

Multimedia Communication

Multimedia Systems

Summary:

- ❑ Multimedia Networking Applications: Requirements
- ❑ Current Networks
 - Limitations & Evolution
- ❑ RTSP

Sources:

- ❑ Chapter 6 from "Computer Networking: A Top-Down Approach Featuring the Internet", by Kurose and Ross

Multimedia Networking Applications

❑ Characteristics:

- *Highly delay-sensitive:*
Packets that incur a sender-to-receiver delay of more than a few hundred milliseconds (for Internet Telephony) to a few seconds (for streaming of stored media) are essentially useless.
- *Loss Tolerant:*
Occasional losses only cause occasional glitches in the audio/video playback, and these losses can be partially of fully concealed.
- ❑ Compare these against "elastic" applications like WWW (text/images), e-mail, ftp and Telnet.
 - Long delays are annoying but not harmful; Integrity of data is paramount.

Example Category 1

Streaming, Stored Audio and Video

- ❑ Clients request on-demand compressed audio or video files that are stored on servers
- ❑ Characteristics:
 - *Stored Media:* Content is prerecorded, the client may pause, rewind, fast-forward or index; The time from when a client makes a request until the action manifests itself at the client should be of the order of 1 to 10 seconds
 - *Streaming:* Client begins playout of the media after (a short delay, maybe zero) it begins receiving the file from server.
 - *Continuous Playout:* Once playout of the media begins, it should proceed according to the original timing of the recording; The end-to-end delay constraints for streaming, stored media are typically less stringent than those of live, interactive applications like Internet telephony and video conferencing.

Example Category 2

Streaming of Live Audio and Video

- ❑ Similar to traditional broadcast Radio/TV, except the transmission takes place over the Internet
- ❑ Characteristics:
 - No fast-forward of media; May use local storage for rewind and pause functionality
 - Distribution to large audiences can be done using Multicast
 - Continuous playout is required, although the timing constraints are less stringent than live interactive media.
 - Delays of up to tens of seconds from when the user requests the delivery/playout of a live transmission to when the playout begins can be tolerated.

Example Category 3

Real-Time Interactive Audio and Video

- ❑ Allows people to use audio and video to communicate with each other.
- ❑ Internet Phone (Audio), Microsoft's Netmeeting (Video)
- ❑ For a conversation with interaction among multiple speakers, the delay from when a user speaks or moves until the action is manifested at the receiving hosts should be less than a few hundred milliseconds
- ❑ For voice,
 - Delays ≤ 150 ms are not perceived by a human listener
 - Delays between 150 and 400 ms are acceptable
 - Delays ≥ 400 ms can be frustrating if not completely unintelligible.

Internet: Limitations

- ❑ Best-Effort Service
 - No guarantee on end-to-end delay
 - No promises about variation of packet delay within a packet stream (Packet Jitter)
 - No class based service: Totally Egalitarian
 - Some techniques to adapt:
 - Use UDP for audio/video communication, to circumvent TCP's low throughput when it enters its slow-start phase (recall TCP congestion control).
 - Delayed (say 100 ms) playback
 - Timestamp packets
 - Prefetch stored media when extra bandwidth is available
 - Send redundant information to mitigate network-induced packet loss

How to Evolve the Internet

There are three camps of thought

- ❑ Camp 1
 - *Philosophy*: No need to make any fundamental changes to the best-effort service, just add more bandwidth to the links.
 - *Arguments against*: Additional bandwidth might be too costly, further it will soon be eaten up by bandwidth hungry applications like high-definition video on demand.
- ❑ Camp 2
 - *Philosophy*: Fundamental changes are needed to allow applications to explicitly reserve end-to-end bandwidth.
 - Need a protocol for reservation
 - Modify scheduling policies at intermediate routers to honor bandwidth reservations
 - Applications must quantitatively describe their requirements and the network must police the applications to make sure they adhere to this.
 - The network must have the means to ascertain whether a request can be satisfied given its current guarantees (Admission Control).

How to Evolve the Internet (Contd.)

- ❑ Camp 3
 - *Philosophy*: Differentiated Service, Make relatively small changes at the network and transport layers and introduce simple pricing and policing schemes at the edge of the network (i.e., at the interface between the user and the user's ISP)
 - The idea is to
 - introduce a small number of service classes
 - assign each datagram to one of the service class
 - intermediate routers prioritize datagrams according to their class
 - charge accordingly

Streaming Stored Media

What has led to the popularity of this class?

- Cheap Storage => More stored media content on the Internet
- Improvements in Internet infrastructure: ADSL, Cable, Caching of Video, new QoS-oriented Internet Protocols
- Enormous demand for high-quality video streaming, a killer app that combines two existing technologies: television and on-demand Web.


How?

- Client:
 - Media Player: Does decompression, jitter removal, error correction, and provide a GUI.
- Server
 - Web Server: Using HTTP (TCP)
 - Streaming Server

Streaming Server

Options for delivering audio/video from a streaming server to a media player:

1. Audio and Video is sent over UDP at a constant rate equal to the drain rate at the receiver
2. Same as above, however, the media player delays playout for 2-5 seconds in order to delay network-induced jitter.
3. Media sent over TCP. The server pushes the media file as quickly as it can into the TCP socket;
 - the media player reads from the TCP socket as quickly as it can and places the compressed video into a player buffer.
 - the player reads from the buffer at a rate of d and performs decompression and playback.



Real Time Streaming Protocol (RTSP)

Henning Schulzrinne's site: <http://www.cs.columbia.edu/~hqs/rtsp/>

Not to be confused with RTP; No relation whatsoever

RTSP: To allow a user (of a media player) to control playback of streamed media, the media player and the server need a protocol for exchanging playback control information.

What RTSP is not:

- RTSP does not define compression schemes
- RTSP does not define how media is encapsulated in packets for transmission over the network (RTP or like may provide this)
- RTSP does not restrict how streamed media is transported; it can be transported using UDP or TCP
- RTSP does not restrict how (or if) the media player buffers media. The media player may choose to play, immediately, delayed or, completely download and then play.

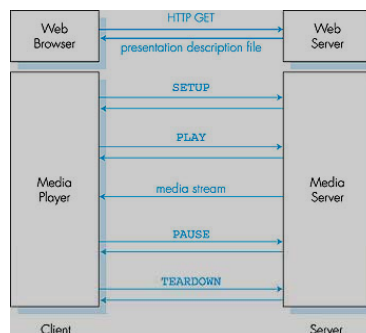
RTSP is:

- A protocol that allows a media player to control the transmission of a media stream
- It's a out-of-band protocol; uses port # 544 (UDP or TCP).

RTSP Example

The Web browser first requests a presentation description file from a Web server.

- The presentation description file can have references to several continuous-media files as well as directives for synchronization of the continuous-media files.
- Each reference to a continuous-media file begins with the URL method, `rtsp://`.



Sample presentation description file

```
<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>
```

In this presentation, an audio and video stream are played in parallel and in lip sync (as part of the same "group").

For the audio stream, the media player can choose ("switch") between two audio recordings, a low-fidelity recording and a high-fidelity recording.

RTSP Operation

- ❑ The player sends an RTSP SETUP request, and the server sends an RTSP SETUP response with a session identifier.
- ❑ Each RTSP session has a session identifier, which is chosen by the server.
- ❑ The player sends an RTSP PLAY request, say, for low-fidelity audio, and the server sends an RTSP PLAY response.
- ❑ At this point, the streaming server pumps the low-fidelity audio into its own in-band channel.
- ❑ Later, the media player may send an RTSP PAUSE request to which, the server responds with an RTSP PAUSE response.
- ❑ The client repeats the session identifier for each request, until the client closes the session with the TEARDOWN request.

RTSP Example Session

- ❑ The following is a simplified example of an RTSP session between a client (C:) and a sender (S:).

```
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
```

- ❑ Notice that in this example, the player chose not to play back the complete presentation, but instead only the low-fidelity portion of the presentation.
- ❑ The RTSP protocol is actually capable of doing much more than described in this brief introduction.
- ❑ In particular, RTSP has facilities that allow clients to stream toward the server (for example, for recording).