Network Administration

Building and running the network

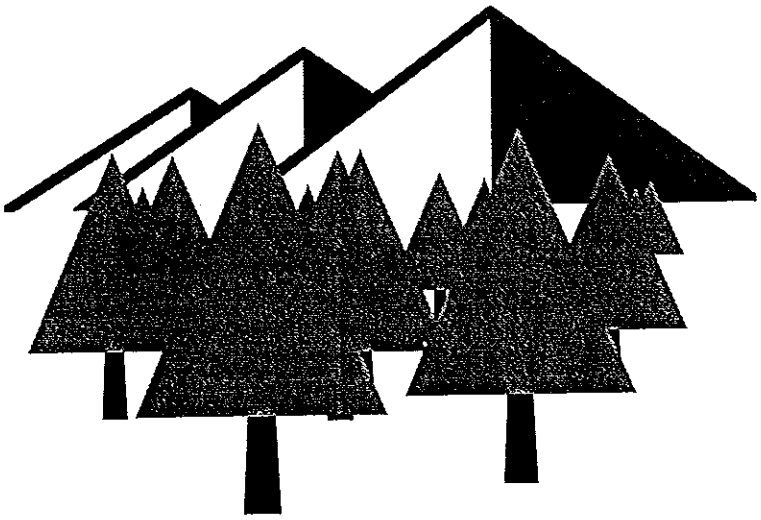

Protocols
Ethernet hardware
IP network administration
Names
Interoperability
Managing connections

WANs *vs.* LANs

Wide Area Networks

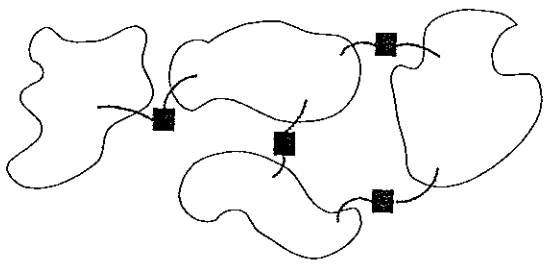

- Connect distant sites
- Require public carriers
- Generally have higher error rates
- Expensive interconnection hardware
- *Not likely to be your concern*

Local Area Networks

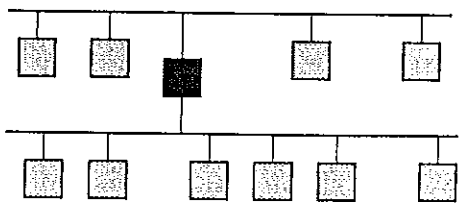

- Restricted to a building or campus
- Inexpensive connection points
- Subject to "environmental" damage
- *Will be your concern*

Protocols

Syntax

Format of transmitted data

Semantics

Meaning of transmitted data

Synchronization

Rules for exchanging data

Examples of protocols

Ethernet
Token Ring
AppleTalk
DECNET
SNA (System Network Architecture)
IP (Internet Protocol)
TCP (Transmission Control Protocol)
Sun RPC (Remote Procedure Call) Protocol

Addresses and Names

Many are needed

Hardware addresses

Specific to a "real" connection point

E. g., Ethernet addresses

Software addresses

Independent of hardware

For example,

IP addresses

`152.18.18.80`

Port numbers

`25`, for mail delivery

Names

For users and applications

E. g., `penrose.cs.unca.edu`

The naming of objects in a network system is *still* an active area of research and development.

ISO OSI Reference Model

ISO = International Standards Organization
OSI = Open Systems Interconnection

Application	Client-to-Server X.400 (mail), X.500 (white pages)
Presentation	Data representation *e. g.*, translating ASCII to EBCDIC
Session	Session creation
Transport	Reliable host-to-host
Network	Connects diverse networks X.25
Link	Reliability in a single net IEEE 802.2 (logical link control) LLC (Logical Link Control) MAC (Medium Access Control)
Physical	Bit encoding, *etc.* *part* of Ethernet, FDDI, ...

Internet protocol

The *software* protocol for Unix networking

Supports two forms of communication

Datagrams

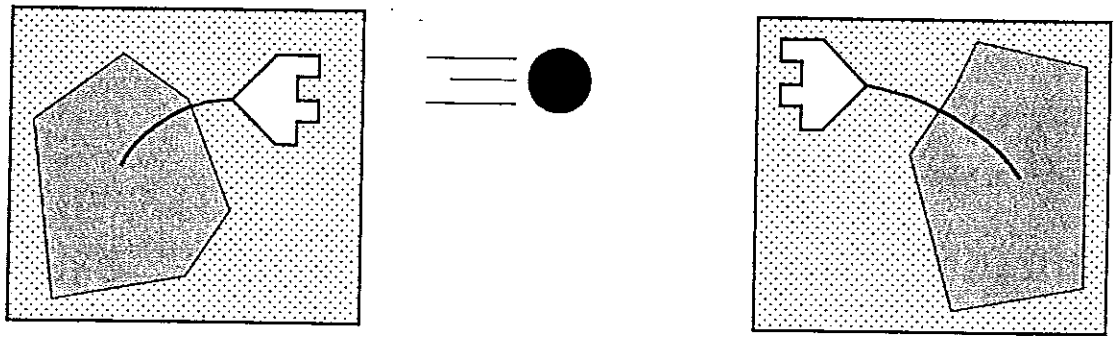

Streams

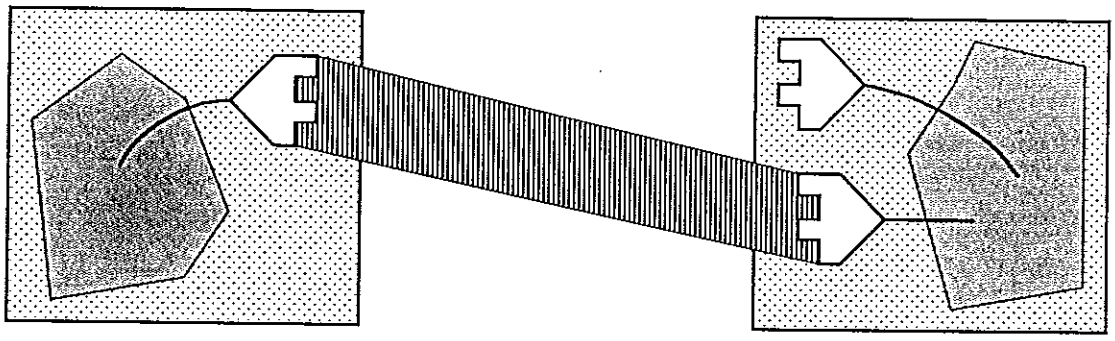

ARPANET Reference Model

Four layer model

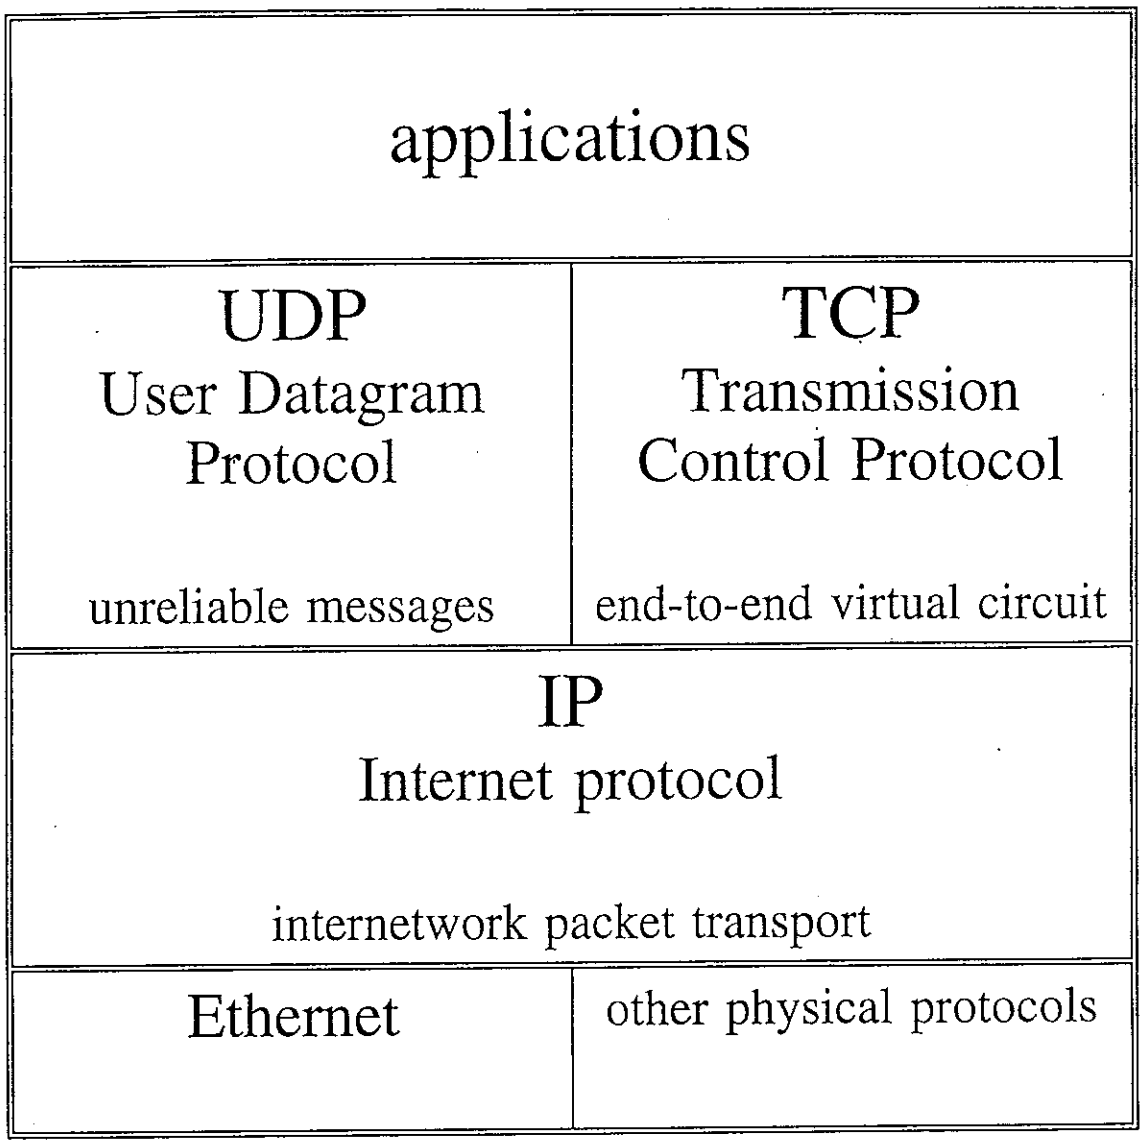

Ethernet
Network segments

Thickwire 10Base5

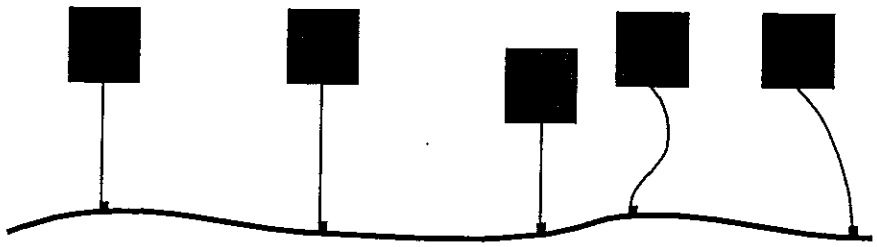

Thinwire 10Base2

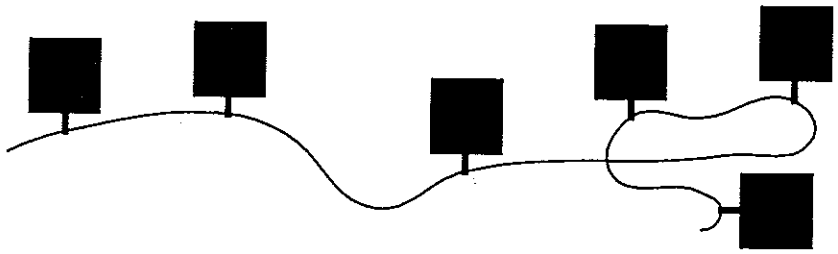

Twisted Pair 10BaseT

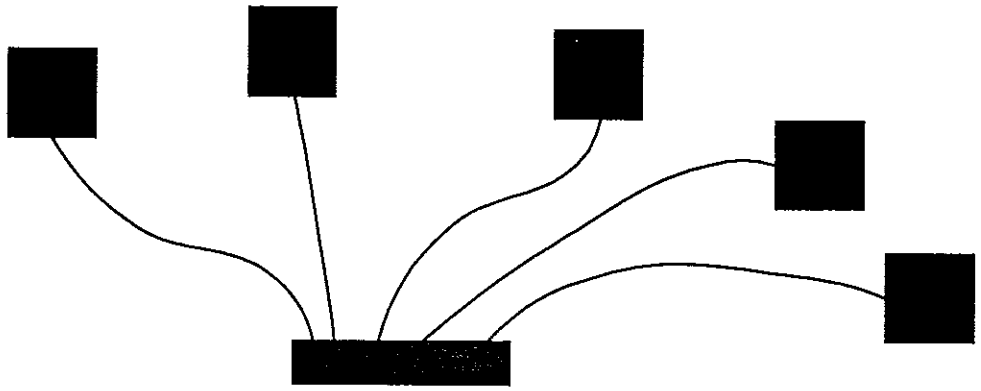

Connectors

Transceivers

- For Thick and Thin
  - "Drives" the daisy chain bus
- For Twisted Pair
  - One end of a wire to the hub
- For Optical fiber
  - Usually ends of point-to-point links
    - May connect to a hub
- Most recent workstation have "built-in" twisted pair transceivers!

Hub

- Joins several twisted pair segments
- May be *cascaded* to form "trees" of hubs

Multi-port transceiver

- "Ethernet in a box"
- *Hub* for Thickwire

Repeater

- Amplifier for the LAN
- Extends LAN beyond length limitations
- Two may operate as a point-to-point link

Addressing

Interfaces have 48 bit addresses
- 24 bits identify manufacturer
  - `08:00:20` Sun
  - `08:00:2b` DEC
- 24 bits assigned by manufacturer

Broadcast address is for all hosts
- `FF:FF:FF:FF:FF:FF`

Packet format

Preamble	10101010....1011	64
Destination		48
Source		48
Length		16
Data	packet can be from 64 to 1514 bytes in length	
Cyclic Redundancy Check		32

CSMA/CD

Carrier Sense CS

Before transmitting,
host verifies the net is free.

Multiple Access MA

Any host may transmit, if net is free.

Collision Detection CD

When two hosts transmits simultaneously
a *collision* occurs.
Both hosts see the collision and *jam*.
Retransmission occurs after a random wait.

Exponential Binary Backoff

Maximum wait doubles after each collision
up to 10

Expected maximum throughput is $10/e$ Mbps.

Restrictions

Mostly a result of the CSMA/CD *delay budget*

Distance limitations

- 500 m for 10Base5
- 185 m for 10Base2
- ~100 m for 10BaseT runs
- 1000 m for repeater links
- 2500 m between most distant stations

Limitations on number of hosts

- ≤25 hosts on Thinwire segment

Limitations on number of repeaters

- No more than four repeaters between any two hosts

Play it safe!

Network problems are difficult to debug

Bridges

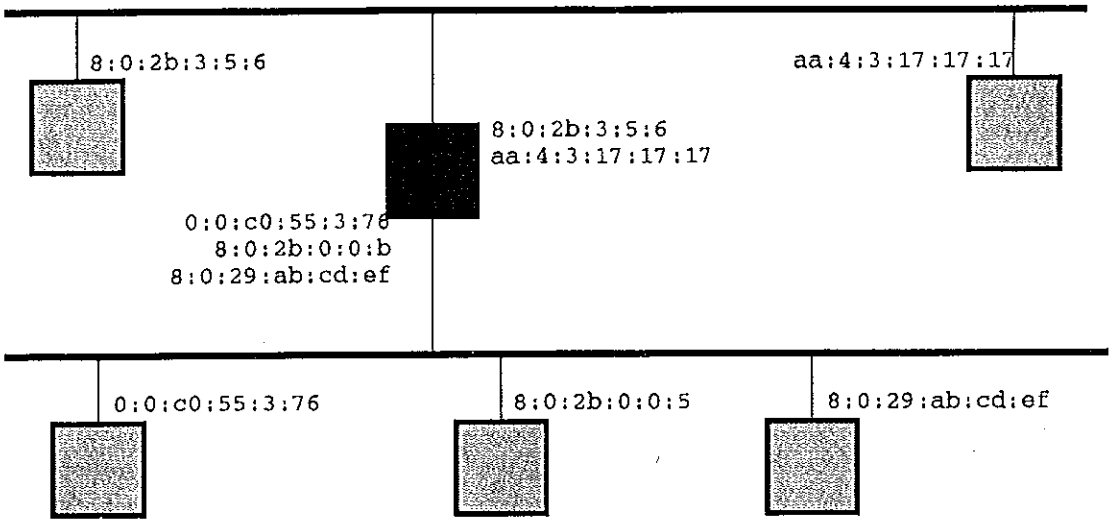

Bridges

- Isolate local network traffic
- "Learn" the hosts located at each port
  - A host is located once it transmits

Spanning tree bridges

- May be connected in cycles
- Form a *tree* within a *mesh*

Bridges are easy to install.
Bridges solve and avoid many problems.

LAN growth

First CSCI LAN

About one-half of campus ''backbone''

Second CSCI LAN

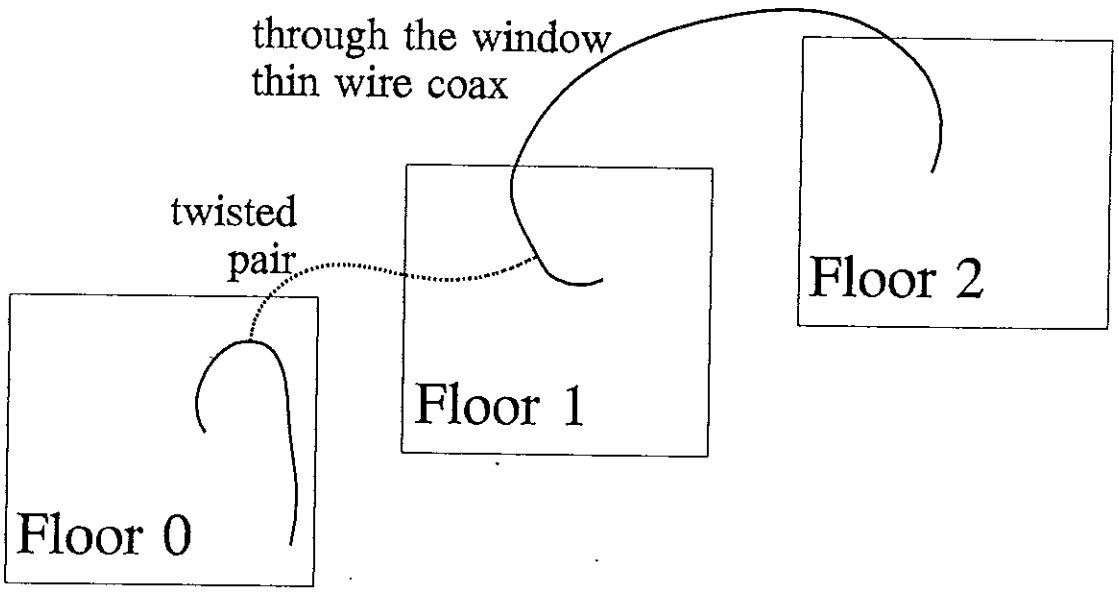

Third CSCI LAN

One thinwire segment connected to campus backbone with a bridge

Fall 1994 CSCI LAN

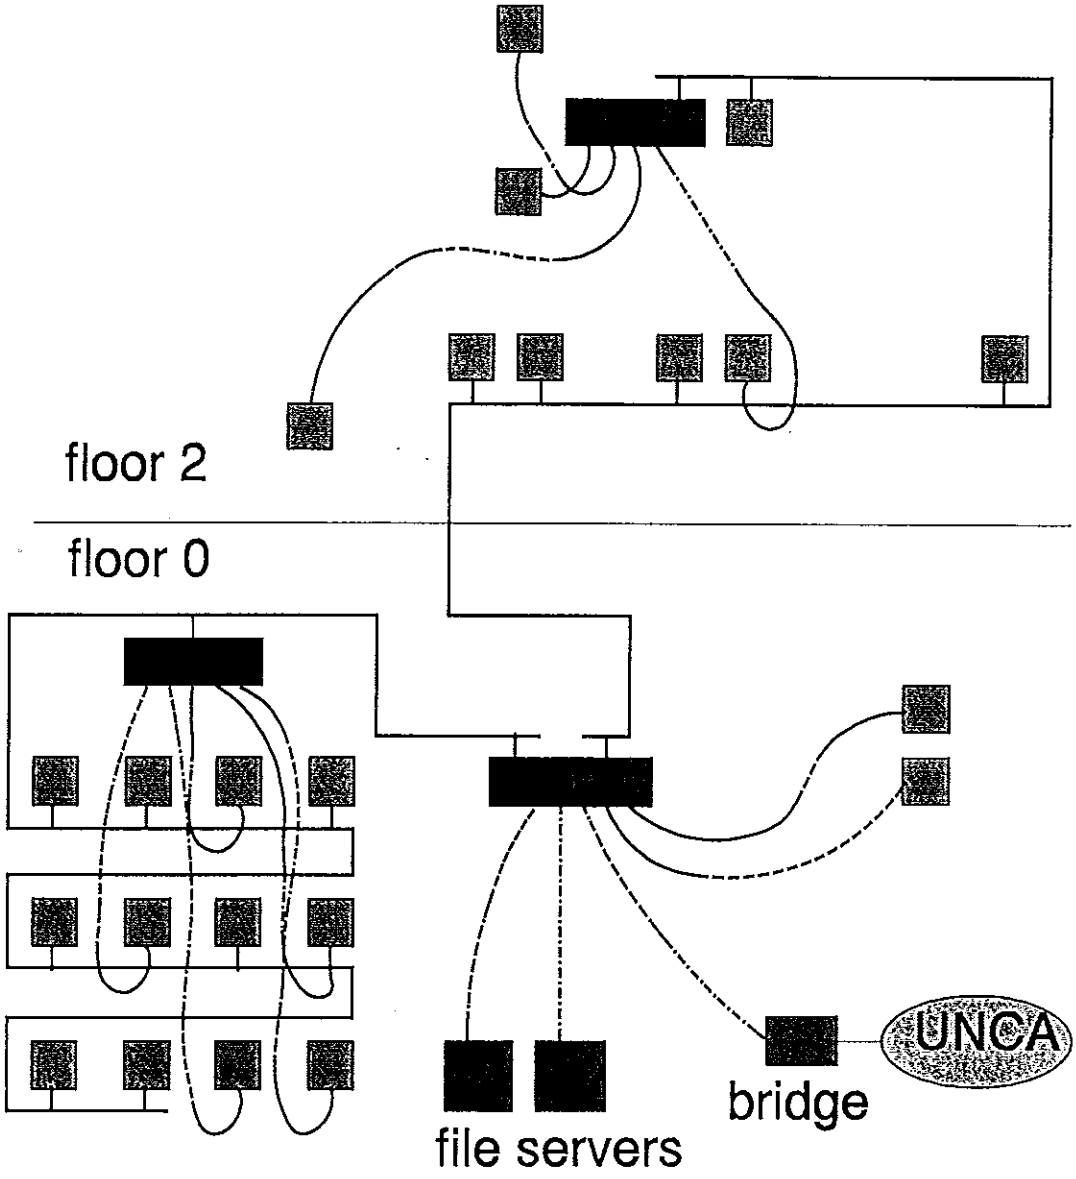

UNCA LAN

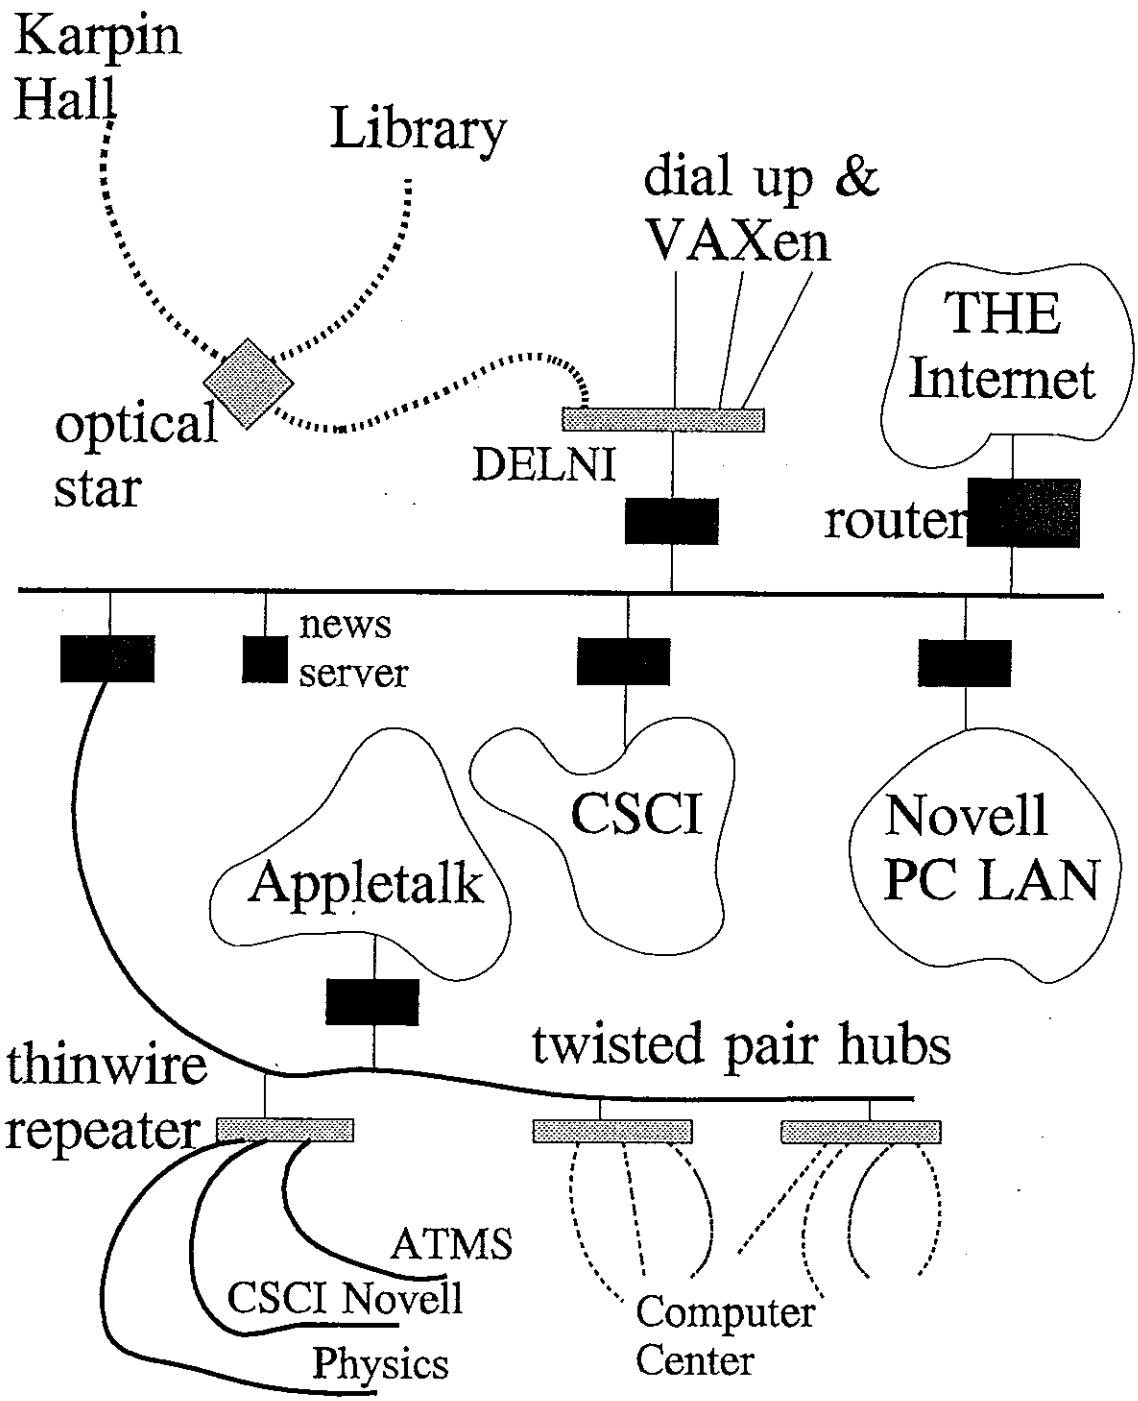
Karpin
Hall
Library
dial up &
VAXen
THE
Internet
optical
star
DELNI
router
news
server
Appletalk
CSCI
Novell
PC LAN
thinwire
repeater
twisted pair hubs
ATMS
CSCI Novell
Physics
Computer
Center

After Ethernet

Limitations of Ethernet

10 Mbps is too slow

Performance decreases as load increases

10/*e* Mbps is even slower

No guarantees of bandwidth

Not suitable for multimedia applications

The next big LAN?

Should run on *category 5* twisted pair

If the connector problem is solved

Should "win" in two years

Or maybe three or five or seven or ...

FDDI

Fiber Distributed Data Interface

Most developed *next* network interface

- ANSI X3T9.5 standard
- 100 Mbps
- Up to 2 km between stations
- Token-based allocation
- Will run over category 5 twisted pair
  - TP-PMD -- Twisted Pair Physical Layer Medium
- FDDI Hubs
  - *E. g.*, Digital's Gigaswitch

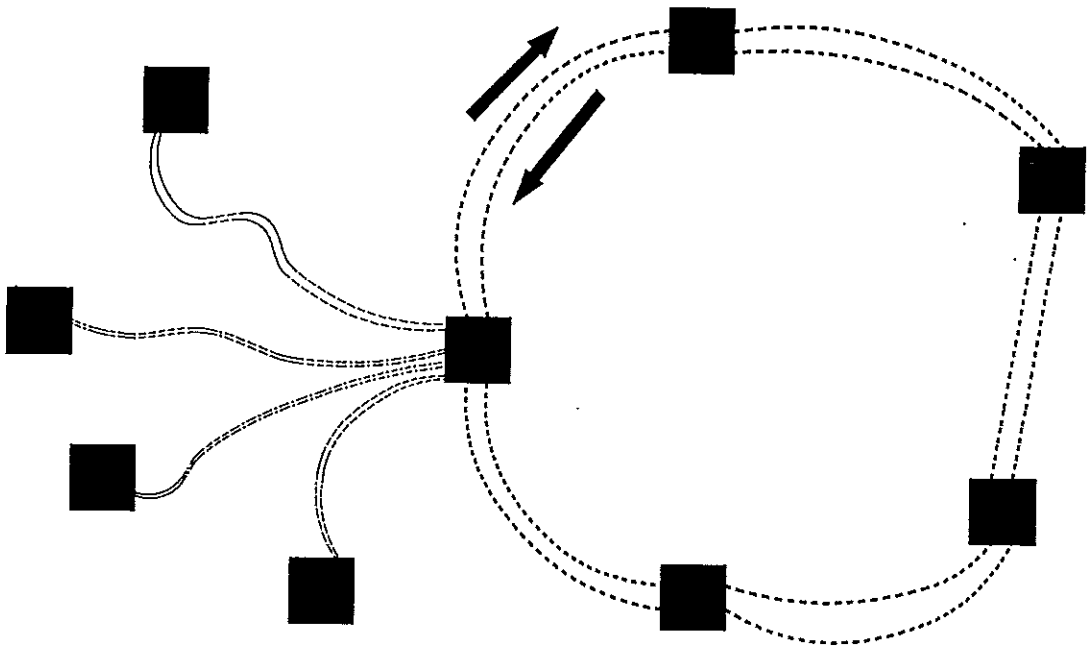

100 VG-AnyLAN

Standardization

Developed by Hewlett-Packard and AT&T

Will be IEEE 802.12

Draft 7 in April 1995

Guided by 100VG-AnyLAN Forum

Packet Format

Uses standard Ethernet packet *and* may also use token ring packet

Wires and medium access

Hub based topology

Quartet signalling

Uses all four twisted pairs

Demand priority transmission with round-robin allocation

Fast Ethernet

Standardization

Developed by Grand Junction

Will be IEEE 802.3u -- 100BaseT

Guided by Fast Ethernet Alliance
3COM, Cabletron, SMC, Digital, ...

Packet Format

Standard Ethernet packet

Wires and medium access

Hub based

Auto-Negotiation switches
Support both 100 Base T and 10 Base T
along with ''Full Duplex mode''

T Twisted pair

100 Base TX for two-pair Category 5
100 Base T4 for four-pair Category 3

revised CSMA/CD

ATM
Asynchronous Transfer Mode

Standardization

Embraced by **The Phone Company**

Guided by ATM Forum

Packet Format

Uses *frame relay*

53 byte cells containing 48 bytes of data

Wires

Both WAN and LAN

Presently runs on optical fiber

Workstation interfaces are available

Should run on category 5 twisted pair

at 155 Mbps

IP networks and numbers

32 bit numbering scheme

"dotted decimal" notation

`152.18.52.5` ⇒ `0x98123405`

IP address = network + local

Class A nets 0.0.0.0 to 127.255.0.0

net	local
0.....	24 bits

Class B nets 128.0.0.0 to 191.255.0.0

net	local
10....	16 bits

Class C nets 192.0.0.0 to 223.255.255.0

net	local
110...	8 bits

IP subnets

At many universities,
a Class B network is subnetted

Semi-hypothetical subnet

IP addresses 128.52.0.0 to 128.52.31.255

Third *octet* in binary

`00000000` → `00011111`

To select out network number use *mask*

`11111111111111111110000000000000`, or

`255.255.224.0`

Do not try this at home

You **must** ask your network manager

IP has been around a while and sometimes shows its age

ARP
Address Resolution Protocol

Each host maintains an ARP cache

152.18.18.4	08:00:2b:0d:c2:59
152.18.1.1	aa:00:04:00:01:c0
152.18.18.50	08:00:2b:29:00:00
152.18.18.1	08:00:20:0a:c0:3f

If needed IP address not in ARP cache,
host *broadcasts* an ARP request.

> IP # 152.18.18.52! Where are you?

Target host should respond.

> IP # 152.18.18.52 has
> Ethernet # 08:00:2b:28:ff:5c

New mapping will be added to ARP cache.

IP routing

Route based on IP (or similar) addresses

IP fans often call these "gateways."

IP *style* routing tables

for the middle host

net	152.18	ln0
net	199.17.83	200.18.5.3
default		152.18.254.254
net	200.18.5	ln1

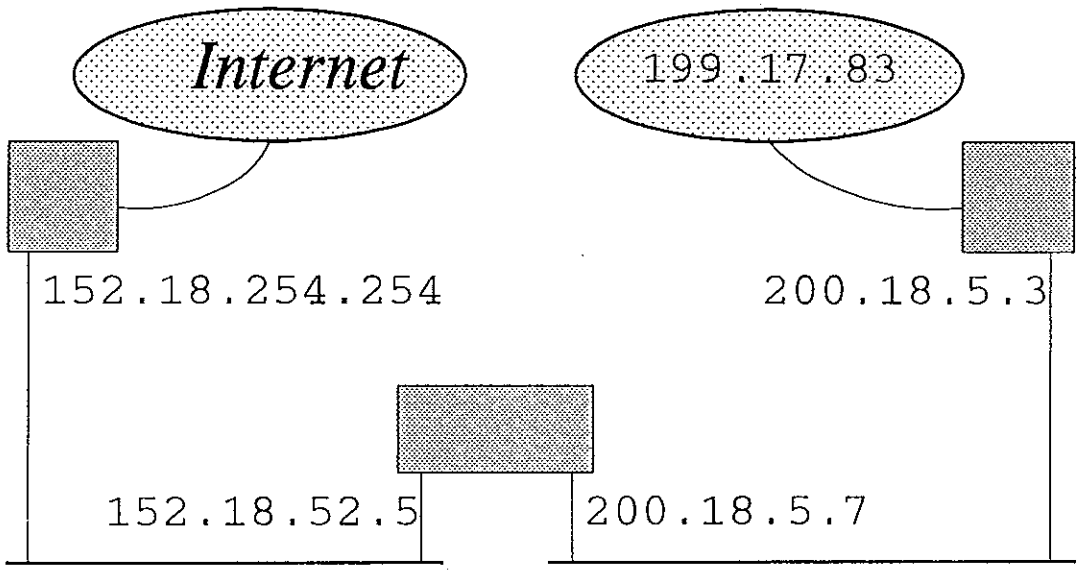

Most campuses have a default router

Ports

Applications identified by ports

16-bit integers

Unix systems "reserve" 0..1023

Some *well-known* ports

from `/etc/services`

```
ftp       21/tcp
telnet    23/tcp
smtp      25/tcp      mail
domain    53/tcp      nameserver
domain    53/udp      nameserver
finger    79/tcp
pop       109/tcp     postoffice
sunrpc    111/tcp
sunrpc    111/udp
login     513/tcp
who       513/udp     whod
shell     514/tcp     cmd
```

Snapshot of an IP packet

A packet for a `talk` connection from `inria.inria.fr` to `penrose.cs.unca.edu` going through the UNCA Ether

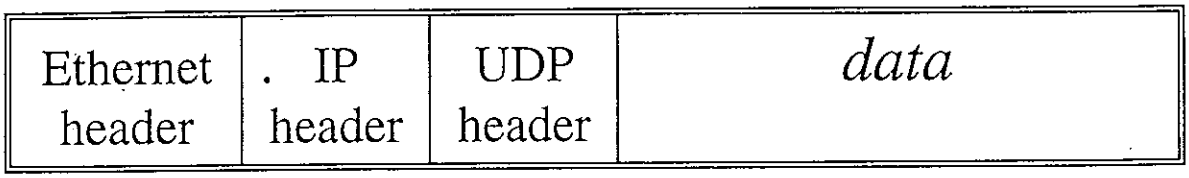

Ethernet header

destination	`aa:00:04:00:06:ac`, UNCA router
source	`08:00:2b:39:f4:58`, penrose
type	`2048`, IP protocol

IP header

destination	`192.93.2.1`, inria
source	`152.18.18.83`, penrose
protocol	`17`, UDP
fragmentation information	

UDP header

port	unpredictable, but identifies talk *peer*

Domain names ⇒ IP numbers

`/etc/hosts`

```
127.0.0.1 localhost
152.18.18.5 ivy.cs.unca.edu ivy
152.18.18.4 enka.cs.unca.edu enka
# Decstation clients
152.18.18.50 celo.cs.unca.edu   celo
152.18.18.51 etowah.cs.unca.edu etowah
152.18.18.52 canton.cs.unca.edu canton
152.18.18.53 arden.cs.unca.edu  arden
```

Non-local lookups use *name server*

- Distributed database
  - Contains "resource records"
  - Organized hierarchically
    - from ~10 root servers
  - Redundancy for reliability
  - Caching for performance

`tryon.cs.unca.edu`

primary server for `cs.unca.edu`

secondary server for `unca.edu`

Distributed naming

Networked workstations require distributed password schemes.

Network Information Services (NIS)

Yellow pages (yp)

Master server

Contains the *one* permanent copy

Slave servers

Master *pushes* copies of the database to slave servers

Clients

Request *maps* and *matches* from servers

BIND/Hesiod

Similar scheme from Project Athena

Master → Primary

Slave → Secondary

Maintaining ~~YP~~ NIS

Hosts are organized into domains
- domains ≈ administrative unit
- Example domains
  - `uncacsci` or `uncaatms`
- This is *not* the Internet domain name
  - At least, it shouldn't be.

Generally used to share passwords

Install with setup scripts
- `ypinit` or `ypsetup`

Database stored on primary server
- generally `/var/yp/`*`domainname`*

Binary database is `maked` from ASCII files
- perhaps stored in `/etc`

Maintaining DNS

Domain Name Service

This *should* be maintained by your university!

Domain *resource records*

Administrative

```
cs.unca.edu.IN   SOA tryon.cs.unca.edu. (
                 postmaster.cs.unca.edu.
                 63   ; serial
                 14400    ; refresh
                 3600     ; retry
                 1209600  ; expire
                 86400 )  ; minimum
             IN  NS       tryon.cs.unca.edu.
             IN  NS       ncnoc.concert.net.
             IN  NS       balsam.unca.edu.
             IN  MX       10 tryon
```

name →number

```
ivy      IN  A       152.18.18.5
ftp      IN  CNAME   ivy
tryon    IN  A       152.18.18.6
www      IN  CNAME   tryon
```

number → name

```
32       IN  PTR alexander.cs.unca.edu.
33       IN  PTR pensacola.cs.unca.edu.
```

If you must see more,

try `/etc/namedb`

on `tryon.cs.unca.edu`

Using BIND

/etc/resolv.conf

Used to find the local servers

Example contents

```
domain          cs.unca.edu
nameserver      152.18.52.6
nameserver      152.18.52.5
```

Test with

`nslookup`

Important tips

Use more than one server

Try to locate one on your own Ethernet

If performance is a problem, try using a local *caching only* server

Choosing the method

Three sources of names

local files, *e. g.*, `/etc/hosts`

~~YP~~ NIS

Domain name servers, BIND

Most vendors do

(1) local files

(2) YP or DNS

Some vendors allow you to choose

from `/etc/svc.conf` under DEC OSF/1

```
hosts=local,bind
passwd=local,yp
```

Some vendors put YP in charge

from `/var/yp/Makefile` under SunOS

```
# Set the following variable to
# "-b" to have NIS servers use the
# domain name resolver for hosts
# not in the current domain.
B=-b
```

Security

/etc/hosts.equiv

List of *trusted* hosts

Advantages

- User convenience
- Reduces net-snoopable passwords

Disadvantages

- Some reduction of security
  - Particularly for unattended terminals

Watch out for "+ +" on Suns

/.rhosts

List of *really* trusted hosts

May prove useful *but*

- use with extreme care!

Network File System

Client machines mount remote file systems

Runs on top of Sun RPC protocol

- Generally, with UDP as the transport protocol

Servers are *stateless*

- Servers retain *no* information for clients
- If case of server crash, client can just wait

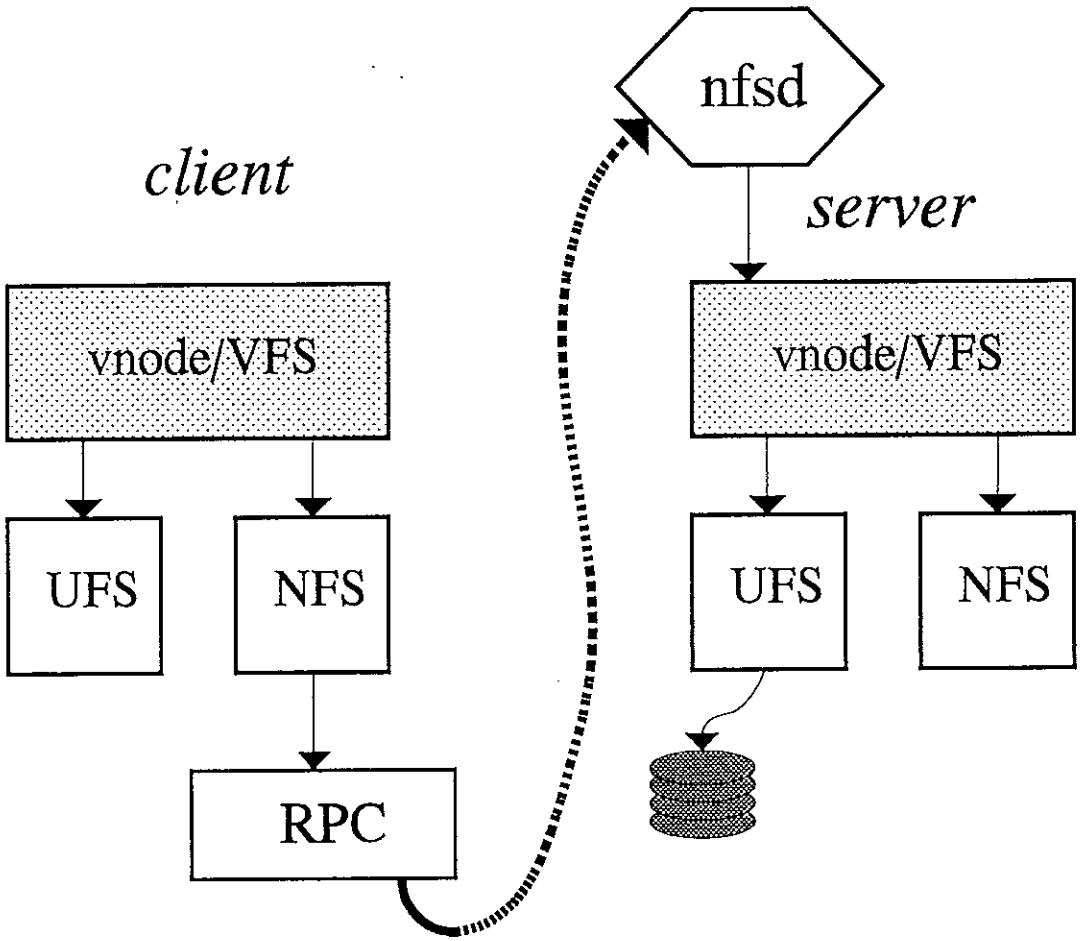

Managing NFS

Exporting file systems

From `/etc/exports`

```
/users enka celo etowah
```

Export file system to three machines

```
/cdrom -o
```

Exports file system read-only anywhere

```
/melville -r=1819 ivy
```

If I/O request is from superuser, pretend request from user 1819.

Importing file systems

From `/etc/fstab` on Ultrix

```
/src@tryon:/tryon/src:rw:....
```

or from `/etc/vfstab` on Solaris 2

```
tryon:/src - /tryon/src nfs ....
```

Directory `/tryon/src` is really directory `/src` on `tryon`.

Hints for maintaining NFS

Make mounts identify remote machine

I. e., mount on */host/filesys*

`/mipsloc` on `tryon`
mounted as `/tryon/mipsloc`

Apply symbolic links liberally

On MIPS machines
`/usr/local` → `/tryon/mipsloc`

On Alpha AXP machines
`/usr/local` → `/tryon/alphaloc`

Use NFS options wisely

For exports

`ro`	read-only

For imports

`nosuid`	no set-uid
`nodev`	no devices
`bg`	background
`intr`	interruptable

Printing

Use the setup script!

BSD style remote printing

from `/etc/printcap`

```
rbh006|lp0:\
    :lp=:\
    :rm=uncavx.unca.edu:\
    :rp=sys$print:\
    :sd=/usr/spool/lpd-rbh006:
rbh005|lp3:\
    :lp=:\
    :rm=marion:\
    :rp=rbh005:\
    :sd=/usr/spool/lpd-rbh005:
```

rm=*hostname*

Gives the remote host

rp=*printer*

Gives name of printer on remote host

Many popular MS/DOS networking packages support BSD style printing.

Useful tools

included on almost every system

`ping inria.inria.fr`
Sends an echo request to remote host

`arp -a`
Prints contents of the ARP cache

`ifconfig ln0`
Prints configuration of interface `ln0`

`netstat -i`
Prints statistics for interfaces

`netstat -a`
Prints open process (socket) connections

`netstat -r`
Prints routing table

`nslookup hatteras.cs.unc.edu`
Queries the name server

`rpcinfo -p ivy`
Gives information about Sun RPC servers

`showmount`
Gives information about exported and mounted files. Options vary. Read `man`.

`nfsstat`
Gives statistics about NFS clients and servers. For advanced users.

Useful tools

not so standard

`traceroute ftp.uu.net`

Follows route to a remote site

`etherman`

Graphical display of ethernet traffic

"*sniffer*"

Software to capture and display packets

Sniffer is a trademark of Network General

network management tools

Programs that use SNMP (Simple Network Management Protocol) to obtain statistics from network devices

`esum`

Summarizes ethernet traffic over extended periods

Interoperability

NCSA telnet

Easily FTPable public domain TCP/IP networking for PC's and Macs

Fairly easy to install

`telnet`
`ftp`
`lpr` utilities
r-utilities

News readers

Mail programs

X-terminals

MS/DOS networking

Goals

- Redirection for files
  - `M:\` → my Unix home directory
- Redirection for printing
  - `LPT2` → `rbh221` print queue

Perils

- Many different interface cards
- Different "styles" of protocol layering
  - packet drivers
    - Specification of FTP Software
  - NDIS
    - Network Driver Interface Specification
    - Microsoft and 3Com
  - ODI
    - Open Data-link Interface
    - Novell
- Different transports
  - Digital's Pathworks
  - Sun's PC-NFS
  - *and so on and so on*

AppleTalk

Well-defined and universally accepted

if universe = Macintosh

cf MS/DOS networking

Most Unix vendors will sell you AppleTalk

CAP -- Columbia AppleTalk Package

almost Public domain AppleTalk for Unix

Supports

EtherTalk, phase 1 and phase 2

IPTalk

with appropriate local gateways

Application level protocols

AppleShare 2.0

File service

Printer Access Protocol

Printing to Apple printers

World Wide Web

Client programs

Anyone can install

`netscape`, `Mosaic`, `lynx`, `chimera`

Server programs

Usually installed by the super-user

But may be installed by any user at non-standard port

`httpd`

FTP from the usual sites

Read the directions

Set up your server's home page

Probably `/usr/local/etc/http/htdocs/index.html`

Modify system startup to start `httpd`

Gopher

Similar to World Wide Web installation

Client programs

`gopher`, `xgopher`

Server program

`gopherd`

Setting up the gopher database

Information is stored in Unix files, *but*

- `.Links` can point to other services

```
Type=8
Name=D. H. Ramsey Library
Path=dhramsey
Host=uncavx.unca.edu
Port=0
```

- `.names` can give better descriptions

```
Path=./ComputerInfo
Name=Computer and Network Information
```

FTP

To set up anonymous FTP

Read `man ftpd` carefully

Create the FTP directory

`/usr/local/ftp`

Create FTP subdirectories

`~ftp/etc`
with *false* password and group files

`~ftp/bin`
with `ls` program

`~ftp/pub`
with careful attention to protection

Create FTP user

with `NoLogin` stored in password field

Mail

/*checkyourdir*/sendmail.cf

The *rules* for mail

```
R$*<@$*.bitnet> $#smtp $@uncavx.unca.edu $:$1<@$2.bitnet>
```

You don't want to touch this

Find someone in your situation

The Book -- all 792 pages

Bryan Costales with Eric Allman & Neil Ricket, *sendmail*, O'Reilly & Associates, 1993, ISBN: 1-56592-056-2.

/*checkyourdir*/aliases

For local mail aliases

```
uncaww:brock, stigle
stigleman:stigle
```

News

News requires Gigabytes of disk

Find the on-campus news server
and `setenv NNTPSERVER`

Otherwise,

Find student news junkie
Point student toward Internet
`INND` is a popular server
Send in PO for new disks
Figure out a policy for new groups
For when the local newspaper calls you

News reader clients

All may be obtained via FTP

NNTP clients require `nntp` package
for `inews` program