

Chapter 4 Network Layer

A note on the use of these ppt slides:


We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you see the animations; and can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a lot of work on our part. In return for use, we only ask the following:

- ♦ If you use these slides (e.g., in a class) that you mention their source (after all, we'd like people to use our books!)
- ♦ If you post any slides on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

© All material copyright 1996-2012
J.F. Kurose and K.W. Ross, All Rights Reserved

The course notes are adapted for Bucknell's CSCI 363
Xiannong Meng
Spring 2016



**Computer
Networking: A Top
Down Approach**
6th edition
Jim Kurose, Keith Ross
Addison-Wesley
March 2012

Application Layer 2-1

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 4-2

Distance vector algorithm

Bellman-Ford equation (dynamic programming)

let

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


then

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

\downarrow cost from neighbor v to destination y
 \downarrow cost to neighbor v
 \downarrow min taken over all neighbors v of x

Network Layer 4-3

Bellman-Ford example



Assume we have computed,
 $d_v(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

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 achieving minimum (in our case, x) is next hop in shortest path, used in forwarding table

Network Layer 4-4

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 4-5

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 4-6

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:



Network Layer 4-7



Network Layer 4-8




Network Layer 4-9

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 x.

Network Layer 4-10

Distance vector: link cost changes

link cost changes:

- node detects local link cost change
- bad news travels slow** - "count to infinity" problem, e.g., $c(x,y)$ is changed from 4 to 60
- Initially, $D_y(x) = 4$, $D_z(x) = 5$, so next update leads to $D_y(x) = 6$, $D_z(x) = 7$, next, $D_y(x) = 8$, $D_z(x) = 9$, ...



the "count-to-infinity problem"

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 4-11

Other means to tackle count-to-infinity

- Split-horizon: if a node X learns the route to Y through Z, the X should not be part of the advertising to node Z to reach Y.
- Hold-down timer: a router starts a hold-down-timer when new routing information is available. It doesn't update its own routing information until the timer expires.

<https://www.techopedia.com/definition/14850/split-horizon>
<https://www.techopedia.com/definition/25073/hold-down-timer>

Network Layer 4-12

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 4-13

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 4-14

Hierarchical routing

our routing study thus far - idealization

- ❖ all routers identical
- ❖ network "flat"
- ... *not true in practice*

scale: with 600 million destinations:

- ❖ can't store all dest's in routing tables!
- ❖ routing table exchange would swamp links!

administrative autonomy

- ❖ internet = network of networks
- ❖ each network admin may want to control routing in its own network

Network Layer 4-15

Hierarchical routing

- ❖ aggregate routers into regions, "**autonomous systems**" (AS)
- ❖ routers in same AS run same routing protocol

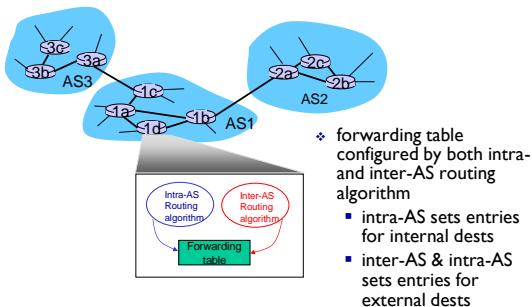
gateway router:

- ❖ at "edge" of its own AS
- ❖ has link to router in another AS

- "intra-AS" routing protocol
- routers in different AS can run different intra-AS routing protocol

Network Layer 4-16

Interconnected ASes



Network Layer 4-17


Inter-AS tasks

- ❖ suppose router in AS1 receives datagram destined outside of AS1:
 - router should forward packet to gateway router, but which one?

AS1 must:

1. learn which dests are reachable through AS2, which through AS3
2. propagate this reachability info to all routers in AS1


job of inter-AS routing!



Network Layer 4-18


Example: setting forwarding table in router 1d

- ❖ suppose AS1 learns (via inter-AS protocol) that subnet x is reachable via AS3 (gateway 1c), but not via AS2
 - inter-AS protocol propagates reachability info to all internal routers
- ❖ router 1d, which has three out-going links (1 for c, 2 for a, 3 for b), determines from intra-AS routing info that its interface 1 is on the least cost path to 1c
 - installs forwarding table entry $(x, 1)$



Example: choosing among multiple ASes

- ❖ now suppose AS1 learns from inter-AS protocol that subnet x is reachable from AS3 *and* from AS2.
- ❖ it is the job of inter-AS routing protocol to determine which gateway it should forward packets towards for dest x



Example: choosing among multiple ASes

- ❖ in the presence of multiple ASes to reach a destination, to configure forwarding table, router 1d must determine towards which gateway it should forward packets for dest x
- ❖ **hot potato routing**: send packet towards closest of two routers. (closest \rightarrow smallest cost)

