


Computer Networks

Routing Algorithms

Based on Computer Networking, 4th Edition by Kurose and Ross


Stan Kurkovsky

Interplay between routing, forwarding



Stan Kurkovsky


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

Stan Kurkovsky

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

Stan Kurkovsky

Routing Algorithm classification

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
- "distance vector" algorithms

Static or dynamic?

Static:

- routes change slowly over time

Dynamic:

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

Stan Kurkovsky

A Link-State Routing Algorithm

Dijkstra's algorithm

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

Notation:

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

Stan Kurkovsky


Dijkstra'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'**

Stan Kurkovsky

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					



Stan Kurkovsky

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)

Stan Kurkovsky


Dijkstra's algorithm, discussion

Algorithm complexity: n nodes

- each iteration: need to check all nodes, w, not in N
- $n(n+1)/2$ comparisons: $O(n^2)$
- more efficient implementations possible: $O(n \log n)$

Oscillations possible:

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



Stan Kurkovsky

Distance Vector Algorithm

Bellman-Ford Equation (dynamic programming)

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


- Then

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

where min is taken over all neighbors v of x

Stan Kurkovsky

Bellman-Ford example



Clearly, $d_v(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

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 that achieves minimum is next
hop in shortest path → forwarding table

Stan Kurkovsky

Distance Vector Algorithm

- $D_x(y)$ = estimate of least cost from x to y
- Distance vector: $\mathbf{D}_x = [D_x(y): y \in N]$
- Node x knows cost to each neighbor v: $c(x,v)$
- Node x maintains $\mathbf{D}_x = [D_x(y): y \in N]$
- Node x also maintains its neighbors' distance vectors
 - For each neighbor v, x maintains $\mathbf{D}_v = [D_v(y): y \in N]$

Basic idea:

- Each node periodically sends its own distance vector estimate to neighbors
- When a node x receives new DV estimate from neighbor, it updates its own DV using B-F equation:
 $Dx(y) \leftarrow \min_v \{c(x,v) + Dv(y)\}$ for each node $y \in N$
- Under minor, natural conditions, the estimate $Dx(y)$ converge to the actual least cost $dx(y)$

Stan Kurkovsky

Distance Vector Algorithm

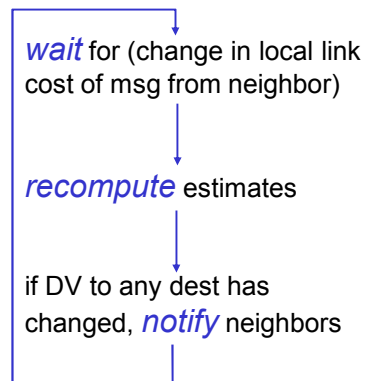
Iterative, asynchronous: each local iteration caused by:

- 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:




Stan Kurkovsky

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”


- At time t_0 , y detects the link-cost change, updates its DV, and informs its neighbors.
- At time t_1 , z receives the update from y and updates its table. It computes a new least cost to x and sends its neighbors its DV.
- At time t_2 , y receives z 's update and updates its distance table. y 's least costs do not change and hence y does *not* send any message to z .

Stan Kurkovsky

Distance Vector: link cost changes

Link cost changes:

- good news travels fast
- bad news travels slow - “count to infinity” problem!
- 44 iterations before algorithm stabilizes

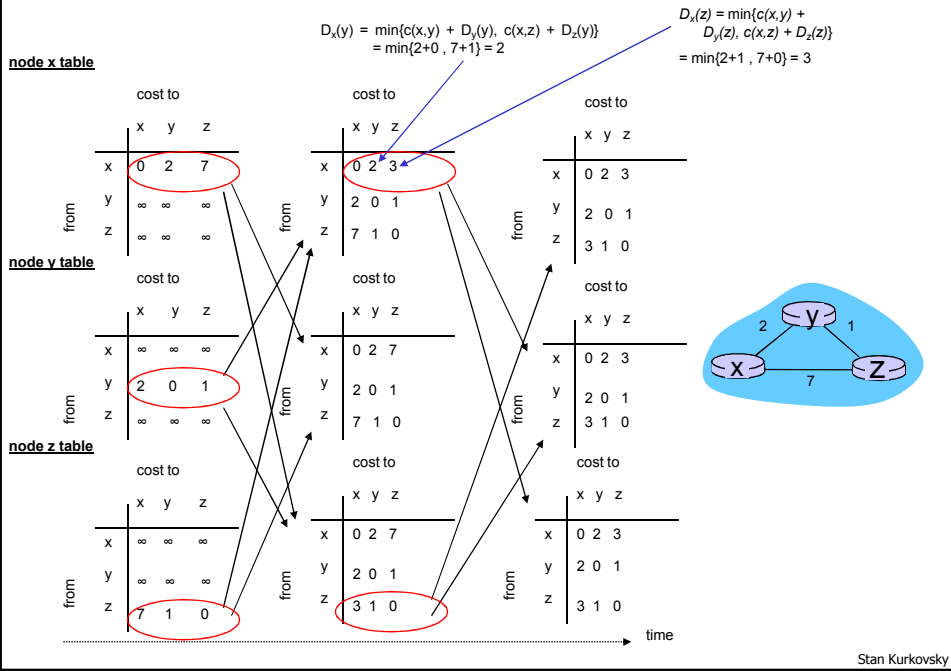


Poisoned reverse:

- 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?

Stan Kurkovsky

Distance Vector: example



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(n²) 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

Stan Kurkovsky

Hierarchical Routing

Our routing study thus far - idealization

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

scale: with 200 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

- aggregate routers into regions, "**autonomous systems**" (AS)

- routers in same AS run same routing protocol

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

Gateway router


- Direct link to router in another AS

Stan Kurkovsky

Interconnected ASes

- Forwarding table is configured by both intra- and inter-AS routing algorithm

- Intra-AS sets entries for internal dests
- Inter-AS & Intra-As sets entries for external dests



Stan Kurkovsky


Inter-AS tasks

- Suppose router in AS1 receives datagram for which dest is outside of AS1
 - Router should forward packet towards one of the gateway routers, but which one?

AS1 needs:

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

Job of inter-AS routing!



Stan Kurkovsky