CS 725/825 & IT 725

Lecture 4

Networking Fundamentals

September 9, 2024

Anatomy of a router/switch

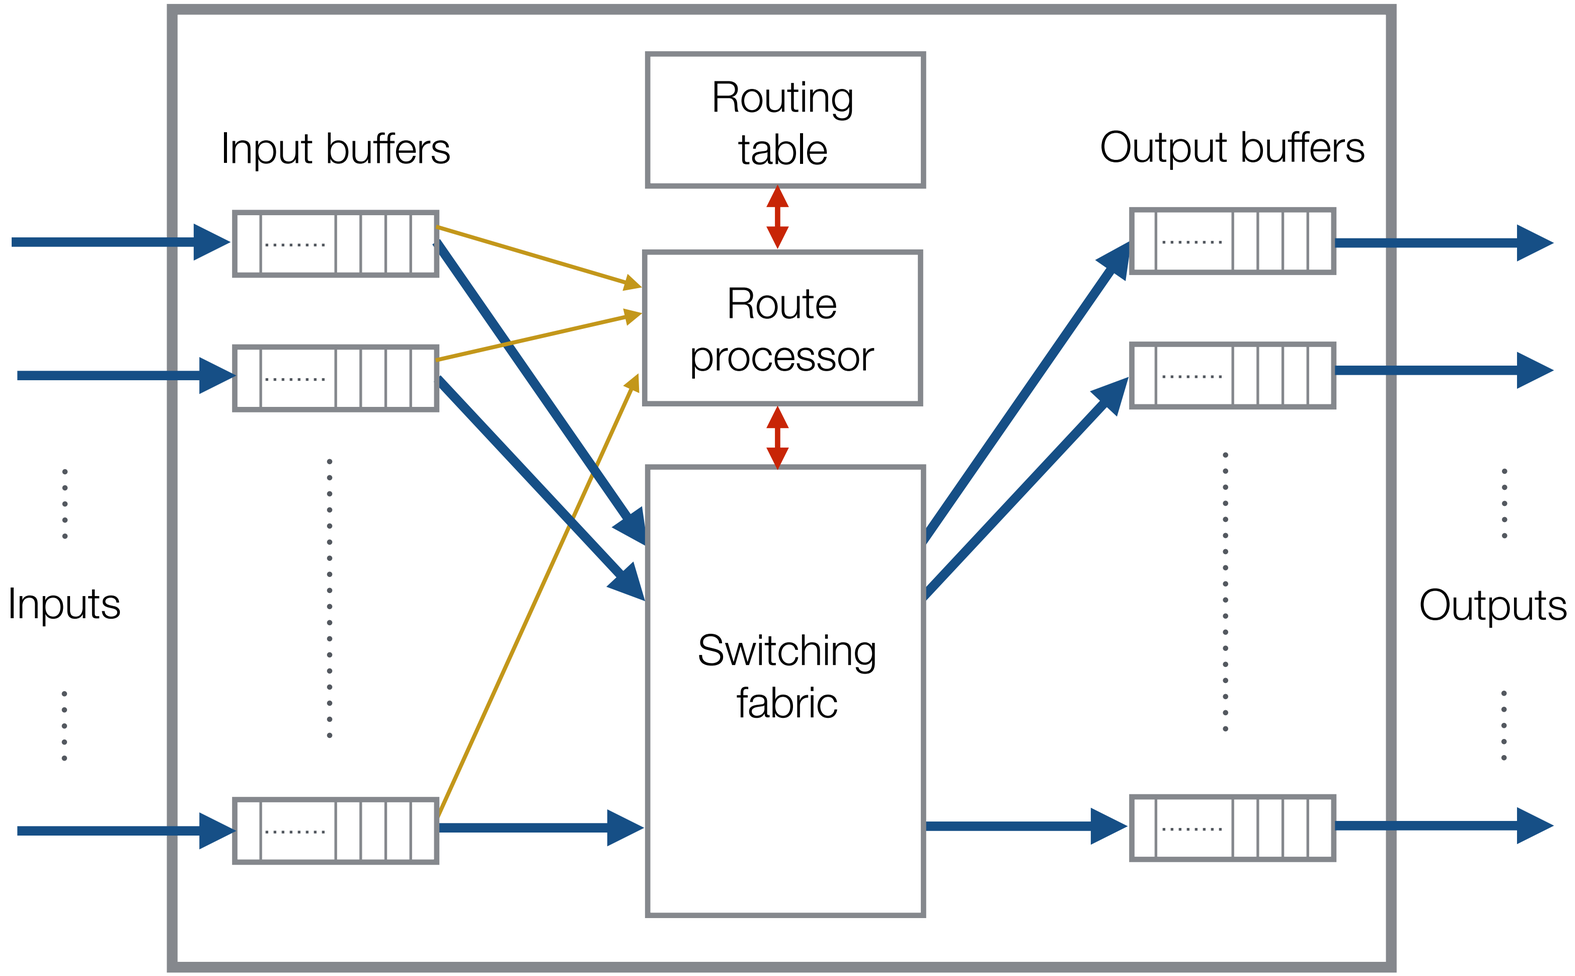

Store and Forward

- Intermediate nodes receive, store, and forward packets
  - storing and retransmission adds fixed delay
- Output conflict - multiple packets waiting to be forwarded on the same output
  - queues/buffers used to store packets waiting transmission
  - queuing adds variable delay
  - potential for packet loss due to buffer overflow

A bit of history...

- Local area networks (late 80's, early 90's)
  - (then) based on broadcast & select medium

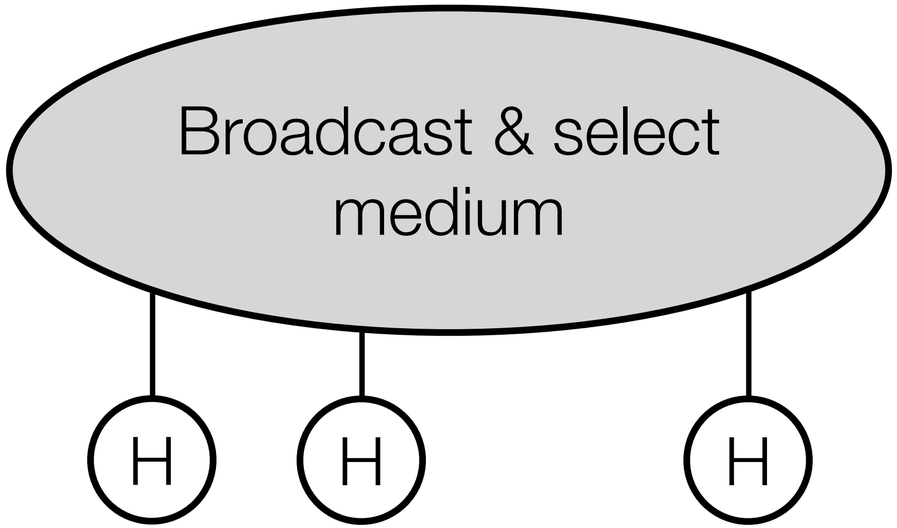

  - (today) a network of Ethernet (L2) switches

Comparison

- Routed networks

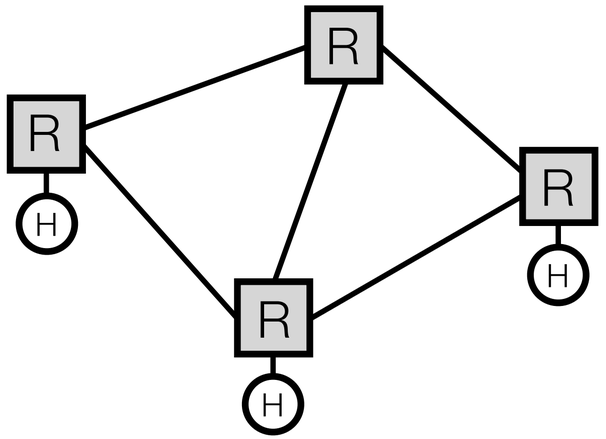

- topology driven by geography
- long distances (high latency)
- need for scalability
- location-related addresses
- routing

➡ **Network Layer (L3)**

- Broadcast & select

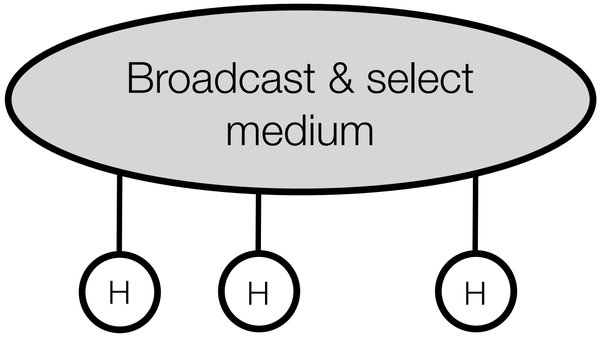

- everyone connected to everyone
- short distances (low latency)
- lesser need for scalability
- arbitrary addresses
- address discovery

➡ **Link Layer (L2)**

Today

Internet - a network to INTERconnect NETworks

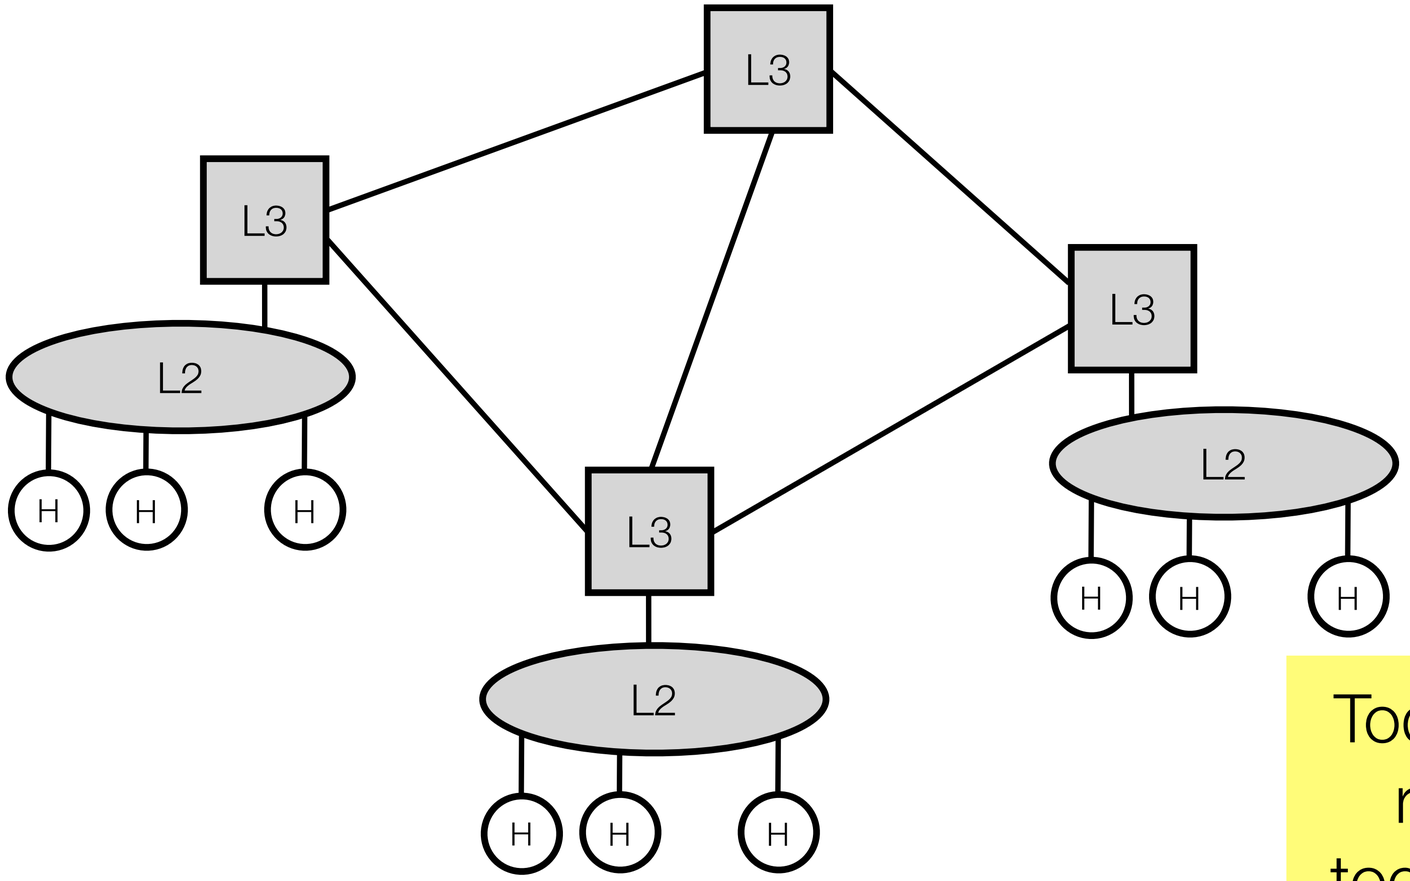

Today, L2 switching and L3 routing are often mixed together in devices capable of processing multiple layers

Networking Basics

- Goals: IP/MAC addresses, fundamental protocols, subnetting, basics of routing/switching
- Assumptions:
  - packet switched network
  - nodes attached to a L2 broadcast-and-select network
  - each node “has” a 6-byte MAC and a 4-byte IP addresses

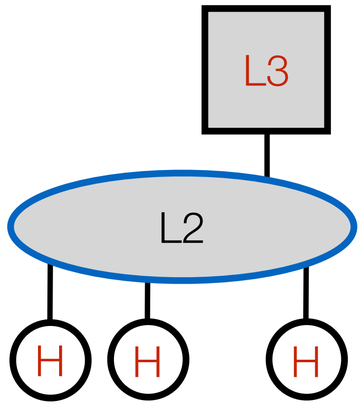

Address Resolution

- Problem: Find MAC address of a node with a given IP address

A has a packet with
IP destination address B,
A needs B's MAC address
to deliver the packet

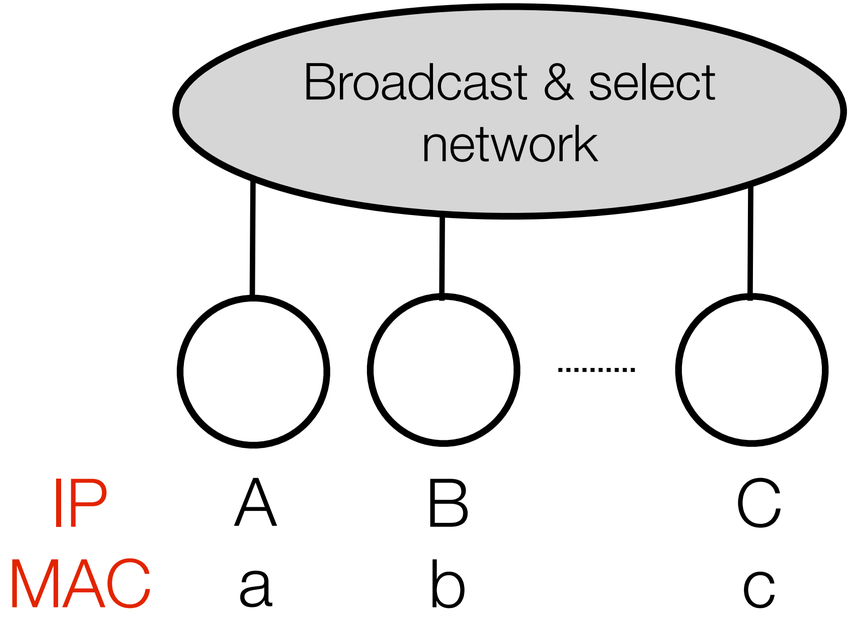

- Solution: ARP - Address Resolution Protocol

ARP

A has a packet with IP destination address B, A needs B's MAC address to deliver the packet

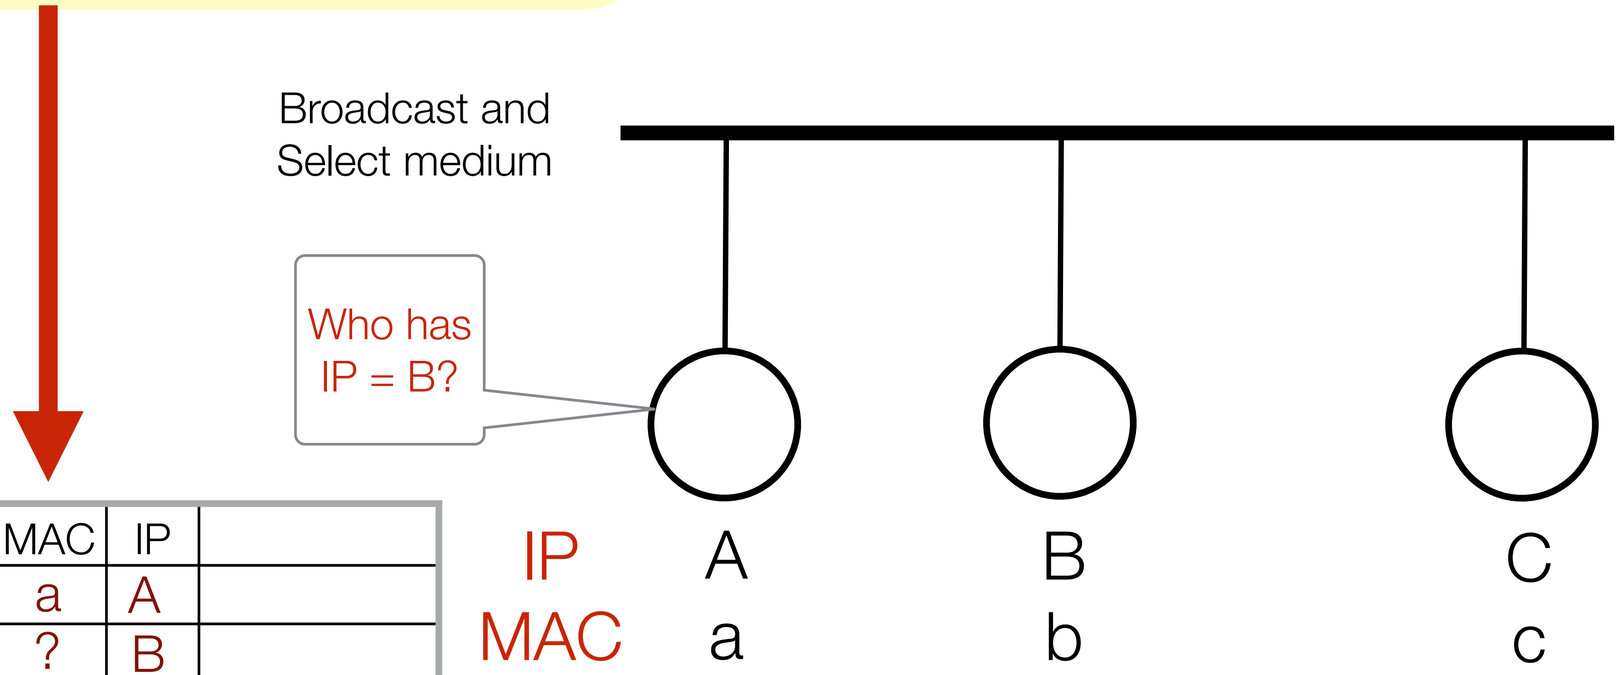

	MAC	IP	
SRC	a	A	
DST	?	B	

ARP

- Step 1: “Who has …” broadcasted to everyone by A

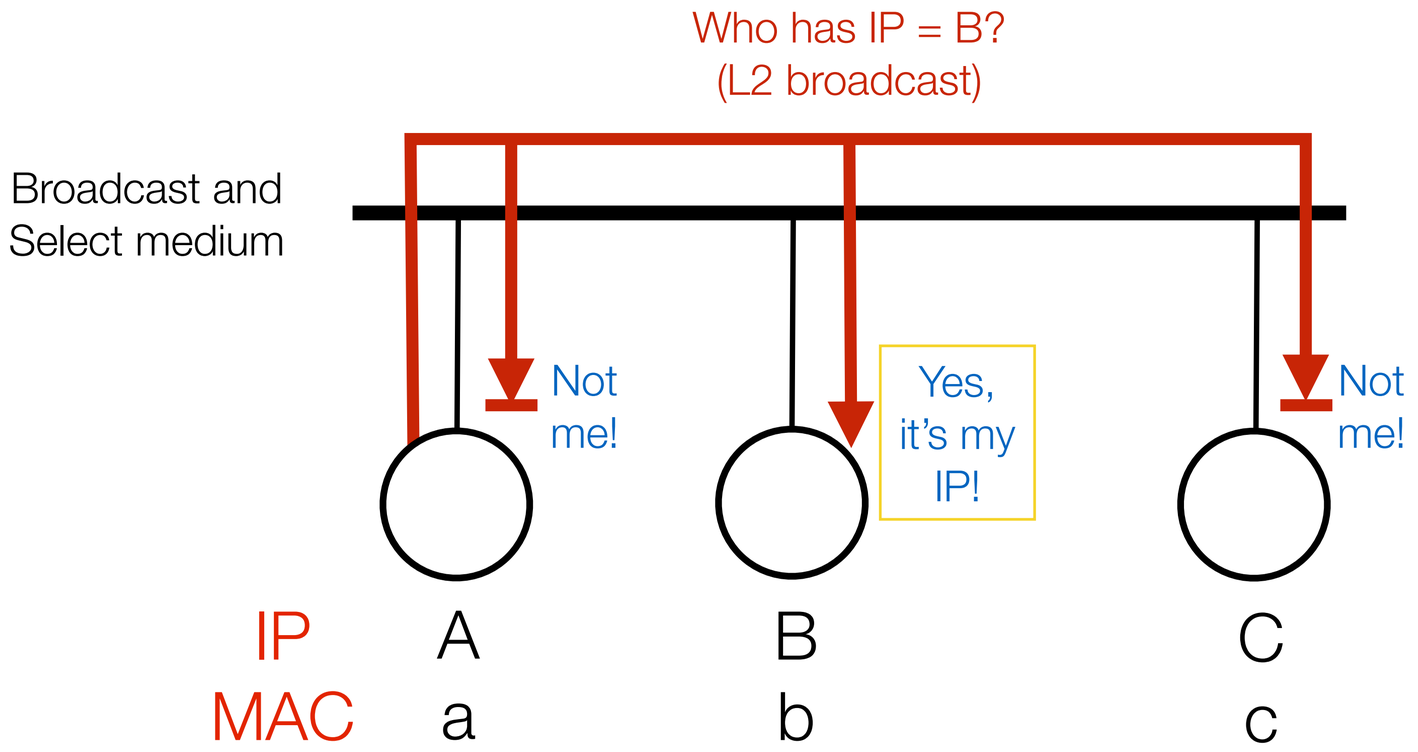

ARP packet processing

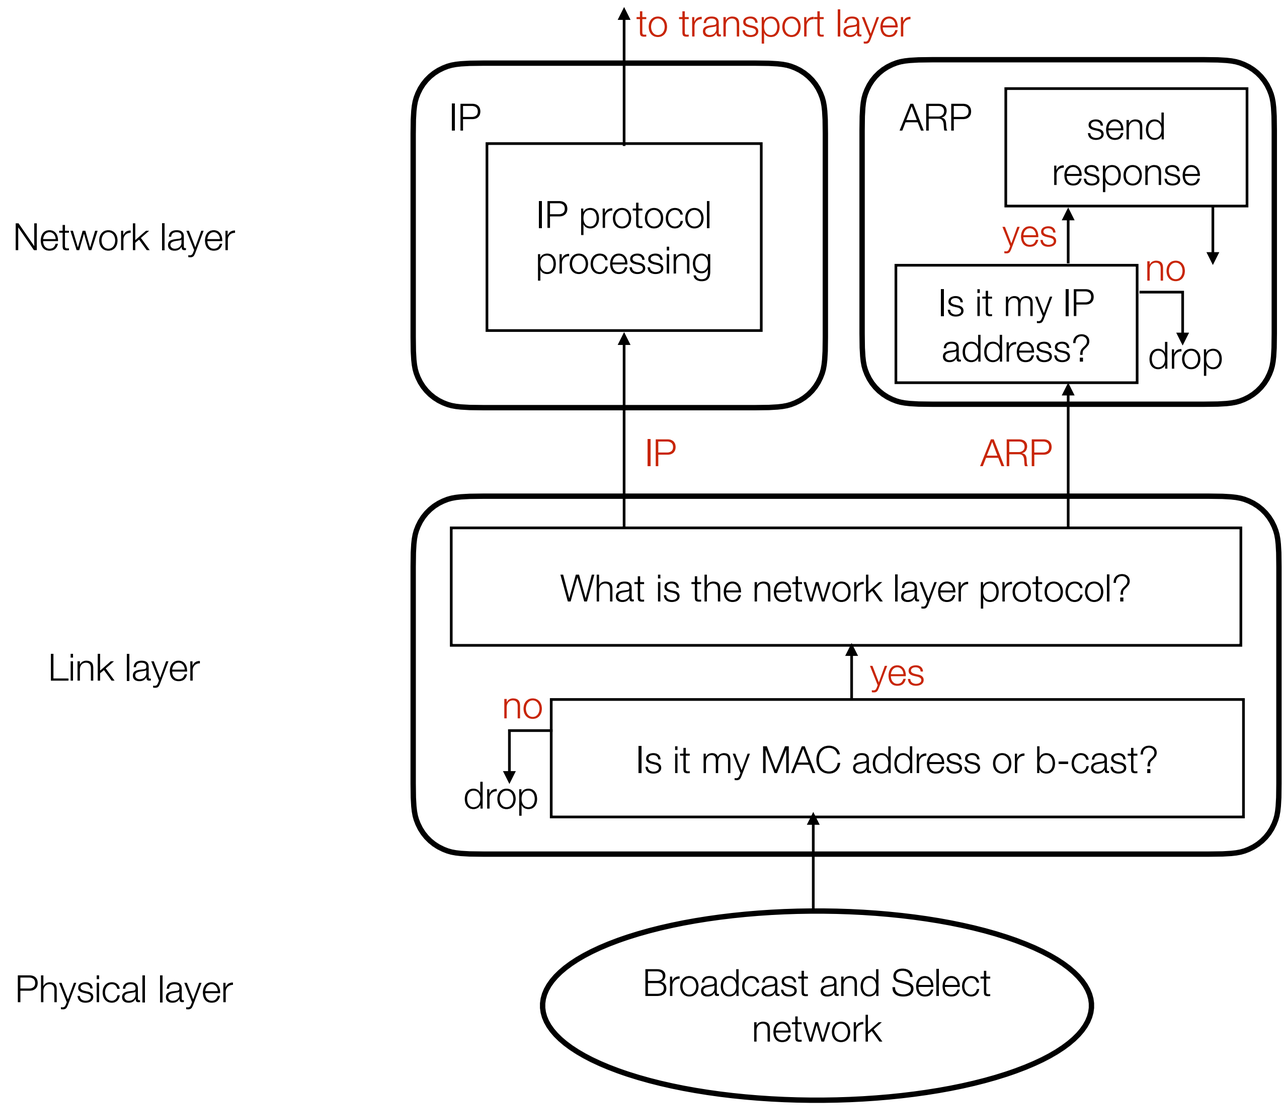

ARP packet processing

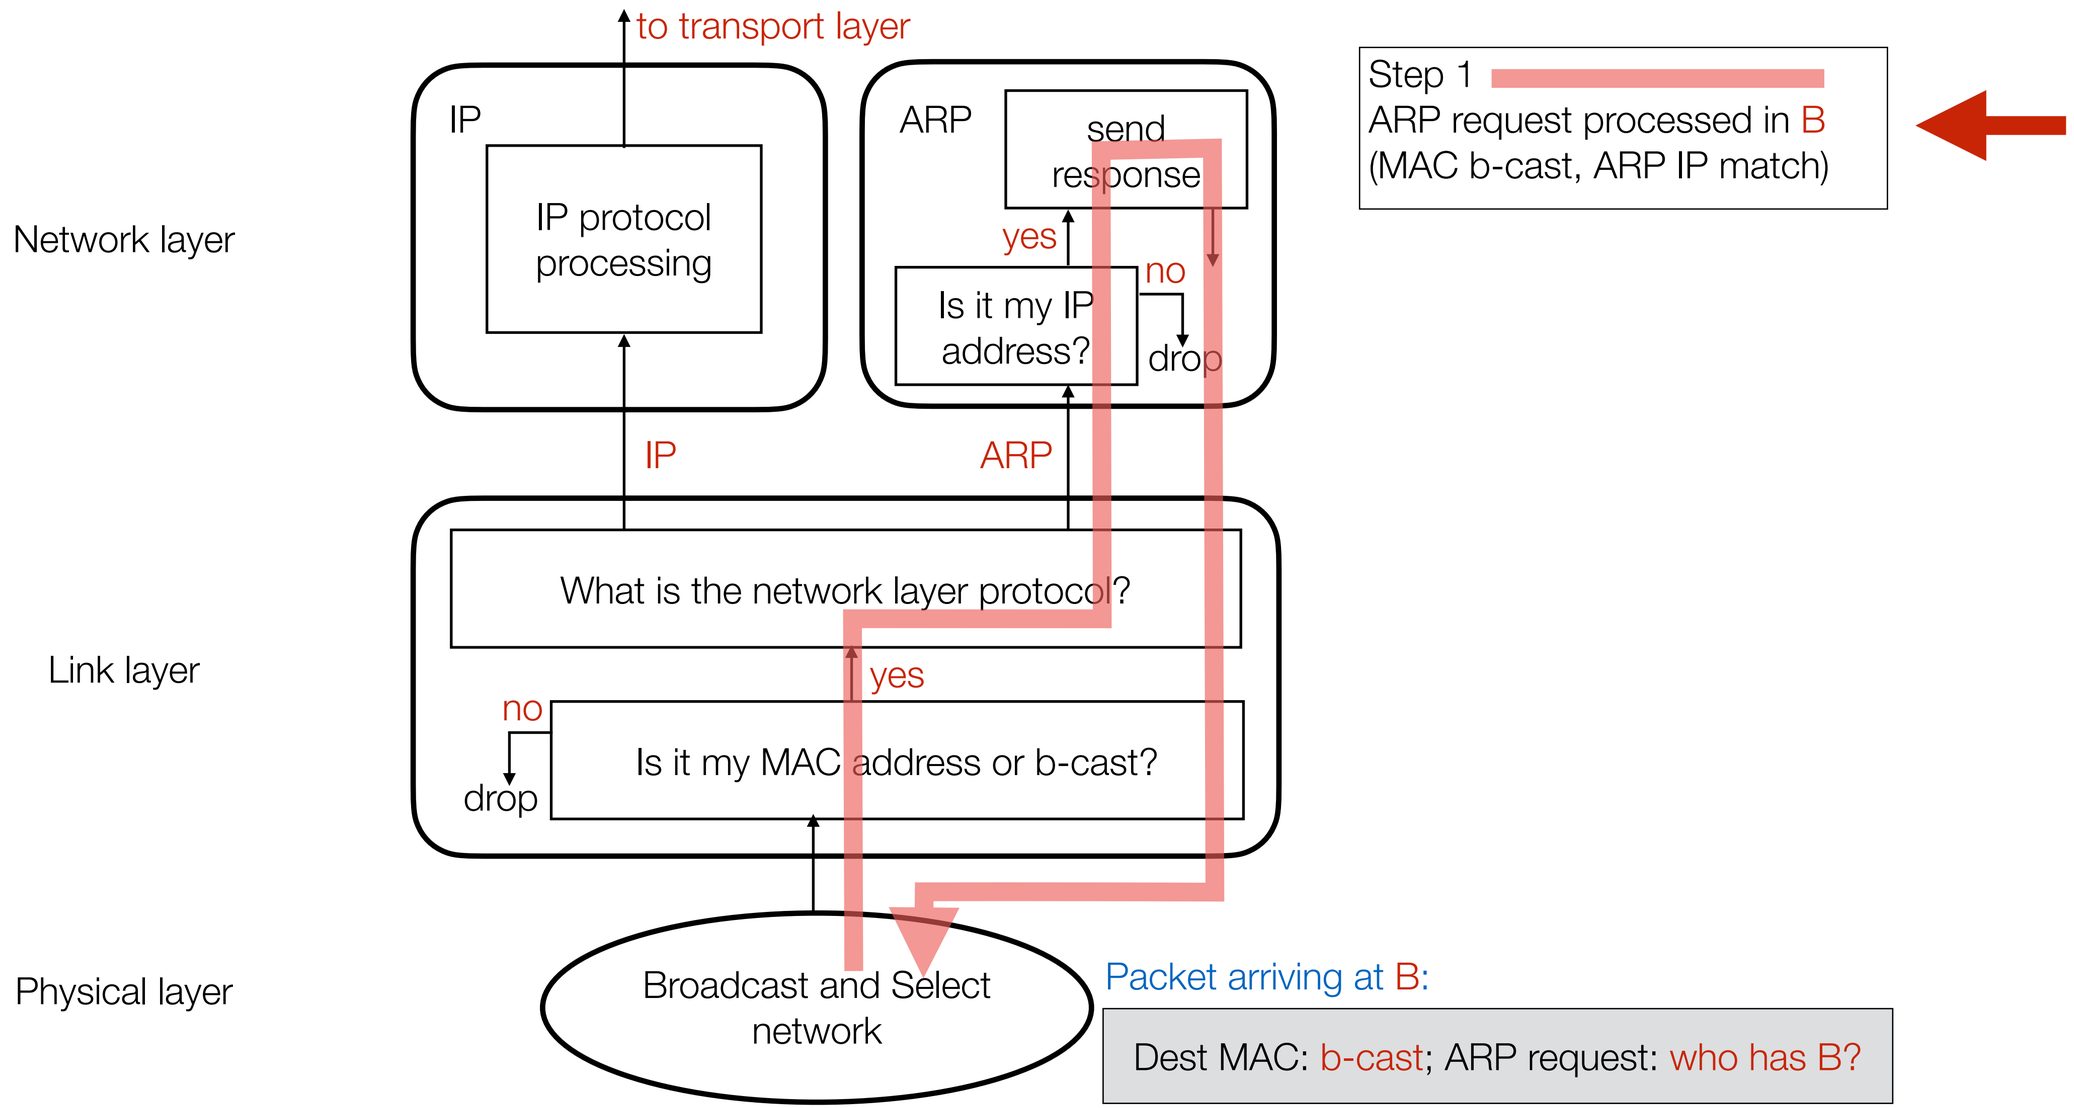

ARP packet processing

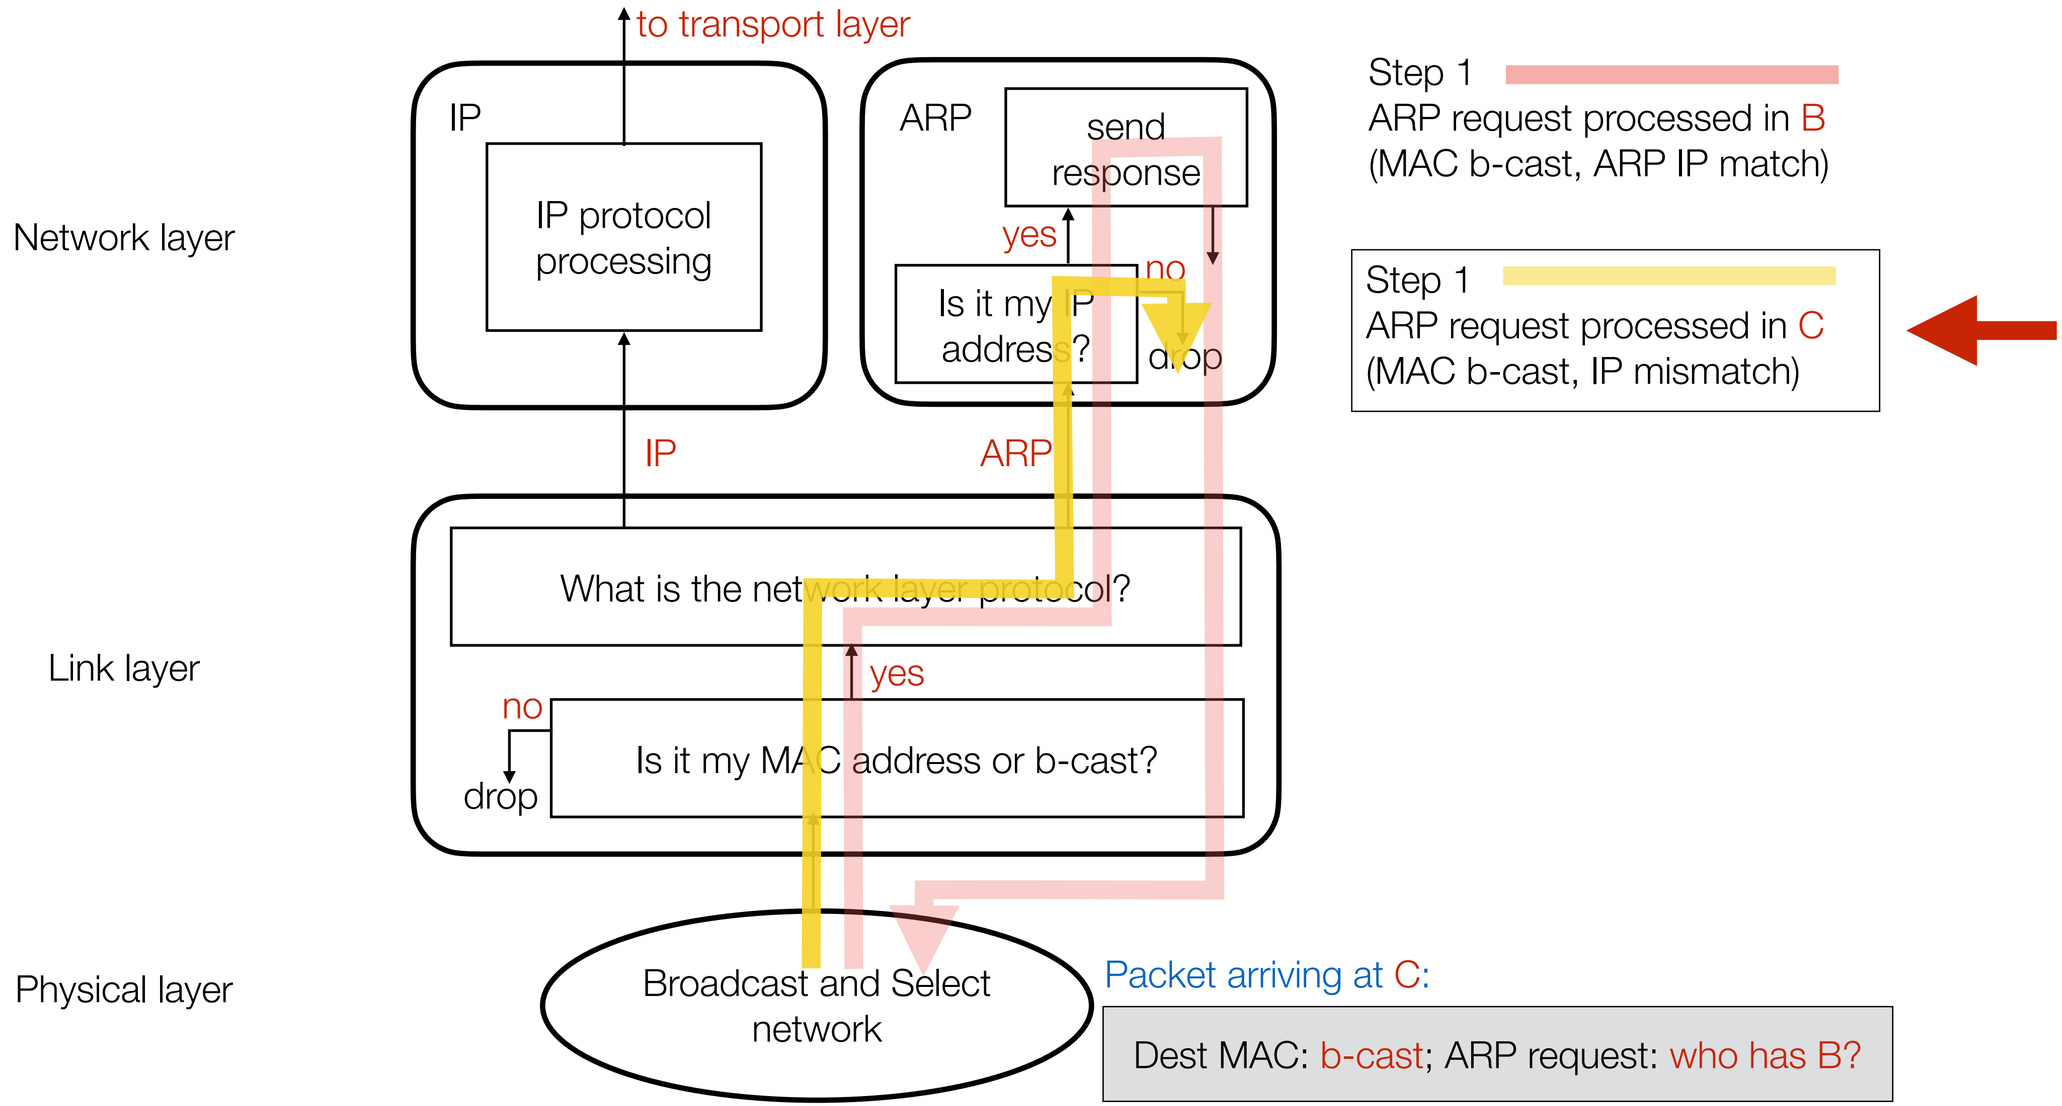

ARP

- Step 2: B sends L2 unicast response to A (MAC a)

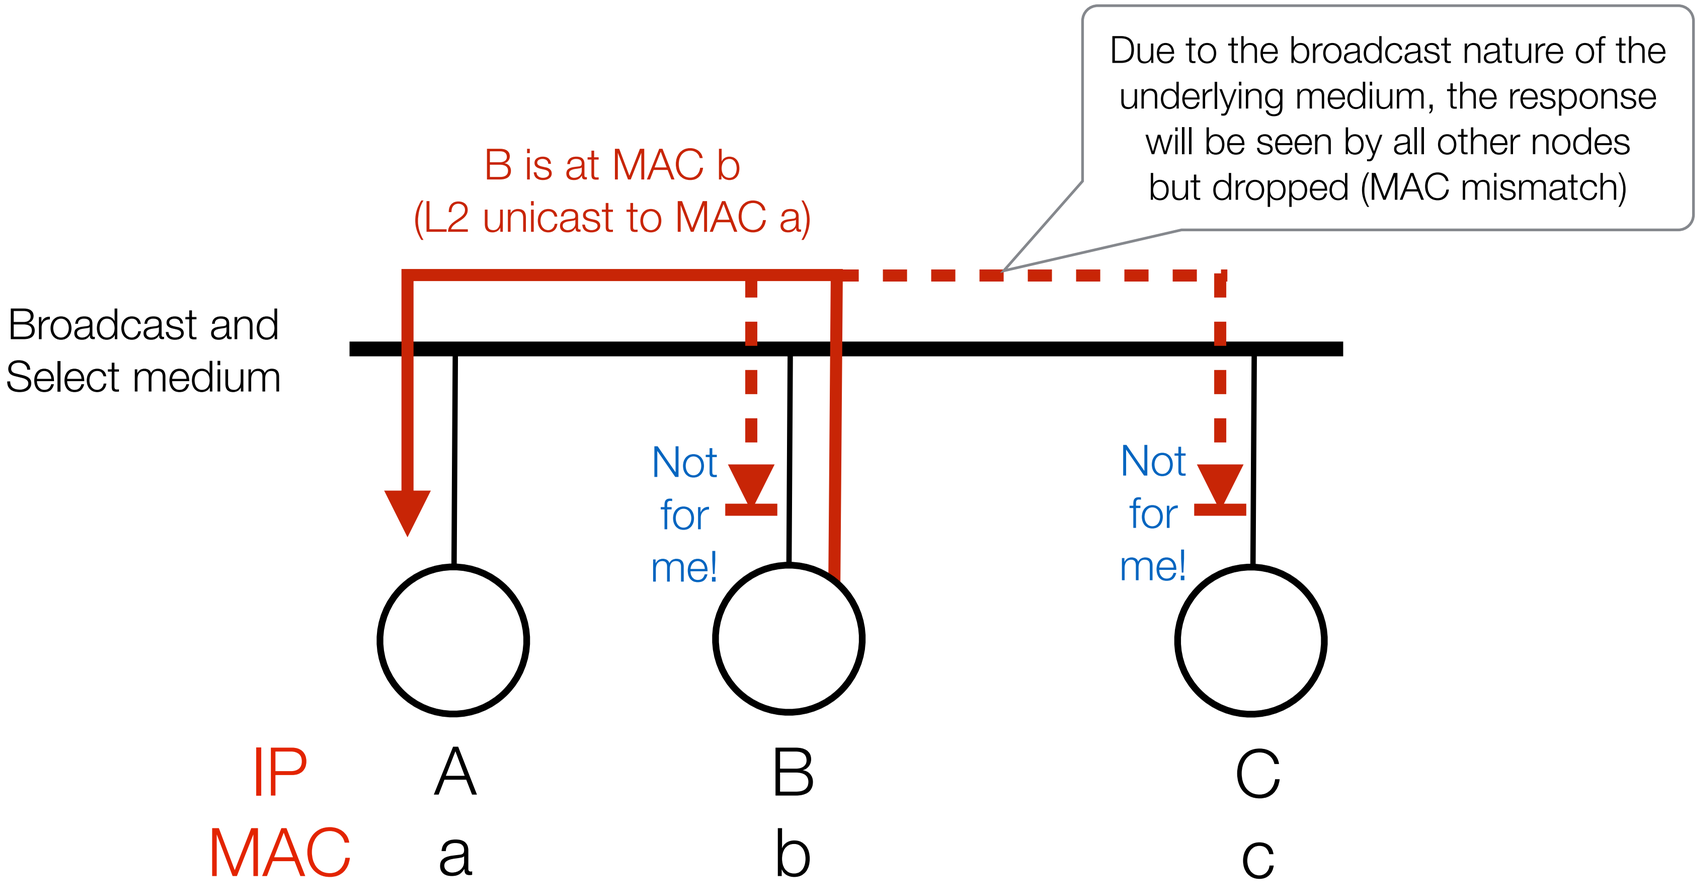

ARP packet processing

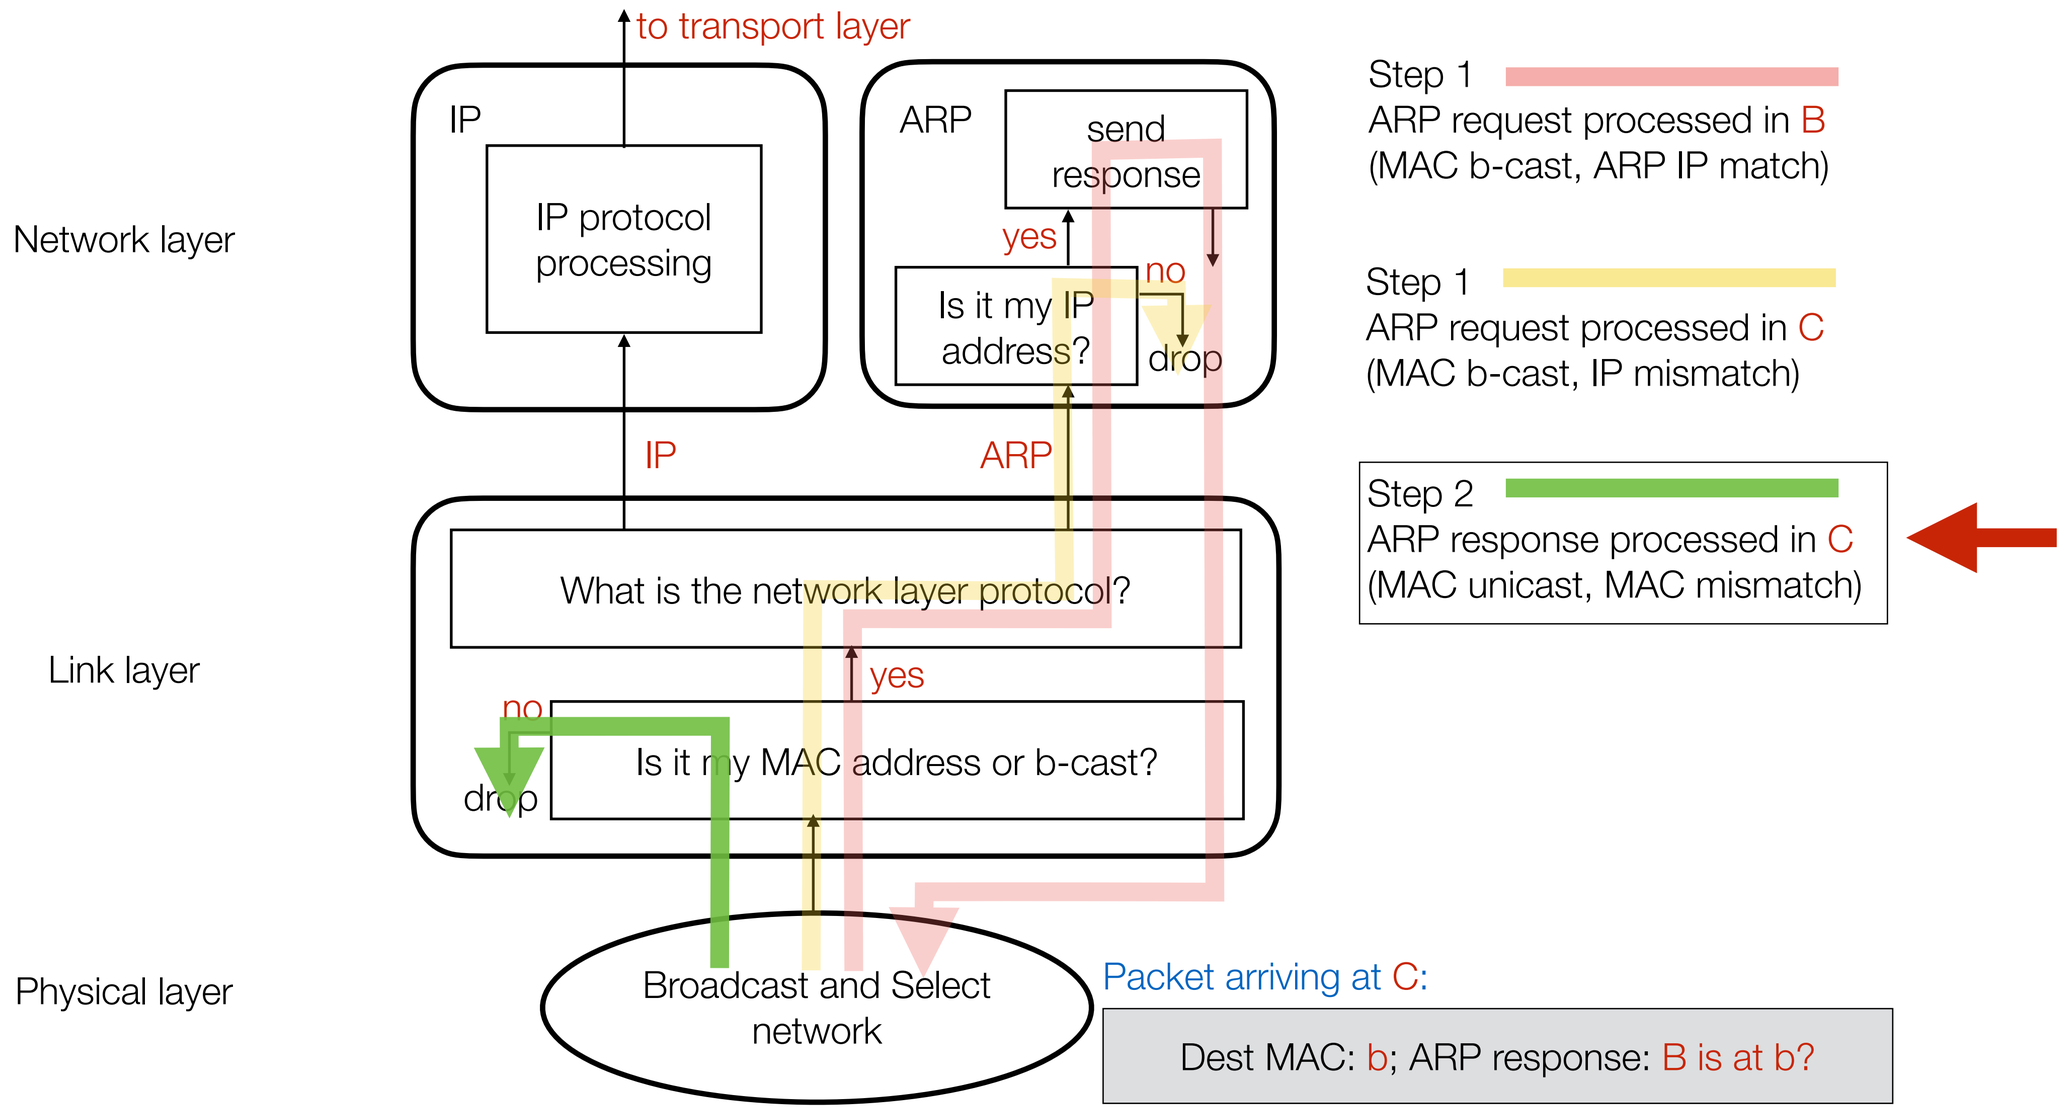

ARP

- Step 3: A sends data using L2 unicast to MAC b

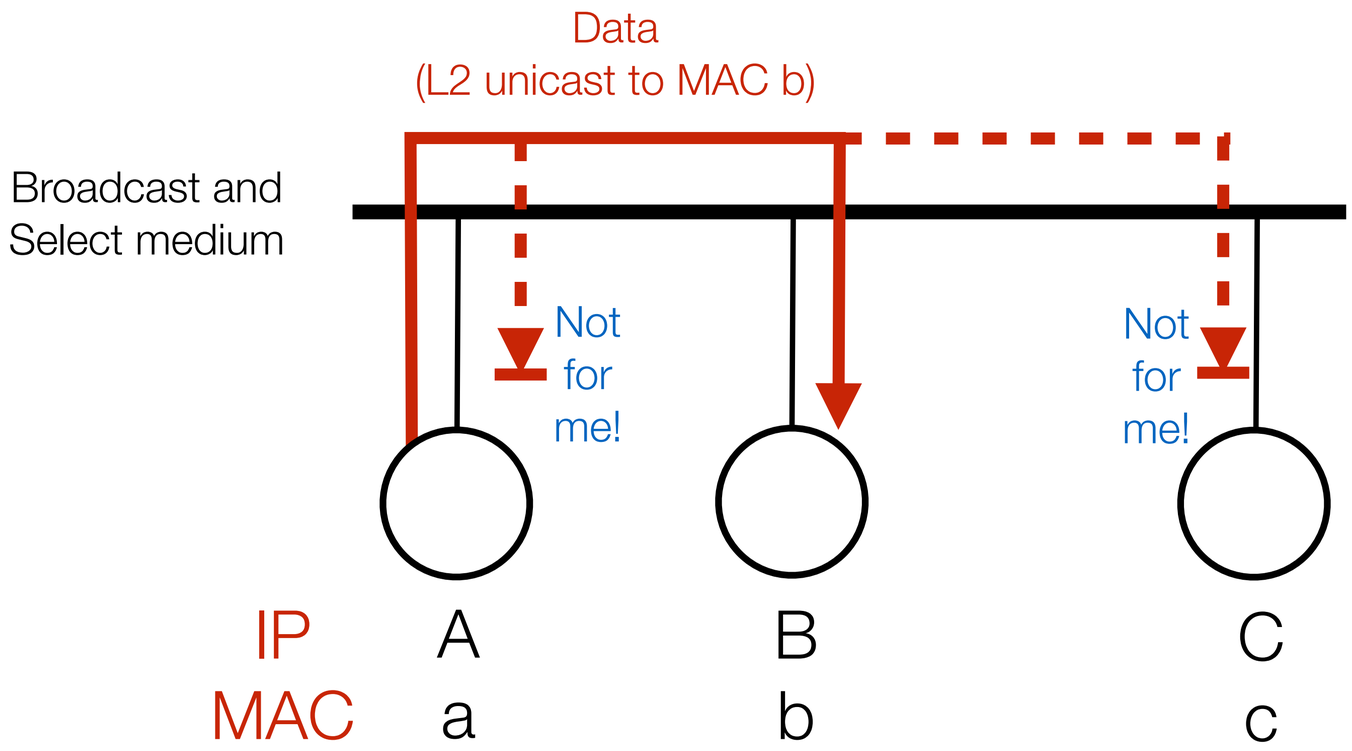

ARP packet processing

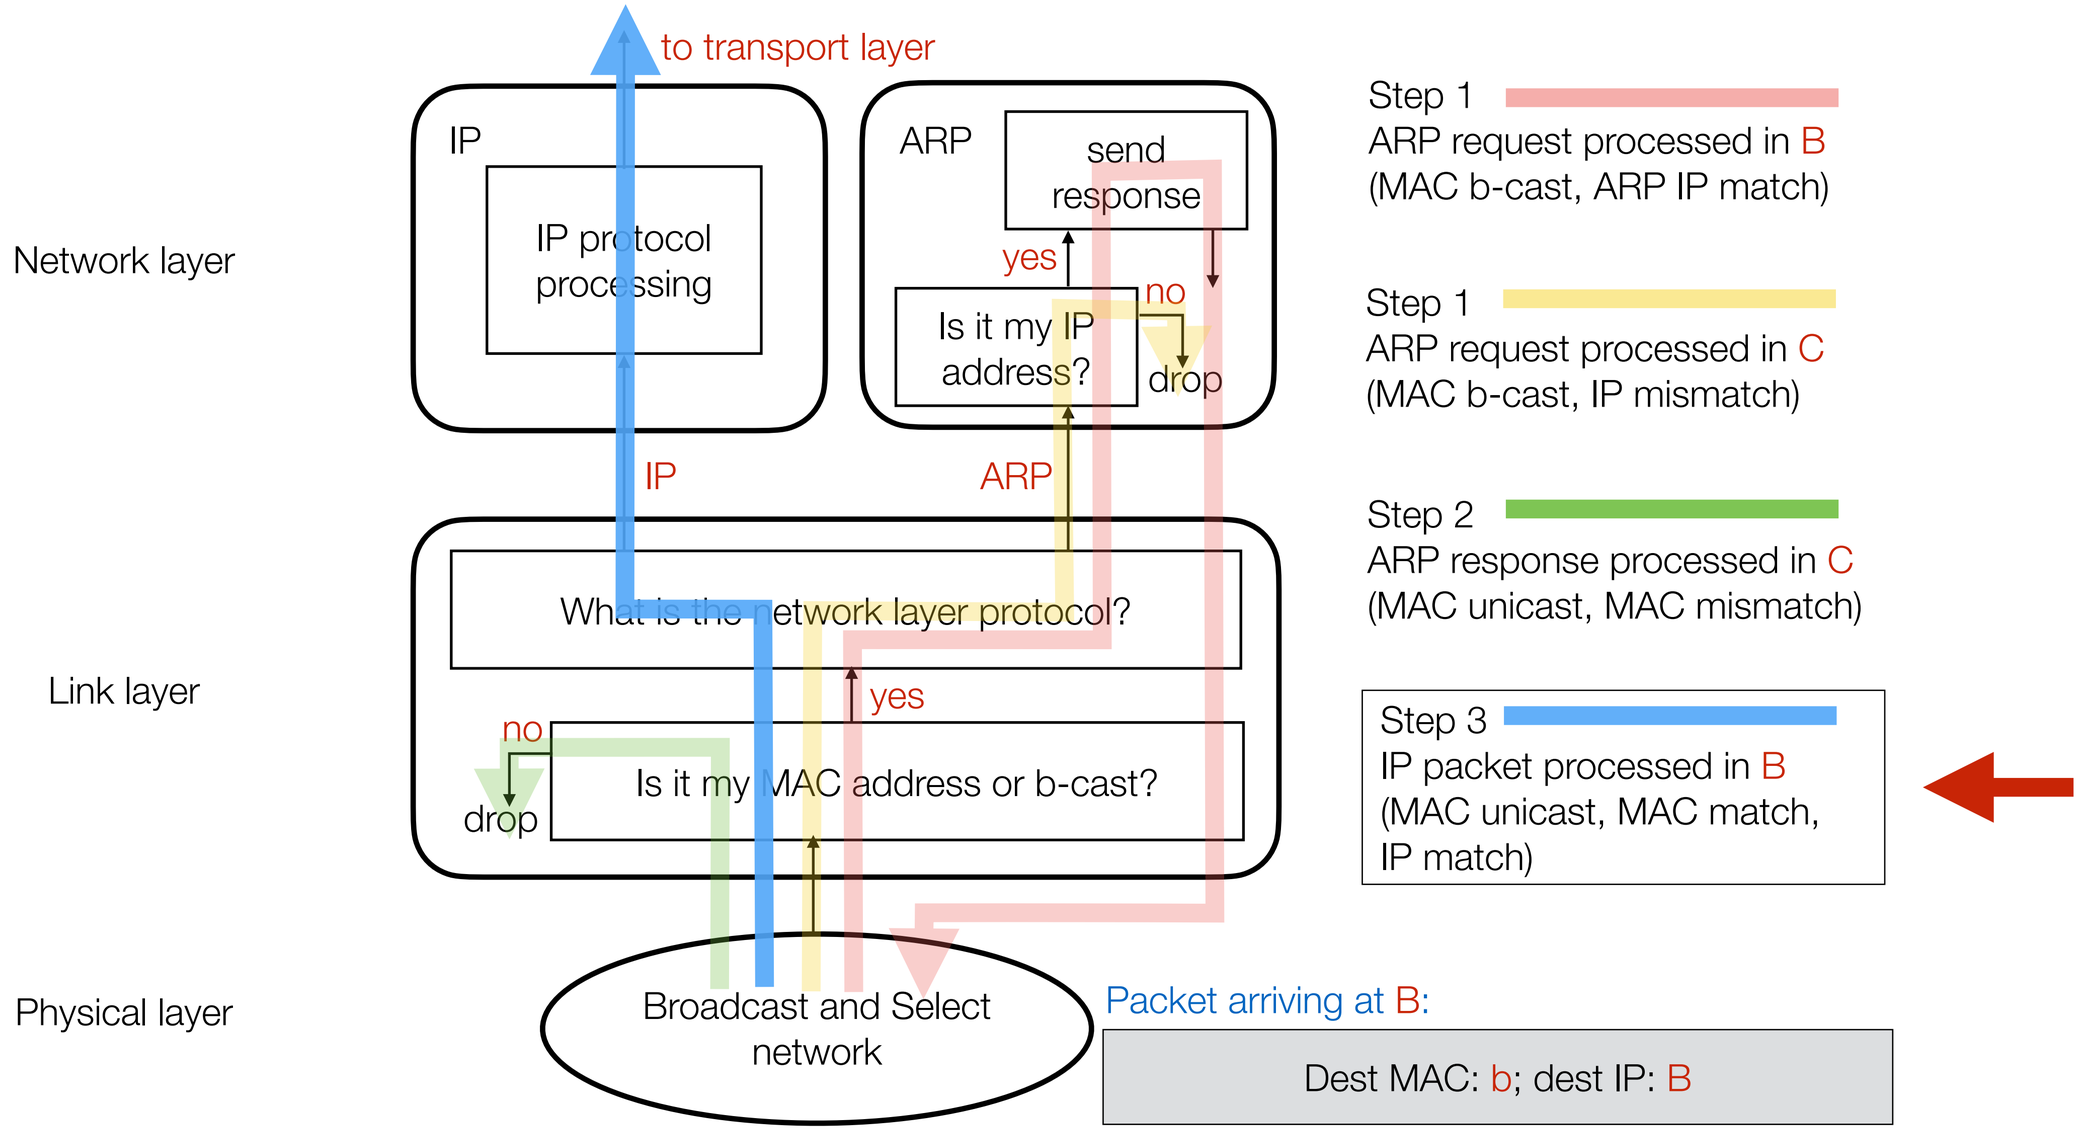

Subnetting

- Problem: deciding whether an IP address belongs to a specific subset of IP addresses
- Solution: nodes on a subnet (and only those) have IP addresses within a specific range

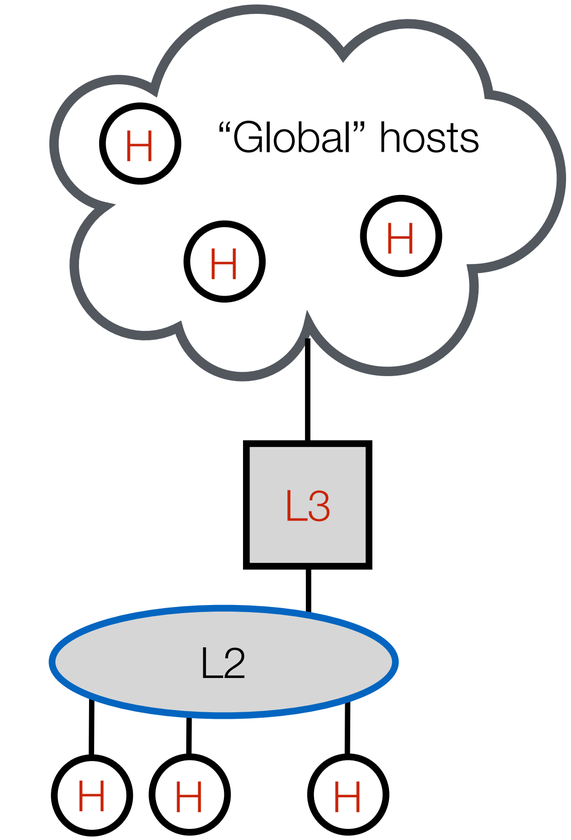

Hosts on a local subnet

Subnetting

- Problem: deciding whether an IP address belongs to a specific subset of IP addresses
- Solution: nodes on a subnet (and only those) have IP addresses within a specific range
- Simplifying HW implementation: subnets (ranges of IP addresses that can be placed only in a specific, constrained way):

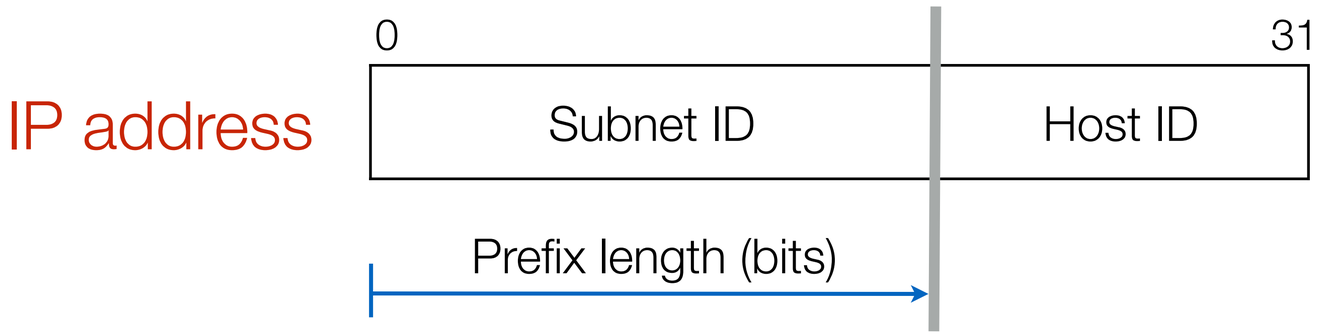

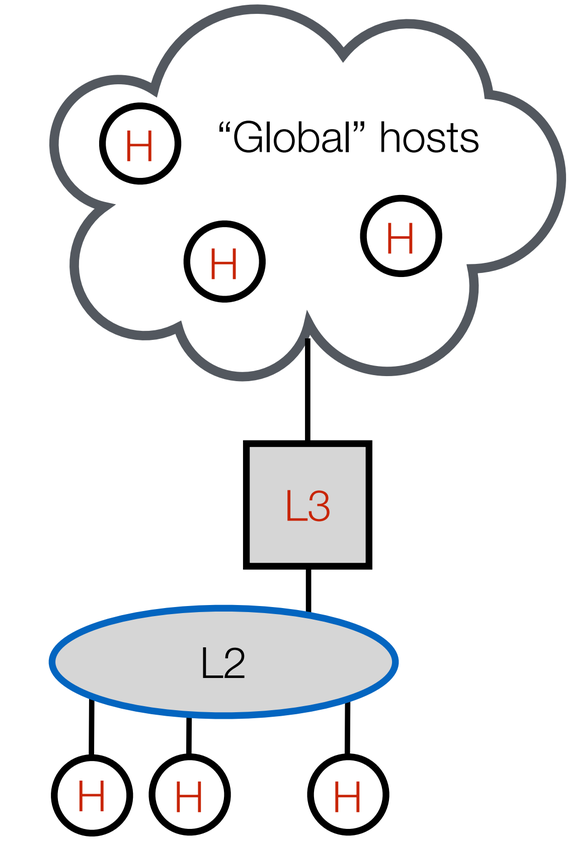

Hosts on a local subnet

IP Prefix

- specifies a range of consecutive IP addresses
- consists of a *subnet id* and a *length* (len)
  e.g., 132.177.4.0/22
- An IP address belongs to a range specified by a prefix if its first len bits are equal to those in the *subnet id* ($s_i = b_i$ for all $i \leq$ len):

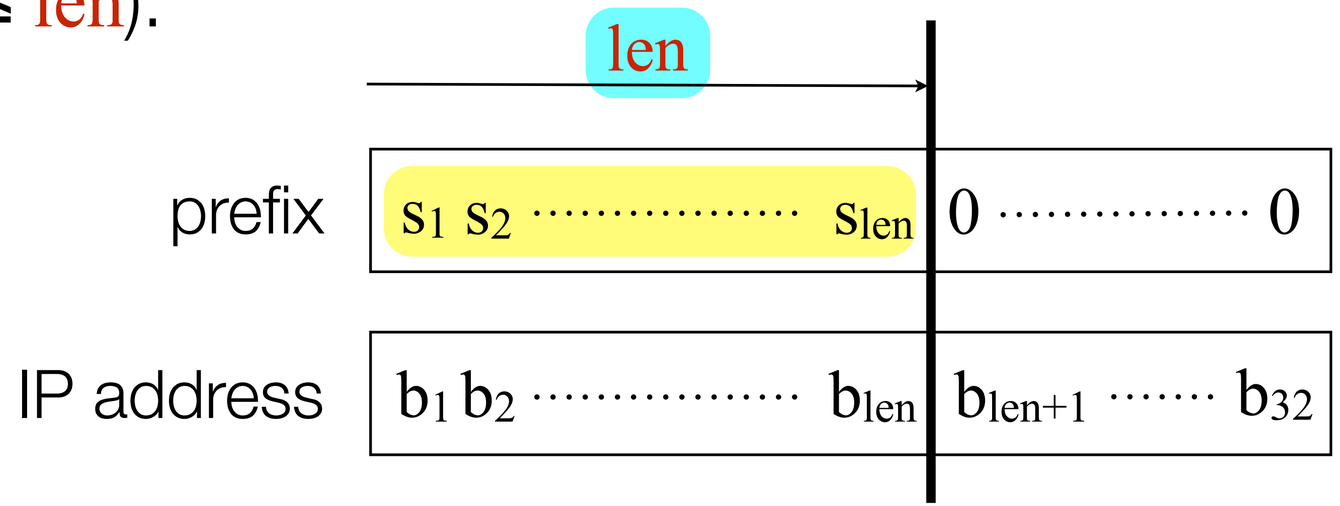

Example questions

- Range of 132.177.4.0/24?

  132.177.4.0 → 132.177.4.255

- Prefix for range 132.177.0.0 → 132.177.255.255?

  132.177.0.0/16

Example questions

- Range of 132.177.4.0/26?

  132.177.4.0 → 132.177.4.63

- Prefix for range 132.177.2.192 → 132.177.2.223?

  132.177.2.192/27

Combining prefixes

132.177.0.0/24	Range 132.177.0.0 - 132.177.0.255
132.177.1.0/24	Range 132.177.1.0 - 132.177.1.255

Combining prefixes

132.177.0.0/24
132.177.1.0/24

132.177.0.0/23

Range 132.177.0.0 - 132.177.1.255

Combining prefixes

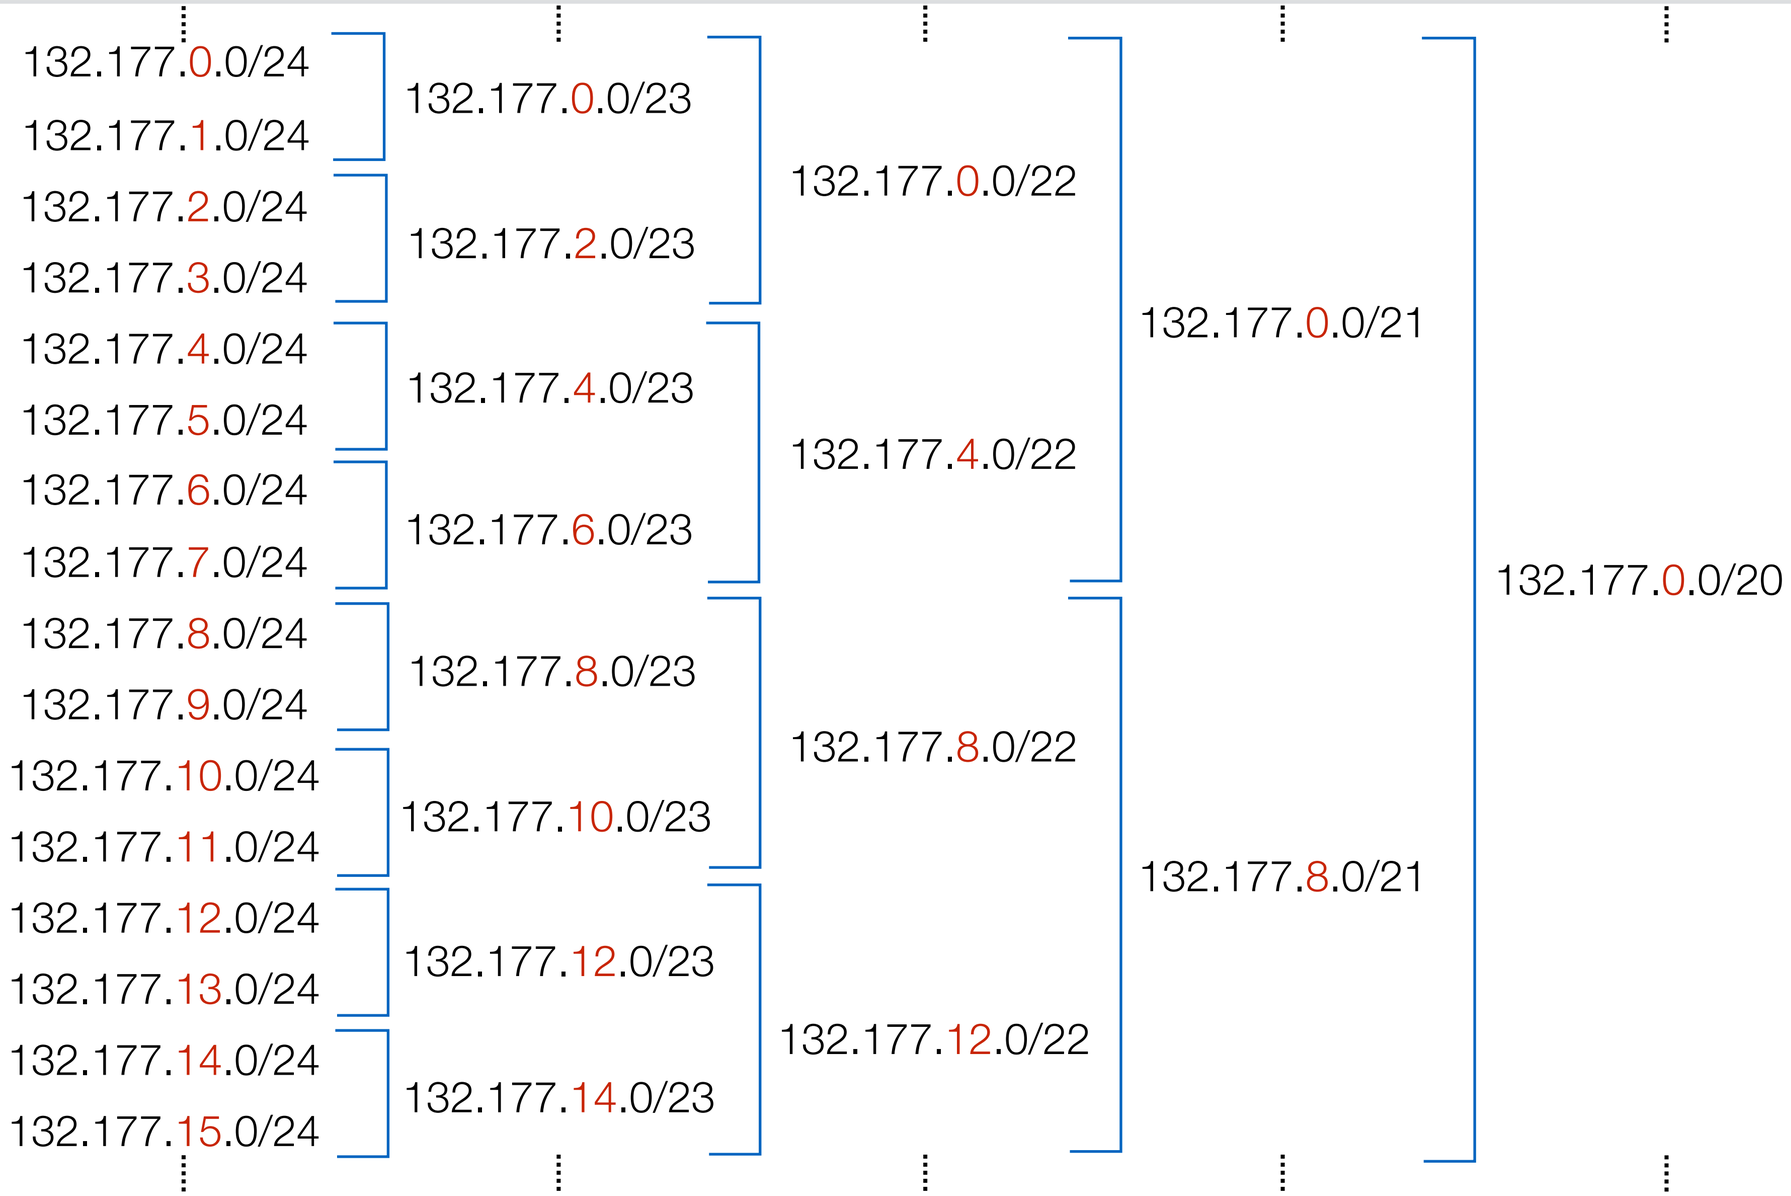

Subnet Mask

- Another way to specify prefix length
- A 32 bit, IP address-like value whose binary representation has binary ones in in bits corresponding to the subnet id bits.

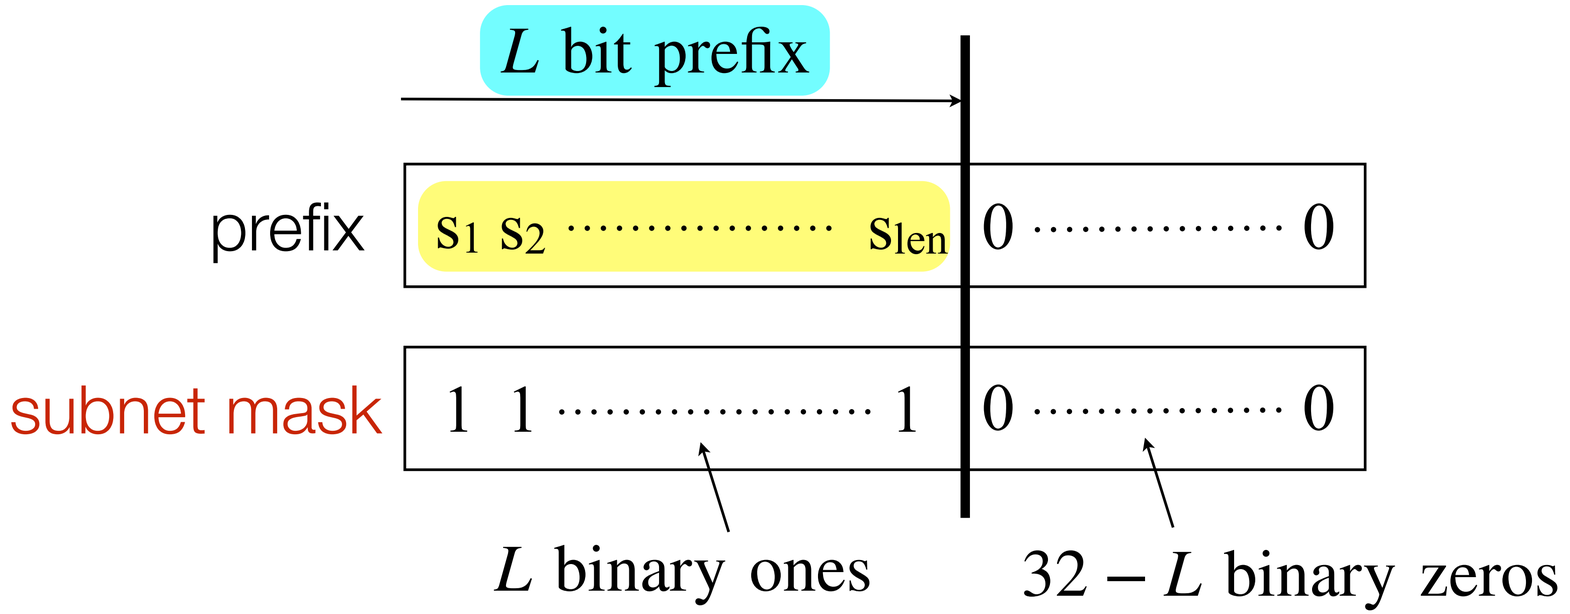

Subnet Mask Representation

- Typically represented using *decimal dotted notation*:

  255.255.0.0

- Subnet length and subnet mask (netmask) are for all practical purposes equivalent:

132.177.4.0 with netmask 255.255.255.0

is equivalent to 132.177.4.0 /24

Example questions

- Subnet mask for prefix length /17?

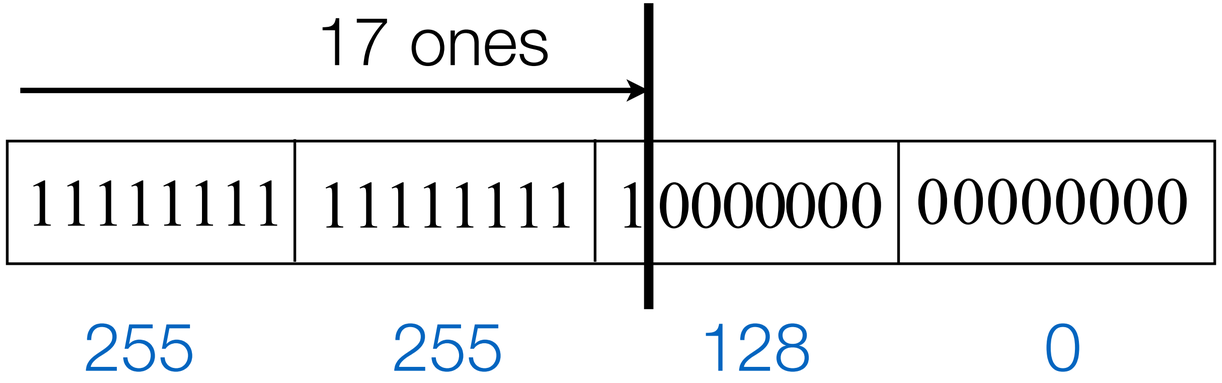

- Prefix length for subnet mask 255.255.255.192?

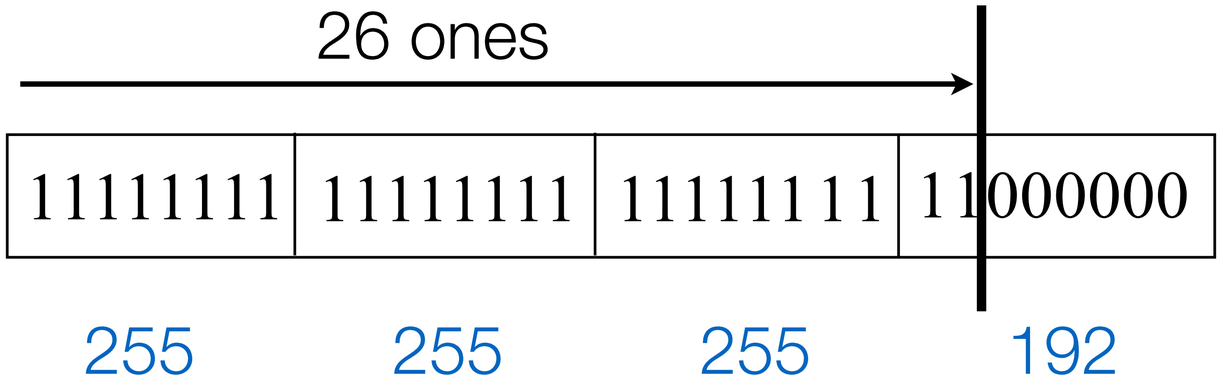

Domain Name Service

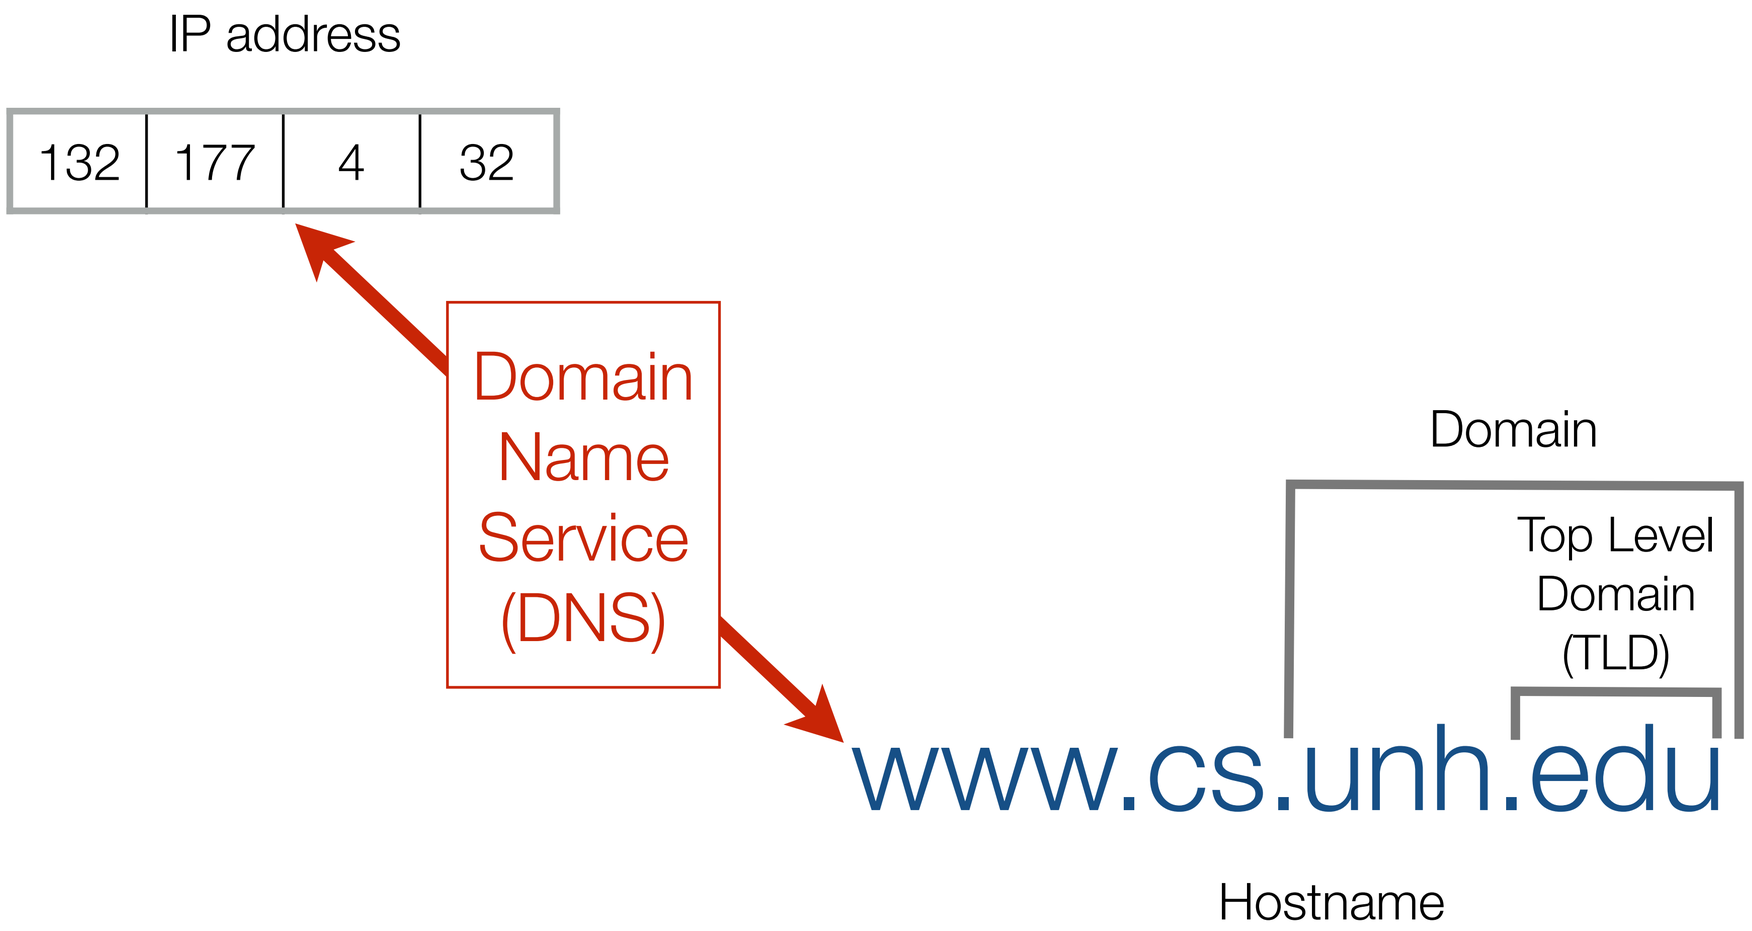

Domain Name Service

- Mapping between hostnames and IP addresses:
  - one-to-one, one-to-many, many-to-one, or many-to-many?
  - mapping in both directions
- Possible solutions:
  - centralized database
  - fully distributed database