

Last time

- IP addressing
 - ◆ addressing, subnets, CIDR
 - ◆ address aggregation
- ARP
 - ◆ Learning other hosts' MAC addresses
 - ◆ Same LAN only
- DHCP
 - ◆ Learning your own IP address
- ICMP
 - ◆ Internet “error messages”
 - ◆ How traceroute works
- IPv6
 - ◆ Differences from IPv4

This time



- Transitioning to IPv6
- Routing
 - ◆ Link-state routing
 - ◆ Distance-vector routing

Transitioning From IPv4 To IPv6

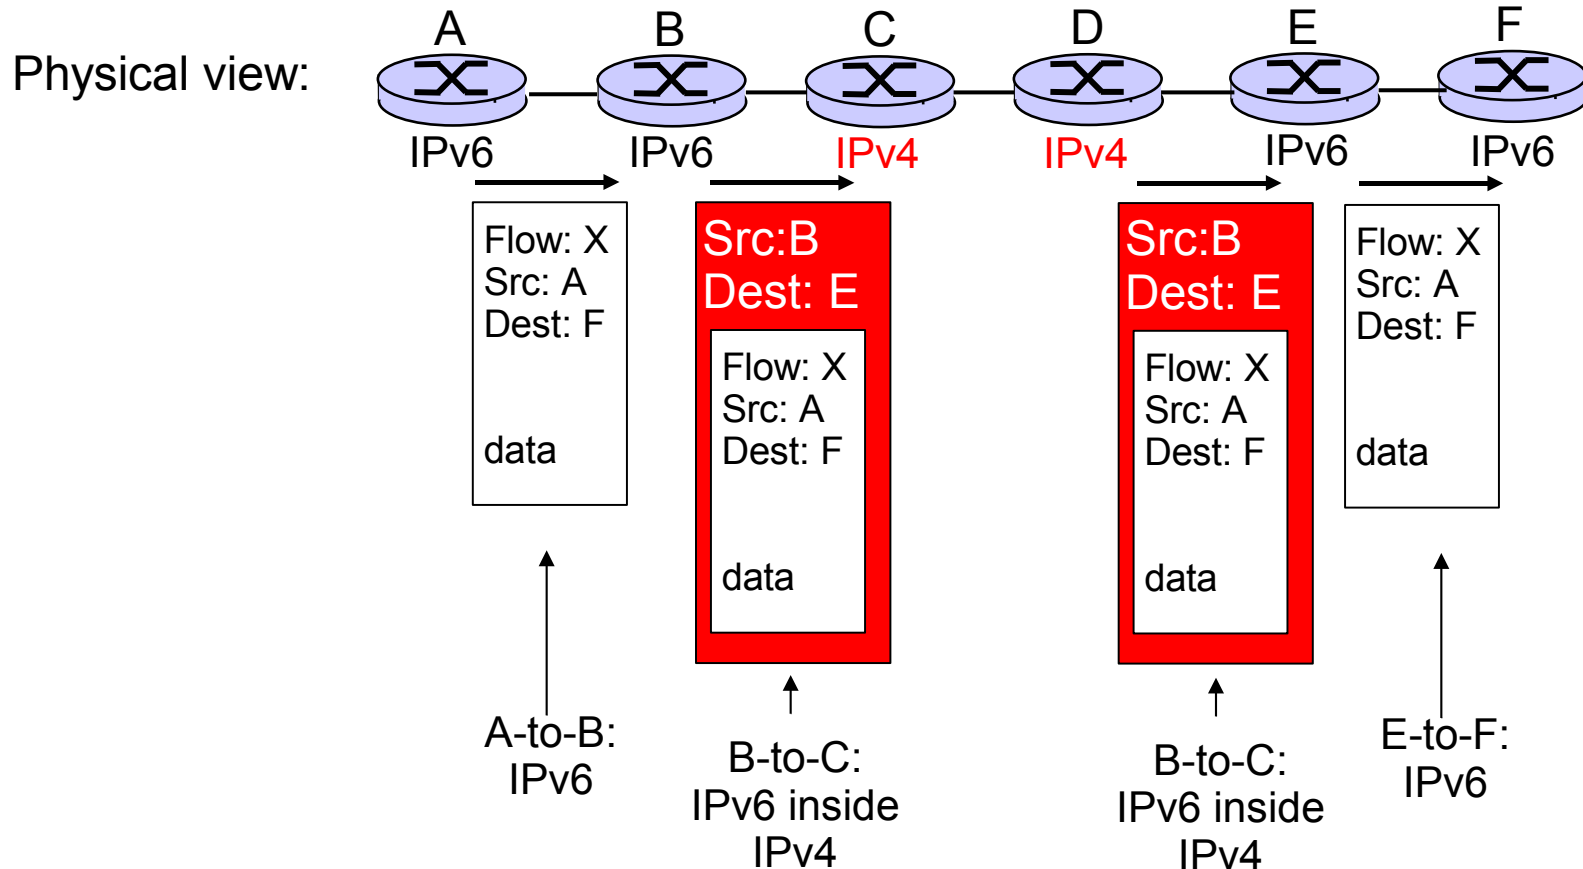

- Not all routers can be upgraded simultaneously
 - ◆ no “flag days”
 - ◆ How will the network operate with mixed IPv4 and IPv6 routers?

- *Tunneling*: IPv6 carried as payload in IPv4 datagram among IPv4 routers

Tunneling



Tunneling




Gateway

- Interaction between IPv6 clients and IPv4 servers – or vice versa?
 - ◆ tunneling does not help
- Solution: special “IPv4” address range in IPv6 addresses
- Need protocol gateway to convert between IPv4 and IPv6
- More complicated than tunneling


Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- 4.4 IP: Internet Protocol
 - ◆ Datagram format
 - ◆ IPv4 addressing
 - ◆ ICMP
 - ◆ IPv6
- 4.5 **Routing algorithms**
 - ◆ Link state
 - ◆ Distance Vector
 - ◆ Hierarchical routing
- 4.6 Routing in the Internet
 - ◆ RIP
 - ◆ OSPF
 - ◆ BGP
- 4.7 Broadcast and multicast routing

Interplay between routing, forwarding



Graph abstraction



Graph: $G = (N, E)$


$N =$ set of routers = $\{ u, v, w, x, y, z \}$

$E =$ set of links = $\{ (u, v), (u, x), (v, x), (v, w), (x, w), (x, y), (w, y), (w, z), (y, z) \}$

Remark: Graph abstraction is useful in other network contexts

Example: P2P, where N is set of peers and E is set of TCP connections

Graph abstraction: costs



- $c(x,x')$ = cost of link (x,x')
 - e.g., $c(w,z) = 5$
- cost could always be 1, or inversely related to bandwidth, or directly related to congestion

Cost of path $(x_1, x_2, x_3, \dots, x_p) = c(x_1, x_2) + c(x_2, x_3) + \dots + c(x_{p-1}, x_p)$

Question: What's the least-cost path between u and z ?

Routing algorithm: algorithm that finds least-cost path

Routing Algorithm classification

Global or decentralized information?

Global:

- all routers have complete topology, link cost info
- “link state” algorithms

Decentralized:

- router knows physically-connected neighbours, link costs to neighbours
- iterative process of computation, exchange of info with neighbours
- “distance vector” algorithms

Static or dynamic?

Static:

- routes change slowly over time

Dynamic:

- routes change more quickly
 - ◆ periodic update
 - ◆ in response to link cost changes

Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- 4.4 IP: Internet Protocol
 - ◆ Datagram format
 - ◆ IPv4 addressing
 - ◆ ICMP
 - ◆ IPv6
- 4.5 Routing algorithms
 - ◆ Link state
 - ◆ Distance Vector
 - ◆ Hierarchical routing
- 4.6 Routing in the Internet
 - ◆ RIP
 - ◆ OSPF
 - ◆ BGP
- 4.7 Broadcast and multicast routing

A Link-State Routing Algorithm

Dijkstra's algorithm

- net topology, link costs known to all nodes
 - ◆ accomplished via “link state broadcast”
 - ◆ all nodes have same info
- computes least cost paths from one node (“source”) to all other nodes
 - ◆ gives **forwarding table** for that node
- iterative: after k iterations, know least cost path to k destinations

Notation:

- $c(x,y)$: link cost from node x to y ; $= \infty$ if not direct neighbours
- $D(v)$: current value of cost of path from source to destination v
- $p(v)$: predecessor node along path from source to v
- N' : set of nodes whose least cost path definitively known

Dijkstra's Algorithm

1 **Initialization:**

2 $N' = \{u\}$

3 for all nodes v

4 if v adjacent to u

5 then $D(v) = c(u,v)$

6 else $D(v) = \infty$

7

8 **Loop**

9 find w not in N' such that $D(w)$ is a minimum

10 add w to N'

11 update $D(v)$ for all v adjacent to w and not in N' :

12 $D(v) = \min(D(v), D(w) + c(w,v))$

13 /* new cost to v is either old cost to v or known


14 shortest path cost to w plus cost from w to v */

15 **until all nodes in N'**




Dijkstra's algorithm: example

Step	N'	D(v),p(v)	D(w),p(w)	D(x),p(x)	D(y),p(y)	D(z),p(z)
0	u	2,u	5,u	1,u	∞	∞
1	ux	2,u	4,x		2,x	∞
2	uxy	2,u	3,y			4,y
3	uxyv		3,y			4,y
4	uxyvw					4,y
5	uxyvwz					



Dijkstra's algorithm: example (2)

Resulting shortest-path tree from u:



Resulting forwarding table in u:

destination	link
v	(u,v)
x	(u,x)
y	(u,x)
w	(u,x)
z	(u,x)





Dijkstra's algorithm, discussion

Algorithm complexity: n nodes

- each iteration: need to check all nodes, w, not in N'
- $n(n+1)/2$ comparisons: $O(n^2)$
- more efficient implementations possible: $O(n \log n)$

Oscillations possible:

- e.g., link cost = amount of carried traffic



Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- 4.4 IP: Internet Protocol
 - ◆ Datagram format
 - ◆ IPv4 addressing
 - ◆ ICMP
 - ◆ IPv6
- 4.5 Routing algorithms
 - ◆ Link state
 - ◆ Distance Vector
 - ◆ Hierarchical routing
- 4.6 Routing in the Internet
 - ◆ RIP
 - ◆ OSPF
 - ◆ BGP
- 4.7 Broadcast and multicast routing

Distance Vector Algorithm

Bellman-Ford Equation (dynamic programming)

Define


$d_x(y) :=$ cost of least-cost path from x to y

Then

$$d_x(y) = \min_v \{c(x,v) + d_v(y)\}$$

where min is taken over all neighbours v of x

Bellman-Ford example



Clearly, $d_v(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

B-F equation says:

$$\begin{aligned} d_u(z) &= \min \{ c(u,v) + d_v(z), \\ &\quad c(u,x) + d_x(z), \\ &\quad c(u,w) + d_w(z) \} \\ &= \min \{ 2 + 5, \\ &\quad 1 + 3, \\ &\quad 5 + 3 \} = 4 \end{aligned}$$

Node that achieves minimum is next hop in shortest path → forwarding table

Distance Vector Algorithm

- $D_x(y)$ = estimate of least cost from x to y
- Node x knows cost to each neighbour v:
 $c(x,v)$
- Node x maintains distance vector
 $D_x = [D_x(y): y \in N]$
- Node x also maintains its neighbours' distance vectors
 - ◆ For each neighbour v, x maintains
 $D_v = [D_v(y): y \in N]$

Distance vector algorithm (4)

Basic idea:

- Each node periodically sends its own distance vector estimate to neighbours
- When a node x receives new DV estimate from neighbour, it updates its own DV using B-F equation:

$$D_x(y) \leftarrow \min_v \{c(x,v) + D_v(y)\} \quad \text{for each node } y \in N$$

- Under minor, natural conditions, the estimate $D_x(y)$ converges to the actual least cost $d_x(y)$

Distance Vector Algorithm (5)

Iterative, asynchronous:

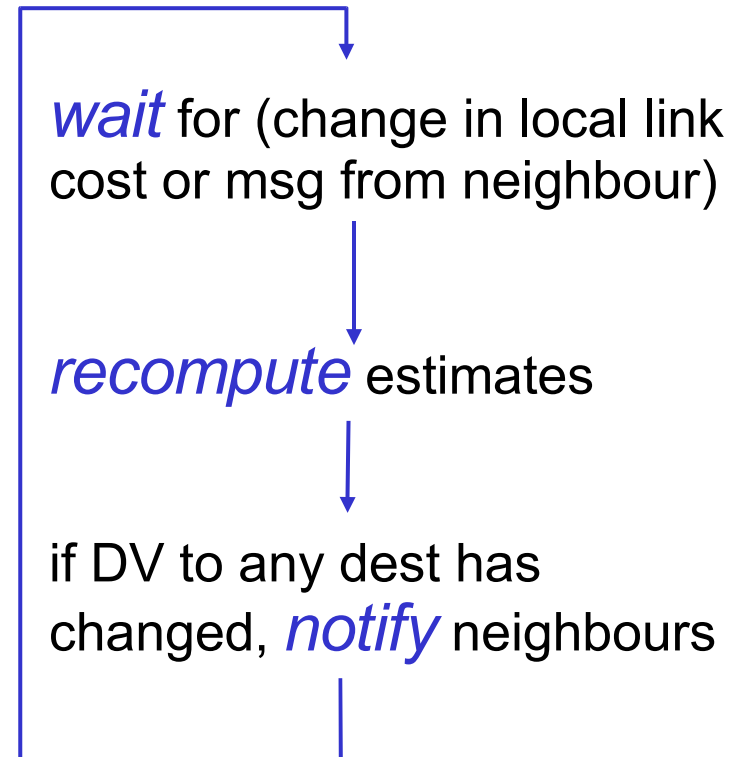
each local iteration caused by:

- local link cost change
- DV update message from neighbour

Distributed:

- each node notifies neighbours *only* when its DV changes
 - ◆ neighbours then notify their neighbours if necessary

Each node:



$$D_x(y) = \min\{c(x,y) + D_y(y), c(x,z) + D_z(y)\}$$

$$= \min\{2+0, 7+1\} = 2$$

$$D_x(z) = \min\{c(x,y) + D_y(z), c(x,z) + D_z(z)\}$$

$$= \min\{2+1, 7+0\} = 3$$

node x table

		cost to		
		x	y	z
from	x	0	2	7
	y	∞	∞	∞
	z	∞	∞	∞


		cost to		
		x	y	z
from	x	0	2	3
	y	2	0	1
	z	7	1	0

node y table

		cost to		
		x	y	z
from	x	∞	∞	∞
	y	2	0	1
	z	∞	∞	∞

node z table

		cost to		
		x	y	z
from	x	∞	∞	∞
	y	∞	∞	∞
	z	7	1	0



time

$$D_x(y) = \min\{c(x,y) + D_y(y), c(x,z) + D_z(y)\}$$

$$= \min\{2+0, 7+1\} = 2$$

$$D_x(z) = \min\{c(x,y) + D_y(z), c(x,z) + D_z(z)\}$$

$$= \min\{2+1, 7+0\} = 3$$

node x table

		cost to		
		x	y	z
from	x	0	2	7
	y	∞	∞	∞
	z	∞	∞	∞

node y table

		cost to		
		x	y	z
from	x	∞	∞	∞
	y	2	0	1
	z	∞	∞	∞

node z table

		cost to		
		x	y	z
from	x	∞	∞	∞
	y	∞	∞	∞
	z	7	1	0

		cost to		
		x	y	z
from	x	0	2	3
	y	2	0	1
	z	7	1	0


		cost to		
		x	y	z
from	x	0	2	7
	y	2	0	1
	z	7	1	0

		cost to		
		x	y	z
from	x	0	2	7
	y	2	0	1
	z	3	1	0

		cost to		
		x	y	z
from	x	0	2	3
	y	2	0	1
	z	3	1	0

		cost to		
		x	y	z
from	x	0	2	3
	y	2	0	1
	z	3	1	0

		cost to		
		x	y	z
from	x	0	2	3
	y	2	0	1
	z	3	1	0



time →

Recap

- Transitioning to IPv6
 - ◆ Tunneling
 - ◆ Gateways

- Routing
 - ◆ Graph abstraction
 - ◆ Link-state routing
 - Dijkstra's Algorithm
 - ◆ Distance-vector routing
 - Bellman-Ford Equation

Next time

- Distance vector link cost changes
- Hierarchical routing
- Routing protocols