University of Glasgow

Data Link Layer (2)

Networked Systems 3
Lecture 7

Lecture Outline

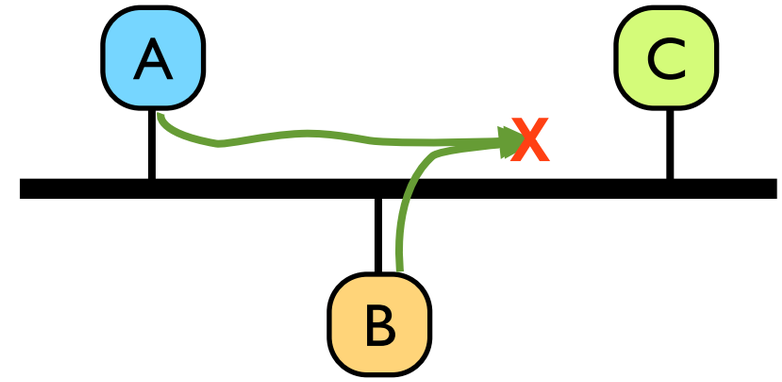
- Media access control
  - Contention, Token, Slotted
- Addressing
- Bridging

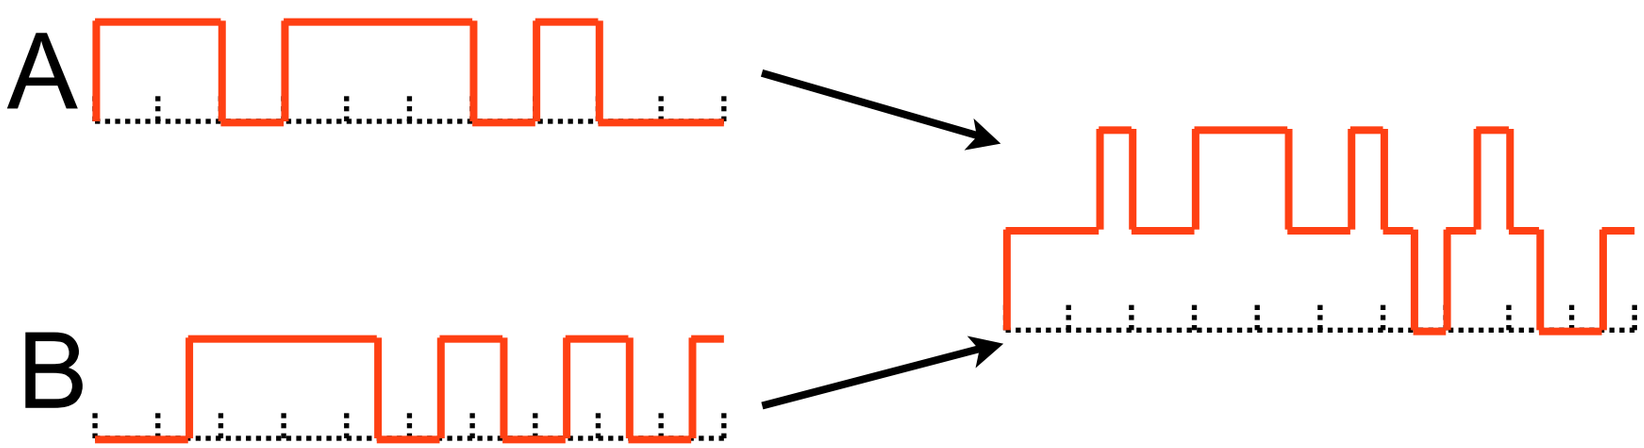
Media Access Control

- Links may be point-to-point or multi-access
- How to arbitrate access to the link?
  - Point-to-point links typically two unidirectional links
    - Separate physical cables for each direction
    - Need framing in each direction, but there is no contention for the link
    - ARQ with stop-and-wait or sliding-window for flow control
  - Multi-access links typically share a bidirectional link
    - A single physical cable – nodes contend for access to the link
    - A single radio frequency

Link Contention

- A *collision* occurs when two nodes transmit at once
- Signals overlap: only garbage received





Media Access Control

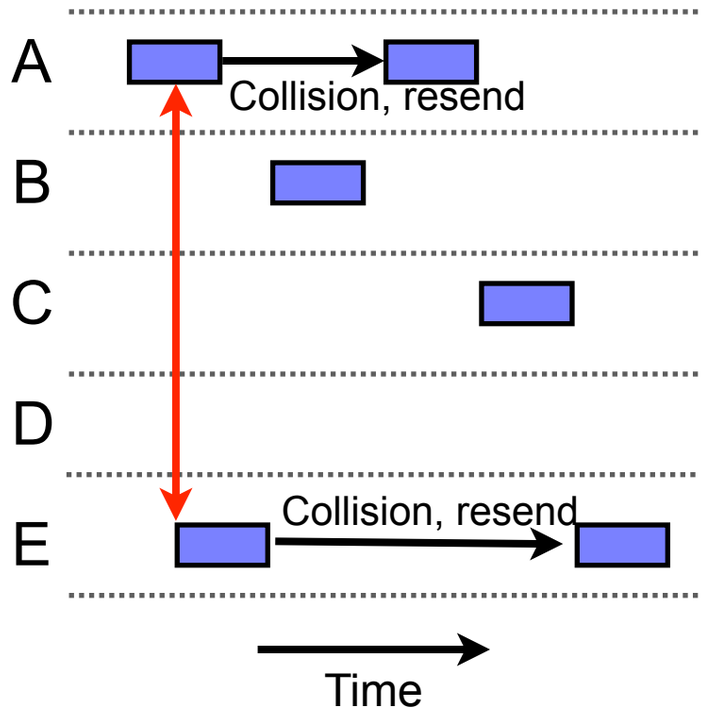
- Avoid link contention with *media access control* (MAC) protocol
  - Contention based (ALOHA, CSMA, etc.)
  - Token based (Token ring)
  - Slotted (TDMA)
- Different degrees of fairness, access policies, etc.

Contention-based MAC

- If multiple users share the channel in a way that can lead to conflicts, system is *contention-based*
  - Detect that a collision is occurring/will occur
    - By listening to the channel while/before sending
  - Back-off and retransmit according to some algorithm
    - Often randomised back-off delay
    - Might give priority to certain nodes or users
  - Gives probabilistic, variable latency, access to channel

The ALOHA Network

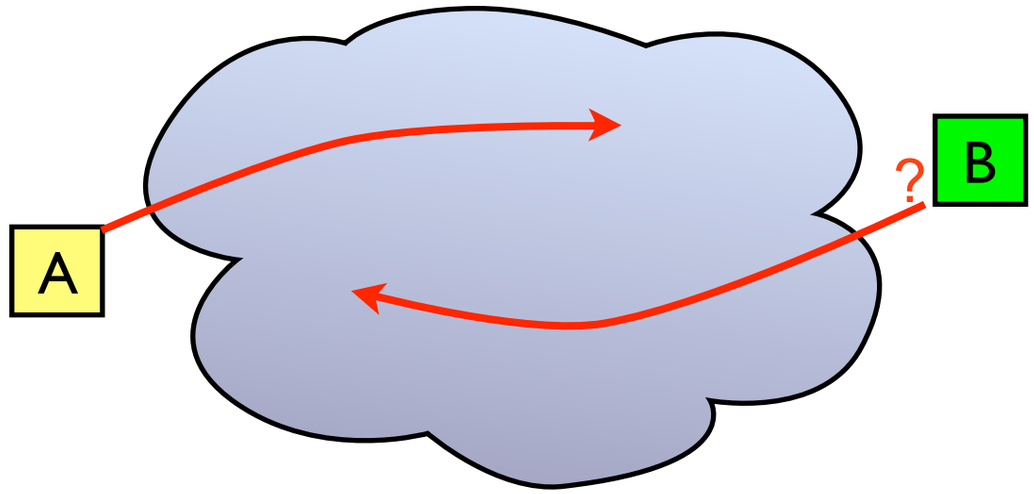
- Wireless network developed at the University of Hawaii (1970)
  - The first wireless packet switched network
- Simplest contention-based MAC
  - Try to transmit whenever data is available
  - If a collision occurs, wait random amount of time then retransmit; repeat until successful
- Simple, but poor performance
  - Low channel utilisation; long delays



Carrier Sense Multiple Access

- When *propagation delay* low, listen before trying to send ("CSMA")
  - If another transmission is active, back-off as if collision occurred – but without sending anything
  - Improves utilisation, since active transmission not disrupted by a collision occurring
  - Only one node backs-off and retransmits, rather than both nodes in ALOHA

Why does propagation delay matter?



A starts transmitting

B listens, hears no traffic since A's message has not reached it yet

B starts transmitting

Collision occurs, as messages overlap in transit; smaller propagation delay → less likely to occur

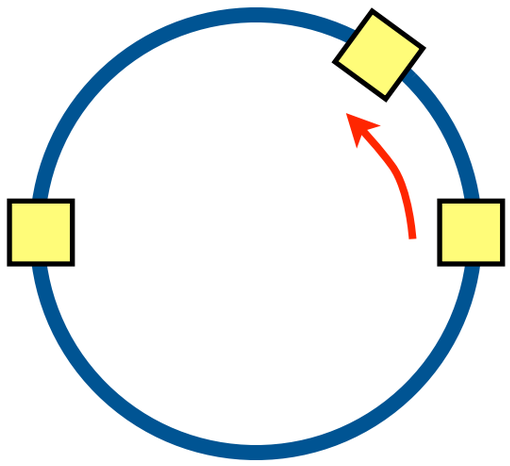
CSMA/CD

- Carrier Sense Multiple Access with Collision Detection (CSMA/CD)
  - Listen to channel before and while transmitting data
  - If collision occurs, immediately stop sending, back-off and retransmit
  - Reduces time channel is blocked due to collisions
    - Better performance than plain CSMA
- Examples: Ethernet, Wi-Fi

CSMA/CD: How to Back-Off?

- Contention protocols rely on receivers *backing-off* when a collision happens
- How long should the back-off time be?
  - Random – to avoid deterministic repeated collisions
  - Increasing with the number of collisions that affect a transmission – repeated collisions signal congestion, need to reduce transmission rate to allow network to recover

Token-based MAC

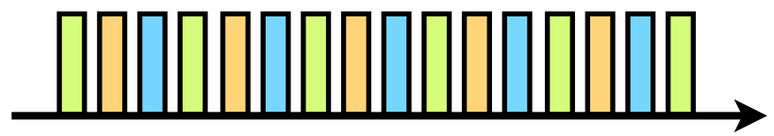


Token Ring

- Contention-based MAC protocols have poor worst case performance
  - Potentially long transmission delays, high jitter
- Token-based protocols more predictable
- Example: IBM token ring
  - Nodes circulate token on the ring if no data to send
  - Node wanting to send removes token, sends one packet, then restarts the token circulating
  - Token circulation enforces round robin transmission, and bounds maximum wait time before a node can transmit

Slotted MAC

- Split the channel into time (TDM) or frequency (FDM) slots
- Assign each sender a slot
  - Provides a limited, but guaranteed, capacity per-sender
  - QoS and fairness guarantees
- Example: GSM mobile phones



124 frequency bands, each 200 kHz wide
8 slots TDM within each frequency band

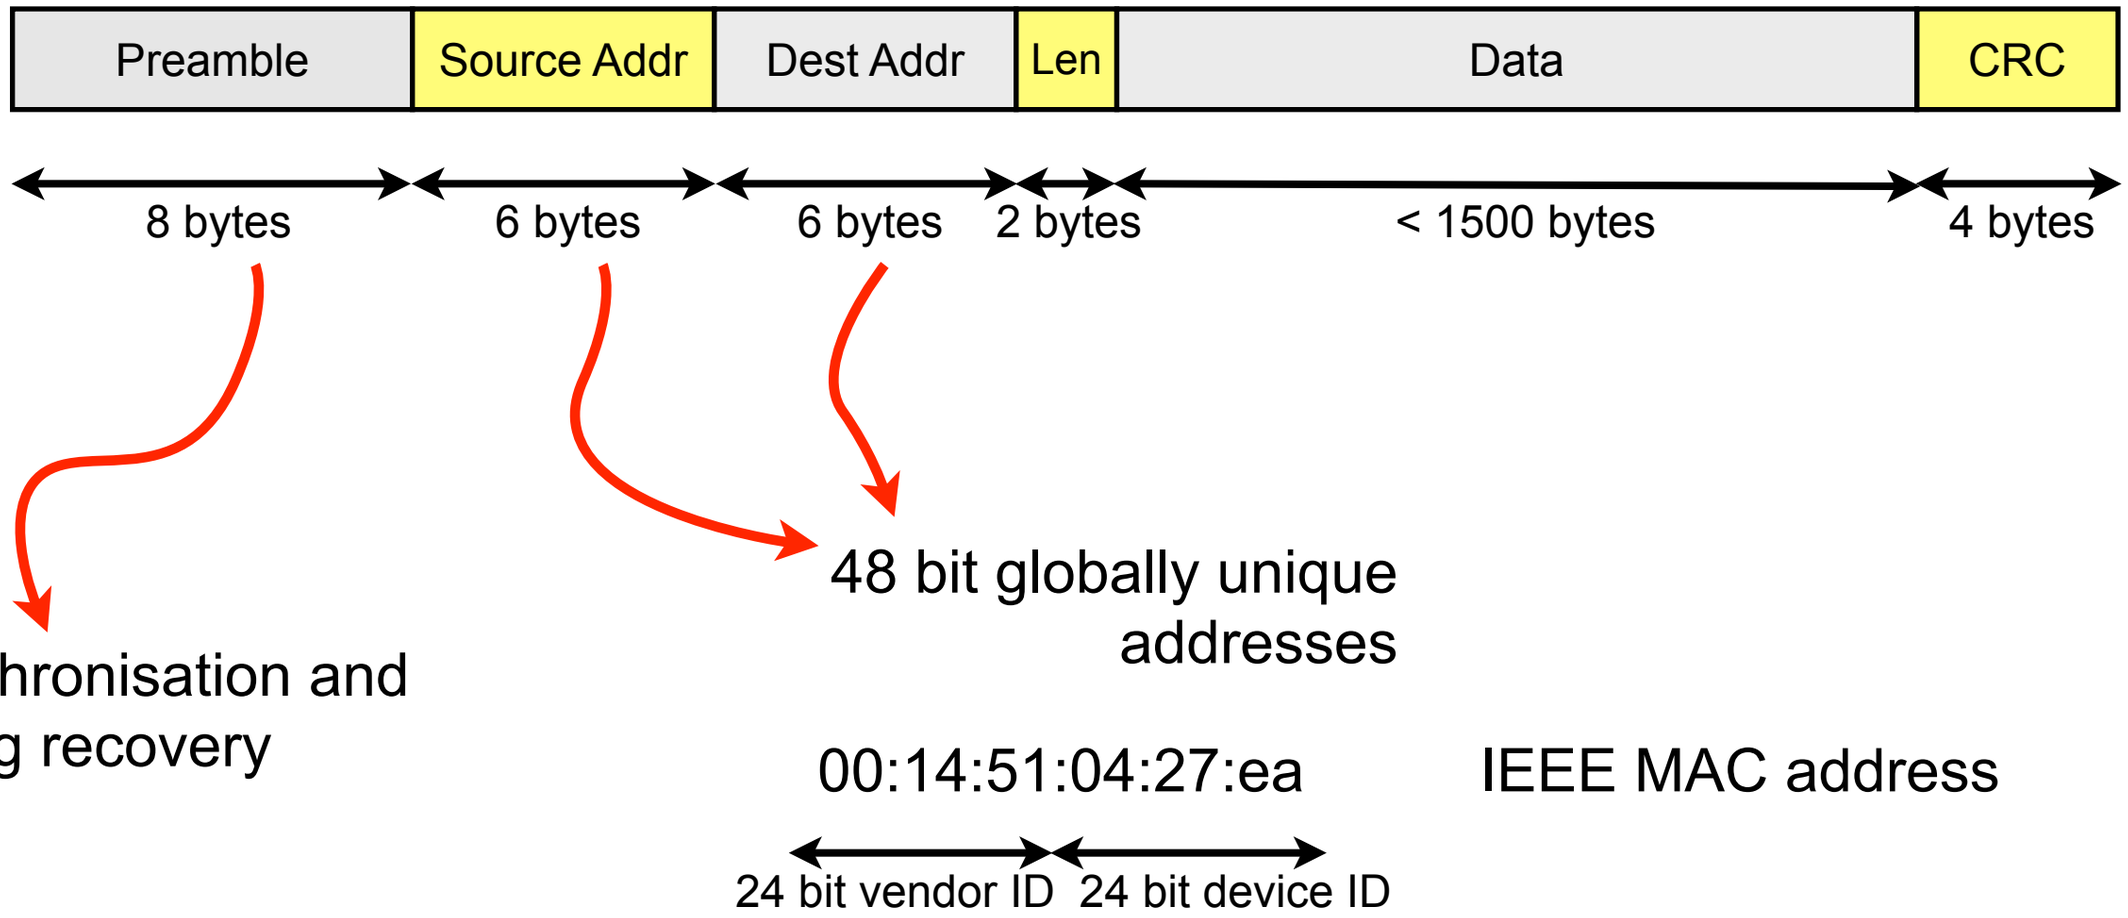
Design Trade-Off

- Which is best – contention, token, or slotted MAC?
  - As usual, it depends on the application
  - Contention protocols cheaper, hardware more readily available, but cannot guarantee real-time performance
  - Token-based or slotted MAC protocols offer stricter performance guarantees, but are more expensive and need less widely available hardware

Addressing

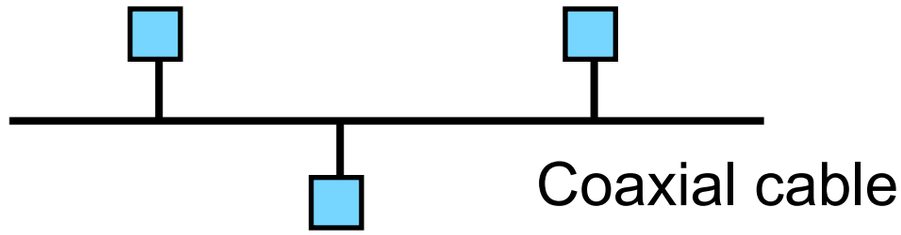
- Multiple nodes on a shared link → addressing
- Addresses may have *link local* or *global* scope
  - Only have to identify hosts on a single link
  - IEEE standard link layers (e.g. Ethernet, 802.11) use a global scope address
    - Easier to manage equipment
    - Don’t have to change address when moving devices

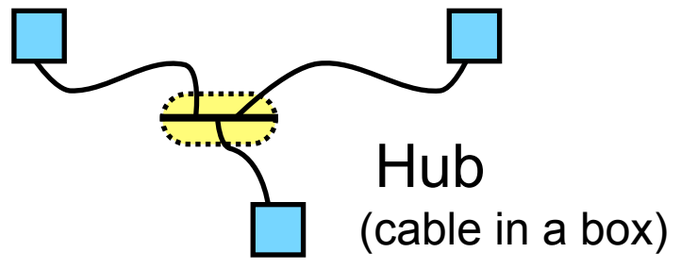
Example: Ethernet



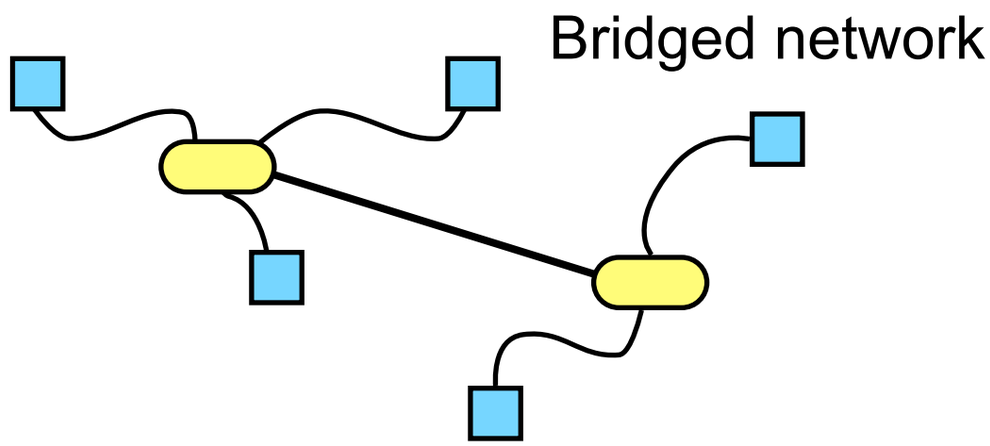
Bridging

- Ethernet Topology Evolution
  - Coaxial cable → hub for robustness
  - Hub → switch extends range of local area network (LAN)



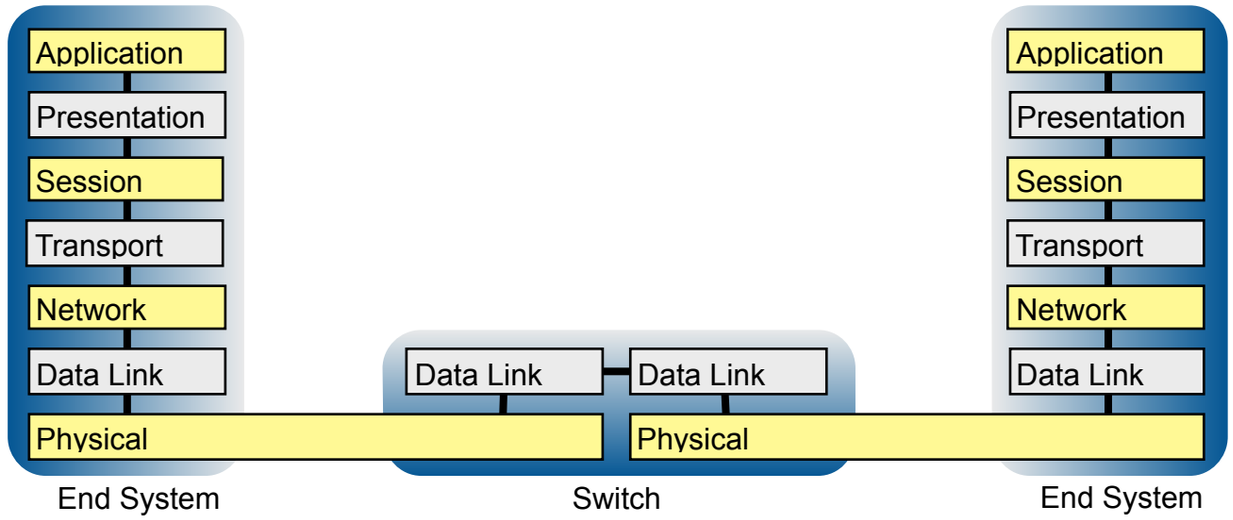


- *Bridging* links multiple LAN segments of the same type
  - Same conceptual network model
  - More nodes over a larger area than can be supported by a single link



Bridging Concepts

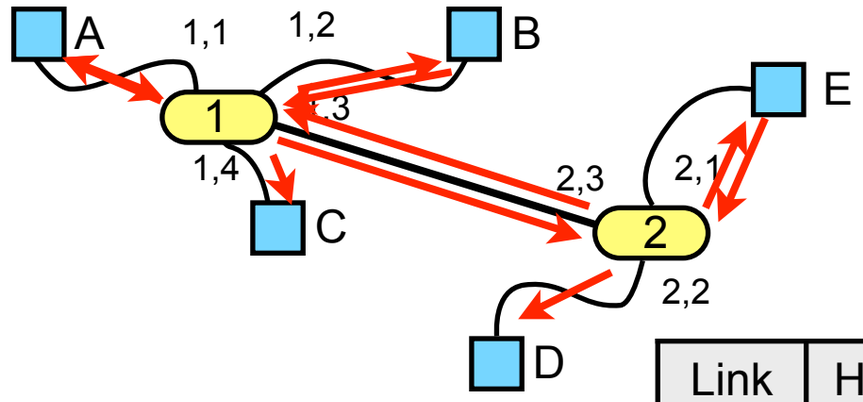
- Transparent connection of multiple LANs at the data link layer
- Automatic – needs zero configuration



- Two approaches:
  - Amplify and forward all traffic at the physical layer → hub
  - Process link-layer frames, identify location of hosts, forward only those frames of interest → bridge (a.k.a. “Ethernet switch”)
  - Bridges are more complex than hubs, but *much* more scalable

Basic Bridge Operation

- Learn addresses on each LAN
  - Observe source addresses of packets
  - *Soft state* time-out allows for graceful response to failure and node mobility
- Forward traffic as appropriate
  - Unicast traffic based on host locations (hash from address to destination link, flooding packets to unknown hosts)
  - Multicast based on group membership
  - Broadcast traffic



A sends to B

Packet flooded since B not known; bridges learn location of A

B responds; packet delivered without flooding since A is known

E sends to A; packet also delivered without flooding, since A is known due to previous flood

Link	Host
1,1	A
1,2	B
1,3	
1,4	

Link	Host
2,1	E
2,2	
2,3	A

Loops in Bridged Networks

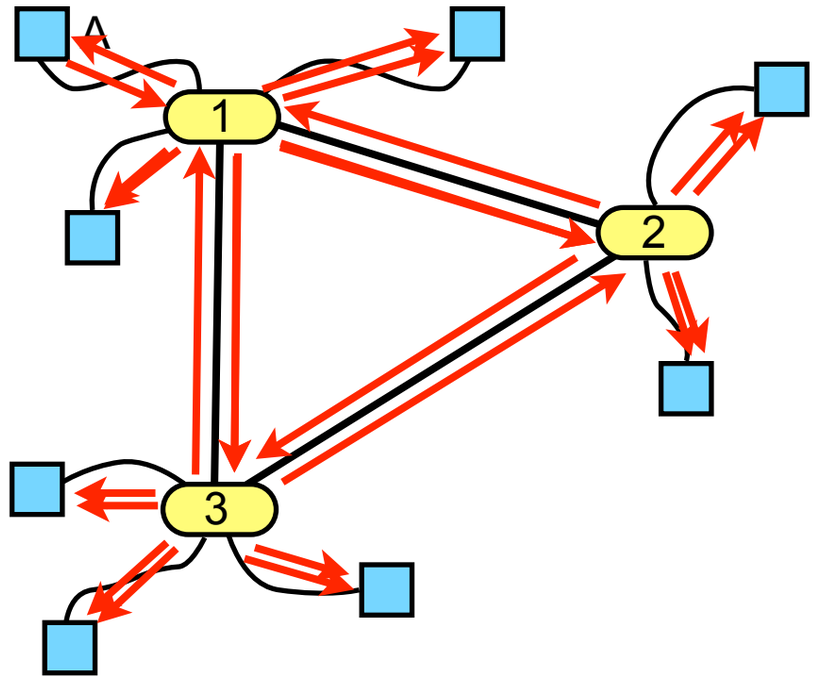
A sends packet to an unknown host

Bridge 1 floods the packet on all its links, hoping to find destination

Bridges 2 and 3 receive the flooded packet, and themselves flood it out on all their links

Bridges 3 and 2 receive the flooded packet from bridges 2 and 3, and themselves flood it out on all their links

Bridge 1 receives the flooded packet, and itself floods it, hoping to find the destination… The loop begins, and the packets cycle forever!



Loops in Bridged Networks

- Solution: build a *spanning tree* over the network, forward packets along this tree
  - Model network as an undirected graph, *G*
  - A *spanning tree* over that graph is a *tree* comprising all the vertices and some of the edges of *G*
    - Edges are removed eliminating loops, to leave the minimal set of edges that connect all vertices

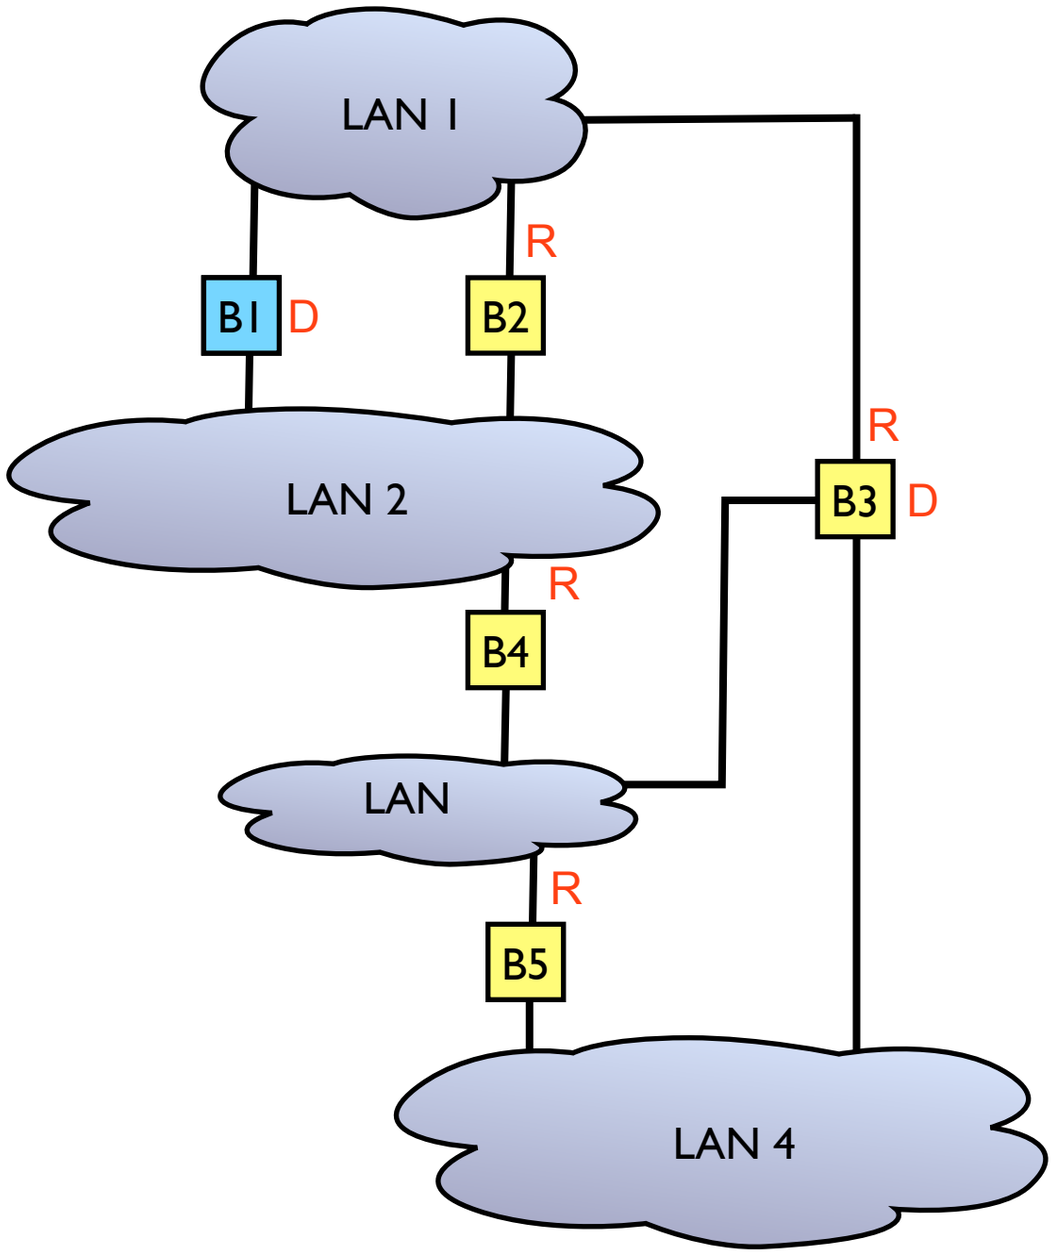
Spanning Tree Algorithm



Source: Sun Microsystems

- Distributed algorithm to build a spanning tree due to Radia Perlman:
  - Bridges broadcast their globally unique serial number
  - Bridge with lowest serial number becomes the root bridge
  - Determine the root port of each bridge except the root
    - Root port = port with shortest path to the root bridge
  - Select designated bridge for each LAN
    - Designated bridge = the bridge with the shortest path from the LAN to the root bridge
    - Designated port connects the LAN and the designated bridge
  - All root ports and all designated ports enabled; all other ports are disabled and cannot forward data

Spanning Tree Example



Bridge B1 elected as root bridge

Root ports selected for every other bridge

Designated bridge selected for each LAN

Designated and root ports enabled, others disabled

Resulting links allow all LANs to be reached via a loop-free path

Questions?