

COMP 790-088

Networked and Distributed Systems

Internet Routing

Jasleen Kaur

September 16, 2009

COMP 790-088

1

Forwarding vs. Routing

Local vs. Distributed

- ◆ Both datagram and virtual-circuit based networks need to know how to construct forwarding tables
- ◆ Forwarding:
 - » Select an output port based on destination address and routing table
 - » Simple and well-defined process performed locally at a given node
 - » Forwarding table used while forwarding packets
 - ❖ eg: (network number, interface id, MAC address)
- ◆ Routing:
 - » Building the forwarding/routing table
 - » A complex, distributed algorithm problem
 - » Routing table may be a precursor to a forwarding table
 - ❖ eg: (network number, nextHopIP, CostToDestination)

Why do we need two tables?

COMP 790-088

2


Intra-domain Routing

Formulation

- ◆ Intra-domain routing \Rightarrow ~ 100 routers
- ◆ Given:
 - » **Graph**: where nodes are routers and edges are links
 - » **Cost**: associated with each link

- ◆ Find:
 - » **Lowest-cost path** between any two nodes

- ◆ Requirements:
 - » Self-healing, traffic-sensitive, scalable



Need dynamic, distributed algorithms!
Two classes: based on "distance-vector" and "link-state"

COMP 790-088

3

Distance-vector Routing

Basic Idea

- ◆ Each node:
 - » Constructs a vector of distances to all other nodes
 - ❖ Distance vector
 - » Distributes to immediate neighbors
- ◆ Neighbors use the distributed information to update their distance vectors
- ◆ This distributed exchange-update-exchange should lead to globally consistent distance vectors (and routing tables)

COMP 790-088

4

Distance Vector Routing

Distance table data structure

- ◆ Each node has its own table with a...
 - » Row for each possible destination
 - » Column for each directly-attached adjacent node (neighbor)
- ◆ Each table entry gives cost to reach destination via that adjacent node
 - » Distance = Cost

| | | Cost to Destination via | | |
|-------------|---|-------------------------|----|---|
| $D^E()$ | | A | B | D |
| Destination | A | 1 | 14 | 5 |
| | B | 7 | 8 | 5 |
| | C | 6 | 9 | 4 |
| | D | 4 | 11 | 2 |

$$D^X(Y,Z) = \text{distance from } X \text{ to } Y \text{ via } Z \text{ as first hop}$$

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


$w = \{\text{neighbors of } Z\}$

COMP 790-088

5

Distance Vector Routing

Distance table example



| | | Cost to Destination via | | |
|-------------|---|-------------------------|----|---|
| $D^E()$ | | A | B | D |
| Destination | A | 1 | 14 | 5 |
| | B | 7 | 8 | 5 |
| | C | 6 | 9 | 4 |
| | D | 4 | 11 | 2 |

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

$$= 2 + 2 = 4$$

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

$$= 2 + 3 = 5 \quad \text{A loop?!}$$

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


$$= 8 + 6 = 14 \quad \text{Loop!}$$

COMP 790-088

6

Distance Vector Routing

Distance table example



Distance Table

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

Routing Table

| $D^E()$ | |
|---------|------|
| A | A, 1 |
| B | D, 5 |
| C | D, 4 |
| D | D, 2 |

Adjacent Node to Use, Cost

- ◆ The distance table gives the routing table

- » Just take the minimum cost per destination

COMP 790-088

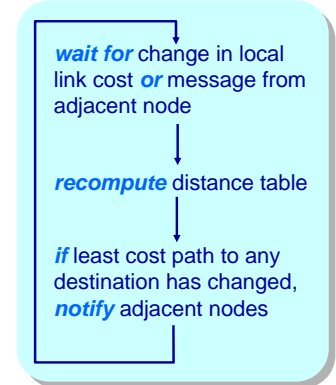
7

Distance Vector Routing

Algorithm

- ◆ Iterative, asynchronous:
 - each local iteration caused by:
 - » Local link cost change, or
 - » Message from adjacent node that its least cost path to some destination has changed
- ◆ Distributed:
 - » Each node notifies adjacent nodes *only* when its least cost path to some destination changes
 - » Adjacent nodes then notify their adjacent nodes if this update changes a least cost path

Each node:

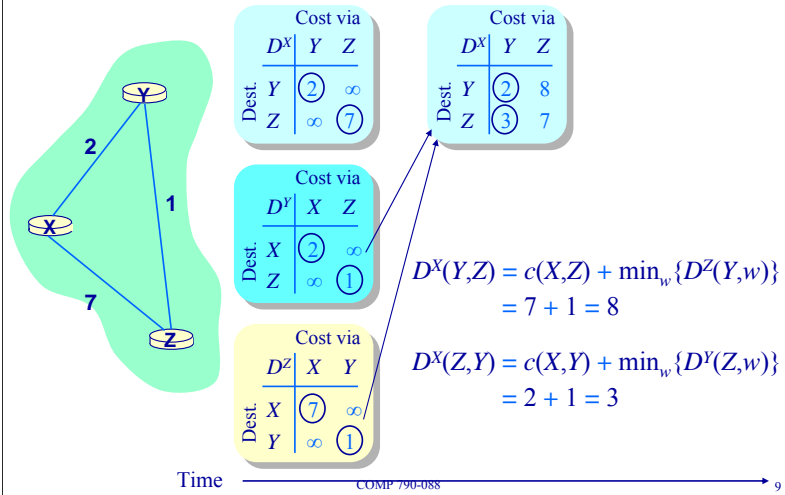


COMP 790-088

8

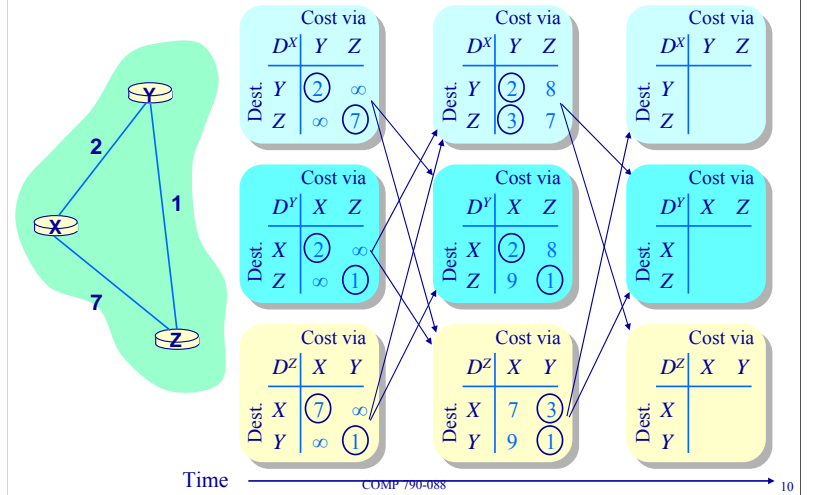
Distance Vector Algorithm

Example



Distance Vector Algorithm


Example



Distance Vector Algorithm


Link cost changes

- ◆ When a node detects a local link cost change:
 - » The nodes updates its distance table
 - » If the least cost path changes, the node notifies its neighbors



How long does convergence take?

“Good news travels fast”



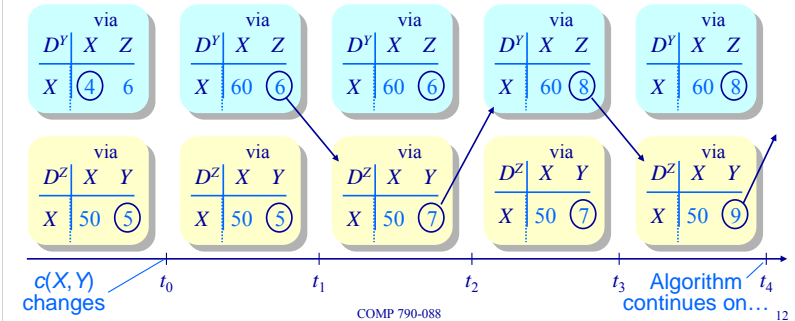

11

Distance Vector Algorithm

Link cost changes

- ◆ Good news travels fast, but...
- ◆ “Bad news” travels slow!
 - » The “count to infinity” problem

**Routing Loop!
Does it Terminate?**




12


The Count to Infinity Problem

The “poisoned reverse” technique

- ◆ If Z routes through Y to get to X:
 - » Then Z tells Y that Z’s distance to X is infinite



Initialization...

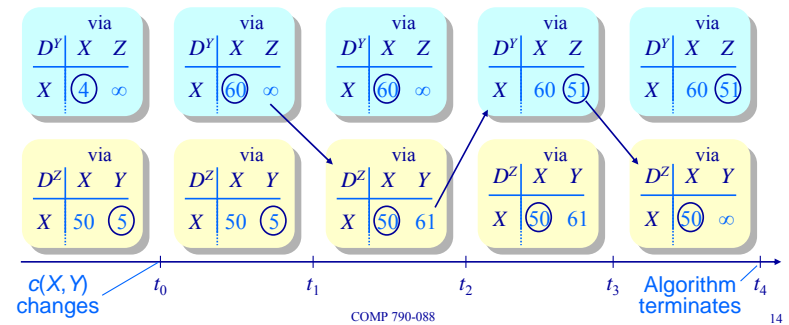



13

The Count to Infinity Problem

The “poisoned reverse” technique

- ◆ If Z routes through Y to get to X:
 - » Then Z tells Y that Z’s distance to X is infinite
- ◆ (Will this completely solve the problem?)
 - » How about delaying advertising alternate routes after link failures?



14