COMP2303/COMP7306 Network and OS Principles

Week 12 - Friday

Internet & Ethernet

School of Information Technology and Electrical Engineering
The University of Queensland

Outline

- Ethernet
  - Hubs/Bridges/Switches/Routers
  - Ethernet Frames and MAC addresses
- IP Header fields (revisited)
- IP addresses

- Tuesday
  - Remainder of material
  - TEVAL-Joel
- Credits:
  - Tanenbaum, "Computer Networks"
  - Bryant and O'Halloran, "Computer Systems: A Programmer's Perspective"

IEEE 802 LAN Standards

- Several standards for LANs/MANs
  - Ethernet (802.3)
  - WLAN (802.11)
  - Bluetooth (802.15)
  - Many others

Ethernet

- Name "Ethernet" now refers to a family of networking standards
- Ethernet is a CSMA/CD protocol
  - **CS** = Carrier Sense
    - Nodes can sense transmissions currently taking place (check no one else is sending before transmitting)
  - **MA** = Multiple Access
    - Multiple nodes use the channel
  - **CD** = Collision Detect
    - Nodes can detect collision (multiple nodes transmitting) as soon as it happens and can abort transmission immediately

---

802.3 - History

- Early 70's – Xerox PARC, Metcalfe and Boggs built a 2.94Mbps CSMA/CD network
  - 100 workstations
  - 1km cable
  - called "Ethernet"
  - published 1976
- 1979-1982 – DEC, Intel, Xerox standardised 10Mbps "thick-wire" Ethernet
  - later known as 10base5
  - "base" refers to baseband transmission
- 1983 – IEEE defined 802.3 protocol
  - very similar to Ethernet

802.3 Standards

Include

- 802.3 (1983)– **10Base-5** – original Ethernet
  - Thick coaxial cable
- 802.3a (1985) – **10Base-2** – thin Ethernet
  - "Cheapernet" – thinner cheaper coax.
- 802.3i (1990) – **10Base-T** – 10Mbps over Cat 3 UTP or Cat 5 UTP
- 802.3u (1995) – **100Base-T**
  - "Fast Ethernet"
  - several media supported (Cat-5 UTP, Optical Fibre)
- 802.3z (1998) – **1000Base-X**
  - "Gigabit Ethernet"
  - several media supported
- 802.3ab (1999) – **1000Base-T**
  - Gigabit Ethernet over UTP
- 802.3ae (2002) – **10GBase-...** - 10-Gigabit Ethernet
- Faster Ethernet under development

Ethernet Collision Domains

- Only one ethernet host on a segment may transmit at a time
- Otherwise, frames **collide**
  - Collisions mean frames unintelligible to receivers (network cards)
  - Ethernet uses binary exponential backoff when collision occurs
- An ethernet segment forms a **collision domain** – only one host may transmit
- More hosts ➔ more chance of collision

7

Binary Exponential Backoff

- What happens after a collision?
- Time divided into slots
  - Slot – 64 bytes wide
    - Allows for worst case propagation delay
  - Wait 0 or 1 slots
- After second collision
  - Wait 0, 1, 2 or 3 slots
- After i'th collision – wait 0 … $2^i - 1$ slots
  - Maximum 1023
- After 16 collisions
  - Give up

8

Ethernet – with hubs

- Ethernet segment consists of a collection of hosts connected by wires (twisted pairs) to a **hub**
- Spans room or floor in a building

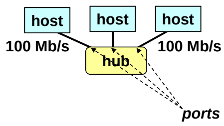

- Operation
  - Each Ethernet adapter (network card) has a unique 48-bit address (**Ethernet address** or **MAC address**)
  - Hosts send bits to any other host in chunks called **frames**.
  - Hub copies each bit from each port to every other port
    - Every host sees every bit
    - Network cards filter out frames not for that host

9

Bridges

- Bridges cleverly learn which hosts are reachable from which ports and then *selectively* copy frames from port to port
  - bridges learn MAC (Ethernet) addresses
  - e.g. traffic from A to B never makes it here

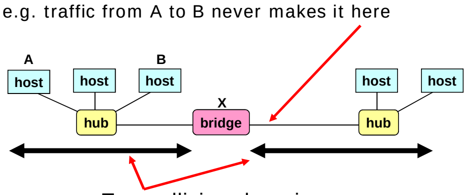

Two collision domains
• improved network capacity

10

Switches

- Bridge – usually connects LANs
- **Switches** – like bridges, but connect individual computers and filter frames based on frame addresses
  - Each port is its own collision domain
    - Frames never lost to collisions
  - Network has much greater capacity
    - But switch can discard frames if overloaded
- Example – A can communicate with B whilst C communicates with D

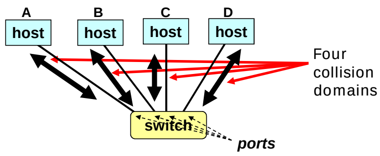

11

Conceptual View of LANs

- For simplicity, hubs, bridges, switches and wires for a single LAN are often shown as a collection of hosts attached to a single wire:

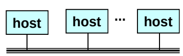

- Even if drawn like this, note that most modern "Ethernets" are switch based – there isn't a medium shared between all hosts

12

Next Level: internets

- Multiple incompatible LANs can be physically connected by specialized computers called *routers*
- The connected networks are called an *internet* (lower case i)

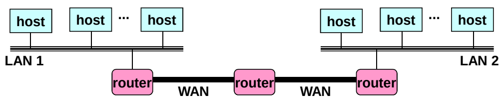

- LAN 1 and LAN 2 might be completely different, totally incompatible LANs (e.g., Ethernet and ATM)
  - could be the same too

13

Short Break

- Stand up and stretch

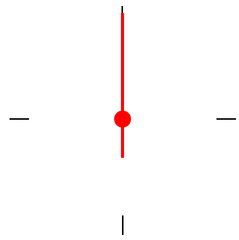

14

Ethernet Frames

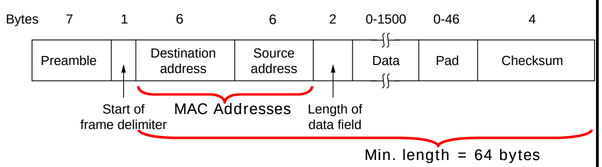

- Length
  - Number of *bytes* in data field
  - 0-1500
- Padding
  - So that minimum size of frame = 64 bytes

15

MAC Addresses

- 6 byte (48 bit) MAC address
  - Often expressed as a 12 hexadecimal digits
    - e.g. 00-14-22-BC-91-72
  - Addresses are globally unique, i.e. should correspond to exactly one network interface card (NIC)
  - Each manufacturer is assigned a block of MAC addresses
- Given a MAC address, you can find the manufacturer of the NIC: http://coffer.com/mac_find/
- Some cards allow the MAC address to be changed via software!!!
  - This can circumvent some security mechanisms that are based on MAC addresses

16

Transferring Data Over an internet

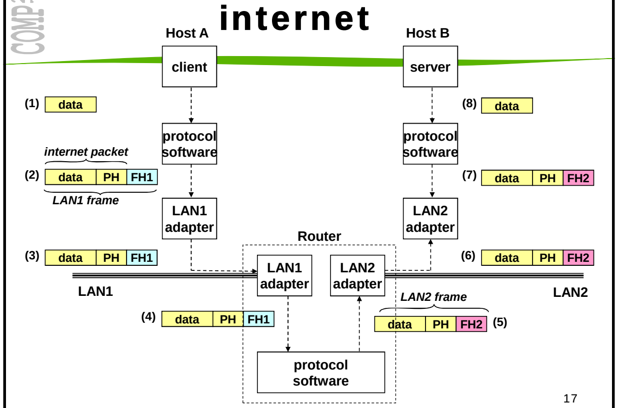

17

Internet Protocol (IP) Header (from Week 9)

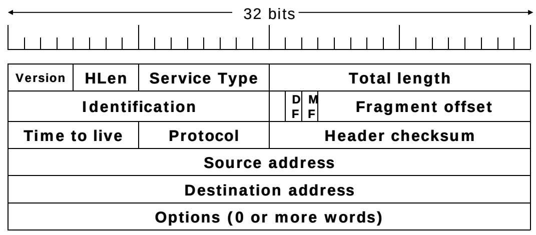

- We'll look at a few more header fields today

18

IP Header: Version

- 4 bits
- Specifies protocol version
- Allows protocol versions to change
- Currently v4
  - v6 around also (uses different header)

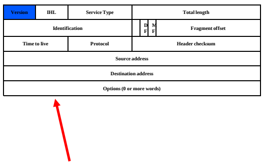

You're not expected to read these words. The figure just indicates which part of the header we're talking about (highlighted). (Refer to previous slide.)

19

IP Header: Total Length

- 16 bits
- Specifies total datagram length
  - includes header and data
- Maximum length = 2^{16}-1 = 65535 bytes
- BUT, in reality, datagrams are often limited by the underlying physical network, e.g. Ethernet

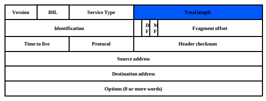

20

Question

- What's the largest amount of **data** that can be sent in an IP datagram?

21

Fragmentation

- IP designed to work over many network technologies
- Different frame sizes
  - Ethernet – up to 1500 bytes
  - FDDI – up to 4500 bytes
  - (MTU = Maximum Transfer Unit)
- Sometimes need **fragmentation** of packets
  - Reassembly happens at destination
  - Discard incomplete packets after timeout
- IP requires all machines/networks to accept fragments of 576 bytes or less

22

IP Header: Protocol

- 8 bits
- Range: 0 to 255
- Identifies higher level protocol to which packet should be passed, e.g.
  - 6 = TCP
  - 17 = UDP

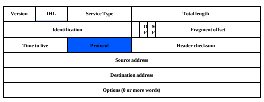

- RFC 1700 – defines protocol numbers

23

IP Header: Time-to-live

- 8 bits
- Range: 0 to 255
- Limits packet lifetimes
- Decremented at each router
  - Was supposed to count seconds
  - In practice – counts hops
- Packet discarded when reaches zero

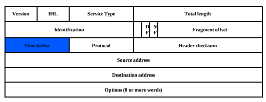

24

Time-to-live

- Why limit packet lifetimes?

25

Aside – traceroute

- **traceroute** uses the ICMP protocol (built on IP) with varying Time-to-live values (starting from 1)
- The expiry messages returned allow the route through the network to be determined

26

Break

27

IP Header: Source and Dest. Addresses

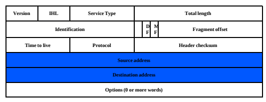

- 32 bits each
- Dest. address
  - Used to route packet to intended destination
- Source address
  - Destination can choose whether to receive
  - Destination knows whom to reply to

28

IP Addresses – classful addressing

- Each host (network interface) has a unique 32 bit address
- Hierarchical format: (network number plus host number)

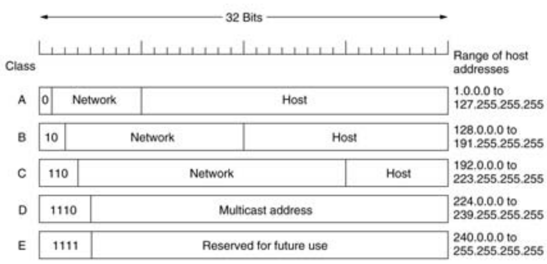

Why?

- Why are IP addresses hierarchical?
- Why don't we just use Ethernet (MAC) addresses?

30

IP Address Classes

- **A** – up to 126 (2^7-2) networks
  - 16 million (2^{24}-2) hosts each
- **B** – up to 16382 (2^{14}-2) networks
  - 64k (2^{16}-2) hosts each
- **C** – up to 2 million (2^{21}-2) networks
  - 254 hosts each

31

Dotted Notation

- IP addresses written in **dotted-decimal** notation
  - each of 4 bytes written in decimal
    - e.g. 130.102.2.15
- Addresses with all 0's and all 1's have special meanings
  - 0 = this network or host
  - 1 = broadcast – all hosts on indicated network

32

Special IP addresses

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		This host
0 0 . . . 0 0	Host	A host on this network
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1		Broadcast on the local network
Network	1 1 1 1 . . . 1 1 1 1	Broadcast on a distant network
127	(Anything)	Loopback

33

Subnets

- **Subnet** – split a network into parts
- Still appears as single network to outside world
- Example – class B network:

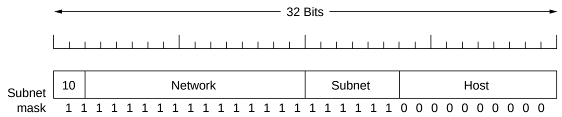

- Doesn't have to happen on 8-bit boundary
- This example – 6 bit subnet ID, 10 bit host ID - allows 62 LANs, each with up to 1022 hosts

34

Forwarding / Routing

- Routers have (incomplete) listings of
  - network-num.0
  - this-network-num.host-num

  addresses
- Each entry associated with network interface to send packet out on
  - Process is called **forwarding**
  - (Note: Each router interface has different IP address)
  - (**Routing** is the process of building up the tables of information)
- If network not listed – packet sent to some default router
- Routers only need to know about
  - local hosts
  - some other networks

35

Forwarding and Subnets

- Forwarding table entries – slightly different form
  - this-network-num.subnet-num.0
  - this-network-num.this-subnet-num.host
- Router on subnet k knows how to
  - send packets to other subnets
  - send packets to hosts on subnet k
- Subnet k router *doesn't* need to know about hosts on other subnets
- Forwarding table sometimes called routing table

36

Subnet masks

- Mask that removes host-id when ANDed with address
  - Bitwise AND
- Example address: 130.50.15.6
  Subnet mask: 255.255.252.0
- Destination subnet: ___.___.___.___
- Subnet address is looked up in forwarding table

37

UQ Network

- Class B
  - 130.102.*.*
- Most subnets used to be at 8 bit boundary
- Now have been rationalized
- Old ITEE (then CSEE) subnets
  - 130.102.16
  - 130.102.48
  - 130.102.64
  - 130.102.65
  - 130.102.96
  - 130.102.180
  - 130.102.192
  - 130.102.193
- Current ITEE subnets
  - 130.102.64-67
    - (10 bit host number)
    - Staff/postgrad network
  - 130.102.72-75
    - (10 bit host number)
    - Student network
  - 130.102.79
    - (8 bit host number)
- Main UQ routers don't need to know all hosts on campus
- Just need to know about subnets

38

Hosts and Forwarding

- Hosts with multiple network interfaces do forwarding too
- PC example
  - `route print`

39

CIDR – Classless Inter-Domain Routing

- A,B,C classful networks too inflexible
- IP address blocks now allocated in a classless manner using a hierarchy of registries
- Networks expressed as
  base IP address/N
  where N is the number of bits identifying the network part of the address
- Examples
  - 58.0.0.0/15 belongs to Fujitsu
  - 125.128.0.0/11 belongs to Korea Telecom
  - 203.48.0.0/14 belongs to Telstra

[From www.apnic.net]

40

Registry Hierarchy

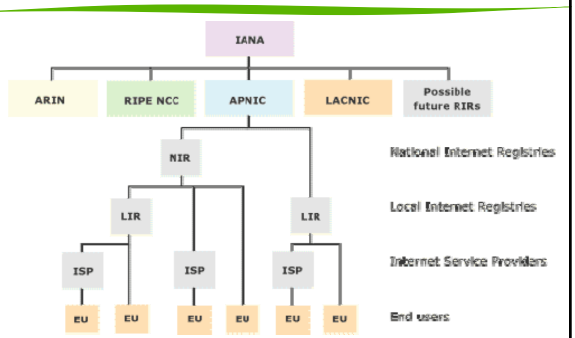

[From www.apnic.net]

41

CIDR and Routing

- ISPs can allocate blocks of addresses within the blocks that have been allocated to them
- Routers outside the ISP only need to know about the common prefix
- This is called **routing prefix aggregation**, or **supernetting** or **route summarization**

42

NAT - Network Address Translation

- One approach to deal with shortage of IP addresses
- Basic idea
  - Assign an entity (organisation) a single IP address
  - Use unique, private IP addresses within organisation
    - These same addresses can be used within multiple organisations
  - Change private IP address into organisation's IP address when packet leaves network
- Three ranges of private IP addresses exist
  - Details in class
- Private IP addresses must not appear on Internet

43

NAT - Network Address Translation

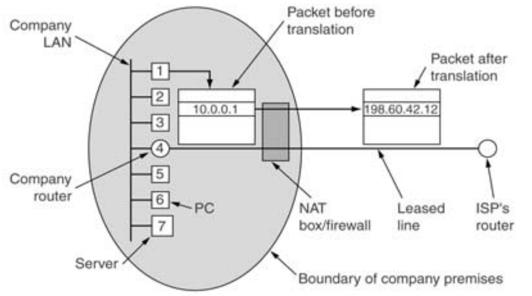

44

Short Break

- Stand up and stretch

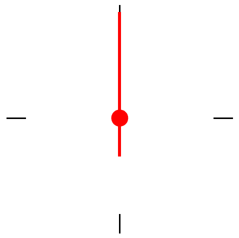

45

Addresses on the Network

- Ethernet doesn't understand IP addresses
- Actually need to send info with MAC address
- How do we map IP addresses to LAN addresses?
  - Static – have a configuration file or table
  - Dynamic – ask over the network

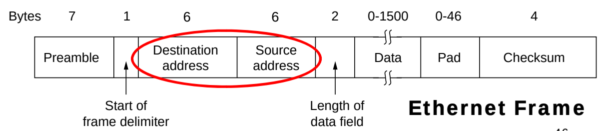

Ethernet Frame

46

ARP: Address Resolution Protocol

- Example: Host needs to send to 192.31.65.5 (already known to be on local network)
  - Sends broadcast packet:
    - "Who owns IP address 192.31.65.5?"
  - ONE reply should come back
- Variations
  - Cache
    - Entries should expire every few minutes
  - Supply own details when make request
  - Broadcast on boot
    - "Who owns my IP address?"

47

Example

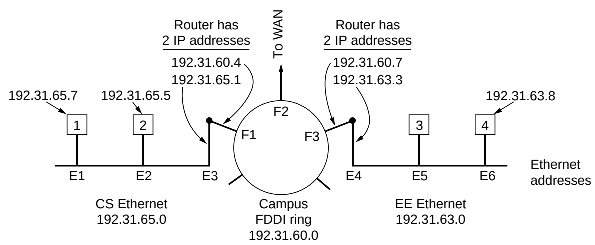

- Several class C networks (how do we know this?)
- How does an IP packet get from
  - host 1 to 2? host 1 to 4?

48

Coming Up

- Friday
  - Course recap – don't miss this
  - Course evaluation

49