Reti di Elaboratori

Corso di Laurea in Informatica

Università degli Studi di Roma "La Sapienza"

Canale A-L

Prof.ssa Chiara Petrioli

Chapter 2: outline

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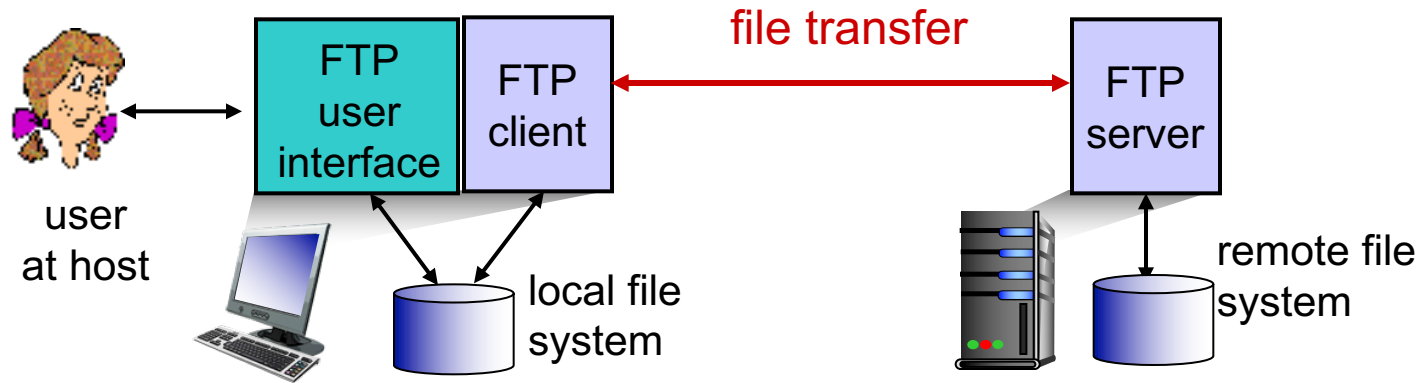
- ❖ transfer file to/from remote host
- ❖ client/server model
  - ▪ *client:* side that initiates transfer (either to/from remote)
  - ▪ *server:* remote host
- ❖ ftp: RFC 959
- ❖ ftp server: port 21

FTP: separate control, data connections

- FTP client contacts FTP server at port 21, using TCP

- client authorized over control connection

- client browses remote directory, sends commands over control connection

- when server receives file transfer command, *server* opens 2*nd* TCP data connection (for file) *to* client

- after transferring one file, server closes data connection

TCP control connection,
server port 21

TCP data connection,
server port 20

FTP
client

FTP
server

- 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**

Chapter 2: outline

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., Outlook, Thunderbird, iPhone mail client

❑ outgoing, incoming messages stored on server

outgoing message queue

user mailbox

user agent

mail server

SMTP

SMTP

mail server

user agent

user agent

SMTP

mail server

user agent

user agent

user agent
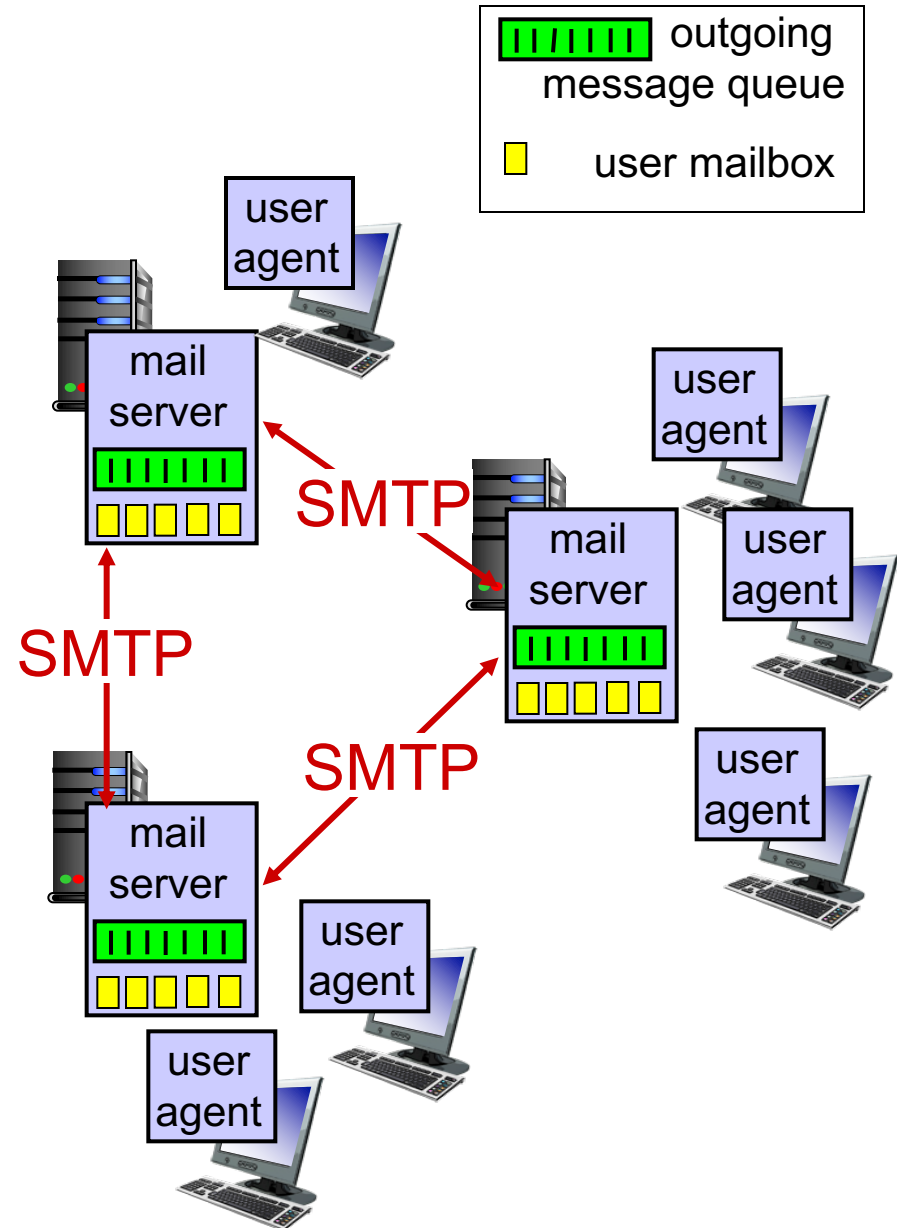
Scenario: Alice sends message to Bob

1) Alice uses UA to compose message "to" `bob@someschool.edu`
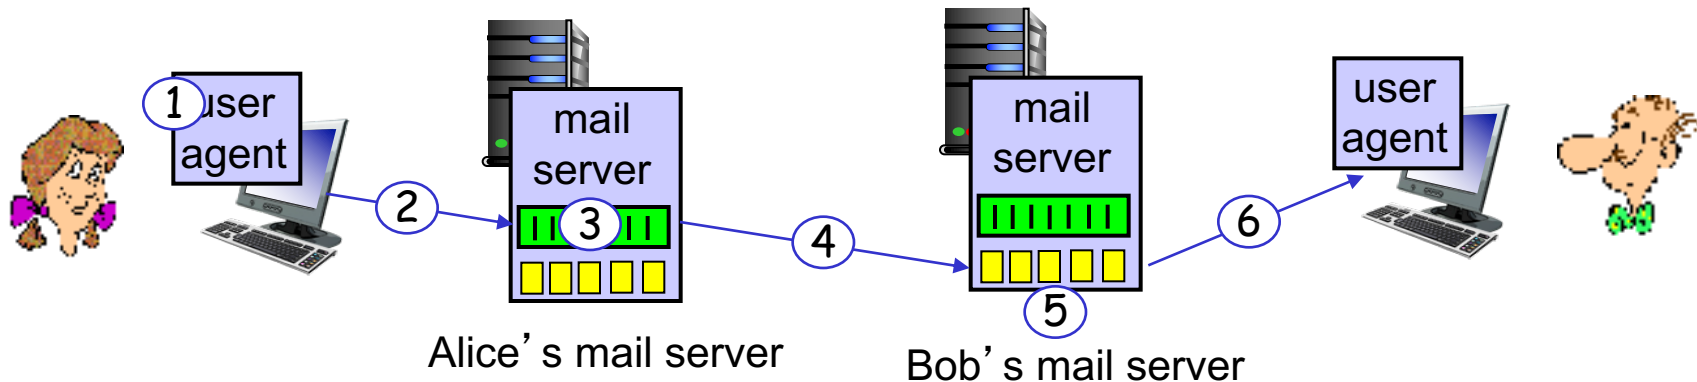
2) Alice's UA sends message to her mail server; message placed in message queue

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

Alice's mail server

Bob's mail 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

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 (like HTTP, FTP)
  ○ commands: ASCII text
  ○ response: status code and phrase

❑ messages must be in 7-bit ASCI

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
```

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:
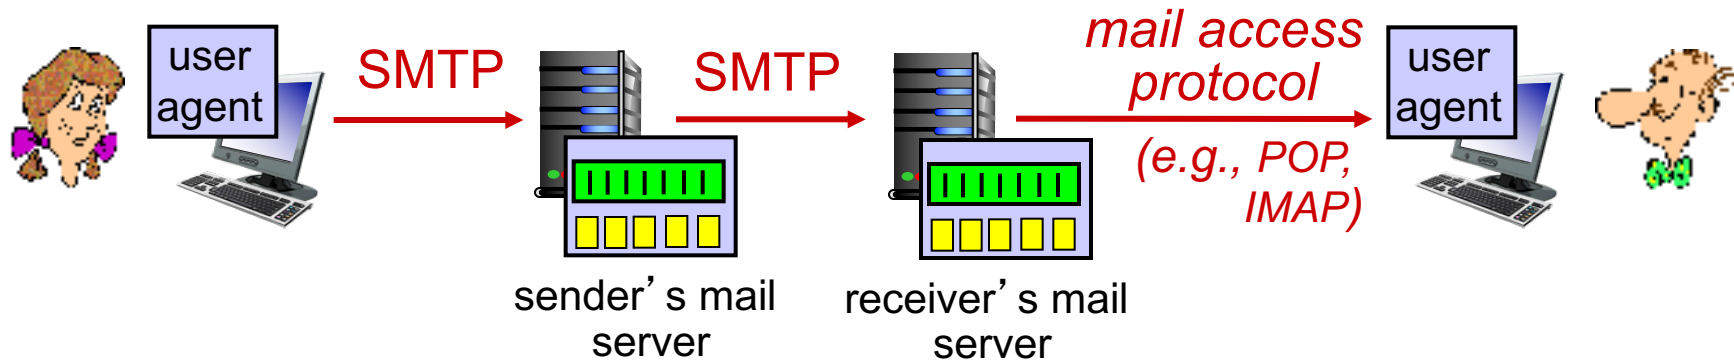
❑ header lines, e.g.,
  ○ To:
  ○ From:
  ○ Subject:

  different from SMTP MAIL FROM, RCPT TO: commands!

❑ Body: the "message"
  ○ ASCII characters only

header

blank line

body

Mail access protocols

- ❑ **SMTP:** delivery/storage to receiver's server
- ❑ mail access protocol: retrieval from server
  - ○ **POP:** Post Office Protocol [RFC 1939]: authorization, download
  - ○ **IMAP:** Internet Mail Access Protocol [RFC 1730]: more features, including manipulation of stored msgs on server
  - ○ **HTTP:** gmail, 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
```

POP3 (more) and IMAP

more about POP3

- previous example uses POP3 "download and delete" mode
  - Bob cannot re-read e-mail if he changes client
- POP3 "download-and-keep": copies of messages on different clients
- POP3 is stateless across sessions

IMAP

- keeps all messages in one place: at server
- allows user to organize messages in folders
- keeps user state across sessions:
  - names of folders and mappings between message IDs and folder name

Chapter 2: outline

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
  - ○ server for original requesting client
  - ○ client to origin 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 (so too does P2P file sharing)

Caching example:

assumptions:

- avg object size: 100K bits
- avg request rate from browsers to origin servers: 15/sec
- avg data rate to browsers: 1.50 Mbps
- RTT from institutional router to any origin server: 2 sec
- access link rate: 1.54 Mbps

consequences: *problem!*

- LAN utilization: 0,15%
- access link utilization = 99%
- total delay   = Internet delay + access delay + LAN delay
   = 2 sec + minutes + usecs

origin servers

public Internet

1.54 Mbps access link

institutional network

1 Gbps LAN

Caching example: fatter access link

assumptions:
- avg object size: 100K bits
- avg request rate from browsers to origin servers: 15/sec
- avg data rate to browsers: 1.50 Mbps → 154 Mbps
- RTT from institutional router to any origin server: 2 sec
- access link rate: 1.54 Mbps → 9.9%

consequences:
- LAN utilization: 0,15%
- access link utilization = 9,9%
- total delay = Internet delay + access delay + LAN delay
  = 2 sec + minutes + usecs → msecs

Cost: increased access link speed (not cheap!)

origin servers

public Internet

1.54 Mbps access link → 154 Mbps

institutional network

1 Gbps LAN

# Caching example: install local cache

assumptions:

- avg object size: 100K bits
- avg request rate from browsers to origin servers:15/sec
- avg data rate to browsers: 1.50 Mbps
- RTT from institutional router to any origin server: 2 sec
- access link rate: 1.54 Mbps

consequences:

- LAN utilization: 0,15%
- access link utilization = 100%
- total delay = Internet delay + access delay + LAN delay
  = 2 sec + n ? ?

How to compute link utilization, delay?

Cost: web cache (cheap!)

origin servers

public Internet

1.54 Mbps access link

institutional network

1 Gbps LAN

local web cache

Caching example: install local cache

Calculating access link

utilization, delay with cache:

❑ suppose cache hit rate is 0.4
  - ○ 40% requests satisfied at cache,

    60% requests satisfied at origin
  - ■ access link utilization:
    - ▪ 60% of requests use access link
  - ■ data rate to browsers over access link
    = 0.6*1.50 Mbps = .9 Mbps
    - ▪ utilization = 0.9/1.54 = .58
  - ■ total delay
    - ▪ = 0.6 * (delay from origin servers) +0.4 * (delay when satisfied at cache)
    - ▪ = 0.6 (2.01) + 0.4 (~msecs) = ~ 1.2 secs
    - ▪ less than with 154 Mbps link (and cheaper too!)

origin servers

public Internet

1.54 Mbps access link

institutional network

1 Gbps LAN

local web cache

Content distribution networks

❑*challenge:* how to stream content (selected from millions of videos) to hundreds of thousands of *simultaneous* users?

❑*option 1:* single, large "mega-server"
  ❍ single point of failure
  ❍ point of network congestion
  ❍ long path to distant clients
  ❍ multiple copies of video sent over outgoing link

….quite simply: this solution *doesn't scale*

Content Delivery Networks

☐ We have seen the extensive use of caching for reducing latencies in resolving names and accessing web content

☐ Is this enough?
- Origin servers may still have to be accessed to maintain consistency

☐ Caching
- What to cache
- How to maintain consistency
- How to invalidate or update in case an inconsistency is detected

☐ More here:http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.73.586&rep=rep1&type=pdf

Content Delivery Networks


(a) Traditional centralized architecture   (b) Distributed CDN architecture

- improving client-perceived response time by bringing content closer to the network edge, and thus closer to end-users

- off-loading work from origin servers by serving larger objects, such as images and multimedia, from multiple CDN servers

- reducing content provider costs by reducing the need to invest in more powerful servers or more bandwidth as user population increases

Content Delivery Networks

(a) Traditional centralized architecture

(b) Distributed CDN architecture

- improving site availability by replicating content in many distributed locations

Content Delivery Networks

(a) CDN architectural elements

(b) CDN request-routing

Content Delivery Networks

☐ HTTP Redirect

☐ DNS Redirect

Content distribution networks

❑ *challenge:* how to stream content (selected from millions of videos) to hundreds of thousands of simultaneous users?

❑ *option 2:* store/serve multiple copies of videos at multiple geographically distributed sites *(CDN)*

  ○ *enter deep:* push CDN servers deep into many access networks
    • close to users
    • used by Akamai, 1700 locations
  ○ *bring home:* smaller number (10's) of larger clusters in POPs near (but not within) access networks
    • used by Limelight

Content Distribution Networks (CDNs)

- CDN: stores copies of content at CDN nodes
  - e.g. Netflix stores copies of MadMen

- subscriber requests content from CDN
  - directed to nearby copy, retrieves content
  - may choose different copy if network path congested

Content Distribution Networks (CDNs)

"over the top"

Internet host-host communication as a service
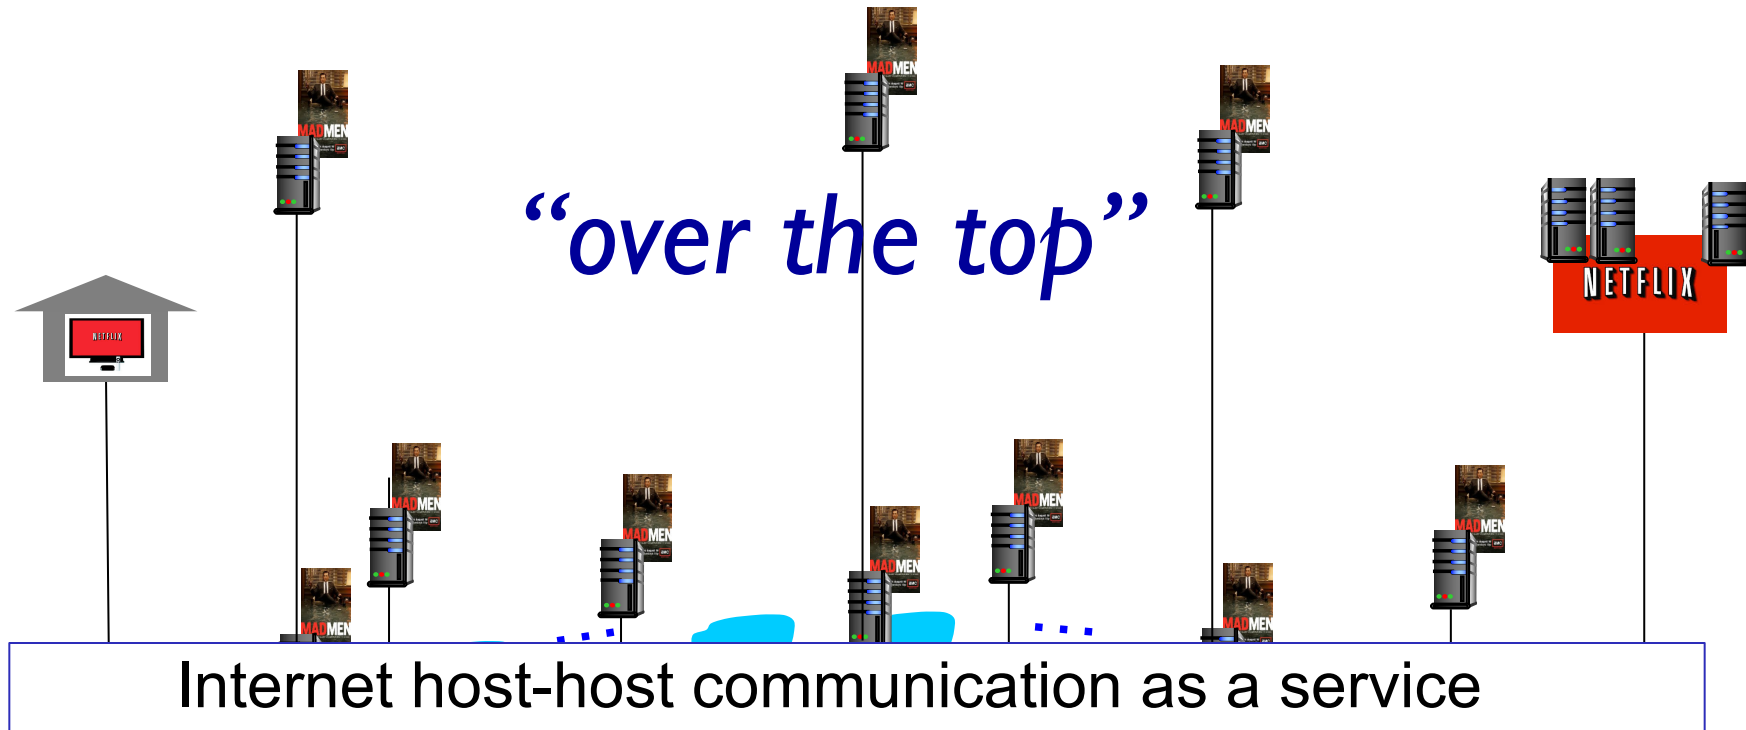
OTT challenges: coping with a congested Internet
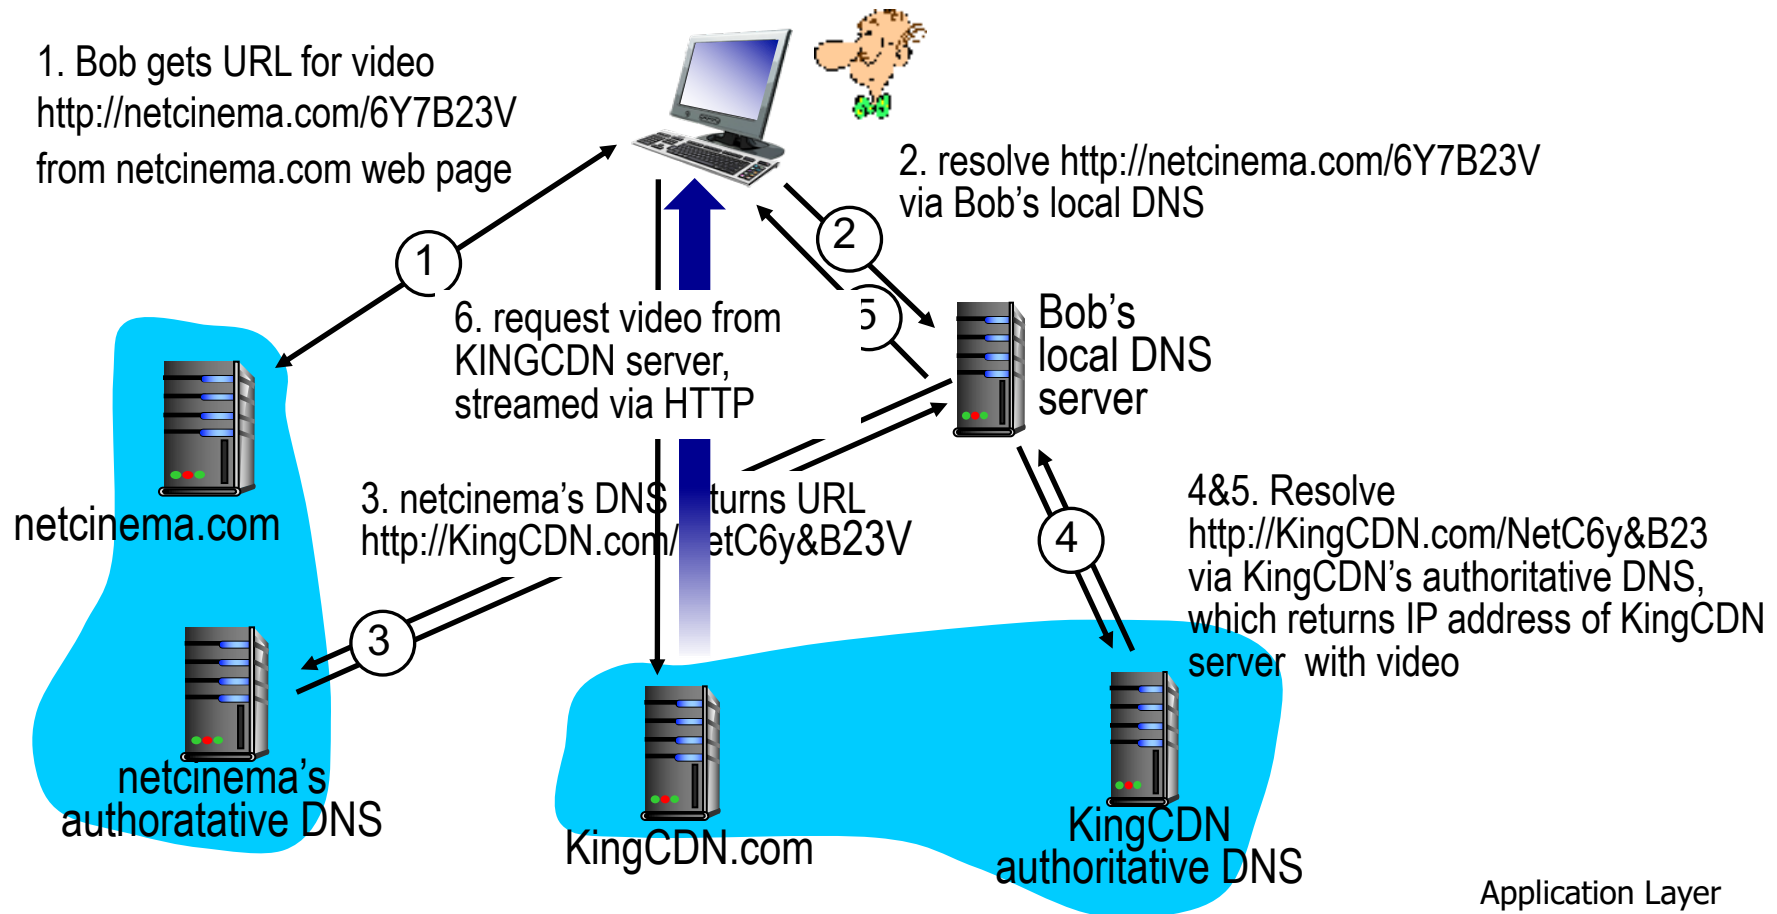- from which CDN node to retrieve content?
- viewer behavior in presence of congestion?
- what content to place in which CDN node?

more .. in chapter 7

CDN content access: a closer look

Bob (client) requests video http://netcinema.com/6Y7B23V
- video stored in CDN at http://KingCDN.com/NetC6y&B23V

1. Bob gets URL for video
http://netcinema.com/6Y7B23V
from netcinema.com web page

2. resolve http://netcinema.com/6Y7B23V
via Bob's local DNS

6. request video from
KINGCDN server,
streamed via HTTP

Bob's
local DNS
server

netcinema.com

3. netcinema's DNS returns URL
http://KingCDN.com/NetC6y&B23V

4&5. Resolve
http://KingCDN.com/NetC6y&B23
via KingCDN's authoritative DNS,
which returns IP address of KingCDN
server with video

netcinema's
authoratative DNS

KingCDN.com

KingCDN
authoritative DNS

Case study: Netflix

Netflix registration, accounting servers

Amazon cloud

upload copies of multiple versions of video to CDN servers

CDN server

CDN server

CDN server

1. Bob manages Netflix account

2. Bob browses Netflix video

3. Manifest file returned for requested video

4. DASH streaming

Video Streaming and CDNs: context

- video traffic: major consumer of Internet bandwidth
  - Netflix, YouTube: 37%, 16% of downstream residential ISP traffic
  - ~1B YouTube users, ~75M Netflix users
- challenge: scale - how to reach ~1B users?
  - single mega-video server won't work (why?)
- challenge: heterogeneity
  - different users have different capabilities (e.g., wired versus mobile; bandwidth rich versus bandwidth poor)
- *solution:* distributed, application-level infrastructure

Multimedia: video

- video: sequence of images displayed at constant rate
  - e.g., 24 images/sec
- digital image: array of pixels
  - each pixel represented by bits
- coding: use redundancy *within* and *between* images to decrease # bits used to encode image
  - spatial (within image)
  - temporal (from one image to next)

spatial coding example: instead of sending *N* values of same color (all purple), send only two values: color value (*purple*) and number of repeated values (*N*)

frame *i*

temporal coding example: instead of sending complete frame at i+1, send only differences from frame i


frame *i+1*

Multimedia: video

- **CBR: (constant bit rate):** video encoding rate fixed

- **VBR: (variable bit rate):** video encoding rate changes as amount of spatial, temporal coding changes

- **examples:**
  - MPEG 1 (CD-ROM) 1.5 Mbps
  - MPEG2 (DVD) 3-6 Mbps
  - MPEG4 (often used in Internet, < 1 Mbps)

spatial coding example: instead of sending *N* values of same color (all purple), send only two values: color value (*purple*) and number of repeated values (*N*)


frame *i*

temporal coding example: instead of sending complete frame at i+1, send only differences from frame i

frame *i+1*

Streaming stored video:

simple scenario:

video server
(stored video)

Internet

client

Streaming multimedia: DASH

❑ *DASH: D*ynamic, *A*daptive *S*treaming over *H*TTP

❑ *server:*

  ○ divides video file into multiple chunks

  ○ each chunk stored, encoded at different rates

  ○ *manifest file:* provides URLs for different chunks

❑ *client:*

  ○ periodically measures server-to-client bandwidth

  ○ consulting manifest, requests one chunk at a time

   • chooses maximum coding rate sustainable given current bandwidth

   • can choose different coding rates at different points in time (depending on available bandwidth at time)

Streaming multimedia: DASH

❑ *DASH: Dynamic, Adaptive Streaming over HTTP*

❑ *"intelligence"* at client: client determines

- ○ *when* to request chunk (so that buffer starvation, or overflow does not occur)

- ○ *what encoding rate* to request (higher quality when more bandwidth available)

- ○ *where* to request chunk (can request from URL server that is "close" to client or has high available bandwidth)

Case study: Netflix

Amazon cloud

upload copies of multiple versions of video to CDN servers

CDN server

CDN server

CDN server

Netflix registration, accounting servers

3. Manifest file returned for requested video

2. Bob browses Netflix video

②
③

①

1. Bob manages Netflix account

4. DASH streaming