

AIR EQUALIZER: The Hidden Node

TECHNICAL REPORT

Report

Wireless LANS

1 IEEE 802.11 Wireless LAN

The IEEE (Institute of Electrical and Electronics Engineers) 802.11 Committee was formed as an independent standards group within IEEE 802.4L in 1990, and in 1991 the first IEEE wireless LAN workshop took place to discuss standardising the burgeoning 802.11 Wireless LAN family [PPC95]. From this point, the 2.4 GHz frequency was defined for unlicensed use, so the 802.11 wireless technology found its feet, resulting in a number of 802.11 Wireless standards with differing data rates and frequency ranges.

The success of 802.11 WLAN is evident from the extensive roll-out of pockets of high speed WLAN access. These coverage areas have enjoyed a type of vernacular branding and are commonly referred to as ‘HotSpots’. University campuses, cafés, restaurants and airports now provide IEEE 802.11 access to consumers. Another indication of WLAN’s success is the backing it has garnered from industrial parties. WLAN has received a form of branding through The Wireless Ethernet Compatibility Alliance (WECA) [wecge] who initiated a Wi-Fi (“wireless fidelity”) certification program allowing any 802.11 vendor to have its products tested for interoperability. Wi-Fi compatible products and HotSpots receive a certified Wi-Fi branding. Wireless LANs have gained strong popularity in a number of vertical markets, including health-care, retail, and manufacturing [CZ06].

This section provides an overview of the IEEE 802.11 Wireless LAN technology [80203a], in terms of the different standards, architectures and QoS mechanisms.

1.1 IEEE 802.11 Wireless LAN Standards

The 802.11 WLAN lineage is firmly rooted in the IEEE 802 standards family, which comprises a series of specifications for local area networks. Figure 1 shows where 802.11 fits in the 802 family tree. The 802.11 specifications are closely related to the 802.3 Ethernet standards for wired lans, and WLAN works comfortably with existing Ethernet systems. The standardisation of 802.11 was an effort by the IEEE to consolidate the various proprietary standards that were being produced by industrial vendors into a single homogeneous system. The 802.11 standard defines only the MAC and PHY layers, as the 802.2 logical link

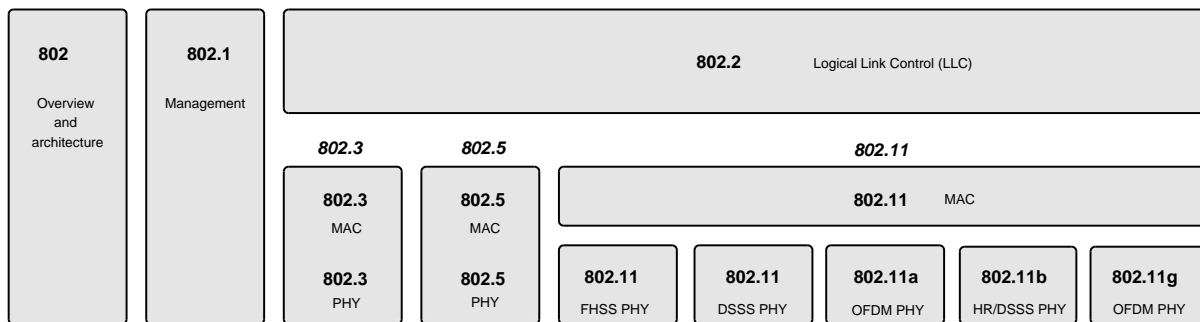


Figure 1: IEEE 802 Family

layer enables the interworking of WLAN and the 802 family of standards. This also allows for the support of various higher layers in accordance with the OSI (Open Systems Interconnect) model. Since its inception the 802.11 standard group has evolved to incorporate higher bandwidth flavours of WLAN in addition to a large number of other supporting specifications including measurements, QoS and radio resource management.

In 1997, the IEEE released 802.11 as the first internationally ratified standard for wireless LANs. The legacy 802.11 standard specified speeds of 1 and 2 Mbps to be transmitted via infrared (IR) signals or by either Frequency Hopping Spread Spectrum (FHSS) or Direct-sequence Spread Spectrum (DSSS) in the Industrial Scientific Medical (ISM) frequency band at 2.4 GHz. Two years later, in September 1999, the

| IEEE Standard | Description |
|---------------|--|
| 802.11a | 5 GHz, 54 Mbps |
| 802.11b | 2.4 GHz, 11 Mbps |
| 802.11d | Multiple Regulatory Domains |
| 802.11e | Quality of Service |
| 802.11g | 2.4 GHz, 54 Mbps |
| 802.11h | Dynamic Frequency Selection and Transmit Power Control |
| 802.11i | Security |
| 802.11j | Japan 5 GHz Channels |
| 802.11k | Radio Resource Management |
| 802.11n | High-Speed MIMO |
| 802.11p | Wireless Access for Vehicular Environment |
| 802.11r | Fast Roaming |
| 802.11s | ESS Mesh Networking |
| 802.11u | Interworking with non-802 networks |
| 802.11v | Wireless Network Management |
| 802.11w | Protected Management Frames |
| 802.11y | 3650-3700 Operation in the U.S. |

Table 1: IEEE 802.11 Specifications

original 802.11 standard was supplemented with ‘High Rate’ amendment [80299b] to include the higher data rates of 11 Mbps and 5.5 Mbps, and this became known as 802.11b.

1999 also saw the 802.11a [80299a] specification sanctioned with the introduction of new frequency bands in the unlicensed 5GHz range. 802.11a delivers a theoretical 54 Mbps data rate by employing an Orthogonal Frequency Division Multiplexing (OFDM) phy layer. 802.11a has 12 non-overlapping channels, 8 dedicated to indoor and 4 to point to point. It is not interoperable with 802.11b, except if using equipment that implements both standards.

In 2003 a third flavour of 802.11 was introduced in the form of 802.11g [80203b]. By applying the same phy modifications as 802.11a to the lower 2.4Ghz band, theoretical data rates of 54 Mbps were made possible, with the added benefit of being backwards compatible with 802.11b devices.

1.2 WLAN Architecture

A contributory factor of the success of WLAN is its relative architectural simplicity. The roll-out of a WLAN network does not demand the installation of a multitude of devices and supporting network entities. Since WLAN sits comfortably with wired Ethernet networks that have become the *de facto* standard for office and home networking, the backbone of WLAN networks is generally in place and WLAN enables a wireless connection point or bridge to existing IP infrastructures. The description of WLAN architecture can be split along the lines of the entities that comprise the network (network elements), and the topologies to which those elements conform.

1.2.1 WLAN Network Elements

The fundamental elements that comprise a WLAN network are shown in Figure 2. Central to the operation of the wireless LAN is the Access Point (AP). The AP acts as the coordinator of wireless resources over the bandwidth spectrum the WLAN is using. Chief among the responsibilities of the AP is the bridging of the wireless network to the wired backbone. Since the clients communicate with the access point via WLAN frames, the AP must convert these frames to packets for delivery over the network to which it is connected. The AP is also responsible for authentication and de-authentication of clients, association and re-association, beacon sending and frame acknowledgement and signal and frame buffering.

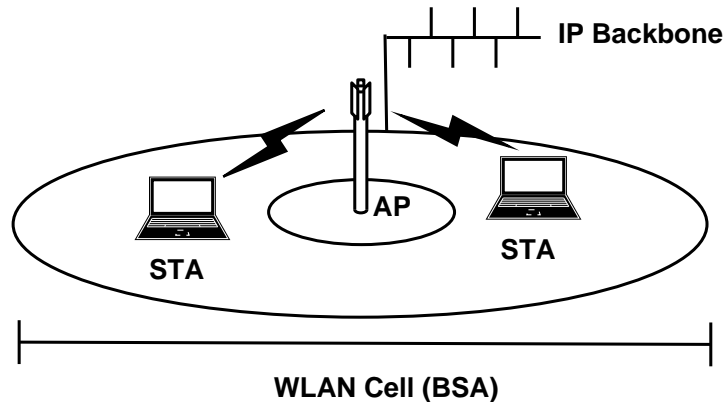


Figure 2: IEEE WLAN Network Elements

The clients that connect to the AP can differ in terminology but are generically referred to as Wireless Stations or STAs. Stations are devices, ranging from laptops, PDAs, Desktops and smart phones, that have a wireless network interface for communicating with the AP. STAs provide the human interface to the WLAN and depending on the device can support a range of IP service from Internet, Email, Video Conference and file transfer. More recently consumer devices such as cameras and portable music players are sporting WLAN interfaces for quick uploading and downloading of data.

The radius of radio coverage that AP transmits over is known as a WLAN cell or Basic Service Area (BSA). Depending on the flavour of 802.11 being used the cell footprint can range from a few meters to a few hundred meters. Since the widespread deployment of 802.11 wireless cells, the phrase 'HotSpot' has been coined to describe the pockets of WLAN cells that are rapidly becoming available in public places such as airports, libraries and university campuses.

1.2.2 WLAN Network Topologies

The arrangement of WLAN APs and the client STAs is known as a WLAN network topology. The topology is dependent on the mode of operation of the STAs and as a precursor to describing the WLAN network topologies, the two modes of STA operation must be detailed. In *Infrastructure Mode* the STAs are connected to an AP, and the AP is connected to a backbone network. All the communication in the network flows through the AP and the AP coordinates the radio access for each STA. Figure 2 is an example of Infrastructure mode.

In *Ad-Hoc* mode STAs communicate with each other directly, without the need for a central coordinator. In this mode all the STAs perform the dual function of the being a client and an AP at the same time, providing a point of connection for other devices within range. When multiple STAs form a communication path between each other it is known as an Ad-Hoc Network.

The basic building block of a WLAN topology is the Basic Service Set (BSS). A BSS is simply grouping of at least two or more STAs that communicate with each other. The area over which they can communicate is the Basic Service Area (BSA).

Independent BSS (IBSS)

Generally, when a BSS supports peer-to-peer communication but has no connection to an external network, it is known as an Independent BSS (IBSS). This is the simplest form of WLAN Network, and these networks communicate on an ad-hoc basis. Stations must be within range of the stations they wish to communicate with. This is illustrated in Figure 3 (a).

Infrastructure BSS

When one of the stations in the BSS is an Access Point (AP), the network becomes an Infrastructure BSS. This creates cell-like coverage areas, with each station in the BSS associated with the AP. The AP provides similar functions to a Base Station or Node B, because each station must communicate through the AP. The AP regulates the transmission of information and the access to the medium. This is shown in Figure 3 (b).

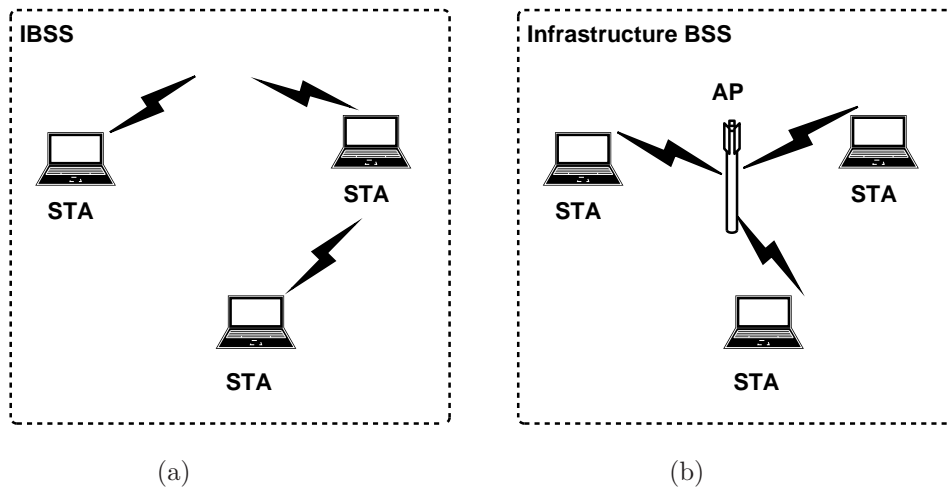


Figure 3: WLAN IBSS and Infrastructure BSS Topologies

Extended Service Set ESS

An Extended Service Set (ESS) is formed when one Infrastructure BSS or more are connected to a Distribution System (DS), as illustrated in Figure 4. The DS can be wired or wireless and supports connectivity between Infrastructural BSSs and to the outside world. There are no restrictions from the IEEE 802.11 standard on how to implement the DS, simply on the services that it provides.

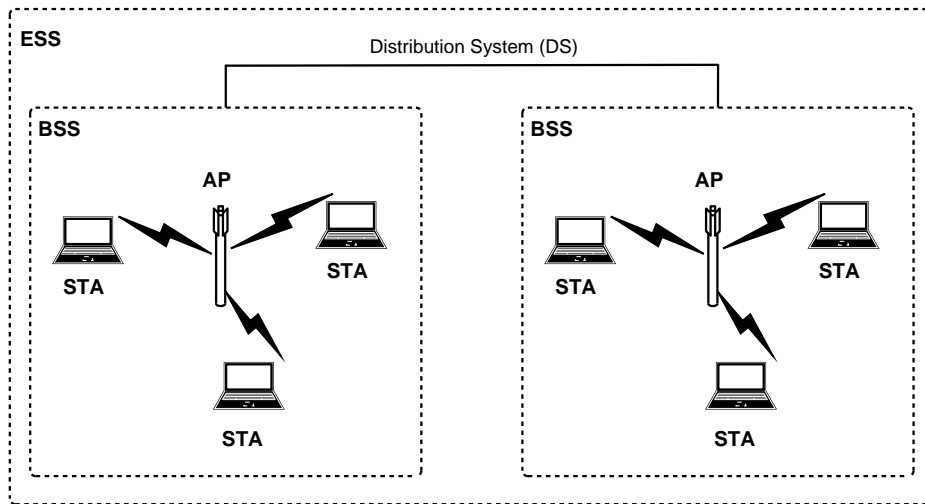


Figure 4: WLAN ESS Topology

1.3 WLAN Protocol Architecture

Following in the OSI [DZ83] vein of layered architectures, 802.11 WLAN has a series of layers, sub-layers and corresponding stacked protocols that control the flow of data and signalling throughout WLAN subsystems. Figure 5 illustrates the most important layers in the 802.11 protocol stack. The *physical layer* is responsible for the transmission of data via the radio interface using varying modulation schemes according to the flavour of 802.11 being used. The *Data Link Layer* (DLL) is made up of two parts; the Medium Access (MAC) Layer and the *Logical Link Controller* (LLC). The MAC layer is responsible for deciding when a station can access the transmission medium while the LLC, which is derived from the IEEE 802.2 common to both wired and wireless LANs, is responsible for detection and re-transmission of dropped packets. Although the MAC layer is found in both wired and wireless Ethernet devices, the functions it performs are very different due to the environmental differences between radio and fixed line mediums. This section introduces the MAC layer and the protocols it implements to ensure reliable packet delivery over the 802.11 wireless interface.

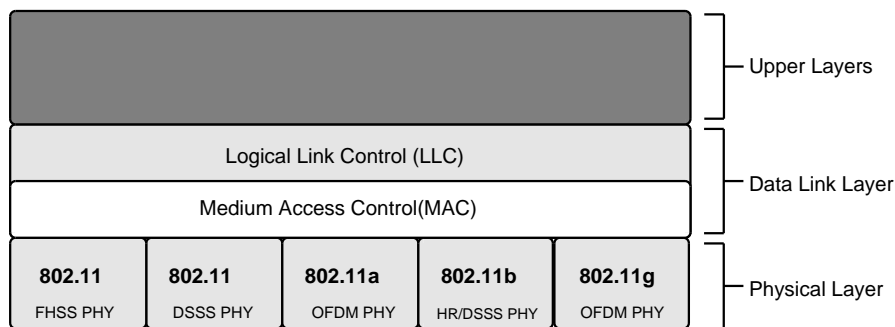


Figure 5: 802.11 Protocol Stack

1.3.1 Medium Access Control (MAC)

Radio bandwidth is a finite resource and the number of available channels is limited. WLAN then must share the available channels and divide the bandwidth fairly among multiple users. The protocols used to determine who get access on a multi-access channel belong to a sublayer of the DLL called the MAC (Medium Access Control). When dealing with an ever fluctuating transmission medium like radio, the MAC must be able to adapt to the wireless conditions in order to deliver packets for all users. This represent a particularly challenging aspect to MAC design. The MAC implements a number of mechanisms to deliver packets as reliably as possible over radio.

Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA)

When using a wired Ethernet channel, once a packet is transmitted, it is reasonable to assume that the packet reaches its destination. Since Ethernet users must shared a fixed line channel, Ethernet uses a protocol known as Carrier Sense Multiple Access (CSMA), to ensure that packet collisions do not occur.

Using CSMA, when a station wishes to transmit some data, it first *listens* to the channel to see if any other station is transmitting at that moment. If no other station is detected as transmitting, the channel is deemed idle, and the station transmits its data frame. This allows one user to user the channel at a time, to reduce the number of packet collisions.

The possibility does exist however, for one station to falsely detect that the medium is idle (does not receive indication that another station wished to use the channel), and transmit data. When this happens packet collisions occur and the data frames may become corrupt. An extension to the CSMA protocol is CSMA/CD (Carrier Sense Multiple Access/Collision Detection). Using this protocol, if two station detect the channel idle and start transmitting at the same time, they will both detect the collision of packets.

Rather than finishing sending their frames, which will inevitably become lost or corrupt, they immediately stop transmitting. This reduces the number of frames that need to be re-transmitted.

Using radio however, introduces a new level of instability when it comes to packet transmissions, especially when using the heavily populated, unlicensed ISM bands which are plagued by RF interference. Inheriting some of the functionality from CSMA/CD, a new protocol was adopted to work with WLAN, with the aim of avoiding packet collisions altogether. This protocol is known as CSMA/CA (Carrier Sense Multiple Access/Collision Avoidance) [80299b]. This multiple access scheme senses the medium after a random time to detect transmissions. If the medium is not in use, the station can transmit on the channel. If the medium is busy, the station waits until the end of the current transmission. It then waits for a random number of slots in the contention window corresponding to the backoff counter, before sensing the medium again. If transmissions are ongoing at this time, the station freezes the backoff counter, otherwise it transmits. If the medium is busy, the station waits until the transmissions end and then resumes the backoff counter. When it reaches zero, the station transmits.

1.4 Hidden Node

Relative to an existing transmission between a sender A and a receiver AP, a hidden node B is one that is within the interfering range of the receiver AP but out of the sensing range of the sender A. The nature of the hidden terminal is such that the sender A can not detect B's existence, but transmission from B can cause collisions at the receiver AP (even if AP is not the intended receiver by B) and hence disrupt the existing communication between A and AP, and ultimately degrade the network throughput.

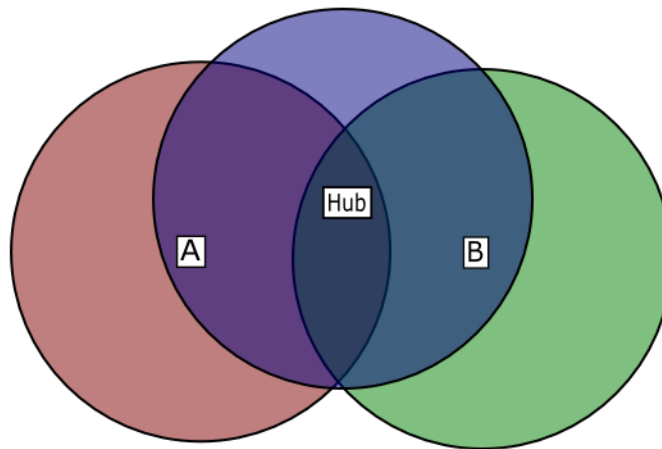


Figure 6: Hidden Node

Hidden node problem is a hard problem because it is difficult to detect hidden nodes in a wireless network, esp., in a multi-hop mesh network. First of all, carrier sensing is typically done at the transmitter, but by definition the transmitter alone can't detect the existence of the hidden nodes. Secondly, the interference range of the receiver is a logical concept but it is not a fixed range because it varies based on a number of variables including the distance between A and AP, transmission power at A, etc. All these variables can change from packet to packet. It is especially worthwhile to point out that interference range is not the same as transmission range. Transmission range of node X is the range within which all nodes can hear the transmission from node X and is able to decode the packet.

Hidden node problem exists in the traditional BSS networks centered around Access Points, and 802.11 virtual carrier sensing with RTS/BTS handshake is designed to mitigate the hidden node problem. RTS/BTS works reasonably well for such one-hop networks by informing the nodes in the 2-hop transmission range

neighbourhood (around sender and receiver), and hence reducing the chance of hidden nodes. However, the use of RTS/BTS does not completely solve the hidden node problem when multiple BSS networks co-exist in the same physical environment. Because interference range may be larger than the transmission range, and hence a hidden node B may still exist in a different BSS than the one A and AP belong to. Typically, non-overlapping channels are assigned for neighbouring BSS cells to mitigate such problems.

| Parameter | Value |
|---------------|--------|
| Max Penalty | 140 |
| Penalty Unit | 5 |
| Ratio | 85 |
| Trunk Up | 192000 |
| Trunk Down | 192000 |
| Ancient | 20 |
| Inactive Tics | 200 |
| Moving Avg | 8 |
| Brain Size | 3000 |
| Drop Count | 19 |
| Hog Min | 12000 |
| Hog Max | 32000 |

Table 2: AirEqualizer Parameters

Evaluating the Hidden Node

2 Hidden Node Scenario

In this section a scenario is developed to evaluate the hidden node problem in a reliable test environment, using real IEEE 802.11 equipment. Employing real equipment builds confidence in accuracy of results that correspond to a real-world setting. The hidden node problem is examined in a controllable test platform giving results that are consistent and repeatable. Nodes are placed in isolation chambers which allow virtual distancing to achieve the appropriate setup as in Figure 7.

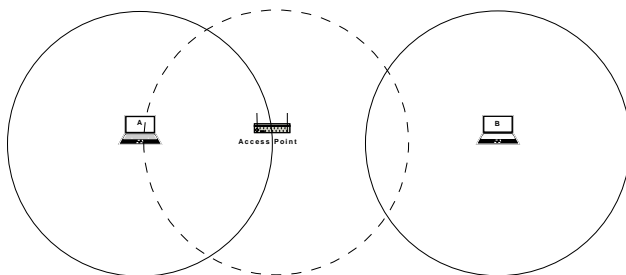


Figure 7: Hidden Node

Node A is positioned in close proximity to the access point and has a stronger signal than Node B based farther away. Both Node A and Node B can communicate with the access point, however Node A and Node B cannot communicate with each other.

The test scenario entails Node A initiating a session where it downloads a large file (649Mb), after 10 seconds Node B also attempts to download the same file. Tests are carried out with the AirEqualizer turned off and on.

Following a series of configuration tests varying the different parameters of the AirEqualizer, mainly Ratio, Trunk Up/Down and the size of the Penalty unit, optimum values were calculated and are tabulated in Table 2.

The tools used include the Azimuth test environment, two Asus Eee PCs and the AirEqualizer.

2.1 Azimuth Test Environment

A practical test platform is developed taking into account several design considerations outlined by De et al. in [DRSc05]. The purpose of the test platform is to provide repeatable laboratory performance evaluation of wireless networks. The test platform provides a dynamic framework for exploring a range of tests, one of which is the impact of mobility on a wireless network. Predominantly, network performance evaluation studies are done by means of stochastic simulation. However, studies have shown that simulators do not always reflect reality [SSG07] if not configured correctly. Therefore, using real-world devices (e.g., IEEE 802.11) in a controlled test environment, which our test platform provides, builds confidence in performance results and outcomes. Another example of such as testbed is presented by Metreaud et al. ??.

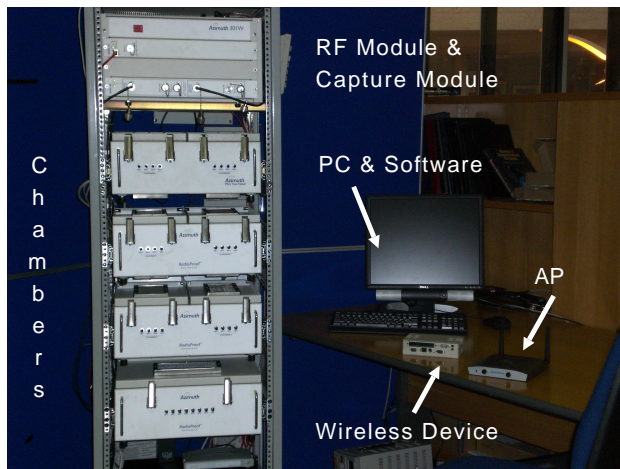


Figure 8: RF Isolation test platform, pictured with Access Point (AP) and Wireless Device (WD)

2.2 Architecture

At the centre of the test platform's architecture is the Azimuth W-Series Wireless LAN analyser [AS07]. This part of the test platform allows advanced testing and measurement of IEEE 802.11 wireless devices in RF-isolated chambers. The test platform allows creation of an IEEE 802.11 test environment in which real devices, contained in the RF-isolated chambers, can be virtually distanced from each other using controlled channel attenuation. For example, consider a Wireless Device (WD) and Access Point (AP) in Fig. 8. They are enclosed in separate RF-isolation chambers which provide 120dB isolation. The antenna of each device is connected directly to the controllable channel. The pair communicate with each other over this controlled attenuation channel. A variation in the channel attenuation corresponds to a variation of the virtual distance between the wireless devices which correlates to device movement. Thus, by varying the controlled attenuated channel it is possible to map the movement of a user according to a distinct mobility pattern. The test platform does not emulate the wireless channel of a real WLAN. It simulates changes in distance via digitally controlled attenuators. Essentially, the test platform consists of RF-isolation chambers, a chassis backplane where modules can plug into, and management software. The following sections explain these in further detail.

2.2.1 RF-Isolation Chambers

The isolation chambers are radio proof enclosures that provide an isolated environment to accommodate and attach IEEE 802.11 devices for real-world testing. Each independent isolation chamber can house up to two network enabled devices (representing users standing side by side). The discrete RF-isolated

| Parameter | Value |
|------------|----------|
| Technology | 802.11 |
| Band | 11b |
| Channel | 1 |
| Frequency | 2412 MHz |

Table 3: Test Environment Parameters

chambers provide both radiated and conducted isolation, prohibiting unwanted radio frequency interference from either entering or exiting the chambers. A filter module controls conducted path emissions which allows the desired DC and Ethernet signals to pass on these lines while filtering undesired signals to each chamber. The minimisation of RF leakage on these sources ensures RF integrity.

2.2.2 Chassis Backplane

The test platform's chassis backplane contains slots which house modules that combine and split RF signals to emulate fading in a real wireless network. Each chassis module connects to the chassis's backplane which enables a variety of test scenarios to be implemented.

2.2.3 RF Module

The Radio-Frequency module consists of a number of attenuators used to create sophisticated wireless network configurations. It is possible to virtually roam clients attached to the module's ports using dynamic range attenuation. Extensive RF filtering, which prevents RF energy from leaking into adjacent test modules and external devices, is employed in the specific module design to ensure RF isolation.

2.2.4 Capture Module

The test platform contains a multi-purpose module capable of capturing packets in the wireless network. The *Capture* mode enables the gathering of accurate timestamps of wireless traffic for analysis using the software analysis tool Ethereal.

2.2.5 Management Software

A software application acts as a centre for command and control of the test platform, supporting a flexible means of configuration for the hardware modules via a tool command language (TCL) scripting engine. The software provides complete management of the test platform via three separate and distinct networks:

- Test-Network: Network under test; runs only test traffic.
- Bus-Network: Management, control, configuration and statistics gathering of devices under test.
- Pub-Network: Functions as an Ethernet Local Area Network connection to networks/servers that exist outside of the test environment, including Internet, Intranet and external servers.

2.3 Screen Shot and Settings of Azimuth Director

Figure 9 displays a screenshot of the Azimuth Director. The left column lists the connected wireless devices; Node A, Node B and the AirEqualizer. The topology view also shows these devices, including the TMM which is configured to capture all wireless traffic in the network.

Table 3 lists the test environment settings for the experiments. IEEE 802.11b technology is employed in this work.

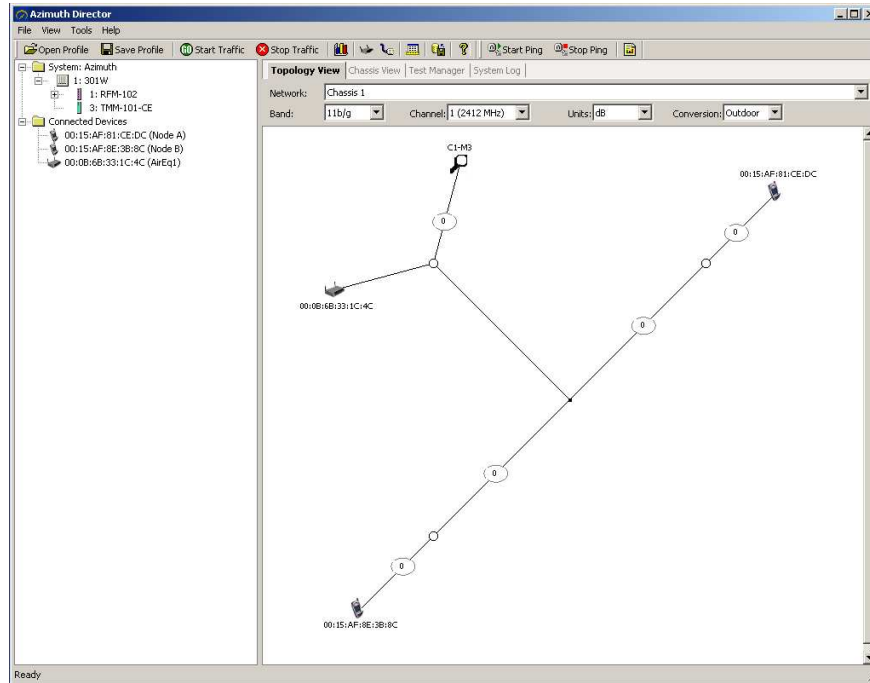


Figure 9: Azimuth Director Screenshot

| RF Path | IEEE 802.11b 2400 - 2497 MHz | |
|---------------------------------|---------------------------------|---------|
| | RF Port 1A to RF Backplane Port | 24.0 dB |
| RF Port 2A to RF Backplane Port | 24.0 dB | 3.5 dB |
| RF Port 1A to RF Port 2A | 42.0 dB | 3.0 dB |

Table 4: System Fixed Loss for RFM Components

To achieve the hidden node setup all three wireless devices are placed in separate isolation chambers and allowed communicate over controllable channels. The combiners/splitters facilitate virtual distancing through varying common and port attenuators (view Figure 10).

Node A is connected to RF Port 1A, Node B is connected to RF Port 2A, and the AirEqualizer is connected to RF Backplane Port. The fixed loss between these devices can be seen in Table 4. The devices are connected by coaxial cables with specifications in Table 5.

| Component | Description |
|------------------------------|--|
| External RF Connection Cable | Coaxial cable with at least 100 dB ingress and 50 ohm characteristic impedance Total RF loss < 0.5 db |

Table 5: RF Connection Cable Specification

Additional photographs of the test chamber, AirEqualizer and clients are in the Appendix of this report.

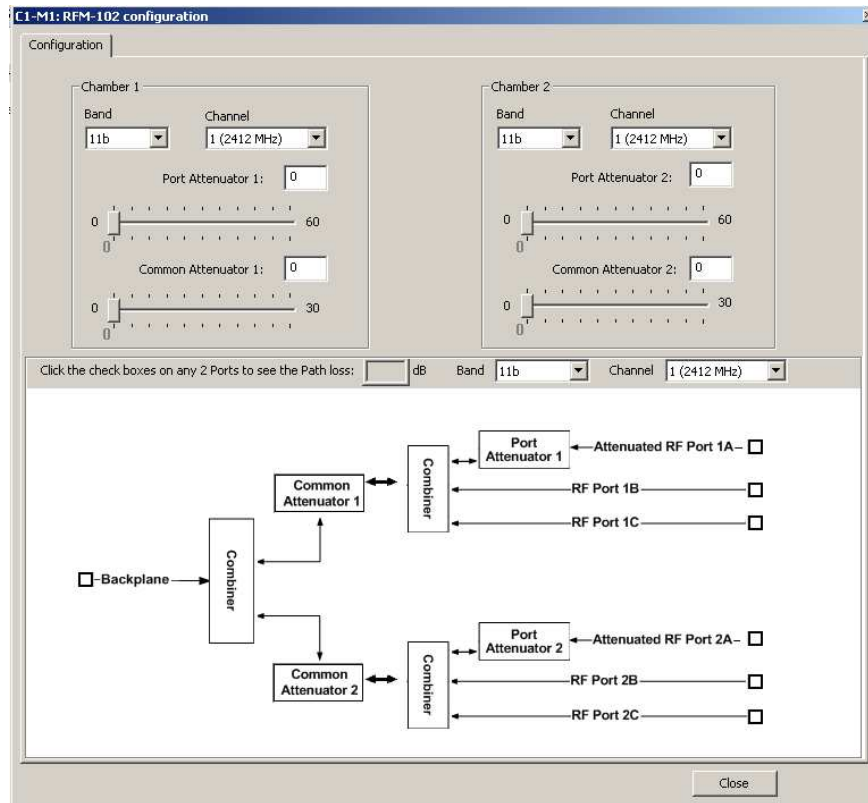


Figure 10: RFM Configuration Screenshot

3 Observations

The hidden node problem is evaluated in a reliable test environment. Node A is virtually positioned in close proximity to the access point and has a stronger signal than Node B based farther away. Node A initiates a large file download, following a fixed period of time Node B also attempts to download the same file. Tests are undertaken with the AirEqualizer turned off and on.

With the AirEqualizer turned off, as illustrated in Figure 11, the observations show Node A receiving as much bandwidth as needed for its download. However, due to Node A hogging all the bandwidth Node B, with its weaker signal at a greater distance, is locked out and cannot access the network. Thus, the wireless network is not fair to both users. This means Node B may have to wait until Node A has finished its download before accessing the network. In the real world, this is unacceptable for the user of a modern wireless communication network.

With the AirEqualizer turned on, as illustrated in Figure 12, the observations indicate a fairness for both users on the network. AirEqualizer determines detrimental flows from normal ones by taking the a number of questions into consideration: 1) How persistent is the flow? 2) How many active flows are there? 3) How long has the flow been active? 4) How much total congestion is currently on the trunk? 5) How much bandwidth is the flow using relative to the link size? In this case, the AirEqualizer penalises Node A when its bandwidth reaches 85 % (Ratio value) of the network capacity. This prohibits Node A from hogging all the bandwidth which allows Node B to access the network and begin downloading. Further calculations determine whether to continue penalising Node A or to remove the penalty so as allow both users an adequate amount of bandwidth.

The AirEqualizer shows fairness when it is turned on by prohibiting offending flows by adding latency, forcing them to back off and allow potentially hidden nodes to establish communications, thus reducing

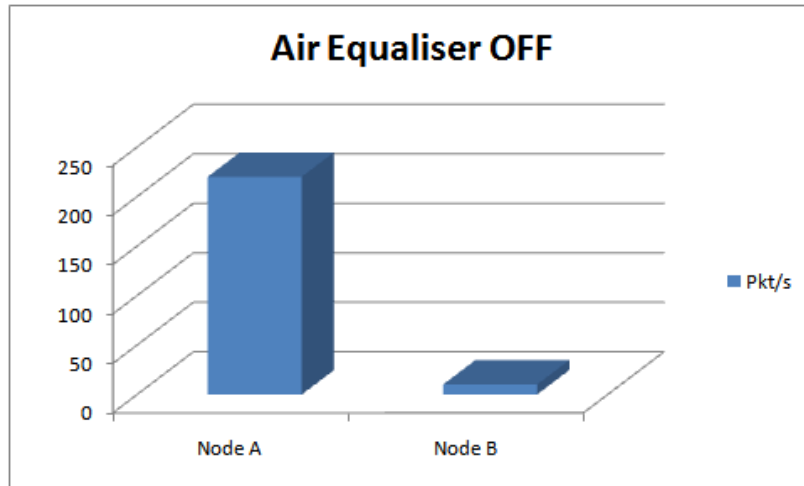


Figure 11: Air Equaliser OFF

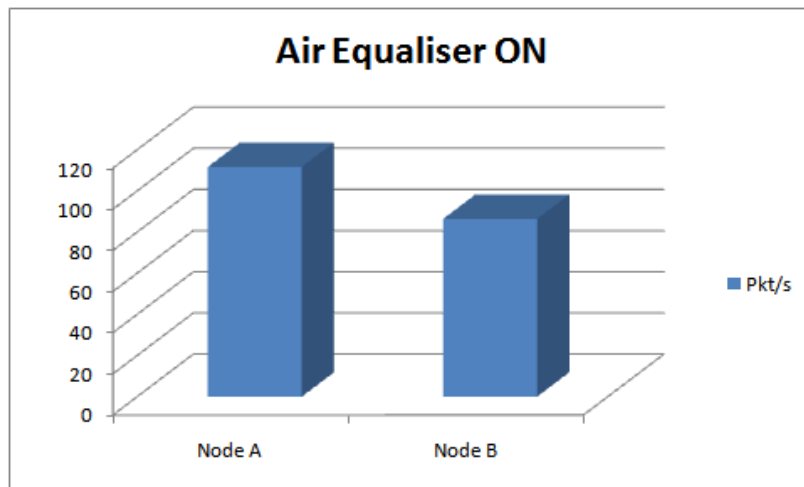


Figure 12: Air Equaliser on

network disruption.

References

- [80299a] 802.11a. Supplement to IEEE standard for information technology telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements. Part 11: wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: high-speed physical layer in the 5 GHz band," IEEE Std 802.11a-1999, 1999.
- [80299b] 802.11b. Supplement To IEEE Standard For Information Technology- Telecommunications And Information Exchange Between Systems- Local And Metropolitan Area Networks- Specific Requirements- Part 11: Wireless LAN Medium Access Control (MAC) And Physical Layer (PHY) Specifications: Higher-speed Physical Layer Extension In The 2.4 GHz Band, 1999.
- [80203a] *IEEE Std 802.11, 1999 Edition (R2003). IEEE Standard for Information Technology; Telecommunications and Information Exchange between Systems; Local and Metropolitan Area Networks; Specific Requirements; Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, June 2003.
- [80203b] 802.11g. IEEE standard for information technology- telecommunications and information exchange between systems- local and metropolitan area networks- specific requirements Part II: wireless LAN medium access control (MAC) and physical layer (PHY) specifications," IEEE Std 802.11g-2003 (Amendment to IEEE Std 802.11, 1999 Edn. (Reaff 2003) as amended by IEEE Stds 802.11a-1999, 802.11b-1999, 802.11b-1999/Cor 1-2001, and 802.11d-2001), 2003.
- [Abr85] N. Abramson. Development of alohanet. *IEEE Transactions on Information Theory*, 31(2):119–123, March 1985.
- [AS07] Inc. Azimuth Systems. Azimuth w-series wlan analysis platform. Documentation Release 4.2, Rev. v2.0.5, 2007.
- [BDP05] T. Bautts, T. Dawson, and G. N. Purdy. *Linux Network Administrators Guide*. O'Reilly and Associates, 2005.
- [CZ06] M.C. Chuah and Q. Zhang. *Design and Performance of 3G Wireless Networks and Wireless LANs*. Springer Science+Business Media Ltd, 2006.
- [DRScC05] Pradipta De, Ashish Raniwala, Srikant Sharma, and Tzi cker Chiueh. Design considerations for a multihop wireless network testbed. *IEEE Communications Magazine*, 43:102– 109, 2005.
- [DZ83] J.D. Day and H. Zimmerman. The osi reference model. *Proc. of the IEEE*, 71(12):1334–1340, 1983.
- [PPC95] K. Pahlavan, T. H. Probert, and M. E. Chase. Trends in local wireless networks. *IEEE Communications Magazine*, 33(3):88–95, March 1995.
- [SSG07] Ronan Skehill, Padriag Scully, and Sean Mc Grath. Characteristics, results and findings of ieee 802.11 in an rf isolated testbed. In *Proceedings of IEEE Personal, Indoor and Mobile Radio Communications*, pages 1–5, September 2007.
- [Tuc93] B. Tuch. Wavelan, an ism band wireless lan. *AT and T Technical Journal*, 72(4):27–37, August 1993.
- [wecge] WECA Homepage. <http://www.wi-fi.org>.

Section 3

Appendices



Figure 13: AirEqualizer in RF Chamber



Figure 14: RF Chamber: Closed



Figure 15: A client in the RF isolation Chamber