# Computer Networks X_400487

Lecture 7

Chapter 5: The Network Layer—Part 1

Lecturer: Jesse Donkervliet

VU Vrije Universiteit Amsterdam





## 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?

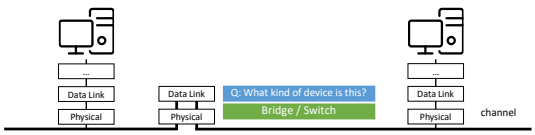






Image source: NASA



## 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*?

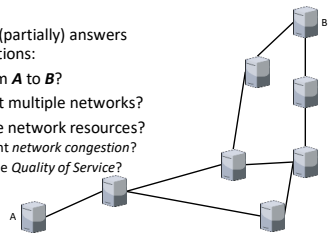

## The Network Layer

TCP/IP form the "narrow waist" of the Internet

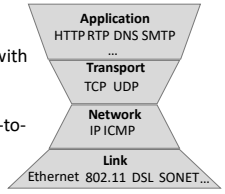


Heart of the networking model, together with the transport layer (next week)

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?

Router performs *store-and-forward* packet switching

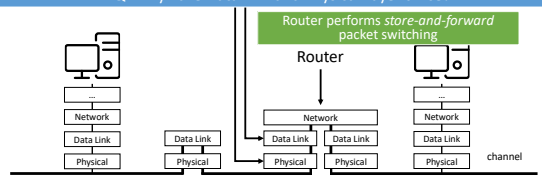


 

## The Network Layer

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

Next lecture

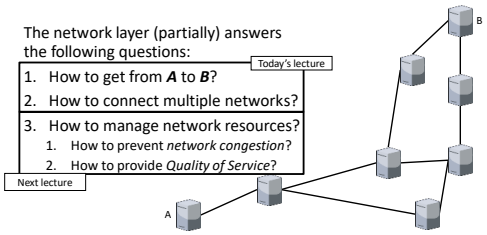


 

## Today's Lecture

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 table for C.

| To | Distance | Line |
|---|---|---|
| A | 7 | A |
| B | 59 | A |
| C | 0 | - |
| D | 75 | E |

We are directly connected to A

We are indirectly connected to B

If packet should go to D, minimum distance is 75. Forward to E




## Routing tables

For each packet, we need to know on which link to forward it
For this we use a routing table

Routing table for C.

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

We are **indirectly** connected to B




## Distance Vector Routing

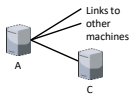


1. Send your *distance vector* to your neighbors
2. Update your *routing table* based on incoming *distance vectors*

Distance from A to B is 1

Distance from C to A is 7

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

| 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

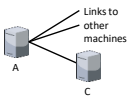


1. Send your *distance vector* to your neighbors
2. Update your *routing table* based on incoming *distance vectors*

*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

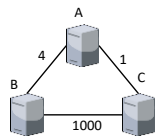


*Distance vector A*
A, -
B, 4
C, 1

*Distance vector B*
A, 4
B, -
C, 5

*Distance vector C*
A, 1
B, 5
C, -

## Distance Vector Routing
## Good News Propagation

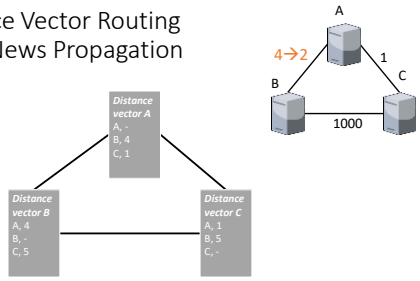





## Distance Vector Routing
## Good News Propagation

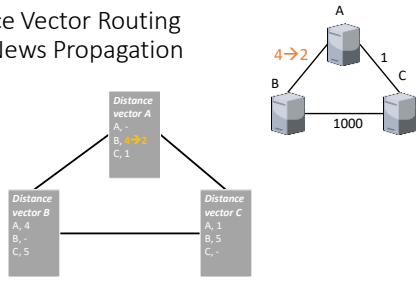





## Distance Vector Routing
## Good News Propagation

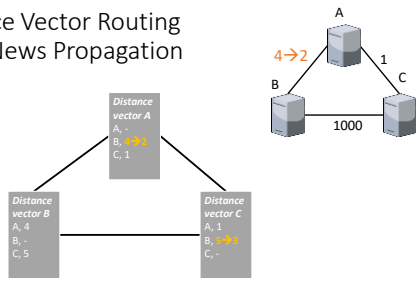





## Distance Vector Routing
## Bad News Propagation

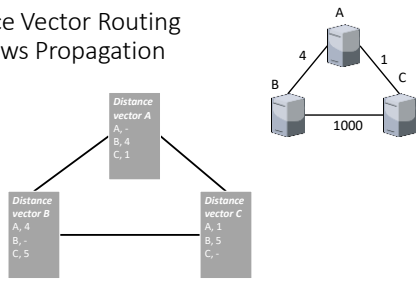





## Distance Vector Routing
## Bad News Propagation

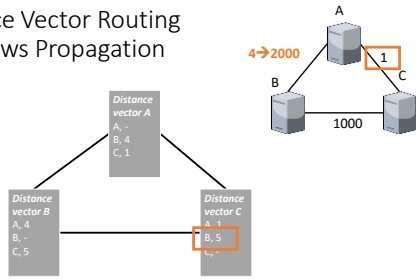





## Distance Vector Routing
## Bad News Propagation

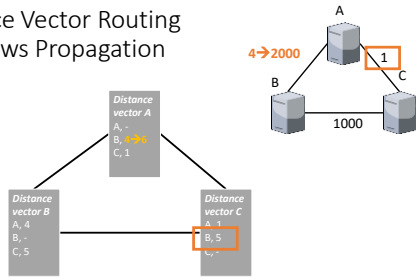

## Distance Vector Routing
## Bad News Propagation

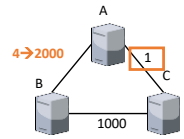


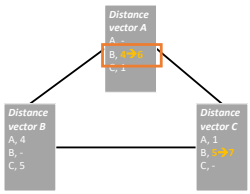




## Distance Vector Routing
## Bad News Propagation

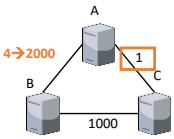


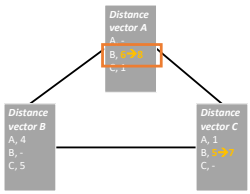




## Count to infinity problem
## When machine fails

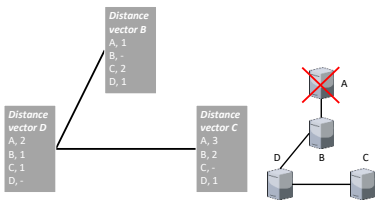




## Count to infinity problem
## When machine fails

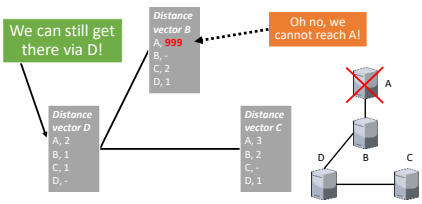




## Count to infinity problem
## When machine fails

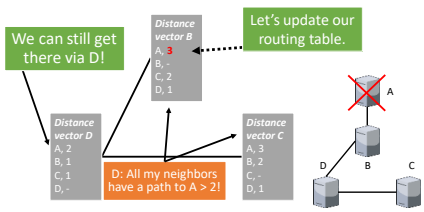




## Count to infinity problem
## When machine fails

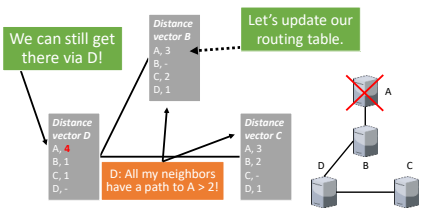

## Count to infinity problem
## When machine fails

This will go on for a while…

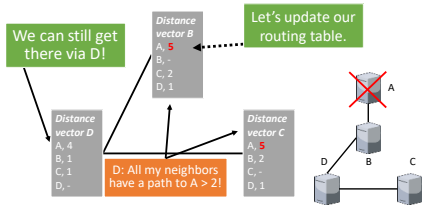





## Link state routing

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

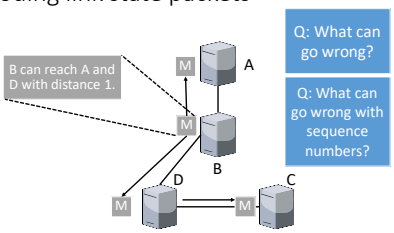


Q: What can go wrong?

Q: What can go wrong with sequence numbers?




## Link state routing
## Building a map of the network

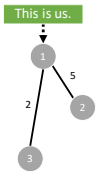





## Link state routing
## Building a map of the network

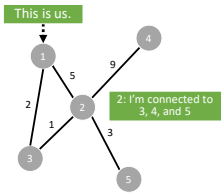





## Link state routing
## Building a map of the network

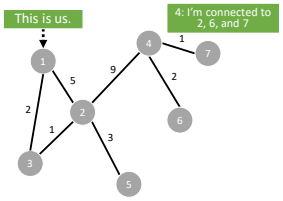

## Link state routing
## Building a map of the network

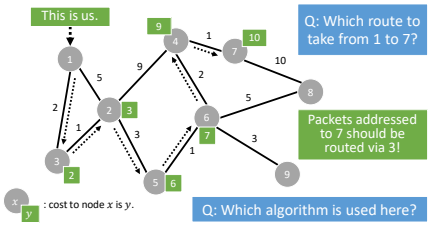





## Hierarchical routing

Reducing routing table sizes for large networks.

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

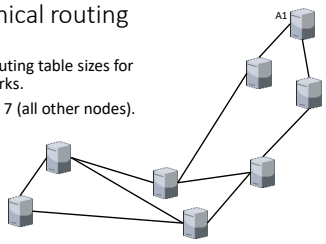





## Hierarchical routing

Reducing routing table sizes for large networks.

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

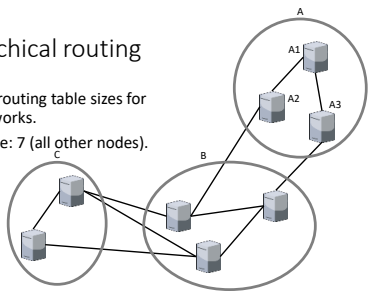





## Hierarchical routing

Reducing routing table sizes for large networks.

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

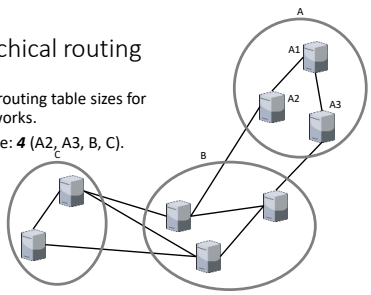





## Today's Lecture

1. Routing Algorithms
2. **Internetworking**
   1. **Real-World Routing in and between Autonomous Systems**
   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

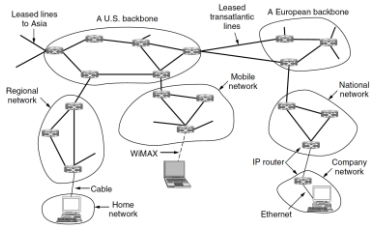


Tanenbaum, Andrew S., and David J. Wetherall. "Computer networks fifth edition." In *Pearson Education, Inc.* Prentice Hall, 2011.




## We've almost come full circle: A Hierarchical Topology

A multitree?

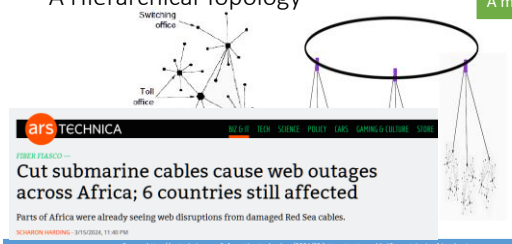


Source: https://arstechnica.com/information-technology/2024/03/internet-outages-hit-13-countries-in-africa-due-to-undersea-cable-damage/




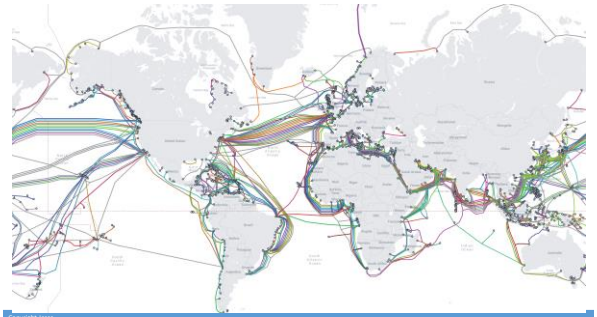

Source: https://www.submarinecablemap.com/




## 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)

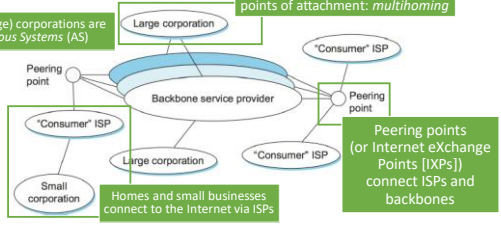


Homes and small businesses connect to the Internet via ISPs

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

Peterson, Larry L., and Bruce S. Davie. *Computer networks: a systems approach*. Elsevier, 2012.




## 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

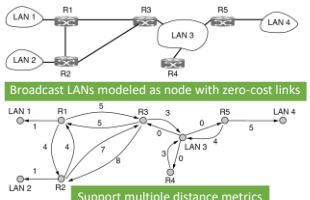


Broadcast LANs modeled as node with zero-cost links

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




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.

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

Area border router exchanges summarized routing info between areas

Routers only know topology of the area they are in

Q: Why support hierarchical routing?

For destinations not in area: go through backbone

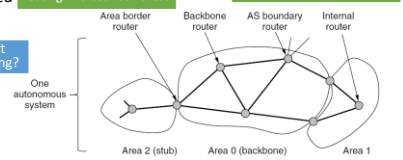


All areas connected to Area 0 (backbone)




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*

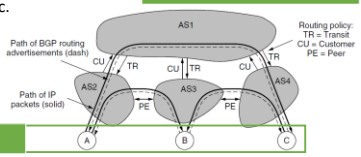


Autonomous Systems connect at IXPs

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***)

Path Vectors

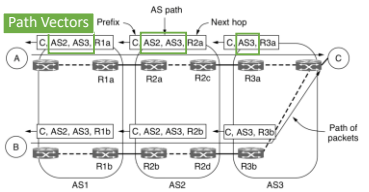


Routing based on policies

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.

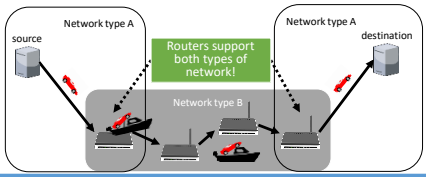

## Business as usual
## Packets in packets in packets in …

Data wrapped in headers from multiple networking layers.

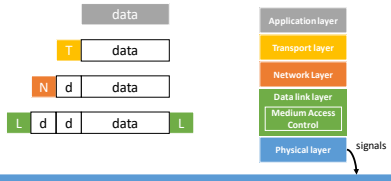





## Tunneling
## Packets in packets in packets in …

Data wrapped in headers from multiple networking layers.

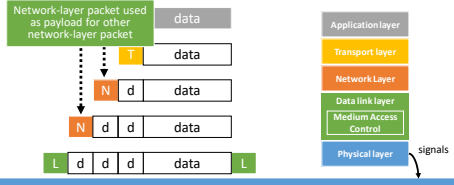





## Packet fragmentation
## Transparent fragmentation

Q: What can cause packet size limits?

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

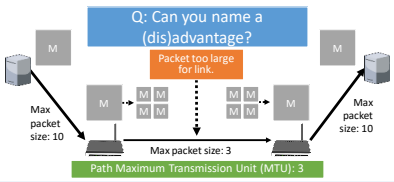





## Packet fragmentation
## Nontransparent fragmentation

Used in IP!

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

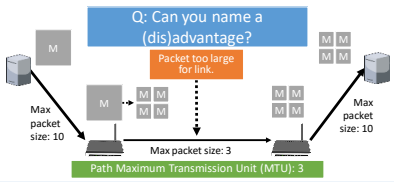





## ***Avoiding*** packet fragmentation
## MTU discovery

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

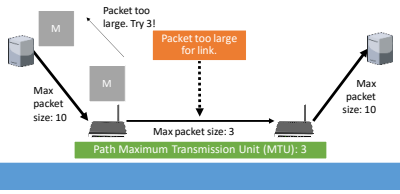





## ***Avoiding*** packet fragmentation
## MTU discovery

Used in IP!

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