Computer Networks X_400487

Lecture 5

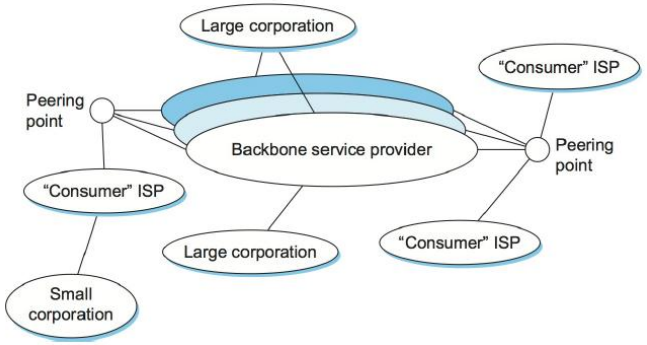

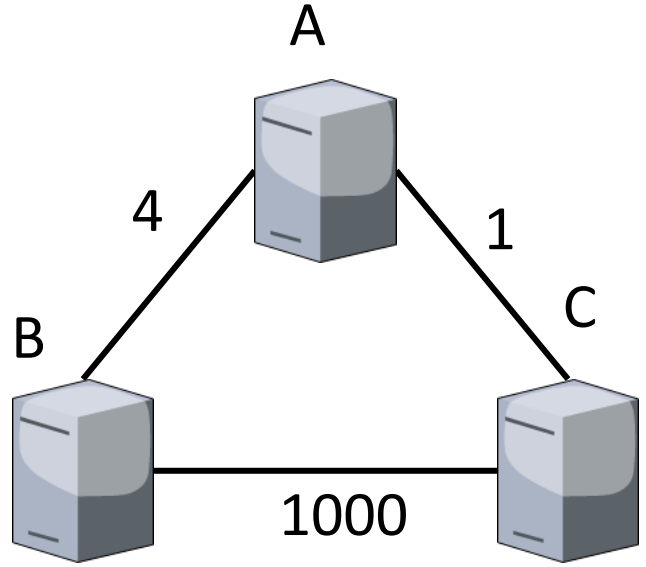

Chapter 5: The Network Layer—Part 1

Lecturer: Jesse Donkervliet

VU VRIJE UNIVERSITEIT AMSTERDAM

Roadmap of the Computer Networks Course

Lets distributed applications communicate

Application layer

Sends segments from one **process** to another within and across networks

Transport layer

Sends packets from one **machine** to another within and across networks

Network Layer

You are here

Midterm

Sends frames from one machine to another over a single link / local network

Data link layer

Medium Access Control

Sends bits over a physical medium

Physical layer

Basic principles and course overview

Introduction

Recap Data Link Layer

Data link layer (partially) answers the following questions:

1. How create frames from bits/bytes?
2. How to detect/correct transmission errors?
3. How to efficiently multiplex frames from multiple stations over a single channel? MAC sublayer Q: What kind of efficiency?

Recap Data Link Layer

Data link layer (partially) answers the following questions:

1. How create frames from bits/bytes?
2. How to detect/correct transmission errors?
3. How to efficiently multiplex frames from multiple stations over a single channel? MAC sublayer Q: What kind of efficiency?

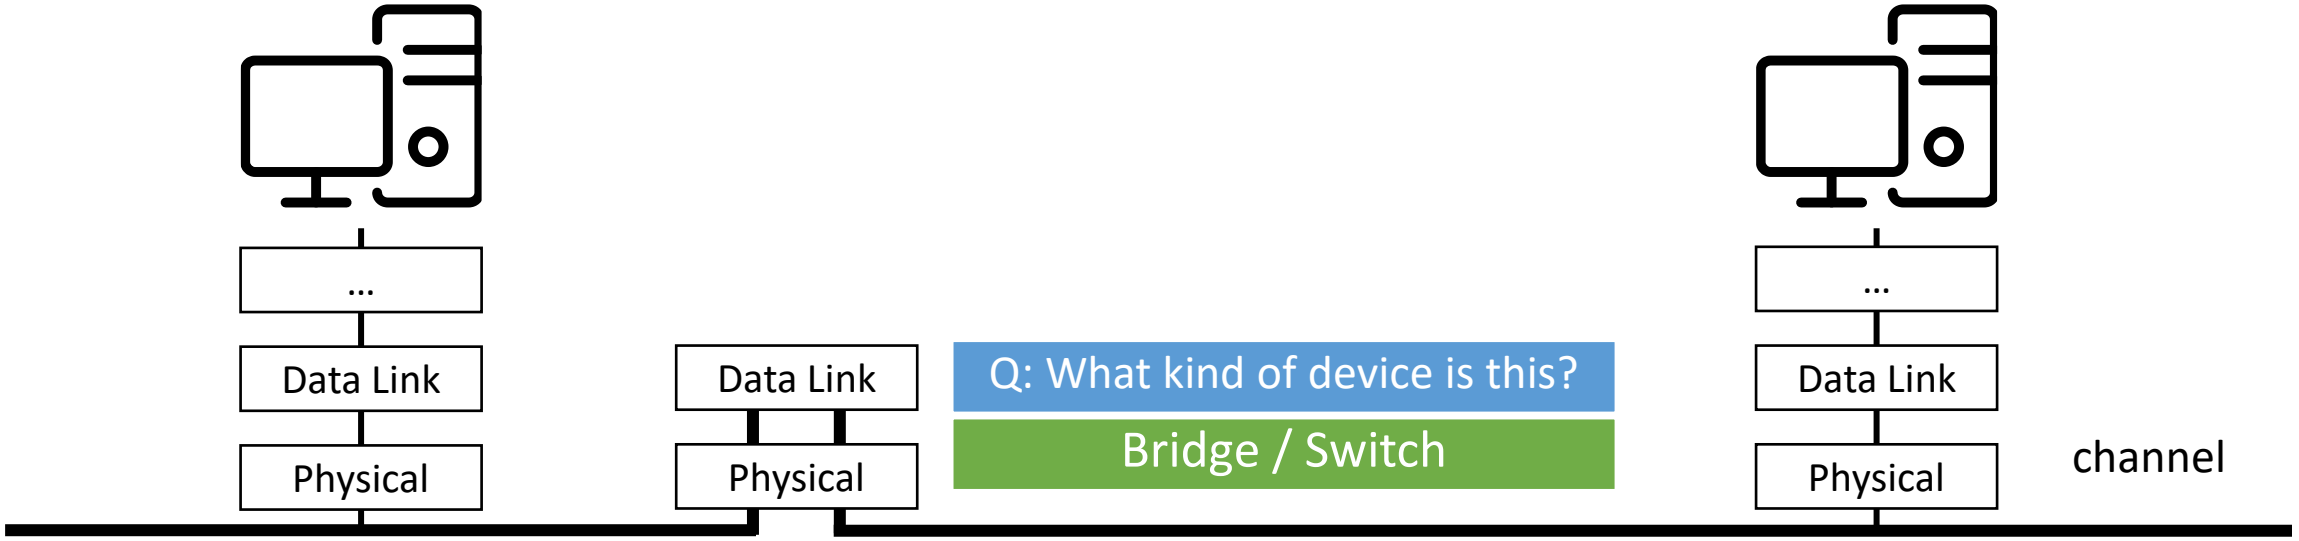

Image source: NASA

What Else Do We Need?

Data link layer not enough for a world-wide *internet* (=network of networks)

1. Switches not built for large networks (e.g., hash table size)

   Hashmap of all addresses requires very large memory. Searching memory reduces performance

2. Protocols tailored to physical medium:
   bad idea to use one such protocol for all types of networks

Q: How to solve this?

Add new layer of abstraction: the Network Layer

The Network Layer

Lowest layer concerned with "end-to-end" delivery

The network layer (partially) answers the following questions:

1. How to get from ***A*** to ***B***?
2. How to connect multiple networks?
3. How to manage network resources?
    1. How to prevent *network congestion*?
    2. How to provide *Quality of Service*?

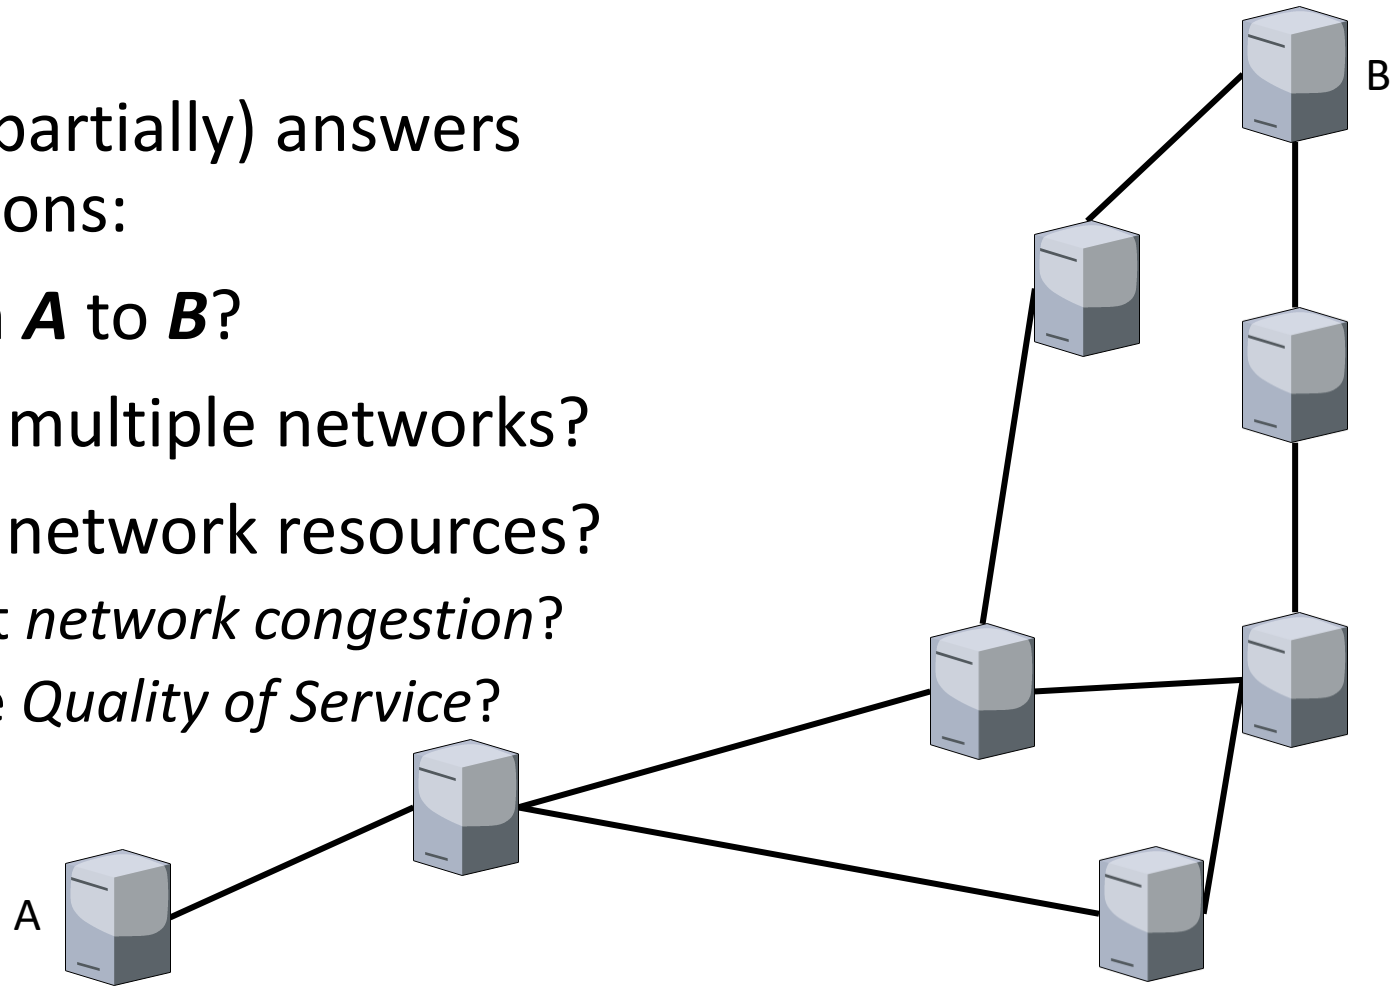

The Network Layer

TCP/IP form the “narrow waist” of the Internet

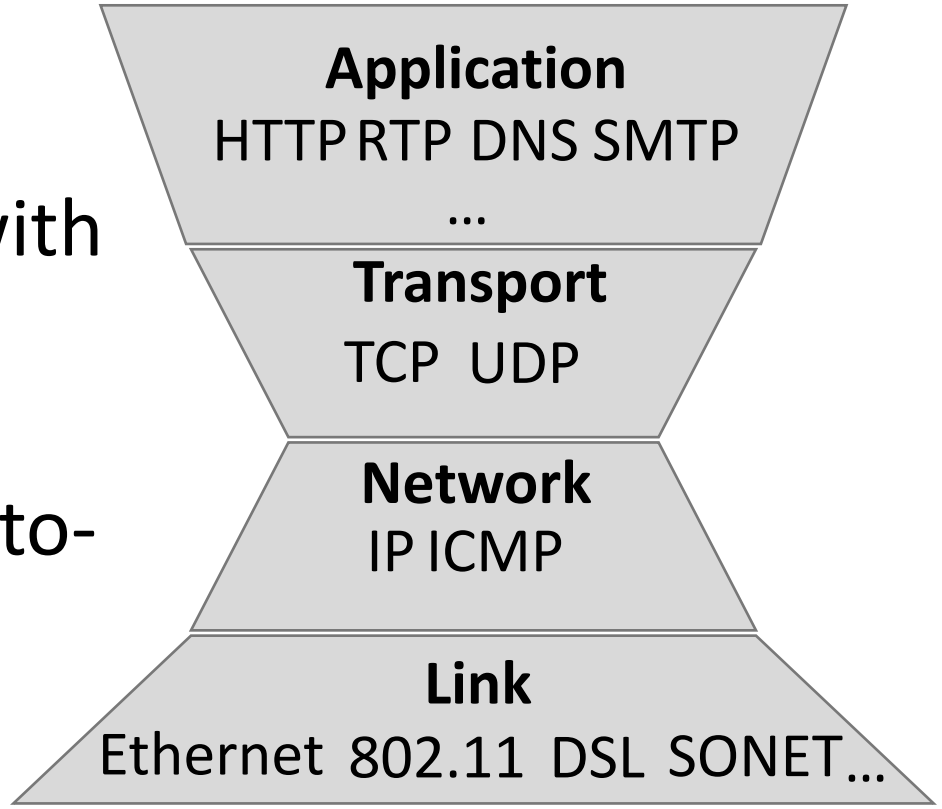

Heart of the networking model, together with the transport layer (next week)

1. Layers below it do not know about end-to-end delivery
2. Layers above it do not know about the topology of the network
3. Both layers above and below do not know about routing

*Although switching is similar to routing…

The Network Layer

Network layer provides new features through *new protocols* and a *new device* (the router)

Q: Why have Data Link and Physical Layer twice?

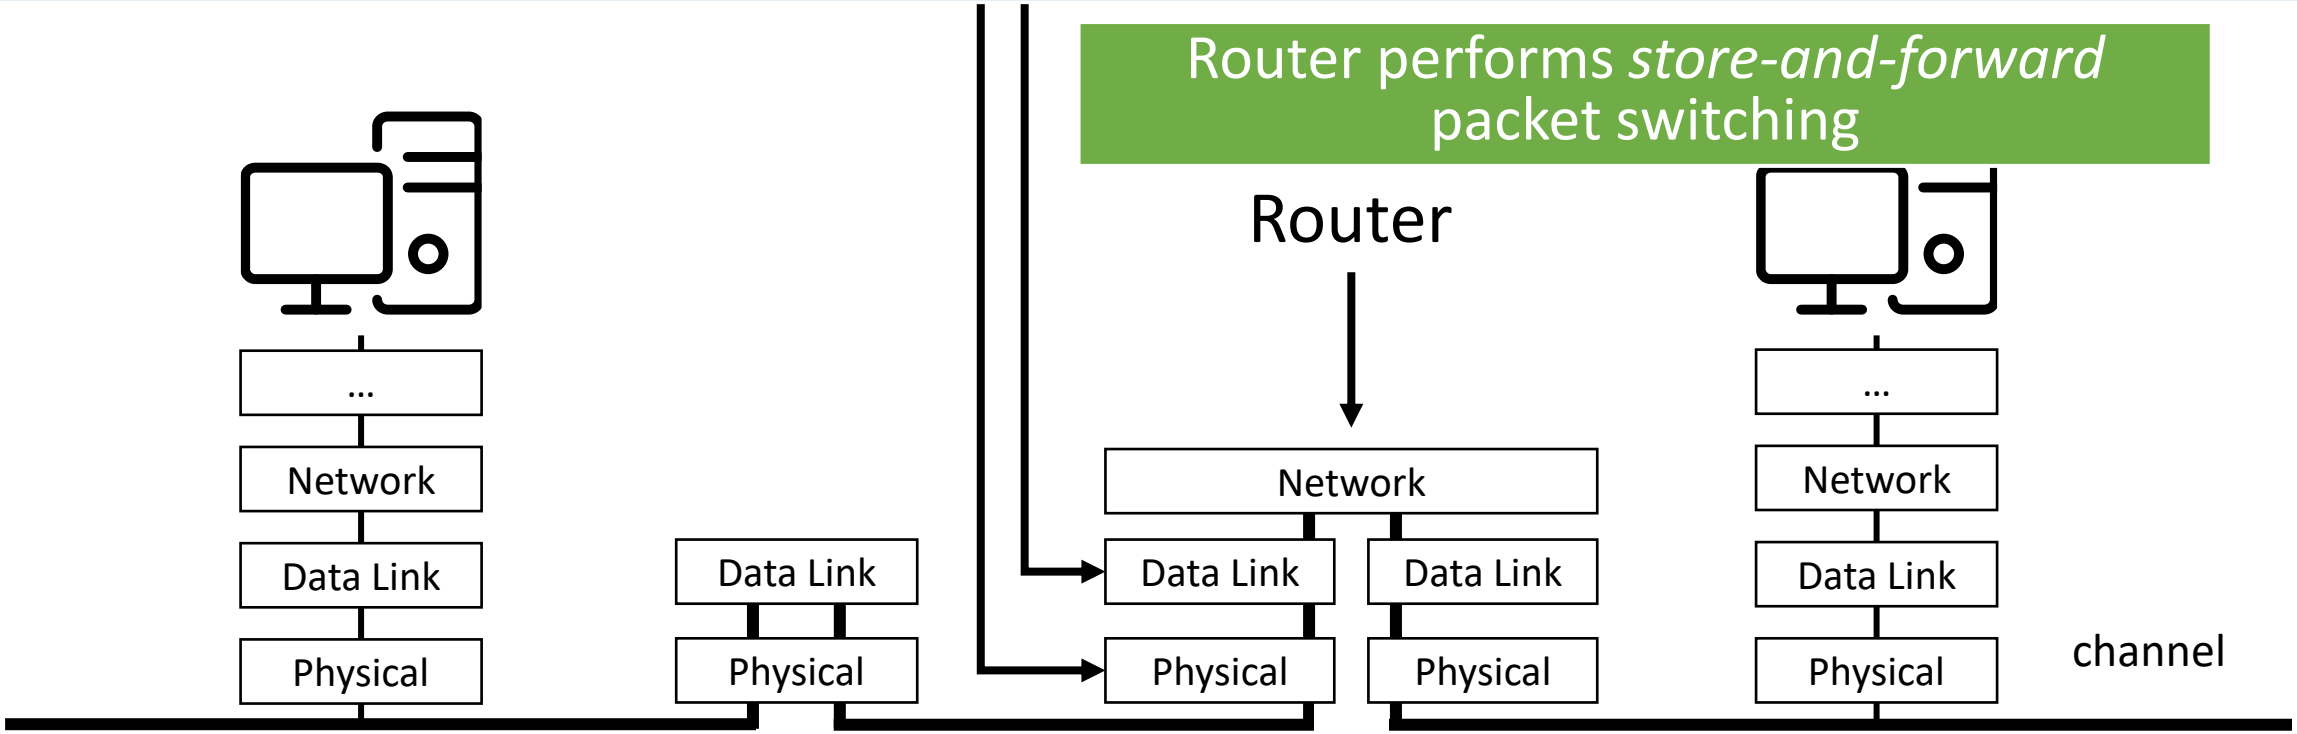

The Network Layer

Lowest layer concerned with "end-to-end" delivery

The network layer (partially) answers the following questions:

Today's lecture

1. How to get from ***A*** to ***B***?
2. How to connect multiple networks?

Next lecture

3. How to manage network resources?
    1. How to prevent *network congestion*?
    2. How to provide *Quality of Service*?

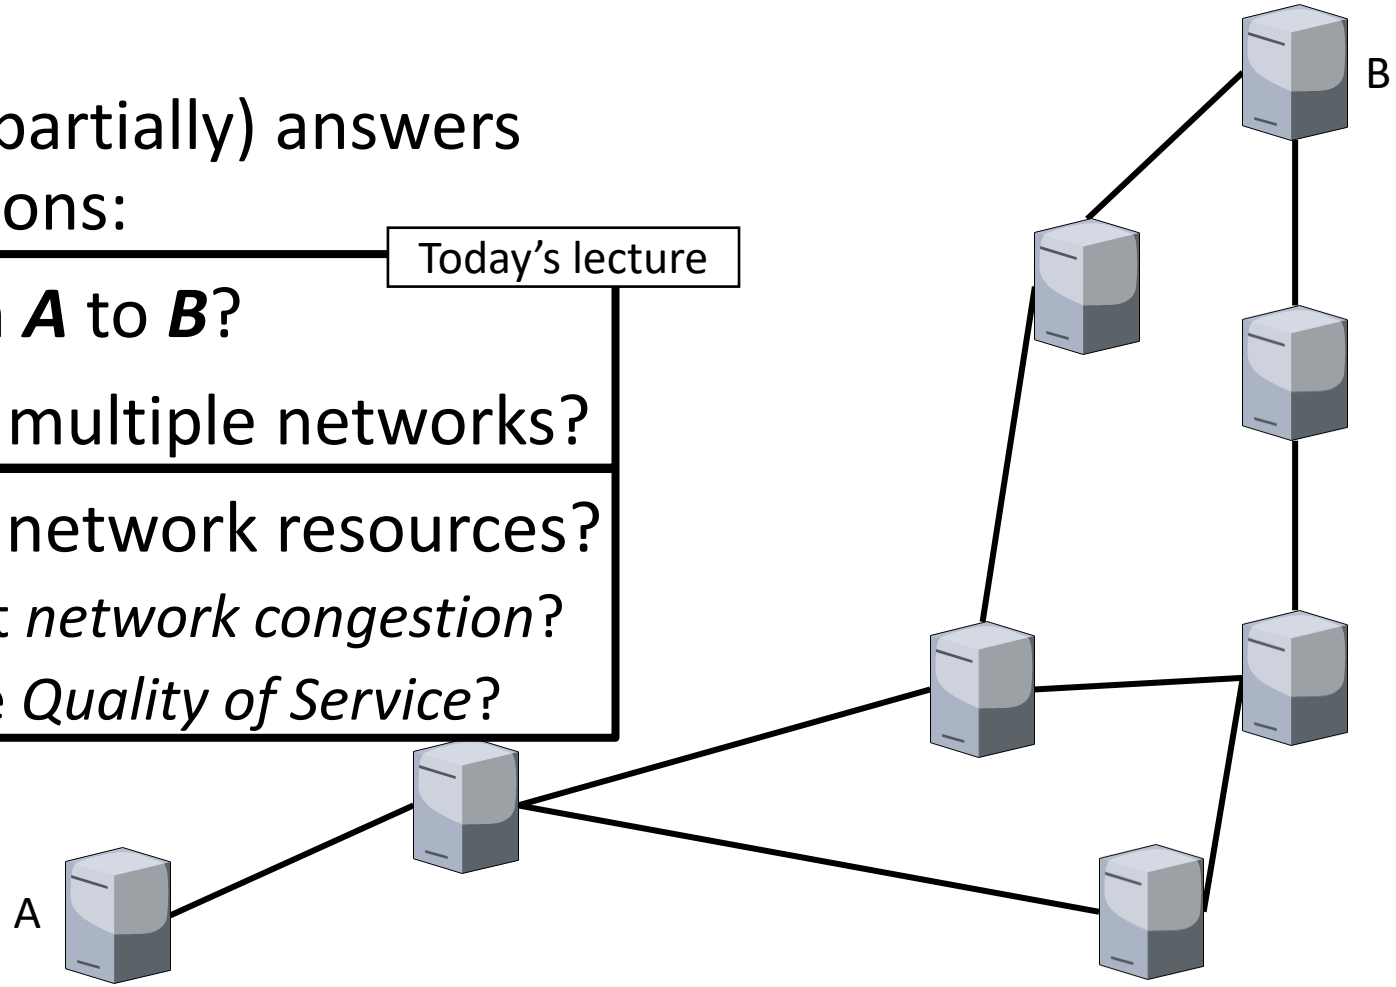

Today's Lecture

1. **Routing Algorithms**
2. Internetworking

Routing

Finding a path through a network

Important Properties for Routing

1. Correctness
2. Simplicity
3. Robustness
4. Stability
5. Fairness
6. Efficiency

These properties can (and often do) conflict with each other.

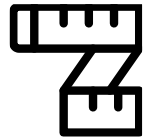

How to find a route?

We will look at three key ideas:

1. Distance vector routing
2. Link state routing
3. Hierarchical routing

Routing tables

For each packet, we need to know on which link to forward it.

For this we use a routing table

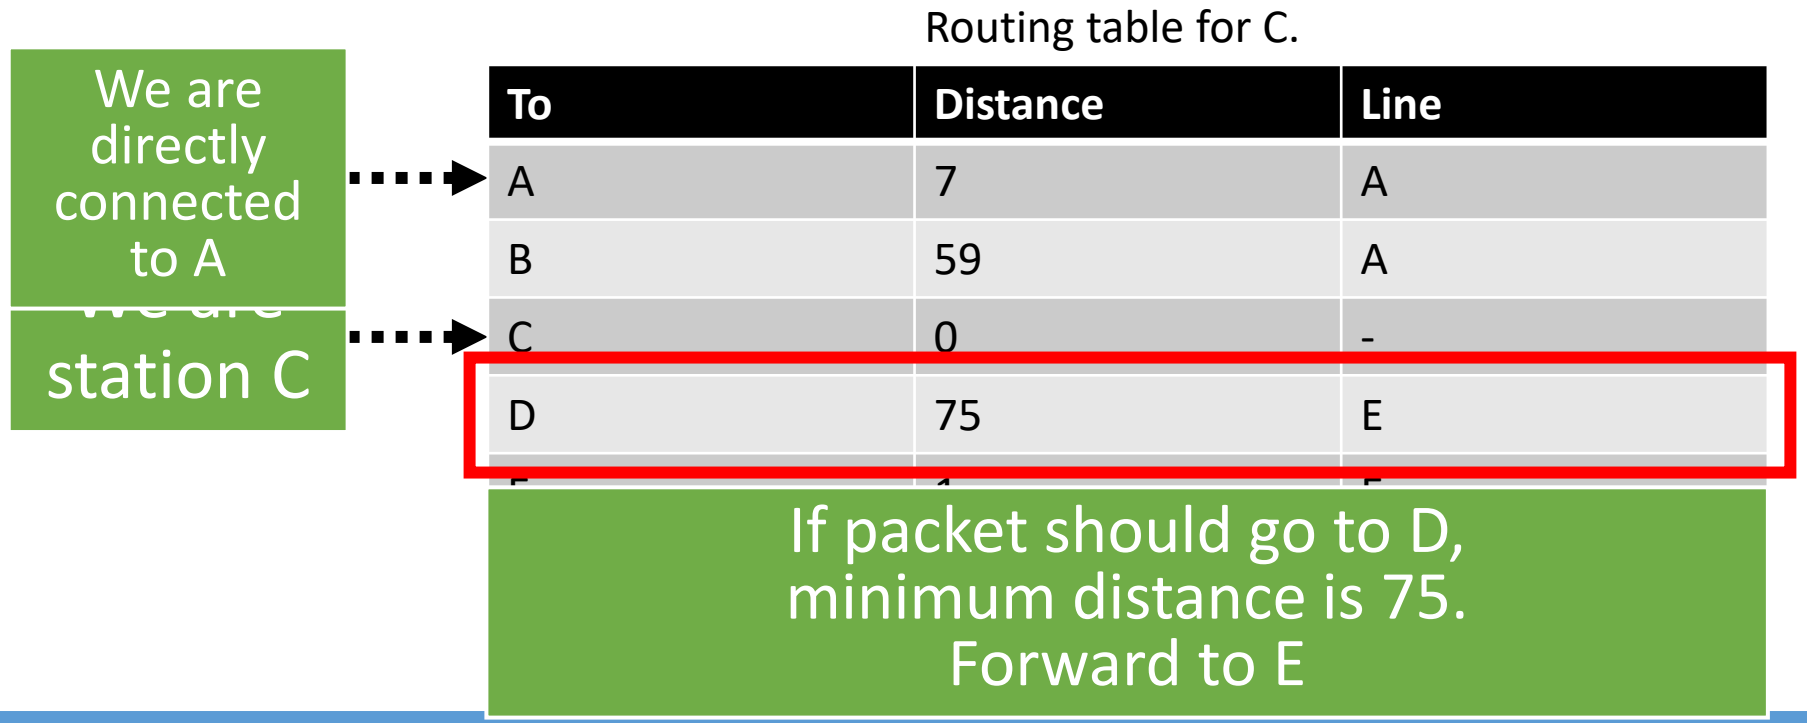

Routing tables

For each packet, we need to know on which link to forward it

For this we use a routing table

We are **indirectly** connected to B

Routing table for C.

To	Distance	Line
A	7	A
B	59	A
C	0	-
D	75	E
E	1	E
F	103	E

Distance Vector Routing

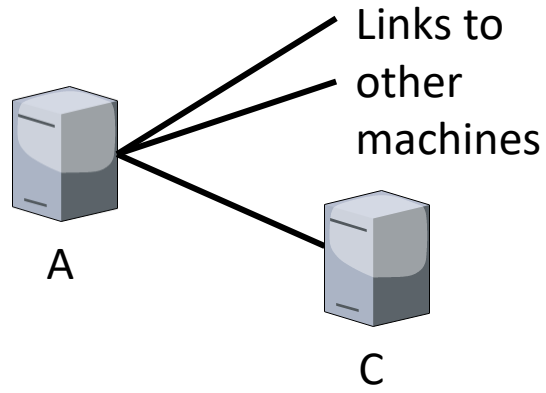

1. Send your *distance vector* to your neighbors
2. Update your *routing table* based on incoming *distance vectors*

Distance from A to B is 1

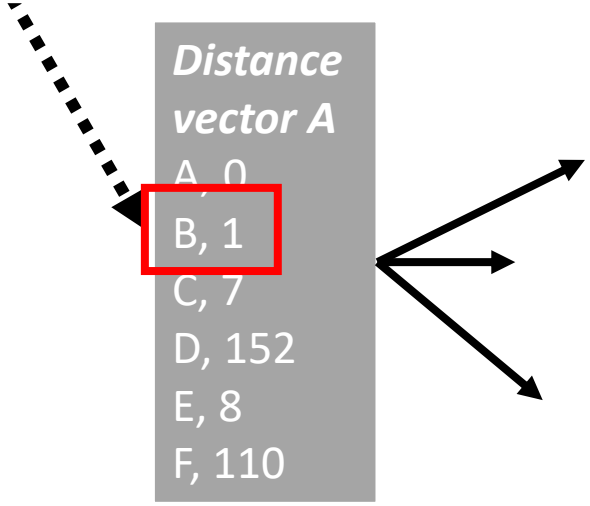

Distance from C to A is 7

To	Distance	Line
A	7	A
B	59	A
C	0	-
D	75	
E	1	
F	103	

Q: How to update our routing table?

Distance Vector Routing

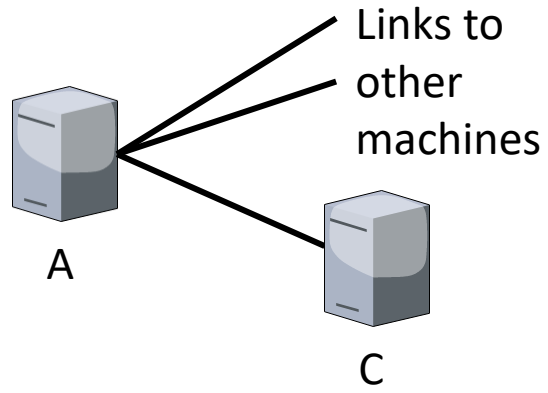

1. Send your *distance vector* to your neighbors
2. Update your *routing table* based on incoming *distance vectors*

Distance vector A
A, 0
B, 1
C, 7
D, 152
E, 8
F, 110

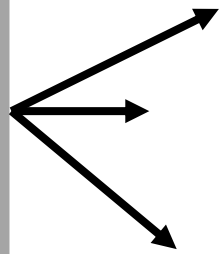

Routing table for C.

To	Distance	Line
A	7	A
B	**59 → 8**	A
C	0	-
D		
E		
F		

Q: What can go wrong when using this algorithm?

Distance Vector Routing
Good News Propagation

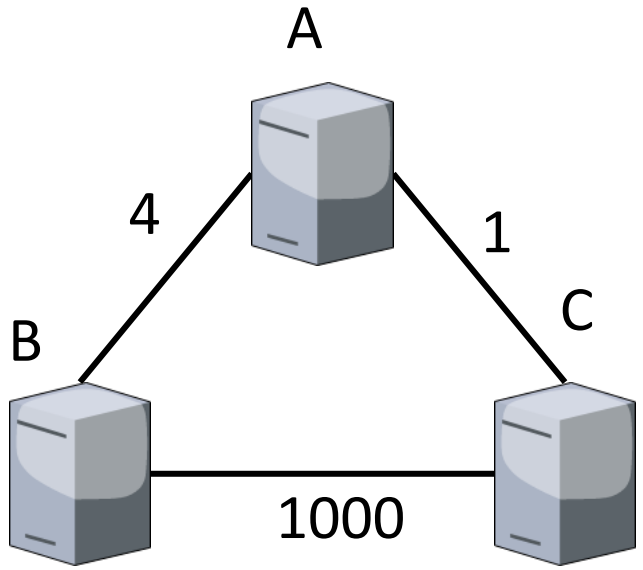

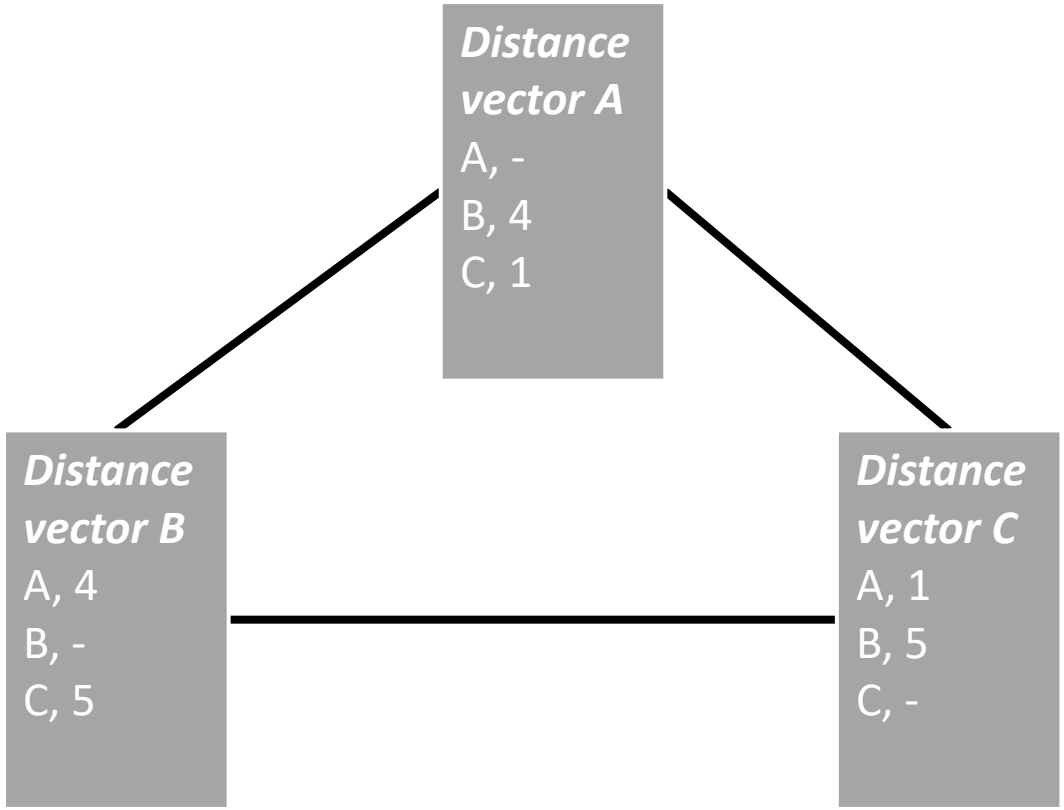

Distance Vector Routing
Good News Propagation

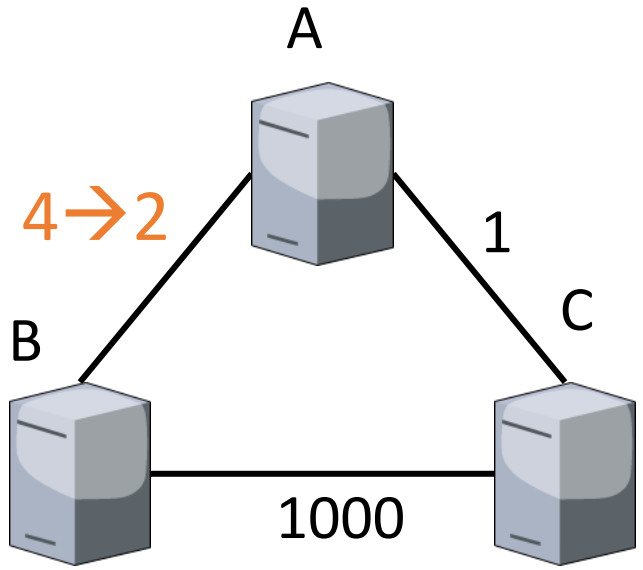

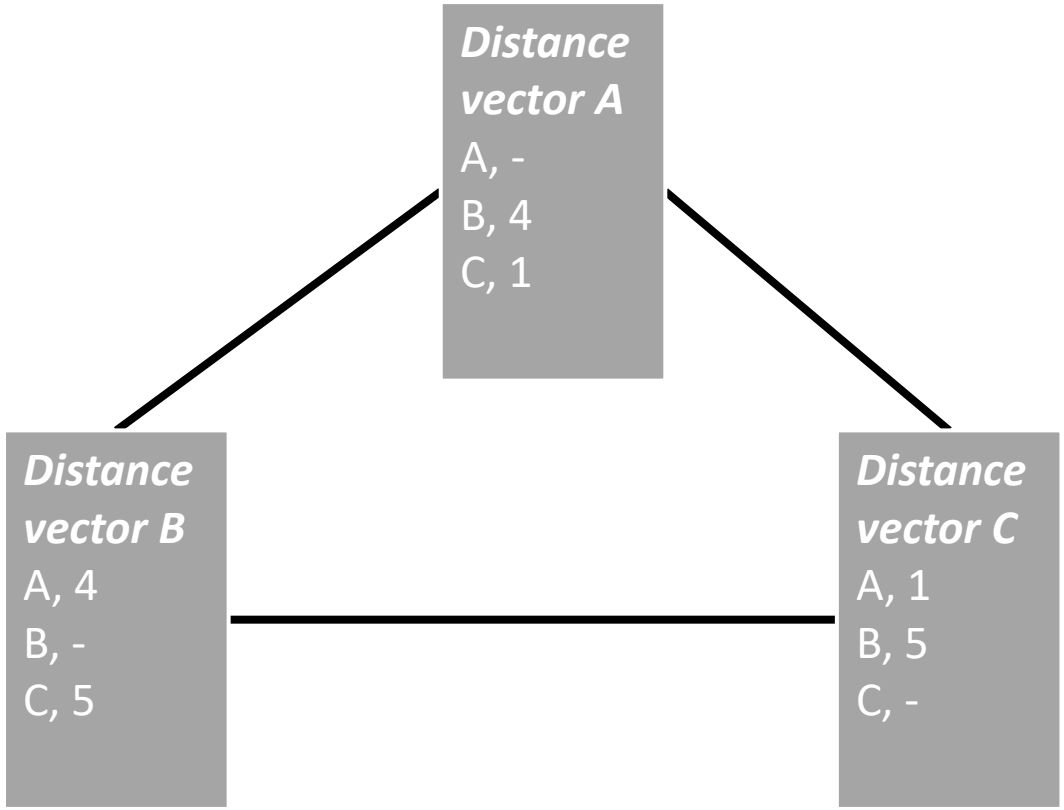

Distance Vector Routing
Good News Propagation

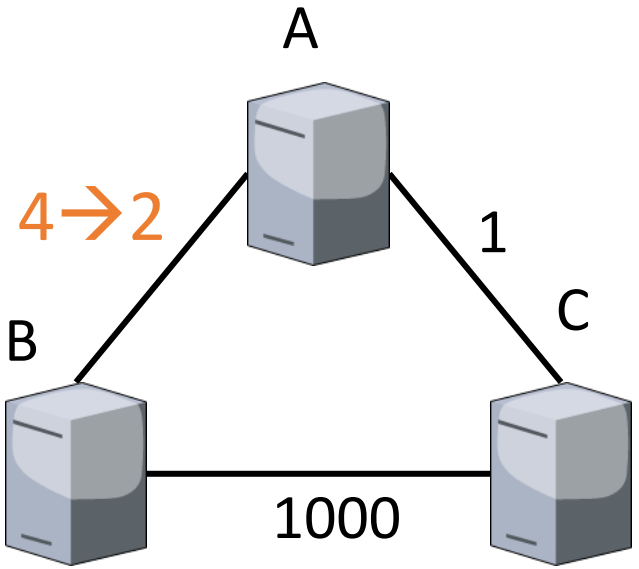

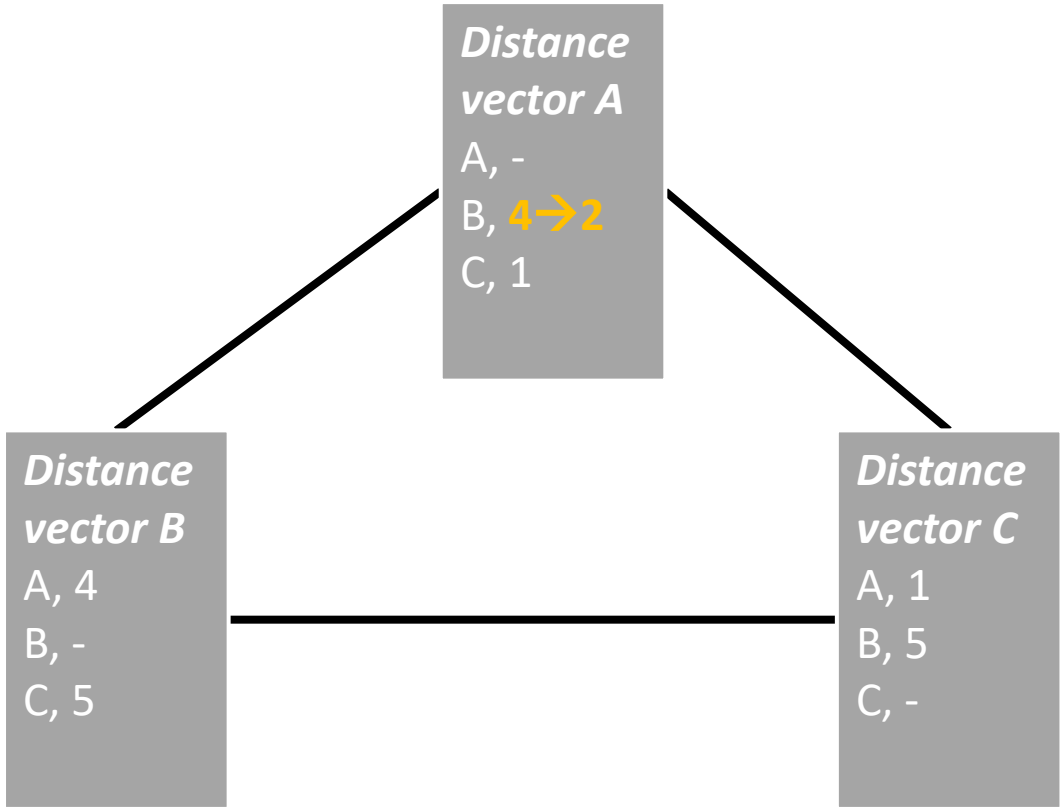

Distance Vector Routing
Good News Propagation

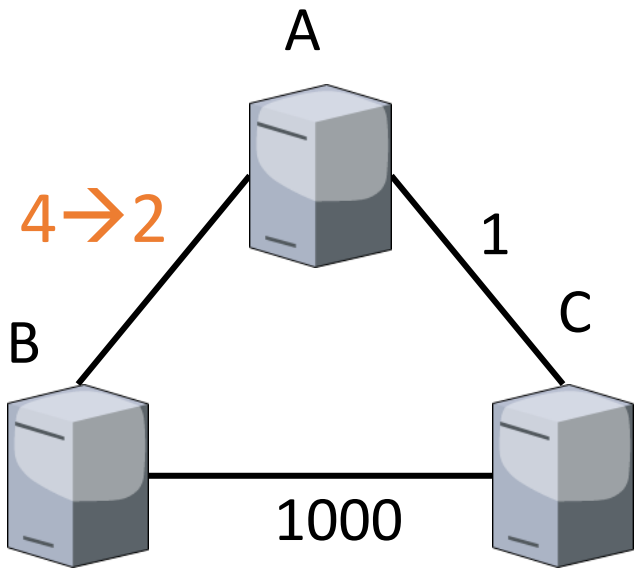

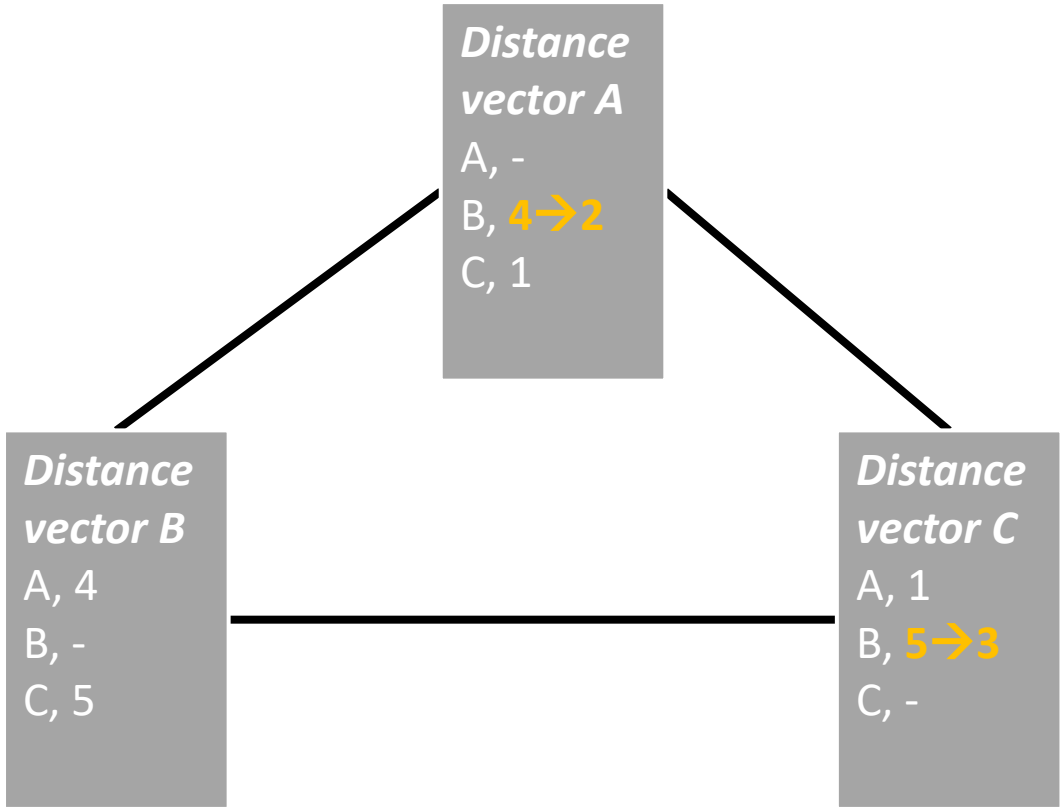

Distance Vector Routing
Bad News Propagation

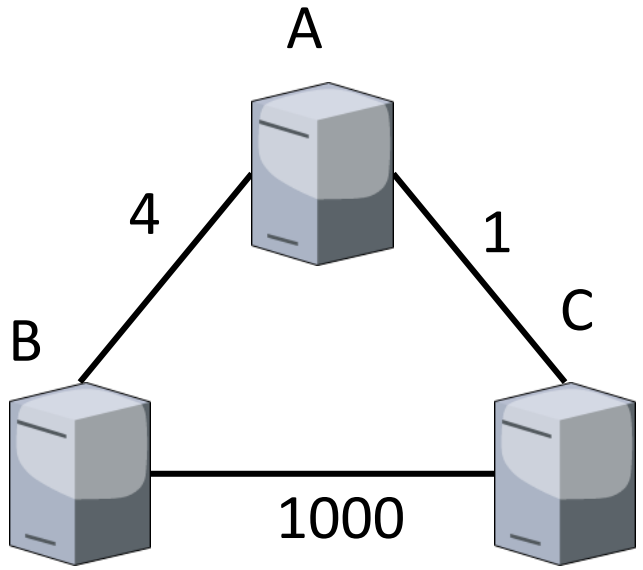

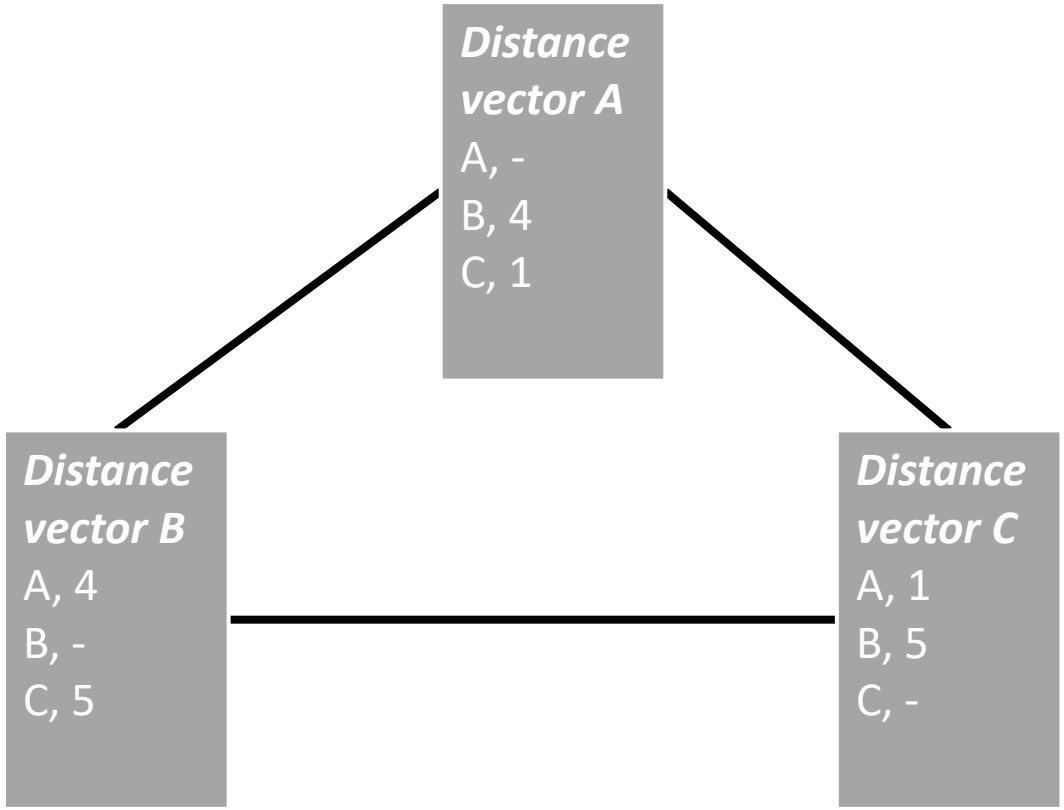

Distance Vector Routing
Bad News Propagation

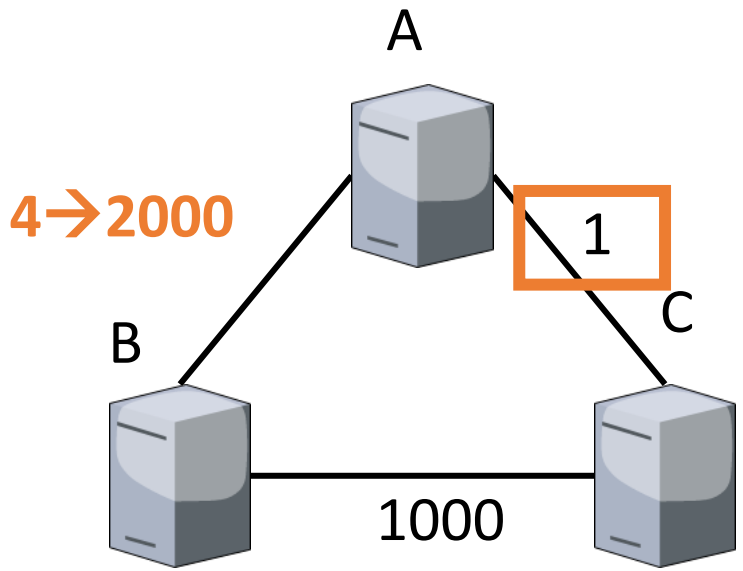

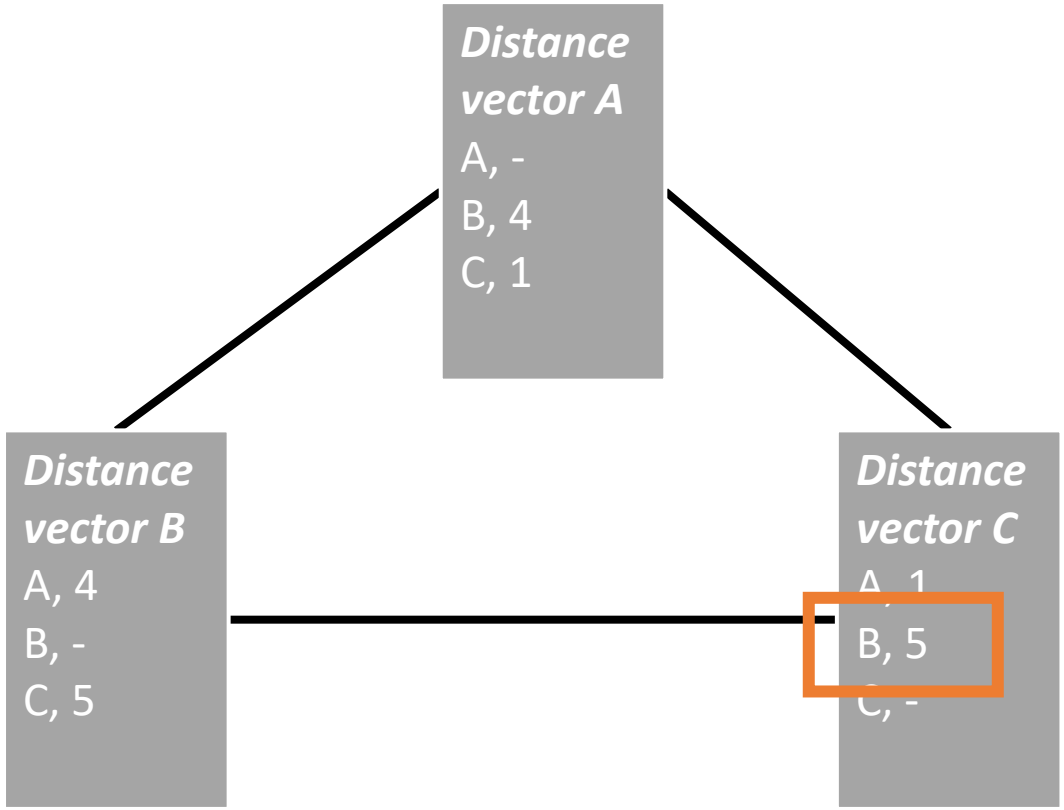

Distance Vector Routing
Bad News Propagation

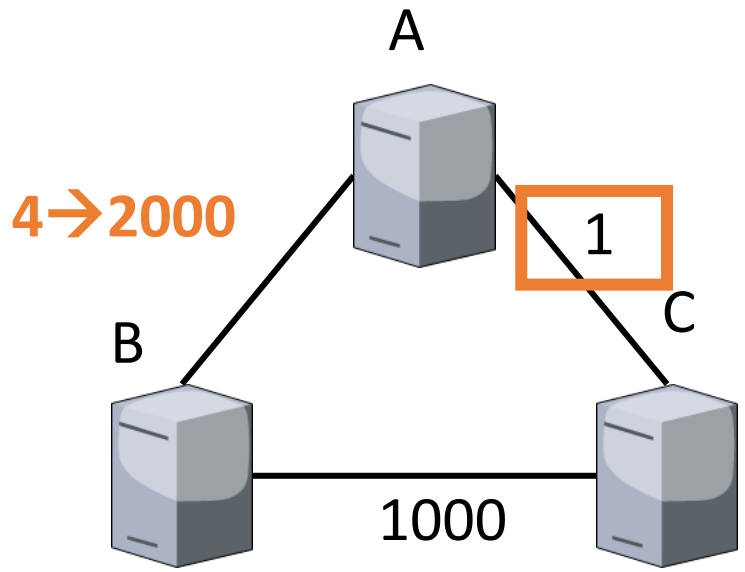

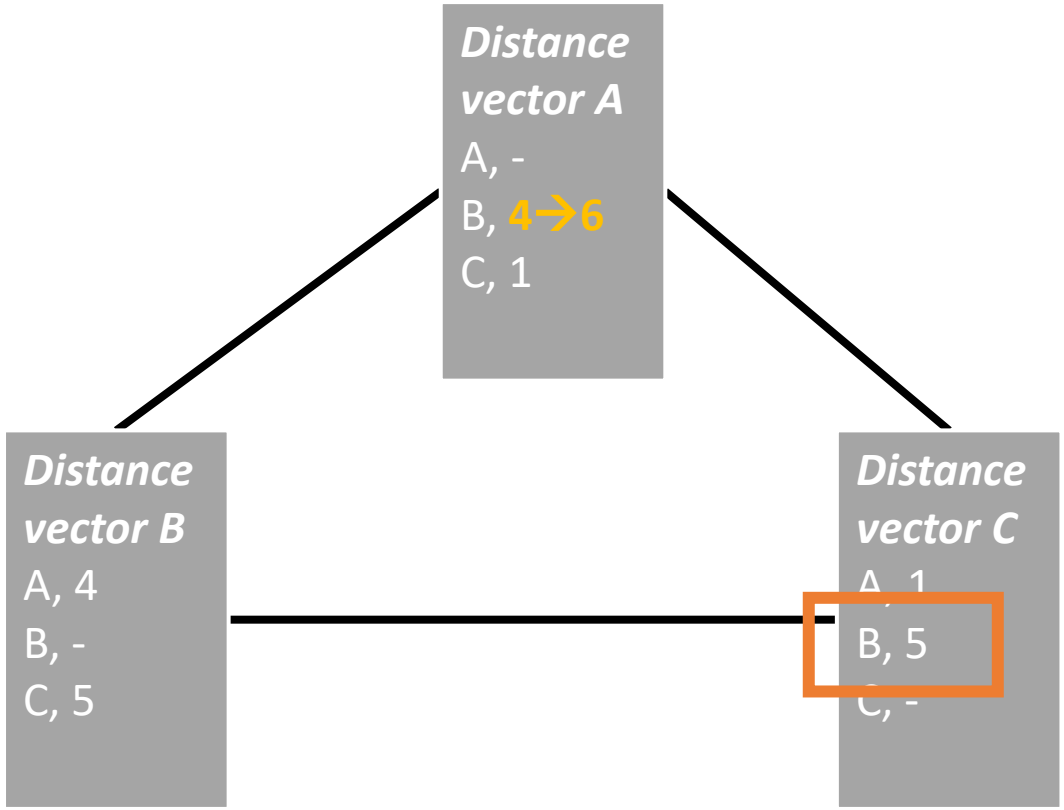

Distance Vector Routing
Bad News Propagation

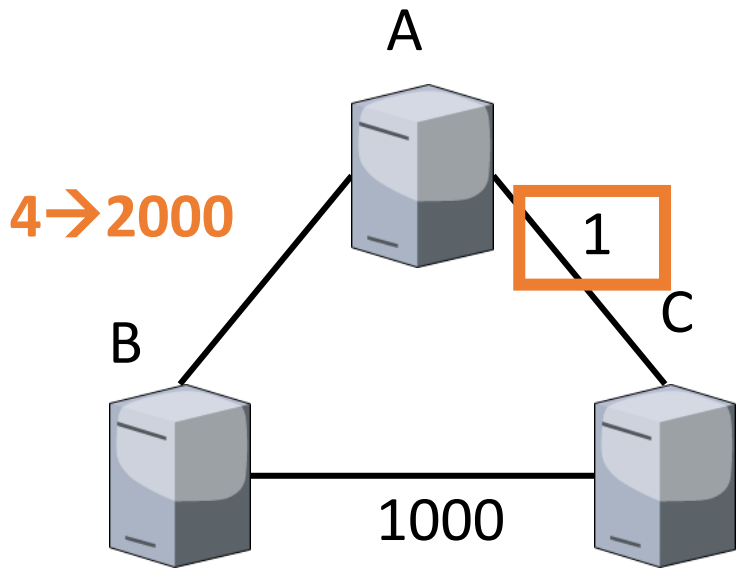

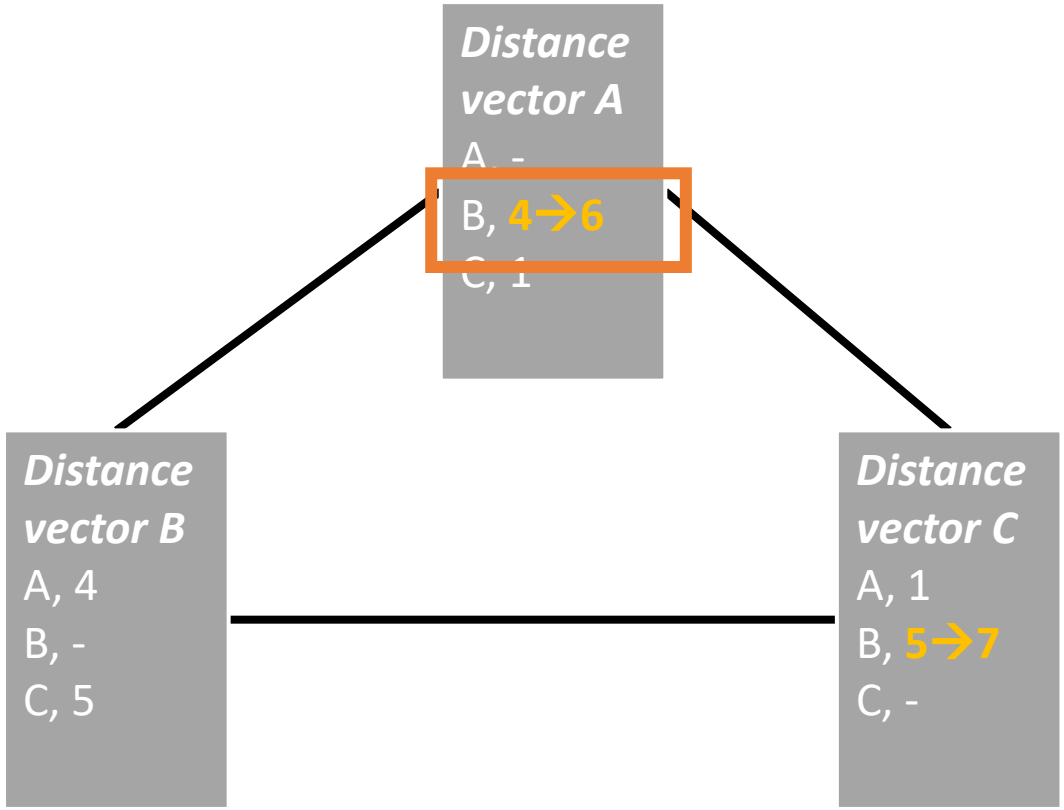

Distance Vector Routing
Bad News Propagation

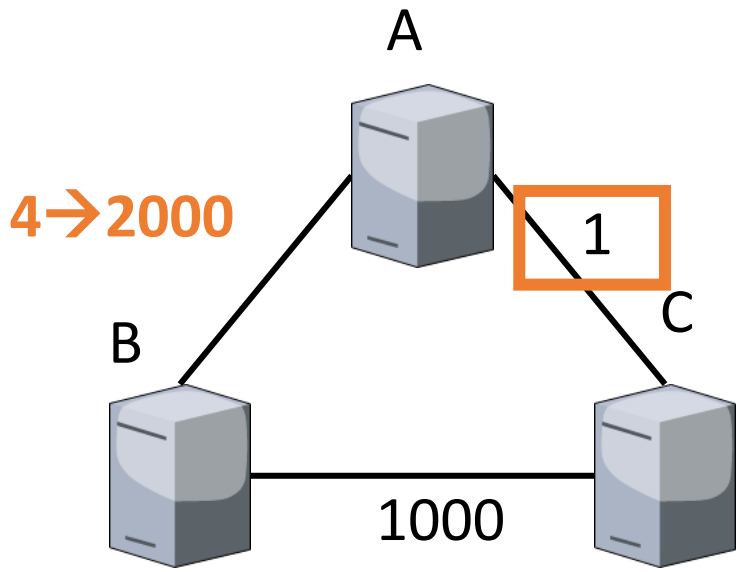

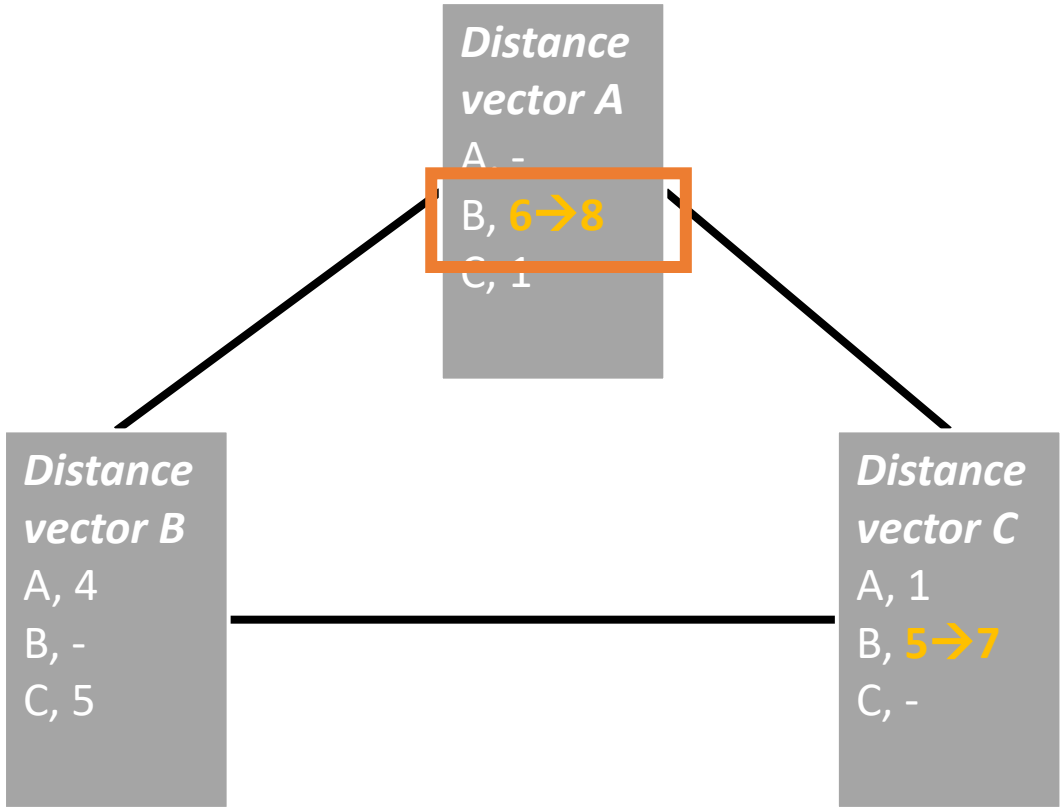

Count to infinity problem
When machine fails

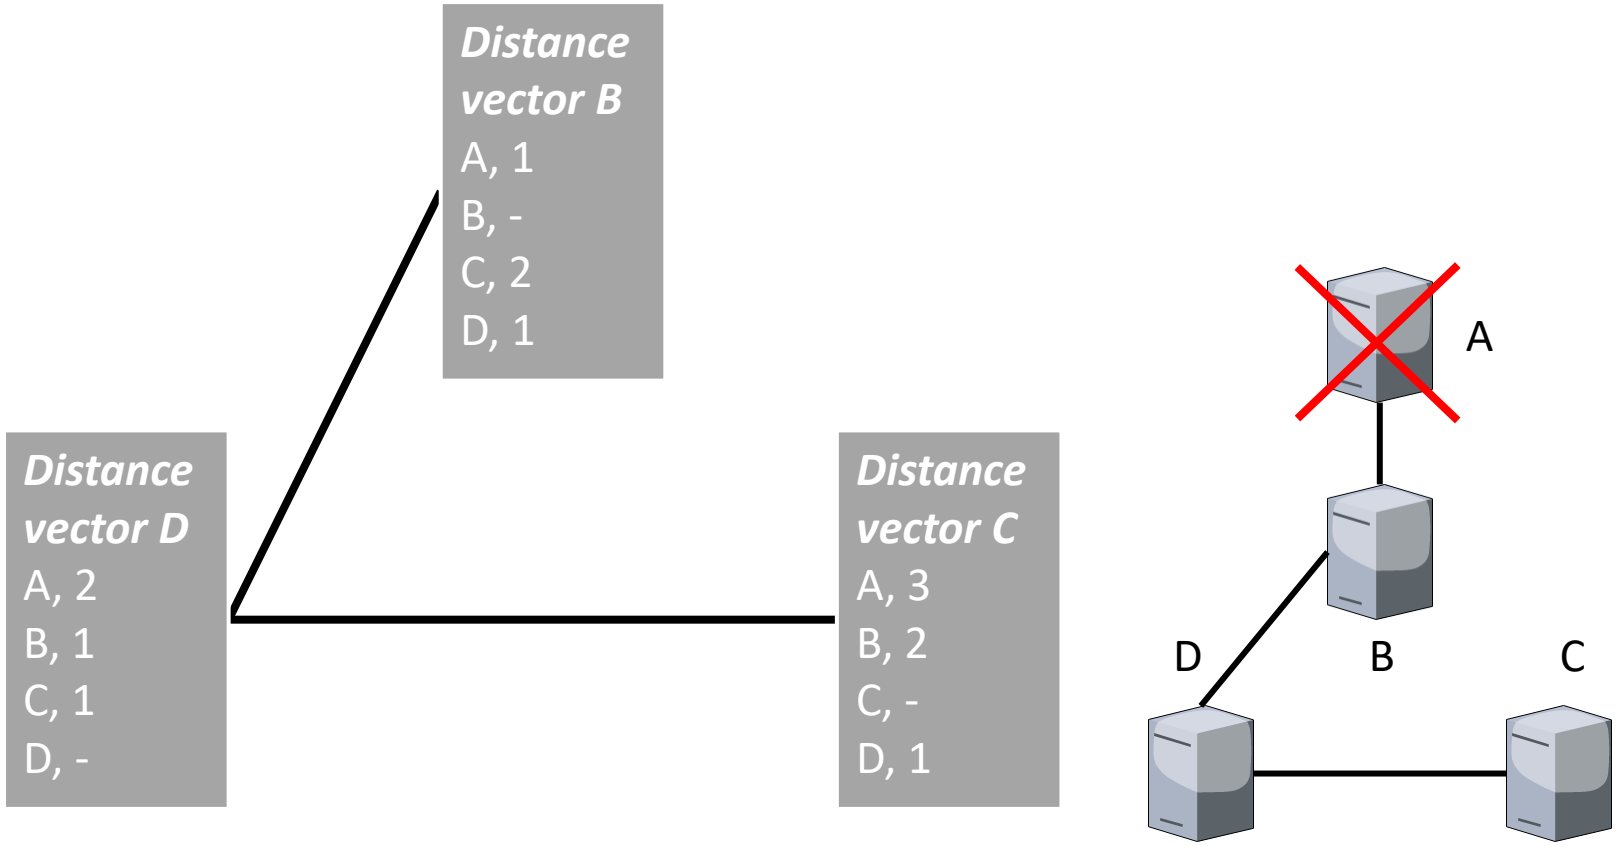

Count to infinity problem
When machine fails

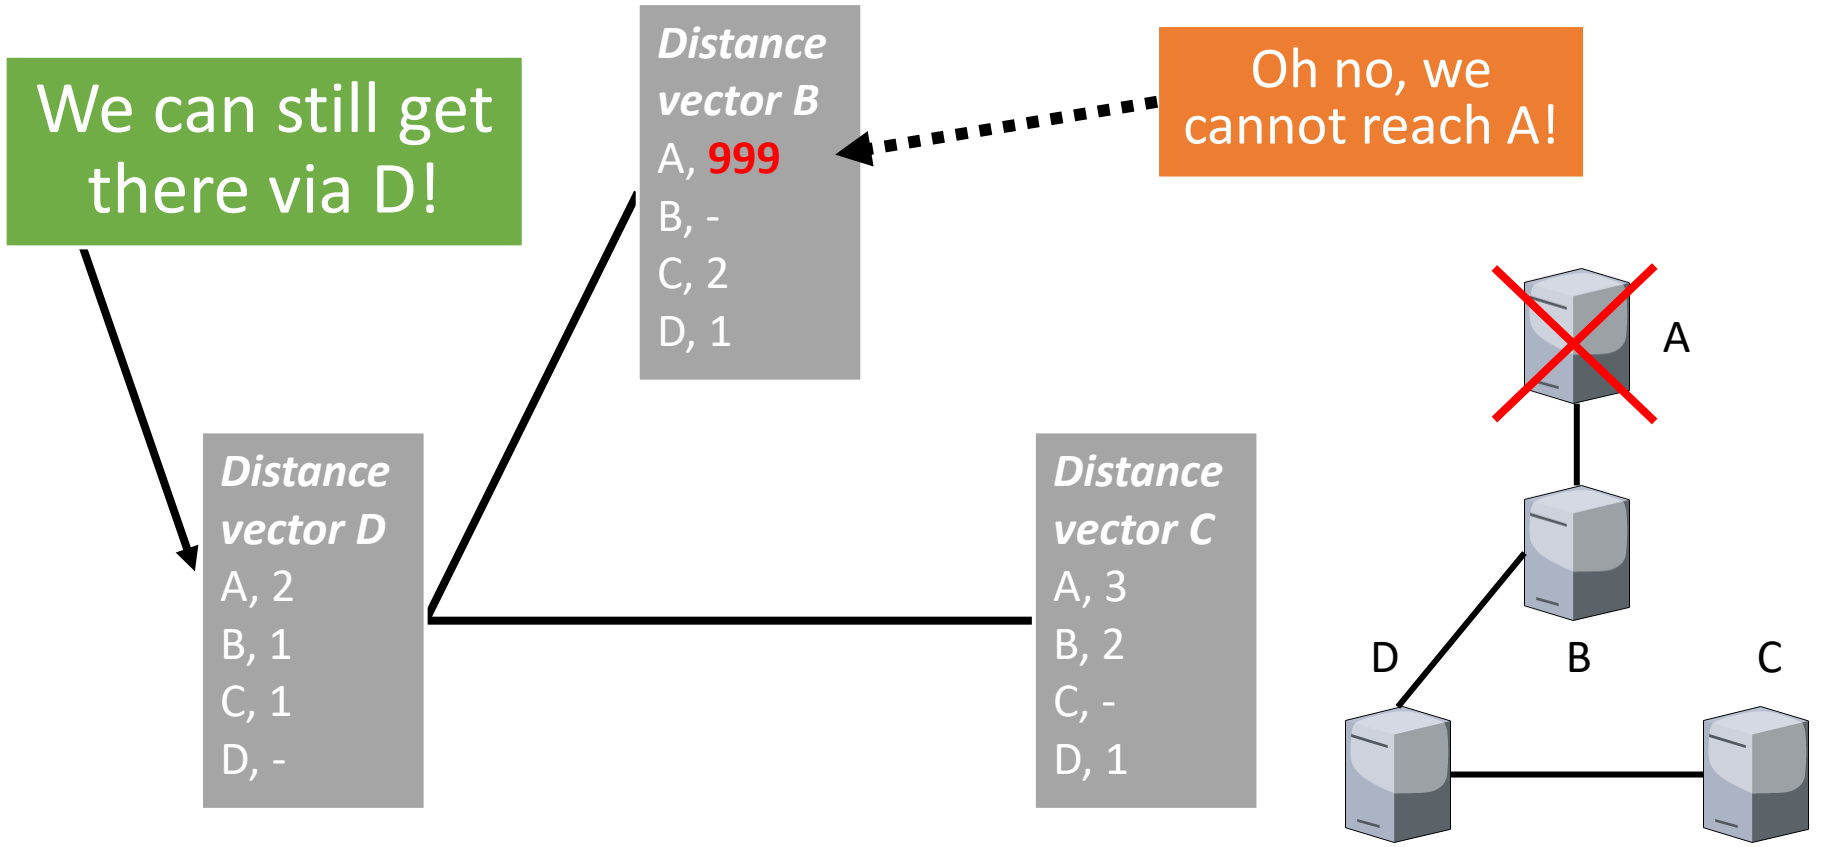

Count to infinity problem
When machine fails

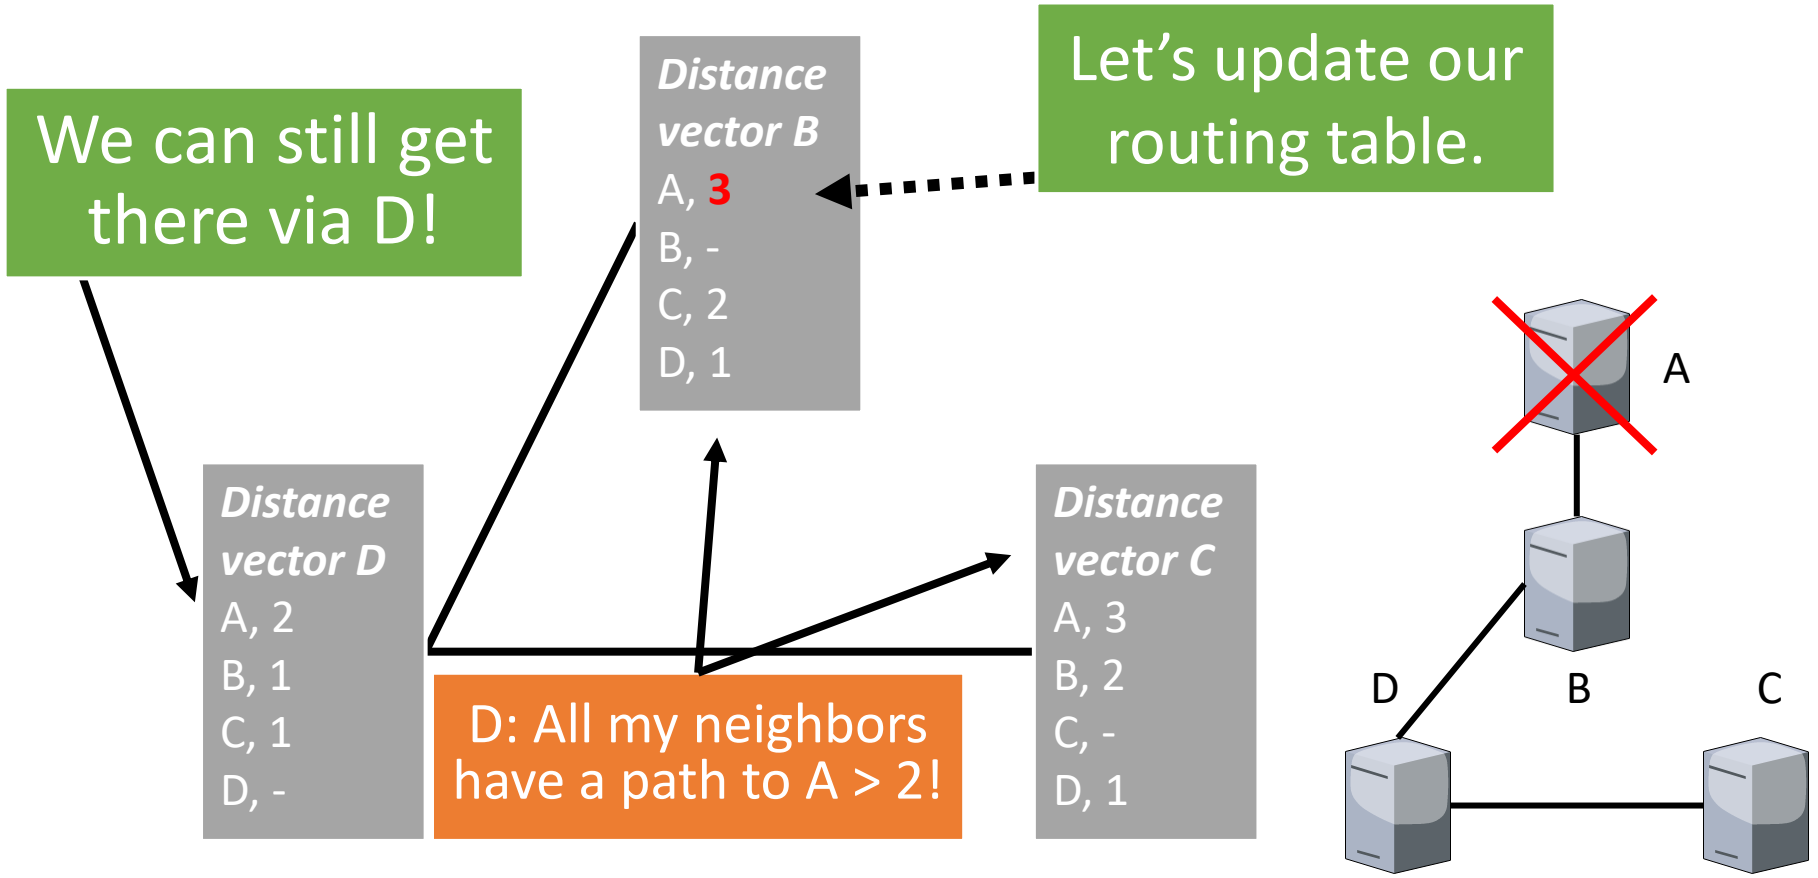

Count to infinity problem
When machine fails

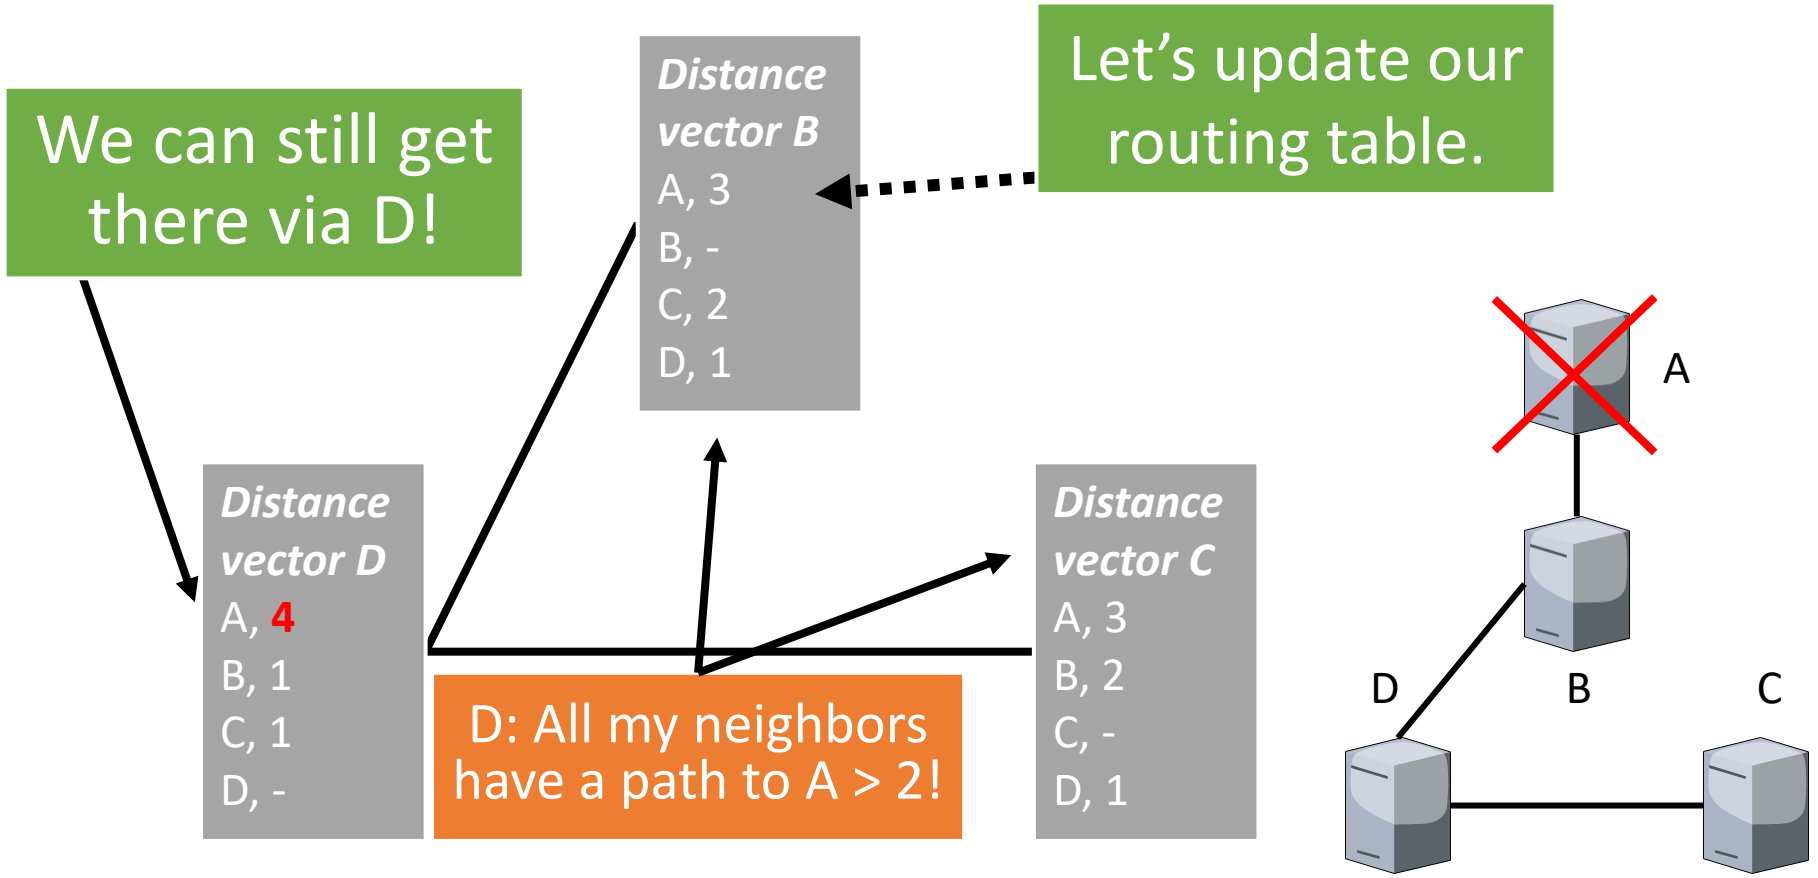

Count to infinity problem When machine fails

This will go on for a while…

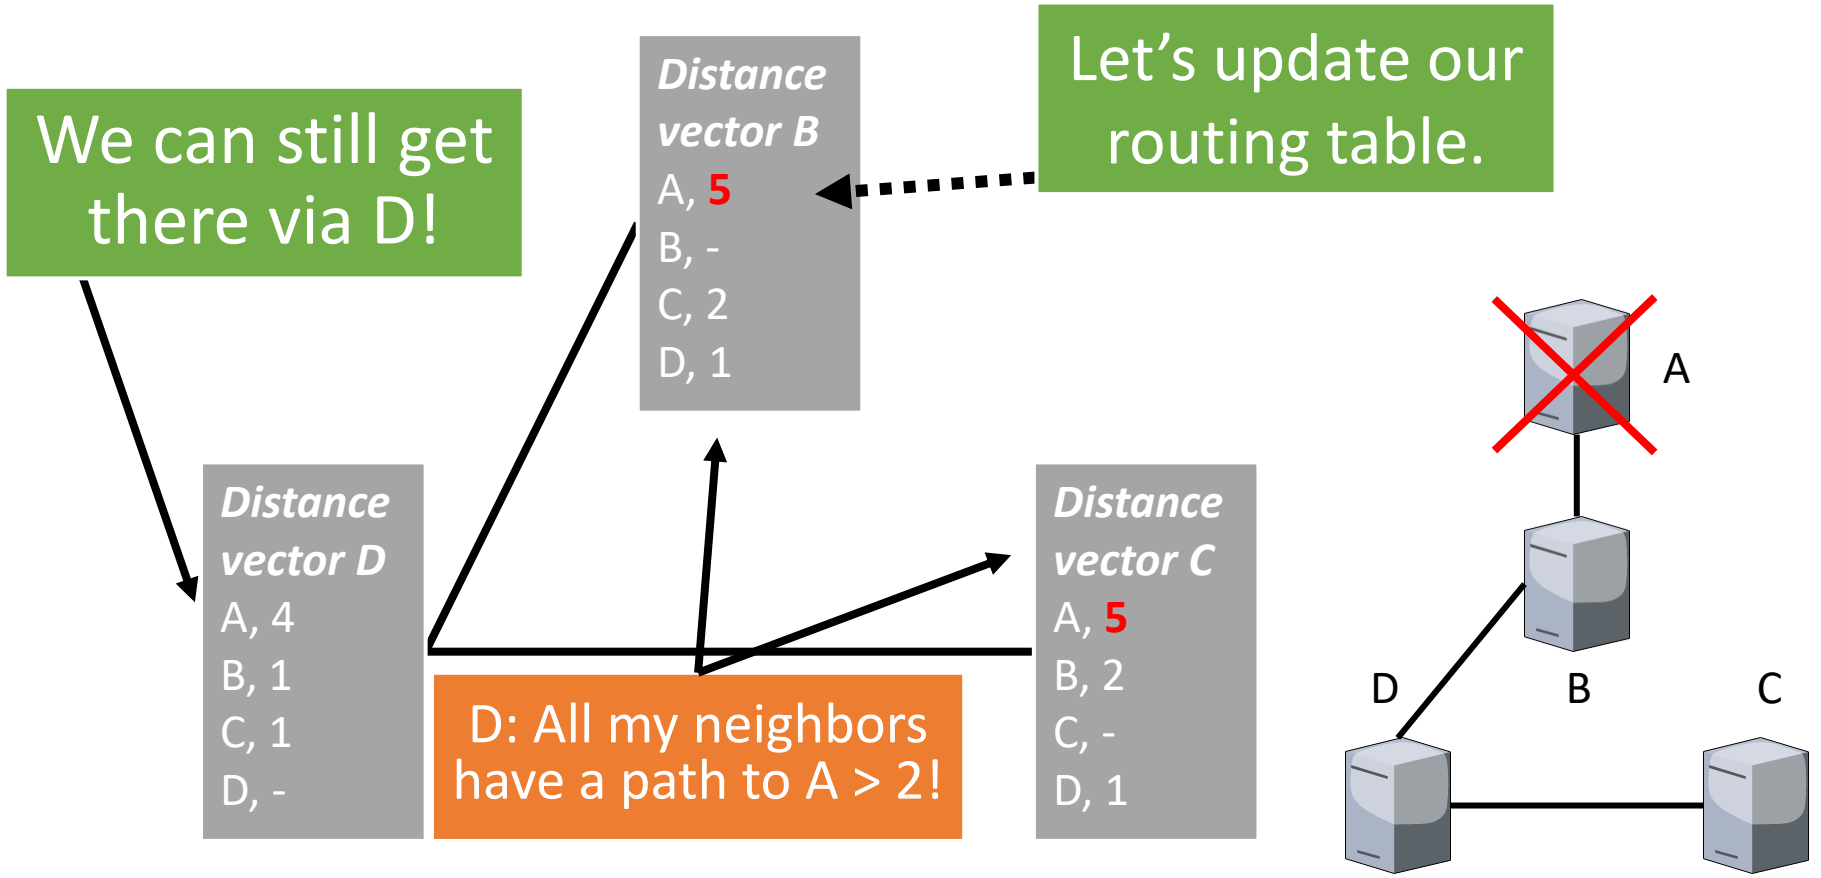

Link state routing

Replaced distance vector routing in ARPANET in 1979

Does not suffer from the *count to infinity problem*, but is more complex

Uses a *shortest path algorithm*

1. Routers only send packets with information about their direct neighbors
2. These packets are *flooded* over the network
3. Routers built an overview of the network using these packets and run a shortest path algorithm

Flooding link state packets

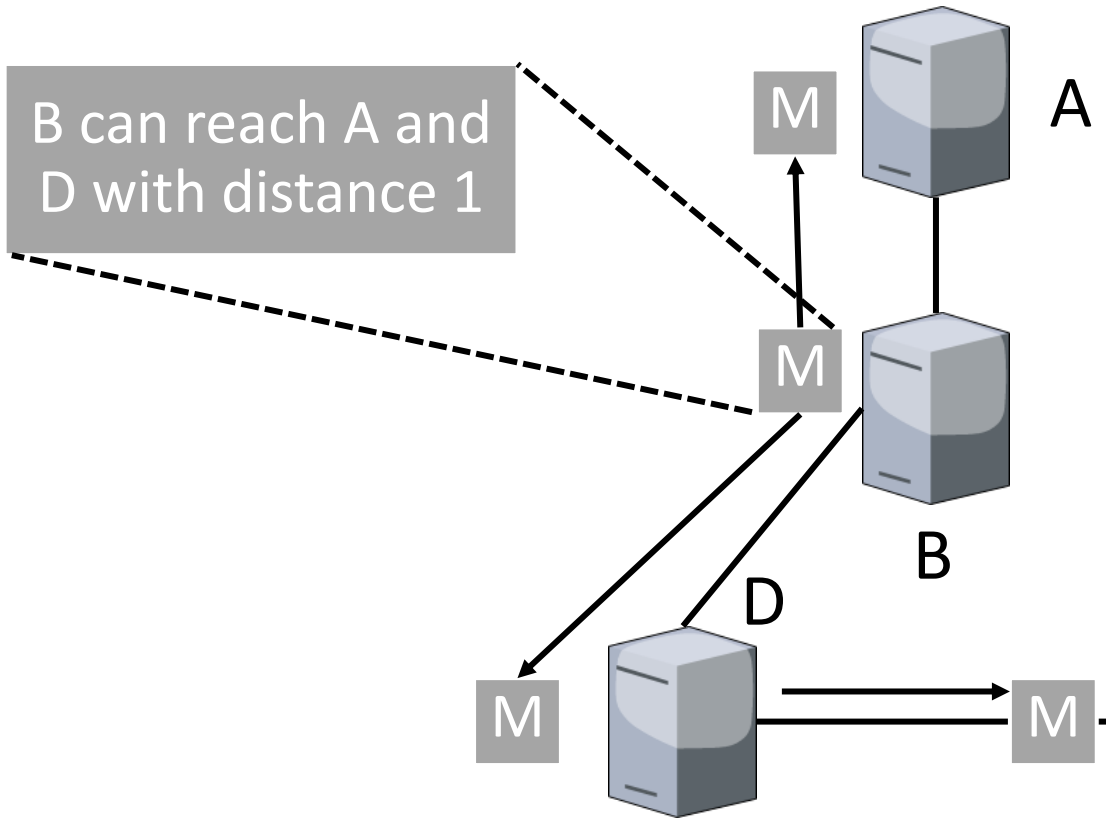

Q: What can go wrong?

Q: What can go wrong with sequence numbers?

Link state routing
Building a map of the network

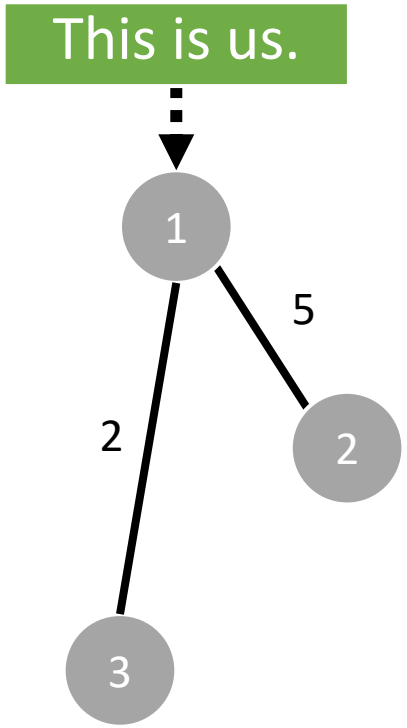

Link state routing
Building a map of the network

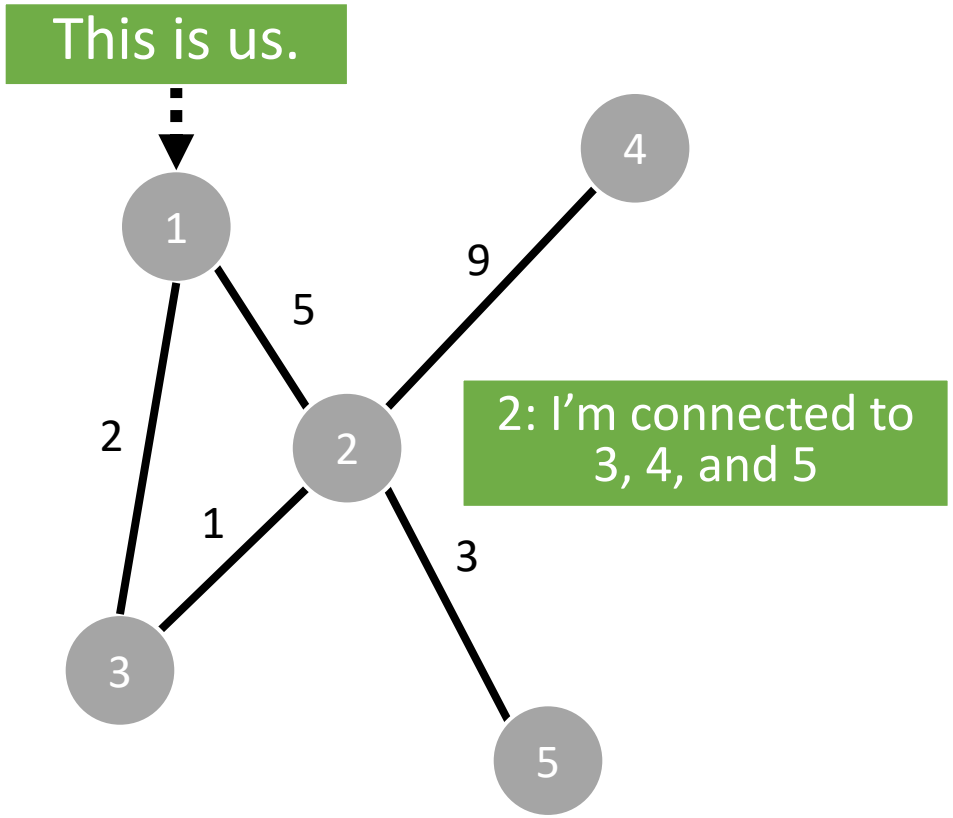

Link state routing
Building a map of the network

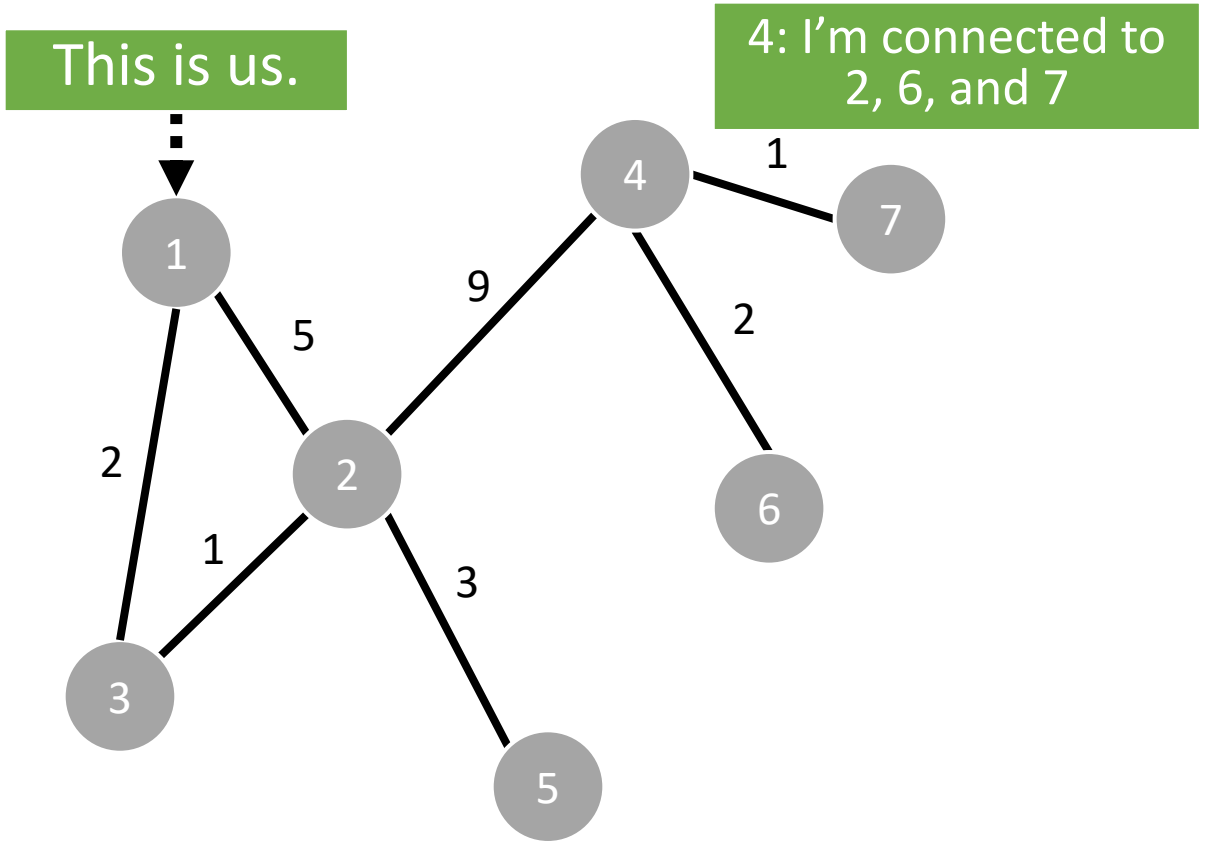

Link state routing
Building a map of the network

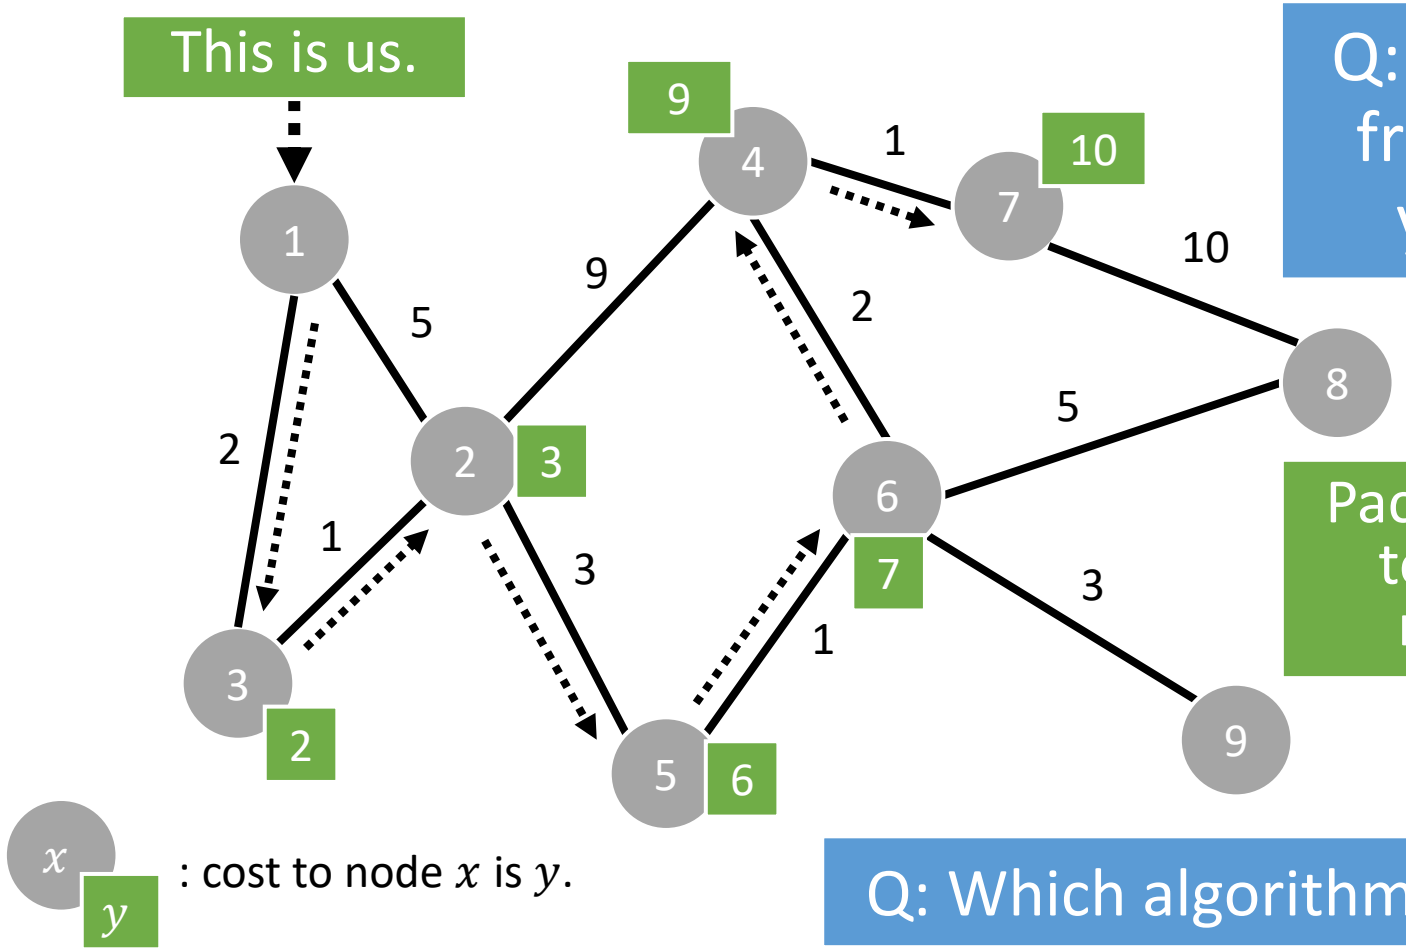

Q: Which route to take from 1 to 7? How did you compute this?

Packets addressed to 7 should be routed via 3!

Q: Which algorithm is used here?

Hierarchical routing

Reducing routing table sizes for large networks.

A1’s RT size: 7 (all other nodes).

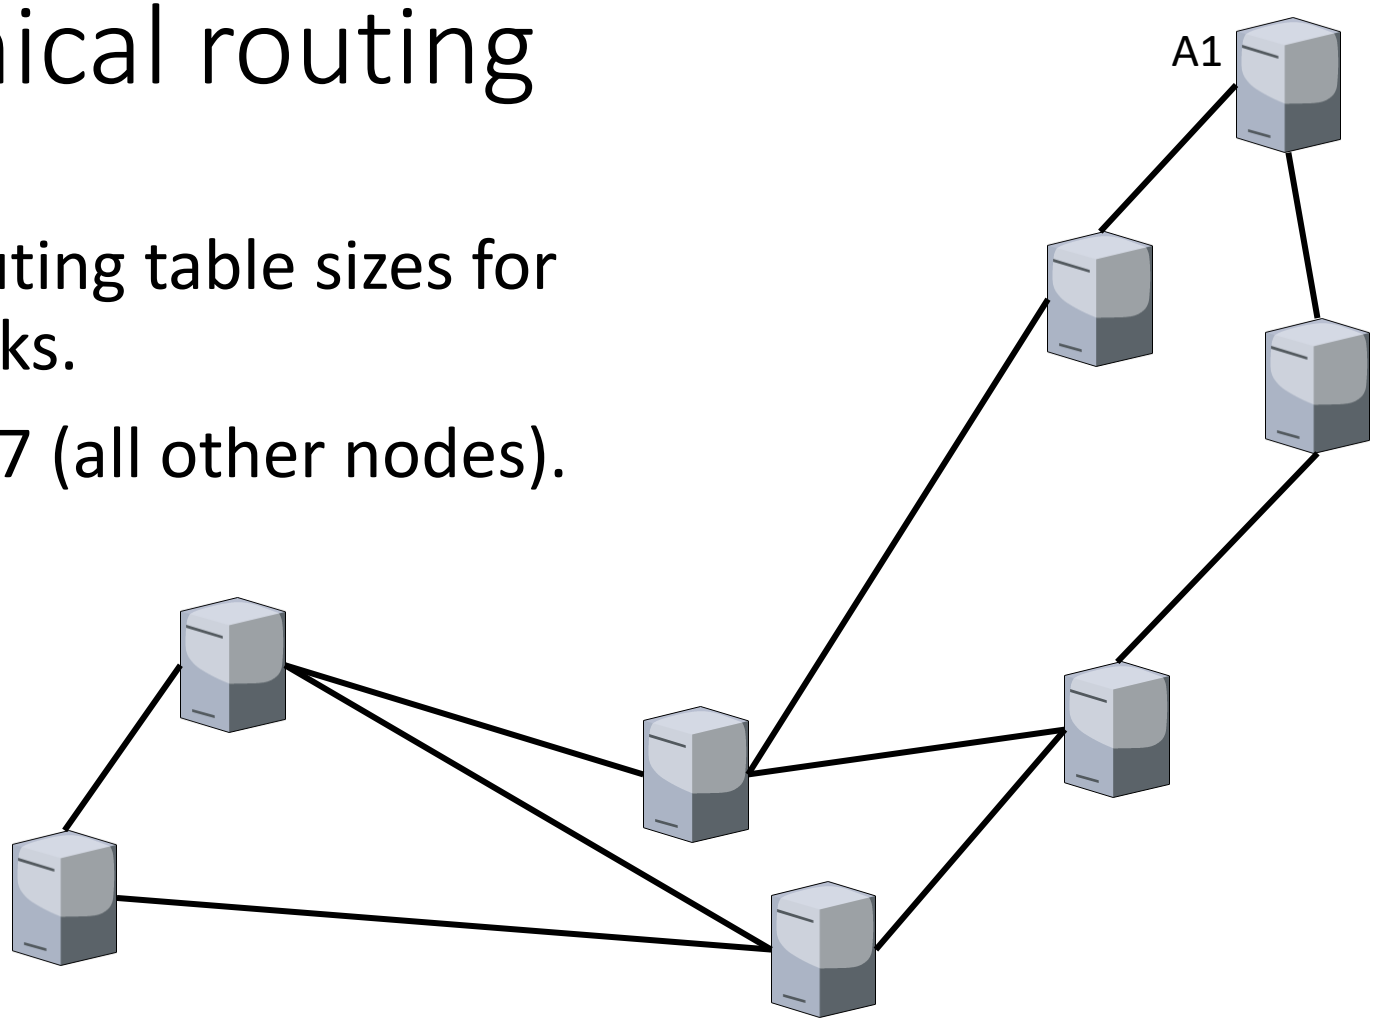

Hierarchical routing

Reducing routing table sizes for large networks.

A1's RT size: 7 (all other nodes).

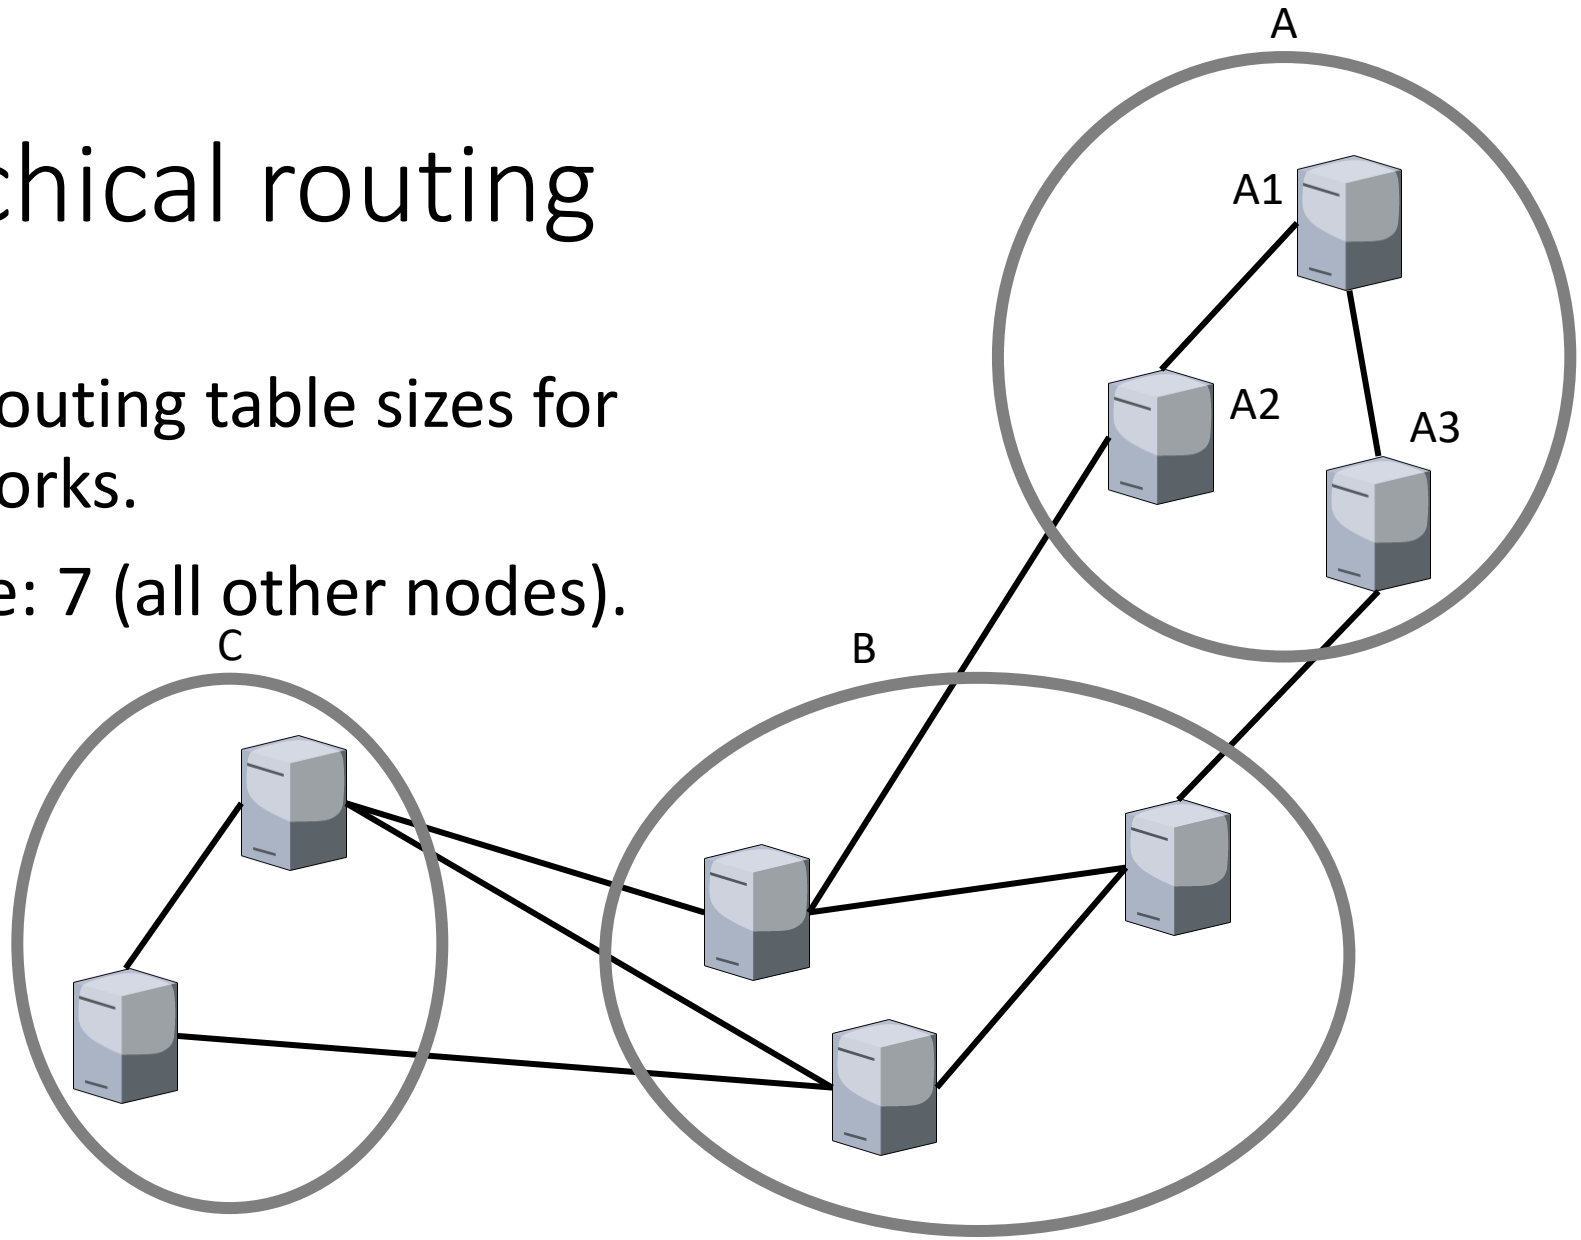

Hierarchical routing

Reducing routing table sizes for large networks.

A1's RT size: ***4*** (A2, A3, B, C).

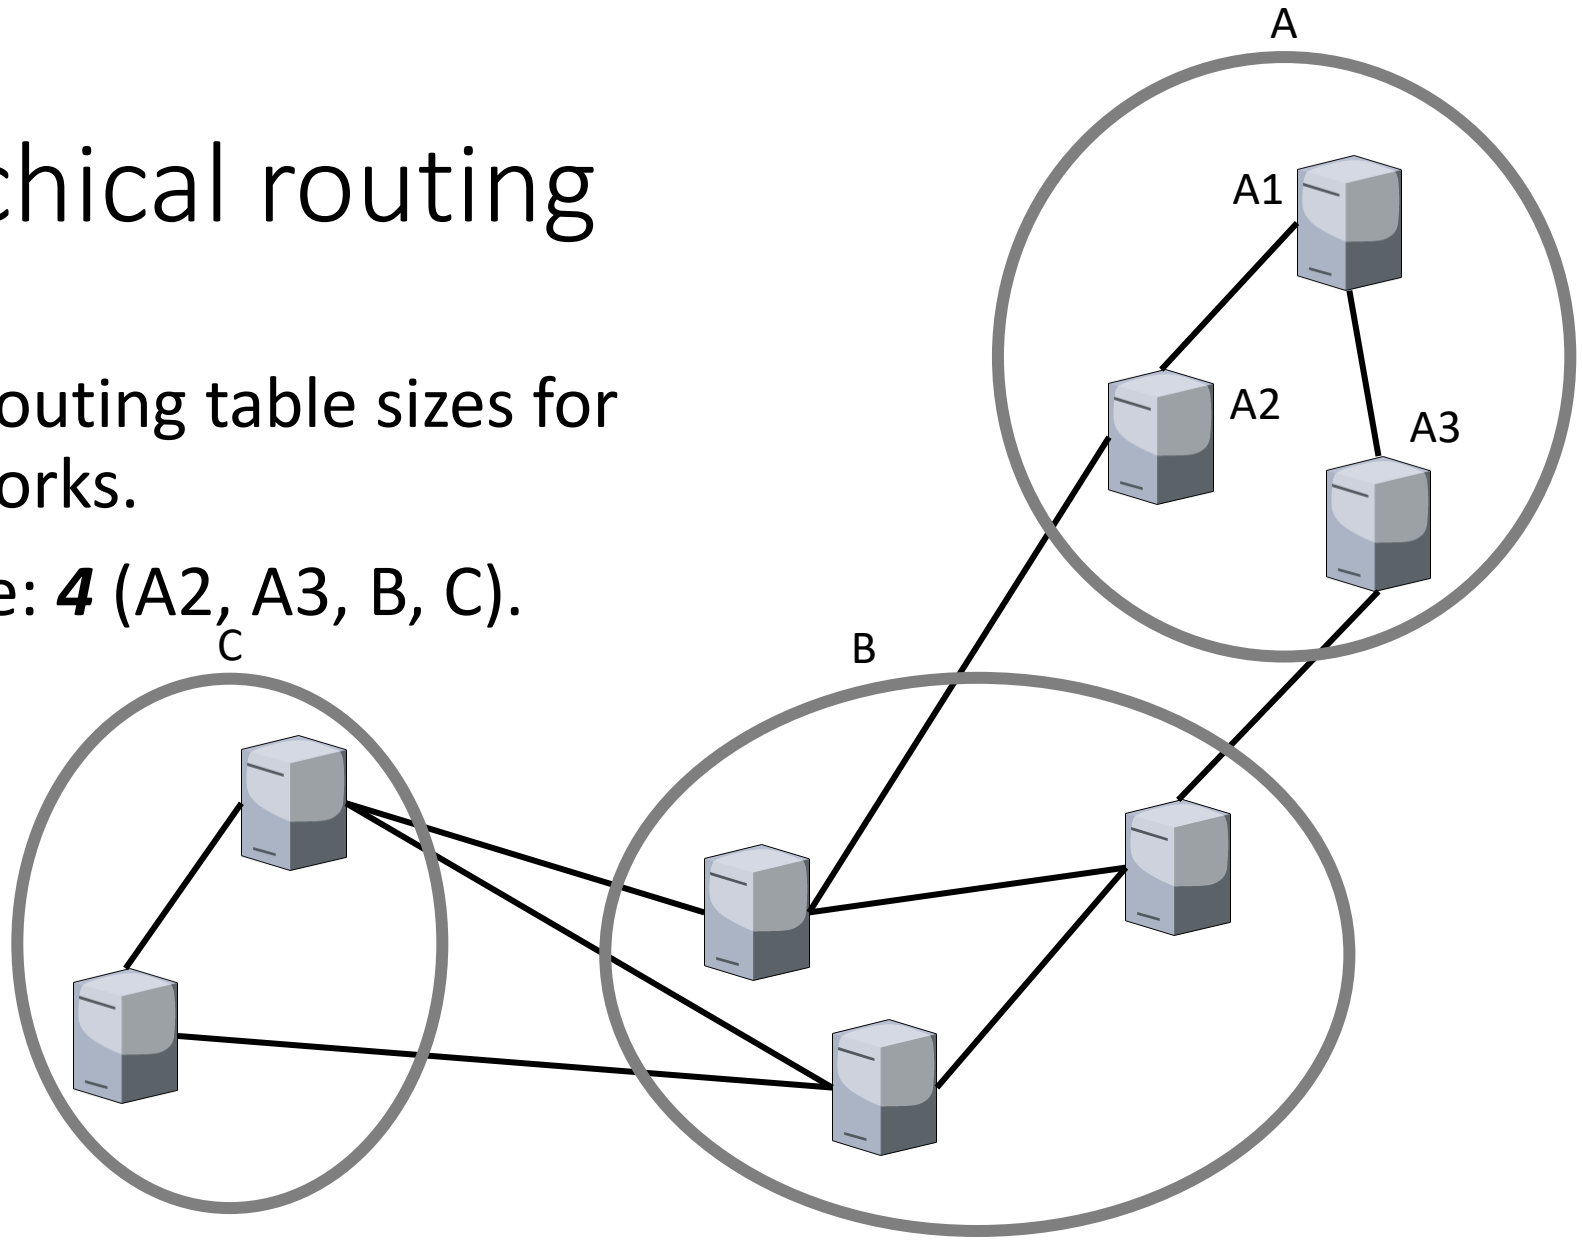

Today's Lecture

1. Routing Algorithms
2. **Internetworking**
    1. **Routing in and between Autonomous Systems: OSPF and BGP**
    2. **Tunneling**
    3. **Packet fragmentation**

Q: Difference between a single network and a collection of networks?

Internetworking

Getting packets to their destination across multiple networks

Internetworking

Challenges for sending packets end-to-end over multiple networks:

1. Technological
   - **Different protocols**
   - **Different maximum packet sizes**
   - Different QoS guarantees
2. Political
   - **Different costs**
   - Privacy concerns
   - Competition/disputes

Structure of the Internet: A Network of Networks

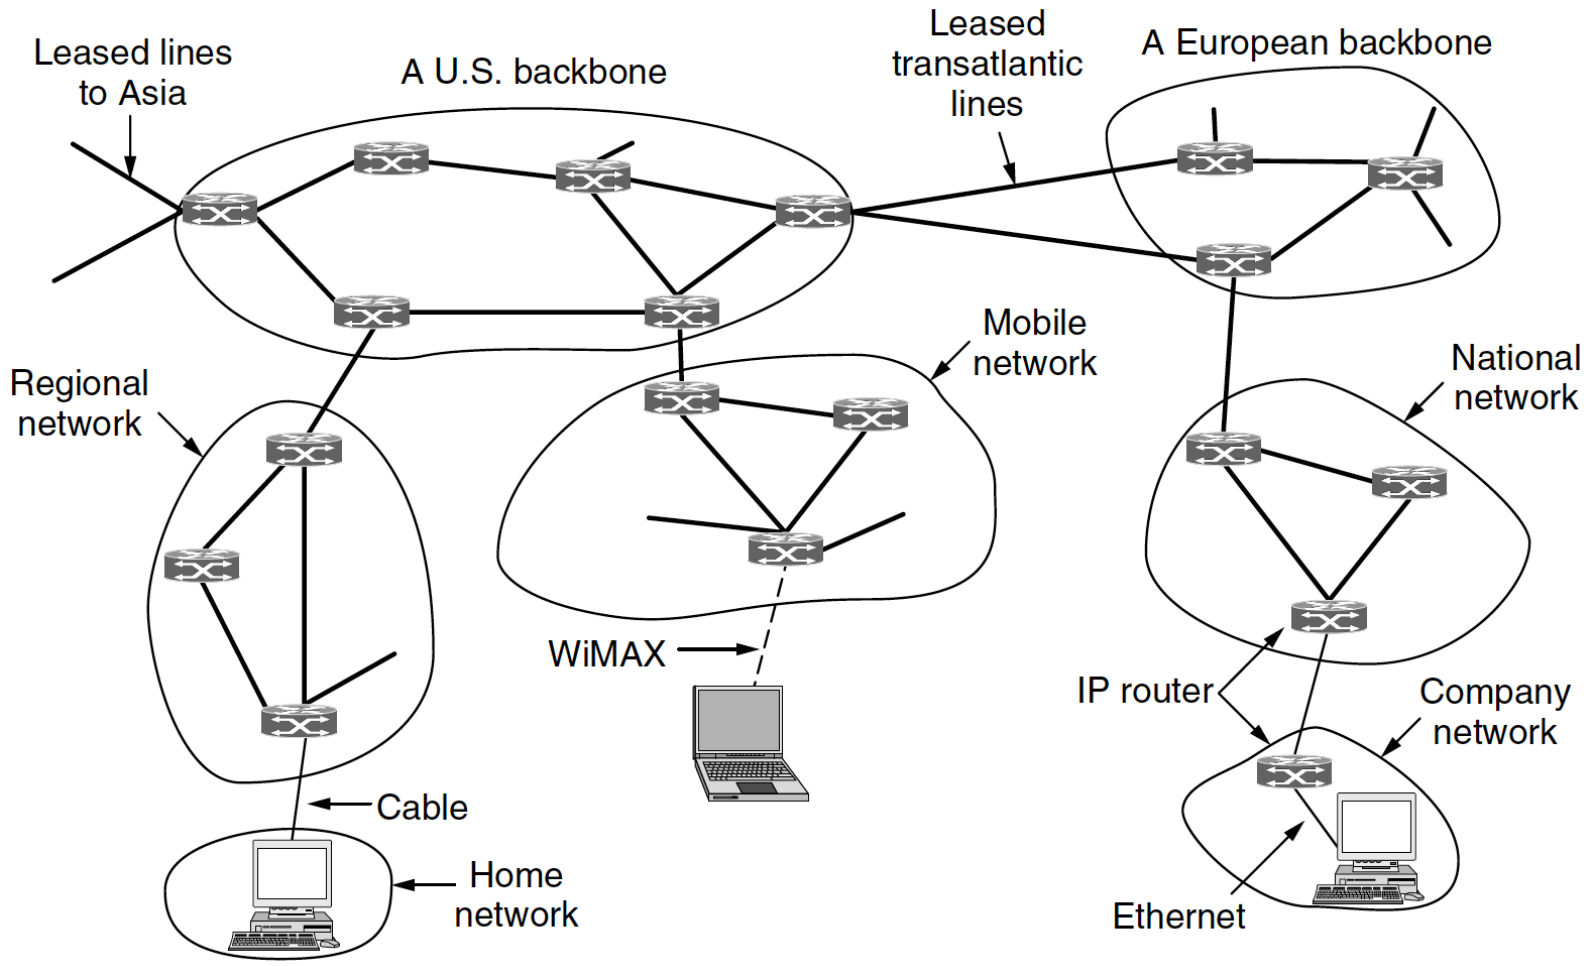

Tanenbaum, Andrew S., and David J. Wetherall. "Computer networks fifth edition." In *Pearson Education, Inc.* Prentice Hall, 2011.

We've almost come full circle: A Hierarchical Topology

A multitree?

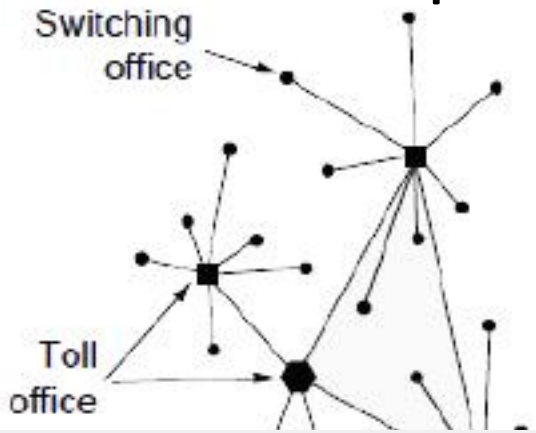

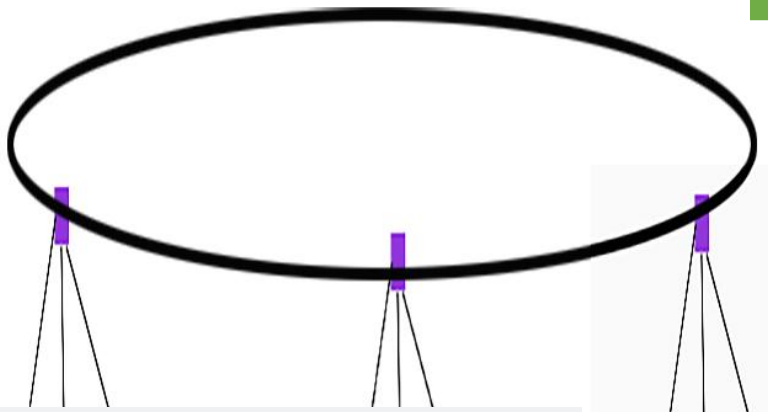

ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE STORE

FIBER FIASCO —

Cut submarine cables cause web outages across Africa; 6 countries still affected

Parts of Africa were already seeing web disruptions from damaged Red Sea cables.

SCHARON HARDING - 3/15/2024, 11:40 PM

Source: https://arstechnica.com/information-technology/2024/03/internet-outages-hit-13-countries-in-africa-due-to-undersea-cable-damage/

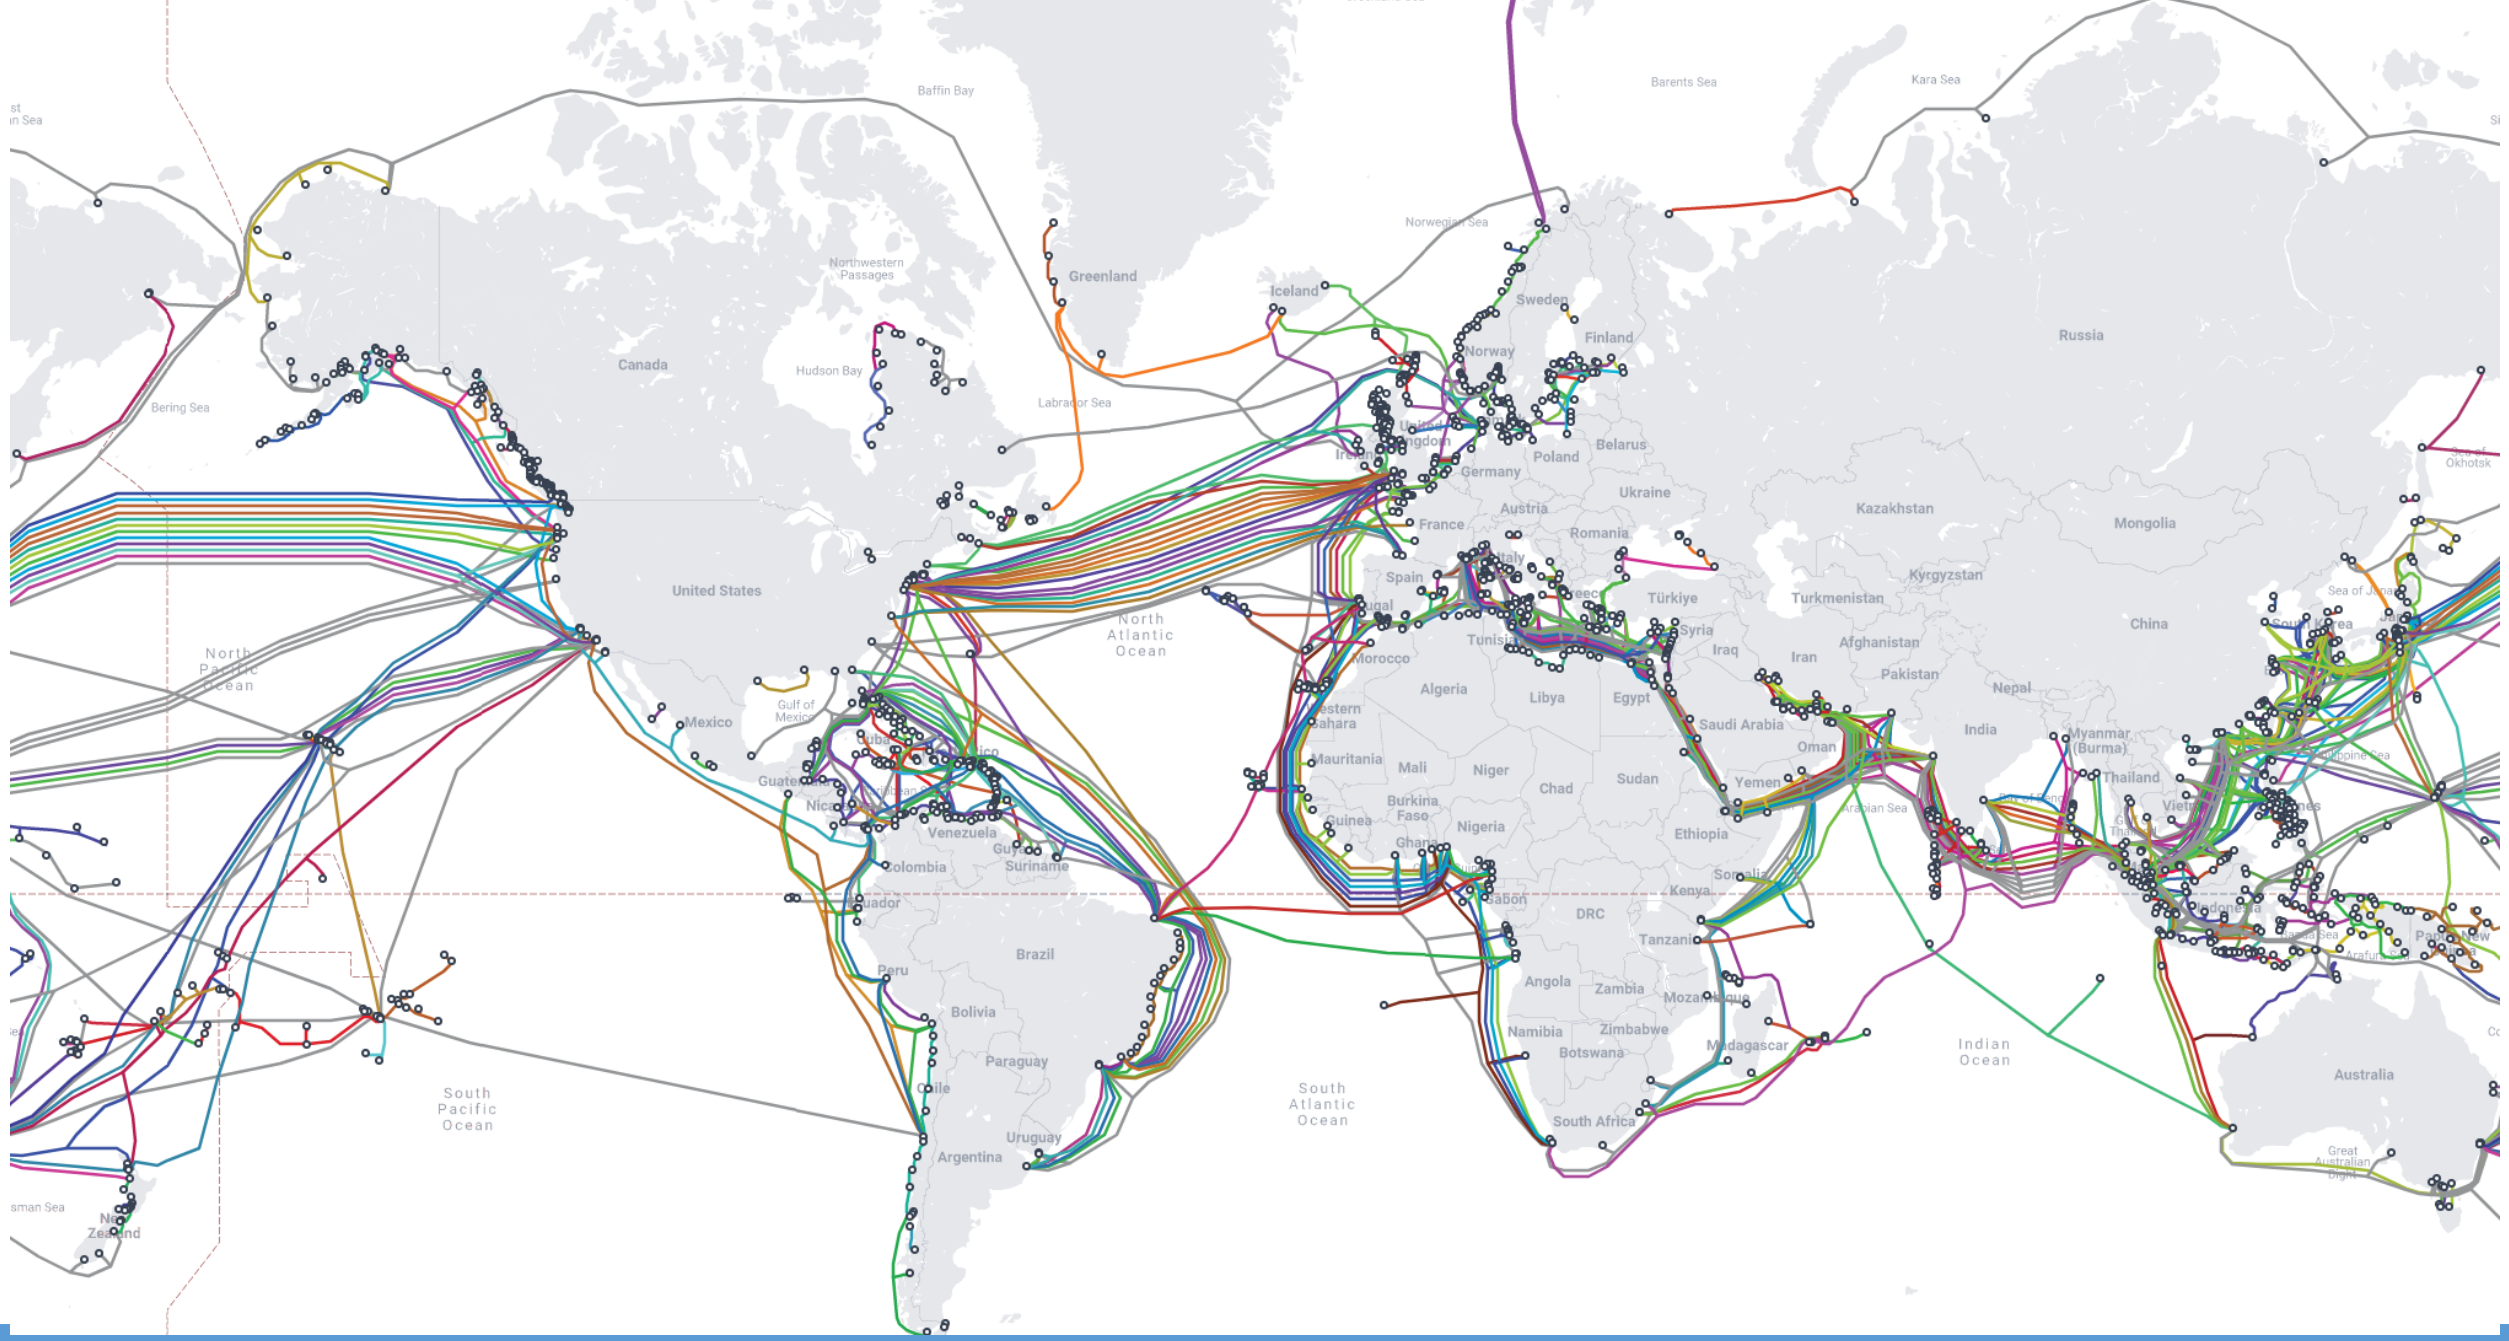

Source: https://www.submarinecablemap.com/

Structure of the Internet: A Network of Networks

Network of networks: yet another layer of abstraction

A large corporation using multiple points of attachment: *multihoming*

ISPs and (large) corporations are *Autonomous Systems* (AS)

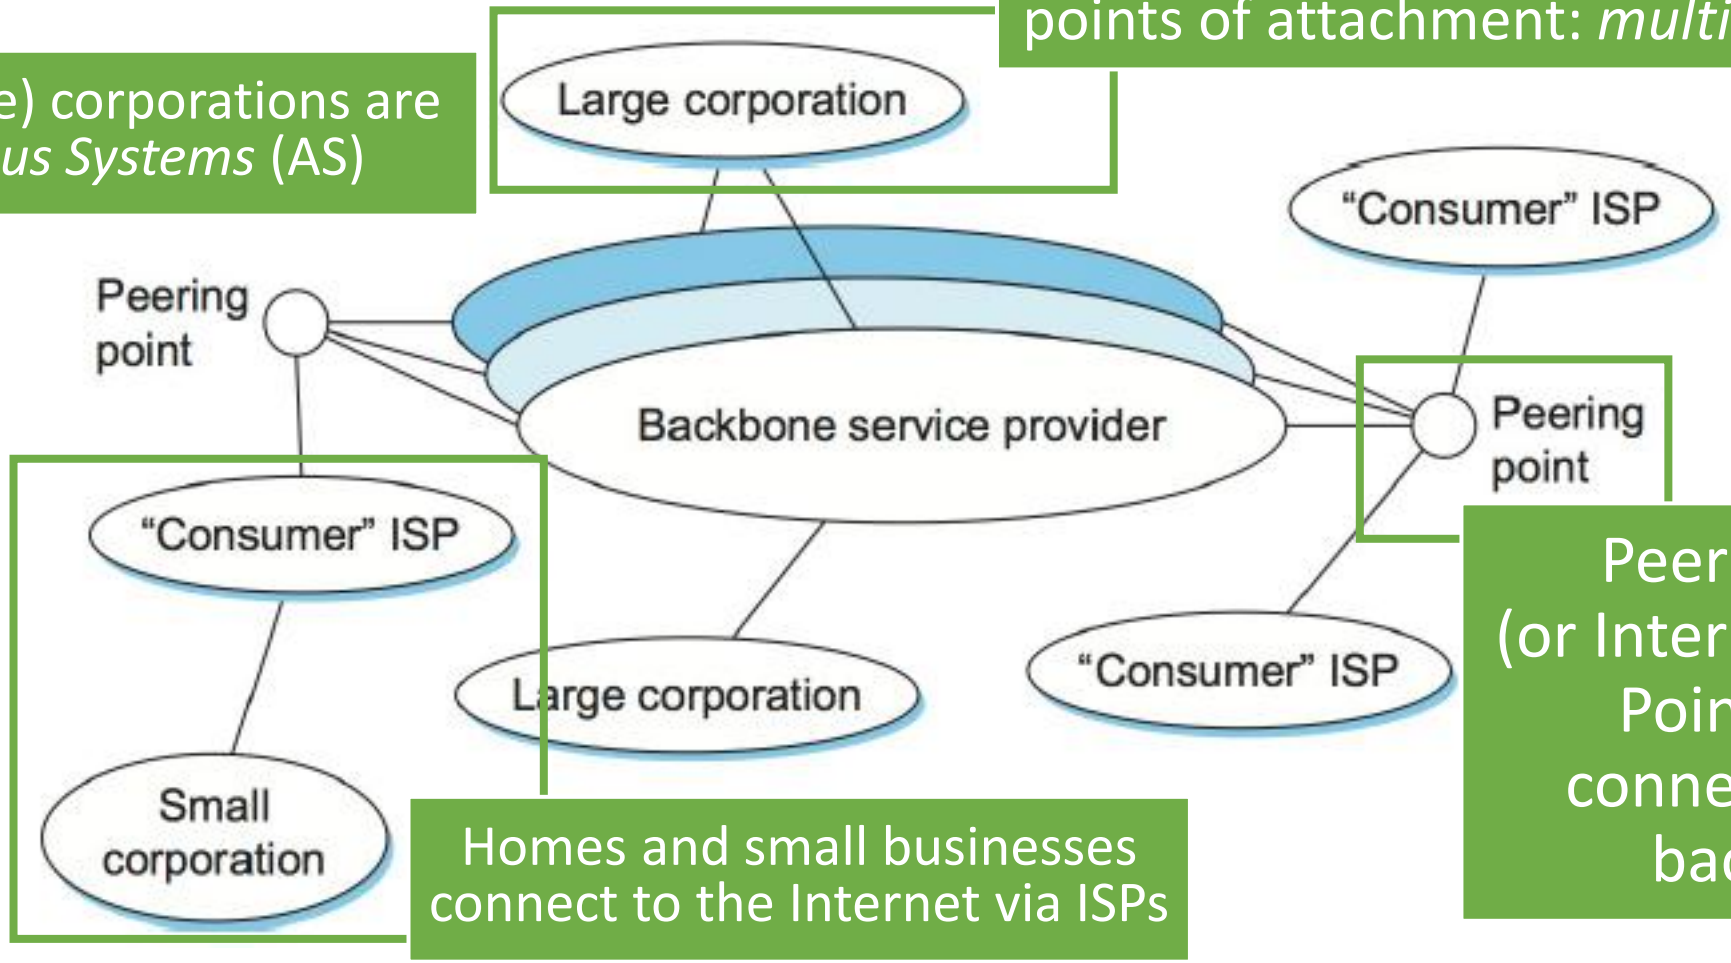

Homes and small businesses connect to the Internet via ISPs

Peering points (or Internet eXchange Points [IXPs]) connect ISPs and backbones

Internetworking with Autonomous Systems

Involves two key ideas:

1. Routing ***inside*** an Autonomous System
   - *Intradomain* routing. Uses an **Interior Gateway Protocol**

     We will look at the **Open Shortest Path First (OSPF)** protocol

2. Routing ***between*** Autonomous Systems
   - *Interdomain* routing. Uses an **Exterior Gateway Protocol**

     We will look at the **Border Gateway Protocol (BGP)**

Example of an Interior Gateway Protocol

Open Shortest Path First (OSPF)

Routing ***within*** an Autonomous System.

Uses a form of
link state routing

Builds a graph
representation of the network

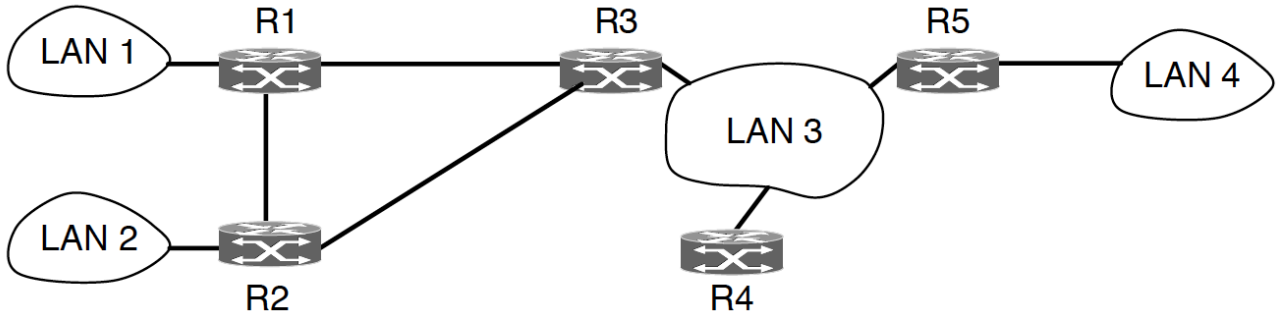

Broadcast LANs modeled as node with zero-cost links

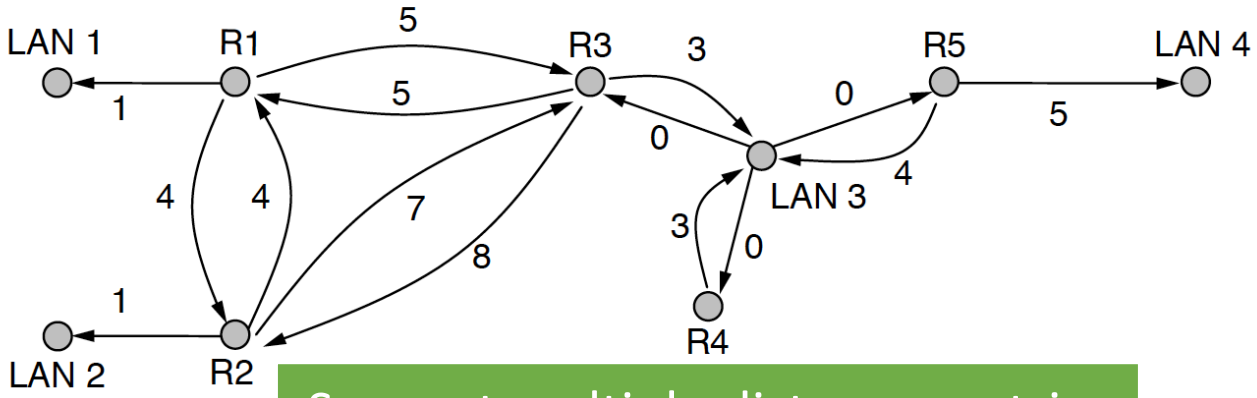

Support multiple distance metrics
(e.g., delay, physical distance)

Open Shortest Path First (OSPF)

Finds all shortest paths. Uses **Equal Cost MultiPath (ECMP):** traffic distributed over shortest paths. Example: traffic from R3 to LAN 2

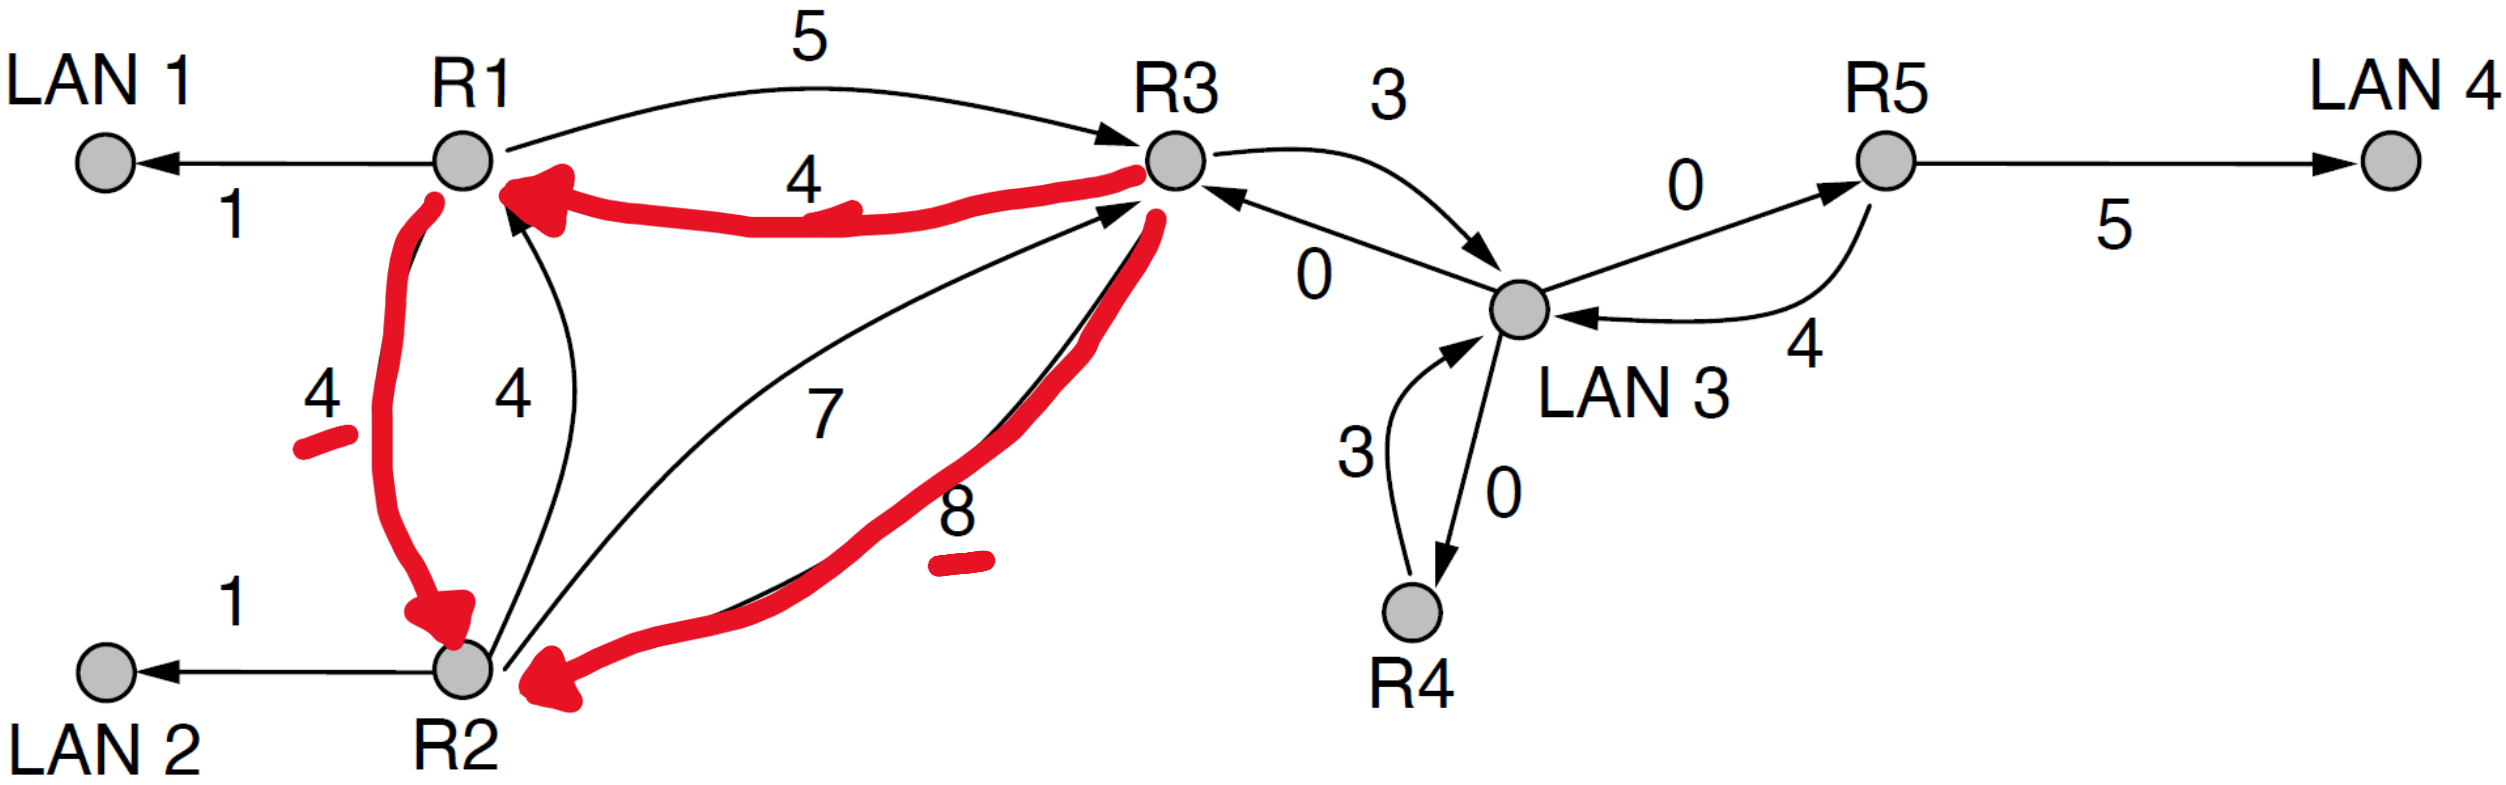

OSPF Message Types

Message type	Description
Hello	Used to discover who the neighbors are
Link state update	Provides the sender's costs to its neighbors
Link state ack	Acknowledges link state update
Database description	Announces which updates the sender has
Link state request	Requests information from the partner

OSPF messages are exchanged between routers; are wrapped in IP packets

Example of an Interior Gateway Protocol

Open Shortest Path First (OSPF)

Boundary routers connect to another AS. Uses both an Interior and Exterior Gateway Protocol

Routing ***within*** an Autonomous System.

Area border router exchanges summarized routing info between areas

Routers only know topology of the area they are in

Uses a ***hierarchy*** called "areas" to manage large networks

Q: Why support hierarchical routing?

For destinations not in area: go through backbone

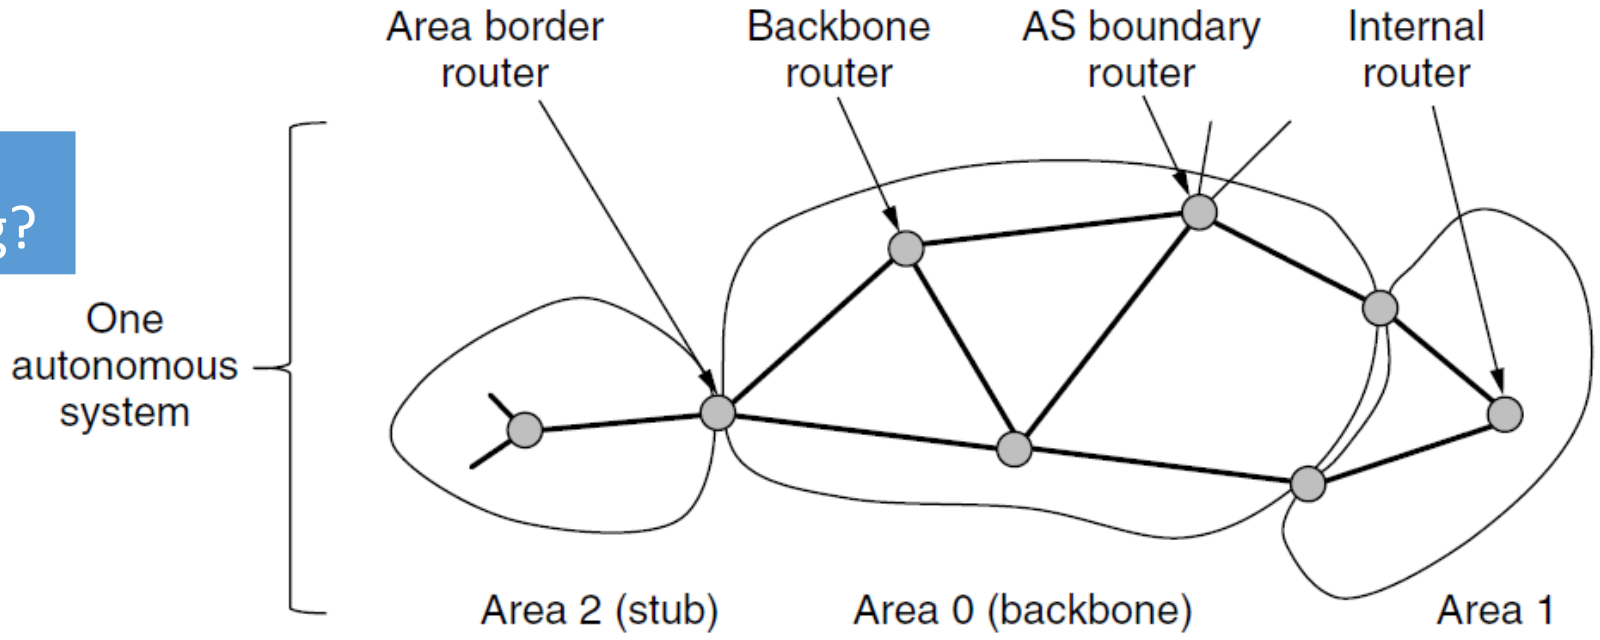

All areas connected to Area 0 (backbone)

Example of an Exterior Gateway Protocol

Border Gateway Protocol (BGP)

Routing ***between*** (large) independent networks.

Q: Can you think of examples of such policies?

Supports arbitrary policies put in place by ISPs, companies, organizations, etc.

AS1 provides *transit service*

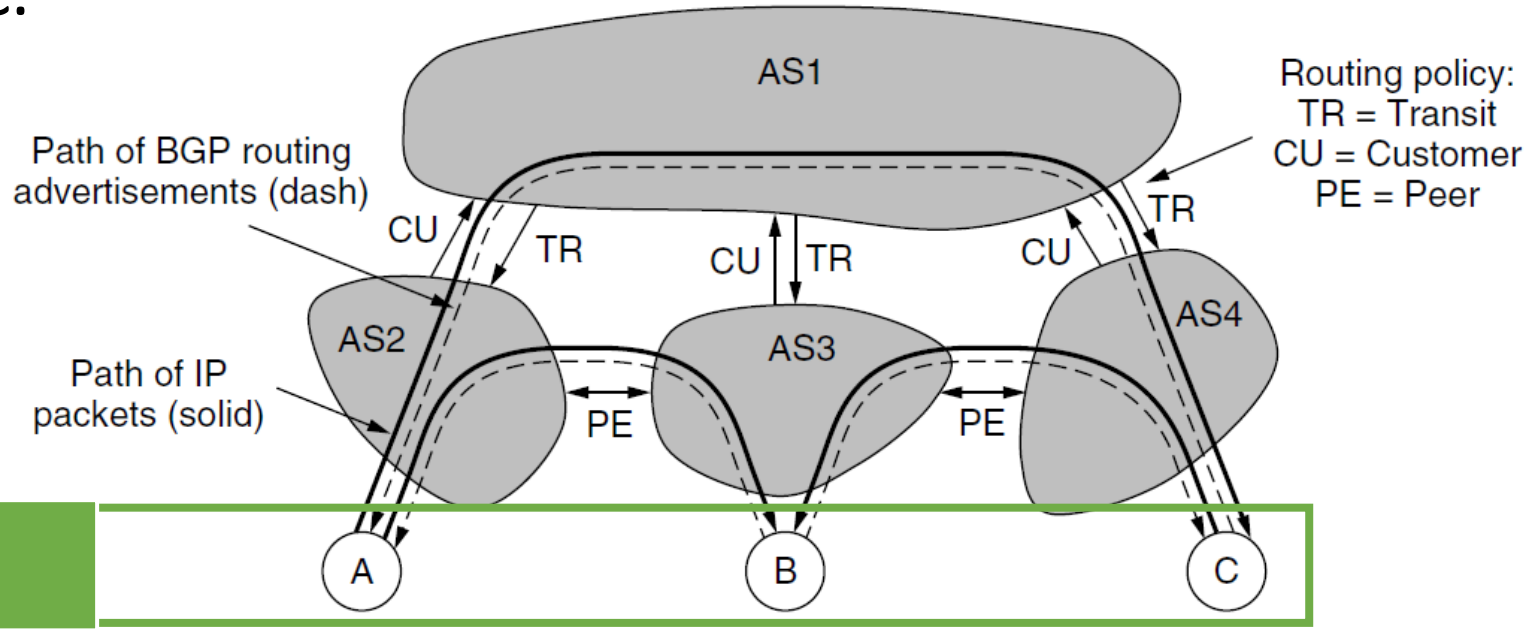

Autonomous Systems connect at IXPs

Stub networks do not need BGP:
only one way to send outgoing packets

Example of an Exterior Gateway Protocol

Border Gateway Protocol (BGP)

Routing ***between*** (large) independent networks.

Uses a ***path vector protocol*** (form of ***distance vector routing***)

Q: A, B, C? How are autonomous systems identified in practice in BGP?

AS identified by address prefix e.g., A is all IP addresses starting with 123

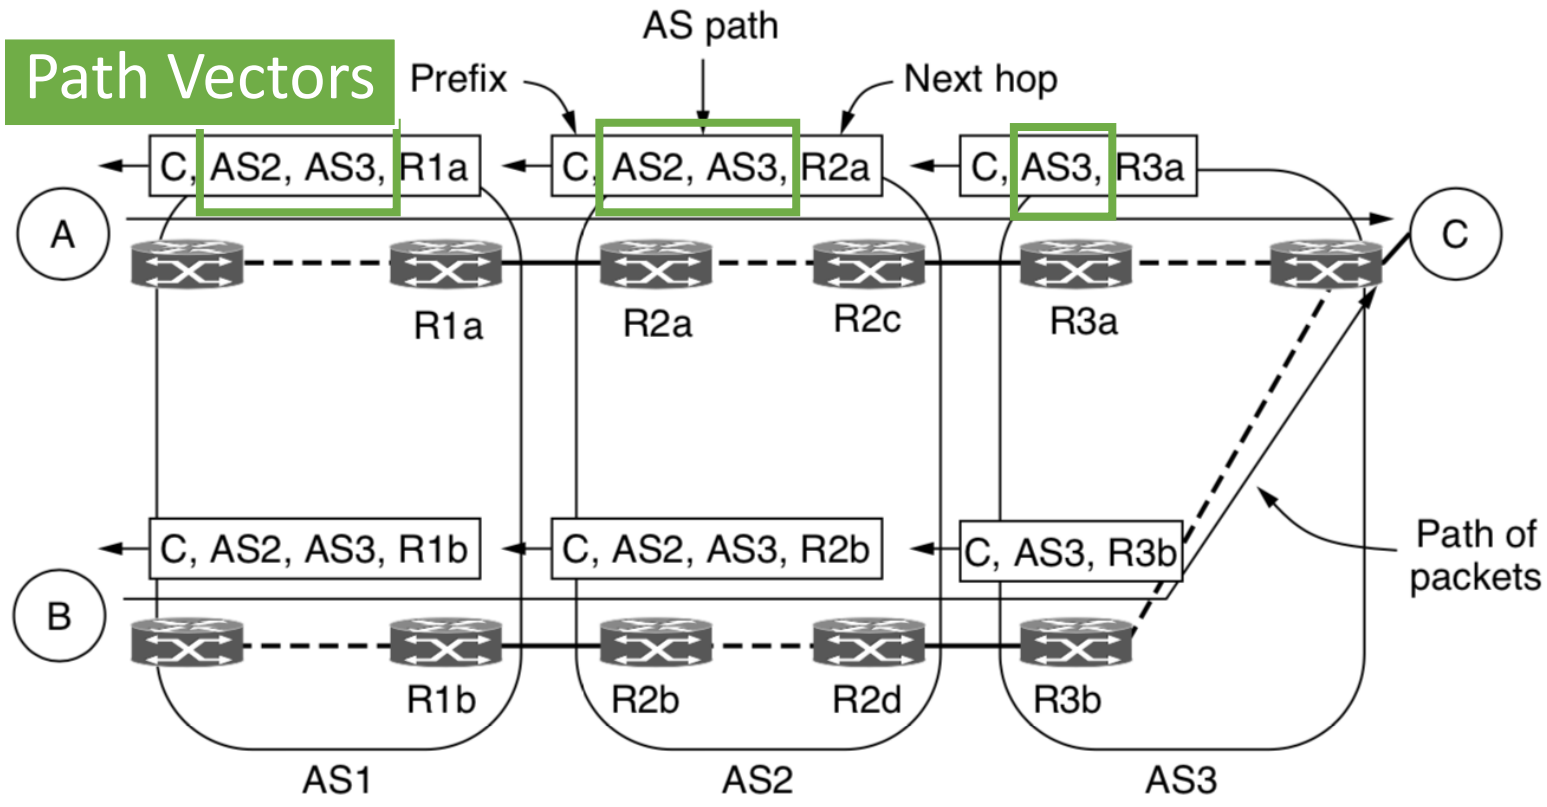

Tanenbaum, Andrew S., and David J. Wetherall. "Computer networks fifth edition." In *Pearson Education, Inc.* Prentice Hall, 2011.

Border Gateway Protocol (BGP)

Routing based on policies

Examples:

1. Peer above transit
2. Hot potato routing
3. Shortest AS path

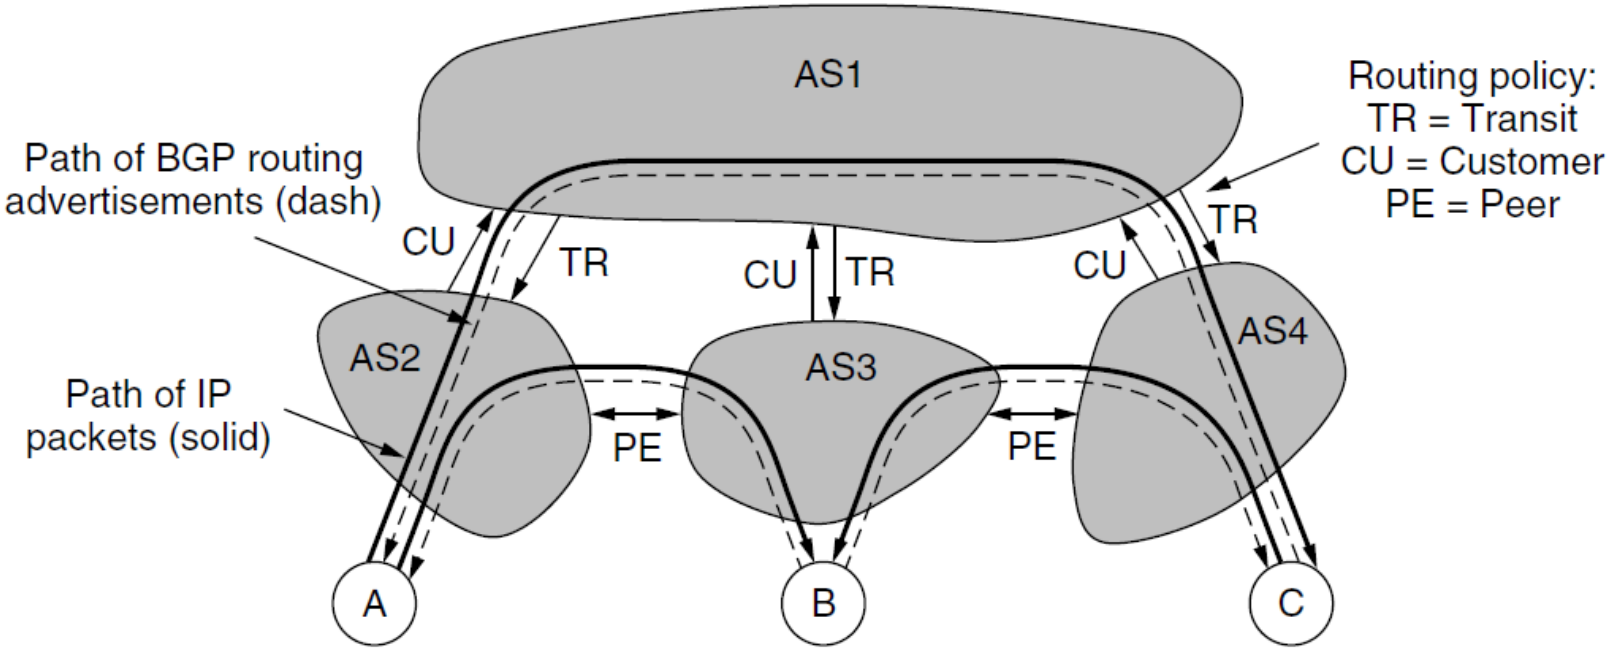

Connecting Networks with Different Protocols

If source and destination networks use different protocols, they cannot communicate.

Network A:
Uses ‘cars’ protocol.

Network B:
Uses ‘boats’ protocol.

Network C:
Uses ‘planes’ protocol.

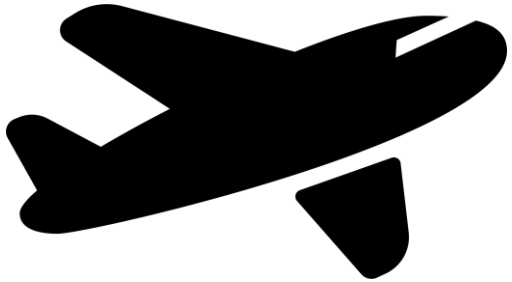

Tunneling

Used to route IPv6 packets over IPv4 networks

Q: Can you name a (dis)advantage?

If an intermediate network uses different protocols, they can communicate by tunneling.

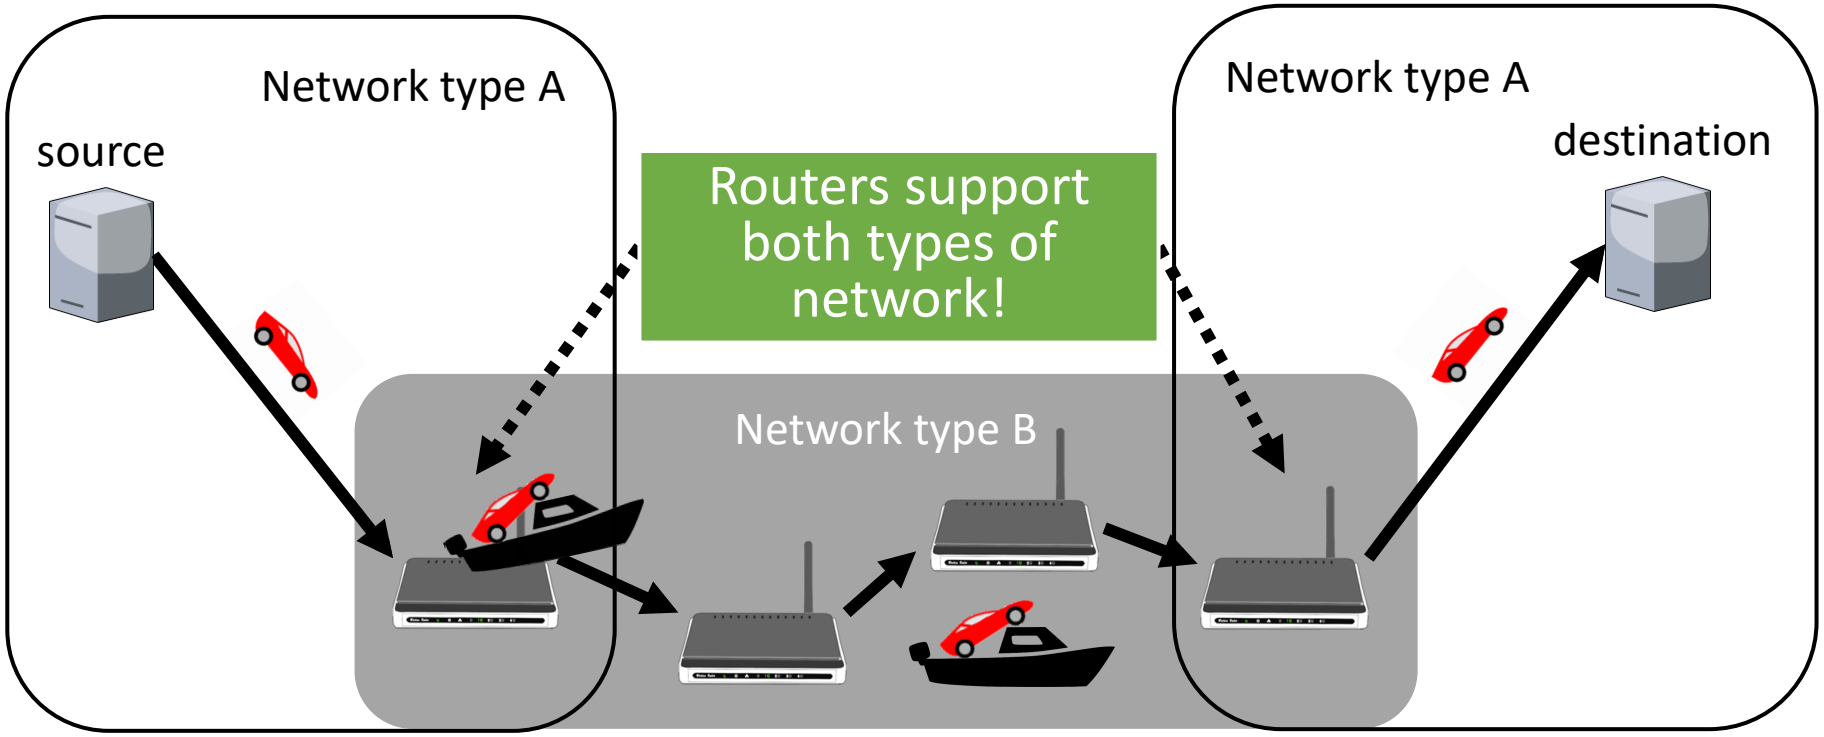

Business as usual
Packets in packets in …

Data wrapped in headers from multiple networking layers.

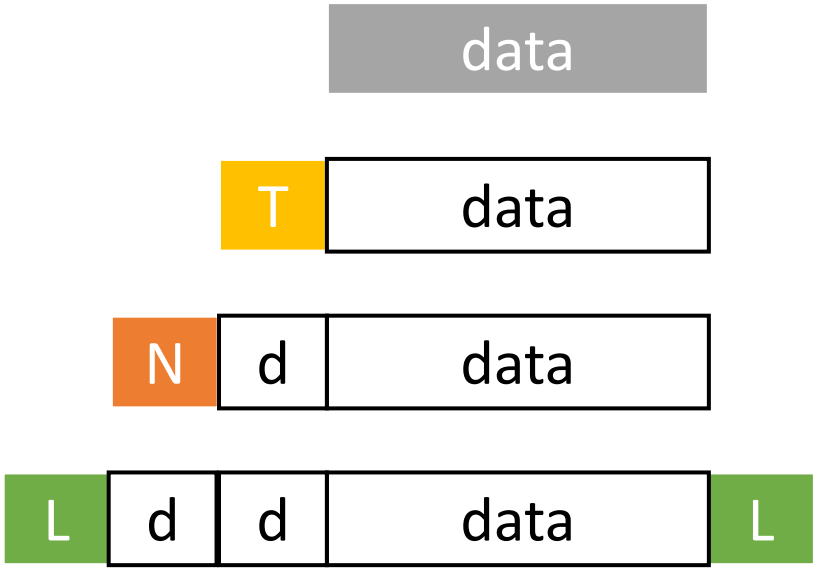

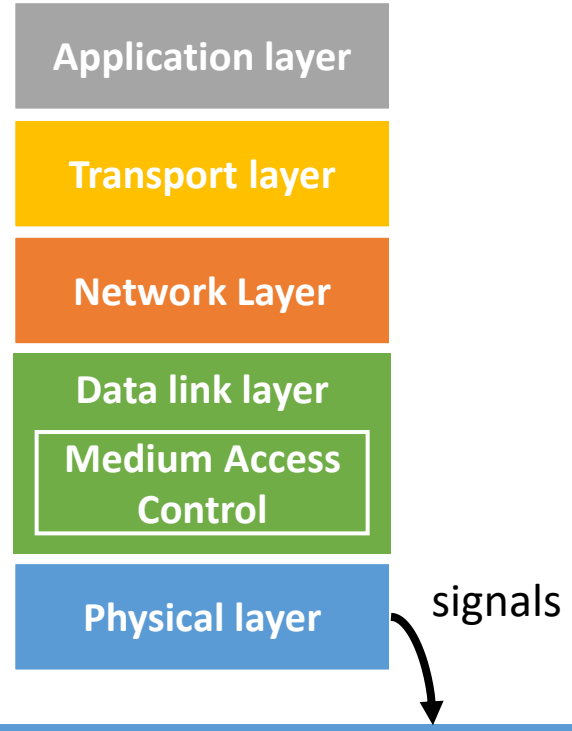

Tunneling
Packets in packets in packets in …

Data wrapped in headers from multiple networking layers.

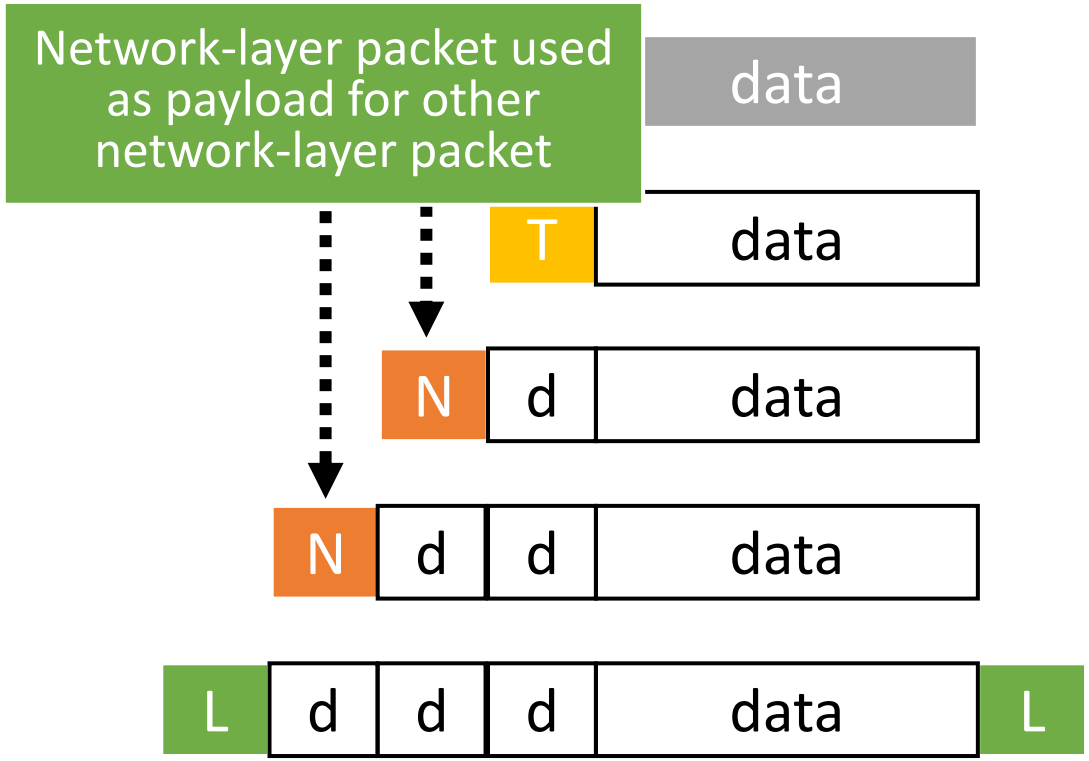

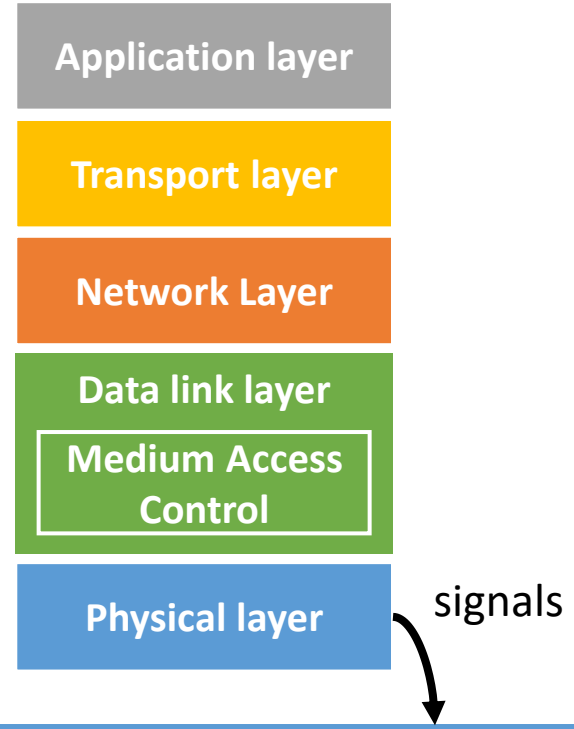

Packet fragmentation
Transparent fragmentation

Q: What can cause packet size limits?

Packet size can be limited by hardware, software, protocols, law, etc.

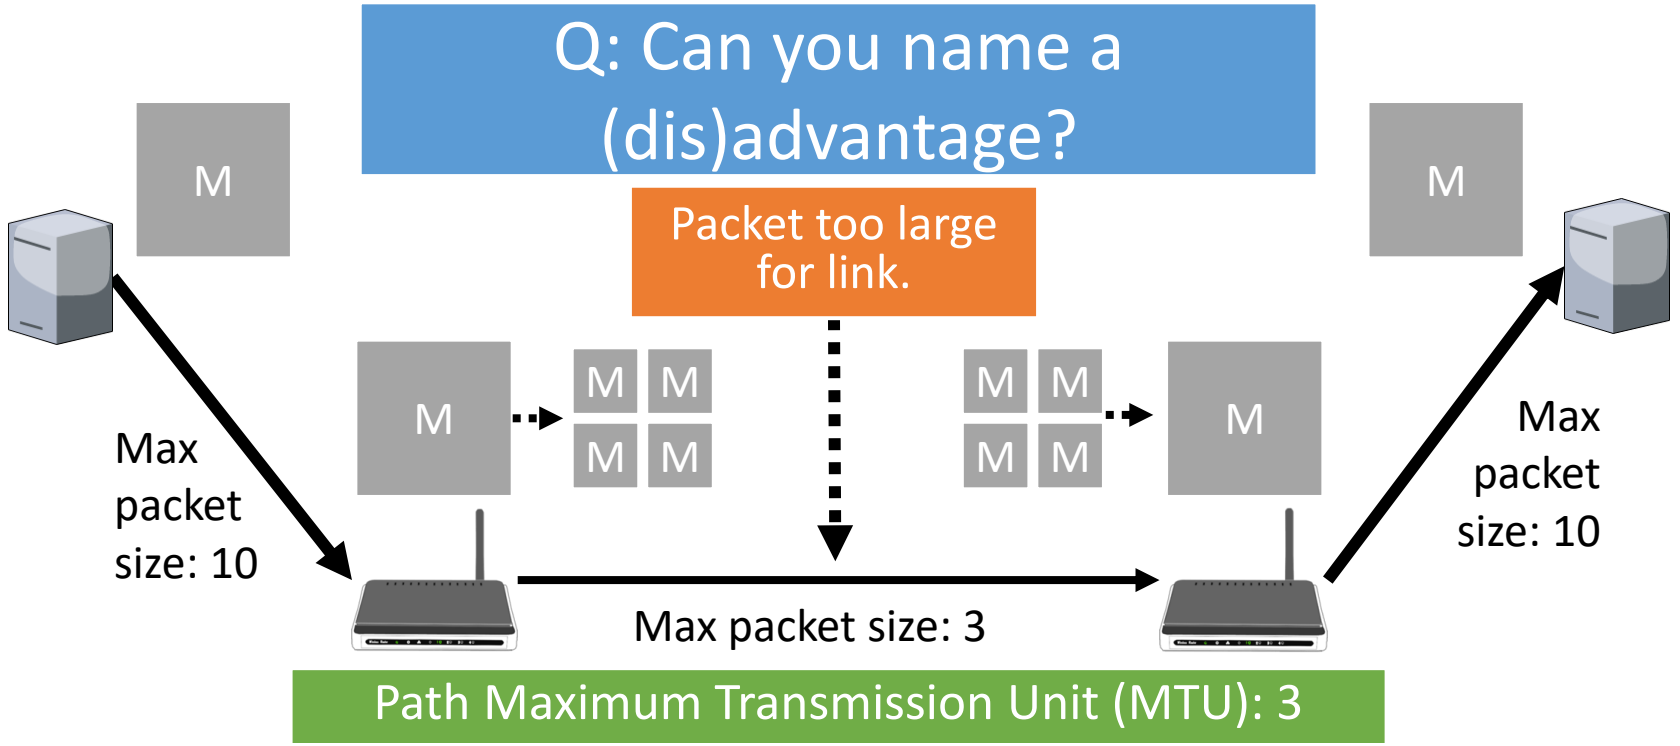

Used in IP!

Packet fragmentation
Nontransparent fragmentation

Packet size can be limited by hardware, software, protocols, law, etc.

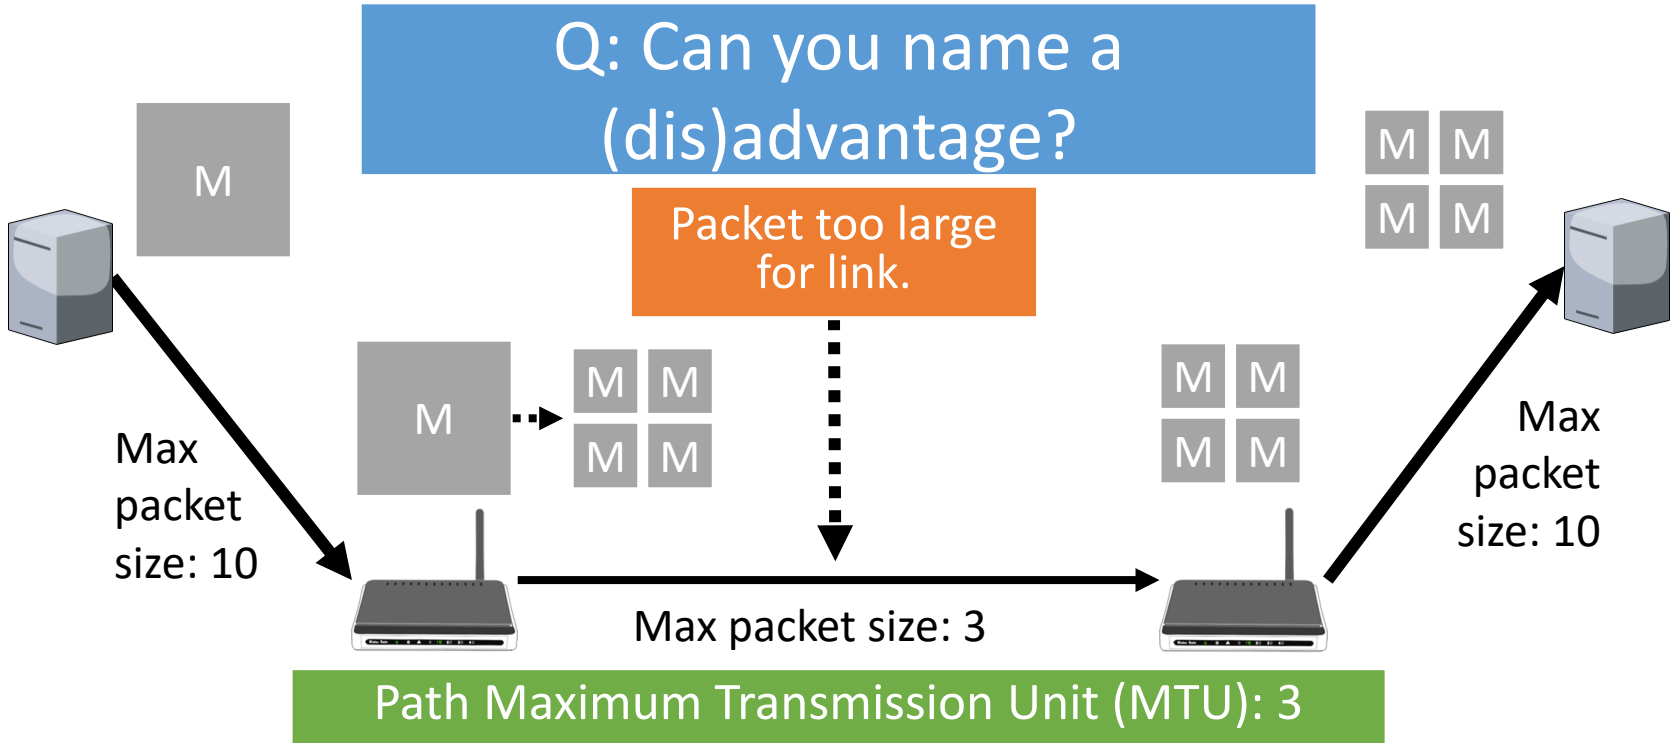

Avoiding packet fragmentation
MTU discovery

Packet size can be limited by hardware, software, protocols, law, etc.

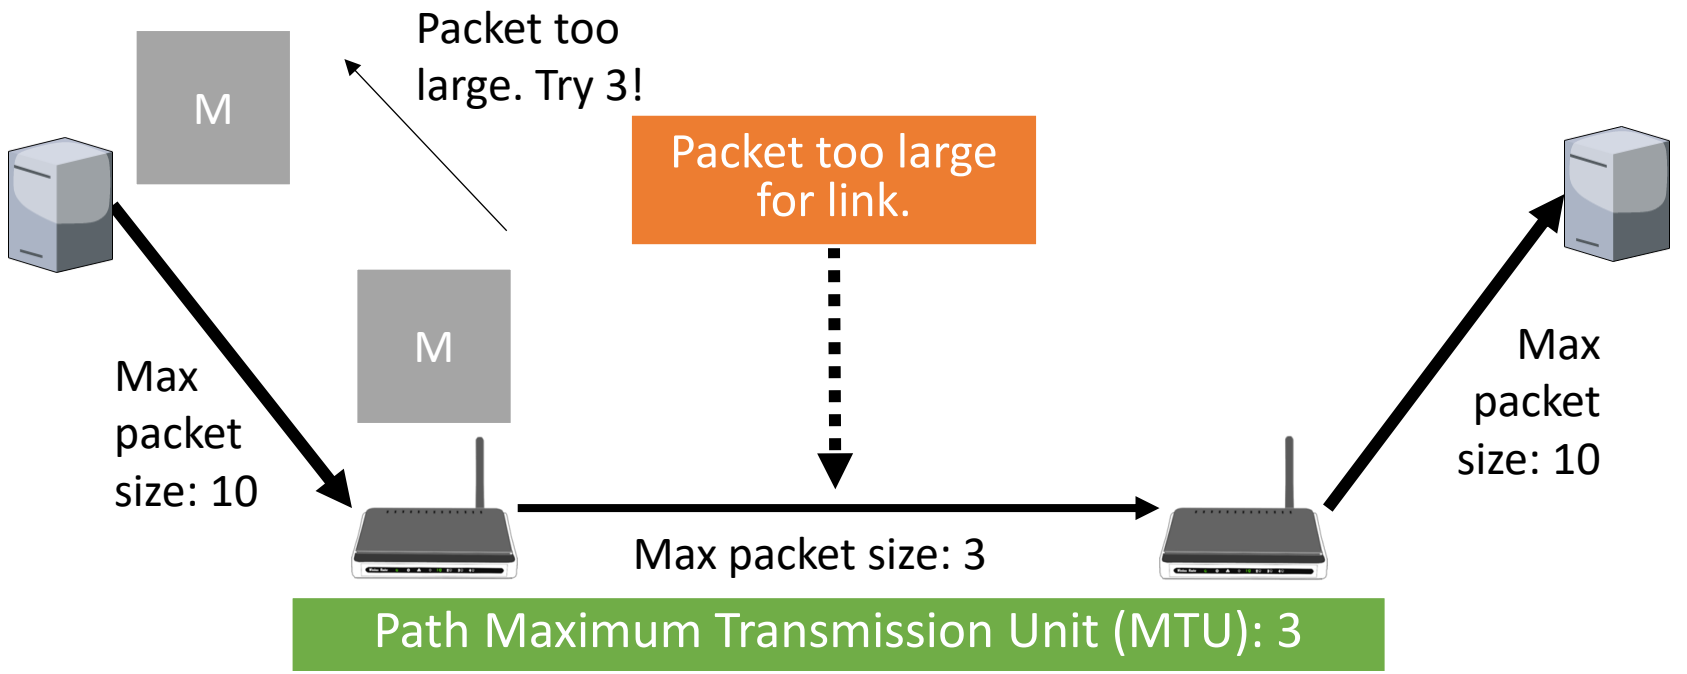

Used in IP!

Avoiding packet fragmentation
MTU discovery

Packet size can be limited by hardware, software, protocols, law, etc.

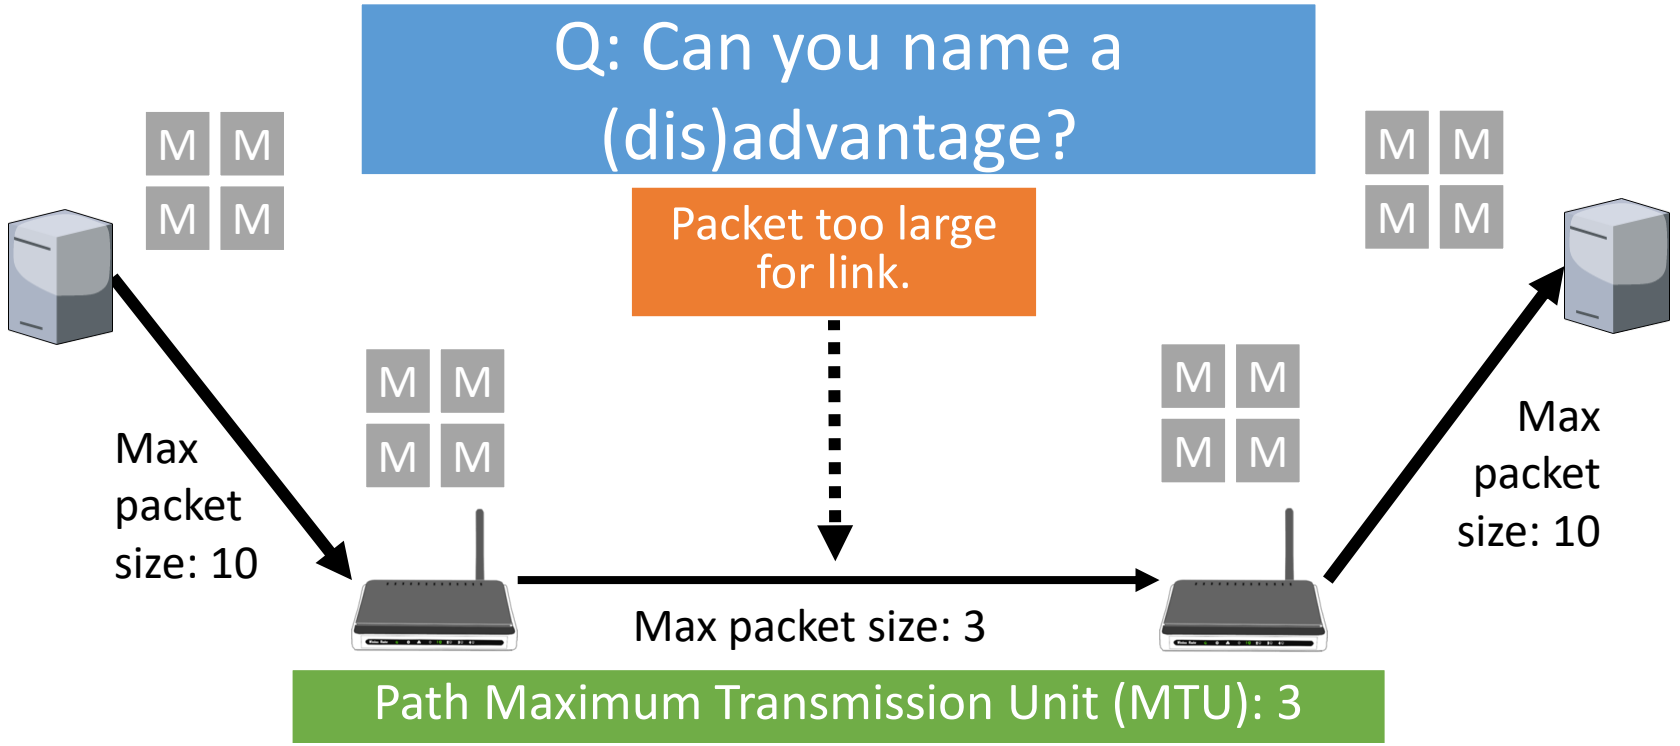

Extra Slide(s)

MultiProtocol Label Switching (MPLS)

Q: Why replace the labels at each hop?

Supports multiple networks protocols

Simplifies decision making at routers

Q: Why useful?