


More on Network Routing and Internet Protocol

Kai Shen

Dept. of Computer Science, University of Rochester

Distance Vector Routing



- **routing table** (at each host): the next hop for each destination in the network
- **distance vector routing**: the routing table is constructed from a distance vector at each node
- distance vectors can be maintained in a decentralized fashion


Distance Vector Routing: Constructing Distance Vector

At node X, each entry $D(X,Y,Z)$ includes:

- the distance from X to neighbor Y
- the shortest distance from Y to Z

Distance vector can be built on information about local links and neighbors' routing tables!

Distance Vector Algorithm: An Example



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

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


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

$$D(X,Z,Y) = c(X,Z) + \text{shortest}(Z,Y) = 7+1 = 8$$

$$D(X,Y,Z) = c(X,Y) + \text{shortest}(Y,Z) = 2+1 = 3$$

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

Distance Vector Algorithm: An Example



D ^X		cost via		D ^X		cost via		D ^X		cost via	
d	e	Y	Z	d	e	Y	Z	d	e	Y	Z
∞	∞	2	∞	2	8	2	8	2	8	2	8
∞	7	∞	1	∞	9	∞	1	∞	9	∞	1

D ^Y		cost via		D ^Y		cost via		D ^Y		cost via	
d	e	X	Z	d	e	X	Z	d	e	X	Z
∞	∞	2	∞	2	8	2	8	2	8	2	8
∞	1	∞	1	∞	9	∞	1	∞	9	∞	1

D ^Z		cost via		D ^Z		cost via		D ^Z		cost via	
d	e	X	Y	d	e	X	Y	d	e	X	Y
∞	∞	7	∞	7	3	7	3	7	3	7	3
∞	1	∞	1	∞	9	∞	1	∞	9	∞	1

9/26/2007

CSC 257/457 - Fall 2007

5

Distance Vector Routing: Maintaining Distance Vector

Information:

- local link cost
- distance vector → routing table

Recompute caused by:

- local link cost change
- message from neighbor: its routing table has changed

Notify neighbor:

- if my own routing table has changed, notify my neighbors

Each node:

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

recompute distance vector

if the local routing table has changed, notify neighbors

9/26/2007


CSC 257/457 - Fall 2007

6

Distance Vector: Link Cost Changes

Link cost changes:

- node detects local link cost change
- recompute - updates local distance vector
- if cost change in least cost path, notify neighbors



"good news travels fast"

D ^Y		via		D ^Y		via		D ^Y		via	
d	e	X	Z	d	e	X	Z	d	e	X	Z
x	∞	4	6	x	1	1	6	x	1	1	3
x	50	5	5	x	50	5	5	x	50	2	2

algorithm terminates

time t_0 t_1 t_2

9/26/2007


CSC 257/457 - Fall 2007

7

Distance Vector: Link Cost Changes

Link cost changes:

- good news travels fast
- bad news travels slow - "count to infinity" problem!



D ^Y		via		D ^Y		via		D ^Y		via	
d	e	X	Z	d	e	X	Z	d	e	X	Z
x	∞	4	6	x	60	6	6	x	60	8	8
x	50	5	5	x	50	5	5	x	50	7	7

algorithm continues on!

time t_0 t_1 t_2 t_3 t_4

9/26/2007



CSC 257/457 - Fall 2007

8

Distance Vector: Poisoned Reverse

If Z routes through Y to get to X :

- Z tells Y its (Z's) distance to X is infinity (so Y won't route to X via Z)
- Will this completely solve count to infinity problem?




9/26/2007

CSC 257/457 - Fall 2007

9

Routing Loops

- In packet switching networks, each node computes its own routing table independently
- What if?
 - Y's next hop to X is Z;
 - and Z's next hop to X is Y.



- Routing loops
 - during count-to-infinity
 - unsynchronized information updates in dynamic networks

9/26/2007

CSC 257/457 - Fall 2007

10

Comparison of LS and DV Algorithms

Speed of convergence

- LS:**
 - O(n log n) algorithm in local computation
 - may have oscillations
- DV:**
 - convergence time at startup?
 - count-to-infinity problem at link cost change

Robustness: what happens if router malfunctions?

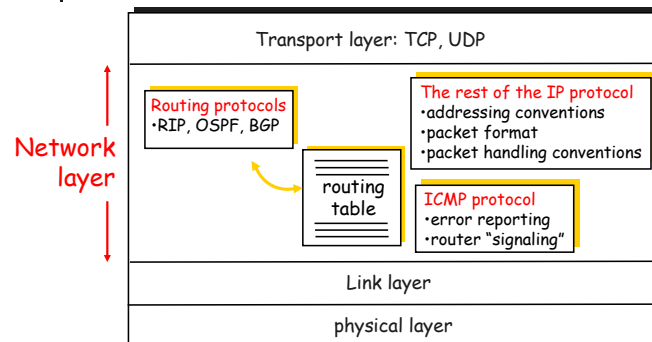
- LS:**
 - node can advertise incorrect link cost
- DV:**
 - node can advertise incorrect routing table

9/26/2007

CSC 257/457 - Fall 2007

11

The Internet Network Layer



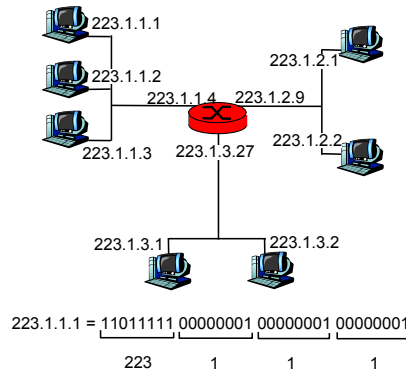
9/26/2007

CSC 257/457 - Fall 2007

12

IP Addressing: Introduction

- **IP address:**
 - 32-bit identifier for each host, router interface
- **interface:** connecting point into each data link
 - routers typically have multiple interfaces
 - hosts often have single interfaces




9/26/2007

CSC 257/457 - Fall 2007

13

IP Network and Hierarchical Addressing

- **What's an IP network?**
 - can physically reach each other without intervening router (intervening switches?)
- **IP address:**
 - network part (high order bits); host part (low order bits)
 - devices with same network part of IP address are in the same IP network



9/26/2007

CSC 257/457 - Fall 2007

14

IP Addresses: Original Standard

Given notion of "network", let's re-examine IP addresses:

"class"-based addressing:

class	network	host	range
A	0		10.0.0 to 127.255.255.255
B	10		128.0.0 to 191.255.255.255
C	110		192.0.0 to 223.255.255.255
D	1110		224.0.0 to 239.255.255.255

← 32 bits →


9/26/2007

CSC 257/457 - Fall 2007

15

IP Addressing: CIDR

- "Class"-based addressing:
 - inefficient use of address space, address space exhaustion
 - e.g., class B net allocated enough addresses for 65K hosts, even if only 2K hosts in that network
- **CIDR: Classless InterDomain Routing**
 - network portion of address of arbitrary length
 - address format: a.b.c.d/x, where x is # bits in network portion of address



9/26/2007

CSC 257/457 - Fall 2007

16

DHCP: Dynamic Host Configuration Protocol

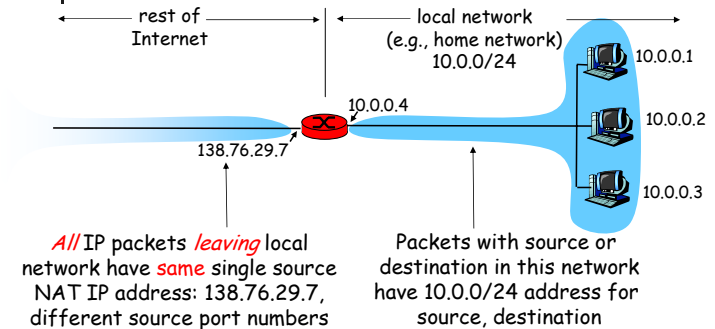
Goal: allow host to *dynamically* obtain its IP address from network server when it joins network

- Can renew its lease on address in use
- Allows reuse of addresses (only hold address while connected)

DHCP protocol messages:

- host broadcasts "DHCP discover" msg
- DHCP server responds with "DHCP offer" msg
- host requests IP address: "DHCP request" msg
- DHCP server sends address: "DHCP ack" msg

NAT: Network Address Translation



Disclaimer

- Parts of the lecture slides contain original work of James Kurose, Larry Peterson, and Keith Ross. The slides are intended for the sole purpose of instruction of computer networks at the University of Rochester. All copyrighted materials belong to their original owner(s).