

Multimedia streaming

Media streaming

- Media streaming is the distribution of audio and video contents, in general synchronized
- The basic feature of media streaming is that content is used while it is transferred through the network, instead of first downloading the content and then playing it on the user's client
- This is the most important difference between media streaming and the classic file transfer service

Media streaming

- Voice/audio and video codecs play an important role in media streaming
- Some codecs provide just one transmission speed, while layered codecs and multiple-description codecs provide a set of multiple transmission speeds, usually corresponding to different quality of the media
- Layered codecs logically divide the media stream into a set of separate substreams, called layers
- The user must receive the basic layer to be able to reproduce the content, at the lower quality level
- Additional layers can provide better quality, if the user's access bandwidth is enough to transport them
- In this way, layered codecs are flexible as it is possible to adapt media quality to user's resources
- Disadvantage: usually a layered codec consumes more bandwidth than a non-layered codec, given an equal quality level

Media streaming: communication costs

- Media transport costs are a relevant issue for carrier-grade multimedia streaming
- Cost targets vary: a reasonable figure is that the cost of the transport of an average film coded with MPEG2@5 Mbit/s should not exceed 1 USD
- This results into about 0.00023 USD/Mbyte for a content with average duration
- In order to reduce costs of transport, architectural solutions such as caching can be adopted
- Also multicasting is an efficient solution

Media streaming

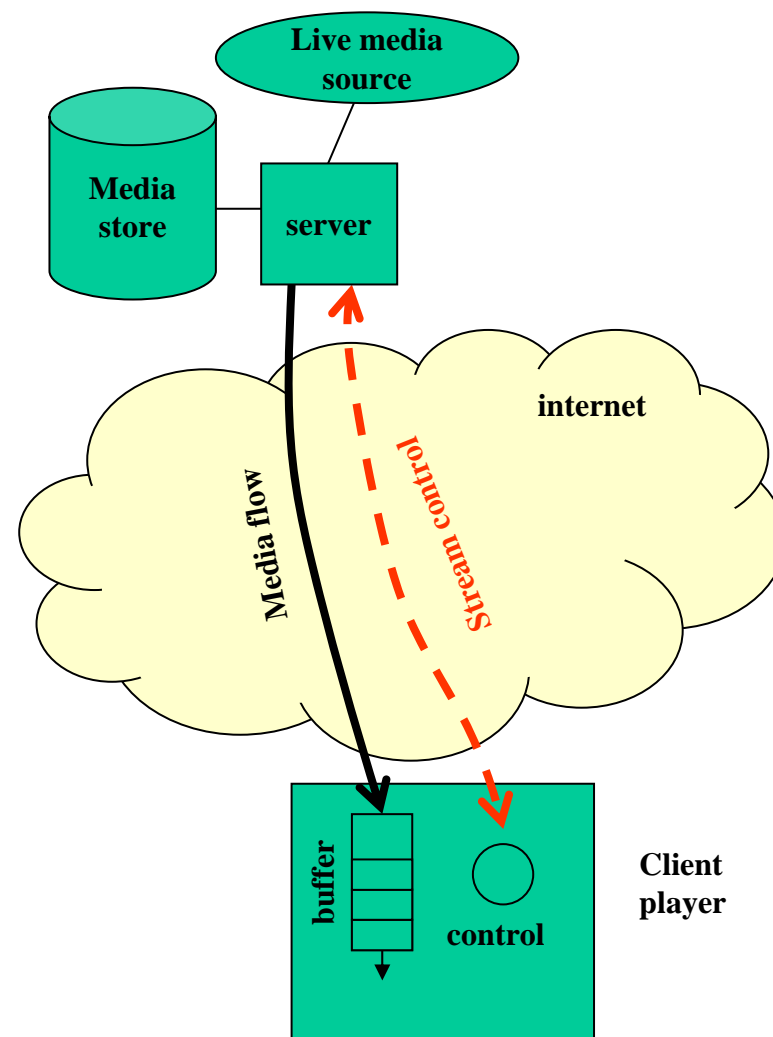
- A content distribution service adopting media streaming as a transport technology can be provided in two main fashions:
 - Pull: the user explicitly requires a specific content
 - Push: contents are provided, through a distribution channel, according to a schedule
- A second characterization of a content distribution service is concerned with delay:
 - Live events (for example a soccer match) require very small delays, on the order of 1 s
 - Films are much more tolerant, as delay is concerned

Playout buffer

- In the user's client, a part of the content is accumulated in the playout buffer before starting the replay of the content
- The playout buffer compensate the varying speed of the network
- The larger the playout buffer, the smaller is the frequency of buffer underflows, that degrade media quality
- However, delay increases as the size of the playout buffer grows
- A tradeoff is to be sought between quality (small number of buffer underflows) and delay
- The tradeoff varies depending on the type of content


Streaming unicast

- Unicast streaming is provided in a classic client-server fashion
- At least two flows are established between client and server. A distribution flow to deliver the media and a control flow to provide to the server quality feedbacks



Streaming unicast

- High hit-rate contents can be placed into peripheral servers (proxies) and then distributed through media streaming from proxies to clients
- The distribution of contents to proxies can be carried out through classic ftp



Streaming multicast

- Multicasting is an efficient way to distribute the same content to many users
- The basic problem is that IP provides primitive multicasting capabilities, insufficient for providing a scalable distribution of contents
- Technologies such as MPLS are quite effective for the implementation of multicasting
- In p2p systems, multicasting is performed at the overlay layer

Real Time Streaming Protocol

- The Real Time Streaming protocol (RTSP) is a client-server protocol for the distribution of multimedia contents over IP networks
- RTSP has been originally developed jointly by Progressive Networks, Netscape Communications and Columbia University

RTSP

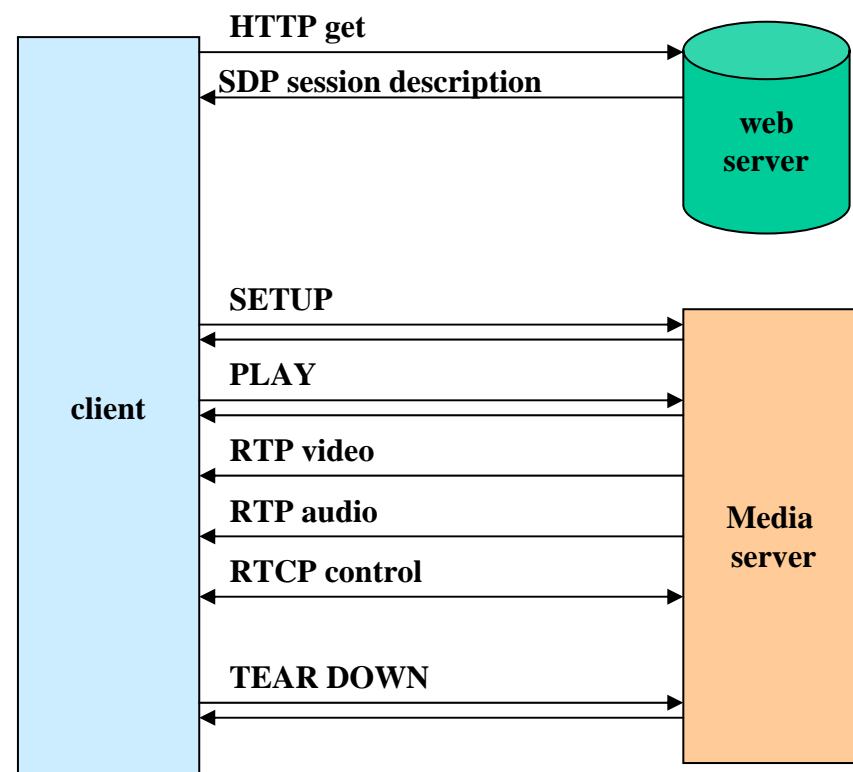
- RTSP can interwork with popular media-related internet protocols:
 - Real Time Protocol and Real Time Control Protocol
 - ReSerVation Protocol
 - Session Initiation Protocol
 - HyperText Transfer Protocol
 - Session Description Protocol

RTSP

- RTSP provides commands to control the media replay:
 - Start
 - Pause
 - Jump
 - Fast forward
 - Fast reverse

RTSP

- Through HTTP the user can obtain the SDP (Session Description Protocol) description of the session
- After the HTTP phase, the streaming of media contents can start



Example (1)

```
// client to web server
GET /Mission_to_Mars.sdp HTTP/1.1
Host: www.movieinfo.com
Accept: application/sdp
```

```
// web server to client
HTTP/1.1 200 OK
Content-Type:application/sdp
V=0
O=2890844526 2890842807 IN IP4 192.16.24.202
S=RTSP session
M=audio 0 RTP/AVP 0
A=control:rtsp://audio.source.com/Mission_to_Mars/audio
M=video 0 RTP/AVP 32
A=control:rtp://video.source.com/Mission_to_Mars/video
```

```
//client to audio server
SETUP rtsp://audio.source.com/Mission_to_Mars/audio RTSP/1.0
CSeq:1
Transport: RTP/AVP/UDP;unicast;client_port=3056-3057
```

Example (2)

```
//client to audio server
SETUP rtsp://audio.source.com/Mission_to_Mars/audio RTSP/1.0
CSeq:1
Transport: RTP/AVP/UDP;unicast;client_port=3056-3057

//audio server to client
RTSP/1.0 200 OK
CSeq:1
Session 12345678
Transport: RTP/AVP/UDP;unicast;client_port=3056-3057;server_port=5000-5001

//client to video server
SETUP rtsp://video.source.com/Mission_to_Mars/video RTSP/1.0
CSeq:1
Transport: RTP/AVP/UDP;unicast;client_port=3058-3059

//video server to client
RTSP/1.0 200 OK
CSeq:1
Session 23456789
Transport: RTP/AVP/UDP;unicast;client_port=3058-3059;server_port=5002-5003
```

Example (3)

```
//play instructions to servers
```

```
PLAY rtsp://video.source.com/Mission_to_Mars/video RTSP/1.0
```

```
CSeq:2
```

```
Session: 23456789 #range: smte=0:10:00-
```

```
RTSP/1.0 200 OK
```

```
CSeq:2
```

```
Session 23456789 #range: smte=0:10:00-
```

```
RTP-info: url=rstp://video.source.com/Mission_to_Mars/video;seq=123122322;rtptime=78712811
```

```
PLAY rtsp://audio.source.com/Mission_to_Mars/audio.it RTSP/1.0
```

```
CSeq:2
```

```
Session: 12345678 #range: smte=0:10:00-
```

```
RTSP/1.0 200 OK
```

```
CSeq:2
```

```
Session 12345678 #range: smte=0:10:00-
```

```
RTP-info: url=rstp://audio.source.com/Mission_to_Mars/audio.it;seq=876655;rtptime=10325465
```


Systems for real time streaming

- In the 90s VOD (Video on Demand) was considered as the “killer application” for broadband networks
- However, the prediction was not true
- The basic problem for setting up a full VoD service through the network is not technological: it is a business problem
- In fact, providing true VoD services is very costly and it is difficult to devise a business model in which costs are compensated by fares that users are willing to pay to obtain the service

Systems for real time streaming

- Since the 90s technology has constantly improved and the main advances have been in the areas of:
 - Advanced codec at relatively lower bit rates
 - More efficient video servers
 - Large Bandwidth user access links

Multimedia servers


- Video (multimedia) server, must store a large number of contents, this calls for a very large storage capacity, on the order of many Terabytes
- The server's storage must also provide large access bandwidth, to accommodate concurrent requests from a large number of users

Multimedia servers

- The cost efficiency of media storage is a critical issue
- Common techniques are hierarchical storage architectures, where low hit rate contents are stored on optical disks or even tapes
- Disks and disk arrays are used massively for high hit rate contents
- Striping is a common solution


Multimedia servers

- RAID (Redundant Array of Inexpensive Disks) storage is a cost-efficient architecture
- A large number of inexpensive disks is used to obtain a total large volume of storage
- Stripes (i.e. chunks of the same media) are stored on different disks



The long tail of contents

- The term “long tail” has been firstly defined by Chris Anderson, the director of the Wired magazine, in 2004



Fonte: Anderson, 2006.

The long tail of contents

- The digital distribution model allows providers to build virtually infinite catalogs of contents, in contrast to, for example, a DVD shop which is mainly limited by the physical space into which DVDs must be kept
- Google, Amazon, iTunes, are examples of log-tailed markets, in which digital contents with high hit rate are relatively few, compared to the vast set of low hit rate contents that the provider can distribute to its customers

The long tail of contents

- A relevant example is that of the book market
- Traditionally, a physical library must choose books complying with the constraint of the available physical space
- In such a business, naturally the provider concentrates on high hit rate objects, leaving the virtually infinite set of low hit rate object untouched

YouTube

- YouTube is a typical example of log-tailed distribution of digital contents
- You Tube provides both contents with high hit rate and contents with very small hit rates
- The number of contents distributed by You Tube is virtually unlimited
- You Tube also provides user-generated contents, a business that cannot be reached by classic content providers