Medium Access Control Sublayer

Chapter 4

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

Revised: August 2011

The MAC Sublayer

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

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

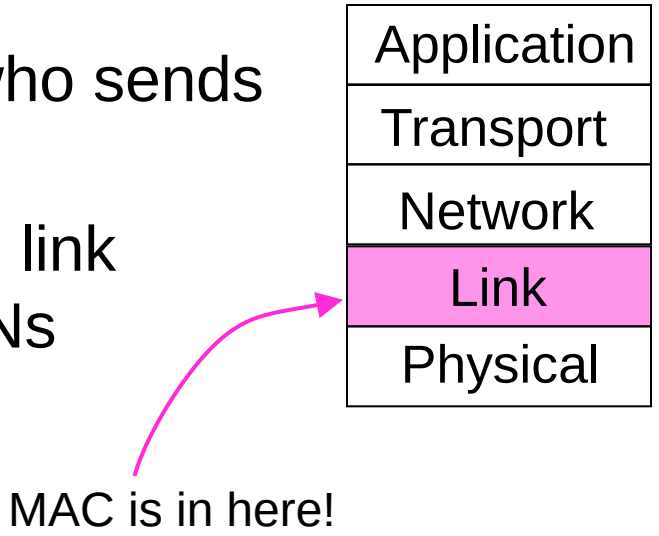

Channel Allocation Problem (1)

For fixed channel and traffic from N users

- Divide up bandwidth using FDM, TDM, CDMA, etc.
  - FDM and TDM problematic with large # of senders or bursty traffic
- These are static allocations, e.g., FM radio
- Works well for voice

This static allocation performs poorly for bursty traffic

- Most data transmissions are inherently bursty
- Allocation to any given user will sometimes go unused = wasteful

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**
Independent traffic	Often not a good model, but permits analysis
Single channel	No external way to coordinate senders
Observable collisions (2+ sending simultaneously)	Needed for reliability; mechanisms vary
Continuous or slotted time	Slotting (time divided up into discrete intervals) may improve performance
Carrier sense	Can improve performance if available

Multiple Access Protocols

Two basic strategies for channel acquisition in a broadcast network:

1. Contention (e.g., Aloha, CSMA) – preferable for low load because of its low delay characteristics
2. Collision Free Protocols – preferable at high load

Each strategy can be rated as to how well it does with respect to the two important performance measures:

1. Delay at low load
2. Channel efficiency at high load

- ALOHA »
- CSMA (Carrier Sense Multiple Access) »
- Collision-free protocols »
- Limited-contention protocols »
  - Attempts to be middle ground between contention and collision-free
- Wireless LAN protocols »

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

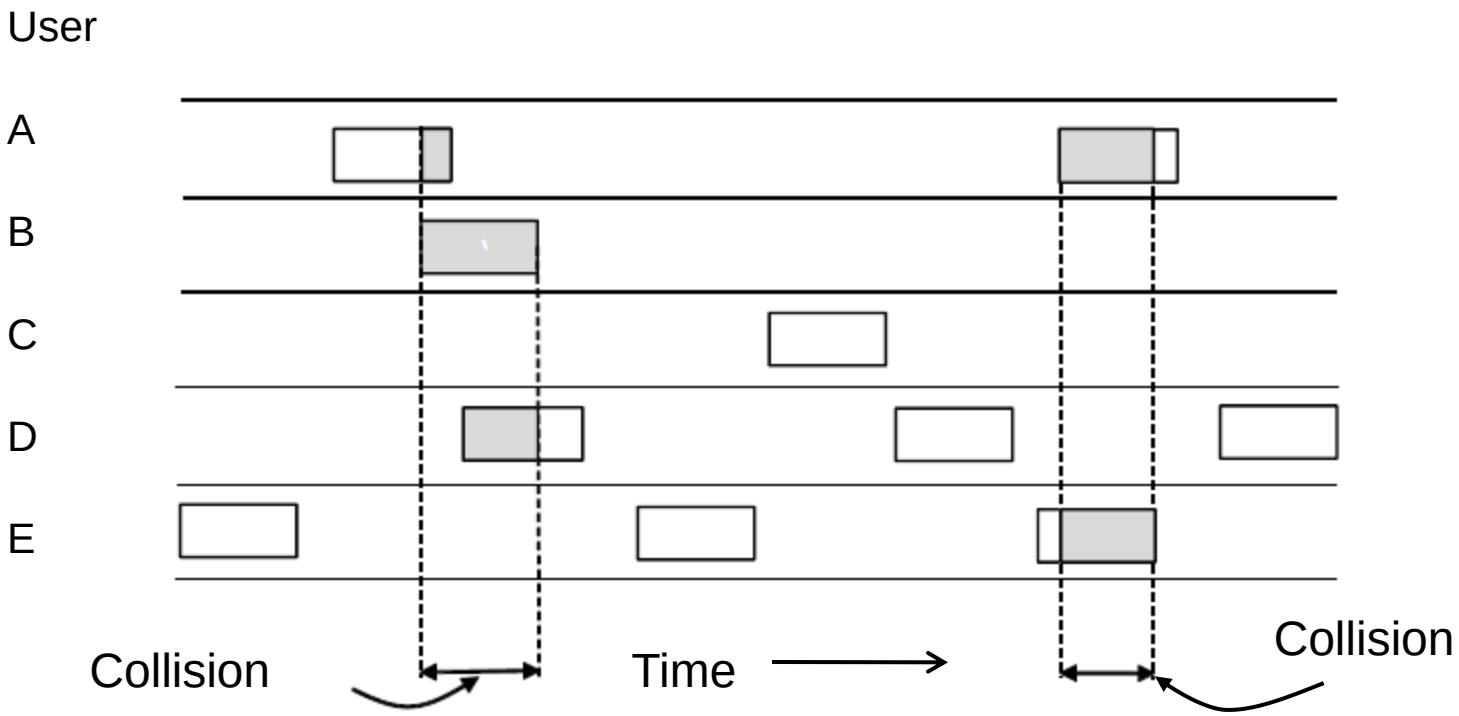

ALOHA (2)

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

- Synchronizing senders to slots can reduce collisions

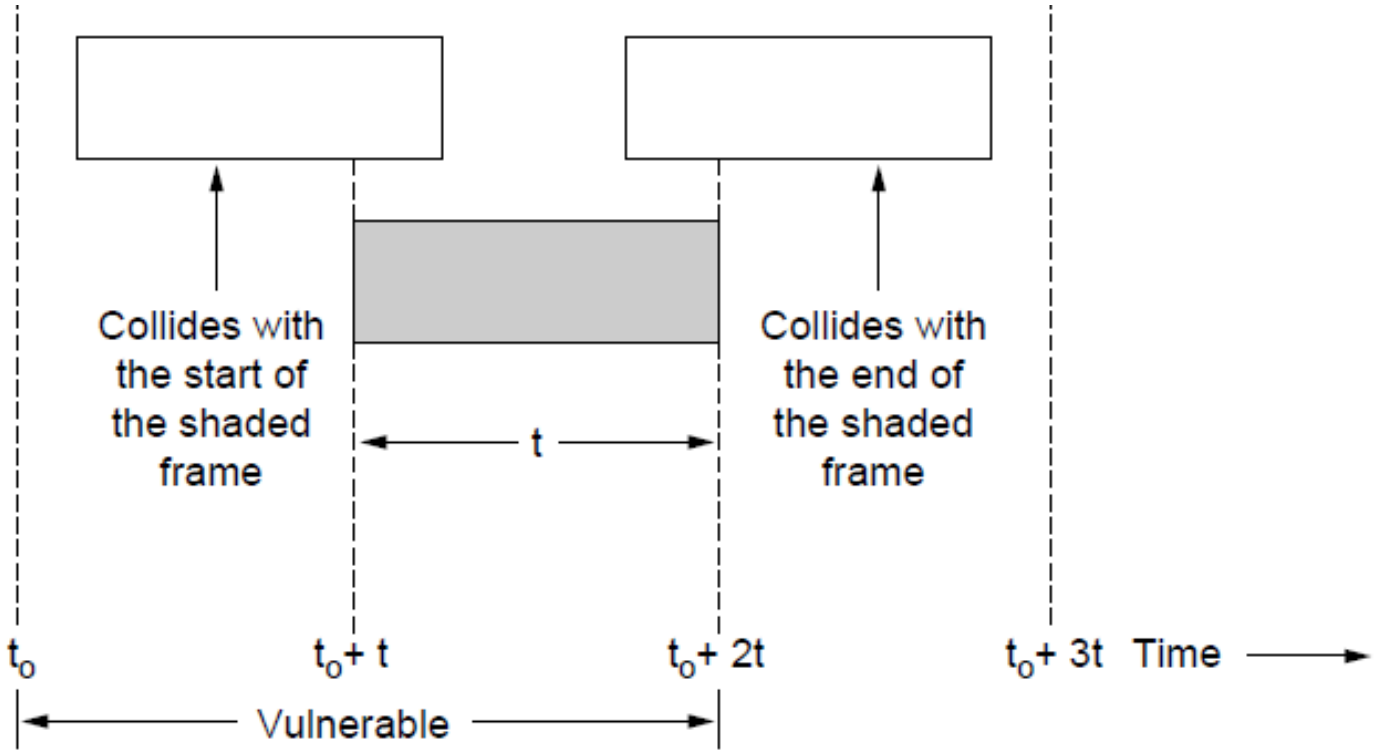

ALOHA (3)

Slotted ALOHA is twice as efficient as pure ALOHA

- How to synch slots? Participants must agree in advance (e.g., one station emit a pip at the start of each interval like a clock)
- Low load wastes slots, high loads causes collisions
- Efficiency up to 1/e (37%) for random traffic models

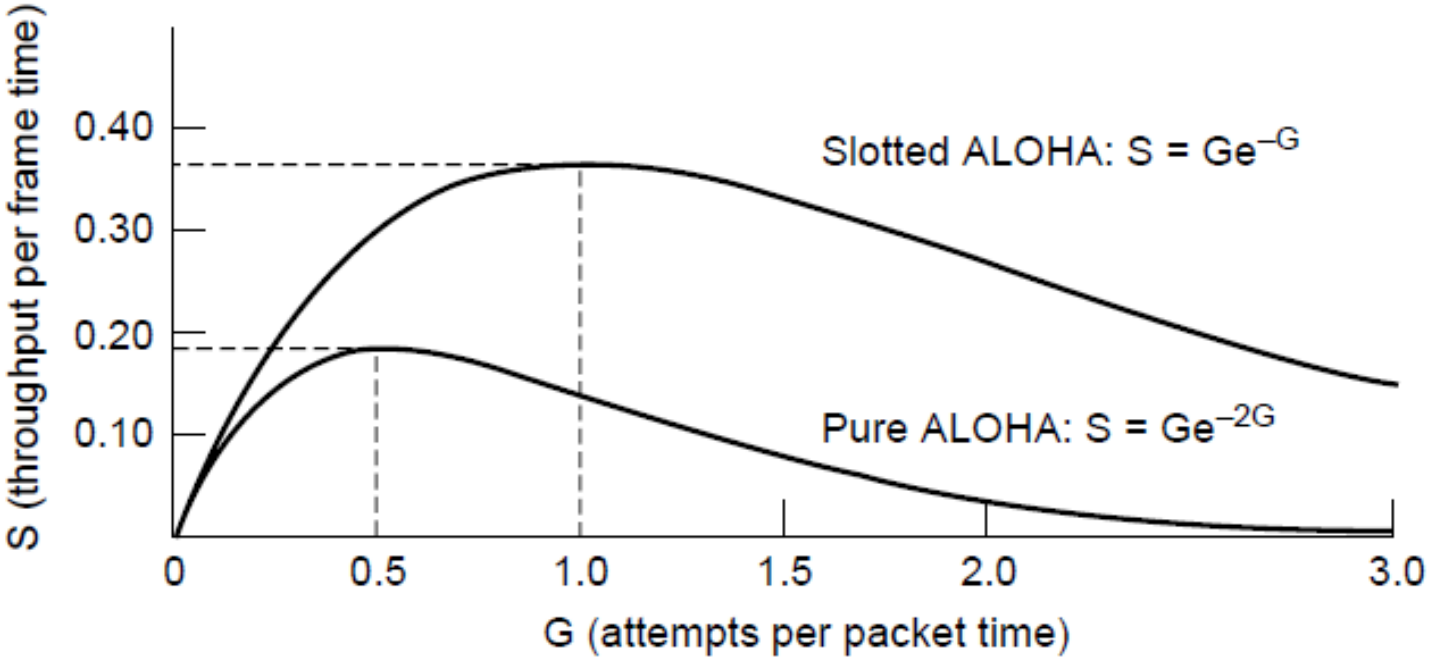

CSMA (1)

Carrier Sense Multiple Access (CSMA) improves on ALOHA by sensing the channel!

- User doesn't send if it senses someone else is already sending (i.e., listens first before sending)

Variations (within CSMA) on what to do if the channel is busy:

1. 1-persistent (greedy) sends as soon as idle
   - Approach used by (classic) Ethernet v2
2. Non-persistent waits a random time then tries again
3. p-persistent sends with probability p when idle
   - Leverages slotted channels to determine probabilities
   - Binary exponential backoff variant used by IEEE 802.3 (modern Ethernet; see page 285)

CSMA (2) – Persistence

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

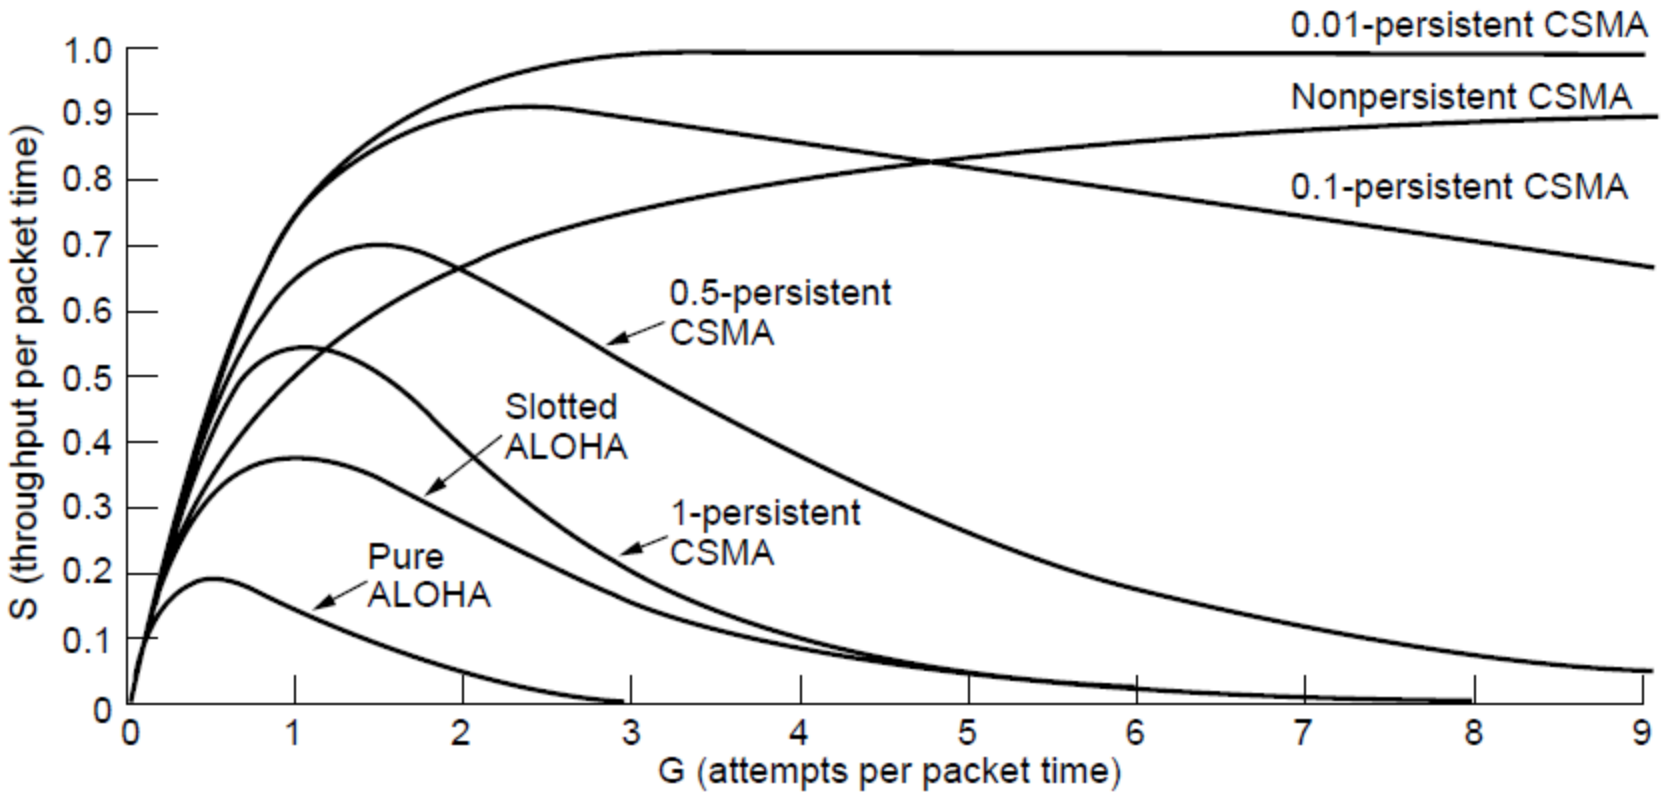

CSMA (3) – Collision Detection

CSMA/CD improvement is to detect/abort collisions

- If station detects a collision, it aborts sending and waits a random period of time before trying again
- Key issue: what if 2 stations transmit simultaneously? How long before they realize this? Worst case is 2*T where T=max propagation time on that LAN
- CSMA/CD is like Slotted Aloha with a slot width of 2T (1km coax T=5 usec)
- Reduced contention times improve performance

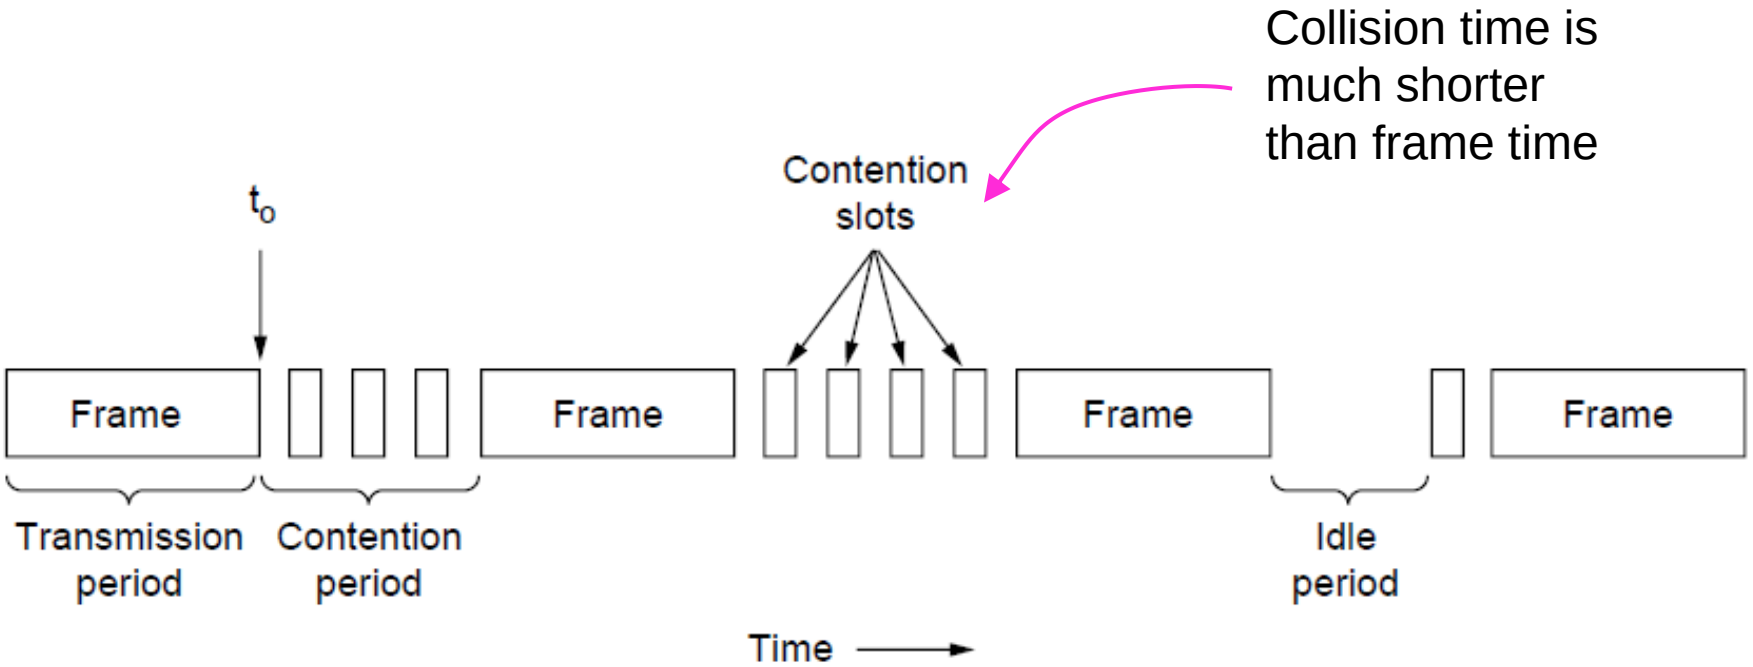

Collision-Free (1) – Bitmap

While CSMA/CD dramatically reduces collisions, any collision reduces bandwidth and make time-to-send variable, potentially impacting real-time traffic

Collision-free protocols avoid collisions entirely

- Key: Senders must know when it is their turn to send

The basic bit-map protocol is **a reservation protocol**:

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

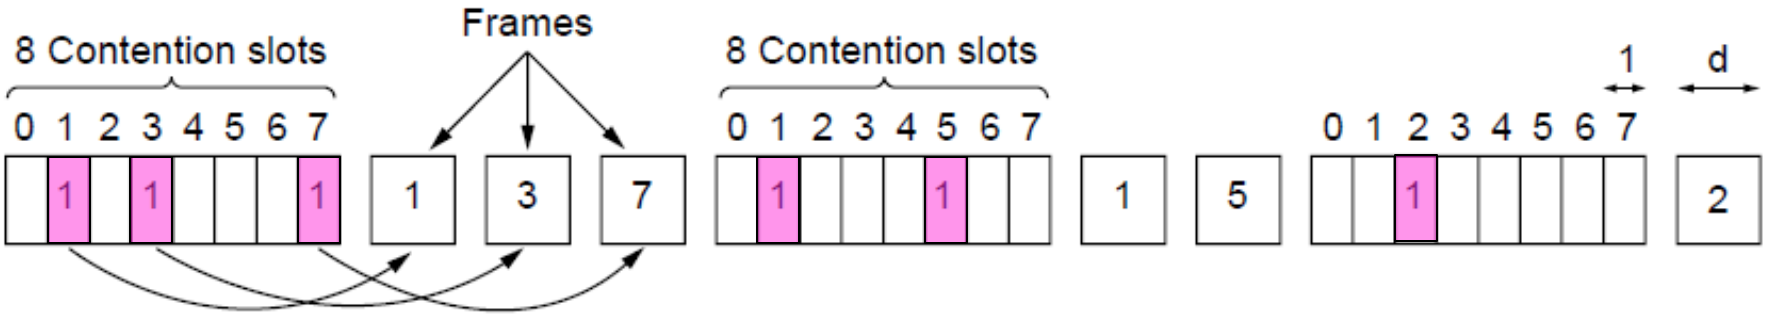

Collision-Free (2) – Token Ring

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

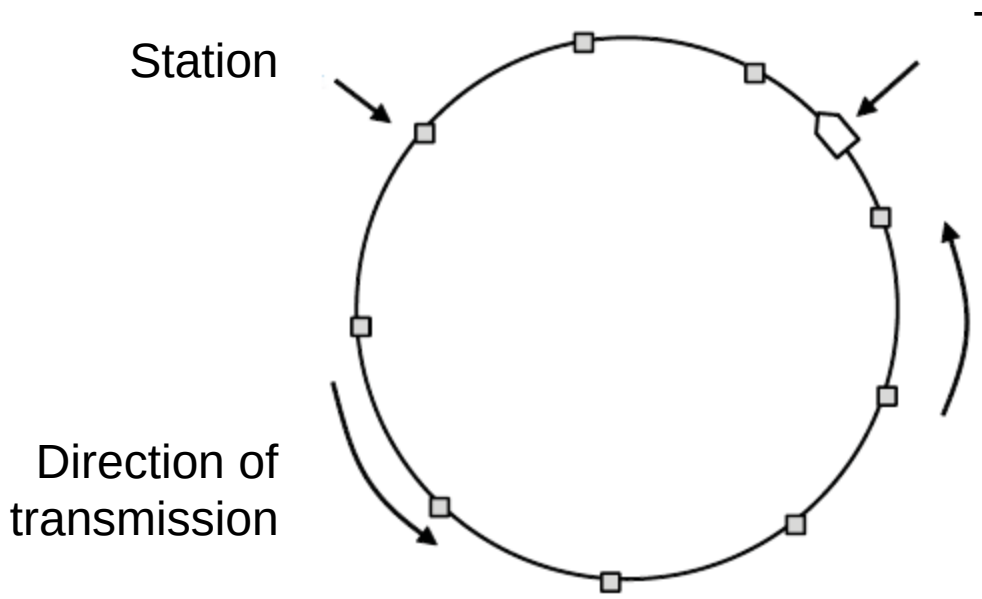

Token Ring:
- IEEE 802.5
- FDDI

Token Bus:
- IEEE 802.4

Collision-Free (3) – Countdown

Problem with bitmap and token ring is that it does not scale to LANs with thousands of nodes. Binary countdown improves on the bitmap protocol for systems having minimal transmission delays

- Stations send their address in contention slot (log N bits instead of N bits)
- Medium ORs bits for simultaneous transmissions; stations give up when their addr has a “0” but see a “1”
- Station that sees its full address is next to send
- Higher numbered stations intrinsically have a higher priority

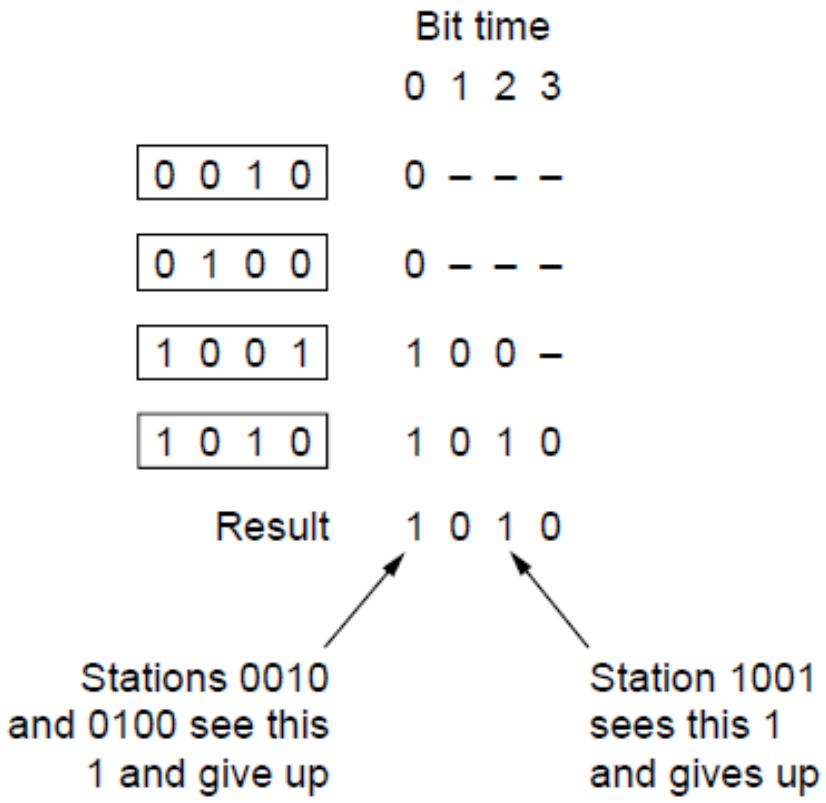

Limited-Contention Protocols (1)

Middle ground between Contention and Collision Free approaches:

- Contention at low load is preferential because it provides low delay
- Collision-free at high load is preferential because it provides good channel efficiency

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

- Apply contention protocols within those small groups
- Avoids wastage due to idle periods and collisions

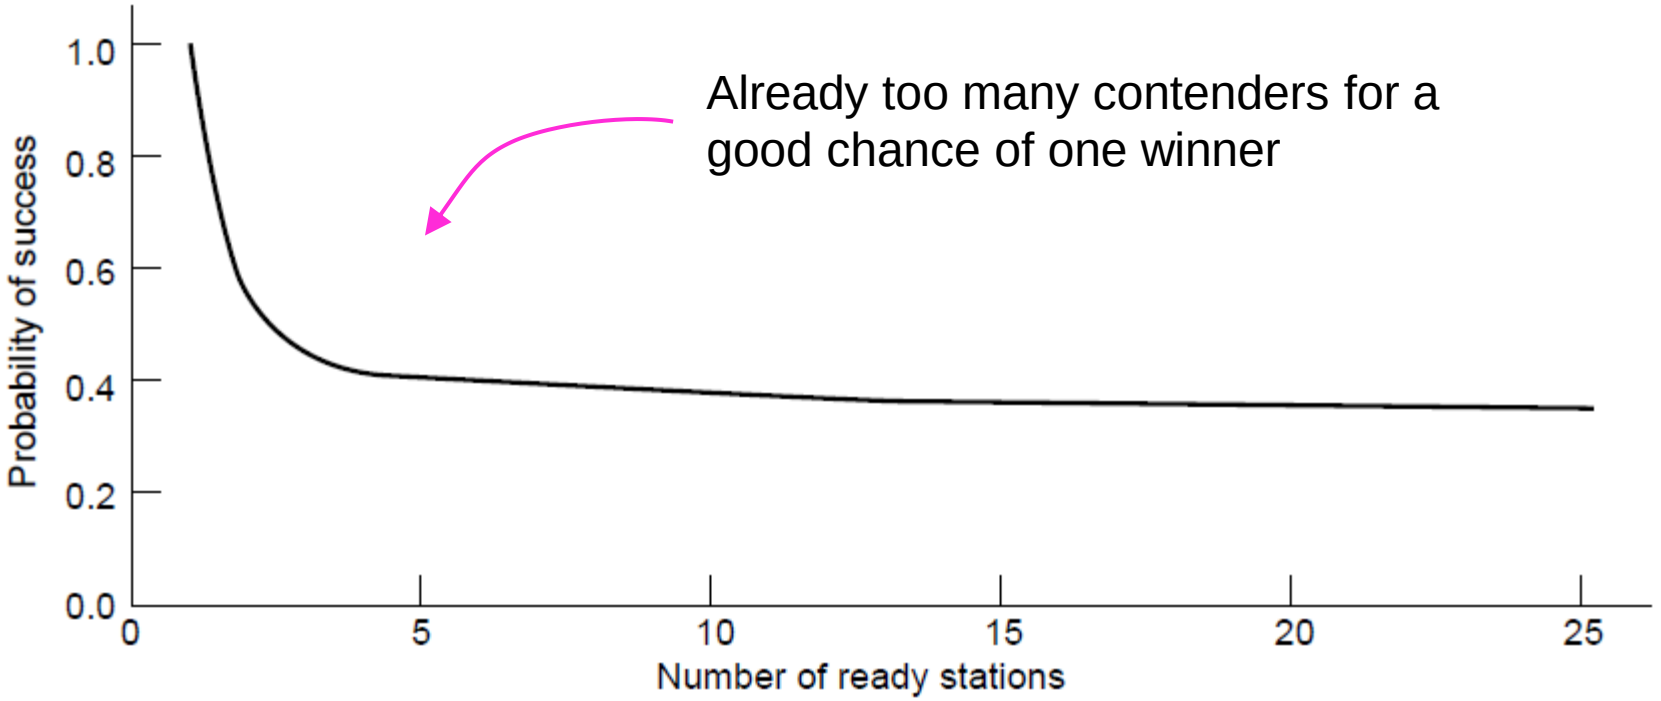

Limited Contention (2) –Adaptive Tree Walk

Tree divides stations into groups (nodes) to poll

- Depth first search under nodes with poll collisions – recursive function
  1. All stations under level 0 are permitted to try to acquire the channel. If only one does then good, but if collision then goto point 2
  2. Only nodes under 2 in level 1 can compete. If only one does then good, but if collision then goto point 3. If none transmit then those under 3 to transmit
  3. Only nodes under 4 in level 2 can compete. If only one does then good, but if collision then choose 1 to transmit. If none then choose 1 under 5 to transmit

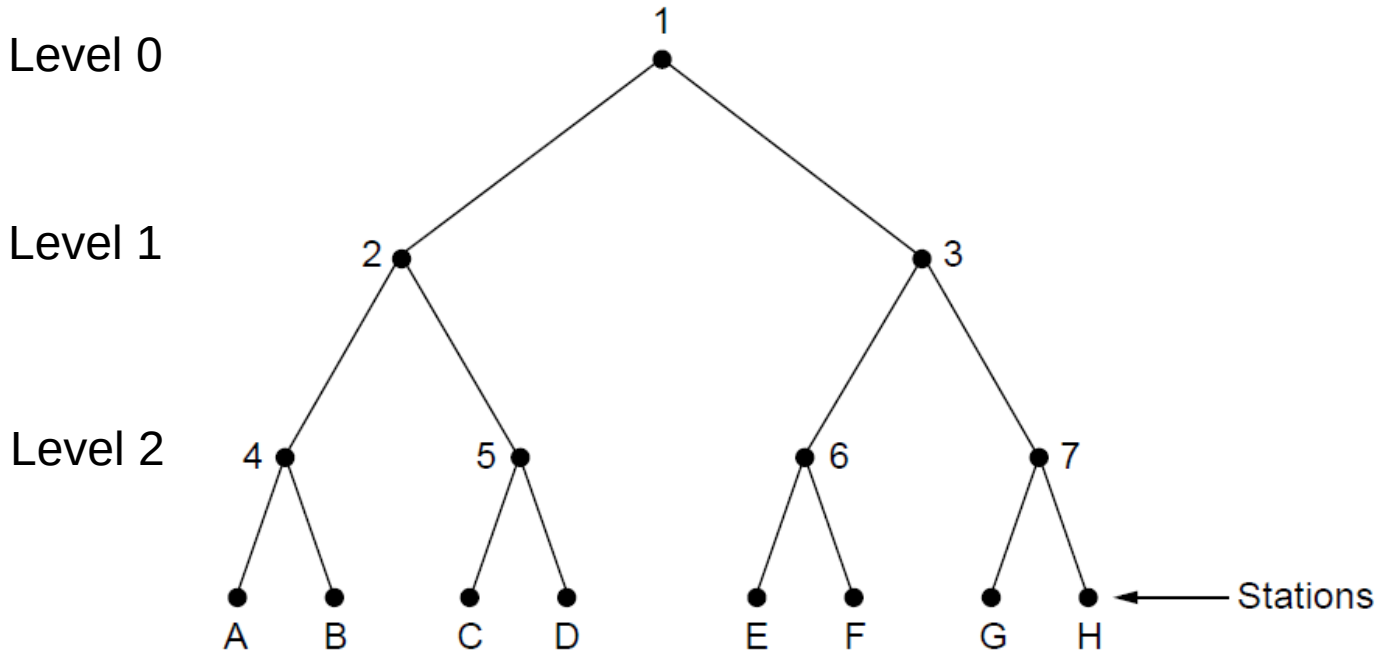

Wireless LAN Protocols (1)

Wireless has complications compared to wired.

Many WLANs use strategically placed Access Points (APs)

- APs dual homed with wired connections to each other and usually/often connect to larger networks
- Nodes can move in WLAN; APs in fixed locations
- WLANs can be configured to resemble cell tech except each cell has only one channel that is shared by all of the stations/APs

WLAN nodes may have different coverage regions

- Some nodes in same WLAN may be in range, others may not be
- Leads to hidden and exposed terminals – described on next slides

Nodes often can't detect collisions, i.e., cannot normally sense a collision while it is happening

- Makes collisions expensive and to be avoided
- ACKs are used to discover collisions after the fact

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

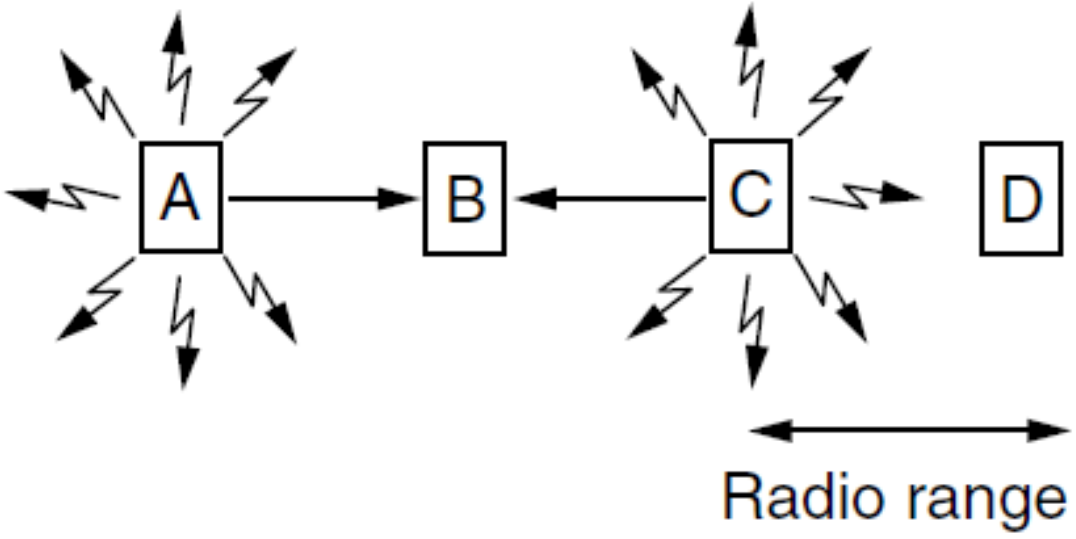

Wireless LANs (3) – Exposed terminals

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

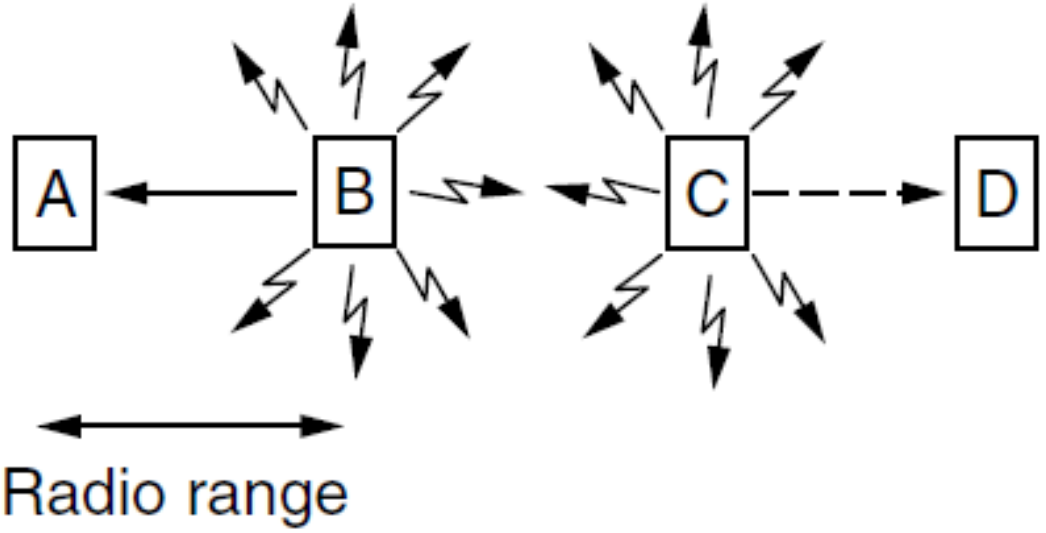

Foundational problem for both hidden and exposed terminals: The transmitter really wants to know whether there is radio activity around the receiver but CSMA merely tells it whether there is activity near the transmitter by sensing the carrier.

Wireless LANs (4) – MACA

Multiple Access with Collision Avoidance (MACA) protocol grants access for A to send to B. Sender stimulates receiver into outputting a short frame so stations can detect this transmission and avoid transmitting for the duration of the upcoming (large) data frame:

- A sends RTS to B [left]; B replies with CTS [right]
  - Request to Send (RTS) and Clear to Send (CTS) are short packets
- A can send with exposed but no hidden terminals

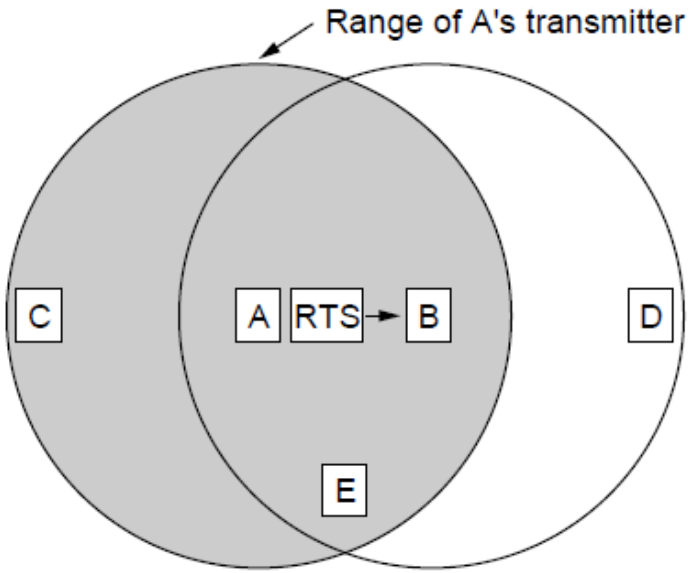

A sends RTS to B; C and E hear and defer for CTS

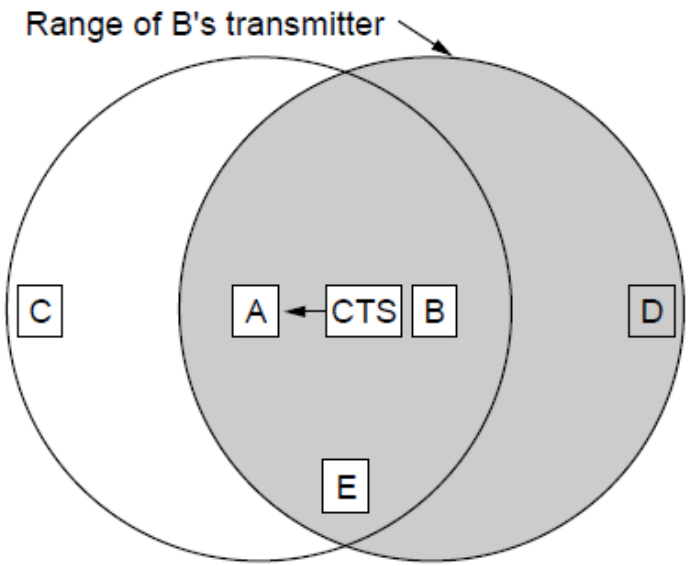

B replies with CTS; D and E hear and defer for data

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

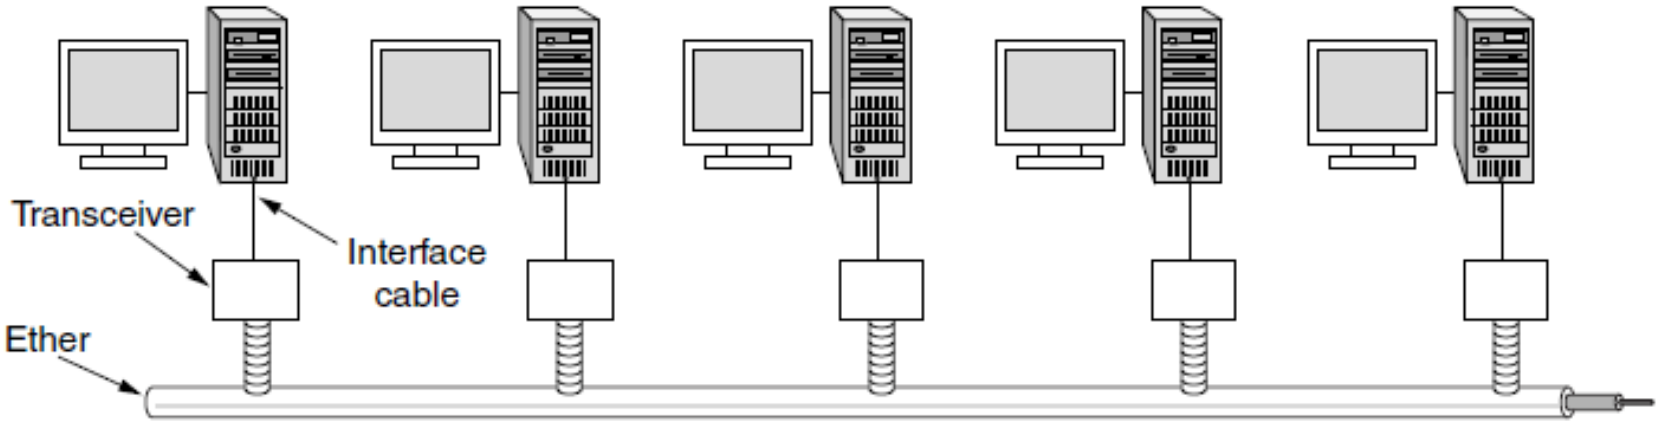

Ethernet V2

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.

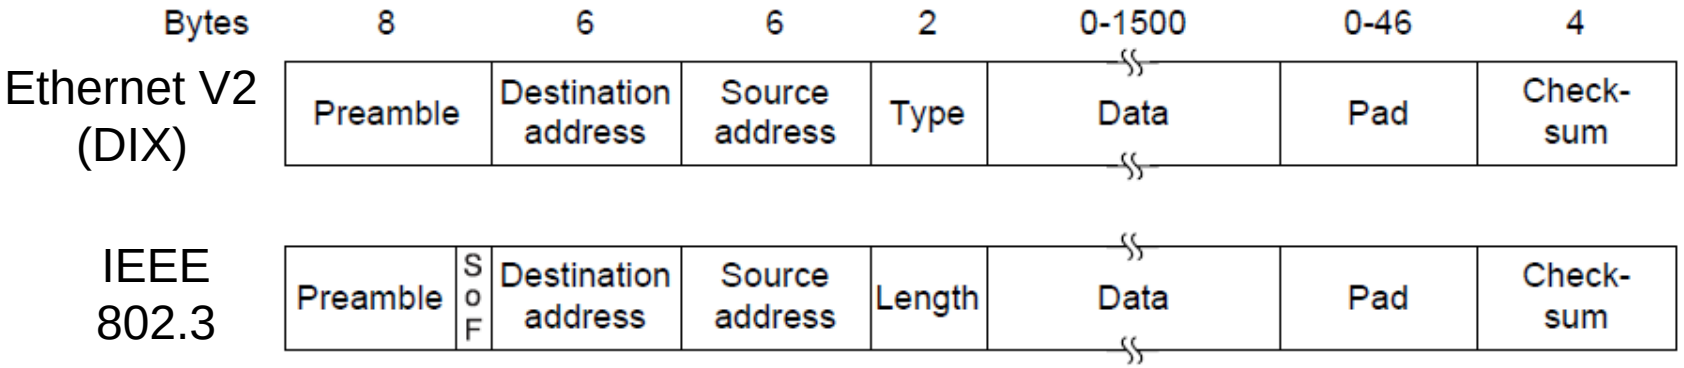

Question: Can Ethernet V2 directly communicate with IEEE 802.3?

Classic Ethernet (3) – 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

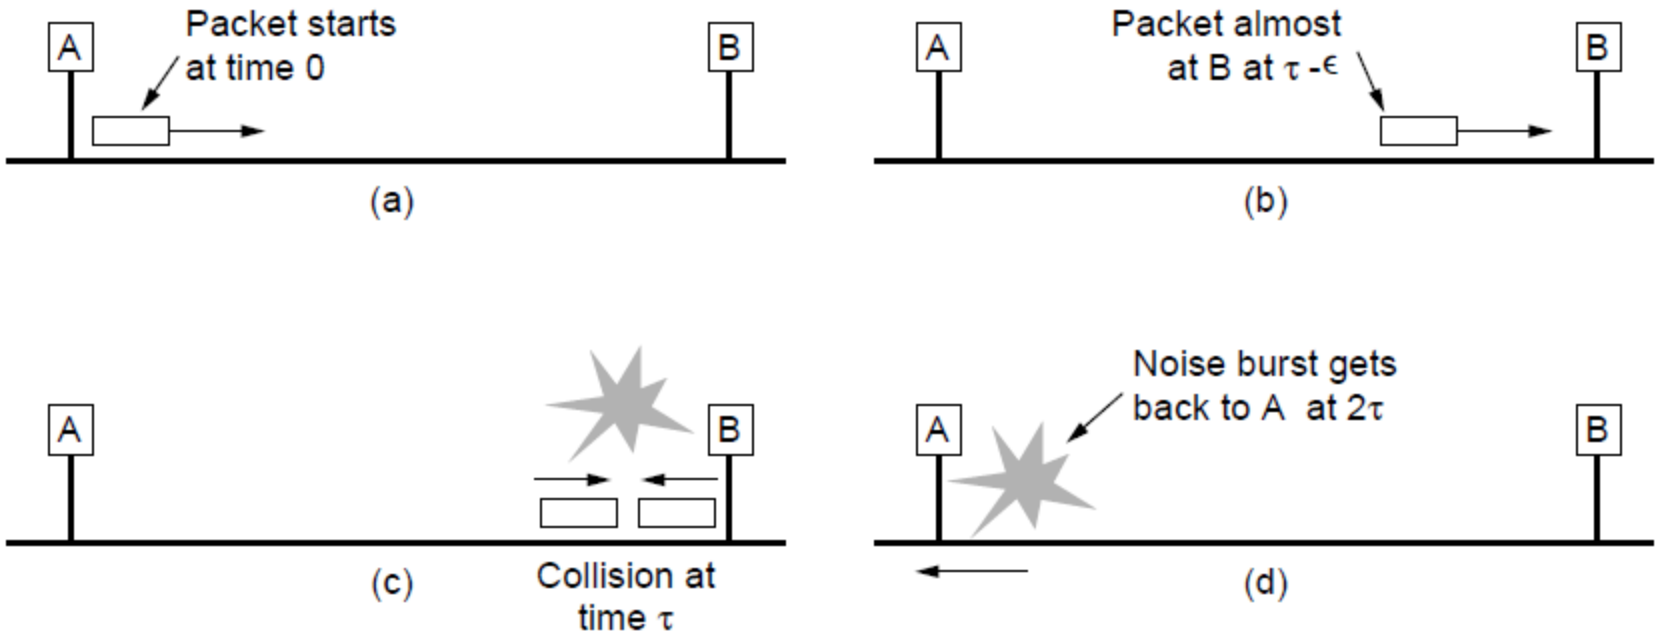

Classic Ethernet (4) – Performance

Efficient for large frames, even with many senders

- Degrades for small frames (and long LANs)

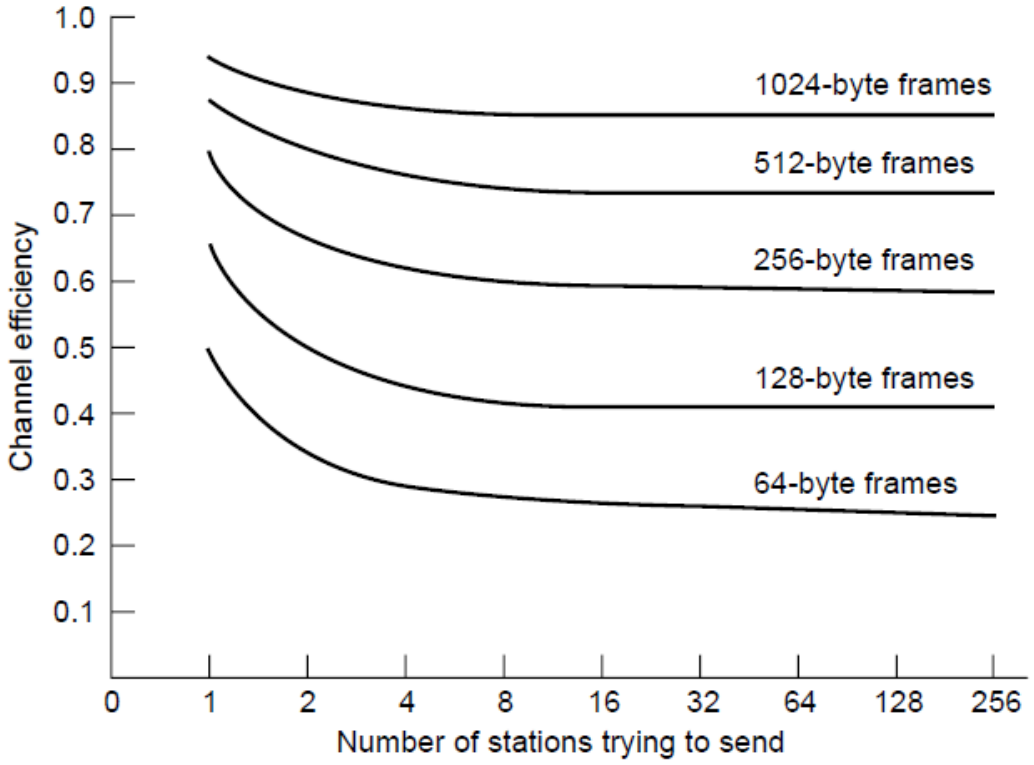

10 Mbps Ethernet,
64 byte min. frame

Switched/Fast Ethernet (1)

IEEE 802.3

- 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

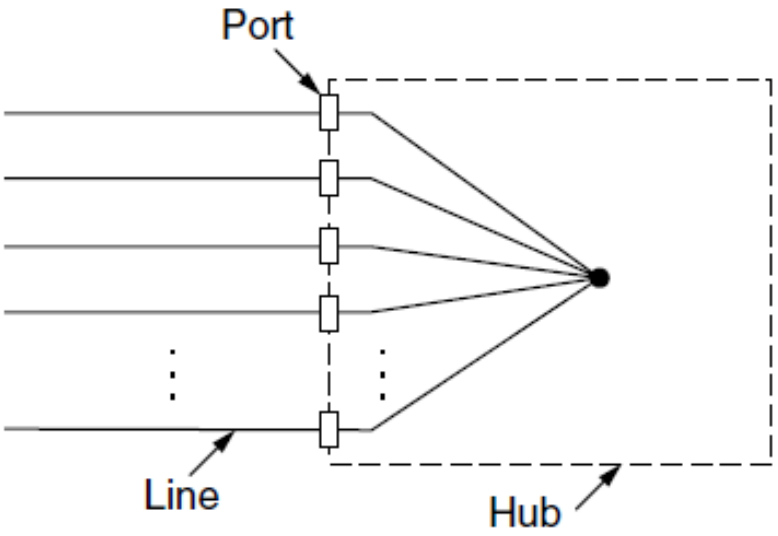

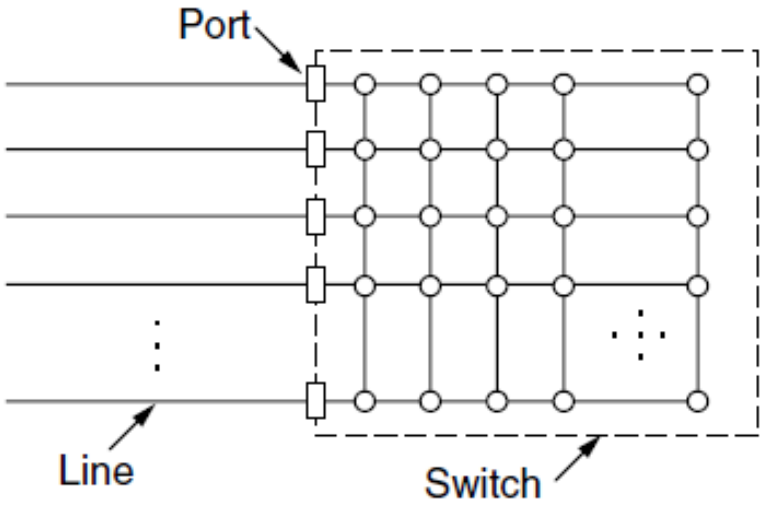

Switched/Fast Ethernet (2)

Switches can be wired to computers, hubs and switches

- Hubs concentrate traffic from computers
- More on how to switch frames the in Section 4.8

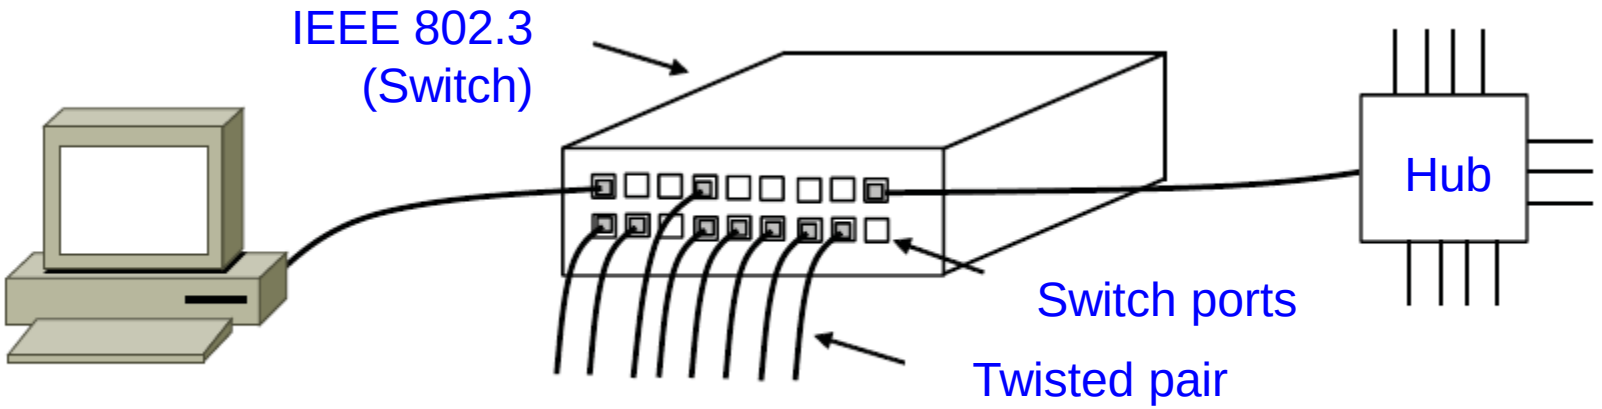

- **Hub** – Input lines are joined electrically. Frames arriving on any of the lines are sent out on all of the others. If 2 frames arrive at same time, they will collide.
- **Bridge** – Like a hub, a bridge also has multiple ports. Unlike a hub, each port is isolated to be its own collision domain. When a frame arrives on a port, the bridge extracts the destination address from the frame header and looks it up in a table to see which port to send the frame out on.

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

Note Well

Pay close attention to Section 4.3.8 in the textbook (pages 298-299) which explains why Ethernet was a major deployment success while many alternative LAN technologies failed.

In a previous reading you learned that Metcalfe's Law helps predict product success. This section explains other factors that also influence product success.

Gigabit / 10 Gigabit Ethernet (1)

Switched Gigabit Ethernet is now the garden variety

- With full-duplex lines between computers/switches

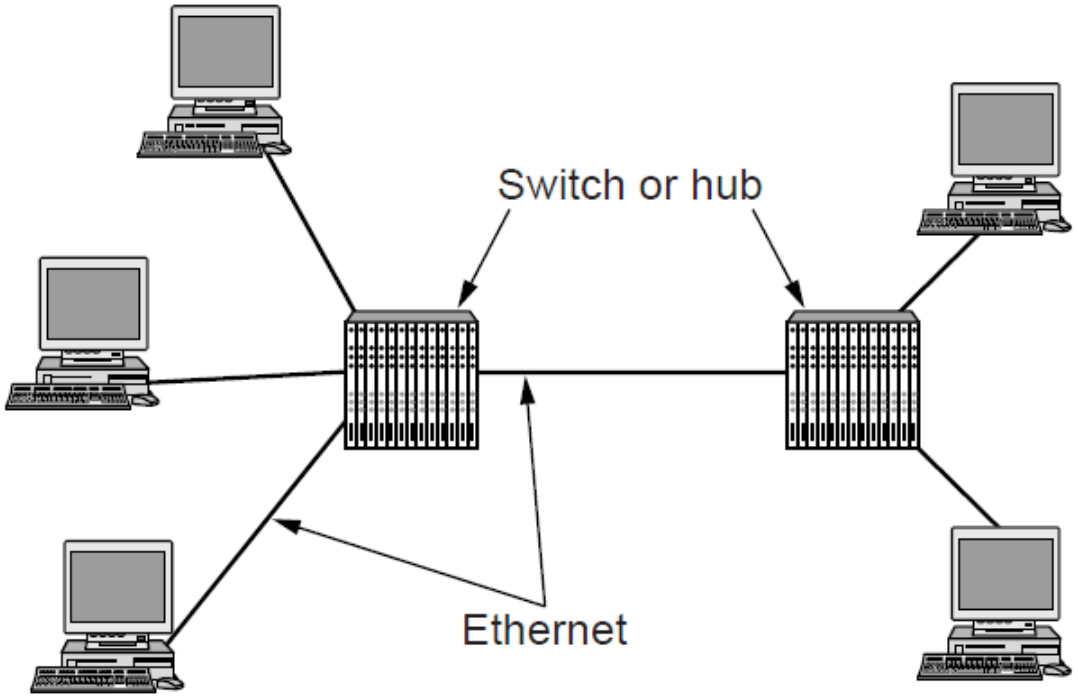

Gigabit / 10 Gigabit Ethernet (1)

- 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, but that is rare.

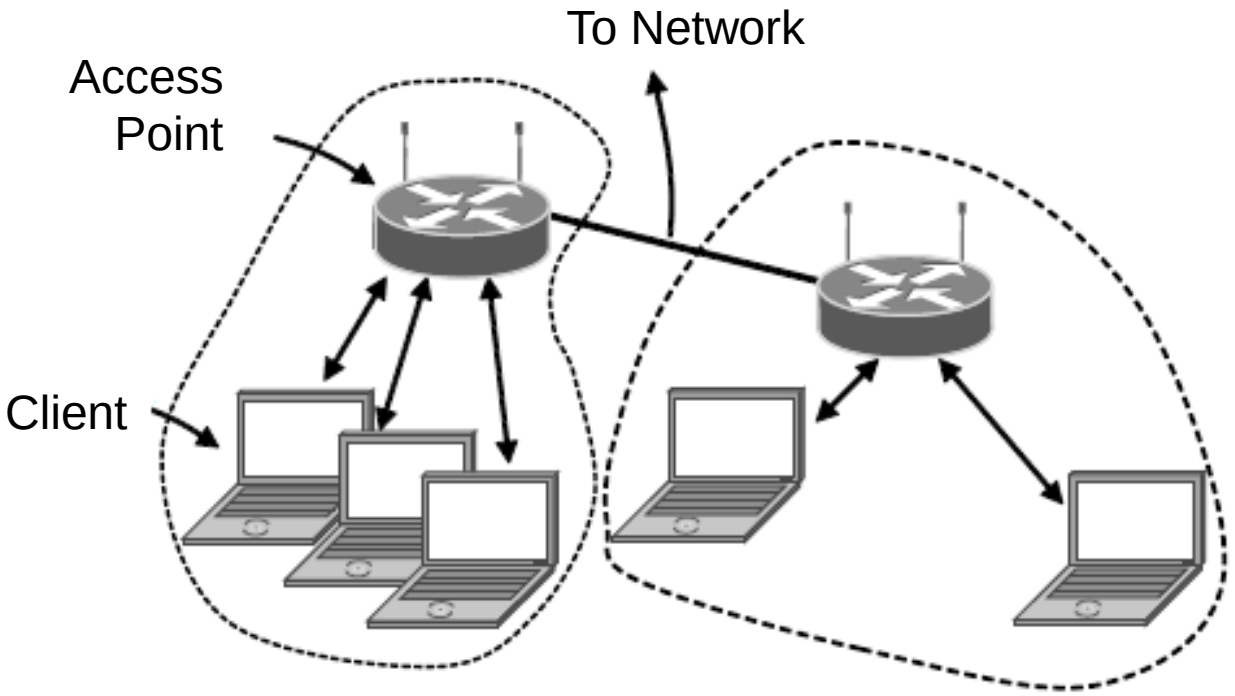

802.11 Architecture/Protocol Stack (2)

MAC is used across different physical layers

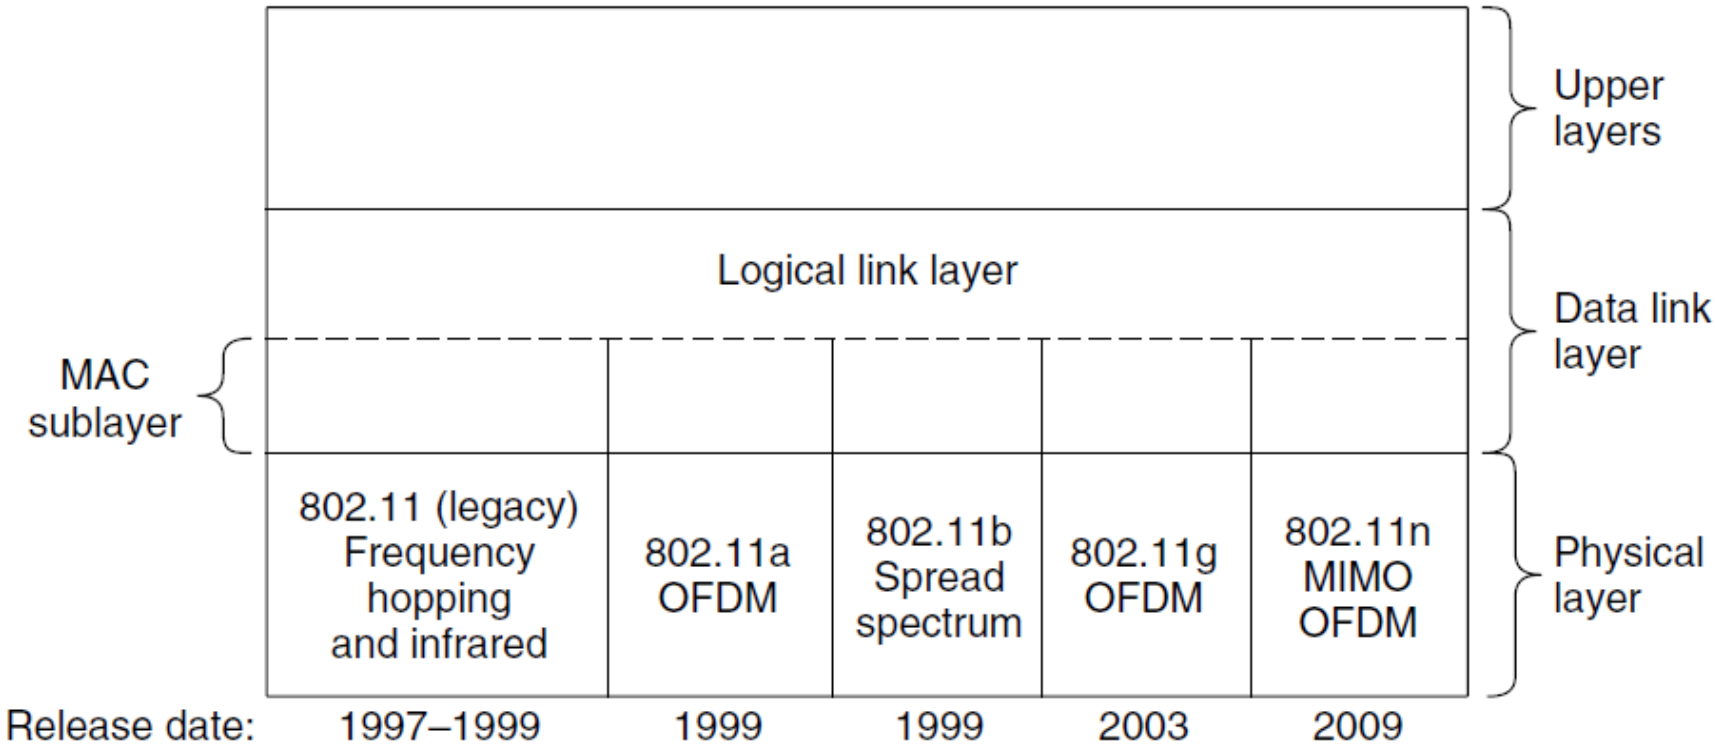

MIMO = Multiple Input Multiple Output (see page 303)

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

802.11 MAC (1)

- CSMA/CA (CSMA with Collision Avoidance) inserts backoff slots to avoid collisions
  - CSMA = Channel sensing before sending and exponential backoff after collisions
  - CA = 1) starts random backoffs early to avoid collisions and 2) ACKs are used to infer collisions (see Slide 20: RTS/CTS to addr hidden and exposed terminals)
- MAC uses ACKs/retransmissions for wireless errors

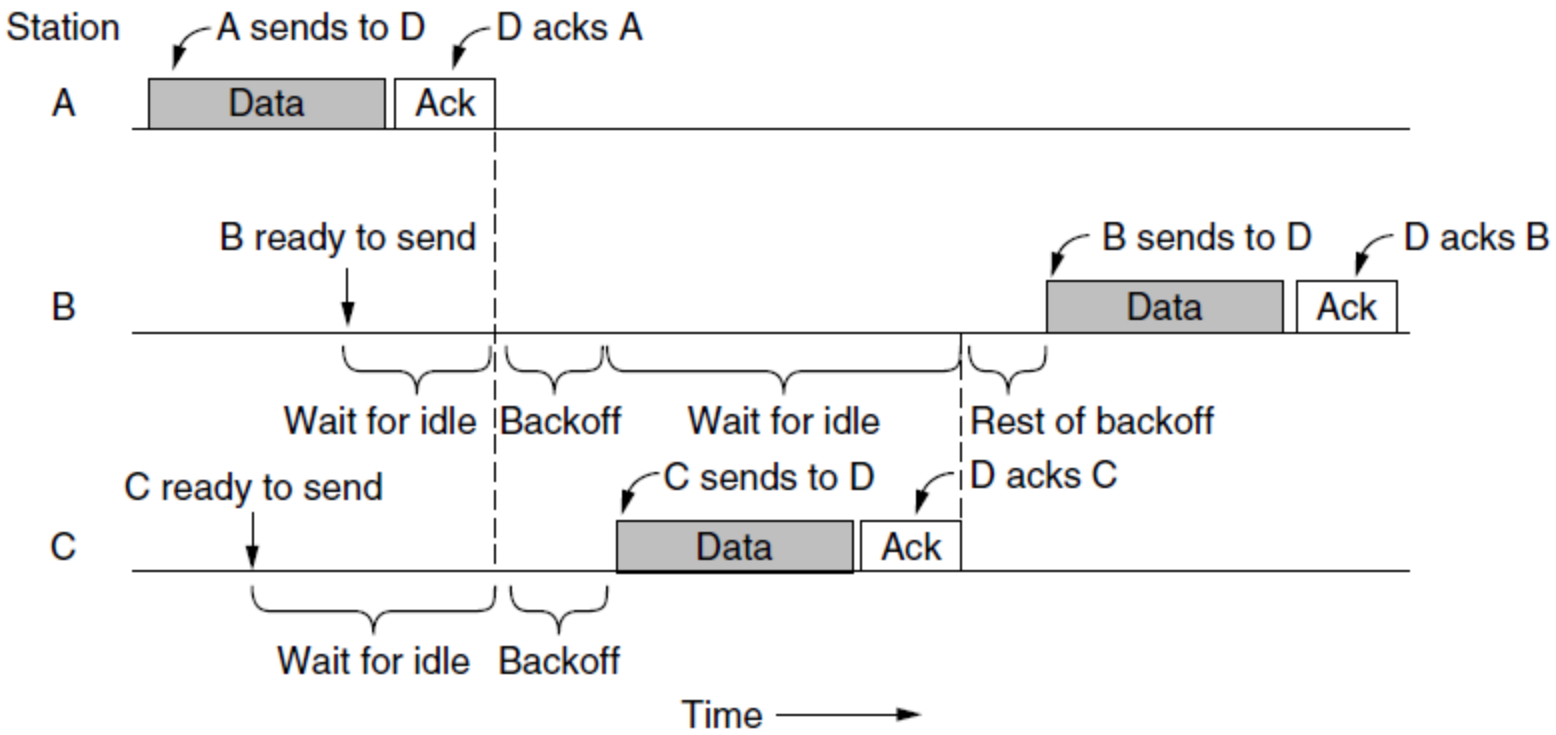

802.11 MAC (2)

To reduce ambiguities (e.g., diff stations have diff transmission ranges) about which station is sending, 802.11 defines channel sensing to consist of both physical sensing and virtual sensing

- Physical – checks medium to see if a valid signal
- Virtual – each station keeps a logical record of when the channel is in use by tracing the Network allocation Vector (NAV) in frame header which indicates expected use duration

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

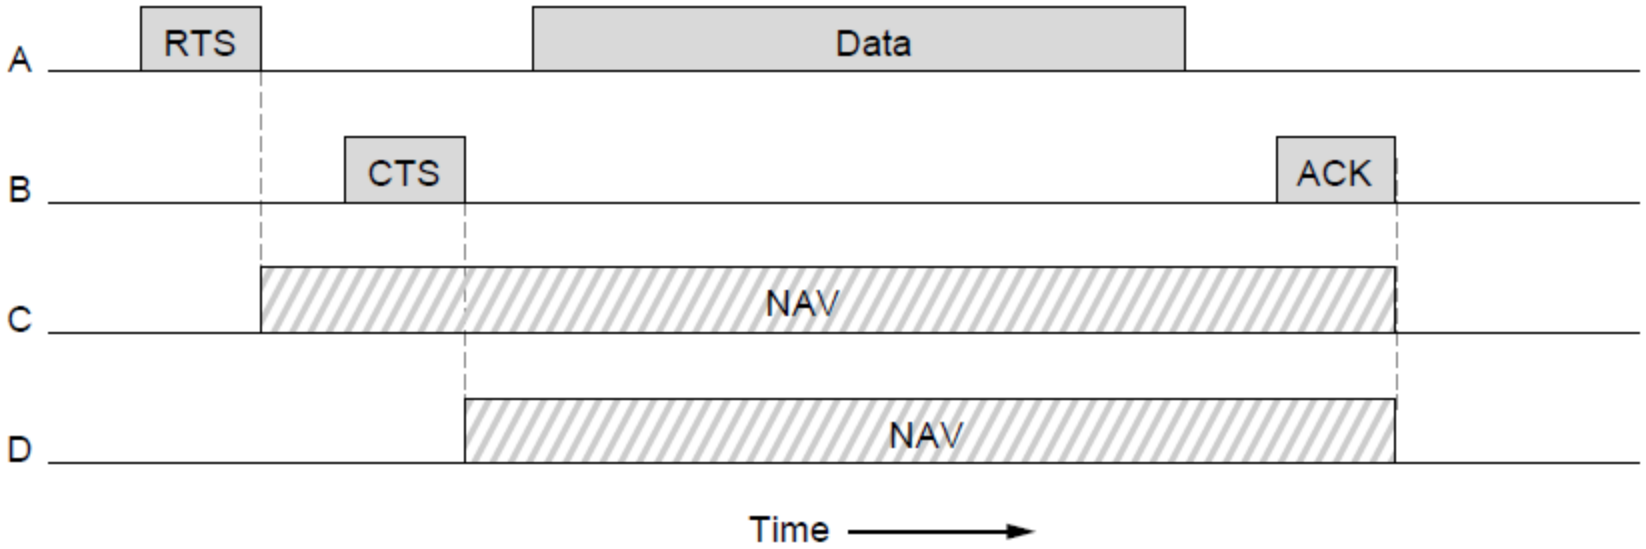

A wants to send to B, C in range of B, D in range of B but not A

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

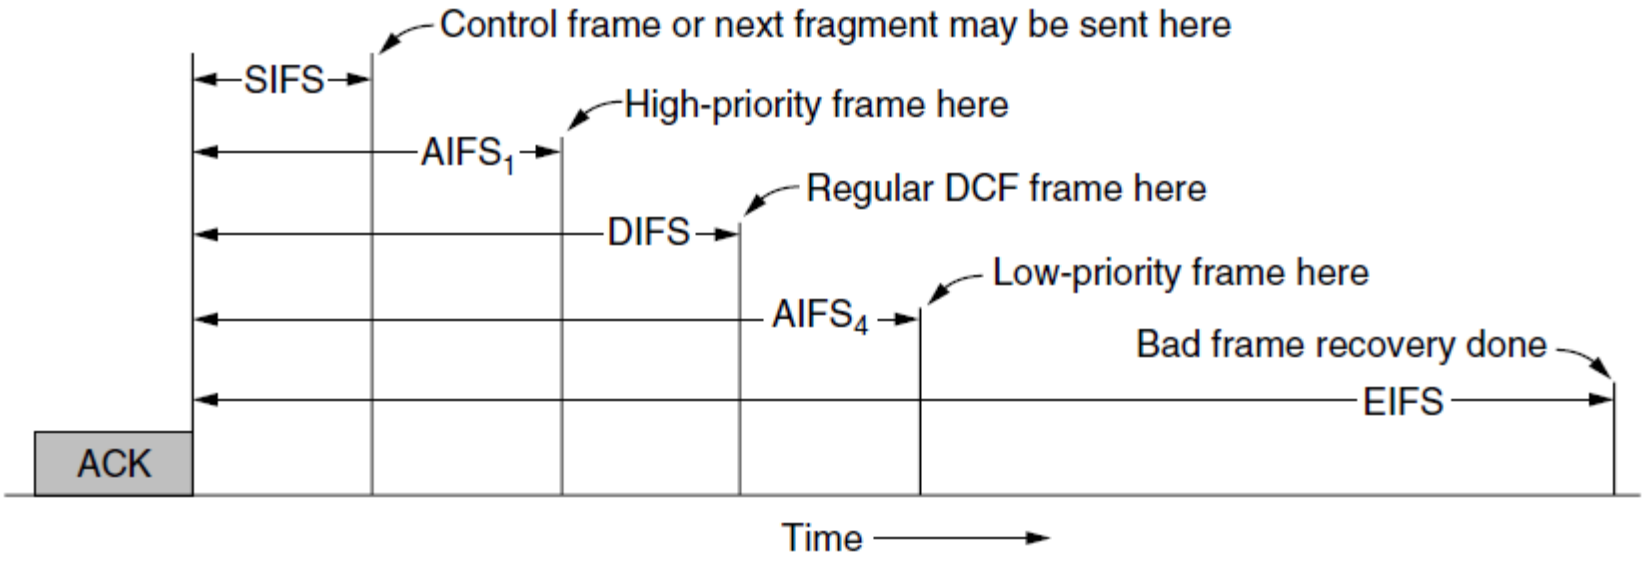

SIFS = Short InterFrame Spacing
AIFS = Arbitration InterFrame Space
DIFS = DCF (distributed coordination function) InterFrame Space
EIFS = Extended InterFrame Spacing

802.11 Frames

- Frames vary depending on their type (Frame control)
- Data frames have 3 addresses to pass via APs

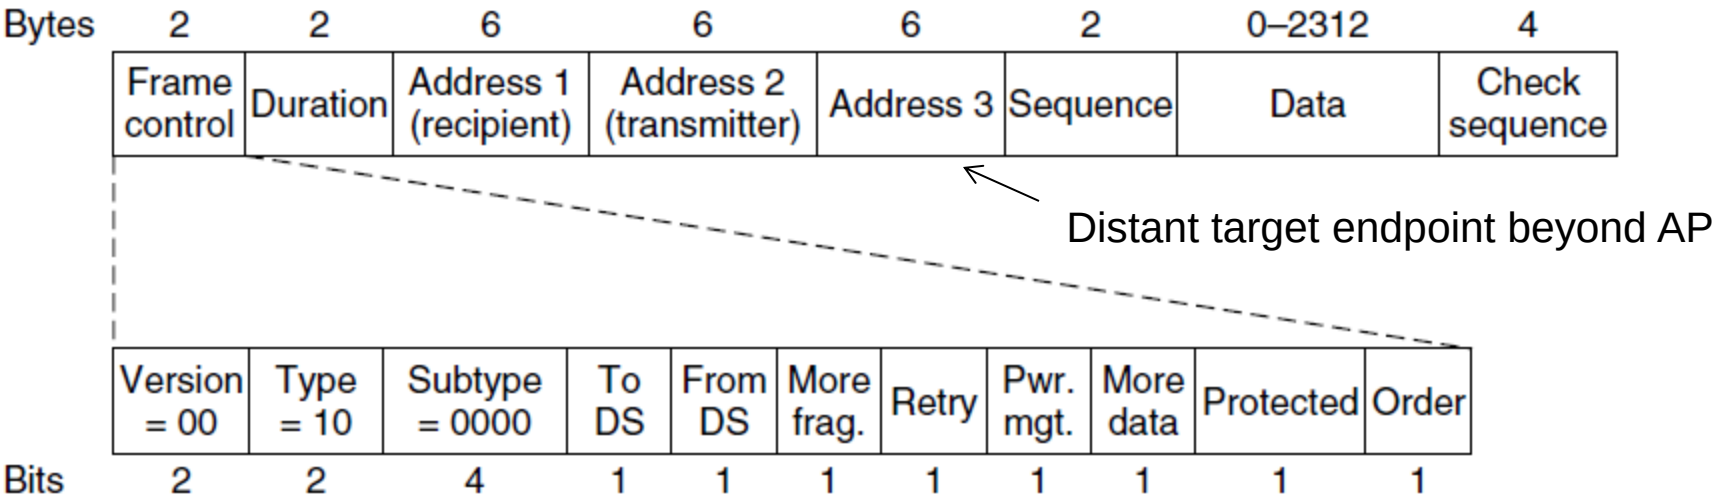

Addresses are standard IEEE 802 six octet MAC addresses
See Pages 309-310 in textbook for a description of the header fields.

CS 450

CS 450 is skipping the material in sections 4.5 through 4.7 (Pages 313-331)

While this material is important, it is of lower priority than the material the class is covering.

Should you need to know this material at work, it will be available to you should you keep the textbook.

Data Link Layer Switching

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

Uses of Bridges

Why Bridge use – 1) link geographically spread orgs, 2) Load balancing (split LANs to spread load), 3) scaling situations where more devices than ports, 4) bridges act like firewalls to protect a single berserk node from bringing down the entire system

Common setup is a building with centralized wiring

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

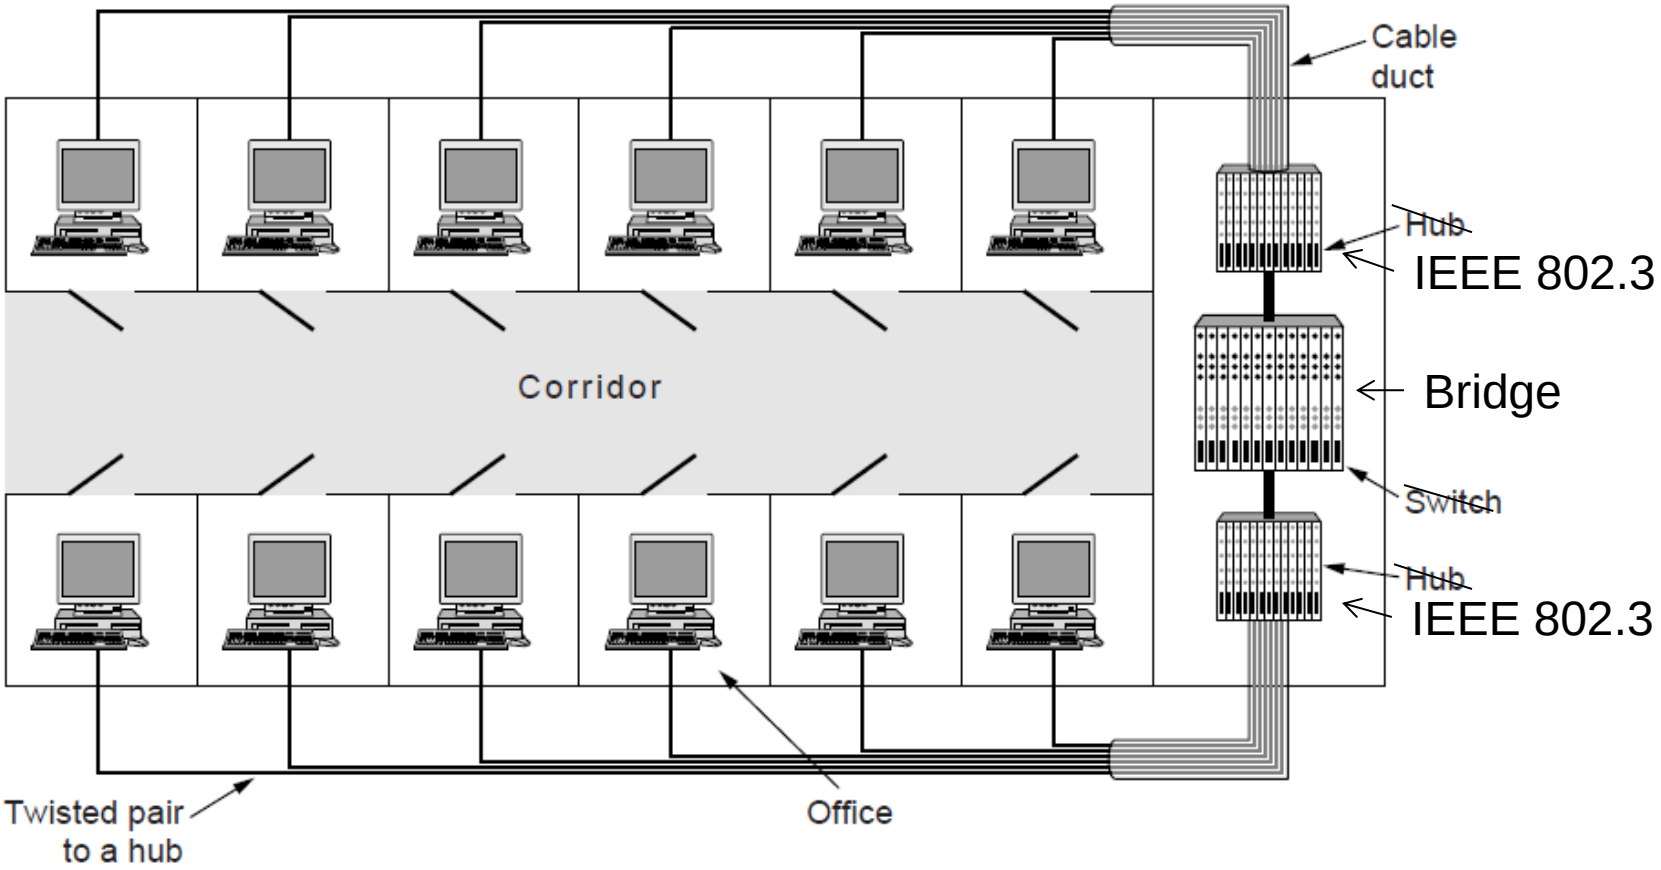

Learning Bridges (1)

Transparent bridges are plug-and-play ready. They use a backward learning algorithm to stop traffic being sent where it is not needed and a spanning tree algorithm to break inadvertent routing loops (e.g., hardware misconfigurations)

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

- Computers, bridges, and hubs connect to its ports

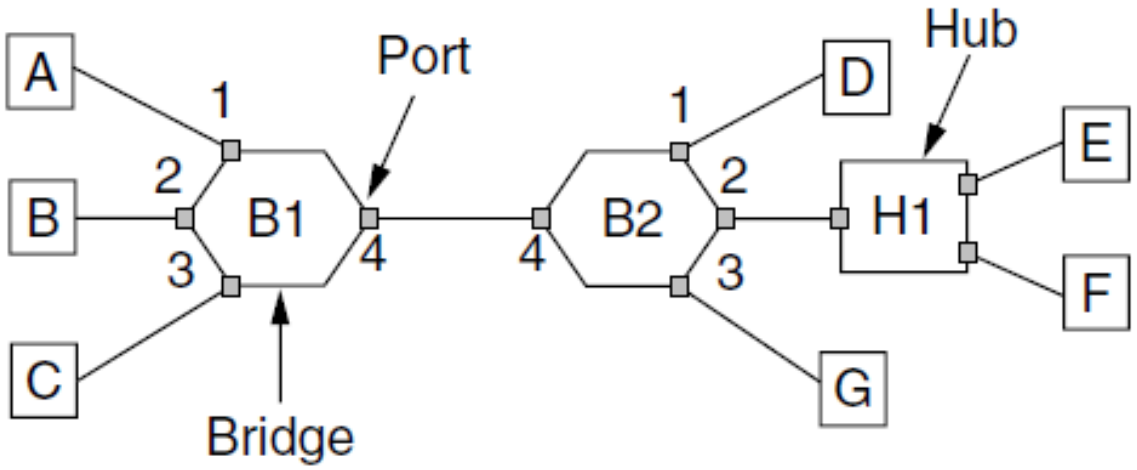

Learning Bridges (2)

When bridges are first plugged in, all of its hash table entries are empty. The bridge doesn't know where the destinations are, so it broadcasts frames to all ports except the one it came in on to ensure frame delivery. Once bridge learns which port is used for which destination, it solely forwards to that port.

Generic Hash Table (source: Wikipedia):

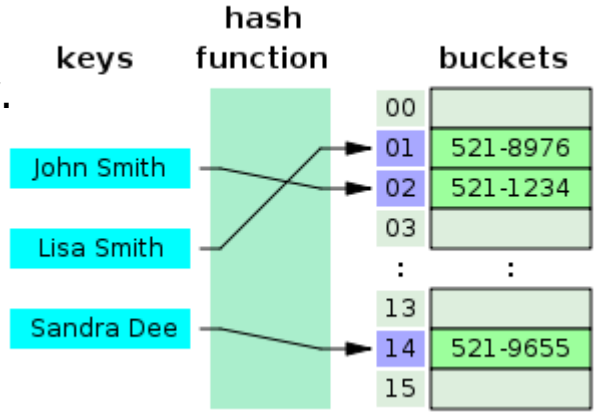

Backward learning algorithm picks the output port:

- 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 or any other higher layer information

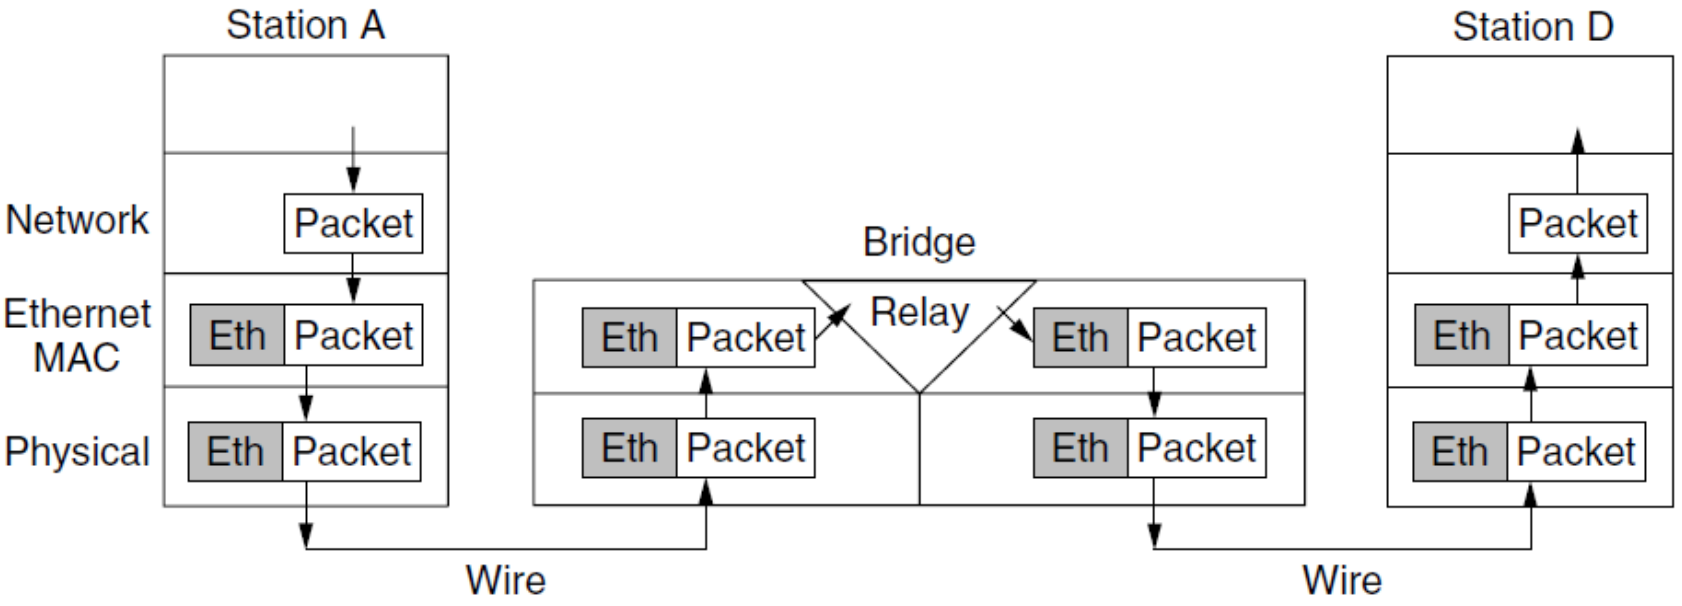

Spanning Tree (1) – Problem

To increase reliability, redundant links can be used between bridges but this may cause problems during initial backward learning (i.e., frame sent back to other bridge via redundant link). It is also possible for poor port configuration choices to introduce forwarding loops in the topology.

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

- Need spanning tree support to solve problem

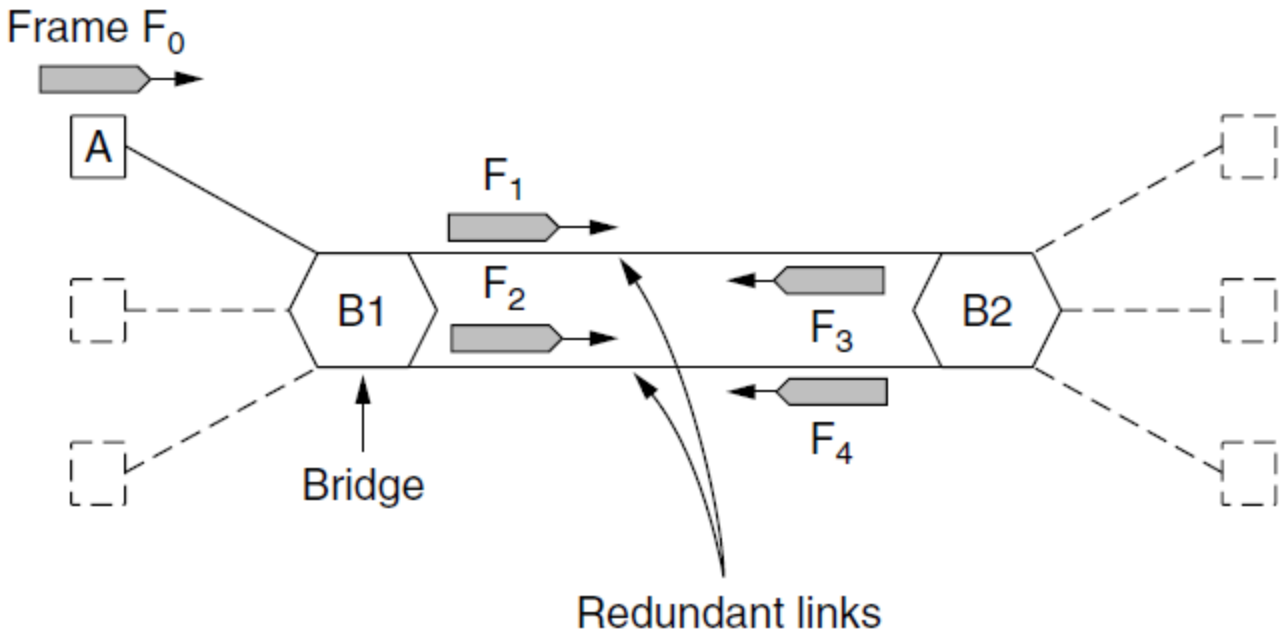

Spanning Tree (2) – Algorithm

Spanning tree – bridges communicate with each other and overlay the actual topology with a spanning tree that reaches every bridge. Bridges then turn off ports that are not part of the shortest path from the root to that location.

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

After the algorithm runs:

- B1 is the root, two dashed links are turned off
  - » Root is the bridge with the lowest IEEE 802 MAC address
- B4 uses link to B2 (lower than B3 also at distance 1)
- B5 uses B3 (distance 1 versus B4 at distance 2)

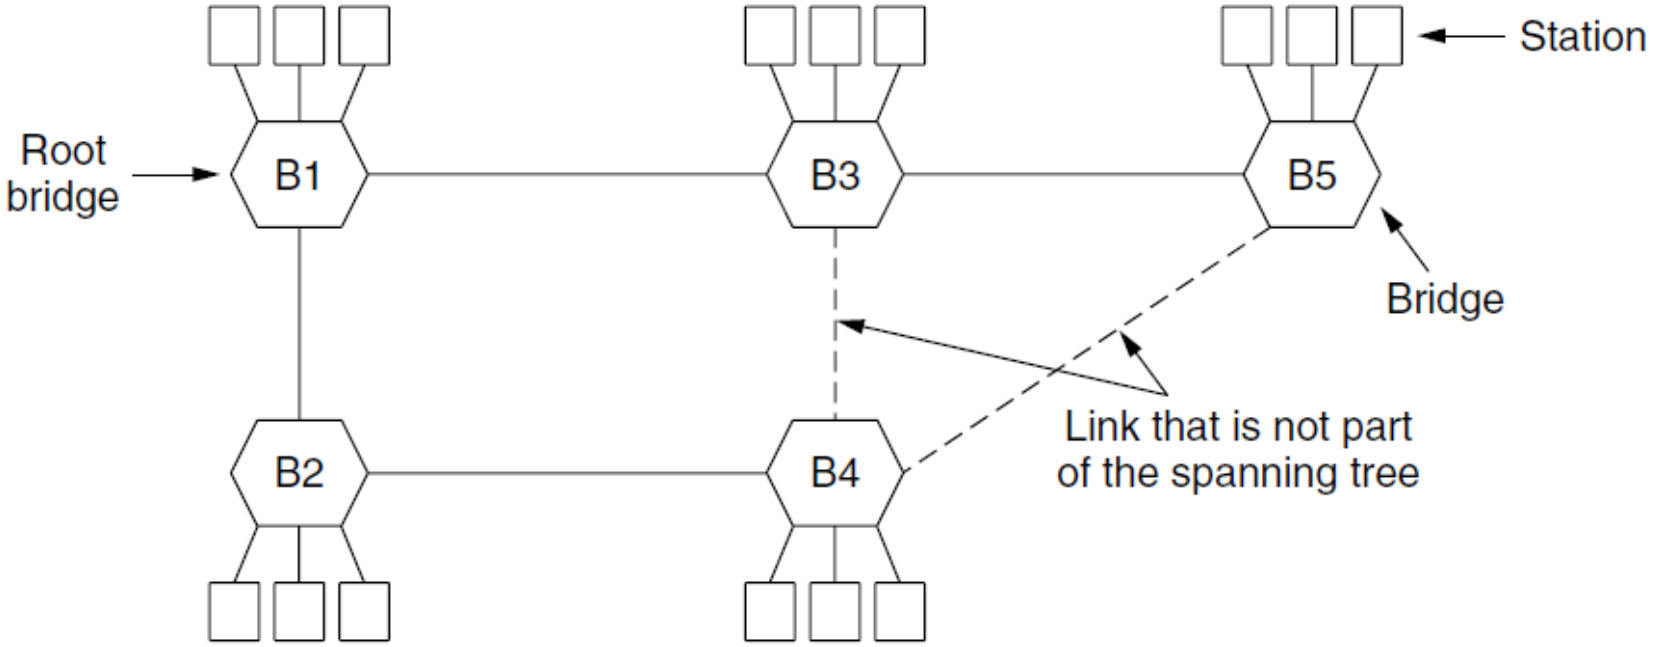

Repeaters, Hubs, Bridges, Switches, Routers, & Gateways

Devices are named according to the layer they process

- A bridge or LAN switch operates in the Link layer

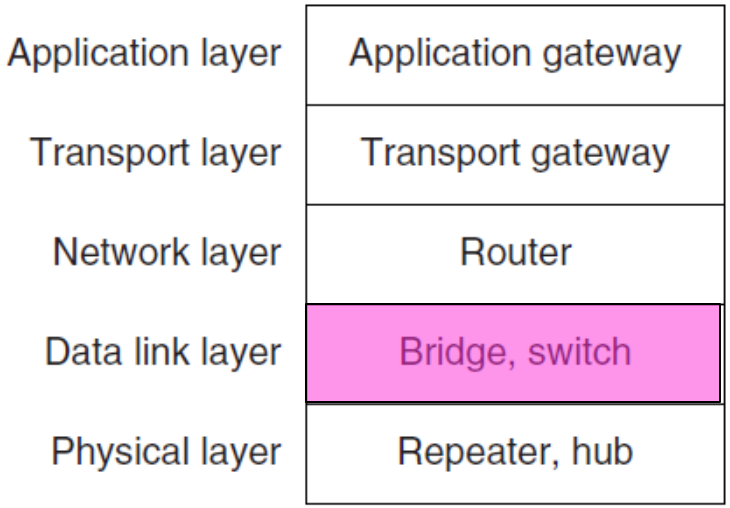

Virtual LANs (1)

Why Virtual LANs – 1) Admins may want the LANs to reflect the Org structure perhaps for security reasons (e.g., not let servers be accessible from outside that LAN), 2) load distribution to offload LANs with heavy traffic, 3) reduce impact of broadcast traffic (e.g., when destinations are not known)

VLANs (Virtual LANs) splits one physical LAN into multiple logical LANs to ease management tasks

- Ports are “colored” according to their VLAN
  - Configuration tables tell which VLANs are accessible by which ports

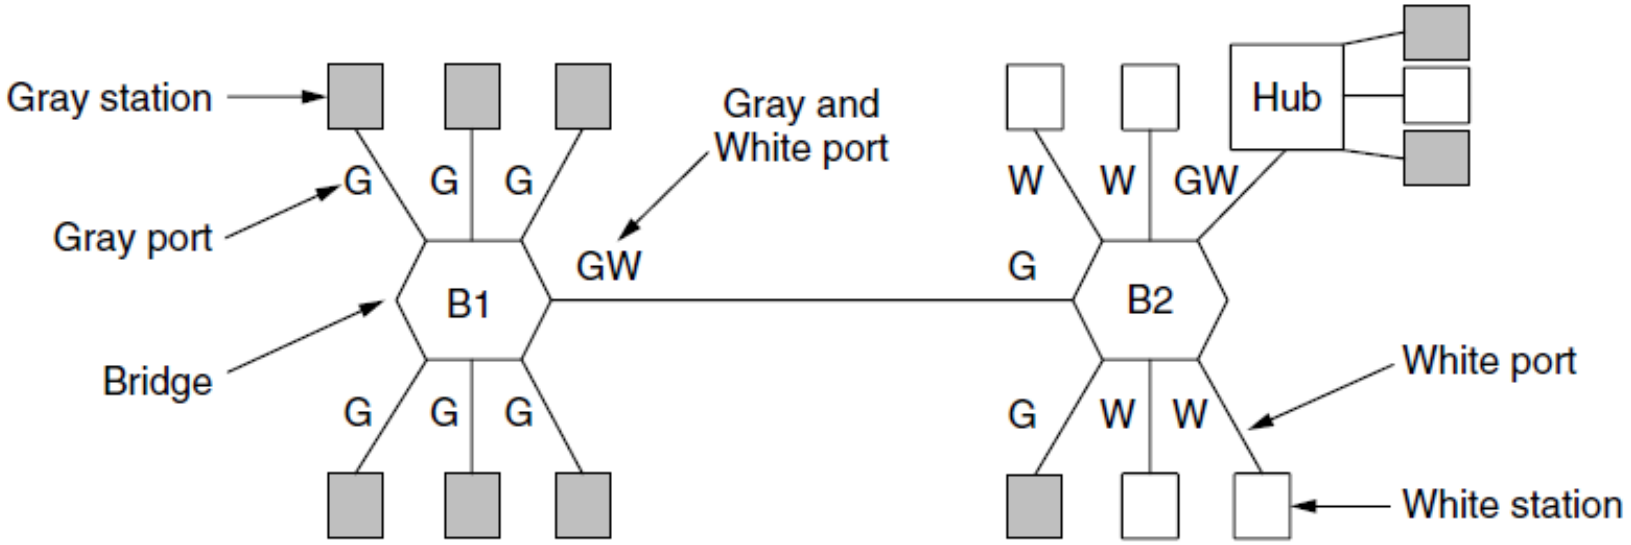

Virtual LANs (2) – IEEE 802.1Q

Bridges need to be aware of which VLAN an incoming frame belongs.

- In 1998 IEEE 802 changed the Ethernet header to add in a VLAN tag (802.1Q), which is only used by bridges and switches and not end user machines.

- In 802.1Q, frames are tagged with their "color"
  - All machines on a port must belong to the same VLAN
- Legacy switches with no tags are supported

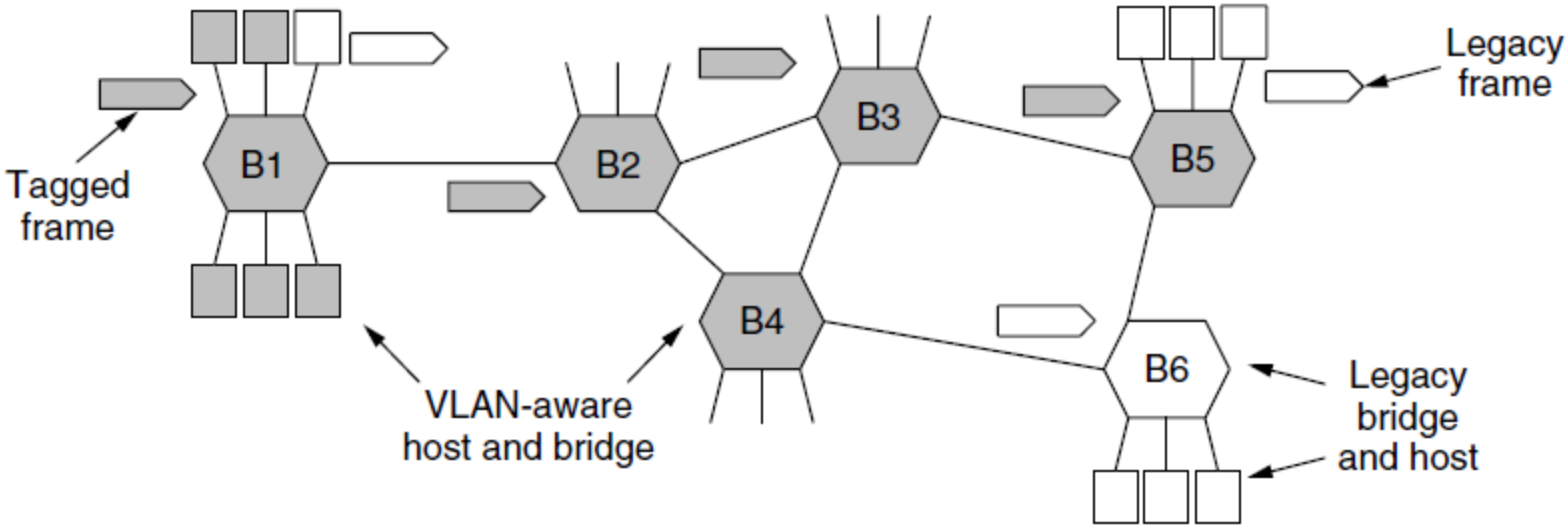

Virtual LANs (3) – IEEE 802.1Q

802.1Q frames carry a color tag (VLAN identifier)

- Length/Type value is 0x8100 for VLAN protocol

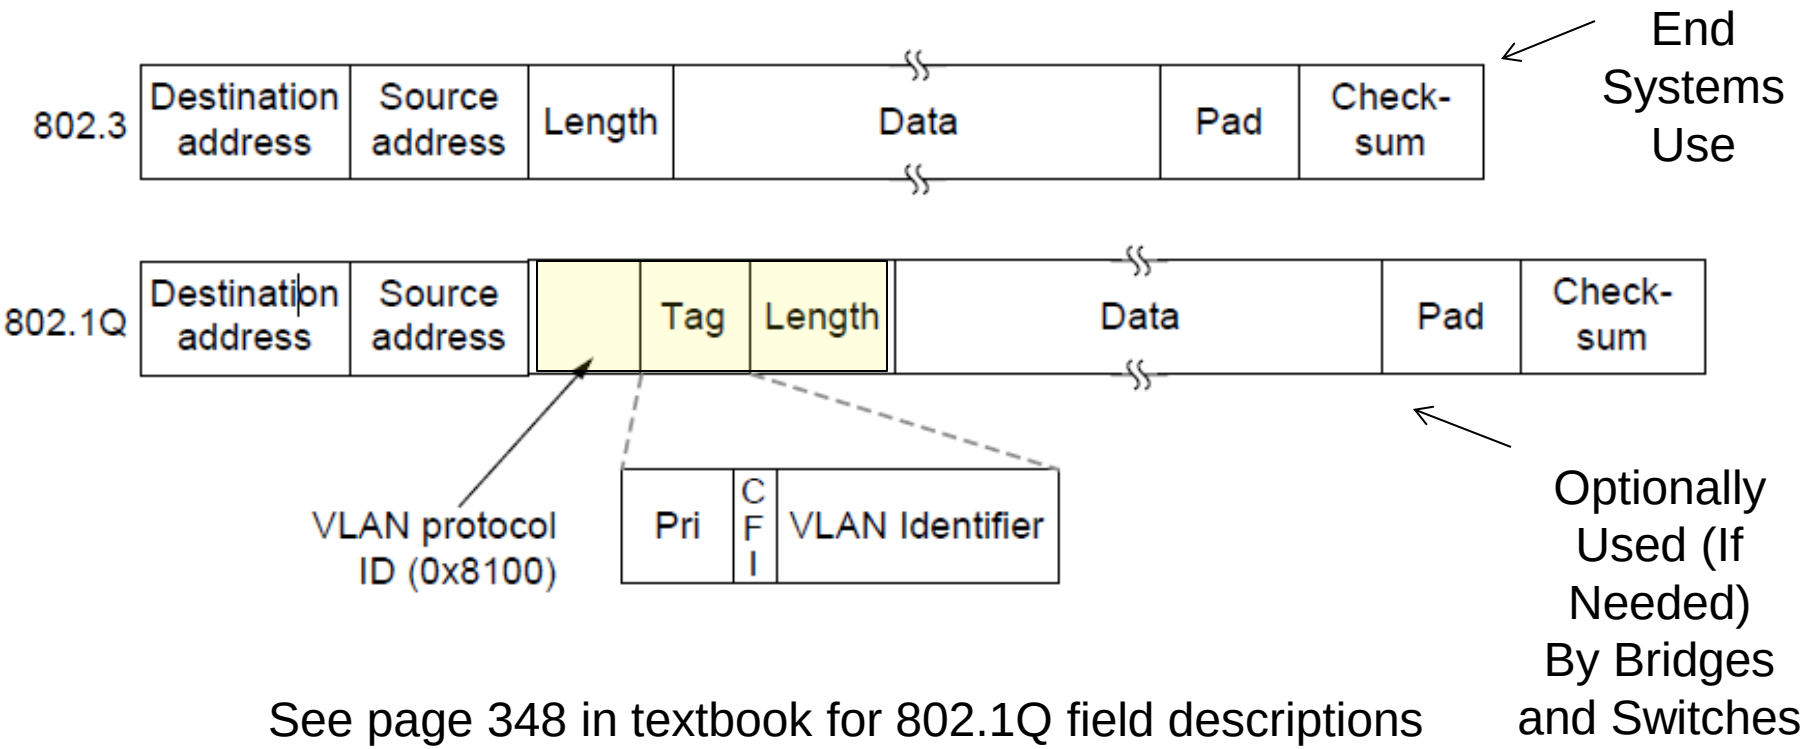

See page 348 in textbook for 802.1Q field descriptions

End

Chapter 4