Chapter 4

(Week 8)

The Medium Access Control Sublayer (CONTINUATION)

ANDREW S. TANENBAUM

COMPUTER NETWORKS

FOURTH EDITION

PP. 292-338

Networks can be divided into two categories:

1) Those using point-to-point connections and

2) Those using broadcast channels.

This chapter deals with broadcast networks and their protocols.

In any broadcast network, the key issue is how to determine who gets to use the channel when there is competition for it.

To make this point clearer, consider a conference call in which six people, on six different telephones, are all connected together so that each one can hear and talk to all the others.

It is very likely that when one of them stops speaking, two or more will start talking at once, leading to chaos.

In a face-to-face meeting, chaos is avoided by external means, for example, at a meeting, people raise their hands to request permission to speak.

When only a single channel is available, determining who should go next is much harder.

Many protocols for solving the problem are known and form the contents of this chapter.

In the literature, broadcast channels are sometimes referred to as multi-access channels or random access channels.

The protocols used to determine who goes next on a multi-access channel belong to a sublayer of the data link layer called the MAC (Medium Access Control) sublayer.

Technically, the MAC sublayer is the bottom part of the data link layer.

The MAC sublayer is especially important in LANs, nearly all of which use a multi-access channel as the basis of their communication.

WANs, in contrast, use point-to-point links, except for satellite networks.

Because multi-access channels and LANs are so closely related, in this chapter we will discuss LANs in general, as well as satellite and some other broadcast networks

4.1. THE CHANNEL ALLOCATION PROBLEM

4.2. MULTIPLE ACCESS PROTOCOLS

4.3. ETHERNET

4.4. WIRELESS LANS

4.5. BROADBAND WIRELESS

4.6. BLUETOOTH

4.7. DATA LINK LAYER SWITCHING

4.8. SUMMARY

4.4. WIRELESS LANs (1)

Almost as soon as notebook computers appeared, many people had a dream of walking into an office and magically having their notebook computer be connected to the INTERNET.

Consequently, various groups began working on ways to accomplish this goal.

WIRELESS LANs (2)

The most practical approach is to equip both the office and the notebook computers with short-range radio transmitters and receivers to allow them to communicate.

This work rapidly led to WIRELESS LANS being marketed by a variety of companies.

WIRELESS LANs (3)

Although Ethernet is widely used, it is about to get some competition.

Wireless LANs are increasingly popular, and more and more office buildings, airports, and other public places are being outfitted with them.

IEEE 802 Standards

Number	Topic
802.1	Overview and architecture of LANs
802.2 ↓	Logical link control
802.3 *	Ethernet
802.4 ↓	Token bus (was briefly used in manufacturing plants)
802.5	Token ring (IBM's entry into the LAN world)
802.6 ↓	Dual queue dual bus (early metropolitan area network)
802.7 ↓	Technical advisory group on broadband technologies
802.8 †	Technical advisory group on fiber optic technologies
802.9 ↓	Isochronous LANs (for real-time applications)
802.10 ↓	Virtual LANs and security
802.11 *	Wireless LANs
802.12 ↓	Demand priority (Hewlett-Packard's AnyLAN)
802.13	Unlucky number. Nobody wanted it
802.14 ↓	Cable modems (defunct: an industry consortium got there first)
802.15 *	Personal area networks (Bluetooth)
802.16 *	Broadband wireless
802.17	Resilient packet ring

WIRELESS LANs (4)

Wireless LANs can operate in one of two configurations:

1) With a base station and

2) Without a base station

The IEEE 802.11 standard takes this into account and makes provision for both arrangements.

WIRELESS LANs (5)

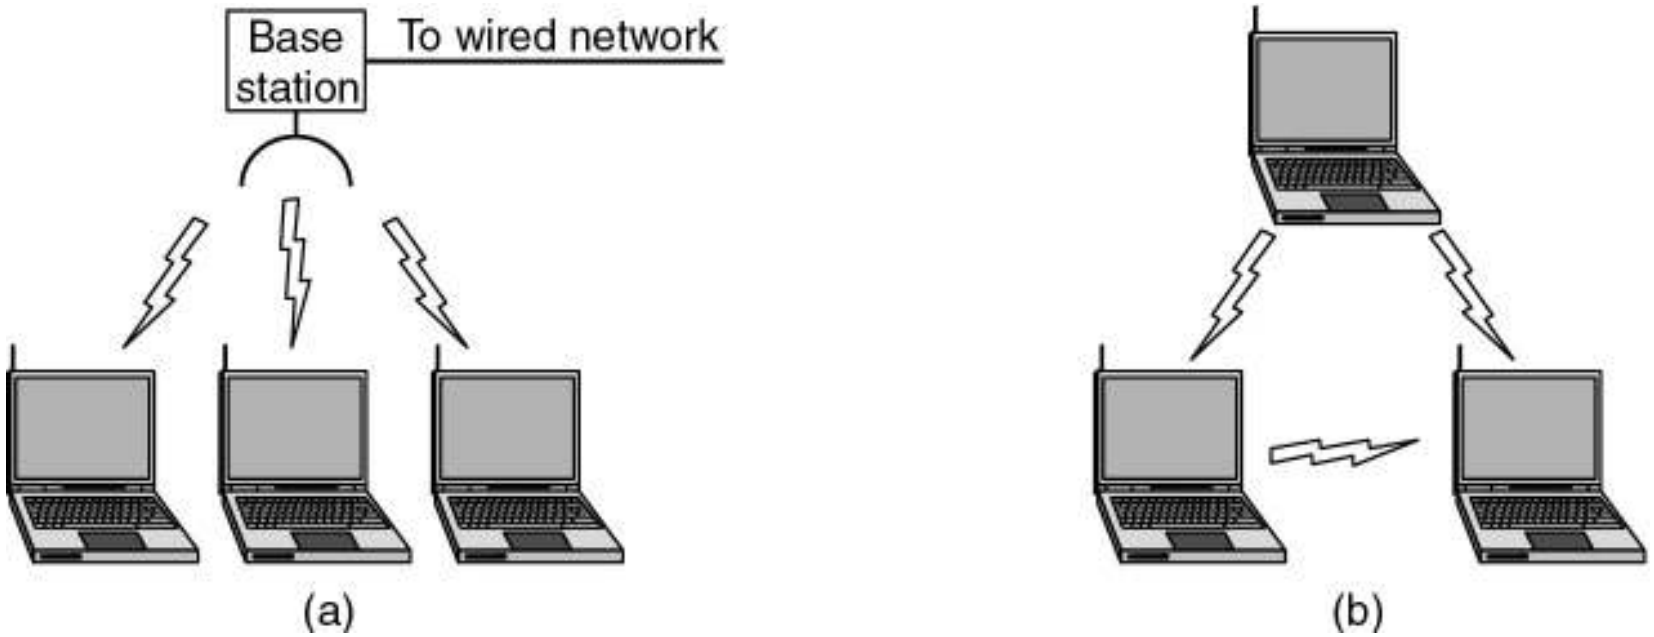

(a) Wireless networking with a base station.

(b) Ad hoc networking.

WIRELESS LANs (6)

In (a), all communication was to go through the base station, called an access point in 802.11 terminology.

In (b), the computers would just send to one another directly. This mode is now sometimes called ad hoc networking.

WIRELESS LANs (7)

At the time the standardization process started (1990), ETERNET had already come to dominate local area networking, so the committee decided to make 802.11 compatible with ETERNET above the data link layer.

WIRELESS LANs (8)

802.11a (1999) standard uses a wider frequency band and runs at speed up to 54 Mbps.

802.11b (1999) standard uses the same frequency band as 802.11, but uses a different modulation technique to achieve 11 Mbps.

WIRELESS LANs (9)

Some people see this as psychologically important since 11 Mbps is faster than original wired ETHERNET.

802.11g standard uses the modulation technique of 802.11a but the frequency band of 802.11b

Now is the time to take a closer look at the technology of 802.11 standard.

Wireless LANs (10)

- The 802.11 Protocol Stack
- The 802.11 Physical Layer
- The 802.11 MAC Sublayer Protocol
- The 802.11 Frame Structure
- Services

The 802.11 Protocol Stack (1)

The protocols used by all 802 variants, including Ethernet, have a certain commonality of structure.

A partial view of the 802.11 protocol stack is given in following figure.

The 802.11 Protocol Stack (2)

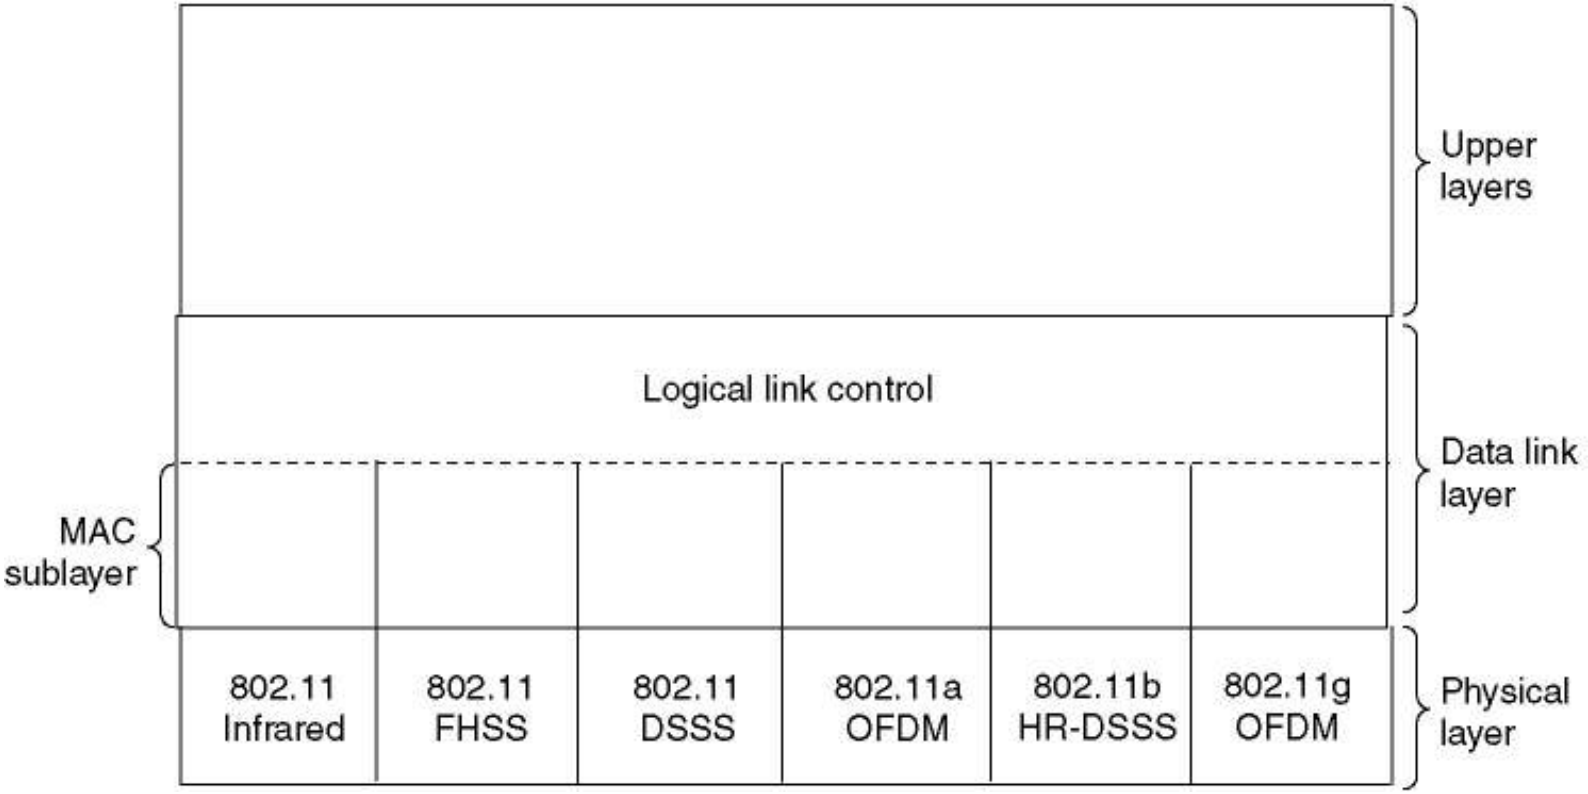

Part of the 802.11 protocol stack.

FHSS - Frequency Hopping Spread Spectrum

DSSS - Direct Sequence Spread Spectrum

OFDM - Orthogonal Frequency Division Multiplexing

HR-DSSS - High Rate DSSS

The 802.11 Physical Layer (1)

The 1997 802.11 standard specifies three groups of transmission techniques allowed in the physical layer.

1) Infrared method (1 technique)
2) Short – range radio method (2 techniques)
3) Higher bandwidth method (2+1 techniques)

The 802.11 Physical Layer (2)

Each of the five permitted transmission techniques makes it possible to send a MAC frame from one station to another.

They differ, however, in the technology used and speeds achievable.

The 802.11 Physical Layer (3)

1) Infrared method

- This method uses much the same technology as television remote controls do.
- The infrared option uses diffused (i.e., not line of sight) transmission at 0.85 or 0.95 microns.
- Two speeds are permitted: 1 Mbps and 2 Mbps

The 802.11 Physical Layer (4)

2) Short – range radio method (2 techniques)

a) FHSS (Frequency Hopping Spread Spectrum) uses 79 channels, each 1-MHz wide, starting at the low end of the 2.4-GHz ISM (Industrial, Scientific, Medical) band.

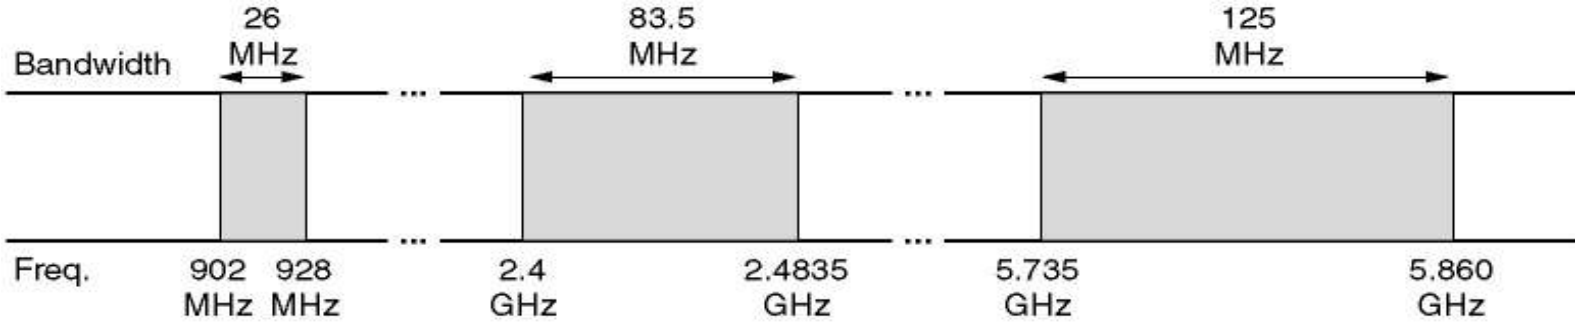

A pseudorandom number generator is used to produce the sequence of frequencies hopped to.

The 802.11 Physical Layer (5)

2) Short – range radio method (2 techniques)

b) DSSS (Direct Sequence Spread Spectrum), is also restricted to 1 or 2 Mbps.

- The scheme used has some similarities to the CDMA (Code Division Multiple Access) system.
- Each bit is transmitted as 11 chips, using what is called a Barker Sequence.

DSSS (Direct Sequence Spread Spectrum)

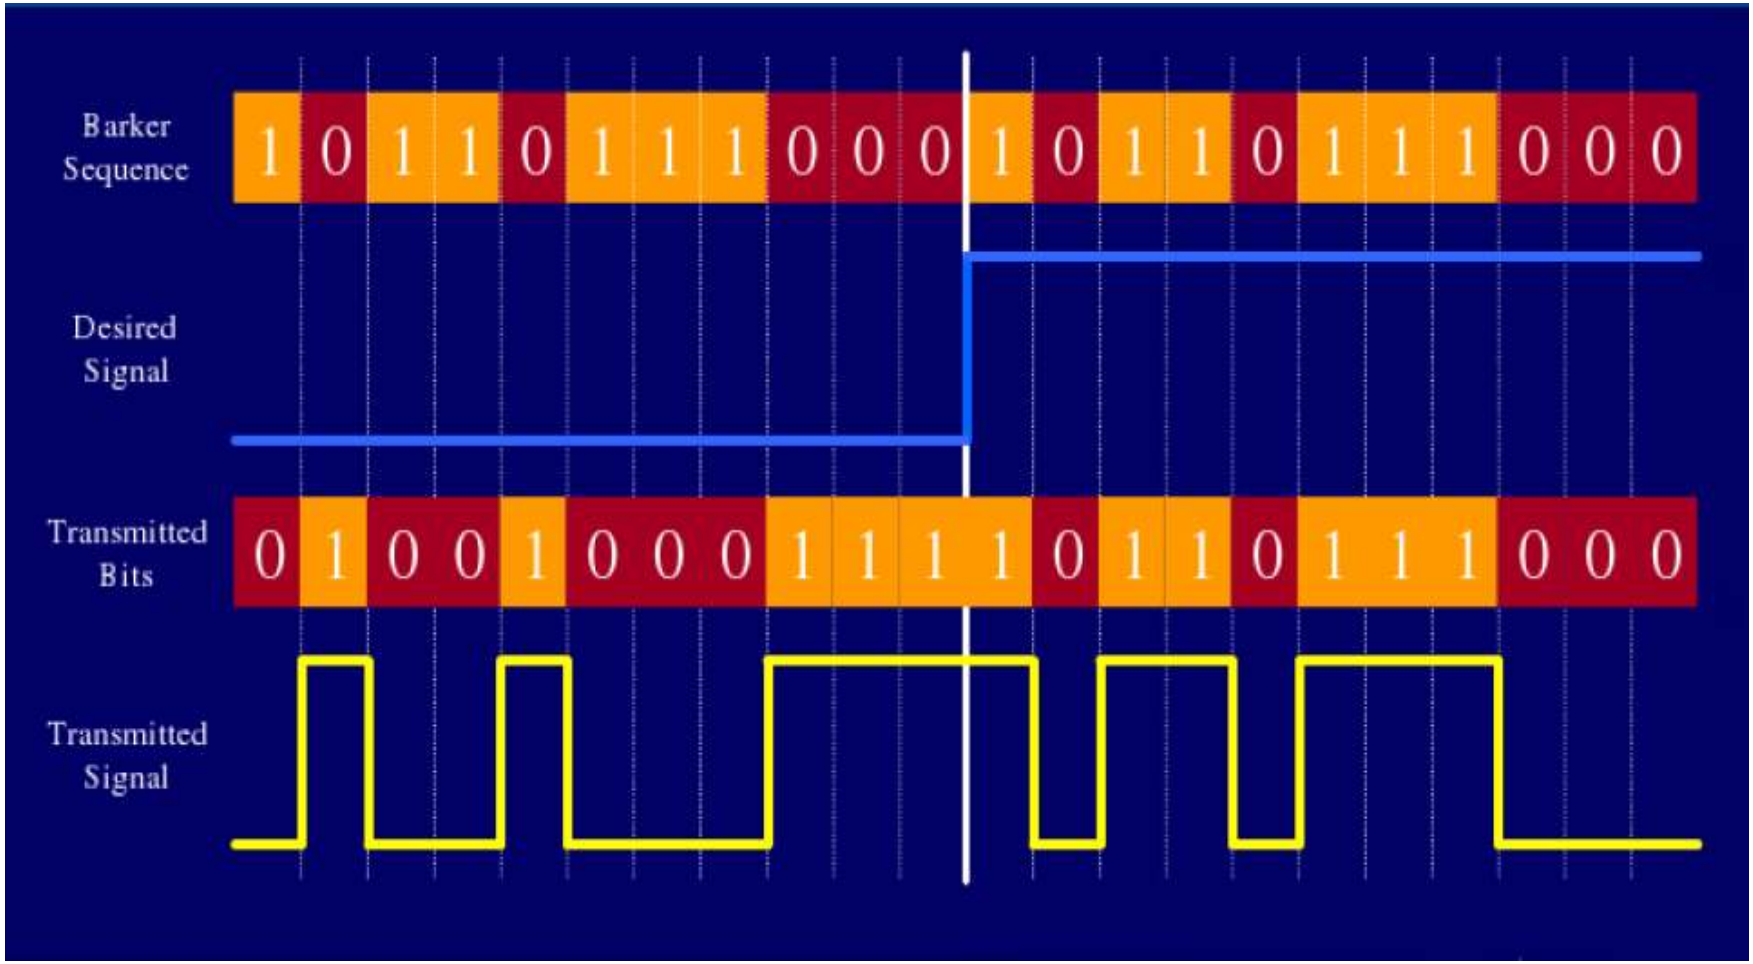

The 802.11 Physical Layer (6)

2) Short – range radio method (2 techniques)

Both of these (FHSS and DSSS) use a part of the spectrum that does not require licensing (the 2.4-GHz ISM band). Radio-controlled garage door openers also use this piece of the spectrum, so your notebook computer may find itself in competition with your garage door.

The 802.11 Physical Layer (7)

3) Higher bandwidth method (3 techniques)

a) 802.11a, OFDM (Orthogonal Frequency Division Multiplexing) is used to deliver up to 54 Mbps in the wider 5 GHz ISM band.

Splitting the signal into many narrow bands has some key advantages over using a single wide band, including better immunity to narrowband interference

The 802.11 Physical Layer (8)

3) Higher bandwidth method (3 techniques)

b) 802.11b, HR-DSSS (High Rate Direct Sequence Spread Spectrum) uses 11 million chips/sec to achieve 11 Mbps in the 2.4-GHz band.

Although 802.11b is slower than 802.11a, its range is about 7 times greater, which is more important in many situations.

The 802.11 Physical Layer (9)

3) Higher bandwidth method (3 techniques)

c) 802.11g is an enhanced version of 802.11b. It uses the OFDM (Orthogonal Frequency Division Multiplexing) modulation method of 802.11a but operates in the narrow 2.4-GHz ISM band along with 802.11b. In theory it can operate at up to 54 Mbps. It is not yet clear whether this speed will be realized in practice.

The 802.11 MAC Sublayer Protocol (1)

In 802.11,

- The MAC (Medium Access Control) sublayer determines how the channel is allocated, that is, who gets to transmit next.
- Above it is the LLC (Logical Link Control) sublayer, whose job it is to hide the differences between the different 802 variants and make them indistinguishable as far as the network layer is concerned.

The 802.11 MAC Sublayer Protocol (2)

- The 802.11 MAC sublayer protocol is quite different from that of Ethernet due to inherent complexity of the wireless environment compared to that of a wired system.
- With Ethernet, a situation just waits until the ether goes silent and starts transmitting. If it does not receive a noise burst back within the first 64 bytes, the frame has almost assuredly been delivered correctly.

The 802.11 MAC Sublayer Protocol (3)

- With wireless, this situation does not hold.
- There is the hidden station problem.
- Station C is transmitting to station B. If A sense the channel, it will not hear anything and falsely concludes that it may now start transmitting to B.

The 802.11 MAC Sublayer Protocol (4)

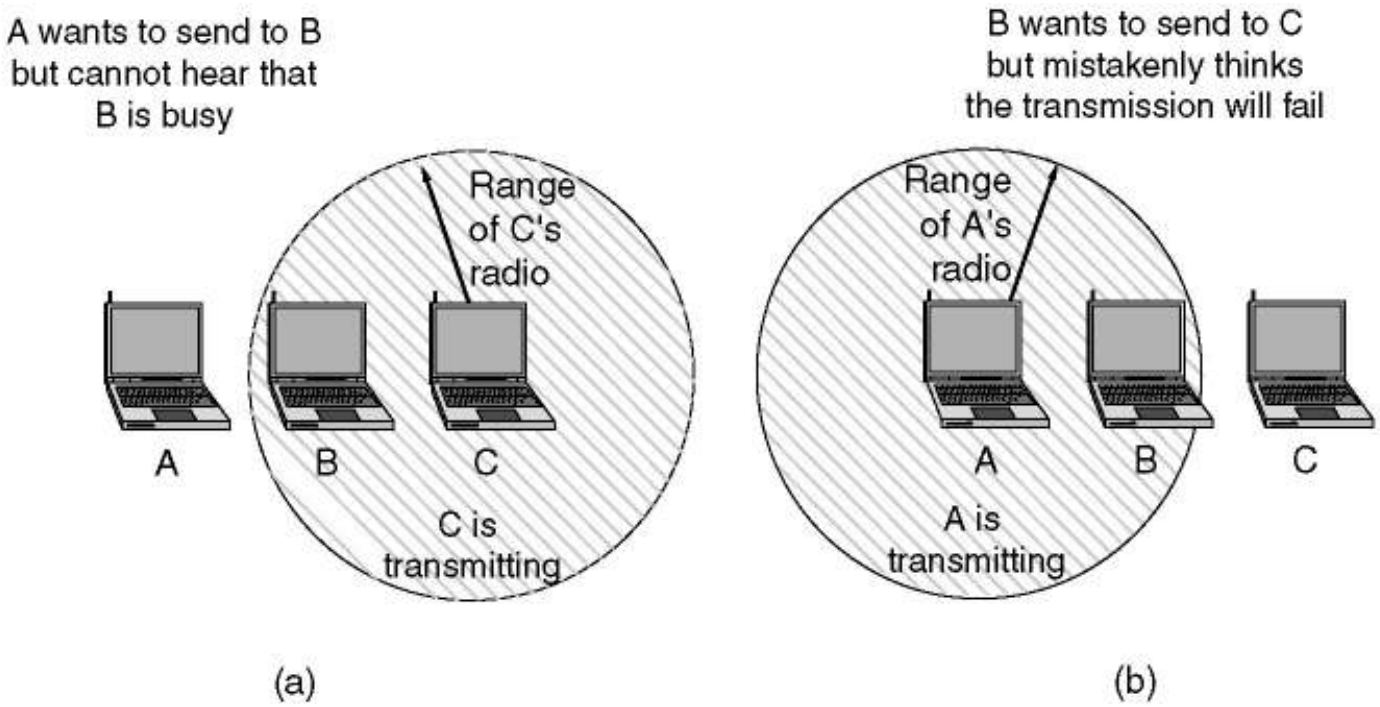

(a) The hidden station problem.

(b) The exposed station problem.

The 802.11 MAC Sublayer Protocol (5)

- There is inverse problem, the exposed station problem.
- Station B wants to send to C so it listens to the channel. When it hears a transmission, it falsely concludes that it may not send to C, even though A may be transmitting to D (not shown).

The 802.11 MAC Sublayer Protocol (6)

- In addition, most radios are half duplex, meaning that they cannot transmit and listen for noise bursts at the same time on a single frequency.
- As a result of these problems, 802.11 does not use CSMA/CD (Carrier Sense Multiple Access Protocols with Collision Detection), as Ethernet does.

The 802.11 MAC Sublayer Protocol (7)

To deal with this problem, 80.11 supports two modes of operation.

1) DCF (Distributed Coordination Function) does not use any kind of central control (in that respect, similar to Ethernet).
2) PCF (Point Coordination Function) uses the base station to control all activity in its cell.

The 802.11 MAC Sublayer Protocol (8)

- When DCF is employed, 802.11 uses a protocol called CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance)
- In this protocol, both physical channel sensing and virtual channel sensing are used.
- Two methods of operation are supported by CSMA/CA

The 802.11 MAC Sublayer Protocol (9)

- In the first method, when a station wants to transmit, it senses the channel.
- If it is idle, it just starts transmitting.
- It does not sense the channel while transmitting but emits its entire frame, which may well be destroyed at the receiver due to interference there.

The 802.11 MAC Sublayer Protocol (10)

- If the channel is busy, the sender defers until it goes idle and then starts transmitting.
- If a collision occurs, the colliding stations wait a random time, using the Ethernet binary exponential backoff algorithm, and then try again later.

The 802.11 MAC Sublayer Protocol (11)

- The other mode of CSMA/CA operation is based on MACAW (Multiple Access with Collision Avoidance for Wireless) and uses virtual channel sensing, as illustrated in following figure.
- In this example, A wants to send to B. C is a station within range of A (and possibly within range of B, but that does not matter). D is a station within range of B but not within range of A.

The 802.11 MAC Sublayer Protocol (12)

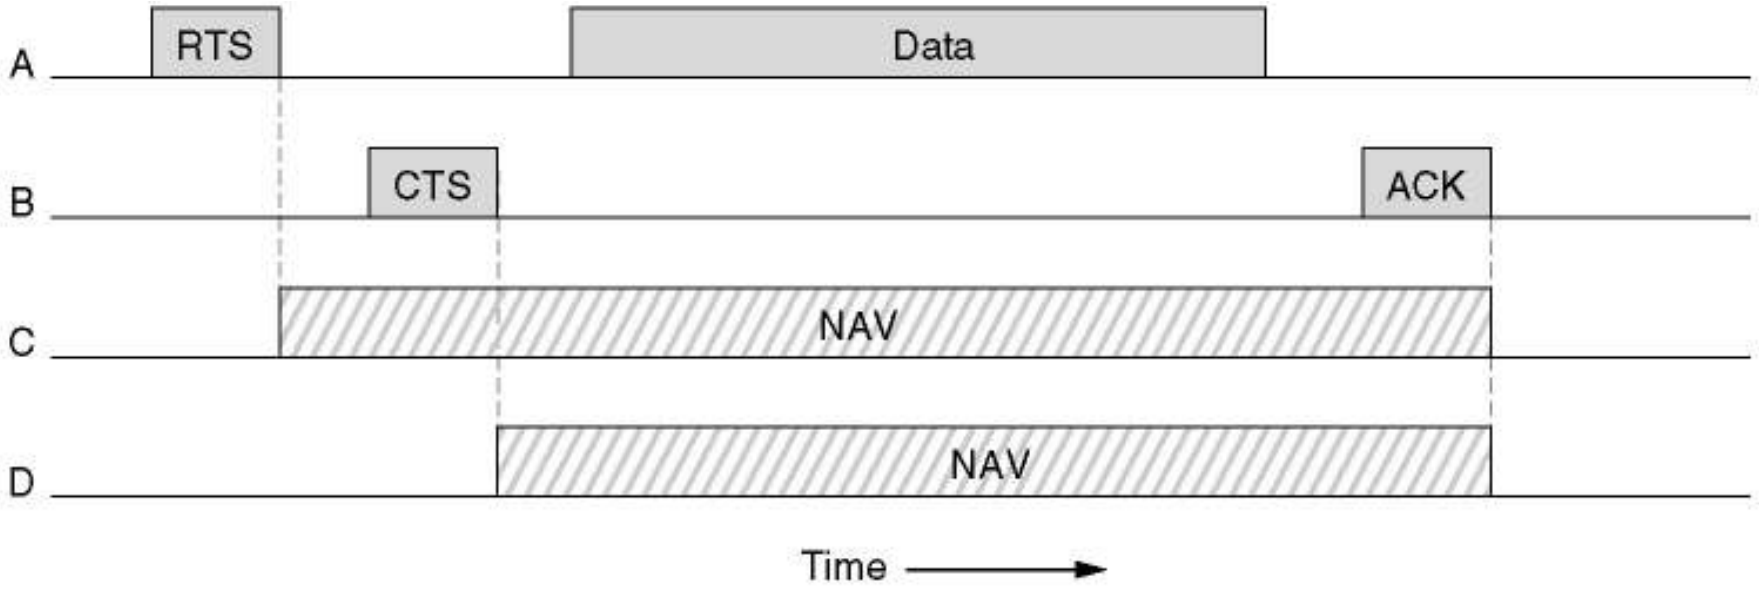

The use of virtual channel sensing using CSMA/CA.

The 802.11 MAC Sublayer Protocol (13)

- The protocol starts when A decides it wants to send data to B. It begins by sending an RTS (Request To Send) frame to B to request permission to send it a frame.
- When B receives this request, it may decide to grant permission, in which case it sends a CTS (Clear To Send) frame back.
- Upon receipt of the CTS, A now sends its frame and starts an ACK timer.

The 802.11 MAC Sublayer Protocol (14)

- Upon correct receipt of the data frame, B responds with an ACK frame, terminating the exchange.
- If A's ACK timer expires before the ACK gets back to it, the whole protocol is run again.

The 802.11 MAC Sublayer Protocol (15)

- Now let us consider this exchange from the viewpoints of C and D. C is within range of A, so it may receive the RTS frame. If it does, it realizes that someone is going to send data soon, so for the good of all it desists from transmitting anything until the exchange is completed.

The 802.11 MAC Sublayer Protocol (16)

- From the information provided in the RTS request, it can estimate how long the sequence will take, including the final ACK, so it asserts a kind of virtual channel busy for itself, including by NAV (Network Allocation Vector).
- D does not hear the RTS, but it does hear the CTS, so it also asserts the NAV signal for itself.

The 802.11 MAC Sublayer Protocol (17)

- To deal with the problem of noisy channel, 802.11 allows frames to be fragmented into smaller pieces, each with its own checksum.
- The fragments are individually numbered and acknowledged using a stop-and-wait protocol (i.e., the sender may not transmit fragment k+1 until it has received the acknowledgment for fragment k)

The 802.11 MAC Sublayer Protocol (18)

- Once the channel has been acquired using RTS and CTS, multiple fragments can be sent in row, sequence of fragments is called a fragment burst.
- Fragmentation increase the throughput by restricting retransmissions to the bad fragments rather than the entire frame.

The 802.11 MAC Sublayer Protocol (19)

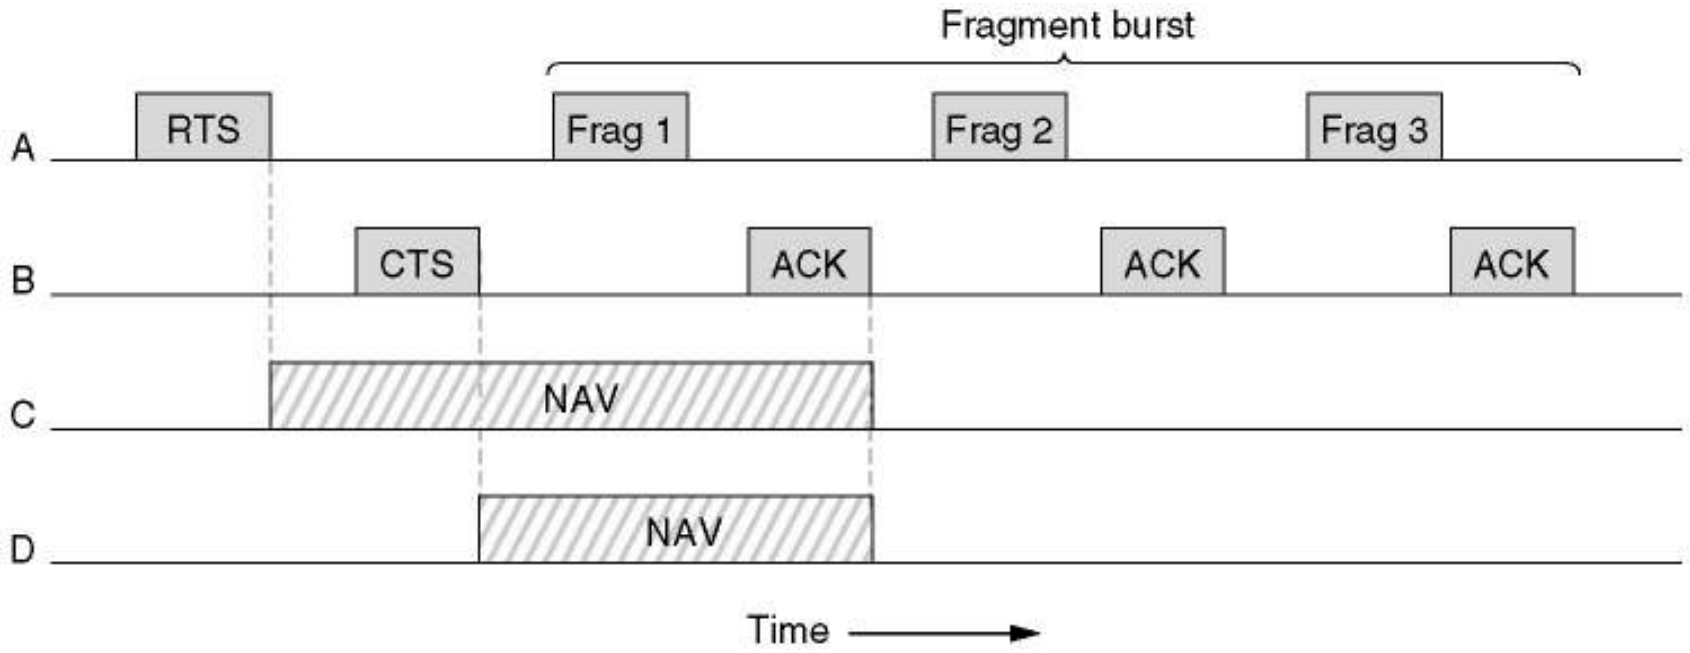

A fragment burst.

The 802.11 MAC Sublayer Protocol (20)

- PCF (in which the base station polls the other stations, asking them if they have any frames to send) and DCF (in which there is no central control) can coexist within one cell.
- At first it might seem impossible to have central control and distributed control operating at the same time, but 802.11 provides a way to achieve this goal.

The 802.11 MAC Sublayer Protocol (21)

- It works by carefully defining the interframe time interval.
- After a frame has been sent, a certain amount of dead time is required before any station may send a frame.
- Four different intervals are defined, each for a specific purpose.

The 802.11 MAC Sublayer Protocol (22)

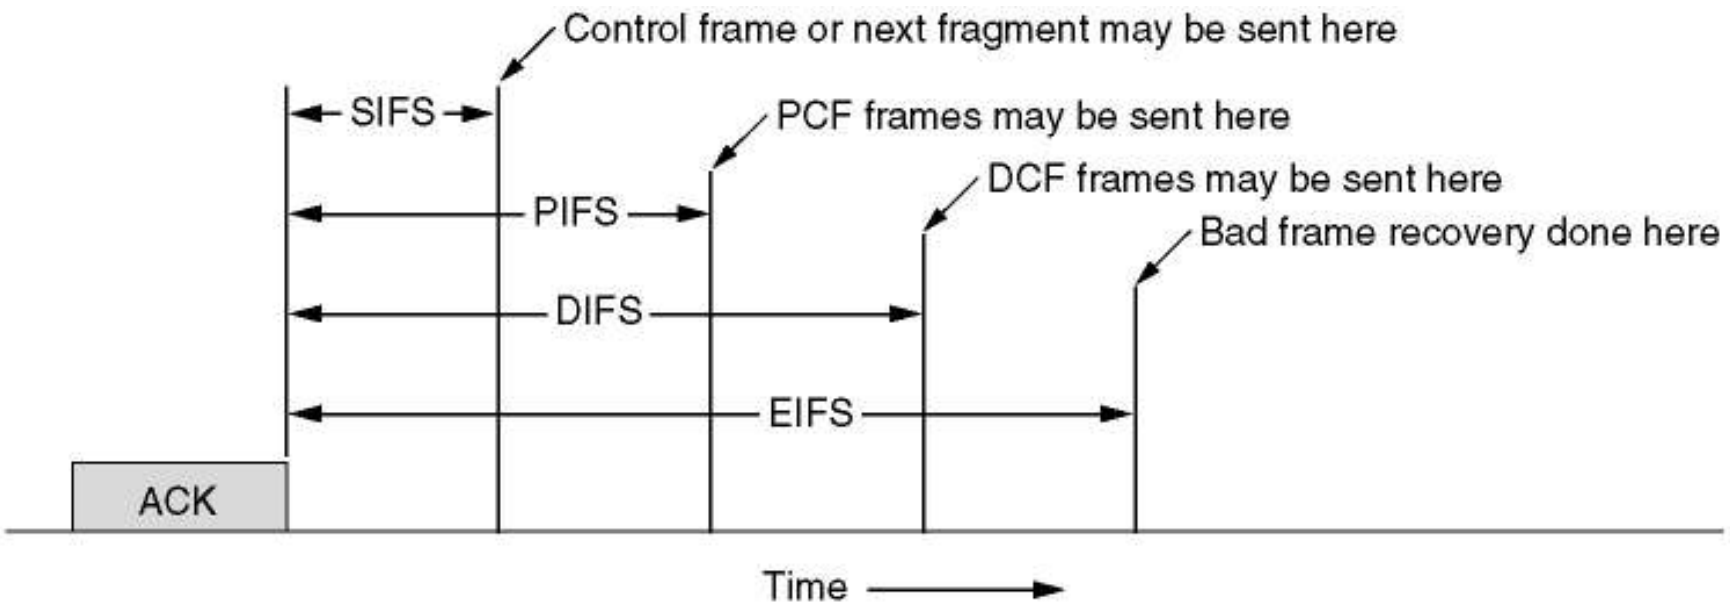

Interframe spacing in 802.11.

The 802.11 MAC Sublayer Protocol (23)

- SIFS (Short InterFrame Spacing) is used to allow the parties in a single dialog the chance to go first.
- PIFS (PCF InterFrame Spacing)
- DIFS (DCF InterFrame Spacing)
- EIFS (Extended InterFrame Spacing)

The 802.11 Frame Structure (1)

- The 802.11 standards defines three different classes of frames on the wire: data, control and management.
- Each of these has a header with a variety of fields used within the MAC sublayer.
- In addition, there are some headers used by the physical layer but these mostly deal with the modulation techniques used.

The 802.11 Frame Structure (2)

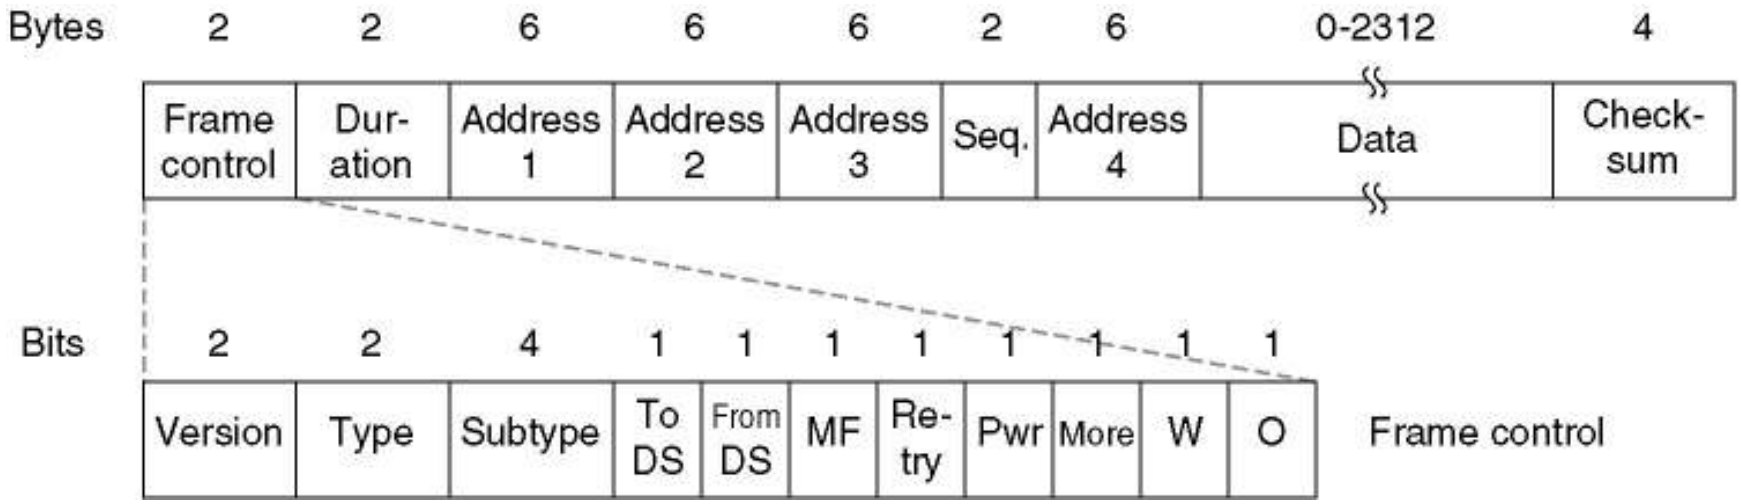

The 802.11 data frame.

802.11 Services (1)

- The 802.11 standard states that each conformant wireless LAN must provide nine services.
- These services are divided into two categories:

a) Five distribution services

b) Four station services.

802.11 Services (2)

- The distribution services relate to managing cell membership and interacting with stations outside the cell.
- The five distribution services are provided by the base station and deal with station mobility as they enter and leave cells, attaching themselves to and detaching themselves from base stations.

802.11 Services (3)

Distribution Services

- Association
- Disassociation
- Reassociation
- Distribution
- Integration

802.11 Services (4)

- Association

This service is used by mobile stations to connect themselves to base stations. The base station may accept or reject the mobile station. If the mobile station is accepted, it must then authenticate itself.

802.11 Services (5)

- Disassociation

Either the station or the base station may disassociate, thus breaking the relationship.

- Reassociation

A station may change its preferred base station using this service

802.11 Services (6)

- Distribution

This service determines how to route frames sent to the base station.

- Integration

If a frame needs to be sent through a non - 802.11 network with a different addressing scheme or frame format, this service handles the translation from the 802.11 format to the format required by the destination network

802.11 Services (7)

- Station services relate to activity within a single cell.
- They are used after association has taken place and are as follows.

802.11 Services (8)

Intracell Services

- Authentication
- Deauthentication
- Privacy
- Data Delivery

802.11 Services (9)

- Authentication

Because wireless communication can easily be sent or received by unauthorized station, a station must authenticate itself before it is permitted to send data.

- Deauthentication

When a previously authenticated station wants to leave the network, it is deauthenticated.

802.11 Services (10)

- Privacy

For information sent over a wireless LAN to be kept confidential, it must be encrypted. This service manages the encryption and decryption.

- Data Delivery

Finally, data transmission is what it is all about, so 802.11 naturally provides a way to transmit and receive data.

4.5. BROADBAND WIRELESS (1)

- Let us now go outside and see if any interesting networking is going on there.
- With the deregulation of the telephone system in many countries, competitors to the entrenched telephone company are now often allowed to offer local voice and high-speed Internet service.
- There is certainly plenty of demand.

Broadband Wireless (2)

- The problem is that running fiber, coax, or even category 5 twisted pair to millions of homes and businesses is prohibitively expensive.
- What is a competitor to do?
- The answer is broadband wireless.
- Erecting a big antenna on a hill just outside of town and installing antennas directed at it on customers' roofs is much easier and cheaper than digging trenches and stringing cables.

Broadband Wireless (3)

- In April 2002, IEEE completed the 802.16 Standard named "Air Interface for Fixed Broadband Wireless Access Systems"
- However, some people prefer to call it a wireless MAN or a wireless local loop

Broadband Wireless (4)

- Comparison of 802.11 and 802.16
- The 802.16 Protocol Stack
- The 802.16 Physical Layer
- The 802.16 MAC Sublayer Protocol
- The 802.16 Frame Structure

Broadband Wireless (5)

- Comparison of 802.11 and 802.16
- The environments in which 802.11 and 802.16 operate are similar in some ways, primarily in that they were designed to provide high-bandwidth wireless communications
- But they also differ in some major ways.

Broadband Wireless (6)

- 802.16 provides service to buildings, and buildings are not mobile. They do not migrate from cell to cell often.
- 802.16 uses full-duplex communication
- 802.11 was designed to be mobile Ethernet, whereas 802.16 was designed to be wireless, but stationary, cable television

The 802.16 Protocol Stack

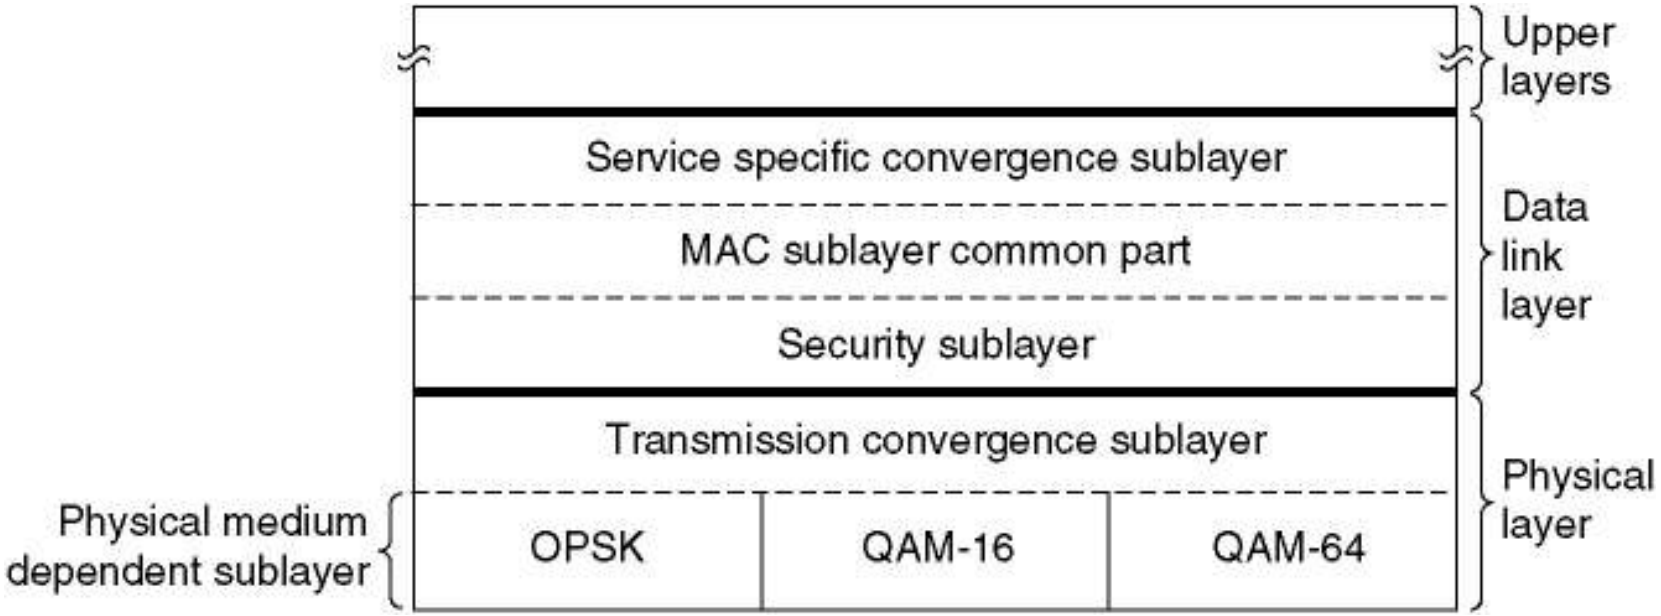

The 802.16 Protocol Stack.

Modems

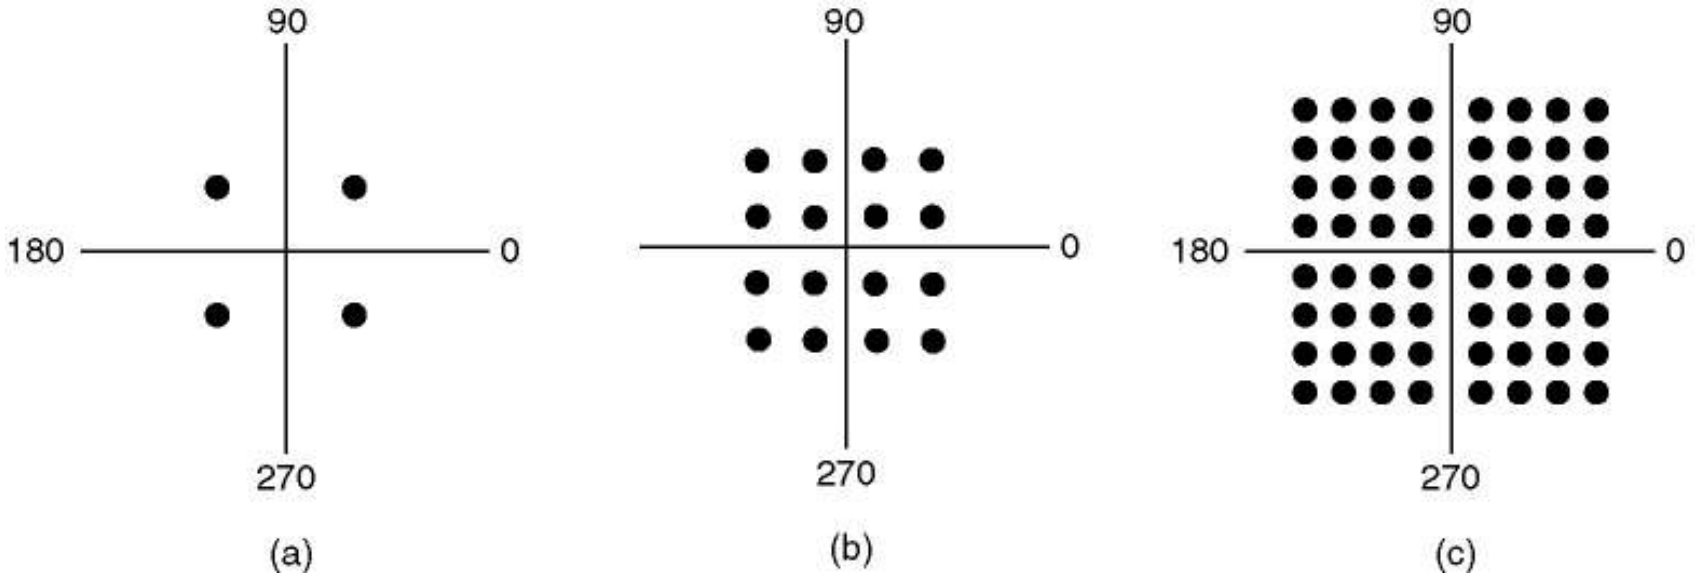

(a) QPSK.

(b) QAM-16.

(c) QAM-64.

The 802.16 Physical Layer

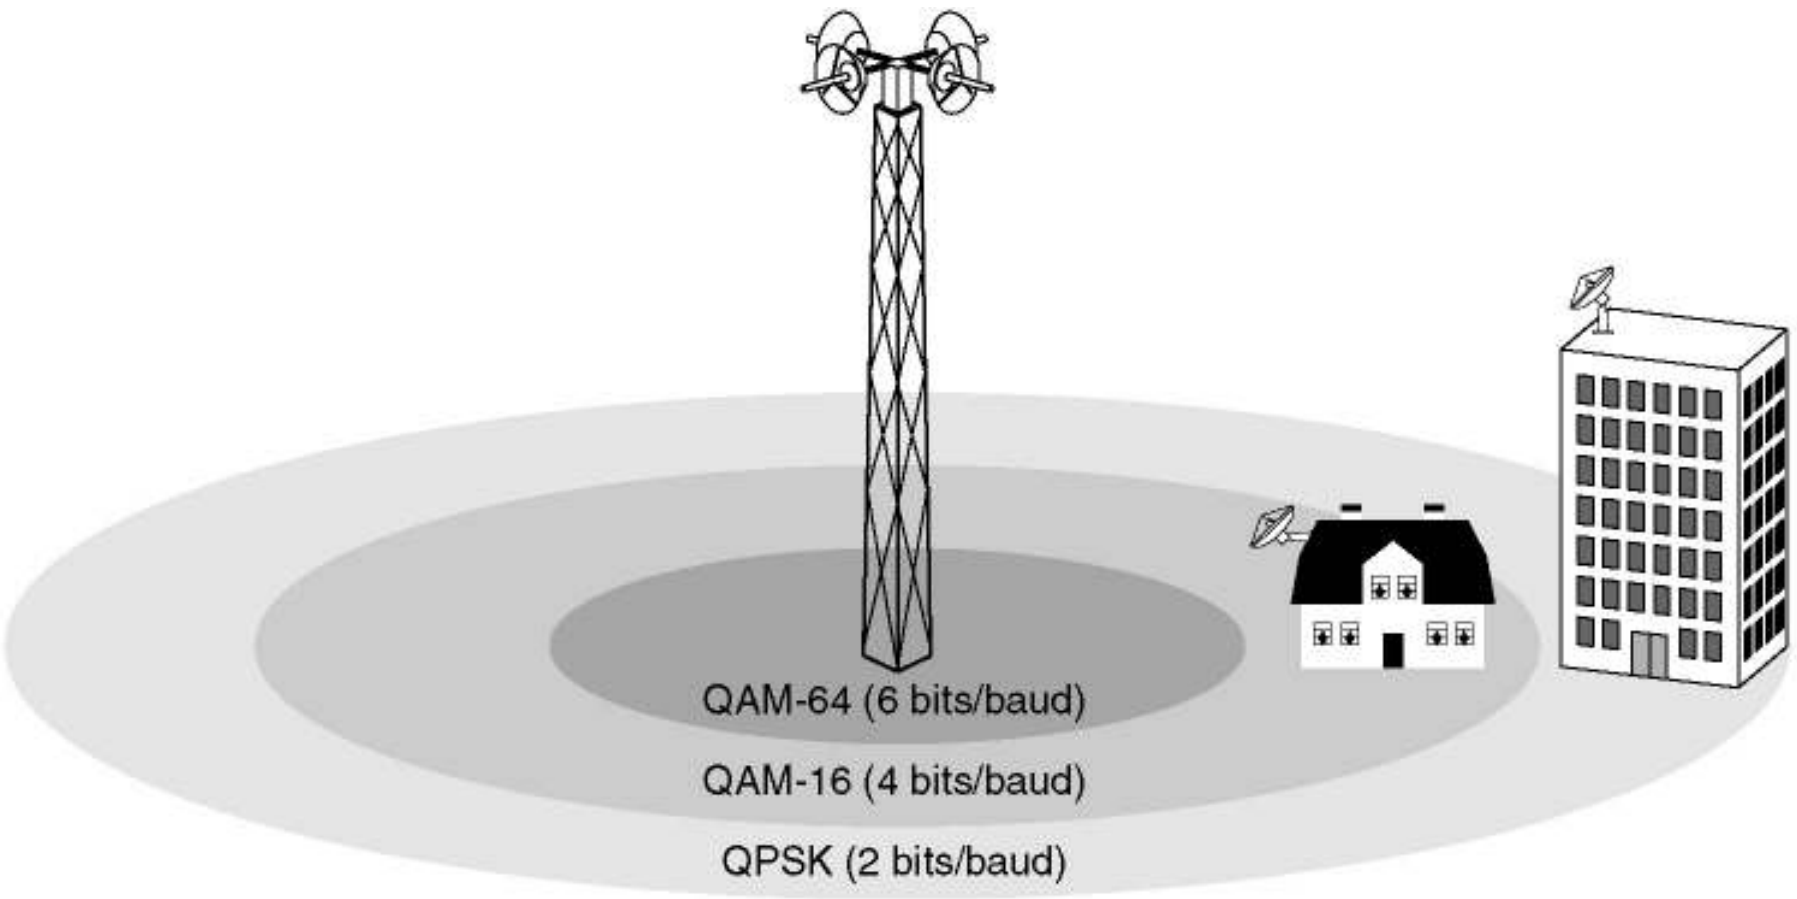

The 802.16 transmission environment.

The 802.16 Physical Layer (2)

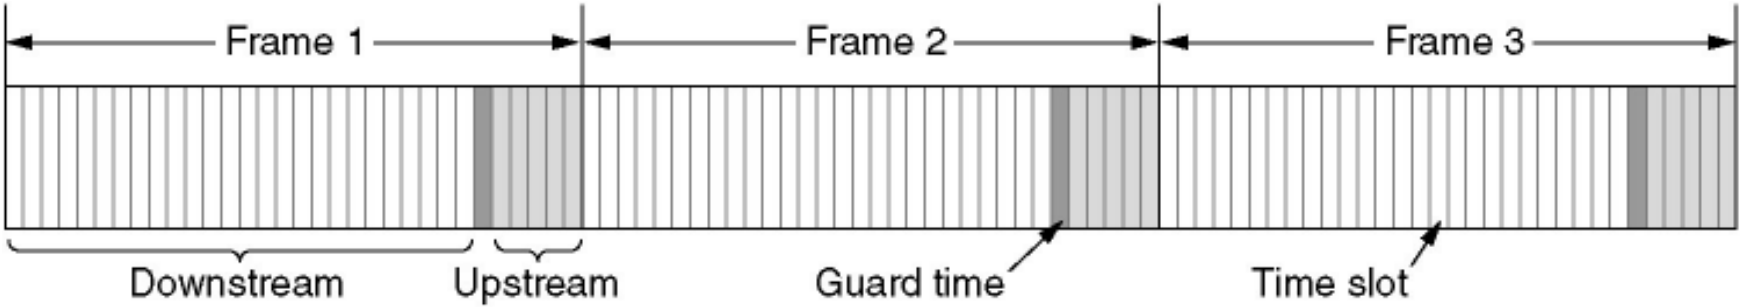

Frames and time slots for time division duplexing.

The 802.16 MAC Sublayer Protocol

Service Classes

- Constant bit rate service
- Real-time variable bit rate service
- Non-real-time variable bit rate service
- Best efforts service

The 802.16 Frame Structure

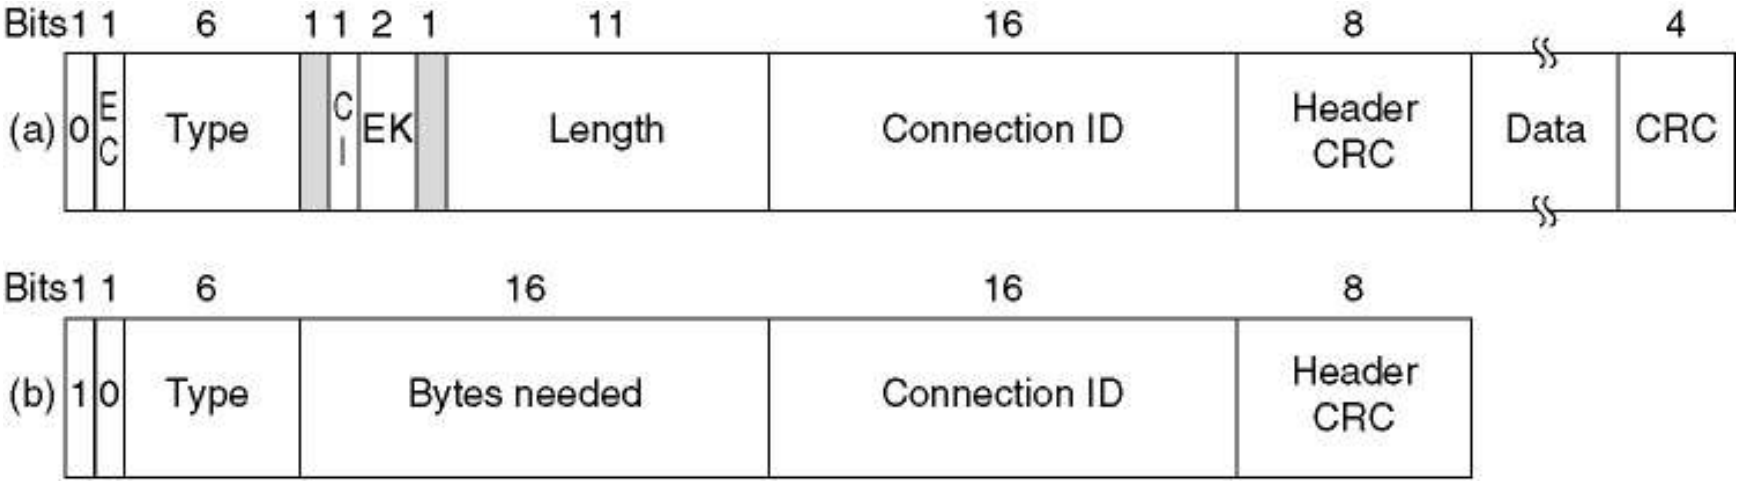

(a) A generic frame. (b) A bandwidth request frame.

4.6. BLUETOOTH (1)

- Bluetooth is a wireless standard for interconnecting computing and communication devices and accessories using short-range, low-power, inexpensive wireless radios.
- In July 1999 the Bluetooth SIG issued a 1500 page specification of V1.0.
- Shortly thereafter, IEEE has designed 802.15 – wireless personnel area networks

Bluetooth (2)

- Bluetooth Architecture
- Bluetooth Applications
- The Bluetooth Protocol Stack
- The Bluetooth Radio Layer
- The Bluetooth Baseband Layer
- The Bluetooth L2CAP Layer
- The Bluetooth Frame Structure

Bluetooth Architecture

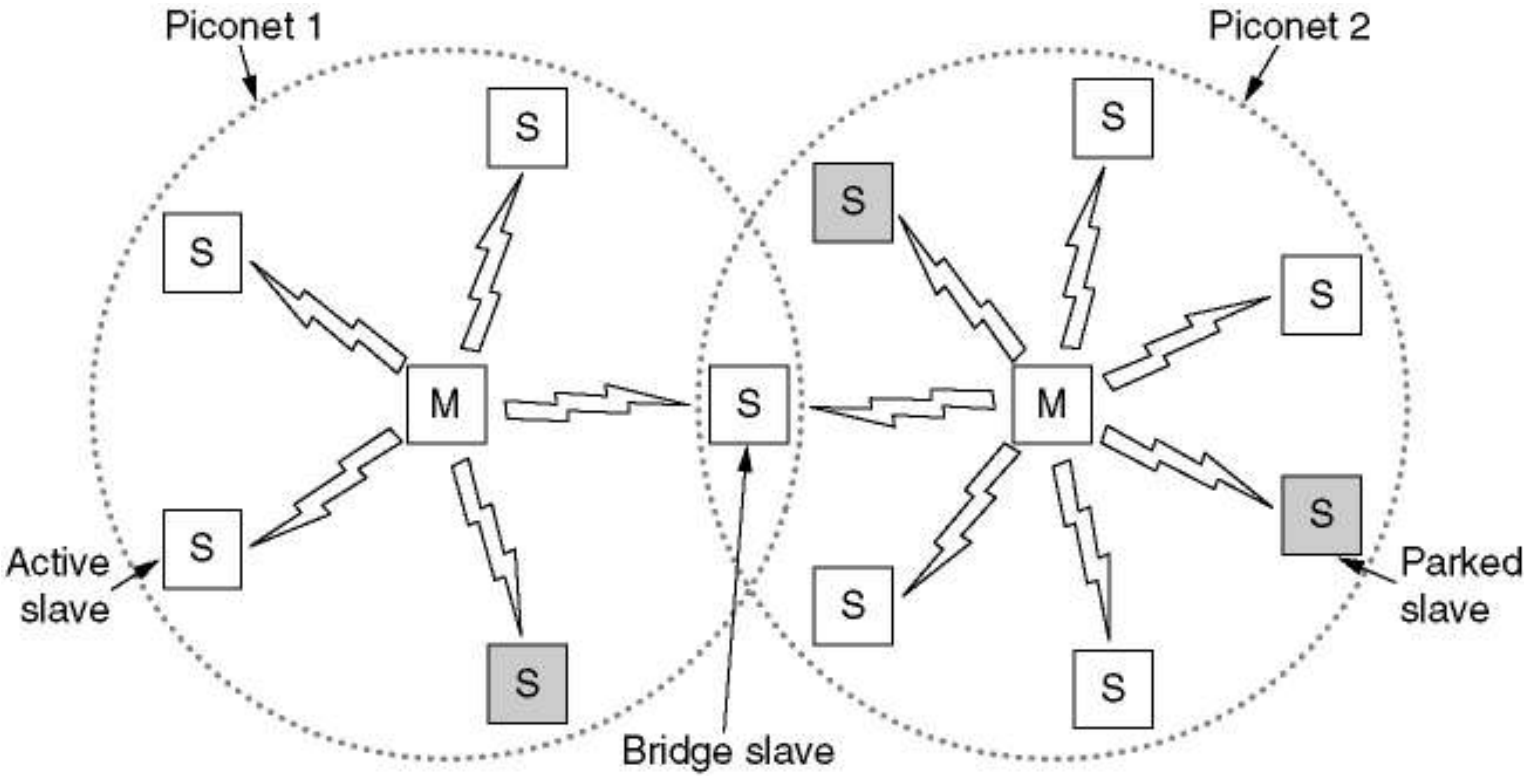

Two piconets can be connected to form a scatternet.

Bluetooth Applications

Name	Description
Generic access	Procedures for link management
Service discovery	Protocol for discovering offered services
Serial port	Replacement for a serial port cable
Generic object exchange	Defines client-server relationship for object movement
LAN access	Protocol between a mobile computer and a fixed LAN
Dial-up networking	Allows a notebook computer to call via a mobile phone
Fax	Allows a mobile fax machine to talk to a mobile phone
Cordless telephony	Connects a handset and its local base station
Intercom	Digital walkie-talkie
Headset	Intended for hands-free voice communication
Object push	Provides a way to exchange simple objects
File transfer	Provides a more general file transfer facility
Synchronization	Permits a PDA to synchronize with another computer

The Bluetooth profiles.

The Bluetooth Protocol Stack

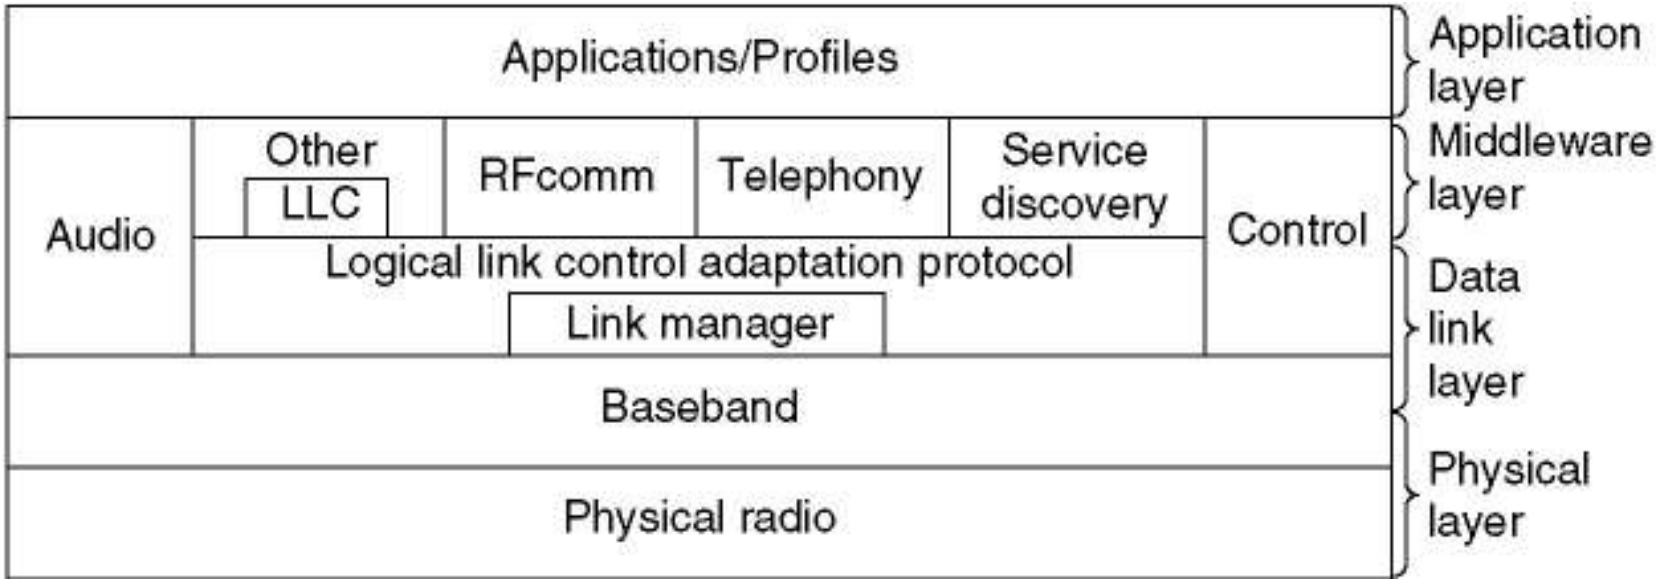

The 802.15 version of the Bluetooth protocol architecture.

The Bluetooth Frame Structure

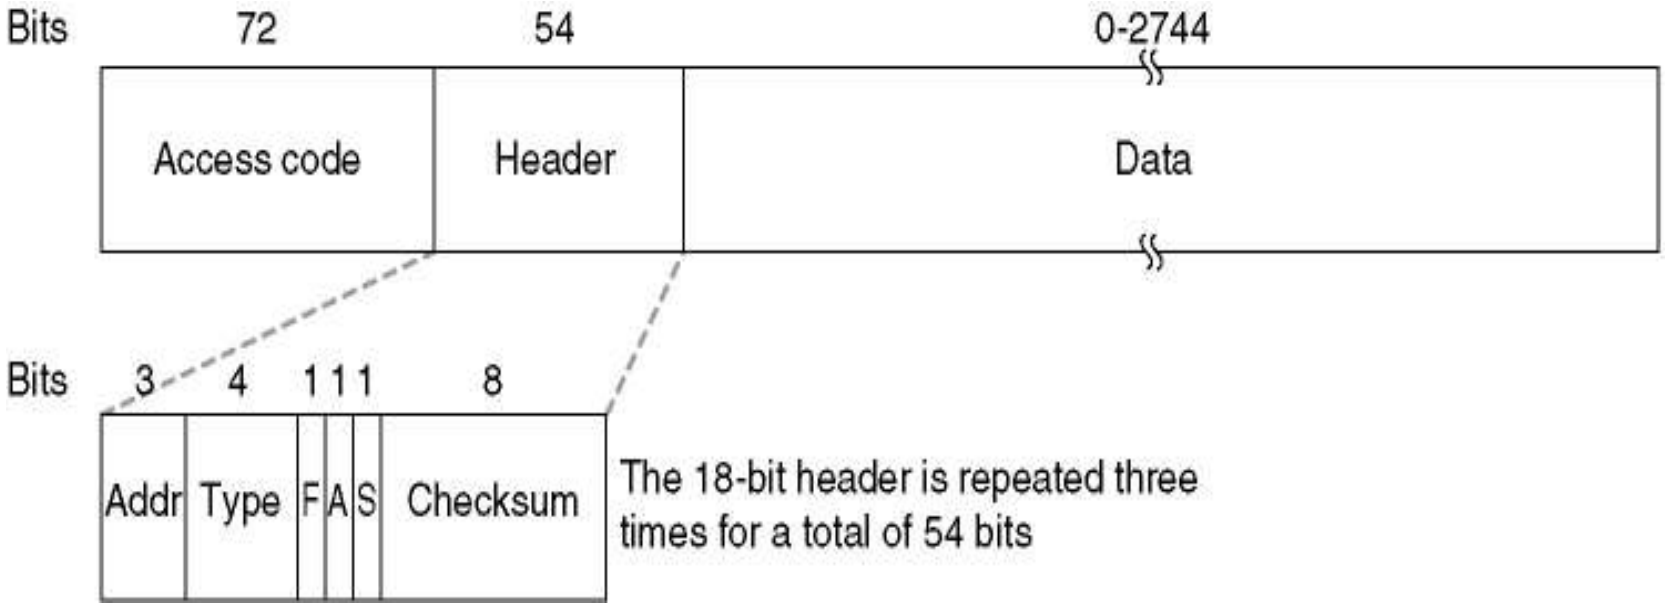

A typical Bluetooth data frame.

4.7.DATA LINK LAYER SWITCHING (1)

- Many organizations have multiple LANs and wish to connect them.

- LANs can be connected by devices called bridges, which operate in the data link layer.

- Bridges examine the data layer link addresses to do routing

Data Link Layer Switching (2)

- Since they are not supposed to examine the payload field of the frames they route, they can transport IPv4 (used in the Internet now), IPv6 (will be used in the Internet in the future), AppleTalk, ATM, OSI, or any other kinds of packets.
- In contrast, routers examine the addresses in packets and route based on them.

Data Link Layer Switching (3)

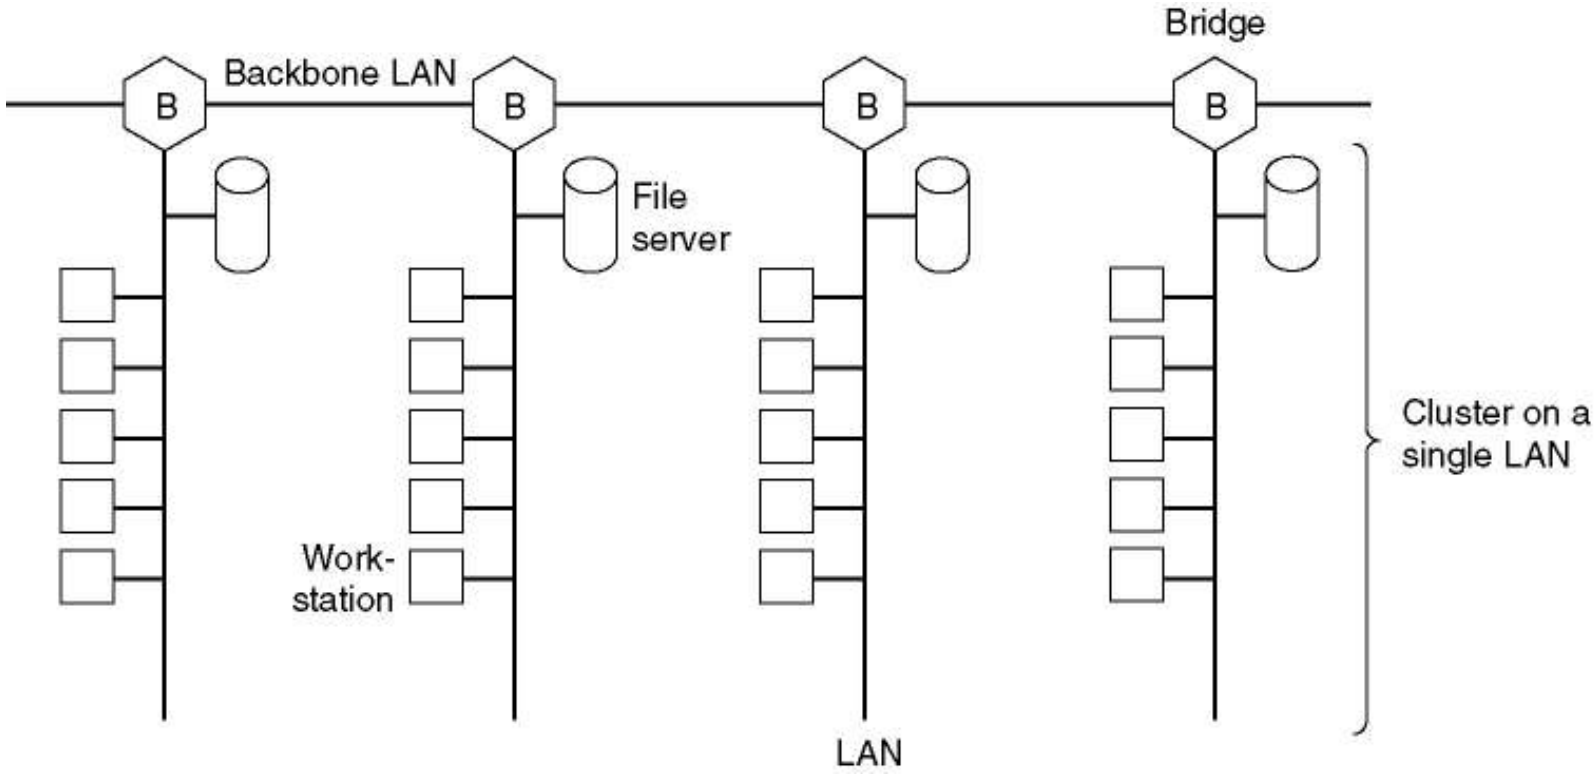

Multiple LANs connected by a backbone to handle a total load higher than the capacity of a single LAN.

Data Link Layer Switching (4)

The technology of bridges

- Bridges from 802.x to 802.y
- Local Internetworking
- Spanning Tree Bridges
- Remote Bridges
- Repeaters, Hubs, Bridges, Switches, Routers, Gateways
- Virtual LANs

Bridges from 802.x to 802.y (1)

- Following figure illustrates the operation of a simple two-port bridge.
- Host A on a wireless (802.11) LAN has a packet to send to a fixed host, B, on an (802.3) Ethernet to which the wireless LAN is connected.
- Note that a bridge connecting k different LANs will have k different MAC sublayers and k different physical layers, one for each type.

Bridges from 802.x to 802.y (2)

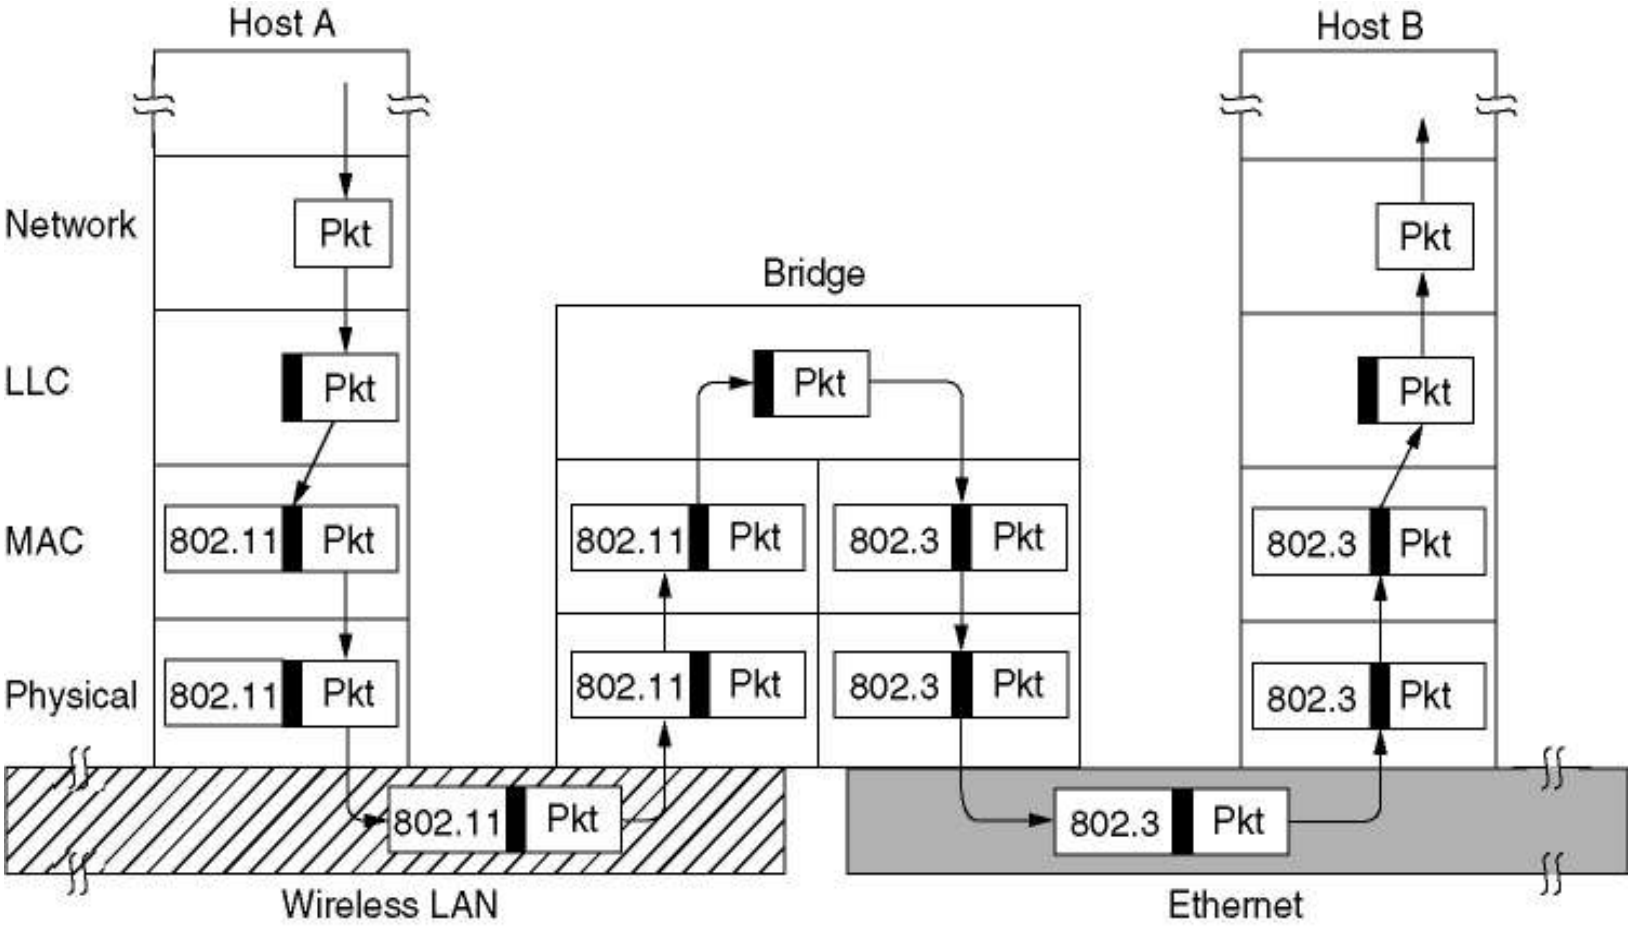

Operation of a LAN bridge from 802.11 to 802.3.

Bridges from 802.x to 802.y (3)

- There are many difficulties that one encounters when trying to build a bridge between the various 802 LANs (and MANs)
- Each of the LANs uses a different frame format

Bridges from 802.x to 802.y (4)

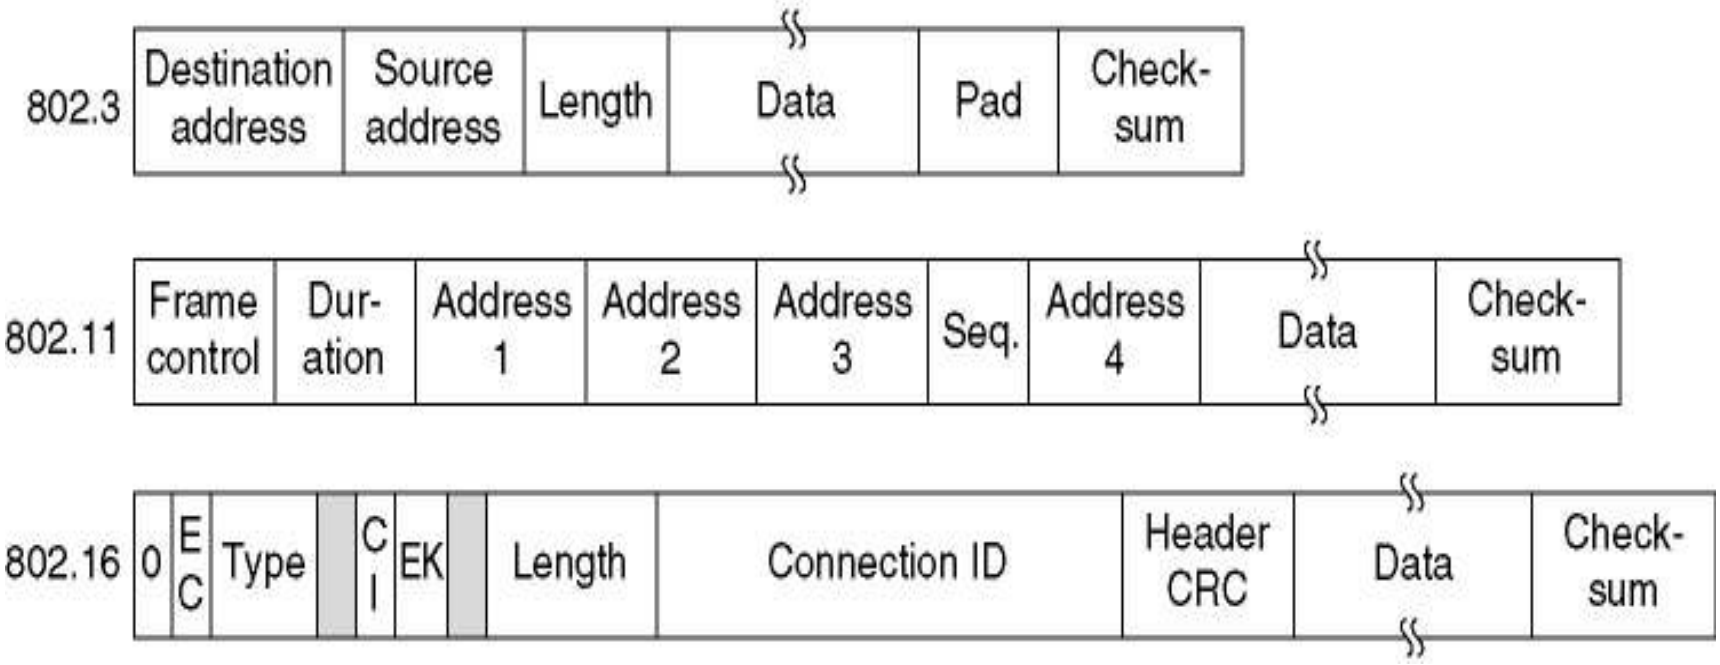

The IEEE 802 frame formats. The drawing is not to scale.

Bridges from 802.x to 802.y (5)

Some problems are :

- The first is any copying between different LANs requires reformatting
- The second is data rate problem
- The third is different 802 LANs have different maximum frame lengths.
- Another is the security problem
- Last is quality of service.

Local Internetworking (1)

- The bridges should be completely transparent (invisible to all the hardware and software)
- A transparent bridge operates in promiscuous mode, accepting every frame transmitted on all the LANs to which it is attached.

Local Internetworking (2)

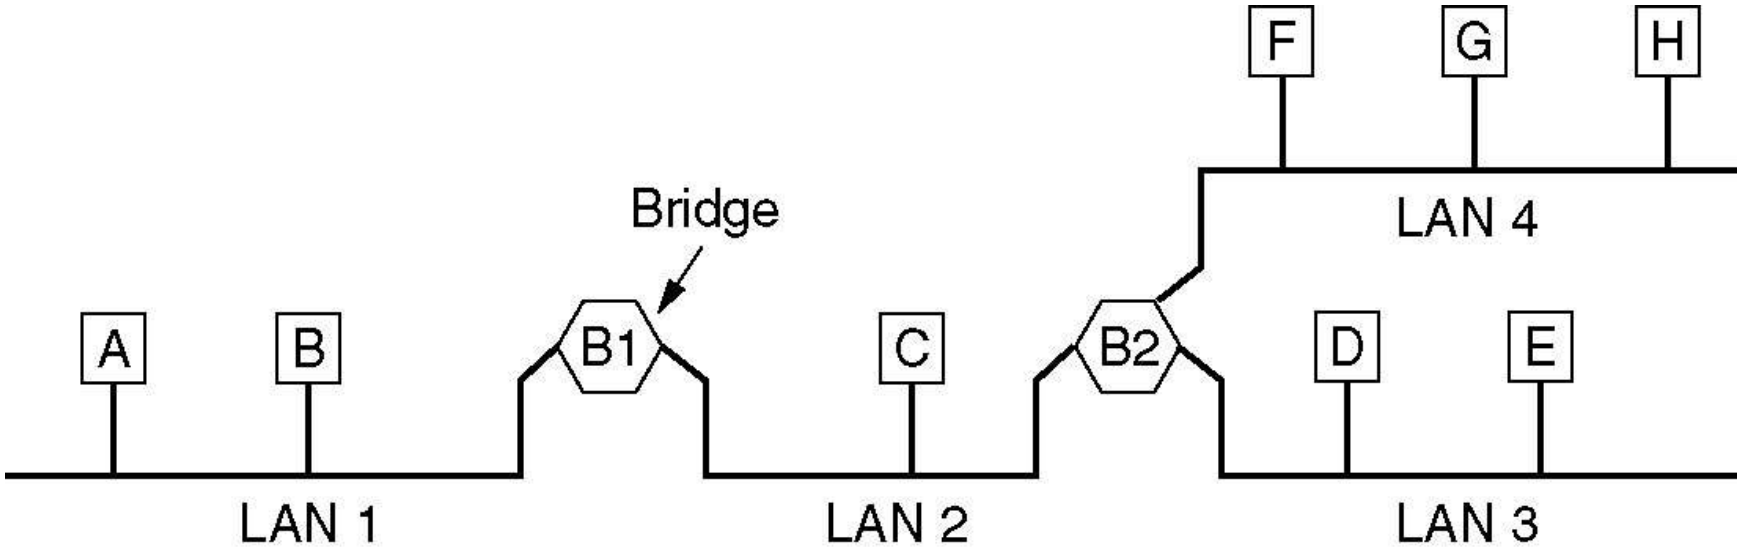

A configuration with four LANs and two bridges.

Local Internetworking (3)

As each frame arrives, this algorithm must be applied (by using a big table inside the bridge by special purpose VLSI chip):

- If destination and source LANs are the same, discard the frame.
- If destination and source LANs are the different, forward the frame.
- If destination LAN is unknown, use flooding.

Spanning Tree Bridges (1)

- To increase reliability, some sites use two or more bridges in parallel between pairs of LANs.

- This arrangement, however, also introduces some additional problems because it creates loops in the topology.

Spanning Tree Bridges (2)

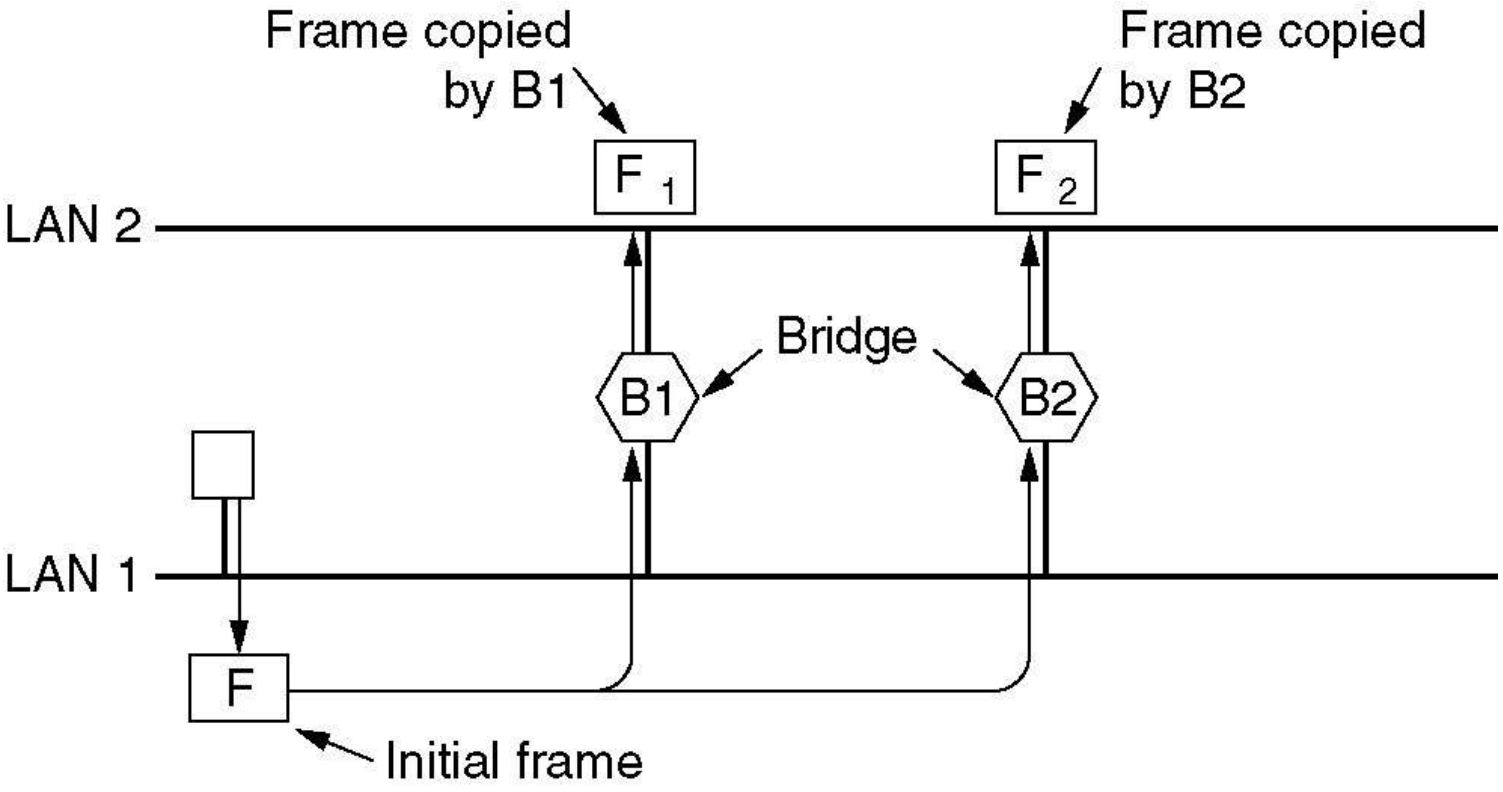

Two parallel transparent bridges.

Spanning Tree Bridges (3)

- Let us see how a frame, F, with unknown destination is handled.
- Each bridge uses flooding, which in this example just means copying it to LAN 2.
- Shortly thereafter, bridge 1 sees F2, a frame with an unknown destination, which it copies to LAN 1, generating F3 (not shown)
- Similarly, bridge 2 copies F1 to LAN 1 generating F4 (also not shown) ….

Spanning Tree Bridges (4)

- This problem can be solved by using a spanning tree that reaches every LAN.
- The following configuration can be abstracted into a graph with the LANs as the nodes.
- An arc connects any two LANs that are connected by a bridge
- This graph can be reduced to a spanning tree by dropping the arcs shown as dotted lines

Spanning Tree Bridges (5)

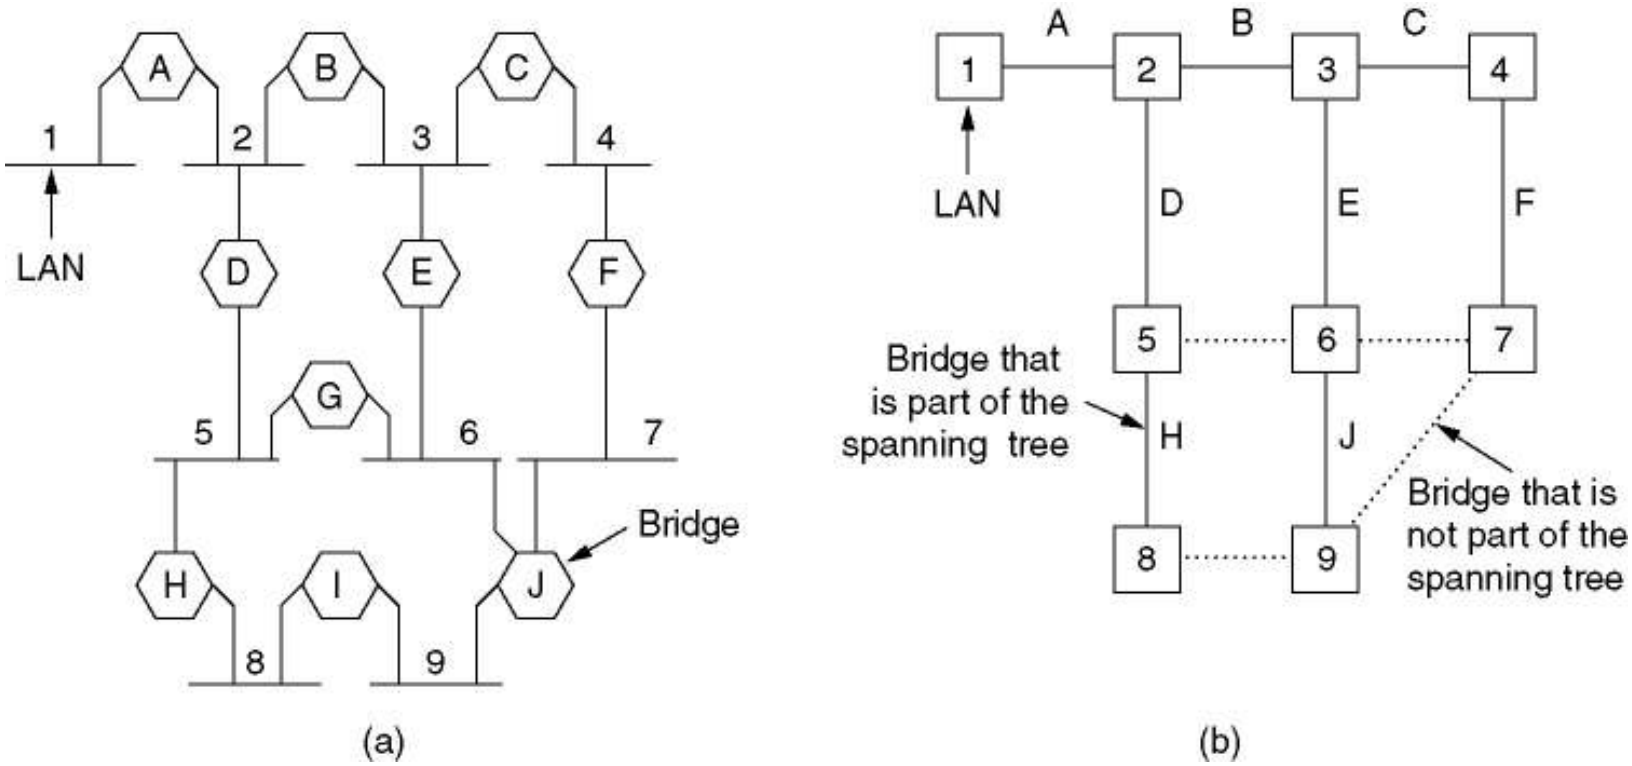

(a) Interconnected LANs. (b) A spanning tree covering the LANs. The dotted lines are not part of the spanning tree.

Remote Bridges (1)

- All the LANs belonging to a company should be interconnected, so the complete system acts like one large LAN.
- This goal can be achieved by putting a bridge on each LAN and connecting the bridges pairwise with point-to-point lines (e.g., lines leased from a telephone company).

Remote Bridges (2)

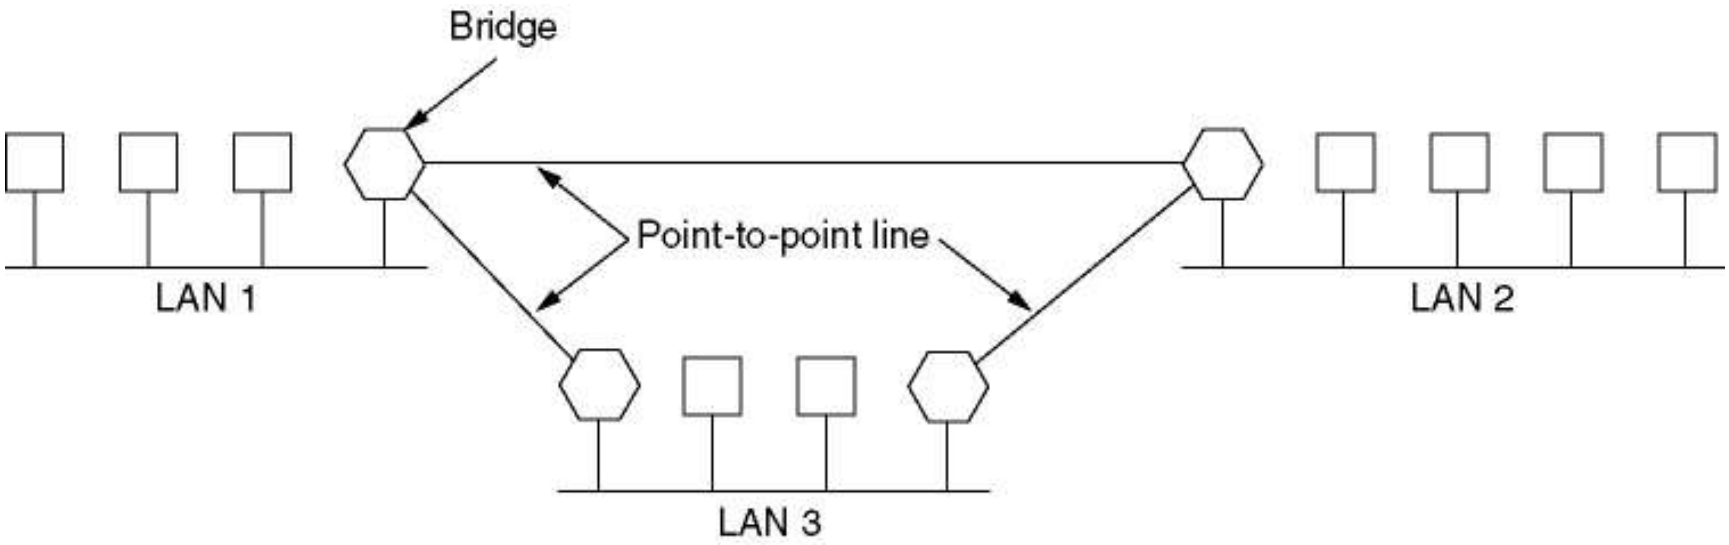

Remote bridges can be used to interconnect distant LANs.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (1)

- So far in this book we have looked at a variety of ways to get frames and packets from one cable segment to another.
- We have mentioned repeaters, bridges, switches, hubs, routers, and gateways.
- These devices operate in different layers.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (2)

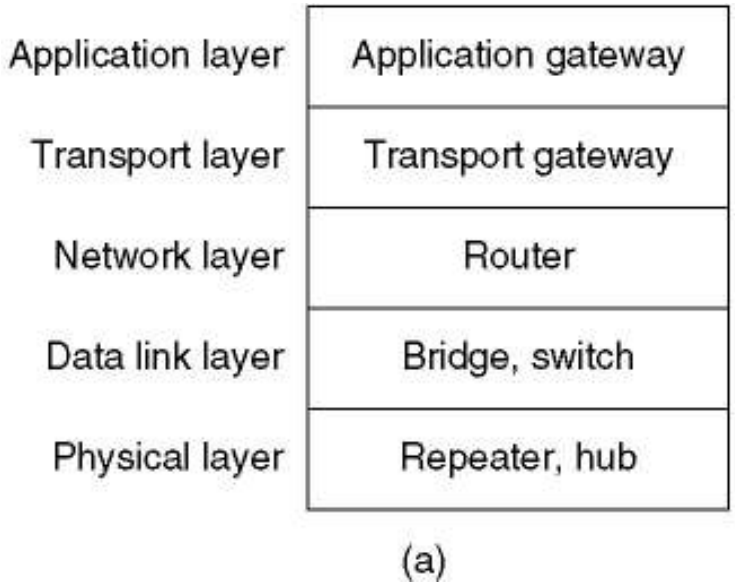

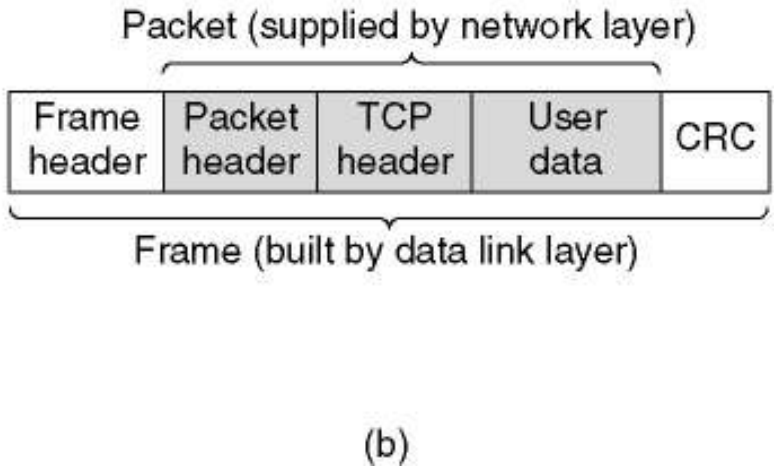

(a) Which device is in which layer.

(b) Frames, packets, and headers.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (3)

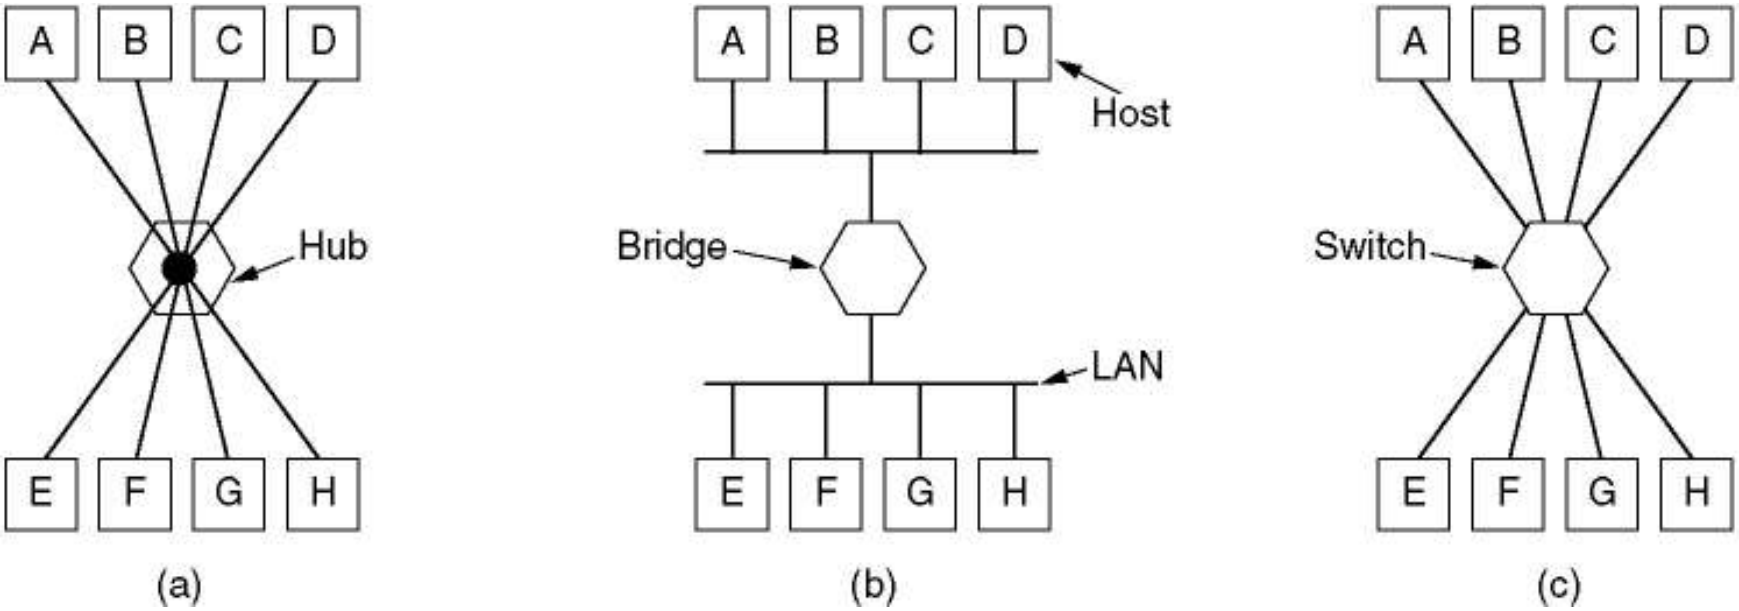

(a) A hub. (b) A bridge. (c) a switch.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (4)

- Repeaters are analog devices that are connected to two cable segments. A signal appearing on one of them is amplified and put out on the other.
- Hubs have a number of inputs lines that it joins electrically. Frames arriving on any of the lines are sent out on all the others.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (5)

- Bridges connect two or more LANs. When a frame arrives, software in the bridge extracts the destination address from the frame header and looks it up in a table to see where to send the frame.
- Switches are similar to bridges in that both route on frame addresses. The main difference is that a switch is most often used to connect individual computers.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (6)

- Routers are different from all of the above. When a packet comes into a router, the frame header and trailer are stripped off and the packet located in the frame's payload field is passed to the routing software. This software uses the packet header to choose an output line.

Repeaters, Hubs, Bridges, Switches, Routers and Gateways (7)

- Transport gateways connect two computers that use different connection – oriented transport protocols. The transport gateway can copy the packets from one connection to the other, reformatting them as need be.
- Application gateways understand the format and contents of the data and translate messages from one format to another.

Virtual LANs (1)

- With the advent of 10Base-T and hubs in the 1990s, buildings were rewired to rip out all the yellow garden hoses and install twisted pairs from every office to central wiring closets at the end of each corridor or in a central machine room.

Virtual LANs (2)

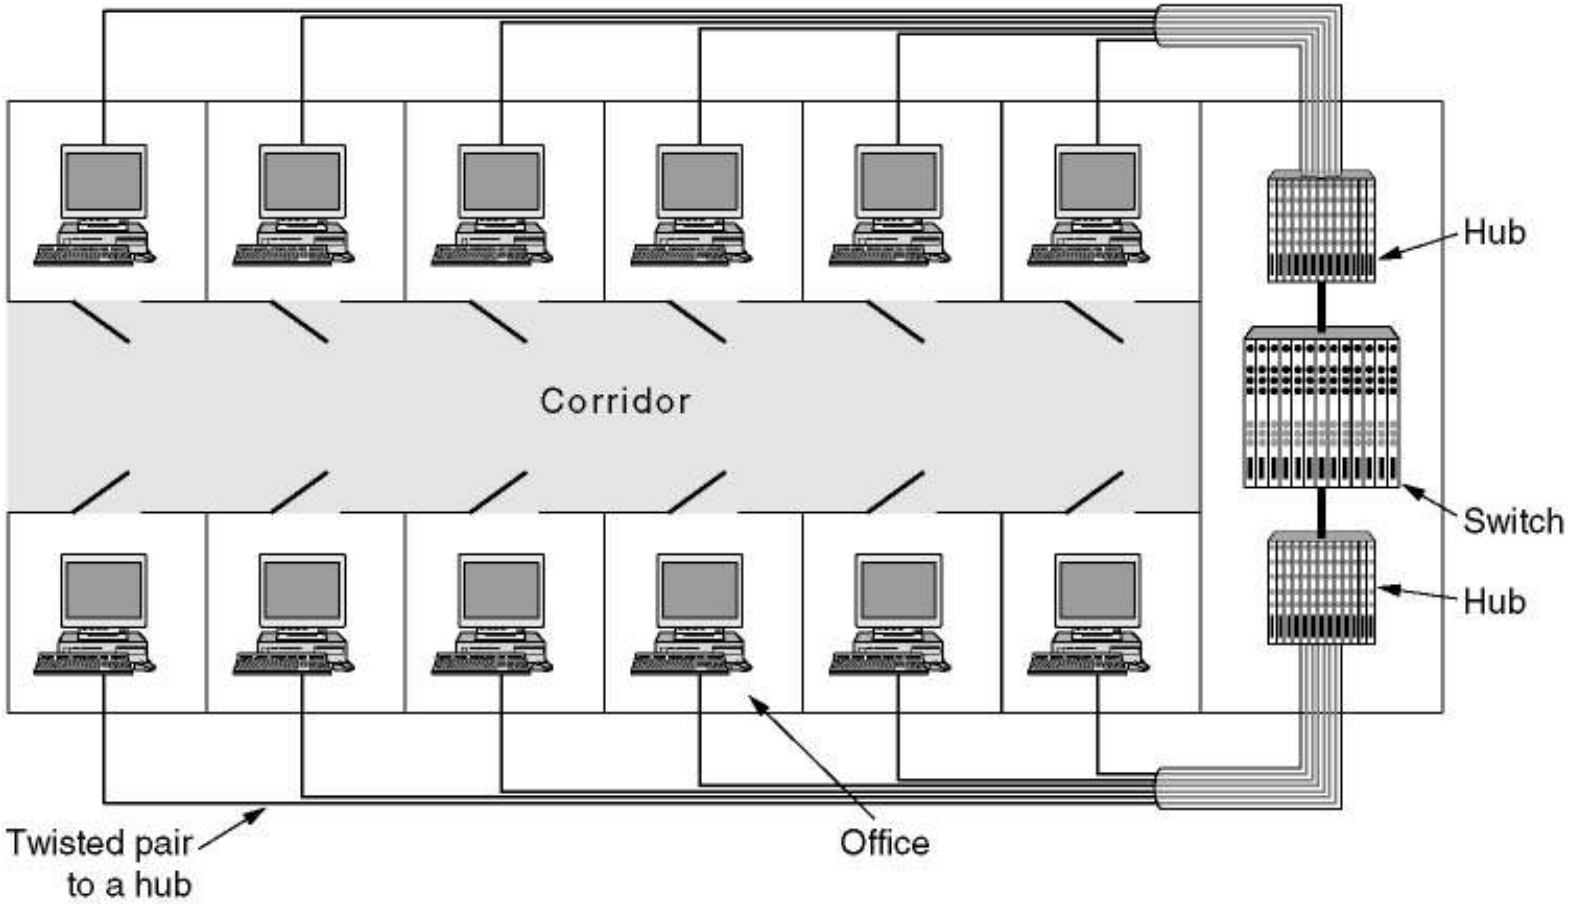

A building with centralized wiring using hubs and a switch.

Virtual LANs (3)

- In response to user requests for more flexibility, network vendors began working on a way to rewire buildings entirely in software.
- The resulting concept is called a VLAN (Virtual LAN) and has even been standardized by the 802 committee.
- VLANs are based on specially – designed VLAN – aware switches

Virtual LANs (4)

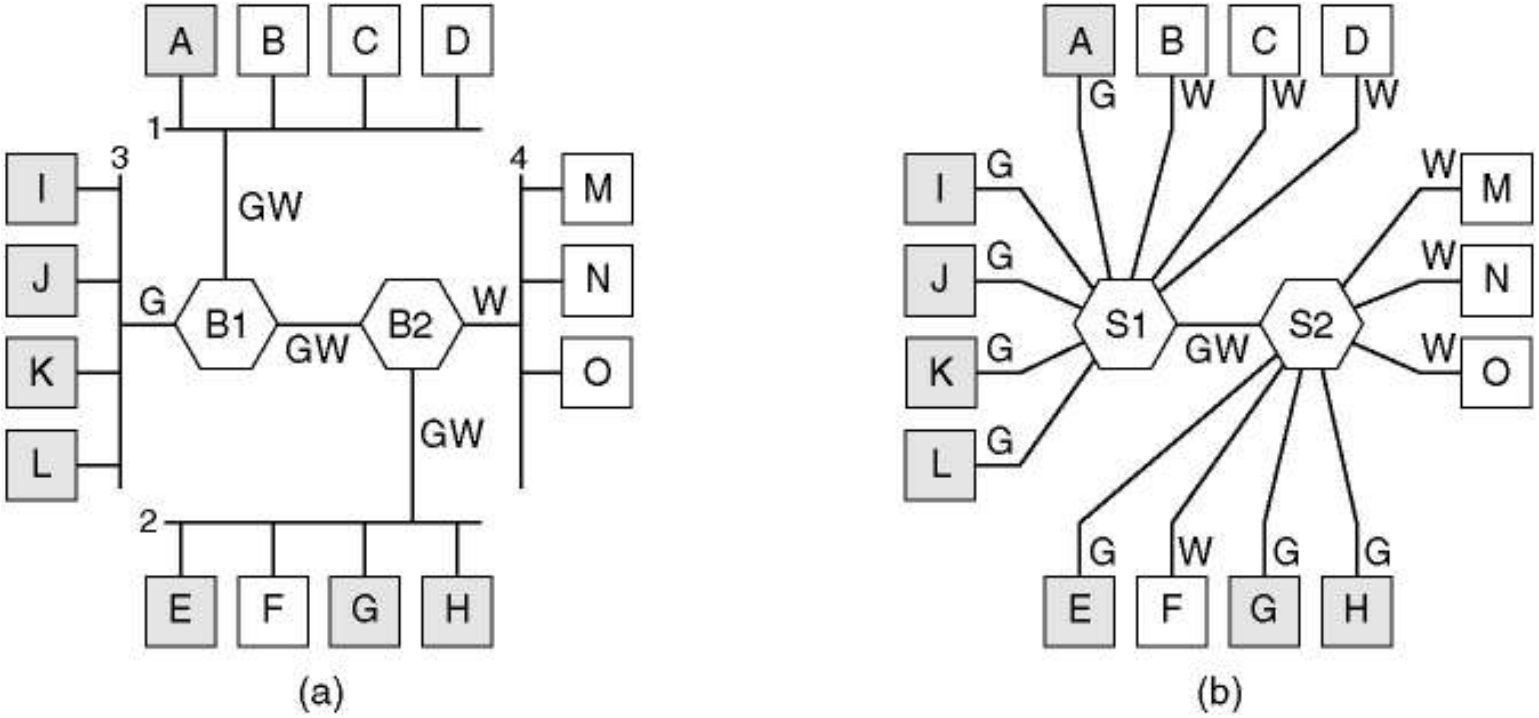

(a) Four physical LANs organized into two VLANs, gray and white, by two bridges. (b) The same 15 machines organized into two VLANs by switches.

The IEEE 802.1Q Standard (1)

- In VLAN the actually matters is the VLAN of the frame itself, not the VLAN of the sending machine.
- If there were some way to identify the VLAN in the frame header, then the need to inspect the payload would vanish.
- For a new LAN, such as 802.11 or 802.16, it would have been easy enough to just add a VLAN field in the header.

The IEEE 802.1Q Standard (2)

- In fact, the *Connection Identifier* field in 802.16 is somewhat similar in spirit to a VLAN identifier.
- But what to do about Ethernet, which is the dominant LAN, and does not have any spare fields lying around for the VLAN identifier?
- This problem was solved by changing the Ethernet header and the new format was published in IEEE Standard 802.1Q.

The IEEE 802.1Q Standard (3)

- The new format contains a VLAN tag; we will examine it shortly.
- During the transition process, many installations will have some legacy machines (typically classic or fast Ethernet) that are not VLAN aware and others (typically gigabit Ethernet) that are.

The IEEE 802.1Q Standard (4)

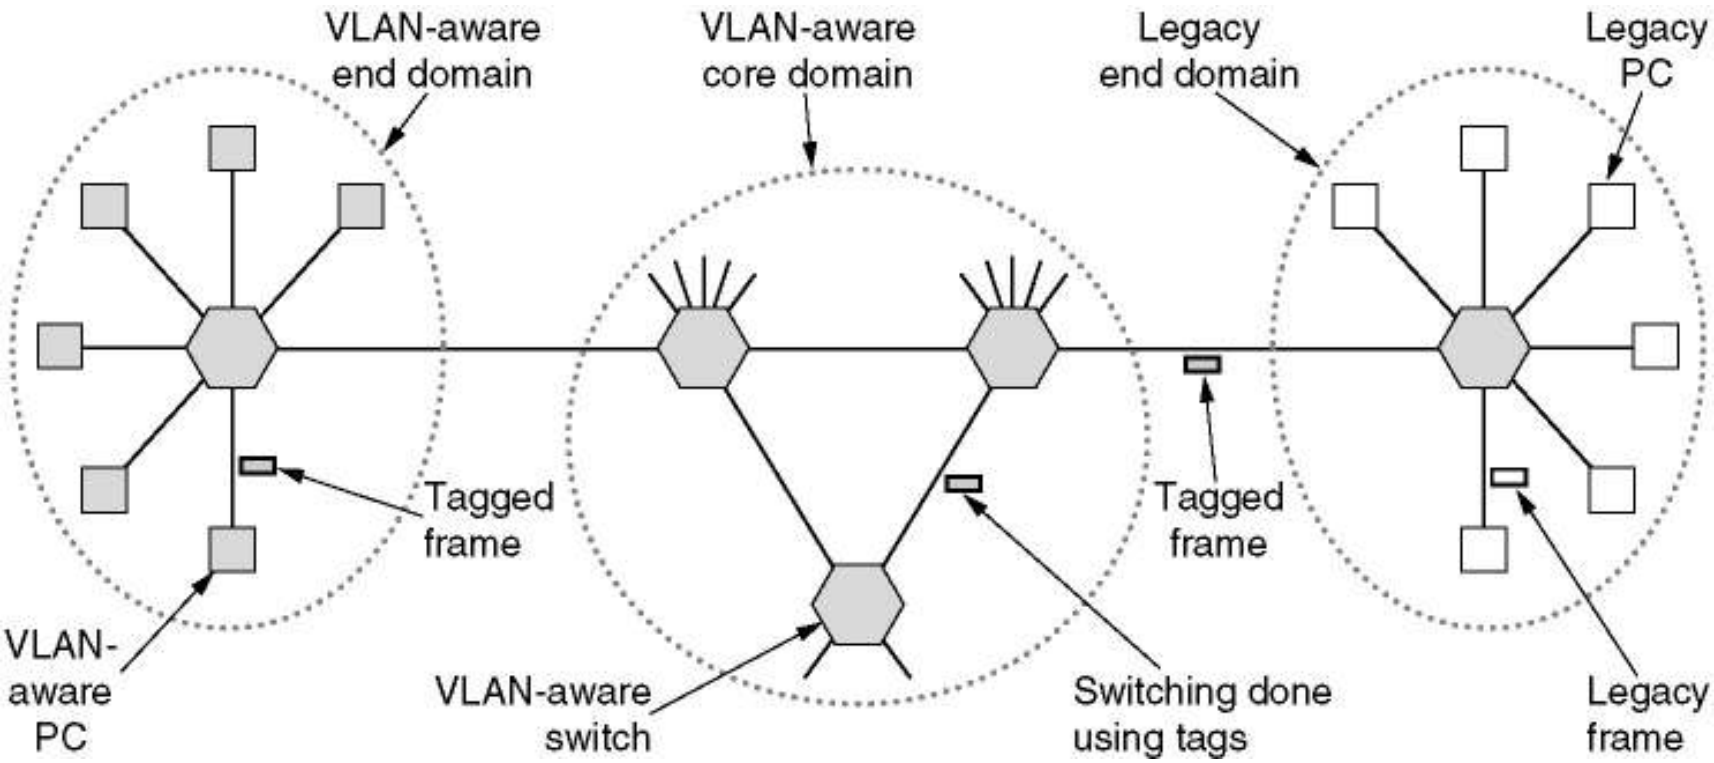

Transition from legacy Ethernet to VLAN-aware Ethernet. The shaded symbols are VLAN aware. The empty ones are not.

The IEEE 802.1Q Standard (5)

- When a legacy PC sends a frame to a VLAN-aware switch, the switch builds a new tagged frame based on its knowledge of the sender's VLAN (using the port, MAC address, or IP address).
- From that point on, it no longer matters that the sender was a legacy machine.
- Similarly, a switch that needs to deliver a tagged frame to a legacy machine has to reformat the frame in the legacy format before delivering it.

The IEEE 802.1Q Standard (6)

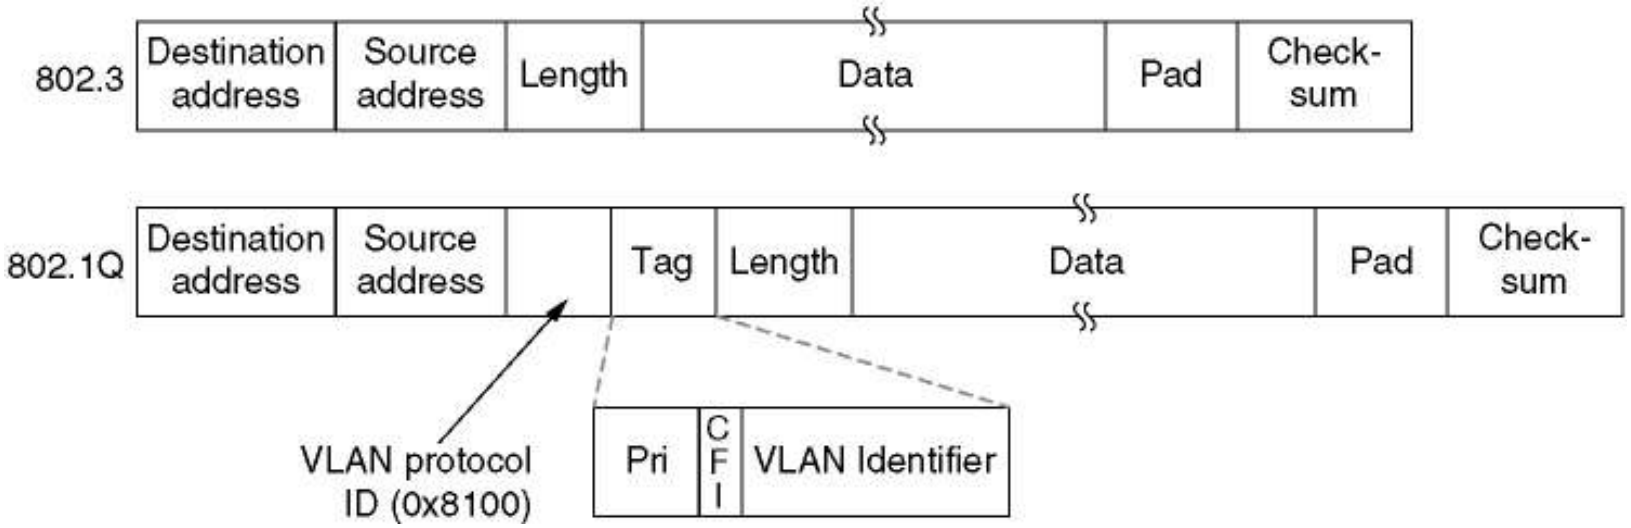

The 802.3 (legacy) and 802.1Q Ethernet frame formats.

The IEEE 802.1Q Standard (7)

- The only change is the addition of a pair of 2-byte fields.
- The first one is the *VLAN protocol ID* . It always has the value 0x8100. Since this number is greater than 1500, all Ethernet cards interpret it as a type rather than a length.
- What a legacy card does with such a frame is moot since such frames are not supposed to be sent to legacy cards.

The IEEE 802.1Q Standard (8)

- The second 2-byte field contains three subfields.
- The main one is the *VLAN identifier*, occupying the low-order 12 bits. This is what the whole thing is about – which VLAN does the frame belong to?
- The *3-bit Priority* field has nothing to do with VLANs at all. This field makes it possible to distinguish hard real-time traffic from soft real-time traffic .

The IEEE 802.1Q Standard (9)

- The last bit, *CFI (Canonical Format Indicator)* should have been called the CEI (Corporate Ego Indicator)
- It was originally intended to indicate little-endian MAC addresses versus big-endian MAC addresses, but that use got lost in other controversies.

Summary

Method	Description
FDM	Dedicate a frequency band to each station
WDM	A dynamic FDM scheme for fiber
TDM	Dedicate a time slot to each station
Pure ALOHA	Unsynchronized transmission at any instant
Slotted ALOHA	Random transmission in well-defined time slots
1-persistent CSMA	Standard carrier sense multiple access
Nonpersistent CSMA	Random delay when channel is sensed busy
P-persistent CSMA	CSMA, but with a probability of p of persisting
CSMA/CD	CSMA, but abort on detecting a collision
Bit map	Round robin scheduling using a bit map
Binary countdown	Highest numbered ready station goes next
Tree walk	Reduced contention by selective enabling
MACA, MACAW	Wireless LAN protocols
Ethernet	CSMA/CD with binary exponential backoff
FHSS	Frequency hopping spread spectrum
DSSS	Direct sequence spread spectrum
CSMA/CA	Carrier sense multiple access with collision avoidance

Channel allocation methods and systems for a common channel.