15-213

Network technology

November 12, 1998

Topics

- Overview
- Telephone system
- Ethernet
- ATM

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

Computer system

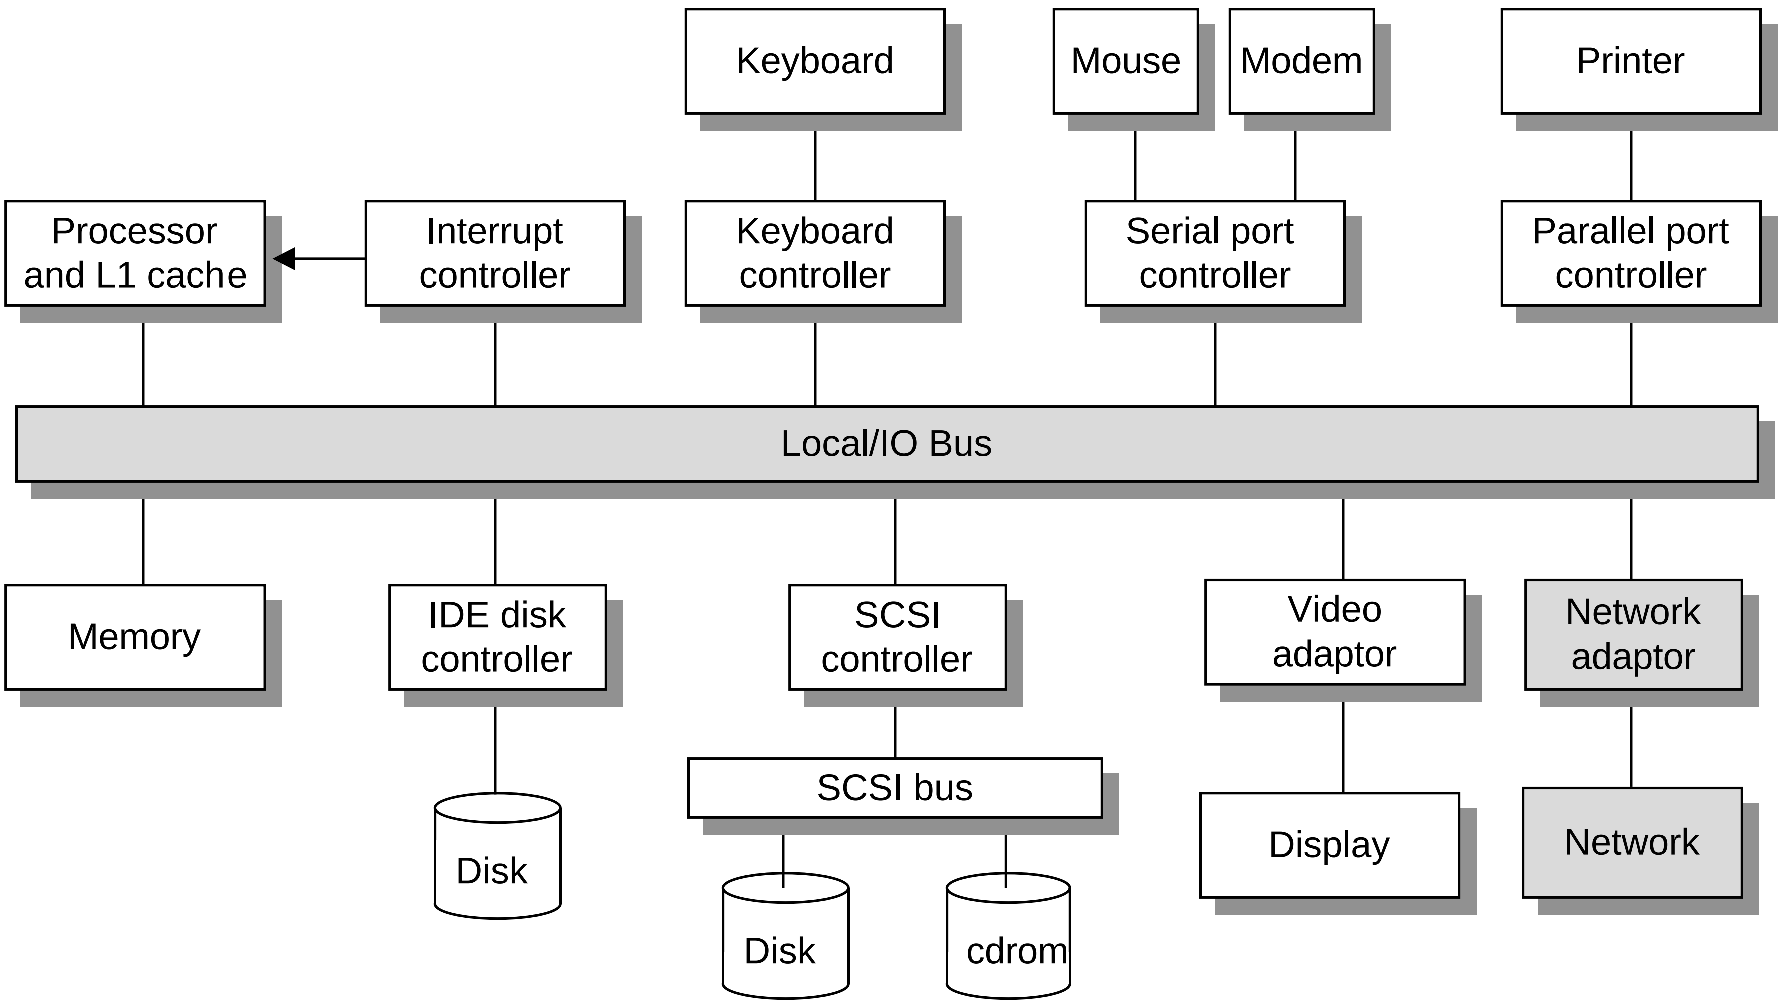

Simple example

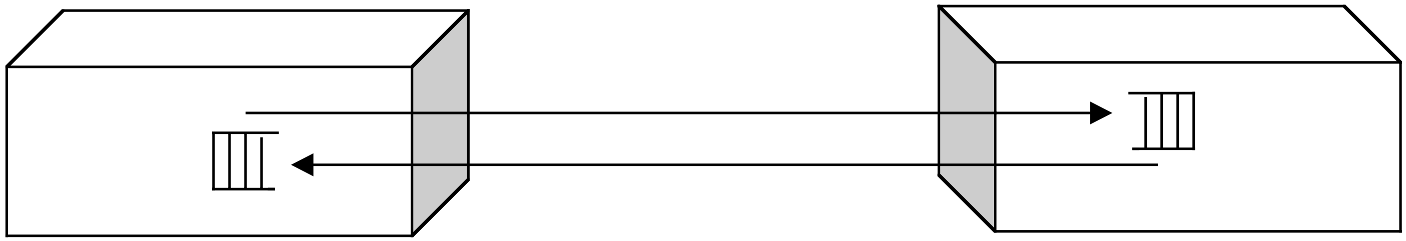

Starting Point: Want to send bits between 2 computers

- **FIFO queue 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 computer address field in packet?**

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/machine?

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

Warning: You are entering a buzzword-rich environment!!!

Generic network

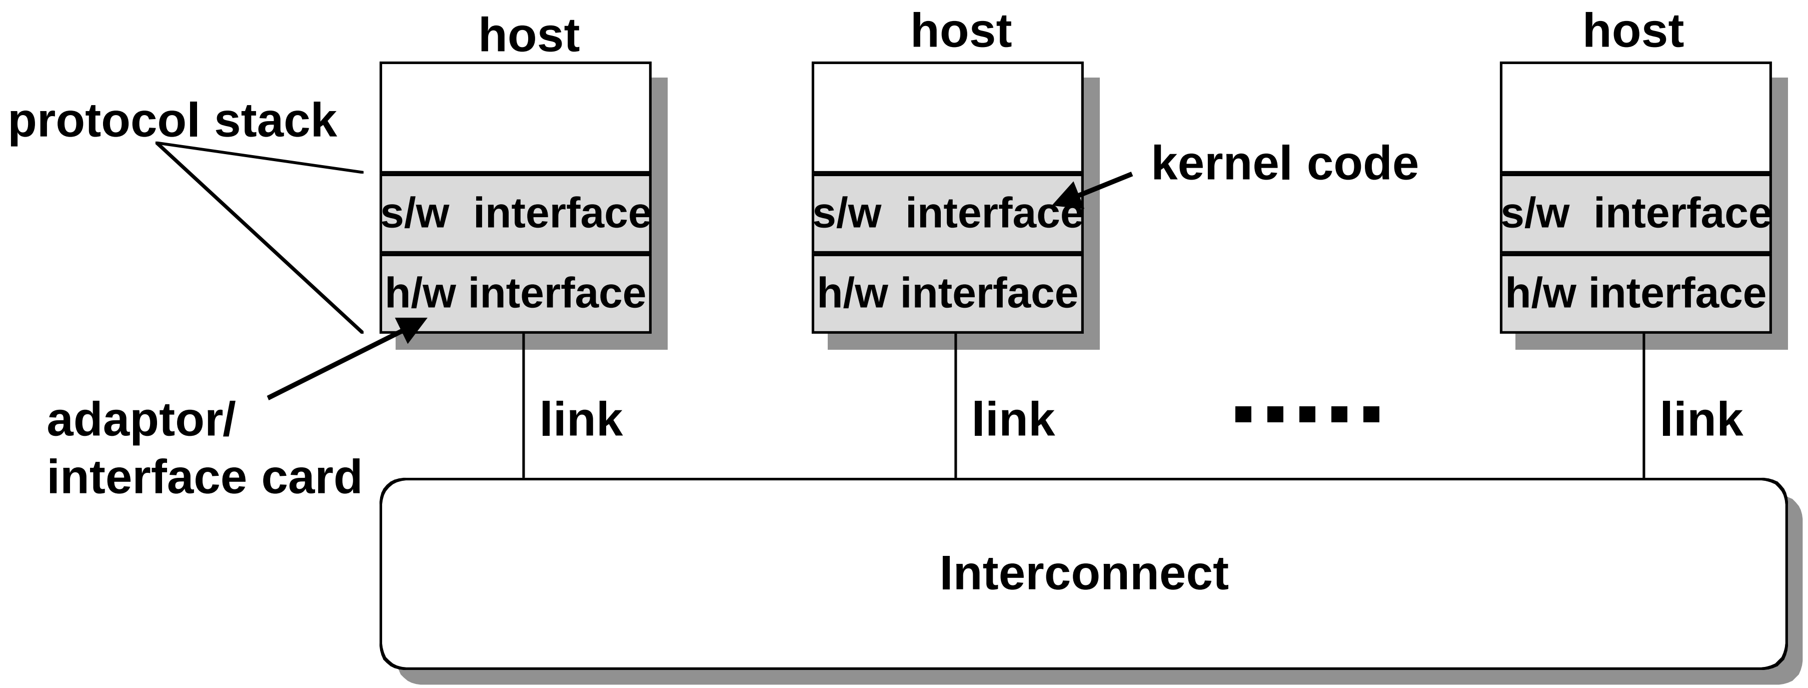

Protocols

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)**

Protocol layering

Protocols provide specialized services by building on services provided by other protocols.

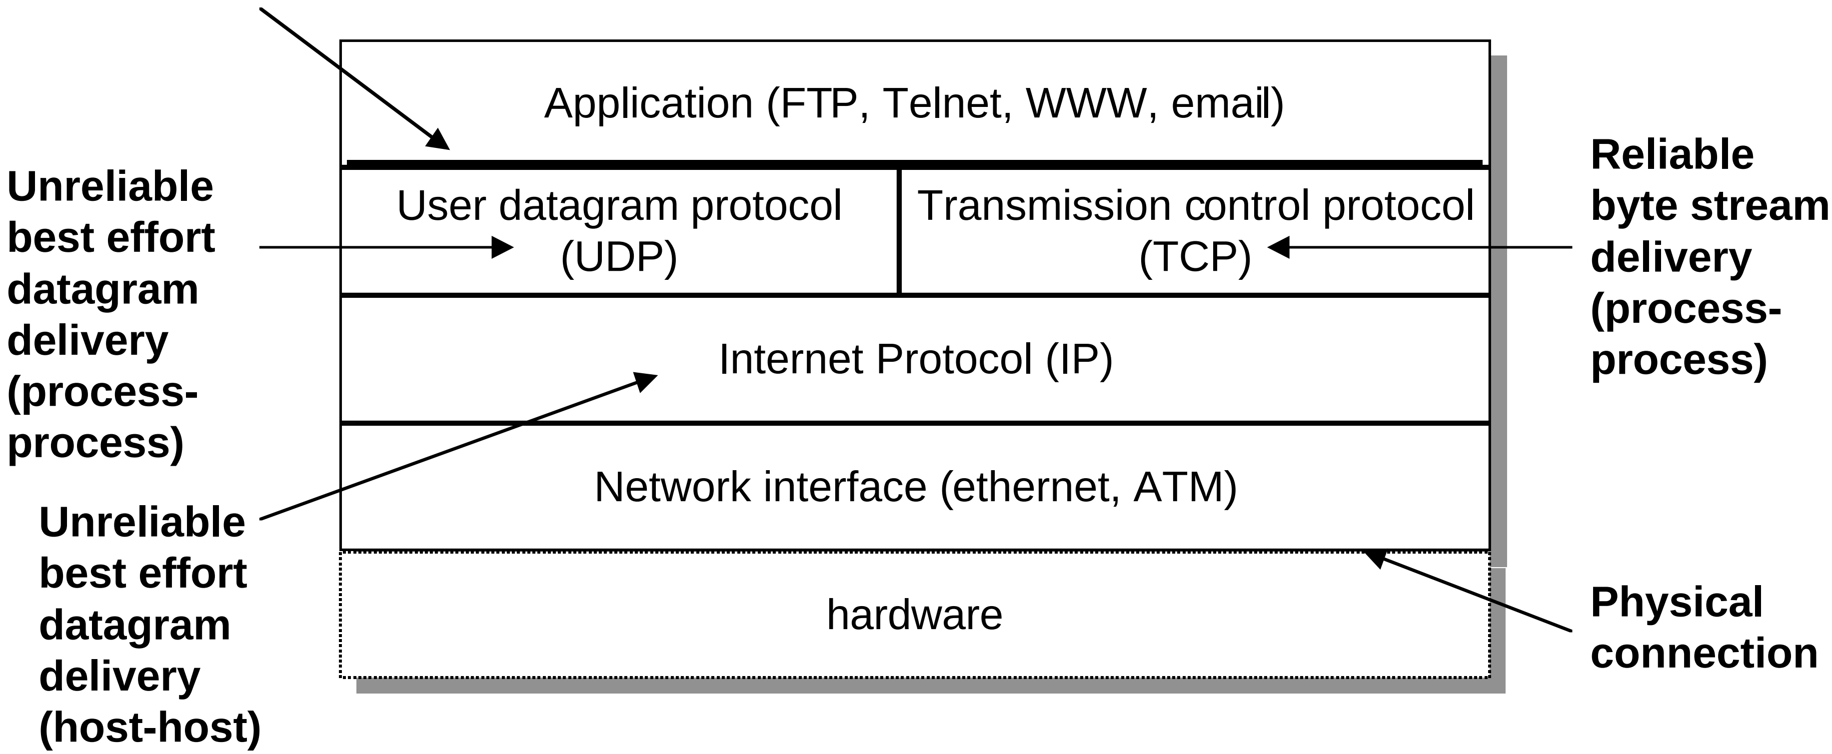

Encapsulation

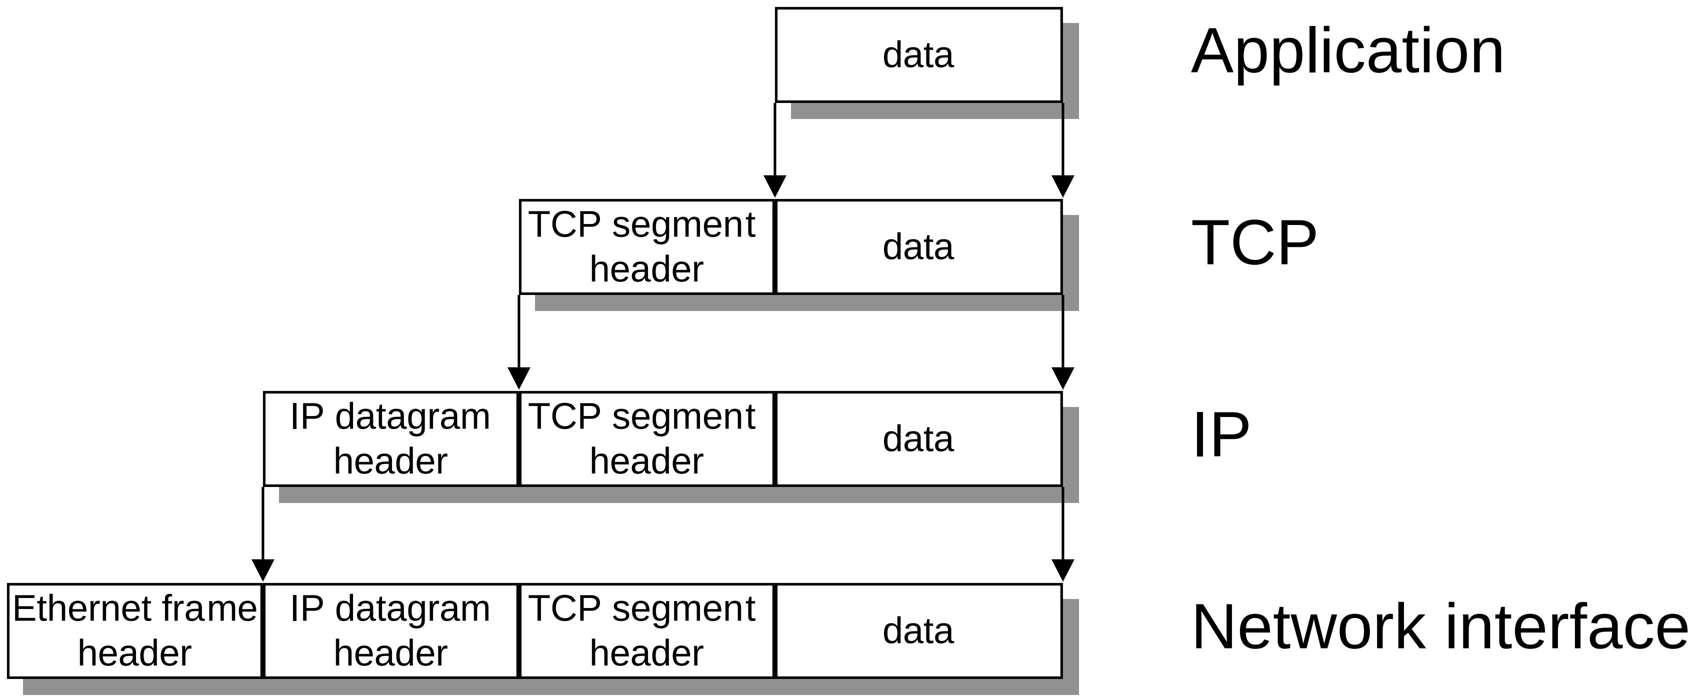

Protocol stacks

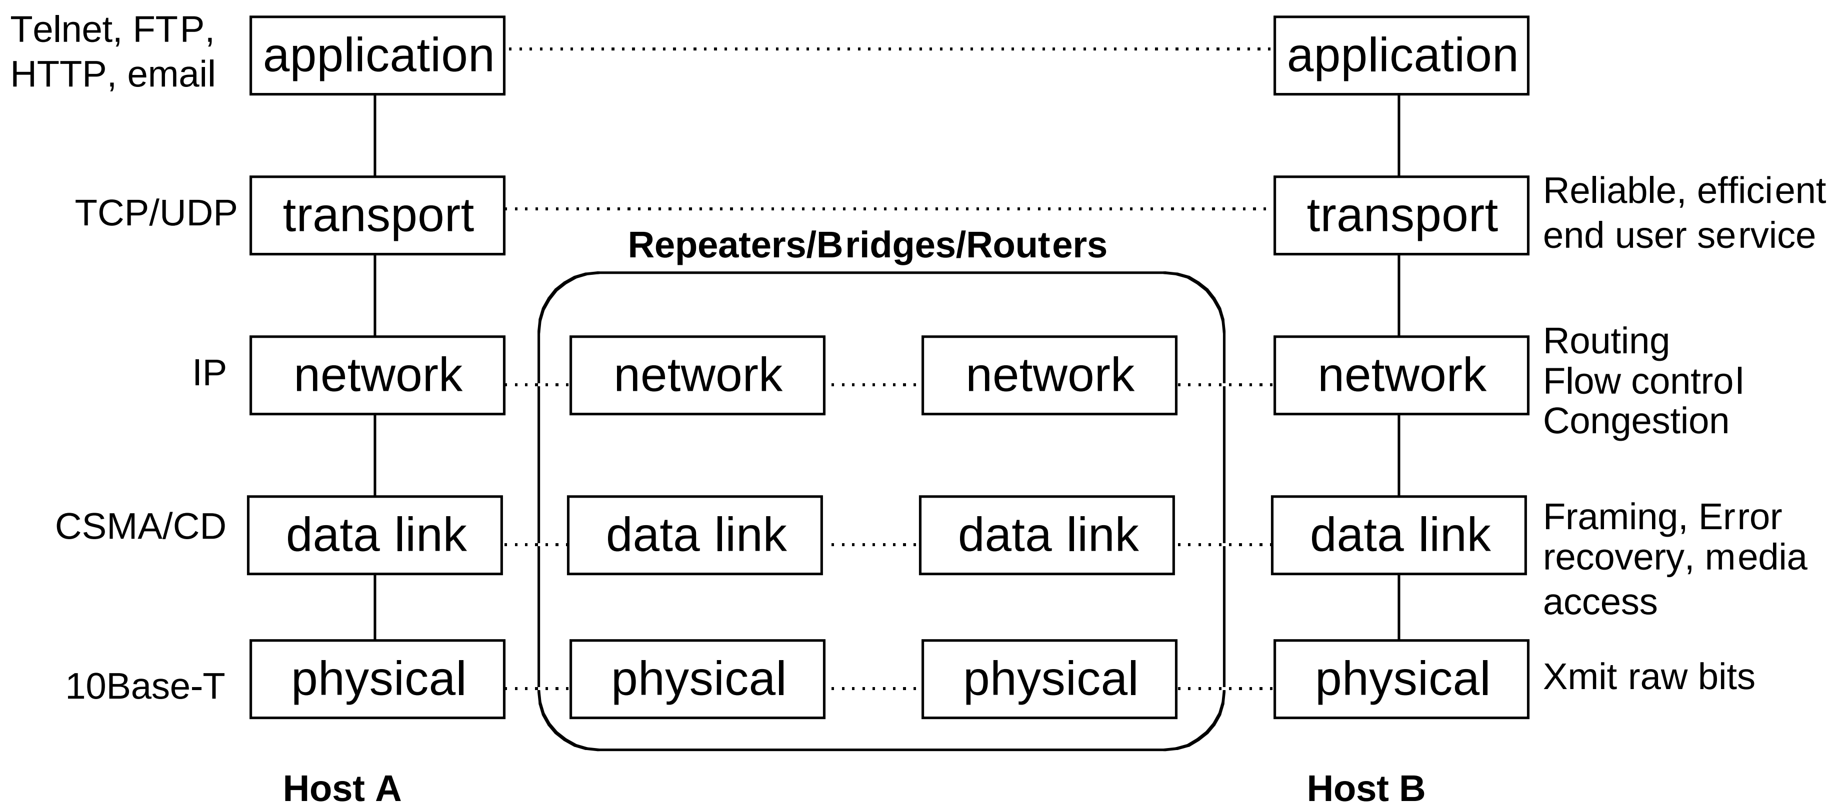

Transmission media

twisted pair: (1-2 Mb/s at 1 km)

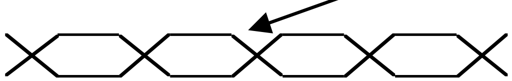

fiber: (100-200 Gb/s at 1 km)

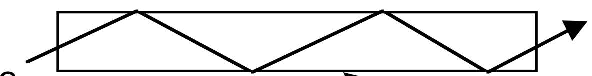

coaxial cable: (1-2 Gb/s at 1 km)

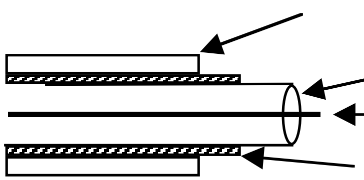

station wagon full of mag tapes hurtling down the highway:

(15 Gb/s at 1 hour)

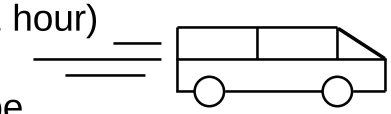

7 GBytes/tape
1000 tapes/station wagon (50x50x50cm)
7,000 GBytes total
7,000 GBytes/3600 minutes = 15 Gb/s

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

Shared vs switched media

Shared media (e.g., Ethernet)

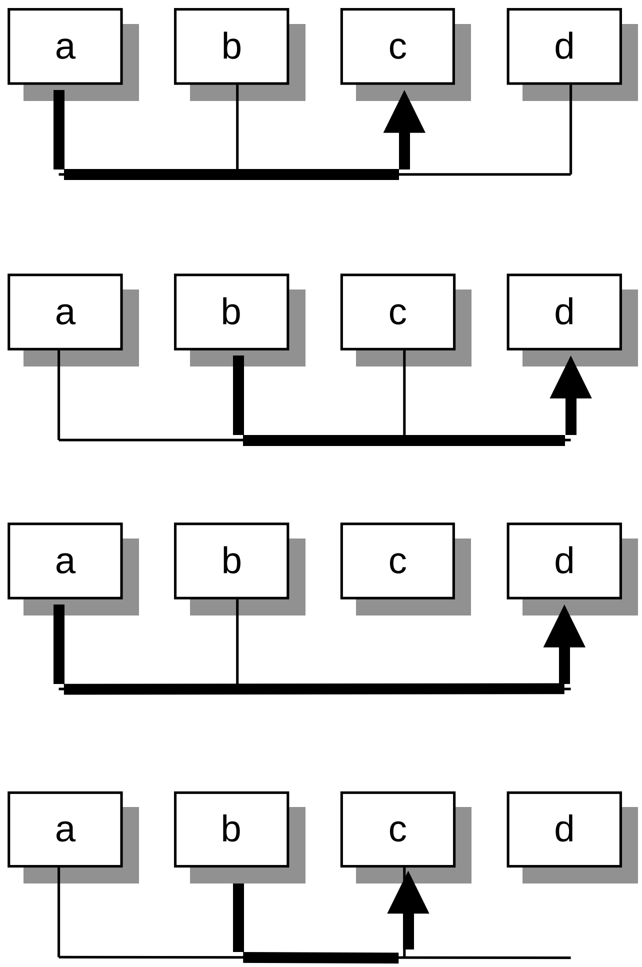

Switched media (e.g., ATM)

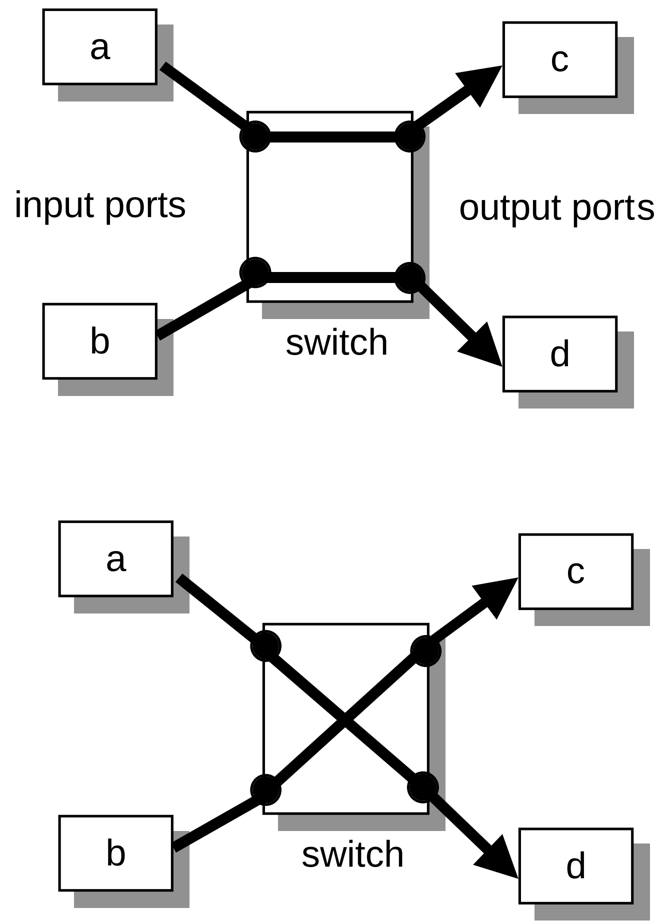

Network performance measures

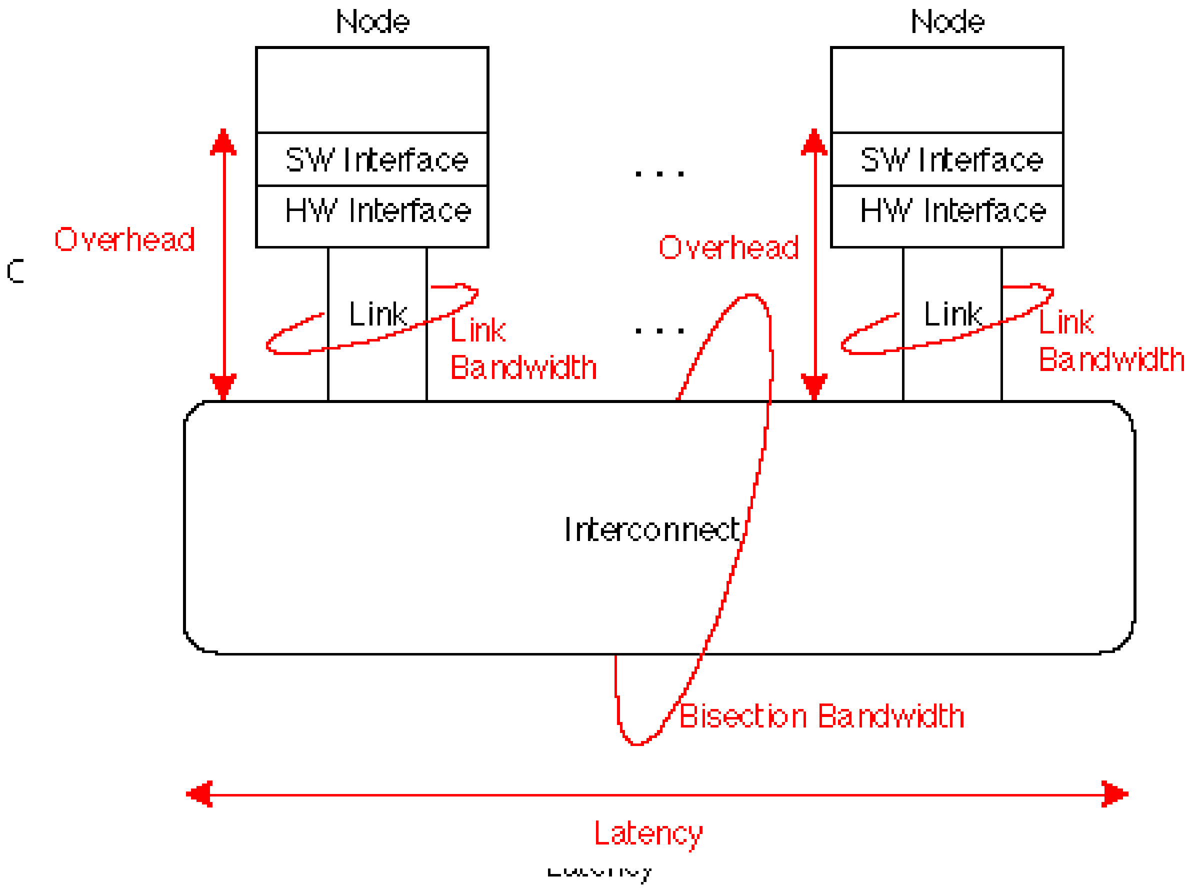

Example performance measures

Interconnect	SAN	LAN	WAN
Example	CM-5	Ethernet	ATM
Bisection BW	N x 5MB/s	1.125 MB/s	N x 10 MB/s
Int./Link BW	20 MB/s	1.125 MB/s	10 MB/s
Latency	5 µsec	15 µsec	50 to 10,000 µs
HW Overhead to/from	0.5/0.5 µs	6/6 µs	6/6 µs
SW Overhead to/from	1.6/12.4 µs	200/241 µs	207/360 µs
		(TCP/IP on LAN/WAN)	

Importance of Overhead (+ Latency)

Ethernet / SS10:	**9 Mb/s BW,**	**900 μsecs ovhd**
ATM Synoptics**:**	**78 Mb/s BW,**	**1,250 μsecs ovhd.**

NFS trace over 1 week: 95% msgs < 200 bytes

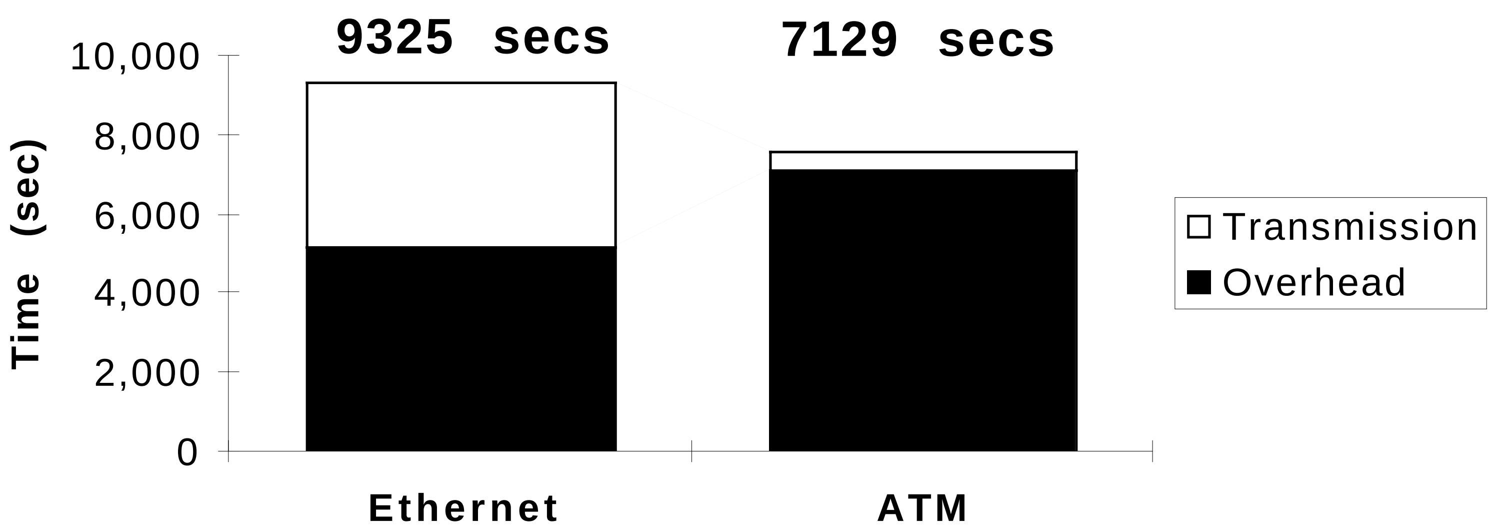

- Link bandwidth is as misleading as MIPS

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

AT&T Telephone Hierarchy

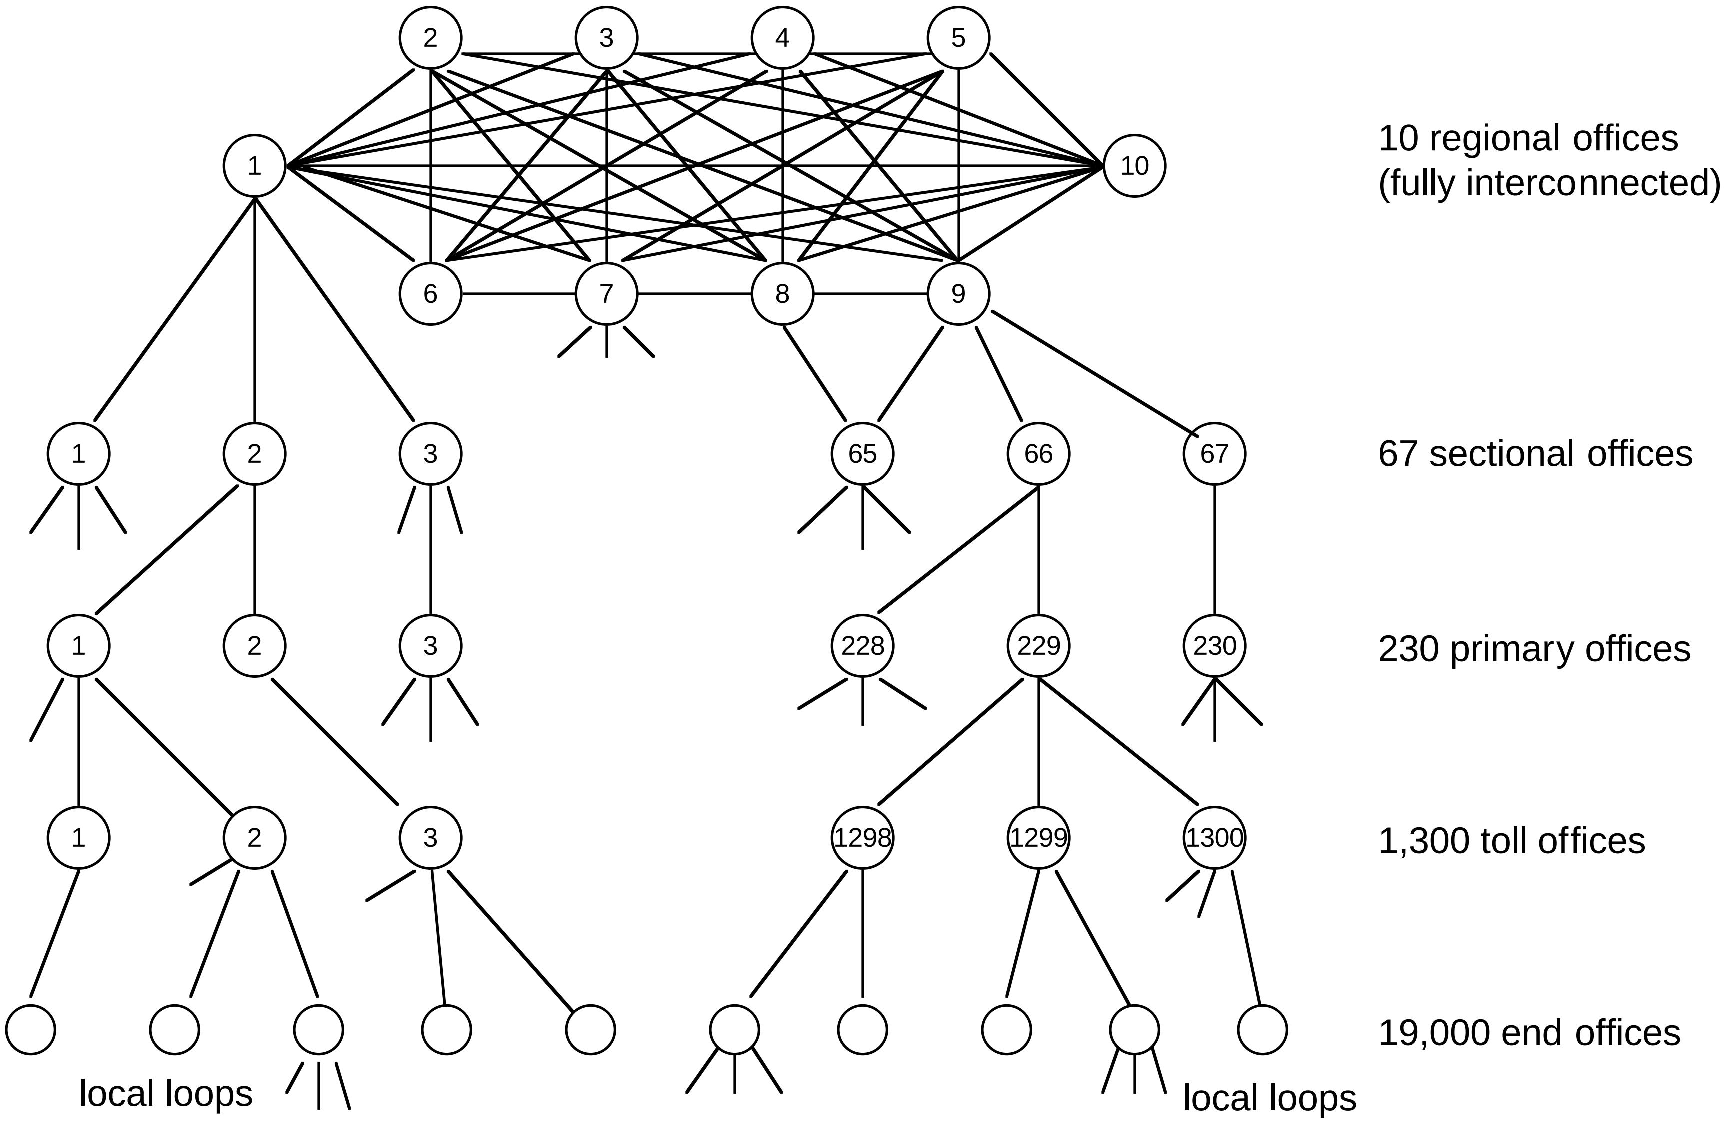

Computer-to-computer calls

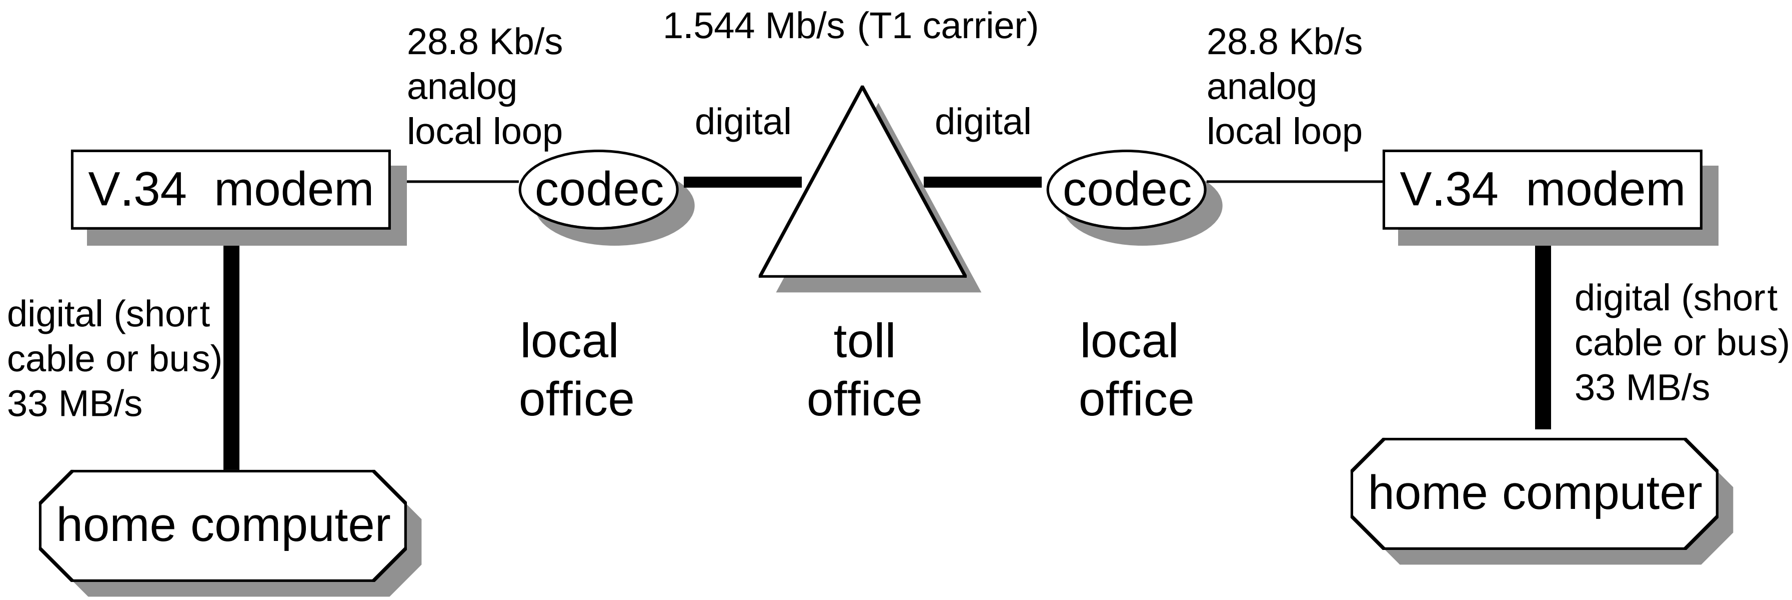

Modulating digital signals

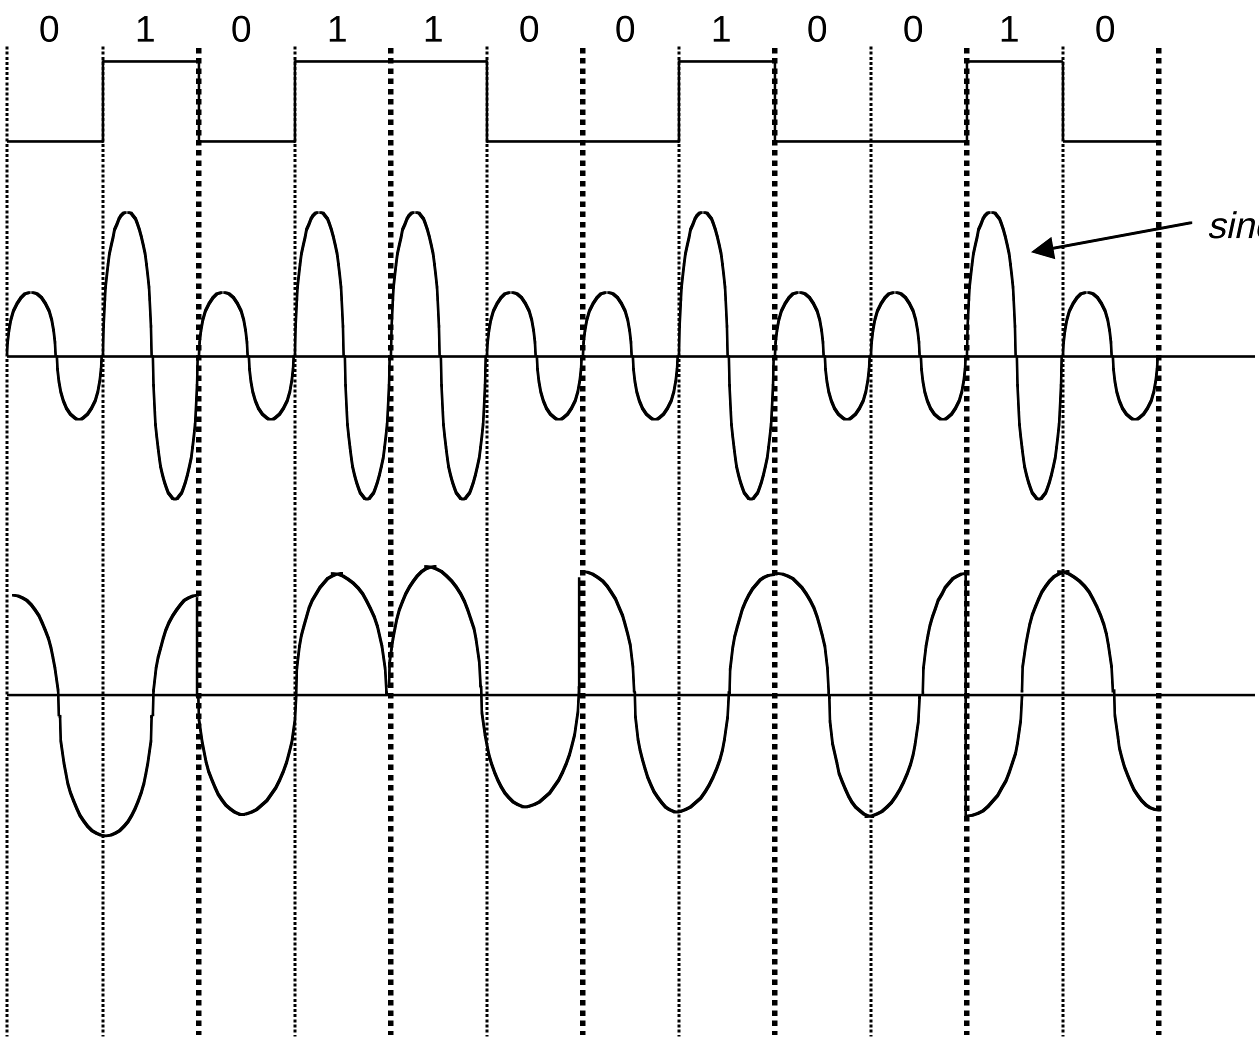

00 : no shift
01: 1/4 shift left
10: 1/2 shift left
11: 3/4 shift left
(shifts are relative to previous wave)

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.

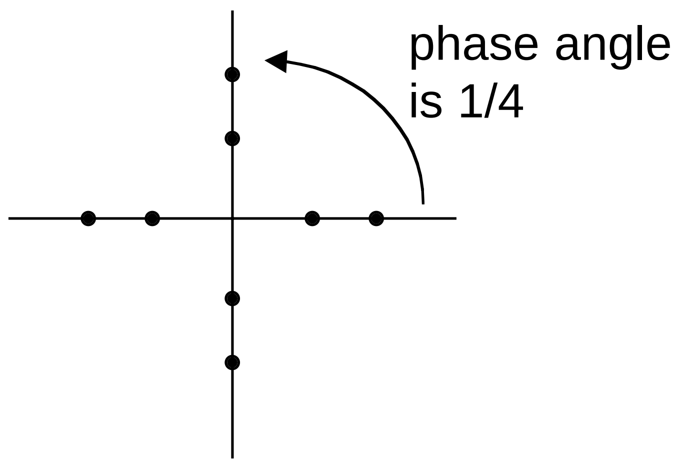

3 bits/symbol QAM modulation
(8 symbols)

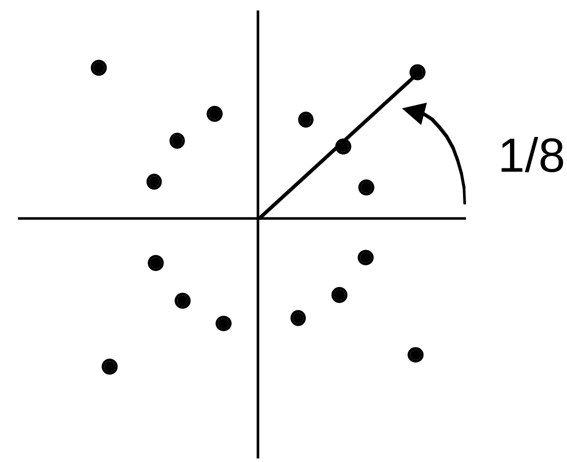

4 bits/symbol QAM modulation
(16 symbols)

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 bdwdth and S/N is signal to noise ratio. For phone network, H~3,600 and S/N is 30 dB. Thus max rate is ~35 Kb/s.

T1 carrier (1.544 Mb/s)

Digital part of phone system based on the T1 carrier:

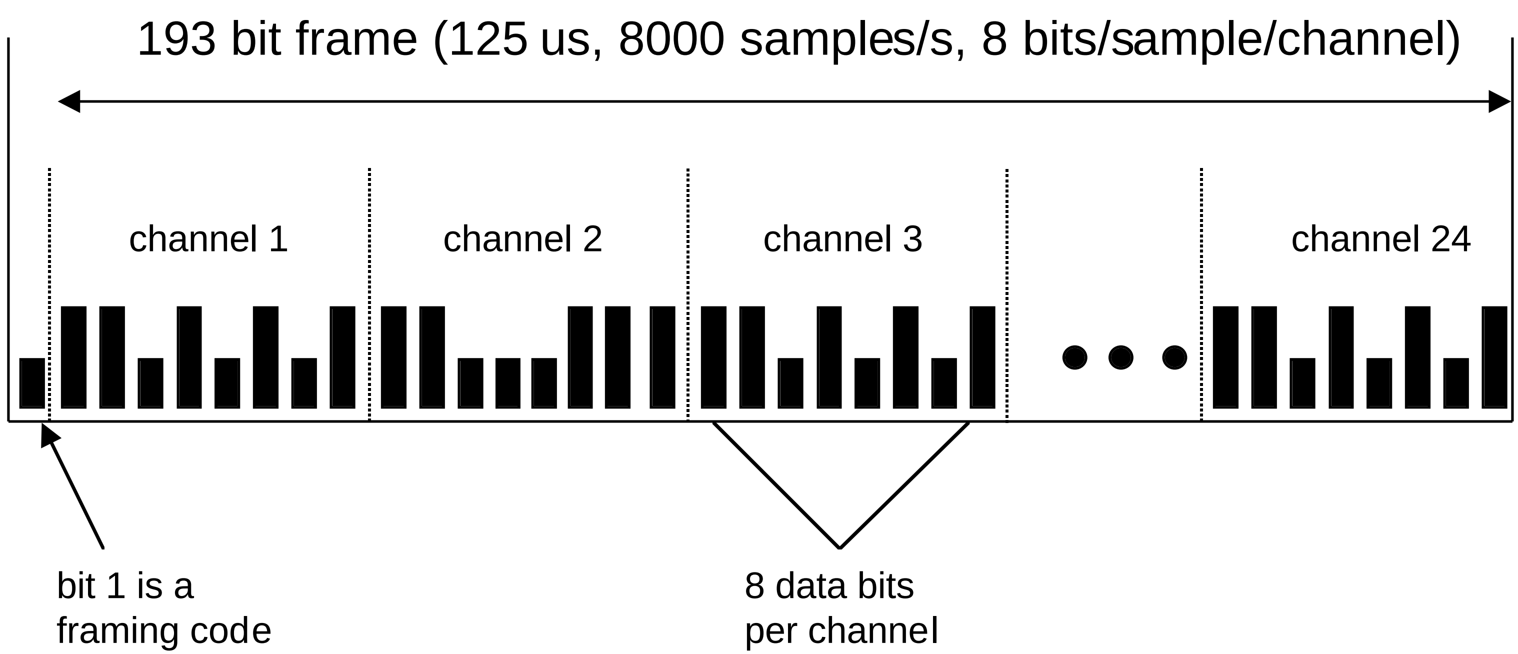

Each channel has a data rate
of 8000 samples/s * 8 bits/channel = 64 Kb/s

56KB “Modems”

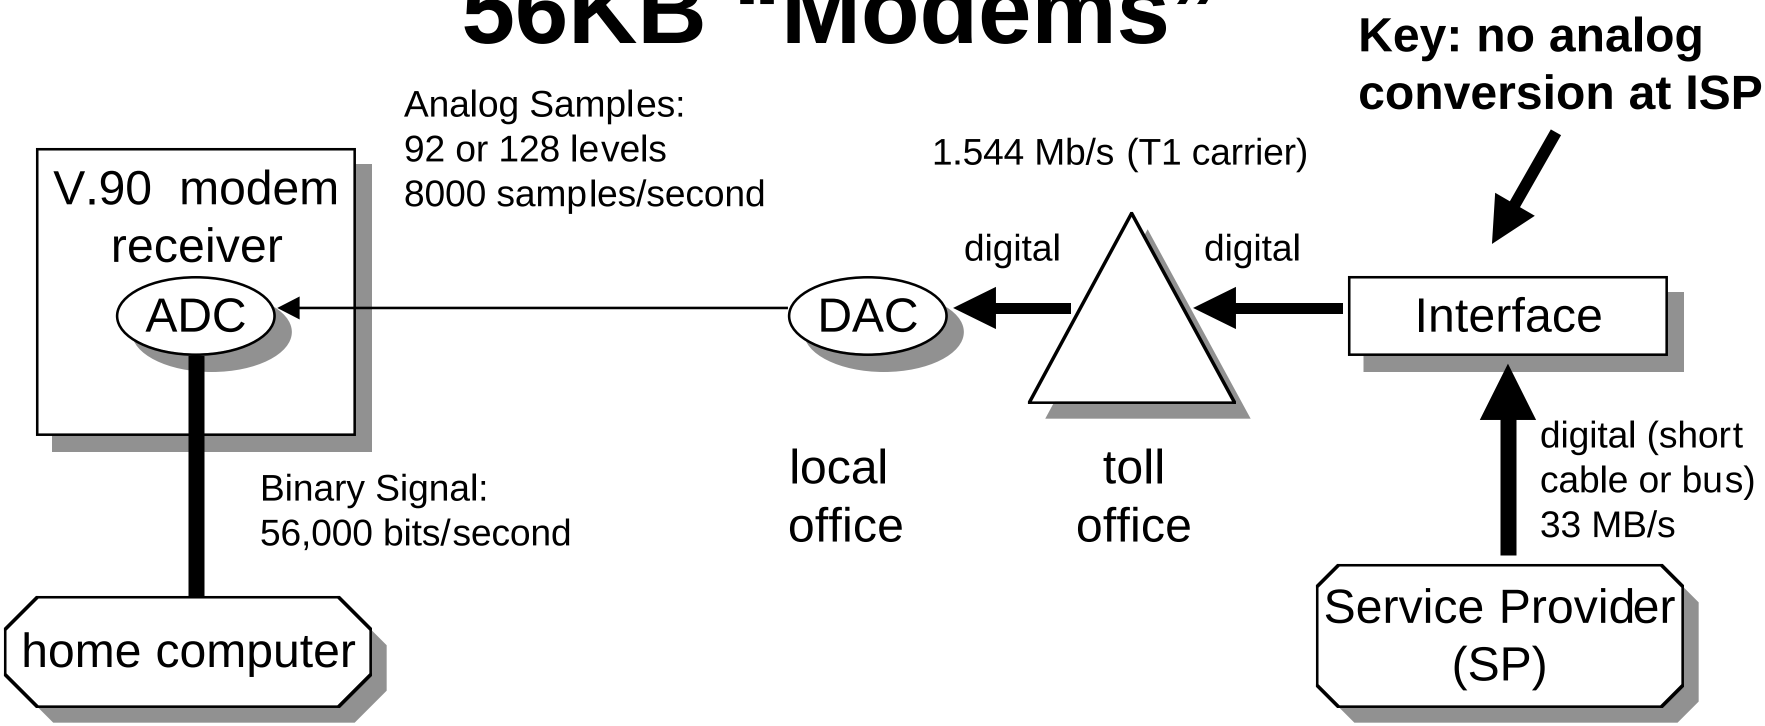

- **Asymmetric: home to SP uses conventional v.34 modem**
- **SP has digital connection into phone system**
  - Channel sending 8000 samples / second, up to 8-bits/sample
- **DAC encodes each sample with 92 or 128 voltage levels**
  - Not enough precision on analog side to handle finer resolution
- **Receiver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

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 (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

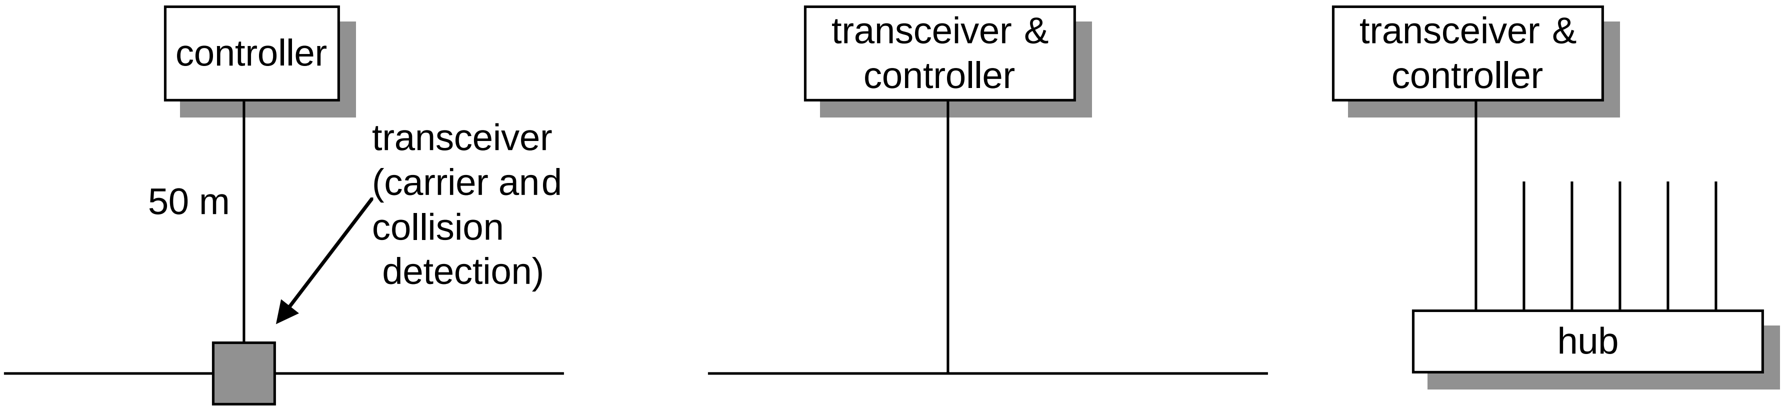

10Base5
("thick ethernet")

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

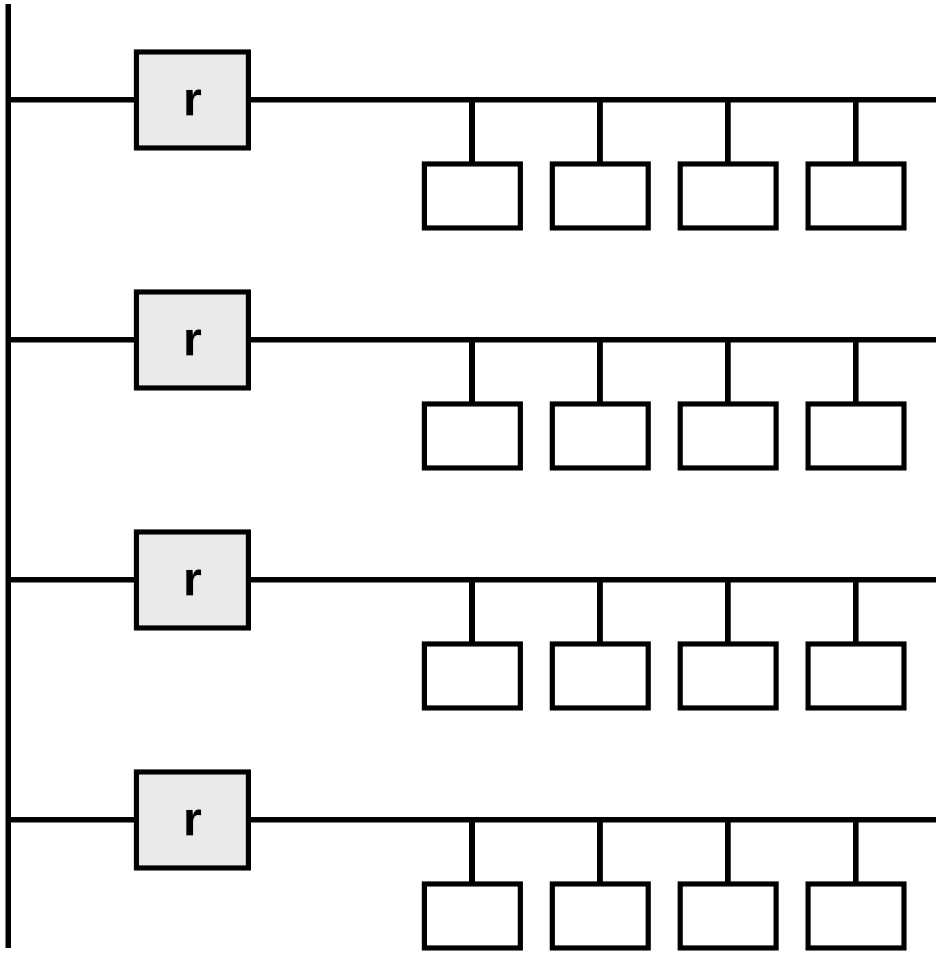

Repeaters directly transfer their inputs to their outputs.

Bridges

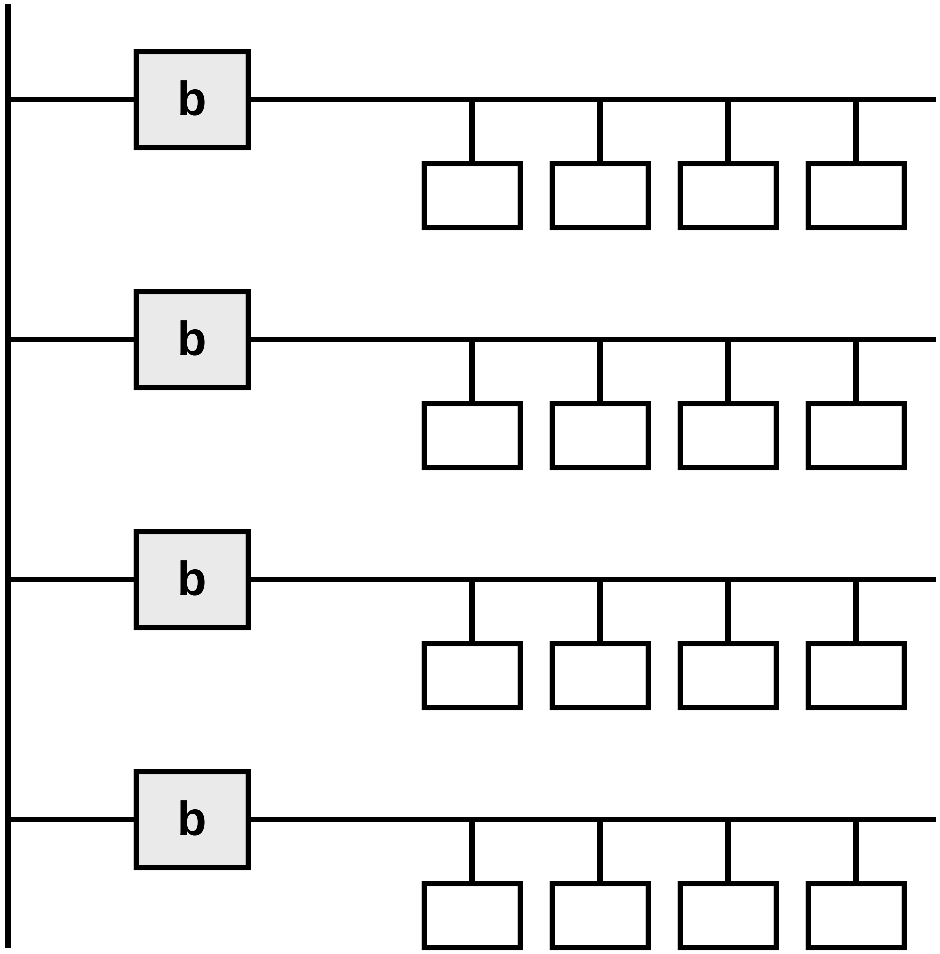

Bridges maintain a cache of hosts on their input segments.

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 adaptor 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. Adaptor has frame to send and line is idle:
   - then send frame immediately
2. When adaptor 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

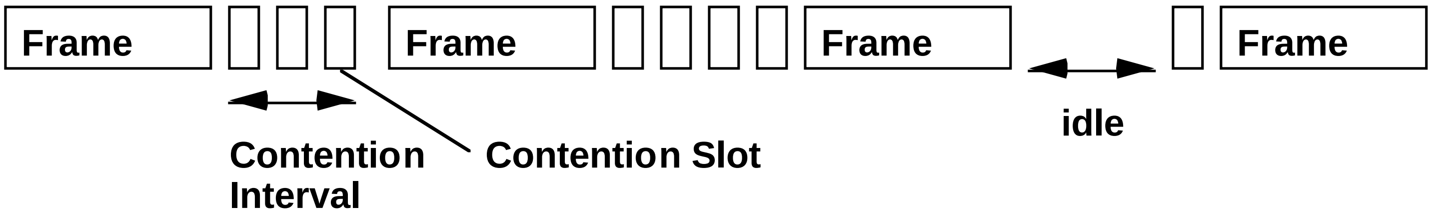

Binary exponential backoff algorithm:

- **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 propogation delay from A to B is tau.

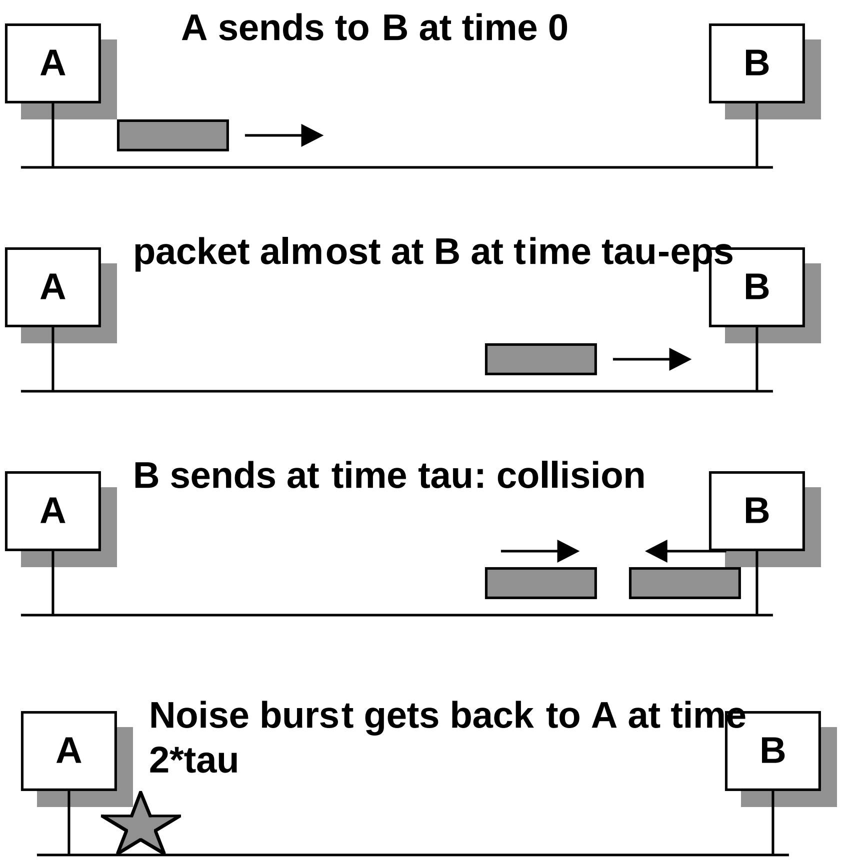

Conclusions: Senders must take more than 2*tau seconds to send their packets.

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

Rough estimate: propogation 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

Asynchronous transfer mode (ATM)

History

- **1988- proposed by international ATM forum**
- **telecomunications and computer vendors**

Goal:

- **mechanism for integrated transport of bit streams with different performance and reliability requirements (quality of service)**
  - video: 1.5 Mbits/sec, latency and variance sensitive, some bit loss OK
  - voice: 64 Kbits/sec, latency and variance sensitive, some bit loss OK
  - data: high data rates, latency and variance insensitive, bit loss not OK

ATM overview (cont)

Bandwidths

- **OC-1: 51.84 Mbits/sec**
- **OC-3: 155.52 Mbits/sec (current LAN rate)**
- **OC-12: 622.08 Mbits/sec (current LAN rate)**
- **OC-24: 1244.16 (Gigabit network)**

Key features:

- **virtual connections (VC's)**
  - allow bandwidth reservation
- **fixed cell (frame) size of 53 bytes**
  - simplifies high-speed switching
- **small cell size**
  - allows fine-grained allocation of network bandwidth

ATM cell format

53 bytes (fixed)

Generic flow ctl	VCI/ VPI	Payload type	Priority	Header checksum	Payload
4 bits	24 bits	2 bits	2 bits	8 bits	48 bytes

VCI: virtual connection (channel, circuit) identifier

VPI: virtual path identifier

payload: data

ATM cell routing

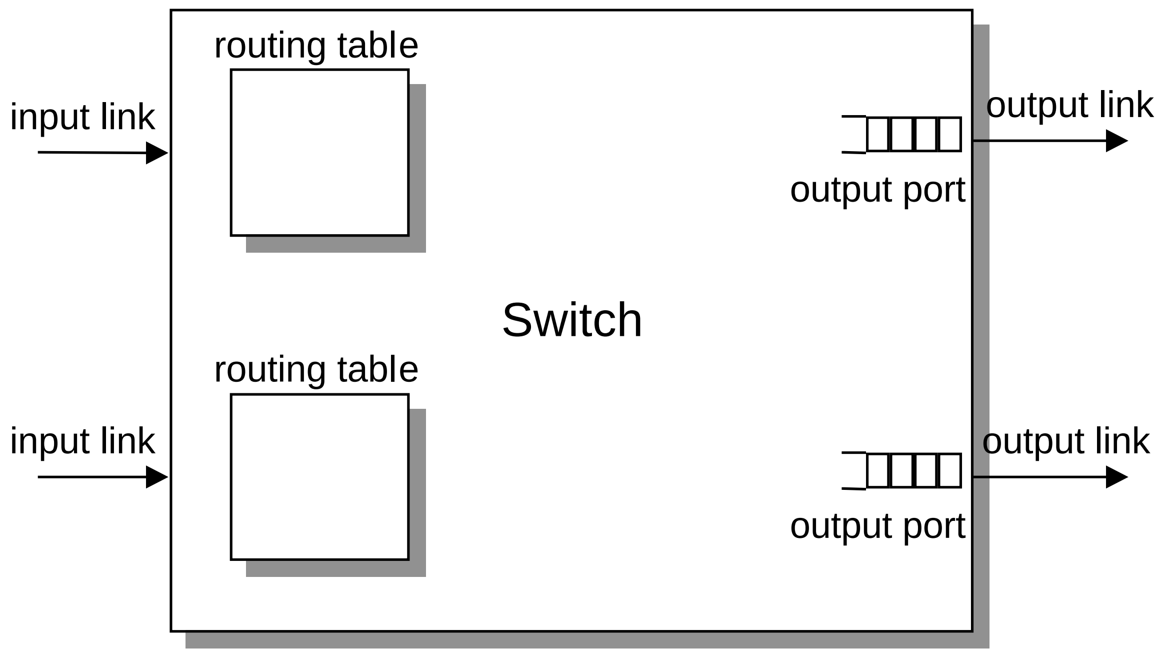

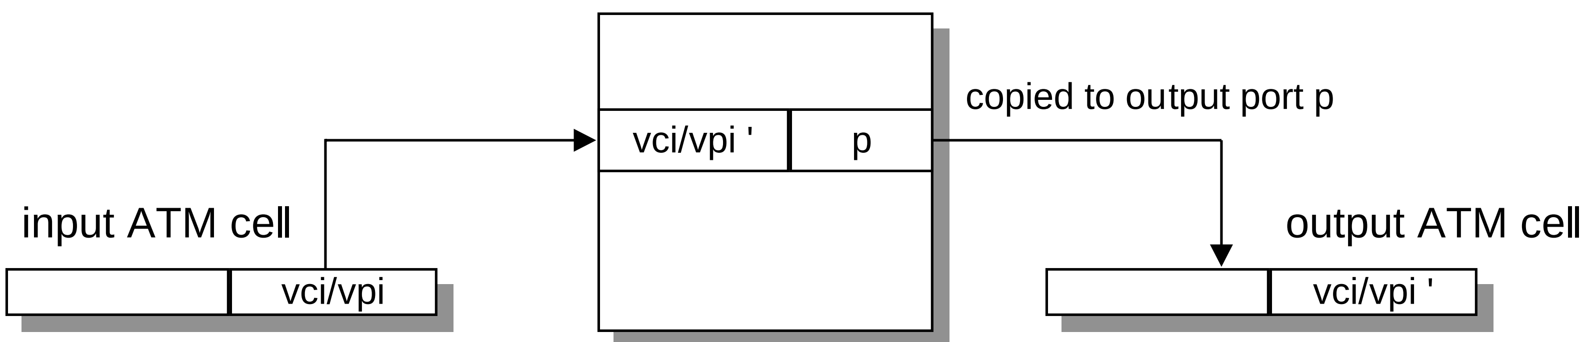

Input port routing table

ATM pros and cons

Pros:

- bandwidth can be reserved (connections)
- scalable aggragate bandwidth (wide range of supported bit rates)
- support for network traffic with different quality of service requirements (small, fixed, easily multiplexed cells)
- potential for high speed switching (small fixed-size cells)

Cons:

- maximum user bandwidth still limited by link bandwidth
- connections make broadcast and multicast more difficult
- quality of service is still a research issue