Wireless LANs

ITS 413 – Internet Technologies and Applications

Contents

- Wireless Communications
  - Characteristics and Challenges
  - Wireless Transmission
- IEEE 802.11 Wireless LANs
  - Details of how wireless LANs work
    - Physical layer characteristics
    - MAC layer protocol
    - Important design considerations

Wireless Communications

- Benefits
  - Untethered communications (no wires)
    - In some cases, can enable quick installation
  - Mobility of users and devices
- Challenges
  - Wireless channel is not as robust as wires
    - More errors, therefore more losses and retransmissions, less throughput
    - Higher delays, therefore must wait long time for retransmissions, less throughput
    - Varying conditions due to mobility and environment
      - Example: timeout based retransmissions can lead to poor performance
  - Radio spectrum is limited (cannot just add more wires)
    - Therefore must efficiently “share” the spectrum amongst all users
  - Many Internet protocols designed assuming a “perfect link”
    - For examples, sometimes TCP may perform poorly over wireless link
  - Physical security is difficult (e.g. cannot easily limit the transmissions to a building)
    - Hence, extra network security is needed

Wireless Transmission

- A simple model of wireless transmission:

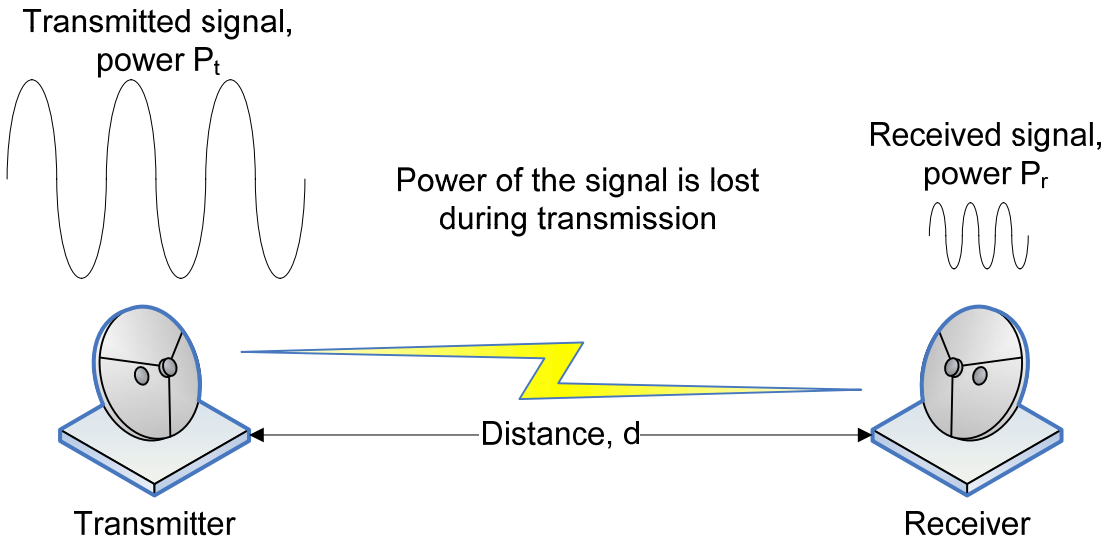

  - The amount of power lost between transmitter and receiver depends on:
    - Distance, frequency, size of antenna, directionality of antenna, obstructions
  - The encoding of bits (0's and 1's) into an analog signal, and decoding at receiver, determines the data rate that can be used it particular environment
  - A receiver can only successfully decode ("understand") a signal received above a certain power level

Wireless Transmission

- An even simpler model of wireless transmission:

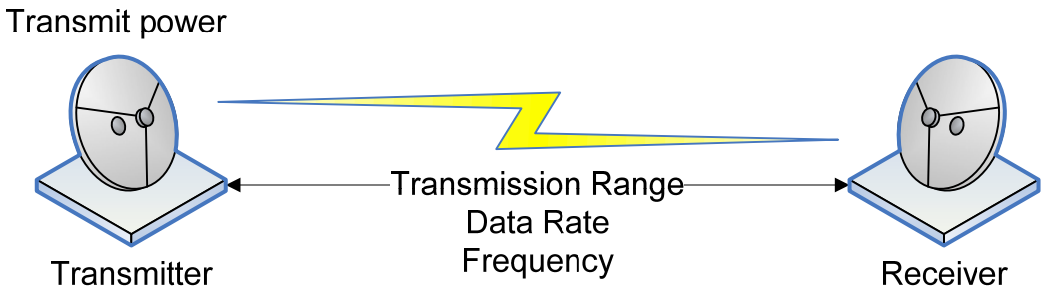

- As IT professionals, we are interested in:
  - Data Rate: how fast can we send the data?
  - Transmission Range: how far can we send the data?
  - Frequency: is it free or licensed? Who else may interfere?
  - Transmit power: how much battery of our wireless device will it use?
  - (and of course, cost: different technologies will have different costs)

Spectrum, Frequency and Bandwidth

- A signal is sent at some frequency *f* with bandwidth *b*
  - The set of all frequencies available is called the *spectrum*
- Why is the frequency (and bandwidth) important?
  - Data rate
    - A higher bandwidth (and frequency) generally leads to higher data rate
  - Transmission range
    - Higher frequency leads to shorter range
    - Different frequency signals are affected by obstacles in different ways
      - E.g. some frequencies are affected by rain, some frequencies will pass through walls, others wont, …
  - Interference
    - If other people/technologies use the same frequency, they may interfere, causing lower data rates
      - E.g. some cordless home phones may interfere with wireless LAN
  - Cost
    - The spectrum is limited and managed by national/international organisations
    - Some frequencies are free to use by anybody (within some rules)
      - E.g. most wireless LANs operate at the free Industrial Scientific Medical (ISM) frequency
    - Other frequencies you need a license to use
      - The license may be expensive, e.g. companies in Germany spent 2 trillion Baht (2,000,000,000,000) on licenses to use spectrum for 3G mobile networks

Spectrum, Frequency and Bandwidth

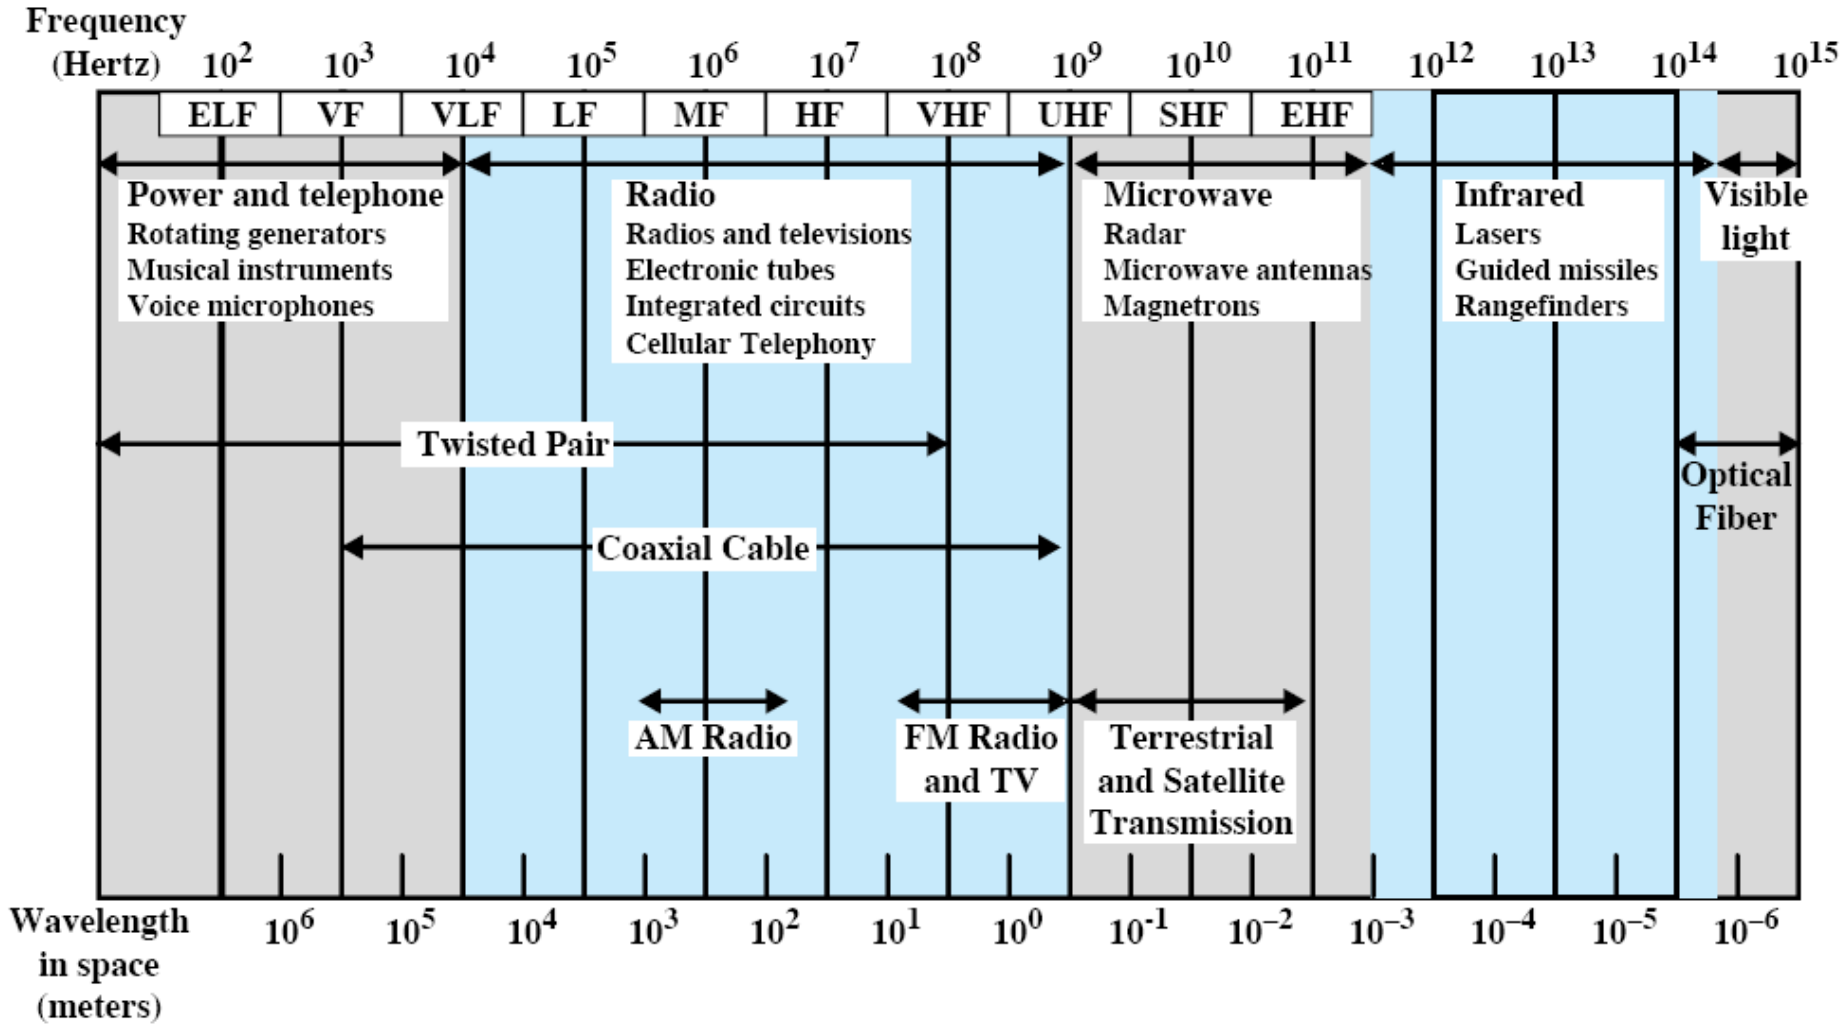

Transmission Topology

- Point-to-point
  - Transmit antenna points at receive antenna: directional
  - Signal power is concentrated between transmitter and receiver

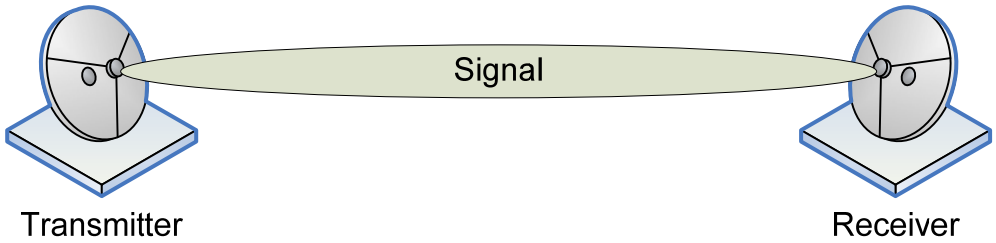

- Broadcast Radio
  - Transmitter sends signal in every direction: omni-directional
  - Anyone “within range” can receive the signal

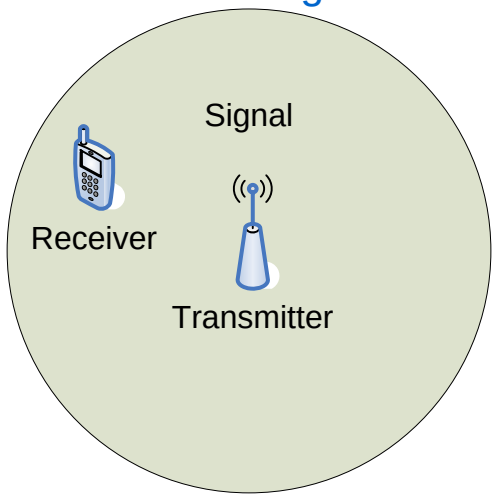

Characteristics of Broadcast Radio

- Some assumptions about broadcast transmissions:
  - A radio device's transmission can be received by all other devices within the *transmission range*
    - For example, if the transmission range is 30m, then all nodes within 30m of the transmitter will receive the transmission. A node that is 31m (or more) from the transmitter will not receive the transmission (its not that simple in practice though)
    - The transmission goes in all directions: *omni-directional*
  - A radio device cannot transmit and receive at the same time
    - Due to interference, it is difficult to implement a transmitter/receiver (transceiver) that operate at same time
    - Therefore, we have *half-duplex* operation
    - (Whereas most wired networks today, such as Ethernet, allow transmission and reception at same time: full-duplex)
  - A radio device cannot successfully receive transmissions from two or more sources at the same time
    - The two transmissions interfere with each other, making it almost impossible for the receiver to determine what the two original signals were
    - So we assume that if two (or more) devices transmit at the same time, then a receive cannot successfully receive either of them. Often called a *collision* between packets

Classifying Wireless Networks

Name	Abbrev.	Range	Applications	Examples
Wireless Personal Area Network	WPAN	Several metres	Connecting peripherals	Bluetooth, IrDA
Wireless Local Area Network	WLAN	10's to 1000's of metres	Office, home, street communications	IEEE 802.11
Wireless Metropolitan Area Network	WMAN	Km's	Inter-office and building connections,	IEEE 802.16, IEEE 802.20
Wireless Wide Area Network	WWAN	Km's to regional to global	City, nation wide telecommunications; connections between cities	3G(UMTS),GSM, Satellite

IEEE 802.11 Wireless LANs

A Quick Intro to Wireless LANs

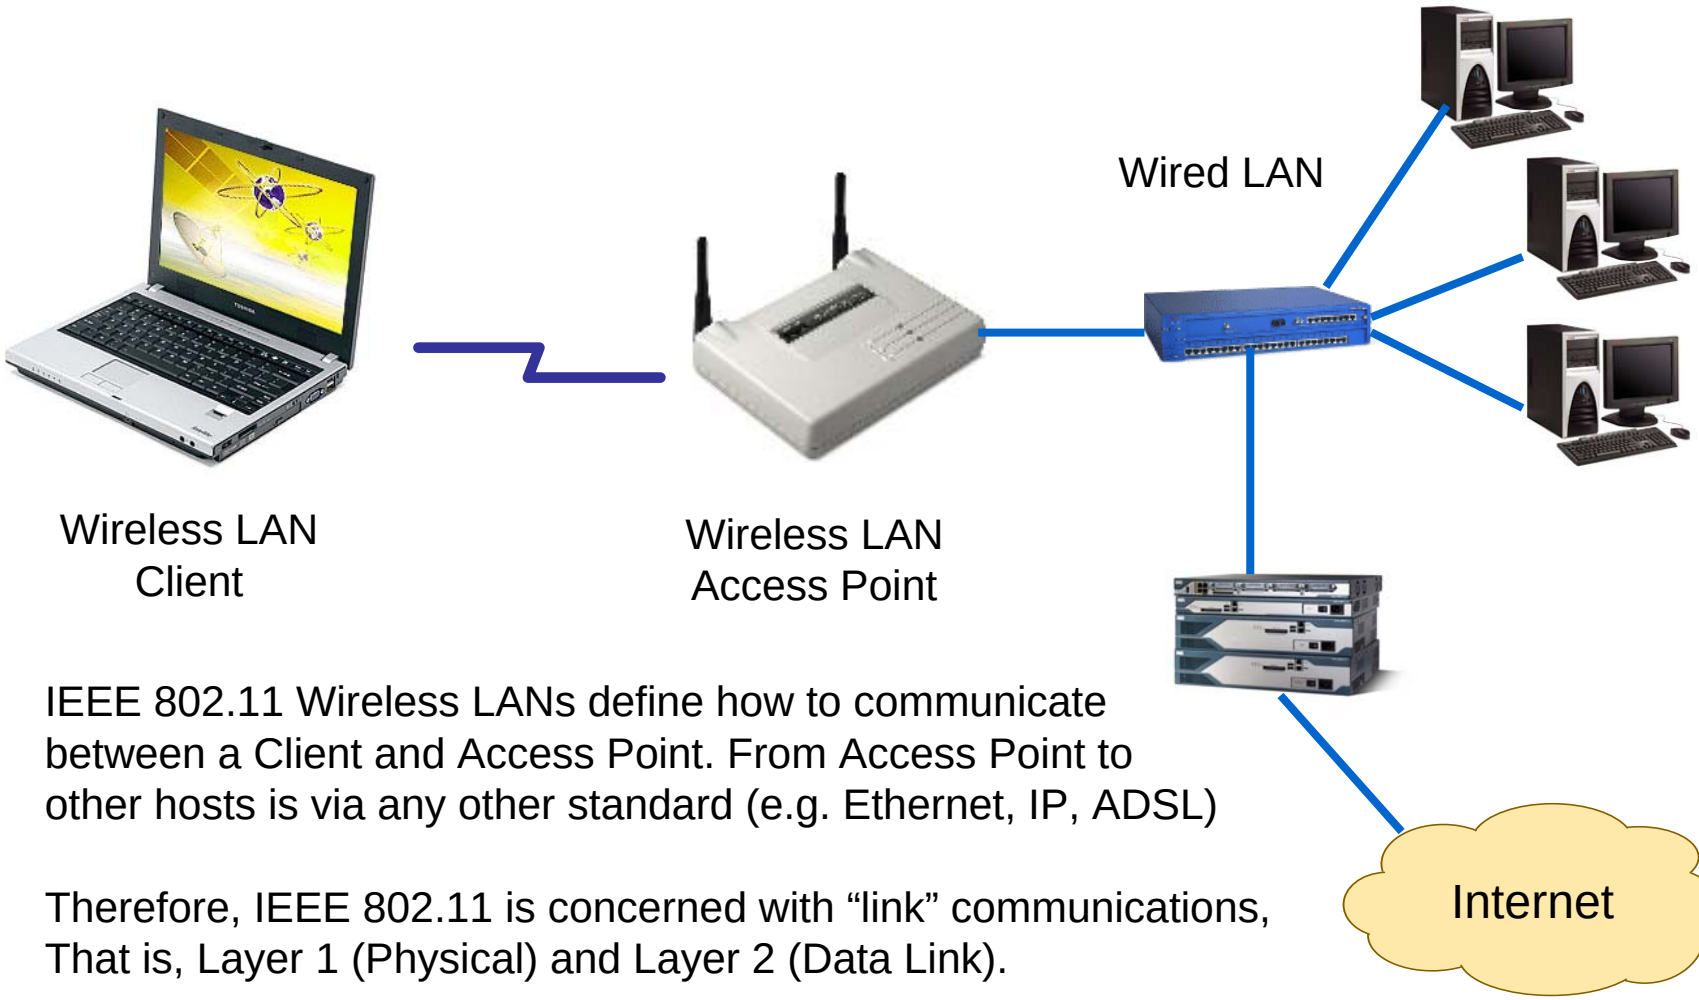

IEEE 802.11 Wireless LANs define how to communicate between a Client and Access Point. From Access Point to other hosts is via any other standard (e.g. Ethernet, IP, ADSL)

Therefore, IEEE 802.11 is concerned with “link” communications, That is, Layer 1 (Physical) and Layer 2 (Data Link).

A Quick Intro to Wireless LANs

- How fast?
  - Most common standards are: 11Mb/s and 54Mb/s
  - Throughput as around half (5 to 25Mb/s), but shared!
- How far?
  - 10’s of metres indoor, 100’s of metres outdoor
- How much power?
  - 1mW up to 100mW
- What frequency?
  - Most common is 2.4GHz, an unlicensed band
- Are they secure?
  - Satisfactory link level security is provided by WPA

Terminology and Concepts

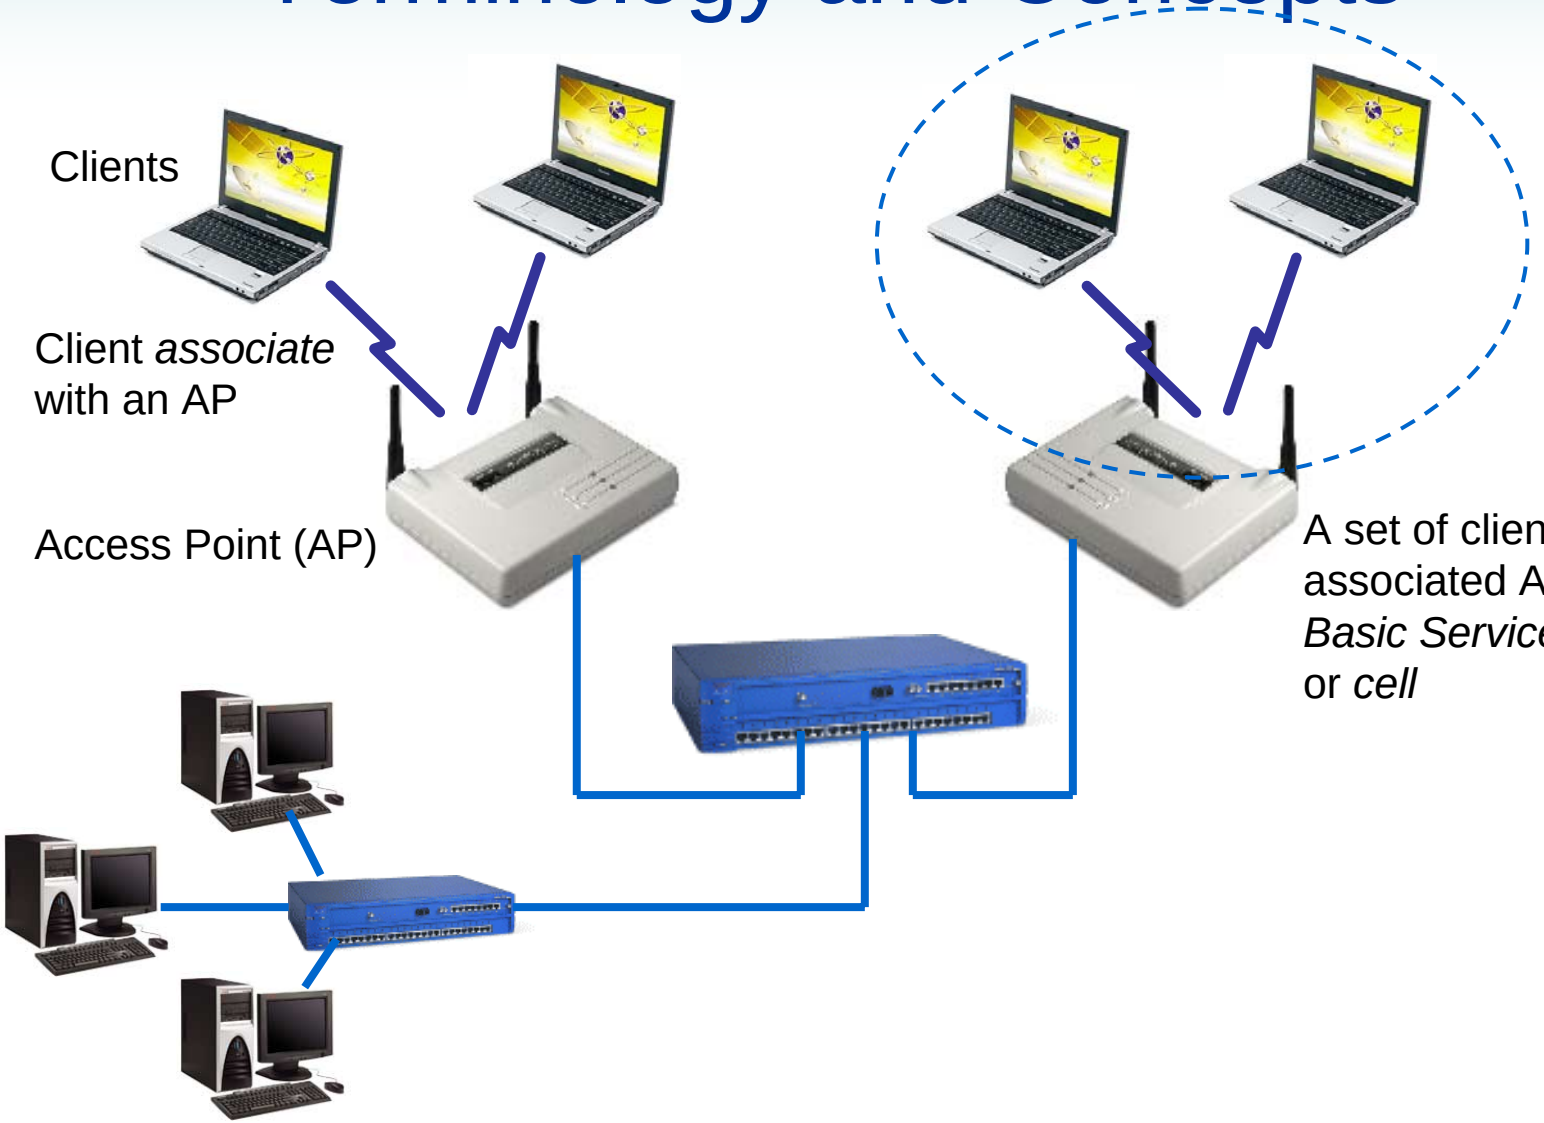

Terminology and Concepts

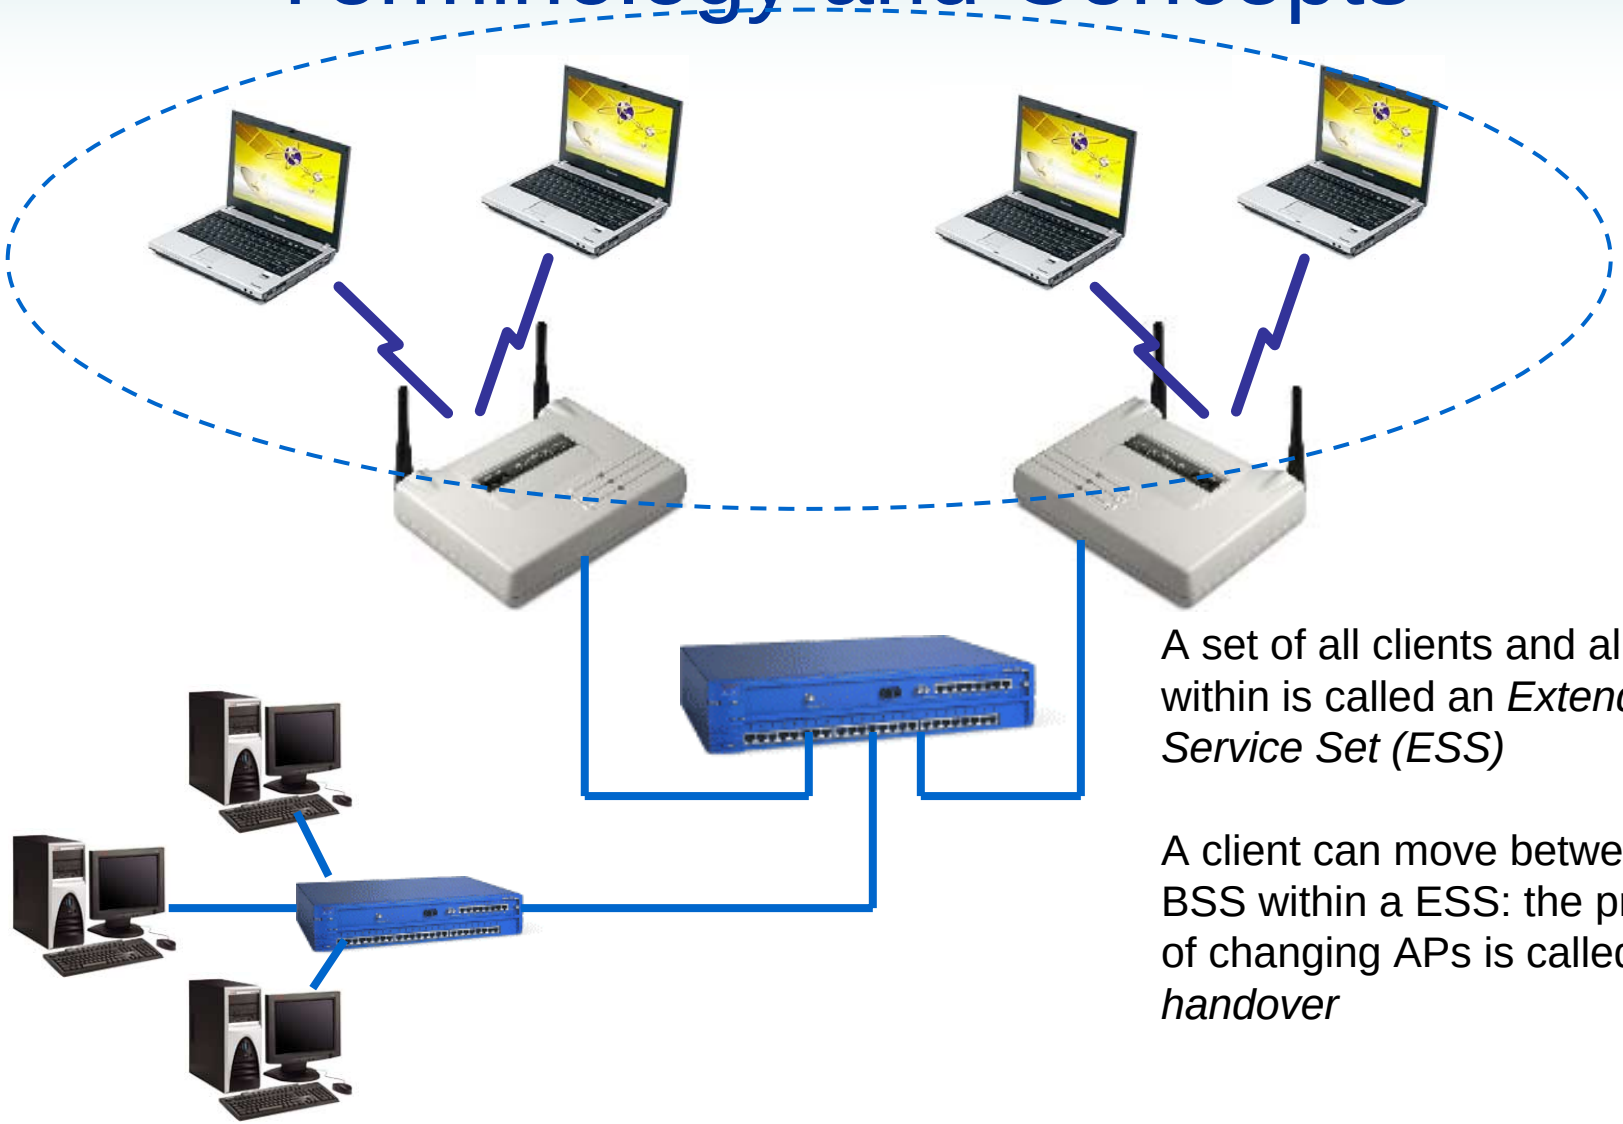

A set of all clients and all APs within is called an *Extended Service Set (ESS)*

A client can move between BSS within a ESS: the process of changing APs is called *handover*

IEEE 802.11 and TCP/IP Stack

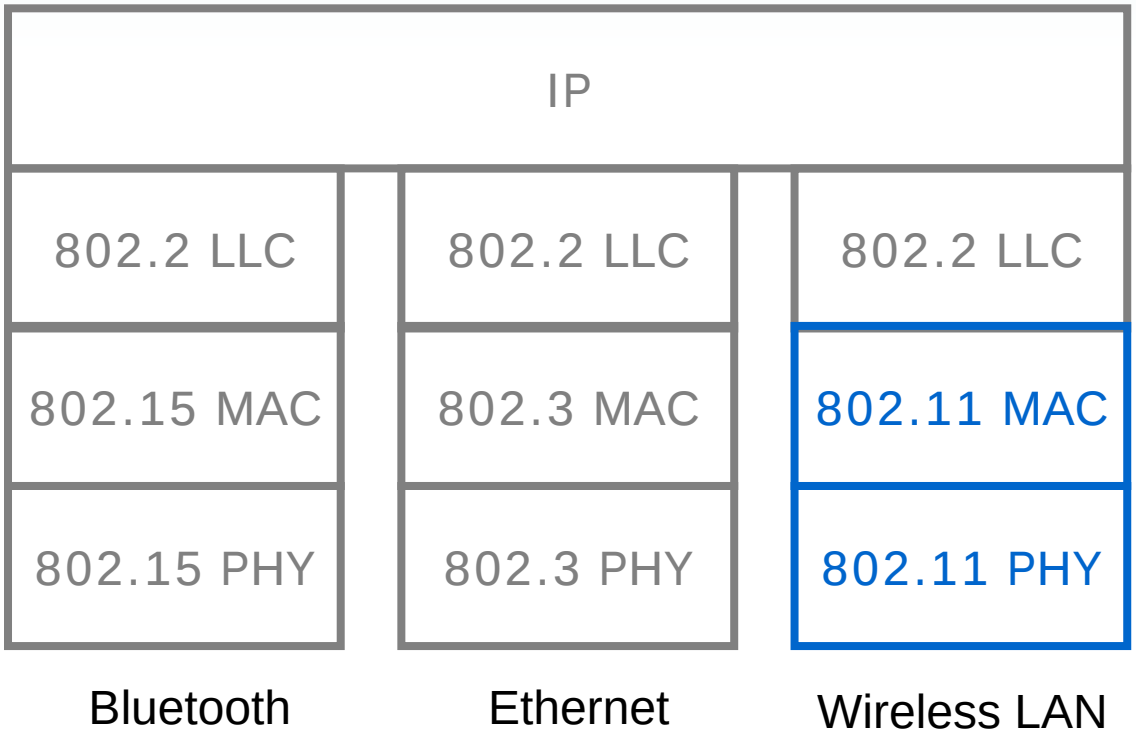

3. Network Layer

2. Data Link Layer

1. Physical Layer

This shows an example set of stacks for a single laptop that supports wireless LAN, Ethernet and Bluetooth. Note that 802.2 is a common protocol for all 802 series LAN protocols. The MAC and PHY layers are specific to each LAN protocol

LLC = Logical Link Control
MAC = Medium Access Control
PHY = Physical

IEEE 802.11 Standards

- IEEE 802.11 Working Group started in 1990
  - Began developing a standard for wireless LANs – finished the first standard in 1997!
  - Enhancements and new features are continually added
    - These are identified by the Task Groups who develop them
    - Task Groups b, g, n – developed new Physical layers
    - Task Groups e, i, h – developed enhancements for quality of service, security, mobility
  - Summary:
    - The basic operations of the MAC layer is the same
    - New (faster) Physical layer standards have been developed
    - Optional features for security and quality of service are available

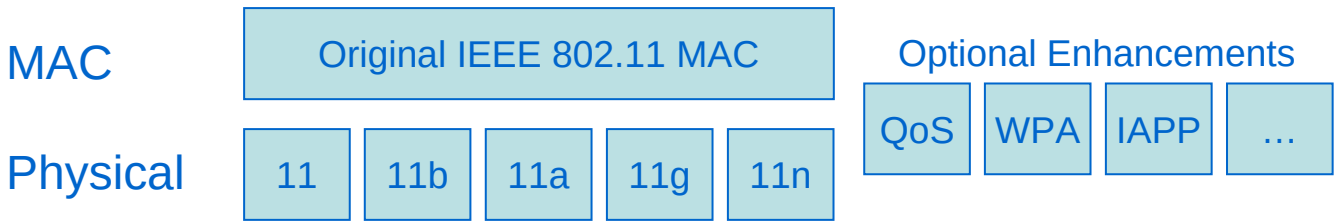

IEEE 802.11 Physical Layer

- Remember: the Physical layer defines how to send 0's and 1's as a physical signal (e.g. radio waves)
- What does the IEEE 802.11 Physical layer define:
  - Modulation: "shape" the analog signal into "efficient" form
  - Frequency: the part of frequency and bandwidth the signal is sent at
  - Timing: how to synchronise sender/receiver
- What are the practical characteristics of different Physical layers
  - Data rate: the speed at which data is sent [bits per second]
  - Transmission range: the maximum distance that two nodes can communicate [metres]
    - Depends on transmit power and receiver sensitivity
  - Frequency: determines how transmissions travel in an environment, and potential with interference with other sources [Hertz]
  - Non-overlapping channels: the number of channels (frequencies) that can be used at the same time without causing interference

IEEE 802.11 Physical Layers

	11	11b	11a	11g	11n
Released	1997	1999	1999	2003	2007
Frequency	2.4GHz	2.4GHz	5GHz	2.4GHZ	5GHz
Modulation	DSSS	DSSS	OFDM	OFDM	OFDM, MIMO
Channels	3	3	8	3	4-8
Data Rate	2Mb/s	11Mb/s	54Mb/s	54Mb/s	300Mb/s
Range	20-300m	20-300m	15-30m	25-75m	20-60m

IEEE 802.11 MAC

IEEE 802.11 MAC Layer

- IEEE 802.11 Physical layer provides means for sending data over wireless medium
- IEEE 802.11 MAC layer:
  - Defines management procedures for discovering, joining and leaving BSS/ESS
    - How does your laptop find an AP? How does it connect to the AP?
  - Defines protocol for efficient and robust communication over wireless medium
    - How does your laptop share access with other surrounding laptops?
  - Is common across different physical layers
    - Same MAC used for 11a, 11b, 11g (although some parameter values change)
  - Uses Hardware (MAC) addresses to identify clients/APs
    - The format of Hardware addresses is same as Ethernet: 48-bit addresses

IEEE 802.11 MAC Management

- When you turn on your laptop, it needs to find an AP to associate with. How does it find it? Two approaches are used:
  - APs periodically broadcast *Beacon* frames advertising itself
    - Any client that receives the Beacon frame will know that the AP exists (and can choose to associate with the AP)
  - A client can broadcast a *Probe Request* frame, searching for an AP
    - Any AP that receives the Probe Request frame may respond with a *Probe Response* frame
    - The client, upon receiving the Probe Response, can choose to associate with the AP
- Once your laptop knows about an AP, how does it associate?
  - Authentication
    - Allows AP to check whether client is allowed to connect to network
  - Association
    - Registers the client with the AP
- Once associated, a client and AP can exchange data. When the client shuts down or moves, what happens?
  - Either client or AP can de-authenticate or disassociate
  - If moving from one AP to another AP, a special *re-association* can occur
    - There are ways to make the *handover* fast

MAC Management

Client

AP

Periodic broadcast by AP;
Probe request/response:
active discovery by client

Discovery

Beacon

Joining

Authenticate Req.

Authenticate Resp.

Re-association possible.
Moves association from
AP to AP

Associate Request

Associate Response

Either client or AP can
disassociate and/or de-
authenticate.

Disassociate

Leaving

Deauthenticate

MAC Management Frames

- Discovery Frames
  - Beacon: periodic broadcast by AP (e.g. 10/sec)
  - Probe Request: active discovery by client
  - Probe Response: APs response to Probe Request
- Joining Frames
  - Authenticate Request (client) / Response (AP)
  - Associate Request (client) / Response (AP)
  - Reassociate Request (client) / Response (AP)
- Leaving Frames
  - De-authenticate (client or AP)
  - Disassociate (client or AP)
  - (these are notifications, not requests; no responses needed)

IEEE 802.11 MAC Data Transfer

- The aim of MAC is to provide efficient, robust and fair data transfer
  - Efficient: minimises overheads in the bandwidth scarce wireless medium
  - Robust: can handle errors
  - Fair: all stations get an “equal share” of the wireless medium
- Remember: operate in a broadcast radio environment. If two nodes transmit at the same time, then collision (lost data). Therefore aim of MAC is:
  - Only one node transmits at a time

- There were two approaches for transferring data in the original IEEE 802.11 standard: DCF and PCF
  - Today, DCF is most commonly used (along with enhancements developed in IEEE 802.11e etc.)

- Distributed Coordination Function (DCF)
  - Carrier Sense Multiple Access (CSMA)
  - Collision Avoidance (CA)
  - DCF is mandatory in 802.11
- Point Coordination Function (PCF)
  - Polling based transfer: the AP indicates a schedule for the clients to transmit

Distributed Coordination Function

- Distributed Coordination:
  - Aim is that only one node transmits at a time
  - Each node implements the function, so that they determine when they are allowed to transmit in a *distributed* manner
- In DCF, clients and APs are the same
  - Lets refer to them as *stations*
- Two modes of operation:
  - Basic Access mode
  - RTS/CTS mode
  - Depending on the size of the Data to transmit, the mode will be chosen
    - For a smaller Data, use Basic Access; for larger Data use RTS/CTS

DCF Frames

- DATA

  - Sent from Sender to Receiver
  - Data (or payload) up to 2312 bytes (typically limited to 1500 bytes, and often varies)
  - Header+Trailer: typically 34 bytes
- ACK

  - Acknowledgement from Receiver to Sender
  - Typically 14 bytes
- RTS

  - Request To Send, from Sender to Receiver
  - Typically 20 bytes
- CTS

  - Clear To Send, from Receiver to Sender
  - Typically 14 bytes

802.11 DATA Frame

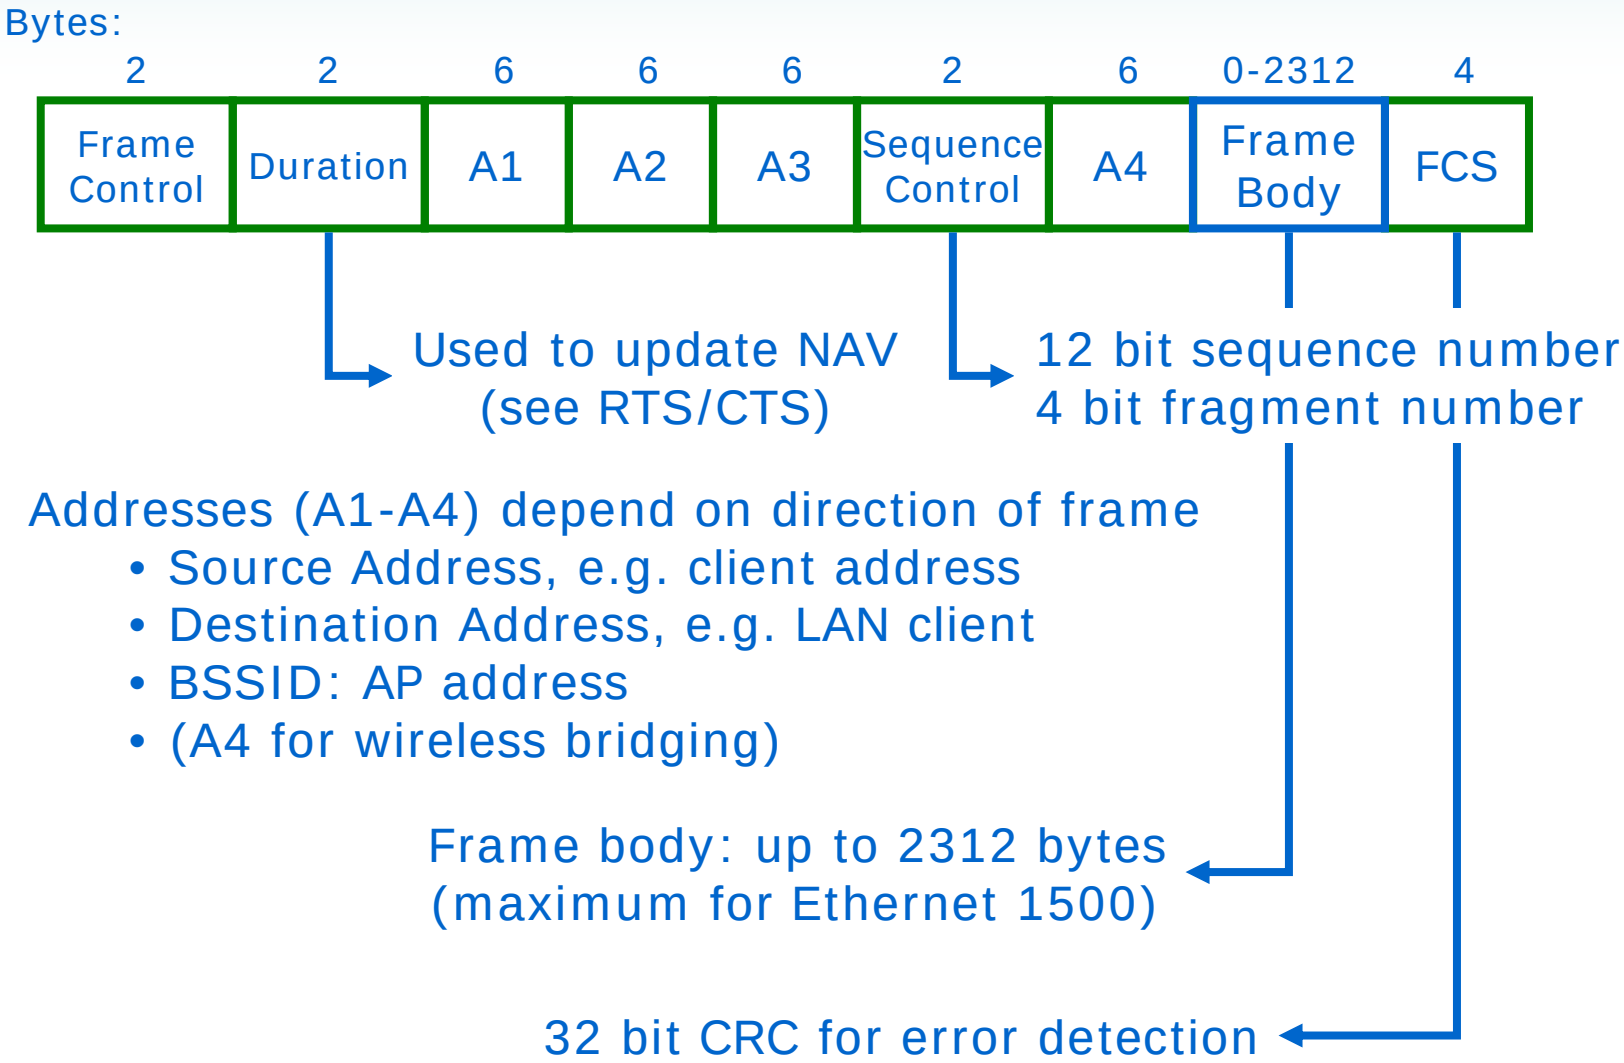

DATA Frame Control field

Field	Bits	Description
Protocol Version	2	Value: 0
Type	2	Control, data, management
Subtype	4	Probe Req., Data, Ack, …
To DS	1	00: Ad hoc; 10: Client to AP; 01: AP to client; 11: AP-AP (bridge)
From DS	1	
More Frag	1	More fragments to follow
Retry	1	Retransmission
Pwr Mgt	1	Power save mode
More Data	1	Power save or CFP
WEP	1	On or Off
Order	1	StrictlyOrdered/OrderableMulticast

Addressing in 802.11

- There are four address fields that can be carried in a frame
  - Each carry IEEE 48-bit Hardware addresses
  - Three of the four addresses are commonly used:
    - Client address: the 802.11 client involved in the transfer
    - AP address: the AP involved in the transfer
    - LAN address: the client on the “other” LAN involved in the transfer

Addressing in 802.11 – To LAN

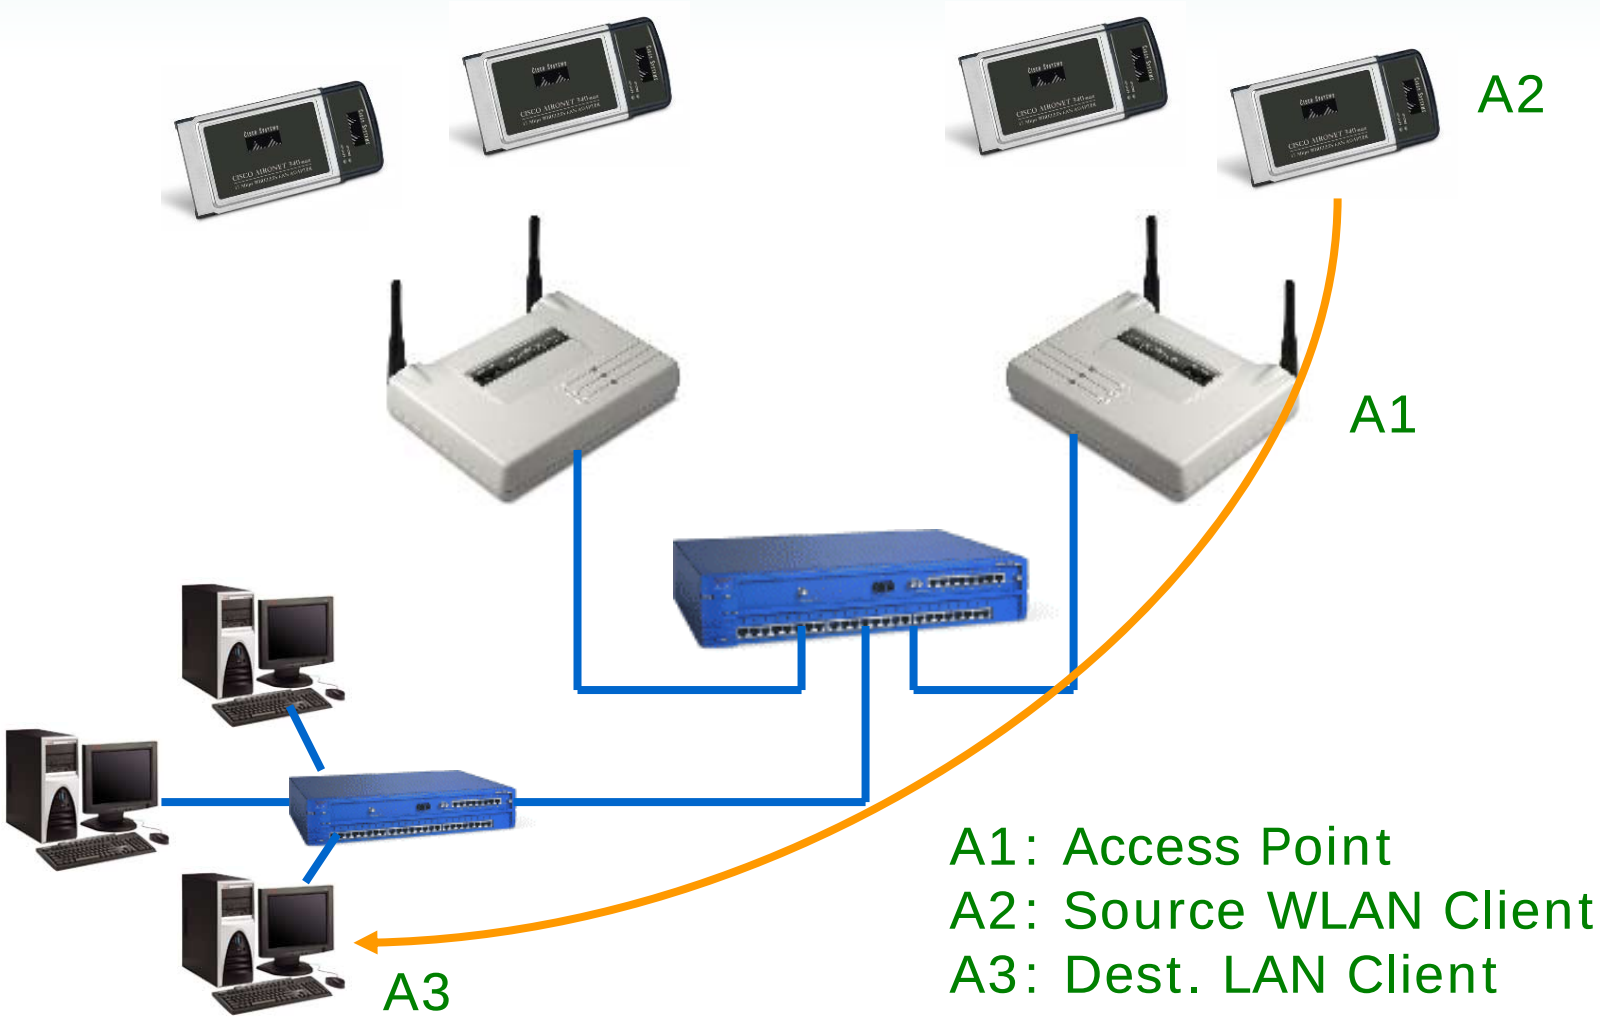

Addressing in 802.11 – From LAN

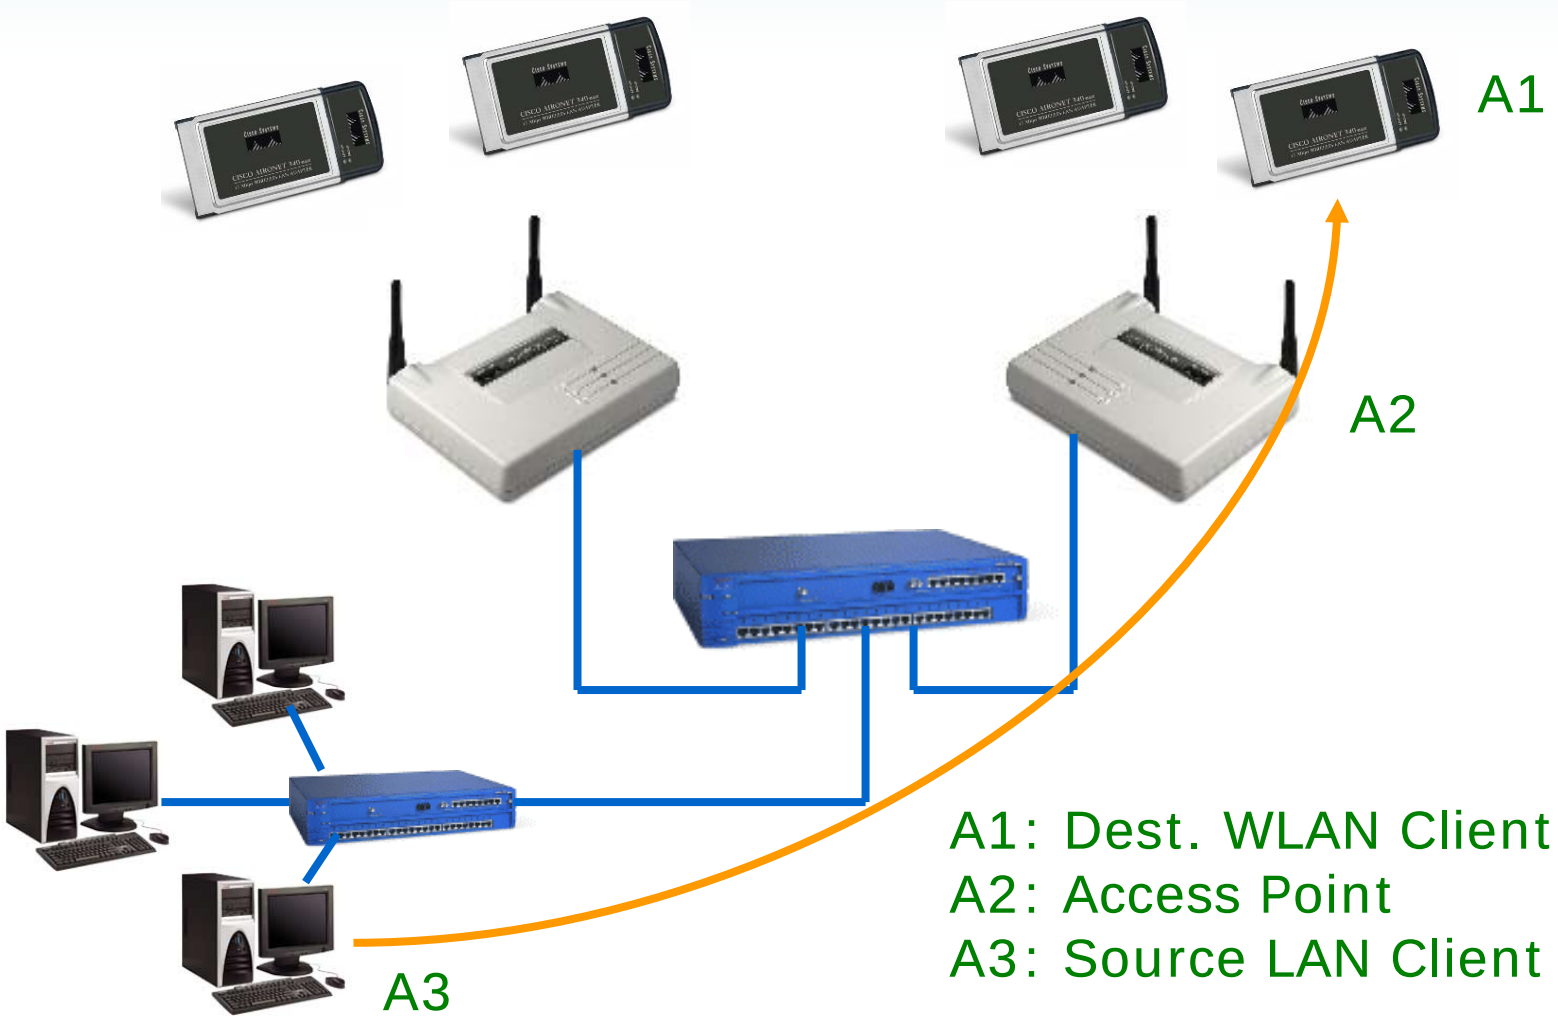

Addressing in 802.11 – To WLAN

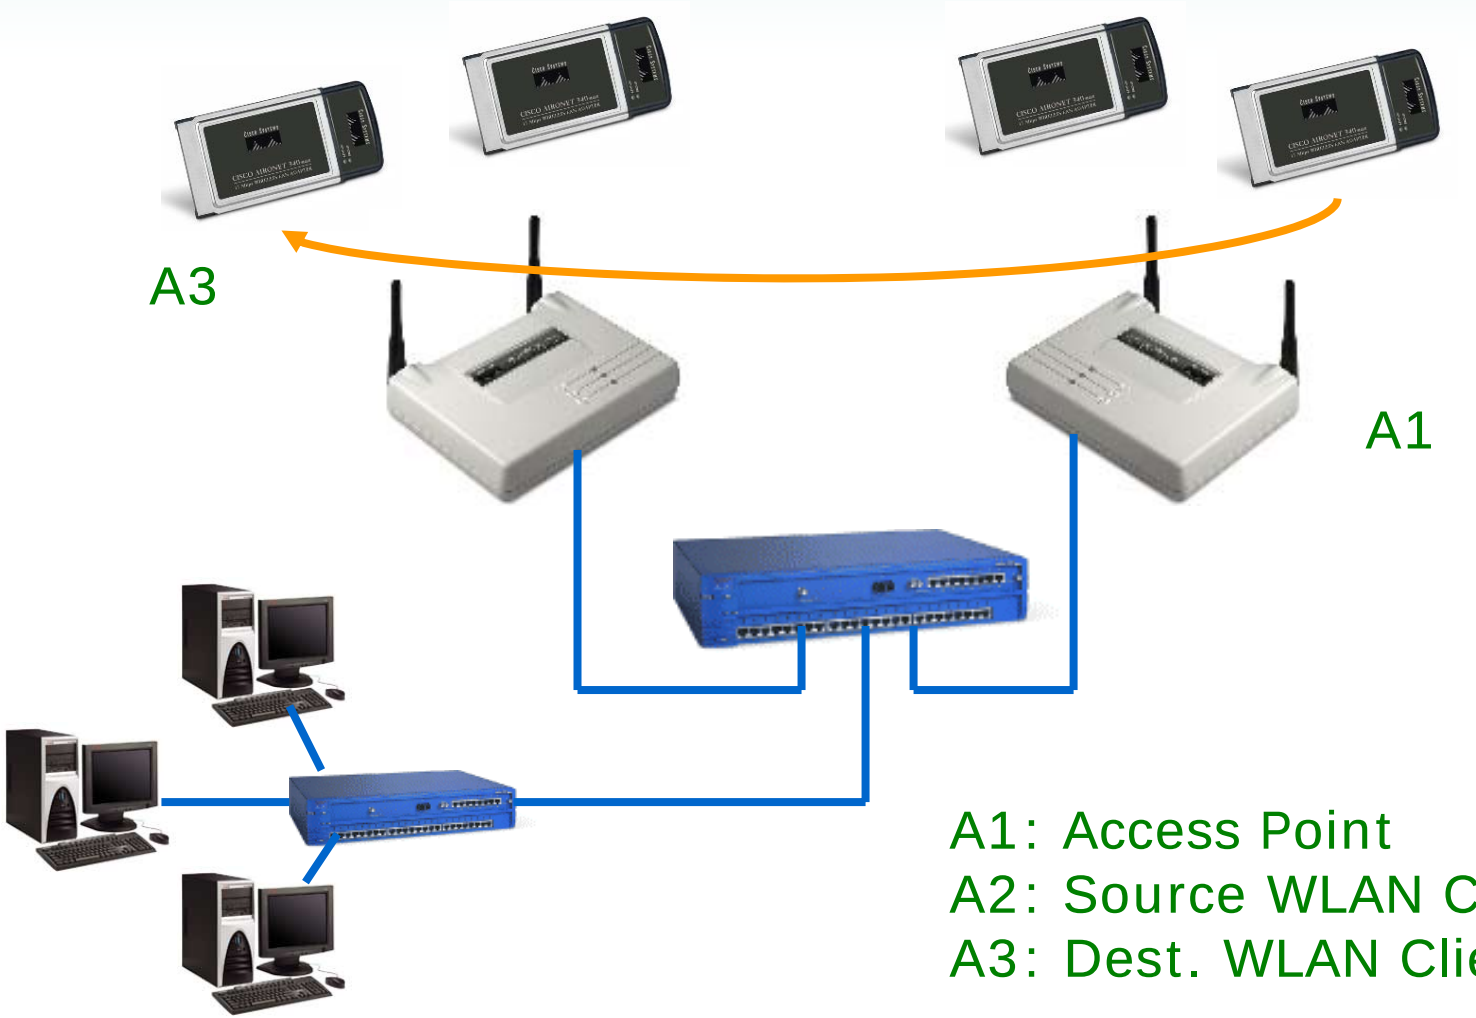

A1: Access Point
A2: Source WLAN Client
A3: Dest. WLAN Client

DCF Basic Access

- Concept:
  - A station will transmit if the medium is idle (no-one else transmitting)
  - If medium is busy (someone else transmitting), then wait until they have finished and then transmit
- Operation:
  - When station has data ready to send:
    1. Medium must be idle for period of DCF Inter Frame Space (DIFS)
    2. After DIFS, medium must be idle for Backoff period
    3. When backoff complete, transmit DATA frame
    4. Upon receipt of ACK frame, data transfer is complete
  - If medium becomes busy during DIFS:
    - Wait until idle, then restart from point 1 above
  - If medium becomes busy during Backoff:
    - Suspend backoff counter, wait until idle, then restart from point 1 above
    - Continue backoff from where it was suspended

DCF Basic Access

- Interframe Spaces
  - DCF IFS (DIFS): period that the medium must be sensed idle before starting backoff
  - Short IFS (SIFS): period to wait between frame transmissions during data transfer
    - E.g. Receiver waits SIFS before sending ACK
    - SIFS is always less than DIFS
- Backoff Period
  - R = random integer between 0 and CW-1
  - Backoff Period = R x SlotTime
  - CW: Contention Window size, initially CWmin
  - Choosing a random Backoff period minimises collisions after two or more stations defer
  - Provides fair access to all nodes (on average, every station gets same chance of winning access)
- Parameter values are defined for each Physical layer
  - DIFS, SIFS, SlotTime, CWmin, CWmax

DCF Basic Access Example

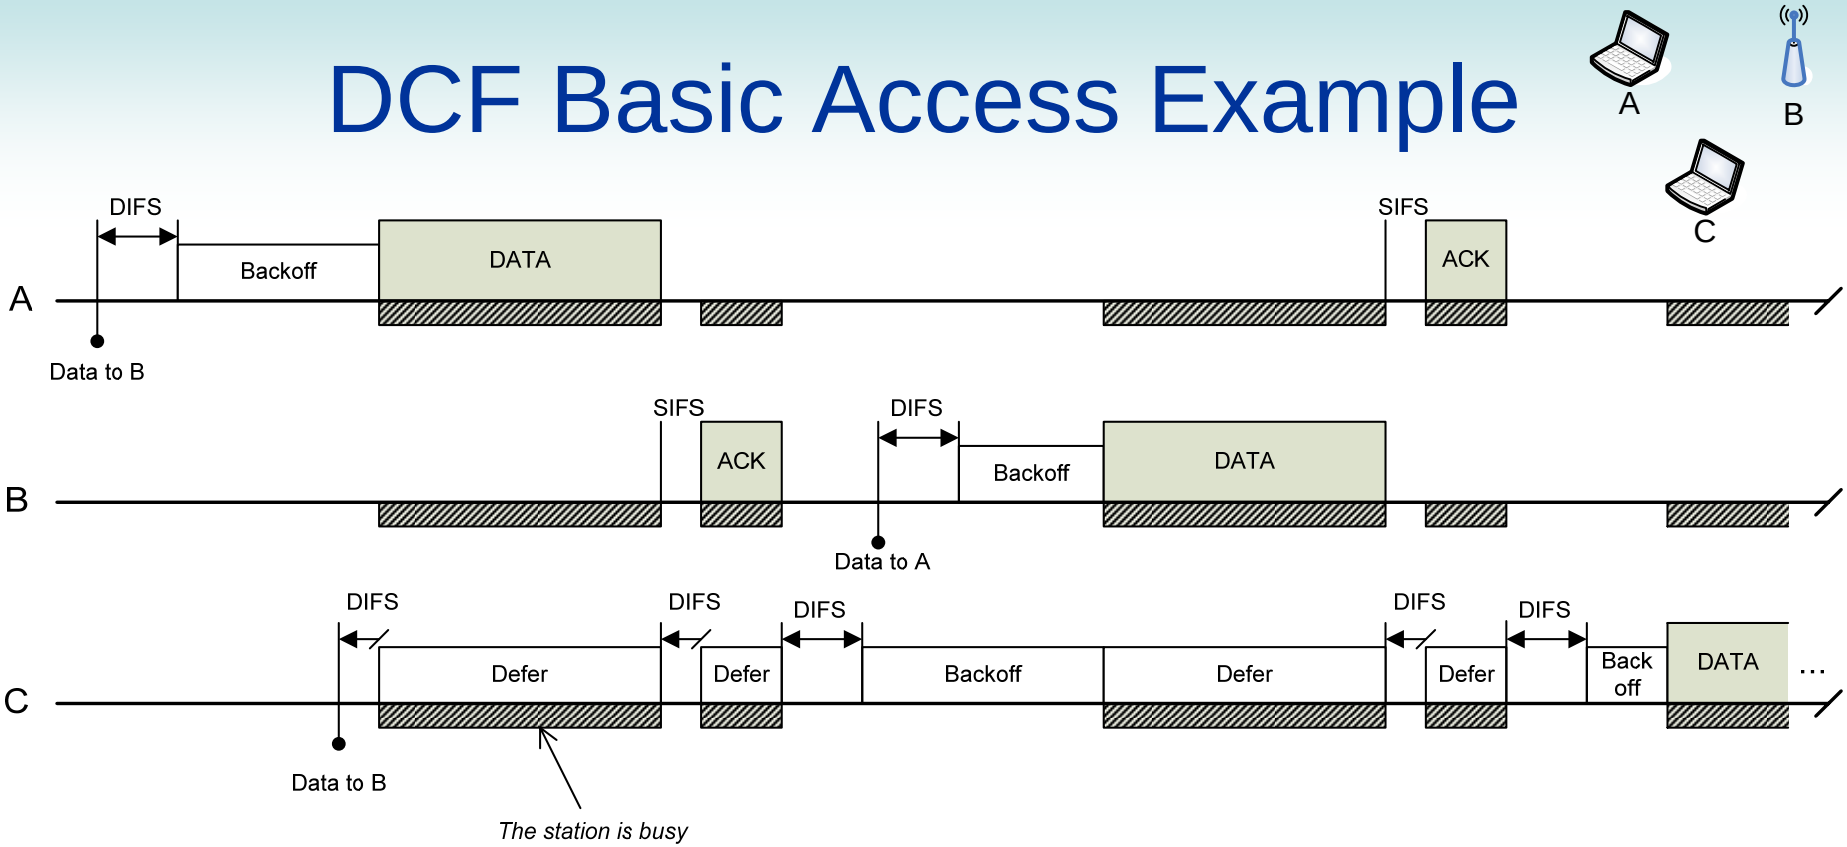

- A has data to send to B: the medium is idle for DIFS plus Backoff period, so sends the DATA and receives an ACK
- C has data to send to B: the medium becomes busy during DIFS, so defer until idle again. Finally, after medium is idle for DIFS, C selects a Backoff
- B has data to sent to A: the medium is idle for DIFS, selects a Backoff
  - In this example, it assumes B completes its Backoff before C (that is, B chose a smaller random number). Therefore B transmits the DATA, while C has to defer again
- Finally, after B is complete, C finishes its Backoff and transmits the DATA

DCF Error Control

- After sending a DATA frame, a station waits for an ACK
  - If no ACK received within ACKTimeout, then assume an error
    - Double CW (up until CWmax)
    - Retransmit the DATA frame, applying the normal rules
    - ACKTimeout: assume it's the time for SIFS + ACK
  - If too many retransmissions of a DATA frame are attempted, the station will abort and return an error to the user
    - "Too many" is usually 7 (but can be modified)

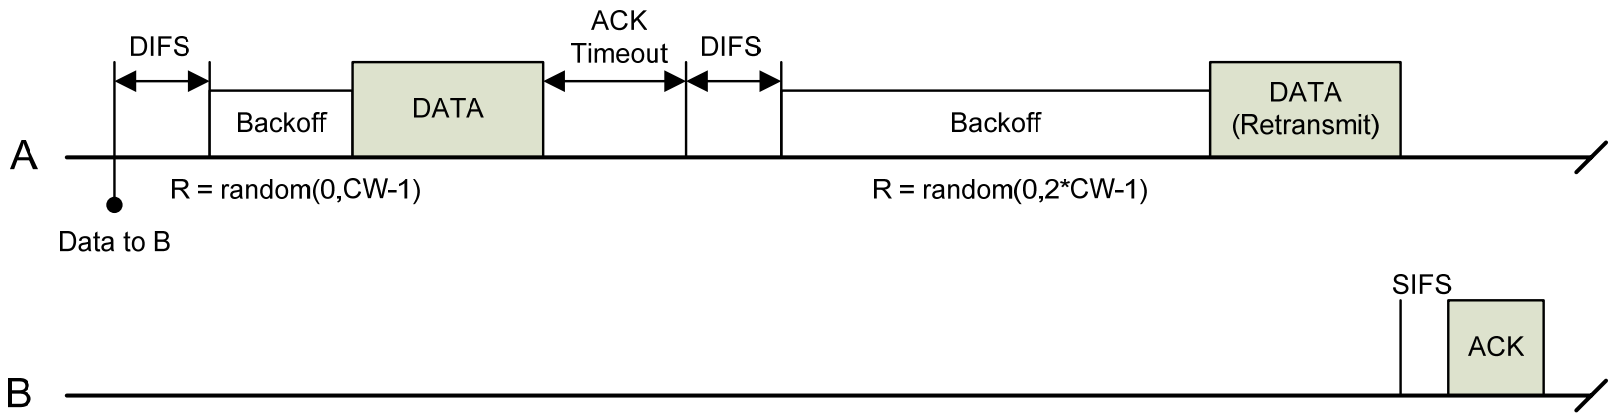

DCF Contention Window

- The size of the contention window (CW) determines how long a station waits until it transmits (backoff)
  - Backoff Period = R x SlotTime, where R is random number between 0 and CW-1
  - For each new data frame, CW = CWmin
  - For a retransmission, CW is doubled (up until CWmax)

Example: CWmin = 32, CWmax = 1024

R is randomly chosen between 0 and CW-1

The figure shows CW-1 for a retransmitted frame

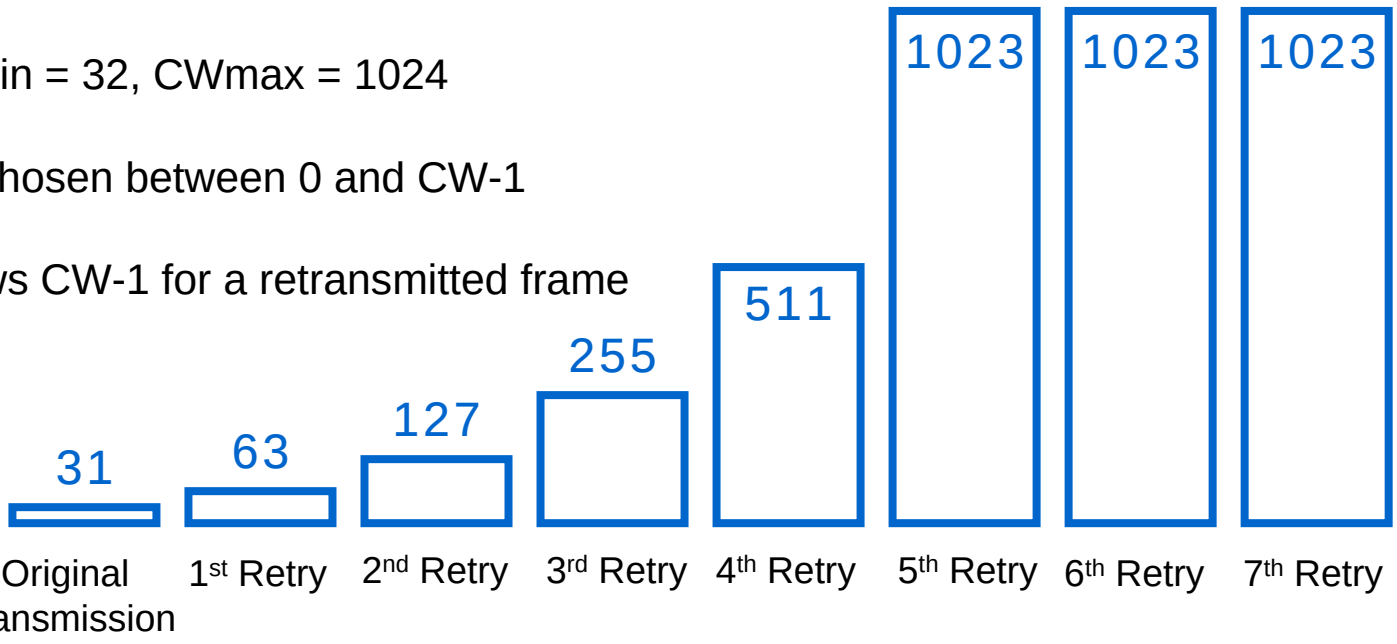

DCF Contention Window

- Purpose:
  - If two stations want to transmit, then which one does?
    - By selecting a random time to backoff, in most cases, one station will choose a shorter backoff than the other
    - The station with the shortest backoff transmits first (the other will have to wait). This implements the fairness – the chance a station gets to transmit is random (on average, all station with have equal chance to transmit)
    - But what if the two stations choose the same random number? They both transmit and collision occurs
  - The larger the value of CW, the less chance too stations will choose the same random number. Therefore, less chance of collision
    - If a collision does occur, then CW is increased (doubled) so there is *even less chance* of a collision for the retransmitted frame
  - However, the larger the value of CW, the longer time spent in backoff. Therefore, less efficient

A Problem: Hidden Stations

- Consider the scenario below:

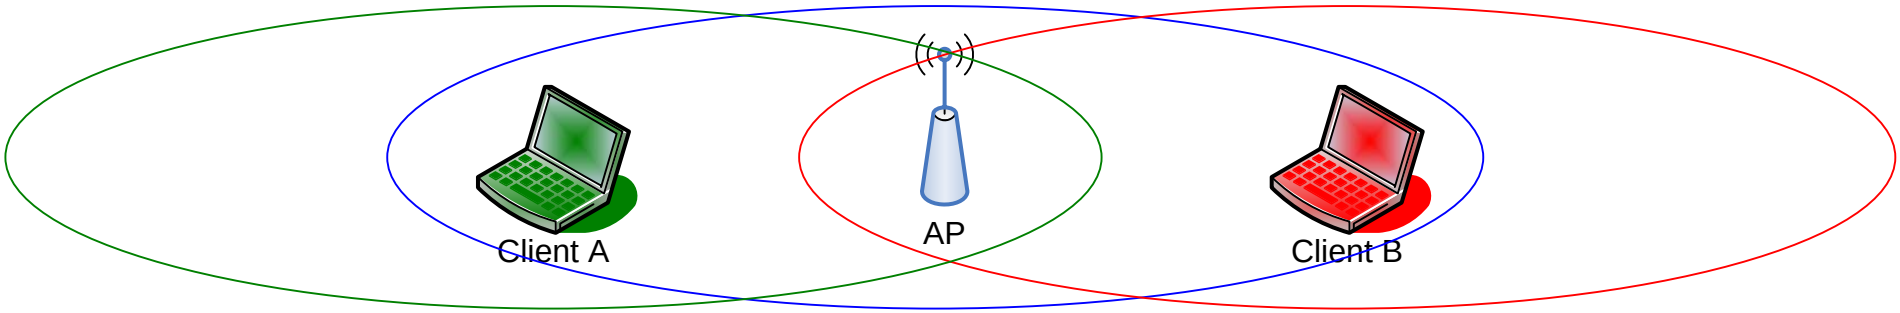

  - Client A is within range of AP
  - Client B is within range of AP
  - Client A is out of range of Client B (or Client A is *hidden* from Client B)

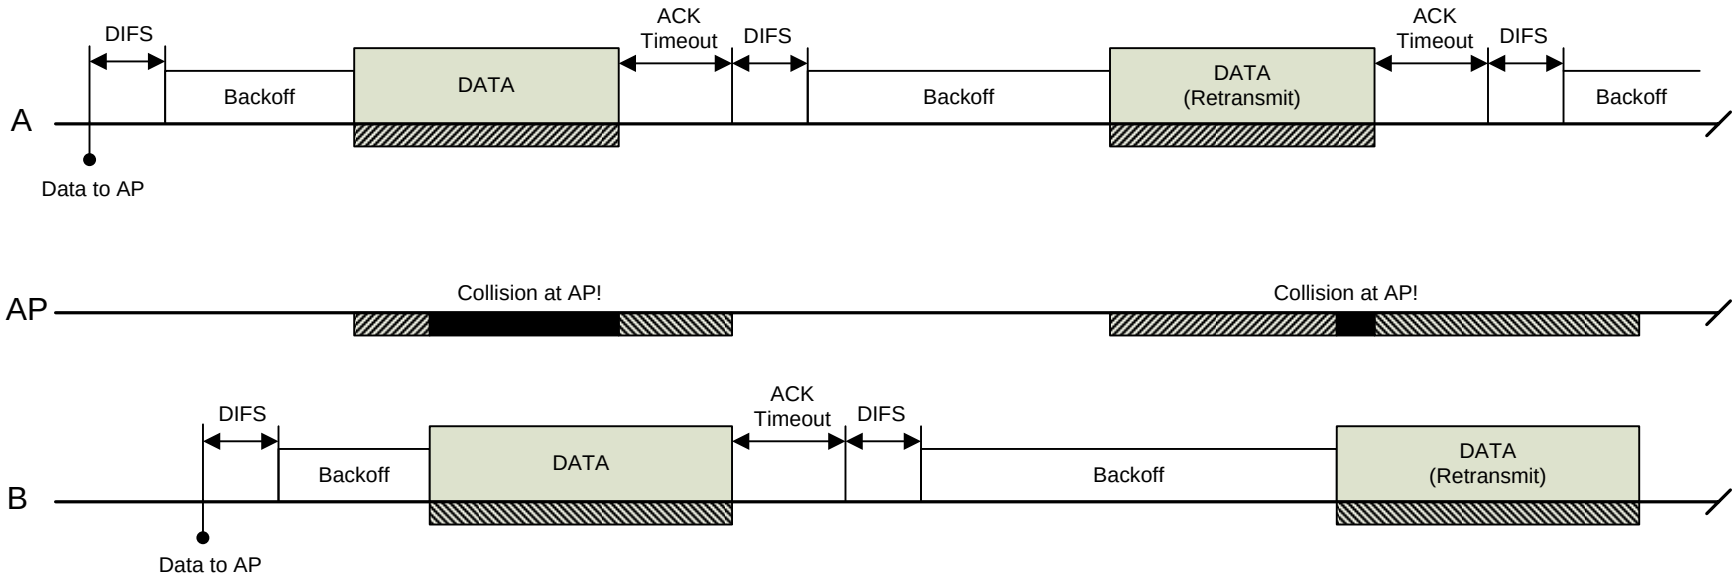

A Problem: Hidden Stations

- When two stations are *hidden* from each other (e.g. out of range, cannot hear each other), then very high chance of collisions at a receiver (e.g. AP)
  - Collisions lead to retransmissions; even chance of further collisions on retransmissions
  - Retransmissions lead to significantly lower throughput (takes a long time to successfully send a single DATA frame)
- This is called the *hidden station (or terminal) problem*
- A solution:
  - Before sending the DATA frame, ask the receiver if they are busy (Request To Send)
  - The receiver responds if they are not busy (Clear To Send)
    - If receiver is busy (receiving from someone else), then will not respond
  - When sender receives the Clear To Send, they send the DATA frame
- In IEEE 802.11 DCF there are two modes of operation
  - Basic Access: DATA then ACK
  - RTS/CTS: RTS then CTS then DATA then ACK

Basic Access and Hidden Stations

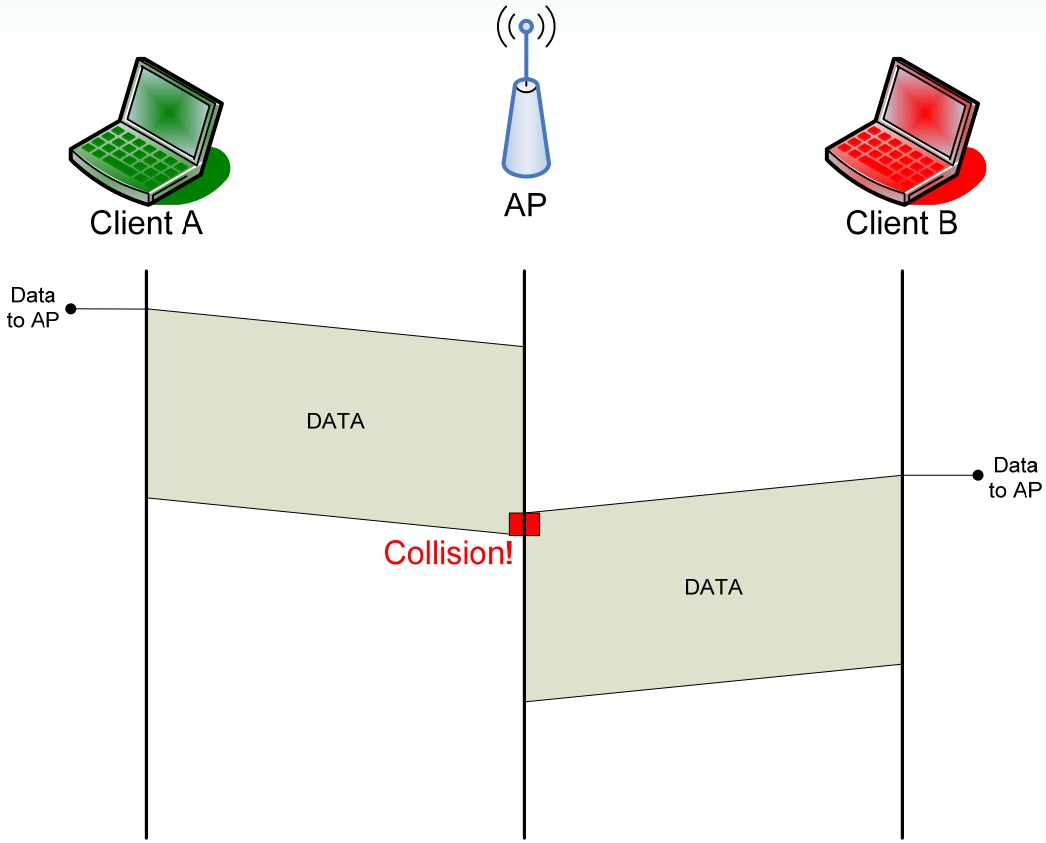

There is a high probability of collisions if using Basic Access in the presence of hidden stations. Client B starts transmitting because it believes the medium is idle (because Client B cannot hear Client A's transmission)

RTS/CTS and Hidden Stations

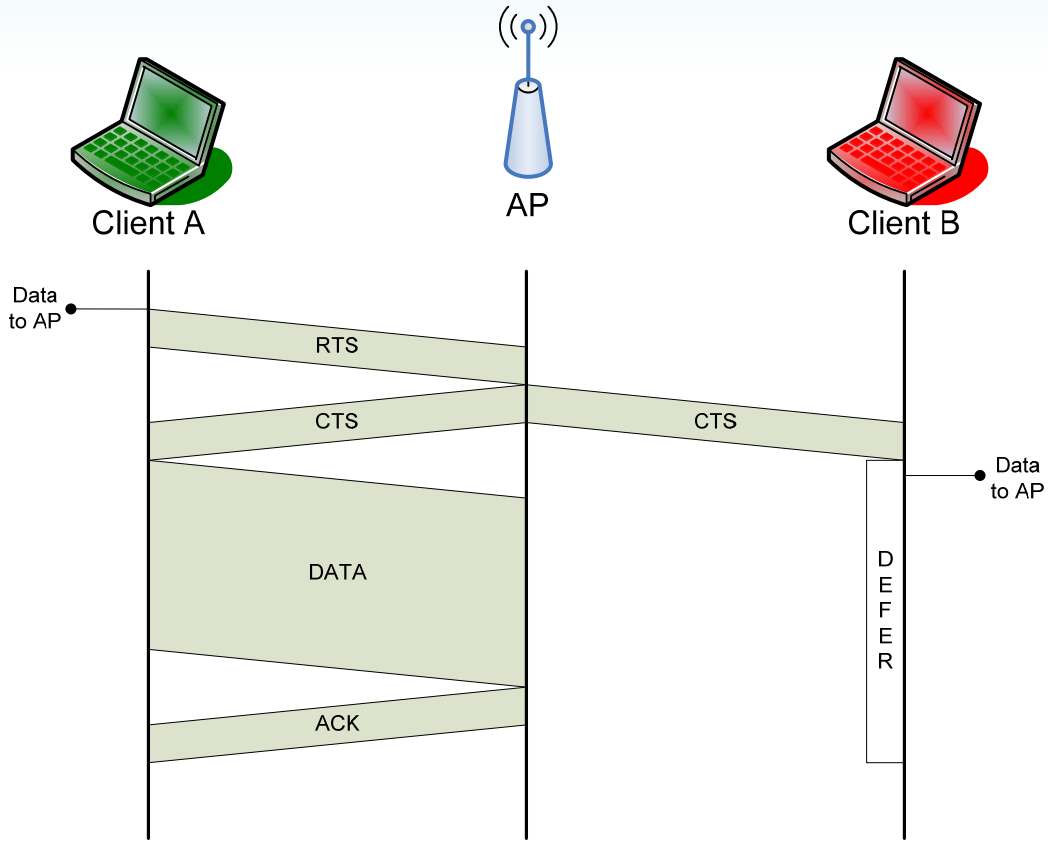

Using RTS/CTS before sending DATA reduces the probability of collision in the presence of hidden stations. When the AP responds with a CTS, Client B hears the CTS and immediately *defers* because someone else is about to transmit.

RTS/CTS Frames

- RTS Frame
  - Request To Send
  - Sent to intended recipient of DATA
  - Notifies all stations of upcoming DATA frame
  - Size ~ 20 bytes
- CTS Frame
  - Clear To Send
  - Response from the recipient of RTS
  - Notifies all stations of upcoming DATA frame
  - Size ~ 14 bytes
- DATA and ACK also used

RTS/CTS Operation

- Before sending DATA, a successful RTS and CTS handshake is needed
  - Sender sends RTS to receiver: Request To Send
    - Everyone else who hears the RTS is aware a transmission is about to take place – they defer
  - Receiver responds with CTS (if it is ready)
    - Everyone else who hears the CTS is aware a transmission is about to take place – they defer
  - Sender sends DATA to receiver, and ACK is sent in response
- The same access procedures as Basic Access applied to sending RTS frames
  - e.g. sense medium idle for DIFS then backoff
- SIFS before sending CTS, DATA, ACK
- All stations receiving RTS or CTS set their Network Allocation Vector (NAV) based on Duration field in the RTS or CTS frame
  - NAV keeps track of when the medium is in use
  - After the NAV period, other stations can attempt transmission (normal backoff rules apply)

RTS/CTS Example

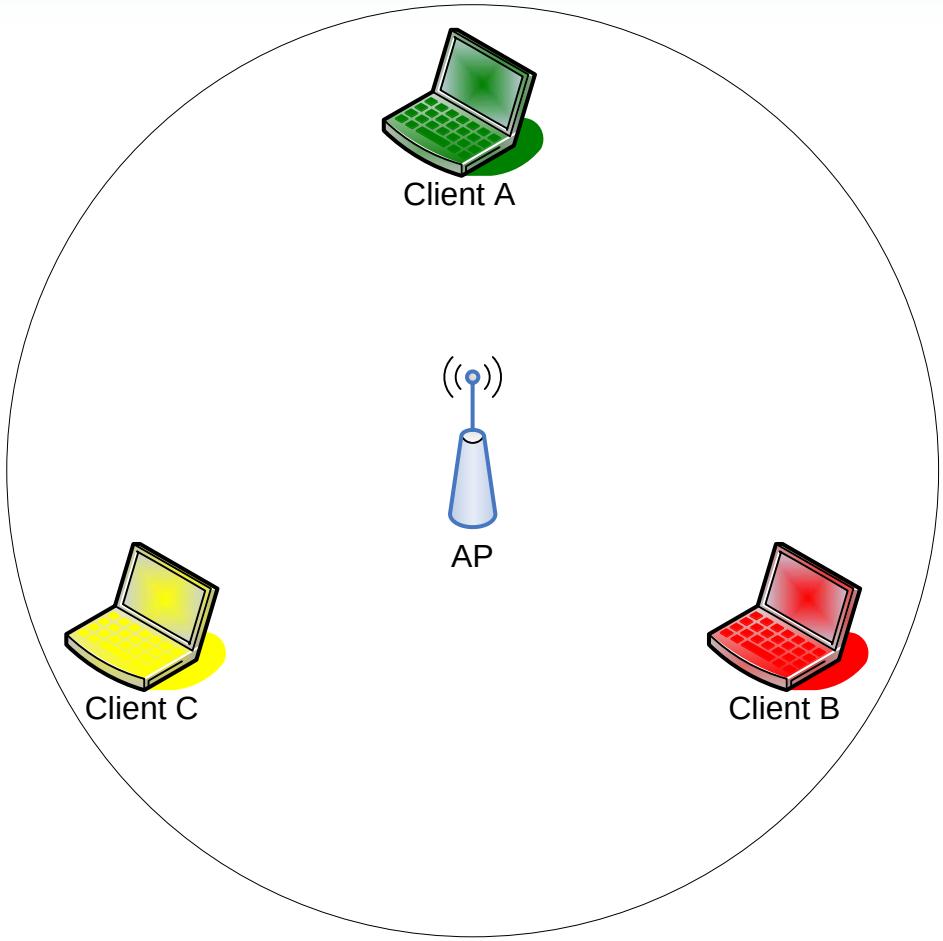

Lets assume all clients are within range of the AP, but no within range of any other client (e.g. A cannot hear B or C, B cannot hear C)
All clients have data to transmit to the AP

RTS/CTS Example

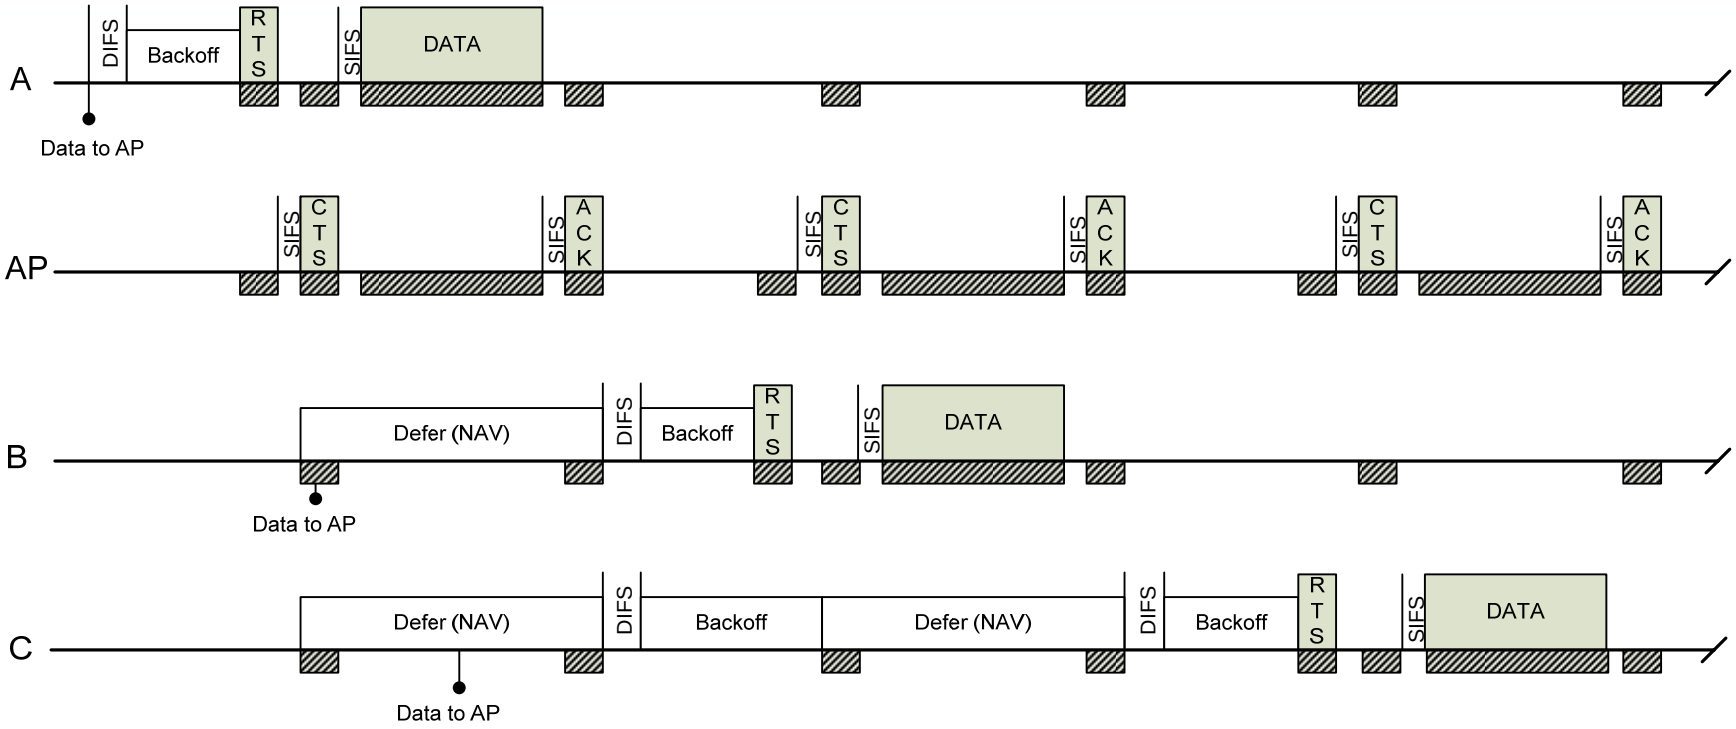

Even though the clients are hidden from each other, the use of the RTS/CTS means each client becomes aware of the other clients' transmissions (since each client hears the CTS from the AP). As a result there are no collisions: one station transmits at a time.

Basic Access vs RTS/CTS

- In practice, Basic Access is used for small DATA frames, RTS/CTS used for large DATA frames
  - RTSThreshold is a parameter in wireless LAN devices
    - If the DATA frame is smaller than RTSThreshold, then use Basic Access
    - If the DATA frame is larger than (or equal to) RTSThreshold, use RTC/CTS
- Why?
  - RTS/CTS is good for solving hidden terminal problem (and in general, avoiding collisions), but introduces an extra overhead
    - With small frames, the overhead is very significant
- Other factors that determine RTSThreshold:
  - Amount of traffic/nodes in network: with higher traffic it is better to use RTS/CTS more (hence lower RTSThreshold)
  - Use of RTS/CTS may lead to *exposed terminal problem*, which reduces the throughput!

IEEE 802.11 MAC Performance

- Physical layer offers raw data rate (e.g. 54Mb/s in 802.11g)
- MAC introduces overheads to provide addressing, reliability and management:
  - Frame headers
  - Control frames: ACK, RTS, CTS, …
  - Interframe spaces
  - Backoffs
  - Collisions and retransmissions
- We are interested in throughput at MAC layer
  - MAC layer throughput is the rate at which user data is successfully delivered to destination
  - An approximate way to measure the throughput is to count the total data (minus MAC headers) delivered over a period of time

IEEE 802.11 MAC Parameters

Parameter	802.11b	802.11a	802.11g
DIFS	50μs	34 μs	28 μs
SIFS	10 μs	16 μs	10 μs
SlotTime	20 μs	9 μs	9 μs
CWmin	31	15	15
CWmax	1023	1023	1023

Simple Throughput Calculation

- Assume 1500 byte payload using Basic Access, no collisions or deference!
- Best case for 802.11g (54Mbit/s)
  - Time = DIFS + AverageBackoff + DATA + SIFS + ACK
  - Time = 28 + 7.5*9 + (1500+34)*8/54 + 10 + 14*8/54
    - = 334 usec
  - Throughput = 1500bytes/334 usec
    - = 35.93Mbit/s
- RTS/CTS: need to add 2 x SIFS + RTS + CTS time
  - 11g throughput: 33.61Mb/s

- But we need to consider collisions and deference!

Theoretical Throughput

802.11b Example
10 clients per AP
Basic Access
1000 byte payload
600kb/s per client

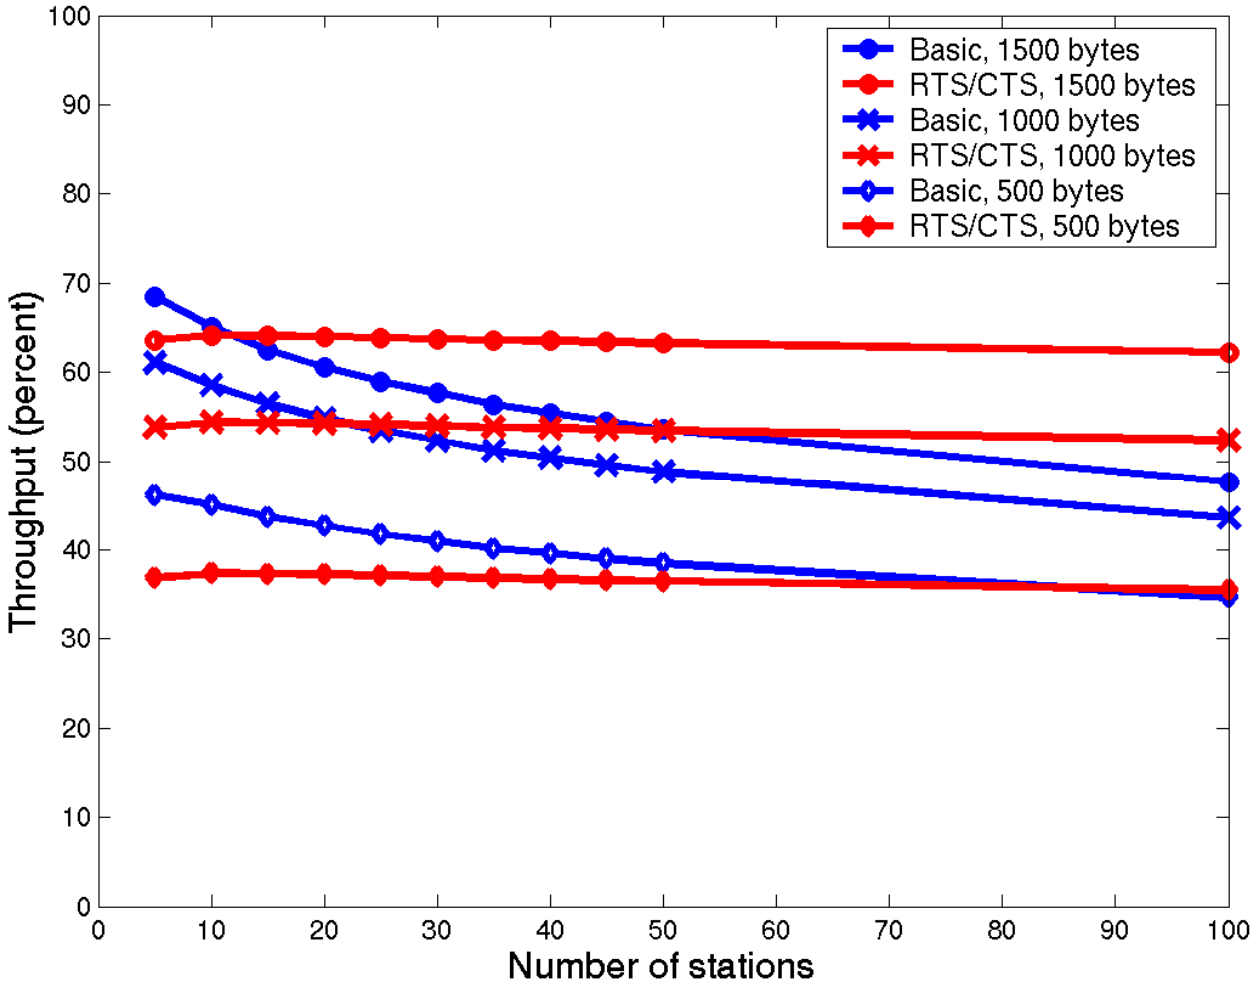

Realistic Throughput

- Take into account:
  - Collisions, retransmissions
  - IP, TCP and other protocol overheads
  - Varying sizes of payload
  - About 10 nodes per AP
- All IEEE 802.11b clients; 11b AP
  - 3 to 5 Mb/s per cell (Basic Service Set)
- Mixture of IEEE 802.11g and 11b clients; 11g AP
  - 10 to 15 Mb/s per cell

Security in Wireless LANs

- Original 802.11
  - Authentication
    - Ensure the client has permission to access the network
    - Originally used a shared secret key (Wired Equivalent Privacy, WEP)
      - Client and AP must be pre-configured with the same secret key
  - Confidentiality
    - Ensure the communications between client and AP cannot be overheard
      - WEP shared secret key also used for encryption
  - WEP has several limitations
    - In practice, if an attacker can collect several GB of traffic between a client and AP, it can discover the secret key
- Enhanced Wireless LAN Security
  - Wireless Protected Access (WPA): increase key size and solve WEP problems
  - IEEE 802.11i: complete security architecture that can use other network security mechanisms

Wireless LAN Design Issues

- How many users per Access Point?
  - Performance per user drops as number of users increase
  - But we want to minimise number of APs
    - Costly devices, costly to install and manage
    - Handover between APs may become inefficient
- Basic Access versus RTS/CTS – What RTS threshold?
  - Basic Access is more efficient if few collisions (unless hidden terminals)
  - RTS/CTS helps avoid hidden terminals
- How to cover a large area?
  - Cellular coverage: many small cells or a few large cells?
  - Avoid interference between cells
    - Use different frequencies, but only 3 non-overlapping frequencies available

Wireless LAN Design Issues

- How do we secure the network?
  - Need to authenticate users (usually to a central network authentication server)
  - Need encryption: Layer 2 (802.11 WEP, WPA, 11i) or other layer (IPsec, VPN, …)
- How do we give priority to users and applications?
  - Voice calls get priority over data traffic
  - Quality of service management on APs; but what about network wide?