


The Link Layer and LANs



Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:


<http://www.cse.wustl.edu/~jain/cse473-20/>



1. Datalink Services
2. Error Detection
3. Multiple Access
4. Bridging
5. MPLS

Note: This class lecture is based on Chapter 6 of the textbook (Kurose and Ross) and the figures provided by the authors.

Link Layer Services



- ❑ Link = One hop
- ❑ Framing: Bit patterns at begin/end of a frame
- ❑ Multiple Access: Multiple users sharing a wire
- ❑ Optional (On Lossy wireless links)
 - Flow Control
 - Error Detection/Correction
 - Reliable Delivery
- ❑ Duplex Operation

Line Duplexity

- ❑ Simplex: Transmit or receive, e.g., Television



- ❑ Full Duplex: Transmit and receive simultaneously, e.g., Telephone



- ❑ Half-Duplex: Transmit and receive alternately, e.g., Police Radio





Error Detection

- ❑ Parity Checks
- ❑ Check Digit Method
- ❑ Modulo 2 Arithmetic
- ❑ Cyclic Redundancy Check (CRC)
- ❑ Popular CRC Polynomials

Parity Checks

1 0 1 1 1 1 0 1 0

Odd Parity

1 0 1 1 1 1 0 1 0 0

1 2 3 4 5 6 7 8 9

0 0 1 1 1 1 0 1 0 0

1 2 3 4 5 6 7 8 9

1-bit error

0 0 0 1 0 0 1 0 0

1 2 3 4 5 6 7 8 9

3-bit error

0 0 0 1 1 1 0 1 0 0

1 2 3 4 5 6 7 8 9

2-bit error


Even Parity

1 0 1 1 1 1 0 1 1 0

1 2 3 4 5 6 7 8 9

Two Dimensional Parity

- Detect and correct single bit errors



1	0	1	0	1	1
1	1	1	1	0	0
0	1	1	1	0	1
0	0	1	0	1	0

no errors

1	0	1	0	1	1
1	0	1	1	0	0
0	1	1	1	0	1
0	0	1	0	1	0

parity error

*correctable
single bit error*

Check Digit Method

- Make a number divisible by 9

Example: 823 is to be sent

1. Left-shift: 8230
2. Divide by 9, find remainder: 4
3. Subtract remainder from 9: $9-4=5$
4. Add the result of step 3 to step 1: 8235
5. Check that the result is divisible by 9.

Detects all single-digit errors: 7235, 8335, 8255, 8237

Detects several multiple-digit errors: 8765, 7346

Does not detect some errors: 7335, 8775, ...

Does not detect transpositions: 2835

Credit card numbers are protected via a similar method called “Luhn Algorithm” which detects most transpositions.

Ref: http://en.wikipedia.org/wiki/Luhn_algorithm

Modulo 2 Arithmetic

$$\begin{array}{r}
 1111 \\
 +1010 \\
 \hline
 0101
 \end{array}$$

$$\begin{array}{r}
 11001 \\
 \times 11 \\
 \hline
 11001 \\
 11001 \\
 \hline
 101011
 \end{array}$$

$$\begin{array}{r}
 110 \\
 \hline
 11 \mid 1010 \\
 / 11 \downarrow \\
 \hline
 x11 \\
 11 \\
 \hline
 x00 \\
 00 \\
 \hline
 x0
 \end{array}$$

010	2	
011	3	
---	--	
001	1	Mod 2
101	5	Binary

Cyclic Redundancy Check (CRC)

□ Binary Check Digit Method

□ Make a number divisible by $P=110101$ ($n+1=6$ bits)

Example: $M=1010001101$ is to be sent

1. Left-shift M by n bits $2^n M = 101000110100000$

2. Divide $2^n M$ by P , find remainder: $R=01110$

~~3. Subtract remainder from P ← Not required in Mod 2~~

4. Add the result of step 2 to step 1 :

$T=101000110101110$

5. Check that the result T is divisible by P .

Modulo 2 Division

$$Q = \underline{1101010110}$$

$$P = 110101 \mid 101000110100000 = 2^n M$$

110101

111011

110101

011101

000000

111010

110101

011111

000000

111110

110101

010110

000000

101100

110101

110010

110101


001110

000000

01110 = R

Checking At The Receiver

$$\begin{array}{r} \underline{1101010110} \\ 110101)101000110101110 \\ \underline{110101} \\ 111011 \qquad \qquad \qquad 010111 \\ \underline{110101} \qquad \qquad \qquad \underline{000000} \\ 011101 \qquad \qquad \qquad 101111 \\ \underline{000000} \qquad \qquad \qquad \underline{110101} \\ 111010 \qquad \qquad \qquad 110101 \\ \underline{110101} \qquad \qquad \qquad \underline{110101} \\ 011111 \qquad \qquad \qquad 00000 \\ \underline{000000} \qquad \qquad \qquad 00000 \\ 111110 \\ \underline{110101} \end{array}$$



Error Detection: Review

1. **Parity bits** can help detect/correct errors
2. Remainder obtained by dividing by a **prime** number provides good error detection
3. **CRC** uses mod 2 division

Homework 6A: CRC

- [4 points] Find the CRC of 1001100 using a generator 1011. Use mod 2 division. Show all steps including the checking at the receiver.



Multiple Access Links and Protocols

1. Multiple Access
2. CSMA/CD
3. IEEE 802.3 CSMA/CD
4. CSMA/CD Performance
5. Cable Modem Access

Multiple Access




(a) Aloha Multiple Access



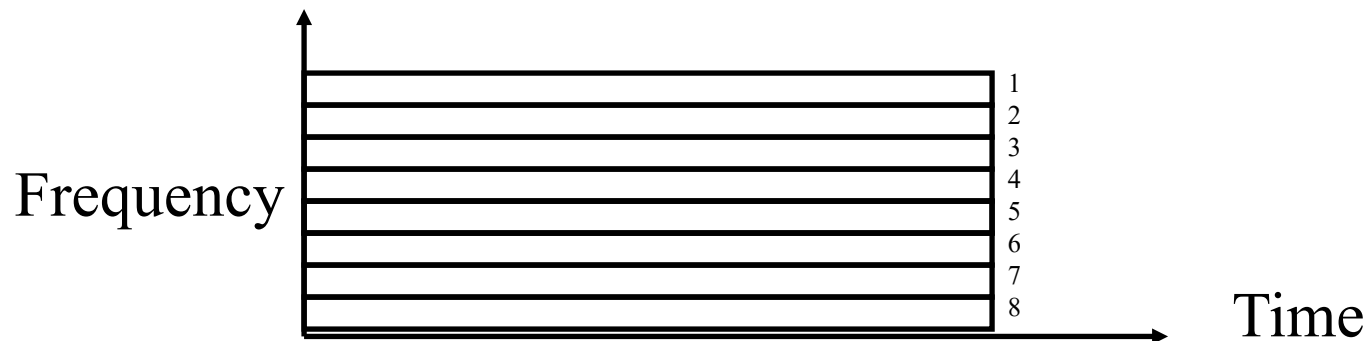
(b) Carrier-Sense Multiple Access with Collision Detection

Multiple Access

- ❑ How multiple users can share a link?
- ❑ **Time Division Multiplexing**



- ❑ **Frequency Division Multiplexing**



CSMA/CD



- ❑ **Aloha** at Univ of Hawaii:
Transmit whenever you like
Worst case utilization = $1/(2e) = 18\%$
- ❑ **Slotted Aloha**: Fixed size transmission slots
Worst case utilization = $1/e = 37\%$
- ❑ **CSMA**: Carrier Sense Multiple Access
Listen before you transmit
- ❑ **p-Persistent CSMA**: If idle, transmit with probability p . Delay by one time unit with probability $1-p$
- ❑ **CSMA/CD**: CSMA with Collision Detection
Listen while transmitting. Stop if you hear someone else

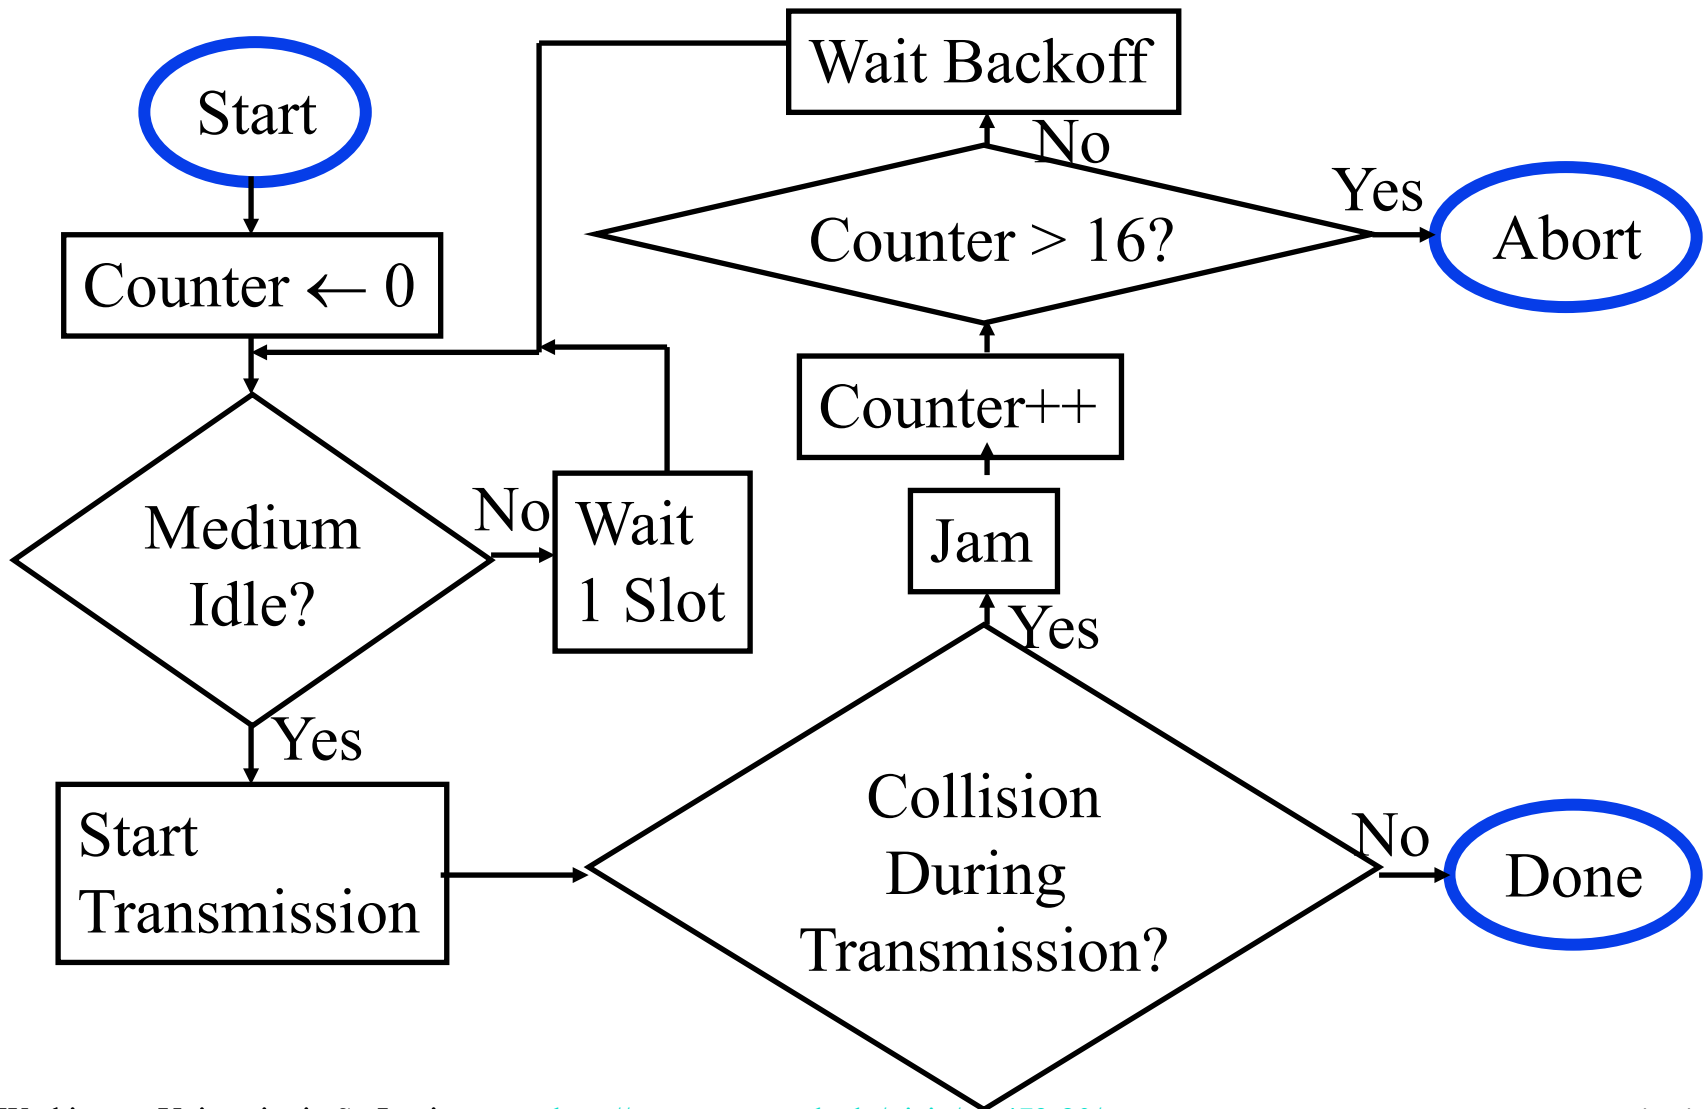
IEEE 802.3 CSMA/CD

- ❑ If the medium is idle, transmit (1-persistent).
- ❑ If the medium is busy, wait until idle and then transmit immediately.
- ❑ If a collision is detected while transmitting,
 - Transmit a **jam** signal for one slot
(= $51.2 \mu\text{s} = 64$ byte times)
 - Wait for a random time and reattempt (up to **16** times)
 - Random time = Uniform[$0, 2^{\min(k, 10)} - 1$] slots

Truncated Binary Backoff

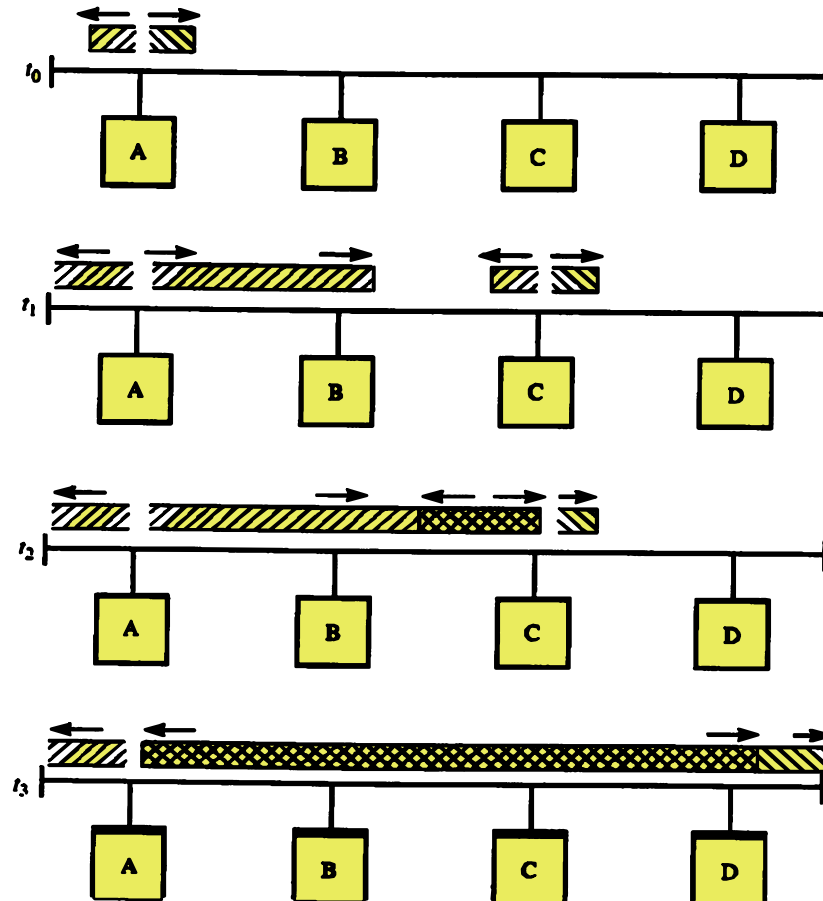
- ❑ Collision detected by monitoring the voltage
High voltage \Rightarrow two or more transmitters \Rightarrow Collision
 \Rightarrow Length of the cable is limited to 2 km

IEEE 802.3 CSMA/CD Flow Chart



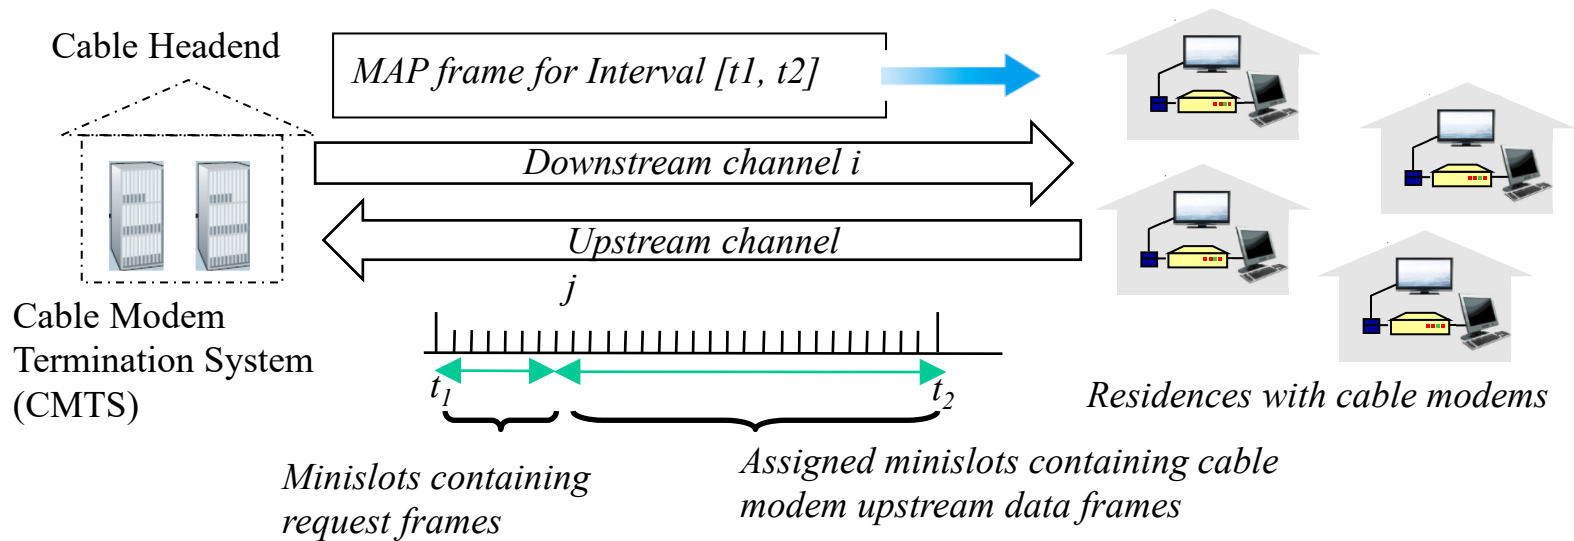
CSMA/CD Operation

- Collision window = $2 \times$ One-way Propagation delay = $51.2 \mu\text{s}$




One way delay
= $25.6 \mu\text{s}$
Max Distance
< 2.5 km

Cable Access Network



- ❑ **DOCSIS**: Data Over Cable Service Interface Specification
- ❑ Frequency Division Multiplexed channels over upstream and downstream
- ❑ Time Division Multiplexed slots in each upstream channel:
 - Some slots assigned, some have contention
 - Downstream **MAP** frame: Assigns upstream slots
 - Request for upstream slots (and data) transmitted random access (binary backoff) in selected slots

Multiple Access Links and Protocols: Review




1. Multiple users can share using **TDMA** or **FDMA**
2. Random access is better for data traffic.
3. Aloha has an efficiency of $1/2e$. Slotted Aloha makes it $1/e$.
4. Carrier sense and collision detection improves the efficiency further.
5. IEEE 802.3 uses **CSMA/CD** with **truncated** binary exponential backoff
6. DOCSIS used in cable access networks has **frequency division** multiplexed channels. With each channel **time division** multiplexed with some slots reserved for random access.



Switched Local Area Networks

1. Ethernet Standards
2. IEEE 802 Address Format
3. Address Resolution Protocol
4. Bridging
5. Virtual LANs

Ethernet Frame Structure



- ❑ **Preamble:** 7 bytes with pattern 10101010 followed by one byte with pattern 10101011. To synchronize receiver, sender clocks
- ❑ **Addresses:** 6 byte source, destination MAC addresses
- ❑ **Type:** indicates higher layer protocol
 - ❑ IP : 0x0800
 - ❑ ARP: 0x0806
- ❑ **CRC:** Cyclic Redundancy Check
 - ❑ If error detected: frame is silently dropped at the receiver
- ❑ **Connectionless:** No need to ask the receiver
- ❑ **Unreliable:** No ack, nack, or retransmissions

Ethernet Standards

- ❑ **10BASE5**: 10 Mb/s over coaxial cable (ThickWire)
- ❑ **10BROAD36**: 10 Mb/s over broadband cable, 3600 m max segments
- ❑ **1BASE5**: 1 Mb/s over 2 pairs of UTP
- ❑ **10BASE2**: 10 Mb/s over thin RG58 coaxial cable (ThinWire), 185 m max segments
- ❑ **10BASE-T**: 10 Mb/s over 2 pairs of UTP
- ❑ **100BASE-T4**: 100 Mb/s over 4 pairs of CAT-3, 4, 5 UTP
- ❑ **100BASE-TX**: 100 Mb/s over 2 pairs of CAT-5 UTP or STP
- ❑ **1000BASE-T**: 1 Gbps (Gigabit Ethernet)
- ❑ **10GBASE-T**: 10 Gbps
- ❑ **40GBASE-T**: 40 Gbps

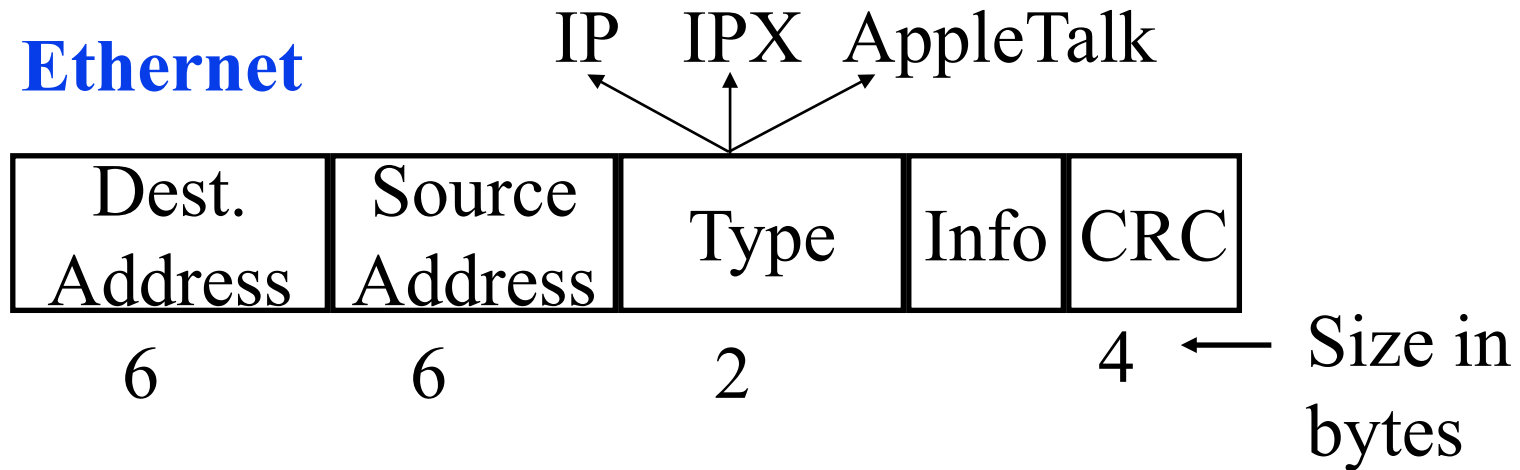
Ethernet vs. IEEE 802.3

IP	IPX	IP	IPX
Ethernet		Logical Link Control (LLC)	
		Media Access Control (MAC)	


- ❑ In 802.3, datalink was divided into two sublayers:
LLC and MAC
- ❑ **LLC** provides protocol multiplexing. MAC does not.
- ❑ **MAC** does not need a protocol type field.

Ethernet and 802.3 Frame Formats

□ Ethernet



□ IEEE 802.3



□ Length > 1518 ⇒ It is a protocol type ⇒ Ethernet

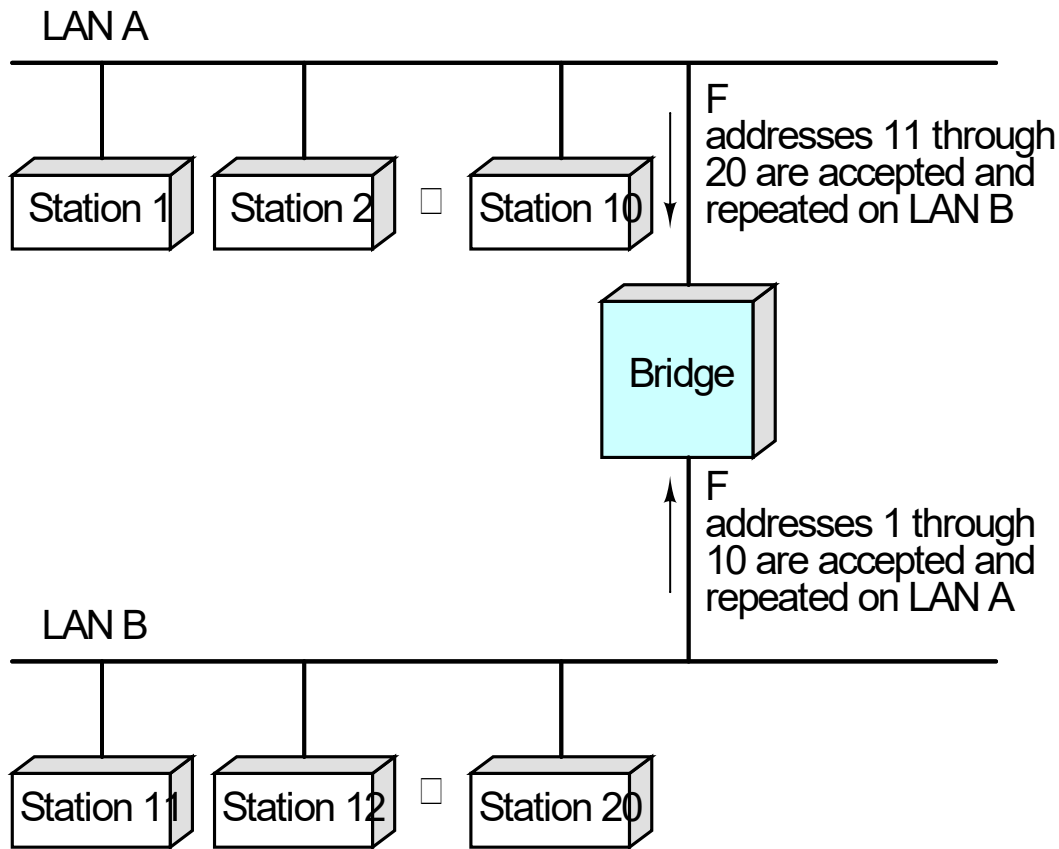
IEEE 802 Address Format

- 48-bit: 1000 0000 : 0000 0001 : 0100 0011
 : 0000 0000 : 1000 0000 : 0000 1100
 = 80:01:43:00:80:0C

Organizationally Unique Identifier (OUI)		24 bits assigned by OUI Owner
Individual/Group	Universal/Local	
1	1	22
		24

- Multicast = “To all bridges on this LAN”
- Broadcast = “To all stations”
 = 111111...111 = FF:FF:FF:FF:FF:FF


Bridges



Bridge: Functions

- ❑ Monitor all frames on LAN A
- ❑ Pickup frames that are for stations on the other side
- ❑ Retransmit the frames on the other side
- ❑ Knows or learns about stations are on various sides
Learns by looking at source addresses ⇒ **Self-learning**
- ❑ Makes **no modification** to content of the frames ⇒ **Transparent**
May change headers.
- ❑ Provides storage for frames to be forwarded
- ❑ Improves reliability (less nodes per LAN)
- ❑ Improves performance (more bandwidth per node)
- ❑ Security (Keeps different traffic from entering a LAN)
- ❑ May provide flow and congestion control (in Token Rings)


Full-Duplex Ethernet



- ❑ Uses point-to-point links between **TWO** nodes
- ❑ Full-duplex bi-directional transmission \Rightarrow Transmit any time
- ❑ Standardized in IEEE 802.3-2018
- ❑ All vendors are shipping switch/bridge/NICs with full duplex
- ❑ No collisions \Rightarrow 50+ km on fiber.
- ❑ Between servers and switches or between switches
- ❑ CSMA/CD is no longer used (except in old 10/100 hubs)
- ❑ 1G Ethernet standard allows CSMA/CD but not implemented.
- ❑ 10G and higher speed Ethernet standards do not allow CSMA/CD

Interconnection Devices

- ❑ **Repeater:** PHY device that restores data and collision signals
- ❑ **Hub:** Multiport repeater + fault detection, notification and signal broadcast
- ❑ **Bridge:** Datalink layer device connecting two or more collision domains
- ❑ **Router:** Network layer device (does not propagate MAC multicasts)

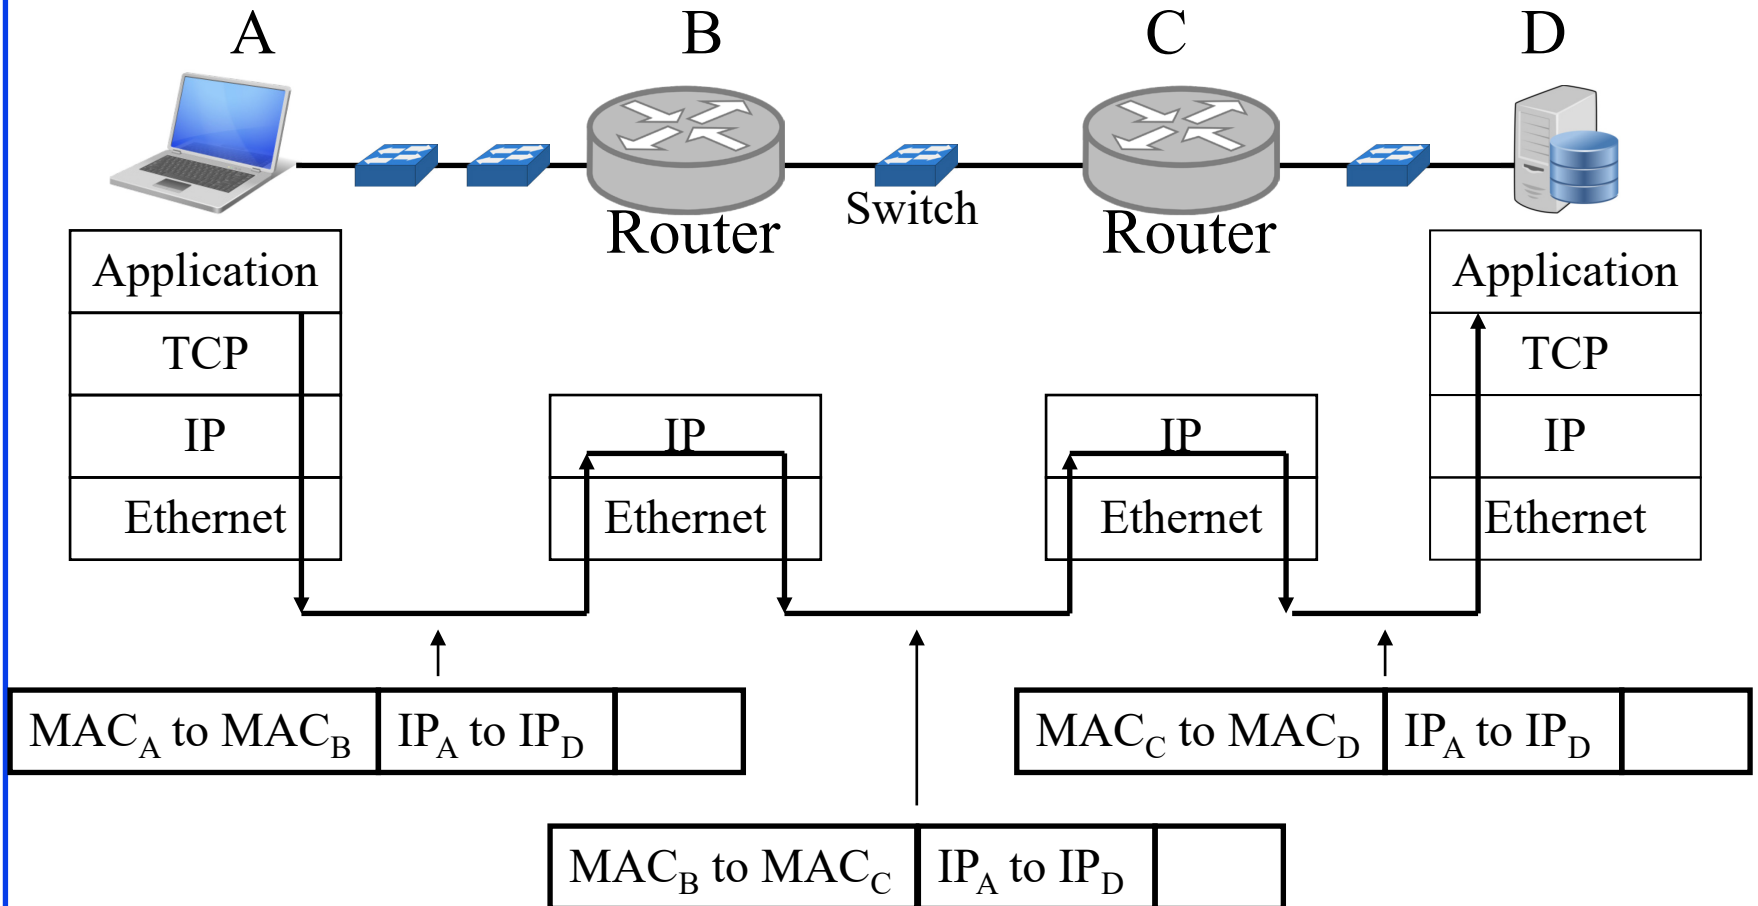


Address Resolution Protocol

- ❑ Problem: Given an IP address find the MAC address
- ❑ Solution: Address Resolution Protocol (**ARP**)
- ❑ The host broadcasts a request (Dest MAC=FFFFFFFF):
“What is the MAC address of 127.123.115.08?”
- ❑ The host whose IP address is 127.123.115.08 replies back:
“The MAC address for 127.123.115.08 is
8A:5F:3C:23:45:56₁₆”
- ❑ Nodes cache the MAC-IP mapping in a “ARP table”
You can list ARP table using “arp -a” command
- ❑ Frame Format: Hardware (HW): 0x0001 = Ethernet,
 - Protocol (Prot): 0x0800 = IP,
 - Operation: 1 = Request, 2=Response

HW Type	Prot Type	HW Addr Length	Prot Addr Length	Operation	Sender HW Addr	Sender Prot Addr	Target HW Addr	Target Prot Addr
16b	16b	8b	8b	16b	48b	32b	48b	32b


IP over Multiple Hops



- Switches = Transparent Bridges \Rightarrow No changes to frames
- ARP required only for nodes on the **same** “subnet”

Homework 6B: Collision Detection


- [6 Points] Suppose nodes A and B are on the same 10 Mbps Ethernet bus, and the propagation delay between the two nodes is 325 bit times. Suppose node A begins transmitting a frame and, before it finishes, node B begins transmitting a frame. Can A finish transmitting before it detects that B has transmitted? Why or why not? In the worst case when does B's signal reach A? (Minimum frame size is 512+64 bits).



Lab 6: Ethernet and ARP

[32 points] Download the Wireshark traces from <http://gaia.cs.umass.edu/wireshark-labs/wireshark-traces.zip>

Open *ethernet--ethereal-trace-1* in Wireshark. Select **View → Expand All**. This trace shows a HTTP exchange between end host A and Server B via Router R as shown below:



1. Examine HTTP request Frame 10. Answer the following questions.

- A. What is the 48-bit Ethernet source address? Who does it belong to: A, B, or R?
- B. What is the 48-bit Ethernet destination address? Who does it belong to: A, B, or R?
- C. What is the hexadecimal value for the two-byte Frame type field. What upper layer protocol does this correspond to?

Lab 6 (Cont)

- D. How many bytes from the very start of the Ethernet frame does the ASCII “G” in “GET” appear in the Ethernet frame? How many bytes are used up in Ethernet header, IP header, and TCP header before this first byte of HTTP message.
2. Examine Frame **16**. This is the HTTP OK response.
- A. What is the Ethernet source address? Who does it belong to: A, B, or R?
- B. What is the destination address in the Ethernet frame? Who does it belong to: A, B, or R?
- C. What is the hexadecimal value for the two-byte Frame type field. What upper layer protocol does this correspond to?
- D. How many bytes from the very start of the Ethernet frame does the ASCII “O” in “OK” appear in the Ethernet frame? How many bytes are used up in Ethernet header, IP header, and TCP header before the first byte of HTTP message.

Lab 6 (Cont)

3. Examine Frame 1. This is an ARP request.

- A. What are the hexadecimal values for the source and destination addresses in the Ethernet frame containing the ARP request message?
- B. What is the hexadecimal value for the two-byte Frame type field. What upper layer protocol does this correspond to?
- C. How many bytes from the very beginning of the Ethernet frame does the ARP opcode field begin?
- D. What is the value of the opcode field within the ARP-payload?
- E. What is the IP address of the sender?
- F. What is the target MAC and IP addresses in the ARP “question”?


Lab 6 (Cont)

4. *Examine Frame 2. This is the ARP response.*

- A. What are the hexadecimal values for the source and destination addresses in the Ethernet frame containing the ARP response message?
- B. What is the hexadecimal value for the two-byte Frame type field. What upper layer protocol does this correspond to?
- C. How many bytes from the very beginning of the Ethernet frame does the ARP opcode field begin?
- D. What is the value of the opcode field within the ARP-payload?
- E. What is the IP address of the sender?
- F. What is the target MAC and IP addresses in the ARP “answer”?

For all questions of this lab, please provide **numerical answers only**. No need to add screen captures.


What is a LAN?




- ❑ LAN = **Single broadcast domain** = Subnet
- ❑ No routing between members of a LAN
- ❑ Routing required between LANs

What is a Virtual LAN


Physical View



Logical View



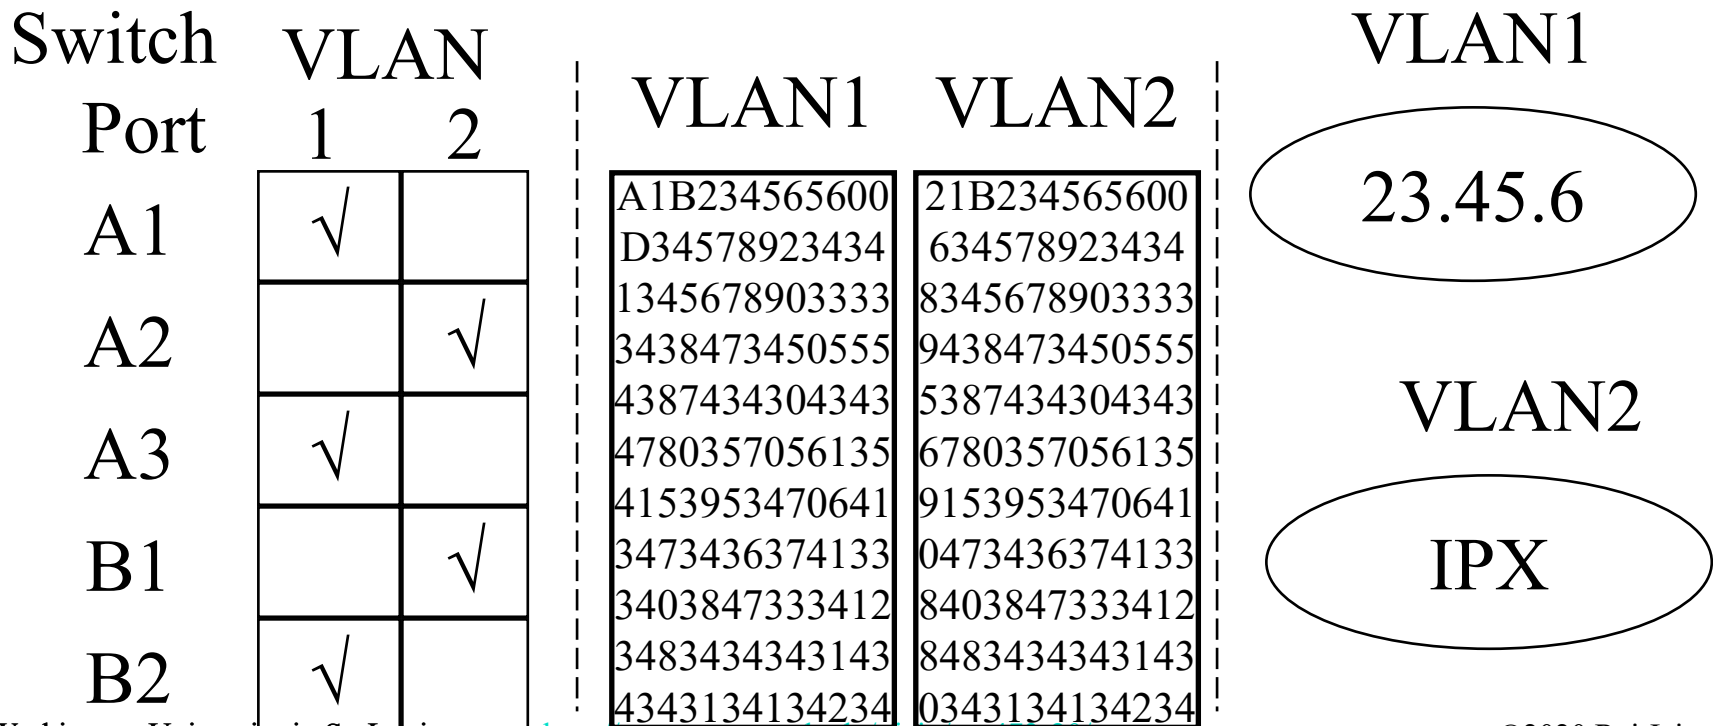
Virtual LAN



- ❑ Virtual LAN = Broadcasts and multicast goes only to the nodes in the virtual LAN
- ❑ LAN membership defined by the network manager
⇒ Virtual

Types of Virtual LANs

- Layer-1 VLAN = Group of Physical ports
- Layer-2 VLAN = Group of MAC addresses
- Layer-3 VLAN = IP subnet



IEEE 802.1Q-2011 Tag

- ❑ **Tag Protocol Identifier (TPI)**
- ❑ **Priority Code Point (PCP):** 3 bits = 8 priorities 0..7 (High)
- ❑ **Canonical Format Indicator (CFI):** 0 \Rightarrow Standard Ethernet, 1 \Rightarrow IBM Token Ring format (non-canonical or non-standard)
- ❑ CFI now replaced by Drop Eligibility Indicator (DEI)
- ❑ VLAN Identifier (12 bits \Rightarrow 4095 VLANs)
- ❑ Switches forward based on MAC address + VLAN ID
Unknown addresses are flooded.

Untagged
Frame



32b IEEE 802.1Q-2011 Header

Tagged
Frame



Ref: Canonical vs. MSB Addresses, http://support.lexmark.com/index?page=content&id=HO1299&locale=en&userlocale=EN_US
Ref: G. Santana, "Data Center Virtualization Fundamentals," Cisco Press, 2014, ISBN:1587143240
Washington University in St. Louis <http://www.cse.wustl.edu/~jain/cse473-20/>

©2020 Raj Jain



Switched Local Area Networks : Review

1. IEEE 802.3 uses a *truncated binary exponential backoff*.
2. Ethernet uses 48-bit addresses of which the first bit is the unicast/multicast, 2nd bit is universal/local, 22-bits are OUI (Organizationally unique identifier).
3. Ethernet bridges are **transparent** and **self-learning** using source addresses in the frame
4. Bridges are layer 2 devices while routers are layer 3 devices and do not forward layer 2 broadcasts
5. Address Resolution Protocol (**ARP**) is used to find the MAC address for a given IP address and vice versa.
6. IEEE **802.1Q tag** in Ethernet frames allows a LAN to be divided in to multiple VLANs. Broadcasts are limited to each VLAN and you need a router to go from one VLAN to another.




Multiprotocol Label Switching

Connection-oriented IP: Paths set up in advance


Borrowed from the Telephone networks

- ❑ Multiprotocol Label Switching (MPLS)
- ❑ Label Switching Example
- ❑ MPLS Forwarding Tables
- ❑ MPLS versus IP Paths
- ❑ MPLS Label Format

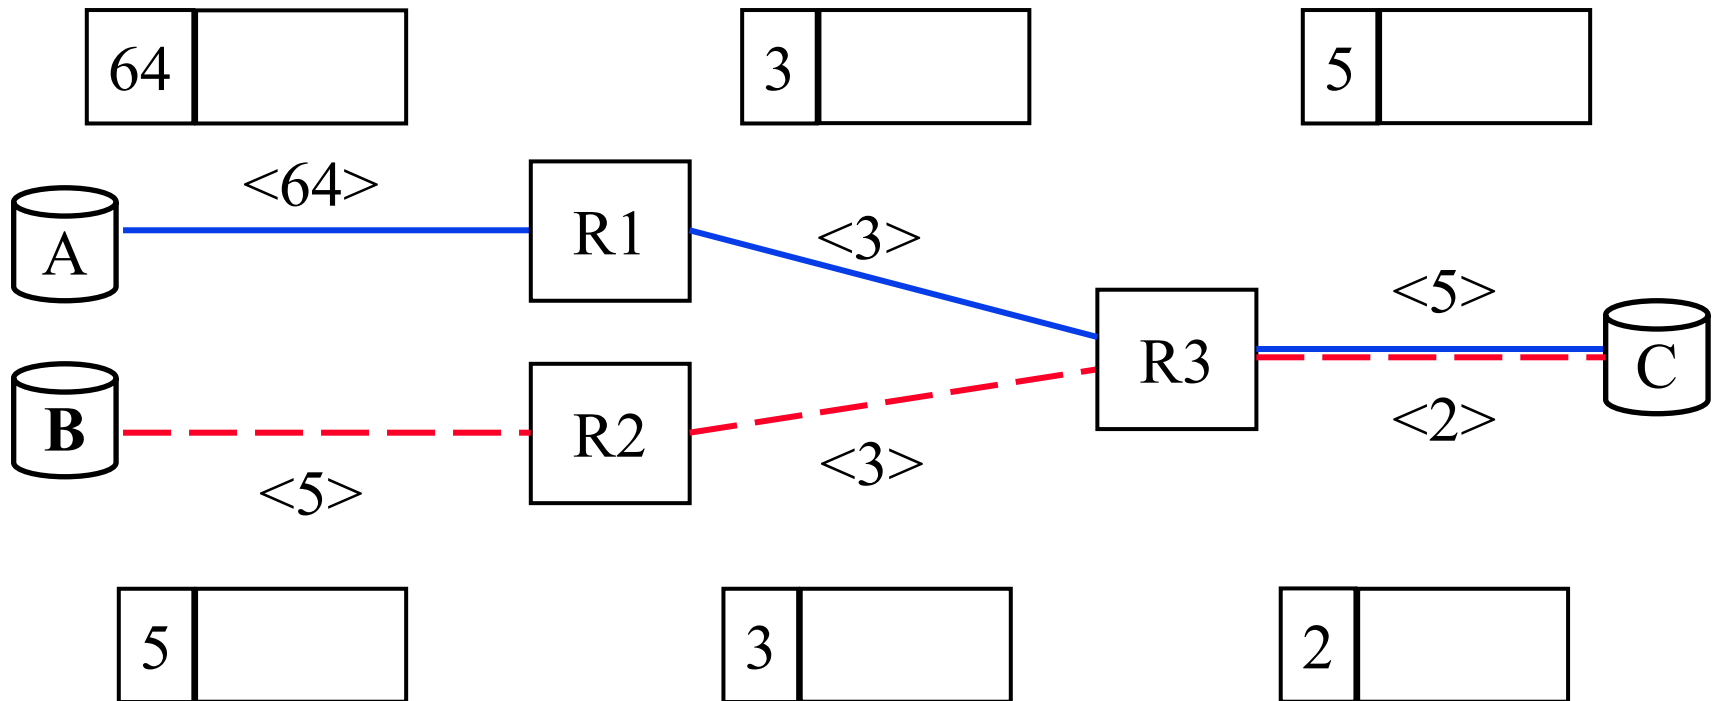

Multiprotocol Label Switching (MPLS)



- ❑ Allows virtual circuits in IP Networks (May 1996)
- ❑ Each packet has a **virtual circuit ID** called 'label'
- ❑ Label determines the packet's queuing and forwarding
- ❑ Circuits are called **Label Switched Paths (LSPs)**
- ❑ LSP's have to be set up before use
- ❑ **Label switching routers (LSRs)** allows traffic engineering



Label Switching Example



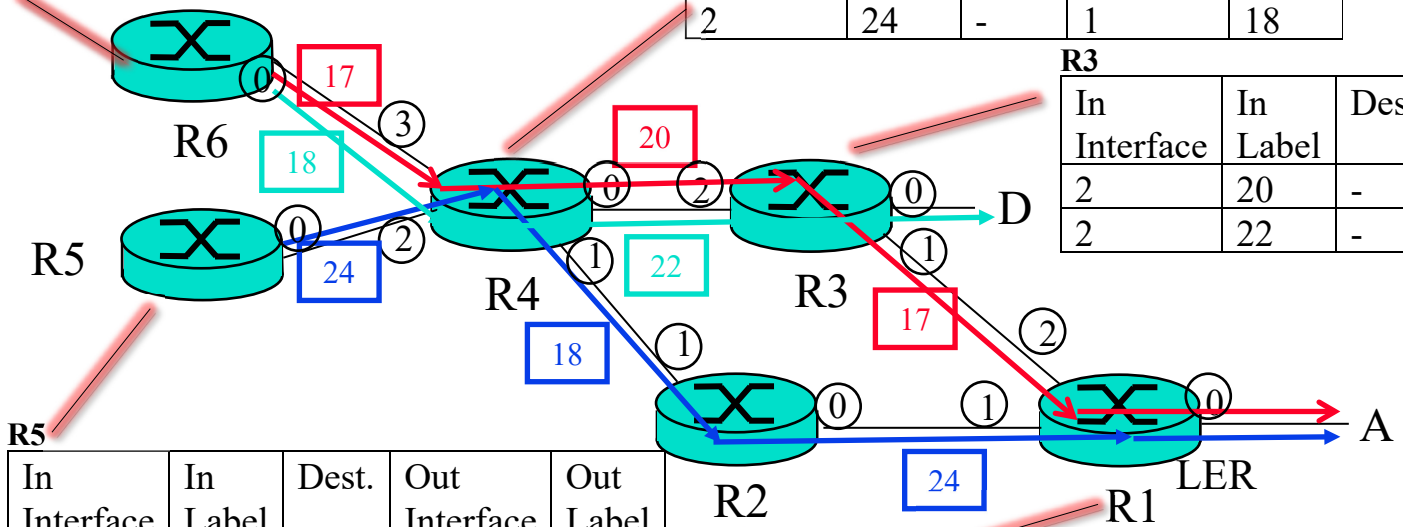
MPLS Forwarding Tables

□ Interface numbers are in circles. Label IDs are in rectangles.

In Interface	In Label	Dest.	Out Interface	Out Label
-	None	A	0	17
-	None	D	0	18

In Interface	In Label	Dest.	Out Interface	Out Label
3	17	-	0	20
3	18	-	0	22
2	24	-	1	18

In Interface	In Label	Dest.	Out Interface	Out Label
2	20	-	1	17
2	22	-	0	None



In Interface	In Label	Dest.	Out Interface	Out Label
-	None	A	0	24

In Interface	In Label	Dest.	Out Interface	Out Label
1	18	-	0	24

In Interface	In Label	Dest.	Out Interface	Out Label
2	17	-	0	None
1	24	-	0	None

Notes

❑ Error in the Book:

- The tables are per interface not per router.
- For compatibility, we have kept table per router but added the input interface column.
- The book lists no input interface in the table.
- Same label #s are allowed to be used in different interfaces of the same router. For example, See Router R3 in the “Label Switching Example” slide.
- The textbook notation will not allow this possibility.

❑ Only one direction of circuits is shown for clarity.

- There are equal number of reverse circuits that have their own labels not related to forward labels.


❑ Out Label=None ⇒ MPLS Tag is removed.

In Label=None ⇒ Packet arrives with no MPLS tag

MPLS Label Switched Paths (LSPs)

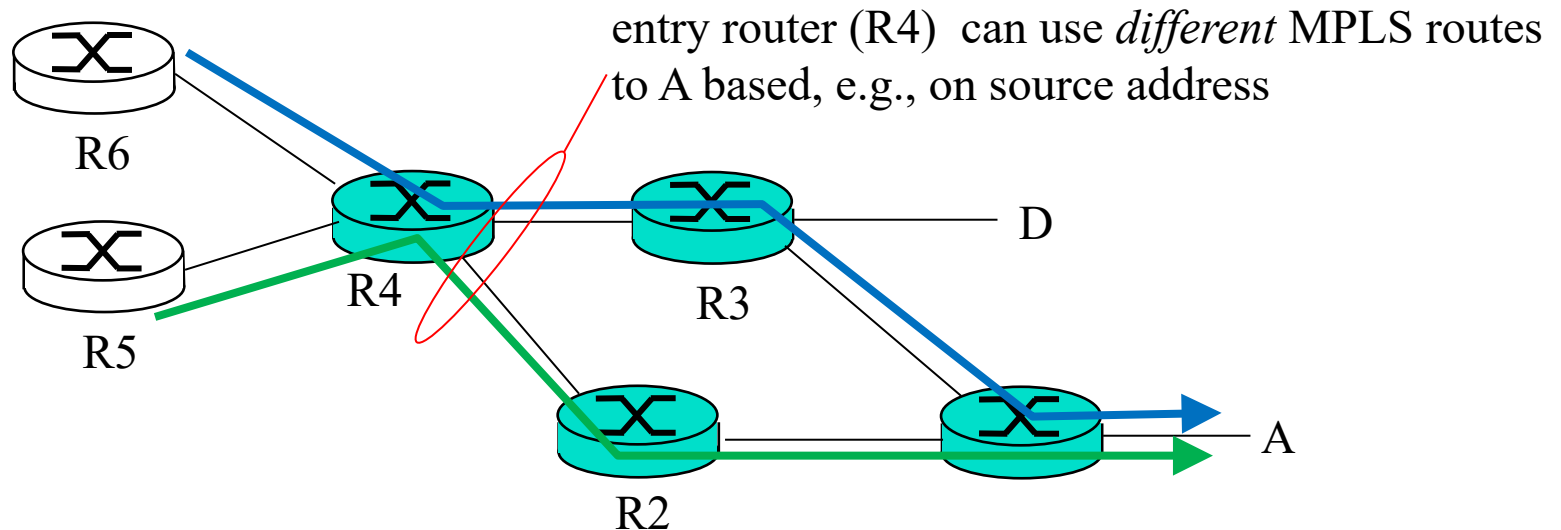
- ❑ Label switched paths (LSPs) are set up before use.
⇒ Connection oriented
- ❑ During set up each router tells the previous router what label it should put on the frames of that LSP.
- ❑ The label is actually an index in the MPLS forwarding table.
- ❑ Indexing in MPLS table is much faster than searching in IP tables.
- ❑ Although speed was one reason for using MPLS but the main reason is that the bandwidth can be reserved along the path.
- ❑ Labels are local. The same label ID may be used by different routers for different LSPs.
- ❑ The label ID changes along various links of the same LSP.
- ❑ Label IDs are 20-bit long ⇒ $2^{20}-1$ Labels.
Labels 0-15 are reserved.

Label Edge Routers (LERs)




- ❑ Routers connected to non-MPLS routers or nodes or routers of other MPLS domains are called Label Edge Routers (**LERs**)
- ❑ LERs add labels to frames coming from non-MPLS nodes or remove their labels if forwarding to non-MPLS nodes or other domains.
- ❑ The labels added by LERs **may be** based on destination address along with other considerations, such as source address, QoS, etc.
- ❑ Other LSRs forward based solely on the label and the interface the frame came in. They **do not** look at the destination address field.

MPLS versus IP Paths




- ❑ **IP Routing:** Path determined by destination address alone
- ❑ **MPLS Routing:** Path can be based on source and destination address, flow type, ...
 - **Fast reroute:** Precompute backup routes in case of link failure



MPLS Label Format

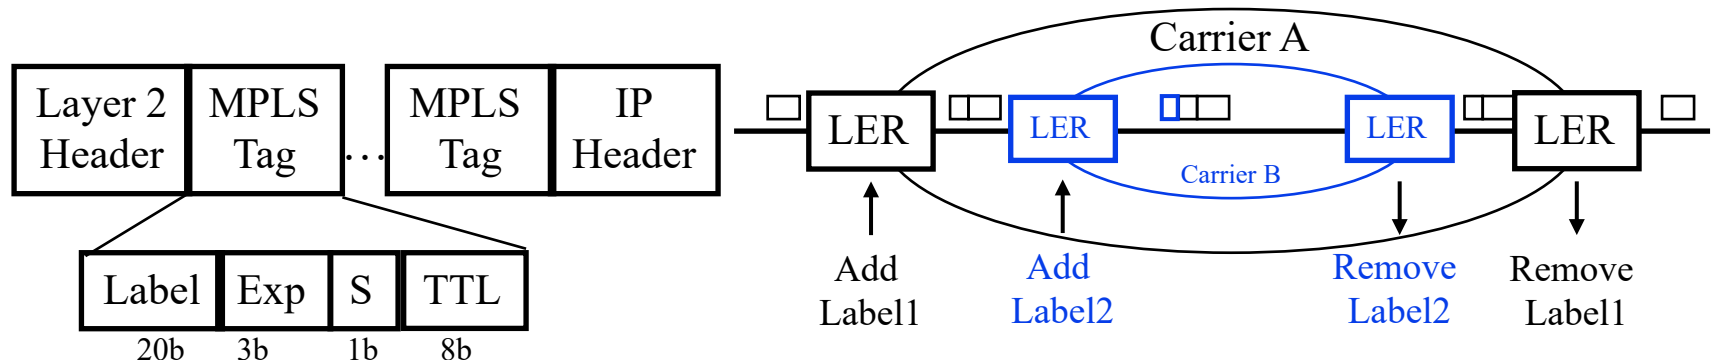
- ❑ MPLS label is inserted after layer 2 header but before layer 3 header \Rightarrow MPLS is **Layer 2.5**


- 20 bit label
- 3 bit Experimental: Class of Service
- 1 bit end-of-stack. A packet may have a stack of labels to allow carrier nesting.



- ❑ TTL field is decremented for all forwarded packets.
When adding label TTL field from IP header is copied to the MPLS tag.
When removing label TTL field from MPLS tag is copied to IP Header.

- ❑ MPLS Signaling:
 - OSPF has been extended to help prepare label tables
 - There are several other “*Label Distribution Protocols*”





MPLS: Review

1. Multiprotocol Label Switching (MPLS) allows virtual circuits called “**Label Switched Paths (LSPs)**” in IP
2. Each packet has a Layer 2.5 **MPLS tag** which includes a 20-bit label ID
3. Label switching routers (**LSRs**) forward based on input interface and the label
4. Label table is prepared by a “**Label Distribution Protocol.**” OSPF is one example of a LDP.
5. MPLS tags can be **stacked** to allow network nesting

Ref: Section 6.5

Homework 6C: MPLS

- [6 points] Consider the MPLS network shown in “MPLS Forwarding Tables” slide. Suppose that we want to perform traffic engineering so that packets from R6 destined for A are switched to A via R6-R4-R2-R1 and packets from R5 destined for A are switched via R5-R4-R3-R1. Show the updated MPLS table in R4 that would make this possible. For simplicity, use the same label values as shown currently. Only LSP paths change and the table at Router R4.

Homework 6C (Cont)

□ Interface numbers are in circles. Label IDs are in rectangles.

R6


In Interface	In Label	Dest.	Out Interface	Out Label
-	None	A	0	17
-	None	D	0	18

R4

In Interface	In Label	Dest.	Out Interface	Out Label

R3

In Interface	In Label	Dest.	Out Interface	Out Label
2	20	-	1	17
2	22	-	0	None



R5

In Interface	In Label	Dest.	Out Interface	Out Label
-	None	A	0	24

R2

In Interface	In Label	Dest.	Out Interface	Out Label
1	18	-	0	24

R1

In Interface	In Label	Dest.	Out Interface	Out Label
2	17	-	0	None
1	24	-	0	None

Homework 6D: MPLS

- [28 points] The next figure shows the flows on an MPLS network with the reverse direction flows. Using the Labels shown, fill in all the tables.

Homework 6D (Cont)

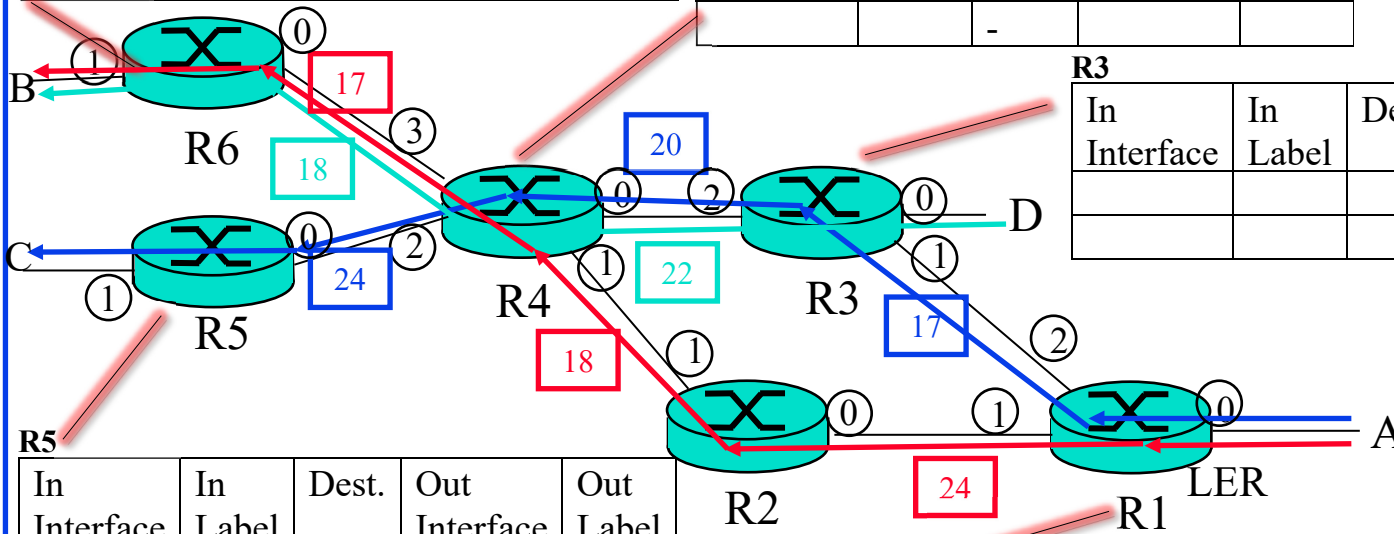
- Interface numbers are in circles. Label IDs are in rectangles.

R6

In Interface	In Label	Dest.	Out Interface	Out Label

R4

In Interface	In Label	Dest.	Out Interface	Out Label
		-		



R3

In Interface	In Label	Dest.	Out Interface	Out Label

R5

In Interface	In Label	Dest.	Out Interface	Out Label

R2

In Interface	In Label	Dest.	Out Interface	Out Label

R1

In Interface	In Label	Dest.	Out Interface	Out Label

Google's Data Center




Source: <http://webodyssey.com/technologyscience/visit-the-googles-data-centers/>
Washington University in St. Louis <http://www.cse.wustl.edu/~jain/cse473-20/>

©2020 Raj Jain

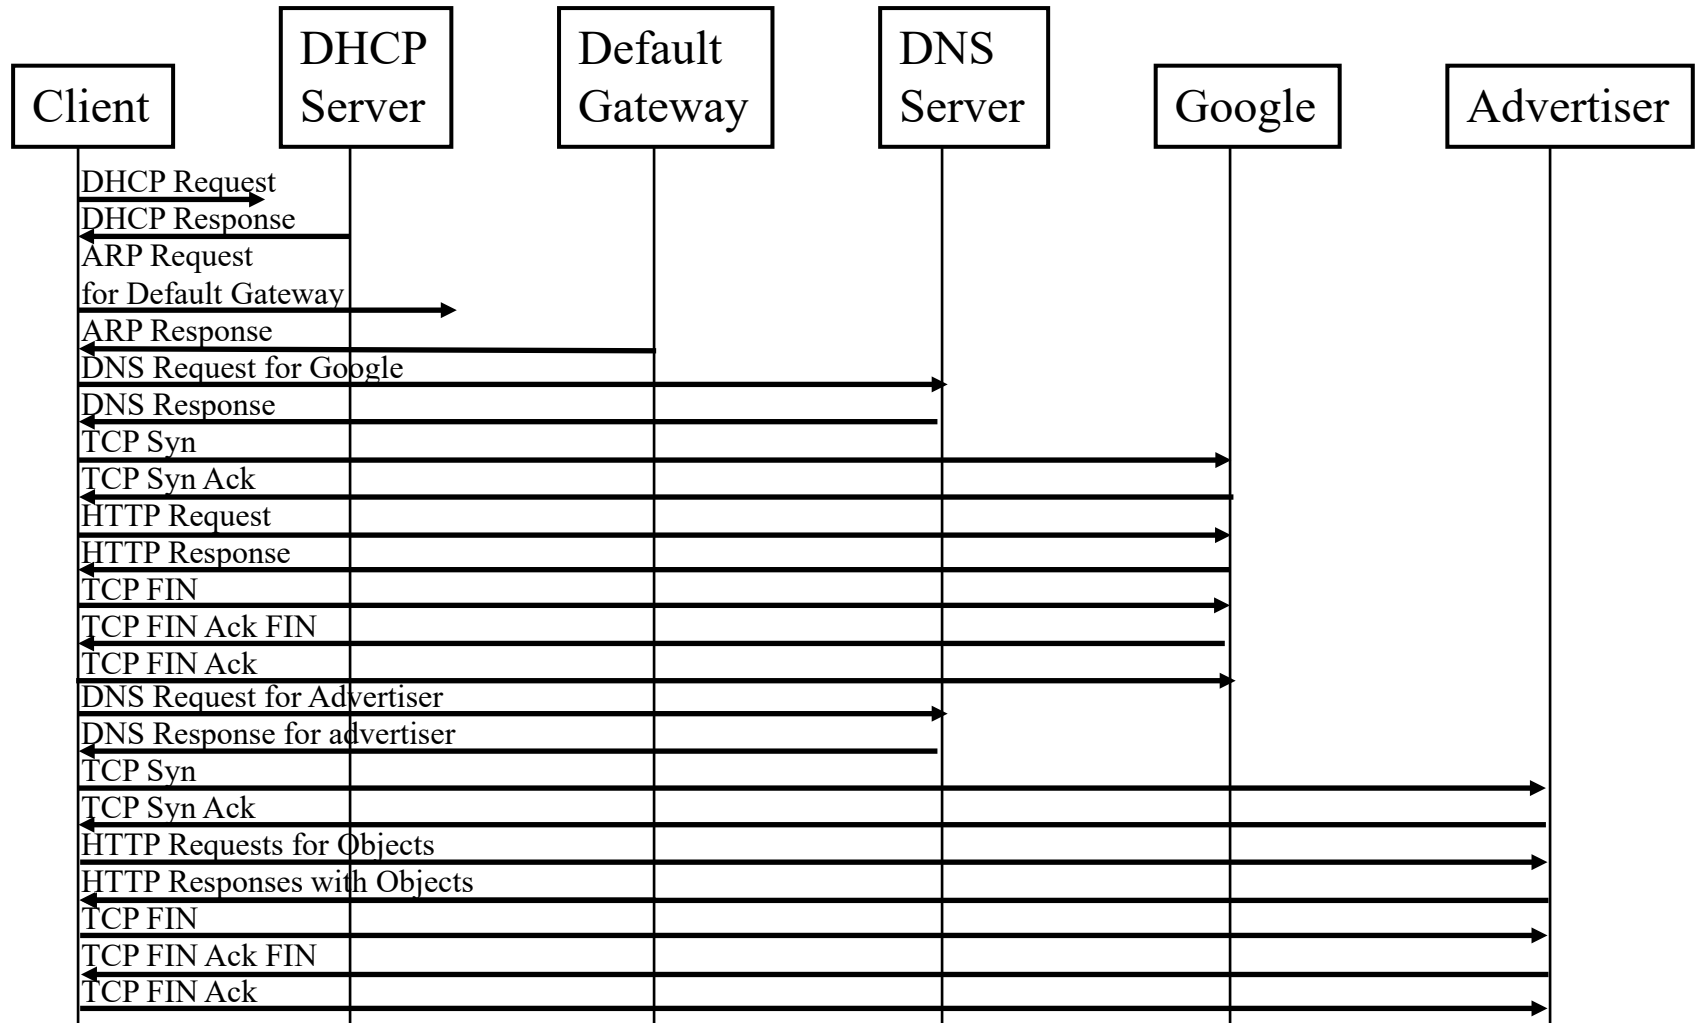
Data Center Networks Topology

- ❑ **3-Tier Architecture: Top-of-Rack, Aggregation, Core**
- ❑ **Middle boxes:** Load balancer, Firewall, Intrusion detection, ...
- ❑ Rich Interconnection between switches




Protocols: Complete Picture

Task: Connect and search in www.google.com



Summary



1. CRC uses **mod-2 division** using specially selected numbers
2. IEEE 802.3 uses a *truncated binary exponential backoff*.
3. Ethernet uses 48-bit universal addresses
4. Ethernet bridges are **transparent** and **self-learning**
5. **802.1Q** allows several **virtual LANs** inside a LAN.
6. Address Resolution Protocol (**ARP**) is used to find the MAC address for a given IP address and vice versa.
7. MPLS allows virtual circuits (**LSPs**) on IP networks.
8. Data centers use a **multi-tier switching** architecture with redundancy

Acronyms

- ❑ ARP Address Resolution Protocol
- ❑ ASCII American Standard Code for Information Exchange
- ❑ CAT Category
- ❑ CD Collision Detection
- ❑ CRC Cyclic Redundancy Check
- ❑ CSMA Carrier Sense Multiple Access
- ❑ DA Destination Address
- ❑ DEI Drop Eligibility Indicator
- ❑ DHCP Dynamic Host Control Protocol
- ❑ DNS Domain Name Server
- ❑ DOCSIS Data over Cable Service Interface Specification
- ❑ FDMA Frequency Division Multiple Access
- ❑ HTTP Hypertext Transfer Protocol
- ❑ ID Identifier
- ❑ IEEE Institution of Electrical and Electronic Engineers

Acronyms (Cont)

- ❑ IP Internet Protocol
- ❑ IPX Internetwork Packet Exchange
- ❑ LAN Local Area Network
- ❑ LDP Label Distribution Protocol
- ❑ LLC Logical Link Control
- ❑ LSP Label Switched Path
- ❑ MAC Media Access Control
- ❑ MAP Map
- ❑ MPLS Multiprotocol Label Switching
- ❑ MSB Most Significant Byte First
- ❑ NIC Network Interface Card
- ❑ OSPF Open Shortest Path First
- ❑ OUI Organizationally Unique Identifier
- ❑ PBX Private Branch Exchange
- ❑ PCP Priority Code Point
- ❑ PHY Physical Layer

Acronyms (Cont)

- ❑ SA Source Address
- ❑ STP Shielded Twisted Pair
- ❑ TCP Transmission Control Protocol
- ❑ TDMA Time Division Multiple Access
- ❑ TOR Top of the Rack
- ❑ TPI Tag Protocol Identifier
- ❑ TTL Time to live
- ❑ TX Transmit
- ❑ UTP Unshielded Twisted Pair
- ❑ VLAN Virtual Local Area Network

Scan This to Download These Slides



Raj Jain

<http://rajjain.com>

http://www.cse.wustl.edu/~jain/cse473-20/i_6lan.htm

Related Modules



CSE 567: The Art of Computer Systems Performance Analysis

https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),

https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcg5e_10TiDw



CSE 570: Recent Advances in Networking (Spring 2013)

<https://www.youtube.com/playlist?list=PLjGG94etKypLHyBN8mOgwJLHD2FFIMGq5>

CSE571S: Network Security (Spring 2011),

<https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u>



Video Podcasts of Prof. Raj Jain's Lectures,

<https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw>