Ethernet

Outline

- Multiple Access and Ethernet Intro
- Ethernet Framing
- CSMA/CD protocol
- Exponential backoff

Shared Access Networks are Different

- Shared Access Networks assume multiple nodes on the same physical link
  - Bus, ring and wireless structures
  - Transmission sent by one node is received by all others
  - No intermediate switches
- Need methods for moderating access (MAC protocols)
  - Fairness
  - Performance
  - How can this be done?

Multiple Access Methods

- Fixed assignment
  - Partition channel so each node gets a slice of the bandwidth
  - Essentially circuit switching – thus inefficient
  - Examples: TDMA, FDMA, CDMA (all used in wireless/cellular environments)
- Contention-based
  - Nodes contends equally for bandwidth and recover from collisions
  - Examples: Aloha, Ethernet
- Token-based or reservation-based
  - Take turns using the channel
  - Examples: Token ring

Coordinated/Fixed sharing approaches

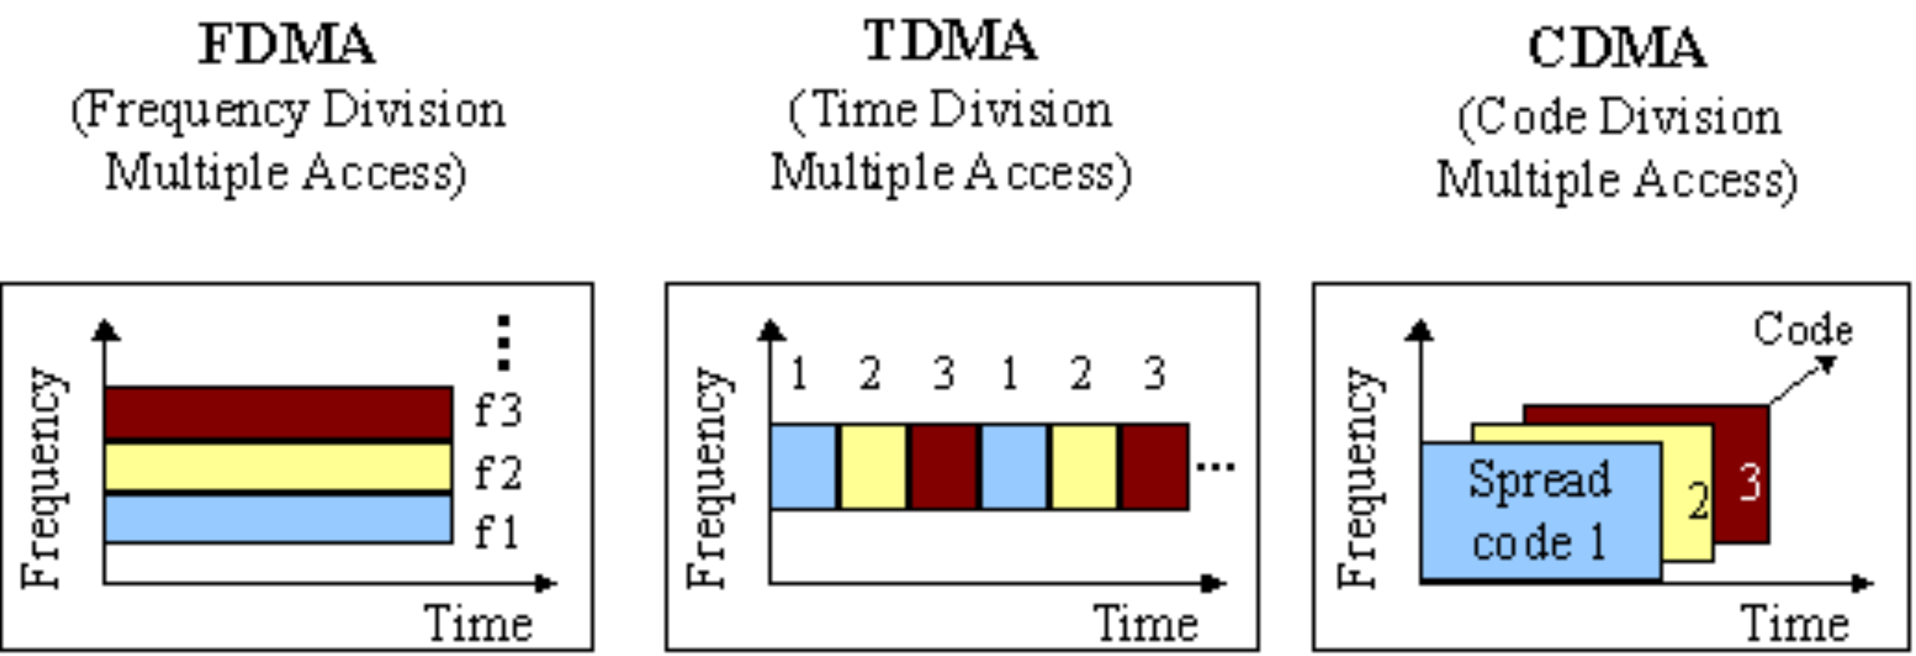

TDMA

- Examples
  - GSM
  - IS136
  - iDen
  - DECT
  - Satellite communications

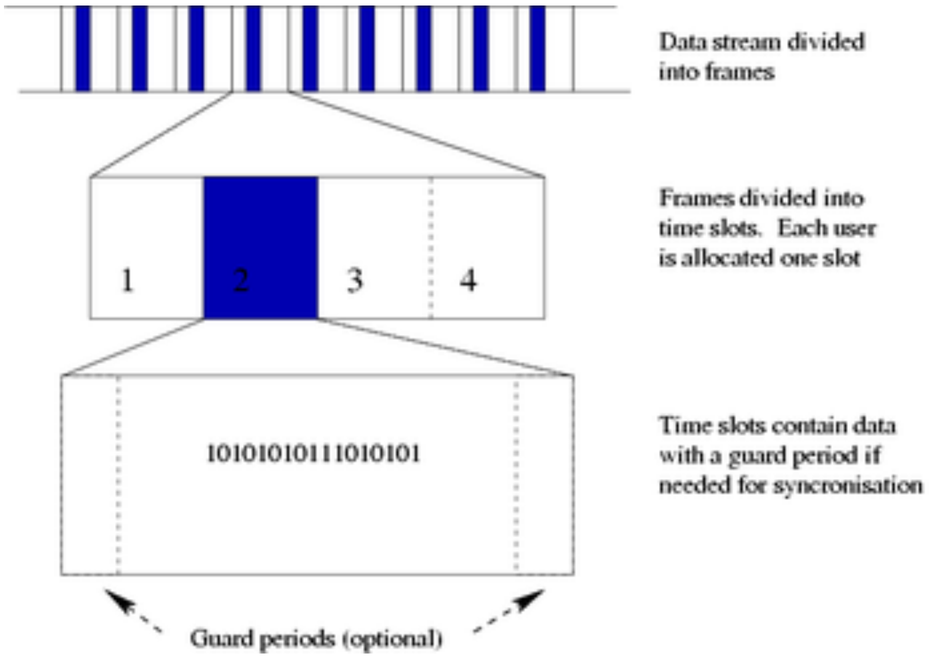

TDMA

- Advantages
  - Active only during own time slot, can do anything else in other time slots
    - Measure the channel, search for other transmitters in different frequencies, etc.
    - Facilities inter-frequency handoffs efficiently (imagine a macrocell and a microcell operating in the same region with different frequencies)
- Disadvantages
  - Need guard slots to protect between two transmitters (to avoid very tight synchronization)
  - This wastes capacity of the channel

A Quick Word about Token Ring

- Developed by IBM in early 80's as a new LAN architecture
  - Consists of nodes connected into a ring (typically via concentrators)
  - Special message called a token is passed around the ring
    - When nodes gets the token it can transmit for a limited time
    - Every node gets an equal opportunity to send
  - IEEE 802.5 standard for Token Ring
- Designed for predictability, fairness and reliability
  - Originally designed to run at either 4Mbps and 16Mbps
- Still used and sold but beaten out by Ethernet

Our Focus is Ethernet

- History
  - Developed by Bob Metcalfe and others at Xerox PARC in mid-1970s
  - Roots in Aloha packet-radio network
  - Standardized by Xerox, DEC, and Intel in 1978
  - LAN standards define MAC and physical layer connectivity
    - IEEE 802.3 (CSMA/CD - Ethernet) standard – originally 2Mbps
    - IEEE 802.3u standard for 100Mbps Ethernet
    - IEEE 802.3z standard for 1,000Mbps Ethernet
- CSMA/CD: Ethernet's Media Access Control (MAC) policy
  - CS = carrier sense
    - Send only if medium is idle
  - MA = multiple access
  - CD = collision detection
    - Stop sending immediately if collision is detected

Ethernet Standard Defines Physical Layer

- 802.3 standard defines both MAC *and* physical layer details

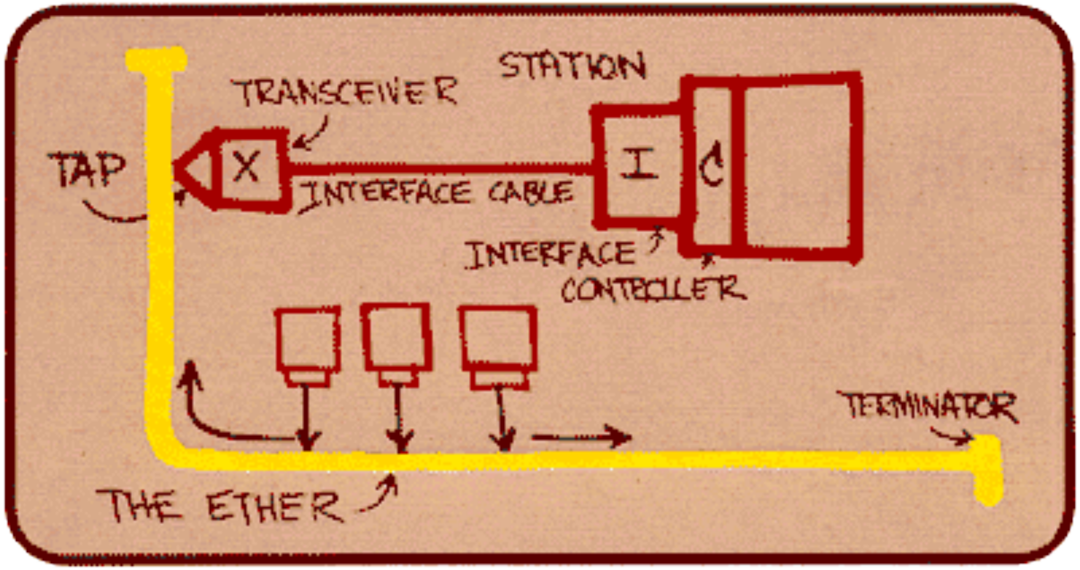

Metcalfe's original Ethernet Sketch

Ethernet Technologies: 10Base2

- 10: 10Mbps; 2: under 185 (~200) meters cable length
- Thin coaxial cable in a bus topology

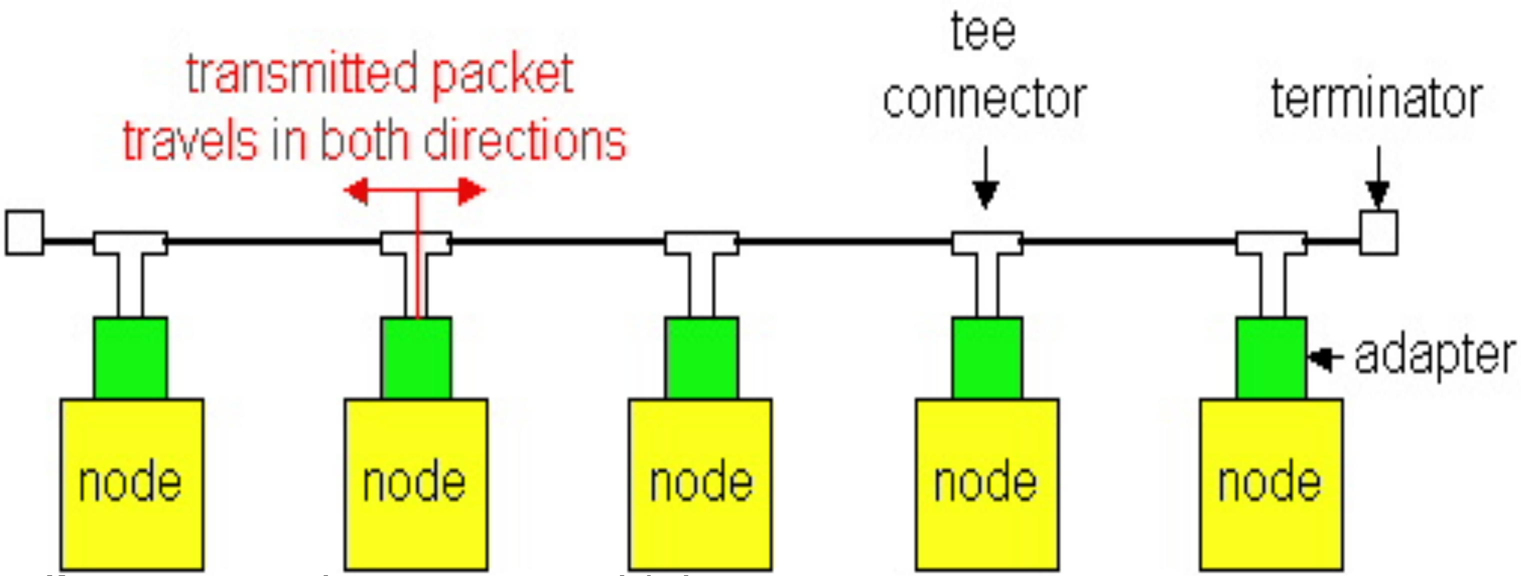

- Repeaters used to connect multiple segments
  - Repeater repeats bits it hears on one interface to its other interfaces: physical layer device only!

10BaseT and 100BaseT

- 10/100 Mbps rate
- T stands for Twisted Pair
- Hub(s) connected by twisted pair facilitate "star topology"
  - Distance of any node to hub must be < 100M

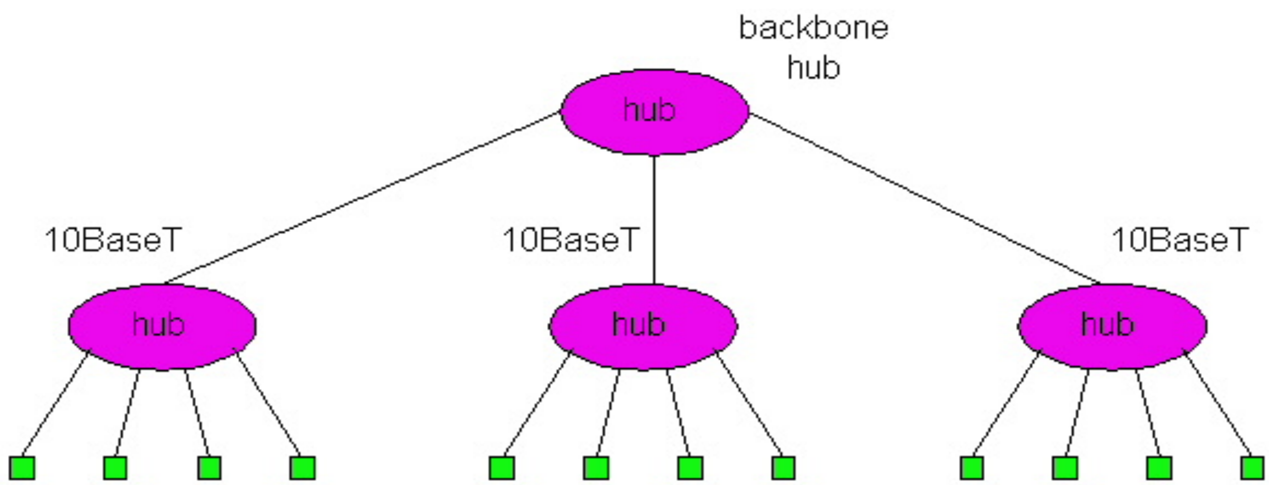

Physical Layer Configurations for 802.3

- Physical layer configurations are specified in three parts
- Data rate (10, 100, 1,000)
  - 10, 100, 1,000Mbps
- Signaling method (base, broad)
  - Baseband
    - Digital signaling
  - Broadband
    - Analog signaling
- Cabling (2, 5, T, F, S, L)
  - 5 - Thick coax (original Ethernet cabling)
  - F – Optical fiber
  - S – Short wave laser over multimode fiber
  - L – Long wave laser over single mode fiber

Ethernet Overview

- Most popular packet-switched LAN technology
- Bandwidths: 10Mbps, 100Mbps, 1Gbps
- Max bus length: 2500m
  - 500m segments with 4 repeaters
- Bus and Star topologies are used to connect hosts
  - Hosts attach to network via Ethernet transceiver or hub or switch
    - Detects line state and sends/receives signals
  - Hubs are used to facilitate shared connections
  - All hosts on an Ethernet are competing for access to the medium
    - Switches break this model
- Problem: Distributed algorithm that provides fair access

Ethernet Overview (contd.)

- Ethernet by definition is a broadcast protocol
  - Any signal can be received by all hosts
  - Switching enables individual hosts to communicate
- Network layer packets are transmitted over an Ethernet by encapsulating
- Frame Format

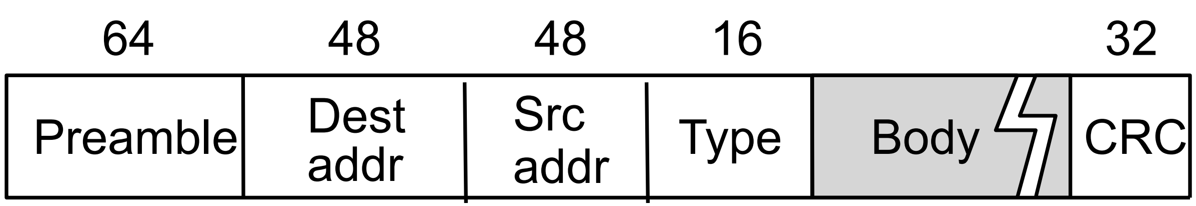

Switched Ethernet

- Switches forward and filter frames based on LAN addresses
  - It's not a bus or a router (although simple forwarding tables are maintained)
- Very scalable
  - Options for many interfaces
  - Full duplex operation (send/receive frames simultaneously)
- Connect two or more "segments" by copying data frames between them
  - Switches only copy data when needed
    - key difference from repeaters
- Higher link bandwidth
  - Collisions are completely avoided
- Much greater aggregate bandwidth
  - Separate segments can send at once

Ethernet Frames

- Preamble is a sequence of 7 bytes, each set to "10101010"
  - Used to synchronize receiver before actual data is sent
- Addresses
  - unique, 48-bit unicast address assigned to each adapter
    - example: **`8:0:e4:b1:2`**
    - Each manufacturer gets their own address range
  - broadcast: all **`1`**s
  - multicast: first bit is **`1`**
- Type field is a demultiplexing key used to determine which higher level protocol the frame should be delivered to
- Body can contain up to 1500 bytes of data

A Quick Word about Aloha Networks

- Developed in late 60's by Norm Abramson at Univ. of Hawaii (!!) for use with packet radio systems
  - Any station can send data at any time
  - Receiver sends an ACK for data
  - Timeout for ACK signals that there was a collision
    - What happens if timeout is poorly timed?
  - If there is a collision, sender will resend data after a random backoff
- Utilization (fraction of transmitted frames avoiding collision for N nodes) was pretty bad
  - Max utilization = 18%
- Slotted Aloha (dividing transmit time into windows) helped
  - Max utilization increased to 36%

Ethernet's MAC Algorithm

- In Aloha, decisions to transmit are made without paying attention to what other nodes might be doing
- Ethernet uses CSMA/CD – listens to line before/during sending
- If line is idle (no carrier sensed)
  - send packet immediately
  - upper bound message size of 1500 bytes
  - must wait 9.6us between back-to-back frames
- If line is busy (carrier sensed)
  - wait until idle and transmit packet immediately
    - called *1-persistent* sending
- If collision detected
  - Stop sending and jam signal
  - Try again later

State Diagram for CSMA/CD

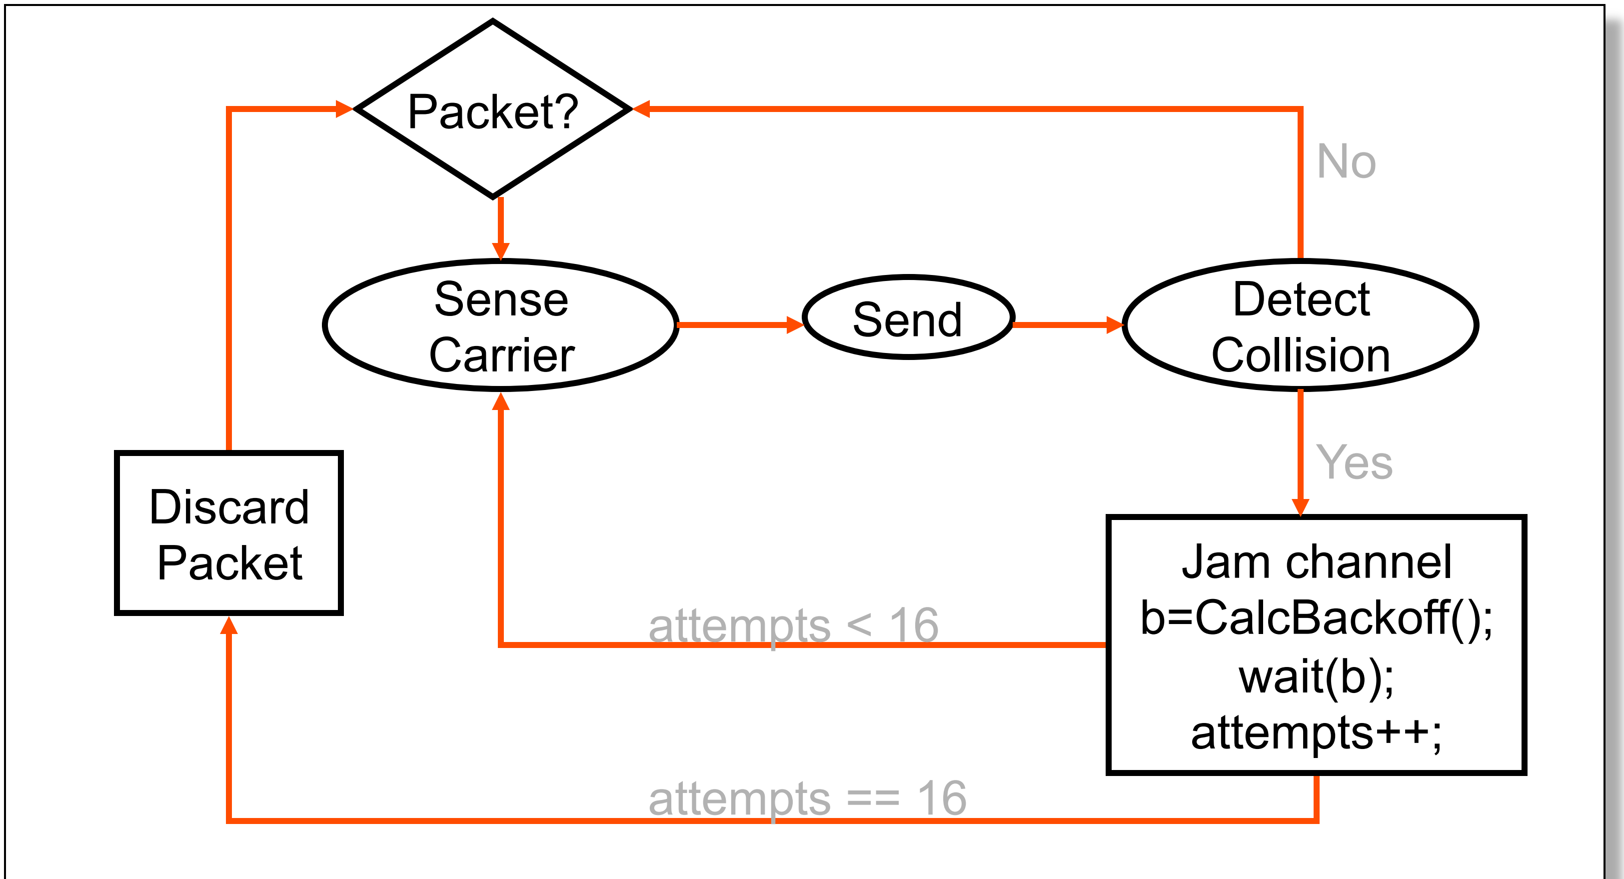

Collisions

Collisions are caused when two adaptors transmit at the same time (adaptors sense collision based on voltage differences)

- Both found line to be idle
- Both had been waiting to for a busy line to become idle

A starts at time 0

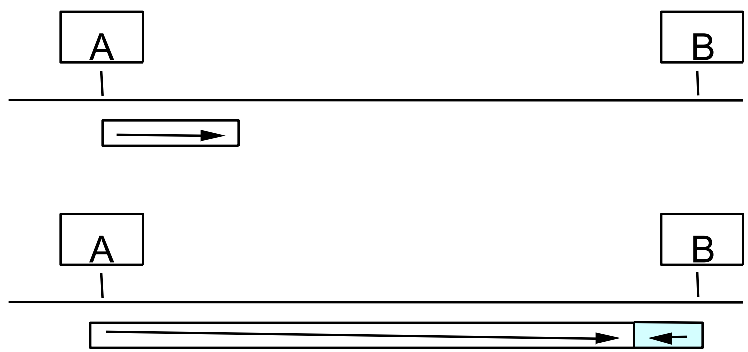

Message almost there at time T when B starts – collision!

How can we be sure A knows about the collision?

Collision Detection

- How can A know that a collision has taken place?
  - There must be a mechanism to insure retransmission on collision
  - A's message reaches B at time T
  - B's message reaches A at time 2T
  - So, A must still be transmitting at 2T
- IEEE 802.3 specifies max value of 2T to be 51.2us
  - This relates to maximum distance of 2500m between hosts
  - At 10Mbps it takes 0.1us to transmit one bit so 512 bits (64B) take 51.2us to send
  - So, Ethernet frames must be at least 64B long
    - 14B header, 46B data, 4B CRC
    - Padding is used if data is less than 46B
- Send jamming signal after collision is detected to insure all hosts see collision
  - 48 bit signal

Collision Detection contd.

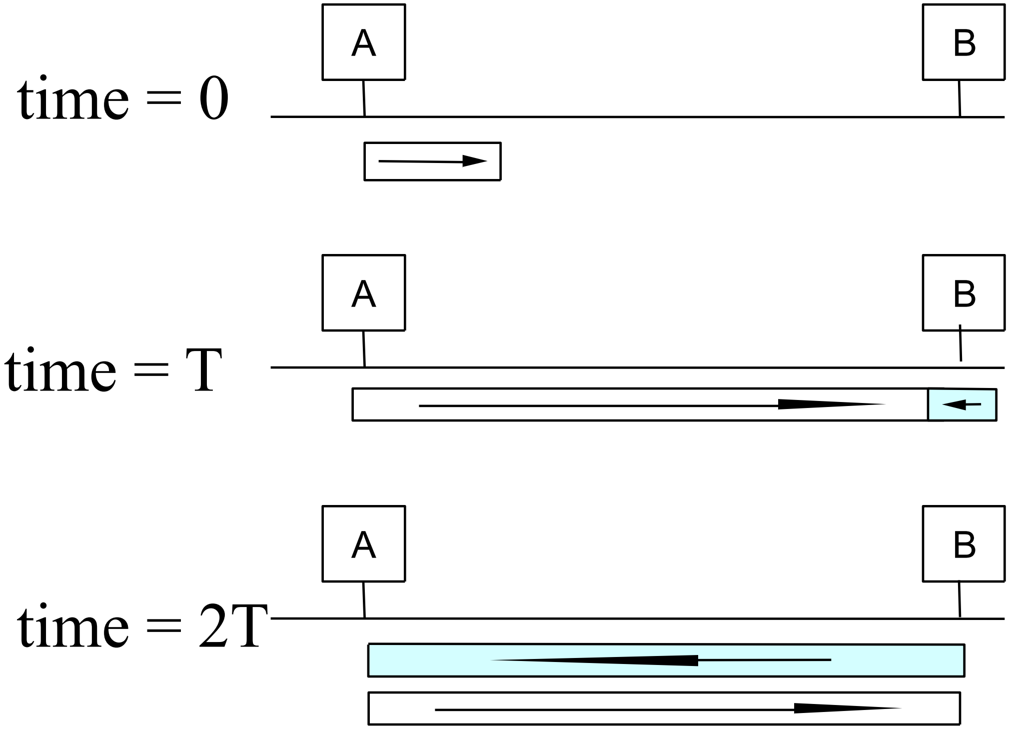

Exponential Backoff

- If a collision is detected, delay and try again
- Delay time is selected using binary exponential backoff
  - 1st time: choose K from {0,1} then delay = K * 51.2us
  - 2nd time: choose K from {0,1,2,3} then delay = K * 51.2us
  - *nth* time: delay = $K \times 51.2$us, for $K=0..2^n - 1$
    - Note max value for k = 1023
  - give up after several tries (usually 16)
    - Report transmit error to host
- If delay were not random, then there is a chance that sources would retransmit in lock step
- Why not just choose from small set for K
  - This works fine for a small number of hosts
  - Large number of nodes would result in more collisions

CSMA/CD

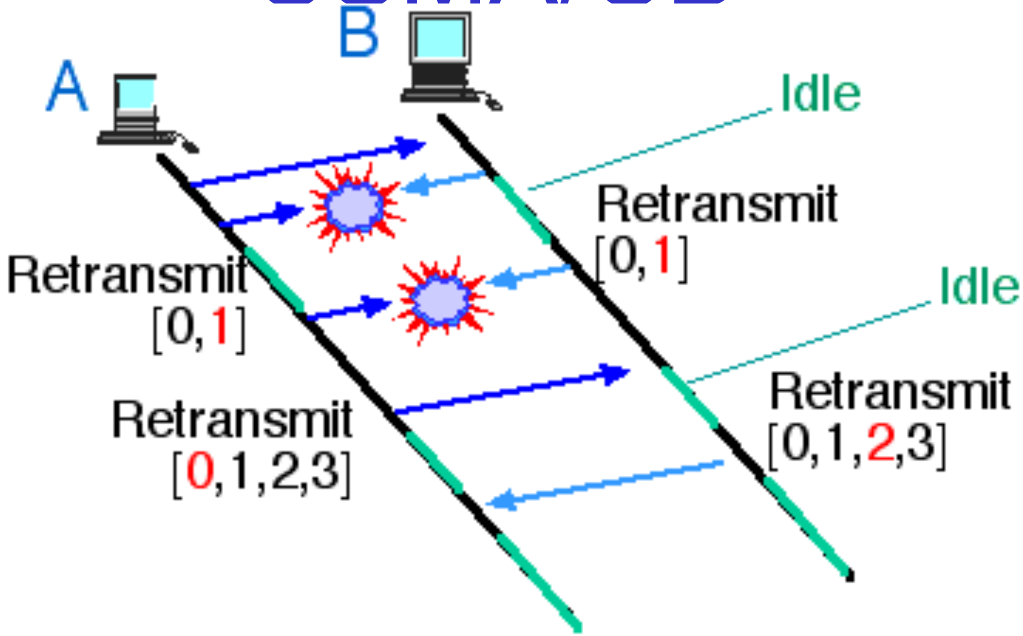

- CSMA stands for Carrier Sense Multiple Access
- CD stands for Collision Detection
- Allows for Retransmission Back-Off

MAC Algorithm from the Receiver Side

- Senders handle all access control
- Receivers simply read frames with acceptable address
  - Address to host
  - Address to broadcast
  - Address to multicast to which host belongs
  - All frames if host is in promiscuous mode

Fast and Gigabit Ethernet

- Fast Ethernet (100Mbps) has technology very similar to 10Mbps Ethernet
  - Uses different physical layer encoding (4B5B)
  - Many NIC's are 10/100 capable
    - Can be used at either speed
- Gigabit Ethernet (1,000Mbps)
  - Compatible with lower speeds
  - Uses standard framing and CSMA/CD algorithm
  - Distances are severely limited
  - Typically used for backbones and inter-router connectivity
  - Becoming cost competitive
  - How much of this bandwidth is realizable?

Experiences with Ethernet

- Ethernets work best under light loads
  - Utilization over 30% is considered heavy
    - Network capacity is wasted by collisions
- Most networks are limited to about 200 hosts
  - Specification allows for up to 1024
- Most networks are much shorter
  - 5 to 10 microsecond RTT
- Transport level flow control helps reduce load (number of back to back packets)
- Ethernet is inexpensive, fast and easy to administer!

Ethernet Problems

- Ethernet's peak utilization is pretty low (like Aloha)
- Peak throughput worst with
  - More hosts
    - More collisions needed to identify single sender
  - Smaller packet sizes
    - More frequent arbitration
  - Longer links
    - Collisions take longer to observe, more wasted bandwidth
  - Efficiency is improved by avoiding these conditions

Why did Ethernet Win?

- There are LOTS of LAN protocols
- **Price**
- Performance
- Availability
- Ease of use
- Scalability

Wireless networking

- No wires!
- Shared media
- Interference and noise
- Bandwidth and range limits
- Variable performance
- Mobility
- Implications for higher network layers

Wireless Spectrum (1)

Broadcast TV
- VHF: 54 to 88 MHz, 174 to 216 MHz
- UHF: 470 to 806 MHz

30 MHz 300 MHz 3 GHz 30 GHz

FM Radio
- 88 to 108 MHz

Digital TV
- 54 to 88 MHz, 174 to 216 MHz, 470 to 806 MHz

Wireless Spectrum (2)

3G Broadband Wireless
- 746-794 MHz, 1.7-1.85 GHz, 2.5-2.7 GHz

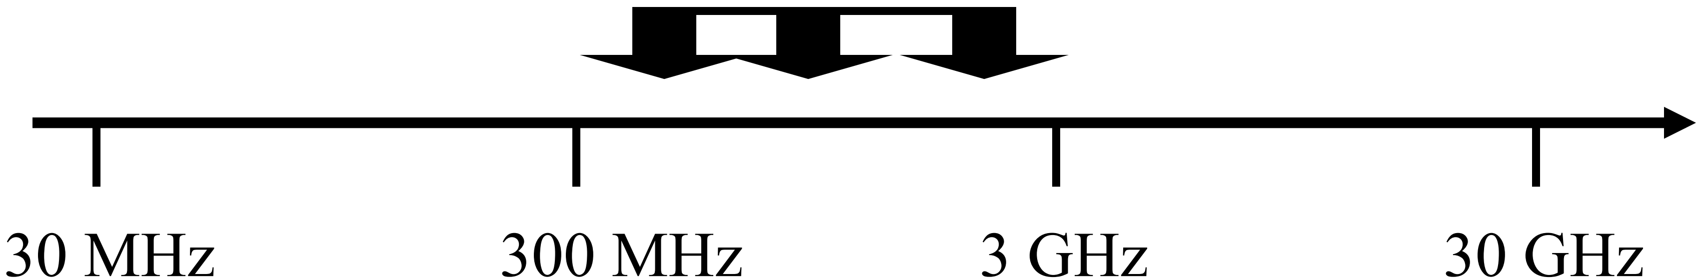

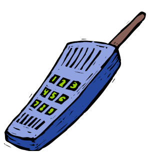

Cellular Phone
- 800-900 MHz

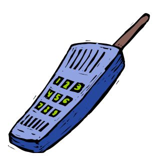

Personal Communication Service (PCS)
- 1.85-1.99 GHz

Wireless Spectrum (3)

Wireless LAN
(IEEE 802.11b/g)
• 2.4 GHz

Wireless LAN
(IEEE
802.11a)
• 5 GHz

30 MHz　300 MHz　3 GHz　30 GHz

Bluetooth
• 2.45 GHz

Local Multipoint
Distribution Service
(LMDS)
• 27.5-31.3 GHz

Signal attenuation

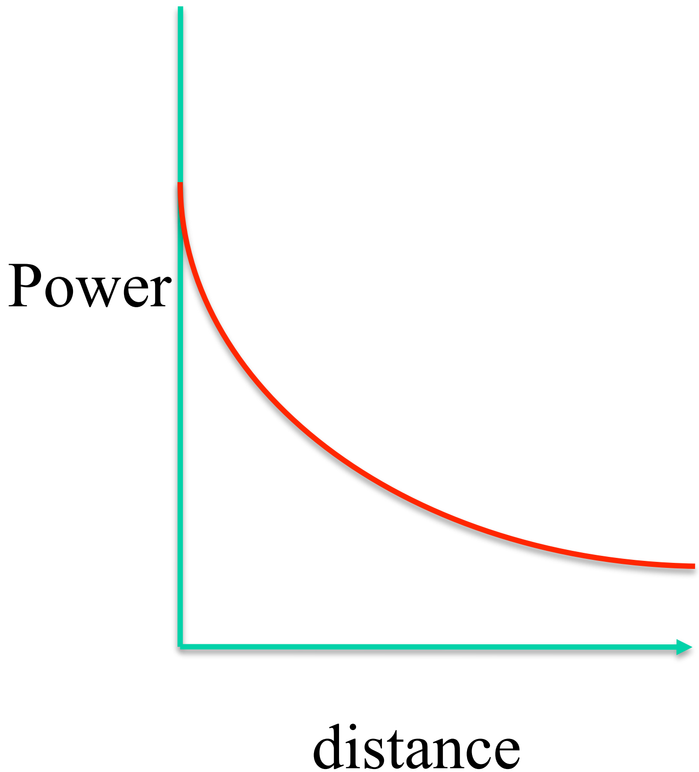

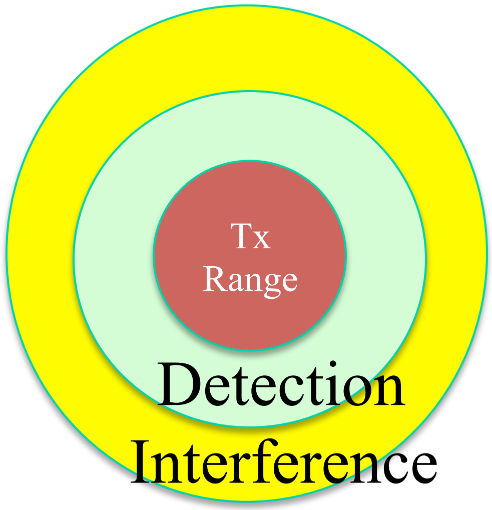

Signal attenuation

- A simple model of attenuation in free space
- Assumes
  - Energy radiates in all directions
  - Consider a sphere (4 *pi *r ^2)
  - Signal attenuates as 1/r^2
- SINR = Signal to Interference + Noise Ratio

Passing through a door

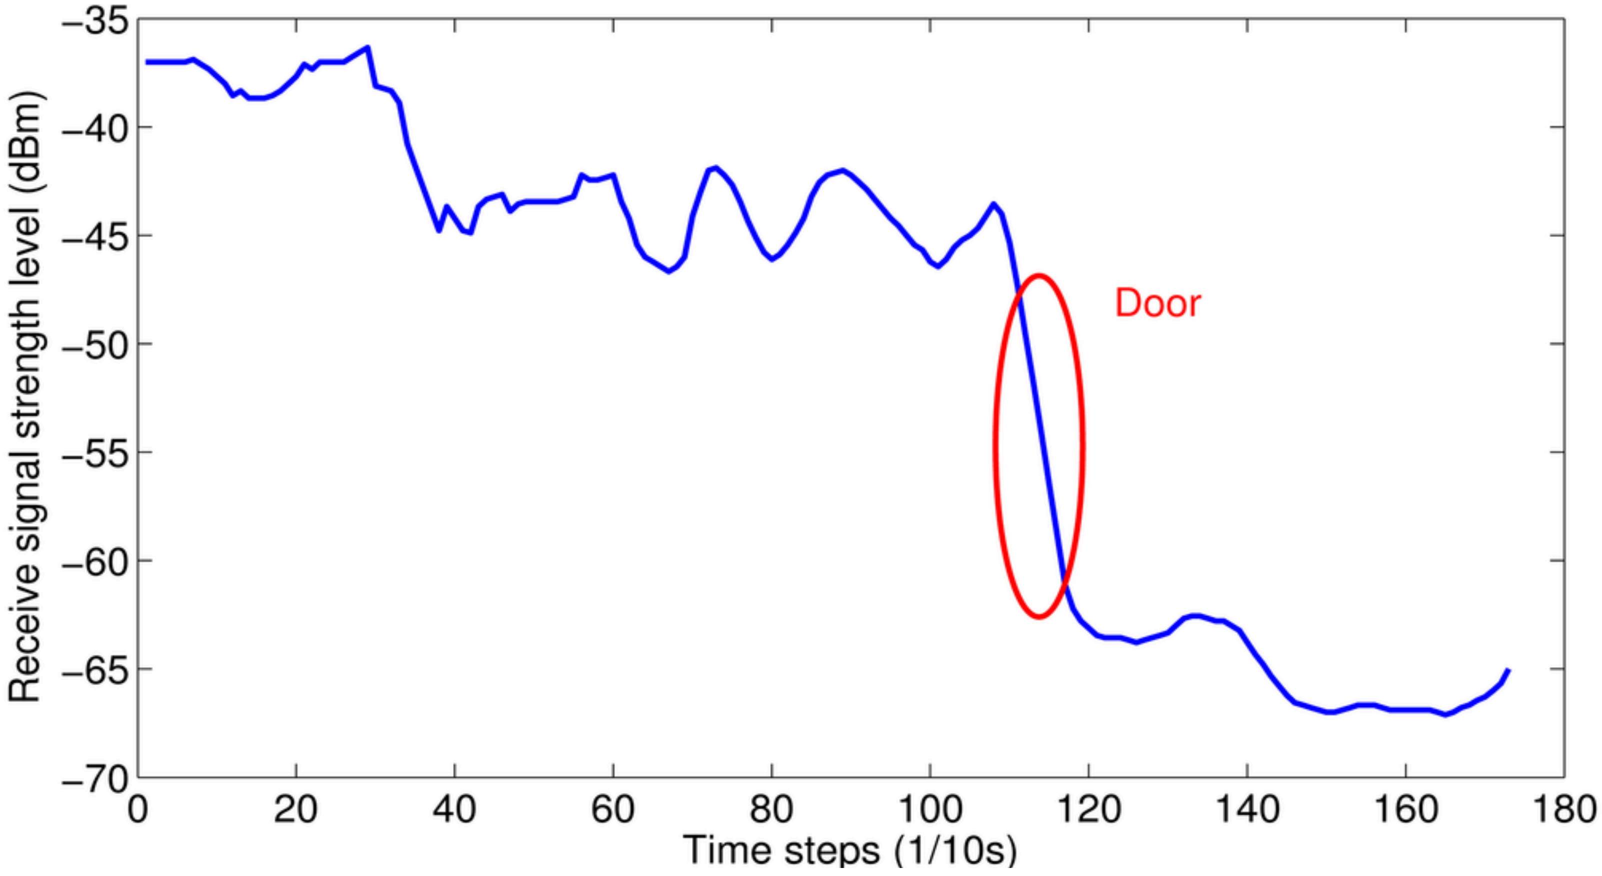

Bandwidth, Capacity, Throughput

- Bandwidth: Width of the channel on which we are operating (units: Hz)
  - WiFi channels are ~20 MHz, or 40 MHz, or 80 MHz
  - TV channels are 6 MHz

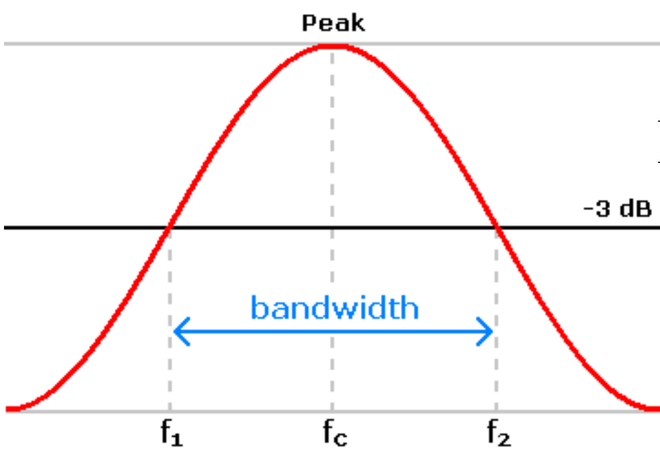

Bandwidth
=
Cross
section

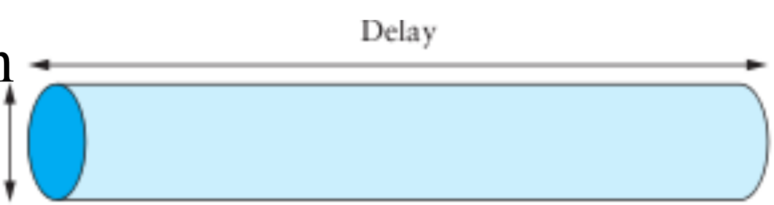

Bandwidth, Capacity, Throughput

- Capacity: An ideal notion of (max) rate at which data can be sent across a channel (units: bits/sec)
  - Analogy, capacity of a water pipe = what is the maximum rate at which water can flow through

Capacity
=
Max rate
of water
flowing
through

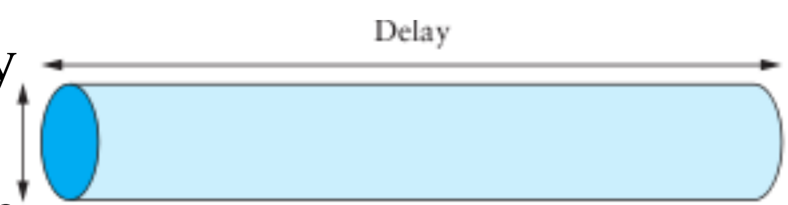

Bandwidth, Capacity, Throughput

- Throughput: What is the rate at which data is being sent across the channel (units: bits/sec)
  - Depends on real factors, e.g., amount of interference, signal quality, antenna properties, etc.
  - Analogy, throughput of a water pipe = rate at which water is actually flowing through it, which depends on blockages inside and other factors

Thruput
=
how much
Water
Actually

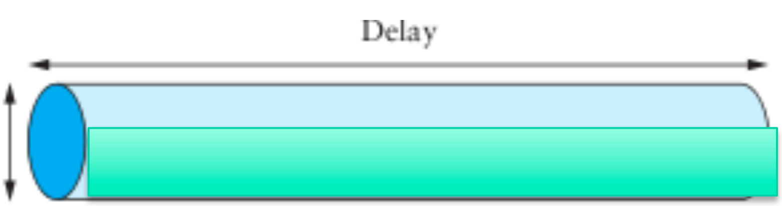

Channel capacity

- Shannon’ law
  - C = B log (1 + SINR)

- SINR: Signal to Interference + Noise Ratio

$$\mathrm{SINR} = \frac{S}{I + N}$$

  - S: Received power of signal at the receiver
  - I: Interference power
  - N: Noise power

Units dor SINR (or SNR)?

- SINR is a ratio, so unitless

- However, we commonly use a unit Bels (or deciBels) to refer to SINR or SNR values

- SINR (in Bels) = log_10 (actual SINR value)
- SINR (in deciBels) = 10 log_10 (actual SINR value)

dB

- dB = deciBels

- dBm = 10 log10 (power in mW)

- dBW = 10 log10 (power in W)

- Express 1 W in dBW and dBm
- Express 1 mW in dBW and dBm
- Expres 10 mW in dBW and dBm

Some values for SNR

- 802.11a/n, Tx power = 40 mW = 14 dBm
- 802.11b, Tx power = 100 mW = 20 dBm
- 802.11g, Tx power = 50 mW = 17 dBm

Free space loss
100m ~ 80 dB

Tx power + antenna gain = EIRP
(17 dBm) + (6 dBi) = (23 dBm)

Received signal
23 – 80 = -57 dBm

- If noise = -95 dBm, then SNR = -57+95 = 38 dB

Channel capacity

- $C = B \log\left(1 + \frac{S}{I+N}\right)$

- How to increase it?
  - Increase B
  - Increase S

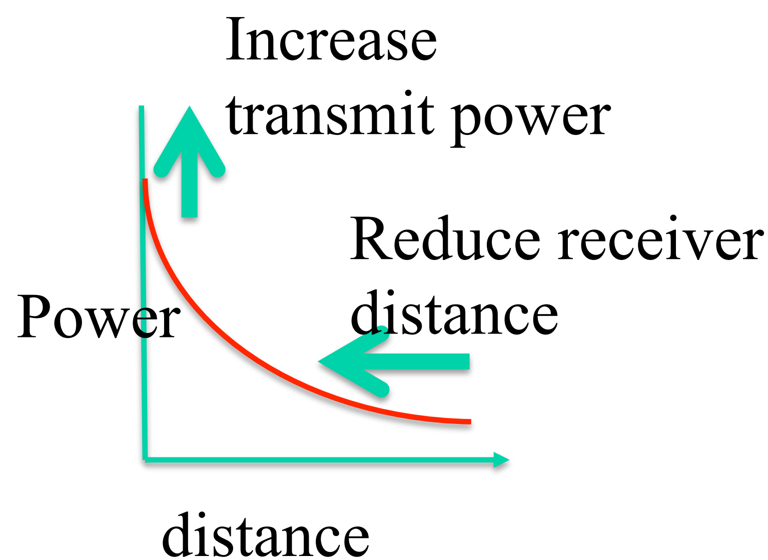

Efficiency of a wireless system

- Spectral efficiency --- throughput per unit frequency : units of b/s/Hz
- A good way to compare technologies

- Check http://en.wikipedia.org/wiki/Spectral_efficiency for some example spectral efficiencies

Example 1

Consider an extremely noisy channel in which the value of the signal-to-noise ratio is almost zero. In other words, the noise is so strong that the signal is faint. For this channel the capacity C is calculated as

$$C = B \log_2 (1 + \text{SNR}) = B \log_2 (1 + 0) = B \log_2 1 = B \times 0 = 0$$

This means that the capacity of this channel is zero regardless of the bandwidth. In other words, we cannot receive any data through this channel.

Example 2

We can calculate the theoretical highest bit rate of a regular telephone line. A telephone line normally has a bandwidth of 3000 Hz. The signal-to-noise ratio is usually 3162. For this channel the capacity is calculated as

$$C = B \log_2 (1 + \text{SNR}) = 3000 \log_2 (1 + 3162) = 3000 \log_2 3163$$
$$= 3000 \times 11.62 = 34{,}860 \text{ bps}$$

This means that the highest bit rate for a telephone line is 34.860 kbps. If we want to send data faster than this, we can either increase the bandwidth of the line or improve the signal-to-noise ratio.

Example 3

The signal-to-noise ratio is often given in decibels. Assume that SNR_{dB} = 36 and the channel bandwidth is 2 MHz. The theoretical channel capacity can be calculated as

$$SNR_{dB} = 10 \log_{10} SNR \rightarrow SNR = 10^{SNR_{dB}/10} \rightarrow SNR = 10^{3.6} = 3981$$

$$C = B \log_2 (1+ SNR) = 2 \times 10^6 \times \log_2 3982 = 24 \text{ Mbps}$$

Example 4

For practical purposes, when the SNR is very high, we can assume that SNR + 1 is almost the same as SNR. In these cases, the theoretical channel capacity can be simplified to

$$\boxed{C = B \times \frac{\mathrm{SNR_{dB}}}{3}}$$

For example, we can calculate the theoretical capacity of the previous example as

$$\boxed{C = 2\ \mathrm{MHz} \times \frac{36}{3} = 24\ \mathrm{Mbps}}$$

Example 5

We have a channel with a 1-MHz bandwidth. The SNR for this channel is 63. What is the appropriate bit rate?

Solution

We use the Shannon formula to find the upper limit.

$$\boxed{C = B \log_2 (1 + \text{SNR}) = 10^6 \log_2 (1 + 63) = 10^6 \log_2 64 = 6 \text{ Mbps}}$$

IEEE 802.11

- Characteristics
- Modes of operation
- Association, authentication and privacy

IEEE 802.11 Standard

- Final draft approved in 1997
- Operates in the 2.4 GHz industrial, scientific and medical (ISM) band
- Standard defines the physical (PHY) and medium access control (MAC) layers
  - Note that the 802.11 MAC layer also performs functions that we usually associated with higher layers (e.g., fragmentation, error recovery, mobility management)
- Initially defined for operation at 1 and 2 Mbps
  - DSSS, FHSS or infrared
  - Extensions (IEEE 802.11b, IEEE 802.11a, etc.) allow for operation at higher data rates and (in the case of 802.11a) different frequency bands

Reference Model (1)

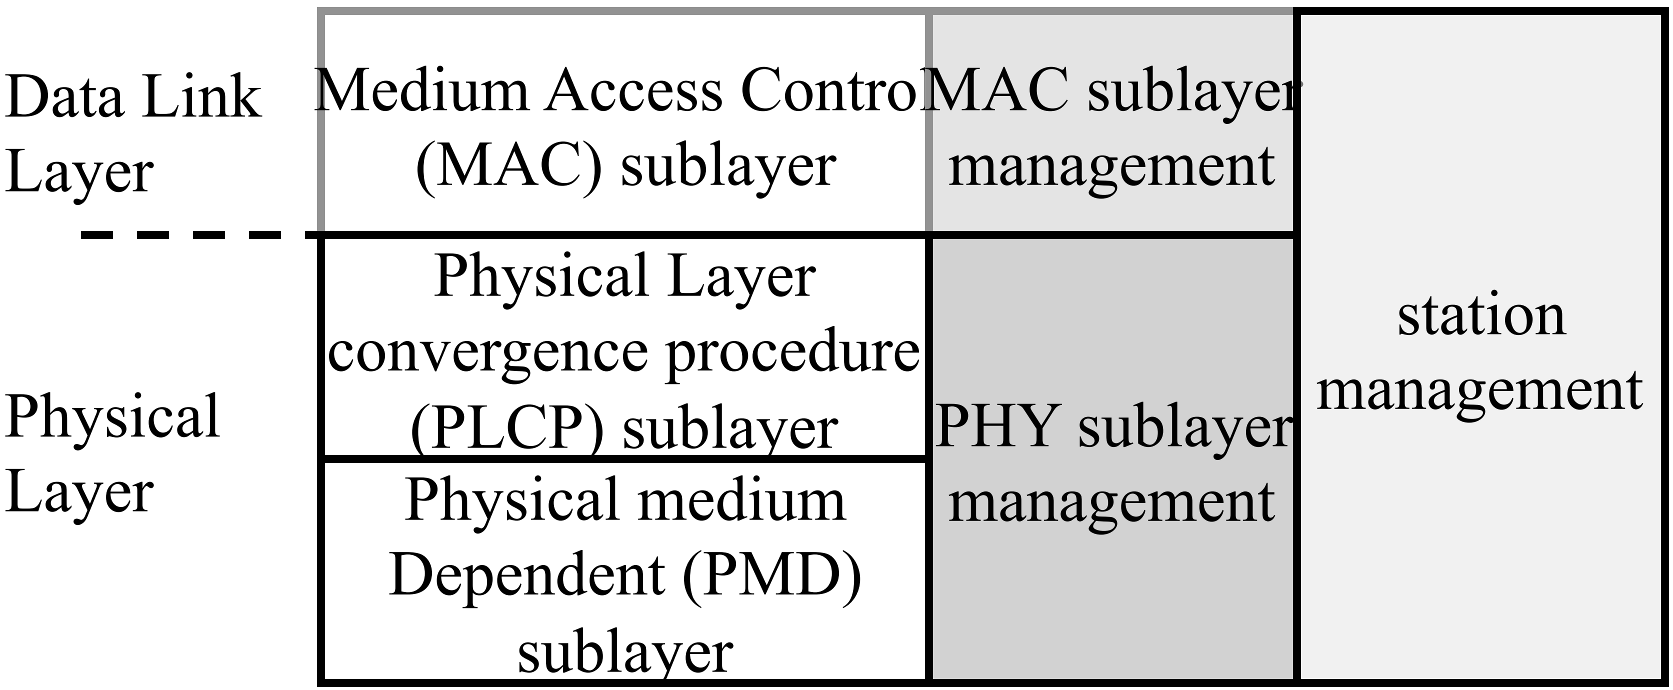

Reference Model (2)

- Physical Medium Dependent (PMD) sublayer
  - Defines a method for transmitting and receiving data through the medium, including modulation and coding
  - Dependent on whether DSSS, FHSS or IR is used
- Physical Layer Convergence Procedure (PLCP) sublayer
  - Maps MAC layer PDUs into a packet suitable for transmission by the PMD sublayer
  - Performs carrier sensing
- MAC sublayer
  - Defines access mechanism, based on CSMA
  - Performs fragmentation and encryption of data packets

IEEE 802.11b

- Standard released in 1999
- 2.4 – 2.483 GHz band
- Uses DSSS
- Data rates of up to 11 Mbps
  - Data rates are automatically adjusted for noisy conditions, so can operate at 1, 2, 5.5 or 11 Mbps
- Modes of operation
  - Infrastructure-based
  - Ad-hoc

Infrastructure Mode (1)

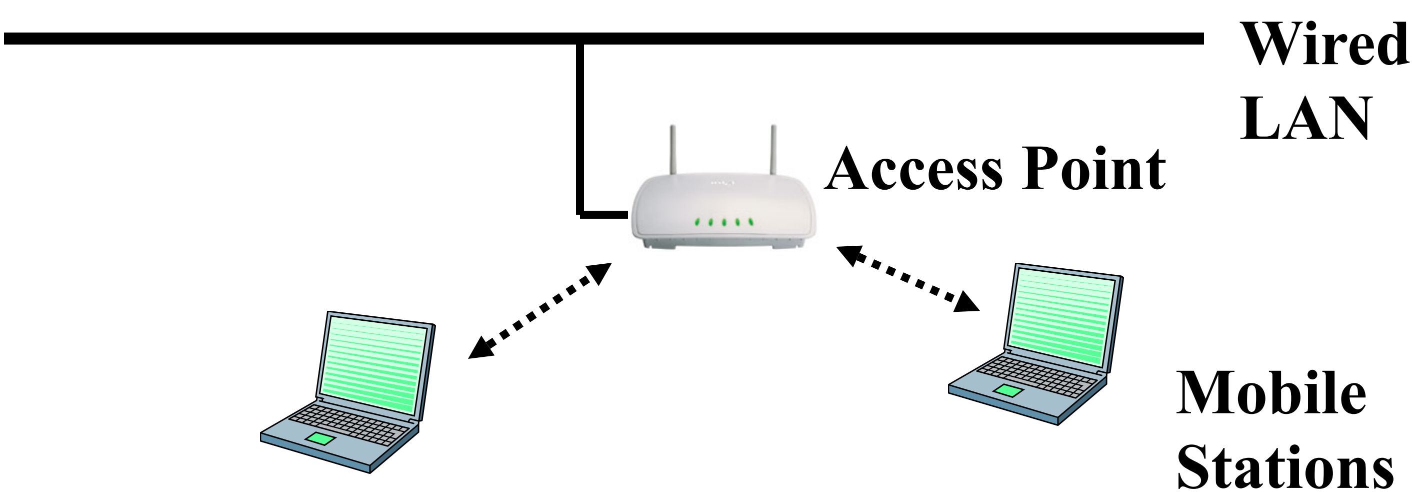

- Basic Service Set (BSS)
- Access point serves as a local bridge
- Stations communicate through the access point, which relays frames to/from mobile stations

Infrastructure Mode (2)

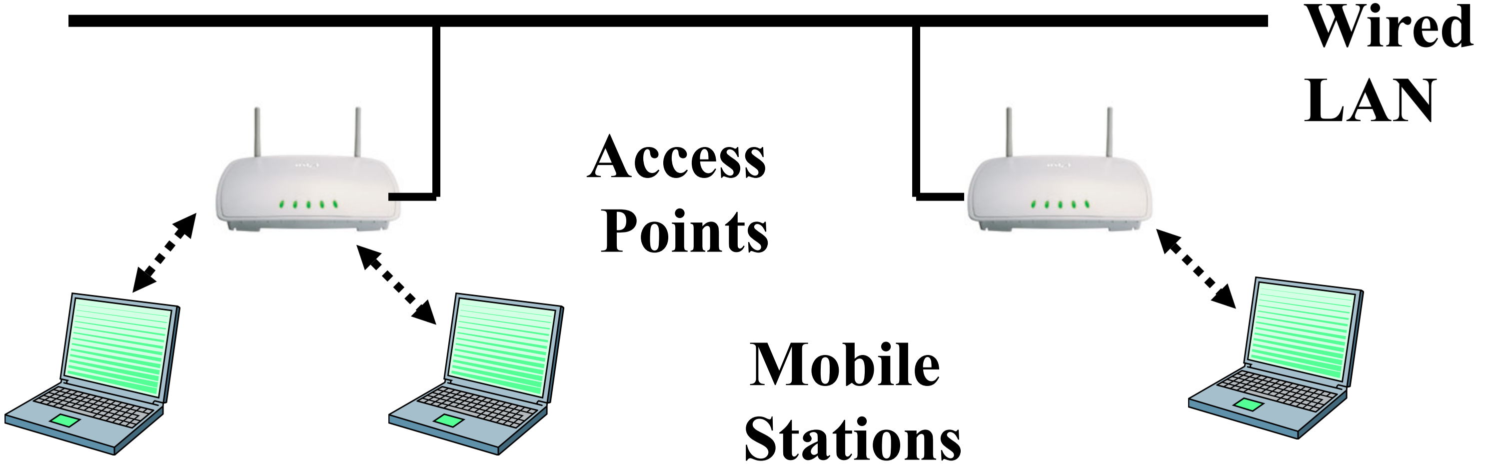

- Extended Service Set (ESS)
- A set of infrastructure BSSs
- Access points communicate among themselves to forward frames between BSSs and to facilitate movement of stations between BSSs

Ad Hoc Mode

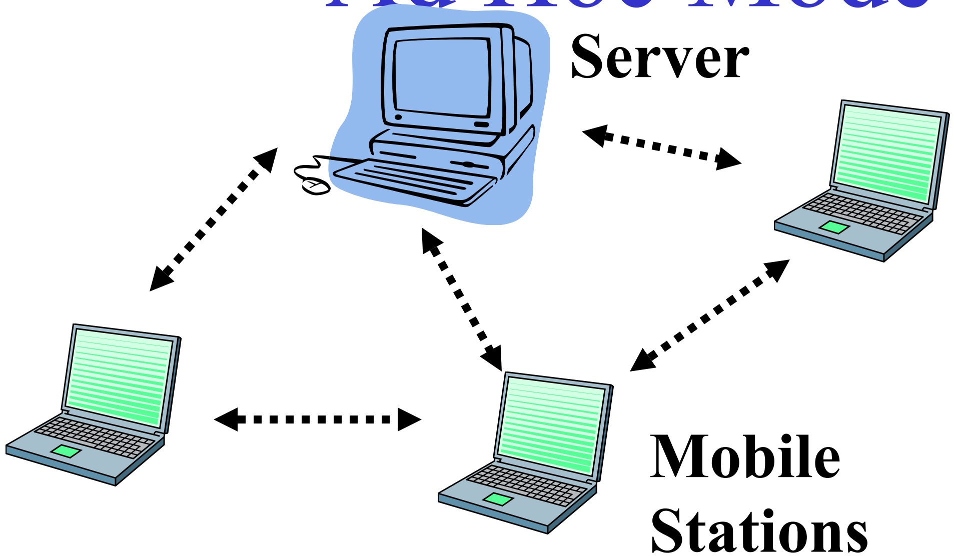

- Independent Basic Service Set (IBSS) or Peer to Peer
- Stations communicate directly with each other
- When no direct link is feasible between two station, a third station may act as a relay (multi-hop communications)

Distribution Systems

- The architectural component used to interconnect BSSs is the distribution system (DS)
- DS enable mobile device support
  - Address-to-destination mapping
  - Seamless integration of several BSSs
- In practice, an access point implements DS services

Distribution Systems and Access Points

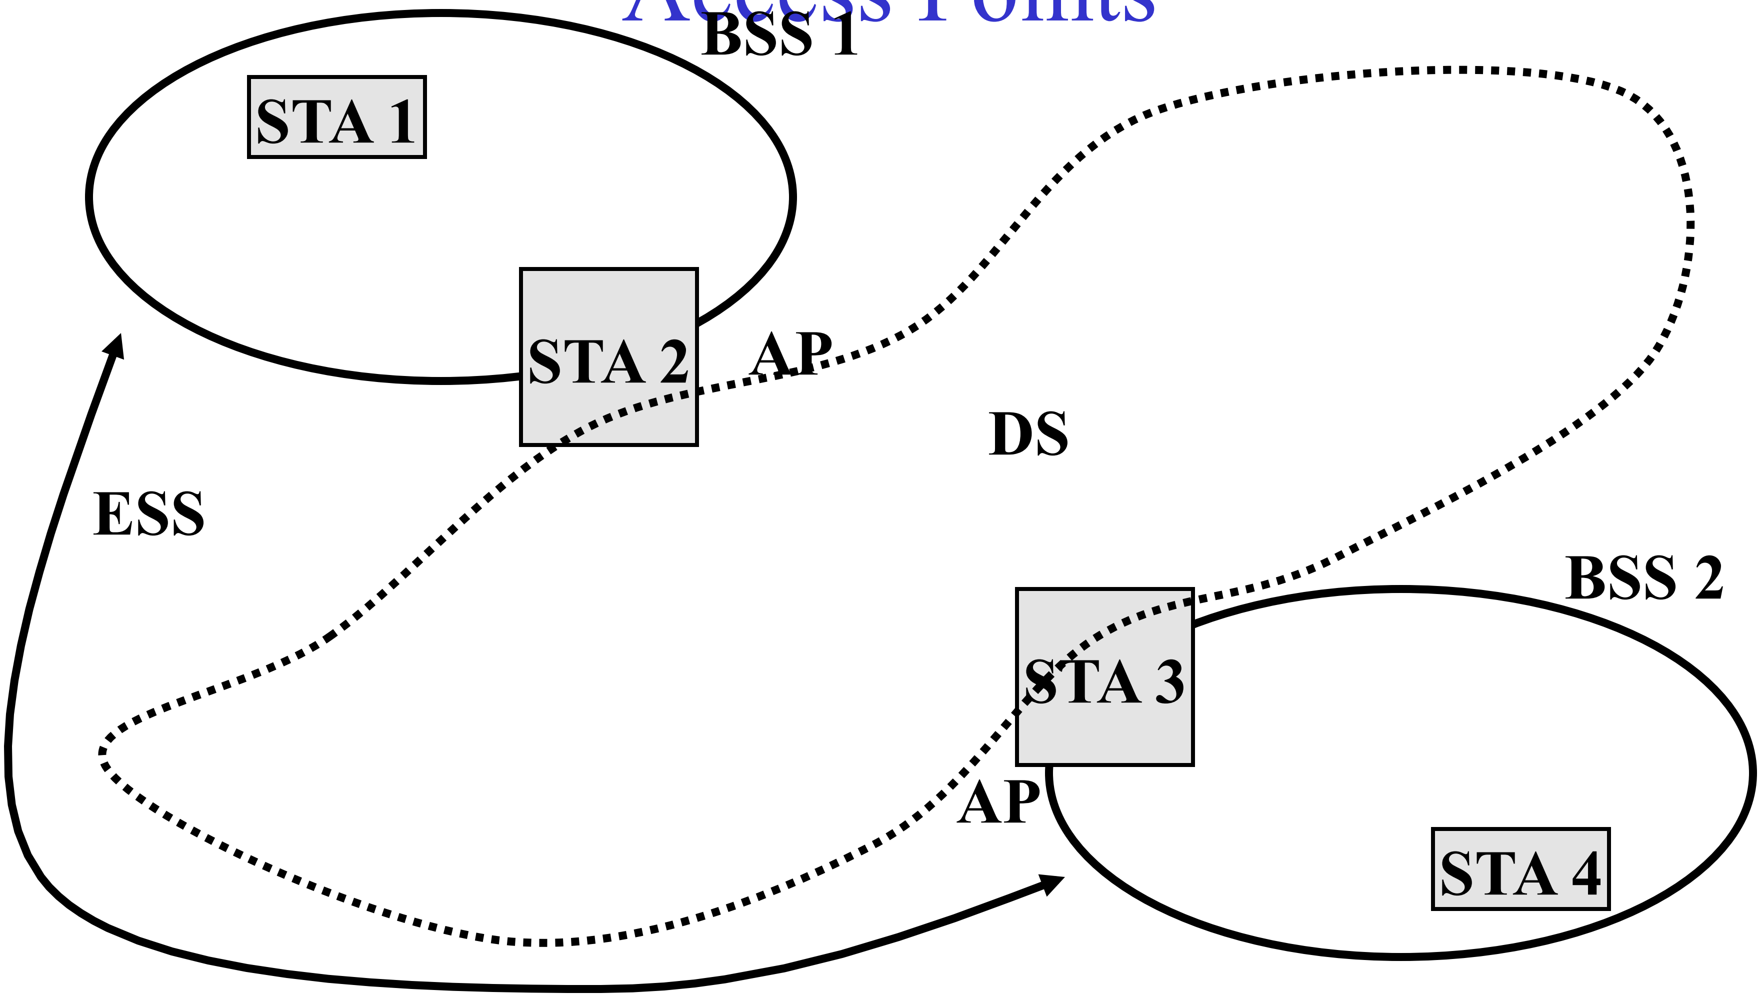

Integration with Wired LANs

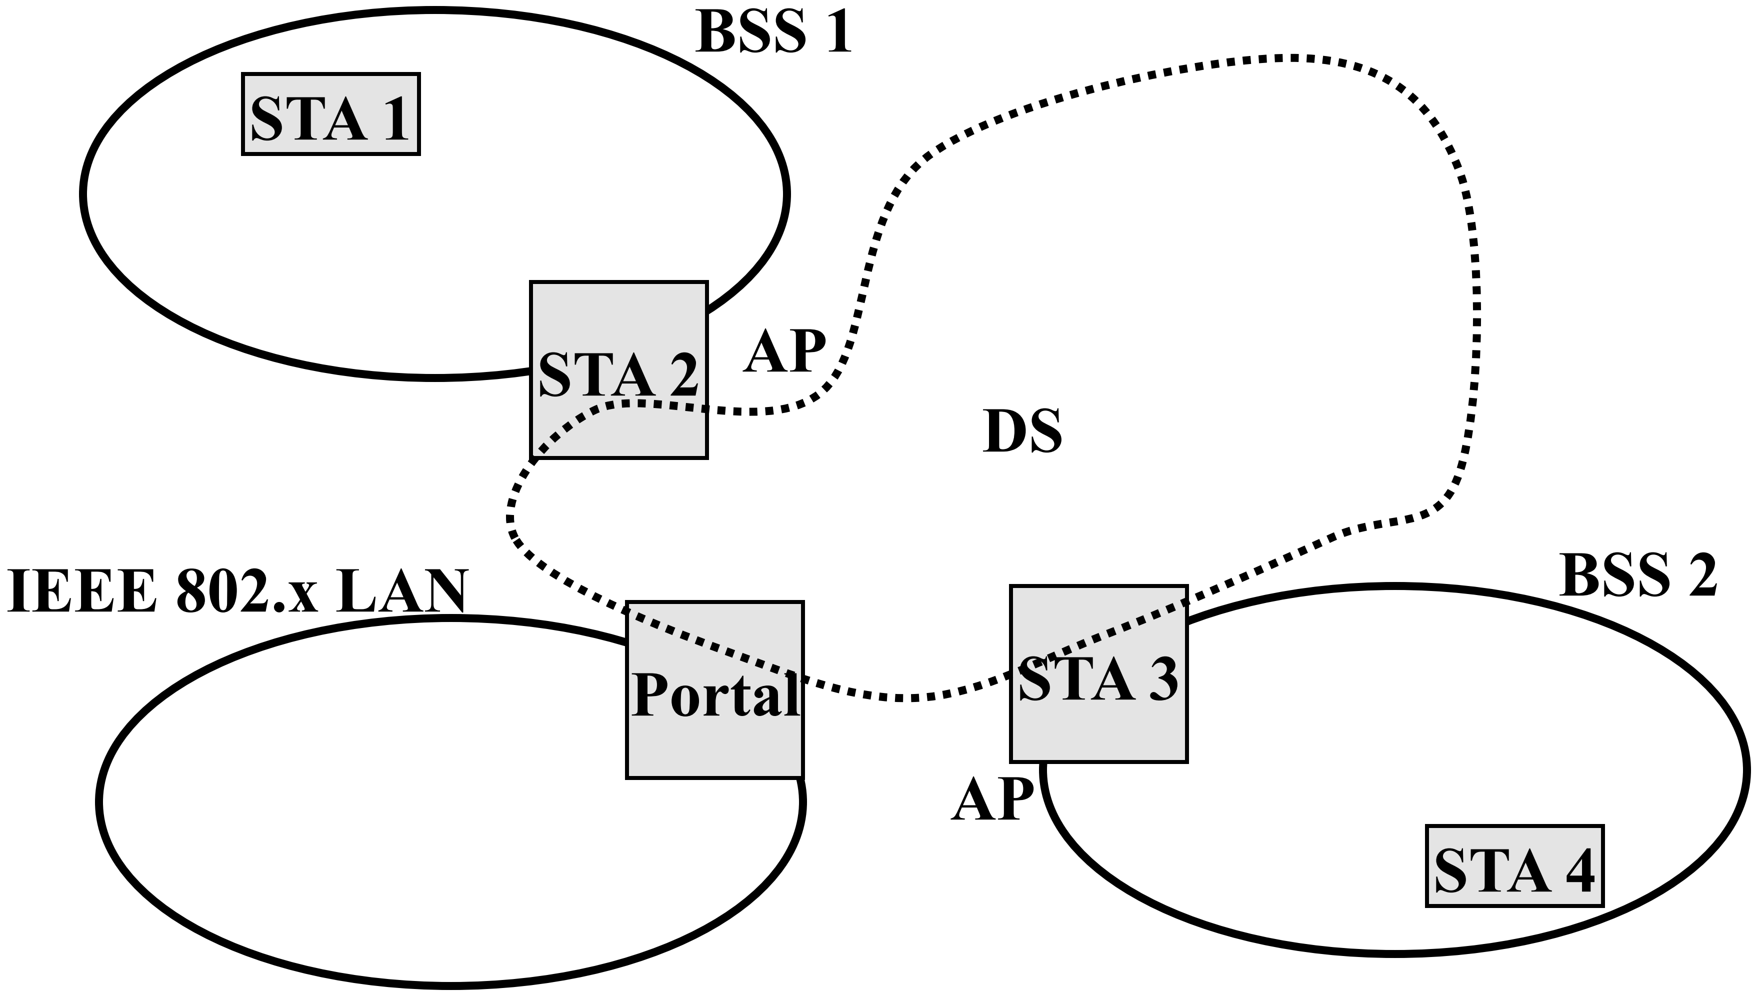

Steps to Building an 802.11 Connection

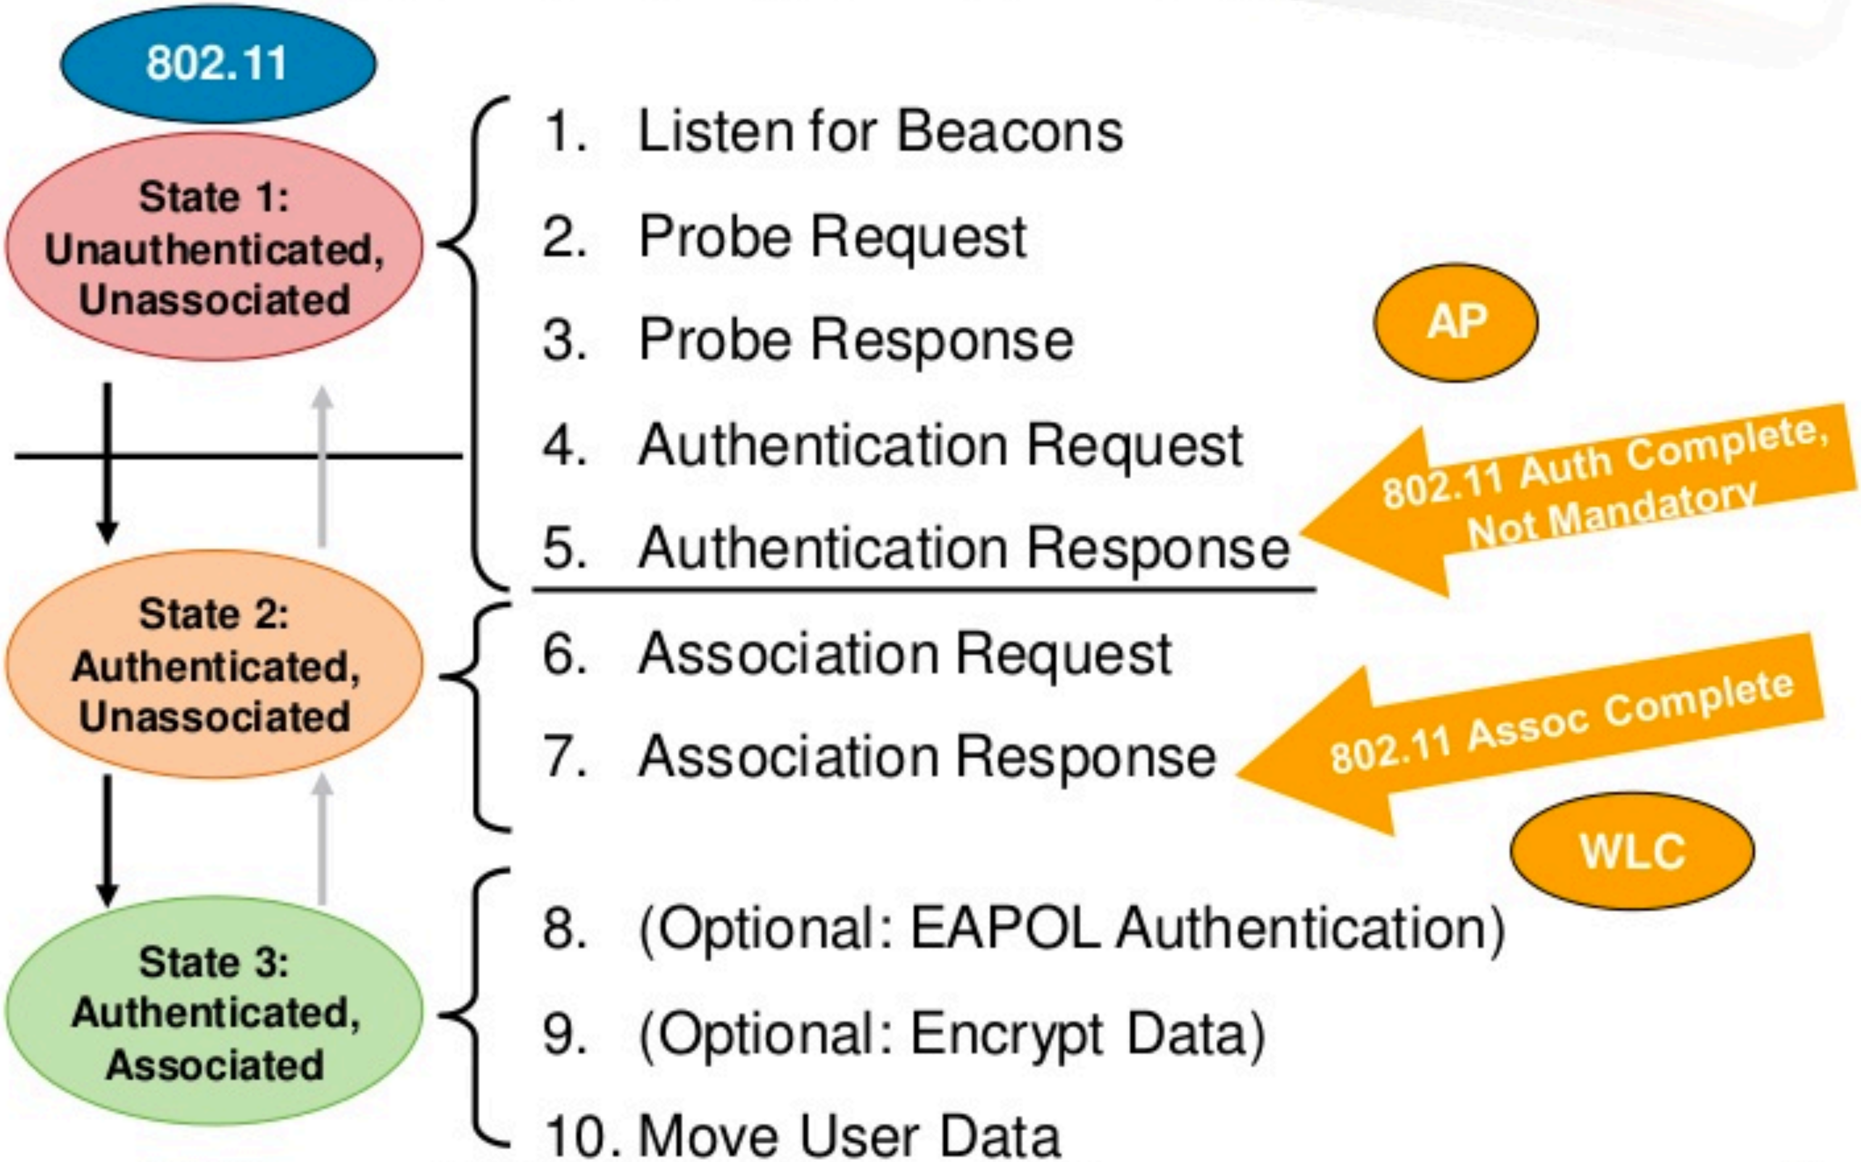

Authentication

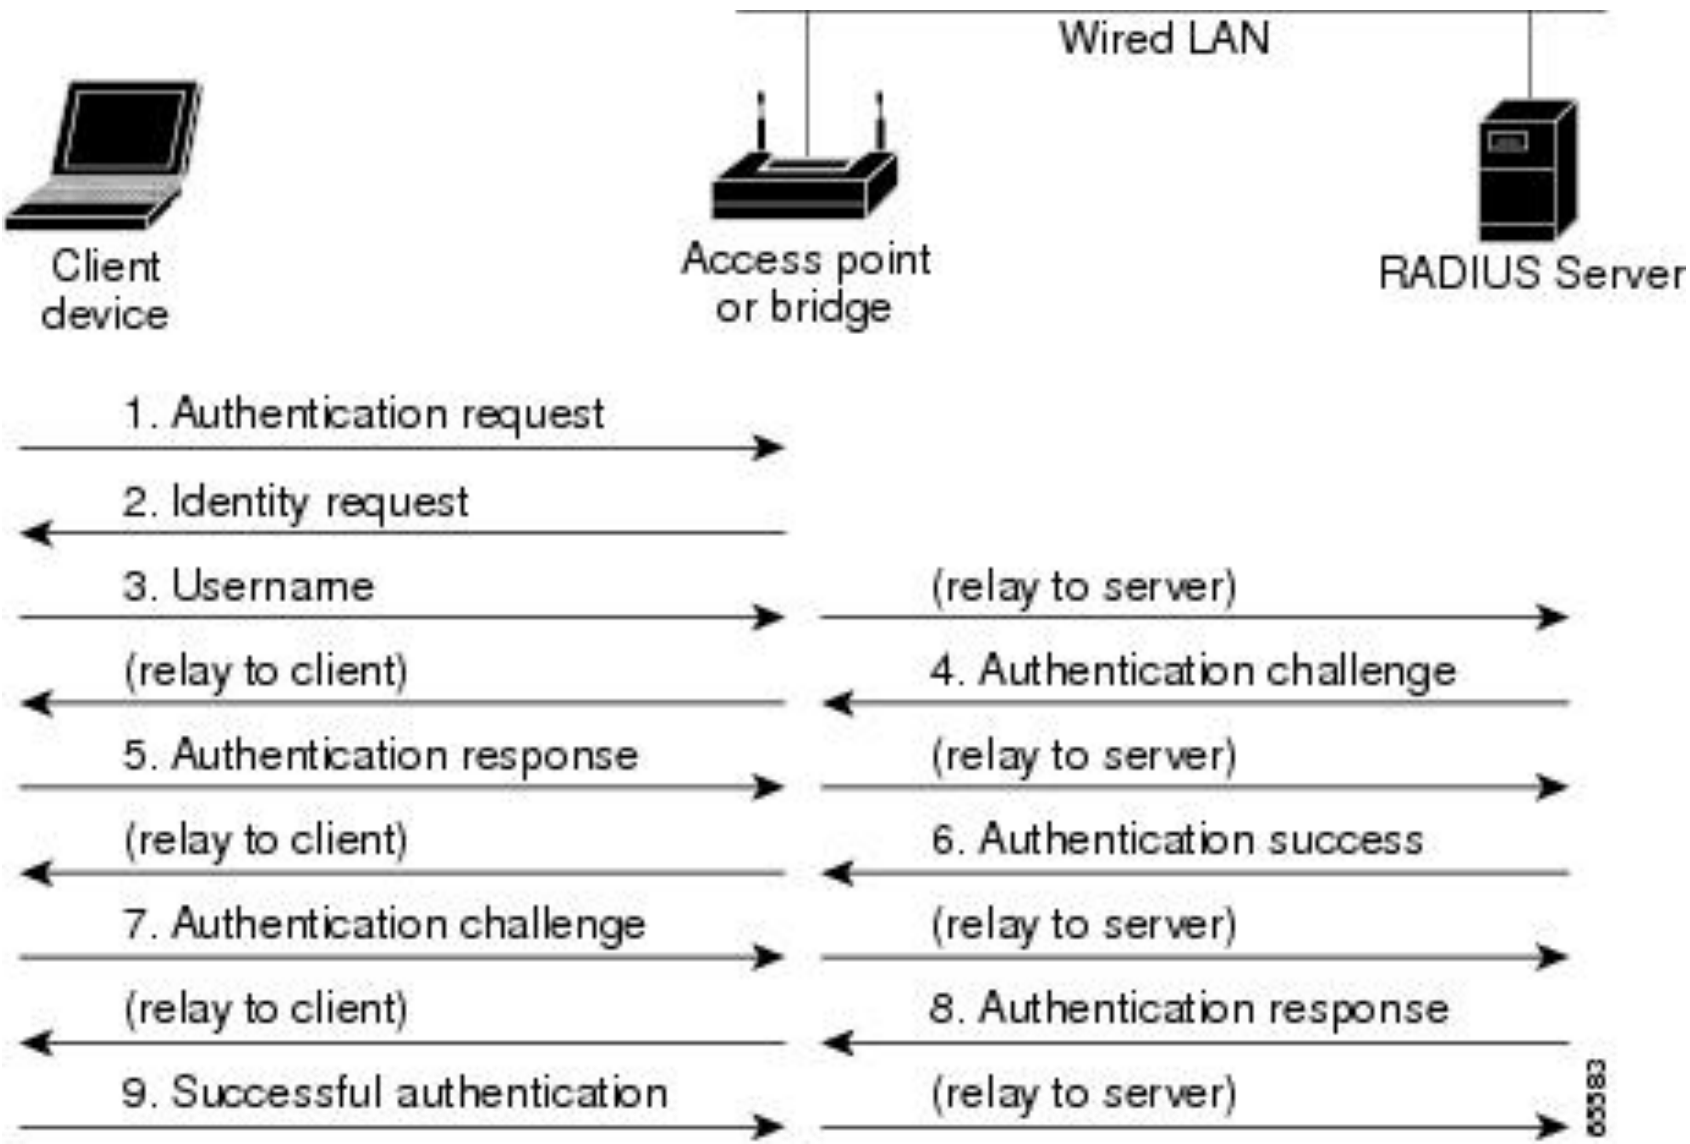

Association

- To deliver a message within the DS, must know which AP to access for a given mobile station
- Before a station is allowed to send a message through an AP, it must <u>associate</u> itself with that AP
  - At any given time, a station must be associated with no more than one AP
  - An AP may be associated with multiple stations
- As it moves between BSSs, a mobile station may <u>reassociate</u> itself with a different AP

Authentication

- 802.11 provides link-level authentication between stations
- 802.11 also supports shared key authentication
  - Requires that wired equivalent privacy (WEP) be enabled
  - Identity is demonstrated by knowledge of a shared, secret, WEP encryption key
- Typically, authentication is performed at association with an AP

Channel contention
(Medium Access)

MAC in Detail

- Channel access mechanism
  - Distributed Coordination Function (DCF)
    - Carrier sense multiple access (CSMA) with immediate MAC-level ACK
    - RTS/CTS(4-way handshaking) exchange (optional)
      - RTS stands for Request-to-Send
      - CTS stands for Clear-to-Send
  - Point Coordination Function (PCF)
    - Polled access through AP and distributed access
    - Contention-free period (CFP) and contention period (CP)
    - Seldom implemented in practice

Slide taken from: http://www.cse.ucsc.edu/classes/cmpe257/Spring05/lecture/lecture3.pdf#search='802.11e%20MAC%20protocol%2C%20powerpoint

Channel contention (CSMA/CA)

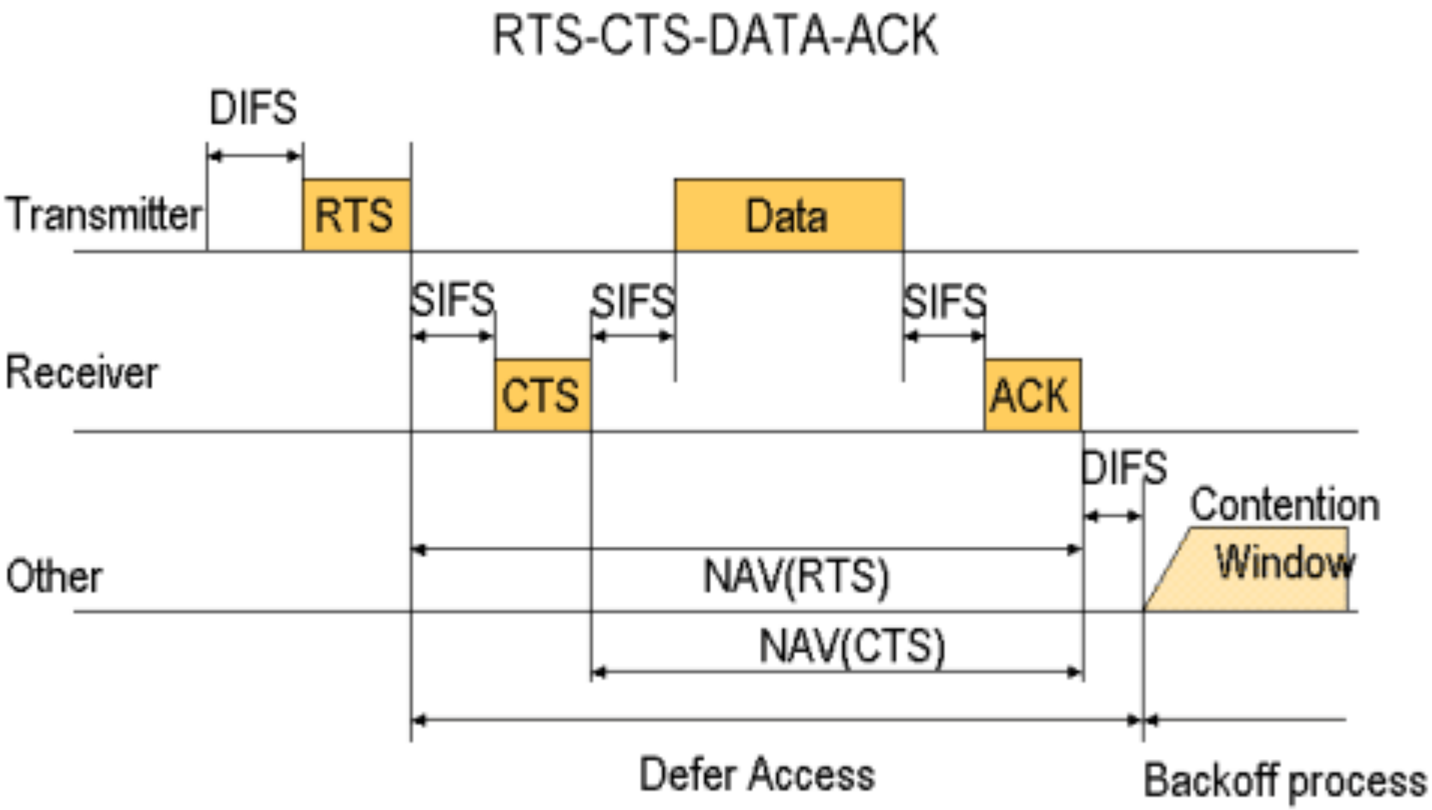

DIFS: Distributed IFS
RTS: Request To Send
SIFS: Short IFS
CTS: Clear To Send

ACK: Acknowledgement
NAV: Network Allocation Vector
DCF: Distributed Coordination Function

CSMA/CA and RTS/CTS

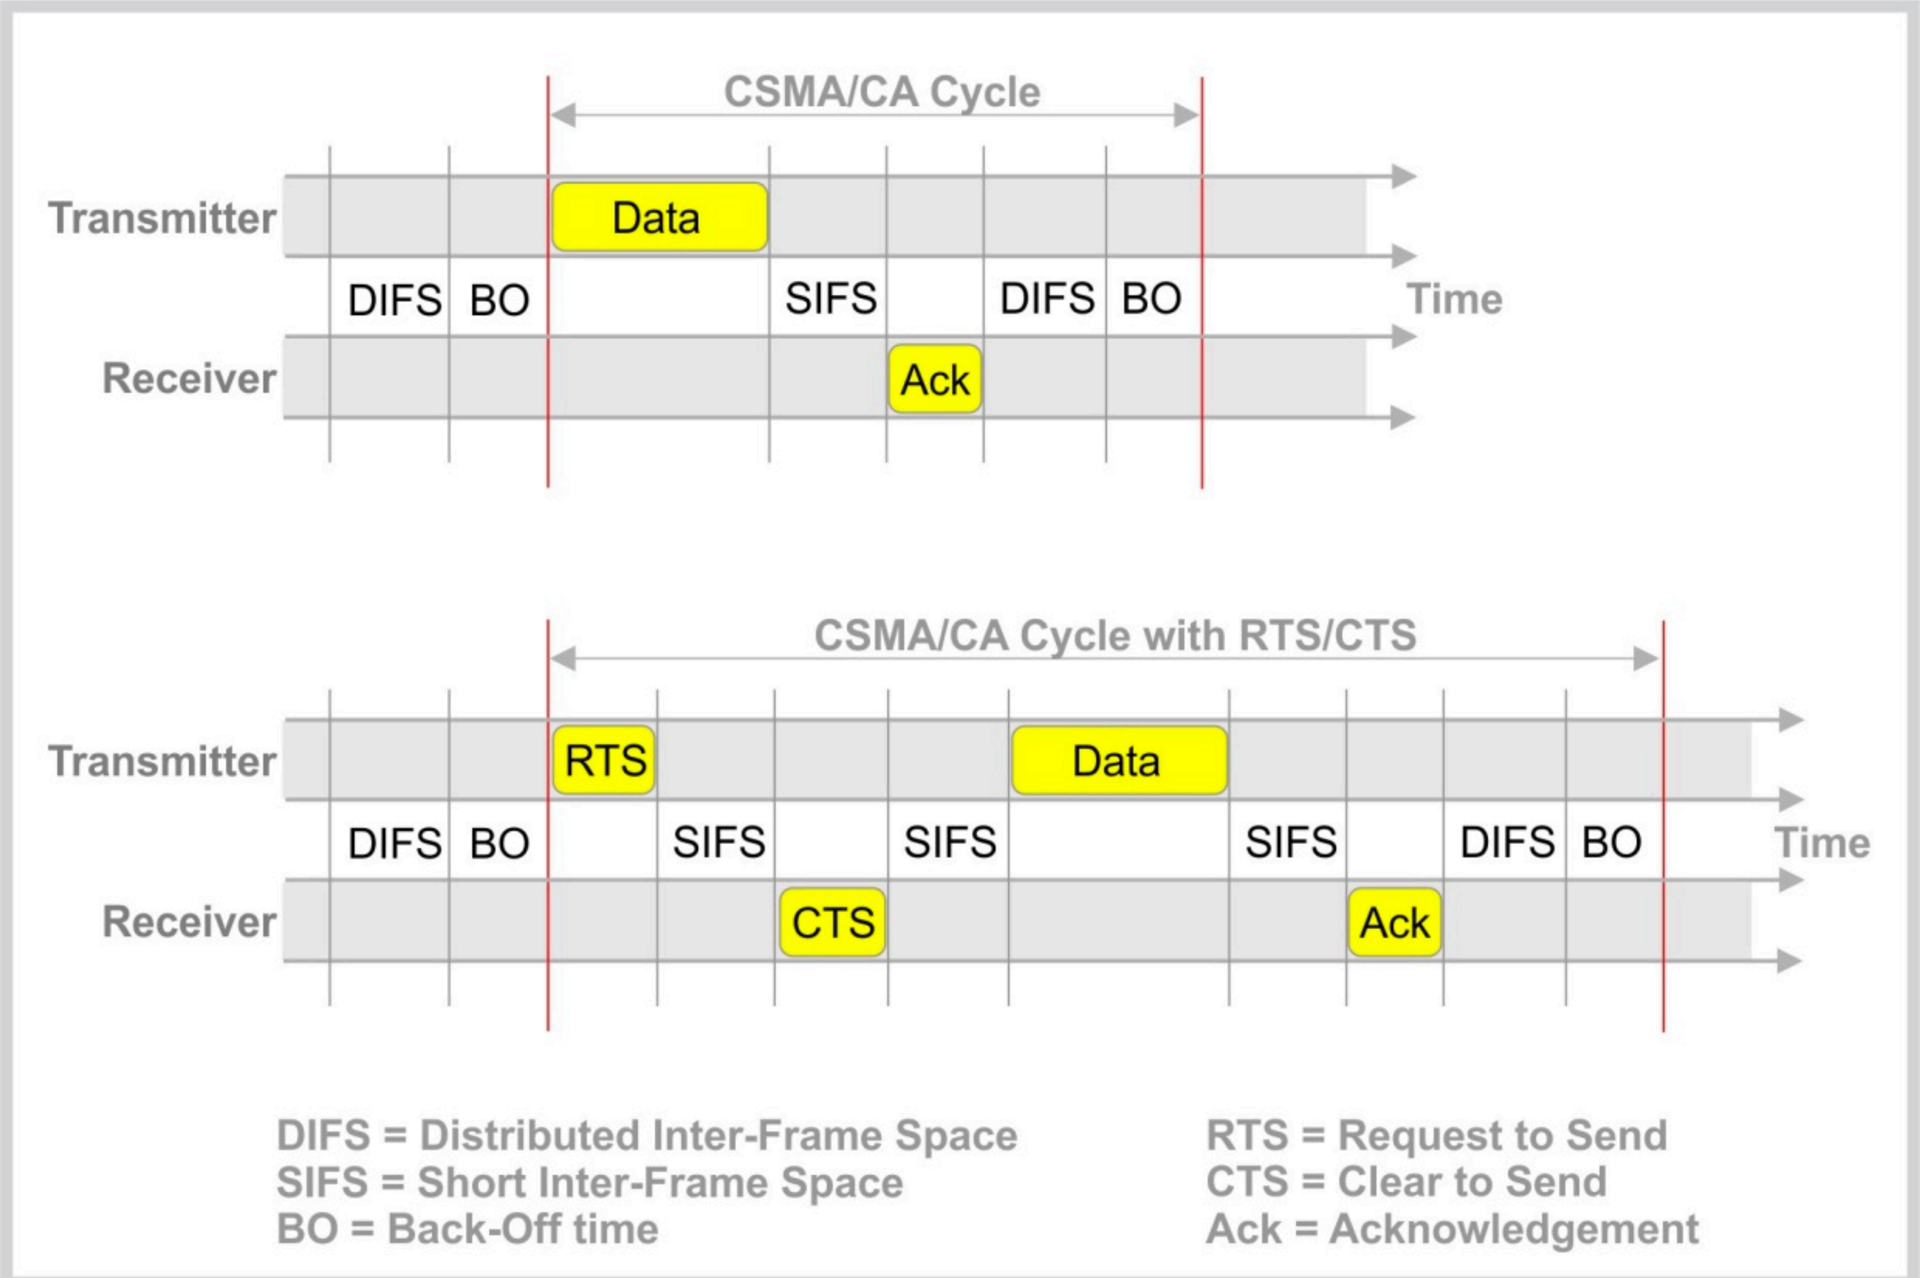

Inter-frame spacing

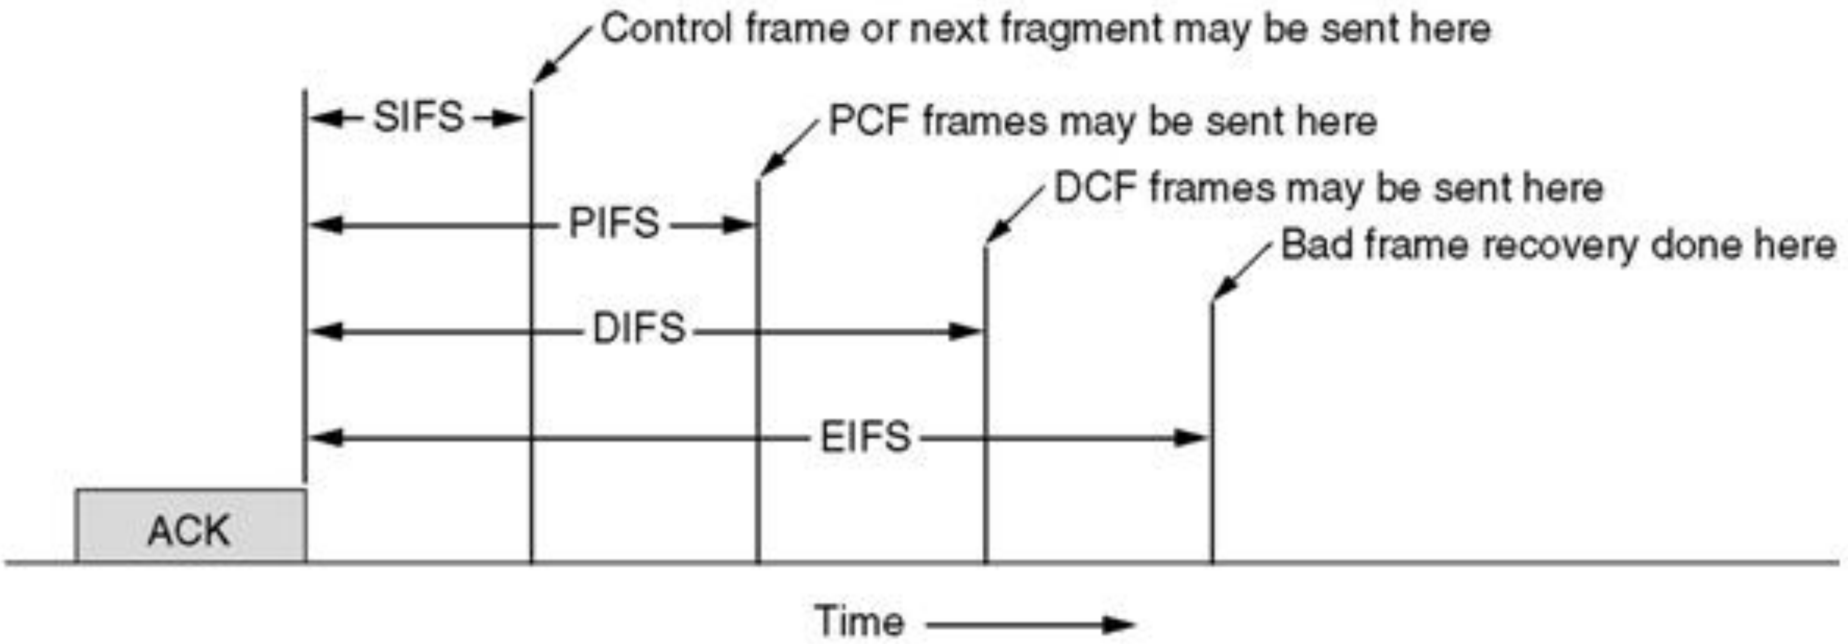

802.11 channels

- 11 channels in 2.4 GHz (in the US)

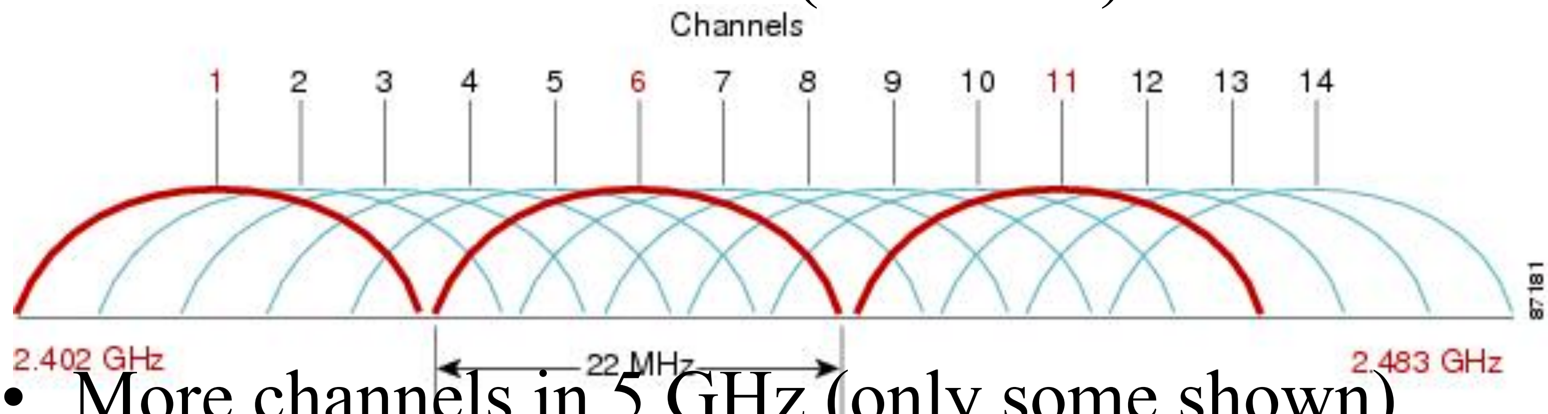

- More channels in 5 GHz (only some shown)

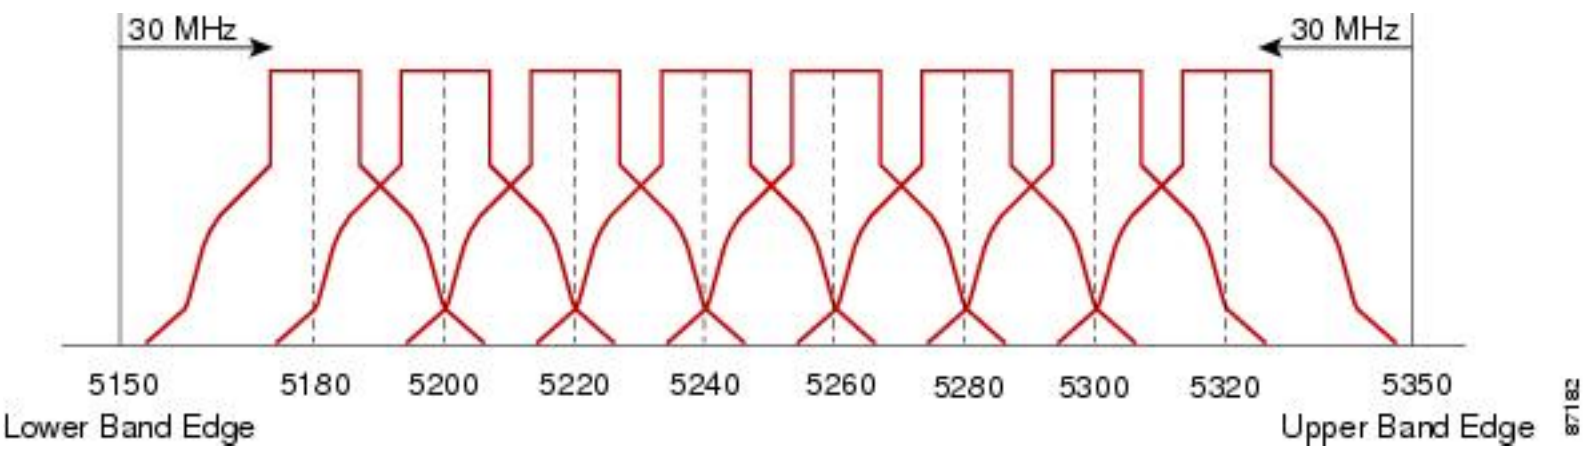

802.11 PHY rates

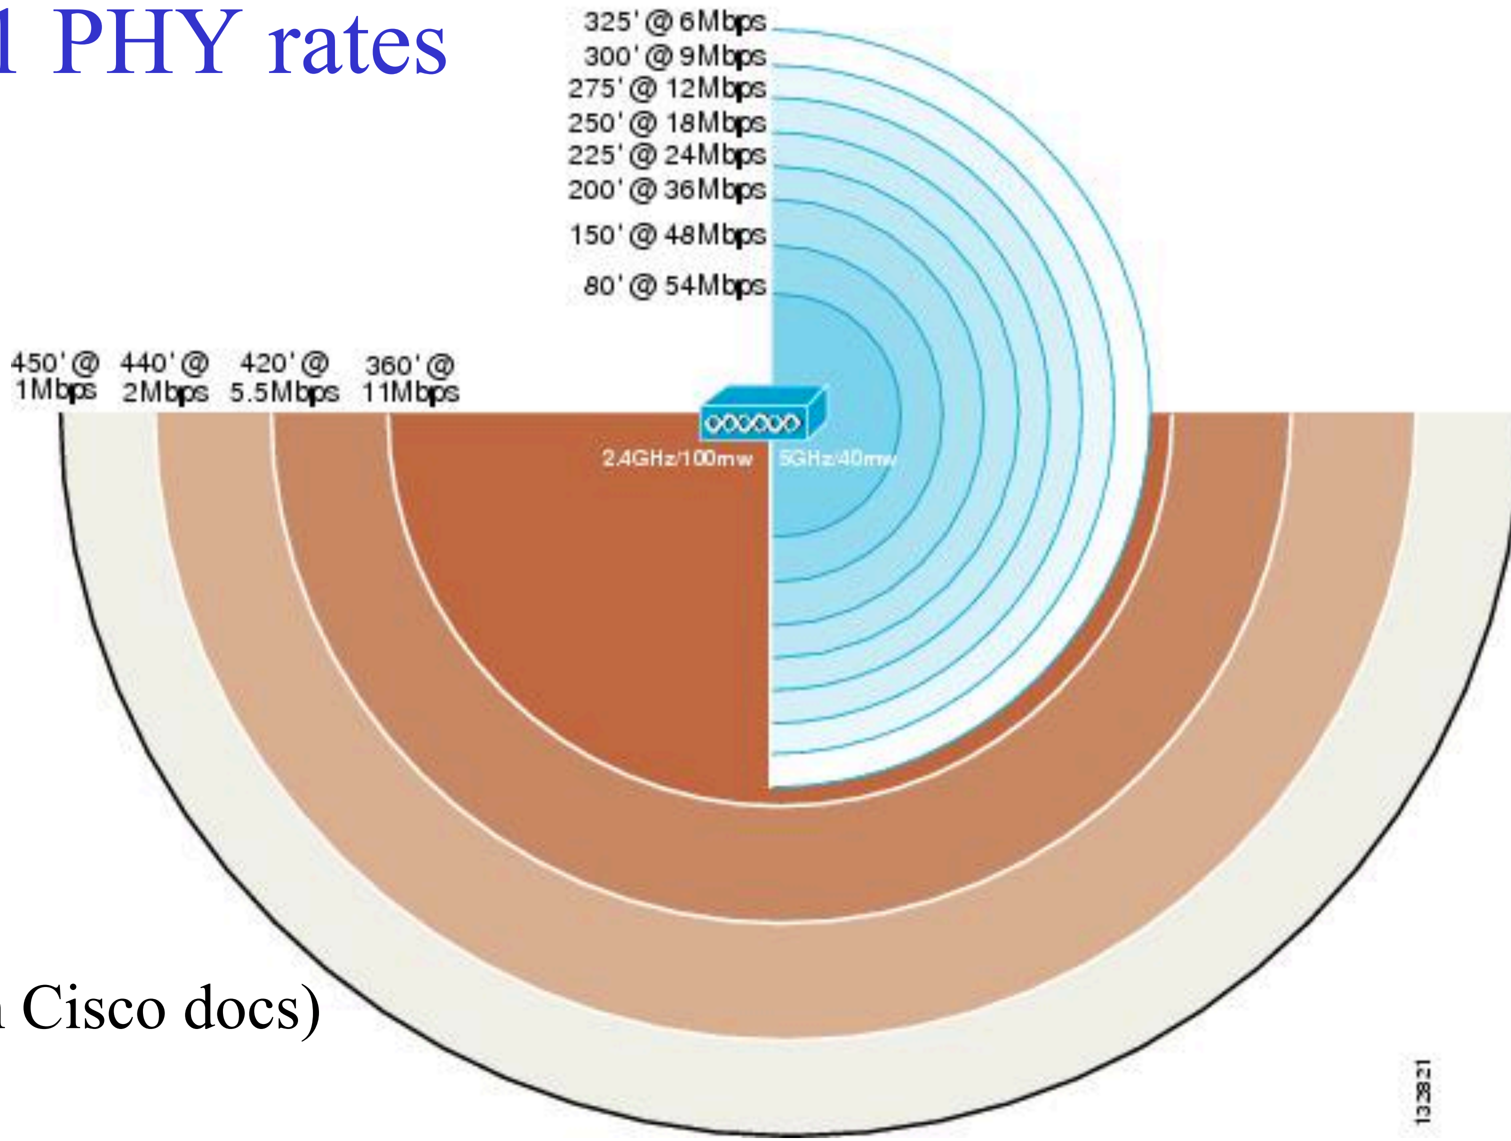

(From Cisco docs)

- PHY rate depends on SNR (which often depends on distance from transmitter)

How to cover a space

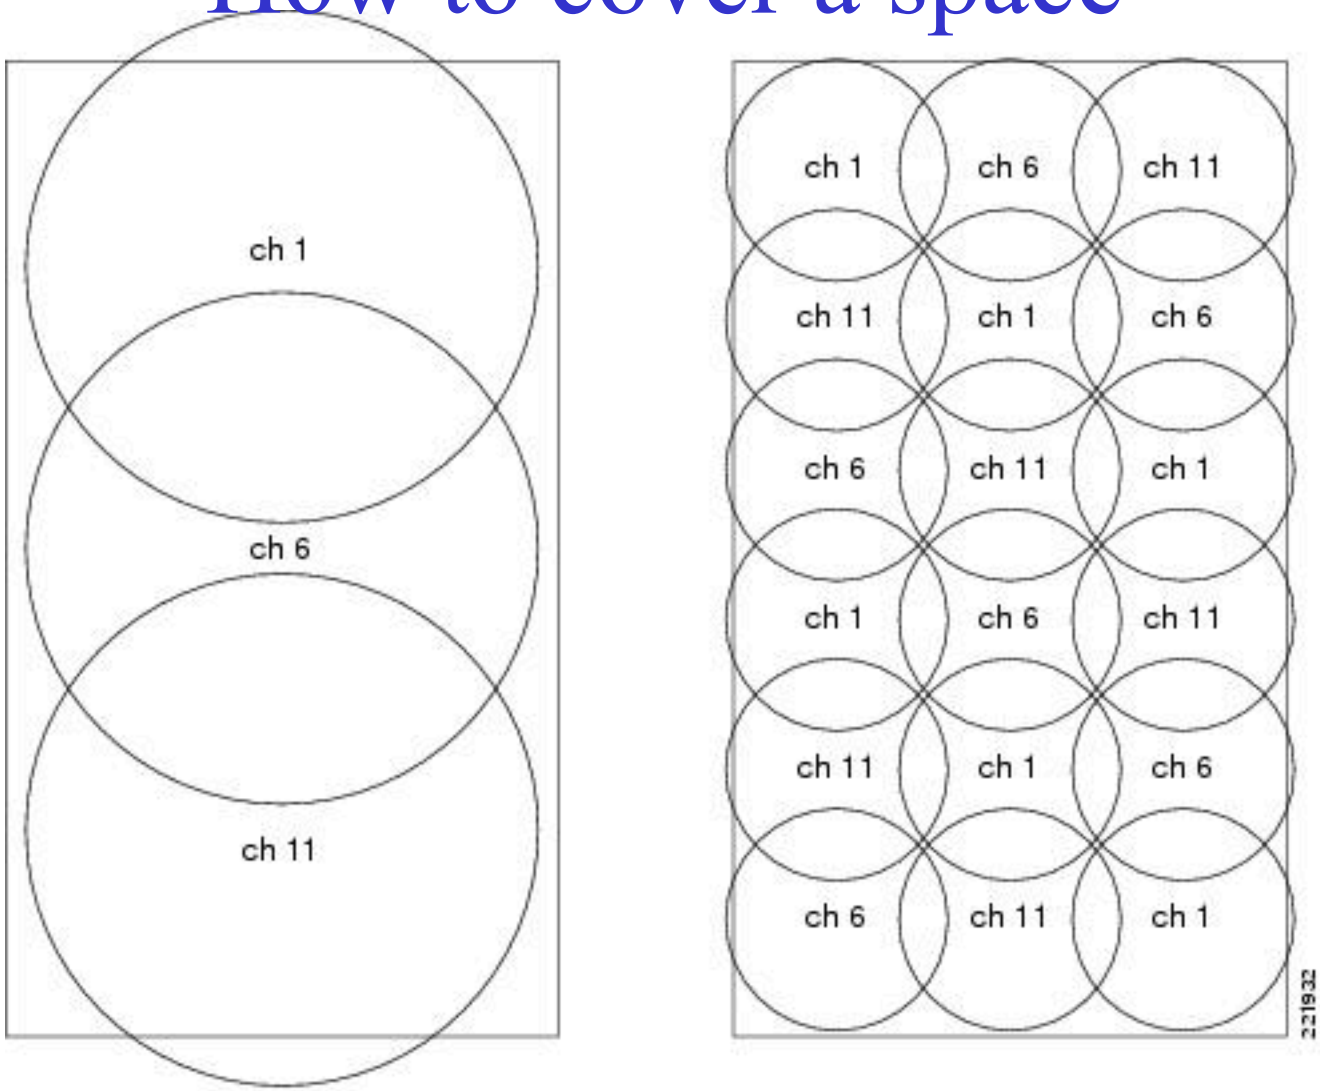

180 Users per floor
30 mW transmitter power
3 Accss Points
60 users per AP
11 Mbps data rate

180 Users per floor
5 mW transmitter power
18 Accss Points
10 users per AP
11 Mbps data rate

Cellular vs WiFi

Cost: Expensive licensed spectrum

Range: 1 to 20 km

Tx power: 1-10 W

Protocols: Highly coordinated

Throughputs: ~ 10 Kbps – 2 Mbps

Cost: Cheap and unlicensed use

Range: ~100 m

Tx power: ~0.1 W

Protocols: Uncoordinated

Throughputs: ~ 1 Mbps – 100 Mbps

Cellular network planning

patterns

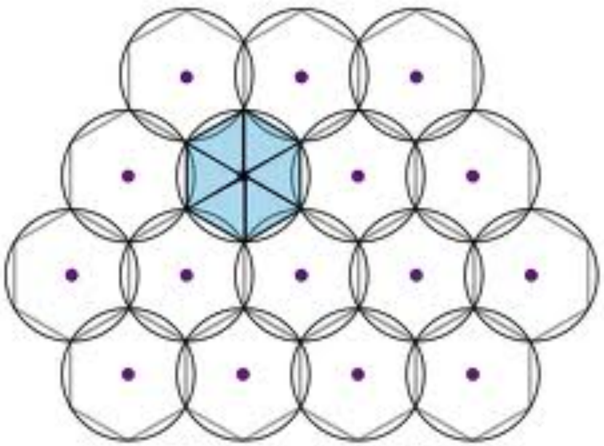

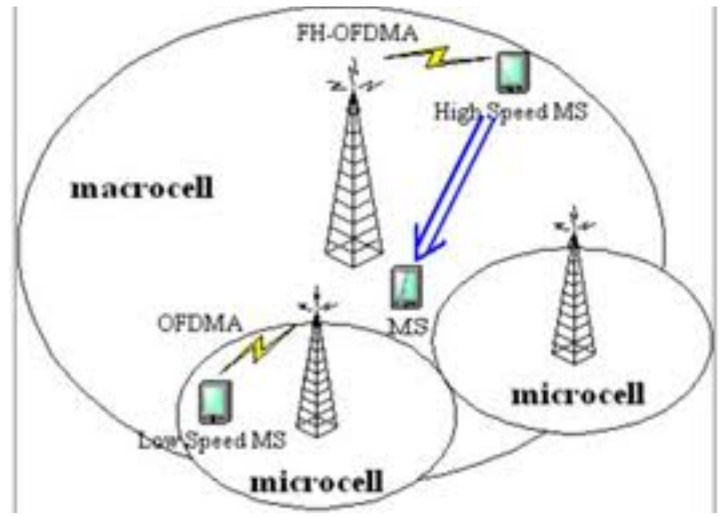

» In real life – Macro and Microcells

Spatial re-use

- Macrocell vs Microcell vs Pico/Femtocells