, including a plenary session from 17:00 to 19:30. During the plenary there were no other sessions in progress in other rooms and more than 500 clients attended the plenary in this room. This made Chantilly an ideal place to study congested environments.

In conducting this study, we were not interested in traffic analysis, but rather in characterizing unusual behaviors due to the highly congested environment.

The rest of the paper is organized as follows: Section II gives an overview of other studies done in highly congested environments; in Section III we give an overview of the IETF wireless network and of our measurement setup, we also give some statistics on the use of the network; Section IV shows how a load balancing algorithm based on the number of users rather than on per-client bandwidth, can achieve good performance in highly congested environments. In Section V we analyze the handoff behavior of the wireless clients finding, for example, that Apple wireless clients behave better than other vendors' clients; Section VI looks at some of the consequences of deploying many adjacent APs on the same channel, discovering that this introduces high interference and a lot of network inefficiencies such as broadcast and multicast packet duplication. Section VII concludes the paper.

II. Related Work

In [1] the authors monitored the wireless traffic at SIGCOMM 2001. The conference was held in U.C. San Diego in August 2001 and was attended by about 200 participants. Four IEEE 802.11b APs were deployed using channels 1, 4, 7 and 11. The authors found that the throughput on each channel was not proportional to the number of clients on that channel but rather was proportional to the bandwidth use of each client. Load balancing algorithms should, therefore, take into consideration the bandwidth used by each client and not just the number of clients.

Jardosh et al. [2] analyzed the wireless network deployed at the 62nd IETF meeting (March 2005). The IEEE 802.11b wireless network comprised 38 APs that used channels 1, 6 and 11. The APs would dynamically decide which of the three channels to use according to some non-specified proprietary load balancing policy. The study showed data transmissions at lower data-rates are more likely to succeed than at higher data-rates. They also propose to calculate link reliability using the beacon loss rate and estimate channel congestion using the correlation between retransmission rate and data transmission rate.

Rodrig et al. [3] monitored wireless traffic at SIGCOMM 2004. The conference was attended by roughly 550 participants. The wireless network serving the conference was an IEEE 802.11b network. Only five APs were installed and they used channels 1, 8 and 11. Some of the main results of this study were the high overhead of the 802.11 protocol, frequent retransmissions and changes in client data-rates. The data transmission rate was analyzed in detail finding that low data-rate had lower probability loss than higher data-rate, although with a minor difference.

Other studies, [4], [5], analyze users' behaviors like roaming patterns and average number of visited APs rather than focusing on network issues like throughput, interference and packet loss.

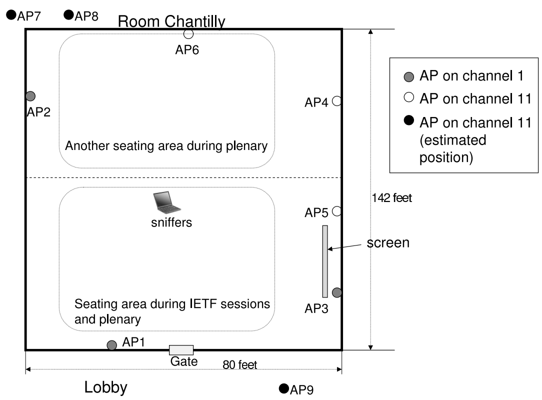

Fig. 1. Measurement and network setup in Chantilly

III. The Wireless Network at the IETF Meeting

In this section, we describe the wireless network environment at the IETF meeting, the measurement setup, and the usage of the wireless network.

A. Wireless Network Setup

The 65th IETF meeting was held at the Hilton Anatole hotel in Dallas, TX. The hotel had conference rooms located on two different levels. All the conference rooms already had an 802.11b wireless coverage. However, in order to support the large number of participants, the IETF Network Operations Center (NOC) decided to deploy more APs. The NOC installed a total of 91 IEEE 802.11a/b Cisco Aironet 1200 APs around the various conference rooms on the first and second floor of the hotel conference center in order to increase the capacity and coverage of the wireless network. IEEE 802.11a allowed NOC to install multiple APs in the same area without any interference among APs due to the large number of non-overlapping channels, while the IEEE 802.11b network was meant to be used as backup. The 802.11b APs were set up to use only channels 1 and 11 since all the hotel APs used channel 6.

No wireless security was enabled in the wireless network, and the whole wireless network formed one Extended Service Set (ESS) with ESSID *ietf65*. The hotel APs and the APs deployed by NOC formed one single subnet.

B. Measurement Setup

In our measurements we used four IBM T42 Think Pad laptops as sniffers, each with one Proxim ORiNOCO 11a/b/g combo wireless card. We used Airopeek NX [6], a commercial network analyzer, as a wireless sniffer. Airopeek can capture both data and 802.11 management frames such as 802.11 Acknowledgement (ACK) frames, beacon frames, probe requests and responses. It also allows to monitor signal strength and transmission data-rate on a per-packet basis.

Although the 802.11b network was supposed to be used as a backup, we found out, with preliminary measurements, that 802.11b was the most used network, and hence decided to focus our measurements on 802.11b traffic. We configured three of the four sniffers to monitor channels 1, 6 and 11, one sniffer per channel.

As said in Section I, we focus our analysis on the measurements taken in room Chantilly ($142' \times 80'$, total capacity of about 600 persons). On March 22nd there was one session in the morning, two sessions in the afternoon, and a plenary session in the evening. During the plenary, there were no other IETF sessions ongoing in other rooms, and most of the attendees of the day participated in the plenary. This allowed us to measure very large scale traffic on the wireless network with more than 500 clients in the same room.

Fig. 1 shows the position of the APs, clients and sniffers in Chantilly. Only half of the room was used during the three regular IETF sessions, while the whole room was used for the plenary session. Because of this and given the large number of APs used, we set the sniffers at the center of the room to capture the maximum number of frames from APs and clients during the regular IETF sessions and the plenary. Since we were interested in the detection of anomalies in the network, this positioning of the sniffers allowed us to get the same "view" of the medium than normal clients experience without affecting the correctness of our results. We located three APs on channel 1 and three on channel 11, inside Chantilly; three other APs on channel 11 had been positioned outside of Chantilly, although they were rarely used. From our measurements we also detected 14 hotel APs on channel 6, six of which appeared to be installed inside Chantilly. We do not know the position of the hotel APs.

With separate measurements we also found that the range of each AP was large enough to cover the whole room, confirming that positioning our sniffers the way we did, allowed us to capture most of the frames from the various APs. Nevertheless, as we will discuss later in Section VI, such a broad coverage introduces a significant amount of interference among the APs.

IV. Load Balancing

In large crowded wireless networks load balancing becomes critical for achieving fair distribution of resources and bandwidth among clients. At the IETF meeting, no load balancing algorithm was used. We observed some problems that a good load balancing algorithm could have prevented, allowing a higher degree of fairness in the utilization of network resources.

A. Distribution of Wireless Clients

Since load balancing was not used at the IETF meeting, clients connected to a particular AP only according to their proximity to the AP.

Table I shows the distribution of clients on channel 1. AP1 on channel 1 had more clients on average than AP2 and AP3 because AP1 was the closest AP to most of the clients during the IETF sessions. The number of clients on AP1 and AP2 became comparable during the plenary as the whole room, and not just half of it, was used to accommodate people. This allowed for a more even distribution of clients (see Fig. 1). The number of clients associated to AP3 was very small in all IETF sessions even though AP3 was in close proximity of AP2, covering roughly the same area. From Fig. 1 we can see that AP3 was located behind a projection screen. In general, these kinds of screens contain a significant amount of metal

TABLE I

NUMBER OF CLIENTS ON CHANNEL 1

	AP1	AP2	AP3	Total
Session 1 (AM)	95	56	28	163
Lunch	35	42	18	102
Session 1 (PM)	85	52	36	167
Plenary	158	132	30	320

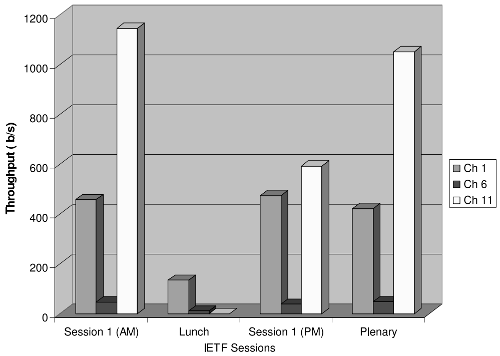

Fig. 2. Average throughput per client

to make the projected image brighter and higher in contrast. Because of this, the screen severely attenuated AP3' signal strength, which caused less clients to be able to associate with it.

B. Throughput

Fig. 2 shows the average throughput per client on the three 802.11b channels. The throughput is calculated considering IP packets transmitted and received by clients and APs: we compute the total size of each IP packet transmitted every second (B/s) on each channel and calculate the average.

Looking at Fig. 2 the average throughput per client seems to be extremely low. This, however, is not true. In particular, Fig. 2 shows the throughput per client averaged on the *whole* session. This means that Fig. 2 must not be used to compute the actual throughput experienced by clients. As we can see from Fig. 3, the total throughput averaged every minute reaches peaks of 3.2 Mb/s which is close to the maximum throughput achievable in IEEE 802.11b networks. What we want to show in Fig. 2 is the significant difference in throughput between channels, and between networks, which indicates the importance of a load balancing algorithm in highly congested wireless networks.

Balanchandran et al. found [1] that the number of clients does not correlate with the throughput and argued that the

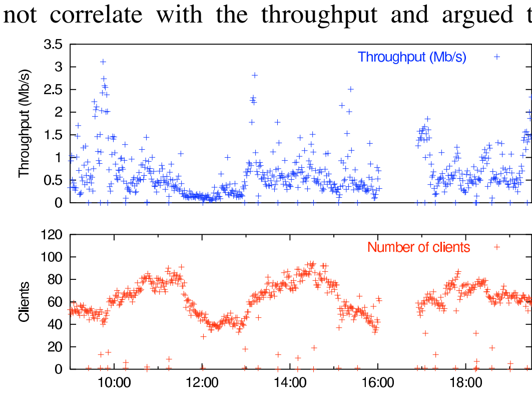

Fig. 3. Total throughput and number of clients on channel 6

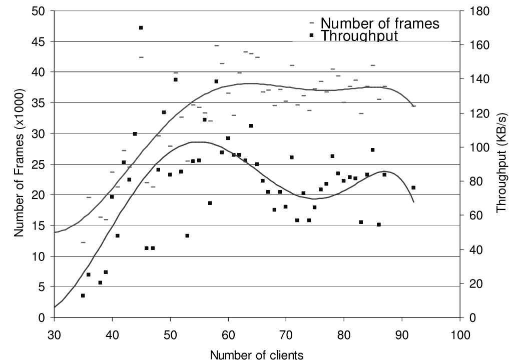

Fig. 4. Throughput and num. of frames vs. num. of clients on channel 6

TABLE II

FRACTION OF APPLICATION PROTOCOLS PER 802.11B CHANNEL

Channel	HTTP	SSH	DNS	IMAP	NB-NAME
1	25.1%	5.7%	4.3%	3.8%	5.6%
6	20.9%	5.1%	5.8%	1.4%	9.1%
11	24.2%	3.7%	4.4%	4.4%	6.0%

throughput per client represents a better metric for load balancing. In our measurements, we found a reasonable correlation between the number of clients and the traffic load in each channel. This would suggest that in highly congested environments, the number of clients still represents a good metric for load balancing and it is much simpler to adopt.

Fig 4 shows the correlation between number of clients, number of frames and throughput (KB/s) on channel 6, the most congested channel. We can see that, initially, as the number of clients increases, the total throughput increases. However, after the number of clients reaches a certain value, the throughput starts decreasing. As we can see from Fig. 4, at the IETF meeting this certain value was about 55 clients and it represents the maximum number of clients the channel can handle, that is, the channel capacity. Once the number of clients exceeds capacity, collisions and retries increase bringing down the overall throughput. In Fig. 4 we can see that when the number of clients exceeds 55, the overall throughput starts decreasing even though the total number of frames remains roughly the same.

In highly crowded environments we can assume that different types of traffic are evenly distributed between channels, that is, on average, the network utilization per client is the same between all the clients. Under this assumption, doing load balancing according to the number of clients rather than to the throughput per client, achieves good results with less complexity. Table II shows the distribution of different traffic types among the three 802.11b channels used at the IETF meeting. As we expected, they appear to be evenly distributed among the channels.

V. HANDOFF ANALYSIS

Because of the particular configuration of the APs and because of the large number of clients these APs had to serve, we were able to observe non-typical handoff behaviors. The following sections show the main factors responsible for such behaviors.

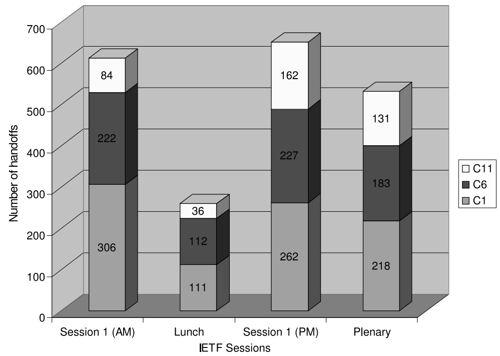

Fig. 5. Total number of handoffs

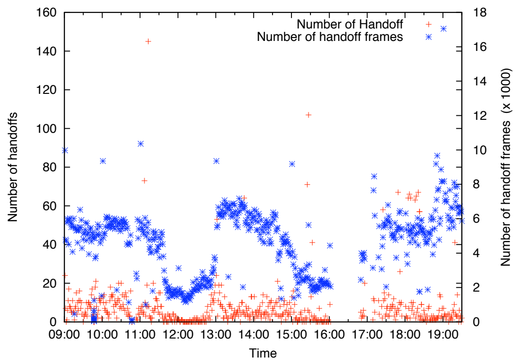

Fig. 6. Total number of management frames detected on channel 1 in Chantilly vs. the total number of handoffs on the same channel

A. Handoff Behavior

Generally speaking, in a highly congested environment the first thing to notice is the very high number of handoffs performed by clients. Usually, in highly crowded environments, most of the handoffs are triggered by congestion, that is, by a significant packet loss [7]. Packet loss is mainly caused by collisions due to medium access and by poor channel conditions.

Fig. 5 shows the total number of handoffs observed on each channel, per IETF session. We can see that the highest number of handoffs was performed by clients during the morning session, followed by plenary session and afternoon session. In all cases most handoffs occurred on channel 1 and channel 6. Less handoffs were observed on channel 11. This is consistent with the distribution of clients over the three channels.

In Fig. 6, we show the total number of handoffs on channel 1 and the number of corresponding management frames, averaged on the minute. As we can see, there are peaks of up to 150 handoffs and 17,000 handoff frames per minute. This, together with the fact that handoffs were triggered by poor channel conditions and not by client's movement, clearly shows the large impact that handoffs had on the congestion of the network and on the overall user experience.

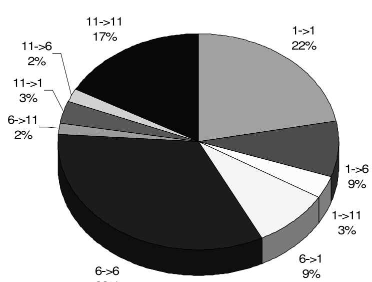

Fig. 7. Handoffs between channels

TABLE III

PERCENTAGE OF HANDOFFS PERFORMED TO THE CURRENT AP

	S1 (AM)	Lunch	S1 (PM)	Plenary
Handoffs (%)	30.5	33.0	30.2	54.7

TABLE IV

SESSION TIME: TIME BETWEEN HANDOFFS

Session time	< 1 min	< 5 min	< 10 min	>= 10 min
Percentage of handoffs	22.8%	34.0%	11.5%	31.7%
Percentage of clients	23.8%	11.9%	5.5%	58.8%

As we can see from Fig. 7, the vast majority of handoffs was performed between APs on the same channel. In particular, handoffs between APs on channel 6 are responsible for 33% of the total handoffs, handoffs between APs on channel 11 are responsible for 17% of the total handoffs and handoffs between APs on channel 1 are responsible for 22% of the total handoffs. About 72% of the total handoffs were performed between APs on the same channel. Furthermore, in the worst case scenario, 54.7% of the total handoffs were performed to the current AP - that is, to the same AP the client just disconnected from. Table III shows the percentage of handoffs in which current AP and next AP are the same.

Performing a handoff to the current AP is useless, and also performing handoffs between different APs on the same channel does not help at all. A client moving between two APs on the same channel experiences the same level of congestion, throughput and packet loss before and after the handoff. The channel is the same, the channel conditions are the same and the number of contentions on the channel is the same. Potentially, this can lead to a situation where the client is repeatedly and frequently performing handoffs to the same channel, as if it was "trapped" on that particular channel. We have observed this anomalous behavior and it is shown in Fig. 7 and Table IV. The session time shown in Table IV is defined as the time between two consecutive 802.11 Association Response frames for a particular client, that is, the time in between handoffs for that client. As we can see from Table IV, the time in between handoffs is less than one minute for 22.8% of the handoffs and 34% of the handoffs are performed within one to five minutes. The percentage of handoffs performed within ten minutes or more is 31.7%. This clearly shows that handoffs happened very frequently and most of them between APs on the same channels. Table IV also shows the percentage of clients that performed a handoff within the specified times.

Having clients performing frequent handoffs to APs on the same channel or to the same AP over and over, causes disruptions in the network connection without introducing any advantage. Also, this represents a problem not just for those clients performing the handoff, but for all the clients on that channel. Every time a handoff happens, management frames are exchanged between the station performing the handoff and the target AP. IEEE 802.11 management frames are always

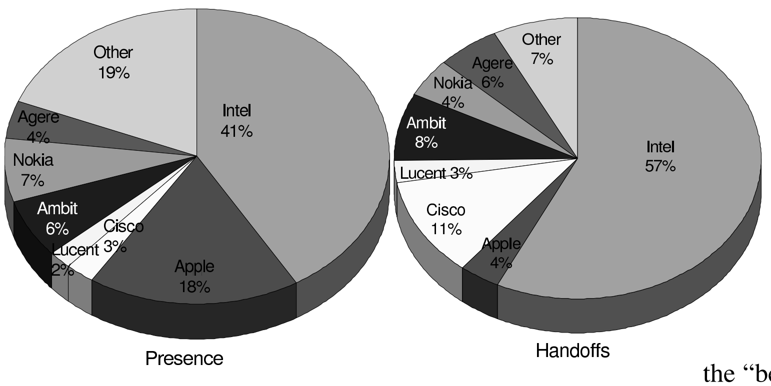

Fig. 8. Handoff statistics per wireless card vendor

TABLE V

PERCENTAGE OF HANDOFFS WITHIN 1 MINUTE AND WITHIN 5 MINUTES

Vendor	Nokia	Intel	Agere	Lucent	Ambit	Apple	Cisco
<= 1 min	30.3%	24.1%	26.3%	21.0%	17.4%	1.2%	24.3%
<= 5 min	49.5%	57.5%	64.3%	38.7%	75.8%	3.5%	83.0%

transmitted at the lowest available bit-rate, thus keeping the medium busy for longer and preventing other stations from accessing the medium. The time needed to send a Probe Request[1] is 816 μs while the time needed to send a data frame of size equal to the MTU size of 1500 bytes at 11 Mb/s is 1307.6 μs. As we can see, the time needed to send the shortest management frame is about $2/3$ of the time needed to send the longest data frame at the maximum bit-rate of 11 Mb/s. This clearly shows the large overhead involved in sending management frames. Because of this, unnecessary handoffs degrade network performance by increasing network congestion for all the clients on a particular channel.

At the IETF meeting, probe requests and responses were responsible for 10.4% of the total network traffic, with probe requests taking only 1.5% of the traffic. This big difference between the number of probe requests and probe responses is mainly due to the fact that there were many retries for probe responses and many APs on the same channel would all answer to the same probe request. The high number of retries for probe responses was mainly because of the high degree of congestion in the network, that is, the high number of collisions. The percentage of probe requests and responses shown above is computed based on the total number of frames and not on the total number of bytes. It is important to consider the number of frames because the overhead to transmit frames in a congested channel is large and, as we have seen earlier, the total transmission time of a probe request frame is almost as long as the time needed to send an MPDU data frame.

From the previous results we can see how today there are many problems with the way MNs select the AP to connect to. In particular, the AP is selected according to the link signal strength and Signal-to-Noise Ratio (SNR) levels. Other factors such as effective throughput, number of retries, number of collisions, packet loss, bit-rate or Bit-Error-Rate (BER) are ignored. When a client needs to perform a handoff, it has to look for a different AP to connect to. Unfortunately, with a very high probability, the client will pick the same AP it was connected to because its link signal strength and SNR are still

[1]Probe requests are the shortest management frames in IEEE 802.11b networks.

TABLE VI

DISTRIBUTION OF BROADCAST TRAFFIC

802.11b Channel	ARP Requests	Beacon Frames	Probe Requests	Redundant Broadcasts
1	6.8%	35.2%	12.8%	45.1%
6	18.8%	32.7%	12.7%	35.6%
11	3.8%	45.1%	17.5%	33.5%

TABLE VII

PERCENTAGE OF BROADCAST TRAFFIC

Channel	Redundant Broadcasts	Total Broadcasts
1	3.4%	7.5%
6	6.0%	17.0%
11	3.1%	9.3%

the "best" available. The information regarding the congestion of the AP is completely ignored and this bad behavior keeps repeating itself. This behavior can create situations where users end up connecting all to the "best" AP creating the scenario depicted earlier and at the same time leaving other APs under-utilized [2].

All of this clearly shows the need for better handoff algorithms and more robust implementations.

B. *Vendors and Handoff*

There were about 1200 attendees with cards from many different vendors, dominated by Intel wireless cards (see Fig. 8). Most of the different vendors had similar handoff policies and algorithms as they behaved pretty much in the same way (Fig. 8), except for Apple, whose cards were used by 18% of attendees yet only caused 4% of the handoffs. Apple has the lowest number of handoffs per client among the different card vendors. On average, an Apple wireless client performed no handoff at all during the day of meetings. Furthermore, looking at Table V, we can see that while cards of other vendors performed poorly by having a lot of unnecessary handoffs, the percentage of handoffs performed by an Apple client within 1 minute and within 5 minutes was 1.2% and 3.5% respectively, clearly showing the adoption of a conservative handoff algorithm. A conservative handoff algorithm proved to be the best choice for highly congested environments such as the IETF meeting.

On average, all clients from all vendors stayed connected to the network for the same amount of time. In our analysis, we assumed same deviation of usage across each vendor.

VI. BROADCAST AND MULTICAST TRAFFIC

As we said earlier, at this IETF meeting APs were deployed so that adjacent APs used the same channel and covered roughly the same area. In addition to the problems discussed earlier, this also introduces problems for broadcast and multicast traffic and increases interference.

Broadcast and multicast traffic represented 10.5% of the total traffic. We discovered significant overhead in the IETF network introduced by broadcast and multicast frames. When a node in the network sends a broadcast frame, this frame is duplicated by all the APs in the subnet. If this frame is an ARP request, for example, this is the correct behavior as any node in the subnet might be the one that has to respond to the request. However, things are different if the broadcast frame is, for example, a DHCP request. In this case, the target of such

TABLE VIII

PERCENTAGE OF MULTICAST TRAFFIC

Channel	IPv6 Multicast	IPv4 Multicast
1	0.8%	1.8%
6	2.0%	3.4%
11	0.8%	1.6%

a frame is a DHCP server[2] and not other clients. Nevertheless, the DHCP request is sent to all the clients of all available APs, thus introducing unnecessary traffic. This situation becomes even more critical when multiple adjacent APs use the same channel. In this case, we have unnecessary traffic even with legitimate broadcast frames such as ARP requests. The ARP request is sent over the same channel a number of times equal to the number of APs serving that channel. This means that if we have three APs on channel 1, for example, the same ARP request will be sent three times to the clients on channel 1, furthermore the three APs will each have to contend access to the medium in order to send such a frame.

We have categorized broadcast frames and the respective protocols in redundant and non-redundant, depending on who should receive these frames and who actually receives them. For example, ARP requests are non-redundant as the reply to the ARP request could come from any client connected to any AP. On the other hand, DHCP requests are redundant as sending these frames to other clients is useless since the target of such packets is a DHCP server. Other non-redundant frames are beacons and probe requests. The first ones are sent by an AP to its clients and the latter ones are sent by clients to APs which do not propagate them any further. To summarize, in regard to broadcast traffic, in our measurements we have encountered and classified the following frame types and protocols:

- *Redundant*: NetBios, UDP, Apple Talk (NBP lookup, ZIP), DHCP, TiVO.
- *Non-redundant*: ARP Requests, Beacons, Probe Requests.

From Table VI we can see that redundant broadcasts are 45.1%, 35.6% and 33.5% of the total broadcast traffic on channel 1, 6 and 11, respectively. From Table VII we can see that, on channel 6, 17% of the traffic is broadcast traffic. The reason for such a high percentage of broadcast traffic on channel 6 is the larger number of clients connected to the APs on channel 6. As we can see from Table VII, the percentage of redundant broadcast frames on channel 6 is 6% of the network traffic which is almost twice the amount of redundant broadcast traffic on the other two channels. This significant difference with channels 1 and 11 is due to the larger number of adjacent APs using channel 6.

Similarly, *all* multicast frames are forwarded to all the APs in the network. From our measurements, Bonjour DNS queries are responsible for more than 90% of all multicast traffic, followed by IGMP frames making up almost all of the other multicast traffic. Table VIII shows statistics for multicast, showing the presence of some IPv6 traffic as well.

All of this superfluous traffic significantly contributes to the congestion level of the wireless network.

[2]The DHCP server was located in the fixed network.

VII. CONCLUSIONS

We have analyzed the data collected in the wireless network at the 65th IETF meeting held in Dallas, TX, from March 19th to March 24th, 2006. About 1200 engineers attended the meeting, giving us the opportunity to study IEEE 802.11a/b wireless networks in a highly congested environment.

In our measurements we observed a very large number of handoffs. These handoffs were triggered by packet loss due to congestion since all clients were stationary during the sessions. About 72% of them were performed between APs on the same channel and, during the plenary, the number of handoffs from and to the same AP reached 54.7%. Handoffs also occured very frequently, with 24% of them happening within one minute and 12% happening between 1 and 5 minutes. Furthermore, the percentage of probe request and response frames reached 10.4% of the total network traffic. This is a far-from-optimal behavior with a lot of unnecessary handoffs that cause disruptions in users' connectivity and increase congestion in the network. 41% of wireless cards were Intel wireless cards followed by Apple cards with 18%. Apple clients performed particularly well in terms of number of handoffs, being responsible for only 4% of the total handoffs. Cisco clients, 3% of the total, contributed 11% of the handoffs.

Installing multiple APs on the same channel covering the same area introduces considerable overhead. In particular, broadcast and multicast packets are duplicated at each AP, thus wasting bandwidth and contributing to the high level of congestion.

We found a clear correlation between the number of clients and network utilization among channels which would suggest that, in highly populated wireless networks, a load balancing algorithm based on the number of clients rather than on the throughput per client, would achieve good performance while keeping complexity low.

Solutions to many of the problems presented in this paper are reserved for future study.

REFERENCES

[1] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless LAN," in *SIGMETRICS '02*. New York, NY, USA: ACM Press, 2002, pp. 195–205.

[2] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding link-layer behavior in highly congested IEEE 802.11b wireless networks," in *E-WIND '05*. New York, NY, USA: ACM Press, 2005, pp. 11–16.

[3] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *E-WIND '05*. New York, NY, USA: ACM Press, 2005, pp. 5–10.

[4] D. Kotz and K. Essien, "Analysis of a campus-wide wireless network," in *MobiCom '02*. New York, NY, USA: ACM Press, 2002, pp. 107–118.

[5] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *MobiCom '04*. New York, NY, USA: ACM Press, 2004, pp. 187–201.

[6] WildPackets. (2006) AiroPeek NX: Expert Wireless LAN Analyzer. [Online]. Available: http://www.wildpackets.com

[7] G. Karlsson and H. Velayos, "Techniques to reduce the IEEE 802.11b handoff time," in *ICC'04*, June 2004, pp. 3844–3848.