

# Networking Applications

**Dr. Ayman A. Abdel-Hamid**


College of Computing and Information Technology

Arab Academy for Science & Technology and  
Maritime Transport

**The Web and HTTP**

The material is adapted from slides for  
Chapter 2 and Chapter 7 of this textbook

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



*Computer Networking:  
A Top Down Approach,  
5<sup>th</sup> edition.*

*Jim Kurose, Keith Ross  
Addison-Wesley, April  
2009.*

# History of The Web

- ❑ World Wide Web, "Web", "WWW"
- ❑ Tim Berners-Lee at CERN in 1991
  - ❖ Demonstrated prototype at a conf. in '91
  - ❖ Text-based
- ❑ Marc Andreessen developed the first graphical Web browser in 1993: Mosaic
- ❑ Andreessen founds Netscape Communications
- ❑ Browser war starts around 1995-96
- ❑ America Online buys Netscape in 1998

# 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 defines communication protocol between client and sever
- ❑ HTTP 1.0: RFC 1945
- ❑ HTTP 1.1: RFC 2616 (starting 1998)



# 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

### Protocols that maintain "state" are complex! aside

- ❑ 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

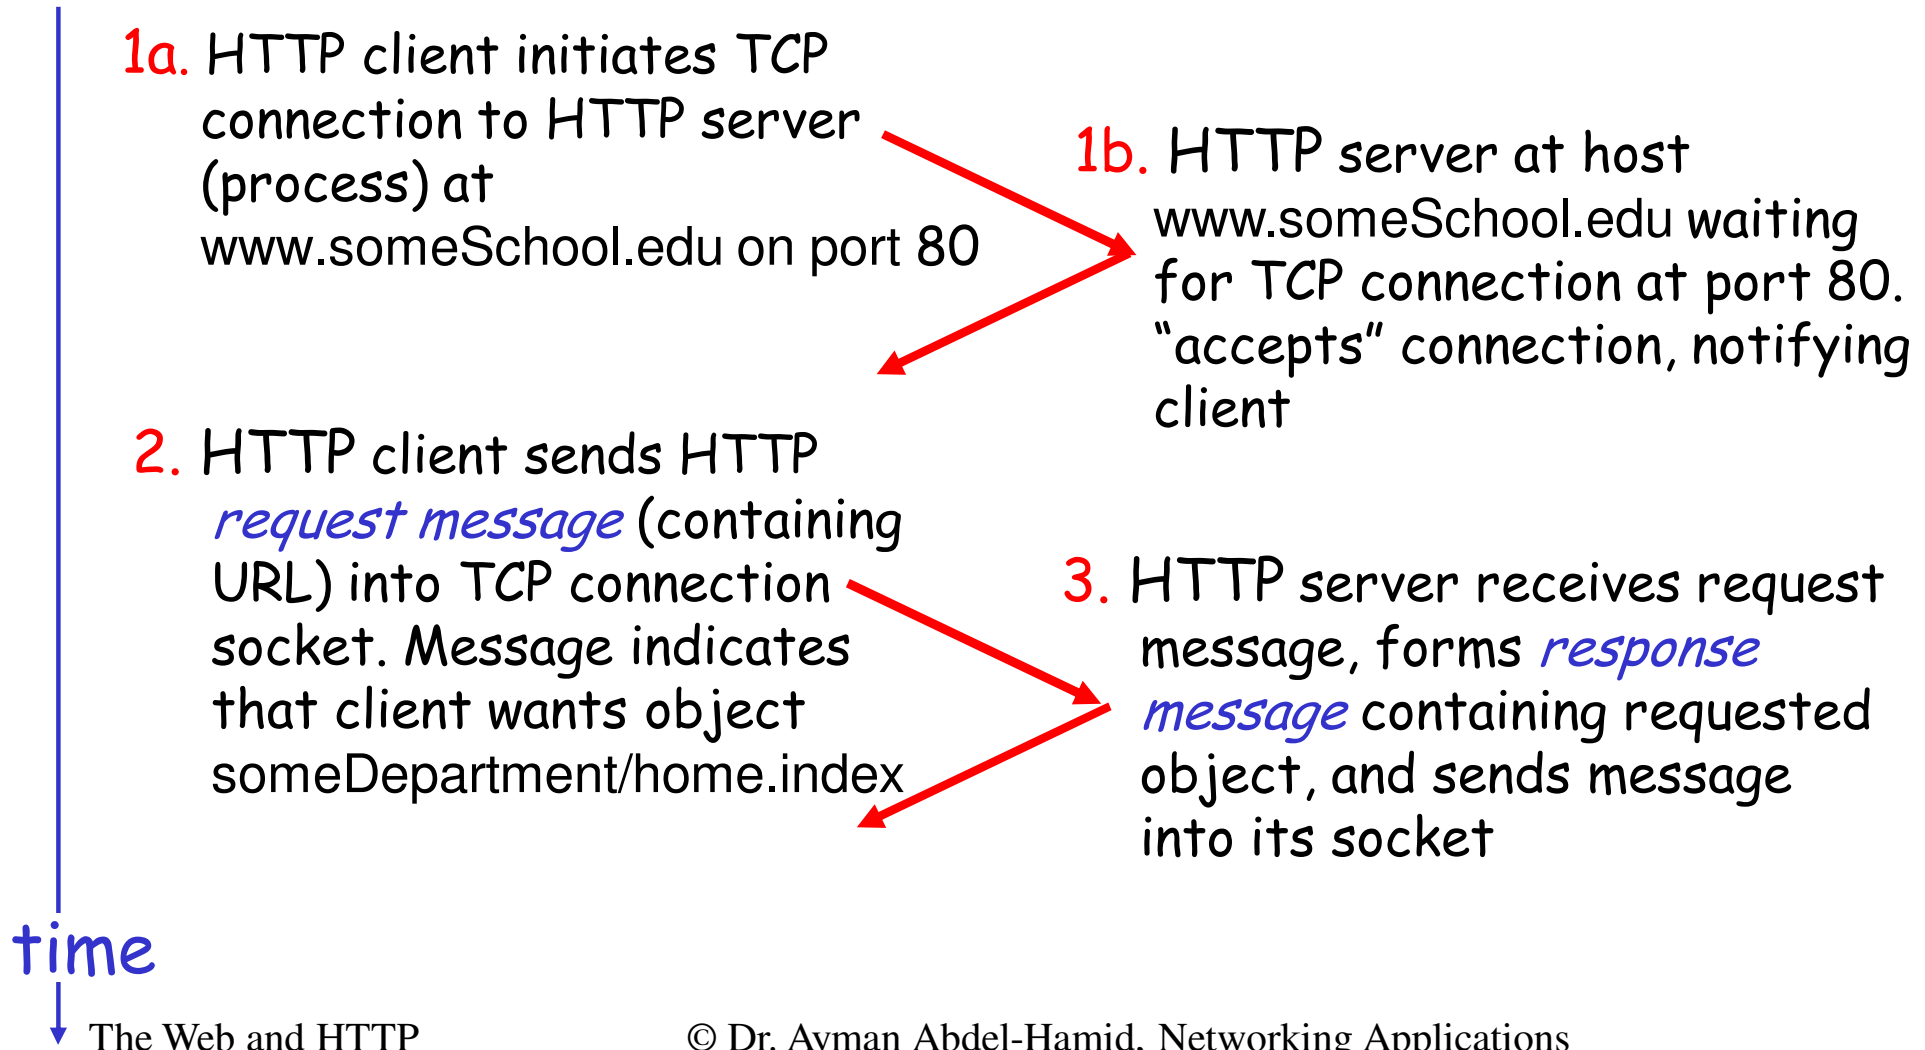
- ❑ Multiple objects can be sent over single TCP connection between client and server.
- ❑ HTTP/1.1 uses persistent connections in default mode

# Nonpersistent HTTP

Suppose user enters URL

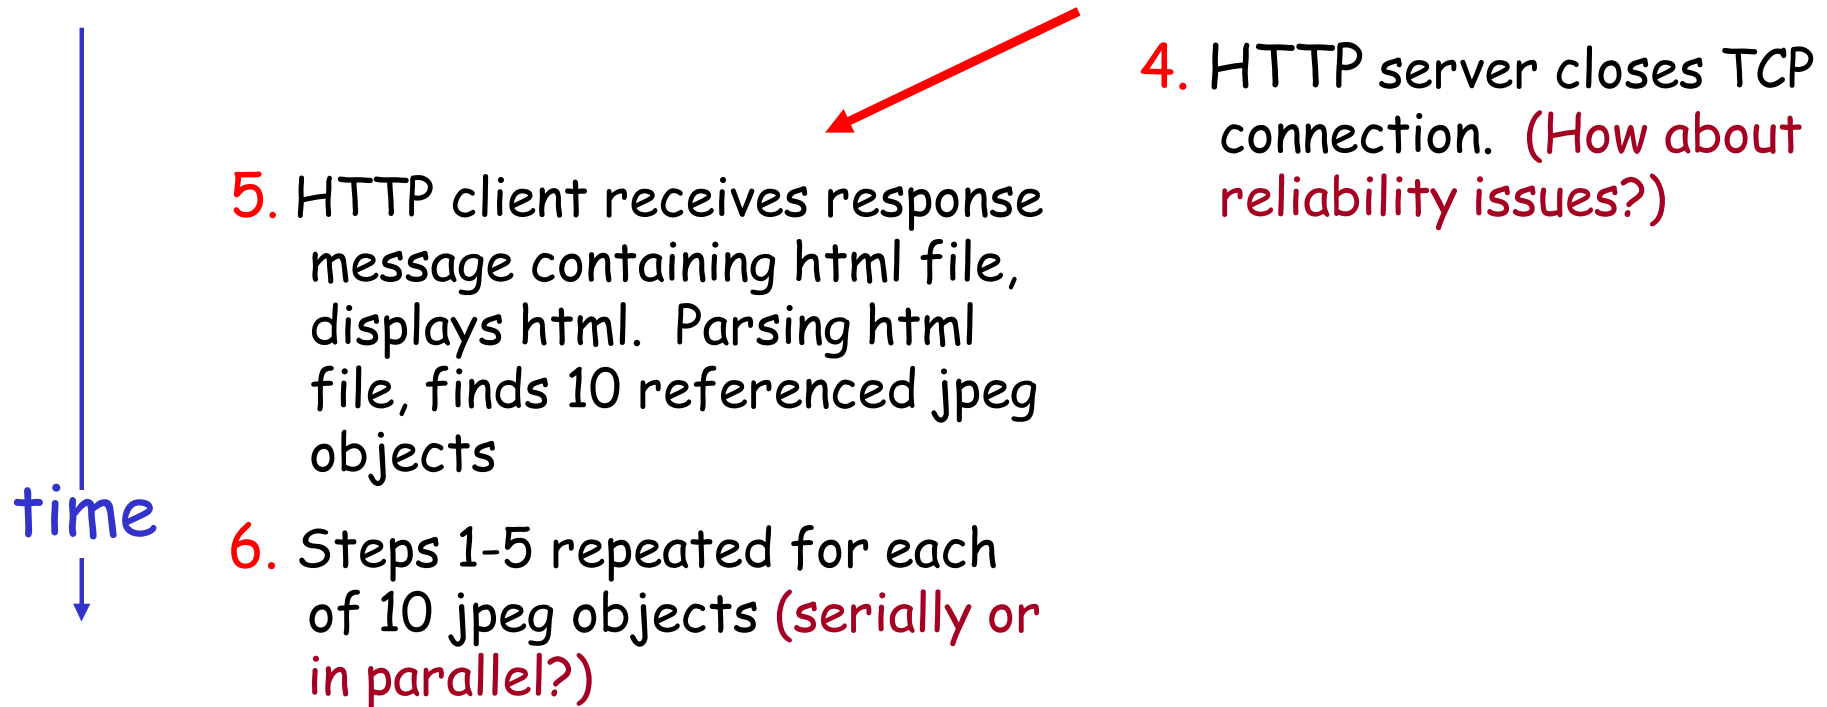
`www.someSchool.edu/someDepartment/home.index`

(contains text,  
references to 10  
jpeg images)





# Nonpersistent HTTP (cont.)




# Non-Persistent HTTP: Response time

**Definition of 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 =  $2RTT + \text{transmit time}$**



# Persistent HTTP

## Nonpersistent HTTP issues:

- ❑ requires 2 RTTs per object
- ❑ OS overhead for *each* TCP connection (buffers, etc...)
- ❑ 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
- ❑ Close connection when not used for a certain time

## 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)

request line  
(GET, POST,  
HEAD commands)

Used when browser requests an object

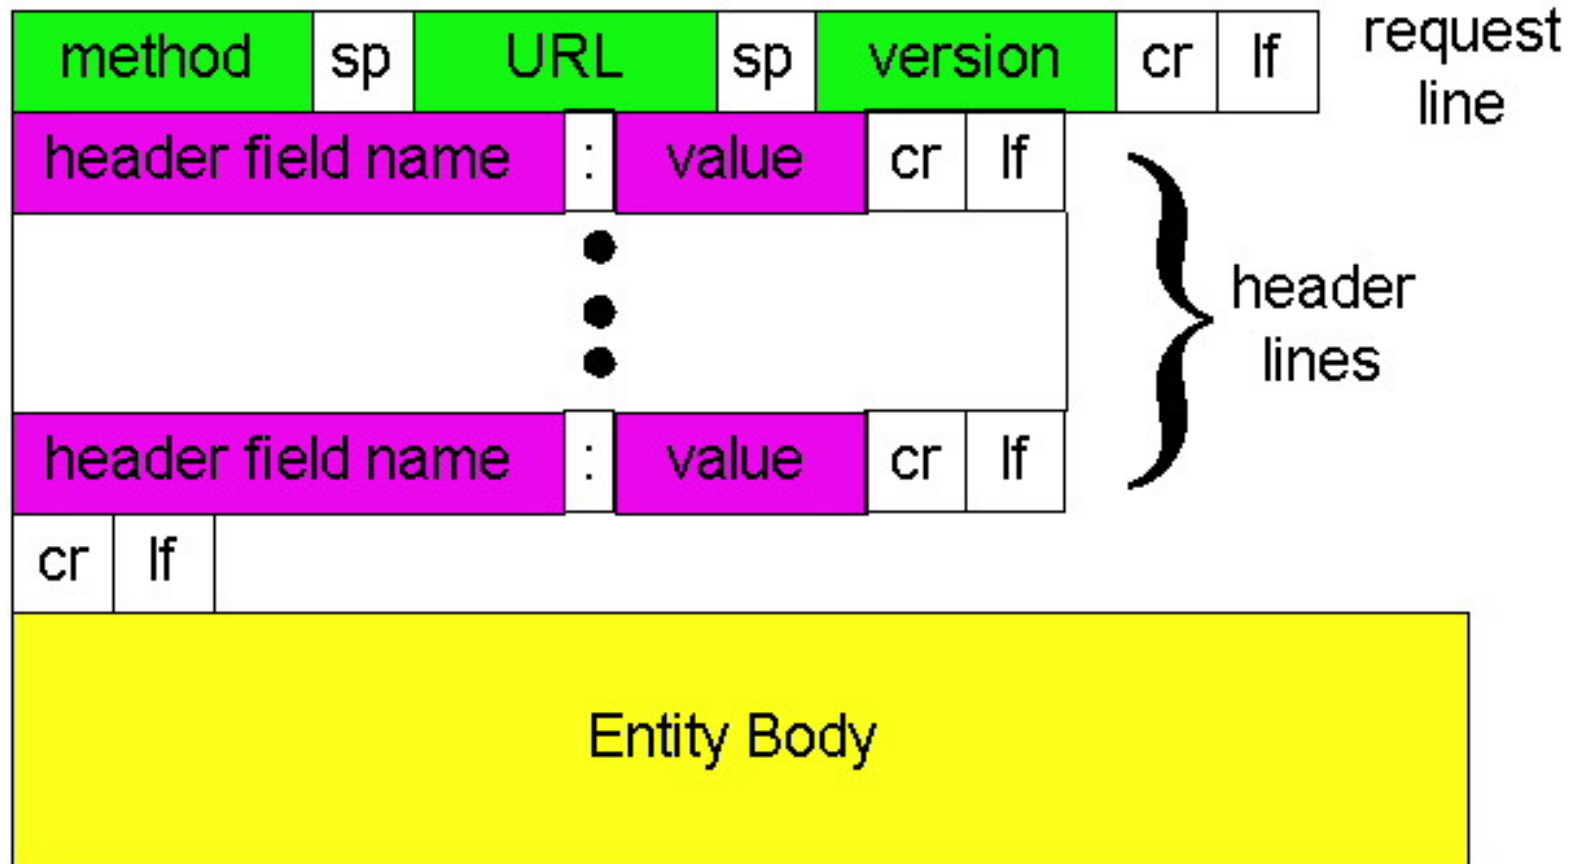
header  
lines

```
GET /somedir/page.html HTTP/1.1
Host: www.someschool.edu
User-agent: Mozilla/4.0
Connection: close
Accept-language: fr
```

Carriage return,  
line feed  
indicates end  
of message

(extra carriage return, line feed)

# 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

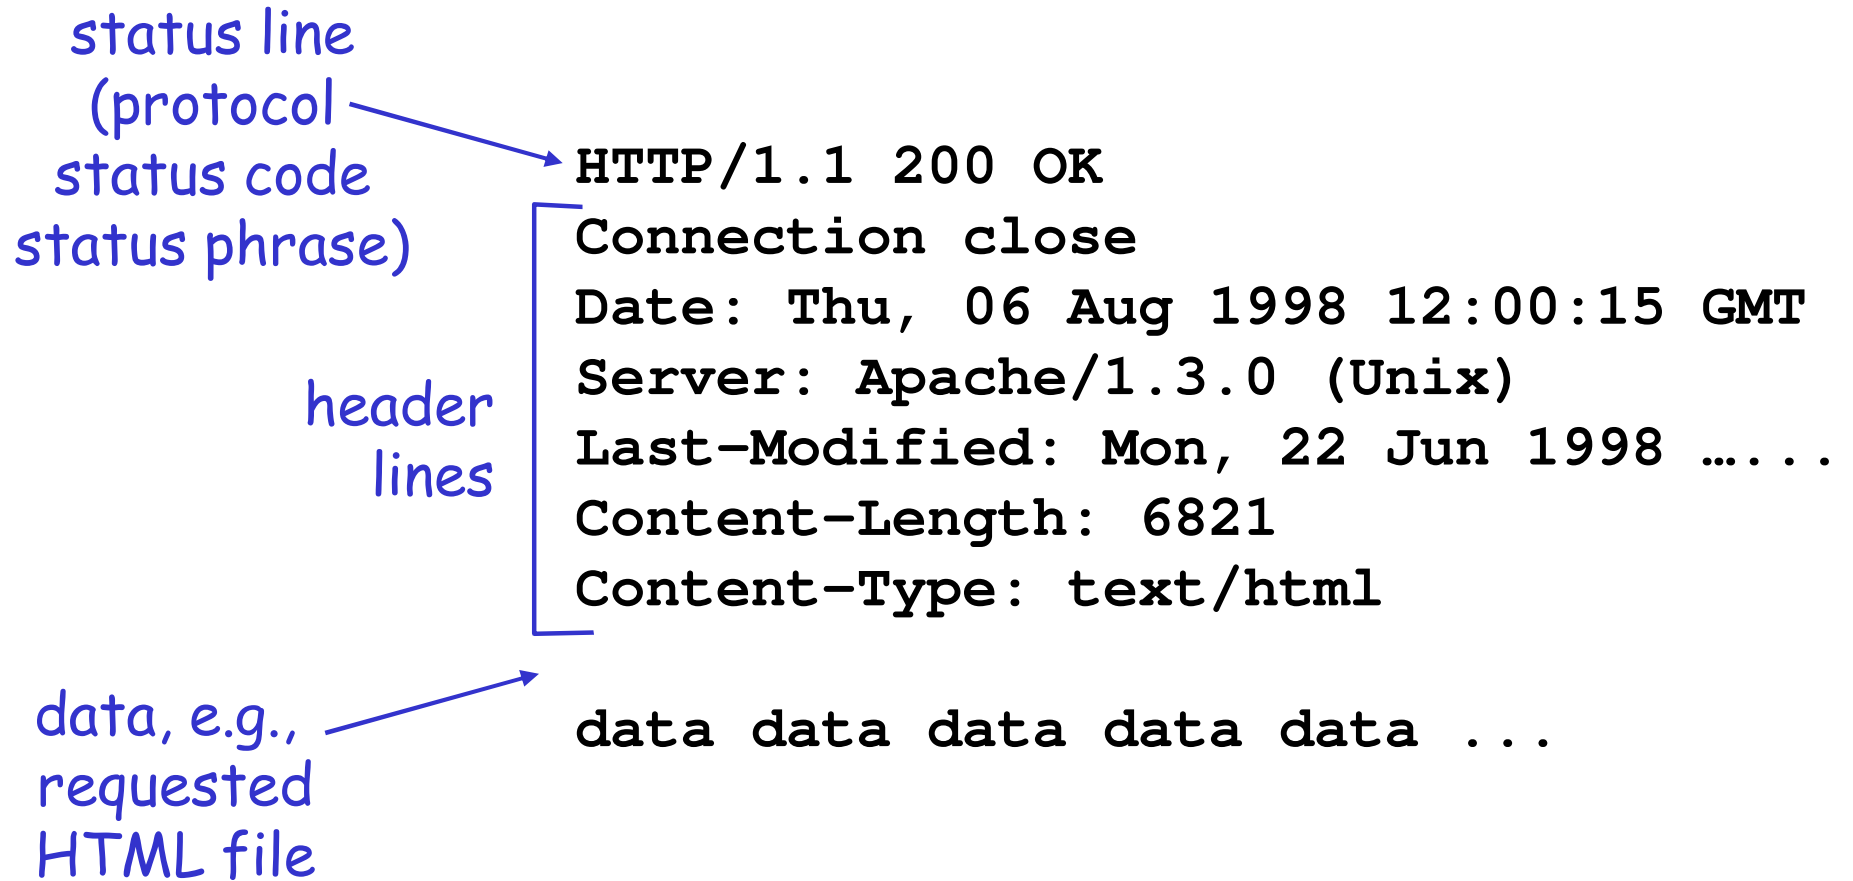
## HTTP/1.0

- ❑ 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!

# User-server state: cookies (RFC 2109)

Many major Web sites  
use cookies

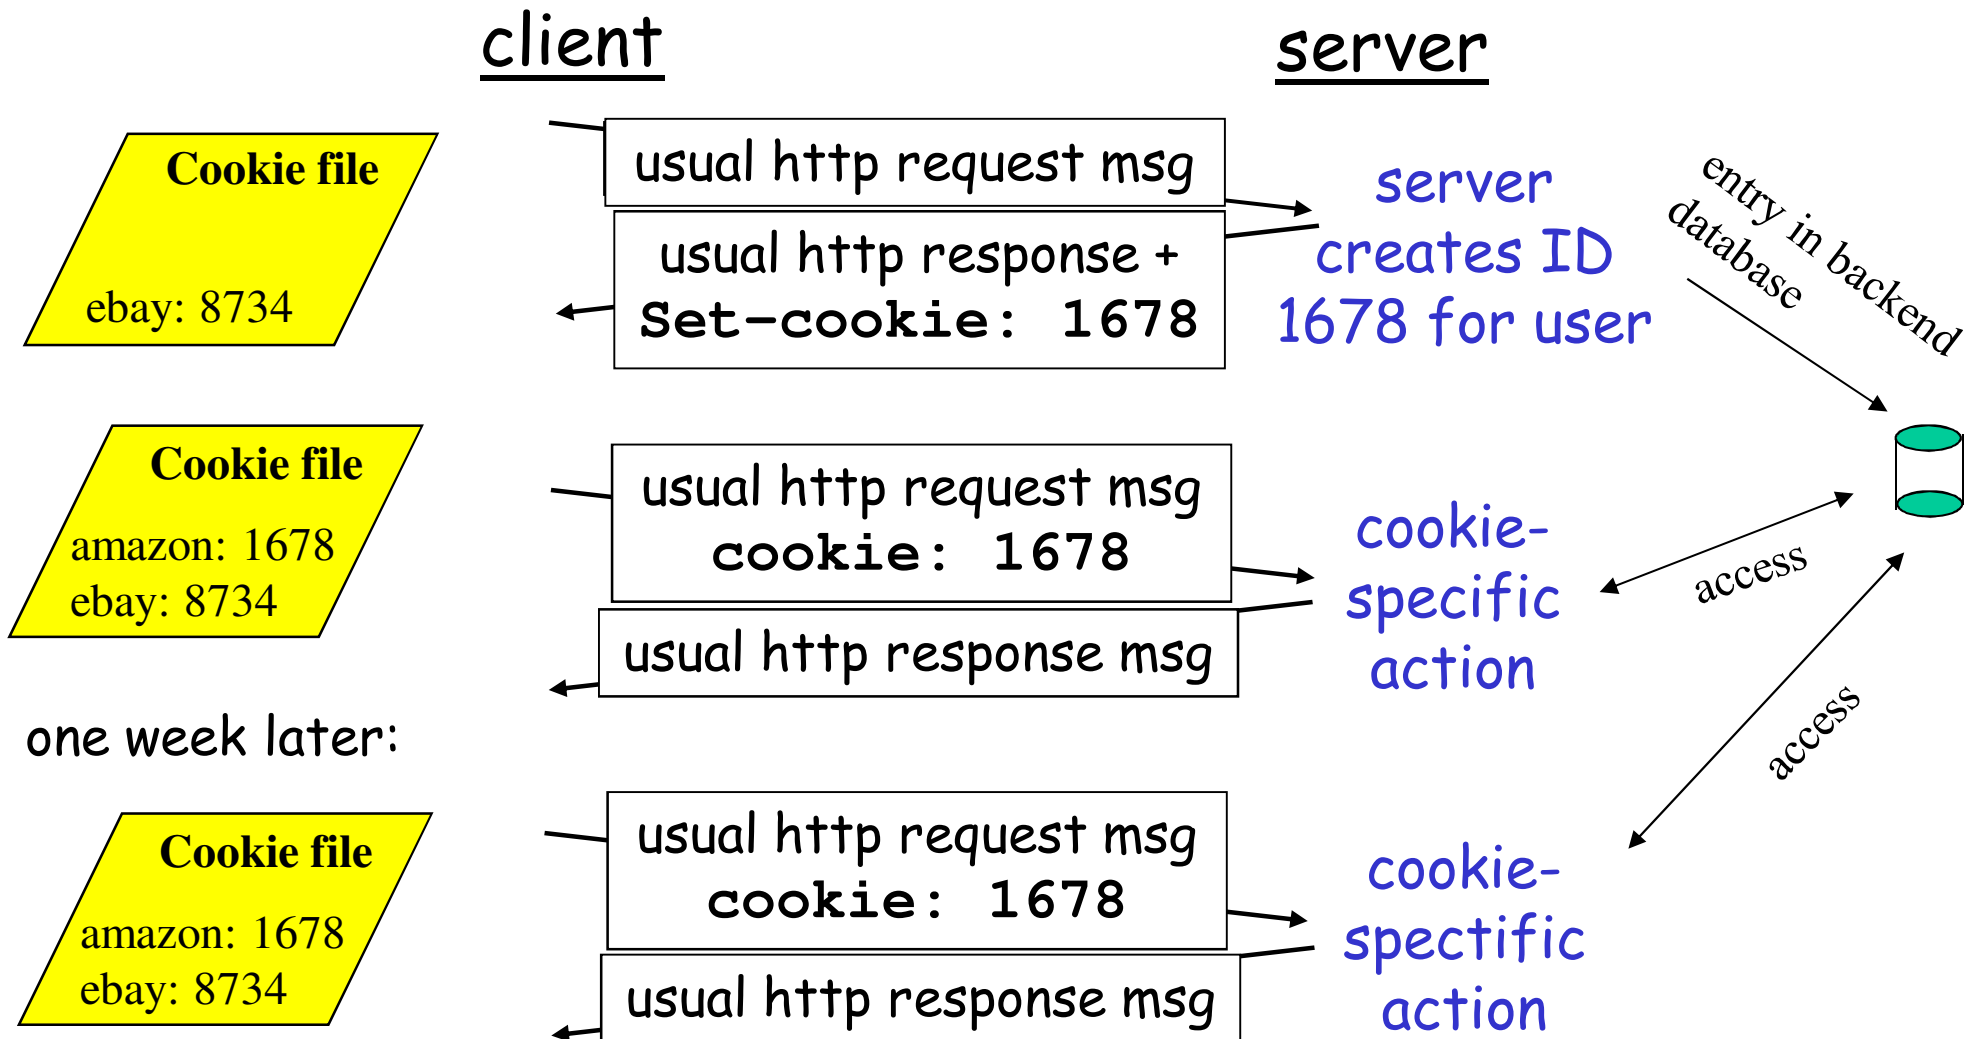
## Four components:

- 1) cookie header line of  
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)

## Cookies and privacy: aside

- ☐ cookies permit sites to learn a lot about you
- ☐ you may supply name and e-mail to sites


## How to keep "state":

- ☐ Protocol endpoints: maintain state at sender/receiver over multiple transactions
- ☐ cookies: http messages carry state

# 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 = 100,000 bits
- avg. request rate from institution's browsers to origin servers = 15/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, 10 Mbps

## Consequences

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




# Caching example (cont)

## Install cache

- suppose hit rate is .4

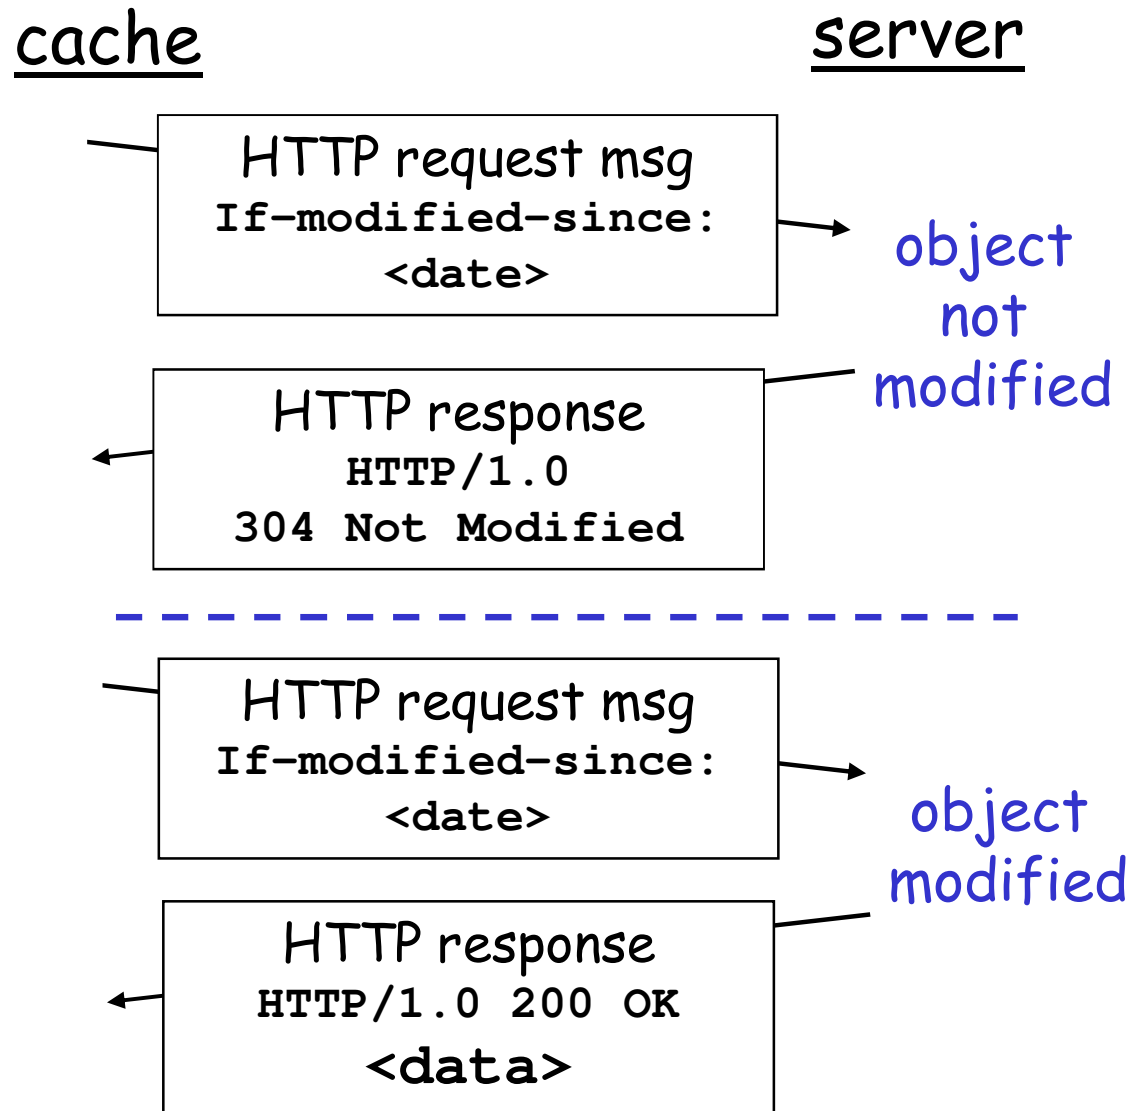
## Consequence

- 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  
 $= .6 * (2.01) \text{ secs} + .4 * \text{milliseconds} < 1.4 \text{ secs}$



# 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



# Content distribution networks (CDNs)

## Content replication

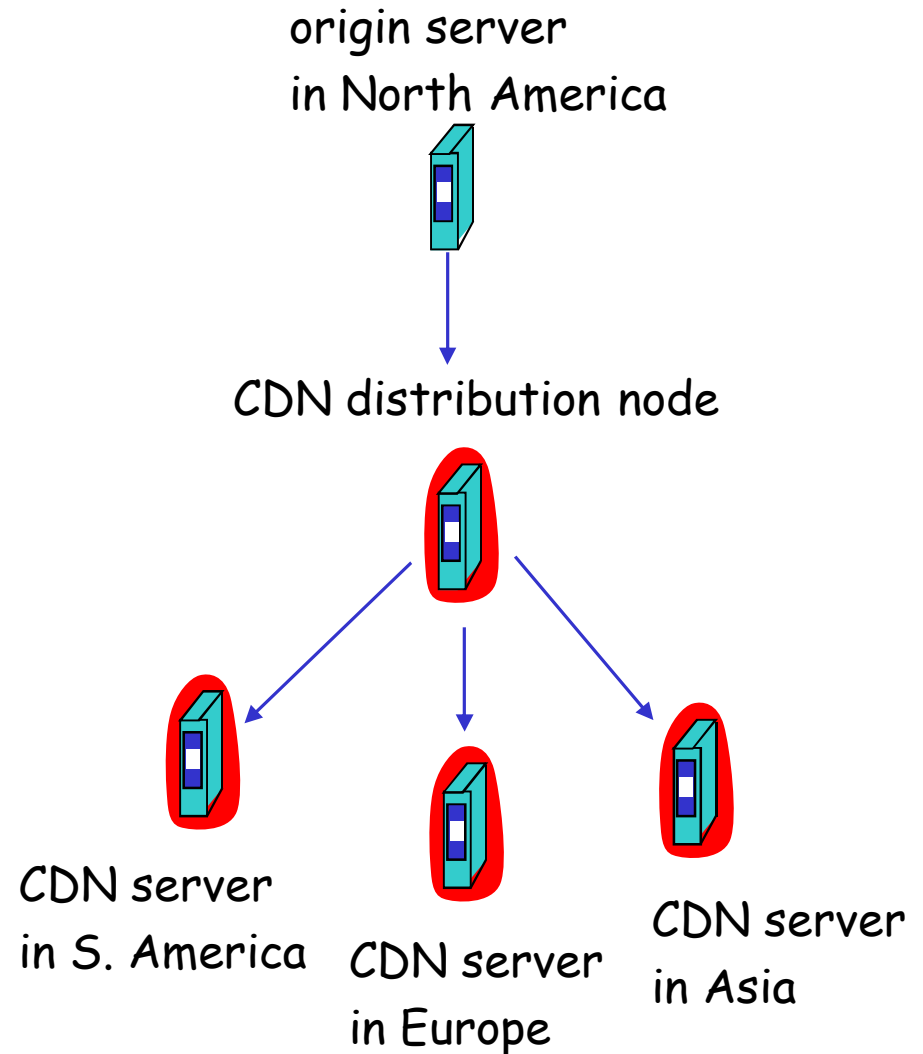
- ❑ Challenging to stream large files (e.g., video) from single origin server in real time
- ❑ Solution: replicate content at hundreds of servers throughout Internet
  - ❖ content downloaded to CDN servers ahead of time
  - ❖ placing content "close" to user avoids impairments (loss, delay) of sending content over long paths
  - ❖ CDN server typically in edge/access network



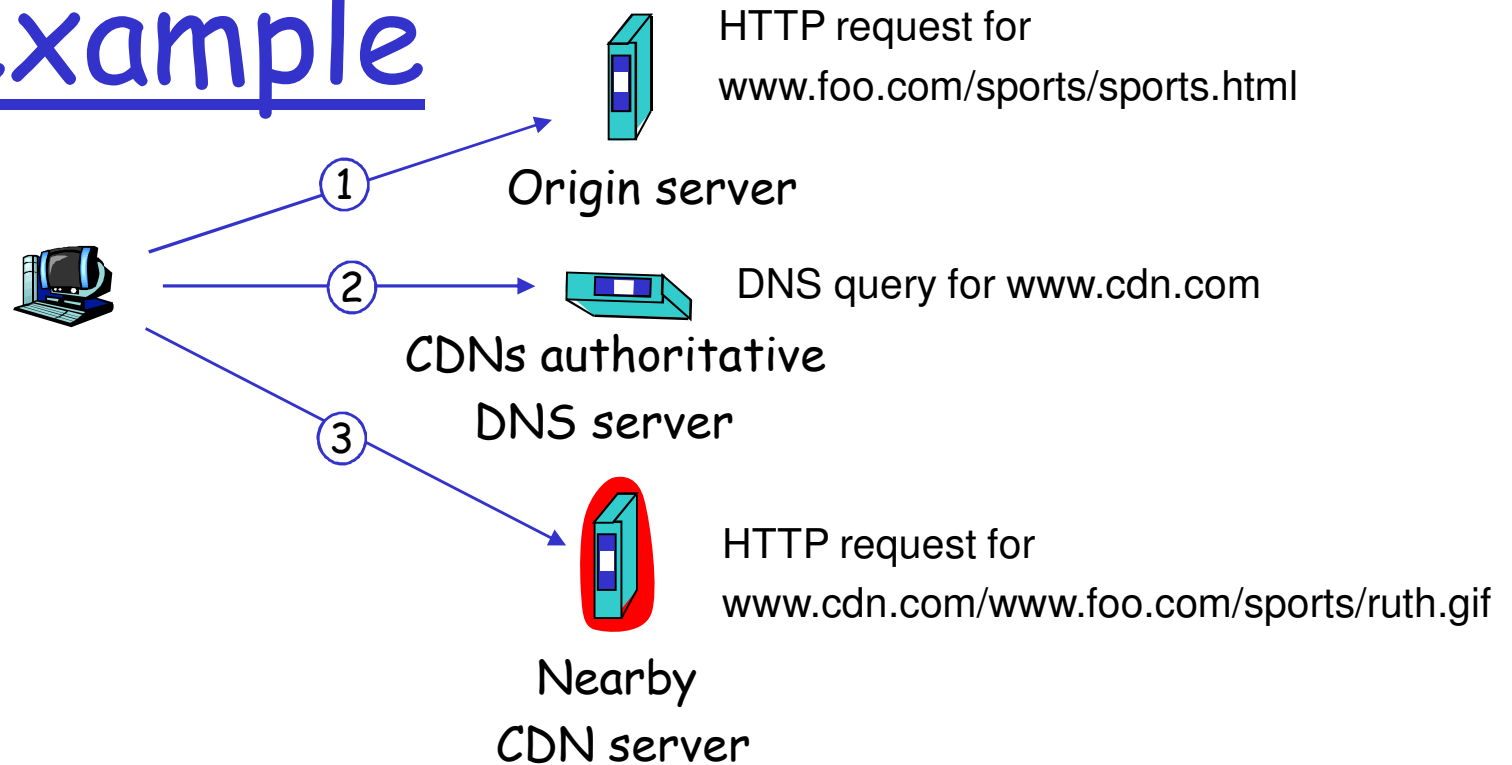
# Content distribution networks (CDNs)

## Content replication

- ❑ CDN (e.g., Akamai) customer is the content provider (e.g., CNN)
- ❑ CDN replicates customers' content in CDN servers. When provider updates content, CDN updates servers



# CDN example



origin server (www.foo.com) CDN company (cdn.com)

❑ distributes HTML

❑ replaces:

`http://www.foo.com/sports.ruth.gif`

with

`http://www.cdn.com/www.foo.com/sports/ruth.gif`

❑ distributes gif files

❑ uses its authoritative DNS server to route redirect requests

# More about CDNs

## routing requests

- ❑ CDN creates a "map", indicating distances from leaf ISPs and CDN nodes
- ❑ when query arrives at authoritative DNS server:
  - ❖ server determines ISP from which query originates
  - ❖ uses "map" to determine best CDN server
- ❑ CDN nodes create application-layer overlay network

# Further Information

- RFC 1945: HTTP 1.0, May 1996
- RFC 2616: HTTP 1.1, June 1999
- RFC 2109: HTTP State Management Mechanism, February 1997