


# Computer Networks Homework #5

David Koll

November 29th 2012

# Dijkstra's algorithm

- Given the following network, use Dijkstra's algorithm to find the least cost paths from node u. Please provide a table showing the steps of the algorithm, a graph showing the resulting shortest-path tree from u and the final forwarding table of u.




# Dijkstra's algorithm (cont'd)


Step	N'	D(v), p(v)	D(w), p(w)	D(x), p(x)	D(y), p(y)	D(z), p(z)

# Dijkstra's algorithm (cont'd)

Step	N'	D(v), p(v)	D(w), p(w)	D(x), p(x)	D(y), p(y)	D(z), p(z)
0	u	3,u	$\infty$	$\infty$	3,u	2,u
1	uz	3,u	6,z	$\infty$	3,u	
2	uzy	3,u	5,y	8,y		
3	uzyv		5,y	8,y		
4	uzyvw			6,w		
5	uzyvwx					




# Dijkstra's algorithm (cont'd)



Dest.	Link.
z	z
y	y
v	v
w	y
x	y

# Distance Vector algorithm

- Given the following network, use the Distance Vector algorithm to find the least cost paths for all nodes. Fill the provided tables and indicate with arrows between the tables when a node sends a distance vector to another node.



# Distance Vector algorithm

		Node				
		w	x	y	z	
from	w					
	x					
	y					
	z					

		Node				
		x	w	x	y	z
from	w					
	x					
	y					
	z					

		Node				
		y	w	x	y	z
from	w					
	x					
	y					
	z					

		Node				
		z	w	x	y	z
from	w					
	x					
	y					
	z					

# Distance Vector algorithm

Node w		cost to			
		w	x	y	z
from	w	0	6	$\infty$	25
	x	$\infty$	$\infty$	$\infty$	$\infty$
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	$\infty$	$\infty$	$\infty$	$\infty$

Node w		cost to			
		w	x	y	z
from	w	0	6	8	13
	x	6	0	2	7
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	25	7	$\infty$	0

Node w		cost to			
		w	x	y	z
from	w	0	6	8	13
	x	6	0	2	7
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	13	7	9	0

Node w		cost to			
		w	x	y	z
from	w				3
	x				
	y				
	z				

Node x		cost to			
		w	x	y	z
from	w	$\infty$	$\infty$	$\infty$	$\infty$
	x	6	0	2	7
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	$\infty$	$\infty$	$\infty$	$\infty$

Node x		cost to			
		w	x	y	z
from	w	0	6	8	13
	x	6	0	2	7
	y	8	2	0	9
	z	13	7	9	0

Node x		cost to			
		w	x	y	z
from	w				
	x				
	y				
	z				

Node y		cost to			
		w	x	y	z
from	w	$\infty$	$\infty$	$\infty$	$\infty$
	x	$\infty$	$\infty$	$\infty$	$\infty$
	y	$\infty$	2	0	$\infty$
	z	$\infty$	$\infty$	$\infty$	$\infty$

Node y		cost to			
		w	x	y	z
from	w	$\infty$	$\infty$	$\infty$	$\infty$
	x	6	0	2	7
	y	8	2	0	9
	z	$\infty$	$\infty$	$\infty$	$\infty$


Node y		cost to			
		w	x	y	z
from	w	$\infty$	$\infty$	$\infty$	$\infty$
	x				
	y				
	z				

Node z		cost to			
		w	x	y	z
from	w	$\infty$	$\infty$	$\infty$	$\infty$
	x	$\infty$	$\infty$	$\infty$	$\infty$
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	25	7	$\infty$	0

Node z		cost to			
		w	x	y	z
from	w	0	6	$\infty$	25
	x	6	0	2	7
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	13	7	9	0

Node z		cost to			
		w	x	y	z
from	w	0	6	8	13
	x	6	0	2	7
	y	$\infty$	$\infty$	$\infty$	$\infty$
	z	13	7	9	0

Node z		cost to			
		w	x	y	z
from	w				
	x				
	y				
	z				






# Comparison LV vs. DV


- Scalability
  - LS uses broadcasts to disseminate complete knowledge about all links to entire network
  - DV only sends (local) information to neighboring nodes. Convergence time and DV size still increase with network size
- Robustness
  - LS: every router does its own calculations
  - DV: wrong DV will be used by neighboring nodes and further propagate the error

# Count-to-infinity problem

- Q: Explain the count-to-infinity problem using a simple example. How can this problem be avoided?
- Consider the following example:




# Count-to-infinity problem (con't)



- b sends DV (a,4) to c
- c computes that it can reach a in 5 hops via b and sends DV (a,5) to b
- Now the cost of the a-b changes to 30
- b recomputes its DV to a
  - Using the (old) DV (a,5) from c it computes the DV (a,6) and sends it to c
- c recomputes its DV to a
  - Using the DV from b it computes the DV (a,7) and sends it to b

# Count-to-infinity problem (con't)



- The last two steps repeat with an increasing DV to reach a until b sends the DV (a,20)
- c recomputes its DV to a
  - Using the DV from b it determines that DV (a,20) using the link c-a is less costly and sends it back to b
- b recomputes its DV to a
  - Using the DV (a,20) from c it computes the DV (a,21) and sends it to c
- Now the system is finally stable again

# Count-to-infinity problem (con't)

- Count-to-infinity problem can be avoided using the poisoned reverse technique.
  - Router A will advertise a distance as infinite to Router B if Router B is on the advertised path
  - In the example: In its advertisements to Router B, Router C will advertise the cost to reach Router A as infinite as long as it routes packets to A via B
  - Poisoned reverse will only prevent routing loops that involve just two gateways. It is still possible to end up with patterns in which three gateways are engaged in mutual deception. E.g. A may believe it has a route through B, B through C, and C through A.

# Routing policies

- Q: How are routing policies used in BGP. Give one example.
- Routing policies determine ...
  - ... which BGP advertisements to regard
  - ... which routes to advertise
- Example
  - AS x is connected to AS y and AS z
  - Policy : AS x does not want AS y to route traffic via AS x to AS z
  - Therefore, AS x does not advertise any route to reach AS z to AS y

# Intra- vs. inter-AS routing

- Q: Why are different inter-AS and intra-AS protocols used in the Internet?
- Different policies
  - Inter-AS: control over how (foreign) traffic is routed via the own network
  - Intra-AS: control over how traffic is routed within the the own network
- Scale
  - Hierarchical routing saves table size, reduced update traffic

# Intra- vs. inter-AS routing

- Performance
  - Intra-AS: can focus on performance
  - Inter-AS: policy may dominate over performance



# Thank you

Any questions?