

Chapter 2

Application Layer

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

Pure P2P architecture

- ◆ No always-on server
- ◆ Arbitrary end systems directly communicate
- ◆ Peers are intermittently connected and change IP addresses


Three topics:

- ◆ file distribution
- ◆ searching for information
- ◆ case Study: Skype




File Distribution: Server-Client vs P2P

Question : How much time to distribute file from one server to N peers?




File distribution time: server-client

- server sequentially sends N copies:
 - NF/u_s time
- client i takes F/d_i time to download



File distribution time: server-client

- server sequentially sends N copies:
 - NF/u_s time
- client i takes F/d_i time to download




Time to distribute F
to N clients using
client/server approach

$$= d_{cs} = \max \left\{ NF/u_s, F/\min_i(d_i) \right\}$$

increases linearly in N
(for large N)


File distribution time: P2P

- ♦ server must send one copy: F/u_s time
- ♦ client i takes F/d_i time to download
- ♦ NF bits must be downloaded (aggregate)
 - fastest possible upload rate: $u_s + \sum u_i$



File distribution time: P2P


- ♦ server must send one copy: F/u_s time
- ♦ client i takes F/d_i time to download
- ♦ NF bits must be downloaded (aggregate)
 - fastest possible upload rate: $u_s + \sum u_i$



$$d_{P2P} = \max \left\{ F/u_s, F/\min(d_i), NF/(u_s + \sum u_i) \right\}$$

Server-client vs. P2P: example

Client upload rate = u , $F/u = 1$ hour, $u_s = 10u$, $d_{\min} \geq u_s$

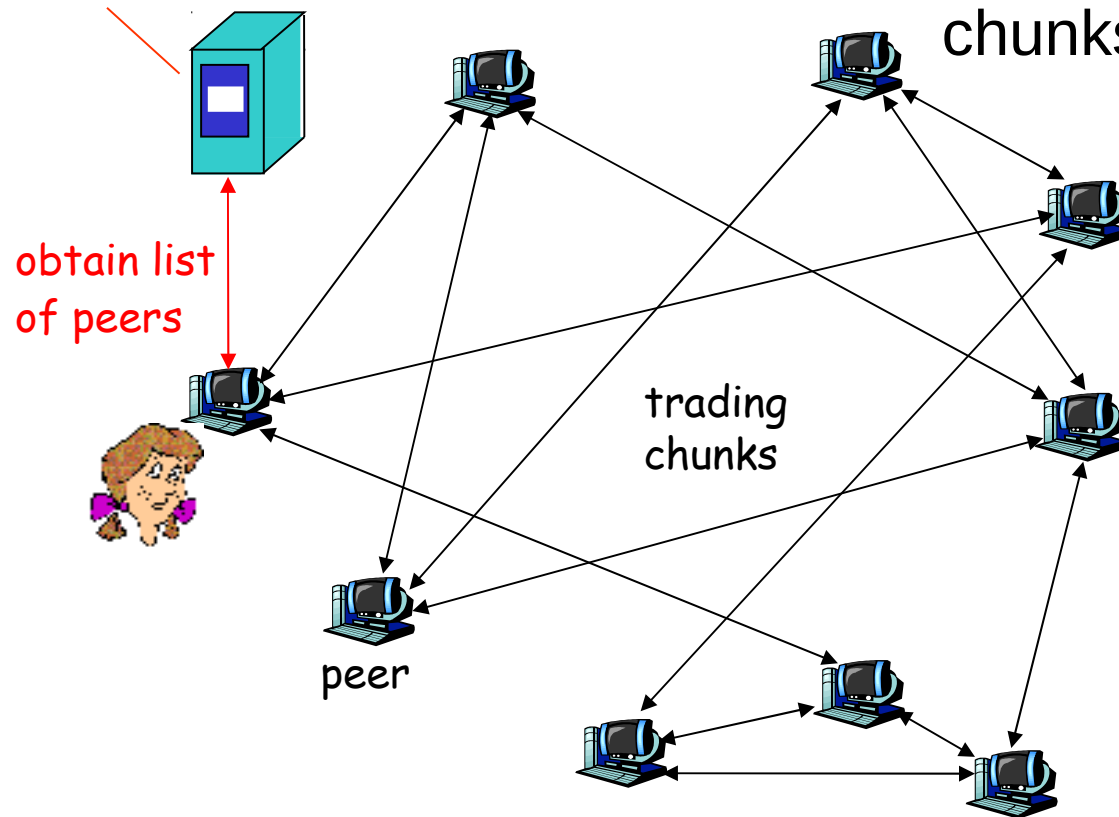


File distribution: BitTorrent


P2P file distribution

tracker: tracks peers participating in torrent

torrent: group of peers exchanging chunks of a file



BitTorrent (1)



- ♦ file divided into 256KB *chunks*.
- ♦ peer joining torrent:
 - ♦ has no chunks, but will accumulate them over time
 - ♦ registers with tracker to get list of peers, connects to subset of peers (“neighbors”)
- ♦ while downloading, peer uploads chunks to other peers.
- ♦ peers may come and go
- ♦ once peer has entire file, it may (selfishly) leave or (altruistically) remain

BitTorrent (2)

Pulling Chunks


- ◆ at any given time, different peers have different subsets of file chunks
- ◆ periodically, a peer (Alice) asks each neighbor for list of chunks that they have.
- ◆ Alice sends requests for her missing chunks
 - ◆ rarest first

Sending Chunks: tit-for-tat

- ❖ Alice sends chunks to four neighbors currently sending her chunks *at the highest rate*
 - re-evaluate top 4 every 10 secs
- ❖ every 30 secs: randomly select another peer, starts sending chunks
 - newly chosen peer may join top 4
 - “optimistically unchoke”


BitTorrent: Tit-for-tat

(1) Alice “optimistically unchokes” Bob




BitTorrent: Tit-for-tat

- (1) Alice “optimistically unchokes” Bob
- (2) Alice becomes one of Bob’s top-four providers;
Bob reciprocates



BitTorrent: Tit-for-tat

- (1) Alice “optimistically unchokes” Bob
- (2) Alice becomes one of Bob’s top-four providers;
Bob reciprocates
- (3) Bob becomes one of Alice’s top-four providers



Distributed Hash Table (DHT)

- ◆ DHT: distributed P2P database
- ◆ database has (key, value) pairs;
 - ◆ key: ss number; value: human name
 - ◆ key: content type; value: IP address
- ◆ peers query DB with key
 - ◆ DB returns values that match the key
- ◆ peers can also insert (key, value) peers


DHT Identifiers

- ◆ Assign integer identifier to each peer in range $[0, 2^n - 1]$.
 - ◆ Each identifier can be represented by n bits.
- ◆ Require each key to be an integer in **same range**.
- ◆ To get integer keys, hash original key.
 - ◆ e.g., $\text{key} = h(\text{"Led Zeppelin IV"})$
 - ◆ this is why they call it a distributed "hash" table

How to assign keys to peers?


- ♦ central issue:
 - ♦ assigning (key, value) pairs to peers.
- ♦ rule: assign key to the peer that has the **closest** ID.
- ♦ convention in lecture: closest is the **immediate successor** of the key.
- ♦ e.g.,: $n=4$; peers: 1,3,4,5,8,10,12,14;
 - ♦ key = 13, then successor peer = 14
 - ♦ key = 15, then successor peer = 1

Circular DHT (1)




- ◆ each peer *only* aware of immediate successor and predecessor.
- ◆ “overlay network”

Circular DHT (2)




Define closest
as closest
successor

Circular DHT (2)



Define closest
as closest
successor


Circular DHT (2)



Define closest
as closest
successor


Circular DHT (2)

$O(N)$ messages
on avg to resolve
query, when there
are N peers




Define closest
as closest
successor

Circular DHT with Shortcuts




- ◆ Each peer keeps track of IP addresses of predecessor, successor, short cuts.

Circular DHT with Shortcuts




- ◆ Each peer keeps track of IP addresses of predecessor, successor, short cuts.
- ◆ Reduced from 6 to 2 messages.
- ◆ Possible to design shortcuts so $O(\log N)$ neighbors, $O(\log N)$ messages in query

Peer Churn



- ❖ To handle peer churn, require each peer to know the IP address of its two successors.
- ❖ Each peer periodically pings its two successors to see if they are still alive.

Peer Churn



- ❖ To handle peer churn, require each peer to know the IP address of its two successors.
- ❖ Each peer periodically pings its two successors to see if they are still alive.

- ❖ peer 5 abruptly leaves
- ❖ Peer 4 detects; makes 8 its immediate successor; asks 8 who its immediate successor is; makes 8's immediate successor its second successor.

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

Goal: learn how to build client/server application that communicate using sockets

Socket API

- ◆ introduced in BSD4.1 UNIX, 1981
- ◆ explicitly created, used, released by apps
- ◆ client/server paradigm
- ◆ two types of transport service via socket API:
 - ◆ unreliable datagram
 - ◆ reliable, byte stream-oriented


socket

a *host-local*, *application-created*, *OS-controlled* interface (a “door”) into which application process can *both send and receive* messages to/from another application process

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

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

Client/server socket interaction: TCP

Server (running on `hostid`)

Client

```
create socket,  
port=x, for  
incoming request:  
welcomeSocket =  
    ServerSocket()
```



Client/server socket interaction: TCP

Server (running on `hostid`)

Client

```
create socket,  
port=x, for  
incoming request:  
welcomeSocket =  
    ServerSocket()  
    ↓  
wait for incoming  
connection request  
connectionSocket =  
welcomeSocket.accept()
```

Client/server socket interaction: TCP

Server (running on `hostid`)

```
create socket,  
port=x, for  
incoming request:  
welcomeSocket =  
    ServerSocket()  
    ↓  
wait for incoming  
connection request  
connectionSocket =  
    welcomeSocket.accept()
```

Client

```
create socket,  
connect to hostid, port=x  
clientSocket =  
    Socket()
```

Client/server socket interaction: TCP

Server (running on `hostid`)

Client

create socket,
port=`x`, for
incoming request:
`welcomeSocket =`
`ServerSocket()`



wait for incoming
connection request
`connectionSocket =`
`welcomeSocket.accept()`

← TCP
connection setup →

create socket,
connect to `hostid`, port=`x`
`clientSocket =`
`Socket()`

Client/server socket interaction: TCP

Server (running on `hostid`)

Client

create socket,
port=`x`, for
incoming request:
`welcomeSocket =`
`ServerSocket()`

wait for incoming
connection request
`connectionSocket =`
`welcomeSocket.accept()`

TCP
connection setup

create socket,
connect to `hostid`, port=`x`
`clientSocket =`
`Socket()`

send request using
`clientSocket`



Client/server socket interaction: TCP

Server (running on `hostid`)

Client

create socket,
port=`x`, for
incoming request:
`welcomeSocket =`
`ServerSocket()`

wait for incoming
connection request
`connectionSocket =`
`welcomeSocket.accept()`


read request from
`connectionSocket`

write reply to
`connectionSocket`

TCP
connection setup

create socket,
connect to `hostid`, port=`x`
`clientSocket =`
`Socket()`

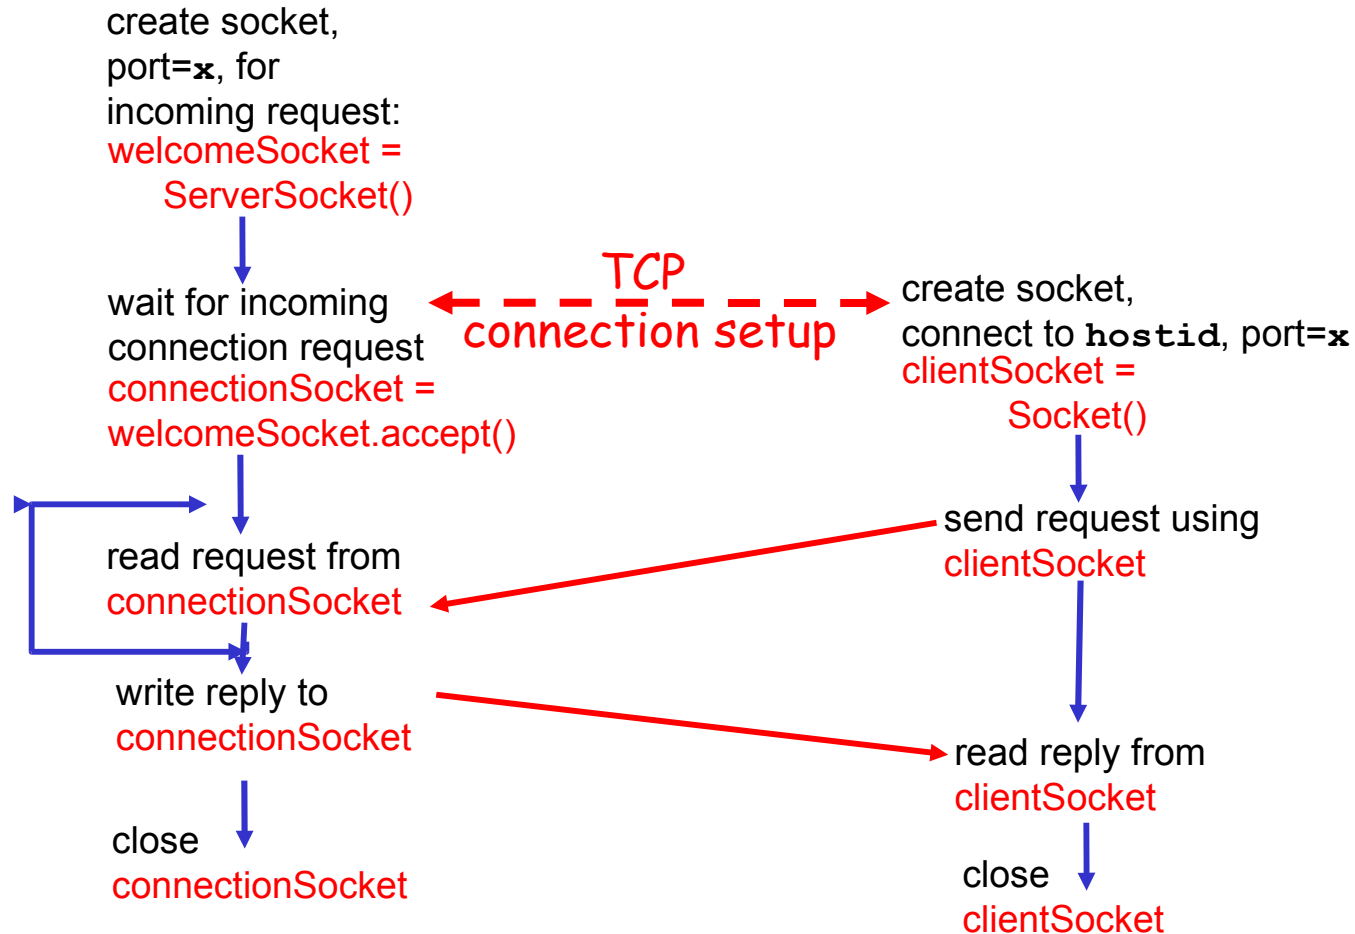
send request using
`clientSocket`



Client/server socket interaction: TCP

Server (running on `hostid`)

Client

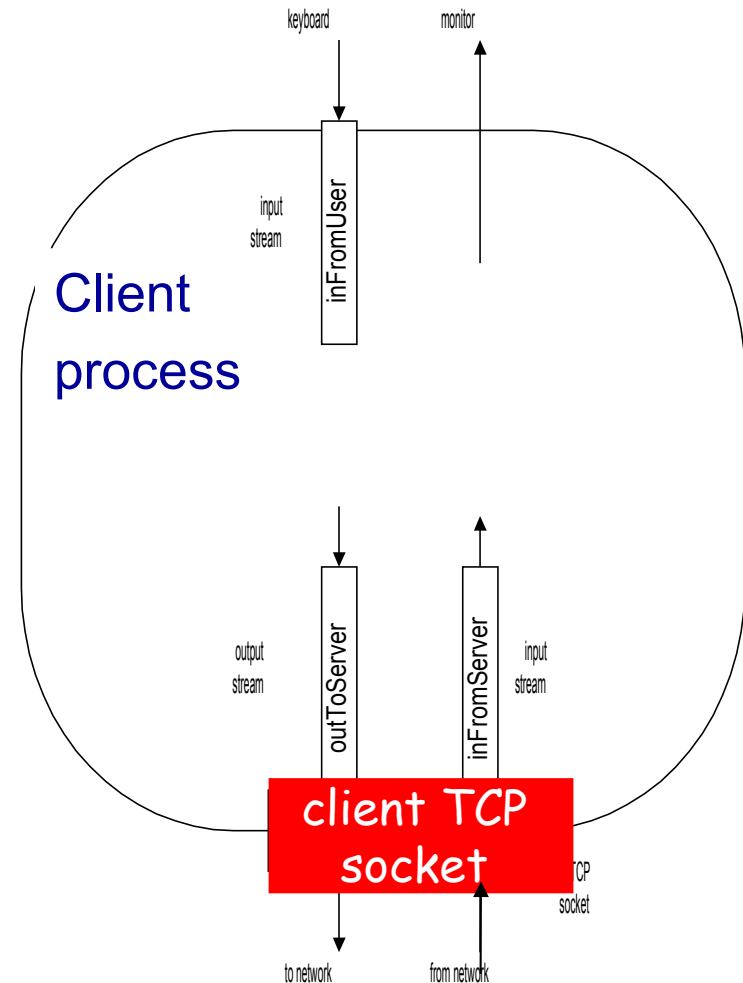


Stream jargon

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.*; ← This package defines Socket()
                    and ServerSocket() classes
class TCPCClient {

    public static void main(String argv[]) throws Exception
    {
        String sentence;
        String modifiedSentence;

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

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

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

Example: Java client (TCP)

```
import java.io.*;
import java.net.*; ← This package defines Socket()
                    and ServerSocket() classes
class TCPCClient {
```

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

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

```
        create input stream → BufferedReader inFromUser =
                               new BufferedReader(new InputStreamReader(System.in));
```

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

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

Example: Java client (TCP)

```
import java.io.*;
import java.net.*; ← This package defines Socket()
                    and ServerSocket() classes
class TCPCClient {
```

```
    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);
```

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

Example: Java client (TCP)

```
import java.io.*;
import java.net.*; ← This package defines Socket()
                    and ServerSocket() classes
class TCPCClient {
```

```
    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 #

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

Example: Java client (TCP)

```
import java.io.*;
import java.net.*; ← This package defines Socket()
                    and ServerSocket() classes
class TCPCClient {
```

```
    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
attached to socket → `BufferedReader inFromServer =
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
(clean up behind yourself!) → `clientSocket.close();`

`}`
`}`

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);

        while(true) {

            Socket connectionSocket = welcomeSocket.accept();

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

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);`

```
        while(true) {
```

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

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

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

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);`

```
        while(true) {
```

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

→ `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`)

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

Client/server socket interaction: UDP

Server (running on `hostid`)

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

Client

create socket,
`clientSocket =`
`DatagramSocket()`



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



Client/server socket interaction: UDP

Server (running on `hostid`)

Client

create socket,
port= x.

`serverSocket =
DatagramSocket()`



read datagram from
`serverSocket`

create socket,

`clientSocket =
DatagramSocket()`



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



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`

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

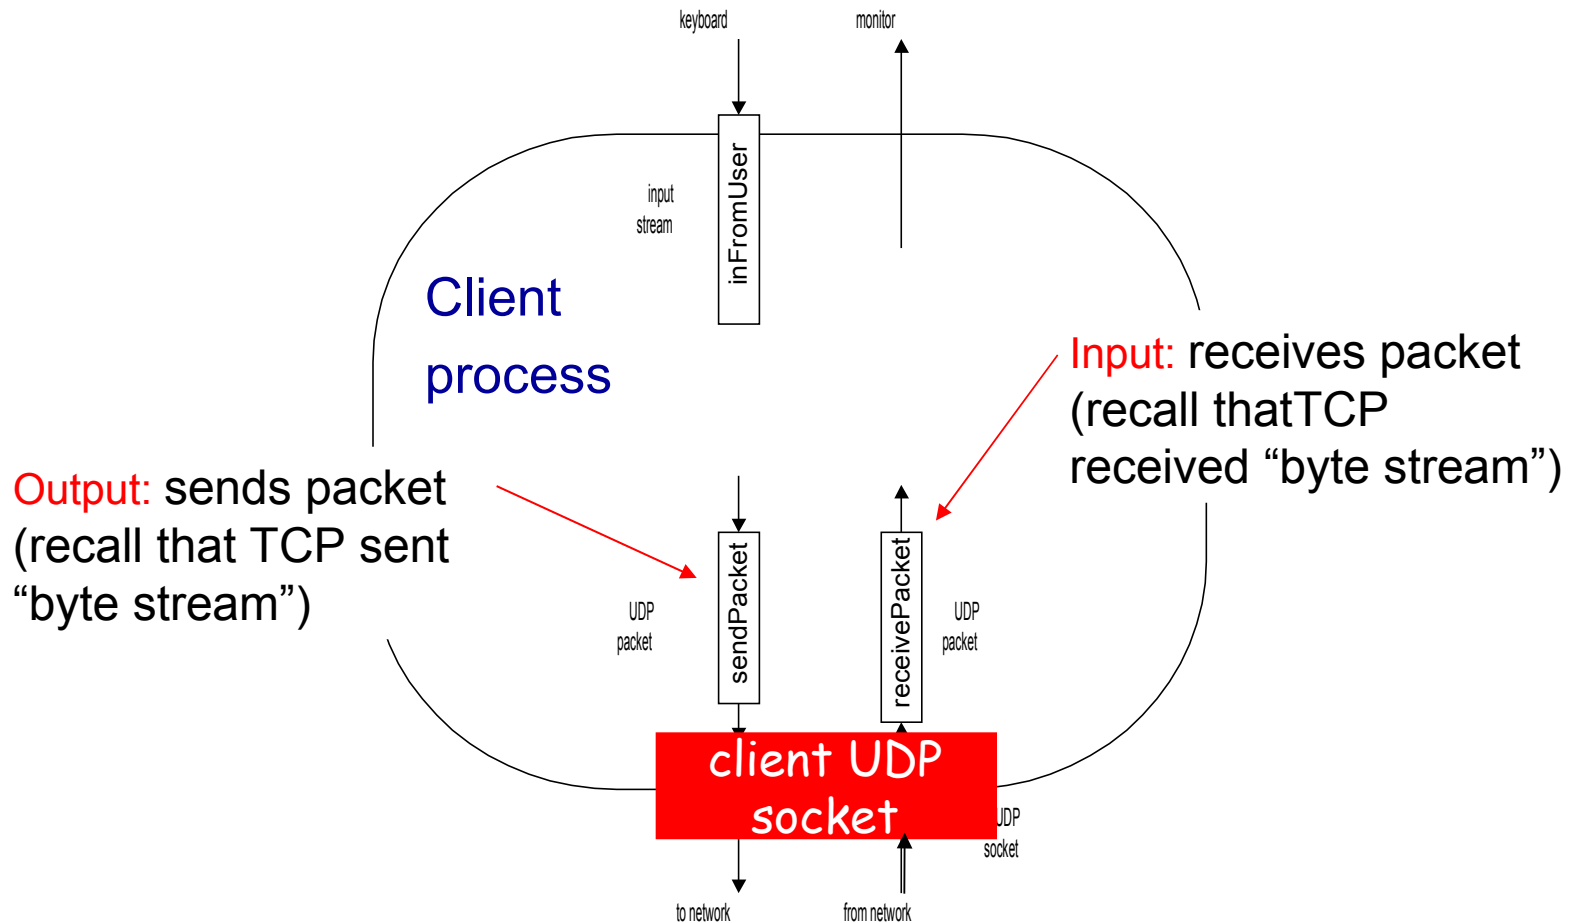
create socket,
`clientSocket =`
`DatagramSocket()`

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

send datagram
to server

read datagram
from server

```
DatagramPacket sendPacket =  
    new DatagramPacket(sendData, sendData.length, IPAddress, 9876);  
  
clientSocket.send(sendPacket);  
  
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);
```