

# IS 450/IS 650– Data Communications and Networks

## Application Layer

Nirmalya Roy

Department of Information Systems

University of Maryland Baltimore County

# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Chapter 2: Application Layer

## Our goals:

- Principles of network application design
  - transport-layer service models
  - client-server paradigm
  - peer-to-peer paradigm

- Popular protocols through case studies
  - ❖ HTTP
  - ❖ FTP
  - ❖ SMTP / POP3 / IMAP
  - ❖ DNS
- Network programming
  - ❖ socket API

# Some network apps

- E-mail
- Web
- Instant messaging
- Remote login
- P2P file sharing
- Multi-user network games
- Streaming stored video clips
- Internet telephone
- Real-time video conference
- Massive parallel computing
- 
- 
- 

Next generation: The network will be the computer. Most Applications will run over the network. Local PC minimally required Example: Google spread sheet

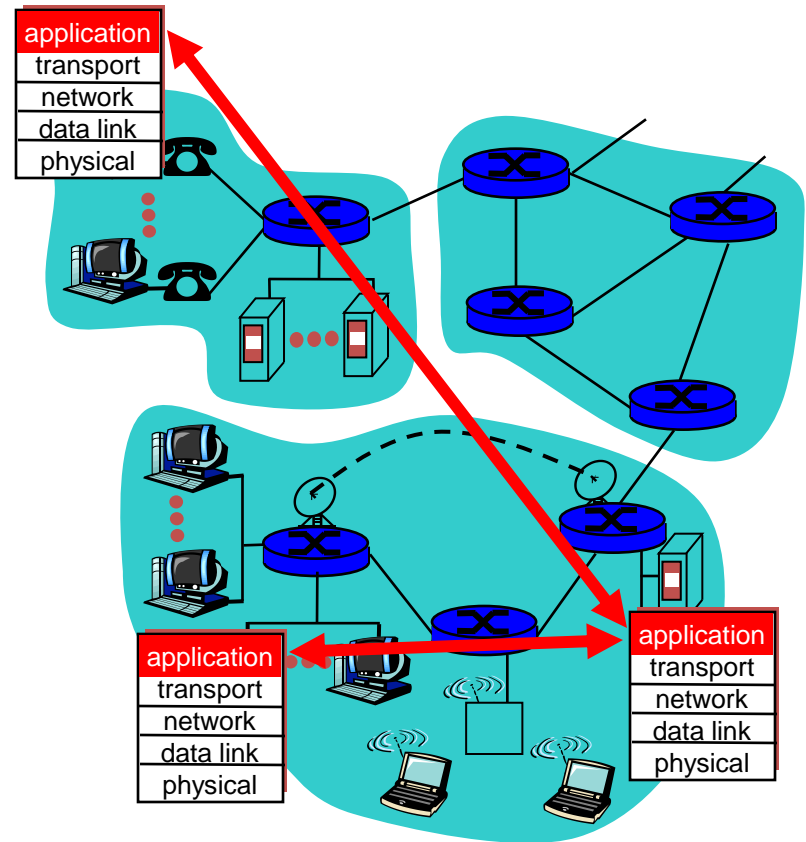
# Creating a network app

## Write programs that

- run on different end systems
- communicate over a network
- e.g., Web: Web server software communicates with browser software

## Little software written for devices in network core

- network core devices do not run user application code
- application on end systems allows for rapid app development, propagation



# Chapter 2: Application layer

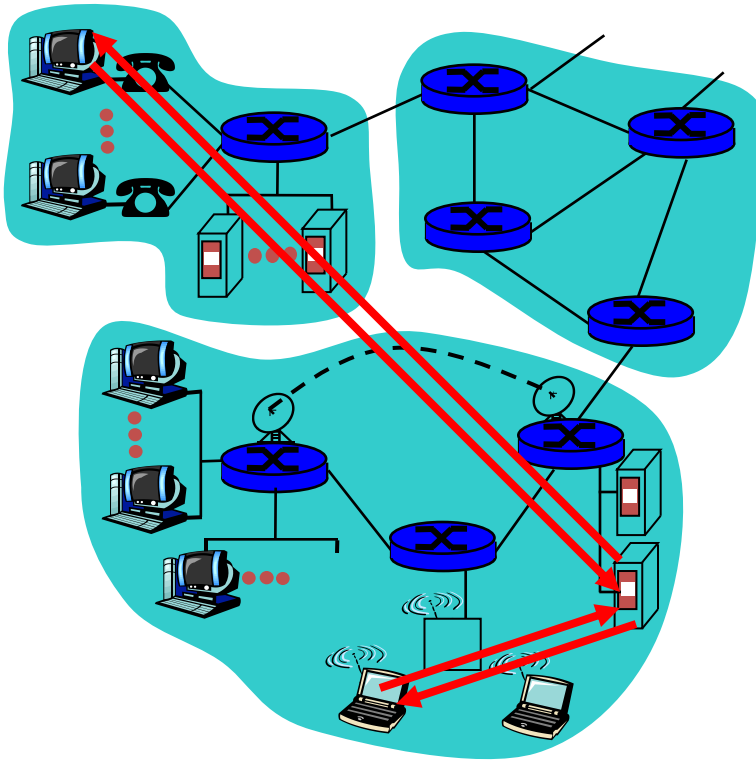
- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Application architectures

---

- Client-server
- Peer-to-peer (P2P)
- Hybrid of client-server and P2P

# Client-server architecture



## server:

- always-on host
- permanent IP address
- server farms for scaling
- data centers

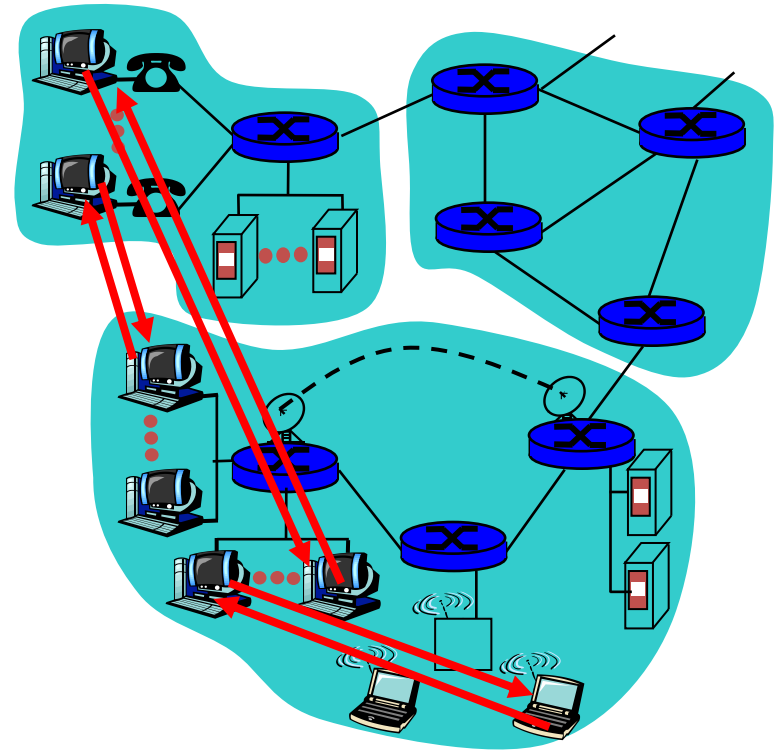
## clients:

- communicate with server
- may be intermittently connected
- may have dynamic IP addresses
- do not communicate directly with each other



# Pure P2P architecture

- no always-on server
- arbitrary end systems directly communicate
- peers are intermittently connected and change IP addresses
- Major Challenges: ISP friendly, Security, Incentives
- example: Gnutella (peer-to-peer file sharing network)



Highly scalable but difficult to manage

# Hybrid of client-server and P2P

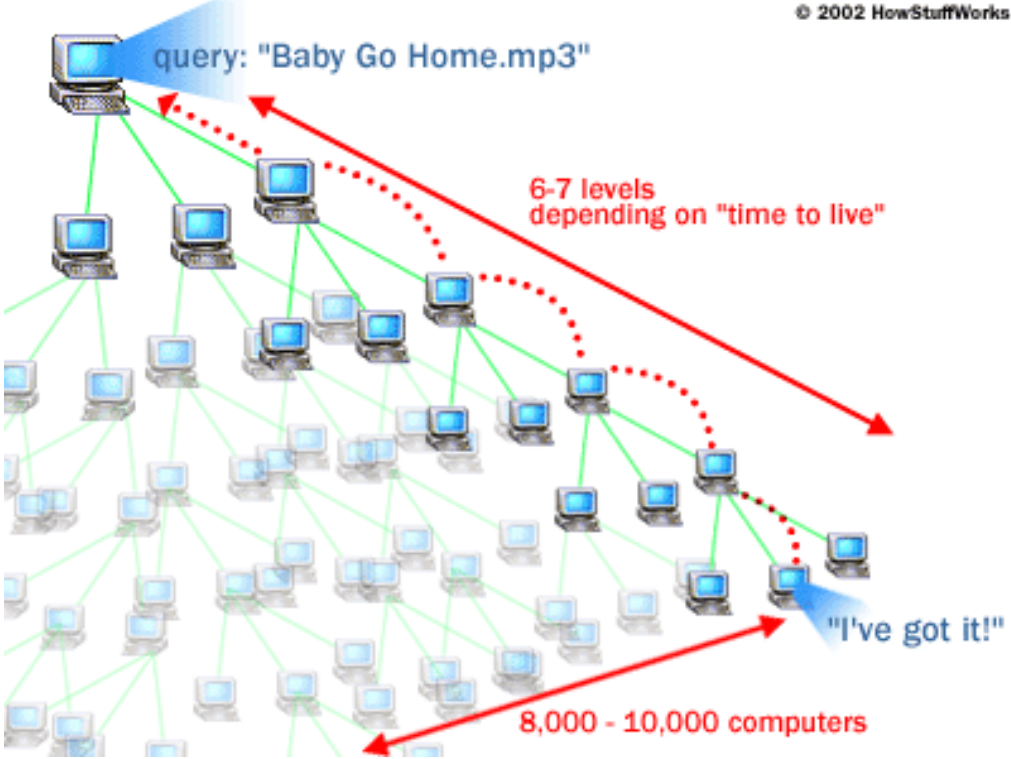
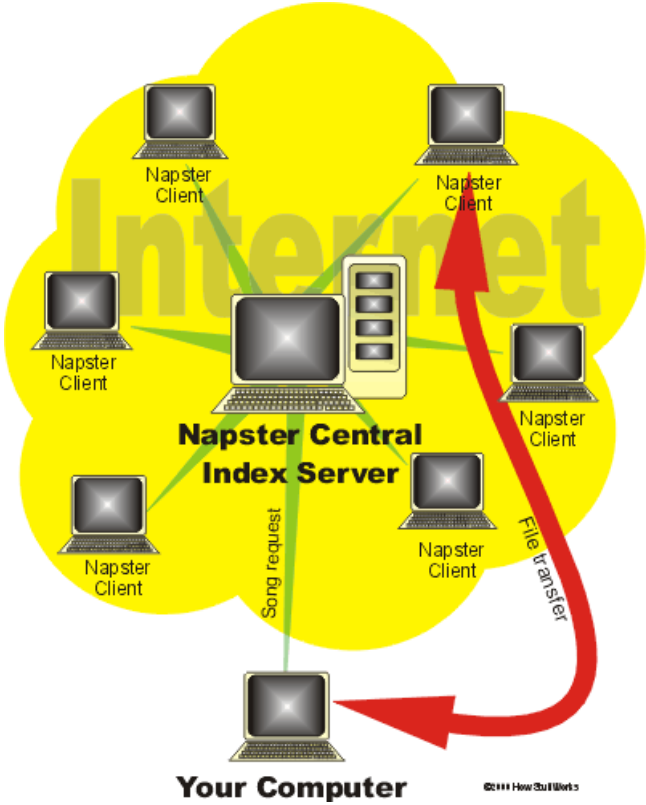
## Skype

- Internet telephony app
- Finding address of remote party: centralized server(s)
- Client-client connection is direct (not through server)

## Instant messaging

- Chatting between two users is P2P
- Presence detection/location centralized:
  - User registers its IP address with central server when it comes online
  - User contacts central server to find IP addresses of buddies

# Case Study: Napster Vs Gnutella



Any problem with this architecture?

# Processes communicating

- Process:** program running within a host
- within same host, two processes communicate using **inter-process communication** (defined by OS)
  - processes in different hosts communicate by exchanging **messages**

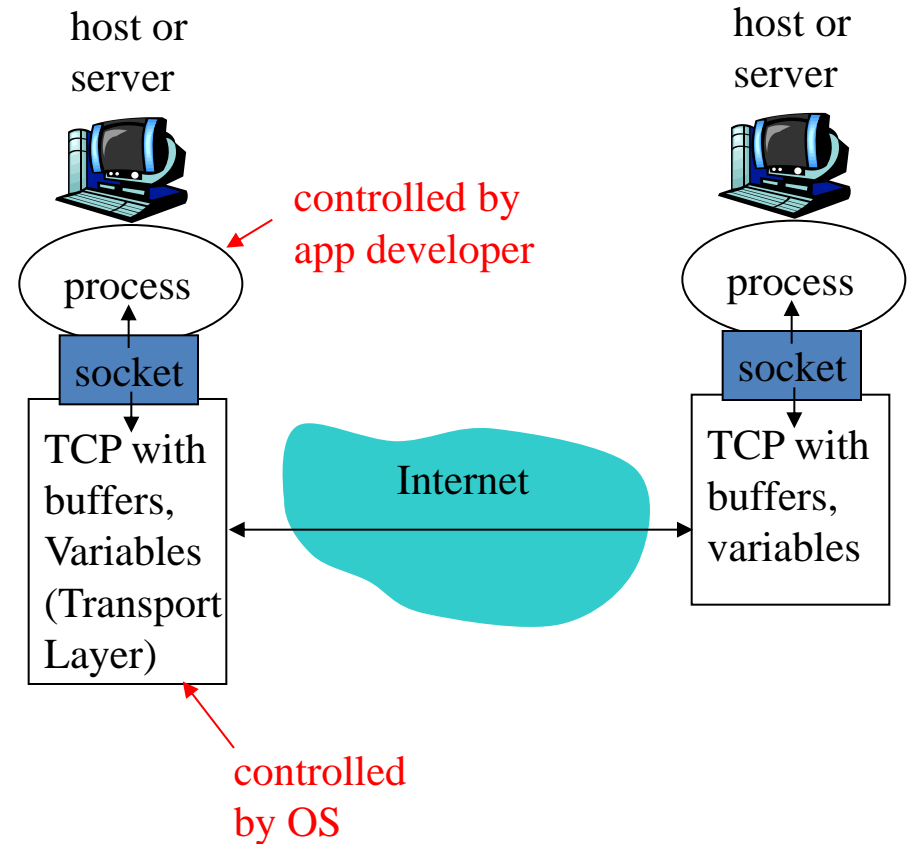
**Client process:** process that initiates communication

**Server process:** process that waits to be contacted

- Note: applications with P2P architectures have client processes & server processes

# Sockets (Interface)

- process sends/receives messages to/from its **socket**
- socket analogous to door
  - sending process shoves message out door
  - sending process relies on transport infrastructure on other side of door which brings message to socket at receiving process



- API: (1) choice of transport protocol; (2) ability to fix a few parameters (**lots more on this later**)

# Addressing processes

- to receive messages, process must have *identifier*
- host device has unique 32-bit IP address
- **Q:** does IP address of host on which process runs suffice for identifying the process?

# Addressing processes

- to receive messages, process must have *identifier*
- host device has unique 32-bit IP address
- **Q:** does IP address of host on which process runs suffice for identifying the process?
  - **Answer:** NO, many processes can be running on same host
- *identifier* includes both **IP address** and **port numbers** associated with process on host
- Example port numbers:
  - HTTP server: 80
  - Mail server: 25
- to send HTTP message to **gaia.cs.umass.edu web server:**
  - IP address: 128.119.245.12
  - Port number: 80
- more shortly...

# Requirements for Message Transport:

## Data loss (Reliable Data Transfer)

- some apps (e.g., audio) can tolerate some loss
- other apps (e.g., file transfer, telnet) require 100% reliable data transfer

## Timing

- some apps (e.g., Internet telephony, interactive games) require low delay to be “effective”

## Bandwidth (Throughput)

- some apps (e.g., multimedia) require minimum amount of bandwidth to be “effective”
- other apps (“elastic apps”) make use of whatever bandwidth they get

## Security

- a transport protocol can encrypt all data

Why is **bandwidth** different from **timing** constraints?



# Transport service requirements of common apps

<b>Application</b>	<b>Data loss</b>	<b>Bandwidth</b>	<b>Time Sensitive</b>
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video:10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

# Internet transport protocols services

## TCP service:

- *connection-oriented*: setup required between client and server processes
- *reliable transport* between sending and receiving process
- *flow control*: sender won't overwhelm receiver
- *congestion control*: throttle sender when network overloaded
- *does not provide*: timing, minimum bandwidth guarantees

## UDP service:

- unreliable data transfer between sending and receiving process
- does not provide: connection setup, reliability, flow control, congestion control, timing, or bandwidth guarantee

Q: Why bother? Why is there a UDP?

# Internet apps: application, transport protocols

<b>Application</b>	<b>Application layer protocol</b>	<b>Underlying transport protocol</b>
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Vonage, Dialpad)	typically UDP

# Chapter 2: Application layer

- 2.1 Principles of network applications
  - app architectures
  - app requirements
- **2.2 Web and HTTP**
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Message Format:

App-layer protocol defines

- Types of messages exchanged
  - e.g., request, response
- Message syntax:
  - what fields in messages & how fields are delineated
- Message semantics
  - meaning of information in fields
- Rules for when and how processes send & respond to messages

## Public-domain protocols:

- defined in RFCs
- allows for interoperability
- e.g., HTTP, SMTP

## Proprietary protocols:

- e.g., KaZaA, Skype

# Web and HTTP

## First some jargon

- **Web page** consists of **objects**
- Object can be HTML file, JPEG image, Java applet, audio file,...
- Web page consists of **base HTML-file** which includes several referenced objects
- Each object is addressable by a **URL**
- Example URL:

`www.someschool.edu/someDept/pic.gif`

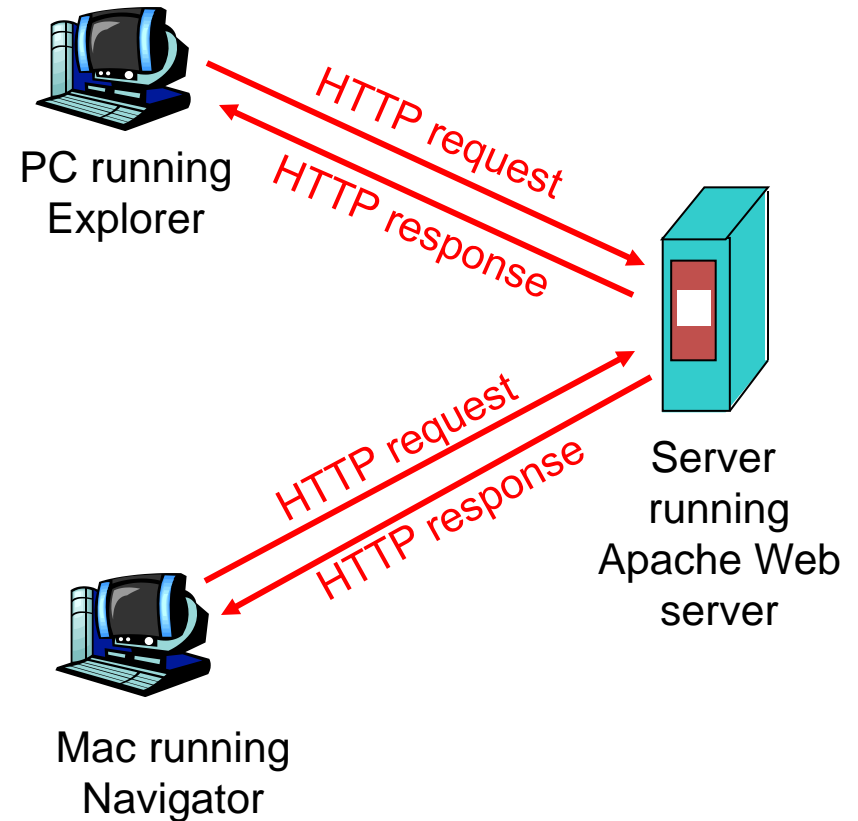
host name

path name

# HTTP overview

## HTTP: hypertext transfer protocol

- Web's application layer protocol
- client/server model
  - *client*: browser that requests, receives, "displays" Web objects
  - *server*: Web server sends objects in response to requests
- HTTP 1.0: RFC 1945
- HTTP 1.1: Obsoletes RFC 2068  
RFC 2616



# HTTP overview (continued)

## Uses TCP:

- client initiates TCP connection (creates socket) to server, port 80
- server accepts TCP connection from client
- HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server)
- TCP connection closed

## HTTP is “stateless”

- server maintains no information about past client requests

aside  
Protocols that maintain “state” are complex!

- ❑ past history (state) must be maintained
- ❑ if server/client crashes, their views of “state” may be inconsistent, must be reconciled



# HTTP connections

## Nonpersistent HTTP

- At most one object is sent over a TCP connection
- HTTP/1.0 uses nonpersistent HTTP

## Persistent HTTP

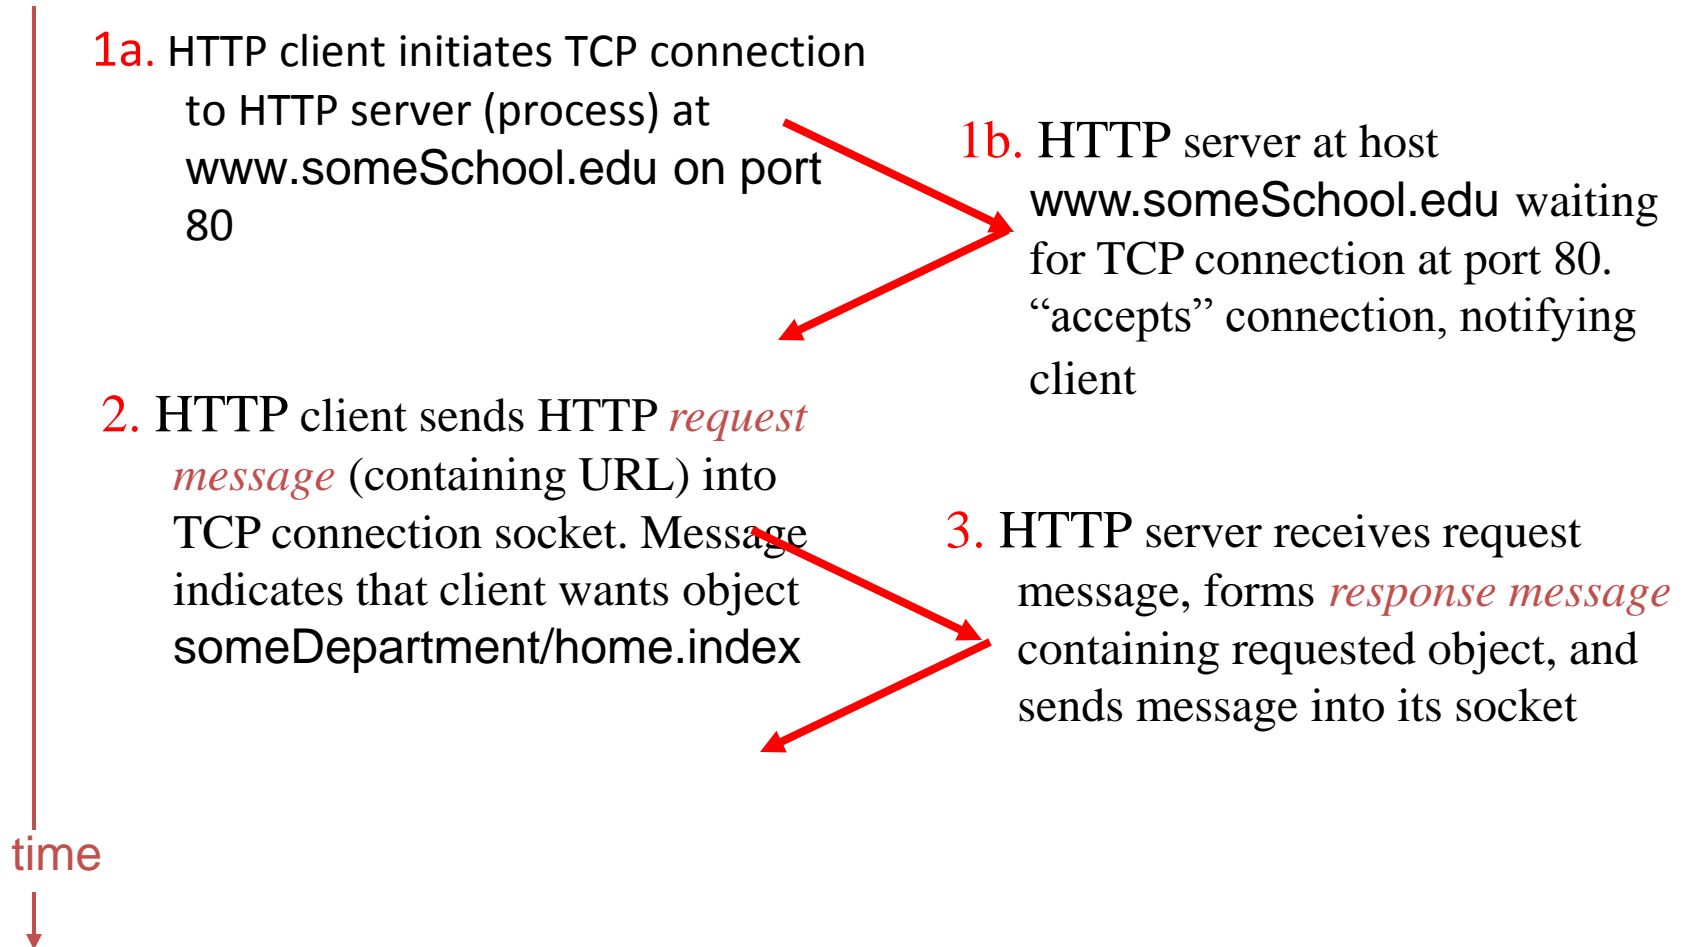
- Multiple objects can be sent over single TCP connection between client and server
- HTTP/1.1 uses persistent connections in default mode

# Nonpersistent HTTP

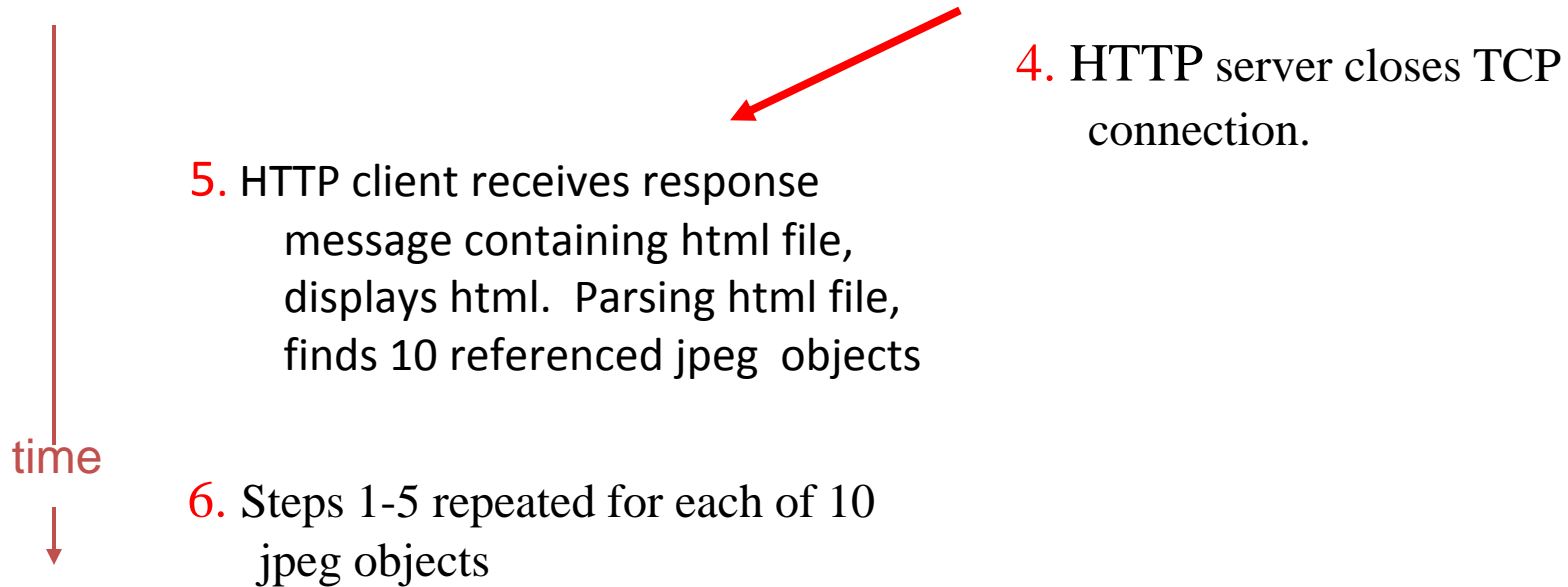
(contains text,  
references to 10  
jpeg images)

Suppose user enters URL

`www.someSchool.edu/someDepartment/home.index`



# Nonpersistent HTTP (cont.)



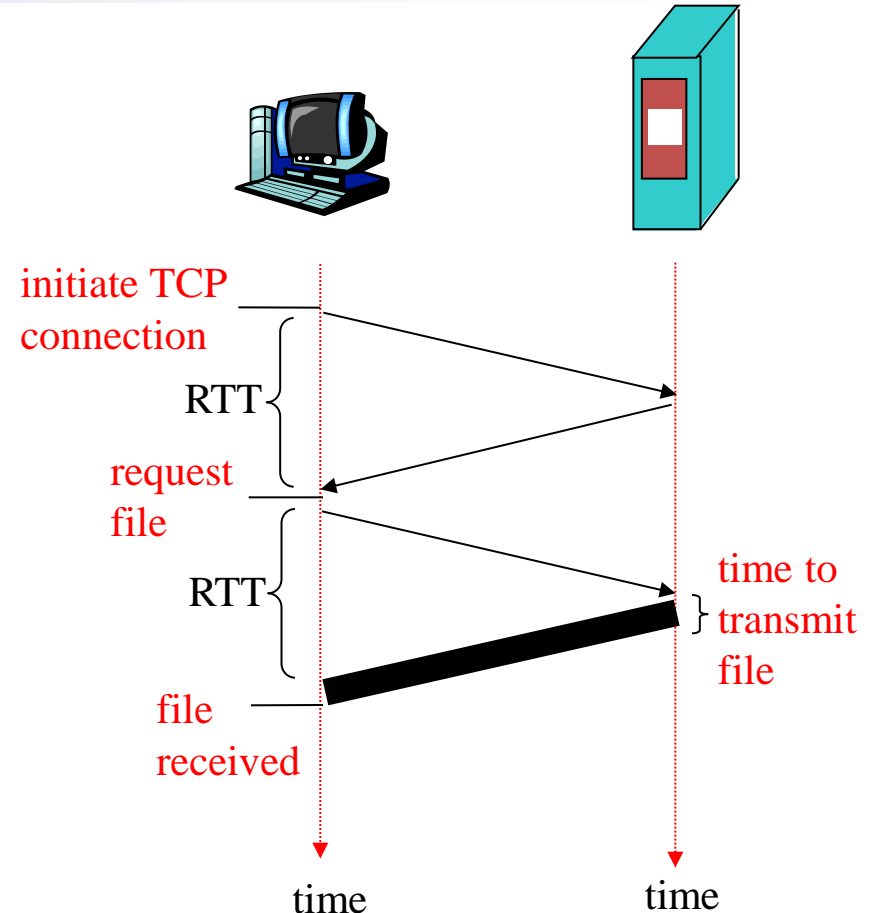
# Non-Persistent HTTP: Response time

**Round Trip Time (RTT)** = time to send a small packet to travel from client to server and back

## Response time:

- one RTT to initiate TCP connection
- one RTT for HTTP request and first few bytes of HTTP response to return
- file transmission time

**total =  $2RTT + \langle \text{file transmit time} \rangle$**



# Persistent HTTP

## Nonpersistent HTTP issues:

- requires 2 RTTs per object
- OS overhead for *each* TCP connection
- browsers often open parallel TCP connections to fetch referenced objects

## Persistent HTTP

- server leaves connection open after sending response
- subsequent HTTP messages between same client/server sent over open connection

## Persistent *without* pipelining:

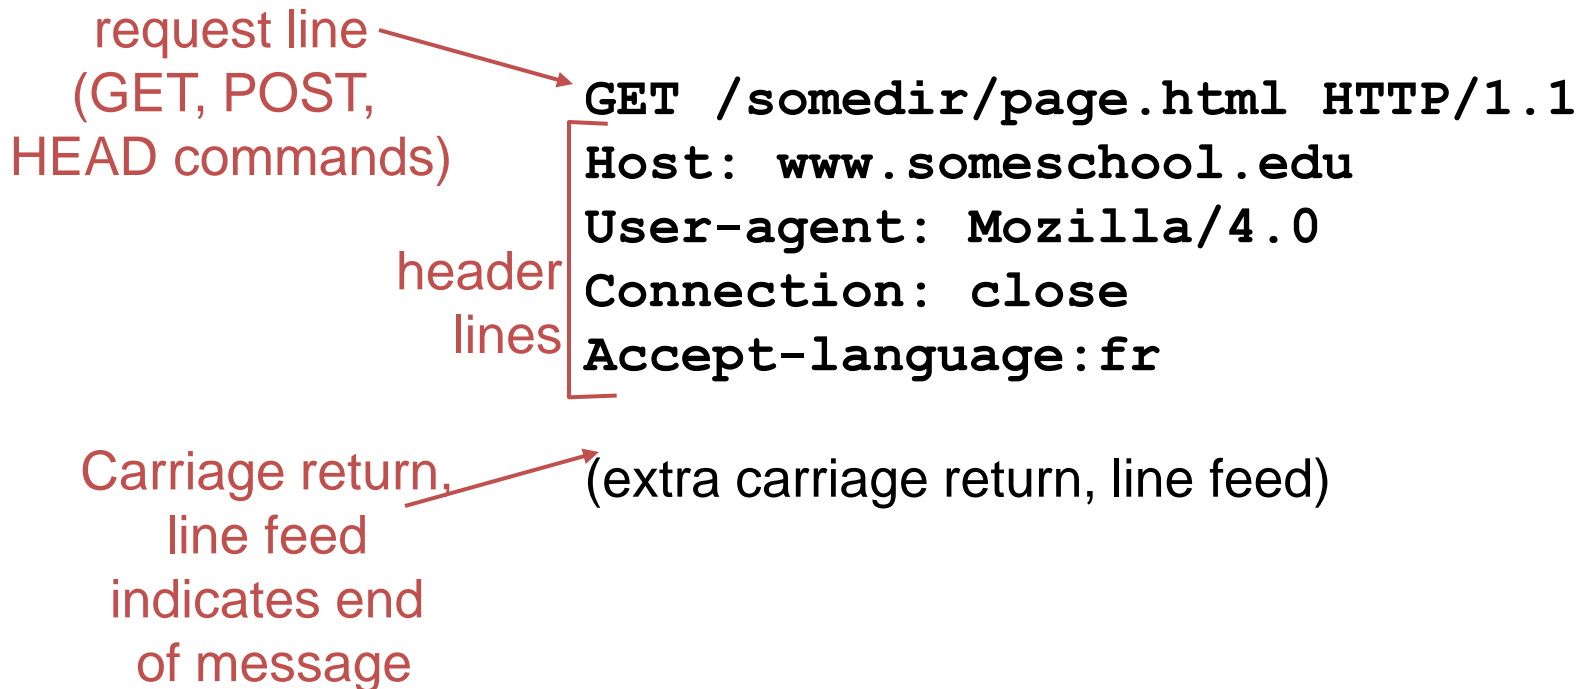
- client issues new request only when previous response has been received
- one RTT for each referenced object

## Persistent *with* pipelining:

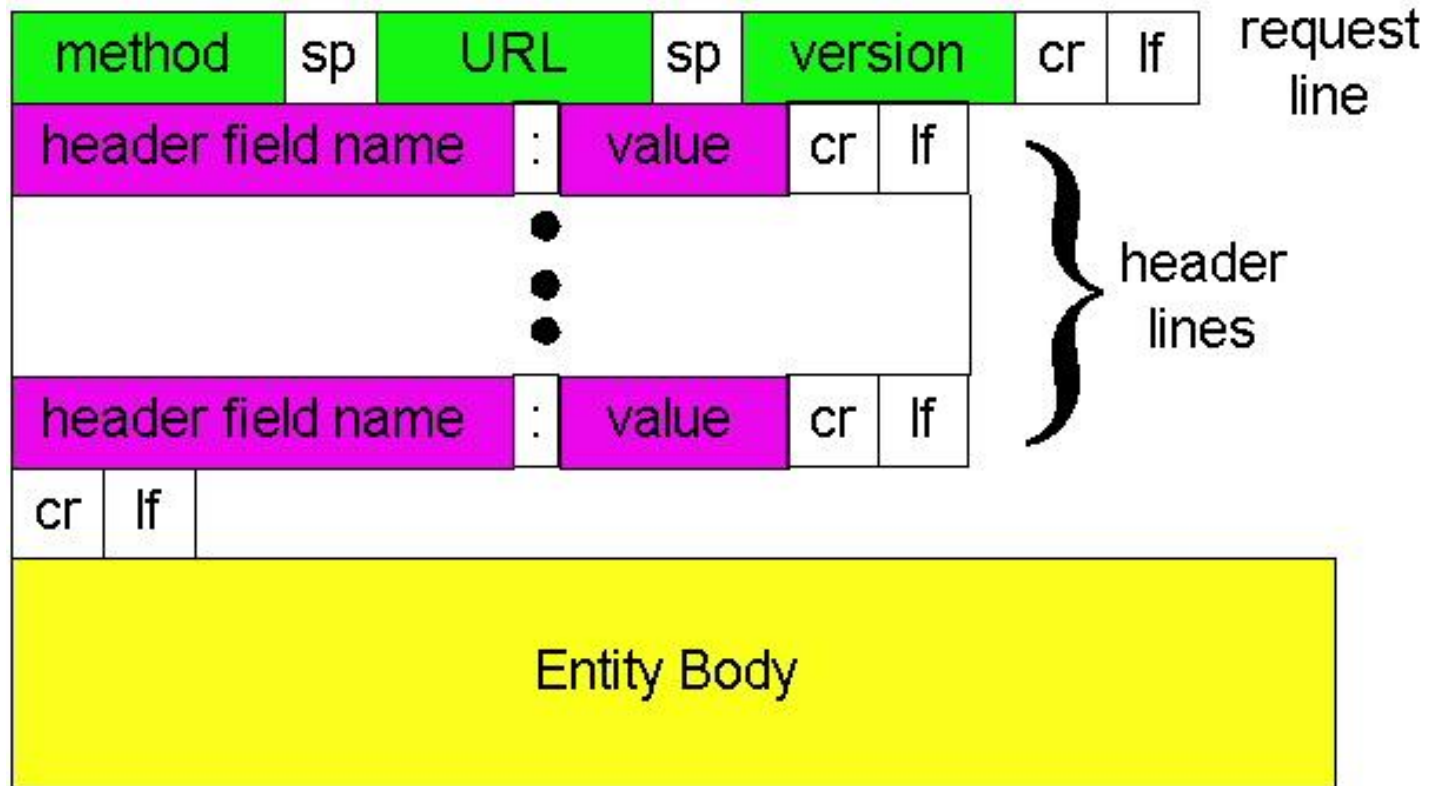
- default in HTTP/1.1
- client sends requests as soon as it encounters a referenced object
- as little as one RTT for all the referenced objects

# HTTP request message

- two types of HTTP messages: *request, response*
- HTTP request message:
  - ASCII (human-readable format)



# HTTP request message: general format



# Uploading form input

## Post method:

- Web page often includes form input
- Input is uploaded to server in entity body

## URL method:

- Uses GET method
- Input is uploaded in URL field of request line:

`www.somesite.com/animalsearch?monkeys&banana`



# Method types

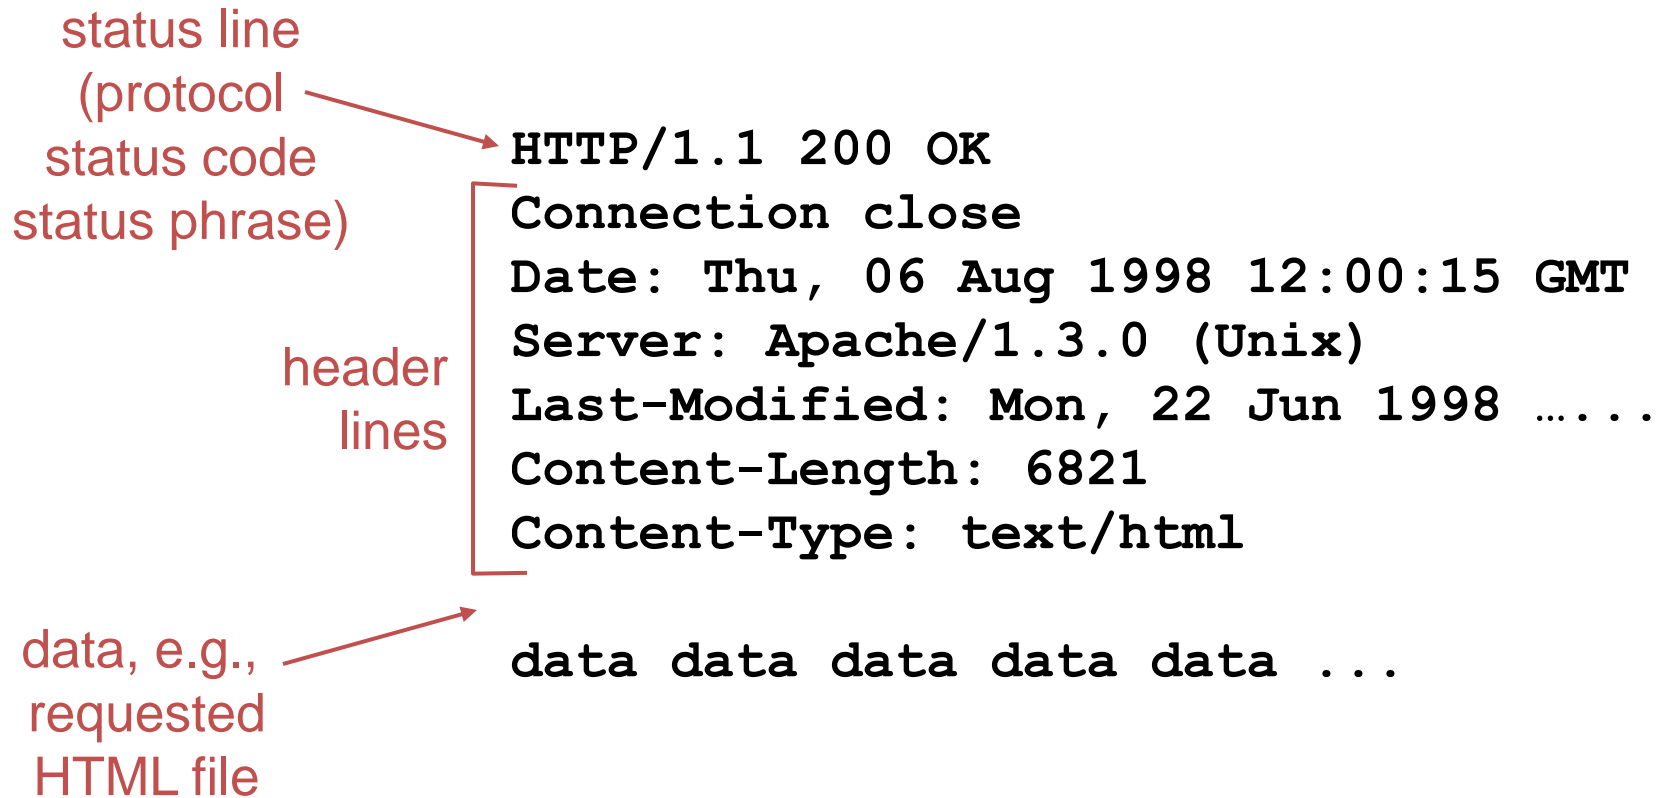
## HTTP/1.0

- GET
- POST
- HEAD
  - asks server to leave requested object out of response

## HTTP/1.1

- GET, POST, HEAD
- PUT
  - uploads file in entity body to path specified in URL field
- DELETE
  - deletes file specified in the URL field

# HTTP response message



# HTTP response status codes

In first line in server->client response message.

A few sample codes:

## **200 OK**

- request succeeded, requested object later in this message

## **301 Moved Permanently**

- requested object moved, new location specified later in this message (Location:)

## **400 Bad Request**

- request message not understood by server

## **404 Not Found**

- requested document not found on this server

## **505 HTTP Version Not Supported**

# Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

```
telnet cis.poly.edu 80
```

Opens TCP connection to port 80 (default HTTP server port) at cis.poly.edu. Anything typed in sent to port 80 at cis.poly.edu

2. Type in a GET HTTP request:

```
GET /~ross/ HTTP/1.1  
Host: cis.poly.edu
```

By typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. Look at response message sent by HTTP server!

# Chapter 2: Application layer

- 2.1 Principles of network applications
  - app architectures
  - app requirements
- **2.2 Web and HTTP**
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# User-server state: cookies

Many major Web sites use cookies

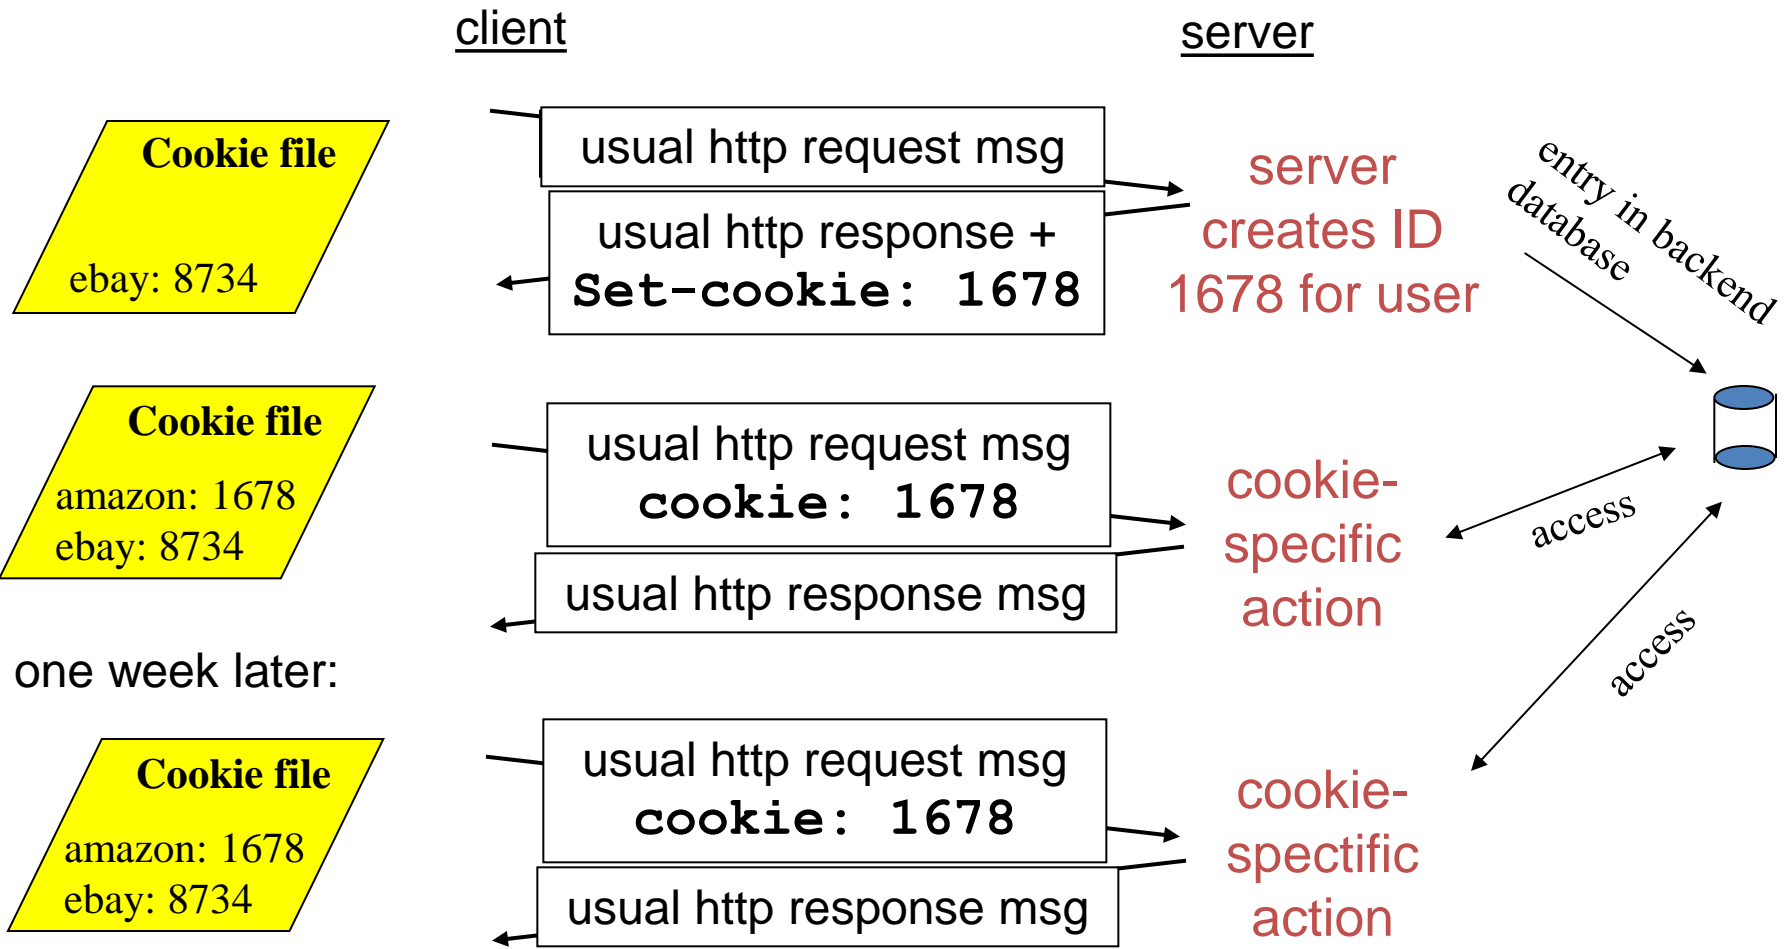
## Four components:

- 1) cookie header line in HTTP *response* message
- 2) cookie header line in HTTP *request* message
- 3) cookie file kept on user's host, managed by user's browser
- 4) back-end database at Web site

## Example:

- Susan access Internet always from same PC
- She visits a specific e-commerce site for first time
- When initial HTTP requests arrives at site, site creates a unique ID and creates an entry in backend database for ID

# Cookies: keeping "state" (cont.)



# Cookies (continued)

## What cookies can bring:

- authorization
- shopping carts
- recommendations
- user session state (Web e-mail)

— aside —

## Cookies and privacy:

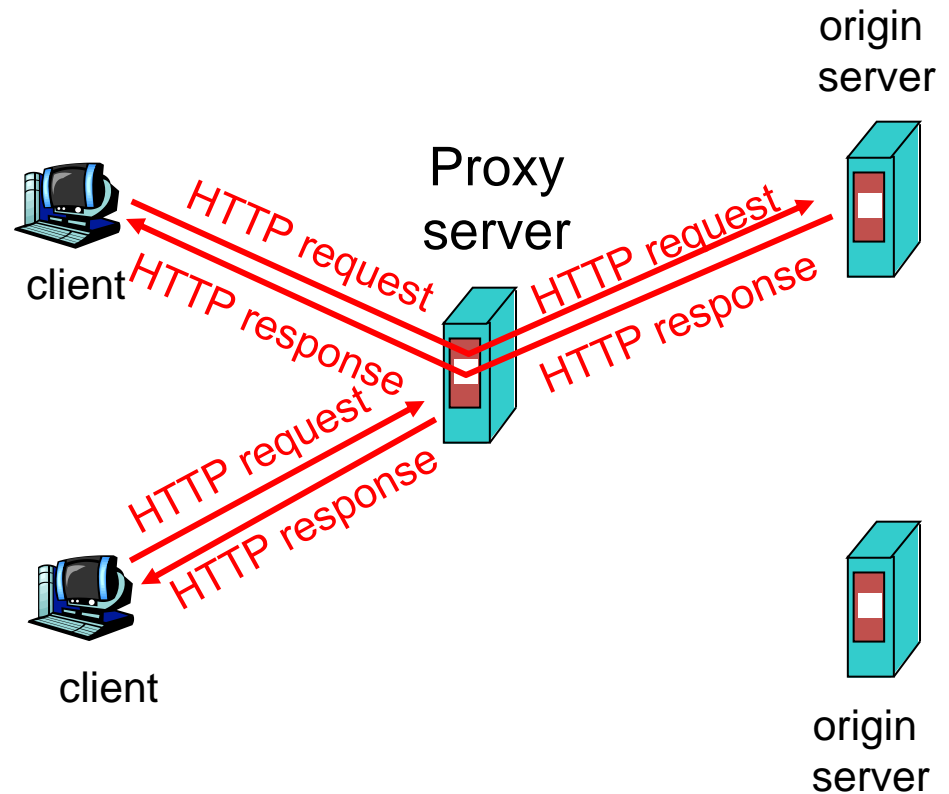
- cookies permit sites to learn a lot about you
- you may supply name and e-mail to sites
- search engines use redirection & cookies to learn yet more
- advertising companies obtain info across sites



# Web caches (proxy server)

**Goal:** satisfy client request without involving origin server

- user sets browser: Web accesses via cache
- browser sends all HTTP requests to cache
  - object in cache: cache returns object
  - else cache requests object from origin server, then returns object to client



# More about Web caching

- Cache acts as both client and server
- Typically cache is installed by ISP (university, company, residential ISP)

## Why Web caching?

- Reduce response time for client request.
- Reduce traffic on an institution's access link.
- Internet dense with caches enables "poor" content providers to effectively deliver content (but so does P2P file sharing)

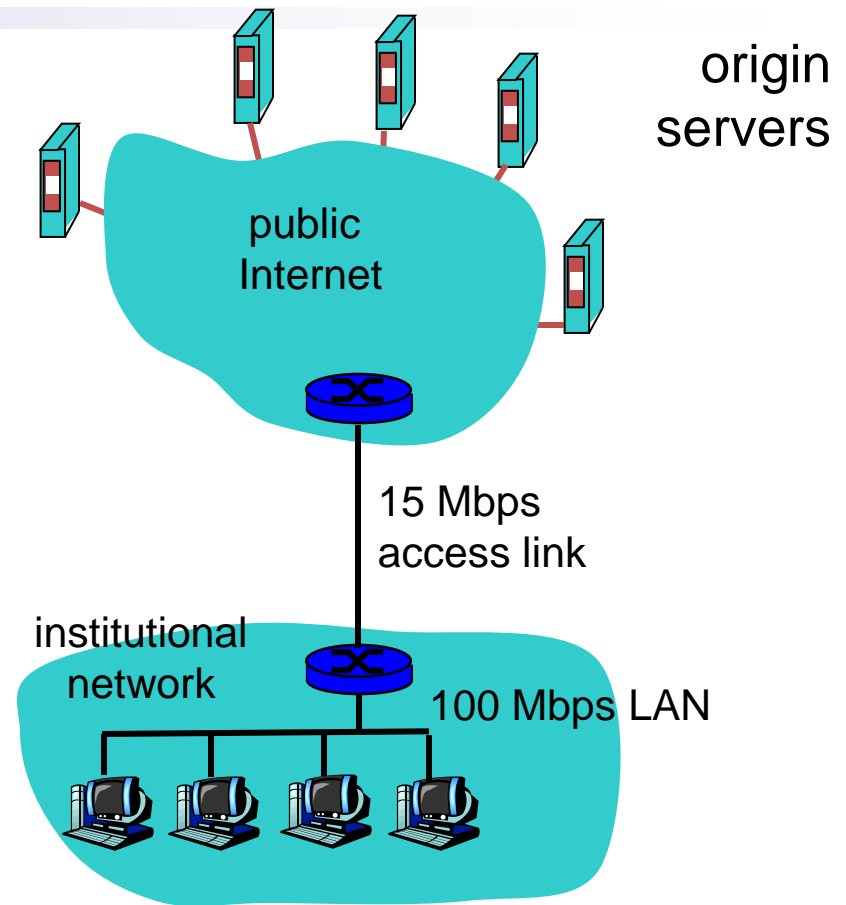
# Caching example

## Assumptions

- average object size = 100,0000 bits  
= 1 Mbits
- avg. request rate from institution's browsers to origin servers = 15 requests/sec
- delay from institutional router to any origin server and back to router = 2 sec

## Consequences

- utilization on LAN = 15%
- utilization on access link = 100%
- total delay = Internet delay + access delay + LAN delay  
= 2 sec + minutes + milliseconds



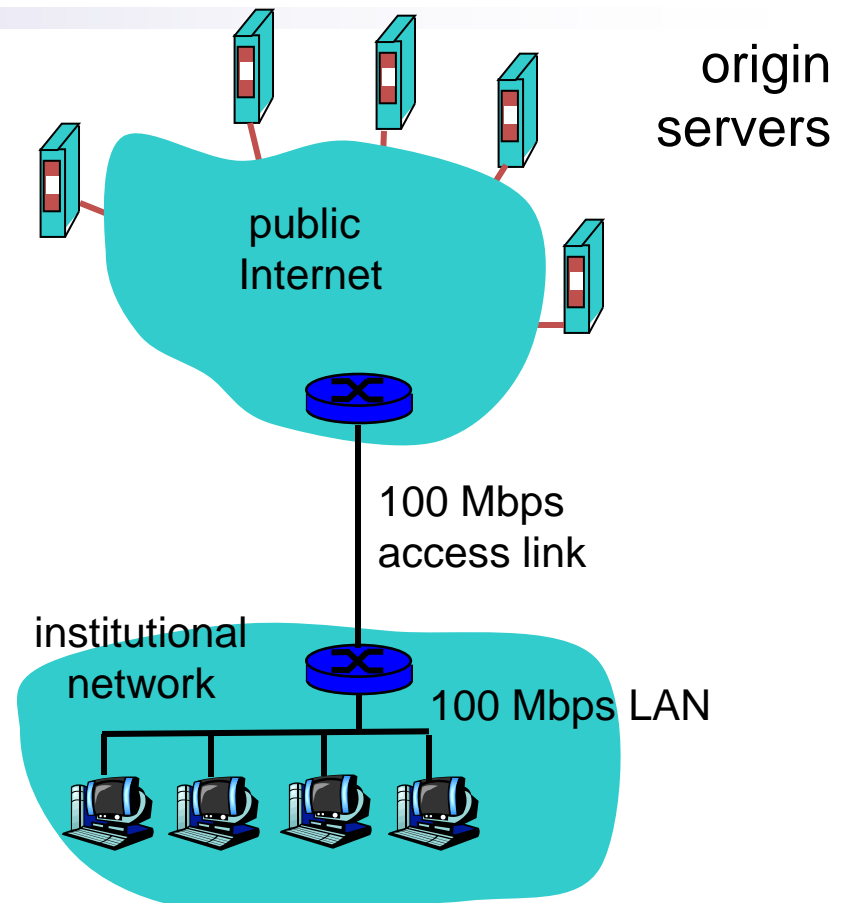
# Caching example (cont)

## Possible solution

- increase bandwidth of access link to, say, 100 Mbps

## Consequences

- utilization on LAN = 15%
- utilization on access link = 15%
- Total delay = Internet delay + access delay + LAN delay  
= 2 sec + msec + msec
- often a costly upgrade



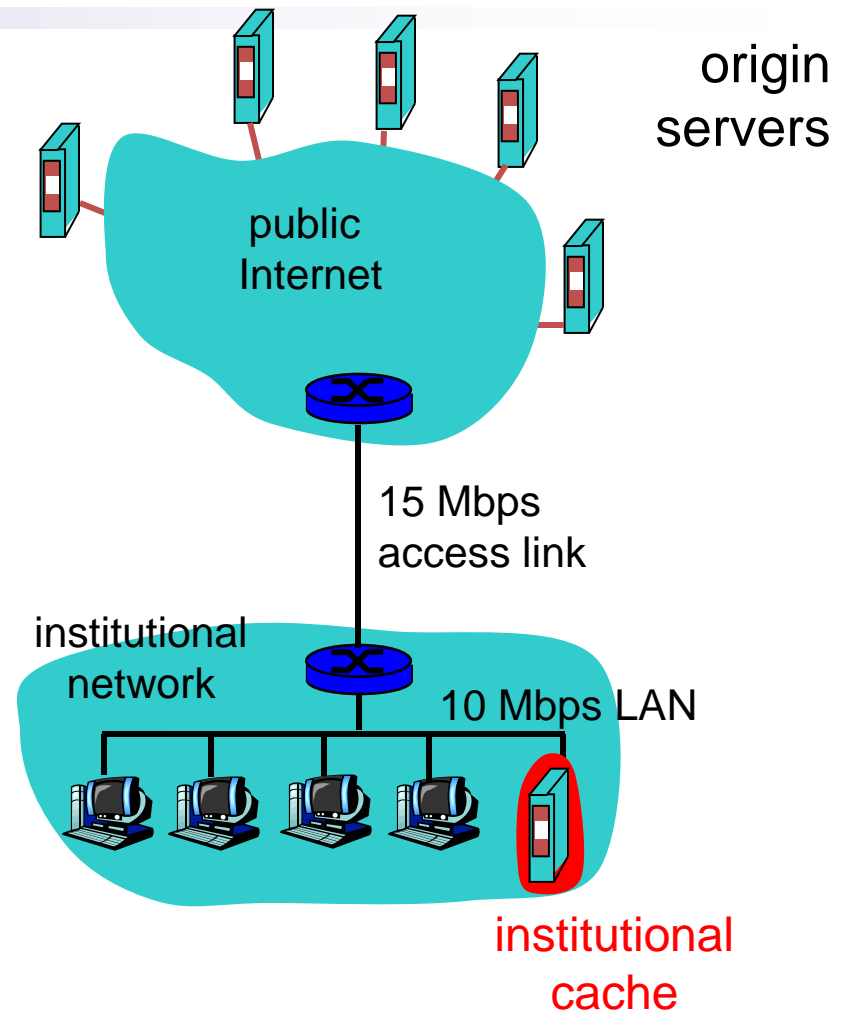
# Caching example (cont)

## Install cache

- suppose hit rate is 0.4

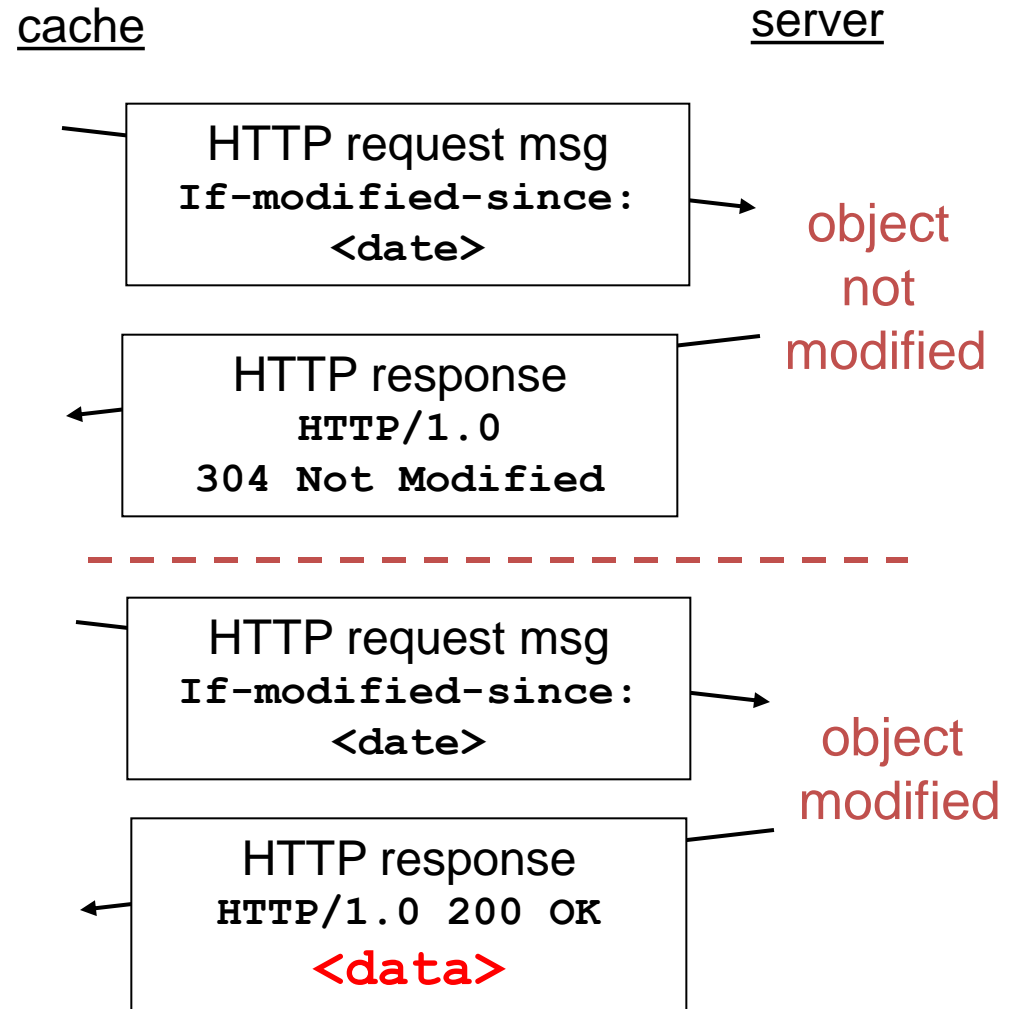
## Consequence

- 40% requests will be satisfied almost immediately
- 60% requests satisfied by origin server
- utilization of access link reduced to 60%, resulting in negligible delays (say 10 msec)
- total avg delay = Internet delay + access delay + LAN delay =  
 $0.6 \cdot (2.0) \text{ secs} + 0.6 \cdot (0.01) \text{ secs} + 0.4 \cdot (\text{milliseconds}) > 1.2 \text{ secs}$



# Conditional GET

- **Goal:** don't send object if cache has up-to-date cached version
- cache: specify date of cached copy in HTTP request  
`If-modified-since: <date>`
- server: response contains no object if cached copy is up-to-date:  
`HTTP/1.0 304 Not Modified`



---

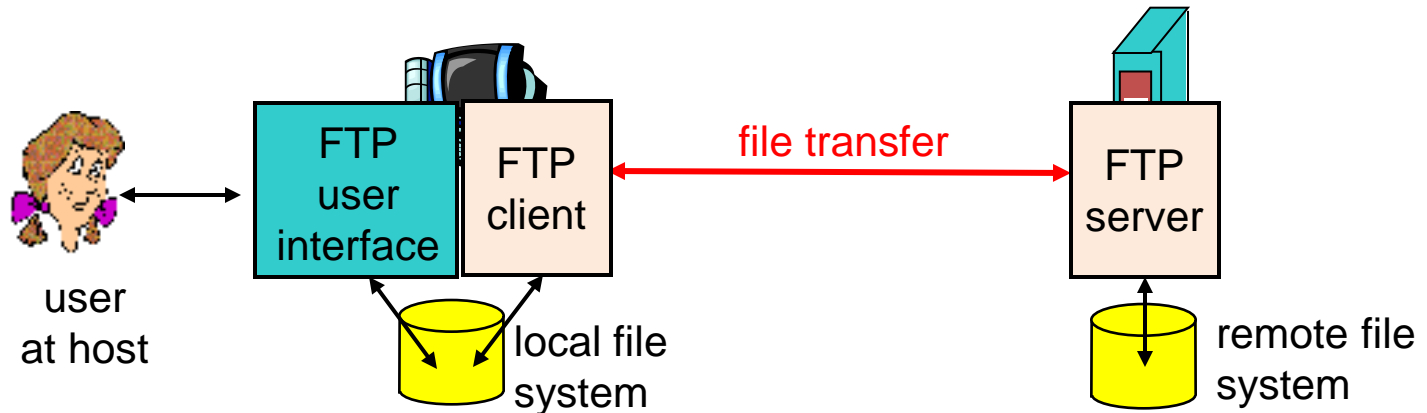
Questions?

# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server



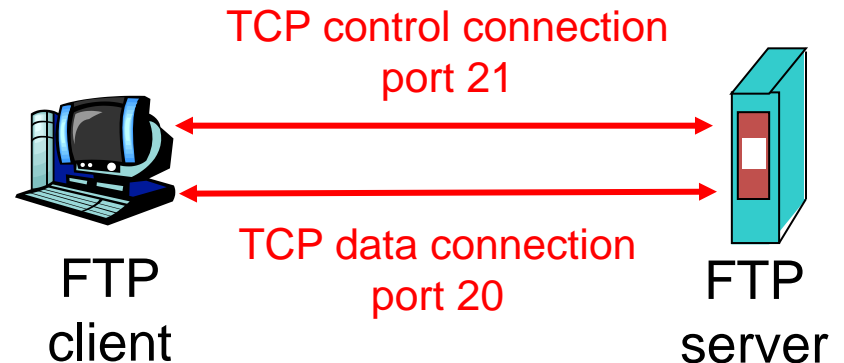
# FTP: the file transfer protocol



- transfer file to/from remote host
- client/server model
  - *client*: side that initiates transfer (either to/from remote)
  - *server*: remote host
- ftp: RFC 959
- ftp server: port 21

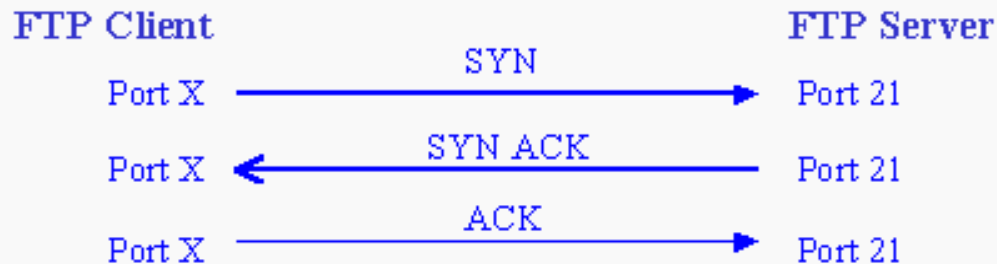
# FTP: separate control, data connections

- FTP client contacts FTP server at port 21, specifying TCP as transport protocol
  - Client obtains authorization
- Client browses remote directory by sending control commands
- When server receives a command, opens *TCP data connection* to client
- After transferring one file, server closes connection

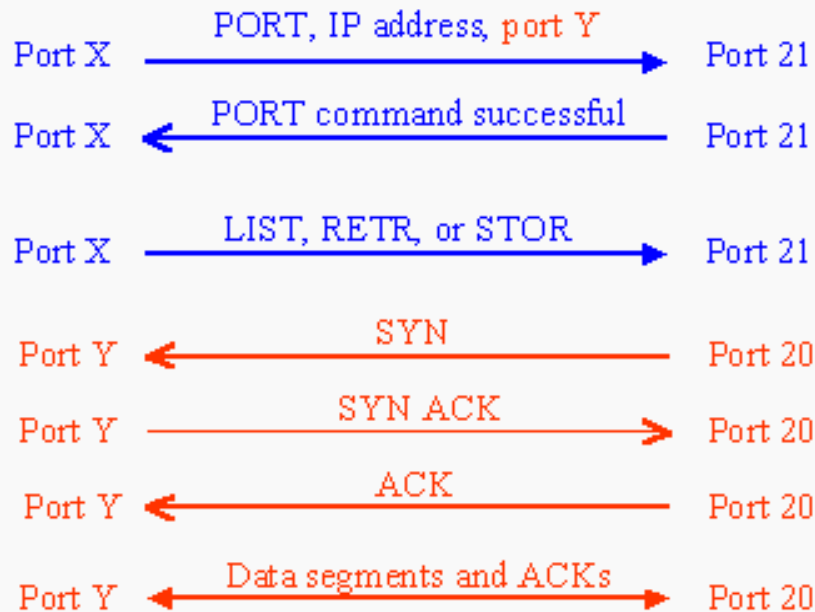


- Server opens a second TCP data connection to transfer another file.
- Control connection: “out of band”
- FTP server maintains “state”: current directory, earlier authentication
- Data Connections: non-persistent

# FTP Timeline



User lists directory or gets or puts a file



# FTP commands, responses

## Sample commands:

- sent as ASCII text over control channel
- **USER *username***
- **PASS *password***
- **LIST** return list of file in current directory
- **RETR *filename*** retrieves (gets) file
- **STOR *filename*** stores (puts) file onto remote host

## Sample return codes

- status code and phrase (as in HTTP)
- **331 Username OK, password required**
- **125 data connection already open; transfer starting**
- **425 Can't open data connection**
- **452 Error writing file**

# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

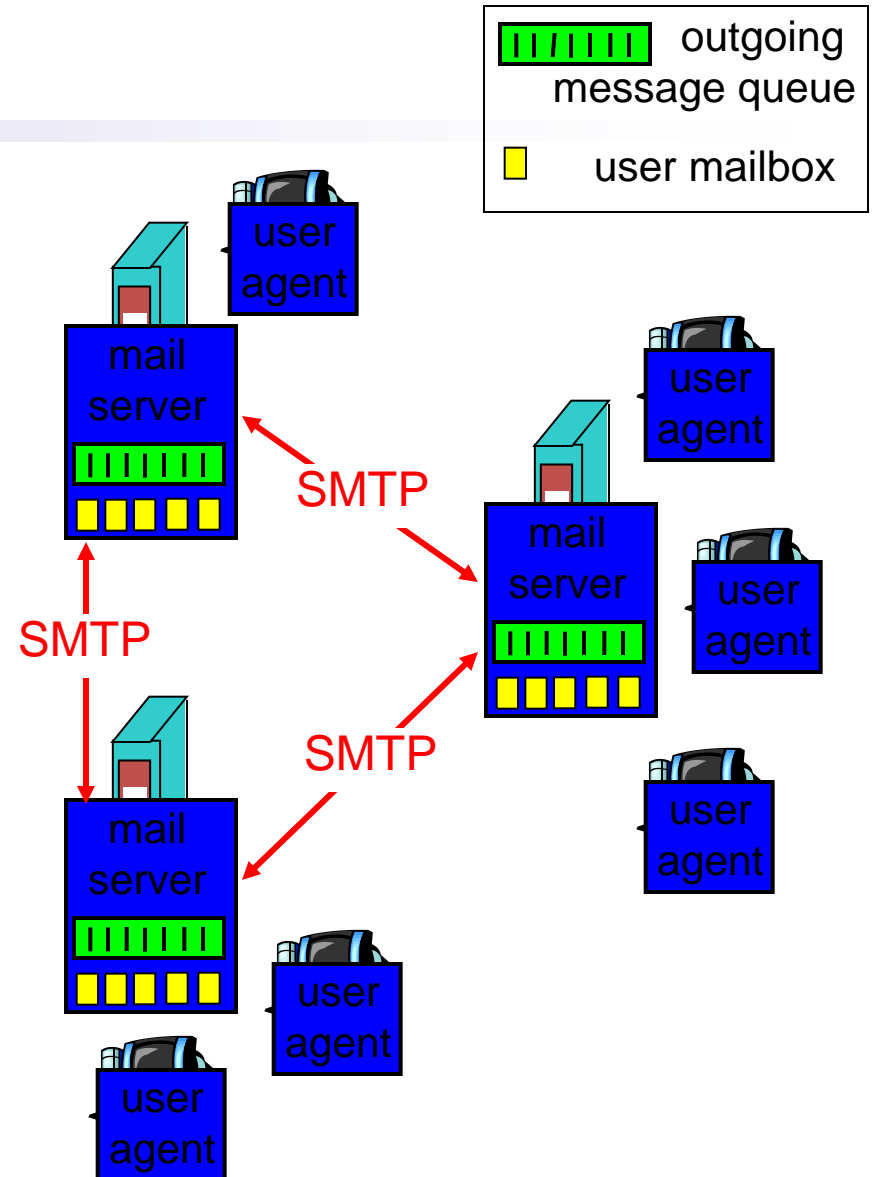
# Electronic Mail

## Three major components:

- user agents
- mail servers
- simple mail transfer protocol: SMTP

## User Agent

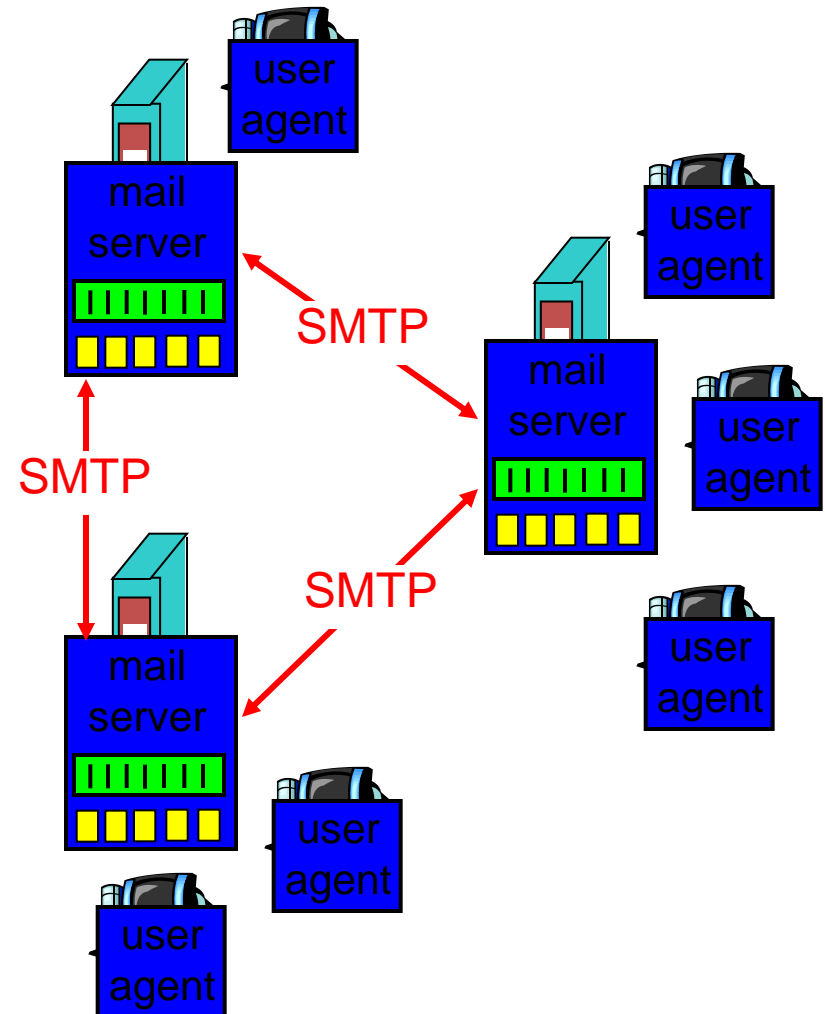
- a.k.a. “mail reader”
- composing, editing, reading mail messages
  - e.g., Eudora, Outlook, elm, Netscape Messenger
- outgoing, incoming messages stored on server



# Electronic Mail: mail servers

## Mail Servers

- **mailbox** contains incoming messages for user
- **message queue** of outgoing (to be sent) mail messages
- **SMTP protocol** between mail servers to send email messages
  - client: sending mail server
  - “server”: receiving mail server



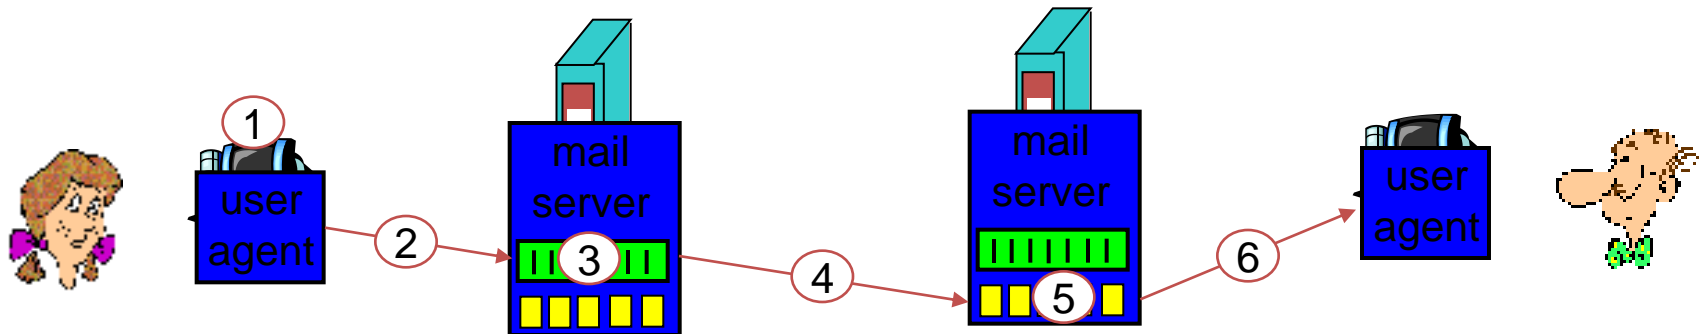
# Electronic Mail: SMTP [RFC 2821]

- uses TCP on port 25 to reliably transfer email
- direct transfer: sending server to receiving server
- three phases of transfer
  - handshaking (greeting)
  - transfer of messages
  - Closure
- command/response interaction
  - **commands:** ASCII text
  - **response:** status code and phrase



# Scenario: Alice Emails Bob

- 1) Alice uses UA to compose message and “to” bob@some school.edu
- 2) Alice’s UA sends message to her mail server; message placed in message queue
- 3) Client side of SMTP opens TCP connection with Bob’s mail server
- 4) SMTP client sends Alice’s message over the TCP connection
- 5) Bob’s mail server places the message in Bob’s mailbox
- 6) Bob invokes his user agent to read message



# SMTP Commands to send email

- Telenet into port 25
  - HELO hostname
  - MAIL FROM:
  - RCPT TO
  - RCPT TO ...
  - DATA
  - ... text ...
  - .
  - QUIT
- You can try doing this yourself

# Try SMTP interaction for yourself:

- `telnet servername 25`
  - see 220 reply from server
  - enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands
- above lets you send email without using email client (reader)

# SMTP: final words

- SMTP uses persistent connections
- SMTP requires message (header & body) to be in 7-bit ASCII
- SMTP server uses `CRLF.CRLF` to determine end of message

## Comparison with HTTP:

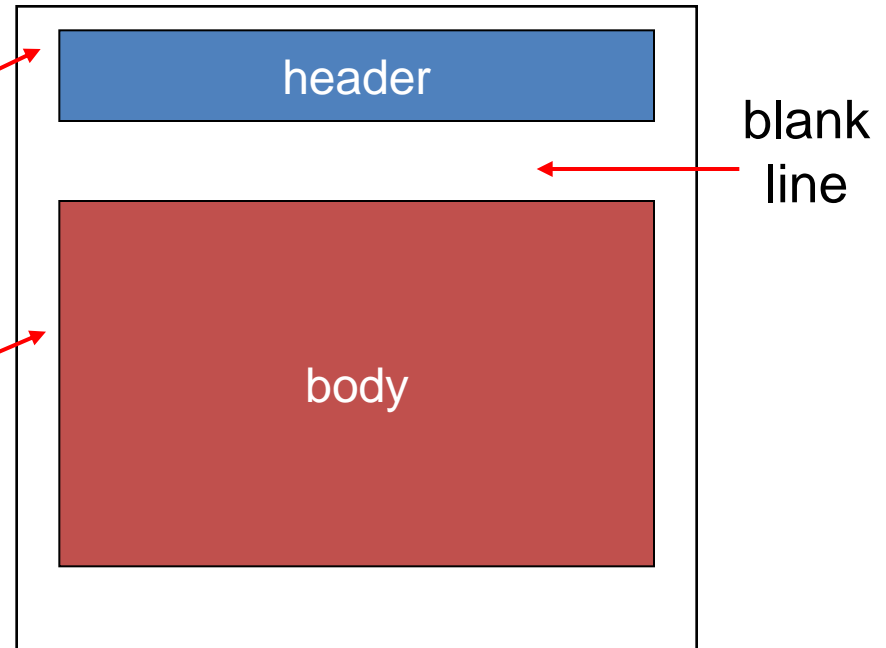
- HTTP: pull
- SMTP: push
- both have ASCII command/response interaction, status codes
- HTTP: each object encapsulated in its own response msg (such as text, image or other media types)
- SMTP: multiple objects sent in multipart msg

# Mail message format

SMTP: protocol for exchanging email  
msgs

RFC 822: standard for text message  
format:

- header lines, e.g.,
  - To:
  - From:
  - Subject:*different from SMTP commands!*
- body
  - the “message”, ASCII characters only



# Message format: multimedia extensions

- MIME: multimedia mail extension, RFC 2045, 2056
- additional lines in msg header declare MIME content type
  - Think of image attachments with your email

MIME version

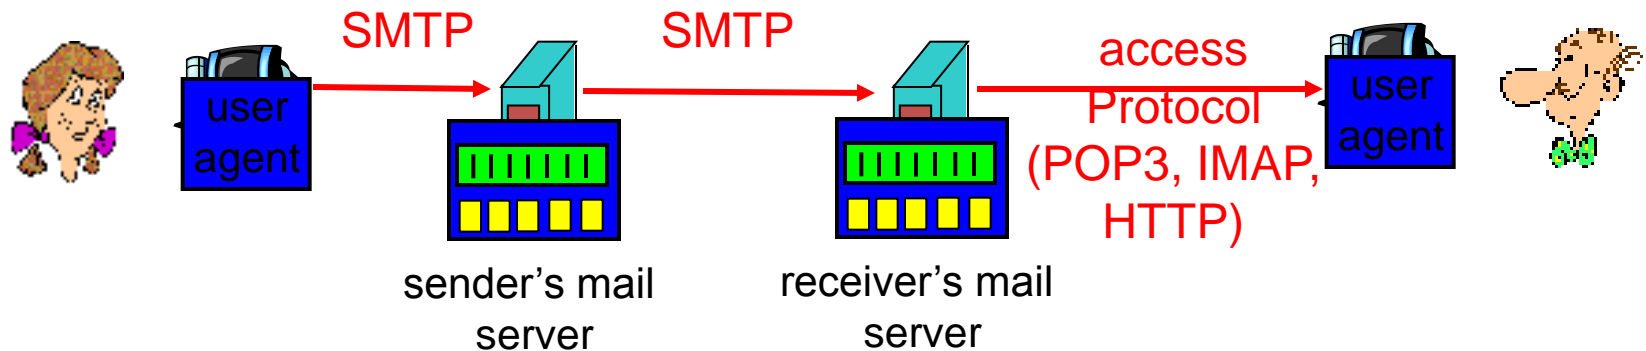
method used  
to encode data

multimedia data  
type, subtype,  
parameter declaration

encoded data

```
From: alice@crepes.fr
To: bob@hamburger.edu
Subject: Picture of yummy crepe.
MIME-Version: 1.0
Content-Transfer-Encoding: base64
Content-Type: image/jpeg
base64 encoded data .....
.....
.....base64 encoded data
```

# Mail access protocols



- SMTP: delivery/storage to receiver's server
- Mail access protocol: retrieval from server
  - POP: Post Office Protocol [RFC 1939]
    - authorization (agent <-->server) and download
  - IMAP: Internet Mail Access Protocol [RFC 1730]
    - more features (more complex)
    - manipulation of stored msgs on server
  - HTTP: Hotmail , Yahoo! Mail, etc.

What's the Difference?

# POP3 protocol

## authorization phase

- client commands:
  - **user**: declare username
  - **pass**: password
- server responses
  - **+OK**
  - **-ERR**

## transaction phase, client:

- **list**: list message numbers
- **retr**: retrieve message by number
- **dele**: delete
- **Quit** (update phase)

```
S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on
C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 1 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off
```



# POP3 (more) and IMAP

## More about POP3

- Previous example uses “download and delete” mode.
- Bob cannot re-read e-mail if he changes client
- “Download-and-keep”: copies of messages on different clients
- POP3 is stateless across sessions

## IMAP

- Keep all messages in one place: the server
- Allows user to organize messages in folders
- IMAP keeps user state across sessions:
  - names of folders and mappings between message IDs and folder name

## Web-based Email

- Hotmail, Gmail, Yahoo
  - HTTP then SMTP then HTTP

# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# DNS: Domain Name System

- Imagine a world without DNS
- You would have to remember the IP addresses of
  - Every website you want to visit
  - Your bookmarks will be a list of IP addresses

- You will speak like

*“I went to 167.33.24.10, and there was an awesome link to 153.11.35.81... “*

# DNS: Domain Name System

**People:** many identifiers:

- SSN, name, passport #

**Internet hosts, routers:**

- IP address (32 bit) - used for addressing datagrams
- “name”, e.g.,  
www.yahoo.com - used by humans

**Q:** map between IP addresses and name ?

**Domain Name System:**

- ***distributed database*** implemented in hierarchy of many *name servers*
- *application-layer protocol* host, routers, name servers to communicate to *resolve* names (address/name translation)
  - note: core Internet function, implemented as application-layer protocol
  - complexity at network’s “edge”

# DNS Services

## DNS services

- Hostname to IP address translation
- Host aliasing
  - Canonical and alias names
  - Alias: `enterprise.com` or [www.enterprise.com](http://www.enterprise.com)  
  
Canonical: `relay1.westcoast.enterprise.com`
- Load distribution
  - Replicated Web servers: set of IP addresses for one canonical name

- Mail server aliasing
  - `bob@hotmail.com`  
Canonical hostname: `relay1.westcoast.hotmail.com`

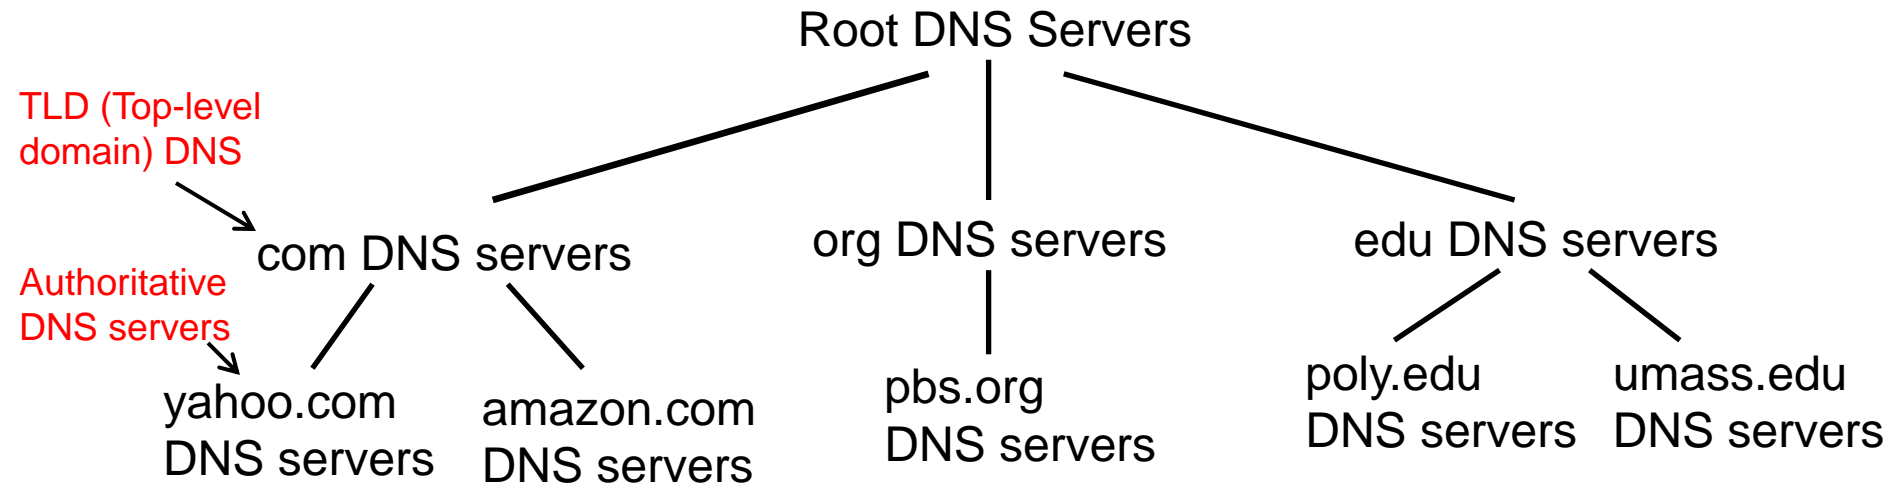
# DNS Implications

## Why not centralize DNS?

- single point of failure
- traffic volume
- distant centralized database
- Maintenance

doesn't *scale!*

# Distributed, Hierarchical Database

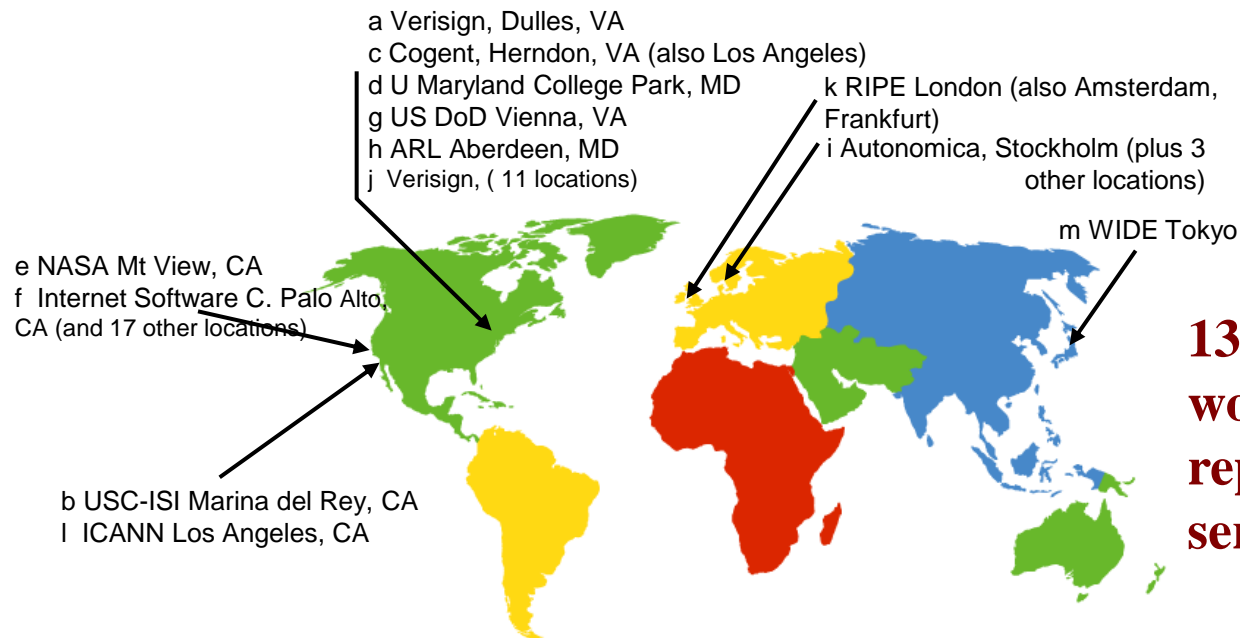


Client wants IP for [www.amazon.com](http://www.amazon.com); 1<sup>st</sup> approx:

- Client queries a root server to find [.com](#) DNS server
- Client queries com DNS server to get [amazon.com](#) DNS server
- Client queries amazon.com DNS server to get IP address for [www.amazon.com](http://www.amazon.com)

# DNS: Root name servers

- contacted by local name server that can not resolve name
- root name server:
  - contacts **authoritative name server** if name mapping not known
  - gets mapping
  - returns mapping to local name server



**13 root name servers  
worldwide: A network of  
replicated servers (247 root  
servers)**



# TLD and Authoritative Servers

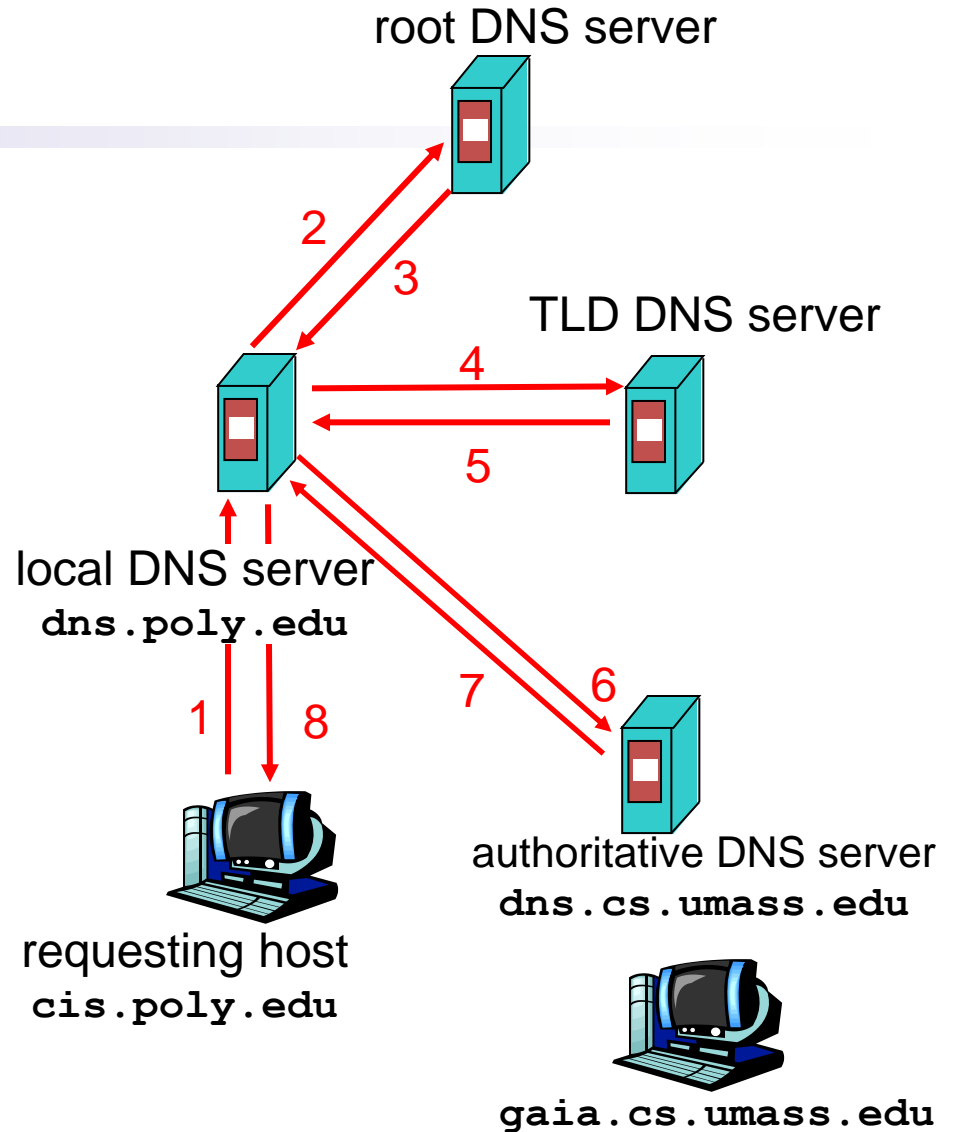
- **Top-level domain (TLD) servers:**
  - responsible for com, org, net, edu, etc.
  - all top-level country domains uk, fr, ca, jp.
  - Network solutions maintains servers for com TLD
  - Educause for edu TLD
- **Authoritative DNS servers:**
  - An organization's DNS servers,
    - providing authoritative hostname to IP mappings for organization's servers (e.g., Web and mail).
  - Can be maintained by organization or service provider

# Local Name Server

- Local DNS server
  - Does not strictly belong to hierarchy
- Each ISP (residential, company, univ) has one
  - Also called “default name server”
- When a host makes a DNS query
  - query is sent to its local DNS server
  - Acts as a proxy, forwards query into hierarchy

# Example

- **Iterative Querying**  
Host at cis.poly.edu  
wants IP address for  
gaia.cs.umass.edu



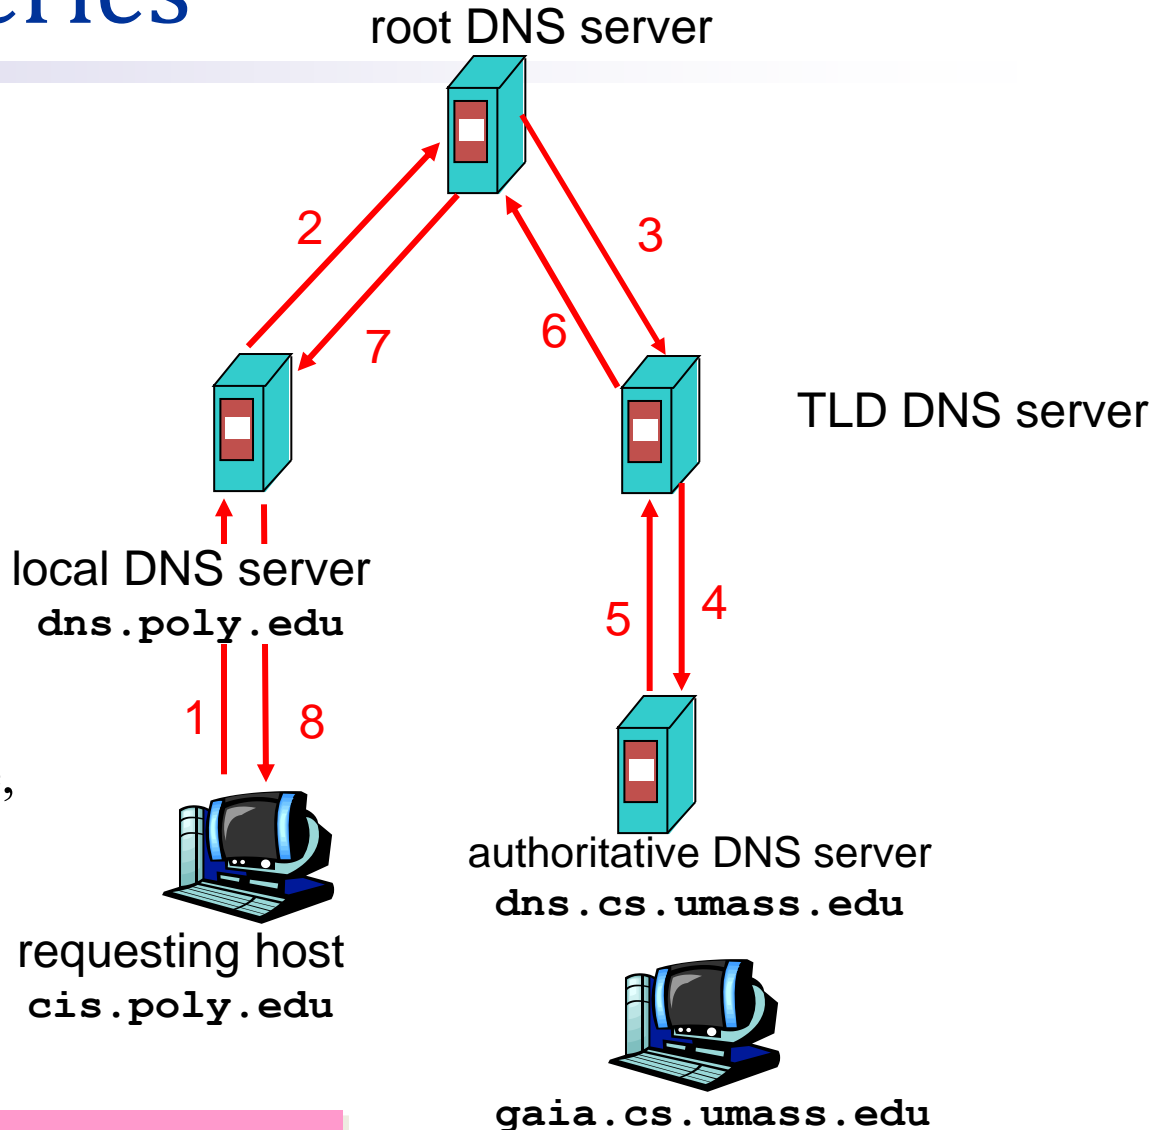
# Recursive queries

## recursive query:

- puts burden of name resolution on contacted name server
- heavy load?

## iterated query:

- contacted server replies with name of server to contact
- “I don’t know this name, but ask this server”



Which is a better design choice?

# DNS: caching and updating records

- Once (any) name server learns mapping, it *caches* mapping
  - cache entries timeout (disappear) after some time
  - TLD servers typically cached in local name servers
    - Thus root name servers not often visited
- Update/notify mechanisms under design by IETF
  - RFC 2136
  - <http://www.ietf.org/html.charters/dnsind-charter.html>

# DNS records

DNS: distributed db storing resource records (RR)

RR format: (**name**, **value**, **type**, **ttl**)

- Type=A
  - ❖ **name** is hostname
  - ❖ **value** is IP address
- Type=NS
  - **name** is domain (e.g. foo.com)
  - **value** is hostname of authoritative name server for this domain (e.g. dns.foo.com)
- Type=CNAME
  - ❖ **name** is alias name for some “canonical” (the real) name  
www.ibm.com is really servereast.backup2.ibm.com
  - ❖ **value** is canonical name
- Type=MX
  - ❖ **value** is name of mailserver associated with **name** (e.g. foo.com, mail.bar.foo.com, MX)

# DNS protocol, messages

DNS protocol: *query* and *reply* messages, both with same *message format*

msg header

❑ **identification**: 16 bit # for query, reply to query uses same #

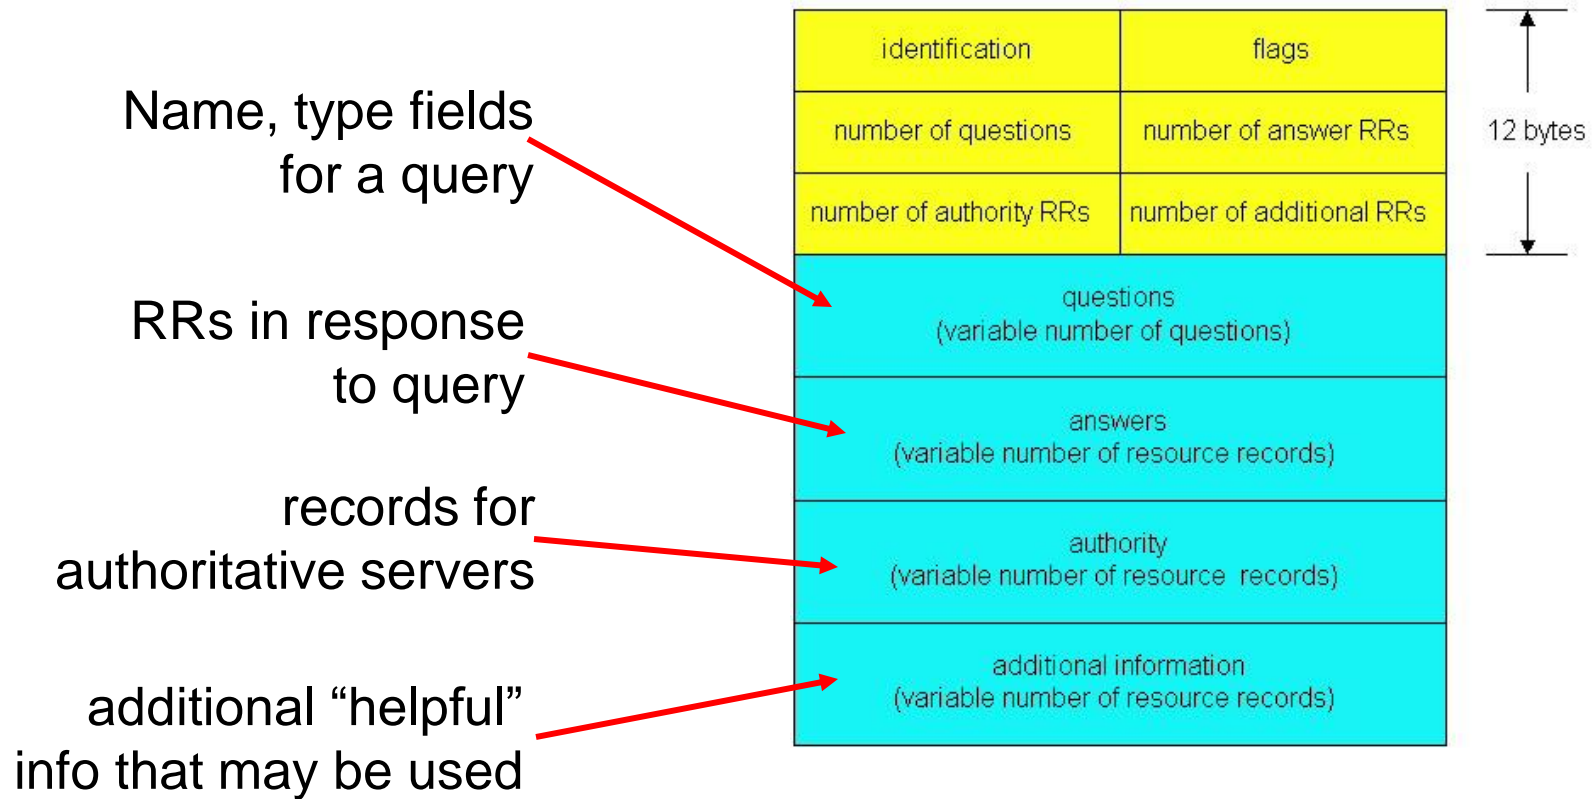
❑ **flags**:

- ❖ query or reply
- ❖ recursion desired
- ❖ recursion available
- ❖ reply is authoritative

identification	flags
number of questions	number of answer RRs
number of authority RRs	number of additional RRs
questions (variable number of questions)	
answers (variable number of resource records)	
authority (variable number of resource records)	
additional information (variable number of resource records)	



# DNS protocol, messages





# Inserting records into DNS

- Example: just created startup “Network Utopia”
- Register name networkutopia.com at a registrar (e.g., Network Solutions)
  - Need to provide registrar with names and IP addresses of your authoritative name server (primary and secondary)
  - Registrar inserts two RRs into the com TLD server:

```
(networkutopia.com, dns1.networkutopia.com, NS)  
(dns1.networkutopia.com, 212.212.212.1, A)
```

- Put in authoritative server Type A record for www.networkuptopia.com and Type MX record for mail.networkutopia.com
- How do people get the IP address of your Web site?

---

Questions till now ?

# Chapter 2: Application layer

- 2.1 Principles of network applications
  - app architectures
  - app requirements
- 2.2 Web and HTTP
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# P2P file sharing

## Example

- Alice runs P2P client application on her notebook computer
- Intermittently connects to Internet; gets new IP address for each connection
- Asks for “Hey-Jude.mp3”
- Application displays other peers that have copy of Hey Jude.
- Alice chooses one of the peers, Bob.
- File is copied from Bob’s PC to Alice’s notebook: HTTP
- While Alice downloads, other users uploading from Alice.
- Alice’s peer is both a Web client and a transient Web server.

All peers are servers = highly scalable!

# P2P: centralized directory

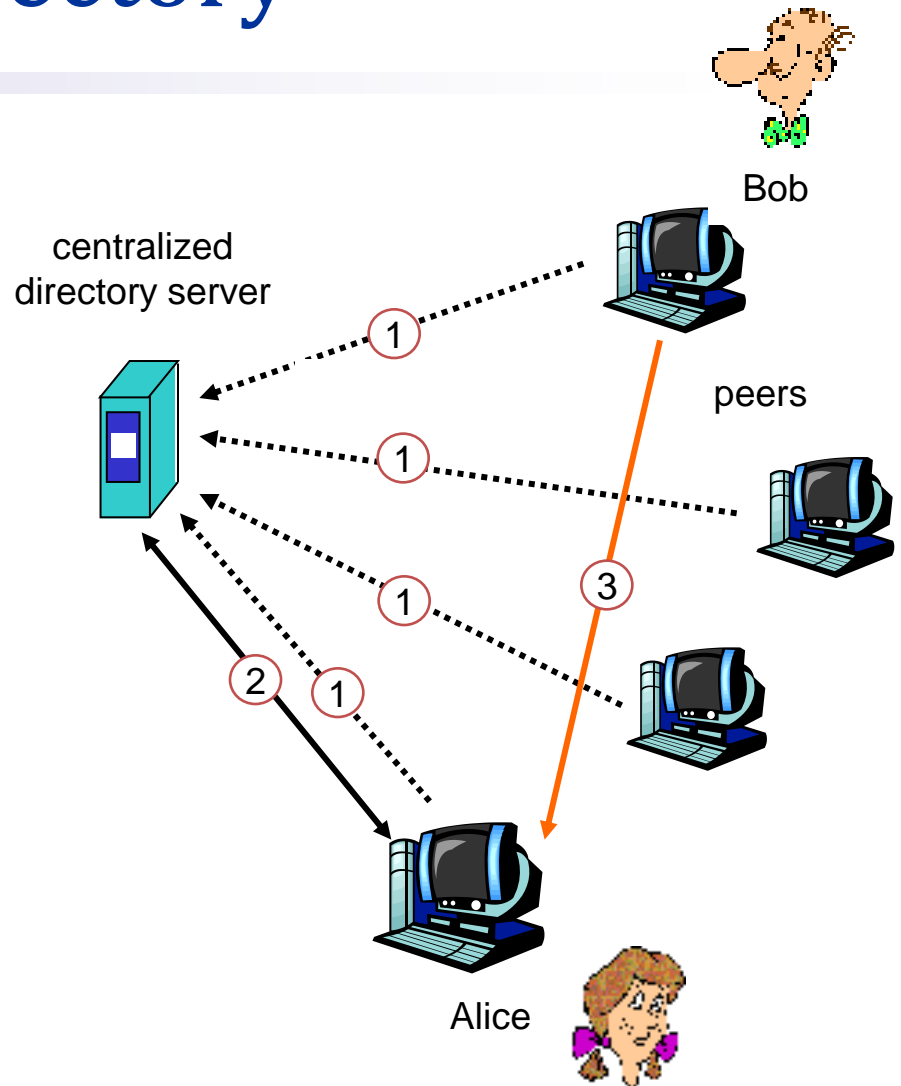
original “Napster” design

1) when peer connects, it informs central server:

- IP address
- content

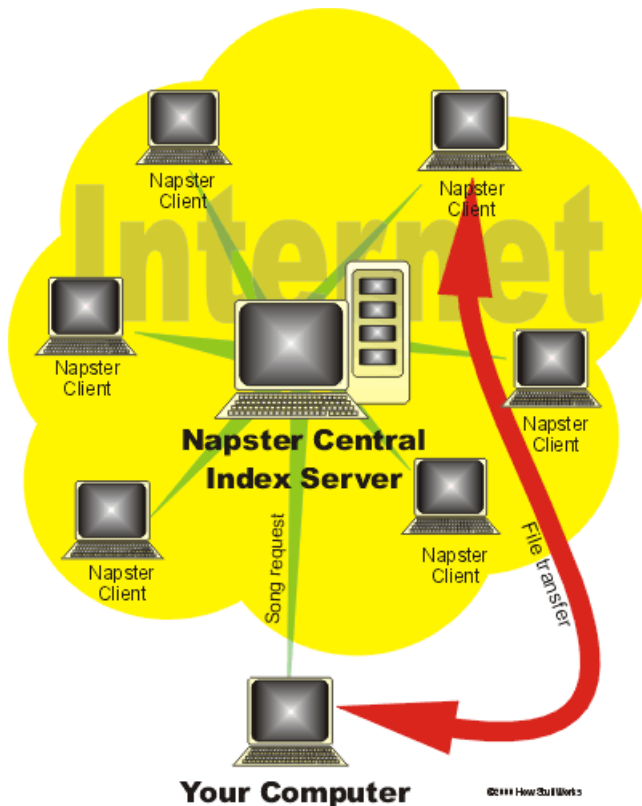
2) Alice queries for “Hey Jude”

3) Alice requests file from Bob



# P2P: problems with centralized directory

- Single point of failure
- Performance bottleneck
- Copyright infringement



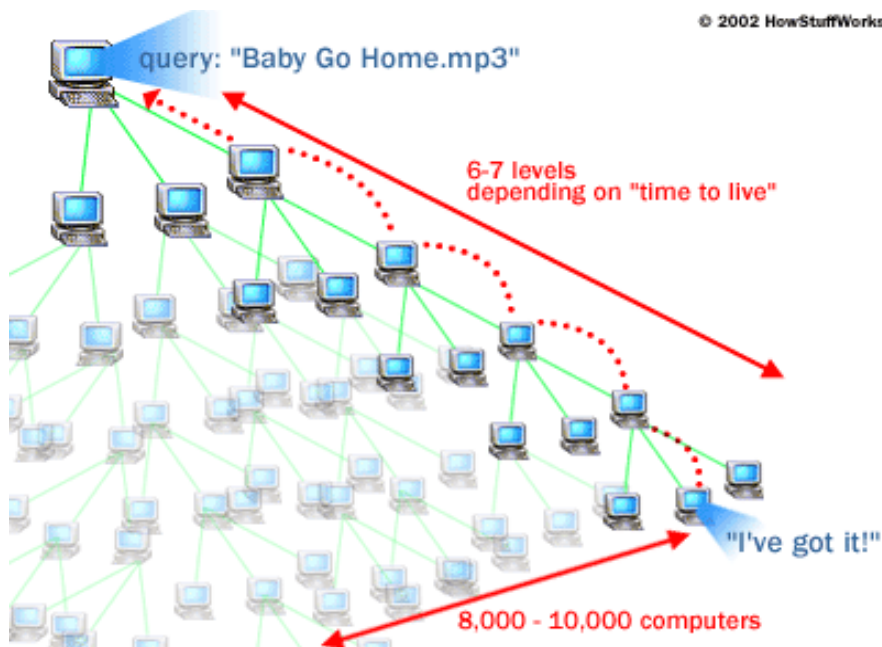
file transfer is decentralized,  
but locating content is highly  
centralized

# Query flooding: Gnutella

- fully distributed
  - no central server
- public domain protocol
- many Gnutella clients implementing protocol

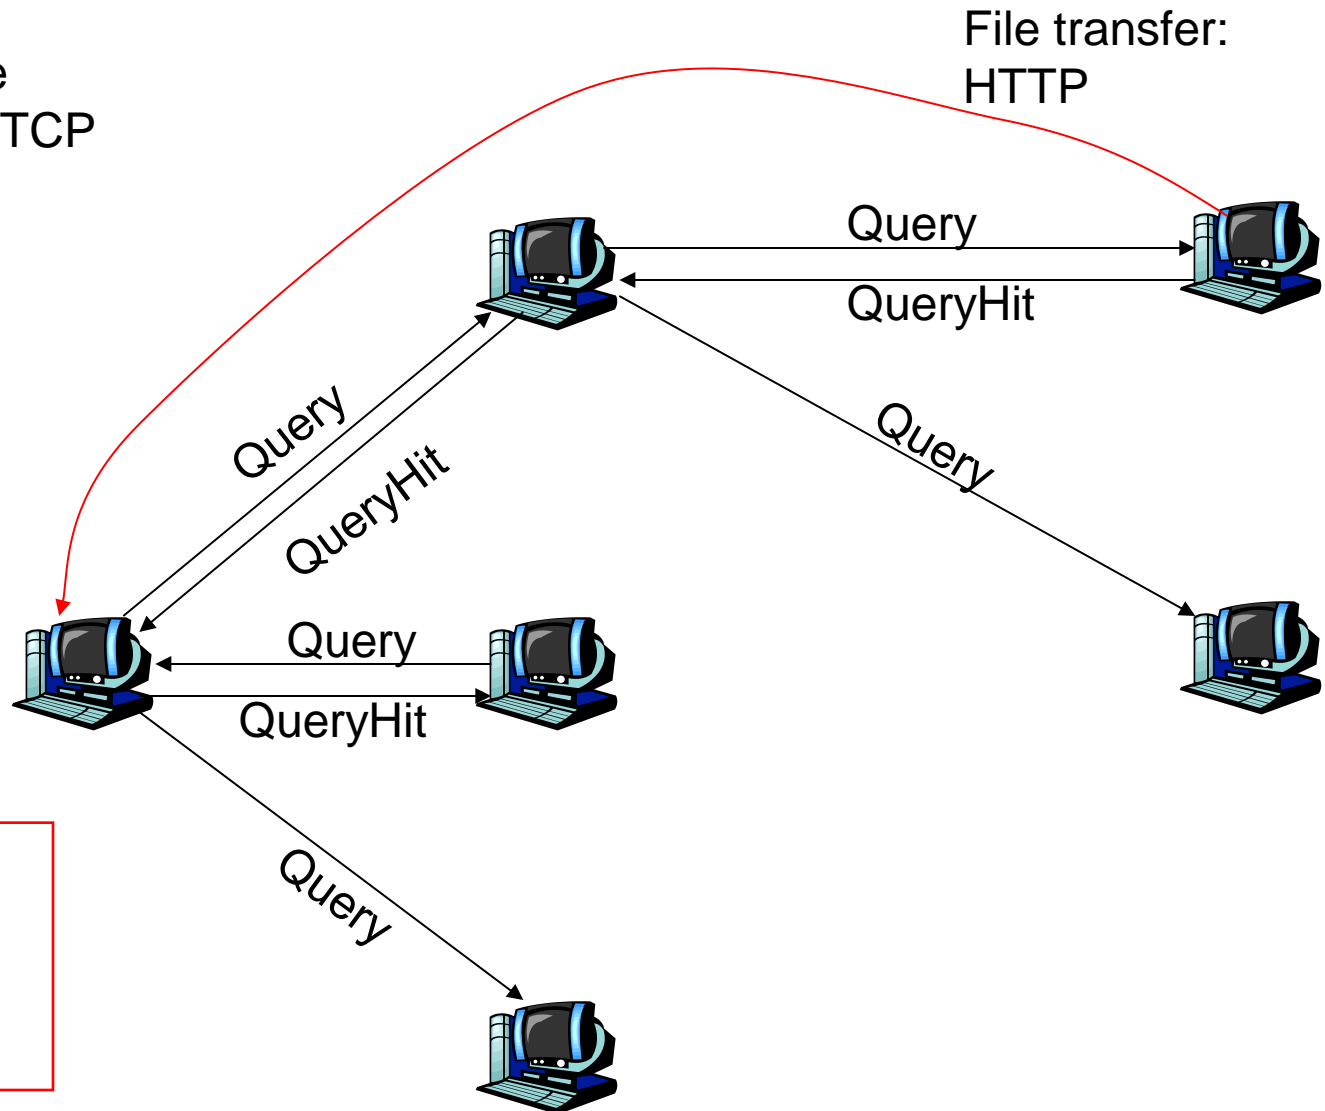
## overlay network: graph

- edge between peer X and Y if there's a TCP connection
- all active peers and edges is overlay net
- Edge is not a physical link
- Given peer will typically be connected with  $< 10$  overlay neighbors



# Gnutella: protocol

- ❑ Query message sent over existing TCP connections
- ❑ peers forward Query message
- ❑ QueryHit sent over reverse path



Scalability:  
limited scope  
flooding



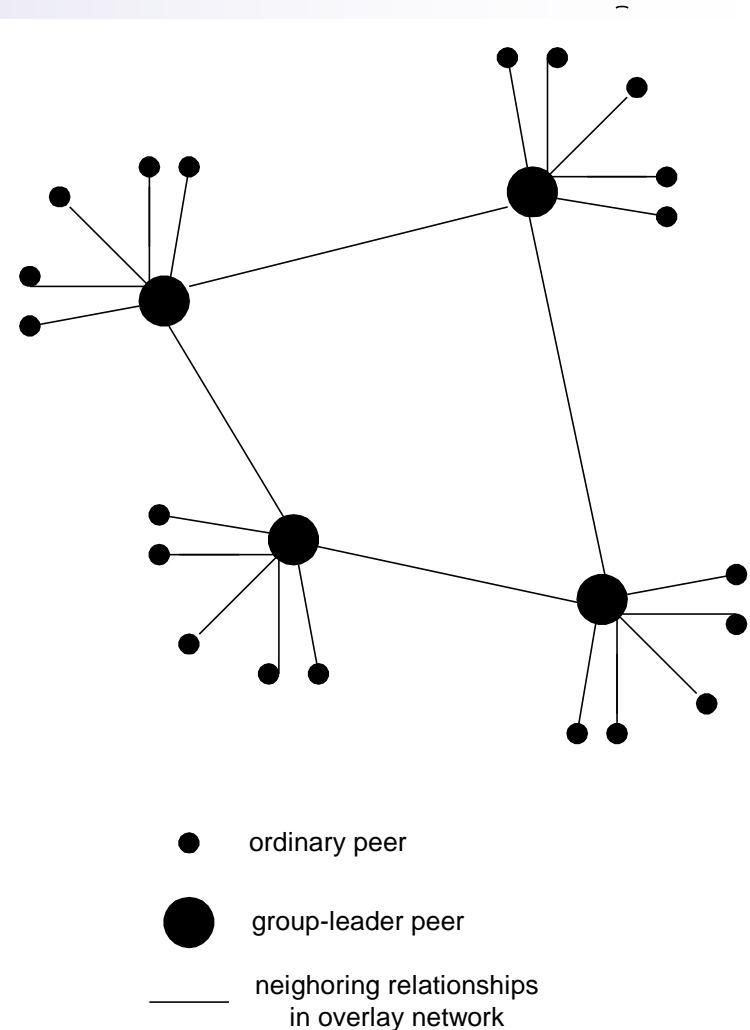
# Gnutella: Peer joining

1. Joining peer X must find some other peer in Gnutella network: use list of candidate peers
2. X sequentially attempts to make TCP with peers on list until connection setup with Y
3. X sends Ping message to Y; Y forwards Ping message.
4. All peers receiving Ping message respond with Pong message
5. X receives many Pong messages. It can then setup additional TCP connections

What happens when peer leaves: find out as an exercise!

# Exploiting heterogeneity: KaZaA

- Each peer is either a group leader or assigned to a group leader.
  - TCP connection between peer and its group leader.
  - TCP connections between some pairs of group leaders.
- Group leader tracks the content in all its children.



# KaZaA: Querying

- Each file has a hash and a descriptor
- Client sends keyword query to its group leader
- Group leader responds with matches:
  - For each match: metadata, hash, IP address
- If group leader forwards query to other group leaders, they respond with matches
- Client then selects files for downloading
  - HTTP requests using hash as identifier sent to peers holding desired file

# KaZaA tricks

- Limitations on simultaneous uploads
- Request queuing
- Incentive priorities
- Parallel downloading

For more info:

□ J. Liang, R. Kumar, K. Ross, “Understanding KaZaA,”  
(available via [cis.poly.edu/~ross](http://cis.poly.edu/~ross))

# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Socket programming

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

## Socket API

- introduced in BSD 4.1 UNIX, 1981
- explicitly created, used, released by apps
- client/server paradigm
- two types of transport service via socket API:
  - unreliable datagram (UDP)
  - reliable, byte stream-oriented (TCP)

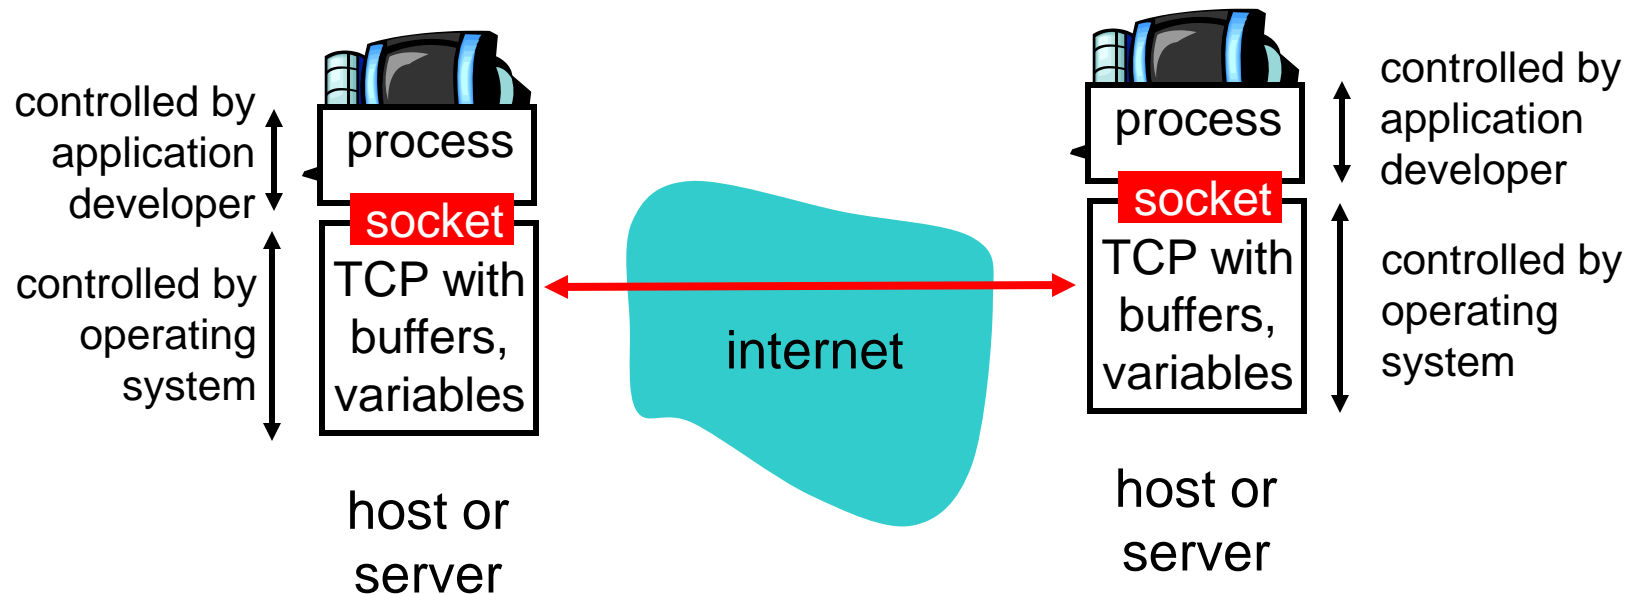
socket

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 (UDP 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*



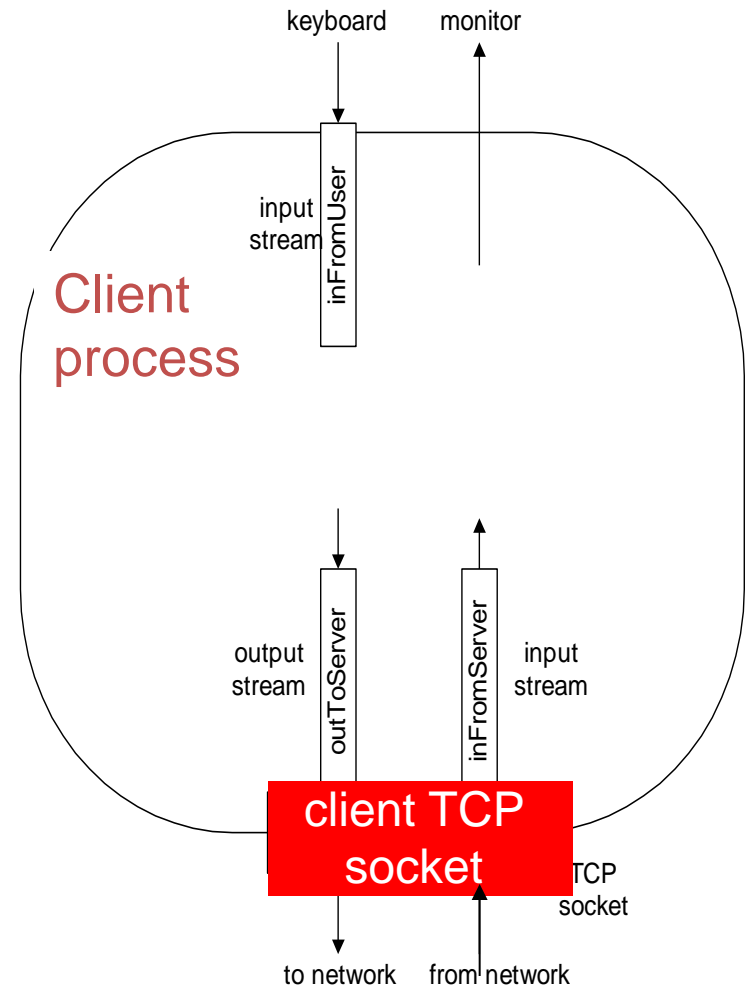
# Stream jargon

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

# Socket programming with TCP

## Example client-server app:

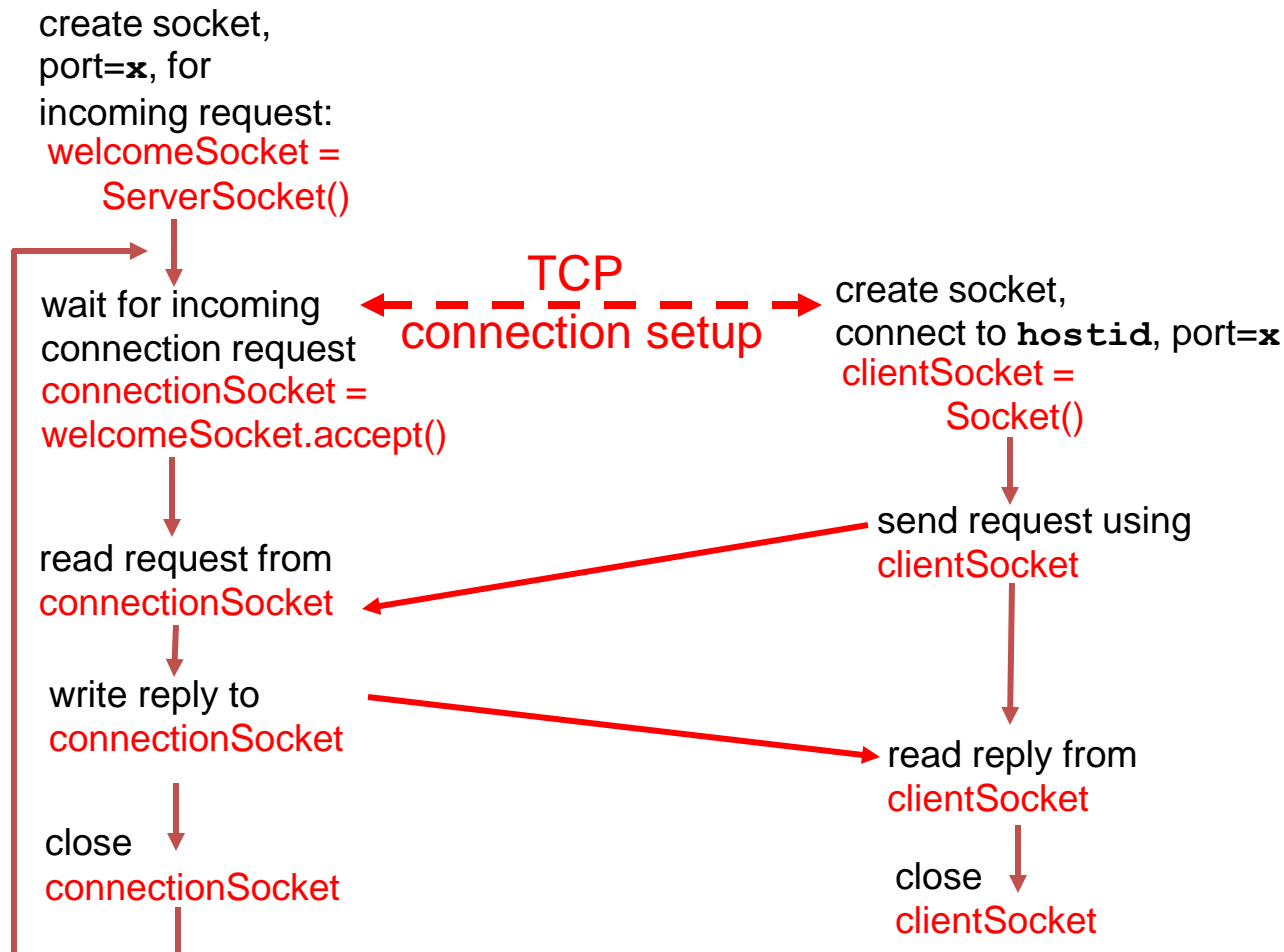
- 1) client reads line from standard input (**inFromUser** stream), sends to server via socket (**outToServer** stream)
- 2) server reads line from socket
- 3) server converts line to uppercase, sends back to client
- 4) client reads, prints modified line from socket (**inFromServer** stream)



# Client/server socket interaction: TCP

Server (running on `hostid`)

Client



# Example: Java client (TCP)

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

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

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

Create  
input stream

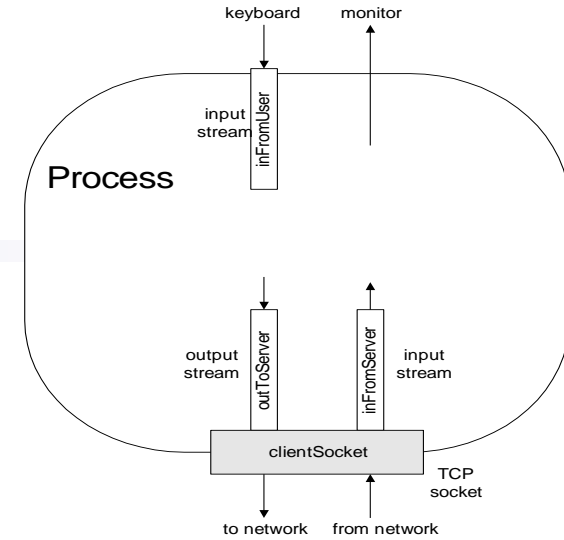
```
        BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));
```

Create  
client socket,  
connect to server

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

Create  
input stream  
attached to socket

```
        BufferedReader inFromServer =
            new BufferedReader(new
                InputStreamReader(clientSocket.getInputStream()));
```



# Example: Java client (TCP), cont.

Create  
output stream  
attached to socket

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

Send line  
to server

```
sentence = inFromUser.readLine();  
outToServer.writeBytes(sentence + '\n');
```

Read line  
from server

```
modifiedSentence = inFromServer.readLine();
```

```
System.out.println("FROM SERVER: " + modifiedSentence);
```

```
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);
```

Wait, on welcoming  
socket for contact  
by client

```
        while(true) {
```

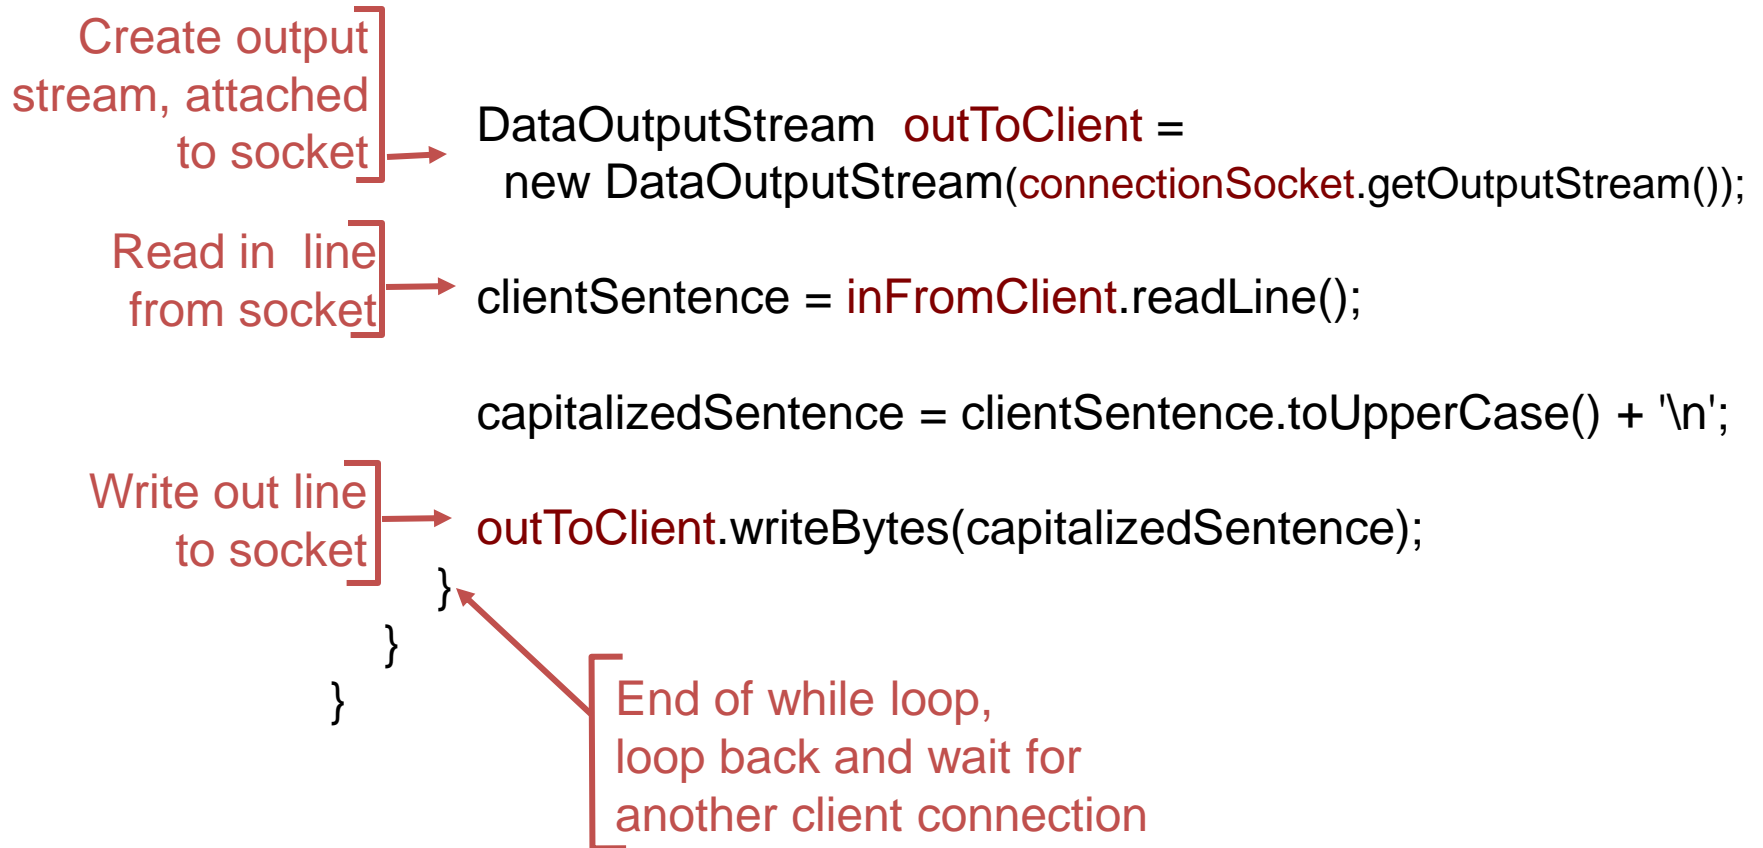
```
            Socket connectionSocket = welcomeSocket.accept();
```

Create input  
stream, attached  
to socket

```
            BufferedReader inFromClient =
```

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

# Example: Java server (TCP), cont



# 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 file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server



# 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

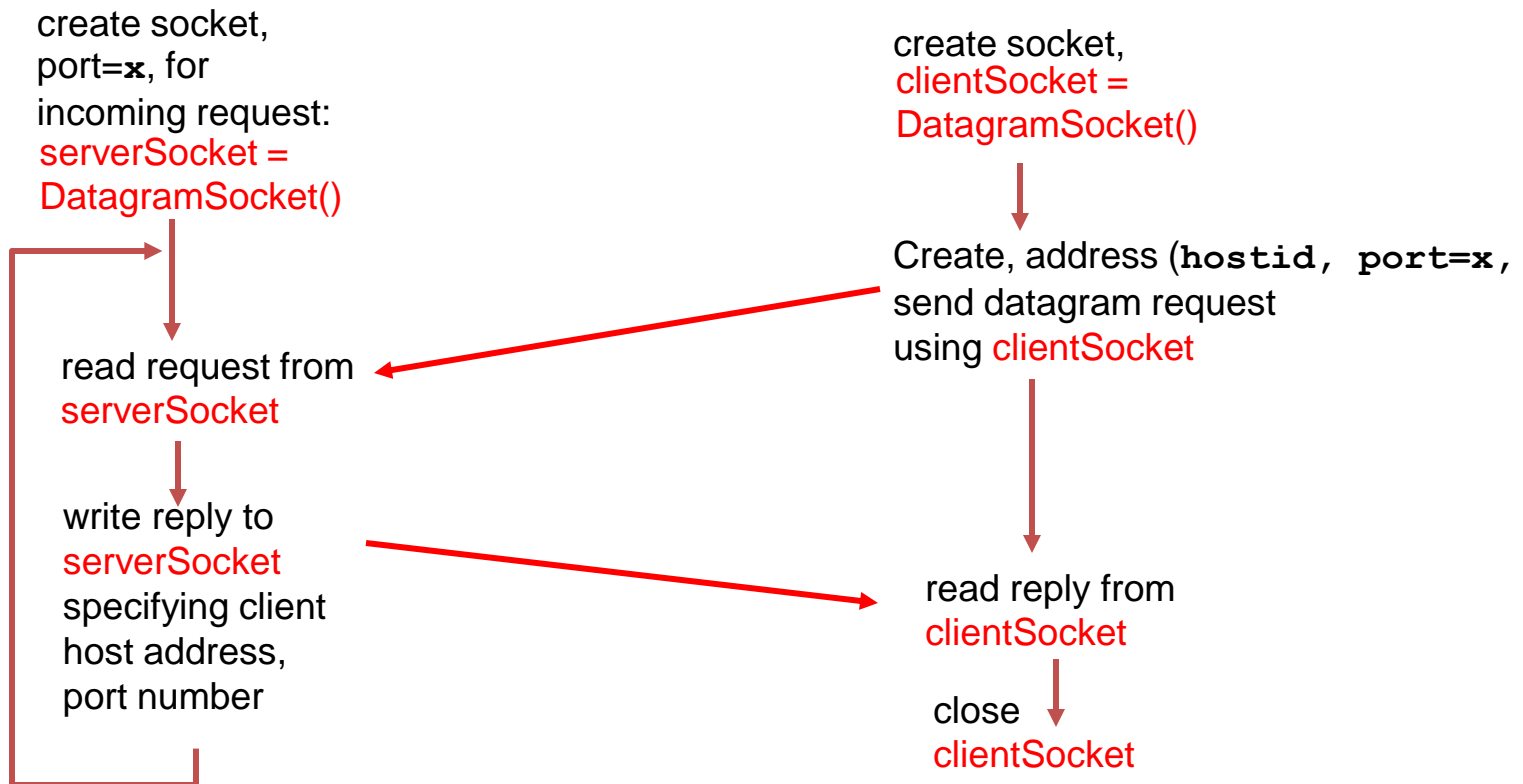
application viewpoint

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

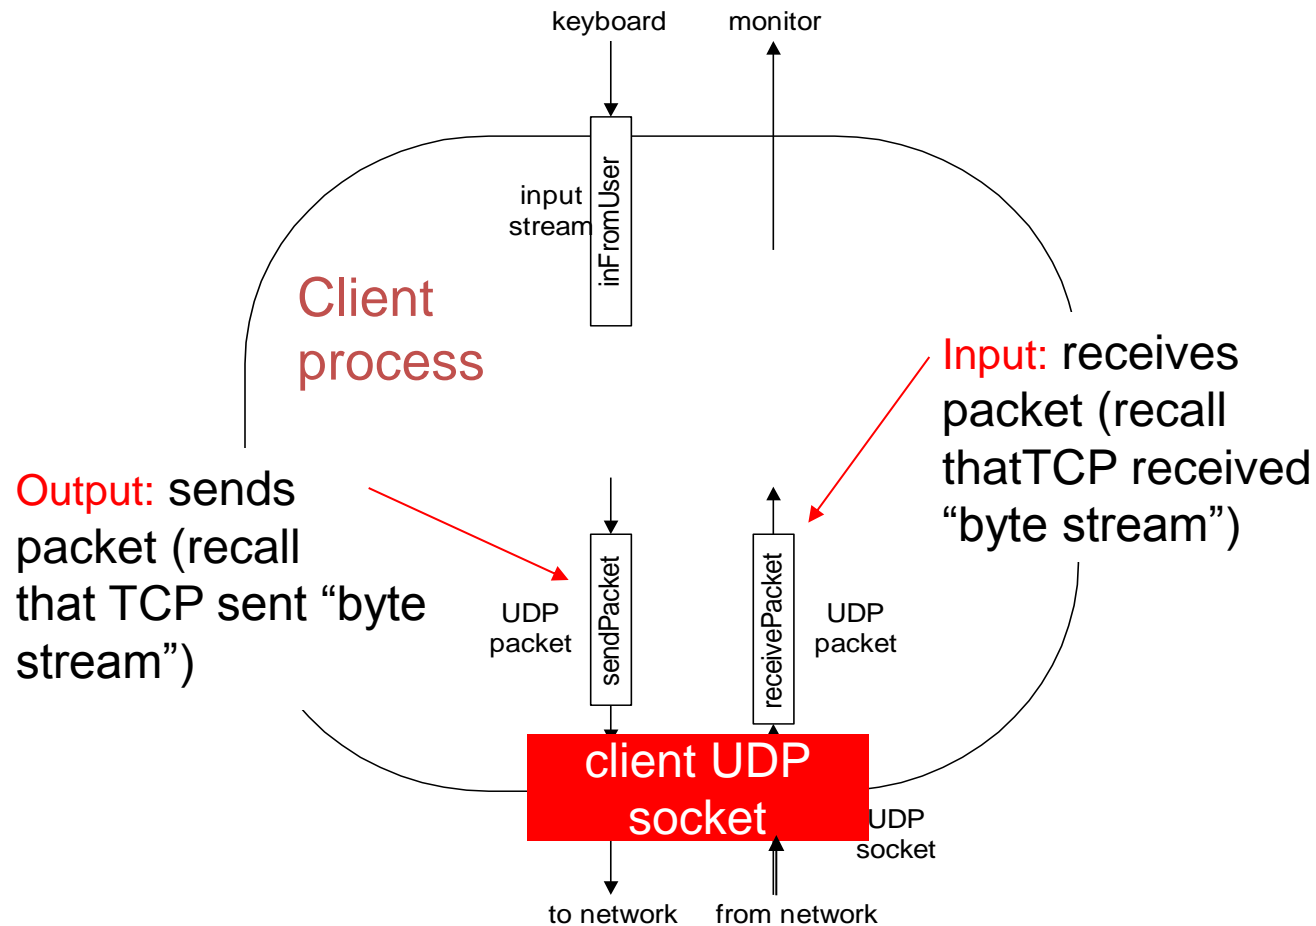
# Client/server socket interaction: UDP

Server (running on `hostid`)

Client



# 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 =
```

Create  
client socket

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

```
        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();
```

Which address is this?

```
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

# Sockets and Ports

What is the difference between sockets and ports?

Sockets are physical telephones

Ports are extension numbers

IP address is the phone number

1. If you are interested in a quick Java tutorial targeted specifically at socket programming, check out "Socket Programming in Java: a tutorial," by Q. Mahmoud, Javaworld, Dec. 1996, [http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets\\_p.html](http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets_p.html)
2. For implementing Multi-threaded Client/Server Applications you may read through the Java example source code at <http://www.ase.md/~aursu/ClientServerThreads.html>



# Chapter 2: Application layer

- 2.1 Principles of network applications
  - app architectures
  - app requirements
- 2.2 Web and HTTP
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Building a simple Web server

- handles one HTTP request
- accepts the request
- parses header
- obtains requested file from server's file system
- creates HTTP response message:
  - header lines + file
- sends response to client
- after creating server, you can request file using a browser (e.g., IE explorer)
- see text for details

# 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

### □ 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
- control vs. data msgs
    - centralized vs. decentralized
  - stateless vs. stateful
  - reliable vs. unreliable msg transfer
  - “complexity at network edge”

# Questions?

---