CS4450

Computer Networks: Architecture and Protocols

Lecture 9
Recap: Spanning Tree Protocol
Fundamentals of Routing

Rachit Agarwal

Goals for Today's Lecture

- Recap **Spanning Tree Protocol**
- **Why** do we need network layer?
  - **Why** not just use switched Ethernet across the Internet?
- **Fundamentals** of network layer
  - Routing tables
  - The **right** way to think about routing tables
- But, before that …..

Exam 1 Updates

- **I am SO proud of you all!**

- Full marks **50/50: ~0.01%** of the class
- More than **45/50: ~27%** of the class
- More than **40/50: ~47%** of the class
    - **Absolutely amazing!**

- **Mean: ~38.69 (last time I taught: 36)**

- **Median: ~39 (last time I taught: 36.5)**

- **Std. Dev.: ~7.06 (last time I taught: 11)**

Exam 1 Discussions

- **I am here for you.**

- If you would like to go through your exam copy
    - I will make time for each and every one of you
        - To discuss how/where we can improve
    - Send an email to cs4450-prof@cornell.edu to set up a meeting
    - **Please send me your availability**

Recap of Link Layer so far

Recap: Link layer

- **Traditional Link Layer: Broadcast Ethernet**
- **CSMA/CD**
    - Random access on a broadcast channel
    - Exponential Backoff
- **Why Frames?**
    - To incorporate sentinel bits for identifying frame start/end
    - To incorporate link layer source and destination names
    - To incorporate CRC for checking correctness of received frames
- **Modern Link Layer: Switched Ethernet**
    - Why? Scalability limits of traditional Ethernet
        - Why? Detecting collisions on a broadcast channel

Recap: Switched Ethernet

- **Hosts connect to broadcast (Ethernet) buses**
    - Each bus has a maximum length and/or minimum frame size
- **Multiple broadcast buses connected via relays/switches**
    - Can now scale to arbitrarily large lengths
- **How to transfer data across broadcast buses connected via relays**
    - Cannot simply forward the data across relays
    - The topology may have loops
    - **Recall: broadcast storm problem!**
- **Core idea in switched Ethernet: Spanning Tree Protocol**
    - Switches create a Spanning Tree
    - Using THE Spanning Tree Protocol

Recap: Spanning Tree definition

- **Subgraph that includes all vertices but contains no cycles**
  - Links not in the spanning tree are not used in forwarding frames

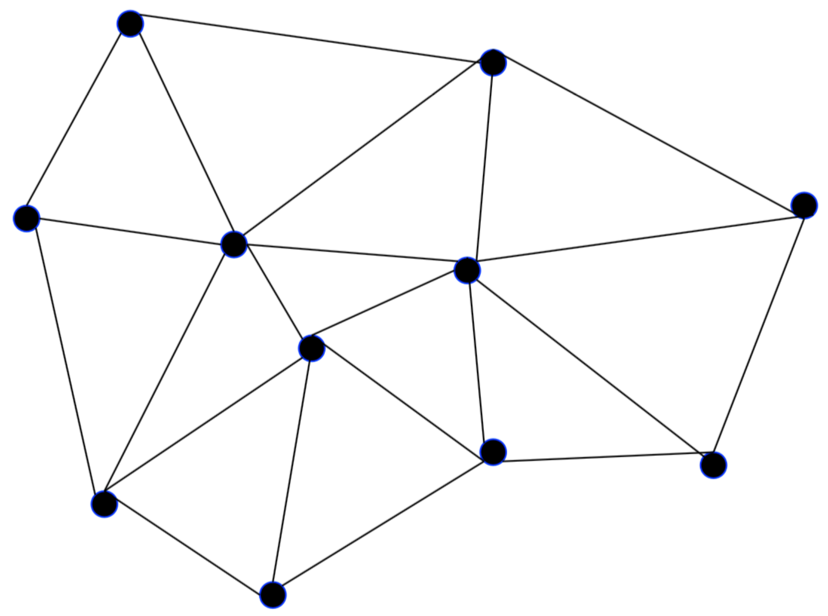

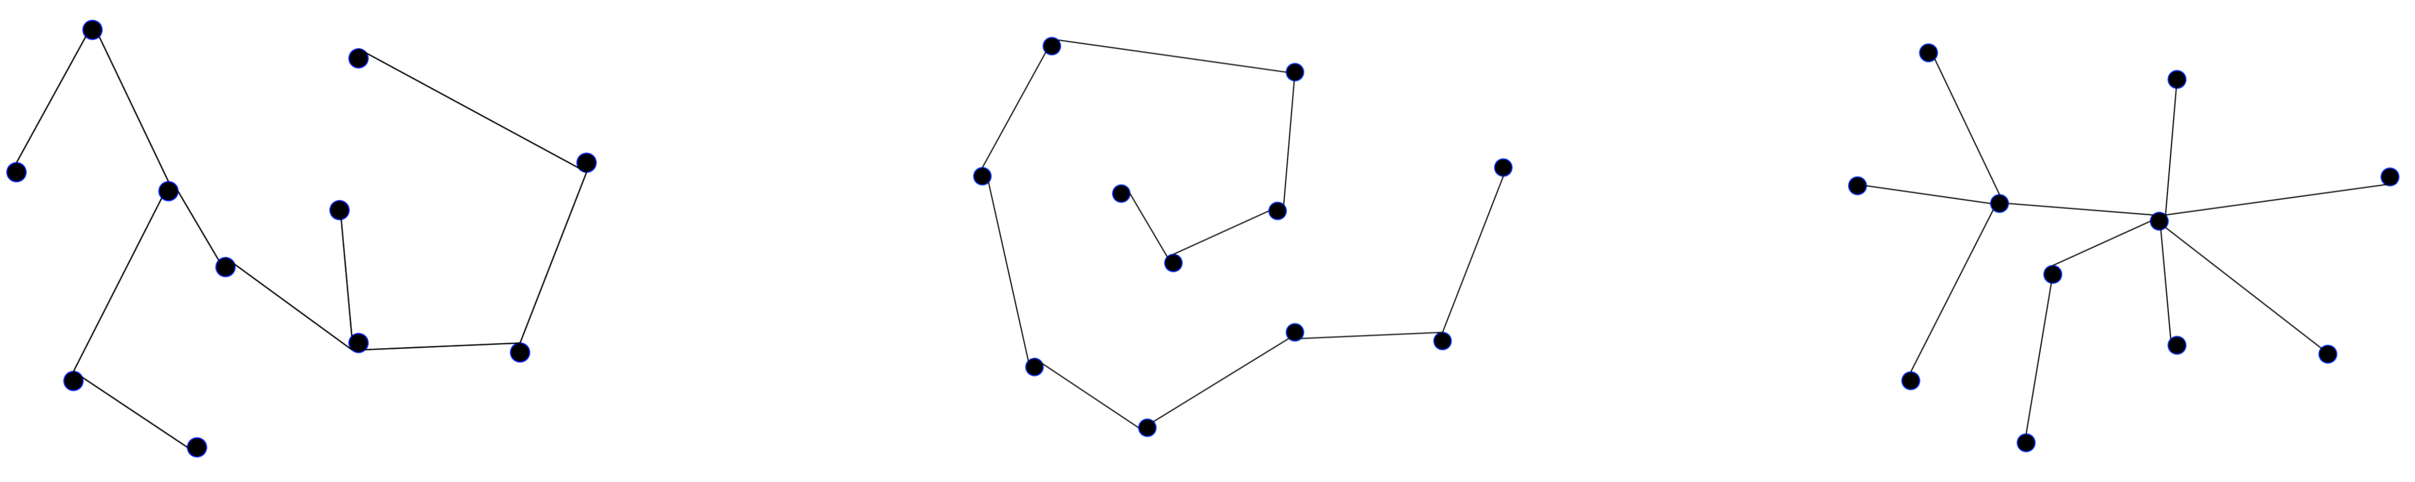

Recap: Spanning Tree Protocol

- **Messages (Y,d,X)**
  - Proposing root Y; from node X; advertising a distance d to Y

- Initially each switch proposes itself as the root
  - that is, switch X announces (X,0,X) to its neighbors

- At each switch Z:

  WHENEVER a message (Y,d,X) is received from X:
  - IF Y's id < current root
    - THEN set root = Y; next-hop = X
  - IF Shortest distance to root > d + distance_from_X
    - THEN set shortest-distance-to-root = d + distance_from_X
  - IF **root changed OR shortest distance to the root changed:**
    - Send all neighbors message (Y, shortest-distance-to-root, Z)

We ran the Spanning Tree Protocol on this example

(assume all links have "distance" 1)

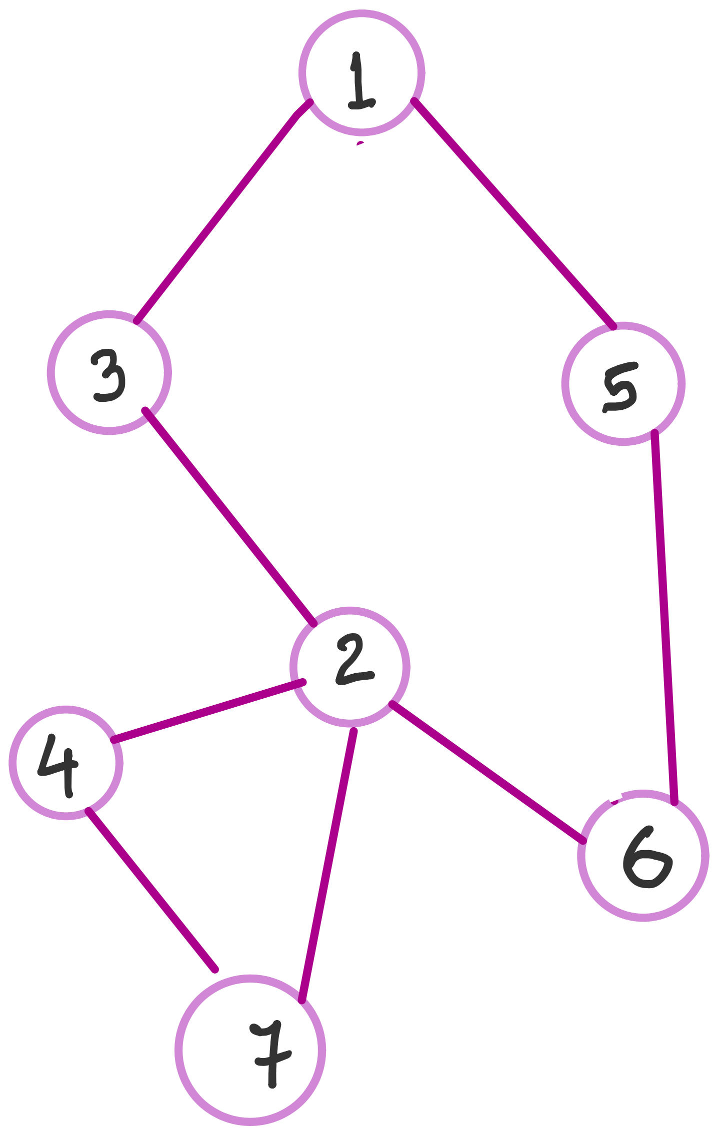

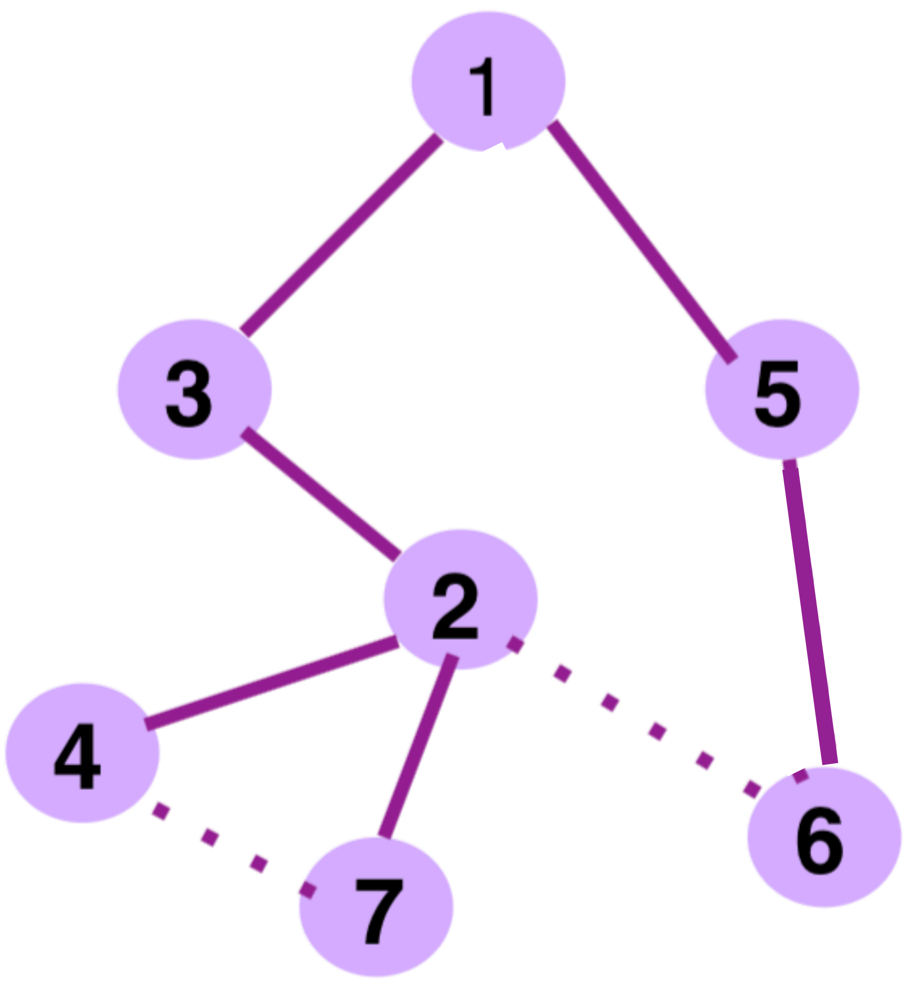

Another example: Spanning Tree Protocol

(assume all links have "distance" 1)

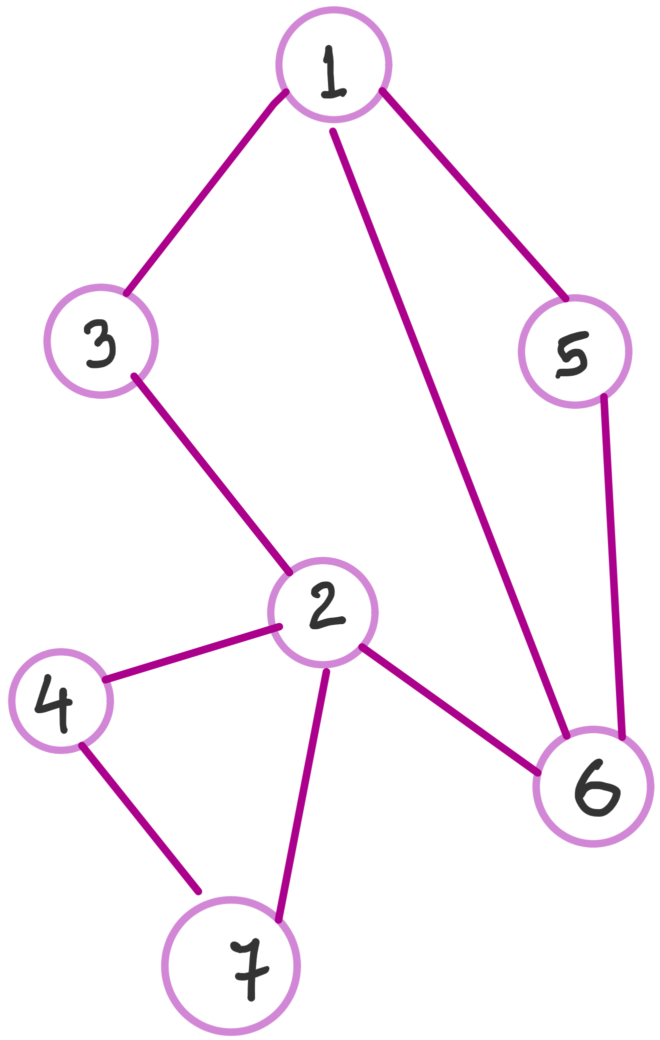

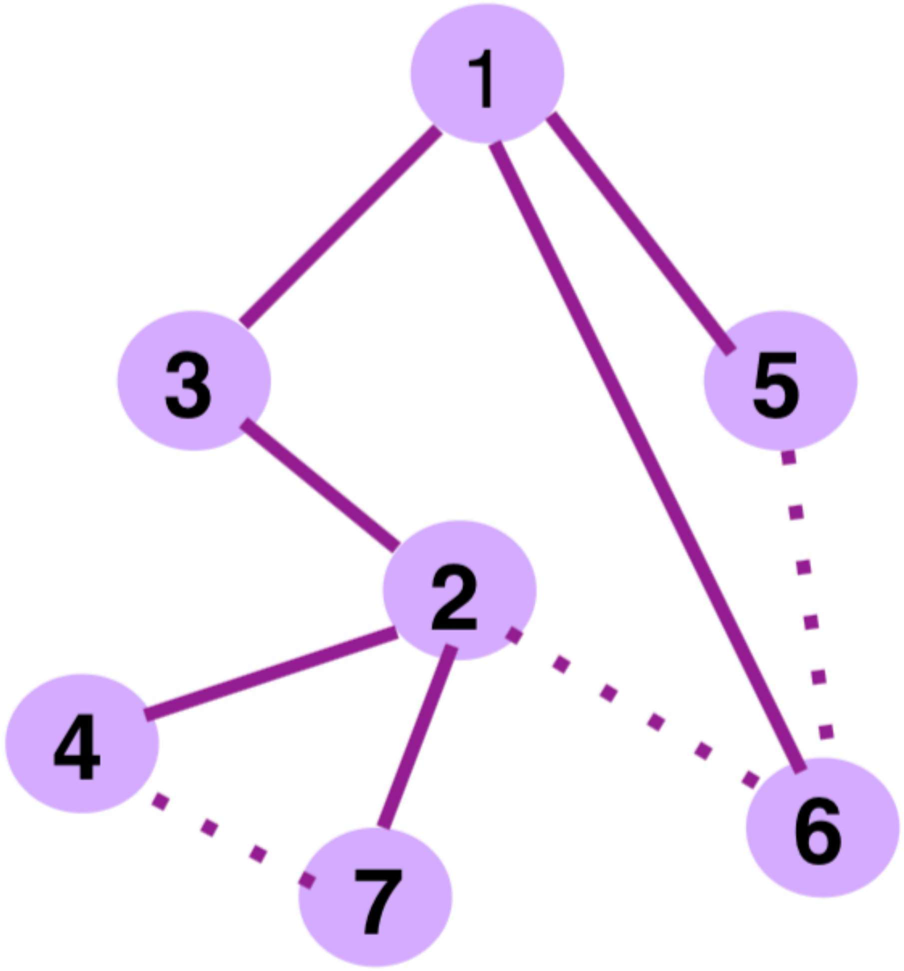

After Round 5: We have our Spanning Tree

- 3-1
- 5-1
- 6-1
- 2-3
- 4-2
- 7-2

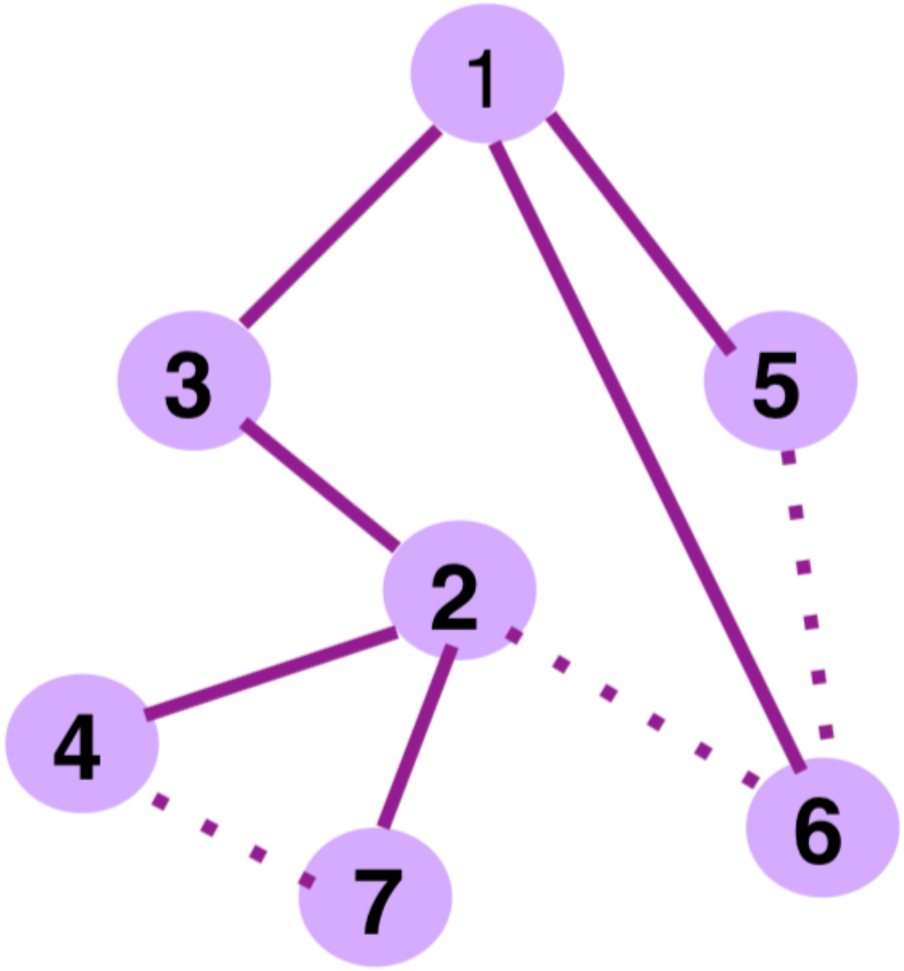

Spanning Tree Protocol ++ (incorporating failures)

- Protocol must react to failures
    - Failure of the root node
    - Failure of switches and links

- **Root node sends periodic announcement messages**
    - Few possible implementations, but this is simple to understand
    - Other switches continue forwarding messages

- Detecting failures through timeout (soft state)
    - If no word from root, time out and send a (Y, 0, Y) message to all neighbors (in the graph)!

- **If multiple messages with a new root received, send message (Y, d, X) to the neighbor sending the message**

Example: Suppose link 2-4 fails

- 4 will send (4, 0, 4) to all its neighbors
    - 4 will stop receiving announcement messages from the root
    - Why?
- At some point, 7 will respond with (1, 3, 7)
- 4 will now update to (1, 4, 4) and send update message
- New spanning tree!

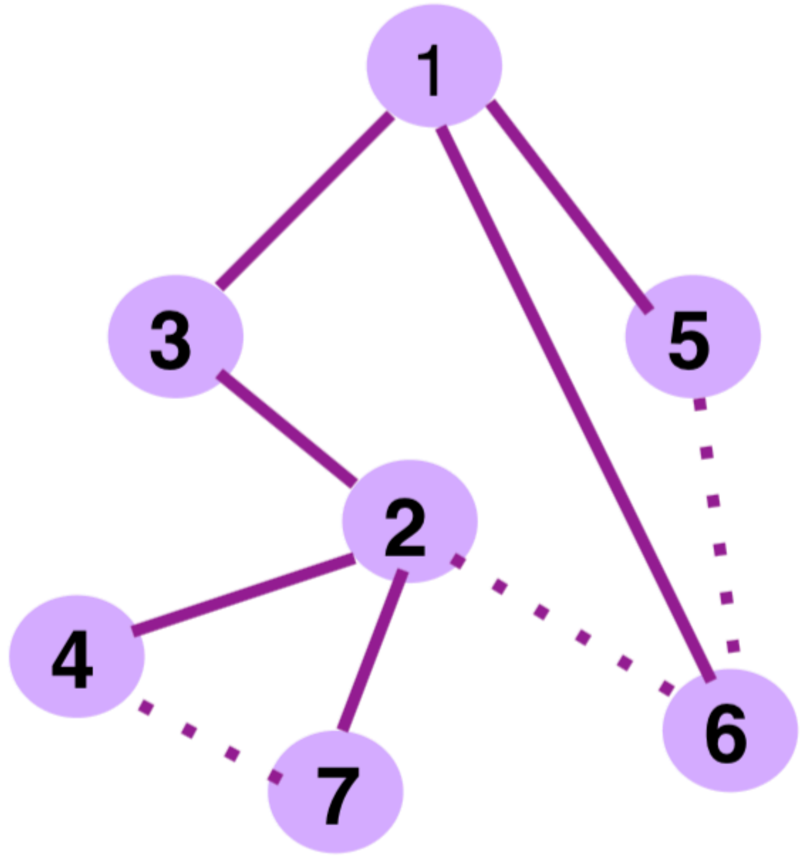

Questions?

The end of Link Layer ….

And the beginning of network layer!

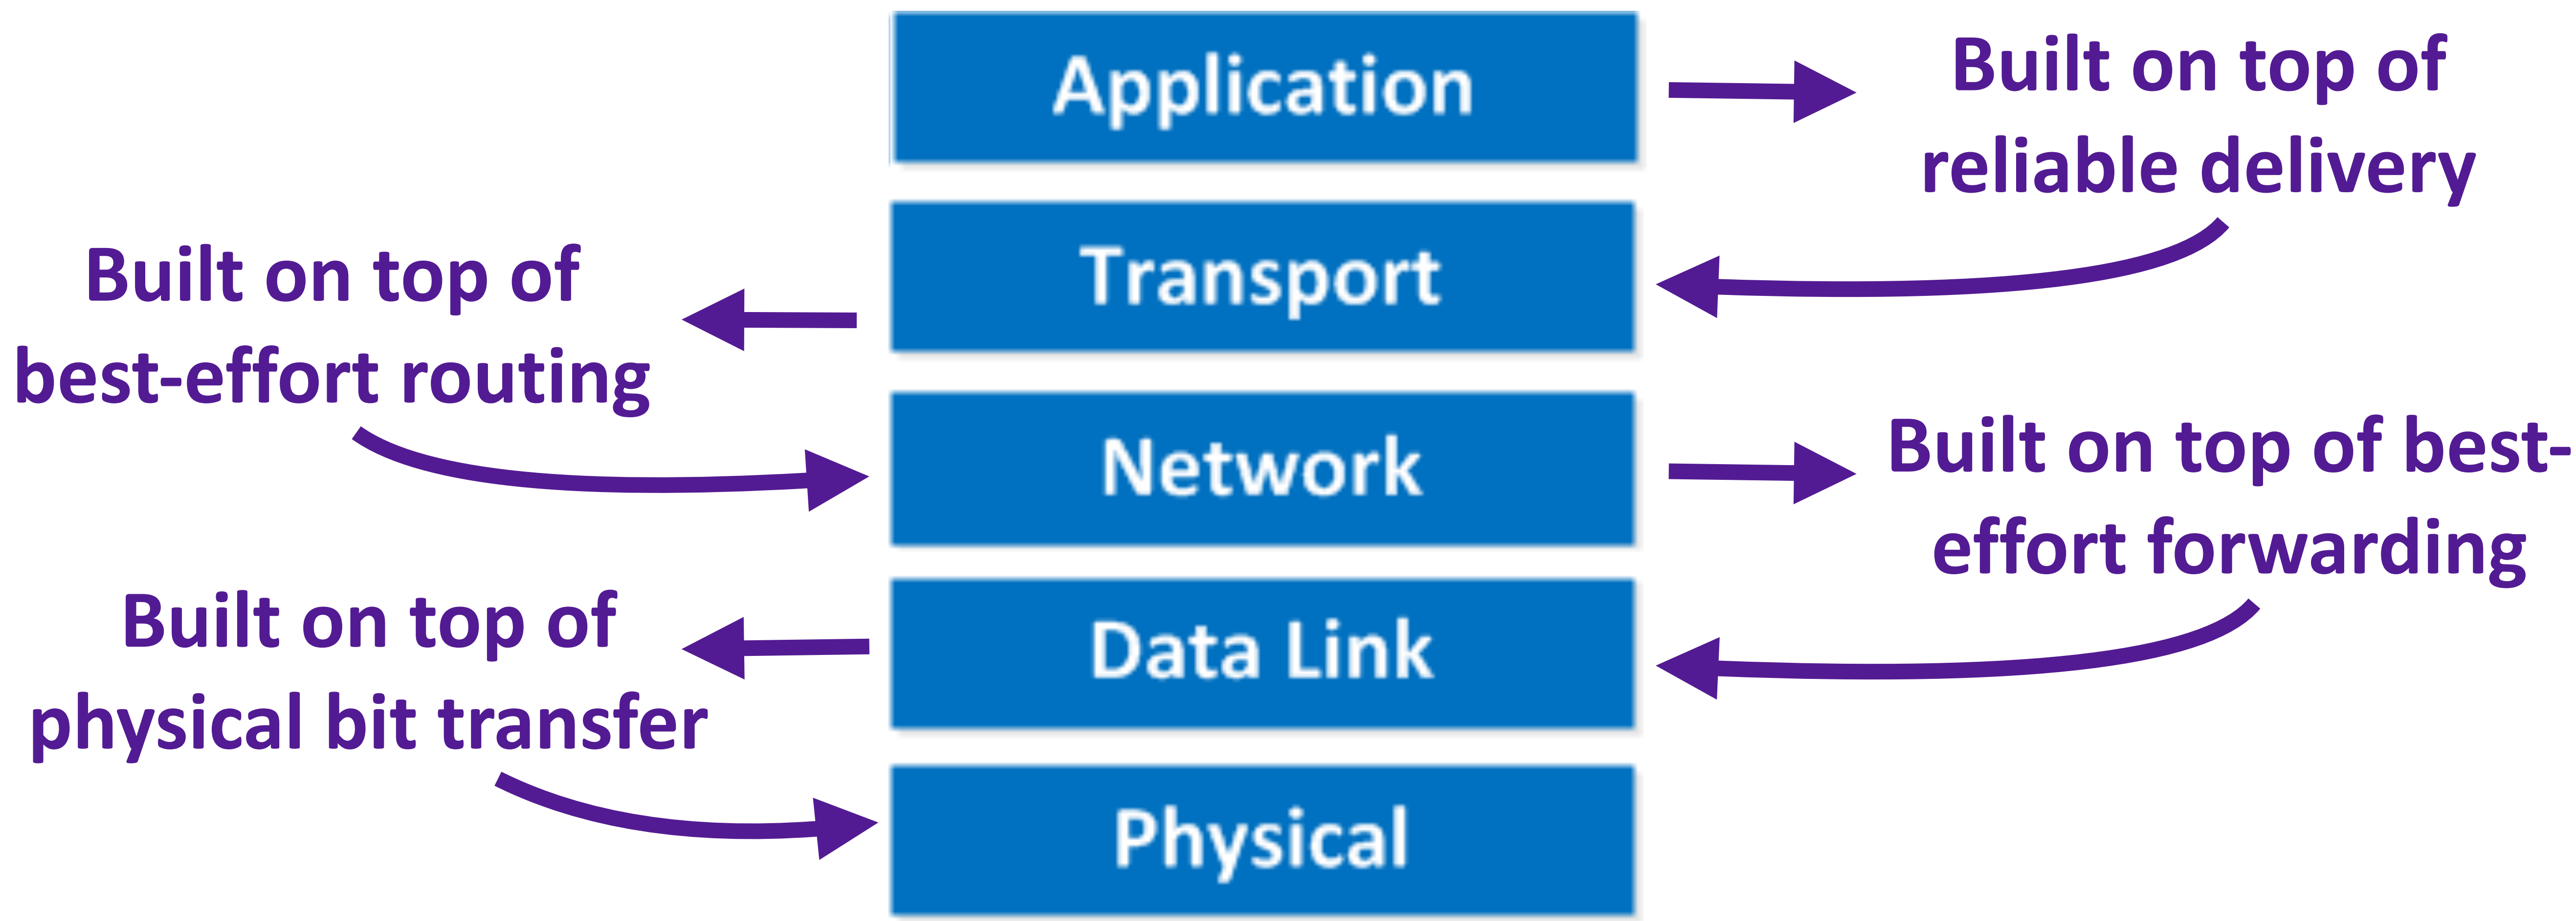

Why do we need a network layer?

- Why not just use spanning trees across the entire network?

- Easy to design routing algorithms for (spanning) trees
    - **Nodes can “flood” packet to all other nodes**

Flooding on a Spanning Tree

- Sends packet to ***every*** node in the network

- **Step 1**: Ignore the links not belonging to the Spanning Tree

- **Step 2**: Originating node sends "flood" packet out every link (on spanning tree)

- **Step 3**: Send incoming packet out to all links **other than the one that sent the packet**

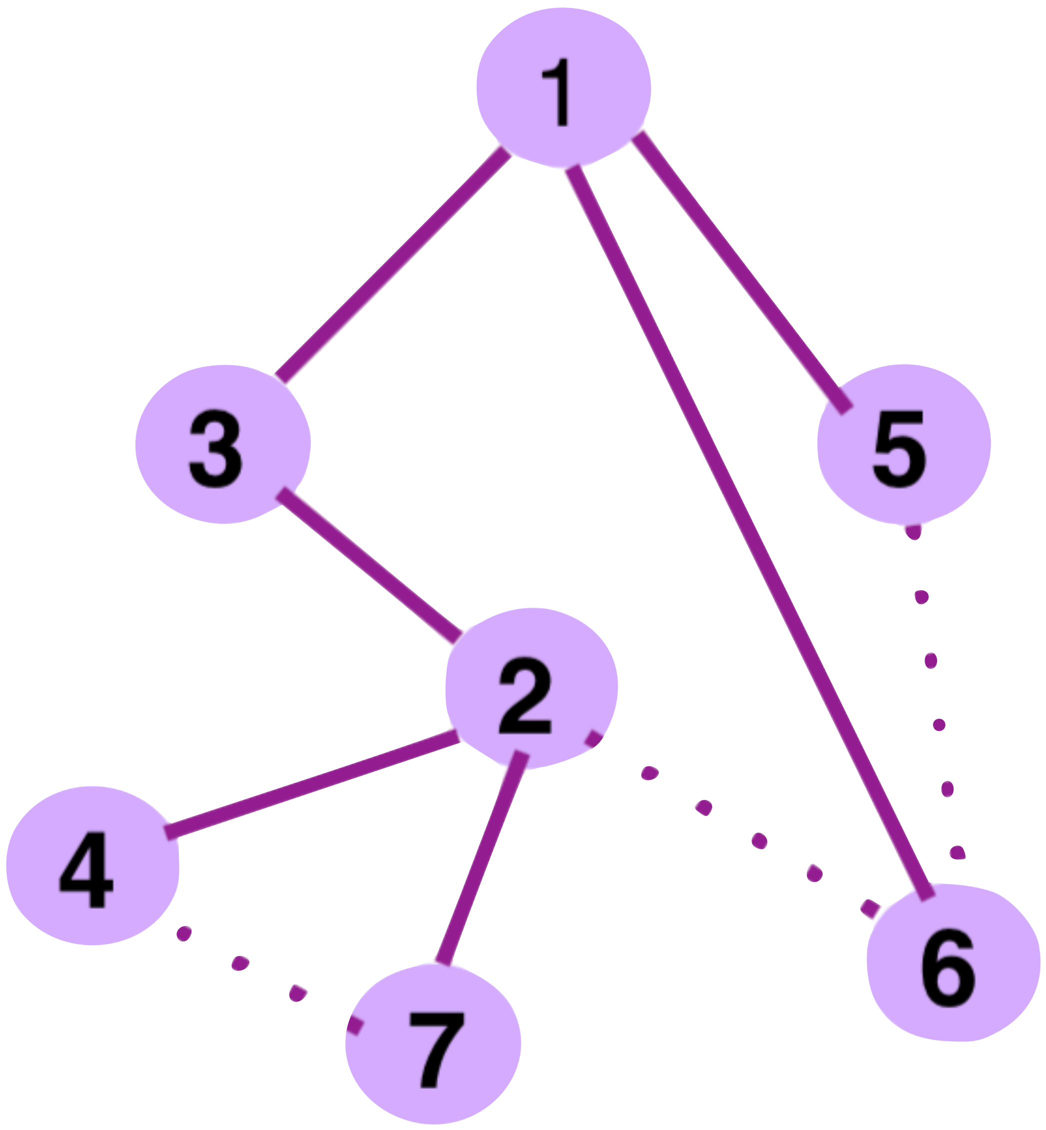

Flooding Example

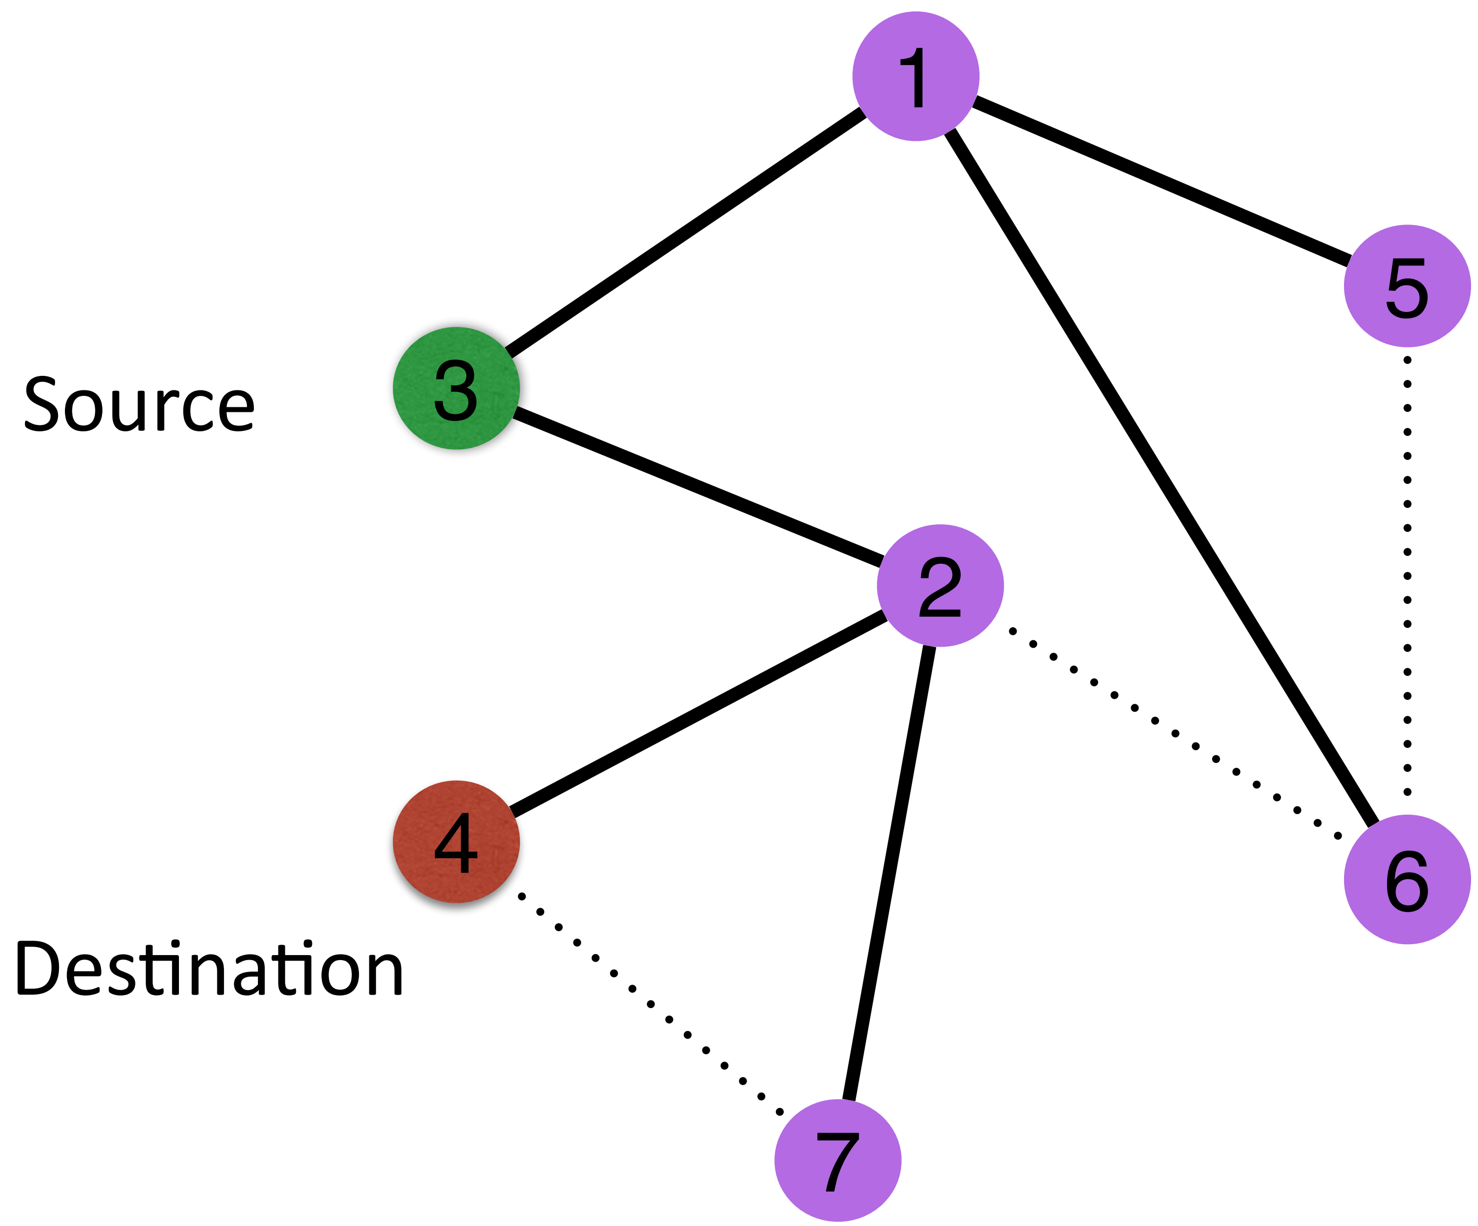

Flooding Example

Eventually all nodes are covered

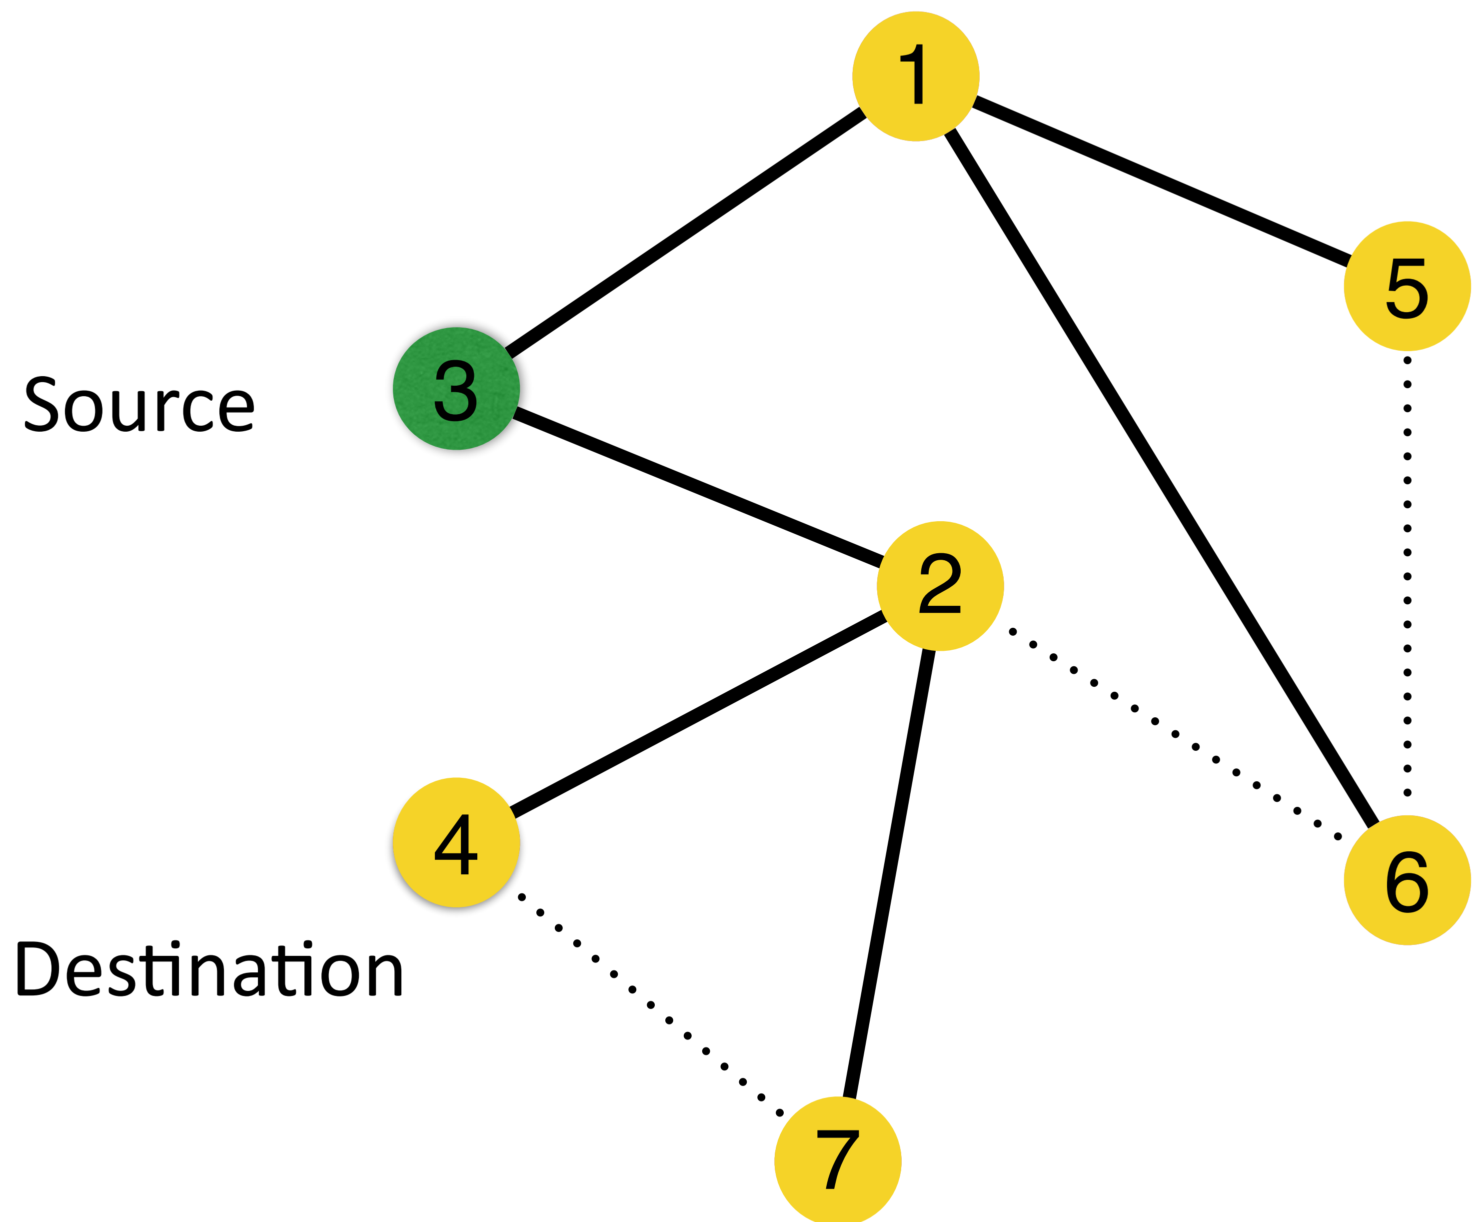

One copy of packet delivered to destination

Routing via Flooding on Spanning Tree …

- Easy to design routing algorithms for trees
  - **Nodes can "flood" packet to all other nodes**

- Amazing properties:
  - No routing tables needed!
  - No packets will ever loop.
  - At least (and exactly) one packet must reach the destination
    - Assuming no failures

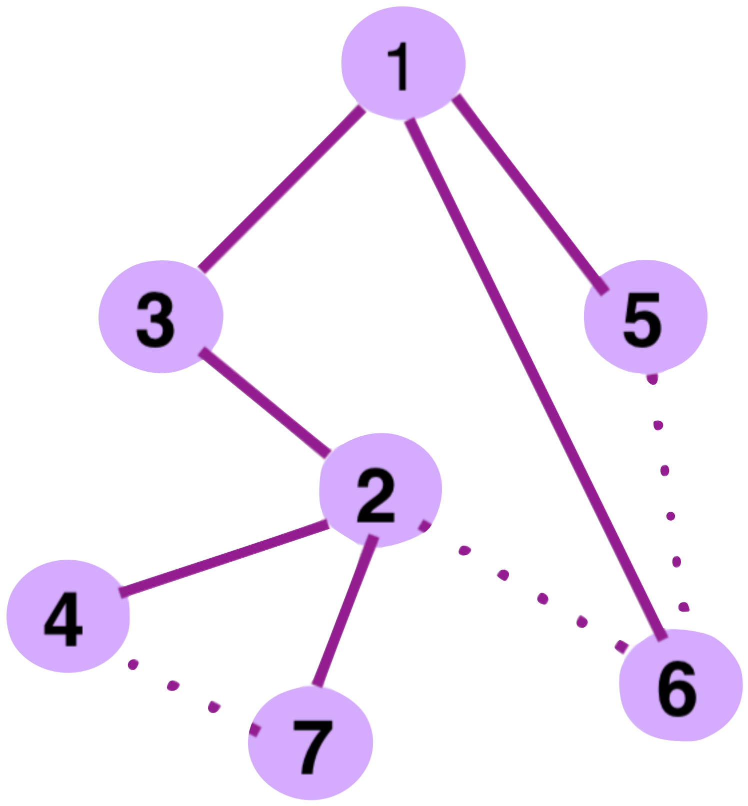

Three fundamental issues!

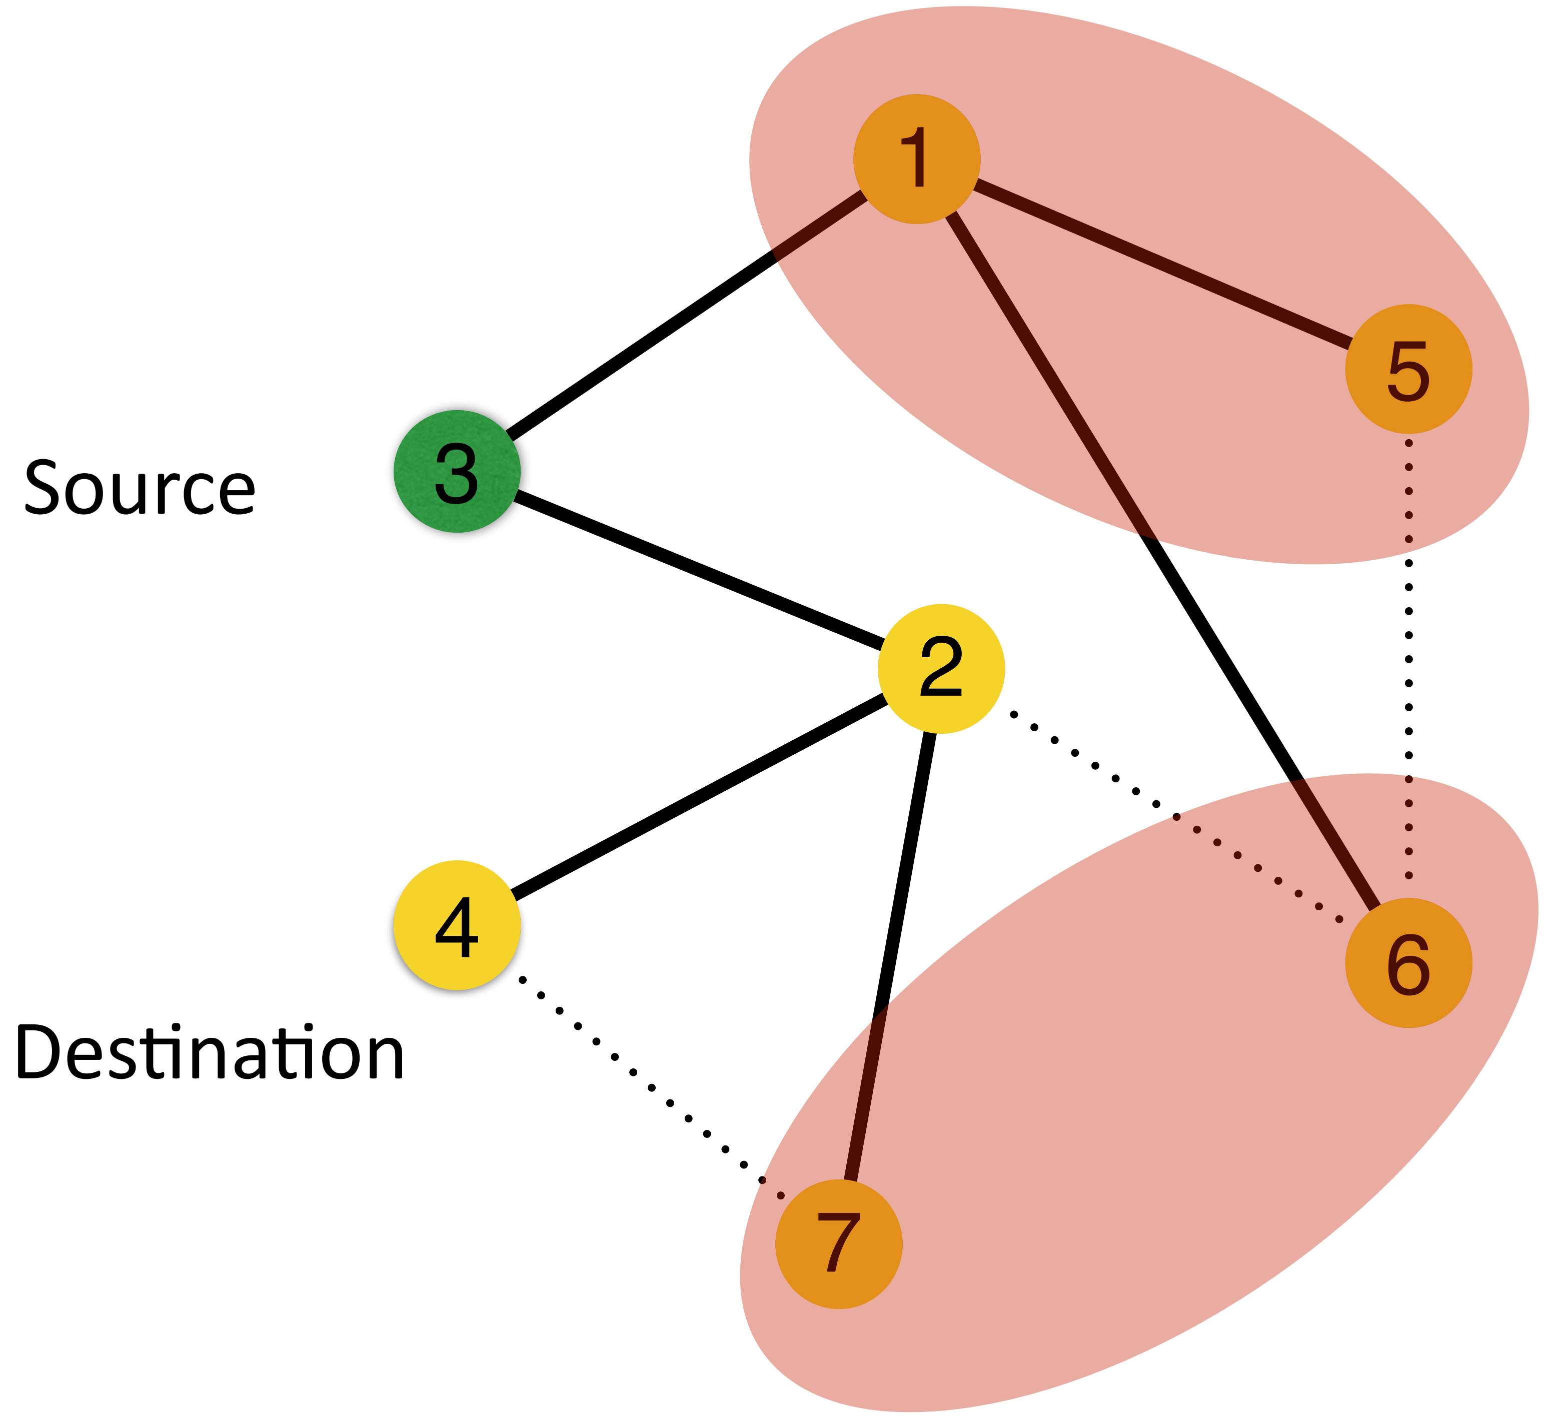

Issue 1: Each host has to do unnecessary packet processing!
(to decide whether the packet is destined to the host)

Three fundamental issues!

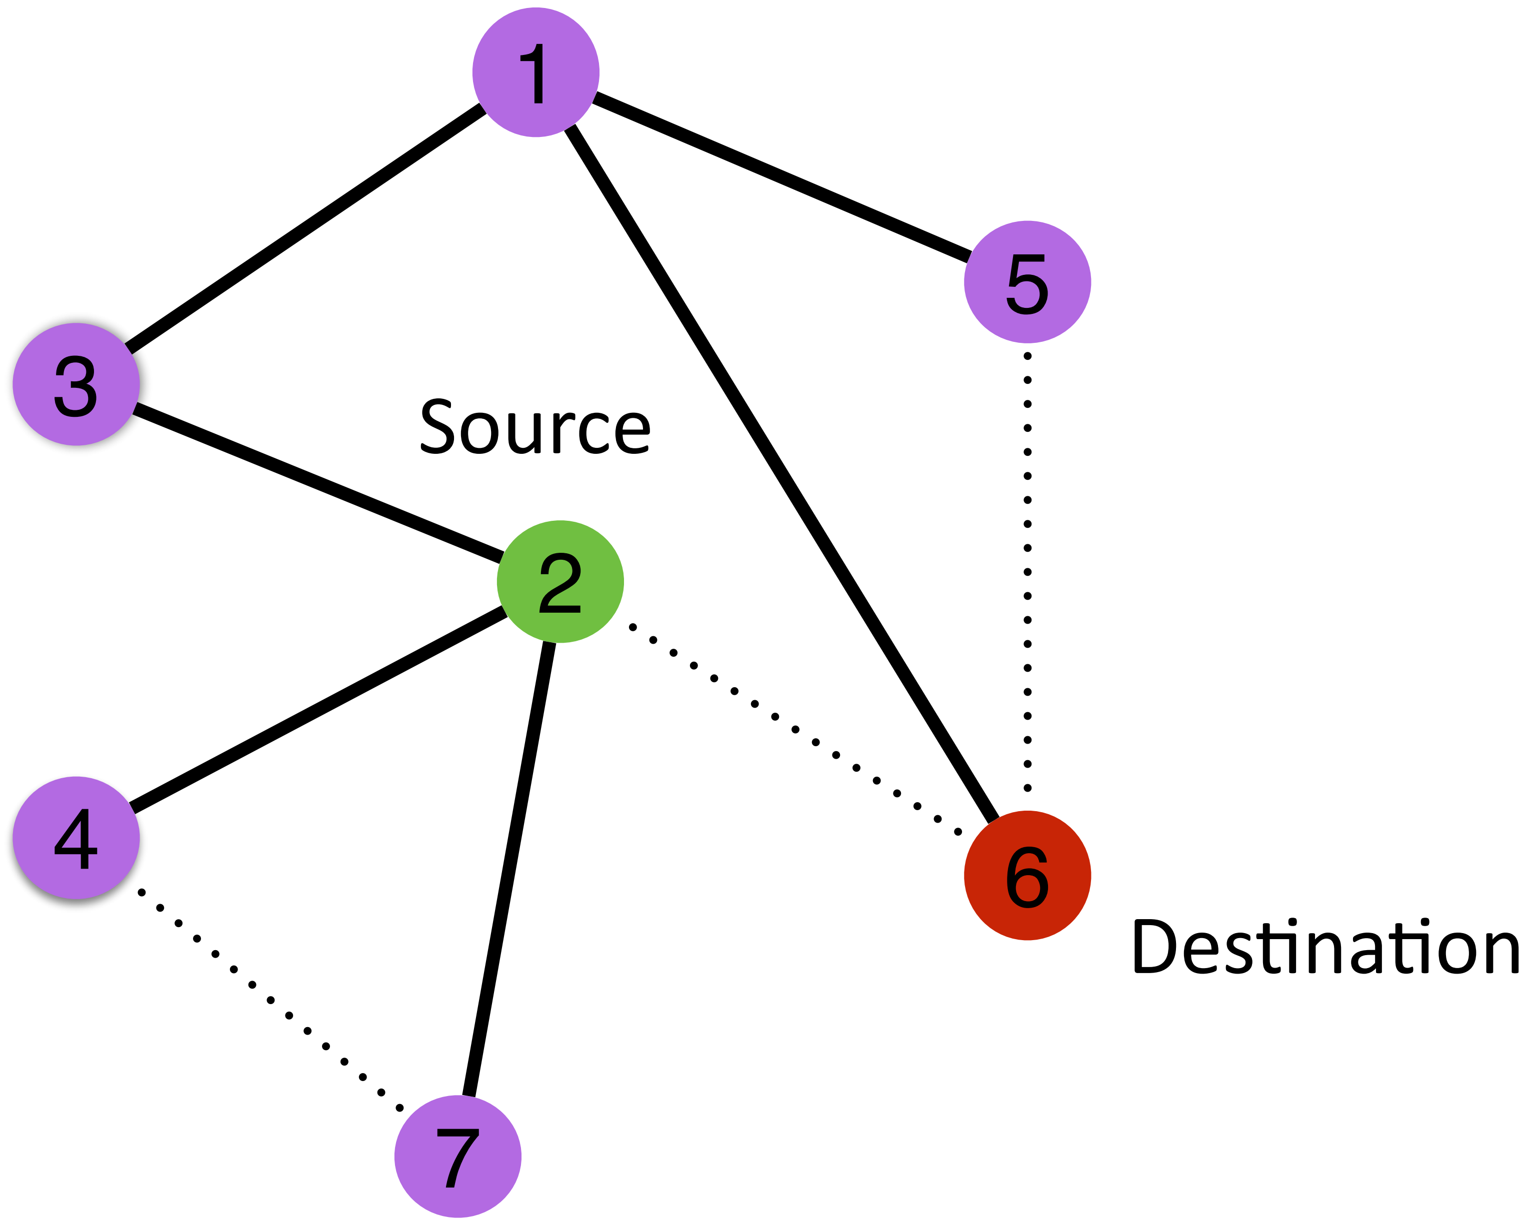

Issue 2: Higher latency!
(The packets unnecessarily traverse much longer paths)

Three fundamental issues!

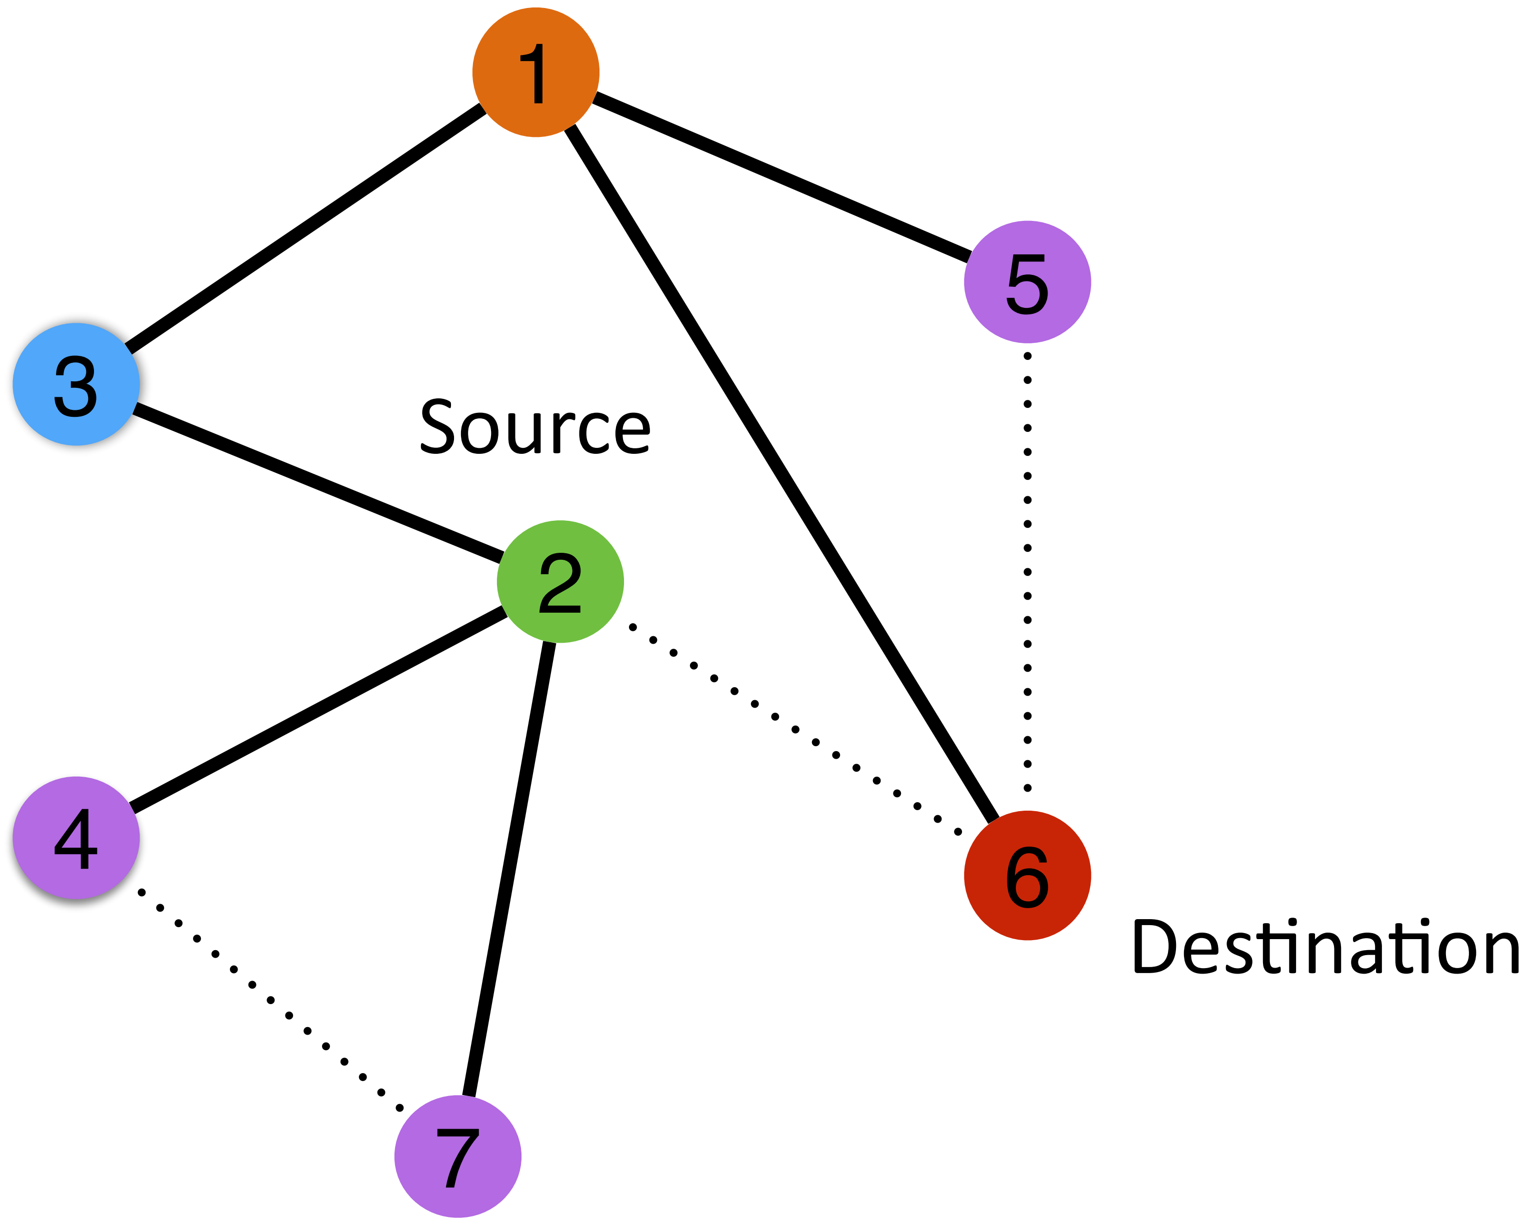

Issue 3: Lower bandwidth availability!
(2-6 and 3-1 packets unnecessarily have to share bandwidth)

Questions?

Why do we need a network layer?

- Network layer performs “routing” of packets to alleviate these issues
- Uses routing tables
- Lets understand routing tables first

Routing Packets via Routing Tables

- **Routing tables allow finding path from source to destination**

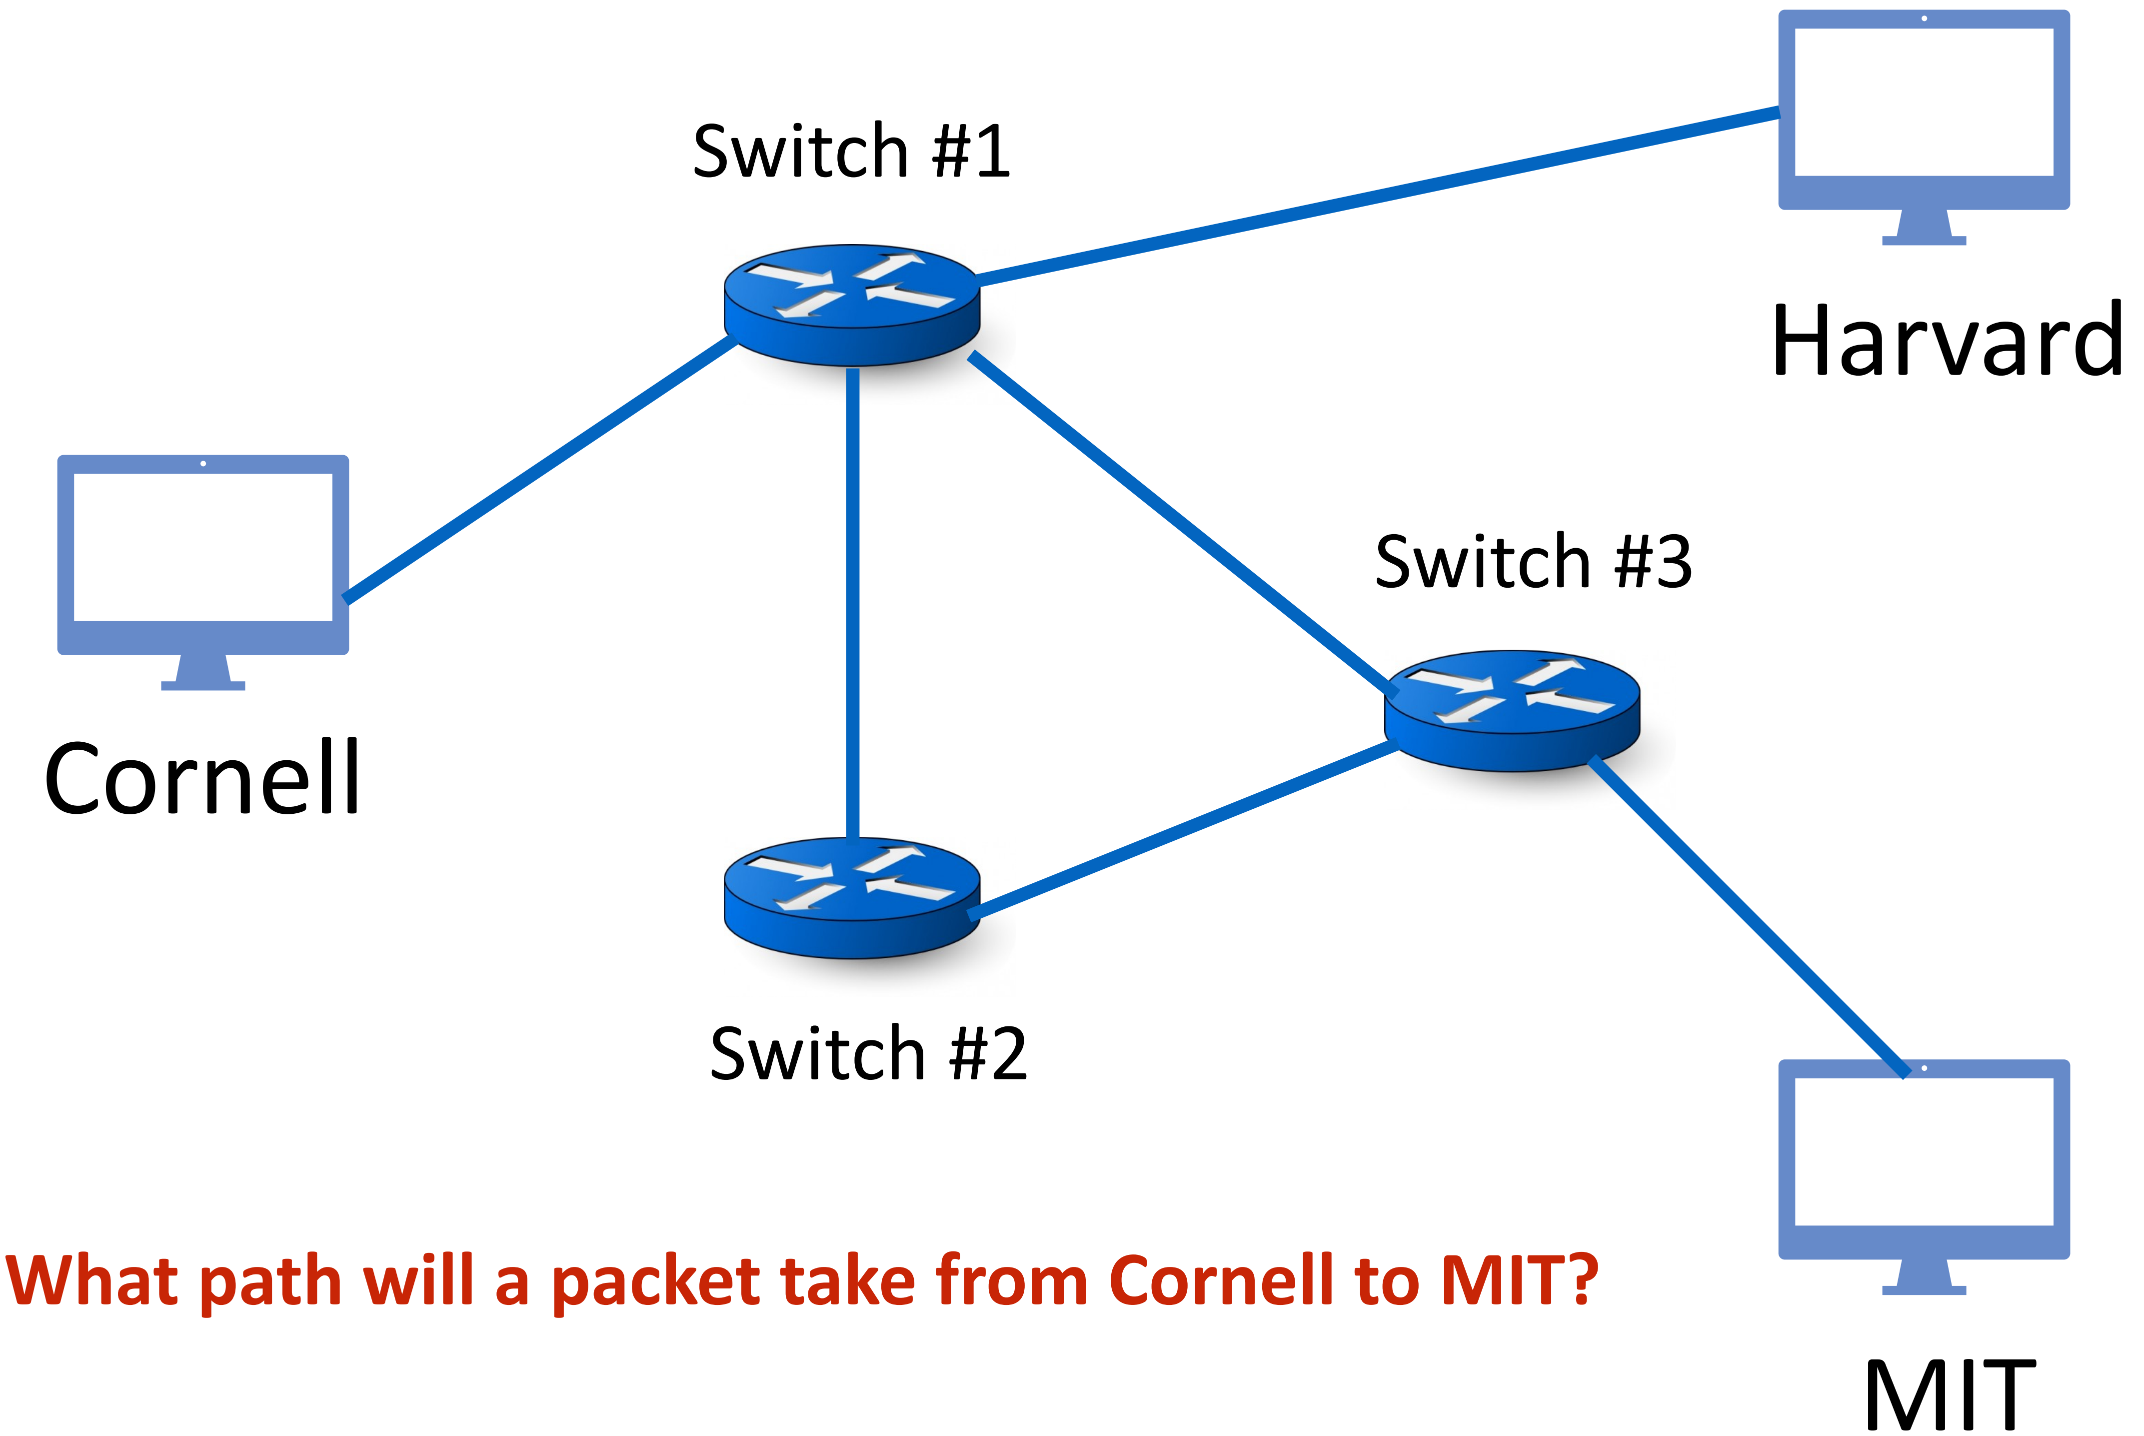

What path will a packet take from Cornell to MIT?

Routing Packets via Routing Tables

- **Finding path for a packet from source to destination**

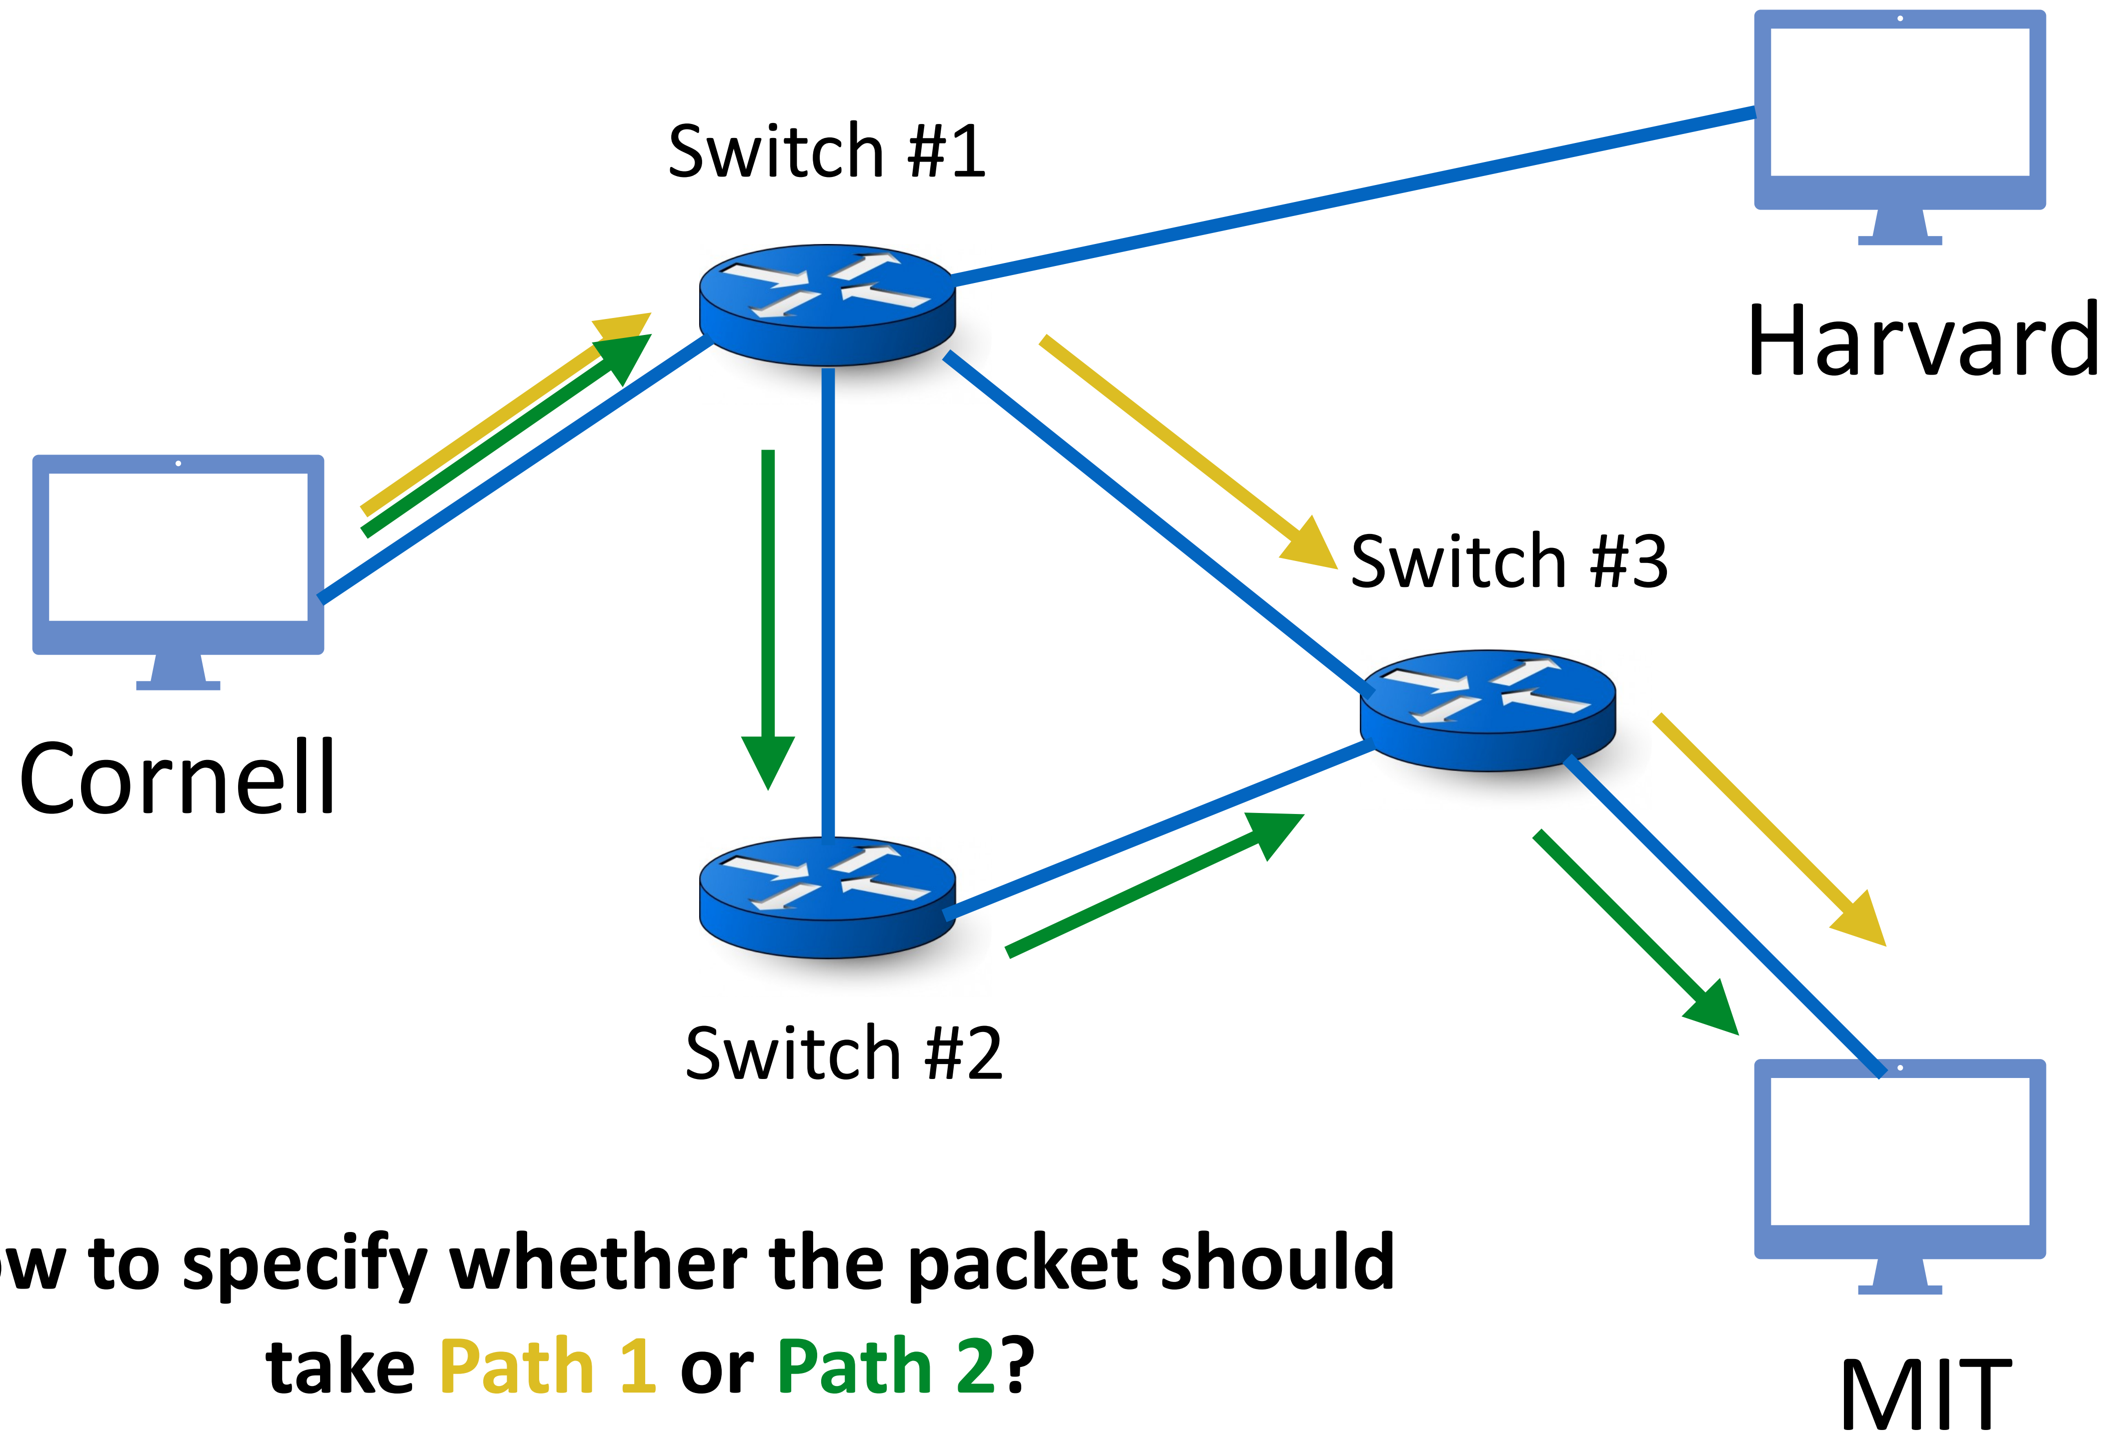

How to specify whether the packet should take Path 1 or Path 2?

Routing Table

- Suppose packet follows Path 1: Cornell - S#1 - S#3 - MIT

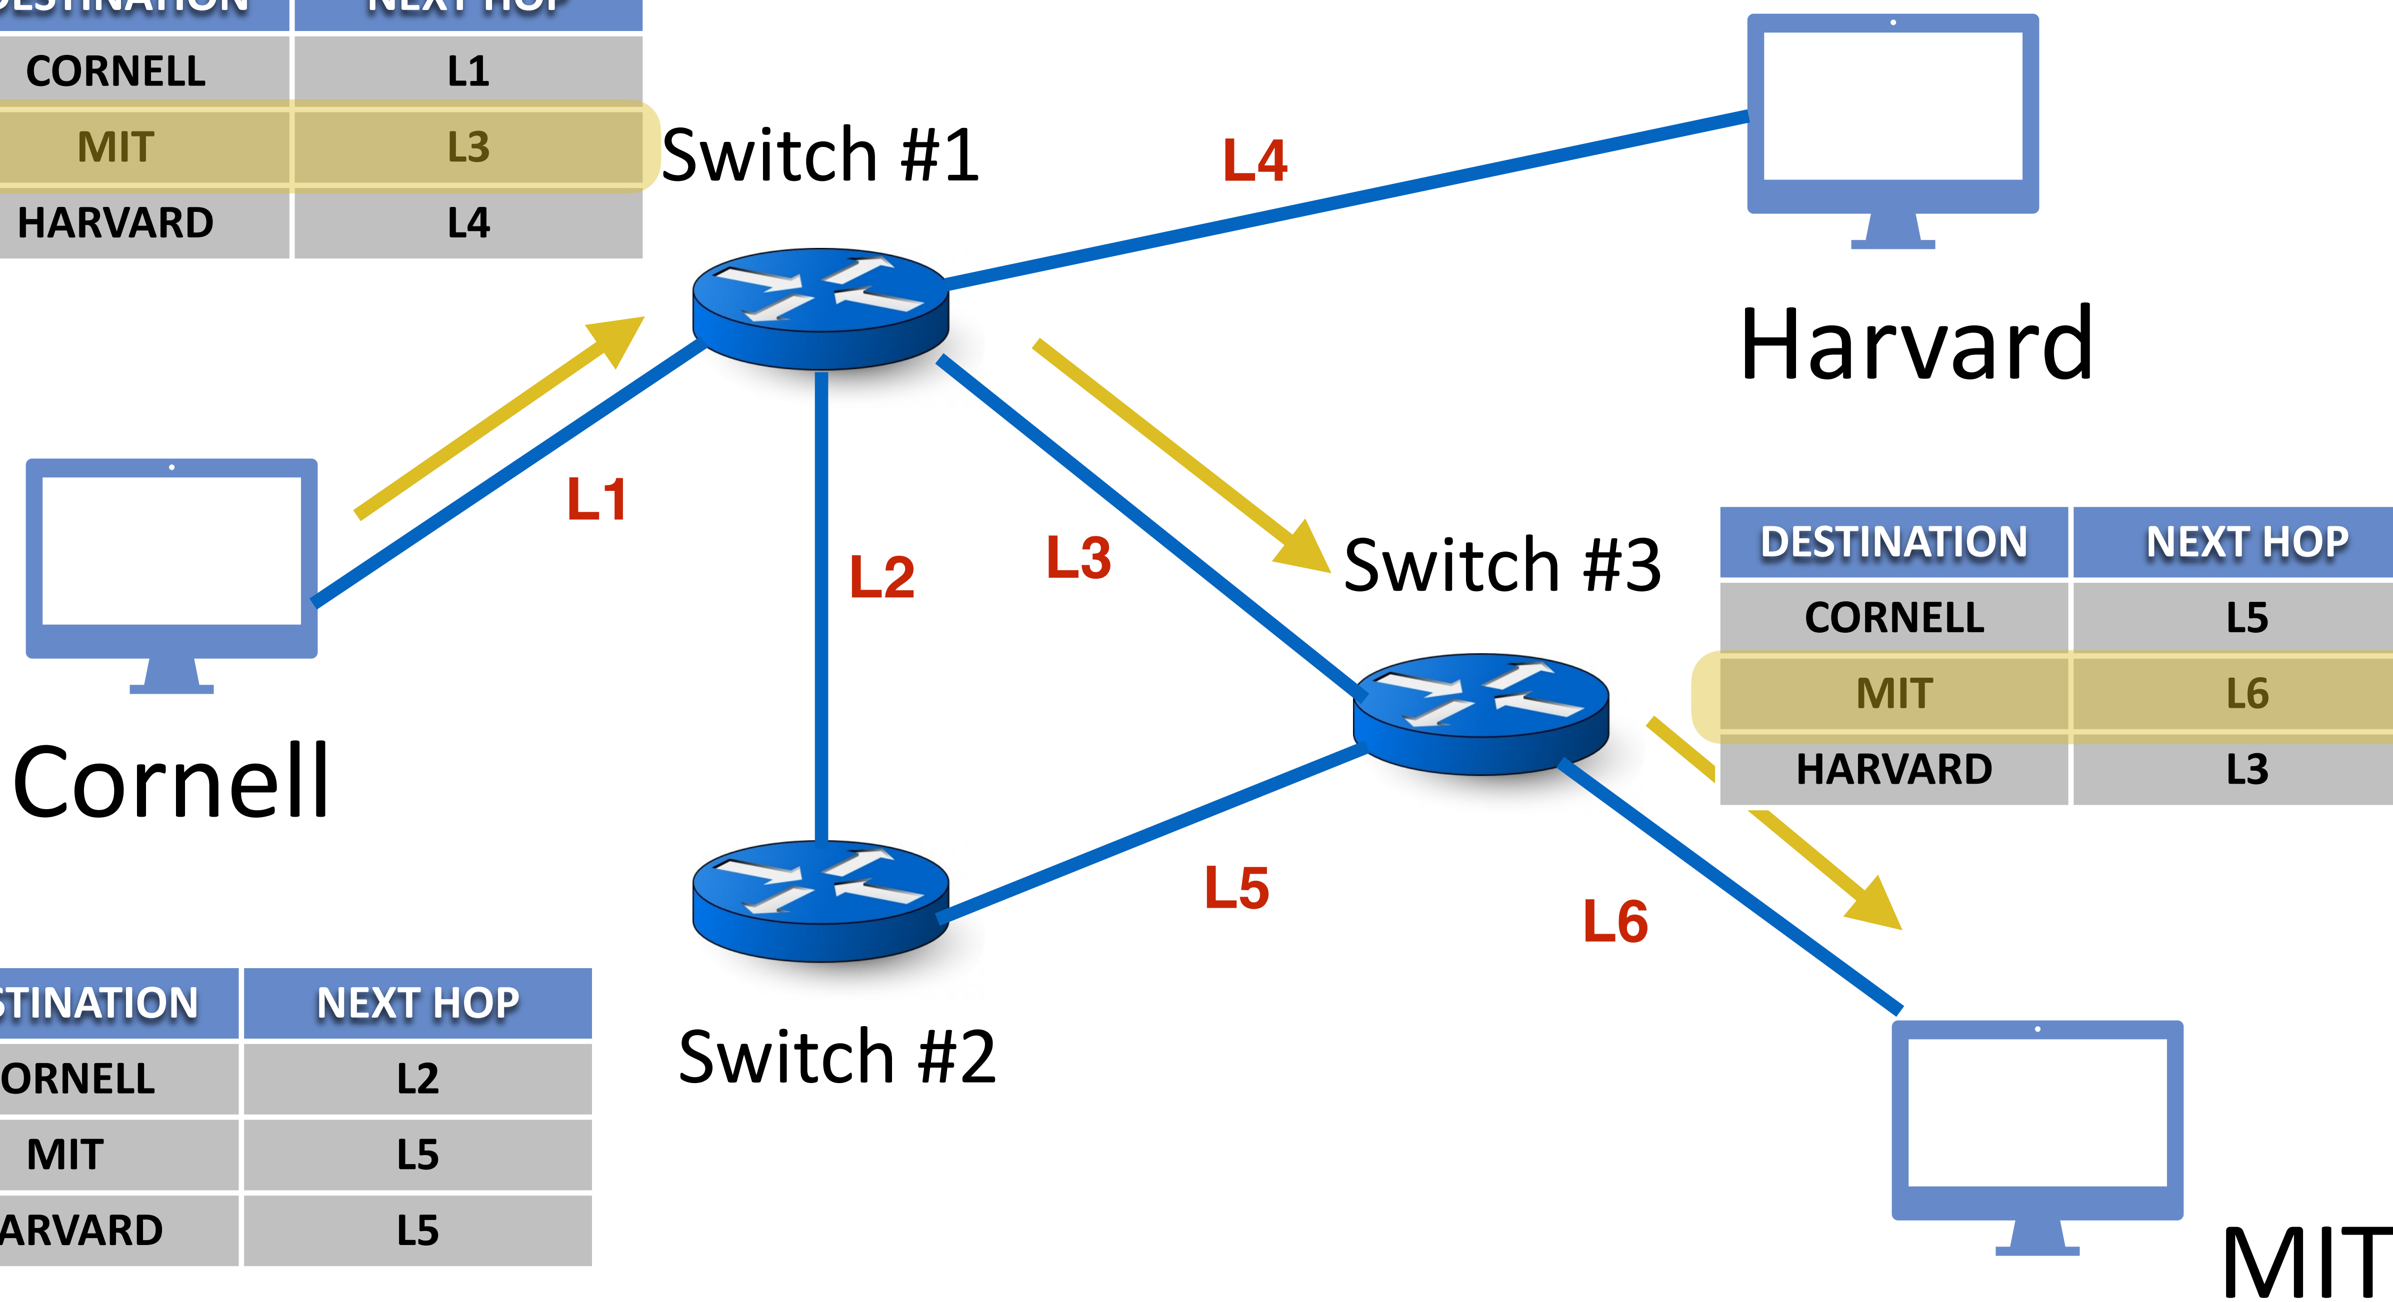

DESTINATION	NEXT HOP
CORNELL	L1
MIT	L3
HARVARD	L4

DESTINATION	NEXT HOP
CORNELL	L5
MIT	L6
HARVARD	L3

DESTINATION	NEXT HOP
CORNELL	L2
MIT	L5
HARVARD	L5

Each Switch stores a table indicating the next hop for corresponding destination of a packet (called a routing table)

Routing Table: The right way to think about them

- Lets focus on one destination - MIT

DESTINATION	NEXT HOP
CORNELL	L1
MIT	L3
HARVARD	L4

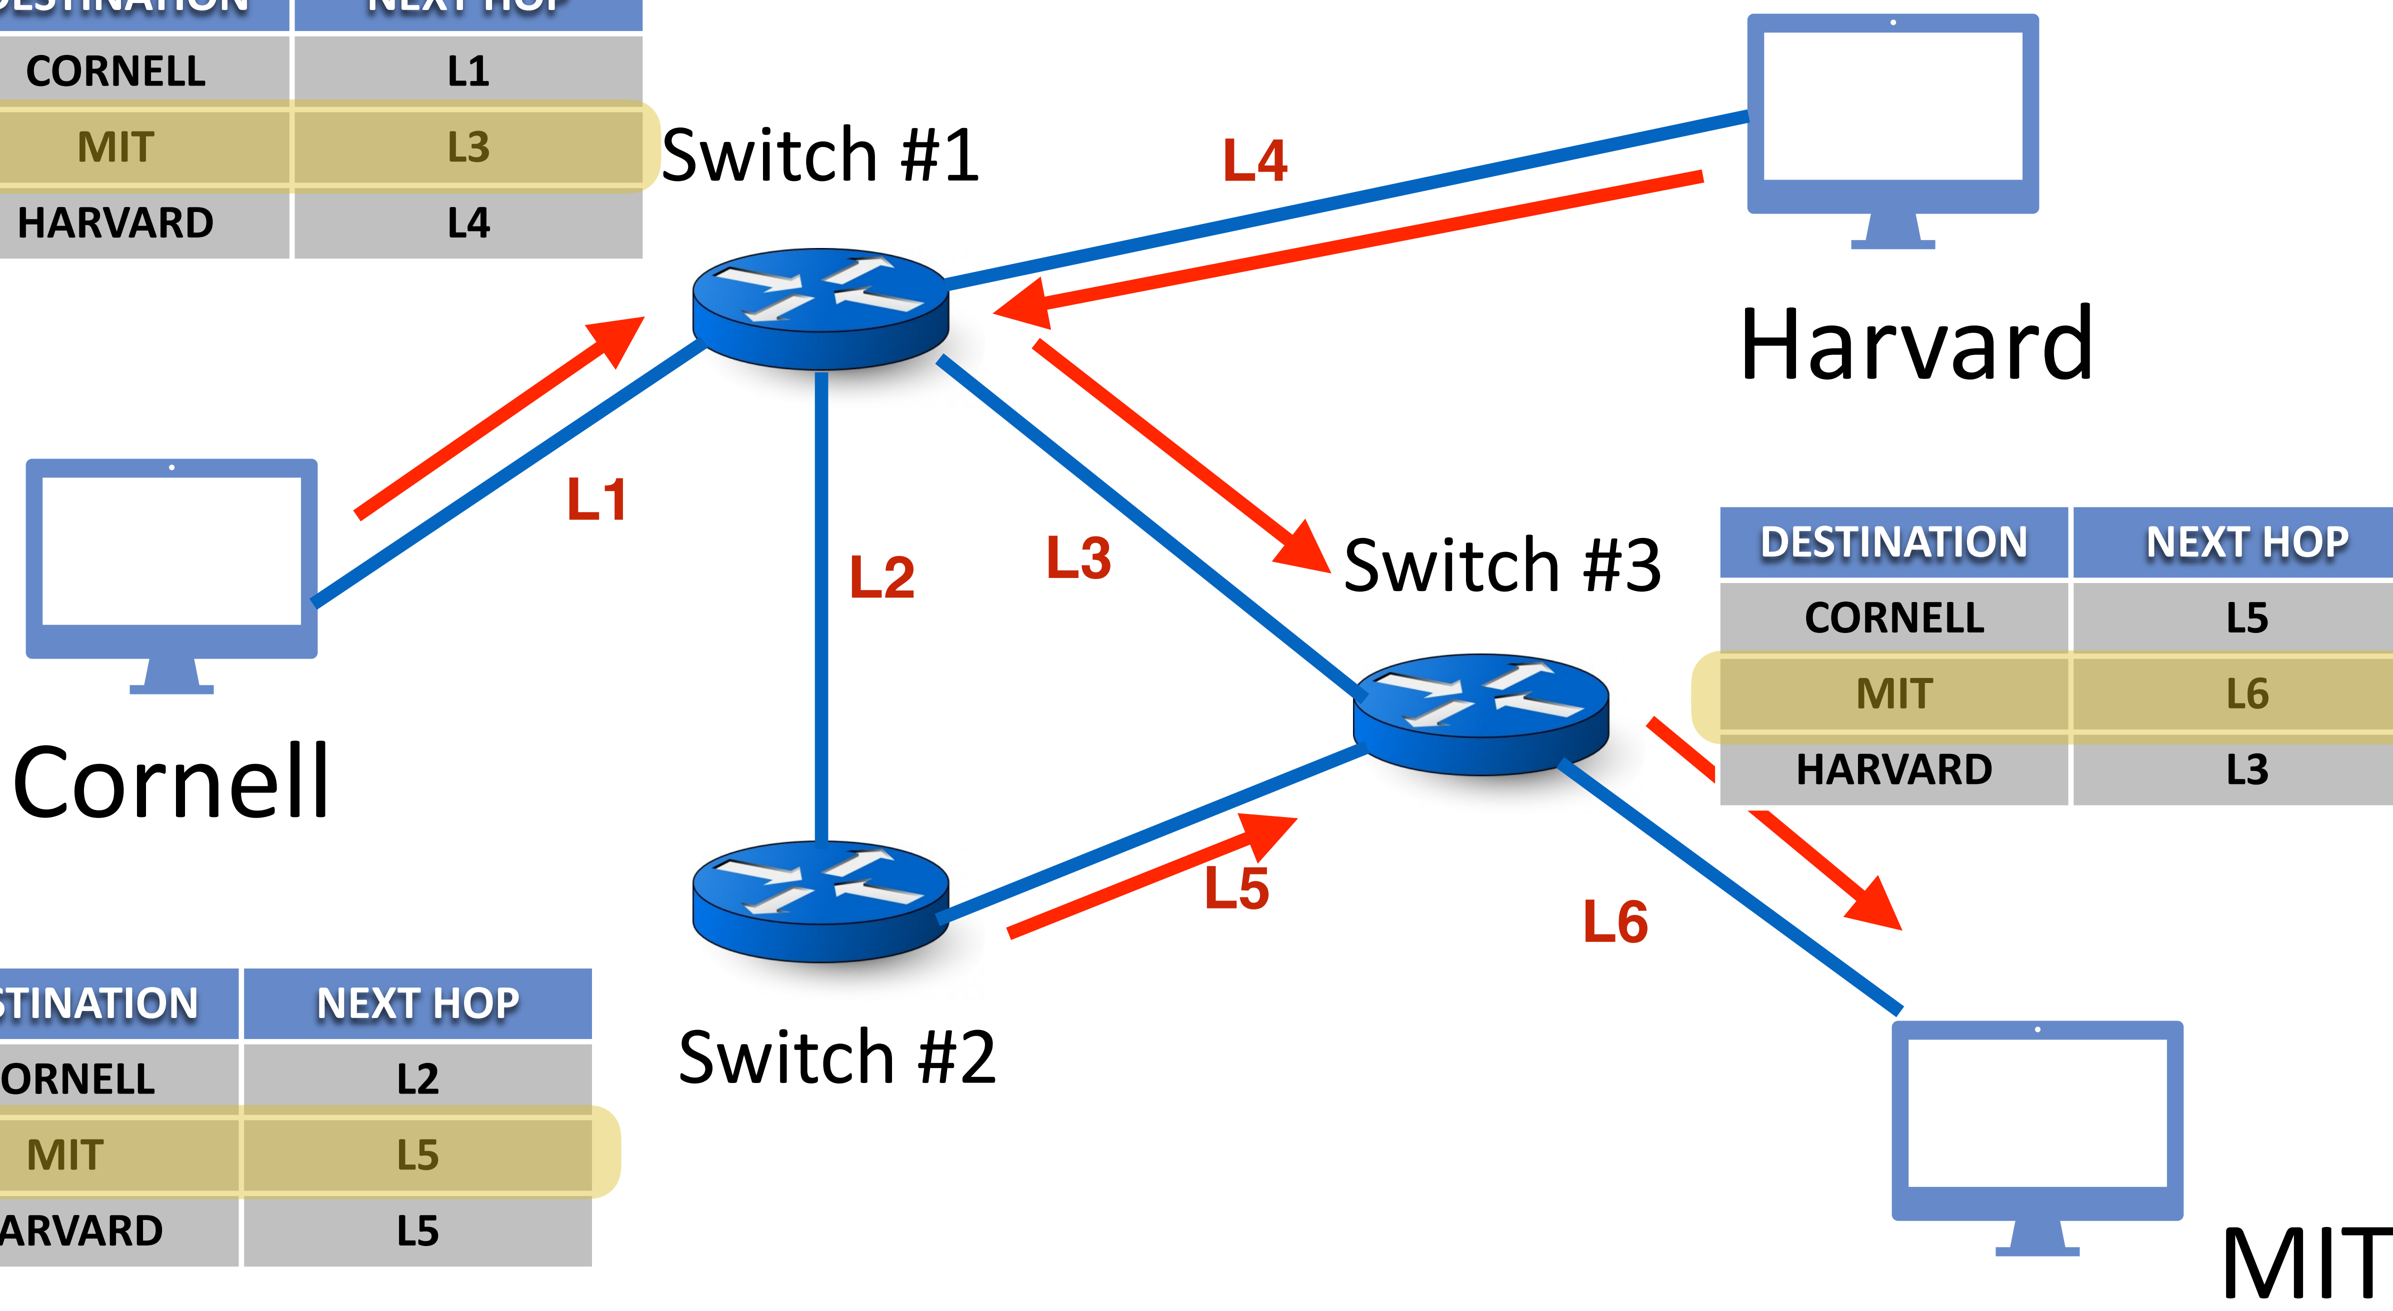

DESTINATION	NEXT HOP
CORNELL	L5
MIT	L6
HARVARD	L3

DESTINATION	NEXT HOP
CORNELL	L2
MIT	L5
HARVARD	L5

See something interesting?

Routing Table: The right way to think about them

- Lets focus on one destination - MIT

DESTINATION	NEXT HOP
CORNELL	L1
MIT	L3
HARVARD	L4

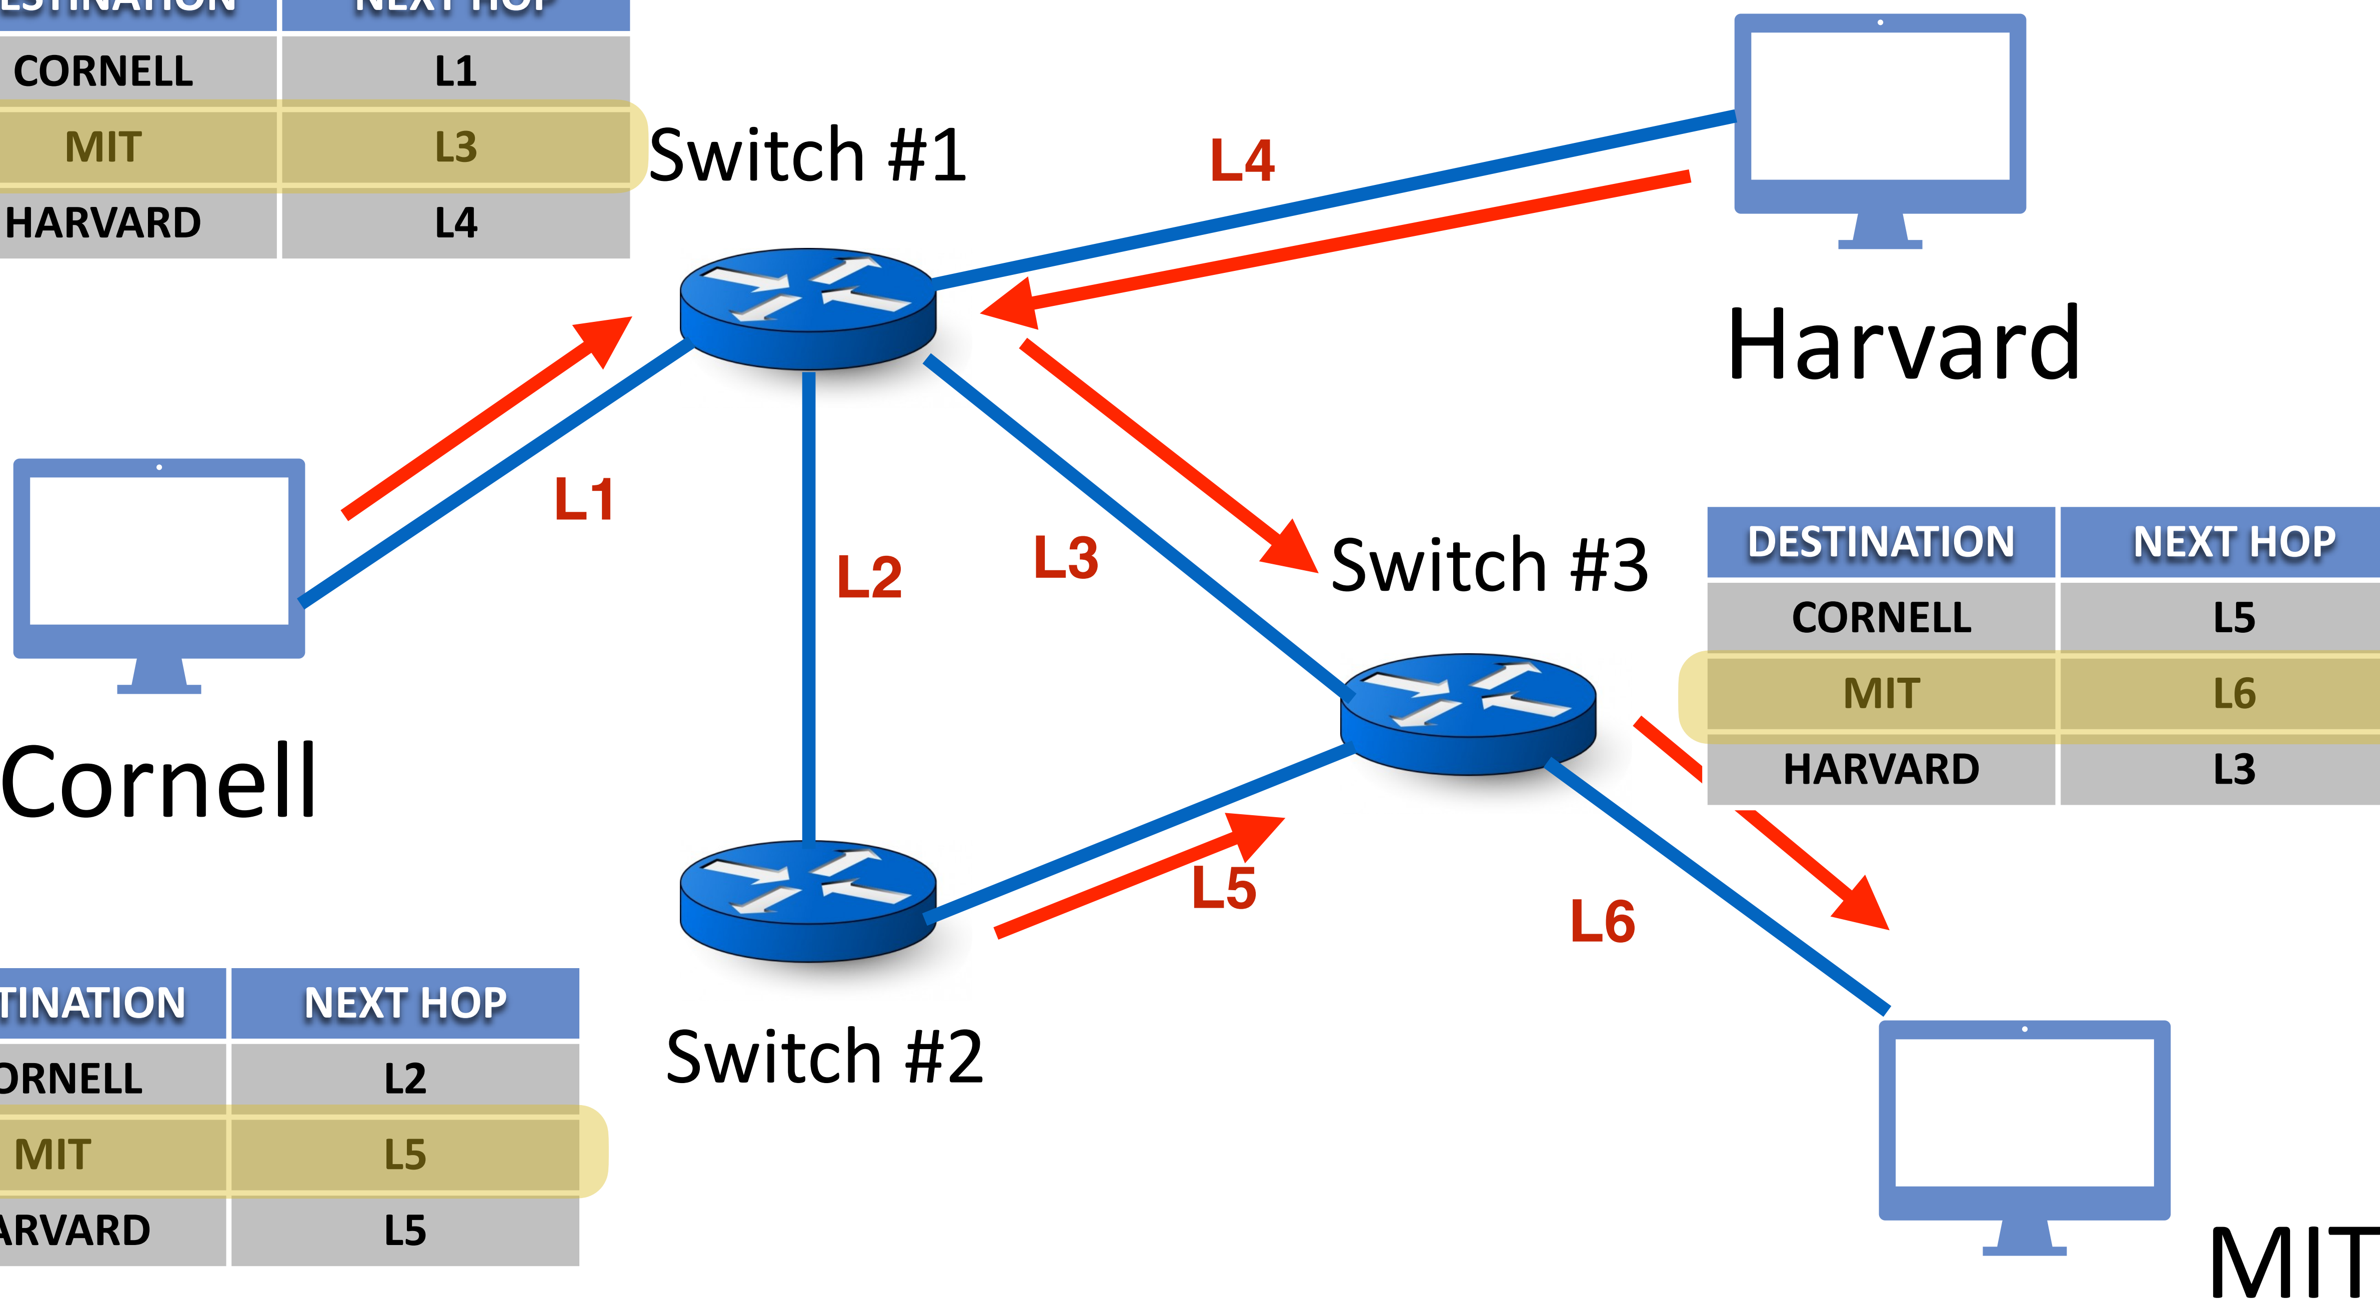

DESTINATION	NEXT HOP
CORNELL	L5
MIT	L6
HARVARD	L3

DESTINATION	NEXT HOP
CORNELL	L2
MIT	L5
HARVARD	L5

Routing table entries for a particular destination form a (directed) spanning tree with that destination as the root!!!!

Routing Table: The right way to think about them

- Routing tables are nothing but ….
  - A collection of (directed) spanning tree
  - One for each destination
- **Routing Protocols**
  - "n" spanning tree protocols running in parallel

Next lecture!