ETH zürich

Networked Systems
ETH Zürich — seit 2015

Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 3 – Routing Concepts, Ethernet & Switching

Routing Concepts

3.1 Distance Vector

The figure on the left shows a weighted graph representing a network topology with 7 nodes. The nodes in the network use a distance vector algorithm to compute the shortest-paths in a distributed way. It takes one time step for a distance vector message to be sent from one node to another on a link. A node can send the distance vector message on multiple links at the same time.

In case paths have the same weight, the node picks the path traversing the smaller number of links. In case there is still a tie, the node picks the path of the neighbor with the lower identifier (alphabetical order).

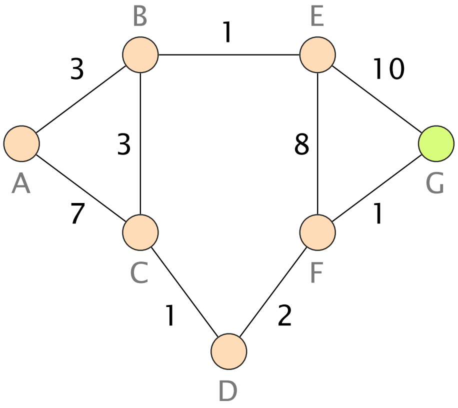

Weighted graph representing a network topology.

#	A	B	C	D	E	F	G
0	∅	∅	∅	∅	∅	∅	0
1	∅	∅	∅	∅	10	1	0
2	∅	11	∅	3	9	1	0
3	14	10	4	3	9	1	0
4	11	7	4	3	9	1	0
5	10	7	4	3	8	1	0
6							

a) Compute the paths from any node in the network to G. Use the provided table to fill in the state of each node at every time step. Stop when a stable state is reached. The first time step is provided as an example.

Solution: cf. table on the left

b) Highlight the actual paths taken in the graph.

Solution:

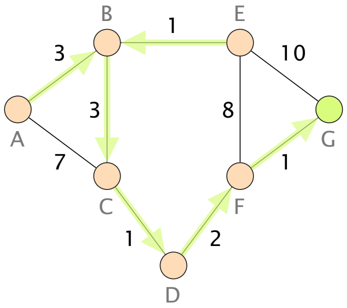

c) The network operator realizes that there is a potential bottleneck as all traffic is crossing the following links: *C*-*D*, *D*-*F*, and *F*-*G*. She prefers to balance the traffic across the available links in the network. Therefore, she would like to have all traffic from the nodes *A*, *B*, *E* to go across the link *E*-*G* and the traffic of the remaining nodes to go across *F*-*G*.

(i) If she can only change the weight of the link *E*-*G*, what should she change it to?

Solution: 6 or below

(ii) If she cannot change the weight of the link *E*-*G*, what should she change instead? Propose a change that requires to change the weights of as few links as possible.

Solution: She could set the weight of *F*-*G* to a value in the range from 5 to 10.

3.2 Dijkstra's Algorithm with Link Failure

The routers in the network on the top left use Dijkstra's Algorithm to find the shortest path towards destination d. You can assume that every router knows the entire network graph. In case of a failure, the routers directly affected by it inform the other routers by flooding this information in the network.

The blue arrows indicate how the routers forward the traffic: $R2$, for example, sends packets for destination d to router $R3$.

Now the link between router $R3$ and d fails (network at the bottom left) and $R3$ can no longer send packets towards destination d. As $R3$ is directly connected to the failed link, it detects the failure immediately. It starts to flood this information, such that all the routers can update their network view and recompute Dijkstra's algorithm.

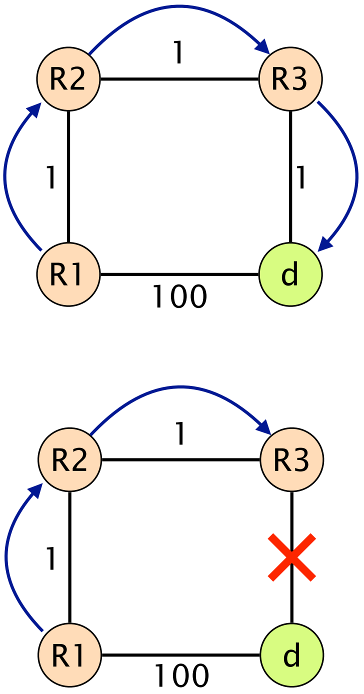

Dijkstra's algorithm with a link failure

a) What is the new shortest-path from $R3$ towards destination d?

Solution: $R3, R2, R1, d$

b) Assume now that the computation of the new shortest-path is *very* fast and finishes before $R3$ starts the flooding of the messages announcing the link failure. $R3$ sends a packet towards d using the new shortest-path. Will the packet reach its destination? Which path will it take?

Solution: No, $R2$ does not yet know about the link failure and did not update its shortest-path towards d. It will send the packet back towards $R3$. The packet is therefore stuck in a forwarding loop.

c) Can you find a sequence of link failure messages and shortest-path computations such that the problem discovered in the previous task is observed between $R1$ and $R2$? The *only* link failure is still between router $R3$ and d.

Solution: $R2$ did receive a link failure message and updated its shortest-path. $R2$ then sends a packet towards d to the next-hop $R1$ before $R1$ can update its shortest-path.

d) As we have discovered, the order in which the link failure messages are processed and the new shortest-paths are computed is crucial for a correct forwarding behavior in the network. In which order should the router update their forwarding tables, such that the previously observed problems will not occur? Can you find a more general "rule" for a safe ordering of forwarding rule updates?

Solution: Good order: $R1, R2, R3$. To prevent forwarding loops, the routers should update their forwarding tables based on their distance to the destination. The router nearest to the destination should update its forwarding table first.

Ethernet & Switching

3.3 Duplicate MAC Address

Consider three hosts Alice, Bob, and Eve connected through the network below composed of 3 Layer 2 (Ethernet) switches.

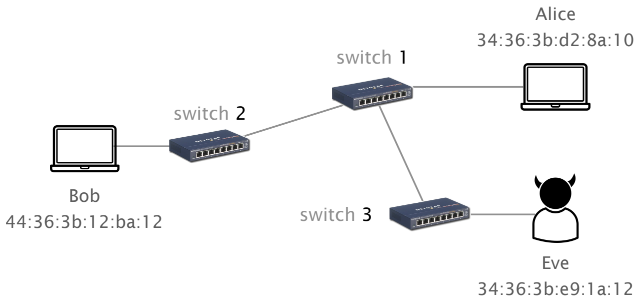

In the beginning the tables of the learning switches are still empty. Bob starts sending Ethernet frames to Alice. Eve is curious and wants to know what Bob is sending to Alice. Assume that Bob and Alice know the MAC address of each other.

a) What is the source and destination address in the Ethernet header for frames sent from Bob to Alice?

Solution: Source address: 44:36:3b:12:ba:12
Destination address: 34:36:3b:d2:8a:10

b) What do the switches do when they receive the frames?

Solution: Each switch adds a new entry to its table with the source MAC address and the incoming port. As the address of Alice is not yet in any of the switch tables, each switch floods the frame on all ports, but the port the packet came in on. This means the frame is sent to both Alice and Eve.

c) Due to the flooding, the frames are sent to both Alice and Eve. Does Eve actually receive the frames? (*hint:* promiscuous mode).

Solution: As long as Eve's Ethernet adapter is not set to promiscuous mode, the frame is not decapsulated and Eve will not receive it.

Alice starts acknowledging the received frames by sending frames to Bob.

d) Is Eve able to eavesdrop either on the frames being sent from Alice to Bob or on new frames sent from Bob to Alice? Explain.

Solution: No. The frames from Alice to Bob will not be flooded as the switches already know the path. After the first frame from Alice reaches Bob, the switches have also learned over which ports Alice can be reached. Frames from Bob to Alice are therefore no longer flooded.

e) Can you think of a way for Eve to redirect the frames destined to Alice again to herself?

Solution: Eve can send an Ethernet frame destined to Bob with the source address set to the MAC address of Alice. The switches will update their tables and Eve will receive the frames for Alice as long as Alice does not send a packet.

3.4 Impostor

The three hosts Bob, Alice and Eve are all connected to the same network, which has a DHCP server.

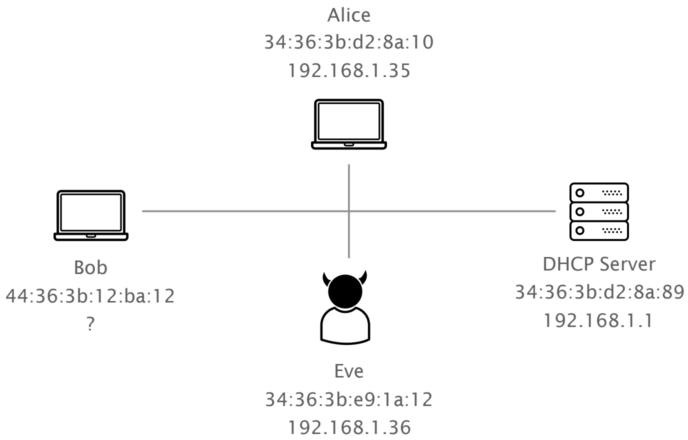

Bob just connected to the network and wants to send important IP packets to Alice. Bob only knows the IP address of Alice (`192.168.1.35`) and his laptop is not yet configured with an IP address.

a) Explain all the steps that are necessary such that Bob's computer can finally send packets to Alice.

Solution: Please note that the lecture slides introduce a simplified version of the DHCP protocol which only shows the first two steps (discovery and offer). This is enough to solve the question, i.e. afterwards Bob is able to communicate with Alice as he knows which IP to use. However, in reality we also have a request and ack step which are also shown in the table below. This way Bob tells the DHCP server that he accepts the IP address and the server sends an acknowledgement back. It now also knows that the given IP is currently used.

You might wonder why Bob uses the broadcast address as DST MAC in the DHCP request step instead of the MAC address which belongs to the DHCP server (known from the previous DHCP offer step). In bigger networks, you often have multiple DHCP servers, e.g. for redundancy. After the discovery message each of the DHCP servers will send an offer to Bob. Afterwards Bob selects one offer and sends the corresponding DHCP request. By broadcasting this message, all DHCP servers in the network will know if their offer was either picked (in this case they will send a DHCP ack back) or not picked, in which case they can use the offered IP address again for the next discovery message they get (they will not send an ACK back).

SRC MAC address	DST MAC address	Message type	Message content
44:36:3b:12:ba:12	ff:ff:ff:ff:ff:ff	DHCP discovery	I need an IP address
34:36:3b:d2:8a:89	44:36:3b:12:ba:12	DHCP offer	use 192.168.1.37
44:36:3b:12:ba:12	ff:ff:ff:ff:ff:ff	DHCP request	I want the offered IP
34:36:3b:d2:8a:89	44:36:3b:12:ba:12	DHCP ack	Lease duration & configuration
44:36:3b:12:ba:12	ff:ff:ff:ff:ff:ff	ARP request	Who has 192.168.1.35 Tell 192.168.1.37
34:36:3b:d2:8a:10	44:36:3b:12:ba:12	ARP reply	192.168.1.35 is at 34:36:3b:d2:8a:10

b) Eve is very interested to find out what Bob is sending to Alice. What could she do to intercept Bob's packets?

Solution: When Bob sends the ARP request to learn the MAC address of Alice, Eve also receives it as it is destined to the MAC broadcast address (`ff:ff:ff:ff:ff:ff`). If Eve can send a fake reply to Bob before Alice does so, she can make Bob believe that her MAC address is the one of Alice. This is called ARP spoofing.

3.5 MAC-Learning (Exam question from 2021)

Consider the Local Area Network (LAN) made up of 4 Ethernet switches in the figure below. Several hosts (A, B, C, D, E) are connected to the switches. The MAC tables of all switches are still empty.

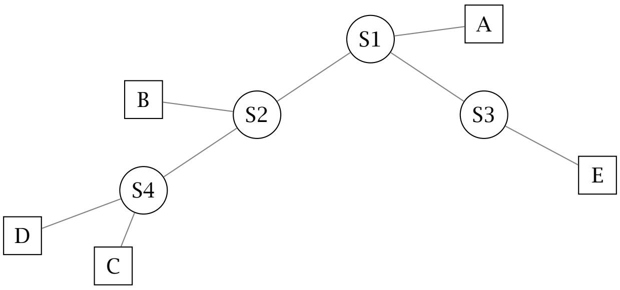

a) Host A sends a packet to host B. List below all the hosts that will receive the packet. In addition, fill in the MAC tables of all switches with the learned information.

Hosts receiving the packet:

Solution: All hosts receive the packet since the MAC tables are still empty and all the switches simply flood the packet.

S1 MAC-Table	
dst	next hop
A	connected

S2 MAC-Table	
dst	next hop
A	S1

S3 MAC-Table	
dst	next hop
A	S1

S4 MAC-Table	
dst	next hop
A	S2

b) Host C sends a packet to host A. Again, list all the hosts that receive the packet and update the MAC tables with the learned information. The entries from task a) are still available.

Hosts receiving the packet:

Solution: Only A will receive the packet as all the switches have learned through which port they can reach A.

S1 MAC-Table	
dst	next hop
A	connected
C	S2

S2 MAC-Table	
dst	next hop
A	S1
C	S4

S3 MAC-Table	
dst	next hop
A	S1

S4 MAC-Table	
dst	next hop
A	S2
C	connected

c) After some time, the switches have full MAC-tables (i.e., they have an entry for each host in the network). Host B wants to hijack all the packets destined to host A. By only sending packets, how can host B manipulate the switches in the network to receive all that traffic? How many "manipulation" packets are minimally necessary and to which addresses does host B have to send them? Explain your approach, state the required number of manipulated packets, and list the source and destination addresses of all manipulated packets.

Note: The hosts are not aware of the other hosts and do not know the network's topology.

Solution: B can send a packet destined to almost any address in the network (not any because if the switches already learned the address, not all switches will be reached) with the source address set to the MAC address of A. The switches will update their tables and B will receive the frames for A as long as A does not send a packet. Therefore, B needs to send minimally **1** packet, which is destined to a random MAC address that is not present in the network.