

cs/ee/ids 143 Communication Networks

Chapter 4 Internetworking

Text: Walrand & Parekh, 2010

Steven Low
CMS, EE, Caltech





Warning

These notes are **not** self-contained,
probably **not** understandable,
unless you also were in the lecture

They are supplement to, not replacement for class attendance



Recap: Internet overview

Some basic mechanisms

- Packet switching
- Addressing
- Routing
 - hierarchical (AS), forwarding, shortest path routing, software defined networking
- Transport
 - congestion control, error recovery
- Medium access control
- Internetworking

Project



Recap: Internet overview

Some basic concepts

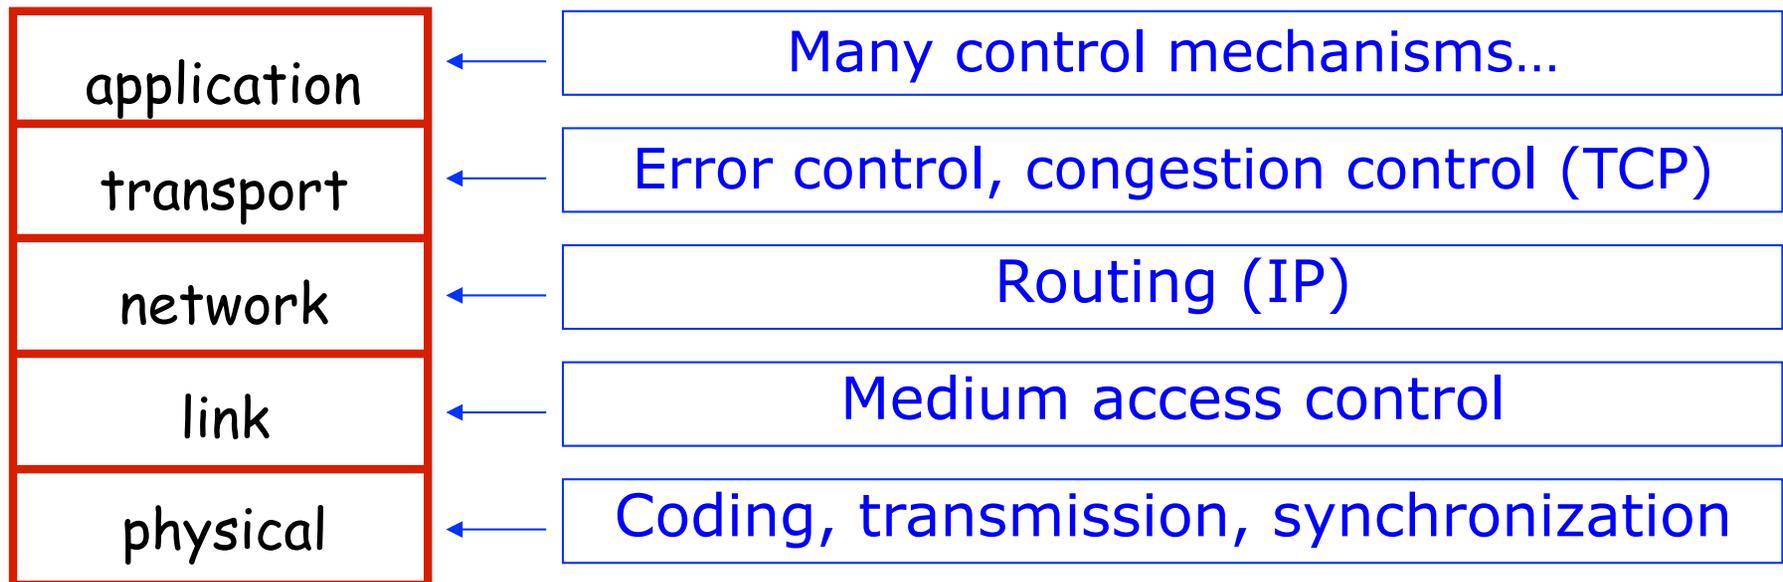
- Performance metrics
 - Throughput, line rate (bandwidth), line capacity
 - Delay, delay jitter
- Scalability
 - location-based routing, hierarchical
 - best-effort service, end-to-end principle
- Layering



Protocol stack

Network mechanisms implemented as protocol stack

Each layer designed separately, evolves asynchronously





Recap: Internet overview

Some basic analytic tools

- Convex optimization
 - We will use it to understand equilibrium properties of TCP congestion control
- Control and dynamical system
 - We will use it to understand stability properties of TCP congestion control
- Queueing theory
 - We will use it to understand statistical properties of wireless MAC



Recap: Routing

Covered layer 3 routing

- Autonomous systems (AS)
 - Defined by administrative domains
- Inter-AS: BGP
 - Policy based
- Intra-AS: Dijkstra, Bellman-Ford
 - Shortest-path routing

Error recovery

- Can be used in link, transport, or application layer
- Parity check, FEC, network coding



This week

Internetworking

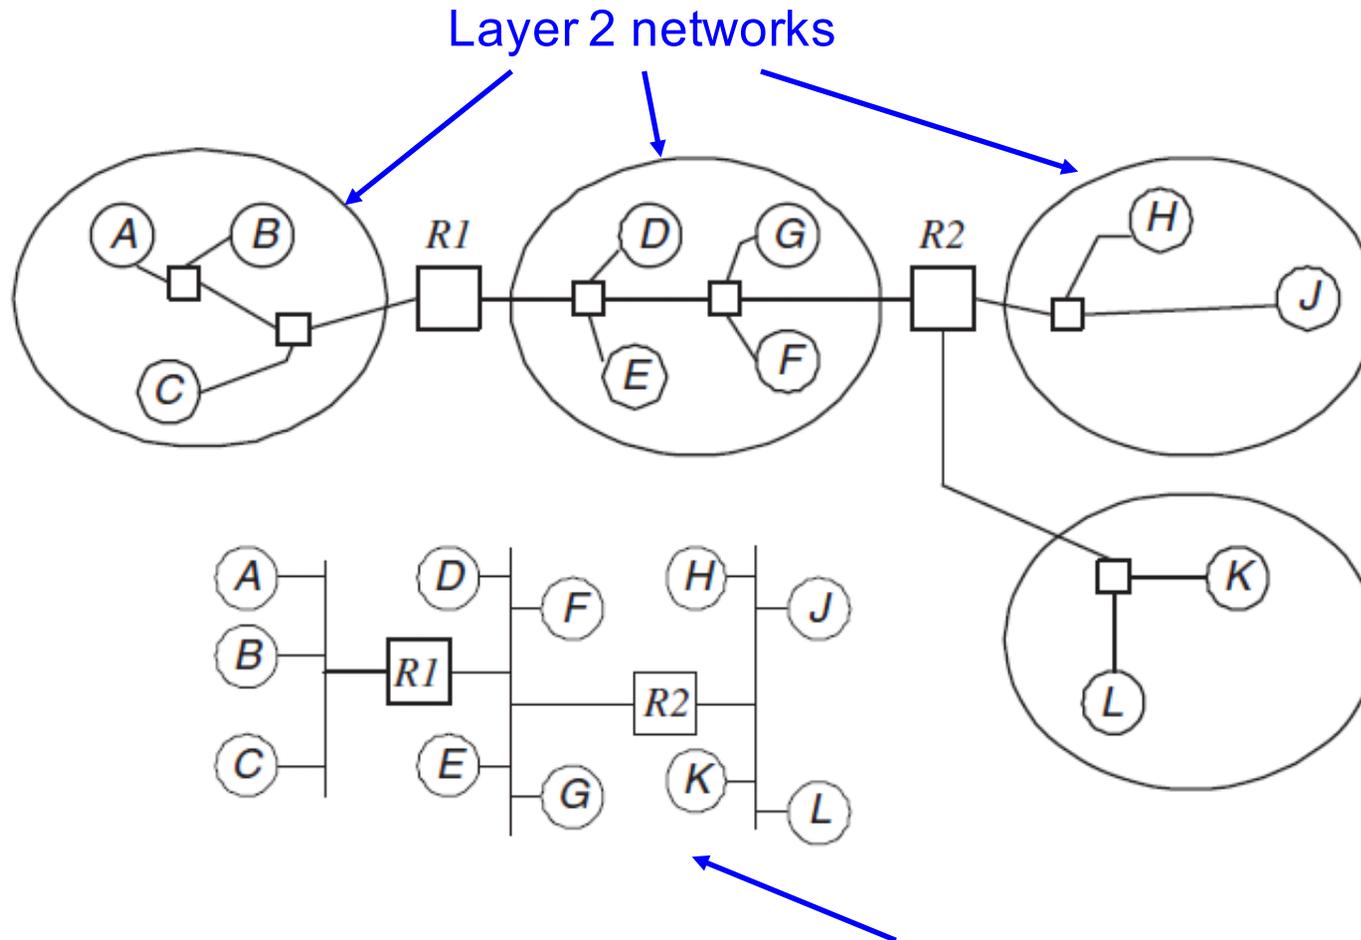
- Routing across LANs, layer2-layer3
- DHCP
- NAT

Transport layer

- Connection setup
- Error recovery: retransmission
- Congestion control



Internetworks



Layer 3 view: ignore the details of subnetworks

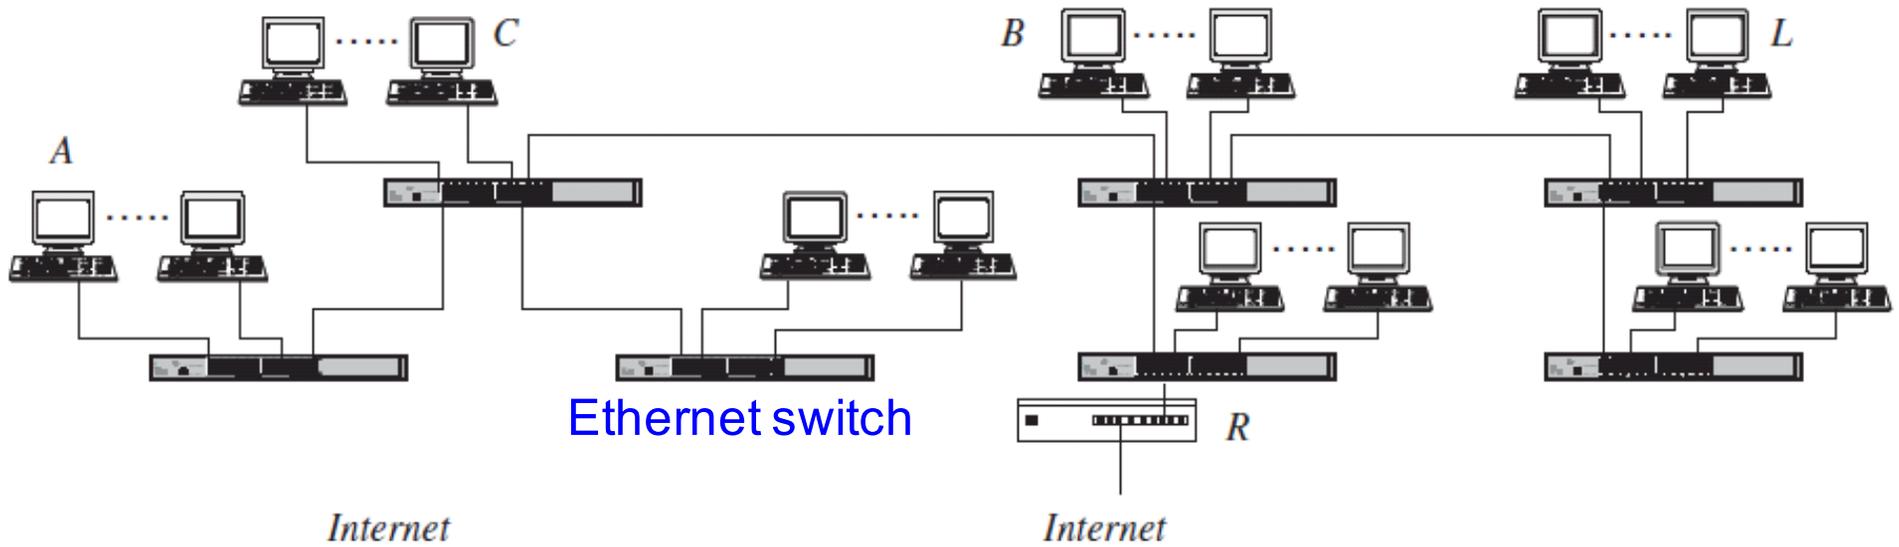
Each (layer 2) network

- **full connectivity**: every node can reach every other node
- **broadcast capable**: every node (inc. router) can broadcast to all other nodes



Internetworks

an Ethernet network (layer 2)



Ethernet switch



Layer 3 view



Addresses, Subnet mask, Gateway

Network Connection Details

Network Connection Details:

Property	Value
Connection-specific DN...	gateway.2wire.net
Description	Intel(R) Wireless WiFi Link 4965AGN
Physical Address	00-1D-E0-00-C1-95
DHCP Enabled	Yes
IPv4 Address	192.168.1.108
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Tuesday, October 11, 2011 11:11:15 PM
Lease Expires	Wednesday, October 12, 2011 11:11:17 P
IPv4 Default Gateway	192.168.1.254
IPv4 DHCP Server	192.168.1.254
IPv4 DNS Server	192.168.1.254
IPv4 WINS Server	
NetBIOS over Tcip En...	Yes
Link-local IPv6 Address	fe80::3ca0:7431:7e6:261e%11
IPv6 Default Gateway	
IPv6 DNS Server	

Close

Network layer on a host knows:

1. Own IP address
2. Subnet mask
3. Gateway router IP
4. DNS server IP



Addresses, Subnet mask, Gateway

Network Connection Details

Network Connection Details:

Property	Value
Connection-specific DN...	gateway.2wire.net
Description	Intel(R) Wireless WiFi Link 4965AGN
Physical Address	00-1D-E0-00-C1-95
DHCP Enabled	Yes
IPv4 Address	192.168.1.108
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Tuesday, October 11, 2011 11:11:15 PM
Lease Expires	Wednesday, October 12, 2011 11:11:17 P
IPv4 Default Gateway	192.168.1.254
IPv4 DHCP Server	192.168.1.254
IPv4 DNS Server	192.168.1.254
IPv4 WINS Server	
NetBIOS over Tcip En...	Yes
Link-local IPv6 Address	fe80::3ca0:7431:7e6:261e%11
IPv6 Default Gateway	
IPv6 DNS Server	

Close

Network layer on a host knows:

1. Own IP address
2. Subnet mask
3. Gateway router IP
4. DNS server IP



Addresses, Subnet mask, Gateway

Network Connection Details

Network Connection Details:

Property	Value
Connection-specific DN...	gateway.2wire.net
Description	Intel(R) Wireless WiFi Link 4965AGN
Physical Address	00-1D-E0-00-C1-95
DHCP Enabled	Yes
IPv4 Address	192.168.1.108
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Tuesday, October 11, 2011 11:11:15 PM
Lease Expires	Wednesday, October 12, 2011 11:11:17 P
IPv4 Default Gateway	192.168.1.254
IPv4 DHCP Server	192.168.1.254
IPv4 DNS Server	192.168.1.254
IPv4 WINS Server	
NetBIOS over Tcip En...	Yes
Link-local IPv6 Address	fe80::3ca0:7431:7e6:261e%11
IPv6 Default Gateway	
IPv6 DNS Server	

Close

Network layer on a host knows:

1. Own IP address
2. Subnet mask
3. Gateway router IP
4. DNS server IP

Tells the host the range of IP addresses in the same Layer 2 subnet



Addresses, Subnet mask, Gateway

Network Connection Details

Network Connection Details:

Property	Value
Connection-specific DN...	gateway.2wire.net
Description	Intel(R) Wireless WiFi Link 4965AGN
Physical Address	00-1D-E0-00-C1-95
DHCP Enabled	Yes
IPv4 Address	192.168.1.108
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Tuesday, October 11, 2011 11:11:15 PM
Lease Expires	Wednesday, October 12, 2011 11:11:17 P
IPv4 Default Gateway	192.168.1.254
IPv4 DHCP Server	192.168.1.254
IPv4 DNS Server	192.168.1.254
IPv4 WINS Server	
NetBIOS over Tcip En...	Yes
Link-local IPv6 Address	fe80::3ca0:7431:7e6:261e%11
IPv6 Default Gateway	
IPv6 DNS Server	

Close

Network layer on a host knows:

1. Own IP address
2. Subnet mask
3. Gateway router IP
4. DNS server IP

Router (interface) that connects host to the rest of the internet



Addresses, Subnet mask, Gateway

Network Connection Details

Network Connection Details:

Property	Value
Connection-specific DN...	gateway.2wire.net
Description	Intel(R) Wireless WiFi Link 4965AGN
Physical Address	00-1D-E0-00-C1-95
DHCP Enabled	Yes
IPv4 Address	192.168.1.108
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Tuesday, October 11, 2011 11:11:15 PM
Lease Expires	Wednesday, October 12, 2011 11:11:17 P
IPv4 Default Gateway	192.168.1.254
IPv4 DHCP Server	192.168.1.254
IPv4 DNS Server	192.168.1.254
IPv4 WINS Server	
NetBIOS over Tcip En...	Yes
Link-local IPv6 Address	fe80::3ca0:7431:7e6:261e%11
IPv6 Default Gateway	
IPv6 DNS Server	

Close

Network layer on a host knows:

1. Own IP address
2. Subnet mask
3. Gateway router IP
4. DNS server IP

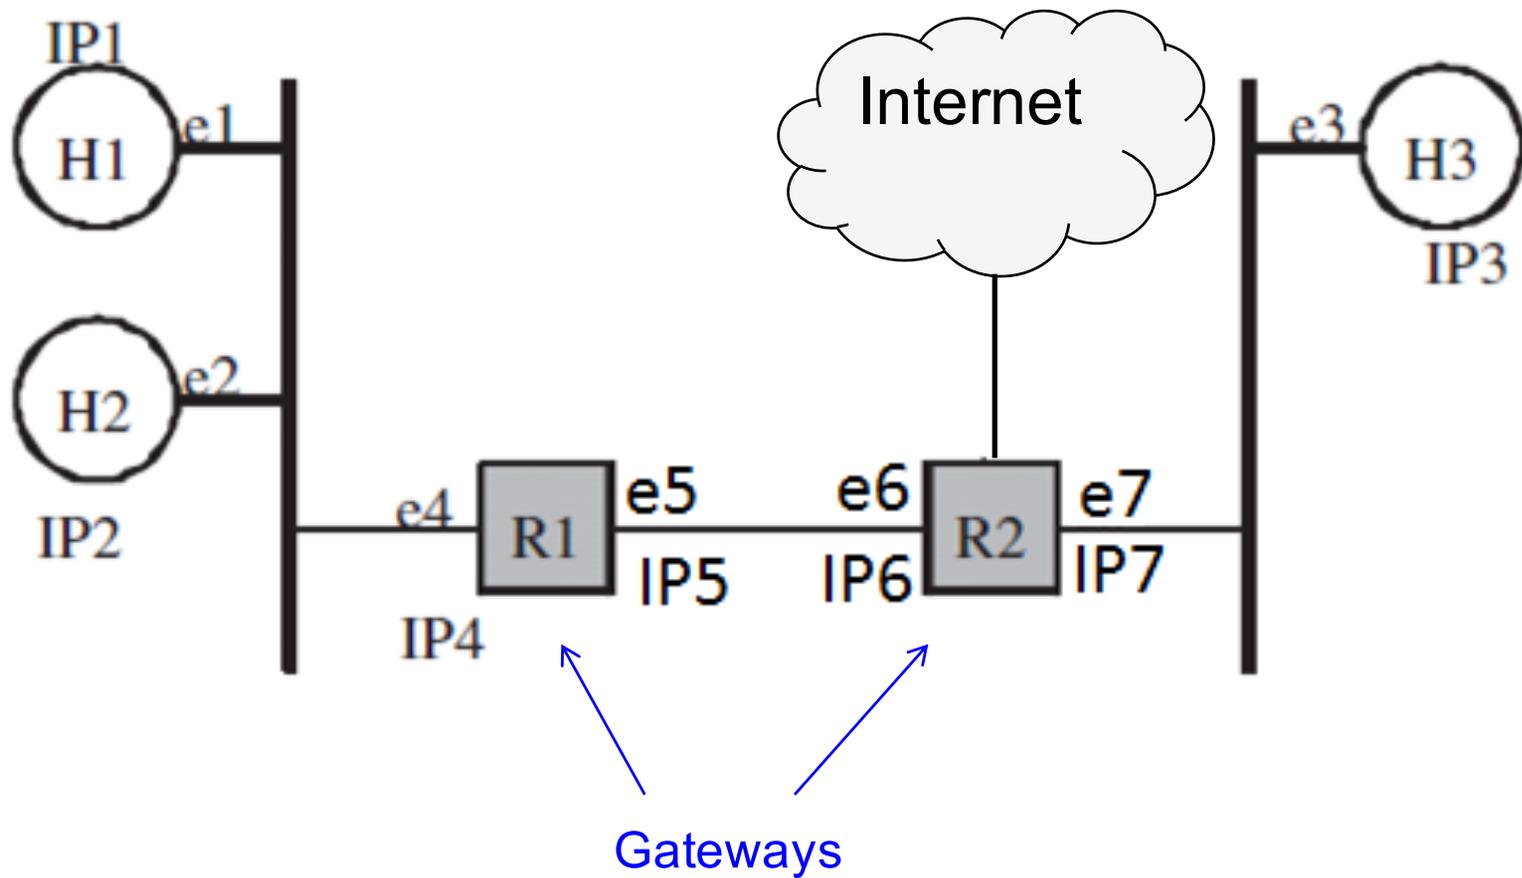
This server returns IP addresses for a domain name



Example

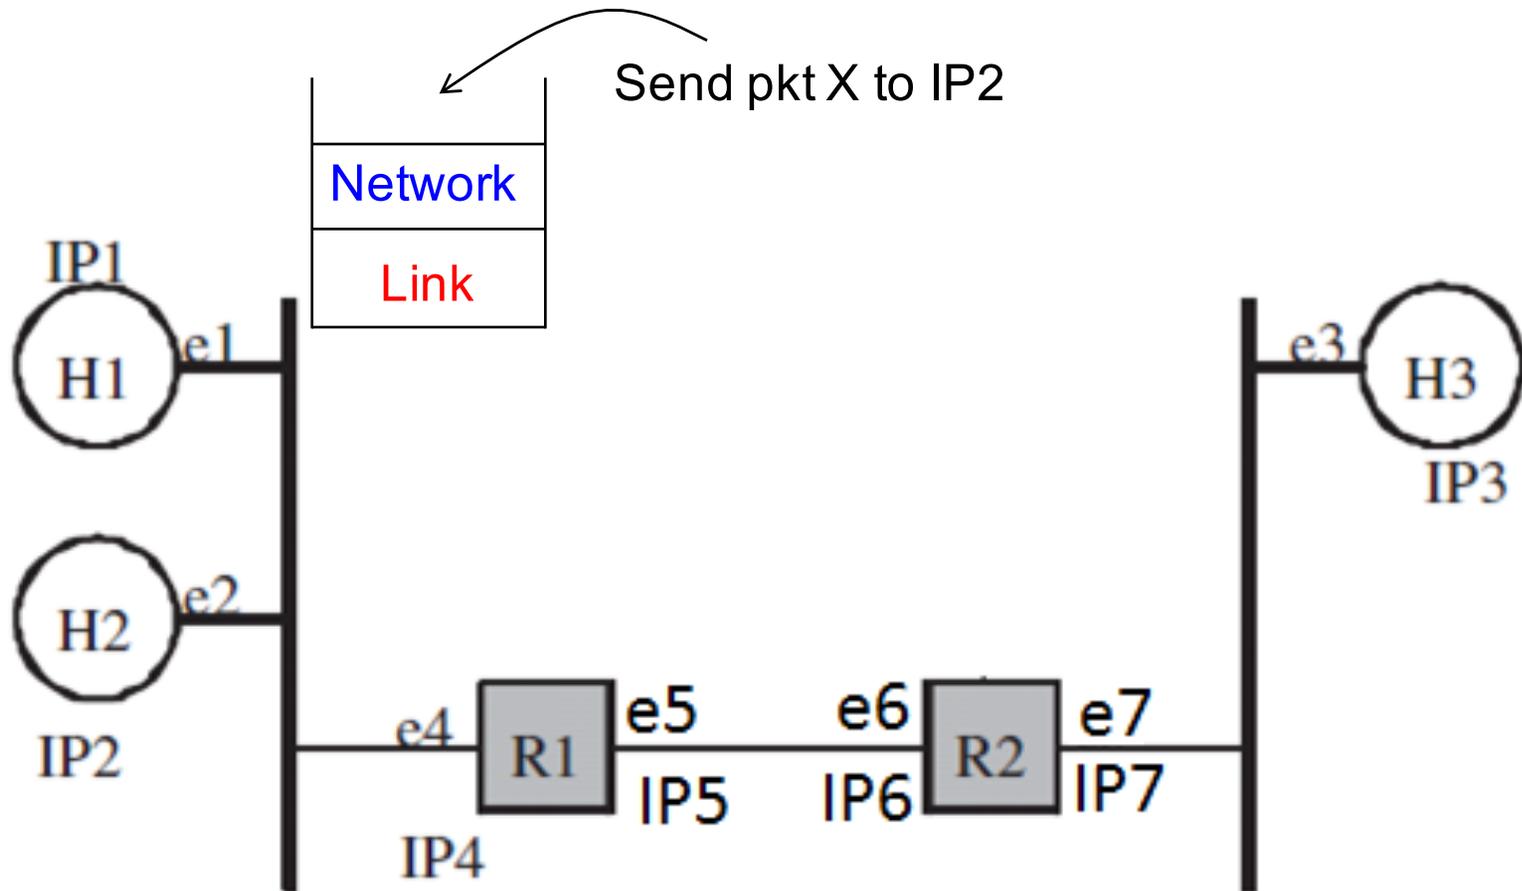
Ethernet sub-network 1

Ethernet sub-network 2





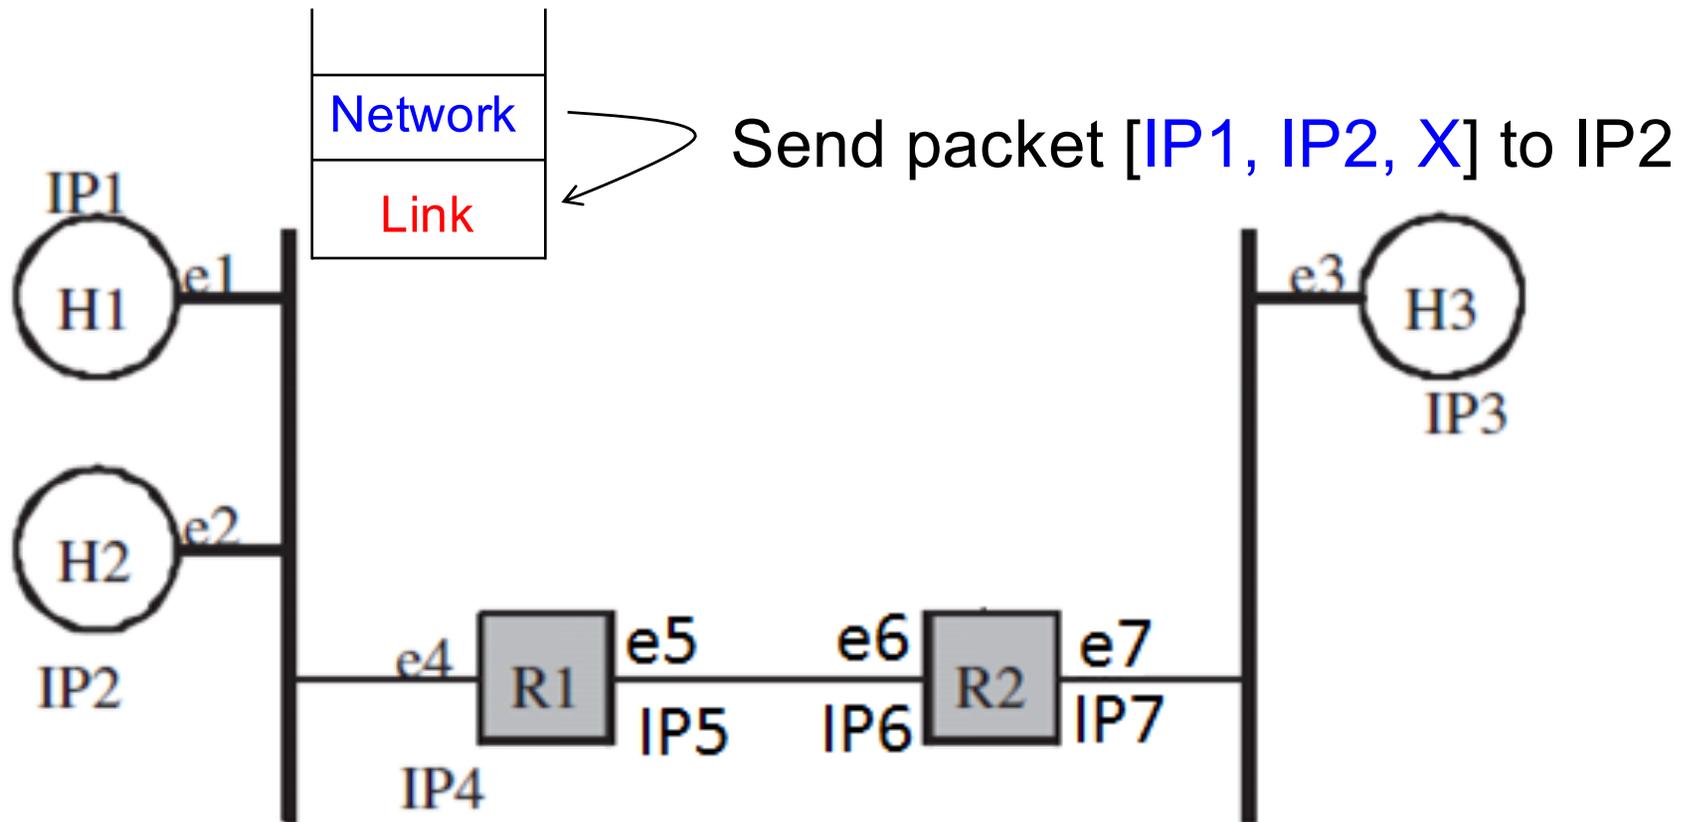
Example: H1 wants to send packet to H2



Network layer on H1 uses subset mask to determine that IP2 is within its subnet => next (layer 3) hop is IP2



Example: H1 wants to send packet to H2

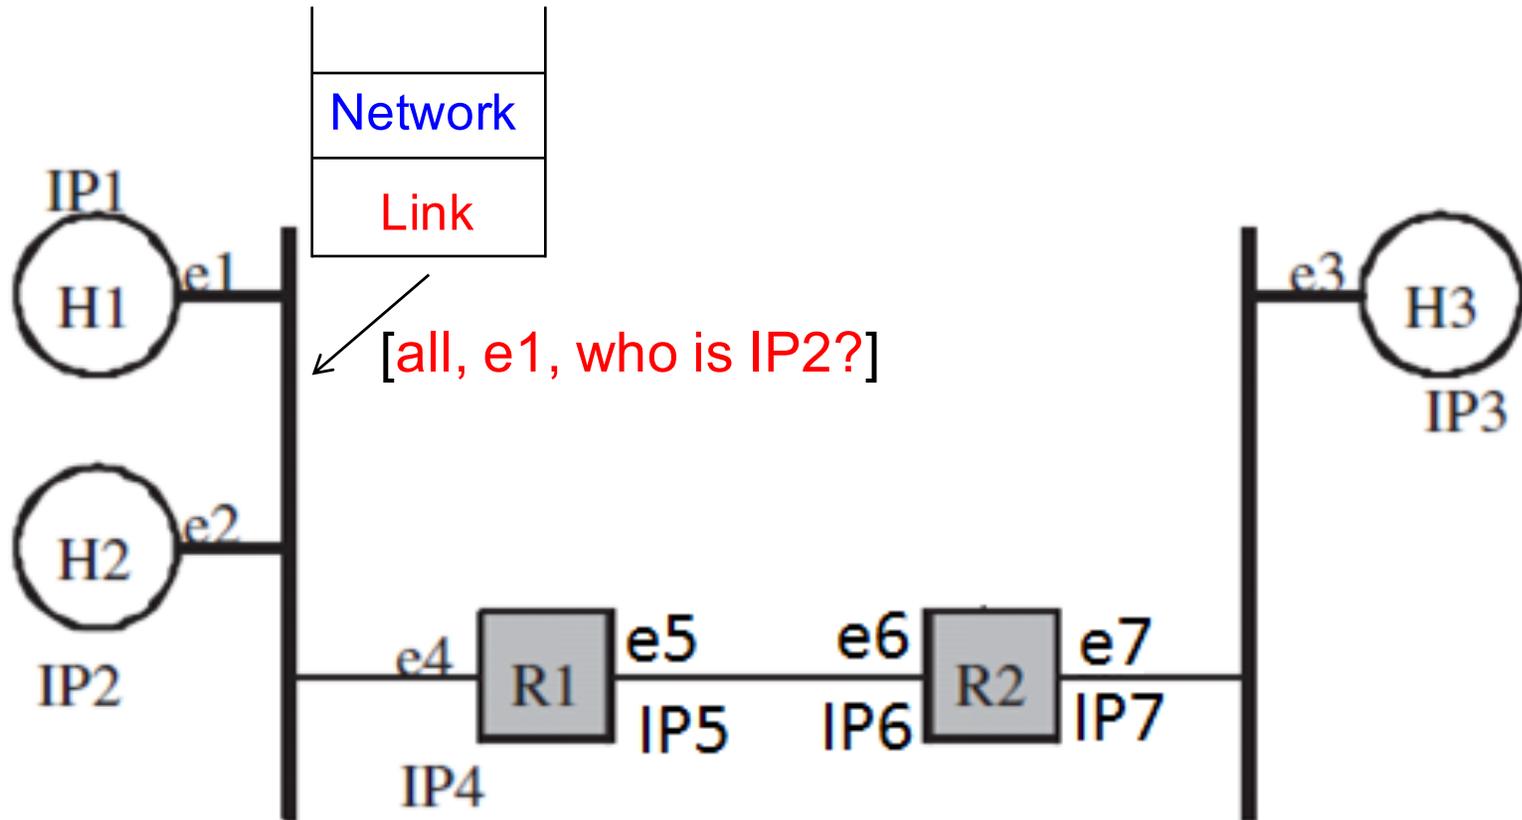


Link layer on H1 needs the MAC address e2 corresponding to IP2 to deliver the packet:

ARP (Address Resolution protocol)



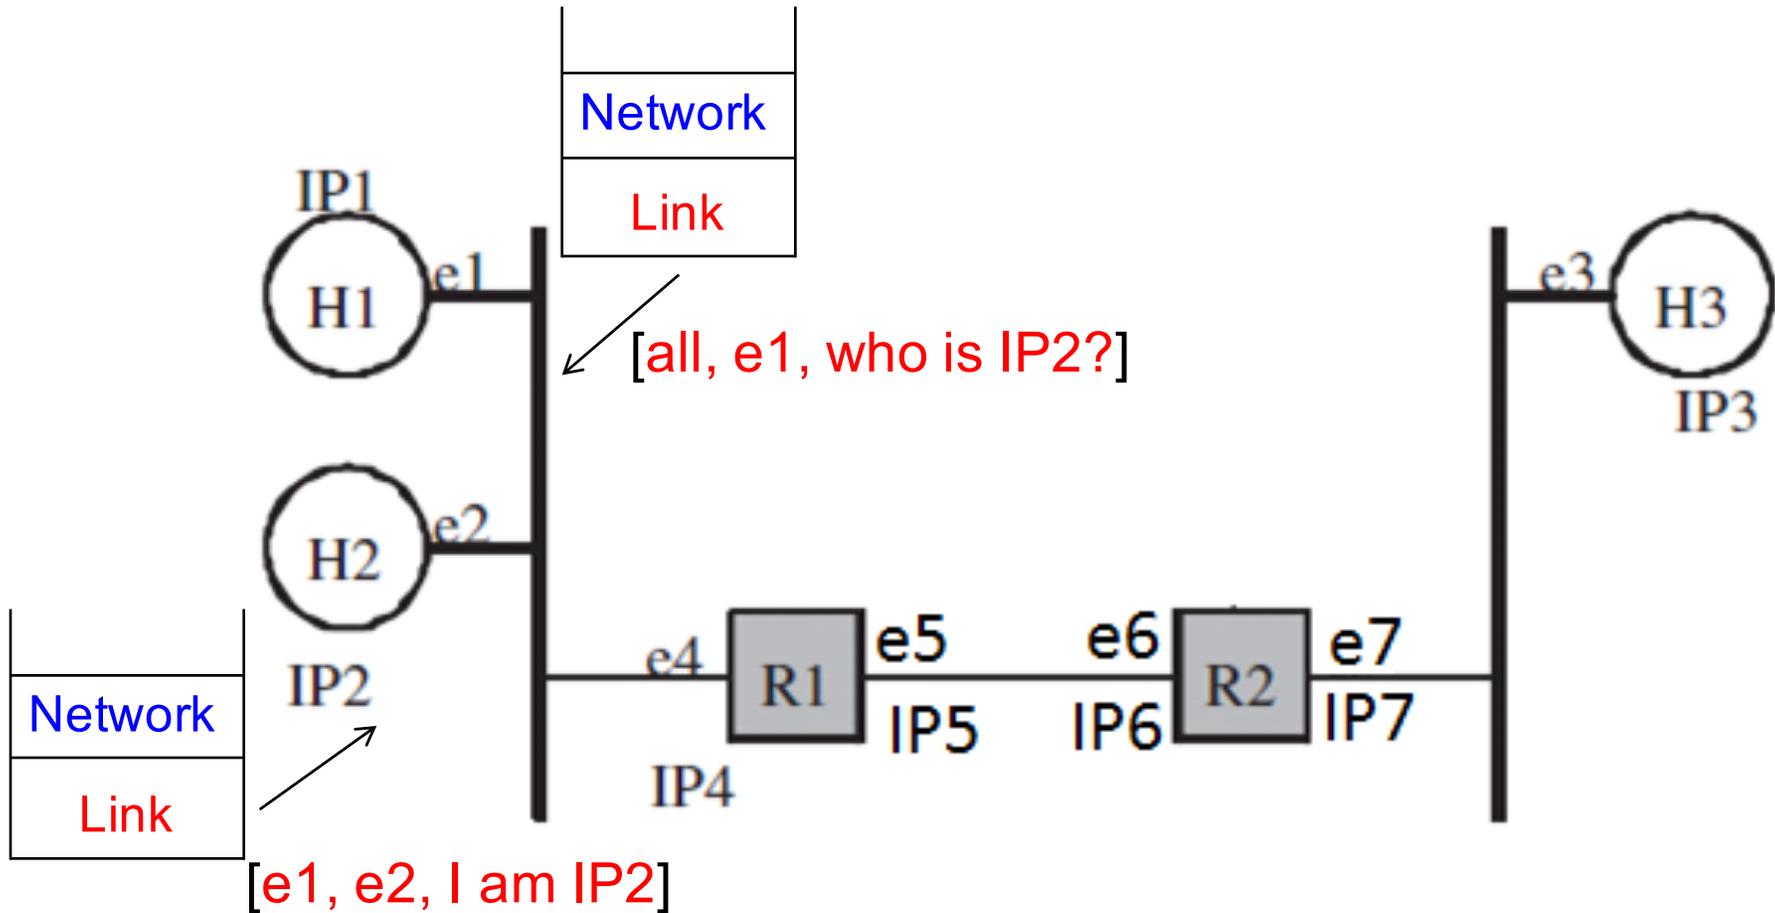
Example: H1 wants to send packet to H2



Link layer on H1 broadcasts a message (ARP query) on its layer 2 network asking for the MAC address corresponding to IP2



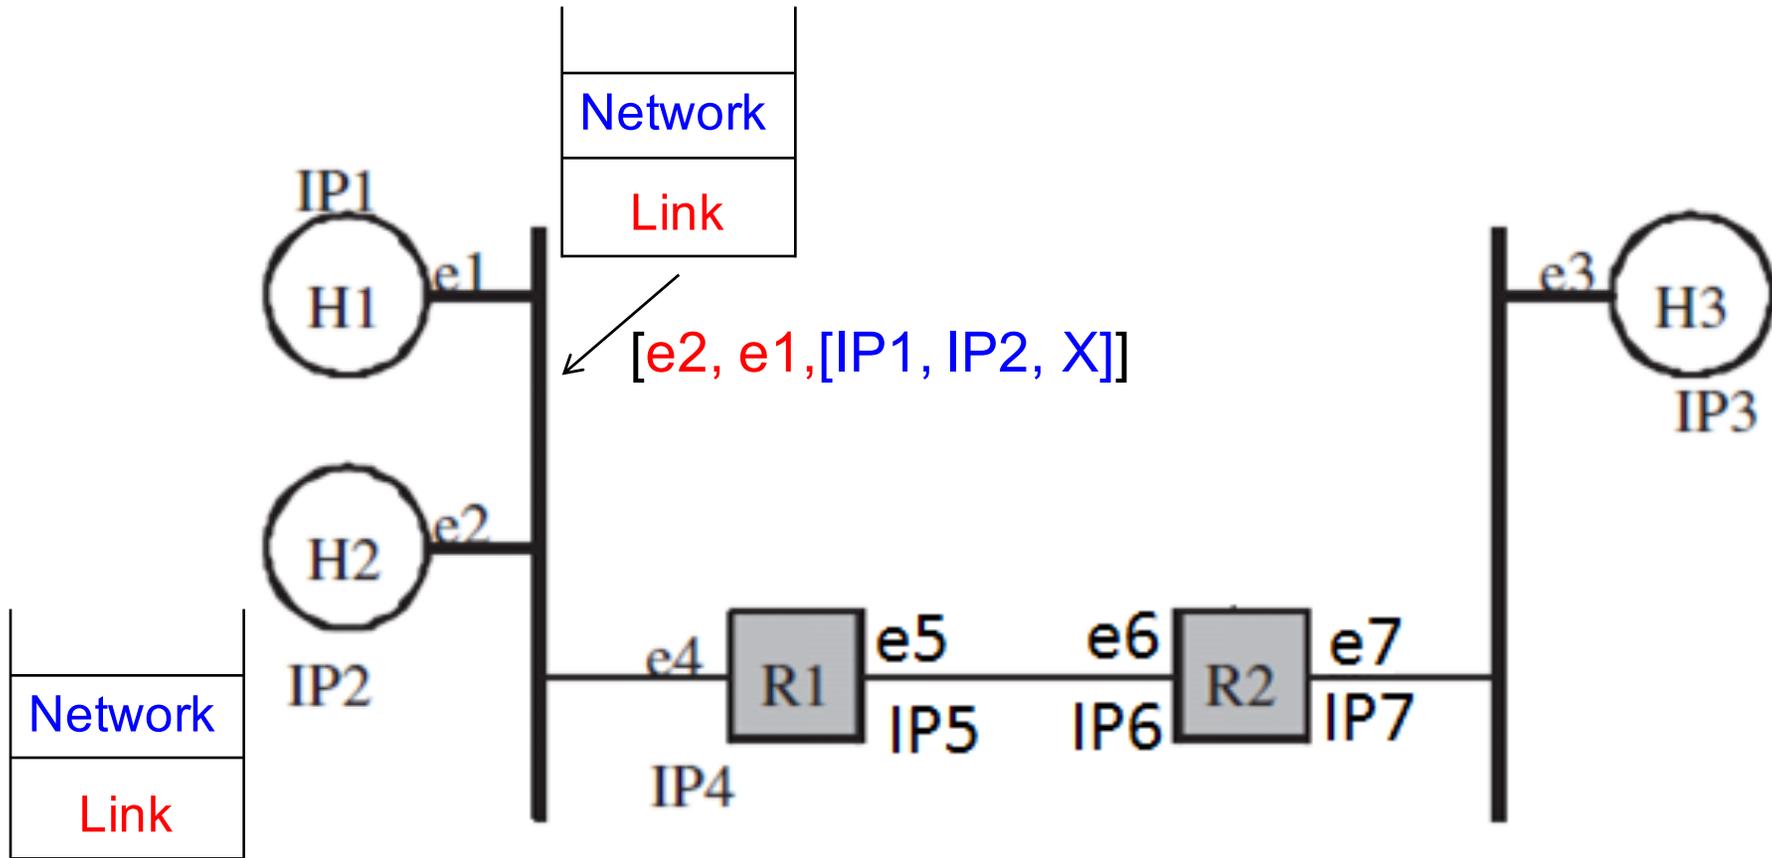
Example: H1 wants to send packet to H2



Link layer on H2 responds to the ARP query with its MAC address



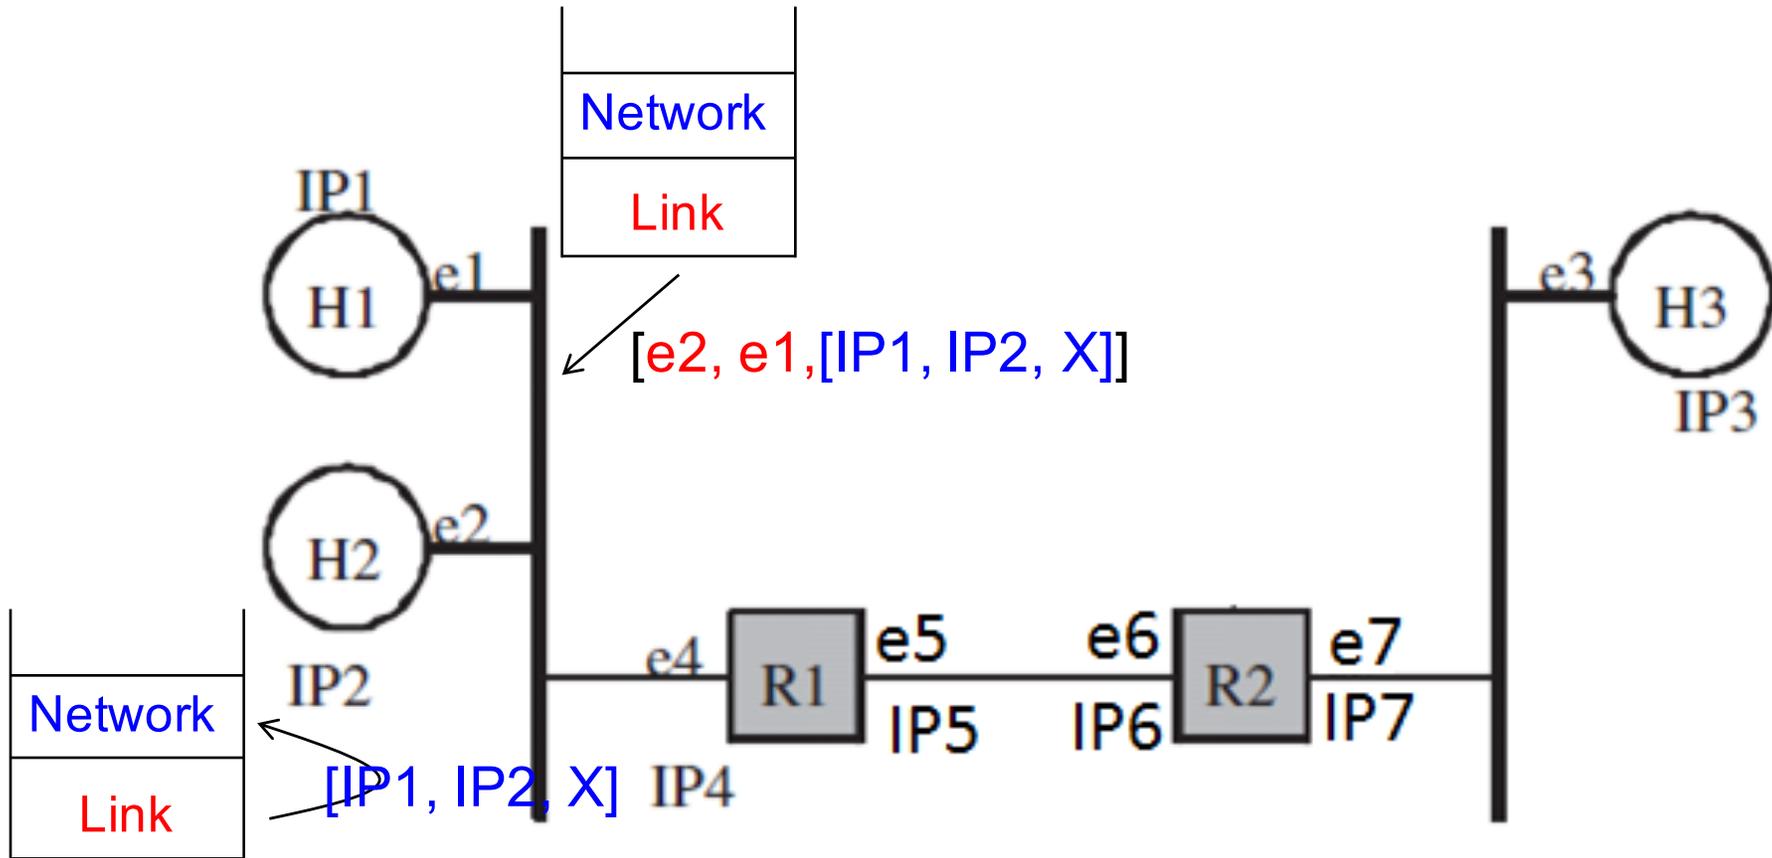
Example: H1 wants to send packet to H2



Once the link layer on H1 knows e2, it can now send the original message



Example: H1 wants to send packet to H2

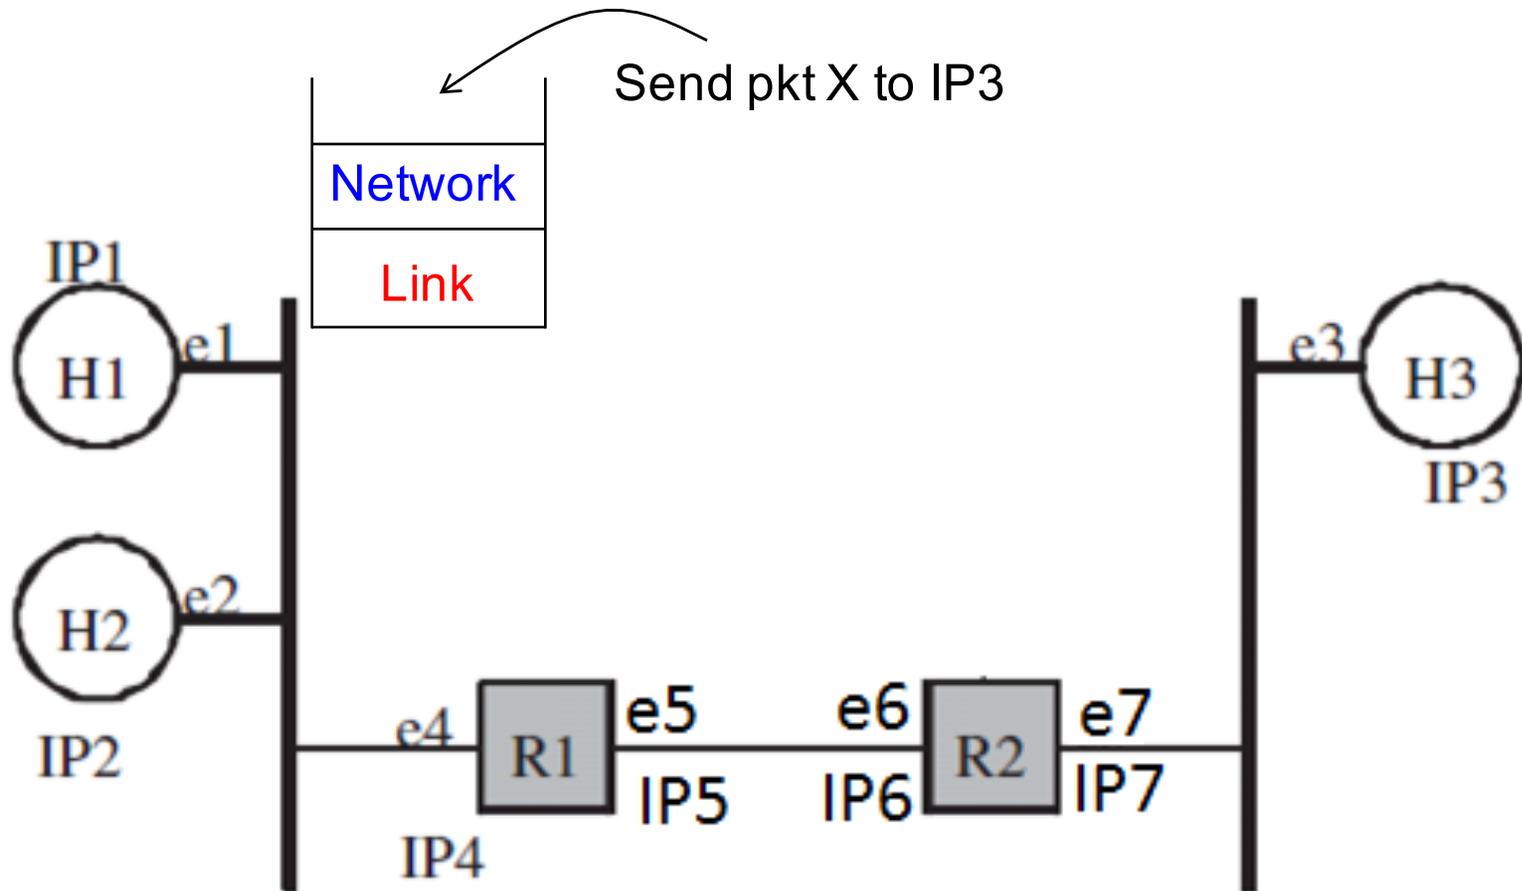


$[e2, e1, [IP1, IP2, X]]$

Link layer on H2 delivers the packet to the network layer on H2

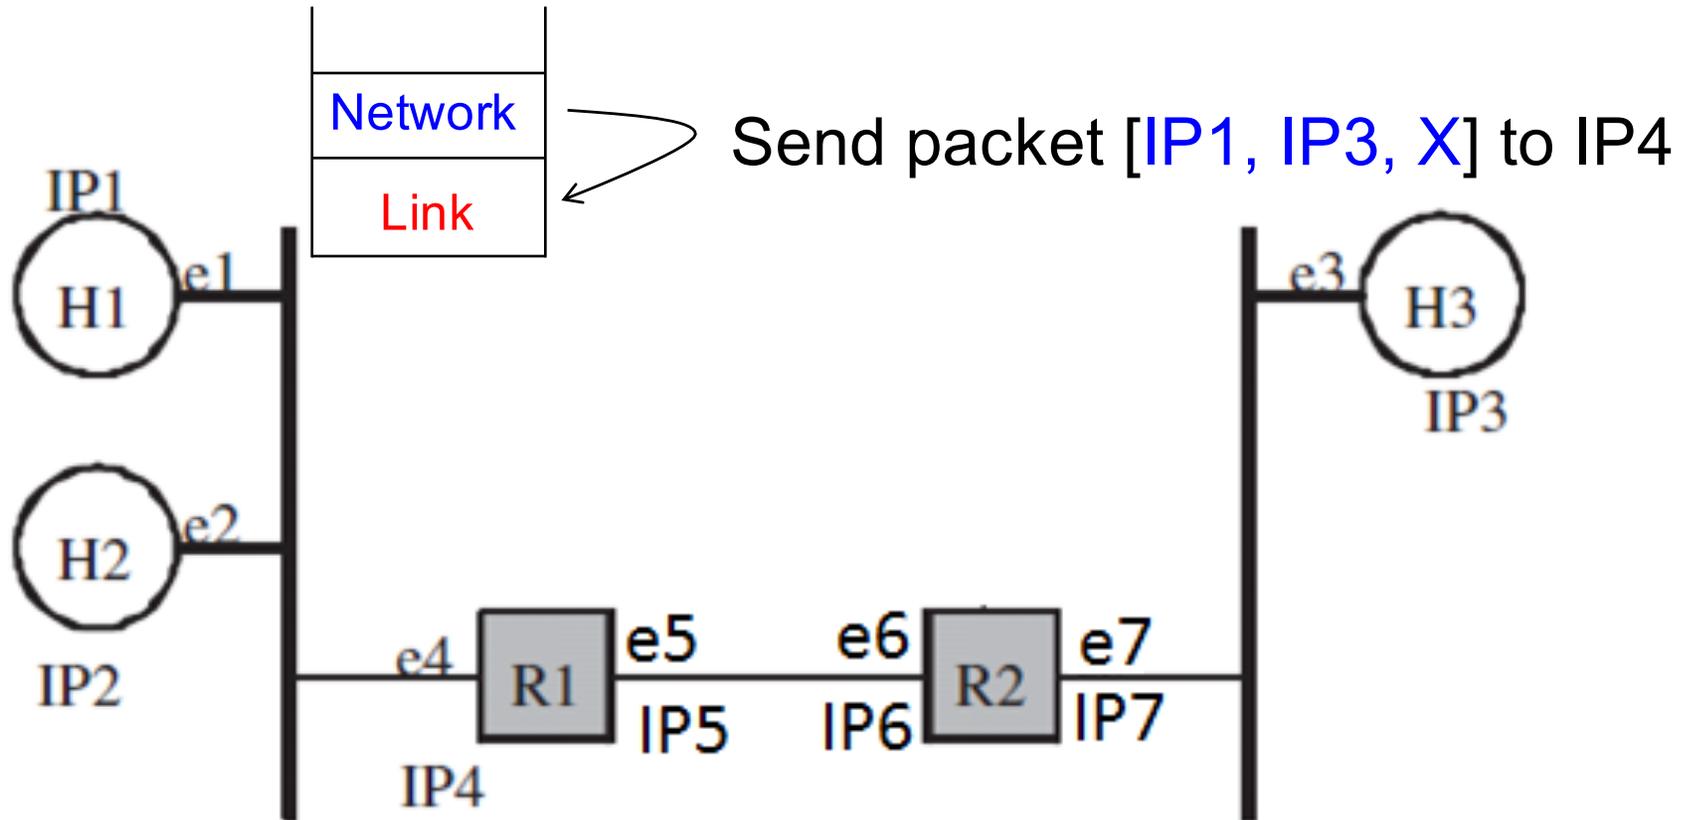


Example: H1 wants to send packet to H3





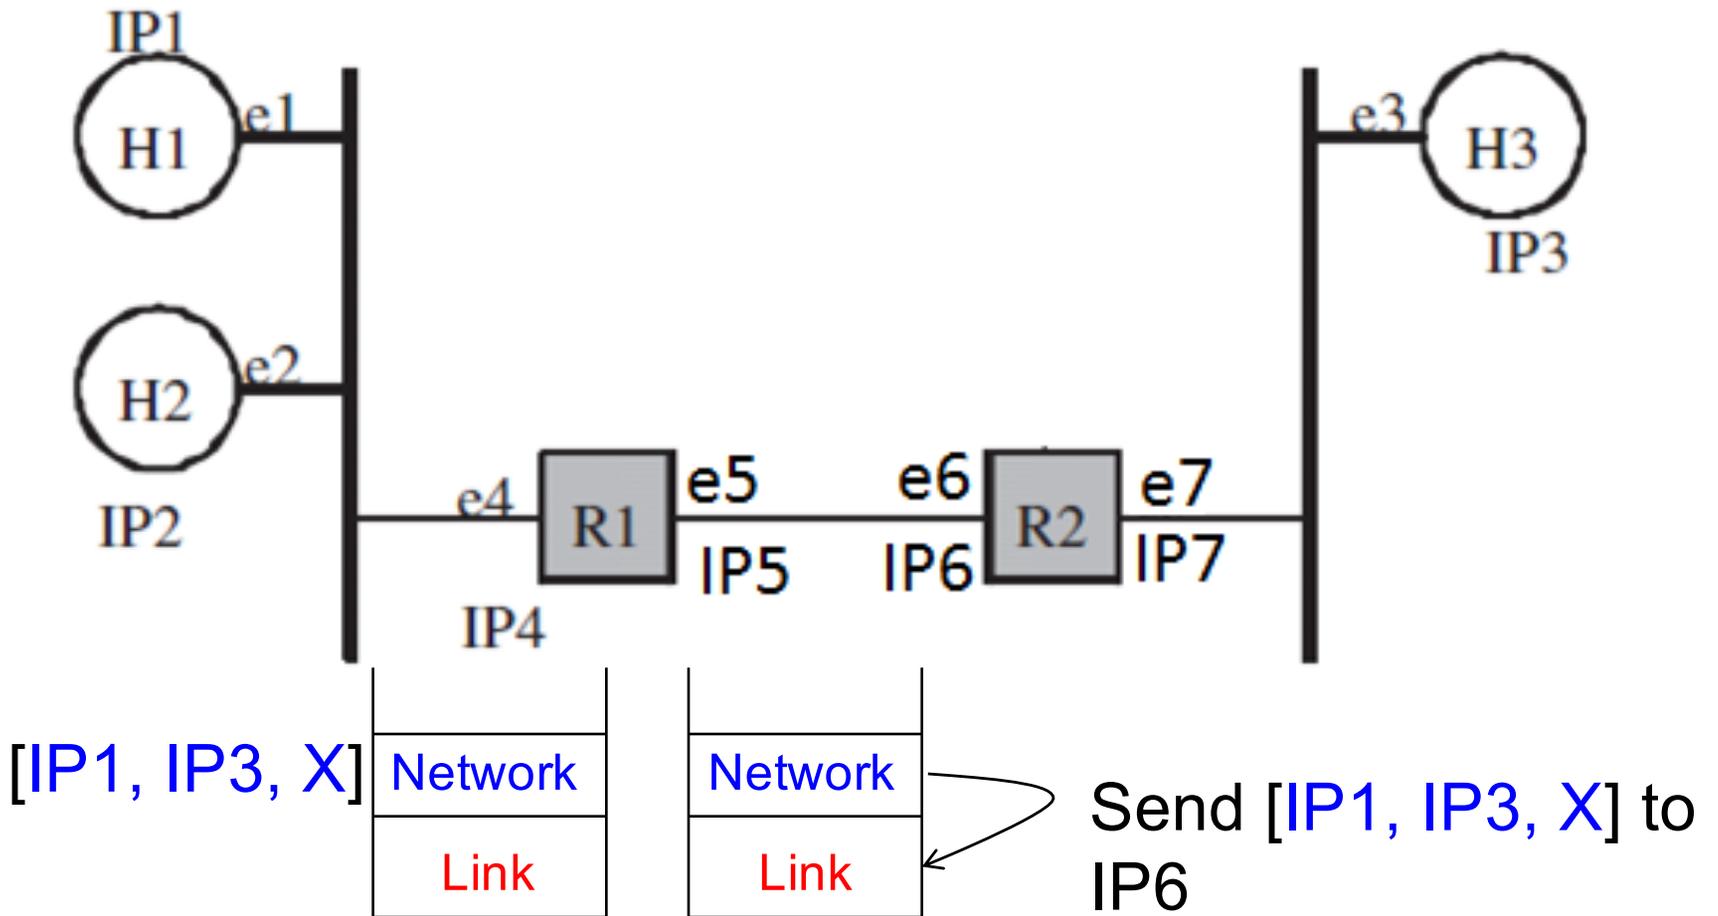
Example: H1 wants to send packet to H3



Network layer on H1 checks that IP3 is not in its subnet, so it routes it to its gateway IP4

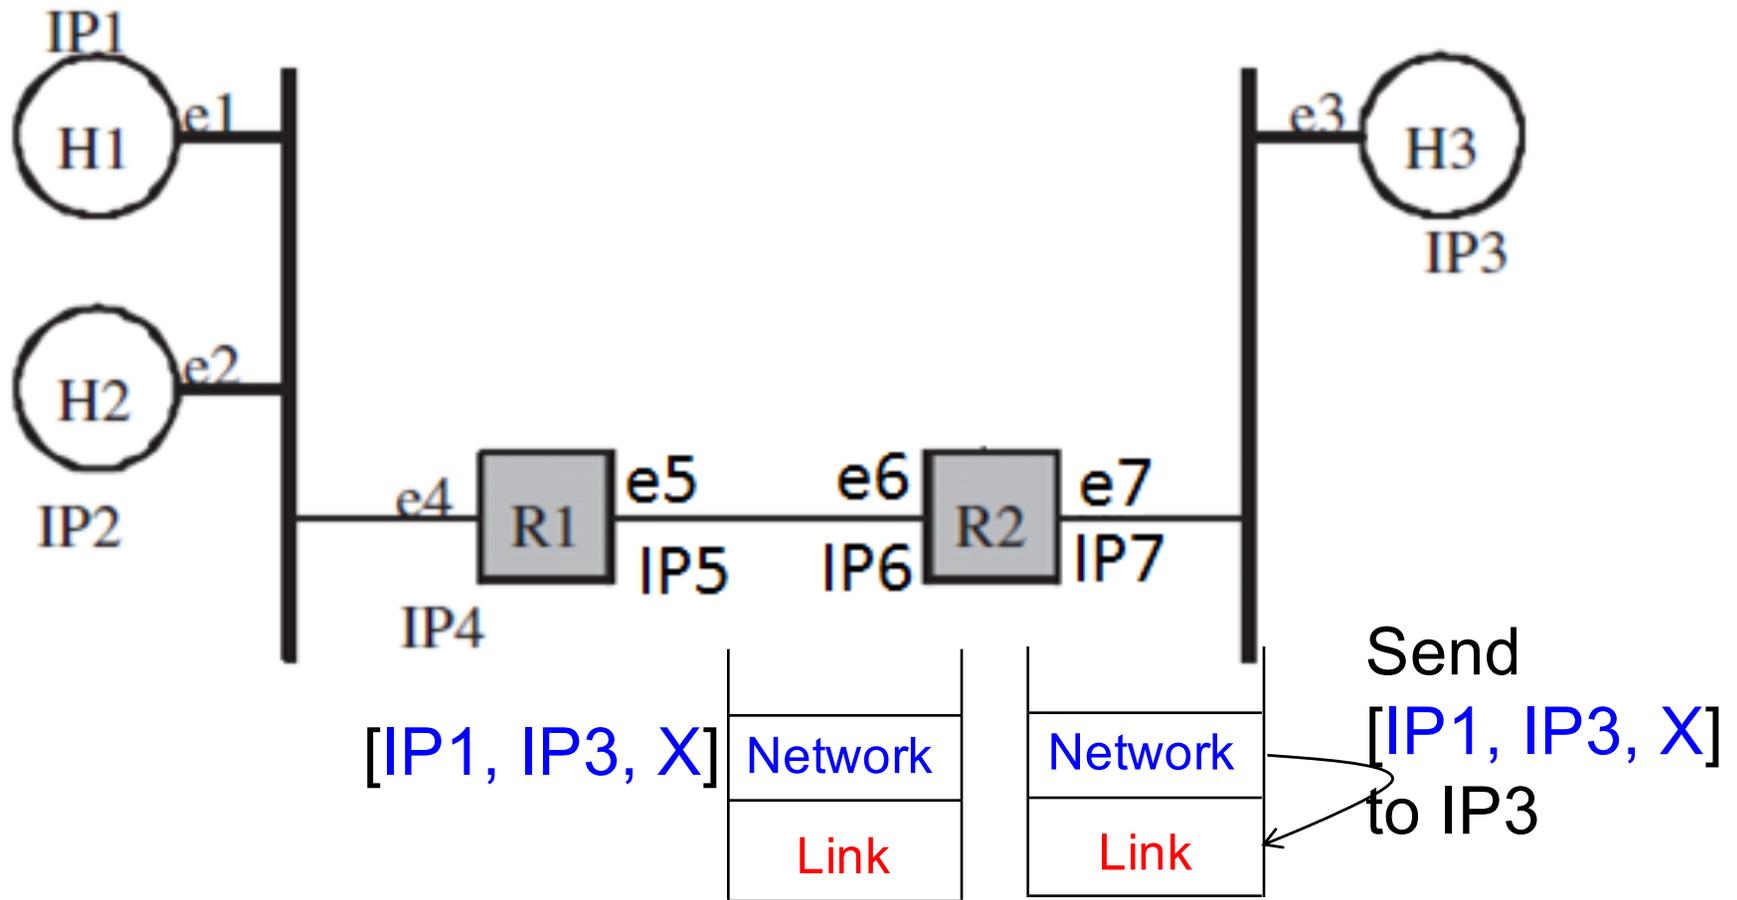


Example: H1 wants to send packet to H3





Example: H1 wants to send packet to H3





Summary

□ Network layer (Layer 3)

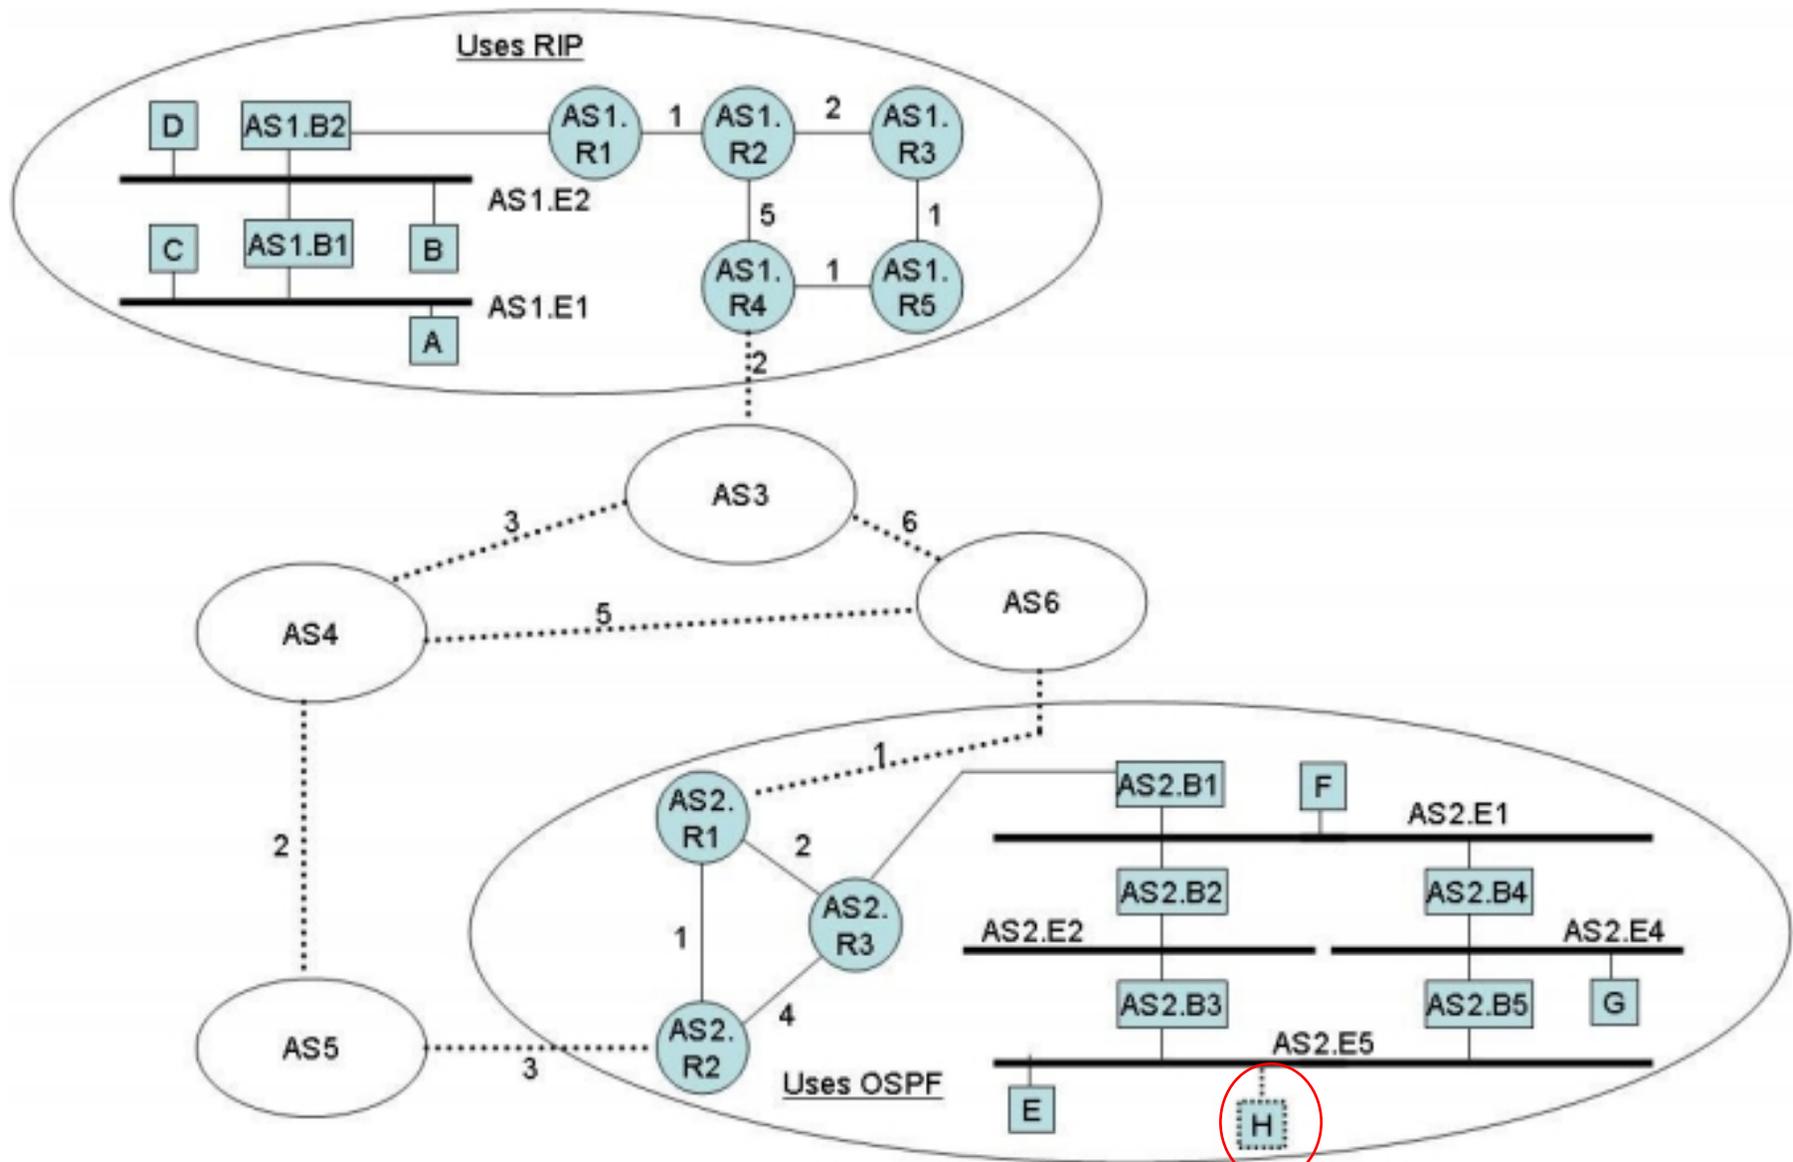
- Provides host-to-host communication service by **finding a path of routers connecting any two hosts**
- Main function: routing

□ Link layer (Layer 2)

- Provides **host-router and router-router communication** by utilizing the physical communication links
=> provides a service to the network layer



Putting it all together

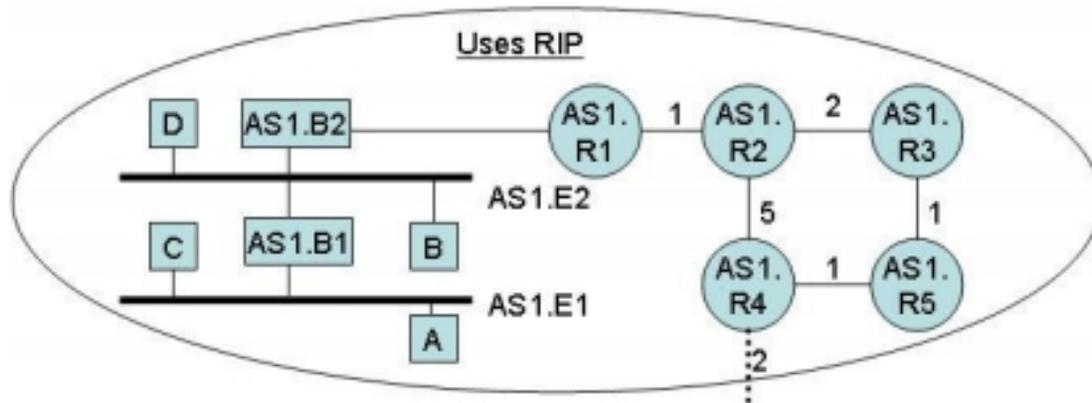


initially disconnected

Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



Putting it all together

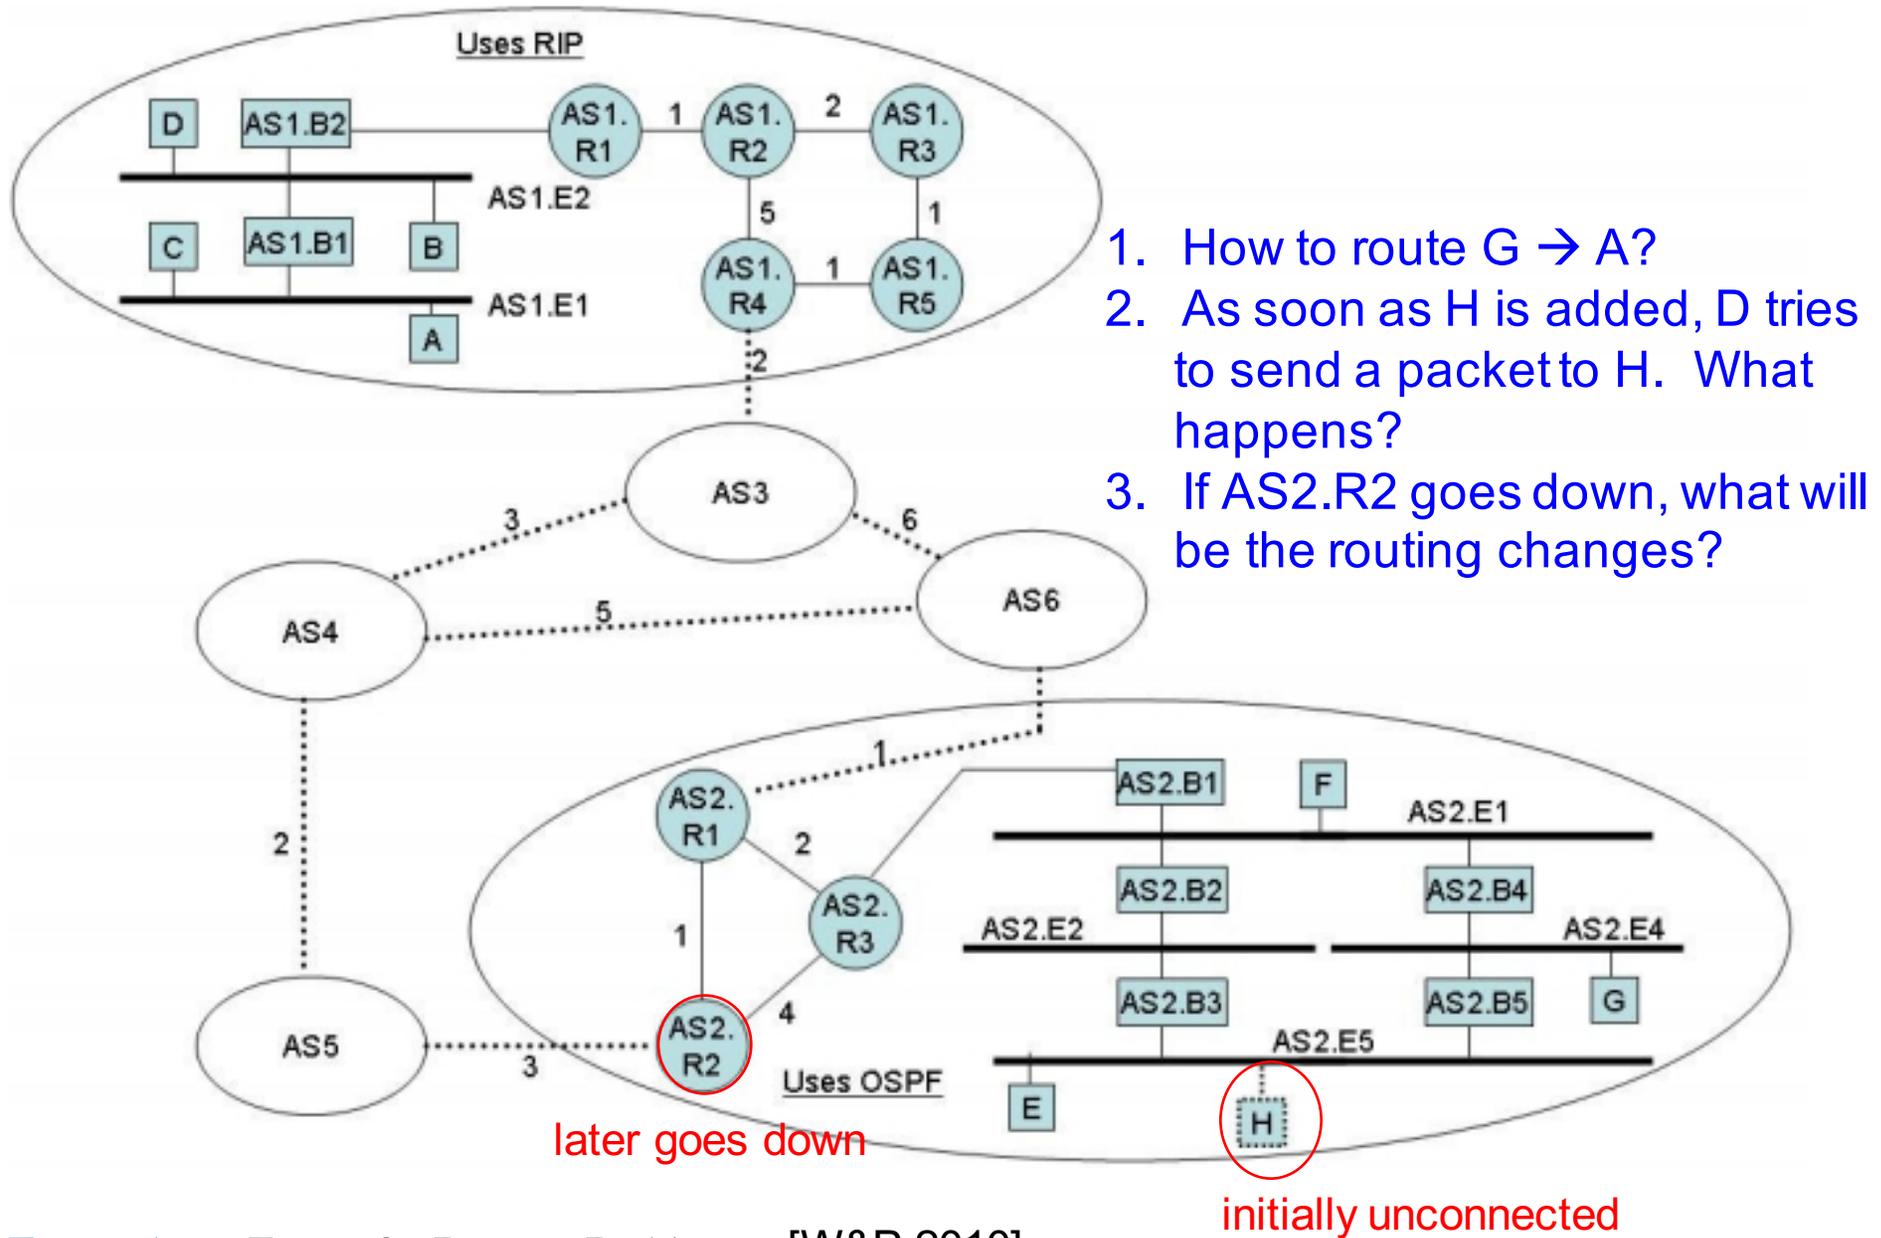


- The network addresses of nodes are given by $\langle \text{AS} \rangle . \langle \text{Network} \rangle . 0 . \langle \text{node} \rangle$, e.g., node A has the address AS1.E1.0.A,
- The bridge IDs satisfy $B1 < B2 < B3 \dots$,
- H is not connected to AS2.E5 for part (a),
- The BGP Speakers use the least-next-hop-cost policy for routing (i.e., among alternative paths to the destination AS, choose the one that has the least cost on the first hop), and
- The network topology shown has been stable for a long enough time to allow all the routing algorithms to converge and all the bridges to learn where to forward each packet.

Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



Putting it all together



1. How to route $G \rightarrow A$?
2. As soon as H is added, D tries to send a packet to H. What happens?
3. If AS2.R2 goes down, what will be the routing changes?

Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



1. compute spanning tree

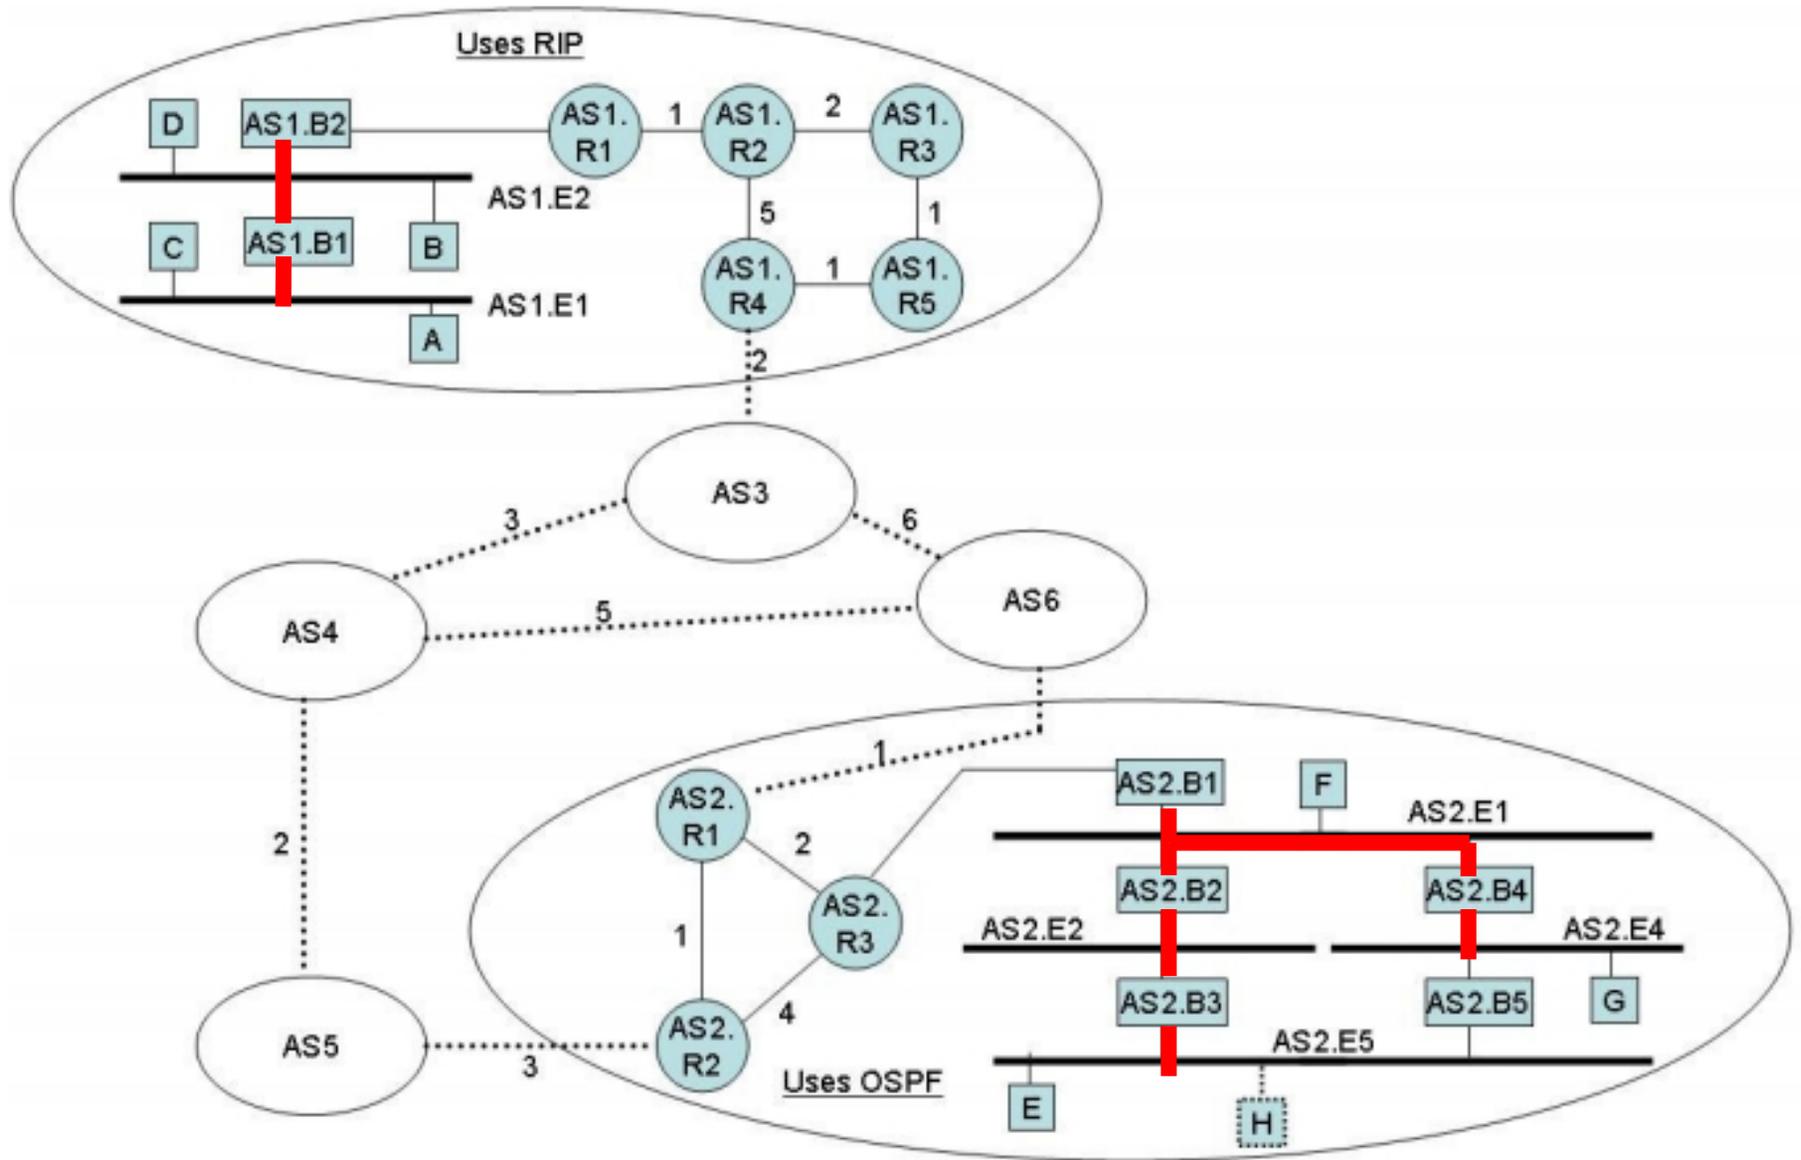
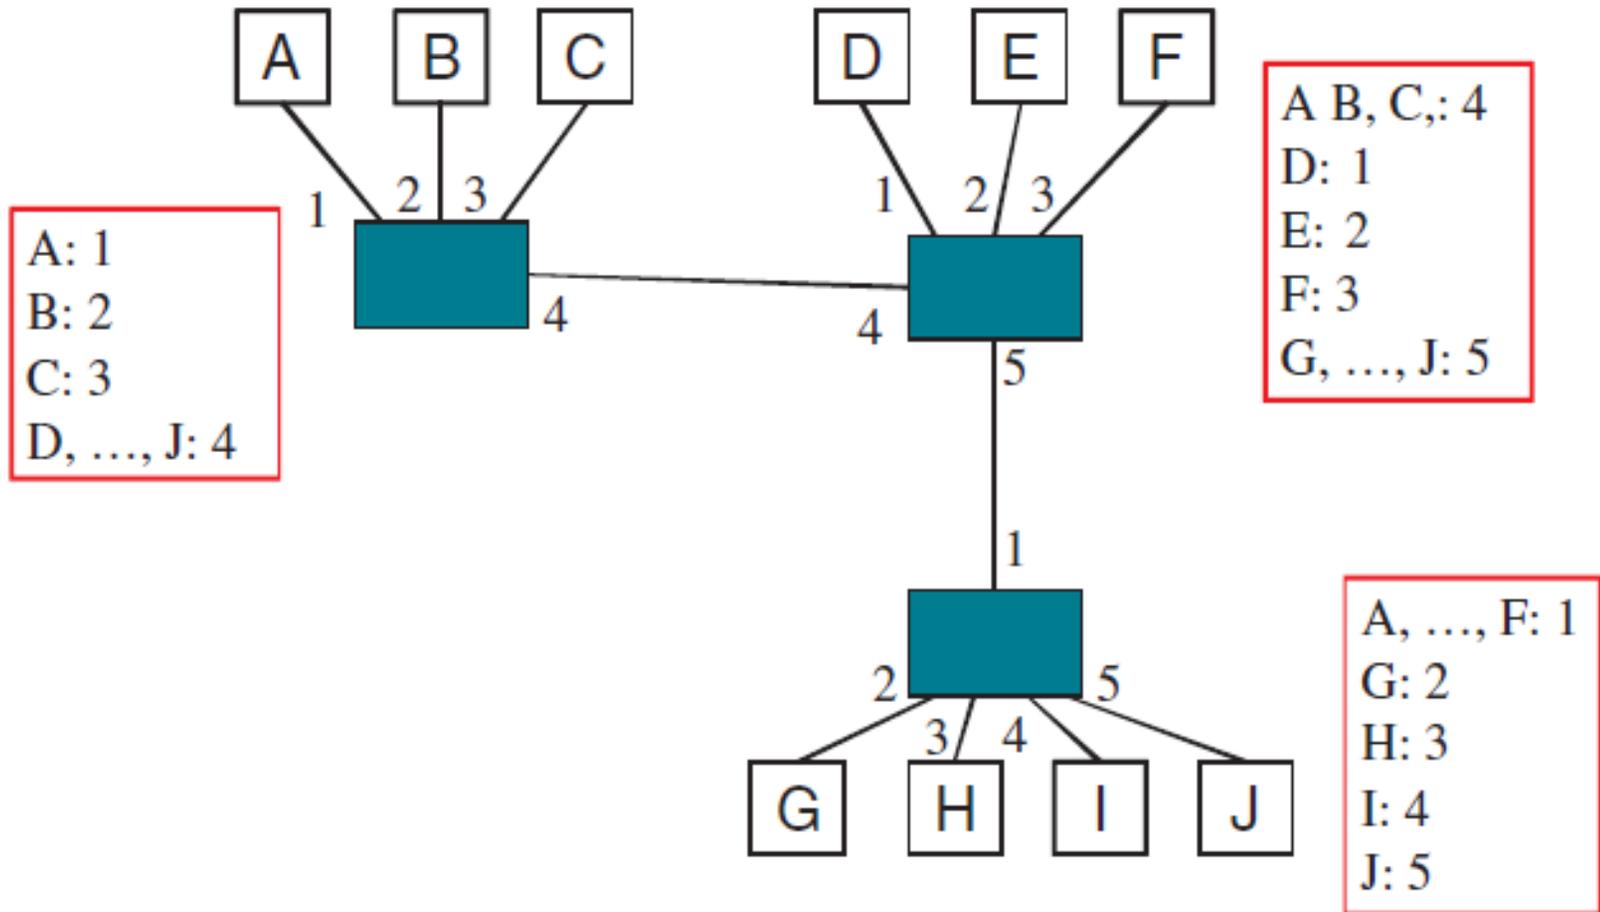


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



Recap: STP routing on subnet



$x \rightarrow y: [y | x | \text{data}]$



2. compute intra-AS routing

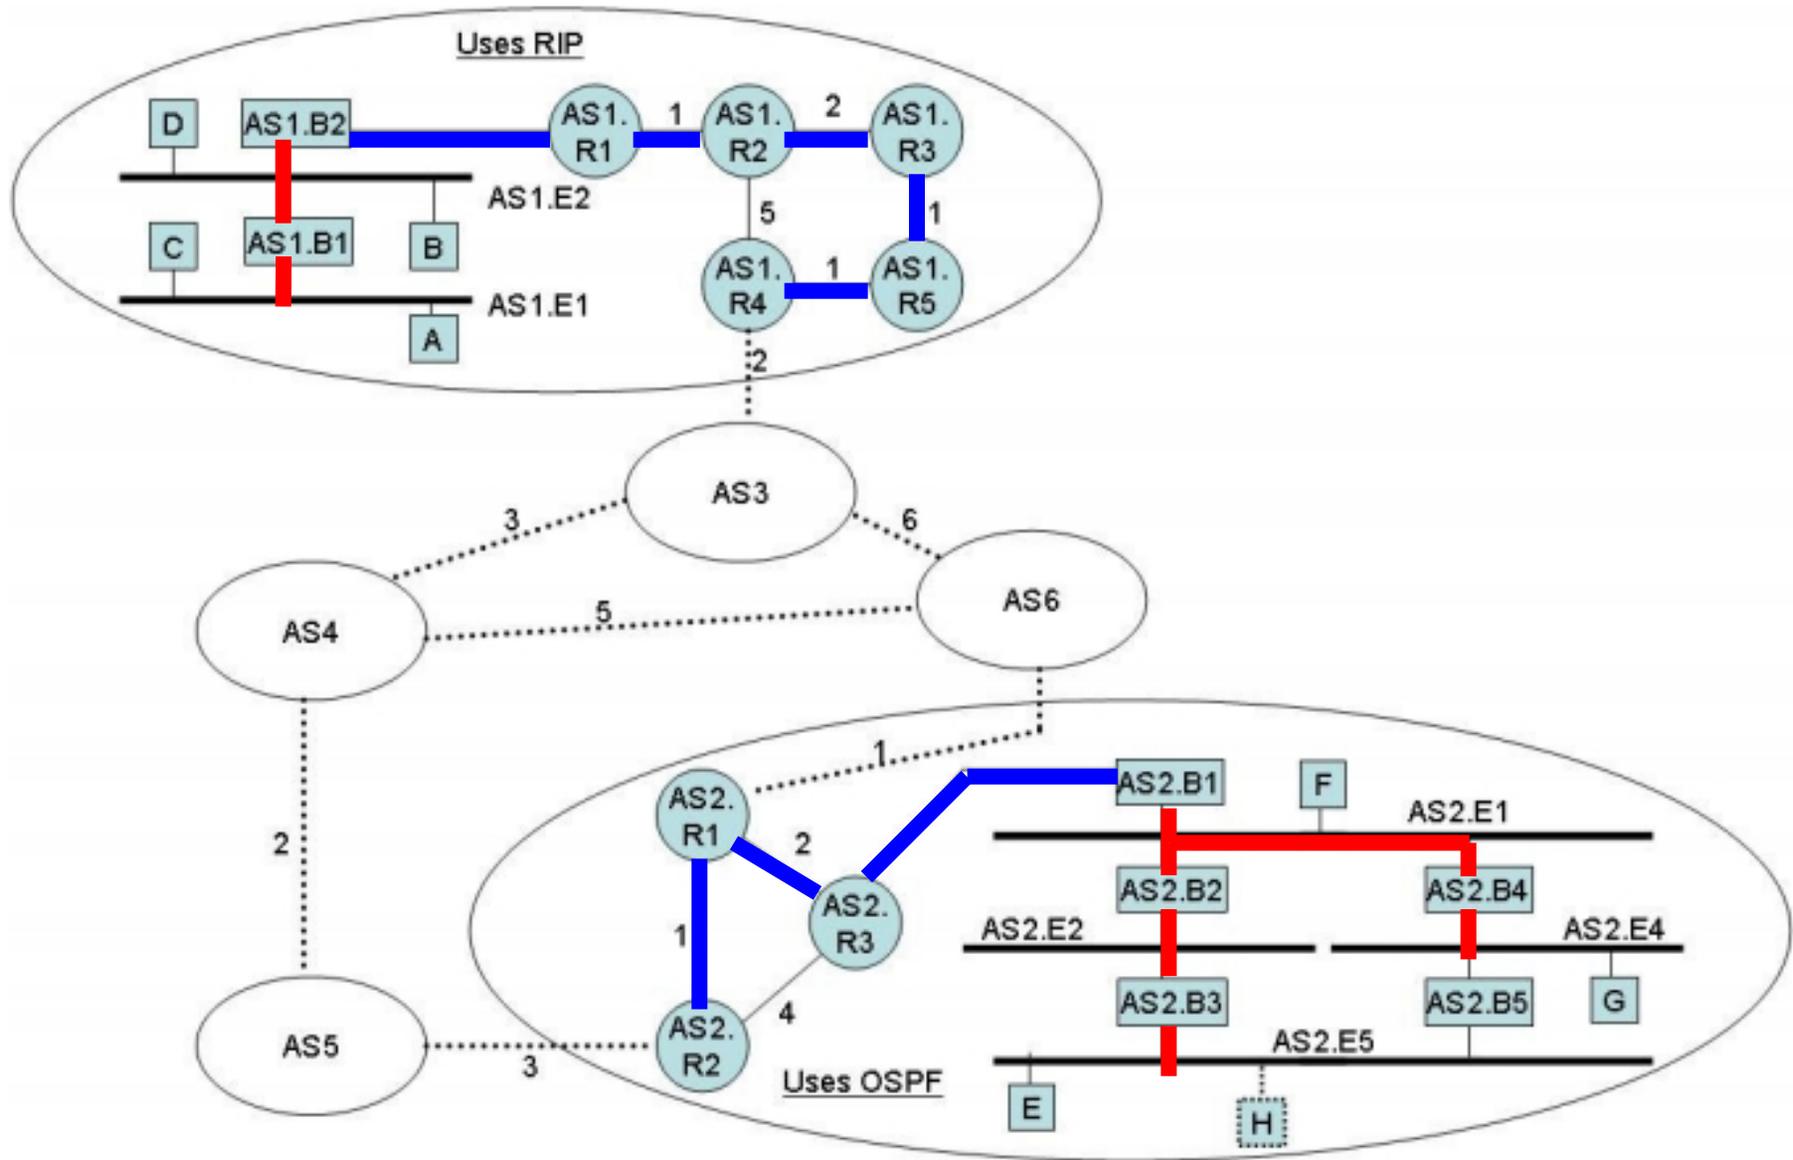
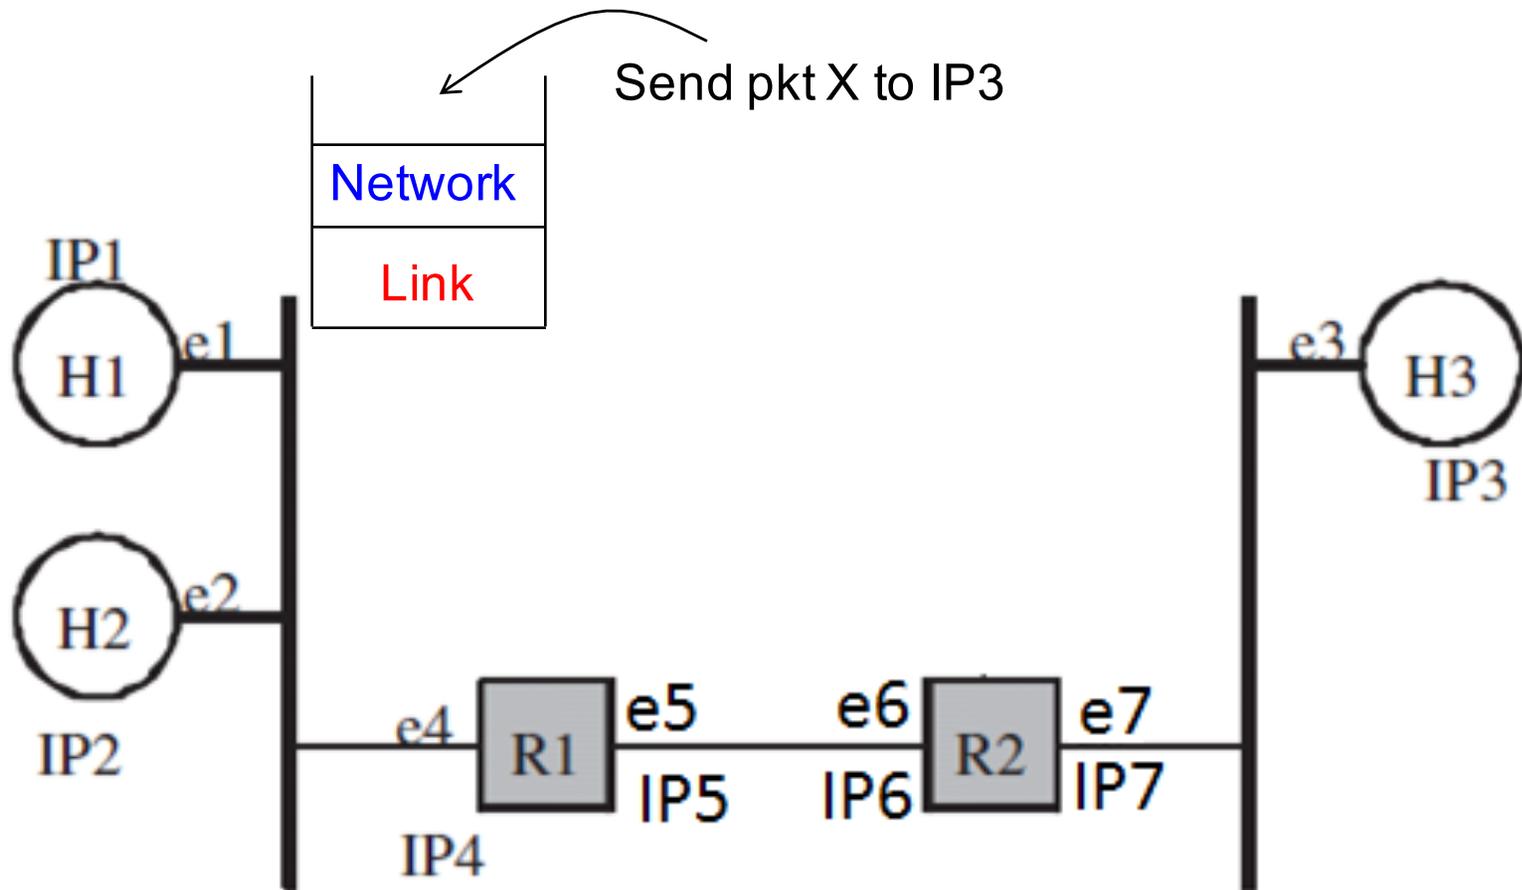


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



Recap: intra-AS routing





3. compute inter-AS routing

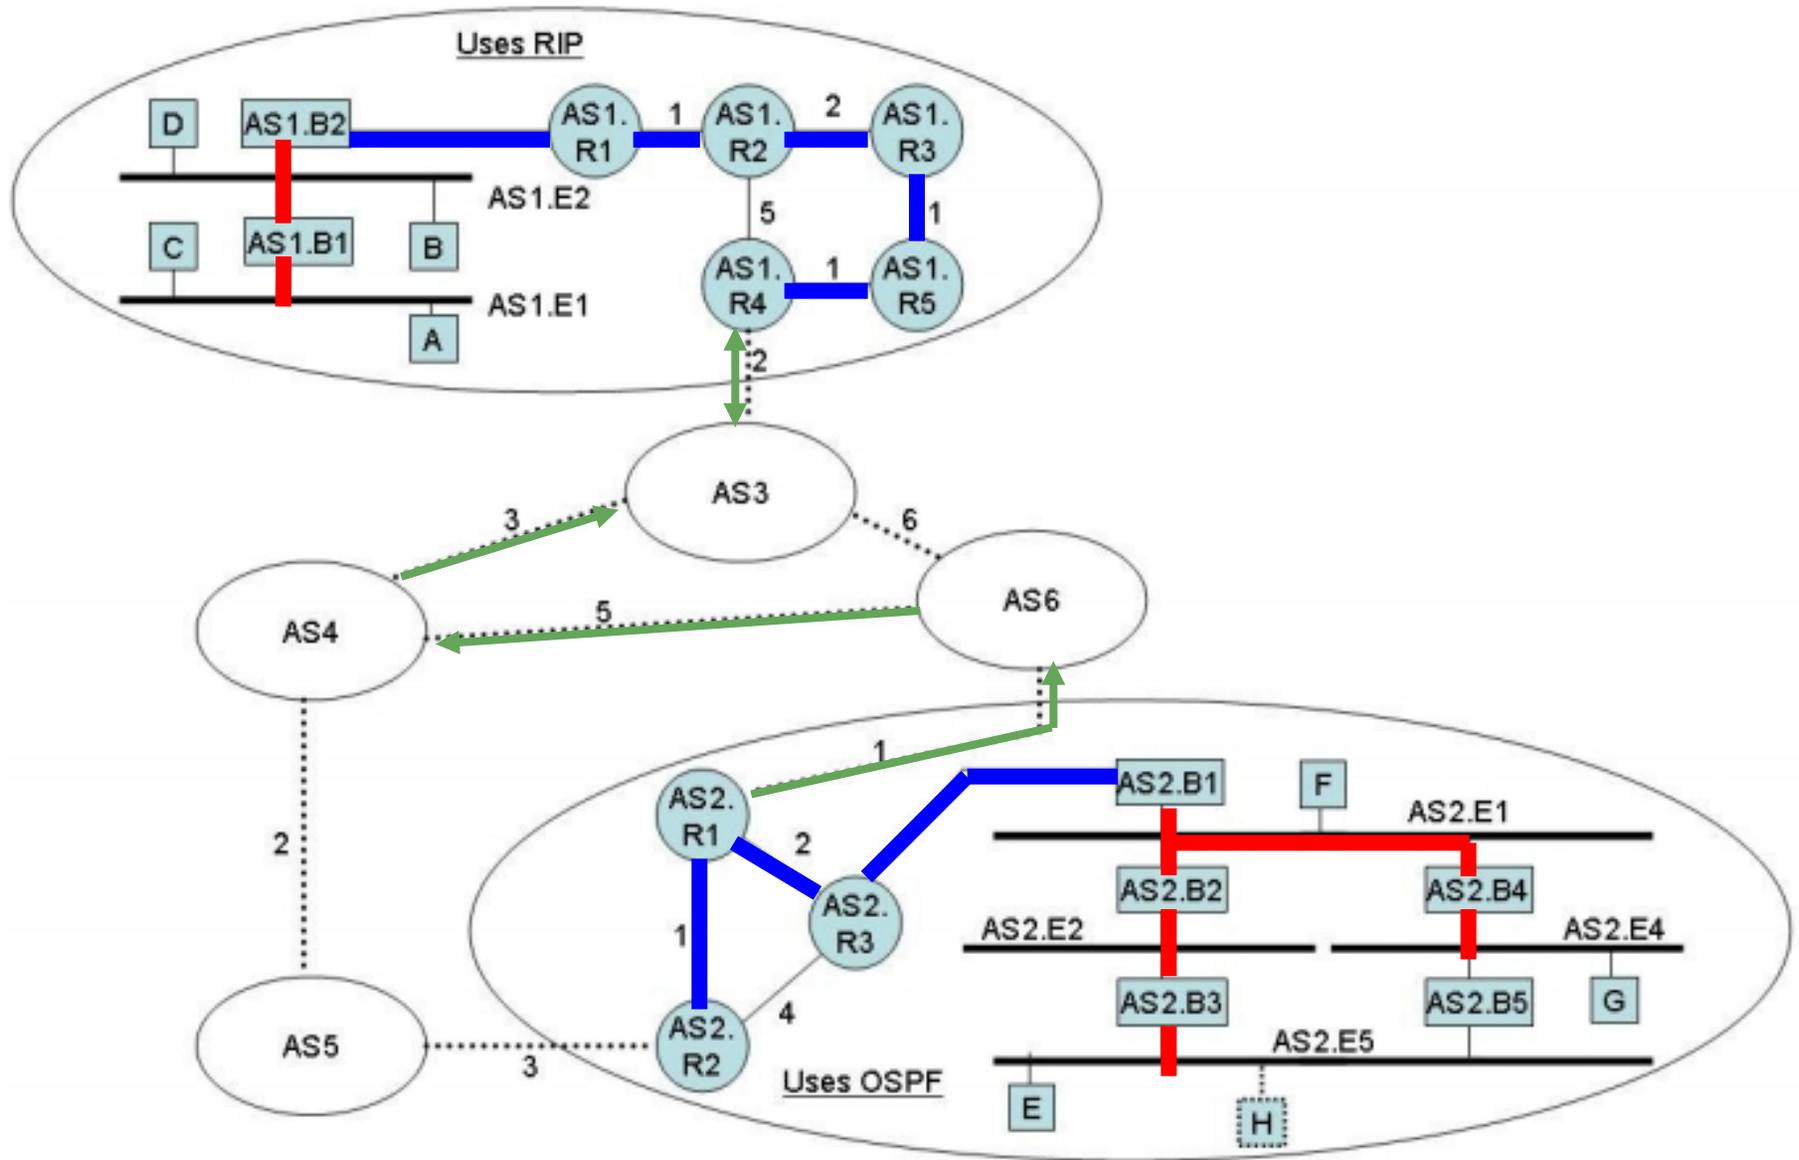
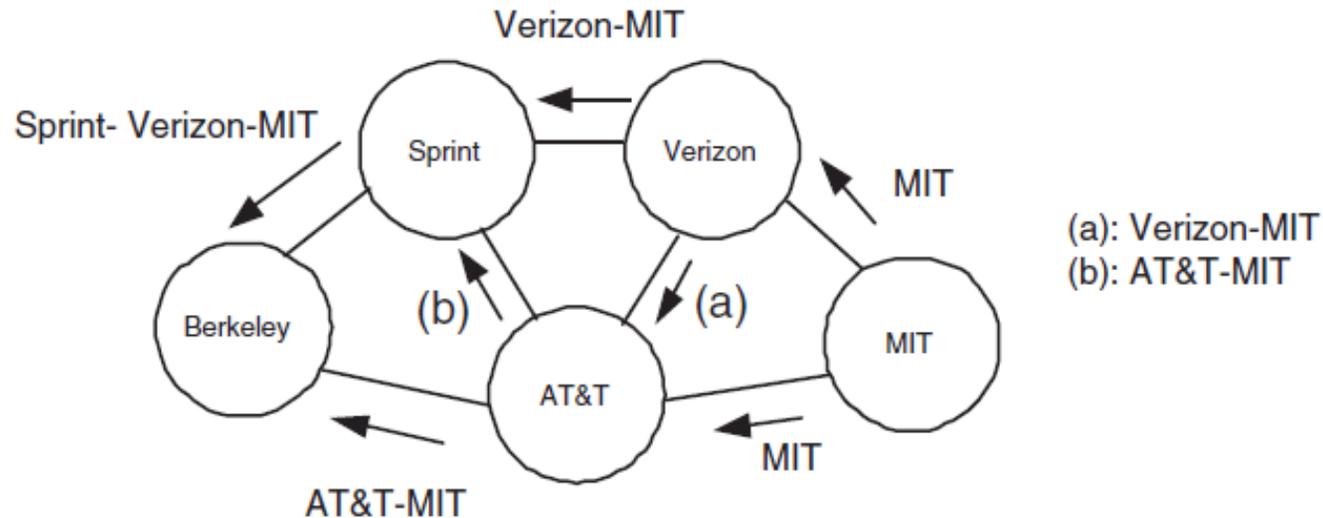


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



Recap: inter-AS routing



Example

BGP policy at Berkeley:

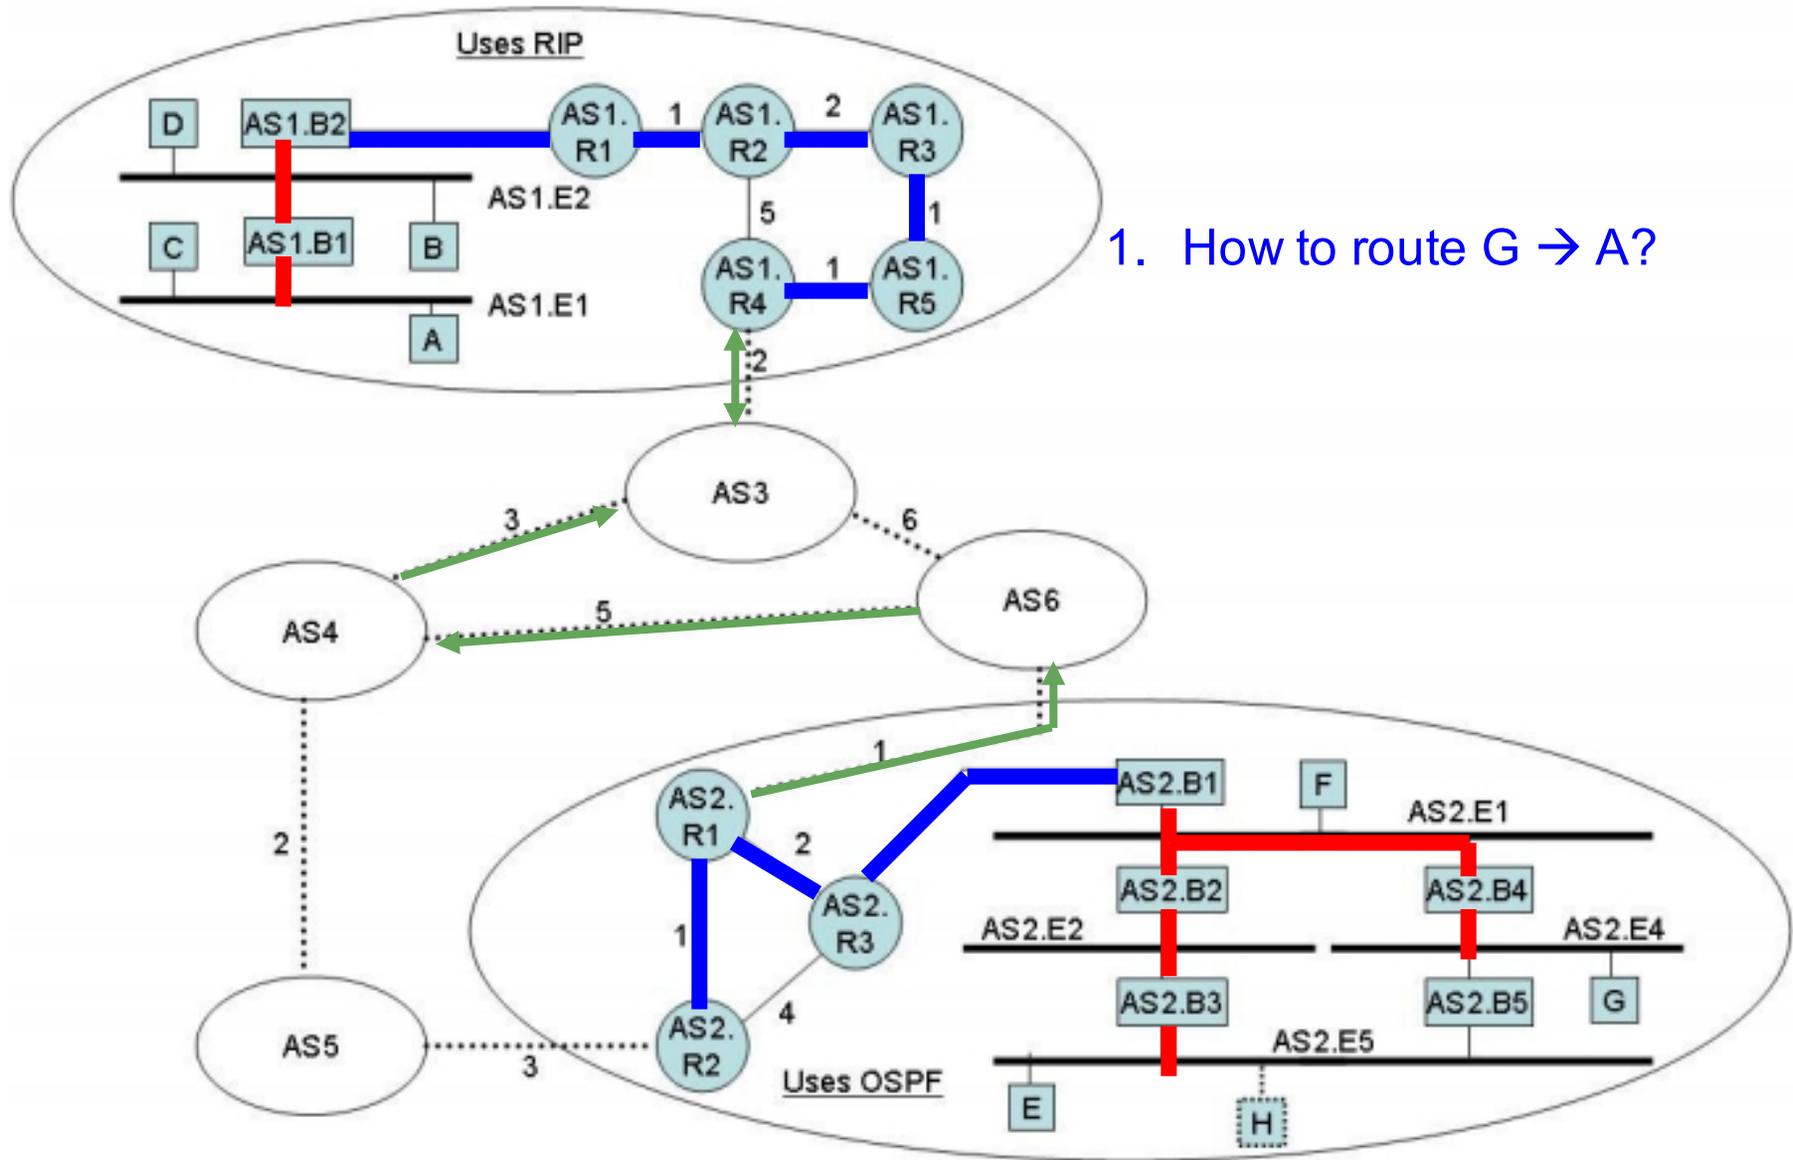
1. If possible, avoid AT&T
2. Choose path with smallest #hops
3. Alphabetical

Berkeley decision:

use path Sprint-Verizon-MIT to reach MIT



3. compute inter-AS routing

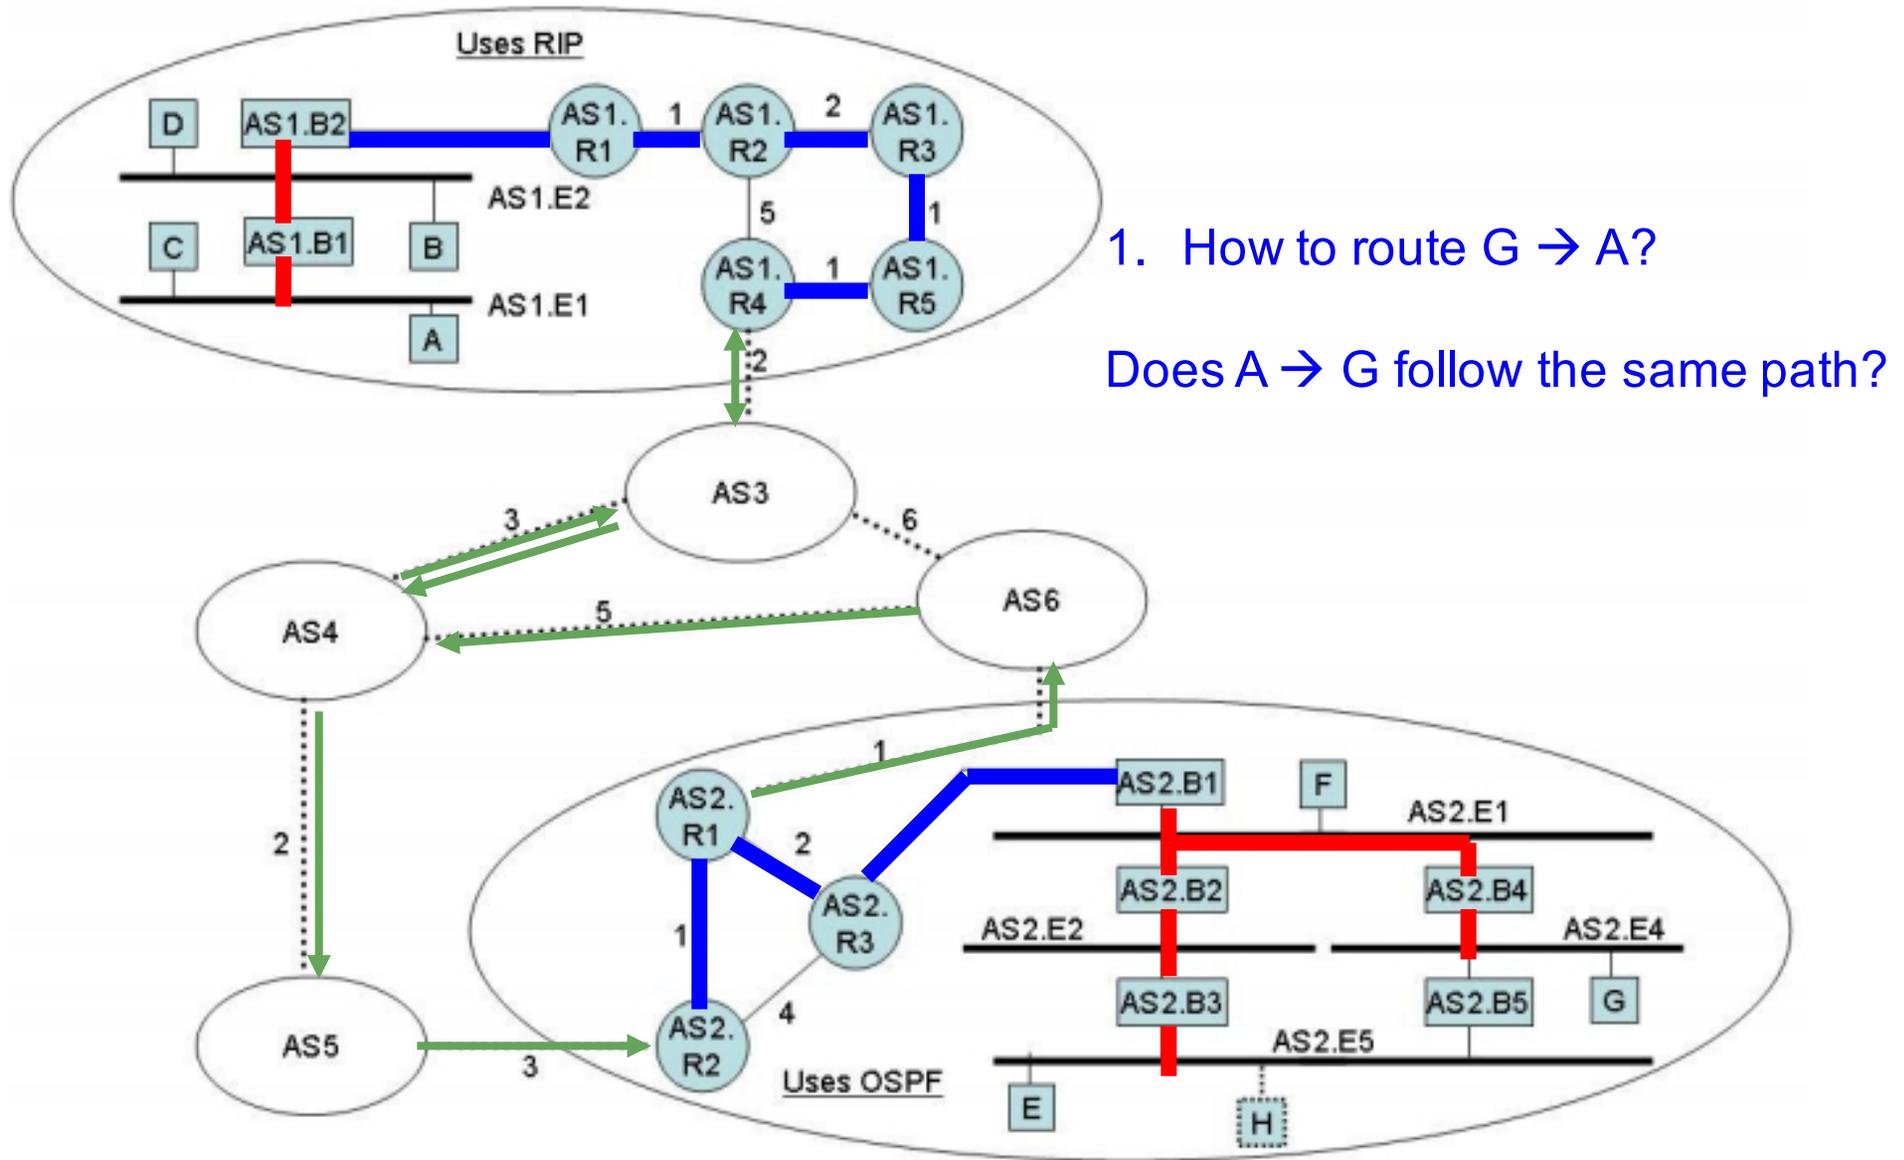


1. How to route G → A?

Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



3. compute inter-AS routing



1. How to route $G \rightarrow A$?

Does $A \rightarrow G$ follow the same path?

Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



4. Address resolution protocol

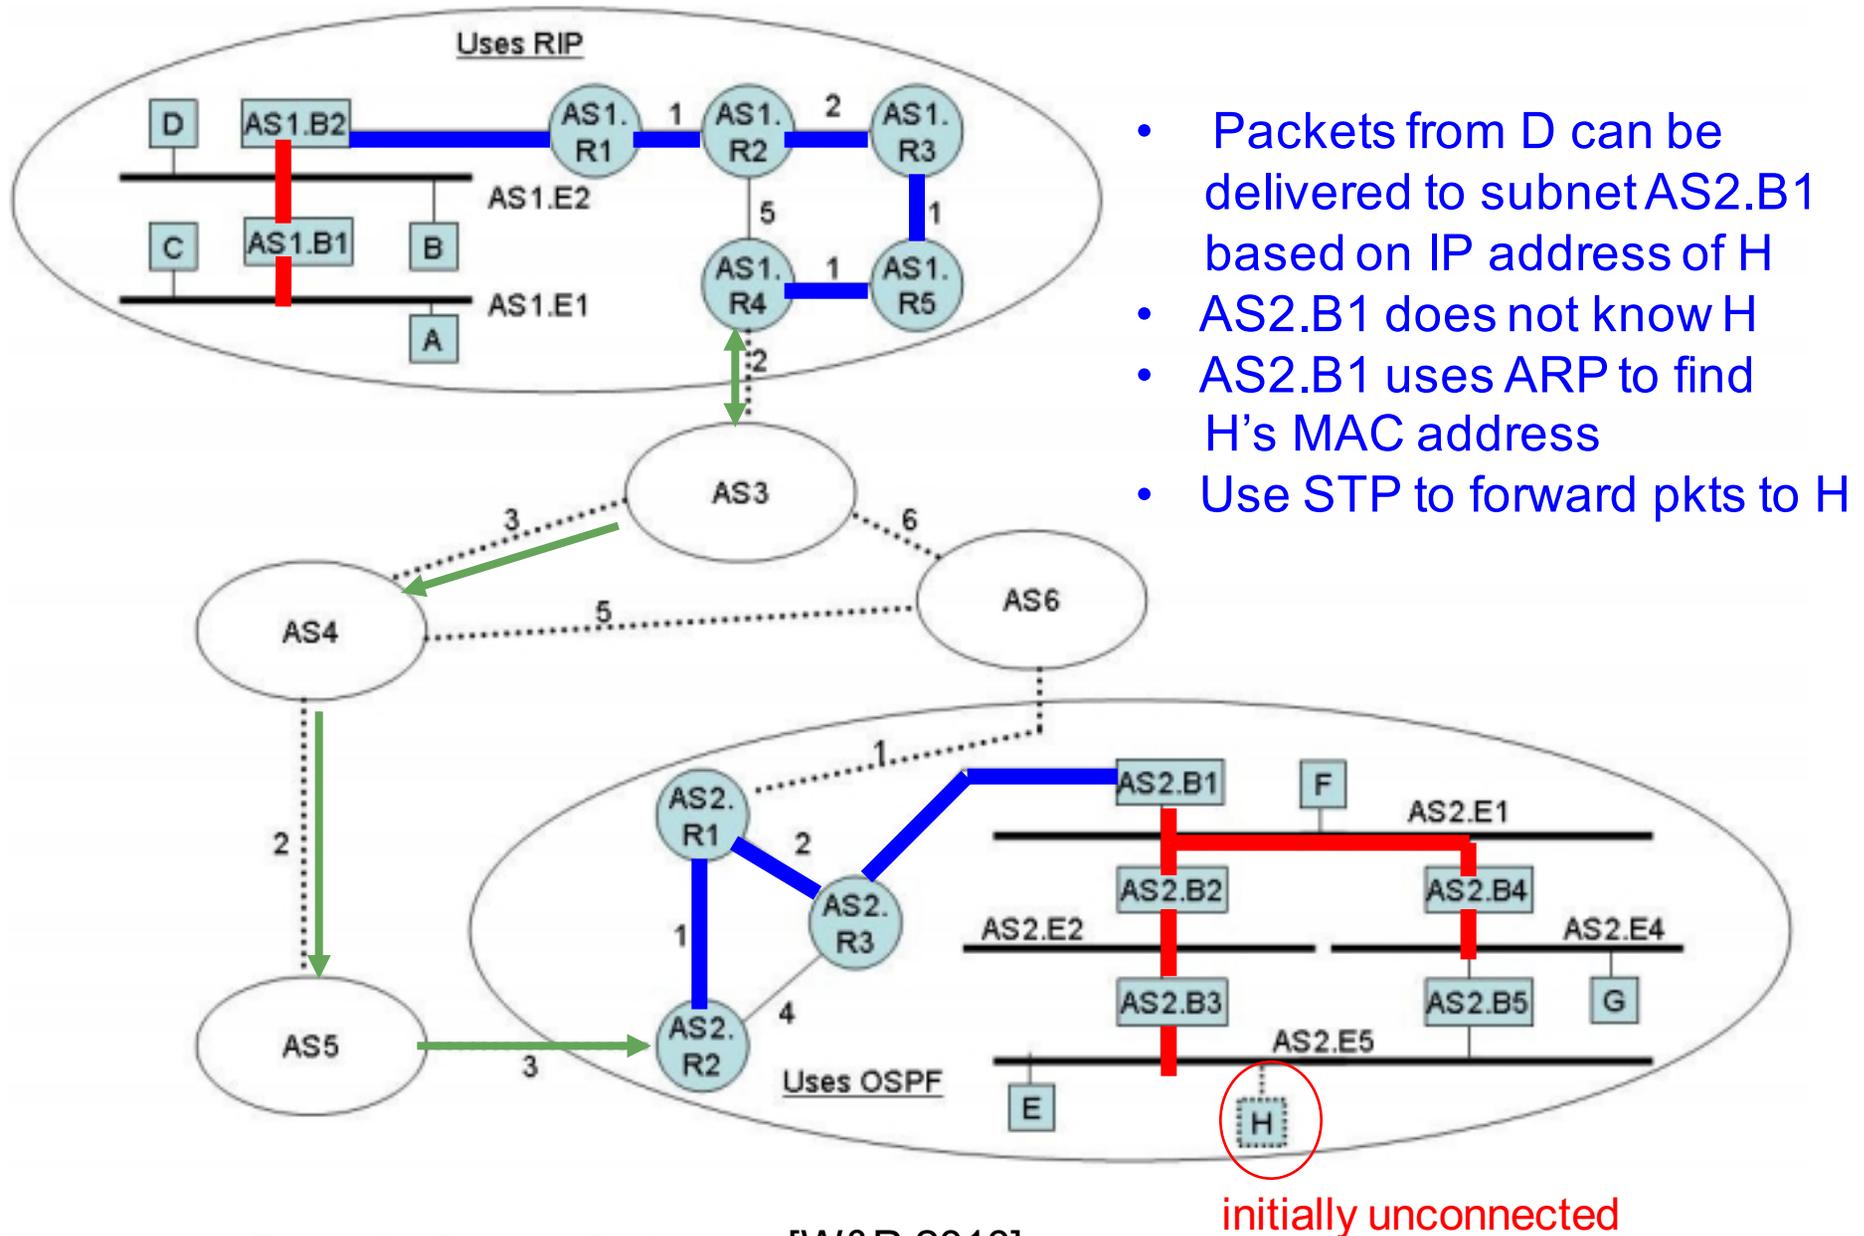
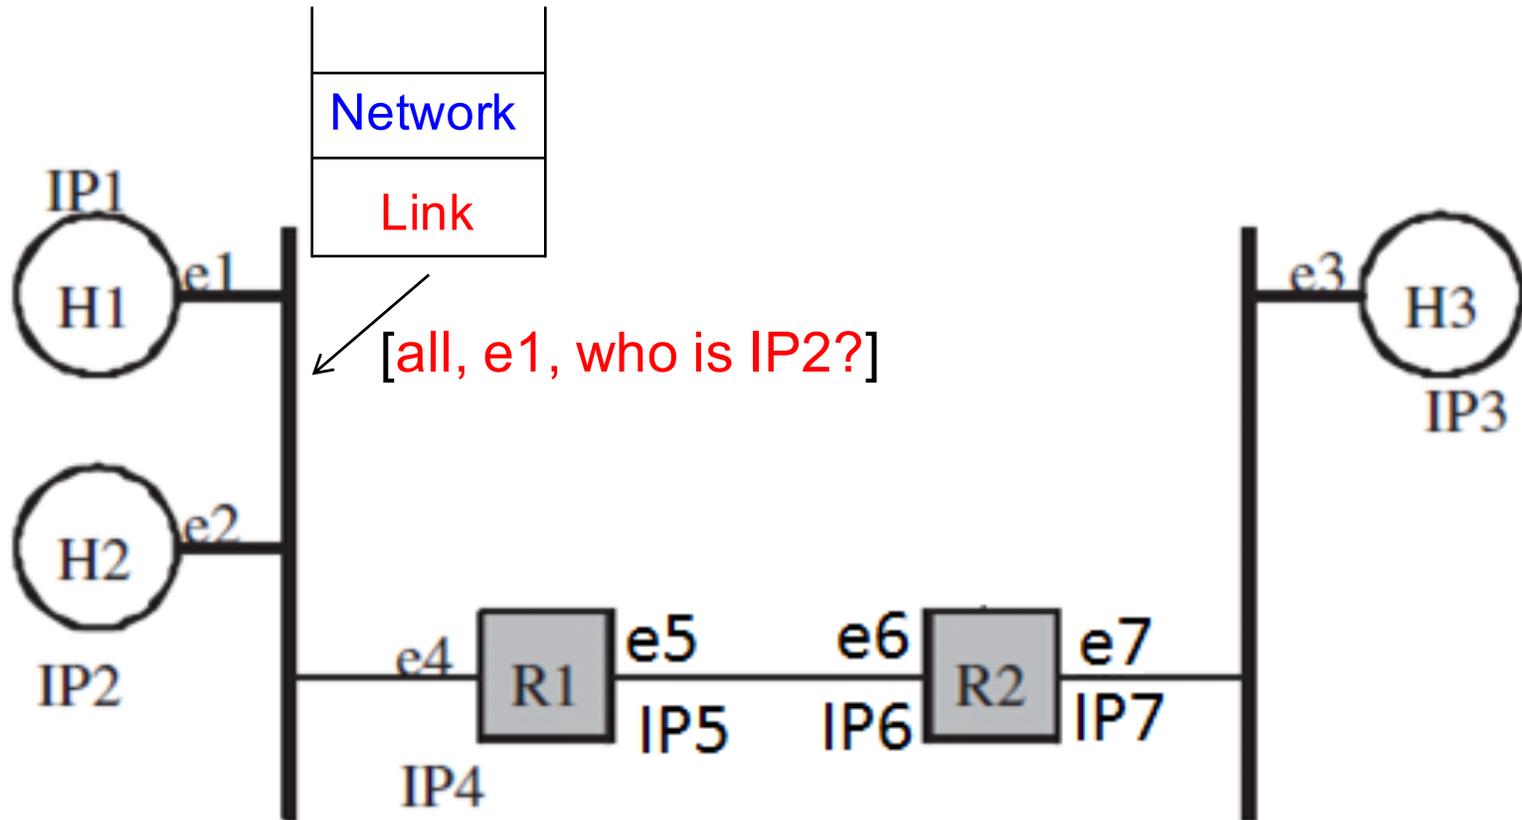


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



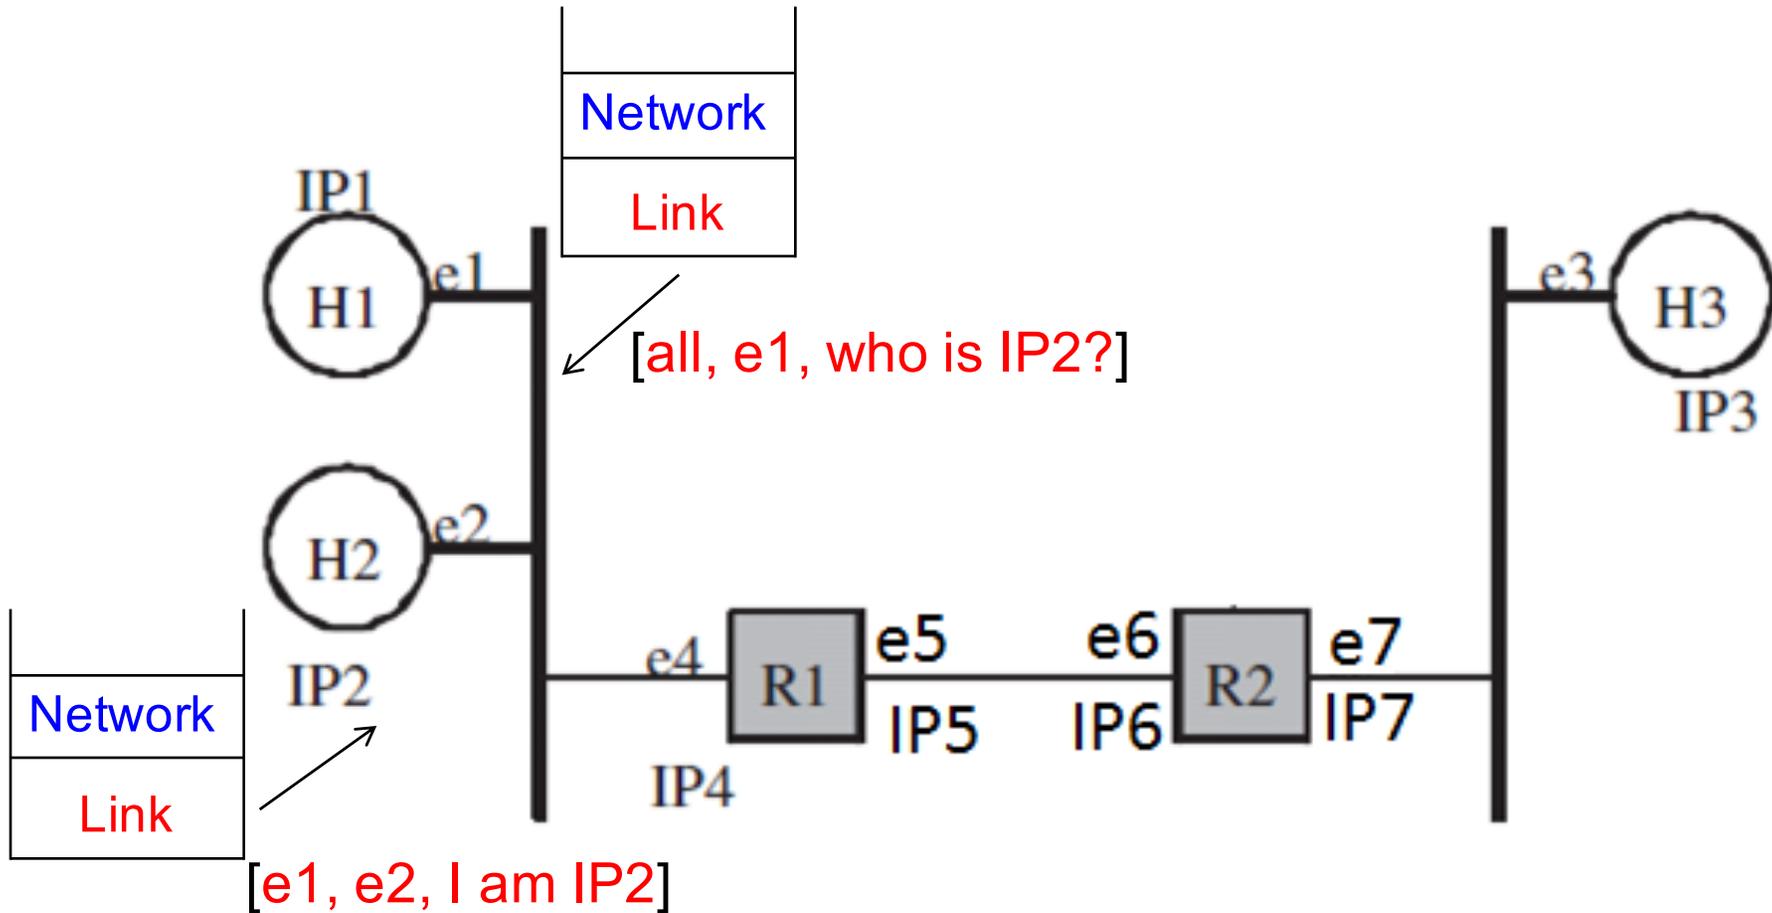
Example: H1 wants to send packet to H2



Link layer on H1 broadcasts a message (ARP query) on its layer 2 network asking for the MAC address corresponding to IP2



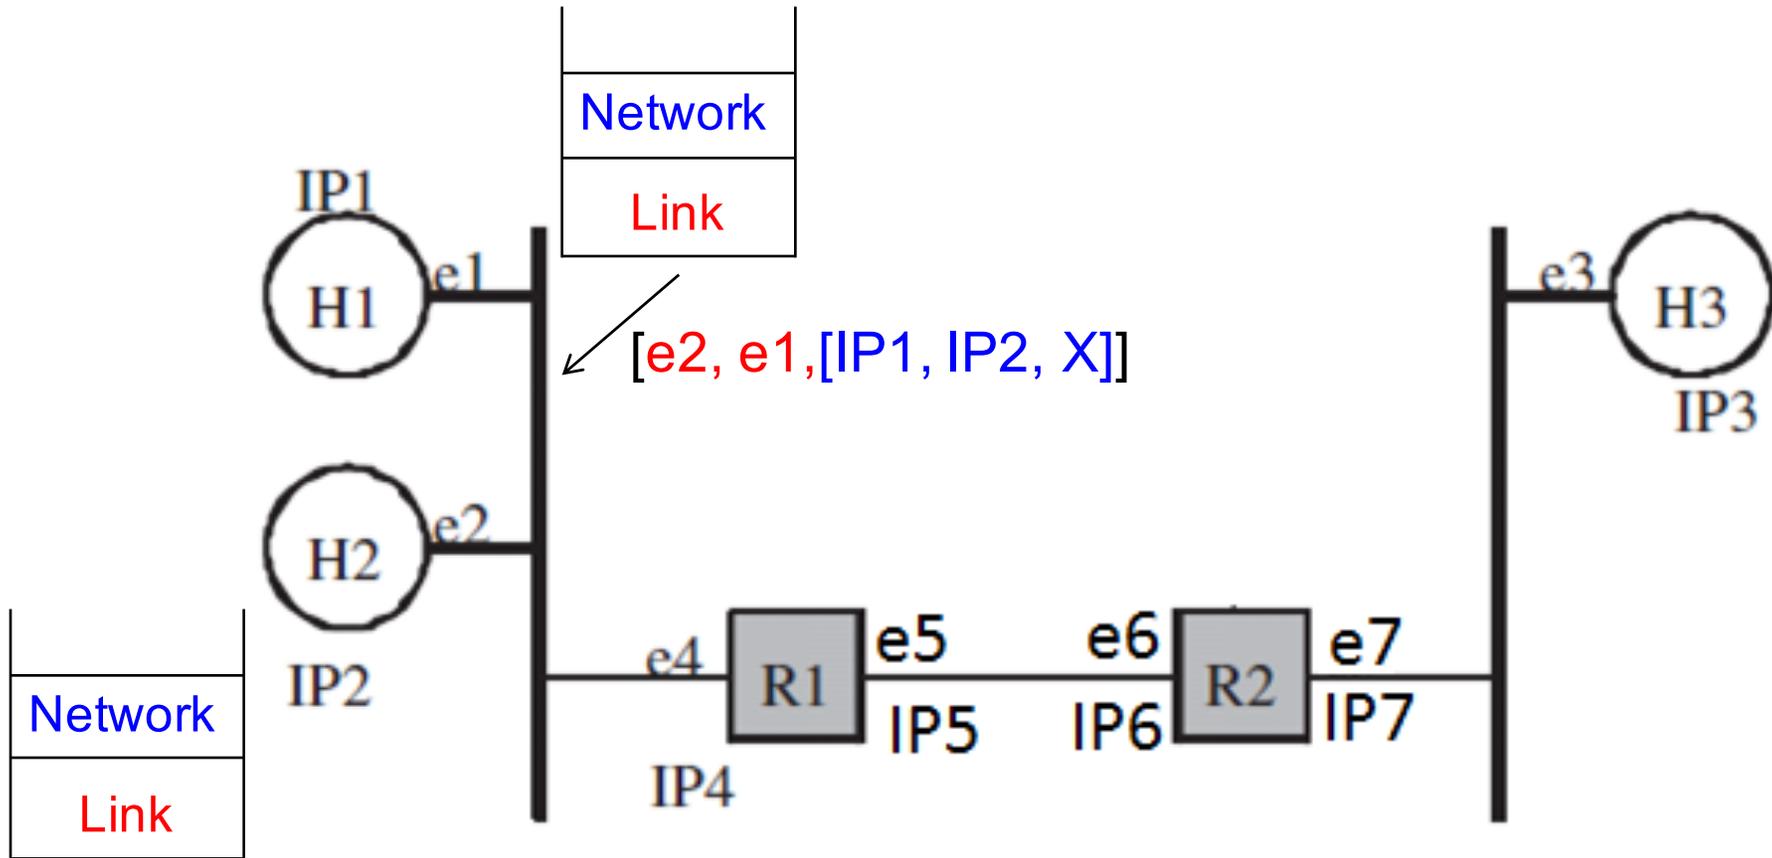
Example: H1 wants to send packet to H2



Link layer on H2 responds to the ARP query with its MAC address



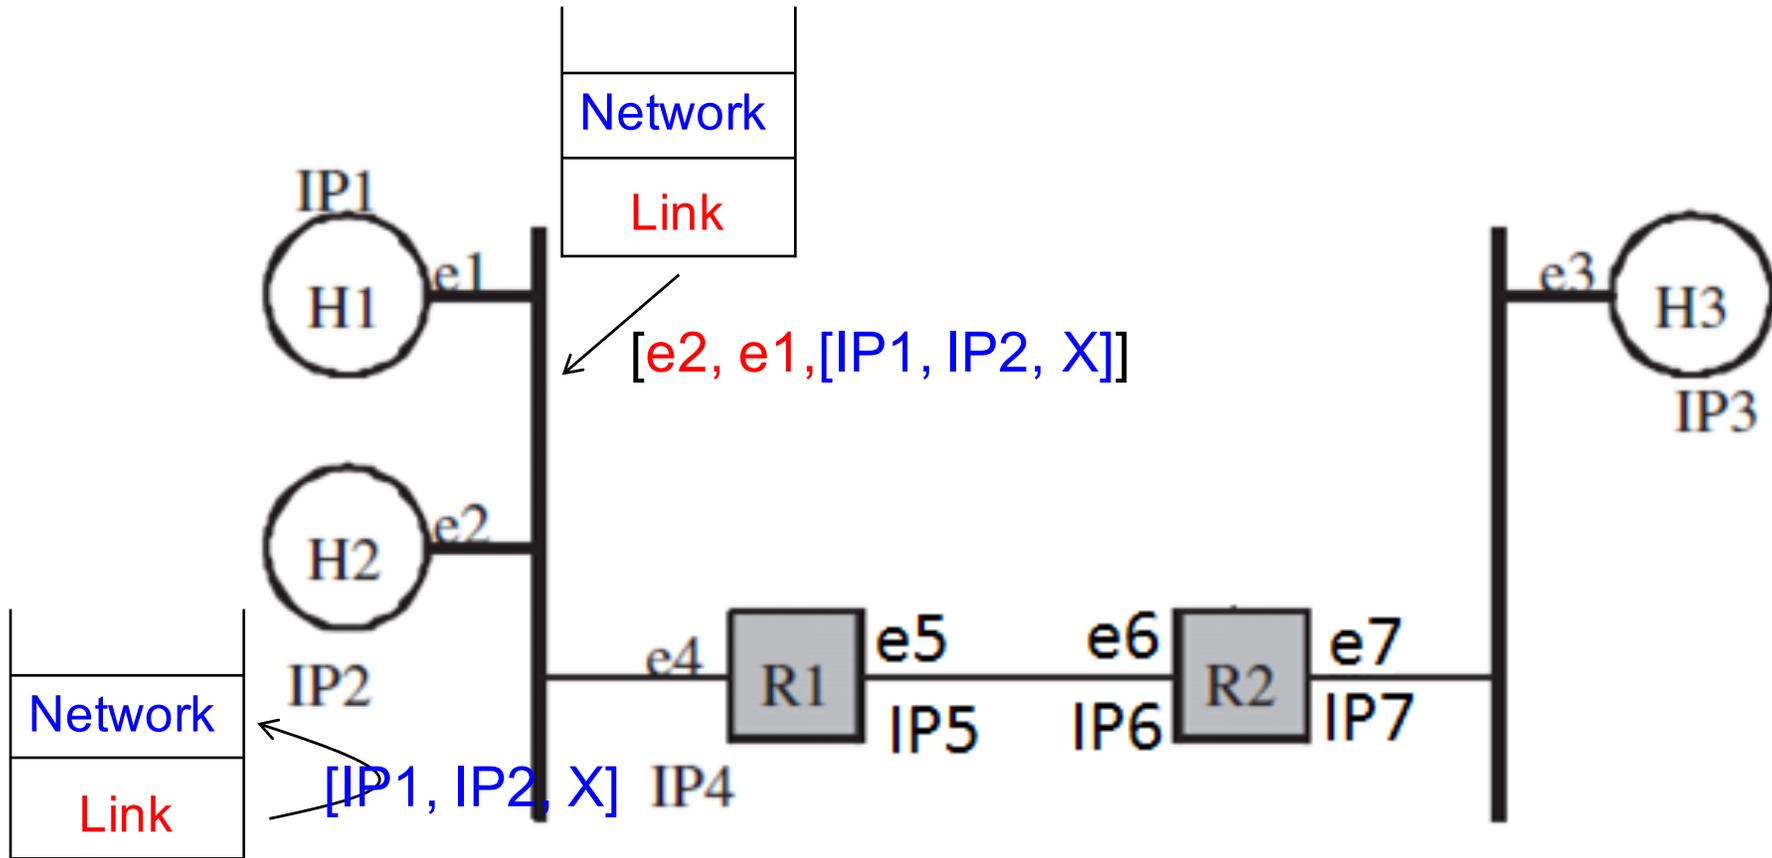
Example: H1 wants to send packet to H2



Once the link layer on H1 knows e2, it can now send the original message



Example: H1 wants to send packet to H2



$[e2, e1, [IP1, IP2, X]]$

Link layer on H2 delivers the packet to the network layer on H2



5. re-compute routing table

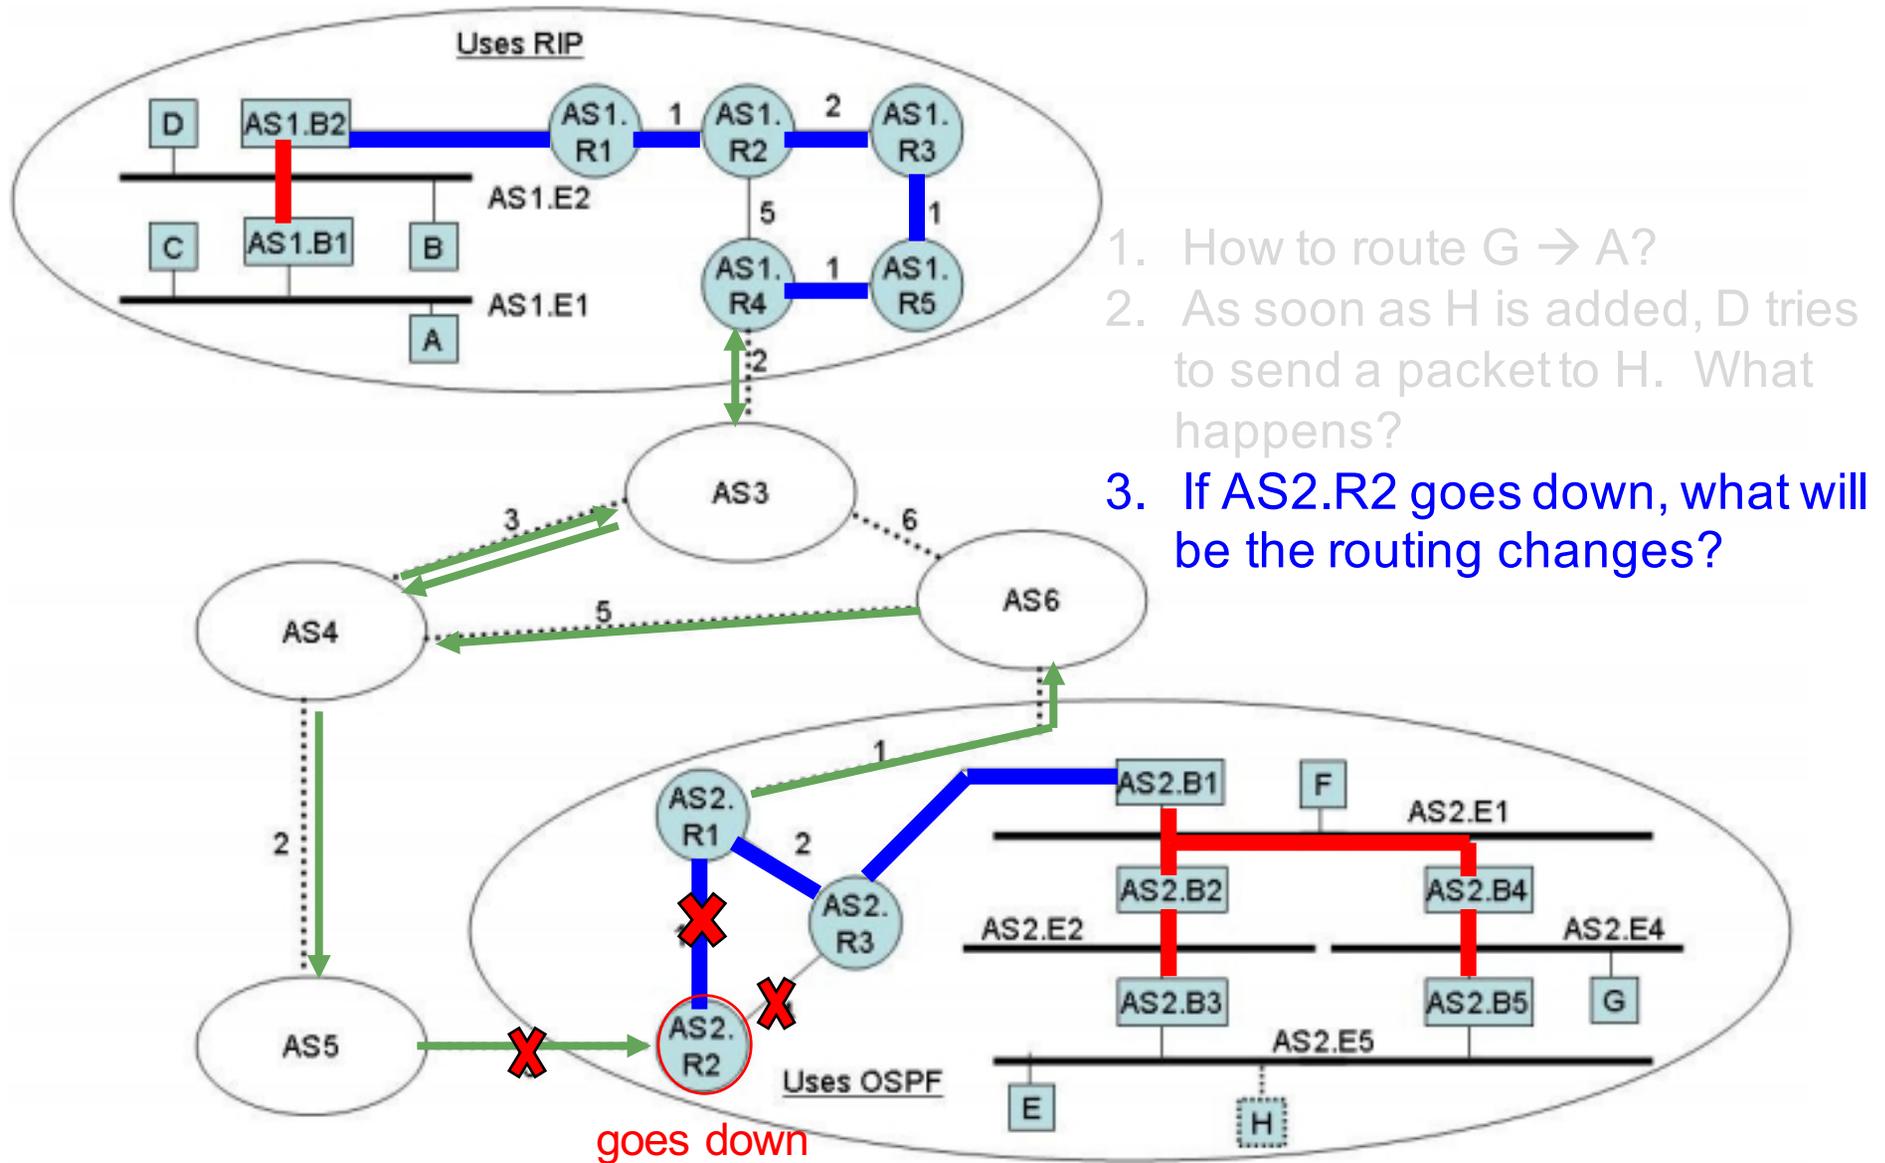


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



5. re-compute routing tables

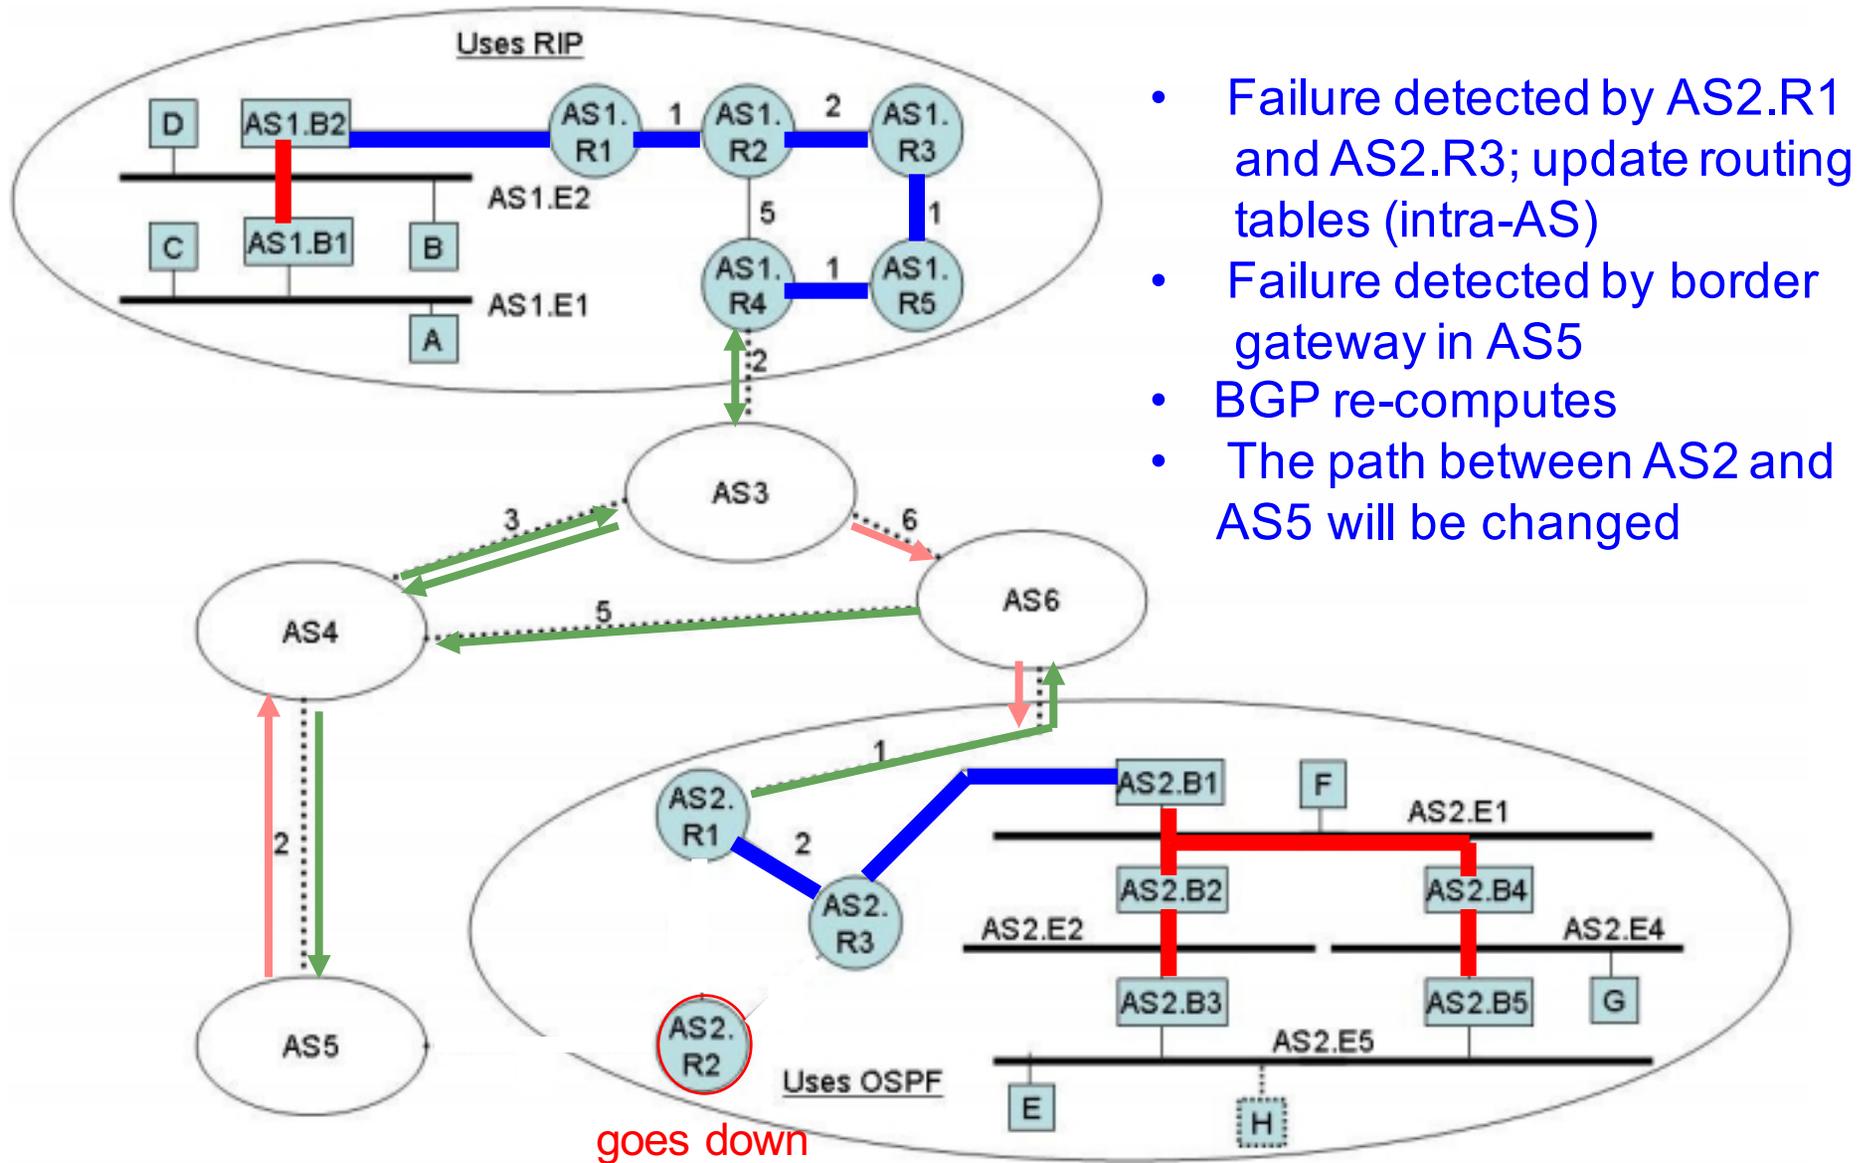


Figure 5.19: Figure for Routing Problem 3. [W&P 2010]



DHCP

- ❑ Motivations
 - Allows reuse of IP addresses
 - Enhances mobility, e.g., a laptop gets a new (local) IP address whenever it connects to a different LAN
- ❑ A DHCP server maintains a pool of available IP addresses
- ❑ A host that wants to connect to LAN requests a **dynamic** IP address of the DHCP server
- ❑ Once granted, the host periodically renew its IP
- ❑ When the host leaves and the IP address expires, the DHCP server puts the IP back into the available pool
- ❑ Host using dynamic IP cannot be a server



NAT

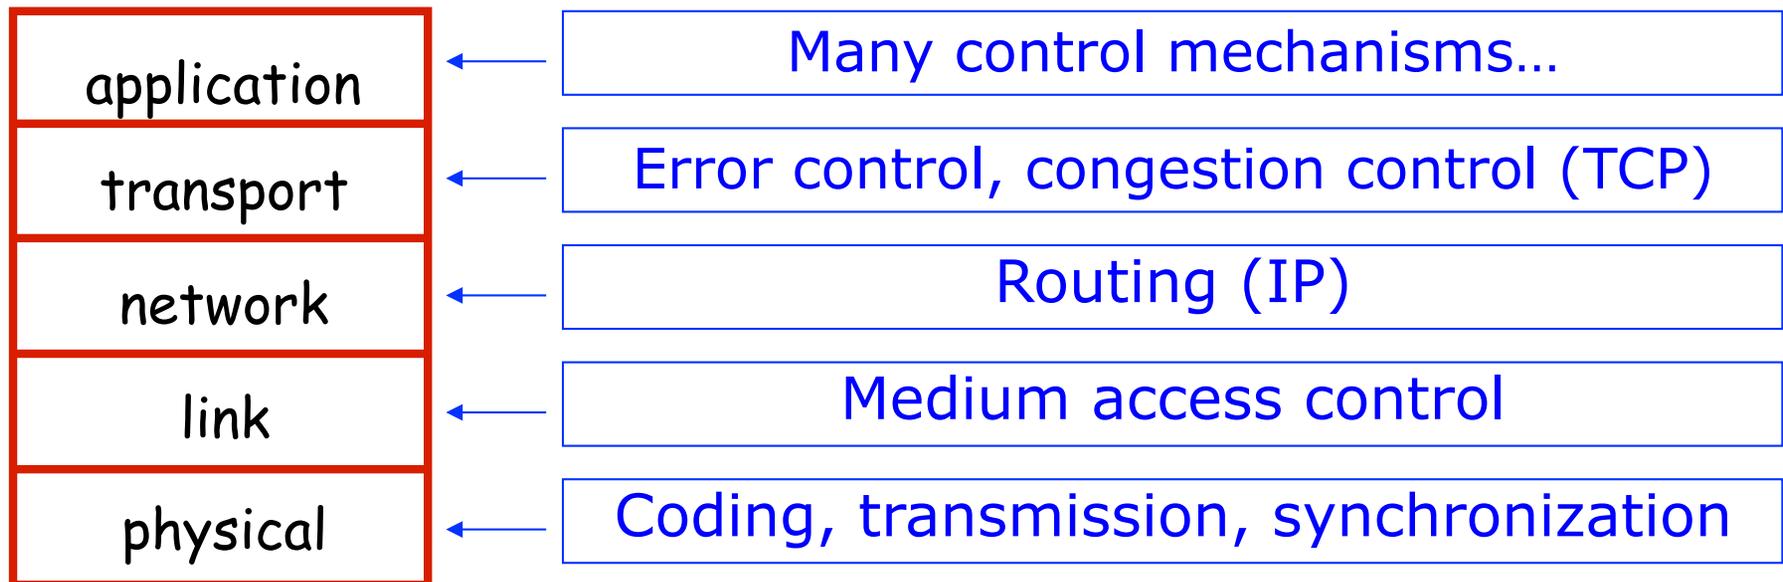
- ❑ Most home routers implement NAT (Network Address Translation)
- ❑ Motivations
 - Reuse of IP addresses
 - Security
- ❑ Devices on LAN (e.g. in home) use private IP addresses that are local in scope
- ❑ Outside world only knows the (static) IP address of the NAT box
- ❑ NAT translates to maintain end-to-end connection of a TCP flow
- ❑ How?



Protocol stack

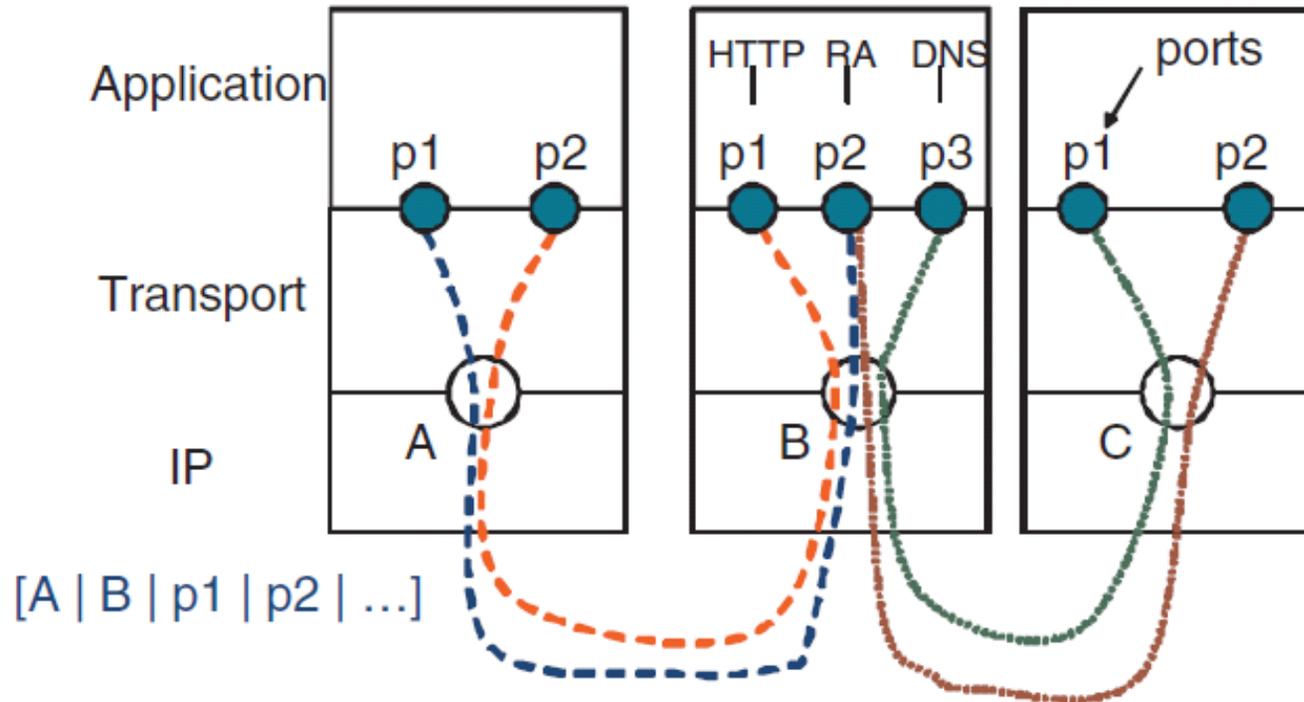
Network mechanisms implemented as protocol stack

Each layer designed separately, evolves asynchronously





Transport services



UDP

- Datagram service
- No congestion control
- No error/loss recovery
- Lightweight

TCP

- Connection oriented service
- Congestion control
- Error/loss recovery
- Heavyweight

[source IP | destination IP | source port | destination port | ... | data].



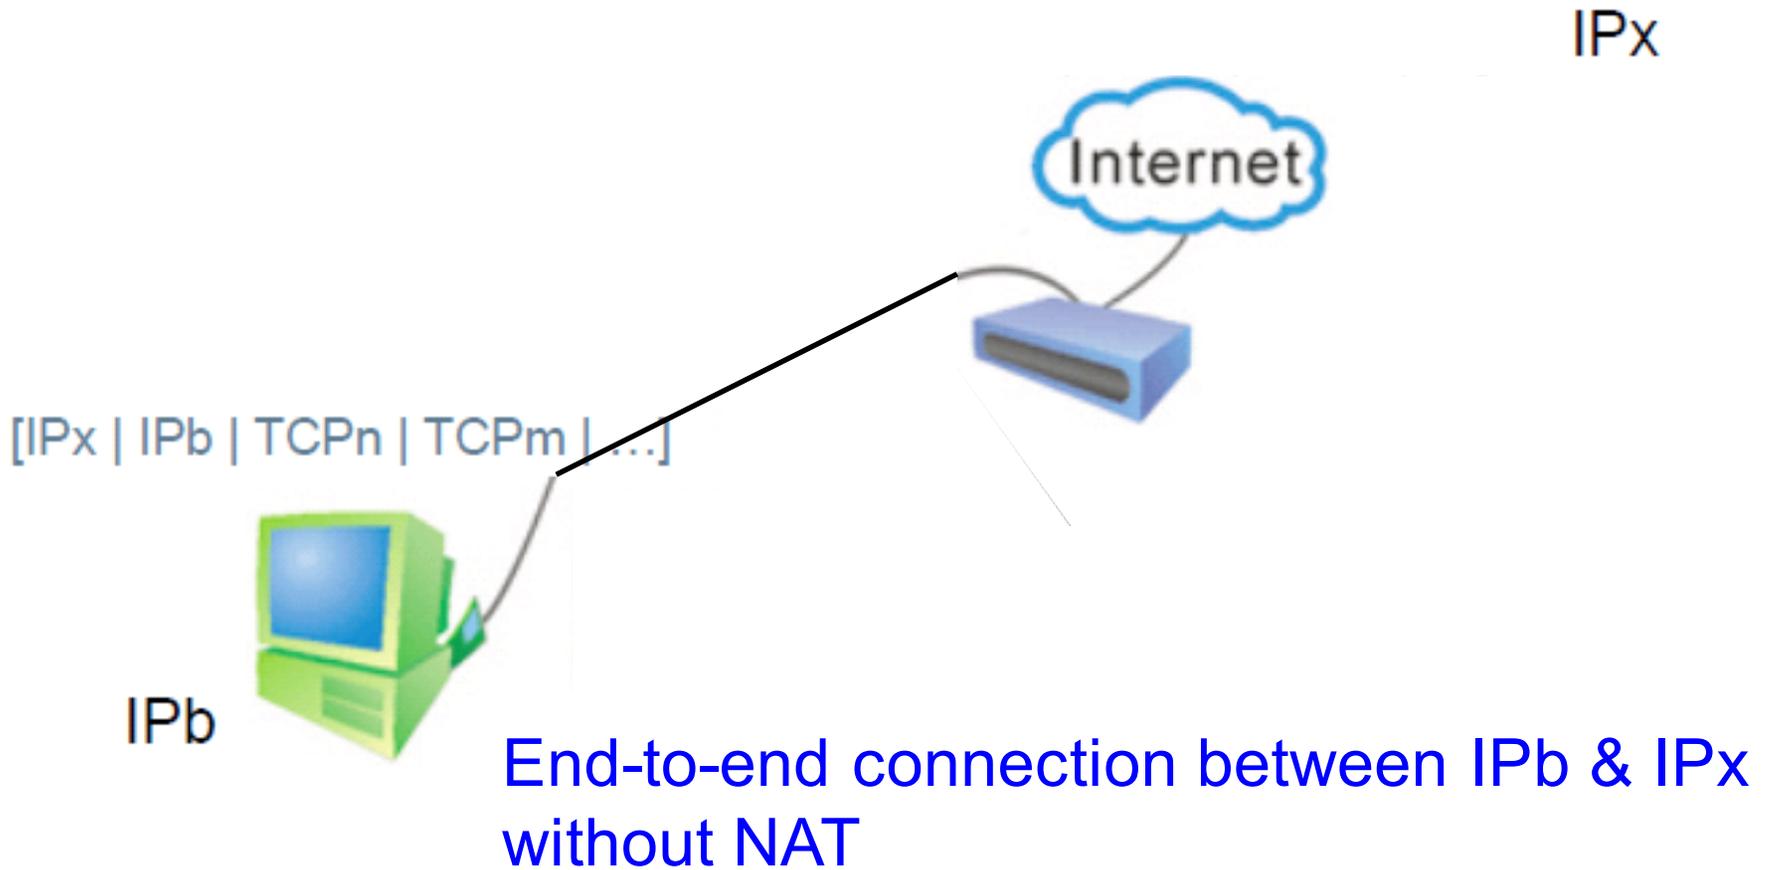
NAT

Key idea

- A flow is specified by IP address + TCP port #
- NAT uses TCP port# to hide private IP



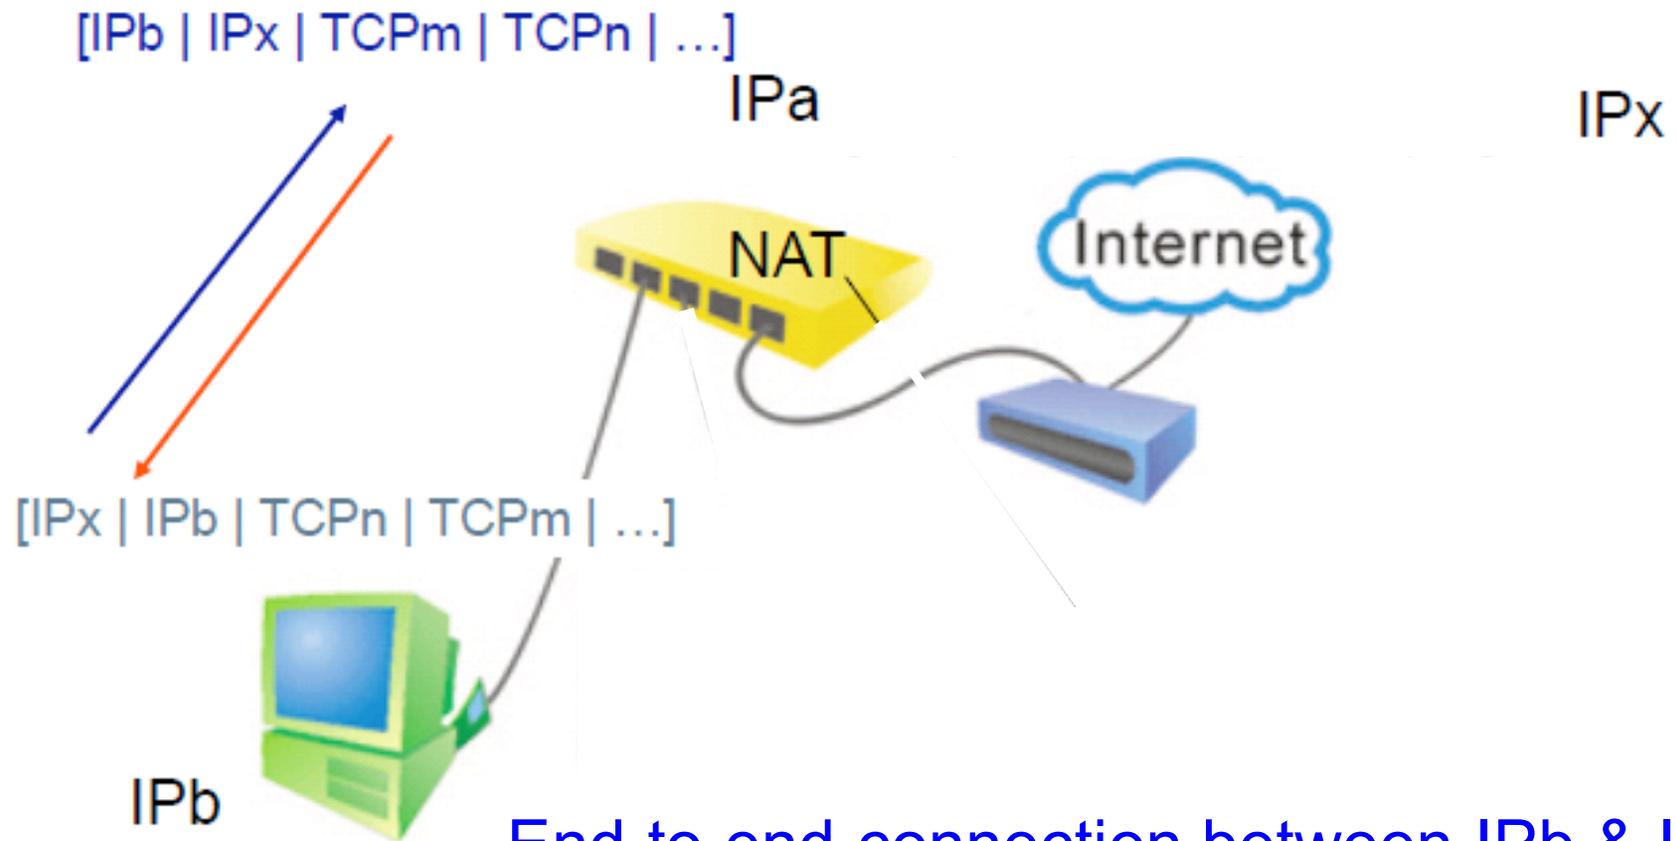
NAT



[source IP | destination IP | source port | destination port | ... | data]



NAT

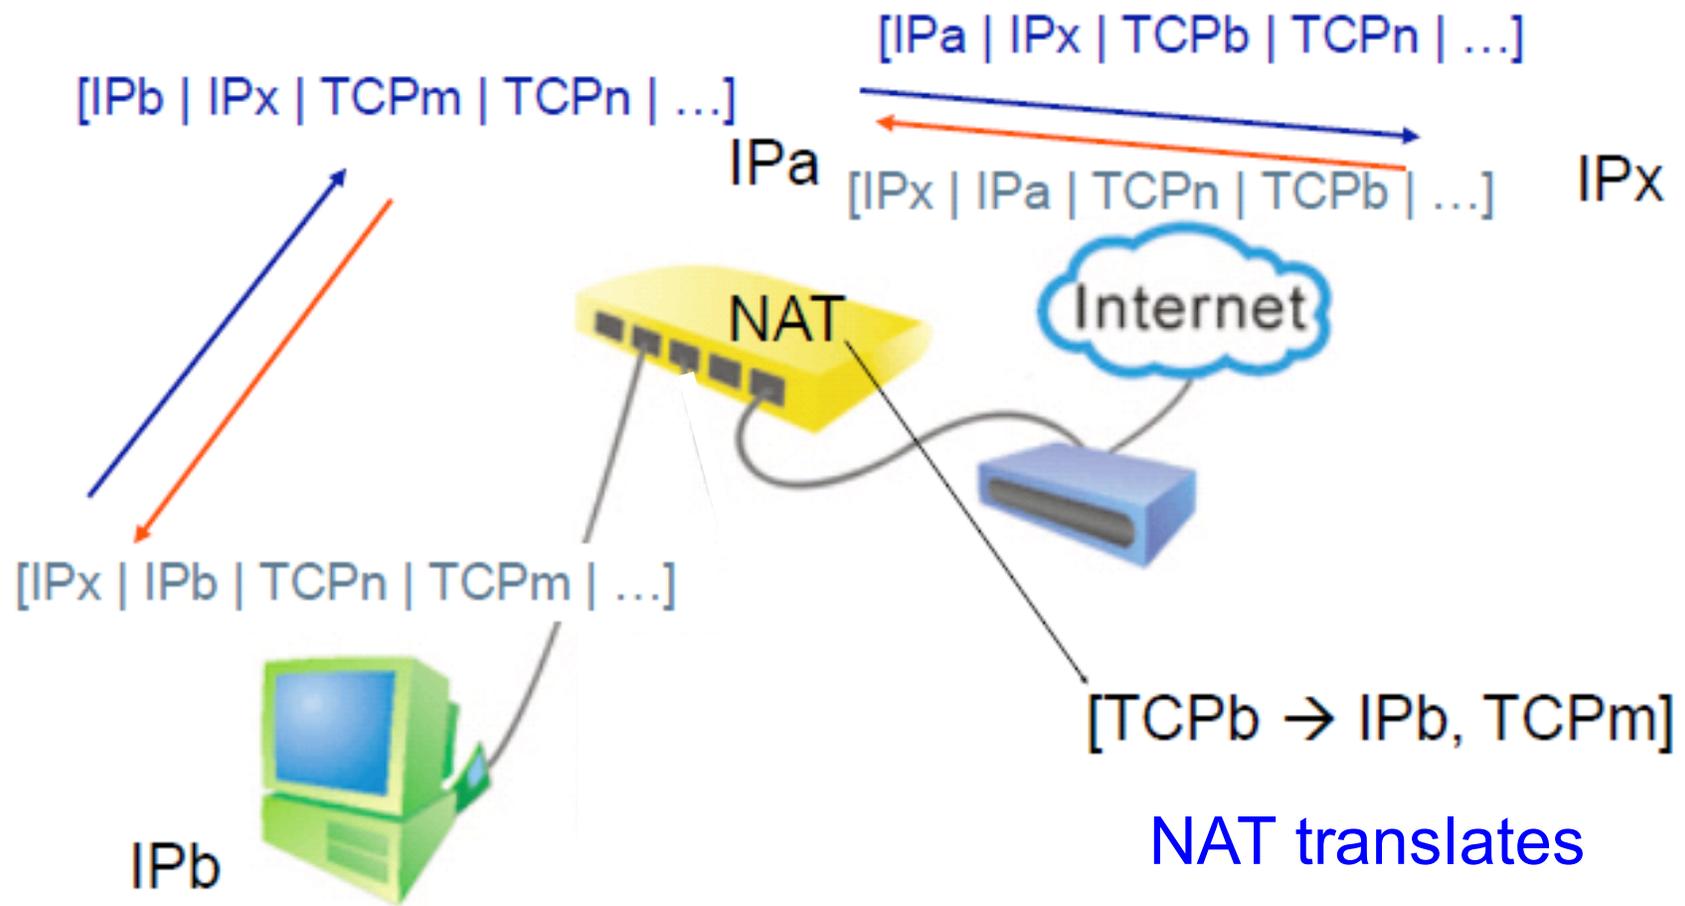


End-to-end connection between IPb & IPx with NAT

[source IP | destination IP | source port | destination port | ... | data]



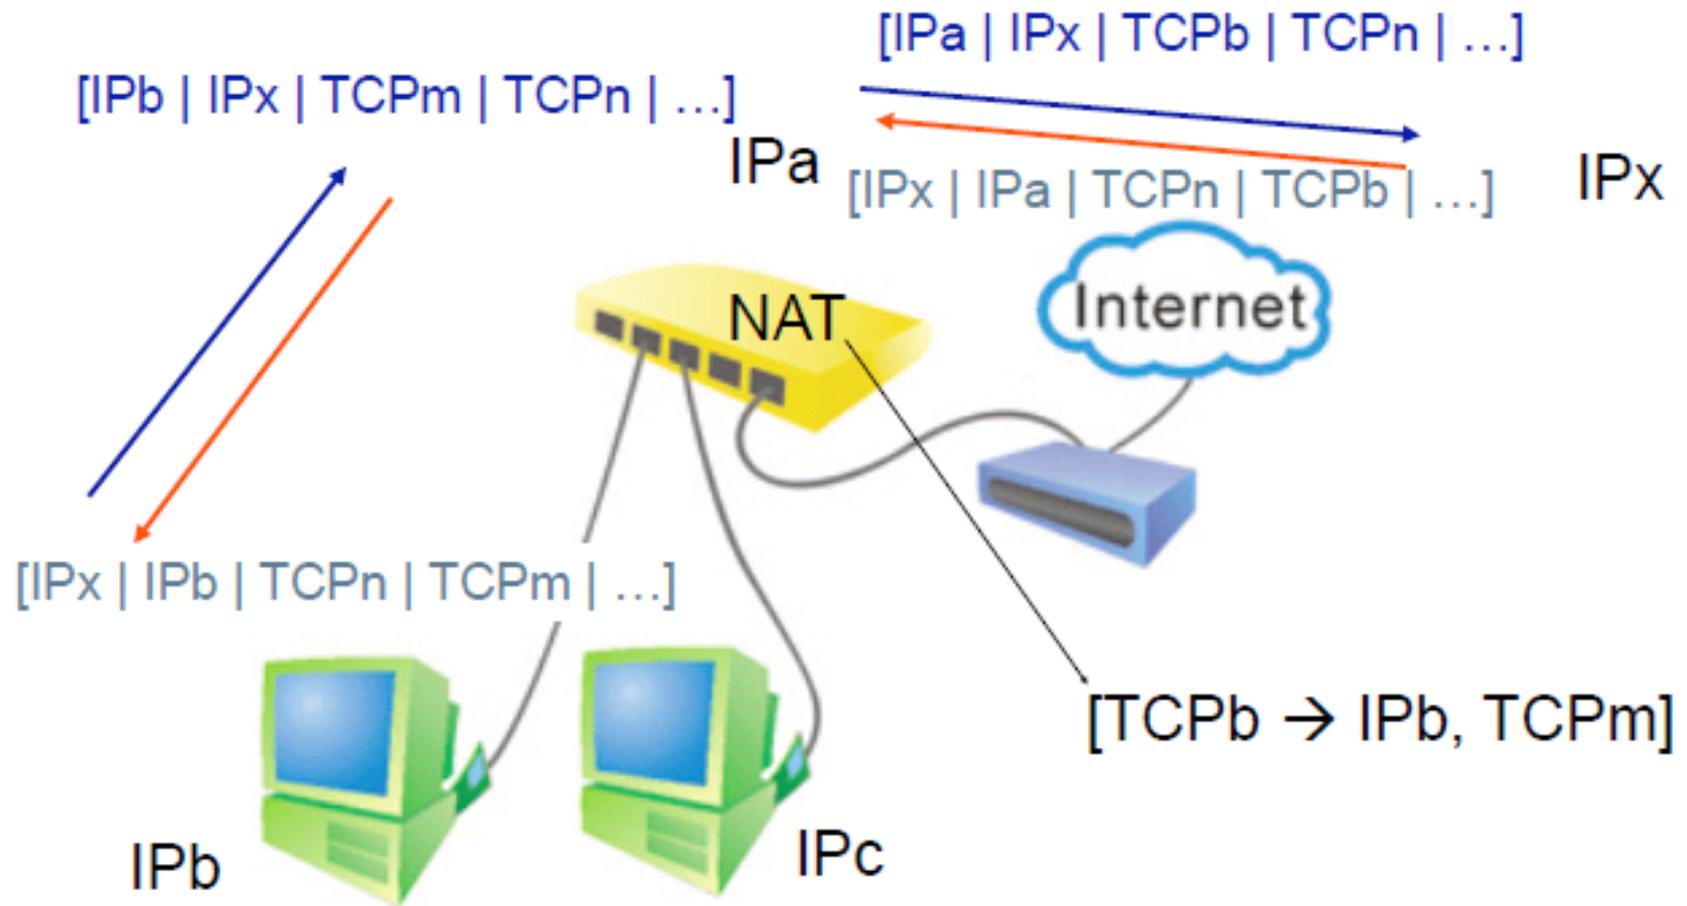
NAT



[source IP | destination IP | source port | destination port | \dots | data].



NAT





NAT

- ❑ Most home routers implement NAT (Network Address Translation)
- ❑ Motivations
 - Reuse of IP addresses
 - Security
- ❑ Devices on LAN (e.g. in home) use private IP addresses that are local in scope
- ❑ Outside world only knows the (static) IP address of the NAT box
- ❑ NAT translates to maintain end-to-end connection of a TCP flow
- ❑ Only home device can initiate connection to outside world (server)