CSE/EE 461 – Lecture 7

Bridging LANs

David Wetherall
djw@cs.washington.edu

Last Two Times …

- Medium Access Control (MAC) protocols
  - Part of the OSI Data Link Layer
  - Local Area Networks (LANs)
- How do multiple parties share a wire or the air?
  - Random access protocols (CSMA variants, Ethernet)
  - Contention-free protocols (turn-taking like FDDI, reservations like DQDB)
  - Wireless protocols (hidden/exposed terminals solutions like CSMA/CA and MACA)

This Time

- Focus:
  - What to do when one shared LAN isn't big enough?
- Interconnecting LANs
  - Bridges and LAN switches
  - A preview of the Network layer

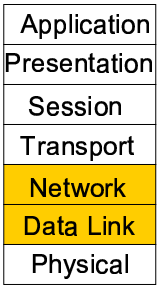

Limits of a LAN

- One shared LAN can limit us in terms of:
  - Distance
  - Number of nodes
  - Performance

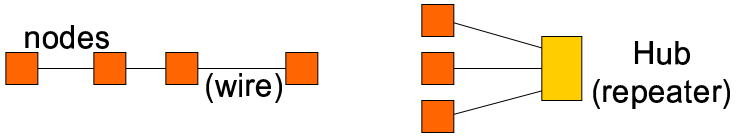

- How do we scale to a larger, faster network?
  - We must be able to interconnect LANs

Bridges and Extended LANs

- "Transparently" interconnect LANs with bridge
  - Receive frames from each LAN and forward to the other
  - Each LAN is its own collision domain; bridge isn't a repeater
  - Could have many ports or join to a remote LAN

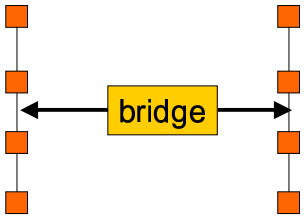

Backward Learning Algorithm

- To optimize overall performance:
  - Shouldn't forward A→B or C→D, should forward A→C and D→B

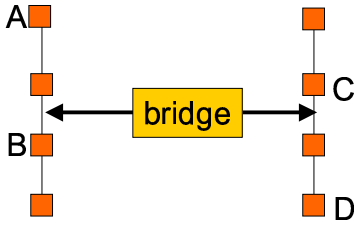

- How does the bridge know?
  - Learn who is where by observing source addresses and prune
  - Forward using destination address; age for robustness

Why stop at one bridge?

- But to avoid loops we must forward only on select bridge ports!
- The Spanning Tree algorithm does this
  - Separate from backward learning

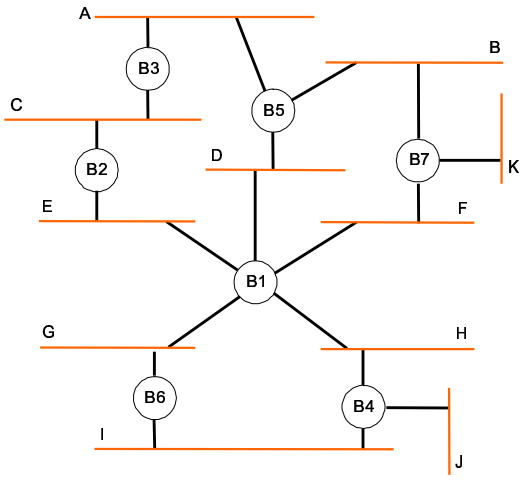

Spanning Tree Example

- Spanning tree uses select bridges so there are no cycles
  - Prune some ports
  - Only one tree
- Q: How do we find a spanning tree?
  - Automatically

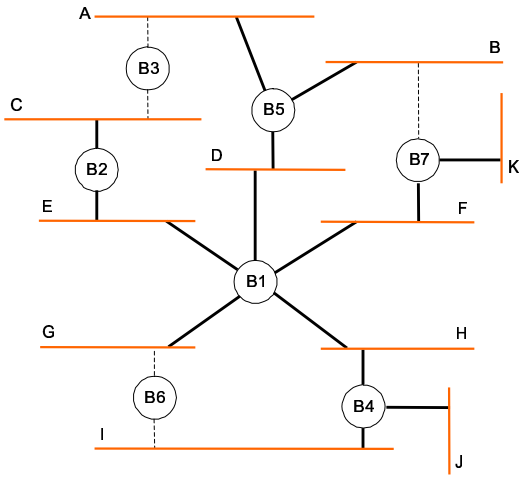

Spanning Tree Algorithm

- Distributed algorithm to compute spanning tree
  - Robust against failures, needs no organization
- Outline:
  - Goal is to turn some bridge ports off
  1. Elect a root node of the tree (lowest address)
  2. Grow tree as shortest distances from the root (using lowest address to break distance ties)
  - All done by bridges sending periodic configuration messages over ports for which they are the "best" path
  - Then turn off ports that aren't on "best" paths

Algorithm continued

- Each bridge sends periodic messages to others containing:
  - Its address, address of the root bridge, and distance (in hops) to root
- Each bridge receives messages, updates "best" config.
  - Smaller root address is better, then shorter distance
  - To break ties, bridge with smaller address is better
- Initially, each bridge thinks it is the root
  - Sends configuration messages on all ports
- Later, bridges send only "best" configs
  - Add 1 to distance, send configs where still "best" (called the designated bridge)
  - Turn off forwarding on ports except those that send/receive "best"

Algorithm Example

- Message format: (root, dist to root, bridge)
- Sample messages sequences to and from B3:
  1. B3 sends (B3, 0, B3) to B2 and B5
  2. B3 receives (B2, 0, B2) and (B5, 0, B5) and accepts B2 as root
  3. B3 sends (B2, 1, B3) to B5
  4. B3 receives (B1, 1, B2) and (B1, 1, B5) and accepts B1 as root
  5. B3 wants to send (B1, 2, B2) but doesn't as its nowhere "best"
  6. B3 receives (B1, 1, B2) and (B1, 1, B5) again … stable
     - Data forwarding is turned off to the LAN A

Some other tricky details

- Configuration information is aged
  - If the root fails a new one will be elected
- Reconfiguration is damped
  - Adopt new spanning trees slowly to avoid temporary loops

LAN Switches

- LAN switches are multi-port bridges
  - Modern, high performance form of bridged LANs
  - Looks like a hub, but frames are switched, not shared
  - Every host on a separate port, or can combine switches

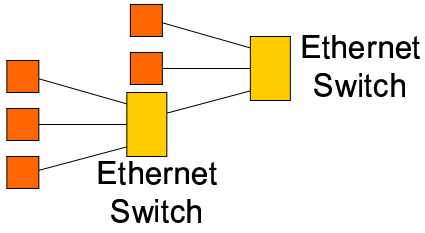

Limitations of Bridges/Switches

- LAN switches form an effective small-scale network
  - Plug and play for real!
- Why can't we build a large network using bridges?
  - Little control over forwarding paths
  - Size of bridge forwarding tables grows with number of hosts
  - Broadcast traffic flows freely over whole extended LAN
  - Spanning tree algorithm limits reconfiguration speed
  - Poor solution for connecting LANs of different kinds

Key Concepts

- We can overcome LAN limits by interconnection
  - Bridges and LAN switches
  - But there are limits to this strategy …
- Next Topic: Routing and the Network layer
  - How to grow large and really large networks