15-213

Network technology

April 11, 2000

Topics

- **Fundamental concepts**
  - protocols, layering, encapsulation, network types
- **Wide area networks**
  - phone lines and modems
  - Internet backbones
- **Local area networks**
  - Ethernet

Course Theme

Abstraction is good, but don't forget reality!

Earlier courses to date emphasize abstraction

- **Abstract data types**
- **Asymptotic analysis**

These abstractions have limits

- **Especially in the presence of bugs**
- **Need to understand underlying implementations**

Useful outcomes

- **Become more effective programmers**
  - Able to find and eliminate bugs efficiently
  - Able to tune program performance
- **Prepare for later "systems" classes**
  - Compilers, Operating Systems, Networks, Computer Architecture

"Harsh Realities" of Computer Science

- **Int's are not integers; float's are not reals**
  - Must understand characteristics of finite numeric representations
- **You've got to know assembly**
  - Basis for understanding what really happens when execute program
- **Memory matters**
  - Memory referencing bugs especially difficult
    - » Violates programming language abstraction
  - Significant performance issues
    - » E.g., cache effects
- **There's more to performance than asymptotic complexity**
  - Constant factors also matter
- **Computers do more than execute programs**
  - Get data in and out
  - Communicate with each other via networks

Typical computer system

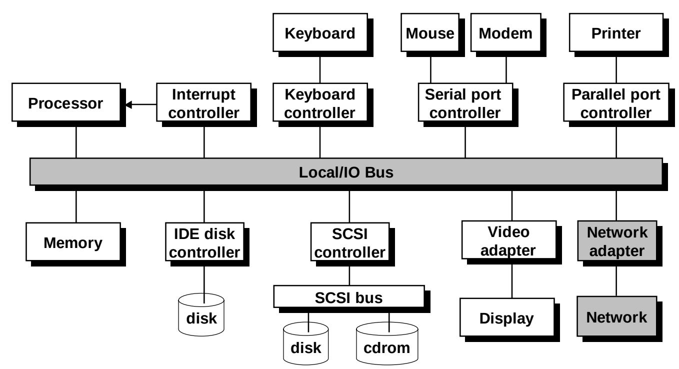

Simple example

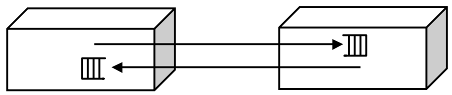

Starting Point: Want to send bits between 2 computers

- FIFO (First-in First-out) queue (buffer) on each end
- Can send both ways ("full duplex")
- Name for standard group of bits sent: "packet"
- Packet format and rules for communicating them ("protocol")

Simple *request/response* protocol and packet format:

header	payload
0/1	data/address

0: please send the data word at "address"
1: here is the data word you asked for.

Questions about simple example

What if more than 2 computers want to communicate?

- Need an interconnect? Need computer address field in packet?

What if the machines are far away?

- WAN vs LAN

How do multiple machines share the interconnect?

- multiple paths? arbitration? congestion control?

What if a packet is garbled in transit?

- Add error detection field in packet?

What if a packet is lost?

- More elaborate protocols to detect loss?

What if multiple processes per machine?

- one queue per process? separate field in packet to identify process?

<u>Warning: You are entering a buzzword-rich environment!!!</u>

Generic network

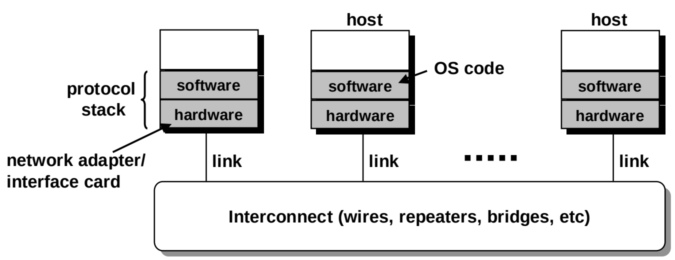

Protocols

A *protocol* defines the format of packets and the rules for communicating them across the network.

Different protocols provide different levels of service:

- simple error correction (ethernet)
- uniform name space, unreliable best-effort datagrams (host-host) (IP)
- reliable byte streams (TCP)
- unreliable best-effort datagrams (process-process) (UDP)
- multimedia data retrieval (HTTP)

<u>Crucial idea: protocols leverage off of the capabilities of other protocols.</u>

Protocol layering

Protocols provide specialized services by relying on services provided by lower-level protocols (i.e., they *leverage* lower-level services).

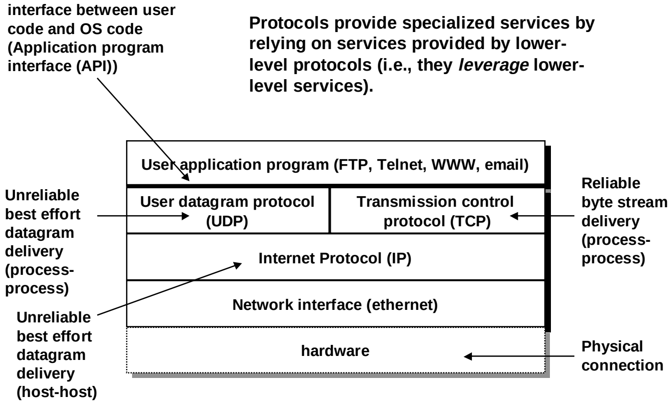

Encapsulation

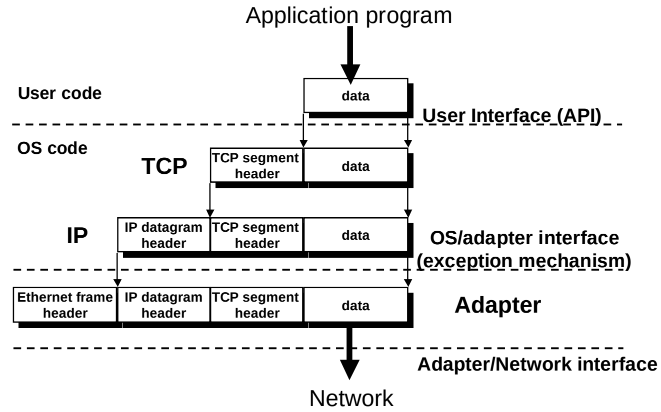

Transmission media

twisted pair: (10 Mb/s at 5 km)

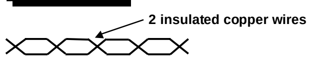

fiber: (100-200 Gb/s at 1 km)

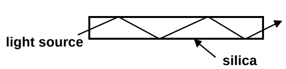

coaxial cable: (1-2 Gb/s at 1 km)

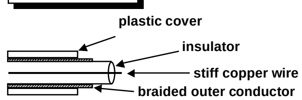

station wagon full of mag tapes hurtling down the highway every hour:

(15 Gb/s at 1 hour)

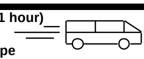

7 GBytes/tape
1000 tapes/station wagon (50x50x50cm)
7,000 GBytes total
7,000 GBytes/3600 seconds = 15 Gb/s

$5/tape reused 10 times -> $500 tape cost
$200 for shipping ->10 cents /GByte

Basic network types

System area network (SAN)
- same room (meters)
- 300 MB/s Cray T3E

Local area network (LAN)
- same bldg or campus (kilometers)
- 10 Mb/sEthernet
- 100 Mb/s Fast Ethernet
- 100 Mb/s FDDI
- 150 Mb/s OC-3 ATM
- 622 Mb/s OC-12 ATM

Metropolitan area network (MAN)
- same city (10's of kilometers)
- 800 Mb/s Gigabit Nectar

Wide area network (WAN)
- nationwide or worldwide (1000's of kilometers)
- telephone system
- 1.544 Mb/s T1 carrier
- 44.736 Mb/s T3 carrier

We'll look at WANs and LANs.

AT&T Telephone Hierarchy

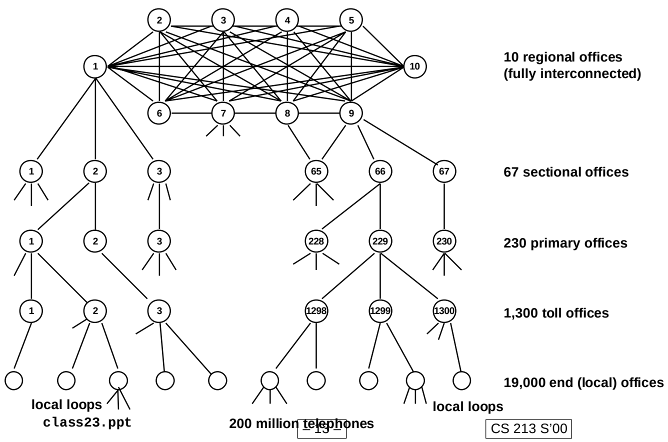

Connecting distant computers with modems

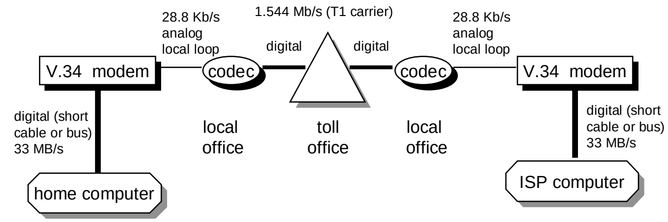

Modulating digital signals

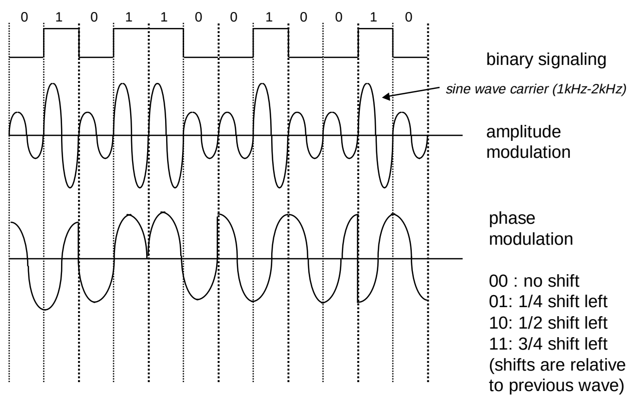

00 : no shift
01: 1/4 shift left
10: 1/2 shift left
11: 3/4 shift left
(shifts are relative to previous wave)

Quadrature amplitude modulation (QAM)

Modern modems use a combination of of amplitude and phase modulation to encode multiple bits per "symbol", i.e. amplitude/phase pair.

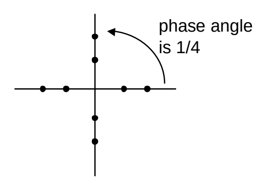

3 bits/symbol QAM modulation
(8 symbols)

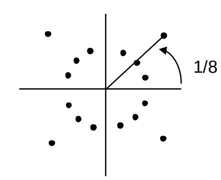

4 bits/symbol QAM modulation
(16 symbols)

Conventional Modems

MOdulate: convert from digital to analog
DEModulate: convert from analog to digital

modem standards:

type	symbols/sec	bits/symbol	Kb/s
v.32	2400	4	9.6
v.32.bis	2400	6	14.4
v.34	3200	9	28.8

Theoretical limit for modulated signals is approx 35 Kb/s:
Shannon's law: max bits/s = H log2(1 + S/N), where H is bandwidth and S/N is signal to noise ratio. For phone network, H ~ 3,600 and 10*log10(S/N) ~ 30 dB, which implies S/N ~ 1000. Thus max rate is ~35 Kb/s.

So what's the deal with 56K modems?

T1 carrier (1.544 Mb/s)

Digital part of phone system based on the T1 carrier:

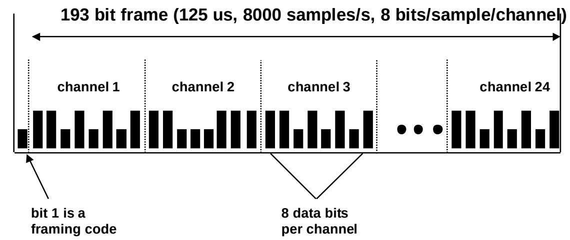

Each channel has a data rate of 8000 samples/s * 8 bits/channel = 64 Kb/s

56KB "Modems"

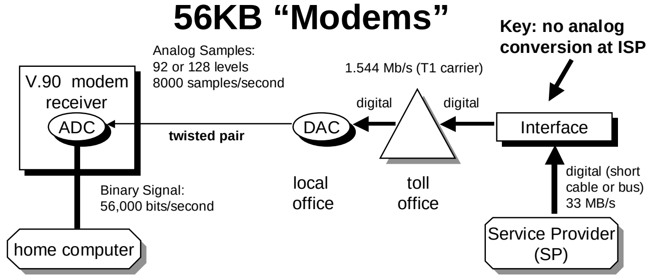

- **Asymmetric: home to SP uses conventional v.34 modem**
- **SP has digital connection into phone system**
  - Channel sending 8000 samples / second, up to 8-bits/sample
- **DAC encodes each sample with 92 or 128 voltage levels**
  - Not enough precision on analog side to handle finer resolution
- **Receiver converts samples back to digital values**
  - Must match frequency & phase of senders DAC
  - Establish using "training" signals from sender

Comparison with other connection technologies

technology	media access	downstream	upstream
modem	**dedicated**	**56 Kb/s**	**33 Kb/s**
ADSL (Assym. Digital Subscriber Line)	**dedicated**	**1.5 -- 9 Mb/s**	**16 -- 640 Kb/s**
cable modem	**shared**	**27 -- 56 Mb/s**	**3 Mb/s**

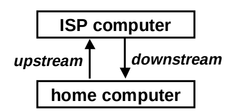

Basic Internet components

An Internet *backbone* is a collection of routers (nationwide or worldwide) connected by high-speed point-to-point networks.

A *Network Access Point (NAP)* is a router that connects multiple backbones (sometimes referred to as *peers*).

Regional networks are smaller backbones that cover smaller geographical areas (e.g., cities or states)

A *point of presence (POP)* is a machine that is connected to the Internet.

Internet Service Providers (ISPs) provide dial-up or direct access to POPs.

The Internet circa 1993

In 1993, the Internet consisted of one backbone (NSFNET) that connected 13 sites via 45 Mbs T3 links.

- Merit (Univ of Mich), NCSA (Illinois), Cornell Theory Center, Pittsburgh Supercomputing Center, San Diego Supercomputing Center, John von Neumann Center (Princeton), BARRNet (Palo Alto), MidNet (Lincoln, NE), WestNet (Salt Lake City), NorthwestNet (Seattle), SESQUINET (Rice), SURANET (Georgia Tech).

Connecting to the Internet involved connecting one of your routers to a router at a backbone site, or to a regional network that was already connected to the backbone.

The Internet backbone (circa 1993)

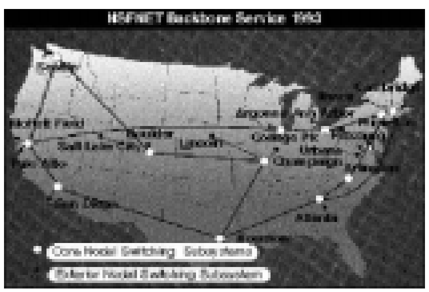

Current NAP-based Internet architecture

In the early 90's commercial outfits were building their own high-speed backbones, connecting to NSFNET, and selling access to their POPs to companies, ISPs, and individuals.

In 1995, NSF decommissioned NSFNET, and fostered creation of a collection of NAPs to connect the commercial backbones.

Currently in the US there are about 50 commercial backbones connected by ~12 NAPs (peering points).

Similar architecture worldwide connects national networks to the Internet.

Internet connection hierarchy

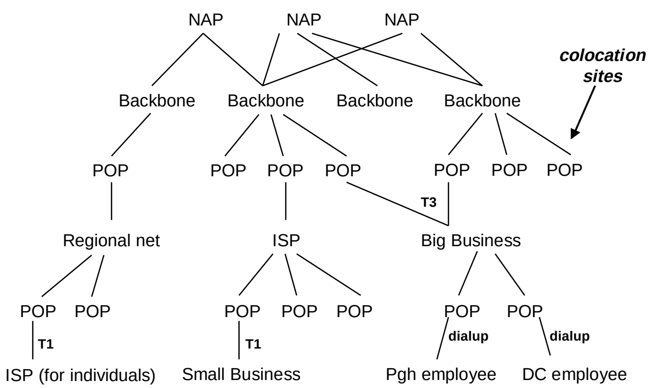

Network access points (NAPs)

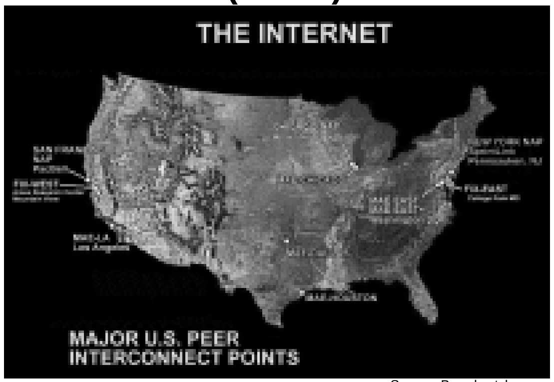

Source: Boardwatch.com

MCI/WorldCom/UUNET Global Backbone

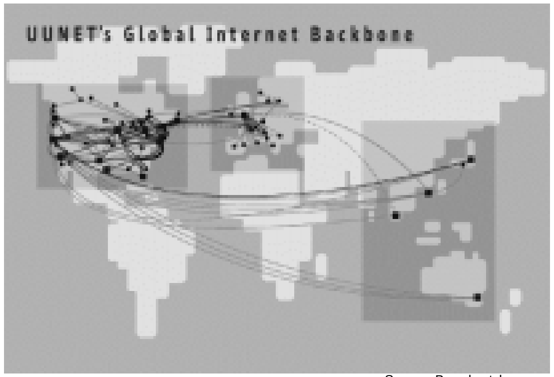

Source: Boardwatch.com

Cost of Frame Relay connections

56 Kbps frame relay:

Availability: All U.S. backbone cities
Setup: $495

Monthly: $595

Recommended Equipment:
Cisco 2524 router with 5IN1 Card &
Kentrox 56K CSU/DSU: Total $2,395

Source: Boardwatch.com (MCI/Worldcom)

Cost of T1 connections

Burstable 1.544 Mbps T-1 service:

Monthly charge based on 95 percent usage level
Availability: All U.S. backbone cities
Average Installation Time: 4-6 weeks
Setup: $5,000
Recommended Equipment: Cisco Integrated T-1
CSU/DSU - $995, Cisco 2524 router - $1,950

Bandwidth	Monthly
0-128 Kbps	$1,295
128 Kbps-256 Kbps	$1,895
256 Kbps-384 Kbps	$2,495
384 Kbps-512 Kbps	$2,750
512 Kbps-1.544 Mbps	$3,000

95/5 pricing model: sample bandwidth every 5 minutes. Set monthly price for smallest bandwidth that is greater than 95% of the samples.

Source: Boardwatch.com (MCI/Worldcom)

Cost of T3 connections

Burstable 45 Mbps T-3 service:

Monthly price based on 95th percentile usage level.
Availability: All U.S. backbone cities
Average Install Time: 8-10 weeks
Setup: $6,000

Bandwidth	Monthly
up to 6 Mbps	$12,000
6.01 Mbps-7.5 Mbps	$14,000
7.51 Mbps-9 Mbps	$17,000
9.01 Mbps-10.5 Mbps	$19,000
10.51 Mbps-12 Mbps	$22,000
12.01 Mbps-13.5 Mbps	$26,000
3.51 Mbps-15 Mbps	$29,000
15.01 Mbps-16.5 Mbps	$32,000
16.51 Mbps-18.01 Mbps	$37,000
18.01 Mbps-19.5 Mbps	$43,000
19.51 Mbps-21 Mbps	$48,000
21.01 Mbps-45 Mbps	$55,500

Recommended Equipment: Cisco 7204 router

Source: Boardwatch.com (MCI/Worldcom)

Ethernet

History
- **1976- proposed by Metcalfe and Boggs at Xerox PARC**
- **1978 - standardized by Xerox, Intel, DEC**

Bandwidth
- **10 Mbits/sec (old) , 100 Mbits/sec (standard), 1 Gbits/s (new)**

Key features
- **broadcast over shared bus (the ether)**
  - no centralized bus arbiter
  - each adapter sees all bits
- **each adapter has a unique (for all time!) 48-bit address**
- **variable length frames (packets) (64 - 1518 bytes)**

Ethernet cabling

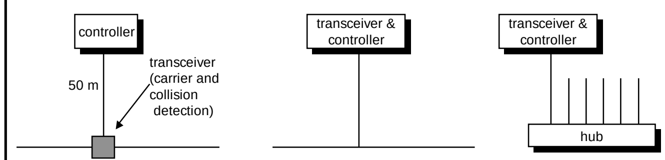

10Base5 ("thick ethernet") 10Base2 ("thin ethernet") 10Base-T

name	cable	max segment	nodes/segment	advantages
10Base5	thick coax	500 m	100	good for backbones
10base2	thin coax	200 m	30	cheapest
10Base-T	**twisted pair**	**100 m**	**1024**	**easy maintenance**
10Base-F	fiber	2000 m	1024	best between bldgs

Repeaters vs bridges

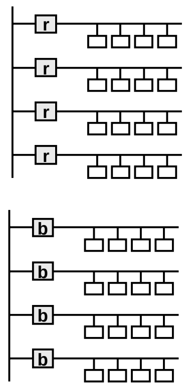

Repeaters directly transfer their inputs to their outputs.

Bridges maintain a cache of hosts on their input segments.

Selectively transfer packets from their inputs to their outputs.

Ethernet packet (frame) format

64 - 1518 bytes

Preamble	Dest addr	Src addr	Frame type	Payload	CRC	Postamble
64 bits	48 bits	48 bits	16 bits	368-12000 bits	32 bits	8 bits

visible from the host

Preamble: 101010101 (synch)

dest and src addr: unique ethernet addresses

payload: data

CRC: cyclic redundancy check (error detection/correction)

Ethernet receiving algorithm

Ethernet adapter receives all frames.

Accepts:

- frames addressed to its own address
- frames addressed to broadcast address (all 1's).
- frames addressed to multicast address (1xxx...), if it has been instructed to listen to that address
- all frames, if it has placed in promiscuous mode

Passes to the host only those packets it accepts.

Ethernet sending algorithm (CSMA/CD)

Problem: how to share one wire without centralized control.

Ethernet solution: Carrier Sense Multiple Access with Collision Detection (CSMA/CD):

1. Adapter has frame to send and line is idle:
- then send frame immediately

2. When adapter has frame to send and line is busy:
- wait for line to become idle, then send frame immediately.

3. If "collision" (simultaneous sends) occurs during transmission:
- send at least 1024 bits
- send "jam signal" to notify receivers
- wait some period of time (binary exponential backoff)
- retry

Binary exponential backoff

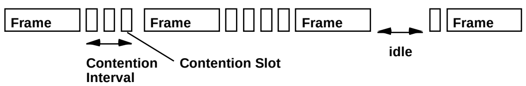

Binary exponential backoff algorithm:

- after 1st collision, wait 0 or 1 slots, at random.
- after 2nd collision, wait 0, 1, 2, 3 slots at random.
- etc up to 1023 slots.
- after 16 collisions, exception.

Why the 64 byte minimum packet size?

Assume propagation delay from A to B is tau microseconds (us).

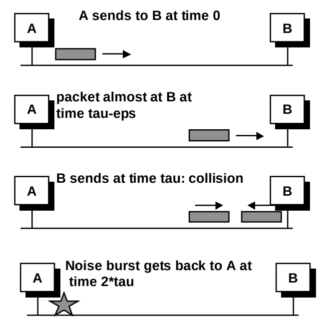

Conclusion: Senders must take more than 2*tau seconds to send their packets.

For ethernet, 2*tau is specified by standard (2500 m cable w/ 4 repeaters) to be 51.2 us, which at 10 Mb/s is 512 bit times, or 64 bytes.

Rough estimate: propagation through copper is about 20 cm/ns. With a 2500 m cable, tau is 12.5 us and 2*tau is 25 us.

As speeds increase there are two possibilities:
1. increase packet sizes
2. decrease maximum cable length

Neither is particularly appealing.

Ethernet pros and cons

Pros:

- simple
- robust
- cheap ($50/adapter in 1998)

Cons:

- no quality of service guarantees
  - OK for data
  - not OK for real-time bit streams like video or voice
- fixed bit rate
  - not keeping up with faster processors
  - workstation can produce data at 10-50 MBytes/sec
- prone to congestion
  - processors getting faster
  - bridged Ethernets can help some