

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/cse473-22/>

Student Questions



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.

Student Questions



Network Application Architectures

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

Student Questions

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			

Student Questions


Network Application Architectures

- ❑ Client-Server
- ❑ Peer-to-Peer

Student Questions

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



Student Questions

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 (BitTorrent, eMule, LimeWire), Telephony (Skype)
- ❑ Highly scalable
- ❑ Highly symmetric traffic
⇒ ISP unfriendly
- ❑ Difficult to authenticate ⇒ Insecure
- ❑ Need incentives to share



Student Questions

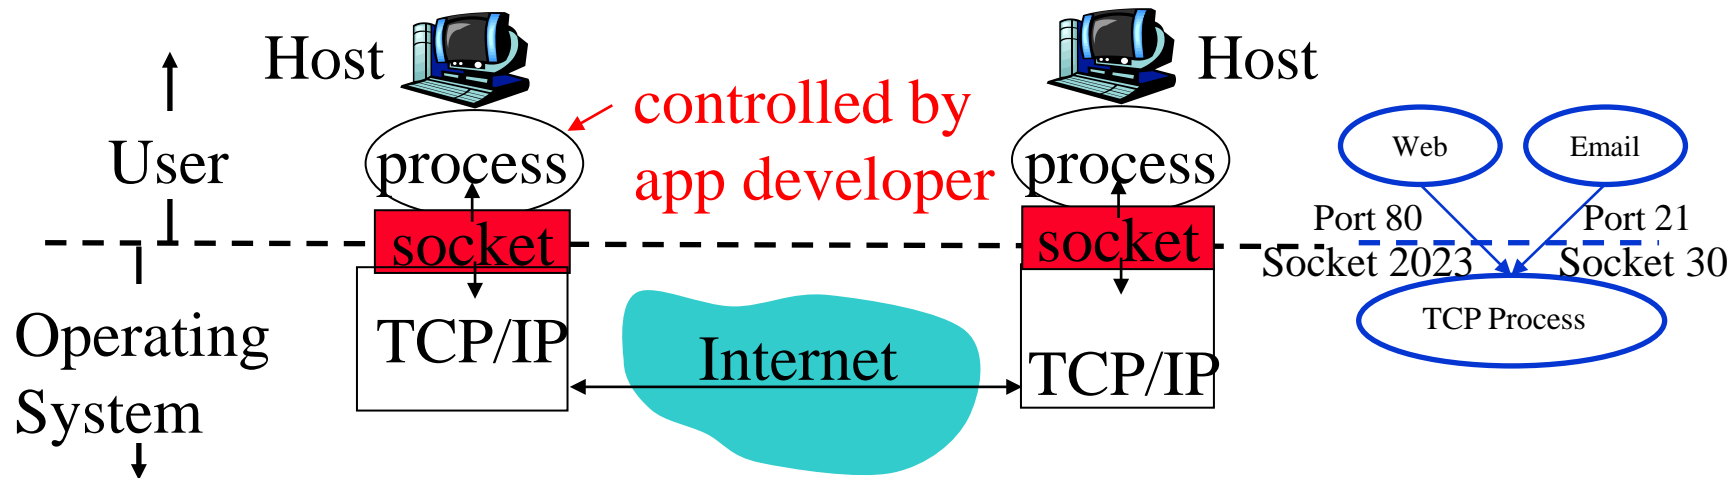
- ❑ Does IPFS p2p protocol solve authenticity problem by hashing the files?
- ❑ When would client-server be more appropriate than P2P.

Don't know much about authenticity in IPFS. Hashing is not an authentication mechanism. It is an integrity mechanism. To be discussed in Chapter 8.

Web server, DNS, Mail Server, etc.

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



Student Questions

- ❑ Are ports for external communication and sockets for communication within the OS?
- ❑ What is the distinction between a port and a socket if any?
- ❑ Can you explain again what ports are and how they are used?

Ports are Networking/TCP concept. Socket is a OS/Unix concept. Ports are standardized so that applications can run any computer. Sockets are not standard. Each process opens sockets and numbers them as needed. For example, port 80 is for Web. Web running on different servers may use different sockets but same port 80.

See above.


- ❑ On page 2-8 where is the port on that picture? In the lecture you pointed to it, but we can't see where you pointed since this is just a screen recording.

A new picture has been added.


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



Student Questions

- ❑ Are sockets only used for IPC? Does each process have its own fixed set of sockets or is it a pool shared by processes? Are ports only used for inter-host communication (as opposed to IPC)?

Yes, sockets are used for inter-process communication. Even non-networking processes need and use sockets. Ports are exclusively for networking. The networking module on the computer knows which port is handled by which process on what socket.

Transports

TCP	UDP
Reliable data transfer	Unreliable Data Transfer
Packet Sequence # required	No Sequence #s
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

Student Questions


- Will TCP wait to transmit packet #6 until receipt of packet #5 is acknowledged?

No. We will discuss the details in Chapter 3.

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 (Class of applications)
- ❑ Skype
- ❑ Each application has its own protocol, message format, semantics of fields

Student Questions



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

Student Questions

Ref: Read Section 2.1 full. Try R1-R9




HTTP

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

Student Questions

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`



Student Questions

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**)

Student Questions

- ❑ Can you explain what Persistent and Non-Persistent are?


Non-Persistent = Connection is closed immediately after a query has been answered. Next query will require a new TCP connection.

- ❑ For persistent HTTP: even if the client has closed their end of the connection, does the server still see the connection as open prior to timeout? If one end closes the connection, do both ends see it? Is it the client or the server that decides whether the TCP connection is persistent or non-persistent?

Closing a TCP connection will be discussed in Chapter 3. Any side can tell the other side and close a connection.

Sample Web Page

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



Raj Jain

Student Questions

Sample HTTP Request Message

GET /~jain/cse473-16/sample.htm HTTP/1.1

Host: www.cse.wustl.edu

Connection: close

User-agent: Mozilla/4.0

Accept-Language: en

- ❑ **Method** = Get
- ❑ **URL** = /~jain/cse473-16/sample.htm
- ❑ **Version** = HTTP/1.1
- ❑ **Header Fields** = Host, Connection, User-agent, ...

Student Questions


- ❑ Does a URL never include the www.cse.wustl.edu (using this slide as an example)?

URL consists of two parts: Host and Directory. These are separately indicated in HTTP. Host is case insensitive. Directory is case sensitive.

URL=Uniform Resource Locator.

Here directory is the URL/location on the host.

HTTP Request Message Format



Student Questions

- Why hasn't there been a push to remove legacy characters such as CR and LF from the HTTP Messaging Format?

You can manually type the whole request using a keyboard. This is a feature not a bug.

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

Student Questions

HTTP Response Message Format



Student Questions

Hands-on HTTP

```
telnet www1.cse.wustl.edu 80
GET /~jain/cse473-19/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu
```

```
HTTP/1.1 200 OK
Date: Tue, 13 Sep 2019 23:39:53 GMT
Server: Apache/2.2.3 (CentOS)
Accept-Ranges: bytes
Content-Length: 233
Connection: close
Content-Type: text/html; charset=ISO-8859-1
```

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

This is a sample text.

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

- Where does the content length start and end? Does the content length only refer to the requested data or also include the rest of the response message?

“Content Length” does not include anything other than the content (data).

- What does "Connection: close" mean
Please close the TCP connection after responding to this request.

Hands-on HTTP (cont)

```
telnet www1.cse.wustl.edu 80
GET /~jain/cse473-08/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu
```

```
HTTP/1.1 404 Not Found
Date: Tue, 13 Sep 2019 23:42:48 GMT
Server: Apache/2.2.3 (CentOS)
Content-Length: 307
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

```
<!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/2.2.3 (CentOS) Server at www1.cse.wustl.edu Port 80</address>
</body></html>
```

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

Student Questions

- Is there any way to prevent/catch small typos in the url from returning 404, such as redirecting or checking capitalization?

Yes. But, this is not a part of the protocol. This can be done by the application clients (Firefox, Chrome,...). The directory part of the URL is case sensitive on Linux and Case insensitive on Windows. People make money by buying misspelt domain names, e.g., citibnk.com ⇒ Credit Alert Systems

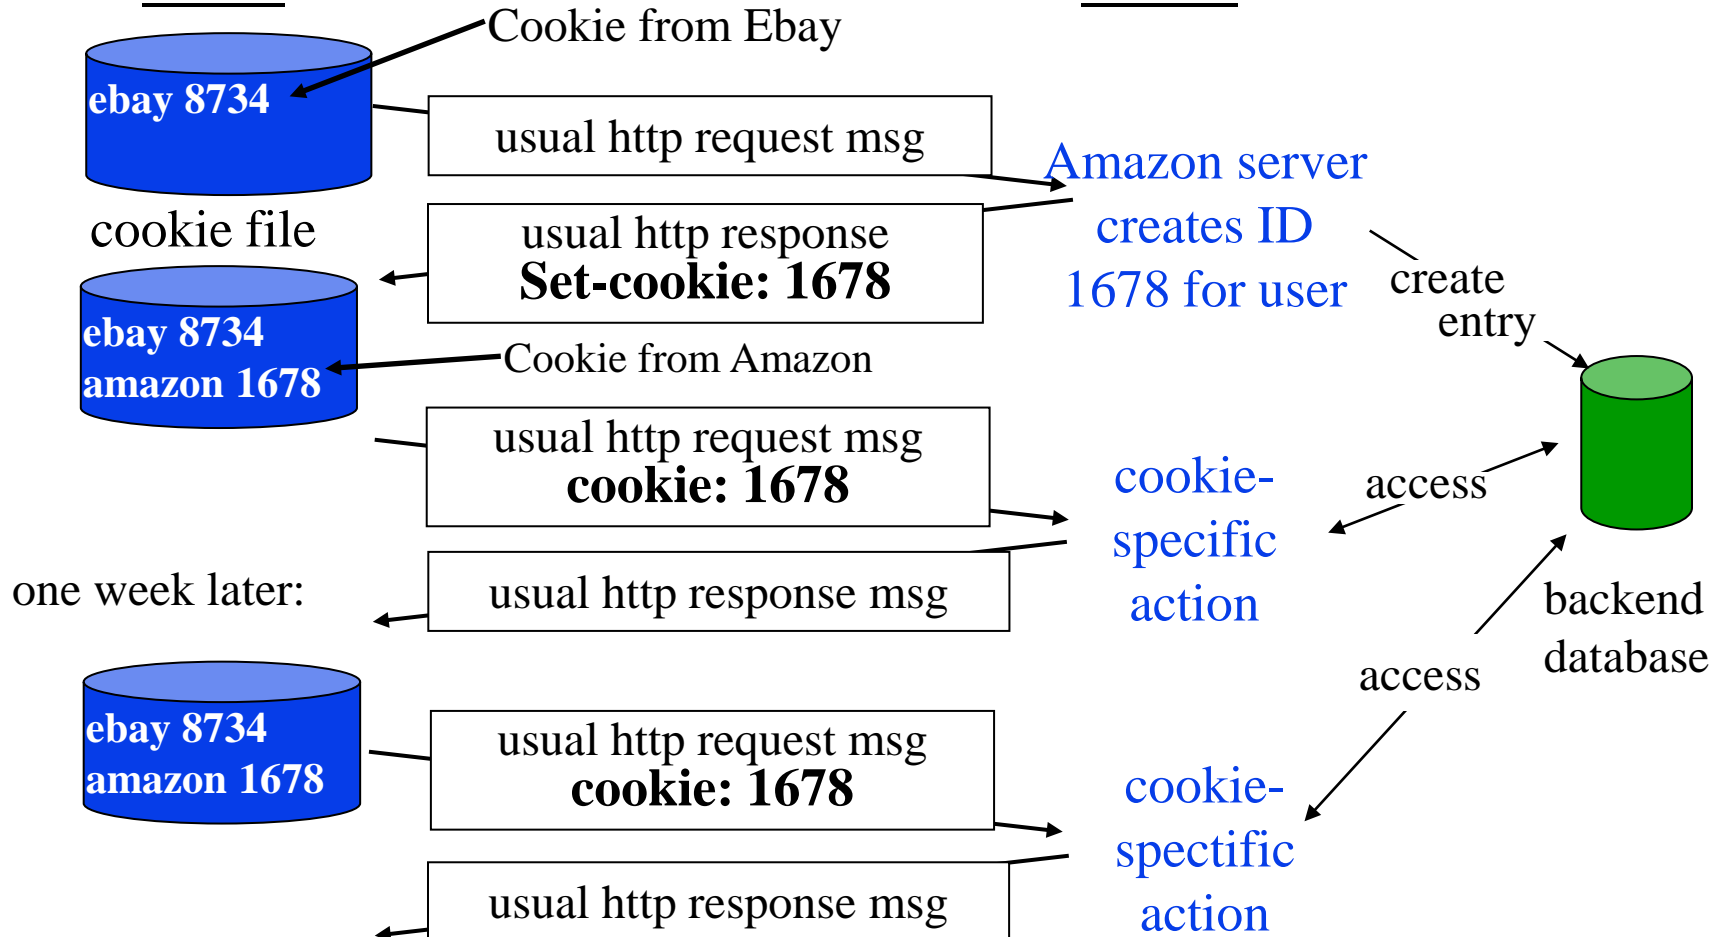
Cookies

- Allow servers to remember previous information

C:\Documents and Settings\Raj Jain\Cookies\

client

server




Student Questions

- Can only the server determine when it wants to delete a cookie, and are cookies stored on the browser level or website level?

Cookies are stored on your computer. Cookies are used by servers as one factor to identify you. But they are misused by advertisers to track you. You can delete them anytime or set browser option to not store them. But then you may have to do two-factor authentication on some servers.

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




Student Questions

- ❑ Is there an effect or significance to URLs starting with 'www'?
- WWW is for convenience. HTTP does not require it. You can use any name for your webserver.
<http://raj.jain.com> is a valid URL for a web server.*

Conditional GET



Student Questions



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

Ref: **Read Section 2.2 Full. Try R10-R14.**

Student Questions

- Are conditional GETs only used with proxy servers?

This is not necessary.

Homework 2A: HTTP

[10 points] 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 2019 14:39:45GMT

Server: Apache/2.0.52 (Fedor)

Last-Modified: Sat, 5 Jan 2019 19:27:46 GMT

Etag: "526c3-f22-a88a4c80"

Accept-ranges: bytes

Content-Length: 5071

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?
- E. Did the server agree to a persistent connection?

Student Questions

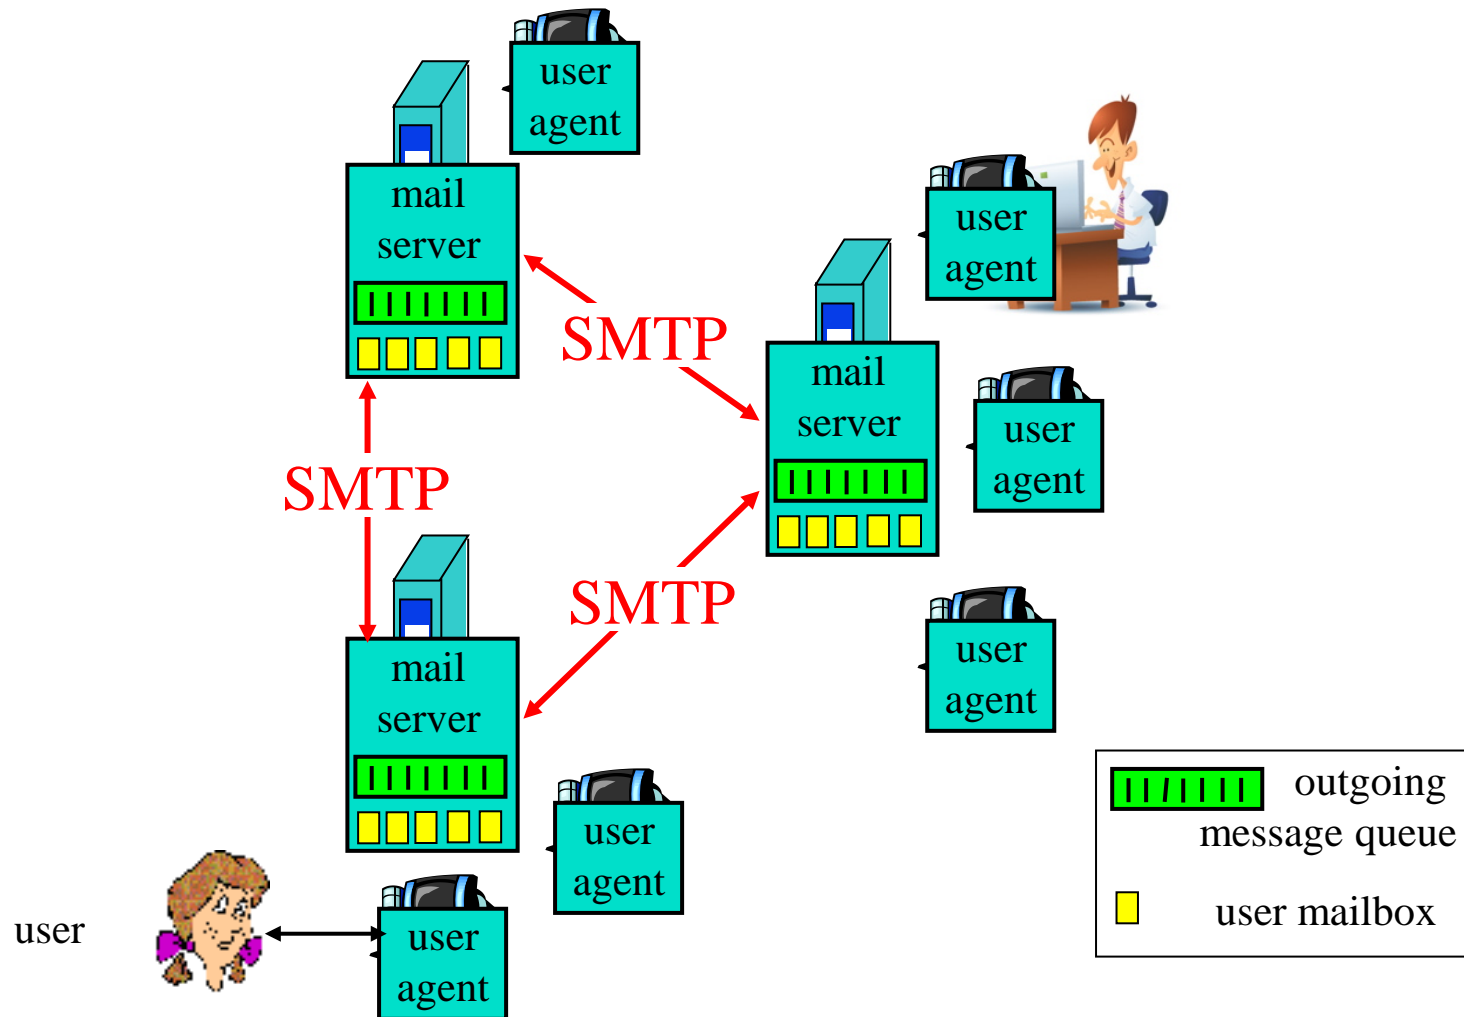
Lab 2A: Domains

[10 points] Submit answers for the following: (See hints in the parenthesis.)

1. Find the IP addresses of www.google.com and www.yahoo.com (ping)
2. Modify the hosts file to map www.google.com to yahoo's IP address and ping to www.google.com. Notice what address it is pinging to. Remove the modification to the host file, open a new command window and repeat. (Windows: c:\windows\system32\drivers\etc\hosts)
3. Find the domain name and country of 128.252.165.7 (<http://www.webyield.net/domainquery.html>)
4. Find the owner of wustl.edu domain (<http://www.networksolutions.com/whois/index.jsp>)
5. Find the name server of wustl.edu domain (<http://www.networksolutions.com/whois/index.jsp>)

Student Questions

Electronic Mail



Student Questions

- How does SMTP know how to deliver mail to the correct server?

Mail servers use DNS to find the destination mail server. DNS is discussed later in this chapter.

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

Student Questions

- ❑ Why hasn't SMTP been upgraded to accept more than 7 bit ASCII messages?

There are mappings that allow representing all characters in all languages to ASCII, e.g., UTF-8.

Sample SMTP Exchange

C:telnet mail.seas.wustl.edu 25

S: 220 POSTOFFICE.seas.wustl.edu Microsoft ESMTP MAIL Service, Version: 6.0.3790.46
75 ready at Tue, 13 Sep 2011 18:34:56 -0500

C:HELO acm.org

S: 250 POSTOFFICE.seas.wustl.edu Hello [128.252.19.232]

C:MAIL FROM: jain@acm.org

S: 250 2.1.0 jain@acm.org....Sender OK

C:RCPT TO: jain@wustl.edu

S: 250 2.1.5 jain@wustl.edu

C:DATA

S: 354 Start mail input; end with <CRLF>.<CRLF>

C:This is test email.

This serves as an exmaple for CS473 class.

S: 250 2.6.0 <MAIL2j97vPYGrN7kf0V00000aff@POSTOFFICE.seas.wustl.edu> Queued mail
for delivery

C:QUIT

S: 221 2.0.0 POSTOFFICE.seas.wustl.edu Service closing transmission channel

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

- Do attached images also go through SMTP protocol or require some other protocol?

Bit strings can also be mapped as ASCII.

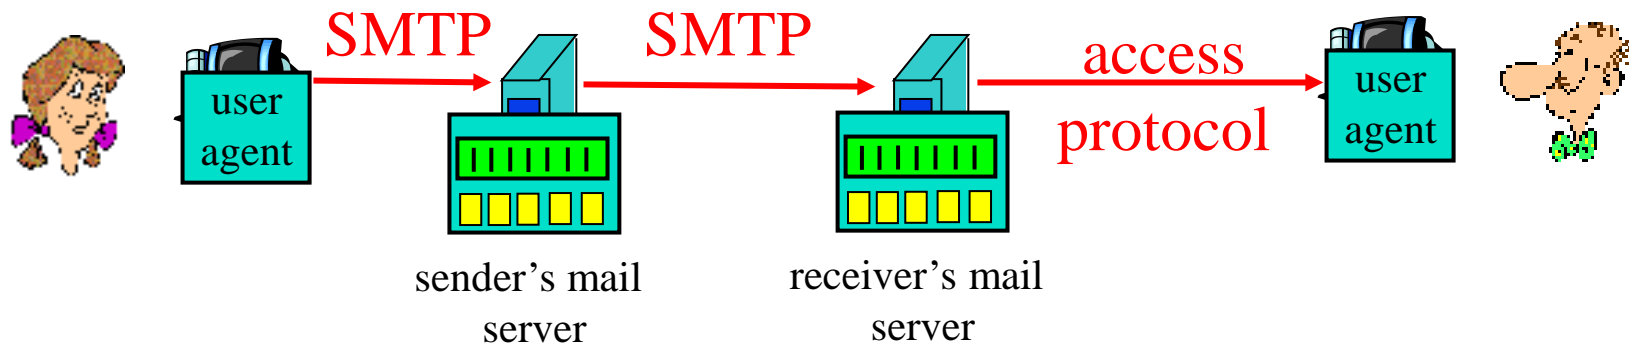
HTTP vs. SMTP

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

Student Questions

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



Student Questions

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
```

Student Questions

- Are there any common ways of cross protocol communication.


Not common. But yes. See

https://en.wikipedia.org/wiki/Cross-layer_optimization

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

Student Questions



Mail: Summary

1. SMTP is the protocol to **send** email
2. SMTP uses only **7-bit ASCII** messages
3. POP3, IMAP, or HTTP is used to **receive** email

Student Questions

- Why are SMTP TCP connections persistent? Shouldn't email messages only just send the one item?

SMTP sets up a TCP connection, sends all pending emails, and then disconnects. Each mail is often several megabytes and so it is not a good idea to disconnect after each segment. Therefore, the connections are persistent but they are not on when there is no email to send.

Ref: Read Section 2.3 Full. Try R15-R20

Homework 2B: Mail

[12 points] Consider accessing your e-mail with POP3.

a) Suppose you have configured your POP mail client to operate in the download and keep mode.

Complete the following transaction to retrieve both messages, and sign off. Show the complete sequence of messages. (Fill in ? and successive messages)

```
C: list
S: 1 500
S: 2 901
S: .
C: retr 1
S: blah blah ...
S: ... Blah
S: .
?
?
```

b) Repeat part a if you have programmed your POP client in download and delete mode.

c) Suppose 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 for both options a and b above.

Student Questions



Domain Name Service

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

Student Questions

DNS

- ❑ Domain Name Service
- ❑ DNS servers translate a host name to IP address
E.g., www.wustl.edu ⇒ 128.252.87.149
- ❑ 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

Student Questions

- ❑ Does host Aliasing save different URLs that direct to the same website?

Host aliasing simply provides the same IP address for different names.

DNS Example

F:\>**nslookup www.wustl.edu**

Server: ns00.ip.wustl.edu

Address: 128.252.0.1

Name: www.wustl.edu

Address: 128.252.87.149

F:\>**nslookup www.google.com**

Server: ns00.ip.wustl.edu

Address: 128.252.0.1

Non-authoritative answer:

Name: www.l.google.com

Addresses: 74.125.225.48, 74.125.225.52, 74.125.225.50, 74.125.225.49
74.125.225.51


Aliases: www.google.com

Student Questions

- ❑ Is 'www' the hostname of the DNS server?
No. www is usually the hostname of Web server. It could be anything else too.
- ❑ Can you explain the difference between sockets and ports again? Are they on different layers?
Sockets are in the operating system. Ports are for transport layer.
- ❑ What part is the www in www.facebook or wustl in wustl.instructure.com? is it another child authoritative server below the facebook or instructure authoritative server?

These are names in those domains. Wustl.instructure.com is a name in instructure.com. Several names may map to a single server at instructure.com

DNS Hierarchy



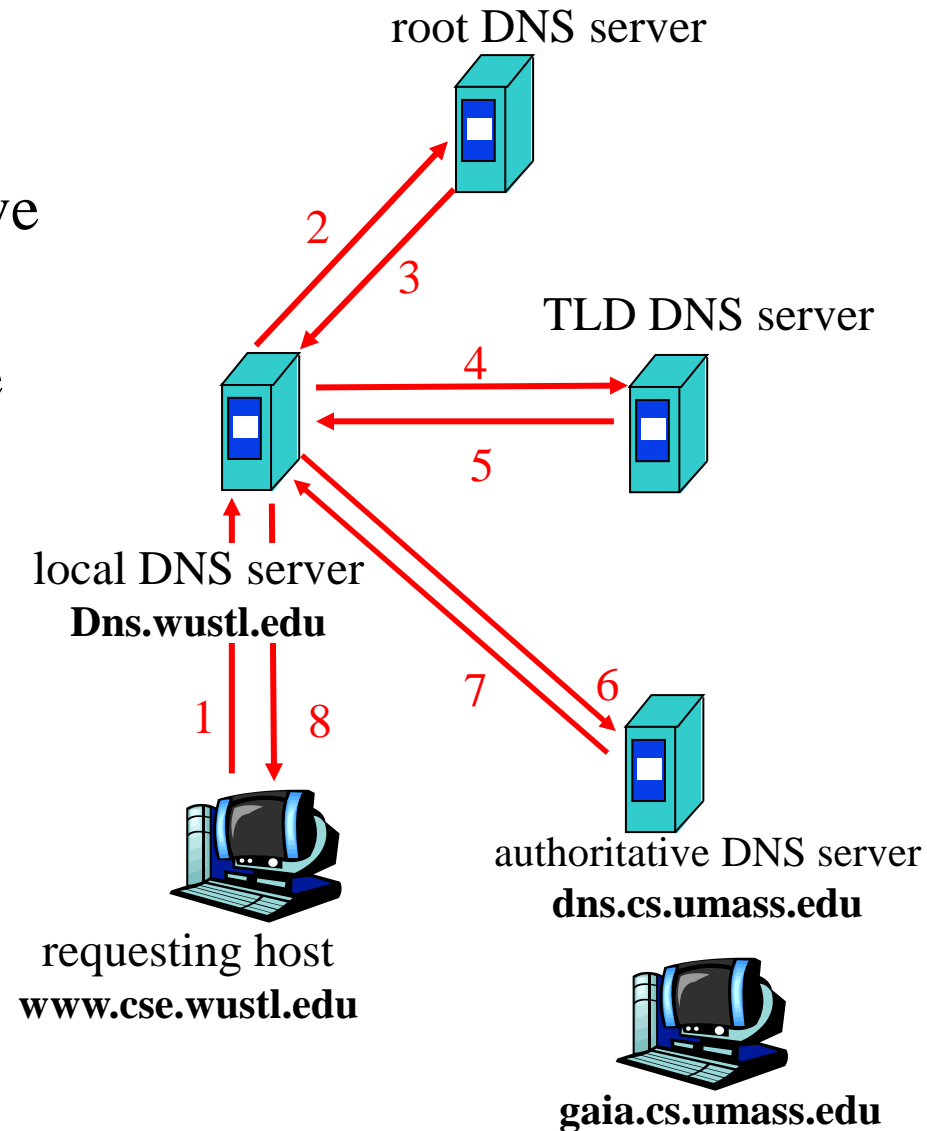
- ❑ Root DNS Servers
- ❑ Top-level Domain (TLD) servers
- ❑ Authoritative DNS Servers

Student Questions

- ❑ What is the "seas." prefix that's discussed?
SEAS=School of Engineering and Applied Science at WUSTL. Now renamed to McKelvey School of Engineering.
- ❑ Why is the authoritative DNS server umass? How does that relate to WUSTL?
Umass will be authoritative server for names ending in Umass.edu and not for wustl.edu.
- ❑ Do TLDs actually change anything about how the protocol or data is managed? Like, are .gov and .edu servers managed differently?
No. They manage more names and so need more computer capacity.

How DNS Works?

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



Student Questions

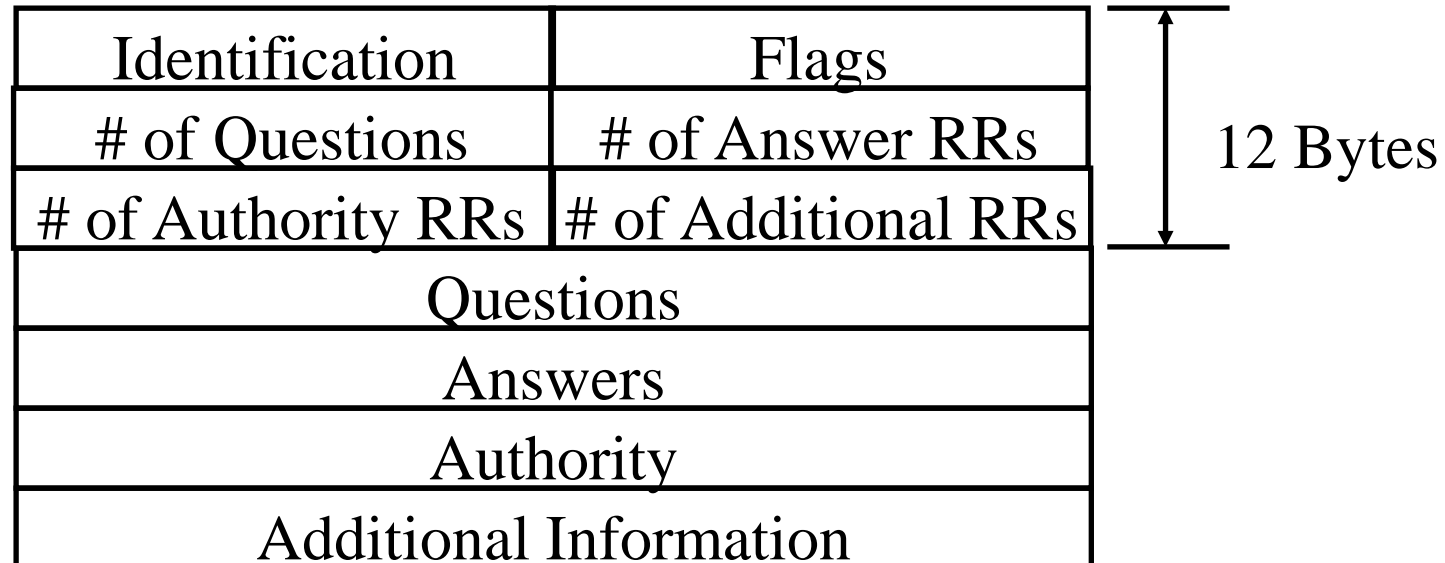
DNS Records

- ❑ 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

Student Questions

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



Student Questions

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


Student Questions

- ❑ Can you re-explain DNS cache poisoning?
Some server gives wrong answers to a DNS request. The client sends its packets to wrong address until it decides to check the address again.
-

DNS Vulnerability

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

Student Questions



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

Student Questions

Ref: **Read Section 2.4 Full.**

Homework 2C: DNS

- ❑ [4 points]
- ❑ 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?

Student Questions



Peer-to-Peer Applications

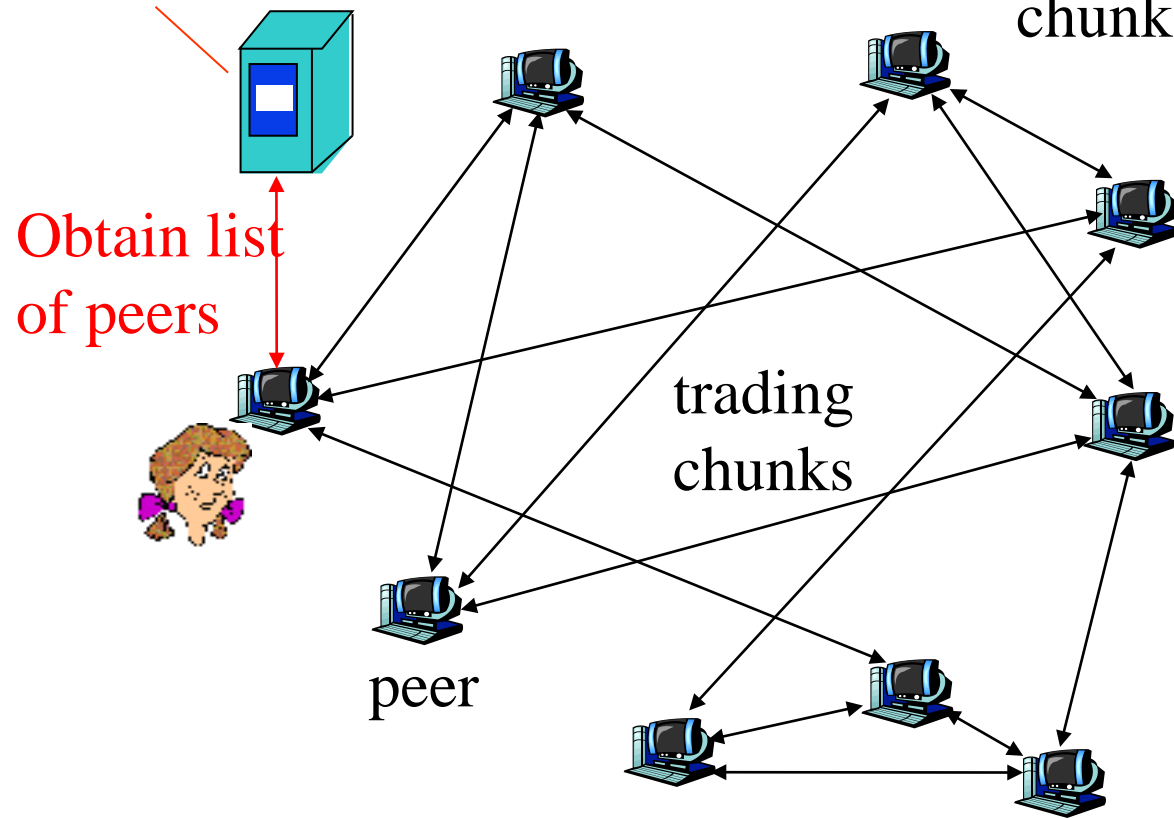
1. Client Server vs. P2P Scalability
2. P2P File Distribution (BitTorrent)

Student Questions

P2P File Distribution (BitTorrent)

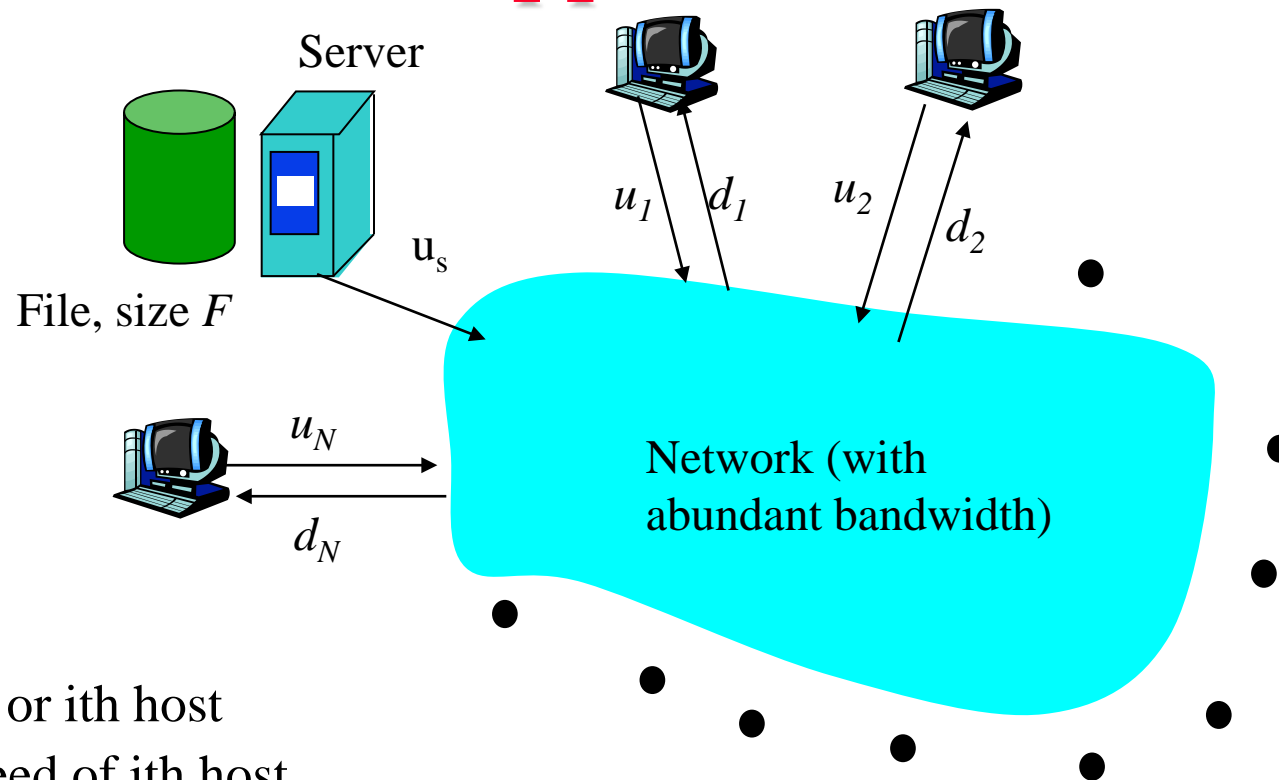
Tracker: tracks peers participating in torrent

Torrent: group of peers exchanging chunks of a file



Student Questions

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
- $d_{\min} = \min\{d_1, d_2, \dots, d_N\}$
- $D_{cs} \geq \max\{NF/u_s, F/d_{\min}\}$
- $D_{P2P} \geq \max\{F/u_s, F/d_{\min}, NF/(u_s + \sum u_i)\}$

Server, Client, Network


Student Questions

- Are we required to know Asymptotic Complexities of Server Distribution Algorithms vs P2P Algorithms?

Yes. These formulas should be on your cheat sheet.

Client Server vs. P2P Scalability

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

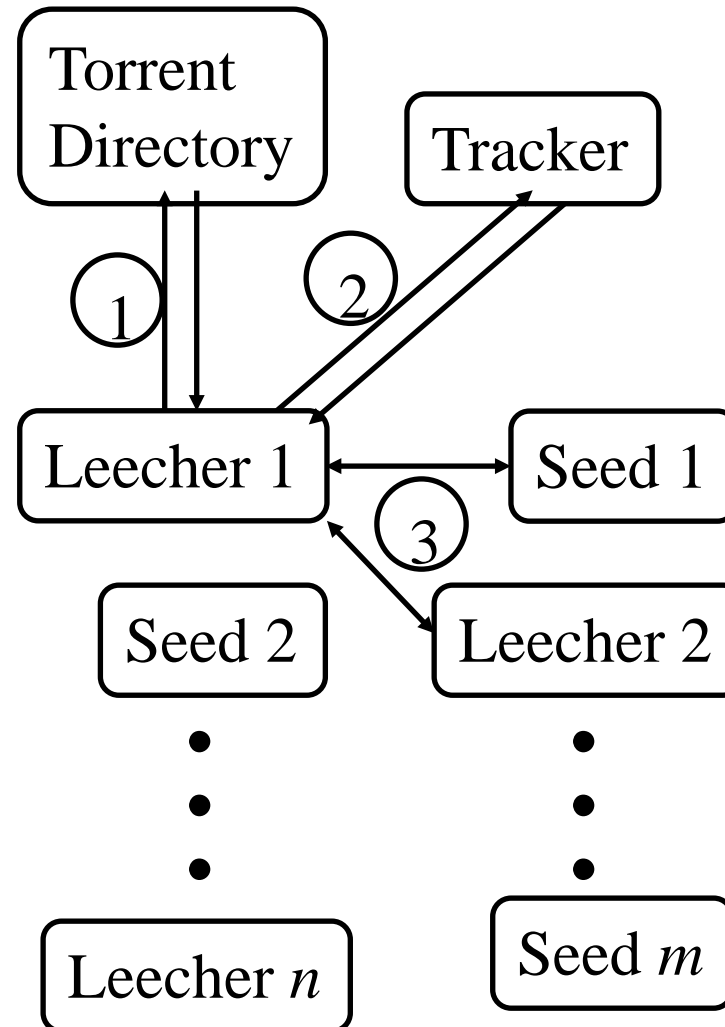


Student Questions

- Can you go over scenarios in which it would be advantageous to use a client/server design rather than P2P?

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



Student Questions


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)

Student Questions

- ❑ Can you explain unchoking/choking?

Choking = Not sending the packets



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**

Ref: Read Section 2.5 full. Try R21-R23.

Student Questions

Homework 2D: P2P

[4 points] 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?

Student Questions

Streaming Video

- ❑ Video traffic is 80% of consumer traffic
- ❑ Video: 25-30 Frames/sec
- ❑ Video can be compressed:
 - Spatial: next pixel is similar to this
 - Temporal: Pixel in the next frame is similar to this
- ❑ Variable bit rate (VBR)/Constant bit rate (CBR)
 - Motion Picture Expert Group (MPEG) 1: 1.5 Mbps
 - MPEG2: 3-6 Mbps
 - MPEG4 (.mp4): Less than 1 Mbps



Student Questions

- ❑ What is the difference between VBR and CBR?
The bit rate depends on the changes in the frame and therefore, on motion. So the choice is either to keep decrease the quality (CBR) or to increase the bit rate (VBR) during motion.

Ref: Cisco Visual Networking Index: Forecast and Methodology, 2014-2019 White Paper,

http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html


Washington University in St. Louis

<http://www.cse.wustl.edu/~jain/cse473-22/>

©2022 Raj Jain

Dynamic Adaptive Streaming over HTTP (DASH)

- ❑ DASH provides an efficient method for video streaming
- ❑ Standard Web Servers: No changes required to servers, Content Distribution Networks (CDN), or HTTP protocol.
- ❑ Mobile client controls what is downloaded using a “**media presentation description (MPD)**” file defined by DASH
- ❑ MPD contains URLs for segments
- ❑ Client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.



Student Questions

- ❑ For the DASH protocol, I'm confused about what benefit using multiple URLs brings. Homework 2E implies that every segment of every video quality has a distinct URL, but wouldn't a single URL per video quality be enough?

Page 149 of the textbook states:

"With DASH, each video version is stored in the HTTP server, each with a different URL..."


The client then selects one chunk at a time by specifying a URL and a byte range in an HTTP GET

Each video is stored in multiple rates and each copy is partitioned in multiple segments. Each segment needs a URL.

of URLs = # of rates × # of segments/rate

Content Distribution Networks (CDN)

- ❑ To reduce latency to worldwide users, the data is replicated at many sites
- ❑ Users are directed to nearby site by DNS
- ❑ [netflix.com](https://www.netflix.com) -> cdn_stl.com or cdn_sfo.com, ...




Student Questions

Homework 2E: DASH

- [2 points] A DASH system stores video at 8 different qualities (rates) and 15 minute segments. How many URLs will be required for a 2-hour movie?

Student Questions

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
6. **DASH** allows clients to request different video segments as needed
7. **CDN**'s directs users to to nearby copy via DNS

Ref: In addition to previous readings, read Sections 2.6.1-2.6.3. Try R24-R25.

Student Questions

Lab 2B: UDP Pinger

- ❑ [50 points] In this lab, you will learn the basics of socket programming for UDP in Python. You will learn how to send and receive datagram packets using UDP sockets and also, how to set a proper socket timeout. Throughout the lab, you will gain familiarity with a Ping application and its usefulness in computing statistics such as packet loss rate.
- ❑ You will first study a simple Internet ping server written in the Python, and implement a corresponding client. The functionality provided by these programs is similar to the functionality provided by standard ping programs available in modern operating systems. However, these programs use a simpler protocol, UDP, rather than the standard Internet Control Message Protocol (ICMP) to communicate with each other. The ping protocol allows a client machine to send a packet of data to a remote machine, and have the remote machine return the data back to the client unchanged (an action referred to as echoing). Among other uses, the ping protocol allows hosts to determine round-trip times to other machines.
- ❑ You are given the complete code for the Ping server below. Your task is to write the Ping client.

Student Questions

Lab 2B (Cont)

Server Code

The following code fully implements a ping server. You need to compile and run this code before running your client program. *You do not need to modify this code.*

In this server code, 30% of the client's packets are simulated to be lost. You should study this code carefully, as it will help you write your ping client.

```
# UDPPingerServer.py
# We will need the following module to generate randomized lost packets
import random
from socket import *
# Create a UDP socket
# Notice the use of SOCK_DGRAM for UDP packets
serverSocket = socket(AF_INET, SOCK_DGRAM)
# Assign IP address and port number to socket
serverSocket.bind(('', 12000))
```

Student Questions

Lab 2B (Cont)

while True:

Generate random number in the range of 0 to 10

rand = random.randint(0, 10)

Receive the client packet along with the address it is coming from
message, address = serverSocket.recvfrom(1024)

Capitalize the message from the client

message = message.upper()

If rand is less is than 4, we consider the packet lost and do not respond

if rand < 4:

continue

Otherwise, the server responds

serverSocket.sendto(message, address)

The server sits in an infinite loop listening for incoming UDP packets. When a packet comes in and if a randomized integer is greater than or equal to 4, the server simply capitalizes the encapsulated data and sends it back to the client.

Student Questions

Lab 2B (Cont)

Packet Loss

UDP provides applications with an unreliable transport service. Messages may get lost in the network due to router queue overflows, faulty hardware or some other reasons. Because packet loss is rare or even non-existent in typical campus networks, the server in this lab injects artificial loss to simulate the effects of network packet loss. The server creates a variable randomized integer which determines whether a particular incoming packet is lost or not.

Client Code

You need to implement the following client program.

The client should send 10 pings to the server. Because UDP is an unreliable protocol, a packet sent from the client to the server may be lost in the network, or vice versa. For this reason, the client cannot wait indefinitely for a reply to a ping message. You should get the client wait up to one second for a reply; if no reply is received within one second, your client program should assume that the packet was lost during transmission across the network. You will need to look up the Python documentation to find out how to set the timeout value on a datagram socket.

Student Questions

Lab 2B (Cont)

Specifically, your client program should

- (1) send the ping message using UDP (Note: Unlike TCP, you do not need to establish a connection first, since UDP is a connectionless protocol.)
- (2) print the response message from server, if any
- (3) calculate and print the round trip time (RTT), in seconds, of each packet, if server responses
- (4) otherwise, print “Request timed out”

During development, you should run the UDPPingerServer.py on your machine, and test your client by sending packets to *localhost* (or, **127.0.0.1**). After you have fully debugged your code, you should see how your application communicates across the network with the ping server and ping client running on different machines.

Message Format

The ping messages in this lab are formatted in a simple way. The client message is one line, consisting of ASCII characters in the following format:

Ping *sequence_number* *time*

where *sequence_number* starts at 1 and progresses to 10 for each successive ping message sent by the client, and *time* is the time when the client sends the message.

Student Questions

Lab 2B (Cont)

What to Hand in

You will hand in the complete client code and screenshots at the client verifying that your ping program works as required.

Student Questions

Reading List

- ❑ Read Chapter 3 of the textbook for the next lecture.

Student Questions

Acronyms

- ❑ ASCII American Standard Code for Information Interchange
- ❑ CBR Constant bit rate
- ❑ CDN Content Distribution Network
- ❑ DASH Dynamic Adaptive Streaming
- ❑ DNS Domain Name System
- ❑ FTP File Transfer Protocol
- ❑ GMT Greenwich Mean Time
- ❑ HTML Hyper-Text Markup Language
- ❑ HTTP Hyper-Text Transfer Protocol
- ❑ ICANN International Corporation for Assigned Names and Numbers
- ❑ ID Identifier
- ❑ IMAP Internet Message Access Protocol
- ❑ IP Internet Protocol
- ❑ ISO International Standards Organization
- ❑ ISP Internet Service Provider
- ❑ kB Kilo Byte

Student Questions

Acronyms (Cont)

- ❑ MPD Media Presentation Description
- ❑ MPEG Moving Picture Expert Group
- ❑ NAT Network Address Translator
- ❑ NS Name Service
- ❑ PC Personal Computer
- ❑ POP Point of Presence
- ❑ RR Resource Record
- ❑ SMTP Simple Mail Transfer Protocol
- ❑ TCP Transmission Control Protocol
- ❑ TLD Top Level Domain
- ❑ TTL Time to Live
- ❑ UDP Universal Data Protocol
- ❑ URL Uniform Resource Locator
- ❑ VBR Variable bit rate

Student Questions

Scan This to Download These Slides



Raj Jain

<http://rajjain.com>

Student Questions

http://www.cse.wustl.edu/~jain/cse473-22/i_2app.htm

Related Modules



CSE 567: The Art of Computer Systems Performance Analysis
https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),
https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcg5e_10TiDw



CSE 570: Recent Advances in Networking (Spring 2013)
<https://www.youtube.com/playlist?list=PLjGG94etKypLHyBN8mOgwJLHD2FFIMGq5>

CSE571S: Network Security (Spring 2011),
<https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u>



Video Podcasts of Prof. Raj Jain's Lectures,
<https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw>

Student Questions

MIME Encoding

From: "Saved by Internet Explorer 11"
Subject:
Date: Wed, 3 Feb 2021 12:08:58 -0600
MIME-Version: 1.0
Content-Type: multipart/related;
 type="text/html";
 boundary="-----_NextPart_000_0000_01D6FA25.59EF7E90"
X-MimeOLE: Produced By Microsoft MimeOLE

This is a multi-part message in **MIME** format.

-----_NextPart_000_0000_01D6FA25.59EF7E90

Content-Type: text/html;
 charset="Windows-1252"
Content-Transfer-Encoding: quoted-printable
Content-Location: file:///D:/u/jain.htm

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">
<HTML><HEAD><META content=3D"IE=3D5.0000" =
http-equiv=3D"X-UA-Compatible">

<META http-equiv=3D"Content-Type" content=3D"text/html; =
charset=3Dwindows-1252">
<META name=3D"GENERATOR" content=3D"MSHTML 11.00.10570.1001"></HEAD>=20
<BODY><IMG src=3D"file:///D:/u/jain.jpg"> <BR>Raj Jain </BODY></HTML>
```

Washington University in St. Louis

<http://www.cse.wustl.edu/~jain/cse473-22/>

©2022 Raj Jain

Student Questions

*MIME = Multipurpose
Internet Mail Extensions
MHTML = MIME Encoded
HTML*



MIME Encoding (Cont)

-----=_NextPart_000_0000_01D6FA25.59EF7E90

Content-Type: image/jpeg

Content-Transfer-Encoding: **base64**

Content-Location: file:///D:/u/**jain.jpg**

/9j/4AAQSkZJRgABAQEAYABgAAD/2wBDAAoHBwgHBgoICAgLCgoLDhgQDg0NDh0VFhEYIx8lJCIf
liEmKzcvJik0KSEiMEEhNDk7Pj4+JS5ESUM8SDc9Pjv/2wBDAQoLCw4NDhwQEbw7KCloOzs7Ozs7
Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozv/wAARCACWAGYDICA
AhEBAxEB/8QAGwAAAQUBAQAAAAAAAAAAAAAAAAABgACAwQFAQf/xAA3EAABAwIFAQYCCQDAAAAAAAAAB
AAIDBBEFBhIhMUETIIFhcYGRsRQjMjNCUqHB0Rvi4fAHNFP/xAAZAQADAQEBAAAAAAAAAAAAAAAAAAA
AQMEAgX/xAAiEQACAgICAgIDAAAAAAAAAAAAAAAAAQIRAyEEMRJBeyIUMmH/2gAMA5QAAhEDEQA/AK11

.....
0zhMzuPdcSDp6rQgZdhdYEHghOMYvpIuPMJUkbaeU6S7S7h19h7KRQY8Fs19PeHCuUcvZVUM4P3b
D3Wk/uLg+6Kc5SXqaMX5Dv2QkKB+IYfiFG5319O/6TE//wBBax+Q91r1lccToMHnbyPhJd6iwJ+
IKb/AF2cV9qRXgLo3gg7cLQhDnEaXb+XRUnAM72nU09FagmDC022U0M26bFZ6ezJ7ys8fxD+VrMn
jnj7SJ4cCh0EEBzRcLjJ5KWXESL8joV0AQE7pKrT1bapmoCzhy3wSTEBjagmudDb7LNNTT5qYu2A
aLW3v6pJKZQjcdZueCo2m5HTdJJJaJK6ok/psVGx7mtlmOq3UBpNvksmW0H1cY07bkJLjJ2jdxk
vBv+mtgL5mNebmQGMn+1wt/HwUD6V+GYjFhxk1tpqcAEeJJcfmkkqQf0MvISWTRba8OdpI2sugN
hfYgljr2APCSSEQZfw6QvYWngHb0KkmAvb4JLLoQ6lkdDd7TvwUkkl0I/9k=

-----=_NextPart_000_0000_01D6FA25.59EF7E90--

Student Questions