

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/cse473-23/>

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

- Why Ethernet and point-to-point are paratactic?

All protocols of the same layer are paratactic.

Paratactic=Placed next to each other

Network Application Architectures

- ❑ Client-Server
- ❑ Peer-to-Peer


Student Questions

- ❑ Can one connection consist of both client-server and peer-to-peer links? As in one hop is client-server while the next is peer-to-peer.

The client is served by the server. Peer-to-peer does not imply any service. Both client-server and peer-to-peer connections can be 1-to-1, 1-to-many, many-to-1, or many-to-many.

Client-Server

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



Student Questions

- ❑ There is a B/S Architecture. What is the difference between C/S and B/S?
Can I consider B/S as a type of C/S?

B/S=Browser-Server


Browser is a client.

Therefore, yes,

$B/S \subseteq C/S$.

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 the IPFS p2p protocol solve the authenticity problem by hashing the files?

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.

- ❑ When would client-server be more appropriate than P2P?

Web server, DNS, Mail Server, etc.

- ❑ Since web 3.0 is decentralized. Is it going to be a P2P network?

Web3=Web using blockchain

- ❑ Based on my understanding, is it correct to say that the Bitcoin network is an example of a peer-to-peer network?

Bitcoin does not have a network. It uses the Internet. Mining nodes are a kind of server. So not P2P.

- ❑ Did P2P only serve file transfer?

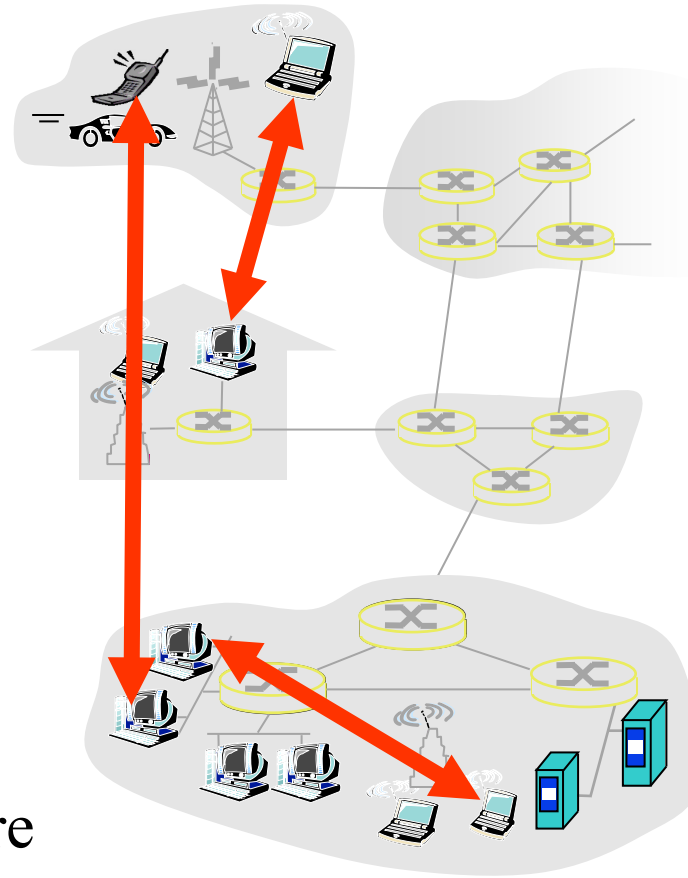
No. You can make almost any other application.

- ❑ Is zoom considered to be P2P or client-to-server?

Client to Server

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



Student Questions

- ❑ Are there any applications that make use of both peer-to-peer and client-server?

File transfer

- ❑ The peer-to-peer seems more efficient than the client-server. Is that true?

Yes. See Slide 2-52

- ❑ How could torrenting and peer-to-peer be made safer?

Using authenticated peers

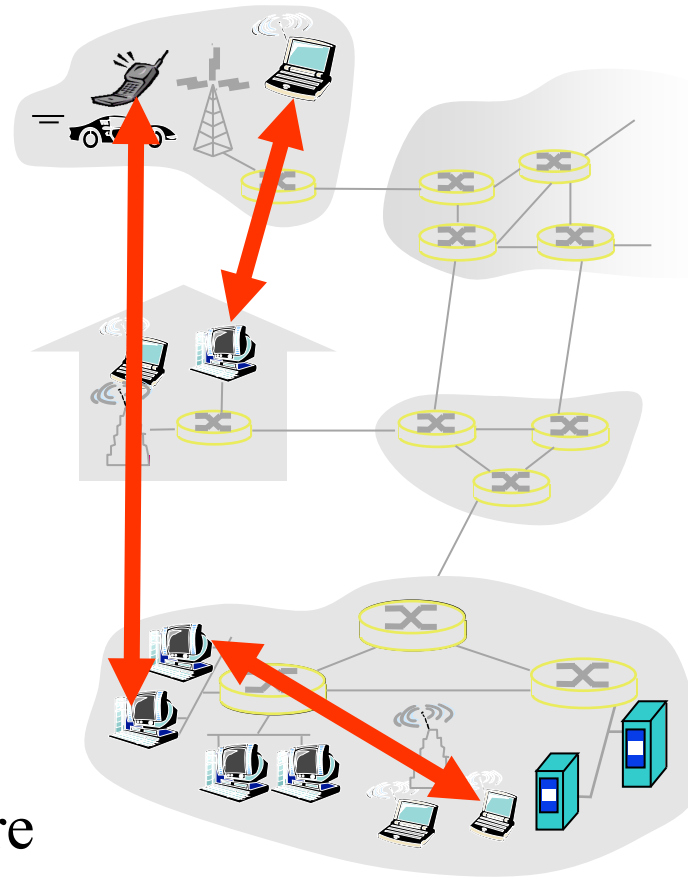
- ❑ Does Peer-to-Peer require a server (to see who is seeding or sending which part of the file etc.)?

The servers are used to discover seeders and leechers, but they do not keep track of the part of the file. If there are no servers, one could discover by broadcast, but that would result in too much traffic.

- ❑ Is it true that P2P does not involve any servers at all?
- ❑ P2P=Anyone can serve anyone.

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



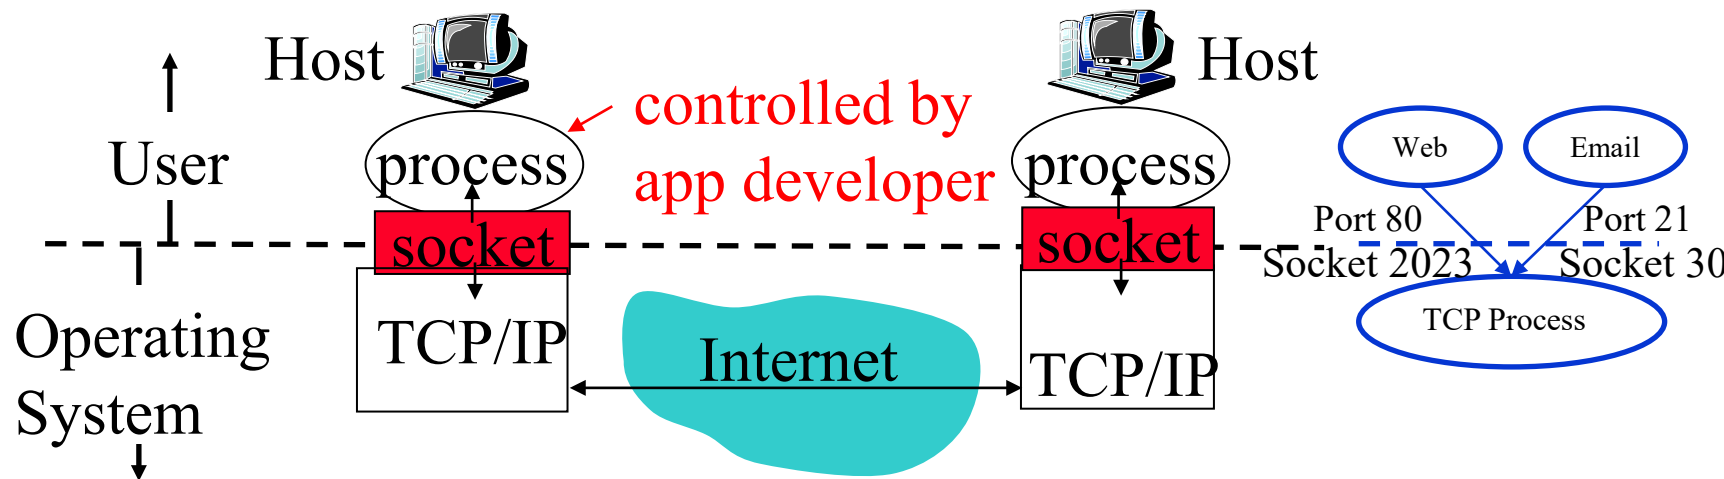
Student Questions

- ❑ Why are ISPs throttling P2P traffic, and is there any way to get around it?

VPN

Process Communications

- ❑ Inter-Process Communication on the Same Host
⇒ Operating system provides message passing.
- ❑ Unix provides an 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 concepts. A socket is an OS/Unix concept. Ports are standardized so that applications can run on any computer. Sockets are not standard. Each process opens sockets and numbers them as needed. For example, port 80 is for the Web. Web running on different servers may use different sockets but the same port, 80.

See above.

- ❑ On pages 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.

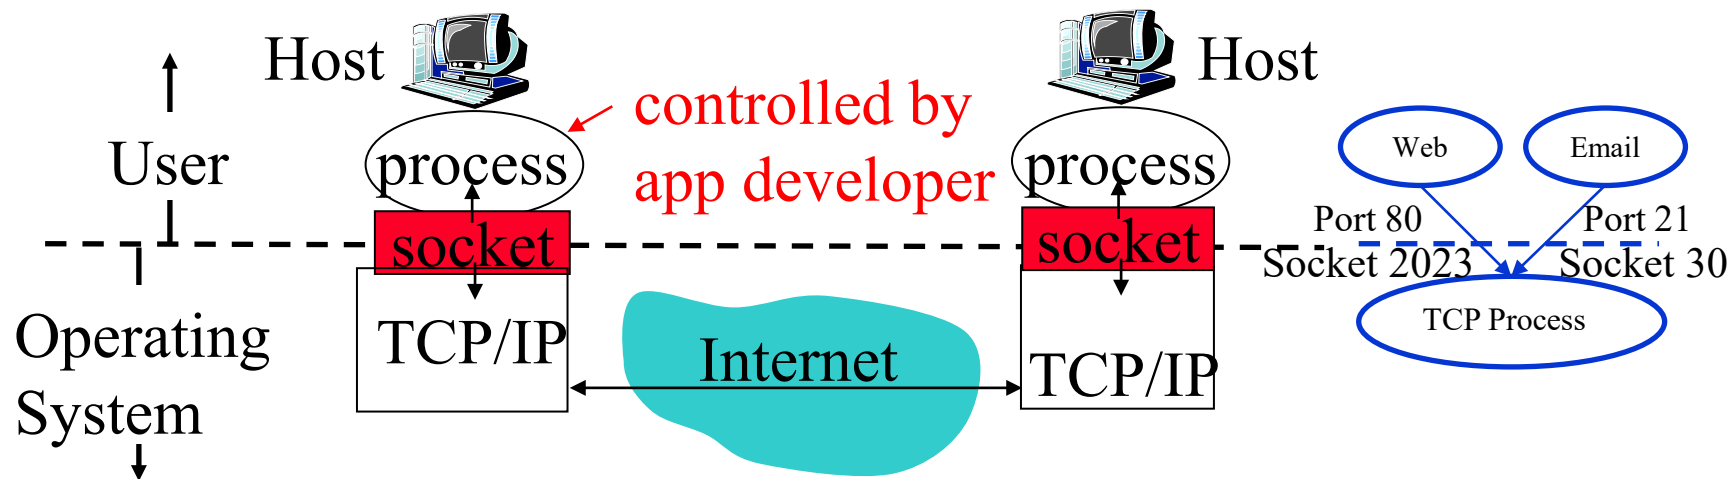
A new picture has been added.

- ❑ Is the socket an SDU or PDU?
- The socket is the UNIX name for API*
- ❑ How can we know which port of a server is open?

There are utilities/tools.

Process Communications

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



Student Questions

- ❑ Can we expose multiple applications to the same port? *No.*
- ❑ Are programming languages responsible for giving developers the socket interface to make web requests? If so, how do languages support different OSs?

Unix provides a Socket interface.


- ❑ Will the port open for TCP/IP protocols or applications or both? *Both. Any OS service.*

- ❑ I saw port 22 is always being used. Could you explain why port 22 is usually preferred?

Most apps used port numbers standardized by the Internet Engineering Task Force (IETF).

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



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.

- ❑ Are all the ports identical, other than the number?

No, each port provides a different service.

- ❑ You mentioned that we could change the default port for HTTP from 80 to whatever we want. Is that also true for other applications? For example, can we change the port for FTP from 21 to something else?

Yes. The clients must be told to connect to the new port numbers.


- ❑ How does a name turn into an IP address?

Names are easier to remember and static.

Numbers are dynamic.

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



Student Questions

- ❑ Who picked these commonly used port numbers? (80 for HTTP, 443 for HTTPS) *IETF*
- ❑ What are the differences between a port and a socket?

The socket is the Unix name for a port.

- ❑ What is the difference between using the default port vs. any other port numbers? Does the user have to specify the port at any time?
- ❑ Is there a predetermined number of ports on a computer? Does that number change between computers?

1 to $2^{16}-1$

- ❑ Are all the ports exist virtually, or do they physically exist on chips or other hardware?

They are just numbers.

- ❑ What is the Physical Address (MAC) in our IP information? Is that identical to any computer in the world?

Discussed in Chapter 6.


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

- ❑ The data are sent to the computer's port by socket, then sent to an external TCP/IP port?

Each port has a socket assigned to it.

- ❑ For HTTPS, isn't the standard port 443, and you can connect to that port directly instead of connecting to port 80 and then upgrading the security?

Sure. Both are OK.

- ❑ How can you connect to multiple websites on different tabs all through port 80? *Like party-line group calls.*

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.

- What makes TCP reliable for data transfer?

Discussed in Chapter 3

- Are transports a single layer up from the physical media used for communication?

Transports are layer 4

Physical media is layer 1.

- How exactly do transports work?

Discussed in Chapter 3

- Does connection-oriented mean the destination for the packets is the same, but the path for each packet to get to the destination might be different?

Connection => Setup before talking

- Do calling applications that use the internet (like Discord) use TCP or UDP?

UDP

- Is the TCP communicating by sending PDUs?

Yes.

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

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, the semantics of fields.

Student Questions

- ❑ Is the control plane mentioned here the same as an SDN control plane?
Yes. But the Control plane existed much before SDN.
-



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

- ❑ When building a web page, Apache can be downloaded to a device in order to serve the web page. Does this download establish a connection?

Apache is an open-source web server software. Firefox is an example of an open-source web client.

HTTP

- ❑ Uses TCP
- ❑ **Stateless**: The server does not remember the previous history
- ❑ **Non-Persistent**: Open a new TCP connection, get one object, and close
- ❑ **Persistent**: Open one TCP connection, get all objects, and close
The server leaves the connection open after sending an object and closes on the 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. The 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.

- ❑ Does this mean that HTTP can be switched between Persistent and Non-Persistent?

Yes. Originally it was non-persistent. Now it is mostly persistent.

HTTP

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

Student Questions


- ❑ In general, what happens during the process of the server forgetting its previous history (becoming stateless)?

It will not remember your name, address, or credit card numbers.

- ❑ Is non-persistent safer than persistent?
Security is independent of persistence.

Sample Web Page

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



Raj Jain

Student Questions

- How much HTML experience do we need for this course?

Not much.

- Can one HTTP session have multiple open TCP connections?

They generally do.

- Why does persistent HTTP need to close on timeout? Is it due to issues with too many people trying to connect to the server?

No, the client may have walked out. We need the resources that were committed for the client.

- What kinds of HTTP connections are non-persistent, and what kinds are persistent?

Very few elements on a page

⇒ Nonpersistent is OK

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 ever include www.cse.wustl.edu (using this slide as an example)?

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

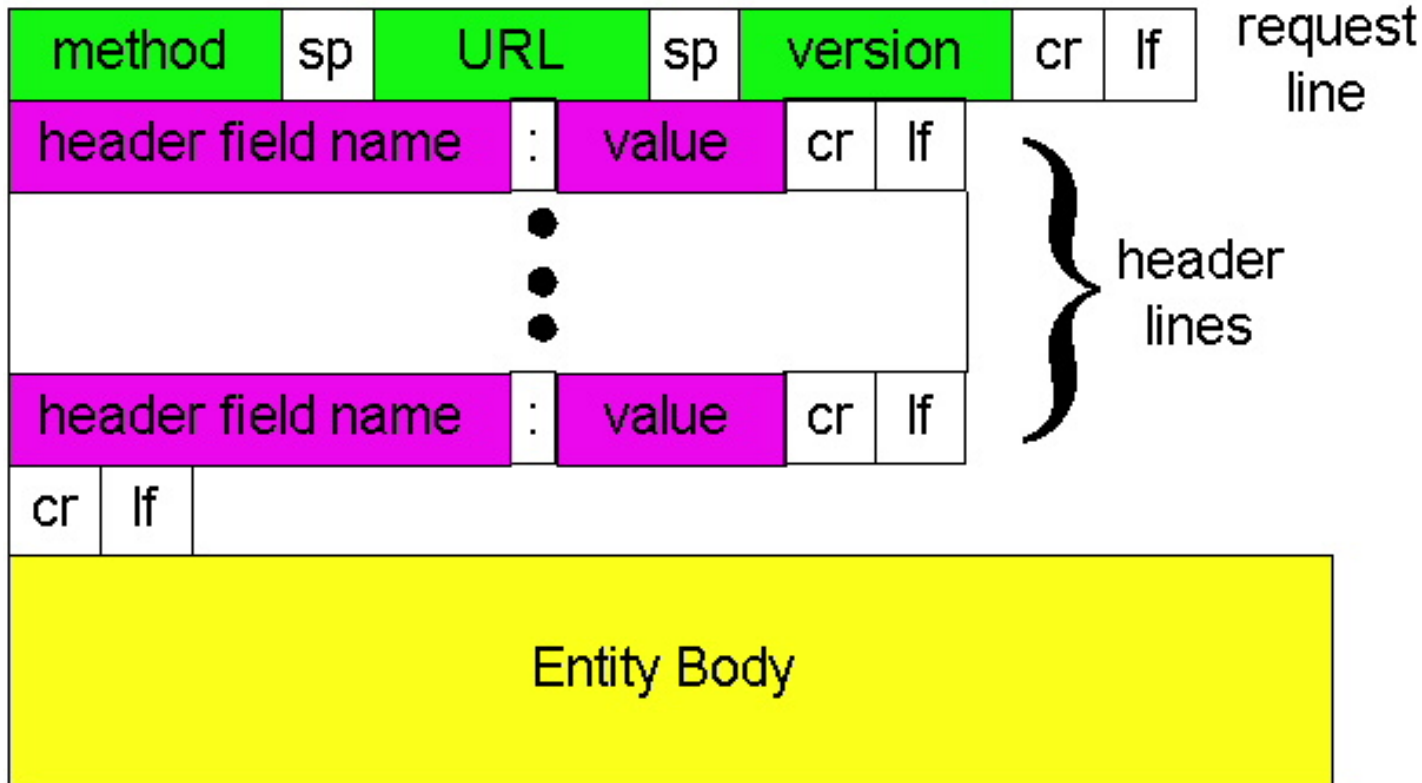
URL=Uniform Resource Locator.

Here directory is the URL/location on the host.

- ❑ It seems like there are many versions of HTTP, like 0.9, 1.0, 1.1, and 2.0. What would happen if I used another version than 1.1? Can my browser still interpret this response? What stops people from using the latest version besides habits?

Servers and Browsers try to be up-to-date and backward compatible.

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.

- Is there a default if a version is not provided?
No. If a version-specific feature is used, it is not clear what the server/browser will do.

- Why does the HTTP request message format require "cr", "lf", or "sp" to separate the information?

CR=Carriage Return

LF=Line Feed

SP=Space bar

WEB was designed by non-CS (physics) people. They wanted it as close to English as possible. It was one of the first ASCII protocols. Now there are many. They are easier to debug.

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

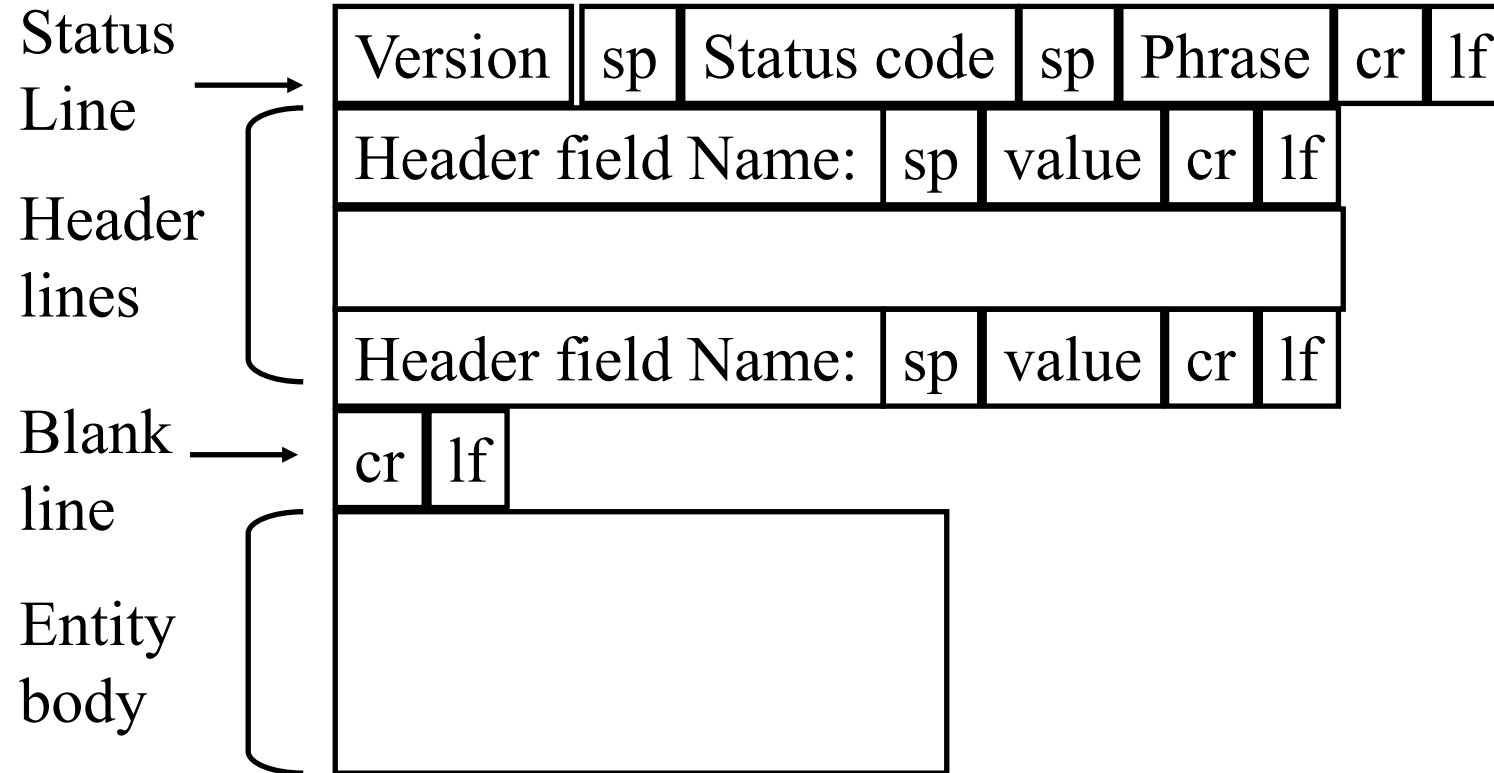
- What's the difference between a host and a server? Does the host decide what type of server it will use?

Host = Client or server

- What is 200 OK and 301 Moved Permanently?

Error messages have two parts: The first part can be read by computers, 2nd the part by humans.

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.

- What does "Accept-Ranges: bytes" mean?
Unit of length

- How do images/sound work?
They are binary files or streams.

- In the first line(telnet www1.cse.wustl.edu 80), can we access a different port by physically entering, let's say, 123 instead of 80? If that's not possible, how can we access a different port on someone else's machine, and how do ports ensure security?

Initially, telnet access was allowed on all ports. Now it is blocked for security reasons.

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

What is Apache?

Apache is an open-source web server.

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 misspelled domain names, e.g., citibnk.com.
⇒ *Credit Alert Systems*

- ❑ Can someone modify or add cookies manually to access other users' data?
You can modify cookies. But they are generally encrypted for security. If you modify any encrypted information, the receiver will know that the message has been modified and will reject it.

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

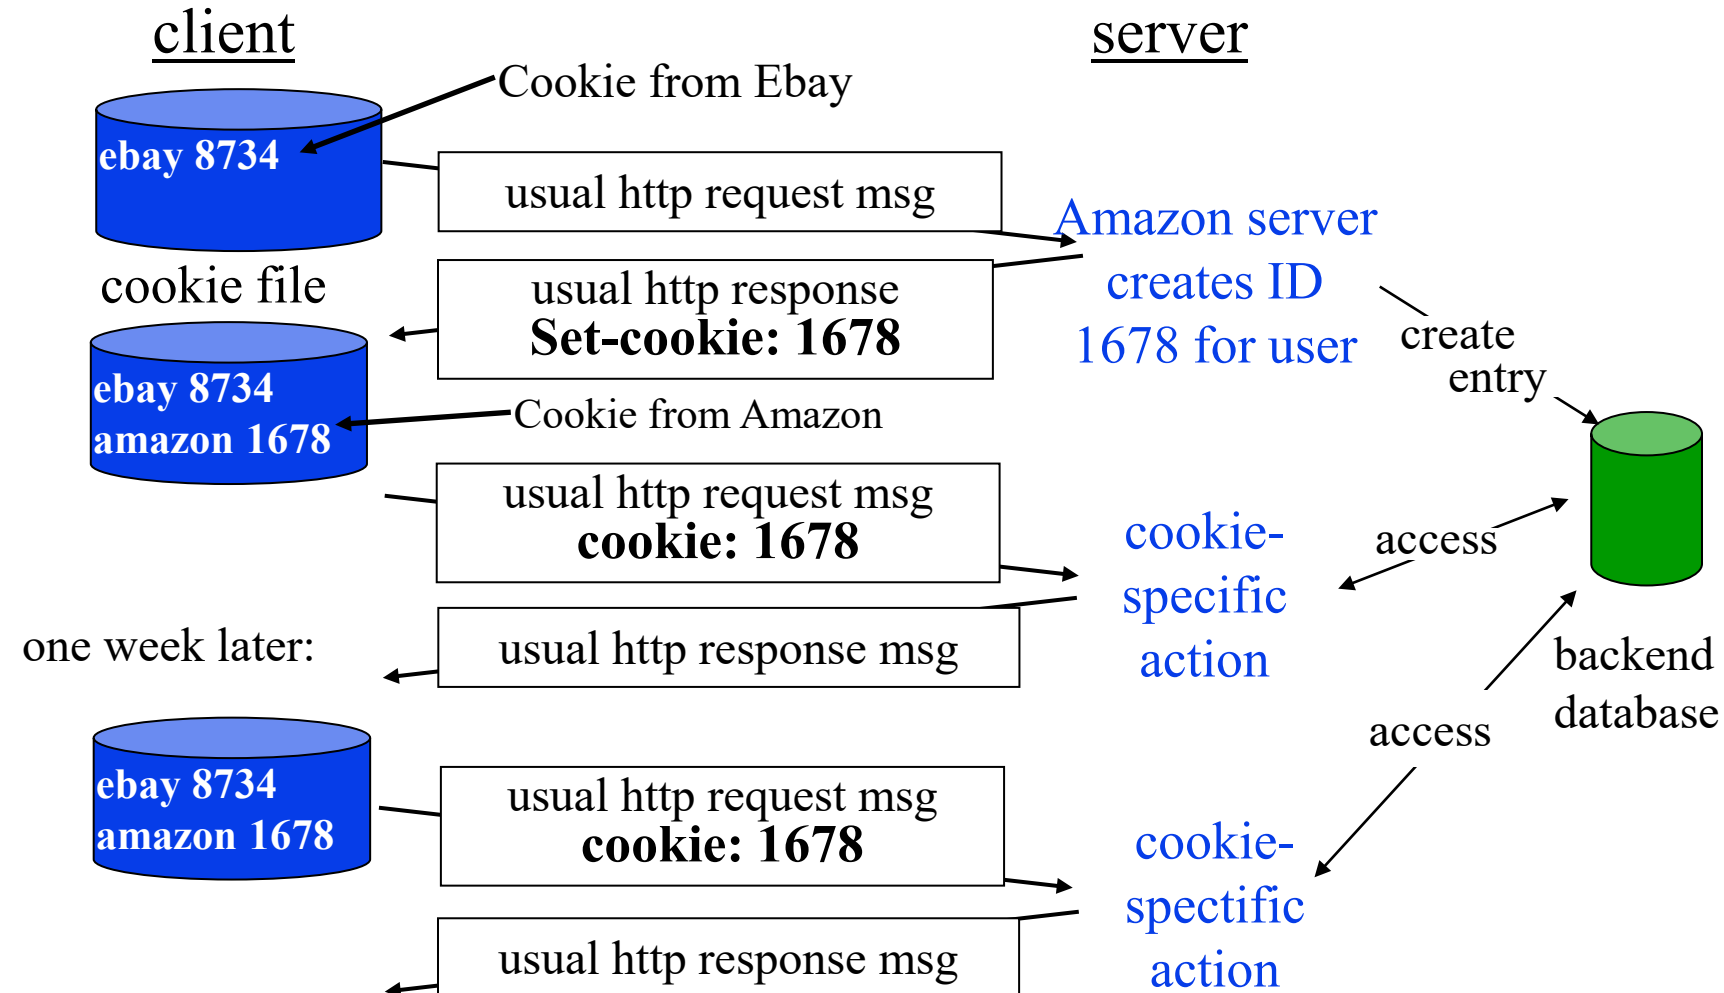
How does an ad appear on a server when I searched for the thing on another server previously?

They all share the info like credit reports

Cookies

- Allow servers to remember previous information

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



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 the browser option to not store them. But then you may have to do two-factor authentication on some servers.

- Are there web servers that require a username and password that do not use cookies? If so, how do those servers "remember" user information?

They start without prior information. That is when they will ask you to verify using 2-factor authentication.

- What do you think about Google's decision to move away from cookies? *They use cookies.*

- Where do clients store cookies? *Wherever browsers store bookmarks*


- Does each website own a cookie, or does some specific content on the website own a cookie?

Each server issues a cookie. A page may have info from many servers, e.g., an Ad server.

Cookies

- Allow servers to remember previous information

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




Student Questions

- Is sensitive information (passwords, etc.) ever stored in a cookie? Is it possible for a malicious user to create a "fake" cookie to impersonate another user?

Passwords are never transmitted on the wire. There are methods that allow password authentication without transmitting the password.

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 web server. <http://raj.jain.com> is a valid URL for a web server.

- ❑ So if we use a proxy server, the data saved by the proxy server has already been got by us before? *It could be for someone else that needed the same info.*
- ❑ Is the proxy server the local cache of the origin server? *Yes*
- ❑ Is being fast a requirement for a proxy server, or is it just desirable? If a proxy server is slow, would it not become a bottleneck?

Anything on the path between the client and server could become a bottleneck if slow.


Conditional GET



Student Questions

- ❑ Can a conditional get be used to launch a DoS attack? You mentioned that multiple get requests could be sent, and only a few might return (or return with the intended response).

Clients can retransmit the request.



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 a **state**.
5. Proxy servers improve performance by **caching** frequently used pages
6. **Conditional gets** allow 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

- Is the end of HTML content indicated by End-of-File or CR/LF

Two CR/LFs.

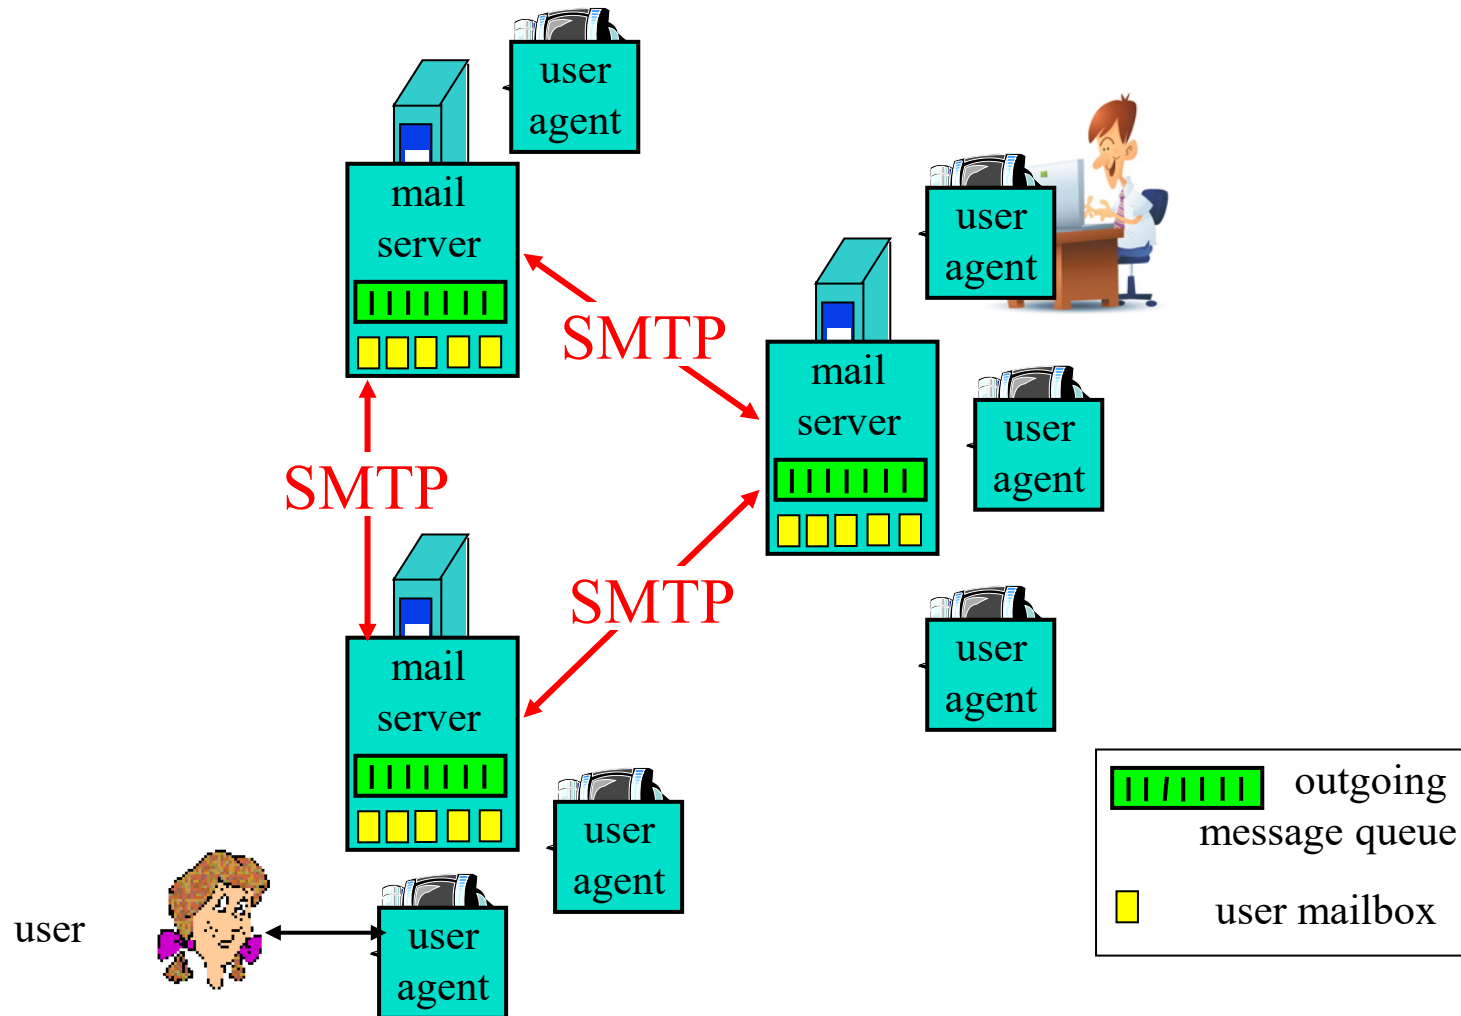
Lab 2A: Domains

[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 the wustl.edu domain (<http://www.networksolutions.com/whois/index.jsp>)
5. Find the name server of the 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.*
- What protocols are used by social media to send messages, such as Facebook, Whatsapp, etc.? *Each app uses many protocols, including some proprietary protocols.*
- Does SMTP also require the destination's IP address to deliver emails?

SMTP requires the destination mail server's IP address. Mail applications find that out from the destination user's email address and DNS.

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 example 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 the attached images also go through SMTP protocol or require some other protocol?

Bit strings can also be mapped as ASCII.

- Why is telnet considered insecure? Why has it not been updated and made more secure?

Telnet is no longer used. Similar and better functionality is offered by newer protocols.

- Using telnet, in the MAIL FROM section, can I type in any random email address? How do we verify the sender since we are not logging in?

Email senders are not verified or authenticated unless you use secure email.

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

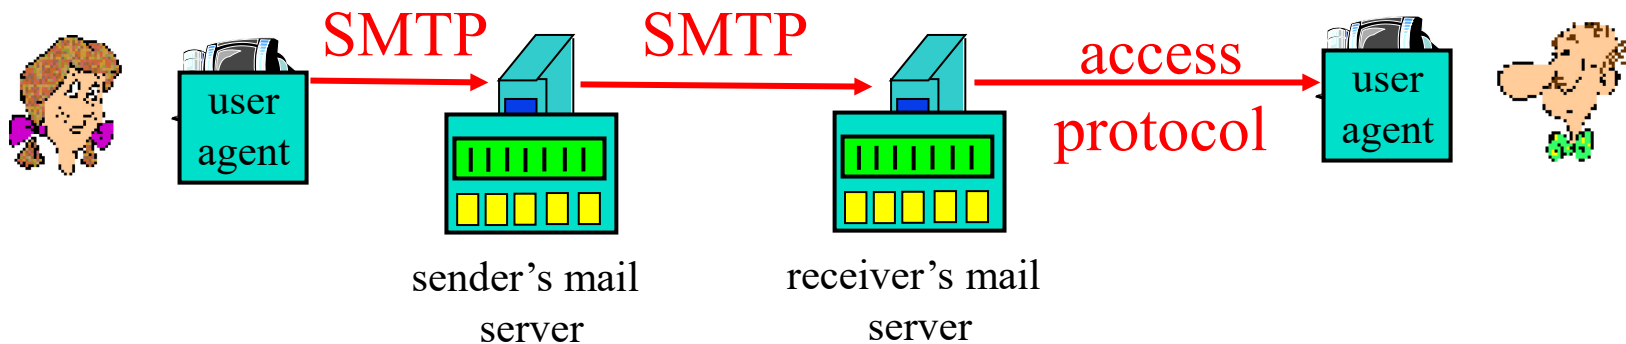
- The table indicates that HTTP only allows one object per response. It was also suggested previously that one website is composed of several objects: sound, image, video, text, etc. Then does it mean that one website might require several HTTP responses to completely transfer the data objects to the client?

Persistent connections return multiple responses. One object per response.

- Can we use HTTP to replace SMTP?
Yes. Now a days, HTTP is used to transport any protocol. Discussed in advanced networking courses.
-

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

- ❑ Is SMTP now obsolete?
No. I use it every day.
- ❑ Is HTTP also a mailing protocol, or does it work alongside/on top of a mailing protocol to deliver emails?
HTTP is used underneath the email protocol.

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

- Is there any reason that there is only one protocol for sending emails but multiple for receiving?

SMTP has been updated many times as the needs change. Receiving required new protocols to satisfy different people's needs.

- So between mail servers, they use SMTP, and between a user and mail server, POP is used?

POP is for receiving by the user.

- In the textbook, it says a user can send an email by using HTTP and also retrieve an email using HTTP, but that only happens between a user agent and the mail server. Why can it not happen between two mail servers?

SMTP is much faster and more efficient. HTTP is easier for the user, but it requires a lot more bits.

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

- For POP3 protocol, does the server always need to send back "OK" to clients' requests, or it's optional?
For servers, it is required. If the client doesn't wait for it, some messages may be lost.
- Since I had already deleted the first message, the second message became the current first message. Why would the client still ask for "retr 2"?

The serial numbers are not changed in the middle of a connection.

- Does POP have any previous versions? If it has any predecessor, what is the main drawback?
POP3 is 3rd version of POP.
- For a server, how to know the address of the server serves the user that the outgoing email is sent to? Is this using the feature of aliasing of DNS protocol? If so, it seems the DNS server needs to maintain a highly verbose list of aliases because there exist huge numbers of email accounts.

Yes. DNS can resolve all addresses in the world. See Slide 2-42

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

- Is the user urgent for the client in this condition?


In what condition?

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

- ❑ When a message is deleted on the phone, can we still access the message on the laptop?
You can program each receiver to either delete the email or keep it even after reading it. So if you program your phone to not delete it from the server, it will be there.
- ❑ Does more advanced protocol mean lower efficiency? *No.*
- ❑ Can you go over IMAP again? *Sure.*



Mail: Summary

1. SMTP is the protocol for **sending** 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, 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 a new e-mail. Suppose that in the five-minute interval, no new messages have been sent to you. Provide a transcript of this second POP session for both options a and b above.

Student Questions

- For problem c, do we need to include the authorization phase? Or just the transaction phase? *Just the transaction phase.*
-



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

- How does DNS work when the service is distributed across multiple servers? How does it know to resolve to the closest?

See Slide 2-42

- How does DNS translate a hostname to IP?

See Slide 2-42

- To clarify, can you send an email from a non-existing mail server (i.e., president@whitehouse.com)?

Yes. But nowadays, the servers authenticate the user and verify the domain name.

- Some mail servers still don't care what from address is used. If I claim to have sent mail from an address, would the actual owner of that address be able to find out (without the recipient telling them)?

Yes, only the name is faked. The IP address is still your IP address.



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

- Is it risky to change DNS? *Yes.*
-

DNS

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

- ❑ Do DNS servers violate your privacy by knowing every site you wish to visit?

Yes.

- ❑ Could you please explain more about the idea of a distributed database?

See Slide 2-41

- ❑ We can get an IP address by DNS. But how can we get the port number? Should I use the default port number if I want to talk with a server? *Yes. Unless the serving company told you to use another port #.*

Are the host aliasing and load distribution mean that the hostname and IP address are not one-to-one mappings?

Yes. Google.com is thousands of servers.

Are DNS numbers fixed?

What is a DNS number?

Why do websites have multiple IP addresses?

For load distribution

DNS

- ❑ Domain Name Service
- ❑ DNS servers translate a hostname to IP address
E.g., www.wustl.edu ⇒ 128.252.87.149
- ❑ A 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.
 - ❑ Do DNS servers violate your privacy by knowing every site you wish to visit? *Yes.*
 - ❑ Could you please explain more about the idea of a distributed database? *See Slide 2-41*
 - ❑ We can get an IP address by DNS. But how can we get the port number? Should I use the default port number if I want to talk with a server? *Yes. Unless the serving company told you to use another port #.*
-

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 a 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 the transport layer.

- What part is the www in www.facebook or wustl in wustl.instructure.com? Is it another child authoritative server below facebook or an 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

- A data center has a lot of hosts. Do these hosts share the same IP?

Not necessarily. Address allocation is discussed in Chapter 4.

- Where can I find my public IP?
Whatismyip.com

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

Do all hosts under a Wi-Fi network share the same public IP?

Not necessarily. Address allocation is discussed in Chapter 4.

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 the 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.
- ❑ What organization manages the root DNS server?
Internet Assigned Numbers Authority (IANA)
- ❑ So if we want to give a name like 'seas.', should we have a DNS server?
Seas.wustl.edu would resolve all names ending in seas.wustl.edu

DNS Hierarchy




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

Student Questions

- ❑ Will the 'poly.edu' DNS server still can have its children's DNS server?
Yes. There are no restrictions on the levels of hierarchy.

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

- ❑ What's the typical cache policy of DNS servers?
Each entry has a lifetime.
- ❑ Can you explain the example of the picture again? What exactly is being requested, and why does UMass provide an answer to WUSTL?

A server in WUSTL needs the address of `gaia.cs.umass.edu`.

- ❑ How does the computer choose between sending recursive queries or iterative queries?

Lower-level users just want the answer. Higher-level DNS servers may use the hint.

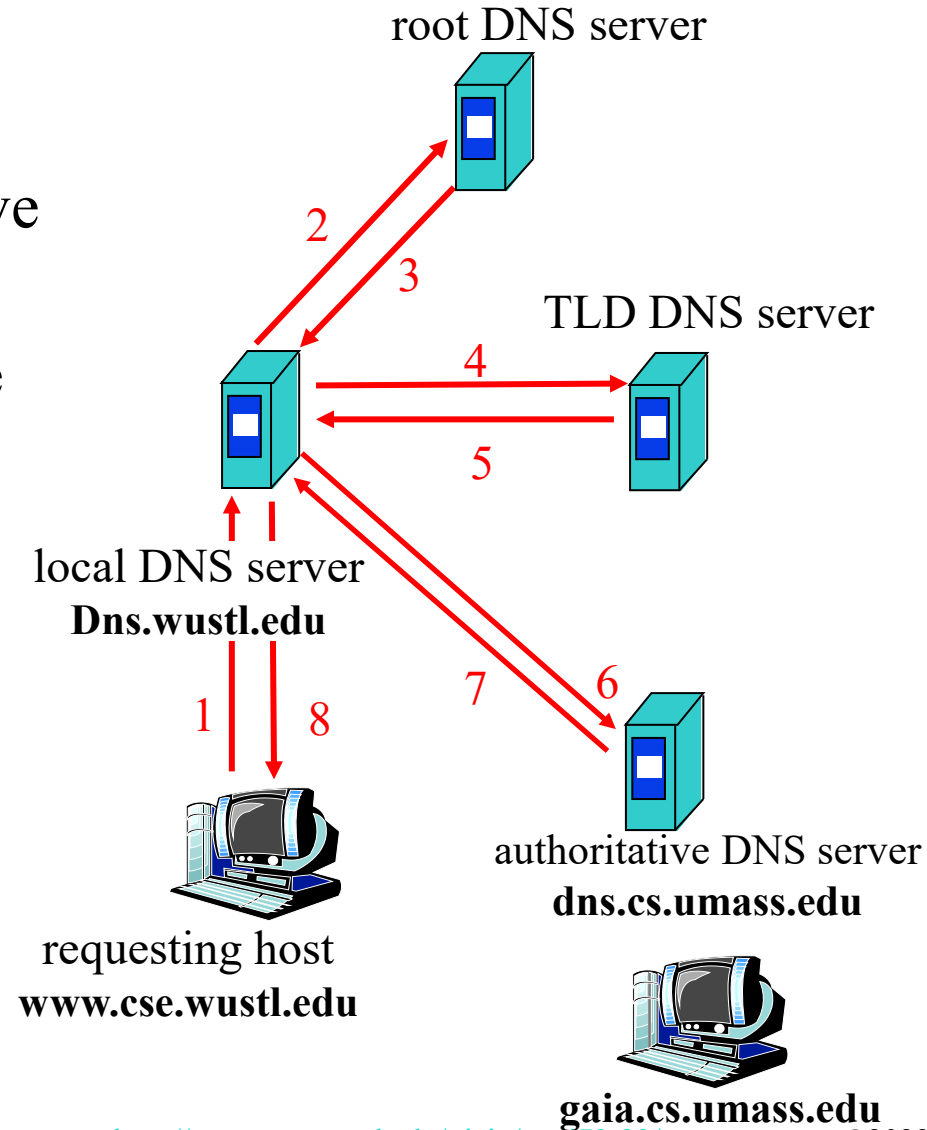
- ❑ Why would you want a hint and not an answer? Can you explain this part again?

You may want to resolve many addresses in that domain.

- ❑ What may cause iterative queries?
You may want to resolve many addresses in that domain.

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

- ❑ What's the difference between a local DNS server and an authoritative DNS server?

Local is near you. Authoritative is near the name you are trying to resolve.

- ❑ What's the cache policy of the DNS server?

Controlled by the network manager. Refresh after a while or refresh only if unreachable.

- ❑ Which of the requests on the picture are recursive, and which are iterative?

Anyone can be iterative or recursive. Most clients will make recursive requests. Most servers will make an iterative request.

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

- ❑ Say if a mail system and web page both use DNS. Then how does DNS differentiate between these two requests?

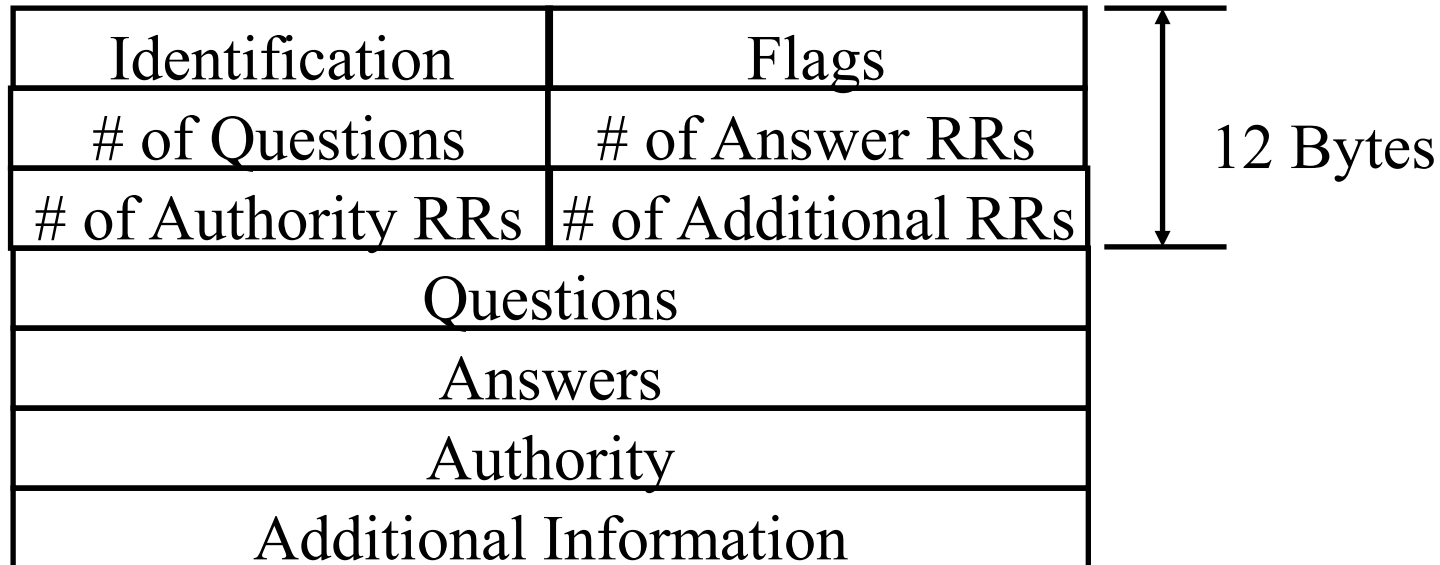
The answer is returned to the sending port.

- ❑ Is there a specific reason mail servers use type MX? If it is a canonical name, why not just use type CNAME?

The mail server for cse.wustl.edu may be different from the Web server. Each computer has a canonical name.

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

- ❑ What is the authority section used for in most cases? What is an example?
Seas.wustl.edu is the authority for all seas.wustl.edu names. But WUSTL.edu may also serve as an authority.
- ❑ What is the relationship between DNS Records and the DNS Message Format?
Records are kept on the disk. Messages are sent over the network. Messages need additional headers.

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 the 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 the wrong answer

Student Questions

- ❑ Could you explain how DNS cache poisoning works again?

You get a wrong answer from someone. Then you keep sending your packets to the wrong place.



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

- Does DNS consist of a single database, or could there be multiple databases?
Multiple databases. However, there is only one authoritative database.
- What would happen if I changed the DNS of my PC? *You will be sending your queries to this different server.*
- As far as I understand, every OS has a DNS cache. Would disabling it also disable access to the world wide web (meaning the browsers or something else strongly depends on the cache), or would it just make accessing it a lot harder?

You can resolve names without caching them just by asking someone else.

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

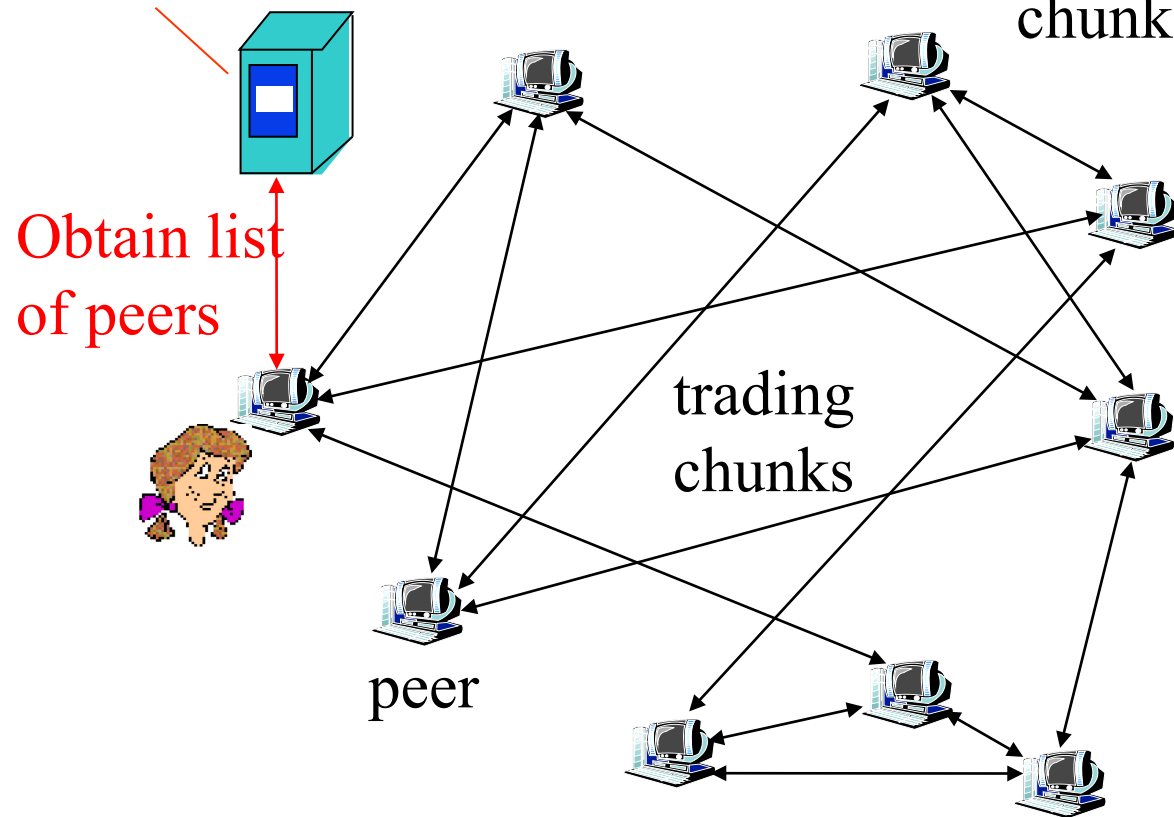
1. Client Server vs. P2P Scalability
2. P2P File Distribution (BitTorrent)

Student Questions

P2P File Distribution (BitTorrent)

Tracker: tracks peers participating in torrent


Torrent: group of peers exchanging chunks of a file



Student Questions

- Can the communication in p2p be blocked or monitored by other clients?
Yes. They can become members of the torrent but throw away your requests.
- Is P2P vulnerable to attack because of its decentralized nature? *Security is independent of the P2P vs. Client-Server.*
- For peer-to-peer, if you have the file, can anyone get the file from you automatically, or do you have to take some action to make the file available to others? *You have to put it in the directory where your torrent client can access it. All files in the directory where downloaded files are kept are automatically served.*

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 the 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?
- If Peer to peer networks are able to scale much better than the client-server model, Why do we still use the client-server model?


For companies, it is easier to manage central servers. Charging and billing are also easier with C/S than with P2P.

- P2P has a high scalability benefit over client-server. Is P2P only less popular because it has been used for nefarious purposes? *See above*
- What exactly is meant by "Distribution time" on the y-axis of the graph?

Time to finish the distribution of files to all clients.

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

- ❑ Like this figure on the left, do Leechers have different roles in different positions?
Leechers are also servers.
- ❑ For a particular movie file, is the list of all peers changing once there are new seeds or leechers? If so, will the tracker change? And if the tracker changes, will we have different torrent files for this particular movie?

The torrent file is simply the "ID" of the file and a list of seeds + leechers. ID is simply a hash of the content.

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers.
The tracker may provide a random subset (say, 50) of 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))

Student Questions

- ❑ Can you explain unchoking/choking?

Choking = Not sending the packets

- ❑ Is Alice sending the same file as what she is receiving?

It may have other files that other users want.

- ❑ Can you elaborate on what "Unchokes"
Unchoke=Starts sending.

- ❑ Does "top 4 senders" refer to the top 4 senders to that particular peer or the top 4 senders in the whole torrent?

To that peer:

- ❑ Is Alice finding high-speed neighbors for herself or the tracker?

For herself:

- ❑ Can we go over Slide 2-54 again, specifically regarding choking/unchoking and how this relates to file distribution?

Sure.

- ❑ How does the tracker keep track of all the available peers?

Whoever asks for information gets into the database.

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers.
The tracker may provide a random subset (say, 50) of 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

- ❑ According to the slides, Alice continuously sends data to the top 4 senders. However, in a scenario where we only have 4 peers, and one peer is not sending any data to the other peers. Would it still continuously receive data from Alice? Previously in the hw, we saw that if a peer (Bob) does not upload any data to any other peers, he was still able to get the file due to optimistic unchoking. However, it could take him a long time to receive the file because the process of receiving it was random. However, since now we only have 4 peers in total, would Bob receive data at the same rate as the other peers who are uploading data?

In this case, random selection will result in the 4th user being selected every time, and yes, it will get the file.

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers.
The tracker may provide a random subset (say, 50) of 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))


Student Questions

- ❑ I don't understand what Alice is "sending" or to whom. It wouldn't make sense for her to be sending chunks of the same file back to senders.

She may have other files that those senders may want.

- ❑ What happens if everyone is missing one of the chunks? Would the chunk be sent from the tracker?

Trackers do not keep any files.



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: Pixels 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 the motion. So the choice is either to keep decreasing the quality (CBR) or to increase the bit rate (VBR) during motion.

Dynamic Adaptive Streaming over HTTP (DASH)

- ❑ DASH provides an efficient method for video streaming
- ❑ Standard Web Servers: No changes are 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
- ❑ The client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.



Student Questions

- ❑ For the DASH protocol, I'm confused about what benefits 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 into multiple segments. Each segment needs a URL.

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

Dynamic Adaptive Streaming over HTTP (DASH)

- ❑ DASH provides an efficient method for video streaming
- ❑ Standard Web Servers: No changes are 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
- ❑ The client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.




Student Questions

- ❑ If it happens that a high-quality video will drop to a low-quality video. Would video drop from 720p to 360p. Are we intentionally dropping some packets of the video at each frame or simply loading frames with lower quality?

Both.

Content Distribution Networks (CDN)

- ❑ To reduce latency to worldwide users, the data is replicated at many sites
- ❑ Users are directed to the 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 5 different qualities (rates) and 10-minute segments. How many URLs will be required for a 3-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 allow servers to maintain states
4. **Proxy** servers improve performance by caching
5. BitTorrent is a **P2P** file distribution protocol and uses trackers to keep a list of peers.
6. **DASH** allows clients to request different video segments as needed
7. **CDN** directs users to nearby copies via DNS

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

Student Questions

- Why email and file transfers are using different protocols? Why not treat email also as a file?

Files are usually “pulled,” while emails are usually “pushed.”

- Since proxy servers cache from the original server, does this mean it's a secure way to hide actions on the Internet?

Nothing is hidden. The fact that the proxy server took it from the original server is visible to the world.

- Can you go over what cache is?

Computer Science 101.

- Are all five homework questions and two labs due on Monday?

Yes.

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 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

- ❑ Will the exam involve the use of the Socket APIs?

The exam does not require any coding. So no APIs. But it may ask about standard port numbers.

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 that 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 to 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 the 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>

http://www.cse.wustl.edu/~jain/cse473-23/i_2app.htm

Student Questions

- When are Lab 1 and Lab 2 due? Where do we submit them?

Monday, February 7th. Submit on Canvas.

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/cse473-23/>

©2023 Raj Jain

Student Questions

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



Raj Jain

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
Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozv/wAARCACWAGYDICA
AhEBAxEB/8QAGwAAAQUBAQAAAAAAAAAAAAAAAAABgACAwwQFAQf/xAA3EAABAwIFAQYCCQQA
AAAAAAAAAABAAIDBBEFBhIhMUETIIFhcYGRsRQjMjNCUqHB0RVi4fAHNFP/xAAZAQADAQEBA
AAAAAAAAAAAAAAAAAAQMEAgX/xAAiEQACAgICAglDAAAAAAAAAAAAAAAAAQIRAyEEMRJBeyIUMmH/2gAMA5QAAhEDEQA/AK11

.....
0zhMzuPdcSDp6rQgZdhdYEHghOMYvpIuPMJUkbaeU6S7S7h19h7KRQY8Fs19PeHCuUcvZVUM4P3b
D3Wk/uLg+6Kc5SXqaMX5Dv2QkKB+IYfiFG5319O/6TE//wBBax+Q91r1lccToMHncbyPhJd6iwJ+
IKb/AF2cV9qRXgLo3gg7cLQhDnEaXb+XRUnAM72nU09FagmDC022U0M26bFZ6ezJ7ys8fxD+VrMn
jnj7SJ4cCh0EEBzRcLjJ5KWXESL8joV0AQE7pKrT1bapmoCzhy3wSTEBjagmudDb7LNTT5qYu2A
aLW3v6pJKZQjcdZueCo2m5HTdJJJaJK6ok/psVGx7mtlmOq3UBpNvksmW0H1cY07bkJLjJ2jdxk
vBv+mtgL5mNebmQGMn+1wt/HwUD6V+GYjFhxk1tpqcAEeJJcfmkkqQf0MvISWTRba8OdpI2sugN
hfYgljr2APCSSEQZfw6QvYWngHb0KkmAvb4JJLoQ6lkdDd7TvwUkk10I/9k=

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

Student Questions

- What does the first line of the MIME message signify? (containing NextPart + some numbers)

Not really sure. The numerical part could be a unique serial number in hex. Such numbers are generally obtained by hashing.