

9: Intro to Routing Algorithms

Last Modified:
3/24/2003 2:08:40 PM

4: Network Layer 4a-1

Routing

- IP Routing - each router is supposed to send each IP datagram one step closer to its destination
- How do they do that?
 - Static Routing
 - Hierarchical Routing - in ideal world would that be enough? Well its not an ideal world
 - Dynamic Routing
 - Routers communicate amongst themselves to determine good routes (ICMP redirect is a simple example of this)
 - Before we cover specific routing protocols we will cover principles of dynamic routing protocols

4: Network Layer 4a-2

Routing Algorithm classification: Static or Dynamic?

Choice 1: Static or dynamic?

Static:

- routes change slowly over time
- Configured by system administrator
- Appropriate in some circumstances, but obvious drawbacks (routes added/removed? sharing load?)
- Not much more to say?

Dynamic:

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

4: Network Layer 4a-3

Routing Algorithm classification: Global or decentralized?

Choice 2, if dynamic: 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 (gossip)
- "distance vector" algorithms

4: Network Layer 4a-4

Roadmap

- Details of Link State
- Details of Distance Vector
- Comparison

4: Network Layer 4a-5


Routing

Routing protocol

Goal: determine "good" path (sequence of routers) thru network from source to dest.

Graph abstraction for routing algorithms:

- graph nodes are routers
- graph edges are physical links
 - link cost: delay, \$ cost, or congestion level



□ "good" path:


- typically means minimum cost path
- other definitions possible

4: Network Layer 4a-6

Global Dynamic Routing

See the big picture; Find the best Route

What algorithm do you use?



4: Network Layer 4a-7

A Link-State Routing Algorithm

Dijkstra's algorithm

- Know complete network topology with link costs for each link is known to all nodes
 - accomplished via "link state broadcast"
 - In theory, all nodes have same info
- Based on info from all other nodes, each node individually computes least cost paths from one node ("source") to all other nodes
 - gives routing table for that node
- iterative: after k iterations, know least cost path to k dest.'s

4: Network Layer 4a-8

Link State Algorithm: Some Notation

Notation:

- $c(i,j)$: link cost from node i to j . cost infinite if not direct neighbors
- $D(v)$: current value of cost of path from source to dest. V
- $n(v)$: next hop from this source to v along the least cost path
- N : set of nodes whose least cost path definitively known

4: Network Layer 4a-9


Dijkstra's Algorithm

- 1 **Initialization** – know $c(l,j)$ to start:
- 2 $N = \{A\}$
- 3 for all nodes v
- 4 if v adjacent to A
- 5 then $D(v) = c(A,v)$
- 6 else $D(v) = \text{infty}$
- 7
- 8 **Loop**
- 9 find w not in N such that $D(w)$ is a minimum (optional?)
- 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**

4: Network Layer 4a-10

Dijkstra's algorithm: example

Step	start N	D(B),n(B)	D(C),n(C)	D(D),n(D)	D(E),n(E)	D(F),n(F)
→ 0	A	2,B	5,C	1,A	infinity	infinity
→ 1	AD	2,B	4,D		2,D	infinity
→ 2	ADE	2,B	3,D			4,D
→ 3	ADEB		3,D			4,D
→ 4	ADEBC					4,D
→ 5	ADEBCF					



4: Network Layer 4a-11

Dijkstra's Algorithm gives routing table

	Outgoing Link
A	
A	$n(A) = A$
B	$n(B) = B$
C	$n(C) = D$
D	$n(D) = D$
E	$n(E) = D$
F	$n(F) = D$

4: Network Layer 4a-12


Complexity of Link State

Algorithm complexity: n nodes

- each iteration
 - Find next w not in N such that $D(w)$ is a minimum
 - Then for that w , check its best path to other destinations
 - $\Rightarrow n*(n+1)/2$ comparisons: $O(n^2)$
- more efficient implementations possible using a heap: $O(n \log n)$

Oscillations

- Assume:
 - Link cost = amount of carried traffic
 - Link cost is not symmetric
 - B and D sending 1 unit of traffic; C send e units of traffic
- Initially start with slightly unbalanced routes
- Everyone goes with least loaded, making them most loaded for next time, so everyone switches
- Herding effect!



Initially start with almost equal routes ... B and C go clockwise to A ... B, C and D go counterclockwise ... B, C, D go clockwise

Preventing Oscillations

- Avoid link costs based on experienced load
 - But want to be able to route around heavily loaded links...
- Avoid "herding" effect
 - Avoid all routers recomputing at the same time
 - Not enough to start them computing at a different time because will synchronize over time as send updates
 - Deliberately introduce randomization into time between when receive an update and when compute a new route

Distance Vector Routing Algorithm

distributed:

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

iterative:

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

asynchronous:

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

Distance Vector Routing Algorithm

Distance Table data structure

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

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


Column only for each neighbor

$$D^X() \quad \begin{array}{c} \text{cost to destination via} \\ \underline{\hspace{1cm}} \\ Z \end{array}$$

$$\begin{array}{c} \text{destination} \\ Y \end{array} \quad D^X(Y,Z)$$

Rows for each possible dest !

Example: Distance Table for E



Column only for each neighbor

$$D^E() \quad \begin{array}{c} \text{cost to destination via} \\ \underline{\hspace{1cm}} \\ \text{neighbor} \end{array}$$

	A	B	D
A	1	14	5
B	7	8	5
C	6	9	4
D	4	11	2

destination

D^E (row, col)

$$D^E(C,D) = c(E,D) + \min_w \{D^D(C,w)\} = 2+2 = 4$$

$$D^E(A,D) = c(E,D) + \min_w \{D^D(A,w)\} = 2+3 = 5 \quad \text{Loop back through E!}$$

$$D^E(A,B) = c(E,B) + \min_w \{D^B(A,w)\} = 8+6 = 14 \quad \text{Loop back through E!}$$

Rows for each possible dest !

Distance table gives routing table

cost to destination via ○ = least cost

D ^E ()		cost to destination via			Outgoing link to use, cost	
		A	B	D		
destination	A	1	14	5	A	A,1
	B	7	8	5	B	D,5
	C	6	9	4	C	D,4
	D	4	11	2	D	D,4

Distance table → Routing table

Distance Vector Routing: overview


Iterative, asynchronous:
each local iteration caused by:

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

Distributed:

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

Each node:



Distance Vector Algorithm:

At all nodes, X:

- 1 Initialization (don't start knowing link costs for all links in graph):
- 2 for all adjacent nodes v:
- 3 $D^X(*,v) = \text{infty}$ /* 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 */

Then in steady state...


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

To start just know directly connected links...tell neighbors



Distance Vector Algorithm: example


In steady state, when have good news tell neighbor




Distance Vector: link cost changes

Link cost changes:

- node detects local link cost change
- updates distance table (line 15)
- if cost change in least cost path, notify neighbors (lines 23,24)



"good news travels fast"



algorithm terminates



Anyone see a problem?

4: Network Layer 4a-25

Distance Vector: link cost changes

Link cost changes:

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





algorithm continues on!

4: Network Layer 4a-26

Distance Vector: poisoned reverse

If Z routes through Y to get to X :

- Originally, Z tells Y its (Z's) distance to X is infinite (so Y won't route to X via Z)
- In end, Y tells Z infinity
- will this completely solve count to infinity problem?




algorithm terminates

4: Network Layer 4a-27

Bigger Loops and Poison Reverse


$$D^E(A,D) = c(E,D) + \min_w \{D^D(A,w)\} \\ = 2 + 3 = 5$$



Loop back through E! Poison reverse will fix this D tells E infinity because D's route to A through E

$$D^E(A,B) = c(E,B) + \min_w \{D^B(A,w)\} \\ = 8 + 6 = 14$$



Loop back through E! Poison reverse will not fix this B's route to A is through E but B doesn't know that so does not tell E infinity B's route is through C so no poison reverse E will try to send through B



4: Network Layer 4a-28

Count to Infinity Example with Bigger Loop

- B will learn bad news
- C will have told B infinity because its route to A is through B, so B won't reroute through C
- E however will have told B about a good route to A through D (cost 6)
- B will choose that route instead and advertise it as the new best to C (cost 6+8 = 14); it will be worse than the old one it advertised to C (old cost = 1)
- C will propagate this updated "best" route to D (cost 15)
- D will propagate this new "best" route to E (cost 17)
- E will update the "best" route to B (cost 19)
- Last time it advertised cost 6 to B
- It will loop around adding 13 each time (cost of loop)
- Will continue until cost E advertises to B is bigger than 500



4: Network Layer 4a-29

Comparison of LS and DV algorithms

Message complexity

- LS: nodes send info on directly connections to all other nodes
 - More, smaller messages
- DV: nodes send info on best paths to all destinations to neighbors
 - Fewer, larger messages

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

Speed of Convergence

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

4: Network Layer 4a-30