



Computer Networking

Local Area Networks

Prof. Andrzej Duda
duda@imag.fr

<http://duda.imag.fr>

1

The **data-link layer** is responsible for transferring packets across a link which is the communication channel connecting two adjacent hosts or routers.

Examples of link-layer protocols include Ethernet, wireless lans such as 802.11, and PPP.

LANs


Our goals:

- understand principles behind LANs:
 - sharing a broadcast channel: multiple access
 - link layer addressing
 - LAN interconnection
- instantiation and implementation of various LAN technologies

Overview:

- multiple access protocols
- example LANs:
 - Ethernet
 - 802.11
 - token ring
 - token bus
- link layer addressing
- LAN interconnection
 - hubs, bridges, switches

Characteristics



Metcalfe's Ethernet sketch

- Short distances (100 m - 1 km)
- High bit rate (10 Mb/s, 100 Mb/s, 1 Gb/s)
- Shared communication channel
- Used in a distributed environment
 - shared equipment, shared data

3


Today, Ethernet is by far the most prevalent LAN technology, and is likely to remain so for the foreseeable future. There are many reasons for Ethernet's success. First, Ethernet hardware (in particular, network interface cards) has become a commodity and is remarkably cheap. This low cost is also due to the fact that Ethernet's multiple access protocol, CSMA/CD, is completely decentralized, which has also contributed to a simple design. Ethernet is easy to install and manage than token LANs or ATM. Moreover, Ethernet was the first widely deployed high-speed LAN, therefore familiar to many network administrators reluctant to switch to new technologies. Finally, Ethernet is an evolving technology. In the past only 10 Mbps Ethernet was available, but currently so called fast Ethernet allows a nominal bandwidth of 100 Mbps and even 1000 Mbits (1 Gbps).

Data link layer in LANs

- Shared channel
 - multiplexing (TDM, FDM, or CDM)
 - fixed allocation: wasted bandwidth if no active sources
 - statistical multiplexing (multiple access)
 - suitable for bursty traffic - channel used at the full capacity
- Most of LANs
 - no retransmission (up to upper layers)
- WLANs
 - ACK of delivery

Multiple Access protocols

- single shared communication channel
- two or more simultaneous transmissions by nodes: interference
 - only one node can send successfully at a time
- *multiple access protocol*:
 - distributed algorithm that determines how stations share channel, i.e., determine when station can transmit
 - communication about channel sharing must use channel itself!
 - what to look for in multiple access protocols:
 - synchronous or asynchronous
 - information needed about other stations
 - robustness (e.g., to channel errors)
 - performance



5

In presence of a shared medium, it can happen that some nodes transmit at the same time and that frames collide or interfere. It is therefore necessary to find a protocol for sharing a broadcast medium. **Multiple access protocols** regulate nodes transmission onto the shared broadcast channel. Moreover, also the communication due to the coordination of the transmission must use the channel itself.

Multiple Access Protocols

Three broad classes:

- Random Access (Ethernet, 802.11)
 - allow collisions
 - "recover" from collisions
- Tokens - "Taking turns" (Token Ring, FDDI)
 - tightly coordinate shared access to avoid collisions
- Distributed Queue (DQDB)
 - use the channel in the arrival order


- Goal: efficient, fair, simple, decentralized

6

Multiple access protocols can be classified as belonging to one of three categories: **random access protocols**, token based, and distributed queue.

LAN technologies

- Data link layer:
 - services, multiple access
- LAN technologies
 - addressing
 - Ethernet, 802.11
 - repeaters, hubs, bridges, switches
 - virtual LANs




7

Multiple access protocols are extensively used in **local area networks (LANs)**. A LAN is a broadcast channel, which provides to its host access to the Internet through a router. The LAN is a single "link" between each user host and the router, where each node sends frames to each other over a broadcast channel; it therefore uses a link-layer protocol, part of which is a multiple access protocol. The transmission rate, R , of most LANs is very high (up to 1 Gbps).

However, despite the broadcast capability, in general a node in the LAN doesn't want to send a frame to *all* of the other LAN nodes but instead wants to send to some *particular* LAN node. Therefore, the nodes need LAN addresses (in reality this adapter has a LAN address) and the link-layer frame needs a field to contain such a destination address. In this manner, when a node receives a frame, it can determine whether the frame was intended for it or for some other node in the LAN. Note that, with the introduction of layer 2 addresses, broadcast must be explicitly addressed. Additionally, some LANs need to be interconnected together, and this can be obtained with different type of devices: repeaters, hubs, bridges, switches. This interconnection takes place at layer 2. Finally, several geographically distant LANs can be interconnected only at physical layer and "virtually" interconnected at layer 2 in a so called virtual LAN.

LAN Reference model




- LLC - Logical Link Control: IEEE 802.2 (ISO 8802.2)
- MAC - Medium Access Control
 - IEEE 802.3 (ISO 8802.3): CSMA/CD
 - IEEE 802.4 (ISO 8802.4): token bus
 - IEEE 802.5 (ISO 8802.5): token ring
 - IEEE 802.11: CSMA/CA

8

Today, Ethernet is by far the most prevalent LAN technology, and is likely to remain so for the foreseeable future. There are many reasons for Ethernet's success. First, Ethernet hardware (in particular, network interface cards) has become a commodity and is remarkably cheap. This low cost is also due to the fact that Ethernet's multiple access protocol, CSMA/CD, is completely decentralized, which has also contributed to a simple design. Ethernet is easy to install and manage than token LANs or ATM. Moreover, Ethernet was the first widely deployed high-speed LAN, therefore familiar to many network administrators reluctant to switch to new technologies. Finally, Ethernet is an evolving technology. In the past only 10 Mbps Ethernet was available, but currently so called fast Ethernet allows a nominal bandwidth of 100 Mbps and even 1000 Mbits (1 Gbps).

IEEE 802.3 - Ethernet



Variants

10: bit rate in Mb/s

BASE: modulation: BASE ou BROAD

5: maximal segment size in 100 m

Variant	Cable	Segment	Stations	Coverage
10 BASE 5	thick	500m	100	2500m
10 BASE 2	thin	200m	30	1000m
10 BASE T	pair	100m	1024	400m
10 BASE FX	fiber	2000m	1024	2000m

Segment

limited to 500 m

Two repeaters between any two stations at most

Transceiver cable


limited to 50 m

Distance between any two stations 2500 m

Round trip time of the signal between two stations

limited to 45 μ s

Coding



- Synchronous transmission
 - receiving station locks on 10 MHz - preamble
- Manchester coding

Random Access protocols

- When node has packet to send
 - transmit at full channel data rate R .
 - no *a priori* coordination among nodes
- two or more transmitting nodes -> "collision",
- random access protocol specifies:
 - how to detect collisions
 - how to recover from collisions (e.g., via delayed retransmissions)
- Examples of random access protocols:
 - ALOHA, slotted ALOHA
 - CSMA, CSMA/CD (Ethernet), CSMA/CA (802.11)

11

In a random access protocol, a transmitting node always transmits at the full rate of the channel, namely, R bps. When there is a collision, each node involved in the collision repeatedly retransmits its frame until the frame gets through without a collision. But when a node experiences a collision, it doesn't necessarily retransmit the frame right away. *Instead it waits a random delay before retransmitting the frame.* Each node involved in a collision chooses independent random delays. Because after a collision the random delays are independently chosen, it is possible that one of the nodes will pick a delay that is sufficiently less than the delays of the other colliding nodes and will therefore be able to sneak its frame into the channel without a collision.

ALOHA is the basis of all non-deterministic access methods. The ALOHA protocol requires acknowledgements and timers.

In this scheme a station wishing to transmit, does so at will. As a result, two or more frames may overlap in time, causing a collision. Collisions occur, and if a packet is lost, then sources have to retransmit; but they must stagger their attempts randomly, following some collision resolution algorithm, to avoid colliding again.

The maximum utilization can be proven to be 18%. This is assuming an ideal retransmission policy that avoids unnecessary repetitions of collisions.

With slotted ALOHA, time is divided into slots of equal size M that is the time necessary to transmit one frame and nodes start to transmit frames only at the beginnings of slots. Nodes need to be synchronized so that each node knows when the slots begin. With this expedient the maximum throughput is doubled.

CSMA improves on Aloha by requiring that stations listen before transmitting (compare to CB radio). Some collisions can be avoided, but not completely. This is because of propagation delays. Two or more stations may sense that the medium (= the channel) is free and start transmitting at time instants that are

CSMA/CD (Collision Detection)

- CSMA/CD (*Carrier Sense Multiple Access/ Collision Detection*)
 - carrier sensing, deferral if ongoing transmission
 - collisions *detected* within short time
 - colliding transmissions aborted, reducing channel wastage
 - persistent transmission
- collision detection:
 - easy in wired LANs: measure signal strengths, compare transmitted, received signals
 - difficult in wireless LANs: receiver shut off while transmitting

12

CSMA/CD is the protocol used by Ethernet. In addition to CSMA, it requires that a sending station monitors the channel and detects a collision. The benefit is that a collision is detected within a propagation round trip time. These mechanisms give CSMA/CD much better performance than slotted ALOHA in a LAN environment. In fact, if the maximum propagation delay between stations is very small, the efficiency of CSMA/CD can approach 100%. Collisions may still occur.

CSMA/CD algorithm


```
i = 1
while (i <= maxAttempts) do
    listen until channel is idle
    transmit and listen
    wait until (end of transmission) or
               (collision detected)
    if collision detected then
        stop transmitting, send jam bits (32 bits)
    else
        wait for interframe delay (9.6 μs)
        leave
    wait random time
    increment i
end do
```

3

CSMA/CD is the protocol used by Ethernet. In addition to CSMA, it requires that a sending station monitors the channel and detects a collision. The benefit is that a collision is detected within a propagation round trip time. These mechanisms give CSMA/CD much better performance than slotted ALOHA in a LAN environment. In fact, if the maximum propagation delay between stations is very small, the efficiency of CSMA/CD can approach 100%. Collisions may still occur.

CSMA / CD Collision

- A senses idle channel, starts transmitting
- shortly before T , B senses idle channel, starts transmitting




14

If the adapter in A senses that the channel is idle (that is, there is no signal energy from the channel entering the adapter), it starts to transmit the frame. However, due to the transmission time T , the adapter in B can sense that the channel is idle as well, even if A has started the transmission.

In this case there is a collision.

CSMA / CD Jam Signal

- B senses collision, continues to transmit the jam signal (32-bit)
- A senses collision, continues to transmit the jam signal



15

If the adapter detects signal energy from other adapters while transmitting, it stops transmitting its frame and instead transmits a jam signal. Jam signals are simply there to make sure the collision is long enough to be detected by the hardware.

Random retransmission interval

$r = \text{random}(0, 2^k - 1)$

$k = \text{min}(10, \text{AttemptNb})$

$$t_r = r \cdot 51.2 \mu\text{s}, \quad r \in [0, 2^k - 1]$$

- **slot time** = 51.2 μ s
- 1st collision, $r = 0, 1$
- 2nd collision, $r = 0, 1, 2, 3$
- 10th, $r = 0, 1, \dots, 1023$
- 15th, stop

16

After aborting (that is, transmitting the jam signal), the adapter enters an **exponential backoff** phase. Specifically, when transmitting a given frame, after experiencing the n th collision in a row for this frame, the adapter chooses a value for K at random from $\{0, 1, 2, \dots, 2^m - 1\}$ where $m = \text{min}(n, 10)$. The adapter then waits $K \cdot 512$ bit times and then returns to sense the channel.

Slot time

Round trip time limits the interval during which collisions may occur

slot

$45 \mu\text{s} + 3.2 \mu\text{s} < 51.2 \mu\text{s}$ - transmission of 512 bits

channel is *acquired* after 51.2 μ s

non-valid frames (results of collisions) < 512 bits \square minimal frame size (data field \geq 46 bytes)

unit of the retransmission interval

CSMA/CD performance

- Maximum utilization of Ethernet (approximation)

$$S \approx 1 / (1 + C \tau)$$

where $\tau = 2Db / L$,

D = propagation delay, b = bit rate,

L = frame size

C is a constant:

- C = 3.1 is a pessimistic value;
- C = 2.5 is an approximate value based on simulations

18

For a large network, $2Db$ is close to 60 bytes; for traffic with small frames ($L = 64$ bytes), the utilization is less than 30 %.

For large frames (1500 Bytes), it is around 90%.

Key for high utilization is:

bandwidth delay product \ll frame size (small τ !)

Frame format (Ethernet v.2)

preamble	dest	source	type	data	CRC
8 bytes	6 bytes	6 bytes	2 bytes	46 - 1500 bytes	4 bytes

- Preamble
 - synchronization : 10101010....0101011
- Addresses
 - unique, unicast and multicast (starts with the first bit 1)
 - broadcast: 11111...11111
- Type
 - upper layer protocol (IP, IPX, ARP, etc.)

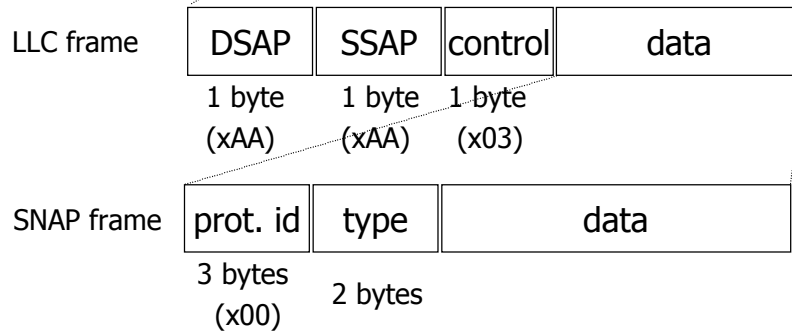

19

An Ethernet LAN can have a bus topology or a star topology. An Ethernet LAN can run over coaxial cable, twisted-pair copper wire, or fiber optics. Furthermore, Ethernet can transmit data at different rates, specifically, at 10 Mbps, 100 Mbps, and 1 Gbps.

The structure of an Ethernet frame is as follows:

• *Preamble (8 bytes)*. The Ethernet frame begins with an eight-byte preamble field. Each of the first seven bytes of the preamble has a value of 10101010; the last byte is 10101011. The first seven bytes of the preamble serve to "wake up" the receiving adapters and to synchronize their clocks to that of the sender's clock. Why should the clocks be out of synchronization? Keep in mind that adapter A aims to transmit the frame at 10 Mbps, 100 Mbps, or 1 Gbps, depending on the type of Ethernet LAN. However, because nothing is absolutely perfect, adapter A will not transmit the frame at exactly the target rate; there will always be some *drift* from the target rate, a drift which is not known *a priori* by the other adapters on the LAN. A receiving adapter can lock onto adapter A's clock by simply locking onto the bits in the first seven bytes of the preamble. The last two bits of the eighth byte of the preamble (the first two consecutive 1s) alert adapter B that the "important stuff" is about to come. When host B sees the two consecutive 1s, it knows that the next six bytes are the destination address. An adapter can tell when a frame ends by simply detecting absence of current.

Frame format (802.3)



- SNAP (Subnet Access Protocol) used in bridge management (any length of data: 0 - 1492)

20

• *Destination Address (6 bytes)*. This field contains the destination address. If a node receives a frame with an address *other* than its own MAC address, or the LAN broadcast address, it discards the frame. Otherwise, it passes the contents of the data field to the network layer.

• *Source Address (6 bytes)*. This field contains the LAN address of the source.


• *Data Field (46 to 1500 bytes)*. This field carries the IP datagram. The maximum transfer unit (MTU) of Ethernet is 1500 bytes. The minimum size of the data field is 46 bytes. This means that if the IP datagram is less than 46 bytes, the data field has to be "stuffed" to fill it out to 46 bytes. Data on Ethernet is transmitted least significant bit of first octet first (a bug dictated by Intel processors). Canonical representation thus inverts the order of bits inside a byte (the first bit of the address is the least significant bit of the first byte).

• *Type Field (2 bytes)*. The type field permits Ethernet to distinguish the network-layer protocols.

• *Cyclic Redundancy Check (CRC) (4 bytes)*. To detect whether any errors have been introduced into the frame.

Addressing

- MAC address: 48 bits = adapter identifier
- sender puts destination MAC address in the frame
- all stations read all frames; keep only if destination address matches
- all 1 address (**FF:FF:FF:FF:FF:FF**) = broadcast



21

- Ethernet addresses are known as MAC addresses. Every Ethernet interface has its own MAC address, which is in fact the serial number of the adapter, put by the manufacturer. MAC addresses are 48 bit-long. The 1st address bit is the individual/group bit, used to differentiate normal addresses from group addresses. The second bit indicates whether the address is globally administered (the normal case, burnt-in) or locally administered. Group addresses are always locally administered.
- When A sends a data frame to B, A creates a MAC frame with source addr = A, dest addr = B. The frame is sent on the network and recognized by the destination.
- Some systems like DEC networks require that MAC addresses be configured by software; those are so-called locally administered MAC addresses. This is avoided whenever possible in order to simplify network management.
- Data on Ethernet is transmitted least significant bit of first byte first (a bug dictated by Intel processors). Canonical representation thus inverts the order of bits inside a byte (the first bit of the address is the least significant bit of the first byte); examples of addresses:

01:00:5e:02:a6:cf (a group address)
 08:00:20:71:0d:d4 (a SUN machine)
 00:00:c0:3f:6c:a4 (a PC)
 00:00:0c:02:78:36 (a CISCO router)
 FF:FF:FF:FF:FF:FF the broadcast address

Addressing

- Data on Ethernet is transmitted least significant bit of first byte first (a bug dictated by Intel processors)
- Canonical representation thus inverts the order of bits inside a byte (the first bit of the address is the least significant bit of the first byte)
- examples of addresses:
 - `01:00:5e:02:a6:cf` (a group address)
 - `08:00:20:71:0d:d4` (a SUN machine)
 - `00:00:c0:3f:6c:a4` (a PC)
 - `00:00:0c:02:78:36` (a CISCO router)
 - `FF:FF:FF:FF:FF:FF` the broadcast address

22

48 bits : 24 bits delegated to a manufacturer and 24 bits of serial number

Interconnecting LANs

Why not just one big LAN?

- Limited amount of supportable traffic: on single LAN, all stations must share bandwidth
- limited distance
- large "collision domain" (can collide with many stations)
- processing broadcast frames

LAN evolution


- increase the bit rate: 10Mb/s, 100Mb/s, 1 Gb/s
- from hubs to switches

23

In principle, Internet could be implemented as one big LAN. However, there are several limitations to this solution: (1) the cables used for LANs are usually limited in length, therefore intercontinental distance could not be covered; (2) LANs use shared technologies, therefore the bandwidth is shared among all the station participating to the LAN; (3) statistically, if the number of stations increases, the number of collisions augments.

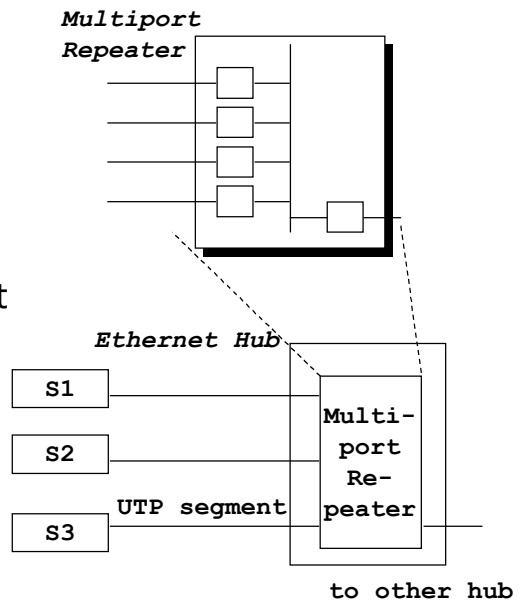
Repeaters

- Function of a simple, 2 port repeater:
 - repeat bits received on one port to other port
 - if collision sensed on one port, repeat random bits on other port
- One network with repeaters = **one collision domain**
- Repeaters perform only physical layer functions (bit repeaters)




From Repeaters to Hubs

- Multiport repeater (n ports), logically equivalent to:
 - n simple repeater
 - connected to one internal Ethernet segment
- Multi-port repeaters make it possible to use point-to-point segments (Ethernet in the box)
 - ease of management
 - fault isolation



10 BASE T Hubs




- Tree topology (star)
 - hub (*répéteur multipoint*)
 - max. 4 hubs

26

10BaseT and 100BaseT Ethernet are similar technologies. The first transmits at 10 Mbps and 100BaseT Ethernet transmits at 100 Mbps. 100BaseT is also commonly called "fast Ethernet". Both 10BaseT and 100BaseT Ethernet use a star based topology cabling. There is a central device called a **hub** (also sometimes called a concentrator.) Each adapter on each node has a direct, point-to-point connection to the hub. This connection consists of two pairs of twisted-pair copper wire, one for transmitting and the other for receiving. At each end of the connection there is a connector that resembles the RJ-45 connector used for ordinary telephones. The "T" in 10BaseT and 100BaseT stands for "twisted pair." For both 10BaseT and 100BaseT, the maximum length of the connection between an adapter and the hub is 100 meters; the maximum length between any two nodes is thus 200 meters. A hub is a repeater: when it receives a bit from an adapter, it sends the bit to all the other adapters. In this manner, each adapter can (1) sense the channel to determine if it is idle, and (2) detect a collision while it is transmitting. But hubs are popular because they also provide network management features. When a node has a problem the hub will detect the problem and internally disconnect the malfunctioning adapter.


10 BASE T



- Two pairs
 - emission
 - reception
- RJ-45 jack
- Hub - host
 - straight cable
- Hub - hub
 - inversed cable

10BaseT and 100BaseT

- 10/100 Mbps rate; latter called "fast ethernet"
- T stands for Twisted Pair
- Hub to which nodes are connected by twisted pair, thus "star topology"
- CSMA/CD supported by hubs



28

10BaseT and 100BaseT Ethernet are similar technologies. The first transmits at 10 Mbps and 100BaseT Ethernet transmits at 100 Mbps. 100BaseT is also commonly called "fast Ethernet". Both 10BaseT and 100BaseT Ethernet use a star based topology cabling. There is a central device called a **hub** (also sometimes called a concentrator.) Each adapter on each node has a direct, point-to-point connection to the hub. This connection consists of two pairs of twisted-pair copper wire, one for transmitting and the other for receiving. At each end of the connection there is a connector that resembles the RJ-45 connector used for ordinary telephones. The "T" in 10BaseT and 100BaseT stands for "twisted pair." For both 10BaseT and 100BaseT, the maximum length of the connection between an adapter and the hub is 100 meters; the maximum length between any two nodes is thus 200 meters. A hub is a repeater: when it receives a bit from an adapter, it sends the bit to all the other adapters. In this manner, each adapter can (1) sense the channel to determine if it is idle, and (2) detect a collision while it is transmitting. But hubs are popular because they also provide network management features. When a node has a problem the hub will detect the problem and internally disconnect the malfunctioning adapter.

Gigabit Ethernet

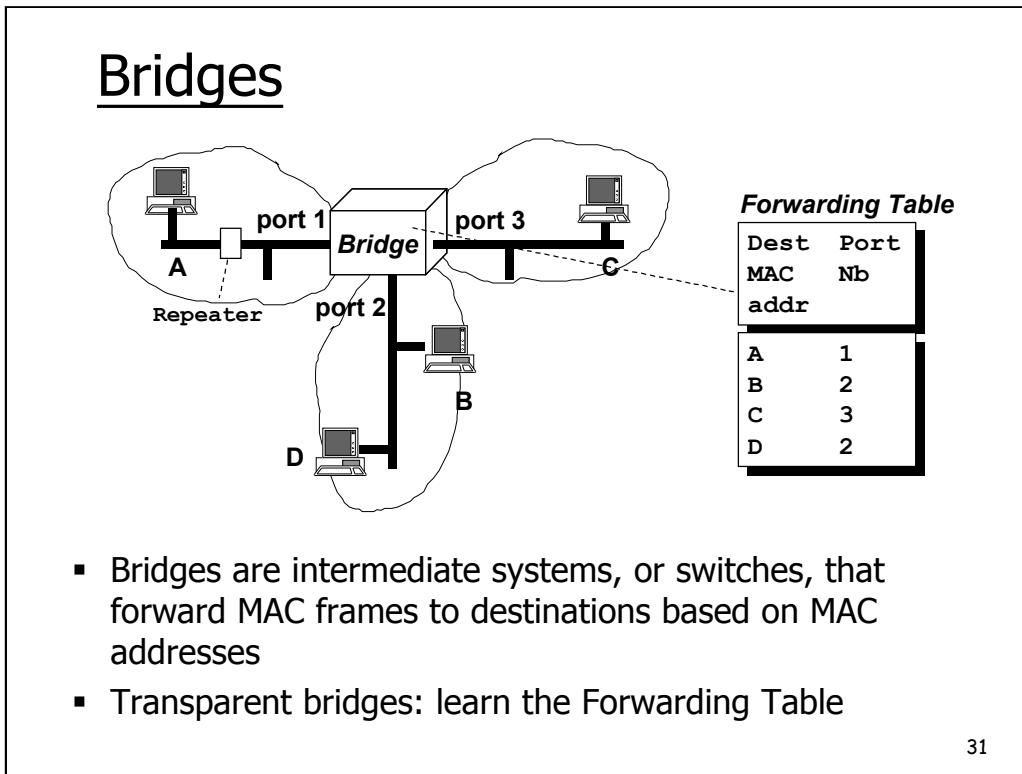
- use standard Ethernet frame format
- allows for point-to-point links and shared broadcast channels
- in shared mode, CSMA/CD is used; short distances between nodes to be efficient
- Full-Duplex at 1 Gbps for point-to-point links

29

Gigabit Ethernet is an extension to a raw data rate of 1,000 Mbps. Gigabit Ethernet is backward compatible with 10BaseT and 100BaseT technologies. It allows for point-to-point links as well as shared broadcast channels. Point-to-point links use switches whereas broadcast channels use hubs. Gbit Ethernet uses CSMA/CD for shared broadcast channels. In order to have acceptable efficiency, the maximum distance between nodes must be severely restricted. It allows for full-duplex operation at 1,000 Mbps in both directions for point-to-point channels.

Gigabit Ethernet

- 1000 BASE T
 - over twisted pair (25 m)
- 1000 BASE SX
 - short wavelength (850 nm) over multimode (500 m)
- 1000 BASE LX
 - long wavelength (1300 nm) over multimode (550 m) and single-mode fiber (10 km)
- 1000 BASE LH (Long Haul)
 - greater distance over 10 μ m single-mode (500 m)
- 1000 BASE ZX
 - extended wavelength (1550 nm) over 10 μ m single-mode (70 km)



A bridge is an intermediate system for the MAC layer. It receives MAC frames and forwards them further.

Bridges – interconnection at layer 2


- Link Layer devices: operate on Ethernet frames, examining frame header and selectively forwarding frame based on its destination
- Bridge isolates collision domains since it buffers frames
- When needs to forward a frame on a segment, bridge uses CSMA/CD to access the segment and transmit
- Can connect different type Ethernets, since it is a buffering device
- Two main types of bridges: transparent bridges and spanning tree bridges (guarantee no loops)

32

Bridges operate on Ethernet frames and thus are layer-2 devices. In fact, **bridges** are full-fledged packet switches that forward and filter frames using the LAN destination addresses. When a frame comes into a bridge interface, the bridge does not just copy the frame onto all of the other interfaces. Instead, the bridge examines the layer-2 destination address of the frame and attempts to forward the frame on the interface that leads to the destination. First, bridges permit isolates collision. Second, bridges can interconnect different LAN technologies, including 10 Mbps and 100 Mbps Ethernets. Third, there is no limit to how large a LAN can be when bridges are used to interconnect LAN segments; in theory, using bridges, it is possible to build a LAN that spans the entire globe.

Bridges vs. Routers

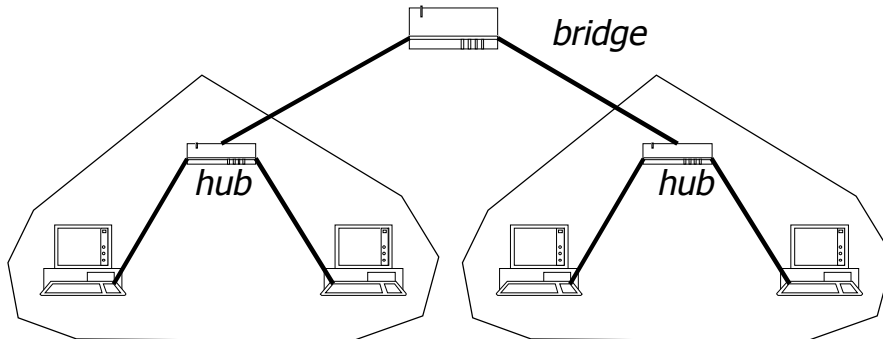
- both store-and-forward devices
 - routers: network layer devices (examine network layer headers)
 - bridges are Link Layer devices (look into MAC headers)
- routers are more complex
- bridges are plug-and-play



33

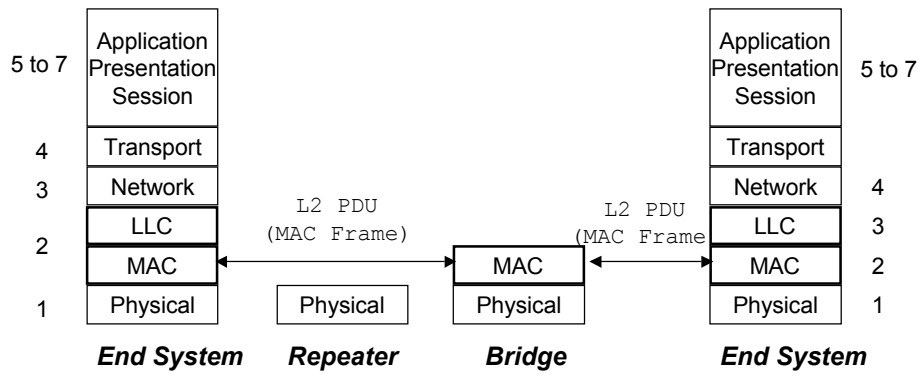
Routers are store-and-forward packet switches that forward packets using network-layer addresses. Although a bridge is also a store-and-forward packet switch, it is fundamentally different from a router in that it forwards packets using LAN addresses. Whereas a router is a layer 3 packet switch, a bridge is a layer-2 packet switch.

Collision domains



- Bridges separate collision domains
 - a bridged LAN maybe much larger than a repeated LAN
 - there may be several frames transmitted in parallel in a bridged LAN


Repeaters and Bridges in OSI Model



- Bridges are layer 2 intermediate systems
- Repeaters are in layer 1 intermediate systems
- Routers are layer 3 intermediate systems (IP routers)

Ethernet Switches – layer 2


- layer 2 (frame) forwarding, filtering using LAN addresses
- Switching: A-to-B and A'-to-B' simultaneously, no collisions
- large number of interfaces
- often: individual hosts, star-connected into switch
 - Ethernet, but no collisions!



36

Ethernet **switches** are in essence high-performance multi-interface bridges. As do bridges, they forward and filter frames using LAN destination addresses, and they automatically build forwarding tables using the source addresses in the traversing frames. The most important difference between a bridge and switch is that bridges usually have a small number of interfaces (that is, 2-4), whereas switches may have dozens of interfaces. A large number of interfaces generates a high aggregate forwarding rate through the switch fabric, therefore necessitating a high-performance design (especially for 100 Mbps and 1 Gbps interfaces). When a host has a direct connection to a switch (rather than a shared LAN connection), the host is said to have **dedicated access**.

Ethernet Switches (more)



Switching

- *Store-and-forward*
 - receive frame, check if valid, retransmit
 - 50 μ s delay for a 64 bytes frame
- *Cut through*
 - address read, retransmit
 - 20 μ s delay for a 64 bytes frame
 - transmission of non-valid frames

Full duplex Ethernet

- A shared medium Ethernet cable is half duplex
- Full duplex Ethernet = a point to point cable, used in both directions
 - no access method, no CSMA/CD
- 100 Mb/s and Gigabit Ethernet switches use full duplex links to avoid distance limitations and to guarantee bandwidth for stations
- Requires full duplex adapters at stations

Gigabit Ethernet

- 1000 BASE T
 - over twisted pair (25 m)
- 1000 BASE SX
 - short wavelength (850 nm) over multimode (500 m)
- 1000 BASE LX
 - long wavelength (1300 nm) over multimode (550 m) and single-mode fiber (10 km)
- 1000 BASE LH (Long Haul)
 - greater distance over 10 μ m single-mode (500 m)
- 1000 BASE ZX
 - extended wavelength (1550 nm) over 10 μ m single-mode (70 km)

Wireless LAN: 802.11b

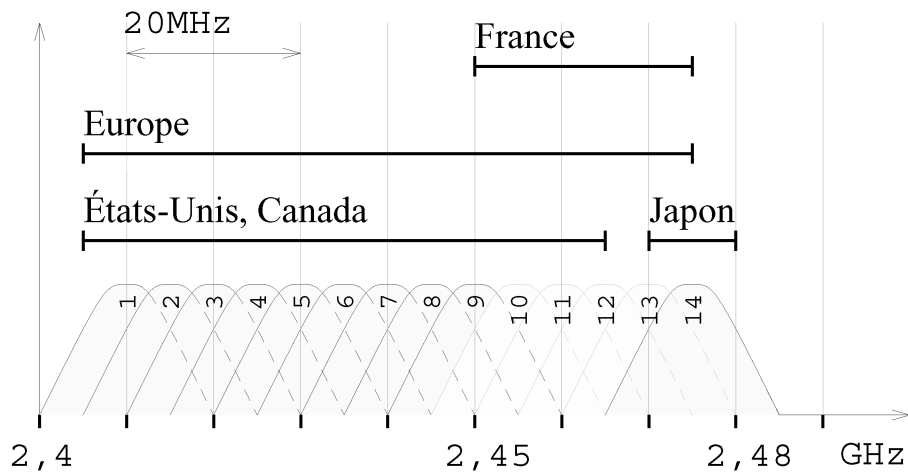
- 802.11b: wireless LAN
 - nominal bit rate of 11 Mb/s, degraded to 5.5, 2, 1 Mb/s
 - 6.5 Mb/s at application layer (file transfer)
 - shared radio channel, 2.4 GHz band, 13 channels (3 non overlapping of 22 MHz)
 - DSSS (*Direct Sequence Spread Spectrum*), 1 bit \square chipping sequence
 - coverage 50m, open air 100m
- MAC layer
 - DCF (Distributed Coordination Function)
 - CSMA/CA (Carrier Sense Multiple Access/Collision Avoidance), similar to Ethernet, no collision detection
 - PCF (Point Coordination Function)
 - polling, optional

802.11 - Physical layer


- 802.11b
 - frequency band of 2.4 GHz: [2,4 GHz ; 2,48 GHz]
 - nominal bit rate of 11 Mb/s
 - passes through concrete
- 802.11g
 - frequency band of 2.4 GHz
 - nominal bit rate of > 22 Mb/s
- 802.11a
 - frequency band of 5 GHz: [5,15 GHz ; 5,825 GHz]
 - nominal bit rate of 54 Mb/s
 - 6, 9, 12, 18, 24, 36, 48, 54 Mb/s, (6, 12, 24 Mb/s mandatory)
 - LOS - Line-of-Sight (no obstacles)

802.11 - Physical layer


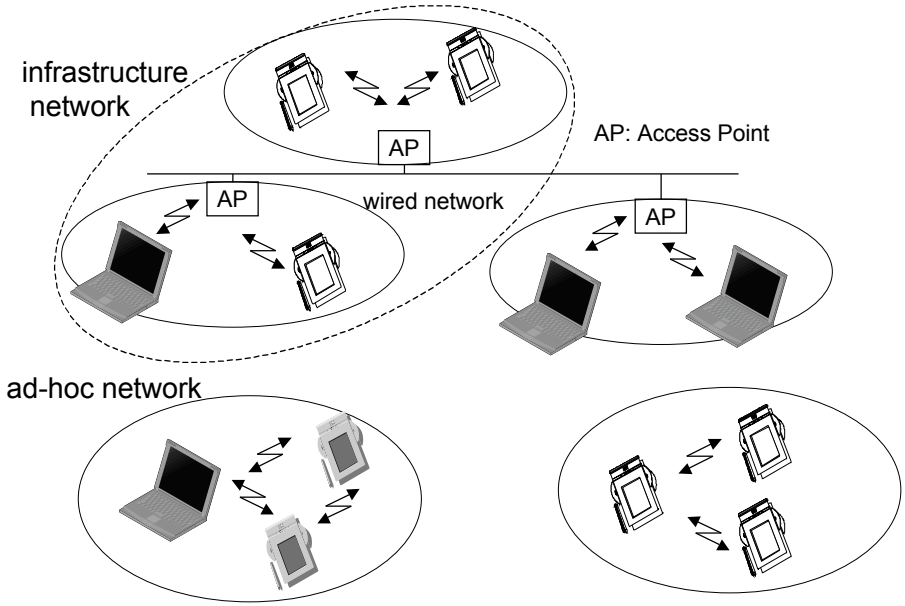
802.11b : canaux



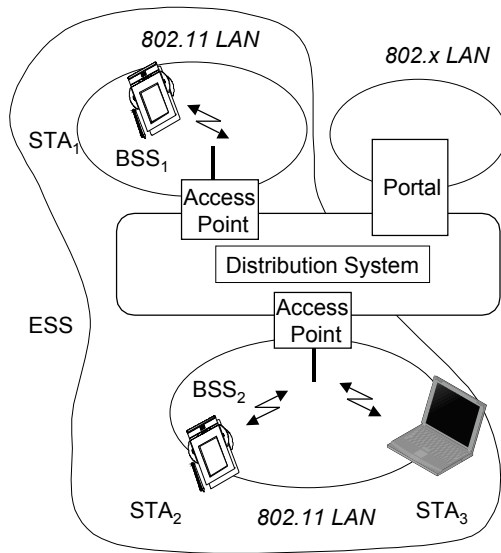
Channel selection



Infrastructure vs. ad-hoc




802.11 - infrastructure



- Station (STA)
 - terminal with access mechanisms to the wireless medium and radio contact to the access point
- Basic Service Set (BSS)
 - group of stations using the same radio frequency
- Access Point
 - station integrated into the wireless LAN and the distribution system
- Portal
 - bridge to other (wired) networks
- Distribution System
 - interconnection network to form one logical network


802.11

- Inter-frame spacing
 - SIFS (Short Inter Frame Spacing)
 - 10 μ s, for ACK, CTS, polling response
 - PIFS (PCF IFS)
 - for time-bounded service using PCF
 - DIFS (DCF IFS)
 - 50 μ s, for contention access



47

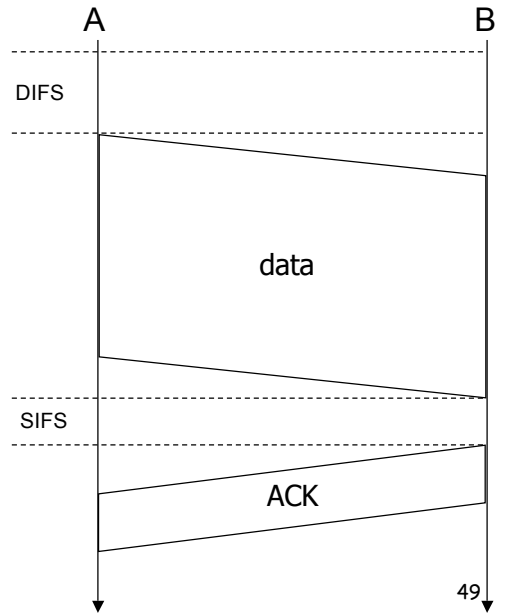
802.11 DCF - CSMA/CA



- Channel idle during DIFS, transmit frame
- If the medium is busy, wait for a free DIFS and a random back-off time (collision avoidance, multiple of slot-time)
- If another station uses the medium during the back-off time of the station, the back-off timer stops (fairness)

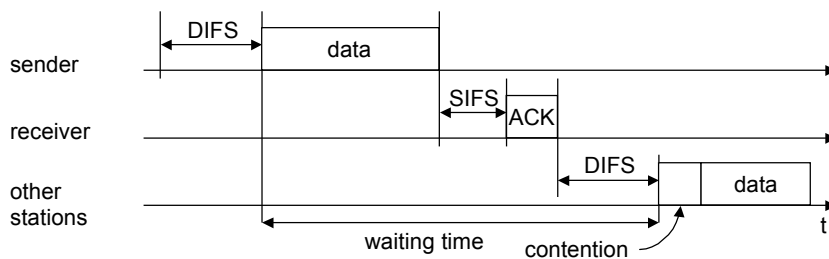
CSMA/CA (Collision Avoidance)

- Channel idle during DIFS, transmit frame
- Frame received correctly, wait SIFS, and send ACK




802.11 - CSMA/CA

- Sending unicast packets
 - station has to wait for DIFS before sending data
 - receivers acknowledge at once (after waiting for SIFS) if the packet was received correctly (CRC)
 - automatic retransmission of data packets in case of transmission errors



50

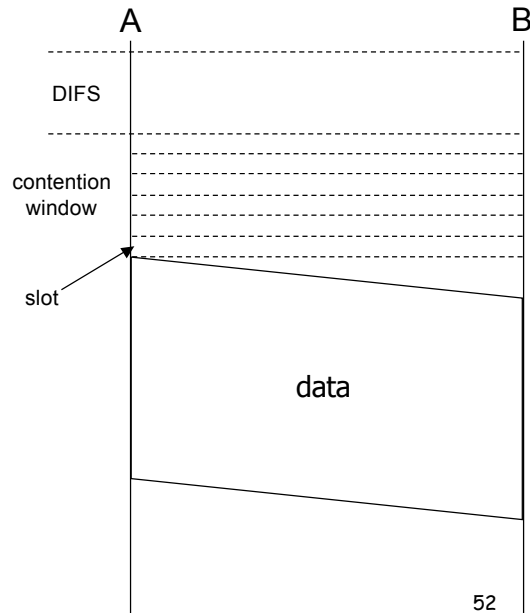
Contention



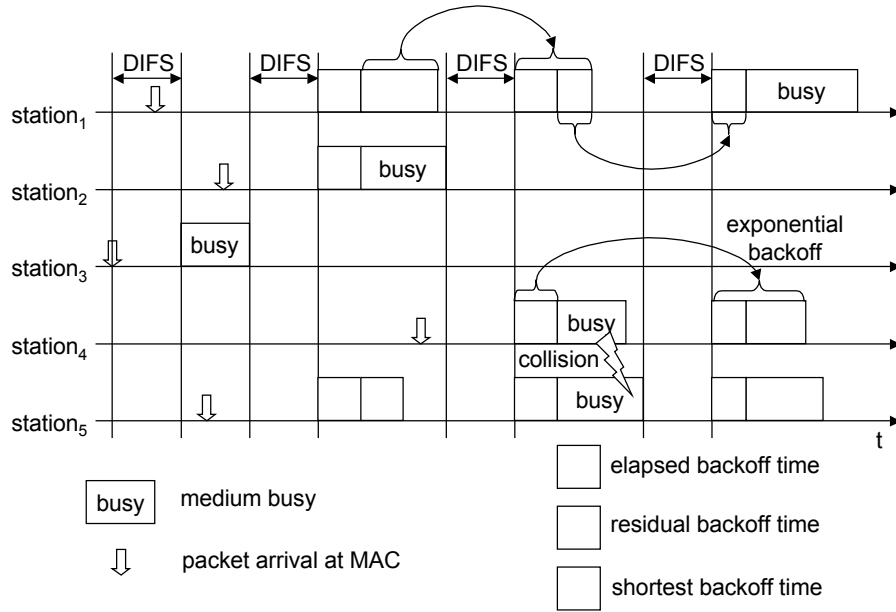
- Backoff time - random interval
 - Contention Window: uniform distribution $[0, CW] * \text{SLOT}$
 - CW: $CW_{\min} = 31, CW_{\max} = 1023$
 - $\text{SLOT} = 20 \mu\text{s}$
- $T(N)$ should also include time wasted in collisions

CSMA/CA (Collision Avoidance)

- If channel busy, defer. Then, if idle during DIFS, wait random interval (multiple of the slot) and transmit
- If channel busy, wait again until medium idle for at least DIFS
- Contention window doubles with each collision - exponential back-off

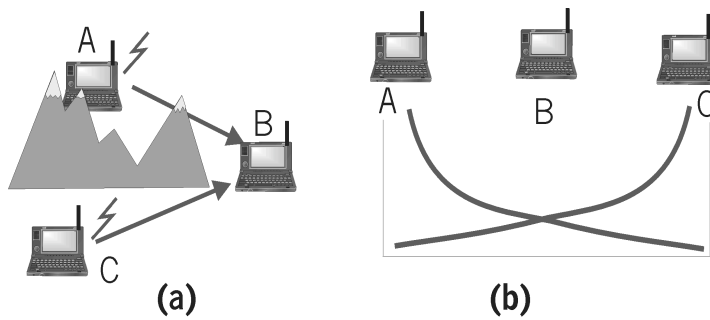


802.11 - contention



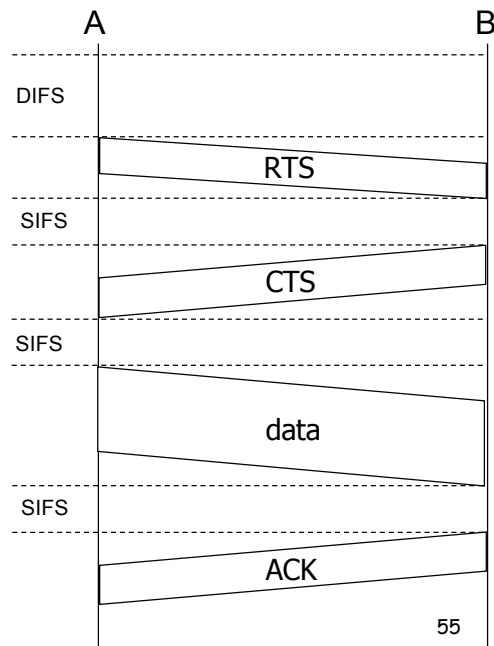
Hidden Terminal effect

- Hidden terminals: A and B cannot hear each other because of obstacles or signal attenuation; so, their packets collide at B

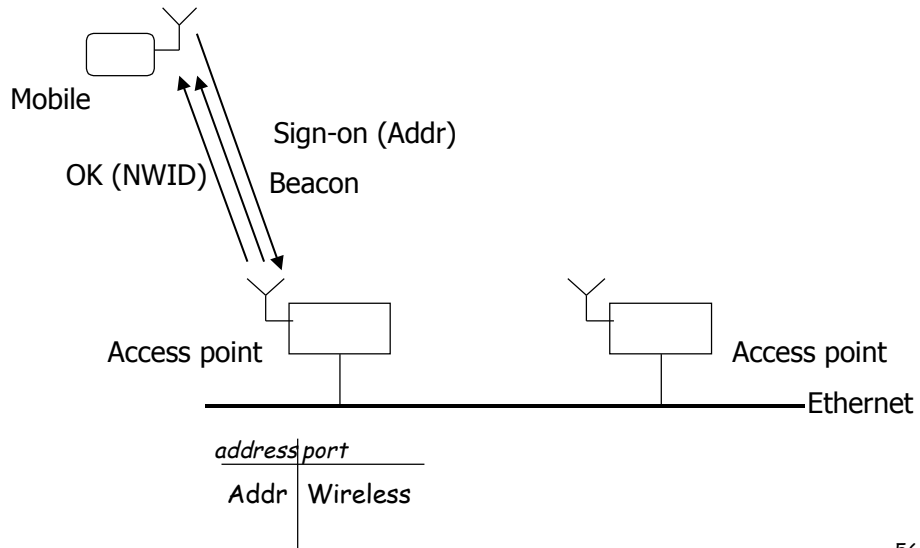


RTS/CTS Extension

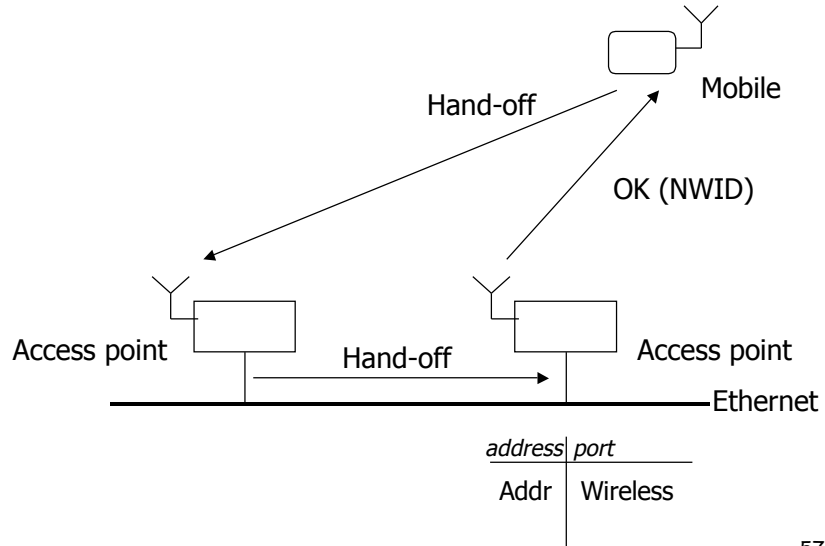
- CTS (*Clear To Send*) "freezes" stations within range of receiver (hidden from transmitter); this prevents collisions by hidden station during data transfer
- RTS (*Request To Send*) and CTS are very short: collisions are very unlikely (the end result is similar to Collision Detection)



Register to Access Point



Hand-off



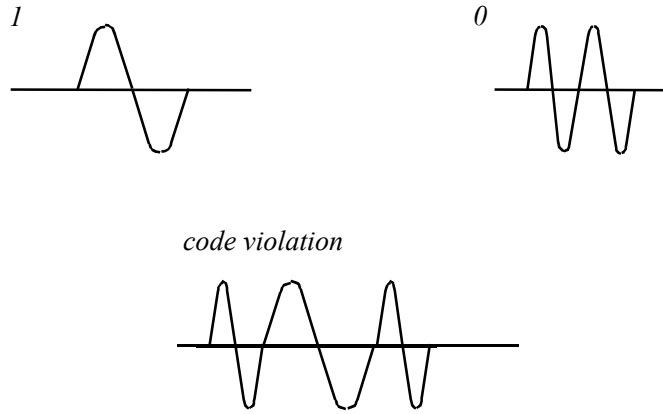
Bluetooth

- Replaces cables
 - short range (10m), low power, cheap
 - 2.4 GHz band
 - FHSS (*Frequency Hopping Spread Spectrum*)
 - piconet
 - all devices share the same hopping sequence
 - one master, seven slaves
 - bit rate: around 1 Mb/s
 - symmetric connections - 432.6 kb/s
 - asymmetric - 721 kb/s, 57.6 Kb/s
 - access method: polling, reservation


IEEE 802.4

- Token Bus
 - industrial LAN
- Physical layer
 - modulation (*broadband*)
 - coaxial cable 75 Ω
 - 1, 5, 10 Mb/s bit rate
- Access method
 - token on a virtual ring

Physical layer



Topology

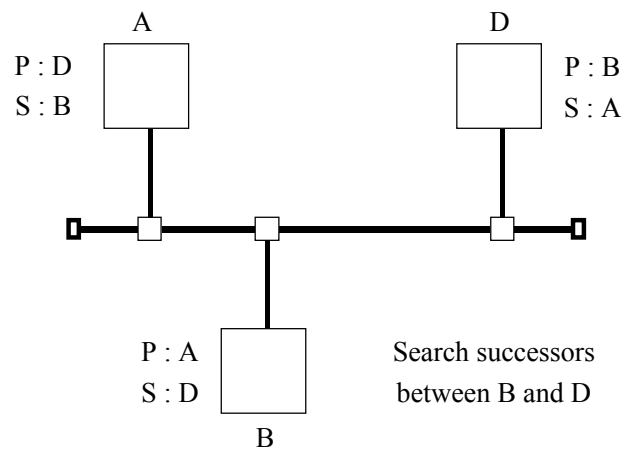


- Physical bus, virtual ring

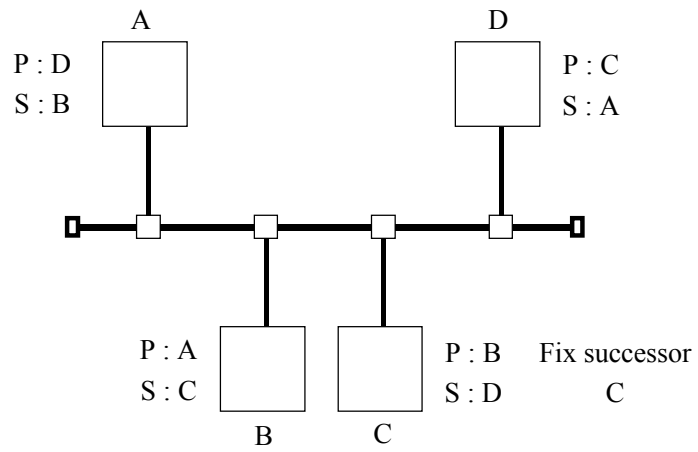
Access method

- Token
 - station can send one or several frames during the token holding interval
 - several priorities per station
- Virtual ring
 - two addresses: Successor, Predecessor
 - token holder passes it to its successor
 - ring maintenance:
 - each N tours, invite to enter


Adding a station



Adding a station



Departure of a station



Frame format

preamble	start	FC	dest	source	data	CRC	end
----------	-------	----	------	--------	------	-----	-----

≥ 1 bytes 1 byte 1 byte 2-6 bytes 2-6 bytes 0 - 8191 bytes 4 bytes 1 byte

- Preamble
 - synchronization
- Start and End
 - frame delimitation: NN0NN000, N - code violation
- FC - *Frame Control*
 - type of a frame: Token, Search Successor, Fix Successor


IEEE 802.5

- Token Ring
- Physical layer
 - differential Manchester coding
 - bits: H-L, L-H
 - violation: H-H, L-L
 - bit rate 4, 16 Mb/s
- Access method
 - token on a physical ring

Topology


- Physical ring
 - repeater
 - 1 bit shift register, on the fly modification
- Twisted pair cabling
 - star topology - wiring concentrator MAU (*Multistation Access Unit*), max. 8 stations
 - one pair - reception; one pair - transmission
- Coverage
 - station - MAU: 300 m, if one MAU; 100 m, if several MAU
 - MAU - MAU: 200 m

Ring




Repeater

- Listen
 - address/token recognition
 - copy/repeat
 - modify one bit (token hold)



- Transmission
 - buffer insertion
 - remove frame



Access method

- Token
 - token holding time limited to 10 ms
 - variants
 - 4 Mb/s: transmitting station generates token after removing the frame
 - 16 Mb/s: transmitting station generates token after the end of the frame (*daisy chain*)

Access method

- Priorities
 - token with different priorities (0 - 7)
 - priority reservation
 - a station can request generation of a token with a given priority
 - global priorities (vs. local priorities in 802.4)

Maintenance

- Monitoring station
 - elected at power up based on the address
 - every station may become monitor
 - initialize the ring
 - inserts a register of 24 bits (3 bytes) - token frame
 - monitor the ring:
 - presence of the token
 - absence of multiple tokens
 - purge if a frame is not removed

Problems

- Lost token
 - no token during an interval, purge the ring and regenerate the token
 - abandoned frames
 - monitoring station sets bit M in each frame
 - if frame received with M set, it is an abandoned frame
 - purge and regenerate the token

Frame format

start	AC	FC	dest	source	data	CRC	end	FS
1 byte	1 byte	1 byte	2-6 bytes	2-6 bytes	□ variable	4 bytes	1 byte	1 byte

- Start
 - frame delimitation - code violation
- AC - Access Control
 - token (1 bit)
 - priority (3 bits)
 - priority reservation (3 bits)
 - bit M - monitor (1 bit)

Frame format

- FC - *Frame Control* - type of frame
 - Claim Token (station wants to become monitor)
 - Purge (initialize the ring)
 - Monitor Present (if no such a frame, a station will try to become a monitor station)
- Data
 - token holding time: 10 ms
 - 4 Mb/s - 4464 bytes
 - 16 Mb/s - 17914 bytes

Frame format

- CRC
 - on FC ... data
- End
 - code violation
- FS - *Frame Status*
 - bit C: frame accepted
 - bit A: address recognized


FDDI (*Fiber Distributed Data Interface*)

- Dual fiber ring
 - multi-mode fiber
 - up to 500 stations
 - 100 km per ring (MAN - Metropolitan Area Network)
- Coding
 - 125 MHz clock, 100 Mb/s bit rate
 - 4B5B coding
 - 4 bits coded as 5 binary symbols
 - some symbols used for delimitation
 - NRZI signal

Access method

- Token ring, similar to 802.5
 - *daisy chain*
- Frame format similar to 802.5, 4352 bytes of data
- FDDI-II
 - synchronous traffic
 - monitoring station transmits a special frame every 125 μ s
 - up to 96 PCM voice channels

802.6 - DQDB (Distributed Queue Dual Bus)



- Dual bus
 - 160 km at 44 Mb/s (T3), 155 Mb/s

Access method

- Controller
 - generates a train of 53 bytes cells
- Cell format
 - addresses, *Request* bit, *Busy* bit,
 - 44 bytes of data

Access method


- **Distributed queue of transmission requests**
 - before transmit, set *Request* bit in a cell on the opposite bus
 - upper stations learn the request and leave one empty cell per request
 - set *Busy* bit in the first empty cell and insert data
- **Advantages**
 - no overhead, good throughput
- **Drawback**
 - not symmetric topology

LLC (*Logical Link Control*)

- IEEE 802.2
 - used in some LAN protocols (SNAP)
- HDLC family (PPP)
- Three types of services
 - 1: datagram
 - 2: connected mode (similar to X.25 LAPB)
 - 3: acknowledged datagram

VLAN - Virtual LAN

- Keep the advantages of Layer 2 interconnection
 - auto-configuration (addresses, topology - Spanning Tree)
 - performance of switching
- Enhance with functionalities of Layer 3
 - extensibility
 - spanning large distances
 - traffic filtering
- Limit broadcast domains
- Security
 - separate subnetworks




84

A **Virtual LAN** is a subset of stations physically connected in a LAN that are logically connected. The procedure of logically connecting a group of stations can be seen as a colouring procedure that is managed by a manager generally implemented in a switch.

Virtual LANs

- No traffic between different VLANs
- VLANs build on bridges or switches

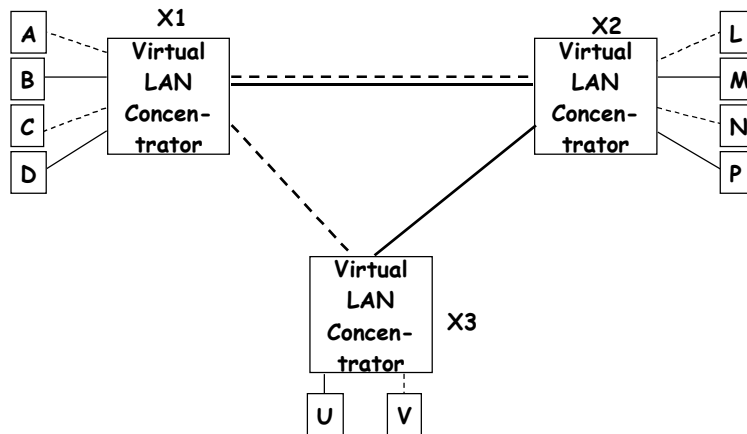


VLANs

- How to define which port belongs to a VLAN?
 - per port
 - simple, secure, not flexible for moving hosts (one host per port)
 - per MAC address
 - several hosts per port, flexible for moving hosts, not secure, difficult to manage, problems with protocols Layer 3 (should be coupled with dynamic address negotiation - DHCP)
 - per Layer 3 protocol
 - allows to limit frame broadcast (VLAN1: IP, VLAN2: IPX)
 - per Layer 3 address
 - one VLAN per IP subnetwork
 - flexible for moving hosts
 - may be less efficient (requires inspecting packets)

Remote VLANs

- works at layer 2
- uses an interconnection network (ATM) or a proprietary protocol



87

The picture shows two virtual LANs: (ACLNV) and (BDMPU). For each of the virtual LANs, there exists one or more collision domains per concentrator, plus one per inter-concentrator link. The concentrators perform bridging between the different collision domains of the *same* virtual LAN.

Between X1 and X2, the two virtual LANs use the same physical link. The advantage is that physical location becomes independent of LANs. For example, all servers and routers can be concentrated in the same rooms (ex: U and V). There is no communication between the different virtual LANs at layer 2.

Summary

- Original Ethernet is a shared medium: one collision domain per LAN
- Bridges are connectionless intermediate systems that interconnect LANs
- Using bridging, we can have several collision domains per LAN
- Ethernet switches use bridging
- State of the art
 - switched 100 Mb/s Ethernet to the host
 - 1 Gb Ethernet between switches
- Wireless LANs become increasingly popular
 - WiFi, Bluetooth

