CS168 Fall 2014 Discussion Section 2: Routing

Inspired by EE122 Fall 2013 Discussion Section 2

Problem 1: Link-State Routing

The following is a network of routers using Link-State routing to communicate with each other. The numbers adjacent to each link represent the cost to traverse the link.

Rows represent the iteration in each table, and columns represent destinations.

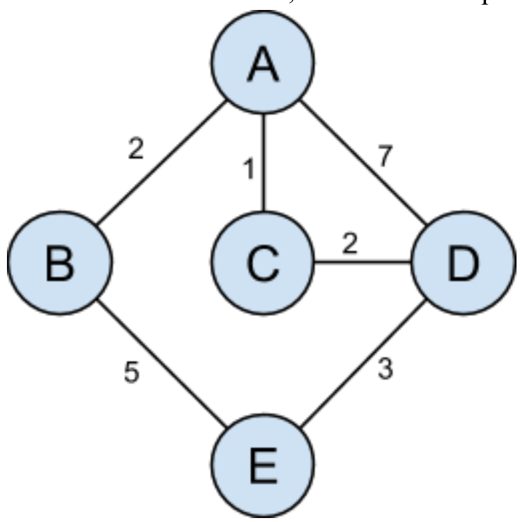

(a) After all routers have the global view of the network topology, run Dijkstra's algorithm on each node and fill up the following tables. Rows represent the iteration in each table, and columns represent destinations. Use the notation (cost, previous node) for each cell and specify *S* (set of nodes whose least cost path definitively known).

Node A's table has been filled for you. Note that these are NOT the node's routing tables; only the last row in each table matters in the end. Highlighted cell is chosen to be added to S.

i \ *Dest*	*B*	*C*	*D*	*E*	*S*
1	(2, A)	(1, A)	(7, A)	∞	AC
2	(2, A)	(1, A)	(3, C)	∞	ACB
3	(2, A)	(1, A)	(3, C)	(7, B)	ACBD
4	(2, A)	(1, A)	(3, C)	(6, D)	ACBDE

Node A

i \ *Dest*	*A*	*C*	*D*	*E*	*S*
1	(2, B)	∞	∞	(5, B)	BA
2	(2, B)	(3, A)	(9, A)	(5, B)	BAC
3	(2, B)	(3, A)	(5, C)	(5, B)	BACD
4	(2, B)	(3, A)	(5, C)	(5, B)	BACDE

Node B

i \ Dest	A	B	D	E	S
1	(1, C)	∞	(2, C)	∞	CA
2	(1, C)	(3, A)	(2, C)	∞	CAD
3	(1, C)	(3, A)	(2, C)	(5, D)	CADB
4	(1, C)	(3, A)	(2, C)	(5, D)	CADBE

Node C

i \ Dest	A	B	C	E	S
1	(7, D)	∞	(2, D)	(3, D)	DC
2	(3, C)	∞	(2, D)	(3, D)	DCA
3	(3, C)	(5, A)	(2, D)	(3, D)	DCAE
4	(3, C)	(5, A)	(2, D)	(3, D)	DCAEB

Node D

i \ Dest	A	B	C	D	S
1	∞	(5, E)	∞	(3, E)	ED
2	(10, D)	(5, E)	(5, D)	(3, E)	EDB
3	(7, B)	(5, E)	(5, D)	(3, E)	EDBC
4	(6, C)	(5, E)	(5, D)	(3, E)	EDBCA

Node E

(b) Now node B wants to send to D. What path does the packet traverse given the routing tables? What is the cost associated with this path?

Go through B's routing table (last row), we get B->A->C->D

The cost is 5

(c) The cost of link CD suddenly shoots up to 20, where may transient forwarding loops happen?

C and D change their routing tables immediately, but before their updates propagate, there is transient loop between A and C for destination D.

So if A sends to D, then the packet goes ACACACACAC… until the routing table converges.

Problem 2: Distance-Vector Routing

For the same network topology in Problem 1, consider the nodes communicating with each other using Distance-Vector routing.

Below are the initial routing tables before any routing updates are received. Rows represent the neighbors and columns represent the destination. An adjacent table indicates link costs. Highlighted cell indicates shortest path. For each node, the row corresponding to the same node also indicates the next hop for the shortest path that is chosen.

Nbr	Cost
A	0
B	2
C	1
D	7

From \ To	A	B	C	D
A	0, A	2, A	1, A	7, A
B	-	0	-	-
C	-	-	0	-
D	-	-	-	0

Node A

Nbr	Cost
A	2
B	0
E	5

From \ To	A	B	E
A	0	-	-
B	2, B	0, B	5, B
E	-	-	0

Node B

Nbr	Cost
A	1
C	0
D	2

From \ To	A	C	D
A	0	-	-
C	1, C	0, C	2, C
D	-	-	0

Node C

Nbr	Cost
A	7
C	2
D	0
E	3

To / From	A	C	D	E
A	0	-	-	-
C	-	0	-	-
D	7, D	2, D	0, D	3, D
E	-	-	-	0

Nbr	Cost
B	5
D	3
E	0

To / From	B	D	E
B	0	-	-
D	-	0	-
E	5, E	3, E	0, E

Node D **Node E**

Answer following questions, which indicate events that happen consecutively. Assume no packet exchanges other than the ones specified.

(a) C sends its update to A and D.

(a.i) What information is contained in C's update?

C sends a distance vector: a vector containing the shortest distance from C to each other node. At this point in time, C's distance vector is (A: 1, D: 2).

(a.ii) What do the routing tables for A and D look like after receiving C's update (You may not need to fill in all columns)?

Updated entries are shown in red font.

Nbr	Cost
A	0
B	2
C	1
D	7

To / From	A	B	C	D
A	0, A	2, A	1, A	3, C
B	-	0	-	-
C	1	-	0	2
D	-	-	-	0

Nbr	Cost
A	7
C	2
D	0
E	3

To / From	A	C	D	E
A	0	-	-	-
C	1	0	2	-
D	3, C	2, D	0, D	3, D
E	-	-	-	0

Node A **Node D**

(a.iii) Which nodes among A and D are expected to send routing updates after receiving C's update?

A and D will both send updates, because they all had the shortest path to one or more nodes change.

(b) A sends its update to B, C, and D.

(b.i) What information is contained in A's update?

As in part (i) of (a), A sends a distance vector: (B: 2, C: 1, D: 3)

(b.ii) What do the routing tables for B, C, and D look like after receiving A's update (You may not need to fill in all columns)?

Updated entries are shown in red font.

Nbr	Cost
A	2
B	0
E	5

To / From	A	B	C	D	E
A	0	2	1	3	-
B	2, B	0, B	3, A	5, A	5, B
E	-	-	-	-	0

Node B

Nbr	Cost
A	1
C	0
D	2

To / From	A	B	C	D
A	0	2	1	3
C	1, C	3, A	0, C	2, C
D	-	-	-	0

Node C

Nbr	Cost
A	7
C	2
D	0
E	3

To / From	A	B	C	D	E
A	0	2	1	3	-
C	1	-	0	2	-
D	3, C	9, A	2, D	0, D	3, D
E	-		-	-	0

Node D

(b.iii) At this point, what route does D use to reach B? It knows that it can route to A via C with total distance 3 and that A can reach B with distance 2. Should it use this information to optimize the route to B or should it wait for an update for C?

D->A->B with distance 9. It should not use this information now, as it has no way of knowing whether C is aware of the route to B through A and that it will be able to forward a packet with destination B to A. D needs to wait until it explicitly receives a distance vector from C confirming this, before it can optimize its route to B.

(b.iv) Which nodes among B, C, and D are expected to send routing updates after receiving A's update?

B, C, and D will all send updates, because they all had the shortest path to one or more nodes change.

(c) D sends its update to A, C, and E.

(c.i) What information is contained in D's update?

As in part (i) of (a), D sends a distance vector: (A: 3, B: 9, C: 2, E: 3)

(c.ii) What do the routing tables for A, C, and E look like after receiving D's update (You may not need to fill in all columns)?

Updated entries are shown in bold-face, red font.

Nbr	Cost
A	0
B	2
C	1
D	7

To / From	A	B	C	D	E
A	0, A	2, A	1, A	3, C	10, D
B	-	0	-	-	-
C	1	-	0	2	-
D	3	9	2	0	3

Node A

Nbr	Cost
A	1
C	0
D	2

To / From	A	B	C	D	E
A	0	2	1	3	-
C	1, C	3, A	0, C	2, C	5, D
D	3	9	2	0	3

Node C

Nbr	Cost
B	5
D	3
E	0

To / From	A	B	C	D	E
B	-	0	-	-	-
D	3	9	2	0	3
E	6, D	5, E	5, D	3, E	0, E

Node E

(c.iii) Which nodes among A, C, and E are expected to send routing updates after receiving D's update?

A, C, and E will all send updates, because they all had the shortest path to one or more nodes change.

(d) Have the routing tables converged? Why or why not?

No, because there are still distance vector updates, which are not heartbeat messages, propagating on the network.

Problem 3: Count-To-Infinity Problem

Consider a simple topology:

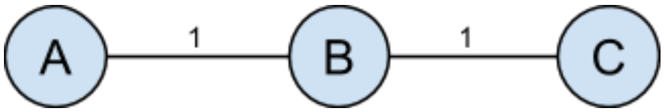

(a) What values will the routing tables have when the system has stabilized (after many rounds)?

Node A

Nbr	*Cost*
A	0
B	1

To / From	*A*	*B*	*C*
A	0, A	1, A	2, B
B	1	0	1

Node B

Nbr	*Cost*
A	1
B	0
C	1

To / From	*A*	*B*	*C*
A	0	1	2
B	1, B	0, B	1, B
C	2	1	0

Node C

Nbr	*Cost*
B	1
C	0

To / From	*A*	*B*	*C*
B	1	0	1
C	2, B	1, C	0, C

(b) Now suppose the link from A to B goes down, such that A is no longer reachable:

(b.i) B notices the link outage and updates its routing table. What does B's updated routing table look like?

Updated entries are shown in red font.

Nbr	*Cost*
A	∞
B	0
C	1

To / From	*A*	*B*	*C*
B	3, C	0, B	1, B
C	2	1	0

Node B

(b.ii) According to its routing table, what is the cost of B's minimum-cost path to A?

through C

(c) B sends an update to C. What is C's routing table after receiving the update?

Updated entries are shown in red font.

Nbr	*Cost*
B	1
C	0

To / From	*A*	*B*	*C*
B	3	0	1
C	4, B	1, C	0, C

Node C

(d) After updating its table, C sends an update to B. What is B's routing table after receiving the update?
Updated entries are shown in red font.

Nbr	*Cost*
A	∞
B	0
C	1

From \ *To*	*A*	*B*	*C*
B	5, C	0, B	1, B
C	4	1	0

Node B

(e) How many updates are exchanged before the tables converge?
Infinitely many

Problem 4: Poison Reverse

One solution to the count-to-infinity problem is "poison-reverse": if you are currently routing through a neighbor, tell that neighbor that your path to the destination has infinite cost.

(a) Continue on the network topology in Problem 3, before the link from A to B goes down, what is B's routing table (assuming that poison reverse was used when exchanging route information)?

Nbr	*Cost*
A	1
B	0
C	1

From \ *To*	*A*	*B*	*C*
A	0	1	∞
B	1, B	0, B	1, B
C	∞	1	0

Node B

(b) B detects the link outage and sends an update to C.
(b.i) What information is contained in B's update?
(A: ∞, C: 1)

(b.ii) What does C's routing table look like after receiving the update?

Updated entries are shown in bold-face, red font.

Nbr	Cost
B	1
C	0

To / From	A	B	C
B	∞	0	1
C	∞	1, C	0, C

Node C

(c) Now consider a more complex topology, with stabilized routing tables for A, B, and C:

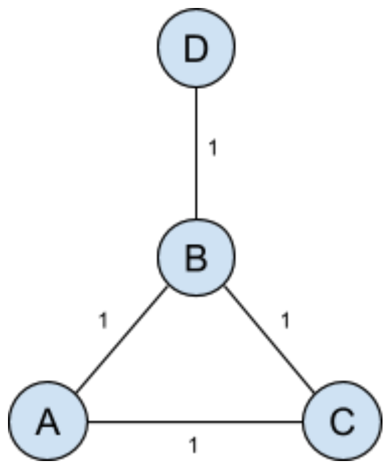

Nbr	Cost
A	0
B	1
C	1

To / From	A	B	C	D
A	0, A	1, A	1, A	2, B
B	1	0	1	1
C	1	1	0	2

Node A

Nbr	Cost
A	1
B	0
C	1
D	1

To / From	A	B	C	D
A	0	1	1	∞
B	1, B	0	1, B	1, B
C	1	1	0	∞
D	∞	1	∞	0

Node B

Nbr	Cost
A	1
B	1
C	0

To / From	A	B	C	D
A	0	1	1	2
B	1	0	1	1
C	1, C	1, C	0, C	2, B

Node C

Suppose the link between B and D goes down. B notices this change and sends an update to A.

(c.i) What information is contained in B's update? What is A's routing table after processing B's update?

B will send the following update: (A: 1, C: 1, D: ∞).

Nbr	Cost
A	0
B	1
C	1

To / From	A	B	C	D
A	0, A	1, A	1, A	3, C
B	1	0	1	∞
C	1	1	0	2

Node A

(c.ii) A then sends an update back to B. What information is contained in A's update? What is B's routing table after processing A's update? And whether there will be count-to-infinity problem?

A will send (B: 1, C: 1, D: 3).

Nbr	***Cost***
A	1
B	0
C	1

To / From	***A***	***B***	***C***	***D***
A	0	1	1	3
B	1, B	0	1, B	4, A
C	1	1	0	∞

Node B

Yes, there will be count-to-infinity problem, because B will send update to C, then C will change it's distance to D to 5, then C will send update to A, then A will change it's distance to D to 6 and so on.

(c.iii) How might you avoid the count-to-infinity problem here altogether?

Use link-state instead [actual reason for the invention of link-state!], use path-vector, avoid topologies that are prone to this problem, route on DAGs, use source routing, buy every company that owns networking equipment and replace their equipment with a single giant super router.