


Lecture 10:

The Link Layer

Katerina Argyraki, EPFL


Link vs. network layer

- Link layer: takes each packet from one end of **one link** to the other end
- Network layer: takes each packet from one end of **the network** to the other end



IP subnet point of view

- Link layer: takes packet from one end of **one physical link** to the other end
- Network layer: takes packet from one end of **the IP subnet** to the other end



IP routers = network-layer switches

Internet point of view

- Link layer: takes packet from one end of **one IP subnet** to the other end
- Network layer: takes packet from one end of **the Internet** to the other end

The “link layer”

- Link layer **of an IP subnet**: takes packet from one end of **one physical link** to the other end
- Link layer **of the Internet**: takes packet from one end of **one IP subnet** to the other end

The “link layer”

- Link layer **of an IP subnet**: takes packet from one end of **one physical link** to the other end
- Link layer of the Internet: takes packet from one end of one IP subnet to the other end


Link-layer services

- Error detection
 - * receiver detects and drops corrupted packets
 - * relies on checksums
- Reliable data delivery
 - * sender/receiver detect corruption and loss, and try to recover
 - * relies on checksums, ACKs, retransmissions, ...
 - * only for error-prone links, typically wireless

Link-layer services

- Medium access control (MAC)
 - * sender manages access to shared medium (typically wireless link)
 - * listens for ongoing transmissions or "collisions"
 - * backs off and retries later

Why reliable data delivery at the link layer?
The transport layer does that anyway.



The “link layer”



- Link layer of an IP subnet: takes packet from one end of one physical link to the other end
- Link layer **of the Internet**: takes packet from one end of **one IP subnet** to the other end



Outline

- Addressing
- Forwarding
- Learning
- Address resolution

Outline

- Addressing
- Forwarding
- Learning
- Address resolution






MAC address

- 48-bit number
 - * typical format: 1A-2B-DD-78-CF-CC
 - * the value of each byte as hexadecimal
- Flat
 - * not hierarchical like IP address
 - * not location dependent

Outline

- Addressing
- Forwarding
- Learning
- Address resolution

MAC address	link
5c:f9:38:a4:00:76	a
ab:ff:4d:a4:bb:af	b
...	...



L2 forwarding

- Local switch process that determines output link for each packet
- Relies on forwarding table
 - * maps destination MAC addresses to output links
- Similar to IP (L3) forwarding, except...

MAC address

- Flat
 - * not hierarchical like IP addresses
 - * not location dependent

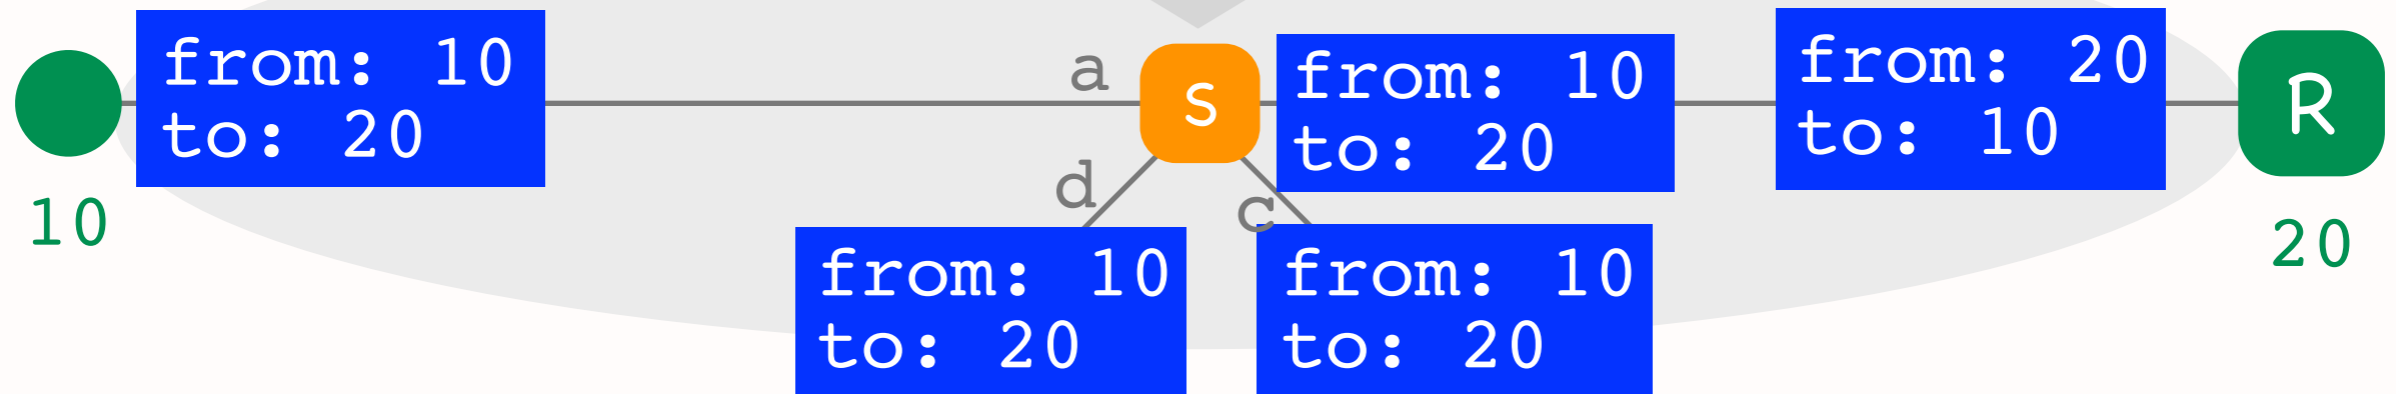
L2 vs. IP forwarding

- L2: relies on flat addresses
 - * no way to group MAC addresses in prefixes
 - * forwarding table size = # of active destination MAC addresses in the IP subnet
- IP (L3): relies on hierarchical addresses
 - * IP addresses grouped in IP prefixes
 - * forwarding table size = # of IP prefixes in the world

Outline

- Addressing
- Forwarding
- Learning
- Address resolution

MAC address	link
10	a
20	b




L2 learning


- Switch learns from traffic
 - * when packet with src MAC x arrives at link y , switch adds $\text{MAC } x \rightarrow \text{link } y$ mapping to forwarding table
- Broadcasts when it does not know
 - * when packet with unknown dst MAC arrives, switch broadcasts the packet
- Serves similar role as IP routing, but...


L2 learning vs. IP routing

- L2 learning: relies on actual traffic
 - * switches do not exchange explicit routing information
- IP routing: relies on routing protocol
 - * routers exchange explicit routing messages

naïve broadcasting = forwarding loops!








Spanning tree

- A subgraph with special properties
 - * includes all nodes + some edges
 - * cannot remove an edge without disconnecting a node
- Useful for loop-free broadcasting
 - * broadcast traffic propagated only along tree
 - * prevents forwarding loops

Outline


- Addressing
- Forwarding
- Learning
- Address resolution

I want to send a packet to IP address 128.178.2.20.
Which destination MAC address should I use?




I want to send a packet to IP address 128.178.2.20.
Which destination MAC address should I use?

Use my MAC address! It is 20.




I want to send a packet to IP address 104.20.228.42.
Which destination MAC address should I use?




I want to send a packet to IP address 104.20.228.42.
Which destination MAC address should I use?

Use my MAC address! It's 20.




My default gateway has IP address 128.178.2.1.
Which is its MAC address?



My default gateway has IP address 128.178.2.1.
Which is its MAC address?

That's me. It's 20.



Address Resolution Protocol (ARP)

- Goal: map IP address to MAC address
 - * Alice knows destination IP address
 - * which destination MAC address to use?
- How: broadcast request, targeted response
 - * Alice broadcasts her request
 - * the right entity responds to Alice
- Serves similar role as DNS, but...

Broadcasting

- Alice sends request to special, broadcast destination MAC address
 - * FF-FF-FF-FF-FF-FF
- Reaches every end-system and router in the local IP subnet

ARP vs. DNS

- ARP: relies on broadcasting
 - * no logically centralized map
 - * each entity knows its own MAC address and knows which requests to respond to
- DNS: relies on DNS infrastructure
 - * logically centralized map
 - * stored in DNS servers

Basic Ethernet elements

- Address Resolution Protocol
 - * resolves IP address to MAC address
- L2 forwarding
 - * based on MAC addresses (which are flat)
- L2 learning
 - * populates switch forwarding table

Basic Ethernet elements


- Address Resolution Protocol [rel. to DNS]
 - * resolves IP address to MAC address
- L2 forwarding [rel. to IP forwarding]
 - * based on MAC addresses (which are flat)
- L2 learning [rel. to IP routing]
 - * populates switch forwarding table

Get rid of IP addresses and IP forwarding?

Get rid of MAC addresses and L2 forwarding?

Three levels of hierarchy

- IP subnet
 - * L2 forwarding
 - * L2 learning
- Autonomous System (AS)
 - * IP (L3) forwarding
 - * intra-domain routing
- Internet
 - * IP (L3) forwarding
 - * inter-domain routing (BGP)



A types `http://www.epfl.ch` in her browser

At least 4 packets:

A's DNS request to local DNS server

local DNS server's response to A

A's HTTP GET request to web server

web server's response to A

A types `http://www.epfl.ch` in her browser

At least 4 packets:

A's DNS request to local DNS server

local DNS server's response to A

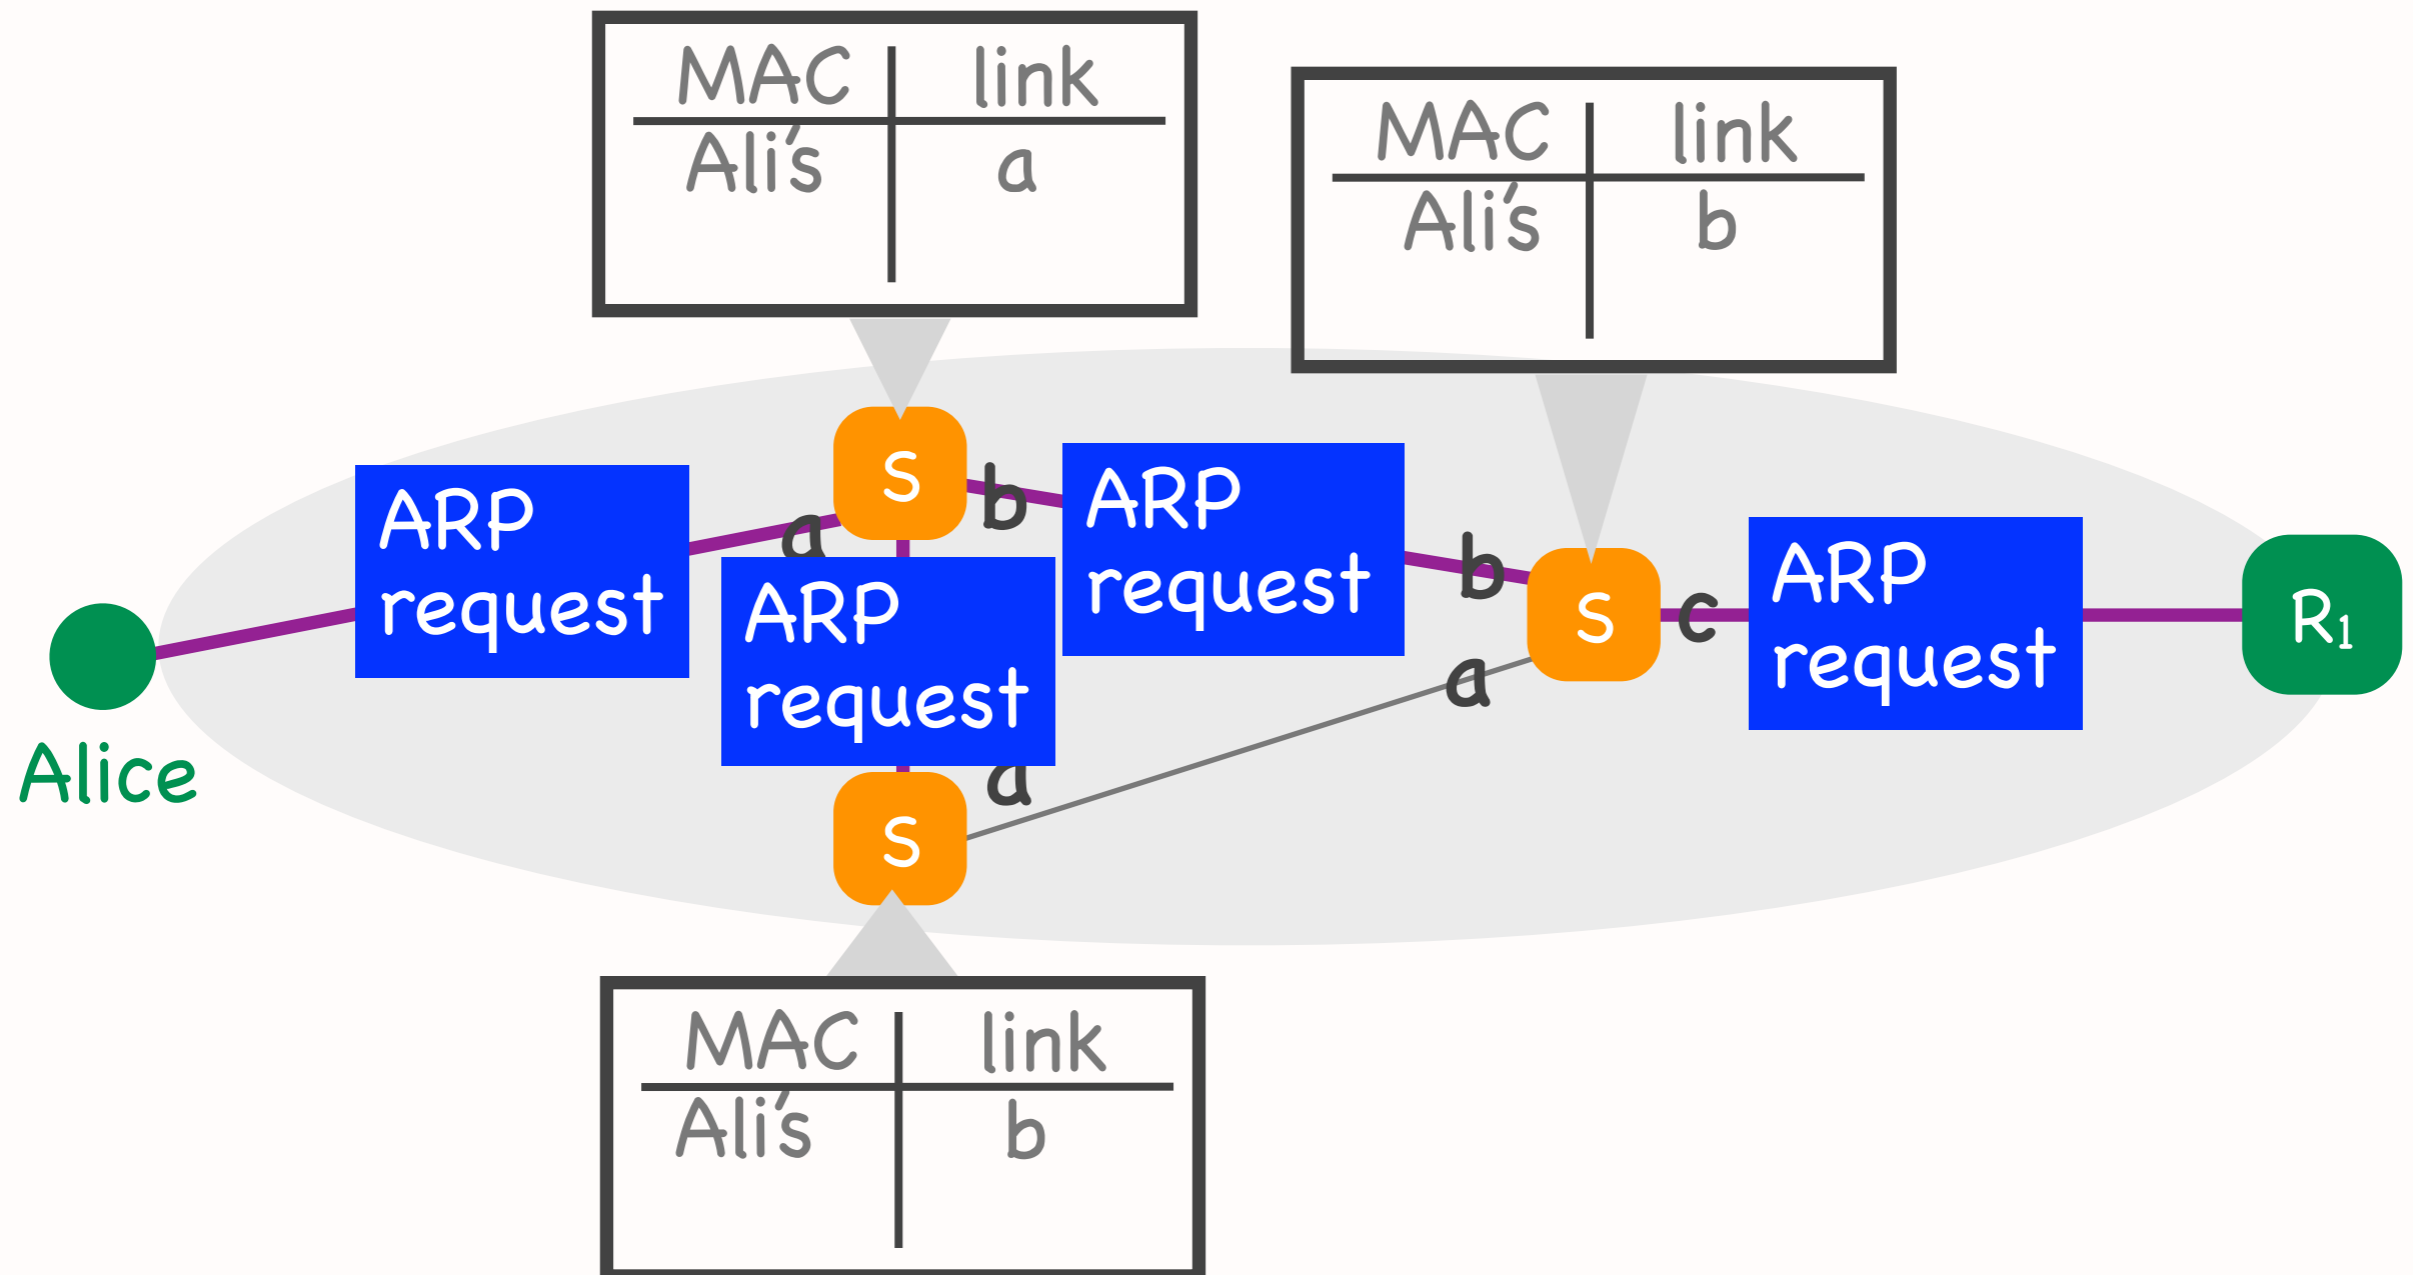
A's HTTP GET request to web server

web server's response to A


1. A's DNS client process creates DNS request
2. Passed down to transport, network layer
 - IP src: A's IP address
 - IP dst: local DNS server's IP address
3. A's network layer sends ARP request
 - to resolve DNS server's IP address

src MAC: Alice's

dst MAC: broadcast



4. R_1 's network layer sends ARP response

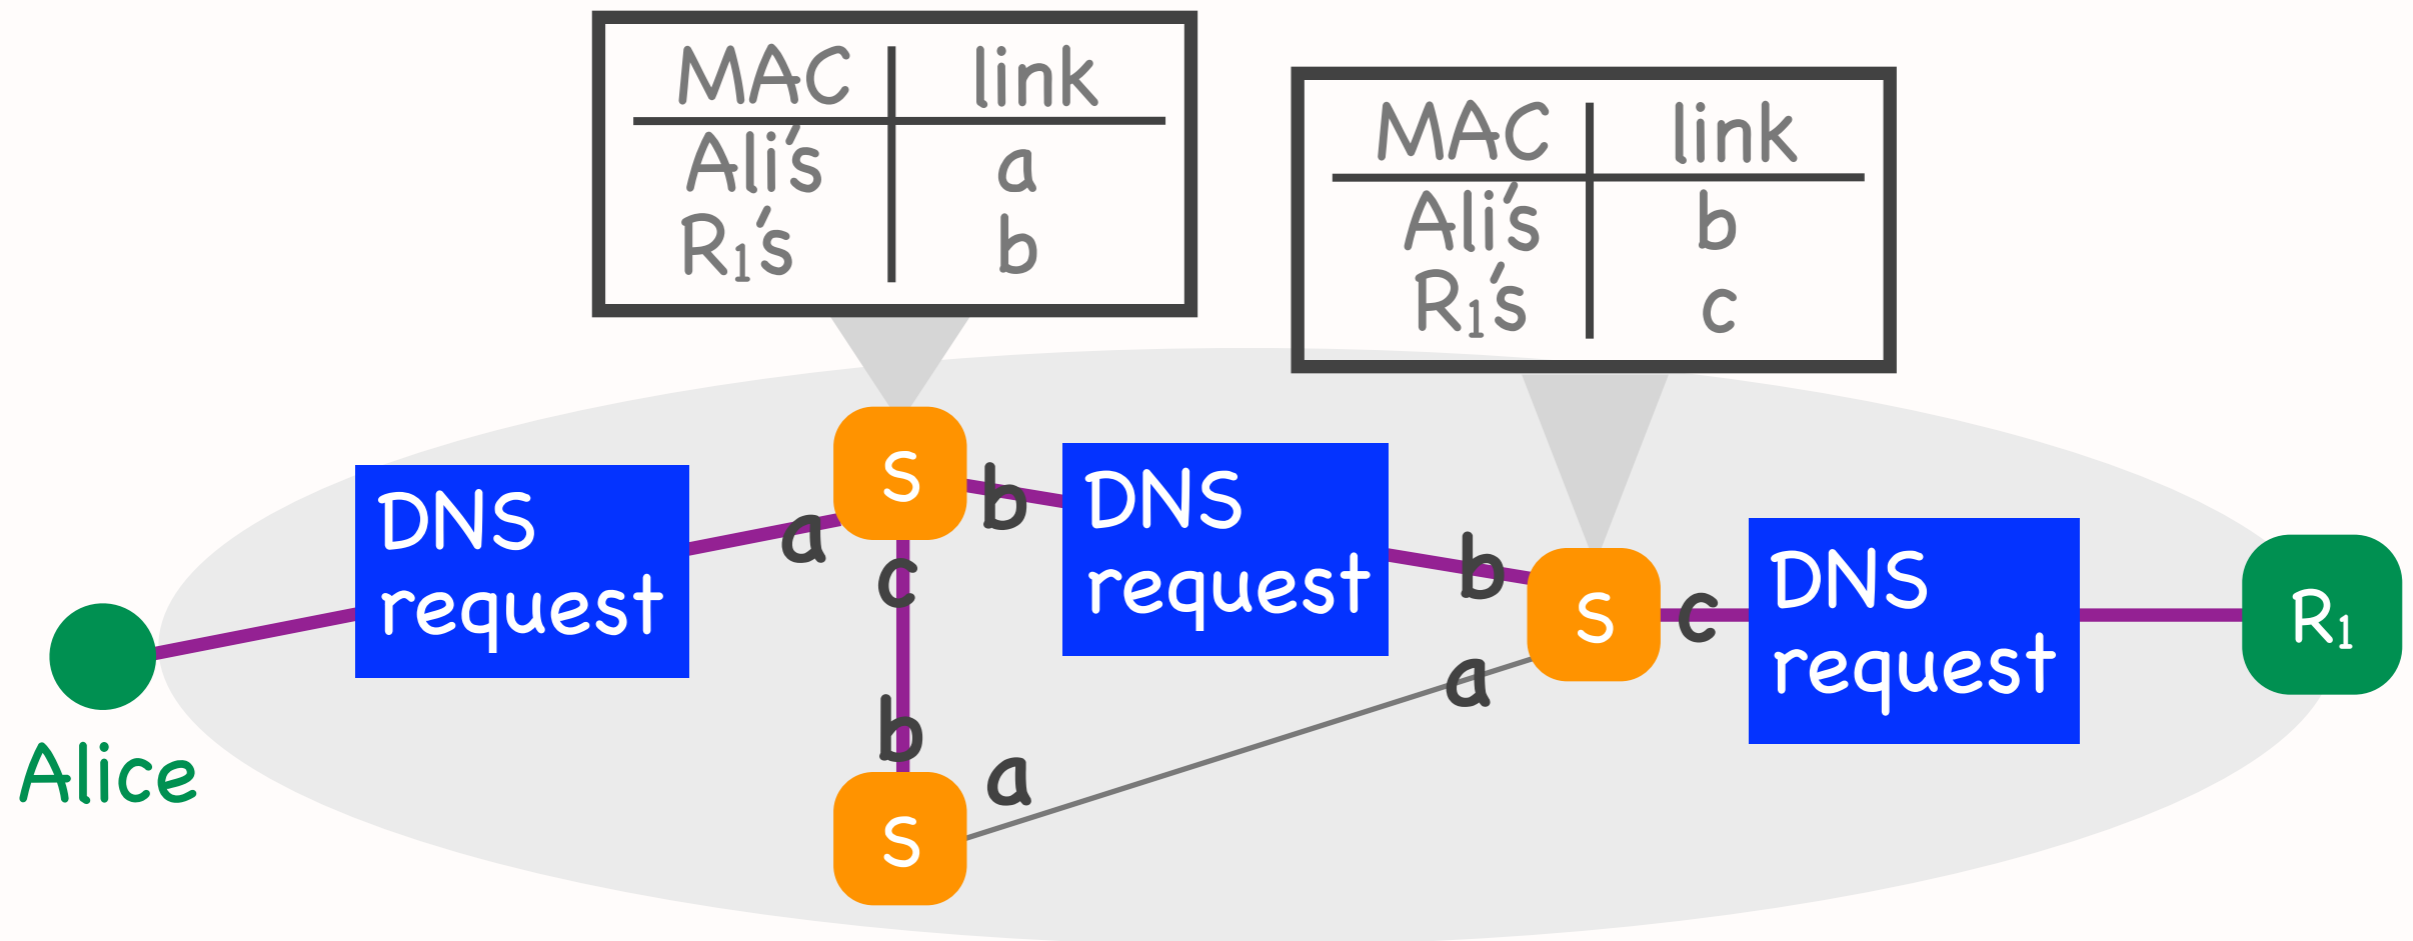


src MAC: R₁'s
 dst MAC: Alice's

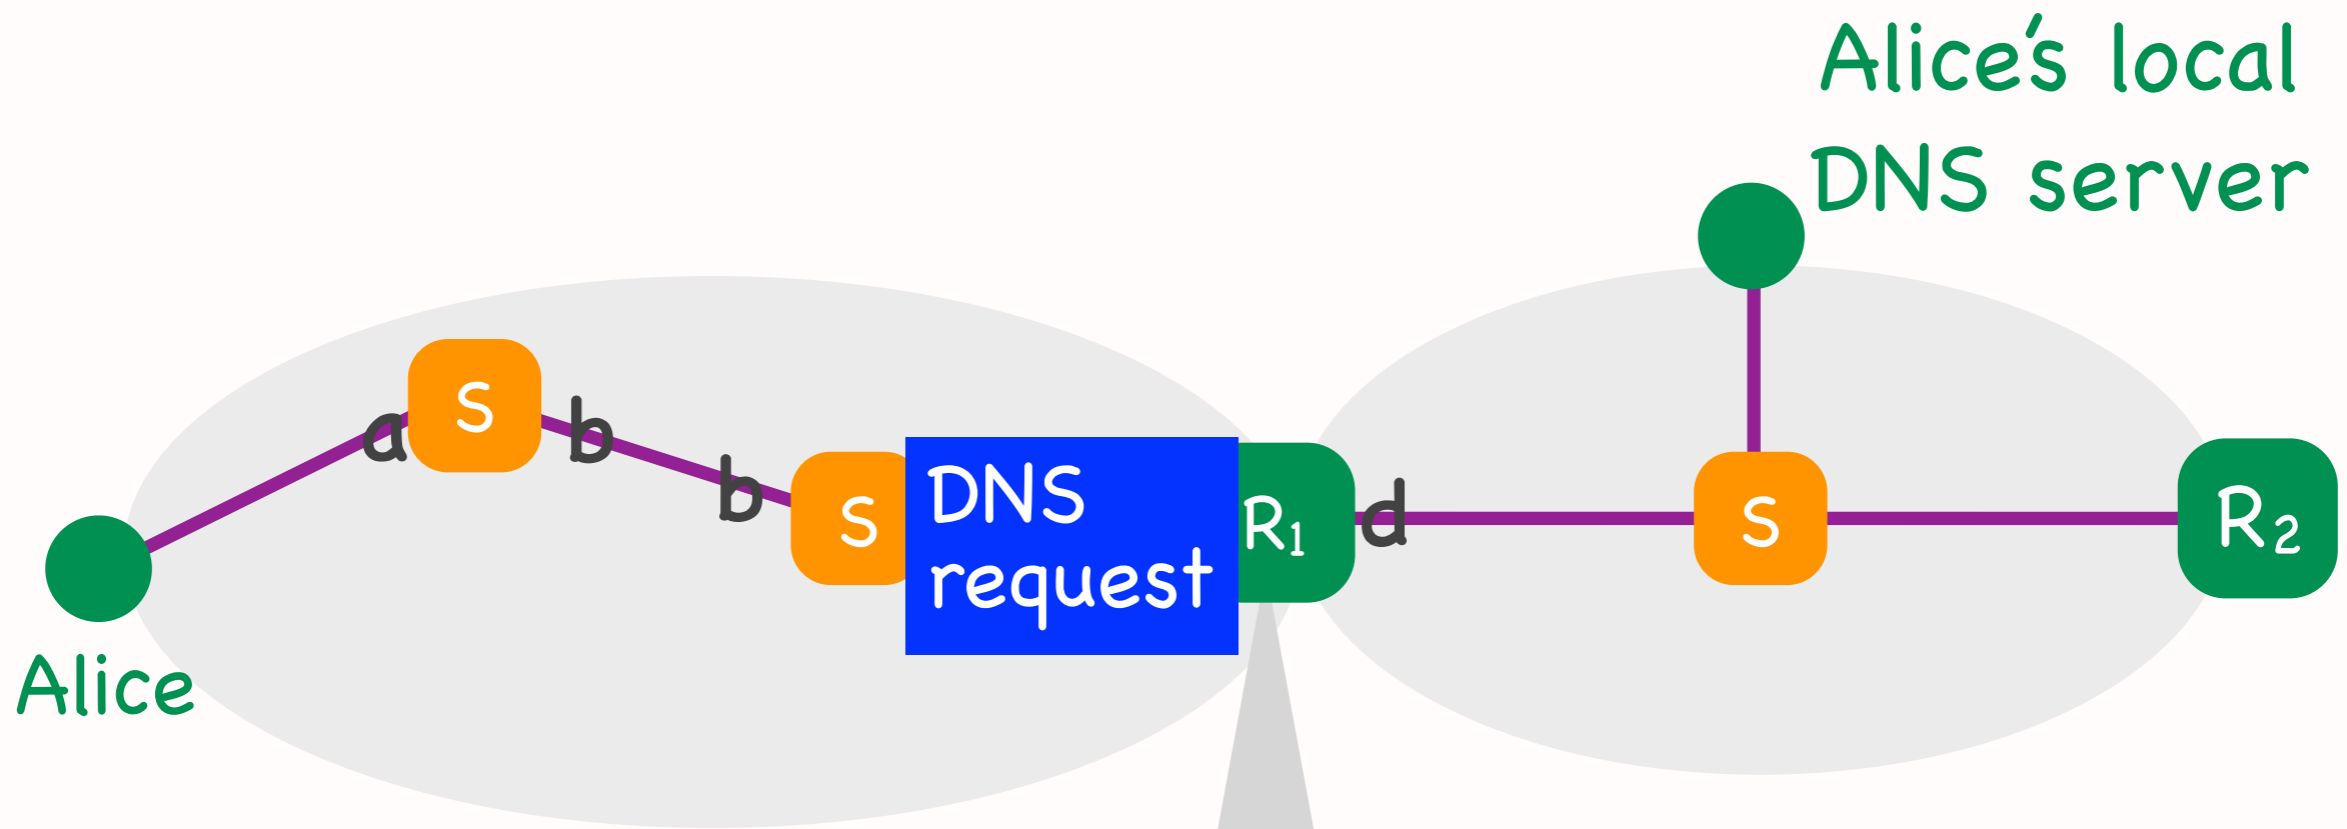
5. As network layer sends DNS request
 - it now knows what dst MAC address to use

src MAC: Alice's
dst MAC: R₁'s

src IP: Alice's
dst IP: DNS server's



6. R_1 's network layer performs IP forwarding



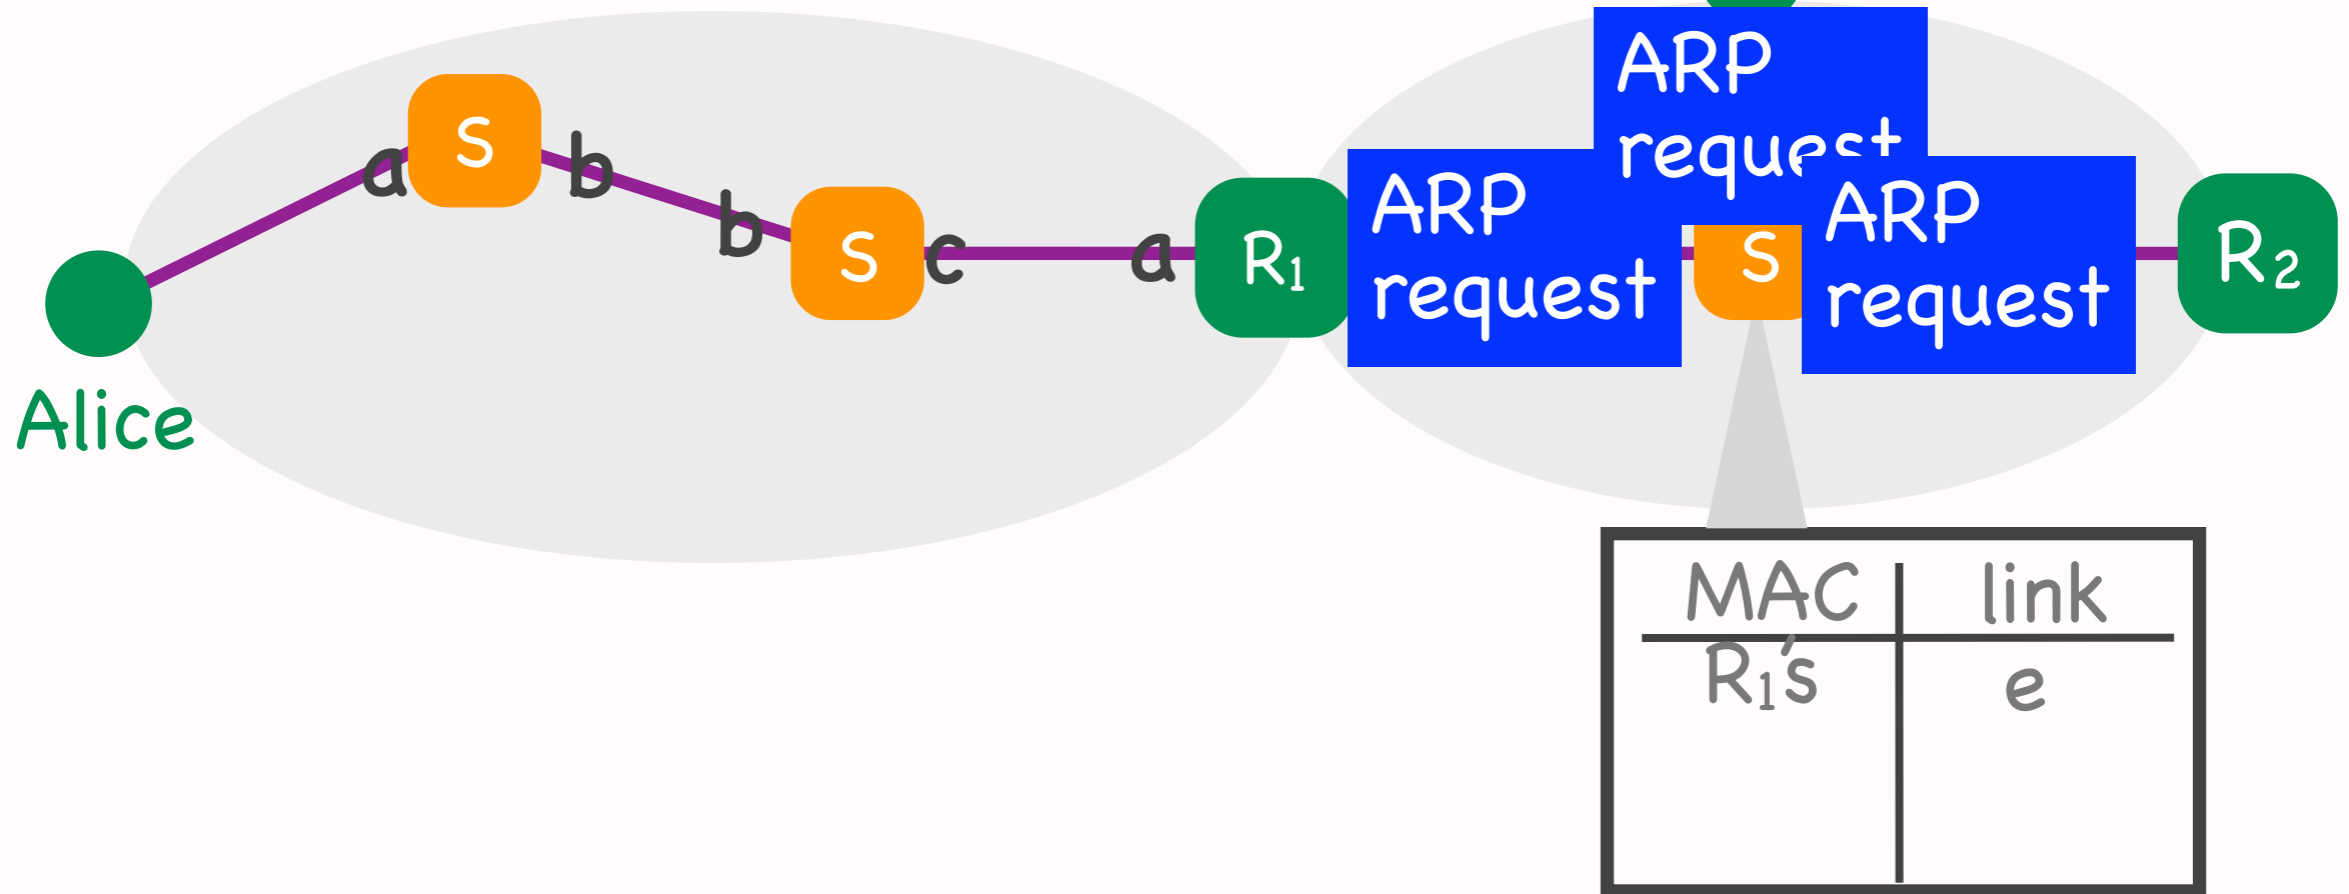
IP prefix	link
11.2.34.0/24	a
8.0.0.0/8	c
19.7.0.0/16	d
...	...

7. R_1 's network layer sends ARP request
 - to resolve DNS server's IP address

src MAC: R_1 's


dst MAC: broadcast

Alice's local
DNS server



8. DNS server's network layer sends ARP response

src MAC: DNS server's
dst MAC: R₁'s




9. R_1 's network layer forwards DNS request
 - it now knows what dst MAC address to use

src MAC: R_1 's

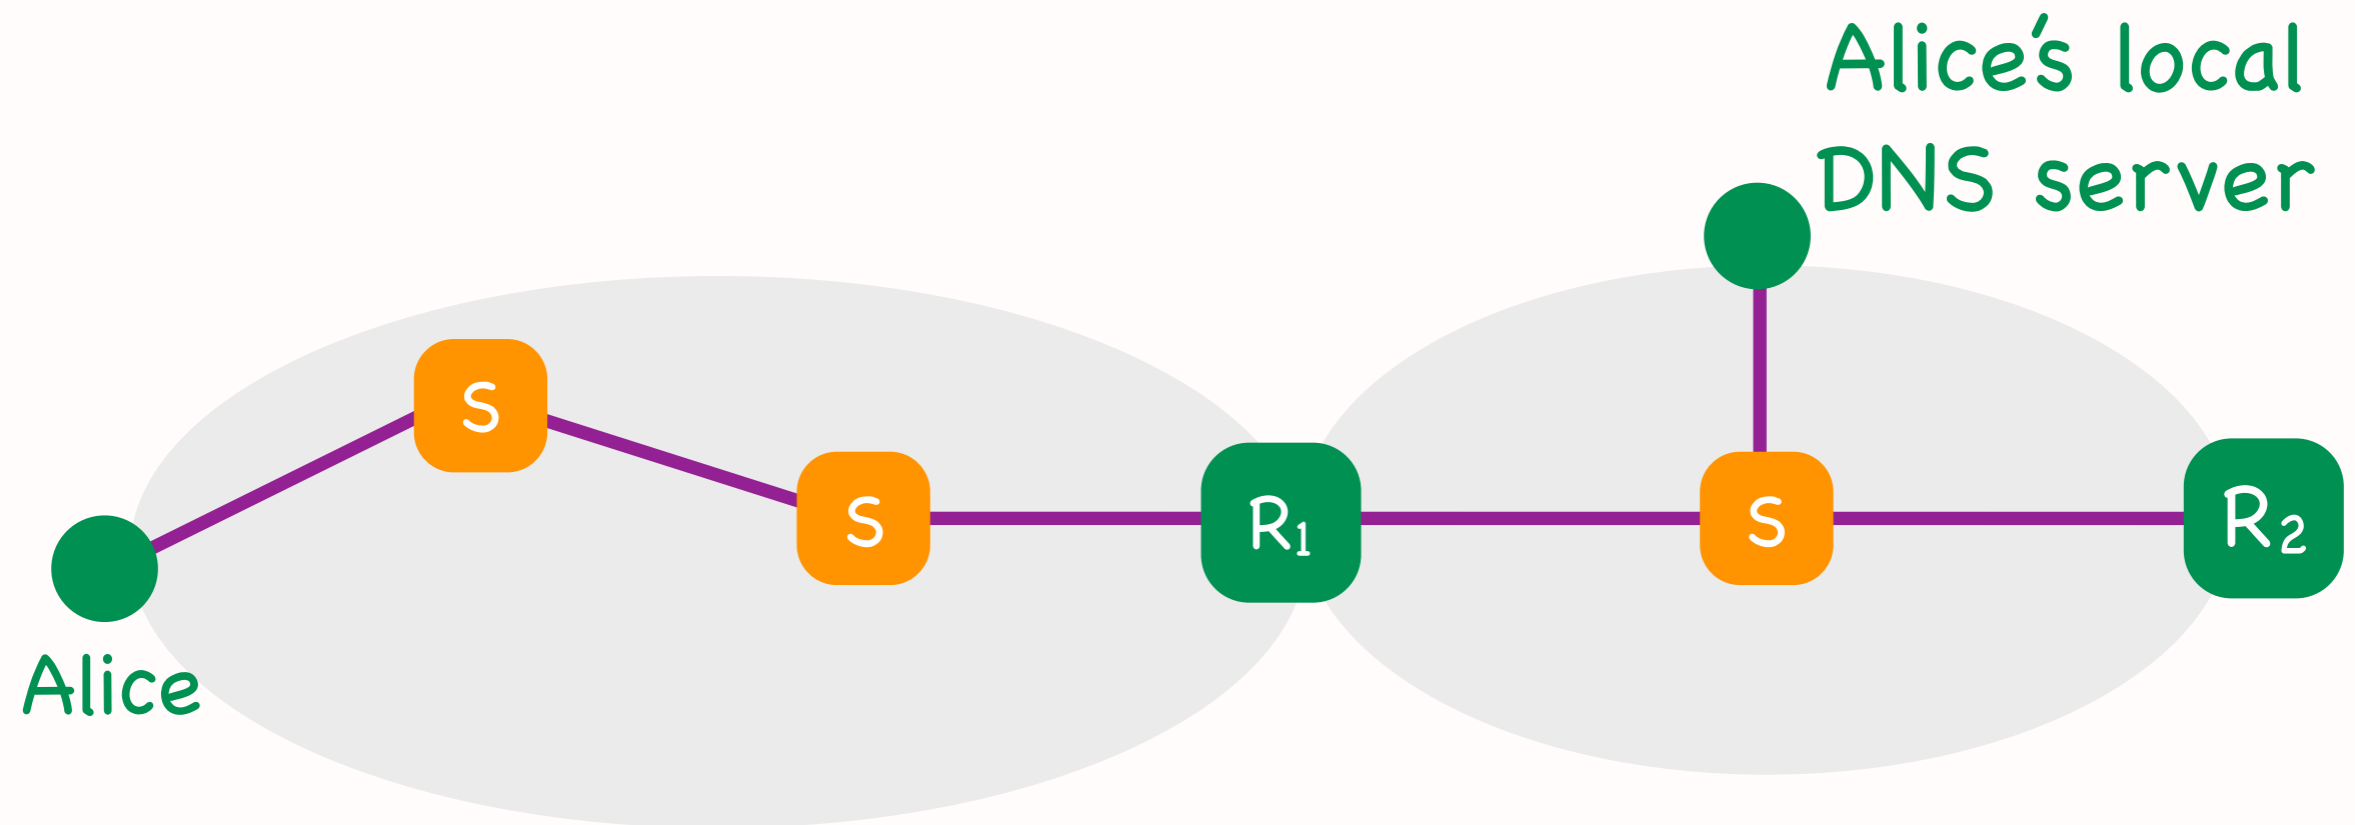
dst MAC: DNS server's

src IP: Alice's

dst IP: DNS server's




The switches forward traffic within local IP subnet between end-systems and routers



End-systems and routers need MAC addresses
Switches do not (*)

The routers forward traffic end-to-end between source and destination end-systems



End-systems need IP addresses
Routers do not (*)

Switch and router addresses

- Yet both switches and routers have both MAC and IP addresses
 - * 1 MAC address + \geq 1 IP address per network interface
- For various practical reasons
 - * to be reachable by an administrator (*)
 - * for link testing (*)
 - * a router needs an IP address to respond to ARP requests
 - * a router that acts as a NAT gateway needs an IP address for NAT