15-441 Computer Networking
Lecture 8 – Switching and Bridging

Peter Steenkiste
Departments of Computer Science and
Electrical and Computer Engineering

15-441 Networking, Spring 2008
http://www.cs.cmu.edu/~dga/15-441/S08

Scale

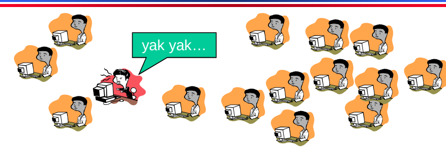

- What breaks when we keep adding people to the same wire?

Scale

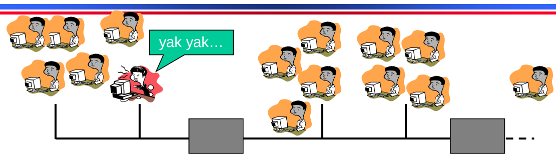

- What breaks when we keep adding people to the same wire?
- Only solution: split up the people onto multiple wires
  - » But how can they talk to each other?

Problem: How to Connect LANs?

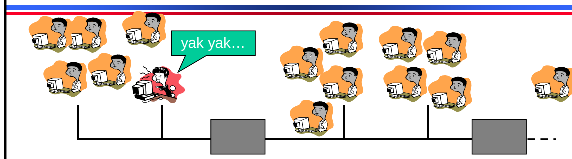

- When should these boxes forward packets between wires?
- How do you specify a destination?
- How does your packet find its way?

Outline

- **Bridging**
- Internetworks
  - » Methods for packet forwarding

Building Larger LANs: Bridges

- Extend reach of a single shared medium
- Connect two or more "segments" by copying data frames between them
  - » Only copy data when needed → key difference from repeaters/hubs
  - » Reduce collision domain compared with single LAN
  - » Separate segments can send at once → much greater bandwidth
- Challenge: learning which packets to copy across links

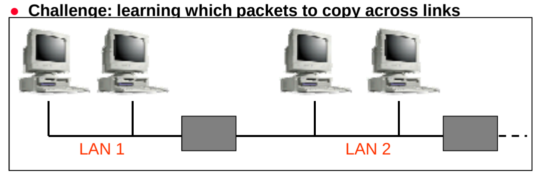

Transparent Bridges

- Design goals:
  - » Self-configuring without hardware or software changes
  - » Bridge do not impact the operation of the individual LANs
- Three parts to making bridges transparent:
  1) Forwarding frames
  2) Learning addresses/host locations
  3) Spanning tree algorithm

Frame Forwarding

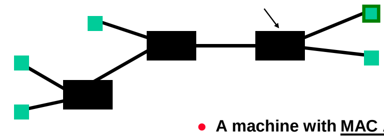

A21032C9A591	1	36
99A323C90842	2	01
8711C98900AA	2	15
301B2369011C	2	16
695519001190	3	11

- A machine with MAC Address lies in the direction of number port of the bridge
- For every packet, the bridge "looks up" the entry for the packets destination MAC address and forwards the packet on that port.
  - » Other packets are broadcast – why?
- Timer is used to flush old entries

Learning Bridges

- Manually filling in bridge tables?
  - » Time consuming, error-prone
- Keep track of source address of packets arriving on every link, showing what segment hosts are on
  - » Fill in the forwarding table based on this information

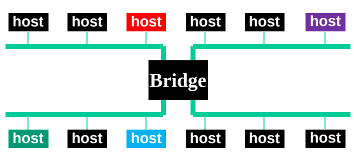

Spanning Tree Bridges

- More complex topologies can provide redundancy.
  - » But can also create loops.
- What is the problem with loops?
- Solution: spanning tree

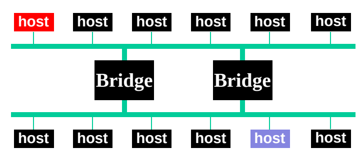

Spanning Tree Protocol Overview

Embed a tree that provides a single unique path to each destination:

1) Elect a single bridge as a root bridge
2) Each bridge calculates the distance of the shortest path to the root bridge
3) Each LAN identifies a *designated bridge*, the bridge closest to the root. It will forward packets to the root.
4) Each bridge determines a *root port*, which will be used to send packets to the root
5) Identify the ports that form the spanning tree

Spanning Tree Algorithm Steps

- Root of the spanning tree is the bridge with the lowest identifier.
  - » All ports are part of tree
- Each bridge finds shortest path to the root.
  - » Remembers port that is on the shortest path
  - » Used to forward packets
- Select for each LAN the designated bridge that has the shortest path to the root.
  - » Identifier as tie-breaker
  - » Responsible for that LAN

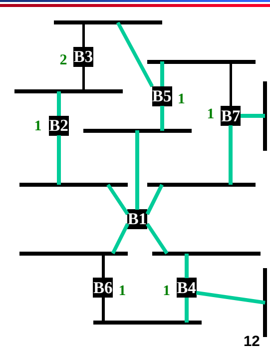

Spanning Tree Algorithm

- Each node sends configuration message to all neighbors.
  - » Identifier of the sender
  - » Id of the presumed root
  - » Distance to the presumed root
  - » E.g. B5 sends (B5, B5, 0)
- When B receive a message, it decide whether the solution is better than their local solution.
  - » A root with a lower identifier?
  - » Same root but lower distance?
  - » Same root, distance but sender has lower identifier?
- After convergence, each bridge knows the root, distance to root, root port, and designated bridge for each LAN.

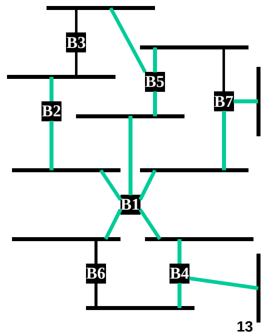

Spanning Tree Algorithm (part 2)

- Each bridge B can now select which of its ports make up the spanning tree:
  - » B's root port
  - » All ports for which B is the designated bridge on the LAN
- Bridges can not configure their ports.
  - » *Forwarding state* or *blocked state*, depending on whether the port is part of the spanning tree
- Root periodically sends configuration messages and bridges forward them over LANs they are responsible for.

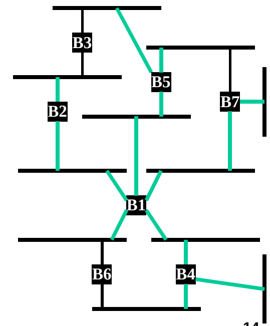

Spanning Tree Algorithm Example

- Node B2:
  - » Sends (B2, B2, 0)
  - » Receives (B1, B1, 0) from B1
  - » Sends (B2, B1, 1) "up"
  - » Continues the forwarding forever
- Node B1:
  - » Will send notifications forever
- Node B7:
  - » Sends (B7, B7, 0)
  - » Receives (B1, B1, 0) from B1
  - » Sends (B7, B1, 1) "up" and "right"
  - » Receives (B5, B5, 0) - ignored
  - » Receives (B5, B1, 1) - better
  - » Continues forwarding the B1 messages forever to the "right"

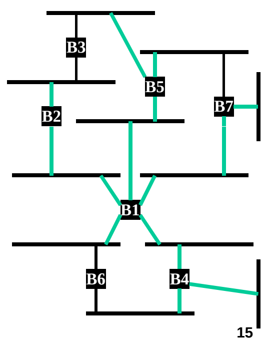

Ethernet Switches

- Bridges make it possible to increase LAN capacity.
  - » Packets are no longer broadcasted - they are only forwarded on selected links
  - » Adds a switching flavor to the broadcast LAN
- Ethernet switch is a special case of a bridge: each bridge port is connected to a single host.
  - » Can make the link full duplex (really simple protocol!)
  - » Simplifies the protocol and hardware used (only two stations on the link) – no longer full CSMA/CD
  - » Can have different port speeds on the same switch
    - – Unlike in a hub, packets can be stored
    - – An alternative is to use cut through switching

Ethernet Evolution

Early Implementations

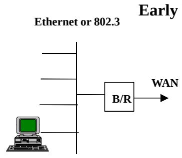

- ·A Local Area Network
- ·MAC addressing, non-routable
- ·BUS or Logical Bus topology
- ·Collision Domain, CSMA/CD
- ·Bridges and Repeaters for distance/capacity extension
- ·1-10Mbps: coax, twisted pair (10BaseT)

Current Implementations

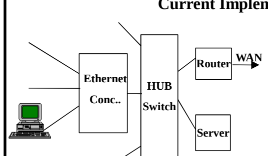

- ·Switched solution
- ·Little use for collision domains
- ·80% of traffic leaves the LAN
- ·Servers, routers 10 x station speed
- ·10/100/1000 Mbps, 10gig coming: Copper, Fiber
- ·95% of new LANs are Ethernet

CSMA Carrier Sense Multiple Access

CD - Collision Detection

Ethernet in a campus network

Today: x 10

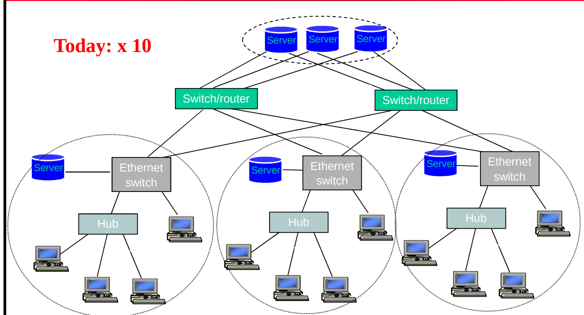

Problem: Bridging Weaknesses

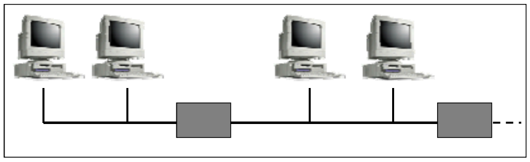

- Doesn't handle incompatible LAN technologies
  - » Is interoperable within 802.* standard
- How well does it scale?

Outline

- Bridging
- Internetworks
  - » Methods for packet forwarding

What is an Internetwork?

- Multiple incompatible LANs can be physically connected by layer 3 switches called ***routers***
- The connected networks are called an ***internetwork***
  - » The "***Internet***" is one (very big & successful) example of an internetwork

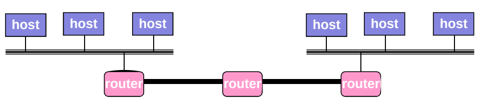

Logical Structure of Internet

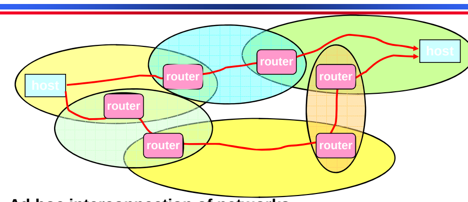

- » Ad hoc interconnection of networks
  - – No particular topology
  - – Vastly different router & link capacities
- » Send packets from source to destination by hopping through networks
  - – Router connect one network to another
  - – Different paths to destination may exist

Internet Protocol (IP)

- Hour Glass Model
  - » Create abstraction layer that hides underlying technology from network application software
  - » Make as minimal as possible
  - » Allows range of current & future technologies
  - » Can support many different types of applications

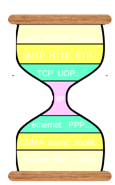

Problem: Internetwork Design

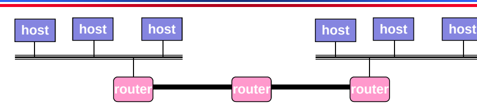

- How do I designate a distant host?
  - » Addressing / naming
- How do I send information to a distant host?
  - » What gets sent?
  - » What route should it take?
- Must support:
  - » Heterogeneity LAN technologies
  - » Scalability → ensure ability to grow to worldwide scale

Getting to a Destination

- How do you get driving directions?
- Intersections→routers
- Roads→links/networks
- Roads change slowly




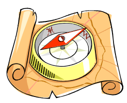

Forwarding Packets

- Table of virtual circuits
  - » Connection routed through network to set up state
  - » Packets forwarded using connection state
- Source routing
  - » Packet carries path
- Table of global addresses (IP)
  - » Routers keep next hop for destination
  - » Packets carry destination address

Simplified Virtual Circuits

- Connection setup phase
  - » Use other means to route setup request
  - » Each router allocates flow ID on local link
- Each packet carries connection ID
  - » Sent from source with 1st hop connection ID
- Router processing
  - » Lookup flow ID – simple table lookup
  - » Replace flow ID with outgoing flow ID
  - » Forward to output port

Simplified Virtual Circuits Example

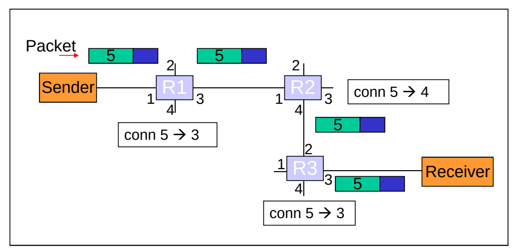

Virtual Circuits

- Advantages
  - » Efficient lookup (simple table lookup)
  - » Can reserve bandwidth at connection setup
  - » Easier for hardware implementations
- Disadvantages
  - » Still need to route connection setup request
  - » More complex failure recovery – must recreate connection state
- Typical use → fast router implementations
  - » ATM – combined with fix sized cells
  - » MPLS – tag switching for IP networks

Source Routing

- List entire path in packet
  - » Driving directions (north 3 hops, east, etc..)
- Router processing
  - » Strip first step from packet
  - » Examine next step in directions
  - » Forward to next step

Source Routing Example

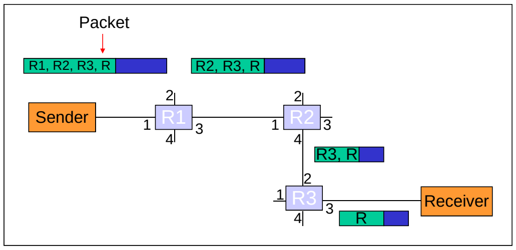

Source Routing

- Advantages
  - » Switches can be very simple and fast
- Disadvantages
  - » Variable (unbounded) header size
  - » Sources must know or discover topology (e.g., failures)
- Typical uses
  - » Ad-hoc networks (DSR)
  - » Machine room networks (Myrinet)

Global Addresses (IP)

- Each packet has destination address
- Each router has forwarding table of destination → next hop
  - » At v and x: destination → east
  - » At w and y: destination → south
  - » At z: destination → north
- Distributed routing algorithm for calculating forwarding tables

Global Address Example

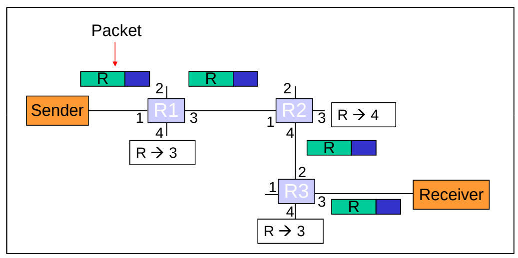

Global Addresses

- Advantages
  - » Stateless – simple error recovery
- Disadvantages
  - » Every switch knows about every destination
    - – Potentially large tables
  - » All packets to destination take same route
  - » Need routing protocol to fill table

Comparison

	Source Routing	Global Addresses	Virtual Circuits
Header Size	Worst	OK – Large address	Best
Router Table Size	None	Number of hosts (prefixes)	Number of circuits
Forward Overhead	Best	Prefix matching (Worst)	Pretty Good
Setup Overhead	None	None	Connection Setup
Error Recovery	Tell all hosts	Tell all routers	Tell all routers and Tear down circuit and re-route