

CSE/EE 461 – Lecture 5

David Wetherall
djw@cs.washington.edu

Last Two Times ...

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

djw // CSE/EE 461, Winter 2000

L5.2

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

Application
Presentation
Session
Transport
Network
Data Link
Physical

djw // CSE/EE 461, Winter 2000

L5.3

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




djw // CSE/EE 461, Winter 2000

L5.4

Bridges and Extended LANs

- “Transparently” interconnect LANs
 - Receive frames from each LAN and forward to the other
 - Each LAN is its own domain; a bridge is not a repeater
 - Could have many ports or join to a remote LAN




djw // CSE/EE 461, Winter 2000

L5.5

Backward Learning Algorithm

- To increase overall performance:
 - Shouldn't forward $A \rightarrow B$ or $C \rightarrow D$, should forward $A \rightarrow C$ and $D \rightarrow B$




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

djw // CSE/EE 461, Winter 2000

L5.6

Why stop at one bridge?

- LANs and bridges form a graph
 - LANs = nodes, bridges = edges
- But to avoid loops we forward only on select bridge ports!
 - Spanning Trees




djw // CSE/EE 461, Winter 2000

L5.7

Spanning Tree Example

- Spanning tree is a subset of the graph that spans it but has no cycles
 - Prune some ports
- Q: How do we find a spanning tree?



djw // CSE/EE 461, Winter 2000

L5.8

Spanning Tree Algorithm

- Distributed algorithm to compute spanning tree
 - Robust against failures, needs no organization
- Outline:
 - Goal is to turn some bridge ports off
 - Bridges send periodic "best" configuration messages
 - Elect a root node of the tree (lowest address)
 - Grow tree as shortest distances from the root
 - Turn off ports that aren't on shortest paths

djw // CSE/EE 461, Winter 2000

L5.9

Algorithm continued

- Each bridge sends periodic messages 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" (designated bridge)
 - Turn off forwarding on ports except those that send/receive "best"

djw // CSE/EE 461, Winter 2000

L5.10

Algorithm Example

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

djw // CSE/EE 461, Winter 2000

L5.11

Some other 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

djw // CSE/EE 461, Winter 2000

L5.12

LAN Switches

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



djw // CSE/EE 461, Winter 2000

L5.13

Limitations of Bridges

- 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 with different MAC protocols

djw // CSE/EE 461, Winter 2000

L5.14

Key Concepts

- We can overcome LAN limits by interconnection
 - Bridges and LAN switches
 - But there are limits to this strategy ...
- Up Next: How to grow large and really large networks
- Network layer functionality
 - Routers (addressing/routing)
 - Scaling the size of the network
 - Combining different LAN technologies

djw // CSE/EE 461, Winter 2000

L5.15