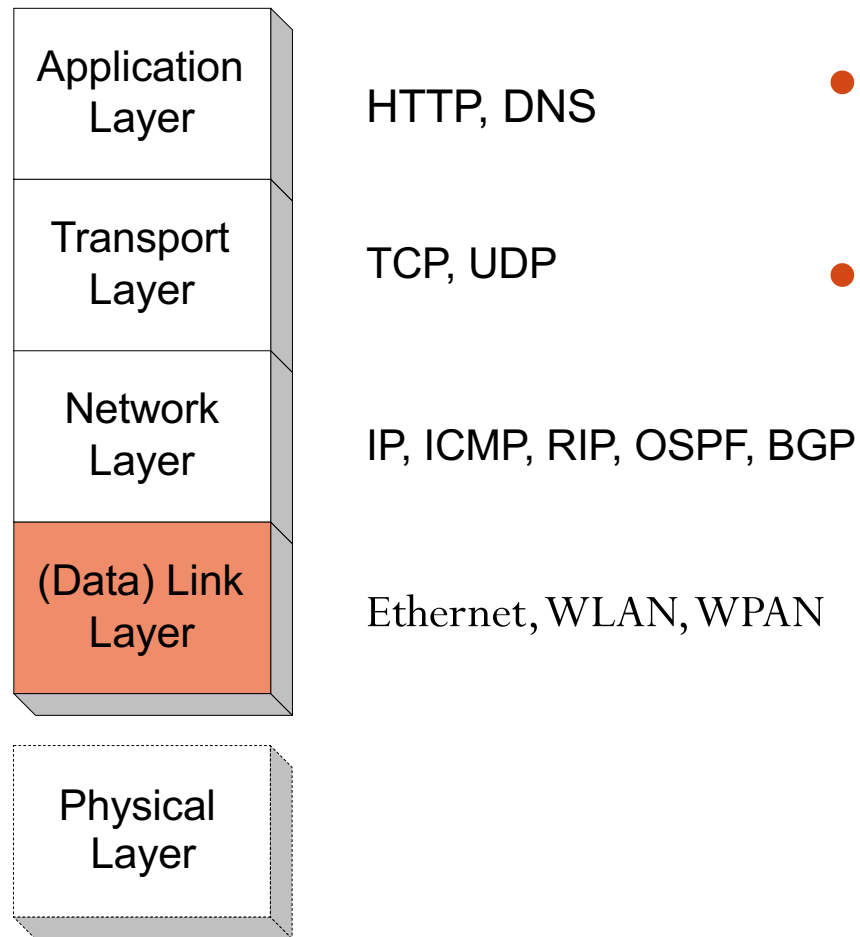


Data Link Layer

Data Link Layer and LANs



- Data link layer services
- Ethernet and Ethernet bridging
- Wireless WLANs

High Diversity

Protocol	Topology	# of Nodes	Segment Length(m)	Bandwidth (bps)	Per node Bitrate (bps)
Modbus	BUS Master/Slave	247	400...1200	9.6K...500K	27K
PROFIBUS	BUS Master/Slave	31 per Segment	100...1200	9.6K...12M	750K
CAN Bus	BUS CSMA	100	40...1000	50K...1M	60K
Ethernet	STAR	N/A	100	100M	6M
USB hub network	BUS Master/Slave	127	5 (between hubs)	2M	100K




Fast Evolution

Wi-Fi Generations

Generation	IEEE Standard	Adopted	Maximum Linkrate (Mbit/s)	Radio Frequency (GHz)
Wi-Fi 7	802.11be	(2024)	1376 to 46120	2.4/5/6
Wi-Fi 6E	802.11ax	2020	574 to 9608 ^[41]	6 ^[42]
Wi-Fi 6		2019		2.4/5
Wi-Fi 5	802.11ac	2014	433 to 6933	5 ^[43]
Wi-Fi 4	802.11n	2008	72 to 600	2.4/5
(Wi-Fi 3)*	802.11g	2003	6 to 54	2.4
(Wi-Fi 2)*	802.11a	1999	6 to 54	5
(Wi-Fi 1)*	802.11b	1999	1 to 11	2.4
(Wi-Fi 0)*	802.11	1997	1 to 2	2.4
*(Wi-Fi 0, 1, 2, 3, are unbranded common usage) ^{[44][45]}				

Link Layer: Introduction

- communication channels that connect adjacent nodes along communication path are **links**
 - wired links
 - wireless links
- layer-2 packet is a **frame**, encapsulates datagram



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

Network interface card, driver

- Network interface card
 - Firmware: program stored on the network card's ROM (BIOS) and configuration information stored in E2PROM.
 - Hardware: ICs, connectors
- Drivers: software interface between the network card hardware/firmware and the operating system

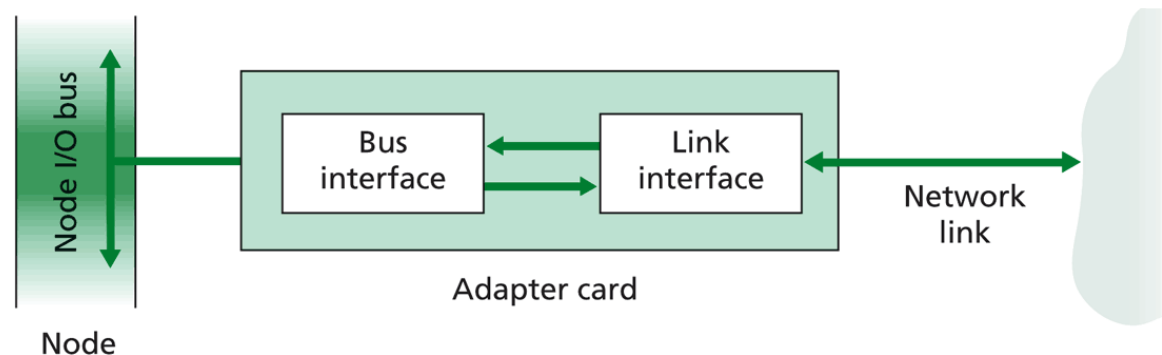



Figure 5.3 ♦ The adapter is a semi-autonomous unit.

Link Layer Services

- **Flow Control:**
 - pacing between adjacent sending and receiving nodes
- **Error Detection:**
 - errors caused by signal attenuation, noise, interference
 - 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



Which other layer provides similar services?

Link Layer Services

- Half-duplex and full-duplex
 - with half duplex, nodes at both ends of link can transmit, but not at same time
- Framing
 - Continuous bit stream at physical layer
 - encapsulate datagram into frame, adding header, trailer
- Link access
 - channel access if shared medium

Channel/Medium Access


Two types of “links”:

- **point-to-point** (dedicated pairwise communication)

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


- **broadcast** (shared wire or medium)

- traditional Ethernet
- need to detect collision
- 802.11 wireless LANs
- Cellular data networks



shared wire
e.g. Ethernet

shared wire
(e.g. WaveL)



satellite



ZZZZZZZZZZZZZZZZZZ



cocktail party

MAC Addresses

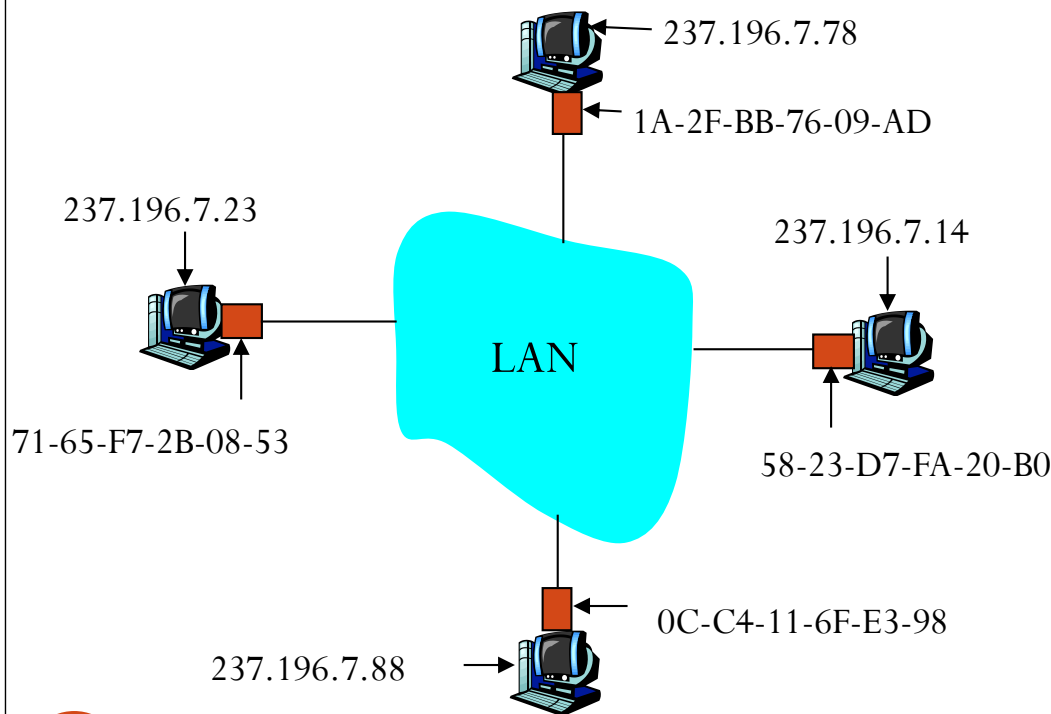
- 32-bit IP address:
 - network-layer address
 - Configured, or learned dynamically
 - used to get datagram to destination IP subnet
- **MAC address:**
 - used to get datagram from one interface to another physically-connected interface (same network)
 - **48 bit** MAC address (e.g. 1A-2F-BB-76-09-AD in **HEX**)
Unique, hard-coded in the adapter ROM when it is built
- **Broadcast address** (FF-FF-FF-FF-FF-FF)
 - Send the frame to **all** adapters

MAC Address (more)

- allocation administered by IEEE
- manufacturer buys portion of MAC address space (to assure uniqueness)
- MAC flat address → portability
 - can move LAN card from one LAN to another
- IP addresses are hierarchical and are NOT portable
 - depends on the IP subnet to which node is attached

ARP: Address Resolution Protocol

Question: how to determine
MAC address of B
knowing B's IP address?



- Each IP node (Host, Router) on LAN has an **ARP** table
- ARP Table: IP/MAC address mappings for some LAN nodes
< IP address; MAC address; TTL >
- TTL (Time To Live): time after which address mapping will be forgotten (typically 20 min)

ARP Request


Argon has the IP address of Router 137 (e.g., via DHCP)

Argon broadcasts an ARP request to all stations on the network:

“What is the MAC address of Router137?”

Argon
128.143.137.144
00:a0:24:71:e4:44

Router137
128.143.137.1
00:e0:f9:23:a8:20




ARP Reply

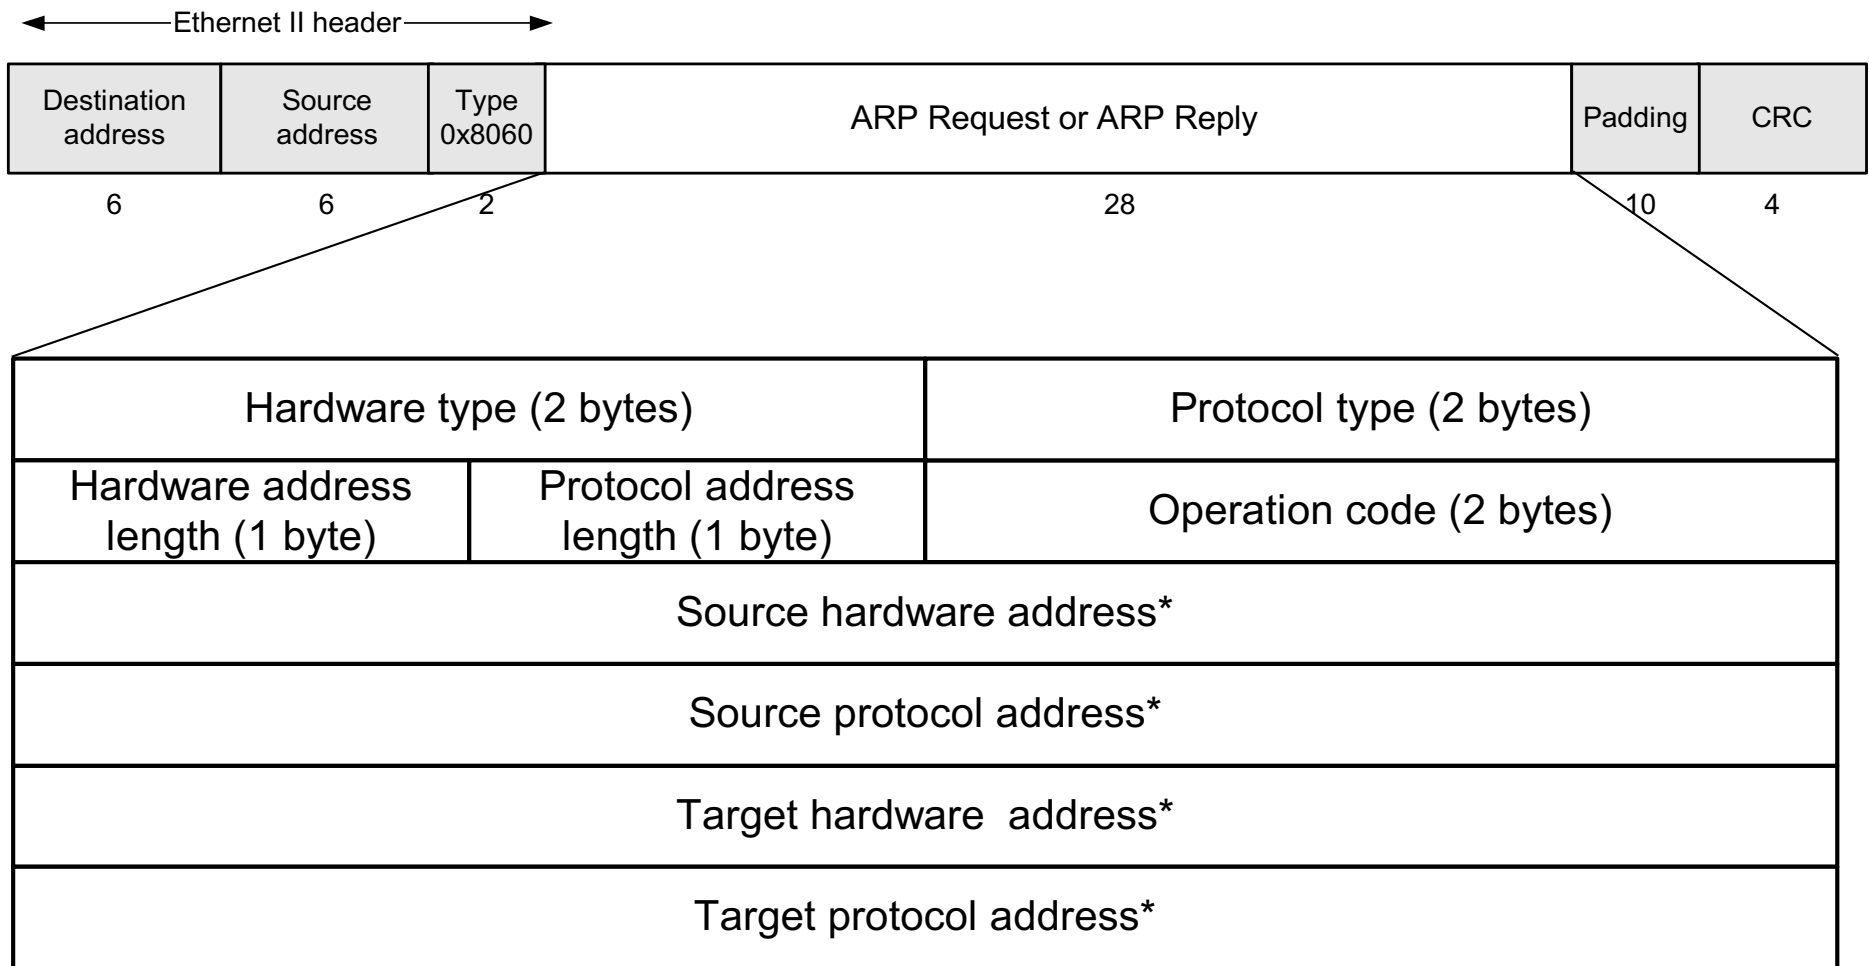
Router 137 responds with an ARP Reply which contains its MAC address

Argon
128.143.137.144
00:a0:24:71:e4:44

Router137
128.143.137.1
00:e0:f9:23:a8:20



ARP Packet Format



* Note: The length of the address fields is determined by the corresponding address length fields

Example

- *ARP Request from Argon:*

Source hardware address:	00:a0:24:71:e4:44
Source protocol address:	128.143.137.144
Target hardware address:	00:00:00:00:00:00
Target protocol address:	128.143.137.1

- *ARP Reply from Router137:*

Source hardware address:	00:e0:f9:23:a8:20
Source protocol address:	128.143.137.1
Target hardware address:	00:a0:24:71:e4:44
Target protocol address:	128.143.137.144

“Specialized” ARP messages

- **ARP Probe** (a special kind of request): check if anyone using an IP address "Is anyone using this address?" (and, "This is the address I hope to use.")
 - *all-zero sender IP address*
 - *sender hardware address set to its own MAC address*
 - *target MAC address set to all zeros*
 - *target IP address set to the address being probed*
- **ARP Announcement** (gratuitous APR): "This is the IP address I am now using."
 - both the sender and target IP address fields contain the IP address being announced

DHCP versus ARP


- DHCP: **Acquiring an IP Address**
 - Broadcast “I need an IP address, please!”
 - Response “You can have IP address 192.168.1.245”
- ARP: **Discovering the Receiver**
 - Broadcast “who has IP address 128.143.137.1?”
 - Response “00-E0-23-6F-A8-20 has 128.143.137.1!”


ARP Cache

- Since sending an ARP request/reply for each IP datagram is inefficient, hosts maintain a cache (ARP Cache) of current entries. The entries expire after **20 minutes**.
- Contents of the ARP Cache:
 - (128.143.71.37) at 00:10:4B:C5:D1:15 [ether] on eth0
 - (128.143.71.36) at 00:B0:D0:E1:17:D5 [ether] on eth0
 - (128.143.71.35) at 00:B0:D0:DE:70:E6 [ether] on eth0
 - (128.143.136.90) at 00:05:3C:06:27:35 [ether] on eth1
 - (128.143.71.34) at 00:B0:D0:E1:17:DB [ether] on eth0
 - (128.143.71.33) at 00:B0:D0:E1:17:DF [ether] on eth0

MAC spoofing & ARP Poisoning

- MAC addresses can be easily modified
 - One liner on OSX: `sudo ifconfig en0 ether 5e:c4:a4:99:b8:e3` [require root access]
 - Implications: **MAC address filtering in WLAN access control isn't really "secure"**
- ARP poisoning
 - No authentication in ARP
 - An attacker use ARP announcements to poison ARP caches





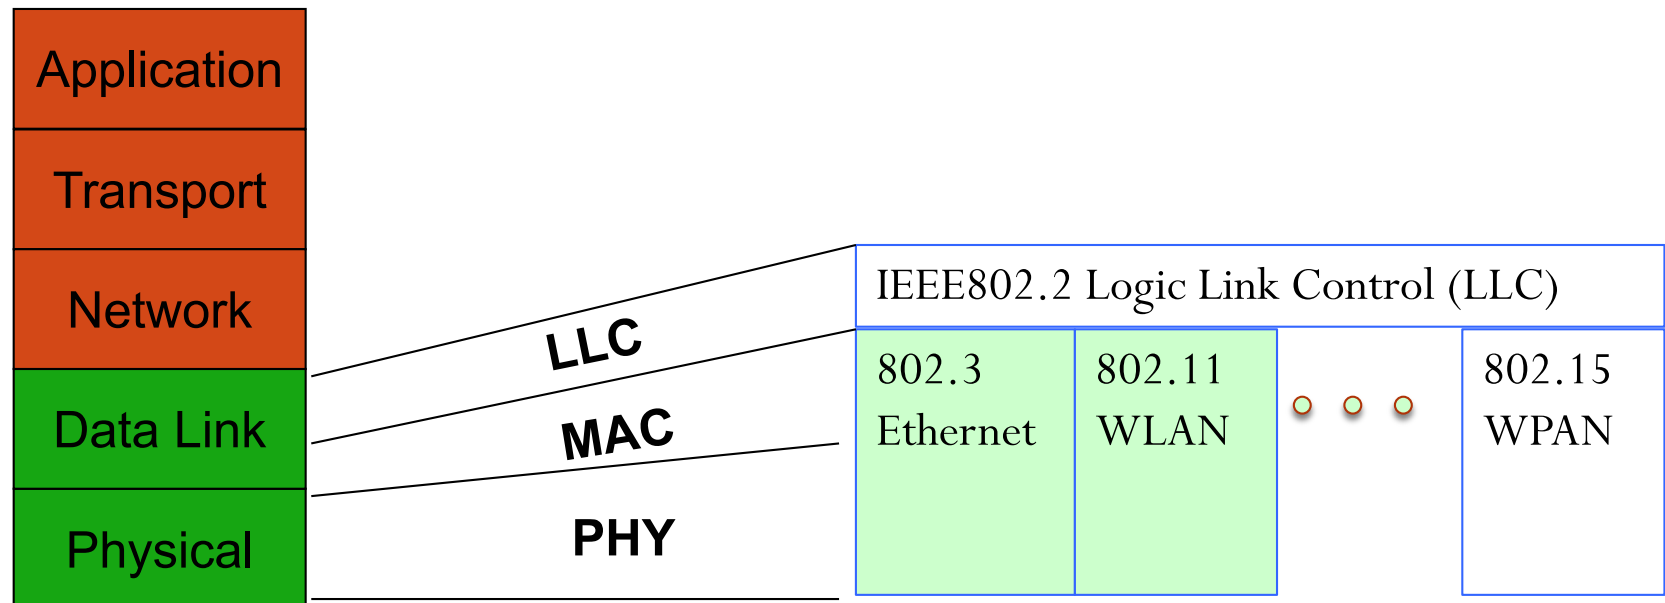
• A -> B

• A -> D

Ethernet

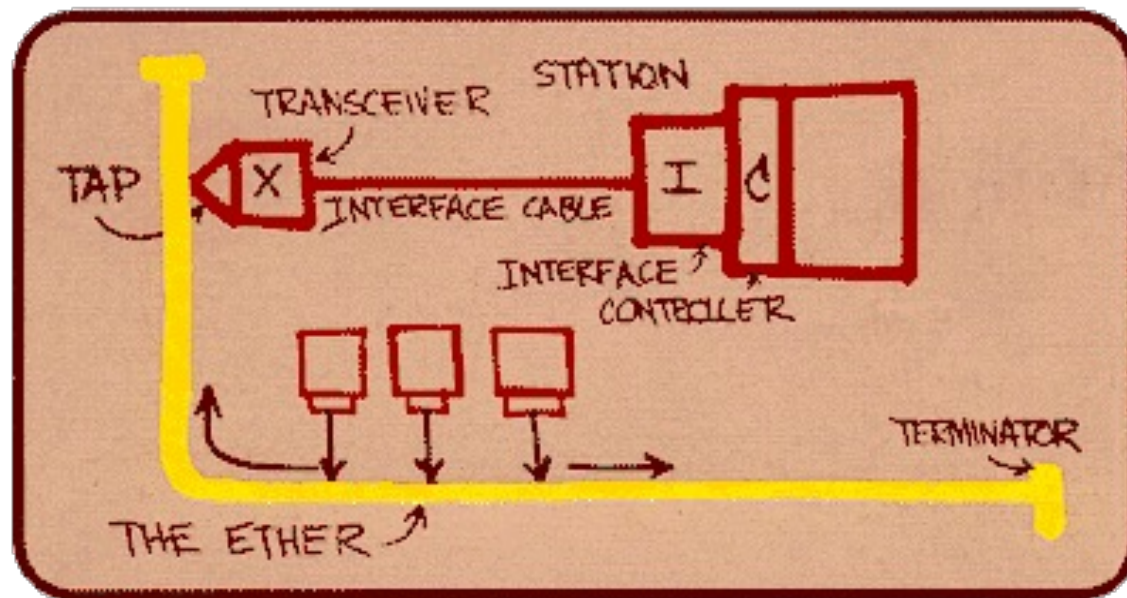
IEEE 802 Protocol Suite

- A family of IEEE standards for body, personal, local area networks and metropolitan area networks
- Specifies **data link & physical layer**
 - 802.2 logical link control (LLC)
 - 802.3 Ethernet
 - 802.11 wireless local area networks (WLAN)
 - 802.15 Wireless personal area networks (WPAN) – bluetooth, zigbee, body area networks, etc.



Ethernet

- “dominant” wired LAN technology:
 - cheap ~\$20 for 1Gbps cards, ~\$100 for 10Gbps cards
 - first widely used LAN technology
- Kept up with speed race: 10 Mbps – 400Gbps




Metcalfe's Ethernet sketch

Bus vs. Star topology

- Bus topology popular through mid 90s





- Now star topology prevails
- Connection choices: hub or switch




Hubs vs Switches

- Both connect segments of LANs
- Hubs are **layer-1** devices that duplicate the messages to other ports
 - Computers connected by hubs are in **the same contention domain**
- Switches are **layer-2** devices that forward the messages to the selected port
 - Computers connected by switches are in the same **broadcast domain**
 - **Point-to-point** link between a computer and a switch




Ethernet Frame Structure

- Encapsulates IP datagram (or other network layer protocol packet, ARP) in Ethernet frame



- **Preamble:**
 - 7 bytes with pattern 10101010 followed by 1 byte with pattern 10101011
 - used to synchronize receiver, sender clock rates

Ethernet Frame Structure (more)



- Addresses: 6 bytes, frames are received by all adapters on a LAN and dropped if dest. address does not match and is not broadcast
- Type: 2 bytes, indicates the types of the Ethernet frame (Ethernet II, 802.2 LLC frame,...)
- CRC: 4 bytes, checked at receiver, if error is detected, the frame is simply dropped
- Data payload: maximum 1500 bytes, minimum 46 bytes
 - If data is less than 46 bytes, pad with zeros to 46 bytes
 - (“Jumbo frame” up to 9000 bytes for 1G/10G/100G)

Unreliable, connectionless service

- **Connectionless**: No handshaking between sending and receiving adapter.
- **Unreliable**: receiving adapter doesn't send acks or nacks to sending adapter
 - stream of datagrams passed to network layer can have gaps
 - gaps will be filled if app is using TCP
 - otherwise, app will see the gaps

Medium Access Control in Ethernet


- Shared medium with bus or star topology with hubs
 - Carrier Sensing Multiple Access with Collision Detection (CSMA/CD)
 - Half-duplex: nodes take turns in transmissions
- Point-to-point in star topology with switches
 - Full duplex: can send and receive at the same time
 - How to learn which port a device is connected?

Switched Ethernet


- Switches are link layer device
 - stores and forwards Ethernet frames
 - examines frame header and selectively forwards frame based on MAC dest address
- transparent
 - hosts are unaware of presence of switches
- plug-and-play, self-learning
 - switches do not need to be configured

Switch: traffic isolation


- switch installation breaks **subnet** into LAN **segments**
- switch filters packets:
 - same-LAN-segment frames not usually forwarded onto other LAN segments
 - segments become separate collision/contention domains



Institutional network



Forwarding



- How to determine onto which LAN segment to forward frame?
- Looks like a routing problem...

Self learning

- A switch has a **switch table**
- entry in switch table:
 - (MAC Address, Interface, TTL)
 - stale entries in table dropped (TTL can be 60 min)
- switch **learns** which hosts can be reached through which interfaces
 - when frame received, switch “learns” location of sender:
incoming LAN segment
 - records sender/location pair in switch table

Learning & Filtering/Forwarding

When switch receives a frame:

record (**src MAC**, interface) in switch table

index switch table using **dest MAC** address

if entry **found** for destination

then {

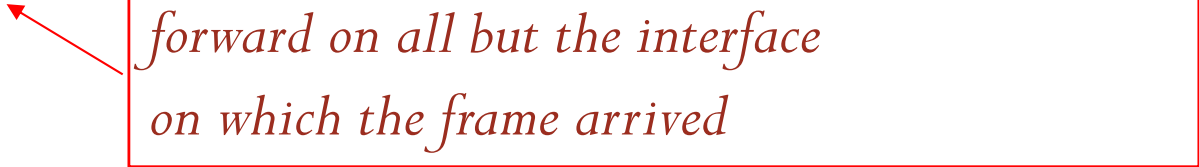
 if **dest** on segment from which frame arrived

 then drop the frame

 else forward the frame on interface indicated

}


else flood



*forward on all but the interface
on which the frame arrived*

Switch example

- Suppose C sends frame to D



Switch receives frame from from C


notes in switch table that C is on interface 1

because D is not in table, switch forwards frame into interfaces 2 and 3

frame received by D

Switch example

- Suppose C sends frame to D



Switch receives frame from from C


notes in switch table that C is on interface 1

because D is not in table, switch forwards frame into interfaces 2 and 3

frame received by D

Switch example

- Suppose D replies back with frame to C.



Switch receives frame from from D


notes in the switch table that D is on interface 2

because C is in table, switch forwards frame only to interface 1

frame received by C

Switch example

- Suppose D replies back with frame to C.



Switch receives frame from from D

notes in the switch table that D is on interface 2


because C is in table, switch forwards frame only to interface 1

frame received by C

Another example

- Initial entries in switch tables are in black
- C sends a frame to D


Switch X	
address	interface
A	1
B	2
C	3



Switch W	
address	interface
A	1
B	1
C	1

Another example


- Initial entries in switch tables are in black
- C sends a frame to D
- D replies back with a frame to C.



Switch X	
address	interface
A	1
B	2
C	3
D	0

Switch W	
address	interface
A	1
B	1
C	1
D	2

Summary: Comparison of Hubs, Switches & Routers



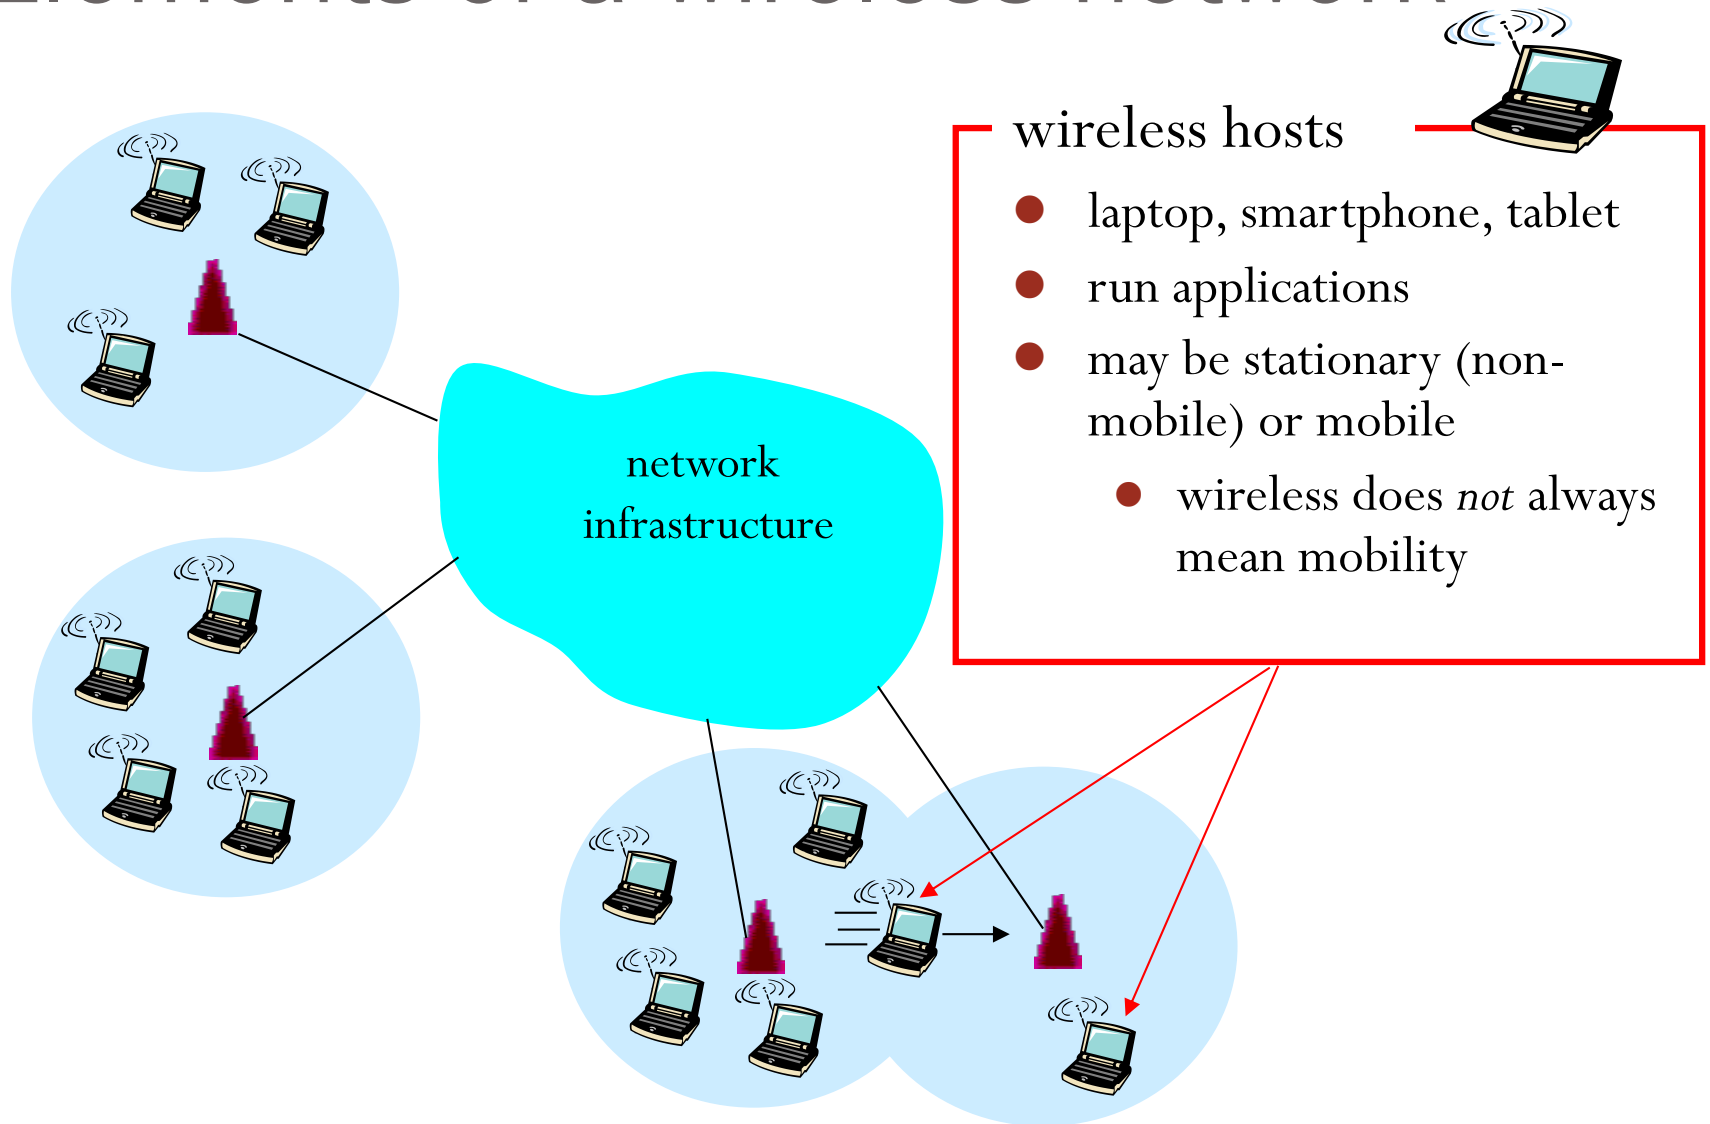
	<u>hubs</u>	<u>routers</u>	<u>switches</u>
traffic isolation	no	yes	yes
plug & play	yes	no*	yes
optimal routing	no	yes	no
cut through	yes	no	yes*

Discussion


- Precondition for self-learning
- Why not apply it to intra-domain routing?

WLANs

Elements of a wireless network



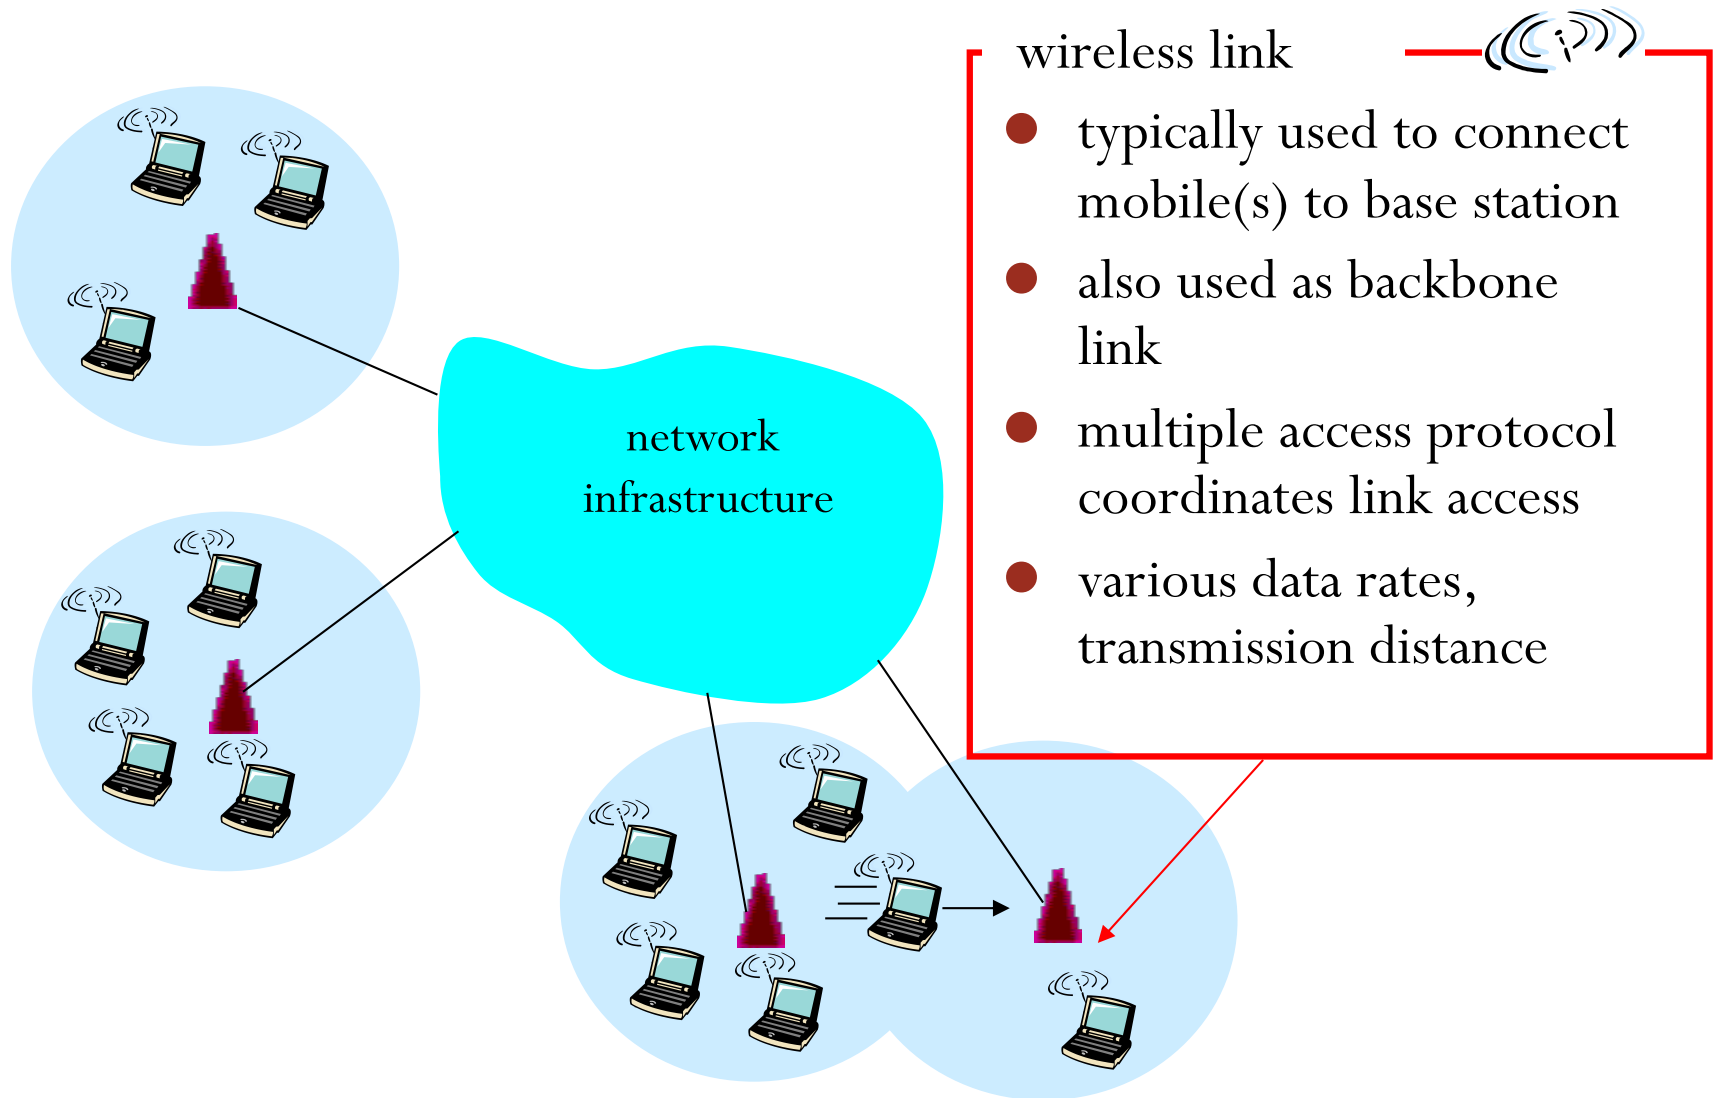
Elements of a wireless network




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



Elements of a wireless 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

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 and access point (AP): base station
 - ad hoc mode: hosts only (IBSS)
- Distribution system (DS)
 - Connects multiple APs
- Extended service set (ESS)
 - Two or more basic service sets interconnected by DS

IEEE 802.11 Specs


Frequency range, or type	PHY	Protocol	Release date ^[18]	Frequency	Bandwidth	Stream data rate ^[19]	Max. MIMO streams	Modulation	Approx. range	
				(GHz)					In-door	Out-door
1–7 GHz	DSSS ^[20] , FHSS ^[A]	802.11-1997	June 1997	2.4	22	1, 2	—	DSSS, FHSS ^[A]	20 m (66 ft)	100 m (330 ft)
	HR/DSSS ^[20]	802.11b	September 1999	2.4	22	1, 2, 5.5, 11	—	CCK, DSSS	35 m (115 ft)	140 m (460 ft)
	OFDM	802.11a	September 1999	5	5, 10, 20	6, 9, 12, 18, 24, 36, 48, 54 (for 20 MHz bandwidth, divide by 2 and 4 for 10 and 5 MHz)	—	OFDM	35 m (115 ft)	120 m (390 ft)
		802.11j	November 2004	4.9, 5.0 ^{[B][21]}					?	?
		802.11y	November 2008	3.7 ^[C]					?	5,000 m (16,000 ft) ^[C]
		802.11p	July 2010	5.9					200 m ^[22]	1,000 m (3,300 ft) ^[22]
		802.11bd	December 2022	5.9, 60					500 m ^[23]	1,000 m (3,300 ft)
	ERP-OFDM ^[23]	802.11g	June 2003	2.4					38 m (125 ft)	140 m (460 ft)
	HT-OFDM ^[24]	802.11n (Wi-Fi 4)	October 2009	2.4, 5	20	Up to 288.8 ^[D]	4	MIMO-OFDM (64-QAM)	70 m (230 ft)	250 m (820 ft) ^[25]
					40	Up to 600 ^[D]				
	VHT-OFDM ^[24]	802.11ac (Wi-Fi 5)	December 2013	5	20	Up to 693 ^[D]	8	DL MU-MIMO OFDM (256-QAM)	35 m (115 ft) ^[26]	?
					40	Up to 1600 ^[D]				
					80	Up to 3467 ^[D]				
					160	Up to 6933 ^[D]				
	HE-OFDMA	802.11ax (Wi-Fi 6, Wi-Fi 6E)	May 2021	2.4, 5, 6	20	Up to 1147 ^[E]	8	UL/DL MU-MIMO OFDMA (1024-QAM)	30 m (98 ft)	120 m (390 ft) ^[F]
					40	Up to 2294 ^[E]				
					80	Up to 5.5 Gbit/s ^[E]				
					80+80	Up to 11.0 Gbit/s ^[E]				
	EHT-OFDMA	802.11be (Wi-Fi 7)	Sep 2024 (est. ^[27])	2.4, 5, 6	80	Up to 11.5 Gbit/s ^[E]	16	UL/DL MU-MIMO OFDMA (4096-QAM)	30 m (98 ft)	120 m (390 ft) ^[F]
					160 (80+80)	Up to 23 Gbit/s ^[E]				
					240 (160+80)	Up to 35 Gbit/s ^[E]				
					320 (160+160)	Up to 46.1 Gbit/s ^[E]				

Non-Overlapping Channels for 2.4 GHz WLAN


802.11b (DSSS) channel width 22 MHz



802.11g/n (OFDM) 20 MHz ch. width - 16.25 MHz used by sub-carriers



802.11n (OFDM) 40 MHz ch. width - 33.75 MHz used by sub-carriers




5GHz (802.11a/h/j/n/ac)

Channel Width	Valid Channel Numbers
20 MHz	36, 40, 44, 48, 52, 56, 60, 64, 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, 140, 144, 149, 153, 161, 165, 169
40 MHz	38, 46, 54, 62, 102, 110, 118, 126, 134, 142, 151, 159
80 MHz	42, 58, 106, 122, 138, 155
160 MHz	50, 114

802.11 frame: addressing

in bytes




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

duration of reserved
transmission time (RTS/CTS)

frame seq #
(for reliable ARQ)



frame type
(RTS, CTS, ACK, data)


Frame Types

- Management frame
 - Beacon
 - (De)association request/respond
 - Announcement traffic indication message
 - Authentication/Deauthentication
- Control frame
 - Poll frame & poll response frame
 - RTS
 - CTS
 - ACK
 - Power save (PS-poll)
- Data frame
 - Limitation on payload size
 - Can be extended to 7395 (with multiple fragments)

Association


- host: must **associate** with an AP
 - scans channels, listening for beacon frames containing service set identifier and AP's MAC address
 - SSID is 32 octets long
 - One SSID per network (BSS or IBSS)
 - selects AP to associate with; initiates association protocol
 - may perform authentication
 - will typically then run DHCP to get IP address in AP's subnet

802.11: Association



passive scanning:

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




active scanning:

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

IEEE 802.11: multiple access


- Half-duplex
- Use Carrier Sensing Multiple Access (CSMA)
 - Random access
 - A device **listens (senses)** to the communication medium before attempting to send data. If the channel is **idle**, the device is allowed to transmit.
 - **Cannot detect collision** — transmit all frames to completion
- **with acknowledgment** — because without collision detection, you don't know if your transmission collided or not
- Avoid collisions -- **CSMA / C(ollision)A(voidance)**

Collisions in WLAN



Propagation delay leads to delay in detecting transmissions in the medium

Hidden Terminal Problem



A, C cannot sense each's transmission. Their transmissions can collide at B


Medium Access Control Logic

802.11 sender

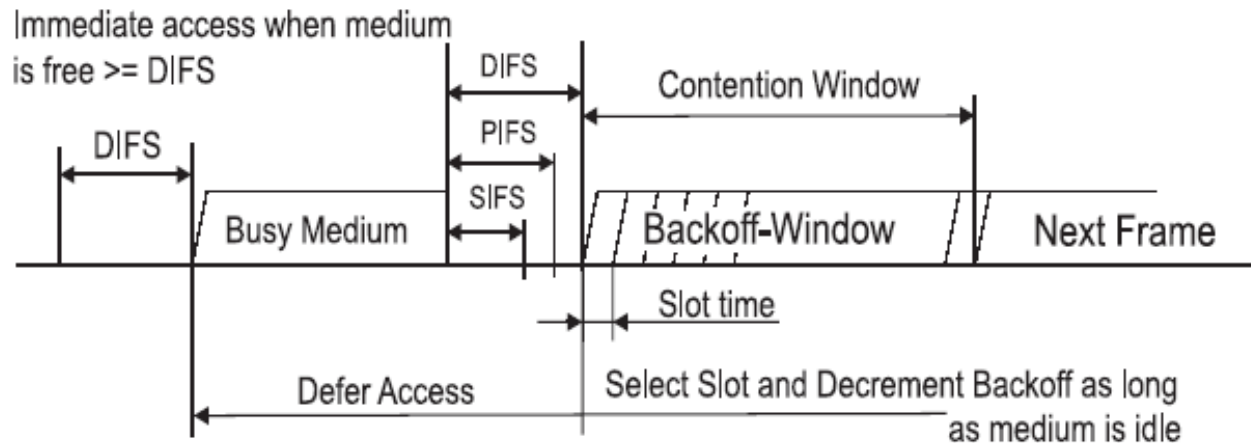
- 1 For a new frame, if sense channel idle for **DIFS** then
transmit entire frame
 - 2 For a retransmitted frame or if sense channel busy then
 - start random backoff time
 - if sense channel idle for DIFS then
 - timer counts down while channel idle
 - transmit when timer expires
- if no ACK for transmission, increase random backoff interval,
repeat 2

802.11 receiver

- if frame received OK
return ACK after **SIFS**




Interframe Space (IFS)



- Short IFS (**SIFS**)
 - Shortest IFS (used for ACK, CTS, poll response)
 - Used for immediate response actions
- Point coordination function IFS (PIFS)
 - Midlength IFS
 - Used by centralized controller in PCF scheme when using polls
- Distributed coordination function IFS (**DIFS**)
 - Longest IFS (data, RTS)
 - Used as minimum delay of asynchronous frames contending for access
- Extended Interframe space (EIFS)
 - Used when received frame containing errors

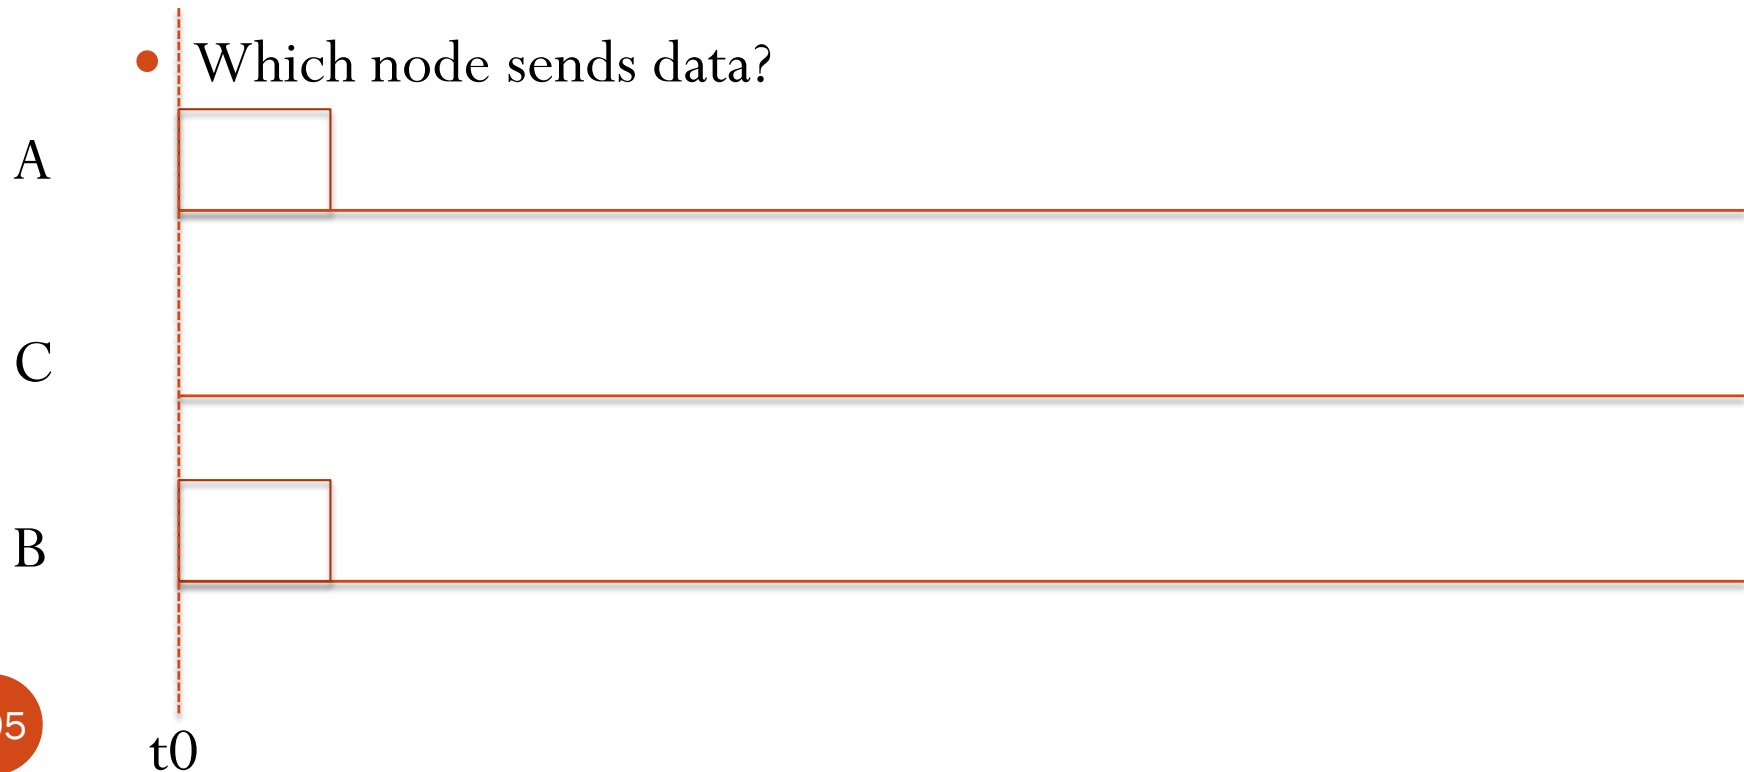
More on Medium Access Control Logic

- Contention window (in slots)
 - Each station maintains a contention window (CW) set to CW_{min} initially
 - Upon collision, $CW' = (CW + 1) * 2 - 1$ (**exponentially backoff**) till it reaches CW_{max} .
 - CW is reset to CW_{min} upon successful delivery
- Backoff time is drawn uniformly from $[0, CW]$




Example

- A pair of nodes **A** and **B** are sending packets to node **C**
- back-off intervals: Node A: 2, 3 and B are 2, 4
- Both nodes count down at t_0 .
- Size of the **backoff-window** in slot before send data?
- Which node sends data?



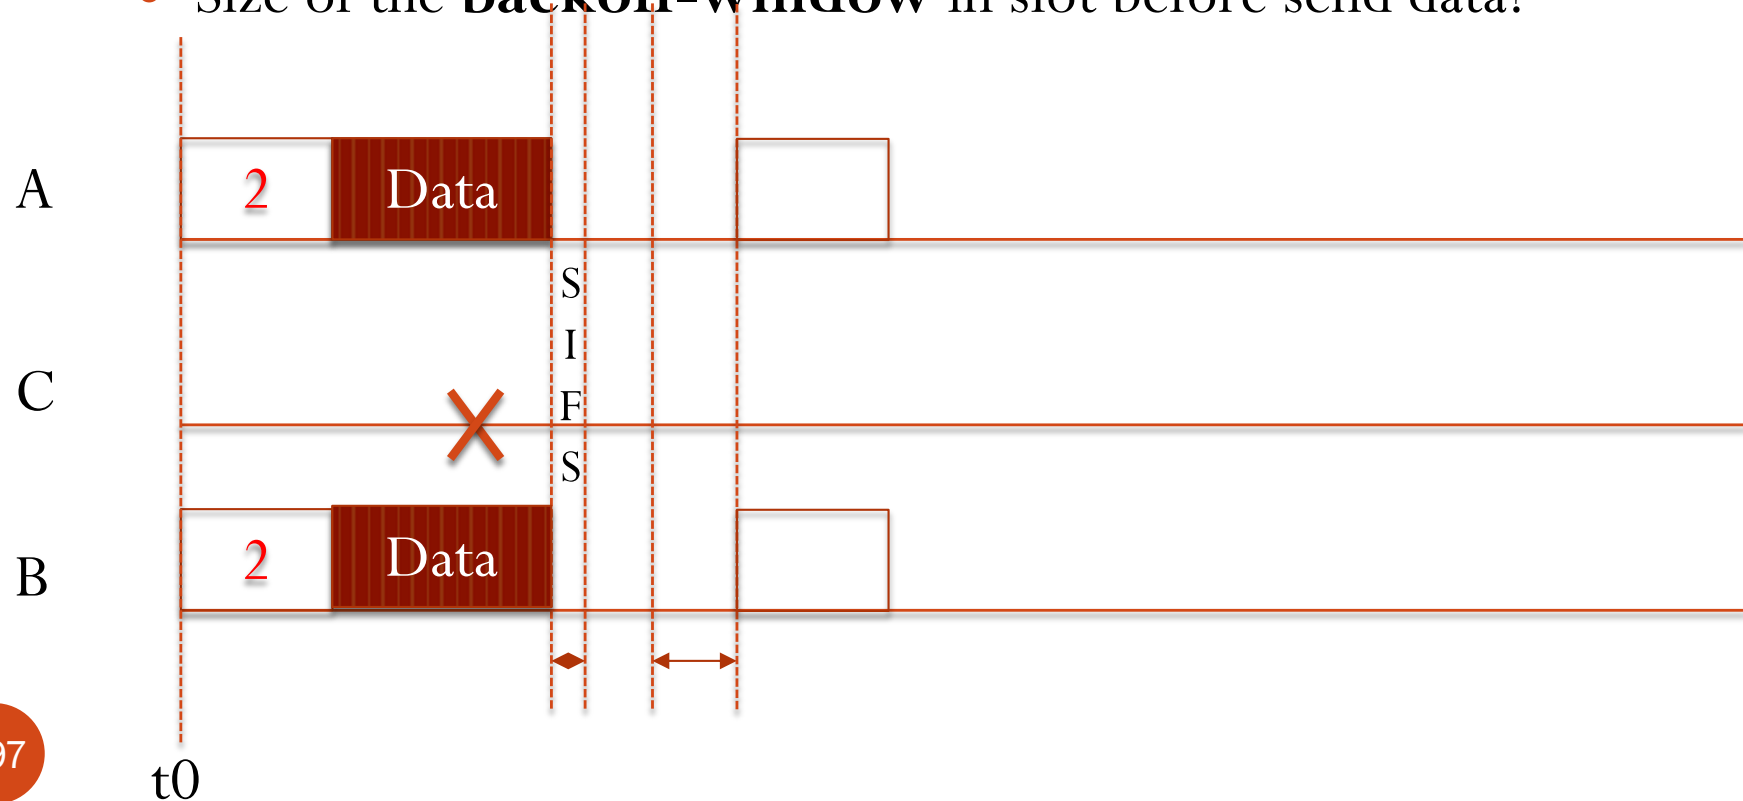
Example

- A pair of nodes **A** and **B** are sending packets to node **C**
- back-off intervals: Node A: 2, 3 and B are 2, 4
- Size of the Interframe Space, will C sends ACK?




Example

- A pair of nodes **A** and **B** are sending packets to node **C**
- back-off intervals: Node A: 2, 4 and B are 2, 3
- Size of the **Interframe Space** before A, B countdown?
- Size of the **backoff-window** in slot before send data?



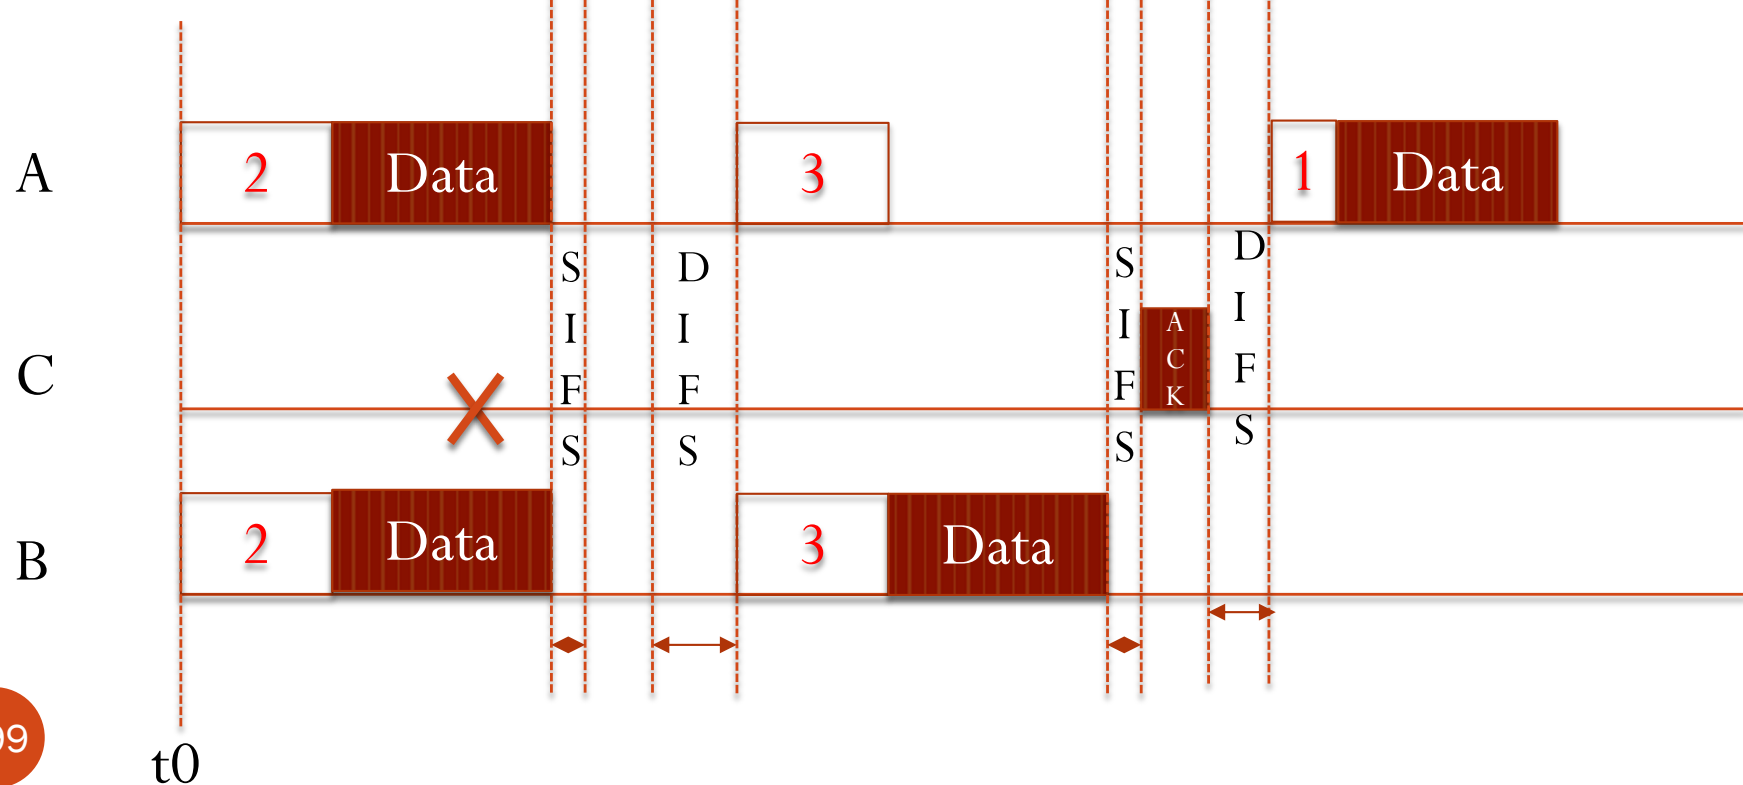
Example

- A pair of nodes **A** and **B** are sending packets to node **C**
- back-off intervals: Node A: 2, 4 and B are 2, 3
- Size of the **Interframe Space** before A countdown?
- Size of the **backoff-window** in slot before sends data?



Example

- A pair of nodes **A** and **B** are sending packets to node **C**
- back-off intervals: Node A: 2, 4 and B are 2, 3
- Size of the **Interframe Space** before A, countdown?
- Size of the **backoff-window** in slot before sends data?



Summary

- Ethernet
 - Frame format
 - Self-learning algorithm
- WLAN
 - DS, BSS, IBSS, ESS
- CSMA/CA
 - Physical and virtual carrier sensing
 - Defer transmission after a busy period
 - Exponential backoff
- IEEE 802.11 frame format

Synthesis: a day in the life of a web request

- journey down protocol stack complete!
 - application, transport, network, link
- putting-it-all-together: synthesis!
 - *goal:* identify, review, understand protocols (at all layers)
involved in seemingly simple scenario: requesting www page
 - *scenario:* student attaches laptop to campus WiFi network,
requests/receives www.google.com