Local-area network technologies

Shared-medium broadcasting

LAN topologies

802 MACs & Ethernet

LANs & TCP/IP

ARP & DHCP

Non-TCP/IP LANs

Overview

- Compared to WANs, local-area networks:
  - serve small areas (reduced geographic scope)
  - serve desktops and departmental servers
  - are generally connected by high-speed communications channels
- Varying philosophies, differing technologies
  - switched point-to-point
  - shared-medium packet broadcasting
- Data link and physical layers
- References:
  - Tanenbaum: Ch1.2.1, 4.1, 4.2.1–4.2.3, 4.3, 4.5.1, 4.5.2, 5.5
  - Kurose & Ross: Ch 5.1, 5.3, 5.5–5.7, 5.4

Switched point-to-point

- Evolution of telephone technology:
  - PBX – private branch exchange
  - digital PBX
  - digital data switch
- Based on a “switch” (hardware or software or hybrid) that creates point-to-point connections
- The distinction between telephony and data communications is blurred
  - telco manufacturers provide “data options” for digital PBXs
  - networking manufacturers creating voice capabilities for data networks

Shared-medium broadcasting

- In shared-medium broadcasting, all hosts transmit & receive in the same medium (channel)
  - broadcasting: every host receives every frame sent by any host
    - channel is *multiple access*
    - analogous to two-way radio with a single frequency
  - this environment requires careful coordination: handled by *Medium Access Control* (MAC) protocols (MAC layer)
  - *contention*: what happens if more than one host wants to transmit?
    - transmission *collisions* might occur: electrical interference that causes each transmissions to be garbled and unusable
  - how to deal with contention
    - by complete avoidance
    - by retransmission (after collision)

… shared-medium broadcasting, 2

- Basic techniques:
  - MAC layer manages *channel allocation*
    - two basic styles: static vs dynamic
  - for static channels, use TDM or FDM to divide the channel in advance
    - class of protocols called *partitioning protocols*
    - suitable for busy networks or networks that require predictable performance, but generally wasteful for non-continuous traffic
  - for dynamic allocation, hosts used entire channel "on demand" whenever they want
    - protocols classified as *random access* or *taking-turn*
      - taking-turn protocols avoid collisions
      - random-access protocols allow hosts to transmit at random and collide
        - collisions are handled by retransmission

Broadcast: contention

- Various methods of resolving contention
- ALOHA:
  - transmit at random, then listen to determine collision
  - wait and retransmit if necessary
- Carrier sense:
  - listen to channel
  - if channel is idle, start transmitting
  - if channel is busy, wait a while and retry
- Carrier sense with collision detection:
  - listen, transmit if channel is available
  - listen during transmission
  - stop transmitting immediately if collision
    - collisions can still occur because of the latency of electrical signals
    - wait and retransmit

Topology and media

- Topology: the logical (as opposed to physical) arrangement of the hosts
  - two types:
    - bus (also called tree)
    - ring
  - star is sometimes considered a topology
- Choices for medium:
  - twisted-pair
  - baseband coaxial
  - broadband coaxial
  - fiber-optic
  - wireless

… topology and media, 2

- In theory, any medium for any topology
  - in practice, only a few combinations are used
    - twisted pair dominates desktop and most installations
    - fiber used for MAN, WAN and backbones (fiber to the desktop still on the horizon)
    - baseband coaxial is legacy only
    - broadband coaxial is used as part of cableco HFC, and in special-purpose industrial applications
- Wireless is an emerging technology
  - wireless LANs for building interiors to eliminate wiring costs & mobile computers
  - wireless for LANs is separate from wireless for hand-held PDAs & similar devices in the cellular world
  - wireless is also used for WAN point-to-point (microwave & related) which is also not related to wireless LANs

Bus & tree

- Bus:

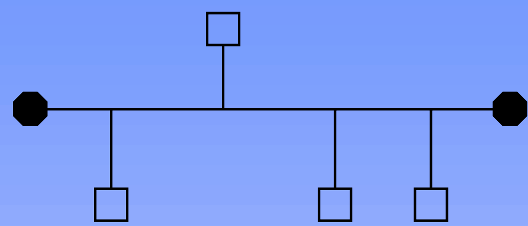

  - data moves in both directions
- Tree:

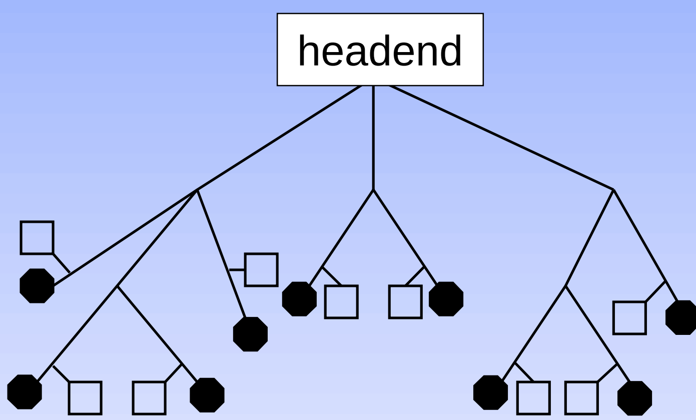

  - data moves in both directions on a branch, “headend” retransmits on other branches
  - each branch is a bus (i.e. a bus is just a single-branch tree)
  - size of branches (number of headends) is a major cost factor
  - analog broadband (cableco) for MAN trees
    - problem: analog signalling is unidirectonal, required FDM for bidirectional

Ring

- Ring is formed from a sequence of point-to-point connections

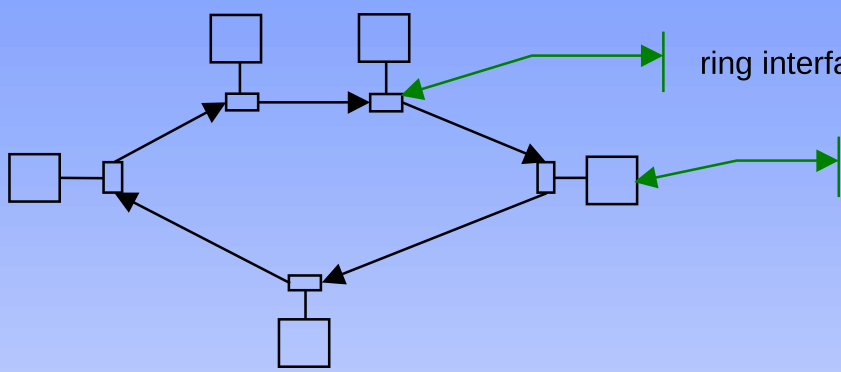

- Connections are really between “ring interface” devices (Tanenbaum)
- Ring interface reads each bit and retransmits it (listen mode)
- Obviously each link is not “shared” in the bus broadcasting sense
  - “shared” because all links must be used by all stations for arbitrary point-to-point
- Data flow is unidirectional (links all go in the same direction)

Star: layout or topology?

- A star-shaped network (all hosts joined at a central location) can be considered just a wiring layout or a full-fledged topology
- A bus can be transformed into a star:

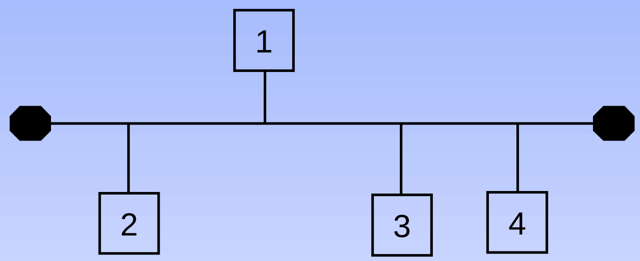

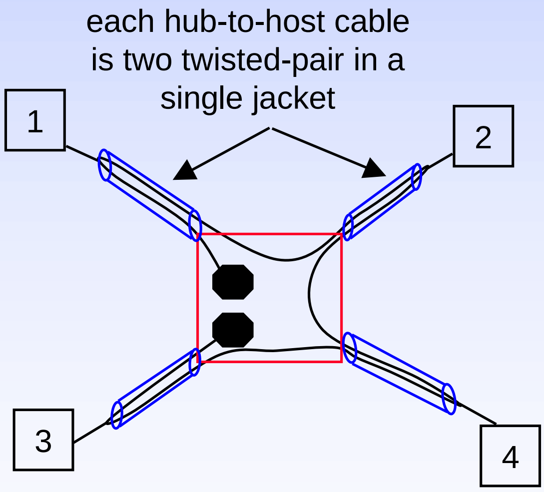

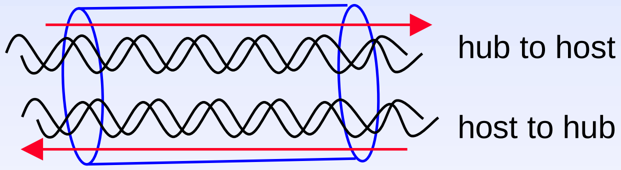

- A ring can be similarly transformed

…star: layout or topology, 2

- In the preceding simple case, the device at the centre is passive (doesn't contain any active circuitry or software)
  - called a *hub*, also, wiring hub, repeater
    - hubs operate at the physical layer
  - doesn't really change the topology, just a layout convenience
- Star/hub wiring can help with network management
  - failure in linear bus wiring hard to find, affects every host
  - failure in star wiring affects just the one host
- But, the hub becomes a single point of failure
  - if it fails, no communication anywhere
- A *switch* is a device that is positioned like a hub, but is entirely different
  - a switch operates at the data link layer; discussion to follow

Commercial MAC implementations

- MAC sublayer is responsible for managing contention and resolving collisions
  - bus: simultaneous transmissions travel along the medium and collide
  - ring: two stations injecting data into the ring will collide
- There is an almost universal family of MAC standards and protocols
  - the 802.x series standardized by IEEE
    - 802.2: LLC – logical link control; equivalent to upper half of OSI DL
    - 802.3 and above: lower half OSI DL and physical layer
  - common addressing scheme (6 byte) and other commonality where feasible
    - called *MAC addresses*
    - addresses are assigned to hardware by manufacturer (6 byte = 48 bit => 2^{48} ~= $3x10^{14}$ (30 trillion))

Ethernet 802.3

- Ethernet™ is a bus topology that uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD):
  - standardized as IEEE 802.3
  - basic operation
    - host waits for idle bus (no carrier)
    - transmits; listens for collision
    - if collision, stop
    - wait random amount of time, repeat
- Cabling options (historic interest only):
  - 10Base5: "thick Ethernet", 10Mbps, baseband coaxial, 500 m/segment
  - 10Base2: "thin Ethernet", 10Mbps, baseband coaxial, 200 m/segment
  - 10Base-F: 10Mbps, fiber, 2km/segment
  - 1Base5: 1Mbps, baseband twisted-pair, 250 m/segment
  - 10Broad36: 10Mbps, broadband coaxial, 3.6km/segment

… Ethernet, 2

- Standard use cabling:
  - 10Base-T: 10Mbps, baseband twisted-pair, 100 m/segment; hub with star layout
    - use is waning, legacy systems
  - some 10Base2 is still used (older legacy)
- Most common new installations are “fast” Ethernet, 802.3u
  - 100 Mbps
  - twisted-pair with central hub or switch
  - cabling options:
    - 100Base-T4: uses 4 twisted pairs; 100 m/segment
    - 100Base-TX: uses 2 “Category 5” (high quality) twisted pair; 100 m/segment
    - 100Base-TF: uses 2 fiber-optic strands; 2000 m/segments
  - note that distances are shorter with copper, inter-building can be a problem

… Ethernet, 3

- In Ethernet, choice of wait period after collision is “random”, but adjusts to number of stations
  - Ethernet is classified as a *random-access dynamic channel allocation* protocol
  - time period increases as number of consecutive collisions increase
    - called the *backoff* period; *exponential backoff* doubles the wait time after each
- 802.3 works well for lightly-loaded networks
  - as traffic increases, probability of collision increases
  - traffic usually increases as the number of hosts connected to the medium increases
- Ethernet has the operational characteristic that is scales poorly

… Ethernet, 4

- Solutions for overloaded 802.3:
  - more speed
    - undesirable price-performance
  - fewer hosts
  - partition the hosts into smaller groups, join groups together in a *switch*
    - the medium is partitioned
    - each partition has fewer hosts
    - fewer hosts means fewer collisions (probably);
    - the *collision domain* has been reduced in size and the probability of collision is reduced
- An 802.3 switch is a device at the centre of a star (same position as a hub) that performs *collision domain partitioning* and improves Ethernet performance

… Ethernet, 5

- Switch operation is a significant product differentiator:
  - there are many product configurations, options and operating characteristics
- Typically:
  - store-and-forward frames at the data-link layer
    - variation: "cut-through switching" starts sending a frame to its destination before it's fully received
    - may support full duplex or only half duplex
  - can handle combinations of 10 and 100 Mbits/sec hosts (and even 1Gbits/sec)
    - may have from 4 to 96 host interfaces of various speeds
  - are implemented with special-purpose hardware and software
    - may be as simple as a collection of hubs interconnected with router-like software, or exceedingly complex

… Ethernet, 6

- Switches can be arranged in a hierarchy
    - two-layer hierarchy of passive hubs connected to small switches can be quite cost effective
    - there are many product configurations, options and operating characteristics
- Ethernet *bridges* perform a similar task to switches
    - a bridge in the generic sense is a device that connects two networks together at the datalink layer
    - this is more or less what a switch does, so the distinction is quite unclear
        - WAN bridges are generally devices at each end of a point-to-point link (two ports)
        - Kurose & Ross simply suggest that Ethernet bridges have fewer ports

… Ethernet, 7

- Both switches and bridges are designed to be transparent (plug & play)
  - they have to learn about the topology and connected devices by inspecting frames and building tables or databases
- Extra confusion: joining LANs is also something that a router can do
  - consider IP subnetting: dividing an IP network into smaller physical networks
  - often done for network administrative autonomy, but could be done for load management
  - each network has a router to the outside world
    - routers are definitely **not** plug & play

Other Ethernet & LAN standards

- Really fast Ethernet:
  - "gigabit Ethernet" (1,000 Mbits/sec) being deployed
    - 802.3z fiber
    - 802.3ab 1000Base-T; Cat5 twisted (4 pairs)
    - not common on the desktop – yet
  - on the drawing board: 802.3ae 10Gbits/sec
- Existing fiber technology is capable of handling transmission rates of 50,000 Gbits/sec (50 Terabits/sec) easily
  - problem is electro-optical conversion
  - research into pure optical computers is crucial

… other standards, 2

- Wireless LANs: 802.11
  - designed to replace wires in LANs
  - traditional deployment has hosts with transponders organized around *access points* (base stations)
    - access points can be connected together with traditional technologies to form larger networks & “wandering” support
  - 802.11 also permits hosts with transponders to form *ad hoc* groups without a defined access point
    - e.g. laptops in a meeting room
  - 802.11 original definition: 1 Mbits/sec and 2 Mbits/sec
    - speeds based on availability of radio frequency spectrum, transmitter power rules, ITU-R bureaucracy
    - faster standards are now available (up to 11 Mbits/sec)
    - an infrared version is also defined

… other standards, 3

- Wireless Personal Area Networks: 802.15
  - WPANs
    - equivalent proprietary systems:
      - Bluetooth: Ericsson *et al*
      - HomeRF: developed Shared Wireless Access Protocol (SWAP), 1998;
  - interoperability of 802.11 devices with the ever-growing set of PDAs, pagers, WAP devices, PCS phones, etc.
  - differing goals from 802.11:
    - cheaper
    - low power consumption
    - simpler devices, lower data rate requirements (initially at least)
  - initial data rate requirements up to 1Mbit/sec

Other 802.x LAN technologies

- Token bus; 802.4
- Developed by GM and others for factory automation applications
- Goals: response-time predictability and cabling robustness
- Physical: broadband coaxial bus
  - 1, 5, 10 Mbps
  - moderate bandwidth, long distances
- MAC contention is handled by avoidance with turn-taking protocols
  - *token passing:* a sender must posses a token before sending
    - no token, no sending!
    - may hold the token for predetermined time period, then must let another host have it
  - token is just circulated between hosts if no host wants to send

… other 802.x, 2

- Token ring; 802.5
- Similar goals and ideas as token bus
  - predictable response times, even in heavy traffic
- Physical: sequence of point-to-point twisted-pair connections, forming a physical ring
  - 4, 16, 100 Mbits/sec
  - wired in a star with a hub
- MAC contention handled by token passing (same ideas as token bus)
- Token ring was popularized by IBM early in the development of the PC
  - gained significant following in IBM corporate and enterprise clientele
  - not for general use now (still applicable for guaranteed-response situations)

LLC – logical link control

- Standardized as 802.2; based on HDLC
- “Semi-layer” above other 802 standards
- Provides typical DL functions like flow-control, error checking
- Three defined services
  - connection-oriented
  - unacknowledged connectionless
  - acknowledged connectionless
- LLC not used in IP networks (IP uses MAC sublayer directly)

Other LAN technologies

- FDDI: Fiber Distributed Data Interface
  - ring topology with token-ring MAC
  - fiber-optic media; 100Mbps
  - uses two rings, one “clockwise”, other “counterclockwise”
    - redundancy for fault tolerance
    - can form one ring out of two semi-rings, if necessary
    - each station has hardware to join rings
  - typical uses: backbone LAN (connecting smaller LANs)
    - up to 200 km with 500-1000 stations

… other LAN technologies, 2

- ATM
  - early in its history, ATM was considered a potential replacement for high-speed Ethernet to the desktop
  - cheap highspeed Ethernet and TCP/IP ended the debate
    - ATM had a APIs defining how application would interface directly with ATM's uppermost layer (roughly a network layer)
      - LANE: emulate 802.3 & 5 over ATM
    - didn't happen, and probably won't
  - ATM is a very strong presence in WAN backbone technologies, especially with the developments of IP directly on ATM
- But, there is also work on running packet-switching directly on SONET (confusingly called POS: Packet Over SONET)
  - recall that SONET is synchronous, so statistical packet multiplexing was not the intended use of SONET

TCP/IP and 802 LANs

- TCP/IP is an internetworking system
  - common use is to connect 802.x LANs
- Conceptually easy:

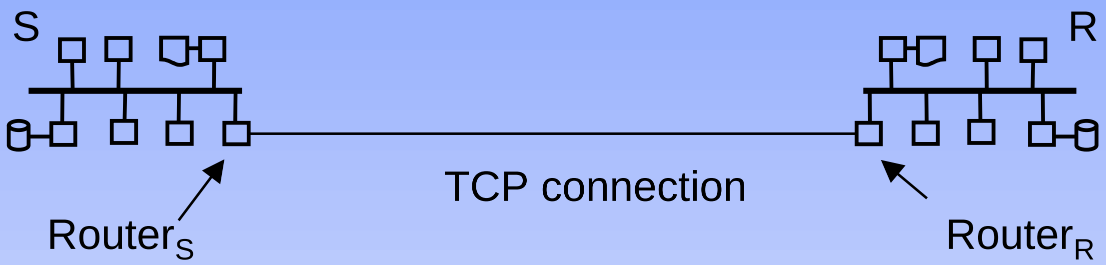

- Problem: host S sends to R, using TCP/IP address (IP address and port)
  - assume that S knows how to get to its router (Router$_S$)
  - how does the receiving router Router$_R$ know which host is R?
    - within Router$_R$'s LAN, R is identified by its MAC address
  - need translations between MAC addresses and IP addresses

… TCP/IP and 802 LANs, 2

- An equivalent problem exists for a single LAN that uses TCP/IP for its transport and higher layers

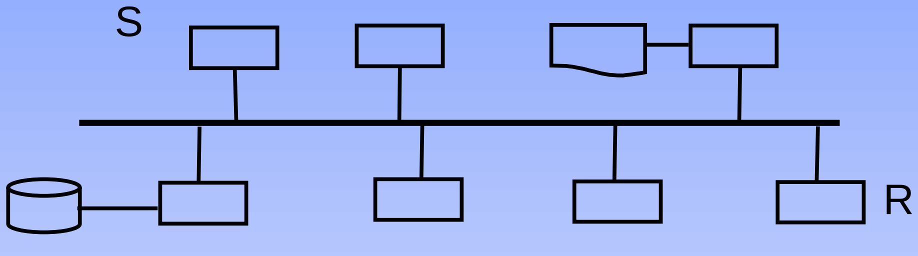

- Within a single LAN, no internetworking is required: hosts only need the LAN MAC addresses
  - if S and R use “internet” names and addresses to identify each other, translation to MAC address is required
  - DNS is not sufficient, since it translates names to IP addresses, not MAC addresses

… TCP/IP and 802 LANs, 3

- LANs need not use TCP/IP addressing internally
  - PCs, Apple both have well-established LAN protocol suites that are not TCP/IP
    - Windows NT networking: NT domains, workgroups, NetBEUI
    - LocalTalk/AppleTalk
    - Novell
    - Banyan
  - all of theses systems handled the naming and addressing of hosts with proprietary protocols
- In many LANs, there are two independent sets of upper-layer protocols that “meet” at the MAC layer
  - or, more commonly now, use IP as a datagram service for proprietary transport & higher

ARP – address resolution protocol

- So, how to solve the IP vs LAN addressing problem?

- Example:
  - S sends to R with IP address (obtained from DNS usually)
    - if S knows only the IP address of Router$_S$, use ARP to get its MAC address
  - Router$_S$ uses normal internet routing to get to Router$_R$ (e.g. by extracting the network address from Router$_R$'s IP)
  - Router$_R$ receives data for a host on its network, but only knows the IP address
  - Router$_R$ uses ARP to discover the MAC address corresponding to R's IP address
- ARP is a TCP/IP host-to-network (link) layer protocol

… ARP, 2

- ARP operation (for $Router_R$):
  - $Router_R$ broadcasts an "ARP request" on its LAN:
    - "would the owner of IP address a.b.c.d please respond"
    - LAN broadcasting is easy
  - station R responds, replies with an "ARP reply"
    - "I am a.b.c.d; my MAC address is z-y-x-w-v-u"
  - $Router_R$ then sends the data directly to R using its MAC address
- For efficiency, ARP information can be cached by routers (and hosts)
- Proxy ARP: a host that answers ARP requests for many IP addresses
  - lets the host act as an IP "funnel"
  - e.g.: dialup terminal server answers ARP request for each connectee

Reversed ARP

- The opposite problem: given a MAC address, how to determine the IP address?
  - in a workstation with secondary storage (disk), use a configuration file or similar repository
  - what about stations that have no such storage (e.g. X terminals, “network computers”)?
- Solution: reversed ARP – each LAN has a server that can respond to broadcasts:
  - “ my MAC address is z-y-x-w-v-u. does anyone know my IP address?”
  - RARP server has a database of MAC–IP mappings: it responds directly with:
    - “your IP address is a.b.c.d”

… reversed ARP, 2

- Problem:
  - LAN broadcasts cannot be routed across networks
  - need a RARP server on each network
- Bigger problem:
  - in general, stations need need more information than an IP address
  - other configuration information
  - boot image
- Solution: BootP – Boot Protocol
  - BootP is an extension of the RARP idea using routable UDP datagrams instead of LAN broadcasting
- BootP can be used to get other information necessary to boot the station
- DHCP (Dynamic Host Configuration Protocol) is an extension of BootP that allows reusable IP addresses

DHCP

- DHCP is designed to make IP address and other configuration information (e.g. host names, subnet masks, DNS addresses) “plug & play”
- A DHCP server can be set up to assign information dynamically, or can reserve based on MAC address
  - DHCP-provided information has an associated *lease time*
  - when a lease expires, the host must reacquire configuration information
    - could be the same information, or different
- DHCP is most applicable to general workstations and especially mobile hosts
  - servers need to be at known addresses
  - DHCP can be used with servers; generally uses reservations of known addresses (management convenience)

… DHCP, 2

- Hosts that use DHCP are not configured with a DHCP server address
  - when host starts, it broadcasts a *DHCP discover* message to see if any DHCP servers are available
  - host choose among responses to obtain information
- DHCP is now common for ISPs to assign configuration information customers
  - even though DSL and cableco offerings are “always on”, DHCP is still used for management purposes
- Most “Internet sharing” (NAT) devices also have a DHCP server to provide information to the interior (private IP)
- DHCP does not coordinate with DNS to create DNS entries assigned addresses
  - newer version of DHCP & DNS do

A non-TCP/IP LAN

- Novell Netware: a LAN based on 802.x standards, without TCP/IP
  - from http://www.protocols.com/pbook/novel.htm

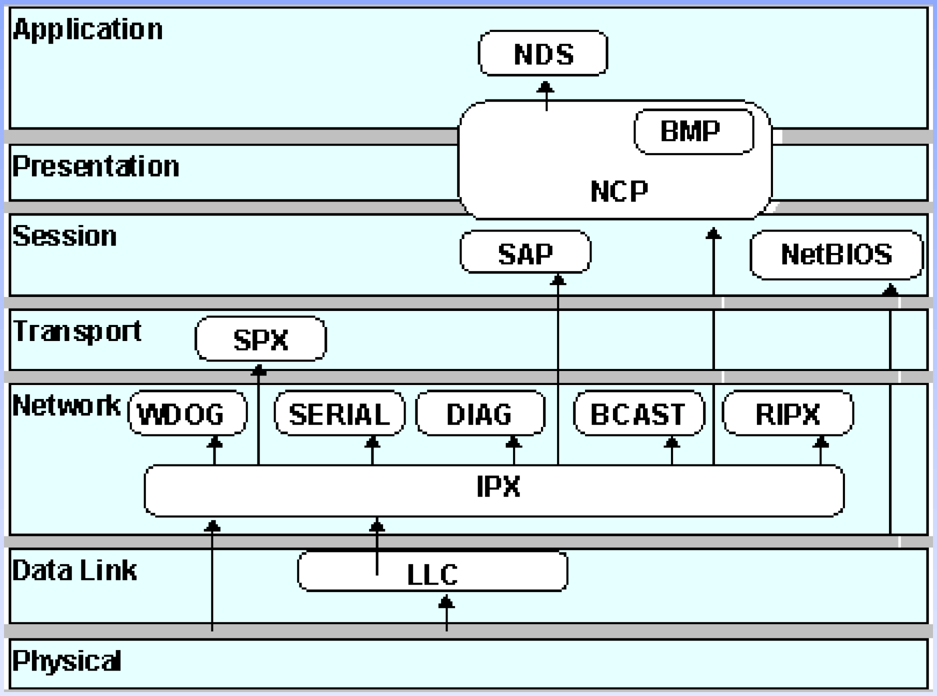

- 4 or 5 layers:
  - layers don’t align exactly with OSI, but close

… Novell netware, 2

- Application to session layers:
  - NCP: NetWare Core Protocol; connection-oriented++
  - NDS: NetWare Directory Services
  - SAP: Service Advertising Protocol
- Transport:
  - SPX: sequenced packet exchange; plain connection-oriented
- Network
  - IPX: internet packet exchange; unreliable connectionless, 12-byte addresses
  - similar in philosophy to IP: internetworking, routable
- DL and physical
  - from 802.x

Another LAN – AppleTalk

- AppleTalk from Apple Corp. protocol suite
  - from http://www.protocols.com/pbook/appletalk.htm

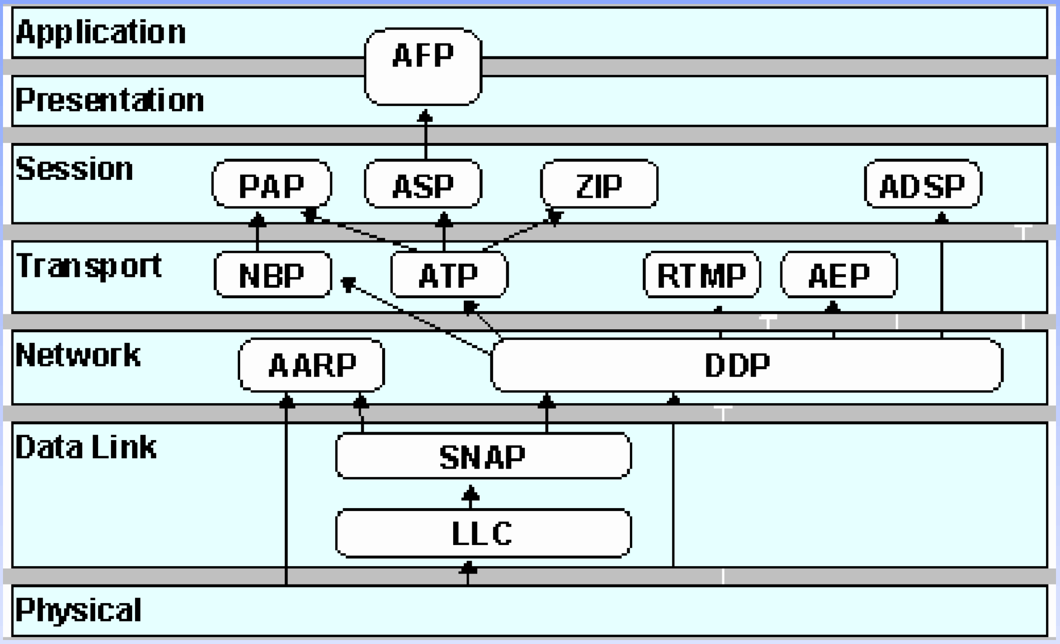

- Physical layer can be LocalTalk, or 802.3
- Internetworking, routable
- **AARP** - AppleTalk Address Resolution Protocol; **DDP** - Datagram Delivery Protocol; **RTMP** - Routing Table Maintenance Protocol; **AEP** - AppleTalk Echo Protocol; **ATP** - AppleTalk Transaction Protocol; **NBP** - Name-Binding Protocol; **ZIP** - Zone Information Protocol; **ASP** - AppleTalk Session Protocol; **PAP** - Printer Access Protocol; **ADSP** - AppleTalk Data Stream Protocol; **AFP** - AppleTalk Filing Protocol

Yet another LAN – Windows NT

- Windows NT networking is a session-to-application suite of protocols inherent in Window NT and related operating systems
  - supports client-server and peer application models
    - file browsing and sharing, printing, authentication
  - provides TCP/IP socket services
  - provides versions of TCP/IP applications an protocols
    - DNS, DHCP, SMTP, FTP
    - not the same behavior as the Unix originals, which is a source of irritation
  - some interoperability with AppleTalk & Novell application layers

… yet another LAN – Windows NT, 2

- Can use any of four transports
  - Novell IPX/SPX
  - TCP/IP
    - including all of the standard TCP/IP protocols: ARP, PPP, DHCP clients
  - NBF (NetBEUI Frame)
    - NetBEU/NBFI is primarily a network layer process, with some reach into transport
    - CO and CL services, sliding window for flow and congestion control
    - uses MAC address directly, **not routable**
    - provides services to the NetBIOS (Network Basic I/O Services) layer; which is a transport/session layer
  - DLC – Data Link Control (from the OSI HDLC, SNA SDLC, X.25 LAP & LAPB family of datalink protocols)
    - used for mainframe, network printers
    - contains an 802.2 LLC interface for driving network interface cards

… yet another LAN – Windows NT, 3

- Pictures:

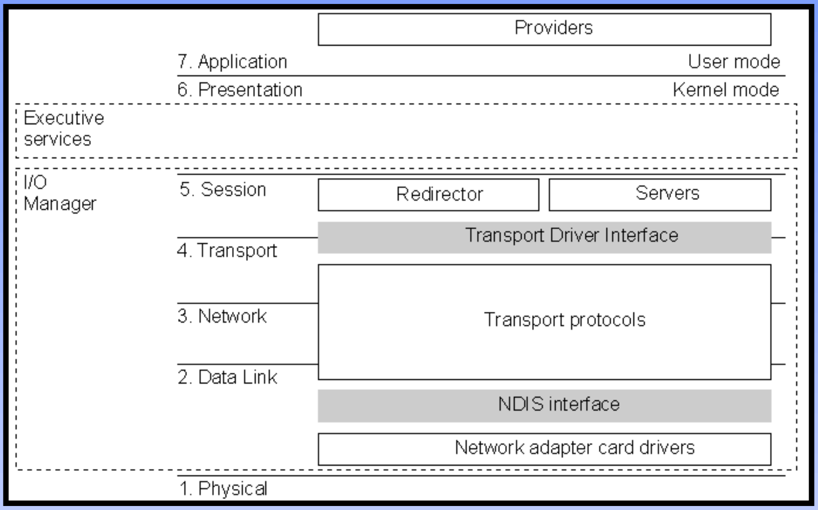

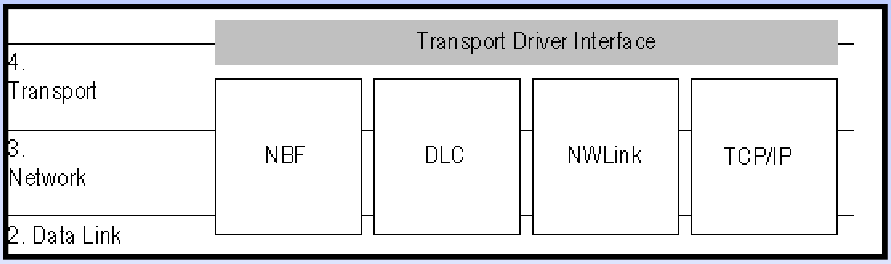

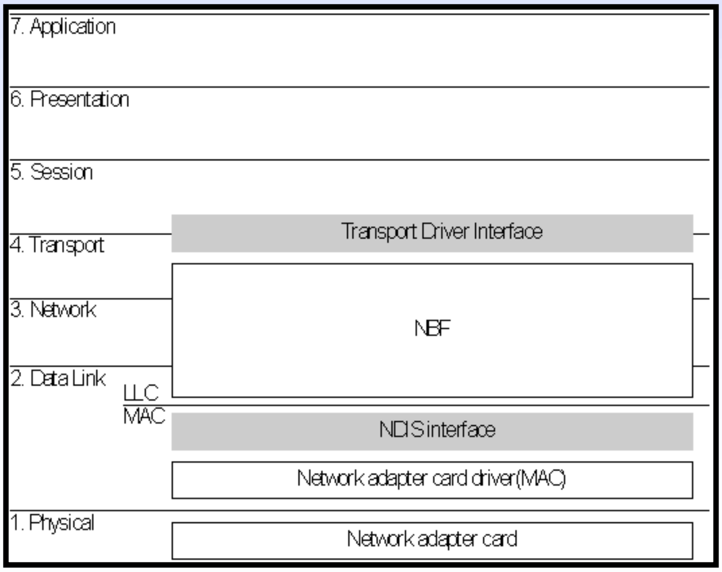

Summary

- LANs are built from collections of hosts connected to shared media
  - requires MAC – medium access control
  - broadcasting to all hosts is easy
- Various hardware implementations exist, but IEEE 802.3 Ethernet CSMA/CD is ubiquitous and the *de facto* standard
  - it has problems under loaded conditions, but ever-increasing data rates and switching technology helps to extend the life of the technology
- TCP/IP was never intended to be the lower-layer protocols within a LAN, so MAC-to-IP address translation is provided by ARP
- Proprietary non-TCP/IP LANs exist and are generally all built on 802.x