Lecture 17 (End-to-End 1)

Ethernet, STP

CS 168, Fall 2026 @ UC Berkeley

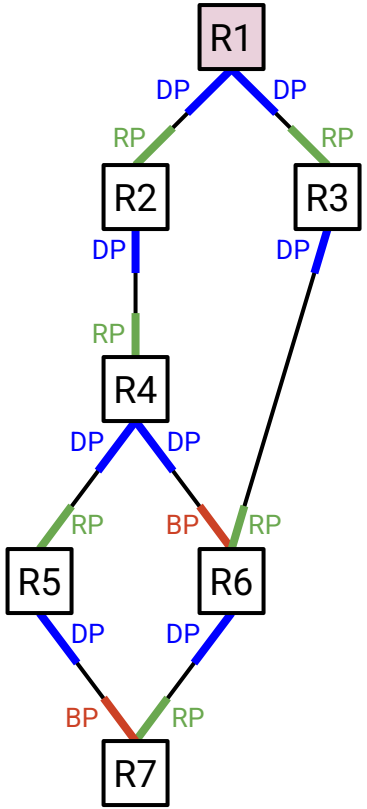

Today: A closer look at Ethernet

Next lecture: end-to-end operation and some important missing pieces (ARP, DHCP)

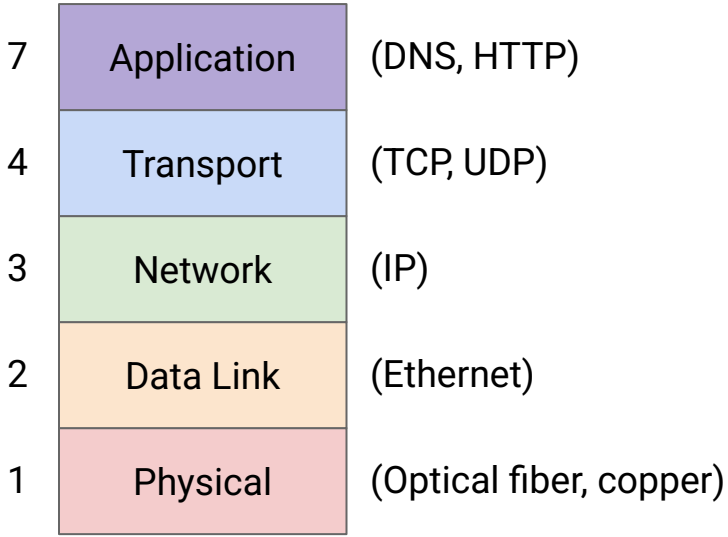

Ethernet: a brief timeline

1970s: Ethernet as a “shared bus” network connecting hosts in a local area network (LAN)

- Key ideas: carrier sense, collision detection, *etc.*

~1990: Ethernet evolves to interconnect individual shared-bus networks via switches

- Key ideas: learning switches, spanning tree protocol

Today: Ethernet has largely evolved to point-to-point links interconnected by switches

- Routing via L3 or L3-like protocols

These ideas persist! Example: in the design of WiFi, zero-config “plug-and-play” systems

Connecting Local Hosts

Lecture 17, CS 168, Fall 2026

Ethernet

- **Connecting Local Hosts**
- Multiple Access Protocols
- Sending Packets
- Scaling Ethernet

Ethernet Routing

- Naive Approach: Flooding
- Learning Switches
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Connecting Local Hosts

So far, we've assumed that every link connects exactly two machines:

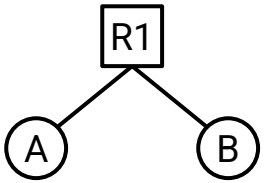

But some “shared media” links can connect multiple machines:

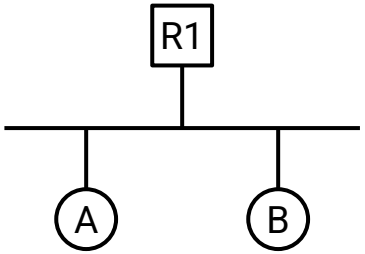

From such a link’s perspective, the router is just a machine like any other:

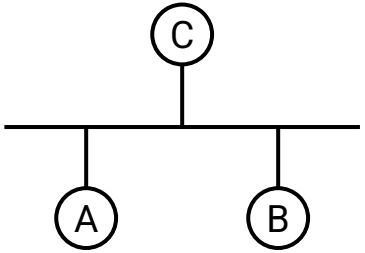

Connecting Local Hosts

So far, we've assumed that every link connects exactly two machines:

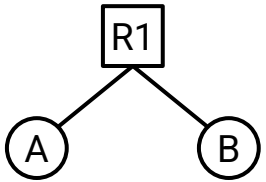

But some links can connect multiple machines:

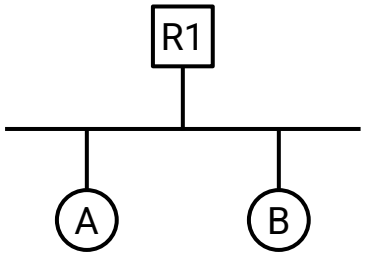

Connecting Local Hosts

How could we connect hosts in a local network?

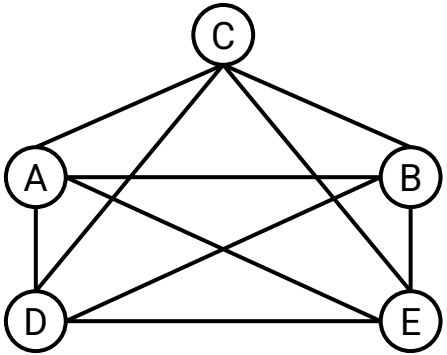

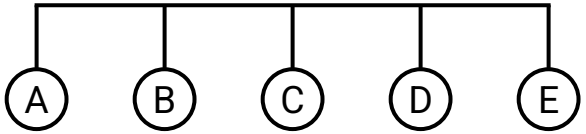

Mesh: Link between every pair of machines.

- For every new host, we have to add new links to every other host.
- Need a lot of (physical) ports per host.

Bus: A single link for all machines.

- Introduces a **shared media**.

Shared Media

Shared media: Many machines using the same link.

- If multiple machines transmit at the same time, signals will *interfere* or *collide*.
- Analogy: People talking simultaneously on a group call.

Note: Shared media is not necessarily a wire.

- Could be radio waves on a shared wireless link.

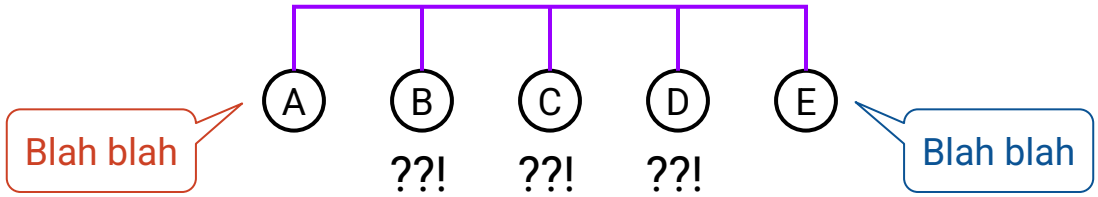

Multiple Access Protocols

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- **Multiple Access Protocols**
- Sending Packets
- Scaling Ethernet

Ethernet Routing

- Naive Approach: Flooding
- Learning Switches
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Multiple Access Protocols

A **multiple access protocol** allocates the shared media to everyone wanting to use it.

- 3 types of approaches.
- Several protocols of each type.

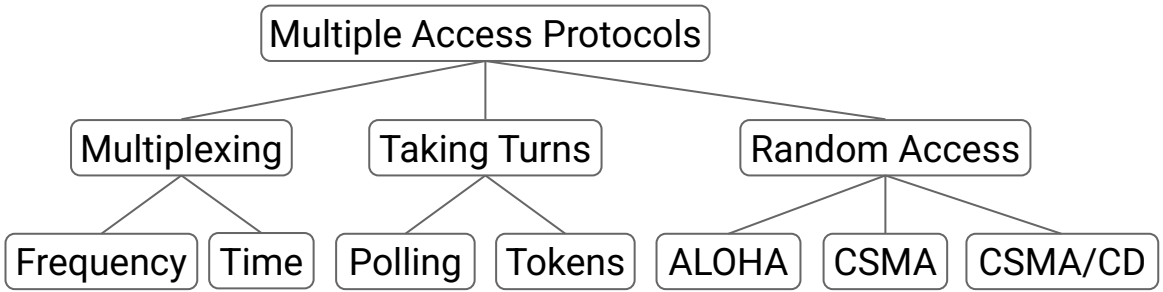

Multiple Access Protocols (1/3) – Multiplexing

Idea: Allocate a fixed slice of resources to each node.

- **Frequency-based** multiplexing: Divide medium into frequency channels.
- **Time-based** multiplexing: Give each node some fixed time slots.

Problem: Can be wasteful.

- There's only so much frequency/time available to allocate.
- If a node has nothing to say, some frequency/time goes unused.

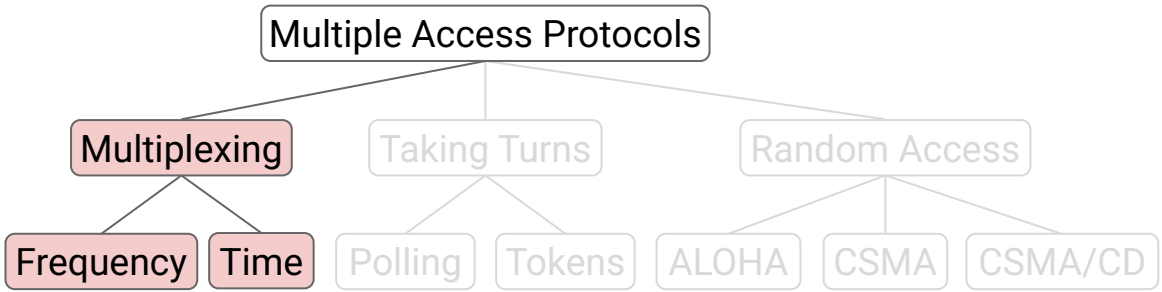

Multiple Access Protocols (2/3) – Taking Turns

Idea: Nodes take turns speaking.

- **Polling protocols**: A coordinator decides when each node can speak
  - Example: I2C
- **Token passing**: Pass a virtual token around. Only the node with the token can speak
  - Example: IBM Token Ring, FDDI

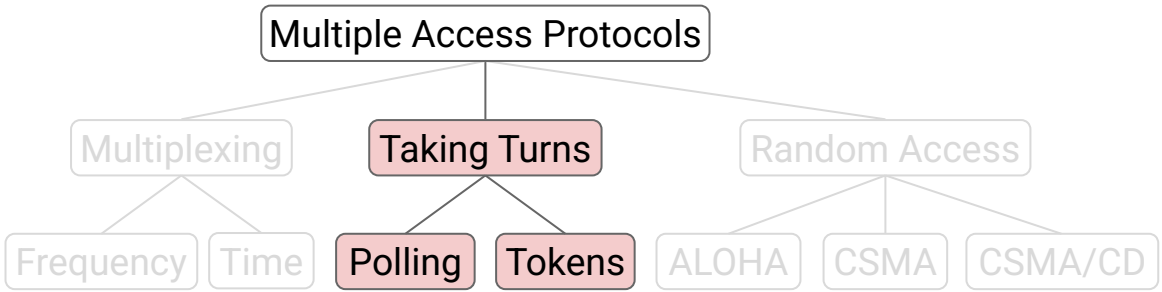

Multiple Access Protocols (2/3) – Taking Turns

Polling protocols: A coordinator decides when each node can speak.

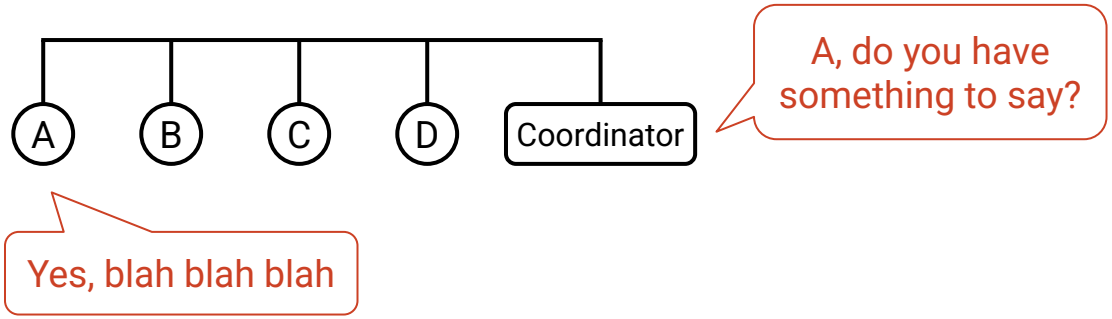

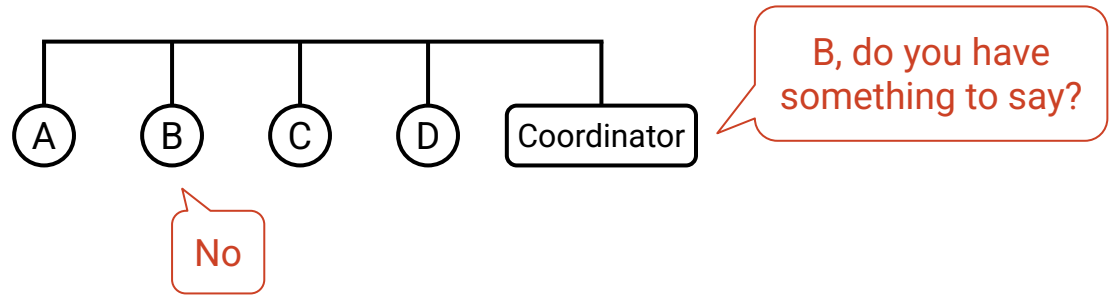

B has nothing to say. Coordinator can immediately move on to C.

Multiple Access Protocols (2/3) – Taking Turns

Token passing: Pass a virtual token around. Only the node with the token can speak.

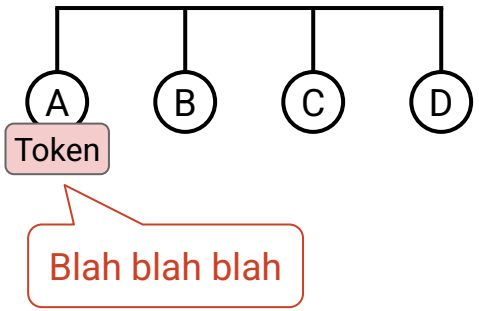

A holds the token.
A can speak.
When A is done, it passes the token to B.

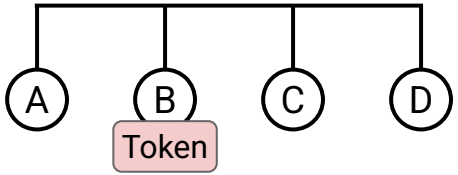

B has nothing to say.
It can immediately pass the token to C.

Multiple Access Protocols (2/3) – Taking Turns

Idea: Nodes take turns speaking.

- Benefit: No more time wasted on idling.
  - If someone has nothing to say, immediately move on.
- Problem: Complexity.
  - Need to implement inter-node communication.
  - How do we elect the central coordinator?
  - What if two nodes both think they have the token?

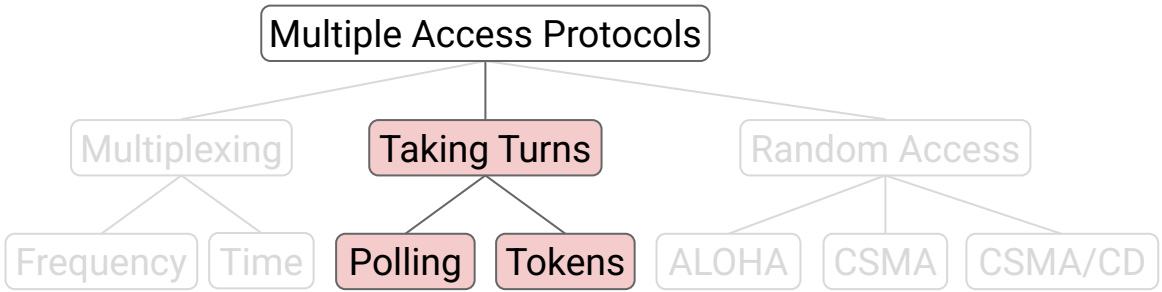

Multiple Access Protocols (3/3) – Random Access

Idea: Nodes talk whenever they have something to say.

- Deal with collisions when they occur.
- Benefit: Simplicity. No coordinators or tokens.

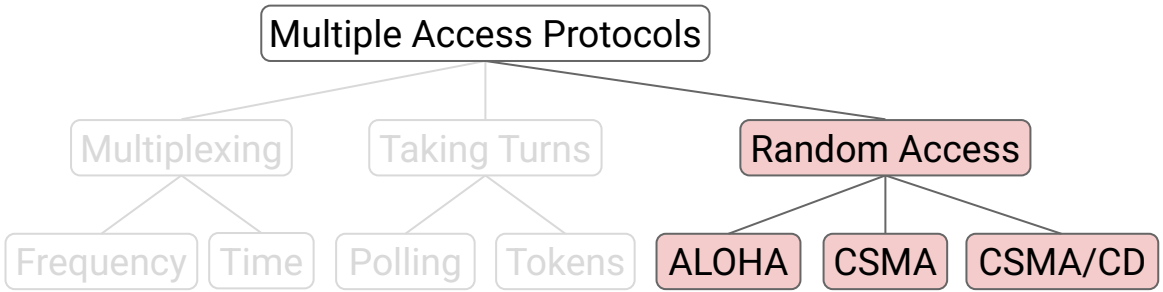

Multiple Access Protocols (3/3) – Random Access – ALOHA

ALOHA random access scheme: "rude" version.

- If you have a packet, just send it.
  - Recipient replies with an ack.
- If two nodes send simultaneously, collision corrupts the packets.
  - No ack!
- If you don't get an ack: Wait some random amount of time, then resend.
  - Randomness helps avoid another collision.

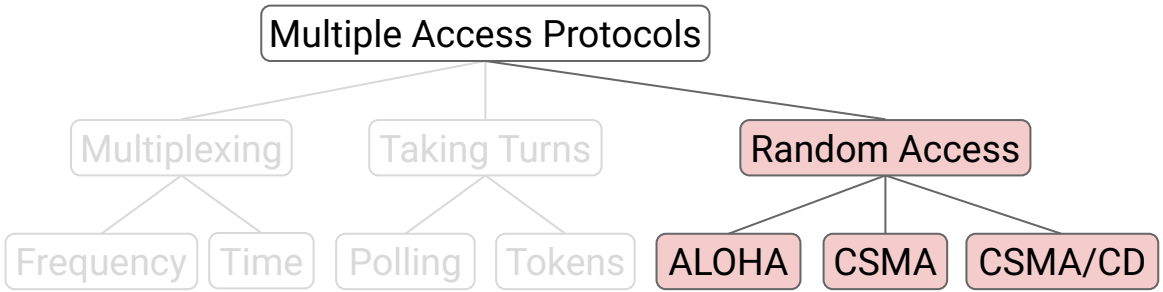

Multiple Access Protocols (3/3) – Random Access – CSMA

CSMA (Carrier Sense Multiple Access): "polite" version.

- First, listen to see if anyone is sending.
- Only start sending when it's quiet.

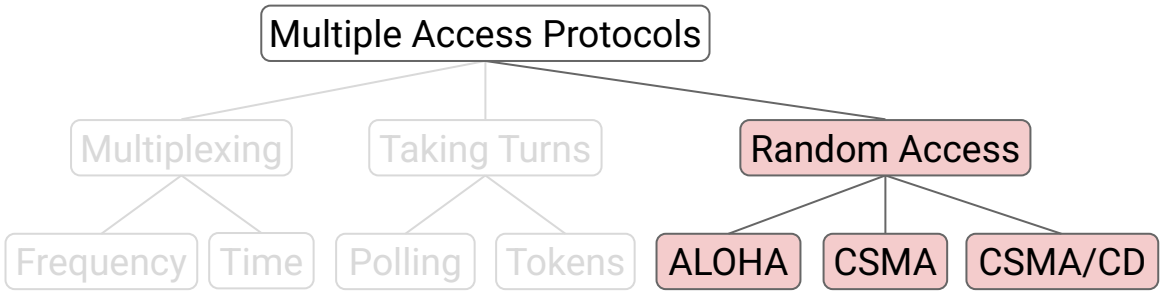

Multiple Access Protocols (3/3) – Random Access – CSMA and Propagation Delay

CSMA does not necessarily avoid collisions, because of **propagation delay**.

- t=0: B starts sending.
  - Signal takes time to reach A, C, D.
- t=2: D wants to send.
  - Signal hasn't reached D yet!
  - D thinks it's quiet and starts sending.
- Result: Collision!

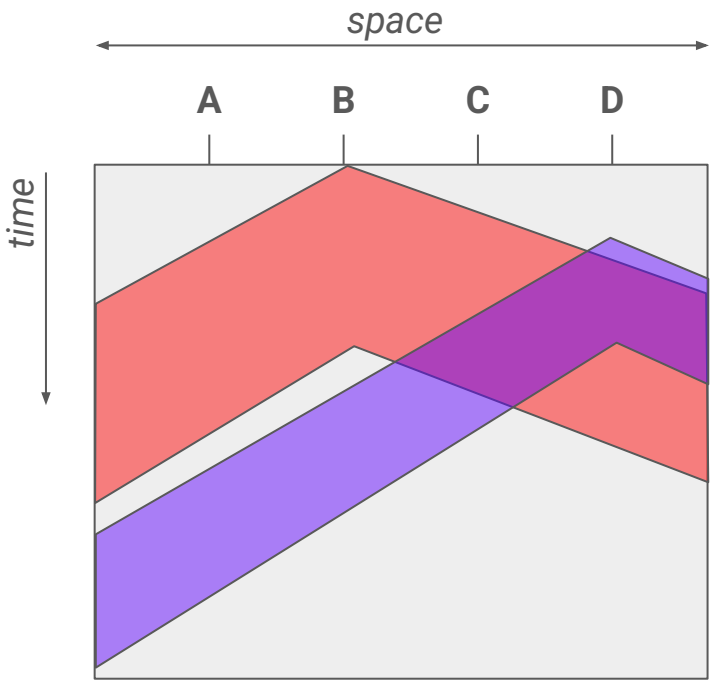

Multiple Access Protocols (3/3) – Random Access – CSMA/CD

CSMA/CD (Carrier Sense Multiple Access with Collision Detection):

- Listen before sending, but also *while* you're sending.
- If you hear someone else sending, stop! Collision detected.

CSMA/CD uses **binary exponential backoff**:

- After every collision, wait up to twice as long before resending.
  - After first collision: Wait between 0–4 seconds.
  - If resend collides again: Wait between 0–8 seconds.
- Resends fast when possible, slowing down when necessary (e.g. many senders).

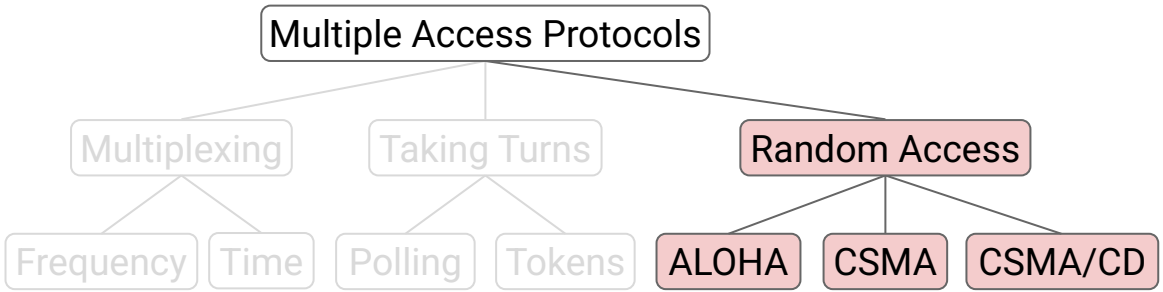

Brief History of Multiple Access Protocols

Why was that protocol called ALOHA, anyway?

- Norman Abramson had a problem at the University of Hawaii in 1968.
- How do we allow users on other islands to access the central hub computer?

Solution: ALOHANet (*Additive Links Online Hawaii Area*). Hugely influential.

- Shared wireless medium between all the computers.

ALOHANet used two frequencies:

- Hub transmits on its own frequency.
  - One sender = no collisions.
  - All remote nodes listen to this frequency.
- All remote sites transmit on the other frequency.
  - May collide. Use ALOHA to deal with collisions.
  - Only the hub listens to this frequency.

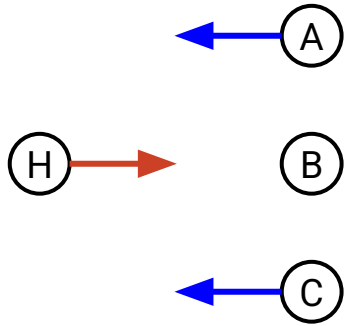

Sending Ethernet Packets

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- **Sending Packets**
- Scaling Ethernet

Ethernet Routing

- Naive Approach: Flooding
- Learning Switches
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Ethernet as LAN Network Protocol

Local Area Networks (LANs) are generally Ethernet.

Machines in the same LAN can exchange messages directly at Layer 2!

- No need for IP headers, routers, forwarding, etc.
- Analogy: If we're in the same room, we can talk without using the postal system.

Ethernet Addressing

At layer 2, each machine has a **MAC address** (*Media Access Control*).

- Can be used even if you don't have an IP address.
- 48 bits, written in hex, e.g. `f8:ff:c2:2b:36:16`.
- Stored permanently on the machine ("burned in").
  - Can be overridden by software.
- Allocated by the manufacturer of the machine.
- Globally unique. You might plug your computer in anywhere.

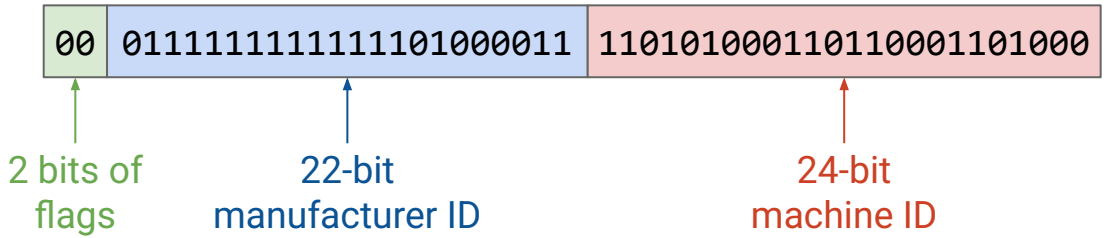

Types of LAN Communication

Ethernet supports three types of communication:

- **Unicast**: Send a packet to a single recipient.
- **Broadcast**: Send a packet to everyone on the local network.
- **Multicast**: Send a packet to everyone in a specific group.
    - Machines in the local network can join groups.

Types of LAN Communication (1/3) – Unicast

Unicast: Send a packet to a single recipient.

- Destination = the recipient's MAC address.

Recall: On a shared medium, everybody gets the signal.

- When you get a packet, check the destination to see if it's meant for you.
- If not, ignore the packet.

Types of LAN Communication (2/3) – Broadcast

Broadcast: Send a packet to everyone on the local network.

- Packet already reaches everybody on the shared medium.
  - We just need to make sure everyone knows it's meant for them.
- Destination = the broadcast address, `FF:FF:FF:FF:FF:FF`.

Types of LAN Communication (3/3) – Multicast

Multicast: Send a packet to everyone in a specific group.

- Again, everyone gets the packet – need to ensure the group knows it's for them.
- Each group has a *group address*.
  - Individual computer's address: First bit 0.
  - Group address: First bit 1.
- Broadcast is just a special case of multicast, where everyone is in a group.

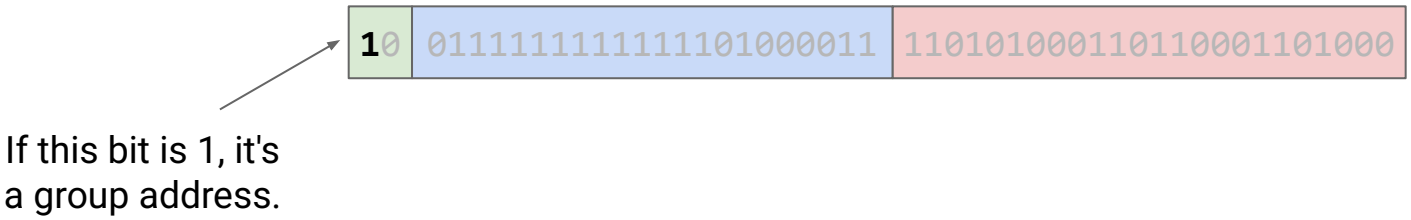

Types of LAN Communication (3/3) – Multicast

Example of using multicast: Bonjour/mDNS.

- Invented by Apple.
- All Apple products join a multicast group.
- Apple devices can multicast to the group.
  - iPhone says: "Hey, local Apple products, can anyone play music?"
  - Apple TV says: "Hey, local Apple products, I'm an Apple TV."

Ethernet Packet Structure

A data packet in Ethernet is often called a *frame*.

- Many fields (destination, source, type, checksum) similar to IP header.
- Need additional fields to separate packets on the wire.

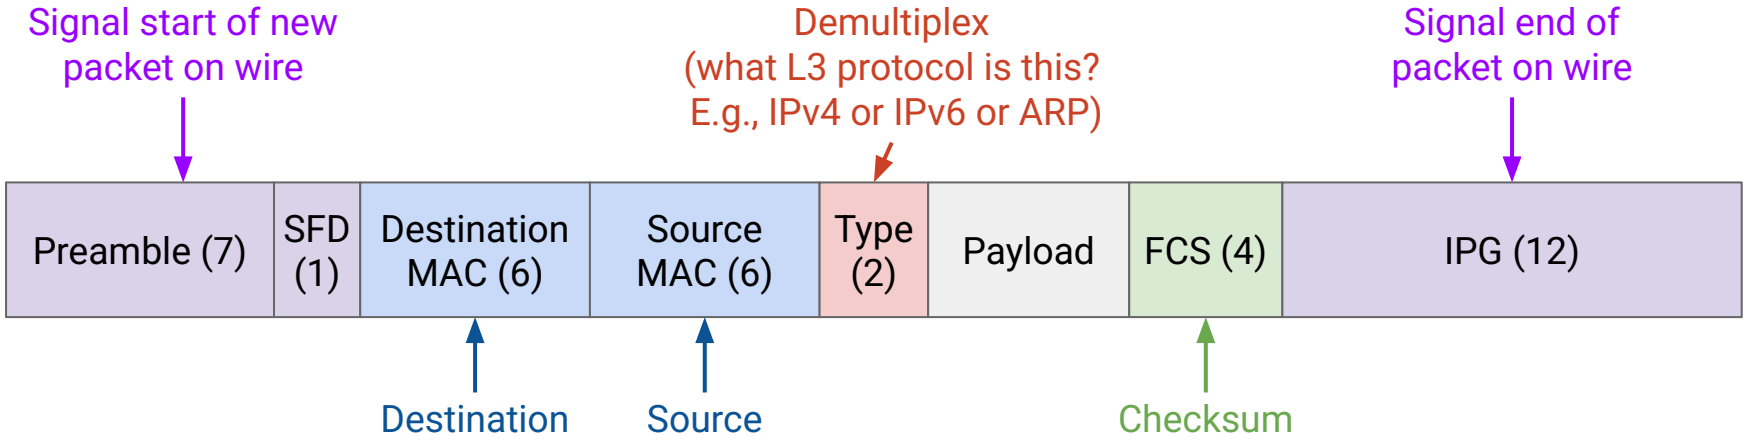

Characteristics of (Shared-Media) Ethernet

A few noteworthy characteristics *(and differences vs. IP)*

- Flat addresses
  - *vs. hierarchical addresses*
- No/minimal control plane
  - *vs. routers/switches, routing protocols, etc.*
- Zero configuration – hosts “just plug in and send”
  - *vs. default/static route config*
  - *vs. careful IP address assignment (DHCP - next lecture)*
  - *vs. waiting for a routing protocol to converge*
- Easily supports unicast, broadcast, multicast
  - *vs. primarily unicast*

Layer 2 Networks

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- Sending Packets
- **Scaling Ethernet**

Layer 2 Routing

- Naive Approach: Flooding
- Learning Switches
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Scaling Ethernet

A single shared link is limited in how many hosts its can scale to

- Number of collisions increases with more hosts and longer propagation delay
- As a result, practice evolved to have multiple independent shared media networks

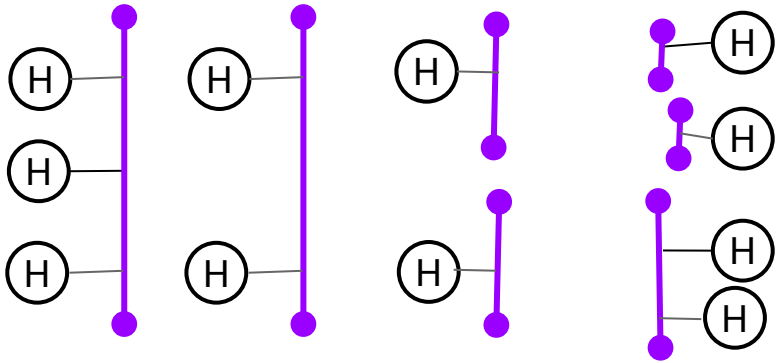

Scaling Ethernet

A single shared link is limited in how many hosts its can scale to

- Number of collisions increases with more hosts and longer propagation delay
- As a result, practice evolved to have multiple independent shared media networks

Is there a better way to scale?

- Idea (~1990): connect individual shared-bus networks with switches

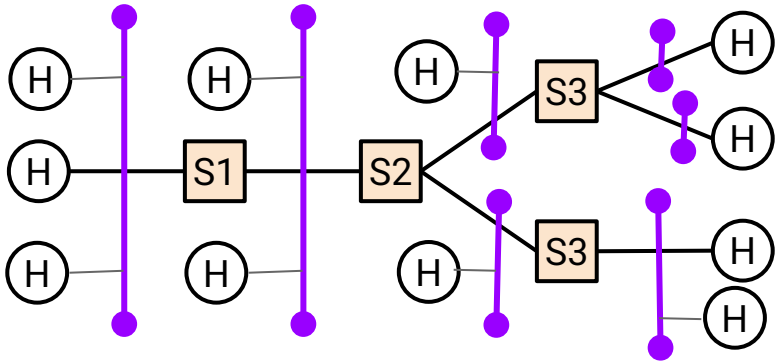

Routing in Switched Ethernet

Challenge: maintain backward compatibility with Ethernet and its characteristics

- Flat addresses, zero-configuration, simple/minimal control plane, *etc.*

We could use more than one wire in a local network.

Switches need to forward packets toward their destination.

- If a switch receives a broadcast packet: Send it to every neighbor.
- Multicast is more complicated. Need to know who's in the group.

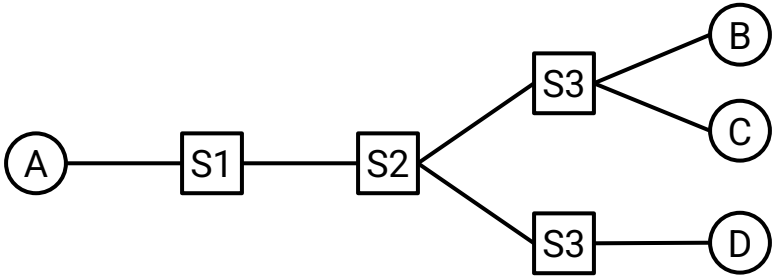

Layer 2 Networks

Routing protocols from Layer 3 can also be used at Layer 2.

- Destinations are MAC addresses, instead of IP addresses.

Problem: MAC addresses can't be aggregated.

- Allocated by manufacturer, not geographically.
- This is why Layer 2 can't scale to the Internet.

S2's Table	
Destination	Next Hop
A	R1
B	R3
C	R3
D	R4

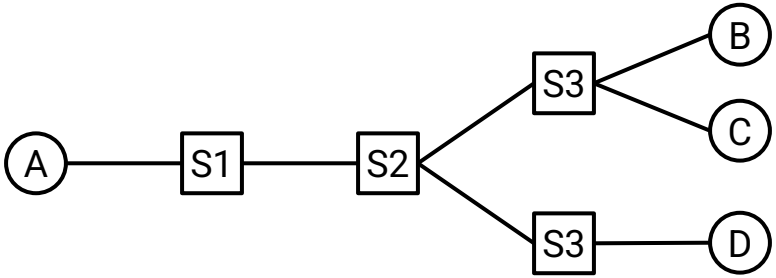

Naive Approach: Flooding

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- Sending Packets
- Scaling Ethernet

Ethernet Routing

- **Naive Approach: Flooding**
- Learning Switches
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Note: Simplification in the following slides

Instead of this figure …

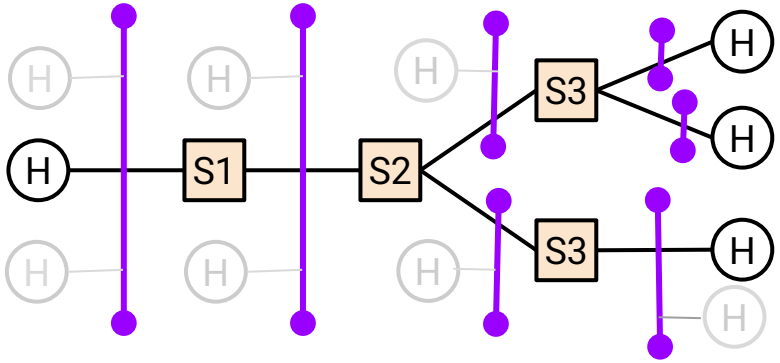

We'll use this simplified one:

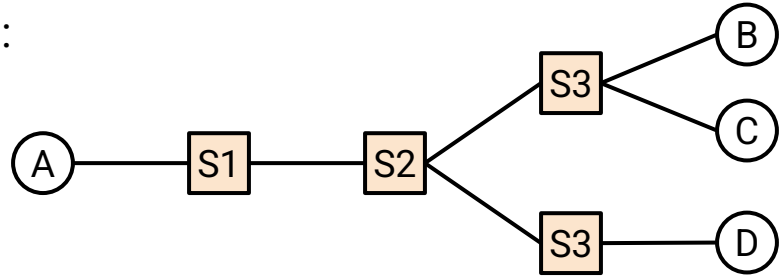

Recall: when S1 sends to S2, any hosts on the shared link between S1-S2 will hear the message

Naive Approach: Flooding

When you receive a packet, send it out of all ports, except the incoming port.

Will this work?

Yes. The packet will reach everyone, including the destination.

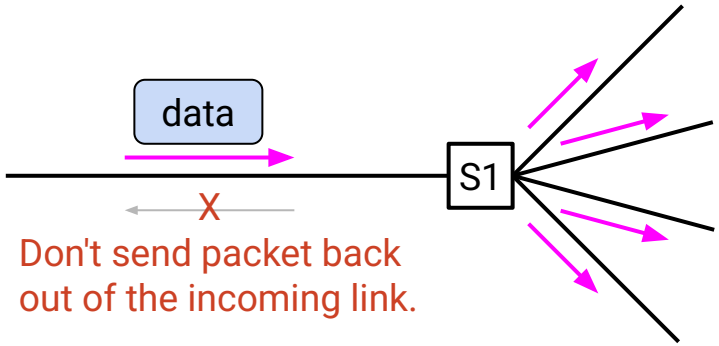

Naive Approach: Flooding

Problems with flooding? Wastes bandwidth in two ways:

1. Packet is sent to other hosts (not the destination).
2. Loops create broadcast storms.

Fixing these 2 problems will result in a good Ethernet routing protocol.

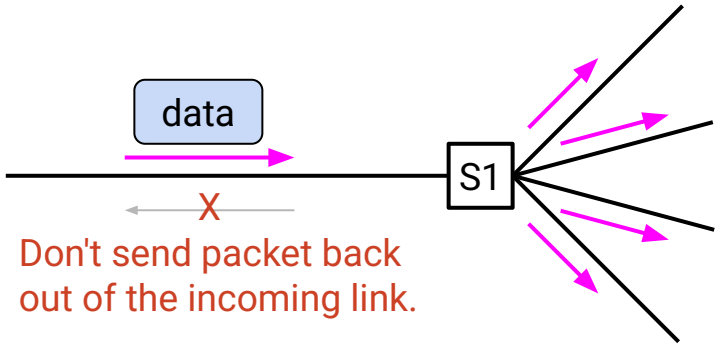

Learning Switches

Lecture 16, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- Sending Packets
- Scaling Ethernet

Layer 2 Ethernet Routing

- Naive Approach: Flooding
- **Learning Switches**
- Spanning Tree Protocol (STP)
- Implementing STP: BPDU Exchanges

Learning Switches: Motivation

Problems with flooding? Wastes bandwidth in two ways:

1. Packet is sent to other hosts (not the destination).
    - Solution: Use learning switches to build forwarding tables.
2. Loops create broadcast storms.

Fixing these 2 problems will result in a good Ethernet routing protocol.

Learning Switches: Routing Process

Key idea: When you receive a **data** packet, you get a clue about where the **sender** is.

- If a packet from host A comes from the west...
- ...then host A must be to my west!

Use this idea to add a table entry for the **sender**.

- As more packets get sent, more table entries get added.
- Note: The packet only helps you learn about the sender, not the receiver.

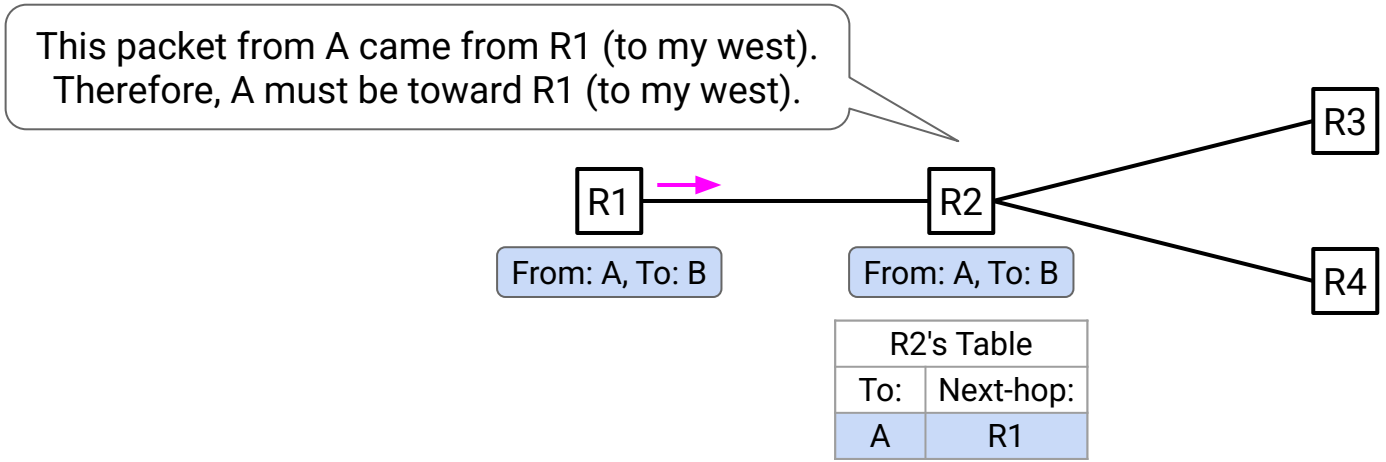

Learning Switches: Forwarding Process

Forwarding is **destination-based**: Use the destination to decide the next-hop.

- If the destination exists in the table: Forward to corresponding next-hop.
- If the destination is not in the table: Flood out of all ports (except incoming port).

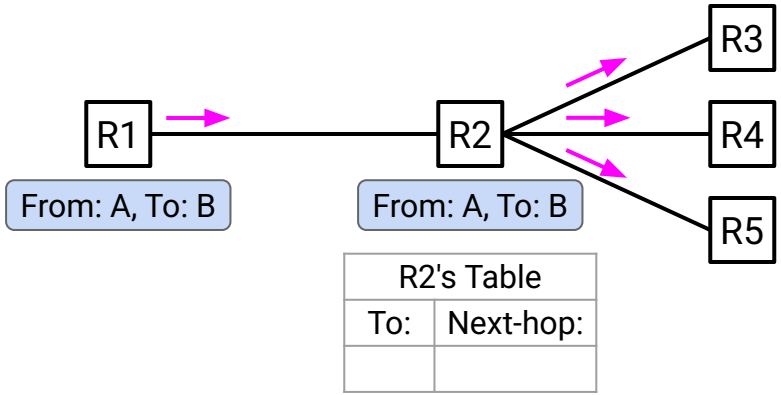

Case 1: No entry for B (destination) in table.

Flood the packet to all ports (except incoming port).

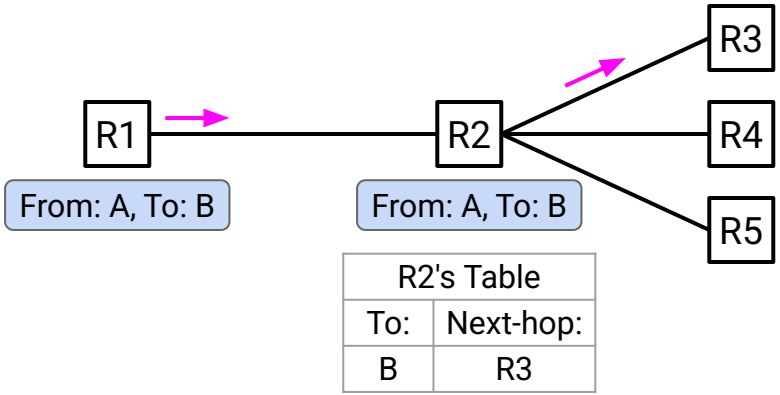

Case 2: Entry for B (destination) is in table.

Use table entry to forward to next-hop.

Learning Switches: Example

All tables start empty. Nobody has an entry for B, so the packet gets flooded.

Along the way, everyone gets a clue about A's location, and adds a table entry for A.

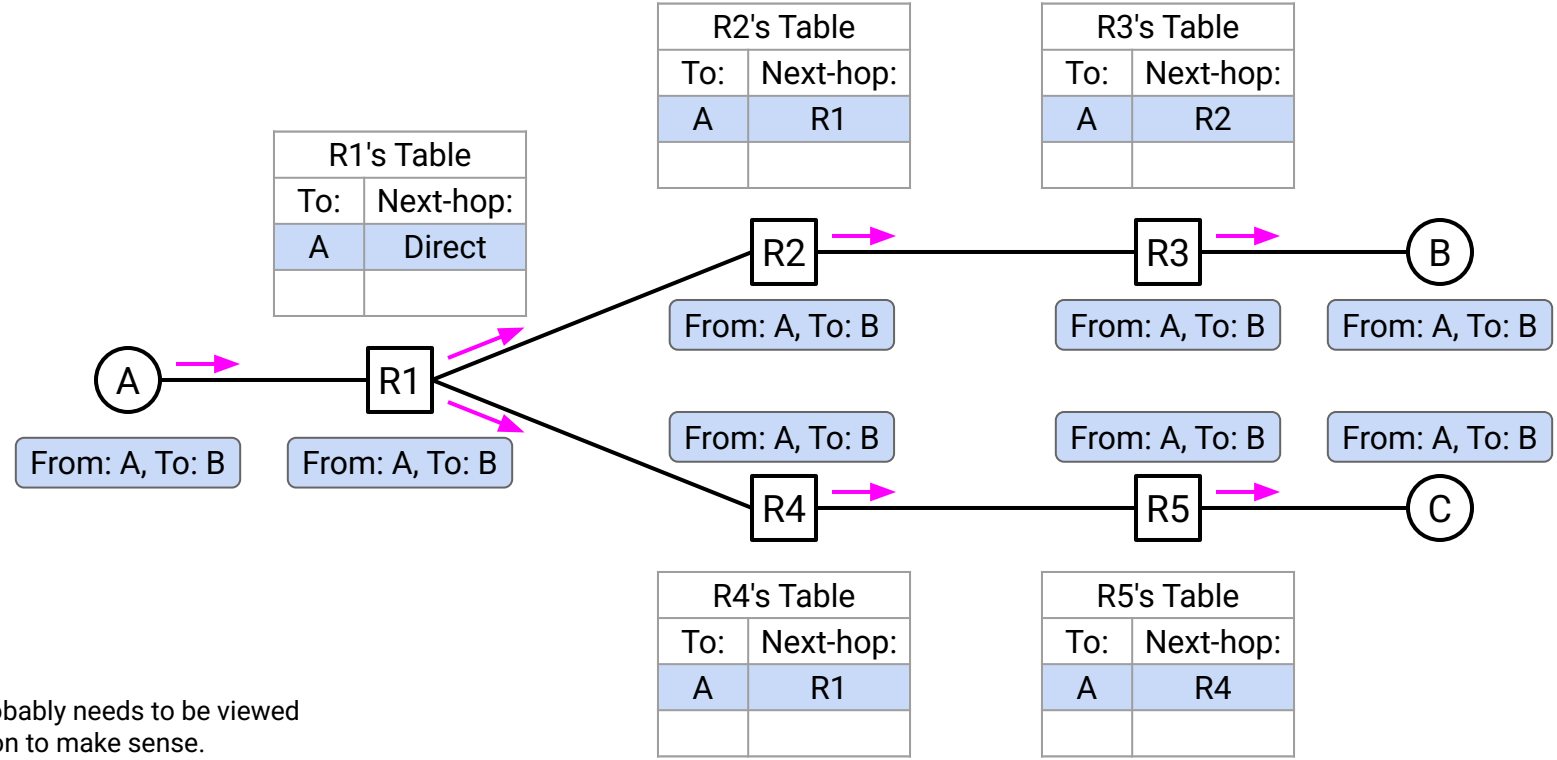

This slide probably needs to be viewed with animation to make sense.

Learning Switches: Example

This packet can get directly forwarded to A (no flooding) using the forwarding tables.

Switches along the B-to-A path get a clue about B's location, and add a table entry for B.

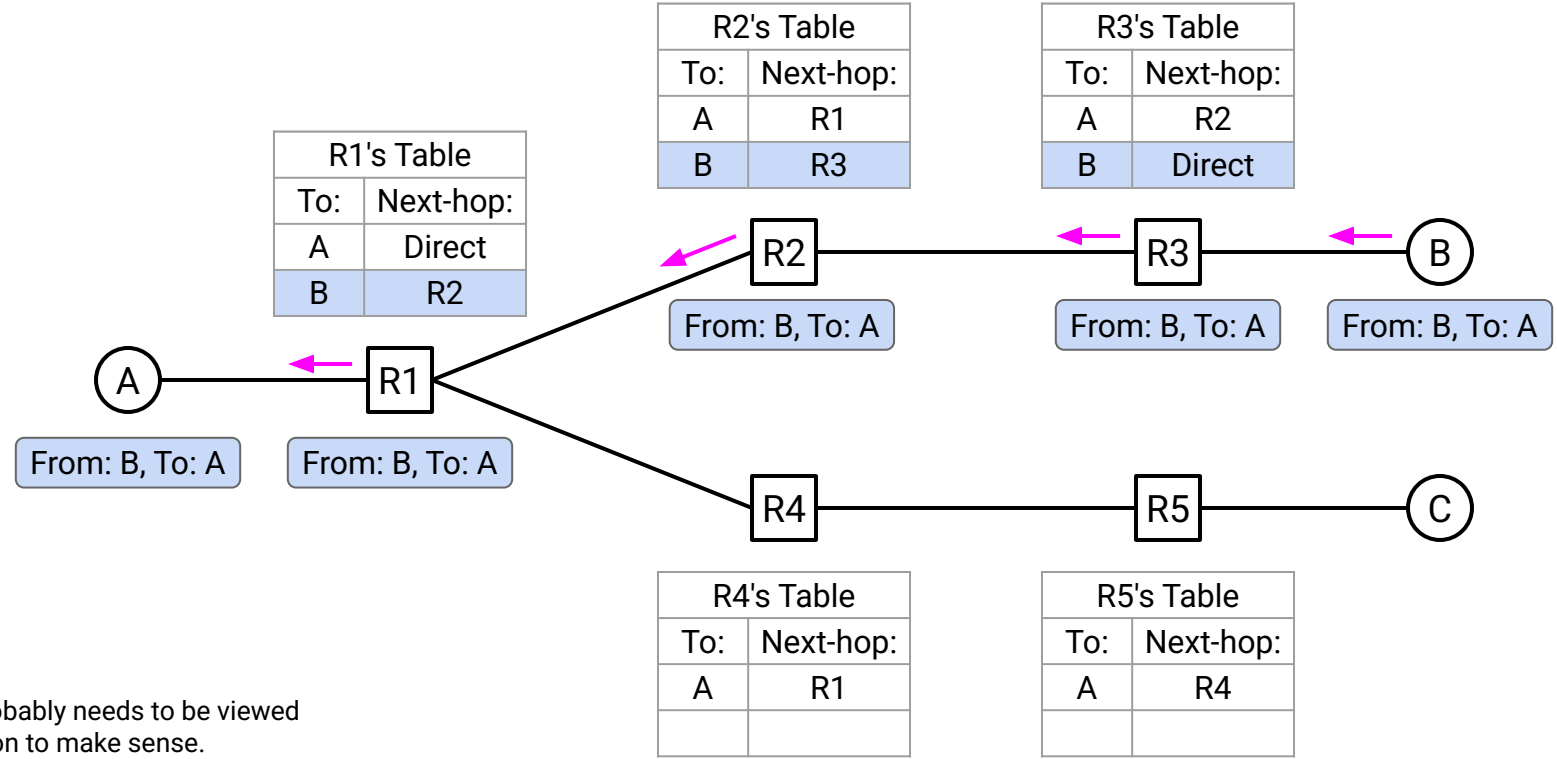

This slide probably needs to be viewed with animation to make sense.

Learning Switches: Expiring Routes

Each table entry is associated with a TTL. If the TTL reaches 0, the entry is deleted.

- This allows invalid routes (e.g. link goes down) to eventually expire.

Spanning Tree Protocol (STP)

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- Sending Packets
- Scaling Ethernet

Layer 2 Ethernet Routing

- Naive Approach: Flooding
- Learning Switches
- **Spanning Tree Protocol (STP)**
- Implementing STP: BPDU Exchanges

Spanning Tree Protocol: Motivation

Problems with flooding? Wastes bandwidth in two ways:

1. Packet is sent to other hosts (not the destination).
   - Solution: Use learning switches to build forwarding tables.
2. Loops create broadcast storms.
   - Solution: Use STP to remove loops.

Fixing these 2 problems will result in a good Layer 2 routing protocol.

We'll now build a protocol for removing loops.

After the loops are removed, we can use learning switches to build forwarding tables and forward packets.

Spanning Tree Protocol: Motivation

Loops can cause **broadcast storms**: Packets get forwarded around a loop forever.

- Forwarding tables don't help. If the tables are empty, everyone floods the packet.
- Learning switches don't help. We can install entries for the sender, but we still don't know where the receiver is, so the packet still gets flooded forever.
- Not flooding out of the incoming port doesn't help.
  Packets could travel in loops of length 3 or more. (See example below.)

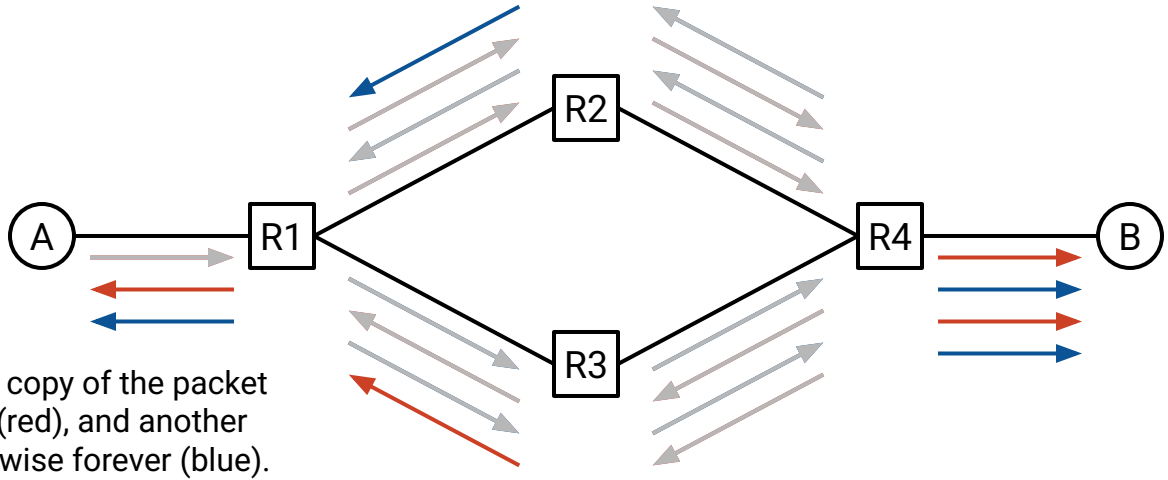

Color is just for clarity. A copy of the packet loops clockwise forever (red), and another copy loops counterclockwise forever (blue).

Spanning Tree Protocol: Motivation

Our goal: "Disable" switch ports to remove all the loops.

- We'll now design the **Spanning Tree Protocol (STP)** to achieve this.
- After the loops are removed, we can use learning switches to route packets.

Some other possible solutions, besides STP:

- Drop packets when TTL expires: Won't work, because Ethernet header doesn't have a TTL field.
- Drop packets we've seen before: Won't work, because Ethernet header doesn't have an ID field for identifying duplicates.

STP Ignores Hosts

Hosts don't participate in STP.

Switches will do the job of **disabling ports** to remove loops.

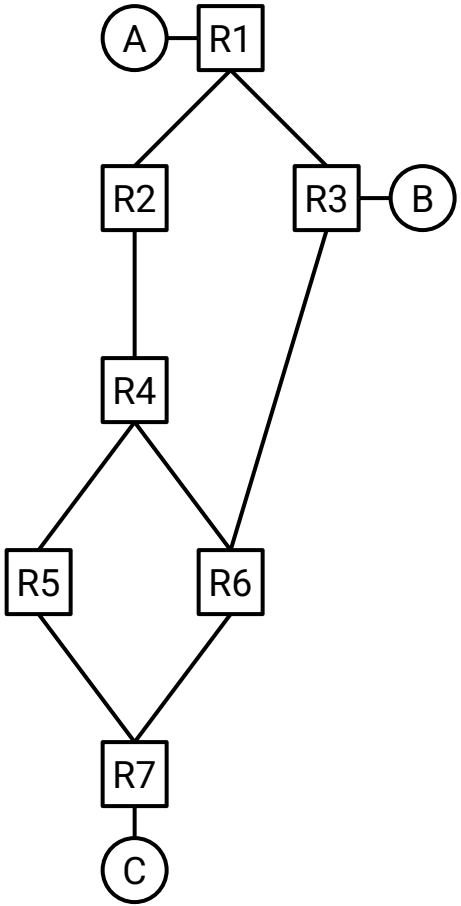

Step 1: Electing a Root

Step 1: Elect one of the switches to be a **root switch**.

- Each switch is assigned an ID.
- The root switch is the switch with the **lowest ID**.

Each switch's ID consists of two values, put together:

- **Priority**: Manually set by operator.
- **MAC address** of switch.
  Used as tiebreaker if two switches have the same priority.

How do we pick a root switch?

- **Manually**: Operator sets a low priority on the chosen root.
- **Default**: All switches have the same default priority value.
  If operator doesn't override the default, the switch with the lowest MAC address is the root.

We'll assume the ID is the switch number.

Thus, R1 is the root.

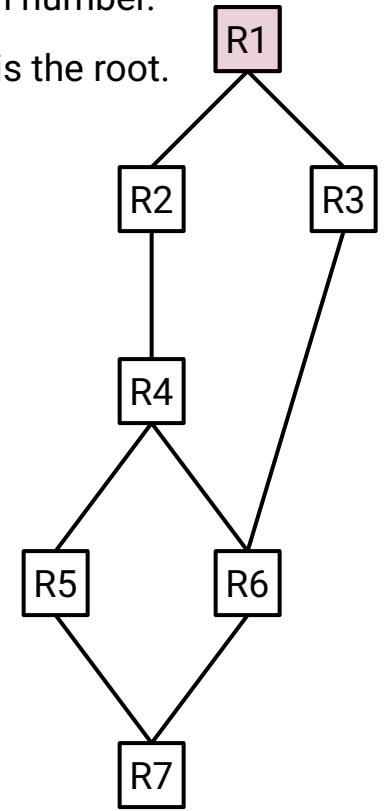

Step 2: Labeling Ports: Each switch labels each of its ports as one of three types

Rule1: For ports to switches that are closer or further away from root than me

Switches distance < *x* away from the root.

R4 is *x* away from root.

Switches distance > *x* away from the root.

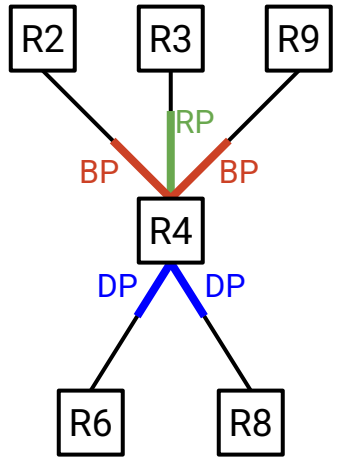

Out of the ports leading to switches *closer* to the root:

- The **Root Port** is the one along the least-cost path to the root.
- **Blocked Ports** are all the other ports (not along least-cost path).

Designated Ports lead to switches *further* from the root.

Step 2: Labeling Ports: : Each switch labels each of its ports as one of three types

Rule 2: Breaking Ties

What if there are two shortest paths to the root?

Use switch ID to tiebreak.

- Lower ID = shortest path to root (**RP**).
- Higher ID = not shortest path to root (**BP**).

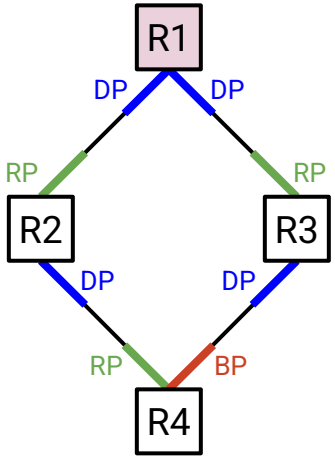

R2 and R3 both yield shortest paths to root.

Tiebreak: Choose R2 (lower ID).

Step 2: Labeling Ports: : Each switch labels each of its ports as one of three types

Rule 3: If the link leads to a switch equally-far from the root as you:

- If their ID is higher, they are **DP** (“further").
- If their ID is lower, they are **BP**

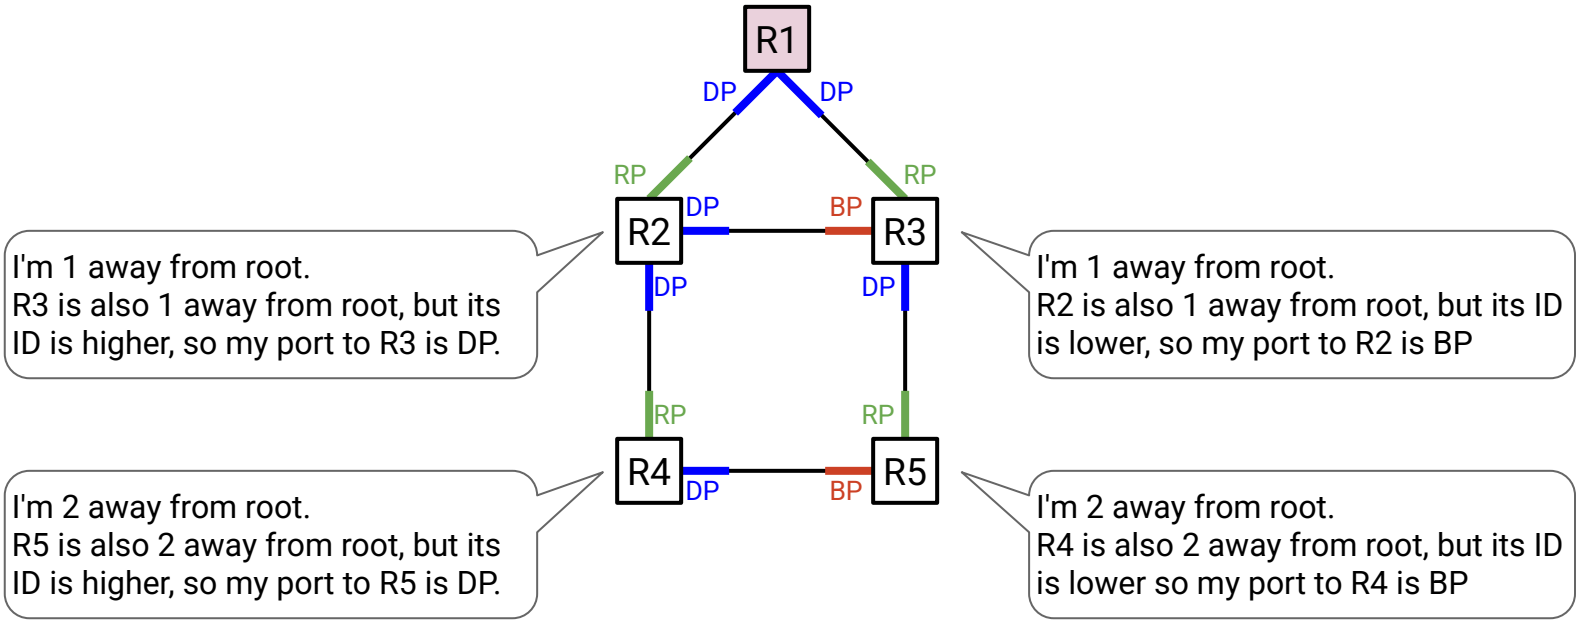

Port Labels: Designated Ports

Designated Ports lead to switches *further* from the root or to switches that are *equidistant from the root but have higher ID*

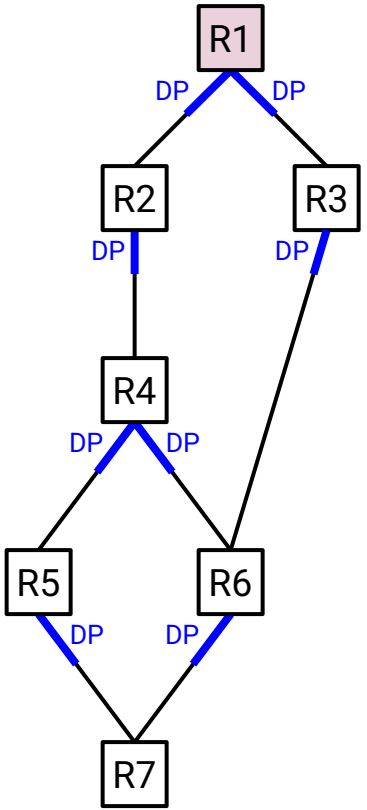

Port Labels: Root Ports

- The **Root Port** is the one along the least-cost path to the root.

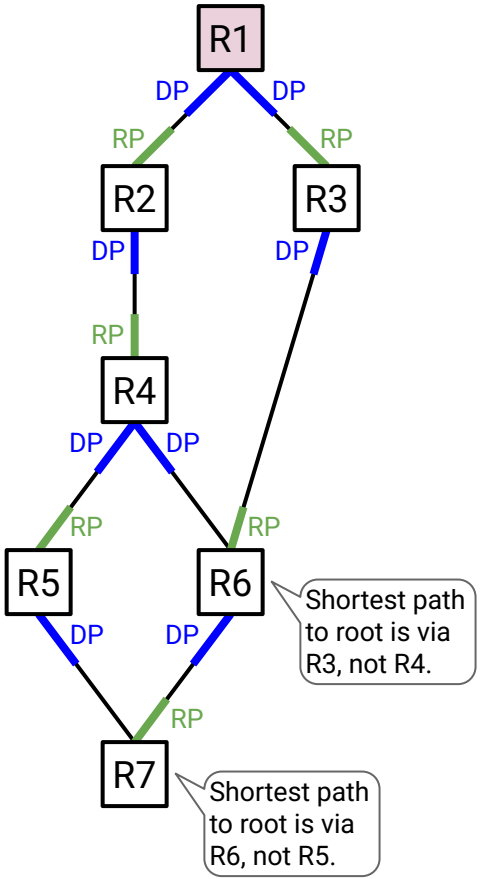

Port Labels: Blocked Ports

- **Blocked Ports**
  - ports to switches with lower cost than mine but not along my least-cost path
  - Ports to switches with equal cost to mine but with lower ID than mine

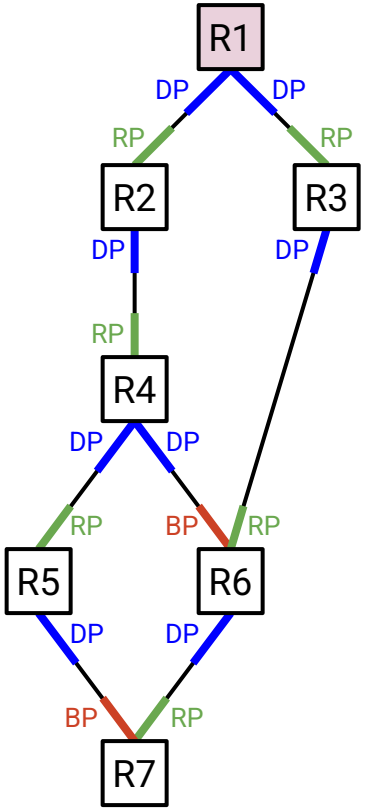

Step 3: Disabling Ports

Each switch can now disable its **blocked ports**.

- To disable a port: Do not send/receive data packets with that port.
- No link is disabled on both sides

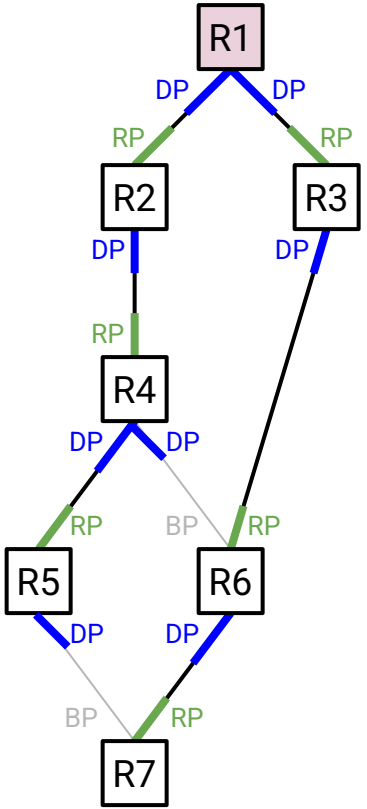

Flooding:

- When a switch receives a packet on an enabled port, send it out over all other enabled ports other than the incoming port

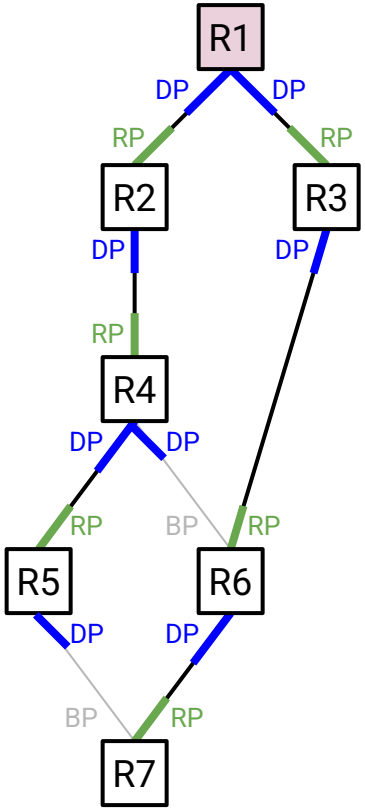

Implementing STP: BPDU Exchanges

Lecture 17, CS 168, Fall 2026

Ethernet

- Connecting Local Hosts
- Multiple Access Protocols
- Sending Packets
- Scaling Ethernet

Layer 2 Ethernet Routing

- Naive Approach: Flooding
- Learning Switches
- Spanning Tree Protocol (STP)
- **Implementing STP: BPDU Exchanges**

Implementing STP with BPDU Exchanges

STP requires global knowledge of the network:

- Who is the root (lowest ID)?
- Does this link lead closer to the root, or further from the root?

To learn this information, switches exchange control-plane messages.

- They have a fancy name: **Bridge Protocol Data Units (BPDUs)**.
- Similar to advertisements in distance-vector.

BPDU Advertisement Rules

Each switch keeps track of two things:

- Who is the best root I know about? *(Recall: "better" = "lower ID".)*
- How far away is this best root?

When you send a BPDU to your neighbors, include these two things.

When you hear a BPDU from your neighbor, accept it and update your state if:

- The advertised root has a lower ID. *(Note: Accept the better root, even if cost is worse.)*
- The advertised root is the same, but the cost to root is better.

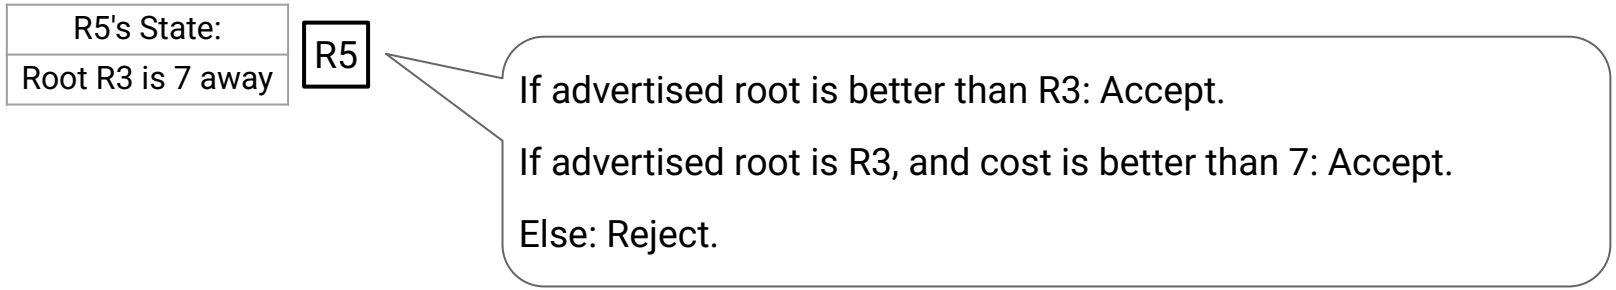

BPDU Exchange Example

Initially, every switch only knows about itself, so it thinks the root is itself.

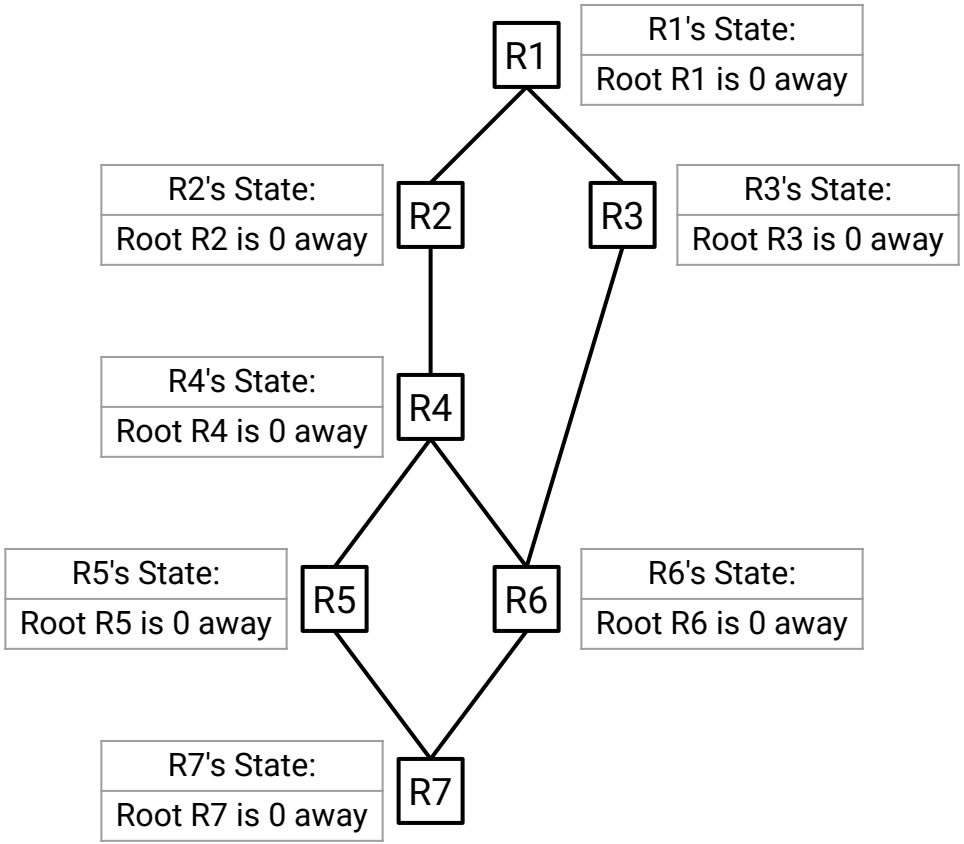

BPDU Exchange Example

Accept if better root, or better cost to same root.

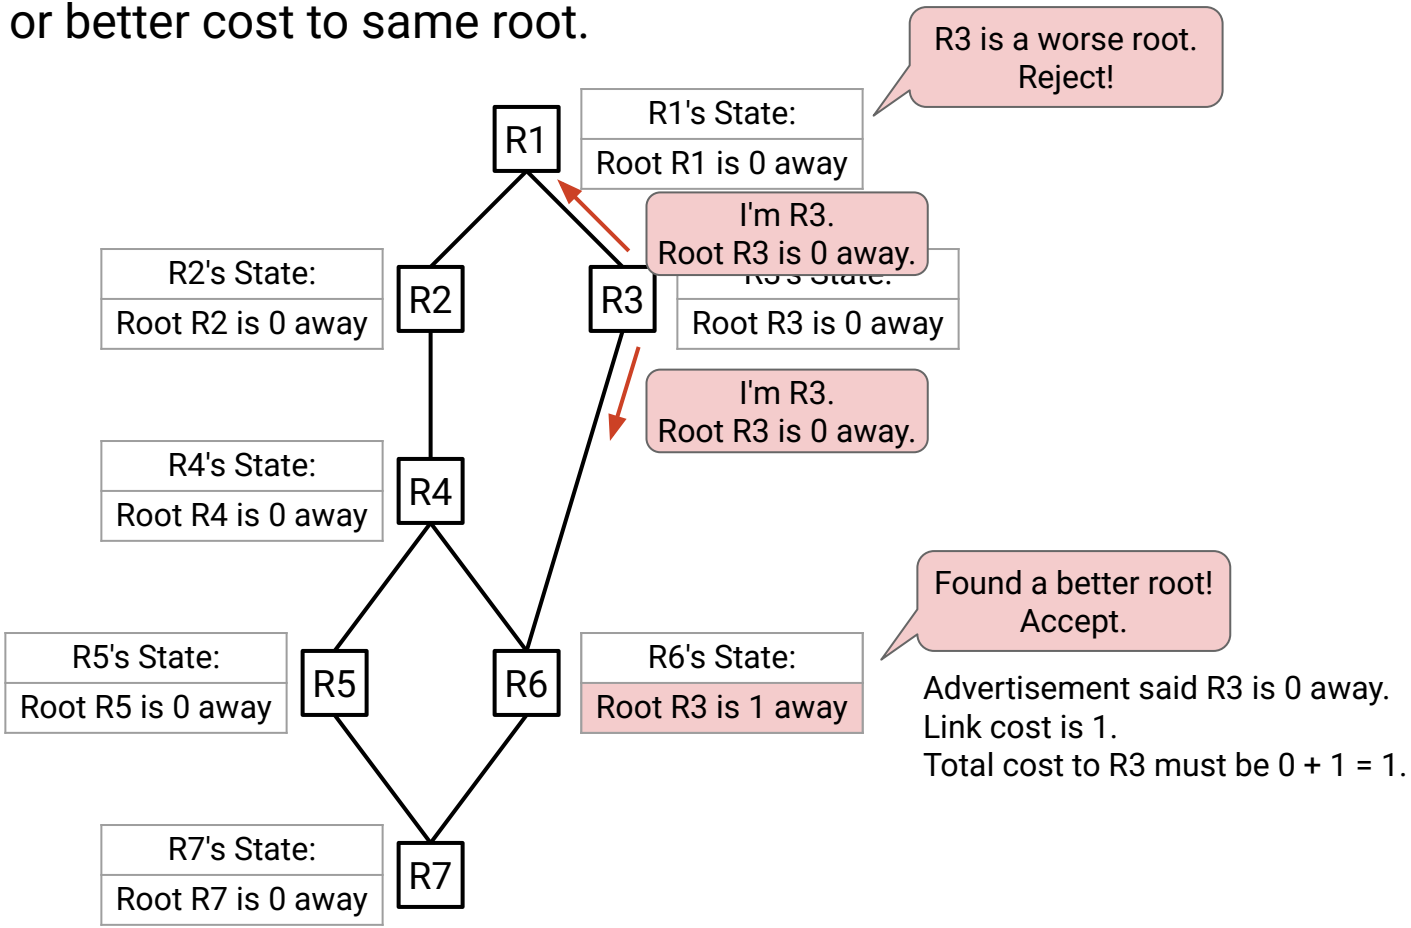

BPDU Exchange Example

Accept if better root, or better cost to same root.

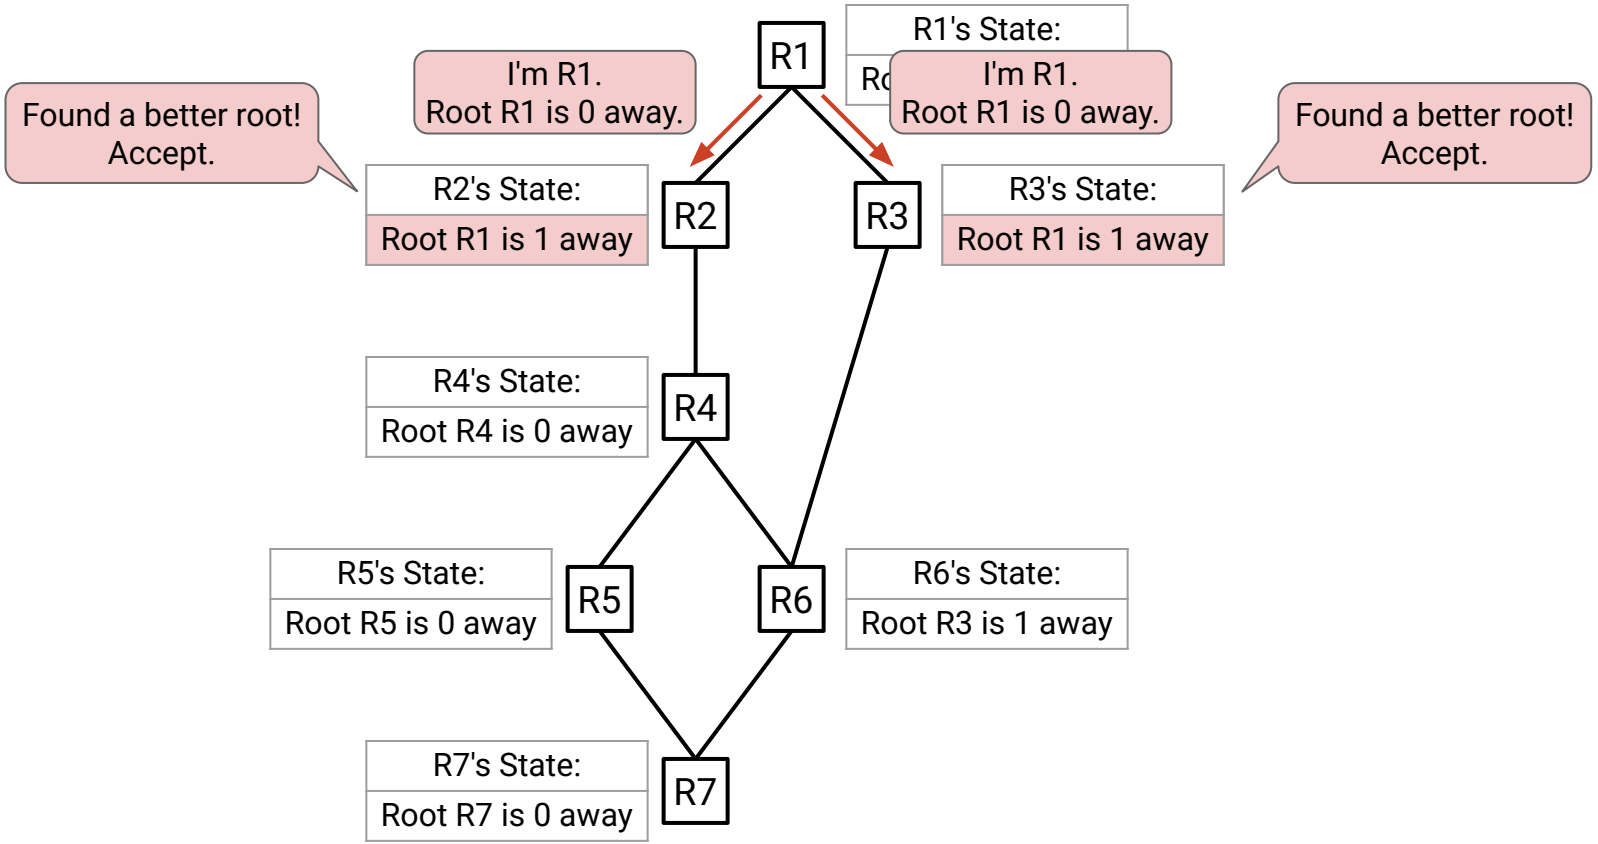

BPDU Exchange Example

Accept if better root, or better cost to same root.

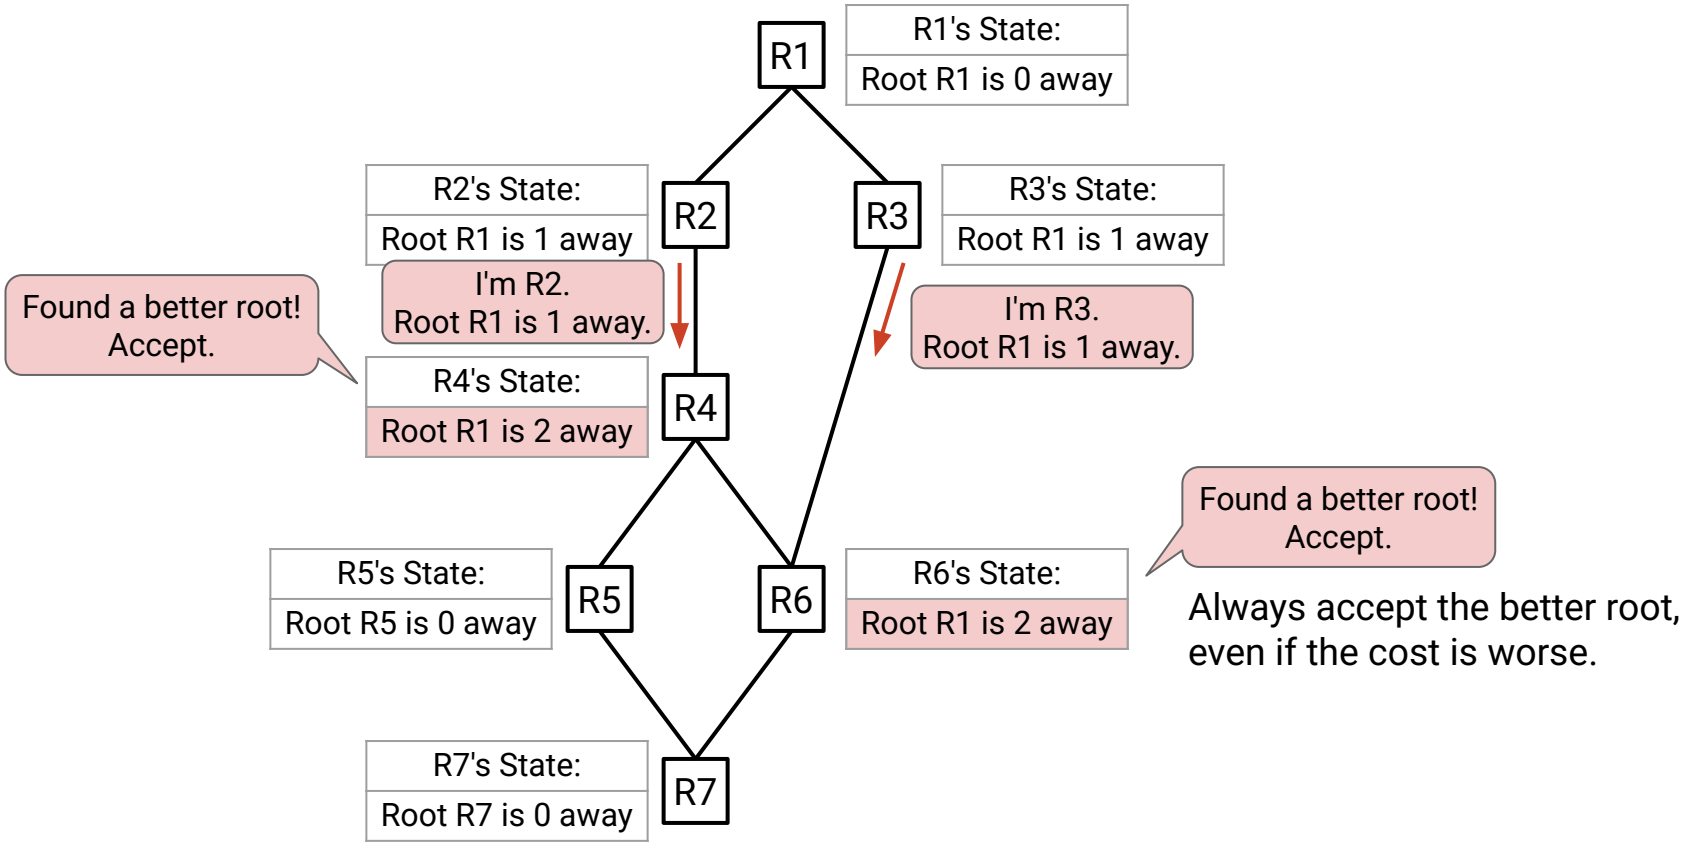

Always accept the better root, even if the cost is worse.

BPDU Exchange Example

Accept if better root, or better cost to same root.

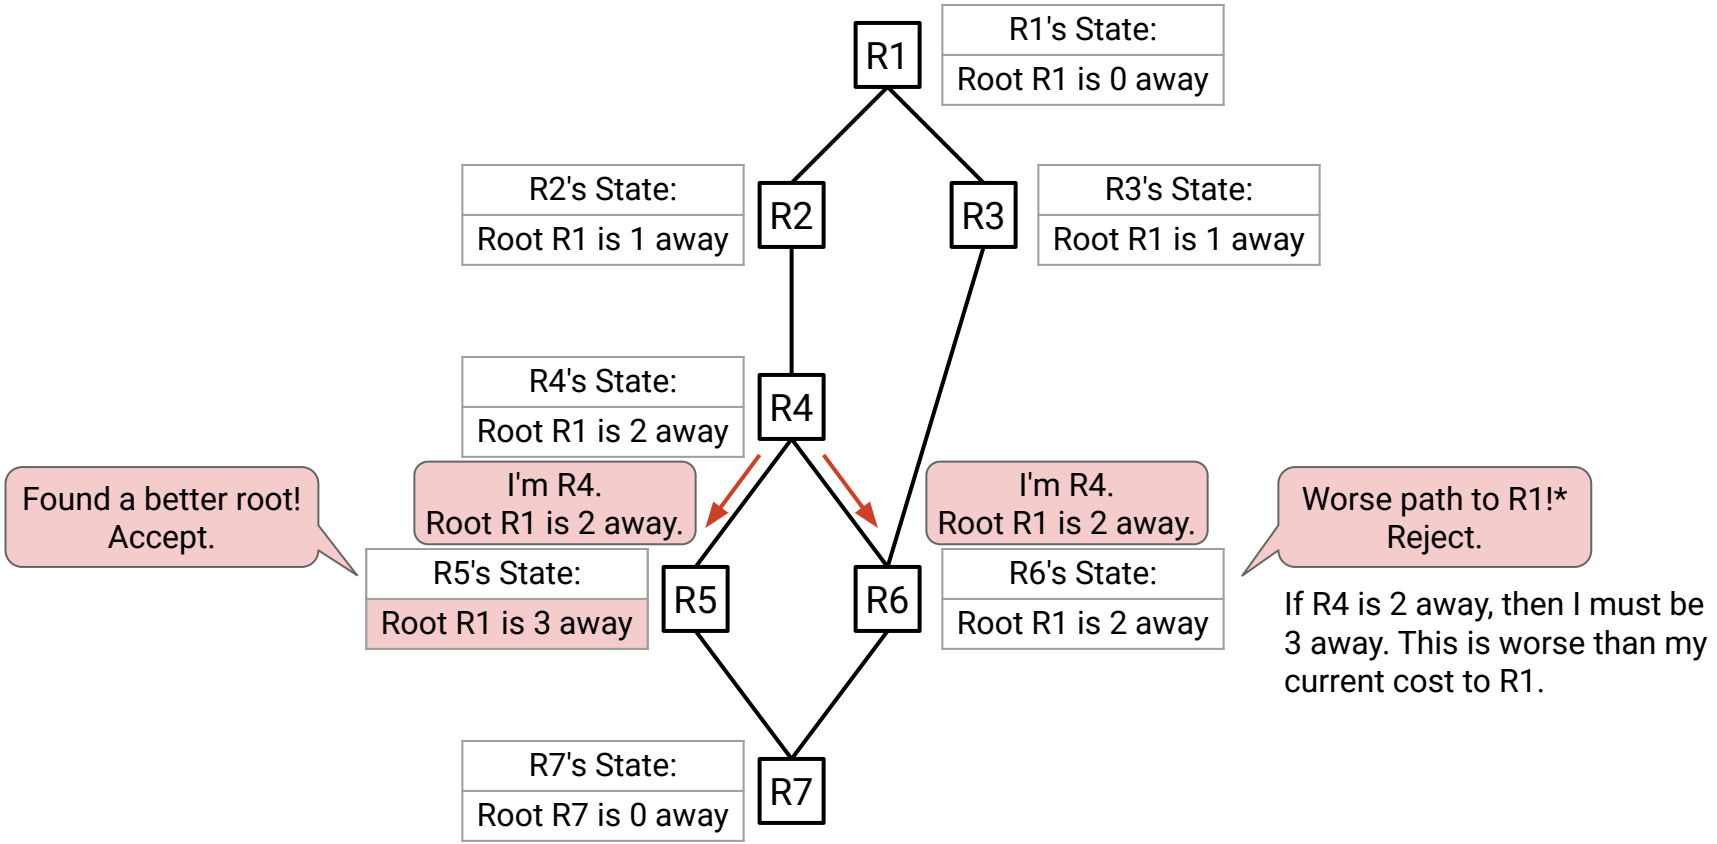

BPDU Exchange Example

Accept if better root, or better cost to same root.

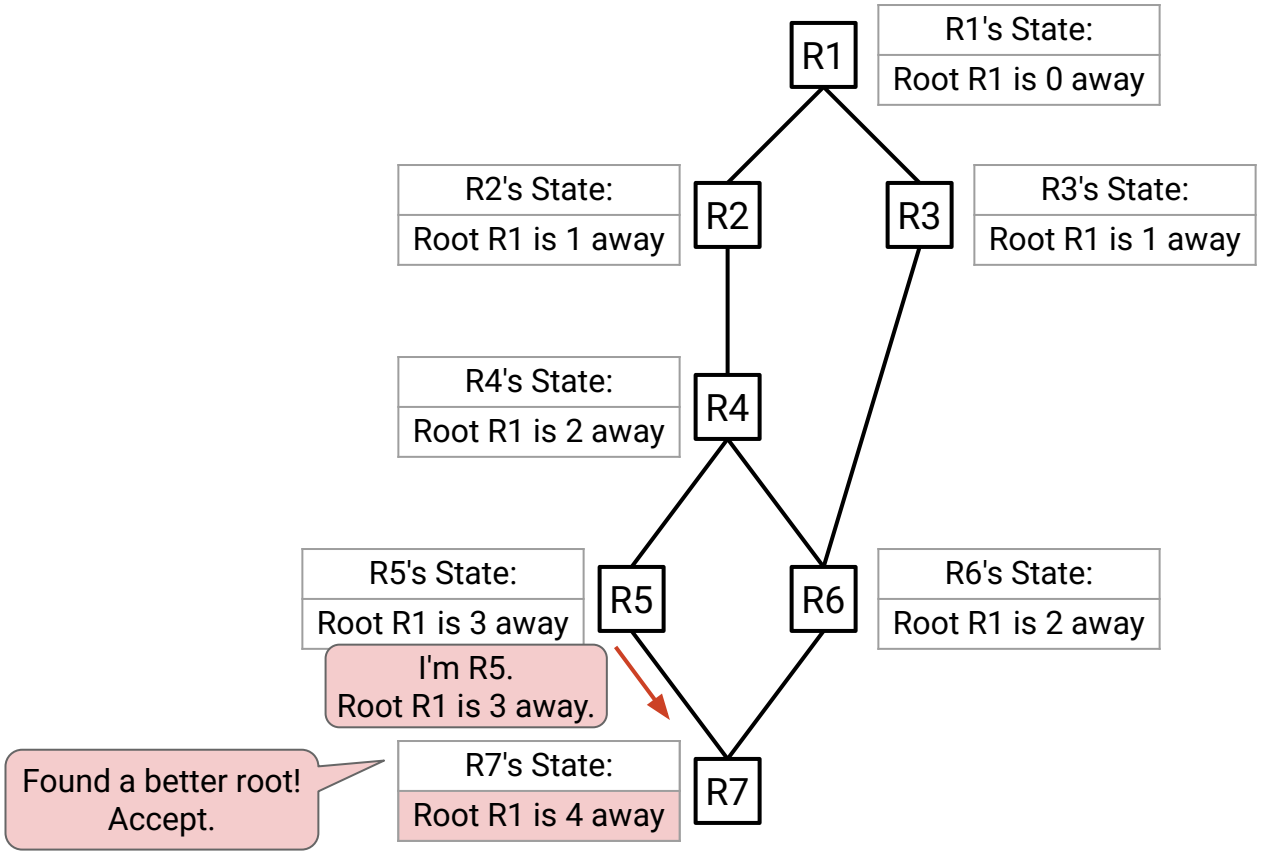

BPDU Exchange Example

Accept if better root, or better cost to same root.

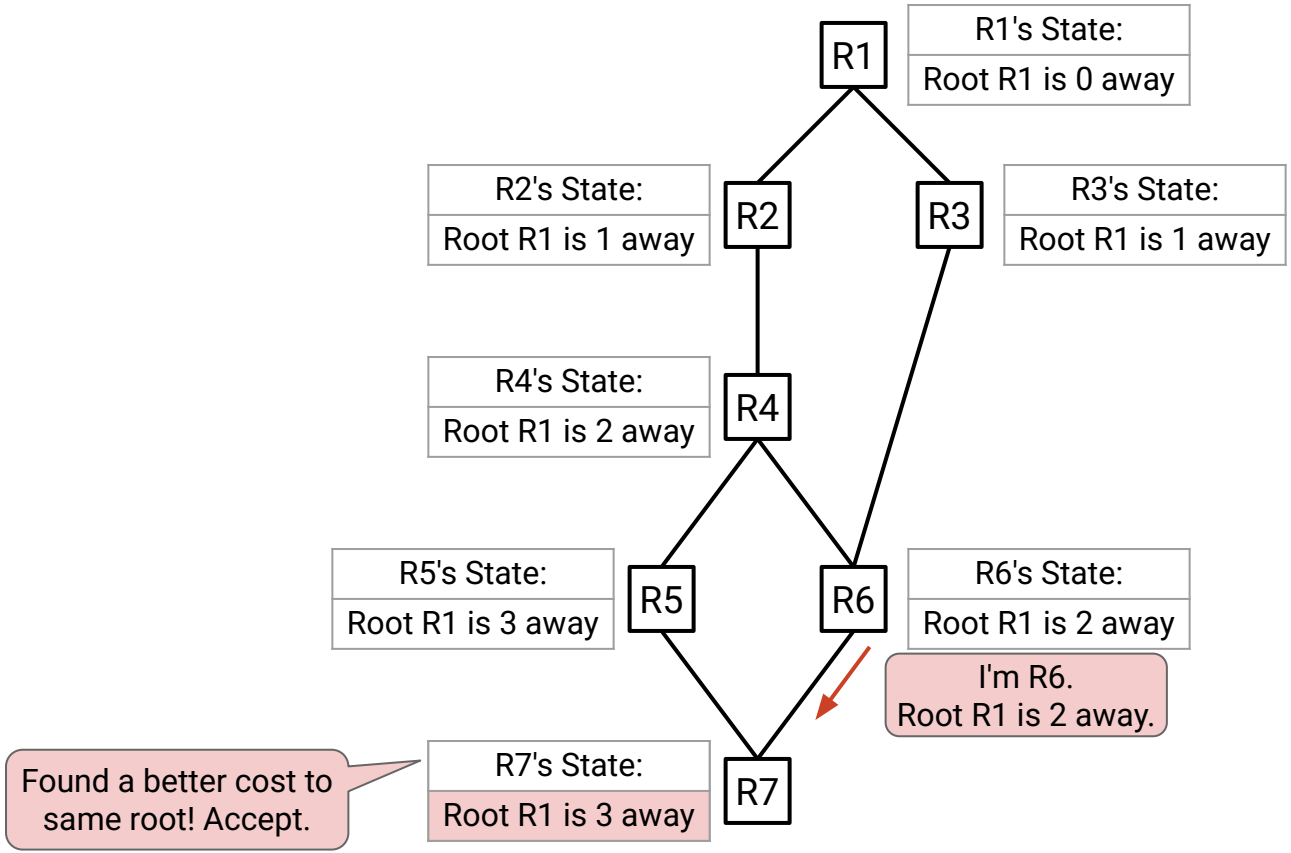

BPDU Exchange Example

After convergence, periodic advertisements allow you to learn your neighbors' cost-to-root values.

Use your neighbors' cost-to-root values to assign DP, RP, BP.

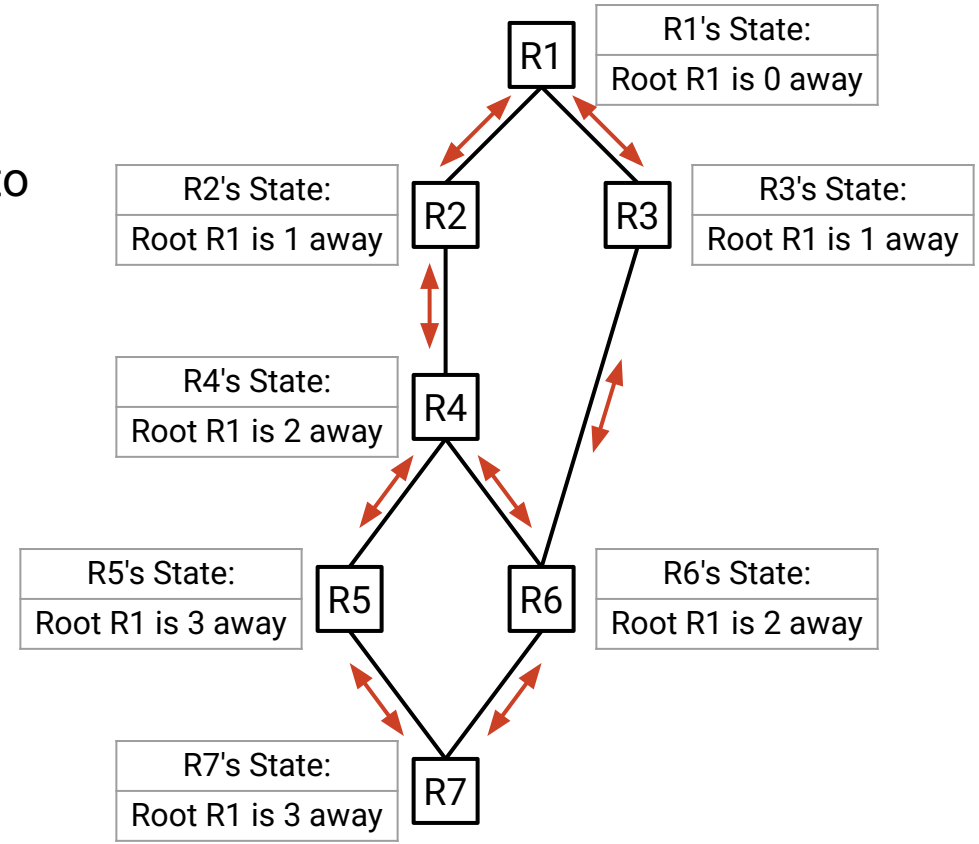

Why BPDU Exchanges Work

Why does this work?

- The true root has the lowest ID, so advertisements with the true root will always be accepted over worse roots.
- Other exchanges might happen (e.g. "I thought root is R7, now I think it's R3"), but eventually, the true root (e.g. R1) will propagate to everybody, along with every neighbor’s cost to that root

Spanning Tree Protocol: Recap

Flooding on a topology with cycles leads to broadcast storms

Idea: remove cycles in the topology

- Pick switch with lowest ID as root
- Construct a spanning tree anchored at the root
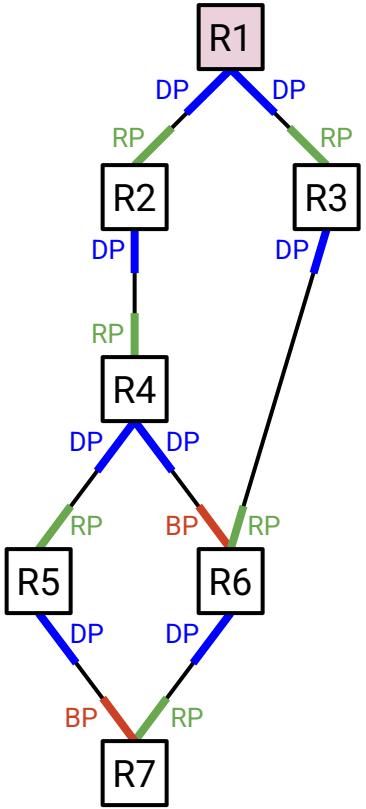
Implemented via BPDU exchange that simultaneously picks root and spanning tree