Final-exam Review

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

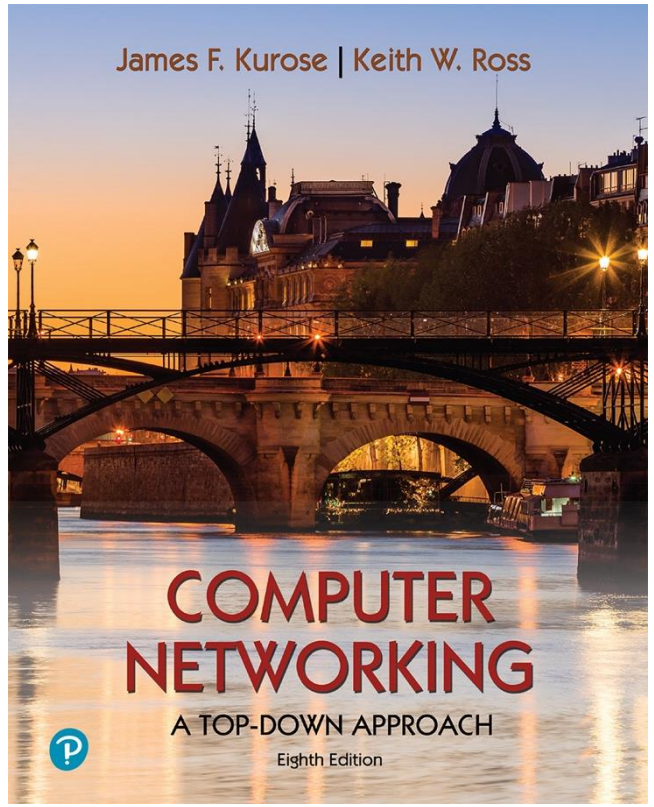

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Adapted from the slides of the book's authors

Time, Place, and Bonus Question

- May 7th, Thursday Knox Hall 110
- 150 mins, so you definitely have enough time
- There will be one bonus question (10 points) about the PA2

Several Points

- Here, I will list all the ***topics*** that I think are important
  - If one topic I didn't mention, then I won't test it
  - It is about the topic, not the slides
    - For each topic, I will try to cover those very critical points
    - If I didn't mention one slides, but I do mention the topic, I probably will cover it
    - There are too many slides if I include every slides about that topic
- It will be fast, I won't teach it again
- Ask questions, if you have any

Chapter 4
Network Layer: Data Plane

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

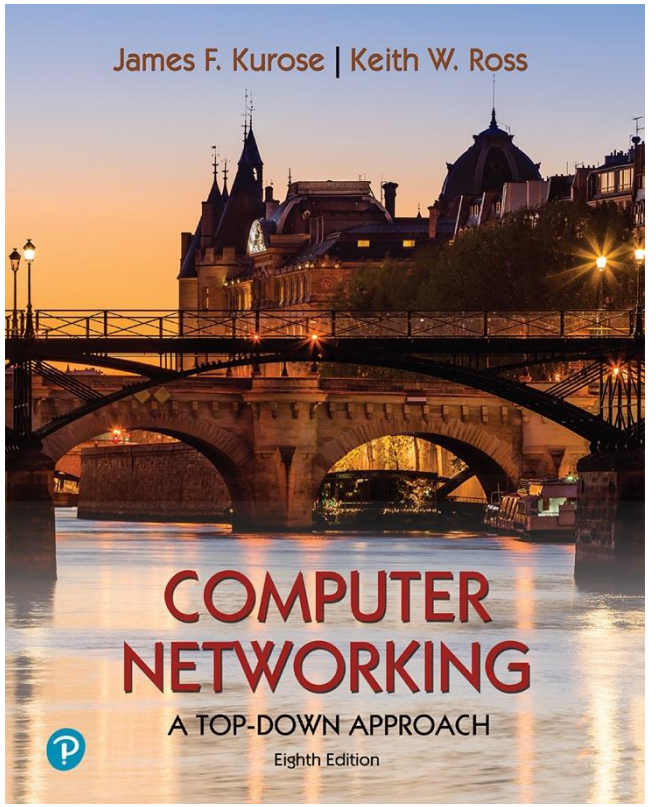

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Adapted from the slides of the book's authors

Network layer: data plane, control plane

Data plane:

- *local*, per-router function
- determines how datagram arriving on router input port is forwarded to router output port

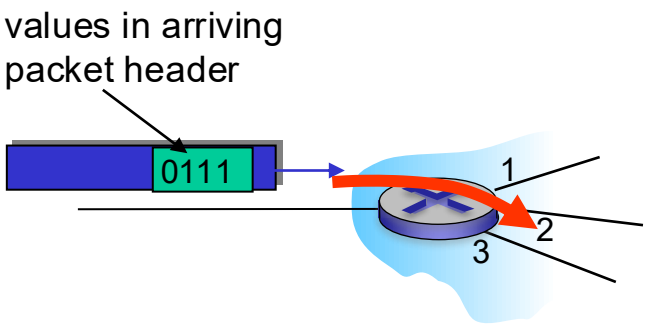

Control plane

- *network-wide* logic
- determines how datagram is routed among routers along end-end path from source host to destination host
- two control-plane approaches:
  - *traditional routing algorithms:* implemented in routers
  - *software-defined networking (SDN)*: implemented in (remote) servers

Per-router control plane

Individual routing algorithm components *in each and every router* interact in the control plane

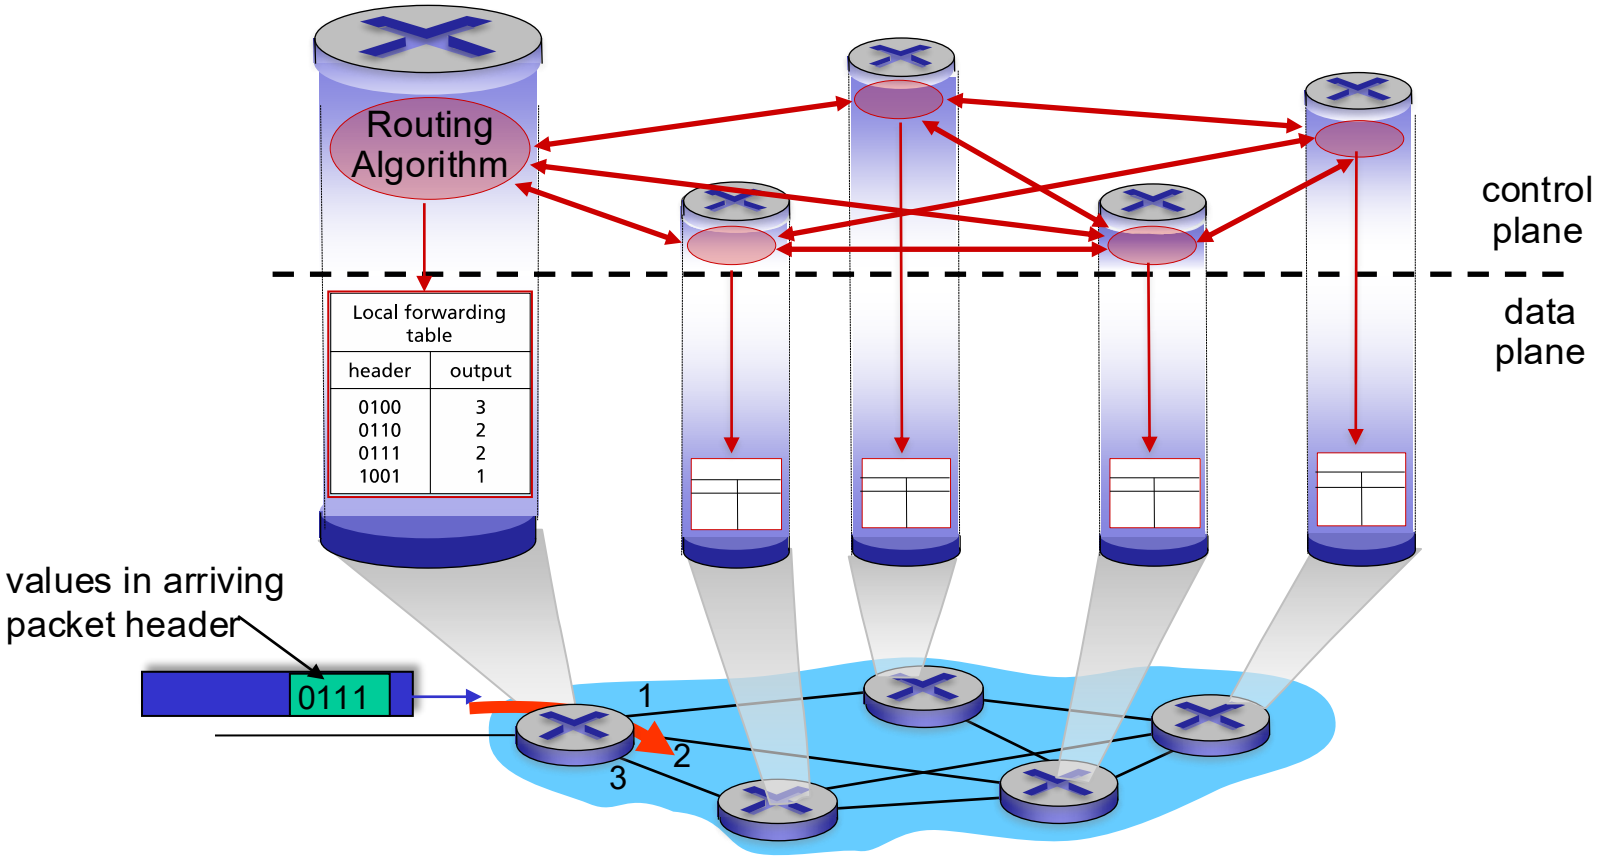

Software-Defined Networking (SDN) control plane

Remote controller computes, installs forwarding tables in routers

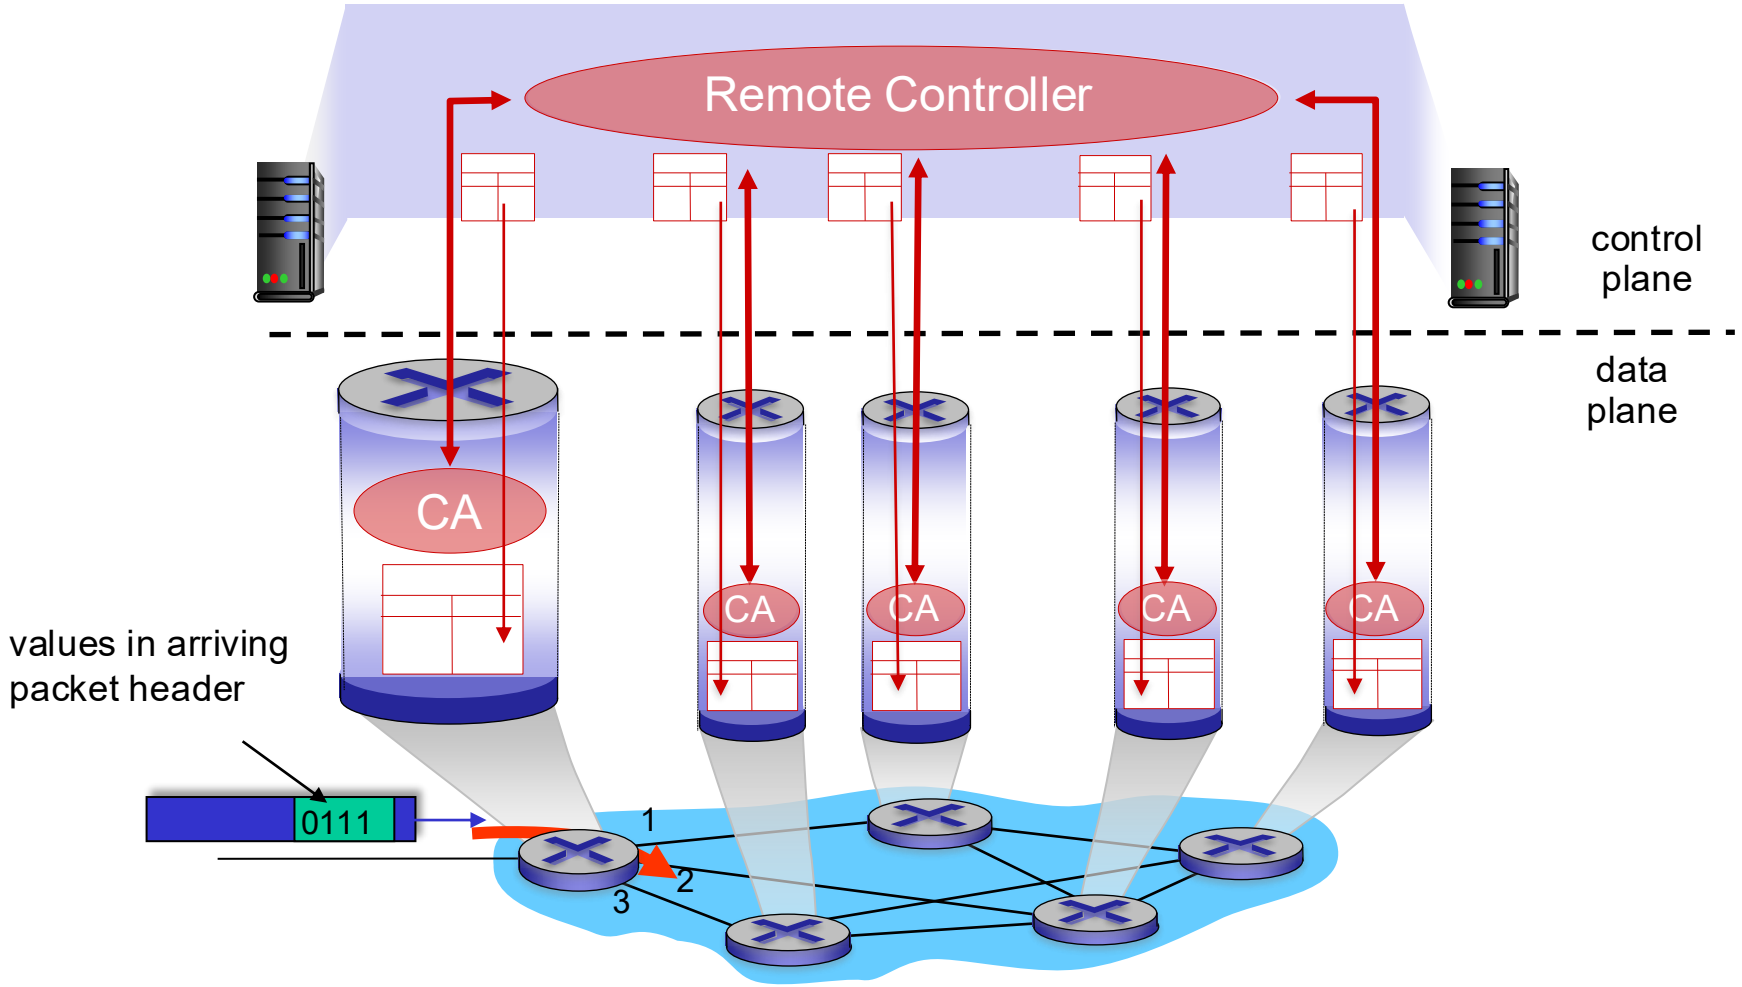

Application and transport layer is end-to-end

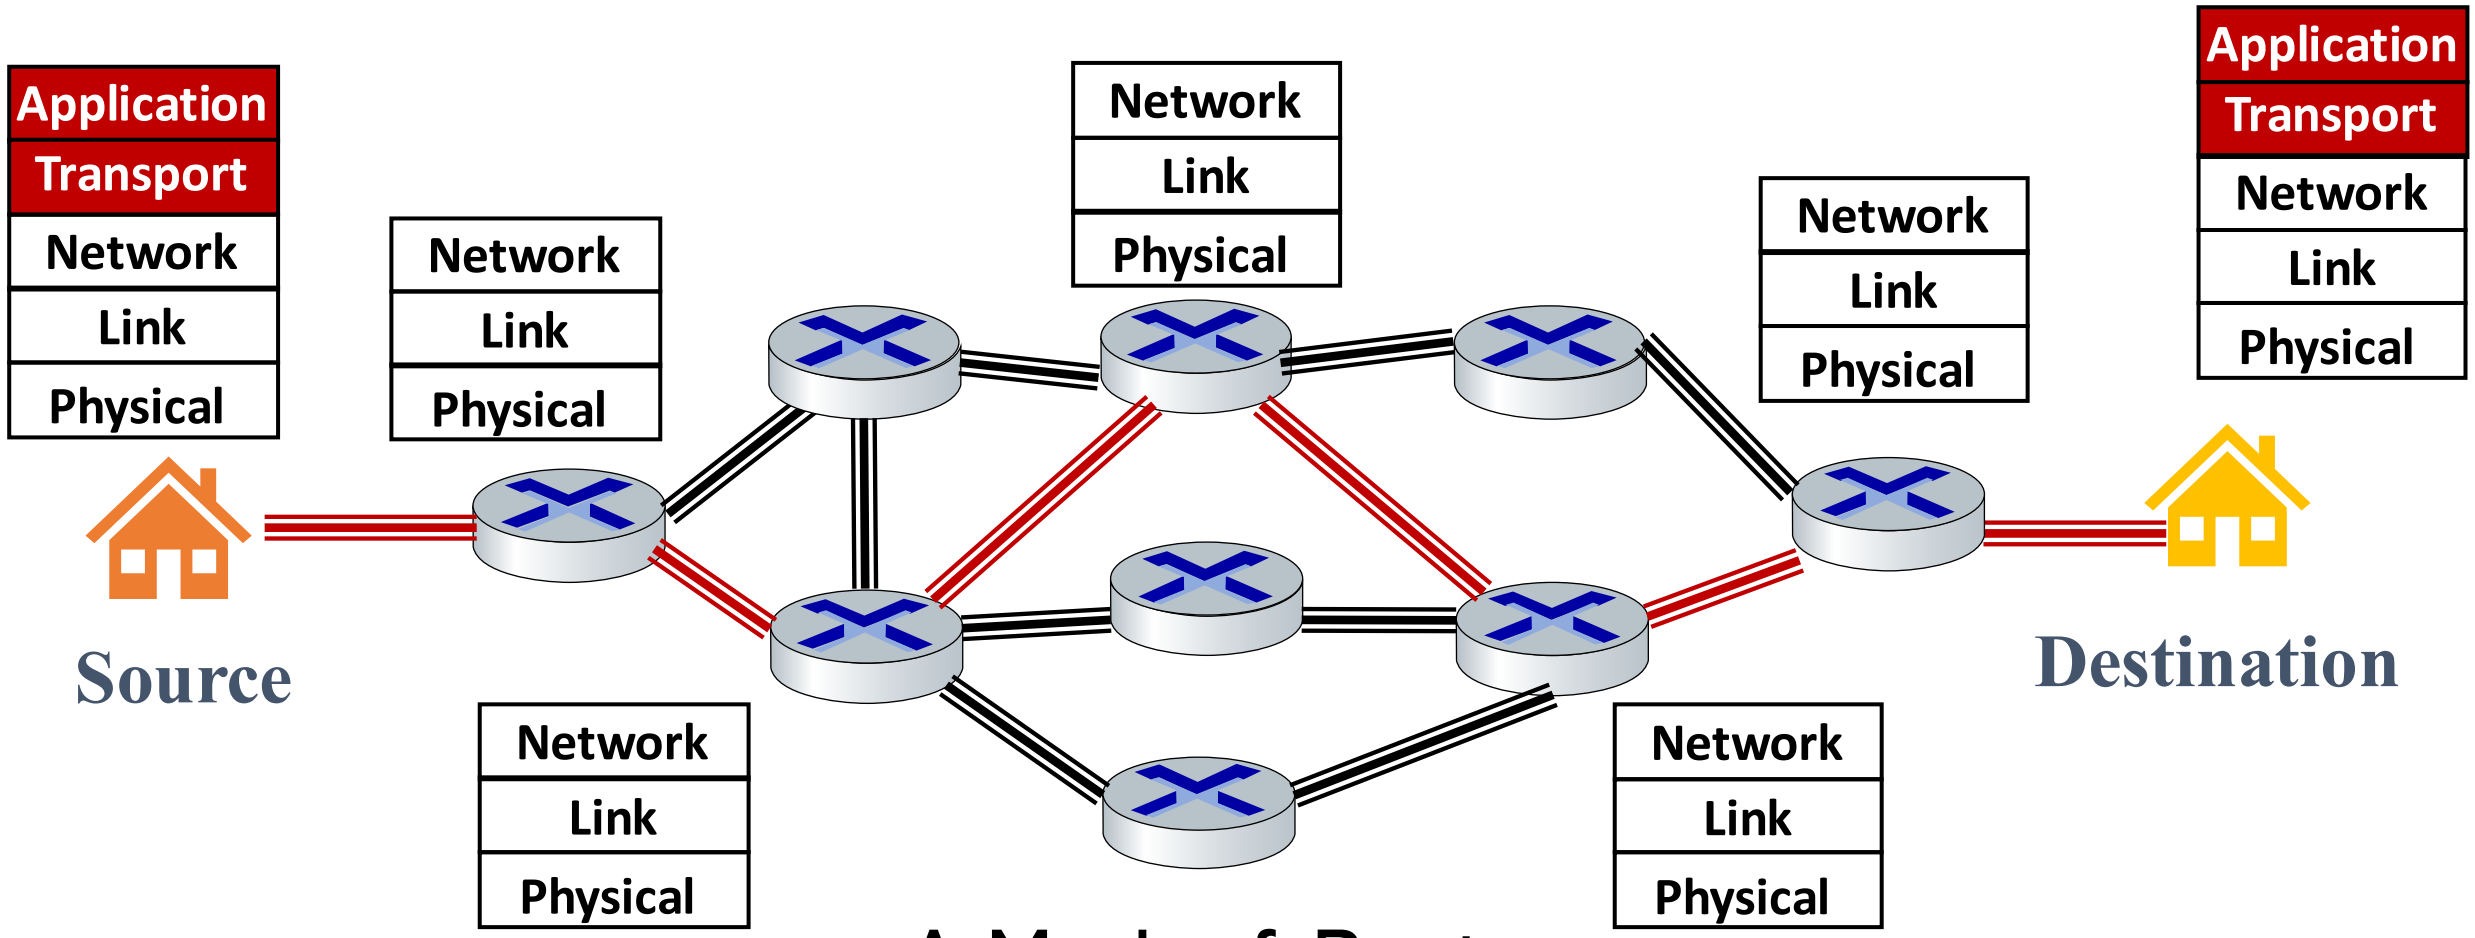

Network-layer is in every network device

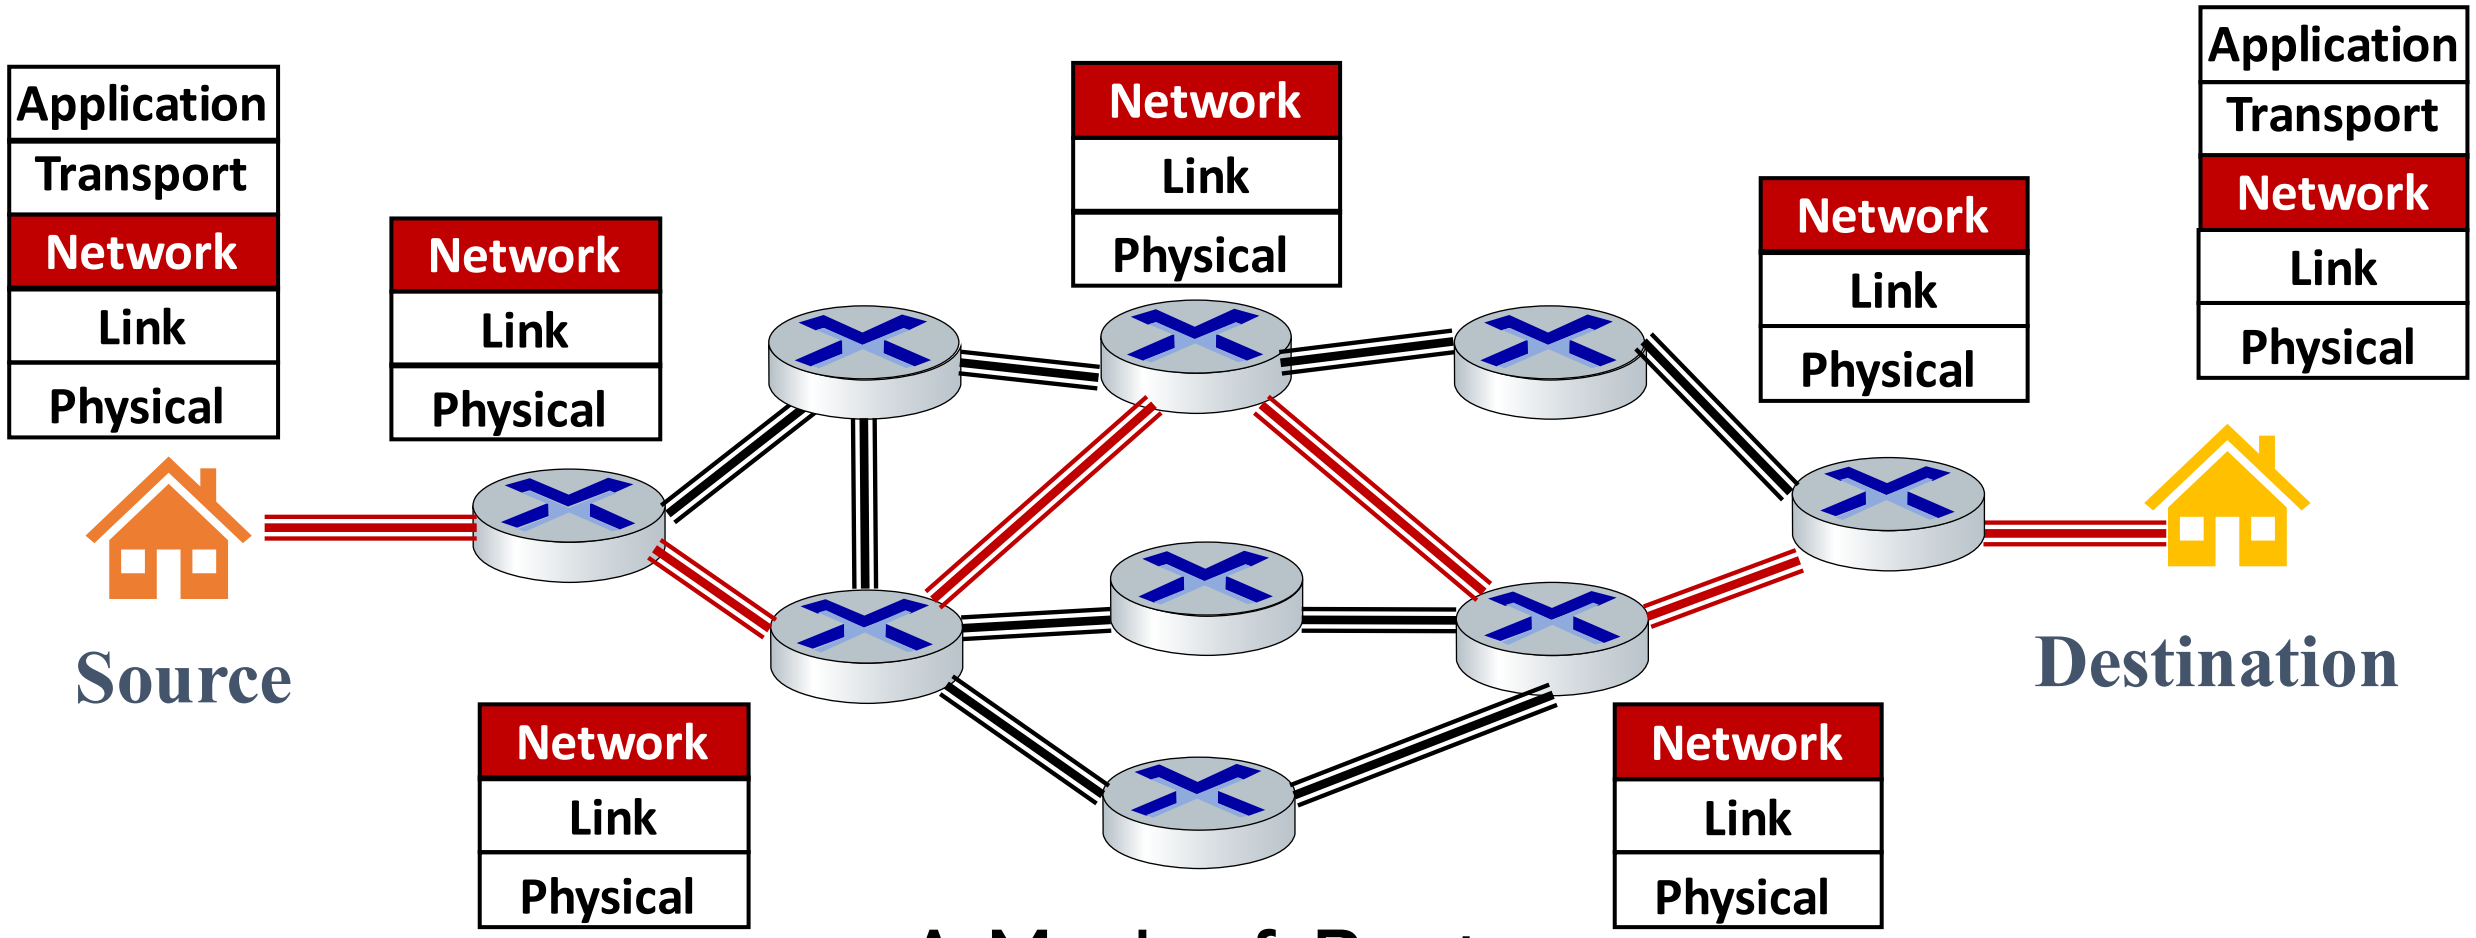

Network layer: "data plane" roadmap

- Network layer: overview
  - data plane
  - control plane
- What's inside a router
  - input ports, switching, output ports
  - buffer management, scheduling
- IP: the Internet Protocol
  - datagram format
  - addressing
  - network address translation
  - IPv6

- Generalized Forwarding, SDN
  - Match+action
  - OpenFlow: match+action in action
- Middleboxes

Packet Scheduling

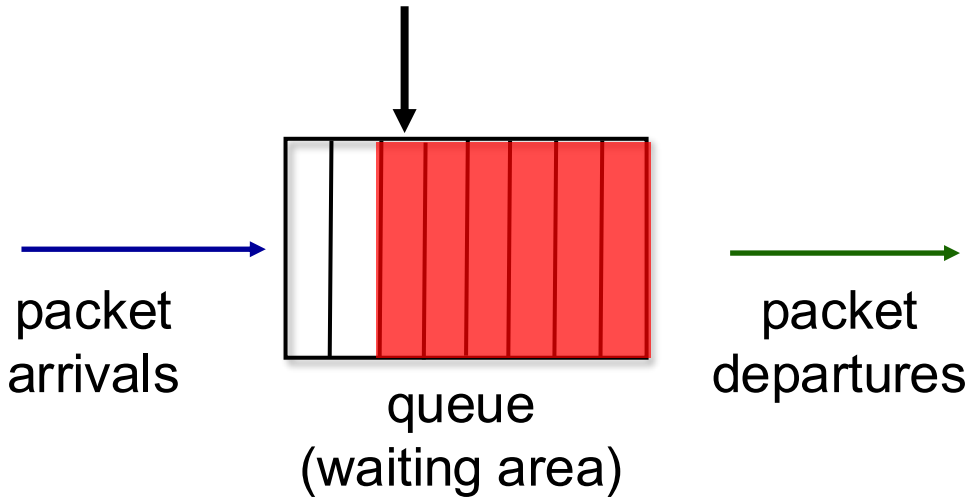

Packet Scheduling

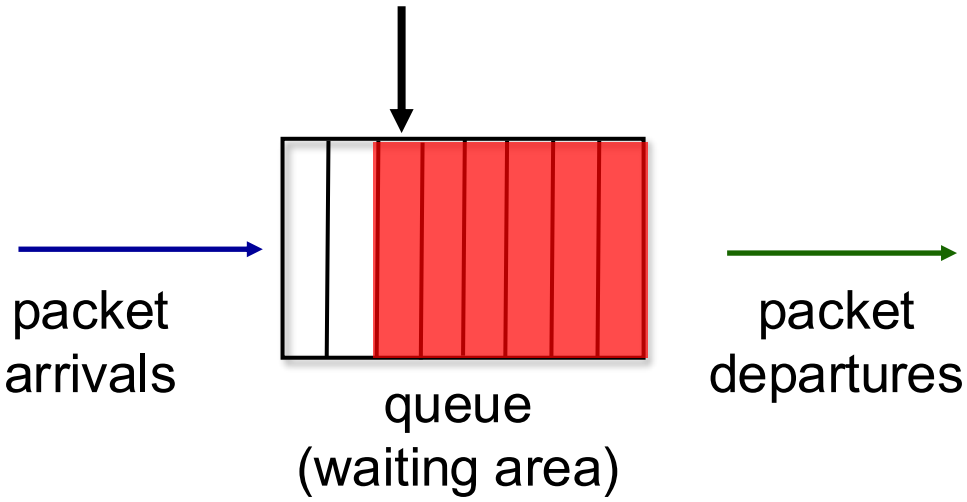

packet scheduling: deciding which packet to send next on link

- first come, first served
- priority
- round robin
- weighted fair queueing

Packet Scheduling: FCFS

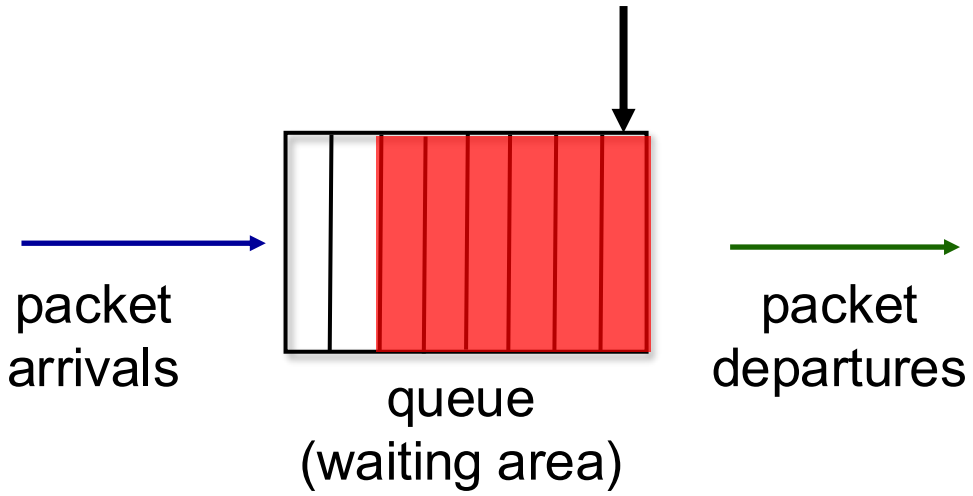

FCFS: packets transmitted in order of arrival to output port

- also known as: First-in-first-out (FIFO)

Packet Scheduling: FCFS

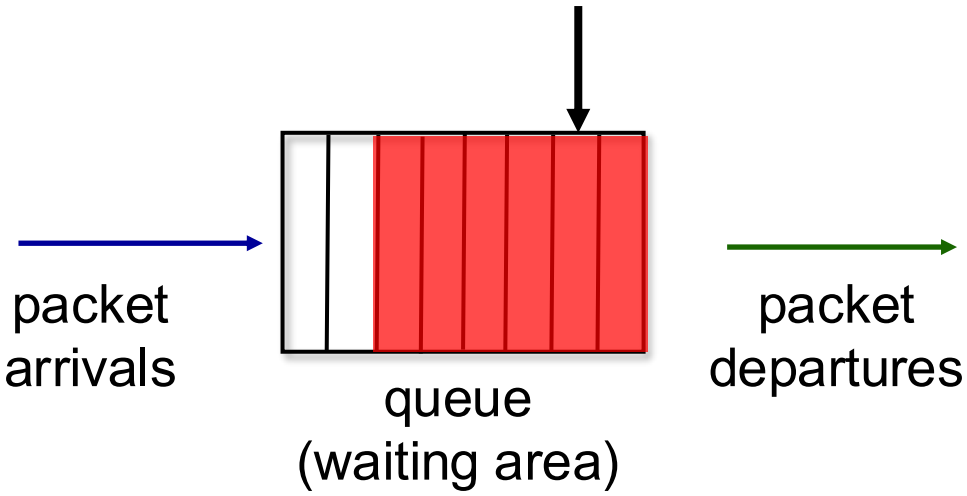

FCFS: packets transmitted in order of arrival to output port

- also known as: First-in-first-out (FIFO)

Packet Scheduling: FCFS

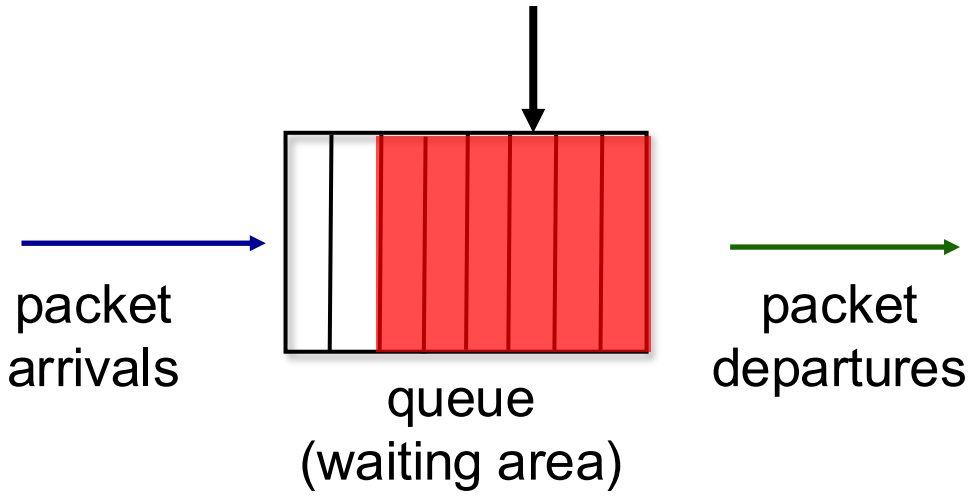

FCFS: packets transmitted in order of arrival to output port

- also known as: First-in-first-out (FIFO)

Packet Scheduling: FCFS

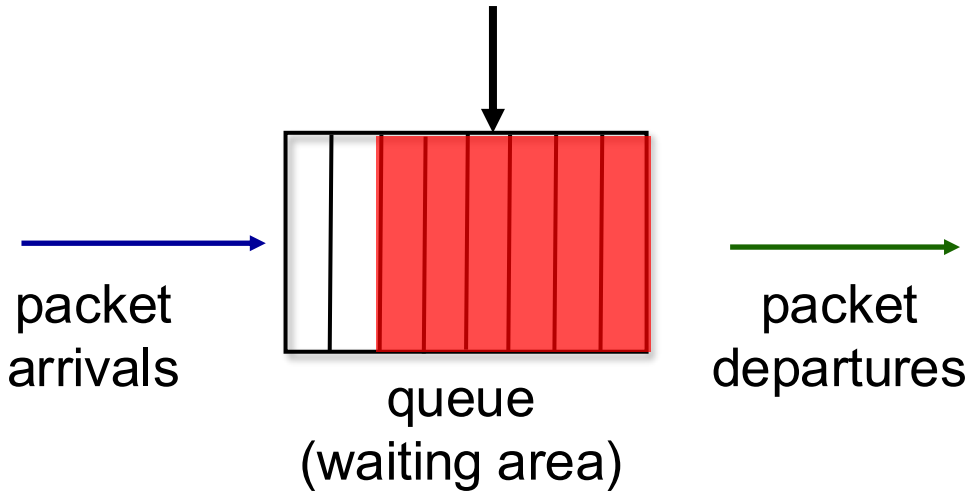

FCFS: packets transmitted in order of arrival to output port

- also known as: First-in-first-out (FIFO)

Packet Scheduling: FCFS

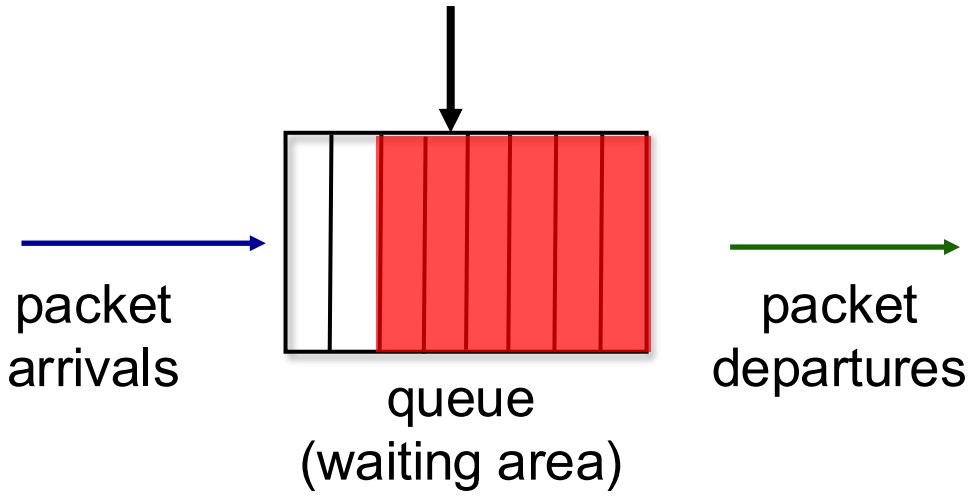

FCFS: packets transmitted in order of arrival to output port

- also known as: First-in-first-out (FIFO)

Scheduling policies: priority

Priority scheduling:

- arriving traffic classified, queued by class
  - any header fields can be used for classification
- send packet from highest priority queue that has buffered packets
  - FCFS within priority class

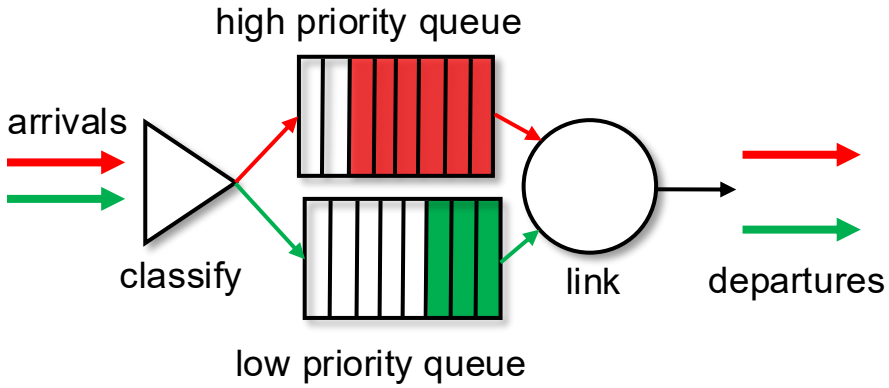

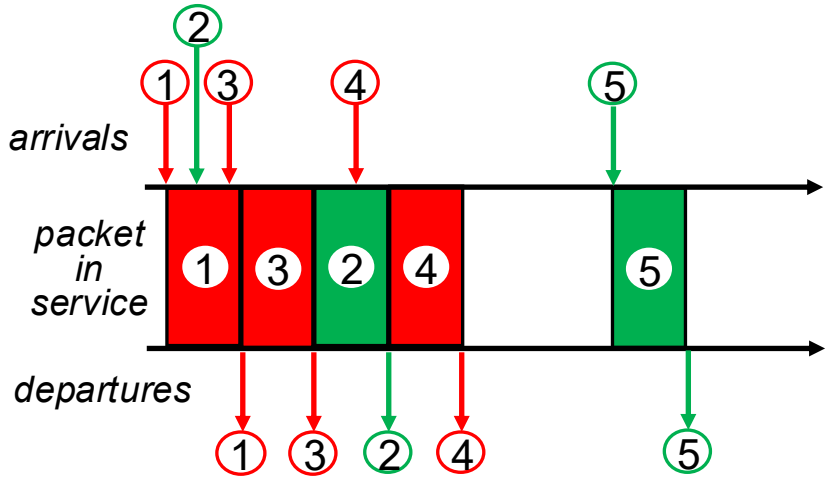

Scheduling policies: round robin

Round Robin (RR) scheduling:

- arriving traffic classified, queued by class
  - any header fields can be used for classification
- server cyclically, repeatedly scans class queues, sending one complete packet from each class (if available) in turn

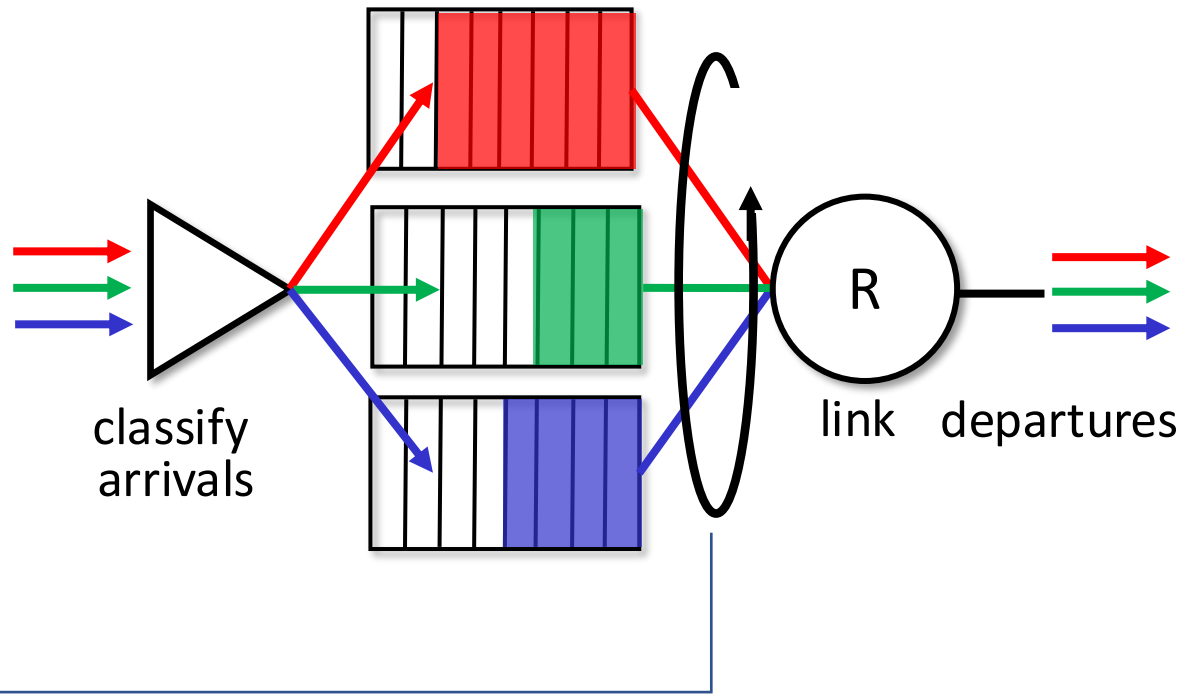

Scheduling policies: weighted fair queueing

Weighted Fair Queuing (WFQ):

- generalized Round Robin
- each class, *i*, has weight, w_i, and gets weighted amount of service in each cycle:

$$\frac{w_i}{\sum_j w_j}$$

- minimum bandwidth guarantee (per-traffic-class)

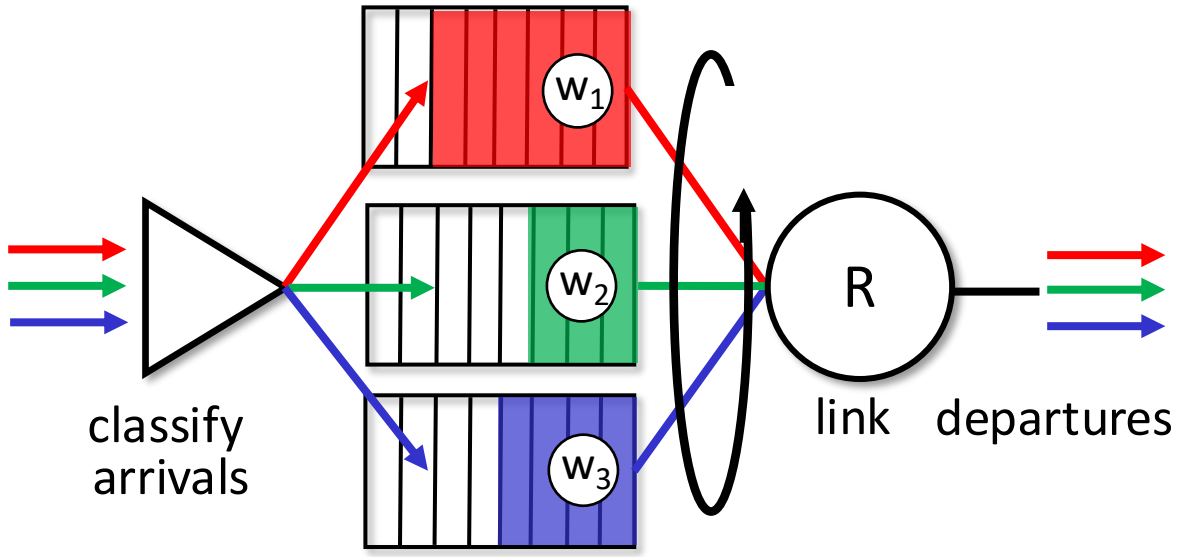

Network layer: "data plane" roadmap

- Network layer: overview
  - data plane
  - control plane
- What's inside a router
  - input ports, switching, output ports
  - buffer management, scheduling
- IP: the Internet Protocol
  - datagram format
  - addressing
  - network address translation
  - IPv6

- Generalized Forwarding, SDN
  - match+action
  - OpenFlow: match+action in action
- Middleboxes

IP addressing: Format

Dotted-decimal IP address notation: **IPv4: 4 bytes**

223.1.1.1 = 11011111 00000001 00000001 00000001

223 1 1 1

2^{32} = 4,294,967,296

IPv6: 16 bytes

2^{128} =

3.4×10^{38}

340,282,366,920,938,463,463,374,607,431,768,211,456

A Very Important Concept:
Subnets

How should we represent the subnets?

Subnets Masks

Network Portion
(fixed for a subnet)

Host Portion
(varies with host)

11000000. 10101000. 00000001. 00000000

11000000. 10101000. 00000001. 01111111

192.168.1.0 to 192.168.1.127

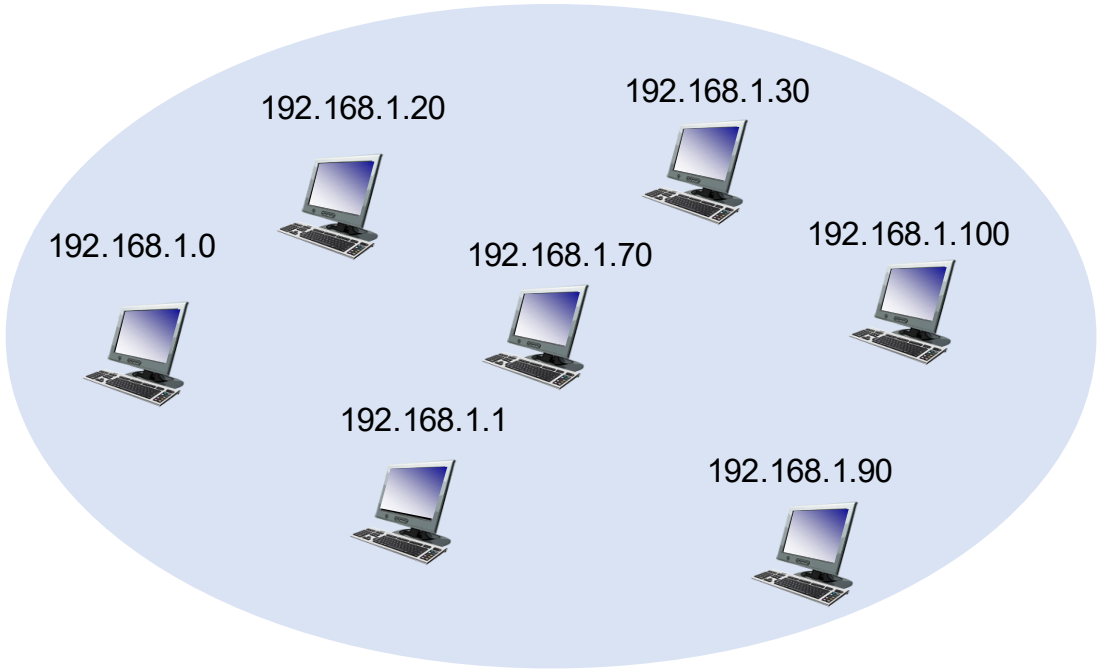

11111111. 11111111. 11111111. 1

Network Portion **Host Portion**

The **sequence of 1s** in the subnet mask indicates which bits of the IP address belong to the **network portion**

The **sequence of 0s** indicates which bits belong to the **host portion**.

Subnets Masks

192.168.1.0 to 192.168.1.127

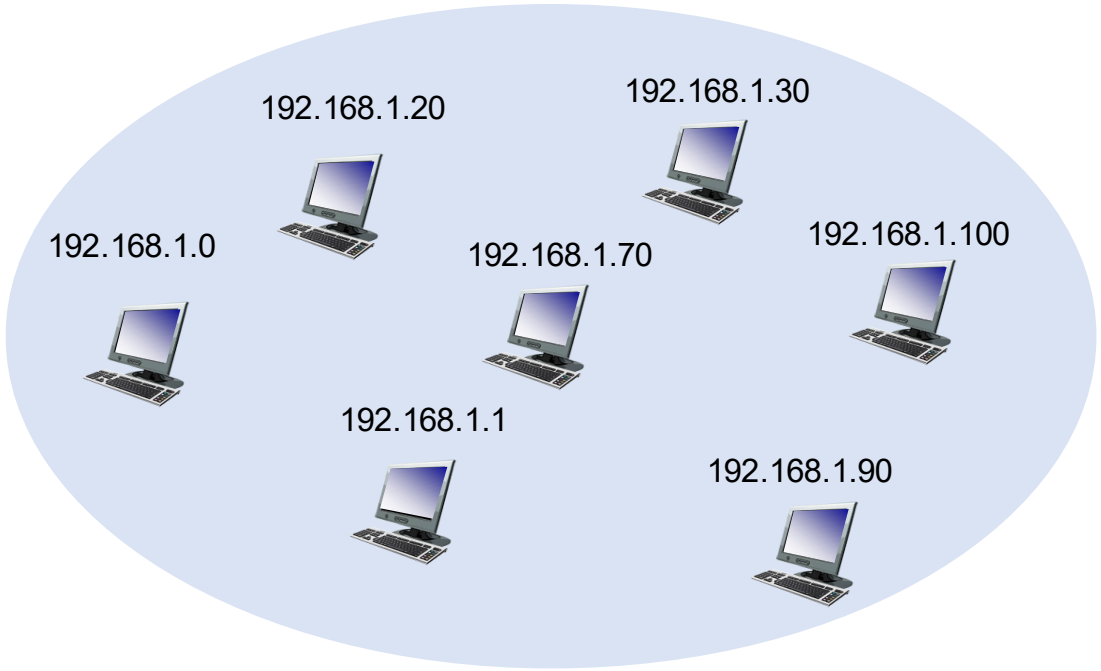

Network Portion (fixed for a subnet)

Host Portion (varies with host)

11000000. 10101000. 00000001. 00000000

11000000. 10101000. 00000001. 01111111

11111111. 11111111. 11111111. 10000000

Network Portion

Host Portion

Subnet Mask: 255. 255. 255. 128

Subnets Masks

192.168.1.0 to 192.168.1.127

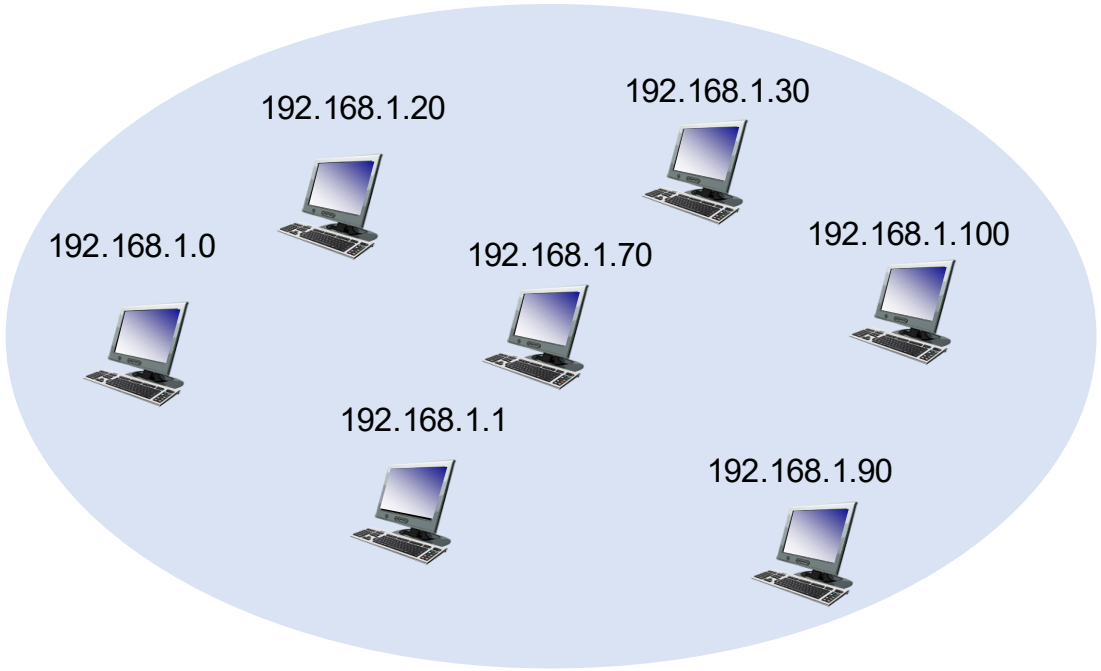

Network Portion
(fixed for a subnet)

Host Portion
(varies with host)

11000000. 10101000. 00000001. 00000000

11111111. 11111111. 11111111. 10000000

Network Portion **Host Portion**

IP:	**192.**	**168.**	**1.**	**0**
Subnet Mask:	**255.**	**255.**	**255.**	**128**

Subnet and CIDR

CIDR: Classless InterDomain Routing (pronounced “cider”)

- address format: a.b.c.d/x, where x is # bits in subnet portion of address

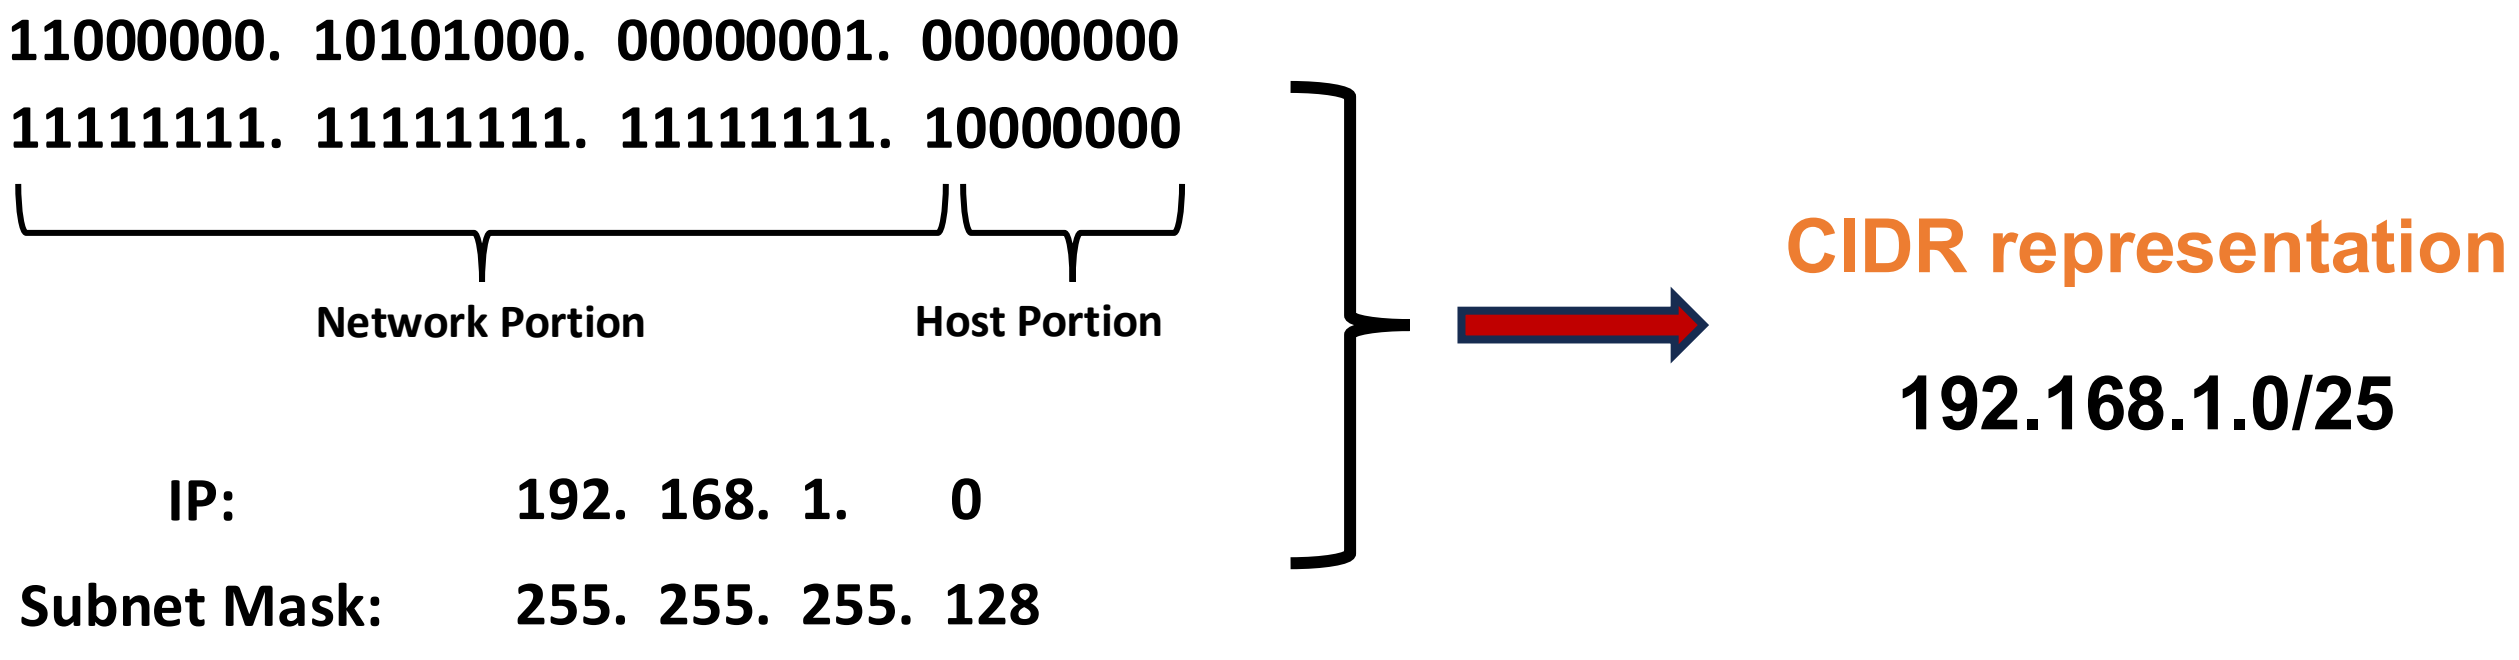

Forwarding Table

IP Address Range		Forwarding Interface
192.168.0.1	192.168.0.20	**1**
192.168.0.40	192.168.0.60	**2**
192.168.0.80	192.168.0.100	**3**

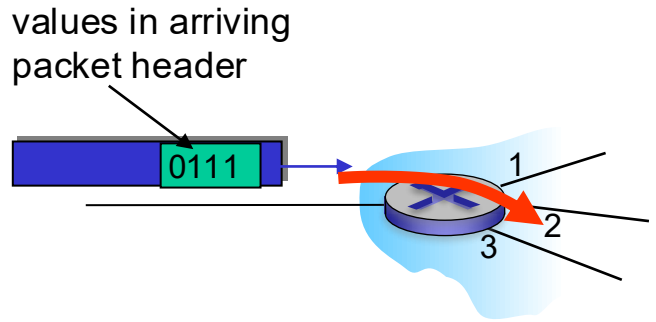

Forwarding Table

IP Address Range	Forwarding Interface
192.168.1.0/25	**1**
192.168.0.0/24	**2**
127.1.0.80/26	**3**

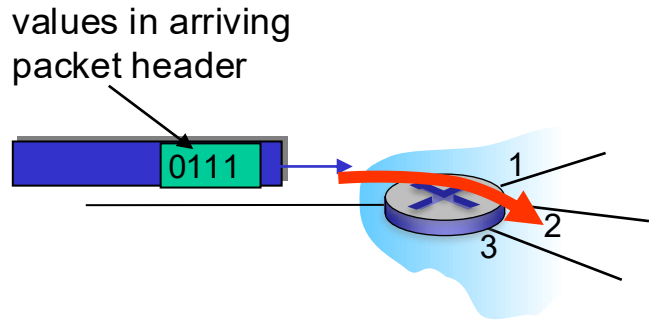

Any digits

11000000. 10101000. 00000001. 00000000

11111111. 11111111. 11111111. 10000000

11000000. 10101000. 00000001. 0*******

Longest prefix matching

longest prefix match

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

examples:

11001000 00010111 00010110 10100001 which interface?

11001000 00010111 00011000 10101010 which interface?

Longest prefix matching

longest prefix match

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

match!

examples:

11001000 00010111 00010110 10100001 which interface?

11001000 00010111 00011000 10101010 which interface?

Longest prefix matching

longest prefix match

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

match!

examples:

11001000 00010110 10100001 which interface?

11001000 00010111 00011000 10101010 which interface?

Longest prefix matching

longest prefix match

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

match!

examples:

11001000 00010111 00010110 10100001 which interface?

11001000 00010111 00011000 10101010 which interface?

Longest prefix matching

longest prefix match

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

examples:

11001000 00010111 00010110 10100001 which interface?

11001000 00010111 00011000 10101010 which interface?

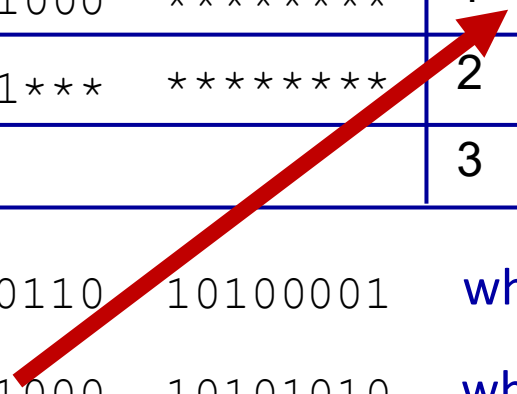

How to obtain the IP address?

Hierarchical addressing: route aggregation

hierarchical addressing allows efficient advertisement of routing information:

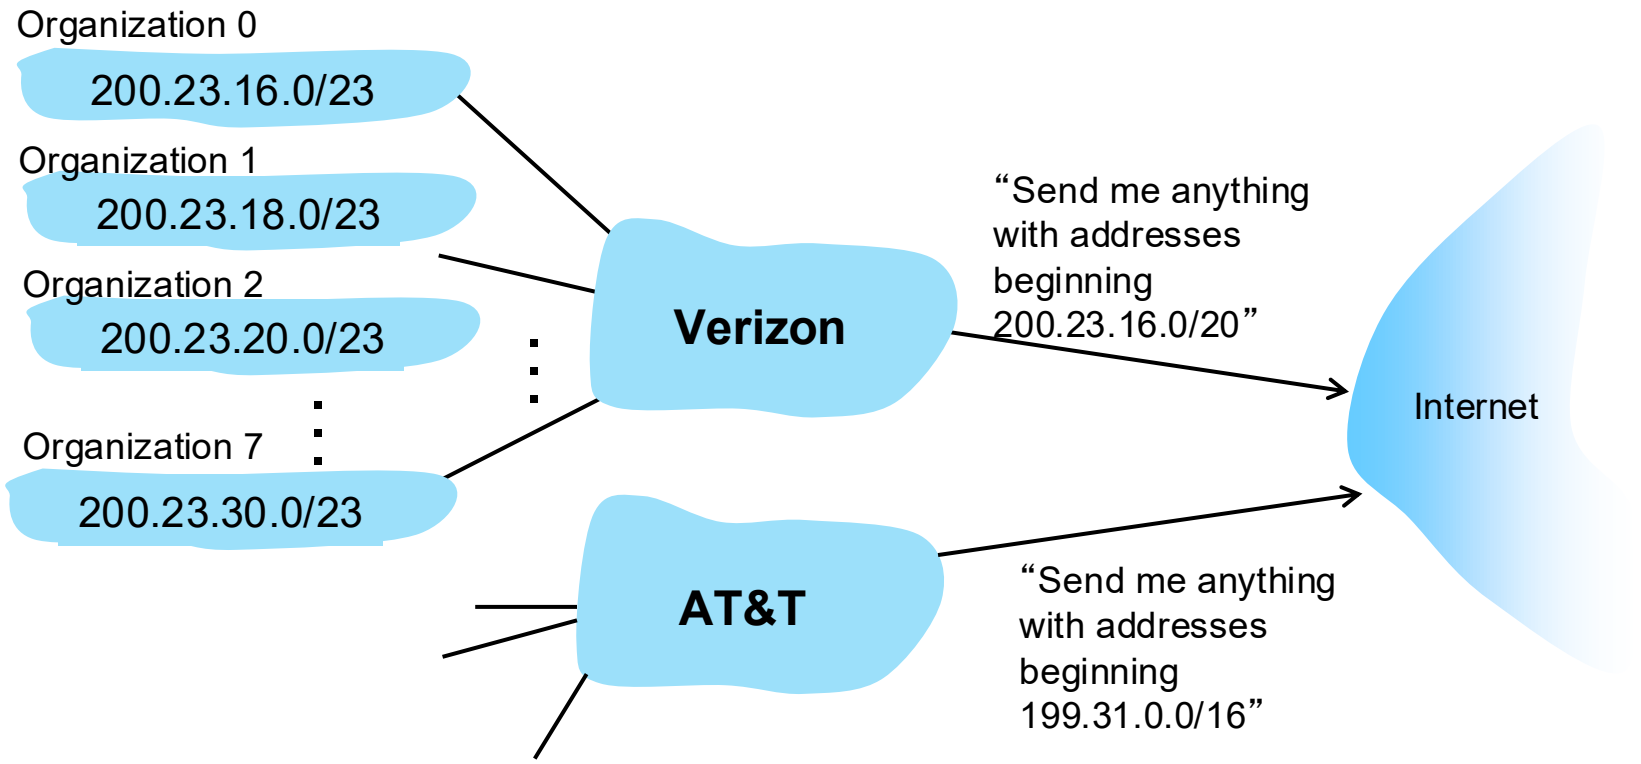

Hierarchical addressing: more specific routes

- Organization 1 moves from ISP-1 (Verizon) to ISP-2 (AT&T)
- ISP-2 (AT&T) now advertises a more specific route to Organization 1

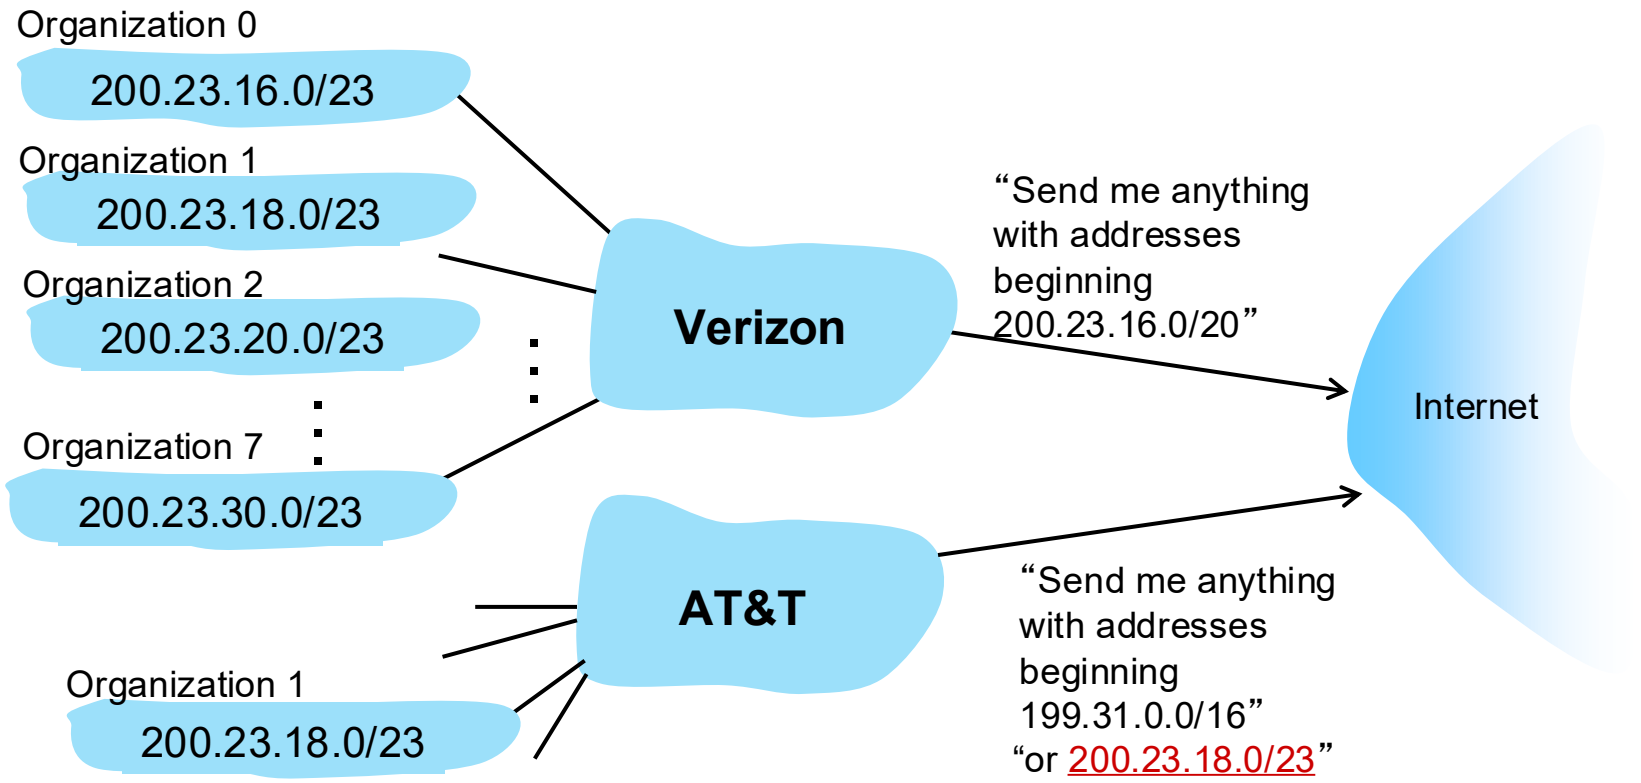

DHCP: Dynamic Host Configuration Protocol

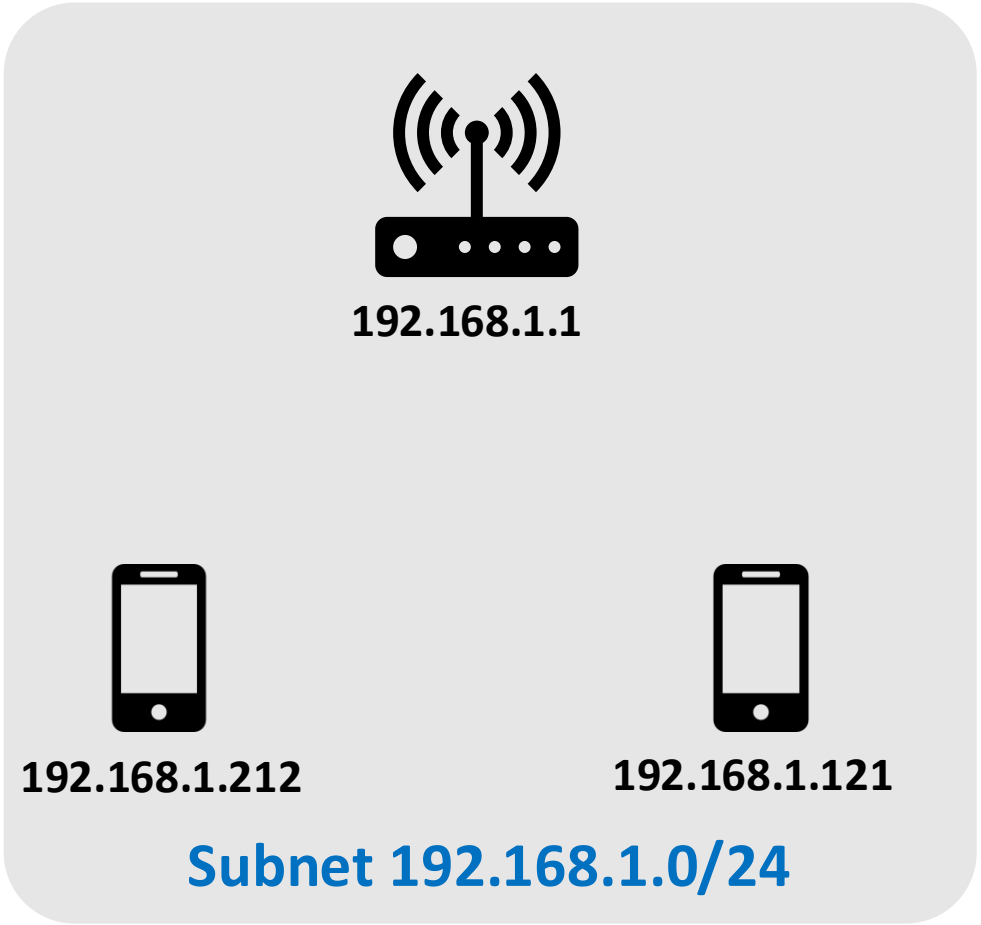

goal: host *dynamically* obtains IP address from network server when it "joins" network

- can renew its lease on address in use
- allows reuse of addresses (only hold address while connected/on)
- support for mobile users who join/leave network

NAT: network address translation

NAT: all devices in local network share just one IPv4 address as far as outside world is concerned

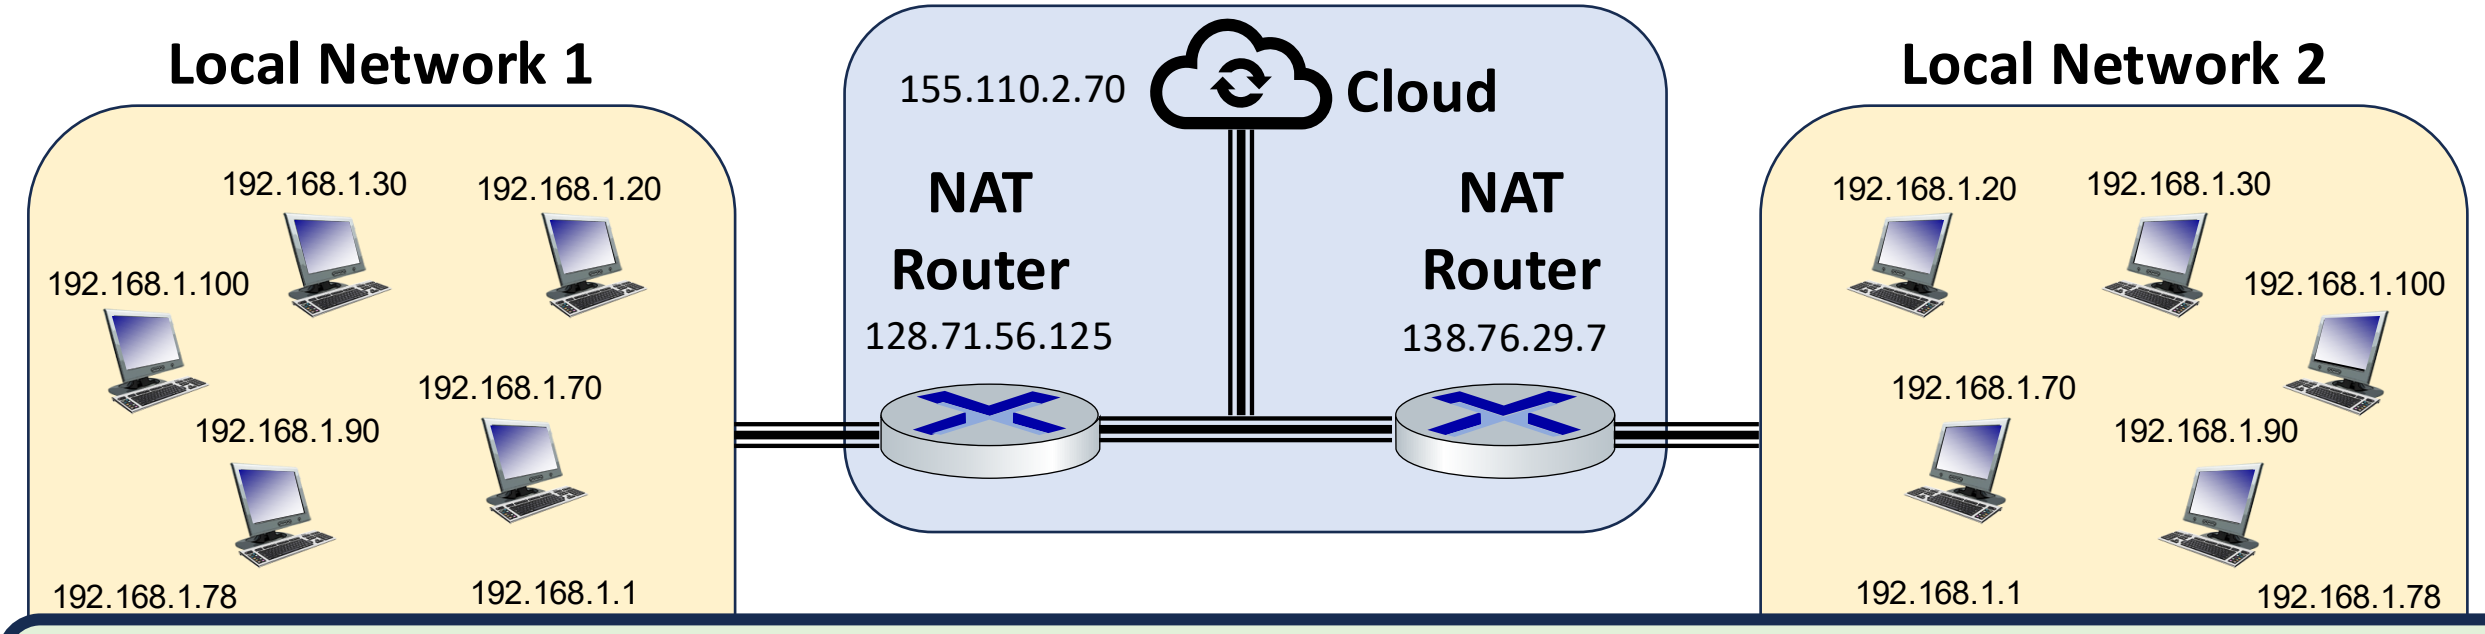

How to we translate the IP address from Private to Public?

Network layer: "data plane" roadmap

- Network layer: overview
  - data plane
  - control plane
- What's inside a router
  - input ports, switching, output ports
  - buffer management, scheduling
- IP: the Internet Protocol
  - datagram format
  - addressing
  - network address translation
  - IPv6

- Generalized Forwarding, SDN
  - Match+action
  - OpenFlow: match+action in action
- Middleboxes

Chapter 5
Network Layer: Control Plane

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

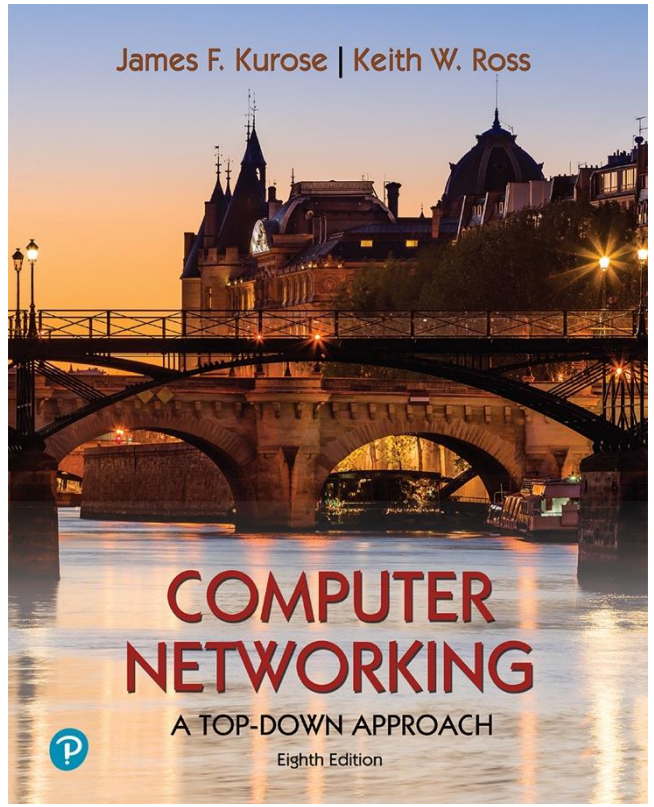

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Adapted from the slides of the book's authors

Network layer: "control plane" roadmap

- introduction
- Per-Router routing protocols
  - link state
  - distance vector

Highly important

- intra-ISP routing: OSPF
- routing among ISPs: BGP
- SDN control plane
- Internet Control Message Protocol
- network management, configuration
  - SNMP
  - NETCONF/YANG

PA3 related topics: Split Horizon, Poising Reverse, Count-to-Infinity

Internet approach to scalable routing

aggregate routers into regions known as "autonomous systems" (AS) (a.k.a. "domains")

intra-AS (aka "intra-domain"):
routing among routers *within same AS ("network")*

- all routers in AS must run same intra-domain protocol
- routers in different AS can run different intra-domain routing protocols
- gateway router: at "edge" of its own AS, has link(s) to router(s) in other AS'es

OSPF

inter-AS (aka "inter-domain"):
routing *among* AS'es

- gateways perform inter-domain routing (as well as intra-domain routing)

BGP

Chapter 6
The Link Layer and LANs

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

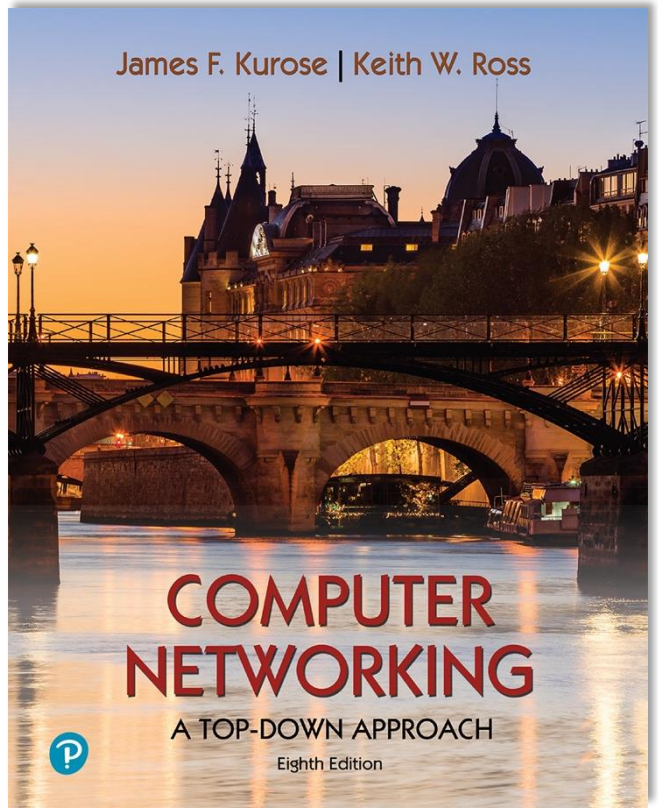

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Adapted from the slides of the book's authors

Error Detection

The Evolution of Redundancy

Parity	Checksum	CRC
	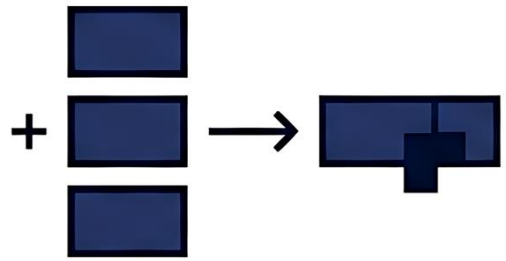	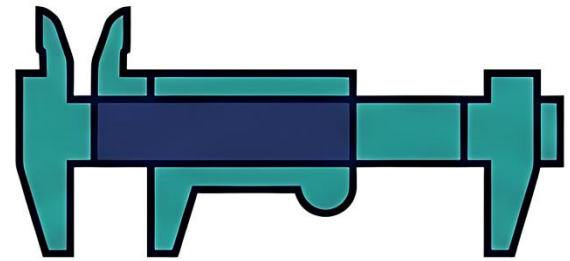
Simple Rule	**Arithmetic Summary**	**Structural Enforcement**
Counts 1s	Adds chunks	Polynomial division
Weak against multi-bit errors	Better, but collisions happen	Extremely resilient to burst errors

All three mechanisms add redundancy. The difference lies entirely in how deeply that redundancy structures the data.

Medium Acces Control

The Design Space: Three Approaches

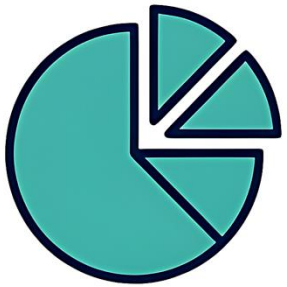

1. Partitioning

Split the resource ahead of time.

2. Random Access

Compete only when needed.

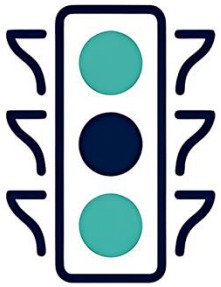

3. Taking Turns

Explicitly decide whose turn it is.

A Very Simple Idea: ALOHA

Send first, recover later: the earliest random access protocol.

No Central Scheduler: Fully distributed, highly simple operation.

Immediate Transmission: Nodes send frames the moment data is available.

Collision Recovery: If frames overlap, nodes wait a randomized duration before trying again.

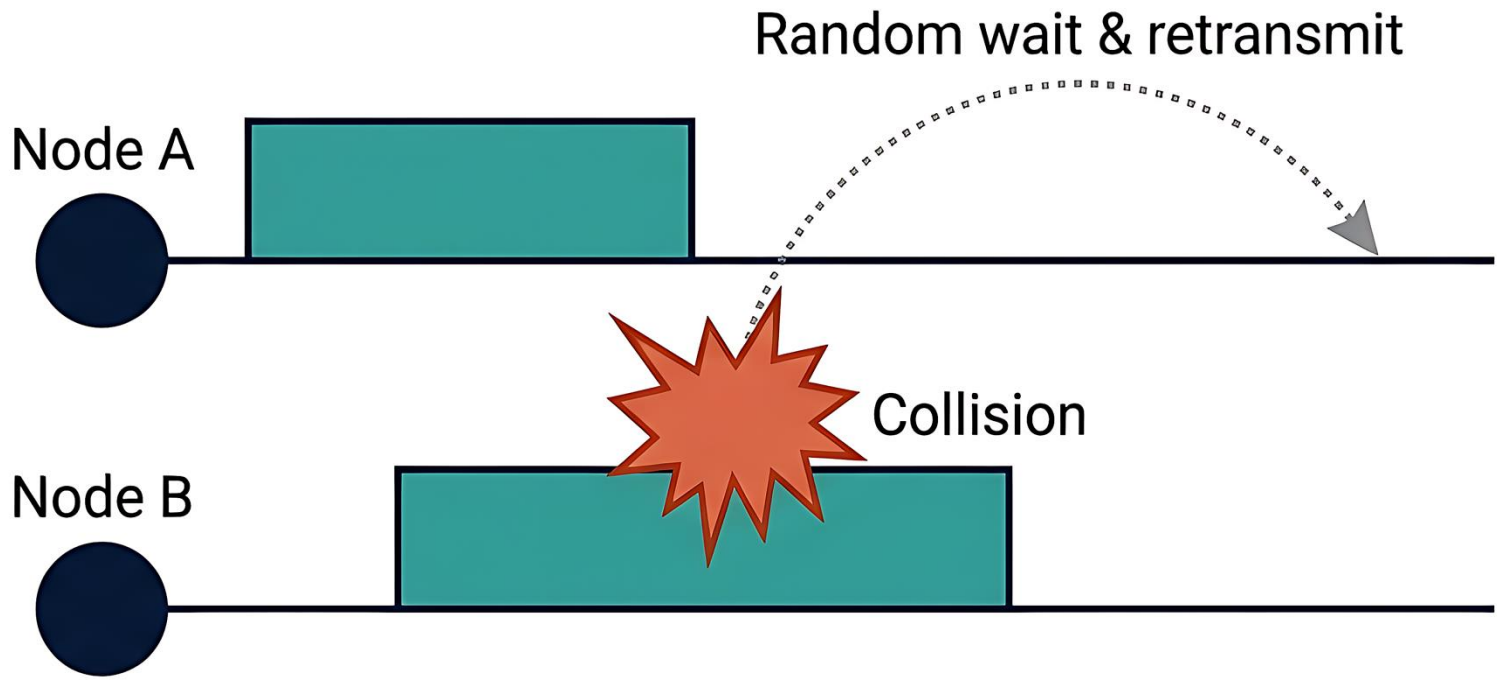

Improvement: Slotted ALOHA

Adding a simple constraint—time slots—drastically reduces collisions.

Aligned Time:
Time is divided into equal, discrete slots.

Strict Boundaries:
Nodes must wait for the beginning of the next slot to transmit.

Eliminates Partial Overlap:
Collisions now only happen if nodes pick the exact same slot.

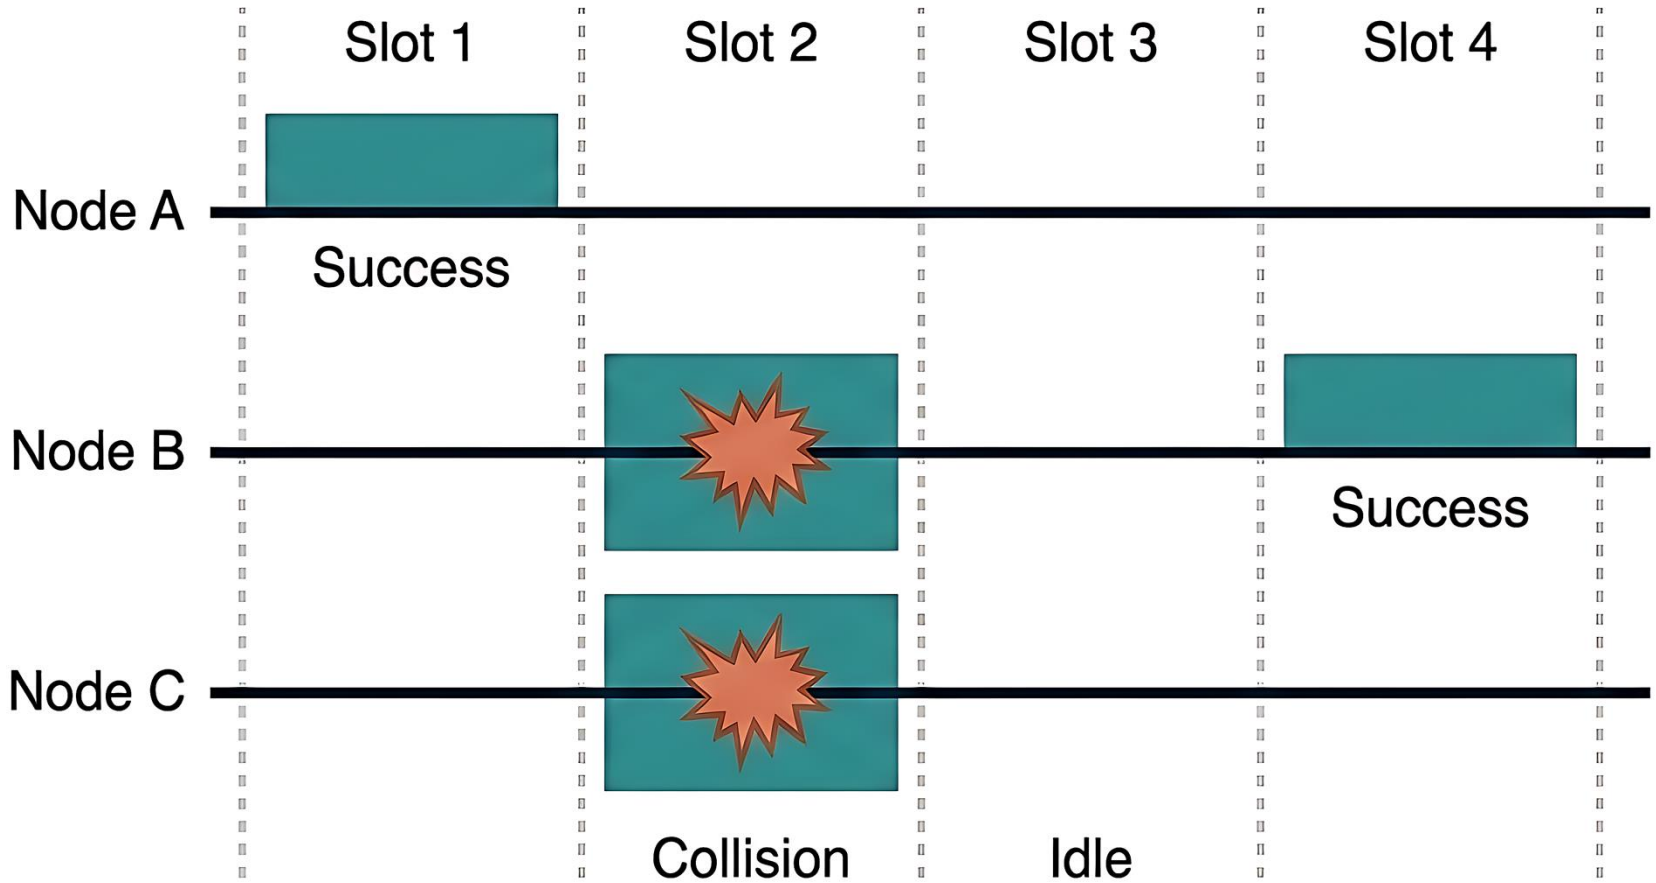

CSMA listens before transmitting

CSMA (Carrier Sense Multiple Access) adds local awareness. If the channel is busy, wait. If it appears idle, transmit.

CSMA is still random access, but smarter.

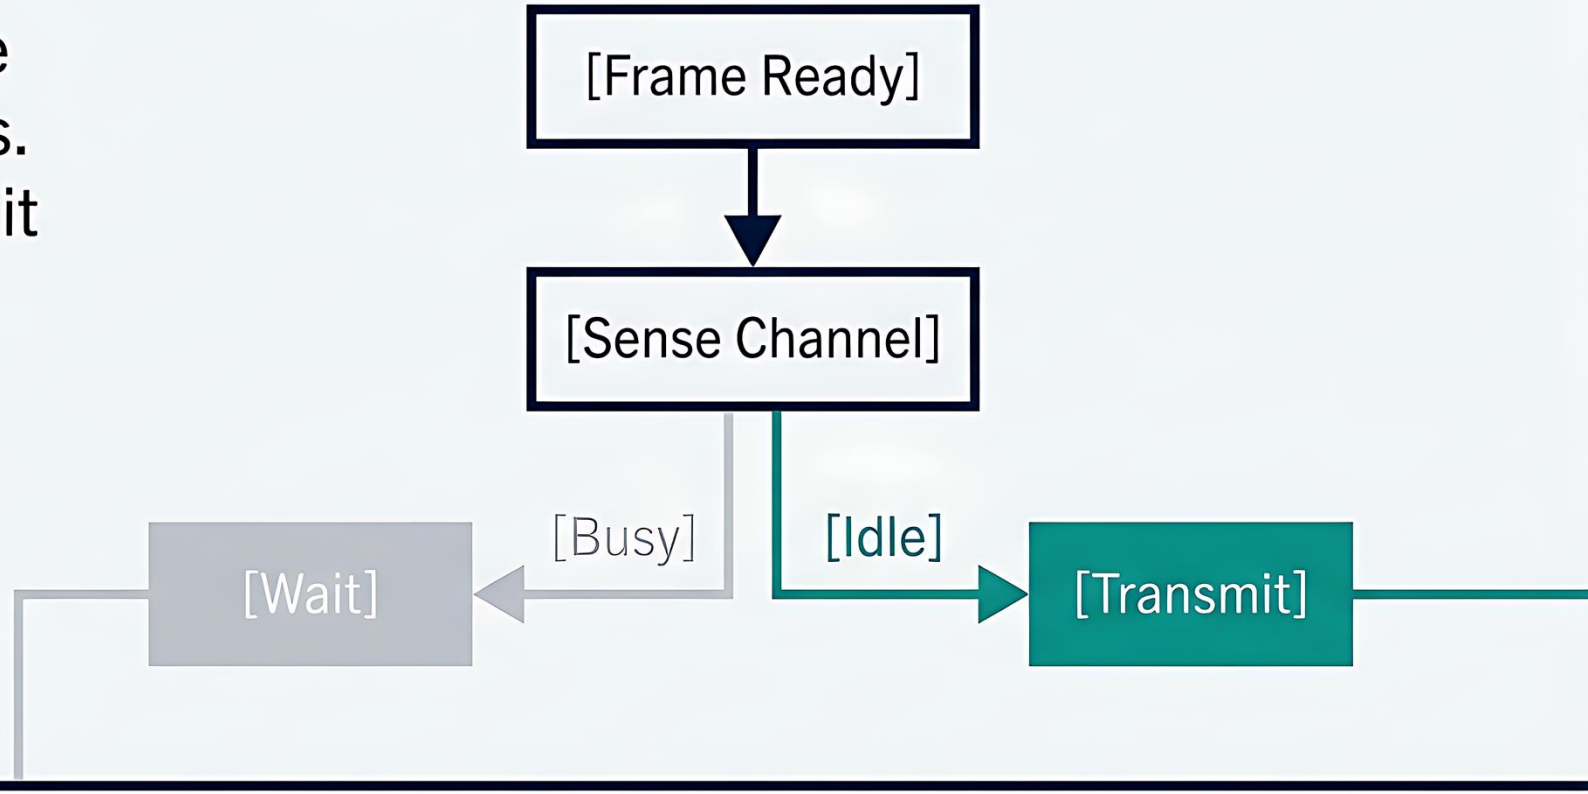

The random-access story is a sequence of refinements

Every new protocol adds structure or information to fix a specific physical or systemic weakness of the previous one.

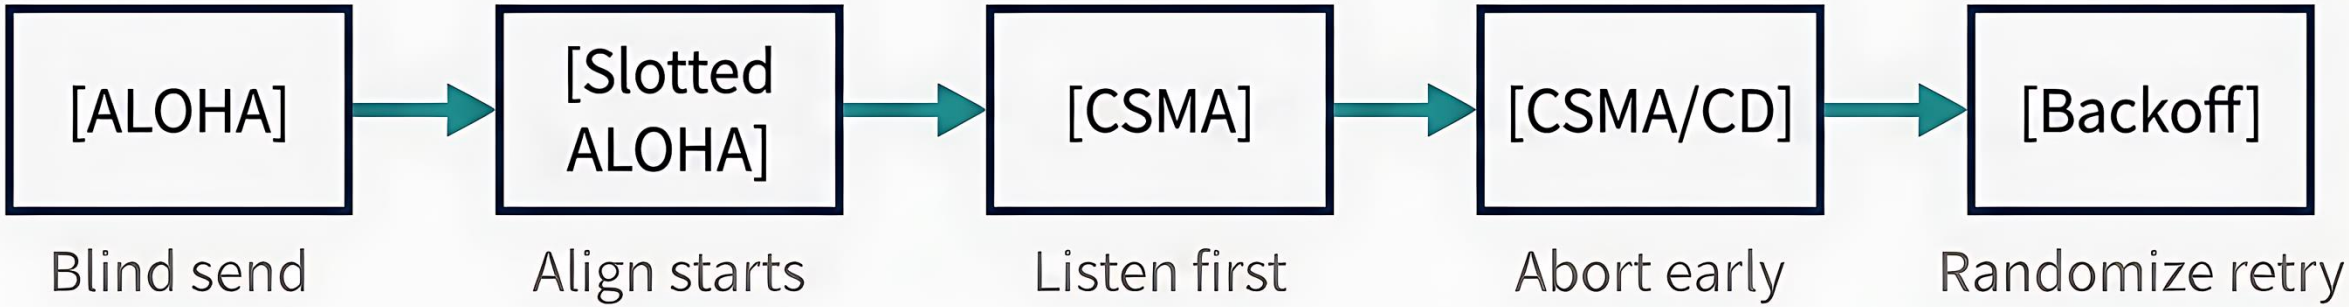

The MAC Comparison Matrix

	Partitioning	Random Access	Taking Turns
Collision Free	Yes	No	Yes
Bursty Traffic Efficiency	Poor	Excellent	Good
Fairness	High	Variable	Excellent
Coordination Overhead	Low	Low	High

MAC Address

Identity on the local link

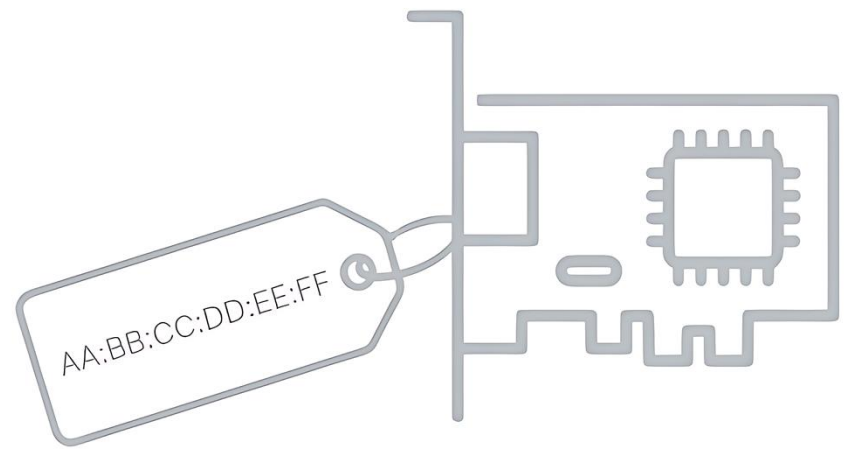

What is a MAC address?

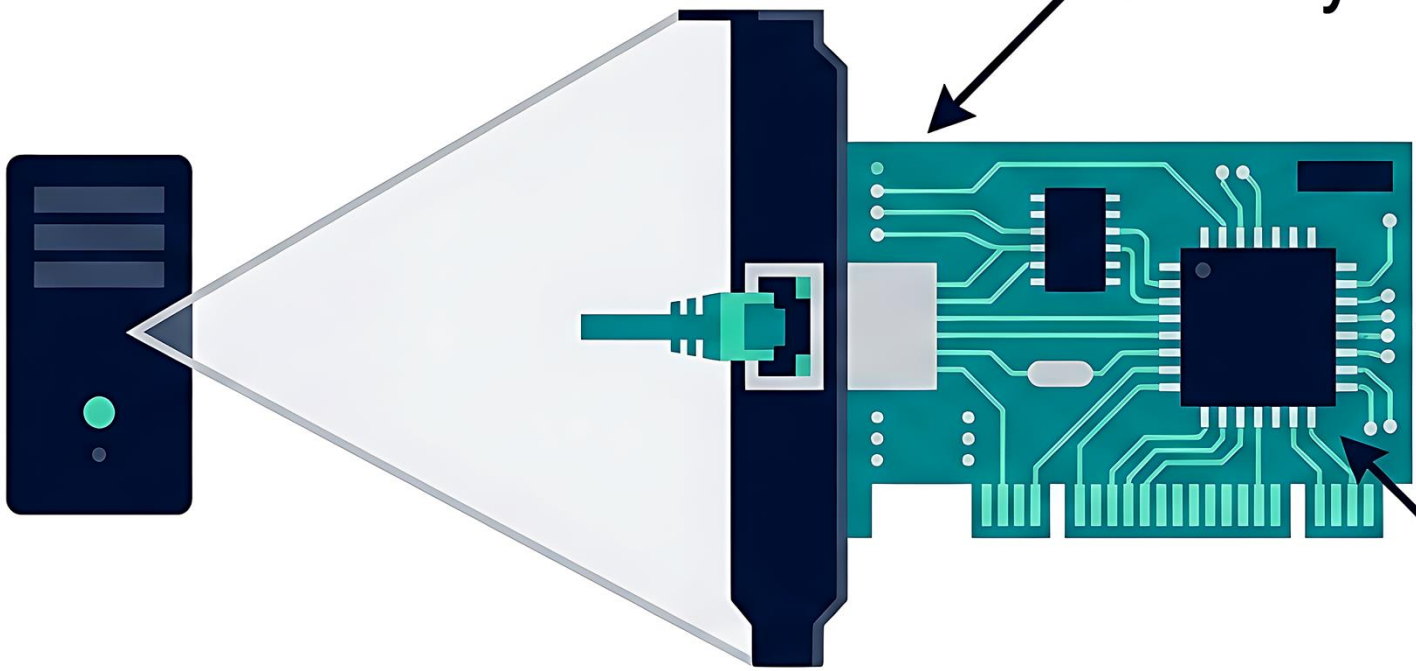

Link-layer address

Usually identifies a network interface (not the whole machine)

Commonly 48 bits (6 bytes)

Example: 08:00:27:1A:2B:3C

Three ways to deliver on a local link

Unicast

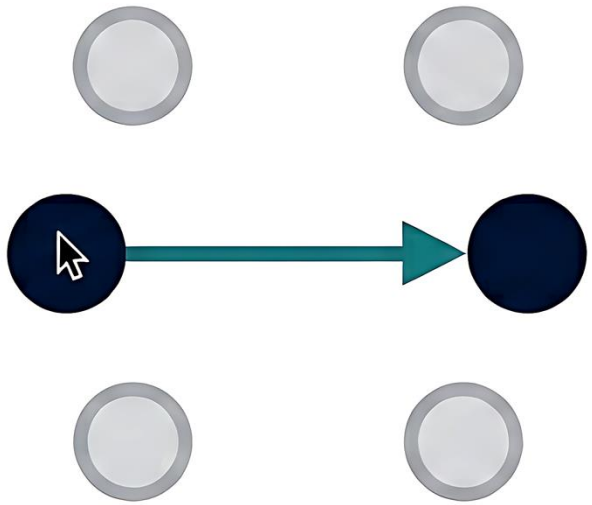

One sender to one receiver.

Broadcast

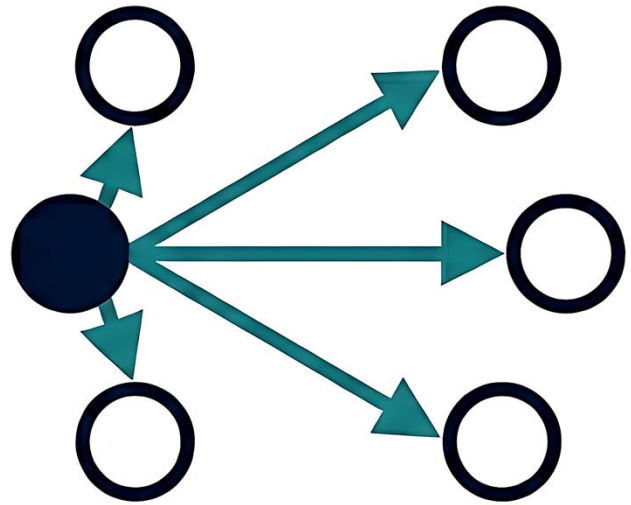

One sender to all nodes on the LAN.

Broadcast MAC: FF:FF:FF:FF:FF:FF

Multicast

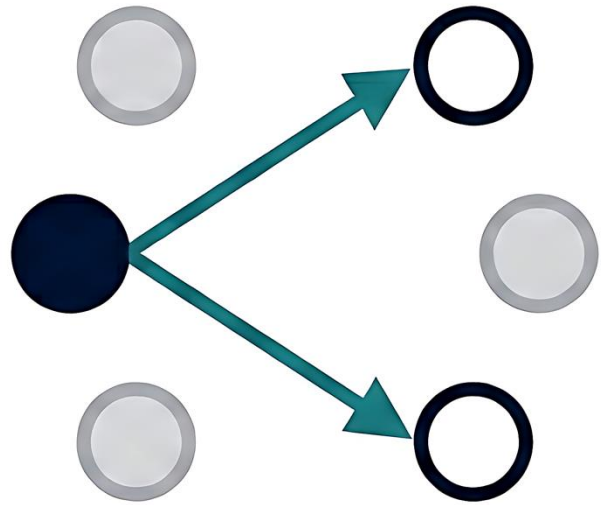

One sender to a group of receivers.

How does a host learn the destination MAC address?

To send an Ethernet frame, **the sender needs a destination MAC address.**

But usually, the sender starts with an IP address.

The Problem:
Which MAC address corresponds to this IP?

The Solution: ARP (Address Resolution Protocol).

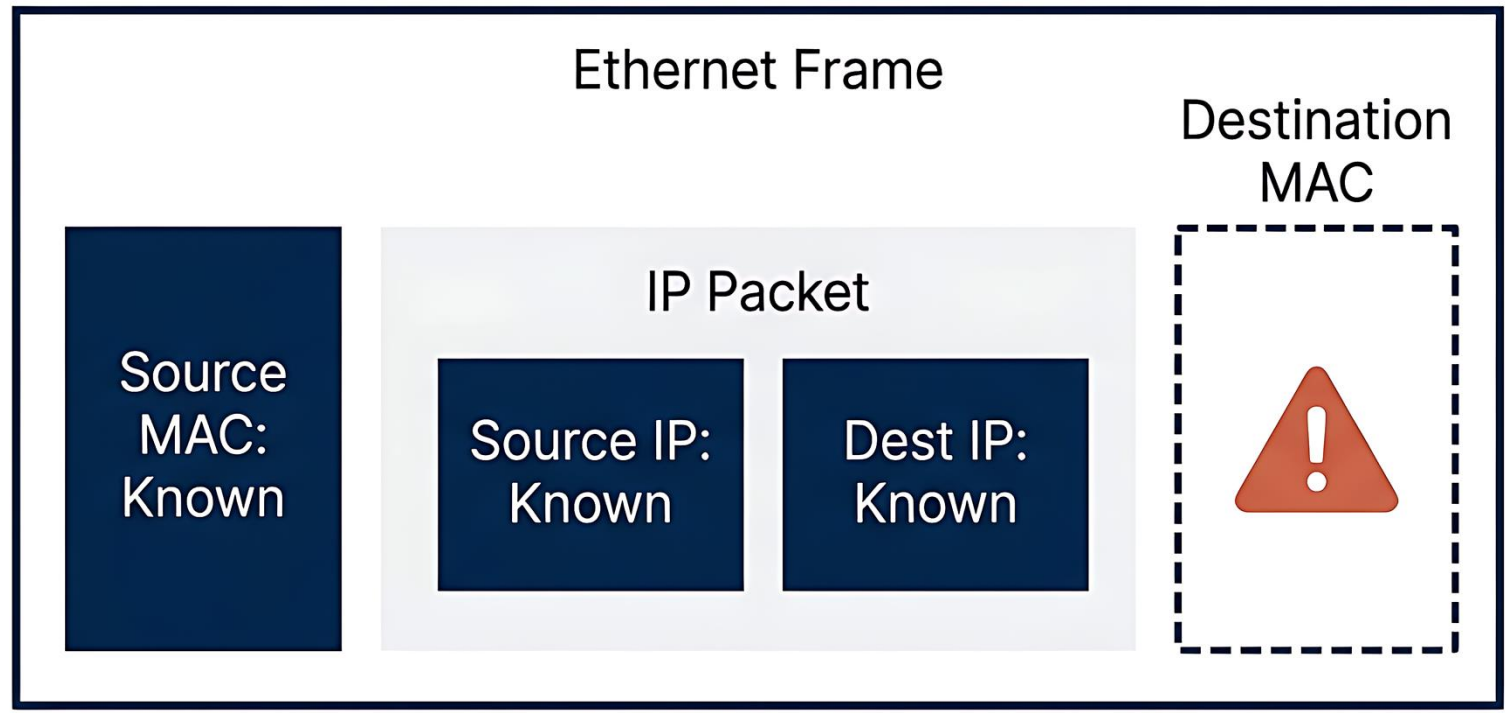

LAN and Switch

From shared media to **switched** networks

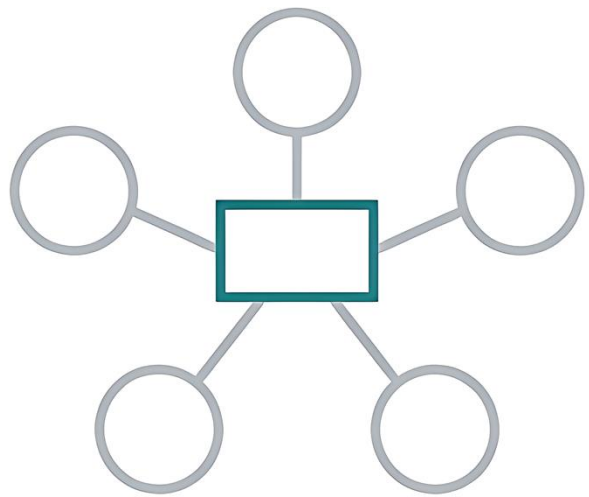

A Key Question: How Does the Switch Know the Right Port?

Checking the destination MAC address is not enough.
The switch must know *where* that address is physically located.

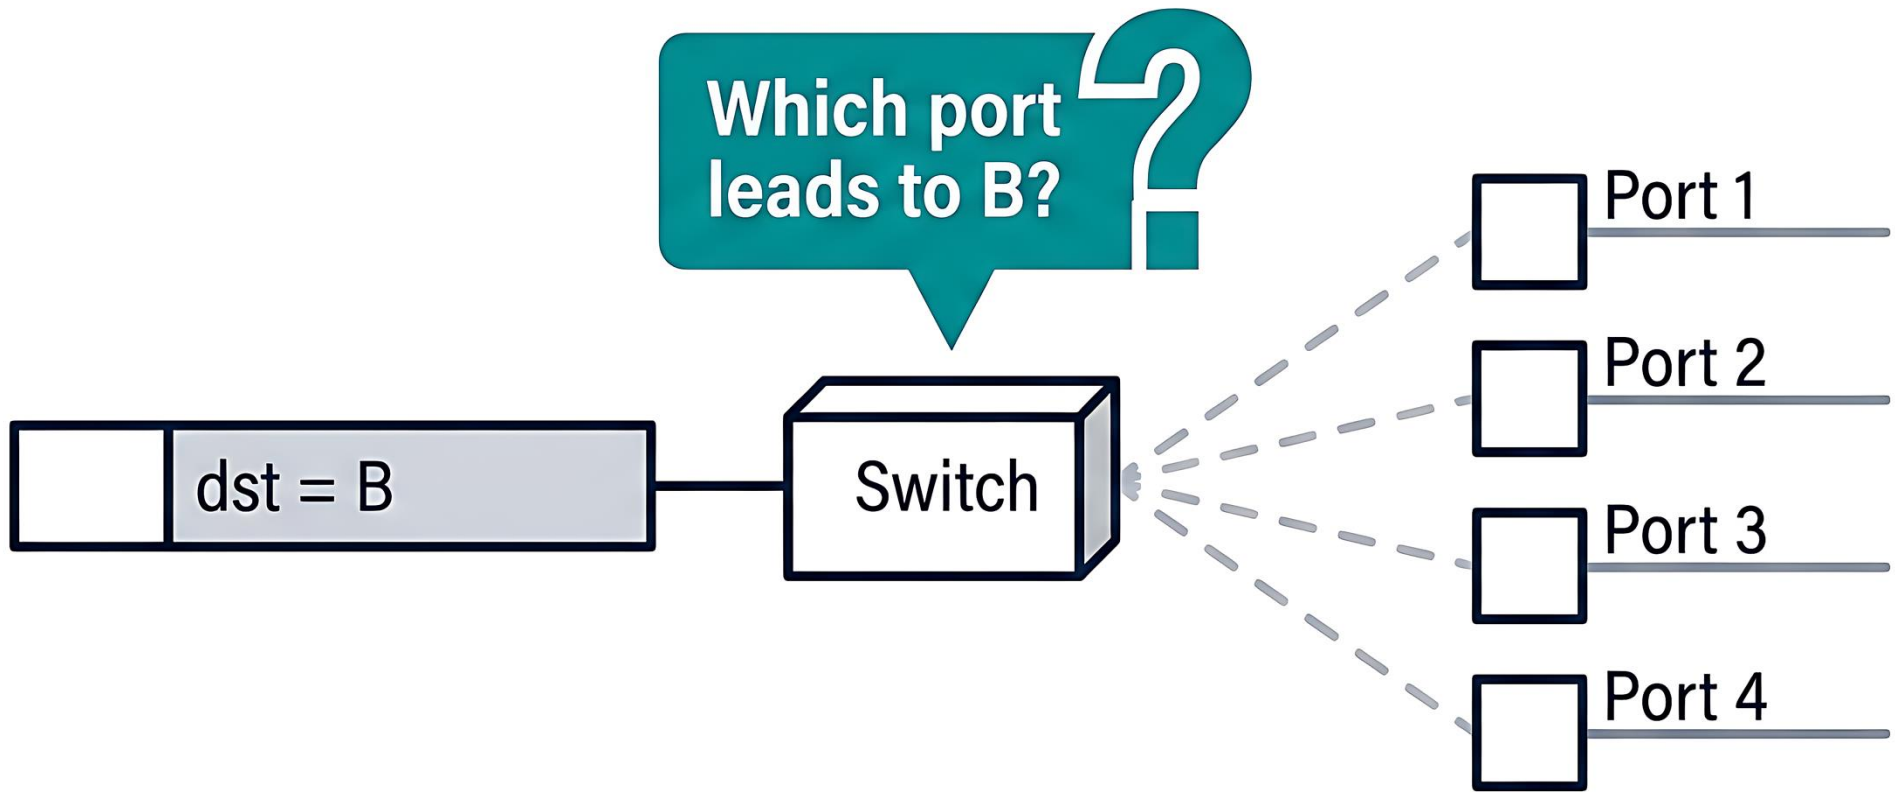

The Switch Needs a MAC-to-Port Table

To forward selectively, the switch maintains an internal mapping between MAC addresses and physical ports.

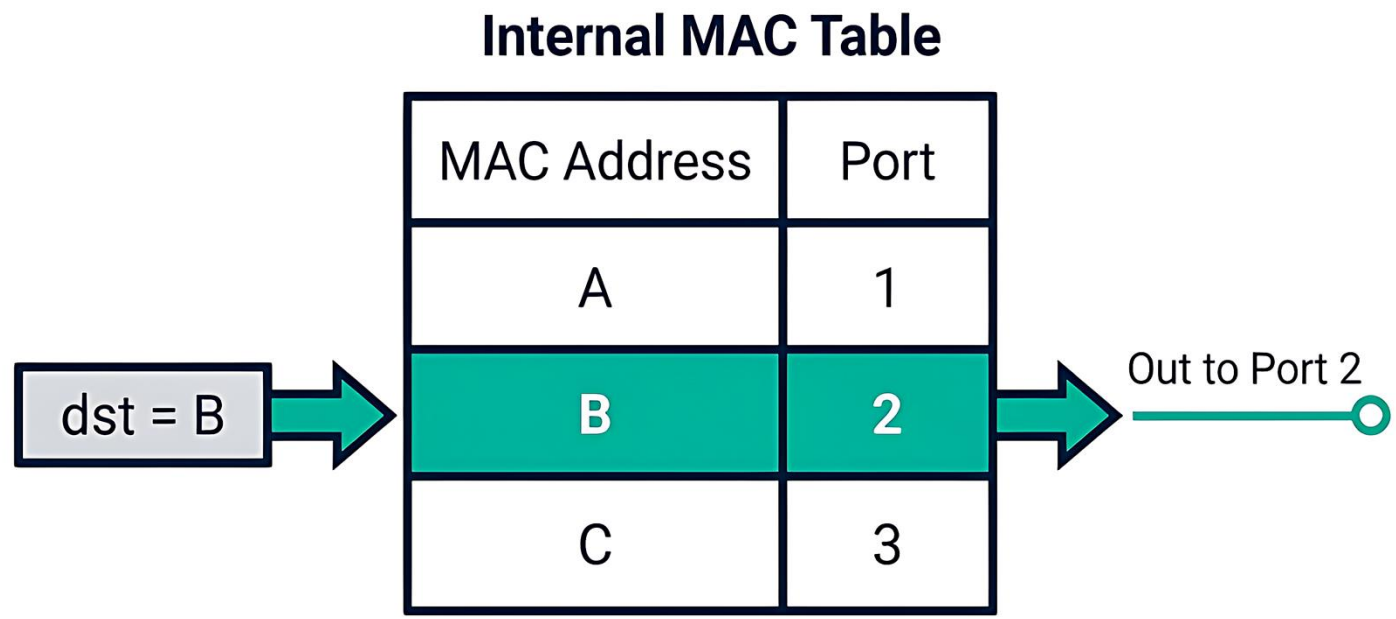

The Key Idea: Learn by Watching Frames

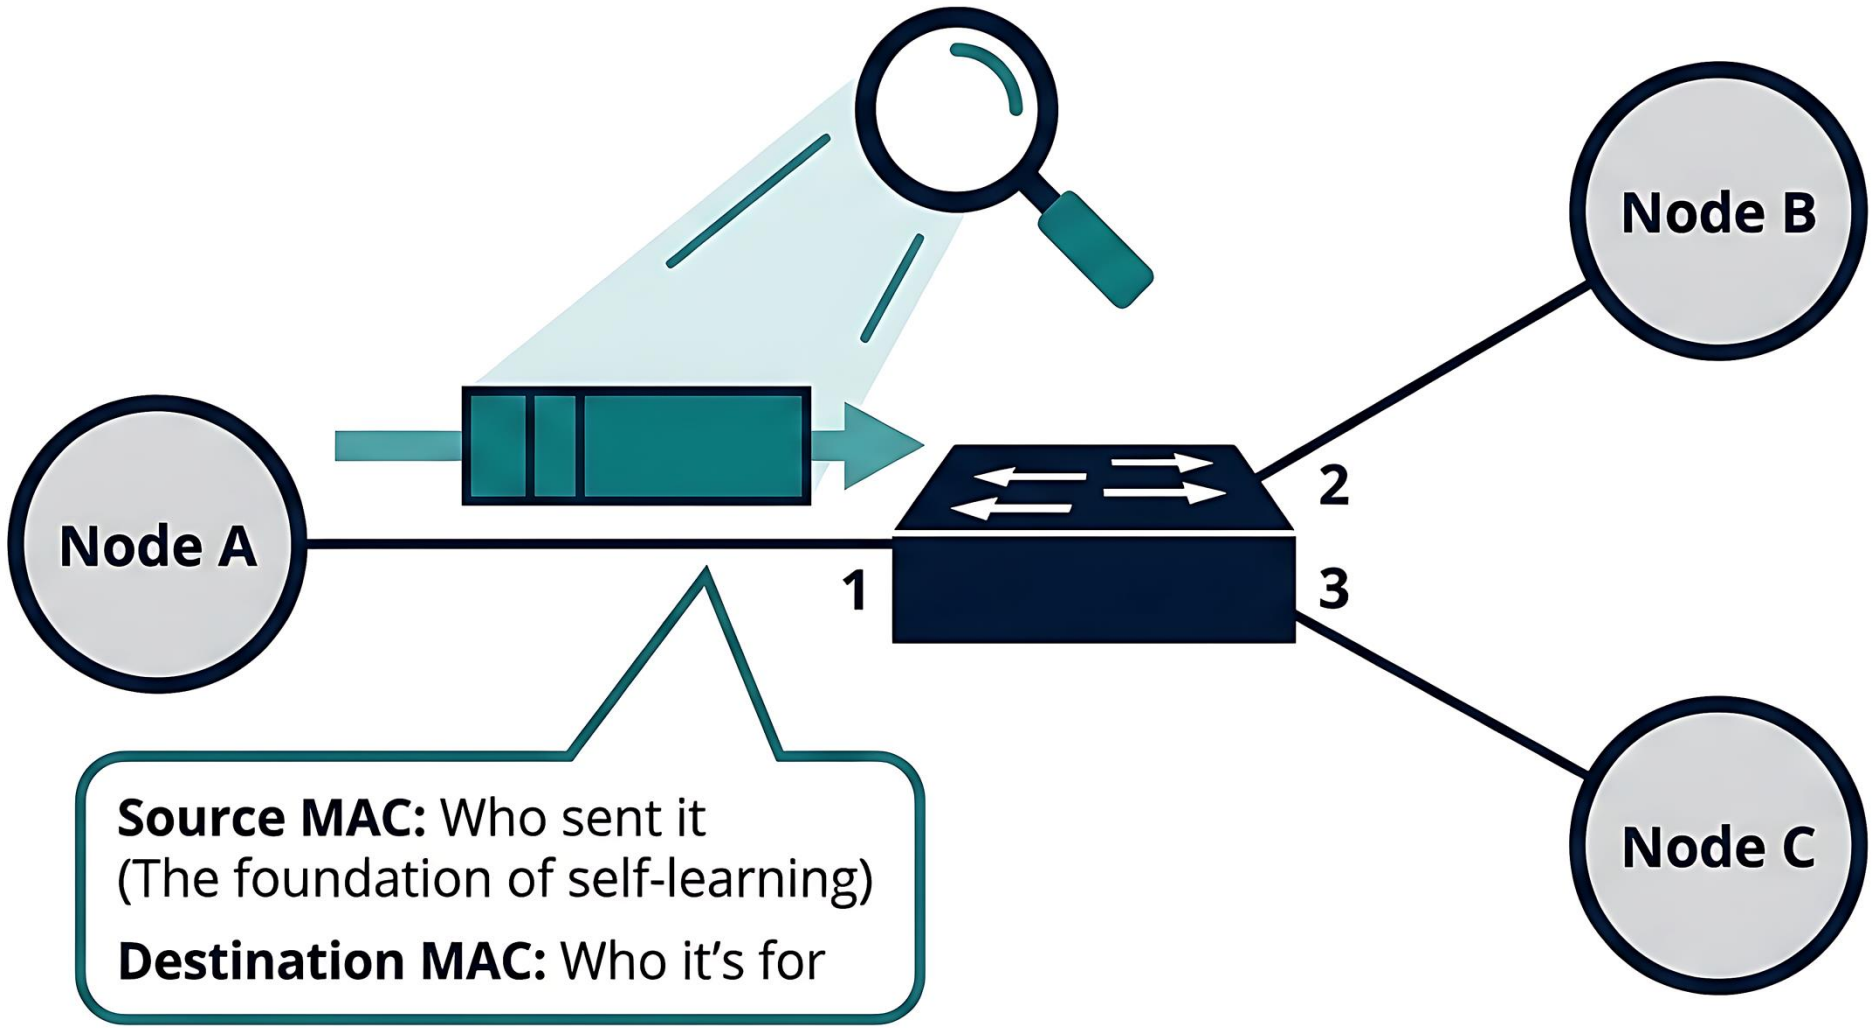

The Learning Rule: Source MAC -> Incoming Port

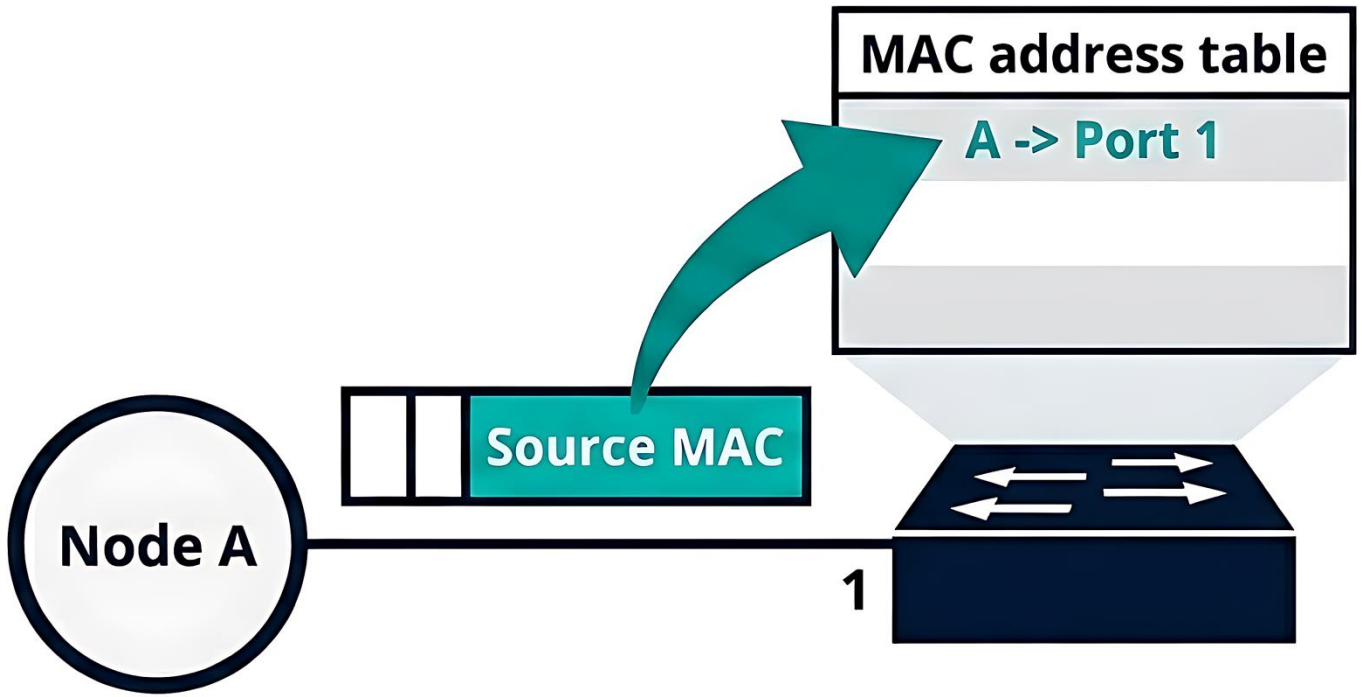

If a frame comes from port X, the sender must be on port X.

What Happens When a Frame Arrives?

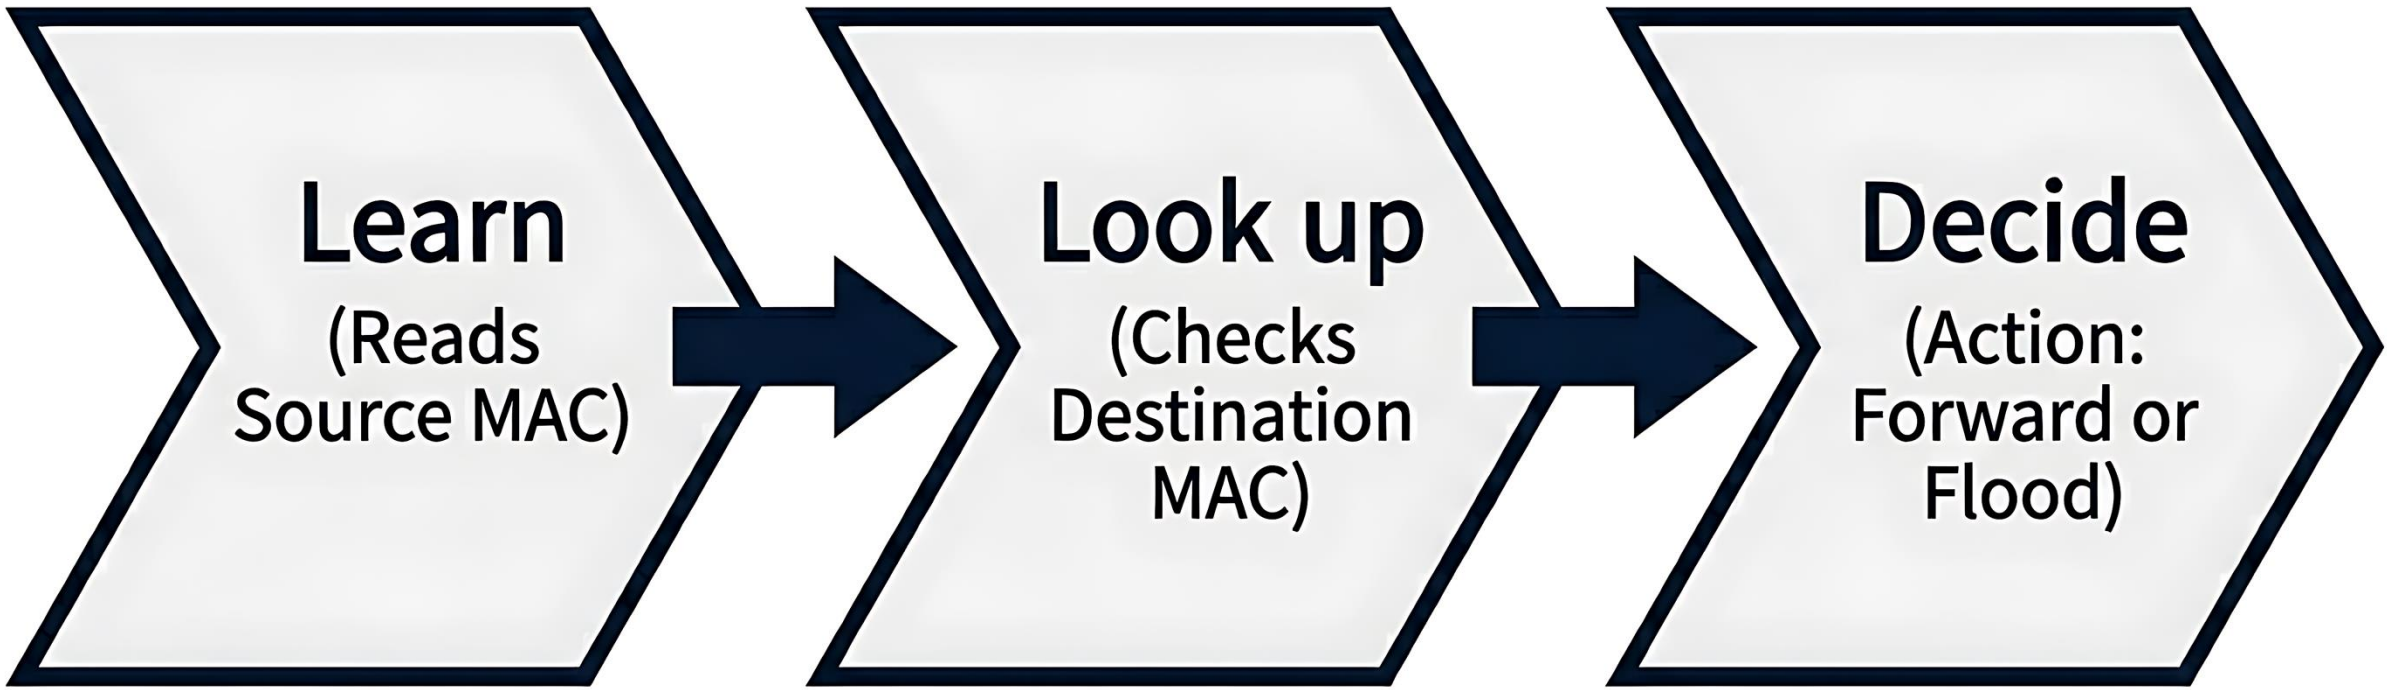

Learning and forwarding happen continuously on every single arriving frame.

Case 1: Destination Found in the Table

- The switch already knows the destination port.
- Forwards the frame selectively.
- Highly efficient.

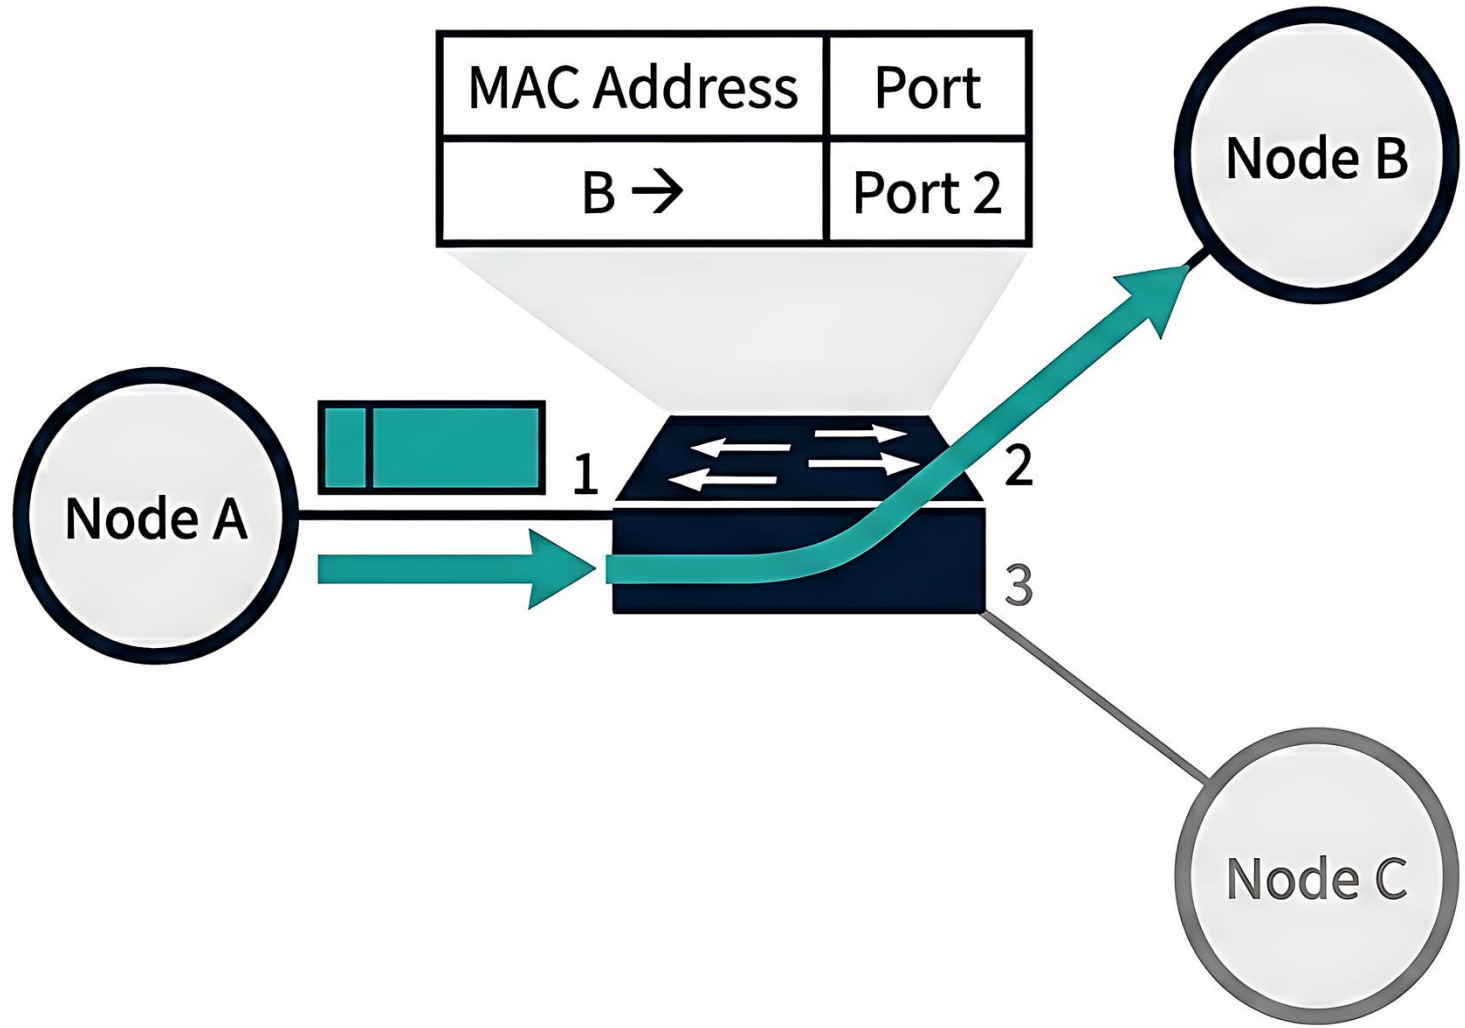

Case 2: Destination Not Found in the Table

- Destination MAC is unknown.
- Switch floods the frame to all ports (except the incoming port).
- Ensures delivery while the system is still learning.

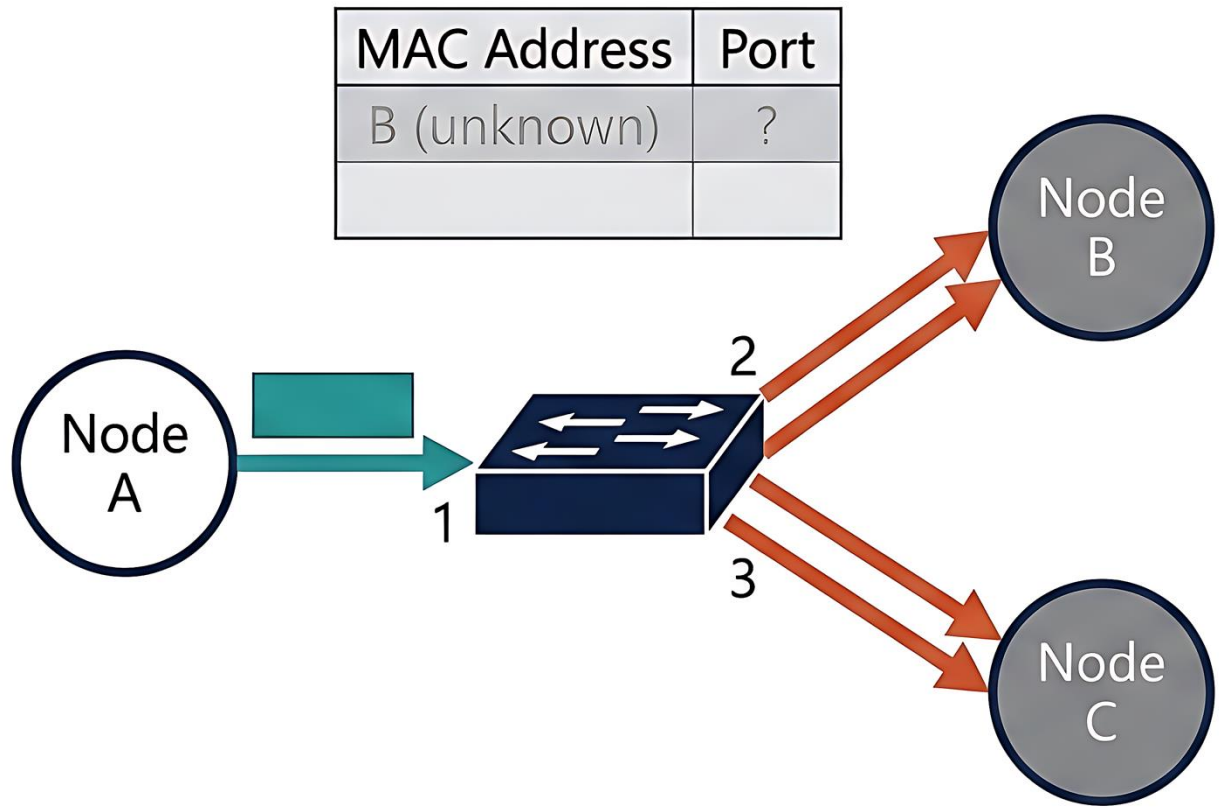

Example: A Sends to B for the First Time

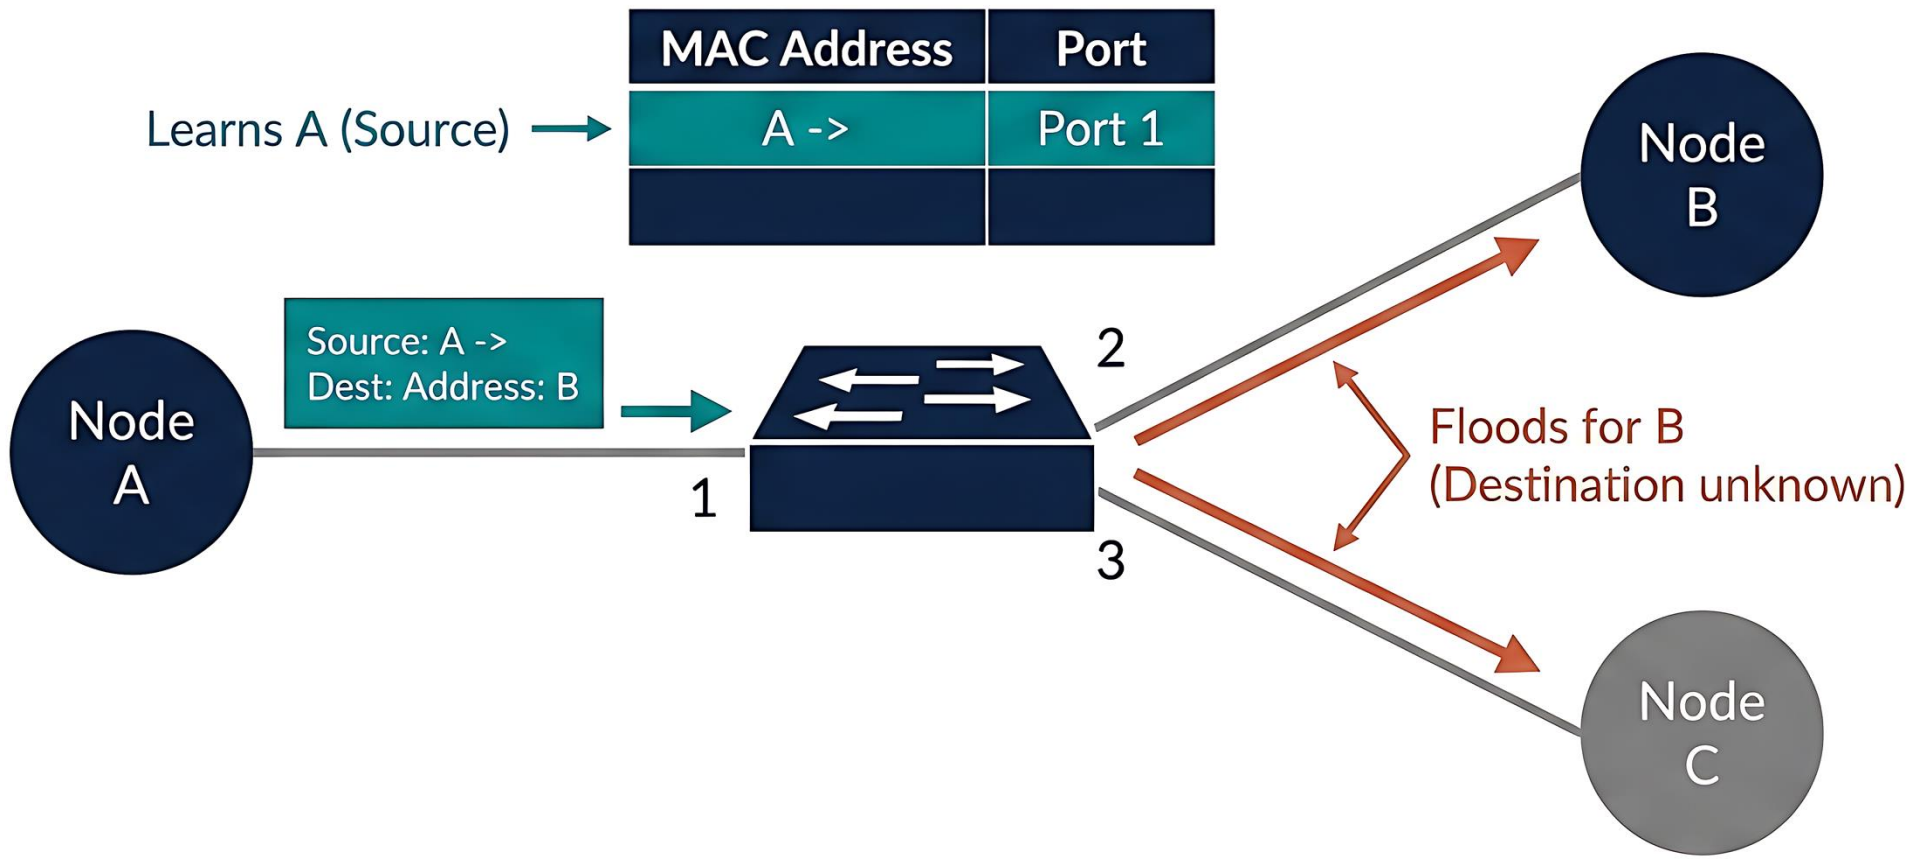

Then B Replies to A

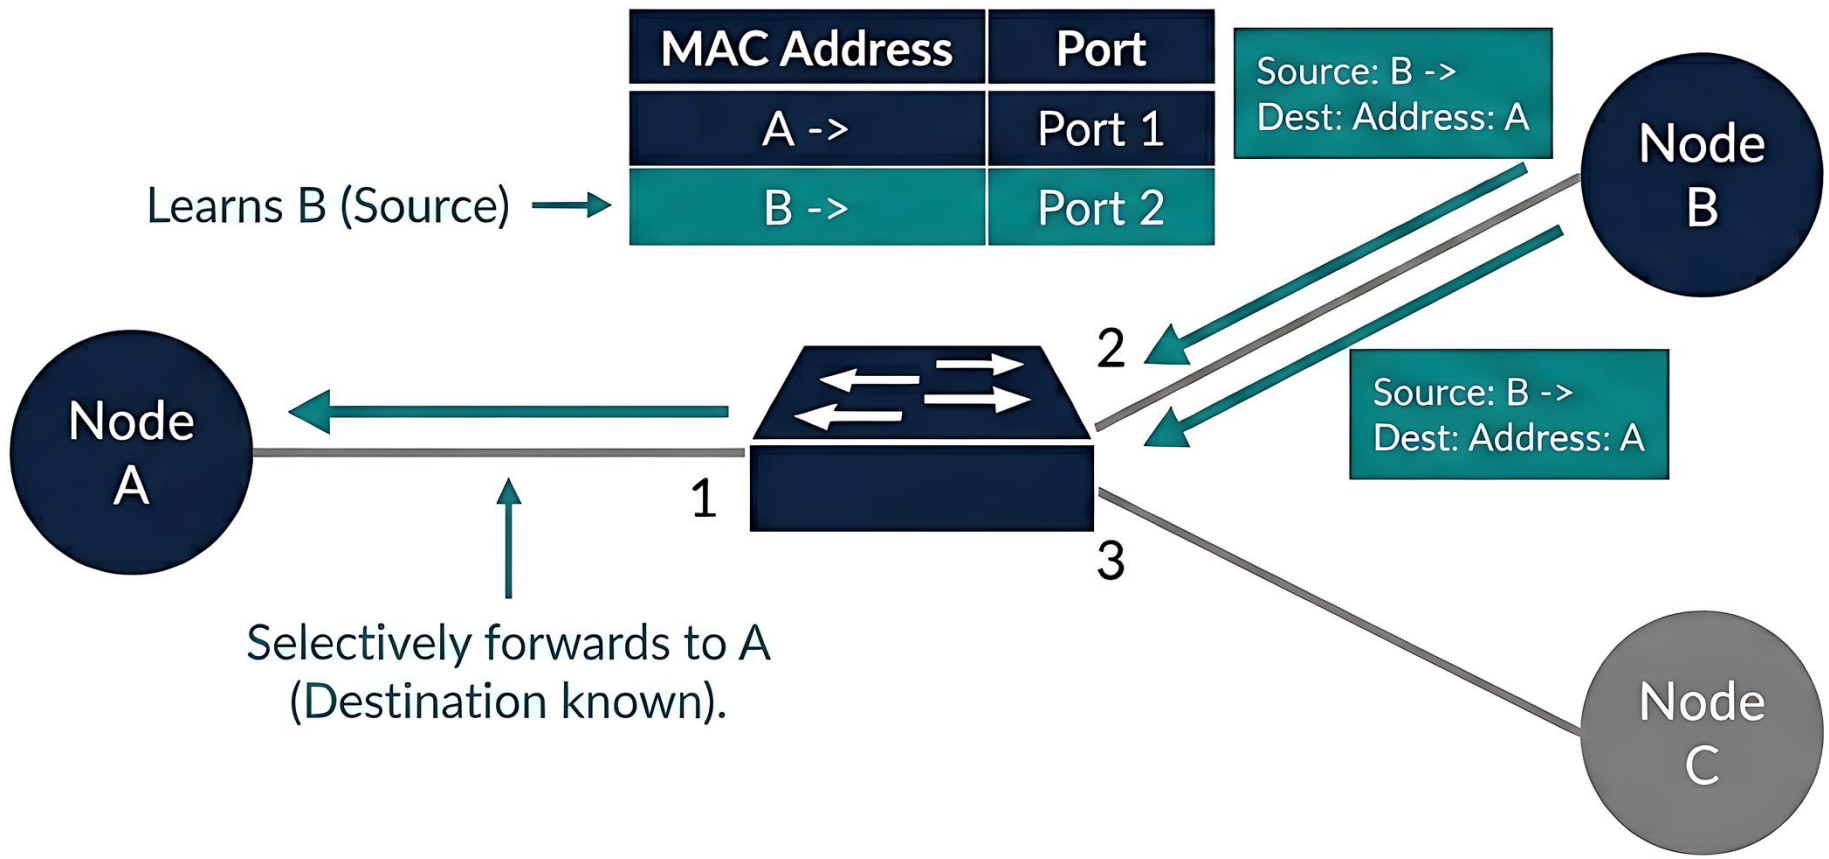

Switch vs. Router: Summary

Both are forwarding devices, but they solve different networking problems.

Aspect	Switch	Router
Main role	Forward inside a LAN	Forward between networks
Main address used	MAC address	IP address
Table used	MAC table	Routing table
Table meaning	MAC → local port	IP Prefix → next hop / interface
Learning method	Self-learning from frames	Static config or routing protocols
Scope	Local network	Across subnets / networks
Typical placement	Inside a LAN	At network boundaries

Both operate at different layers and use distinct information to move data exactly where it needs to go.

Chapter 7
Wireless and Mobile Networks

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

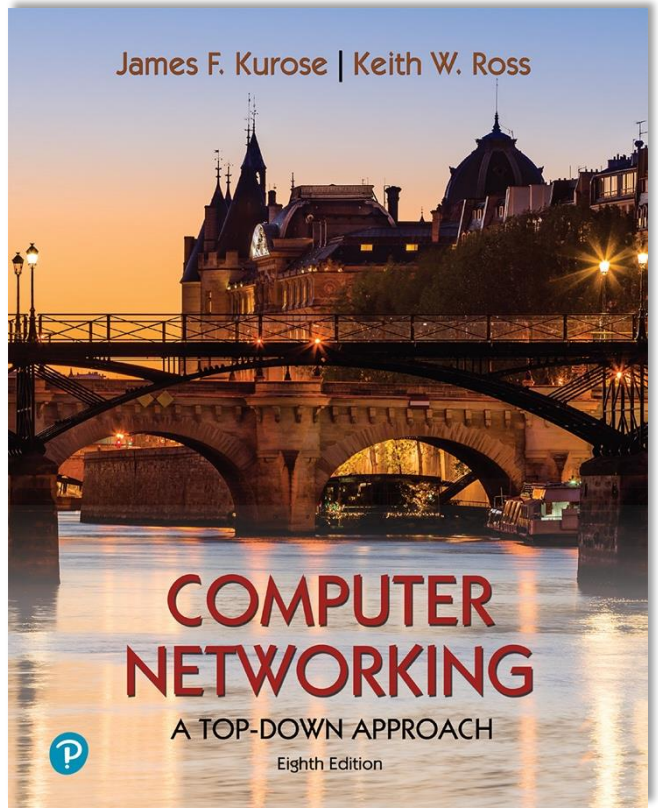

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Adapted from the slides of the book's authors

The Core MAC Question: Who Transmits When?

The MAC layer must answer:

- When should a device transmit?
- When should it wait?
- How can devices avoid transmitting at the same time?
- What should happen after a failed transmission?

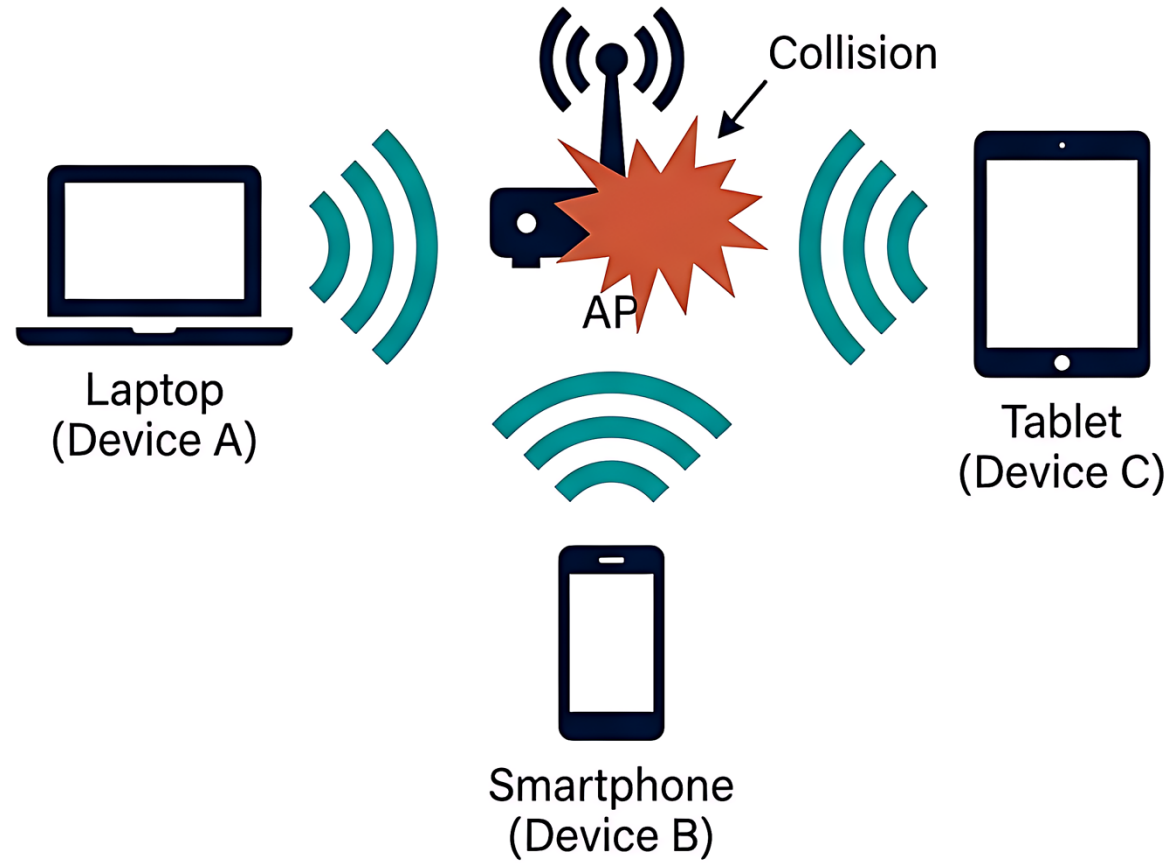

Can We Use CSMA/CD in Wireless Networks?

Wired Ethernet

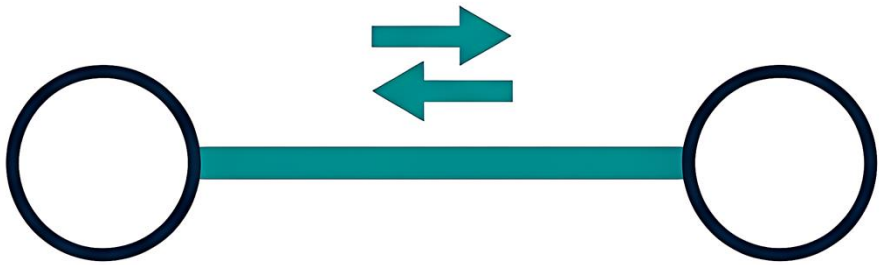

- Transmit immediately when idle

- Detect collision if it happens
- Stop and retry

Wireless Networks

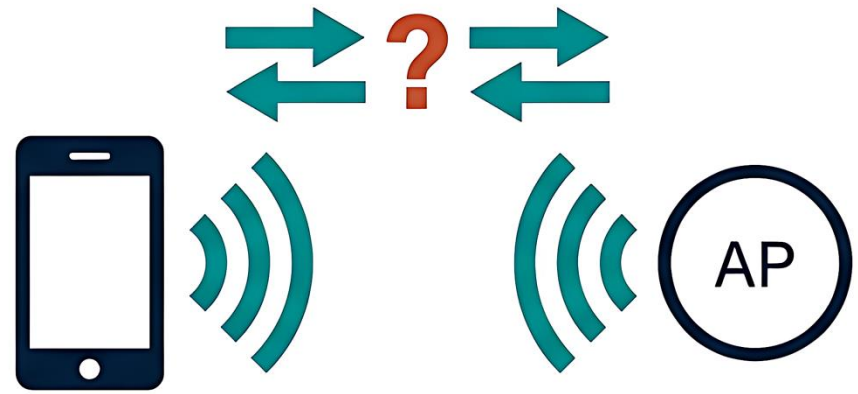

- Can we transmit and listen at the same time?

Can a wireless node detect a collision while transmitting?

Why This Makes Collision Detection Difficult

During transmission:
- Sender's own signal dominates
- Incoming signals are too weak to detect
- Hard to tell if another node is transmitting

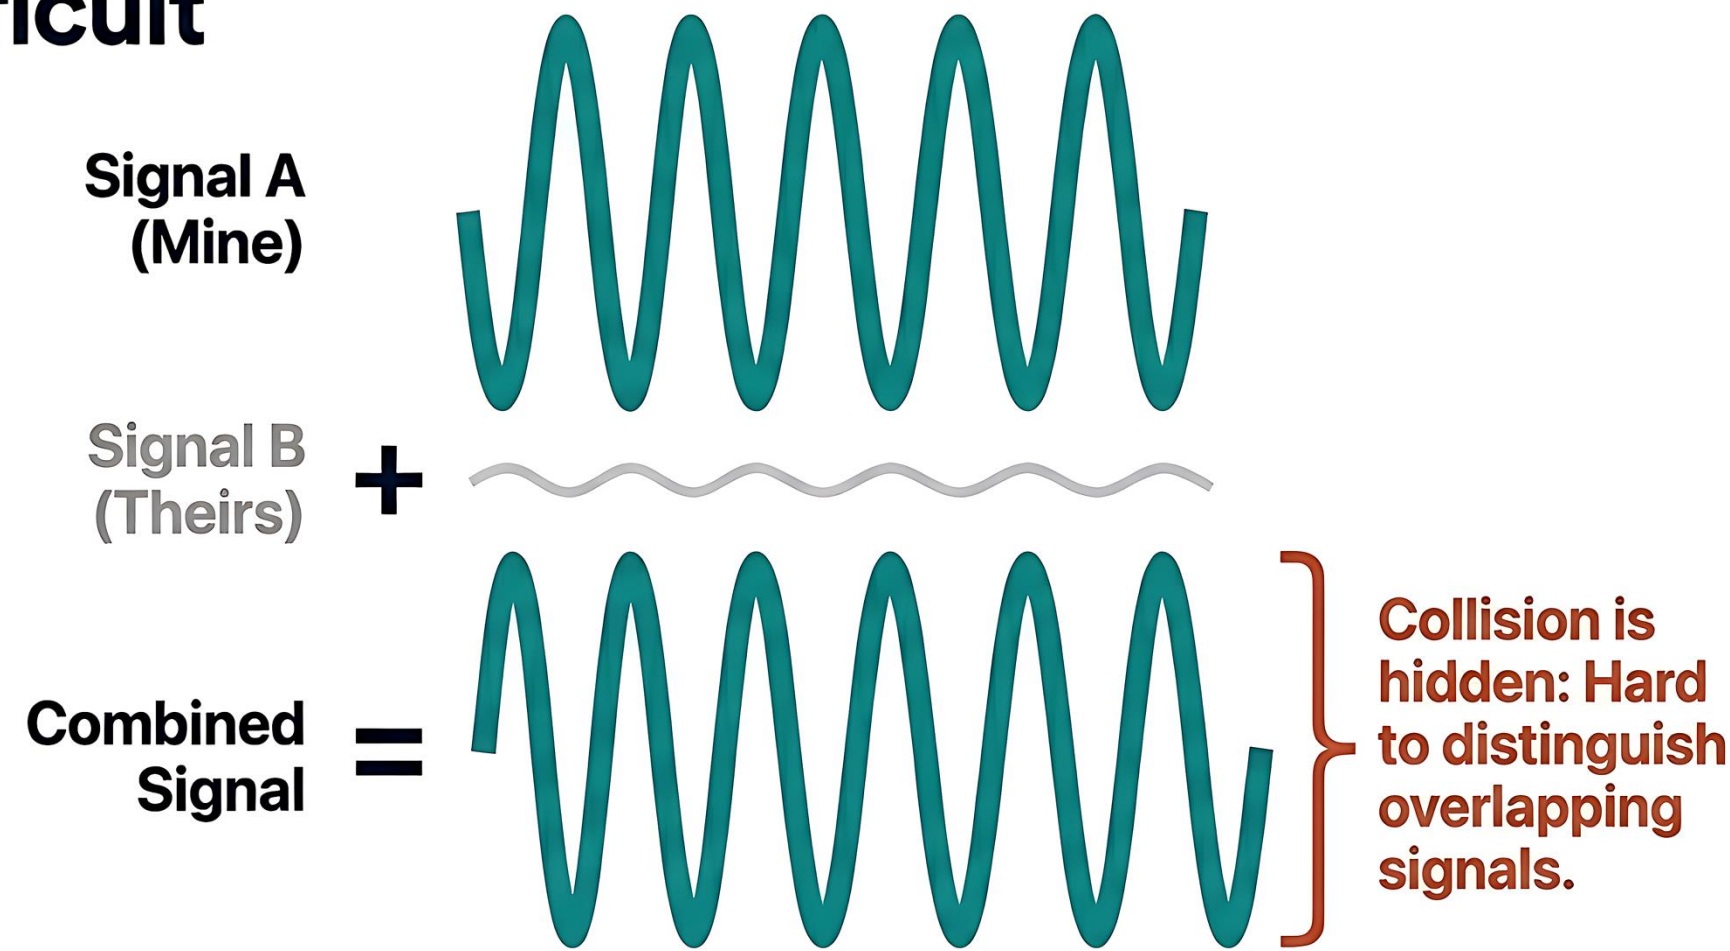

The sender cannot reliably detect a collision while transmitting.

Unreliable Detection Makes Collisions Highly Expensive

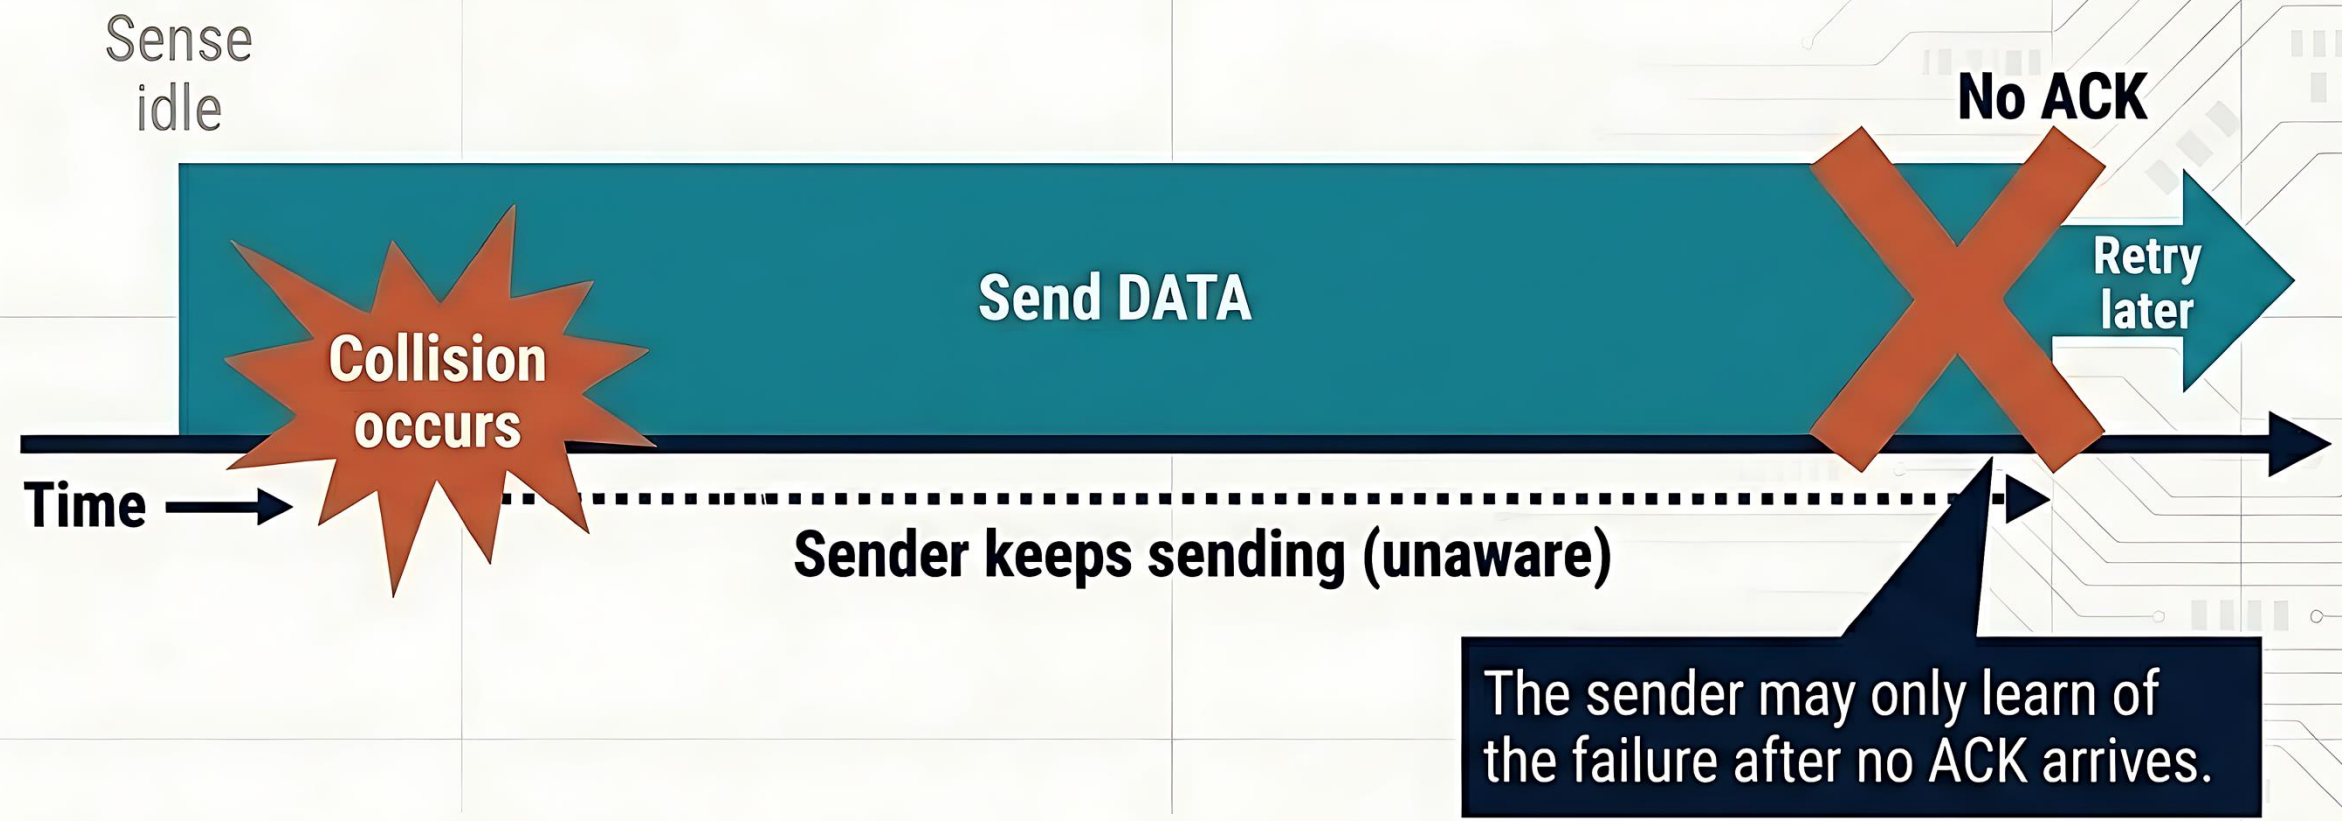

Without reliable detection, collision recovery is slower and vastly more costly.

CSMA/CD Philosophy:

Send first. If collision happens, detect it and recover.

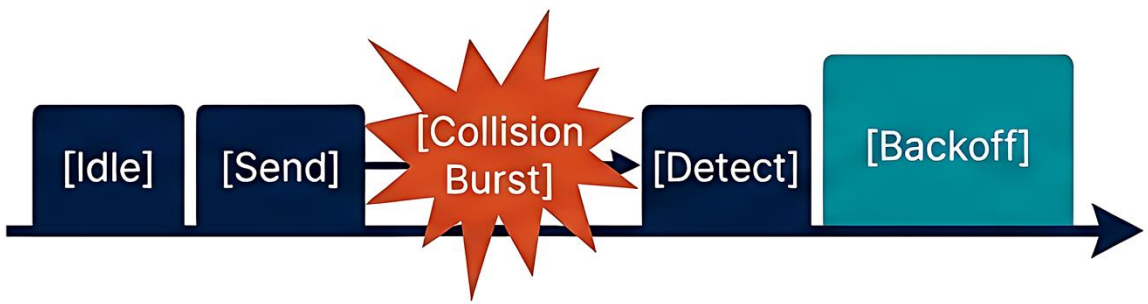

Reactive

Same tool, moved earlier.

CSMA/CA Philosophy:

Since collision is hard to detect, reduce the chance before sending.

Proactive

How an Idle Channel Synchronizes Competitors

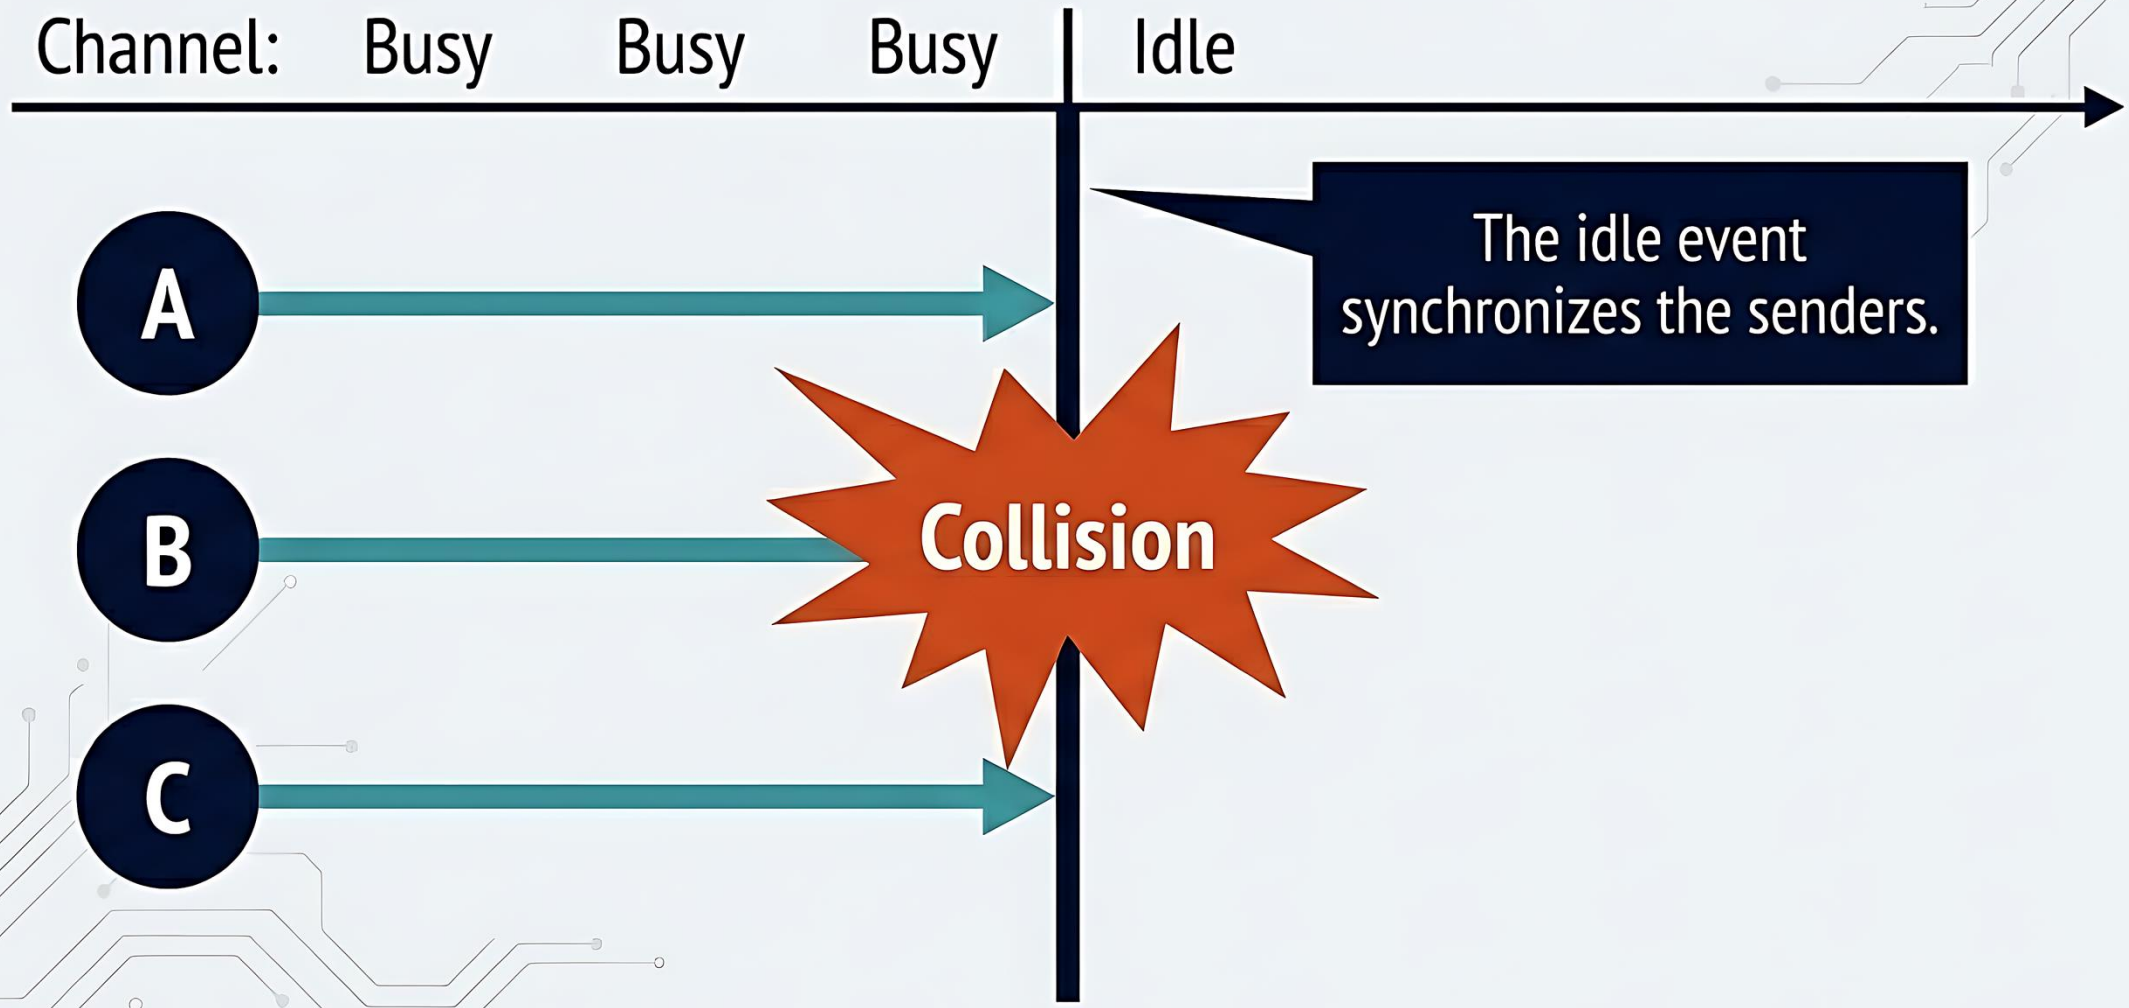

Without pre-backoff, the idle channel triggers many competing senders at once.

Random Backoff Breaks Synchronization

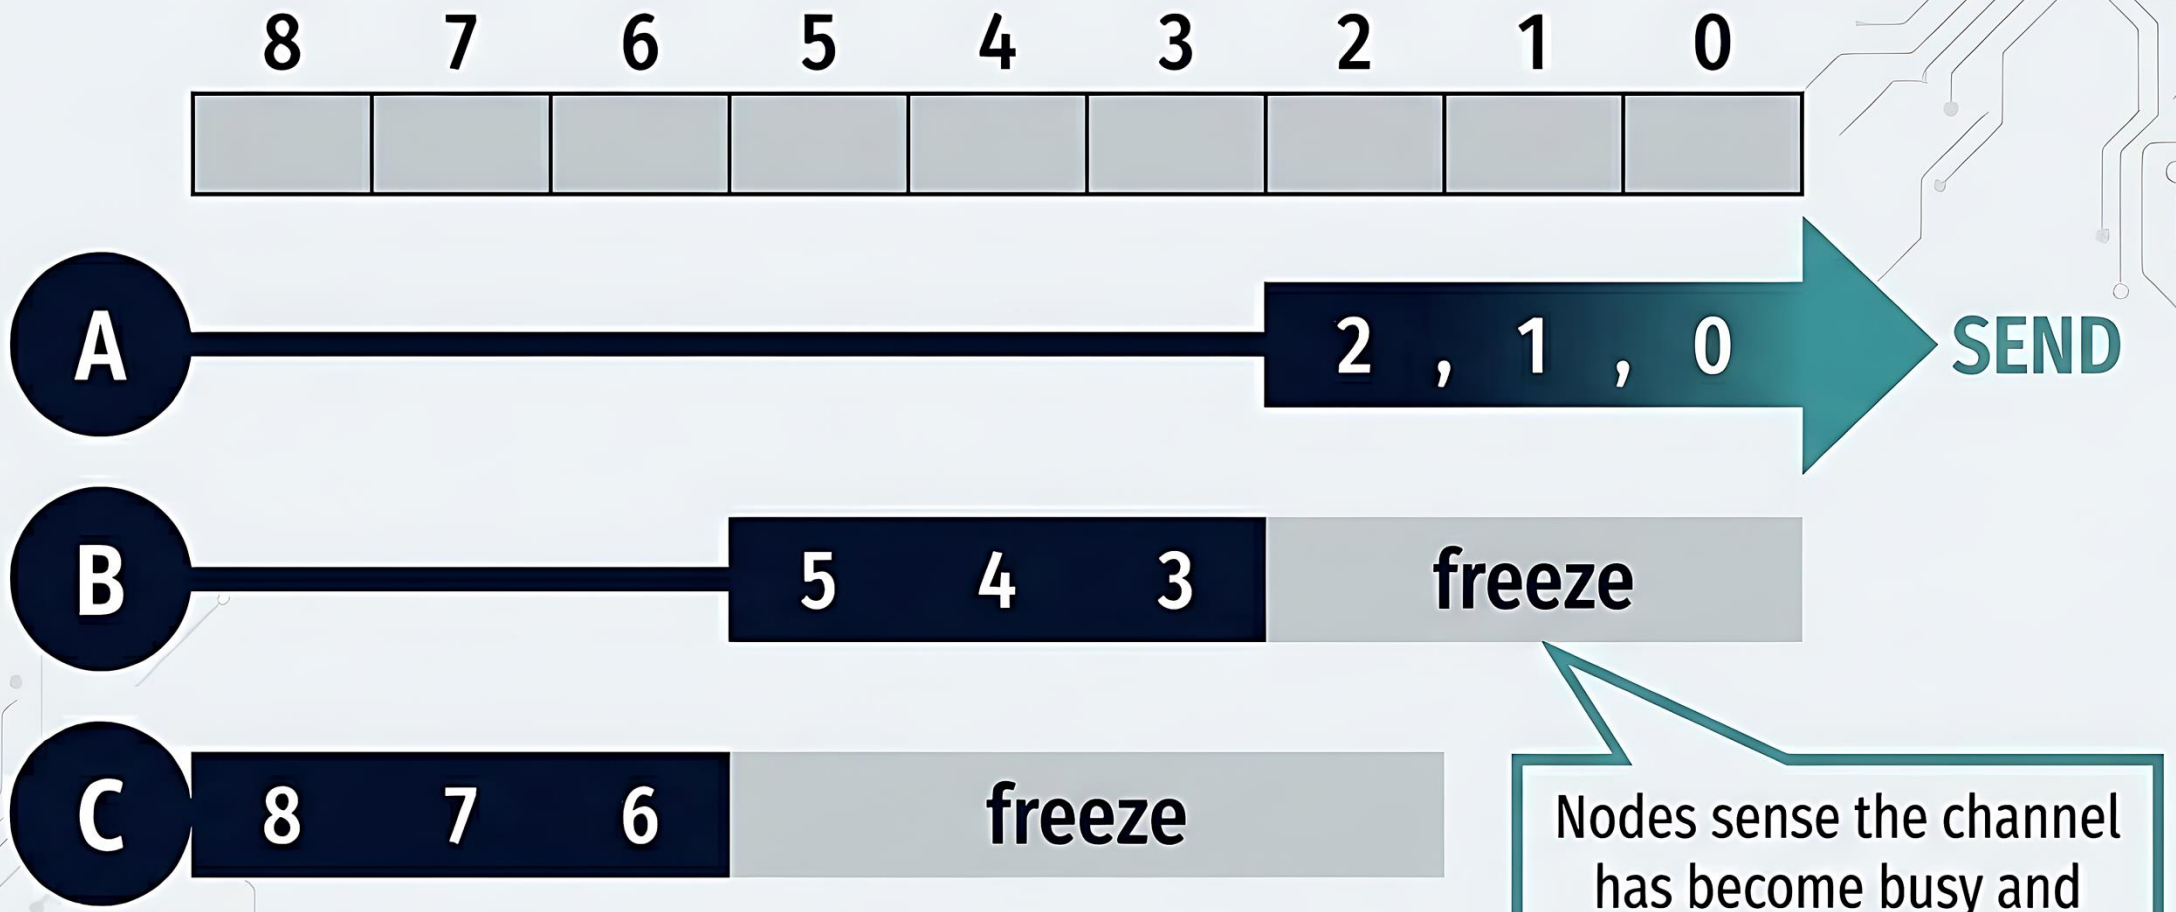

Random waiting desynchronizes competing nodes and spreads attempts over time.

The Hidden Terminal Problem

- A wants to send to AP
- C wants to send to AP
- A cannot hear C, and C cannot hear A

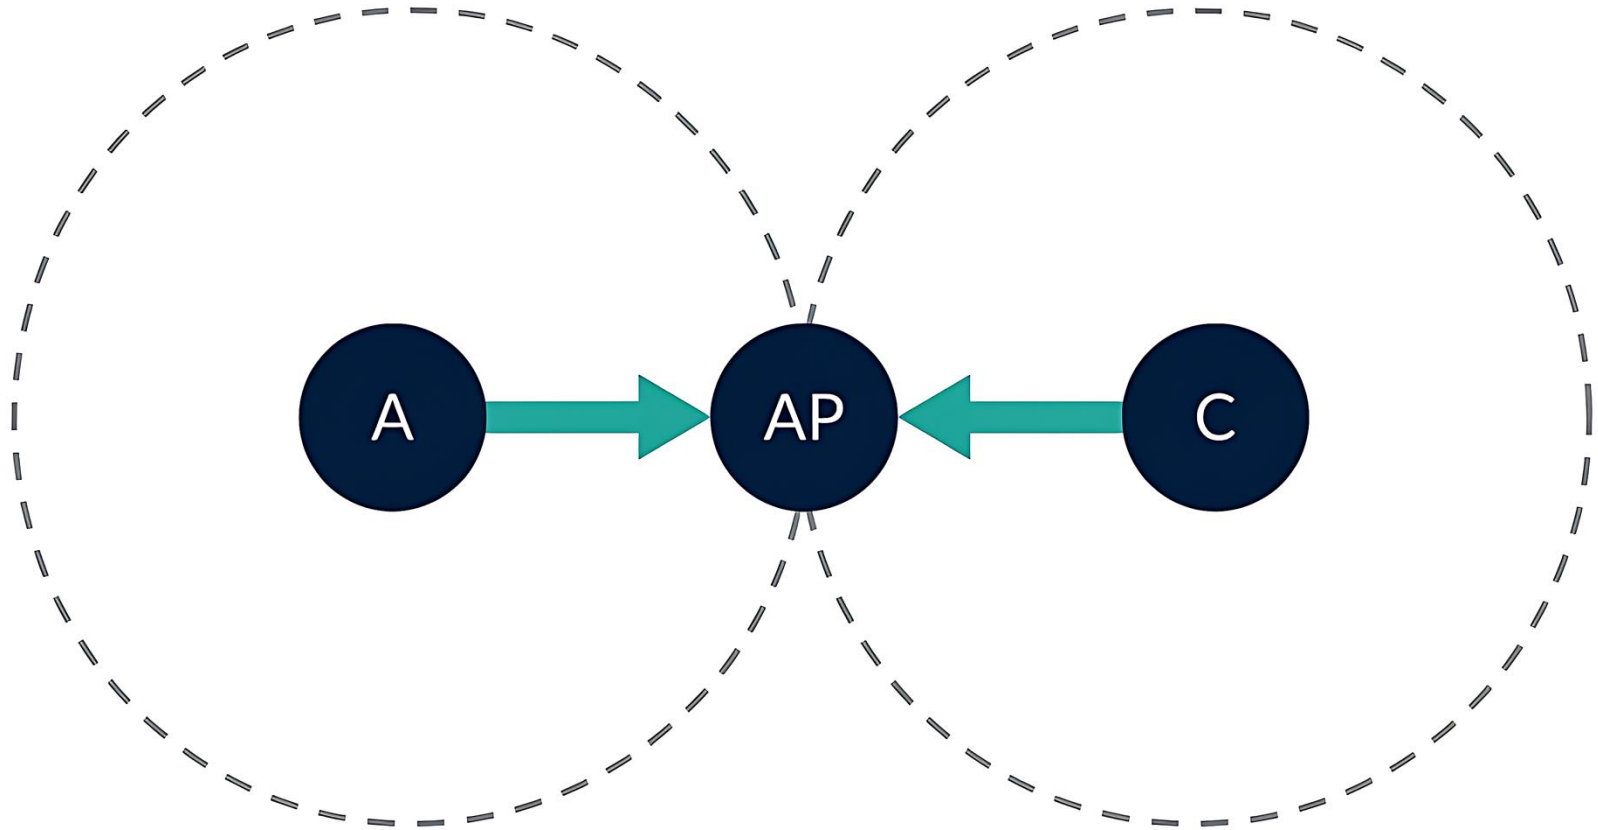

A and C are hidden from each other, but not from the AP.

RTS/CTS: Coordinating Around the Receiver

Let the receiver help reserve the medium using short control frames.

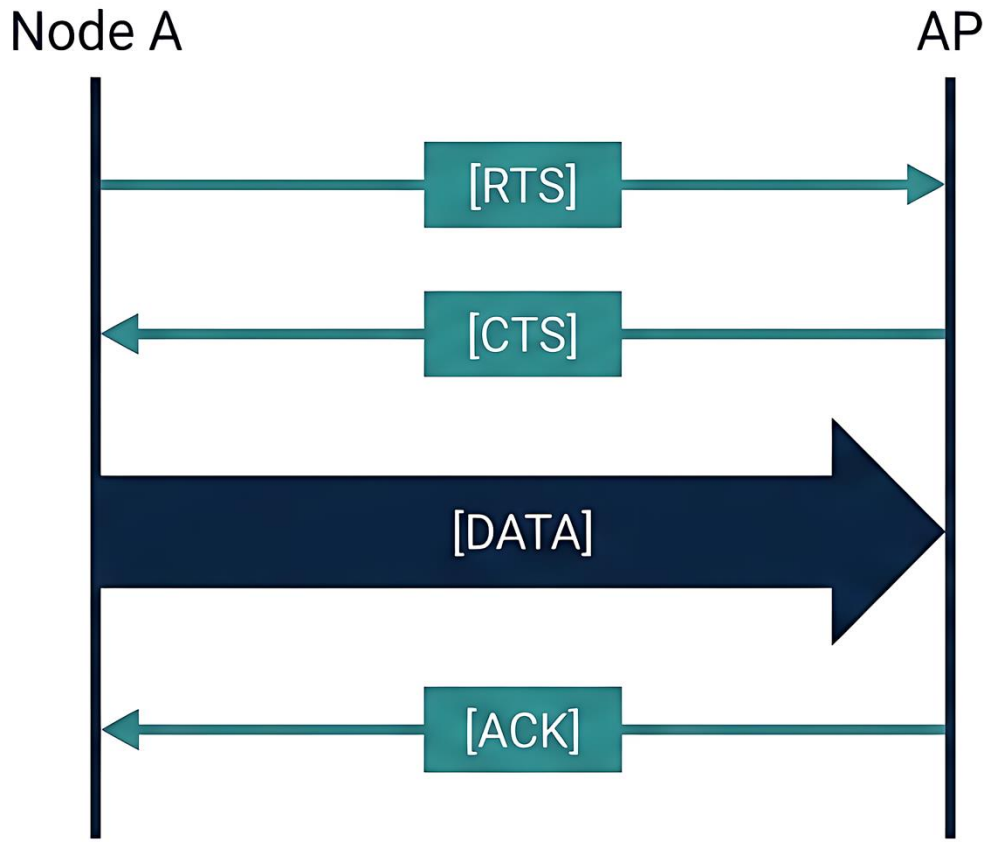

RTS/CTS shifts coordination from the blind sender toward the all-seeing receiver.