

Chapter 7

Multimedia Networking


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 see the animations; and 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) that you mention their source (after all, we'd like people to use our book!)
- ❖ If you post any slides 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 1996-2012
J.F Kurose and K.W. Ross, All Rights Reserved



**Computer
Networking: A Top
Down Approach**
6th edition
Jim Kurose, Keith Ross
Addison-Wesley
March 2012

Multimedia networking: outline

7.1 multimedia networking applications

7.2 streaming *stored* video

~~7.3 voice over IP~~

7.4 protocols for *real-time* conversational applications

~~7.5 network support for multimedia~~

Multimedia networking: outline

7.1 multimedia networking applications

7.2 streaming *stored* video


~~7.3 voice over IP~~

7.4 protocols for *real-time* conversational applications

~~7.5 network support for multimedia~~

Multimedia: audio

- ❖ analog audio 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, e.g., 8 bits for 256 values




Multimedia: audio

- ❖ 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



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 I (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 $i+1$

Multimedia networking: 3 application types

- ❖ *streaming, stored* audio, video
 - *streaming*: can begin playout before downloading entire file
 - *stored (at server)*: can transmit faster than audio/video will be rendered (implies storing/buffering at client)
 - e.g., YouTube, Netflix, Hulu
- ~~❖ *conversational* voice/video over IP~~
 - ~~▪ interactive nature of human-to-human conversation limits delay tolerance~~
 - ~~▪ e.g., Skype~~
- ~~❖ *streaming live* audio, video~~
 - ~~▪ e.g., live sporting event (futbol)~~

Multimedia networking: outline

7.1 multimedia networking applications

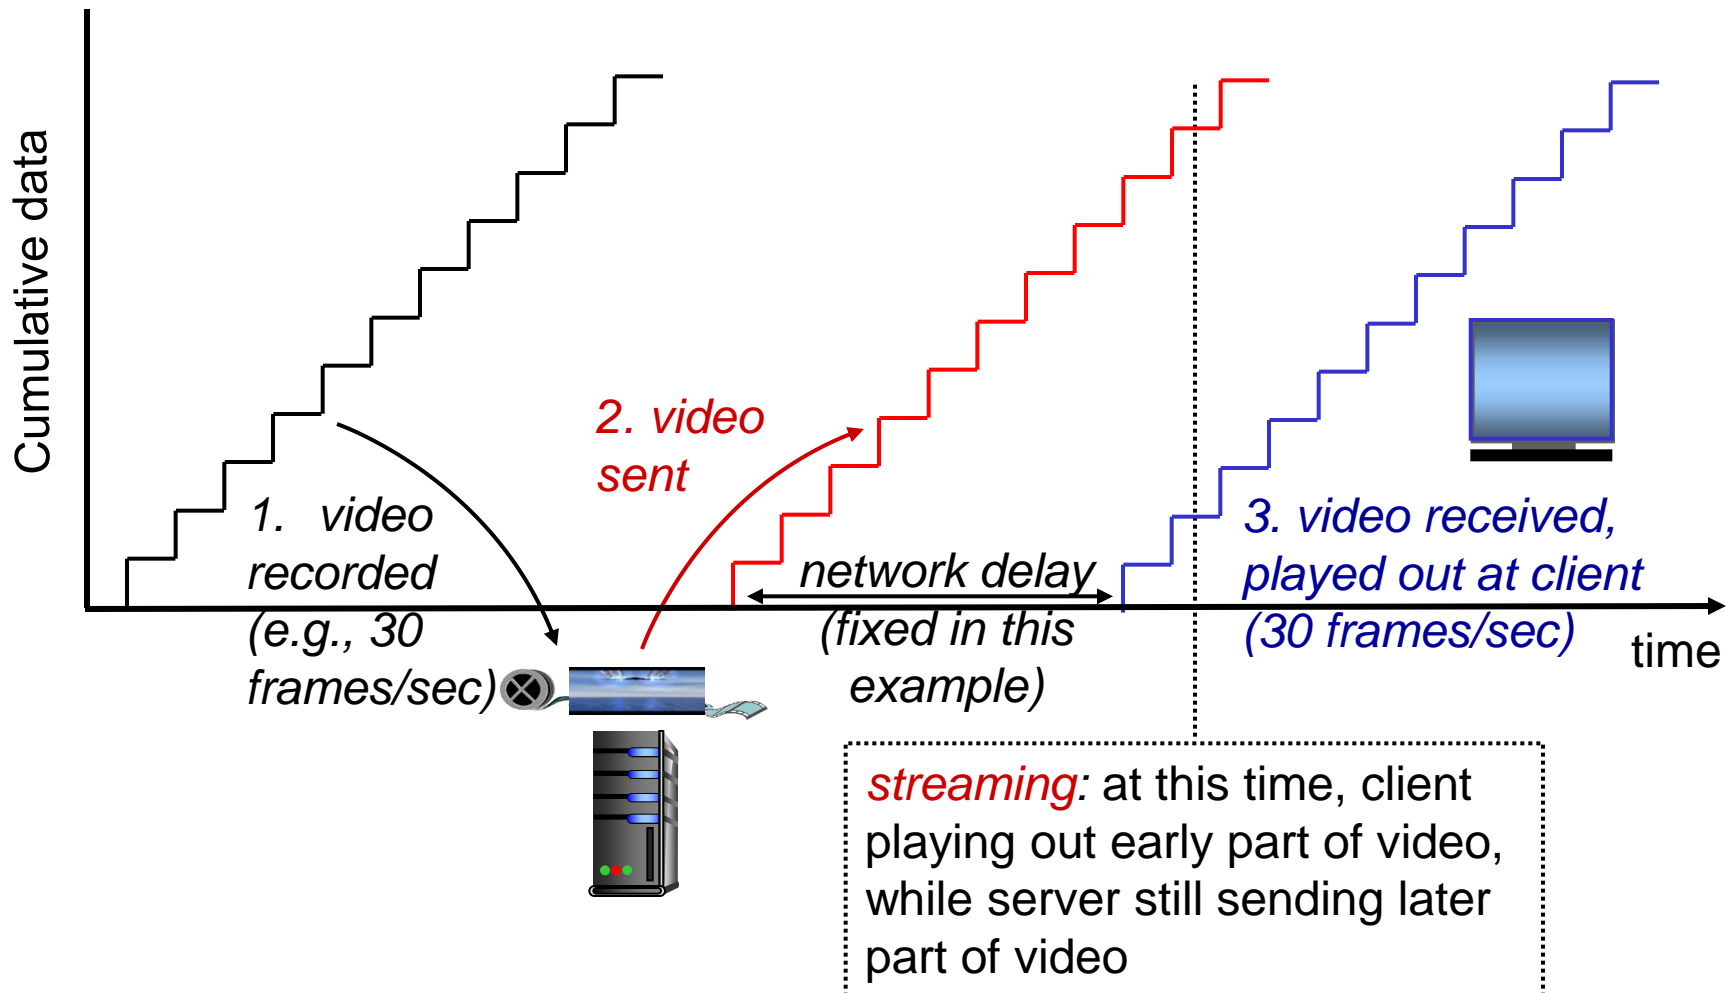
7.2 streaming *stored* video

~~7.3 voice over IP~~

7.4 protocols for *real-time* conversational applications

~~7.5 network support for multimedia~~

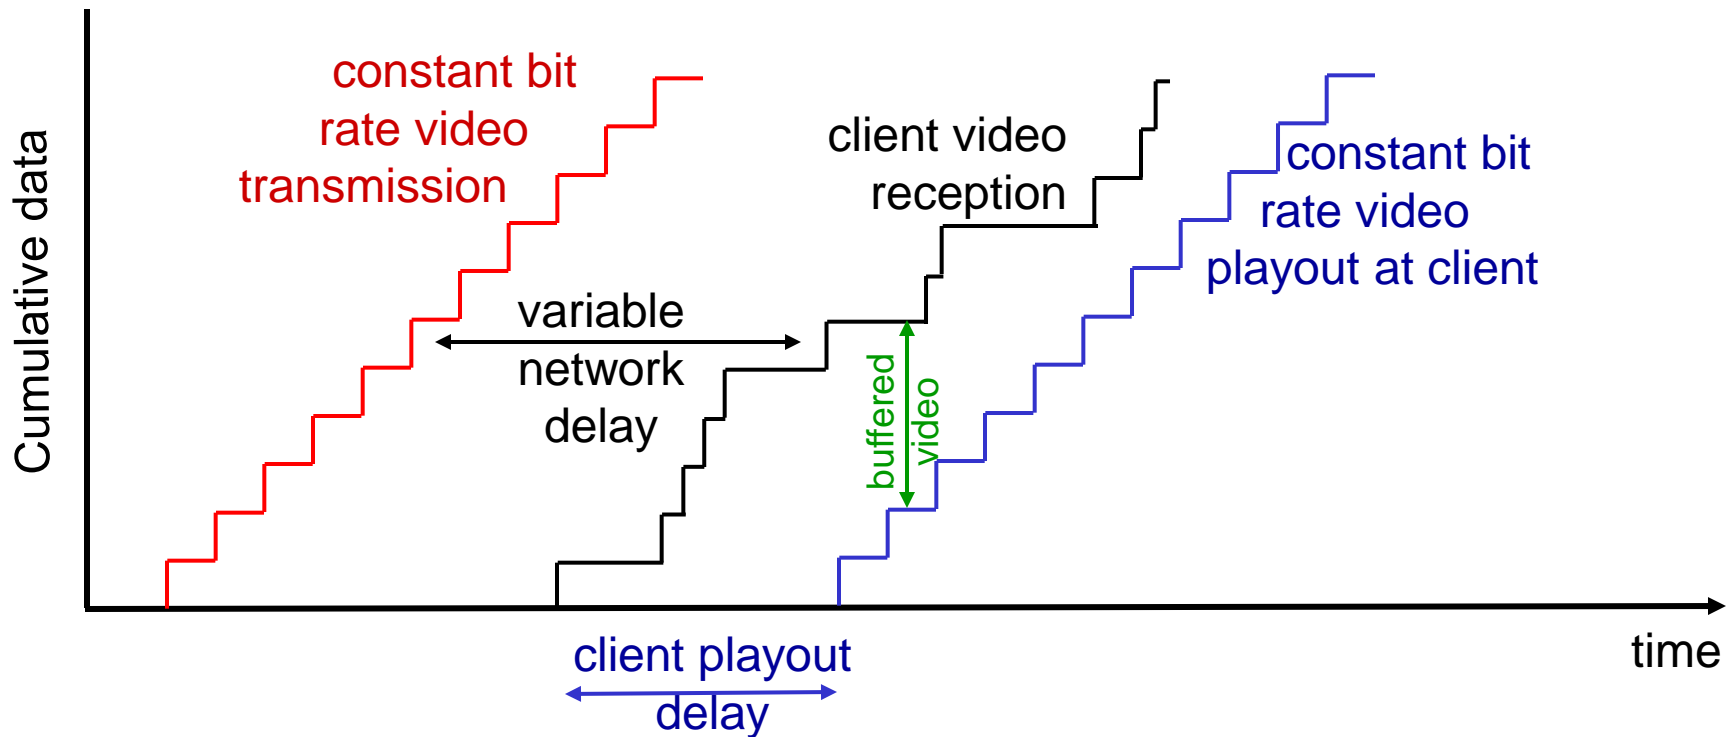
Streaming stored video:



Streaming stored video: challenges

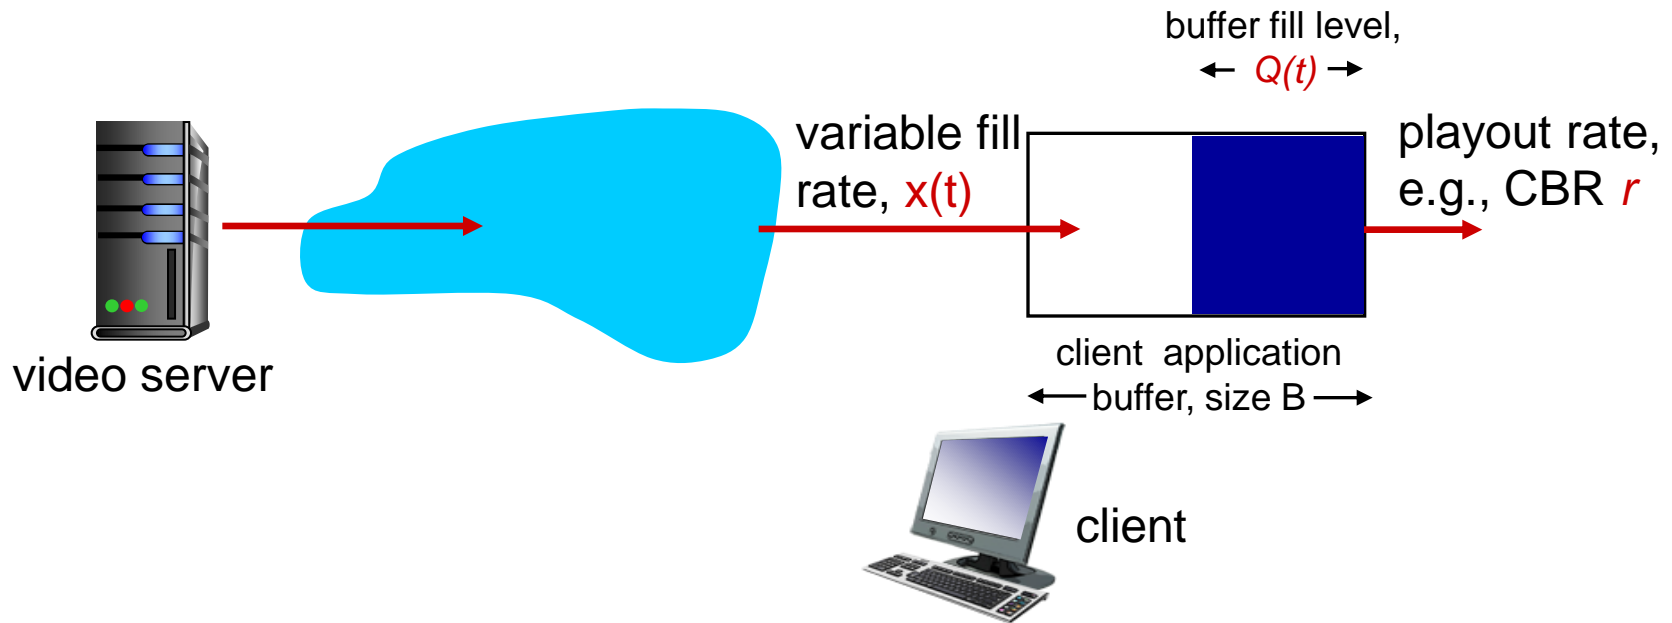
- ❖ *continuous playout constraint*: once client playout begins, playback must match original timing
 - ... but *network delays are variable* (jitter), so will need *client-side buffer* to match playout requirements
- ❖ other challenges:
 - client interactivity: pause, fast-forward, rewind, jump through video
 - video packets may be lost, retransmitted

Streaming stored video: revisited

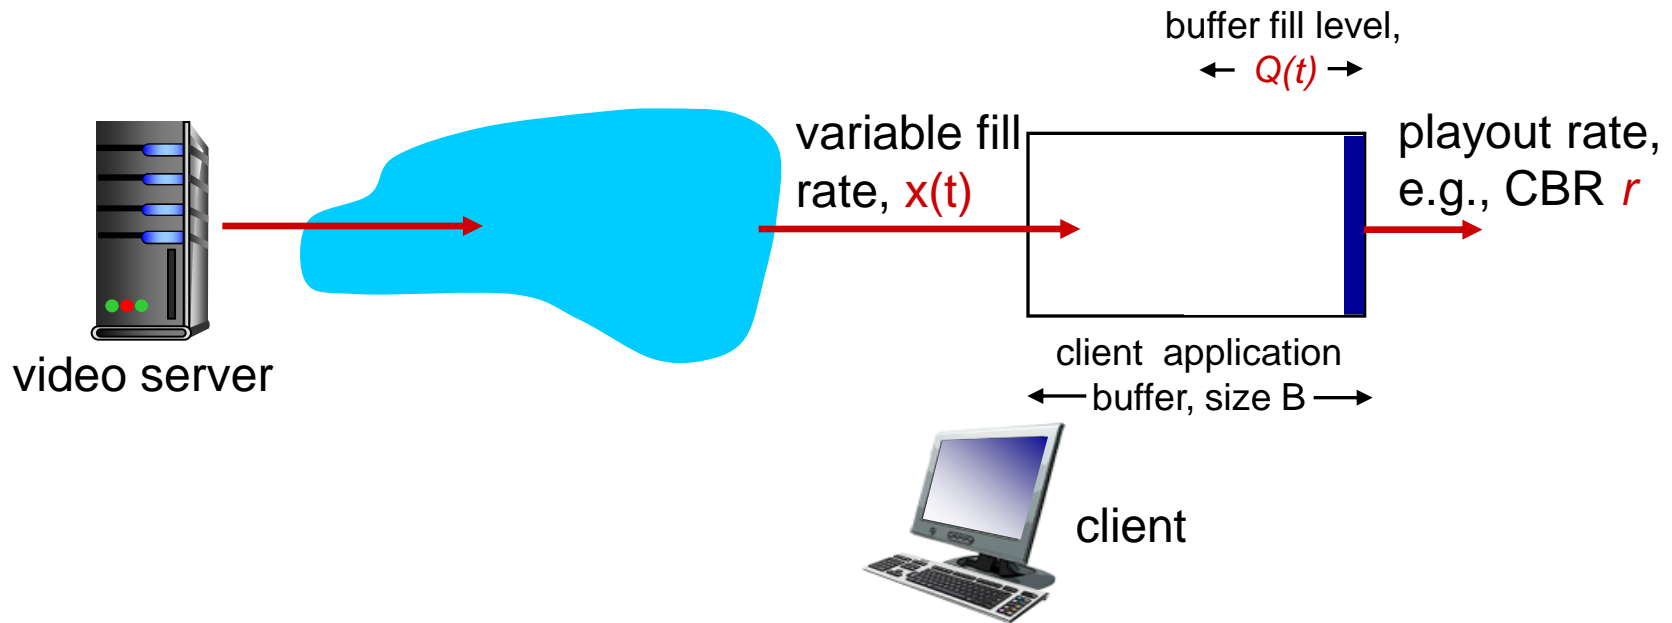


- ❖ *client-side buffering and playout delay*: compensate for network-added delay, delay jitter

Client-side buffering, playout




Client-side buffering, playout



1. Initial fill of buffer until playout begins at t_p
2. playout begins at t_p ,
3. buffer fill level varies over time as fill rate $x(t)$ varies and playout rate r is constant

Client-side buffering, playout



playout buffering: average fill rate (\bar{x}), playout rate (r):


- ❖ $\bar{x} < r$: buffer eventually empties (causing freezing of video playout until buffer again fills)
- ❖ $\bar{x} > r$: buffer will not empty, provided initial playout delay is large enough to absorb variability in $x(t)$
 - *initial playout delay tradeoff*: buffer starvation less likely with larger delay, but larger delay until user begins watching

Streaming multimedia: UDP

- ❖ server sends at rate appropriate for client
 - often: send rate = encoding rate = constant rate
 - transmission rate can be oblivious to congestion levels
- ❖ short playout delay (2-5 seconds) to remove network jitter
- ❖ error recovery: application-level, time permitting
- ❖ RTP [RFC 2326]: multimedia payload types
- ❖ UDP may *not* go through firewalls

Streaming multimedia: HTTP

- ❖ multimedia file retrieved via HTTP GET
- ❖ send at maximum possible rate under TCP



- ❖ fill rate fluctuates due to TCP congestion control, retransmissions (in-order delivery)
- ❖ larger playout delay: smooth TCP delivery rate
- ❖ HTTP/TCP passes more easily through firewalls

Streaming multimedia: DASH

- ❖ *DASH: Dynamic, Adaptive Streaming over HTTP*
- ❖ *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)

Multimedia networking: outline

7.1 multimedia networking applications

7.2 streaming stored video

~~7.3 voice over IP~~

7.4 protocols for real-time conversational applications: RTP, SIP

~~7.5 network support for multimedia~~


Real-Time Protocol (RTP)

- ❖ RTP specifies packet structure for packets carrying audio, video data
- ❖ RFC 3550
- ❖ RTP packet provides
 - payload type identification
 - packet sequence numbering
 - time stamping
- ❖ RTP runs in end systems
- ❖ RTP packets encapsulated in UDP segments
- ❖ interoperability: if two VoIP applications run RTP, 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



RTP example

example: sending 64 kbps PCM-encoded voice over RTP

- ❖ application collects encoded data in chunks, e.g., every 20 msec = 160 bytes in a chunk
- ❖ audio chunk + RTP header form RTP packet, which is encapsulated in UDP segment
- ❖ RTP header indicates type of audio encoding in each packet
 - sender can change encoding during conference
- ❖ 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 only seen at end systems (*not* by intermediate routers)
 - routers provide best-effort service, making no special effort to ensure that RTP packets arrive at destination in timely matter

RTP header

| | | | | |
|---------------------|------------------------|-------------------|----------------------------------|-----------------------------|
| <i>payload type</i> | <i>sequence number</i> | <i>time stamp</i> | <i>Synchronization Source ID</i> | <i>Miscellaneous fields</i> |
|---------------------|------------------------|-------------------|----------------------------------|-----------------------------|

payload type (7 bits): indicates type of encoding currently being used. If sender changes encoding during call, 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 # (16 bits): increment by one for each RTP packet sent

- ❖ detect packet loss, restore packet sequence

RTP header

| | | | | |
|---------------------|------------------------|-------------------|----------------------------------|-----------------------------|
| <i>payload type</i> | <i>sequence number</i> | <i>time stamp</i> | <i>Synchronization Source ID</i> | <i>Miscellaneous fields</i> |
|---------------------|------------------------|-------------------|----------------------------------|-----------------------------|

- ❖ ***timestamp field (32 bits long)***: sampling instant of first byte in this RTP data packet
 - for audio, timestamp clock increments by one for each sampling period (e.g., each 125 usecs for 8 KHz sampling clock)
 - if application generates chunks of 160 encoded samples, timestamp increases by 160 for each RTP packet when source is active. Timestamp clock continues to increase at constant rate when source is inactive.
- ❖ ***SSRC field (32 bits long)***: identifies source of RTP stream. Each stream in RTP session has distinct SSRC

RTSP/RTP programming assignment

- ❖ build a server that encapsulates stored video frames into RTP packets
 - grab video frame, add RTP headers, create UDP segments, send segments to UDP socket
 - include seq numbers and time stamps
 - client RTP provided for you
- ❖ also write client side of RTSP
 - issue play/pause commands
 - server RTSP provided for you

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 I OK
Session 423 I

C: PLAY rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
Session: 423 I
Range: npt=0-

C: PAUSE rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
Session: 423 I
Range: npt=37

C: TEARDOWN rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0
Session: 423 I

S: 200 3 OK

Real-Time Control Protocol (RTCP)

- ❖ works in conjunction with RTP
- ❖ each participant in RTP session periodically sends 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
- ❖ feedback used to control performance
 - sender may modify its transmissions based on feedback