

An Analysis of Handoff in Multi-band 802.11 Networks

David Murray
Murdoch University
D.Murray@murdoch.edu.au

Terry Koziniec
Murdoch University
T.Koziniec@murdoch.edu.au

Michael Dixon
Murdoch University
M.Dixon@murdoch.edu.au

Abstract

The availability of public access WLANs (Wireless LANs) is growing with many cities now announcing municipal city-wide networks. This opens a new realm of possible applications such as telephony and gaming. However, before many of these applications can become widespread, the issue of mobility must be solved. Currently, the time required for wireless clients to handoff or transition their connection between APs (Access Points) is too long, causing call dropouts. This study investigates a specific aspect of handoff known as scanning. Improvements to current scanning mechanisms are proposed and tested in a variety of experiments. In addition, our experimental approach reveals previously unknown scanning issues in 802.11 networks.

1. Introduction

Wireless LANs and VoIP (Voice over IP) are now established and rapidly maturing technologies. The combination of these two technologies is called VoWLAN (Voice over WLAN) and has the potential to offer cell phone like service with the cost and efficiency of VoIP. In addition to cheap telephony, VoWLAN opens the possibility for new services previously impossible over lower speed technologies. Considering the success of cellular technologies, the potential for VoWLAN is huge. In the past, the idea of VoWLAN has been impossible because of the limited availability of WLAN hotspots. However, in recent times, large scale deployments have been attempted in New York, Philadelphia, San Francisco and New Orleans. These city-wide wireless LANs will provide the ubiquitous access required for VoWLAN. Before the technology can become widespread, a number of issues must be resolved. One of these issues is handoff, which is the ability for wireless devices to move between wireless networks. During handoff, wireless devices are unable to send or receive frames which can result in call dropouts. Handoff problems are magnified by the short range of WLAN technology.

Usable ranges between 25m and 50m means that handoff may be performed many times per call. If handoff takes too long, the disruption to connectivity could degrade call quality.

1.2 Handoff in Wireless LANs

Handoff in wireless networks is a broad research area. The IEEE 802.21 working group are investigating handoff between heterogeneous networks such as cellular networks and WiFi networks. Other work has designed mechanisms for IP mobility between subnets [1-2]. This paper focuses on the handoff between two APs in the same subnet. Prior studies [3-6] found that this delay is unacceptably high for time sensitive applications such as telephony and gaming.

Handoff occurs when a wireless client moves from the coverage area of one AP and into the coverage area of another. As APs are deployed on different frequencies or channels, wireless clients moving through wireless networks are unaware of neighboring APs until they begin the handoff process and scan for surrounding APs. When a wireless client is leaving the coverage area of an AP, the handoff process begins. SNR (Signal-to-Noise Ratio) is the commonly used heuristic.

Handoff consists of 3 phases: scanning, authentication and association. Scanning is the process of searching for surrounding APs and can be done passively or actively. Passive scanning involves monitoring the medium for periodically broadcasted messages known as beacons. In active scanning, the client probes for these messages. The APs probe response is used to determine which AP is offering the best SNR. This process typically takes between 70ms and 600ms [3], [5], [6]. Following the completion of the scanning phase, the client must authenticate with the AP offering the best connection. As a variety of authentication mechanisms can be used, the authentication delay can vary. No authentication, often known as open authentication, takes only a few milliseconds as the negotiation requires only a single request and response. More complex authentication

protocols such as 802.1X require 4-way handshaking techniques and can consume in excess of one second [7]. After authentication, clients must associate with the new AP. In the association phase the clients context and sessions are transferred between the APs. The association phase is relatively short, approximately 15ms [3].

To reduce handoff delays, the IEEE working group TGr was created. The goal of the TGr is “to develop a standard specifying fast BSS (Basic Service Set) transitions” [8]. The majority of the work done by TGr focuses on authentication and association phases. They define mechanisms to complete secure 4-way 802.11i authentication and also reservation of QoS (Quality of Service) prior to a roam. TGr places little emphasis on the scanning despite prior work [3], [5], [6] which has highlighted unacceptable scanning delays.

In the future, spectral limitations may require city-wide wireless networks to utilize both 802.11a which operates in the 5GHz spectrum and 802.11b/g which operates in the 2.4GHz spectrum. As there are only 3 usable channels in the 2.4GHz ISM band and up to 24 (depending on regulatory domain) channels in 802.11a networks, scanning delays are likely to be worse than prior work might suggest. Little prior work has investigated scanning in both the 2.4GHz and 5GHz spectrum. This study investigates scanning mechanisms in both 2.4GHz and 5GHz networks.

1.2 The Scanning Phase

To understand scanning, the concept of wireless channels must firstly be understood. Large wireless networks are built with many closely deployed APs. These APs have overlapping coverage areas. To avoid interference, nearby APs are configured on different frequencies or channels. As nearby APs operate on different channels, wireless clients are oblivious to surrounding APs until they scan for them. The scanning process requires wireless clients to switch between and scan each channel independently. Each channel can be searched in two ways, passively or actively.

Passive scanning is the slower of the two scanning mechanisms. It involves switching between channels and waiting for periodically broadcasted beacons. By default, most APs broadcast beacons every 100TUs (Time Units). A TU is equal to 1.024ms. To ensure beacons are received, clients must reside on each channel for a minimum of 100TUs or 102.4ms. Subsequently, with 11 channels in the 2.4 GHz band, a client will require over 1s to scan channels passively and over 2.4s to scan the 24 802.11a channels in the 5GHz band. Such lengthy delays are obviously unacceptable in VoIP networks.

Active scanning is a faster scanning mechanism than passive scanning and is used by the majority of consumer wireless cards. Active scanning involves proactively probing for APs instead of waiting for periodic broadcasts. The client initiates the active scanning process with a probe request. APs reply with a beacon like frame called a probe response. The duration spent actively scanning each channel is controlled by two timers, the minimum channel time and the maximum channel time.

The process begins with the client switching to a new channel and transmitting a probe request. Following the probe, the client starts a timer. If a probe response or any other 802.11 traffic is not detected by the minimum channel time, the channel is declared empty and the wireless client restarts the process on the next channel. If a probe response or any 802.11 traffic is received, the client concludes that one or more APs must exist on this channel. The client will then wait for the maximum channel timer to expire. This timer is sufficiently long to ensure that all AP responses can be collected. Although active and passive scanning is defined in the 802.11 standard [9], the duration of timers is vendor specific. The numerous studies [3], [5], [6] that empirically measured the scanning process found that scanning delays were between 70ms and 600ms.

2. Prior Work

A number of studies [3], [5], [6] have investigated scanning delays in an empirical manner. These experiments were generally performed by moving a wireless client between APs, forcing a handoff while multiple packet captures record traffic for offline analysis. Prior studies investigated scanning in 2.4GHz clients and found delays between 70ms and 600ms [3], [5], [6]. Upon review of the results, it is apparent that there are inconsistencies in the findings both between and also within studies. This is further investigated in section 3.

Scanning delays can be reduced in two ways. One solution is to reduce the time required to scan each channel through the optimization of channel timers. Another solution is to reduce the number of scanned channels to a subset where APs are known to exist. The majority of research falls into one of these two categories.

2.1 Timer Optimization

Scanning timers can be optimized with both passive scanning and active scanning mechanisms. Passive scanning involves collecting beacons

transmitted by APs every 100TUs (102.4ms). Some research [5], [10] has investigated the possibility of reducing the beacon interval. For example, if beacons were broadcasted at 10 TU intervals, instead of every 100TUs, passive scanning timers could likewise be reduced from 100TUs to 10TUs. Two studies modeled these optimizations in NS-2 (Network Simulator 2) [5] and OPNET [10]. Despite producing somewhat divergent results, both studies concluded that the loss of bandwidth as a result of the increased number of beacons was unacceptably high.

A different study [4] optimized passive scanning by synchronizing the broadcast of beacons. Instead of clients waiting 100 TUs to receive a beacon, a time synchronization mechanism enables wireless clients switch to a channel in anticipation of the beacon. However, the design requires clients to regularly perform this scanning operation. Although this mechanism provides better knowledge of the radio environment, the complex design and consistent packet loss from the continual scanning operation is too costly.

Active scanning delays can also be reduced by optimizing timers. Active scanning timers must be set carefully as timers set too low will result in missed APs. Velayos and Karlsson's simulations [5] concluded that the minimum channel time can be set as low as 1TU while maximum channel times should be approximately 10TUs. Pries and Heck [10] also modeled different channel timers and found that in locations with dense AP deployments, maximum channel timers should be as high as 27 TUs whereas in areas where AP coverage is sparse, maximum channel times could be set as low as 7 TUs. While optimized timers are an effective way to reduce scanning delays other research efforts have focused on channel pruning.

2.2 Channel Pruning

A number of studies [10], [11], [12], [13] have investigated the possibility of reducing the number of channels required to be scanned. Channel pruning requires wireless clients to be informed of nearby APs and the channels in use. With this knowledge, wireless clients can simply scan a subset of channels.

The difficulty with channel pruning is developing a way for APs to learn about neighboring APs. A proposed 802.11 amendment known as 802.11k (Radio Resource Measurement Enhancement) aims to provide a mechanism for wireless APs to learn of their radio surroundings. Presently, it is believed that inactive wireless clients will be used to perform site surveys and report information back to the AP. This information could be used to build a list of channels upon which nearby APs reside.

A combination of reducing the number of channels required to be scanned and reducing the amount of time spent scanning per channel can substantially reduce scanning delays. The research in this paper is divided into two sections: scanning in 802.11b/g networks and scanning in 802.11a networks.

3. Experiment 1: Scanning in 802.11b/g networks

3.1 Theory

An analysis of prior empirical work [3], [5], [6] reveals inconsistent results. Not only were results inconsistent between studies but also within studies. Large variations were observed within the same experiment using the same wireless equipment [3]. Considering that the handoff process is based on a series of timers, delays should be predictable and quantifiable. This section investigates the theoretical duration of handoff.

In the theoretical network which we use to calculate scanning delays, two APs exist on different channels in the 2.4GHz spectrum. A client in this network is required to scan 11 channels in total. On two channels the client will detect traffic and wait for the maximum channel timer to expire and on the other nine channels it will wait for the minimum channel timer to expire.

When scanning a channel where no AP exists, calculations of contention times in 802.11b networks indicate that a probe request should be transmitted within 2ms (Fig 1) and following the probe request, the client will start the minimum channel timer. Following the transmission of a probe, clients can quickly assess whether an AP is present. Within 670 μ s [5], the AP should have begun transmitting the response or, if another station has priority over the medium, generated some form of wireless traffic indicating that an AP is present. The time required to assess that a channel is empty requires about 3ms.




Fig 1 Channel Switch estimates and minimum channel time

Clients scanning channels where APs are present will undergo the same process but wait for their maximum channel time to expire. This maximum channel time is set high enough to ensure that even heavily loaded APs can respond. As access to the medium is randomly shared by all wireless stations, response times are dependent on AP load and the number of associated clients competing for the medium. Simulations have revealed that maximum channel times in 802.11b networks can be set between 10TUs and 27TUs [5], [10]. These estimates show that the majority of scanning delays stem from searching channels where APs are known to exist.

Another delay when scanning in wireless LANs is the channel switch delay; which is the time required for wireless cards to switch between and stabilize on a new frequency or channel. Ramani and Savage [4] measured switching delays in Atheros and Intersil wireless chipsets and found the delays to be 5ms and 20ms respectively. Our measurements yielded similar results however we believe that the majority of the delays were introduced by the testing environment. Hardware frequency switches are reportedly 40 μ s [14] however actual delays are higher as a result of drivers. Conservative estimates are that the switching delay is between 1ms and 5ms depending on the chipset and drivers.

The lowest theoretical scan time is $(4\text{ms} \times 9\text{chans}) + (14\text{ms} \times 2\text{chans})$ 64ms while the largest theoretical scan time is $(9\text{ms} \times 9\text{chans}) + (35\text{ms} \times 2\text{chans})$ 142ms. These estimates show that delays should be between 64ms and 142ms. However, experimental studies have shown that actual delays are between 70ms and 600ms [3], [5], [6]. This estimate leads to uncertainty over the source of delays. Although it could be argued that high delays are a result of relaxed minimum and maximum channel timers, Mishra et al, [3] observed a high degree of variation using the same equipment. If scanning delays are based on the minimum and maximum channel timers discussed, what is responsible for the variation in delay? This study investigates active scanning in search of the source of scanning delays.

3.2 Design

To analyze scanning delays, an experiment was designed to capture the scanning process. This would allow empirical measurement and a frame-by-frame analysis to search for sources of delay. Ethereal network analyzer software was used to collect, store and time-stamp packets. Analyzing protocols with Ethereal is usually a simple and common method to troubleshoot problems. However, the 802.11 scanning process is significantly more difficult to capture than

traditional wired Ethernet. Firstly, scanning frames are a type of management frame and are only able to be captured in special promiscuous wireless mode known as RF monitor mode. Interfaces in RF monitor mode are unable to transmit frames and cannot operate as a wireless station. Subsequently, unlike wired Ethernet, a separate interface is required to capture and store management frames sent in wireless networks. Secondly, as the scanning process occurs over many different channels or frequencies, many interfaces are required to capture the entire process. This experiment used four interfaces to capture scanning traffic. Each interface was set to a different non-overlapping channel. The wireless cards were Atheros 802.11a/b/g wireless cards running the Linux MADWiFi driver.

Previous research has shown that scan times are dependent on both the wireless card and the AP [3]. Five different wireless cards were tested including: Cisco Aironet 802.11b, Enterasys RoamAbout 802.11b, Cisco Aironet 802.11a, Cisco 802.11a/b/g and Ubiquiti SRC 802.11a/b/g wireless cards. The APs used were Cisco 1200's with a/b radios and Linksys WRTGS's running OpenWRT with b/g radios. Handoff was induced by physically moving a laptop containing one of the five wireless cards between the APs. The scanning process was also studied with and without background traffic. By adding another client into the wireless network and starting a large FTP file transfer, the effect of heavy traffic loads on scanning delays could also be examined. The purpose of this experiment is not to report on different vendors scanning algorithms, but instead, it is hoped that in-depth analysis may reveal why active scanning delays are so varied and why they do not fit our theoretical estimates.

3.3 Results

The results of the tests using the Cisco Aironet and Enterasys RoamAbout wireless cards are shown in Fig 2, 3, 4, 5 and 6. Fig 2 and Fig 4 show that the Aironet 802.11b wireless card had an average scanning delay of 235ms. Fig 3 and 5 show that the results of the Enterasys RoamAbout wireless card varied between 75ms and 275ms depending on the AP. Few prior studies have investigated scanning in 802.11a environments. One study [12] measured scan times of an 802.11a card however they do not specify their method for capturing 802.11a packets or the wireless client they used. They found that scanning delays were between 900ms and 1000ms. Our results suggest the opposite. Fig 6 shows the Cisco 802.11a card scanning the 8 non-overlapping channels in an average of 46ms, approximately five times faster than the Cisco 802.11b wireless card.




Fig 2. Aironet .11b client and Open WRT AP




Fig 4. Aironet .11b client and Cisco 1200 AP




Fig 3. Enterasys .11b client and OpenWRT AP




Fig 5. Enterasys .11b client and Cisco 1200 AP

A prior study [3] revealed that scanning delays are dependent on the AP. A comparison of Fig 3 and Fig 5 reveals that the scan times of the Enterasys wireless client differed with APs. Fig 5 shows that the Cisco 1200 AP resulted in faster but highly variable scan times whereas Fig 3 shows that the OpenWRT AP caused consistent delays of 250ms.

Another oddity was noticed with the introduction of FTP traffic. The black columns in the graphs indicate scanning with background FTP traffic. The white columns indicate scanning without background traffic. As expected, most wireless clients experienced slightly higher delays when competing for the medium with background FTP traffic, however, the Enterasys card was not capable of roaming to a channel saturated with FTP traffic. This is possibly due to an additional roam criterion or heuristic.

The Cisco 802.11a/b/g and the Ubiquiti SRC 802.11a/b/g wireless cards scanned in a non-standard manner.




Fig 6. Aironet .11a client and Cisco 1200 AP

Following a large initial scan, the wireless cards appeared to cache the channels on which APs were known to exist and subsequently perform frequent scans of those select channels. These smaller, single channel scans consumed approximately 50ms. Prior to these single channel scans, null frames were sent to the

AP. This allows wireless stations to switch into power save mode. Frames destined for the wireless client during this period are buffered at the AP. The length of the scanning delay was consistent for 802.11a and 802.11b scans and appeared to be a function of the switch into and out of power save mode. Although a scanning mechanism whereby the AP caches messages for a scanning client is conceptually ideal, it is unsuitable for highly mobile applications. Subsequently, the results of the Cisco 802.11a/b/g wireless card and the Ubiquiti SRC wireless card are not shown and comparisons have been omitted from the discussion.

3.4 Discussion

The pivotal question which stems from the results is: why are scan times approximately five times lower using the Cisco 802.11a card than the Cisco 802.11b card? A number of factors may contribute to lower scan times in 802.11a wireless cards. Firstly, the legacy Cisco 802.11a wireless card only scanned 8 channels whereas the 802.11b wireless cards scanned 11 channels. Secondly, 802.11b/g wireless cards transmit management frames at 1Mb/s and 802.11a wireless cards transmit management frames at 6Mb/s. Furthermore, the 802.11a standard has lower contention times. High data rates and low contention times allow frames to be serialized onto the medium faster. As less time is required to access the medium as well as transmit and receive frames, it is permissible that maximum channel times could be set lower in 802.11a cards.

Despite these factors, the principal reason scanning delays are lower in 802.11a wireless cards is channel overlap. The division of channels in the 5GHz spectrum is non-overlapping. The 2.4GHz spectrum consists of 11 channels, of which, only 3 are non-overlapping. The IEEE standard [9] specifies that all channels must be scanned as APs can be configured on any channel.

Packet captures show 802.11b APs responding to probe requests transmitted on overlapping channels. Comparatively, each 802.11a AP only responded to one probe request per client scan. The transmission of superfluous probe responses in 802.11b APs is a result of responses to probes on overlapping channels. As discussed in our theoretical estimates, the majority of scanning delays stem from scanning channels where APs are detected. Subsequently, the replies to probe requests on overlapping channels will cause clients to wait the duration of the maximum channel time instead of switching to the next channel after a few milliseconds. The reason that 802.11b scan times are higher than our theoretical estimates is because clients

are waiting the duration of their maximum channel time on channels that overlap with the APs designated channel.

Interpreting this observation is difficult because management frames do not specify the channel on which they were transmitted. This makes it difficult to confirm that, for example, a probe transmitted on channel 2 is being replied to by an AP on channel 1. The only channel information available is provided by the packet capturing wireless interface in RF monitor mode. This identifies which channel the packet was captured on. This is somewhat demonstrative of our theory. Packets sent on one wireless channel, can be processed by interfaces on different channels.

This observation explains how a process based on timers can produce the large variations in delay measured in this and previous experimental studies [3], [5]. As a result of physically where and when a client scans, an AP may, or may not, receive probe requests on any number of overlapping channels. Consequently a wireless card may, or may not, wait for the duration of the maximum channel timer on a given channel.

Wireless card manufacturers use physical layer filtering mechanisms to reduce the sidebands (or width) of RF transmissions however the extent of channel overlap in the 2.4GHz band is large. Furthermore, high transmit powers and close proximity to APs can exacerbate the extent of channel overlap. This novel concept is unique to the scanning phase as associating with an AP mitigates the problem. Once a client has associated with an AP, frames will be filtered based on the AP's MAC address.

A solution to alleviate these problems is to send channel information in the header of probe request frames. APs could read this header and ignore frames transmitted on different channels. Although this may not entirely solve the problem as clients may still detect traffic when scanning overlapping channels it could alleviate the problem in many circumstances. Another way to largely mitigate this problem is to use channel pruning mechanisms discussed in section 2.2. By pruning unused channels, the problem would be avoided because clients would no longer probe overlapping channels.

Comparisons between 802.11b/g and 802.11a scanning revealed previously unknown active scanning issues. However, due to spectrum licensing changes [15], [16], future 802.11a equipment will no longer be able to scan actively. Our second experiment investigates passive scanning in 802.11a networks.

4. Experiment 2: Scanning in 802.11a Networks

In anticipation of the need for more unlicensed spectrum space, 2003/2004 saw the 5GHz spectrum expanded from 12 to 24 channels in the USA [15] and from 4 to 19 channels in most European countries [16]. However, this spectrum expansion and harmonization was conditional. New 5GHz networking equipment is required to implement mechanisms to dynamically avoid and reduce interference to operate around military and aviation transmissions sharing the same spectrum space. To comply with these new regulations, the 802.11h [17] amendment introduced two mechanisms, TPC (Transmit Power Control) and DFS (Dynamic Frequency Selection). In addition to 802.11h, new regulations also require that:

“Slave Devices [wireless clients] shall not transmit before having received an appropriate enabling signal from a Master Device”[16].

The implicit rule in this requirement is that wireless stations may not operate on any channels until a beacon has been detected. Subsequently, wireless clients operating in the 5GHz band must use passive scanning.

4.1 Theory

Passive scanning in 802.11a networks presents a large problem. Under default conditions, APs transmit beacons every 100ms. With 24 channels, passive scanning delays will exceed 2.4s. Prior work [5], [10] has suggested that passive scanning delays can be reduced by increasing the frequency of beacon transmission. Subsequently, the amount of time required for the wireless client to remain on a channel can also be reduced. However, by increasing the number of beacons, the network overhead will also rise. Pries and Heck [10], and Velayos and Karlsson [5] both investigated the bandwidth trade-off associated with different beacon intervals in 802.11b networks. Despite producing somewhat divergent results, the conclusions were unanimous: the network overhead required to support fast passive scanning is unacceptably high. Both prior studies investigated fast passive scanning in 802.11b networks, however, given the new regulatory requirements, this study focuses on fast passive scanning mechanisms in 802.11a networks.

Management frames, such as beacons, are always transmitted at the minimum rate. Hence, in 802.11b/g networks, beacons are sent at 1Mb/s [18], [19]. Comparatively, 802.11a networks transmit beacons at 6Mb/s [20]. Higher data rates enable frames to be

serialized onto the medium faster. This consumes less of the AP’s time and is therefore more bandwidth efficient. This study begins with a theoretical calculation to investigate the efficiency of beacon transmission in 802.11a networks.




Fig 7. Beacon Transmission Times

A theoretical approach is initially used to calculate the amount of bandwidth consumed by beacons at different intervals. The calculation is based on the methods of Gast [21]. Each step is graphically represented in Fig 7. To match the practical experiment that follows this calculation we have used 228 byte beacons which equate to 1824 bits. These bits are serialized onto the medium at different rates depending on the modulation in use. In 802.11b/g networks, beacons are transmitted at 1Mb/s using DSSS (Direct Sequence Spread Spectrum) DBPSK (Differential Binary Phase Shift Keying). DBPSK divides each bit into a symbol and transmits one thousand symbols per millisecond (ms) or 1 symbol per microsecond (μs). Subsequently, 1824 bits are transmitted in 1824μs.

The 802.11a standard transmits beacons at 6Mb/s using OFDM (Orthogonal Frequency Division Multiplexing) BPSK (Binary Phase Shift Keying). BPSK transmits 24 bits per symbol with each symbol consuming 4μs for transmission. Subsequently, the 1824 bit beacon is represented by 76 symbols and is transmitted in 304μs. The transmission times for these beacons are shown in Fig 7.

In addition to the packet serialization delay, the MAC (Medium Access Control) layer delays must also be considered. IFS (Inter-Frame Spacing) is a MAC function used to prioritize different transmissions. Before transmitting a packet, wireless clients are required to wait for their IFS timer to expire. Different network operations have different IFS times to provide priority to critical network activities. Broadcast management frames use SIFS (Short IFS) [9]. The SIFS timer in the 802.11b/g standards is 10μs [18],

[19] whereas the SIFS timer in 802.11a networks is 16 μ s [20].

In addition to the MAC layer IFS delay, a physical layer preamble further adds to the total beacon transmission time. A preamble signals the beginning of a frame transmission and is used to ready wireless radios for communication. There are two preambles in the 802.11b standard, short and long. Short preambles are 96 μ s and long preambles are 192 μ s [18]. This calculation uses the short preamble as it matches the preamble used in later practical experiments. In 802.11a networks, preambles are only 20 μ s [20].

In 802.11b/g networks the preamble, IFS and frame serialization delays require 1930 μ s. Comparatively, in 802.11a networks the beacon transmission is much shorter, 340 μ s (Fig 7).

The aim of this calculation is to measure the throughput loss with different beacon intervals in 802.11 networks. Our measure of impact is the percentage of time consumed by transmitting beacons. Fig 8 graphically depicts the impact of beacons on the bandwidth capacity of APs. The results suggest that beacon intervals can be significantly reduced without radical performance loss. The results also show that lower beacon intervals in 802.11a networks have considerably less impact than in 802.11b/g networks. At 10TUs, the 802.11b standard suffers from an unacceptable 18% reduction in bandwidth. However the 802.11a standard loses only 3% of its potential throughput. To complement these purely theoretical calculations, an empirical experiment is proposed to investigate the impact of beacons in real 802.11a and 802.11b/g networks.




Fig 8. Theoretical: capacity used at different beacon intervals

4.2 Design

This experiment investigates the loss of bandwidth as a result of lower beacon intervals using real 802.11 equipment. This was tested by comparing the FTP download speed of a wireless client with different beacon intervals. On most APs the beacon interval is a

configurable variable. However, AP firmware is often written to ensure variables remain within reasonable parameters. Most APs do not allow beacon intervals of less than 20TUs. The exception is the open source, OpenWRT firmware that allows administrators, to change beacon intervals to any value. However, the OpenWRT firmware only supports 802.11b/g standards. When testing the 802.11a standard, a Linksys WRT55AG with proprietary Linksys firmware was the closest alternative. The Linksys firmware on the WRT55AG restricted the configuration of beacon intervals to 20TUs and above. This experiment consisted of configuring different beacon intervals on the AP and downloading an FTP file from the FTP server. This test was conducted with beacon intervals of 100, 80, 60, 40, 20, 10, 5, 3 and 2TUs with the 802.11a/b/g physical layer standards. The wireless client downloaded a 300MB file from the FTP server. The bandwidth was averaged over the duration of this file transfer.

A concern was that the aggressive behavior of TCP traffic might prevent beacons from being transmitted on the network. Logically, this should not occur, as SIFS (Short Inter-Frame Spacing) should prioritize beacons over data traffic, however as a precautionary measure, a wireless packet capture was used to ensure the number of beacons transmitted per second were congruous with the beacon interval.

4.3 Results

The results in Fig 9 confirm that beacon intervals can be reduced considerably with only minor performance ramifications. The performance loss with beacon intervals between 100TUs and 40TUs is negligible. The impact of beacons on 802.11b and 802.11g networks are similar because beacons require the same amount of airtime. However, as our theoretical calculation showed, beacon transmission is significantly faster in 802.11a networks.

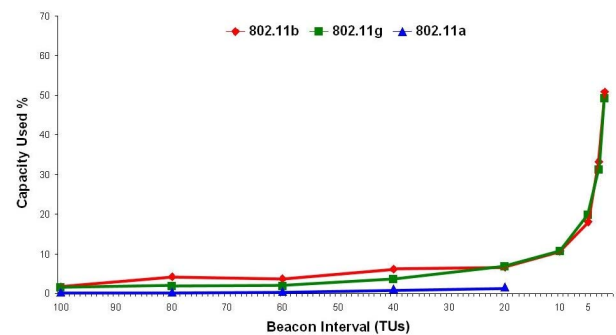


Fig 9. Practical: AP Capacity used at different beacon intervals

Although the firmware of our wireless AP limited beacon intervals between 100TUs and 20TUs, it is apparent that lower beacon intervals in 802.11a networks have a comparatively negligible impact (Fig 9).

4.4 Discussion

The reason that both a theoretical calculation and a practical experiment were performed, was to produce confirmatory results. Fig 10 compares the results of the theoretical (left) and practical (right) studies. It shows that at beacon intervals between 10TUs and 2TUs, the study produces somewhat divergent results. At beacon intervals of 10TUs, the theoretical calculation shows that bandwidth consumption in 802.11b/g networks is approximately 18%. Comparatively the practical experiment suggests that at 10TUs, bandwidth loss is only 10%. The most obvious explanation is that fewer beacons were being transmitted. However, the packet captures confirmed that beacons were being transmitted at the required rate. Despite these unexplained differences, both studies provide similar results until very low beacon intervals. Our theoretical results can be viewed as a worst case scenario and the practical calculation can be viewed as a best case scenario.




Fig 10. Theoretical calculation vs practical experiment

This research aims to propose fast and efficient scanning mechanisms to support voice applications. Considering the theoretical calculation as a worst case scenario, an 802.11a network can use beacon intervals of 5TUs with only a 6% reduction in capacity. Ignoring the channel switch delay, a passively scanning client could scan the 24 802.11a channels in slightly over 120ms - a significant improvement over the default 2400ms delay. Furthermore, this mechanism is over 4 times more efficient than many active scanning mechanisms. To compare our mechanism to active scanning wireless cards measured in section 4, our fast passive scanning mechanism scans over twice the number of channels in less than half the time. The fast

passive scanning mechanism we propose could operate in conjunction with other scanning optimizations such as channel pruning.

5. Conclusion

This paper investigates fast scanning mechanisms in both 802.11b/g and 802.11a networks. In section 3, an experiment was used to measure and analyze the scanning delay under a range of network conditions. The lengthy and varied scanning delays in 2.4GHz 802.11b/g networks were highlighted by comparatively low 5GHz 802.11a scan times. The cause of lengthy and varied scanning delays in the 2.4GHz band is overlapping channels. APs and clients are unable to distinguish between messages transmitted on overlapping channels causing actively scanning clients to detect traffic on multiple channels. Subsequently, wireless clients unnecessarily wait the duration of their maximum channel timer on overlapping channels causing lengthy and unpredictable scanning delays. In section 4, fast passive scanning mechanisms in 802.11a networks are explored and tested. Recent spectrum licensing changes have mandated the use of passive scanning in 802.11a networks. This study used theoretical calculations, reinforced by practical experimentation to explore fast passive scanning. This research shows that the low serialization times of 802.11a management frames allow frequent beacon transmission at minimal cost to the network.

In the future, computers, PDAs and phones will utilize WLANs as a converged mobile communications infrastructure providing voice and data services. To provide acceptable performance, the deployment of wireless LANs will be dense and will fully utilize the available spectrum space in 2.4GHz and 5GHz bands. Mobile devices moving through WLANs must be able to quickly locate and seamlessly transition their connectivity between APs. Lengthy delays will result in unacceptable call quality. This paper presents investigative research in active and passive scanning. It proposes, prototypes and evaluates fast scanning mechanisms for VoIP capable mobile devices.

6. References

- [1] IETF, "Mobility for IPv6." *RFC 3775*, 2006.
- [2] IETF, "IP Mobility Support for IPv4." *Internet Draft RFC3344*, 2006.
- [3] A. Mishra, M. Shin, and W. Arbaugh, "An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process," *ACM SIGCOMM Computer Communications Review*, vol. 33, pp. 93 -102, 2002.

- [4] I. Ramani and S. Savage, "Syncscan: Practical fast handoff for 802.11 Infrastructure Networks," *IEEE INFOCOM The Conference on Computer Communications*, vol. 1, pp. 675 - 678, 2005.
- [5] H. Velayos and G. Karlsson, "Techniques to Reduce IEEE 802.11b MAC Layer Handover Time," *ICC IEEE Conference on Communications*, vol. 1, pp. 3844 - 3848, 2004.
- [6] J-O. Vatn, "An Experimental Study of IEEE 802.11b Handover Performance and its Effect on Voice Traffic," *Tech Rep*, KTH Royal Institute of Technology, 2003.
- [7] A. Mishra, M. Shin, and W. Arbaugh, "Proactive Key Distribution using Neighbor Graphs," *IEEE Wireless Communications*, vol. 11, pp. 26 - 36, 2004.
- [8] IEEE-TGr, "802.11 TGR Just-In-Time Transition Acceleration Proposal (JIT-TAP)," *IEEE 802.11r Working Group Proposal*, 2005.
- [9] IEEE, "802.11 - IEEE Standards for Information Technology -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications." *IEEE Standards*, 1999.
- [10] R. Pries and K. Heck, "Simulative Study of the WLAN Handover Performance," in *OPNETWORK 2005*, (Washington D.C., USA), 8 2005.
- [11] M. Shin, A. Mishra, and W. A. Arbaugh, "Improving the Latency of 802.11 Handoffs using Neighbor Graphs," in *MobiSys '04: Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services*, (New York, NY, USA), pp. 70-83, ACM Press, 2004.
- [12] S. Shin, A. G. Forte, A. S. Rawat, and H. Schulzrinne, "Reducing MAC Layer Handoff Latency in IEEE 802.11 Wireless LANs," in *MobiWac '04: Proceedings of the Second International Workshop on Mobility Management & Wireless Access Protocols*, (New York, NY, USA), pp. 19 - 26, ACM Press, 2004.
- [13] H-S. Kim, S-H. Park, C-S. Park, J-W. Kim, and S-J. Ko, "Selective Channel Scanning for Fast Handoff in Wireless LAN using Neighbor Graph," in *ITC-CSCC2004: The 2004 International Conference on Circuits/Systems, Computers and Communications*, 2004.
- [14] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted Seeded Channel Hopping for Capacity Improvement in IEEE 802.11 Ad-hoc Wireless Networks," in *MobiCom '04: Proceedings of the 10th Annual International Conference on Mobile Computing and Networking*, (New York, NY, USA), pp. 216 - 230, ACM Press, 2004.
- [15] FCC, "Report and Order (FCC 03-278): Revision of Parts 2 and 15 of the Commission's Rules to Permit Unlicensed National Information Infrastructure (UNII) Devices in the 5 GHz band." *FCC Rules*, 2003.
- [16] ECC, "ECC Decision of 12 November 2004 on the Harmonized use of the 5 GHz Frequency Bands for the Implementation of Wireless Access Systems Including Radio Local Area Networks (WAS/RLANs)," *ECC Decision*, 2004.
- [17] IEEE, "802.11h - IEEE Standard for Information Technology -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Spectrum and Transmit Power Managements Extensions in the 5GHz band in Europe", *IEEE Standards*, 2003.
- [18] IEEE, "802.11b IEEE Standard for Information technology -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications— Amendment 2: Higher-speed Physical Layer (PHY) extension in the 2.4 GHz band", *IEEE Standards*, 1999
- [19] IEEE, "802.11g IEEE Standard for Information technology -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications— Amendment 4: Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band", *IEEE Standards*, 2003
- [20] IEEE, 802.11a IEEE Standard for Information technology -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications— Amendment 1: High-speed Physical Layer in the 5 GHz band, *IEEE Standards*, 2003
- [21] M. S. Gast, 802.11 Wireless Networks: The Definitive Guide. O'Reilly, 2005.