An Empirical Analysis of Heterogeneity in IEEE 802.11 MAC Protocol Implementations and its Implications

K N Gopinath, Pravin Bhagwat, AirTight Networks, India
K. Gopinath, CSA Department, IISc, India
{gopinath.kn, pravin.bhagwat}@airtightnetworks.net, gopi@csa.iisc.ernet.in

Abstract: IEEE 802.11 standard specifies a protocol for data link (i.e., medium access control or MAC) level communication in a wireless local area network (WLAN). The WLAN market consists of various products (e.g., access points (APs), client adapters) that have been released by multiple vendors. Certain third party certifications such as those specified by the Wi-Fi alliance have been widely used by vendors to ensure interoperability and basic conformance to the 802.11 standard, thus leading to the expectation that such certified devices exhibit identical MAC level behavior. In this paper, however, we present what we believe to be the first ever set of experimental results that highlight the fact that WLAN devices from different vendors in the market can have heterogeneous MAC level behavior.

Specifically, we demonstrate with examples and data that in certain cases, devices may not be conformant with the 802.11 standard (e.g., due to incorrect implementations), while in other cases, they may differ in details that are not a part of mandatory specifications of the standard, but still important for WLAN operations. We argue that heterogeneous MAC implementations can adversely impact WLAN operations leading to unfair bandwidth allocation, poor network utilization, potential break-down of related MAC functionality and difficulties in provisioning the capacity of a WLAN. However, on the positive side, MAC level heterogeneity can be useful in applications such as vendor/model level device fingerprinting.

1. Introduction

Recently, there has been a proliferation of wireless LAN (WLAN) devices such as access points (APs) and client adapters in the market. The devices are from numerous vendors; to name a few - Cisco Systems Inc. [1], D-link Systems Inc. [2], Intel Corporation [3], Netgear Inc. [4], 3Com Inc. [5], Proxim Inc. [6] etc. WLAN devices implement the IEEE 802.11 standard [7] which specifies a protocol for data link (i.e., medium access control or MAC) level communication in a wireless LAN (WLAN). The 802.11 standard specifies frame formats for MAC level communication and mechanisms for medium access control amongst devices in a WLAN. Further, it specifies a protocol for connection management, data transfer and power management between AP and its clients. Certain third party certifications such as those specified by the "Wi-Fi Alliance" [8] have been widely used by WLAN vendors for interoperability and, also, for basic conformance. Wi-Fi alliance has certified more than 1,500 Wi-Fi products since April 2000 [9] leading to the expectation that such certified devices exhibit identical MAC level behavior. However, the primary mission of the Wi-Fi alliance is to assure a positive user experience through product interoperability [9]. Hence, it is to be noted that although Wi-Fi certification can automatically assure a basic level of conformance to the 802.11 standard, it will not ensure a detailed conformance (e.g., such as those specified by Protocol Implementation Conformance Statement (PICS) [7]).

In this paper, we present experimental results (based on packet trace analysis) that highlight the fact that WLAN devices from different vendors in the market can have heterogeneous MAC level behavior. Specifically, we demonstrate with examples and data that in certain cases, devices may not be conformant with the 802.11 standard (e.g., due to incorrect implementations), while in other cases, they may differ in details that are not a part of mandatory specifications of the standard, but still important for WLAN operations. We argue that heterogeneous MAC implementations can adversely impact WLAN operations leading to unfair bandwidth allocation, poor network utilization, potential break-down of related MAC level functionality and difficulties in provisioning a WLAN. However, on the positive side, MAC level heterogeneity can be useful in other applications such as vendor/model level device fingerprinting.

Certain aspects discussed in this paper such as those related to fairness of bandwidth allocation in an IEEE 802.11 network have been considered in the literature earlier [10, 11]. However, such prior work has modeled all stations in a network to be *identical*. To the best of our knowledge, we believe that our paper presents the first-ever set of experimental results on MAC level heterogeneity amongst devices and its implications on WLAN behavior.

The rest of the paper is organized as follows. Section 2 provides experimental results that highlight the fact that WLAN devices from different vendors in the market can have heterogeneous MAC level behavior.

Section 3 discusses the adverse impact of MAC level heterogeneity on WLAN operations. Section 4 discusses MAC level vendor/model level device fingerprinting, which can be facilitated by heterogeneity in MAC implementations. Section 5 concludes the paper.

2. MAC Level Heterogeneity in WiFi devices

In this section, we present data to demonstrate the heterogeneous behavior of MAC implementations of commonly used Wi-Fi devices. Table 1 and Table 2 summarize devices used in our various experiments. We will start with heterogeneity that is related to conformance aspects of the IEEE 802.11 standard.

2.1 Conformance related heterogeneity

Our experiments indicate that certain devices may not be conformant to parts of the 802.11 specification (e.g., due to incorrect implementations). Hence, they can have heterogeneous MAC behavior with respect to other devices that have a standard conformant implementation.

Device	Manufacturer
Cisco 350	Cisco
Cisco CB21AG	Cisco
Linksys WPC55AG	Cisco
Sparklan WL-360F	Sparklan
Linksys WPC11	Linksys
Centrino	Intel
Orinoco Silver	Lucent

Table 1 : Client adapters/cards

Device	Manufacturer
Cisco 350	Cisco
Cisco 1100	Cisco
Linksys WAP55AG	Cisco
Dlink DWLG730	D-link
Netgear WGR101	Netgear
Asante AP (G Mode)	Asante
Belkin AP (B only)	Belkin
Software AP	HostAP (Linksys)

Table 2 : Access Points (APs)

2.1.1 Random Back-off

IEEE 802.11 standard specifies a CSMA/CA protocol in which a station (client or an AP) should sense the medium before attempting to transmit a packet. Specifically, the standard recommends that, before transmitting a packet, a station should wait for an interval corresponding to Distributed Inter Frame Space (DIFS) and ensure that there is no other transmission on the medium. Further, once it detects that the medium is idle for DIFS interval, the standard recommends the use of a random back-off mechanism (i.e., each station is required to wait for a random amount of time) before transmitting a frame. This backoff mechanism is recommended to avoid collisions between multiple stations.

The random backoff interval is calculated as follows. A station should wait a random backoff time (*Backoff*) as specified by the following equation (1):

*Backoff = Random() * aSlotTime* ([7], page 75) (1)

where *Random*() function generates an integer (called "slot") drawn from a *uniform* distribution over a specified interval called "contention window"; *aSlotTime* is a parameter specific to the physical layer [7]. For example, with DSSS physical layer *aSlotTime* is 20 microseconds (usecs). Further, the contention window used for random number generation is between zero and 31 (assuming that there are no MAC level collisions).

We now describe calculation of the backoff interval of a device in our experiments. The device of interest generates UDP traffic under saturation conditions (i.e., condition in which a station always has a data backlog for transmission) on an isolated radio channel which does not have any device other than the client (and the associated AP). A Prism chipset [12] based sniffer is used to capture packet level trace of the traffic. Each packet in the trace will have a microsecond timestamp called "mactime" associated with it. This timestamp is inserted by the card firmware when it starts receiving the physical layer convergence procedure (PLCP) preamble [8] associated with the frame (Note: The transmission (and reception) of a MAC level frame is always preceded by PLCP related information which help in transmitting the MAC frame over the corresponding physical layer). Let us assume that "mactime1" is the timestamp associated with the reception of MAC level acknowledgement (ACK) frame corresponding to a MAC level data packet associated with the station (Figure 1). Let us assume the time for transmission of PLCP information is P usecs and time to transmit MAC level ACK is t usecs. We calculate the backoff interval used by a card for the *next* data packet (with timestamp "mactime2") as follows:

ACK_Finish_Time = mactime1 + P + t (2)
Backoff = mactime2 – (ACK_Finish_Time + DIFS) (3)

Backoff_slot (*Backoff* in slots) is calculated using the physical layer specific slot time (*aSlotTime*) as:

Backoff_slot = Backoff Interval /aSlotTime (4)

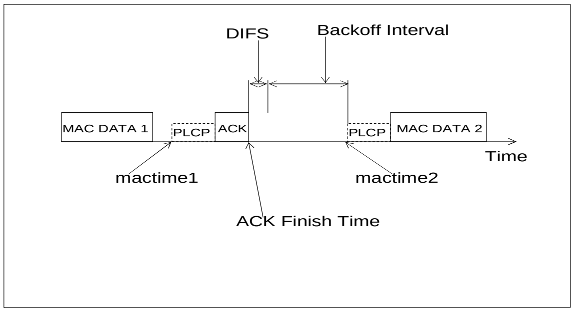

Figure 1: Illustration of backoff interval calculation

Figure 2 represents heterogeneity in the random backoff procedure used by two popular Wi-Fi devices: Cisco 350 series client card and Linksys WPC11 client card. As can be seen, the backoff interval of Cisco 350 series card is skewed towards the bottom of the interval (zero to 10 slots) where as the backoff interval of Linksys WPC11 card is distributed fairly uniformly in the entire interval**.** It should be noted that a few backoff values that are higher than 31 slots are due to MAC level retries (e.g., after a MAC level collision).

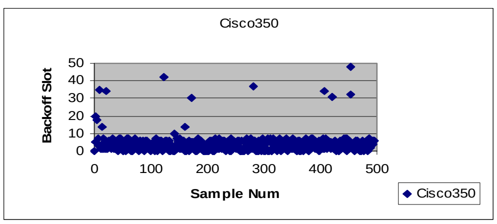

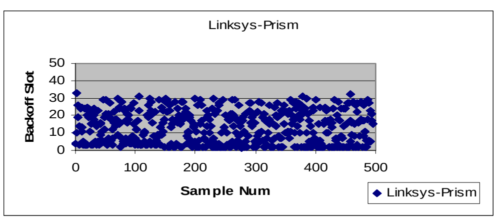

Figure 2 : Heterogeneity in random backoff behavior

2.1.2 Virtual Carrier Sensing

IEEE 802.11 devices implement two types of carrier sensing: physical and virtual. In the former, a device senses the wireless medium to detect if a transmission is currently occurring, and if so, waits for the transmission to complete before attempting to transmit. The second mechanism is based on 'duration' field in the frames. The value in this field can be used by a wireless station (transmitter) to reserve the medium for a specified amount of time (not exceeding 32767 microseconds) for communication with a receiver station. Any other station that receives and decodes this frame refrains from transmitting for a time interval computed based on the value in the 'duration' field in the frame. This way the communication between the transmitter and the receiver can happen without the risk of a collision during the reserved time period.

AP	Cisco 350	D-link DWL G730	Cisco 1100	HostAP (Linksys WPC11)
Honors NAV?	Yes	Yes	Yes	No

Table 3(a): Virtual carrier sensing behavior of APs

Client Card	Linksys WPC55AG	Linksys WPC11	Lucent Orinoco	Intel Centrino
Honors NAV?	Yes	No	No	Yes

Table 3(b): Virtual carrier sensing behavior of Clients

Table 3: Virtual carrier sensing

We now describe the procedure used to determine if an AP honors the duration field in a received packet. We generate a sequence of specially crafted MAC level data packets with the maximum possible duration value (i.e., 32767 microseconds) on the channel in which an AP is operating. We generate such MAC packets using a locally written packet injection tool. An AP honors the duration field if it stops transmitting beacons in response to the generated MAC packets with large duration value. As can be seen from Table 3 (a), Cisco 350 series AP, D-link DWL G730 AP and Cisco 1100 AP honor the duration field in a received packet. However, Host AP based on Linksys WPC11 card does not honor the duration field.

To determine if a client honors the duration field in a packet, we associate the client with an AP and generate certain traffic (e.g., ping traffic from client to AP). Then, using a locally written packet injector, we generate a sequence of artificially crafted MAC level data packets with the maximum possible value (i.e., 32767 microseconds) to see if the client device halts its communication. It should be noted that the AP used in the experiments should not honor duration field in a received packet to unambiguously determine the client's virtual carrier sensing behavior. Hence, we use HostAP based on Linksys WPC11 card for this purpose. As can be seen from Table 3 (b), some of the older generation of client cards such as Linksys WPC11, Lucent Orinoco Silver cards do not honor duration field in a received packet. Newer generation of cards such as Linksys WPC55AG and Intel Centrino honor the duration field in a received packet.

2.1.3 Calculation of Duration field

As mentioned in the previous section, "duration" field in an 802.11 packet is used to reserve the wireless medium. Devices use the duration field to reserve the medium for transmission of MAC level acknowledgement frame (ACK) for any unicast frame. The exact value of the duration field used for reservation of ACK depends on the time required for transmission of the ACK packet, which in turn depends on the rate at which the corresponding unicast frame was transmitted ([7], page 95).

Client Card	Cisco 350	Cisco CB21 AG	Linksys WPC11	Sparklan PCI
Duration Field (usecs)	258	314	64808	62443

Table 4 : Duration Field

However, it should be noted that the value in the duration field of a packet should be less than the maximum value of 32767 as mentioned earlier. As can be seen from Table 4, certain client cards include values greater than 32767 (e.g., Sparklan PCI card, Linksys WPC11 card) in the duration field. Further, we have observed that recipients do not honor such non-conformant values in the duration field of a packet.

2.1.4 Power Management

The 802.11 standard specifies mechanisms for power management between a client and its associated AP. For example, a client can indicate to the AP that it would like to enter power-save (PS) mode (e.g., to conserve battery). Consequently, the AP starts buffering packets that are destined to the client. An AP indicates that it has data stored for a client using *Traffic Indication Map* (TIM) element that is a part of beacons ([7], page 128). A station that operates in PS mode wakes up periodically to listen for beacons and interprets the TIM element in the beacon. Whenever the client determines that the AP has data buffered for it, the client enters the "Active" mode. Further, it should transmit a "PS Poll" frame to the AP to request the buffered data from the AP. It should be noted that the 802.11 standard specifically states that a client should indicate any change in its power management mode to the AP via a successful frame exchange sequence only ([7], page 129). That is, it should *not* indicate power-save mode change using a single frame exchange sequence (e.g., broadcast data/management packet, MAC level ACK packet). The motivation behind the above restriction in the standard seems to be that a client should get a confirmation (e.g., in the form a MAC level ACK) from the AP whenever it is changing its power-save mode (e.g., from active to power-save). Table 5 illustrates the power management behavior of four popular client cards. Specifically, it indicates how Linksys WPC11, Cisco CB21AG, Cisco 350 and Centrino cards indicate their power mode change from active to power-save. As can be noticed in Table 5, certain client cards such as Cisco 350 do use control packets to indicate change in power-save status and hence are non-conformant to the power management specification of the 802.11 standard.

Client Card	Linksys WPC11	Cisco CB21AG	Cisco 350	Centrino
Power save mode change indication	Data NULL	Data NULL	ACK	Data NULL

Table 5 : Power Management behavior

2.2 Heterogeneity related to non-mandatory portions of the 802.11 standard

In this section, we present data to demonstrate MAC level heterogeneous behavior that arises due to differences in implementation of non-mandatory (but, nevertheless important) portions of the 802.11 standard.

2.2.1 Reassociation Latency

The 802.11 standard specifies that a client device should perform a management packet hand-shake before utilizing the services of an AP (e.g., transferring data through the AP). The handshake involves exchange of the following packets in sequence: probe request and probe response (for discovery of AP); authentication request and authentication response (for MAC level authentication); and finally, (re)association request and (re)association response (for creating an association or binding state for the client at the AP) ([7], page 22).

Let us define *reassociation latency* as the time interval required for the connection hand-shake between an AP and a client which is in an "unauthenticated" state ([7], page 22). A client can be forced into unauthenticated state using special MAC level packets such as "deauthentication" packets. Figure 3 illustrates the concept of reassociation latency associated with a client and an AP. As can be seen from the figure, reassociation latency consists of a probe phase (in which a client discovers/scans for APs on one or more channels), authentication phase (in which a client authenticates with an AP at MAC

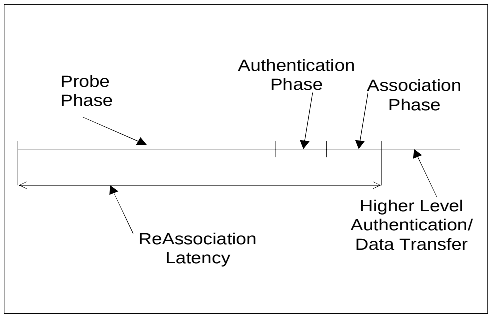

Figure 3 : Reassociation Latency.

level) and an association phase (in which a client performs MAC level association procedure).

Figure 4 illustrates the reassociation latency distribution of four popular client cards measured on a clean channel with no other interfering devices. The data was obtained by measuring the time a card requires to complete the connection handshake starting from an unauthenticated state. The cards were forced to enter the unauthenticated state by transmitting MAC level deauthentication packets. Figure 4 demonstrates the heterogeneity associated with reassociation latency of the cards – Linksys WPC55AG card has a modal value of 1250ms, Cisco 350 series client has a modal value of 450 ms, Cisco CB21AG has a modal value of 2450 ms and Linksys WPC11 has a modal value of 1450 ms. Preliminary analysis of packet traces indicates that the differences in the reassociation latency is primarily due to different heuristics used by cards in the probe phase of MAC level connection handshake. For example, after receiving a deauthentication request, certain cards (e.g., Cisco CB21AG) scan for APs on all channels where as other clients such as the Cisco 350 series card use certain optimizations (e.g., scan the channel on which the card was previously associated before scanning other channels).

2.2.2 Packet Contents

The 802.11 standard allows certain vendor specific extensions to be included in management frames (e.g., beacons, association request, association response). For example, reserved tags can be used in beacon frames to transmit vendor specific information. Vendors can use these tags for exchanging proprietary information between their own APs and clients to achieve some competitive differentiation (e.g., better throughput, load balancing etc.). Thus, devices from different vendors can exhibit such packet content related heterogeneity. Table 6 shows that APs from different vendors behave differently with respect to reserved tags. For example, Cisco 350 series AP contains a reserved tag 133 in its beacon, Cisco 1100 series AP transmits two reserved tags (133, 150) in its beacons, where as, Netgear AP WGR101 and software AP (e.g., Prism chipset based HostAP) do not contain any reserved tags.

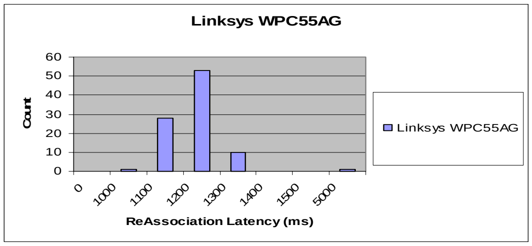

Figure 4(a): Linksys WPC55AG (mode: 1250ms)

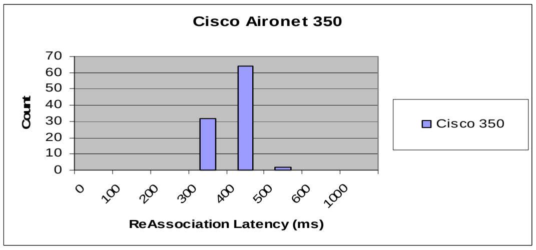

Figure 4(b): Cisco 350 (mode: 450ms)

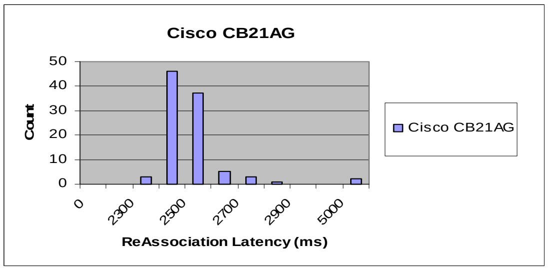

Figure 4(c): Cisco CB21AG (mode: 2450ms)

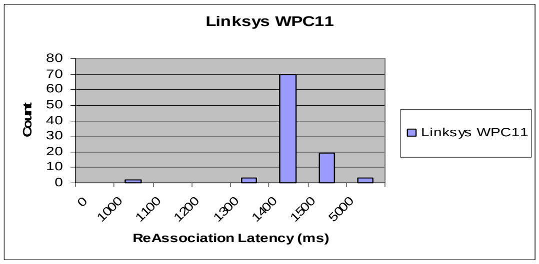

Figure 4(d): Linksys WPC11 (mode: 1450ms)

Figure 4: Heterogeneity associated with Reassociation latency of popular client cards.

AP	Cisco 350	Cisco 1100	Netgear WGR 101	Software AP
Reserved Tag	133	133,150	None	None

Table 6 : Reserved Tags in Management packets

2.2.3 Multiple Rate Support

Physical layer (PHY) schemes that are recommended by the 802.11 standard (e.g., Direct Sequence Spread Spectrum (DSSS)) support multiple data transfer rate capabilities (e.g., 1 Mbps, 2 Mbps etc.). Each of the data rates is realized by a corresponding modulation scheme at the hardware level (e.g., 1Mbps with Differential Binary Phase Shift Keying (DBPSK), 2 Mbps using Differential Quadrature Phase Shift Keying (DQPSK) etc.). Further, each of the modulation schemes exhibit varying levels of robustness to channel conditions (e.g., signal-to-noise ratio (SNR)). For example, a modulation scheme that can transmit at higher rates is relatively more susceptible to channel conditions. This allows implementations to perform dynamic rate switching with the objective of improving performance ([7], page 95). Specifically, stations can dynamically switch the rate of transmission of a data or management frame that has unicast destination (Note: control frames and broadcast management frames are sent at static rates as specified by the 802.11 standard). That is, an 802.11 device can dynamically decrease the rate of packet transmission to take care of poor channel conditions. Similarly, it can switch to a higher rate when the channel conditions improve. Further, there can be significant performance differences depending on the exact rate adaptation algorithm used by a WLAN device [13]. However, the standard does not mandate any specific heuristics for dynamically switching the transmission rates.

We now describe the procedure used to study the rate switching behavior of client cards. On a clean channel that is free of interference from other devices, we associate multiple clients of the same vendor with an AP. Further, we generate a UDP stream (under saturation conditions) from each of the clients to a wired host connected to the AP. A packet trace is collected to find out the percentage of packets that were transmitted at each of the data rates. Figure 5 presents the above data collected with 4 client cards.

Specifically, Figure 5 (a) shows the percentage of packets that are transmitted at each of the supported rates (1, 2, 5.5 and 11Mbps) when Linksys WPC11 cards are used. Figure 5 (b) and Figure 5 (c) show similar data for Lucent Orinoco Silver card and Cisco 350 card respectively. We have conducted the measurements under identical channel conditions. As can be clearly noticed, each card has demonstrated a unique rate switching behavior. Trace analysis indicates that Linksys WPC11 implements a conservative rate switching algorithm where in the card requires a large (i.e., 10) number of successful data packet transmissions before it switches to a higher rate. On the contrary, Cisco 350 series card implements an aggressive rate switching algorithm and switches back to higher rate after few (one of two) successful data packet transmissions.

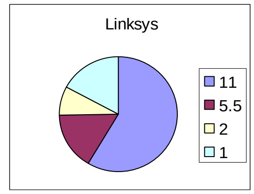

Figure 5(a): Linksys WPC11 card - significant number of packets at lower rates

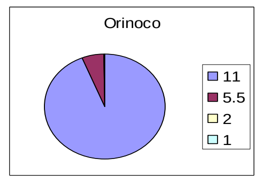

Figure 5(b): Lucent Orinoco Silver card - most of the packets at 11Mbps and 5.5 Mbps rates

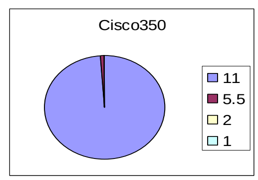

Figure 5(c): Cisco 350 card - most of the packets at the possible highest rate

Figure 5: Multirate support – data transfer rate distribution when 4 cards of each vendor transmit UDP packets (under identical conditions).

AP	Cisco 1100	Asante	Belkin	HostAP (Linksys)
Beacon Frame Rate	1	2	2	1

Table 7 : Beacon Frame Rate (Mbps)

2.2.4 Management Packet Rate

The 802.11 standard specifies that management packets with broadcast destination address (e.g., beacons) should be transmitted at one of the basic rates supported by an AP and its associated clients ([7], page 95). Table 7 indicates that APs from different vendors can transmit beacon frames at different rates (even though all the APs mentioned above were configured identically).

3. Adverse Implications of MAC Heterogeneity on WLAN Operations

In this section, we illustrate with examples that MAC heterogeneity amongst WiFi devices can have several adverse implications on operations of a WLAN.

3.1 Unfair Bandwidth Allocation

MAC level heterogeneity can result in an unfair bandwidth allocation in a WLAN. Figure 6 demonstrates the bandwidth share of a Cisco 350 client and a competing Linksys WPC11 client card (both of which are associated with a Cisco 350 AP). UDP traffic was generated from the clients to a server on the wired side under saturation conditions (through the AP). The figure demonstrates unfair bandwidth allocated to a Cisco 350 client card. This can be explained as follows. First, if a client does not use the standard specified uniform random number generator for backoff, it can access the medium much more frequently than other clients which use a uniform random backoff mechanism. This will contribute to better throughput of Cisco 350 client, which uses a biased random backoff generator (as explained in previous section). Second, Cisco 350 client implements an aggressive rate switching algorithm when compared to Linksys WPC11 (i.e., it switches back to higher transmit rates after fewer successful transmissions). This further contributes to the better throughput of Cisco 350 card, thus, leading to an unfair allocation of bandwidth. Analyzing the individual contributions of each of the above two factors towards the unfairness is one of our immediate research activities.

3.2 Poor Network Utilization

Virtual carrier sensing (coupled with the RTS/CTS mechanism [7]) can be used to alleviate certain network performance issues caused by hidden nodes. A node B is hidden with respect to another node A if B is within the range of node A, but, outside the range of other nodes (e.g., C) communicating with node A. In such a case, transmissions from C and B can occur simultaneously and result in a collision at node A. To avoid such a scenario, node A can “reserve” the medium for a node it will be communicating (e.g., C) by using the duration field in a RTS/CTS frame exchange. However, if a node does not honor duration field in a packet or transmits a non-standard value in the duration field (which the other standard conformant receivers will ignore), it can result in collisions due to hidden nodes and hence, result in poor network utilization. However, it should be noted that RTS/CTS mechanism itself is known not to be efficient in mitigating the hidden node problem. Heterogeneous and non-conformant implementations aggravate the problem further.

3.3 Potential Breakdown of MAC Functionality

If a particular device does not behave as specified in the standard, there can be a potential breakdown in the related MAC level functionality. For example, in the previous section, we have explained that certain clients exhibit non-conformant behavior with respect to power management functionality. Specifically, we have shown that Cisco 350 client uses ACK to indicate change in power-save status. This can be ignored by an AP as this is not a standard specified behavior.

3.4 Difficulties in Provisioning the Throughput Capacity of a WLAN

There is a need to provision the throughput capacity of a WLAN (e.g., determining how many APs need to be used, deciding how many clients can be sustained by an AP, analyzing the throughput received by a client etc.) before the actual deployment. Several theoretical models have been proposed to predict the throughput performance of a WLAN [14, 15]. One can expect that the previously mentioned models can be used to provision a WLAN before the actual deployment. Such models, however, assume all devices to be identical. Hence, with MAC level heterogeneity, the throughput actually realized by a device can be much different from that of the predicted throughput. We illustrate this below.

Figure 7 compares measured throughput of a WLAN with that of the throughput predicted using a locally written packet-by-packet MAC level simulation model based on [14] (which assumes that all devices are identical). Measurements were conducted under identical conditions and up to 4 client cards of the same vendor were used in each trial. Specifically, Cisco 350 series AP and Cisco 350 cards, Linksys WPC-11 cards were used. As can be clearly seen from the figure, even with only 4 clients, the model under-predicts the throughput of Cisco 350 client by 9% and over-predicts the throughput of a Linksys WPC-11 client by 144%. Thus, due to MAC heterogeneity, simple models may not be sufficient to accurately predict the throughput capacity of a WLAN.

4 Advantages of MAC Level Heterogeneity: Device Vendor/Model Level Fingerprinting

MAC level heterogeneity can be used for the purpose of “device fingerprinting”, i.e., uniquely identifying vendor/model of an 802.11 device by observing its traffic. Some examples of possible signatures are as follows: reserved tags used in management frames such as beacons, rate adaptation behavior, reassociation latency, backoff behavior, power save mode behavior, management packet transmission rate, honoring of duration field, calculation of duration field in transmitted packets etc.

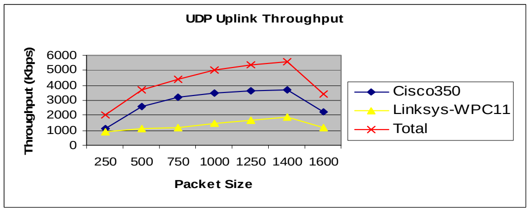

Figure 6 : Non-fair distribution of bandwidth

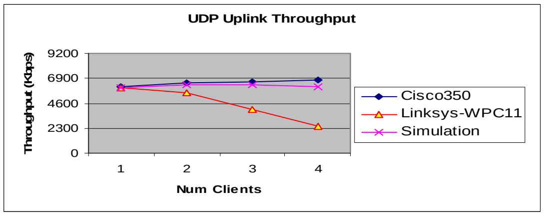

Figure 7: Inaccuracy of Throughput Prediction Tools

Device fingerprinting can be useful in helpful in multiple scenarios. For example, it can be useful in detecting identity thefts using "MAC spoofing". MAC spoofing is an attack where one device changes its MAC level identity and masquerades as another devices to defeat MAC level security mechanisms (e.g., MAC level access control). Conventional techniques to detect MAC spoofing [17] rely on both of the devices being active. However, with device fingerprinting, MAC spoofing across different vendors can be easily detected even if only one of the devices is active. Further, device fingerprinting can also be useful in automatically creating the wireless device inventory of an organization.

5. Conclusions and Future Work

In this paper, we have presented an experimental analysis of MAC implementations of IEEE 802.11 devices from several popular vendors. We believe that our study is the first ever empirical analysis of MAC implementations. We demonstrate that, contrary to common belief, devices implementing the same IEEE 802.11 standard often have differing MAC level behavior. Specifically, we demonstrate that in certain cases, devices may not be conformant with the standard (e.g., due to incorrect implementations), where as, in other cases, they may differ in details that are not a part of mandatory specifications of the 802.11 standard. We have shown that heterogeneous implementations can adversely impact WLAN operations leading to unfair bandwidth allocation, poor network utilization, potential break-down of related MAC level functionality and difficulties in provisioning a WLAN. However, on the positive side, heterogeneity amongst MAC implementations can be useful in vendor/model based device fingerprinting.

We are currently working on some of the interesting issues that arise from the results presented in the paper. First, understanding MAC level behavior of a large number of vendors in a systematic manner is a challenge in itself. There a large number of vendors out in the market and new devices enter the market frequently. We are developing an automated framework for collecting MAC level traces and analyzing the same to extract useful MAC level parameters of the devices in a "knowledge library". Second, vendor/model level fingerprinting of devices is an interesting topic. We are working on building tools that processes the "knowledge library" to obtain signatures for devices from different vendors. Finally, developing a MAC performance prediction model that can predict accurate WLAN performance (even in the presence of MAC level heterogeneity) is another interesting area that we are pursuing.

ACKNOWLEDGEMENTS

We would like to thank Hemant Chaskar, AirTight networks, for his valuable comments on the positioning of the paper. Further, thanks to Bharath. P and Praveen. K for helping with experiments.

References

1. Cisco Systems Inc., www.cisco.com.
2. D-link Systems Inc., www.dlink.com.
3. Intel Corporation, www.intel.com.
4. Netgear Inc., www.netgear.com.
5. 3Com Inc., www.3com.com.
6. Proxim Inc., www.proxim.com.
7. IEEE 802.11 Standard, Wireless LAN Medium Access Control (MAC) and Physical layer (PHY) Specifications, 1999.
8. Wi-Fi Alliance, www.wifialliance.com
9. Wi-Fi CERTIFIED®: The Safe Buy, October 5 2004, http://www.wi-fi.org/white_papers/whitepaper-100504-wifisafebuy/.
10. Hagen Woesner et al., "Modified backoff algorithms for DFWMAC's distributed coordination function, Proc. of 2nd ITG *Fachtagung Mobile Kommunikation '95*, Neu-Ulm, Germany, September 1995.
11. Byung-Jae Kwak et al., "Performance Analysis of Exponential Backoff", IEEE/ACM Transactions on Networking, Vol 13, Issue 2, April 2005
12. Prism Wireless LAN, http://www.intersil.com/globespanvirata/.
13. Sandhya G, K Gopinath, "Design and Analysis of Rate Aware Ad Hoc 802.11 Networks", Technical Report , CSA Dept., Indian Institute of Science, 2004.
14. G. Bianchi, Performance Analysis of the IEEE 802.11 Distributed Coordination Function", IEEE Journal on Selected Areas in Communications, Vol 18, No. 3, March 2000.
15. F. Cali, M. Conti, E. Gregori, "IEEE 802.11 Wireless LAN: Capacity analysis and protocol enhancement", In Proc. of Infocom' 98, San Francisco, USA, 1998.
16. Ping Chung Ng et al., "Experimental Study of Hidde-Node Problem in IEEE802.11 Wireless Networks", Poster Presentation, SIGCOMM 2005, PA, USA
17. Joshua Wright, "Detecting Wireless LAN MAC Address Spoofing", http://www.polarcove.com/whitepapers/detectwireless.pdf.