ELEC / COMP 177 – Fall 2011

Computer Networking → Ethernet

Some slides from Kurose and Ross, *Computer Networking*, 5th Edition

Schedule

- **Project #2** – Due Thursday, Nov 10th
  - *By midnight*
- **Homework #5** – Due Thursday, Nov 17th
- ***Later this semester:***
  - ***Homework #6 - Presentation on security/privacy***
    - ***Topic selection*** *– Due Tuesday, Nov 22nd*
    - **Slides** – Due Monday, Nov 28th
    - **Present!** – Tuesday, Nov 29th (and Thursday?)
  - ***Project #3*** *– Due Tue, Dec 6th*

How to Physically Connect Computers?

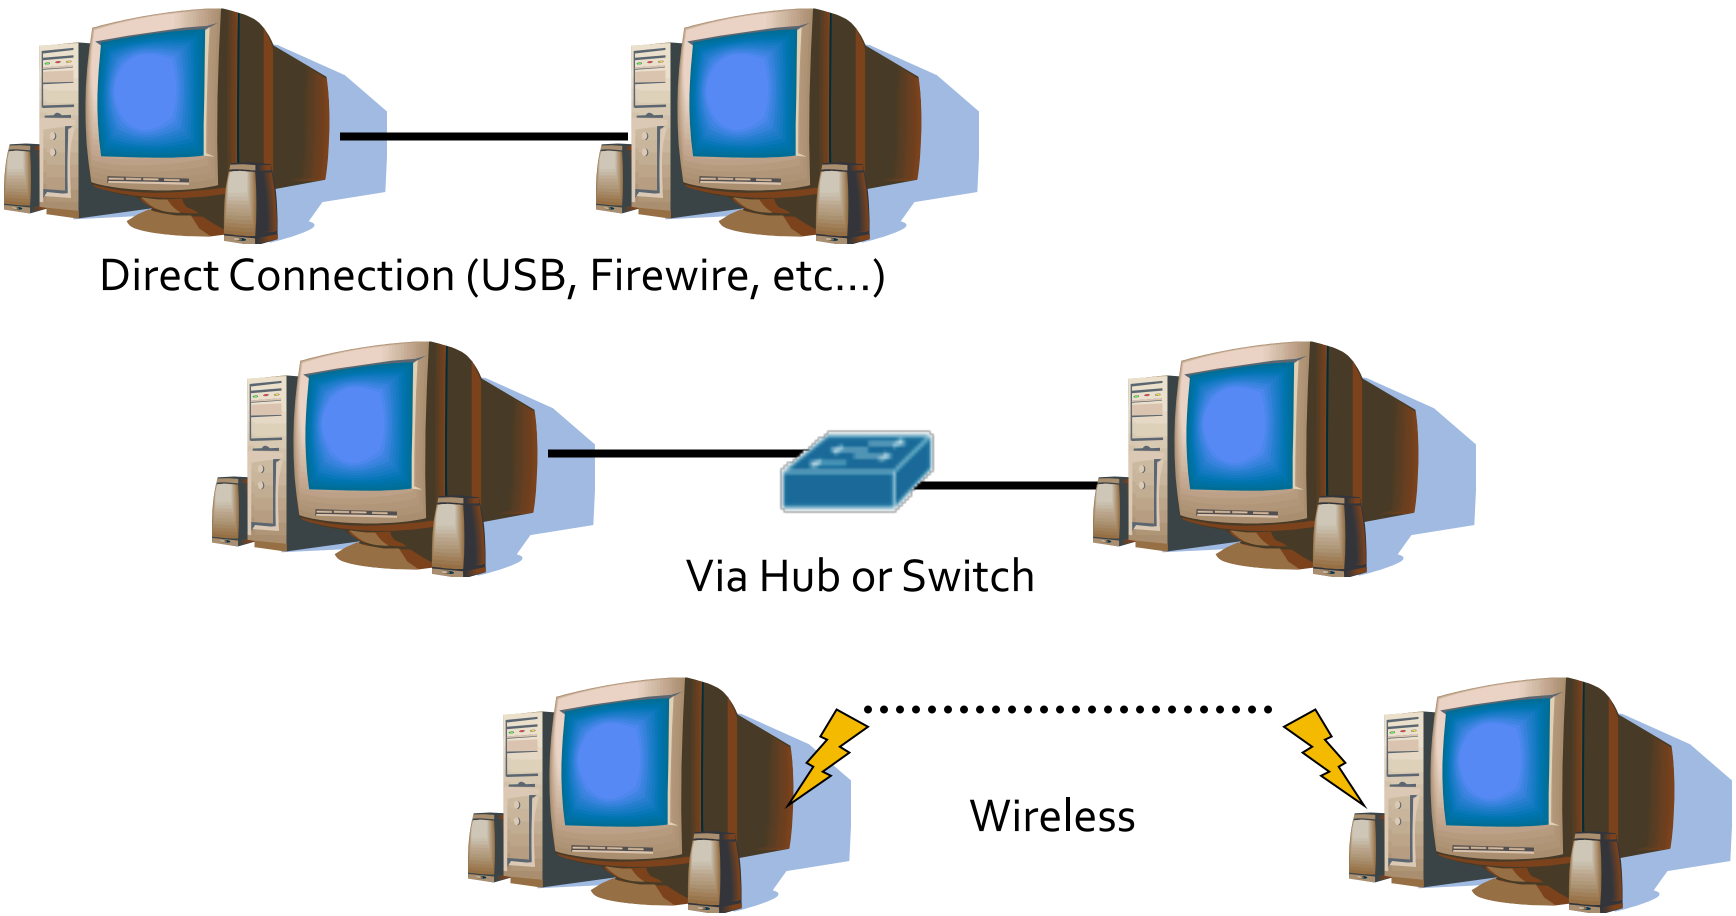

How to Build a Network?

- Four challenges
  - *Encoding* – How to format bits on wire
    - Different solutions for different media (copper, optical, wireless)
  - *Framing* – How to separate sequences of bits into independent message
  - *Error Detection* – How to detect corrupted messages (and possibly repair them)
  - *Media Access Control* – How to share a single wire or frequency among multiple hosts
    - Goals: Fair between users, high efficiency, low delay, fault tolerant

Standards that Solve Challenges

- Many competing standards with varying levels of complexity (for both wired and wireless networks)
  - Token Ring (IEEE 802.5)
  - Ethernet (IEEE 802.3)
  - Wi-Fi (IEEE 802.11 a/b/g/n)
- We focus on Ethernet networks in this course
  - Different standards made different choices, but design principles are similar

The Original Ethernet

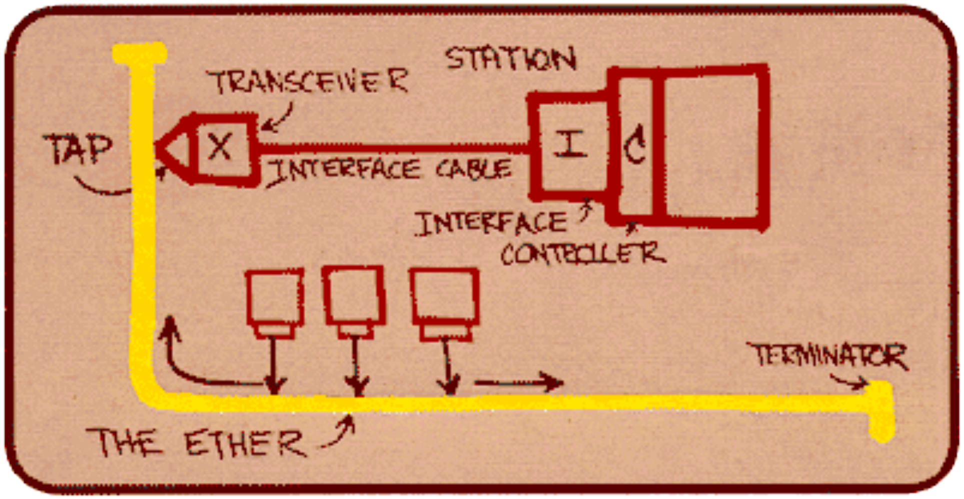

Original picture drawn by Bob Metcalfe, inventor of Ethernet (1972 – Xerox PARC)

Ether – 19th century name for media enabling the propagation of light

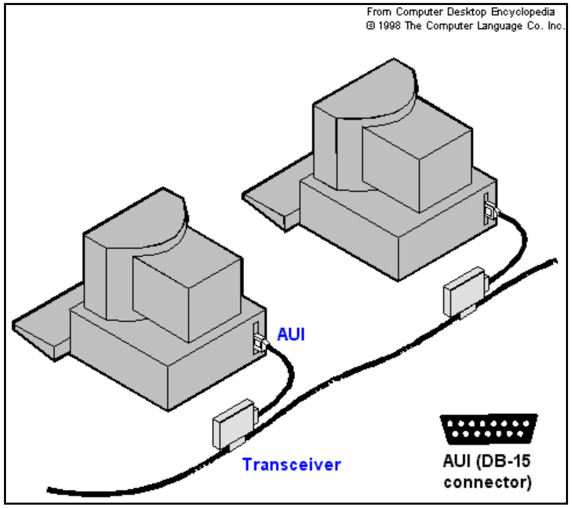

The 10Mb/s Ethernet Standard

- IEEE 802.3
- Common MAC protocol and frame format
- Multiple physical layers to choose from
  - Bus architecture (shared)
    - 10Base-5 : Original Ethernet
      - Thick coaxial cable with taps every 2.5 meters to clamp on network devices
    - 10Base-2 : Thin coaxial cable version with BNC "T" connectors
  - Star architecture (point-to-point)
    - 10Base-F / 10-Base-T standards – Introduced later!

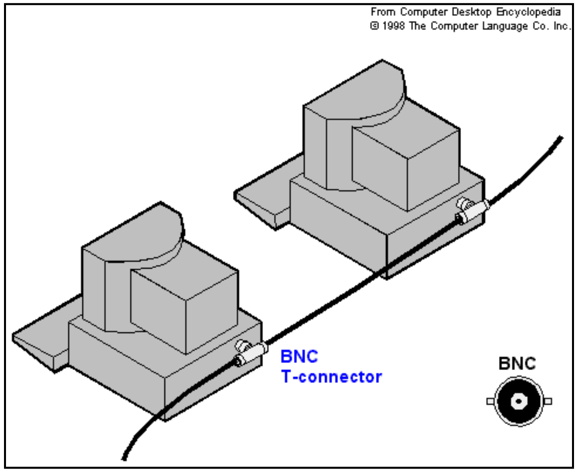

Challenge - Encoding

- How to turn bits into physical signals to send across wire?
- Data transmission across media always distorts data
  - Attenuation (amplitude reduction)
  - Distortion (change in shape)
- Encoding will make transmitted data resilient to these effects

Challenge - Encoding

- Common challenges to recovering data
  (*ignoring any modulation issues, and only using discrete high and low signals*)
  - Baseline Wander – Receiver distinguishes between high and low by keeping an average.
    (Above average = 1, Below average = 0)
    - If data stream contains long periods of 1's or 0's, average can shift!
  - Clock Recovery
    - Too expensive to have dedicated clock wire
    - Must examine incoming data and derive clock
    - Derived clock can skew during long periods of all 1's or 0's

Manchester Encoding (10Mb/s Ethernet)

- Exclusive-OR of input bit and clock
- Pulse decoded by direction of the midpulse transition rather than by its sampled level value
- No long periods without clock transition (prevents receiver clock skew and baseline wander)
- Drawback: Only 50% efficient!
  - *Bit rate* (actual data transfer) is half of *baud rate* (raw channel capacity)

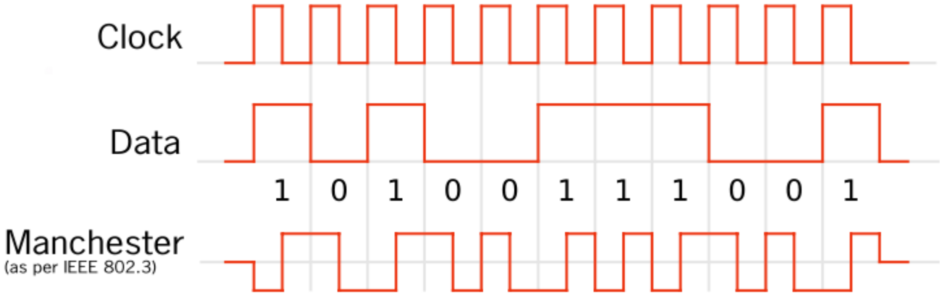

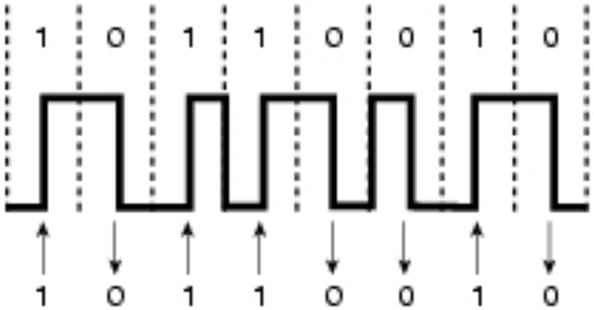

Images from Wikipedia and http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ethernet.htm

Challenge – Ethernet Frame Format

Bytes:	7	1	6	6	2	0-1500	0-46	4	
	Preamble	SFD	DA	SA	Type	Data	Pad	CRC	*Gap…*

- Preamble
  - Alternating 1's and 0's provide data transitions for synchronization
  - Used to train receiver clock-recovery circuit (critical since different transmitters will be using different clocks)
- SFD (Start of Frame Delimiter)
  - Indicates start of frame data. Always 0xAB
- DA (Destination Address) / SA (Source Address)
- Type: Indicates data type or length
- Pad: Zeroes used to ensure minimum frame length of 64 bytes
- CRC (Cyclic Redundancy Check)
- Interframe Gap: Allow time for receiver to recover before next packet
  - Length: 96 times the length of time to transmit 1 bit
  - 9.6 μs for 10 Mbit/s Ethernet, 960 ns for 100 Mbit/s Ethernet, and 96 ns for 1 Gbit/s Ethernet

Ethernet - Addressing

- All Ethernet devices have globally unique 48-bit address assigned by manufacturer
- IEEE assigns OUI (Organizationally Unique Identifier) prefix to each manufacturer
  - Remaining bits are unique per device and chosen by manufacturer
- Example: `0x 00-07-E9-CB-79-4F`
  - `0x 00-07-E9` = Intel Corp (assigned by IEEE)
    - Bit also indicates device is unicast, not multicast
  - `0x CB-79-4F` = Unique address per NIC (picked by Intel)
- Special destination address to broadcast to all devices
  - `0x FF-FF-FF-FF-FF-FF`
- NIC is responsible for filtering packets
  - Address matches (or broadcast)? Send up to host
  - Otherwise, discard

Challenge – Error Detection

- How to detect errors in transmission across copper wire / fiber?
- Ethernet solution
  - 32-bit CRC (Cyclic Redundancy Check) stored in frame header
- n-bit CRC detects all error bursts not longer than n bits, and a $1-2^{-n}$ fraction of all longer error bursts
  - Very useful since most transmission errors on a wire are bursty in nature!

Ethernet CRC

- Limitations of Ethernet CRC
  - No protection against deliberate corruption or alteration of message in transit – Not security!
  - No protection against corruption when packet is transferred through host systems, but only across wire
    - Can still have failures in NIC, memory, data bus (PCI) at <u>either end</u> of the network
  - Insufficient information to recover from error
- Design decision - More efficient to retransmit upon error than to always send enough redundant bits to repair errors
  - Receiver discards invalid packet
  - Transmitter can optionally retransmit packet (but this is not specified in Ethernet protocol)
- Calculated by NIC in real-time as packet is transmitted

Ethernet CRC

- Why do I need a CRC? If Ethernet is unreliable, shouldn't I just let a higher-level protocol detect any errors?

Challenge – Media Access Control

- **CSMA / CD** Protocol for Ethernet
  - *Carrier Sense Multiple Access with Collision Detection*
  - Developed for use with single <u>shared</u> coaxial cable of original Ethernet
  - Decentralized technique – No central arbitration, access tokens, or assigned time slots are needed to manage transmission

Ethernet – Media Access Control

- How to Transmit
  - Prepare frame for transmission
  - Check if media is idle ("Carrier Sense"). If not, wait until idle (plus interframe gap)
  - Transmit frame. Listen for any collisions and enter recovery mode
  - If no collision, finish transmitting
- Max frame size of 1500 bytes prevents one device from monopolizing network

Animation from http://www.datacottage.com/nch/eoperation.htm

Ethernet – Media Access Control

- How to Handle Collisions on Shared "Ether"
  - Continue transmission until minimum packet time is reached to ensure that all receivers detect the collision.
  - Wait random backoff time based on number of collisions
    - Backoff time exponentially increases if >1 collision per frame
  - Restart frame transmission again

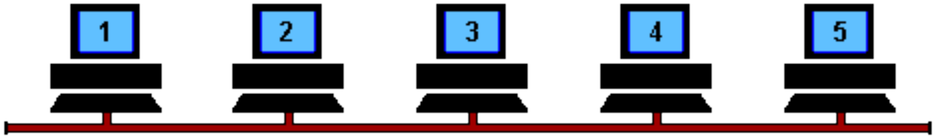

Animation from http://www.datacottage.com/nch/eoperation.htm

Ethernet – Media Access Control

- Example of worst case collision
  - Two most-distant devices send a frame (A and D)
  - D doesn't start transmitting until frame from A has almost arrived
  - D detects collision almost immediately
  - A doesn't detect collision until data propagates all the way down the wire
- Maximum time for collision detection
  - Twice the signal propagation time across entire network (for signal from A to reach D and return with collision)

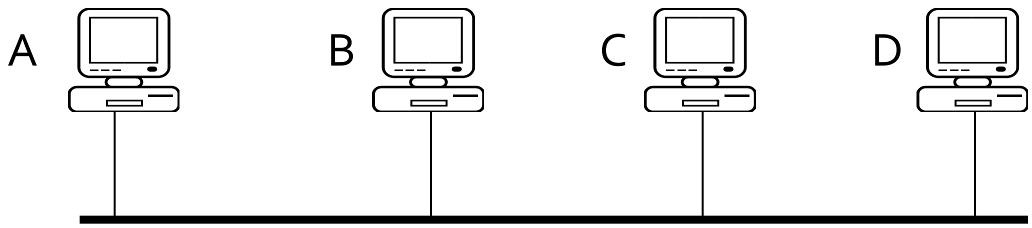

Minimum Packet Size (10Mb Ethernet)

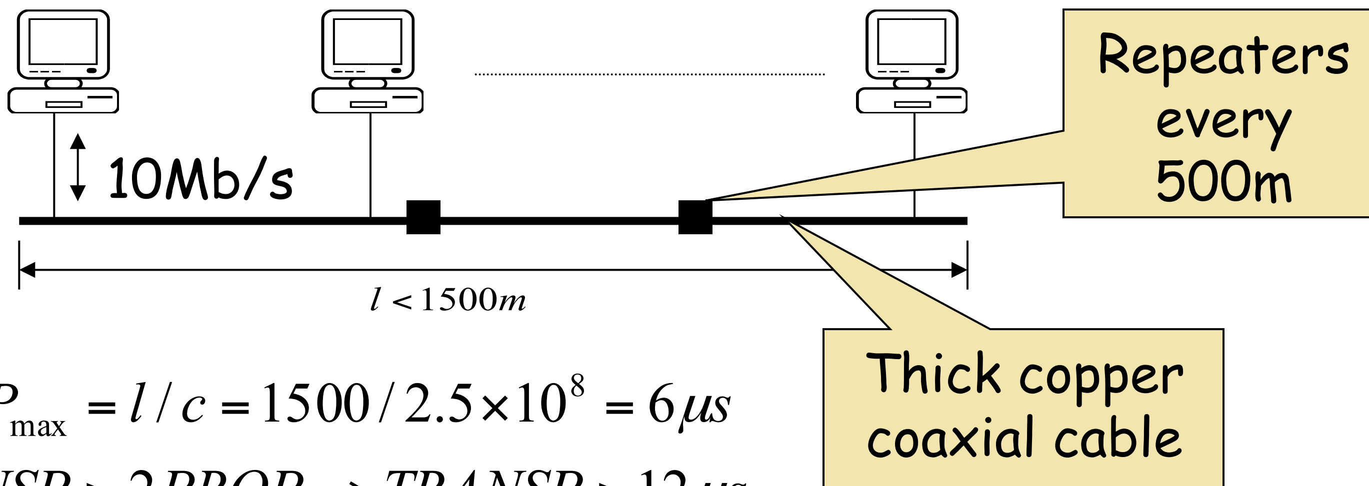

$$PROP_{\max} = l / c = 1500 / 2.5 \times 10^{8} = 6\mu s$$

$$TRANSP > 2PROP \Rightarrow TRANSP > 12\mu s$$

$$\therefore Packetsize \geq (12\mu s) \times 10Mb / s = 120bits$$

- In practice, minimum packet size is 512 bits
  - Allows for extra time to detect collisions
  - Allows for “repeaters” that can boost signal

Ethernet – Media Access Control

- Ethernet device receives frames meeting any of the following conditions:
  - Frames addressed to its own MAC address
  - Frames addressed to the broadcast address
  - Frames addressed to the multicast address (if configured for this device)
  - All frames (in *promiscuous* mode)

Ethernet – MAC Goals

- Previously mentioned design goals – Were they accomplished?
  - Fair between users? (What if users cheat?)
  - High efficiency?
  - Low delay?
  - Fault tolerant?

Recap

- Four Networking Challenges
  - *Encoding* – How to format bits on wire
  - *Framing* – How to separate sequences of bits into independent message
  - *Error Detection* – How to detect corrupted messages (and possibly repair them)
  - *Media Access Control* – How to share a single wire or frequency among multiple hosts

Broader Picture

- Where does Ethernet fit in the big picture?
- Media Access Control
  - Challenge: Who can transmit over a shared channel
  - **Channel Partitioning**
    - Divide channel into smaller “pieces” (time slots, frequency, code)
    - Allocate piece to node for exclusive use
  - **Random Access**
    - Channel not divided, allow collisions
    - “Recover” from collisions
  - **“Taking turns”**
    - Nodes take turns, but nodes with more to send can take longer turns

Broader Picture

- **Channel Partitioning** (time, frequency)
  - Optical networks for telecom (FDMA)
- **Random Access**
  - CSMA
    - Traditional Ethernet
    - 802.11 wireless
- **"Taking turns"**
    - Bluetooth, IBM Token Ring