Packet-Based Transmission

CS442

Frames

- Packet = Small block of data
  - Good for sending data of shared networks, e.g. multiplexed
- Frame = Packet, but in the context of a specific type of network
  - To send a frame, the sender and receiver must agree on the format of the frame
  - i.e. we must know how to specify the beginning and the end of each frame
  - Synchronous - block of data transmitted without start or stop bits, requires synchronized clock (or embed clock signal in data, e.g. Manchester encoding)

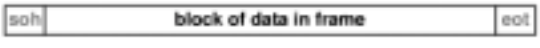

Frames - Overview

- Need to indicate start and end of block
- Use preamble and postamble
  - e.g. block of 11111111 patterns ending in 11111110
- More efficient (lower overhead) than async for large blocks of data
  - Consider a block that requires 8 bytes of overhead (preamble, postamble, other bookkeeping) with 100 information characters, or 800 bits.
  - Transmission Efficiency = 800 / 800+64 = 92.6% efficiency
  - Using async with start/stop bit, 8 data bits: 800 / 800+200 = 80% efficiency
  - Efficiency goes up the larger our blocks are, assuming the overhead bytes stay the same. Why not make our blocks even bigger?

Transmission Efficiency

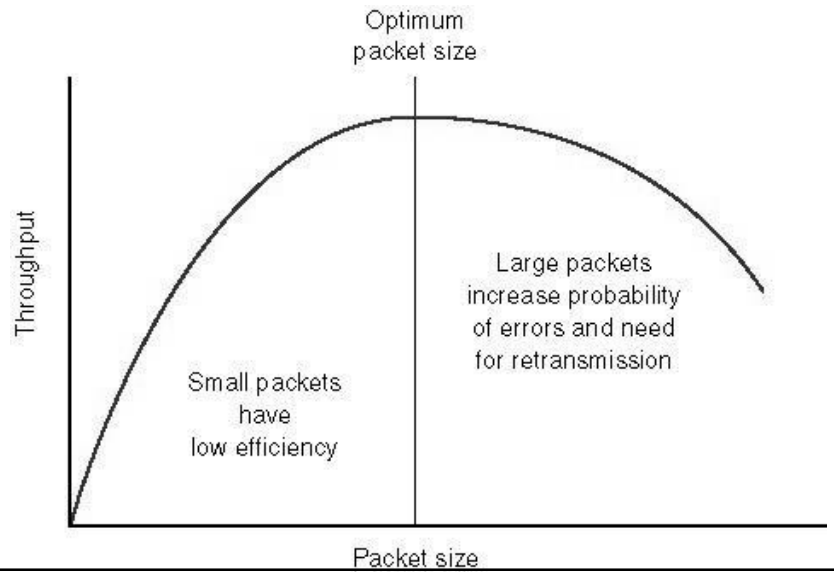

Bit Stuffing with HDLC

- Delimit frame at both ends with 01111110
- May close one frame and open another
- Receiver hunts for flag sequence to synchronize
- **Bit stuffing** used to avoid confusion with data containing 01111110
  - If sender wants to send six 1's as data then it really sends five 1's and then a 0, followed by a 1: 1111101
  - Receiver needs logic to check the next bit after 5 bits are received. If it is 0, then the 0 is deleted and the next bit is read.
  - If receiver reads 1 and seventh bit is 0, accept as flag
  - Known as the **transparency problem**

Byte Stuffing

Same problem at the Byte Level

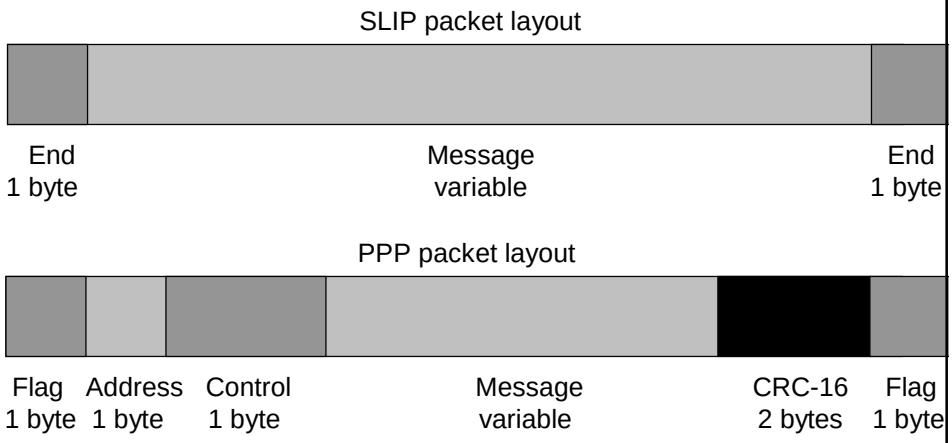

Uses Escape Character to avoid Transparency Problem
More on the frame field formats later

Error Detection

- Additional bits added by transmitter for error detection code
- Parity
  - Value of parity bit is such that character has even (even parity) or odd (odd parity) number of ones
  - Even number of bit errors goes undetected

Polynomial Checking

- Polynomial checking adds 1 or more characters to the end of the message based on a mathematical algorithm.
- With **checksum**, 1 byte is added to the end of the message. It is obtained by summing the message values,and dividing by 255. The remainder is the checksum. (95% effective)
- With CRC (Cyclic Redundancy Check), 8, 16, 24 or 32 bits are added, computed by calculating a remainder to a division problem. (99.969% with 8-bit, 99.99% with 16 bit).
  - See book for more details on implementing CRC

Forward Error Correction

- **Forward error correction** uses codes containing sufficient redundancy to prevent errors by detecting and correcting them at the receiving end without retransmission of the original message.
  - Hamming code
  - Hagelbarger code (corrects up to 6 consecutive bit errors)
  - Bose-Chaudhuri code

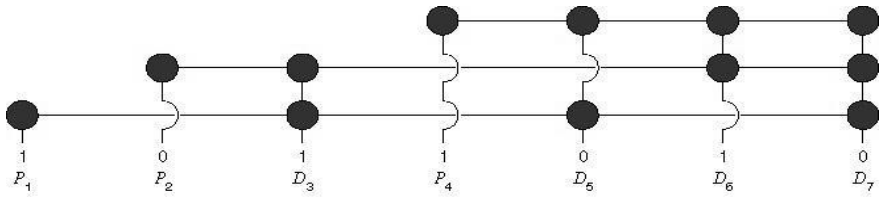

Checking relations between parity bits (P) and data bits (D)

0 = Corresponding parity check is correct 1 = Corresponding parity check fails			Determines in which bit the error occured
P_4	P_2	P_1	
0	0	0	no error
0	0	1	P_1
0	1	0	P_2
0	1	1	D_3
1	0	0	P_4
1	0	1	D_5
1	1	0	D_6
1	1	1	D_7

Interpreting parity bit patterns

Error Control

- Detection and correction of errors
- Lost frames
- Damaged frames
- Automatic Repeat reQuest (ARQ)
  - Error detection
  - Positive acknowledgment
  - Retransmission after timeout
  - Negative acknowledgement and retransmission
- More on this later

LAN Technologies

- Simplest is **point to point**
  - Computers connected by communication channels that each connect exactly two computers
  - Allows flexibility in communication hardware, packet formats, etc.
  - Provides security and privacy because communication channel is not shared
- But it doesn't scale

$$Conn = \frac{n(n-1)}{2}$$

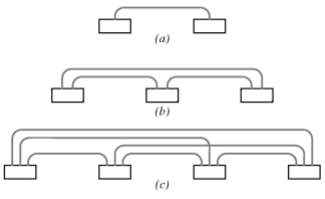

LAN Technologies

- Reduce cost by reducing number of connections via a shared network
- **But** ... attached computers compete for use of shared connection
- Local communication almost exclusively LAN
  - Works out OK though due to locality of reference
  - Principle that we access computers nearby more frequently than those far away
- Long distance almost exclusively point-to-point
  - SMDS
  - ATM
  - Modem / PPP

LAN Topologies

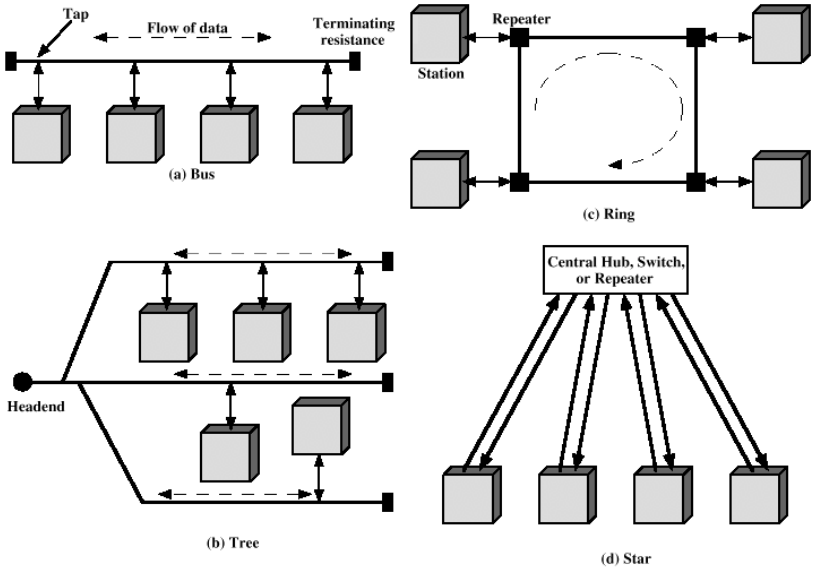

Bus and Tree Properties

- Multipoint medium
- Transmission propagates throughout medium
- Heard by all stations
  - Need to identify target station, each station has unique address
- Full duplex connection between station and tap
  - Allows for transmission and reception
- Need to regulate transmission
  - To avoid collisions
  - To avoid hogging, data in small blocks - frames

Frame Transmission - Bus LAN

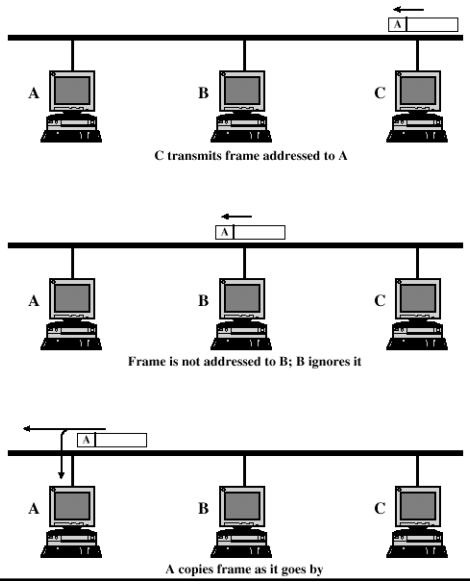

Data broadcast to all computers

Only one system may transmit at a time or we get a **collision**

No transmissions if cable is cut

Typically coax cable

Ring Topology

- Repeaters joined by point to point links in closed loop
  - Receive data on one link and retransmit on another
  - Links unidirectional
  - Stations attach to repeaters
- Data in frames
  - Circulate past all stations
  - Destination recognizes address and copies frame
  - Frame circulates back to source where it is removed
- Media access control determines when station can insert frame

Frame Transmission Ring LAN

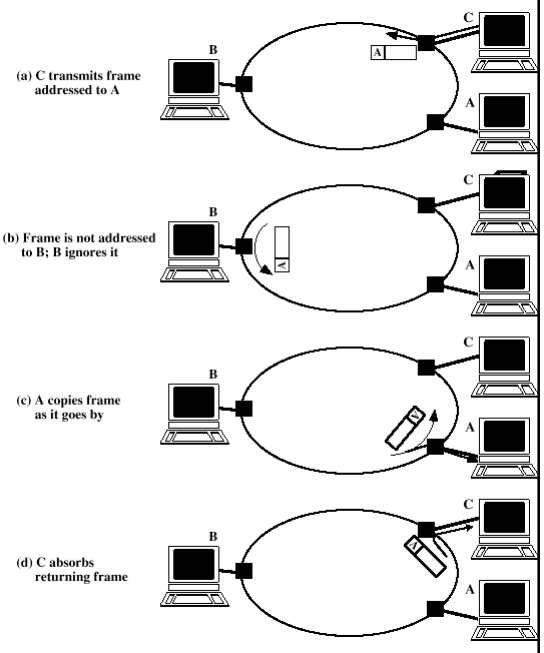

Star Topology

- Each station connected directly to central node
  - Usually via two point to point links
- Central node can broadcast
  - Physical star, logical bus
  - Only one station can transmit at a time
- Central node can act as frame switch
  - Prevents broadcast to other stations except destination, can increase network efficiency
- Typically use twisted pair as the medium

Two Level Star Topology

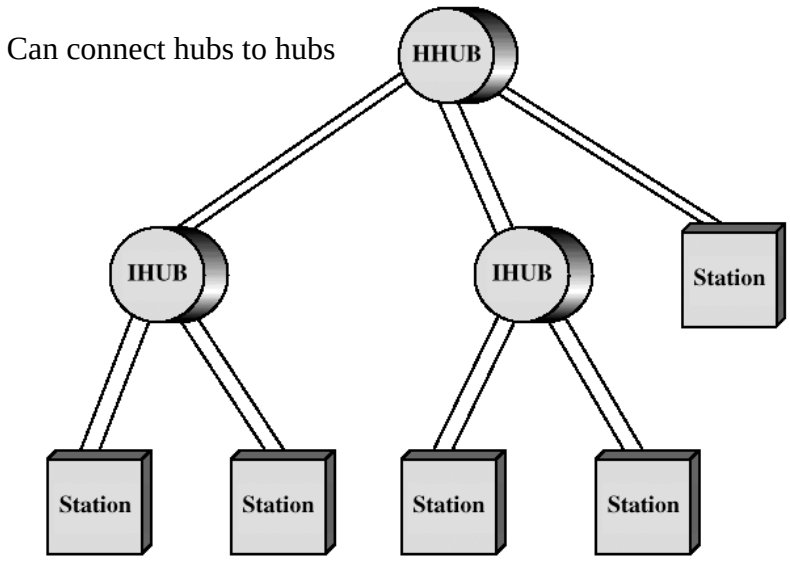

Methods for Media Access Control

- Needed to keep collisions from happening
- Round robin
  - Good if many stations have data to transmit over extended period
- Reservation
  - Good for stream traffic
- Contention
  - Good for bursty traffic
  - All stations contend for time
  - Distributed
  - Simple to implement
  - Efficient under moderate load
  - Tend to collapse under heavy load

Media Access Control - Ethernet (CSMA/CD)

- Original form: 10 Mbps, or 1.25 Mbytes/sec
- Broadcast bus
- Carriers Sense Multiple Access with Collision Detection
  - Distributed access control; i. e., no central control on when you can transmit
  - Stations randomly access medium, contend for access
  - hardware gets every packet
- Xerox - Ethernet
- IEEE 802.3

Ethernet's Predecessor - ALOHA

- Packet Radio
- When station has frame, it sends
- Station listens (for max round trip time) plus small increment
- If ACK, fine. If not, retransmit
- If no ACK after repeated transmissions, give up
- Frame check sequence (as in HDLC)
- If frame OK and address matches receiver, send ACK
- Frame may be damaged by noise or by another station transmitting at the same time (collision)
- Any overlap of frames causes collision
- Max utilization 18% - propagation delay < transmission time, high chance of collision! Can wait/listen for free medium first

CSMA

- Propagation time is much less than transmission time
  - Travel at about 70% the speed of light
- All stations know that a transmission has started almost immediately
- Max utilization depends on propagation time (medium length) and frame length
  - Longer frame and shorter propagation gives better utilization
- First listen for clear medium (carrier sense)
- If medium idle, transmit
- If two stations start at the same instant, collision
- Wait reasonable time (round trip plus ACK contention)
- No ACK then retransmit

If Busy?

- If medium is idle, transmit
- If busy, listen for idle then transmit immediately

- If two stations are waiting, collision

CSMA/CD

- With CSMA, collision occupies medium for duration of transmission
- Stations listen whilst transmitting
- If medium busy
  - listen for idle, then transmit
- Else if medium idle
  - transmit
  - Listen for collision
  - If collision detected, jam then cease transmission
    - After jam, wait random time then start again
    - Binary exponential back off
  - Else wait min idle time - give others nodes a chance
    - Distributed fairness, 51.2us for 10Mbps

CSMA/CD Operation

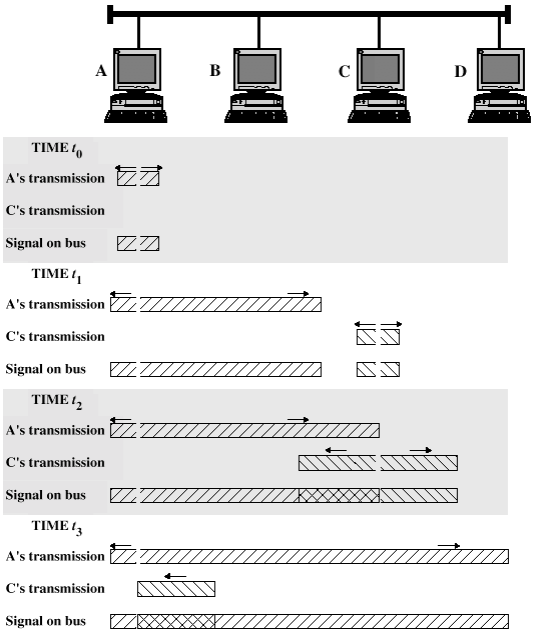

Frames must be long enough to detect collision during transmission!

Collision Retransmission

- Retransmit up to N attempts, e.g. 16
- If collision, must send jam signal, random backoff and retransmit
  - Jam signal == 512 bits (64 bytes), to ensure the end nodes hear the collision
  - The Ethernet minimum frame is 64 bytes (46 data)
- Binary exponential algorithm
  - Wait 1, 2, 4, 8 … 2^m time- slots, etc. where m=Min(10, nth retry)
  - Packets can be lost due to collision, especially if network is heavily used
- Modern network cards can saturate cable; only ~30% utilization under heavy transmission load

Collision Detection

- On baseband bus, collision produces much higher signal voltage than signal
- Collision detected if cable signal greater than single station signal
- Signal attenuated over distance
- Limit distance to 500m (10Base5) or 200m (10Base2)
- For twisted pair (star-topology) activity on more than one port is collision
- Special collision presence signal

Gigabit Ethernet - Differences

- Basically, the same as regular Ethernet but faster
- Carrier extension
  - Frames at least 4096 bit-times long
  - Why? Propagation time is too short, and we want the transmission time longer than the propagation time to detect collisions
- Frame bursting
  - Allows many small frames to be transmitted consecutively up to a limit to avoid carrier extension on short frames

Gigabit Ethernet - Physical

- 1000Base-SX
  - Short wavelength, multimode fiber
- 1000Base-LX
  - Long wavelength, Multi or single mode fiber
- 1000Base-CX
  - Copper jumpers <25m, shielded twisted pair
- 1000Base-T
  - 4 pairs, cat 5 UTP
- Signaling - 8B/10B

Gigabit Ethernet Alliance

- http://www.gigabit-ethernet.org
  - 3Com, Bay Networks, Cisco, Compaq, Granite Systems, Intel, LSI Logic, Packet Engines, Sun Microsystems,UB Networks, VLSI Technology
- Selling points
  - Relatively low cost?
  - Easy migration
  - New emphasis on network applications driving growth

Token Ring (802.5)

- MAC protocol
  - Small frame (token) circulates when idle
  - Station waits for token
  - If no data to transmit, forwards token immediately
  - If data to transmit, “Holds” token temporarily
  - Sends data frame, can send multiple frames up to some maximum number
  - Frame(s) make round trip and are absorbed by transmitting station
  - Station then inserts new token when transmission has finished and leading edge of returning frame arrives
- Under light loads, some inefficiency
- Under heavy loads, round robin

Token Ring Operation

A sends to C

C sends to A and D

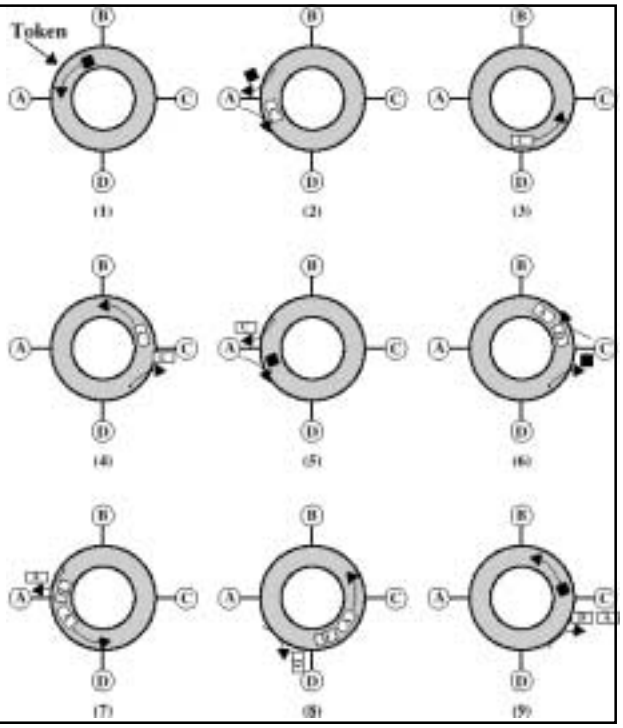

FDDI

- Fiber Distributed Data Interface (Can also run over twisted pair at shorter distances)
- 100Mbps
- LAN and MAN applications
- Token Ring

FDDI MAC Protocol

- As for 802.5 except:
  - Station seizes token by aborting token transmission
  - Once token captured, one or more data frames transmitted
  - New token released as soon as transmission finished
- Two rings for reliability
- Allows data frames from multiple senders to be on the ring at once, higher utilization

FDDI – Multiple Rings

•FDDI uses *counter-rotating* rings in which data flows in opposite directions

•In case of fiber or station failure, remaining stations *loop back* and reroute data through spare ring

•All stations automatically configure loop back by monitoring data ring

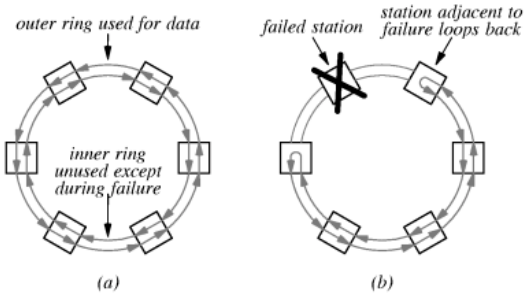

FDDI Operation

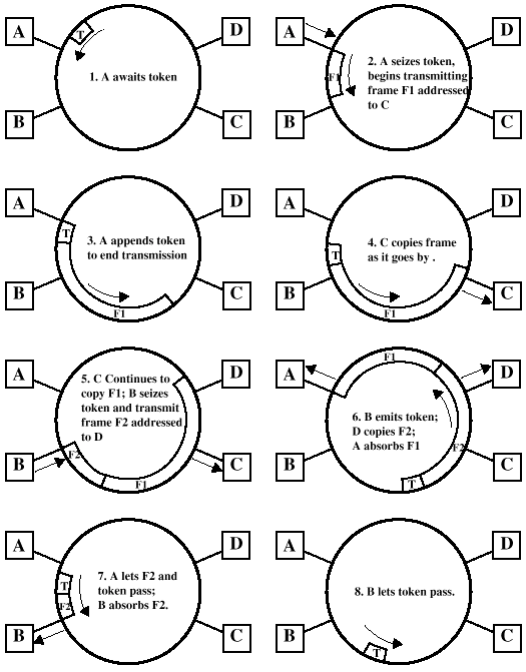

LAN Generations

- First
  - CSMA/CD and token ring
  - Terminal to host and client server
  - Moderate data rates
- Second
  - FDDI
  - Backbone
  - High performance workstations
- Third
  - ATM / Gigabit Ethernet
  - Aggregate throughput and real time support for multimedia applications

ATM LANs

- Asynchronous Transfer Mode
- Virtual circuit
- Types of ATM LANs
  - Gateway to ATM WAN, acts as router
    - Need LAN encapsulation (LANE server or edge switch) if carrying ethernet traffic
    - Maps ethernet address to Virtual Circuit identifier, broadcast
  - Backbone ATM switch
    - Single ATM switch or local network of ATM switches
  - Workgroup ATM
    - End systems connected directly to ATM switch
    - Need to replace ethernet cards, ATM operates at data link layer!
  - Mixed system

Example ATM LAN

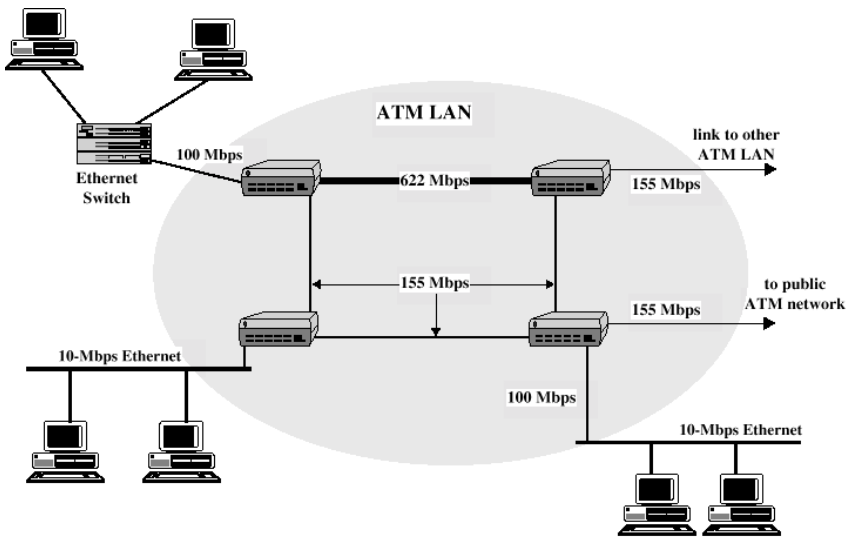

ATM LAN HUB

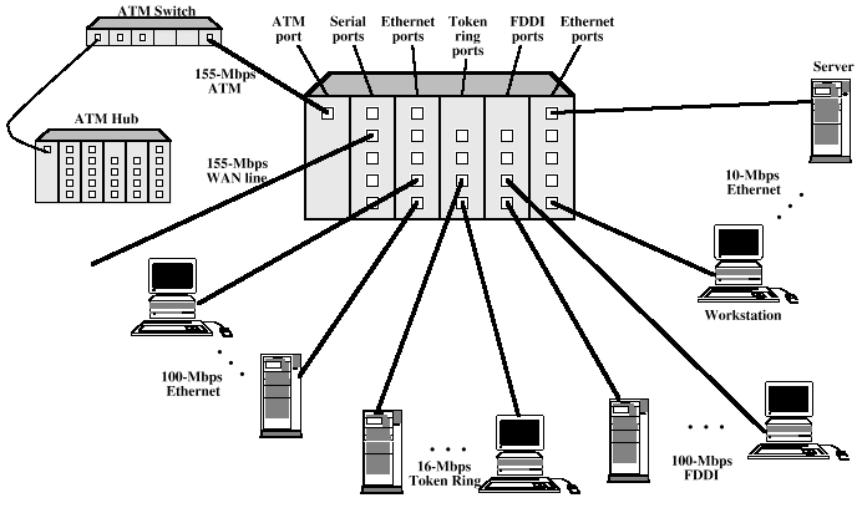

ATM at Dallas Cowboys (www.dallascowboys.com)

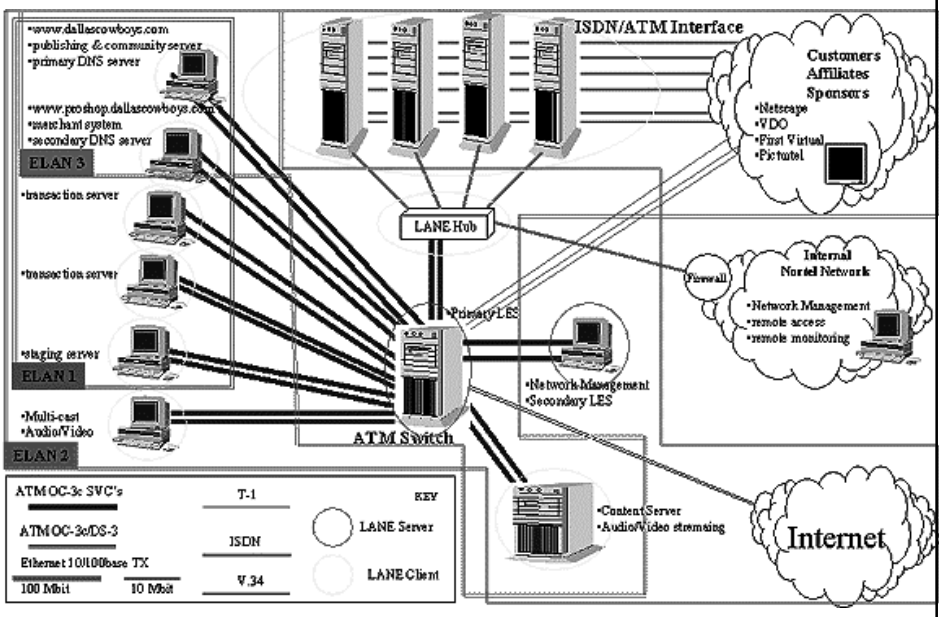

ATM inside Reuters News

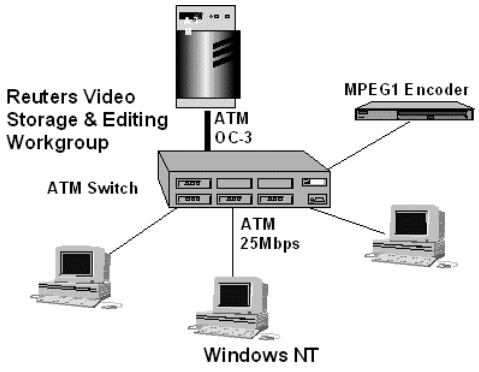

ATM Forum

- Standards body for ATM
  - http://www.atmforum.org
- Some members
  - AT&T
  - Cisco
  - 3Com
  - IBM
  - Lucent
  - LSI Logic

Wireless LANs

- IEEE 802.11
- Basic service set (cell)
  - One or more stations using same MAC protocol, together with an **access point** or **base station**
  - Typically some radius (~200 ft) from the access point
  - Competing to access shared medium
  - May connect to backbone via access point (bridge)
- Extended service set
  - Two or more BSS connected by distributed system
  - Appears as single logic LAN to Data Link layer

Wireless LAN - Physical

- Infrared
  - 1Mbps and 2Mbps
  - Wavelength 850-950nm
- Direct sequence spread spectrum
  - 2.4GHz ISM band
  - Up to 7 channels
  - Each 1Mbps or 2Mbps
- Frequency hopping spread spectrum
  - 2.4GHz ISM band
  - 1Mbps or 2Mbps
- Others under development

802.11 Media Access Control

- Instead of CSMA/CD, uses CSMA/CA
  - Carrier Sense Multiple Access, Collision Avoidance
- No collision detection, aims at avoiding collisions in the first place
- Even with collision detection, if no collision is sensed one could still occur at the receiver
- A, C send to B but no direct signal between A-C

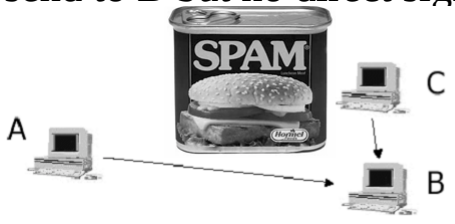

Hidden station problem

802.11 Media Access

- If media is idle for time DIFS (Distributed Inter Frame Space) then allowed to transmit
- If media is busy
  - Wait until media is free for time DIFS
  - Then wait an ADDITIONAL random backoff time before sending instead of immediately as in 802.3
- Receiver might get a collision, but no detection for one – instead the CRC will (hopefully) fail
  - Send an ACK back if data correctly received
  - Needs to send ACK due to lack of collision detection
  - If sender doesn't receive ACK in time period, resends

802.11 Media Access

- Also supports a way to reserve channel access, avoids hidden station problem
- Sender sends short Request To Send (RTS)
- Receiver sends short Clear To Send (CTS) with permission to the Sender
  - All other stations also hear the CTS, and have to withhold sending until transmission is complete
- Sender sends data
- Recipient sends ACK if successful
  - Other stations also hear the ACK, can now issue their own RTS
- Might have collisions with RTS or CTS but not with Data or ACK. If collisions, stations requesting to send will not get a CTS and will have to wait random time to re-submit the RTS

Frame Identification

- Previous section on LAN technology described techniques for providing connectivity between computers
- Need to devise technique for delivering message through LAN medium to single, specific destination computer
- Sending computer uses a *hardware address* to identify the intended destination of a frame
- Sending computer also identifies **type** of data carried in the frame

Specifying a Destination

- Data sent across a shared network reaches all attached stations - for all LAN topologies
  - Not the switched technology though!
- Interface hardware detects delivery of frame and extracts frame from medium
- But ... most applications want data to be delivered to one specific application on another computer - not **all** computers

Hardware Addressing

- Most network technologies have a hardware addressing scheme that identifies stations on the network
- Each station is assigned a numeric *hardware address* or *physical address*
- Sender includes hardware address in each transmitted frame
- Only station identified in frame receives copy of frame
- Most LAN technologies include sender's hardware address in frame, too

LAN Hardware and Packet Filtering

- A little detail about organization of LAN hardware and computer:

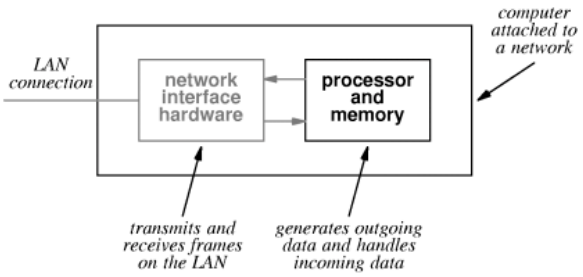

LAN Hardware and Packet Filtering

- LAN interface handles all details of frame transmission and reception
  - Adds hardware addresses, error detection codes, etc. to outgoing frames
  - May use DMA to copy frame data directly from main memory
  - Obeys access rules (e.g., CSMA/CD) when transmitting
  - Checks error detection codes on incoming frames
  - May use DMA to copy data directly into main memory
  - Checks destination address on incoming frames
- If destination address on incoming frame matches the local station's address, a copy of the frame is passed to the attached computer
- Frames **not** addressed to the local computer are ignored and don't affect the local computer in any way

Broadcasting

- Some applications want to *broadcast* messages to all stations on the LAN
- Shared communication channel can make broadcast efficient - message is delivered to all stations
- Special *broadcast address* used to identify broadcast messages, which are captured by *all* stations
  - Station passes data up if the address on the message matches its own address **or** the broadcast address

Multicast

- What if some computers on a LAN want to broadcast for their app, but others don't?
  - Those that don't are stuck processing a bunch of broadcast data they don't care about
  - Broadcast data is not filtered, it is passed up the networking layer stack
  - Solution is **multicasting**
- Set up a multicast address
  - This address passes to the network interface
  - Interface can then choose to pass data up if the frame's address matches the station's address **or** broadcast address **or** multicast address

Ethernet Addressing

- Each NIC has a **unique** MAC address, assigned by IEEE in its ROM
- 48-bit integer, 6 unsigned char bytes
  - Broadcast address: FF FF FF FF FF FF
  - Multicast address: 01 XX XX XX XX XX
  - unicast address the rest, e.g.: 00 02 B3 9F 3A 1B
  - first 3 bytes are manufacturer code
  - Intel: 00 02 B3
  - Sun: 08 00 20
  - 3Com Europe: 00 50 99
- See http://standards.ieee.org/db/oui/ for more

Ethernet Frame Formats

- Two formats!
  - Ethernet 2.0 (Xerox, Unix)
  - IEEE 802.3 (Novell / ISO / Unix)
    - Has the two sub-layers for LLC (logical link control) and MAC

Ethernet Frame Formats

Ethernet 2 Format - Min length=64, Max=1518 bytes

Destination address 6 bytes	Source address 6 bytes	Type 2 bytes	Message variable 46-1500 bytes	CRC-32 4 bytes

Type indicates higher level protocol or control:
IP = 0x0800, AppleTalk over Eth: 0x809B ARP = 0x0806, RARP=0x8035

IEEE 802.3, Min length=64, Max=1518 bytes

Destination address 6 bytes	Source address 6 bytes	Length 2 bytes	LLC 6 bytes	Type 2 bytes	Message variable 38-1492 bytes	CRC-32 4 bytes

Ethernet Frame Formats

- How can an ethernet driver tell the difference between the two frame types in a heterogenous network?
- The frame formats don't overlap (at least not for most types)
  - 802.3 Len >= 46 && <= 1500
  - Ethernet IP Type == 0x800, 2048 in decimal

Other Things to Note

- Ethernet 2 format didn't specify the length of the frame
  - How do we know how large the frame is?
  - Recipient has to read in all the data and then compute that the last 4 bytes are CRC, stuff in between Type and CRC is the payload
- In 802.3 a portion of the LLC and Length fields are called **SNAP** (SubNetwork Attachment Point)

MTU – Maximum Transmission Unit

- We said that Ethernet has a maximum frame size of 1518 bytes (less for the data, depending on Ethernet 2 vs. 802.3)
- Limit on size of frame transmitted at data link layer
- To see MTU on UNIX:
  - netstat -in
  - Ethernet 2: 1500
- 802.3: 1492
- Microsoft's PPP: 1500 (576 better for high-latency modems)
- If the Network Layer has a bigger packet than MTU, it must be fragmented

Ethernet Headers and Trailers

- Both formats have an 8 byte preamble used for synchronization (alternating 1's and 0's, makes a regular square wave)
- CRC code:
  - If computed CRC <> packet's CRC, the packet is thrown away
  - Rely on higher level protocols for data integrity
- No retries, ACKs, or NAKs, so-called "best effort"
  - Ethernet is an **unreliable** protocol
  - Makes it easy and cheap though!
  - what does/doesn't ethernet CRC guarantee you?

Ethernet Packets

- Does CRC catch all bit errors?
- No!

- Are most errors caught?
- Yes!

- Can you guarantee your packet will arrive?
- No! Collision, congestion, Noise

- If going through 10 routers or bridges, can you guarantee reliability?
- Definitely no!

Network Analyzer

- A *network analyzer* or *network monitor* or ``*network sniffer''* is used to examine the performance of or debug a network
- Can report statistics such as capacity utilization, distribution of frame size, collision rate or token circulation time
- Can record and display specific frames, to understand and debug packet transmissions and exchanges

Network Analyzer

- Basic idea is a computer with a network interface that receives **all** frames
- Sometimes called *promiscuous mode*
- Many desktop computers have interface that can be configured for promiscuous mode
  - Combined with software, computer can examine *any* frame on LAN
  - Communication across a LAN is **not** guaranteed to be private!
- Computer receives and displays (but does not respond to) frames on the LAN
- We'll have a homework assignment later using a network analyzer