



*Department of
Information Sciences and Telecommunications*



Computer Networks – TELCOM 2310

Lecture 3

Network Applications
Design Issues and Protocols

Prof. Taieb Znati

Lecture Outline



- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

Network Applications Terminology and Building Blocks



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

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

Applications and application-layer protocols




Application: communicating, distributed processes

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

Application-layer protocols

- ✦ one “component” of an application
- ✦ define messages exchanged by apps and actions taken
- ✦ use communication services provided by lower layer protocols (TCP, UDP)



Application-layer protocol characteristics



✘ It defines:

- ✘ Types of messages exchanged, eg, request & response messages
- ✘ **Syntax** of message types: what fields in messages & how fields are delineated
- ✘ **Semantics** of the fields, ie, meaning of information in fields

- ✘ Rules for when and how processes send & respond to messages

Public-domain protocols:

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

Proprietary protocols:

- ✘ eg, KaZaA

Client-server paradigm




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

Client:

- ✘ initiates contact with server (“speaks first”)
- ✘ typically requests service from server,
- ✘ Web: client implemented in browser; e-mail: in mail reader


Server:

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



Processes communicating across network

- ✘ process sends/receives messages to/from its socket
- ✘ socket analogous to door
 - ✘ sending process inserts message through the door
 - ✘ sending process assumes transport infrastructure on other side of door which brings message to socket at receiving process



- ✘ API: (1) choice of transport protocol; (2) ability to fix a few parameters

Fall '06-02

TELCOM 2310

7

Addressing processes:

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

Fall '06-02

TELCOM 2310

8

What transport service does an application need?



Data loss

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

Timing

- ✘ some apps (e.g., Internet telephony, interactive games) require low delay to be “effective”

Bandwidth

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

Transport service requirements of common apps



Application	Data loss	Bandwidth	Time Sensitive
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video: 10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

Internet transport protocols services



TCP service:

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

UDP service:

- ✦ unreliable data transfer between sending and receiving process
- ✦ does not provide: connection setup, reliability, flow control, congestion control, timing, or bandwidth guarantee

Internet apps: application, transport protocols



Application	Application layer protocol	Underlying transport protocol
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Dialpad)	typically UDP

Lecture Outline



- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

Web and HTTP



Terminology

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


`www.someschool.edu/someDept/pic.gif`

host name path name

HTTP overview

HTTP: hypertext transfer protocol

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



HTTP overview (continued)

Uses TCP:

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

HTTP is "stateless"

- ✘ server maintains no information about past client requests

Protocols that maintain "state" are complex!

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

HTTP connections



Nonpersistent HTTP

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

Persistent HTTP

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

Nonpersistent HTTP



Suppose user enters URL

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

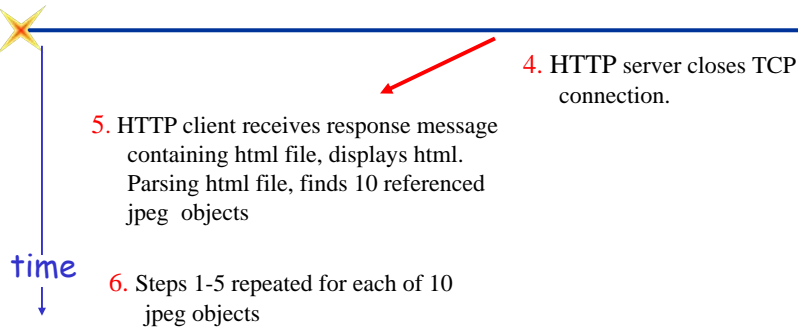
(contains text, references to 10 jpeg images)

- 1a. HTTP client initiates TCP connection to HTTP server (process) at `www.someSchool.edu` on port 80
- 1b. HTTP server at host `www.someSchool.edu` waiting for TCP connection at port 80. "accepts" connection, notifying client
2. HTTP client sends HTTP *request message* (containing URL) into TCP connection socket. Message indicates that client wants object `someDepartment/home.index`
3. HTTP server receives request message, forms *response message* containing requested object, and sends message into its socket

time



Nonpersistent HTTP (cont.)




Response time modeling

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

Response time:

- ✦ one RTT to initiate TCP connection
- ✦ one RTT for HTTP request and first few bytes of HTTP response to return
- ✦ file transmission time

total = 2RTT + transmit time



Persistent HTTP



Nonpersistent HTTP issues:

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

Persistent HTTP

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

Persistent without pipelining:

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

Persistent with pipelining:

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

Fall '06-02

TELCOM 2310

21

HTTP request message



- ✘ two types of HTTP messages: *request, response*
- ✘ **HTTP request message:**
 - ✘ ASCII (human-readable format)

request line
(GET, POST,
HEAD commands)

header
lines

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

Carriage return
line feed
indicates end
of message


(extra carriage return, line feed)

Fall '06-02

TELCOM 2310

22

HTTP request message: general format



Fall '06-02

TELCOM 2310

23

Method types



HTTP/1.0

- ✘ GET
- ✘ POST
- ✘ HEAD
 - ✘ asks server to leave requested object out of response

HTTP/1.1

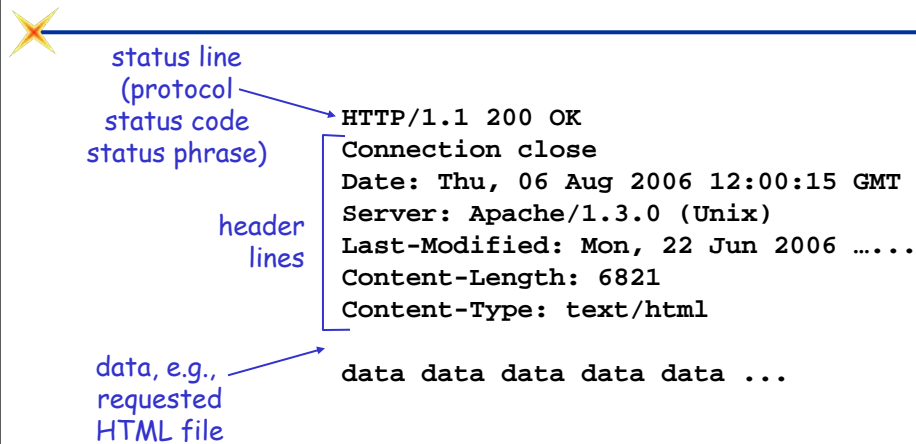
- ✘ GET, POST, HEAD
- ✘ PUT
 - ✘ uploads file in entity body to path specified in URL field
- ✘ DELETE
 - ✘ deletes file specified in the URL field

Fall '06-02

TELCOM 2310

24

HTTP response message



Fall '06-02

TELCOM 2310

25

HTTP response status codes

In first line in server->client response message.

A few sample codes:

200 OK

- * request succeeded, requested object later in this message

301 Moved Permanently

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

400 Bad Request

- * request message not understood by server

404 Not Found

- * requested document not found on this server

505 HTTP Version Not Supported

Fall '06-02

TELCOM 2310

26

Trying out HTTP (client side) for yourself



1. Telnet to your favorite Web server:

```
telnet www.site.fr 80
```

Opens TCP connection to port 80 (default HTTP server port) at www.site.fr. Anything typed in sent to port 80 at www.site.fr

2. Type in a GET HTTP request:

```
GET /~user/index.html HTTP/1.0
```

By typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. Look at response message sent by HTTP server!

Fall '06-02

TELCOM 2310

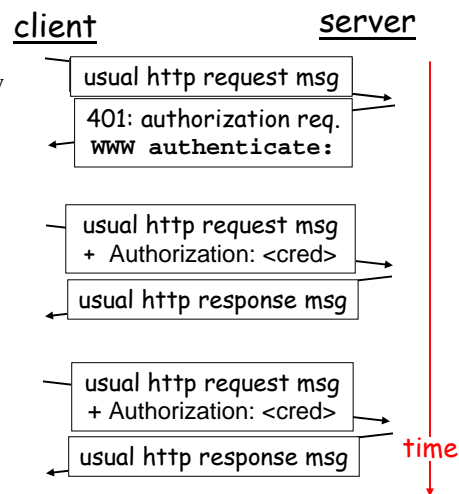
27

User-server interaction: authorization



Authorization : control access to server content

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



Fall '06-02

TELCOM 2310

28

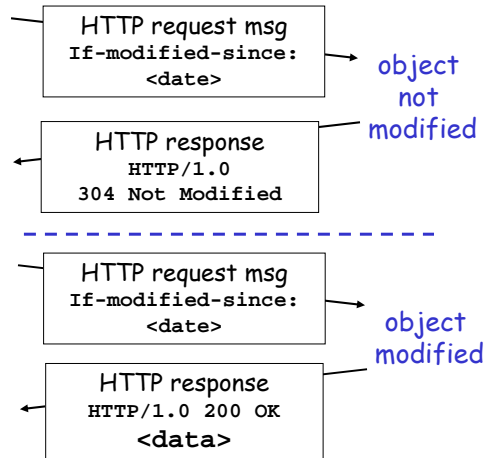
Conditional GET: client-side caching



- ✘ **Goal:** don't send object if client has up-to-date cached version
- ✘ **client:** specify date of cached copy in HTTP request
`If-modified-since: <date>`
- ✘ **server:** response contains no object if cached copy is up-to-date:
`HTTP/1.0 304 Not Modified`

client

server



Fall '06-02

TELCOM 2310

29

Lecture Outline




- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

Fall '06-02

TELCOM 2310

30

FTP: the file transfer protocol



- ✘ transfer file to/from remote host
- ✘ client/server model
 - ✘ *client*: side that initiates transfer (either to/from remote)
 - ✘ *server*: remote host
- ✘ ftp: RFC 959
- ✘ ftp server: port 21


Fall '06-02

TELCOM 2310

31

FTP: separate control, data connections

- ✘ FTP client contacts FTP server at port 21, specifying TCP as transport protocol
- ✘ Client obtains authorization over control connection
- ✘ Client browses remote directory by sending commands over control connection.
- ✘ When server receives a command for a file transfer, the server opens a TCP data connection to client
- ✘ After transferring one file, server closes connection.



- ✘ Server opens a second TCP data connection to transfer another file.
- ✘ Control connection: "out of band"
- ✘ FTP server maintains "state": current directory, earlier authentication

Fall '06-02

TELCOM 2310

32

FTP commands, responses



Sample commands:

- ✘ sent as ASCII text over control channel
- ✘ **USER *username***
- ✘ **PASS *password***
- ✘ **LIST** return list of file in current directory
- ✘ **RETR *filename*** retrieves (gets) file
- ✘ **STOR *filename*** stores (puts) file onto remote host

Sample return codes

- ✘ status code and phrase (as in HTTP)
- ✘ **331 Username OK, password required**
- ✘ **125 data connection already open; transfer starting**
- ✘ **425 Can't open data connection**
- ✘ **452 Error writing file**

Lecture Outline



- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

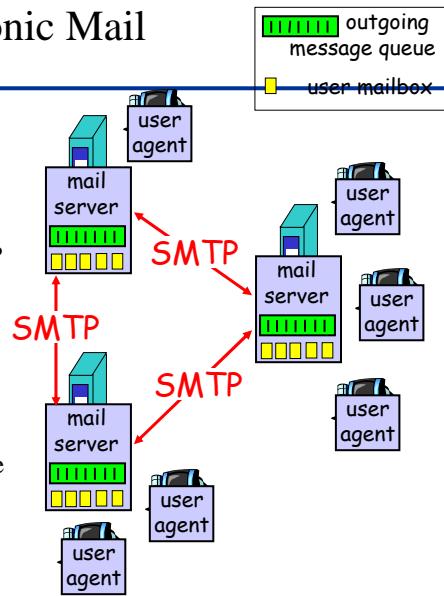
Electronic Mail

Three major components:

- ✘ user agents
- ✘ mail servers
- ✘ simple mail transfer protocol: SMTP

User Agent

- ✘ a.k.a. "mail reader"
- ✘ composing, editing, reading mail messages
- ✘ e.g., Eudora, Outlook, elm, Netscape Messenger
- ✘ outgoing, incoming messages stored on server



Fall '06-02


TELCOM 2310

35

Electronic Mail: mail servers

Mail Servers

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



Fall '06-02

TELCOM 2310

36

Electronic Mail: SMTP [RFC 2821]

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


Fall '06-02

TELCOM 2310

37

Scenario: Alice sends message to Bob

- 1) Alice uses UA to compose message and "to" bob@someschool.edu
- 2) Alice's UA sends message to her mail server; message placed in message queue
- 3) Client side of SMTP opens TCP connection with Bob's mail server
- 4) SMTP client sends Alice's message over the TCP connection
- 5) Bob's mail server places the message in Bob's mailbox
- 6) Bob invokes his user agent to read message



Fall '06-02

TELCOM 2310

38

Sample SMTP interaction



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

Fall '06-02

TELCOM 2310

39

Try SMTP interaction for yourself:



- ✘ **telnet servername 25**
 - ✘ see 220 reply from server
 - ✘ enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands
- above lets you send email without using email client (reader)

Fall '06-02

TELCOM 2310

40

SMTP vs HTTP



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

Comparison with HTTP:

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

Fall '06-02

TELCOM 2310

41


Mail message format



SMTP: protocol for exchanging email msgs

RFC 822: standard for text message format:

- ✘ header lines, e.g.,
 - ✘ To:
 - ✘ From:
 - ✘ Subject:*different from SMTP commands!*
- ✘ body
 - ✘ the “message”, ASCII characters only



Fall '06-02

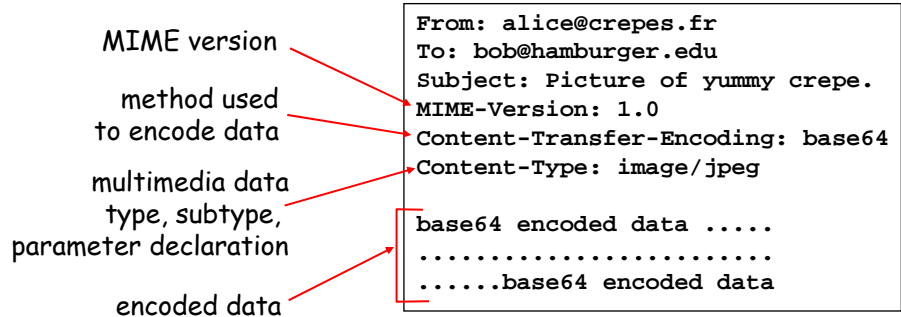
TELCOM 2310

42

Message format: multimedia extensions



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



MIME types

Content-Type: type/subtype; parameters



Text

- ✘ example subtypes: **plain**, **html**

Video

- ✘ example subtypes: **mpeg**, **quicktime**

Image

- ✘ example subtypes: **jpeg**, **gif**

Application

- ✘ other data that must be processed by reader before “viewable”
- ✘ example subtypes: **msword**, **octet-stream**

Audio

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

Multipart Type




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

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



Mail access protocols



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

POP3 protocol



authorization phase

- ✘ client commands:
 - ✘ **user:** declare username
 - ✘ **pass:** password

- ✘ server responses
 - ✘ **+OK**
 - ✘ **-ERR**

transaction phase, client:

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

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

C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 1 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off
```

Fall '06-02

TELCOM 2310

47

POP3 (more) and IMAP



More about POP3

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

IMAP

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

Fall '06-02

TELCOM 2310

48

Lecture Outline



- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

DNS: Domain Name System



People: many identifiers:

- ✘ SSN, name, passport #

Internet hosts, routers:

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

Q: map between IP addresses and name ?

Domain Name System:

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

DNS name servers



Why not centralize DNS?

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

doesn't *scale!*

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

local name servers:

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

authoritative name server:

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

Fall '06-02

TELCOM 2310

51

DNS: Root name servers



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



13 root name servers worldwide

Fall '06-02


TELCOM 2310

52

Simple DNS example

host **surf.eurecom.fr**
wants IP address of
omni.cs.pitt.edu

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



Fall '06-02


TELCOM 2310

53

DNS example

Root name server:

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



Fall '06-02

TELCOM 2310

54

DNS: iterated queries




recursive query:

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

iterated query:

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



Fall '06-02

TELCOM 2310

55

DNS: caching and updating records



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

Fall '06-02

TELCOM 2310

56

DNS records



DNS: distributed db storing resource records (**RR**)

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

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

Fall '06-02

TELCOM 2310

57

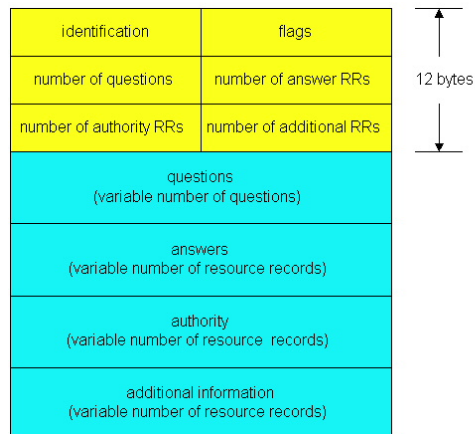
DNS protocol, messages



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

msg header

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



Fall '06-02

TELCOM 2310

58

DNS protocol, messages

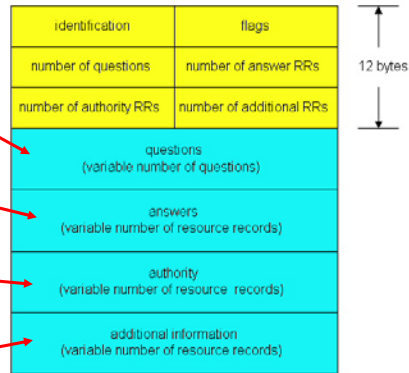


Name, type fields
for a query

RRs in reponse
to query

records for
authoritative servers

additional "helpful"
info that may be used



Lecture Outline



- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

Socket programming



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

Socket API

- ✘ introduced in BSD4.1 UNIX, 1981
- ✘ explicitly created, used, released by apps
- ✘ client/server paradigm
- ✘ two types of transport service via socket API:
 - ✘ unreliable datagram
 - ✘ reliable, byte stream-oriented

socket


a *host-local, application-created, OS-controlled* interface (a "door") into which application process can **both send and receive** messages to/from another application process

Socket-programming using TCP



Socket: a door between application process and end-end-transport protocol (UCP or TCP)

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



Socket programming *with TCP*



Client must contact server

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

Client contacts server by:

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

- ✘ When contacted by client, **server TCP creates new socket** for server process to communicate with client

- ✘ allows server to talk with multiple clients
- ✘ source port numbers used to distinguish clients ([more in Chap 3](#))

application viewpoint

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

Stream terminology




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

Socket programming with TCP

Example client-server app:

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



Fall '06-02

TELCOM 2310

65

Client/server socket interaction: TCP

Server (running on `hostid`)

Client

```
create socket,
port=x, for
incoming request:
welcomeSocket =
ServerSocket()
```

```
wait for incoming
connection request
connectionSocket =
welcomeSocket.accept()
```

```
read request from
connectionSocket
```

```
write reply to
connectionSocket
```

```
close
connectionSocket
```

TCP
connection setup

```
create socket,
connect to hostid, port=x
clientSocket =
Socket()
```

```
send request using
clientSocket
```

```
read reply from
clientSocket
```

```
close
clientSocket
```

Fall '06-02

TELCOM 2310

66

Socket programming *with UDP*

UDP: no “connection” between client and server

- ✘ no handshaking
- ✘ sender explicitly attaches IP address and port of destination to each packet
- ✘ server must extract IP address, port of sender from received packet

application viewpoint

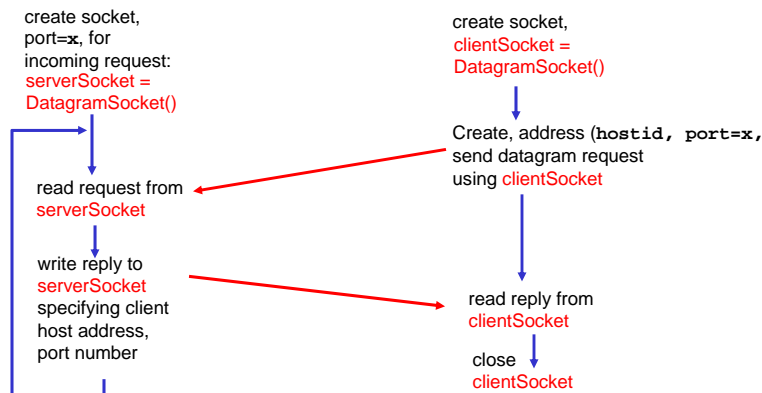
UDP provides *unreliable* transfer of groups of bytes (“datagrams”) between client and server

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

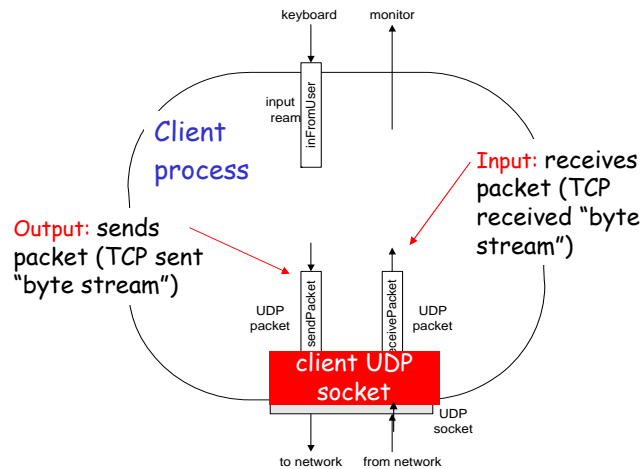
Client/server socket interaction: UDP

Server (running on `hostid`)

Client



UDP Sockets



Fall '06-02

TELCOM 2310

69

Lecture Outline

- ✘ Application Layer Building Blocks and Terminology
 - ✘ Client Server Model
- ✘ Application Layer Examples
 - ✘ HTTP, FTP, SMTP
- ✘ Domain Name System
- ✘ Socket Programming
- ✘ Peer-to-Peer, Overlay Networks
- ✘ Conclusion

Fall '06-02

TELCOM 2310

70

P2P file sharing

Example

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

All peers are servers = highly scalable!

Fall '06-02

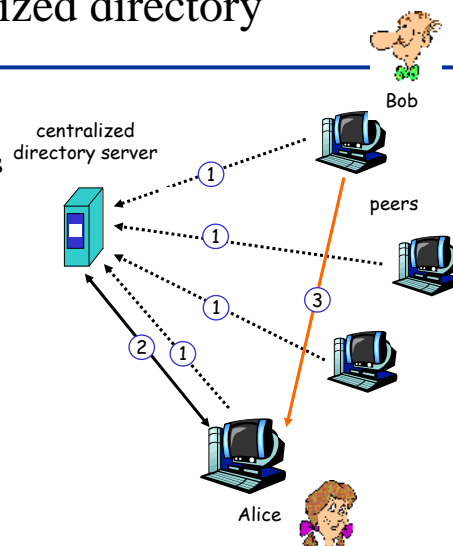
TELCOM 2310

71

P2P: centralized directory

Original “Napster” design

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



Fall '06-02

TELCOM 2310

72

P2P: problems with centralized directory




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

file transfer is decentralized, but locating content is highly decentralized

P2P: decentralized directory



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



● ordinary peer
● group-leader peer
— neighboring relationships in overlay network

More about decentralized directory



overlay network

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

bootstrap node

- ✘ connecting peer is either assigned to a group leader or designated as leader

advantages of approach

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

disadvantages of approach

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

Fall '06-02



TELCOM 2310

75

P2P: Query flooding



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



Fall '06-02

TELCOM 2310

76

P2P: more on query flooding



Pros

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

Cons

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

Summary I



Study of network applications

- ✘ application service requirements:
 - ✘ reliability, bandwidth, delay
- ✘ client-server paradigm
- ✘ Internet transport service model
 - ✘ connection-oriented, reliable: TCP
 - ✘ unreliable, datagrams: UDP
- ✘ specific protocols:
 - ✘ HTTP
 - ✘ FTP
 - ✘ SMTP, POP, IMAP
 - ✘ DNS
- ✘ socket programming
- ✘ content distribution
 - ✘ caches, CDNs
 - ✘ P2P

Summary II



Application protocols and communication paradigms

- ✘ typical request/reply message exchange:
 - ✘ client requests info or service
 - ✘ server responds with data, status code
- ✘ message formats:
 - ✘ headers: fields giving info about data
 - ✘ data: info being communicated
- ✘ control vs. data msgs
 - ✘ in-band, out-of-band
- ✘ centralized vs. decentralized
- ✘ stateless vs. stateful
- ✘ reliable vs. unreliable msg transfer
- ✘ “complexity at network edge”
- ✘ security: authentication