# Multimedia Networking

Computer Networks, Winter 2010/2011

Chapter 5 outline

o 5.1 Multimedia networking applications

o 5.2 Streaming stored audio and video

o 5.3 Making the best out of best effort service

o 5.4 Protocols for real-time interactive applications

    o RTP,RTCP,SIP

MM Networking Applications

Fundamental characteristics

- typically delay sensitive
  - end-to-end delay
  - delay jitter
- loss tolerant: infrequent losses cause minor glitches
- antithesis of data, which are loss intolerant but delay tolerant.

Classes of MM applications

- 1) stored streaming
- 2) live streaming
- 3) interactive, real-time

Jitter is the variability of packet delays within the same packet stream

Streaming Stored Multimedia

Stored streaming:

o media stored at source

o transmitted to client

o *streaming:* client playout begins *before* all data has arrived

  o timing constraint for still-to-be transmitted data: in time for playout

Streaming Stored Multimedia: What is it?

Cumulative data (y-axis) vs time (x-axis)

1. video recorded
2. video sent
3. video received, played out at client

network delay

streaming: at this time, client playing out early part of video, while server still sending later part of video

Streaming *Stored* Multimedia: Interactivity

❑ *VCR-like functionality:* client can pause, rewind, FF, push slider bar

  ❍ 10 sec initial delay OK

  ❍ 1-2 sec until command effect OK

❑ timing constraint for still-to-be transmitted data: in time for playout

Streaming *Live* Multimedia

o Examples:

  o Internet radio talk show

  o live sporting event

o Streaming (as with streaming stored multimedia)

  o playback buffer

  o playback can lag tens of seconds after transmission

  o still have timing constraint

o Interactivity

  o fast forward impossible

  o rewind, pause possible!

Real-Time Interactive Multimedia

o applications: IP telephony, video conference, distributed interactive worlds

o end-end delay requirements:

  o audio: < 150 msec good, < 400 msec OK

    • includes application-level (packetization) and network delays

    • higher delays noticeable, impair interactivity

o session initialization

  o how does callee advertise its IP address, port number, encoding algorithms?

Multimedia Over Today's Internet

TCP/UDP/IP: "best-effort service"

o *no* guarantees on delay, loss

But you said multimedia apps requires QoS and level of performance to be effective!

Today's Internet multimedia applications use application-level techniques to mitigate (as best possible) effects of delay, loss

How should the Internet evolve to better support multimedia?

Integrated services philosophy:

○ fundamental changes in Internet so that apps can reserve end-to-end bandwidth

○ requires new, complex software in hosts & routers

Laissez-faire

○ no major changes

○ more bandwidth when needed

○ content distribution, application-layer multicast
  ○ application layer

Differentiated services philosophy:

○ fewer changes to Internet infrastructure, yet provide 1st and 2nd class service

What's your opinion?

A few words about audio compression

- analog signal sampled at constant rate
  - telephone: 8,000 samples/sec
  - CD music: 44,100 samples/sec
- each sample quantized, i.e., rounded
  - e.g., 2^8=256 possible quantized values
- each quantized value represented by bits
  - 8 bits for 256 values

- example: 8,000 samples/sec, 256 quantized values --> 64,000 bps
- receiver converts bits back to analog signal:
  - some quality reduction

Example rates

- CD: 1.411 Mbps
- MP3: 96, 128, 160 kbps
- Internet telephony: 5.3 kbps and up

A few words about video compression

o video: sequence of images displayed at constant rate
   o e.g. 24 images/sec
o digital image: array of pixels
   o each pixel represented by bits
o redundancy
   o spatial (within image)
   o temporal (from one image to next)

Examples:

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

Research:

o layered (scalable) video
   o adapt layers to available bandwidth

Chapter 5 outline

o 5.1 Multimedia networking applications

o 5.2 Streaming stored audio and video

o 5.3 Making the best out of best effort service

o 5.4 Protocols for real-time interactive applications

  o RTP,RTCP,SIP

Streaming Stored Multimedia

application-level streaming techniques for making the best out of best effort service:

- o client-side buffering
- o use of UDP versus TCP
- o multiple encodings of multimedia

Media Player

- o jitter removal
- o decompression
- o error concealment
- o graphical user interface w/ controls for interactivity

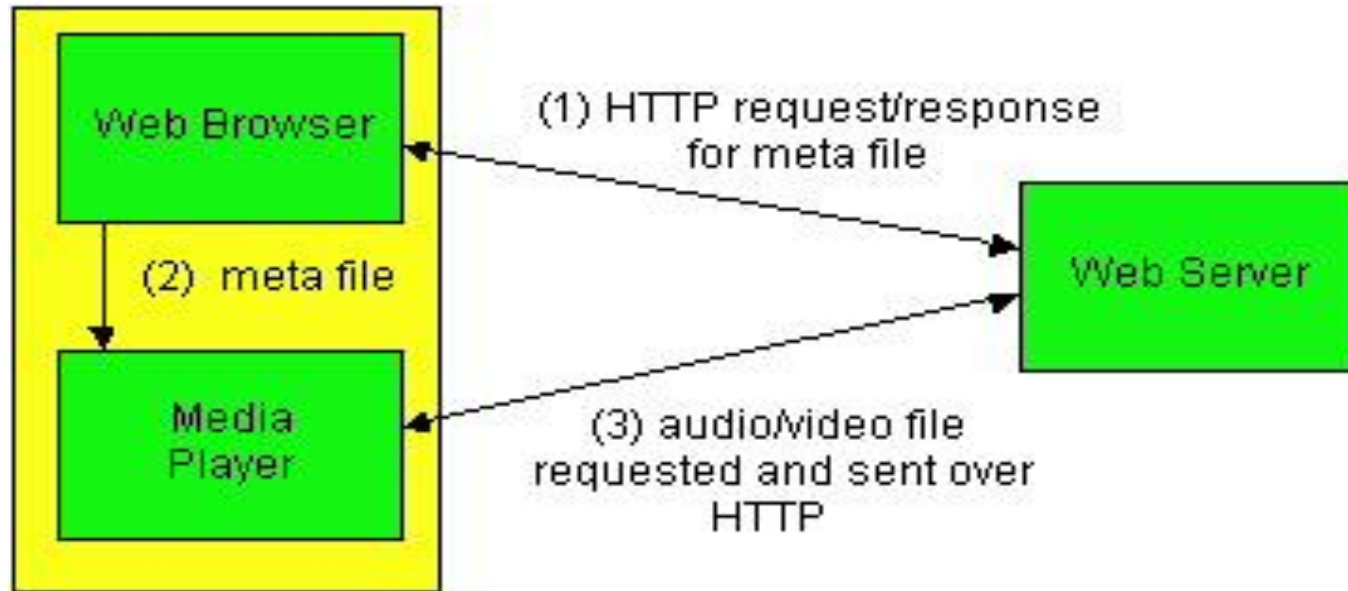# Internet multimedia: simplest approach

client

o audio or video stored in file

o files transferred as HTTP object

   o received in entirety at client

audio, video not streamed:

o no, "pipelining," long delays until playout!

Internet multimedia: streaming approach

❐ browser GETs **metafile**
❐ browser launches player, passing metafile
❐ player contacts server
❐ server streams audio/video to player

Streaming from a streaming server

(1) HTTP request/response for presentation descripton file

(2) presentation description file

(3) audio/video file requested and sent

Web Browser

Web Server

Media Player

Streaming Server
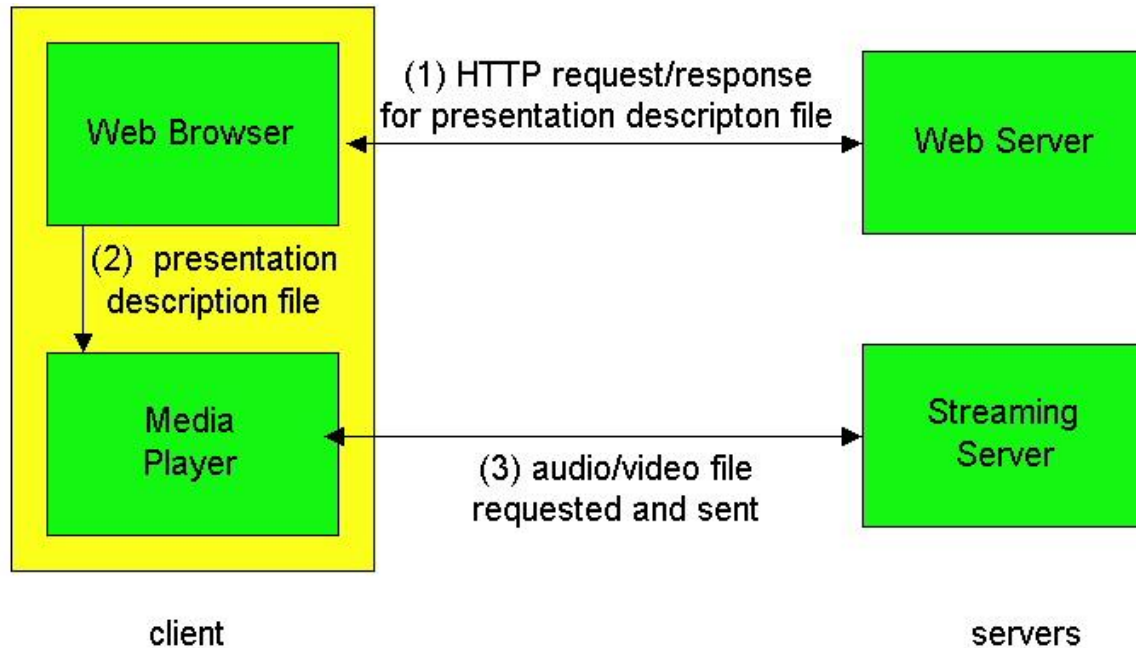
client

servers

- ○ allows for non-HTTP protocol between server, media player
- ○ UDP or TCP for step (3), more shortly
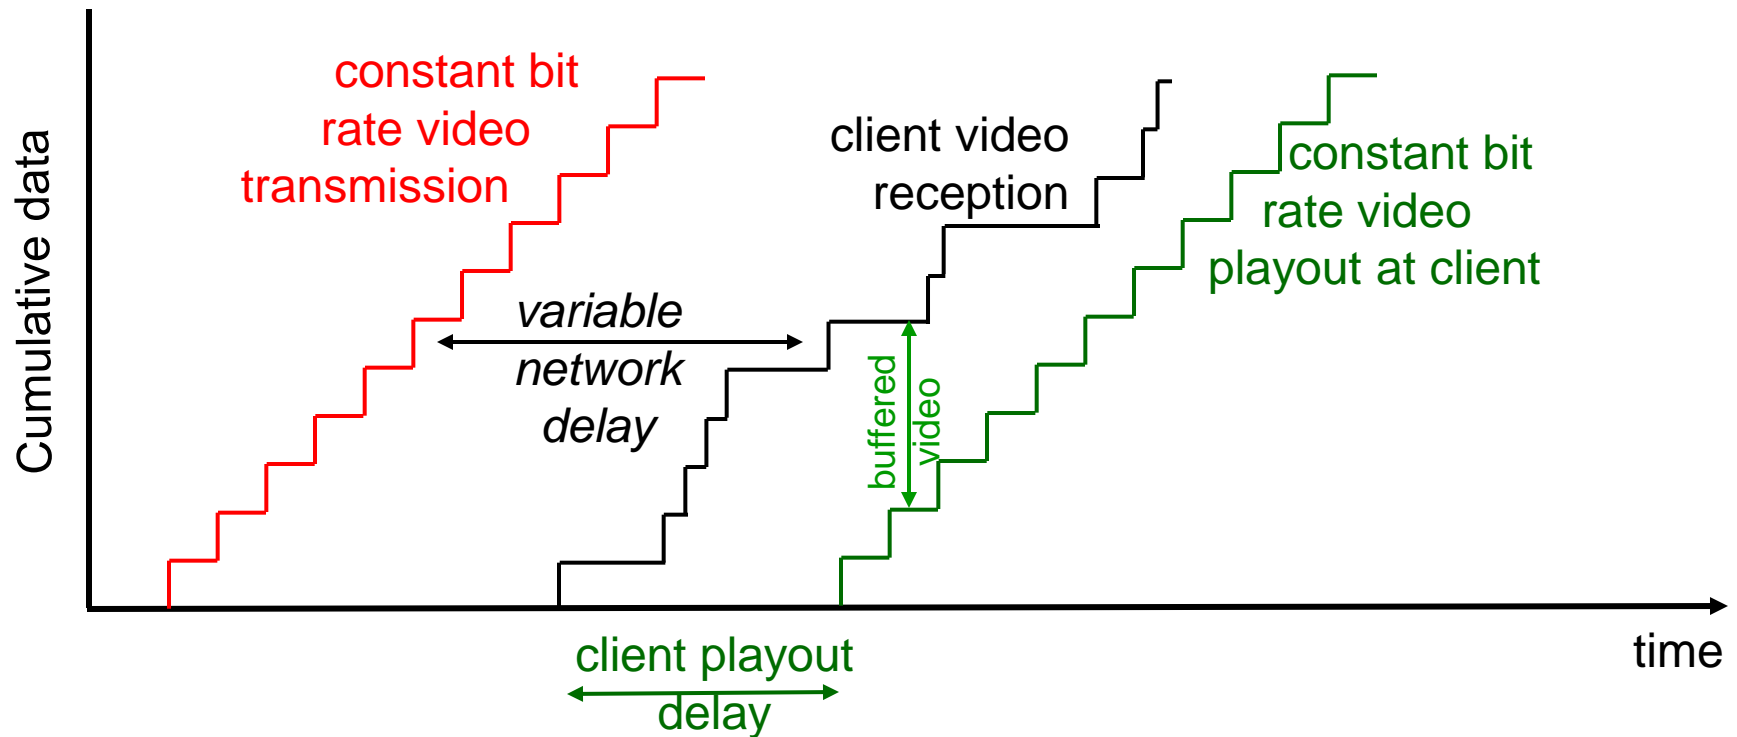
Streaming Multimedia: Client Buffering

- client-side buffering, playout delay compensate for network-added delay, delay jitter

Streaming Multimedia: Client Buffering

constant drain rate, d

variable fill rate, x(t)

from network

client buffer

buffered video

to decompression and playout

○ client-side buffering, playout delay compensate for network-added delay, delay jitter

Streaming Multimedia: UDP or TCP?

o UDP

- o server sends at rate appropriate for client (oblivious to network congestion !)
  - often send rate = encoding rate = constant rate
  - then, fill rate = constant rate - packet loss
- o short playout delay (2-5 seconds) to remove network jitter
- o error recover: time permitting

Streaming Multimedia: UDP or TCP?

o TCP

  o send at maximum possible rate under TCP

  o fill rate fluctuates due to TCP congestion control

  o larger playout delay: smooth TCP delivery rate

  o HTTP/TCP passes more easily through firewalls

Streaming Multimedia: client rate(s)

1.5 Mbps encoding

28.8 Kbps encoding

Q: how to handle different client receive rate capabilities?

- 28.8 Kbps dialup
- 100 Mbps Ethernet

A: server stores, transmits multiple copies of video, encoded at different rates

User Control of Streaming Media: RTSP

HTTP

o does not target multimedia content

o no commands for fast forward, etc.

RTSP: RFC 2326

o client-server application layer protocol

o user control: rewind, fast forward, pause, resume, repositioning, etc…

What it doesn't do:

o doesn't define how audio/video is encapsulated for streaming over network

o doesn't restrict how streamed media is transported (UDP or TCP possible)

o doesn't specify how media player buffers audio/video

RTSP: out of band control

FTP uses an "out-of-band" control channel:

- file transferred over one TCP connection.
- control info (directory changes, file deletion, rename) sent over separate TCP connection
- "out-of-band", "in-band" channels use different port numbers

RTSP messages also sent out-of-band:

- RTSP control messages use different port numbers than media stream: out-of-band.
  - port 554
- media stream is considered "in-band".

RTSP Example

Scenario:

- metafile communicated to web browser

- browser launches player

- player sets up an RTSP control connection, data connection to streaming server

Metafile Example

```
<title>Twister</title>
<session>
        <group language=en lipsync>
                <switch>
                    <track type=audio
                            e="PCMU/8000/1"
                            src = "rtsp://audio.example.com/twister/audio.en/lofi">
                    <track type=audio
                            e="DVI4/16000/2" pt="90 DVI4/8000/1"
                            src="rtsp://audio.example.com/twister/audio.en/hifi">
                </switch>
            <track type="video/jpeg"
                            src="rtsp://video.example.com/twister/video">
        </group>
</session>
```

RTSP Operation

RTSP Exchange Example

C: SETUP rtsp://audio.example.com/twister/audio RTSP/1.0
   Transport: rtp/udp; compression; port=3056; mode=PLAY

S: RTSP/1.0 200 1 OK
    Session 4231

C: PLAY rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
    Session: 4231
    Range: npt=0-

C: PAUSE rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
    Session: 4231
    Range: npt=37

C: TEARDOWN rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
    Session: 4231

S: 200 3 OK

Chapter 5 outline

o 5.1 Multimedia networking applications

o 5.2 Streaming stored audio and video

o 5.3 Making the best out of best effort service

o 5.4 Protocols for real-time interactive applications
  - o RTP,RTCP,SIP

Real-time interactive applications

- PC-2-PC phone
  - Skype
- PC-2-phone
  - Dialpad
  - Net2phone
  - Skype
- videoconference with webcams
  - Skype
  - Polycom

Going to now look at a PC-2-PC Internet phone example in detail

Interactive Multimedia: Internet Phone

- o *Introduce Internet Phone by way of an example*
- o speaker's audio: alternating talk spurts, silent periods.
  - o 64 kbps during talk spurt
  - o pkts generated only during talk spurts
  - o 20 msec chunks at 8 Kbytes/sec: 160 bytes data
- o application-layer header added to each chunk.
- o chunk+header encapsulated into UDP segment.
- o application sends UDP segment into socket every 20 msec during talkspurt

Internet Phone: Packet Loss and Delay

o network loss: IP datagram lost due to network congestion (router buffer overflow)

o delay loss: IP datagram arrives too late for playout at receiver
  o delays: processing, queueing in network; end-system (sender, receiver) delays
  o typical maximum tolerable delay: 400 ms

o loss tolerance: depending on voice encoding, losses concealed, packet loss rates between 1% and 10% can be tolerated.

Delay Jitter
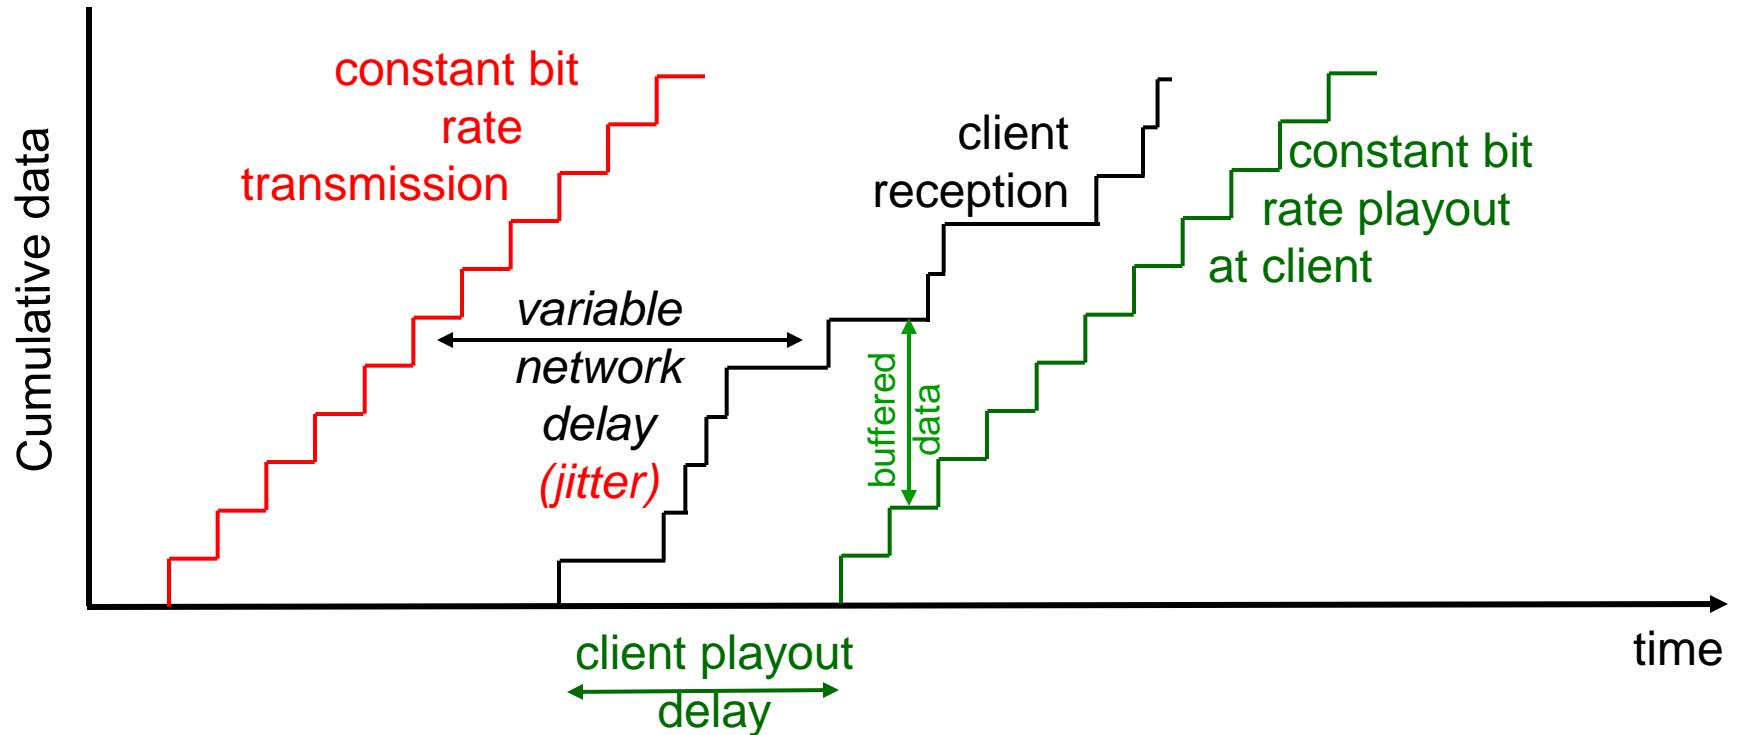
o consider end-to-end delays of two consecutive packets: difference can be more or less than 20 msec (transmission time difference)

Internet Phone: Fixed Playout Delay

o receiver attempts to playout each chunk exactly q msecs after chunk was generated.

- o chunk has time stamp t: play out chunk at t+q .
- o chunk arrives after t+q: data arrives too late for playout, data "lost"

o tradeoff in choosing q:

- o *large q:* less packet loss
- o *small q:* better interactive experience

Fixed Playout Delay

- sender generates packets every 20 msec during talk spurt.
- first packet received at time r
- first playout schedule: begins at p
- second playout schedule: begins at p'

packets

packets generated

packets received

loss

playout schedule p' - r

playout schedule p - r

time

r

p

p'

Adaptive Playout Delay (1)

o <u>Goal:</u> minimize playout delay, keeping late loss rate low

o <u>Approach:</u> adaptive playout delay adjustment:

  o estimate network delay, adjust playout delay at beginning of each talk spurt.

  o silent periods compressed and elongated.

  o chunks still played out every 20 msec during talk spurt.

t_i = timestamp of the ith packet

r_i = the time packet i is received by receiver

p_i = the time packet i is played at receiver

$r_i - t_i$ = network delay for ith packet

d_i = estimate of average network delay after receiving ith packet

dynamic estimate of average delay at receiver:

$$d_i = (1 - u)d_{i-1} + u(r_i - t_i)$$

where u is a fixed constant (e.g., u = .01).

Adaptive playout delay (2)

❑ also useful to estimate average deviation of delay, v_i:

$$v_i = (1 - u)v_{i-1} + u \mid r_i - t_i - d_i \mid$$

❑ estimates d_i, v_i calculated for every received packet
  (but used only at start of talk spurt

❑ for first packet in talk spurt, playout time is:

$$p_i = t_i + d_i + Kv_i$$

   where K is positive constant

❑ remaining packets in talkspurt are played out periodically

# Adaptive Playout (3)

Q: How does receiver determine whether packet is first in a talkspurt?

o if no loss, receiver looks at successive timestamps.

  o difference of successive stamps > 20 msec -->talk spurt begins.

o with loss possible, receiver must look at both time stamps and sequence numbers.

  o difference of successive stamps > 20 msec and sequence numbers without gaps --> talk spurt begins.

Content distribution networks (CDNs)

Content replication

o challenging to stream large files (e.g., video) from single origin server in real time

o *solution:* replicate content at hundreds of servers throughout Internet

   o content downloaded to CDN servers ahead of time

   o *placing content "close" to user avoids impairments (loss, delay) of sending content over long paths*

   o CDN server typically in edge/access network

origin server
in North America

CDN distribution node

CDN server
in S. America

CDN server
in Europe
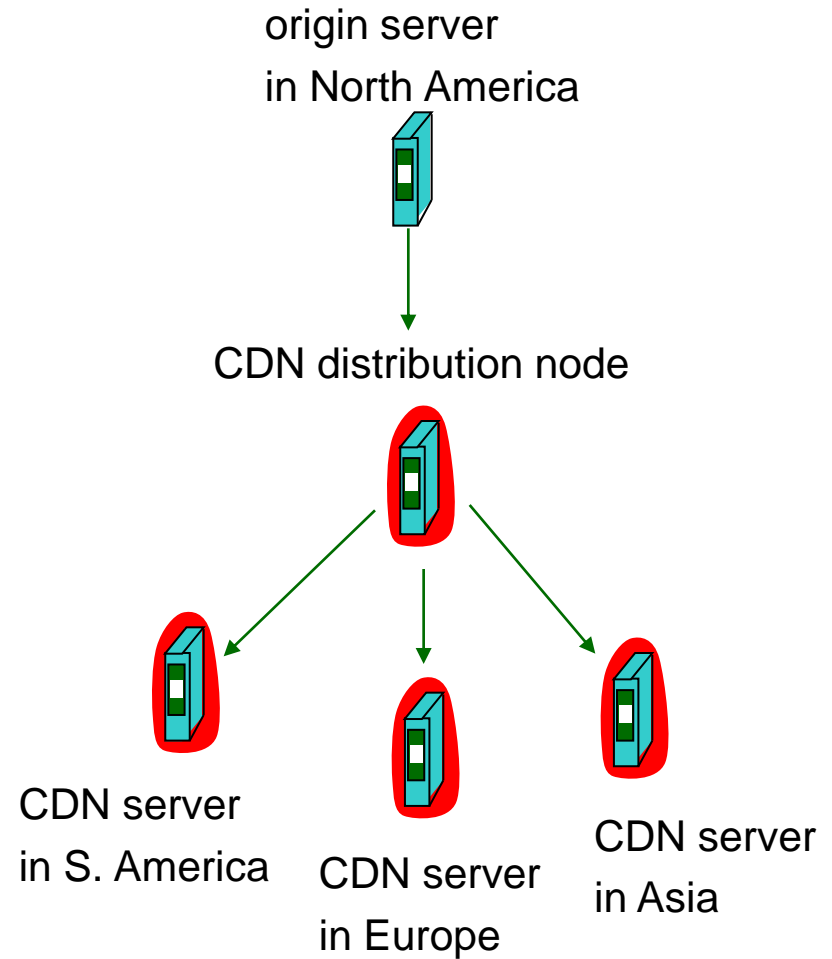
CDN server
in Asia

Content distribution networks (CDNs)

Content replication

- CDN (e.g., Akamai) customer is the content provider (e.g., CNN)
- CDN replicates customers' content in CDN servers.
- when provider updates content, CDN updates servers
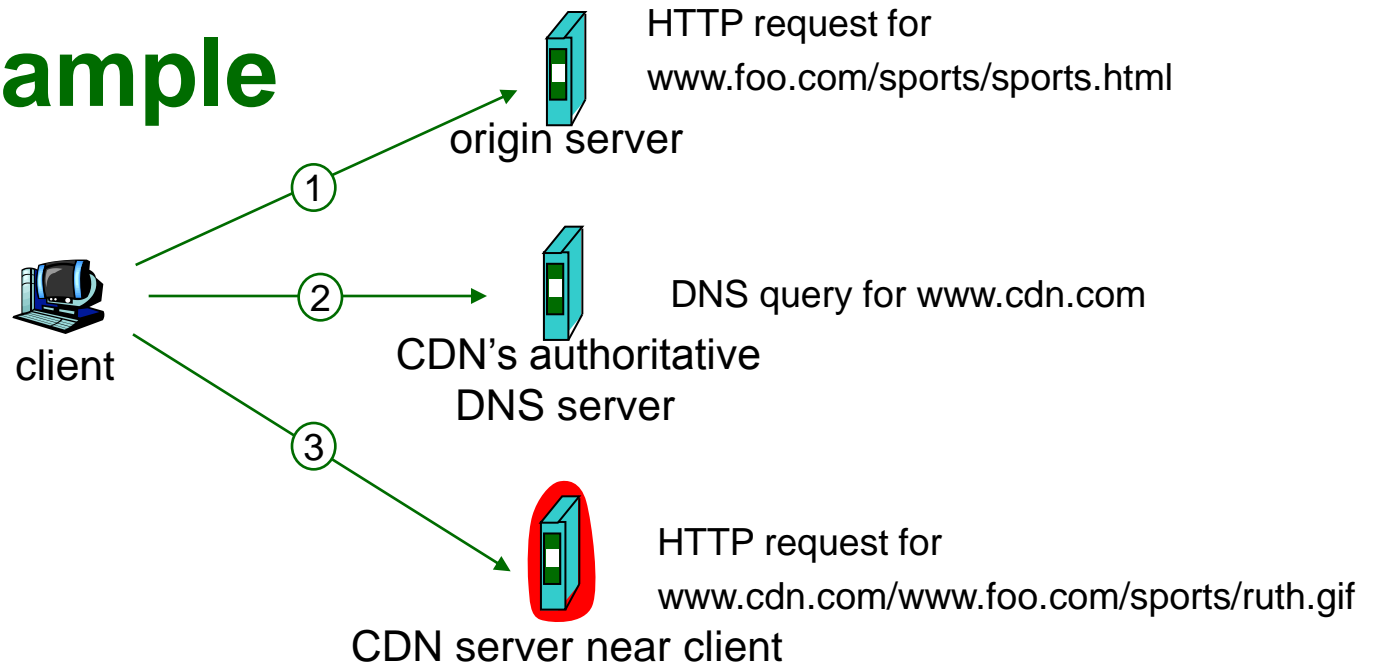
origin server
in North America

CDN distribution node

CDN server
in S. America

CDN server
in Europe

CDN server
in Asia

CDN example

HTTP request for
www.foo.com/sports/sports.html
origin server

①

② DNS query for www.cdn.com
CDN's authoritative
DNS server

③

HTTP request for
www.cdn.com/www.foo.com/sports/ruth.gif
CDN server near client

client
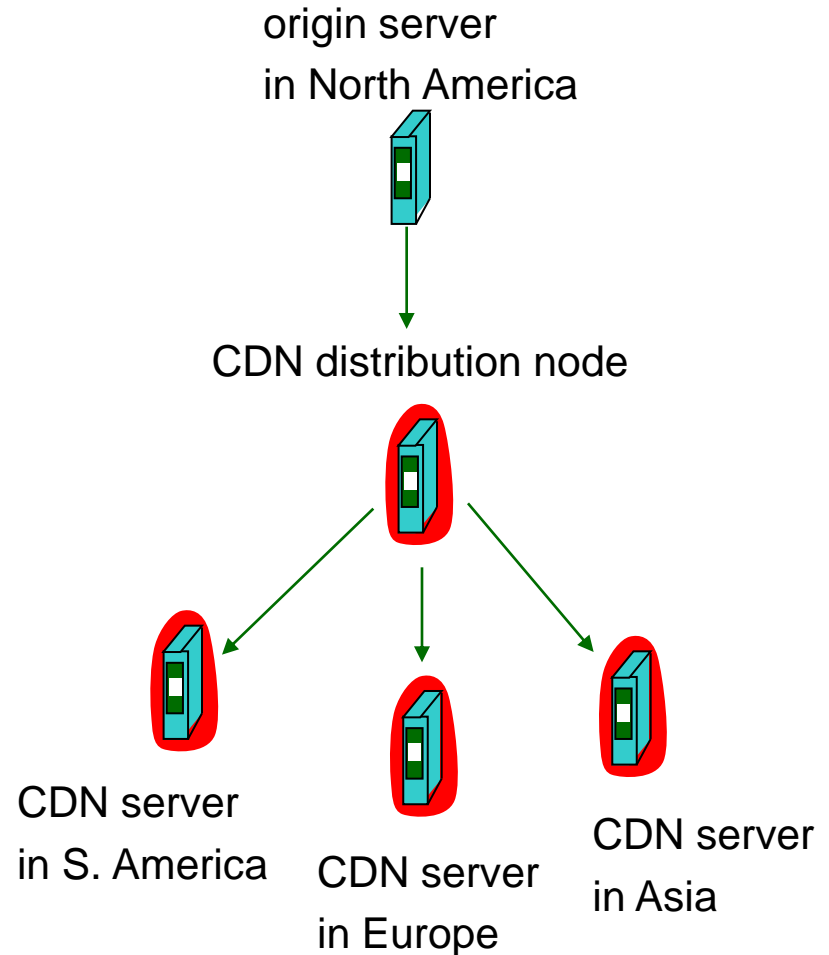
origin server (www.foo.com)

○ distributes HTML

○ replaces:

   http://www.foo.com/sports.ruth.gif

  with

  http://www.cdn.com/www.foo.com/sports/ruth.gif

CDN company (cdn.com)

❐ distributes gif files

❐ uses its authoritative DNS server to route redirect requests

More about CDNs

<u>routing requests</u>

- CDN creates a "map", indicating distances from leaf ISPs and CDN nodes

- when query arrives at authoritative DNS server:
  - server determines ISP from which query originates
  - uses "map" to determine best CDN server

- CDN nodes create application-layer overlay network

Summary: Internet Multimedia: bag of tricks

o use UDP to avoid TCP congestion control (delays) for time-sensitive traffic

o client-side adaptive playout delay: to compensate for delay

o server side matches stream bandwidth to available client-to-server path bandwidth
  - o chose among pre-encoded stream rates
  - o dynamic server encoding rate

o error recovery (on top of UDP)
  - o FEC, interleaving, error concealment
  - o retransmissions, time permitting

o CDN: bring content closer to clients

Chapter 5 outline

- 5.1 Multimedia networking applications
- 5.2 Streaming stored audio and video
- 5.3 Making the best out of best effort service
- 5.4 Protocols for real-time interactive applications
 - RTP,RTCP,SIP

Real-Time Protocol (RTP)

o RTP specifies packet structure for packets carrying audio, video data

o RFC 3550

o RTP packet provides
  o payload type identification
  o packet sequence numbering
  o time stamping

o RTP runs in end systems

o RTP packets encapsulated in UDP segments

o interoperability: if two Internet phone applications run RTP, then they may be able to work together

RTP runs on top of UDP

RTP libraries provide transport-layer interface that extends UDP:

- port numbers, IP addresses
- payload type identification
- packet sequence numbering
- time-stamping

transport layer

Application
RTP
UDP
IP
Data Link
Physical

RTP Example

o consider sending 64 kbps PCM-encoded voice over RTP.

o application collects encoded data in chunks, e.g., every 20 msec = 160 bytes in a chunk.

o audio chunk + RTP header form RTP packet, which is encapsulated in UDP segment

o RTP header indicates type of audio encoding in each packet

o sender can change encoding during conference.

o RTP header also contains sequence numbers, timestamps.

RTP and QoS

○ RTP does **not** provide any mechanism to ensure timely data delivery or other QoS guarantees.

○ RTP encapsulation is only seen at end systems (not) by intermediate routers.

  ○ routers providing best-effort service, making no special effort to ensure that RTP packets arrive at destination in timely matter.

RTP Header

Payload Type	Sequence Number	Timestamp	Synchronization Source Identifer	Miscellaneous Fields

RTP Header

Payload Type (7 bits): Indicates type of encoding currently being used. If sender changes encoding in middle of conference, sender informs receiver via payload type field.

- Payload type 0: PCM mu-law, 64 kbps
- Payload type 3, GSM, 13 kbps
- Payload type 7, LPC, 2.4 kbps
- Payload type 26, Motion JPEG
- Payload type 31. H.261
- Payload type 33, MPEG2 video

Sequence Number (16 bits): Increments by one for each RTP packet sent, and may be used to detect packet loss and to restore packet sequence.

RTP Header (2)

o *Timestamp field (32 bytes long):* sampling instant of first byte in this RTP data packet

  - o for audio, timestamp clock typically increments by one for each sampling period (for example, each 125 usecs for 8 KHz sampling clock)

  - o if application generates chunks of 160 encoded samples, then timestamp increases by 160 for each RTP packet when source is active. Timestamp clock continues to increase at constant rate when source is inactive.

o *SSRC field (32 bits long):* identifies source of t RTP stream. Each stream in RTP session should have distinct SSRC.

Real-Time Control Protocol (RTCP)

○ works in conjunction with RTP.

○ each participant in RTP session periodically transmits RTCP control packets to all other participants.

○ each RTCP packet contains sender and/or receiver reports

   ○ report statistics useful to application: # packets sent, # packets lost, interarrival jitter, etc.

○ feedback can be used to control performance

   ○ sender may modify its transmissions based on feedback

RTCP - Continued

❑ each RTP session: typically a single multicast address; all RTP /RTCP packets belonging to session use multicast address.

❑ RTP, RTCP packets distinguished from each other via distinct port numbers.

❑ to limit traffic, each participant reduces RTCP traffic as number of conference participants increases

RTCP Packets

Receiver report packets:

○ fraction of packets lost, last sequence number, average interarrival jitter

Sender report packets:

○ SSRC of RTP stream, current time, number of packets sent, number of bytes sent

Source description packets:

○ e-mail address of sender, sender's name, SSRC of associated RTP stream

○ provide mapping between the SSRC and the user/host name

Synchronization of Streams

○ RTCP can synchronize different media streams within a RTP session

○ consider videoconferencing app for which each sender generates one RTP stream for video, one for audio.

○ timestamps in RTP packets tied to the video, audio sampling clocks

   ○ *not* tied to wall-clock time

○ each RTCP sender-report packet contains (for most recently generated packet in associated RTP stream):

   ○ timestamp of RTP packet

   ○ wall-clock time for when packet was created.

○ receivers uses association to synchronize playout of audio, video

RTCP Bandwidth Scaling

- RTCP attempts to limit its traffic to 5% of session bandwidth.

Example

- Suppose one sender, sending video at 2 Mbps. Then RTCP attempts to limit its traffic to 100 Kbps.

- RTCP gives 75% of rate to receivers; remaining 25% to sender

- 75 kbps is equally shared among receivers:
  - with R receivers, each receiver gets to send RTCP traffic at 75/R kbps.

- sender gets to send RTCP traffic at 25 kbps.

- participant determines RTCP packet transmission period by calculating avg RTCP packet size (across entire session) and dividing by allocated rate

SIP: Session Initiation Protocol
## [RFC 3261]

SIP long-term vision:

- all telephone calls, video conference calls take place over Internet

- people are identified by names or e-mail addresses, rather than by phone numbers

- you can reach callee, no matter where callee roams, no matter what IP device callee is currently using

SIP Services

○ Setting up a call, SIP provides mechanisms ..

  ○ for caller to let callee know she wants to establish a call

  ○ so caller, callee can agree on media type, encoding

  ○ to end call

○ determine current IP address of callee:

  ○ maps mnemonic identifier to current IP address

○ call management:

  ○ add new media streams during call

  ○ change encoding during call

 ○ invite others

  ○ transfer, hold calls

Setting up a call to known IP address

Alice

167.180.112.24

Bob

193.64.210.89

INVITE bob@193.64.210.89
c=IN IP4 167.180.112.24
m=audio 38060 RTP/AVP 0

port 5060

Bob's terminal rings

200 OK
c=IN IP4 193.64.210.89
m=audio 48753 RTP/AVP 3

port 5060

ACK

port 5060

μ Law audio

port 38060

GSM

port 48753

time

time

❑ Alice's SIP invite message indicates her port number, IP address, encoding she prefers to receive (PCM ulaw)

❑ Bob's 200 OK message indicates his port number, IP address, preferred encoding (GSM)

❑ SIP messages can be sent over TCP or UDP; here sent over RTP/UDP.

❑ default SIP port number is 5060.

Multimedia Networking   7-58

Setting up a call (more)

o codec negotiation:

- o suppose Bob doesn't have PCM ulaw encoder.
- o Bob will instead reply with 606 Not Acceptable Reply, listing his encoders Alice can then send new INVITE message, advertising different encoder

o rejecting a call

- o Bob can reject with replies "busy," "gone," "payment required," "forbidden"

o media can be sent over RTP or some other protocol

Example of SIP message

```
INVITE sip:bob@domain.com SIP/2.0
Via: SIP/2.0/UDP 167.180.112.24
From: sip:alice@hereway.com
To: sip:bob@domain.com
Call-ID: a2e3a@pigeon.hereway.com
Content-Type: application/sdp
Content-Length: 885


c=IN IP4 167.180.112.24
m=audio 38060 RTP/AVP 0
```

❑ Here we don't know Bob's IP address. Intermediate SIP servers needed.

❑ Alice sends, receives SIP messages using SIP default port 506

❑ Alice specifies in Via: header that SIP client sends, receives SIP messages over UDP

Notes:

o HTTP message syntax

o sdp = session description protocol

o Call-ID is unique for every call.

Name translation and user locataion

○ caller wants to call callee, but only has callee's name or e-mail address.

○ need to get IP address of callee's current host:

  ○ user moves around
  ○ DHCP protocol
  ○ user has different IP devices (PC, PDA, car device)

○ result can be based on:

  ○ time of day (work, home)
  ○ caller (don't want boss to call you at home)
  ○ status of callee (calls sent to voicemail when callee is already talking to someone)

Service provided by SIP servers:

○ SIP registrar server

○ SIP proxy server

SIP Registrar

○ when Bob starts SIP client, client sends SIP REGISTER message to Bob's registrar server

(similar function needed by Instant Messaging)

Register Message:

```
REGISTER sip:domain.com SIP/2.0
Via: SIP/2.0/UDP 193.64.210.89
From: sip:bob@domain.com
To: sip:bob@domain.com
Expires: 3600
```

SIP Proxy

o Alice sends invite message to her proxy server

    o contains address sip:bob@domain.com

o proxy responsible for routing SIP messages to callee

    o possibly through multiple proxies.

o callee sends response back through the same set of proxies.

o proxy returns SIP response message to Alice

    o contains Bob's IP address

o proxy analogous to local DNS server

Example

Caller jim@umass.edu
with places a
call to keith@upenn.edu

(1) Jim sends INVITE
message to umass SIP
proxy. (2) Proxy forwards
request to upenn
registrar server.
(3) upenn server returns
redirect response,
indicating that it should
try keith@eurecom.fr

(4) umass proxy sends INVITE to eurecom registrar. (5) eurecom registrar
forwards INVITE to 197.87.54.21, which is running keith's SIP client. (6-8)
SIP response sent back (9) media sent directly
between clients.
Note: also a SIP ack message, which is not shown.

Thank you

Any questions?