Link Layer: Switching

Dr. Yingwu Zhu

Interconnecting LANs

- Bridges (aka Ethernet switches) were introduced to allow the interconnection of several local area networks (LANs) without a router.
- By partitioning a large LAN into multiple smaller networks, there are fewer collisions, and more parallel communications.

• Data center networks

It is now common for the port of an Ethernet switch to connect to just one (or a small number of) hosts.

An Ethernet Network



*Lots of collisions reduces efficiency.

Ethernet Switching



Benefits:

- Number of collisions is reduced. If only one computer per port, no collisions can take place (each cable is now a self-contained point-to-point Ethernet link).
- Capacity is increased: the switch can forward multiple frames to different computers at the same time.

Outside

world

Router

A Real-World Ethernet Switch



Ethernet Switching

- 1. Examines the header of each arriving frame.
- 2. If the Ethernet DA is in its table, it forwards the frame to the correct output port(s).
- 3. If the Ethernet DA is not in its table, it broadcasts the frame to all ports (except the one through which it arrived).
- 4. The table is learned by examining the Ethernet SA of arriving packets.
 - Initially empty
 - Upon received frame, create an entry based on source MAC, incoming interface and time.
 - Purge obsolete entries

Ethernet Switching

Learning addresses





Q: How do we prevent loops?

Preventing loops Spanning Tree Protocol

- The network of switches is a graph.
- The Spanning Tree Protocol finds a a subgraph that spans all the vertices without loops.
 - Spanning => all switches are included.
 - Tree => the topology has no loops.
- □ The distributed protocol runs:
 - To determine which switch is the root of the tree, and
 - Switches only forward on ports that are part of the tree.

Example Spanning Tree



Protocol operation:

- . Picks a root
- 2. For each LAN,
 - picks a designated switch that is closest to the root.
- All switches on a LAN send packets towards the root via the designated switch.

Example Spanning Tree



Spanning Tree Protocol

Periodically, all switches broadcast: 1.

(ID of sender, ID of root, distance from sender to root).

- 2. Initially, every switch claims to be root: sets distance field to 0.
- 3. Every switch broadcasts until it hears a "better" message:
 - a. A root with a smaller ID
 - A root with equal ID, but with shorter distance b
 - Ties broken by smaller ID of sender. С.
- 4. If a switch hears a better message, retransmit message (add 1 to distance).
- 5. If not the designated switch for a LAN, it stops sending configuration messages to that LAN.

Eventually:

- Only the root generates configuration messages, Other switch send configuration messages to LANs for which they are the designated switch.

Switches are not Routers!

Switches forward packets using MAC addresses.