


Chapter 2

Application Layer



Chapter 2: Application layer

- r 2.1 Principles of network applications
- r 2.2 Web and HTTP
- r 2.3 FTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P application

Chapter 2: Application Layer

Our goals:

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

Some network apps

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


Chapter 2: Application layer

- r 2.1 Principles of network applications
- r 2.2 Web and HTTP
- r 2.3 FTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P applications

Application architectures

- r Client-server
- r Peer-to-peer (P2P)
- r 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

Pure P2P architecture

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

Highly scalable but
difficult to manage



Hybrid of client-server and P2P

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

Processes communicating

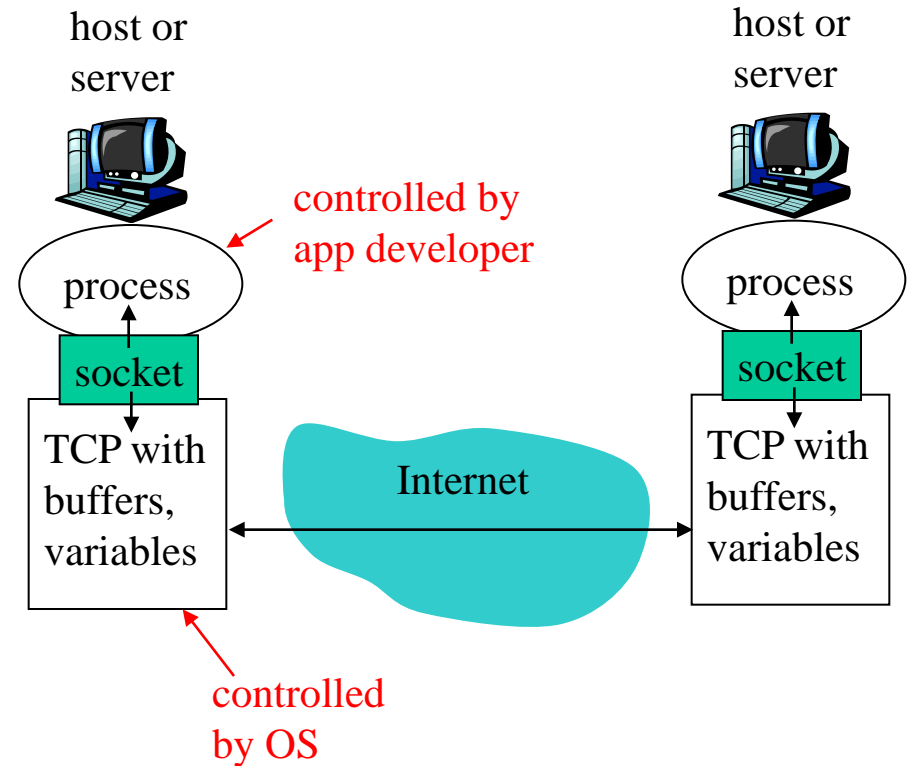
- Process:** program running within a host.
- r within same host, two processes communicate using **inter-process communication** (defined by OS).
 - r processes in different hosts communicate by exchanging **messages**

Client process: process that initiates communication

Server process: process that waits to be contacted

Sockets

- r process sends/receives messages to/from its **socket**
- r API: (1) choice of transport protocol;
(2) ability to fix a few parameters
(lots more on this later)



Addressing processes

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

Addressing processes

- r to receive messages, process must have *identifier*
- r host device has unique 32-bit IP address
- r 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
- r *identifier* includes both IP address and port numbers associated with process on host.
- r Example port numbers:
 - ❖ HTTP server: 80
 - ❖ Mail server: 25
- r to send HTTP message to gaia.cs.umass.edu web server:
 - ❖ IP address: 128.119.245.12
 - ❖ Port number: 80
- r more shortly...

App-layer protocol defines

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

Public-domain protocols:

- r defined in RFCs
- r allows for interoperability
- r e.g., HTTP, SMTP

Proprietary protocols:

- r e.g., Skype

What transport service does an app need?

Data loss

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

Timing

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

Throughput

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

Security

- r Encryption, data integrity, ...

Transport service requirements of common apps

Application	Data loss	Throughput	Time Sensitive
file transfer		elastic	
e-mail		elastic	
Web documents		elastic	
real-time audio/video		audio: 5kbps-1Mbps video:10kbps-5Mbps	
stored audio/video		same as above	
interactive games		few kbps up	
instant messaging		elastic	

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	
interactive games	loss-tolerant	few kbps up	yes, few secs
instant messaging	no loss	elastic	yes, 100' s msec
			yes and no

Internet transport protocols services

TCP service:

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

UDP service:

- r unreliable data transfer between sending and receiving process
- r 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	
remote terminal access		
	Web	
	file transfer	
streaming multimedia		
Internet telephony		

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

Chapter 2: Application layer

- r 2.1 Principles of network applications
 - ❖ app architectures
 - ❖ app requirements
- r 2.2 Web and HTTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P applications
- r 2.7 Socket programming with TCP
- r 2.8 Socket programming with UDP

Web and HTTP

First some jargon

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

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


host name

path name

HTTP overview

HTTP: hypertext transfer protocol

- r Web's application layer protocol
- r 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:

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

HTTP is “stateless”

- r server maintains no information about past client requests

Uploading form input

Post method:

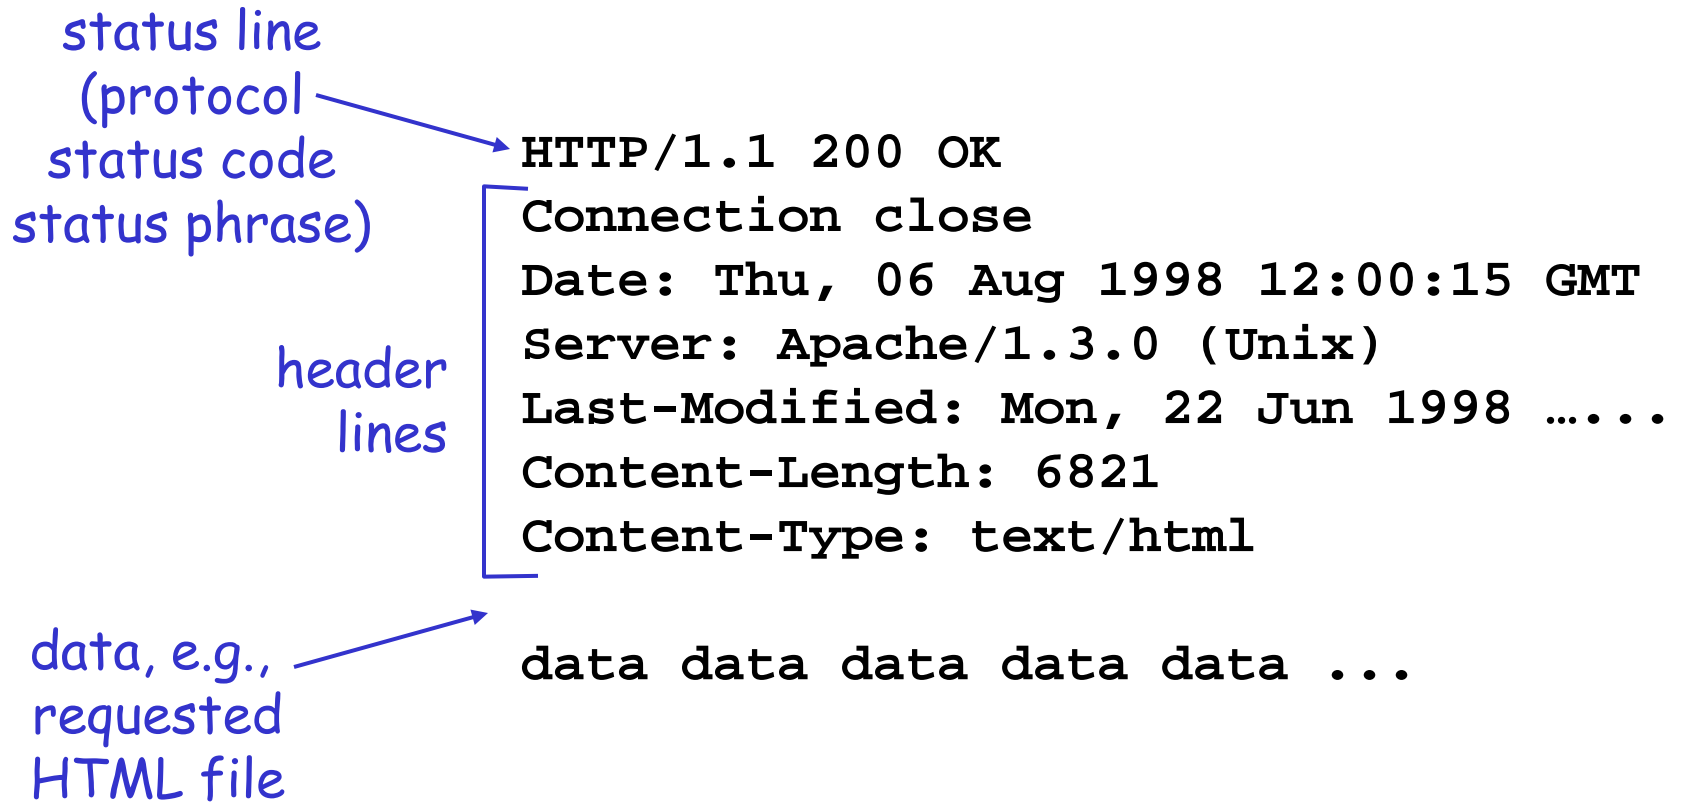
- r Web page often includes form input
- r Input is uploaded to server in entity body

URL method:

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

`www.somesite.com/animalsearch?monkeys&banana`

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

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:

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

Cookies: keeping "state" (cont.)

client

server



usual http request msg

usual http response
Set-cookie: 1678

usual http request msg
cookie: 1678

usual http response msg

usual http request msg
cookie: 1678

usual http response msg

Amazon server
creates ID
1678 for user


cookie-
specific
action

cookie-
specific
action

create
entry

access

access



backend
database

one week later:

Cookies (continued)

What cookies can bring:

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

How to keep “state”:

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

— aside —


Cookies and privacy:

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

Web caches (proxy server)

Goal: satisfy client request without involving origin server

- r user sets browser:
Web accesses via cache
- r 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

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


Why Web caching?

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

Chapter 2: Application layer

- r 2.1 Principles of network applications
- r 2.2 Web and HTTP
- r 2.3 FTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P applications


FTP: the file transfer protocol



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

FTP: separate control, data connections

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



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

Chapter 2: Application layer

- r 2.1 Principles of network applications
- r 2.2 Web and HTTP
- r 2.3 FTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P applications

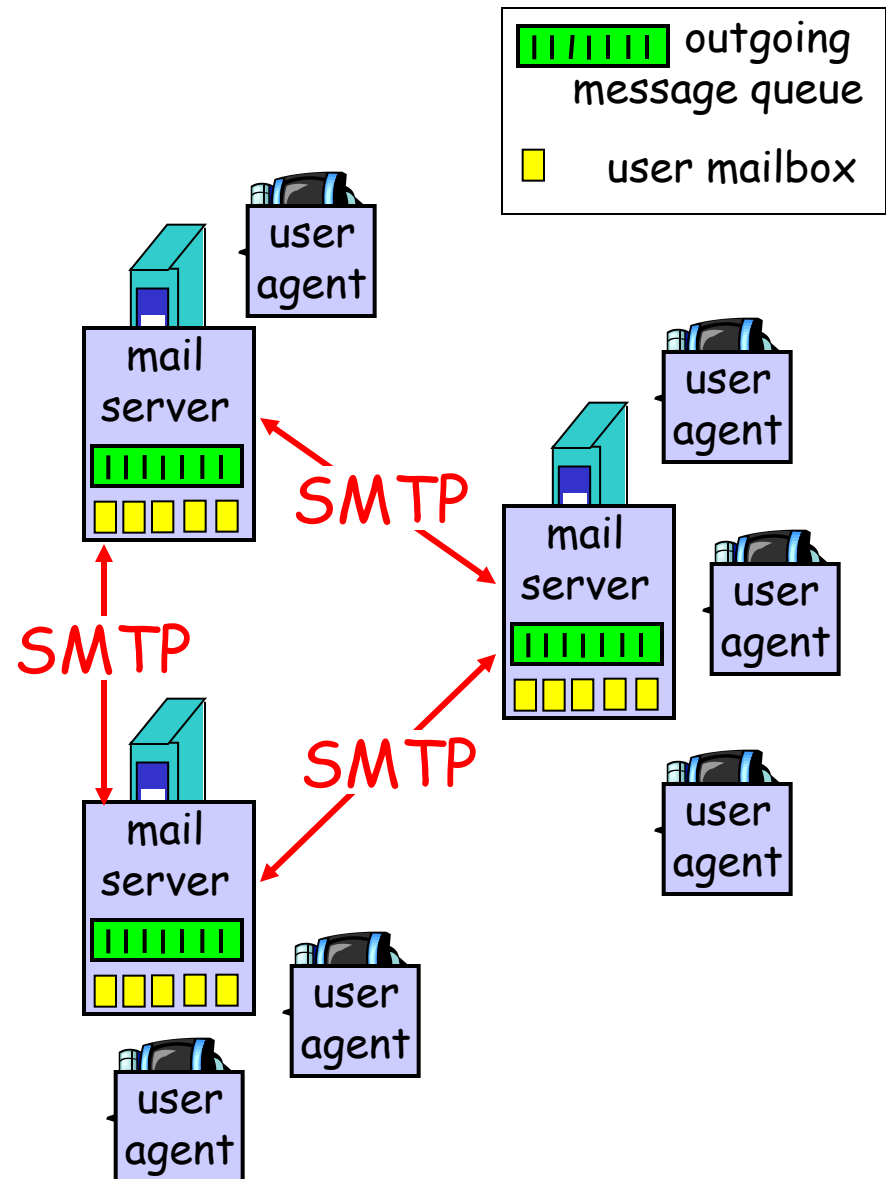
Electronic Mail

Three major components:

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

User Agent


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



Electronic Mail: mail servers

Mail Servers

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

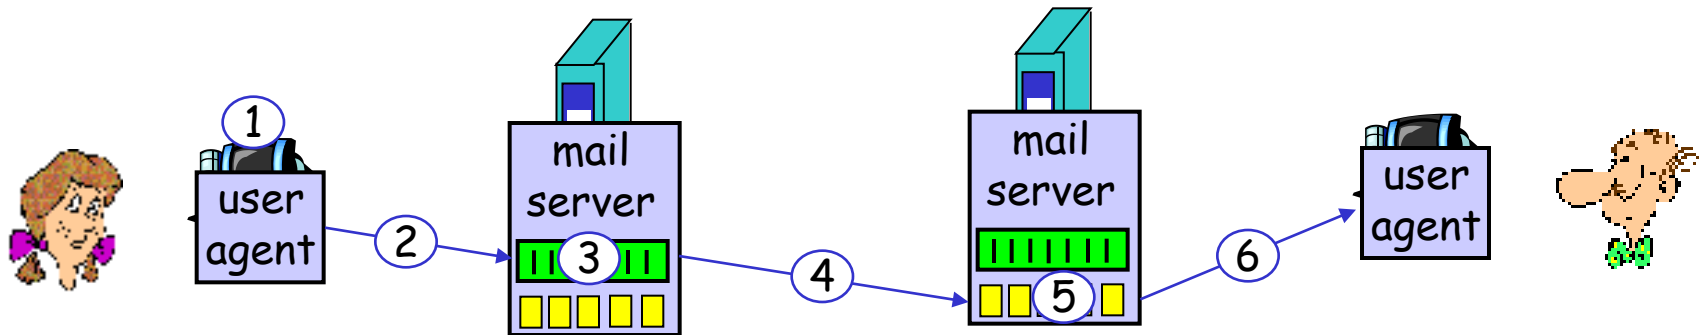


Electronic Mail: SMTP [RFC 2821]

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

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



Mail message format

SMTP: protocol for exchanging email msgs

RFC 822: standard for text message format:


r header lines, e.g.,

- ❖ To:
- ❖ From:
- ❖ Subject:

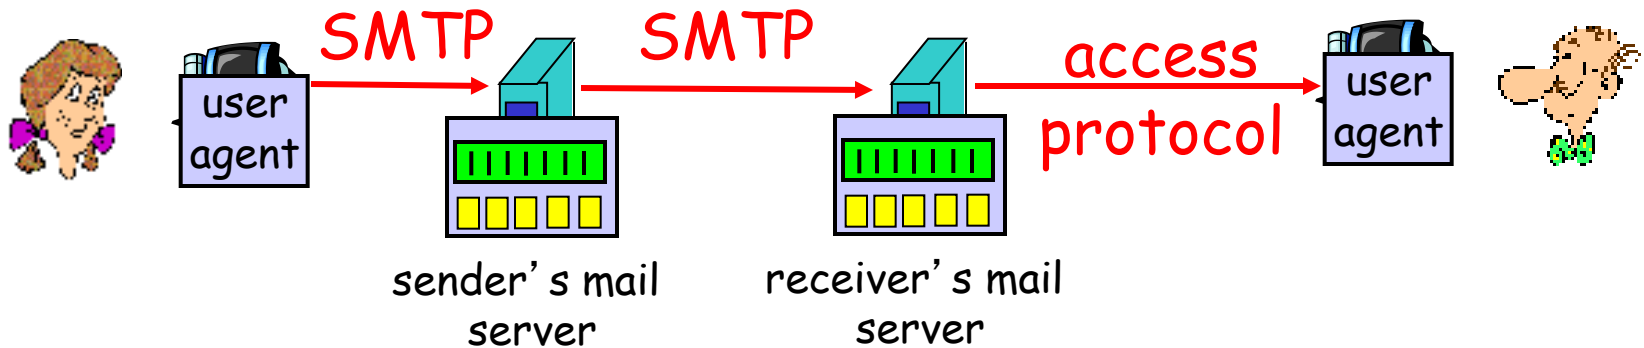
different from SMTP commands!

r body

- ❖ the “message”, ASCII characters only



Mail access protocols



- r SMTP: delivery/storage to receiver's server
- r 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.

Chapter 2: Application layer

- r 2.1 Principles of network applications
- r 2.2 Web and HTTP
- r 2.3 FTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r **2.5 DNS**
- r 2.6 P2P applications

DNS: Domain Name System

People: many identifiers:

- ❖ SSN, name, passport #

Internet hosts, routers:

- ❖ IP address (32 bit) - used for addressing datagrams
- ❖ “name”, e.g.,
ww.yahoo.com - used by humans

Q: map between IP addresses and name ?

Domain Name System:

- r *distributed database*
implemented in hierarchy of many *name servers*
- r *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

DNS services


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

Why not centralize DNS?

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

doesn't *scale!*

Distributed, Hierarchical Database

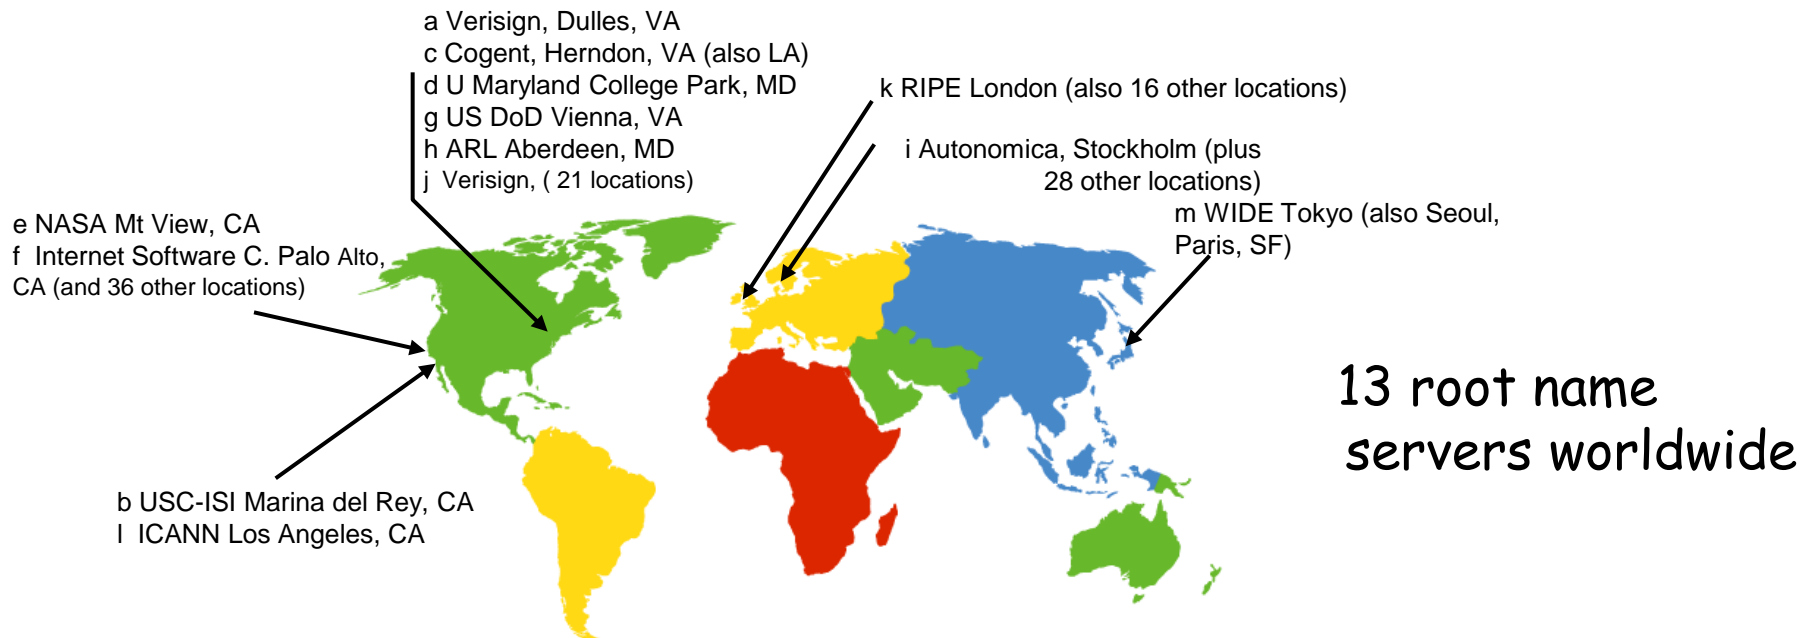


Client wants IP for www.amazon.com; 1st approx:

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

DNS: Root name servers

- r contacted by local name server that can not resolve name
- r 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

TLD and Authoritative Servers

- r **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
- r **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

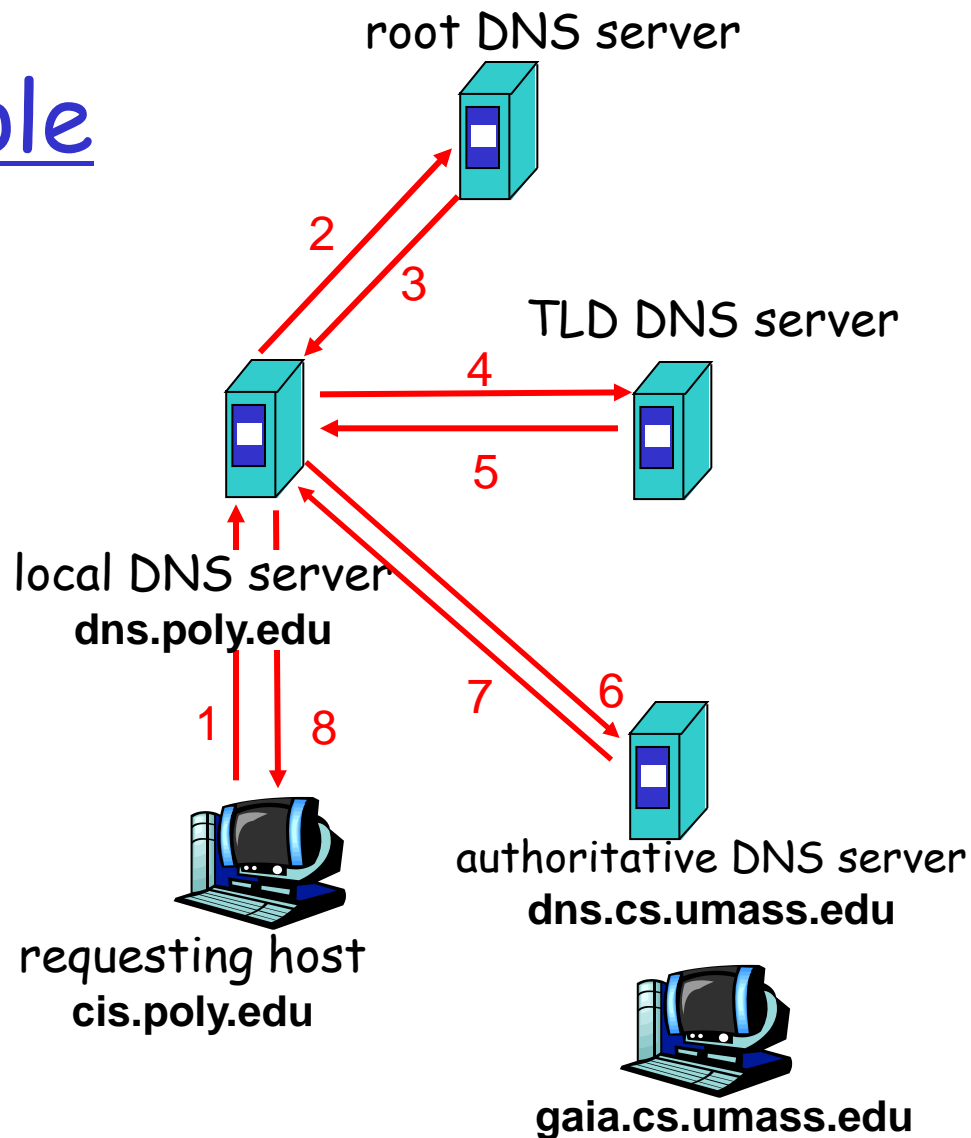
- r does not strictly belong to hierarchy
- r each ISP (residential ISP, company, university) has one.
 - ❖ also called “default name server”
- r when host makes DNS query, query is sent to its local DNS server
 - ❖ acts as proxy, forwards query into hierarchy

DNS name resolution example

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

iterated query:

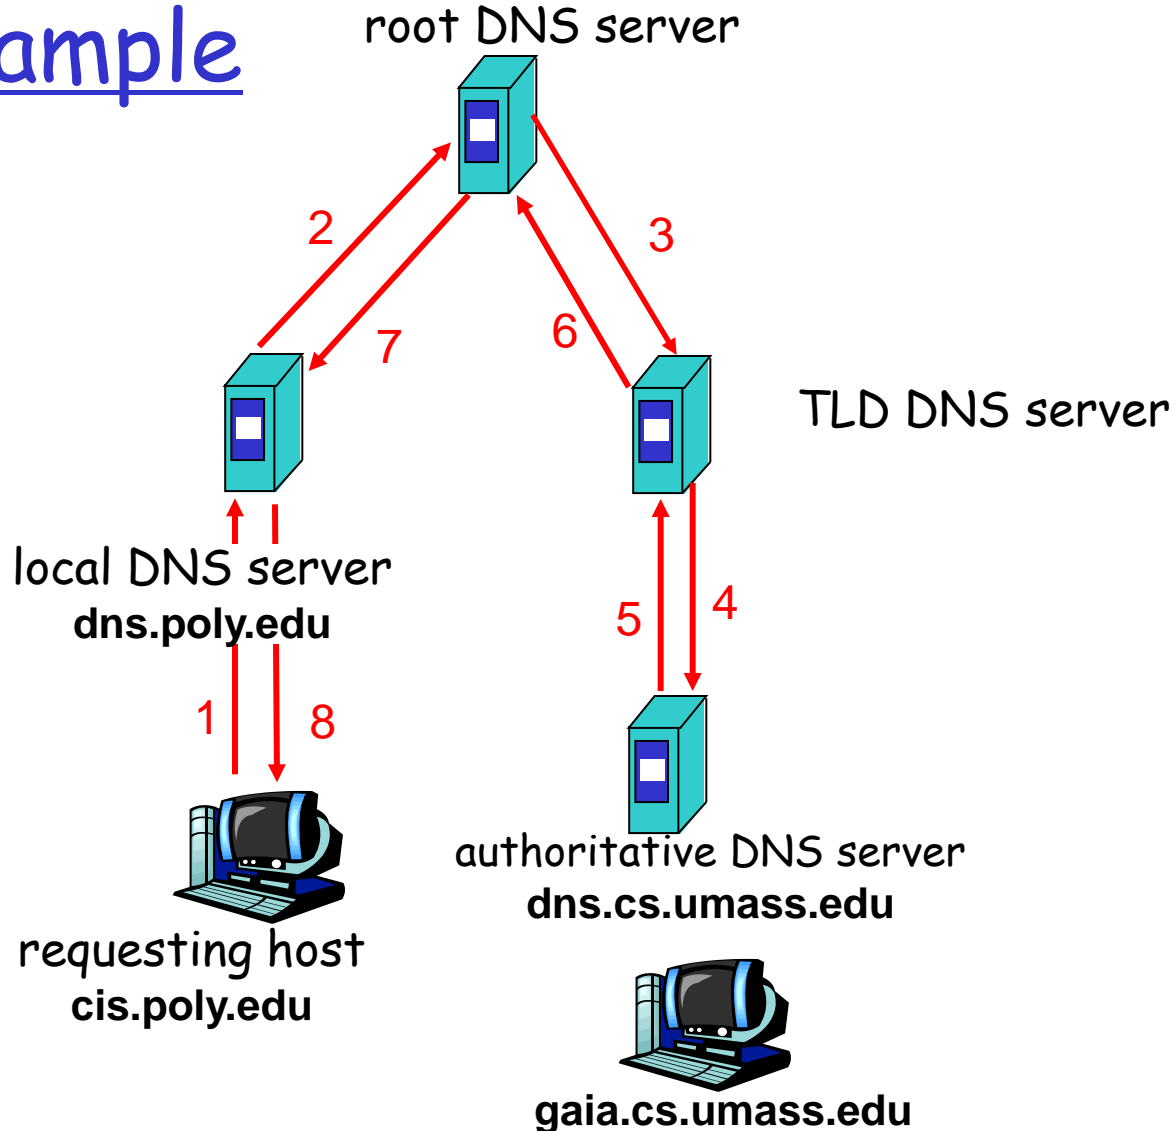
- r contacted server replies with name of server to contact
- r “I don’t know this name, but ask this server”



DNS name resolution example

recursive query:

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



DNS: caching and updating records

- r 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
- r update/notify mechanisms under design by IETF
 - ❖ RFC 2136
 - ❖ <http://www.ietf.org/html.charters/dnsind-charter.html>

Inserting records into DNS

- r example: new startup “Network Utopia”
- r 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)
```
- r create authoritative server Type A record for `www.networkutopia.com`; Type MX record for `networkutopia.com`
- r *How do people get IP address of your Web site?*

Chapter 2: Application layer

- r 2.1 Principles of network applications
 - ❖ app architectures
 - ❖ app requirements
- r 2.2 Web and HTTP
- r 2.4 Electronic Mail
 - ❖ SMTP, POP3, IMAP
- r 2.5 DNS
- r 2.6 P2P applications


Pure P2P architecture

- r *no* always-on server
- r arbitrary end systems directly communicate
- r peers are intermittently connected and change IP addresses
- r Three topics:
 - ❖ File distribution
 - ❖ Searching for information
 - ❖ Case Study: Skype




File Distribution: Server-Client vs P2P

Question: How much time to distribute file from one server to N peers?



File distribution time: server-client

- r server sequentially sends N copies:
 - ❖ NF/u_s time
- r client i takes F/d_i time to download




Time to distribute F to N clients using client/server approach = $d_{cs} = \max \{ NF/u_s, F/\min_i(d_i) \}$

increases linearly in N (for large N)

File distribution time: P2P


- r server must send one copy: F/u_s time
- r client i takes F/d_i time to download
- r NF bits must be downloaded (aggregate)
 - r fastest possible upload rate: $u_s + \sum u_i$



$$d_{\text{P2P}} = \max \left\{ F/u_s, F/\min(d_i)_i, NF/(u_s + \sum u_i) \right\}$$

Server-client vs. P2P: example

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




File distribution: BitTorrent


r P2P file distribution

tracker: tracks peers participating in torrent

torrent: group of peers exchanging chunks of a file



BitTorrent



- r file divided into 256KB *chunks*.
- r peer joining torrent:
 - ❖ has no chunks, but will accumulate them over time
 - ❖ registers with tracker to get list of peers, connects to subset of peers (“neighbors”)
- r while downloading, peer uploads chunks to other peers.
- r peers may come and go
- r once peer has entire file, it may (selfishly) leave or (altruistically) remain

P2P: searching for information

Index in P2P system: maps information to peer location
(location = IP address & port number)

File sharing (eg e-mule)

- r Index dynamically tracks the locations of files that peers share.
- r Peers need to tell index what they have.
- r Peers search index to determine where files can be found

Instant messaging

- r Index maps user names to locations.
- r When user starts IM application, it needs to inform index of its location
- r Peers search index to determine IP address of user.

P2P: centralized index

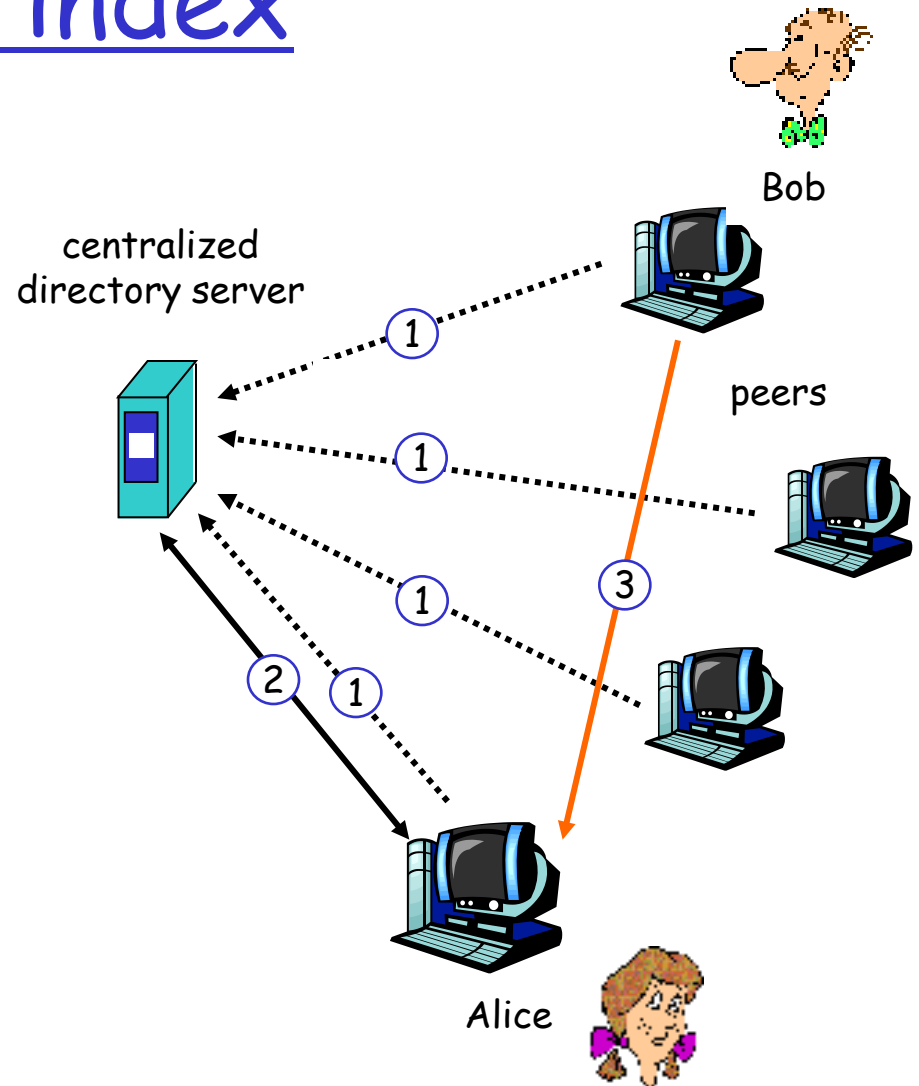
original “Napster” design

1) when peer connects, it informs central server:

- ❖ IP address
- ❖ content

2) Alice queries for “Hey Jude”

3) Alice requests file from Bob



P2P: problems with centralized directory

- r single point of failure
- r performance bottleneck
- r copyright infringement:
“target” of lawsuit is
obvious

file transfer is
decentralized, but
locating content is
highly centralized

Query flooding


- r fully distributed
 - ❖ no central server
- r used by Gnutella
- r Each peer indexes the files it makes available for sharing (and no other files)

overlay network: graph

- r edge between peer X and Y if there's a TCP connection
- r all active peers and edges form overlay net
- r edge: virtual (*not* physical) link
- r given peer typically connected with < 10 overlay neighbors

Query flooding


- r Query message sent over existing TCP connections
- r peers forward Query message
- r QueryHit sent over reverse path



Scalability:
limited scope
flooding


Hierarchical Overlay

- r between centralized index, query flooding approaches
- r each peer is either a *super node* or assigned to a super node
 - ❖ TCP connection between peer and its super node.
 - ❖ TCP connections between some pairs of super nodes.
- r Super node tracks content in its children



P2P Case study: Skype

- r inherently P2P: pairs of users communicate.
- r proprietary application-layer protocol (inferred via reverse engineering)
- r hierarchical overlay with SNs
- r Index maps usernames to IP addresses; distributed over SNs



Chapter 2: Summary

our study of network apps now complete!

r application architectures

- ❖ client-server
- ❖ P2P
- ❖ hybrid

r application service requirements:

- ❖ reliability, bandwidth, delay

r Internet transport service model

- ❖ connection-oriented, reliable: TCP
- ❖ unreliable, datagrams: UDP

r specific protocols:

- ❖ HTTP
- ❖ FTP
- ❖ SMTP, POP, IMAP
- ❖ DNS
- ❖ P2P: BitTorrent, Skype