

Lectures 19, 20 Bridges and switches

Brian Carpenter

COMPSCI 314 S2C 2011

COMPSCI 314

Repeaters to Routing

- How do we connect devices

- In the same room?

- In the next room?

- In the next building?

- In the next department

- In the next campus?

- In the next country?

- In the next continent?

Cabling limit

+ Admin boundary

+ Funding boundary

+ Political boundary

Distance, cost and numbers increase; must share costs and cables

COMPSCI 314

2

Topics (Shay 10.1 - 10.4)


- Repeaters
- Hubs
- Bridges
 - Spanning tree algorithm
- Ethernet switches
 - Virtual LANs
- The need for routing

COMPSCI 314

3

Repeaters

- A repeater simply reshapes, amplifies and sends on the signal
 - Useful when cable length is enough to degrade the signal, or to safely connect two cable segments




- An Ethernet repeater works in both directions

COMPSCI 314

4

Hubs



- A hub is (in the simplest case) simply a multiport repeater.
- A signal coming in on any port is reshaped, amplified, and sent out on all the other ports
- But *hub* is a very general word, so people sometimes speak of *bridging* or *switching* or even *routing* hubs.
- However, if it costs \$20 it's only a repeating hub.

OK, maybe \$27.81...



This handy USB cup warmer plugs into your computer to keep your drink warm. Features include 4 port USB hub, when connected to computer the LCD blue back light is activated, clock, alarm and temperature.

from:
\$27.81 per unit

This handy USB cup warmer plugs into your computer to keep your drink warm. Features include 4 port USB hub

Limits on hubs and repeaters


- Hubs and repeaters are essentially cable extenders.
 - They add slightly to transmission delay.
- In Ethernet, they form part of the collision domain calculation, as we saw earlier.

Advantages and disadvantages of hubs and repeaters

- + Nothing to configure.
- + Cheap and reliable.
- + Regenerate degraded signals.
- Don't avoid Ethernet contention/saturation problems on shared LANs.
- Don't split up the collision domain, therefore subject to distance limits.

Bridges

- Although bridges are little used today, they are important to understand before we discuss switches.
- Consider a situation where we need to connect two Ethernets, but using a simple repeater would exceed the collision domain limit:



- What must we do to avoid joining the two collision domains into one?

COMPSCI 314

9

Bridges (2)


- The trick is to *read* the packet off one Ethernet and then *send* it on the other one quite independently.
 - The bridge pretends to be the destination on network 1 and then pretends to be the source on network 2.
 - Reading a packet on Ethernet is a passive operation (the receiver sends absolutely nothing).
 - When sending the packet out again on network 2, the bridge uses the same source MAC address as the original sender on network 1, instead of its own MAC address.
 - Any station that sees the re-sent packet cannot tell that it came from a bridge.
 - Will any bits be different?
 - Does the bridge need to re-calculate the FCS?

COMPSCI 314

10

Bridges (3)

- How does the bridge know when to resend a packet?
 - In the simple case of two Ethernets, it could just resend everything it sees ("flooding").
 - But with several Ethernets and bridges, that gets too noisy, since every packet will be sent through every bridge to every other Ethernet.



COMPSCI 314


11

Bridges (4)

- Conclusion: bridges must know where each Ethernet station is, and only send packets in that direction.

- Each bridge needs a routing table. For Bridge 1:

From North		From South	
Dest. addr.	Send to:	Dest. addr.	Send to:
A	(Local)	A	North
B	South	B	(Local)
C	South	C	(Local)



'North' and 'South' will be port numbers in a real case.

- No need to send packets again locally

COMPSCI 314

12


How bridges learn

- At start-up, a bridge knows nothing (routing table is empty).
- Suppose it sees a packet on its North port, with destination B and source A.
- It now knows that A is to the north, so it can insert
A (Local) A North
- It still knows nothing of B, so the only choice is 'flooding': re-send the packet on its South port. All bridges do the same.
 - If B was in fact North, the re-sent packet would be wasted.
- If it sees a packet with destination C and source A, it learns nothing new and continues to flood.
- When a reply comes from B to A on its South port, it inserts
B South B (Local)
- Try exercise 2 on Shay page 521. In a real example, the (Local) entries are not needed.

COMPSCI 314

13

Making the problem a little harder



- How do we prevent packets being sent twice or even in loops, when A sends to C?
 - Flooding and learning are not enough: both B2 and B3 will learn that C is to the south.
 - B2 must ignore packets for C on its North port.
 - The bridges must talk to each other, to eliminate duplicate paths.

COMPSCI 314

14


The spanning tree algorithm in brief

- To avoid loops and duplicates we need a map of some kind. Bridges use a *spanning tree*.
 - In graph theory, a spanning tree is a set of edges of a graph that form a tree and includes every vertex.
 - The spanning tree protocol is standardised as IEEE 802.1D
- Bridges send management packets to each other.
- One bridge is "elected" as the *root*.
- Each bridge determines which port is closer to the root.
- The bridges on each LAN elect one bridge as the *designated bridge* which will send frames towards the root. Again, this is based on distance from the root.

COMPSCI 314

15

The spanning tree result



- Bridge 5 is idle, but there are no duplicates or loops.
 - Some shortest paths would use Bridge 5.
 - The spanning tree does *not* optimise paths.
- If Bridge 3 fails, the spanning tree will reconfigure and Bridge 5 will take over.
 - Routing tables and spanning tree must be timed out and refreshed periodically to deal with changes in the network

COMPSCI 314

16

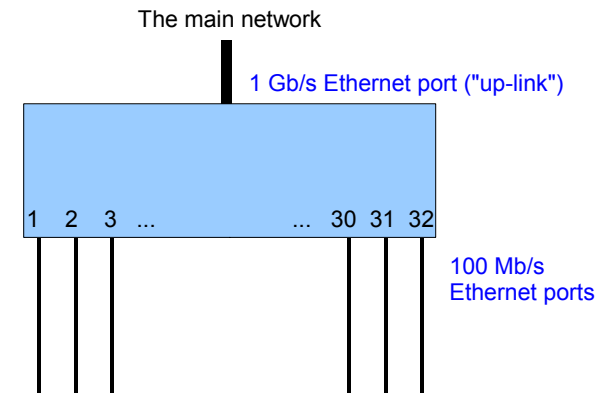
Advantages and disadvantages of bridges

- + Self configuring. (With spanning tree, any old network design will work - the easy way to split up a big LAN.)
- + Fairly cheap and reliable.
- + Overcome distance limit, regenerate degraded frames.
- Don't cure Ethernet saturation on busy LANs.
- Suit old-fashioned shared cable better than star cabling.
- Notorious for allowing “broadcast storms” despite the spanning tree. (All bridges must forward multicast and broadcast packets.)
- Generally hard to operate and manage large bridged networks.

COMPSCI 314

17

Switches




- Functionally, this switch will behave exactly like a 33 port bridge.
 - Routing table with 33 port numbers instead of 2
 - Spanning tree when several switches are connected together

COMPSCI 314

18

Switch internals

- Switches vary widely in internal design (and price)
- The ideal switch can run full duplex at line speed on every port.
 - In our 33 port example, that would mean 32 100 Mb/s flows in and out. 10 each way could be using the 1 Gb/s port.



COMPSCI 314

19

Read the small print

- You can't tell from looking if a device is a switch or a hub.
- You can't tell from looking if a switch can run full speed on all ports, or if its internal design limits it to N frames per second or N simultaneous frames.
- In any case, 32x100 won't go into 1000. What happens if everyone sends to the “up-link” at the same time?
 - Does the switch queue excess packets, discard them, cause a collision, or apply some kind of flow control?

COMPSCI 314

20

Example of small print

- 24 RJ-45, 10/100 Mbps Ethernet LAN ports
- Maximum Bandwidth:
 - Full duplex: 200 Mbps (for 100BASE-TX), 20 Mbps (for 10BASE-T)
 - Half duplex: 100 Mbps (for 100BASE-TX), 10 Mbps (for 10BASE-T)
- Operates at maximum forwarding rate and provides frame filtering and forwarding functions for each port
 - 14,880 packets per second for 10 Mbps ports
 - 148,800 packets per second for 100 Mbps ports
- Store-and-forward scheme to forward packets
- MAC address table: 4096 entries
- Packet buffer: 256 KB
- IEEE 802.3x Flow Control
 - Half-duplex: Back pressure
 - Full-duplex: Pause frame

USRobotics:
"24 Port 10/100 Mbps
Switch" (around 2007)

COMPSCI 314

21

Running the numbers

- "14,880 packets per second for 10 Mb/s ports"
- That means $10,000,000 / 14,880$ bits per packet
- That's 672 bits per packet, or 84 bytes per packet.
- Ethernet packets are up to 1500 bytes, and more typically ~500 bytes.
- If we assume an average of 500 bytes/packet, that's 4000 bits/packet, which at 10 Mb/s means $10,000,000 / 4000$ packets/s
- That's about 2500 packets per second, well below the switch capacity.
- "Packet buffer: 256 KB"
- That will hold about 500 typical packets. Suppose ten ports all fire packets at full speed towards the same port. 500 packets will arrive in 1/50 of a second (20 msec) but will take 200 msec to send on.
- Is the packet buffer big enough?
- In general, whether a switch is suitable depends on the expected traffic load. Obviously, switches with greater throughput cost more money.

COMPSCI 314

22


Advantages and disadvantages of switches

- + Self configuring.
- + Reliable.
- + Overcome distance limit, regenerate degraded frames.
- + Avoid Ethernet contention/saturation problems on shared LANs or simple hubs.
- + Suit modern UTP5 (or optical fibre) cabling.
- Require systematic design and cable management.

COMPSCI 314

23

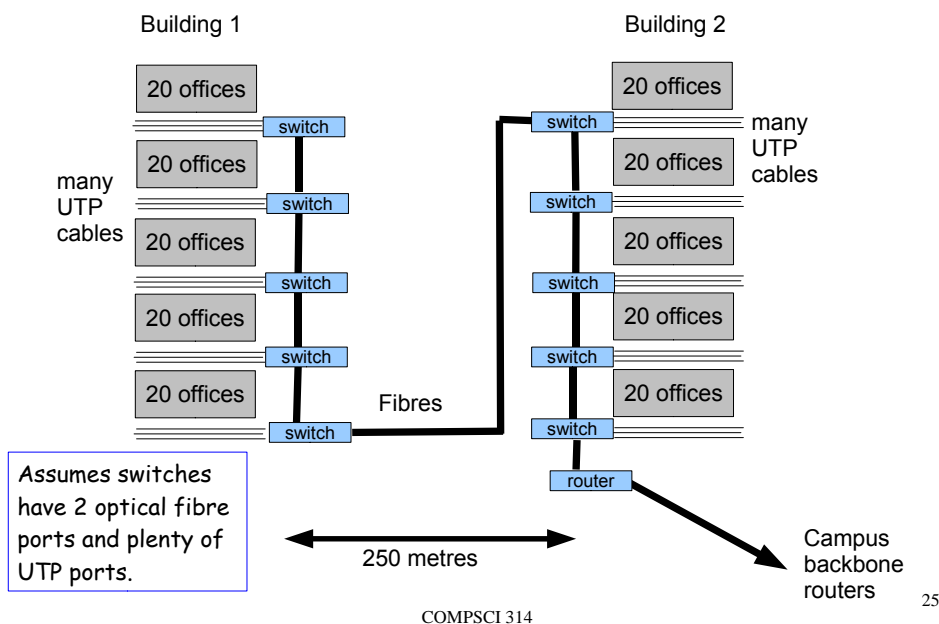
Network for a large office or lab



COMPSCI 314

24

Two buildings



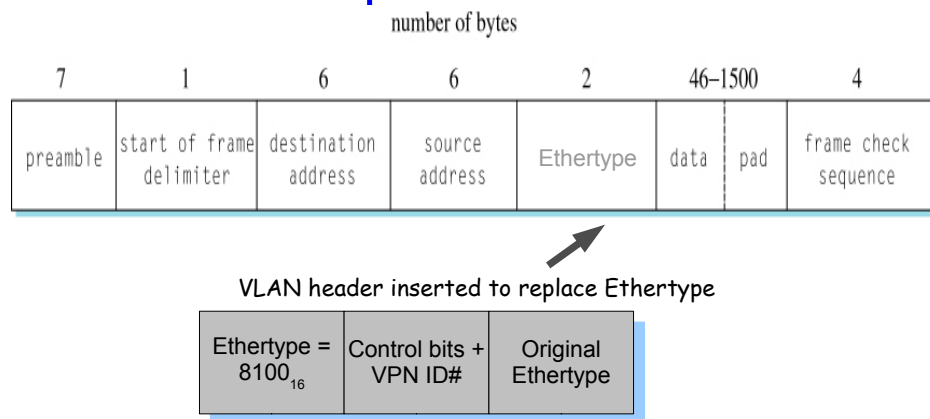
Virtual LANs (VLANs)

- A technique for mixing two or more usages of Ethernet on one physical infrastructure (cables and switches).
- Imagine inserting an extra byte at the beginning of every Ethernet frame that contains "1" or "2".
 - User ports configured in VLAN #1 will send frames with "1" and can only receive frames with "1"
 - Ditto for VLAN #2
 - Switches handle both VLANs
 - Magic! We now have two LANs on one set of cables and switches
- Actual VLANs are a bit more complex than this, as defined in the IEEE 802.1q standard.

COMPSCI 314

26

802.1q VLAN header



- Adds 4 bytes (included in FCS)
- Hidden from software

COMPSCI 314

27

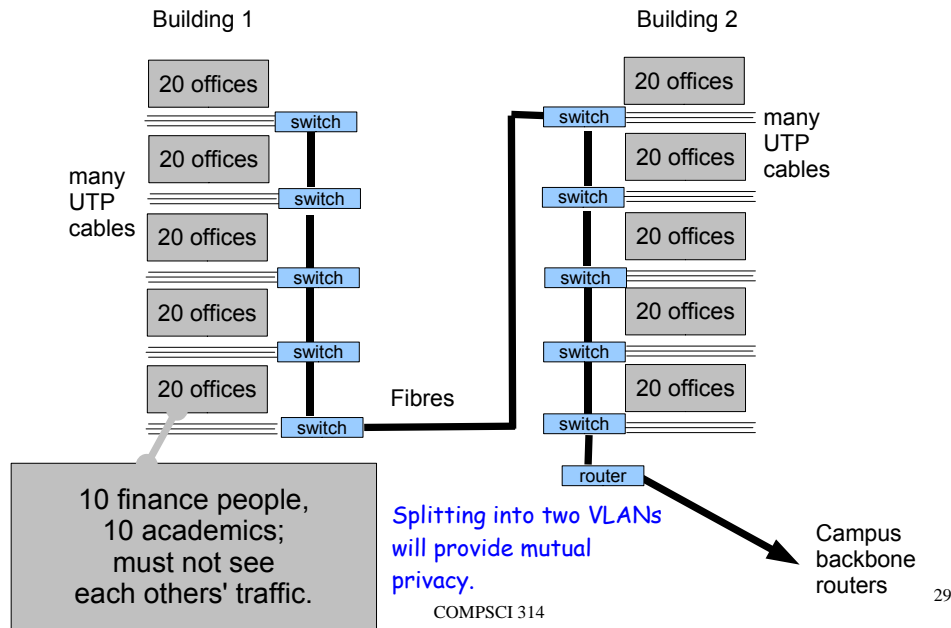
When to use VLANs

- Two populations of users mixed in the same area
 - if there are serious data privacy concerns
 - if one or both is generating a lot of multicast traffic
- Two very different types of traffic in the same area that can easily be separated
 - e.g. Voice over IP versus general usage
- Exceptionally, if there's a need to "stretch" a particular special LAN between buildings, and there no spare cable or equipment
- But always ask whether virtual traffic separation is really going to be an advantage
 - The data are still sharing the same cables and switches

COMPSCI 314

28

Two buildings, two populations



Repeaters to Routing

• How do we connect devices

- In the same room?
- In the next room?
- In the next building?
- In the next department?
- In the next campus?
- In the next country?
- In the next continent?

Physical addressing:
Hubs, switches

Cabling limit

+ Admin boundary

+ Funding boundary

Logical addressing:
routers

+ Political boundary

Distance, cost and numbers increase; must share costs and cables, but separate administration