# TCP/IP Illustrated

One chapter per lecture

L. Aaron Kaplan <aaron@lo-res.org>

# Course layout

- regular meeting of geeks who want to know how this really works
- goal: end at implementation of TCP/IP
- at least one chapter per lecture

# Introduction

# Protocol layers

not 7 but 4

| | |
|---|---|
| **Application** | **Telnet, FTP, e-mail, etc.** |
| **Transport** | **TCP, UDP** |
| **Network** | **IP, ICMP, IGMP** |
| **Link** | **device driver and interface card** |

**Figure 1.1 The four layers of the TCP/IP protocol suite.**

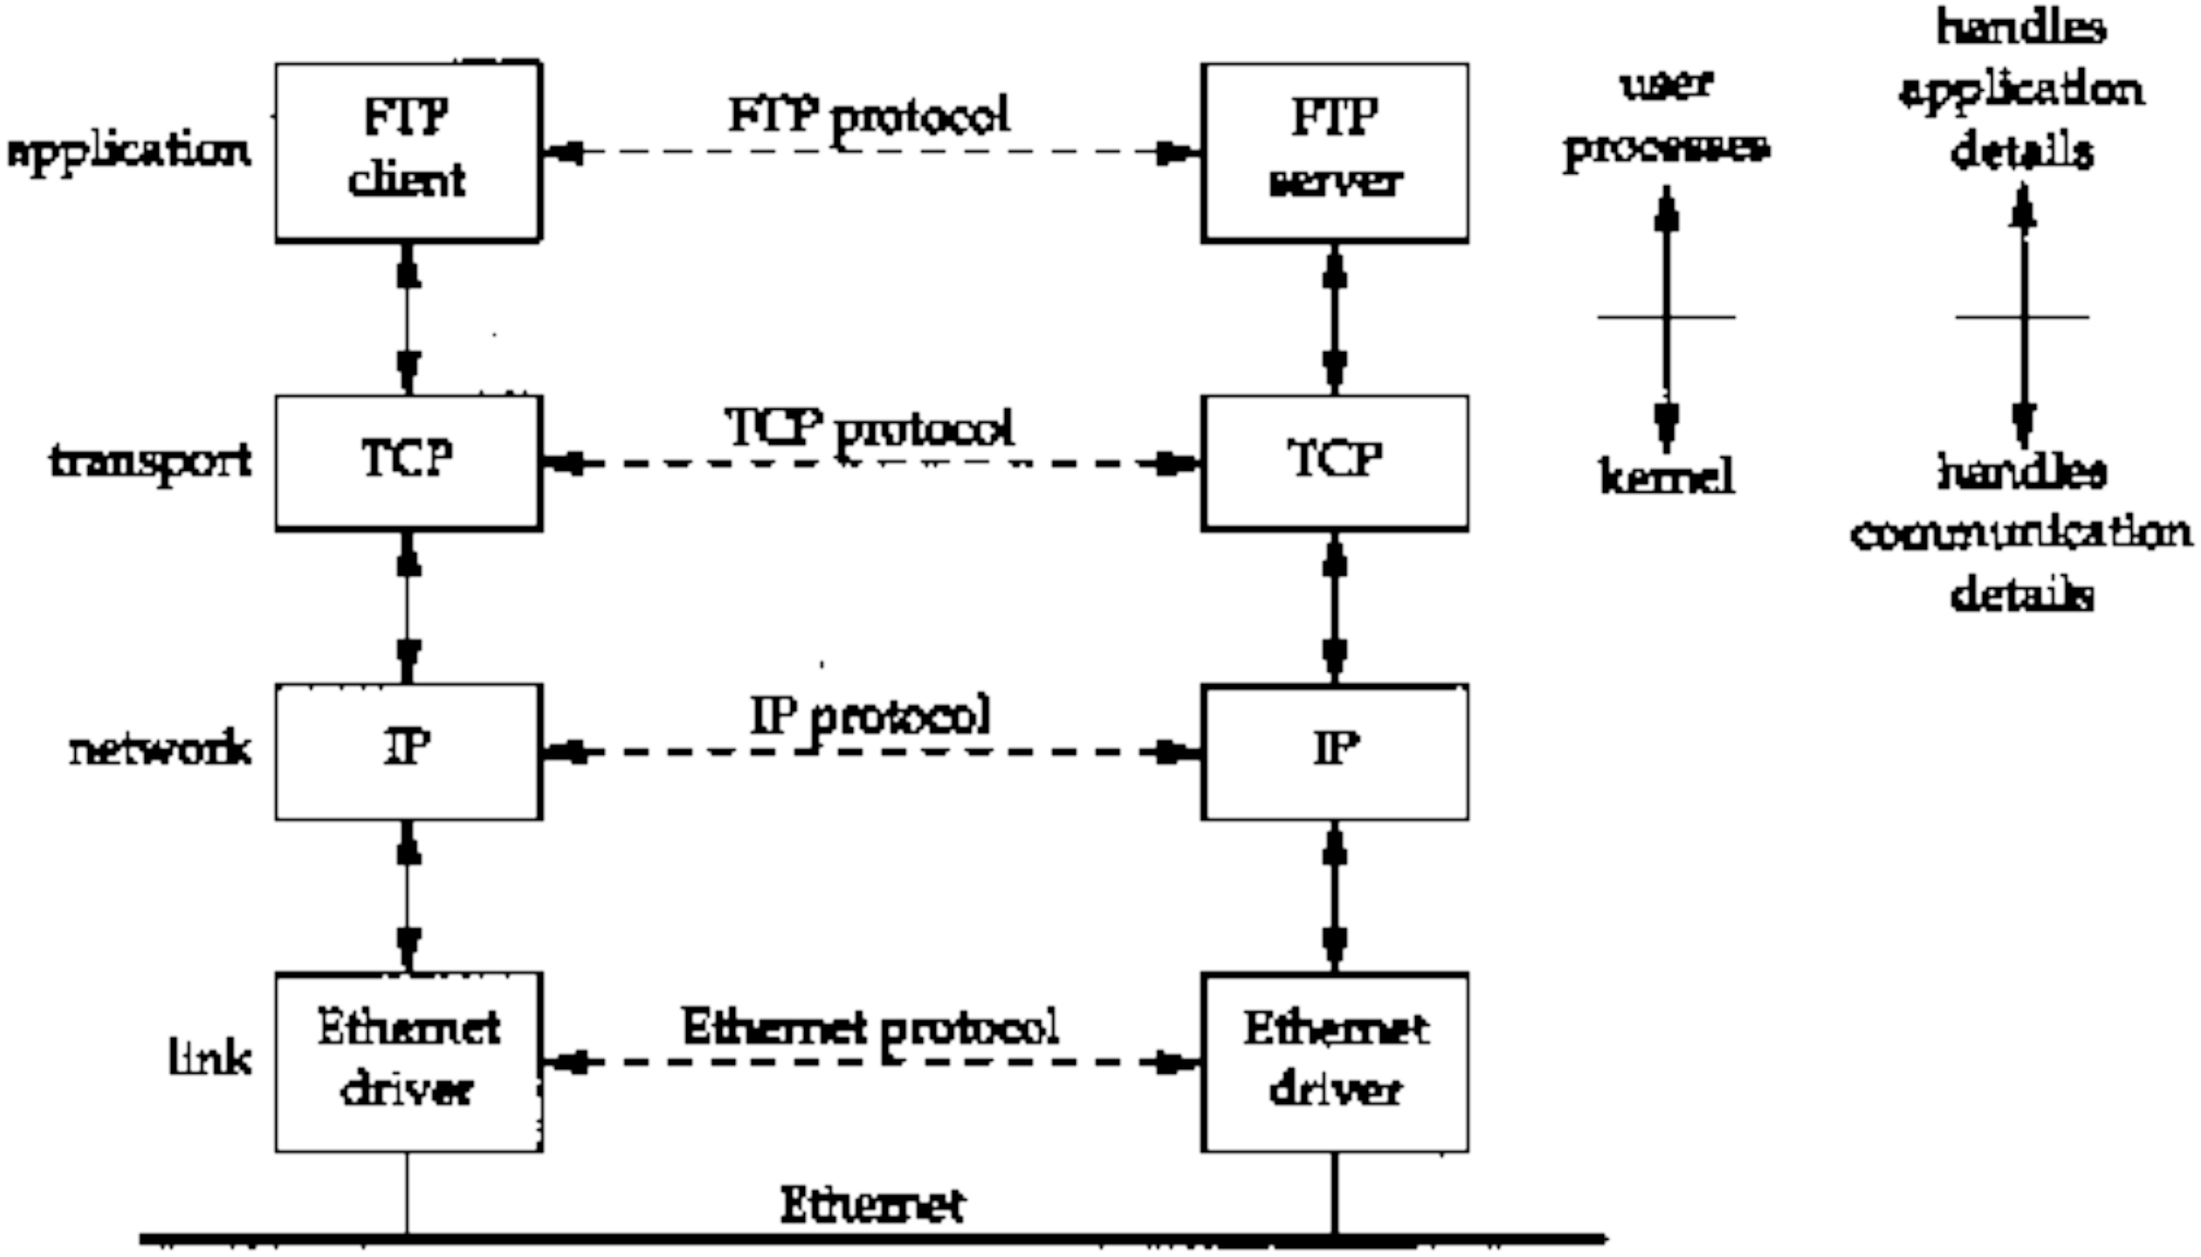
application
FTP client
FTP protocol
FTP server
transport
TCP
TCP protocol
TCP
network
IP
IP protocol
IP
link
Ethernet driver
Ethernet protocol
Ethernet driver
Ethernet
user processes
kernel
handles application details
handles communication details

# Protocol layers
# Layer 1

- PHYSICAL
- Ethernet cable
- WLAN radio
- optical cable
- pigeons ;-)

# Protocol layers
# Layer 2

- data link layer
- device drivers in the OS: ethernet driver, wlan driver
- 802.11a/b/g frames
- Ethernet frames
- Layer 2 1/2: ARP, RARP

# Protocol layers
# Layer 3

- IP = Internet Protocol
  Routing, basic packets, unreliable
- ICMP = Internet Control Protocol
  errors (ICMP unreachable, “ping”,
  “traceroute” (ICMP time exceeded (type 11)),...)
- IGMP = Internet Group Management P.
  membership in multicast groups

# Protocol layers, Layer 4

- TCP = Transmission Control Protocol
  TCP provides a RELIABLE flow of data between two hosts
- UDP = User Datagram Protocol
  UDP, on the other hand, provides a much simpler service to the application layer. It just sends packets of data called datagrams from one host to the other, but there is NO GUARANTEE that the datagrams reach the other end. Any desired reliability must be added by the application layer.

So what is the difference between layer 3 and layer 4?

# Layer 4

- Reliability (TCP)
- bigger fast packets (UDP)

# Layer 3: Routing

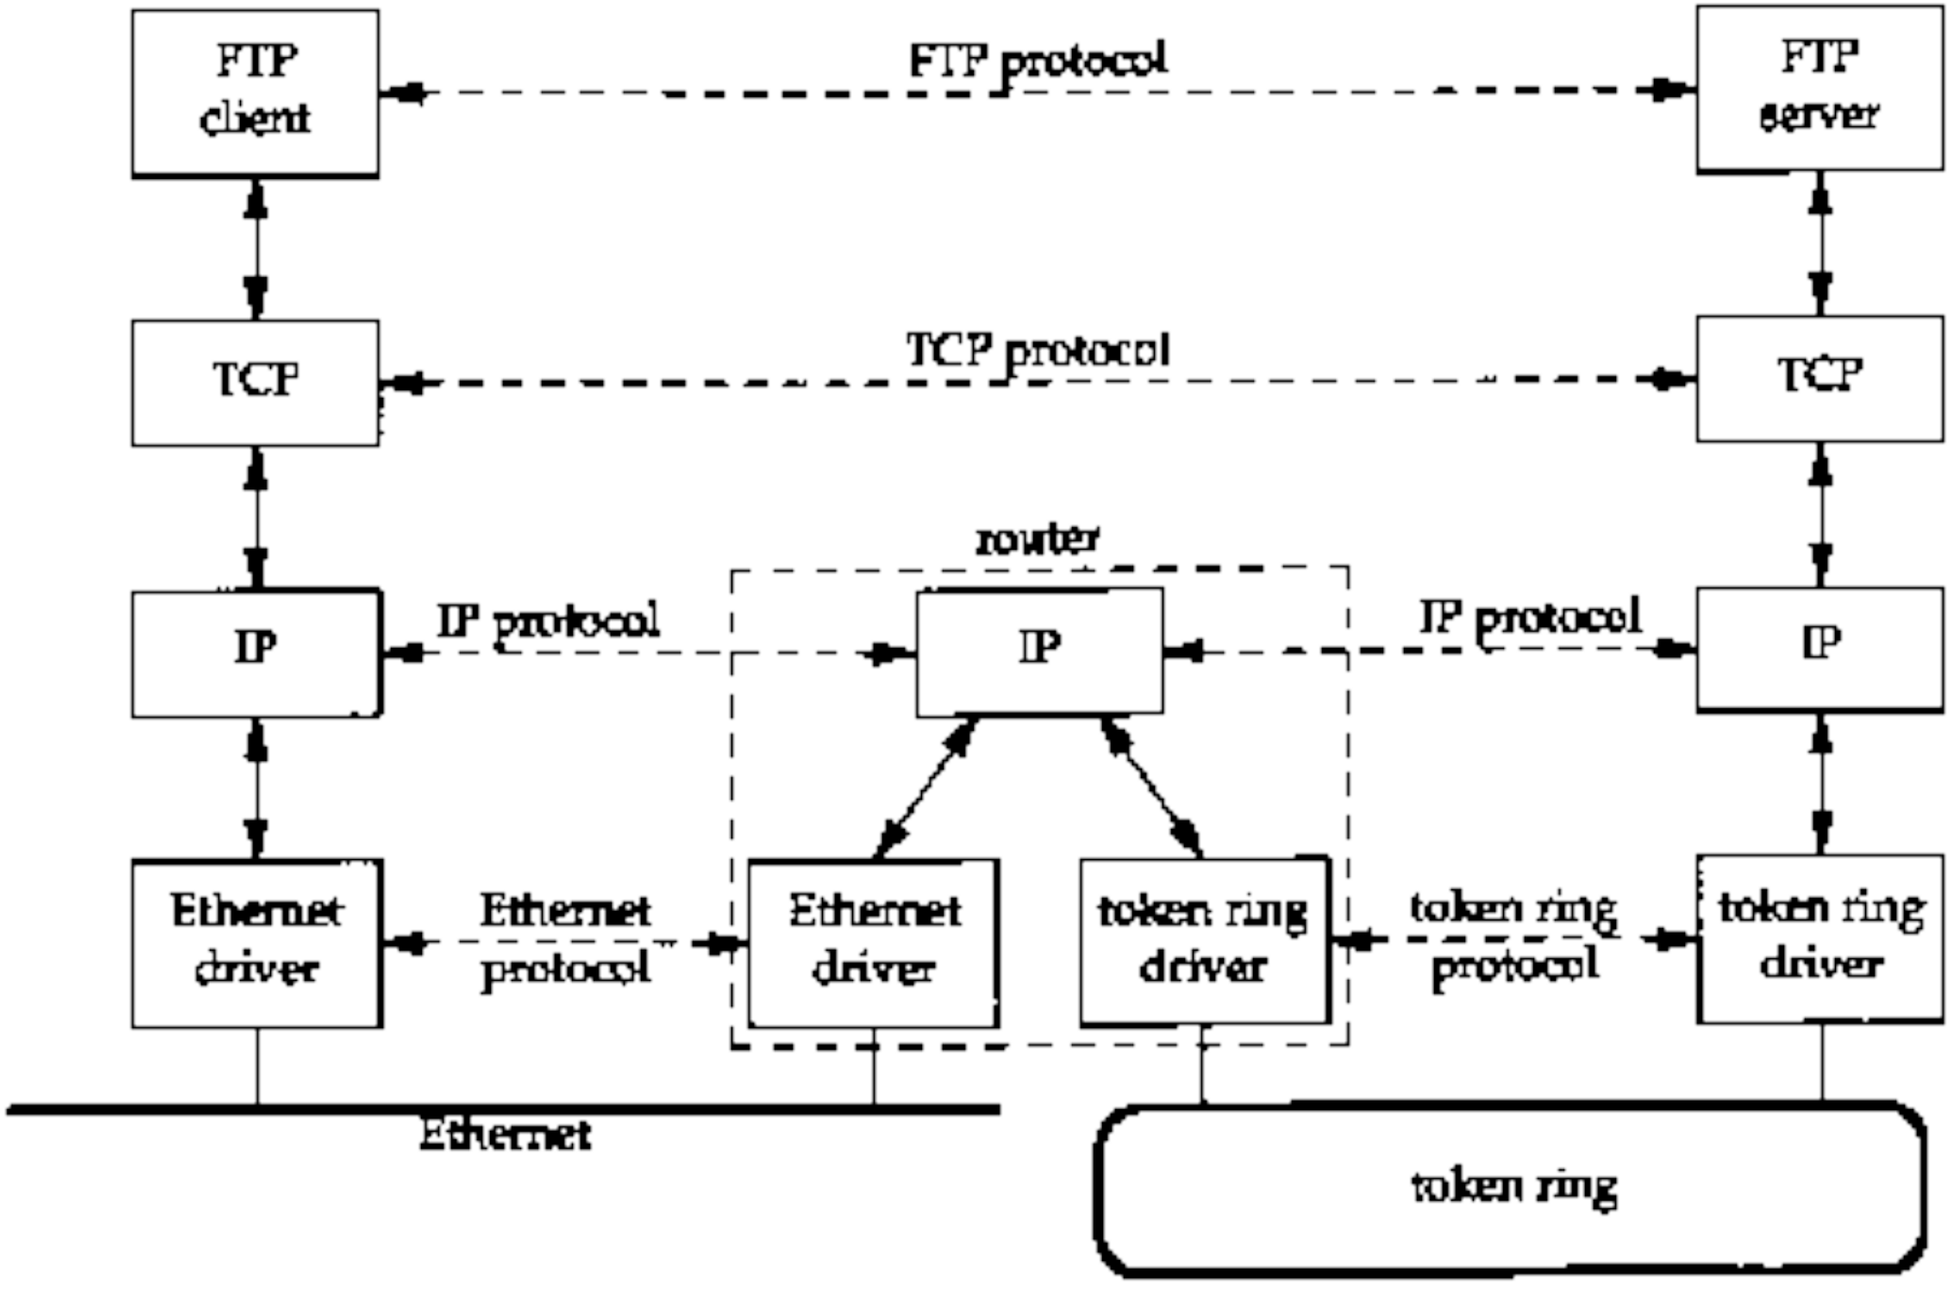

# Layer 3: Routing

multi-homed

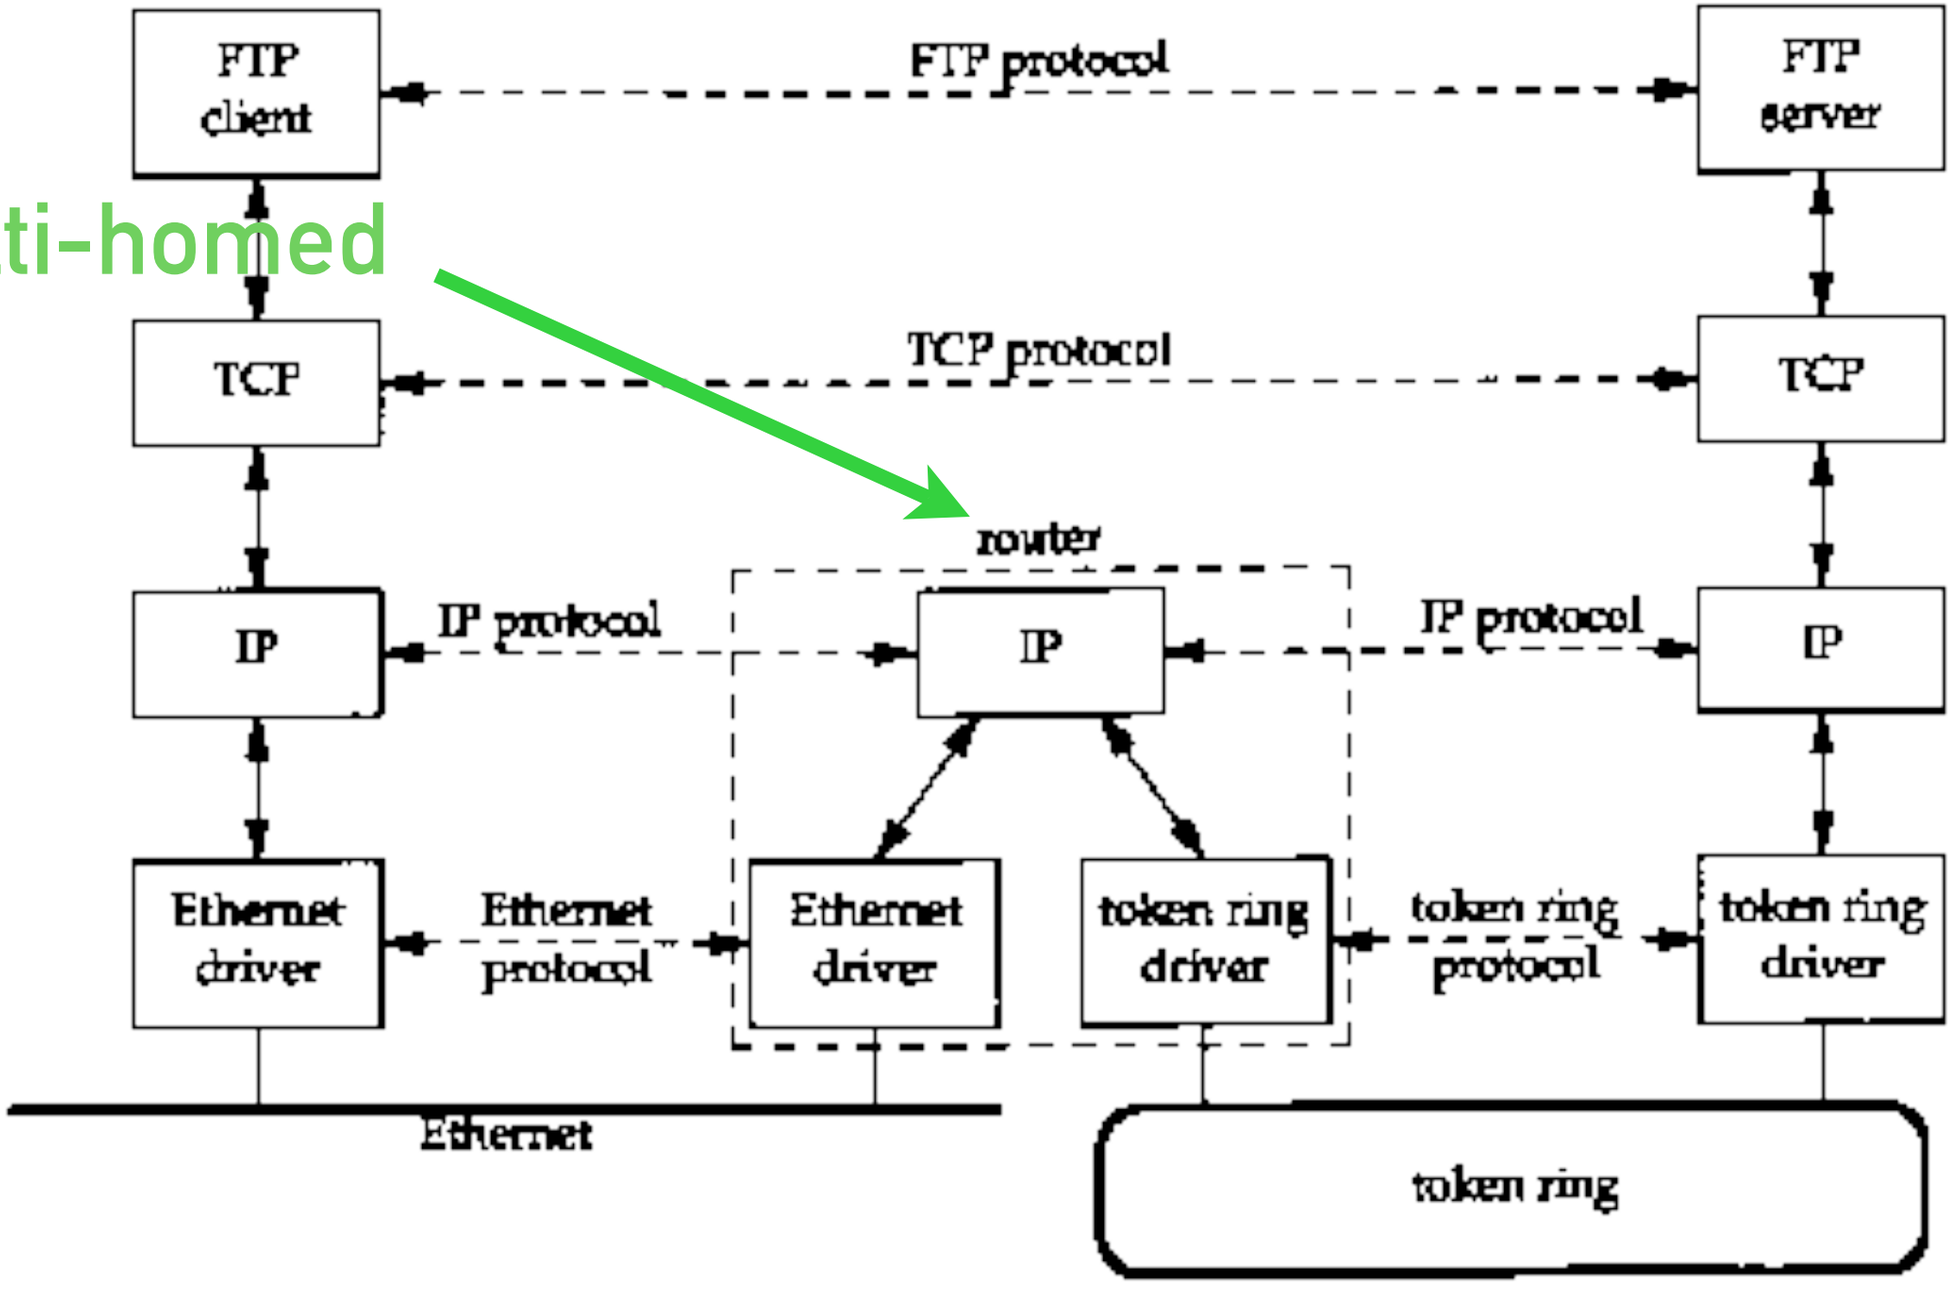

# Layer 3: Routing

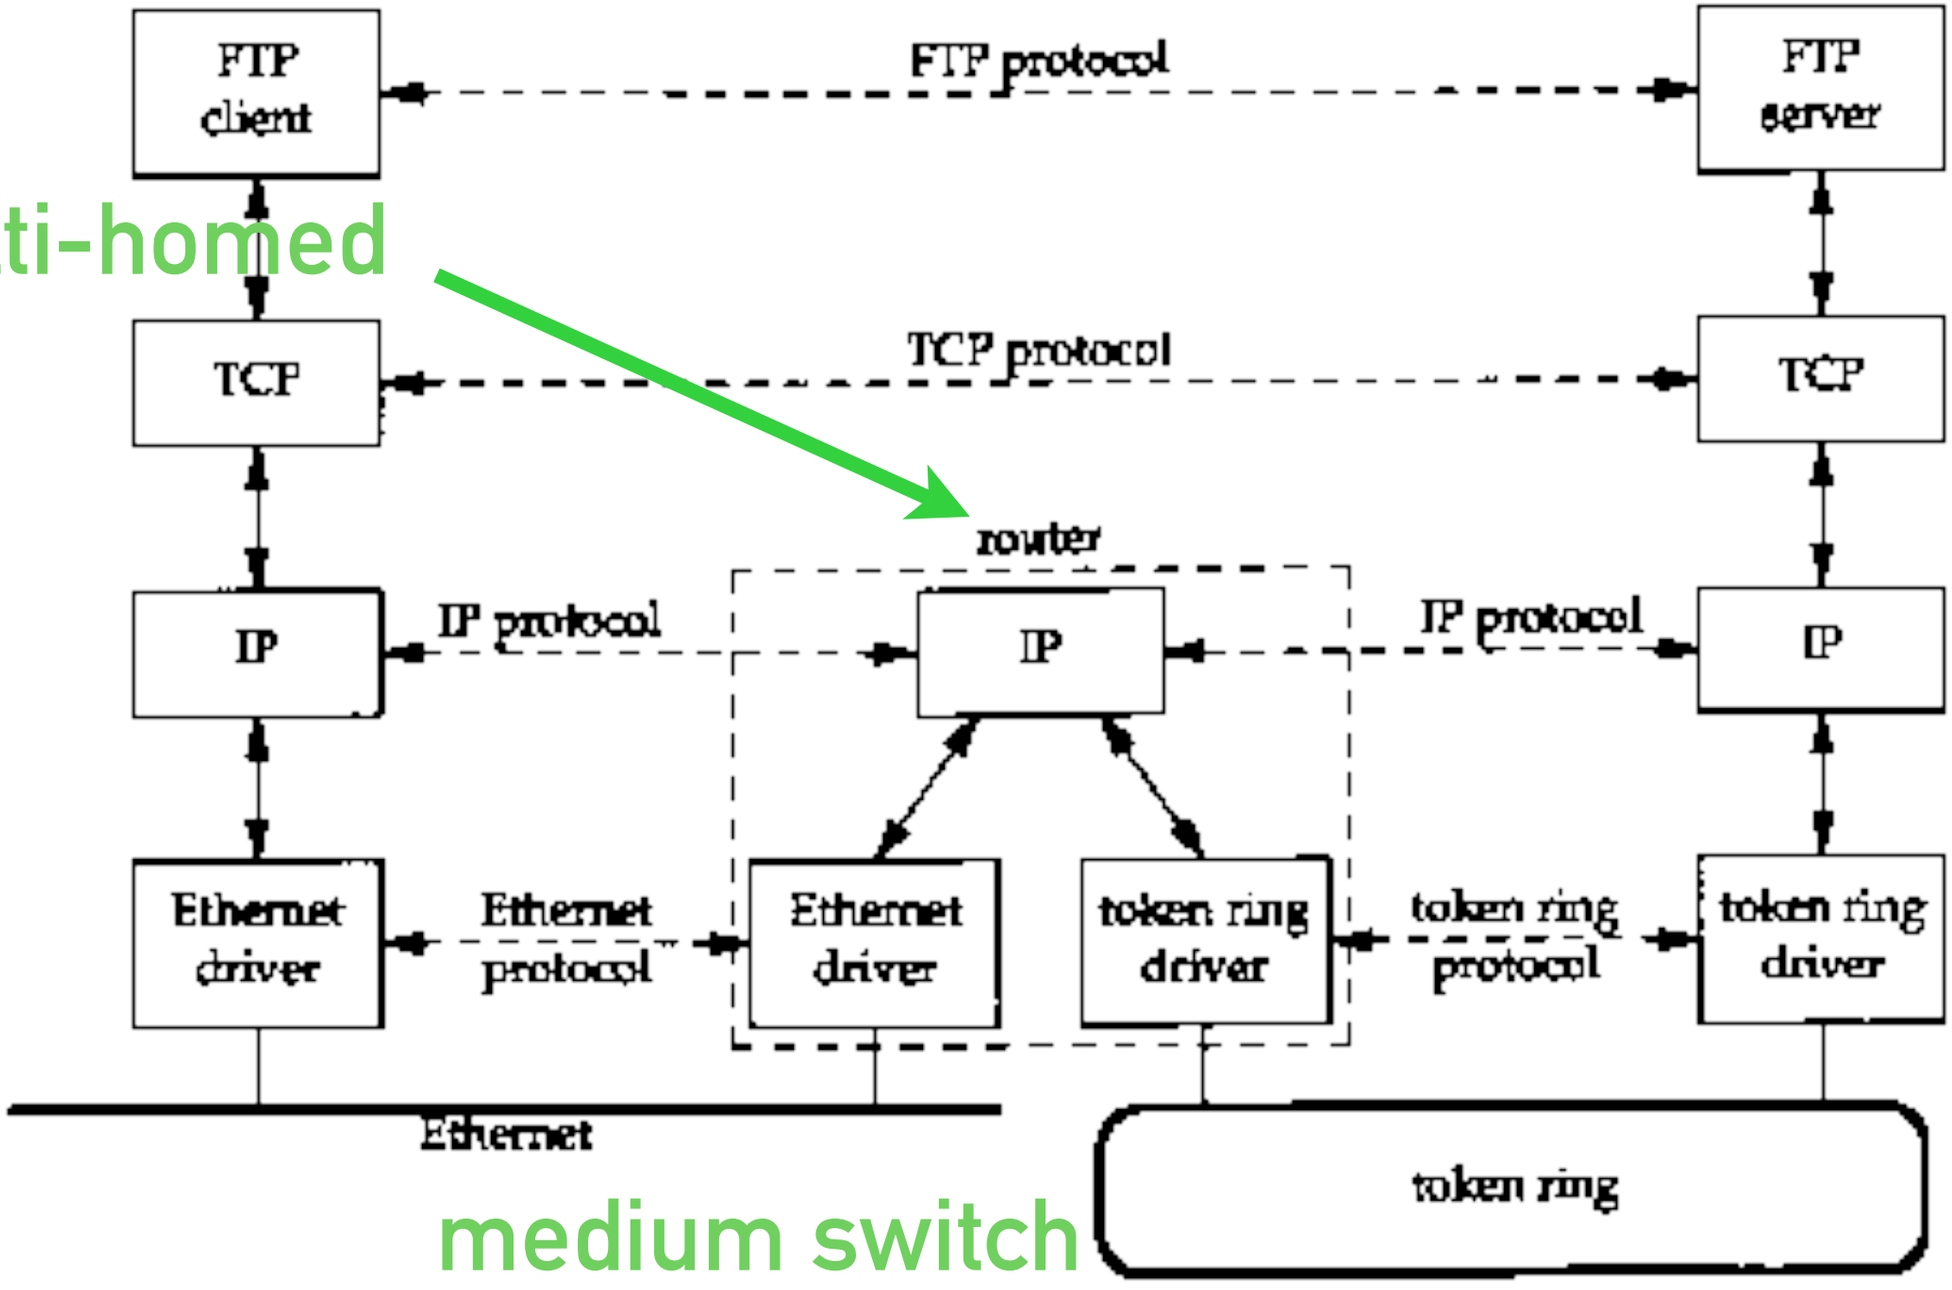


multi-homed

medium switch

# Layer 3: Routing

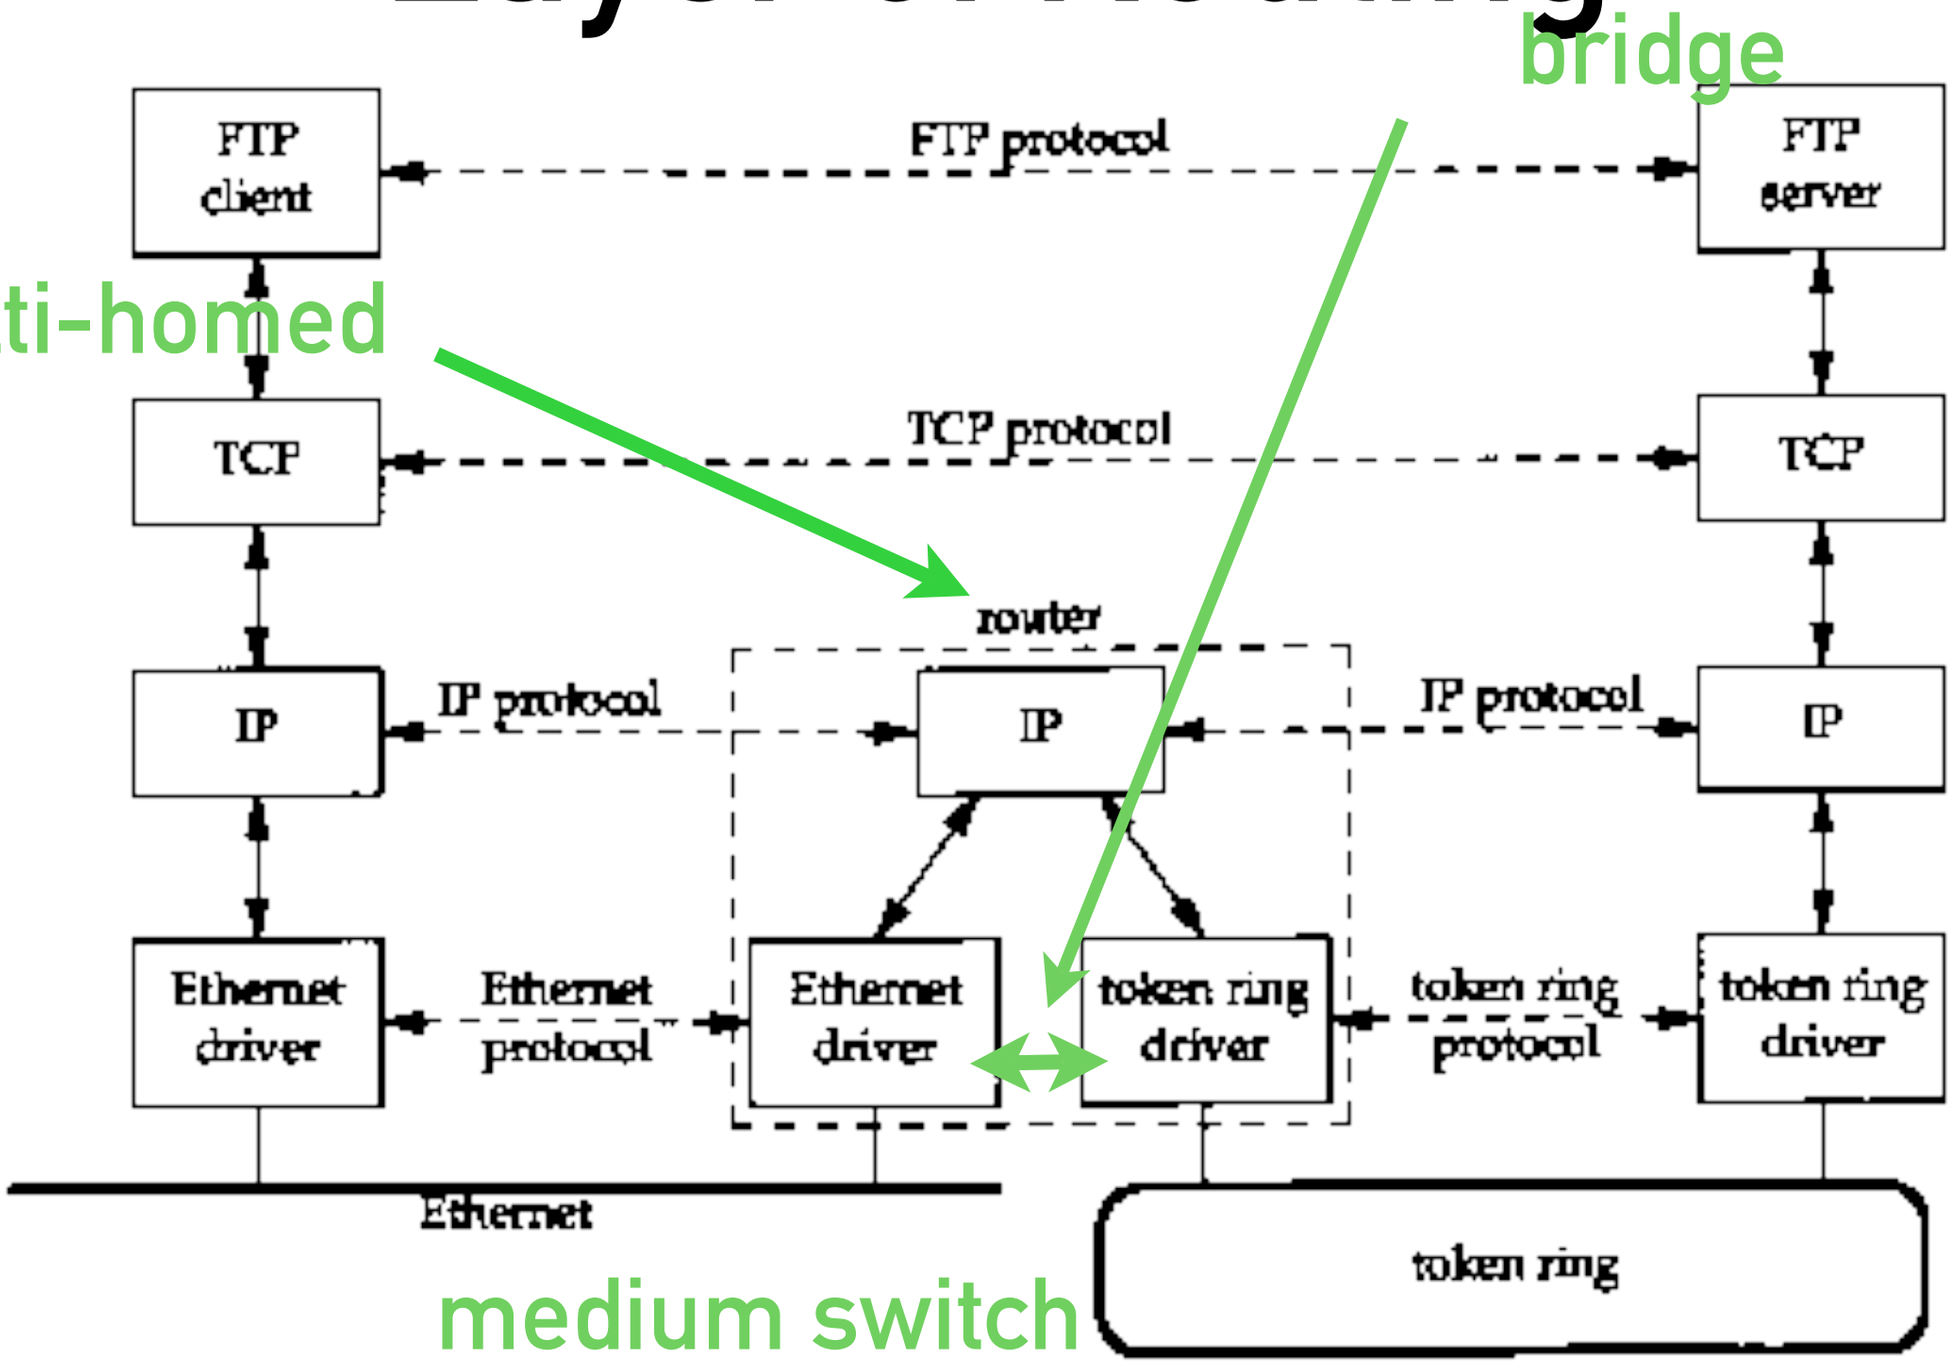

# TCP/IP Layering

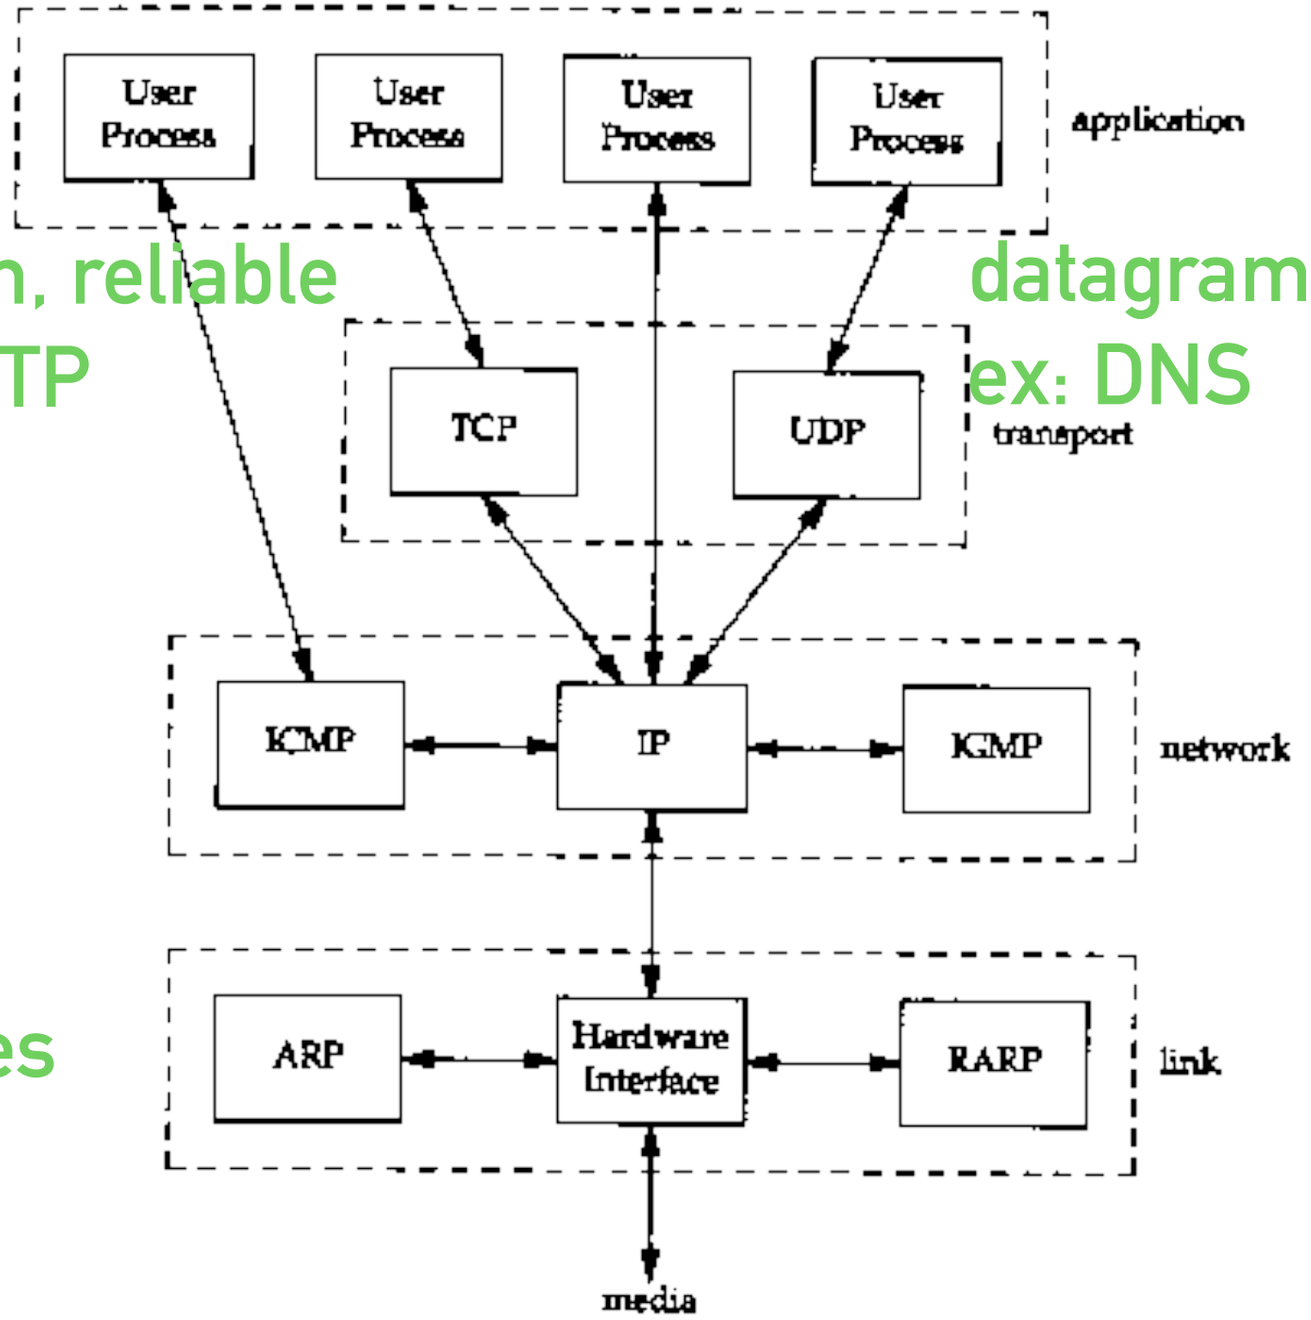


stream, reliable

ex: HTTP

datagram, unreliable

ex: DNS

packets

frames

# Layer 2$^{1/2}$: ARP

- ARP-who has 192.168.100.1?
- ARP-Reply: 192.168.100.1 is at 00:00:c3:fe...
- “glue” between layer 2 and layer 3 addressing

# Addressing

- Every interface on an internet must have a unique Internet address (also called an IP address). These addresses are 32-bit numbers. Instead of using a flat address space such as 1, 2, 3, and so on, there is a structure to Internet addresses.

- Old style: class A, B, C, D

| Class | Range | | |
|---|---|---|---|
| A | **0**.0.0.0 | to | **127**.255.255.255 |
| B | **128**.0.0.0 | to | **191**.255.255.255 |
| C | **192**.0.0.0 | to | **223**.255.255.255 |
| D | **224**.0.0.0 | to | **239**.255.255.255 |
| E | **240**.0.0.0 | to | **247**.255.255.255 |

**Figure 1.6** Ranges for different classes of IP addresses.

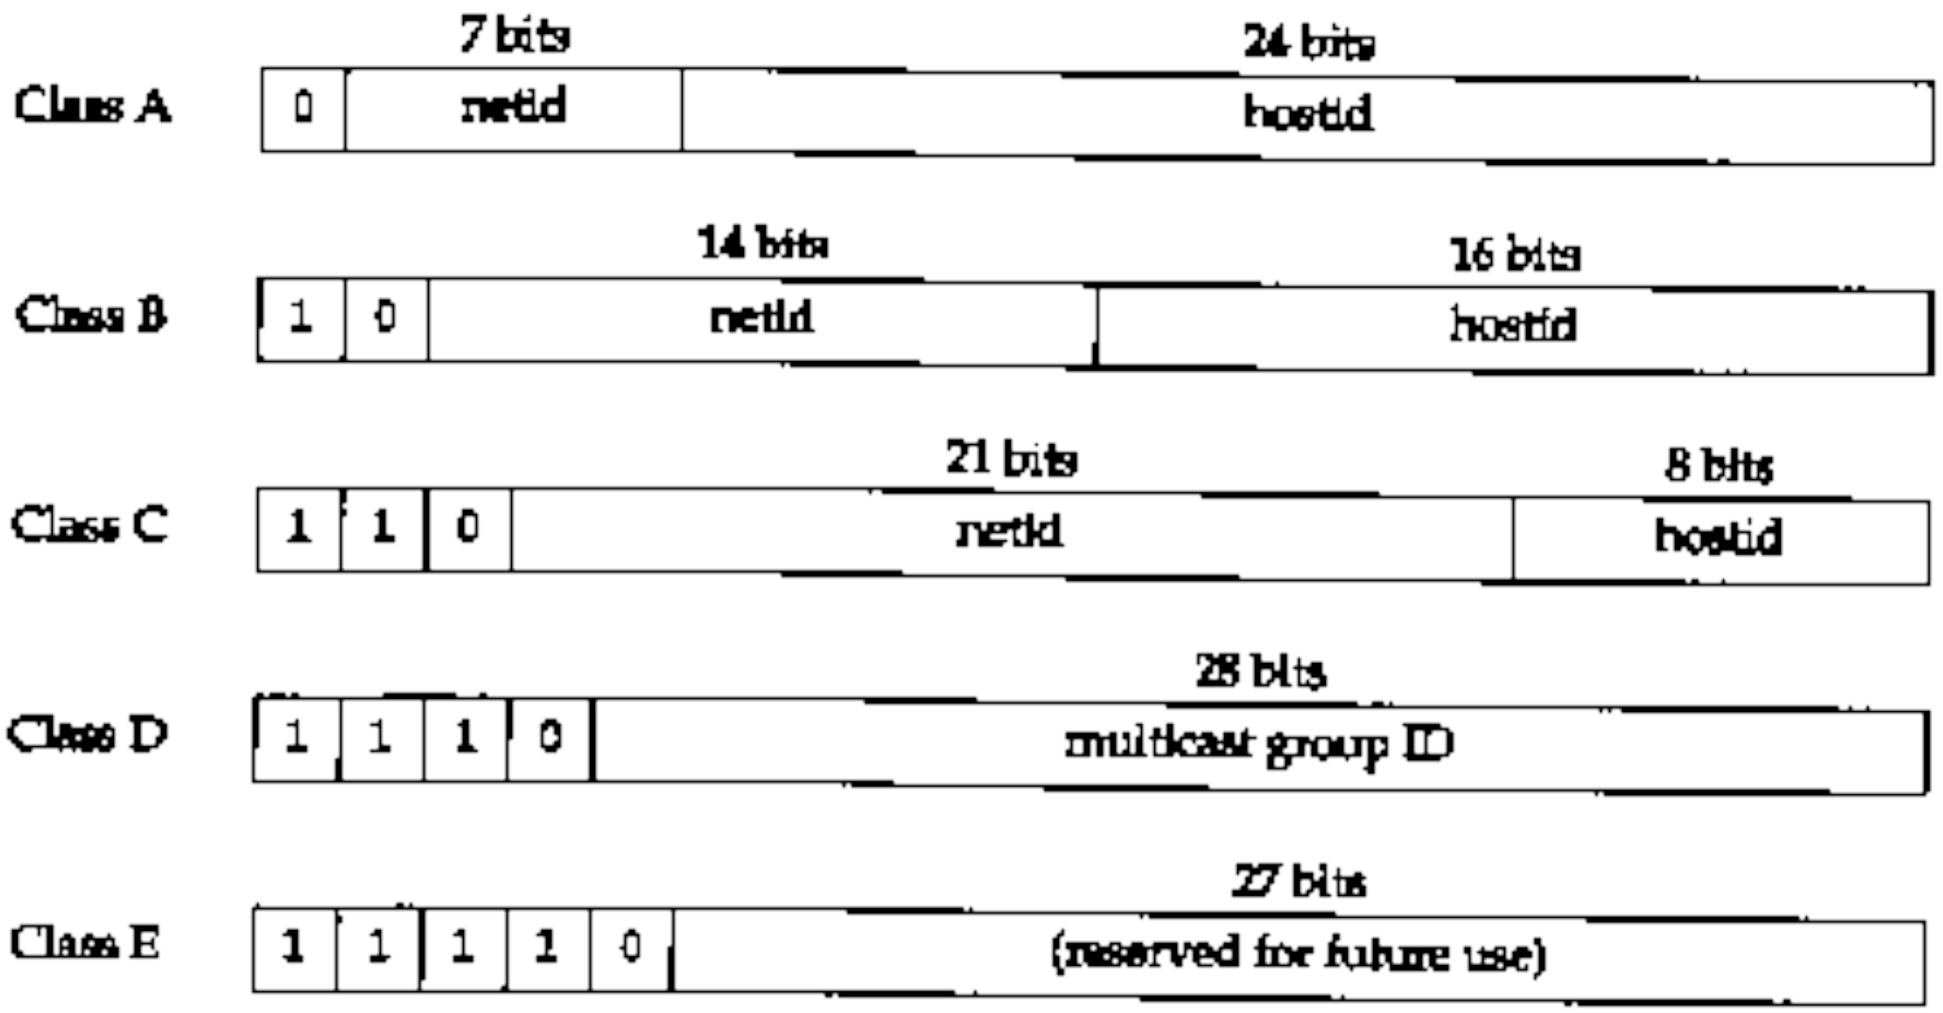
7 bits
24 bits
Class A
0
netid
hostid
14 bits
16 bits
Class B
1
0
netid
hostid
21 bits
8 bits
Class C
1
1
0
netid
hostid
28 bits
Class D
1
1
1
0
multicast group ID
27 bits
Class E
1
1
1
1
0
(reserved for future use)

# MAP OF THE INTERNET

THE IPv4 SPACE, 2006


THIS CHART SHOWS THE IP ADDRESS SPACE ON A PLANE USING A FRACTAL MAPPING WHICH PRESERVES GROUPING -- ANY CONSECUTIVE STRING OF IPs WILL TRANSLATE TO A SINGLE COMPACT, CONTIGUOUS REGION ON THE MAP. EACH OF THE 256 NUMBERED BLOCKS REPRESENTS ONE /8 SUBNET (CONTAINING ALL IPs THAT START WITH THAT NUMBER). THE UPPER LEFT SECTION SHOWS THE BLOCKS SOLD DIRECTLY TO CORPORATIONS AND GOVERNMENTS IN THE 1990's BEFORE THE RIRs TOOK OVER ALLOCATION.

0 1 14 15 16 19 →
3 2 13 12 17 18
4 7 8 11
5 6 9 10

= UNALLOCATED BLOCK

# Addressing CIDR

- Classless Internet domain routing
- “It is worth reiterating that a multihomed host will have multiple IP addresses: one per interface.”
- New style: IP Addr + netmask length 192.168.100.1/24

# Encapsulation

- Down the stack: every layer adds information - headers and trailers
- Unit of TCP -> IP == “TCP segment”
- Unit of UDP, IP == “datagram”
- Unit of link layer == “frame”. Ethernet: between 46 and 1500 Bytes

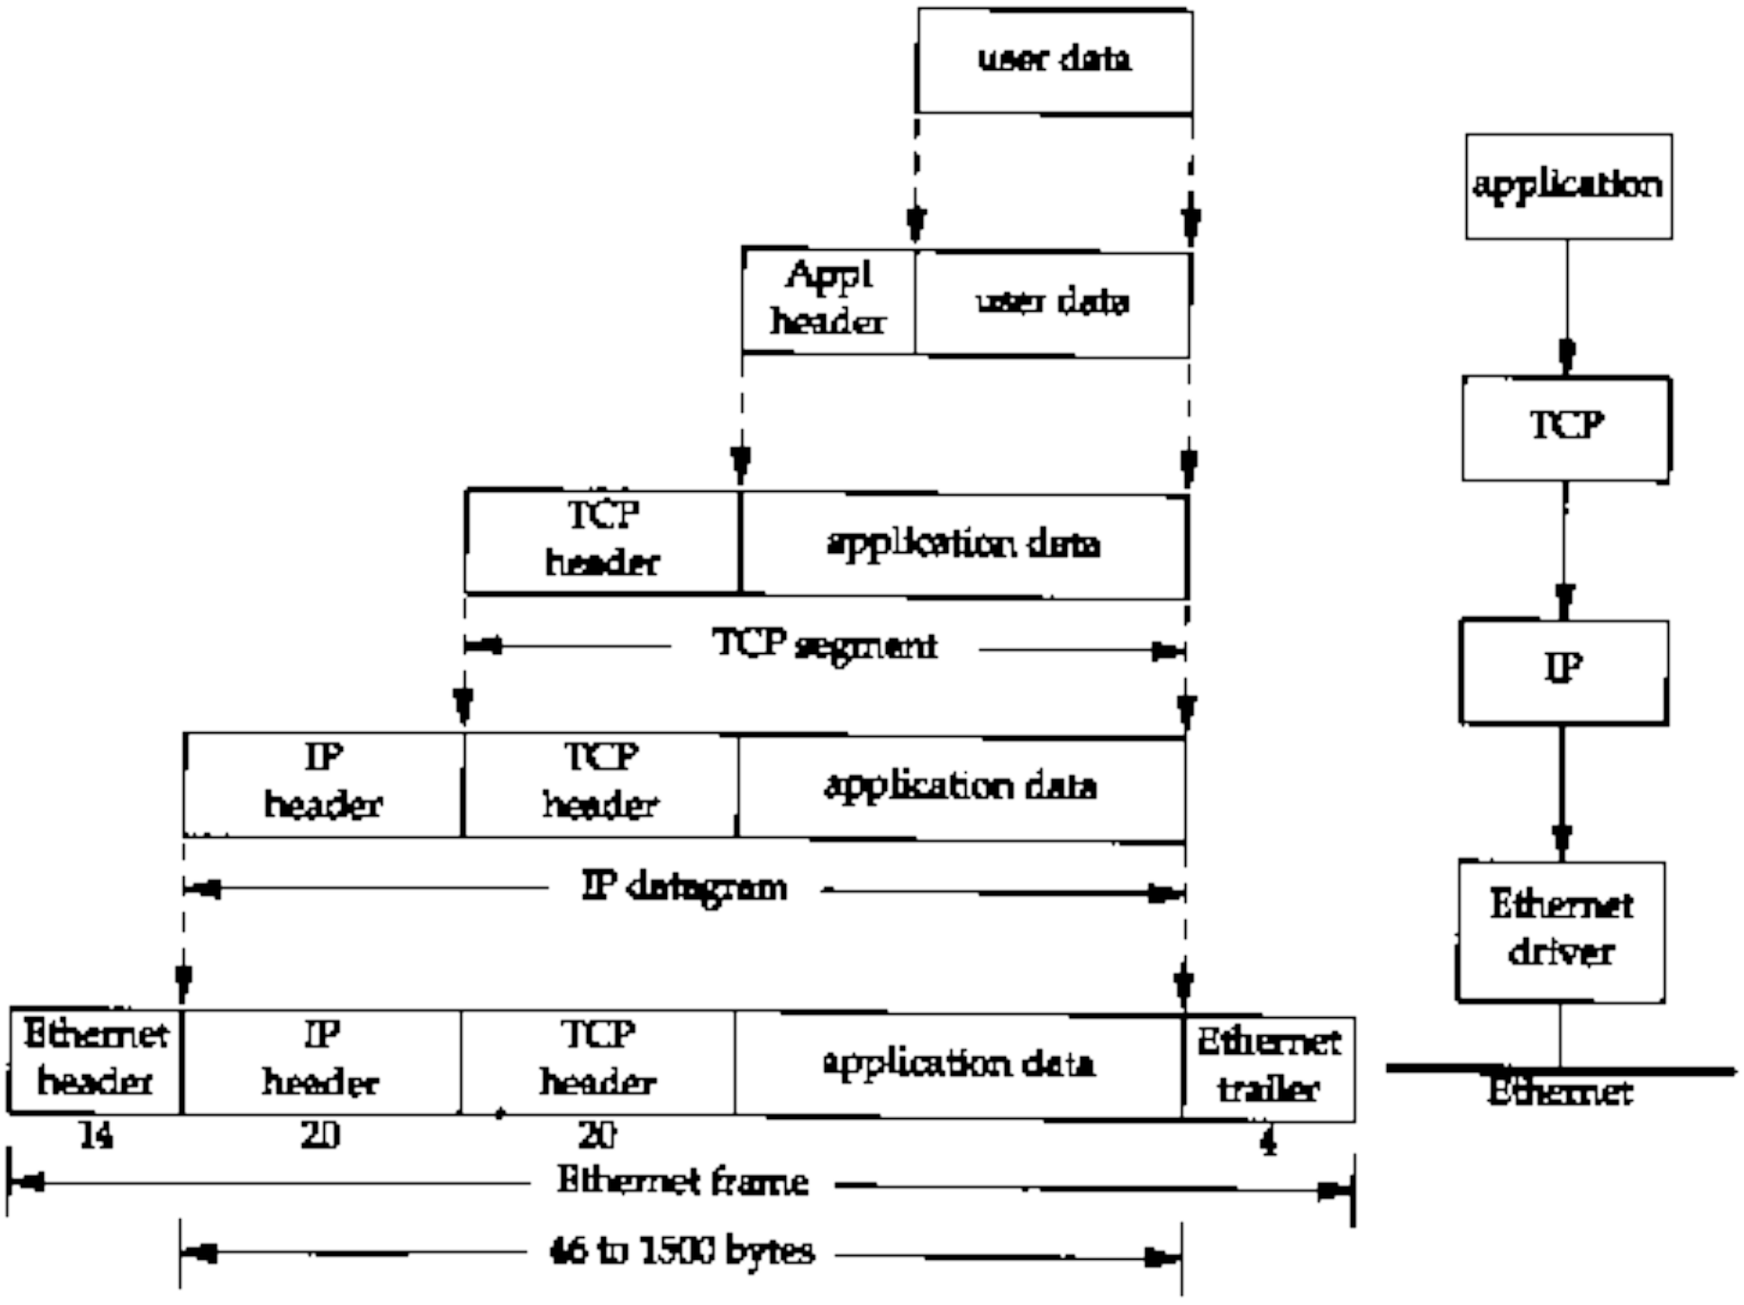
user data
Appl header
user data
TCP header
application data
TCP segment
IP header
TCP header
application data
IP datagram
Ethernet header
IP header
TCP header
application data
Ethernet trailer
14
20
20
4
Ethernet frame
46 to 1500 bytes
application
TCP
IP
Ethernet driver
Ethernet

# What's in the headers?

- Many things
- For finding where stuff belongs:
- IP: protocol field:
  1 = ICMP, 2 = IGMP, 6 = TCP, 17 = UDP
- TCP/UDP: 16 bit port number
- Frames: 16 bit frame type: ARP, IP, RARP,...

# Demultiplexing

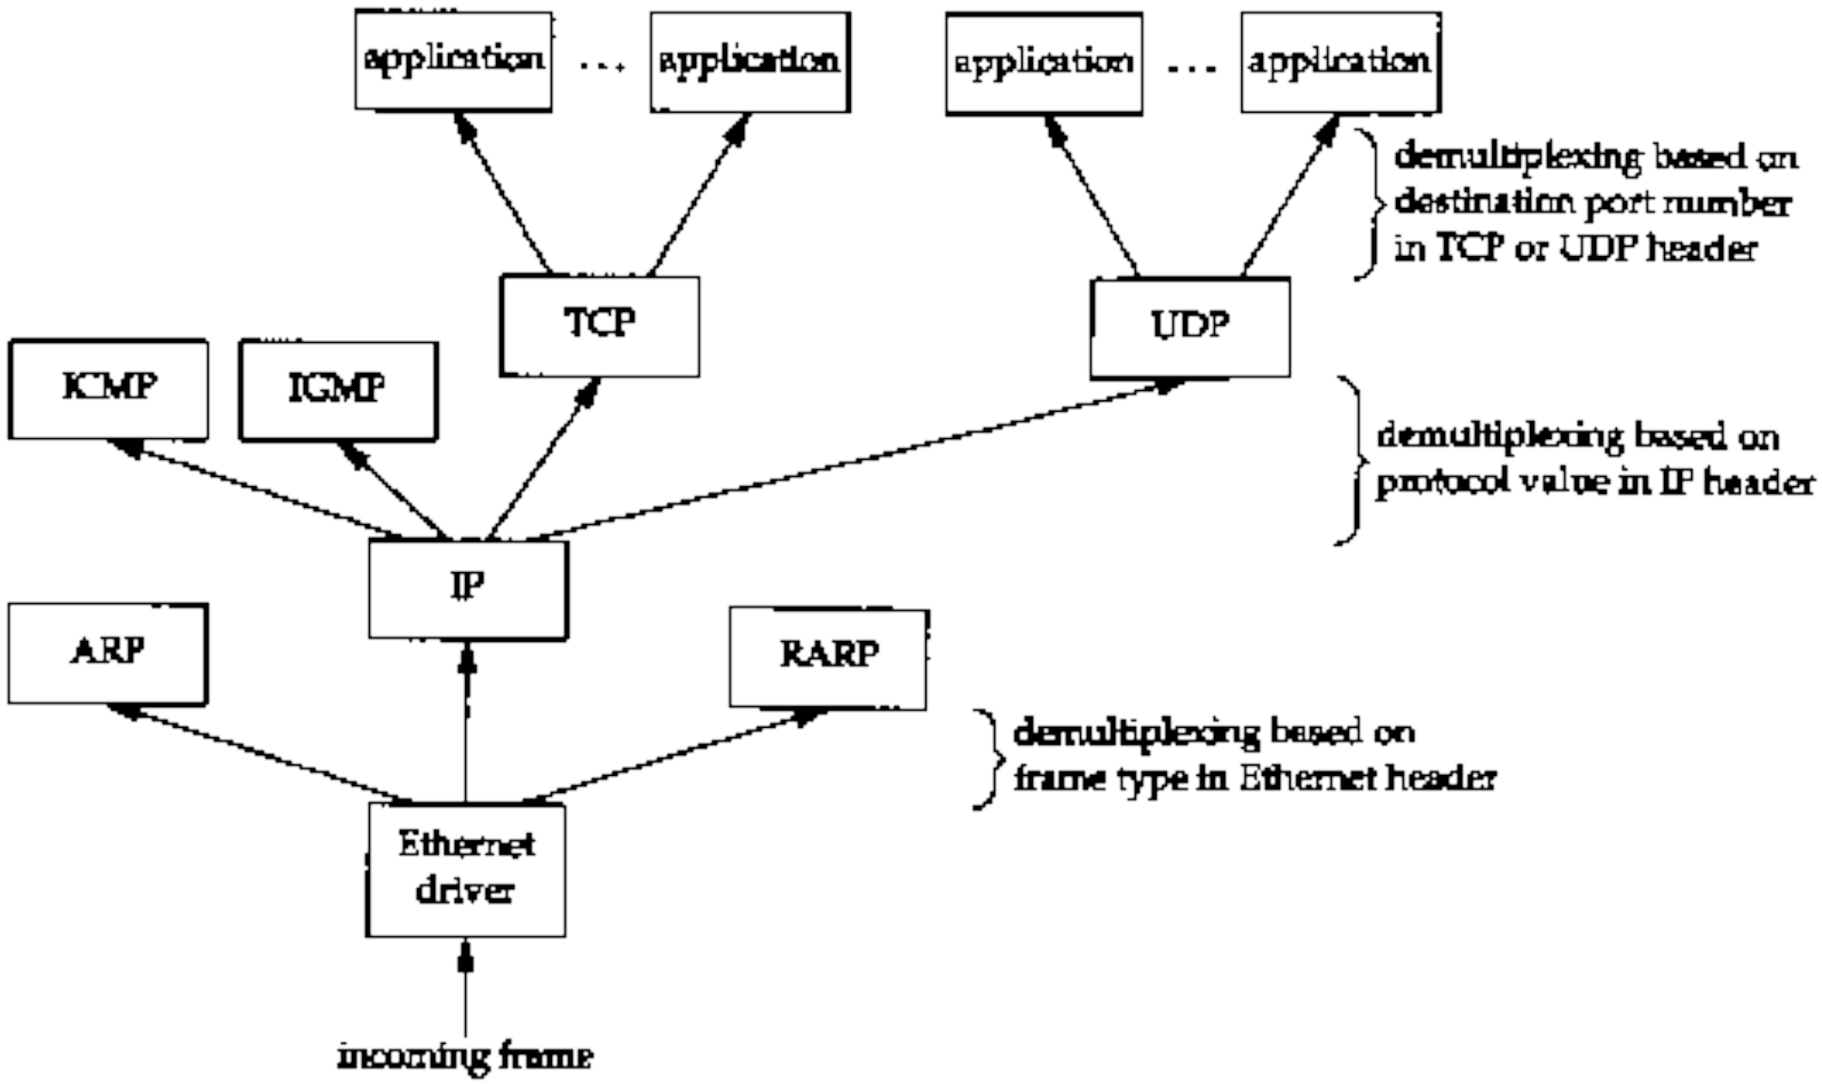

# port numbers

```
sun % grep telnet /etc/services
telnet 23/tcp
sun % grep domain /etc/services
domain 53/udp
domain 53/tcp
```

- reserved ports
- Analogy: district == net, street+house == IP, flat == port

# Standardization process

- ISOC - Internet Society - promo group
- IAB - Internet Architecture Board technical oversight
- IETF - Internet Engineering Task force RFCs
- IRTF - Internet Research Task force long term research goals

# RFCs

- Request for Comments
- undergoes many many drafts
- discussed at IETF meetings
- google:// rfc 3626

# Programing interface

- socket library
- man socket
- man ip

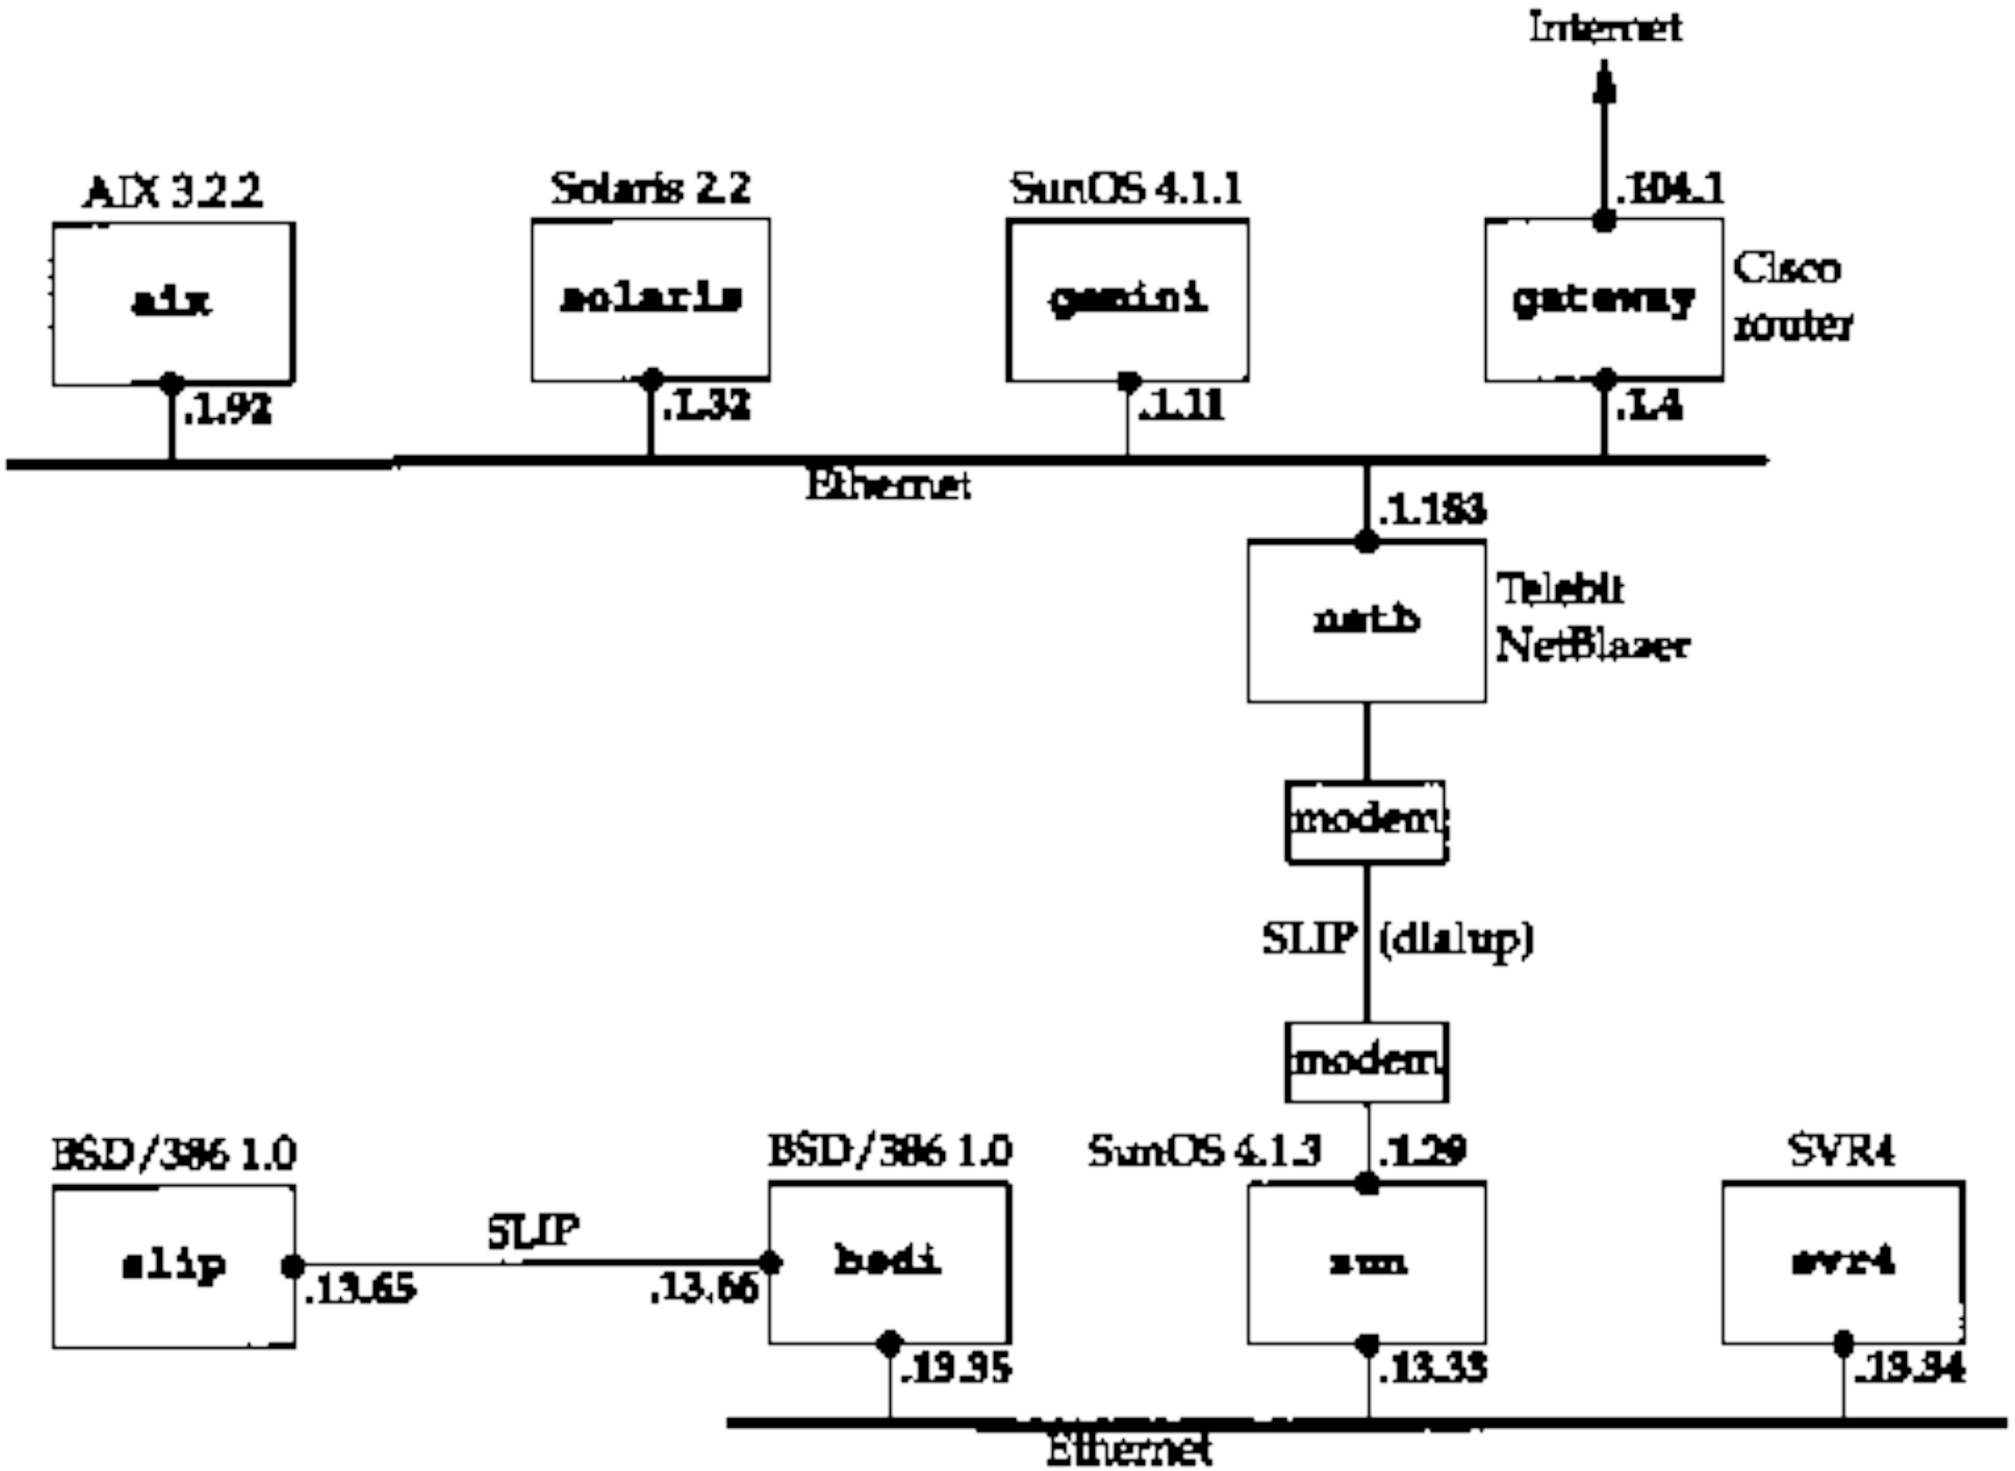
Internet
AIX 3.2.2
Solaris 2.2
SunOS 4.1.1
.134.1
aix
solaris
gemini
gateway
Cisco
router
.1.92
.1.32
.1.11
.1.4
Ethernet
.1.183
netb
Telebit
NetBlazer
modem
SLIP (dialup)
modem
BSD/386 1.0
BSD/386 1.0
SunOS 4.1.3
.1.29
SVR4
slip
SLIP
bodi
sun
svr4
.13.65
.13.66
.13.95
.13.33
.13.34
Ethernet