A Brief Introduction to Ethernet

- A common medium for IP LANs - cheap, fast and easy to build
- Actually a collection of related standards
- The more common ones listed here
- Remember, at layer 1 we are concerned with wiring rules and bits per second only

1

Ethernet layer 1

- 10baseT
  - Connectors: RJ45 (8-pin)
  - Cable: Category 3 or 5 twisted pair
  - Topology: point to point (e.g. device to hub)
  - Maximum cable length: 100m
  - Maximum of 4 hubs between any two devices
  - Speed: 10 megabits per second
  - Although cable has 4 pairs, uses only 2 of them (one for transmit, one for receive)
  - Important to wire so that the correct pairs are twisted together

Ethernet layer 1 (contd)

- 10base2 a.k.a. "thinnet"
  - Connectors: BNC
  - Cable: 50 ohm coaxial cable
  - Topology: single bus with taps
  - Maximum length: 185m from end to end
  - Speed: 10 megabits per second
- 10base5 a.k.a. "thicknet"
  - Heavier cable
  - Maximum length: 500m from end to end
- Both are obsolete. Major problem was that a single device fault would bring down the whole network!

3

Ethernet layer 1 (contd)

- 100baseTX
  - Connectors: RJ45
  - Cable: Category 5 twisted pair
  - Topology: point to point (device to hub)
  - Maximum length: 100m from device to hub
  - Maximum of 4 hubs between any two devices
  - Speed: 100 megabits per second
- 1000baseTX
  - Same connectors and cable, but manages to achieve 1000 megabits per second!
  - Uses all four pairs
  - No hubs, only switches (see later)

Ethernet layer 1 (contd)

- 100baseFX
  - 100 megabits over multimode fibre
- 1000baseSX
  - 1000 megabits over multimode fibre (short haul; 220-550m depending on fibre type)
- 1000baseLX
  - 1000 megabits over single mode fibre (long haul, 5km or more; expensive!)
- Fibre and radio are the ONLY options for linking between buildings, for safety reasons

5

What cabling should you install?

- Within one building, the choice is easy: install CAT5 !
- Works at 10, 100 or 1000Mbps: futureproof
- There is no such thing as "CAT6", "CAT7" etc. Don't be fooled by vendors!
- Make sure your connectors (e.g. patch panels, wall sockets) are CAT5 certified too

Ethernet is a broadcast medium

- Everything sent on the wire is seen by all stations

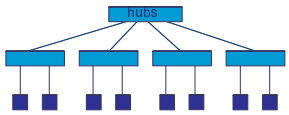

is equivalent to

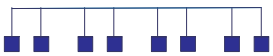

Layer 2: Ethernet sends data in frames

- "MAC frames" (Media Access Control)

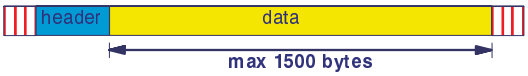

- There is collision detection to prevent two machines transmitting a frame at the same time. "Carrier Sense Multiple Access, Collision Detection" (CSMA/CD)
- "Listen before speaking!"

Each ethernet card has a unique MAC address

- A 48-bit binary number, usually represented in hexadecimal
  - e.g. 00:04:AC:06:13:E2
- Allocated to the card by the manufacturer
  - each manufacturer has a unique range, and every card has a different MAC address
- The frame header contains a source and destination MAC address - i.e. who sent it, and who is intended to receive it

Receivers filter on MAC address

- So they don't need to process frames which are intended for another machine

Note: This is not secure!! User can set "promiscuous mode" to see other people's data

Broadcast Address

FF:FF:FF:FF:FF:FF - "all stations"

Any questions?

- Who has crimped their own RJ45 cables?

How is IP (layer 3) carried over Ethernet (layer 2)?

- Easy: you just put the whole IP datagram in the data portion of an Ethernet frame. This is called "encapsulation"

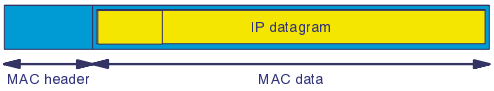

What is the MTU for IP datagrams sent over Ethernet?

Just one problem...

- If you are sending to IP address 1.2.3.4, how do you know what the MAC address of the recipient is?

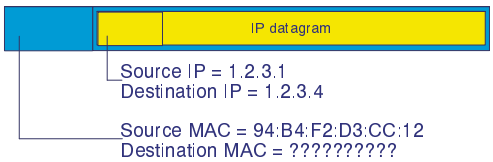

The solution: ARP

- "Address Resolution Protocol"
- Send a broadcast frame to all stations, saying "who has address 1.2.3.4?"
- The machine with that address responds
- The source MAC address of the response tells you the MAC address which corresponds with 1.2.3.4

ARP Example

- Broadcast ARP request

- Machine with that IP responds

- Now send the IP datagram

ARP Cache

- For efficiency, each machine caches the ARP responses it sees
- Entries time out after a few minutes/hours (in case the information changes - e.g. you put a new network card in a device)

1.2.3.1	0A:B3:22:46:C1:99
1.2.3.2	2B:9D:A7:BD:E1:D1
...	...

Note!

- ARP packets are NOT IP datagrams!
- ARP is a completely separate protocol to IP
- ARP packets are never forwarded outside of the local network
- ARP is transparent; you rarely have to mess with it

Watching ARP in action

- Short practical exercise

19

Different ways to build networks

- Use the OSI model to understand your options

Growing your network at layer 1

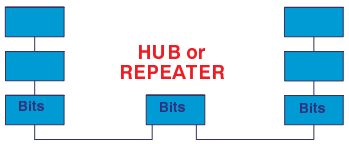

21

What's a hub (repeater)?

- A dumb device which receives a stream of bits on one port, amplifies and resends them on all the other ports
- Receives and sends one bit at a time
  - So a hub could interconnect 10base2 and 10baseT networks, but not 10baseT and 100baseTX
- Ethernet hubs can identify faulty devices (e.g. devices which are transmitting when they should not) and isolate them to protect the rest of the network

Advantages of growing your network at layer 1

- Hubs are simple, cheap, "plug and play"

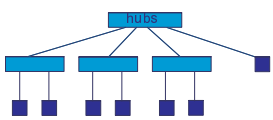

23

Disadvantages

- You can only have 4 hubs between any two machines, which limits the physical size of your network
  - Due to timings of collision detection system
- Can only connect two identical network technologies (e.g. 10M ethernet to 10M ethernet)

More disadvantages

- Only one machine can "talk" at any time; the total available network bandwidth is shared between all the machines
- Because of collision detection, the usable bandwidth on a 10baseT network is about 4Mbps
- Works well for a few tens of machines

Growing your network at layer 2

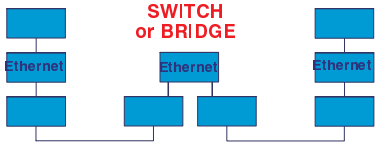

What's a switch (bridge)?

- A device which reads whole ethernet frames, stores and forwards them
- Learns which MAC addresses are connected to which port
  - by monitoring source MAC addresses
- Frames addressed to a particular machine are only sent on the port where that machine is connected
  - Unlike a hub, which sends data on all ports
  - However, broadcast frames are still sent out on every port

Advantages of building your network at layer 2

- Switches are simple drop-in replacements for hubs, and now almost as cheap
- More advanced switches allow remote management (interface status, counters etc)
- Can mix-and-match 10M ethernet, 100M ethernet and gigabit ethernet

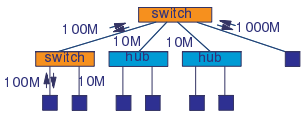

Advantages of switches

- A switch can receive packets on multiple ports simultaneously
  - Machine A can be sending to machine B at the *same time* that machine C is sending to machine D
  - More total bandwidth available
- Individual hosts can transmit and receive simultaneously
  - Known as "full duplex" mode; needs to be agreed by both sides (can auto-negotiate)
  - Eliminates collisions entirely

Disadvantages of switches

- Broadcast packets still appear on every port. Eventually broadcast noise is a problem.
  - e.g. Windows machines send at least one broadcast packet every 30 seconds
- Switches have to build a table with the MAC address of every device on your network. Typical limit is one or two thousand entries.
- Difficult to build loops or multiple paths into your network for redundancy
  - can be done but adds a lot of complexity: e.g. Spanning Tree

Disadvantages of switches

- There is no layer 2 ping or traceroute. With a complex network of switches, debugging becomes a problem
- No isolation
  - Broadcast storms affect everyone
  - One user can 'steal' another user's IP address, accidentally or maliciously
- Auto-detection of half/full-duplex does not always work
  - Cisco devices are particularly bad
- Works well up to a few hundred devices

Growing your network at layer 3

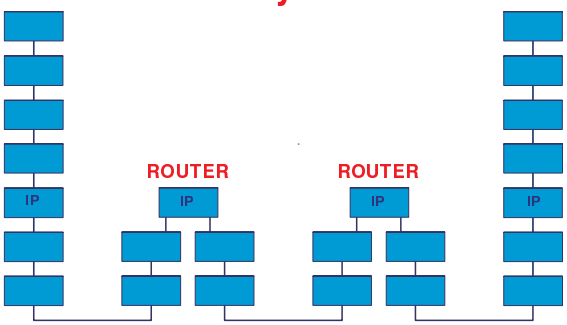

What is a router?

- A host with two or more interfaces *and* which has been configured to forward datagrams *
- Works at layer 3: receives datagrams and forwards them based on the destination IP address
- For a full definition see RFC 1812
- Older documents call it a "Gateway" but this term is normally used for a layer 7 gateway

* Otherwise it is just a multi-homed host

What is a router? (Contd)

- A router is connected to two or more networks
- The Golden Rules say that every network must have its own prefix
- Therefore a router must have two or more IP addresses - one for each interface

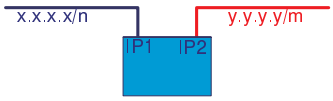

What is forwarding?

- Receiving a datagram on one interface and resending it on another

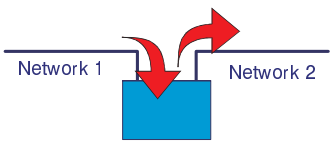

- Why? Because Layer 2 can only send to direct neighbours (on the same link)

What happens when a router receives a datagram?

- First it checks to see if the destination address is local (i.e. the router itself is the final destination)
- Next it decrements the Time To Live (TTL) field in the IP header. If it reaches zero, the datagram is discarded
- Finally it looks up the destination in a *forwarding table* to decide where to send it next
- Could be on a directly-connected network, or could have to send it via another router

Forwarding is hop-by-hop

- Each router can only communicate directly with devices which are on the same networks that it is connected to (Layer 2)
- Each router gets the datagram one hop closer to the destination
- Each router makes an *independent* decision as to the next hop (the route is not pre-planned)
- Each router has a different view of the world so has a different forwarding table

Using prefixes for forwarding

- We don't list every IP number on the Internet - the table would be huge
- Instead, the forwarding table contains prefixes (network numbers)
  - "If the first /n bits matches this entry, send the datagram this way"
  - If more than one prefix matches, the longest prefix wins (more specific route)
  - 0.0.0.0/0 is "default route" - matches any IP address, but only if no other prefix matches
- This is why layer 3 is so scalable

Advantages of routers

- Separates your network into "broadcast domains" - broadcasts do not propagate between them
- Isolates faults - a badly configured machine on one LAN does not affect machines on another
- Globally scalable - works for millions of machines
- Highly robust; can have backup links
- Exchange traffic between different media (e.g. ethernet to PPP dialup)

Disadvantages of routers

- Routed networks are not plug-and-play
  - You need to allocate IP addresses to each network
  - You need to build forwarding tables by hand or configure routing protocols
- Hardware routers are expensive
  - depending on what speed and what sorts of interfaces you need
- PCs running Unix make acceptable routers
  - but not as reliable and harder to configure

Other options for interconnecting networks

- At layer 4: proxies
  - Most common example is SOCKS proxy
- At layer 7: application gateways
  - They understand the application protocol and apply processing to it
  - Example: Squid web cache (HTTP gateway; client requests pages and Squid caches them for future clients)
  - Example: Gateway SMTP to X400 mail

Recommendations

- Build small LANs with routers in between
- Keep different types of users on different LANs (zones of trust)
  - At a university: students on one group of LANs, staff on another, perhaps servers on another
  - At an ISP: infrastructure on one LAN, staff on another, customer hosted servers on another
- Helps your network scale, helps security, helps minimise faults and makes them easier to trace