UNIVERSITY OF WATERLOO

CS 456/656 Computer Networks

Lecture 16: Link Layer – Part 3

Mina Tahmasbi Arashloo and Bo Sun

Fall 2024

A note on the slides

Adapted from the slides that accompany this book.

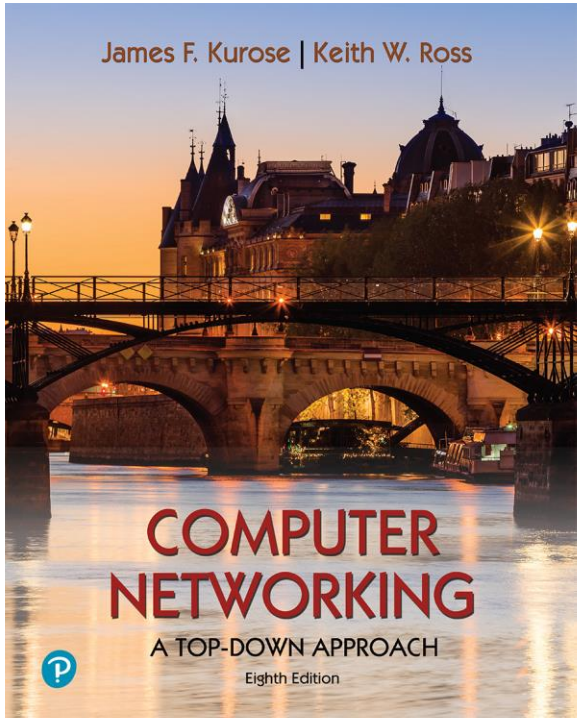

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Link layer: roadmap

- Link layer overview
  - Local Area Networks (LANs)
- Switched LANs
  - Ethernet and Addressing
  - Address Resolution Protocol (ARP)
  - Switches
- Virtual LANs (VLANs)
- Shared LANs and multiple access protocols
  - Channel partitioning
  - Random access
  - “taking turns”

Link layer: local connectivity

Also called a Local Area Network (LAN)

Either “shared link” or a link-layer network

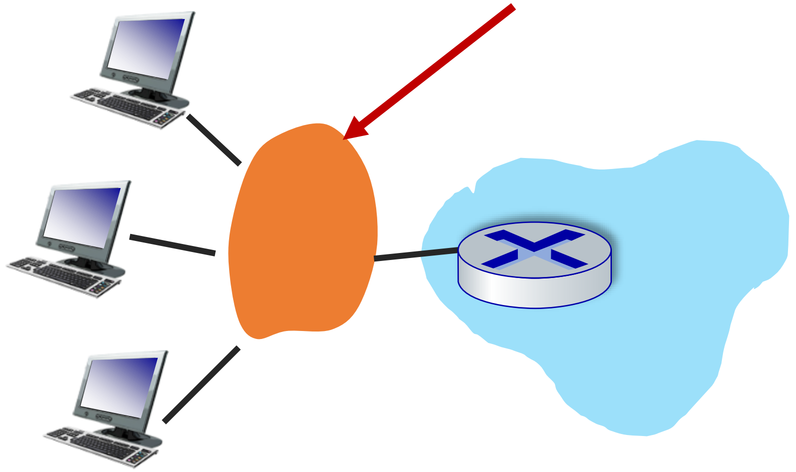

Multiple access links and protocols

two types of "links":

- point-to-point
  - point-to-point link between Ethernet switches and hosts
- shared wire or medium (broadcast)
  - old-school Ethernet
  - 802.11 wireless LAN, 4G/4G. Satellite
  - …

shared wire (e.g., cabled Ethernet)

shared radio: 4G/5G

shared radio: WiFi

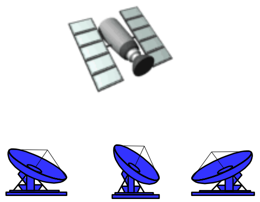

shared radio: satellite

humans at a cocktail party (shared air, acoustical)

Multiple access protocols

- single shared communication channel
- two or more simultaneous transmissions by nodes can lead to interference
  - *collision* if node receives two or more signals at the same time

multiple access protocol

- distributed algorithm that determines how nodes share channel, i.e., determine when node can transmit
- communication about channel sharing must use channel itself!
  - no out-of-band channel for coordination

An ideal multiple access protocol

given*:* multiple access channel (MAC) of rate *R* bps

what we ideally want*:*

1. when one node wants to transmit, it can send at rate *R*.
2. when *M* nodes want to transmit, each can send at average rate *R/M*
3. fully decentralized:
   - no special node to coordinate transmissions
   - no synchronization of clocks, slots
4. simple

MAC protocols: taxonomy

three broad classes:

- channel partitioning
  - divide channel into smaller “pieces” (time slots, frequency, code)
  - allocate piece to node for exclusive use
- random access
  - channel not divided, allow collisions
  - “recover” from collisions
- “taking turns”
  - nodes take turns, but nodes with more to send can take longer turns

Link layer: roadmap

- Link layer overview
  - Local Area Networks (LANs)
- Switched LANs
  - Ethernet and Addressing
  - Address Resolution Protocol (ARP)
  - Switches
- Virtual LANs (VLANs)
- Shared LANs and multiple access protocols
  - Channel partitioning
  - Random access
  - “taking turns”

Channel partitioning MAC protocols: TDMA

TDMA: time division multiple access

- access to channel in "rounds"
- each station gets fixed length slot in each round
- unused slots go idle
- example: 6-station LAN, 1,3,4 have packets to send, slots 2,5,6 idle

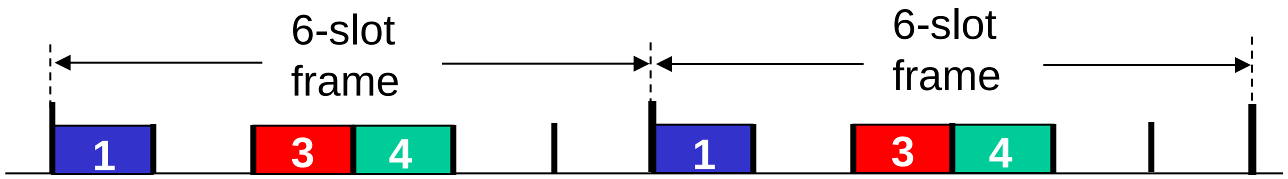

Channel partitioning MAC protocols: FDMA

FDMA: frequency division multiple access

- channel spectrum divided into frequency bands
- each station assigned fixed frequency band
- unused transmission time in frequency bands go idle
- example: 6-station LAN, 1,3,4 have packet to send, frequency bands 2,5,6 idle

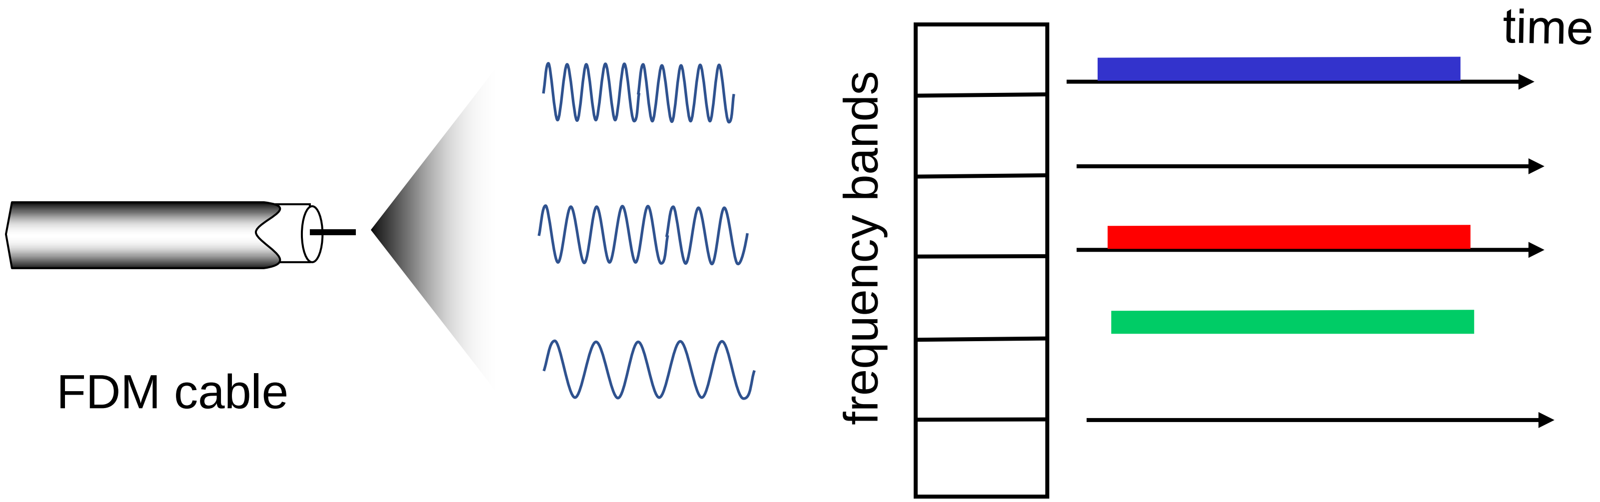

Link layer: roadmap

- Link layer overview
  - Local Area Networks (LANs)
- Switched LANs
  - Ethernet and Addressing
  - Address Resolution Protocol (ARP)
  - Switches
- Virtual LANs (VLANs)
- Shared LANs and multiple access protocols
  - <u>Channel partitioning</u>
  - <u>Random access</u>
  - “taking turns”

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 MAC protocols:
  - ALOHA, slotted ALOHA
  - CSMA, CSMA/CD, CSMA/CA

CSMA (carrier sense multiple access)

simple CSMA: listen before transmit:

- if channel sensed idle: transmit entire frame
- if channel sensed busy: defer transmission

- human analogy: <u>don't interrupt others!</u>

CSMA: collisions

- collisions can *still* occur with carrier sensing:
  - propagation delay means two nodes may not hear each other's just-started transmission
- collision: entire packet transmission time wasted
  - distance & propagation delay play role in determining collision probability

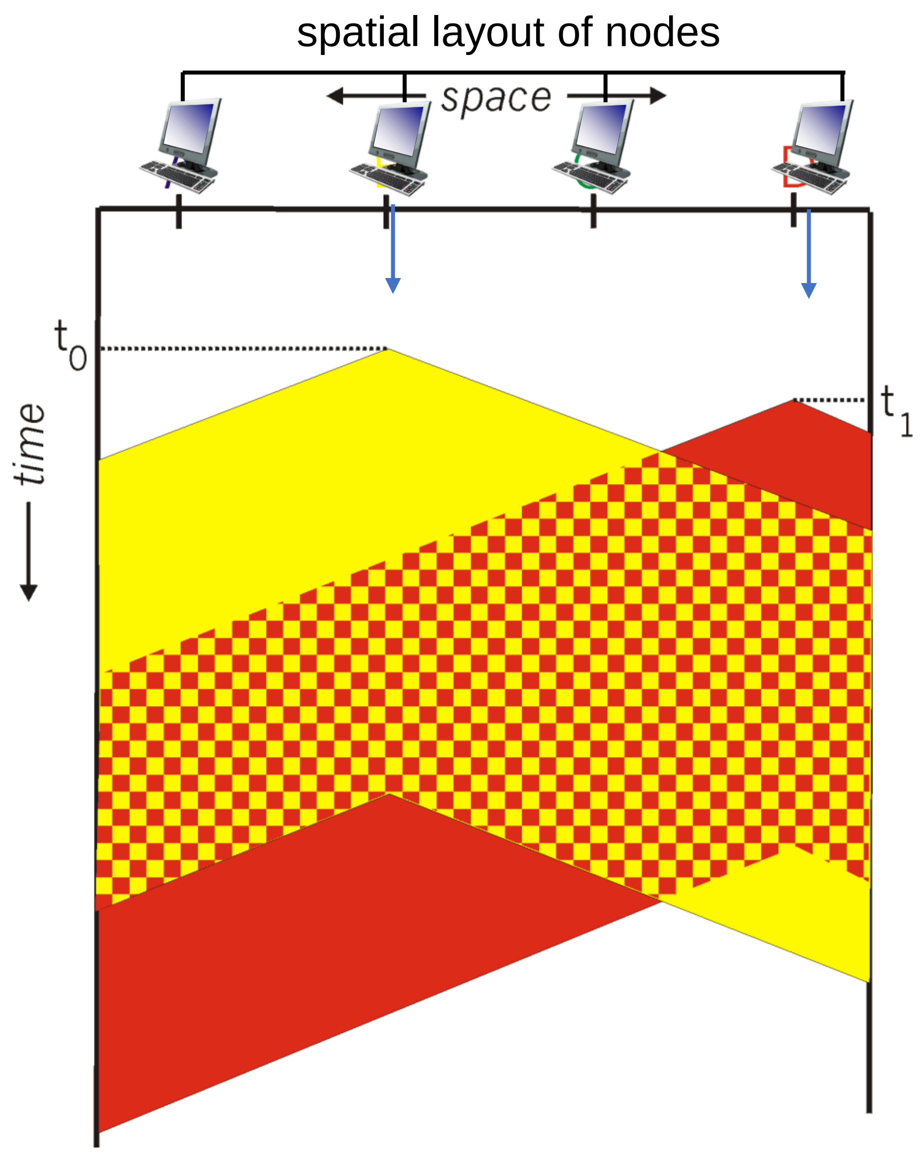

CSMA (carrier sense multiple access)

simple CSMA: listen before transmit:

- if channel sensed idle: transmit entire frame
- if channel sensed busy: defer transmission

- human analogy: don't interrupt others!

CSMA/CD: CSMA with *collision detection*

- monitor for incoming signals while transmitting
- if collision detected: stop sending

- human analogy: If someone else starts talking at the same time, stop talking (the polite conversationalist).

CSMA/CD:

- CSMA/CD reduces the amount of time wasted in collisions
  - transmission aborted on collision detection

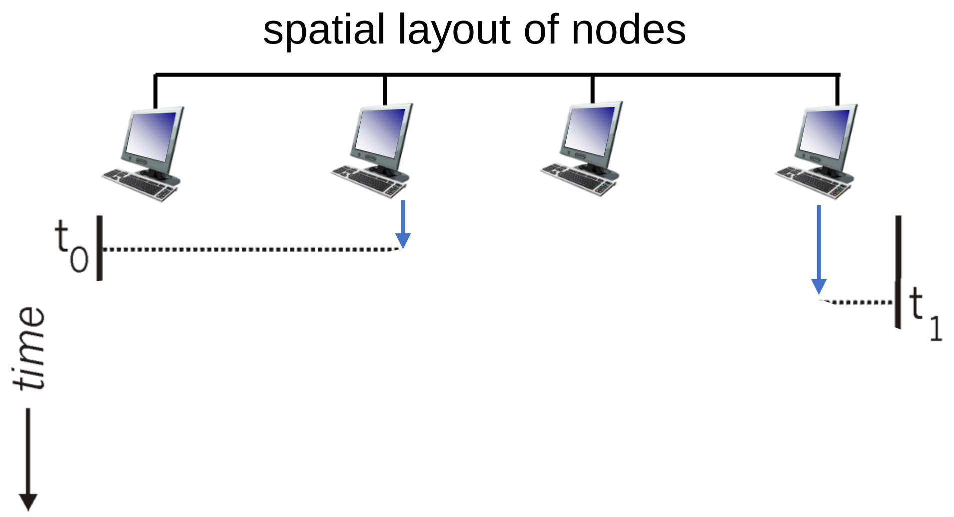

CSMA/CD:

- CSMA/CD reduces the amount of time wasted in collisions
  - transmission aborted on collision detection
- collisions *detected* within short time
- colliding transmissions aborted, reducing channel wastage

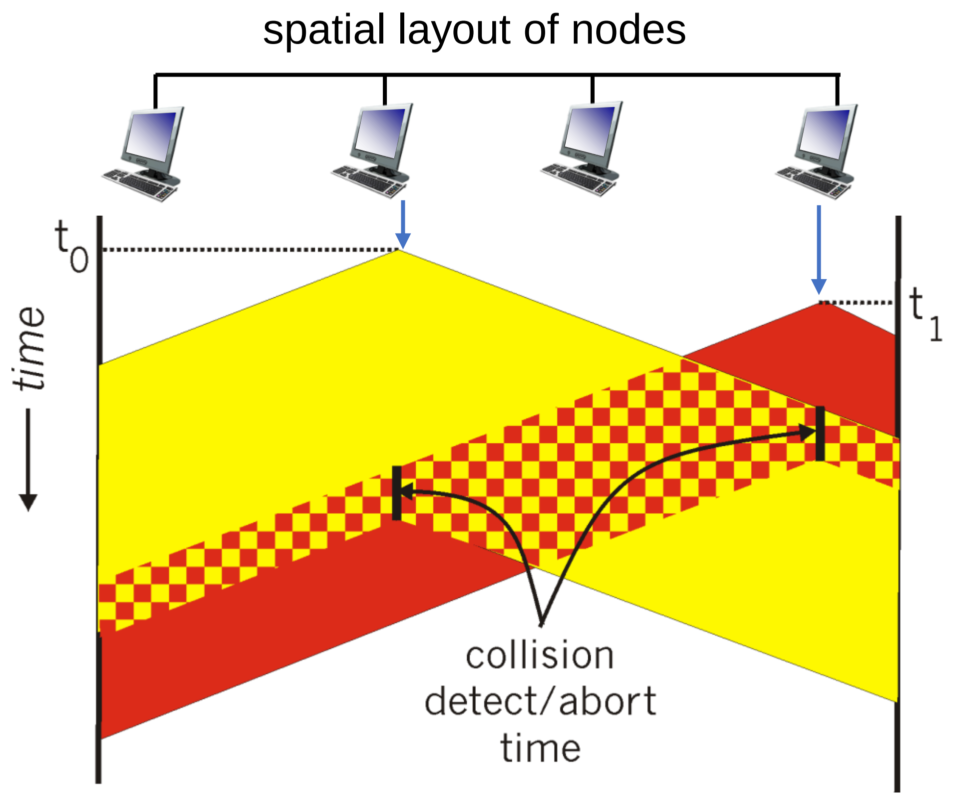

Ethernet CSMA/CD algorithm

1. Ethernet receives datagram from network layer, creates frame
2. If Ethernet senses channel:
   if idle: start frame transmission.
   if busy: wait until channel idle, then transmit
3. If entire frame transmitted without collision - done!

Ethernet CSMA/CD algorithm

4. If another transmission detected while sending: abort
5. After aborting, enter *binary (exponential) backoff:*
   - after *m*th collision, chooses *K* at random from $\{0,1,2, \ldots, 2^m-1\}$.
   - wait $K \cdot 512$ bit times, i.e., K times the amount of time need to send 512 bits
   - returns to Step 2 (sense before sending)

Ethernet CSMA/CD algorithm

4. If another transmission detected while sending: abort
5. After aborting, enter *binary (exponential) backoff:*
   - after *m*th collision, chooses *K* at random from *{0,1,2, …, 2^m-1}*.
   - wait $K \cdot 512$ bit times, i.e., K times the amount of time need to send 512 bits
   - returns to Step 2 (sense before sending)

more collisions: longer backoff interval

Minimum frame size

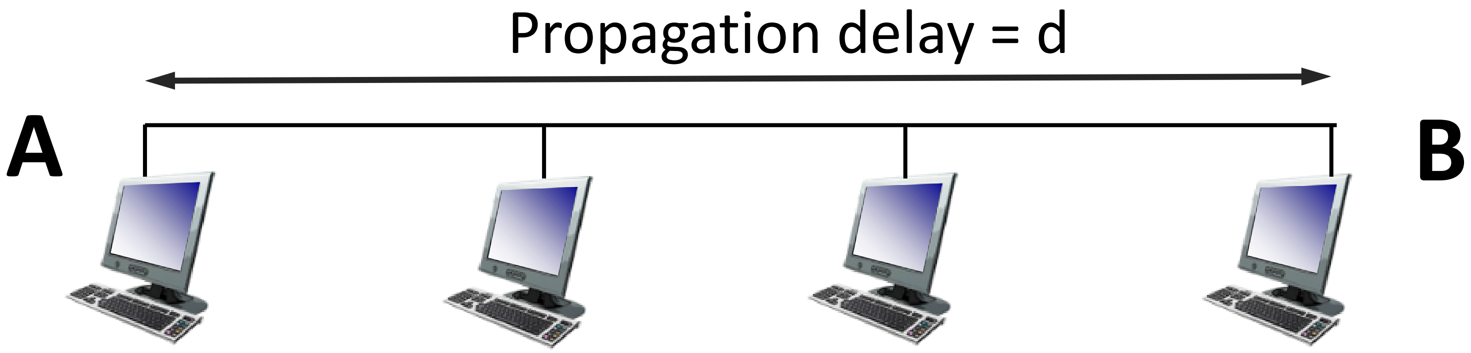

- Suppose A sends a frame at time t
- B sees an idle channel right before t + d and starts transmitting a frame
- A won’t see a collision until t + 2d

Minimum frame size

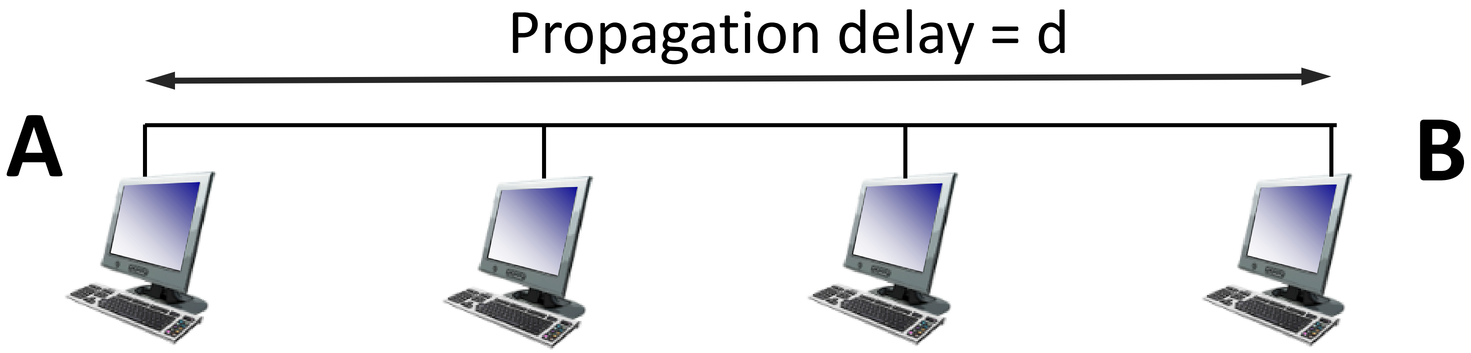

- A should wait for time 2d to detect collision
  - So, it will keep transmitting during this period
- That's why there are restrictions on "classical" Ethernet
  - Maximum length of the wire: 2500 meters
  - Minimum length of the frame: 512 bits (64 bytes)

Link characteristics affect protocol design

- In wired LANs, any two nodes on the shared medium can detect collision easily
  - measure signal strengths, and compare the transmitted and received signals
  - Ethernet uses CSMA/CD
- But, in wireless LANs, collision detection is difficult
  - due to characteristics of wireless links
  - wireless LANs (WiFi) uses CSMA/CA: CSMA with *collision avoidance*

IEEE 802.11 (WiFi) MAC Protocol: CSMA/CA

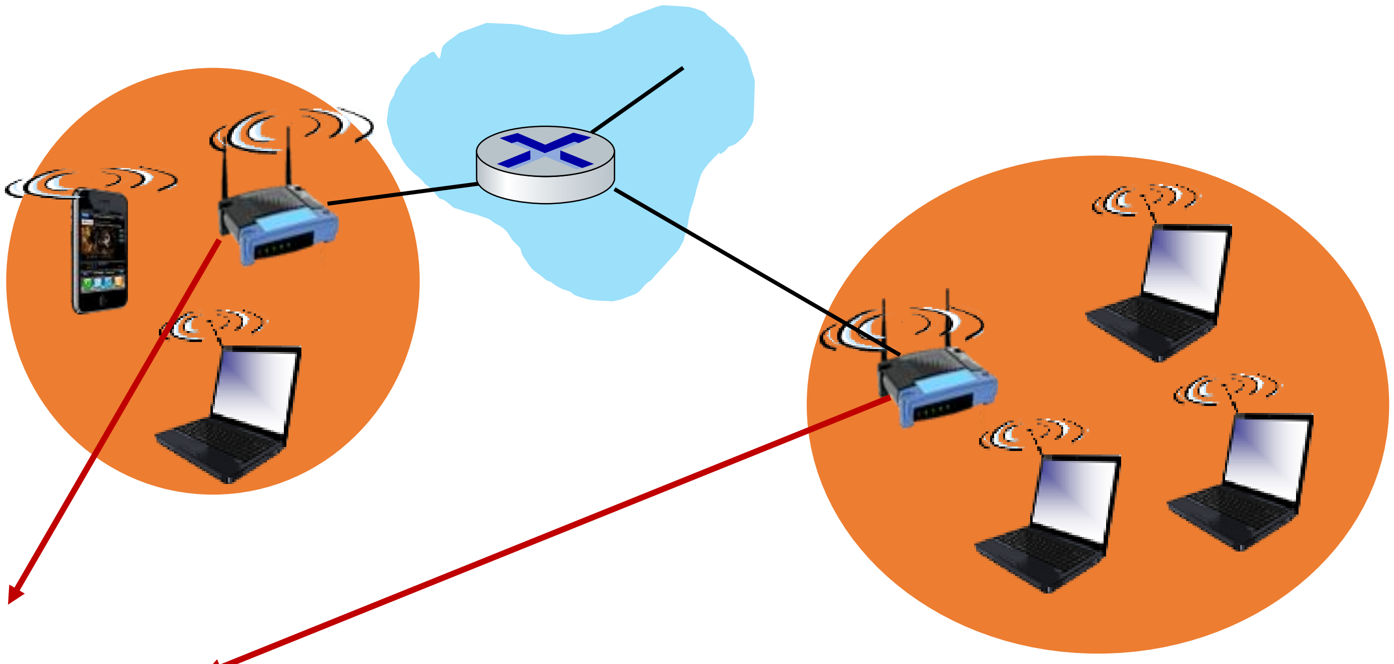

Base stations (or Access point):

- Connects end points via a wireless “link”
  - Shared physical medium
- Connect to the wired network
  - E.g., the Internet
  - Provide connection from user devices to the wired network

Wireless link characteristics

fading (attenuation)

- radio signal attenuates (loses power) as it propagates (free space "path loss")

noise

- received signal is a combination of attenuated original signal and background noise in the environment -> more "lossy" than wired link
- SNR: signal-to-noise ratio
  - larger SNR -> lower bit error rate (BER) -- easier to extract signal from noise (a "good thing")

Wireless link characteristics

Hidden terminal problem

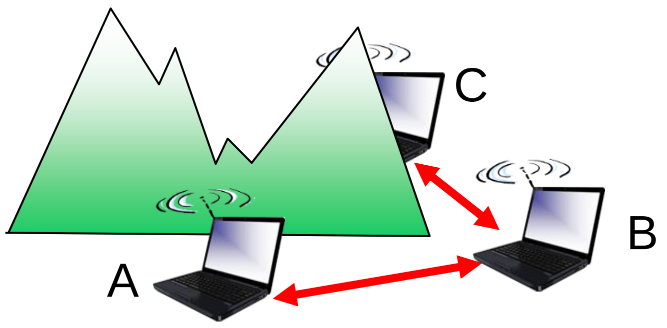

- B, A hear each other
- B, C hear each other
- A, C can not hear each other means A, C unaware of their interference at B

Attenuation also causes "hidden terminals"

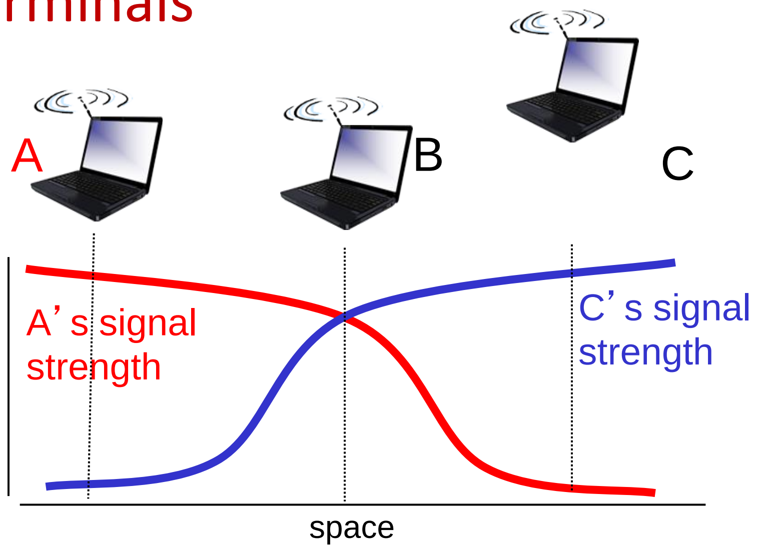

- B, A hear each other
- B, C hear each other
- A, C can not hear each other interfering at B

WiFi CSMA/CA protocol

Wireless link characteristics affect MAC protocol design

- costly to detect a collision (attenuated received signal)
- two nodes may collide but not be able to detect it (hidden terminals)
- high bit error rates

CSMA/CA: CSMA with *collision avoidance*

- use CA technique instead of CD
- use a link-layer RDT (detection and retransmission)

WiFi CSMA/CA protocol

802.11 sender

1 if sense channel idle for **DIFS** then

transmit entire frame (no CD)

2 if sense channel busy then

start random backoff time

timer counts down while channel idle

transmit when timer expires

if no ACK, increase random backoff interval, repeat 2

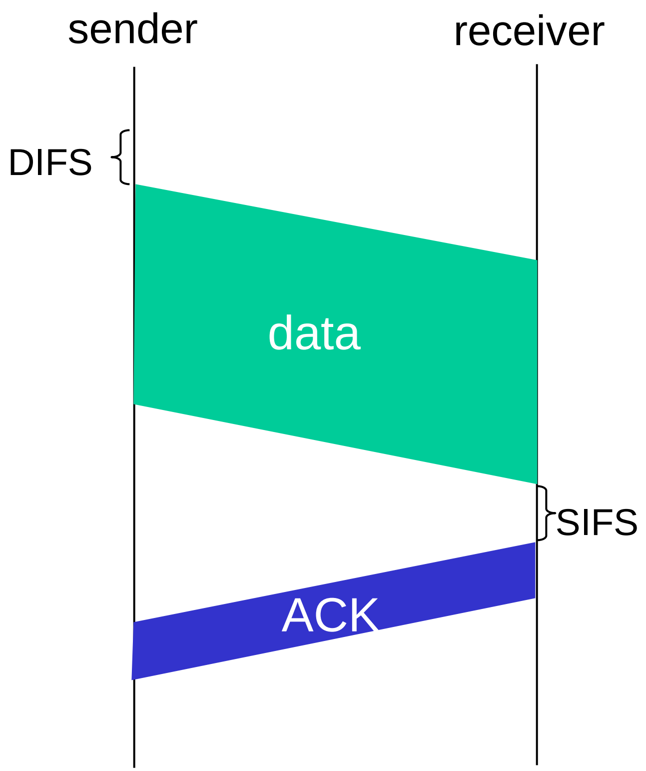

802.11 receiver

if frame received OK

return ACK after **SIFS** (ACK needed due to hidden terminal problem)

WiFi CSMA/CA protocol

802.11 sender

Note the differences from CSMA/CD

1 if sense channel idle for **DIFS** then
  transmit entire frame (no CD)

2 if sense channel busy then
  start random backoff time
  timer counts down while channel idle
  transmit when timer expires
  if no ACK, increase random backoff interval, repeat 2

802.11 receiver

if frame received OK
  return ACK after **SIFS** (ACK needed due to hidden terminal problem)

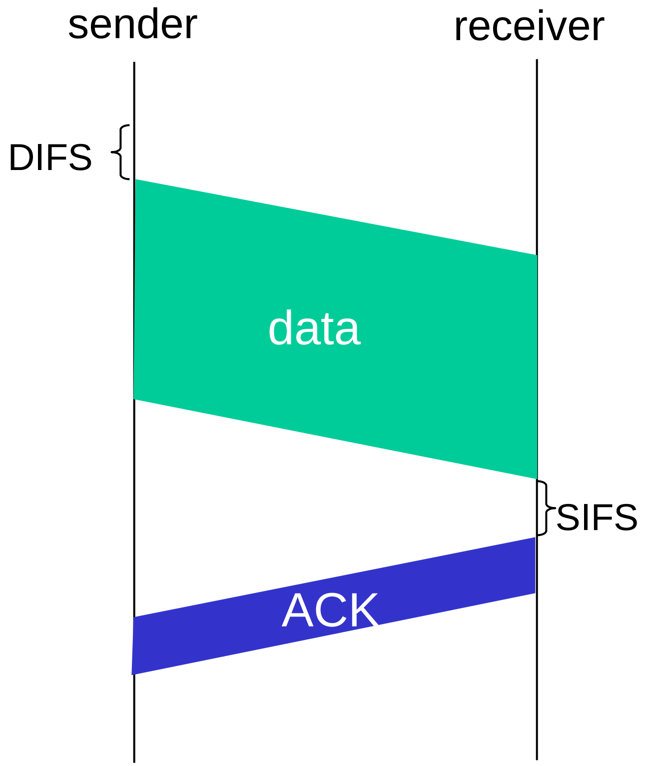

Avoiding collisions: (optional) reservations

idea: sender "reserves" channel use for data frames using small reservation packets

- sender first transmits *small* request-to-send (RTS) packet to BS using CSMA
  - RTSs may still collide with each other (but they're short)
- BS broadcasts clear-to-send CTS in response to RTS
- CTS heard by all nodes
  - sender transmits data frame
  - other stations defer transmissions

Wireless links affect higher-layer protocols

- In wireless LAN, bit errors are much more common than in wired networks. Packets may often be corrupted or lost for reasons other than congestion
  - but TCP will interpret any packet loss as congestion and reduce its send window
- Solutions?
  - Have the "wireless" link layer protocol do retransmissions
  - Provide extra signals to TCP to convey if a loss is due to the nature of the wireless link rather than congestion
  - …

Link layer: roadmap

- Link layer overview
  - Local Area Networks (LANs)
- Switched LANs
  - Ethernet and Addressing
  - Address Resolution Protocol (ARP)
  - Switches
- Virtual LANs (VLANs)
- Shared LANs and multiple access protocols
  - Channel partitioning
  - Random access
  - "taking turns"

“Taking turns” MAC protocols

channel partitioning MAC protocols:

- share channel *efficiently* and *fairly* at high load
- inefficient at low load: delay in channel access, 1/N bandwidth allocated even if only 1 active node!

random access MAC protocols

- efficient at low load: single node can fully utilize channel
- high load: collision overhead

“taking turns” protocols

- look for best of both worlds!

“Taking turns” MAC protocols

polling:

- centralized controller “invites” other nodes to transmit in turn
- typically used with “dumb” devices
- concerns:
  - polling overhead
  - latency
  - single point of failure (controller)
- Bluetooth uses polling

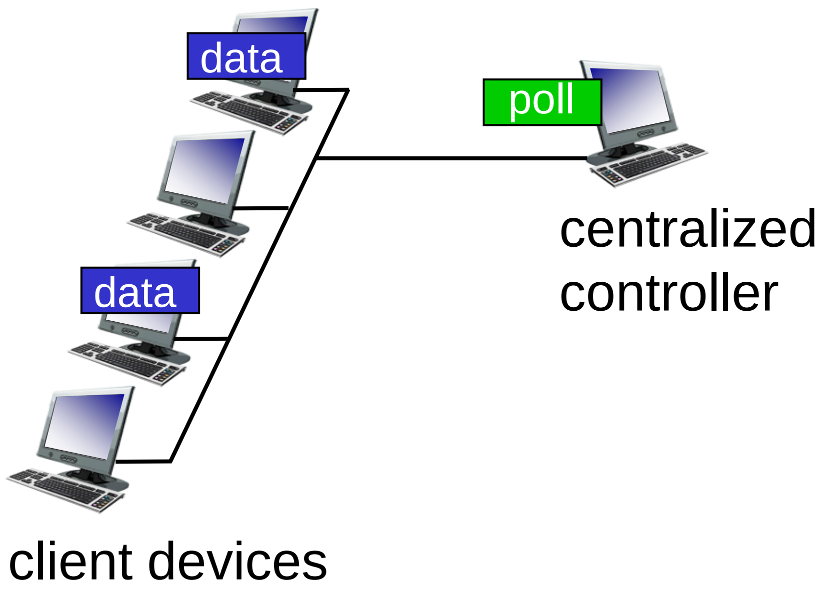

"Taking turns" MAC protocols

token passing:

- control *token* message explicitly passed from one node to next, sequentially
  - transmit while holding token
- concerns:
  - token overhead
  - latency
  - single point of failure (token)

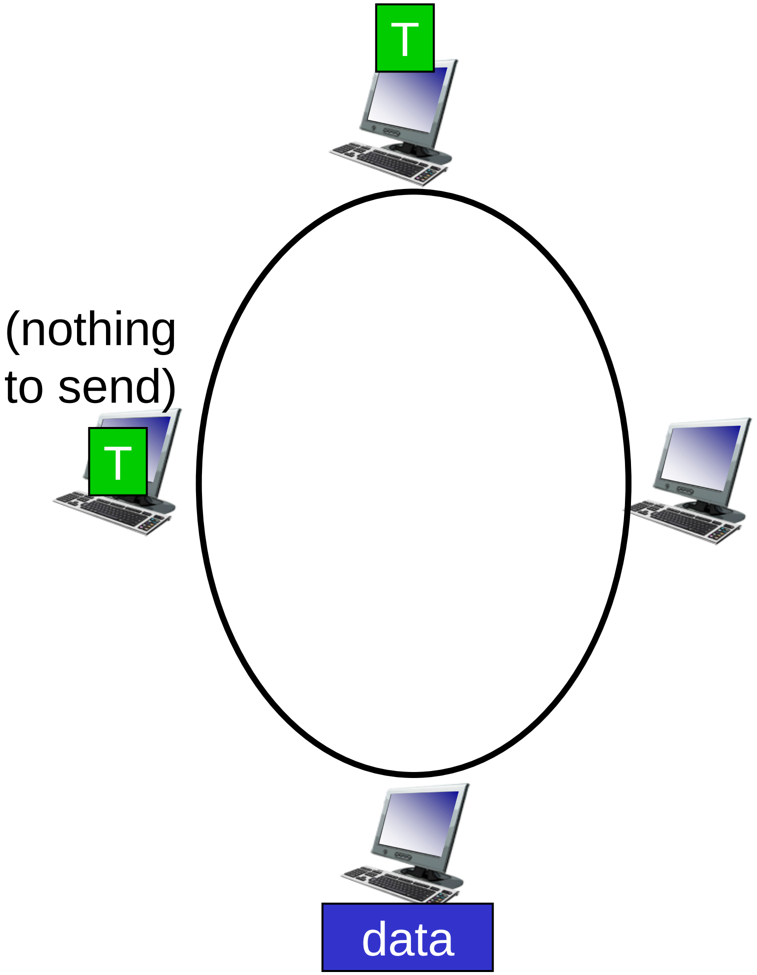

Cable access network: FDM, TDM *and* random access!

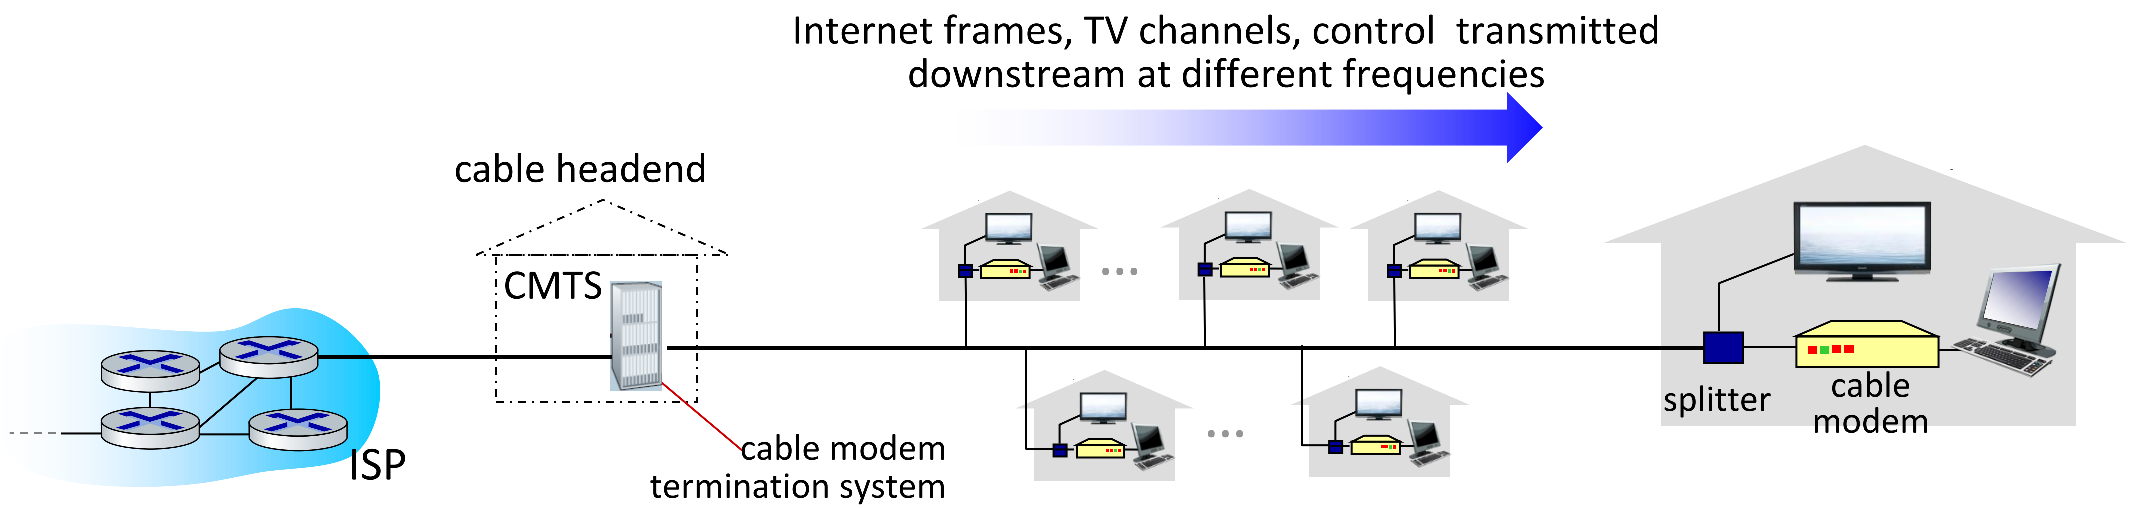

- multiple downstream (broadcast) FDM channels: up to 1.6 Gbps/channel
  - single CMTS transmits into channels
- multiple upstream channels (up to 1 Gbps/channel)
  - multiple access: all users contend (random access) for certain upstream channel time slots; others assigned TDM

What you need to know about multiple access channels

- Know what a multiple access (or shared, or broadcast) channel is.
- Know the main approaches to creating one, and their pros and cons
  - taking turns protocols are not required for exam purpose
- Know the pros and cons of each approach
- Know the details of CSMA/CD
  - E.g., if you are given a scenario with transmissions, you should be able to follow the protocol to figure out when collisions happen, how long a node will back off, and when a frame will finally be transmitted.
- Know the characteristics of wireless links and how they affect protocols designs
  - How does CSMA/CA work?
  - How is TCP affected?

Link layer: roadmap

- Link layer overview
  - Local Area Networks (LANs)
- Switched LANs
  - Ethernet and Addressing
  - Address Resolution Protocol (ARP)
  - Switches
- Virtual LANs (VLANs)
- Shared LANs and multiple access protocols
  - Channel partitioning
  - Random access
  - “taking turns”

Final Remarks on MAC address vs IP addresses

Q: Why use a separate set of addresses in the link layer?
A: ??

- 32-bit IP address:
  - *network-layer* address for interface
  - used for layer 3 (network layer) forwarding
  - e.g.: 128.119.40.136
- MAC (or LAN or physical or Ethernet) address:
  - function: used "locally" to get frame from one interface to another physically-adjacent interface (same subnet, in IP-addressing sense)
  - 48-bit MAC address (for most LANs) burned in the ROM of the interface hardware, also sometimes software settable
  - e.g.: 1A-2F-BB-76-09-AD

hexadecimal (base 16) notation (each "numeral" represents 4 bits)

What use a separate address space in the link layer?

- Network layer and link layer have different goals, hence different requirements
- Network layer: global connectivity
  - Need to aggregate addresses for interfaces close to each other to scale
  - So, IP addresses change when a device moves
- Link layer: local connectivity
  - Much smaller scale -- It is ok to have fixed “random” address for the interface
  - A fixed address makes it easier to bootstrap (we can still talk with the interface until it gets its IP address)
- Also, each local network can have its own way to forward traffic
  - And still be able to connect to different kinds of networks…
  - through IP, or any other network layer protocol that they all agree on.
- Any other thoughts?

Link layer: roadmap

- Link layer overview
  - Lo
- Swi
  - Et
  - Ad
  - Sw
- Virtu
- Mult
  - Ch
  - Ra
  - "taking turns"

We are done with the link layer!

Next Up: Naming and Addressing