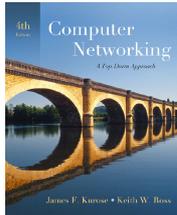


## Chapter 2 Application Layer



*Computer Networking:  
A Top Down Approach,  
4th edition.*  
Jim Kurose, Keith Ross  
Addison-Wesley, July  
2007.

### A note on the use of these ppt slides:

We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a lot of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) in substantially unaltered form, that you mention their source (after all, we'd like people to use our book!)
- If you post any slides in substantially unaltered form on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

All material copyright 1996-2007  
J.F. Kurose and K.W. Ross, All Rights Reserved

## Chapter 2: Application Layer

### Our goals:

- conceptual, implementation aspects of network application protocols
  - ❖ transport-layer service models
  - ❖ client-server paradigm
  - ❖ peer-to-peer paradigm
- learn about protocols by examining popular application-level protocols
  - ❖ HTTP
  - ❖ FTP
  - ❖ SMTP / POP3 / IMAP
  - ❖ DNS
- programming network applications
  - ❖ socket API

## Position of application layer

Gives services to the users

- Client-server paradigm
- Addressing
- Different services

Receives services from the transport layer  
Reliability Delay Throughput



## Some network apps

- e-mail
- web
- instant messaging
- remote login
- P2P file sharing
- multi-user network games
- streaming stored video clips
- voice over IP
- real-time video conferencing
- grid computing
- 
- 
- 

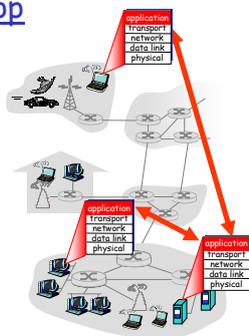
## Creating a network app

### write programs that

- ❖ run on (different) end systems
- ❖ communicate over network
- ❖ e.g., web server software communicates with browser software

### No need to write software for network-core devices

- ❖ Network-core devices do not run user applications
- ❖ applications on end systems allows for rapid app development, propagation



## Network applications: some jargon

**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

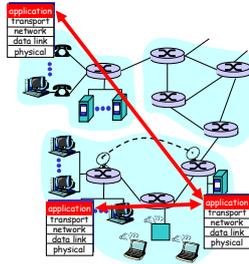
## Applications and application-layer protocols

### 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 "piece" of an app
- ❖ define messages exchanged by apps and actions taken
- ❖ use communication services provided by lower layer protocols (TCP, UDP)



## App-layer protocol defines

- ❑ Types of messages exchanged,
  - ❖ e.g., request, response
- ❑ Message syntax:
  - ❖ what fields in messages & how fields are delineated
- ❑ Message semantics
  - ❖ meaning of information in fields
- ❑ Rules for when and how processes send & respond to messages

### Public-domain protocols:

- ❑ defined in RFCs
- ❑ allows for interoperability

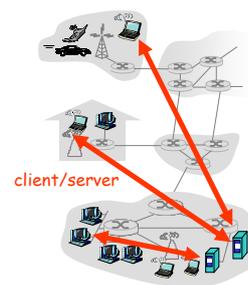
### Proprietary protocols:

- ❑ e.g., HTTP, SMTP
- ❑ e.g., Skype

## Application architectures

- ❑ Client-server
- ❑ Peer-to-peer (P2P)
- ❑ Hybrid of client-server and P2P

## Client-server architecture



### server:

- ❖ always-on host
- ❖ permanent IP address
- ❖ server farms for scaling

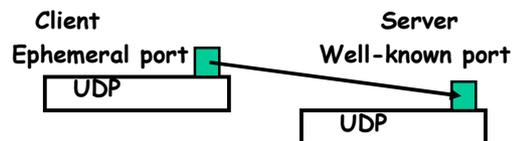
### clients:

- ❖ communicate with server
- ❖ may be intermittently connected
- ❖ may have dynamic IP addresses
- ❖ do not communicate directly with each other

## Concurrency in servers

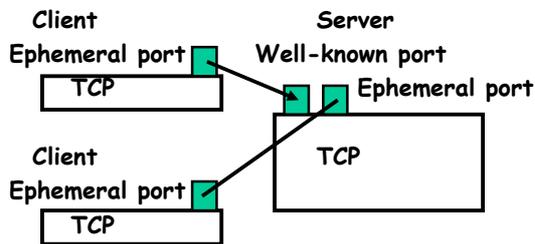
- ❑ Iterative server
- ❑ Concurrent server

## Connectionless iterative server



One client at a time

## Connection-oriented concurrent server



## Processes

Program is code - defines all the variables and actions to be performed on those variables

A process is an instance of a program - when the operating system executes a program, an instance of the program, a process, is created

## Inter-process Communication (IPC)

- Mechanism for processes to communicate and to synchronize their actions.
- Message system - processes communicate with each other without resorting to shared variables.
- IPC facility provides two operations:
  - ❖ `send(message)` - message size fixed or variable
  - ❖ `receive(message)`
- If  $P$  and  $Q$  wish to communicate, they need to:
  - ❖ establish a *communication link* between them
  - ❖ exchange messages via `send/receive`
- Implementation of communication link

## Implementation Questions

- How are links established?
- Can a link be associated with more than two processes?
- How many links can there be between every pair of communicating processes?
- What is the capacity of a link?
- Is the size of a message that the link can accommodate fixed or variable?
- Is a link unidirectional or bi-directional?

## Direct Communication

- Processes must name each other explicitly:
  - ❖ `send(P, message)` - send a message to process  $P$
  - ❖ `receive(Q, message)` - receive a message from process  $Q$
- Properties of communication link
  - ❖ Links are established automatically.
  - ❖ A link is associated with exactly one pair of communicating processes.
  - ❖ Between each pair there exists exactly one link.
  - ❖ The link may be unidirectional, but is usually bi-directional.

## Indirect Communication

- Messages are directed and received from mailboxes (also referred to as ports).
  - ❖ Each mailbox has a unique id.
  - ❖ Processes can communicate only if they share a mailbox.
- Properties of communication link
  - ❖ Link established only if processes share a common mailbox
  - ❖ A link may be associated with many processes.
  - ❖ Each pair of processes may share several communication links.
  - ❖ Link may be unidirectional or bi-directional.

## Indirect Communication

- Operations
  - ❖ create a new mailbox
  - ❖ send and receive messages through mailbox
  - ❖ destroy a mailbox
- Primitives are defined as:
  - send**(*A, message*) - send a message to mailbox *A*
  - receive**(*A, message*) - receive a message from mailbox *A*

2: Application Layer 19

## Synchronization

- Message passing may be either blocking or non-blocking.
- **Blocking** is considered **synchronous**
- **Non-blocking** is considered **asynchronous**
- **send** and **receive** primitives may be either blocking or non-blocking.

2: Application Layer 20

## Buffering

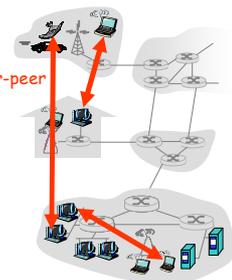
- Queue of messages attached to the link; implemented in one of three ways.
  1. Zero capacity - 0 messages  
Sender must wait for receiver (rendezvous).
  2. Bounded capacity - finite length of *n* messages  
Sender must wait if link full.
  3. Unbounded capacity - infinite length  
Sender never waits.

2: Application Layer 21

## Pure P2P architecture

- no always-on server
- arbitrary end systems directly communicate **peer-peer**
- peers are intermittently connected and change IP addresses

Highly scalable but difficult to manage



2: Application Layer 22

## Hybrid of client-server and P2P

### Skype

- ❖ voice-over-IP P2P application
- ❖ centralized server: finding address of remote party:
- ❖ client-client connection: direct (not through server)

### Instant messaging

- ❖ chatting between two users is P2P
- ❖ centralized service: client presence detection/location
  - user registers its IP address with central server when it comes online
  - user contacts central server to find IP addresses of buddies

2: Application Layer 23

## Processes communicating

- Process:** program running within a host.
- within same host, two processes communicate using **inter-process communication** (defined by OS).
- processes in different hosts communicate by exchanging **messages**

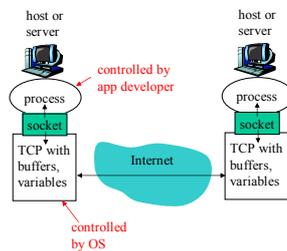
**Client process:** process that initiates communication  
**Server process:** process that waits to be contacted

- Note: applications with P2P architectures have client processes & server processes

2: Application Layer 24

## Sockets

- process sends/receives messages to/from its **socket**
- socket analogous to door
  - sending process shoves message out door
  - sending process relies on 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 (lots more on this later)



## Addressing processes

- to receive messages, process must have **identifier**
- host device has unique 32-bit IP address
- Q:** does IP address of host suffice for identifying the process?

## Addressing processes

- to receive messages, process must have **identifier**
- host device has unique 32-bit IP address
- Q:** does IP address of host on which process runs suffice for identifying the process?
  - A:** No, many processes can be running on same host
- identifier** includes both **IP address** and **port numbers** associated with process on host.
  - Example port numbers:
    - HTTP server: 80
    - Mail server: 25
  - to send HTTP message to gaia.cs.umass.edu web server:
    - IP address:** 128.119.245.12
    - Port number:** 80
- more shortly...

## What transport service does an app 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"

### Throughput

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

### Security

- Encryption, data integrity, ...

## Transport service requirements of common apps

Application	Data loss	Throughput	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 provide:** timing, minimum throughput guarantees, security

### UDP service:

- unreliable data transfer between sending and receiving process
- does not provide: connection setup, reliability, flow control, congestion control, timing, throughput guarantee, or security
- Q:** why bother? Why is there a UDP?

## 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	HTTP (eg Youtube), RTP [RFC 1889]	TCP or UDP
Internet telephony	SIP, RTP, proprietary (e.g., Skype)	typically UDP

## Web and HTTP

### First some jargon

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

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

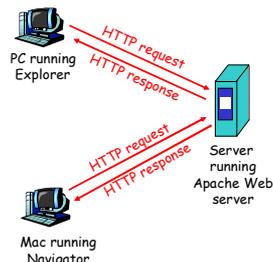
host name

path name

## HTTP overview

### HTTP: hypertext transfer protocol

- Web's application layer protocol
- client/server model
  - client**: browser that requests, receives, "displays" Web objects
  - server**: Web server sends objects in response to requests



## HTTP 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.

### Persistent HTTP

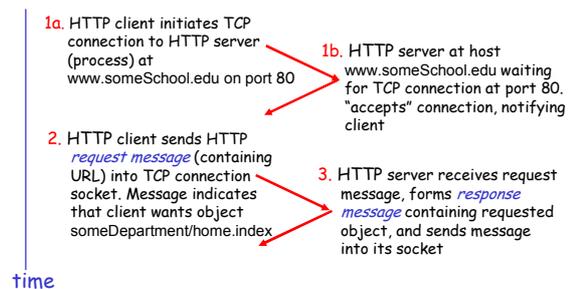
- Multiple objects can be sent over single TCP connection between client and server.

## Nonpersistent HTTP

Suppose user enters URL

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

(contains text, references to 10 jpeg images)



## Nonpersistent HTTP (cont.)

- time ↓
- HTTP client receives response message containing html file, displays html. Parsing html file, finds 10 referenced jpeg objects
  - Steps 1-5 repeated for each of 10 jpeg objects
  - HTTP server closes TCP connection.

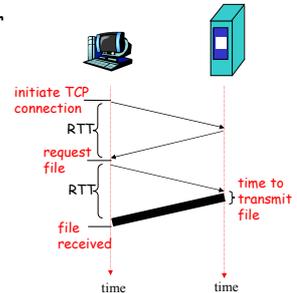
## Non-Persistent HTTP: Response time

**Definition of RTT:** time for 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 overhead for each TCP connection
- browsers often open parallel TCP connections to fetch referenced objects

### Persistent HTTP

- server leaves connection open after sending response
- subsequent HTTP messages between same client/server sent over open connection
- client sends requests as soon as it encounters a referenced object
- as little as one RTT for all the referenced objects

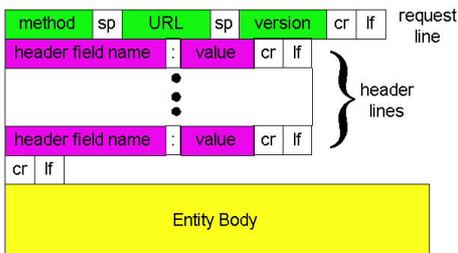
## HTTP request message

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

```

request line
(GET, POST, HEAD commands) GET /somedir/page.html HTTP/1.1
header lines
Host: www.someschool.edu
User-agent: Mozilla/4.0
Connection: close
Accept-language: fr
Carriage return line feed
(extra carriage return, line feed)
    
```

## HTTP request message: general format



## Uploading form input

### Post method:

- Web page often includes form input
- Input is uploaded to server in entity body

### URL method:

- Uses GET method
- Input is uploaded in URL field of request line:

[www.somesite.com/animalsearch?monkeys&banana](http://www.somesite.com/animalsearch?monkeys&banana)

## Method types

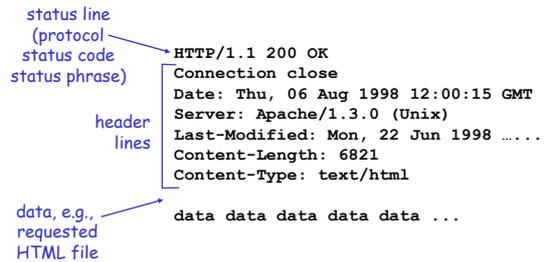
### HTTP/1.0

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

### HTTP/1.1

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

## HTTP response message



## HTTP response status codes

In first line in server→client response message.  
 A few sample codes:

- 200 OK**
  - ❖ request succeeded, requested object later in this message
- 301 Moved Permanently**
  - ❖ requested object moved, new location specified later in this message (Location:)
- 400 Bad Request**
  - ❖ request message not understood by server
- 404 Not Found**
  - ❖ requested document not found on this server
- 505 HTTP Version Not Supported**

## Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

```
telnet cis.poly.edu 80
```

Opens TCP connection to port 80 (default HTTP server port) at cis.poly.edu. Anything typed in sent to port 80 at cis.poly.edu

2. Type in a GET HTTP request:

```
GET /~ross/ HTTP/1.1
Host: cis.poly.edu
```

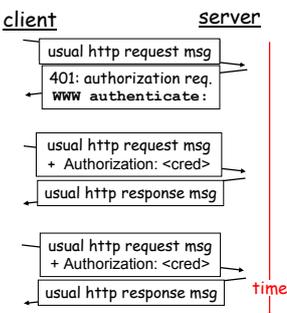
By typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. Look at response message sent by HTTP server!

## User-server 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



## User-server state: cookies

Many major Web sites use cookies

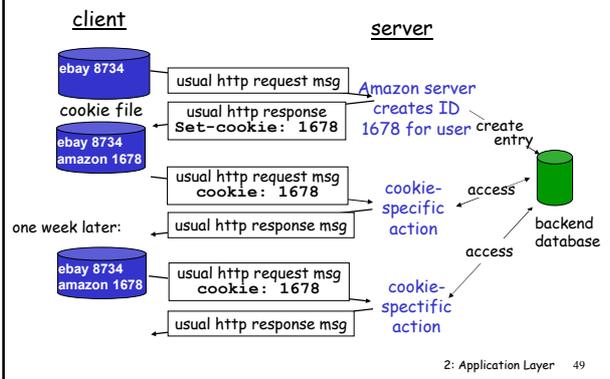
### Four components:

- 1) cookie header line of HTTP *response* message
- 2) cookie header line in HTTP *request* message
- 3) cookie file kept on user's host, managed by user's browser
- 4) back-end database at Web site

### Example:

- Susan always access Internet always from PC
- visits specific e-commerce site for first time
- when initial HTTP requests arrives at site, site creates:
  - ❖ unique ID
  - ❖ entry in backend database for ID

## Cookies: keeping "state" (cont.)



## Cookies (continued)

### What cookies can bring:

- authorization
- shopping carts
- recommendations
- user session state (Web e-mail)

### How to keep "state":

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

### aside

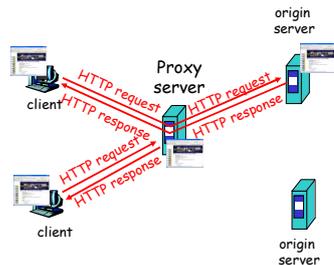
#### Cookies and privacy:

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

## Web caches (proxy server)

Goal: satisfy client request without involving origin server

- user sets browser: Web accesses via cache
- browser sends all HTTP requests to cache
  - ❖ object in cache: cache returns object
  - ❖ else cache requests object from origin server, then returns object to client



## More about Web caching

- cache acts as both client and server
- typically cache is installed by ISP (university, company, residential ISP)

### Why Web caching?

- reduce response time for client request
- reduce traffic on an institution's access link.
- Internet dense with caches: enables "poor" content providers to effectively deliver content (but so does P2P file sharing)

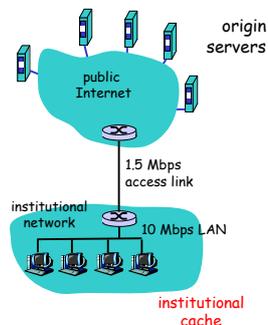
## Caching example

### Assumptions

- average object size = 100,000 bits
- avg. request rate from institution's browsers to origin servers = 15/sec
- delay from institutional router to any origin server and back to router = 2 sec

### Consequences

- utilization on LAN = 15%
- utilization on access link = 100%
- total delay = Internet delay + access delay + LAN delay = 2 sec + minutes + milliseconds



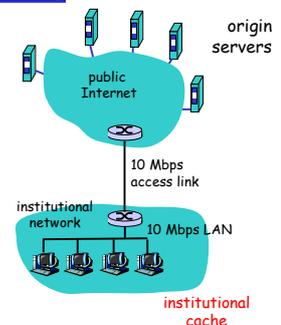
## Caching example (cont)

### possible solution

- increase bandwidth of access link to, say, 10 Mbps

### consequence

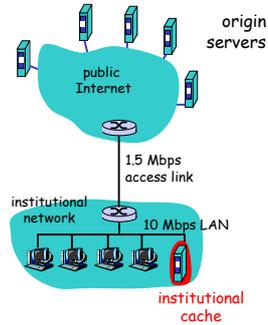
- utilization on LAN = 15%
- utilization on access link = 15%
- Total delay = Internet delay + access delay + LAN delay = 2 sec + msec + msec
- often a costly upgrade



## Caching example (cont)

### possible solution: install cache

- suppose hit rate is 0.4
- consequence**
- 40% requests will be satisfied almost immediately
- 60% requests satisfied by origin server
- utilization of access link reduced to 60%, resulting in negligible delays (say 10 msec)
- total avg delay = Internet delay + access delay + LAN delay =  $.6 * (2.01 \text{ secs} + .4 * \text{milliseconds}) < 1.4 \text{ secs}$



## Web Caching (The issues)

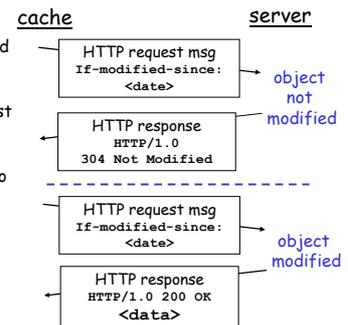
- Cache location, object placement and object replacement. These issues deal with how to locate/fetch cached objects, where to place a cached object in order to achieve optimal performance, what cache objects should be replaced, and how and when to replace them
- Cache organization, distribution, hierarchy and cooperation
- Page Pre-fetching - (network performance?)
- Real-time multimedia traffic, namely streaming media applications.

## Cache Replacement

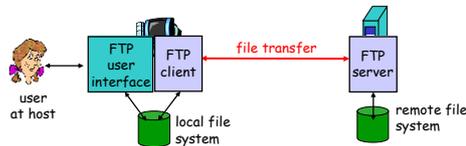
- Page expires
- Policy
  - ❖ Least Frequently Used
  - ❖ Least Recently Used
  - ❖ Page content (priority)

## Conditional GET

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

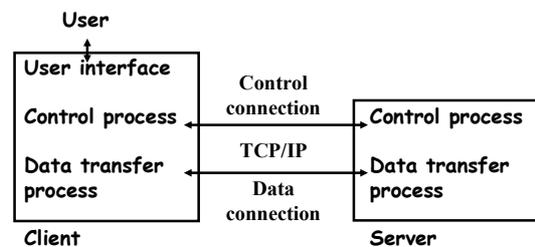


## FTP: the file transfer protocol



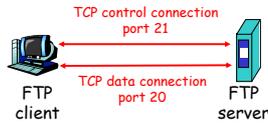
- transfer file to/from remote host
- client/server model
  - ❖ client: side that initiates transfer (either to/from remote)
  - ❖ server: remote host
- ftp: RFC 959
- ftp server: port 21

## FTP



## FTP: separate control, data connections

- FTP client contacts FTP server at port 21, TCP is transport protocol
- client authorized over control connection
- client browses remote directory by sending commands over control connection.
- when server receives file transfer command, server opens 2<sup>nd</sup> TCP connection (for file) to client
- after transferring one file, server closes data connection.



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

## 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

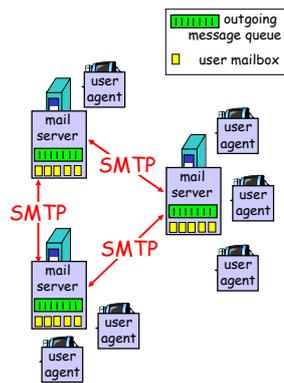
## Electronic Mail

### Three major components:

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

### User Agent

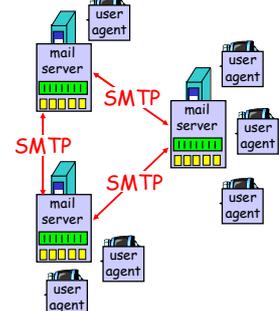
- a.k.a. "mail reader"
- composing, editing, reading mail messages
- e.g., Eudora, Outlook, elm, Mozilla Thunderbird
- outgoing, incoming messages stored on server



## 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



## Format of an e-mail

- Envelope
- Message
  - Header
  - Body

## Addresses

local\_part@domain\_name

rein.paluoja@dcc.ttu.ee

## User Agent (UA)

- ❑ Composing messages
- ❑ Reading messages
- ❑ Replying messages
- ❑ Forwarding messages
- ❑ Handling mailboxes

## 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

2: Application Layer 68

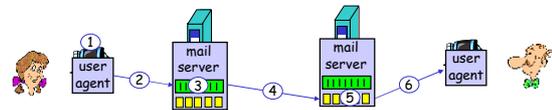
## Electronic Mail: SMTP [RFC 2821]

- ❑ 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

2: Application Layer 69

## 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



2: Application Layer 70

## 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
```

2: Application Layer 71

## 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)

2: Application Layer 72

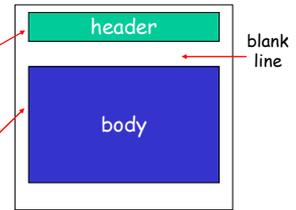
## SMTP: final words

- 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

## Mail message format

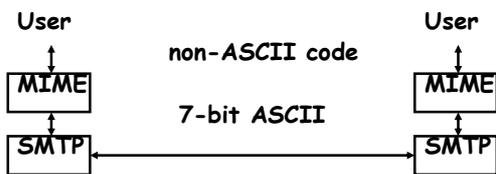
SMTP: protocol for exchanging email msgs  
 RFC 822: standard for text message format:

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



## Multipurpose Internet Mail Extensions (MIME)

... is a supplementary protocol that allows non-ASCII data to be sent through SMTP

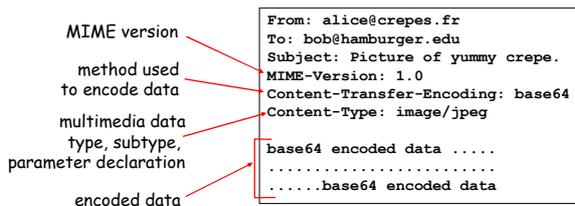


## MIME header

- MIME-version
- Content-type
- Content-transfer-encoding
- Content-id
- Content-description

## 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
- Image**
  - example subtypes: jpeg, gif
- Audio**
  - example subtypes: basic (8-bit mu-law encoded), 32kadpcm (32 kbps coding)
- Video**
  - example subtypes: mpeg, quicktime
- Application**
  - other data that must be processed by reader before "viewable"
  - example subtypes: msworld, octet-stream

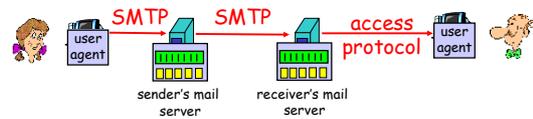
## Multipart Type

```

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

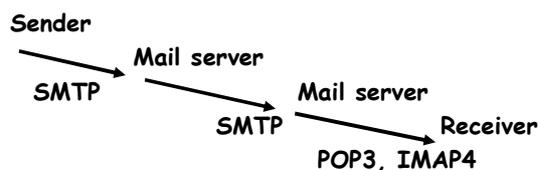
--StartOfNextPart
Dear Bob, Please find a picture of a crepe.
--StartOfNextPart
Content-Transfer-Encoding: base64
Content-Type: image/jpeg
base64 encoded data .....
.....base64 encoded data
--StartOfNextPart
Do you want the recipe?
    
```

## Mail access protocols

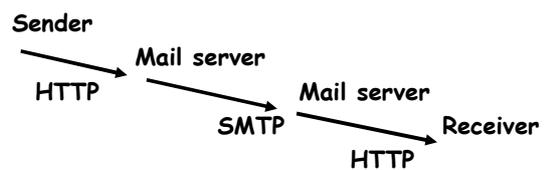


- SMTP: delivery/storage to receiver's server
- Mail access protocol: retrieval from server
  - ❖ POP: Post Office Protocol [RFC 1939]
    - authorization (agent <-->server) and download
  - ❖ IMAP: Internet Mail Access Protocol [RFC 1730]
    - more features (more complex)
    - manipulation of stored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

## Mail delivery



## Web-based mail



## 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
    
```

## 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

## DNS: Domain Name System

**People:** many identifiers:

- ❖ SSN, name, passport #

**Internet hosts, routers:**

- ❖ IP address (32 bit) - used for addressing datagrams
- ❖ "name", e.g., www.yahoo.com - 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"

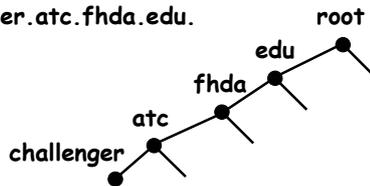
## Name space

- ❑ Flat name space
- ❑ Hierarchical name space

## Domain name

... is a sequence of labels separated by dots

challenger.atc.fhda.edu.



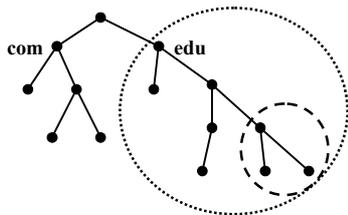
## Domain name

- ❑ Fully Qualified Domain Name (FQDN) - a label is terminated by a null string
- ❑ Partially Qualified Domain Name (PQDN) - a label is not terminated by a null string

partial name = challenger  
suffix = atc.fhda.edu

## Domain

... is a subtree of the domain name space



## DNS

### DNS services

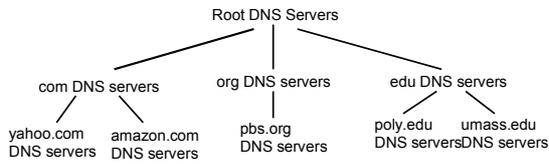
- ❑ hostname to IP address translation
- ❑ host aliasing
  - ❖ Canonical, alias names
- ❑ mail server aliasing
- ❑ load distribution
  - ❖ replicated Web servers: set of IP addresses for one canonical name

### Why not centralize DNS?

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

doesn't *scale!*

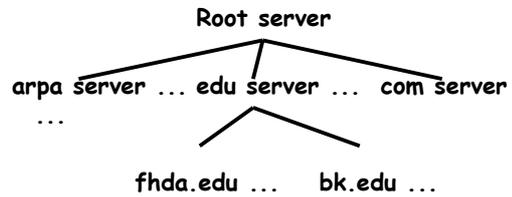
## Distributed, Hierarchical Database



Client wants IP for **www.amazon.com**; 1<sup>st</sup> approx:

- ❑ client queries a root server to find com DNS server
- ❑ client queries com DNS server to get amazon.com DNS server
- ❑ client queries amazon.com DNS server to get IP address for www.amazon.com

## Hierarchy of name servers



## Generic domains

... define registered hosts according to their generic behaviour

- ❑ **com** - commercial organizations
- ❑ **edu** - educational institutions
- ❑ **org** - nonprofit organizations
- ❑ **info** - information service providers
- ❑ **museum** - museums and other nonprofit organizations

## Country domains

- us** - the United States
- ca.us** - California in the United States
- ee** - Estonia
- fi** - Finland

## 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



## TLD and Authoritative Servers

- ❑ **Top-level domain (TLD) servers:**
  - ❖ responsible for com, org, net, edu, etc, and all top-level country domains uk, fr, ca, jp.
  - ❖ Network Solutions maintains servers for com TLD
  - ❖ Educause for edu TLD
- ❑ **Authoritative DNS servers:**
  - ❖ organization's DNS servers, providing authoritative hostname to IP mappings for organization's servers (e.g., Web, mail).
  - ❖ can be maintained by organization or service provider

## Local Name Server

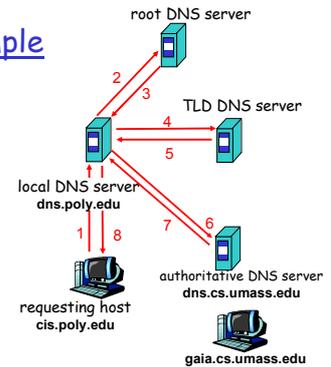
- does not strictly belong to hierarchy
- each ISP (residential ISP, company, university) has one.
  - ❖ also called "default name server"
- when host makes DNS query, query is sent to its local DNS server
  - ❖ acts as proxy, forwards query into hierarchy

## DNS name resolution example

- Host at cis.poly.edu wants IP address for gaia.cs.umass.edu

### iterated query:

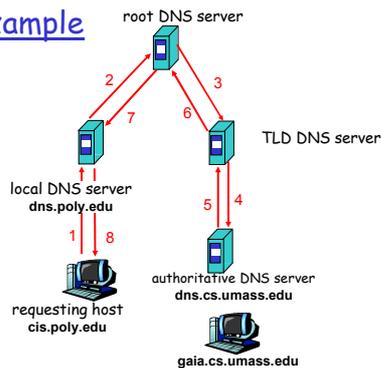
- contacted server replies with name of server to contact
- "I don't know this name, but ask this server"



## DNS name resolution example

### recursive query:

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



## DNS: caching and updating records

- once (any) name server learns mapping, it *caches* mapping
  - ❖ cache entries timeout (disappear) after some time
  - ❖ TLD servers typically cached in local name servers
    - Thus root name servers not often visited
- update/notify mechanisms under design by IETF
  - ❖ RFC 2136
  - ❖ <http://www.ietf.org/html.charters/dnsind-charter.html>

## DNS records

DNS: distributed db storing resource records (RR)

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

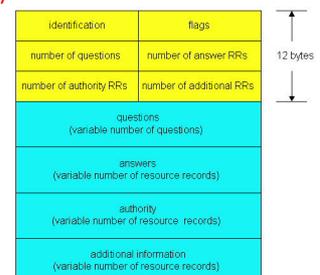
- Type=A
  - ❖ name is hostname
  - ❖ value is IP address
- Type=NS
  - ❖ name is domain (e.g. foo.com)
  - ❖ value is hostname of authoritative name server for this domain
- Type=CNAME
  - ❖ name is alias name for some "canonical" (the real) name
  - ❖ value is canonical name
  - www.ibm.com is really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

## DNS protocol, messages

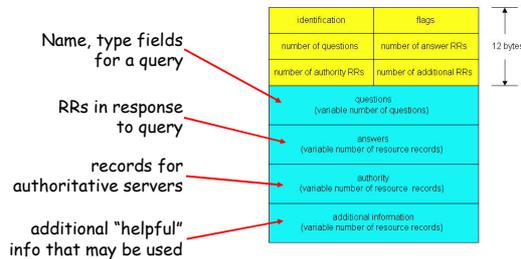
DNS protocol: query and reply messages, both with same *message format*

msg header

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



## DNS protocol, messages



2: Application Layer 103

## Inserting records into DNS

- example: new startup "Network Utopia"
- register name networkutopia.com at *DNS registrar* (e.g., Network Solutions)
  - ❖ provide names, IP addresses of authoritative name server (primary and secondary)
  - ❖ registrar inserts two RRs into com TLD server:

```
(networkutopia.com, dns1.networkutopia.com, NS)
(dns1.networkutopia.com, 212.212.212.1, A)
```

- create authoritative server Type A record for www.networkutopia.com; Type MX record for networkutopia.com
- How do people get IP address of your Web site?

2: Application Layer 104

## 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

2: Application Layer 105

## Sockets

- Family - the protocol group
- Type - stream socket, packet socket or raw socket
- Protocol
- Local socket address - local IP address and the port address of the local application program
- Remote socket address - remote IP address and the port address of the remote application program

## Socket types

- Stream socket - to be used with TCP
- Datagram socket - to be used with UDP
- Raw socket - ICMP ... directly use the services of IP

## Connectionless iterative server

### Server

- Create a socket
- Bind
- Repeat
  - ❖ Receive a request
  - ❖ Process
  - ❖ Send

## Connectionless iterative server

### Client

- Create a socket
- Repeat
  - ❖ Send
  - ❖ Receive
- Destroy

## Connection-oriented concurrent server

### Server

- Create a socket
- Bind
- Listen
- Repeat
  - ❖ Create a child
  - ❖ Create a new socket
  - ❖ Repeating - Read, Process and Write
  - ❖ Destroy socket

## Connection-oriented concurrent server

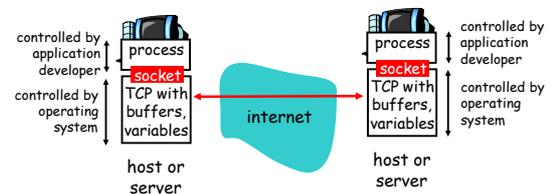
### Client

- Create a socket
- Connect
- Repeat
  - ❖ Write
  - ❖ Read
- Destroy

## 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



2: Application Layer 112

## 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

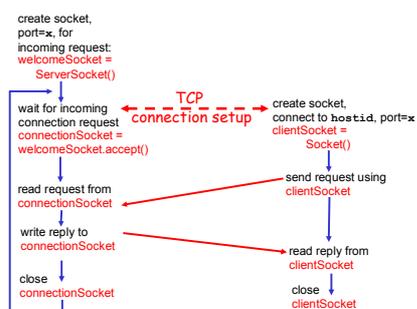
TCP provides reliable, in-order transfer of bytes ("pipe") between client and server

2: Application Layer 113

## Client/server socket interaction: TCP

### Server (running on hostid)

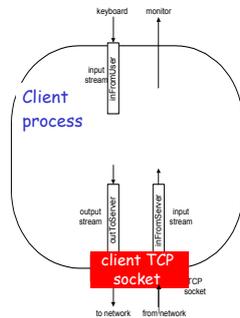
### Client



2: Application Layer 114

## Stream jargon

- 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, e.g., keyboard or socket.
- An **output stream** is attached to an output source, e.g., monitor or socket.



2: Application Layer 115

## 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)

2: Application Layer 116

## Example: Java client (TCP)

```
import java.io.*;
import java.net.*;
class TCPClient {

    public static void main(String argv[]) throws Exception
    {
        String sentence;
        String modifiedSentence;

        Create input stream → BufferedReader inFromUser =
        Create client socket, connect to server → new BufferedReader(new InputStreamReader(System.in));
        Socket clientSocket = new Socket("hostname", 6789);
        Create output stream attached to socket → DataOutputStream outToServer =
        new DataOutputStream(clientSocket.getOutputStream());
```

2: Application Layer 117

## Example: Java client (TCP), cont.

```
        BufferedReader inFromServer =
        Create input stream attached to socket → new BufferedReader(new
        InputStreamReader(clientSocket.getInputStream()));

        sentence = inFromUser.readLine();
        Send line to server → outToServer.writeBytes(sentence + '\n');

        Read line from server → modifiedSentence = inFromServer.readLine();
        System.out.println("FROM SERVER: " + modifiedSentence);
        clientSocket.close();
    }
}
```

2: Application Layer 118

## Example: Java server (TCP)

```
import java.io.*;
import java.net.*;

class TCPServer {

    public static void main(String argv[]) throws Exception
    {
        String clientSentence;
        String capitalizedSentence;

        Create welcoming socket at port 6789 → ServerSocket welcomeSocket = new ServerSocket(6789);
        Wait, on welcoming socket for contact by client → while(true) {
        Socket connectionSocket = welcomeSocket.accept();
        Create input stream, attached to socket → BufferedReader inFromClient =
        new BufferedReader(new
        InputStreamReader(connectionSocket.getInputStream()));
```

2: Application Layer 119

## Example: Java server (TCP), cont

```
        DataOutputStream outToClient =
        Create output stream, attached to socket → new DataOutputStream(connectionSocket.getOutputStream());

        Read in line from socket → clientSentence = inFromClient.readLine();
        capitalizedSentence = clientSentence.toUpperCase() + '\n';
        Write out line to socket → outToClient.writeBytes(capitalizedSentence);
    }
}

End of while loop, loop back and wait for another client connection
```

2: Application Layer 120

## 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

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

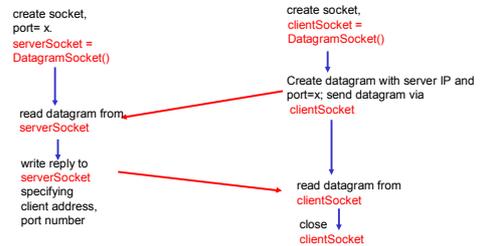
application viewpoint

UDP provides unreliable transfer of groups of bytes ("datagrams") between client and server

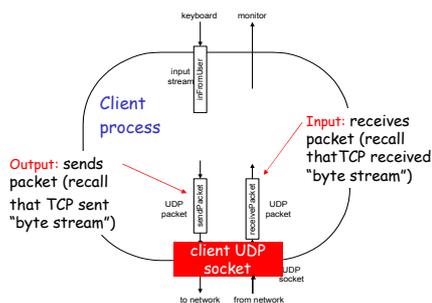
## Client/server socket interaction: UDP

Server (running on `hostid`)

Client



## Example: Java client (UDP)



## Example: Java client (UDP)

```

import java.io.*;
import java.net.*;

class UDPClient {
    public static void main(String args[]) throws Exception
    {
        Create input stream → BufferedReader inFromUser =
        new BufferedReader(new InputStreamReader(System.in));
        Create client socket → DatagramSocket clientSocket = new DatagramSocket();
        Translate hostname to IP address using DNS → InetAddress IPAddress = InetAddress.getByName("hostname");

        byte[] sendData = new byte[1024];
        byte[] receiveData = new byte[1024];

        String sentence = inFromUser.readLine();
        sendData = sentence.getBytes();
    }
}
    
```

## Example: Java client (UDP), cont.

```

Create datagram with data-to-send, length, IP addr, port → DatagramPacket sendPacket =
new DatagramPacket(sendData, sendData.length, IPAddress, 9876);

Send datagram to server → clientSocket.send(sendPacket);

DatagramPacket receivePacket =
new DatagramPacket(receiveData, receiveData.length);

Read datagram from server → clientSocket.receive(receivePacket);

String modifiedSentence =
new String(receivePacket.getData());

System.out.println("FROM SERVER:" + modifiedSentence);
clientSocket.close();
}
    
```

## Example: Java server (UDP)

```

import java.io.*;
import java.net.*;

class UDPServer {
    public static void main(String args[]) throws Exception
    {
        Create datagram socket at port 9876 → DatagramSocket serverSocket = new DatagramSocket(9876);

        byte[] receiveData = new byte[1024];
        byte[] sendData = new byte[1024];

        while(true)
        {
            Create space for received datagram → DatagramPacket receivePacket =
            new DatagramPacket(receiveData, receiveData.length);

            Receive datagram → serverSocket.receive(receivePacket);
        }
    }
}
    
```

## Example: Java server (UDP), cont

```

String sentence = new String(receivePacket.getData());
Get IP addr, port #, of sender → InetAddress IPAddress = receivePacket.getAddress();
int port = receivePacket.getPort();

String capitalizedSentence = sentence.toUpperCase();

Create datagram to send to client → sendData = capitalizedSentence.getBytes();
DatagramPacket sendPacket = new DatagramPacket(sendData, sendData.length, IPAddress, port);

Write out datagram to socket → serverSocket.send(sendPacket);
}
}
}
End of while loop, loop back and wait for another datagram
    
```

2: Application Layer 127

## Chapter 2: Summary

our study of network apps now complete!

- application architectures
  - ❖ client-server
  - ❖ P2P
  - ❖ hybrid
- application service requirements:
  - ❖ reliability, bandwidth, delay
- Internet transport service model
  - ❖ connection-oriented, reliable: TCP
  - ❖ unreliable, datagrams: UDP
- specific protocols:
  - ❖ HTTP
  - ❖ FTP
  - ❖ SMTP, POP, IMAP
  - ❖ DNS
  - ❖ P2P: BitTorrent, Skype
- socket programming

2: Application Layer 128

## Chapter 2: Summary

Most importantly: learned about protocols

- 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
- Important themes:*
  - control vs. data msgs
    - ❖ in-band, out-of-band
  - centralized vs. decentralized
  - stateless vs. stateful
  - reliable vs. unreliable msg transfer
  - "complexity at network edge"

2: Application Layer 129