

Chapter 2

Application Layer


A note on the use of these ppt slides:

We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a *lot* of work on our part. In return for use, we only ask the following:

- ❖ If you use these slides (e.g., in a class) in substantially unaltered form, that you mention their source (after all, we'd like people to use our book!)
- ❖ If you post any slides in substantially unaltered form on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

All material copyright 1996-2010
J.F Kurose and K.W. Ross, All Rights Reserved



*Computer Networking:
A Top Down Approach,
5th edition.
Jim Kurose, Keith Ross
Addison-Wesley, April
2009.*

The Bare minimum

- ❖ To code a socket, you will need at least
 - ACCEPT: *block and wait* for CONNECT PKT
 - CONNECT: *establish* a connection
 - RECEIVE: *block and wait* for a SEND PKT
 - SEND: *actually sending* a PKT on the channel
 - DISCONNECT: *putting an end*

- ❖ These are the functions you'll see
 - C, JAVA, for any connection-oriented transport

A first example


❖ How does it work

- Server LISTEN, wait for CONNECT PKT
- Client send a CONNECT message, and then block until received the answer from server
- Once server received CONNECT message, it becomes unblocked, send an answer, and becomes blocked again in READ
- Once the client received the answer, it becomes unblocked, SENDS a request message, and block again in READ
- The server finally answer with data, and close

Socket-programming using TCP

Socket: a door between application process and end-end-transport protocol (UCP or TCP)

TCP service: reliable transfer of *bytes* from one process to another



Socket programming *with TCP*

Client must contact server

- ❖ server process must first be running
- ❖ server must have created socket (door) that welcomes client's contact

Client contacts server by:

- ❖ creating client-local TCP socket
- ❖ specifying IP address, port number of server process
- ❖ when **client creates socket**: client TCP establishes connection to server TCP

- ❖ when contacted by client, **server TCP creates new socket** for server process to communicate with client
 - allows server to talk with multiple clients
 - source port numbers used to distinguish clients (more in Chap 3)

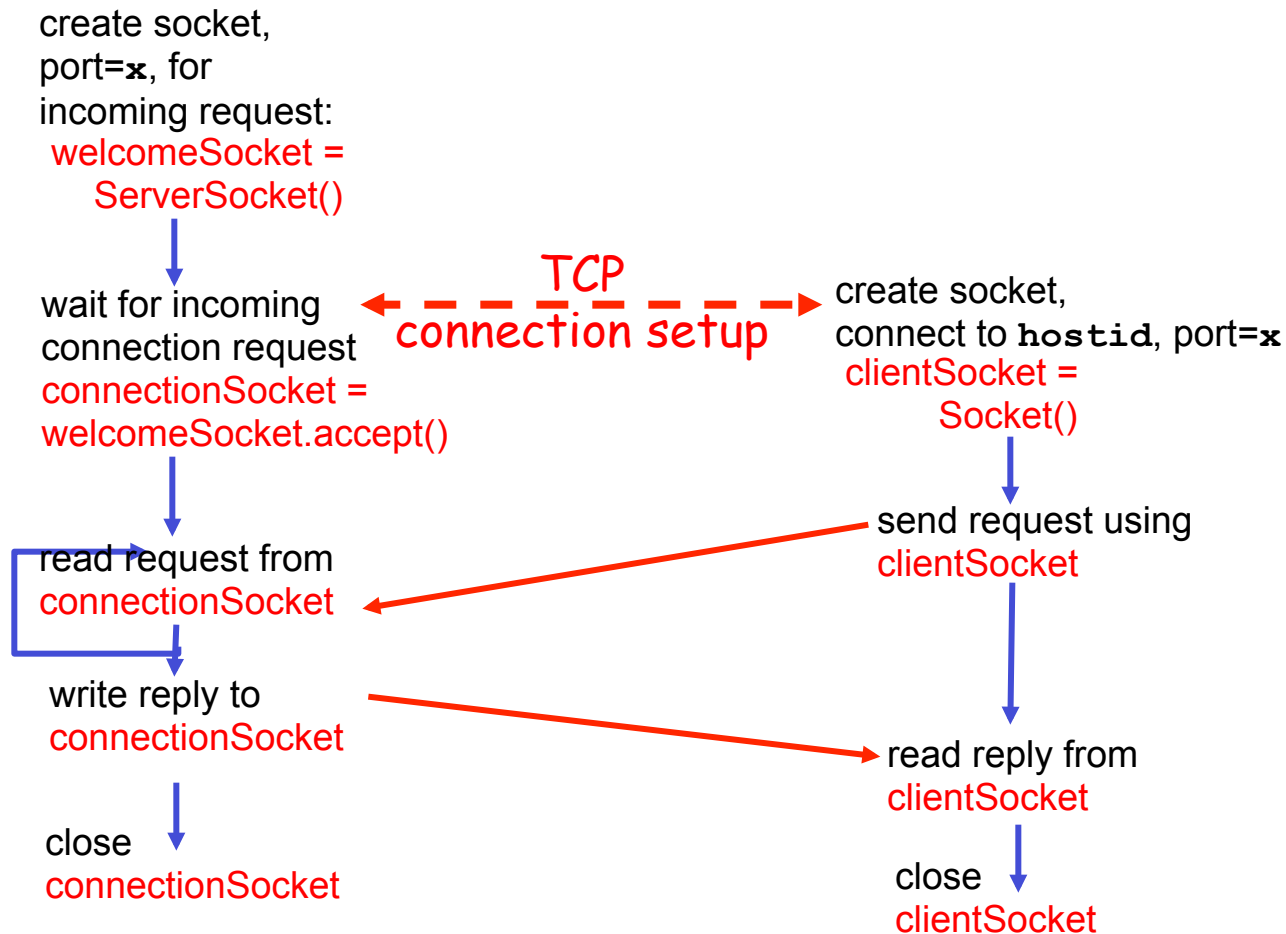
application viewpoint

TCP provides reliable, in-order transfer of bytes ("pipe") between client and server

Client/server socket interaction: TCP

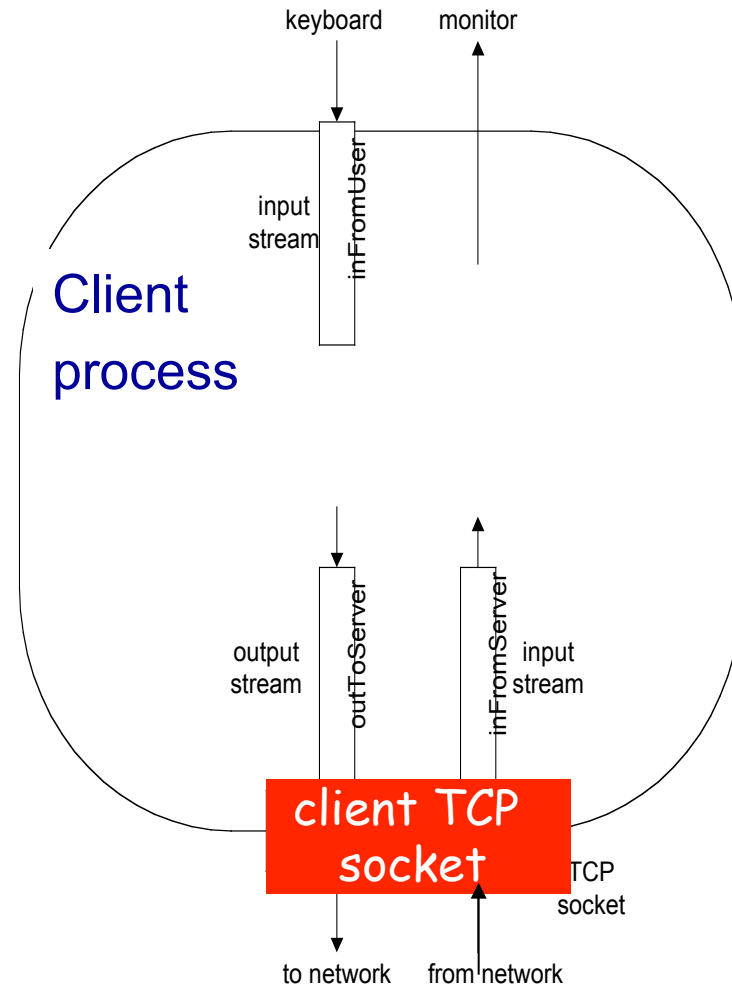
Server (running on `hostid`)

Client



Stream jargon

- ❖ **stream** is a sequence of characters that flow into or out of a process.
- ❖ **input stream** is attached to some input source for the process, e.g., keyboard or socket.
- ❖ **output stream** is attached to an output source, e.g., monitor or socket.



Socket programming with TCP

Example client-server app:

- 1) client reads line from standard input (`inFromUser` stream) , sends to server via socket (`outToServer` stream)
- 2) server reads line from socket
- 3) server converts line to uppercase, sends back to client
- 4) client reads, prints modified line from socket (`inFromServer` stream)

Example: Java client (TCP)

```
import java.io.*;
import java.net.*;
class TCPClient {
```

← This package defines Socket() and ServerSocket() classes

```
    public static void main(String argv[]) throws Exception
    {
```

```
        String sentence;
        String modifiedSentence;
```

create
input stream →

```
        BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));
```

create
clientSocket object
of type Socket,
connect to server →

```
        Socket clientSocket = new Socket("hostname", 6789);
```

server name,
e.g., www.umass.edu
server port #

create
output stream
attached to socket →

```
        DataOutputStream outToServer =
            new DataOutputStream(clientSocket.getOutputStream());
```

Example: Java client (TCP), cont.

```
        create  
        input stream → BufferedReader inFromServer =  
        attached to socket → new BufferedReader(new  
                               InputStreamReader(clientSocket.getInputStream()));  
  
                               sentence = inFromUser.readLine();  
  
        send line  
        to server → outToServer.writeBytes(sentence + '\n');  
  
        read line  
        from server → modifiedSentence = inFromServer.readLine();  
  
                               System.out.println("FROM SERVER: " + modifiedSentence);  
  
        close socket → clientSocket.close();  
        (clean up behind yourself!)  
  
        }  
    }
```

Example: Java server (TCP)

```
import java.io.*;
import java.net.*;
```

```
class TCPServer {
```

```
    public static void main(String argv[]) throws Exception
    {
```

```
        String clientSentence;
        String capitalizedSentence;
```

create
welcoming socket
at port 6789

```
        ServerSocket welcomeSocket = new ServerSocket(6789);
```

wait, on welcoming
socket accept() method
for client contact create,
new socket on return

```
        while(true) {
```

```
            Socket connectionSocket = welcomeSocket.accept();
```

create input
stream, attached
to socket

```
                BufferedReader inFromClient =
```

```
                new BufferedReader(new
                    InputStreamReader(connectionSocket.getInputStream()));
```

Example: Java server (TCP), cont

create output
stream, attached
to socket

→ `DataOutputStream outToClient =
new DataOutputStream(connectionSocket.getOutputStream());`

read in line
from socket

→ `clientSentence = inFromClient.readLine();`

`capitalizedSentence = clientSentence.toUpperCase() + '\n';`

write out line
to socket

→ `outToClient.writeBytes(capitalizedSentence);`

}
}
}

end of while loop,
loop back and wait for
another client connection

Chapter 2: Application layer

2.1 Principles of network applications

2.2 Web and HTTP

2.3 FTP

2.4 Electronic Mail

- SMTP, POP3, IMAP

2.5 DNS

2.6 P2P applications

2.7 Socket programming with TCP

2.8 Socket programming with UDP

Socket programming *with UDP*

UDP: no “connection” between client and server

- ❖ no handshaking
- ❖ sender explicitly attaches IP address and port of destination to each packet
- ❖ server must extract IP address, port of sender from received packet

UDP: transmitted data may be received out of order, or lost

application viewpoint:

UDP provides unreliable transfer of groups of bytes (“datagrams”) between client and server

Client/server socket interaction: UDP

Server (running on `hostid`)

Client

create socket,
port= x.
`serverSocket =
DatagramSocket()`

read datagram from
`serverSocket`

write reply to
`serverSocket`
specifying
client address,
port number

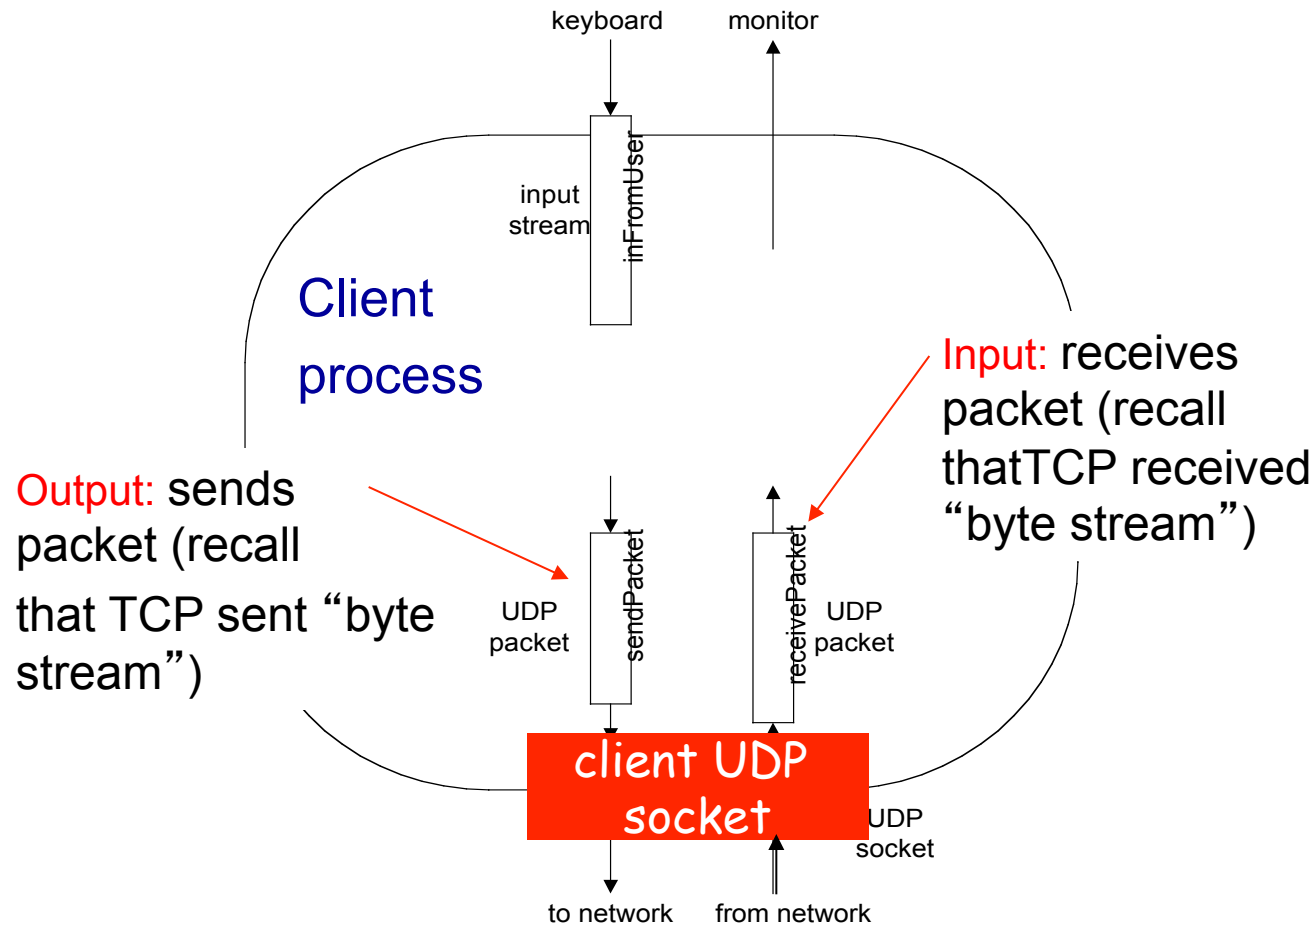
create socket,
`clientSocket =
DatagramSocket()`

Create datagram with server IP and
port=x; send datagram via
`clientSocket`

read datagram from
`clientSocket`

close
`clientSocket`

Example: Java client (UDP)



Example: Java client (UDP)

```
import java.io.*;
import java.net.*;
```

```
class UDPClient {
    public static void main(String args[]) throws Exception
    {
```

create
input stream

```
        BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));
```

create
client socket

```
        DatagramSocket clientSocket = new DatagramSocket();
```

translate
hostname to IP
address using DNS

```
        InetAddress IPAddress = InetAddress.getByName("hostname");
```

```
        byte[] sendData = new byte[1024];
        byte[] receiveData = new byte[1024];
```

```
        String sentence = inFromUser.readLine();
```

```
        sendData = sentence.getBytes();
```

Example: Java client (UDP), cont.

```
create datagram  
with data-to-send,  
length, IP addr, port } DatagramPacket sendPacket =  
                        } new DatagramPacket(sendData, sendData.length, IPAddress, 9876);  
  
send datagram  
to server } clientSocket.send(sendPacket);  
  
read datagram  
from server } DatagramPacket receivePacket =  
              } new DatagramPacket(receiveData, receiveData.length);  
              } clientSocket.receive(receivePacket);  
  
String modifiedSentence =  
    new String(receivePacket.getData());  
  
System.out.println("FROM SERVER:" + modifiedSentence);  
clientSocket.close();  
}  
}
```

Example: Java server (UDP)

```
import java.io.*;  
import java.net.*;
```

```
class UDPServer {  
    public static void main(String args[]) throws Exception  
    {
```

create
datagram socket
at port 9876

```
        DatagramSocket serverSocket = new DatagramSocket(9876);
```

```
        byte[] receiveData = new byte[1024];  
        byte[] sendData = new byte[1024];
```

```
        while(true)  
        {
```

create space for
received datagram

```
            DatagramPacket receivePacket =  
                new DatagramPacket(receiveData, receiveData.length);
```

receive
datagram

```
            serverSocket.receive(receivePacket);
```

Example: Java server (UDP), cont

```
String sentence = new String(receivePacket.getData());
```

get IP addr
port #, of
sender

```
InetAddress IPAddress = receivePacket.getAddress();
```

```
int port = receivePacket.getPort();
```

```
String capitalizedSentence = sentence.toUpperCase();
```

```
sendData = capitalizedSentence.getBytes();
```

create datagram
to send to client

```
DatagramPacket sendPacket =  
    new DatagramPacket(sendData, sendData.length, IPAddress,  
        port);
```

write out
datagram
to socket

```
serverSocket.send(sendPacket);
```

```
}  
}  
}
```

end of while loop,
loop back and wait for
another datagram

Chapter 2: Application layer

2.1 Principles of network applications

2.2 Web and HTTP

2.3 FTP

2.4 Electronic Mail

- SMTP, POP3, IMAP

2.5 DNS

2.6 P2P applications

2.7 Socket programming with TCP

2.8 Socket programming with UDP

+ (Bonus) Same with C

The Bare minimum

- ❖ To code a socket, you will need at least
 - ACCEPT: *block and wait* for CONNECT PKT
 - CONNECT: *establish* a connection
 - RECEIVE: *block and wait* for a SEND PKT
 - SEND: *actually sending* a PKT on the channel
 - DISCONNECT: *putting an end*

- ❖ These are the functions you'll see
 - C, JAVA, etc.

Socket functions overview (C)

- ❖ For TCP with C, the primitives are:
 - **SOCKET**
 - **BIND**
 - **LISTEN:**
 - **ACCEPT:** *block and wait* for CONNECT PKT
 - **CONNECT:** *establish* a connection
 - **RECEIVE:** *block and wait* for a SEND PKT
 - **SEND:** *actually sending* a PKT on the channel
 - **DISCONNECT:** *putting an end*

Socket Creation in C: socket

- ❖ `int s = socket(domain, type, protocol);`
 - `s`: socket descriptor, an integer
 - `domain`: integer, communication domain
 - e.g., `PF_INET` (IPv4 protocol) - typically used
 - `type`: communication type
 - `SOCK_STREAM`: reliable, 2-way, connection-based service
 - `SOCK_DGRAM`: unreliable, connectionless,
 - other values: need root permission, rarely used, or obsolete
 - `protocol`: specifies protocol - usually set to 0
- ❖ NOTE: socket call does not specify where data will be coming from, nor where it will be going to - it just creates the interface!

The bind function

- ❖ associates and (can exclusively) reserves a port for use by the socket
- ❖ `int status = bind(sockid, &addrport, size);`
 - `status`: error status, = -1 if bind failed
 - `sockid`: integer, socket descriptor
 - `addrport`: struct `sockaddr`, the (IP) address and port of the machine (address usually set to `INADDR_ANY` - chooses a local address)
 - `size`: the size (in bytes) of the `addrport` structure
- ❖ bind can be skipped for both types of sockets. When and why?

Skipping the bind

❖ SOCK_DGRAM:

- if only sending, no need to bind. The OS finds a port each time the socket sends a pkt
- if receiving, need to bind

❖ SOCK_STREAM:

- At the client - determined during conn. setup
- don't need to know port sending from (during connection setup, receiving end is informed of port)

Connection Setup (SOCK_STREAM)

- ❖ Recall: no connection setup for SOCK_DGRAM
- ❖ A connection occurs between two kinds of participants
 - passive: waits for an active participant to request connection
 - active: initiates connection request to passive side
- ❖ Once connection is established, passive and active participants are “similar”
 - both can send & receive data
 - either can terminate the connection

Connection setup cont' d

❖ Passive participant

- step 1: **listen** (for incoming requests)
- step 3: **accept** (a request)
- step 4: data transfer

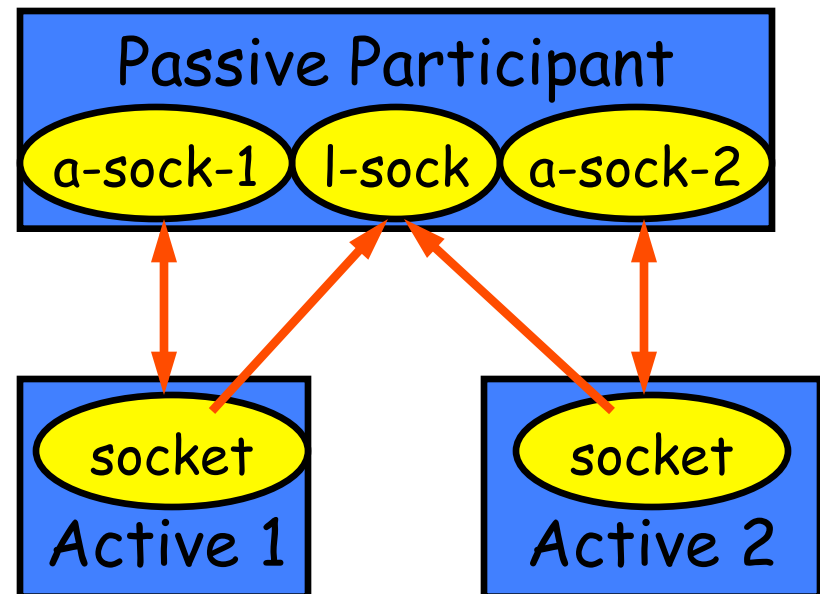
❖ The accepted connection is on a new socket

❖ The old socket continues to listen for other active participants

❖ Why?

❖ Active participant

- step 2: request & establish **connection**
- step 4: data transfer



Connection setup: listen & accept

- ❖ Called by passive participant
- ❖ `int status = listen(sock, queuelen);`
 - `status`: 0 if listening, -1 if error
 - `sock`: integer, socket descriptor
 - `queuelen`: integer, # of active participants that can “wait” for a connection
 - `listen` is **non-blocking**: returns immediately
- ❖ `int s = accept(sock, &name, &namelen);`
 - `s`: integer, the new socket (used for data-transfer)
 - `sock`: integer, the orig. socket (being listened on)
 - `name`: struct `sockaddr`, address of the active participant
 - `namelen`: `sizeof(name)`: value/result parameter
 - must be set appropriately before call
 - adjusted by OS upon return
 - `accept` is **blocking**: waits for connection before returning

connect call

- ❖ `int status = connect(sock, &name, namelen);`
 - `status`: 0 if successful connect, -1 otherwise
 - `sock`: integer, socket to be used in connection
 - `name`: struct `sockaddr`: address of passive participant
 - `namelen`: integer, `sizeof(name)`
- ❖ connect is **blocking**

Sending / Receiving Data

- ❖ With a connection (SOCK_STREAM):
 - `int count = send(sock, &buf, len, flags);`
 - `count`: # bytes transmitted (-1 if error)
 - `buf`: `char[]`, buffer to be transmitted
 - `len`: integer, length of buffer (in bytes) to transmit
 - `flags`: integer, special options, usually just 0
 - `int count = recv(sock, &buf, len, flags);`
 - `count`: # bytes received (-1 if error)
 - `buf`: `void[]`, stores received bytes
 - `len`: # bytes received
 - `flags`: integer, special options, usually just 0
 - Calls are **blocking** [returns only after data is sent (to socket buf) / received]

Sending / Receiving Data (cont' d)

- ❖ Without a connection (SOCK_DGRAM):
 - `int count = sendto(sock, &buf, len, flags, &addr, addrlen);`
 - `count, sock, buf, len, flags`: same as `send`
 - `addr`: `struct sockaddr`, address of the destination
 - `addrlen`: `sizeof(addr)`
 - `int count = recvfrom(sock, &buf, len, flags, &addr, &addrlen);`
 - `count, sock, buf, len, flags`: same as `recv`
 - `addr`: `struct sockaddr`, address of the source
 - `addrlen`: `sizeof(addr)`: value/result parameter
- ❖ Calls are **blocking** [returns only after data is sent (to socket `buf`) / received]

close

- ❖ When finished using a socket, the socket should be closed:
- ❖ `status = close(s);`
 - status: 0 if successful, -1 if error
 - s: the file descriptor (socket being closed)
- ❖ Closing a socket
 - closes a connection (for SOCK_STREAM)
 - frees up the port used by the socket

The struct sockaddr

❖ The generic:

```
struct sockaddr {  
    u_short sa_family;  
    char sa_data[14];  
};
```

- `sa_family`
 - specifies which address family is being used
 - determines how the remaining 14 bytes are used

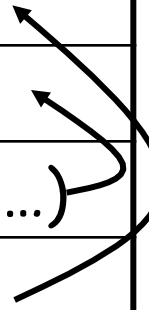
❖ The Internet-specific:

```
struct sockaddr_in {  
    short sin_family;  
    u_short sin_port;  
    struct in_addr sin_addr;  
    char sin_zero[8];  
};
```

- `sin_family` = `AF_INET`
- `sin_port`: port # (0-65535)
- `sin_addr`: IP-address
- `sin_zero`: unused

TCP - Serial Model

Client Side	Server Side
<code>sd=socket(type)</code>	<code>sd=socket(type)</code>
	<code>bind(sd,port)</code>
	<code>listen(sd,len)</code>
<code>connect(sd,dest)</code>	<code>new_sd=accept(sd)</code>
<code>write(sd,...) / send(sd,...)</code>	<code>read(new_sd,...)/recv(new_sd)</code>
<code>read(sd,...)/recv(sd,...)</code>	<code>write(new_sd,...) / send(new_sd,...)</code>
<code>close(sd)</code>	<code>close(new_sd)</code>



TCP - Parallel Model

Client Side	Server Side
<code>sd=socket(type)</code>	<code>sd=socket(type)</code>
	<code>bind(sd,port)</code>
	<code>listen(sd,len)</code>
<code>connect(sd,dest)</code>	<code>new_sd=accept(sd)</code>
	Create another process (e.g., fork)
	<code>close(sd)</code> <code>close(new_sd)</code>
<code>write(sd,...)</code>	<code>read(new_sd,...)</code>
<code>read(sd,...)</code>	<code>write(new_sd,...)</code>
<code>close(sd)</code>	<code>close(new_sd)</code>
	<code>exit()</code>

UDP - Serial Model

Client Side	Server Side
<code>sd=socket(type)</code>	<code>sd=socket(type)</code>
	<code>bind(sd,port)</code>
<code>connect(sd,dest)</code>	
<code>write(sd,...)</code>	<code>recvfrom(sd,...)</code>
<code>read(sd,...)</code>	<code>sendto(sd,...)</code>
<code>close(sd)</code>	<code>close(sd)</code>

Chapter 2: Application layer

2.1 Principles of network applications

2.2 Web and HTTP

2.3 FTP

2.4 Electronic Mail

- SMTP, POP3, IMAP

2.5 DNS

2.6 P2P applications

2.7 Socket programming with TCP

2.8 Socket programming with UDP

+ (Bonus) Same with C

+ (Bonus) A few more functions

Address and port byte-ordering


❖ Address and port are stored as integers

- u_short sin_port; (16 bit)
- in_addr sin_addr; (32 bit)

```
struct in_addr {  
    u_long s_addr;  
};
```

□ Problem:

- different machines / OS' s use different word orderings
 - little-endian: lower bytes first
 - big-endian: higher bytes first
- these machines may communicate with one another over the network



Solution: Network Byte-Ordering


❖ Defs:

- Host Byte-Ordering: the byte ordering used by a host (big or little)
 - Network Byte-Ordering: the byte ordering used by the network - always big-endian
- ❖ Any words sent through the network should be converted to Network Byte-Order prior to transmission (and back to Host Byte-Order once received)
- ❖ Q: should the socket perform the conversion automatically?
- Q: Given big-endian machines don't need conversion routines and little-endian machines do, how do we avoid writing two versions of code?

UNIX's byte-ordering funcs

- ❖ `u_long htonl(u_long x);`
- ❖ `u_short htons(u_short x);`
- ❖ `u_long ntohl(u_long x);`
- ❖ `u_short ntohs(u_short x);`

- ❑ On big-endian machines, these routines do nothing
- ❑ On little-endian machines, they reverse the byte order



- ❑ Same code would have worked regardless of endianness of the two machines

Dealing with blocking calls

- ❖ Many of the functions we saw block until a certain event
 - accept: until a connection comes in
 - connect: until the connection is established
 - recv, recvfrom: until a packet (of data) is received
 - send, sendto: until data is pushed into socket's buffer
 - Q: why not until received?
- ❖ For simple programs, blocking is convenient
- ❖ What about more complex programs?
 - multiple connections
 - simultaneous sends and receives
 - simultaneously doing non-networking processing

Dealing w/ blocking (cont' d)

❖ Options:

- create multi-process or multi-threaded code
- turn off the blocking feature (e.g., using the `fcntl` file-descriptor control function)
- use the `select` function call.

Other useful functions

- ❖ `bzero(char* c, int n)`: 0's n bytes starting at c
- ❖ `gethostname(char *name, int len)`: gets the name of the current host
- ❖ `gethostbyaddr(char *addr, int len, int type)`: converts IP hostname to structure containing long integer
- ❖ `inet_addr(const char *cp)`: converts dotted-decimal char-string to long integer
- ❖ `inet_ntoa(const struct in_addr in)`: converts long to dotted-decimal notation
- ❖ `read()`, `write()`
- ❖ Warning: check function assumptions about byte-ordering (host or network). Often, they assume parameters / return solutions in network byte-order

Release of ports

- ❖ Sometimes, a “rough” exit from a program (e.g., ctrl-c) does not properly free up a port
- ❖ Eventually (after a few minutes), the port will be freed
- ❖ To reduce the likelihood of this problem, include the following code:

```
#include <signal.h>
```

```
void cleanExit(){exit(0);}
```

- in socket code:

```
signal(SIGTERM, cleanExit);
```

```
signal(SIGINT, cleanExit);
```

Final Thoughts

- ❖ Make sure to #include the header files that define used functions
- ❖ Additional info:
 - Ross and Kurose, Computer Networking A Top-Down Approach
 - Comer, Internetworking with TCP/IP, ch. 21
 - Comer and Stevens, Internetworking with TCP/IP - Vol. 3
 - Beej's Guide to Network Programming - <http://www.beej.us/guide/bgnet/>
 - man-pages

Chapter 2: Summary

our study of network apps now complete!

- ❖ application architectures
 - client-server
 - P2P
 - hybrid
- ❖ application service requirements:
 - reliability, bandwidth, delay
- ❖ Internet transport service model
 - connection-oriented, reliable: TCP
 - unreliable, datagrams: UDP
- ❖ specific protocols:
 - HTTP
 - FTP
 - SMTP, POP, IMAP
 - DNS
 - P2P: BitTorrent, Skype
- ❖ socket programming

Chapter 2: Summary

most importantly: learned about protocols

- ❖ typical request/reply message exchange:
 - client requests info or service
 - server responds with data, status code
- ❖ message formats:
 - headers: fields giving info about data
 - data: info being communicated

Important themes:

- ❖ control vs. data msgs
 - ❖ in-band, out-of-band
- ❖ centralized vs. decentralized
- ❖ stateless vs. stateful
- ❖ reliable vs. unreliable msg transfer
- ❖ “complexity at network edge”