COS 461 *Computer Networks*

Class Meeting, Lectures 3 & 4

Kyle Jamieson

Fall 2021

Need to miss 461? Just reach out

Short-Term Remote Teaching and Other Contingency Measures

Under what circumstances might students or faculty need to miss class due to COVID?

Three COVID-related situations would require a student or faculty member to stay in their dorm or at home (as applicable) during the semester:

1. Anyone experiencing **COVID-like symptoms** should stay home/in their dorm room until they are tested and receive a negative test result.
2. Anyone who receives a positive COVID test result will be required to self-isolate. **Isolation** will be a minimum of 10 days; isolation longer than 10 days is rare, but possible. Isolation means restricting activities outside of one's home or isolation dorm, except for obtaining medical care when necessary.

IP Protocol Stack: Key Abstractions

Best-Effort Global Packet Delivery

Circuit Switching (e.g., Phone Network)

- Source establishes connection
  - Reserve resources along hops in the path
- Source sends data
  - Transmit data over the established connection
- Source tears down connection
  - Free the resources for future connections

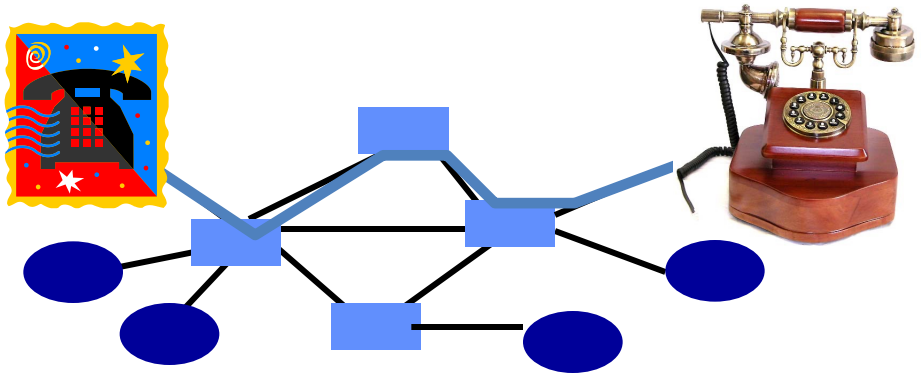

Circuit Switching: Static Allocation

- Time-division
  - Each circuit allocated certain time slots
- Frequency-division
  - Each circuit allocated certain frequencies

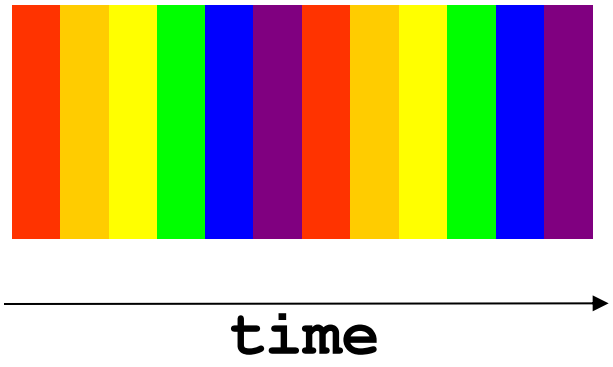

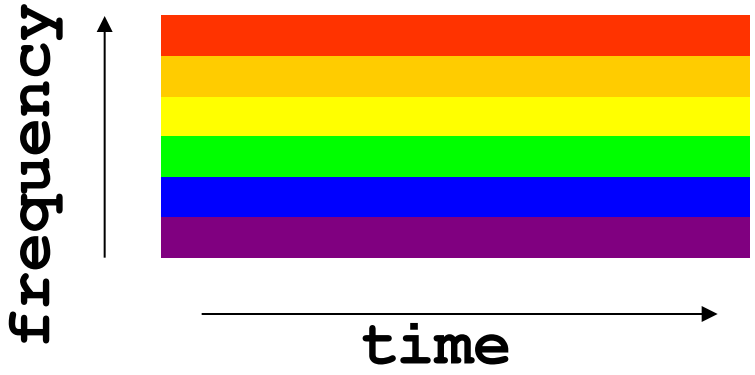

Packet Switching

- Message divided into packets
  - Header identifies the destination address
- Packets travel separately through the network
  - Forwarding based on the destination address
  - Packets may be buffered temporarily
- Destination reconstructs the message

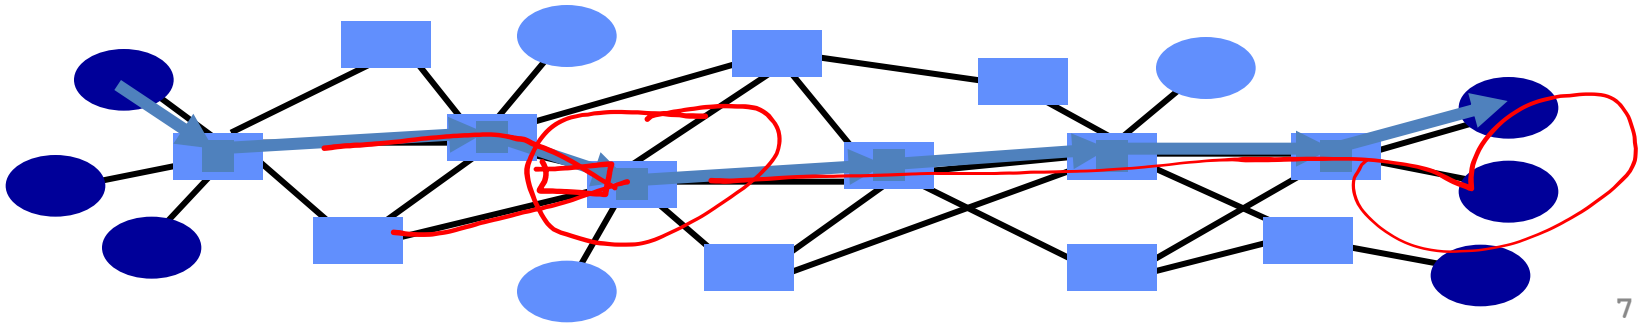

Is Best Effort Good Enough?

- Packet loss and delay
  - Sender can resend
- Packet corruption
  - Receiver can detect, and sender can resend
- Out-of-order delivery
  - Receiver can put the data back in order
- Packets follow different paths
  - Doesn't matter
- Network failure
  - Drop the packet
- Network congestion
  - Drop the packet

Packet (Y) vs. Circuit Switching (A)?

• Predictable performance	Circuit
• Network never blocks senders	Packet
• Reliable, in-order delivery	Circuit
• Low delay to send data	Packet
• Simple forwarding	Circuit
• No overhead for packet headers	Circuit
• High utilization under most workloads	Packet
• No per-connection network state	Packet

Network Addresses

Grouping Related Hosts

- The Internet is an "inter-network"
  - Used to connect networks together, not hosts
  - Need to address a network (i.e., group of hosts)

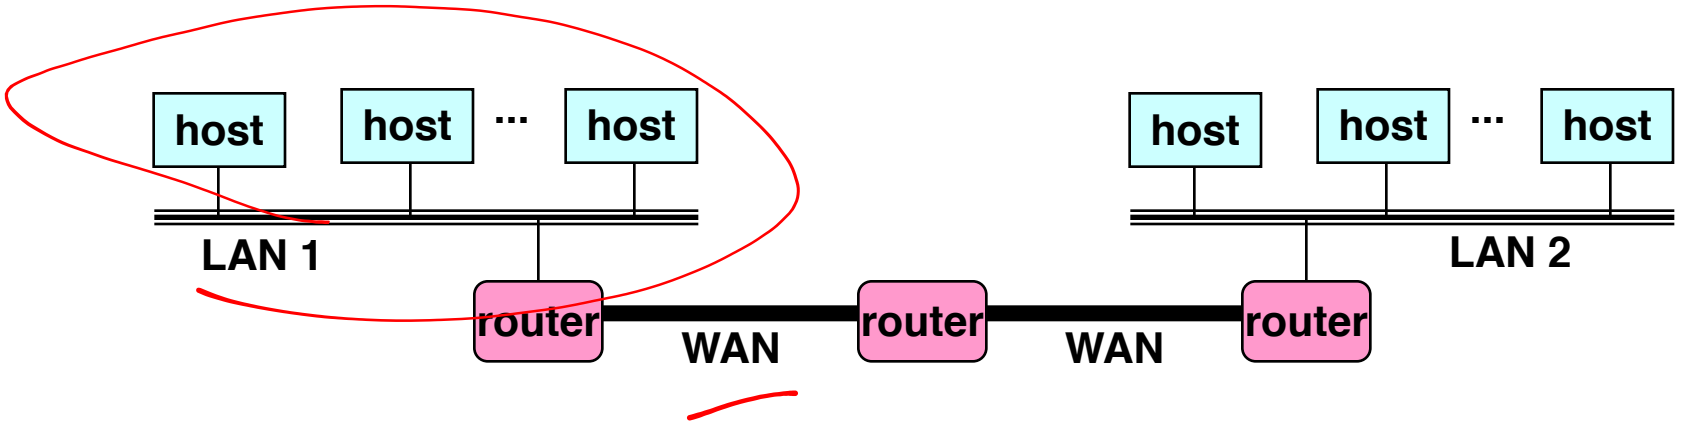

LAN = Local Area Network

WAN = Wide Area Network

Scalability Challenge

- Suppose hosts had arbitrary addresses
  - Then every router would need a lot of information
  - …to know how to direct packets toward every host

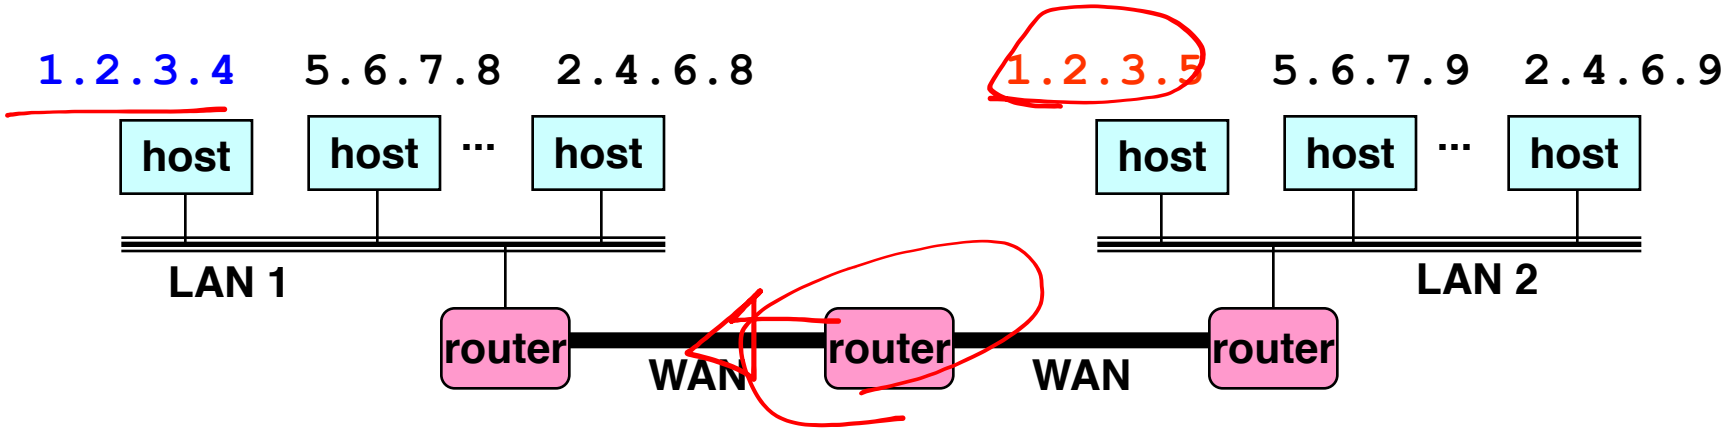

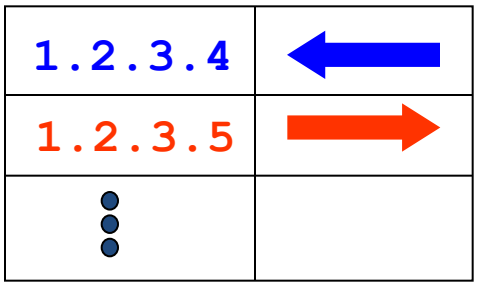

forwarding table

Hierarchical Addressing: IP Prefixes

- Network and host portions (left and right)
- 12.34.158.0/24 is a 24-bit **prefix** with 2^8 addresses

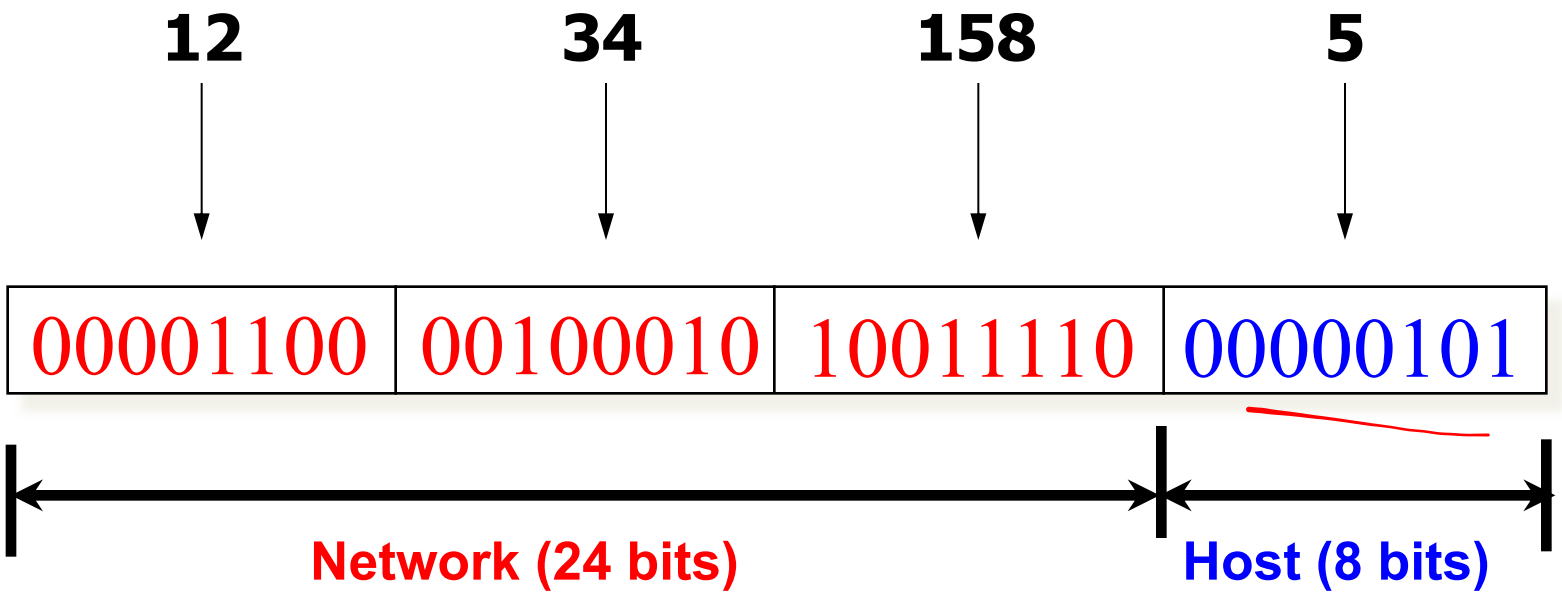

Scalability Improved

- Number related hosts from a common subnet
  - 1.2.3.0/24 on the left LAN
  - 5.6.7.0/24 on the right LAN

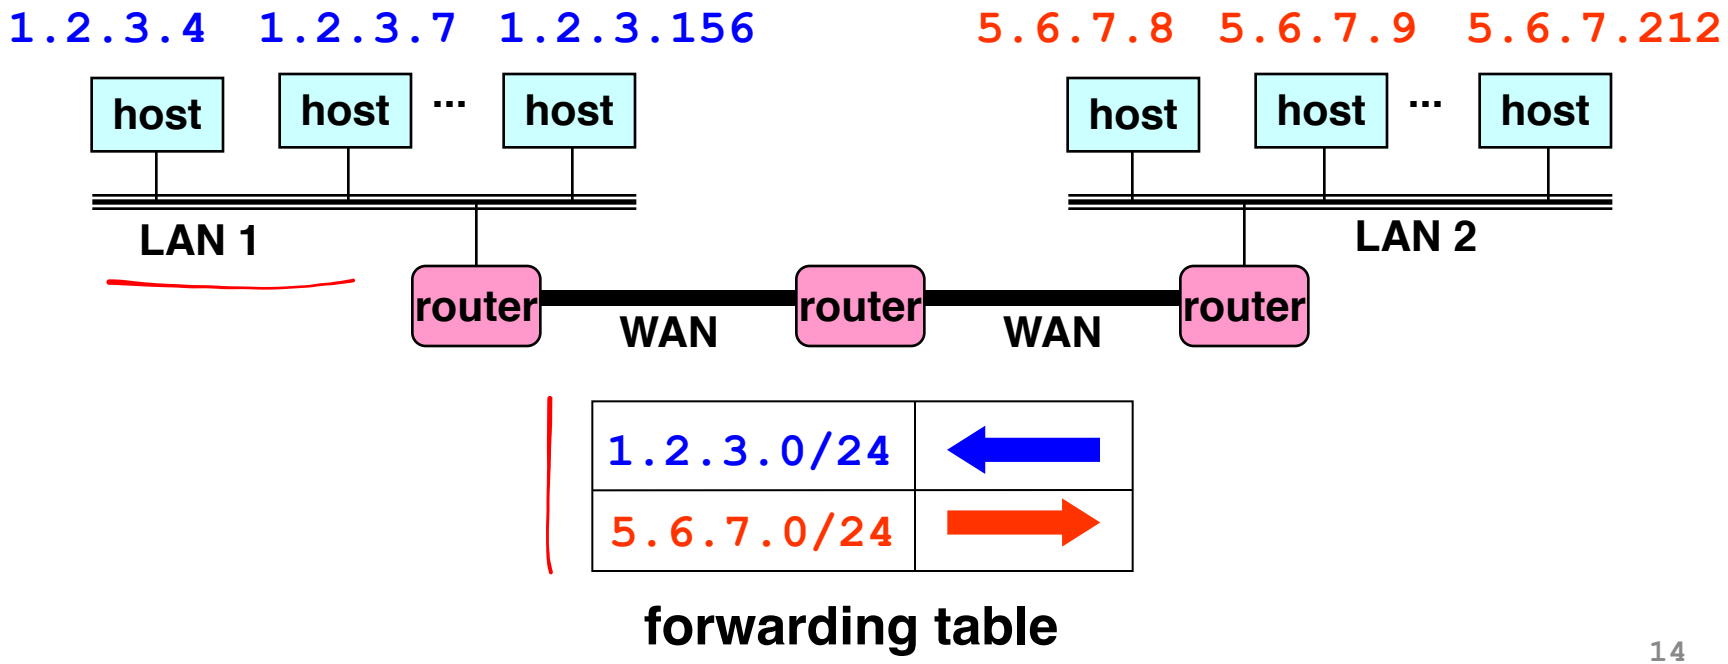

Easy to Add New Hosts

- No need to update the routers
  - E.g., adding a new host 5.6.7.213 on the right
  - Doesn't require adding a new forwarding-table entry

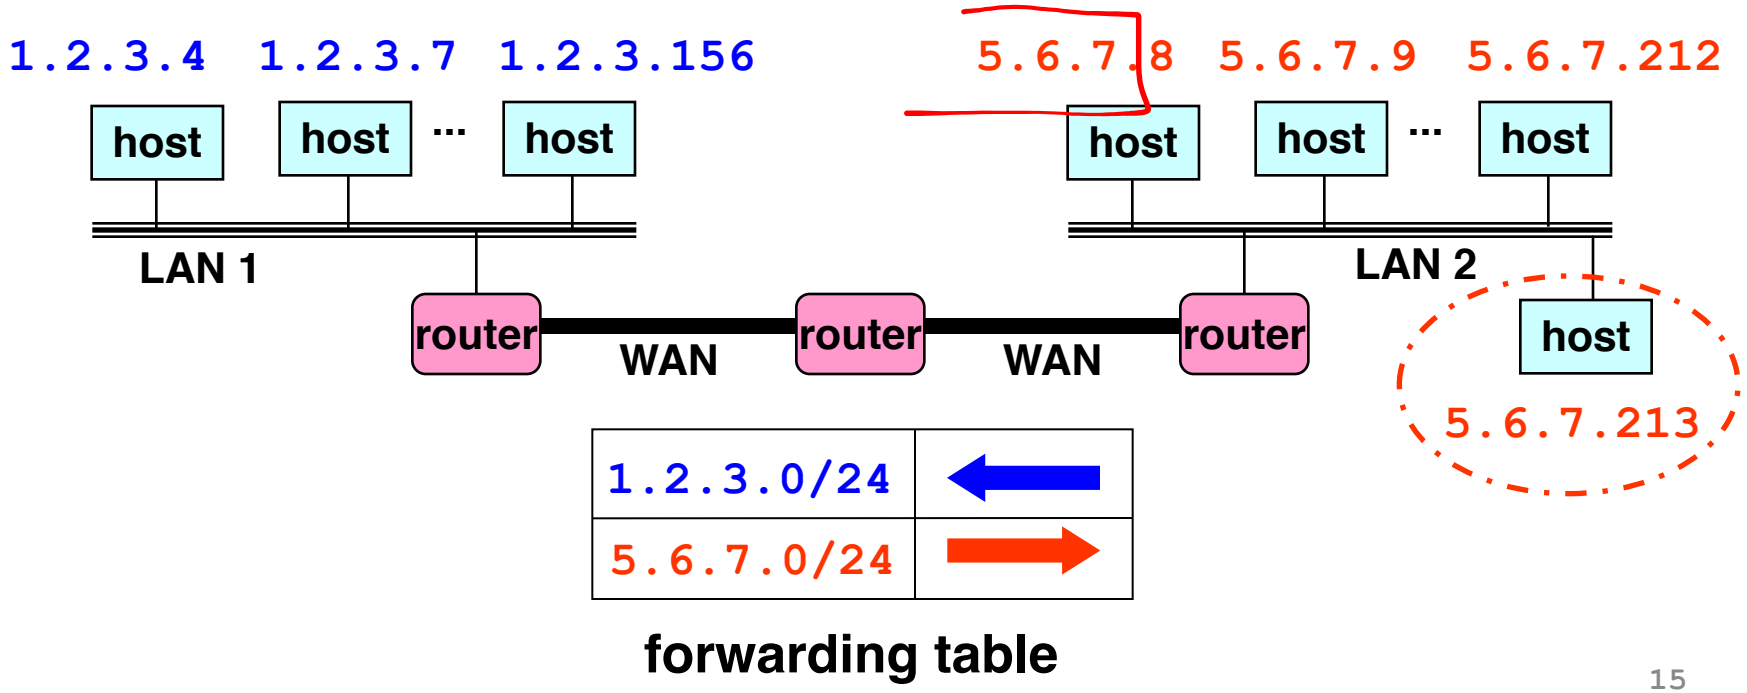

Classless Inter-Domain Routing (CIDR)

- Use two 32-bit numbers to represent network:

  Network number = IP address + Mask

IP Address : 12.4.0.0 **IP Mask: 255.254.0.0**

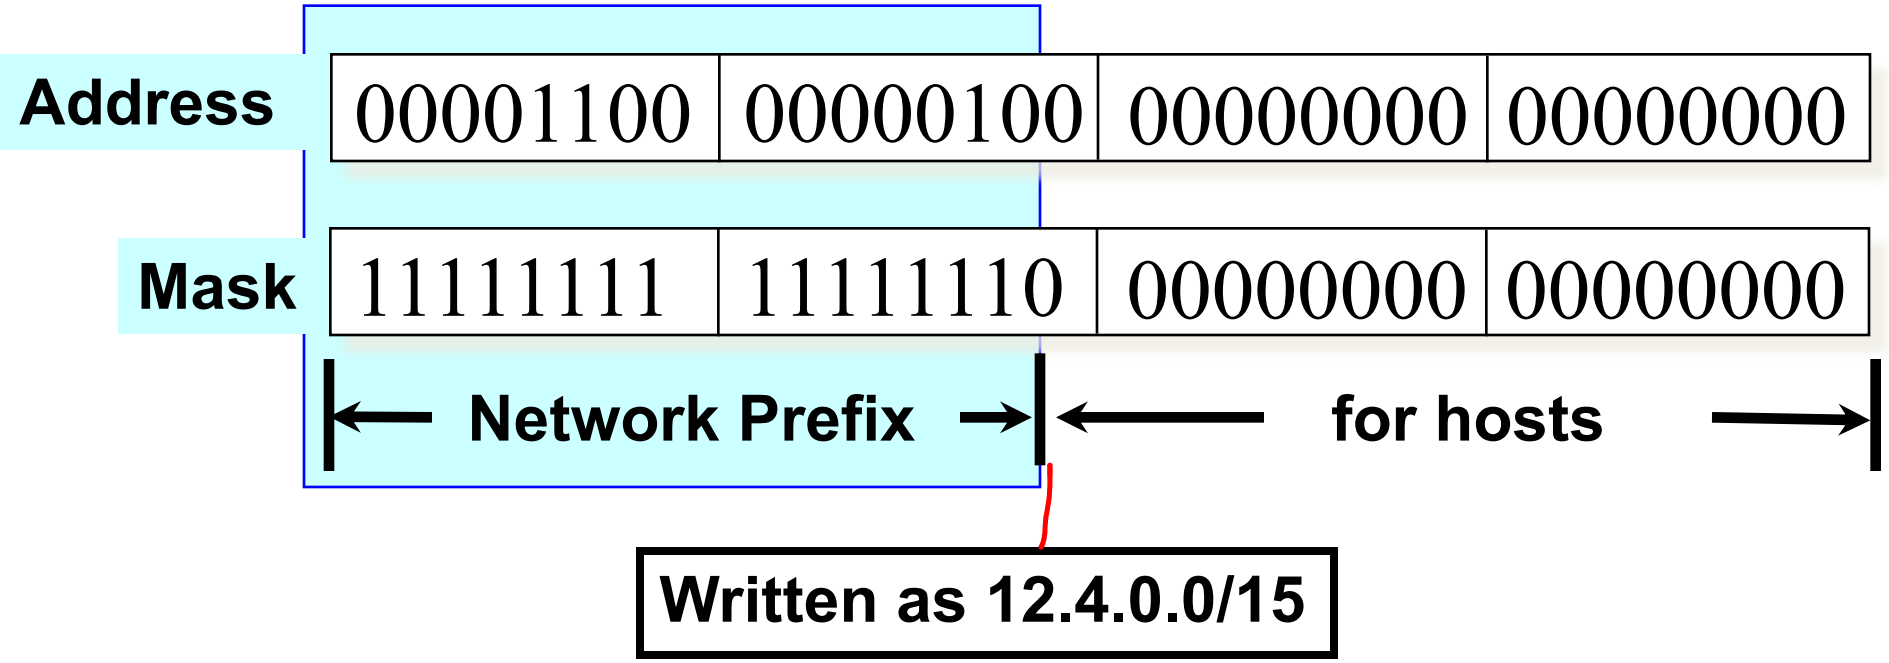

Hierarchical Address Allocation

- Hierarchy is key to scalability
  - Address allocated in contiguous chunks (prefixes)
  - Today, the Internet has about 600-800,000 prefixes

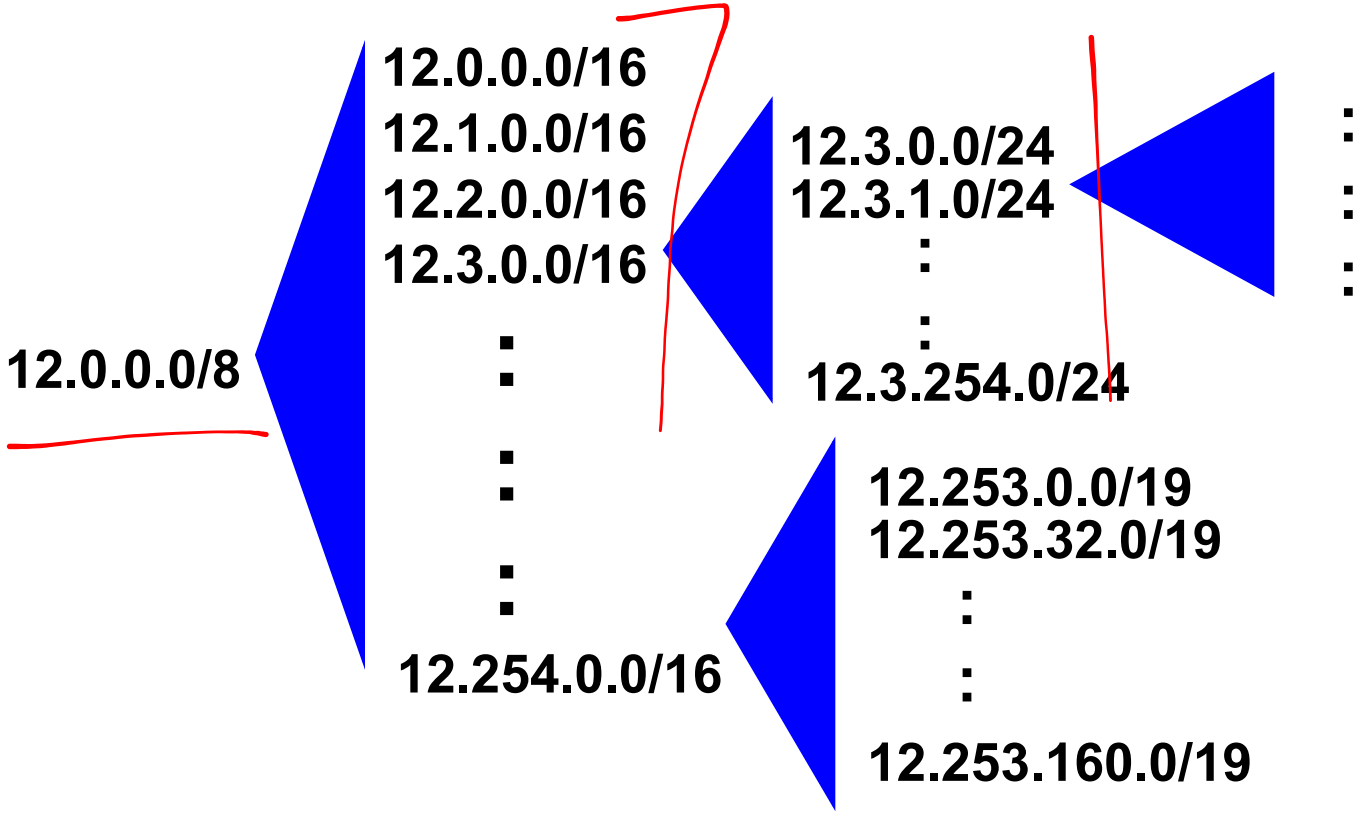

Packet Forwarding

Hop-by-Hop Packet Forwarding

- Each router has a forwarding table
  - Maps destination address to outgoing interface
- Upon receiving a packet
  - Inspect the destination address in the header
  - Index into the table
  - Determine the outgoing interface
  - Forward the packet out that interface
- Then, the next router in the path repeats

Separate Forwarding Entry Per Prefix

- Prefix-based forwarding
  - Map the destination address to matching prefix
  - Forward to the outgoing interface

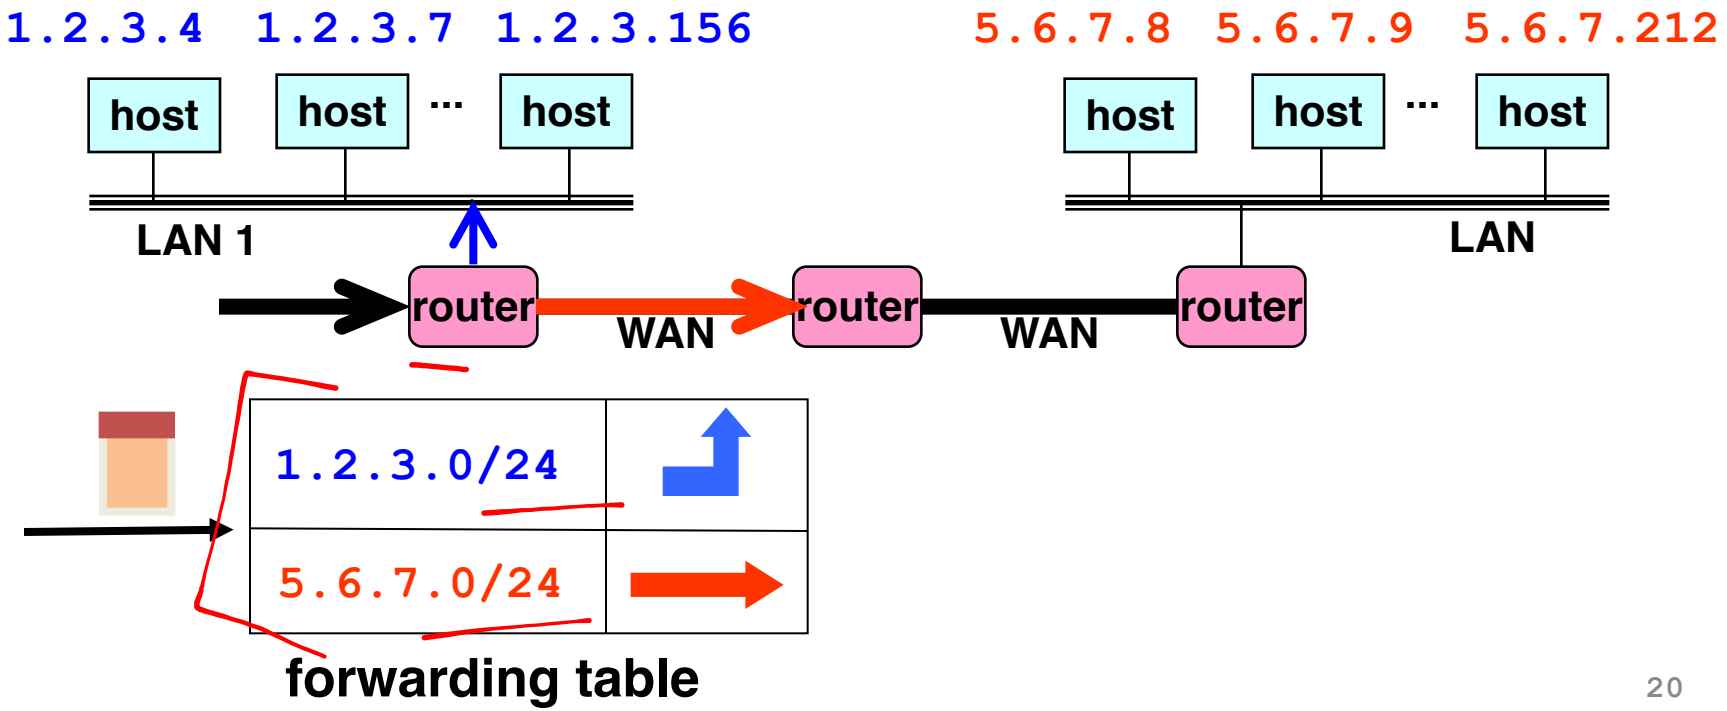

CIDR Makes Packet Forwarding Harder

- Forwarding table may have many matches
  - E.g., entries for 201.10.0.0/21 and 201.10.6.0/23
  - The IP address 201.10.6.17 would match both!

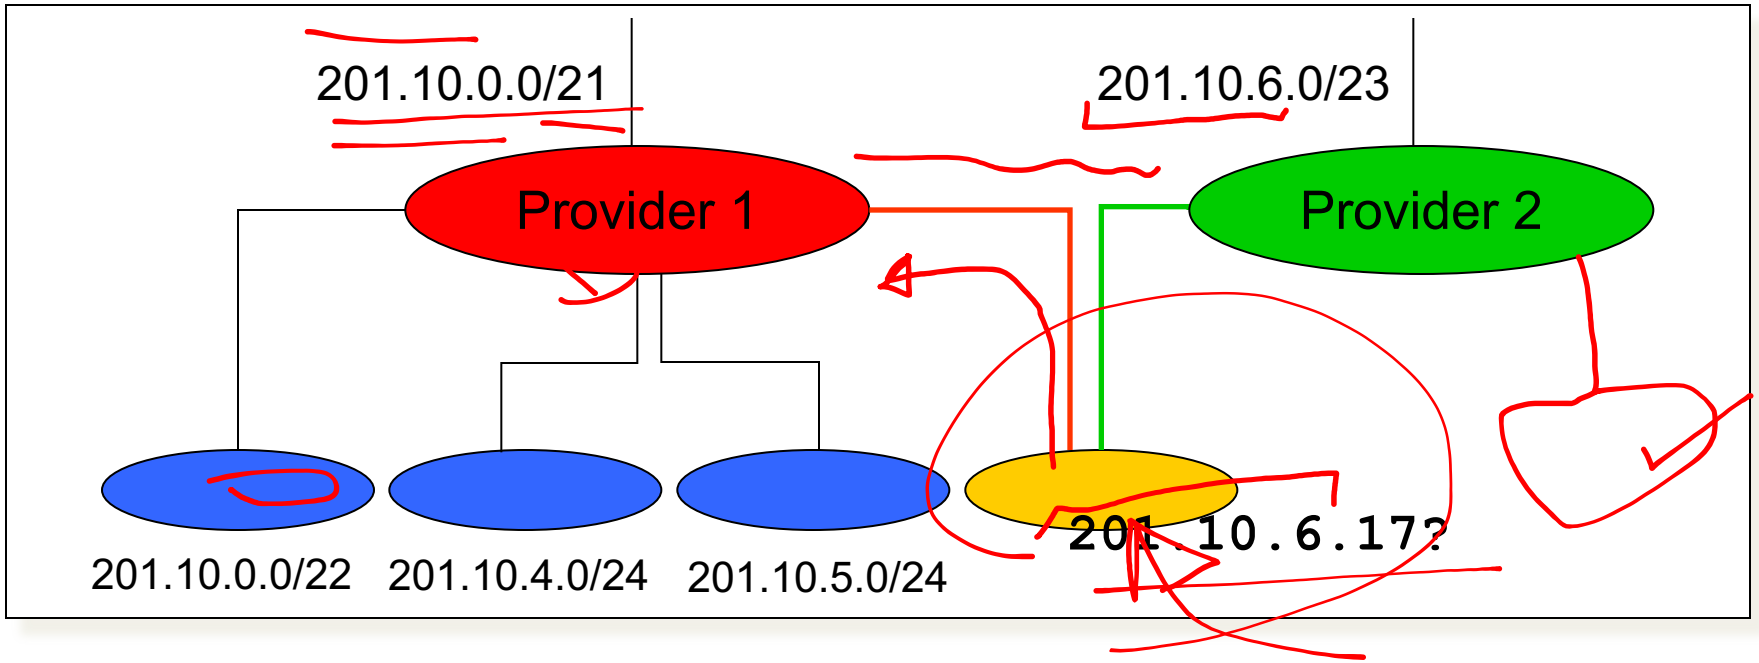

Longest Prefix Match Forwarding

- Destination-based forwarding
  - Packet has a destination address
  - Router identifies longest-matching prefix
  - Cute algorithmic problem: very fast lookups

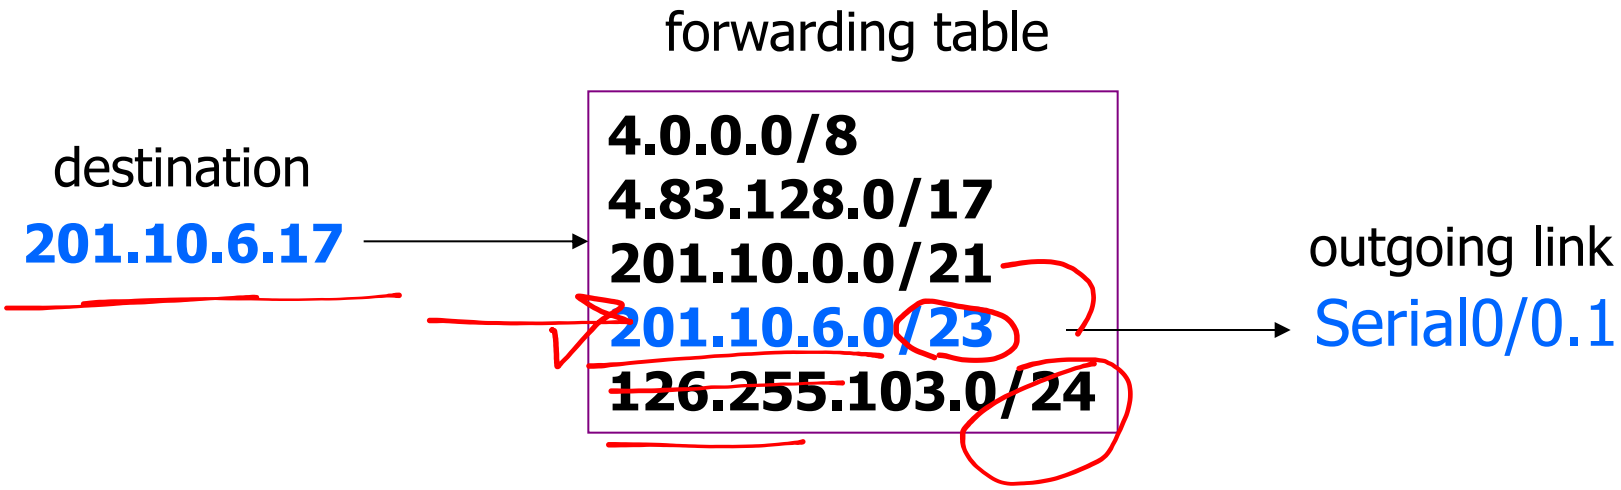

“Layer 2”
Hubs and Switches

Bridges/Switches: Traffic Isolation

- Switch filters packets
  - Frame only forwarded to the necessary segments
  - Segments can support separate transmissions

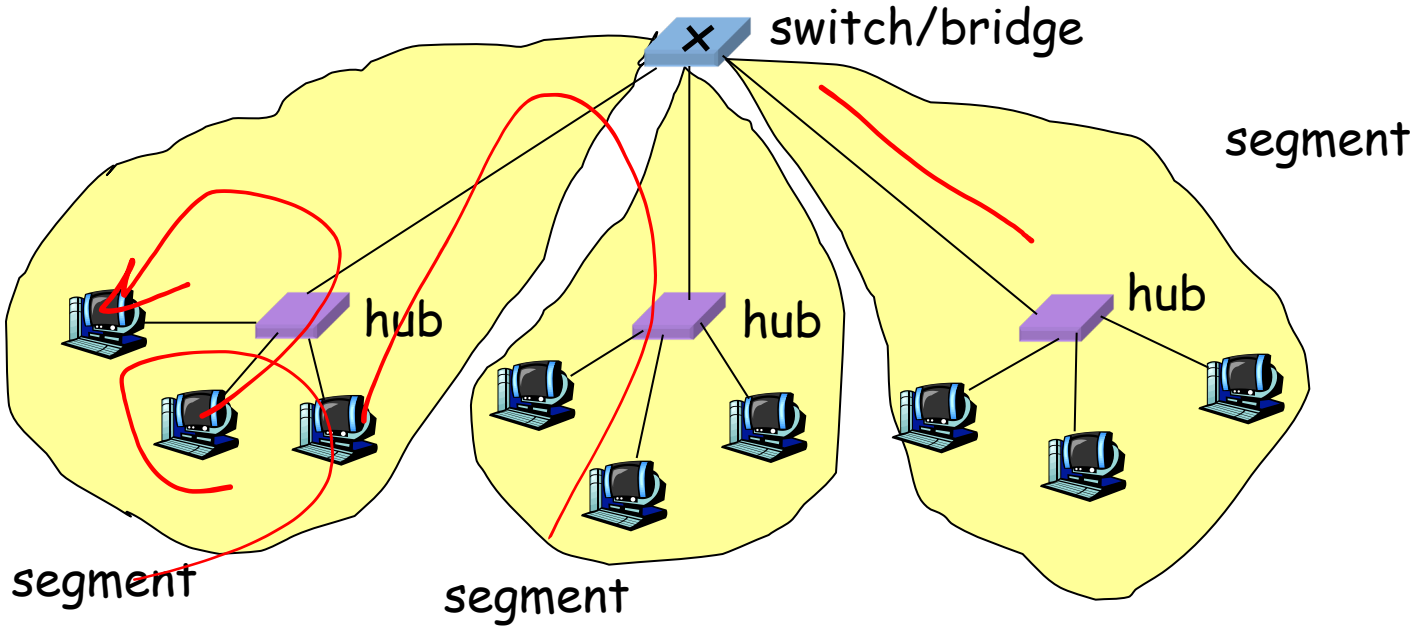

Self Learning: Building the Table

- When a frame arrives
  - Inspect the *source* MAC address
  - Associate the address with the *incoming* interface
  - Store the mapping in the switch table
  - Use a timer to eventually forget the mapping

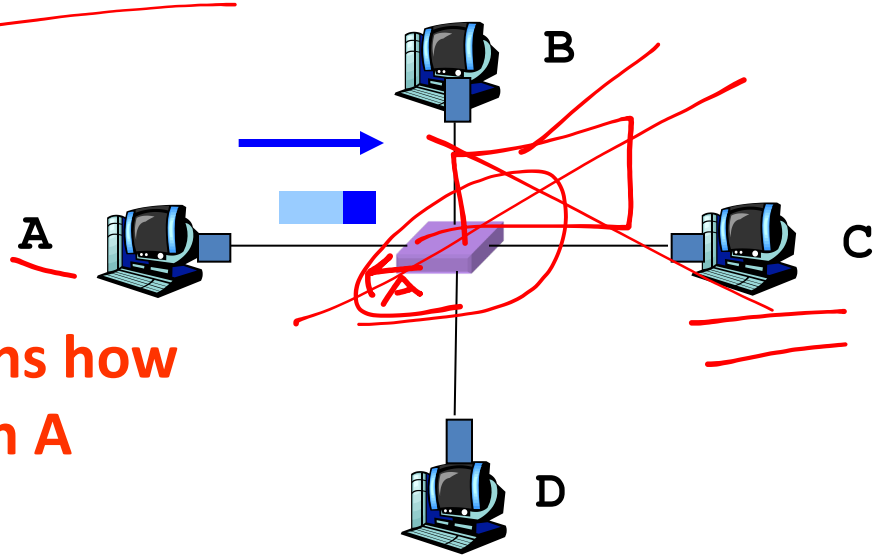

Switch learns how to reach A

Self Learning: Handling Misses

- When frame arrives with unfamiliar destination
  - Forward the frame out all of the interfaces
  - … except for the one where the frame arrived
  - Hopefully, this case won't happen very often!

When in doubt, shout!

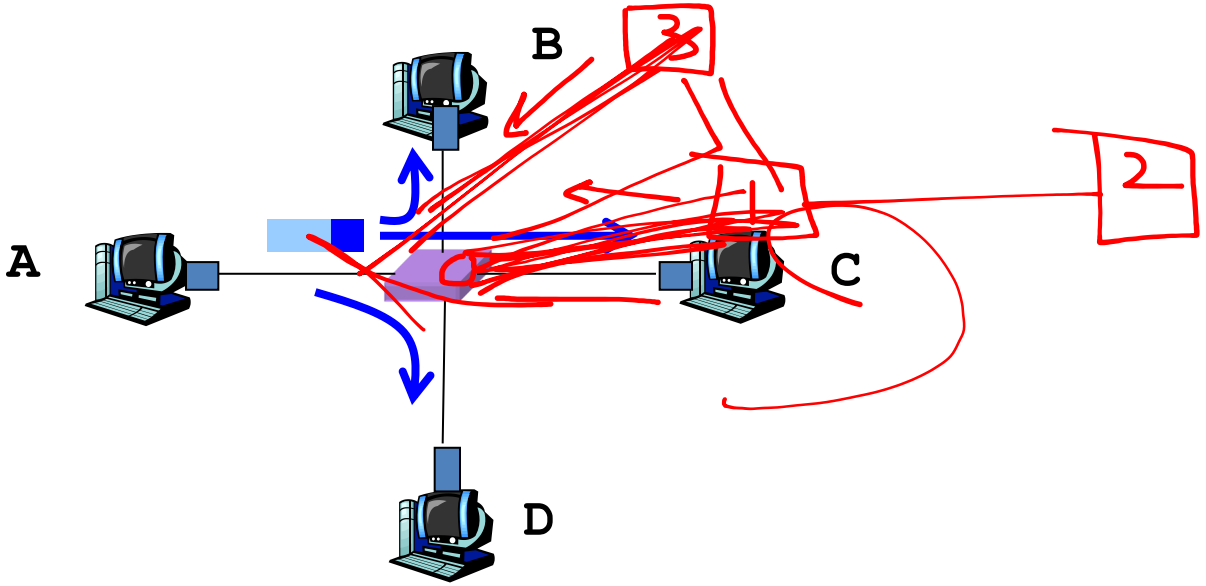

Link Layer: Switches

- Typically connects individual computers
  - A switch is essentially the same as a bridge
  - … though typically used to connect hosts
- Supports concurrent communication
  - Host A can talk to C, while B talks to D

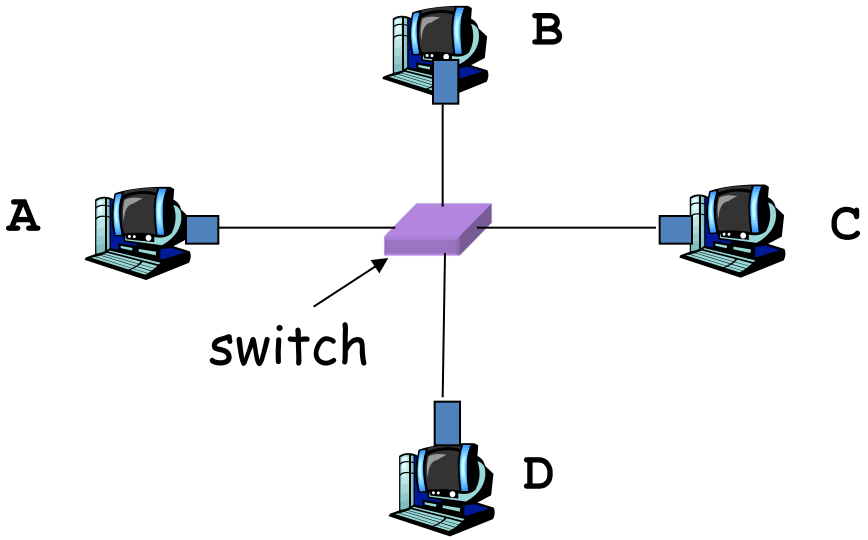

L3 Routers: Looking closer…

Basic Router Architecture

- Each switch/router has a forwarding table
  - Maps destination address to outgoing interface

- Basic operation

1. Receive packet
2. Look at header to determine destination address
3. Look in forwarding table to determine output interface
4. Modify packet header (e.g., decr TTL, update chksum)
5. Send packet to output interface

Basic Router Architecture

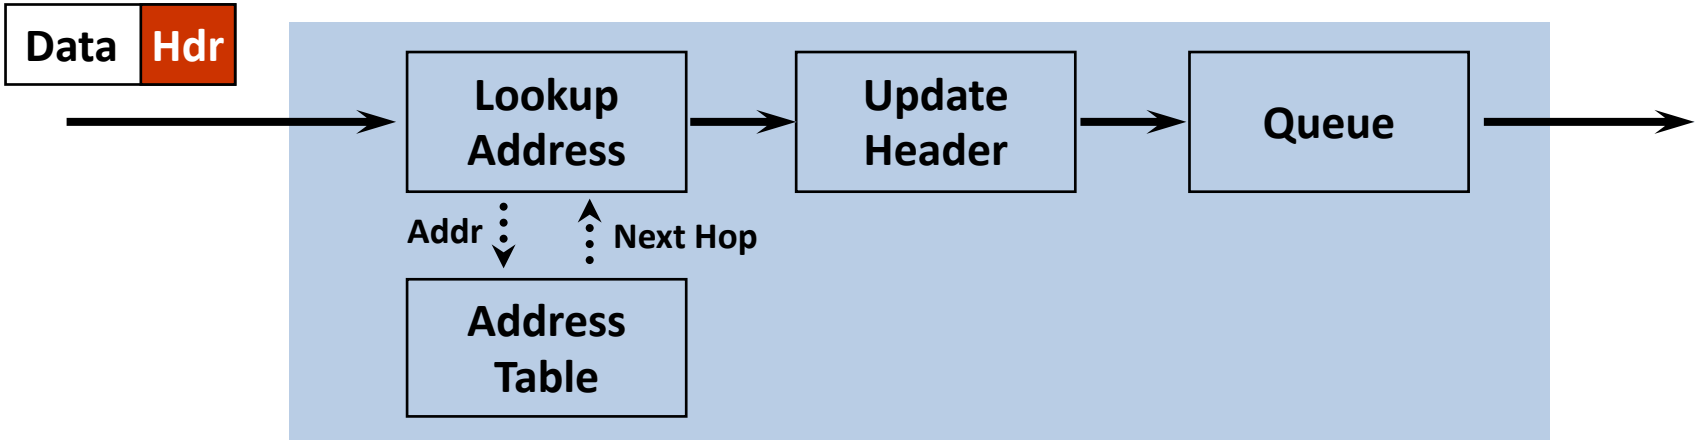

- Basic operation
  1. Receive packet
  2. Look at header to determine destination address
  3. Look in forwarding table to determine output interface
  4. Modify packet header (e.g., decr TTL, update chksum)
  5. Send packet to output interface

Router

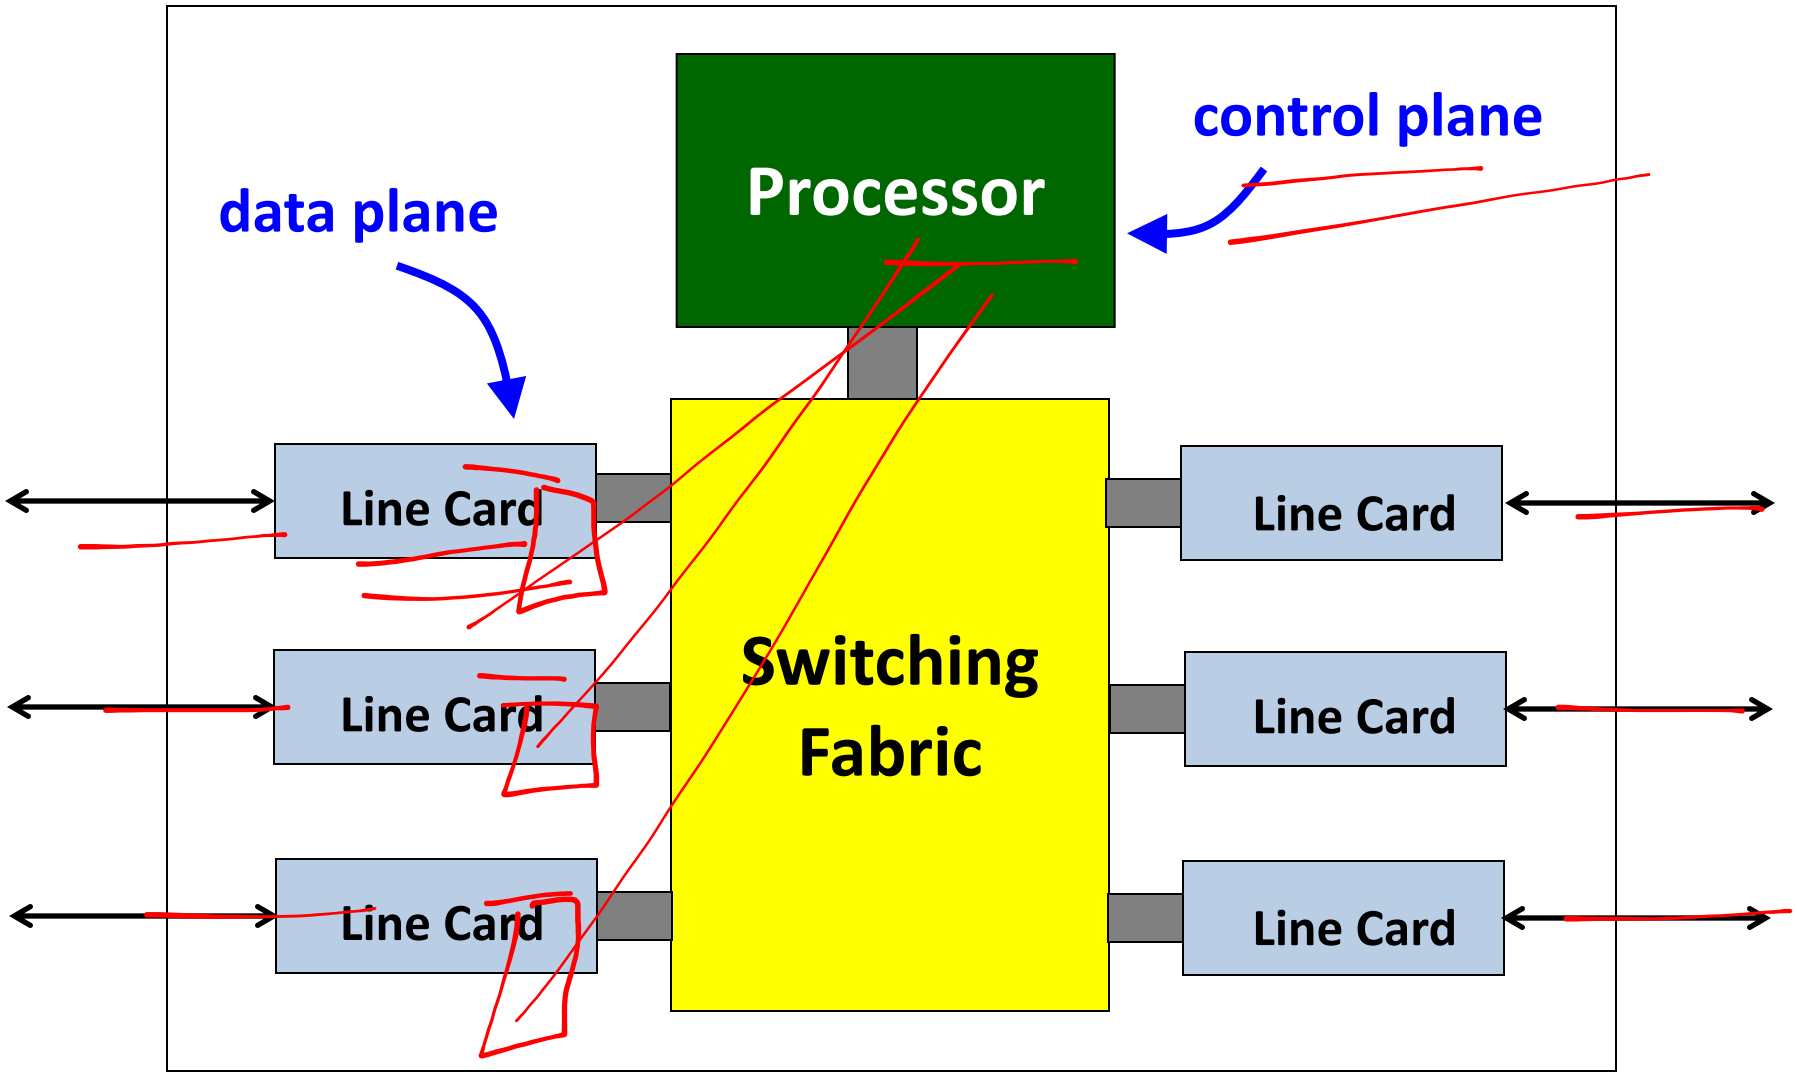
Processor
control plane
data plane
Line Card
Line Card
Line Card
Switching Fabric
Line Card
Line Card
Line Card

LPM: Motivation

- Each packet has destination IP address
- Router looks up table entry that matches address

- Benefits of CIDR allocation and LPM
  - **Efficiency:** Prefixes can be allocated at much finer granularity
  - **Hierarchical aggregation:** Upstream ISP can aggregate 2 contiguous prefixes from downstream ISPs to shorter prefix

Decision: Crossbar switch

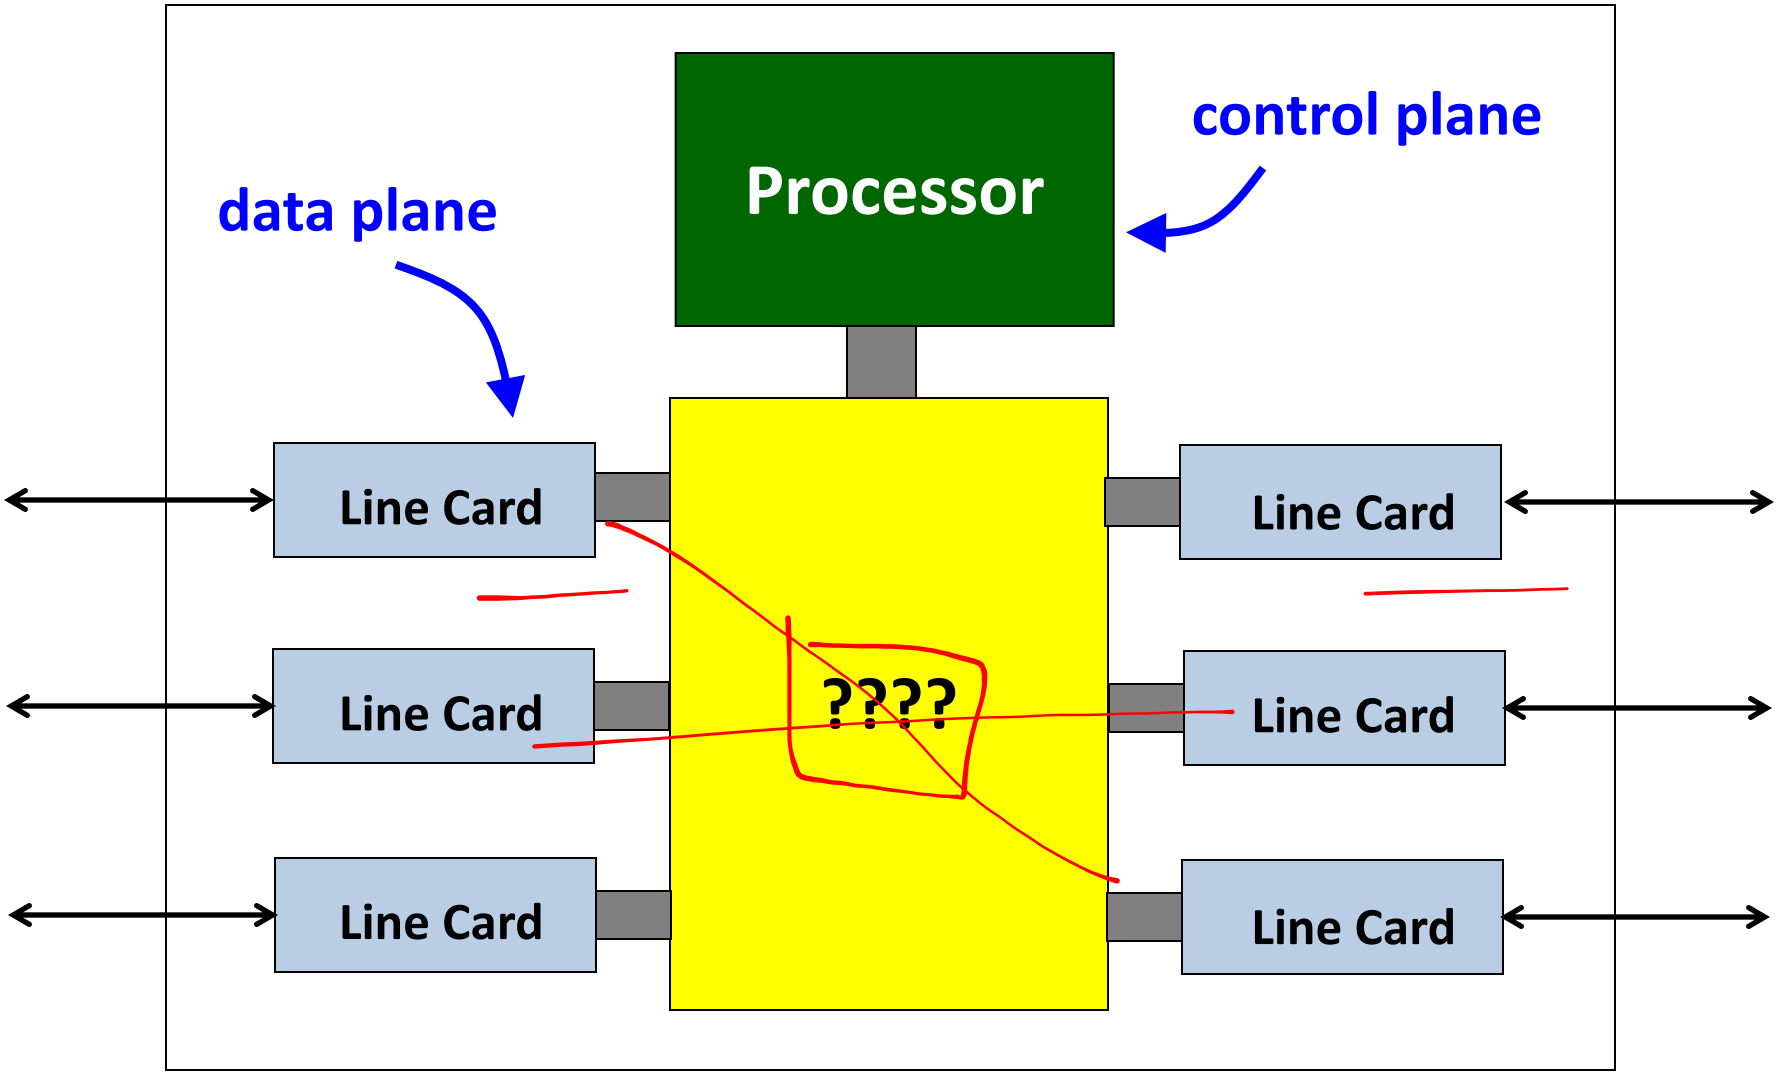

Decision: Crossbar switch

- Shared bus
  - Only one input can speak to one output at a time
- Crossbar switch / switched backplane
  - Input / output pairs that don't compete can send in same timeslot

Crossbar switching

- Every input port has connection to every output port
- In each timeslot, each input connected to zero or more outputs

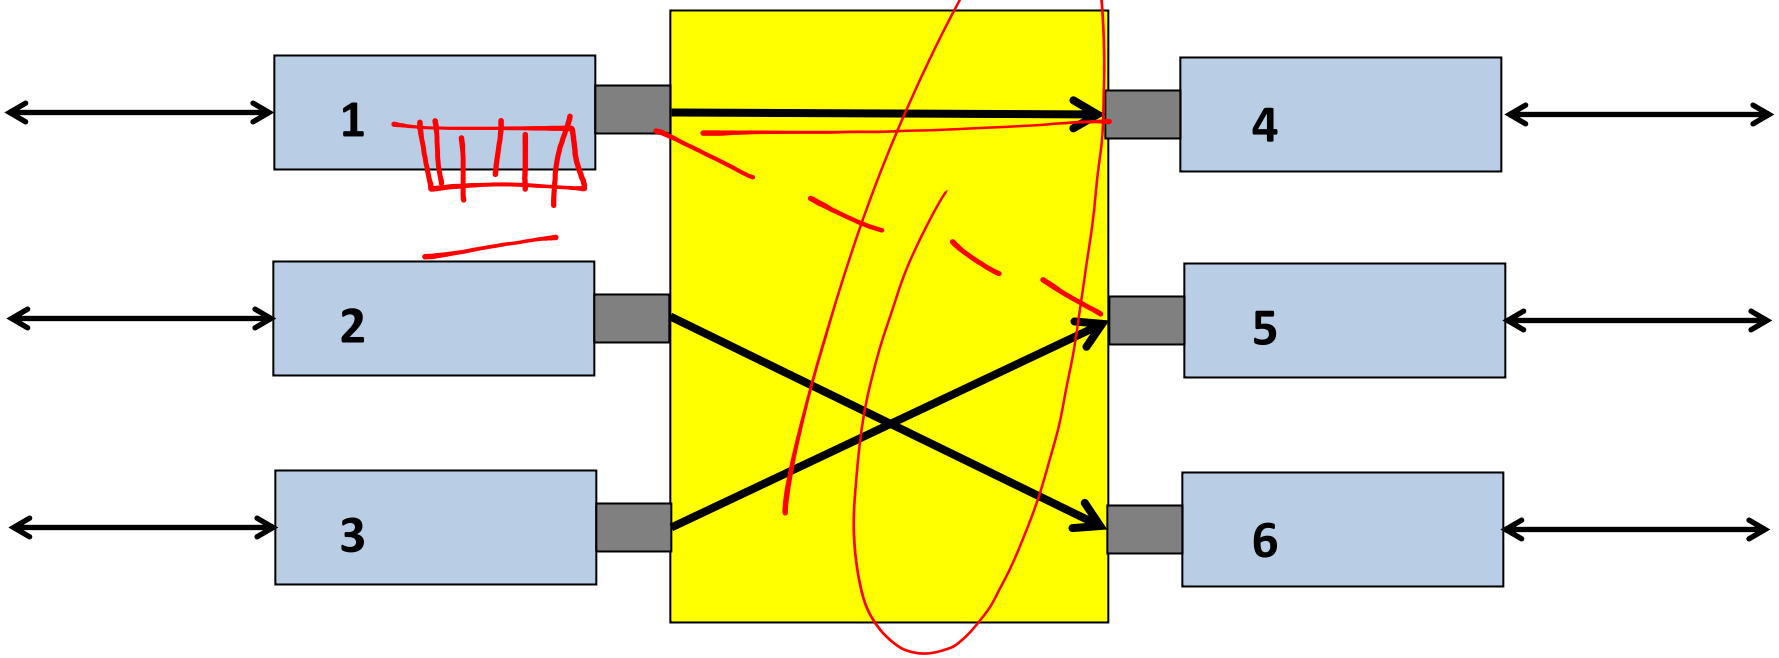

Crossbar switching

- Every input port has connection to every output port
- In each timeslot, each input connected to zero or more outputs

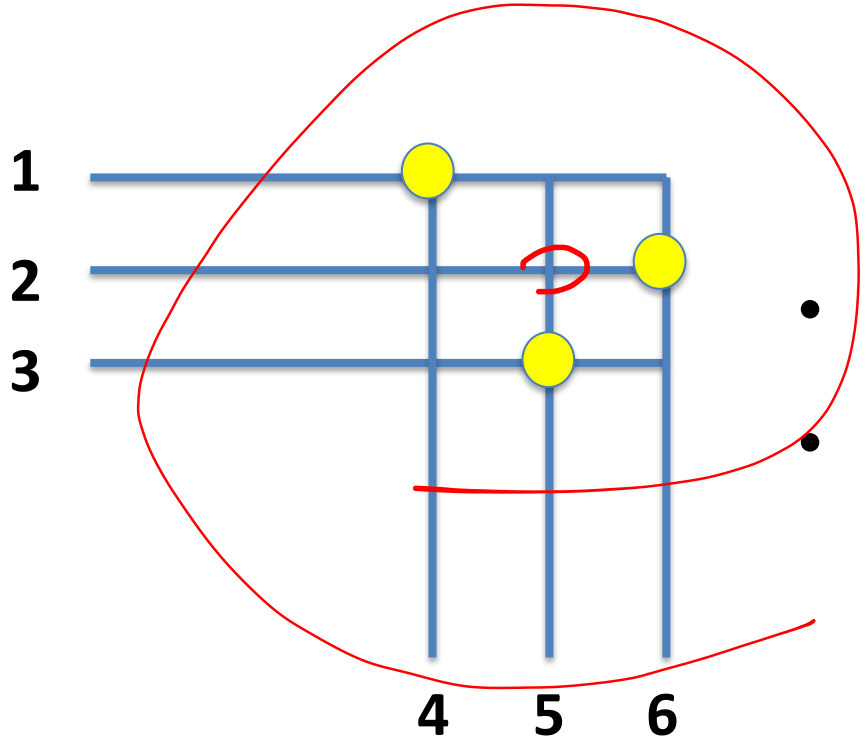

- Good parallelism
- Needs scheduling

Problem: Head-of-line blocking

- Packet at front of queue **blocks** packets behind it from being processed
  - *e.g.:* 1st packet at Input 1 wants to go to Output 5; next packet at Input 1 that wants Output 4 is still blocked

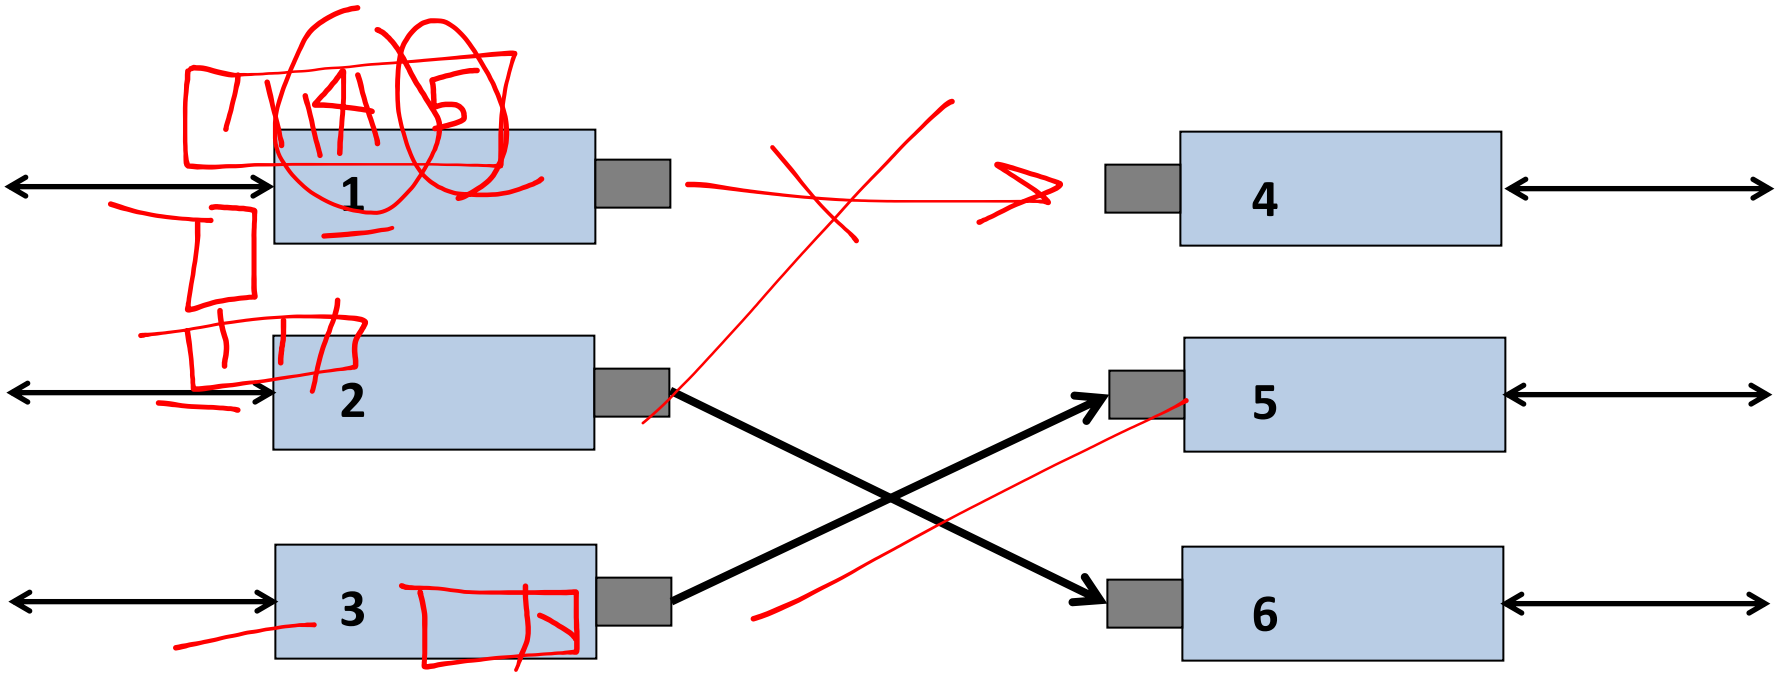

Solution: *Virtual output queues*

- One queue at input, **per output port** (for all inputs)
- So **avoids head-of-line blocking** during crossbar scheduling

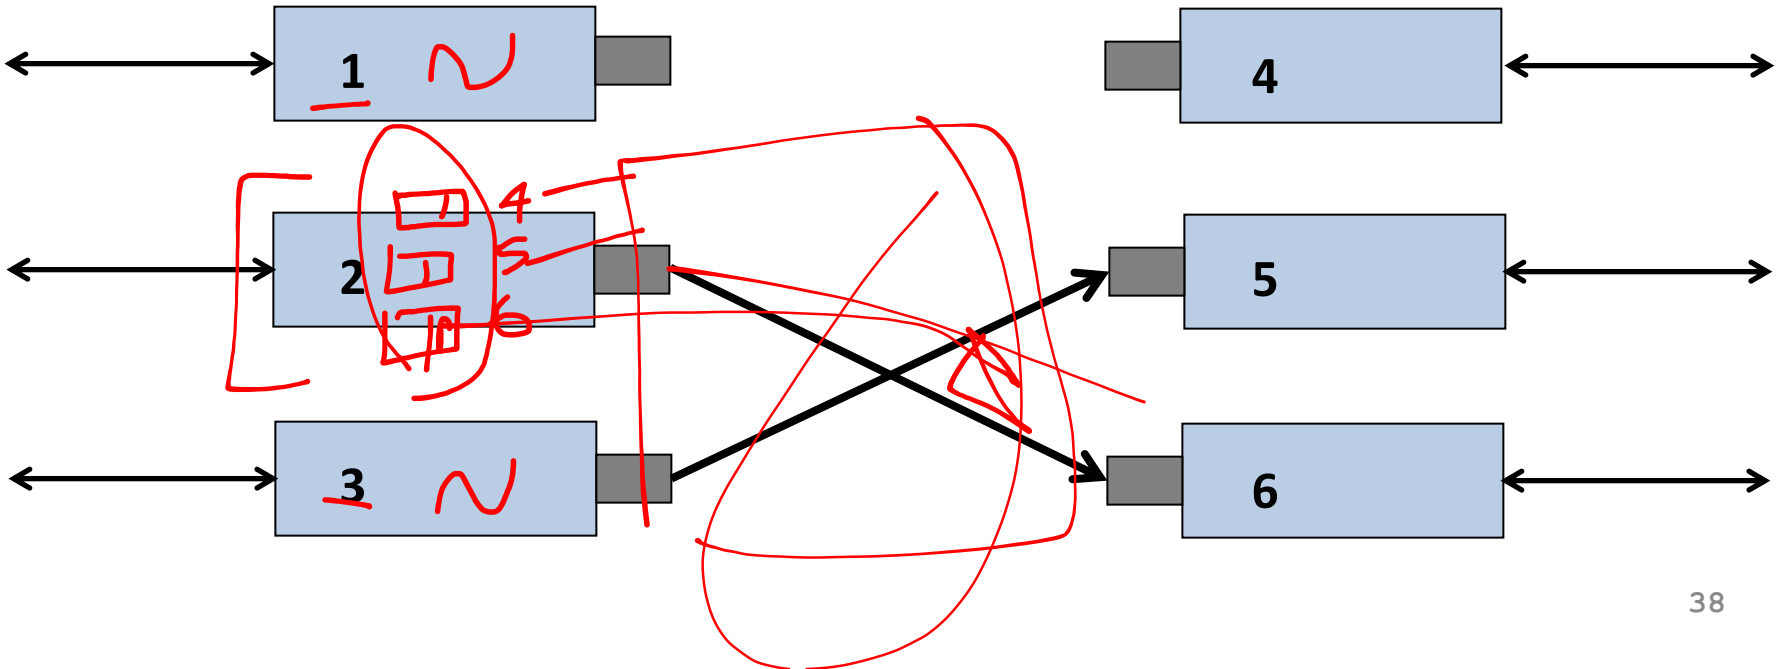

Data, Control, & Management Planes

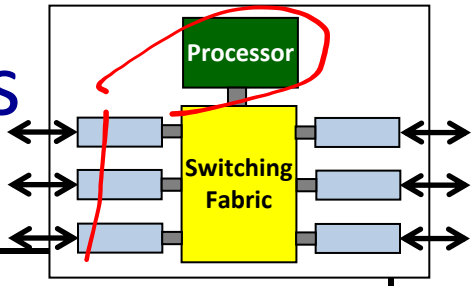

	Data	Control	Management
Time-scale	Packet (ns)	Event (10 ms to sec)	Human (min to hours)
Tasks	Forwarding, buffering, filtering, scheduling	Routing, signaling	Analysis, configuration
Location	Line-card hardware	Router software	Humans or scripts

Coming Up in 461

Next Class Meeting

Lectures 5 (Transport Layer) and 6 (Congestion Control)

Precepts this Thursday and Friday