CHAPTER 2

WIRELESS LOCAL AREA NETWORKS

This chapter introduces several available standards and implementations of wireless networks plus the basic differences between them and also the problems arising from deploying such networks. Rapid advancement in this field of networking has resulted in numerous standards available in the market making it difficult to explain in detail every single aspect of wireless networking, so a basic but hopefully thorough review of the major protocols and functions is given.

2.1 Wireless Local Area Networks

A wireless LAN or wireless local area network (henceforth referred to as WLAN) is a network architecture that uses radio waves as a carrier instead of a physical medium found in traditional wired LANs to provide a means of communication between users/nodes. It is usually used to provide the last link between users and a larger network such as the Internet especially in places where physical cabling is either difficult or impossible (Stallings 2002, Forouzan 2003). After more than a decade since conception, WLANs are only now gaining tremendous momentum mainly due to falling costs and ever improving standards.

Several of the more prominent standards are described below but beforehand, it has to be made clear that although various terminology are used in the text references and sources, they describe the same thing. Common interchangeable terms include '*standards*' with '*protocols*', '*nodes*' with '*stations*' or '*terminals*' and '*band*' with '*range*'. Although consistency is mostly maintained throughout the text, at times the terms are exchanged for clarity purposes which should otherwise remain the same.

2.1.1 Current Wireless Standards

Initial development in WLANs resulted in multiple proprietary standards being developed amongst different companies without much interoperability offered between them. By the end of the 1990s, the confusion was rectified with several fixed standards namely the IEEE 802.11 and HomeRF of which other companies could adhere to. The 802.11 standard was focused towards providing wireless networks in the commercial industry whilst HomeRF targeted smaller home networks by providing cheaper means of implementing wireless networks. However, the adoption of the 802.11 standard in both commercial and home networks finally resulted in the recent disbandment of the HomeRF organisation in early 2003 and is now only used for research comparison purposes.

The original 802.11 standard (IEEE 1999a) provided a bandwidth of either one or two megabits per second (Mbps) in the 2.4GHz frequency range. This involved using either frequency hopping spread spectrum (FHSS) or direct sequence spread spectrum (DSSS) with radio waves to transmit. The 2.4GHz ISM (Industrial, Scientific and Medical) frequency band was chosen as it was unlicensed which allowed the user to broadcast within the frequency without having to register or

obtain a license but strict rules are present for broadcasting in the ISM frequencies to allow peaceful co-habitation with other systems also using the ISM bands.

When released in September of 1997, it was realised that the bandwidth provided by 802.11 was too small and saturated too quickly thus forming a bottleneck in the network especially when wired networks were running at a minimum of 10Mbps with 100Mbps (Fast Ethernet) becoming the norm. It was partly because of the low bandwidth performance that the 802.11 standard never really took off and plans were started to use the 5GHz range to increase throughput. However, the migration from 2.4GHz to 5GHz involved major changes to every part of the network's architecture, did not provide backward compatibility and was a hard sell to customers already using 802.11 devices. Therefore many companies tried to extend the life of their technology by focusing on increasing throughput whilst in the 2.4GHz band.

A little bending of the spread spectrum rules enabled increased data rates of 5 and 11Mbps and the 802.11 standard was saved. In 1999, the standard was ratified and its extension labelled IEEE 802.11HR (High Rate) was released. This standard (IEEE 1999b, 2001), now more commonly known under the IEEE 802.11b name, allowed for data rates up to 11Mbps with a fallback to 5.5, 2 and 1Mbps for backward compatibility with the original standard. Other specifications of the standard remained more or less the same but the FHSS modulation scheme was dropped allowing only DSSS with CCK to be used to modulate the data signal. During this time, the use of OFDM for transmission was not yet approved for use by the FCC in the unlicensed frequency bands and this restricted the maximum speed to 11Mbps due to technical reasons. Nevertheless, soon after 802.11b was released wireless

networking became popular no thanks to adoption by manufacturing companies such as Intel to use in their *Centrino* based mobile devices.

The upcoming standard is IEEE 802.11g (IEEE 2003) which further increases the data rate up to a theoretical 54Mbps whilst still remaining in the 2.4GHz ISM band. This substantial increase in bandwidth is due to the allowance of the FCC to use OFDM in the 2.4GHz range. With OFDM, the problem of multipath interference and fading is reduced and a higher number of bits can be delivered per clock cycle. To maintain compliance to the earlier ISM rules however, 802.11g reverts to using DSSS modulation at lower data rates similar to the 802.11b standard and this contributes to its backward compatibility.

Another competing standard to offer speeds of up to 54Mbps is the IEEE 802.11a (IEEE 1999c). Discussions for the 802.11a standard began early as a replacement to the now obsolete original 802.11 standard and continue to this very day. Using part of the UNII frequencies with OFDM modulation, it is theoretically capable of higher throughput and faces less risk of interference with other devices. The UNII sets of frequencies are located in the higher 5GHz range and are divided into 3 frequency ranges: 5.15-5.25GHz, 5.25-5.35GHz and 5.725-5.825GHz range. Since less devices use this range, the risk of unwanted band interference is reduced.

Unfortunately, the future adoption of 802.11a in WLANs might be a difficult one due to many factors. First is the lack of interoperability with older wireless networks. Any user wishing to deploy 802.11a networks will have to upgrade every single device they own unless they are willing to maintain multiple networks running

simultaneously. The increase in transmission frequency also results in the decrease of coverage area down to approximately 60 feet radius per station because higher frequencies are more 'opaque' to obstacles and are easily blocked. In comparison, 2.4GHz signals can sometimes penetrate brick walls and reach distances over 300 feet. Due to this range decrease, a higher number of access points are also required to cover an equal area making 802.11a more expensive to deploy (1 access point for every 50 feet radius compared with one every 200 feet for 802.11b/g).

The next major problem is manufacturing cost and compatibility. Since the manufacturing of 802.11a devices is still limited and demand is relatively low, the price remains high since they are not produced in bulk. Many OEM manufacturers are shying away from 802.11a chipsets until the demand rises and the cost is reduced to ensure profitability. As of now, reports of interoperability issues have also surfaced where a 802.11a device produced by one OEM is unlikely to work with another OEM's device even though they both comply to the standard (Gain 2002).

The final hurdle to the standard appears to be due to provincialism - the IEEE 802.11a standard only applies to the US territories (and those that follow their standards) and will not work in other countries such as Europe which has its own implementation called *HiperLan/1* and *HiperLan/2* thus making it difficult for companies to justify producing wireless devices for each.

The final wireless standard is *Bluetooth* (BSIG 2003) but this particular standard is meant entirely for a different purpose and not to replace the traditional LAN but alongside it instead. Led by the Bluetooth Special Interest Group (B-SIG) with

industry members including Sony, Ericsson, Nokia, IBM and Intel, Bluetooth technology was designed to create wireless personal area networks (WPAN) for short distances of about 10 metres between personal electronic devices placed in close proximity of each other.

Bluetooth capabilities are often found in small peripheral devices such as printers, mobile phones and input devices such as mice and keyboards. This eliminates the need to have various types of cables and connections plus avoid the mess with port and interrupt settings to connect devices to a workstation since they all now use the same connection which is auto-configured. With Bluetooth, it is possible to create and entirely wire-free computing environment. For example, in the diagram below, the mobile station gets print and Internet services from the main station server (an access point) using 802.11b networking while connections to external peripherals use Bluetooth. The mobile phone provides a connection to the Internet perhaps with GPRS, thus creating a gateway for the mobile station.

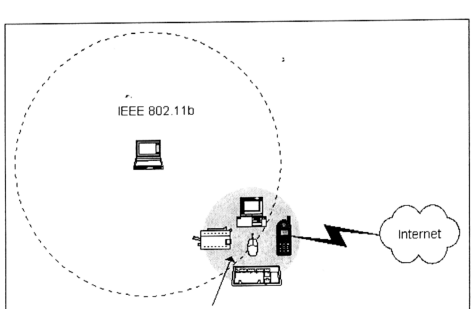

Figure 2.1 – Bluetooth WPANs with IEEE 802.11b WLANs

Special mention has to be made of the Bluetooth standard as it inhabits the same ISM band and uses fairly similar modulation techniques used by several 802.11 standards which result in several problems to be discussed later in this chapter.

A summary of the standards discussed above is listed in the table on the following page. Note that the typical transmission data rate based on actual usage in real network environments is stated when and where available. In addition, the maximum transmission range specifications given are usually affected by environment and physical factors and vary significantly from manufacturer to manufacturer. Thus the values shown are given as an approximation only.

Table 2.1 – Available WLAN standards

Standard	Data rate (typical rate)	Modulation scheme	Max. range	Advantages	Disadvantages
HomeRF	Up to 10Mbps in 2.4GHz band (N/A)	FHSS	150 feet (50m)	- Relatively inexpensive - Good for telephony services	- Less interference - No further development
Bluetooth	Up to 10Mbps in 2.45GHz band (~433 - 720kbps)	FHSS	30 feet (10 m)	- Best suited for connections of short intervals	- No IP or TCP/IP support - Short range
IEEE 802.11	Up to 2Mbps in 2.4GHz band (~2Mbps)	FHSS or DSSS	750 feet (250m)	- See 802.11b	- Low data rates - RF Interference
IEEE 802.11b	Up to 11Mbps in 2.4GHz band (~6Mbps)	DSSS with CCK	300 feet (150m)	- Cheap - High market penetration	- High RF interference - Lack of available channels
IEEE 802.11g	Up to 54Mbps in 2.4GHz band (~25Mbps)	OFDM above 20Mbps, DSSS with CCK otherwise	300 feet (150m)	- Backward compatibility with 'b' plus increased speed	- Still susceptible to RF interference
IEEE 802.11a	Up to 54Mbps in 5GHz band (~30Mbps)	OFDM	60feet (20m)	- Less interference in higher band - Better support for multimedia requirements	- Low market penetration - Not interoperable with 'b' or 'g'

2.1.2 Components in the IEEE 802.11x WLAN Topology

Since this project focuses mainly on the IEEE 802.11x standards which are natively supported and can be simulated using the network simulator, the subsequent sections will only describe those. Although the following text mainly refers to the original 802.11 standard, it applies just as well to the newer 802.11b standard (and to a certain extent, the 802.11g standard too).

As specified in the referenced 802.11 standard (IEEE 1999a), there are several components that interact together to construct a WLAN with support for station mobility transparently to the upper layers. The fundamental component is the basic service set (BSS) which consists a minimum of two stations (any device with suitable wireless networking hardware installed) in communicating range as shown below. The ovals represent the coverage area of which the member stations of a BSS may remain in communication. Once a station moves out of a BSS, it can no longer communicate with other members in the BSS.

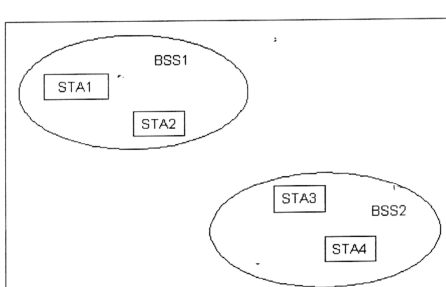

Figure 2.2 – Two independent BSS in an IEEE 802.11 wireless network

An *independent BSS* (IBSS) forms the most basic type of 802.11 WLAN called an *ad-hoc* network (sometimes called peer-to-peer, on-demand or structure-less network). In this type of network, two or more stations within range can set up an independent network without requiring any pre-planning, additional administration or configuration. Mobile stations in ad-hoc networks communicate directly with each other and all connections are considered dynamic and are done in an arbitrary manner. Every station behaves as routers and takes part in the discovery and maintenance of routes to other stations in the network/BSS. Although initial connections between stations may be initiated by one party, once a connection is made, both communicating stations work as peers instead of a server-client type hierarchical relationship.

Instead of existing independently, a BSS may collectively form an *extended service set* (ESS) with the addition of the *distribution system* (DS) component. The DS forms the backbone connection between the BSS components and is typically a wired network (i.e. Ethernet) or in some cases wireless itself. Combining the BSS and DS components allows the construction of ESS networks (sometimes known as infrastructure or base-connected mode) of arbitrary size and complexity.

The access to the DS in WLANs is usually through an access point (AP), which can be a station in itself, that provides the DS services and are addressable entities. These access points are responsible for providing logical services to handle address destination mapping and seamless integration of multiple BSS. The figure on the following page represents how a DS connects two BSSs into an ESS network.

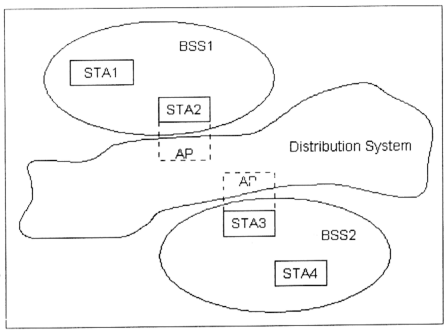

Figure 2.3 – ESS of BSSs using a distribution system

The ESS type of network architecture is usually set up when an increased area of coverage or multiple connectivity with other LANs (wired or wireless) is required. To the stations of the ESS, the network still appears the same to the LLC layer as an IBSS network and stations may communicate and move from one BSS to another (within the ESS) transparently.

Connectivity with wired LANs is usually done through the final component of the WLAN called a *portal* which allows MSDUs originating from a non-IEEE 802.11 LAN to enter the DS. The task of the portal and AP is now commonly combined into a single device such as in a wireless router to offer both functions but it is still referred to as an access point only.

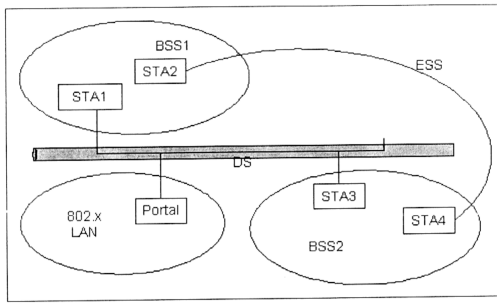

Figure 2.4 – Interconnecting various IEEE 802 LANs

In most WLANs, the AP acts as a bridge to allow communications between mobile stations within its communication range. The AP or base station is usually stationary in the network whilst the stations continue to move geographically all the time communicating through the AP. When a mobile node goes out of range of one base station, it attempts to connect to another (based on characteristics such as signal strength, channel availability) and continues communicating through it in a process called a *handoff*.

2.1.3 Anatomy of IEEE 802.11 Standard

Since the beginning, the IEEE 802.11x standards were designed to co-exist with current networks and appear transparent to the applications running on it. This requires that the station handle the mobility issue within the MAC sublayer and thus require 802.11 to incorporate functionality untraditional for MAC sublayers to meet the reliability assumptions of the LLC layers above. Following any 802.x type

standard, the 802.11 standard therefore only covers the MAC and physical layers only. The relationship between the two layers is illustrated in the figure below.

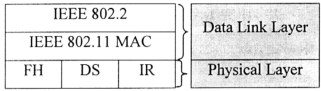

Figure 2.5 – Layers in IEEE 802.11 standard

IEEE 802.11 MAC Layer

The 802.11 standard defines a single MAC type with two access methods to the physical layer namely the Distributed Coordination Function (DCF) and Point Coordination Function (PCF). The DCF is known as *Carrier Sense Multiple Access with Collision Avoidance* (CSMA/CA) and is the fundamental access method for IEEE 802.11 MAC. It is implemented in all wireless stations for use in IBSS and ESS type networks. CSMA type protocols are common in networking where the most popular implementation is in Ethernet with its CSMA/CD protocol (CD for collision detection). Basically, CSMA for WLANs operates under a simple logic:

- A station wishing to transmit senses the transmission medium for activity.
- If the medium is busy (i.e. another station is transmitting) then the station defers its transmission for an arbitrary period.
- After this period, if the medium is free, then the station proceeds with its transmission.
- If the medium remains busy, the waiting period is increased and process repeats until the maximum amount of times is reached.

The CSMA protocols are effective in low traffic situations but if the number of stations and transmissions increase, there is a corresponding increase in chance of multiple stations transmitting at the same time because the stations sensed the medium free and decided to transmit at once. This results in a collision of data packets for the involved stations and thus necessitating retransmissions. Collision situations like these are easily detected on wired networks which then proceed to a retransmission phase with an exponential backoff algorithm before retransmitting the packets. On the other hand, the situation in WLANs is different and collision detection cannot be used because of two reasons (Breezecom 1997):

1. All IEEE 802.11x standards employ only half-duplex radios but the collision detection mechanism requires a full duplex radio transceiver capable of transmitting and receiving at the same time. This is something that would increase the hardware price significantly.
2. Since there is no physical medium, not all stations can hear each other even though they are on the same network. A station wishing to transmit and sensing a medium free doesn't necessarily mean the medium at the receiver is free. Situations like these are called either the "hidden node" or "exposed node" depending on the stations' relative locations.

The hidden node situation occurs when stations are not in receiving range of each other but attempt to transmit simultaneously. Conversely, a node within range of a transmitting node suffers from the exposed node problem where the NAV is set to busy even though the station itself is not transmitting or receiving but because it overhears a transmission. These two situations are illustrated in the figure on the following page. The two source stations A and B are out of range of each other thus

making each node 'hidden' from each other, so when A begins transmission to node C, B does not realise this and might attempt to transmit to C simultaneously. Station D is within range of station A and thus overhears the RTS of A and sets itself as busy for the allocated transmission period. This makes D an exposed node and not able to transmit or receive packets for the duration.

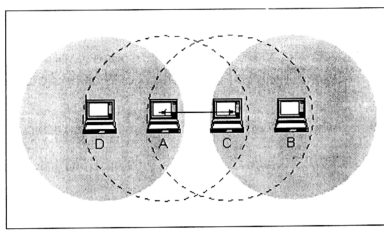

Figure 2.6 – The hidden and exposed node problems

To help overcome these problems, the IEEE 802.11 standard attempts to avoid the collision situation entirely with Collision Avoidance (CA) which basically works as follows:

- A station wishing to transmit senses the medium for any activity.
- If the medium is free, it waits for a specified amount of time (a DIFS period explained below) before re-sensing the medium again for transmission activity.
- If the medium remains free, the station is allowed to transmit.
- If the medium is busy during the second test, the station withdraws the transmission and the time period before the next attempt is increased in the backoff process.

The backoff process computes a random amount of time between zero and a maximum amount called the *Contention Window* (*CW*) and is used to initialise a backoff timer. The *CW* is incremented (left binary shift plus one) each time a station reattempts a transmission until a predetermined maximum value called the *CWMax*. This timer is decremented when the medium is idle again but remains frozen while the medium is busy (i.e. another station transmits). This decrementing period of the timer is called the *Slot time* which corresponds to the maximum roundtrip delay within a BSS.

This CA mechanism is combined with a positive acknowledgement scheme for transmitted packets. The receiving station performs a CRC check on the packet received and sends an acknowledgement packet (ACK) if it was successfully received. If the sender does not receive any acknowledgement, it will retransmit the packet fragment until it receives the ACK or until it reaches its retransmission threshold at which the packet is thrown away. The transmission of the ACK requires a time period equal to SIFS after the end of reception of the previous frame. Since the SIFS period is always smaller than DIFS, the receiving station need not check the medium for any activity before transmitting the ACK.

In situations where stations cannot hear each other, the IEEE 802.11 standard defines a Virtual Carrier Sense (VCS) mechanism to reduce the probability of transmissions from stations colliding. Stations wishing to transmit must first transmit a short control packet called RTS (Request To Send) which includes information about the source and destination stations and the duration required for the transmission. The destination stations replies (if the medium is free) with a response control packet

called CTS (Clear To Send) which includes the same duration information. Stations that are not involved in the transmission but manage to overhear this exchange of RTS/CTS packets set their VCS indicator (called the Network Allocation Vector or NAV) for the given duration of the transmission for use together with the carrier sense mechanism when sensing the medium.

The NAV state combines with the physical CS to indicate the medium is in a busy state. The RTS/CTS exchange also helps avoid the "hidden node" problem when stations cannot hear each other. The figure below illustrates the progression order of packets during a transaction between two stations and the packets overheard by the surrounding neighbours. Neighbours that overhear a RTS or CTS packet not destined for it set their NAV value to busy for time period defined in the packet.

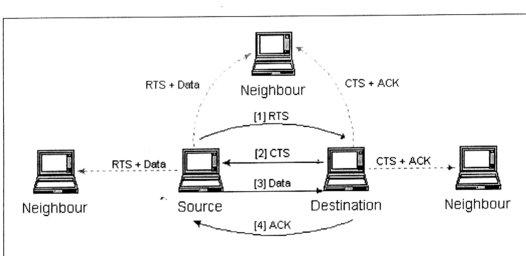

Figure 2.7 – RTS/CTS exchange between nodes

IEEE Inter Frame Periods

The IEEE 802.11 standard defines 4 Inter-frame spaces and one slot time used to provide different times for different priorities that a station waits before transmitting:

- **Slot time** is the amount of time a device waits after a collision before retransmitting a packet. Short slot times decrease the backoff time, which increases throughput.
- *SIFS* or **Short Inter Frame Space** separates transmissions in a single transaction (RTS – CTS – Fragment – ACK) and is the minimum Inter Frame Space available. There is always at most one single station to transmit in this period of time thus having priority over all other stations. The SIFS value is fixed such that a transmitting station has time to switch over to receive mode to decode an incoming packet even though it is waiting to transmit. The IEEE 802.11 standard sets this value at 28 microseconds.
- *PIFS* or **Point Coordination Inter Frame Space** is used by an access point (or Point Coordinator in this case) to gain access to the medium before any other station. The value of PIFS is equal to SIFS plus one slot time.
- *DIFS* represents the **Distributed Inter Frame Space** and is used for stations wishing to start a new transmission. The value it calculated as PIFS plus one slot time (or SIFS plus two slot times).
- *EIFS* is the **Extendcd Inter Frame Space** and is the longest IFS used by a station that received a packet that it could not understand.

The following diagram shows a typical transaction between two stations with the interframe periods and the NAV settings of the affected neighbours that overhear the transmission. The amount of slot time periods depend on the contention window value up to a certain maximum as described earlier.

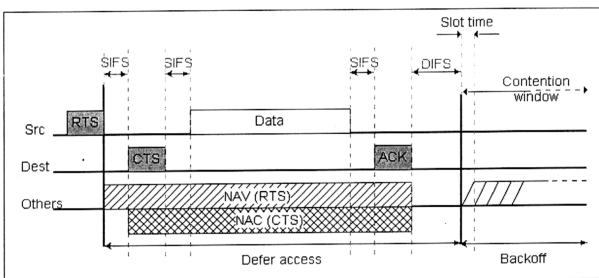

Figure 2.8 – Interframe periods between transmissions

The other mode of MAC access called PCF is only available in infrastructure mode and is essentially a polling method of medium access. Operating from the base station or AP of a ESS, the AP controls access to the channel by determining which station has the right to transmit at a particular time. A poll packet is sent from the AP to a station permitting that station to transmit whilst the others wait to receive its own poll packet. Poll packets can be sent out according to priority, in round robin fashion or based on reservation requests made by stations wishing to transmit but as the number of nodes in a BSS increases, the overhead involved in sending poll and request packets often degrade the network performance.

Both methods of MAC access were designed to coexist such that both can operate concurrently within the same BSS with a station alternating between them but since

CSMA/CA is able to support ad-hoc and infrastructured networking plus it offers comparable performance, it is more commonly implemented in 802.11 devices.

IEEE 802.11 Physical Layer

IEEE 802.11 standard defines three different physical layer characteristics to support the MAC layer: one using infrared diffusion and two using RF transmission methods. Each definition provides the following two protocol functions:

1. *Physical layer convergence* (PLCP) function which maps the incoming PDUs from the MAC layer into a format suitable for transmission between stations on the PMD
2. *Physical medium dependant* (PMD) system whose function is to define the characteristics of and method of transmitting and receiving data through a wireless medium between two or more stations.

The regular method of access in 802.11 networks is using RF and since this project focuses only on RF type transmissions, the PMD systems shall be explained in further detail (Gast 2002).

The 802.11 standard transmissions occupy the ISM 2.4GHz band and is bound to the rules set by the regulating organisations to avoid abusing the band. These organisations include the *FCC* and *IC* for North America and Canada, *ETSI* for Europe, *ARIB* for Japan and *DGPT* for France. The rules include the compulsory use of spread spectrum transmission, maximum transmission power and channel definition but each organisation may determine a subset of the original standard to suit the region it represents. For example, the transmission power in the band is set by the FCC at maximum 1W but the ETSI sets the maximum at only 100mW whilst

the maximum allowed antenna gain remains at 6dB for both. Spread spectrum is enforced to allow devices to share the available bandwidth instead of hogging it all to oneself and also to help reduce the impact of localised interferences on the transmission. This is done by using more bandwidth than necessary for a transmission and prevents any one system from using it to its full capacity. Two methods of spread spectrum are specified namely DSSS and FHSS.

Direct Sequence Spread Spectrum

The ISM 2.4GHz frequency band covers a bandwidth of 83MHz from 2.4 – 2.483GHz. In DSSS, this bandwidth is divided into 14 different channels each separated by 5MHz but not all channels can be used depending on country regulations as shown below.

Table 2.2 – IEEE channel allocation for DSSS in ISM band

Channel number	*Frequency*
1	2412
2	2417
3	2422
4	2427
5	2432
6	2437
7	2442
8	2447
9	2452
10	2457
11	2462
12	2467
13	2472
14	2477

Maximum number of channels allowed
USA, Canada and Taiwan – 11 channels (1-11)
Europe – 13 channels (1-13)
Spain – 2 channels (10-11)
France – 4 channels (10-13)
Japan – 1 channel (14)

DSSS *spreads* a signal over a larger band by combining it with a higher data-rate bit sequence that divides the user data according to a spreading ratio. This bit sequence (sometimes called a chipping code or signature) is a redundant bit pattern for each bit being transmitted to increase the signal's resistance to interference by minimising localised interference and background noise. At the receiving end, a demodulator uses the same chip code to obtain the transmitted signal. Any narrowband interference will appear smaller because the multiplexing at both ends result in a processing gain of about 10dB (Tourrilhes 2000). The initial 802.11 standard used an 11-chip *Barker* code to spread the 2Mb signal over 22MHz of bandwidth. At higher data rates of 11 and 5.5Mbps different chip codes are used but the spread signal still accommodates roughly the same 22MHz bandwidth (Zyren & Petrick 2004).

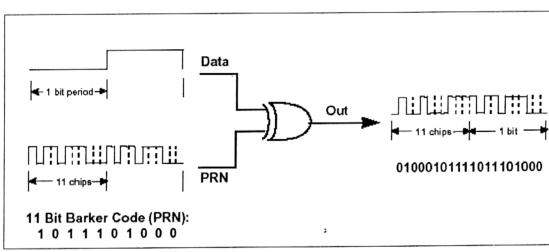

Figure 2.9 – Modulation with 11-chip Barker code

Frequency Hopping Spread Spectrum

FHSS takes a different approach to spreading the signal by using a set of narrow channels and *hopping* through all of them in sequence to transmit. The 2.4GHz ISM band is divided into multiple channels of 1MHz each and periodically, the system hops to a new channel following a predetermined cyclic hopping pattern. The number of channels once again depends on the region and the approval of the

regulatory organisations. The table below provides the number of FHSS operating channels for several major regions as listed in the IEEE 802.11 standard.

Table 2.3 – IEEE channel allocation for FHSS in ISM band

Geography	*Hopping set channels*	*Regulatory range*
USA (North), Canada	79	2.400 – 2.4835 GHz
Europe (except Spain and France)	79	2.400 – 2.4835 GHz
Japan	23	2.471 – 2.497 GHz
Spain	27	2.445 – 2.475 GHz
France	35	2.4465 – 2.4835 GHz

FHSS avoids interference by never staying in a channel throughout the transmission; if a channel is bad, the system simply waits until it hops to a good one with the result of averaging the impact of bad and good channels over time. The hop sequence is pre-determined based on a series of equations to divide the available channels into three non-overlapping sets. A station only uses the channels and hop sequence of a particular set only to transmit to another station also using the same set.

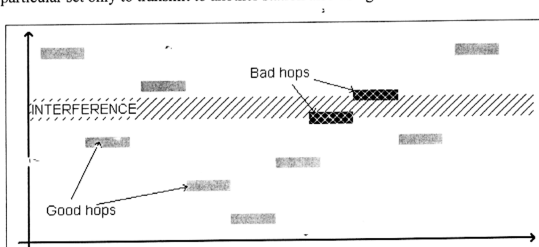

Figure 2.10 – Interference in frequency hopping spread spectrum

The two methods of transmission each have their own advantages and disadvantages. DSSS is simpler to handle at the MAC since transmission only uses a single channel at any time but the number of available channels is limited and overlapping often occurs resulting in a proportional increase in noise. FHSS has the advantage of experiencing poor transmissions only if it hops into a band with interference whilst DSSS continues to experience interference as long as it stays in the same channel range as the interfering band. On the other hand, FHSS is complicated at the MAC layer because it has to scan the network at initialisation and maintain synchronicity of the hops.

Multiple WLANs using FHSS might result in periodic collisions as the hopping patterns collide but this phenomenon is intermittent and dependent on the inter-sequence timing within the FHSS sets. Consequently FHSS can support more than three co-located channels with a reduction in performance but DSSS is unable to do so because co-located channels suffers from constant interference throughout the transmission.

In the 802.11b standard however, the FHSS modulation scheme is dropped at higher data rates but remains valid at 1 and 2Mbps data rates to provide backward compatibility. DSSS with CCK is used instead with a spread code/signature length of 8 based on complementary codes to increase the number of bits multiplexed per time.

2.2 Problems with Wireless Transmissions

Due to the medium-less type of transmission, WLANs suffer from many problems which do not affect the traditional wired LAN. The most obvious of which is the problem of station mobility which is both the feature of having a wireless network and also its downfall. Unlike a wired LAN where each node is presumed to stay at a static location that is addressabie, wireless stations are free to roam and fall into and out of communication range. Station association with an AP is dynamic and broken routes between stations could take ages to recover resulting in unnecessary packet drops and unwanted queue build-ups whilst a route is being re-discovered.

Another obvious problem with WLANs is the security of transmissions since there is no physical medium boundary; the wireless signal is ever present for anyone to intercept (Flickenger 2003a). The process of eavesdropping has nowadays even evolved into the pastime of *war-driving* in which a person drives around in a vehicle with a suitably equipped wireless device and a powerful antenna mounted on the roof scanning for stray unprotected signals. An attempt to increase security was made when Wired Equivalent Privacy (WEP) was incorporated to encrypt the communication signals between stations but this feeble effort was overcome as the encryption was not strong enough (mainly due to 40-bit encryption key limit restriction set by the US government at that time) and easily cracked (Flickenger 2003b, Raghaei 2003). Since then other methods of securing the WLAN have emerged including hardware MAC address filtering to allow only known users access to the network, SSID hiding and WiFi Protected Access (WPA) to authenticate users on the network.

this band and this is unavoidable as long as the equipment complies with Part 15 of the FCC Rules for unlicensed RF devices (FCC 2004). Packets from WLANs that collide with this interference signal are corrupted and cannot be received. Only if the SNR between the packet and interferer is high enough can the packet be successfully captured.

Radio waves attenuate, reflect or diffract when faced with obstacles of different materials much like light is amplified by a magnifier, diffracted by glass and reflected by a mirror. In real world environments, they can either bounce off or go through walls and ceilings depending on thickness and materials. These properties of radio waves make it hard to estimate the actual range of a system and also creates an environment where a signal can come from many different directions because of deflections and with different strengths depending on attenuation. This phenomenon is called multipath and creates the problem of fading and delay spread at the receiver (Tourrilhes 2000). Depending on the number of reflections and delay of the signals arriving at the receiver, the overlapping summation of the signals could cause destructive interferences.

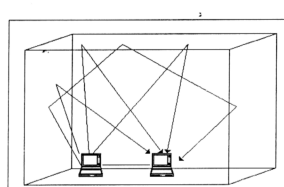

Figure 2.12 – Multipath fading and delay spread

2.3 Performance Issues in Wireless Networks

The advantage of mobility and a medium-less transmission method in wireless networks also becomes its own disadvantage. The performance of mobile networks, especially with TCP transmissions, is known to be poor when compared to their wired counterparts. Even though TCP is a reliable, full-duplex, connection-oriented protocol, its flow and congestion control mechanisms are based upon the assumption that packet loss is an indication of congestion. While this may be true in wired networks, wireless networks suffer additional packet losses due to other reasons too.

Traffic congestion is one of the major factors that lead to poor performance in wireless networks. As more and more wireless networking zones are being set up, the possibility of the zones overlapping will increase and there is bound to be a noticeable increase in the amount of traffic in the 'medium' and thus result in a lower performance due to excessive collisions. This drop is compounded with the fact that the frequency band used by most of the current wireless protocols is unlicensed and is also co-inhabited by many other devices leading to unwanted interference and weaker transmission rates.

Another factor leading to the poor performance is the mobility of wireless stations/nodes themselves. As earlier mentioned, mobility often results in link failures and this in turn causes packet loss. Once the link is considered unavailable (even though it still is but busy) the MAC protocol will report a link failure to the routing layer. This invokes the routing layer to plot new routes to the destinations affected by the failed link. Ideally, this condition only occurs when the nodes move

out of range of each other, but congestion also may induce a similar situation and result in what is known as *false link failures*.

The false link failure occurs when the MAC layer at a node N_A announces to the routing layer that the link to a neighbour N_B is broken even though it is still within its transmission range. Wireless nodes can overhear transmissions that occur up to quite a distance away while it still occurs within the interference range of the node. The nodes are then blocked from performing any form of transmission after the NAV of the node is set to busy and will proceed to ignore any control packets that they may receive. The RTS-CTS handshake then fails to establish any connection because N_B cannot respond to N_A's multiple RTS transmissions after it overhears another transmission within its range.