

The background of the slide is a spiral-bound notebook with a light beige, textured cover. The spiral binding is visible on the left side. The text is centered on the page.

Lecture 8

Berkeley Socket Programming

Berkeley Sockets

[Unix Socket Programming FAQ](#)

[Beej's Guide to Network Programming](#)

Metaphors

- Postal Service

- Address

- Name, Street, City, State, Zip Code

- Return Address

- Network of Post Offices

- Local Post Offices

- Lost Mail

- Indeterminacy of Order

- No confirmation of delivery for regular mail

- No failed delivery notification

Metaphors (continued)

- Toll call from one hotel room to another (circa 1945), (or, London today)
 - call down to local hotel operator
 - tell her the area code and number to call of the remote hotel
 - hotel operator calls long distance operator, who is listening for incoming calls
 - long distance operator calls remote hotel
 - remote hotel operator picks up, as she has been listening for calls, and routes the call to your friend's room
 - You and your friend are now talking directly

You Already Use Sockets

- echo (7), smtp (25), ftp (21), ssh (22)
 - telnet time.mit.edu 13
- cat /etc/services | grep [telnet | ssh | ftp | echo | etc.]
- daytime (13) (telnet time.mit.edu 13)
- email (SMTP) (port 25): (telnet direct to SMTP server)

```
mark@sheik:~$ telnet smpt.cs.uchicago.edu 25
HELO sheik.cs.uchicago.edu
MAIL FROM: <soner@cs.uchicago.edu>
RCPT TO: <cspp51081@mailman.cs.uchicago.edu>
DATA
Hello, this is talking directly to Postfix.
.
QUIT
```

The Fundamentals

- The Computer Systems Research Group (CSRG) at the University of California Berkeley gave birth to the Berkeley Socket API (along with its use of the TCP/IP protocol) with the 4.2BSD release in 1983.
 - A Socket is comprised of:
 - a 32-bit node address (IP address or FQDN)
 - a 16-bit port number (like 7, 21, 13242)
 - Example: 192.168.31.52:1051
 - The 192.168.31.52 host address is in “IPv4 dotted-quad” format, and is a decimal representation of the hex network address 0xc0a81f34

Port Assignments (less /etc/services)

- Ports 0 through 1023 are reserved, *privileged* ports, defined by TCP and UDP well known port assignments (telnet, ssh, http, ftp, etc.)
- Ports 1024 through 49151 are ports *registered* by the IANA (Internet Assigned Numbers Authority), and represent second tier common ports (socks (1080), WINS (1512), kermit (1649), Java RMI Registry (1099)), Microsoft SQL Server (1433), Oracle (1529)
- Ports 49152 through 65535 are *ephemeral* ports, available for temporary client usage

Protocol Stacks

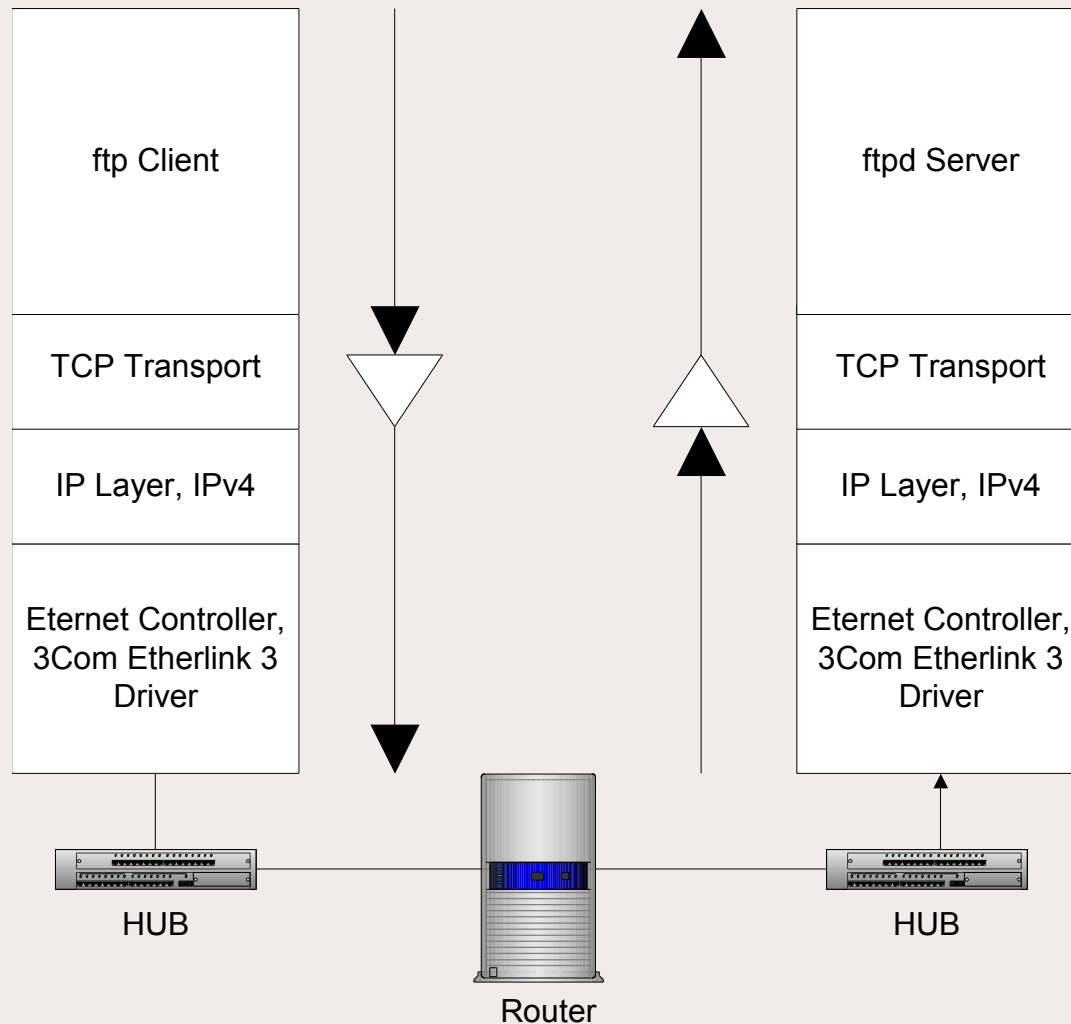
Application (Telnet, ftp, etc.)
Presentation (MIDI, HTML, EBCDIC)
Session (RPC, Netbios, Appletalk, DECnet)
Transport (TCP, UDP)
Network (IPv4, IPv6, IPX)
Datalink (Ethernet, Token Ring, ATM, PPP)
Physical (V.24, 802.3, Ethernet RJ45)

OSI Model
(Tannenbaum, 1988)

Application (Telnet, ftp, etc.)
Transport (TCP, UDP)
IP Layer (IPv4, IPv6)
Device Driver and Hardware (twisted pair, NIC)

Internet Protocol Suite

Protocol Communication



Common Protocols

Application	ICMP	UDP	TCP
Ping	•		
Traceroute	•		
DHCP		•	
NTP		•	
SNMP		•	
SMTP			•
Telnet			•
FTP			•
HTTP			•
NNTP			•
DNS		•	•
NFS		•	•
Sun RPC		•	•

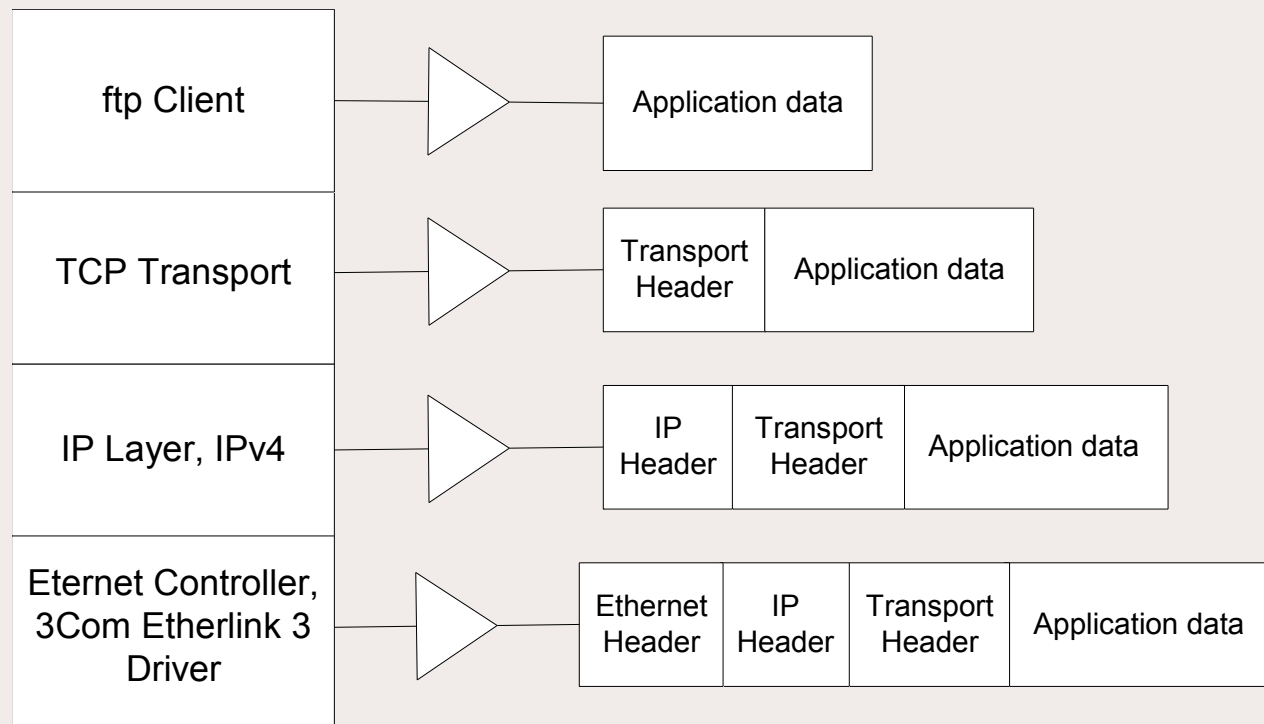
ICMP: Internet Control Message Protocol

UDP: User Datagram Protocol

TCP: Transmission Control Protocol

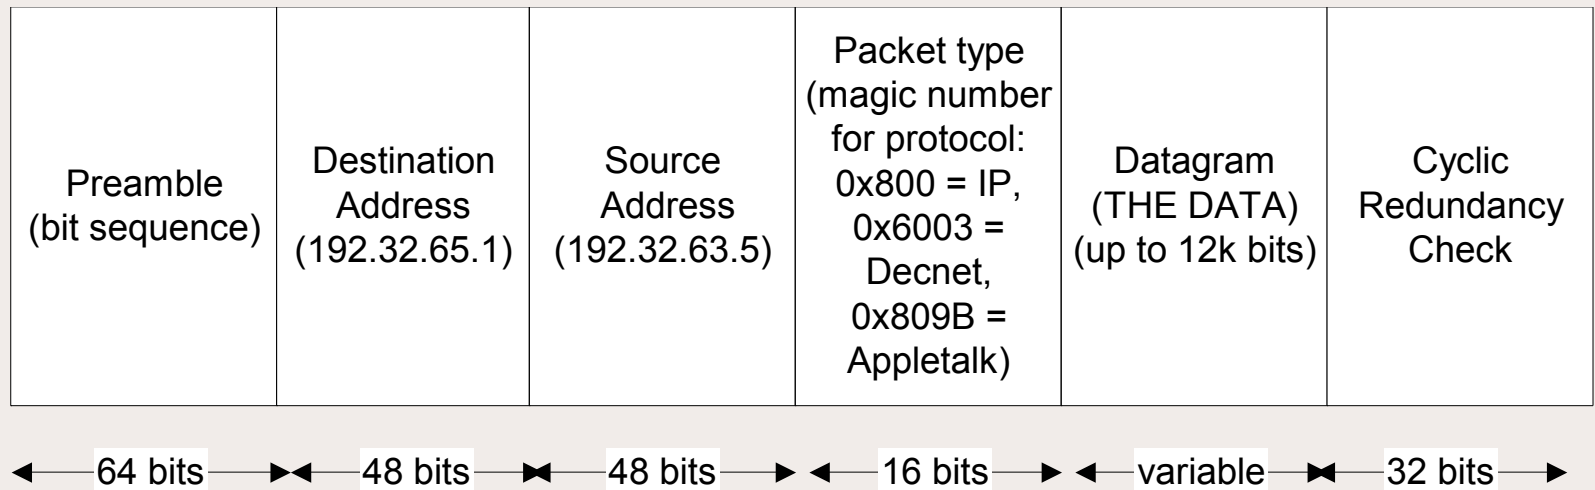
Data Encapsulation

- Application puts data out through a socket
- Each successive layer wraps the received data with its own header:



The Hardware (Ethernet) Layer

- Responsible for transferring frames (units of data) between machines on the same physical network

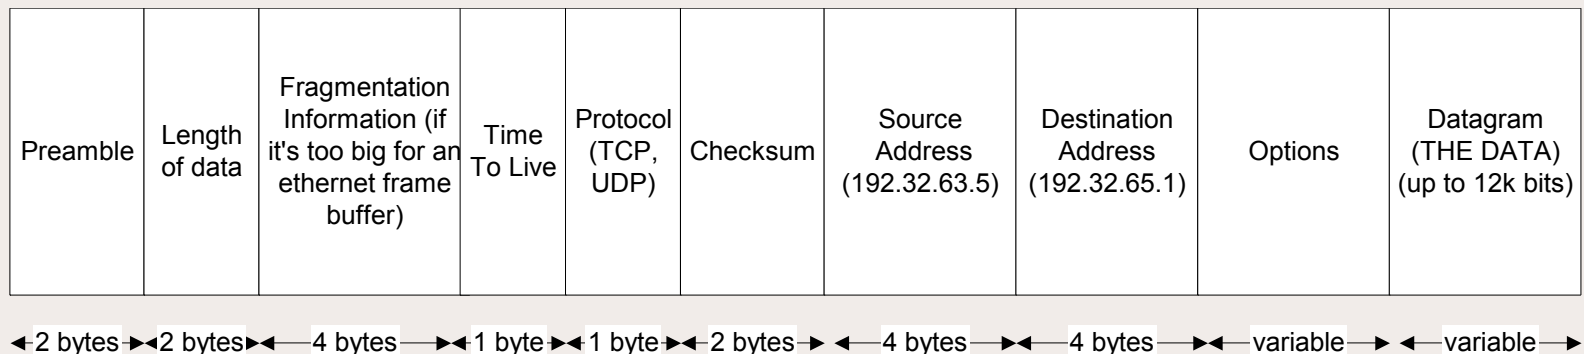


The IP Layer

- The IP layer allows packets to be sent over gateways to machines not on the physical network
- Addresses used are IP addresses, 32-bit numbers divided into a network address (used for routing) and a host address
- The IP protocol is connectionless, implying:
 - gateways route discrete packets independently and irrespective of other packets
 - packets from one host to another may be routed differently (and may arrive at different times)
 - non-guaranteed delivery

IP Datagram Format

- Packets may be broken up, or *fragmented*, if original data is too large for a single packet (Maximum Transmission Unit is currently 12k bits, or 1500 Bytes)
- Packets have a Time To Live, number of seconds/rounds it can bounce around aimlessly among routers until it's killed



The Transport Layer

- Unix has two common transports
 - User Datagram Protocol
 - record protocol
 - connectionless, broadcast
 - *Metaphor*: Postal Service
 - Transmission Control Protocol
 - byte stream protocol
 - direct connection-oriented
 - *Metaphor*: Phone Service circa 1945
 - “Sarah, this is Andy, get me Barney please.”

The Transport Layer:

UDP Protocol

- Connectionless, in that no long term connection exists between the client and server. A connection exists only long enough to deliver *a single packet* and then the connection is severed.
- No guaranteed delivery (“best effort”)
- Fixed size boundaries, sent as a single “fire and forget message”. Think *announcement*.
- No built-in acknowledgement of receipt

The Transport Layer:

UDP Protocol

- No built-in order of delivery, random delivery
- Unreliable, since there is no acknowledgement of receipt, there is no way to know to resend a lost packet
- Does provide checksum to guarantee integrity of packet data
- *Fast* and Efficient

The Transport Layer:

TCP Protocol

- TCP *guarantees delivery* of packets in order of transmission by offering acknowledgement and retransmission: it will automatically resend after a certain time if it does not receive an ACK
- TCP promises *sequenced delivery* to the application layer, by adding a sequence number to every packet. Packets are reordered by the receiving TCP layer before handing off to the application layer. This also aides in handling “duplicate” packets.

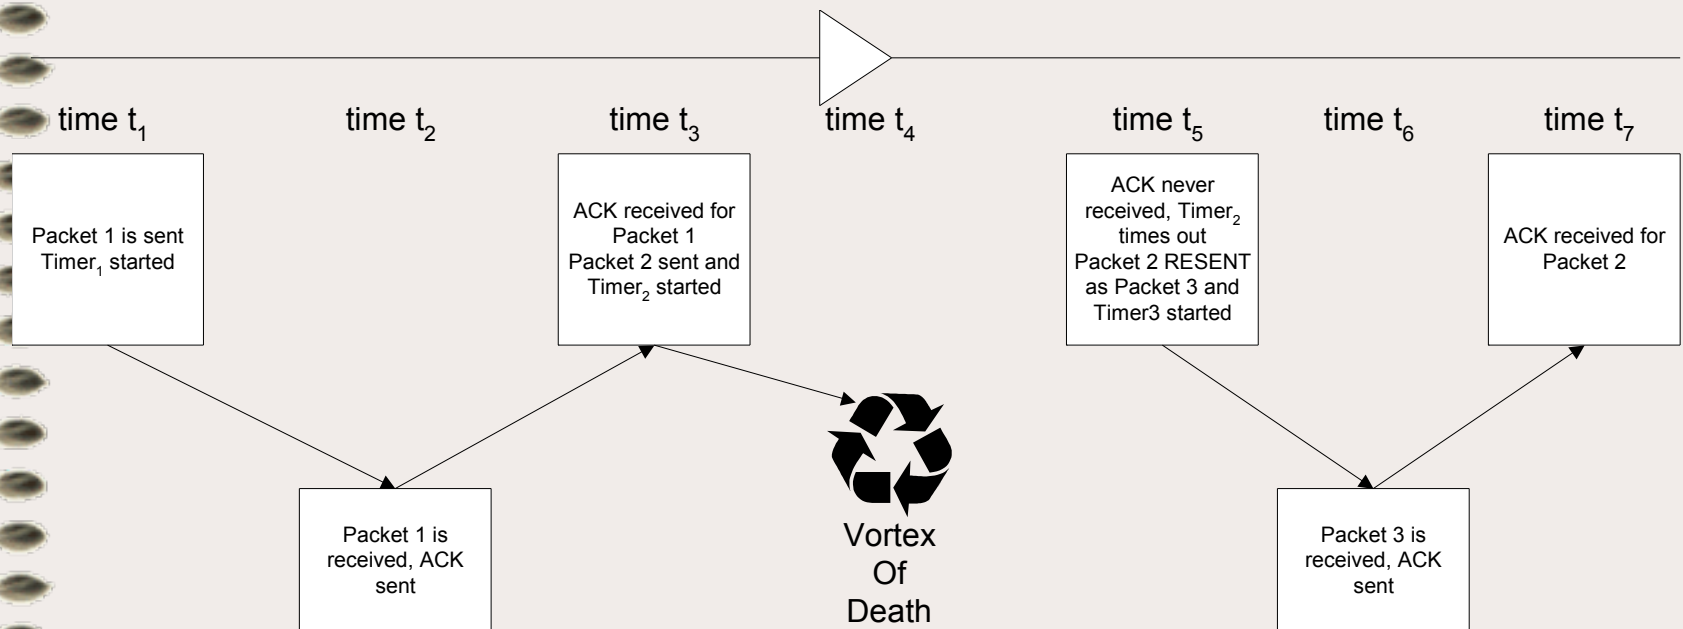
The Transport Layer:

TCP Protocol

- Pure stream-oriented connection, it does not care about message boundaries
- A TCP connection is full duplex (bidirectional), so the same socket can be read and written to (cf. half duplex pipes)
- Provides a checksum that guarantees packet integrity

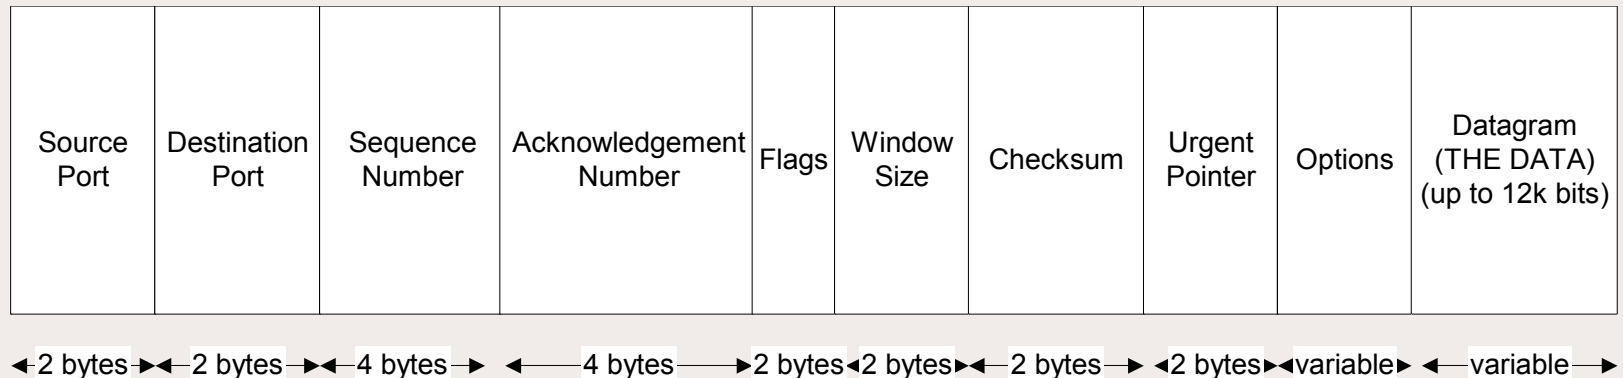
TCP's Positive Acknowledgement with Retransmission

- TCP offers acknowledgement and retransmission: it will automatically resend after a certain time if it does not receive an ACK
- TCP offers *flow control*, which uses a “sliding window” (in the TCP header) will allow a *limited* number of non-ACKs on the net during a given interval of time. This increases the overall bandwidth efficiency. This window is dynamically managed by the recipient TCP layer.



TCP Datagram Format

- Source and Destination addresses
- Sequence Number tells what byte offset within the overall data stream this segment applies
- Acknowledgement number lets the recipient set what packet in the sequence was received ok.



Socket Domain Families

- There are several significant socket domain families:
 - Internet Domain Sockets (AF_INET)
 - implemented via IP addresses and port numbers
 - Unix Domain Sockets (AF_UNIX)
 - implemented via filenames (think “named pipe”)
 - Novell IPX (AF_IPX)
 - AppleTalk DDS (AF_APPLETALK)
 - *Example:* `~mark/pub/51081/sockets/linux/socketpairs.c`

Creating a Socket

```
#include <sys/types.h>
#include <sys/socket.h>
int socket(int domain, int type, int protocol);
```

- **domain** is one of the *Address Families* (AF_INET, AF_UNIX, etc.)
- **type** defines the communication protocol semantics, usually defines either:
 - SOCK_STREAM: connection-oriented stream (TCP)
 - SOCK_DGRAM: connectionless, unreliable (UDP)
- **protocol** specifies a particular protocol, just set this to 0 to accept the default (PF_INET, PF_UNIX) based on the domain

SERVER

Create socket

bind a port to the socket

listen for incoming connections

accept an incoming connection

read from the connection

write to the connection

CLIENT

Create socket

connect to server's port

write to the connection

read from the connection

loop

loop

close connection

- Connection-oriented socket connections
- Client-Server view

Server Side Socket Details

SERVER

Create socket

```
int socket(int domain, int type, int protocol)
sockfd = socket(PF_INET, SOCK_STREAM, 0);
```

bind a port to the socket

```
int bind(int sockfd, struct sockaddr *server_addr, socklen_t length)
bind(sockfd, &server, sizeof(server));
```

listen for incoming connections

```
int listen( int sockfd, int num_queued_requests)
listen( sockfd, 5);
```

accept an incoming connection

```
int accept(int sockfd, struct sockaddr *incoming_address, socklen_t length)
newfd = accept(sockfd, &client, sizeof(client)); /* BLOCKS */
```

read from the connection

```
int read(int sockfd, void * buffer, size_t buffer_size)
read(newfd, buffer, sizeof(buffer));
```

write to the connection

```
int write(int sockfd, void * buffer, size_t buffer_size)
write(newfd, buffer, sizeof(buffer));
```

Client Side Socket Details

CLIENT

Create socket

```
int socket(int domain, int type, int protocol)
sockfd = socket(PF_INET, SOCK_STREAM, 0);
```

connect to Server
socket

```
int connect(int sockfd, struct sockaddr *server_address, socklen_t length)
connect(sockfd, &server, sizeof(server));
```

write to the
connection

```
int write(int sockfd, void * buffer, size_t buffer_size)
write(sockfd, buffer, sizeof(buffer));
```

read from the
connection

```
int read(int sockfd, void * buffer, size_t buffer_size)
read(sockfd, buffer, sizeof(buffer));
```

Setup for an Internet Domain Socket

```
struct sockaddr_in {
    sa_family_t sin_family;
    unsigned short int sin_port;
    struct in_addr sin_addr;
    unsigned char pad[...];
};
```

- `sin_family` is set to *Address Family* `AF_INET`
- `sin_port` is set to the port number you want to *bind* to
- `sin_addr` is set to the IP address of the machine you are binding to (struct `in_addr` is a wrapper struct for an unsigned long). `INADDR_ANY` supports all interfaces (since a given machine may have multiple interface cards)
- ignore padding

Setup for A Unix Domain Socket

```
struct sockaddr_un {  
    sa_family_t sun_family;  
    char sun_path[UNIX_PATH_MAX];  
};
```

- `sun_family` is set to *Address Family* `AF_UNIX`
- `sun_path` is set to a *UNIX pathname* in the filesystem

Reading From and Writing To Stream Sockets

- Sockets, like everything else, are like *files*:
 - low level IO:
 - read() system call
 - write() system call
 - higher level IO:
 - int recv(int socket, char *buf, int len, int flags);
 - blocks on read
 - returns 0 when other connection has terminated
 - int send(int socket, char *buf, int len, int flags);
 - returns the number of bytes *actually sent*
 - where flags may be one of:
 - MSG_DONTROUTE (don't route out of localnet)
 - MSG_OOB (out of band data (think *interruption*))
 - MSG_PEEK (examine, but don't remove from stream)

Closing a Socket Session

- `int close(int socket);`
 - closes read/write IO, closes socket file descriptor
- `int shutdown(int sockfd, int how);`
 - where `how` is:
 - 0: no more receives allowed
 - 1: no more sends are allowed
 - 2: disables both receives and sends (but doesn't close the socket, use `close()` for that)
- *Example:* `hangserver.c` (hangman game)

Host and Network Byte Ordering

- Different computer architectures store numbers differently:
 - Little Endian architectures (like VAX, Intel) store the least significant byte first
 - This means that within a (2-byte) word, the *least* significant byte is stored *first*, that is, at the *lowest* byte address
 - Big Endian architectures (like Sun Sparc, Motorola 68000, PowerPC) store the most significant byte appearing first
 - This means that within a (2-byte) word, the *most* significant byte is stored *first*, that is, at the *lowest* byte address
- *examples: ~mark/pub/51081/byteorder/linux/Endian.sh*
and
~mark/pub/51081/byteorder/solaris/Endian.sh
- *(on nunki.cs.uchicago.edu)*

Why This Matters

- TCP/IP mandates that *big-endian* byte ordering be used for transmitting protocol information
- This means that *little-endian* machines will need to *convert* ip addresses and port numbers into *big-endian* form in order to communicate successfully
- Note that big-endian architectures don't actually have to do anything, because they already meet the specification

What's To Be Done About It?

- Several functions are provided to allow you to easily convert between host and network byte ordering, and they are:
 - to translate 32-bit numbers (i.e. IP addresses):
 - unsigned long htonl(unsigned long hostlong);
 - unsigned long ntohl(unsigned long netlong);
 - to translate 16-bit numbers (i.e. Port numbers):
 - unsigned short htons(unsigned short hostshort);
 - unsigned short ntohs(unsigned short netshort);

UDP Clients and Servers

- Connectionless clients and servers create a socket using `SOCK_DGRAM` instead of `SOCK_STREAM`
- Connectionless servers do not call `listen()` or `accept()`, and *usually* do not call `connect()`
- Since connectionless communications lack a sustained connection, several methods are available that allow you to *specify a destination address with every call*:
 - `sendto(sock, buffer, buflen, flags, to_addr, tolen);`
 - `recvfrom(sock, buffer, buflen, flags, from_addr, fromlen);`
- *Examples: daytimeclient.c*
- Example: `daytimeclient time.mit.edu`

Servicing Multiple Clients

- Two main approaches:
 - forking with `fork()`
 - selecting with `select()`
- `fork()` approach forks a new process to handle each incoming client connection, essentially to act as a “miniserver” dedicated to each new client:
 - must worry about zombies created when parent loops back to `accept()` a new client (ignore `SIG_CHILD` signal)
 - inefficient
- A better approach would be to have a *single* process handle all incoming clients, without having to spawn separate child “server handlers”. Enter `select()`.

select()

```
int select(int numfiledescs, fd_set readfdsset, fd_set  
writefdsset, fd_set errorfdsset, struct timeval * timeout);
```

- The select() system call provides a way for a single server to wait until a set of network connections has data available for reading
- The advantage over fork() here is that no multiple processes are spawned
- The downside is that the single server must handle state management on its own for all its new clients

select() (continued)

- select() will return if *any* of the descriptors in readfdsset and writefdsset of file descriptors are ready for reading or writing, respectively, or, if any of the descriptors in errorfdsset are in an error condition
- The FD_SET(int fd, fd_set *set) function will add the file descriptor *fd* to the set *set*
- The FD_ISSET(int fd, fd_set *set) function will tell you if filedesc *fd* is in the modified set *set*
- select() returns the total number of descriptors in the modified sets
- If a client closes a socket whose file descriptor is in one of your watched sets, select() will return, and your next recv() will return 0, indicating the socket has been closed

Setting the timeout in select()

- If you set the timeout to 0, select() times out *immediately*
- If you set the timeout to NULL, select() will *never* time out, and will block indefinitely until a file descriptor is modified
- If you don't care about a particular file descriptor set, just set it to NULL in the call:
 - `select (max, &readfds, NULL, NULL, NULL);`
 - Here we only care about reading, and we want to block indefinitely until we do have a file descriptor ready to be read
- *examples: multiserver.c, multclient.c*

How to Reuse Addresses

- Local ports are locked from rebinding for a period of time (usually a couple of minutes based on the `TIME_WAIT` state) after a process closes them. This is to ensure that a temporarily “lost” packet does not reappear, and then be delivered to a *reincarnation* of a listening server. But when coding and debugging a client server app, this is bothersome. The following code will turn this feature off:

```
int yes = 1;
server = socket(AF_INET, SOCK_STREAM, 0);

if (setsockopt(server, SOL_SOCKET, SO_REUSEADDR,
&yes, sizeof(int)) < 0)
{
    perror("setsockopt SO_REUSEADDR");
    exit(1);
}
```

Miscellaneous Socket Functions

- `int getpeername(int sockfd, struct sockaddr * addr, int *addrlen);`
 - this tells you the *hostname* of the REMOTE connection embedded in `addr`
- `int gethostname(char * hostname, size_t size);`
 - this tells you the *hostname* of your LOCAL connection
- `int inet_aton(const char * string_address, &(addr.sin_addr));`
 - converts the const ip *string_address* (“192.168.3.1”) into an acceptable numeric form
- `addr.sin_addr = inet_addr(“192.168.3.1”);`
 - does the same thing

More Miscellaneous Functions

- `struct hostent *gethostbyname(const char *hostname);`
 - Does a DNS lookup and returns a pointer to a `hostent` structure that contains the host name, aliases, address type (`AF_INET`, etc.), length, and an array of IP addresses for this host (`hostent.h_addr_list[0]` is usually the one)

(cf. `/etc/nsswitch.conf`)

```
struct hostent {
    char *h_name; /*DNS host name*/
    char **h_aliases; /*alias list*/
    int h_addrtype; /* "AF_INET", etc*/
    int h_length; /* length of addr*/
    char **h_addr_list; /*list of IP
adds*/
};
```

And a Few More

```
struct servent * getservbyname(const char
    *servicename, const char *protocol)
```

```
struct servent * getservbyport(int port, const char
    *protocol)
```

- example:

```
serventptr = getservbyname("daytime", "udp");
struct servent {
    char * s_name; /*official service name*/
    char **s_aliases; /* alias list */
    int s_port; /*port num*/
    char *s_proto; /* protocol: "tcp", "udp"*/
};
```

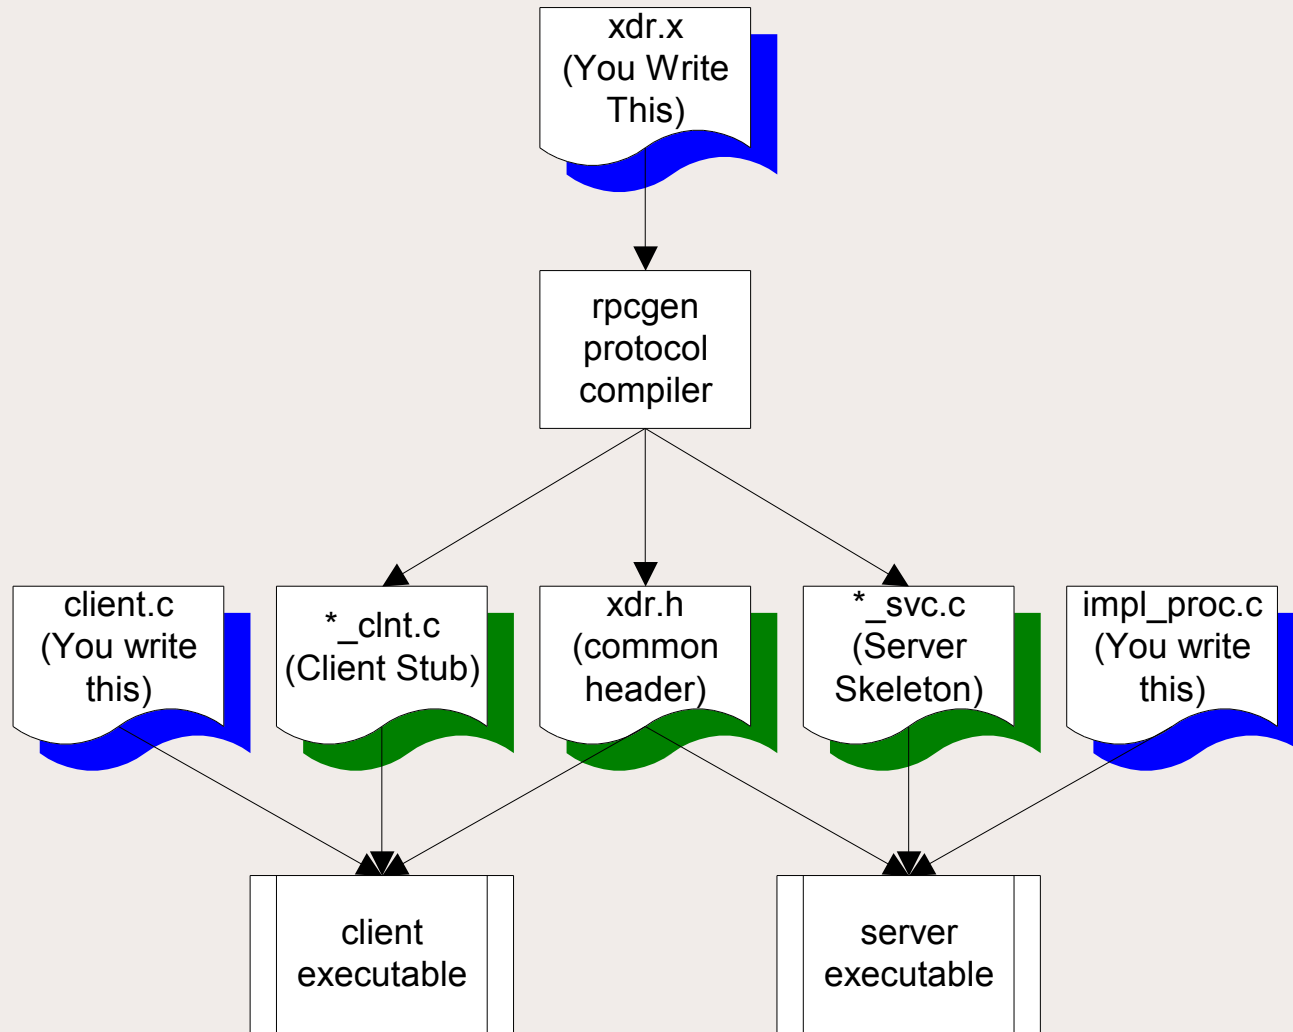
Remote Procedure Calls

Digital ONC RPC

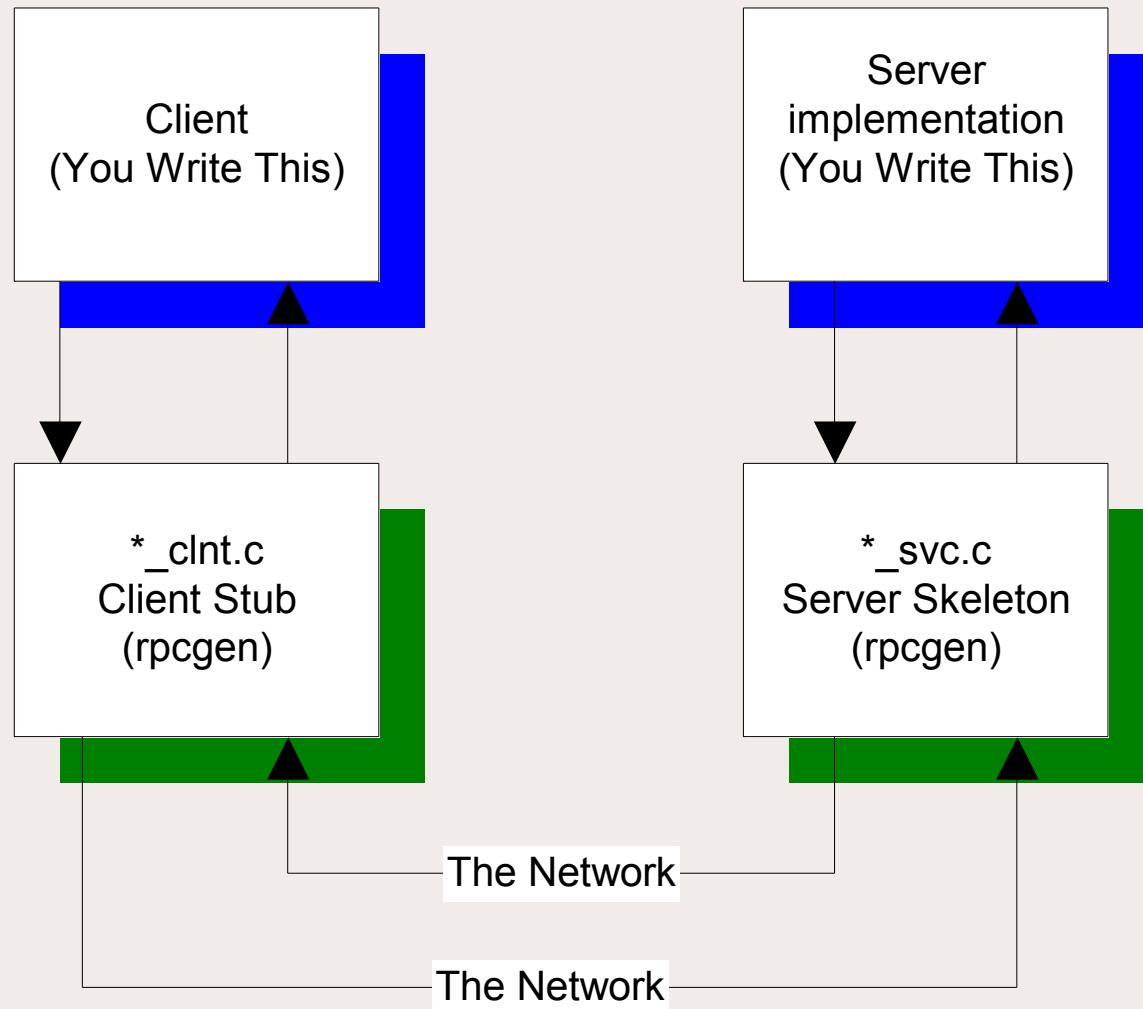
The Point

- “What’s the difference between local and remote procedure calling?”
 - “Very little—that’s the point”
- Remote Procedures generally accept and return *pointers* to data

The Process



Call Sequence



Remote Services

- SUN Remote Procedure Call
 - If the time to transfer the data is more than the time to execute a remote command, the latter is generally preferable.
 - UDP protocol is used to initiate a remote procedure, and the results of the computation are returned.

SUN RPC

- Communication is message-based
- When a server starts, it binds an arbitrary port and publishes that port and the PROGRAM and VERSION with the portmapper daemon (port 111)
- When a client starts, it contacts the portmapper and asks where it can find the remote procedure, using PROGRAM and VERSION ids. The portmapper daemon returns the address and client and server communicate directly.

Sample protocol definition file (.x file)

this XDR file (somefile.x):

```
program NUMPROG
{
    version NUMVERS
    {
        int READNUM(int) = 1; /* procedure number 1 */
    } = 1; /* interface VERSION number */
} = 0x2000002; /* PROGRAM number */
```

is turned into this header file by rpcgen (somefile.h):

```
#define NUMPROG 0x2000002
#define NUMVERS 1

#if defined(__STDC__) || defined(__cplusplus)
#define READNUM 1
extern int * readnum_1(int *, CLIENT *);
extern int * readnum_1_svc(int *, struct svc_req *);
```

RPC Paradigms for Client Server

- Fat Client-DBMS (2 Tier)
 - VB \Leftrightarrow Sybase (ODBC)
 - Motif C++ \Leftrightarrow DBMS (ctlib)
- Fat Client-Application Server-DBMS
 - C Front End \Leftrightarrow C Business Logic \Leftrightarrow DBMS

RPC Under the Hood

- RPC is important because it handles network details for you:
 - Network Details
 - Byte Ordering (Big Endian, Little Endian)
 - Alignment Details
 - 2/4 Byte alignment
 - String Termination (NULL ?)
 - Pointers (how to handle migration of pointers?)

RPC eXternal Data Representation

- XDR provides:
 - Network Transparency
 - Single Canonical Form using Big-Endian
 - 4-Byte alignment
 - XDR passes all data across the wire in a byte stream
 - Filters

XDR Filters

- Integer: int (4 bytes)
- Unsigned Integer: unsigned int (4 bytes)
- char: int (4 byte signed integer)
- Double: double (8 bytes IEEE754 FP)
- Float: float (4 bytes IEEE754 FP)
- int week[7]
- int orders <50> (variable length array)
- opaque data<1000> any data

Building an RPC Application

- Create XDR file (`~mark/pub/518/rpc/[linux|sun]/numdisp.x`)
- run `rpcgen` to create
 - client stub: `numdisp_clnt.c`
 - server skeleton: `numdisp_svc.c`
 - common header: `numdisp.h`
- write `client.c` and `numdisp_proc.c`
- compile client and server (in subdirs)
- run (client on devon, server on orcus)
- *example: ~mark/pub/51081/rpc/linux, ~mark/pub/51081/rpc/sun*