

Chapter 4

Network Layer

Reti di Elaboratori

Corso di Laurea in Informatica

Università degli Studi di Roma "La Sapienza"

Canale A-L

Prof.ssa Chiara Petrioli

Parte di queste slide sono state prese dal materiale associato al libro
Computer Networking: A Top Down Approach, 5th edition.

All material copyright 1996-2009

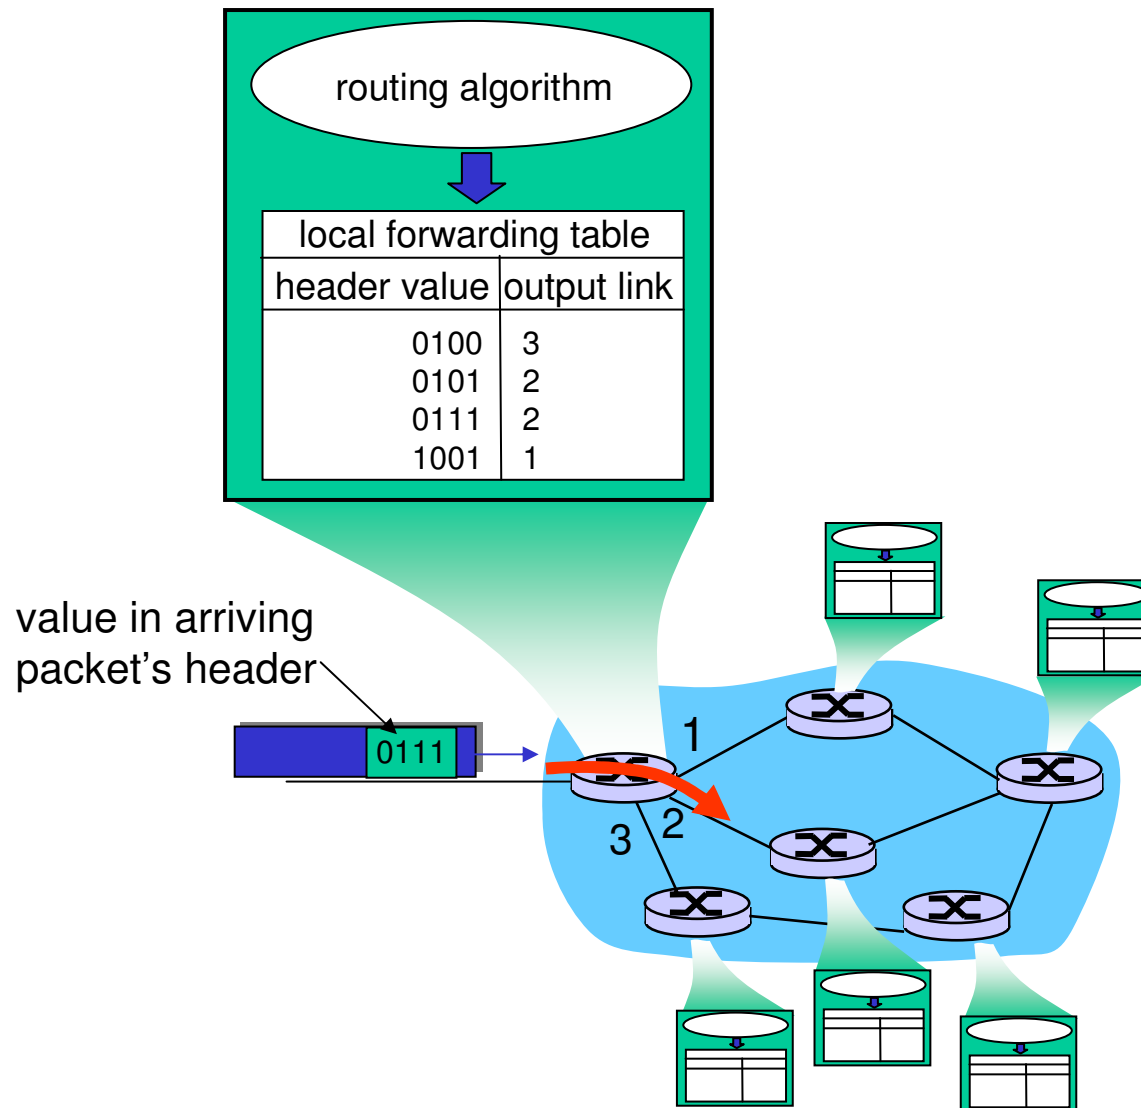
J.F Kurose and K.W. Ross, All Rights Reserved

Thanks also to Antonio Capone, Politecnico di Milano, Giuseppe Bianchi and
Francesco LoPresti, Un. di Roma Tor Vergata


Chapter 4: Network Layer

- r 4.1 Introduction
- r 4.2 Virtual circuit and datagram networks
- r 4.3 What's inside a router
- r 4.4 IP: Internet Protocol
 - m Datagram format
 - m IPv4 addressing
 - m ICMP
 - m IPv6
- r 4.5 **Routing algorithms**
 - m Link state
 - m Distance Vector
 - m Hierarchical routing
- r 4.6 Routing in the Internet
 - m RIP
 - m OSPF
 - m BGP
- r 4.7 Broadcast and multicast routing

Interplay between routing, forwarding



Graph abstraction



Graph: $G = (N, E)$


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

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

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 inversely 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:

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

Decentralized:

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

Static or dynamic?

Static:

- r routes change slowly over time

Dynamic:

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

Chapter 4: Network Layer

- r 4.1 Introduction
- r 4.2 Virtual circuit and datagram networks
- r 4.3 What's inside a router
- r 4.4 IP: Internet Protocol
 - m Datagram format
 - m IPv4 addressing
 - m ICMP
 - m IPv6
- r 4.5 Routing algorithms
 - m Link state
 - m Distance Vector
 - m Hierarchical routing
- r 4.6 Routing in the Internet
 - m RIP
 - m OSPF
 - m BGP
- r 4.7 Broadcast and multicast routing

A Link-State Routing Algorithm

Dijkstra's algorithm


- r net topology, link costs known to all nodes
 - m accomplished via "link state broadcast"
 - m all nodes have same info
- r computes least cost paths from one node ('source") to all other nodes
 - m gives forwarding table for that node
- r iterative: after k iterations, know least cost path to k dest.'s

Notation:

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


Dijsktra'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


r each iteration: need to check all nodes, w , not in N

r $n(n+1)/2$ comparisons: $O(n^2)$

r more efficient implementations possible: $O(n \log n)$

Oscillations possible:

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



Chapter 4: Network Layer

- r 4.1 Introduction
- r 4.2 Virtual circuit and datagram networks
- r 4.3 What's inside a router
- r 4.4 IP: Internet Protocol
 - m Datagram format
 - m IPv4 addressing
 - m ICMP
 - m IPv6
- r 4.5 Routing algorithms
 - m Link state
 - m Distance Vector
 - m Hierarchical routing
- r 4.6 Routing in the Internet
 - m RIP
 - m OSPF
 - m BGP
- r 4.7 Broadcast and multicast routing

Bellman-Ford

Given a graph $G=(N,A)$ and a node s finds the shortest path from s to every node in N .

A shortest walk from s to i subject to the constraint that the walk contains at most h arcs and goes through node s only once, is denoted $\text{shortest}(\leq h)$ walk and its length is D^h_i .

Bellman-Ford rule:

Initiatilization $D^h_s=0$, for all h ; $w_{i,k} = \text{infinity}$ if (i,k) NOT in A ; $w_{k,k} = 0$;
 $D^0_i = \text{infinity}$ for all $i \neq s$.

Iteration:

$$D^{h+1}_i = \min_k [w_{i,k} + D^h_k]$$

Assumption: non negative cycles (this is the case in a network!!)

The Bellman-Ford algorithm first finds the one-arc shortest walk lengths, then the two-arc shortest walk length, then the three-arc...etc. \rightarrow distributed version used for routing

Bellman-Ford

$$D^{h+1}_i = \min_k [w_{i,k} + D^h_k]$$

Can be computed locally.

What do I need?

For each neighbor k , I need to know

-the cost of the link to it (known info)

-The cost of the best route from the neighbor k to the destination
(←this is an info that each of my neighbor has to send to me via messages)

In the real world: I need to know the best routes among each pair of nodes → we apply distributed Bellman Ford to get the best route for each of the possible destinations

Distance Vector Routing Algorithm -Distributed Bellman Ford

iterative:

- r continues until no nodes exchange info.
- r *self-terminating*: no "signal" to stop

asynchronous:

- r nodes need *not* exchange info/iterate in lock step!

Distributed, based on local info:

- r each node communicates *only* with directly-attached neighbors

Distance Table data structure

each node has its own


- r row for each possible destination
- r column for each directly-attached neighbor to node
- r example: in node X, for dest. Y via neighbor Z:

Cost associated to the (X,Z) link

$$D^X(Y,Z) = \text{distance from X to Y, via Z as next hop}$$
$$= c(X,Z) + \min_w \{D^Z(Y,w)\}$$

Info maintained at Z. Min must be communicated

Distance Table: example



Distance table in node E after the algorithm has converged

cost to destination via

$D^E()$	A	B	D
A	1	14	5
B	7	8	5
C	6	9	4
D	4	11	2

destination

$$D^E(C,D) = c(E,D) + \min_w \{D^D(C,w)\}$$

$$= 2+2 = 4$$

$$D^E(A,D) = c(E,D) + \min_w \{D^D(A,w)\}$$

$$= 2+3 = 5$$

$$D^E(A,B) = c(E,B) + \min_w \{D^B(A,w)\}$$

$$= 8+6 = 14$$

loop! Best path from D goes through E

loop!

Path B-C-D-E-A

First example

Distance table gives routing table

cost to destination via

D^E ()	A	B	D
A	1	14	5
B	7	8	5
C	6	9	4
D	4	11	2

destination

Outgoing link to use, cost

A	A,1
B	D,5
C	D,4
D	D,2

destination

Distance table → Routing table

Distance Vector Routing: overview

Iterative, asynchronous:

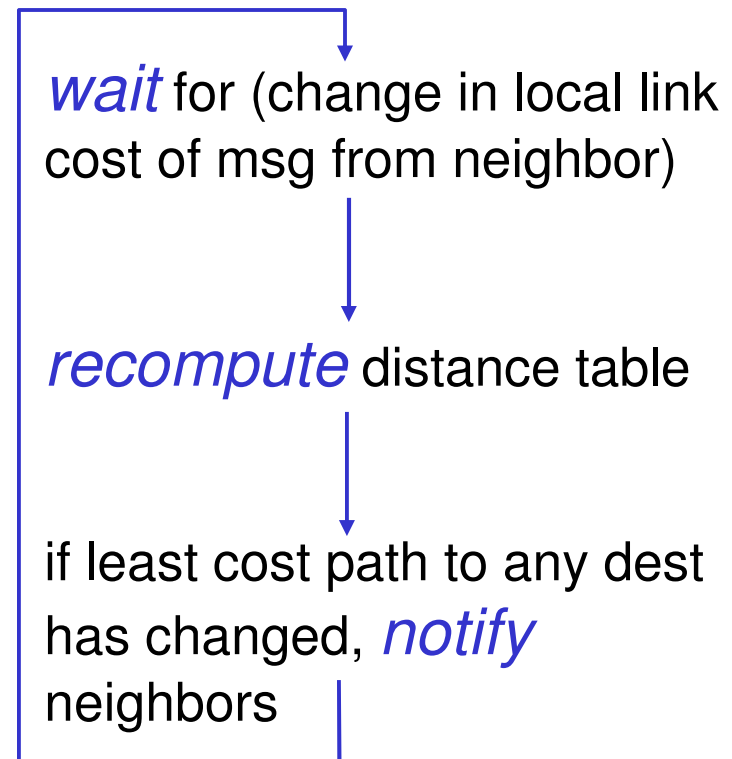
each local iteration caused by:

- r local link cost change
- r message from neighbor: its least cost path change from neighbor

Distributed:

- r each node notifies neighbors *only* when its least cost path to any destination changes
 - m neighbors then notify their neighbors if necessary

Each node:



Distance Vector Algorithm:

At all nodes, X:

- 1 Initialization:
- 2 for all adjacent nodes v:
- 3 $D^X(*,v) = \text{infinity}$ /* the * operator means "for all rows" */
- 4 $D^X(v,v) = c(X,v)$
- 5 for all destinations, y
- 6 send $\min_w D^X(y,w)$ to each neighbor /* w over all X's neighbors */




From the node to whatever destination going through v

Distance Vector Algorithm (cont.):

```
8 loop
9   wait (until I see a link cost change to neighbor V
10      or until I receive update from neighbor V)
11
12  if (c(X,V) changes by d)
13    /* change cost to all dest's via neighbor v by d */
14    /* note: d could be positive or negative */
15    for all destinations y:  $D^X(y,V) = D^X(y,V) + d$ 
16
17  else if (update received from V wrt destination Y)
18    /* shortest path from V to some Y has changed */
19    /* V has sent a new value for its  $\min_w DV(Y,w)$  */
20    /* call this received new value is "newval" */
21    for the single destination y:  $D^X(Y,V) = c(X,V) + \text{newval}$ 
22
23  if we have a new  $\min_w D^X(Y,w)$  for any destination Y
24    send new value of  $\min_w D^X(Y,w)$  to all neighbors
25
26 forever
```

Distance Vector Algorithm: example



		cost via	
		Y	Z
dest	D ^X		
	Y	2	∞
Z	∞	7	

		cost via	
		Y	Z
dest	D ^X		
	Y	2	8
Z	3	7	

		cost via	
		Y	Z
dest	D ^X		
	Y		
Z			

		cost via	
		X	Z
dest	D ^Y		
	X	2	∞
Z	∞	1	

		cost via	
		X	Z
dest	D ^Y		
	X	2	8
Z	9	1	

		cost via	
		X	Z
dest	D ^Y		
	X		
Z			


		cost via	
		X	Y
dest	D ^Z		
	X	7	∞
Y	∞	1	

		cost via	
		X	Y
dest	D ^Z		
	X	7	3
Y	9	1	

		cost via	
		X	Y
dest	D ^Z		
	X		
Y			

Cost updates from the neighbors are used for sake of recomputing
 The best routes and may lead to new cost updates...

Distance Vector Algorithm: example



		cost via	
		Y	Z
d e s t	D ^X		
	Y	2	∞
Z	∞	7	

		cost via	
		X	Z
d e s t	D ^Y		
	X	2	∞
Z	∞	1	

		cost via	
		X	Y
d e s t	D ^Z		
	X	7	∞
Y	∞	1	

		cost via	
		Y	Z
d e s t	D ^X		
	Y	2	8
Z	3	7	

Line 21 of the algorithm description

$$D^X(Y,Z) = c(X,Z) + \min_w \{D^Z(Y,w)\}$$

$$= 7 + 1 = 8$$

$$D^X(Z,Y) = c(X,Y) + \min_w \{D^Y(Z,w)\}$$

$$= 2 + 1 = 3$$