IEEE 802.11 Notes

This document provides a primer on the layered Internet protocol stack and its correspondence to the IEEE 802.11 standard. The components and architecture of an 802.11 network along with the typical services offered are discussed. The IEEE 802.11 frame format and common frame subtypes are described.

1 Local Area Networks

A Local Area Network (LAN) links the devices in a single office, building, or campus [1, 4]. Depending on the needs of an organisation and the type of technology used, a LAN can be as simple as two personal computers and a printer in someone's home office or it can extend throughout an enterprise and include audio and video peripherals. LAN size is typically limited to a few kilometres.

LANs are designed to allow resources to be shared between workstations. The resources to be shared can include hardware (e.g., printer), software (e.g., an application program), or data. For example, in many business environments a LAN links a workgroup of task-related computers, such as engineering workstations or accounting PCs. Generally, a given LAN uses only one type of transmission medium. When two or more networks are connected, they become an inter-network.

2 Protocols

A protocol is a set of rules that governs data communications. A protocol defines what, how, and when data is communicated. The key elements of a protocol are syntax, semantics, and timing [1].

Syntax refers to the structure or format of the data, specifying the order in which they are presented. Semantics refers to the meaning of each section of bits. It dictates how a particular pattern is to be interpreted and what action is to be taken based on that interpretation. Timing relates to flow control, which is often used in protocols to constrain how much data can be sent, and when.

3 Network Model

The layered protocol stack prevalent for data networking today is the five-layer Network reference model. The model is composed of five ordered layers: physical (layer 1), data

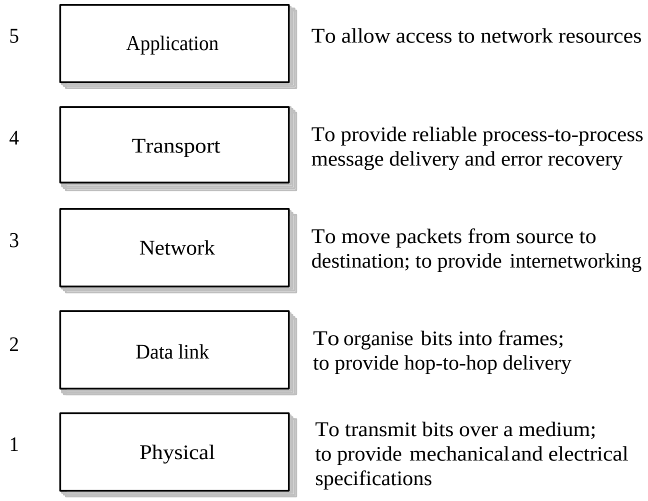
5
Application
To allow access to network resources
4
Transport
To provide reliable process-to-process message delivery and error recovery
3
Network
To move packets from source to destination; to provide internetworking
2
Data link
To organise bits into frames; to provide hop-to-hop delivery
1
Physical
To transmit bits over a medium; to provide mechanical and electrical specifications

Figure 1: Network model layers

link (layer 2), network (layer 3), transport (layer 4), and application (layer 5) [1]. While developing the model, the designers identified networking functions with related issues and collected those functions into discrete groupings that became the layers.

The five layers can be categorised into three subgroups. Layers 1, 2, and 3 collectively deal with the physical aspects of moving data from one device to another (such as electrical specifications, physical connections, physical addressing, timing, and reliability). Layer 5 is a user support layer that enables interoperability among unrelated end systems. Layer 4, the transport layer, connects these two subgroups and ensures that what the lower layers transmit is in a form that the upper layers can use.

The following subsections elaborate on the functionality at each protocol layer. Figure 1 shows the layers in the Network model and their major functions [1].

3.1 Physical Layer

The physical (PHY) layer coordinates the functions required to transmit a bit stream over a physical medium. It deals with the mechanical and electrical specifications of the interface and transmission media. It also defines the procedures and functions that physical devices and interfaces have to perform for the transmission to occur.

3.2 Data Link Layer

The data link layer transforms the physical layer, a raw transmission facility, into a reliable link. It makes the physical layer appear error-free to the upper network layer.

The major functions of the data link layer are as follows:

- Framing: The data link layer divides the stream of bits received from the network layer into manageable data units called frames.

- Physical addressing: If the frames are to be distributed to different systems on the network, the data link layer adds a header to the frame to identify the sender and/or receiver of the frame.

- Flow control: If the rate at which the data are processed by the receiver is less than the rate generated by the sender, the data link layer imposes a flow control mechanism to prevent overwhelming the receiver.

- Error control: The data link layer adds reliability to the physical layer using mechanisms to detect and retransmit damaged or lost frames. It also uses a mechanism to prevent duplication of frames. Error control is normally achieved through a trailer at the end of the frame.

The data link layer is divided into two sub-layers: Logical Link Control (LLC) and Medium Access Control (MAC) [5].

3.2.1 Logical Link Control Sublayer

LLC provides physical-medium-independent data link layer services to the network layer. It also provides error control and flow control services. It enables different protocols from the higher layers of the protocol stack to access different types of physical networks. Hence, the upper layers can function independently without worrying about the type of the physical network in use.

3.2.2 Medium Access Control Sublayer

The MAC sublayer forms the lower half of the data link layer. It directly interfaces with the physical layer. It provides services such as addressing, framing, and medium access control. Unlike the LLC, these services vary with the physical medium in use. Of these, medium access control is considered to be the most important service. It is relevant to networks (such as LANs) where a single broadcast transmission channel needs to be shared by multiple competing machines.

3.3 Network Layer

The network layer is responsible for the source-to-destination delivery of a packet, for example, using IP. Whereas the data link layer oversees the delivery of the packet between two systems on the same network, the network layer ensures that each packet gets from its point of origin to its final destination.

The major functions of the network layer are as follows:

- Logical addressing: The physical addressing implemented by the data link layer handles the addressing problem locally. If a packet passes the network boundary, another addressing system is used to help distinguish the source and destination systems. The network layer adds a header to the packet coming from the upper layer that includes the logical addresses of the sender and the receiver.

- Routing: When independent networks or links are connected to create an inter-network, the connecting devices such as routers or switches direct the packets to their final destination. One of the functions of the network layer is to provide this mechanism.

3.4 Transport Layer

The transport layer is responsible for process-to-process delivery of the messages. Whereas the network layer oversees source-to-destination delivery of individual packets, it does not recognise any relationship between those packets. The transport layer, on the other hand, ensures that the whole message arrives intact and in order, overseeing both error and flow control at the process-to-process level.

Some of the major duties of the transport layer are as follows:

- Port addressing: Computers often run several processes at the same time. For this reason, process-to-process delivery involves delivery not only from one computer to another but also from a specific process on one computer to a specific process on the other. The transport layer header must therefore include a type of address called a port address. The network layer gets each packet to the correct computer; the transport layer gets the entire message to the correct process on that computer.

- Connection management: The transport layer can be either connectionless (e.g., UDP) or connection-oriented (e.g., TCP) . A connectionless transport layer treats each segment as an independent packet and delivers it to the transport layer at the destination machine. A connection-oriented transport layer first makes a logical connection with the transport layer at the destination machine before delivering the packets. Once all the data are transferred, the connection is terminated.

3.5 Application Layer

The application layer allows the user, whether human or software, to access the network. It provides user interfaces and support for services such as electronic mail, remote file access and transfer, remote log-in, access to the World Wide Web, and so on.

4 Packet Encapsulation

Figure 2 illustrates the packet encapsulation process in a network [6]. Each piece of information transmitted on an IEEE 802.x LAN is sent as a packet. A packet is a chunk of data enclosed in one or more wrappers that help deliver the data to the correct destination.

Packets are created at the machine sending the information. The application generating the data on the sending machine passes the data to a protocol stack running on that machine. The protocol stack breaks the data down into chunks and encapsulates each chunk in one or more wrappers that allow the packets to be reassembled in the correct order at the destination. The protocol stack on the sending machine then passes the packets to the NIC.

If the packet's ultimate destination is somewhere off the local network, the header added by the sending machine indicates a router or switch as its destination address. The router

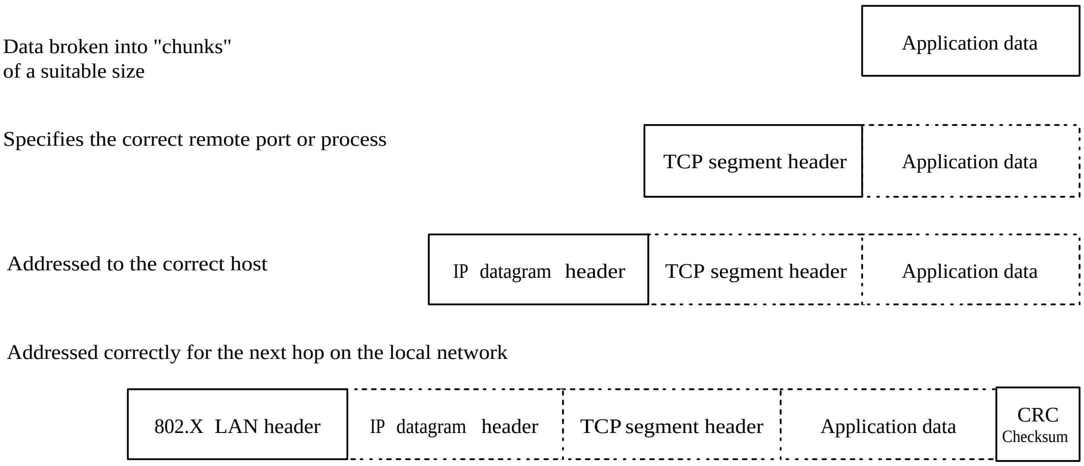

Figure 2: Constructing an IP packet

will process the original wrapper and determine the ultimate destination address. It will then re-wrap the packet, giving it a new header and sends the packets on the next hop of its journey. At the receiving end, this process is reversed.

5 IEEE 802.11 WLAN Standards

In 1997, IEEE approved 802.11, the first sanctioned WLAN standard. This first standard allowed three possible implementations for the physical layer: infrared (IR) pulse position modulation, or radio frequency (RF) signalling in the 2.4 GHz band using either Frequency Hopping Spread Spectrum (FHSS) or Direct Sequence Spread Spectrum (DSSS) [6]. The IR method was never commercially implemented. The RF versions suffered from low transmission speeds. IEEE established two networking groups (A and B) to explore alternate implementations of 802.11. A third group, working group G, was set up after these two.

Group A explored the 5 GHz band, using Orthogonal Frequency Division Multiplexing (OFDM) to achieve transmission rates in the range of 54 Mbps. The 802.11a standard was ratified in 1999. Due to slow availability of cheap 5 GHz components (required for keeping the cost of products low) and international regulations, the 802.11a WLAN standard did not reach the market before mid-2002.

Group B explored more sophisticated DSSS techniques in the original 2.4 GHz band. Their 802.11b WLAN standard, published in 1999, can deliver data rates up to 11 Mbps. Most WLAN systems in the market today follow the 802.11b WLAN standard.

Group G began by exploring a variety of methods to further improve throughput in the 2.4 GHz spectrum used by the 802.11b standard. In 2003, group G ratified the 802.11g standard adopting OFDM, the same signalling method used in the 802.11a WLAN standard. The 802.11g standard provides backward compatibility with the older 802.11b standard, which uses the same spectrum.

Even though 802.11g operates in the same frequency band as 802.11b, it can achieve higher data rates because of its similarities to 802.11a. The maximum range of 802.11g devices is slightly greater than that of 802.11b devices, but the range in which a client can

Table 1: Data rates supported by IEEE 802.11a, b, g

Standard	Data rates (Mbps)
802.11a	6, 9, 12, 24, 36, 48, 54
802.11b	1, 2, 5.5, 11
802.11g	1, 2, 5.5, 6, 9, 11, 12, 22, 24, 36, 48, 54

achieve full data rate speed (54 Mbps) is much shorter than that of 802.11b. The MAC layers of 802.11a, b, and g protocols are identical.

Each portion of the radio spectrum is called a channel. Most designers use one or more channels between 1 and 11 for deploying 802.11 WLANs. To overcome signal degradation, 802.11 WLANs can step down to a slower but more robust transmission rate when conditions are poor, then step back up again when conditions improve. Data rates supported by the 802.11 WLAN standard are shown in Table 1.

6 IEEE 802.11 WLAN Components

IEEE 802.11 networks consist of four major components [2]:

Stations Stations are computing devices with wireless network interfaces. Typically, stations are battery-operated laptop or handheld pocket PCs.

Access points Frames on an 802.11 network must be converted to another type of frame for delivery to a wired network. Devices called access points (AP) perform the wireless-to-wired bridging function.

Wireless medium To move frames from station to station, the standard uses a wireless medium. Typical WLANs utilise an RF physical layer.

Distribution system When several APs are connected to form a large coverage area, they must communicate with each other to handle the movements of mobile stations. The distribution system (DS) is the logical component of 802.11 used to forward frames to their destination. Usually, the DS is implemented as a combination of a bridging engine and a distribution system medium, which is the backbone network used to relay frames between APs. In most cases, Ethernet is used as the backbone network technology.

Most APs operate as bridges. They have at least one wireless network interface and at least one Ethernet network interface. The Ethernet side can be connected to an existing network, and the wireless side becomes an extension of that network.

7 IEEE 802.11 Architecture

The IEEE 802.11 standard defines two kinds of services: the Basic Service Set (BSS) and the Extended Service Set (ESS) [1, 6]. The BSS is the basic building block of a wireless LAN. A BSS consists of stationary or mobile wireless stations and possibly a central base station

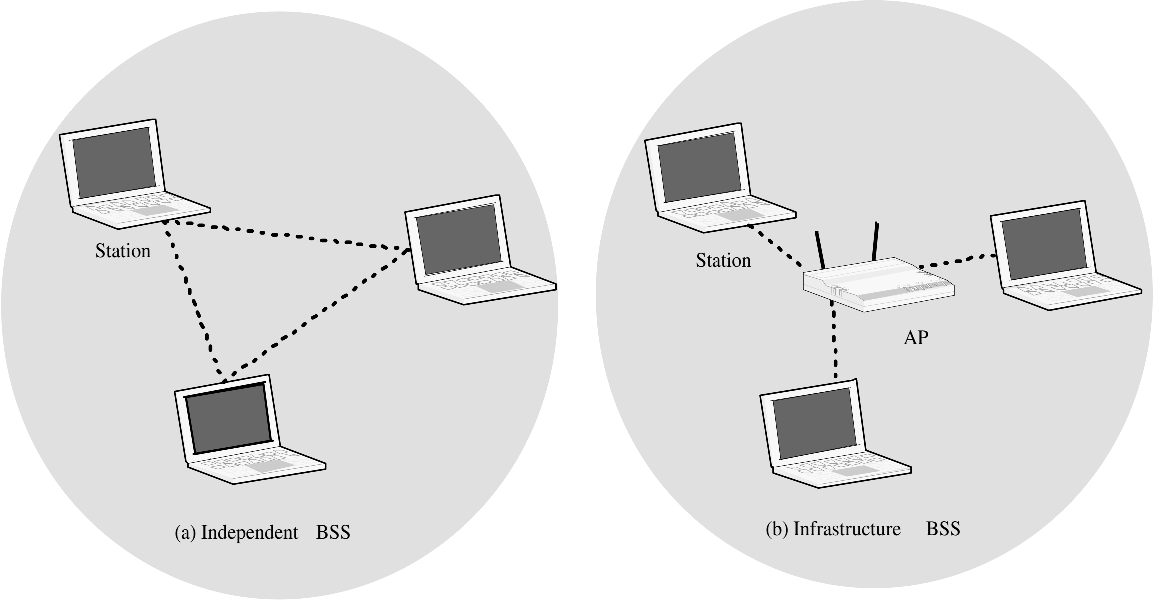

Figure 3: Independent and Infrastructure Basic Service Sets

(e.g., an AP). When a station is in the BSS, it can communicate with the other members of the BSS.

The BSS without an AP is a stand-alone network and cannot send data to other BSSs. Such BSSs are called Independent BSSs (IBSS). See Figure 3 (a). Typically, IBSSs involve a small number of stations set up for a specific purpose and for a short period of time (e.g., creating a short-lived network to support a single meeting in a conference room). IBSSs are also referred to as ad hoc networks.

Infrastructure BSSs are distinguished from ad hoc networks by the use of an AP. See Figure 3 (b). APs are used for all communications in an infrastructure BSS, including communication between mobile nodes in the same service set. An infrastructure BSS is bounded by the coverage distance from the AP. The coverage area of a single AP is called a cell. All mobile stations are required to be within reach of the AP.

802.11 allows wireless networks of arbitrarily large size to be created by linking BSSs into an ESS. An ESS is created by chaining BSSs together with a backbone network. All the APs in an ESS are given the same Service Set Identifier (SSID), which serves as a network name for its users. APs in an ESS operate in a manner such that the outside world can use the station's MAC address to talk to a station without worrying about its location in the ESS.

Figure 4 shows three BSSs corresponding to three APs. There is an equal level of overlap between BSS 1 and BSS 2, and between BSS 2 and BSS 3. Such overlap is necessary to provide stations with seamless connectivity if they move from one BSS to another. In the figure, the router uses the station's MAC address as the destination to deliver frames to a station; only the AP with which that station is associated delivers the frame.

Usually, mobility support is the primary motivation for deploying an 802.11 network. IEEE 802.11 allows mobility between BSSs at the link layer. However, it is not aware of anything that happens above the link layer. When stations move between BSSs, they will find and attempt to associate with an AP with the strongest signal and the least network

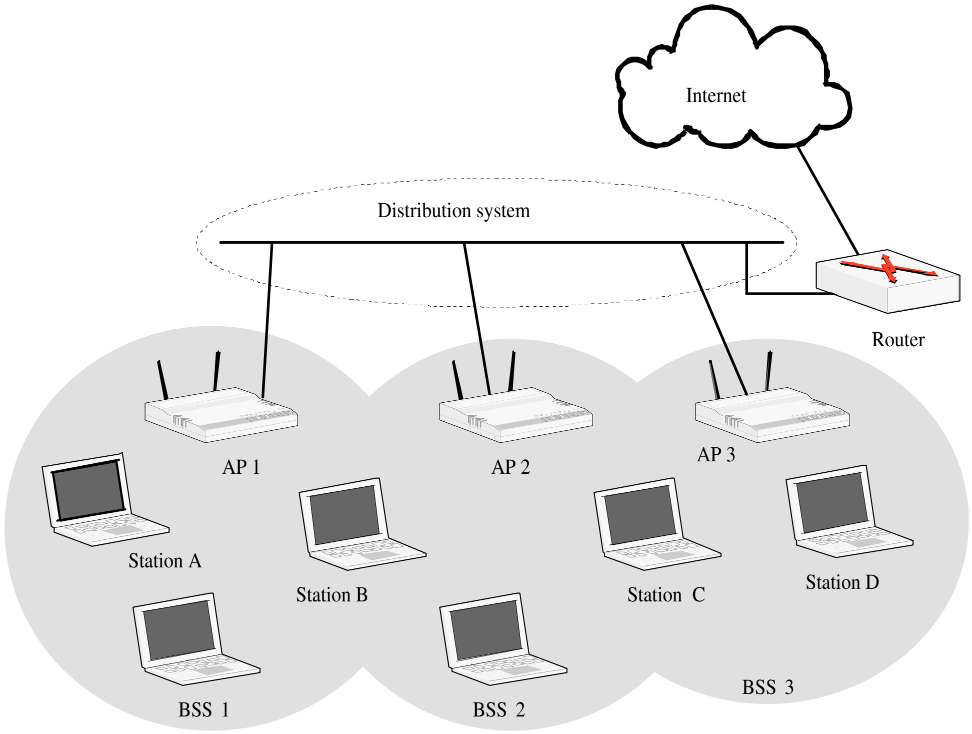

Figure 4: Extended service set

traffic. This way, a mobile station can transition seamlessly from one AP in the network to another, without losing connectivity. This event is often referred to as roaming.

8 IEEE 802.11 Framing

802.11 framing is complex compared to Ethernet framing. This is because the wireless medium requires several management features and frame types not found in wired networks. Every 802.11 frame has a control field that depicts the 802.11 protocol version, frame type, and various indicators, such as whether privacy features are on, power management is active, and so on. In addition, all frames contain MAC addresses of the source and destination station (and AP), a frame sequence number, frame body and frame check sequence (for error detection). The MAC layer frame format consists of nine fields, as shown in Figure 5 [1, 2].

Frame control The frame control is 2 bytes long and defines the type of the frame and control information. The subfields in the frame control field are as follows:

- Protocol version: Indicates which version of 802.11 MAC is contained in the rest of the frame. Only one version of the 802.11 MAC has been developed: it has been assigned a protocol number of 0.
- Type: Defines the type of information carried in the frame body: management (00), control (01), or data (10). These frame types are discussed in detail in

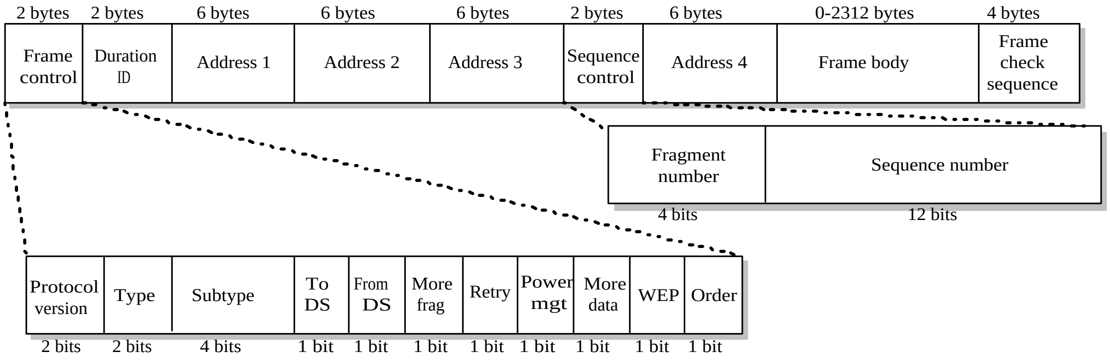

Figure 5: IEEE 802.11 frame format

Section 9.

- Subtype: Defines the subtype of each management, control, or data frame. Common frame subtypes are described in Section 9.
- ToDS: This bit is set to 1 if the frame was sent to the DS.
- FromDS: This bit is set to 1 if the frame was sent from the DS.
- More fragments: When a higher-level packet has been fragmented by the MAC, the initial fragment and any following non-final fragments set this bit to 1.
- Retry: The retry bit is set to 1 if the current packet is a retransmission of a previous attempt. This aids the receiving station in eliminating duplicate packets.
- Power management: To conserve battery life, many small devices have the ability to power down parts of the network interface. A 1 indicates that the station will be in powersave mode, and 0 indicates that the station will be active.
- More data: To accommodate stations in a power saving mode, APs may buffer frames received from the DS. An AP sets this bit to indicate that at least one frame addressed to a “sleeping” station is available.
- WEP: Wireless transmissions are inherently easier to intercept than transmissions on a wired network. The Wired Equivalent Privacy (WEP) bit is set to 1 if the payload of the packet has been encrypted using the WEP algorithm.
- Order: Frames and fragments can be transmitted in any order by both the receiving and sending stations. The bit is set to 1 when the packets must be strictly ordered, for example, for VoIP.

Duration ID The duration field is used to set the Network Allocation Vector (NAV). NAV is used for carrier sensing. The value represents the number of microseconds that the medium is expected to remain busy for the transmission currently in progress. In case of power saving stations, they create an association ID telling the AP their BSSs and where to send the buffered packets.

Address fields There are four address fields, each 6 bytes long.

Table 2: Use of address fields

ToDS	FromDS	Address 1 (receiver)	Address 2 (transmitter)	Address 3	Address 4
0	0	Destination	Source	BSSID	N/A
0	1	Destination	Sending AP	Source	N/A
1	0	Receiving AP	Source	Destination	N/A
1	1	Receiving AP	Sending AP	Destination	Source

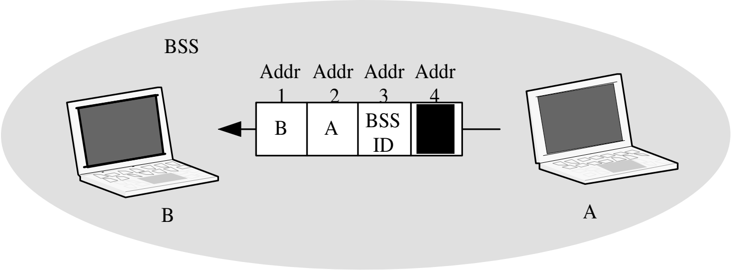

Figure 6: Address field usage in frames in an ad hoc network

The meaning of each address field depends on the value of the ToDS and the FromDS subfields as shown in Table 2. Address 1 is the address of the receiver of the frame. Address 2 is the transmitter address. Address 3 is the final destination station if it is not defined by Address 1. Address 4 is the address of the original source station if it is not the same as Address 2.

When ToDS=0 and FromDS=0, it means that the frame does not pass through a DS. The scenario shown in Figure 6 depicts an ad hoc network where mobile devices communicate amongst themselves without the involvement of APs.

For ToDS=0 and FromDS=1, the frame emerges from the DS and is destined towards a station. Address 3 in this case contains the original sender. Figure 7 shows a situation where the original source of a frame is somewhere in the Internet. In this case, the MAC address of the border gateway would be seen in Address 3, while Address 2 would be reserved for the address of the AP transmitting the frame.

As opposed to the previous case, when ToDS=1 and FromDS=0, the frame originates from a station and is headed to an AP. Address 3 contains the final destination of the frame (in another BSS or network). Figure 8 illustrates a scenario where a station in a WLAN is sending a packet to a device on the Internet.

When both ToDS and FromDS are equal to 1, then it means that the distribution system is also wireless. This scenario is rare and hence not discussed further.

Sequence control The sequence control field is composed of a 4-bit fragment number and a 12-bit sequence number. This field is used for defragmentation and discarding duplicate frames. Higher-level frames are each given a sequence number as they are passed to the MAC layer for transmission. The sequence number subfield operates as a modulo-4096

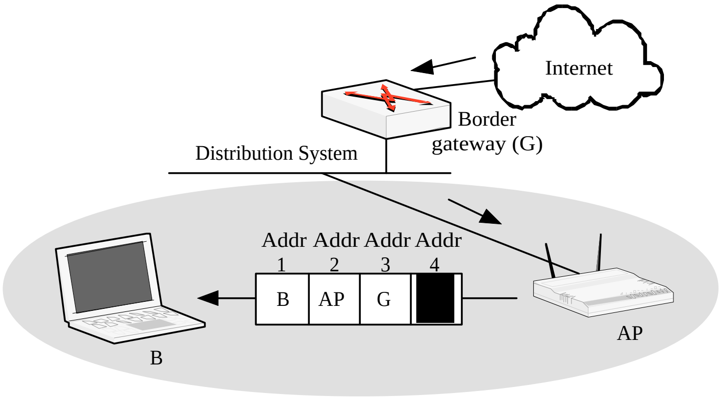

Figure 7: Address field usage in frames from the distribution system

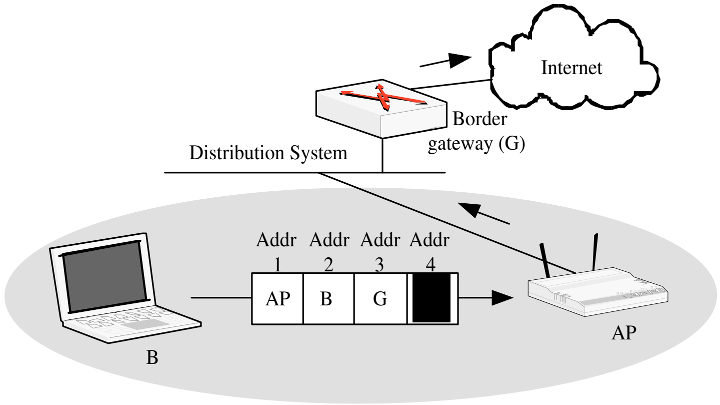

Figure 8: Address field usage in frames to the distribution system

counter of the frames transmitted. It begins at 0 and increments by 1 for each higher-level packet handled by the MAC. If higher-level packets are fragmented, all fragments will have the same sequence number. The first fragment is given a fragment number of 0 and successive fragments have their fragment number subfield incremented by 1. When frames are retransmitted, the sequence number is not changed.

Frame body The frame body moves the higher-layer payload from station to station. 802.11 can transmit frames with a maximum payload of 2304 bytes of data. On IP networks, path maximum transmission unit (MTU) typically prevents the transmission of frames with data larger than 1500 bytes. IEEE 802.11 does not generally pad frames to a minimum length.

Frame check sequence Similar to Ethernet, the 802.11 frame ends with a Frame Check Sequence (FCS), also known as Cyclic Redundancy Check (CRC). The FCS allows stations to check the integrity of received frames.

9 IEEE 802.11 Frame Types

Three major frame types exist. Data frames transport data from station to station. Control frames are used in conjunction with data frames to perform channel acquisition and carrier sensing maintenance functions, and positive acknowledgement of received data. Control and data frames work together to deliver data reliably from station to station. Management frames perform supervisory functions. They are used to join and leave wireless networks and move associations from AP to AP.

9.1 Management Frames

The following are the common management frame subtypes [2, 3]:

- Beacon: Beacon frames announce the existence of a network. They are transmitted at regular intervals to allow mobile stations to find and identify a network and possibly join it. In an infrastructure network, the AP is responsible for transmitting Beacon frames with information such as timestamp, SSID, and other parameters regarding the AP to stations that are within range.

- Authentication: Authentication frames are sent back and forth between the station requesting authentication and the station to which it is attempting to assert its authentic identity. With open system authentication, the station sends only one authentication frame, and the AP responds with an authentication frame as a response indicating acceptance or rejection.

- Deauthentication: This frame is an announcement stating that the receiver is no longer authenticated. It is a one-way communication from the authenticating station and must be accepted. It takes effect immediately.

- Association request: This frame carries information about the station (e.g., supported data rates) and the SSID of the network with which it wishes to associate. After receiving the association request, the AP considers associating with the station, and (if accepted) reserves memory space and establishes an association ID for the station. The sender must already be authenticated to obtain a successful association.

- Association response: An AP sends an association response frame containing an acceptance or rejection notice to the station requesting association. If the AP accepts the station, the frame includes information regarding the association, such as association ID and supported data rates. If the outcome of the association is positive, the station can utilise the AP to communicate with other stations on the network and systems on the DS.

- Reassociation request: If a station roams away from the currently associated AP and finds another AP having a stronger beacon signal, the station will send a reassociation frame to the new AP. The new AP then coordinates the forwarding of data frames that may still be in the buffer of the previous AP waiting for transmission to the station.

- Reassociation response: An AP sends a reassociation response frame containing an acceptance or rejection notice to the station requesting reassociation. Similar to the association process, the frame includes information regarding the association, such as association ID and supported data rates.

- Disassociation: A station sends a disassociation frame to another station if it wishes to terminate the association. For example, a station that is shutting down can "politely" send a disassociation frame to alert the AP that the station is powering off. The AP can then relinquish memory allocations and remove the radio station from the association table.

9.2 Control Frames

Some of the common control frame subtypes are as follows [2, 3]:

- Acknowledgement (ACK): ACK frames are used to send the positive acknowledgements required by the MAC for any data transmission. After receiving a data frame, the receiving station will utilise an error checking process to detect the presence of errors. The receiving station will send an ACK frame to the sending station if no errors are found. If the sending station does not receive an ACK after a period of time, the sending station will retransmit the frame.

- Power-save Poll (PS-Poll): Stations in power save mode wake up periodically to listen to selected Beacons. If they hear that data is waiting for them, they will awake more fully and send a PS-Poll frame to the AP to request the transmission of the waiting data.

10 Defining User Sessions

Users generate sessions. Sessions occur when a user joins the WLAN, uses it for a certain time, and then leaves the network. The session duration is defined as the time spent between the user joining and leaving the network.

Figure 9 illustrates how a session is started and ended by a user. Note that we do not show ACK frames in the figure. A session starts when the user NIC sends an Authentication frame to the AP. After getting a positive response from the AP, the user NIC sends an Association Request frame. If the AP allows the user NIC to associate with it, the AP replies with an Association Response frame. The user NIC is now associated with the AP. Next, DHCP boot request and boot response packets are exchanged between the user station and the DHCP server (not shown in the figure). The DHCP server is part of the DS and hence all wireless packets must pass through the AP. After being assigned an IP address, the user is ready to use the WLAN. When the user decides not to use the network any more, the NIC may send a Disassociation frame to the AP. Usually, NICs do not send out such frames when the stations are shut down. However, after a certain period of inactivity, the AP sends a series of Deauthentication frames indicating the end of a session.

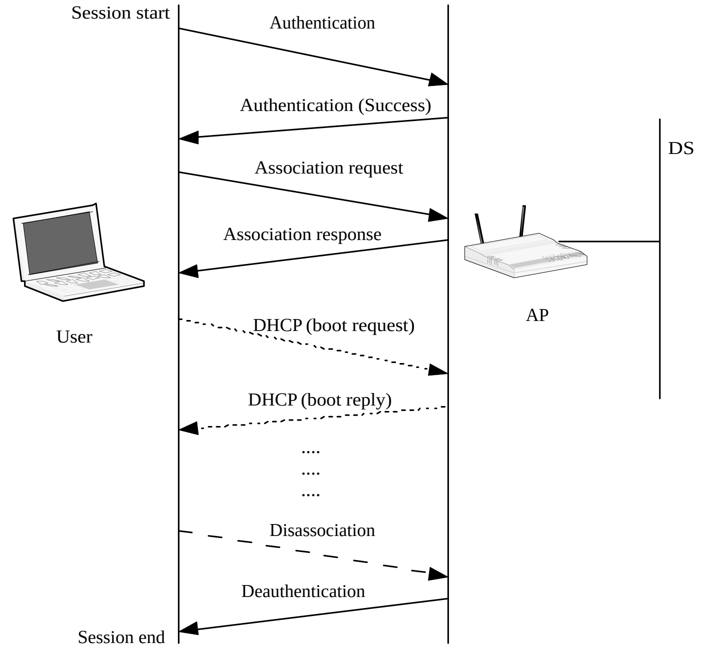

Figure 9: User starting and ending a session in the WLAN

Due to frames missed by the sniffer, the trace may not always capture all the aforementioned packets for us to determine the start and end of user sessions. We analysed every packet from the beginning of the trace. In the absence of Authentication and Association frames, we started a new session whenever a packet from a new user was noticed in the trace. Similarly, in the absence of Disassociation frames, to be able to differentiate between two sessions of the same user, we chose a session timeout of 30 minutes. Thus, if no more packets were sent/received by the user within 30 minutes of the last packet seen for that user, the session was closed. The end time for the session was set to the time of the last packet seen. However, if a new packet was noticed within the 30 minute period, the session was allowed to continue. All active sessions that persisted beyond the end of the trace were ignored.

10.1 Roaming

Users may roam from one AP to another during a session, perhaps within an hour, or during a day. Roaming can be identified by looking at the exchange of Reassociation Request and Reassociation Response frames between the user NIC and the new AP. Again due to missed frames, the trace did not record all such frames. Hence, we maintained two state variables for each unique user, viz., last used AP (represented by its BSSID) and number of roams in a session.

Each time a user NIC used a new AP to either send or receive packets, the roams counter was incremented and the current AP for the user NIC was updated. When the user session ended, the roams counter was saved to indicate the number of roams during the session. Also, we used the 30-minute session timeout for roaming users. Thus, if a user associated with a new AP during the 30-minute time limit, it was assumed that the user had only roamed and not started a new session.

11 Summary

This document presented an overview of the layered Internet protocol stack and its correspondence to the IEEE 802.11 standard. It described three different implementations of the 802.11 standard, viz., a, b, and g. The components of an 802.11 WLAN were described next. The chapter concluded with a description of the 802.11 frame format and common management and control frame subtypes.

References

[1] B. Forouzan. *Data Communications and Networking.* McGraw Hill, 2004.

[2] M. Gast. *802.11 Wireless Networks: The Definitive Guide.* O'Reilly, 2005.

[3] J. Geier. Understanding 802.11 Frame Types, August 2002. `http://www.wi-fiplanet.com/tutorials/article.php/1447501`.

[4] J. Kurose and K. Ross. *Computer Networking: A Top-down Approach Featuring the Internet.* Addison Wesley, 2003.

[5] C. Murthy and S. Manoj. *Ad Hoc Wireless Networks: Architecture and Protocols.* Prentice Hall, 2004.

[6] WildPackets. *Airopeek NX User Manual.* 2003.