

Chapter 4: outline

- 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

Network Layer 4-81

Routing

- ❖ Routing in the Internet – combination of rules and procedures that allow router to:

- Inform one another of "status of" or "changes" in the network
- Determine "best" routing paths in the network
- Transfer packets from a source host to a destination host along the best path

Routing protocol

Routing algorithm

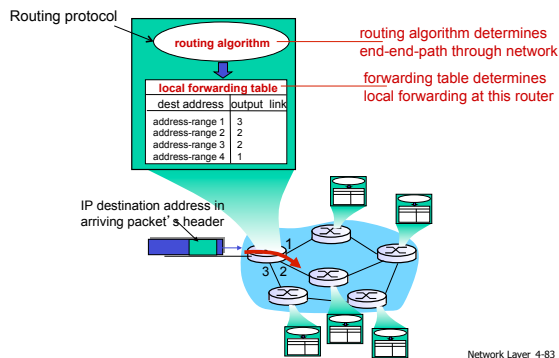
Packet forwarding

- ❖ Internet Routing Goals: accurate, rapid, low cost delivery of packets

- Route packets away from failed and temporarily congested nodes or links
- Avoid routing loops
- Adapt to varying traffic loads
- Low overhead

Network Layer 4-82

Interplay between routing, forwarding



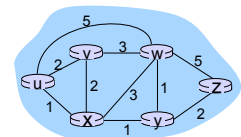
Routing Algorithm

- ❖ Routing Algorithm – heart of routing protocol, determines the best path between any two hosts in the network

- Best path = path that minimizes the *objective function* that the network operator tries to optimize

- Possible *objective functions*:

1. Number of hops
2. End-to-end delay
3. ISP cost, ...



Routing algorithm classification

Q: global or decentralized information?

global:

- ❖ all routers have complete topology, link cost info
- ❖ "link state" algorithms

decentralized:

- ❖ router knows physically-connected neighbors, link costs to neighbors
- ❖ iterative process of computation, exchange of info with neighbors
- ❖ "distance vector" algorithms

Q: static or dynamic?

static:

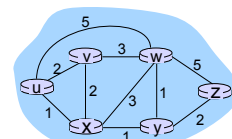
- ❖ routes change slowly over time

dynamic:

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

Network Layer 4-85

Graph abstraction



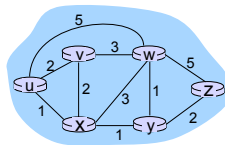
graph: $G = (N, E)$

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

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

Network Layer 4-86

Graph abstraction: costs



$c(x, x') = \text{cost of link } (x, x')$
e.g., $c(w, z) = 5$

$c(x, x) = 0$
 $c(x, y) \geq 0$ if nodes directly connected
Cost could be associated with bandwidth and/or congestion

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

key question: what is the least-cost path between u and z ?
routing algorithm: algorithm that finds that least cost path

Network Layer 4-87

Chapter 4: outline

- 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

Network Layer 4-88

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 neighbors
- ❖ $D(v)$: current value of cost of path from source to v
- ❖ $p(v)$: predecessor node along path from source to v
- ❖ N' : set of nodes whose least cost path definitively known

Network Layer 4-89

Dijkstra's Algorithm

notation:

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

- 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'**

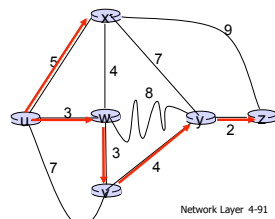
Network Layer 4-90

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	7, u	3, u	5, u	∞	∞
1	uw	6, w	5, u	11, w	∞	∞
2	uwx	6, w	11, w	14, x	∞	∞
3	uwxv	∞	10, v	14, x	∞	∞
4	uwxvy	∞	∞	12, y	∞	∞
5	uwxvyz	∞	∞	∞	∞	∞

notes:

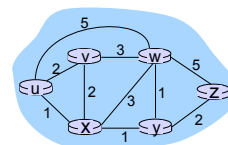
- ❖ construct shortest path tree by tracing predecessor nodes
- ❖ ties can exist (can be broken arbitrarily)



Network Layer 4-91

Dijkstra's algorithm: another 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	∞	∞	∞	∞	∞

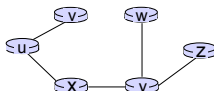


destination	link
v	(u, v)
x	(u, x)
y	(u, y)
w	(u, w)
z	(u, z)

Network Layer 4-92

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)

Network Layer 4-93

Chapter 4: outline

- 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

Network Layer 4-94

Distance vector algorithm

Bellman-Ford equation (dynamic programming)

let

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

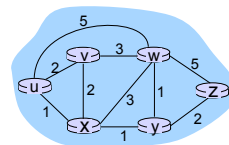
then

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

\min taken over all neighbors v of x
 $c(x,v)$ cost to neighbor v
 $d_v(y)$ cost from neighbor v to destination y

Network Layer 4-95

Bellman-Ford example



clearly, $d_u(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

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 achieving minimum is next hop in shortest path, used in forwarding table

Network Layer 4-96

Distance vector algorithm

- ❖ $D_x(y)$ = estimate of least cost from x to y
 - x maintains distance vector $D_x = [D_x(y): y \in N]$
- ❖ node x:
 - knows cost to each neighbor v: $c(x,v)$
 - maintains its neighbors' distance vectors. For each neighbor v, x maintains $D_v = [D_v(y): y \in N]$

Network Layer 4-97

Distance vector algorithm

key idea:

- ❖ from time-to-time, each node sends its own distance vector estimate to neighbors
- ❖ when x receives new DV estimate from neighbor, it updates its own DV using B-F equation:

$$D_x(y) \leftarrow \min_v \{ c(x,v) + D_v(y) \} \text{ for each node } y \in N$$
- ❖ under minor, natural conditions, the estimate $D_x(y)$ converge to the actual least cost $d_x(y)$

Network Layer 4-98

Distance vector algorithm

iterative, asynchronous:

- each local iteration caused by:
 - ❖ local link cost change
 - ❖ DV update message from neighbor

distributed:

- ❖ each node notifies neighbors *only* when its DV changes
 - neighbors then notify their neighbors if necessary

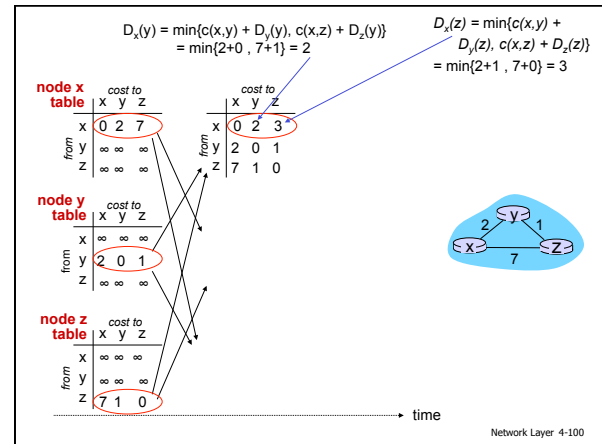
each node:

wait for (change in local link cost or msg from neighbor)

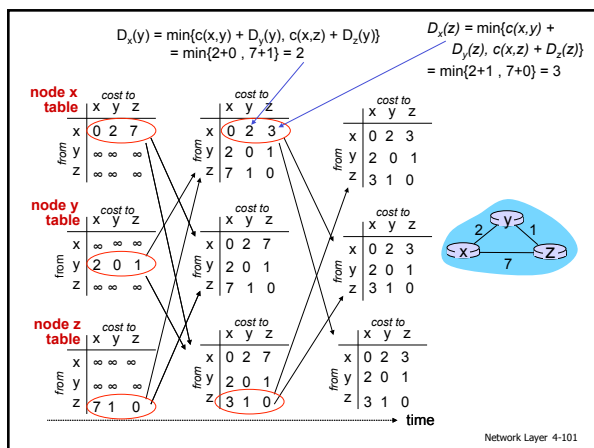
recompute estimates

if DV to any dest has changed, *notify* neighbors

Network Layer 4-99



Network Layer 4-100



Network Layer 4-101

Distance vector: link cost changes

link cost changes:

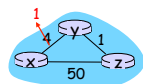
- ❖ node detects local link cost change
- ❖ updates routing info, recalculates distance vector
- ❖ if DV changes, notify neighbors

"good news travels fast"

t_0 : y detects link-cost change, updates its DV, informs its neighbors.

t_1 : z receives update from y, updates its table, computes new least cost to x, sends its neighbors its DV.

t_2 : y receives z's update, updates its distance table. y's least costs do not change, so y does not send a message to z.



Network Layer 4-102

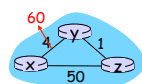
Distance vector: link cost changes

link cost changes:

- ❖ node detects local link cost change
- ❖ *bad news travels slow* - "count to infinity" problem!
- ❖ 44 iterations before algorithm stabilizes: see text

poisoned reverse:

- ❖ If Z routes through Y to get to X:
 - Z tells Y its (Z's) distance to X is infinite (so Y won't route to X via Z)
- ❖ will this completely solve count to infinity problem?



Network Layer 4-103

Comparison of LS and DV algorithms

message complexity

- ❖ **LS**: with n nodes, E links, $O(nE)$ msgs sent
- ❖ **DV**: exchange between neighbors only
 - convergence time varies

speed of convergence

- ❖ **LS**: $O(n^2)$ algorithm requires $O(nE)$ msgs
 - may have oscillations
- ❖ **DV**: convergence time varies
 - may be routing loops
 - count-to-infinity problem

robustness: what happens if router malfunctions?

LS:

- node can advertise incorrect *link* cost
- each node computes only its own table

DV:

- DV node can advertise incorrect *path* cost
- each node's table used by others
 - error propagate thru network

Network Layer 4-104