TCP/IP Protocol Architecture Model

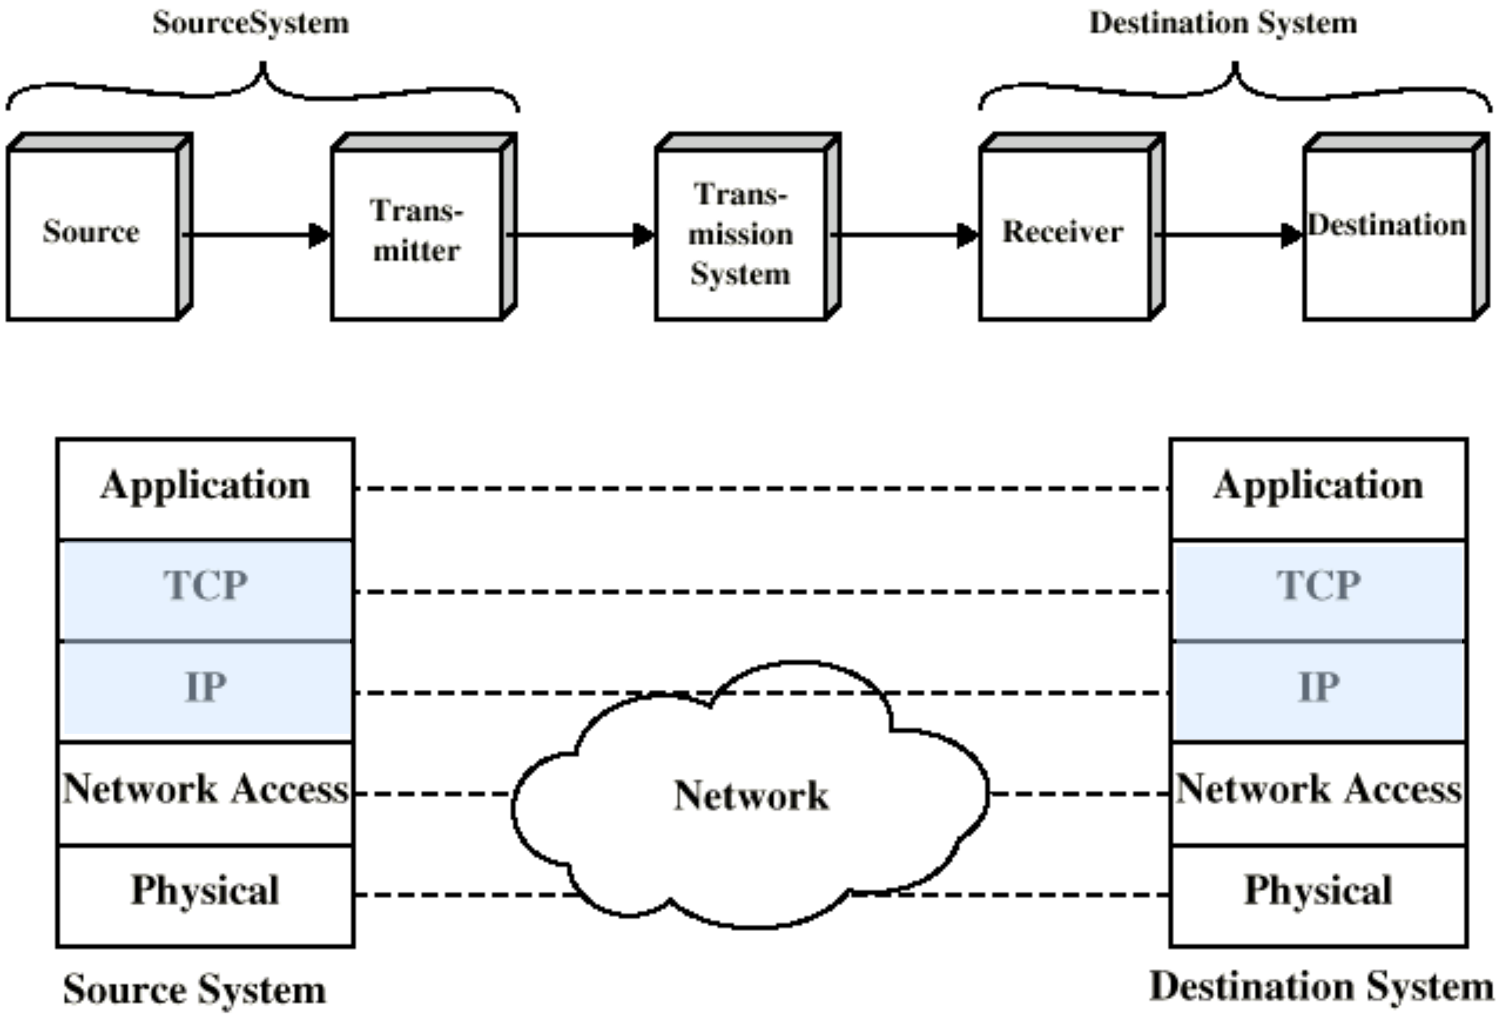

Internet Structure: Network of Networks

- A packet passes through *many* networks from source host to destination host

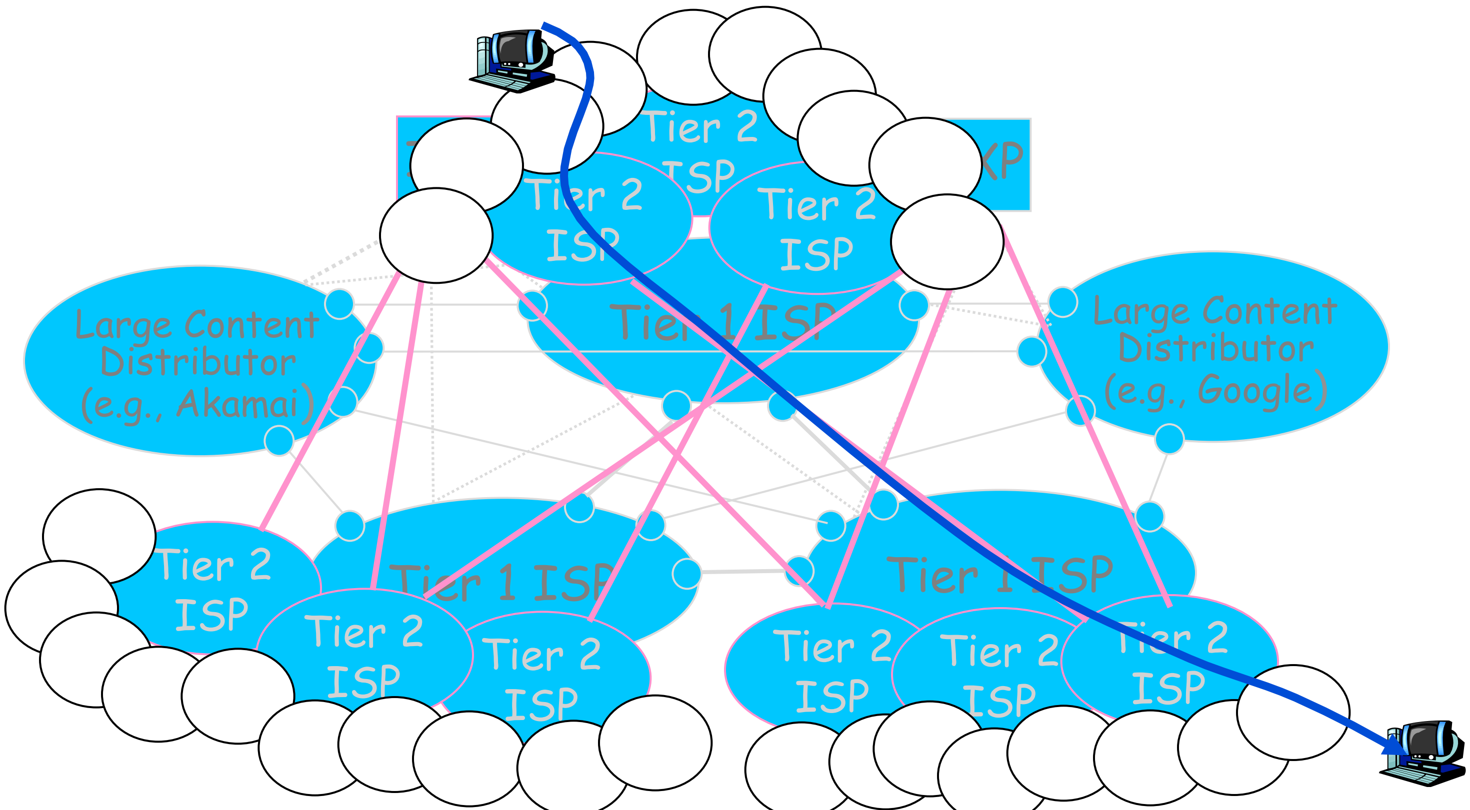

Review: Internet Structure

- A Network of Networks

Local and tier- 3 ISPs are *customers* of higher tier ISPs connecting them to rest of Internet

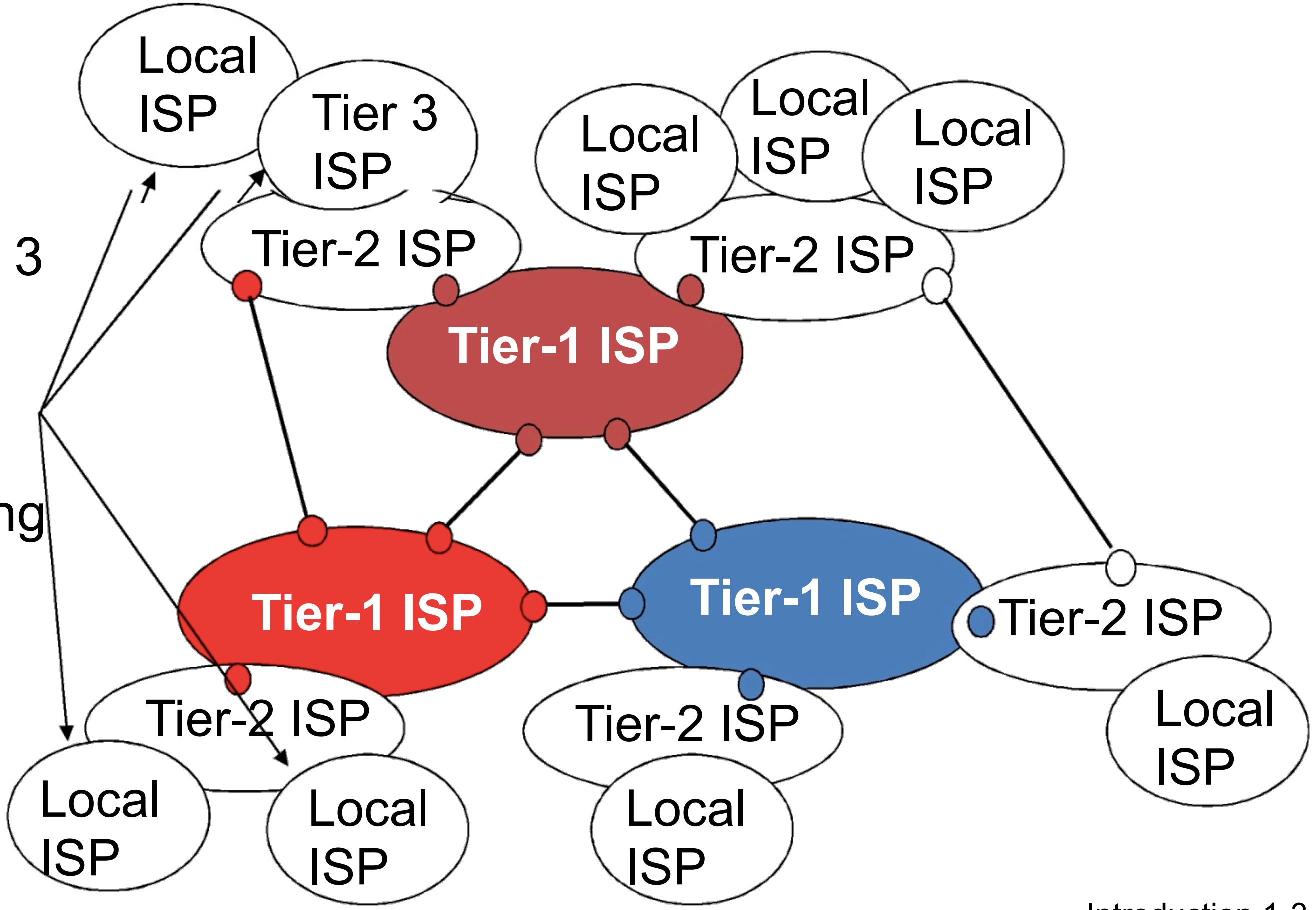

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

- end systems, access networks, links

1.3 Network core

- circuit switching, packet switching, network structure

1.4 Delay, loss and throughput in packet-switched networks

1.5 Protocol layers, service models

1.6 Networks under attack: security

How do Loss and Delay Occur?

- Packets *queue* in router buffers
  - packet arrival rate to link exceeds output link capacity
  - packets queue, wait for turn

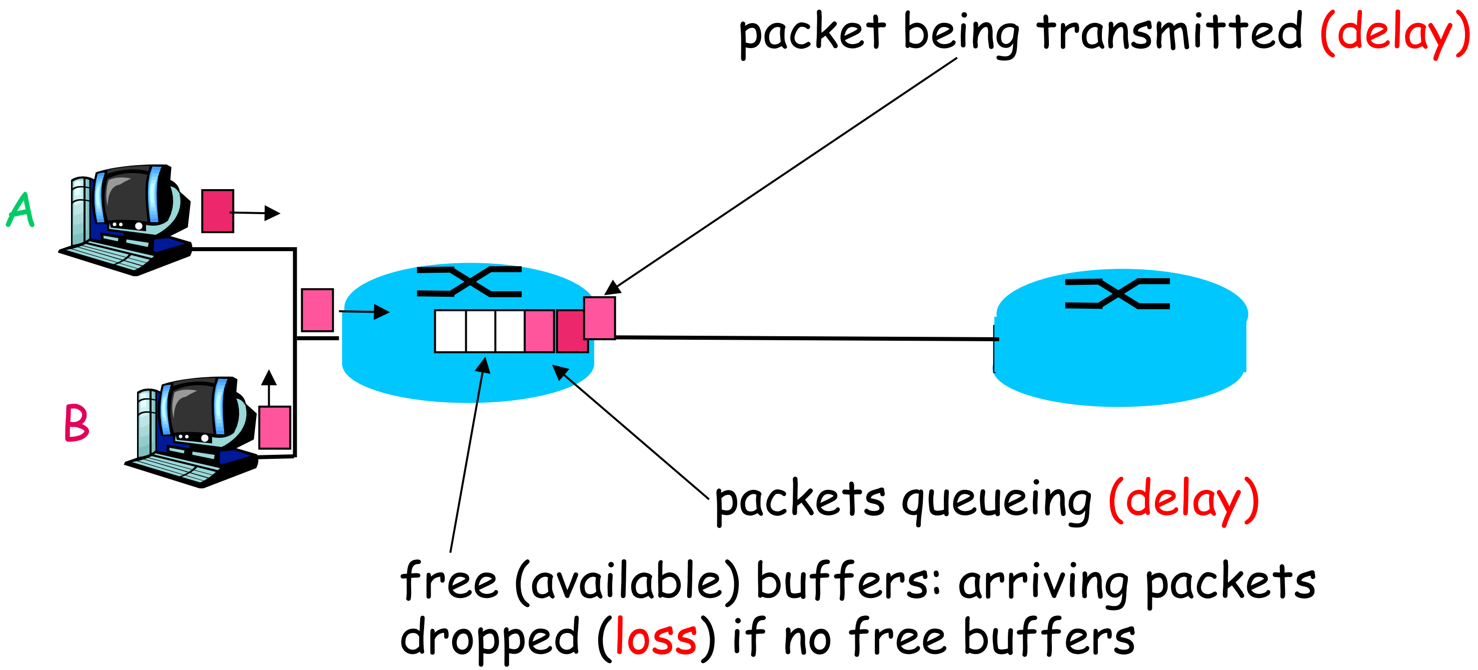

Four Sources of Packet Delay

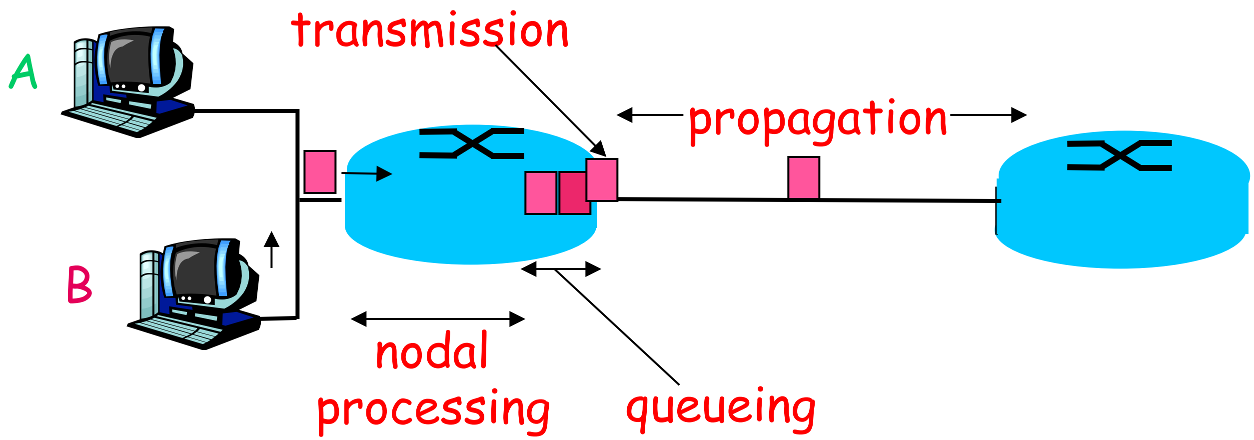

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{proc}: nodal processing

- check bit errors
- determine output link
- typically < msec

d_{queue}: queueing delay

- time waiting at output link for transmission
- depends on congestion level of router

Four Sources of Packet Delay

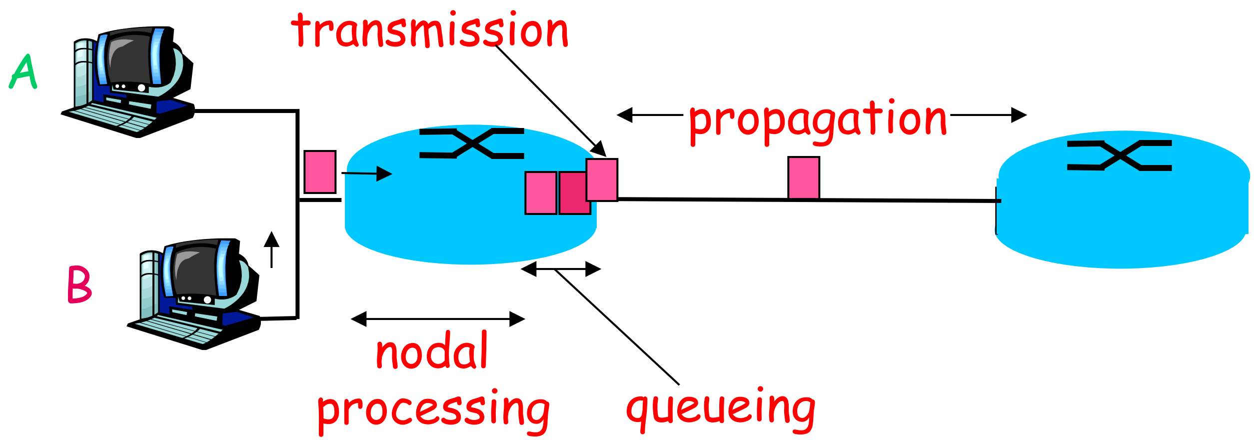

$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$

d_{trans}: transmission delay:

- L: packet length (bits)
- R: link bandwidth (bps)
- $d_{trans} = L/R$

d_{prop}: propagation delay:

- d: length of physical link
- s: propagation speed in medium (~2×10^8 m/sec)
- $d_{prop} = d/s$

d_{trans} and d_{prop} *very different*

Caravan Analogy

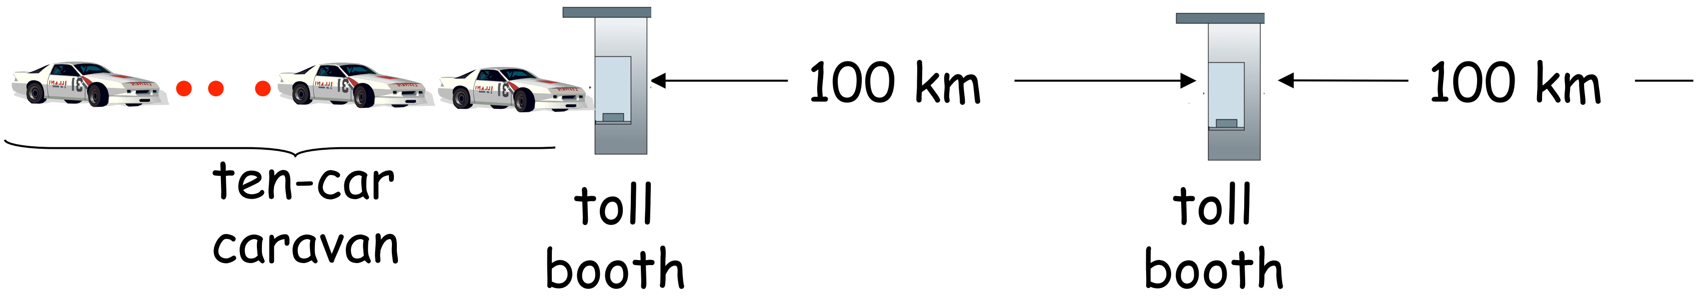

- Cars "propagate" at 100 km/hr
- Toll booth takes 12 sec to service car (transmission time)
- Car~bit; caravan ~ packet
- **Question**: How long until caravan is lined up before 2nd toll booth?

Caravan Analogy (More)

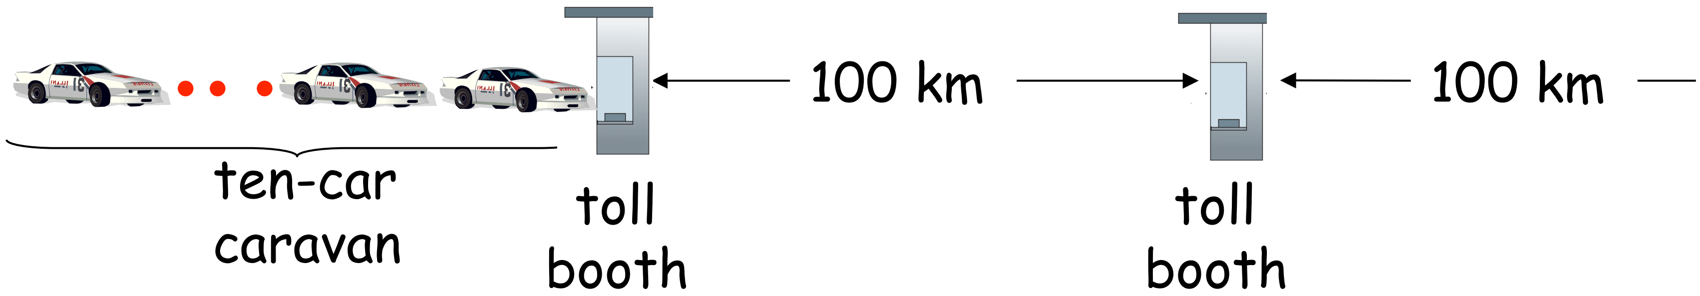

- Cars now "propagate" at 1000 km/hr
- Toll booth now takes 1 min to service a car
- ***Question:*** Will cars arrive to 2nd booth before all cars serviced at 1st booth?

Packet Loss

- Queue (aka buffer) preceding link in buffer has finite capacity
- Packet arriving to full queue dropped (aka lost)
- Lost packet may be retransmitted by previous node, by source end system, or not at all

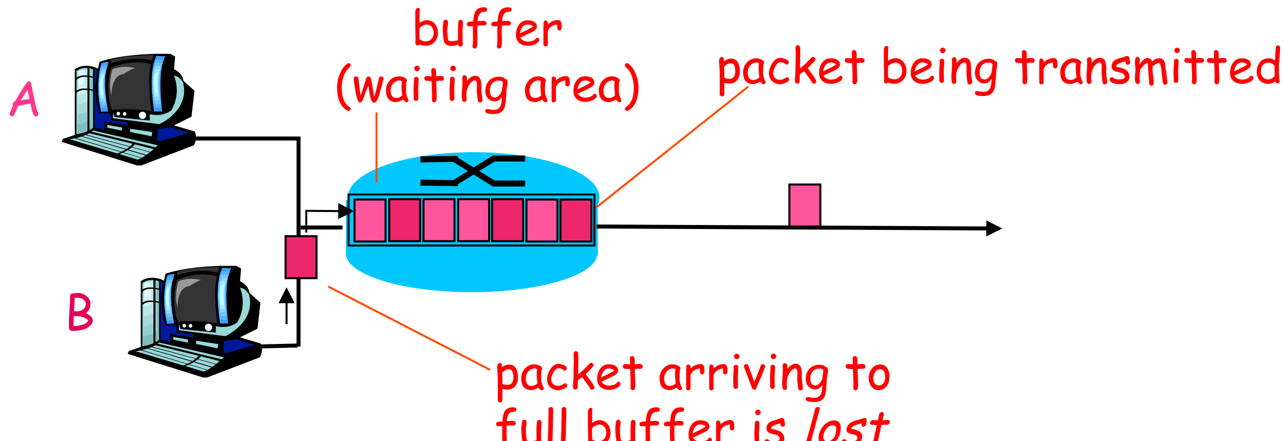

Throughput

- *Throughput:* rate (bits/time unit) at which bits transferred between sender/receiver
  - *instantaneous:* rate at given point in time
  - *average:* rate over longer period of time

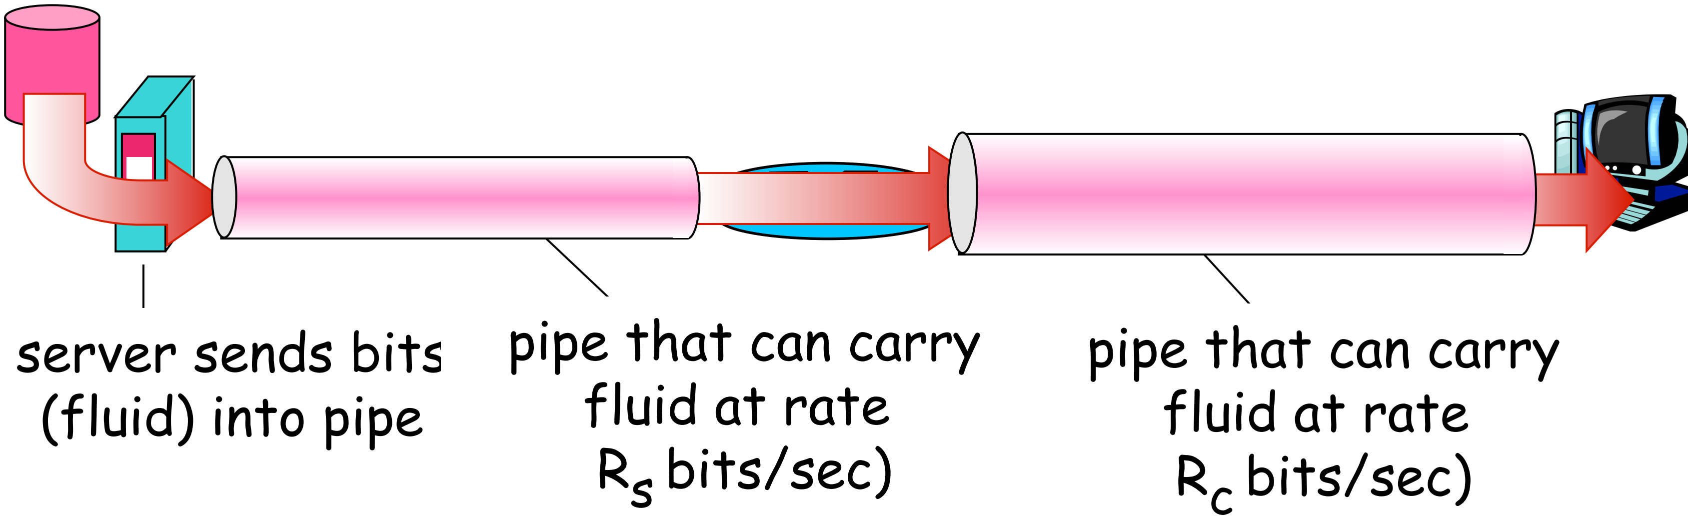

Throughput (more)

- $R_s < R_c$ What is average end-end throughput?

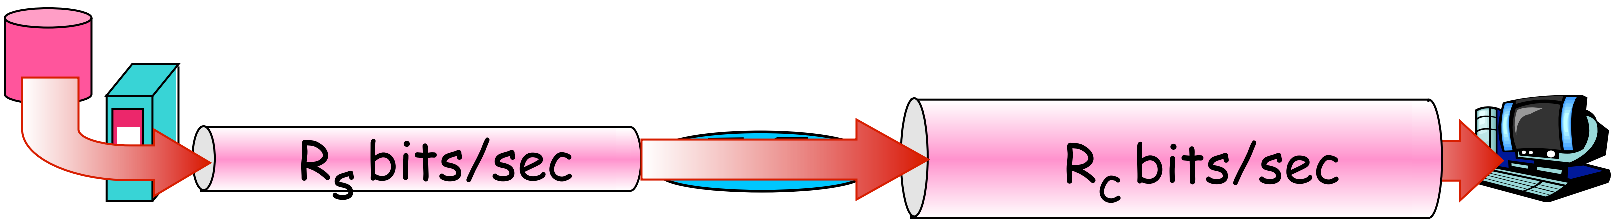

- $R_s > R_c$ What is average end-end throughput?

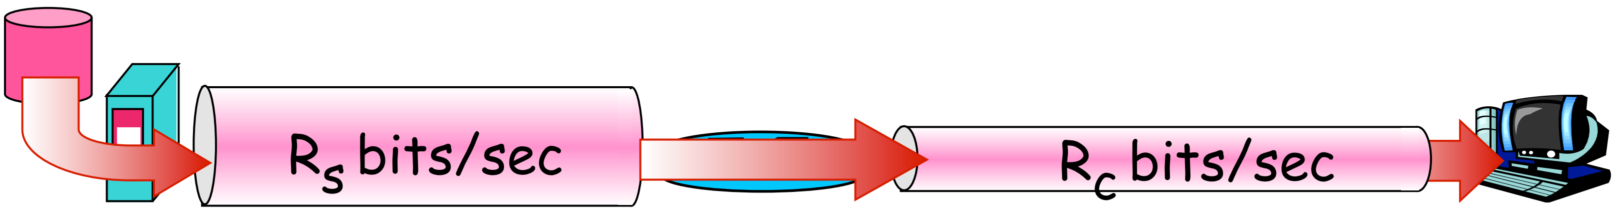

bottleneck link

Link on end-end path that constrains end-end throughput

Throughput: Internet Scenario

- Throughput: rate at which bits transferred between sender/ receiver

- **Question**: 10 connections (fairly) share backbone bottleneck ("goulot d'étranglement") link R bits/ sec.

What is the per-connection end-end throughput?

- min (R_c, R_s, R/10)
- in practice: R_c or R_s is often bottleneck

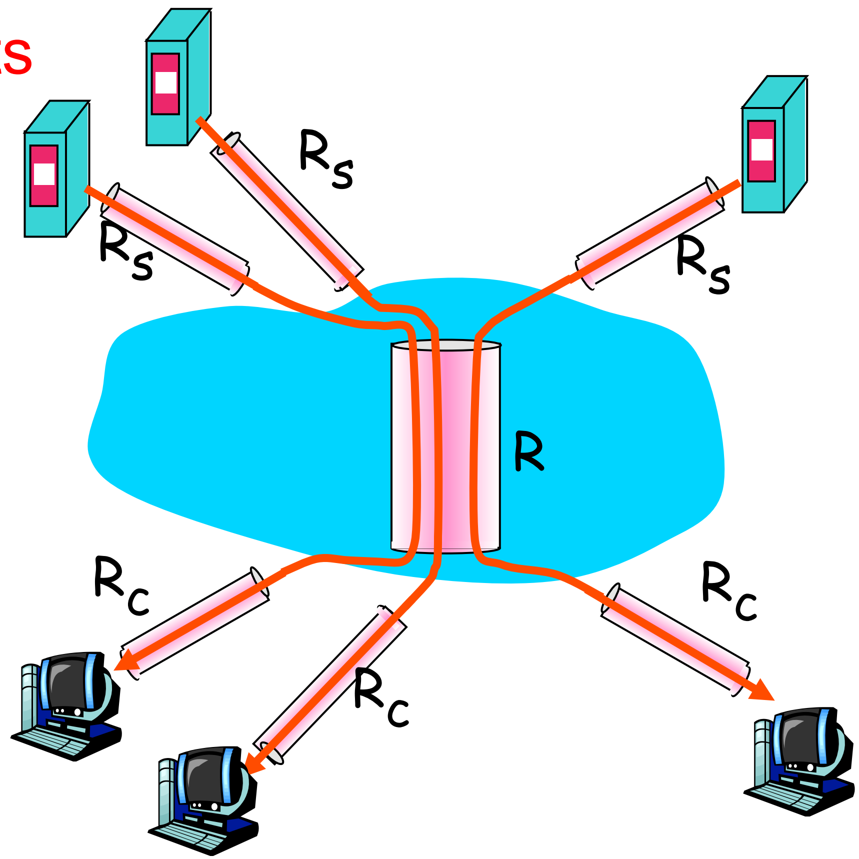

10 connections (fairly) share backbone bottleneck link R bits/sec

TCP Congestion Control

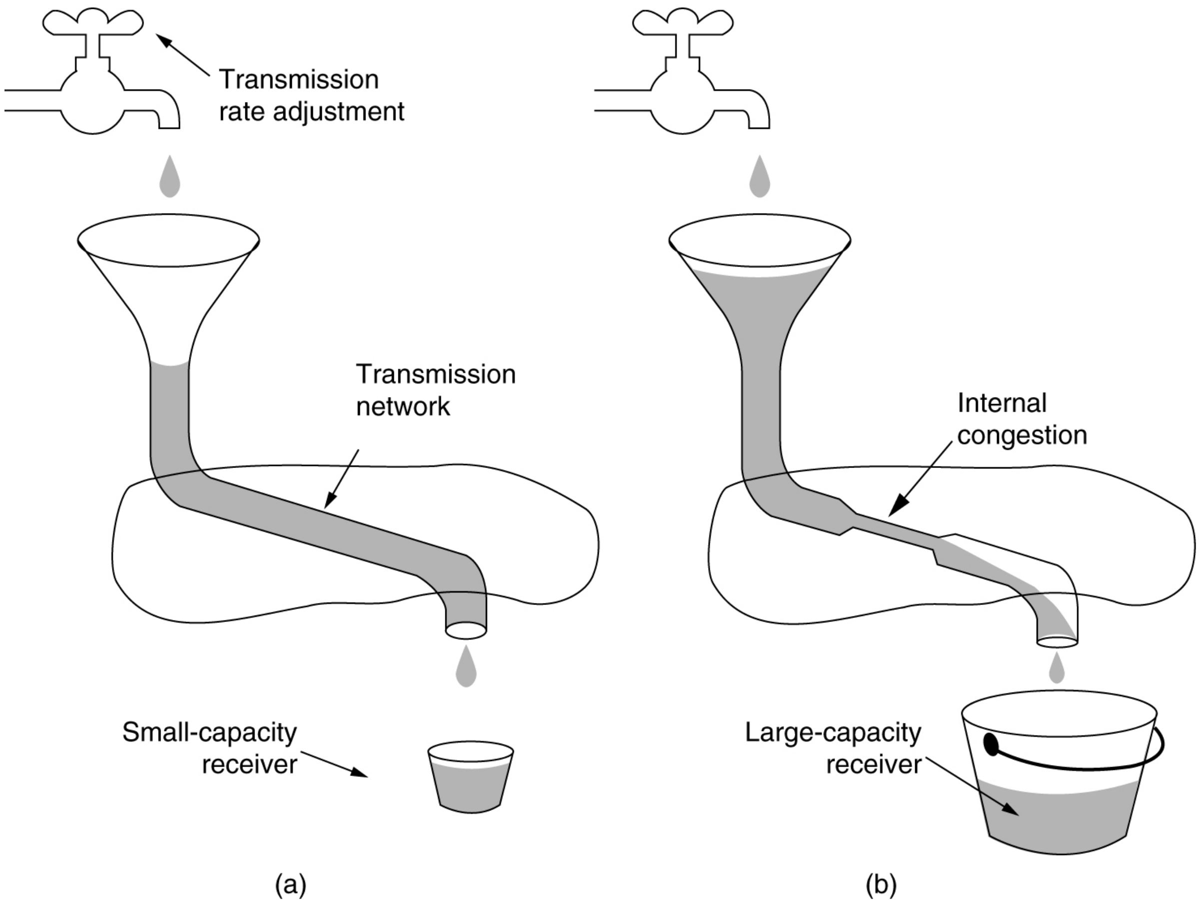

(a) A fast network feeding a low capacity receiver.

(b) A slow network feeding a high-capacity receiver.

TCP/IP Protocol Architecture Model

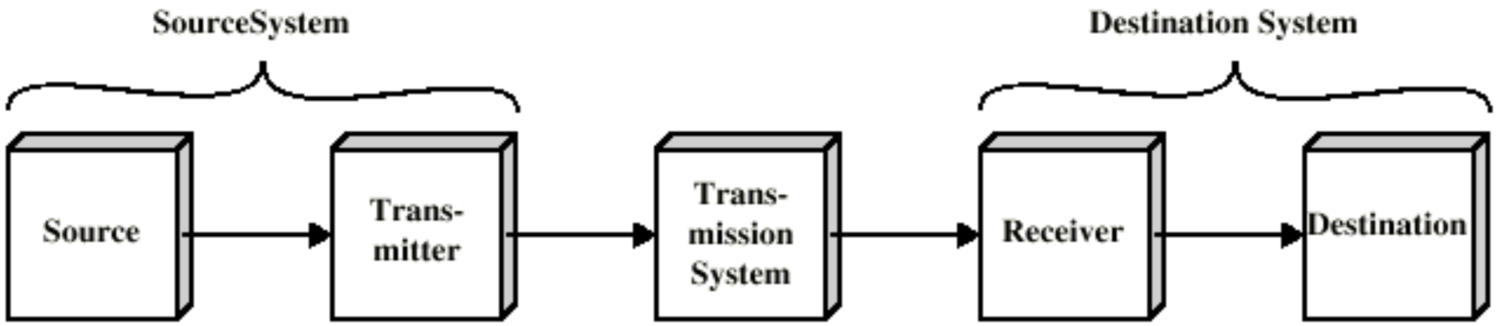

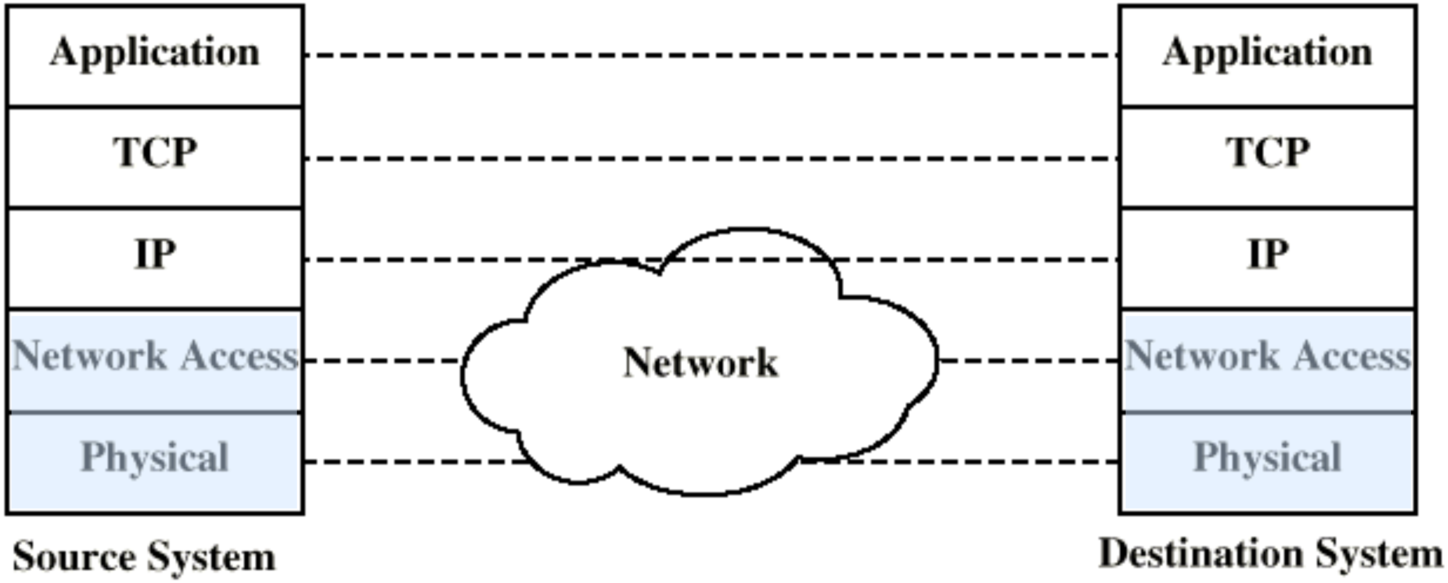

Medium Access Control Sublayer

Our goals understand principles behind :

- Channel Allocation Problem
- Multiple Access Protocols
- Ethernet
- Wireless LANs
- Broadband Wireless
- Bluetooth
- RFID
- Data Link Layer Switching

Revised: August 2011

Link Layer: Introduction

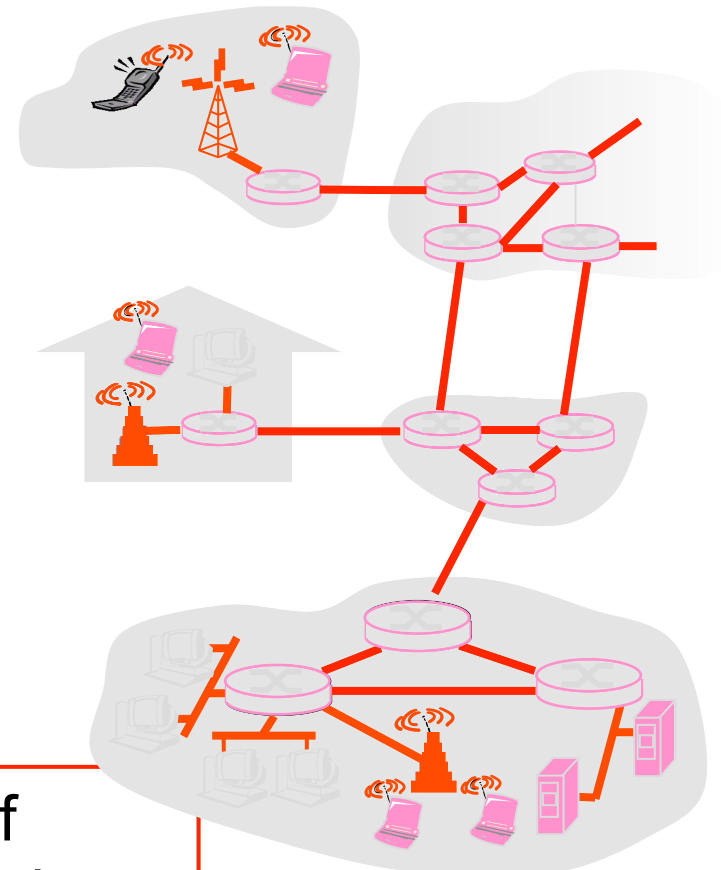

Terminology:

hosts and routers are nodes

communication channels that connect adjacent nodes along communication path are links

- wired links
- wireless links
- LANs

layer-2 packet is a frame, encapsulates datagram

data-link layer has responsibility of transferring datagram from one node to physically adjacent node over a link

Link layer: context

datagram transferred by different link protocols over different links:

- e.g., Ethernet on first link, frame relay on intermediate links, 802.11 on last link

each link protocol provides different services

- e.g., may or may not provide rdt over link

<u>transportation analogy</u>

trip from Princeton to Lausanne

- limo: Princeton to JFK
- plane: JFK to Geneva
- train: Geneva to Lausanne

tourist = datagram

transport segment = communication link

transportation mode = link layer protocol

travel agent = routing algorithm

Link Layer Services

framing, link access:

- encapsulate datagram into frame, adding header, trailer
- channel access if shared medium
- “MAC” addresses used in frame headers to identify source, dest
  - different from IP address!

reliable delivery between adjacent nodes

- we learned how to do this already
- seldom used on low bit-error link (fiber, some twisted pair)
- wireless links: high error rates

Link Layer Services (more)

flow control:

- pacing between adjacent sending and receiving nodes

error detection:

- errors caused by signal attenuation, noise.
- receiver detects presence of errors:
  - signals sender for retransmission or drops frame

error correction:

- receiver identifies *and corrects* bit error(s) without resorting to retransmission

half-duplex and full-duplex

- with half duplex, nodes at both ends of link can transmit, but not at same time

Half-duplex and Full-duplex

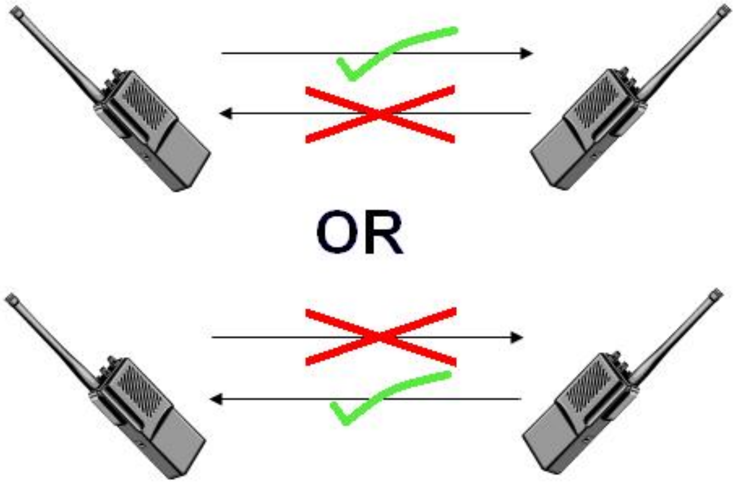

Half-Duplex

Full-Duplex

Where is the link layer implemented?

- in each and every host
- link layer implemented in "adaptor" (aka *network interface card* NIC)
  - Ethernet card, PCMCI card, 802.11 card
  - implements link, physical layer
- attaches into host's system buses
- combination of hardware, software, firmware

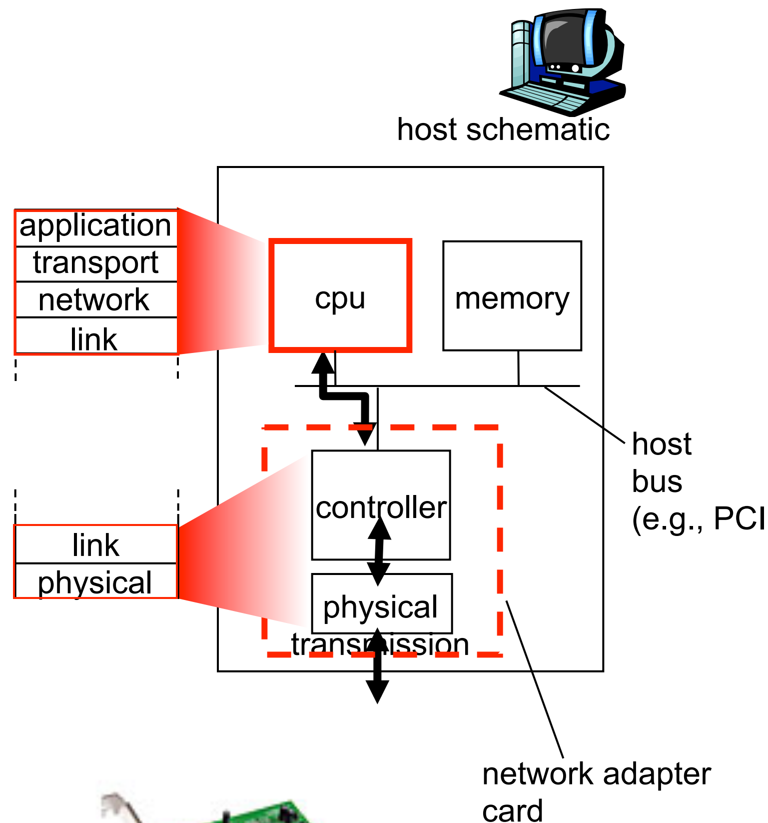

Adaptors Communicating

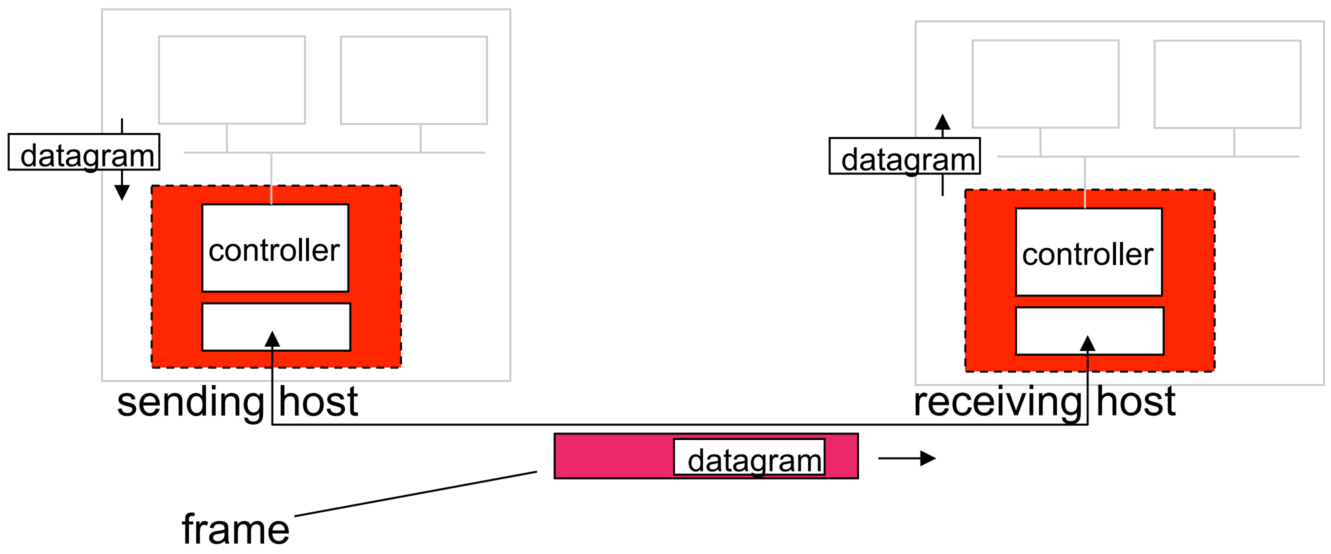

sending side:

- encapsulates datagram in frame
- adds error checking bits, flow control, etc.

receiving side

- looks for errors, flow control, etc
- extracts datagram, passes to upper layer at receiving side

The MAC Sublayer

Responsible for deciding who sends next on a multi-access link

- An important part of the link layer, especially for LANs

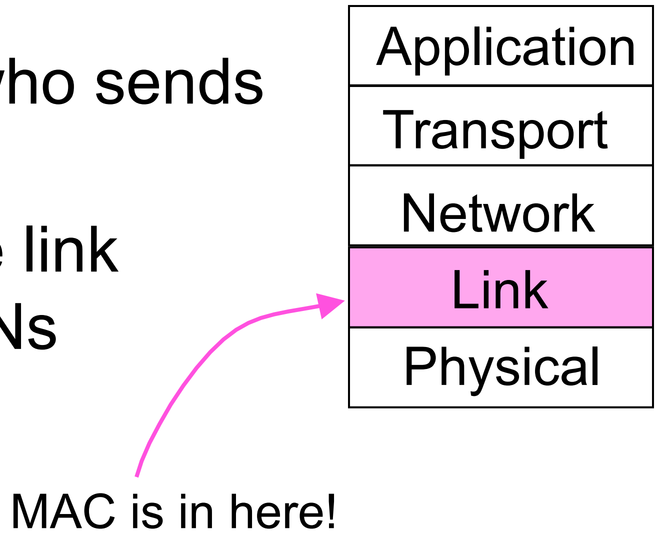

Channel Allocation Problem (1)

For fixed channel and traffic from N users

- Divide up bandwidth using TDM, CDMA, etc.
- This is a static allocation, e.g., FM radio

This static allocation performs poorly for bursty traffic

- Allocation to a user will sometimes go unused

- TDM = Time Division Multiplexing
- CDMA = Code Division Multiple Access

Channel Allocation Problem (3)

Dynamic allocation gives the channel to a user when they need it. Potentially N times as efficient for N users.

Schemes vary with assumptions:

Station Model: The model consists of ***N*** independent stations (e.g., computers, telephones or personal communicators) each with a program or user that generates frames for transmission. Once a frame has been generated, the station is blocked and does nothing until the frame has been successfully transmitted.

Single Channel Assumption: A single channel is available for all communication. All station can transmit on it and all can receive from it;

Collision assumption: If two frames are transmitted simultaneously, they overlap in time and the resulting signal is garbled. This event is called a collision;

Channel Allocation Problem (4)

Dynamic allocation gives the channel to a user when they need it. Potentially N times as efficient for N users.

Schemes vary with assumptions:

Continuous time: Frame transmission can begin at any instant. There is no master clock dividing time into discrete intervals;

Slotted time: Time is divided into discrete intervals (slots). Frame transmissions always begin at the start of a slot. A slot may contain 0, 1 or more frames, corresponding to an idle slot, a successful transmission, or a collision, respectively;

Carrier Sense: Stations can tell if the channel is in use before trying to use it. If the channel is sensed as busy, no station will attempt to use it until goes idle;

No carrier sense: Stations cannot sense the channel before trying to use it. They just go ahead and transmit. Only later they determine whether the transmission was successful.

Channel Allocation Problem (2)

Dynamic allocation gives the channel to a user when they need it. Potentially N times as efficient for N users.

Schemes vary with assumptions:

Assumption	**Implication**
N Independent traffic (station)	Often not a good model, but permits analysis
Single channel	No external way to coordinate senders
Observable collisions	Needed for reliability; mechanisms vary
Continuous or slotted time	Slotting may improve performance
Carrier sense	Can improve performance if available

Example: Congestion in M2M over LTE

The expected number of M2M / MTC devices until 2020 is approximately 20 billions.

This devices are going to be applied in a wide range of applications.

To make this “dream become true” they will need an access network for exchange data/information.

In this context the cellular networks represents a good alternative of access network.

But there is a little problem... Cellular networks were projected for humans, not for machines!

Example M2M: Cellular Networks

Used by mobile networks operators.

Designed for H2H and H2M types of communication.

Key features:

- Ubiquity
- Accessibility
- Security
- Designed for H2H communication

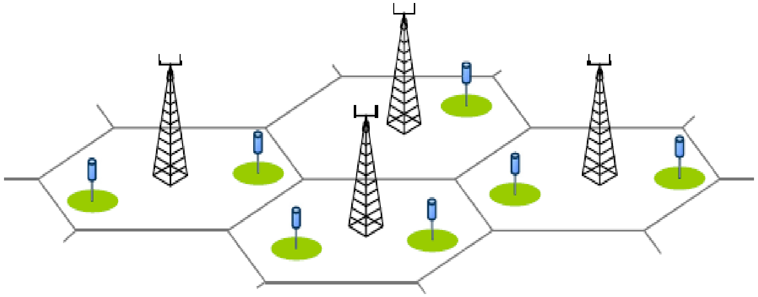

Technology:

GSM, UTMS, CDMA, LTE

Machine-to-Machine Communication

Machine-to-Machine (M2M) communications is a technology that enables one or more autonomous machines to communicate directly with one another without human intervention.

Its main characteristics are:

- Large number of simultaneously connected devices
- Small data volume transmissions
- Vastly diverse quality-of-service (QoS) requirements

Play important role on the Internet of Things (IoT)!

Example: Congestion in M2M over LTE

The congestion of MTC network usually happens in **radio Network** and **core network** because of mass concurrent signaling and data transmissions.

The Occurrence of congestion in LTE network:

- **Radio Access Network (RAN):** large number of devices requesting access to the network to enable \ modify \ disable a connection.

- **Core Network (CN):** excessive signaling flows or data (from several eNB) directed to the same element of the EPC (envolved Packet Core), for example, the S-GW and the MME (Mobility Management Entity).

Example: Congestion in M2M over LTE

Radio Access Network (RAN)

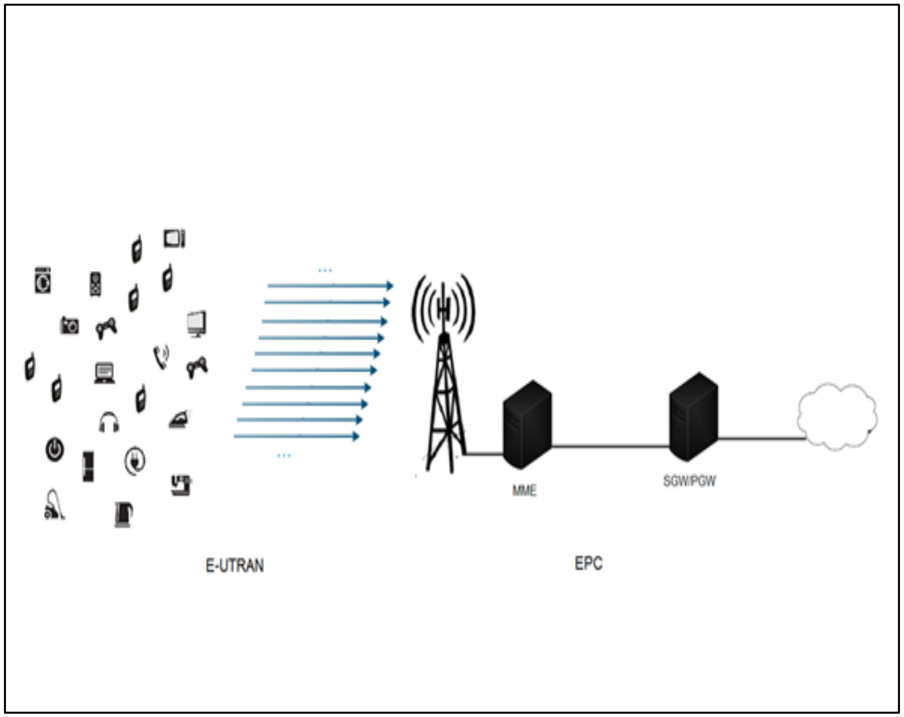

RAN - Congestion

Core Network (CN)

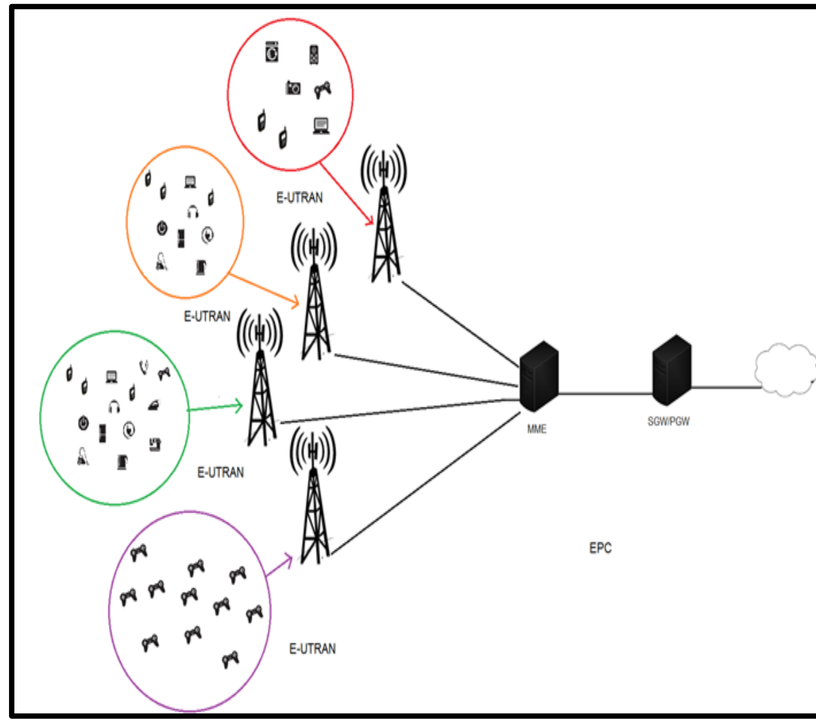

EPC - Congestion

LTE

Home Subscriber Server

Figure 2: Basic EPS architecture with E-UTRAN access

Channel Allocation Problem

Multiple Access Links and Protocols

Two types of "links":

point-to-point

- PPP for dial-up access
- point-to-point link between Ethernet switch and host

broadcast (shared wire or medium)

- old-fashioned Ethernet
- upstream HFC (Hybrid Fiber-Coaxial)
- 802.11 wireless LAN

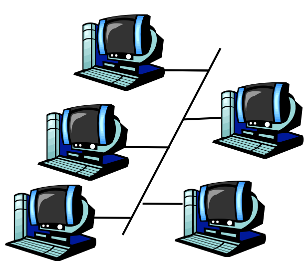

shared wire (e.g., cabled Ethernet)

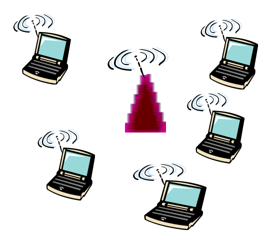

shared RF (Radio frequency) (e.g., 802.11 WiFi)

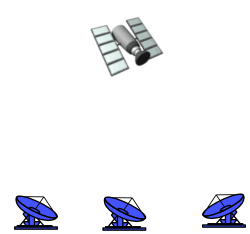

shared RF (satellite)

humans at a cocktail party (shared air, acoustical)

Hybrid Fiber-Coaxial

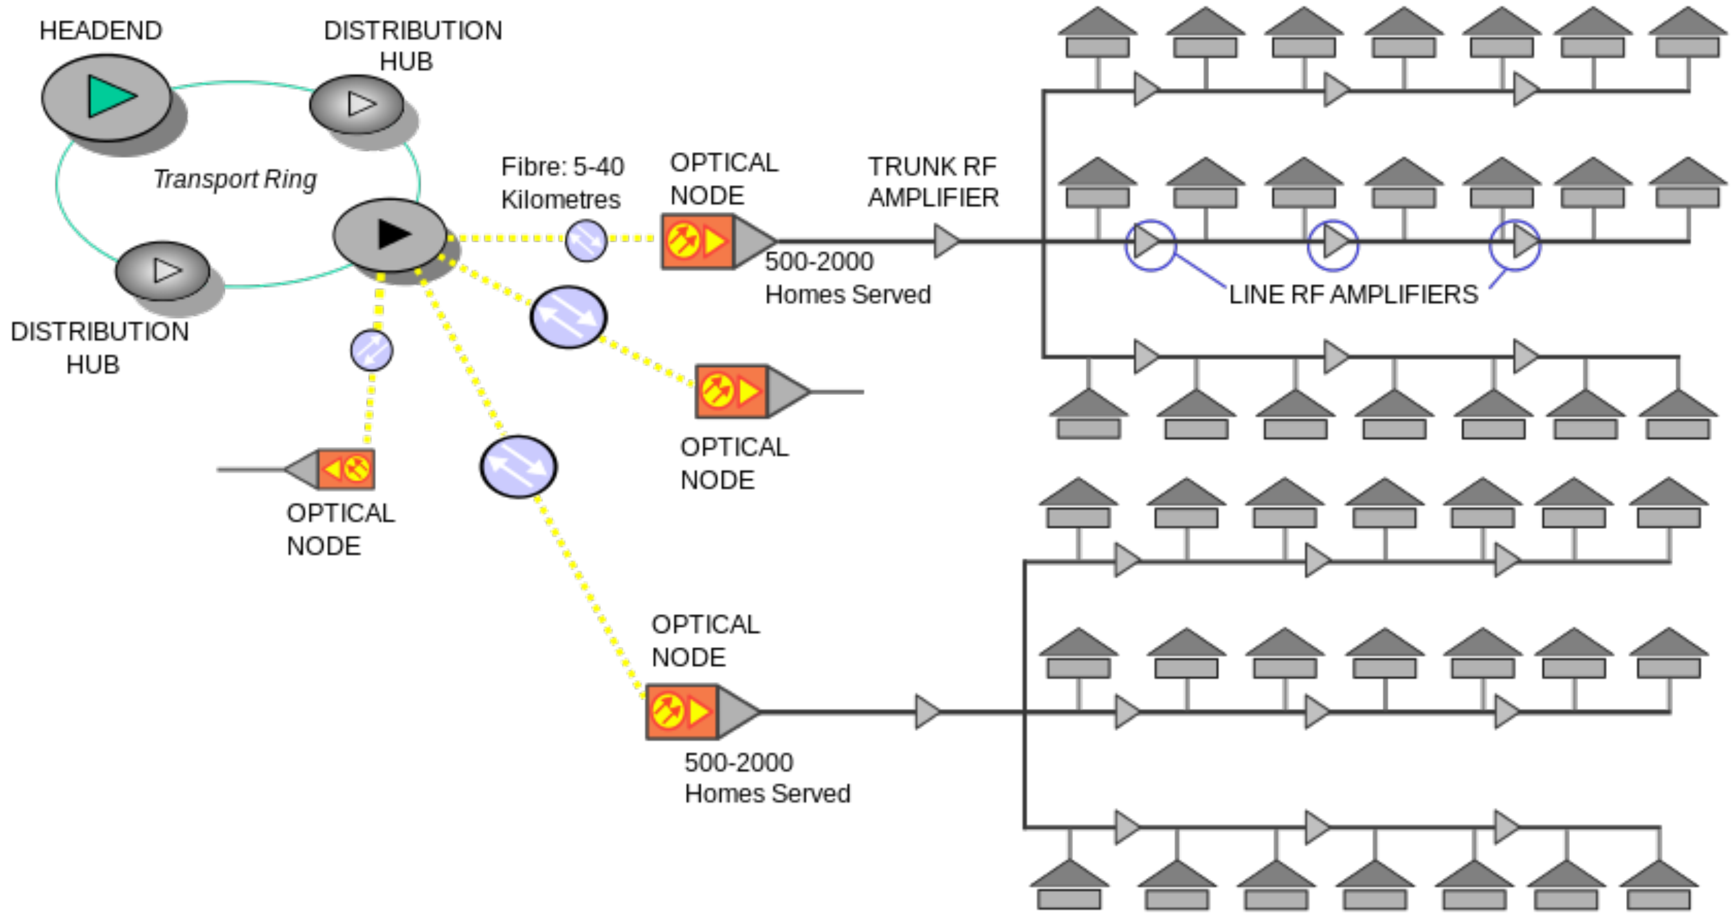

Multiple Access protocols

- single shared broadcast channel
- two or more simultaneous transmissions by nodes: 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

Ideal Multiple Access Protocol

<u>Broadcast channel of rate R bps</u>

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: a 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

Multiple Access Protocols

- ALOHA »
- CSMA (Carrier Sense Multiple Access) »
- Collision-free protocols »
- Limited-contention protocols »
- Wireless LAN protocols »

Channel Partitioning MAC protocols: TDMA

TDMA: Time Division Multiple Access

- access to channel in "rounds"
- each station gets fixed length slot (length = pkt trans time) in each round
- unused slots go idle
- example: 6-station LAN, 1,3,4 have pkt, slots 2,5,6 idle

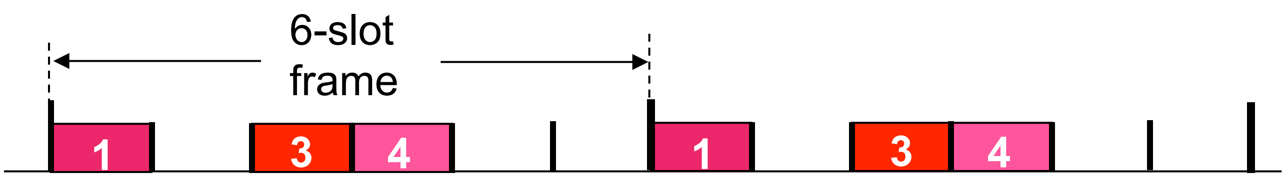

Channel Partitioning MAC protocols: FDMA

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 pkt, frequency bands 2,5,6 idle

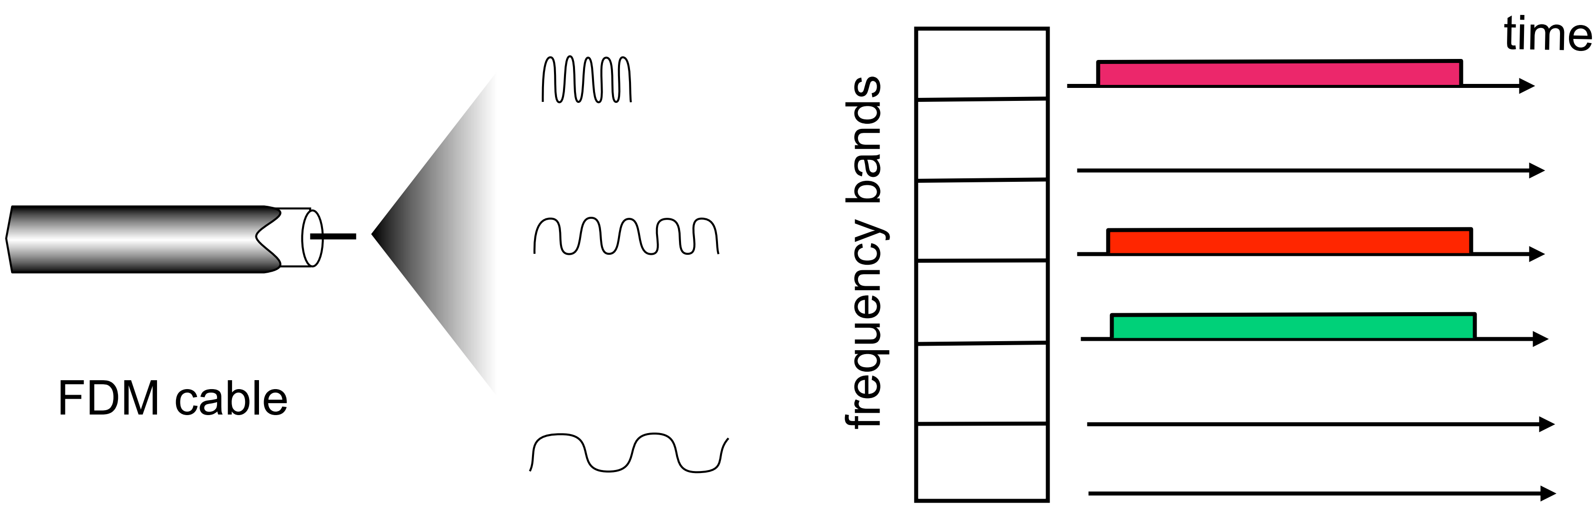

Channel Allocation Problem (1)

For fixed channel and traffic from N users

- Divide up bandwidth using TDM, CDMA, etc.
- This is a static allocation, e.g., FM radio

This static allocation performs poorly for bursty traffic

- Allocation to a user will sometimes go unused

- TDM = Time Division Multiplexing
- CDMA = Code Division Multiple Access

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 MAC 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

ALOHA (1)

In pure ALOHA, users transmit frames whenever they have data; users retry after a random time for collisions

- Efficient and low-delay under low load

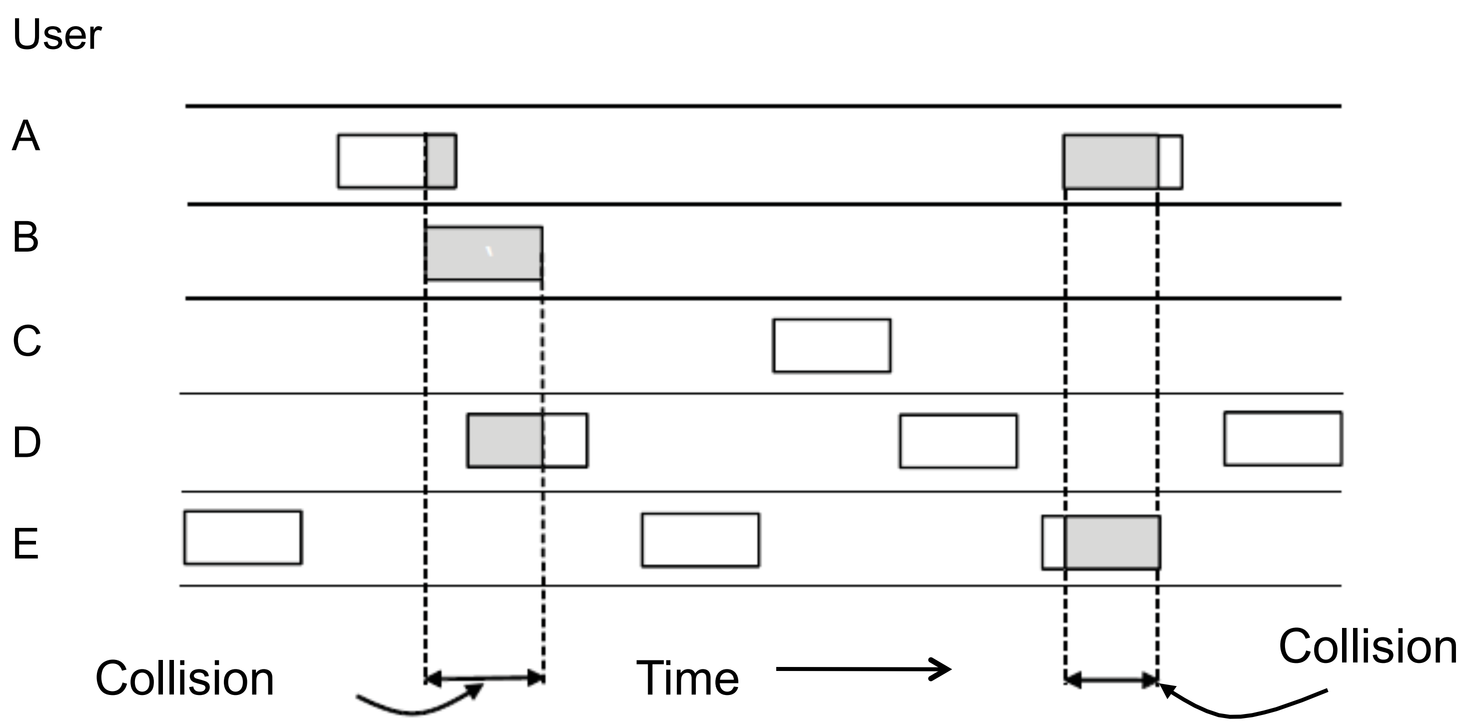

ALOHA (2)

Under what conditions will the shaded frame arrive undamaged?

- Let t be the time required to send a frame
- If any other user has generated a frame between time t_0 and $t_0 + t$, the end of that frame will collide with the beginning of the shaded one
- Similarly, any other frame started between $t_0 + t$ and $t_0 + 2t$ will bump into the end of the shaded frame
- Since the pure ALOHA a station does not listen to the channel before transmitting, it has no way of knowing that another frame was already underway

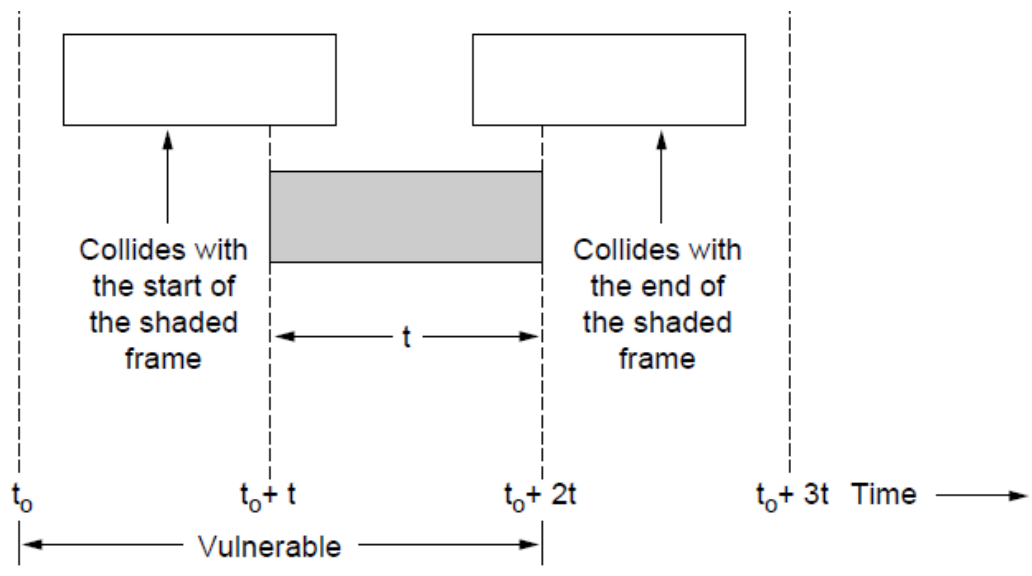

ALOHA (2)

Collisions happen when other users transmit during a vulnerable period that is **twice the frame time**

- Synchronizing senders to slots can reduce collisions

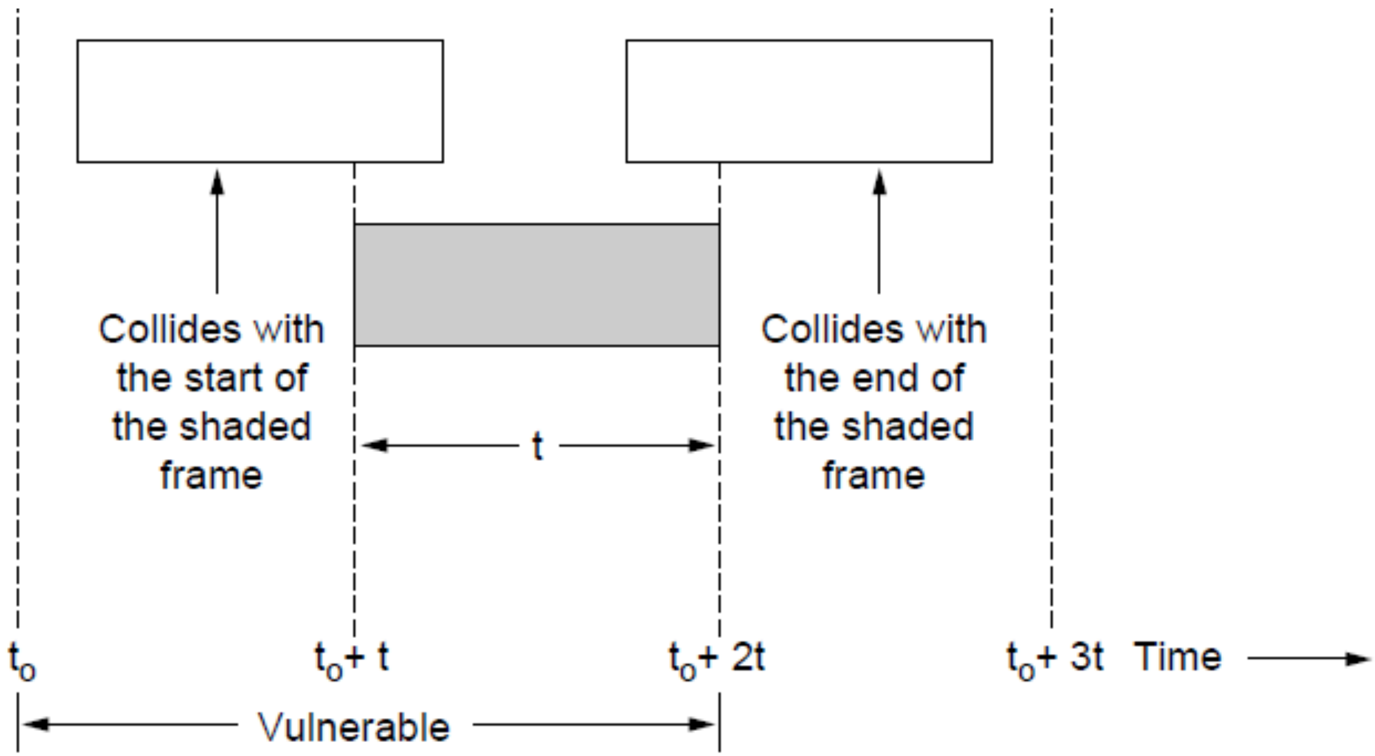

Pure (unslotted) ALOHA

unslotted Aloha: simpler, no synchronization

when frame first arrives

- transmit immediately

collision probability increases:

- frame sent at t_0 collides with other frames sent in $[t_0-1, t_0+1]$

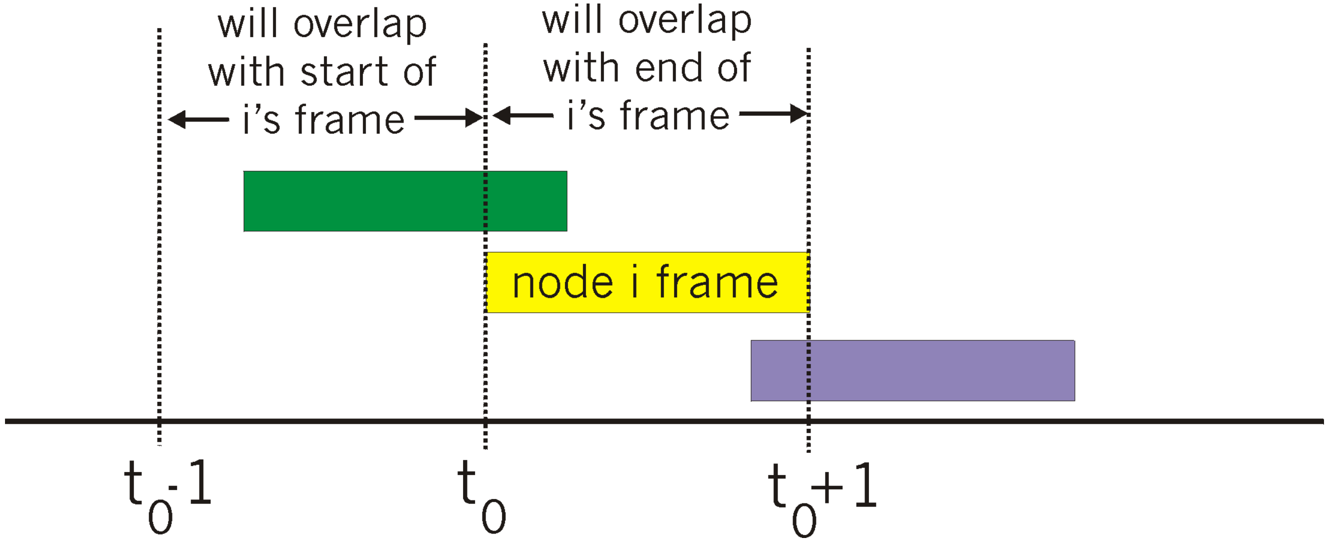

Slotted ALOHA

Assumptions:

- all frames same size
- time divided into equal size slots (time to transmit 1 frame)
- nodes start to transmit only slot beginning
- nodes are synchronized
- if 2 or more nodes transmit in slot, all nodes detect collision

Operation:

when node obtains fresh frame, transmits in next slot

- *if no collision:* node can send new frame in next slot
- *if collision:* node retransmits frame in each subsequent slot with prob. p until success

Slotted ALOHA

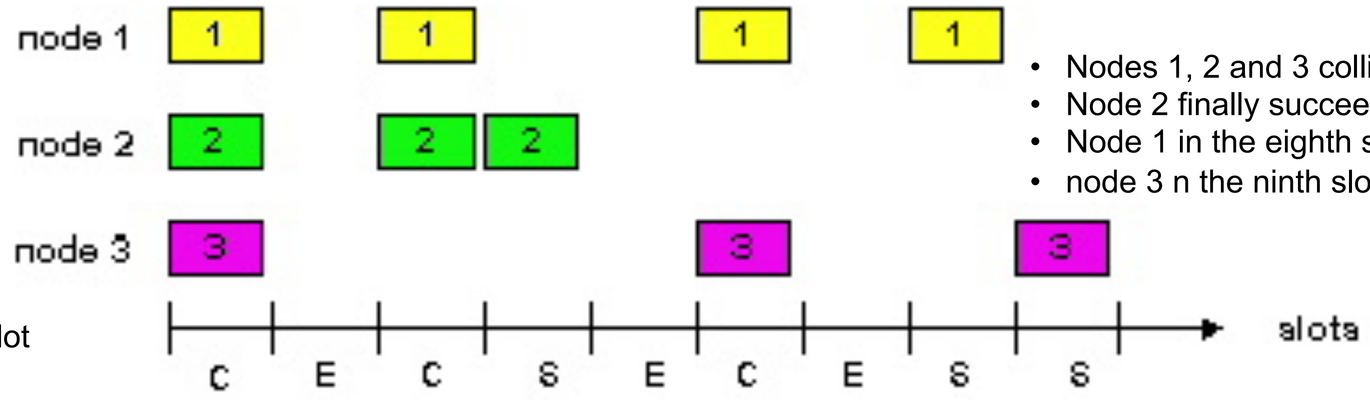

- Nodes 1, 2 and 3 collide in the first slot.
- Node 2 finally succeeds in the fourth slot,
- Node 1 in the eighth slot, and
- node 3 n the ninth slot

Legends:
C = Collision slot
E = Empty slot
S = Successful slot

Pros

1. single active node can continuously transmit at full rate of channel
2. highly decentralized: only slots in nodes need to be in sync
3. simple

Cons

1. collisions, wasting slots
2. idle slots
3. clock synchronization

Slotted Aloha efficiency

Efficiency : long-run fraction of successful slots (many nodes, all with many frames to send)

suppose: N nodes with many frames to send, each transmits in slot with probability *p*

prob that given node has success in a slot = $p(1-p)^{N-1}$

prob that *any* node has a success = $Np(1-p)^{N-1}$

max efficiency: find p* that maximizes $Np(1-p)^{N-1}$

for many nodes, take limit of $Np^*(1-p^*)^{N-1}$ as N goes to infinity, gives:

Max efficiency = 1/e = .37

At best: channel used for useful transmissions 37% of time!

!

Pure Aloha efficiency

P(success by given node) = P(node transmits) ·

P(no other node transmits in $[p_0-1,p_0]$ ·

P(no other node transmits in $[p_0-1,p_0]$

$= p \cdot (1-p)^{N-1} \cdot (1-p)^{N-1}$

$= p \cdot (1-p)^{2(N-1)}$

… choosing optimum p and then letting n -> infty ...

= 1/(2e) = .18

At best: channel used for useful transmissions 18% of time!

!

even worse than slotted Aloha!

ALOHA (3)

Slotted ALOHA is twice as efficient as pure ALOHA

- Low load wastes slots, high loads causes collisions
- Efficiency up to 1/e (37%) for random traffic models

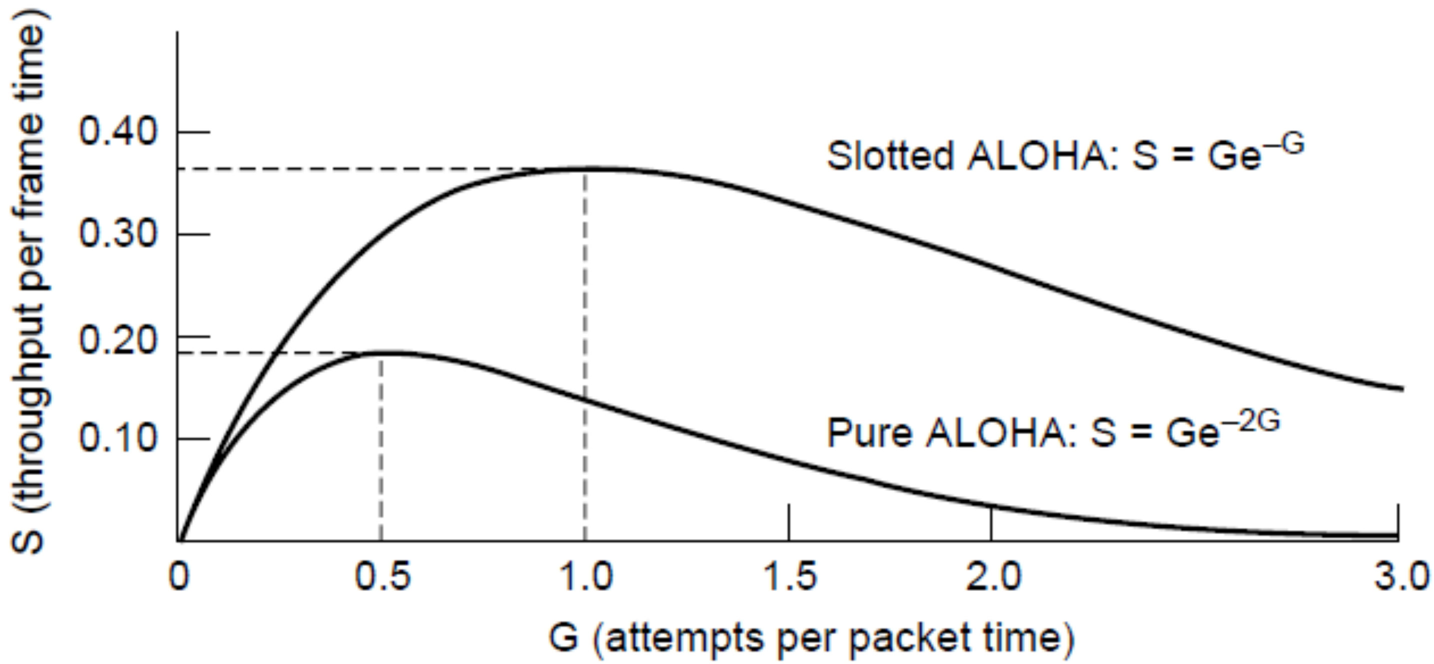

MAC Protocols: a taxonomy

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

CSMA (1)

CSMA improves on ALOHA by sensing the channel!

- User doesn't send if it senses someone else

Variations on what to do if the channel is busy:

- **1-persistent** (greedy) sends as soon as idle
- **Nonpersistent** waits a random time then tries again
- **p-persistent** sends with probability p when idle

CSMA (Carrier Sense Multiple Access)

CSMA: listen before transmit:

1. If channel sensed idle: transmit entire frame

2. If channel sensed busy, defer transmission: waits ("backs off") a random amount of time and then senses the channel.

- human analogy: don't interrupt others!

CSMA collisions

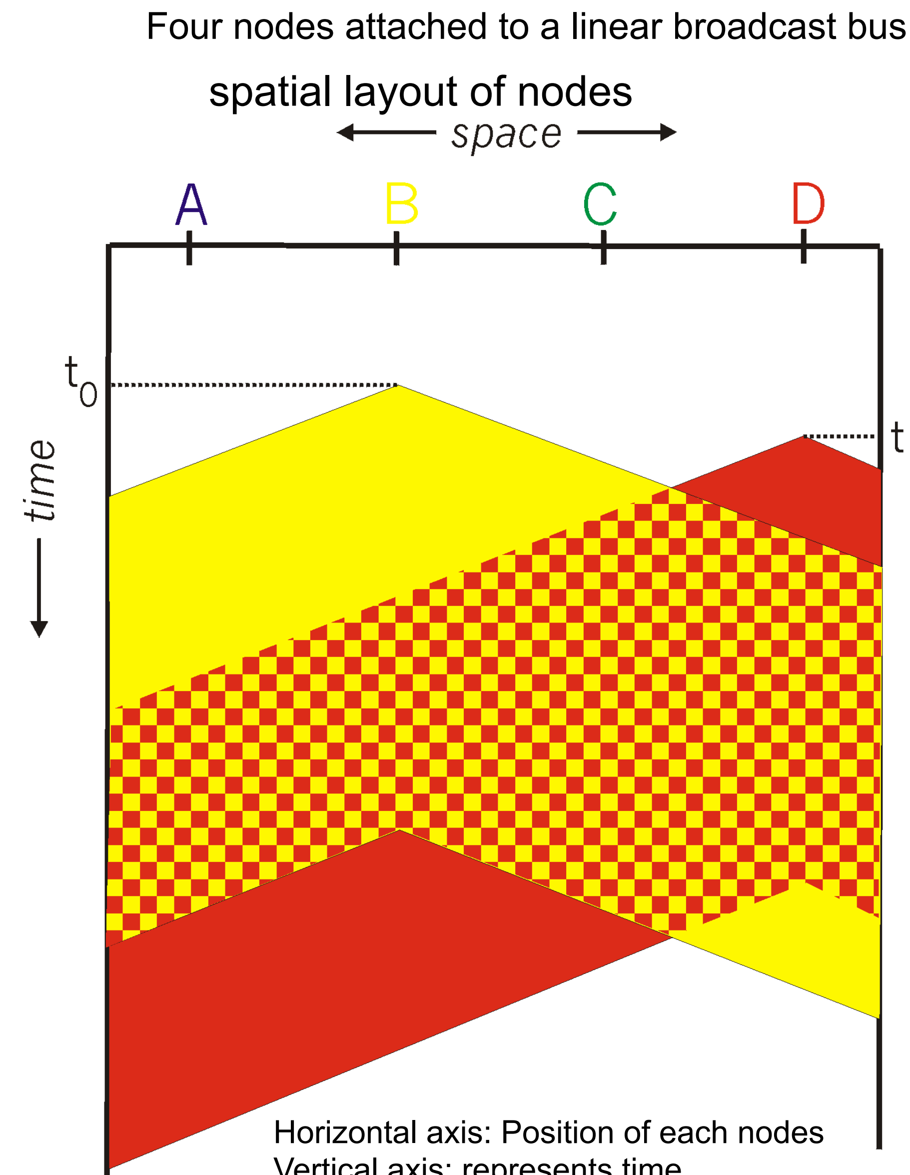

collisions can still occur:
propagation delay means two nodes may not hear each other's transmission

collision:
entire packet transmission time wasted

note:
role of distance & propagation delay in determining collision probability

CSMA/CD (Collision Detection)

CSMA/CD: carrier sensing, deferral as in CSMA

- If someone else begins talking at the same time, stop talking
- collisions *detected* within short time
- colliding transmissions aborted, reducing channel wastage

collision detection:

- easy in wired LANs: measure signal strengths, compare transmitted, received signals
- difficult in wireless LANs: received signal strength overwhelmed by local transmission strength

human analogy: the polite conversationalist

CSMA/CD collision detection

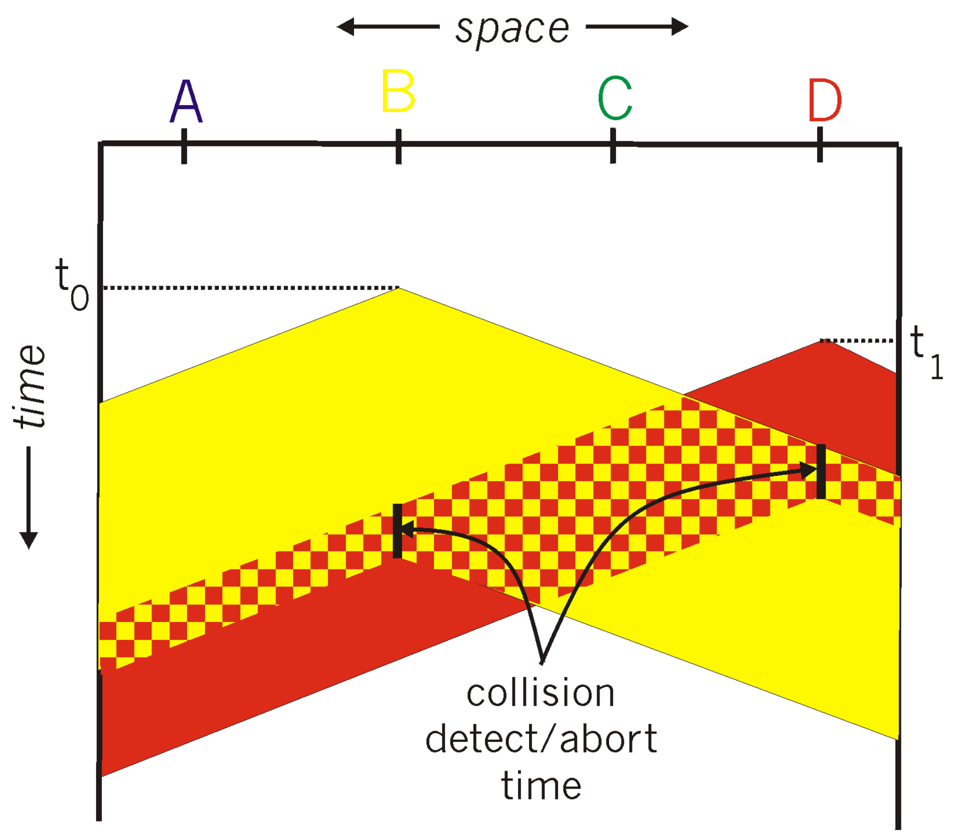

CSMA (2) – Persistence

CSMA outperforms ALOHA, and being less persistent is better under high load

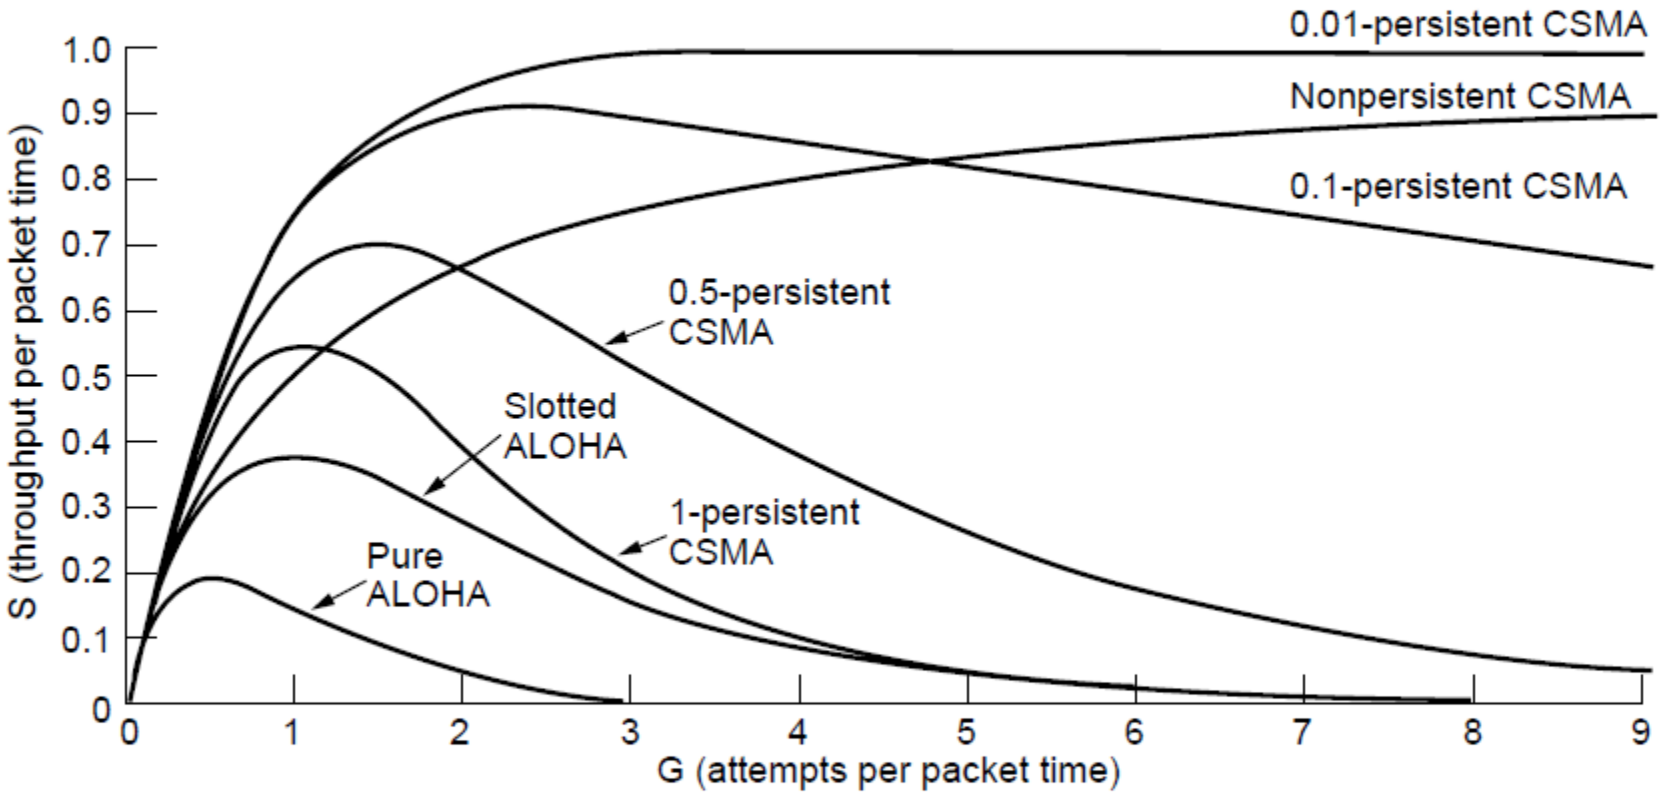

CSMA (3) – Collision Detection

CSMA/CD improvement is to detect/abort collisions

- Reduced contention times improve performance

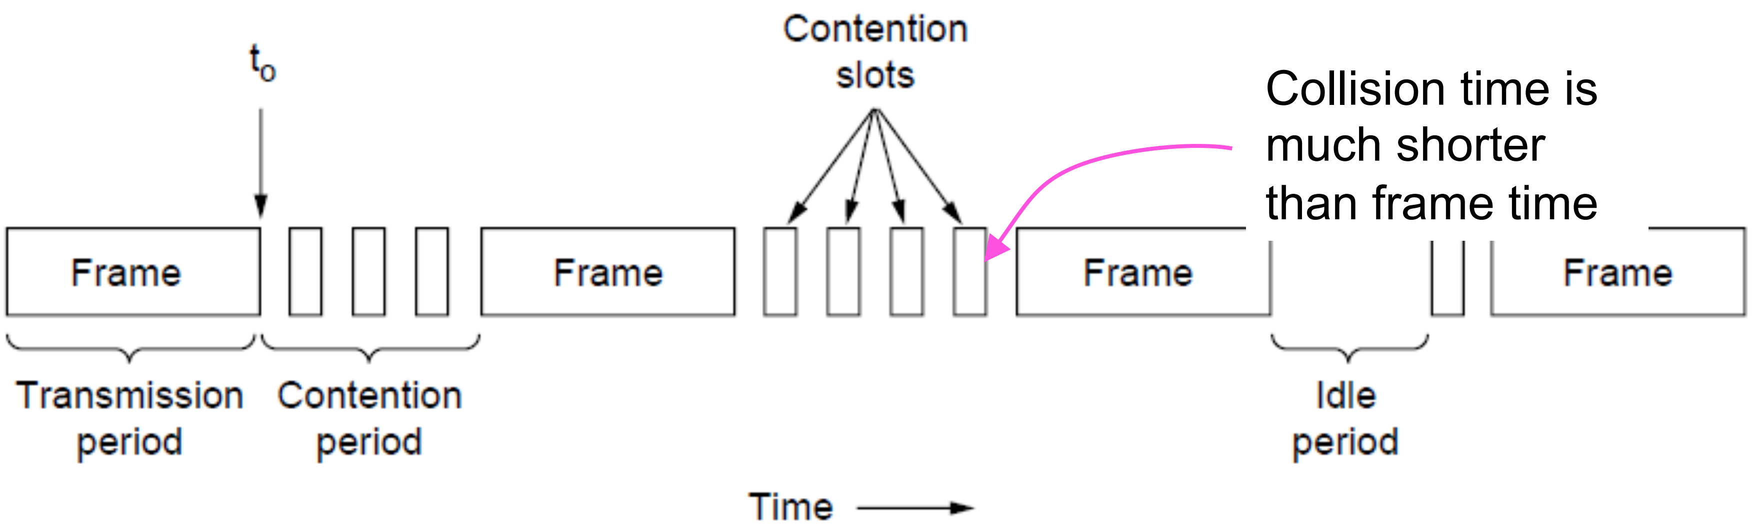

CSMA/CD can be in one of three states: contention, transmission, or idle

Summary of MAC protocols

channel partitioning, by time, frequency or code

- Time Division, Frequency Division

random access (dynamic),

- ALOHA, S-ALOHA, CSMA, CSMA/CD
- carrier sensing: easy in some technologies (wire), hard in others (wireless)
- CSMA/CD used in Ethernet
- CSMA/CA used in 802.11

MAC Protocols: a 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

"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:

master node "invites" slave nodes to transmit in turn

typically used with "dumb" slave devices

concerns:

- polling overhead
- latency
- single point of failure (master)

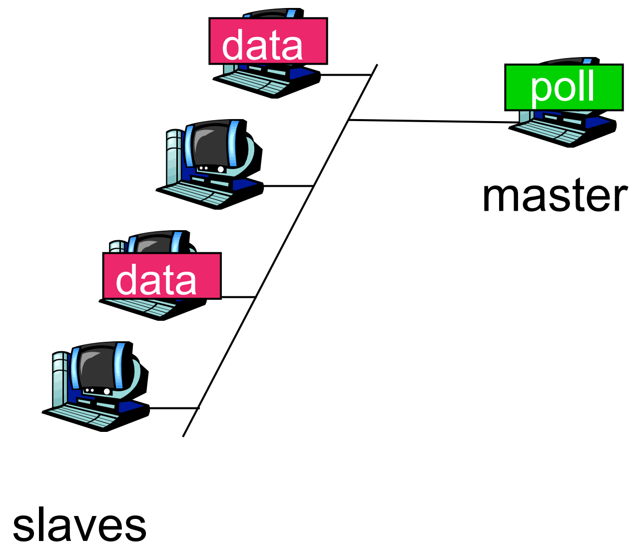

Collision-Free (1) – Bitmap

Collision-free protocols avoid collisions entirely

- Senders must know when it is their turn to send

The basic bit-map protocol:

- Sender set a bit in contention slot if they have data
- Senders send in turn; everyone knows who has data

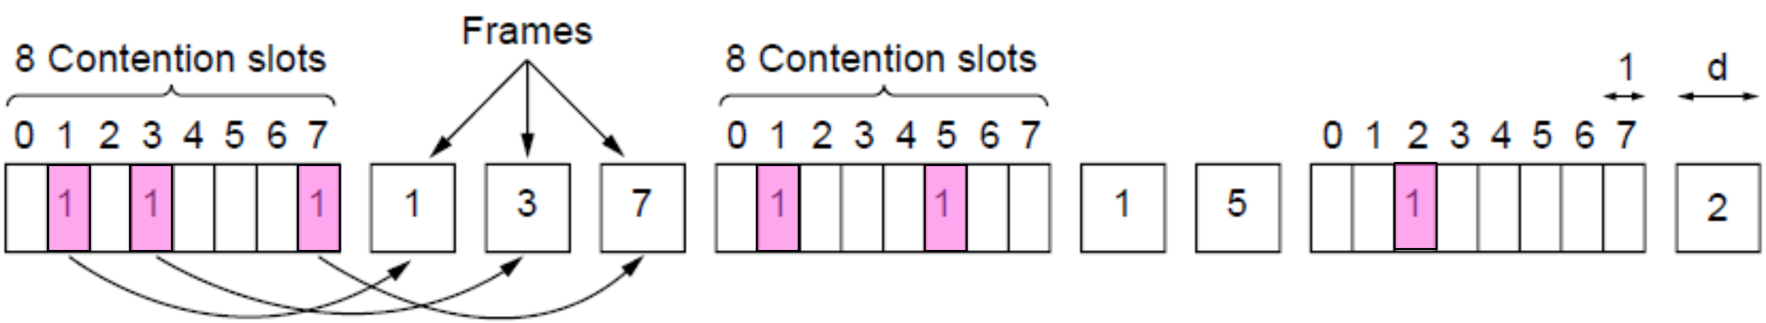

Collision-Free (2) – Token Ring

Token sent round ring defines the sending order

- Station with token may send a frame before passing
- Idea can be used without ring too, e.g., token bus

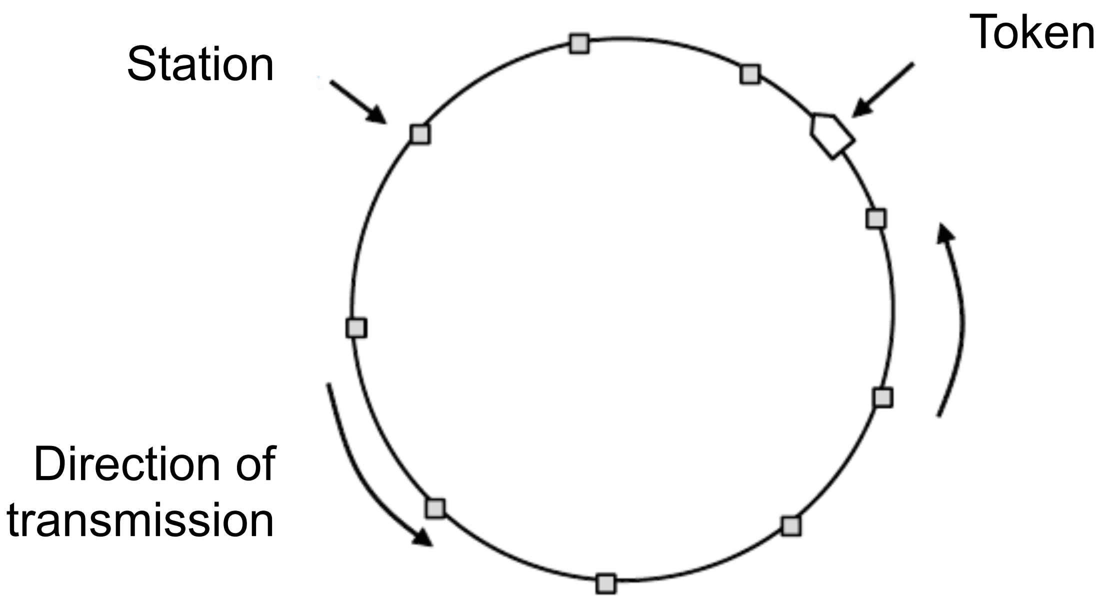

Collision-Free (3) – Countdown

Binary countdown improves on the bitmap protocol

- Stations send their address in contention slot (log N bits instead of N bits)
- Medium ORs bits; stations give up when they send a "0" but see a "1"
- Station that sees its full address is next to send

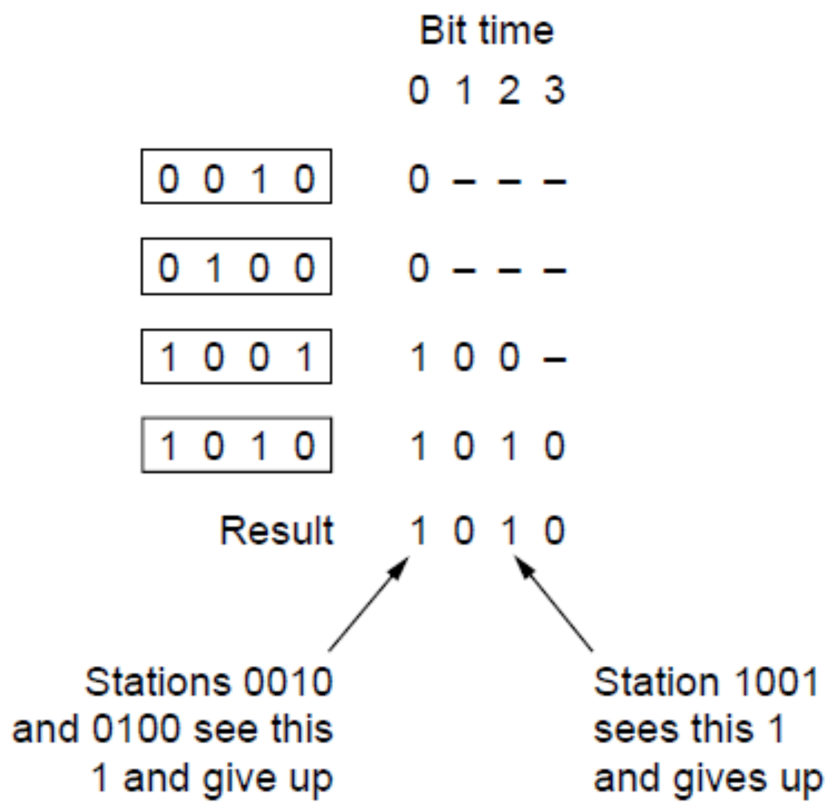

Limited-Contention Protocols (1)

Idea is to divide stations into groups within which only a very small number are likely to want to send

- Avoids wastage due to idle periods and collisions

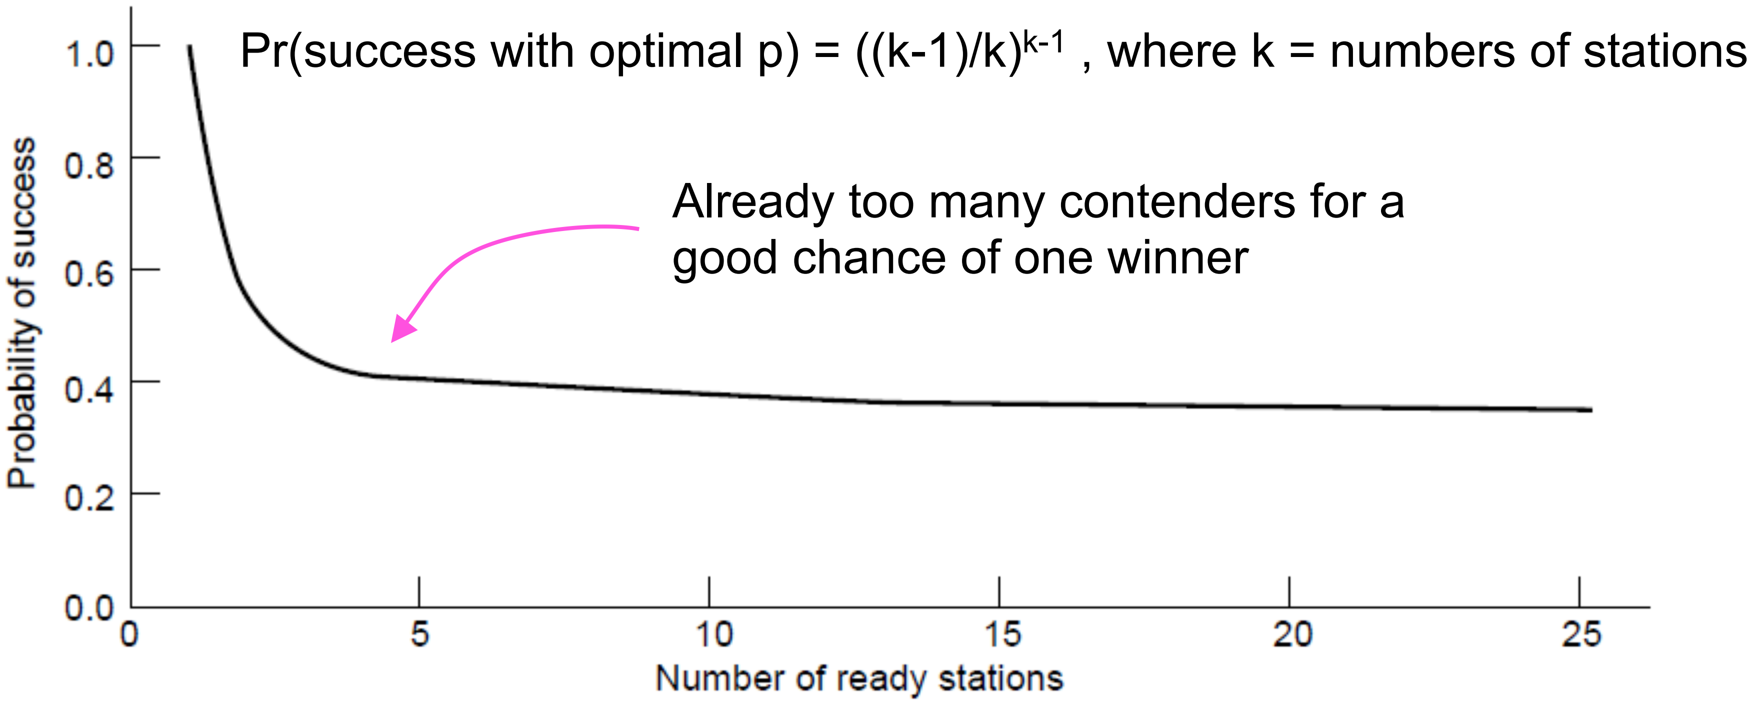

Limited Contention (2) –Adaptive Tree Walk

Tree divides stations into groups (nodes) to poll

- Depth first search under nodes with poll collisions
- Start search at lower levels if >1 station expected

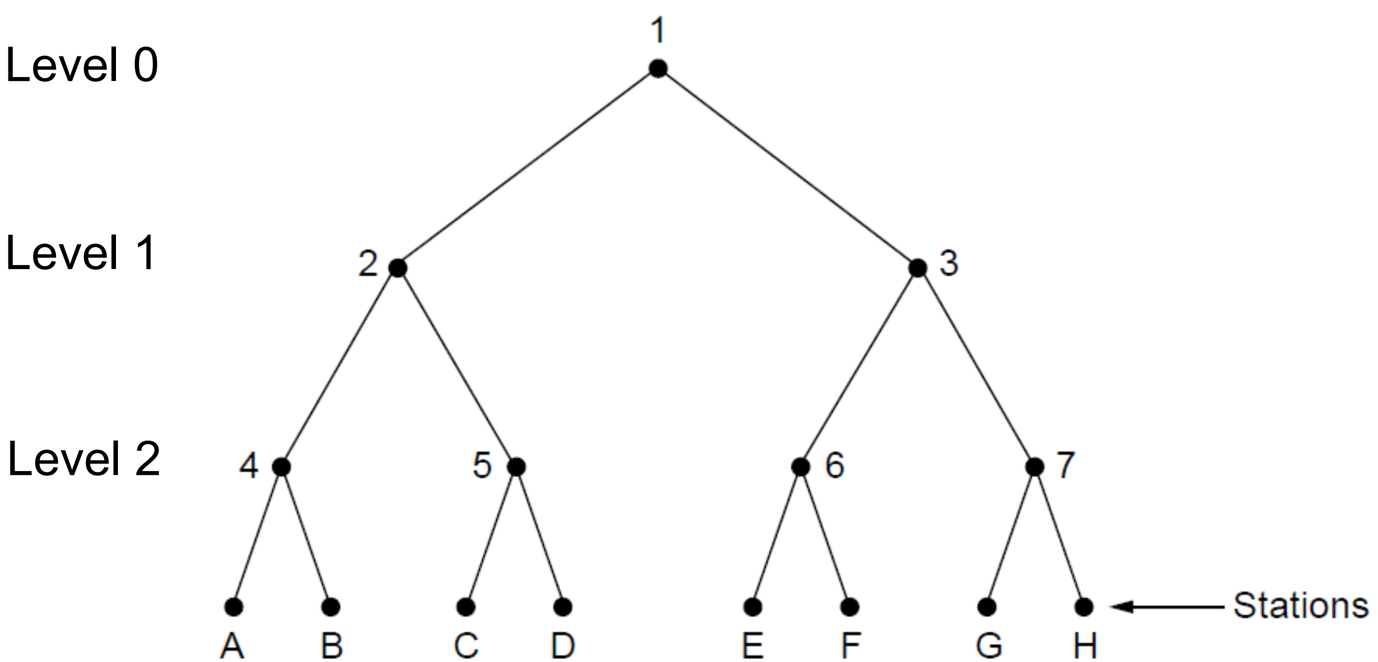

Summary of MAC protocols

channel partitioning, by time, frequency or code

- Time Division, Frequency Division

random access (dynamic),

- ALOHA, S-ALOHA, CSMA, CSMA/CD
- carrier sensing: easy in some technologies (wire), hard in others (wireless)
- CSMA/CD used in Ethernet
- CSMA/CA used in 802.11

taking turns

- polling from central site, token passing
- Bluetooth, FDDI, IBM Token Ring

Wireless and Mobile Networks

Background:

wireless (mobile) phone subscribers now exceeds # wired phone subscribers!

wireless Internet-connected devices soon to exceed # wireline Internet-connected devices

- laptops, Internet-enabled phones promise anytime untethered Internet access

two important (but different) challenges

- *wireless:* communication over wireless link
- *mobility:* handling the mobile user who changes point of attachment to network

Elements of a wireless network

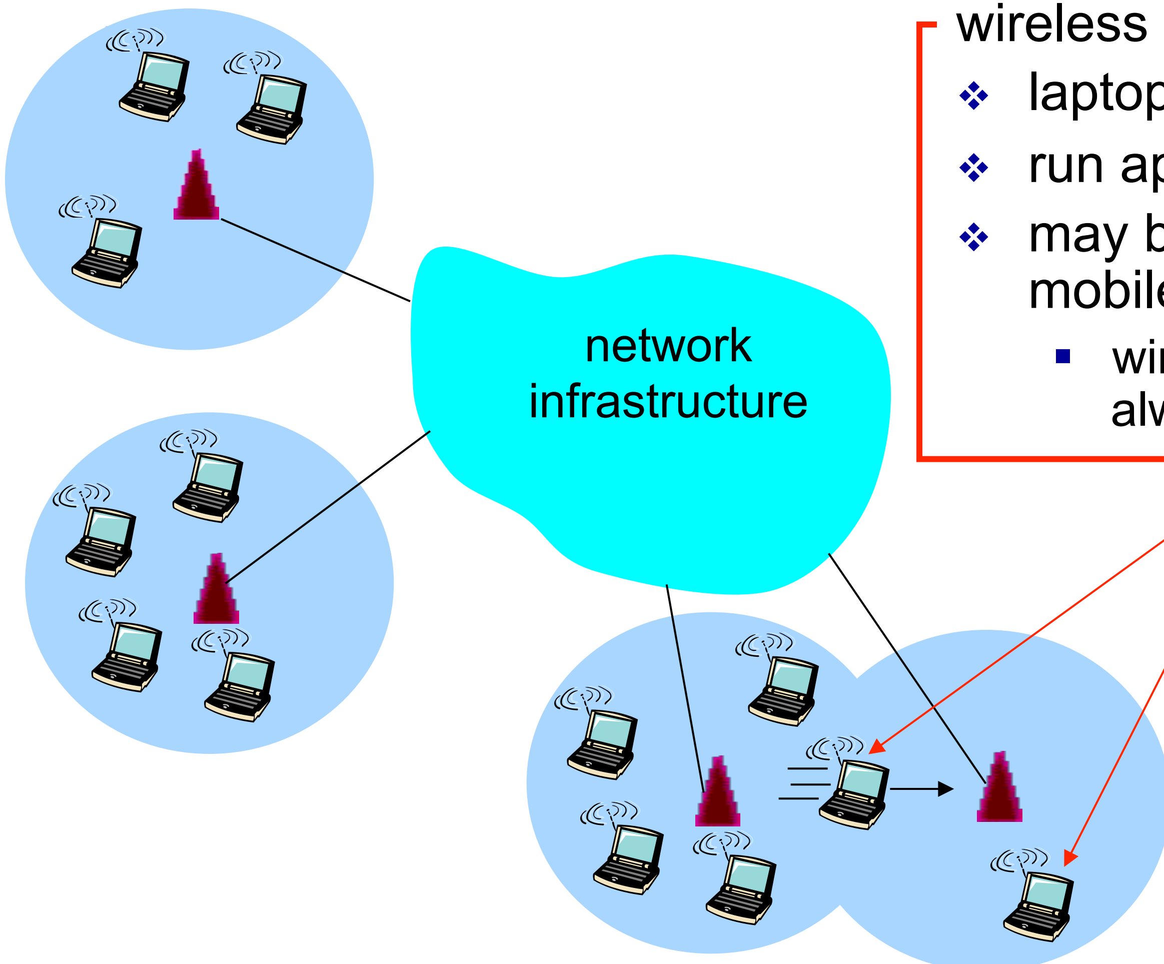

wireless hosts 

- laptop, PDA, IP phone
- run applications
- may be stationary (non-mobile) or mobile
  - wireless does *not* always mean mobility

Elements of a wireless network

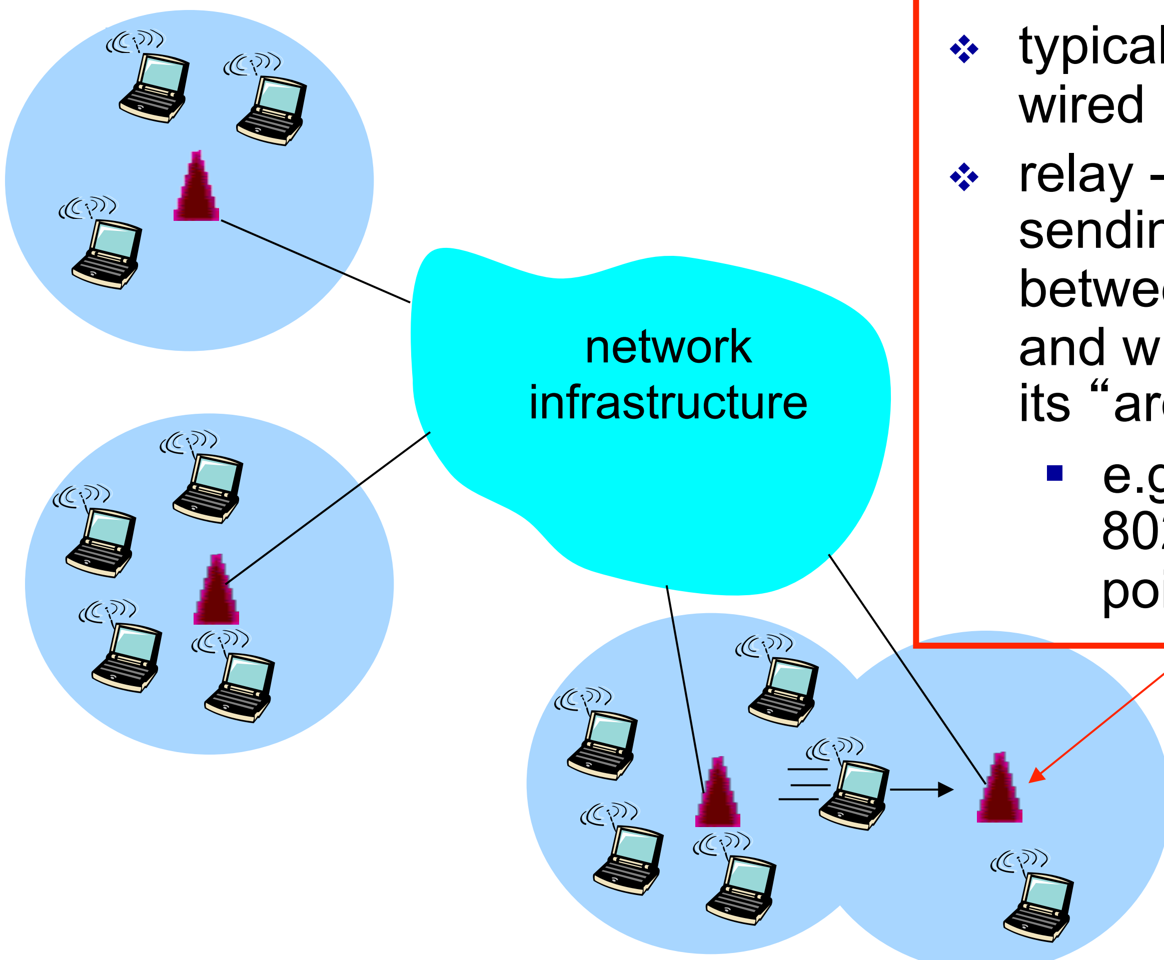

base station

- typically connected to wired network
- relay - responsible for sending packets between wired network and wireless host(s) in its “area”
  - e.g., cell towers, 802.11 access points

Elements of a wireless network

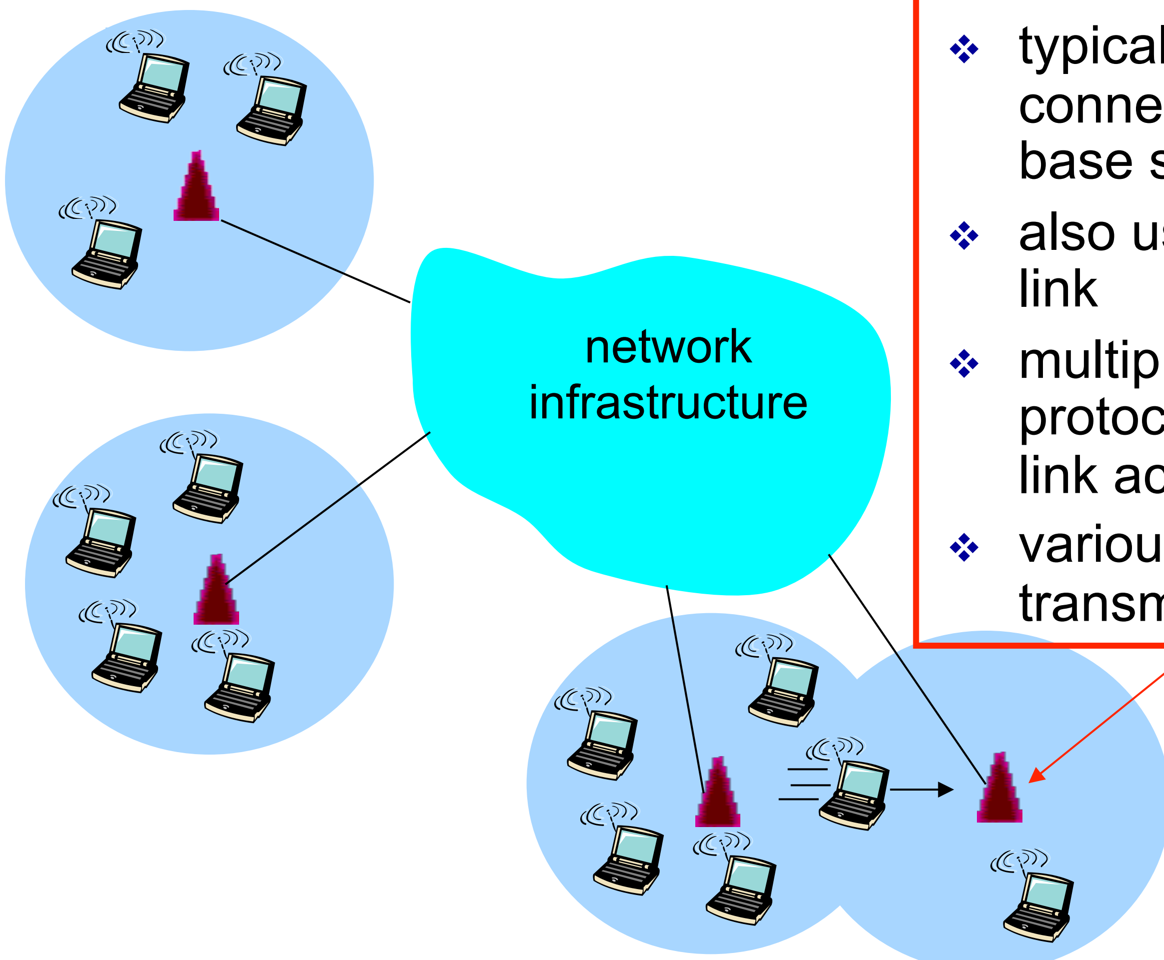

wireless link
- typically used to connect mobile(s) to base station
- also used as backbone link
- multiple access protocol coordinates link access
- various data rates, transmission distance

Characteristics of selected wireless link standards

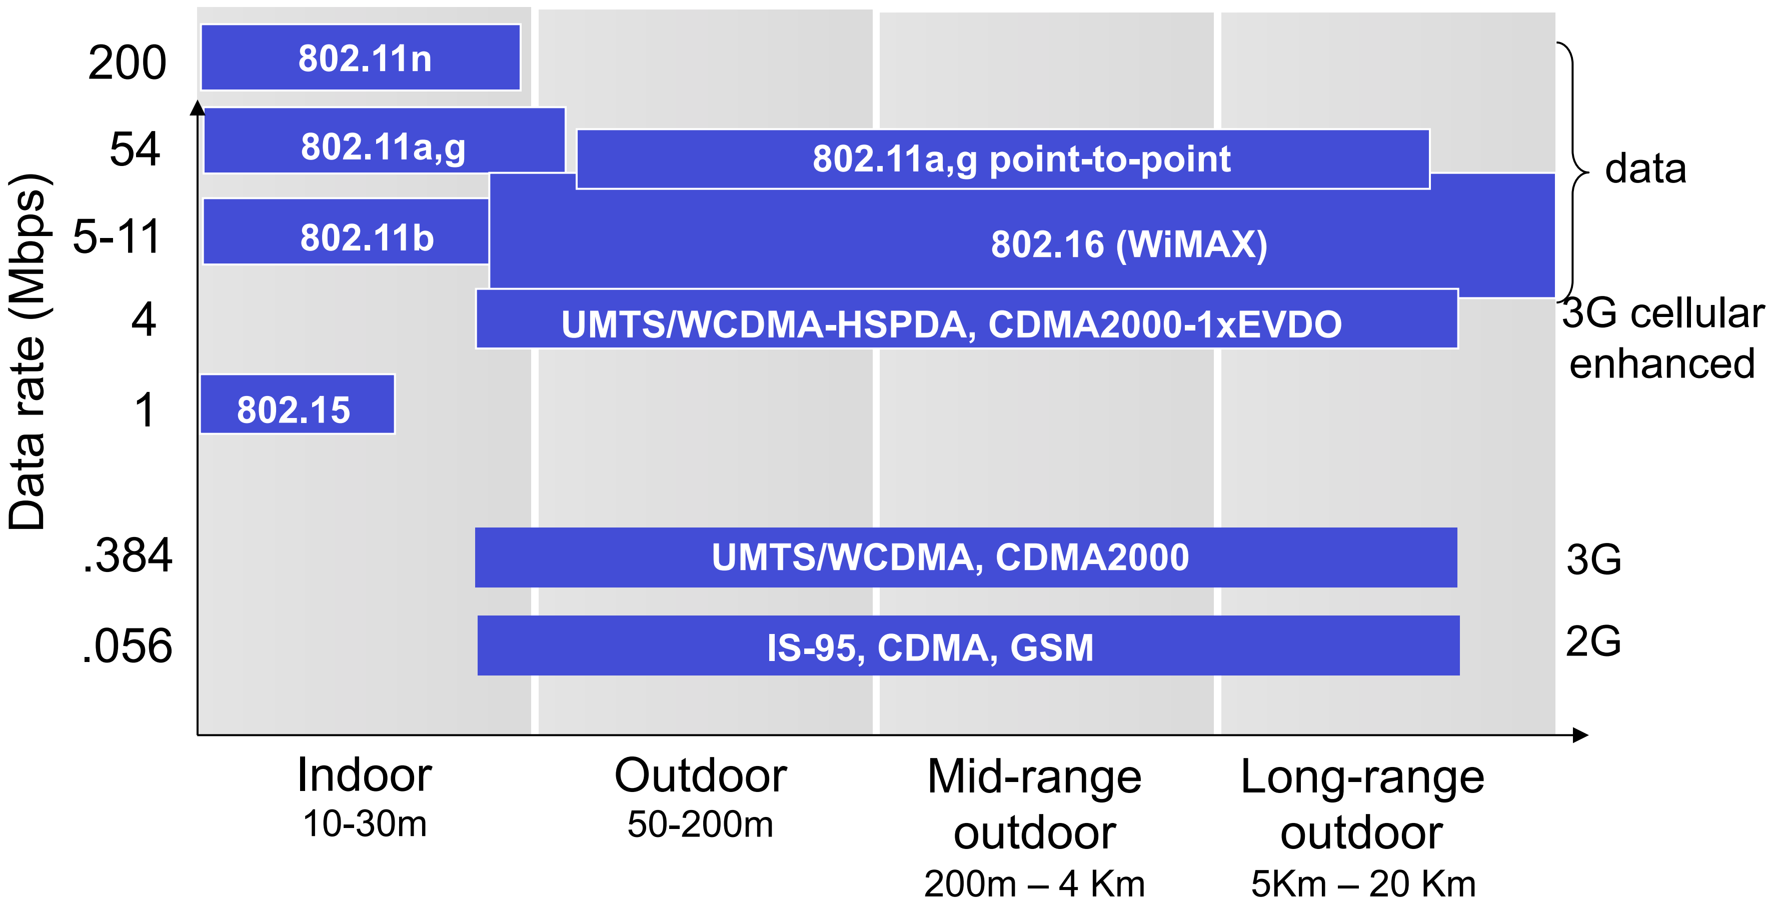

Elements of a wireless network

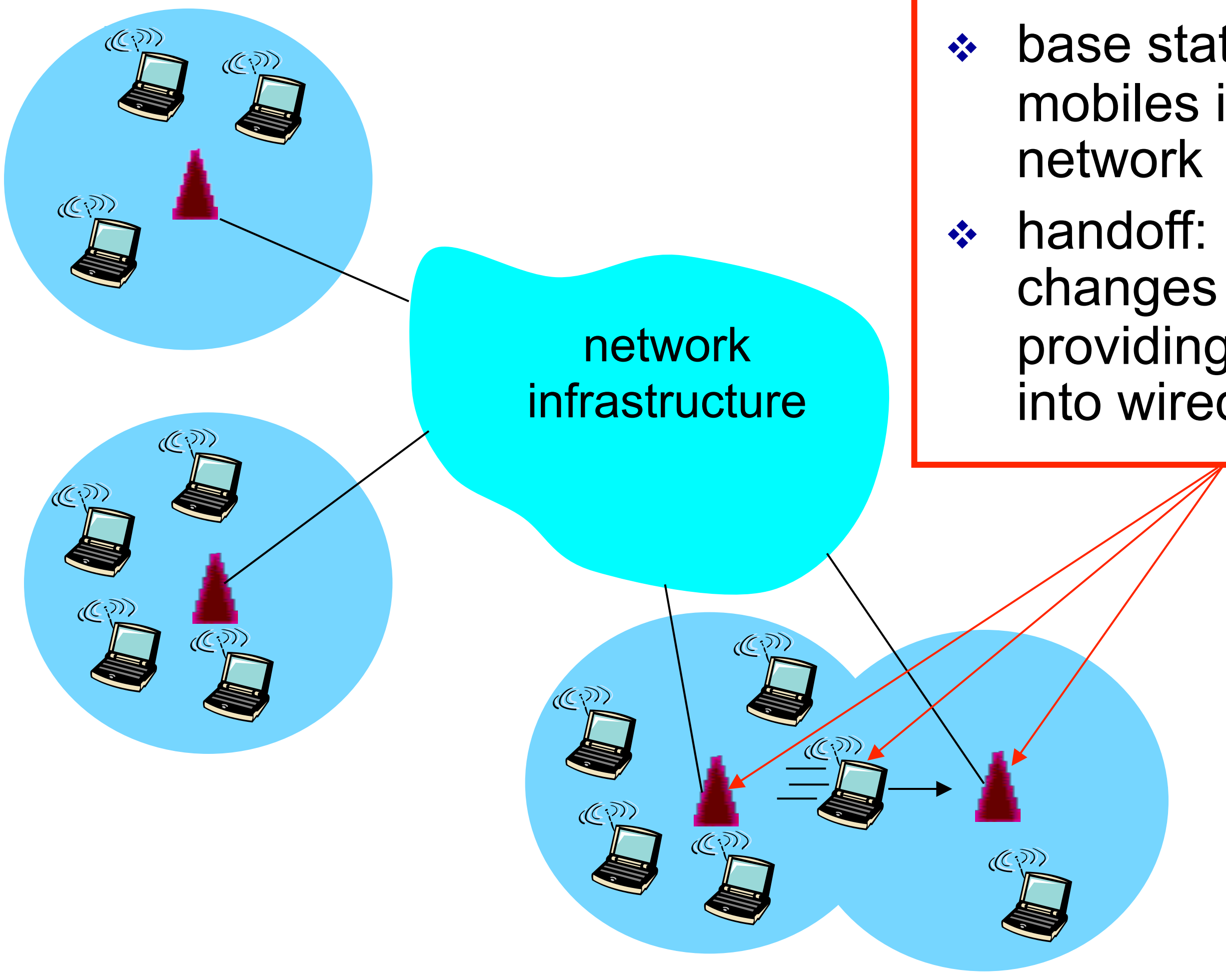

infrastructure mode

- base station connects mobiles into wired network
- handoff: mobile changes base station providing connection into wired network

Elements of a wireless network

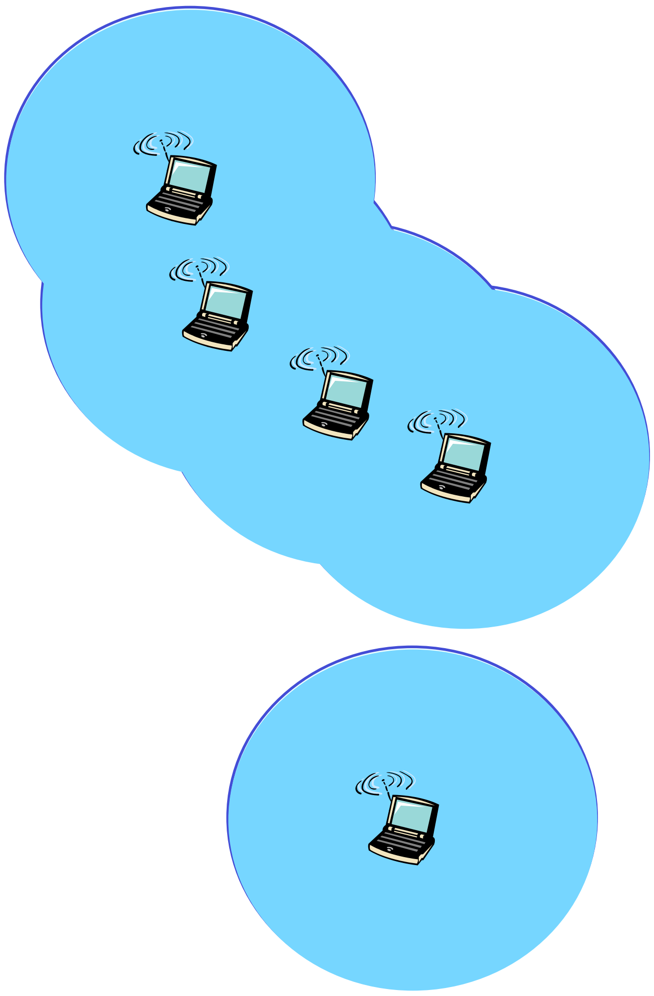

ad hoc mode

- no base stations
- nodes can only transmit to other nodes within link coverage
- nodes organize themselves into a network: route among themselves

Wireless network taxonomy

	single hop	multiple hops
infrastructure (e.g., APs)	host connects to base station (WiFi, WiMAX, cellular) which connects to larger Internet	host may have to relay through several wireless nodes to connect to larger Internet: *mesh net*
no infrastructure	no base station, no connection to larger Internet (Bluetooth, ad hoc nets)	no base station, no connection to larger Internet. May have to relay to reach other a given wireless node MANET, VANET

Wireless Link Characteristics (1)

Differences from wired link ….

- decreased signal strength: radio signal attenuates as it propagates through matter (path loss)
- interference from other sources: standardized wireless network frequencies (e.g., 2.4 GHz) shared by other devices (e.g., phone); devices (motors) interfere as well
- multipath propagation: radio signal reflects off objects ground, arriving ad destination at slightly different times

…. make communication across (even a point to point) wireless link much more "difficult"

Wireless network characteristics

Multiple wireless senders and receivers create additional problems (beyond multiple access):

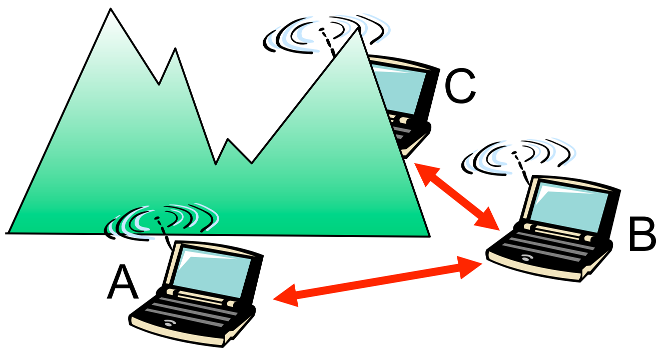

Hidden terminal problem

- 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

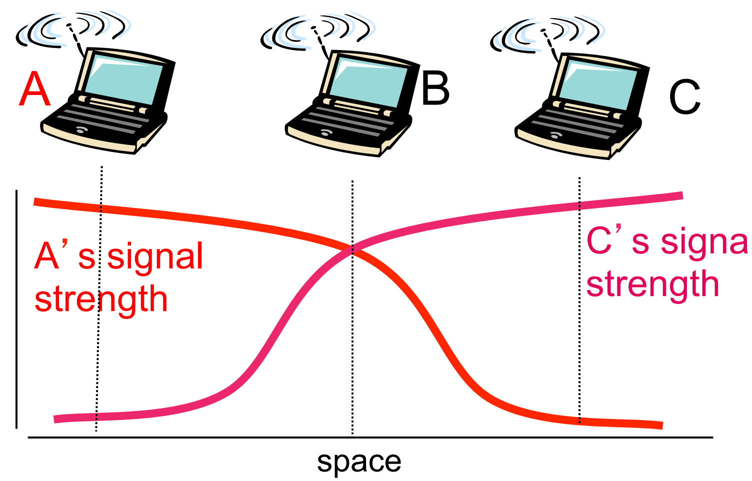

Signal attenuation:

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

Wireless LAN Protocols (1)

Wireless has complications compared to wired.

Nodes may have different coverage regions

- Leads to <u>hidden</u> and <u>exposed</u> terminals

Nodes can't detect collisions, i.e., sense while sending

- Makes collisions expensive and to be avoided

Wireless LANs (2) – Hidden terminals

Hidden terminals are senders that cannot sense each other but nonetheless collide at intended receiver

- Want to prevent; loss of efficiency
- A and C are hidden terminals when sending to B

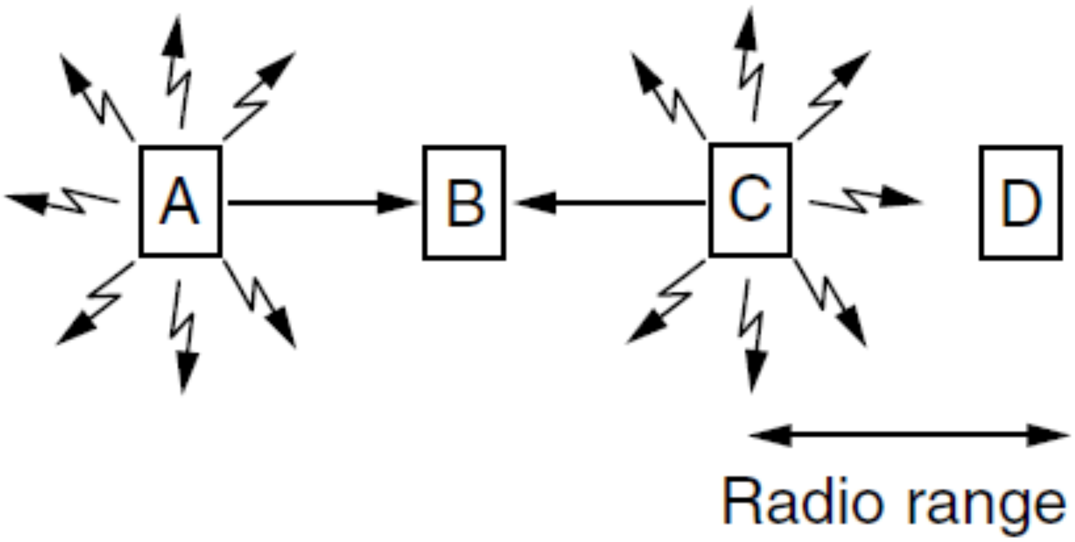

Wireless LANs (3) – Exposed terminals

Exposed terminals are senders who can sense each other but still transmit safely (to different receivers)

- Desirably concurrency; improves performance
- B → A and C → D are exposed terminals

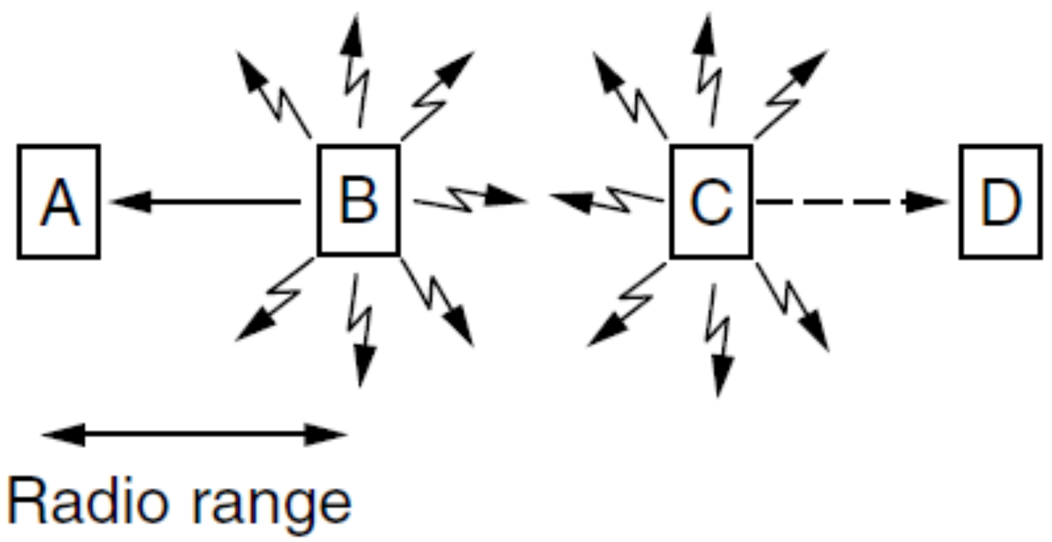

IEEE 802.11: multiple access

avoid collisions: 2^+ nodes transmitting at same time

802.11: CSMA - sense before transmitting

- don't collide with ongoing transmission by other node

802.11: *no* collision detection!

- difficult to receive (sense collisions) when transmitting due to weak received signals (fading)
- can't sense all collisions in any case: hidden terminal, fading
- goal: *avoid collisions:* CSMA/C(ollision)A(voidance)

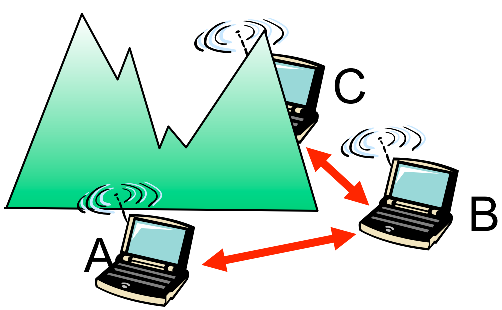

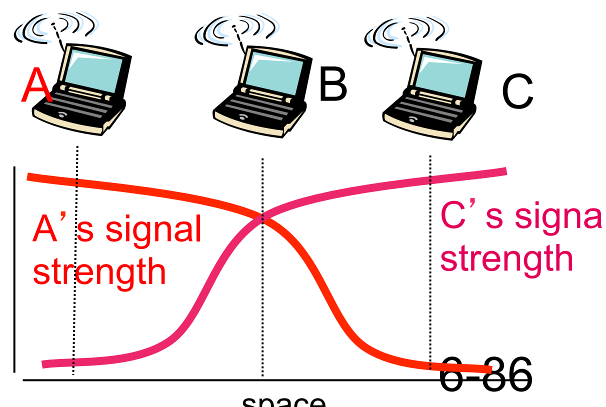

IEEE 802.11 MAC Protocol: CSMA/CA

802.11 sender

1 if sense channel idle for **DIFS** then
transmit entire frame

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

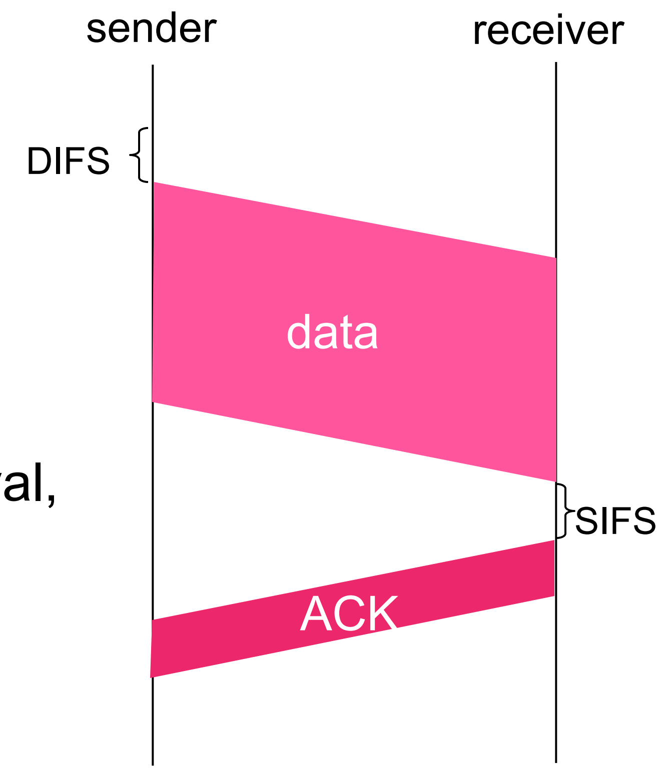

802.11 receiver

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

DIFS (Distributed Inter-frame Spacing)

SIFS (Short Inter-frame Spacing)

802.11 MAC (1)

- CSMA/CA inserts backoff slots to avoid collisions
- MAC uses ACKs/retransmissions for wireless errors

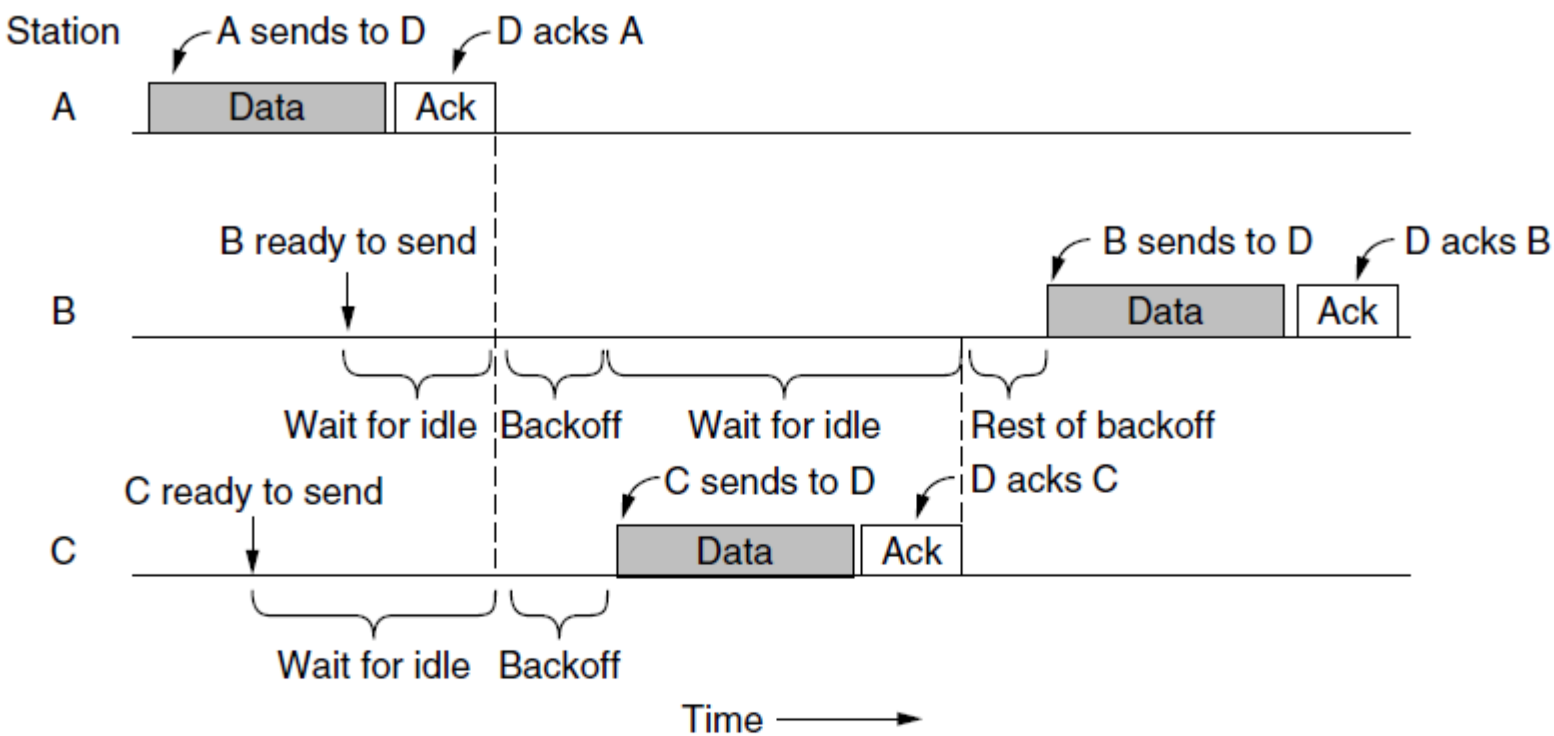

Collision Avoidance: RTS-CTS exchange

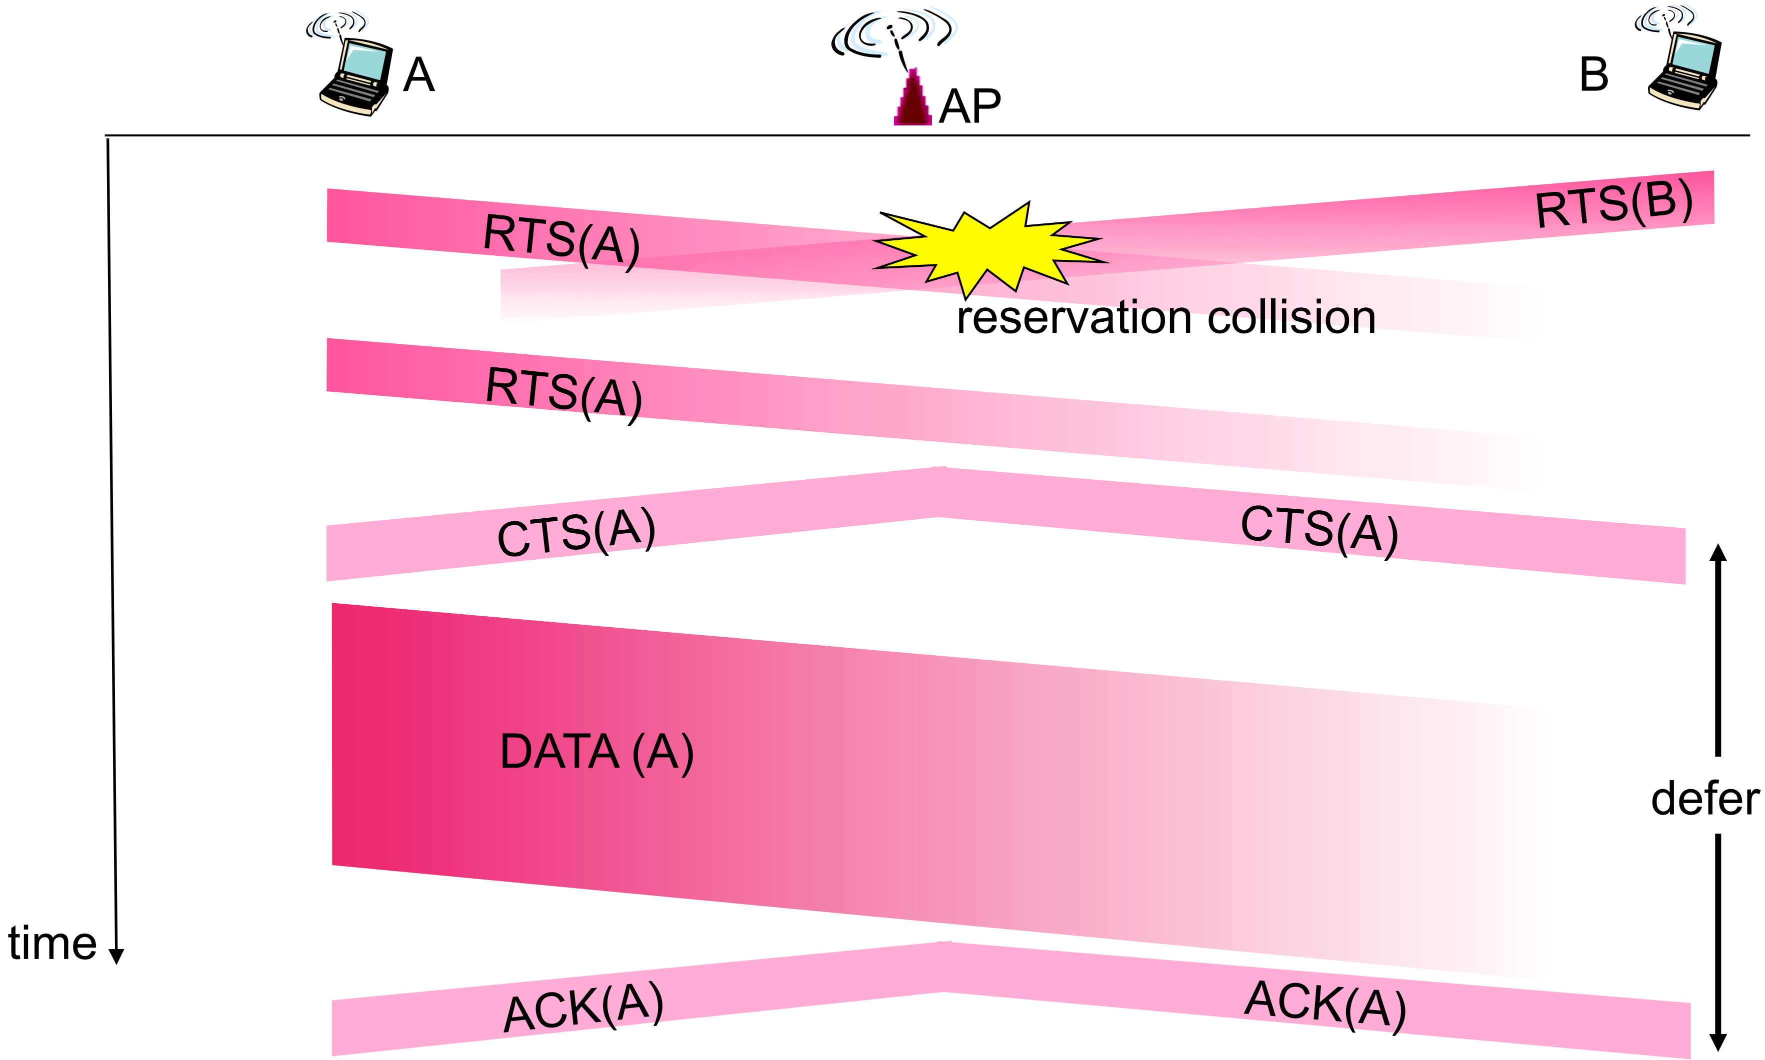

802.11 MAC (2)

Virtual channel sensing with the NAV and optional RTS/CTS (often not used) avoids hidden terminals

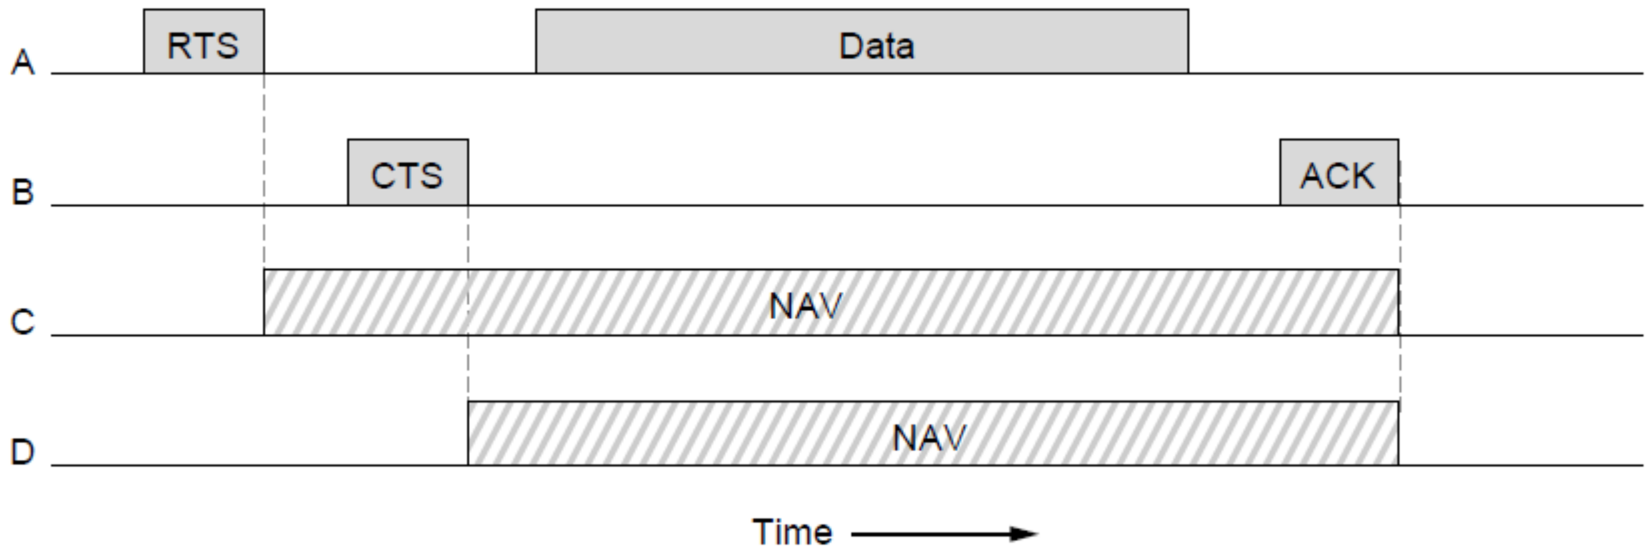

NAV = Network Allocation Vector to keep quiet for a certain period of time

Avoiding collisions (more)

idea: allow sender to "reserve" channel rather than random access of data frames: avoid collisions of long data frames

sender first transmits *small* request-to-send (RTS) packets 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

avoid data frame collisions completely using small reservation packets!

Wireless LANs (4) – MACA

MACA protocol grants access for A to send to B:

- A sends RTS to B [left]; B replies with CTS [right]
- A can send with exposed but no hidden terminals

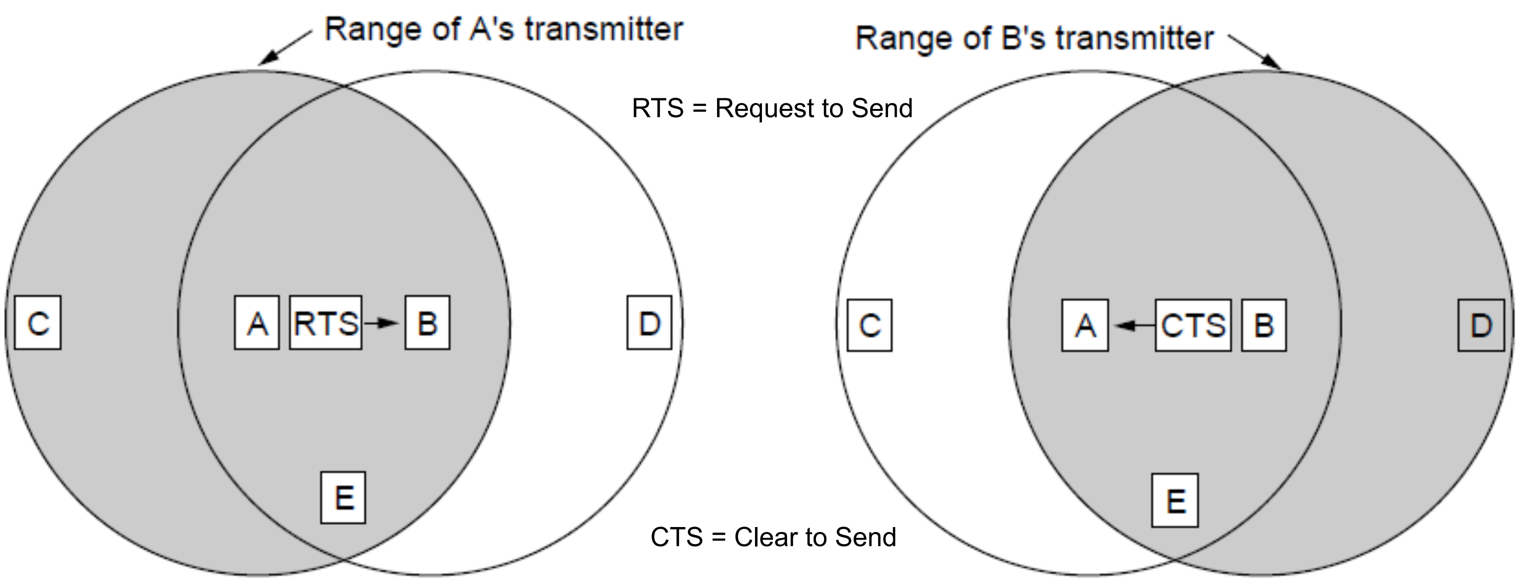

A sends RTS to B; C and E hear and defer for CTS

B replies with CTS; D and E hear and defer for data

802.11 MAC (3)

- Different backoff slot times add quality of service
  - Short intervals give preferred access, e.g., control, VoIP
- MAC has other mechanisms too, e.g., power save

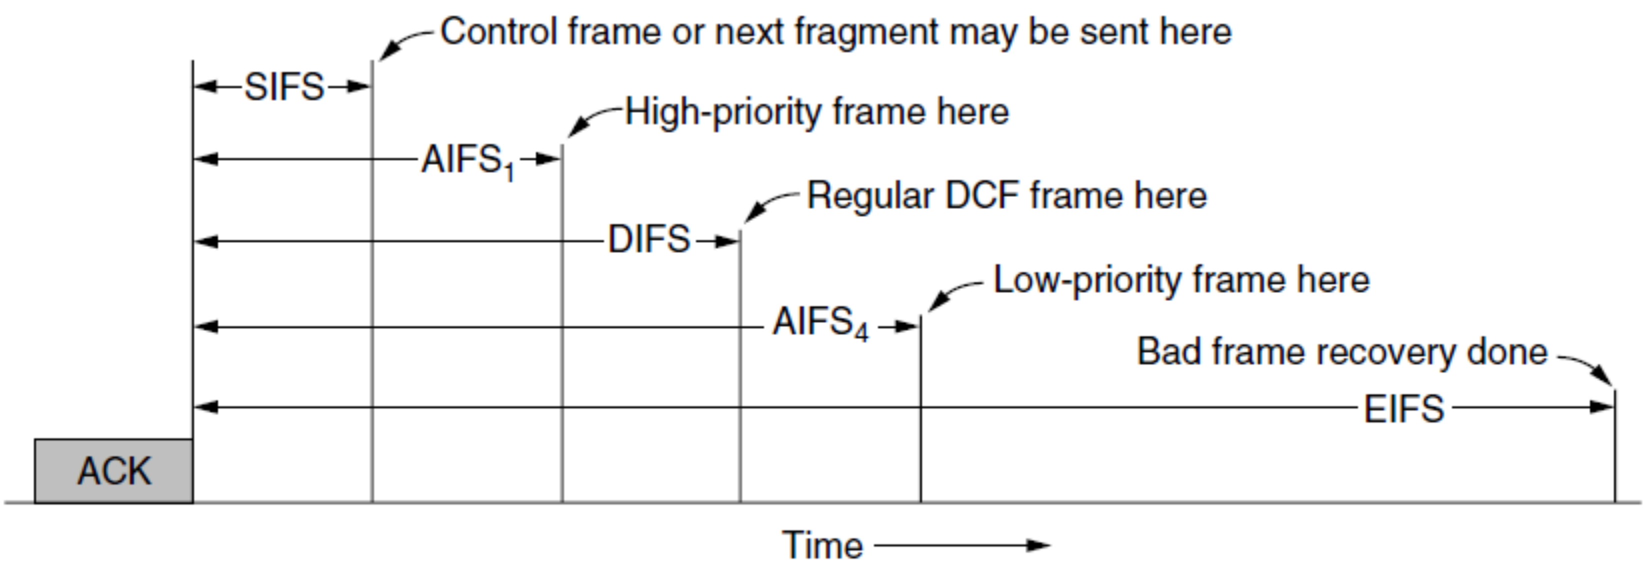

802.11 frame: addressing

2	2	6	6	6	2	6	0 - 2312	4
frame control	duration	address 1	address 2	address 3	seq control	address 4	payload	CRC

Address 1: MAC address of wireless host or AP to receive this frame

Address 2: MAC address of wireless host or AP transmitting this frame

Address 3: MAC address of router interface to which AP is attached

Address 4: used only in ad hoc mode

802.11 frame: addressing

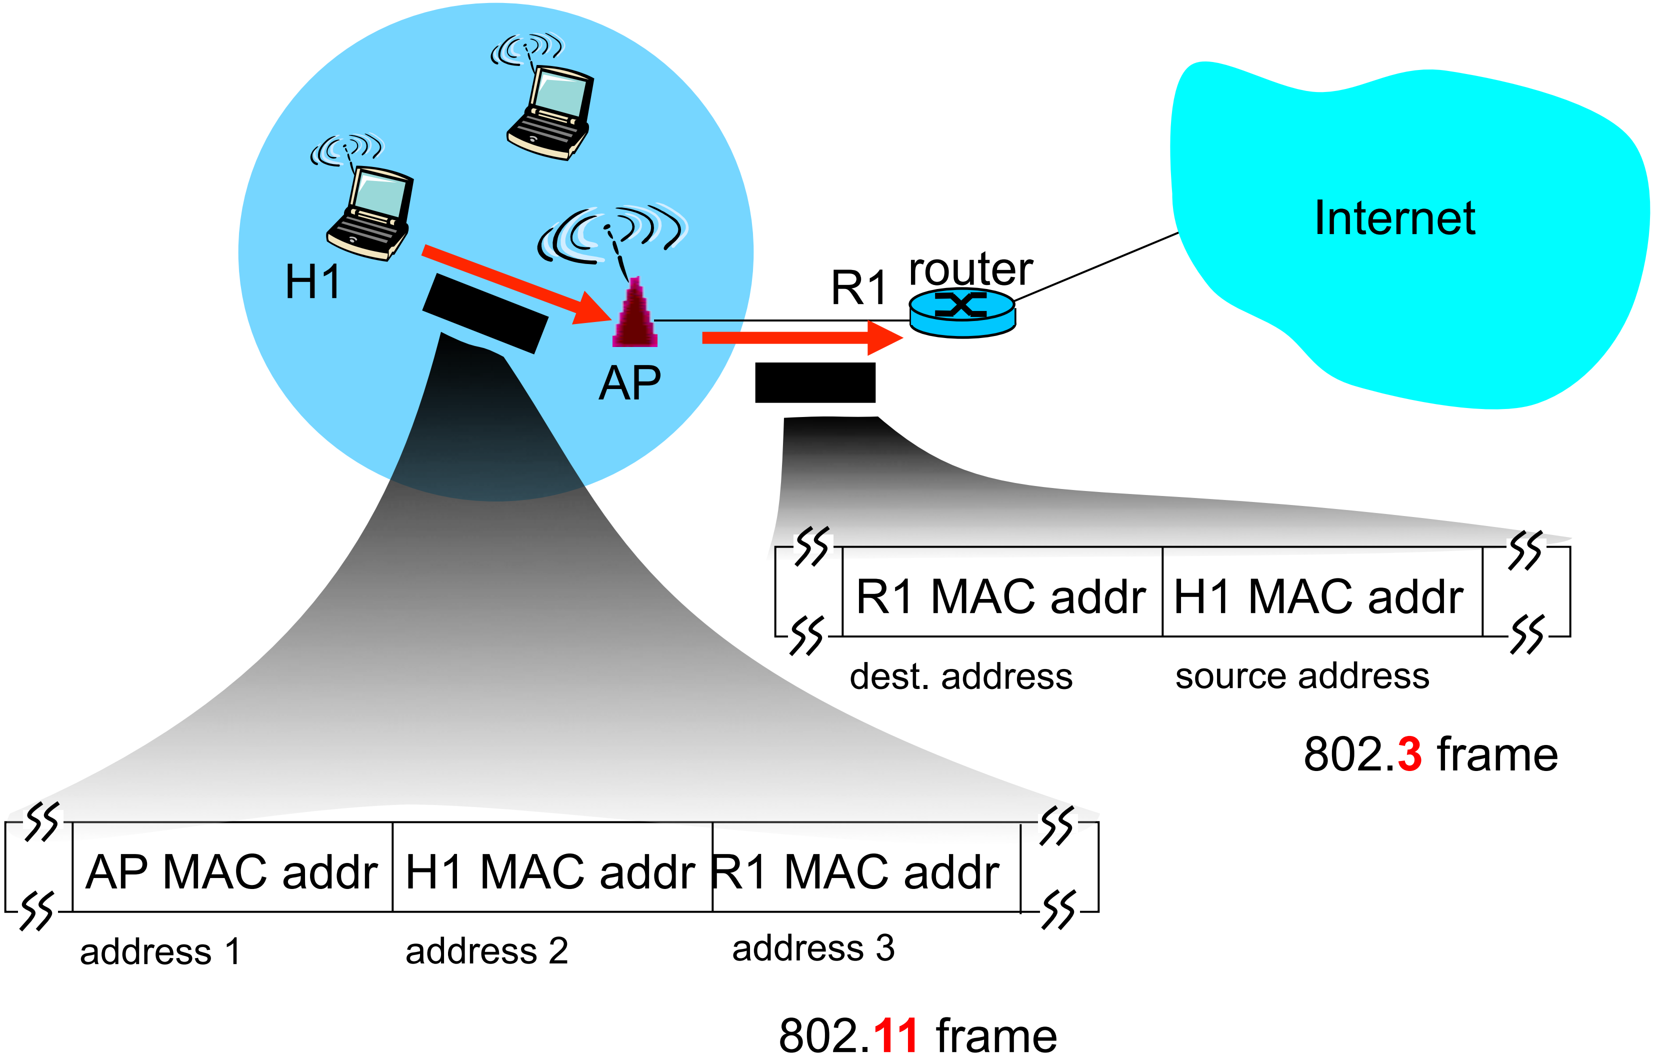

802.11 frame: more

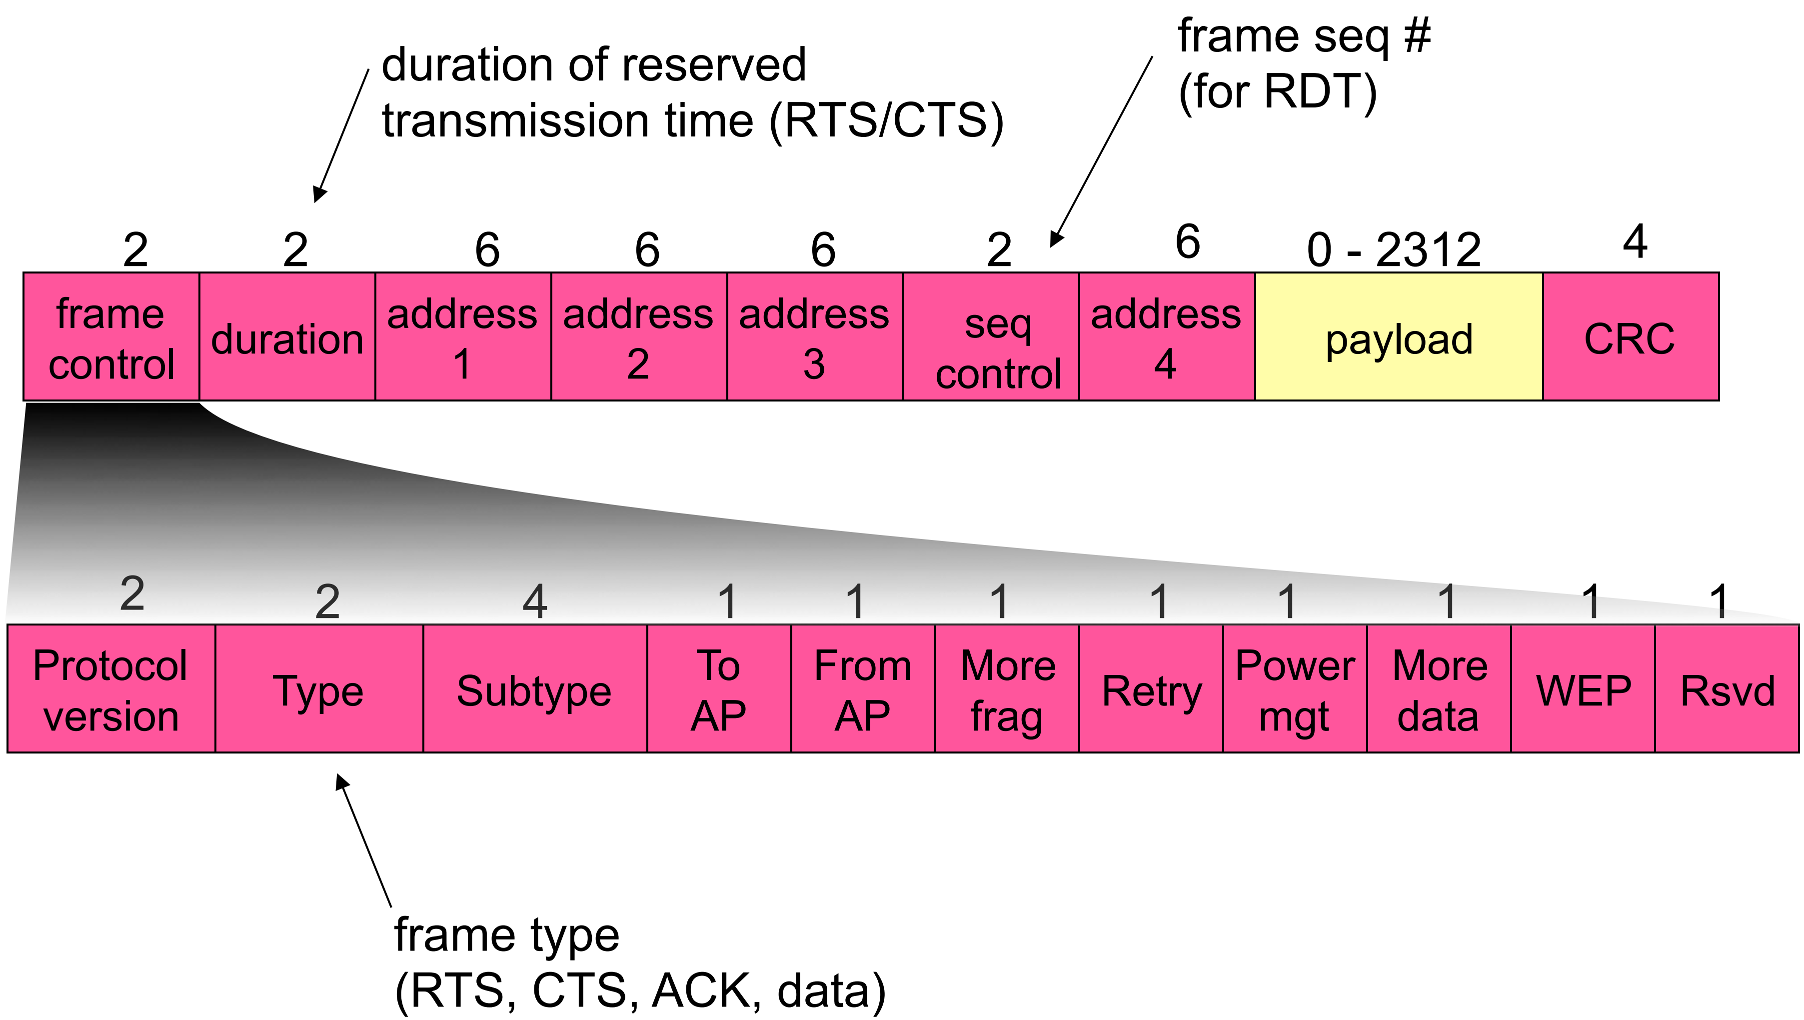

Chapter 6 outline

6.1 Introduction

Wireless

6.2 Wireless links, characteristics

- CDMA

6.3 IEEE 802.11 wireless LANs (“Wi-Fi”)

6.4 Cellular Internet Access

- architecture
- standards (e.g., GSM)

Mobility

6.5 Principles: addressing and routing to mobile users

6.6 Mobile IP

6.7 Handling mobility in cellular networks

6.8 Mobility and higher-layer protocols

6.9 Summary

IEEE 802.11 Wireless LAN

802.11b

- 2.4-5 GHz unlicensed spectrum
- up to 11 Mbps
- direct sequence spread spectrum (DSSS) in physical layer
  - all hosts use same chipping code

802.11a

- 5-6 GHz range
- up to 54 Mbps

802.11g

- 2.4-5 GHz range
- up to 54 Mbps

802.11n: multiple antennae

- 2.4-5 GHz range
- up to 200 Mbps

- all use CSMA/CA for multiple access
- all have base-station and ad-hoc network versions

802.11 LAN architecture

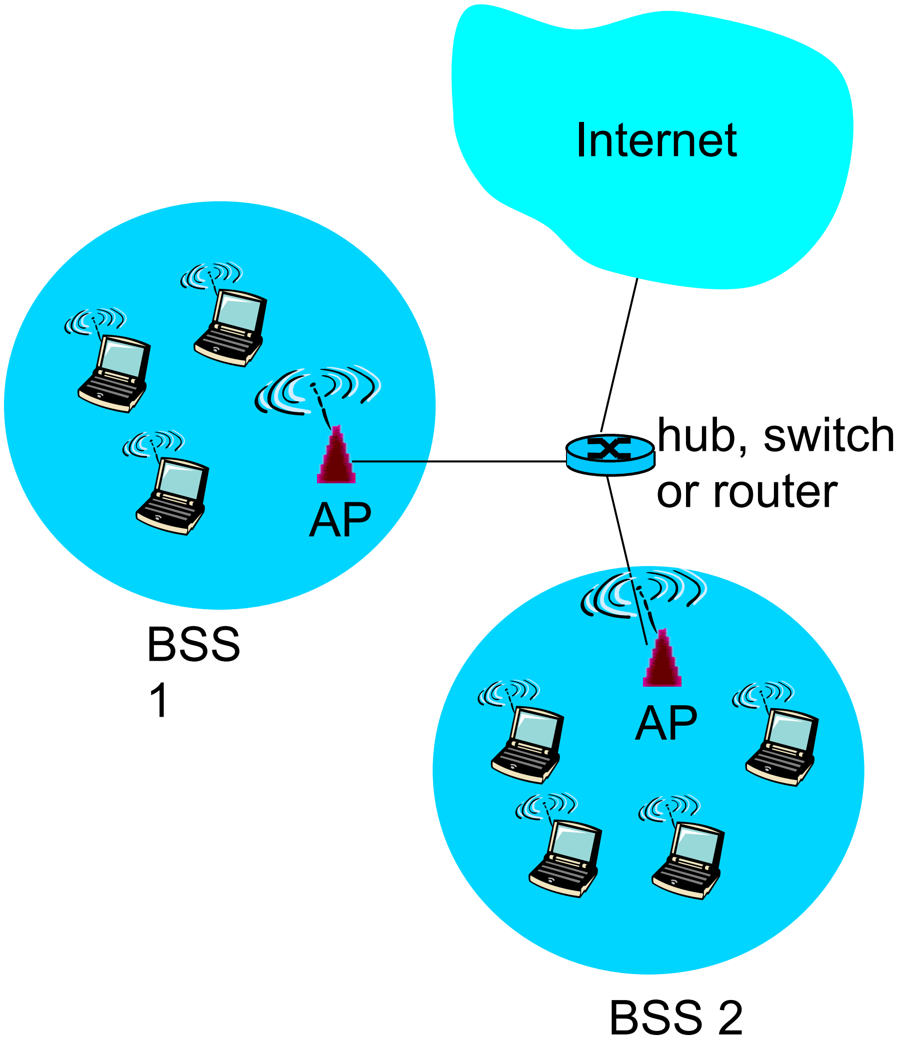

- wireless host communicates with base station
  - base station = access point (AP)
- Basic Service Set (BSS) (aka "cell") in infrastructure mode contains:
  - wireless hosts
  - access point (AP): base station
  - ad hoc mode: hosts only

802.11: Channels, association

802.11b: 2.4GHz-2.485GHz spectrum divided into 11 channels at different frequencies

- AP admin chooses frequency for AP
- interference possible: channel can be same as that chosen by neighboring AP!

host: must *associate* with an AP

- scans channels, listening for *beacon frames* containing AP's name (SSID) and MAC address
- selects AP to associate with
- may perform authentication [Chapter 8]
- will typically run DHCP to get IP address in AP's subnet

802.11: passive/active scanning

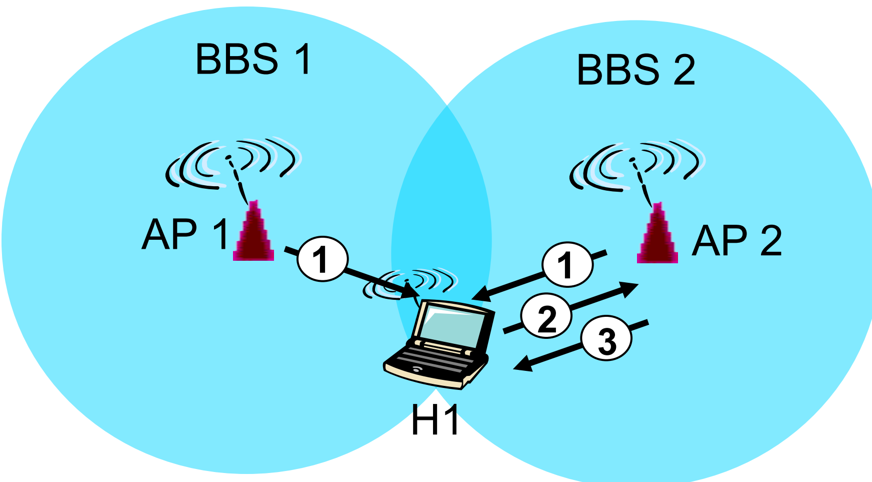

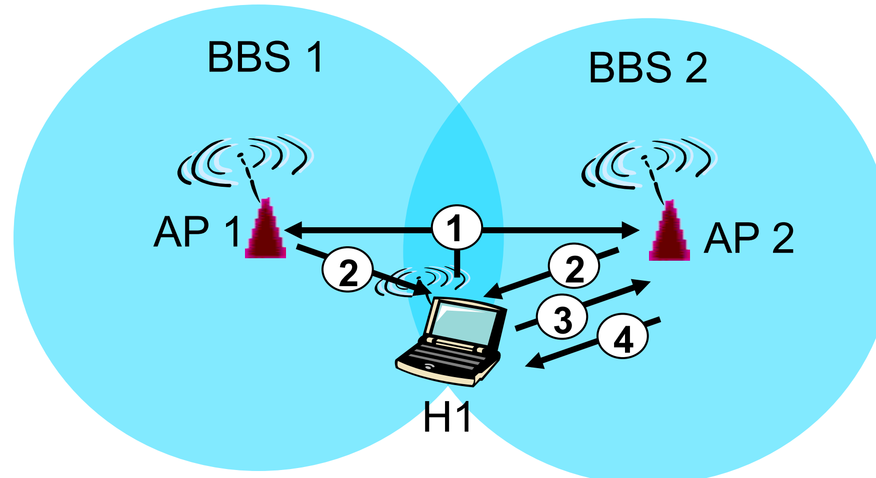

Passive Scanning:

(1) beacon frames sent from APs
(2) association Request frame sent: H1 to selected AP
(3) association Response frame sent: H1 to selected AP

Active Scanning:

(1) Probe Request frame broadcast from H1
(2) Probes response frame sent from APs
(3) Association Request frame sent: H1 to selected AP
(4) Association Response frame sent: H1 to selected AP

802.16: WiMAX

like 802.11 & cellular: base station model

- transmissions to/from base station by hosts with omnidirectional antenna
- base station-to-base station backhaul with point-to-point antenna

unlike 802.11:

- range ~ 6 miles ("city rather than coffee shop")
- ~14 Mbps

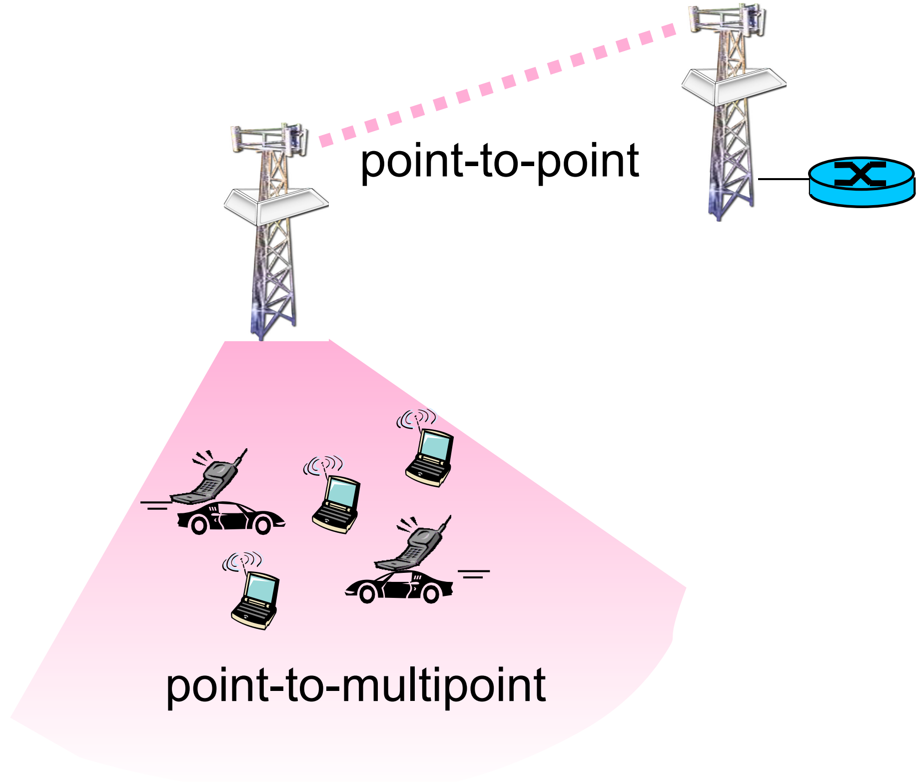

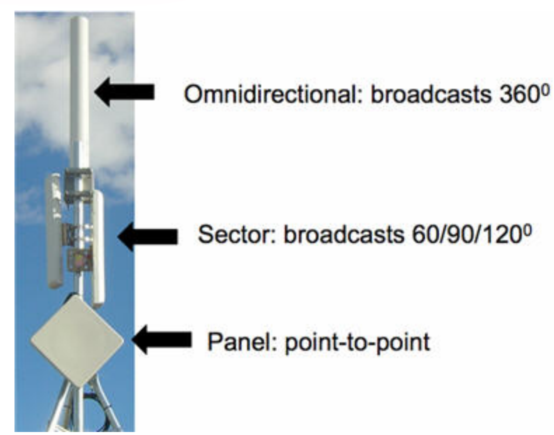

802.16: WiMAX: downlink, uplink scheduling

transmission frame

- down-link subframe: base station to node
- uplink subframe: node to base station

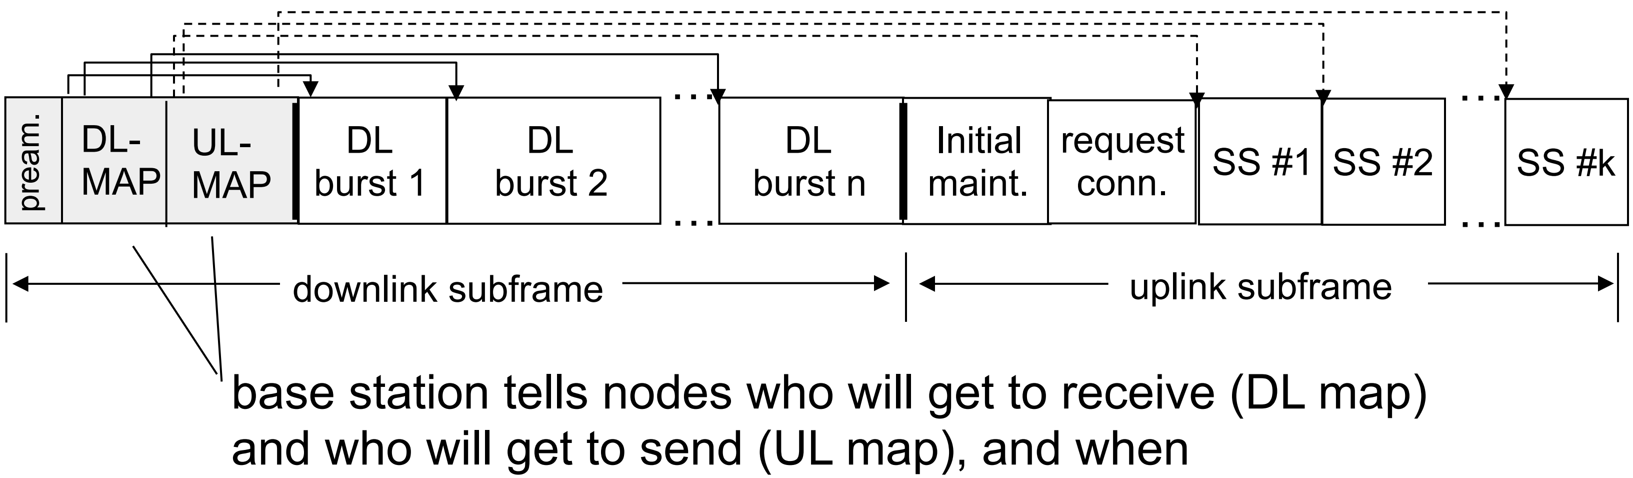

- WiMAX standard provide mechanism for scheduling, but not scheduling algorithm

Chapter 6 outline

6.1 Introduction

Wireless

6.2 Wireless links, characteristics

- CDMA

6.3 IEEE 802.11 wireless LANs (“Wi-Fi”)

6.4 Cellular Internet Access

- architecture
- standards (e.g., GSM)

Mobility

6.5 Principles: addressing and routing to mobile users

6.6 Mobile IP

6.7 Handling mobility in cellular networks

6.8 Mobility and higher-layer protocols

6.9 Summary

Components of cellular network architecture

cell
- covers geographical region
- *base station* (BS) analogous to 802.11 AP
- *mobile users* attach to network through BS
- *air-interface:* physical and link layer protocol between mobile and BS

MSC
- connects cells to wide area net
- manages call setup (more later!)
- handles mobility (more later!)

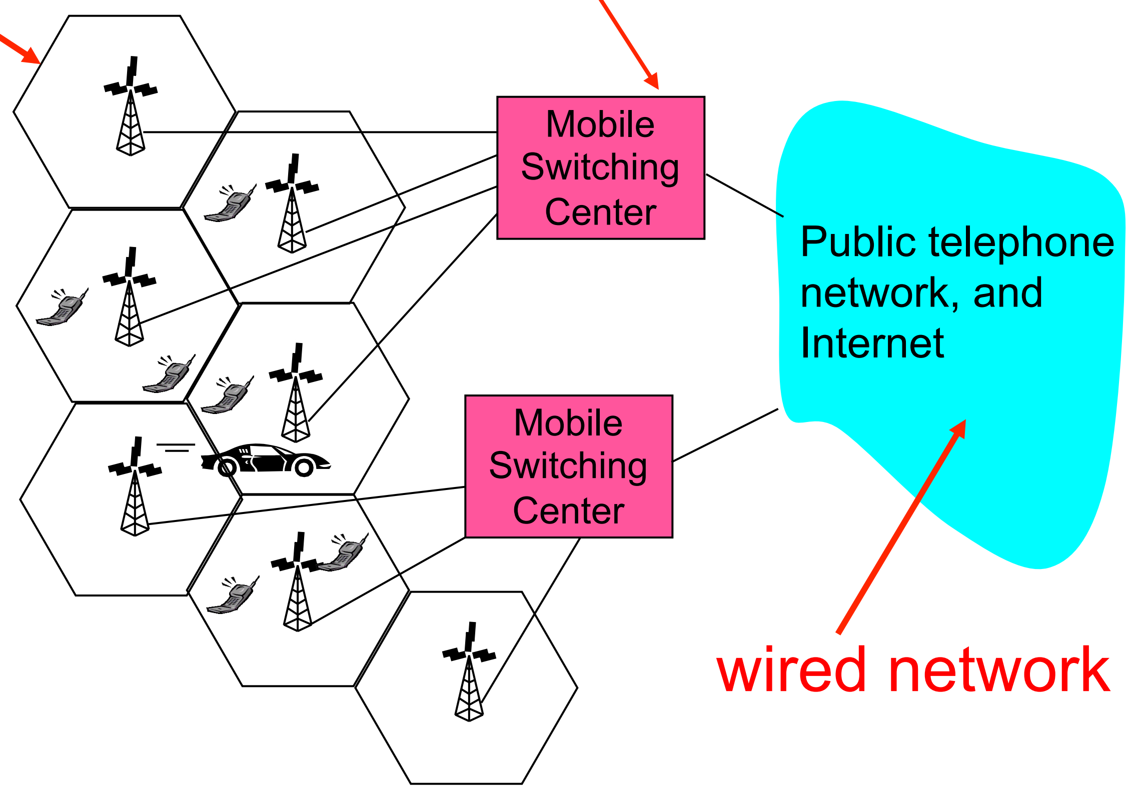

Cellular networks: the first hop

Two techniques for sharing mobile-to-BS radio spectrum

combined FDMA/TDMA: divide spectrum in frequency channels, divide each channel into time slots

CDMA: code division multiple access

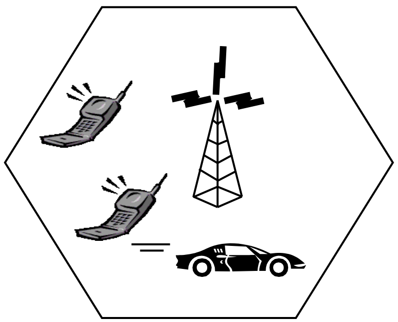

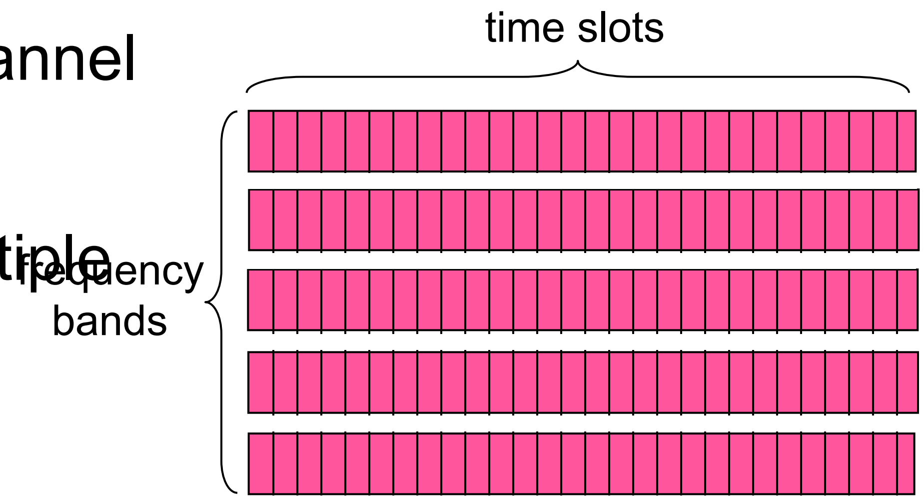

Cellular standards: brief survey

2G systems: voice channels

IS-136 TDMA: combined FDMA/TDMA (North America)

GSM (global system for mobile communications): combined FDMA/TDMA

- most widely deployed

IS-95 CDMA: code division multiple access

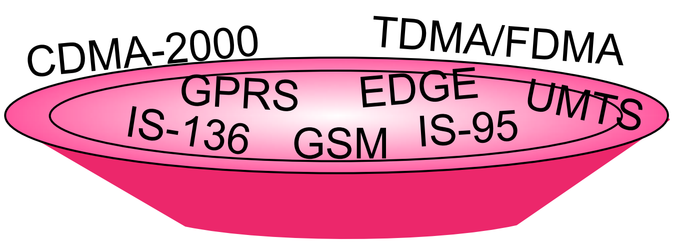

Don't drown in a bowl of alphabet soup: use this for reference only

Cellular standards: brief survey

2.5 G systems: voice and data channels

for those who can't wait for 3G service: 2G extensions

general packet radio service (GPRS)

- evolved from GSM
- data sent on multiple channels (if available)

enhanced data rates for global evolution (EDGE)

- also evolved from GSM, using enhanced modulation
- data rates up to 384K

CDMA-2000 (phase 1)

- data rates up to 144K
- evolved from IS-95

Cellular standards: brief survey

3G systems: voice/data

Universal Mobile Telecommunications Service (UMTS)

- data service: High Speed Uplink/Downlink packet Access (HSDPA/HSUPA): 3 Mbps

CDMA-2000: CDMA in TDMA slots

- data service: 1xEvolution Data Optimized (1xEVDO) up to 14 Mbps

2G (voice) network architecture

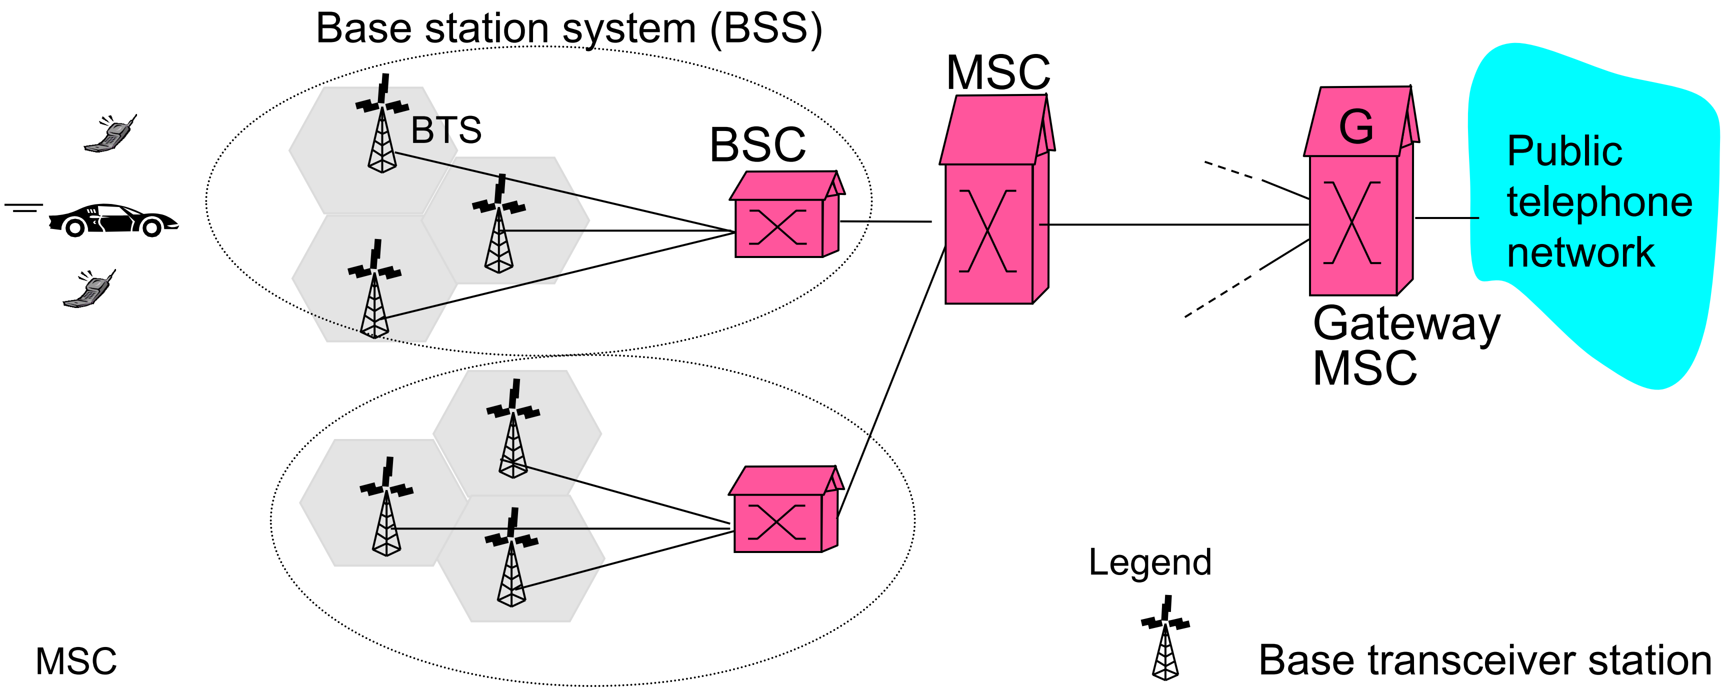

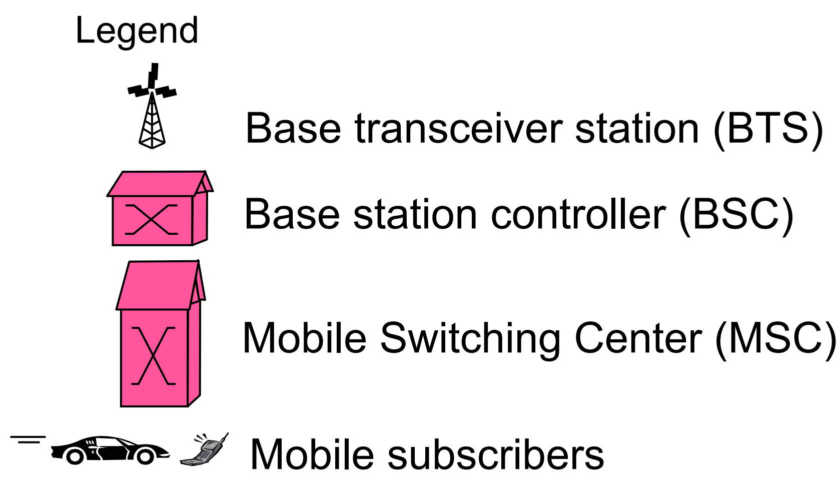

MSC

- User Authorization and accounting
- connects cells to wide area net
- manages call setup
- handles mobility

BSC

- Allocate BTS radio Channel to mobile mobile subscribers
- Performing paging (finding the cell in which a mobile user is resident)
- Perform handoff

2.5G (voice+data) network architecture

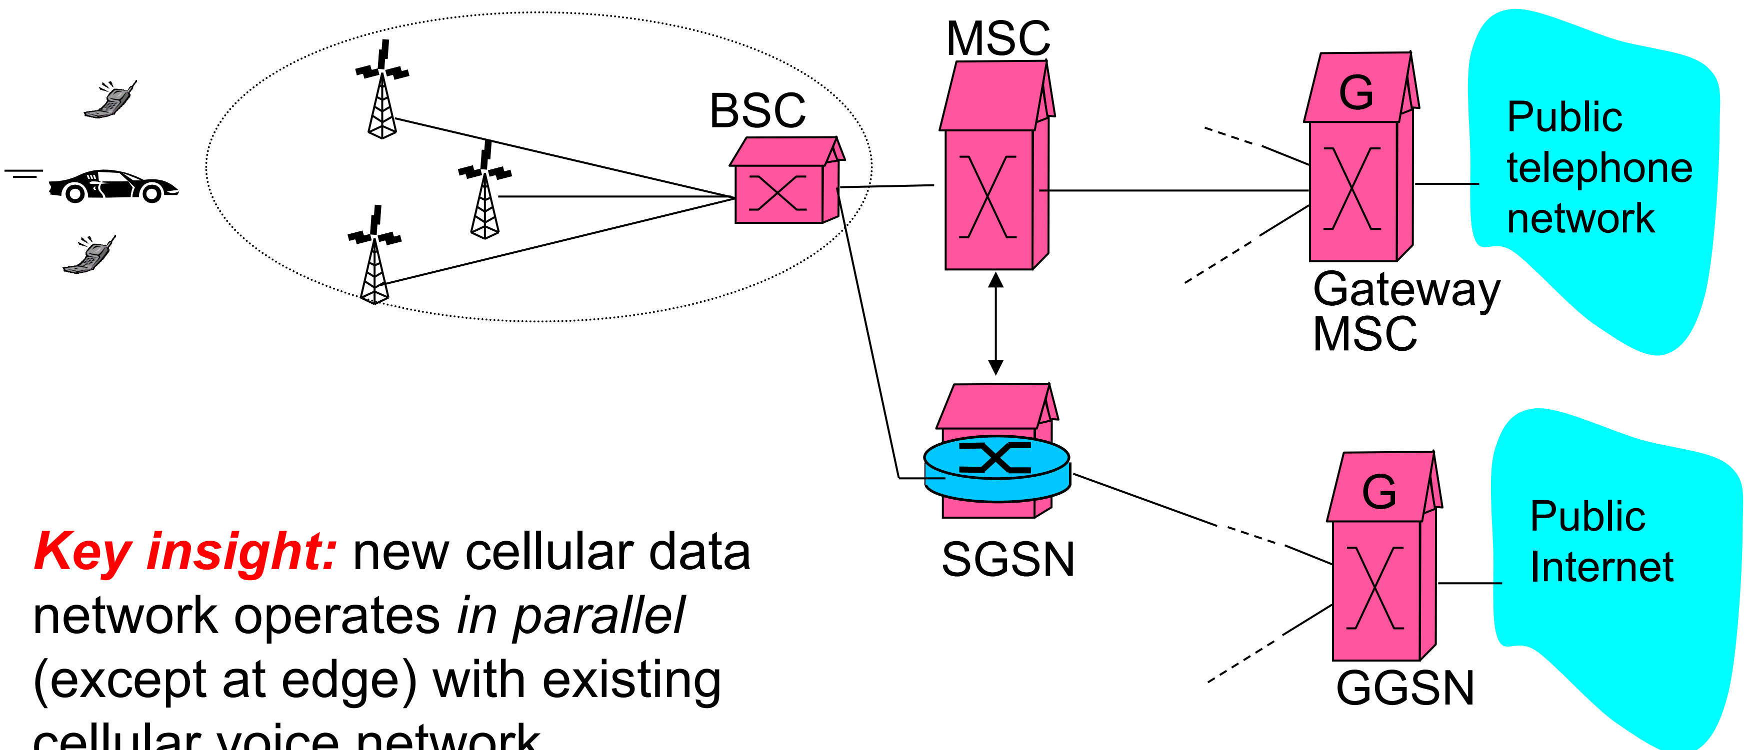

Key insight: new cellular data network operates *in parallel* (except at edge) with existing cellular voice network

- voice network unchanged in core
- data network operates in parallel

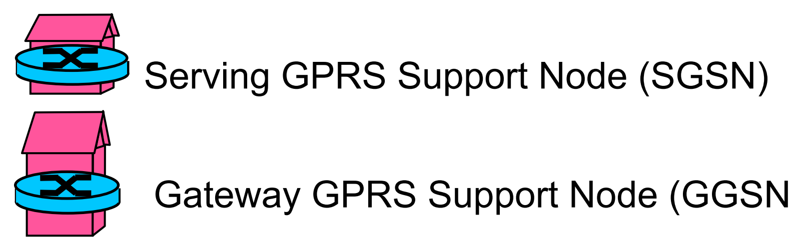

GPRS = General Packet Radio Service

Link Layer

1 Introduction and services

2 Multiple access protocols

3 Error detection and correction

4 Ethernet

Error Detection

EDC= Error Detection and Correction bits (redundancy)
D = Data protected by error checking, may include header fields

- Error detection not 100% reliable!
  - protocol may miss some errors, but rarely
  - larger EDC field yields better detection and correction

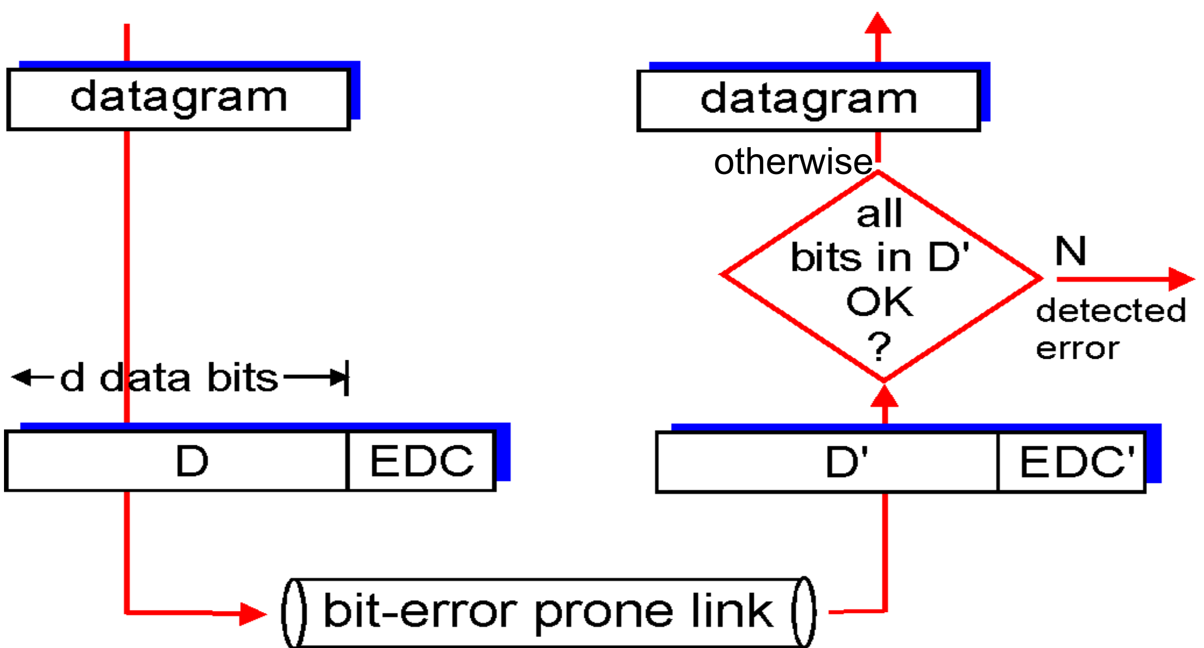

Parity Checking

Single Bit Parity:

Detect single bit errors

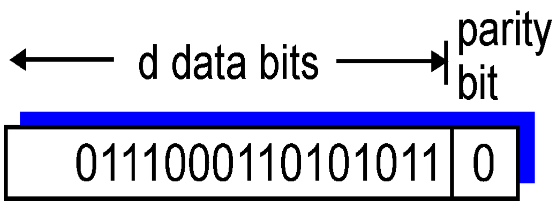

Two Dimensional Bit Parity:

Detect and correct single bit errors

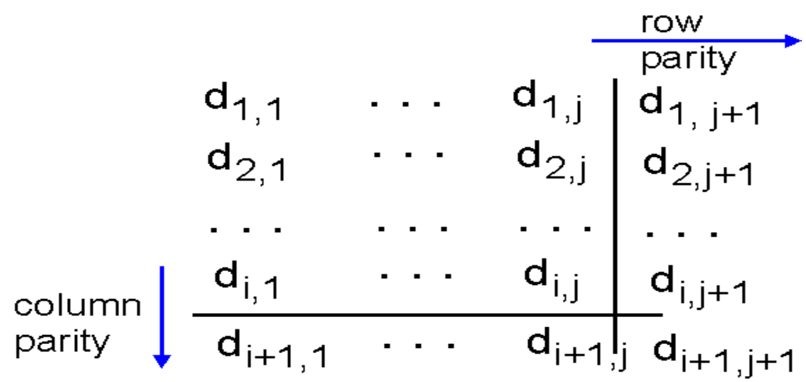

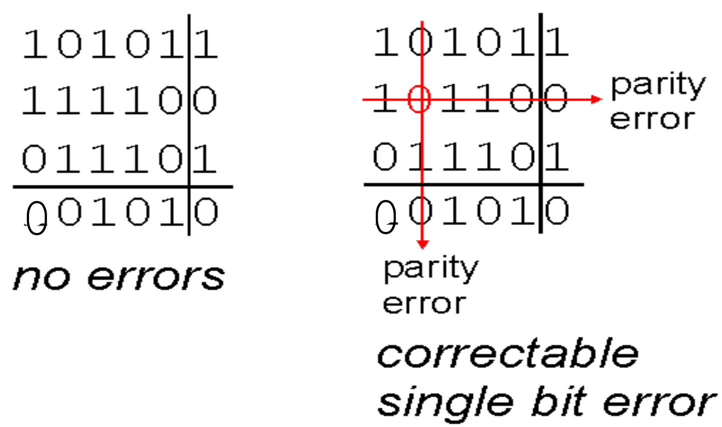

Internet checksum (review)

Goal: detect "errors" (e.g., flipped bits) in transmitted packet (note: used at transport layer only)

Sender:

treat segment contents as sequence of 16-bit integers

checksum: addition (1's complement sum) of segment contents

sender puts checksum value into UDP checksum field

Receiver:

compute checksum of received segment

check if computed checksum equals checksum field value:

- NO - error detected
- YES - no error detected. *But maybe errors nonetheless?*

Checksumming: Cyclic Redundancy Check

view data bits, D, as a binary number

choose r+1 bit pattern (generator), G

goal: choose r CRC bits, R, such that

- <D,R> exactly divisible by G (modulo 2)
- receiver knows G, divides <D,R> by G. If non-zero remainder: error detected!
- can detect all burst errors less than r+1 bits

widely used in practice (Ethernet, 802.11 WiFi, ATM)

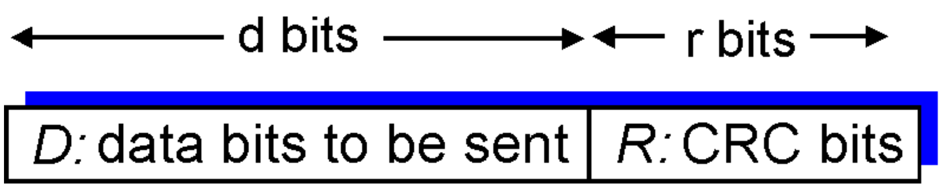

bit pattern

$$D * 2^r \quad XOR \quad R$$

mathematical formula

CRC Example

Want:

$D \cdot 2^r$ XOR R = nG

equivalently:

$D \cdot 2^r$ = nG XOR R

equivalently:

if we divide $D \cdot 2^r$ by G, want remainder R

$$R = \text{remainder}\left[\frac{D \cdot 2^r}{G}\right]$$

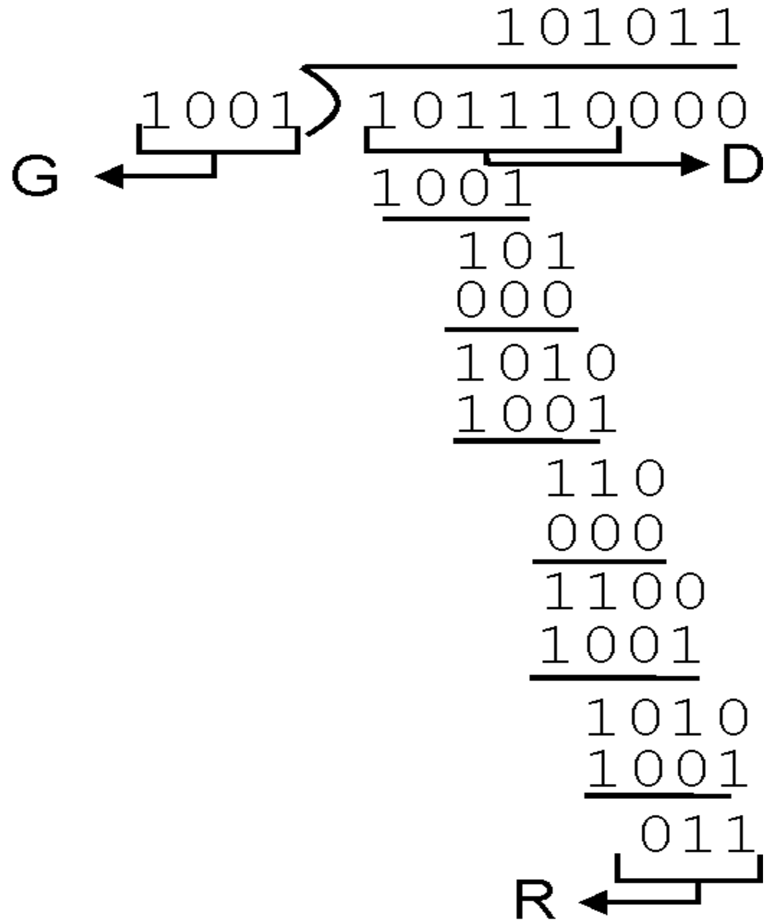

Ethernet

- Classic Ethernet »
- Switched/Fast Ethernet »
- Gigabit/10 Gigabit Ethernet »

Classic Ethernet (1) – Physical Layer

One shared coaxial cable to which all hosts attached

- Up to 10 Mbps, with Manchester encoding
- Hosts ran the classic Ethernet protocol for access

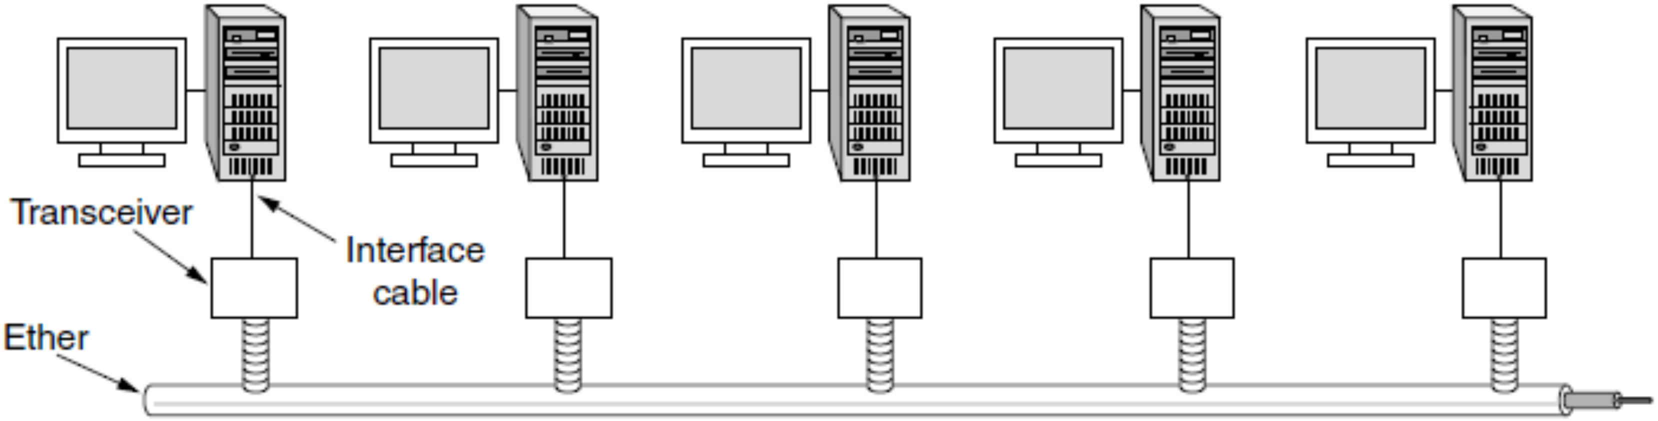

Classic Ethernet (2) – MAC

MAC protocol is 1-persistent CSMA/CD (earlier)

- Random delay (backoff) after collision is computed with BEB (Binary Exponential Backoff)
- Frame format is still used with modern Ethernet.

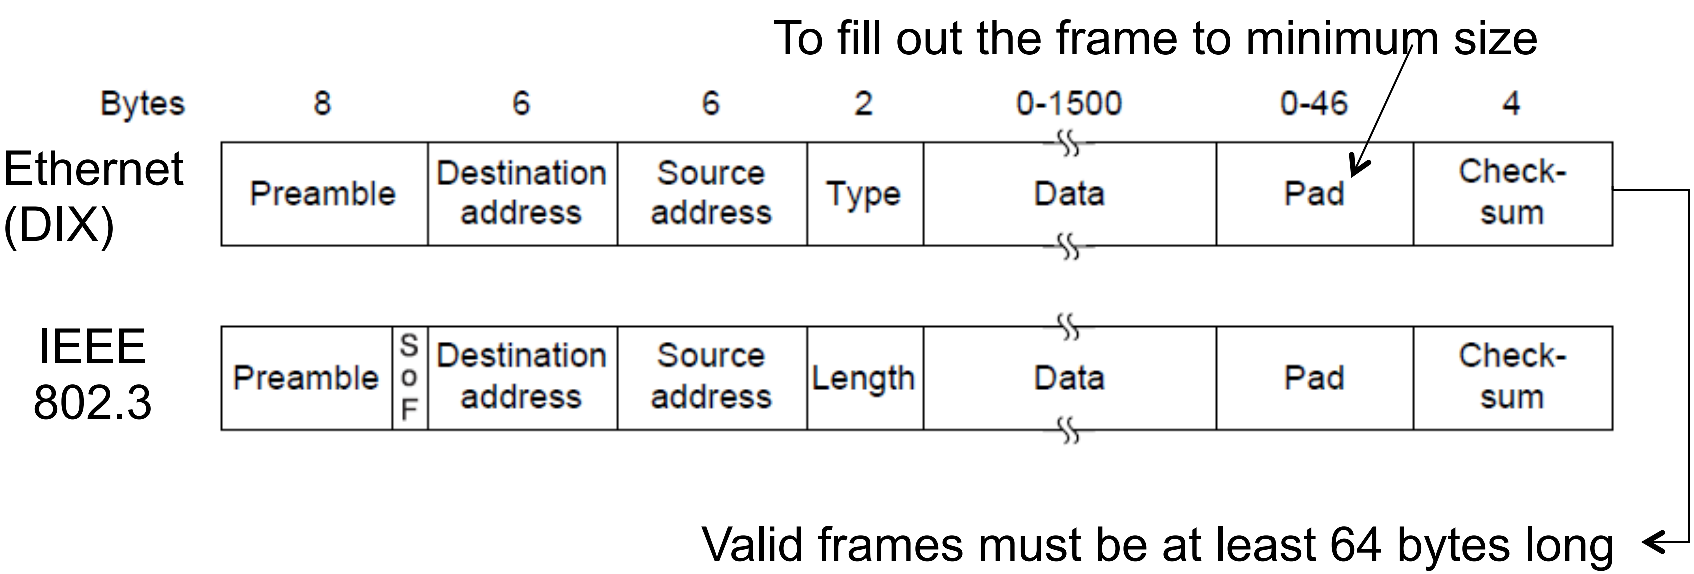

Classic Ethernet (3-1) – MAC

Collisions can occur and take as long as 2τ to detect

- τ is the time it takes to propagate over the Ethernet
- Leads to minimum packet size for reliable detection

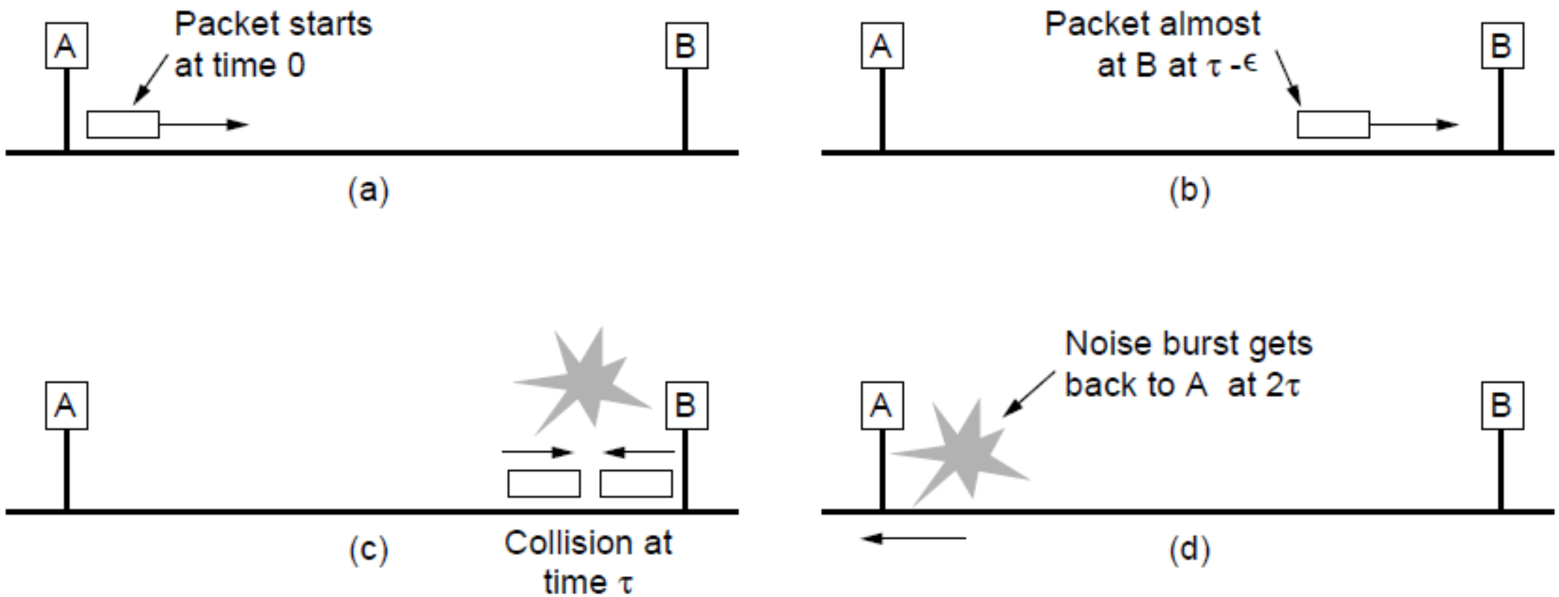

When B detects that a collision has occurred, it generates a 48-bit noise burst to warn all other stations

At about time 2τ , the sender sees the noise burst and aborts its transmission

It then waits a random time before trying again

Classic Ethernet (3-2) – MAC

Collisions can occur and take as long as 2τ to detect

- If a station tries to transmit a very short frame, it is conceivable that a collision occurs
- But…the sender will incorrectly conclude that the frame was successfully sent

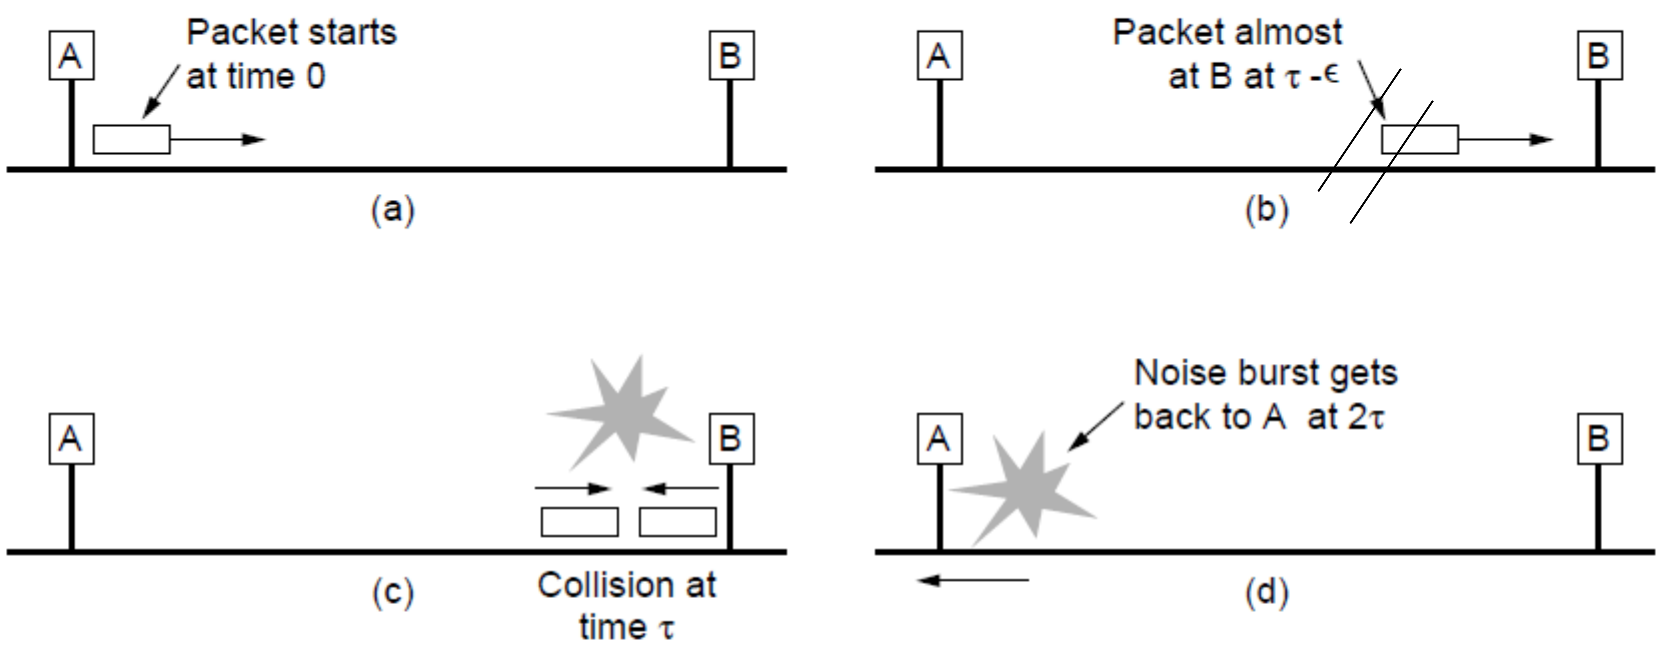

Classic Ethernet (3-3) – MAC

Collisions can occur and take as long as 2τ to detect

- To prevent this situation from occurring, all frames must take more than 2τ to send
- So…transmission is still taking place when the noise burst gets back to the sender

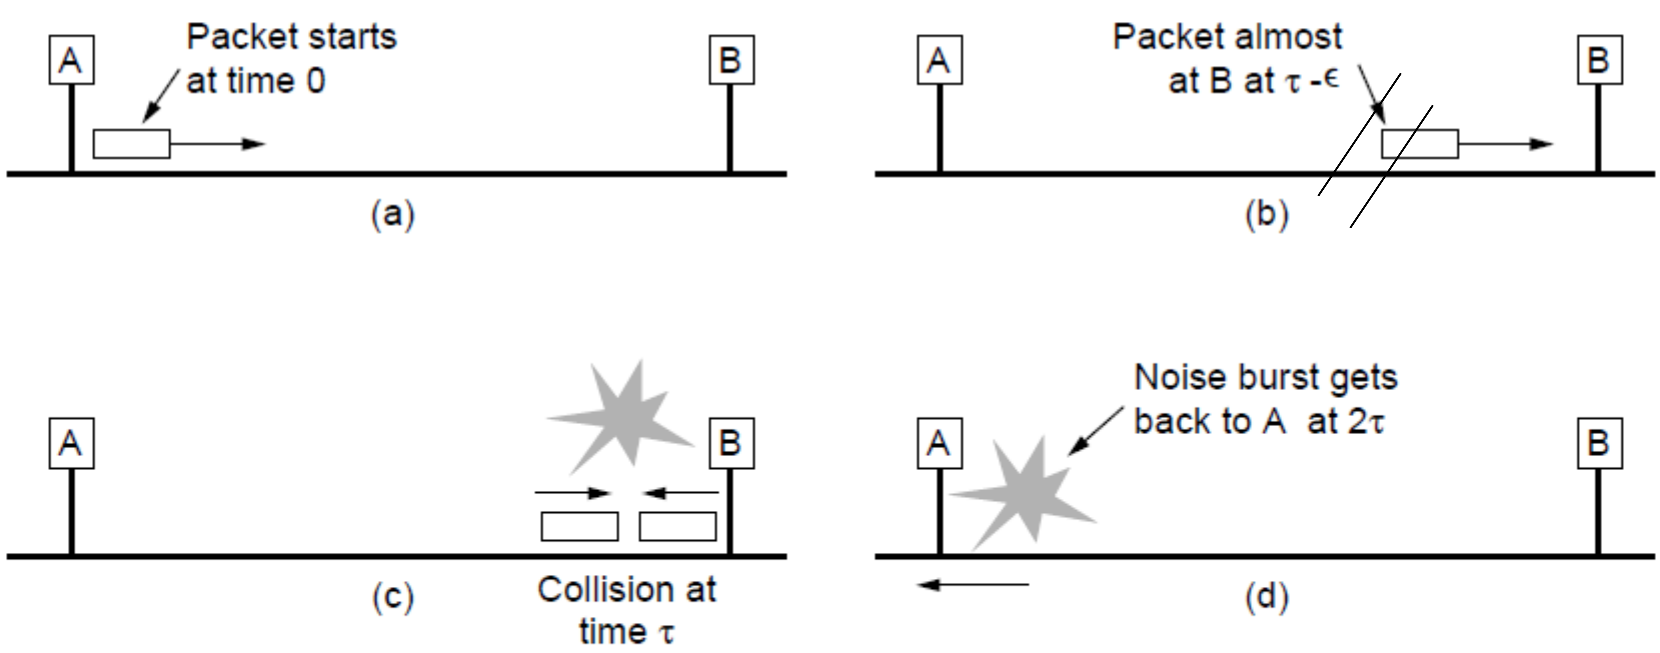

Classic Ethernet (4) – Performance

Efficient for large frames, even with many senders

- Degrades for small frames (and long LANs)

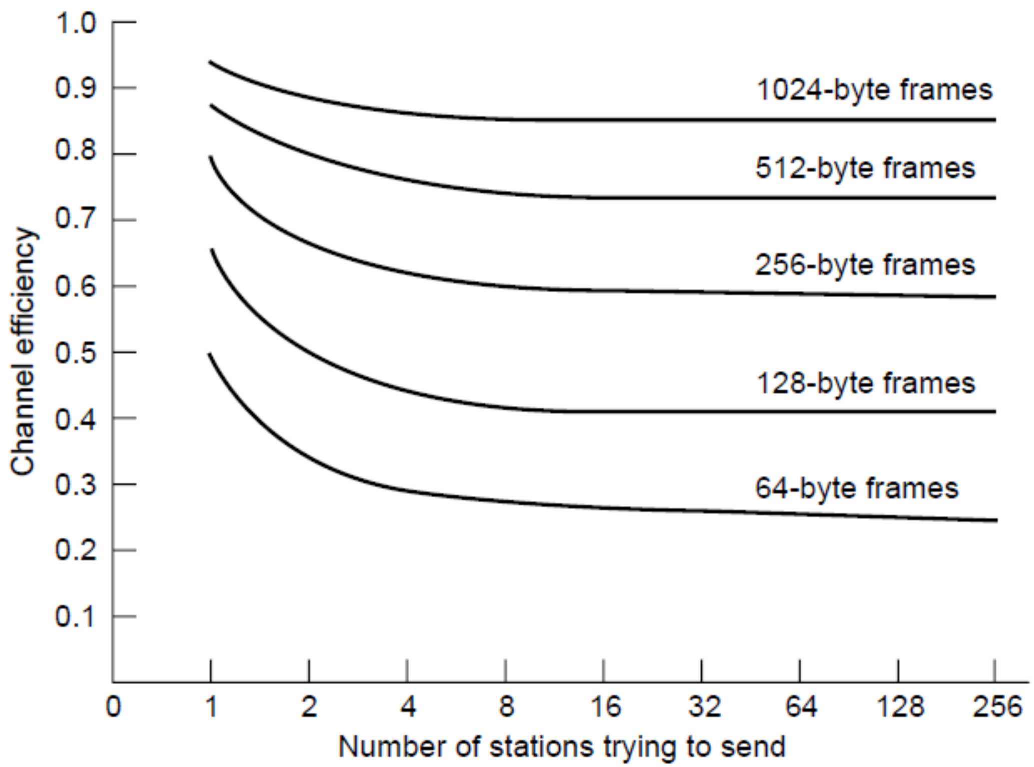

10 Mbps Ethernet,
64 byte min. frame

Switched/Fast Ethernet (1)

- Hubs wire all lines into a single CSMA/CD domain
- Switches isolate each port to a separate domain
  - Much greater throughput for multiple ports
  - No need for CSMA/CD with full-duplex lines

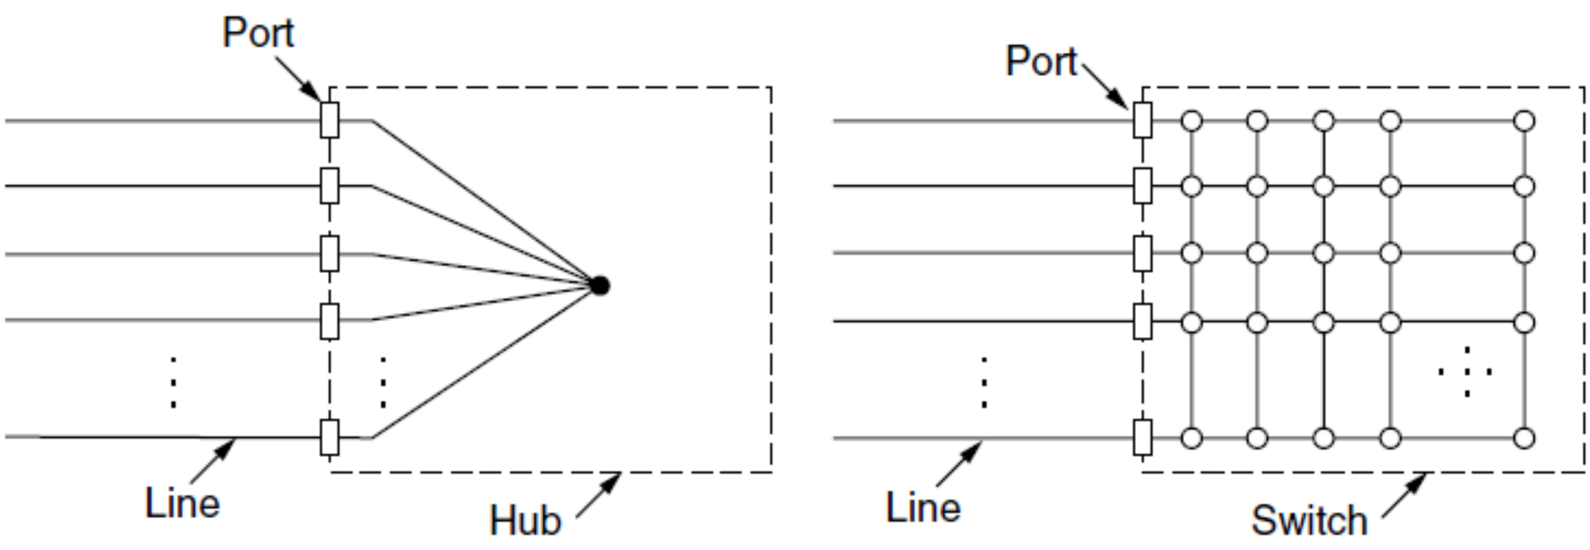

Switched/Fast Ethernet (2)

Switches can be wired to computers, hubs and switches

- Hubs concentrate traffic from computers

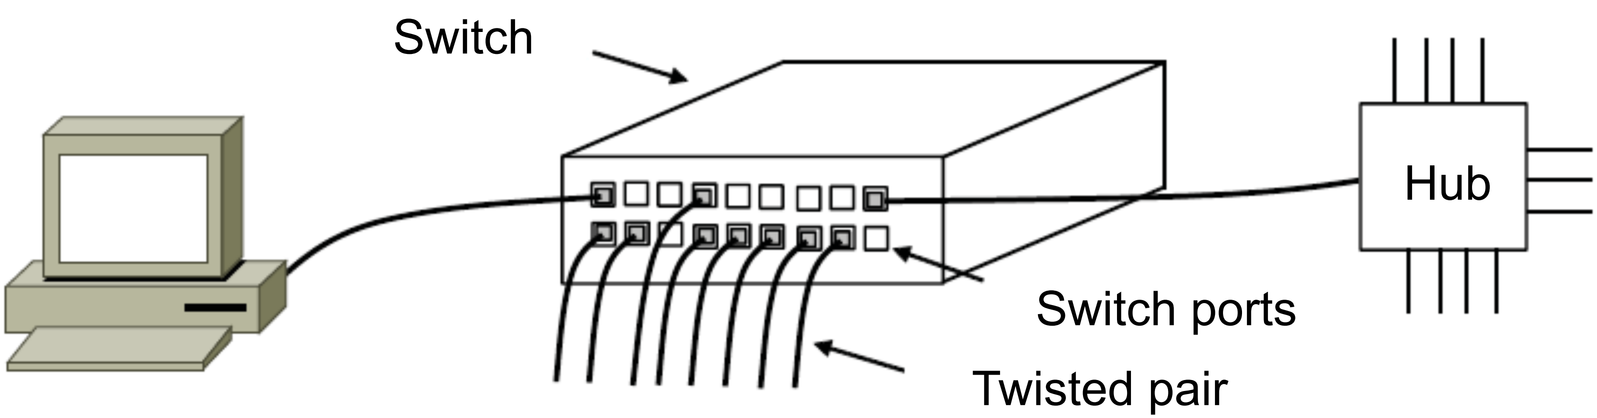

Switched/Fast Ethernet (3)

Fast Ethernet extended Ethernet from 10 to 100 Mbps

- Twisted pair (with Cat 5) dominated the market

Name	Cable	Max. segment	Advantages
100Base-T4	Twisted pair	100 m	Uses category 3 UTP
100Base-TX	Twisted pair	100 m	Full duplex at 100 Mbps (Cat 5 UTP)
100Base-FX	Fiber optics	2000 m	Full duplex at 100 Mbps; long runs

Gigabit / 10 Gigabit Ethernet (1)

Switched Gigabit Ethernet is now the garden variety

- With full-duplex lines between computers/switches

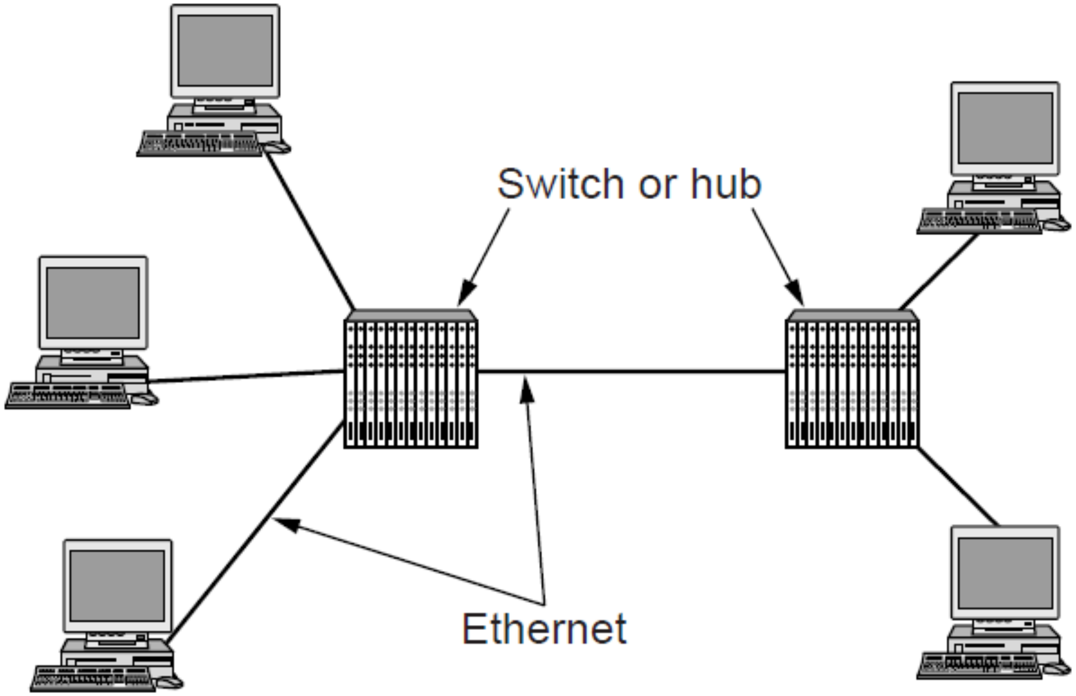

Gigabit / 10 Gigabit Ethernet (1)

- Gigabit Ethernet is commonly run over twisted pair

Name	Cable	Max. segment	Advantages
1000Base-SX	Fiber optics	550 m	Multimode fiber (50, 62.5 microns)
1000Base-LX	Fiber optics	5000 m	Single (10 µ) or multimode (50, 62.5 µ)
1000Base-CX	2 Pairs of STP	25 m	Shielded twisted pair
1000Base-T	4 Pairs of UTP	100 m	Standard category 5 UTP

- 10 Gigabit Ethernet is being deployed where needed

Name	Cable	Max. segment	Advantages
10GBase-SR	Fiber optics	Up to 300 m	Multimode fiber (0.85µ)
10GBase-LR	Fiber optics	10 km	Single-mode fiber (1.3µ)
10GBase-ER	Fiber optics	40 km	Single-mode fiber (1.5µ)
10GBase-CX4	4 Pairs of twinax	15 m	Twinaxial copper
10GBase-T	4 Pairs of UTP	100 m	Category 6a UTP

- 40/100 Gigabit Ethernet is under development

Wireless LANs

- 802.11 architecture/protocol stack »
- 802.11 physical layer »
- 802.11 MAC »
- 802.11 frames »

802.11 Architecture/Protocol Stack (1)

Wireless clients associate to a wired AP (Access Point)

- Called infrastructure mode; there is also ad-hoc mode with no AP.

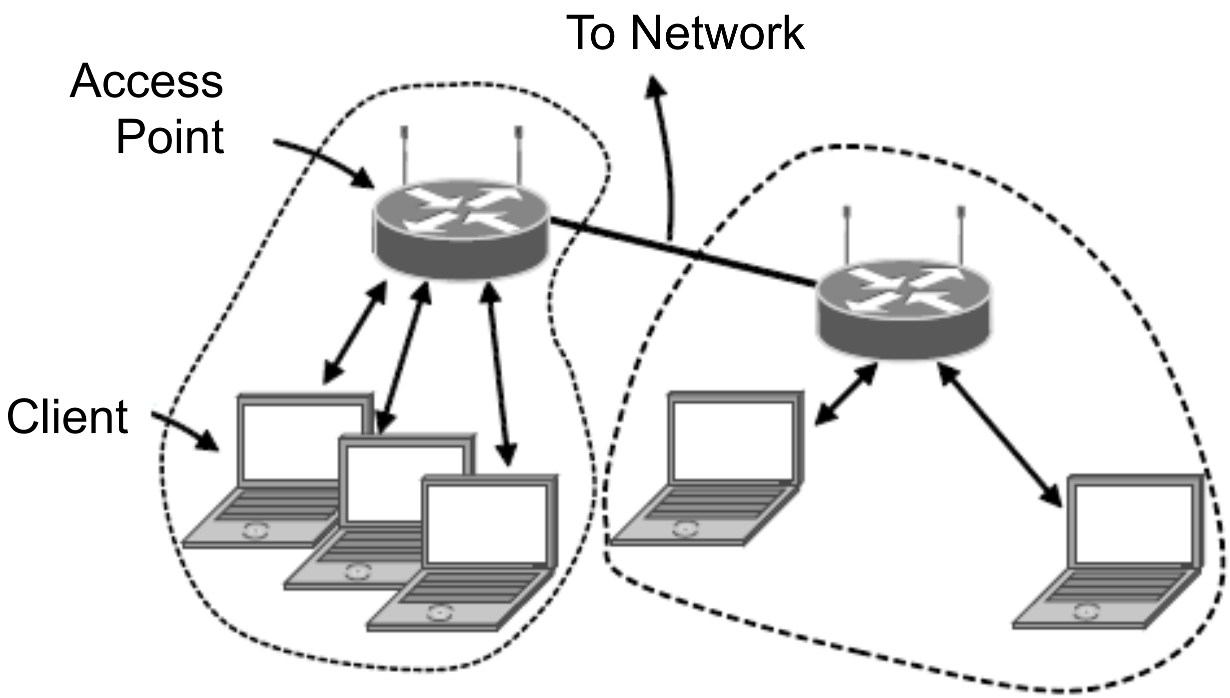

802.11 Architecture/Protocol Stack (2)

MAC is used across different physical layers

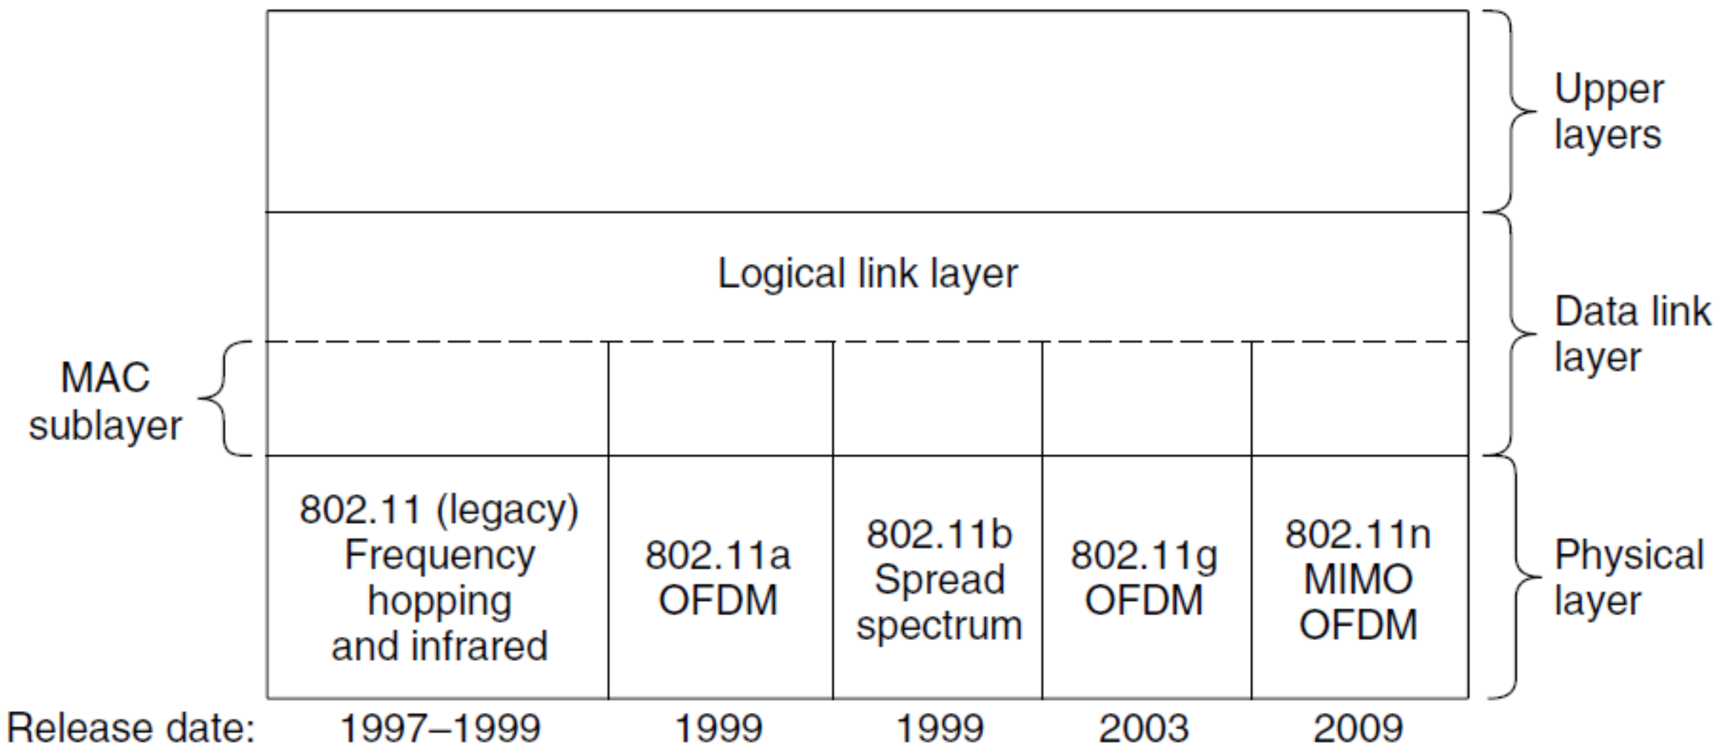

802.11 physical layer

- NICs are compatible with multiple physical layers
  - E.g., 802.11 a/b/g

Name	Technique	Max. Bit Rate
802.11b	Spread spectrum, 2.4 GHz	11 Mbps
802.11g	OFDM, 2.4 GHz	54 Mbps
802.11a	OFDM, 5 GHz	54 Mbps
802.11n	OFDM with MIMO, 2.4/5 GHz	600 Mbps

Data Link Layer Switching

- Uses of Bridges »
- Learning Bridges »
- Spanning Tree »
- Repeaters, hubs, bridges, .., routers, gateways »
- Virtual LANs »

Uses of Bridges

Common setup is a building with centralized wiring

- Bridges (switches) are placed in or near wiring closets

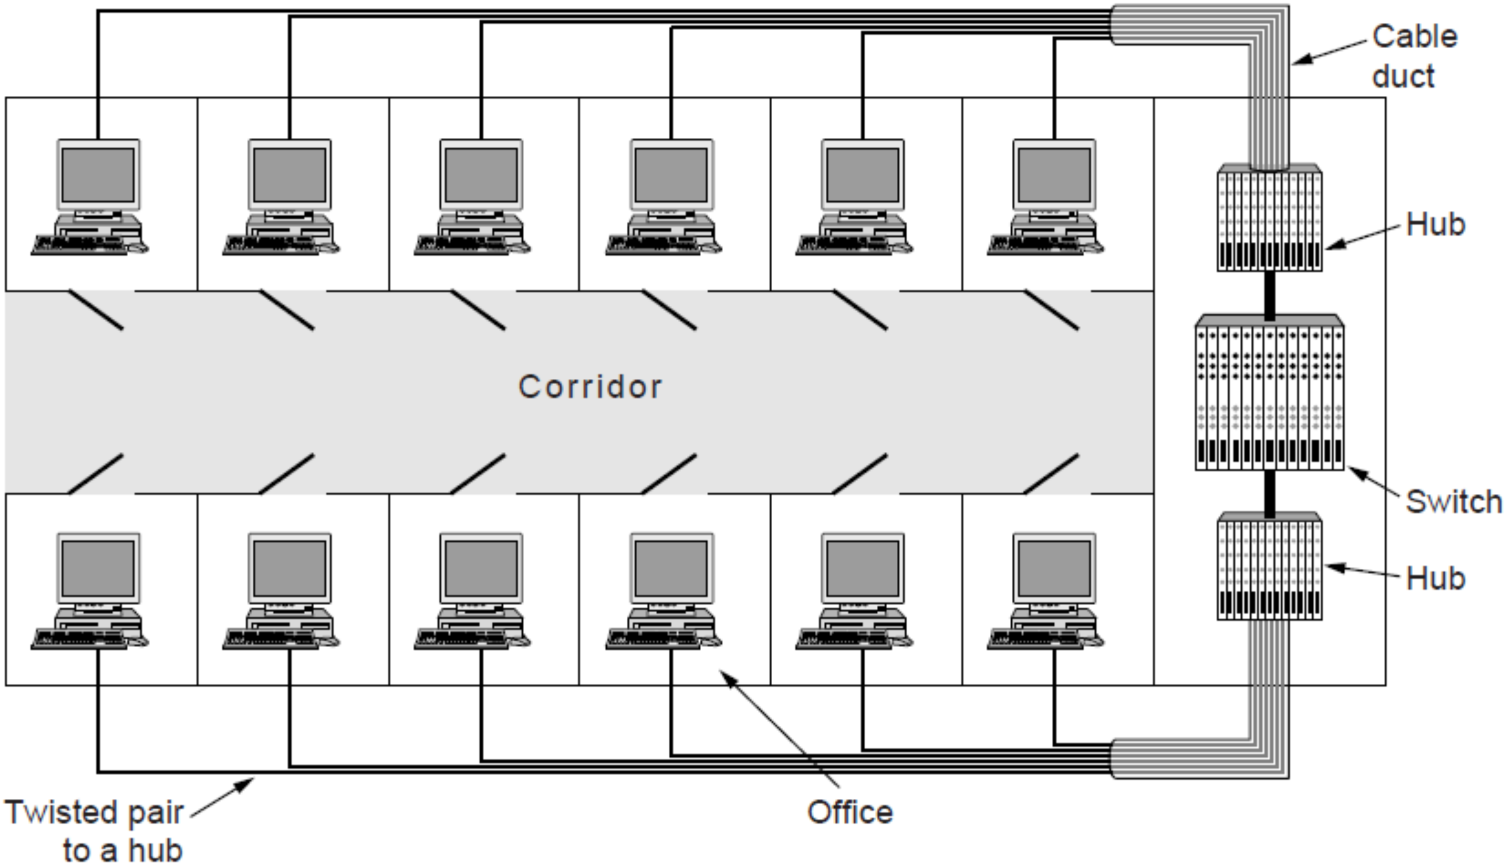

Learning Bridges (1)

A bridge operates as a switched LAN (not a hub)

- Computers, bridges, and hubs connect to its ports

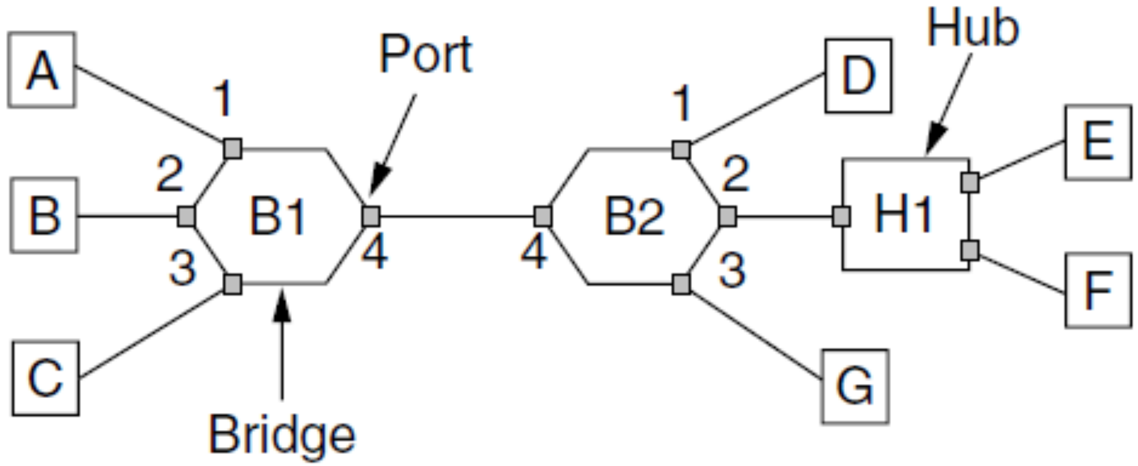

Learning Bridges (2)

Backward learning algorithm picks the output port:

- Associates source address on frame with input port
- Frame with destination address sent to learned port
- Unlearned destinations are sent to all other ports

Needs no configuration

- Forget unused addresses to allow changes
- Bandwidth efficient for two-way traffic

Learning Bridges (3)

Bridges extend the Link layer:

- Use but don't remove Ethernet header/addresses
- Do not inspect Network header

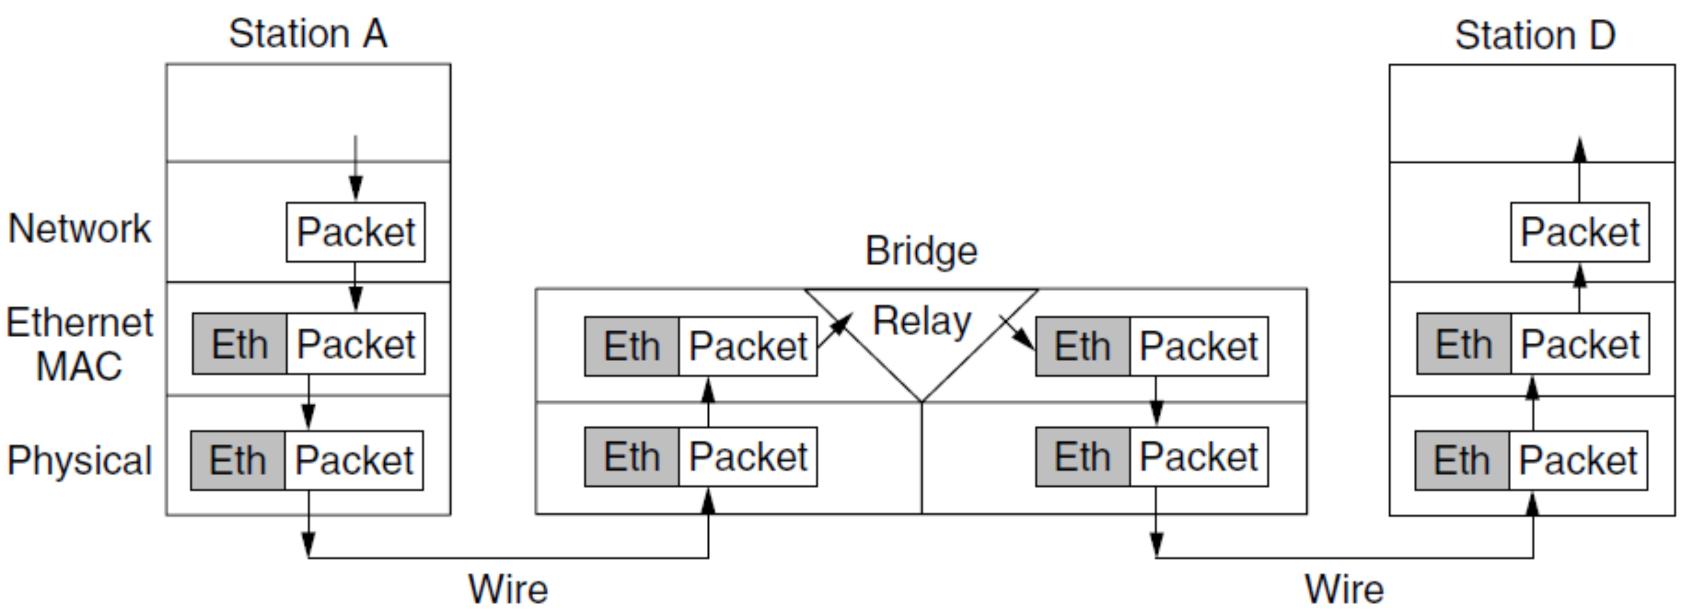

Spanning Tree (1) – Problem

Bridge topologies with loops and only backward learning will cause frames to circulate for ever

- Need spanning tree support to solve problem

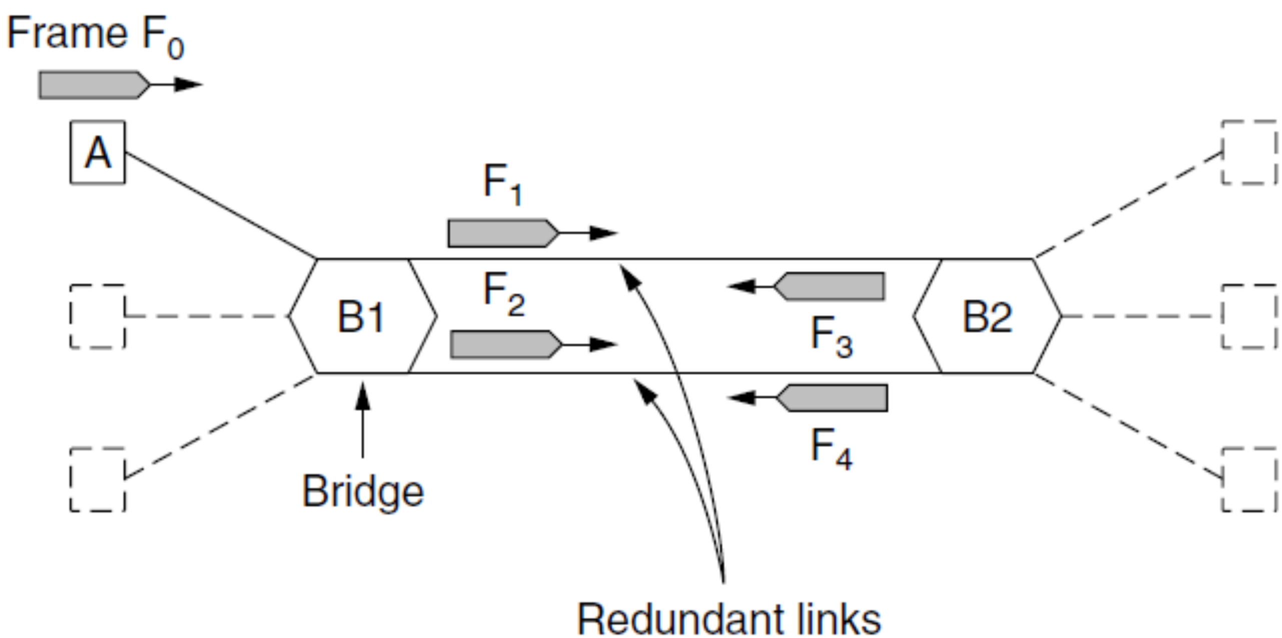

Spanning Tree (2) – Algorithm

- Subset of forwarding ports for data is use to avoid loops
- Selected with the spanning tree distributed algorithm by Perlman

I think that I shall never see
A graph more lovely than a tree.
A tree whose crucial property
Is loop-free connectivity.
A tree which must be sure to span.
So packets can reach every LAN.
First the Root must be selected
By ID it is elected.
Least cost paths from Root are traced
In the tree these paths are placed.
A mesh is made by folks like me
Then bridges find a spanning tree.

– Radia Perlman, 1985.

Spanning Tree (3) – Example

Root: Each node broadcast its serial number. The ones with the lowest serial number becomes the root

After the algorithm runs:

- B1 is the root, two dashed links are turned off
- B4 uses link to B2 (lower than B3 also at distance 1)
- B5 uses B3 (distance 1 versus B4 at distance 2)

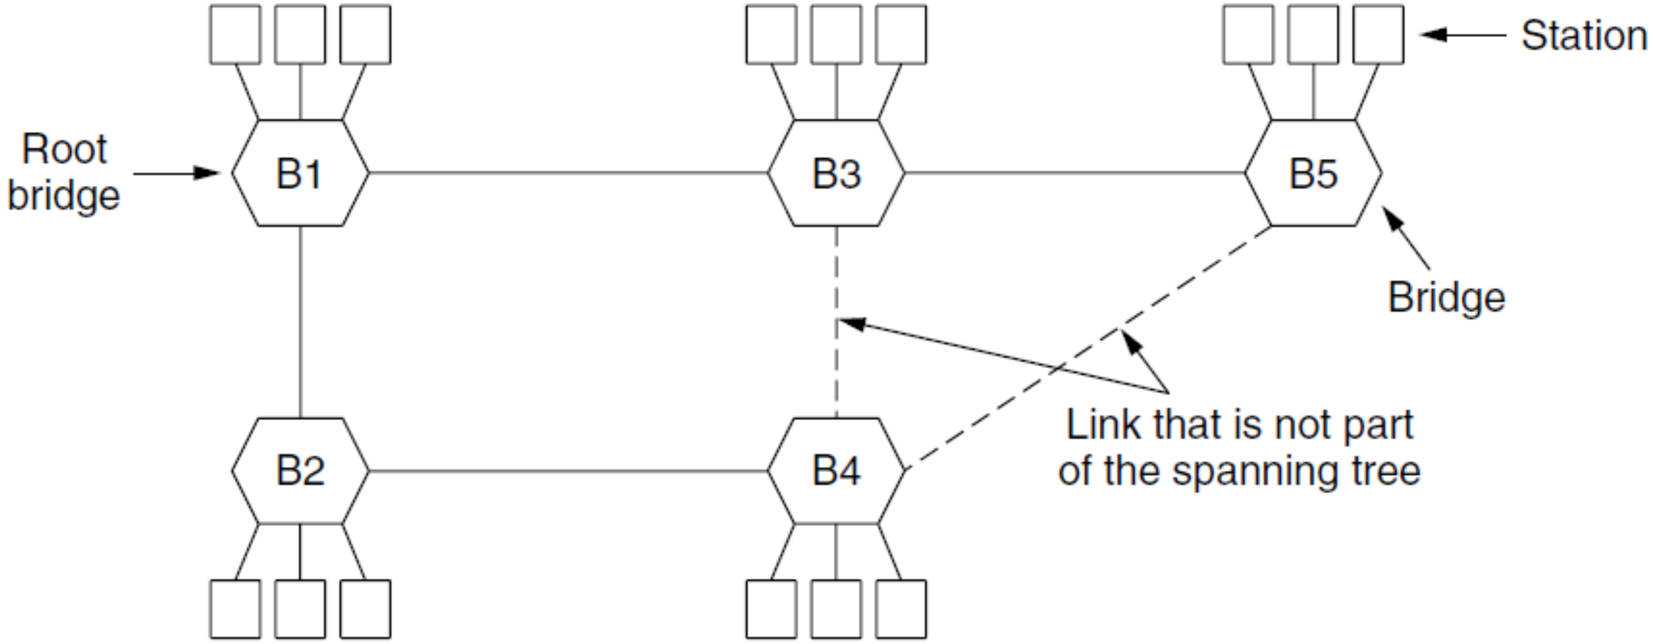

Repeaters, Hubs, Bridges, Switches, Routers, & Gateways

Devices are named according to the layer they process

- A bridge or LAN switch operates in the Link layer

Layer	Device
Application layer	Application gateway
Transport layer	Transport gateway
Network layer	Router
Data link layer	**Bridge, switch**
Physical layer	Repeater, hub