Internetworking II: MPLS, Security, and Traffic Engineering

3035/GZ01 *Networked Systems*

Kyle Jamieson

Department of Computer Science
University College London

Last time: Internetworking

- IP interconnects many heterogeneous networks
  - The Internet is a **datagram** network
    - Each datagram has enough information to allow any switch to decide how to get it to its destination
  - IP is simple and responsible for Internet's success

- But, IP leaves certain questions **unresolved:**

1. What to do about the complexity of the **longest-prefix match (LPM)** for IP address lookup?
2. What about **privacy?**
3. What if we want more control over **where** traffic goes?

Today

- Three topics that address IP's shortcomings:

1. **MPLS**
2. Virtual private networks
3. Traffic engineering in the Internet

Multiprotocol label switching (MPLS)

- Widely-used part of the Internet's architecture, but largely hidden from end-users

- MPLS is a ***virtual circuit (VC)* network**
  - **Unlike IP,** MPLS establishes one or more connections (***circuits***) **before** moving data from **A** to **B**
    - **Unlike IP,** switches keep **connection state**
  - **Like IP,** MPLS sends **packets** over the connection

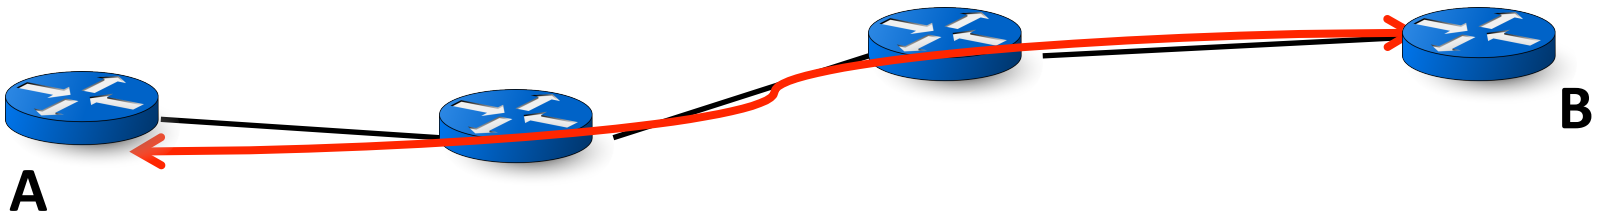

Label-switched forwarding

- MPLS routers forward based on ***labels*** instead of IP address
  - Labels have a **fixed length,** unlike CIDR IP addresses
  - Labels have **local scope,** unlike IP addresses: they only have meaning within one MPLS router

- Where are the labels? Inserted between the link- and network-layer headers, so **encapsulating** the IP datagram:

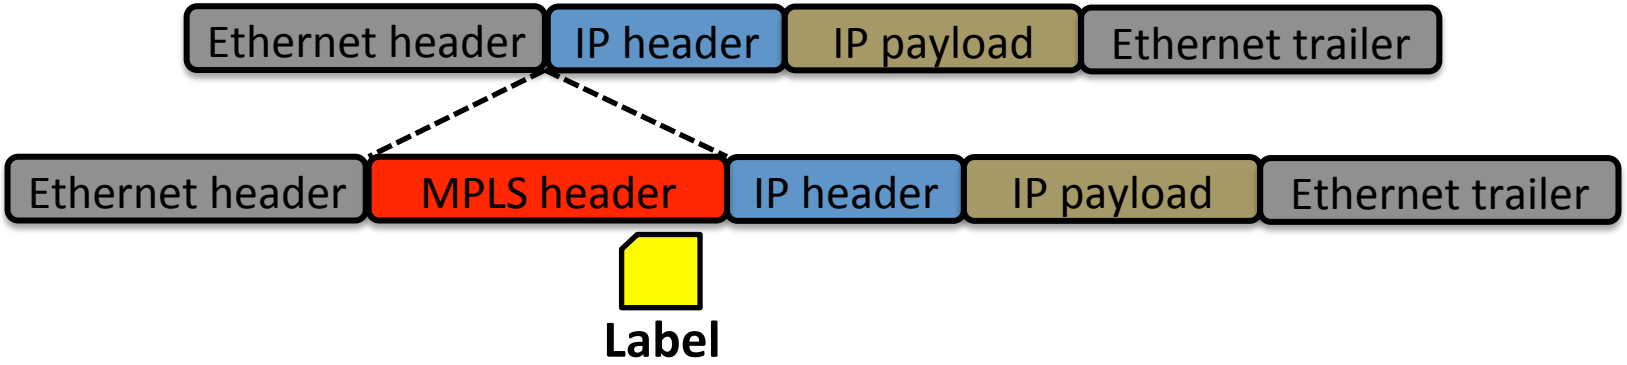

Comparison: IP address-based forwarding

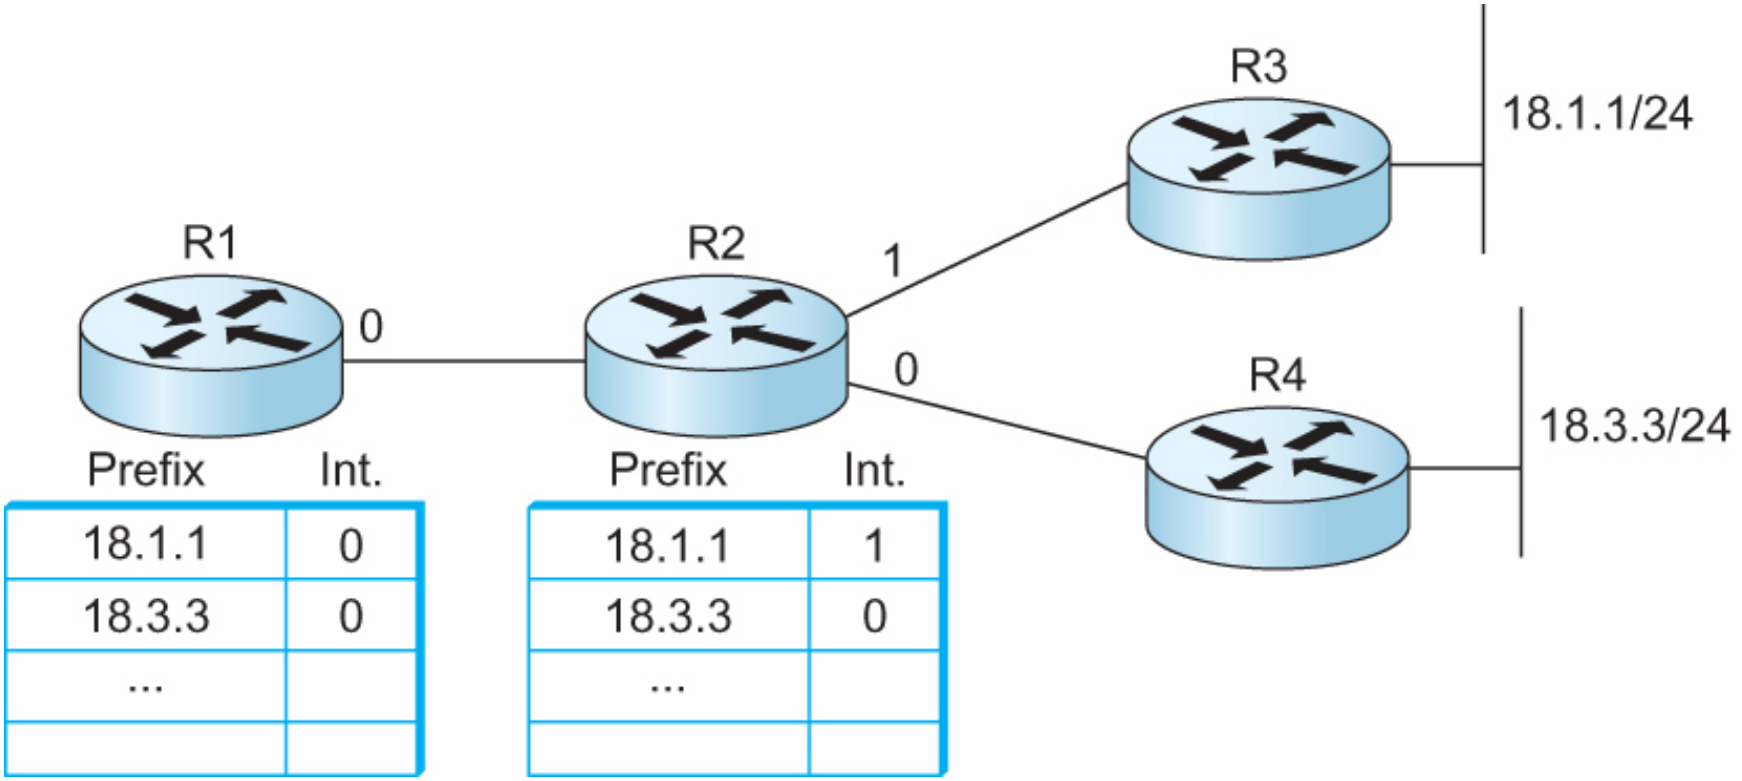

- R3 and R4 each have one connected network
- R1 and R2 have IP routing tables indicating **which outgoing interface** to use for each of the two networks

MPLS label-switched forwarding: Advertising labels

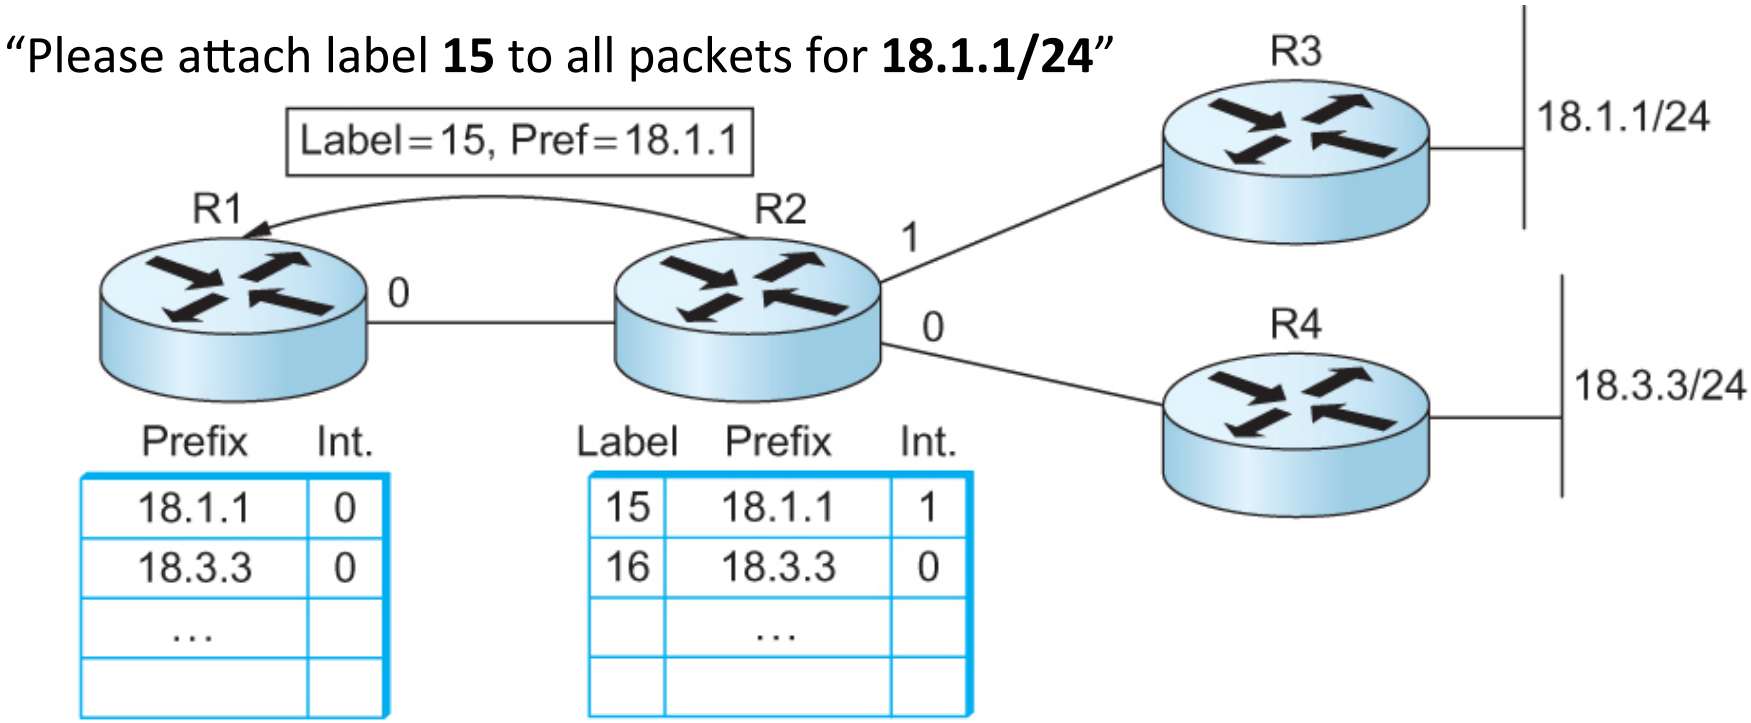

- Routers allocate, advertise a label for each routing table prefix
  - Can think of labels as **indices** into the allocating router's table

MPLS label-switched forwarding: Attaching labels

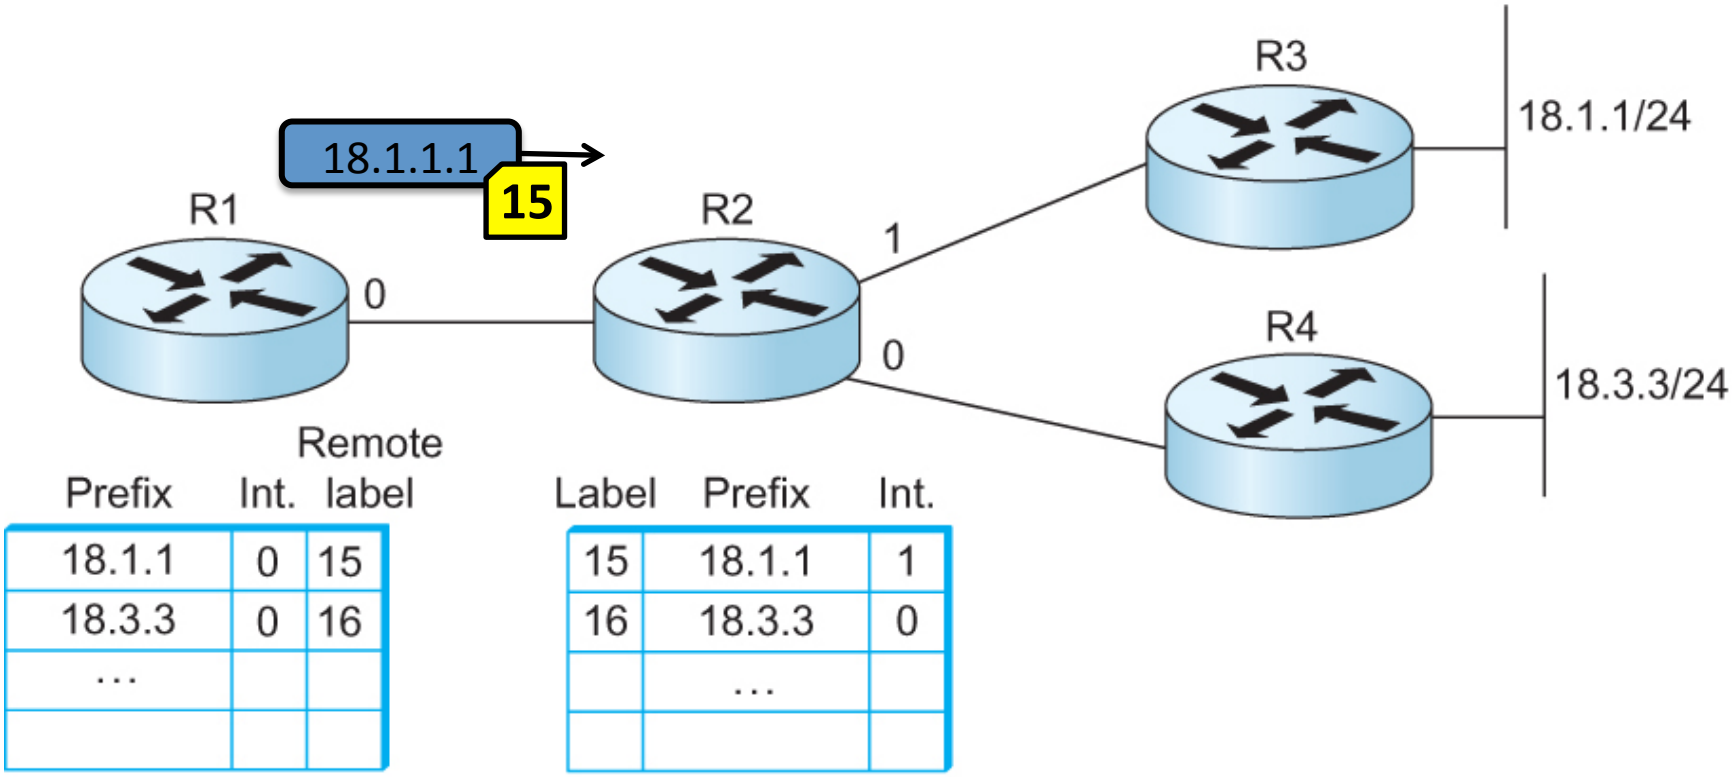

- On hearing advertisement, neighboring router **stores the remote label in its table** alongside the prefix it represents
- Routers **attach** the corresponding label to outgoing packets.

MPLS label-switched forwarding: Forming the virtual circuit

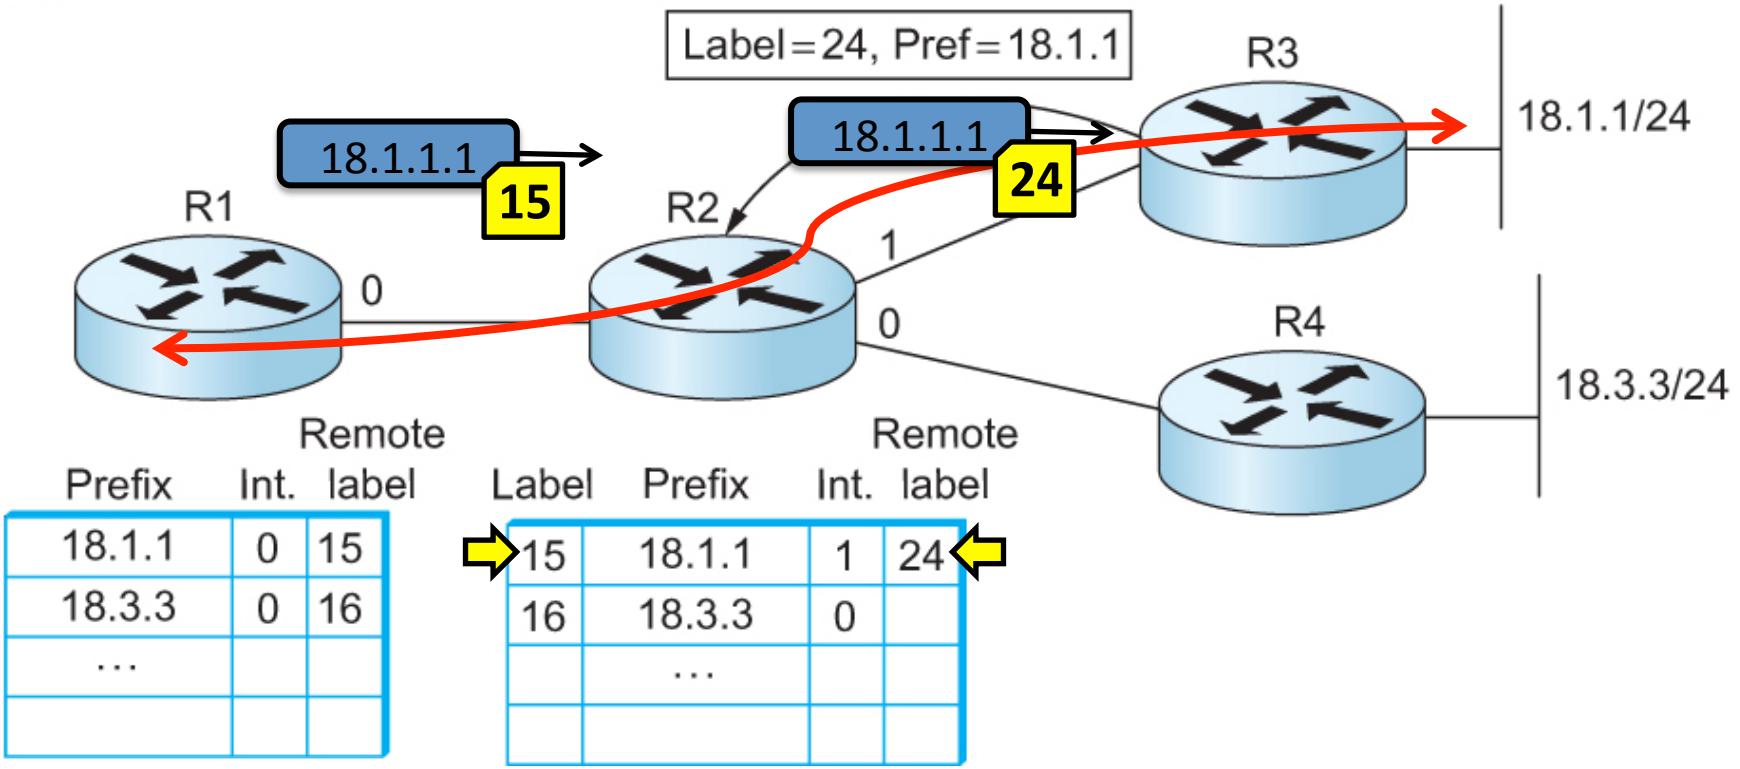

- "Threaded indices" of labels get built up over multiple hops
- **MPLS forwarding rule:** Replace an incoming packet's matching label with the corresponding remote label
- MPLS routers' label state forms a virtual circuit

Label edge routers accept IP packets

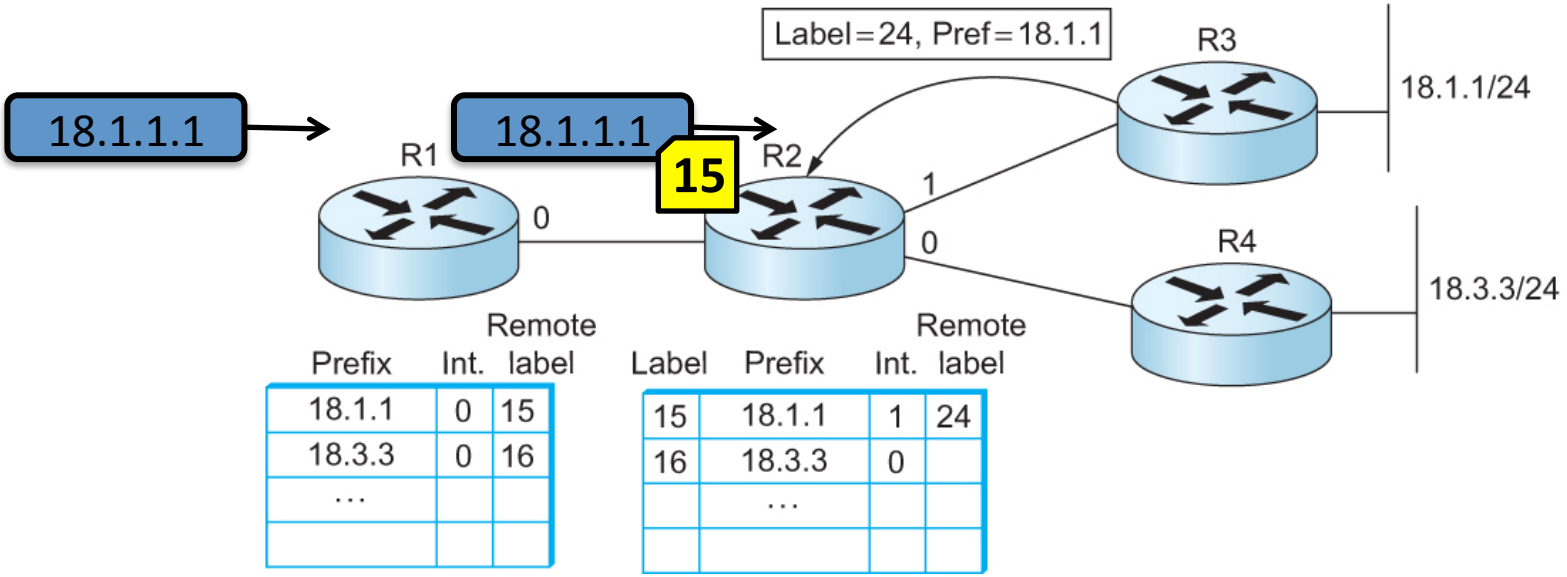

- R1 is a ***label edge router (LER),*** the first MPLS router at which a certain IP packet arrives
- R1 must perform a complete **LPM IP lookup** to apply label 15
- Thereafter, MPLS routers only look at labels, **avoiding LPM**

Today

- Three topics that address IP's shortcomings:

1. MPLS

2. **Virtual private networks**

3. Traffic engineering in the Internet

Private networks

- Internet addresses are **globally routable:** can send an IP packet to any device with a public IP address

- Sometimes, we want to restrict connectivity among nodes in the network as a whole
  - Confidentiality
  - Immunity from attack (denial-of-service, *et al.*)

- Corporations, governments often **lease private lines** and use these to interconnect different sites

Virtual private networks (VPNs)

- **Useful property*:*** VC requires that a circuit be established before data can flow

- VPNs use VCs in the Internet to restrict communication
  - But, the **Internet is a datagram network**
  - So we need a way of creating a VC there

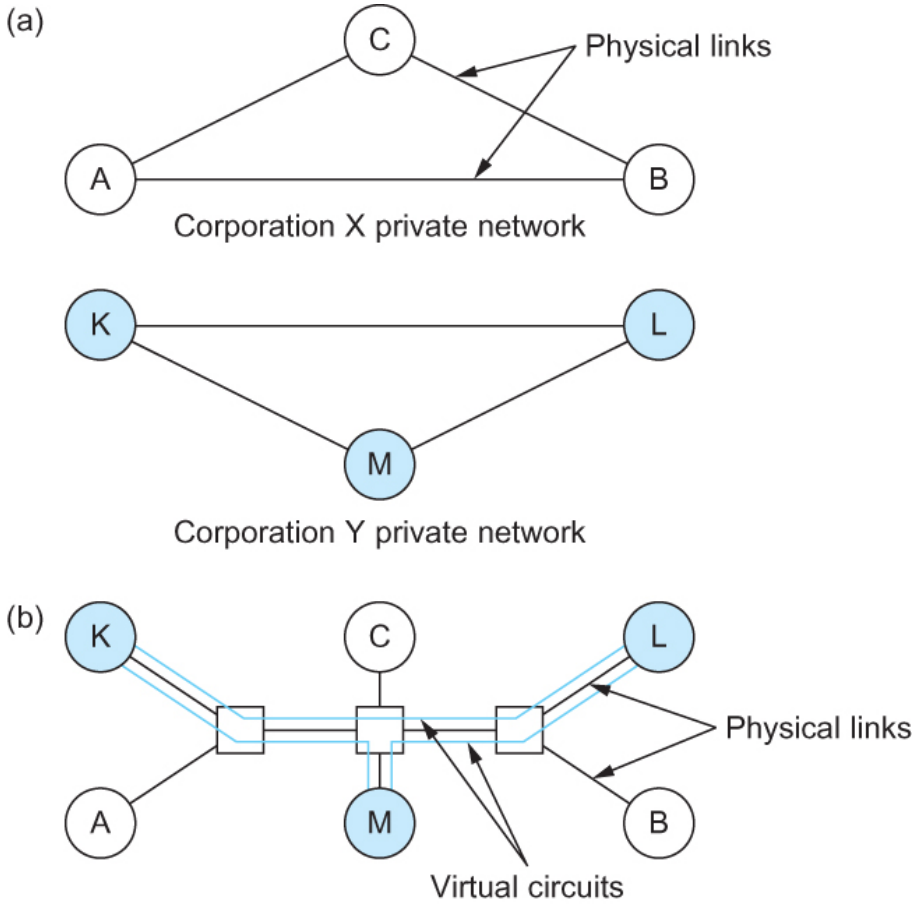

IP tunnels


Network 1
R1
Internetwork
R2
Network 2
18.5.0.1

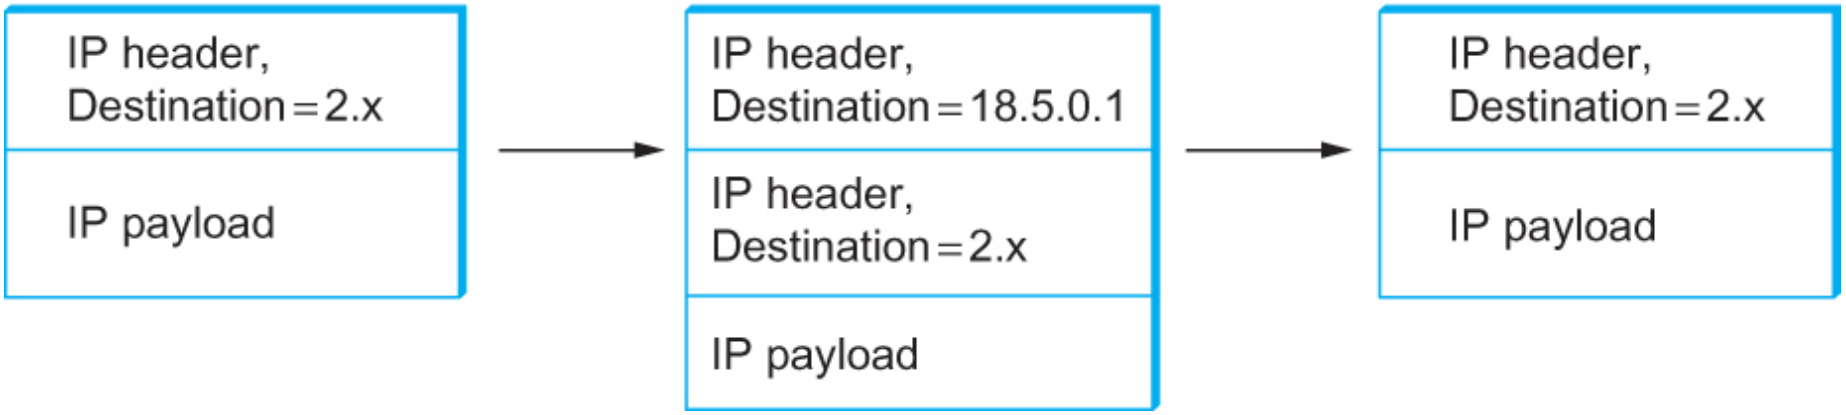
IP header,
Destination = 2.x
IP payload
IP header,
Destination = 18.5.0.1
IP header,
Destination = 2.x
IP payload
IP header,
Destination = 2.x
IP payload

IP tunnels

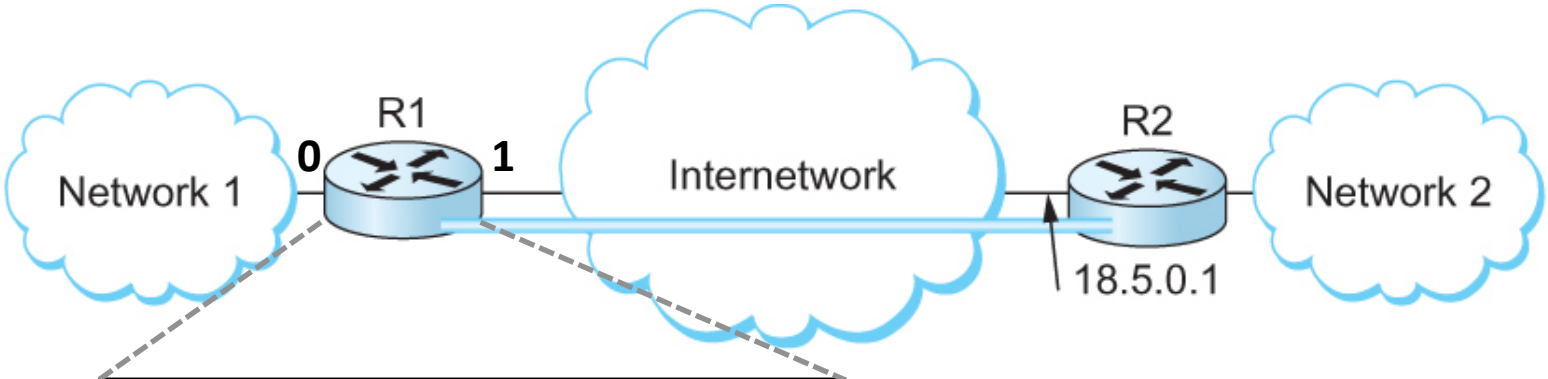

Network number	Interface
1	Interface 0
2	Virtual interface 0
(default)	Interface 1

- To set up the IP tunnel, **encapsulate** IP datagrams leaving virtual interface 0 in an IP datagram addressed to R2
- R2 drops encapsulated IP packets not signed by R1

Today

- Three techniques that address IP's shortcomings:

1. MPLS

2. Virtual private networks

3. **Traffic engineering**
   - **MPLS explicit routing**
   - IP anycast

IP's source routing option

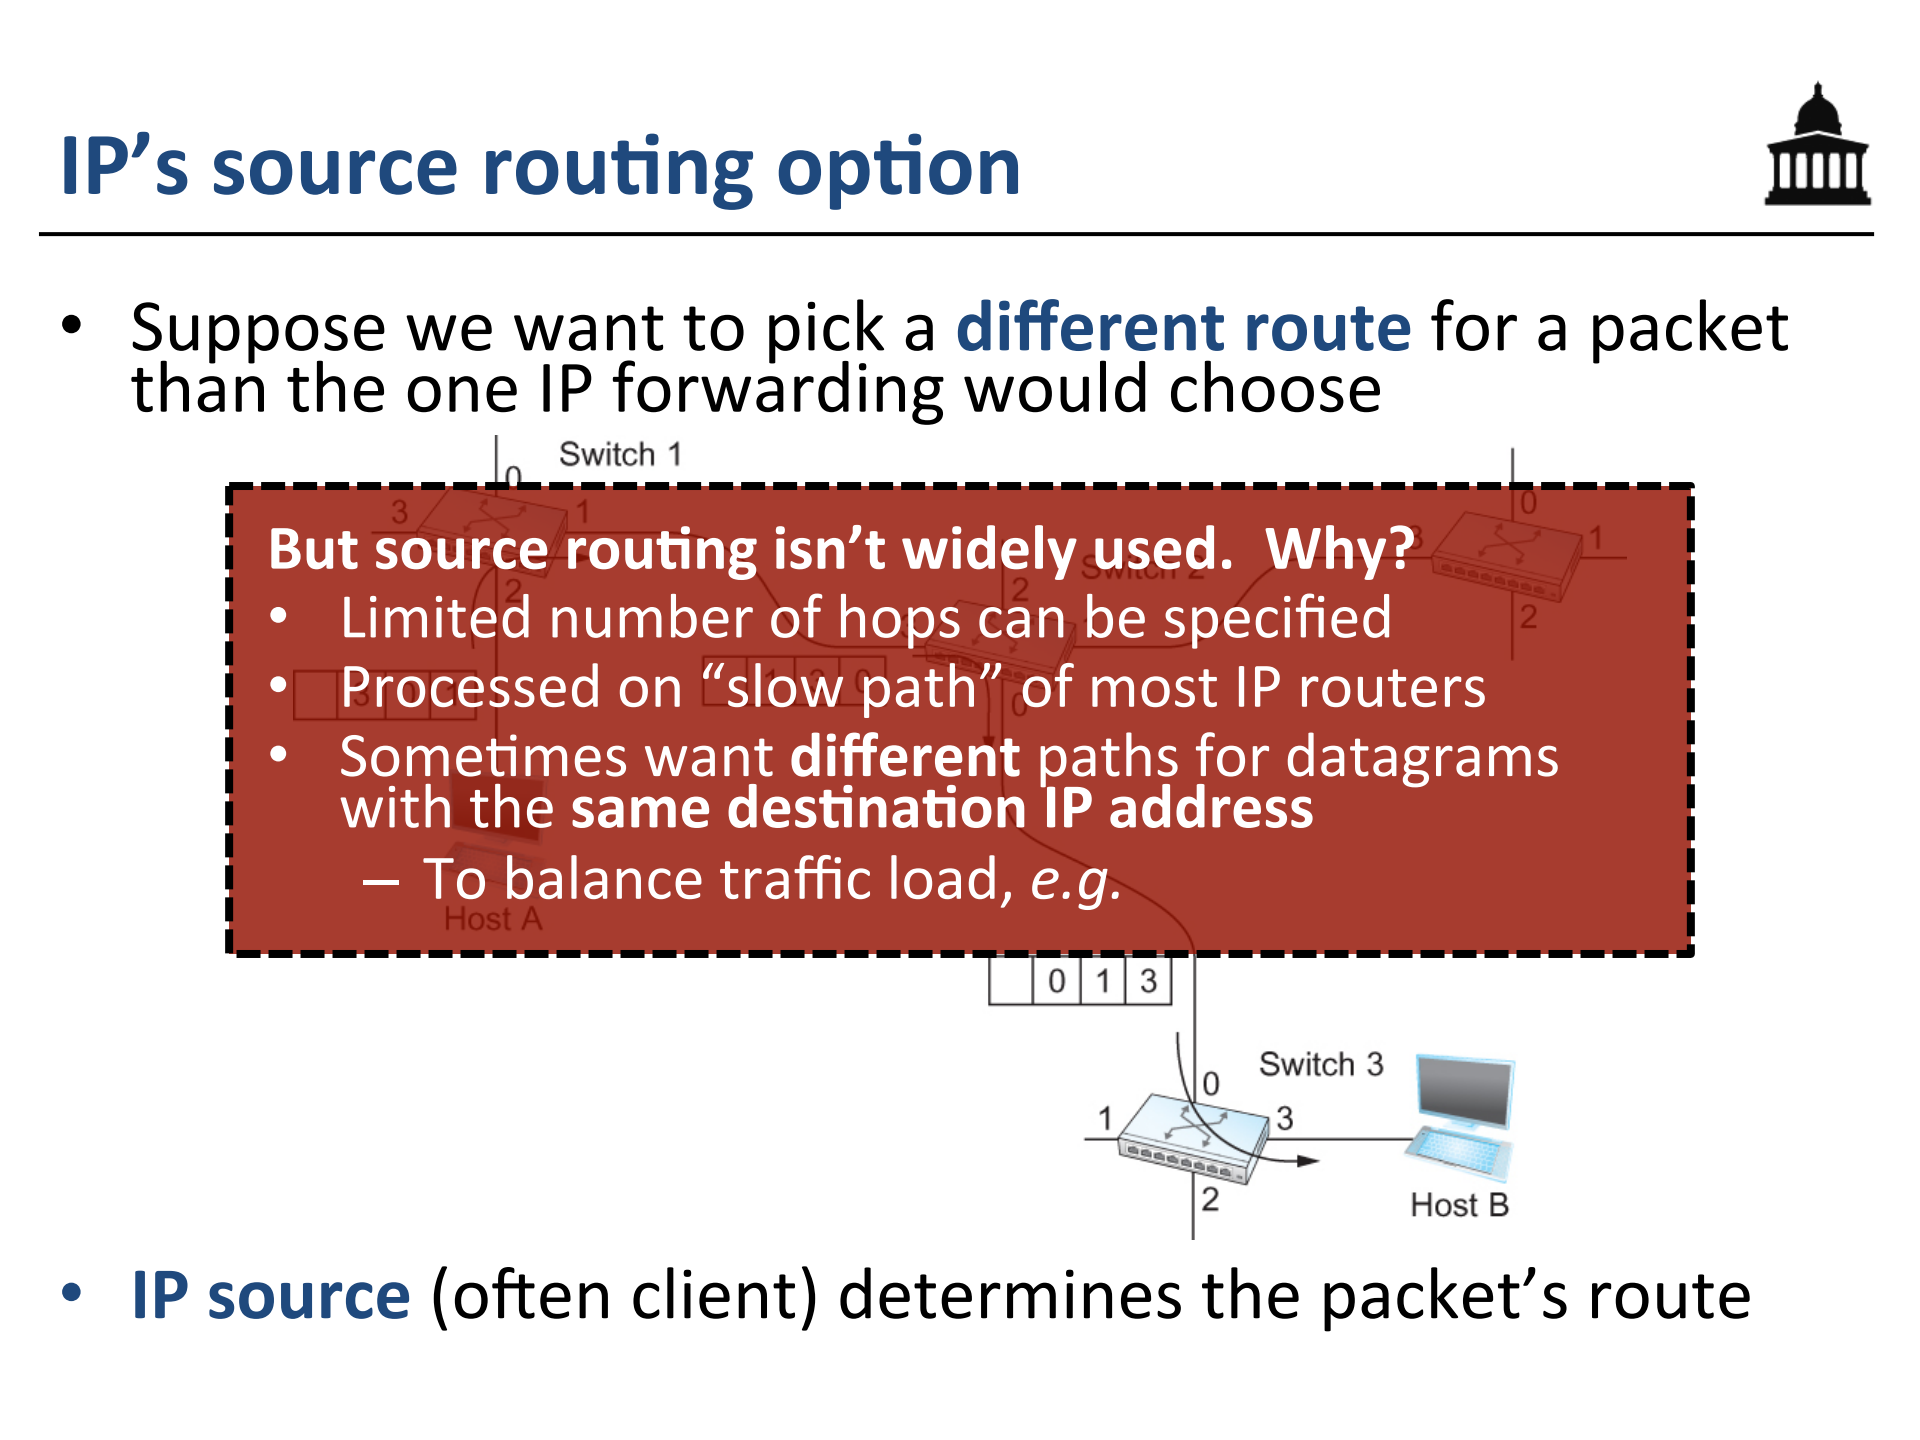

- Suppose we want to pick a **different route** for a packet than the one IP forwarding would choose

But source routing isn't widely used. Why?

- Limited number of hops can be specified
- Processed on "slow path" of most IP routers
- Sometimes want **different** paths for datagrams with the **same destination IP address**
  - To balance traffic load, *e.g.*

- **IP source** (often client) determines the packet's route

Explicit routing with MPLS

- Service provider's LER picks the route, not the IP source
- Suppose we want to load-balance R1 → R7 and R2 → R7 traffic
- Could IP routing handle this?
  - Not here: IP routing only looks at destination, not source
  - Flows from R1 and R2 both have destination R7
- **Solution:** Tag packets at R1, R2 with different MPLS labels
  - Threaded indices then accomplish the desired routing

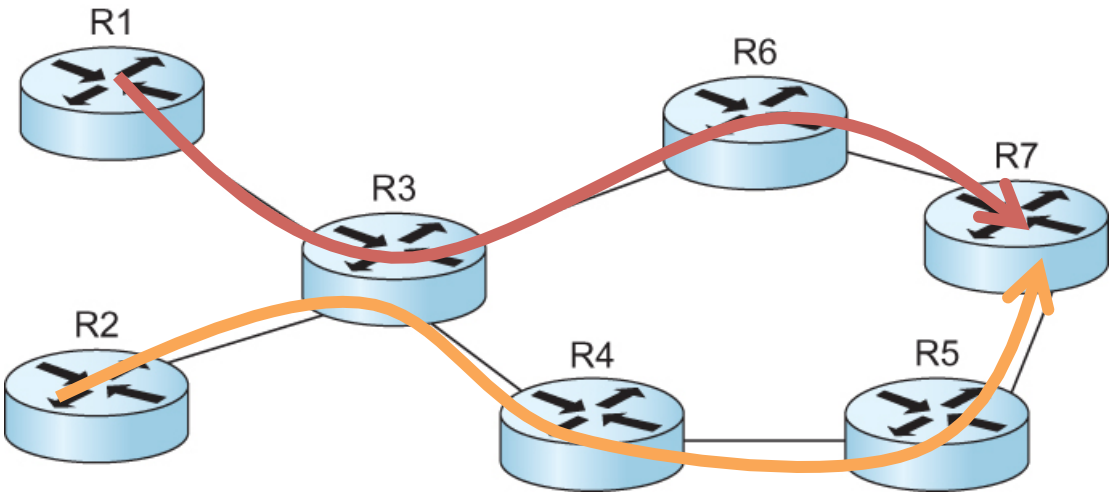

Today

- Three techniques that address IP's shortcomings:

1. MPLS

2. Virtual private networks

3. **Traffic engineering in the Internet**
   - MPLS explicit routing
   - **IP anycast**

Not unicast

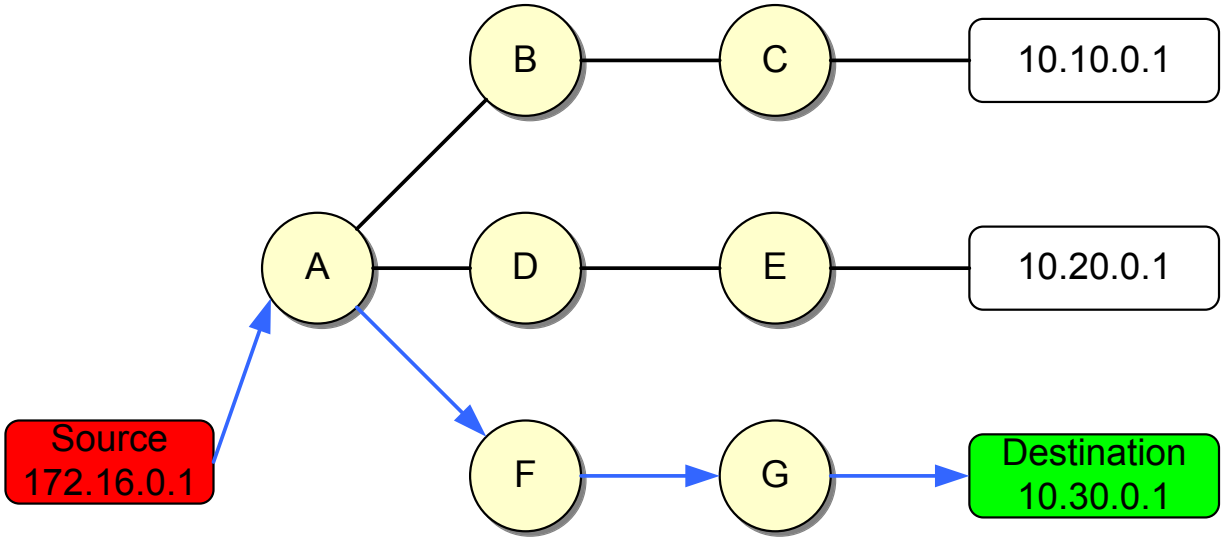

- ***Unicast:*** a single IP host receives all traffic

[Slide credit: CMU Network Group]

Not IP multicast

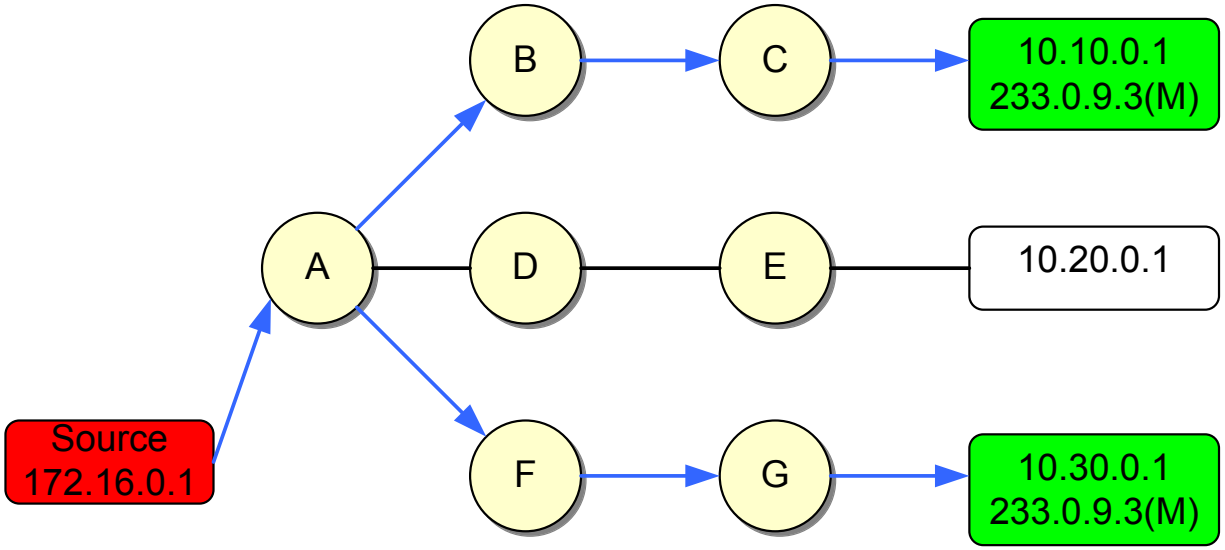

- IP multicast: **Many** hosts receive **all** traffic to a number of hosts (a ***multicast group***)

[Slide credit: CMU Network Group]

IP anycast

- **Multiple** hosts are configured to accept traffic on a **single** IP address

- Usually, just **one host** receives each datagram
  - Datagram **can be dropped** like any other (best effort)
  - Preferably only one node receives packet, but there are no absolute guarantees

- The host that receives a specific datagram is determined by the underlying Internet routing

IP anycast

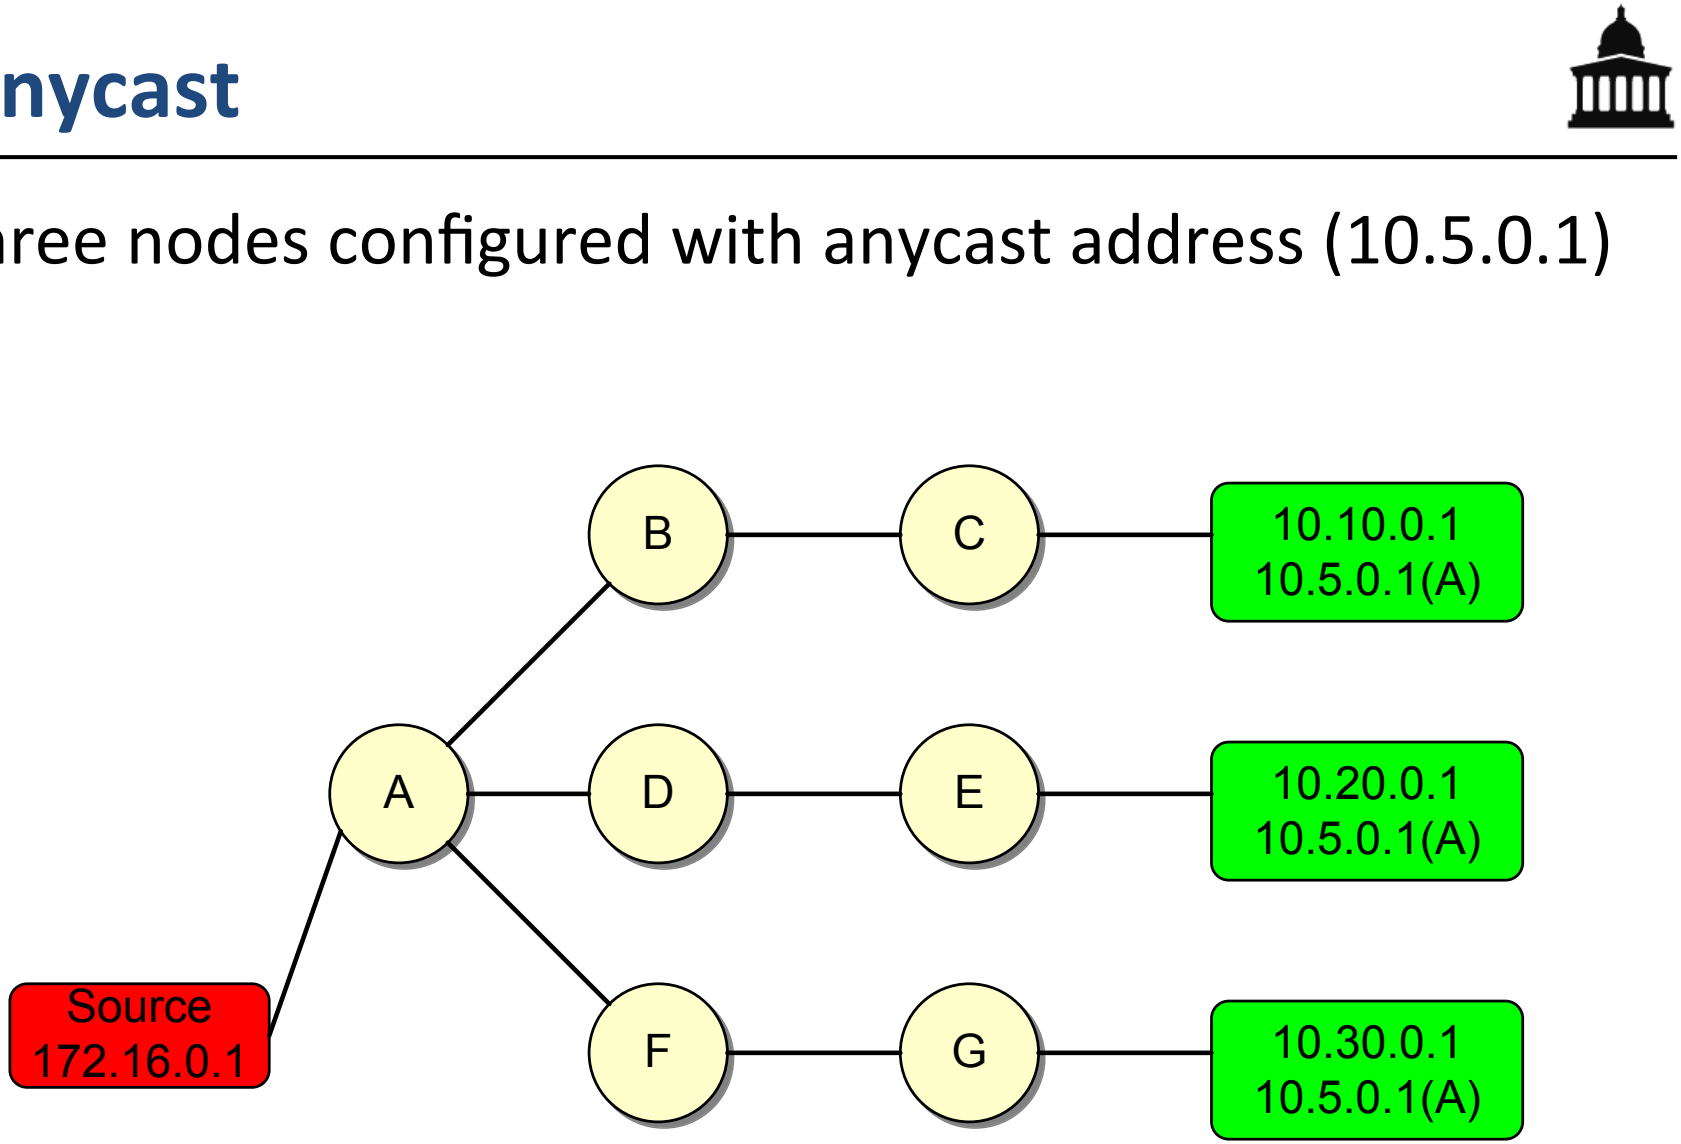

- Three nodes configured with anycast address (10.5.0.1)

[Slide credit: CMU Network Group]

IP anycast

- Paths to different destinations have **equal cost metrics** in A's routing table, so A picks just one next hop

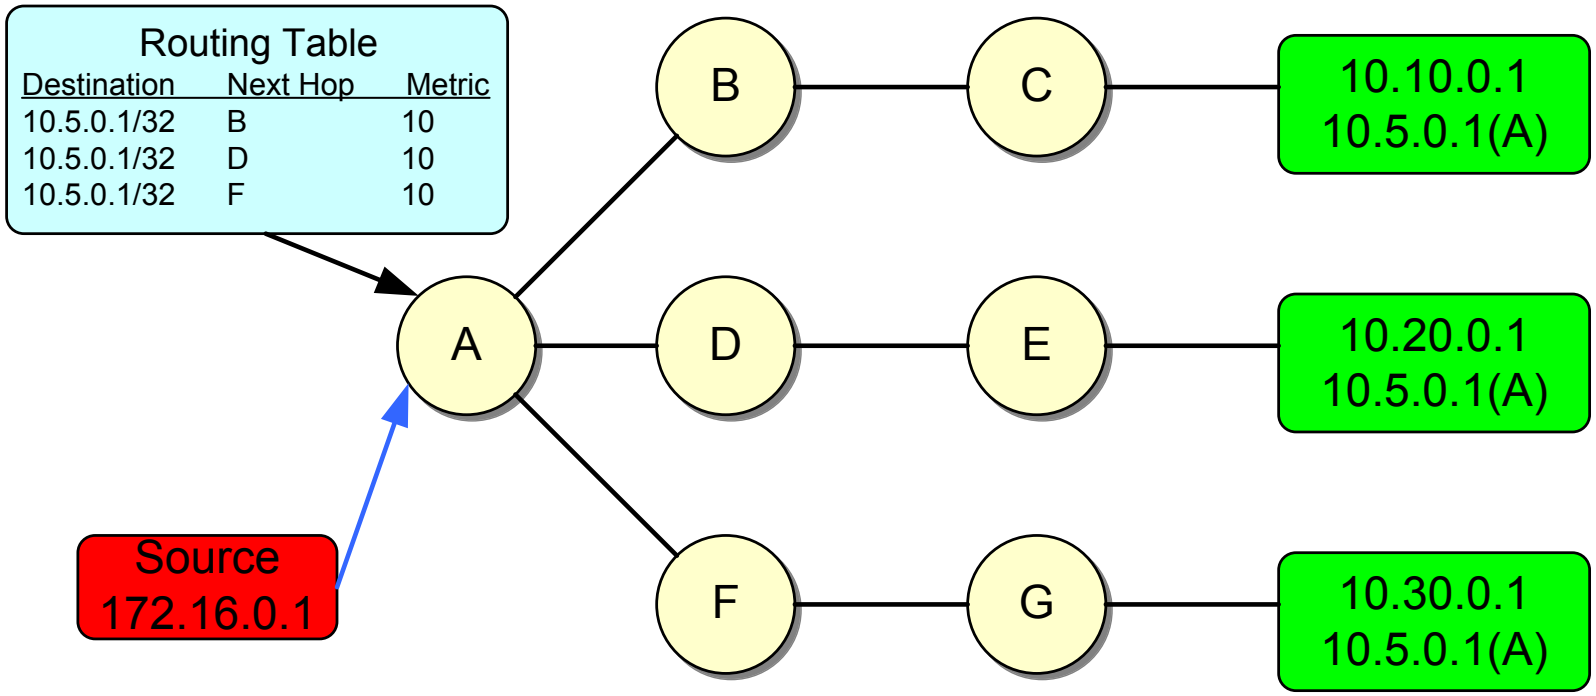

[Slide credit: CMU Network Group]

IP anycast

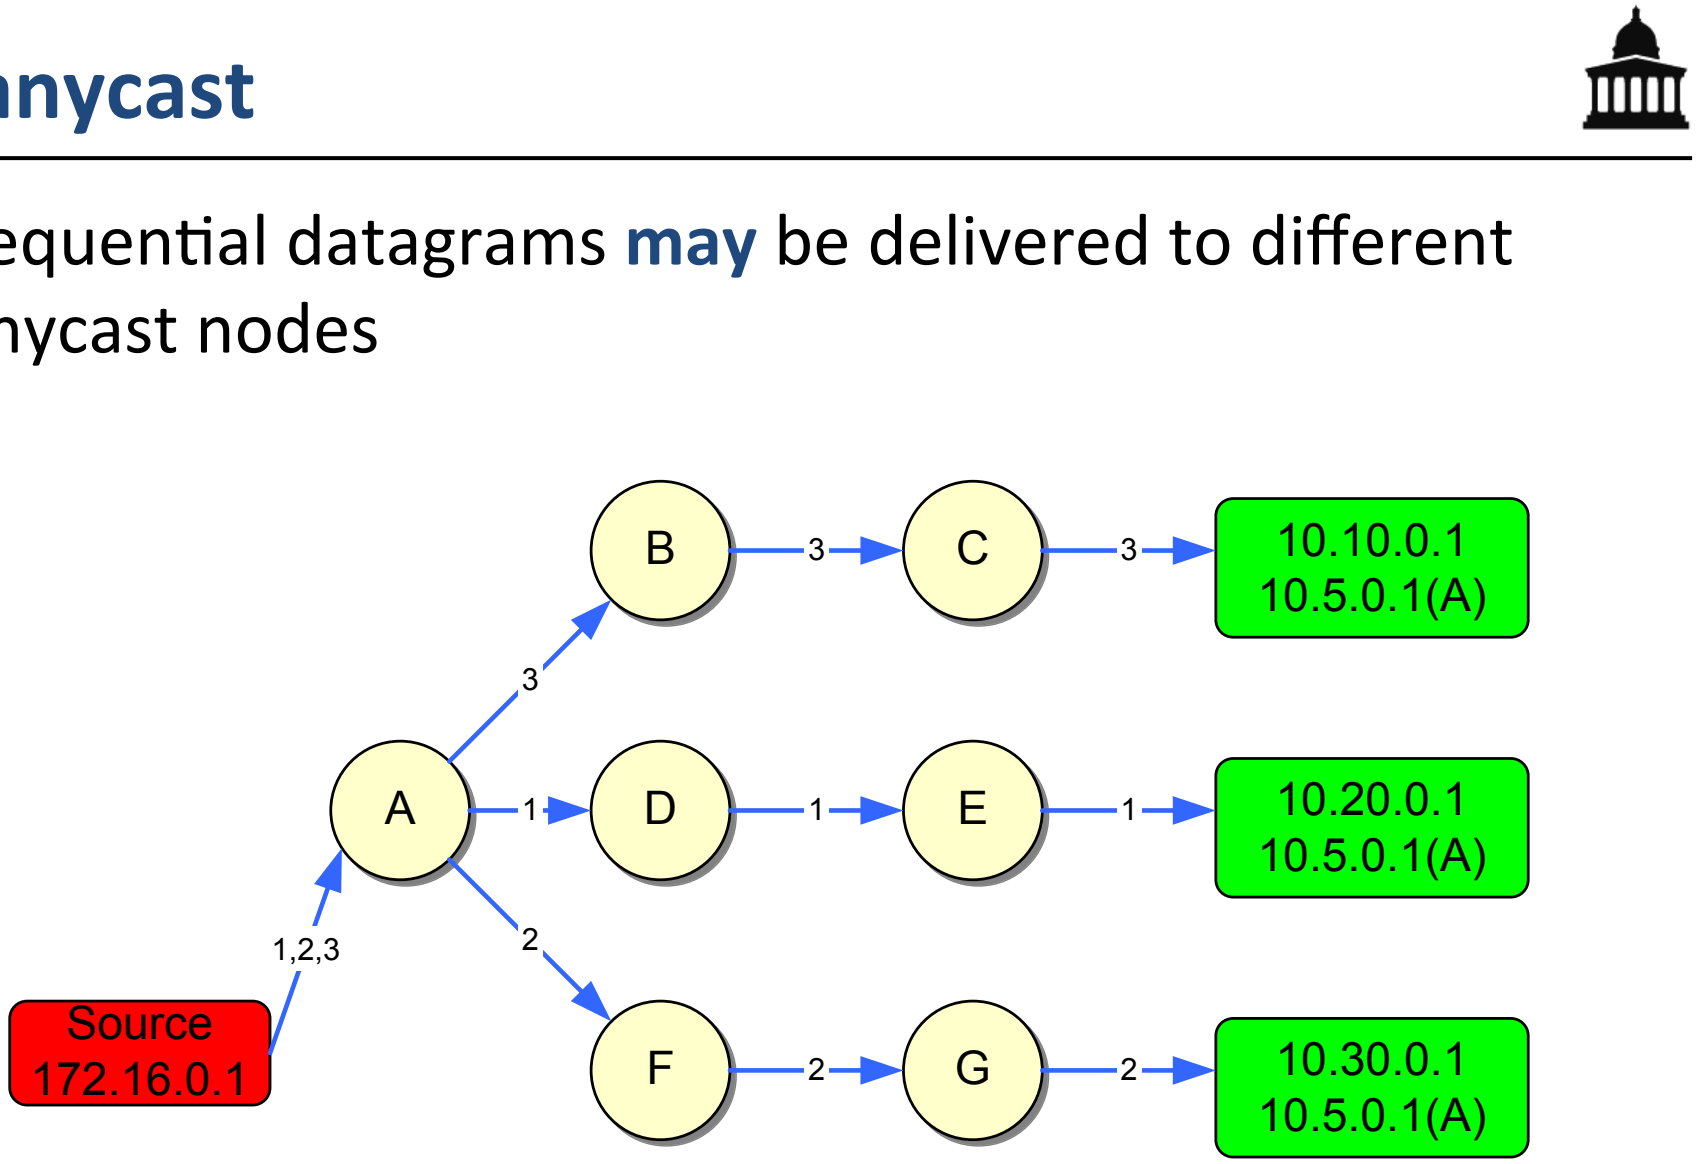

- Sequential datagrams **may** be delivered to different anycast nodes

[Slide credit: CMU Network Group]

IP anycast

- Traffic from different immediately-preceding hops may follow separate paths

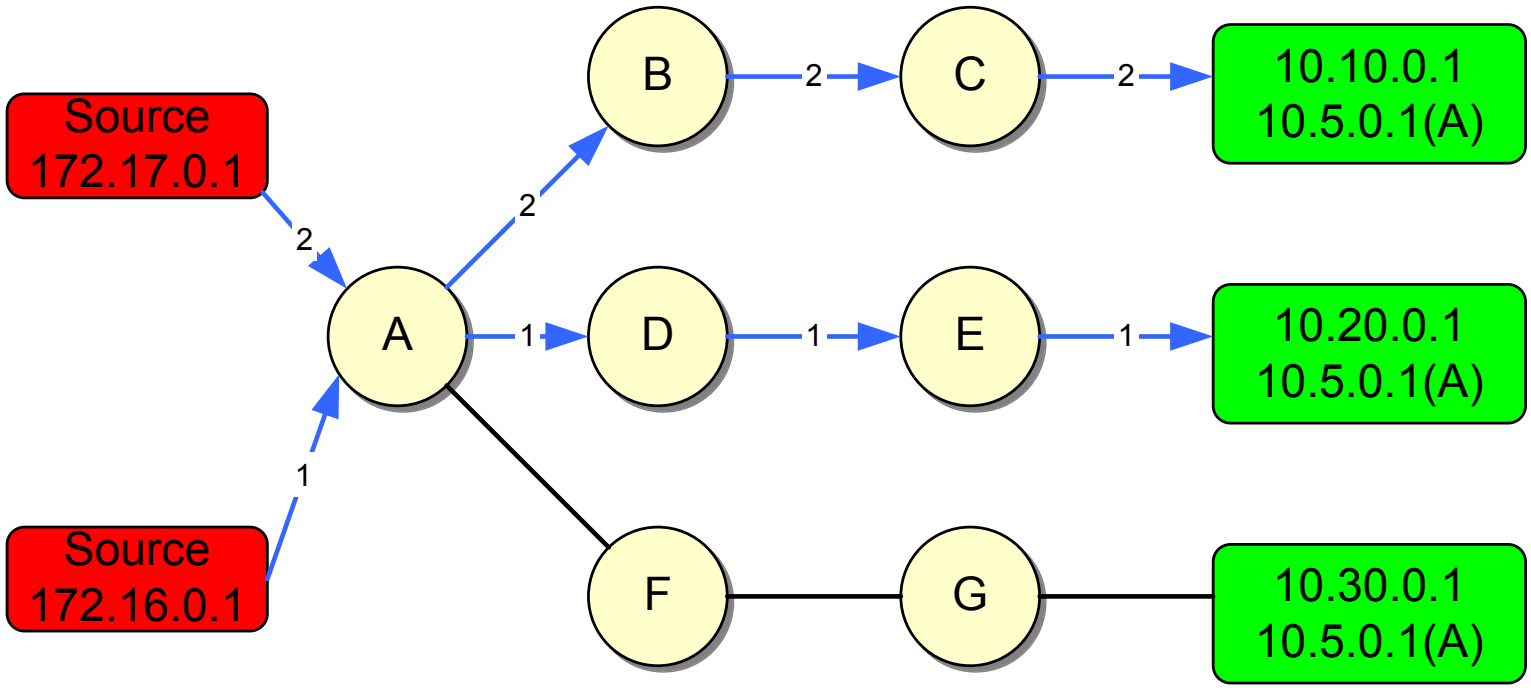

[Slide credit: CMU Network Group]

IP anycast

- **Server receiving a packet is determined by unicast routing**
- Sequential packets from a client to an anycast address may be delivered to different servers
- Best used for single request/response type protocols

[Slide credit: CMU Network Group]