

Lecture 10: Internetworking

CSE 123: Computer Networks
Stefan Savage





Lecture 10 Overview

- Finish bridging
 - ◆ Spanning Tree
- Internet Protocol
 - ◆ Service model
 - ◆ Packet format



Learning bridges recap


- Each bridge keeps a list mapping link-layer *destination* address to port number
(what are the directions to this destination?)
- This list is populated by looking at the *source* address of each packet it receives on a given port and entering those values in the table
(if a packet from A came from port x, then packets to A should be sent on port x)
- If a packet arrives with a destination address not in the table, then send on all ports (except the one it came on)
- Simple, automatic , self healing



Network Topology


- Linear organization

- ◆ Inter-bridge hubs (e.g. CS) are single points of failure
- ◆ Unnecessary transit (e.g. EE $\leftrightarrow$ SE must traverse CS)



- Backbone/tree


- ◆ Can survive LAN failure
- ◆ Manages all inter-LAN communication
- ◆ Requires more ports





An Issue: Cycles


- Learning works well in tree topologies
- But trees are fragile
 - ◆ Net admins like redundant/backup paths
- How to handle Cycles?
 - ◆ Where should *B1* forward packets destined for LAN *A*?





Spanning Tree

- Spanning tree uses *subset* of bridges so there are no cycles
 - ◆ Prune some ports
 - ◆ Only **one** tree
- Q: How do we find a spanning tree?
 - ◆ Automatically!
 - ◆ Elect root, find paths





Spanning Tree Algorithm

- Each bridge sends periodic configuration messages
 - ◆ (RootID, Distance to Root, BridgeID)
 - ◆ All nodes think they are root initially
- Each bridge updates route/Root upon receipt
 - ◆ Smaller root address is better
 - ◆ Select port with lowest cost to root as “root port”
 - ◆ To break ties, bridge with smaller address is better
- Rebroadcast new config to ports for which we’re “best”
 - ◆ Don’t bother sending config to LANs with better options
 - ◆ Add 1 to distance, send new configs on ports that haven’t told us about a shorter path to the root
- Only forward packets on ports for which we’re on the shortest path to root (prunes edges to form tree)




Spanning Tree Example

- Sample messages 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, B3) but doesn't as its nowhere "best"
6. B3 receives (B1, 1, B2) and (B1, 1, B5) again and again...

Data forwarding is turned off for LAN A






Important Details

- What if root bridge fails?
 - ◆ Age configuration info
 - » If not refreshed for MaxAge seconds then delete root and recalculate spanning tree
 - » If config message is received with a more recent age, then recalculate spanning tree
 - ◆ Applies to all bridges (not just root)
- Temporary loops
 - ◆ When topology changes, takes a bit for new configuration messages to spread through the system
 - ◆ Don't start forwarding packets immediately -> wait some time for convergence



Switched Ethernet

- Hosts directly connected to a bridge
 - ◆ learning + spanning tree protocol
- Switch supports parallel forwarding
 - ◆ A-to-B and A'-to-B' simultaneously
 - ◆ Generally full duplex as well
- **Switch** backplane capacity varies
 - ◆ Ideally, **nonblocking**
 - ◆ I.e., can run at full line rate on all ports
- No longer any shared bus
 - ◆ Each link is its own collision domain
 - ◆ Collision detection largely irrelevant





Layer-2 Forwarding

- Create spanning tree across LANs
 - ◆ Learn which ports to use to reach which addresses
- Benefits
 - ◆ Higher link bandwidth (point-to-point links)
 - ◆ Higher aggregate throughput (parallel communication)
 - ◆ Improved fault tolerance (redundant paths)
- Limitations
 - ◆ Requires homogeneous link layer (e.g. all Ethernet)
 - ◆ Harder to control forwarding topology
- What if we want to connect different link layers?



Combing Networks

- Main challenge is heterogeneity of link layers:
 - ◆ Addressing
 - » Each network media has a different addressing scheme
 - ◆ Bandwidth
 - » Modems to terabits
 - ◆ Latency
 - » Seconds to nanoseconds
 - ◆ Frame size
 - » Dozens to thousands of bytes
 - ◆ Loss rates
 - » Differ by many orders of magnitude
 - ◆ Service guarantees
 - » “Send and pray” vs reserved bandwidth




internetworking


- Cerf & Kahn74,
“*A Protocol for Packet Network Intercommunication*”
 - ◆ Foundation for the modern Internet
- Routers forward packets from source to destination
 - ◆ May cross many separate networks along the way
- All packets use a common Internet Protocol
 - ◆ Any underlying data link protocol
 - ◆ Any higher layer transport protocol



TCP/IP Protocol Stack



IP Networking



Routers




- A router is a store-and-forward device
 - ◆ Routers are connected to multiple networks
 - ◆ On each network, looks just like another host
 - ◆ A lot like a switch, but supports multiple datalink layers and makes decisions at the network layer
- Must be explicitly addressed by incoming frames (L2)
 - ◆ Not at all like a switch, which is transparent
 - ◆ Removes link-layer header, parses IP header (L3)
- Looks up next hop, forwards on appropriate network
 - ◆ Each router need only get one step closer to destination



IP Philosophy

- Impose few demands on network
 - ◆ Make few assumptions about what network can do
 - ◆ No QoS, no reliability, no ordering, no large packets
 - ◆ No persistent state about communications; no connections
- Manage heterogeneity at hosts (not in network)
 - ◆ Adapt to underlying network heterogeneity
 - ◆ Re-order packets, detect errors, retransmit lost messages...
 - ◆ Persistent network state only kept in hosts (fate-sharing)
- Service model: **best effort**, a.k.a. *send and pray*


IP Packet Header





Version field


- Which version of IP is this?
 - ◆ Plan for change
 - ◆ Very important!
- Current versions
 - ◆ 4: most of Internet today
 - ◆ 6: new protocol with larger addresses
 - ◆ What happened to 5?
Standards body politics.





Header length


- How big is IP header?
 - ◆ Counted in 32-bit words
 - ◆ Variable length
 - » Options
 - ◆ Engineering consequences of variable length...
- Most IP packet headers are 20 bytes long



Type-of-Service




- How should this packet be treated?
 - ◆ Care/don't care for delay, throughput, reliability, cost
 - ◆ How to interpret, how to apply on underlying net?
 - ◆ Largely unused until 2000 (hijacked for new purposes, ECN & Diffserv)



Length




- How long is whole packet in bytes?
 - ◆ Includes header
 - ◆ Limits total packet to 64K
 - ◆ Redundant?





TTL (Time-to-Live)


- How many more routers can this packet pass through?
 - ◆ Designed to limit packet from looping forever
- Each router decrements TTL field
- If TTL is 0 then router discards packet



Protocol



- Which transport protocol is the data using?
 - ♦ i.e. how should a host interpret the data
- TCP = 6
- UDP = 17






IP Checksum

- Header contains simple checksum
 - ◆ Validates content of header *only*

- Recalculated at each hop
 - ◆ Routers need to update TTL
 - ◆ Hence straightforward to modify



- Ensures *correct* destination receives packet



For Next Time (Monday)

- No class Friday
- Read 3.2.5-6, 9.3.1 in P&D
- Homework 2 due next Wednesday
- Project 1 due next Friday