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. We found the current handoff algorithms to be highly inefficient and in many cases counterproductive, usually leading to wrong AP selection thus increasing the overall congestion of the network. Using multiple APs on the same channel proved to be inefficient for multicast and broadcast traffic, while still increasing the level of co-channel interference and therefore packet loss. Finally, we found the deployment of a load balancing algorithm simply based on the number of users per AP to be efficient enough in such highly congested environments. 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 the per-client bandwidth can achieve good performance in highly congested environments. Section V gives an overview on channel assignments in IEEE 802.11 networks. In Section VI we analyze the handoff behavior of the wireless clients finding, for example, that Apple wireless clients behave better than other vendors' clients; Section VII 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. In Section VIII we discuss some handoff behaviors typical of Windows XP clients and in Section IX we consider interference caused by obstacles such as screens and further analyzed it to prove the destructive effect of such obstacles. Section X concludes the paper.

II. Related Work

Jardosh et al. [1] 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. [2] 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.

In [3] 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.

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.

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, the number of hotel APs was too small to support the large number of participants, therefore 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*.

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,

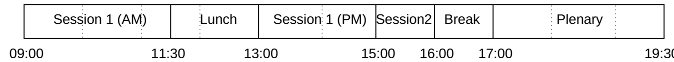

Fig. 1. Timeline of IETF sessions in Chantilly

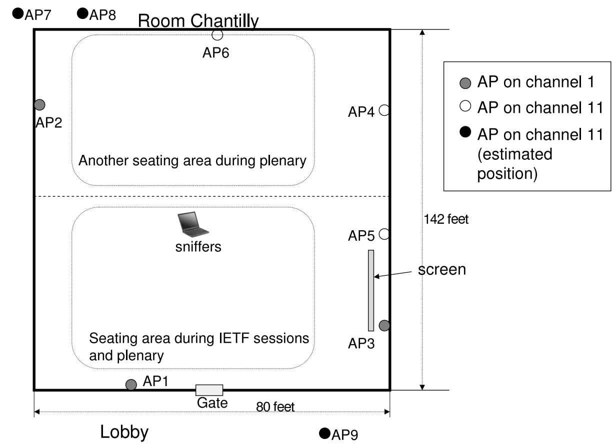

Fig. 2. Measurement and network setup in Chantilly

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. We then used the fourth sniffer to monitor traffic on all of the eight channels used in the 802.11a network.

As we said in Section I, we focus our analysis on the measurements taken in room Chantilly ($142' \times 80'$, total capacity of about 600 persons). Fig. 1 shows the timeline of IETF sessions 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.

Fig. 2 shows the positions 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. 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.

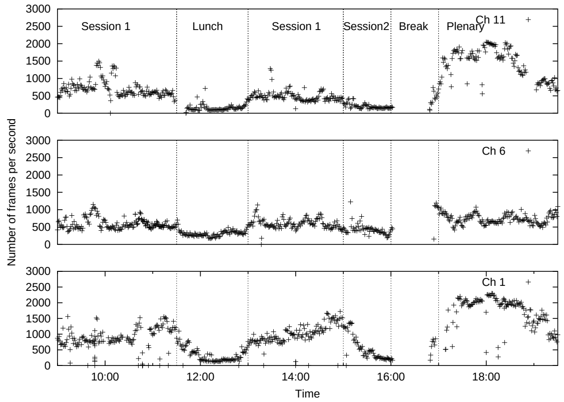

Fig. 3. Total number of frames detected on each 802.11b channel in Chantilly

TABLE I
TOTAL NUMBER OF CLIENTS DETECTED

IETF sessions	S1(AM)	Lunch	S1(PM)	Plenary
Number of clients	395	335	414	536

Nevertheless, as we will discuss later in Section VII, such a broad coverage introduces a significant amount of interference among the APs.

C. Use of the Wireless Network

1) Traffic Volume: Fig. 3 shows the number of frames per second during each IETF session on each of the three 802.11b channels. The number of frames is sampled every minute and shown as the average number of frames per second. In computing the number of frames per second, we considered all the MAC-layer traffic including 802.11 management frames, but did not count corrupted frames.

We can see that channels 1 and 11 are used more than channel 6 and, during the plenary session, the number of frames peaks to about 2500 frames per second for both channel 1 and channel 11. The number of frames per second on channel 6 remains low. As we will explain in more detail in Section VII, this is because clients on channel 6 experience higher congestion than clients on channels 1 and 11.

In the rest of the paper we focus on four time slots: Session 1 (AM), Lunch, Session 1 (PM) and Plenary. We will not consider Session 2 and Break because during these two time slots, Chantilly was not used for meetings. In particular, from Fig. 3 we can see some traffic during Session 2. This traffic was due to the presence of some people in the lobby outside of Chantilly that were not attending any meeting and that connected to the wireless network of Chantilly.

2) Number of Clients: Table I shows the number of wireless clients in Chantilly during each IETF session. Only those clients that transmitted or received at least one error-free IP packet were counted.

3) Protocols: Fig. 4 shows the protocols used by clients and APs. The protocols in the graph are mutually exclusive meaning that, for example, TCP includes only TCP packets

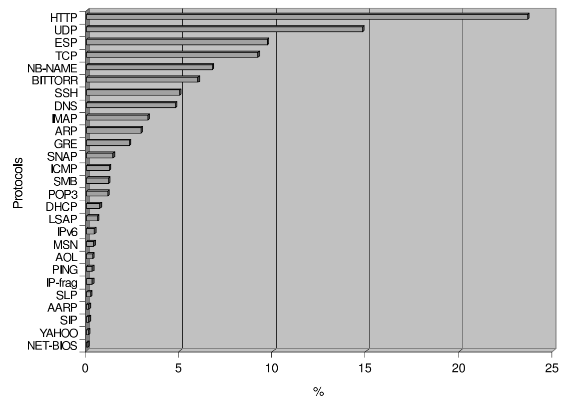

Fig. 4. Protocols used by clients

whose application protocol was not recognized by our sniffers. Therefore, TCP packets do not include HTTP or SSH packets.

We can see from Fig. 4 that users used the wireless network mostly for web surfing and logging in to servers using SSH. Many users were using VPNs as we can see from 10% of the packets being ESP (Encapsulating Security Payload) packets. We also detected BitTorrent traffic, accounting for 6% of the total network traffic. BitTorrent traffic was observed only during the plenary and was responsible for 20% of the traffic on channel 1. Considering the high volume of traffic and high level of congestion during the plenary (see Section VII), traffic such as BitTorrent should have lower priority or available bandwidth should be limited on a per-client basis, preventing a single user from taking a large part of the available bandwidth.

IV. LOAD BALANCING

In large crowded wireless networks load balancing becomes critical for achieving fair distribution of resources and bandwidth among clients. If the number of clients or amount of throughput is not balanced among the different available channels, the clients on the most congested channel will experience high congestion, while clients on the under-utilized channel will experience no congestion at all. Overall, the network throughput will decrease as some channels are under-utilized and other channels experience high congestion. At the IETF meeting, no load balancing algorithm was used. However, 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

Fig. 5 shows the total number of clients on the three 802.11b channels and the number of clients using 802.11a. We can see that channel 6 had the largest number of clients followed by channel 1 and channel 11. Only around 20% of the clients used 802.11a. The low utilization of the 802.11a network can be explained with the much larger number of wireless cards supporting 802.11b/g networks[1]. The big difference in the

[1]Currently, Centrino chipsets support 802.11b/g standards only.

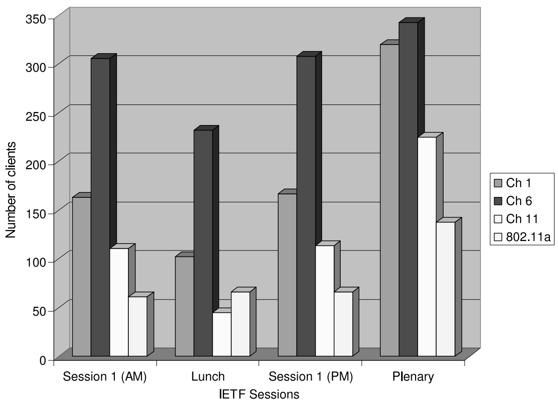

Fig. 5. Number of clients in Chantilly

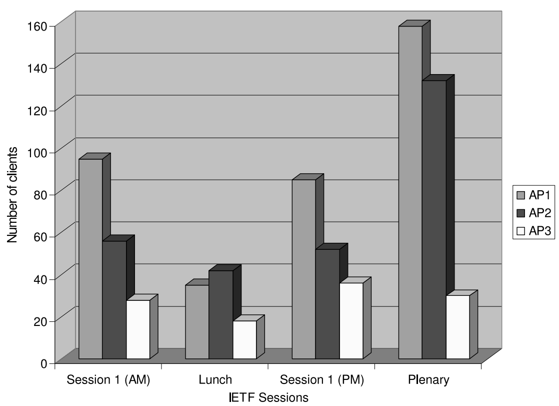

Fig. 6. Number of clients on channel 1

number of clients among the three 802.11b channels indicates the need for a load balancing algorithm. Furthermore, once 802.11a wireless cards become more common, having a load balancing algorithm that would also balance between 802.11a and 802.11b/g networks would be helpful.

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

The hotel APs on channel 6 were located in the ceiling of each room. The distance between these APs and the wireless clients was less, on average, than the distance between the clients and the APs on other channels. Since all the APs had the same transmission power, for many clients the signal strength of the APs on channel 6 was stronger than that of APs on other channels, causing them to connect to that AP.

Fig. 6 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. 2).

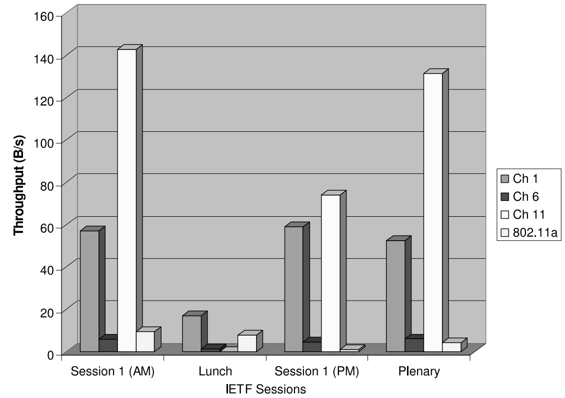

Fig. 7. Average throughput per client

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. 2 we can see that AP3 was located behind a projection screen. In general, these kinds of screens contain a significant amount of metal 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. We performed some measurements on the screen effect and found that the presence of the screen consistently increases the packets' retry rate. The screen effect will be analyzed further in Section IX.

B. Throughput

Fig. 7 shows the average throughput per client on the three 802.11b channels and on all used 802.11a 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. Because we were using only one sniffer for monitoring traffic on the eight 802.11a channels, the wireless card used to monitor 802.11a traffic had to continuously switch channel. This introduced packet loss in the measurements of the 802.11a network. This packet loss contributes to the very low throughput observed in the 802.11a network. However, the significant difference in throughput between channels, and between networks, indicates the importance of a load balancing algorithm in highly congested wireless networks.

Balanchandran et al. found [3] that the number of clients does not correlate with the throughput and argued that the 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 8 shows the correlation between number of clients, number of frames and throughput (KB/s) on channel 6, the most congested

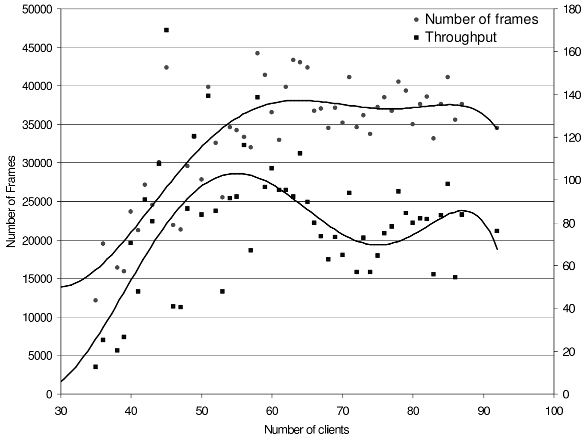

Fig. 8. Throughput and num. of frames vs. num. of clients on channel 6

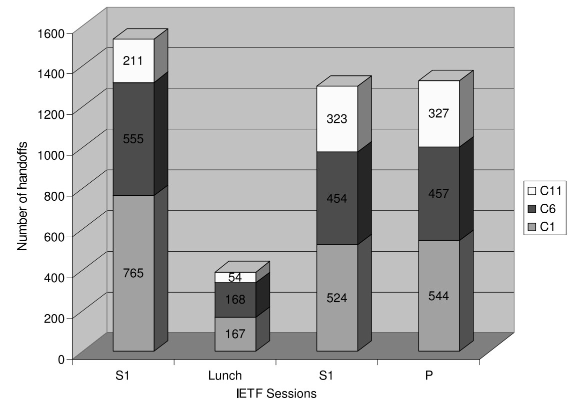

Fig. 9. Total number of handoffs

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. In our measurements, this certain value is about 55 clients (Fig. 8) and it represents the maximum number of clients the channel can handle, that is, the maximum capacity. Once the number of clients exceeds capacity, collisions and retries increase bringing down the overall throughput.

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.

V. Channels in IEEE 802.11 Networks

As mentioned earlier, at the IETF meeting in Dallas, IEEE 802.11b and IEEE 802.11a networks were deployed. The 802.11b standard works in the 2.4 GHz band and offers speeds up to 11 Mb/s while the 802.11a standard works in the 5 GHz band and offers speeds up to 54 Mb/s. One of the most significant differences between the two technologies is the number of non-overlapping channels. In 802.11b there are three non-overlapping channels, namely 1, 6 and 11, while in 802.11a there are twelve non-overlapping channels. Two APs that use two non-overlapping channels do not interfere with each other. In other words, one AP cannot "hear" the other. The number of non-overlapping channels becomes a critical factor in very crowded environments where the large number of clients requires a large number of APs to be closely deployed in order to guarantee a minimum throughput and continuous coverage. Usually, this is not a problem for 802.11a since its twelve non-overlapping channels are more than enough to cover large and small crowded areas. However, this is not the case for 802.11b with its three non-overlapping channels. When using 802.11b in crowded environments, we can either re-use the three non-overlapping channels more than once or we can just decide to use any of the available channels. At the IETF meeting, the network administrators decided to go with the first option and deployed multiple adjacent APs on the same non-overlapping channels, that is, only channels 1, 6 and 11 were used.

In general, deploying different APs on the same channel, can cause interference and degradation of the link quality. Furthermore, all clients will contend access to one channel including all the APs on that channel, hence creating more congestion and introducing higher probability for collisions. This problem becomes even more critical if the APs using the same channel cover roughly the same area - that is, the APs' coverage areas significantly overlap. This last scenario was the one deployed at the IETF meeting in Dallas where, in order to give access to a very large number of clients in a very confined space, multiple adjacent APs using the same channel covered the same space.

VI. 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.

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. Furthermore, Auto Rate Fall-back (ARF), or any other equivalent mechanism, can also contribute to increased congestion by lowering the data-rate when a certain amount of packet loss is experienced by the client [1]. By lowering the data-rate, packets occupy the medium longer, preventing other stations from sending their packets. At the IETF meeting, network administrators fixed the 802.11b APs' bit-rate to 11 Mb/s in order to avoid this last problem. One exception

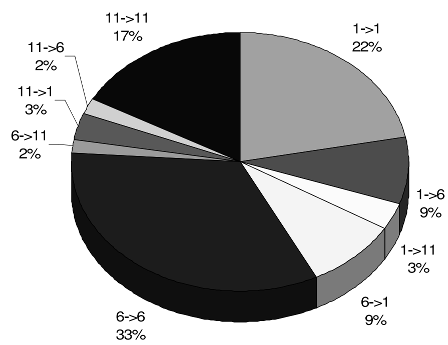

Fig. 10. Handoffs between channels

Session 1 (AM)	30.5%
Lunch	33.0%
Session 1 (PM)	30.2%
Plenary	54.7%

TABLE II

PERCENTAGE OF HANDOFFS PERFORMED TO THE SAME AP

were the hotel APs on channel 6 which supported all the default data rates since the IETF network administrators had no control over them.

Fig. 9 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 (see Fig. 5). However, as we can see in Fig. 10, 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 same AP - that is, to the same AP the client just disconnected from. Table II shows the percentage of handoffs in which current AP and next AP are the same.

Performing a handoff to the same 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. 10 and Table III. The session time shown in Table III 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 III, 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 III also shows the percentage of clients that performed a handoff within the specified times.

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%

TABLE III

SESSION TIME: TIME BETWEEN HANDOFFS

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%

TABLE IV

PERCENTAGE OF HANDOFFS WITHIN 1 MINUTE AND WITHIN 5 MINUTES

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 transmitted at the lowest available bit-rate, thus keeping the medium busy for longer and preventing other stations from accessing the medium. Because of this, unnecessary handoffs degrade network performance by increasing network congestion for all the clients on a particular channel. In particular, 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. At the application level, the overhead introduced by the handoff becomes even more evident since the Operating System (OS) introduces its own overhead. This last point is further analyzed in Section VIII.

From the previous results we can see how today there are many problems with the way MNs select the AP to

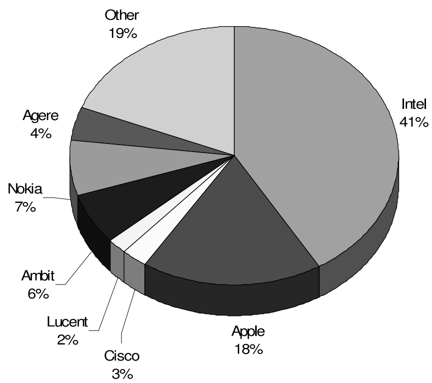

Fig. 11. Presence of wireless card vendors

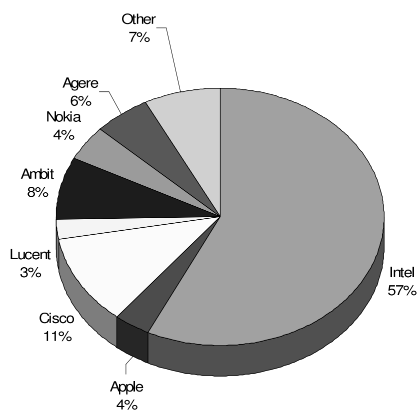

Fig. 12. Percentage of handoffs per wireless card vendor

connect to. In particular, the AP is selected according to the link signal strength and SNR levels. Other factors such as effective throughput, number of retries, number of collisions, packet loss, bit-rate or BER are ignored. When the MN needs to perform a handoff, it has to look for a different AP to connect to. Unfortunately, with a very high probability, the MN will pick the same AP it was connected to because its link signal strength and SNR are still 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 [1].

B. Vendors and Handoff

There were about 1200 attendees with cards from many different vendors, dominated by Intel wireless cards (see Fig. 11). Most of the different vendors had similar handoff policies and algorithms as they behaved pretty much in the same way (Fig. 12), 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 IV, 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 an optimized handoff algorithm.

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.

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 V

DISTRIBUTION OF BROADCAST TRAFFIC

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

TABLE VI

PERCENTAGE OF BROADCAST TRAFFIC

VII. SAME CHANNEL VS. MULTIPLE CHANNELS

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.

A. Broadcast and Multicast Traffic

Broadcast and multicast traffic represent 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 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,

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

Channel	IPv6 Multicast	IPv4 Multicast
1	0.7%	1.8%
6	2.0%	3.4%
11	0.8%	1.6%

TABLE VII

PERCENTAGE OF MULTICAST TRAFFIC

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 V 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 VI 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 VI, 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 VII 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.

B. Interference

Having adjacent APs on the same channel introduces co-channel interference. This means higher BER, hence higher packet loss and number of retries. Overall, this translates to a lower throughput. Of the three 802.11b channels used at the IETF meeting, channel 6 was the one with the largest number of APs. This means that clients on channel 6 experienced the highest co-channel interference. Fig. 7 shows the average

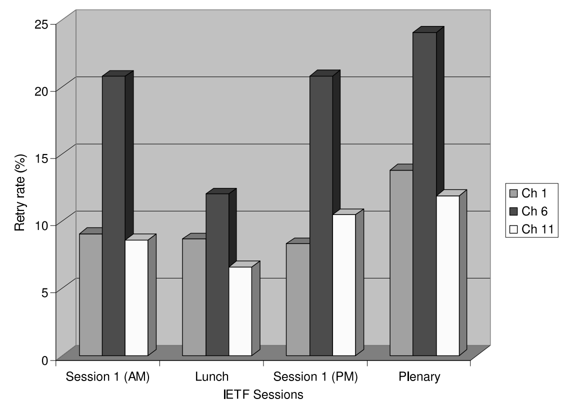

Fig. 13. Retry rate in IEEE 802.11b

throughput per client on each of the three channels. As we can see, on average, clients on channel 11 experienced the highest throughput while clients on channel 6 experienced the lowest throughput. Fig. 13 shows the average retry rate. Clients on channel 6 experienced the worst channel condition with the highest retry rate.

As we have discussed in Section IV, Fig. 5 shows the number of clients per channel. Channel 6 was the most congested channel followed by channel 1 and 11. This high congestion together with co-channel interference, explains the drop in throughput and peak in retry rate experienced by users on channel 6. In particular, during the plenary, the number of users on channel 1 and channel 6 is almost the same. However, the number of retries on channel 6 is significantly higher than the one on channel 1 while the throughput on channel 6 is significantly lower than the throughput on channel 1. The reason for this difference in throughput between channel 1 and channel 6 is the highest degree of co-channel interference on channel 6.

VIII. A CASE STUDY: WINDOWS XP

Having frequent handoffs can lead to a very unsatisfactory user experience. In order to better understand the impact that frequent handoffs have on the user, we analyzed how a user perceives a handoff at the application layer. We used an IBM Thinkpad T42 laptop running Windows XP Professional with Service Pack 2. The laptop was equipped with an Intel Centrino chipset, the embedded wireless card was an Intel(R)PRO/Wireless 2200BG card. We focused on general handoff behaviors.

As general behavior, Windows XP performs a pure L2 handoff if the old and new AP have the same ESSID. If the two ESSIDs mismatch, the OS assumes a L3 handoff. This means that every time the client performs a handoff between two APs with different ESSID, the OS will trigger the DHCP procedure in order to acquire a new IP address based on the assumption that the subnet has changed. This ESSID-based policy introduces significant delay in the handoff if a change in ESSID does not result in a change of subnet. At the same time, such a policy significantly penalizes the user in situations where a change in subnet is not followed by a change in ESSID.

The Intel wireless card in Windows XP has a configuration parameter called ”Roaming Aggressiveness”. From the description of this parameter in Windows XP we read: ”This setting allows you to define how aggressively your wireless client roams to improve connection to an access point. [...]”. In our experiments this parameter was set at the default value of 50%, but still causing unnecessary handoffs as the OS always tried to be connected to the best possible AP all the time, even though the connection to the current AP was marked as ”excellent”. This results in the wireless card scanning the medium almost every minute and connecting to the new best AP even if this triggers a L3 handoff which introduces additional superfluous delay. From a user perspective this is a very undesirable behavior as it causes unnecessary disruptions. From a network perspective, this causes unnecessary traffic that can increase congestion.

A L2 handoff is not much disruptive to applications such as SSH - Secure Shell, Real VnC and PuTTy. There are network connectivity disruptions but the OS does not close the sockets in use. On the other hand, a L3 handoff causes the OS to close all sockets in use, thus terminating any open session for every application. This is another reason why Roaming Aggressiveness can lead to a bad user experience when triggering unnecessary L3 handoffs.

Usually, the handoff can be triggered by either low SNR or high packet loss. We conducted the experiments by switching off the AP to which the laptop was connected to, in order to simulate 100% packet loss. Interestingly enough, the OS detected the loss of connection 2.5 seconds after the last received beacon frame. Furthermore, the average handoff time, measured from the disassociation frame to the association request frame, was on the order of 1.2 seconds on average. Part of this delay was due to the wireless card scanning the same channel multiple times without any apparent reason.

IX. Effect of Obstacles : Screen Effect

Fig. 6 and Fig. 7 show that AP3 is significantly underutilized throughout all the sessions. The number of clients connected to AP3 is much lower than the number of clients using AP2 even though physically the two APs are close to each other and cover the same area, that is, the same clients. In Fig. 2, we can see that AP3 was installed behind a projection screen. In general, these kinds of screens contain a significant amount of metal to make the projected image brighter and higher in contrast. Because of this the screen represented a significant obstacle for AP3 whose signal was severely attenuated. The kind and amount of attenuation introduced by the screen depends on the materials used in building it. This attenuation introduced by the screen was responsible for less clients being able to ”find” AP3. Furthermore, the clients that successfully associated with AP3 experienced a large number of retransmissions. The retry rate for AP3 was almost twice the retry rate of the other APs.

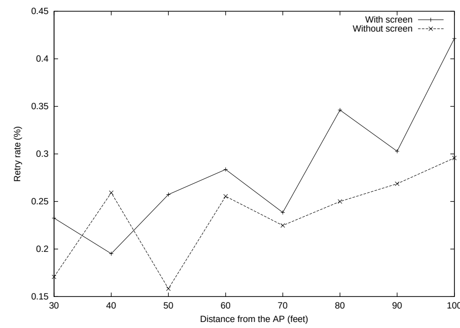

Fig. 14. The retry rate with the AP behind the screen

In order to verify the effect of the screen on wireless networks, we performed some experiments. We performed the experiments on the 7th floor of the Schapiro building in the Columbia University campus. We used a Netgear WG602 v3 access point, one IBM T42 laptop equipped with an Intel Centrino chipset using an Intel(R)PRO/Wireless 2200BG card as wireless client and two other IBM T42 laptops with the same specifications as wireless sniffers. The sniffers used a Proxim ORINOCO 11a/b/g combo card and Airopeek as sniffing software.

For our experiments, we positioned the AP at one end of a long hallway and setup a projection screen in front of it. One sniffer was placed close to the AP so to capture all frames sent and received by the AP and the other sniffer was placed close to the client so to capture all frames sent and received by the client. While associated to the AP, the client transmitted 100 ICMP echo request packets, one every second. The sniffers would capture the requests and the responses on both the AP side and the client side. We performed this measurements at different distances from the AP, from 30 to 100 feet. The same measurements were taken with and without the screen. The captured data was later analyzed to calculate the packet loss rate due to the screen.

Fig. 14 shows the results of our experiments. We can see that in both cases, with and without the screen, the retry rate increases with the increasing of the distance. However, the retry rate with the screen is always higher.

The screen clearly introduces some interference. The amount of this interference is, however, hard to estimate since it depends from the materials used in the screen. Regardless, the effect of this kind of obstacles should be considered in the deployment of a wireless network.

X. 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. 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 behaved 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. Clients experienced a high level of interference and congestion. Interference was mainly caused by co-channel interference while congestion was caused by having multiple APs and correspondent clients contending for the same channel. Having multiple APs also caused broadcast and multicast packets to be duplicated at each AP, thus wasting bandwidth and contributing to the high level of congestion. In the worst case, redundant broadcast packets were responsible for 6% of the network traffic and multicast packets were responsible for 3% of the traffic.

We also observed an uneven distribution of clients and throughput between channels and between APs on the same channel. This was caused by the absence of a load balancing algorithm. We found a clear correlation between the number of clients and network utilization among channels which would suggests 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.

To conclude, current handoff algorithms have proven to be inadequate and, in some cases, counterproductive. Better handoff algorithms and automated adaptation mechanisms are required in highly congested networks. Better design tools can also assist in the planning and deployment of wireless networks in order to minimize misconfiguration and maximize performance. Solutions to many of the problems presented in this paper are reserved for future study.

References

[1] 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: Proceeding of the 2005 ACM SIGCOMM workshop on experimental approaches to wireless network design and analysis*. New York, NY, USA: ACM Press, 2005, pp. 11–16.

[2] 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: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*. New York, NY, USA: ACM Press, 2005, pp. 5–10.

[3] 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: Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*. New York, NY, USA: ACM Press, 2002, pp. 195–205.

[4] D. Kotz and K. Essien, "Analysis of a campus-wide wireless network," in *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*. 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: Proceedings of the 10th annual international conference on Mobile computing and networking*. 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 *the Proceedings of ICC'04: IEEE International Conference on Communications*, June 2004, pp. 3844–3848.