# Multimedia Networking

**Raj Jain**
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:

http://www.cse.wustl.edu/~jain/cse473-10/

 


- Multimedia Networking Applications
- Real-Time Streaming Protocol (RTSP)
- Real-Time Transport Protocol (RTP)
- Session Initiation Protocol (SIP)
- Scheduling Mechanisms

**Note**: This class lecture is based on Chapter 7 of the textbook (Kurose and Ross) and the figures provided by the authors.

# Multimedia Networking Applications

- Streaming Stored Audio and Video
  - Stored Media: Fast rewind, pause, fast forward
  - Streaming: simultaneous play out and download
  - Continuous play out: Delay jitter smoothed by playout buffer
- Streaming Live Audio and Video: IPTV and Internet Radio
  - No fast-forward
- High data rate to large number of users
  $\Rightarrow$ multicast or P2P,
  - delay jitter controlled by caching,
- Real-Time Interactive Audio and Video: Internet Telephone, Video Conferencing
  - Delay<400 ms.

 

# Multimedia on Internet

- Best Effort Service
- TCP not used due to retransmission delays
- Limited packet loss tolerated
- Packet jitter smoothed by buffering
- Hard Guarantee: Min Throughput, Max Delay, Max delay jitter
- Soft Guarantee: Quality of service with a high probability
- Protocol for Bandwidth Reservation and Traffic Description
- Scheduling to honor bandwidth reservation
- High Bandwidth
- Content Distribution Networks: Akamai

# Audio Compression Standards

- 4kHz audio ⇒ Audio sampled at 8000 samples per second
- 256 levels per sample ⇒ 8 bits/sample ⇒ 64 kbps
- Pulse Code Modulation (PCM)
- CD's use 44.1 kSamples/s, 16 b/sample ⇒ 705.6 kbps (mono) or 1.411 Mbps (Stereo)
- GSM Cell phones: 13 kbps
- G.711: 64 kbps
- G.729: 8 kbps
- G.723.3: 6.4 and 5.3 kbps
- MPEG 1 Layer 3 (MP3): 96 kbps, 128 kbps, or 160 kbps

# Video Compression Standards

- Moving Pictures Expert Group (MPEG)
- MPEG 1: CD quality video (1.5 Mbps)
- MPEG 2: DVD quality Video 3-6 Mbps
- MPEG 4: Low-rate high-quality video (.divx or .mp4)
- H.261

# Web Server vs. Streaming Server

- Web Servers sends the whole file as one object
- Streaming Server sends at a constant rate

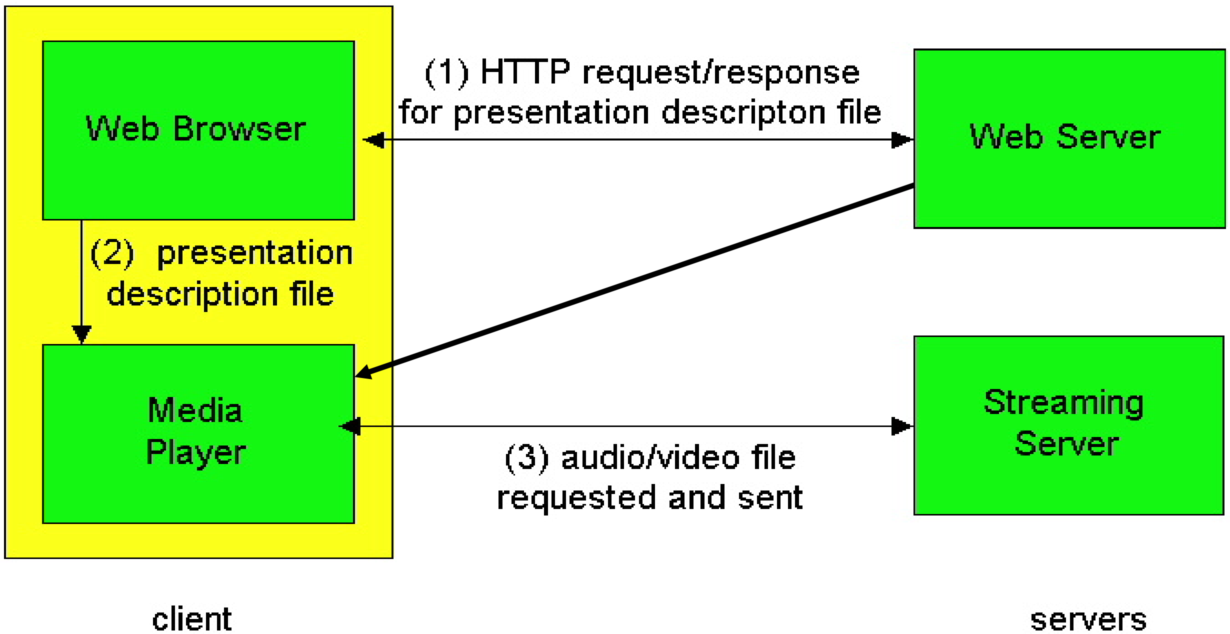

# Real-Time Streaming Protocol (RTSP)

- Protocol to control streaming media
- Allows start, stop, pause, fast forward, rewinding a stream
- Data and control channels
- All commands are sent on control channel (Port 544)
- Specified as a URL in web pages:
  rtsp://www.cse.wustl.edu/~jain/cse473-09/ftp/i_7mmn0.rm

# RTSP Operation

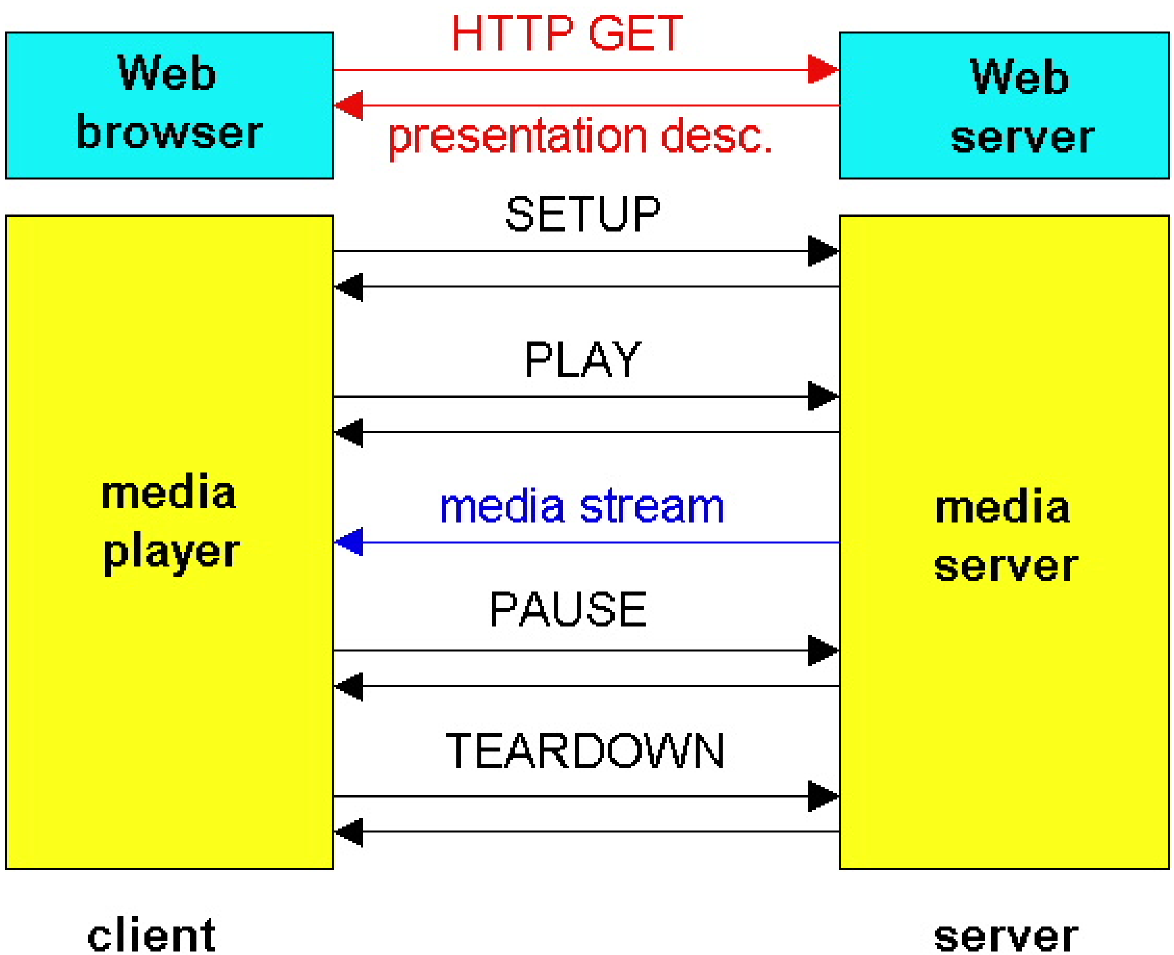

# RTSP Exchange Example

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

# Multimedia with Best Effort Service

- High Compression $\Rightarrow$ Low Rate $\Rightarrow$ Low loss
- 1% to 20% loss can be concealed
- Forward Error Correction (FEC) can be used to overcome loss.
- End-to-end delay limited to 400 ms
- Jitter overcome by play out buffer
- Large jitter $\Rightarrow$ Packets arrive too late $\Rightarrow$ same as Lost
- Each chunk comes with a sequence number and timestamp
- Play out delay can be adaptively adjusted according to measured delay variation

# Playout Buffers

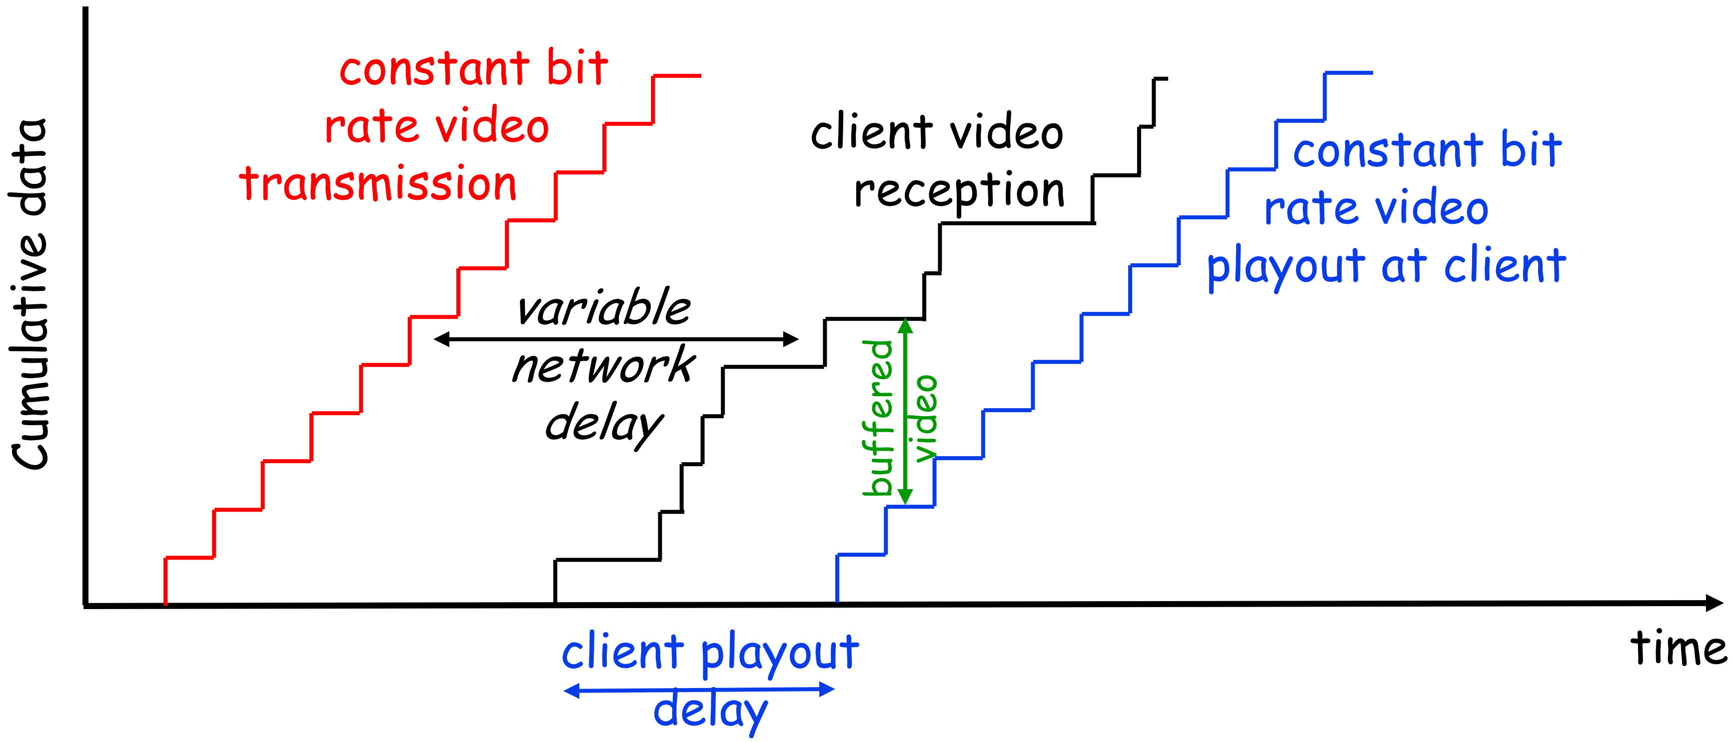


- Playout delay compensates for network delay, delay jitter
- Delay > Playout Delay ⇒ Packet late ⇒ Same as a lost packet

# Adaptive Playout Delay

- $t_i$ =Sending time
- $r_i$= Receiving time
- Measured delay sample = $r_i - t_i$
- $d_i$= Average network delay

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

- $v_i$= Variation of the delay

$$v_i = (1-a)v_{i-1} + a|r_i - t_i - d_i|$$

- $p_i$= Playout time

$$p_i = t_i + d_i + Kv_i$$

- Here K is a constant, say 4.

# Recovering From Packet Loss

- Forward Error Correction
- Send n+1 packets in place of n packets
- Send a lower-resolution stream in addition
- Play out the old syllable
- Busty Loss $\Rightarrow$ Interleave audio/video frames

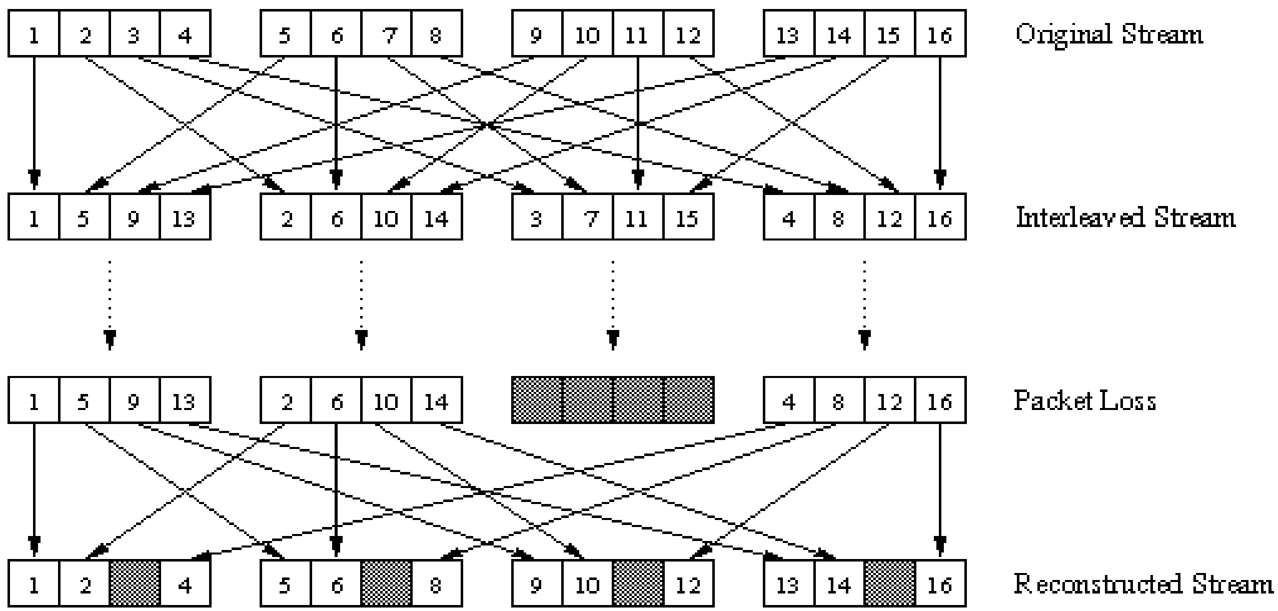

# Content Distribution Networks

- Authoritative DNS server resolves the server address according to the requester's IP address

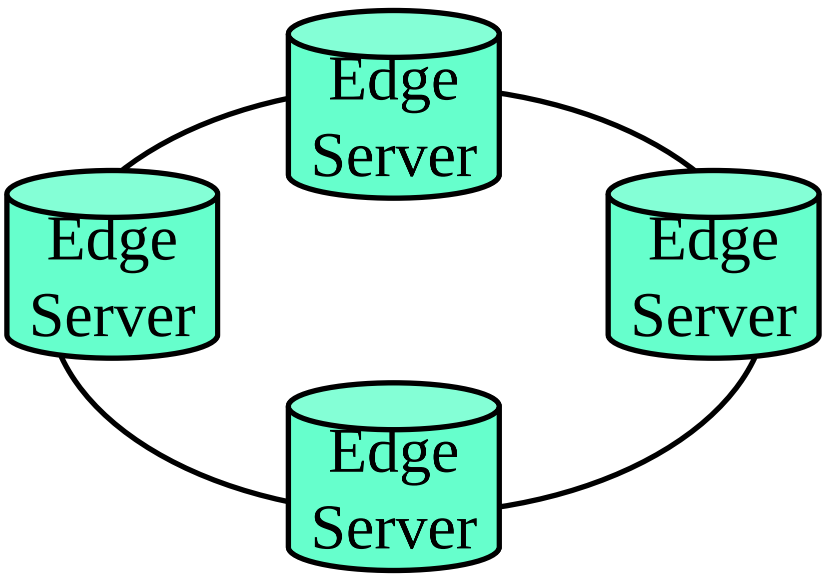

# Real-Time Transport Protocol (RTP)

- Common sublayer between applications and UDP
- Provides sequence numbers, timestamps, and other facilities
- Supports both unicast and multicast

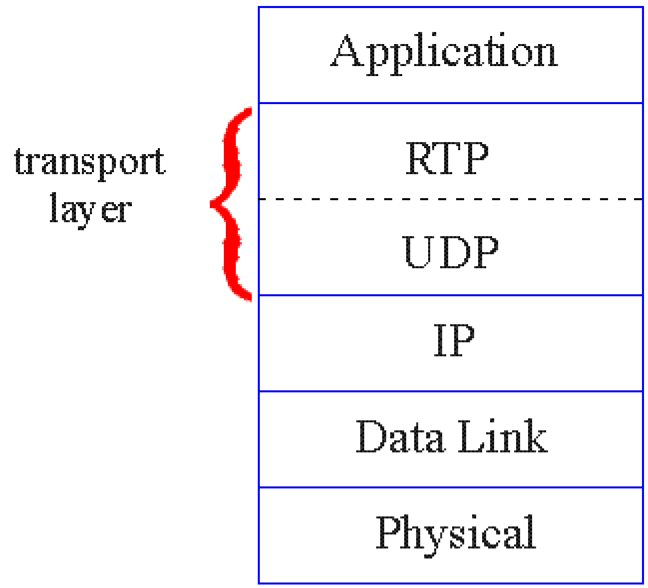

# RTP Packet Format

| Paylcad Type | Sequence Number | Timestamp | Synchron zation Source Identifer | Miscellaneous Fields |
|---|---|---|---|---|
| 8b | 16b | 32b | 32b | |

- SSRC = Synchronization Source Identifier = Stream #

| Payload Type | Coding | Rate |
|---|---|---|
| 0 | PCM mu-law | 64 kbps |
| 3 | GSM | 13 kbps |
| 7 | LPC | 2.4 kbps |
| 26 | Motion JPEG | |
| 31 | H.261 | |
| 33 | MPEG2 video | |

# RTP Control Protocol (RTCP)

- Used to send report about reception quality back to sender
- Also used by sender to report stream information
- Can be used to adjust the transmission speed, quality, or for diagnosis
- SSRC
- Fraction of packets lost
- Last sequence number received
- Inter-arrival jitter
- Receiver report rate is adjusted inversly to number of receivers
- Sender report rate is adjusted inversly to number of senders
- Total RTCP traffic < 5% of media datarate

# Session Initiation Protocol (SIP)

- Application level signaling protocol for voice and video conferencing over Internet
- Allows creating, modifying, terminating sessions with one or more participants
- Carries session descriptions (media types) for user capabilities negotiation
- Supports user location, call setup, call transfers
- Supports mobility by proxying and redirection

# SIP (Cont)

- SIP Uniform Resource Identfiers (URIs):
  Similar to email URLs
  sip:jain@cis.ohio-state.edu
  sip:+1-614-292-3989:123@osu.edu?subject=lecture
- SIP can use UDP or TCP
- SIP messages are sent to SIP servers:
  - Registrar: Clients register and tell their location to it
  - Location: Given name, returns possible addresses for a user. Like Directory service or DNS.
  - Redirect: Returns current address to requesters
  - Proxy: Intermediary. Acts like a server to internal client and like a client to external server

# Locating using SIP

- Allows locating a callee at different locations
- Callee registers different locations with Registrar
- SIP Messages: Ack, Bye, Invite, Register, Redirection, ...

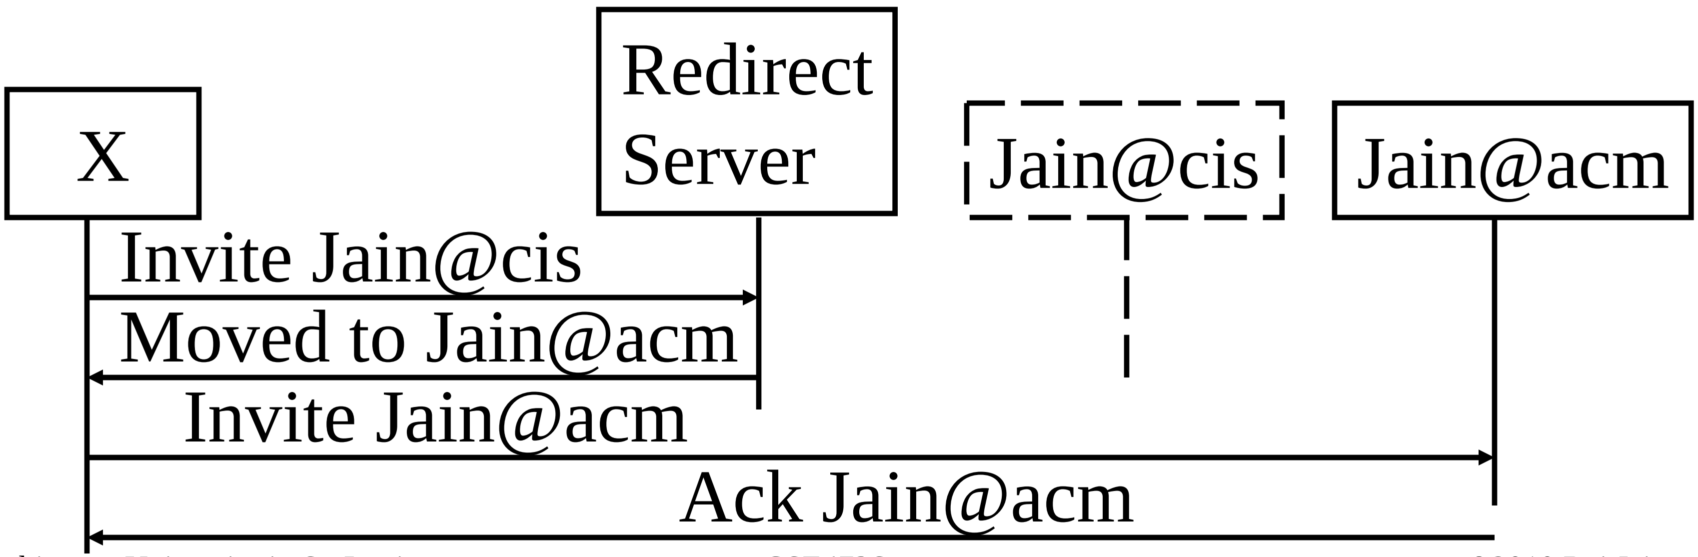

# SIP Proxy

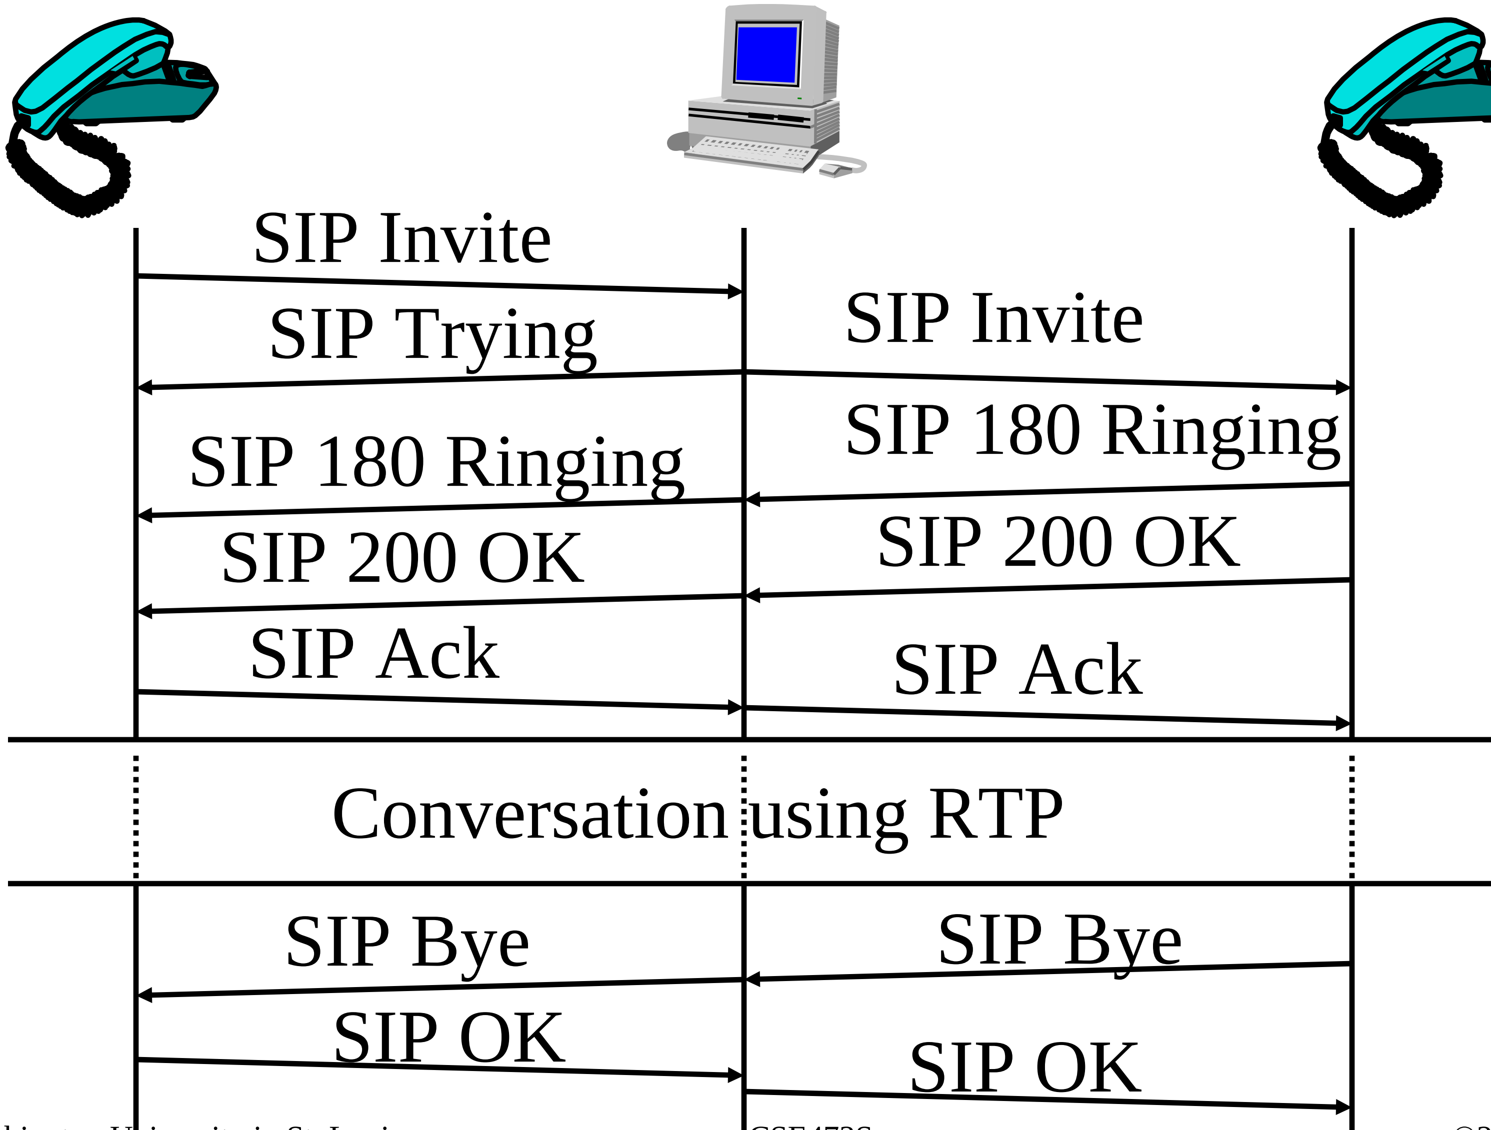

SIP Invite
SIP Trying
SIP Invite
SIP 180 Ringing
SIP 180 Ringing
SIP 200 OK
SIP 200 OK
SIP Ack
SIP Ack
Conversation using RTP
SIP Bye
SIP Bye
SIP OK
SIP OK

# H.323 Protocols

- Multimedia over LANs, V1 (June 96), V2(Feb 98)
- Provides component descriptions, signaling procedures, call control, system control, audio/video codecs, data protocols

<table>
<tr><td>Video</td><td>Audio</td><td colspan="4">Control and Management</td><td>Data</td></tr>
<tr><td>H.261<br>H.263</td><td>G.711, G.722,<br>G.723.1, G.728,<br>G.729</td><td rowspan="2">RTCP</td><td>H.225.0<br>RAS</td><td>H.225.0<br>Signaling</td><td>H.245<br>Control</td><td>T.124</td></tr>
<tr><td colspan="2">RTP</td><td colspan="3">X.224 Class 0</td><td>T.125</td></tr>
<tr><td colspan="3">UDP</td><td colspan="3">TCP</td><td rowspan="2">T.123</td></tr>
<tr><td colspan="6">Network (IP)</td></tr>
<tr><td colspan="7">Datalink (IEEE 802.3)</td></tr>
</table>

# Multiple Classes of Service

- Flow Classification: Based on Source IP, Dest IP, Source Port, Dest Port, Type of Service, ...
- Differentiation: Routers can provide different service to different traffic
- Isolation: One class cannot affect other classes severly

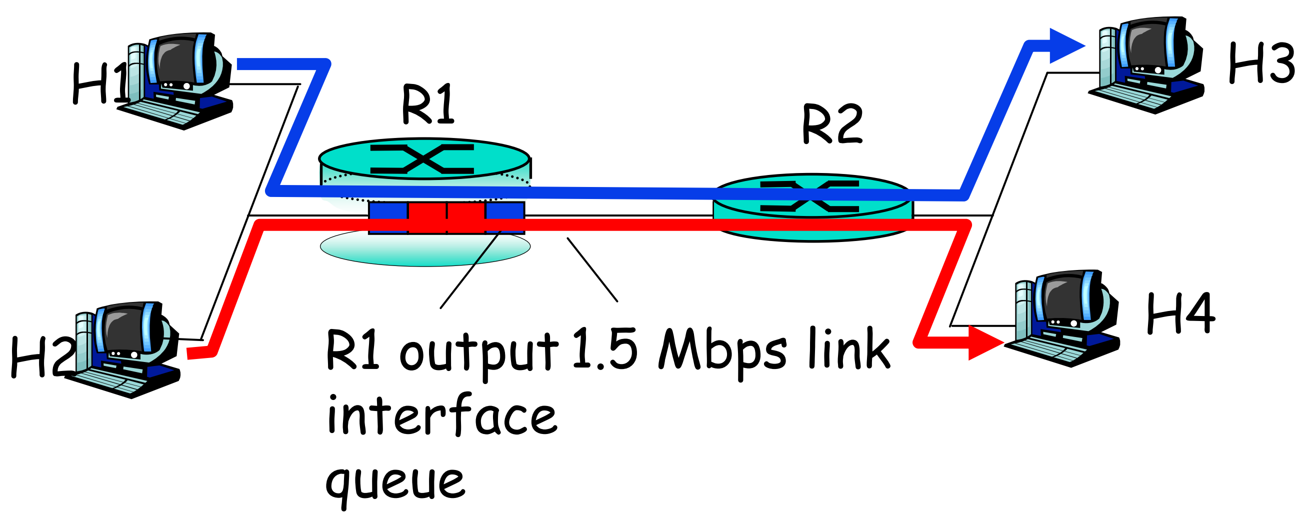

# Scheduling Mechanisms

How to service multiple flows?

- First Come First Served Scheduling
- Priority Queueing
- Round Robin Scheduling
- Generalized Processor Sharing
- Fair Queueing
- Weighted Fair Queueing (WFQ)

Desired Properties:

  - Fair
  - Work-Conserving: Do not waste resources if there is no traffic

# First Come First Served Scheduling

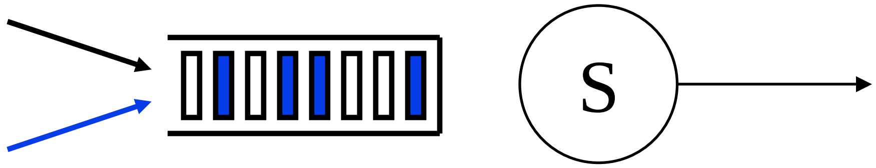


- Unfair: Overloading flows get more service
- No isolation among users

# Priority Queueing

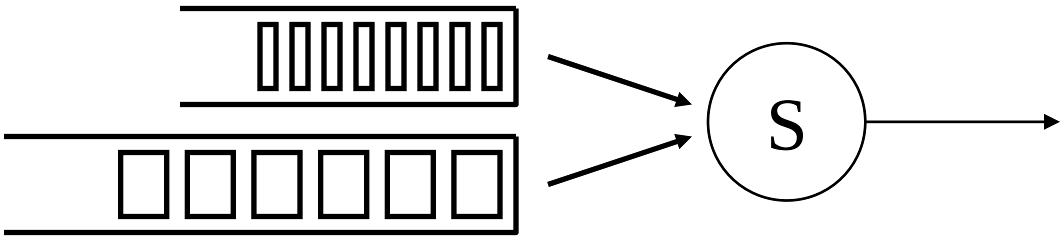


- Priority 0 through n-1
- Priority 0 is always serviced first.
- Priority i is serviced only if 0 through i-1 are empty
- Highest priority has the lowest delay, highest throughput, lowest loss
- Lower priority classes may be starved if higher priority are overloaded

# Round Robin Scheduling

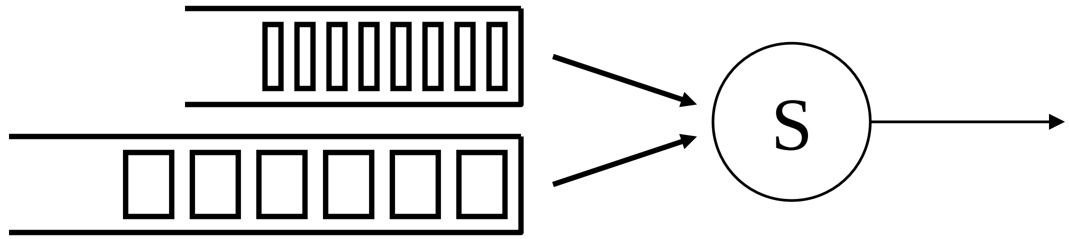


- Round-robin among flows
- Each flow gets the same number of packets
- Flows with larger packets get more bandwidth

# Generalized Processor Sharing

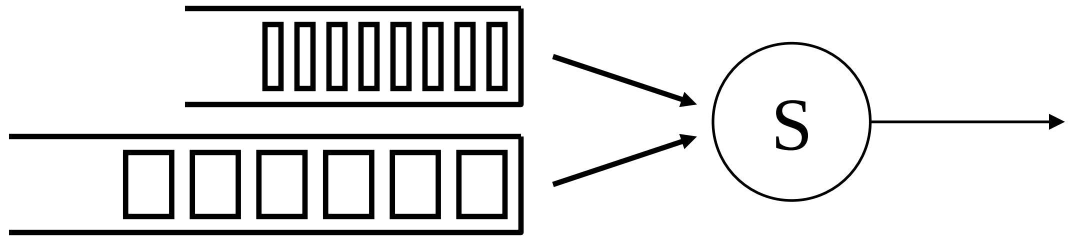


- Bit-level round robin
- Each flow gets the same number of bits/sec
- Too much work

# Fair Queueing

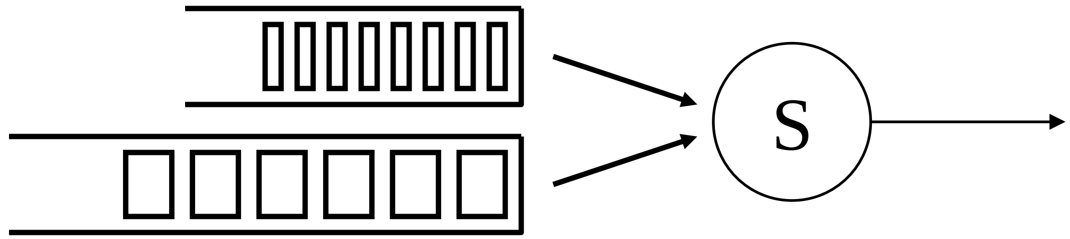


- Bit-level round robin but packet level scheduling
- Count the packet size and determine which packet would finish first. Serve that packet.
- Each flow gets the same number of bits/sec

# Weighted Fair Queueing (WFQ)

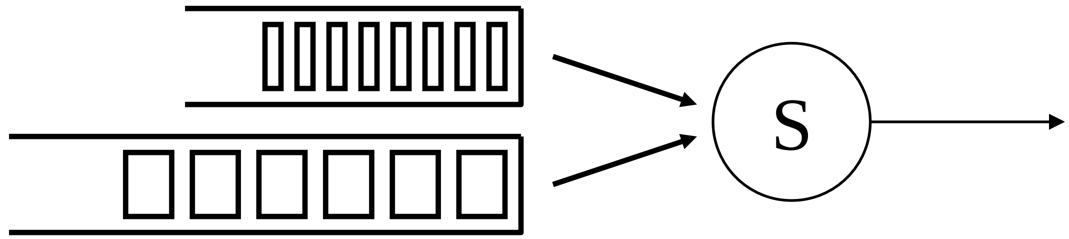


- Fair queueing with different weight for each queue
- Flow 1 gets x bit/sec
- Flow 2 gets y bit/sec
- Flow n gets z bit/sec
- Here, x, y, z are weights

# Summary

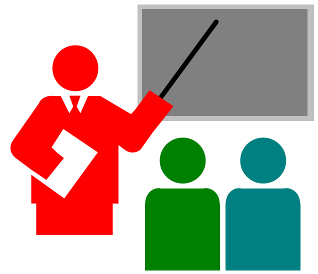

- Multimedia applications require bounded delay, delay jitter, and minimum throughput
- Three Approaches: Service guarantees, Simple priority type service, Increase Capacity
- RTSP allows streaming controls like pause, forward, …
- RTP allows sequencing and timestamping
- SIP allows parameter negotiation and location
- Weighted fair queueing allows packet based fair scheduling

# Review Exercises

- Read Pages 597-657 of the textbook.
- Review Exercises R1-R15
- Problems P2-P4,P9, P11, P16, P19-P22

# Homework 7 (P10)

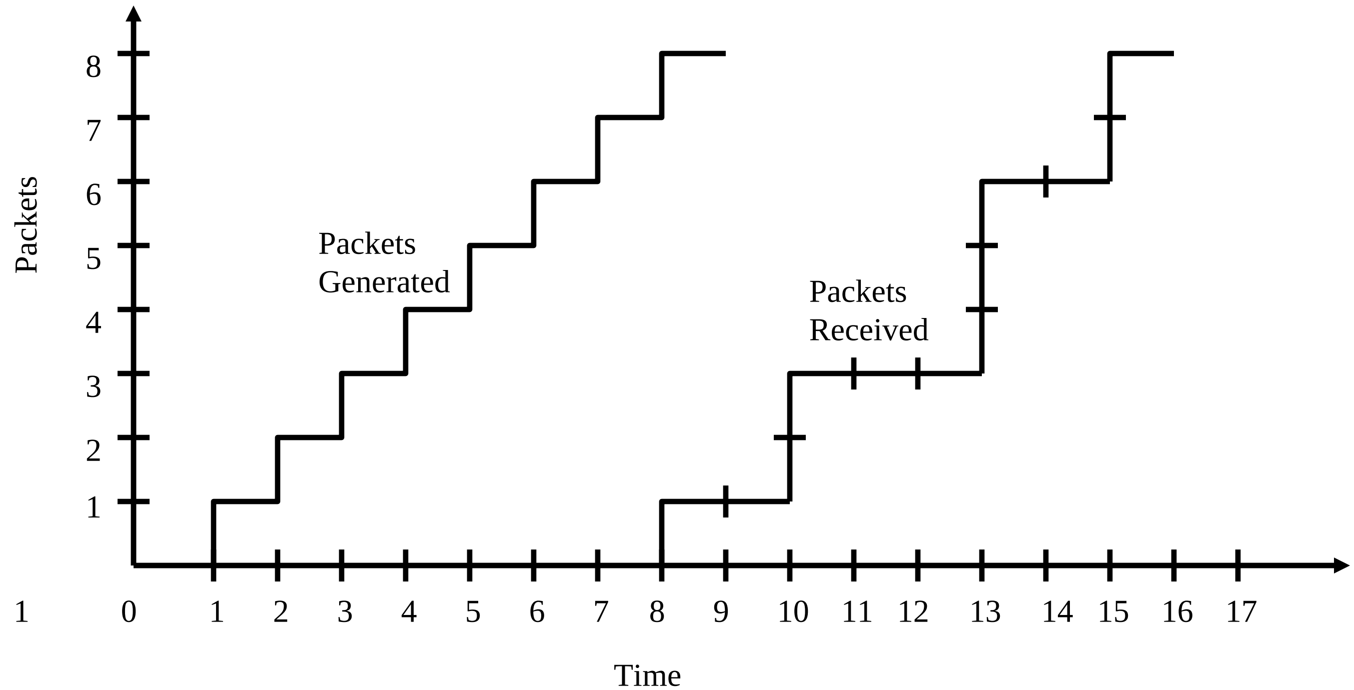


- Consider the packet generation and reception sequence shown below. The first packet is generated at t=1 and is received at t=8.
- A. If Playout delay is zero and playout begins at t=8, which of the packets will not arrive in time?
- B. What is the minimum playout delay at the receiver that result in all of the first eight packets arriving in time for their playout?