

IEEE 802.11 Technical Tutorial

Introduction

The purpose of this document is to give technical readers a basic overview of the new 802.11 Standard, enabling them to understand the basic concepts, principle of operations, and some of the reasons behind some of the features and/or components of the Standard.

The document does not cover the entire Standard and does not provide enough information for the reader to implement an 802.11 compliant device (for this purpose the reader should refer to the Standard itself).

IEEE 802.11 Architecture

Architecture Components

An 802.11 LAN is based on a cellular architecture where the system is subdivided into cells. Each cell (called **Basic Service Set**, or **BSS**, in the 802.11 nomenclature) is controlled by a Base Station (called **Access Point** or, in short, **AP**).

Although a wireless LAN may be formed by a single cell, with a single Access Point, (and as will be described later, it can also work without an Access Point), most installations will be formed by several cells, where the Access Points are connected through some kind of backbone (called **Distribution System** or **DS**). This backbone is typically Ethernet and, in some cases, is wireless itself.

The whole interconnected Wireless LAN, including the different cells, their respective Access Points and the Distribution System, is seen as a single 802 network to the upper layers of the OSI model and is known in the Standard as **Extended Service Set (ESS)**.

The following diagram shows a typical 802.11 LAN including the components described above:




Figure 1: A Typical 802.11 LAN

The standard also defines the concept of a **Portal**. A portal is a device that interconnects between an 802.11 and another 802 LAN. This concept is an abstract description of part of the functionality of a “translation bridge”.

Even though the standard does not necessarily request it, typical installations will have the AP and the Portal on a single physical entity. This is also the case with BreezeCOM’s AP which provides both functions.

IEEE 802.11 Layers Description

As any 802.x protocol, the 802.11 protocol covers the MAC and Physical Layer. The Standard currently defines a single MAC which interacts with three PHYs (all of them running at 1 and 2 Mbit/s) as follows:

- Frequency Hopping Spread Spectrum in the 2.4 GHz Band
- Direct Sequence Spread Spectrum in the 2.4 GHz Band, and
- InfraRed

802.2			Data Link Layer
802.11 MAC			
FH	DS	IR	PHY Layer

Beyond the standard functionality usually performed by MAC Layers, the 802.11 MAC performs other functions that are typically related to upper layer protocols, such as Fragmentation, Packet Retransmissions, and Acknowledges.

The MAC Layer

The MAC Layer defines two different access methods, the Distributed Coordination Function and the Point Coordination Function:

The Basic Access Method: CSMA/CA

The basic access mechanism, called the **Distributed Coordination Function**, is basically a Carrier Sense Multiple Access with Collision Avoidance mechanism (usually known as **CSMA/CA**). CSMA protocols are well-known in the industry, the most popular being the Ethernet, which is a CSMA/CD protocol (CD standing for Collision Detection).

A CSMA protocol works as follows: A station desiring to transmit senses the medium. If the medium is busy (i.e. some other station is transmitting) then the station defers its transmission to a later time. If the medium is sensed free then the station is allowed to transmit.

These kinds of protocols are very effective when the medium is not heavily loaded since it allows stations to transmit with minimum delay. But there is always a chance of stations simultaneously sensing the medium as being free and transmitting at the same time, causing a collision.

These collision situations must be identified so the MAC layer can retransmit the packet by itself and not by upper layers, which would cause significant delay. In the Ethernet case this collision is recognized by the transmitting stations which go into a retransmission phase based on an **exponential random backoff** algorithm.

While these Collision Detection mechanisms are a good idea on a wired LAN, they cannot be used on a Wireless LAN environment for two main reasons:

1. Implementing a Collision Detection Mechanism would require the implementation of a Full Duplex radio capable of transmitting and receiving at once, an approach that would increase the price significantly.
2. In a Wireless environment we cannot assume that all stations hear each other (which is the basic assumption of the Collision Detection scheme), and the fact that a station wants to transmit and senses the medium as free doesn't necessarily mean that the medium is free around the receiver area.

In order to overcome these problems, the 802.11 uses a Collision Avoidance (CA) mechanism together with a Positive Acknowledge scheme, as follows:

1. A station wanting to transmit senses the medium. If the medium is busy then it defers. If the medium is free for a specified time (called Distributed Inter Frame Space (DIFS) in the standard), then the station is allowed to transmit.

- The receiving station checks the CRC of the received packet and sends an acknowledgment packet (ACK). Receipt of the acknowledgment indicates to the transmitter that no collision occurred. If the sender does not receive the acknowledgment then it retransmits the fragment until it receives acknowledgment or is thrown away after a given number of retransmissions.

Virtual Carrier Sense

In order to reduce the probability of two stations colliding because they cannot hear each other, the standard defines a Virtual Carrier Sense mechanism:

A station wanting to transmit a packet first transmits a short control packet called **RTS** (Request To Send), which includes the source, destination, and the duration of the following transaction (i.e. the packet and the respective **ACK**), the destination station responds (if the medium is free) with a response control Packet called **CTS** (Clear to Send), which includes the same duration information.

All stations receiving either the RTS and/or the CTS, set their **Virtual Carrier Sense** indicator (called **NAV**, for **Network Allocation Vector**), for the given duration, and use this information together with the Physical Carrier Sense when sensing the medium.

This mechanism reduces the probability of a collision on the receiver area by a station that is “hidden” from the transmitter to the short duration of the RTS transmission because the station hears the CTS and “reserves” the medium as busy until the end of the transaction. The duration information on the RTS also protects the transmitter area from collisions during the ACK (from stations that are out of range of the acknowledging station).

It should also be noted that, due to the fact that the RTS and CTS are short frames, the mechanism also reduces the overhead of collisions, since these are recognized faster than if the whole packet was to be transmitted. (This is true if the packet is significantly bigger than the RTS, so the standard allows for short packets to be transmitted without the RTS/CTS transaction. This is controlled per station by a parameter called **RTS Threshold**).

The following diagrams show a transaction between stations A and B, and the NAV setting of their neighbors:

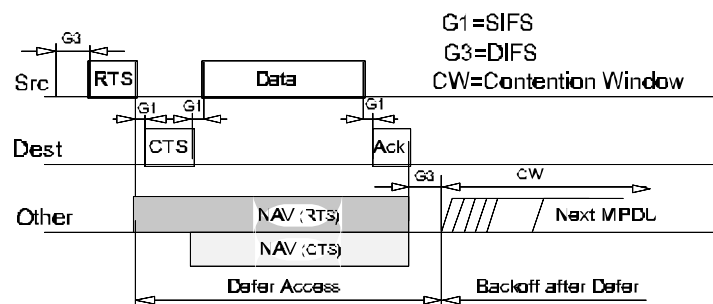


Figure 2: Transaction Between Stations A and B

The NAV State is combined with the physical carrier sense to indicate the busy state of the medium.

MAC Level Acknowledgments

As mentioned earlier in this document, the MAC layer performs Collision Detection by expecting the reception of an acknowledge to any transmitted fragment (Packets that have more than one destination, such as Multicasts, are not acknowledged.)

Fragmentation and Reassembly

Typical LAN protocols use packets several hundred bytes long (the longest Ethernet packet could be up to 1518 bytes long). There are several reasons why it is preferable to use smaller packets in a Wireless LAN environment:

- Due to the higher Bit Error Rate of a radio link, the probability of a packet getting corrupted increases with the packet size.
- In case of packet corruption (either due to collision or noise), the smaller the packet, the less overhead it causes to retransmit it.
- On a Frequency Hopping system, the medium is interrupted periodically for hopping (in our case every 20 milliseconds), so, the smaller the packet, the smaller the chance that the transmission will be postponed after dwell time.

However, it doesn't make sense to introduce a new LAN protocol that cannot deal with packets 1518 bytes long which are used on Ethernet, so the committee decided to solve the problem by adding a simple fragmentation/reassembly mechanism at the MAC Layer.

The mechanism is a simple Send-and-Wait algorithm, where the transmitting station is not allowed to transmit a new fragment until one of the following happens:

1. Receives an ACK for the said fragment, or
2. Decides that the fragment was retransmitted too many times and drops the whole frame

It should be noted that the standard does allow the station to transmit to a different address between retransmissions of a given fragment. This is particularly useful when an AP has several outstanding packets to different destinations and one of them does not respond.

The following diagram shows a frame (MSDU) being divided to several fragments (MPDUs):




Figure 3: Frame Fragmentation

Inter Frame Spaces

The Standard defines 4 types of Inter Frame Spaces, which are used to provide different priorities:

- **SIFS - Short Inter Frame Space**, is used to separate transmissions belonging to a single dialog (e.g. Fragment-Ack), and is the minimum Inter Frame Space. There is always at most one single station to transmit at any given time, therefore giving it priority over all other stations. This value is a fixed value per PHY and is calculated in such a way that the transmitting station will be able to switch back to receive mode and be capable of decoding the incoming packet. On the 802.11 FH PHY this value is set to 28 microseconds
- **PIFS - Point Coordination IFS**, is used by the Access Point (or Point Coordinator, as called in this case), to gain access to the medium before any other station. This value is SIFS plus a Slot Time (defined in the following paragraph), i.e 78 microseconds.
- **DIFS - Distributed IFS**, is the Inter Frame Space used for a station willing to start a new transmission, which is calculated as PIFS plus one slot time, i.e. 128 microseconds.
- **EIFS - Extended IFS**, which is a longer IFS used by a station that has received a packet that it could not understand. This is needed to prevent the station (which could not understand the duration information for the Virtual Carrier Sense) from colliding with a future packet belonging to the current dialog.

Exponential Backoff Algorithm

Backoff is a well known method used to resolve contention between different stations wanting to access the medium. The method requires each station to choose a Random Number (n) between 0

and a given number, and wait for this number of Slots before accessing the medium, always checking if a different station has accessed the medium before.

The **Slot Time** is defined in such a way that a station will always be capable of determining if another station has accessed the medium at the beginning of the previous slot. This reduces collision probability by half.

Exponential Backoff means that each time the station chooses a slot and happens to collide, it will increase the maximum number for the random selection exponentially.

The 802.11 standard defines an **Exponential Backoff Algorithm**, that must be executed in the following cases:

- When the station senses the medium before the first transmission of a packet, and the medium is busy
- After each retransmission, and
- After a successful transmission

The only case when this mechanism is not used is when the station decides to transmit a new packet and the medium has been free for more than DIFS.

The following figure shows a schematic of the access mechanism:

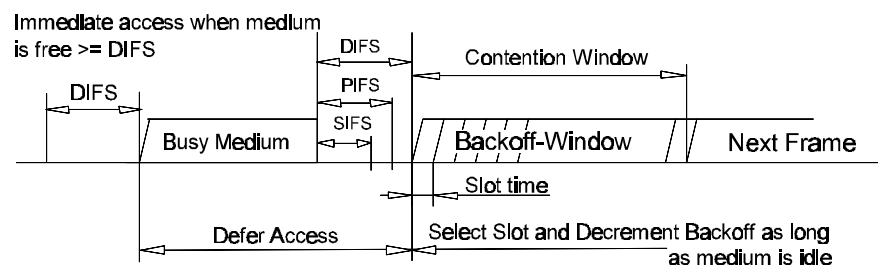


Figure 4: Access Mechanism

How Does a Station Join an Existing Cell (BSS)?

When a station wants to access an existing BSS (either after power-up, sleep mode, or just entering the BSS area), the station needs to get synchronization information from the Access Point (or from the other stations when in ad-hoc mode, which will be discussed later).

The station can get this information by one of two means:

1. **Passive Scanning:** In this case the station just waits to receive a Beacon Frame from the AP, (the beacon frame is a frame sent out periodically by the AP containing synchronization information), or
2. **Active Scanning:** In this case the station tries to locate an Access Point by transmitting Probe Request Frames, and waits for Probe Response from the AP.

Both methods are valid. A method is chosen according to the power consumption/performance trade-off.

The Authentication Process

Once the station has located an Access Point, and decides to join its BSS, it goes through the **Authentication Process**. This is the interchange of information between the AP and the station, where each side proves the knowledge of a given password.

The Association Process

Once the station is authenticated, it then starts the **Association Process**, which is the exchange of information about the stations and BSS capabilities, and which allows the DSS (the set of APs) to know about the current position of the station). A station is capable of transmitting and receiving data frames only after the association process is completed.

Roaming

Roaming is the process of moving from one cell (or BSS) to another without losing connection. This function is similar to the cellular phones' handover, with two main differences:

1. On a packet-based LAN system, the transition from cell to cell may be performed between packet transmissions, as opposed to telephony where the transition may occur during a phone conversation, this makes the LAN roaming a little easier, but
2. On a voice system, a temporary disconnection may not affect the conversation, while in a packet-based environment it significantly reduces performance because retransmission is then performed by the upper layer protocols.

The 802.11 standard does not define how roaming should be performed, but defines the basic tools. These include active/passive scanning, and a re-association process, where a station which is roaming from one Access Point to another becomes associated with the new one¹.

¹. The BreezeNet product line provides an patented enhanced roaming mechanism which allows stations to roam at speeds of 60 Km/h without losing or duplicating packets.

Keeping Synchronization

Stations need to keep synchronization, which is necessary for keeping hopping synchronized, and other functions like Power Saving. On an infrastructure BSS, this is achieved by all the stations updating their clocks according to the AP's clock, using the following mechanism:

The AP periodically transmits frames called **Beacon Frames**. These frames contain the value of the AP's clock at the moment of transmission (note that this is the moment when transmission actually occurs, and not when it is put in the queue for transmission. Since the Beacon Frame is transmitted using CSMA rules, transmission may be delayed significantly).

The receiving stations check the value of their clocks at the moment the signal is received, and correct it to keep in synchronization with the AP's clock. This prevents clock drifting which could cause loss of synch after a few hours of operation.

Security

Security is one of the first concerns that people have when deploying a Wireless LAN. The 802.11 committee has addressed the issue by providing what is called **WEP (Wired Equivalent Privacy)**.

Users are primarily concerned that an intruder should not be able to:

- Access the Network resources by using similar Wireless LAN equipment
- Capture Wireless LAN traffic (eavesdropping)

Preventing Access to Network Resources

This is done by the use of an Authentication mechanism where a station needs to prove knowledge of the current key. This is very similar to Wired LAN privacy, in the sense that an intruder needs to enter the premises (by using a physical key) in order to connect his workstation to the wired LAN.

Eavesdropping

Eavesdropping is prevented by using the WEP algorithm which is a Pseudo Random Number Generator initialized by a shared secret key. This PRNG outputs a key sequence of pseudo-random bits equal in length to the largest possible packet which is combined with the outgoing/incoming packet producing the packet transmitted in the air.

The WEP is a simple algorithm based on RSA's RC4 which has the following properties:

- **Reasonably strong:**

Brute-force attack to this algorithm is difficult because every frame is sent with an Initialization Vector which restarts the PRNG for each frame.

- **Self Synchronizing:**

The algorithm re-synchronizes for each message. This is necessary in order to work in a connection-less environment, where packets may get lost (as any LAN).

Power Saving

Wireless LANs are typically related to mobile applications. In this type of application, battery power is a scarce resource. This is the reason why the 802.11 standard directly addresses the issue of Power Saving and defines an entire mechanism which enables stations to go into sleep mode for long periods of time without losing information.

The main idea behind the Power Saving Mechanism is that the AP maintains a continually updated record of the stations currently working in Power Saving mode, and buffers the packets addressed to these stations until either the stations specifically request the packets by sending a polling request, or until they change their operation mode.

As part of its Beacon Frames, The AP also periodically transmits information about which Power Saving Stations have frames buffered at the AP, so these stations wake up in order to receive the Beacon Frame. If there is an indication that there is a frame stored at the AP waiting for delivery, then the station stays awake and sends a Polling message to the AP to get these frames.

Multicasts and Broadcasts are stored by the AP, and transmitted at a pre-known time (each DTIM), when all Power Saving stations who wish to receive this kind of frames are awake.

Frame Types

There are three main types of frames:

- **Data Frames:** which are used for data transmission
- **Control Frames:** which are used to control access to the medium (e.g. RTS, CTS, and ACK), and
- **Management Frames:** which are frames that are transmitted the same manner as data frames to exchange management information, but are not forwarded to upper layers (e.g. beacon frames).

Each frame type is subdivided into different Subtypes, according to their specific function.

Frame Formats

All 802.11 frames are composed of the following components:

Preamble	PLCP Header	MAC Data	CRC
----------	-------------	----------	-----

Preamble

This is PHY dependent, and includes:

- **Synch:** An 80-bit sequence of alternating zeros and ones, which is used by the PHY circuitry to select the appropriate antenna (if diversity is used), and to reach steady-state frequency offset correction and synchronization with the received packet timing.
- **SFD:** A Start Frame delimiter which consists of the 16-bit binary pattern 0000 1100 1011 1101, which is used to define frame timing.

PLCP Header

The PLCP Header is always transmitted at 1 Mbit/s and contains Logical information used by the PHY Layer to decode the frame. It consists of:

- **PLCP_PDU Length Word:** which represents the number of bytes contained in the packet. This is useful for the PHY to correctly detect the end of packet.
- **PLCP Signaling Field:** which currently contains only the rate information, encoded in 0.5 MBps increments from 1 Mbit/s to 4.5 Mbit/s.
- **Header Error Check Field:** Which is a 16 Bit CRC error detection field.

MAC Data

The following figure shows the general MAC Frame Format. Part of the fields are only present in part of the frames as described later.




Figure 5: MAC Frame Format

Frame Control Field

The Frame Control field contains the following information:




Figure 6: Frame Control Field

Protocol Version

This field consists of 2 bits which are invariant in size and placement across following versions of the 802.11 Standard, and will be used to recognize possible future versions. In the current version of the standard the value is fixed as 0.

Type and Subtype

This 6 bits define the Type and SubType of the frame as indicated in the following table:

Type Value b3 b2	Type Description	Subtype Value b7 b6 b5 b4	Subtype Description
00	Management	0000	Association Request
00	Management	0001	Association Response
00	Management	0010	Association Request
00	Management	0011	Reassociation Response
00	Management	0100	Probe Request
00	Management	0101	Probe Response
00	Management	0110-0111	Reserved
00	Management	1000	Beacon
00	Management	1001	ATIM
00	Management	1010	Disassociation
00	Management	1011	Authentication
00	Management	1100	Deauthentication
00	Management	1101-1111	Reserved
01	Control	0000-0001	Reserved
01	Control	1010	PS-Poll
01	Control	1011	RTS
01	Control	1100	CTS
01	Control	1101	ACK
01	Control	1110	CF End
01	Control	1111	CF End + CF-ACK
10	Data	0000	Data
10	Data	0001	Data + CF-Ack
10	Data	0010	Data + CF-Poll
10	Data	0011	Data + CF-ACK + CF-Poll
10	Data	0100	Null Function (no data)
10	Data	0101	CF-Ack (no data)
10	Data	0110	CF-Poll (no data)
10	Data	0111	CF-Ack + CF-Poll (no data)
10	Data	1000-1111	Reserved
10	Data	0000-1111	Reserved

ToDS

This bit is set to 1 when the frame is addressed to the AP for forwarding to the Distribution System (including the case where the destination station is in the same BSS, and the AP is to relay the frame).

The Bit is set to 0 in all other frames.

FromDS

This bit is set to 1 when the frame is received from the Distribution System.

More Fragments

This bit is set to 1 when there are more fragments belonging to the same frame following the current fragment.

Retry

This bit indicates that this fragment is a retransmission of a previously transmitted fragment. This is used by the receiver station to recognize duplicate transmissions of frames that may occur when an Acknowledgment packet is lost.

Power Management

This bit indicates the Power Management mode that the station will be in after the transmission of this frame. This is used by stations which are changing state either from Power Save to Active or vice versa.

More Data

This bit is used for Power Management as well as by the AP to indicate that there are more frames buffered to this station. The station may decide to use this information to continue polling or even changing to Active mode.

WEP

This bit indicates that the frame body is encrypted according to the WEP algorithm

Order

This bit indicates that this frame is being sent using the Strictly-Ordered service class.²

². The Strictly-Ordered Service Class is defined for users that cannot accept change of ordering between Unicast Frames and Multicast Frames (ordering of unicast frames to a specific address is always maintained). The only known protocol that would need this service class is DEC's LAT.

Duration/ID

This field has two meanings depending on the frame type:

- In Power-Save Poll messages this is the Station ID
- In all other frames this is the duration value used for the NAV Calculation.

Address Fields

A frame may contain up to 4 Addresses depending on the ToDS and FromDS bits defined in the Control Field, as follows:

- **Address-1** is always the Recipient Address (i.e. the BSS station that is the immediate recipient of the packet). If ToDS is set, this is the AP Address, if ToDS is not set then this is the address of the end-station.
- **Address-2** is always the Transmitter Address (i.e. the station which is physically transmitting the packet). If FromDS is set, this is the AP address, if it is not set then it is the Station address.
- **Address-3** is in most cases the remaining, missing address. On a frame with FromDS set to 1, Address-3 is the original Source Address, if the frame has the ToDS set, then Address 3 is the destination Address.
- **Address-4** is used in special cases where a Wireless Distribution System is used, and the frame is being transmitted from one Access Point to another. In such cases, both the ToDS and FromDS bits are set, so both the original Destination and the original Source Addresses are missing.

The following Table summarizes the usage of the different Addresses according to ToDS and FromDS bits setting:

To DS	From DS	Address 1	Address 2	Address 3	Address 4
0	0	DA	SA	BSSID	N/A
0	1	DA	BSSID	SA	N/A
1	0	BSSID	SA	DA	N/A
1	1	RA	TA	DA	SA

Sequence Control

The Sequence Control Field is used to represent the order of different fragments belonging to the same frame, and to recognize packet duplications. It consists of two subfields, Fragment Number and Sequence Number, which define the frame and the number of the fragment in the frame.

CRC

The CRC is a 32 bit field containing a 32-bit Cyclic Redundancy Check (CRC)

Most Common Frame Formats

RTS Frame Format

The RTS frame looks as follows:




Figure 7: RTS Frame Format

The RA of the RTS frame is the address of the STA on the wireless medium that is the intended immediate recipient of the next Data or Management frame.

The TA is the address of the STA transmitting the RTS frame.

The Duration value is the time, in microseconds, required to transmit the next Data or Management frame, plus one CTS frame, plus one ACK frame, plus three SIFS intervals.

CTS Frame Format

The CTS frame looks as follows:




Figure 8: CTS Frame

The Receiver Address (RA) of the CTS frame is copied from the Transmitter Address (TA) field of the immediately previous RTS frame to which the CTS is a response.

The Duration value is the value obtained from the Duration field of the immediately previous RTS frame, minus the time, in microseconds, required to transmit the CTS frame and its SIFS interval.

ACK Frame Format

The ACK frame looks as follows:




Figure 9: ACK Frame Format

The Receiver Address of the ACK frame is copied from the Address 2 field of the immediately previous frame.

If the More Fragment bit was set to 0 in the Frame Control field of the previous frame, the Duration value is set to 0, otherwise the Duration value is obtained from the Duration field of the previous frame, minus the time, in microseconds, required to transmit the ACK frame and its SIFS interval.

Point Coordination Function (PCF)

Beyond the basic Distributed Coordination Function, there is an optional Point Coordination Function, which may be used to implement time-bounded services, like voice or video transmission. This Point Coordination Function makes use of the higher priority that the Access Point may gain by the use of a smaller Inter Frame Space (PIFS).

By using this higher priority access, the Access Point issues polling requests to the stations for data transmission, hence controlling medium access. In order to still enable regular stations to access the



medium, there is a provision that the Access Point must leave enough time for Distributed Access in between the PCF.

Ad-Hoc Networks

In certain circumstances the users may wish to build up Wireless LAN networks without an infrastructure (more specifically without an Access Point). This may include file transfer between two notebooks users, a co-workers meeting outside the office, etc.

The 802.11 Standard addresses this need by the definition of an “ad-hoc” mode of operation. In this case there is no Access Point and part of its functionality is performed by the end-user stations (such as Beacon Generation, synchronization, etc.). Other AP functions are not supported (such as frame-relaying between two stations not in range, or Power Saving).