

## Chapter 2: Application Layer

### Our goals:

- ρ conceptual, implementation aspects of network application protocols
  - μ transport-layer service models
  - μ client-server paradigm
  - μ peer-to-peer paradigm

- ρ learn about protocols by examining popular application-level protocols
  - μ HTTP
  - μ FTP
  - μ SMTP / POP3 / IMAP
  - μ DNS
- ρ programming network applications
  - μ socket API

Application Layer 1

## Network applications: some jargon

- Process:** program running within a host.
  - ρ within same host, two processes communicate using **interprocess communication** (defined by OS).
  - ρ processes running in different hosts communicate with an **application-layer protocol**

- user agent:** interfaces with user "above" and network "below".
  - ρ implements user interface & application-level protocol
    - μ Web: browser
    - μ E-mail: mail reader
    - μ streaming audio/video: media player

Application Layer 2


## Applications and application-layer protocols

### Application: communicating, distributed processes

- μ e.g., e-mail, Web, P2P file sharing, instant messaging
- μ running in end systems (hosts)
- μ exchange messages to implement application

### Application-layer protocols

- μ one "piece" of an app
- μ define messages exchanged by apps and actions taken
- μ use communication services provided by lower layer protocols (TCP, UDP)



Application Layer 3

## App-layer protocol defines

- ρ Types of messages exchanged, eg, request & response messages
- ρ Syntax of message types: what fields in messages & how fields are delineated
- ρ Semantics of the fields, ie, meaning of information in fields
- ρ Rules for when and how processes send & respond to messages

### Public-domain protocols:

- ρ defined in RFCs
- ρ allows for interoperability
- ρ eg, HTTP, SMTP

### Proprietary protocols:

- ρ eg, KaZaA

Application Layer 4

## Client-server paradigm

Typical network app has two pieces: *client* and *server*


### Client:

- ρ initiates contact with server ("speaks first")
- ρ typically requests service from server,

- ρ Web: client implemented in browser; e-mail: in mail reader

### Server:

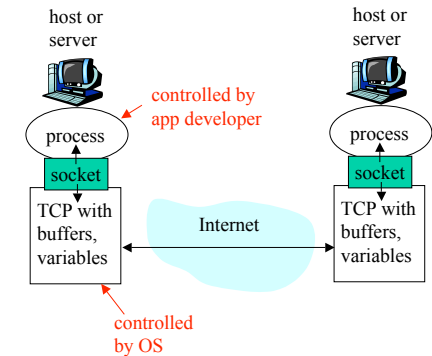
- ρ provides requested service to client
- ρ e.g., Web server sends requested Web page, mail server delivers e-mail



Application Layer 5

## Processes communicating across network

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



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

Application Layer 6

## Addressing processes:

- ρ For a process to receive messages, it must have an identifier
- ρ Every host has a unique 32-bit IP address
- ρ Q: does the IP address of the host on which the process runs suffice for identifying the process?
- ρ Answer: No, many processes can be running on same host
- ρ Identifier includes both the IP address and port numbers associated with the process on the host.
- ρ Example port numbers:
  - μ HTTP server: 80
  - μ Mail server: 25
- ρ More on this later

Application Layer 7

## What transport service does an app need?

### Data loss

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

### Timing

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

### Bandwidth

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

Application Layer 8

## Transport service requirements of common apps

Application	Data loss	Bandwidth	Time Sensitive
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:

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

### UDP service:

- ρ unreliable data transfer between sending and receiving process
  - ρ 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

Application	Application layer protocol	Underlying transport protocol
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., Dialpad)	typically UDP

## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ **2.2 Web and HTTP**
- ρ 2.3 FTP
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ 2.5 DNS
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing

# Web and HTTP

## First some jargon

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


www.someschool.edu / someDept/pic.gif

host name                      path name

# HTTP overview

## HTTP: hypertext transfer protocol

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



# HTTP overview (continued)

## Uses TCP:

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

## HTTP is "stateless"

- ρ server maintains no information about past client requests

aside

Protocols that maintain "state" are complex!

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

# HTTP connections

## Nonpersistent HTTP

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

## Persistent HTTP


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

# Response time modeling

Definition of RRT: time to send a small packet to travel from client to server and back.

**Response time:**

- ρ one RTT to initiate TCP connection
- ρ one RTT for HTTP request and first few bytes of HTTP response to return
- ρ file transmission time
- total = 2RTT+transmit time**



# Persistent HTTP

**Nonpersistent HTTP issues:**

- ρ requires 2 RTTs per object
- ρ OS must work and allocate host resources for each TCP connection
- ρ but browsers often open parallel TCP connections to fetch referenced objects

**Persistent HTTP**

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

**Persistent without pipelining:**

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

**Persistent with pipelining:**

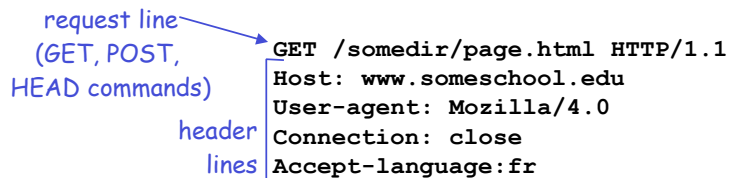
- ρ default in HTTP/1.1
- ρ client sends requests as soon as it encounters a referenced object
- ρ as little as one RTT for all the referenced objects

# HTTP request message

ρ two types of HTTP messages: *request, response*

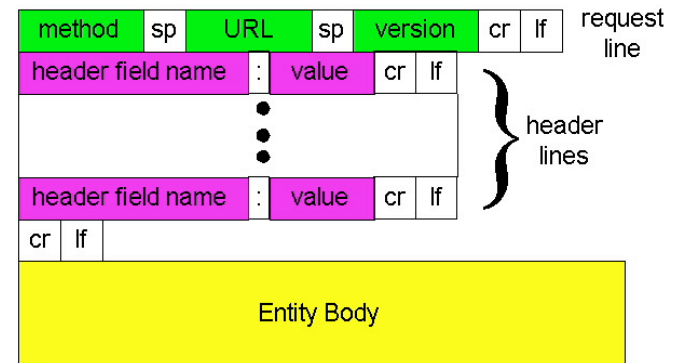
**HTTP request message:**

μ ASCII (human-readable format)



Carriage return, line feed (extra carriage return, line feed) indicates end of message

# HTTP request message: general format



## Uploading form input

### Post method:

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

### URL method:

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

www.somesite.com/animalsearch?monkeys&banana

## Method types

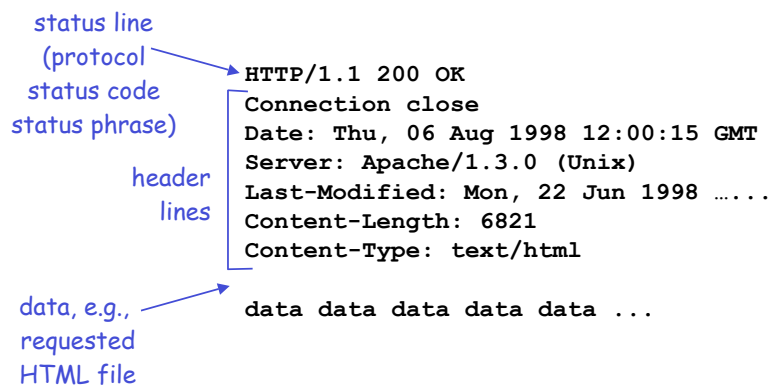
### HTTP/1.0

- ρ GET
- ρ POST
  - μ Like GET but user provides info (ie key words)
- ρ HEAD
  - μ asks server to leave requested object out of response (for debugging)

### HTTP/1.1

- ρ GET, POST, HEAD
- ρ PUT
  - μ uploads file in entity body to path specified in URL field
- ρ DELETE
  - μ 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**

- μ request succeeded, requested object later in this message

### **301 Moved Permanently**

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

### **400 Bad Request**

- μ request message not understood by server

### **404 Not Found**

- μ 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 www.eurecom.fr 80` Opens TCP connection to port 80 (default HTTP server port) at www.eurecom.fr. Anything typed in sent to port 80 at www.eurecom.fr

### 2. Type in a GET HTTP request:

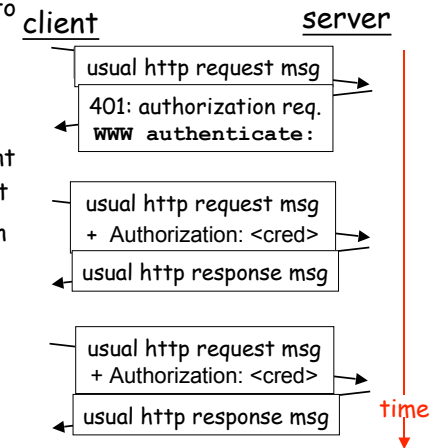
`GET /~ross/index.html HTTP/1.0` By typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

### 3. Look at response message sent by HTTP server!

## User-server interaction: authorization

**Authorization**: control access to server content

- μ authorization credentials: typically name, password
- μ **stateless**: client must present authorization in *each* request
  - μ **authorization**: header line in each request
  - μ if no **authorization**: header, server refuses access, sends `WWW authenticate:` header line in response



## Cookies: keeping "state"

Many major Web sites use cookies

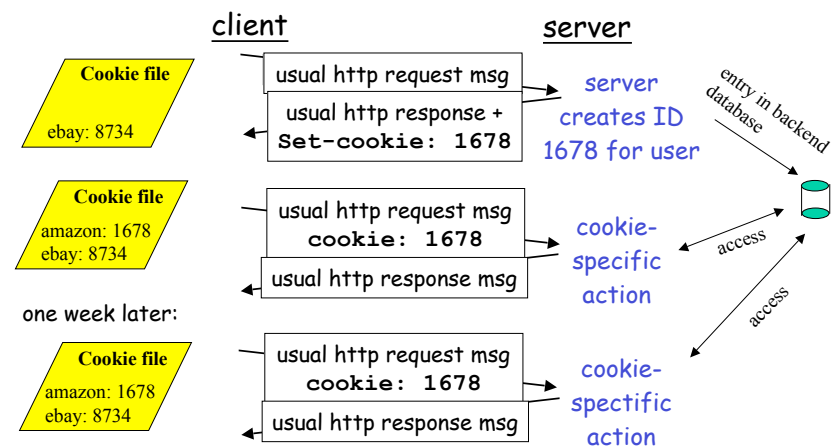
### Four components:

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

### Example:

- μ Susan access Internet always from same PC
- μ She visits a specific e-commerce site for first time
- μ 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:

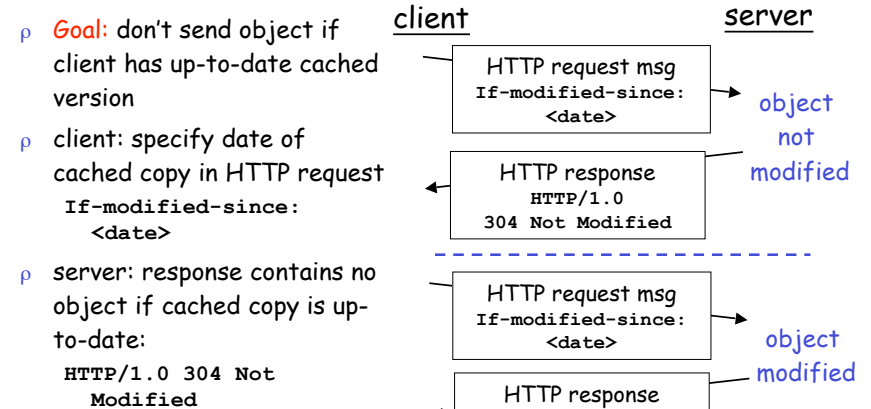
- ρ authorization
- ρ shopping carts
- ρ recommendations
- ρ user session state (Web e-mail)

aside

**Cookies and privacy:**

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


## Conditional GET: client-side caching



## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ 2.2 Web and HTTP
- ρ **2.3 FTP**
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ 2.5 DNS
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing


## FTP: the file transfer protocol



- ρ transfer file to/from remote host
- ρ client/server model
  - μ **client:** side that initiates transfer (either to/from remote)
  - μ **server:** remote host
- ρ ftp: RFC 959
- ρ 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 over control connection
- Client browses remote directory by sending commands over control connection.
- When server receives a command for a file transfer, the server opens a 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

Application Layer 33

## 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

Application Layer 34

## Chapter 2 outline


- 2.1 Principles of app layer protocols
  - clients and servers
  - app requirements
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
  - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 Socket programming with TCP
- 2.7 Socket programming with UDP
- 2.8 Building a Web server
- 2.9 Content distribution
  - Network Web caching
  - Content distribution networks
  - P2P file sharing

Application Layer 35

## 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




Application Layer 36

## Electronic Mail: mail servers

### Mail Servers

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



Application Layer 37


## Electronic Mail: SMTP [RFC 2821]

- ρ uses TCP to reliably transfer email message from client to server, port 25
- ρ direct transfer: sending server to receiving server
- ρ three phases of transfer
  - μ handshaking (greeting)
  - μ transfer of messages
  - μ closure
- ρ command/response interaction
  - μ commands: ASCII text
  - μ response: status code and phrase
- ρ messages must be in 7-bit ASCII

Application Layer 38

## Scenario: Alice sends message to Bob

- 1) Alice uses UA to compose message and "to" bob@someschool.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



Application Layer 39

## Sample SMTP interaction

```

S: 220 hamburger.edu
C: HELO crepes.fr
S: 250 Hello crepes.fr, pleased to meet you
C: MAIL FROM: <alice@crepes.fr>
S: 250 alice@crepes.fr... Sender ok
C: RCPT TO: <bob@hamburger.edu>
S: 250 bob@hamburger.edu ... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Do you like ketchup?
C: How about pickles?
C: .
S: 250 Message accepted for delivery
C: QUIT
S: 221 hamburger.edu closing connection
    
```

Application Layer 40

## Try SMTP interaction for yourself:

- ρ telnet servername 25
- ρ see 220 reply from server
- ρ enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands

above lets you send email without using email client (reader)

## SMTP: final words

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

### Comparison with HTTP:

- ρ HTTP: pull
- ρ SMTP: push
- ρ both have ASCII command/response interaction, status codes
- ρ HTTP: each object encapsulated in its own response msg
- ρ SMTP: multiple objects sent in multipart msg


## Mail message format

SMTP: protocol for exchanging email msgs

RFC 822: standard for text message format:

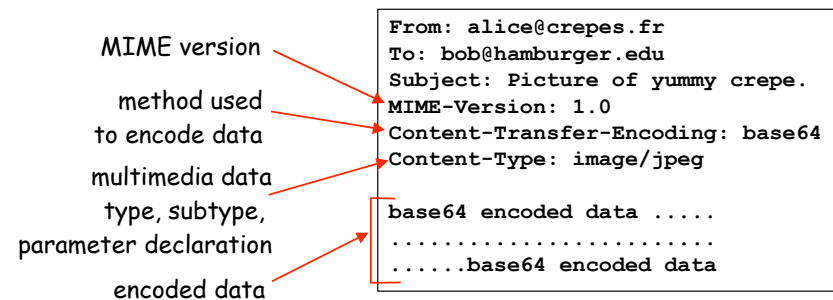
- ρ header lines, e.g.,
  - μ To:
  - μ From:
  - μ Subject:
- different from SMTP commands!*

- ρ body
  - μ the "message", ASCII characters only



## Message format: multimedia extensions

- ρ MIME: multimedia mail extension, RFC 2045, 2056
- ρ additional lines in msg header declare MIME content type



## MIME types

Content-Type: type/subtype; parameters

### Text

- ρ example subtypes: plain, html

### Video

- ρ example subtypes: mpeg, quicktime

### Image

- ρ example subtypes: jpeg, gif

### Application

- ρ other data that must be processed by reader before "viewable"
- ρ example subtypes: msword, octet-stream

### Audio

- ρ example subtypes: basic (8-bit mu-law encoded), 32kadpcm (32 kbps coding)

Application Layer 45

## Multipart Type


```
From: alice@crepes.fr
To: bob@hamburger.edu
Subject: Picture of yummy crepe.
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary=StartOfNextPart
```

```
--StartOfNextPart
Dear Bob, Please find a picture of a crepe.
--StartOfNextPart
Content-Transfer-Encoding: base64
Content-Type: image/jpeg
base64 encoded data .....
.....base64 encoded data
--StartOfNextPart
Do you want the recipe?
```



Application Layer 46

## Mail access protocols



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

Application Layer 47

## POP3 protocol

### authorization phase

- ρ client commands:

```
μ user: declare username
μ pass: password
```

- ρ server responses

```
μ +OK
μ -ERR
```

### transaction phase, client:

- ρ list: list message numbers
- ρ retr: retrieve message by number
- ρ dele: delete
- ρ quit

```
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
```

Application Layer 48

## POP3 (more) and IMAP

### More about POP3

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

### IMAP

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

Application Layer 49

## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ 2.2 Web and HTTP
- ρ 2.3 FTP
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ **2.5 DNS**
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing

Application Layer 50

## DNS: Domain Name System

### People: many identifiers:

- μ SSN, name, passport #

### Internet hosts, routers:

- μ IP address (32 bit) - used for addressing datagrams
- μ "name", e.g., gaia.cs.umass.edu - used by humans

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

### Domain Name System:

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

Application Layer 51

## DNS name servers

### Why not centralize DNS?

- ρ single point of failure
- ρ traffic volume
- ρ distant centralized database
- ρ maintenance

doesn't *scale!*

- ρ no server has all name-to-IP address mappings

### local name servers:

- μ each ISP, company has *local (default) name server*
- μ host DNS query first goes to local name server

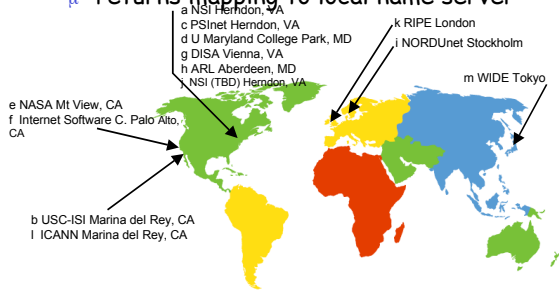
### authoritative name server:

- μ for a host: stores that host's IP address, name
- μ can perform name/address translation for that host's name

Application Layer 52

# DNS: Root name servers

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




13 root name servers worldwide

# Simple DNS example

host `surf.eurecom.fr` wants IP address of `gaia.cs.umass.edu`


1. contacts its local DNS server, `dns.eurecom.fr`
2. `dns.eurecom.fr` contacts root name server, if necessary
3. root name server contacts authoritative name server, `dns.umass.edu`, if necessary



# DNS example

Root name server:

- ρ may not know authoritative name server
- ρ may know *intermediate name server*: who to contact to find authoritative name server




# DNS: iterated queries

recursive query:

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

iterated query:

- ρ contacted server replies with name of server to contact
- ρ "I don't know this name, but ask this server"



## DNS: caching and updating records

- ρ once (any) name server learns mapping, it  *caches*  mapping
  - μ cache entries timeout (disappear) after some time
- ρ update/notify mechanisms under design by IETF
  - μ RFC 2136
  - μ <http://www.ietf.org/html.charters/dnsind-charter.html>

## DNS records

DNS: distributed db storing resource records (RR)

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

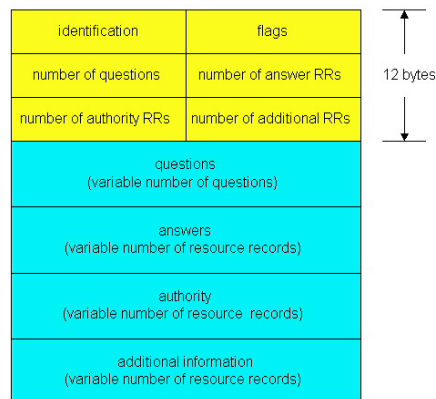
- ρ Type=A
  - μ name is hostname
  - μ value is IP address
- ρ Type=NS
  - μ name is domain (e.g. foo.com)
  - μ value is IP address of authoritative name server for this domain
- ρ Type=CNAME
  - μ name is alias name for some "canonical" (the real) name  
www.ibm.com is really servereast.backup2.ibm.com
- ρ Type=MX
  - μ value is canonical name
  - μ value is name of mailserver associated with name

## DNS protocol, messages

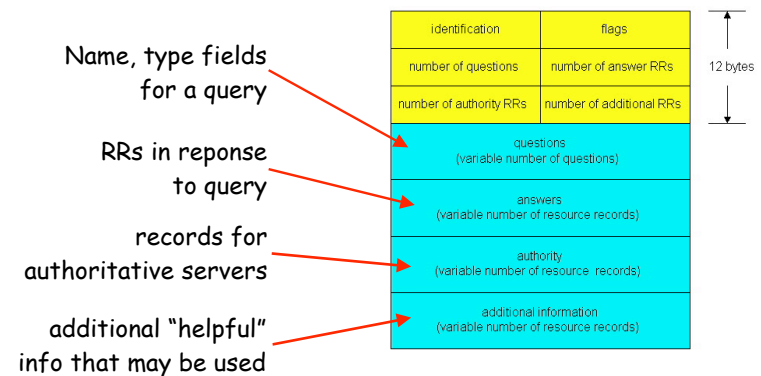
DNS protocol :  *query*  and  *reply*  messages, both with same  *message format*

msg header

- ρ **identification**: 16 bit # for query, reply to query uses same #
- ρ **flags**:
  - μ query or reply
  - μ recursion desired
  - μ recursion available
  - μ reply is authoritative



## DNS protocol, messages



## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ 2.2 Web and HTTP
- ρ 2.3 FTP
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ 2.5 DNS
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing

## Socket programming

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

### Socket API


- ρ introduced in BSD4.1 UNIX, 1981
- ρ explicitly created, used, released by apps
- ρ client/server paradigm
- ρ two types of transport service via socket API:
  - μ unreliable datagram
  - μ 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 (TCP or UDP)

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



## Socket programming *with TCP*

### Client must contact server

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

- ρ When contacted by client, **server TCP creates new socket** for server process to communicate with client
  - μ allows server to talk with multiple clients
  - μ source port numbers used to distinguish clients

### Client contacts server by:

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

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

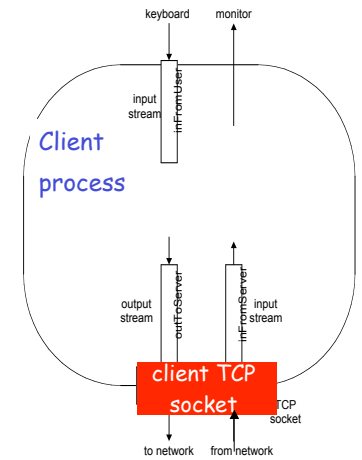
## Stream jargon

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

## Socket programming with TCP

### Example client-server app:

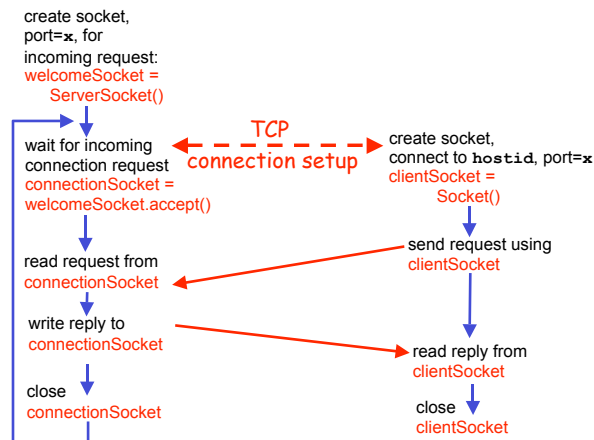
- 1) client reads line from standard input (`inFromUser` stream), sends to server via socket (`outToServer` stream)
- 2) server reads line from socket
- 3) server converts line to uppercase, sends back to client
- 4) client reads, prints modified line from socket (`inFromServer` stream)



## 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 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()));

    Send line
to server ] outToServer.writeBytes(sentence + '\n');

    Read line
from server ] modifiedSentence = inFromServer.readLine();

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

                    clientSocket.close();

                }
            }
    
```

Application Layer 69

## 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()));
    
```

Application Layer 70

## 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
    
```

Application Layer 71

## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ 2.2 Web and HTTP
- ρ 2.3 FTP
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ 2.5 DNS
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing

Application Layer 72

# 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

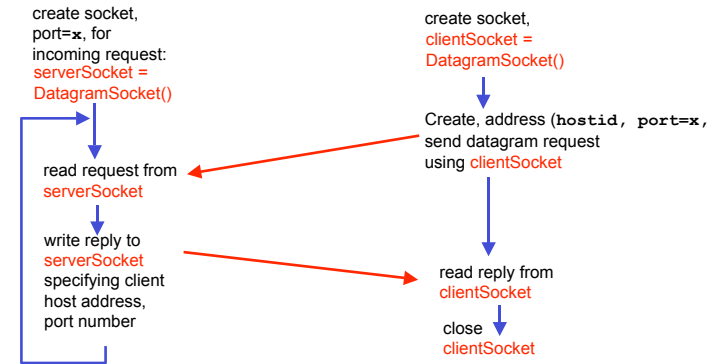
UDP: transmitted data may be received out of order, or lost

**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));
        Translate hostname to IP address using DNS -> DatagramSocket clientSocket = new DatagramSocket();
        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();
}
    
```

Application Layer 77

## 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);
}
}
    
```

Application Layer 78

## 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
    
```

Application Layer 79

## Building a simple Web server

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

Application Layer 80

## Socket programming: references

### C-language tutorial (audio/slides):

- ρ "Unix Network Programming" (J. Kurose),  
<http://manic.cs.umass.edu/~amldemo/courseware/intro>.

### Java-tutorials:

- ρ "All About Sockets" (Sun tutorial),  
<http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets.html>
- ρ "Socket Programming in Java: a tutorial,"  
<http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets.html>


## Chapter 2 outline

- ρ 2.1 Principles of app layer protocols
  - μ clients and servers
  - μ app requirements
- ρ 2.2 Web and HTTP
- ρ 2.3 FTP
- ρ 2.4 Electronic Mail
  - μ SMTP, POP3, IMAP
- ρ 2.5 DNS
- ρ 2.6 Socket programming with TCP
- ρ 2.7 Socket programming with UDP
- ρ 2.8 Building a Web server
- ρ 2.9 Content distribution
  - μ Network Web caching
  - μ Content distribution networks
  - μ P2P file sharing

## Web caches (proxy server)

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

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



## More about Web caching

- ρ Cache acts as both client and server
- ρ Cache can do up-to-date check using `If-modified-since` HTTP header
  - μ Issue: should cache take risk and deliver cached object without checking?
  - μ Heuristics are used.
- ρ Typically cache is installed by ISP (university, company, residential ISP)

### Why Web caching?

- ρ Reduce response time for client request.
- ρ Reduce traffic on an institution's access link.
- ρ Internet dense with caches enables "poor" content providers to effectively deliver content


## Caching example (1)

### Assumptions

- ρ average object size = 100,000 bits
- ρ avg. request rate from institution's browser to origin servers = 15/sec
- ρ delay from institutional router to any origin server and back to router = 2 sec

### Consequences

- ρ utilization on LAN = 15%
- ρ utilization on access link = 100%
- ρ total delay = Internet delay + access delay + LAN delay



Application Layer 85


## Caching example (2)

### Possible solution

- ρ increase bandwidth of access link to, say, 10 Mbps

### Consequences

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



Application Layer 86


## Caching example (3)

### Install cache

- ρ suppose hit rate is .4

### Consequence

- ρ 40% requests will be satisfied almost immediately
- ρ 60% requests satisfied by origin server
- ρ utilization of access link reduced to 60%, resulting in negligible delays (say 10 msec)
- ρ total delay = Internet delay + access delay + LAN delay = .6 \* 2 sec + .6 \* .01 secs + milliseconds < 1.3 secs




Application Layer 87

## Content distribution networks (CDNs)

- ρ The content providers are the CDN customers.

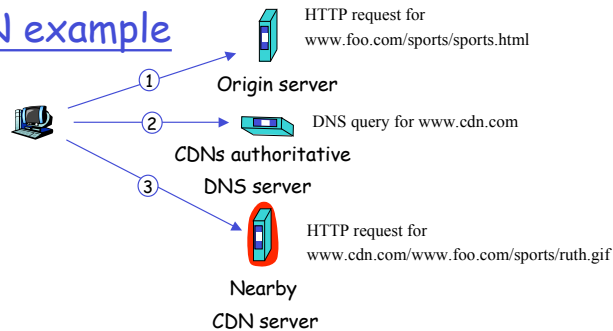
### Content replication

- ρ CDN company installs hundreds of CDN servers throughout Internet
  - μ in lower-tier ISPs, close to users
- ρ CDN replicates its customers' content in CDN servers. When provider updates content, CDN updates servers



Application Layer 88

## CDN example



### origin server

- ρ www.foo.com
- ρ distributes HTML
- ρ Replaces:  
http://www.foo.com/sports.ruth.gif  
with  
http://www.cdn.com/www.foo.com/sports/ruth.gif

### CDN company

- ρ cdn.com
- ρ distributes gif files
- ρ uses its authoritative DNS server to route requests

Application Layer 89

## More about CDNs

### routing requests

- ρ CDN creates a "map", indicating distances from leaf ISPs and CDN nodes
- ρ when query arrives at authoritative DNS server:
  - μ server determines ISP from which query originates
  - μ uses "map" to determine best CDN server

### not just Web pages

- ρ streaming stored audio/video
- ρ streaming real-time audio/video
  - μ CDN nodes create application-layer overlay network

Application Layer 90

## P2P file sharing

### Example

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


All peers are servers = highly scalable

Application Layer 91

## P2P: centralized directory

### original "Napster" design

- 1) when peer connects, it informs central server:
  - μ IP address
  - μ content
- 2) Alice queries for "Hey Jude"
- 3) Alice requests file from Bob



Application Layer 92

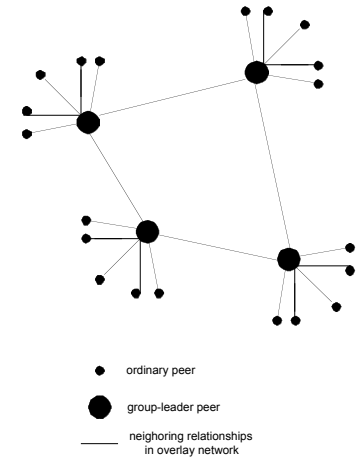
## P2P: problems with centralized directory

- ρ Single point of failure
- ρ Performance bottleneck
- ρ Copyright infringement

file transfer is decentralized, but locating content is highly decentralized

## P2P: decentralized directory

- ρ Each peer is either a group leader or assigned to a group leader.
- ρ Group leader tracks the content in all its children.
- ρ Peer queries group leader; group leader may query other group leaders.



## More about decentralized directory

### overlay network

- ρ peers are nodes
- ρ edges between peers and their group leaders
- ρ edges between some pairs of group leaders
- ρ virtual neighbors

### bootstrap node

- ρ connecting peer is either assigned to a group leader or

### advantages of approach



- ρ no centralized directory server
  - μ location service distributed over peers
  - μ more difficult to shut down

### disadvantages of approach

- ρ bootstrap node needed
- ρ group leaders can get overloaded

## P2P: Query flooding

- ρ Gnutella
- ρ no hierarchy
- ρ use bootstrap node to learn about others
- ρ join message
- ρ Send query to neighbors
- ρ Neighbors forward query
- ρ If queried peer has object, it sends message back to querying peer



## P2P: more on query flooding

### Pros

- ρ peers have similar responsibilities: no group leaders
- ρ highly decentralized
- ρ no peer maintains directory info

### Cons

- ρ excessive query traffic
- ρ query radius: may not have content when present
- ρ bootstrap node
- ρ maintenance of overlay network

## Chapter 2: Summary

### Our study of network apps now complete!

- |  |   |
|--|---|
| <ul style="list-style-type: none"> <li>ρ application service requirements:             <ul style="list-style-type: none"> <li>μ reliability, bandwidth, delay</li> </ul> </li> <li>ρ client-server paradigm</li> <li>ρ Internet transport service model             <ul style="list-style-type: none"> <li>μ connection-oriented, reliable: TCP</li> <li>μ unreliable, datagrams: UDP</li> </ul> </li> </ul> | <ul style="list-style-type: none"> <li>ρ specific protocols:             <ul style="list-style-type: none"> <li>μ HTTP</li> <li>μ FTP</li> <li>μ SMTP, POP, IMAP</li> <li>μ DNS</li> </ul> </li> <li>ρ socket programming</li> <li>ρ content distribution             <ul style="list-style-type: none"> <li>μ caches, CDNs</li> <li>μ P2P</li> </ul> </li> </ul> |
|--|---|

## Chapter 2: Summary

### Most importantly: learned about protocols

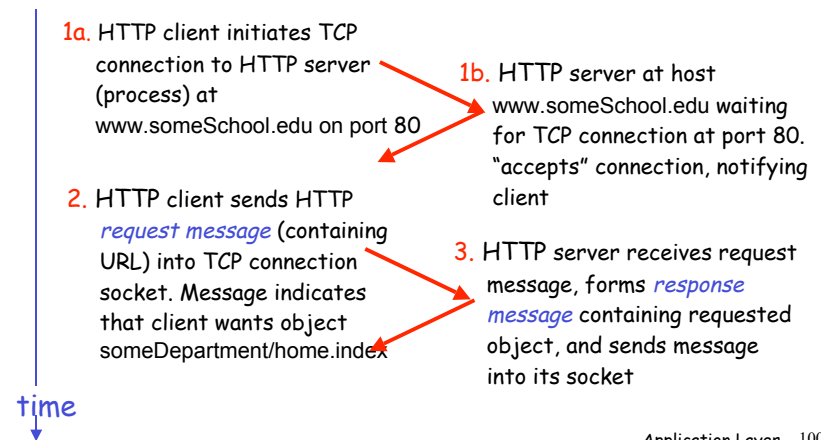
- |   |  |
|---|--|
| <ul style="list-style-type: none"> <li>ρ typical request/reply message exchange:             <ul style="list-style-type: none"> <li>μ client requests info or service</li> <li>μ server responds with data, status code</li> </ul> </li> <li>ρ message formats:             <ul style="list-style-type: none"> <li>μ headers: fields giving info about data</li> <li>μ data: info being communicated</li> </ul> </li> </ul> | <ul style="list-style-type: none"> <li>ρ control vs. data msgs             <ul style="list-style-type: none"> <li>μ in-band, out-of-band</li> </ul> </li> <li>ρ centralized vs. decentralized</li> <li>ρ stateless vs. stateful</li> <li>ρ reliable vs. unreliable msg transfer</li> <li>ρ "complexity at network edge"</li> <li>ρ security: authentication</li> </ul> |
|---|--|

## Nonpersistent HTTP

Suppose user enters URL

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

(contains text, references to 10 jpeg images)



## Nonpersistent HTTP (cont.)

- time ↓
5. HTTP client receives response message containing html file, displays html. Parsing html file, finds 10 referenced jpeg objects
  6. Steps 1-5 repeated for each of 10 jpeg objects
4. HTTP server closes TCP connection.
-