Bridging

Bridging makes networks:

- Bigger
- Smaller

(and it's not a multiple choice question)

Bridges, Spanning Trees & Switches

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

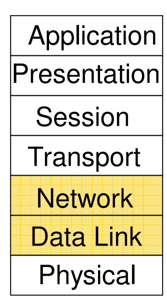

Warning: confusing terminology / pictures

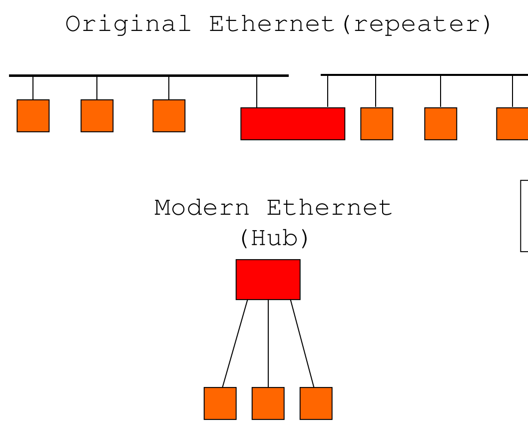

Multi-port “repeaters” – Hubs

That is not all….

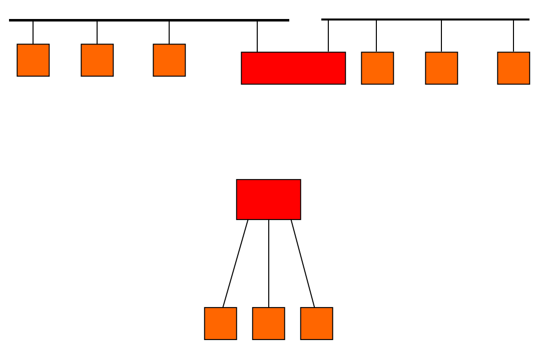

Bridges – well this is historical too …

Why do we need a bridge/switch?

- One shared LAN can limit us in terms of:
  - Distance *(why?)*
  - Number of nodes *(why?)*
- How do we scale to a larger, more efficient networks?
  - We must be able to interconnect LANs

Bridges

- Connected to different LANs over *ports*
- Operate in “promiscuous mode”
  - receives packet on one port and forwards it to the outgoing port
    - Example. From A to D
  - is “never” a communication endpoint itself
- What should each bridge do with an incoming frame?

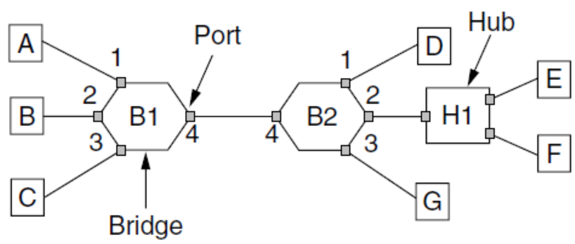

Bridges: What to do with incoming frame

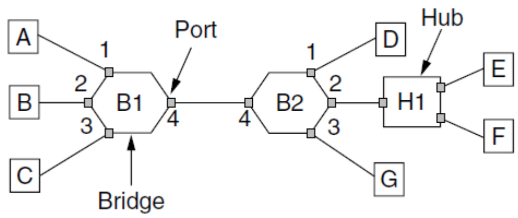

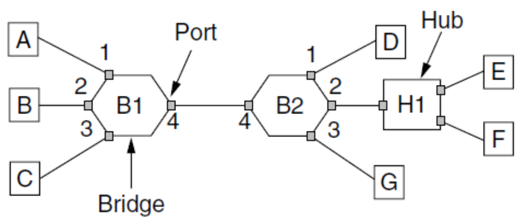

How are bridges configured?

- Backward learning algorithm
  1. At first, bridge doesn't know where any address is. Send incoming frame "everywhere"
  2. But, the incoming frame carries a source address. Remember it.
  3. If an incoming frame has a destination address I've remembered, send it only there

- Once we learn the ports, can we stop the backward learning algorithm?

Learning Bridges

Bridges extend the Link layer:

- Use but don't remove Ethernet header/addresses
- Don't inspect higher layer headers
- May need to store the packet

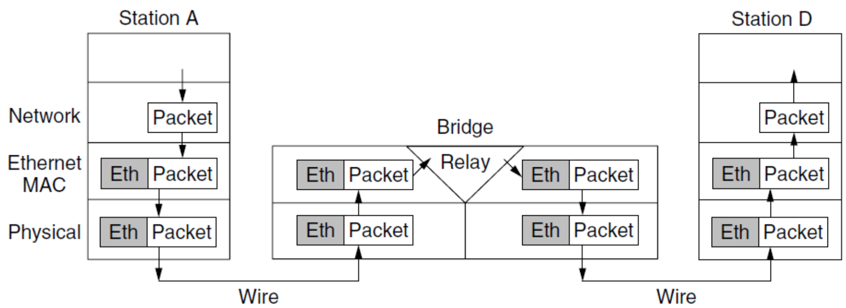

Adding (Un)Reliability

Redundant links

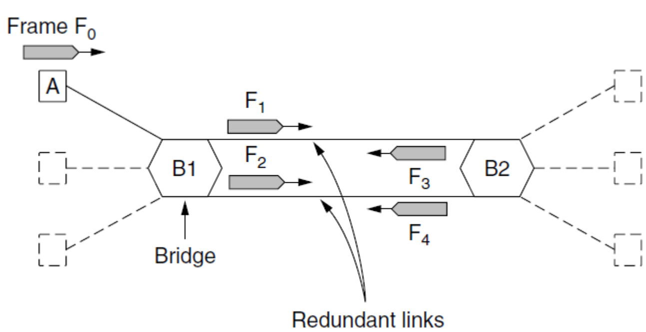

Okay, just don't do that…

Spanning Tree Algorithm

- Goal: prevent forwarding loops, even when there are physical loops, by having only a single bridge responsible for each LAN segment
  - "turn off" many bridge ports
  - Do this dynamically
- Results in a spanning tree for packet routing
  - Robust against failures, needs no organization

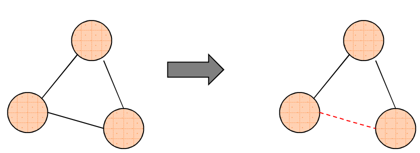

Llan.12

Algorithm Overview

1. Elect a root node for the tree
   - Remember that all bridges are doing this at once, and none know what the network topology is
2. Grow tree as shortest distances from the root
   - Remember that all bridges…
3. Add some details to actually make it work (!)

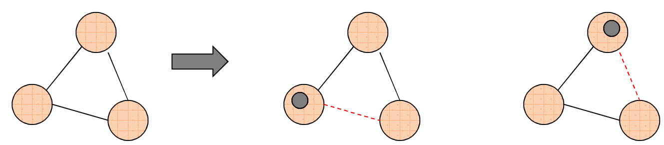

Algorithm continued

- Each bridge sends periodic messages to others containing:
  – [its address,
    address of the bridge it believes is the root,
    distance (in hops) to root]

- A bridge receiving a message may update its view of the "best" bridge on the LAN segment it received it from
  – First criterion: Smaller root address is preferred
  – Second criterion: Smaller distance to root is pre
  – Tie breaker: bridge with smaller address is better

Algorithm continued

- Initial condition: each bridge thinks it is the root
  - Sends configuration messages on all ports
    - [I'm me,
      the root is me,
      I'm at distance 0 from the root]
- Later, bridges send only "best" configs
  - Add 1 to distance, send configs on segments where it's still "best" (designated bridge)
  - Turn off forwarding on ports except those that send/receive "best"

Llan.15

Spanning Tree Example: Final Tree

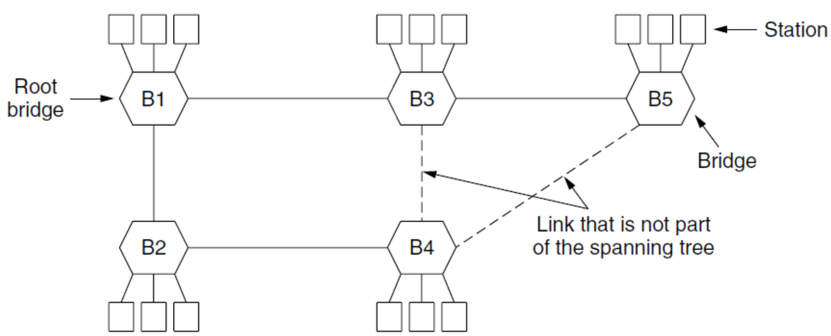

Spanning Tree Example: Transients

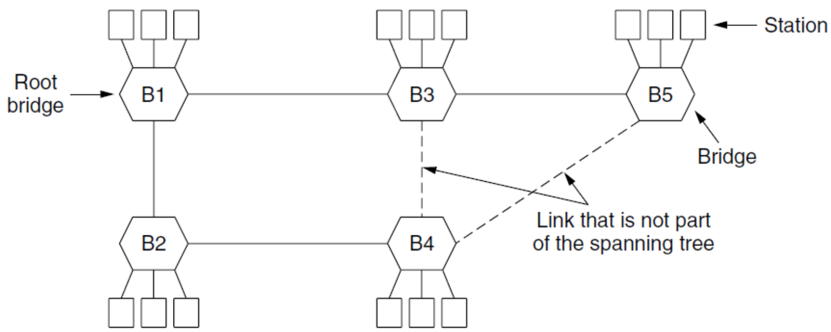

General Design Principles

- All bridges to run the same algorithm
- All bridges start with zero information and operate in parallel
- Bridges send periodic messages about their own state
- State that hasn' t been refreshed in a while is deleted (*soft state*)
- The algorithm converges to a globally consistent state
  - If we all have the same inputs and have implemented it correctly and the inputs aren't changing too faster than we can converge

LAN Switches – what is used today

- LAN switches are multi-port bridges
  - Modern, high performance form of bridged LANs
  - Frames are switched, not shared like hubs
  - Every host on a separate port, can combine switches

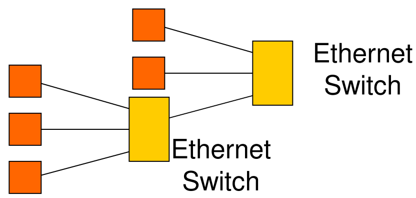

.

Virtual LANs

VLANs (Virtual LANs) split one physical LAN into multiple logical LANs, for management reasons

- Ports are "colored" according to their VLAN

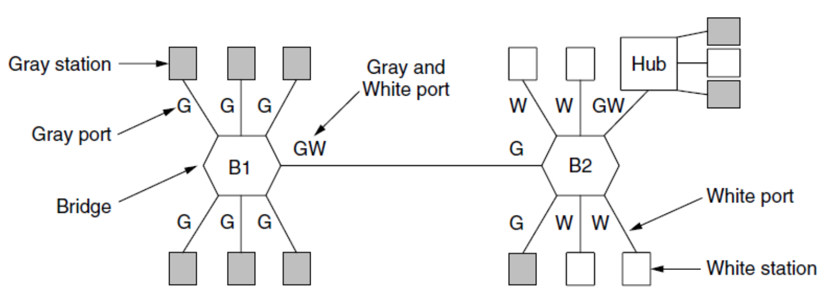

Virtual LANs -- IEEE 802.1Q

802.1Q frames carry a color tag (VLAN identifier)

- Length/Type value is 0x8100 for VLAN protocol

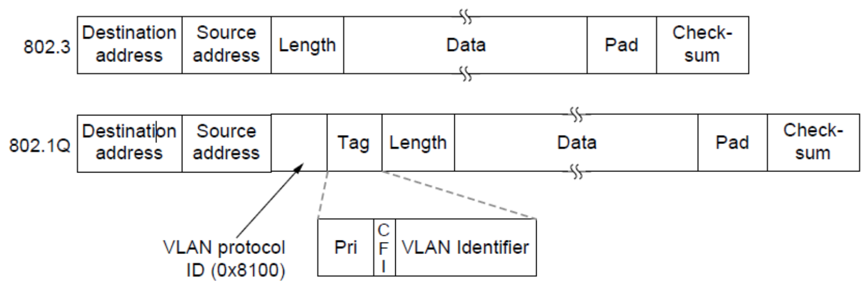

VLAN Breaking News

The CSE wireless network used to be supported by CSE owned APs.

UW-IT has replaced our APs with theirs.

We'll still run an isolated CSE wireless network, and they'll run an isolated UW wireless network, both on their APs.

VLANs…

Perspective on Bridges/Switches

- LAN switches form an effective small-scale network
  - Designed to work transparently with no host changes
  - Plug and play for real!
- But they don't scale well to large networks
  - 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
- We can layer on scalability by changing hosts too.
  - Next up: IP and routing.