

Application Layer

Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:

<http://www.cse.wustl.edu/~jain/cse473-10/>



1. Network Application Architecture
2. HyperText Transfer Protocol (HTTP)
3. File Transfer and Email protocols
4. Domain Name Service
5. Peer-to-Peer Applications

Note: This class lecture is based on Chapter 2 of the textbook (Kurose and Ross) and the figures provided by the authors.



Network Application Architectures

1. Protocol Layers
2. Client-Server vs. Peer-to-Peer
3. Process Communication
4. Names, Addresses, Ports
5. Transports

Protocol Layers

□ Top-Down approach


Application	HTTP	FTP	SMTP	P2P	DNS	Skype
Transport	TCP				UDP	
Internetwork	IP					
Host to Network	Ethernet	Point-to-Point		Wi-Fi		
Physical	Coax	Fiber	Wireless			

Network Application Architectures

- ❑ Client-Server
- ❑ Peer-to-Peer


Client-Server

- ❑ Clients: Request service
- ❑ Server: Provides a service.
Waits for clients
- ❑ Server is always up
- ❑ Clients do not communicate directly with each other
- ❑ Server = Data Center
- ❑ Example: Web Server, Search Engine, Social Networking



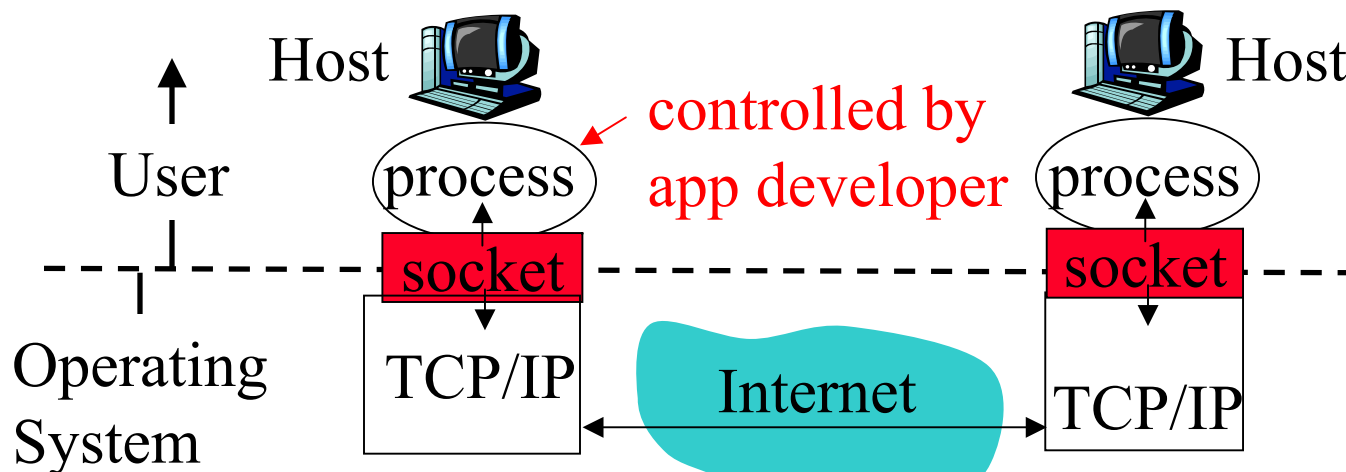
Peer-to-Peer

- ❑ Does not require always-on servers
- ❑ Hosts communicate directly
⇒ Peers
- ❑ Hosts may come on or may go off at any time
- ❑ Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- ❑ Highly scalable
- ❑ Highly symmetric traffic
⇒ ISP unfriendly
- ❑ Difficult to authenticate ⇒ Insecure
- ❑ Need incentives to share



Process Communications

- ❑ Inter-Process Communication on the Same Host
⇒ Operating system provides message passing
- ❑ Unix provides application programming interface called “sockets”
- ❑ Inter-Process Communication on Different Hosts
⇒ Network provides message passing



Names, Addresses, Ports

- ❑ Domain Name System: www.google.com
- ❑ IP Address: 209.85.225.147
- ❑ 4 decimal numbers less than 256=8 bits each
⇒ 32-bits
- ❑ Ports: Entry point (Transport service access points)
- ❑ 21=FTP, 80=HTTP

Port 1




Port 2

Transports

TCP	UDP
Reliable data transfer	Unreliable Data Transfer
Packet Sequence # required	Sequence # optional
Every packet is acked	Not Acked
Lost packets are retransmitted	No Retransmission
May cause long delay	Quick and Lossy
Connection-oriented service	Connection-less Service
Good for Reliable and delay-insensitive applications	Good for loss-tolerant and delay sensitive applications
Applications: email, http, ftp, Remote terminal access	Telephony, Streaming Multimedia

Application Layer Protocols

- ❑ HTTP: HyperText Transfer Protocol
- ❑ FTP: File Transfer Protocol
- ❑ SMTP: Simple Mail Transfer Protocol
- ❑ DNS: Domain Name Server
(Control Plane Application)
- ❑ P2P: Peer-to-Peer Applications
- ❑ Skype
- ❑ Each application has its own protocol, message format, semantics of fields



Application Arch: Summary

1. P2P applications are more scalable than client-server
2. Applications exchanges messages using operating system sockets
3. Applications communicate using host names, addresses, and ports
4. Applications use transports: TCP, UDP, ...
5. TCP is used for reliable communication
UDP for loss-tolerant delay-sensitive applications




HTTP

1. Concepts
2. Sample Web Page
3. HTTP Messages
4. Cookies
5. Proxy Servers
6. Conditional GET

HTTP Concepts

- ❑ **Client**=Browser, e.g., Internet Explorer, Firefox
- ❑ **HTTP Server**, e.g., Microsoft Internet Information Service (IIS), Apache
- ❑ **Web Page**=Group of objects
- ❑ **Object**=Text, Images, files, ...
- ❑ **URL**: Uniform Resource Locator
<http://www.cse.wustl.edu/~jain/cse473-09/sample.htm>




HTTP

- ❑ Uses TCP
- ❑ **Stateless**: Server does not remember previous history
- ❑ **Non-Persistent**: Open new TCP connection, get one object, close
- ❑ **Persistent**: Open one TCP connection, get all objects, close
Server leaves the connection open after sending an object and closes on timeout
- ❑ Web pages are written in HyperText Markup Language (**HTML**)

Sample Web Page

```
<HTML>
<HEAD>
</HEAD>
<BODY>
<img src=jain.jpg>
<BR>
Raj Jain
</BODY>
</HTML>
```



Raj Jain

Sample HTTP Request Message

GET /~jain/cse473-09/sample.htm HTTP/1.1

Host: www.cse.wustl.edu

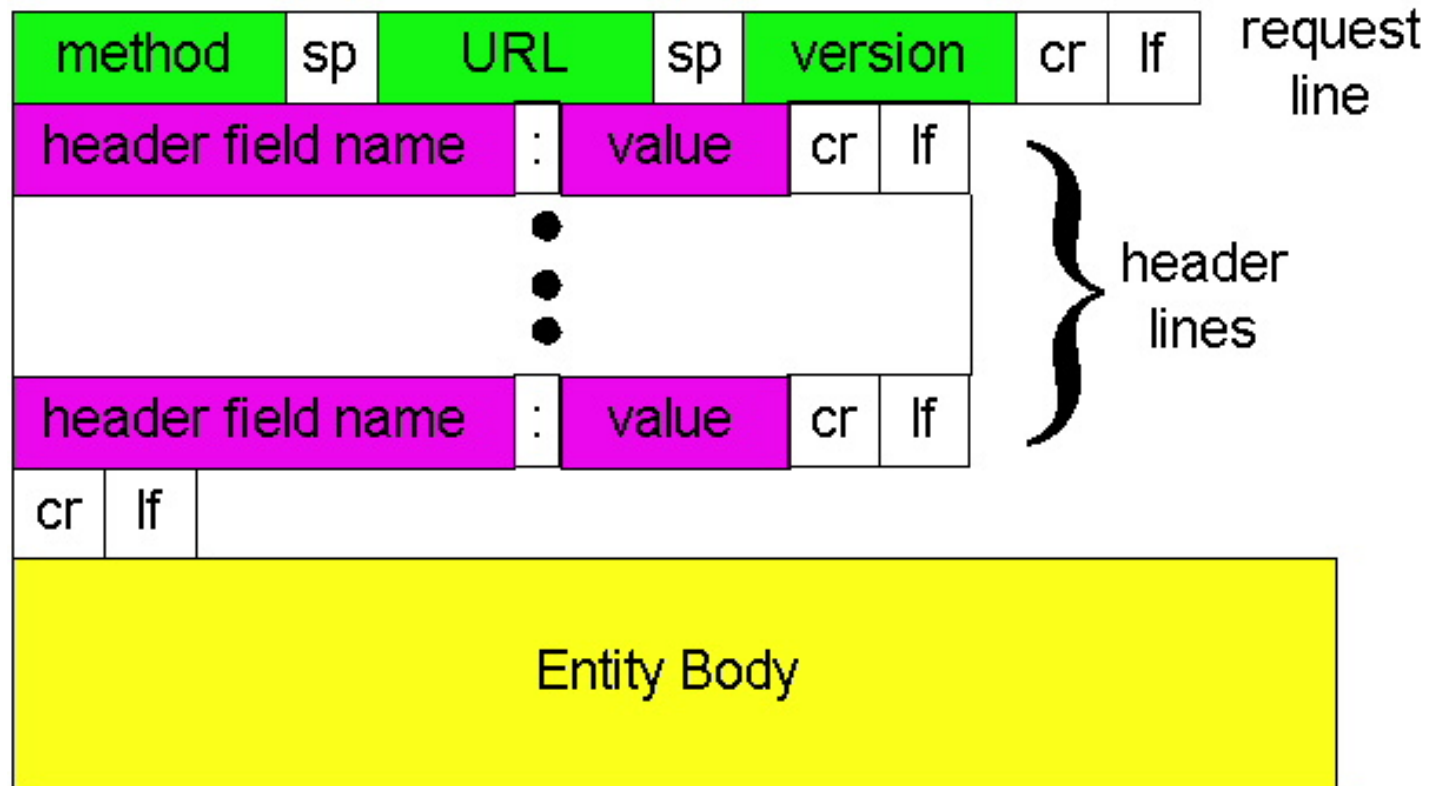
Connection: close

User-agent: Mozilla/4.0

Accept-Language: en

- ❑ **Method** = Get
- ❑ **URL** = /~jain/cse473-09/sample.htm
- ❑ **Version** = HTTP/1.1
- ❑ **Header Fields** = Host, Connection, User-agent, ...

HTTP Request Message Format



Sample HTTP Response Message

HTTP/1.1 200 OK

Connection: close

Date: Tue, 09 Sept 2009 13:00:15 GMT

Server: Apache/1.3.0 (Unix)

Last-Modified: Sun, 6 May 2009 09:23:24 GMT

Content-Length: 6500

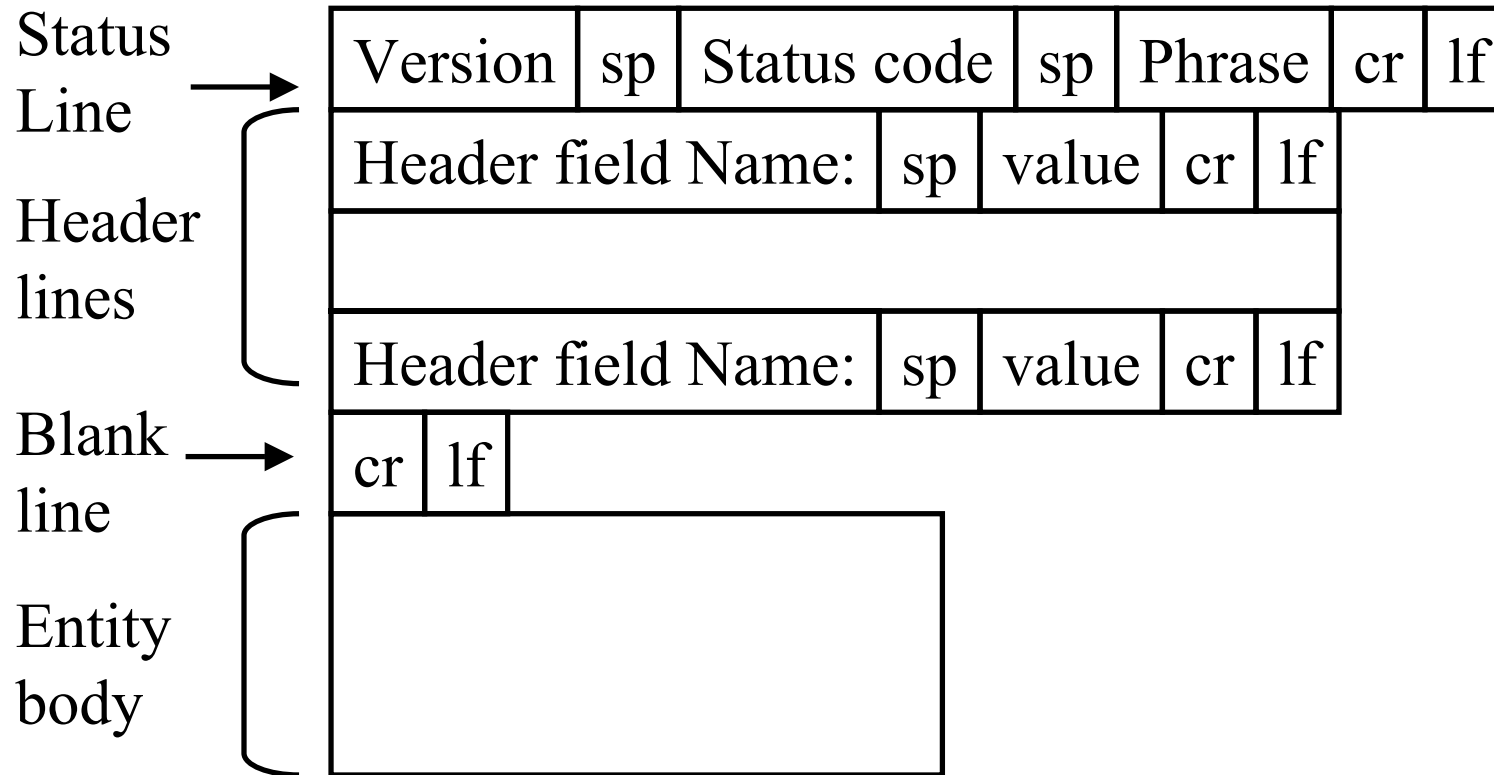
Content-Type: Text/html

Data...

Status Codes:

- ❑ 200 OK
- ❑ 301 Moved Permanently
- ❑ 400 Bad Request
- ❑ 404 Not Found
- ❑ 505 HTTP Version Not Supported

HTTP Response Message Format



Hands-on HTTP

```
telnet www.cse.wustl.edu 80
GET /~jain/cse473-10/sample.htm HTTP/1.1
Host: www.cse.wustl.edu
```

```
HTTP/1.1 200 OK
Date: Tue, 08 Sep 2009 23:39:53 GMT
Server: Apache/2.0.52 (CentOS)
Last-Modified: Tue, 08 Sep 2009 23:38:26 GMT
ETag: "1e329a-4b-767b2080"
Accept-Ranges: bytes
Content-Length: 75
Connection: close
Content-Type: text/html; charset=UTF-8
```

```
<HTML>
<head>
</head>
<body>
This is a sample text.
</body>
</html>
```



This is a sample text.

Hands-on HTTP (cont)

```
telnet www.cse.wustl.edu 80
GET /~jain/cse473-08/sample.htm HTTP/1.1
Host: www.cse.wustl.edu
```

```
HTTP/1.1 404 Not Found
Date: Tue, 08 Sep 2009 23:42:48 GMT
Server: Apache/2.0.52 (CentOS)
Content-Length: 307
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

Not Found

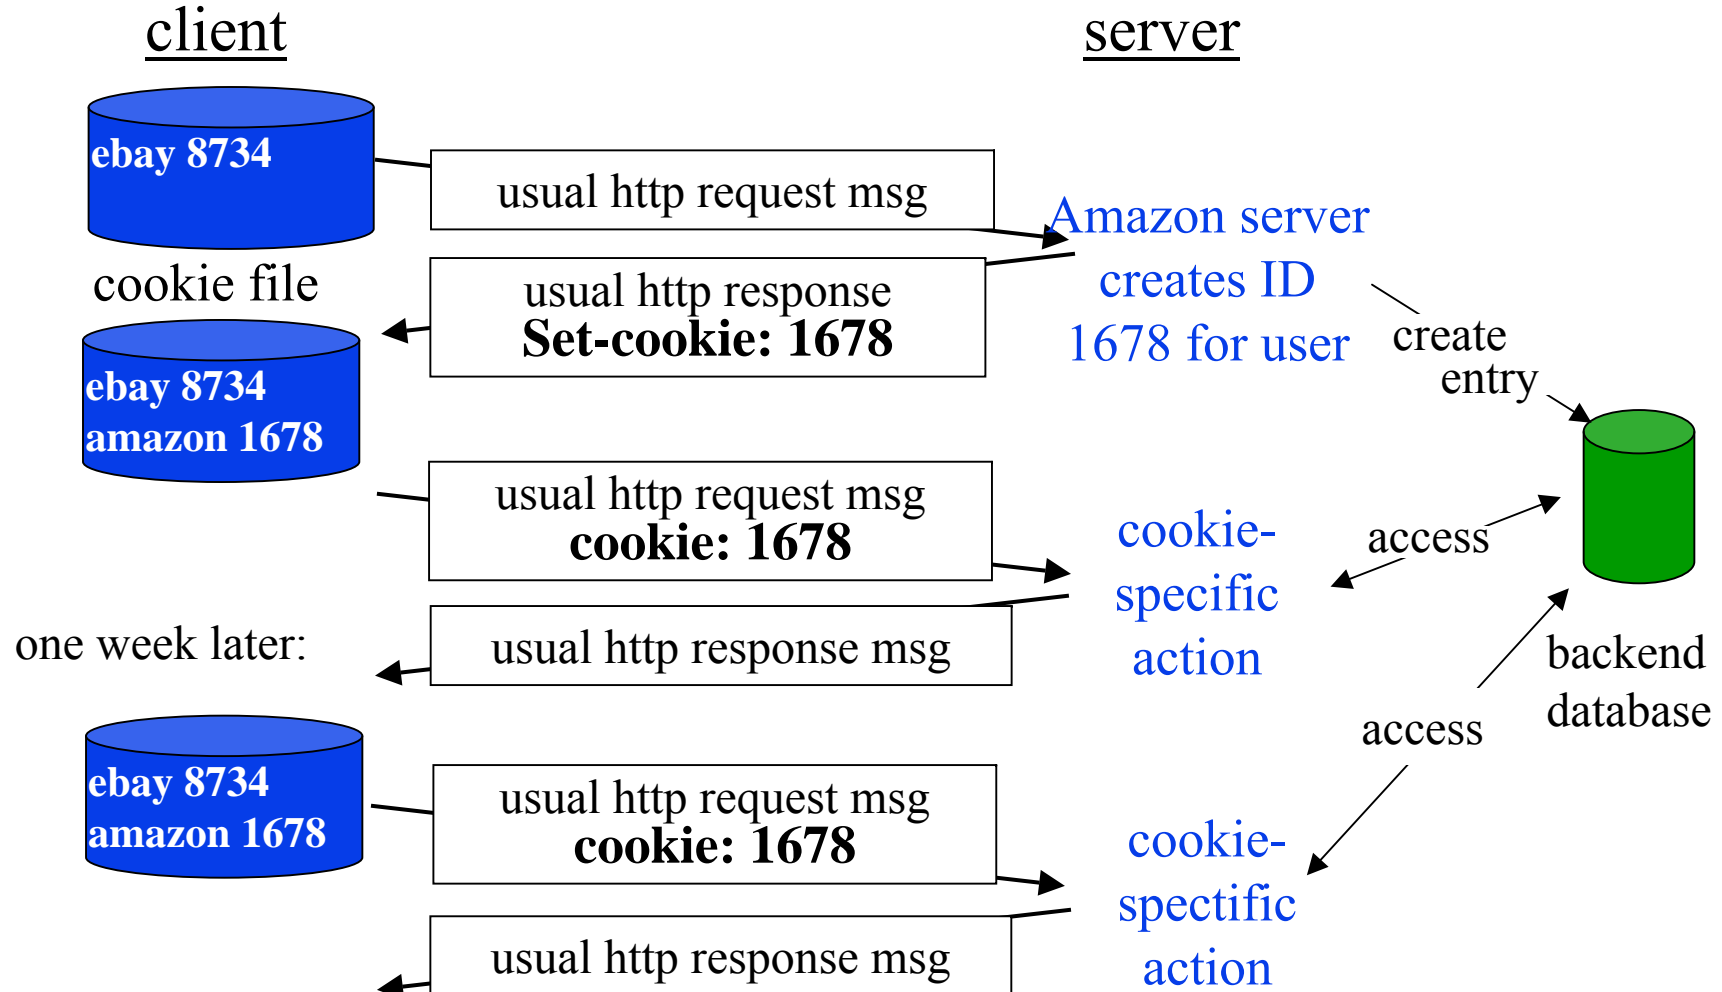
The requested URL /~jain/cse473-08/sample.htm was not found on this server.

Apache/2.0.52 (CentOS) Server at www.cse.wustl.edu Port 80

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>404 Not Found</title>
</head><body>
<h1>Not Found</h1>
<p>The requested URL /~jain/cse473-08/sample.htm was not found on this server.</p>
<hr>
<address>Apache/2.0.52 (CentOS) Server at www.cse.wustl.edu Port 80</address>
</body></html>
```


Cookies

- Allow servers to remember previous information
c:\documents and settings\raj jain\cookies\

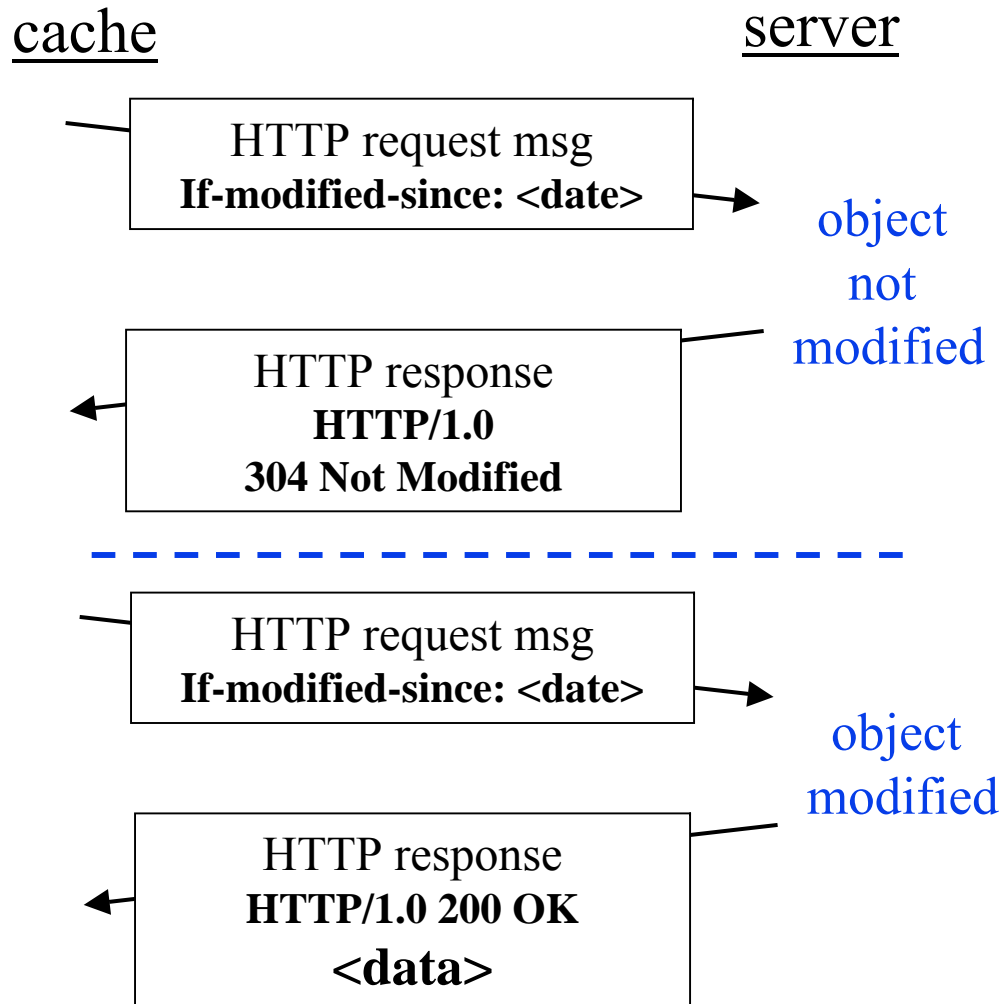



Proxy Server: Web Caching

- ❑ All requests are sent to proxy server
- ❑ Proxy server caches objects
- ❑ Only new objects are requested from origin server
- ❑ Fast, Lower traffic on the link



Conditional GET





HTTP: Summary

1. HTTP is a client-server protocol.
Uses text-based messages
2. Web pages are generally written in HTML
3. HTTP uses non-persistent/persistent TCP connections
4. Cookies allow servers to maintain state
5. Proxy servers improve performance by caching frequently used pages
6. Conditional gets allows proxy servers to reduce Internet traffic

Homework 2A

P5. The text below shows the reply sent from the server in response to the HTTP GET message. Answer the following questions, indicating where in the message below you find the answer.

HTTP/1.1 200 OK

Date: Tue, 07 Mar 2010 12:39:45GMT

Server: Apache/2.0.52 (Fedor)

Last-Modified: Sat, 5 Jan 2007 18:27:46 GMT

Etag: "526c3-f22-a88a4c80"

Accept-ranges: bytes

Content-Length: 4060

Keep-Alive: timeout=max=100

Connection: Keep-Alive

Content-Type: text/html; charset=ISO-8859-1

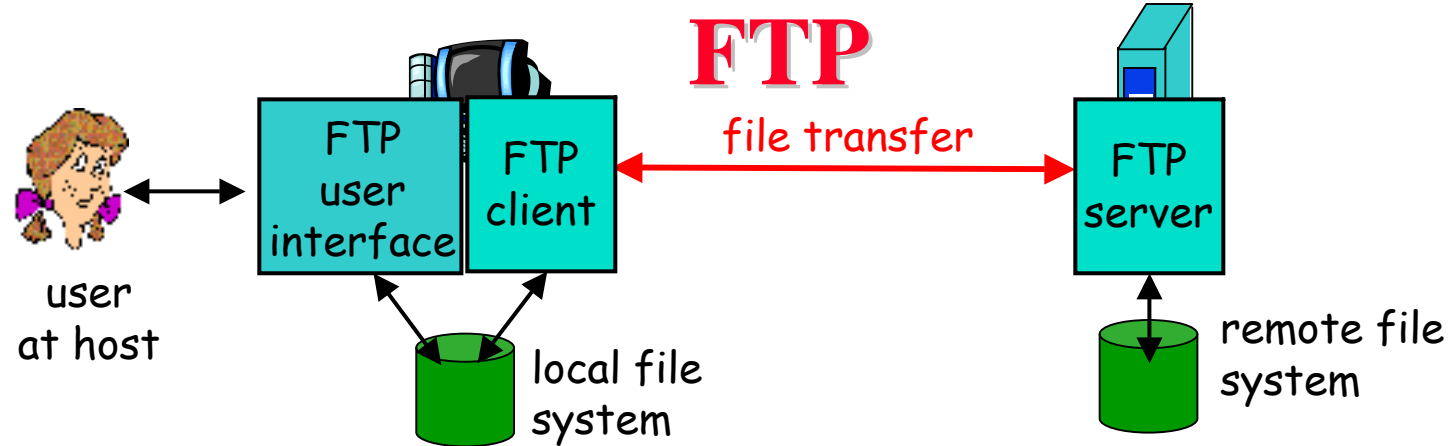
```
<!doctype html publi "-//w3c//dtd html 4.0 transitional//en">
  <html>
  <head>
  <much more document text following here (not shown)>
```

- A. Was the server able to successfully find the document or not? What time was the document reply provided?
- B. When was the document last modified?
- C. How many bytes are there in the document being returned?
- D. What are the first 5 bytes of the document being returned? Did the server agree to a persistent connection?



File Transfer and Electronic Mail

1. File Transfer Protocol (FTP) Commands and Replies
2. Simple Mail Transfer Protocol (SMTP)
3. HTTP vs. SMTP
4. Mail Access Protocols
 - Post-Office Protocol version 3 (POP3) protocol
 - Internet Mail Access Protocol (IMAP)

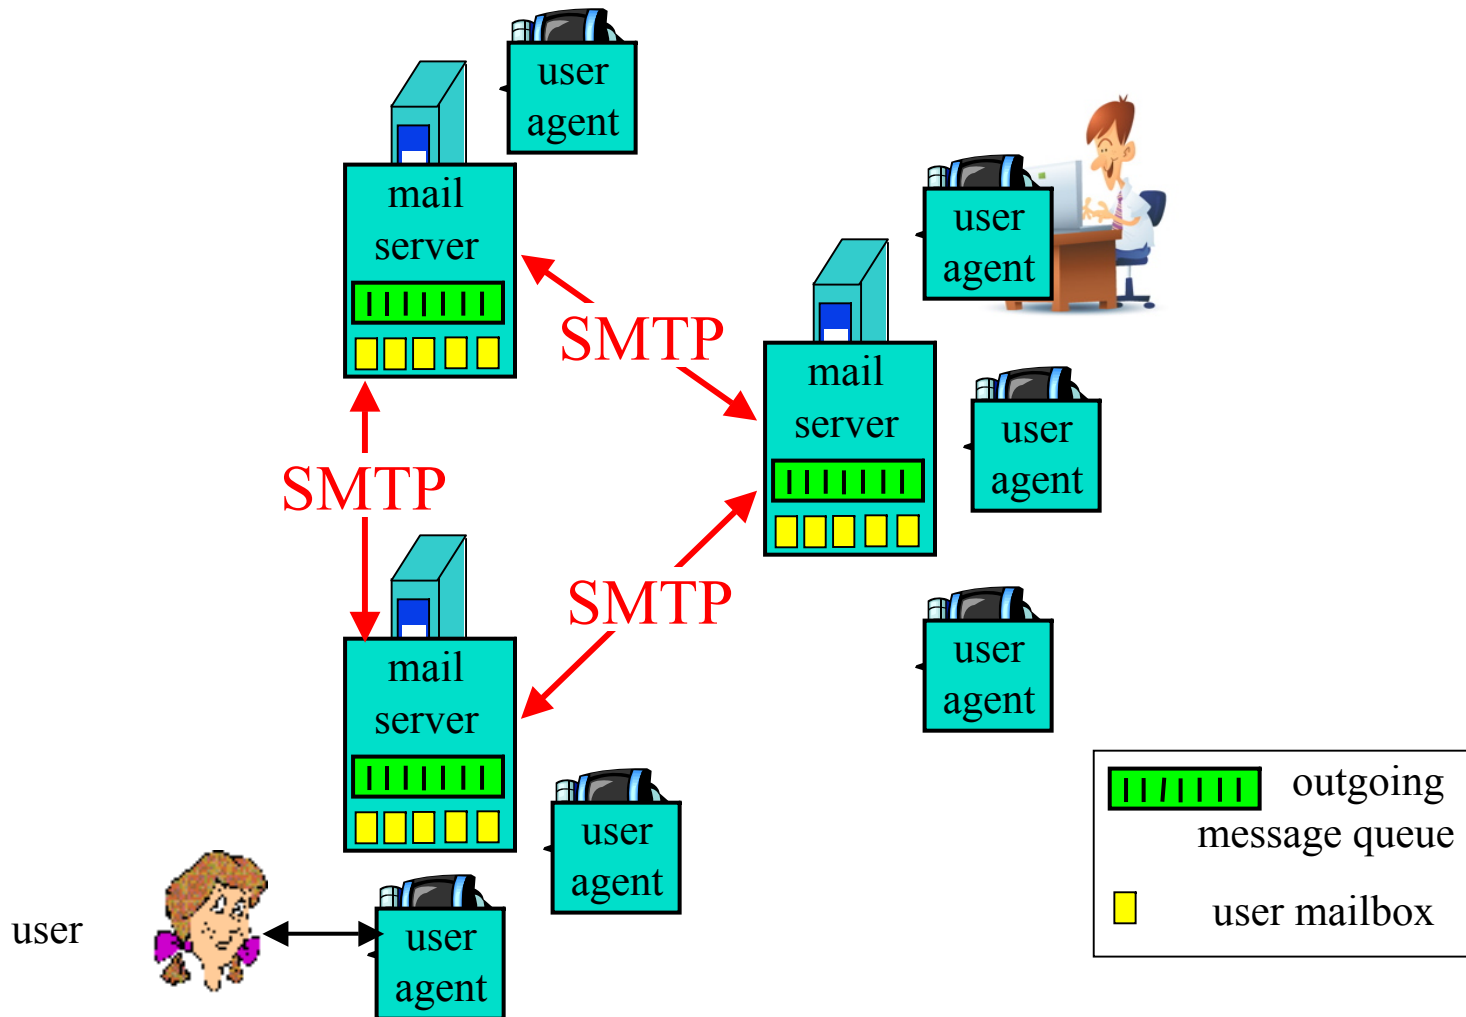


- ❑ File transfer protocol. Uses Port 21
- ❑ Uses two parallel TCP connections: Control and Data
- ❑ Control is out-of-band
- ❑ HTTP uses in-band control
- ❑ Control connection is persistent
- ❑ Data connections are non-persistent: A new TCP connection is opened for each data file
- ❑ Stateful: Server keeps track of user name, directory, ...

FTP Commands and Replies

- ❑ **USER** username
- ❑ **PASS** password
- ❑ **LIST**
- ❑ **RETR** filename
- ❑ **STOR** filename
- ❑ **Status:**
 - ❑ 331 Username OK, password required
 - ❑ 125 Data connection already open; transfer starting
 - ❑ 425 Con't open data connection
 - ❑ 452 Error writing file

Electronic Mail



SMTP

- ❑ Simple Mail Transfer Protocol
- ❑ Old Protocol: Allows only 7-bit ASCII messages
- ❑ All binary objects have to be converted to ASCII
- ❑ Uses port 25 at the server

Sample SMTP Exchange

S: 220 wustl.edu

C: HELO acm.org

S: 250 Hello acm.org, pleased to meet you

C: MAIL FROM: <jain@acm.org>

S: 250 jain@acm.org... Sender ok

C: RCPT TO: <jain@wustl.edu>

S: 250 jain@wustl.edu ... Recipient ok

C: DATA

S: 354 Enter mail, end with "." on a line by itself

C: This is a test email.

C: This serves as an example for class.

C: .

S: 250 Message accepted for delivery

C: QUIT

S: 221 wustl.edu closing connection


Try the above client sequence by *telnet wustl.edu 25*

HTTP vs. SMTP

HTTP	SMTP
Persistent/Non-Persistent TCP	Persistent TCP
Mostly Pull	Mostly Push
Accepts binary objects	Accepts only 7-bit ASCII
One Object/response	Multiple objects/message

Mail Access Protocols

- ❑ SMTP can be used to send messages to destination user agent
⇒ Requires destination to be always accessible
- ❑ Post Office Protocol - Version 3 (POP3)
- ❑ Internet Mail Access Protocol (IMAP)
- ❑ HTTP



POP3 protocol

Authorization phase


```
S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on
```

Transaction phase

```
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 2 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off
```

IMAP

- ❑ Internet Mail Access Protocol
- ❑ More sophisticated than POP3
- ❑ Allows users to maintain folders on the server
- ❑ Messages can be moved from one folder to another
- ❑ Users can get only headers or other components of the message
- ❑ Official IMAP site: www.imap.org



FTP and Mail: Summary

1. FTP is a client-server protocol. Uses Port 21.
2. FTP uses two parallel TCP connections for control and data
3. SMTP is the protocol send email
4. SMTP uses only 7-bit ASCII messages
5. POP3, IMAP, or HTTP is used to receive email

Homework 2B

P17. Consider accessing your e-mail with POP3.

A. Suppose you have configured your POP mail client to operate in the download and delete mode. Complete the following transaction:

C: list

S: 1 498

S: 2 912

S: .

C: retr 1

S: blah blah ...

S: ... Blah

S: .

?

?

B. Suppose you have configured your POP mail client to operate in the download and keep mode. Complete the above transaction.

C. Suppose you have configured your POP mail client to operate in the download-and-keep mode. Using your transcript in part (b), suppose you retrieve messages 1 and 2. Exit POP and then five minutes later you again access POP to retrieve new e-mail. Suppose that in the five-minute interval no new message have been sent to you. Provide a transcript of this second POP session.




Domain Name Service

1. DNS Hierarchy
2. How DNS Works?
3. DNS Records
4. DNS Message Format
5. DNS Registration
6. DNS Vulnerability

DNS

- ❑ Domain Name Service
- ❑ DNS servers translate a host name to IP address
E.g., www.wustl.edu ⇒ 128.252.73.216
- ❑ Distributed database of all hosts in the universe
- ❑ Other Services:
 - ❑ **Host Aliasing**: www.rajjain.com or www.cse.wustl.edu/~jain/
 - ❑ **Mail Server Aliasing**: MX record (e.g., jain@wustl.edu)
 - ❑ **Load Distribution**: Multiple addresses, rotated

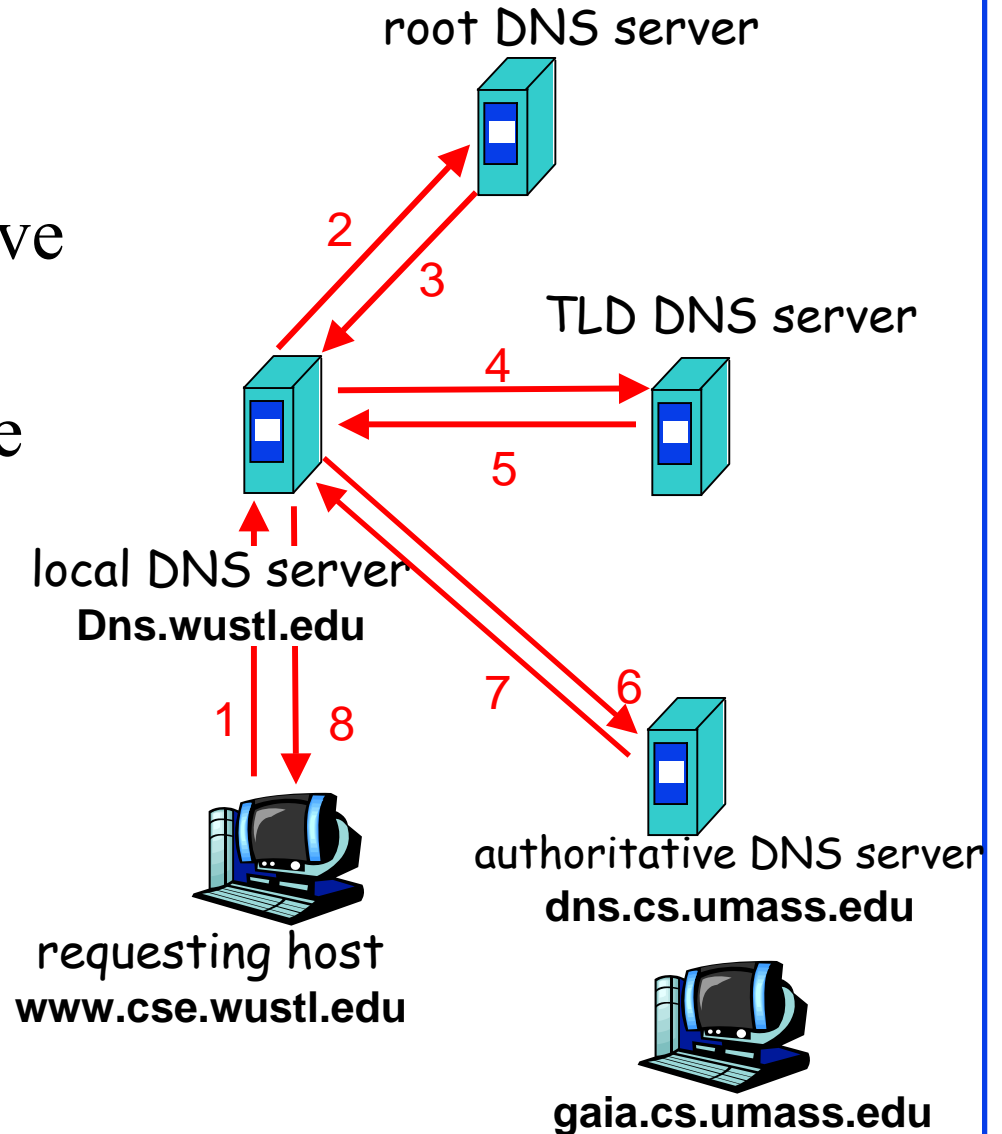
DNS Hierarchy



- ❑ Root DNS Servers
- ❑ Top-level Domain (TLD) servers
- ❑ Authoritative DNS Servers

How DNS Works?

- ❑ Redirects
- ❑ **Recursive queries:** Give me an answer
- ❑ **Iterative queries:** Give me an answer or a hint
- ❑ DNS responses are cached

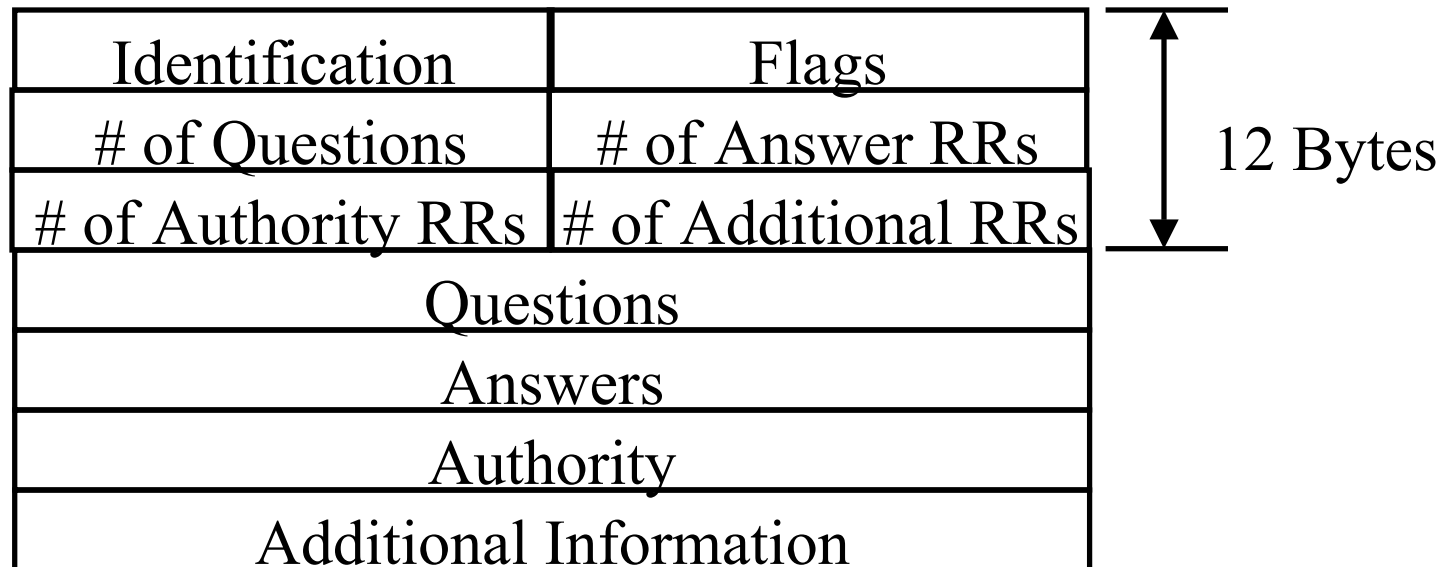


DNS Records

- ❑ Resource Records=(Name, Value, Type, TTL)
- ❑ Type=A: IP Address for the host name
- ❑ Type=NS: Name server for the domain name
- ❑ TYPE=CNAME: Canonical name for a host name
- ❑ Type=MX: Canonical name of mail server

DNS Message Format

- ❑ **Questions:** Name, type
- ❑ **Answers:** Name, type, value, TTL
- ❑ **Authority:** Other authoritative servers
- ❑ **Additional:** Other information, e.g., IP address of canonical name in MX response




DNS Registration

- ❑ Many Registrars
- ❑ Internet Corporation for Assigned Names and Numbers (ICANN) accredits registrars
- ❑ www.internic.net
- ❑ Registrars provide authoritative name servers, A and MX records for the domain

DNS Vulnerability

- ❑ Distributed Denial of service attack on Name server
- ❑ DNS cache poisoning – A server gives wrong answer



DNS: Summary

1. DNS is used to resolve names to IP address
2. Also provides Name aliasing (CNAME), Mail Server (MX) records
3. DNS is a distributed database
⇒ Servers ask other servers for answers when needed
4. Recursive (answer only) or iterative (answer or hint) queries
5. Root Servers, Top level domain servers, Authoritative servers

Homework 2C

R19:

- ❑ Is it possible for an organization's web server and mail server to have exactly the same hostname? (Briefly explain why or why not?)
- ❑ What would be the type of RR that contains the host name of the mail server?




Peer-to-Peer Applications

1. Client Server vs. P2P Scalability
2. P2P File Distribution (BitTorrent)
3. Distributed Hash Tables (DHTs)
4. Peer Churn
5. Network Address Translation (NAT)
6. Skype


P2P File Distribution (BitTorrent)

Tracker: tracks peers participating in torrent

Torrent: group of peers exchanging chunks of a file




Peer-to-Peer Applications



- ❑ $N = \#$ of peers
- ❑ $F =$ File Size
- ❑ $u_i =$ uplink speed of i th host
- ❑ $d_i =$ downlink speed of i th host
- ❑ $d_{\min} = \min \{d_1, d_2, \dots, d_N\}$
- ❑ $D_{cs} \geq \max \{ NF/u_s, F/d_{\min} \}$
- ❑ $D_{P2P} \geq \max \{ F/u_s, F/d_{\min}, NF/(u_s + \sum u_i) \}$

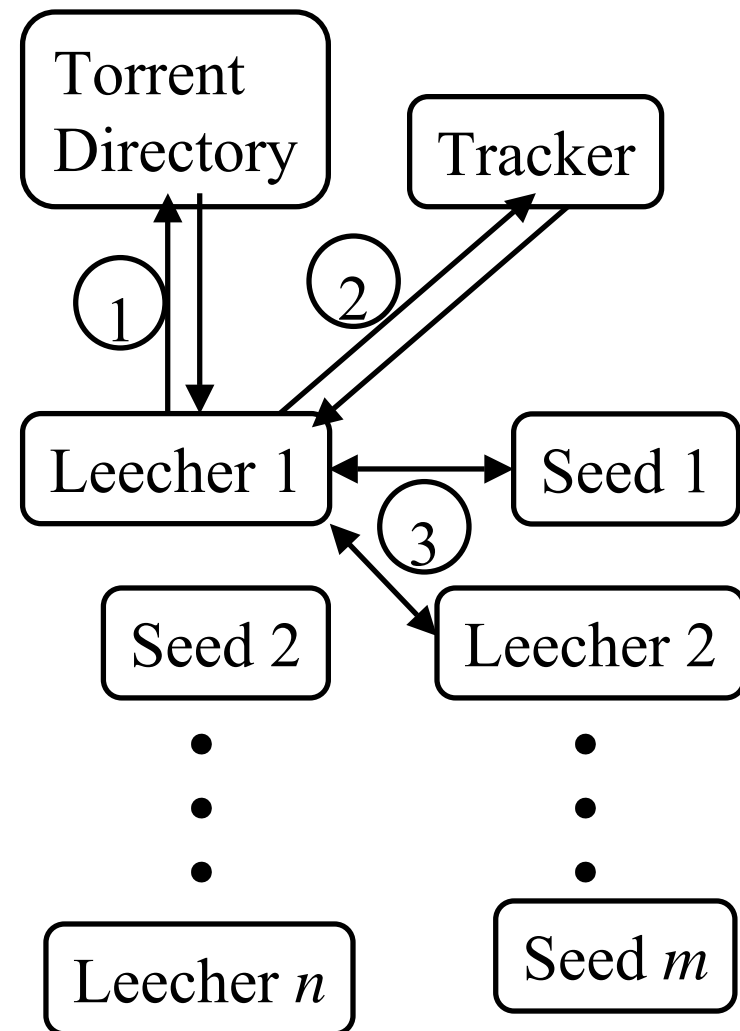
Client Server vs. P2P Scalability

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



BitTorrent P2P File Distribution

- ❑ **Peers**=nodes participating in a file distribution
- ❑ **Torrent**=Set of all peers
- ❑ **Torrent File** =a file containing information about the tracker, object ID, and file
- ❑ Files are segmented into equal size **chunks** (256kB)
- ❑ **Seeds**=Peers that have the complete file
- ❑ **Leechers**=Peers that have incomplete file
- ❑ **Tracker**=Has list of all peers

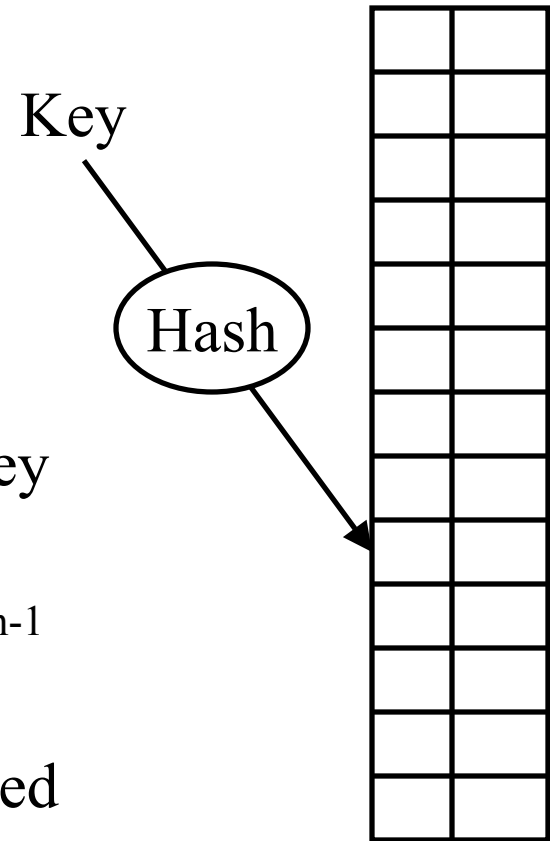


BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contact the tracker to get the current list of peers
Tracker may provide random subset (say 50) peers
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks
 - ❑ Requests least available chunks first (**rarest first**)
 - ❑ Every 10 seconds, Alice calculates the receiving rates
 - ❑ Sends to (**Unchokes**) the top 4 senders
 - ❑ Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
⇒ Helps find high-rate neighbors
 - ❑ Ref: www.bittorrent.org [http://en.wikipedia.org/wiki/BitTorrent_\(protocol\)](http://en.wikipedia.org/wiki/BitTorrent_(protocol))

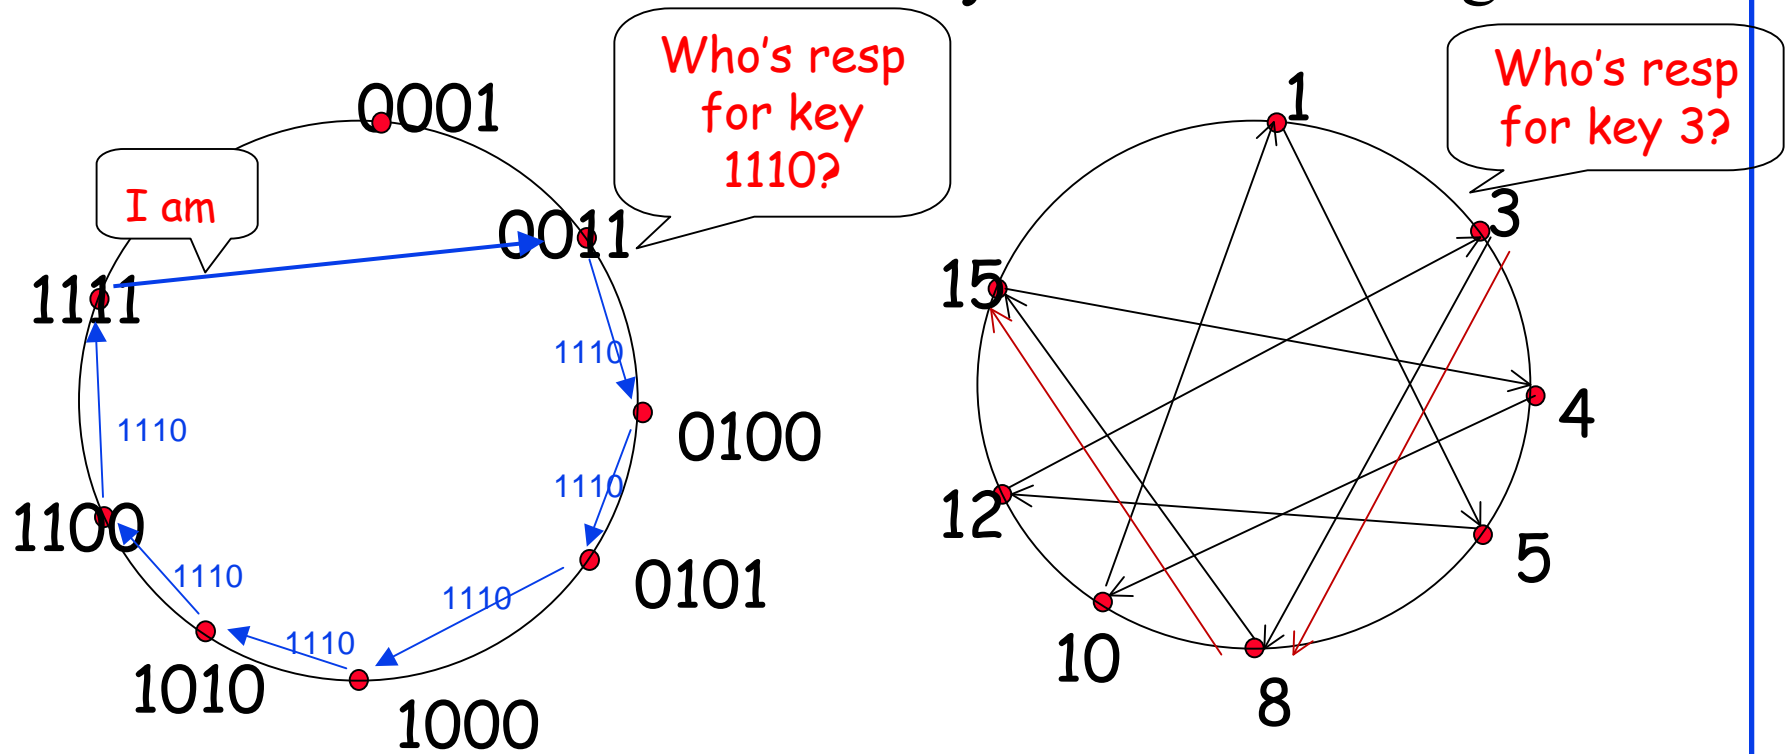
Distributed Hash Tables (DHTs)

- ❑ Hash Tables are used to find an entry in a table of (key, value) pairs in the database
- ❑ Distributed Hash Tables are used for very large databases
- ❑ Database is partitioned and stored in many peers
- ❑ Hash the key to find the peer that has the key
- ❑ Peers are numbered from 0 thru 2^{n-1} (n-bit)
- ❑ Objects (keys) are numbered from 0 thru 2^{n-1} (n-bit)
- ❑ If some peers are missing, the keys are stored at the nearest peer.
- ❑ Ref: http://en.wikipedia.org/wiki/Distributed_hash_table



Circular DHT


- ❑ Circular \Rightarrow $N/2$ messages
- ❑ Shortcuts \Rightarrow Fewer messages
- ❑ Tradeoff between connectivity and # of messages



Peer Churn

- ❑ How to recover when a peer leaves?
- ❑ Each peer keeps track of first and second successors
- ❑ Periodic handshakes to verify connectivity
- ❑ When first successor leaves, connect to the second successor
- ❑ When a new peer joins, it sends a message to peer 1, which gets forwarded until the right predecessor responds

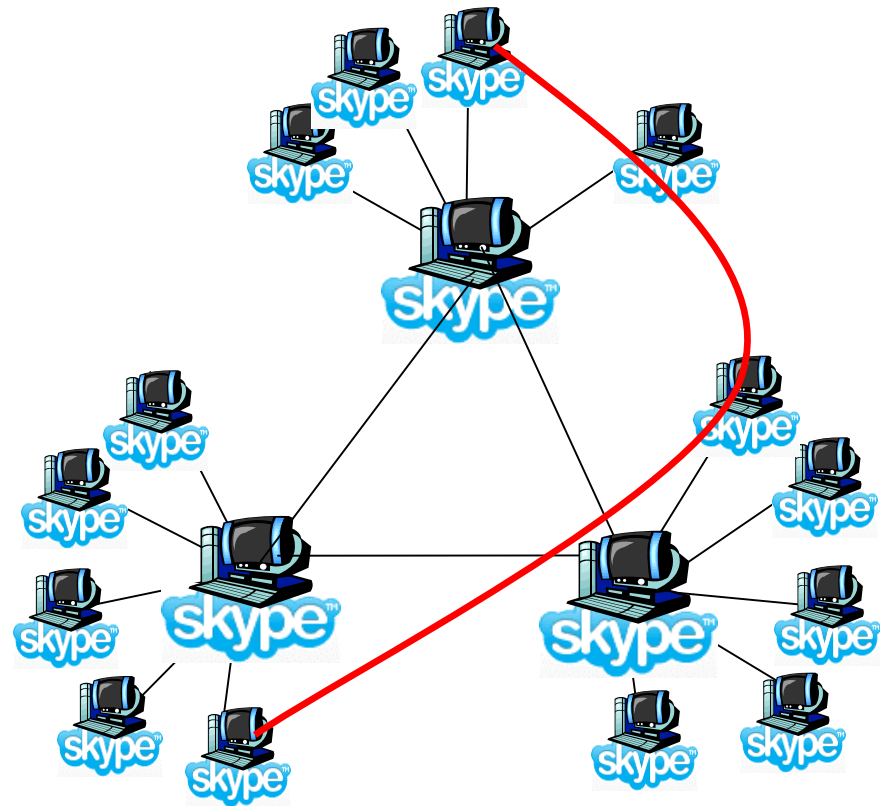
Network Address Translation (NAT)




- ❑ Private IP addresses 192.168.x.x
- ❑ Can be used by anyone inside their networks
- ❑ Cannot be used on the public Internet
- ❑ NAT overwrites source addresses on all outgoing packets and overwrites destination addresses on all incoming packets
- ❑ Only outgoing connections are possible

Skype

- ❑ P2P telephony application
Sky Peer-to-Peer
⇒ Skyper ⇒ Skype
- ❑ Super-peers keep an index (DHT?) of all active Skype users
- ❑ Relays are used to connect NAT'ed nodes






P2P Applications: Summary

1. P2P applications are more scalable
⇒ More efficient when the number of peers is large
2. BitTorrent has peers, trackers, seeds, and leechers
3. BitTorrent unchokes 4 top uploaders and one random node for load balancing
4. Distributed hash tables are used to manage large distributed databases used in P2P applications
5. Skype uses super-nodes to keep track of active users and provides relays for users behind NATs.

Homework 2D

- P26. Suppose Bob joins a BitTorrent torrent, but he does not want to upload any data to any other peers (so called free-riding).
- A. Bob claims that he can receive a complete copy of the file that is shared by the swarm. Is Bob's claim possible? Why or Why not?
 - B. Bob further claims that he can further make his "free-riding" more efficient by using a collection of multiple computers (with distinct IP addresses) in the computer lab in his department. How can he do that?

Application Layer: Summary



1. Applications use TCP/UDP ports for communication
2. HTTP/FTP/SMTP are client-server protocols and use TCP connections
3. HTTP is stateless but cookies allows servers to maintain state
4. Proxy servers improve performance by caching
5. BitTorrent is a P2P file distribution protocol and uses trackers to keep list of peers