

# Outline

- |     |   |   |
|-----|---|---|
| 1.  | Introduction and Motivation                         |   |
| 2.  | Interactive Web Applications                        | Part I:<br>Web Technologies<br>for Interactive MM   |
| 3.  | Web Programming with Java                           |   |
| 4.  | Communities, the Web, and Multimedia                |   |
| 5.  | Digital Rights Management                           |   |
| 6.  | Cryptographic Techniques                            | Part II:<br>Content-Oriented<br>Base Technologies   |
| 7.  | Multimedia Content Description                      |   |
| 8.  | Electronic Books and Magazines                      |   |
| 9.  | Multimedia Content Production and Management        | Part III:<br>Multimedia<br>Distribution<br>Services |
| 10. | Streaming Architectures                             |   |
| 11. | Web Radio, Web TV and IPTV                          |   |
| 12. | Multimedia Conferencing                             | Part IV:<br>Conversational<br>Multimedia Services   |
| 13. | Signaling Protocols for<br>Multimedia Communication |   |
| 14. | Visions and Outlook                                 |   |

# 10 Streaming Architectures

## 10.1 High-Level Streaming Architecture

## 10.2 Push-Model Streaming: Real-Time Data Transport

## 10.3 Pull-Model Streaming: Progressive Download

## 10.4 Scalability and Multicast

## 10.5 Streaming Market

### Literature:

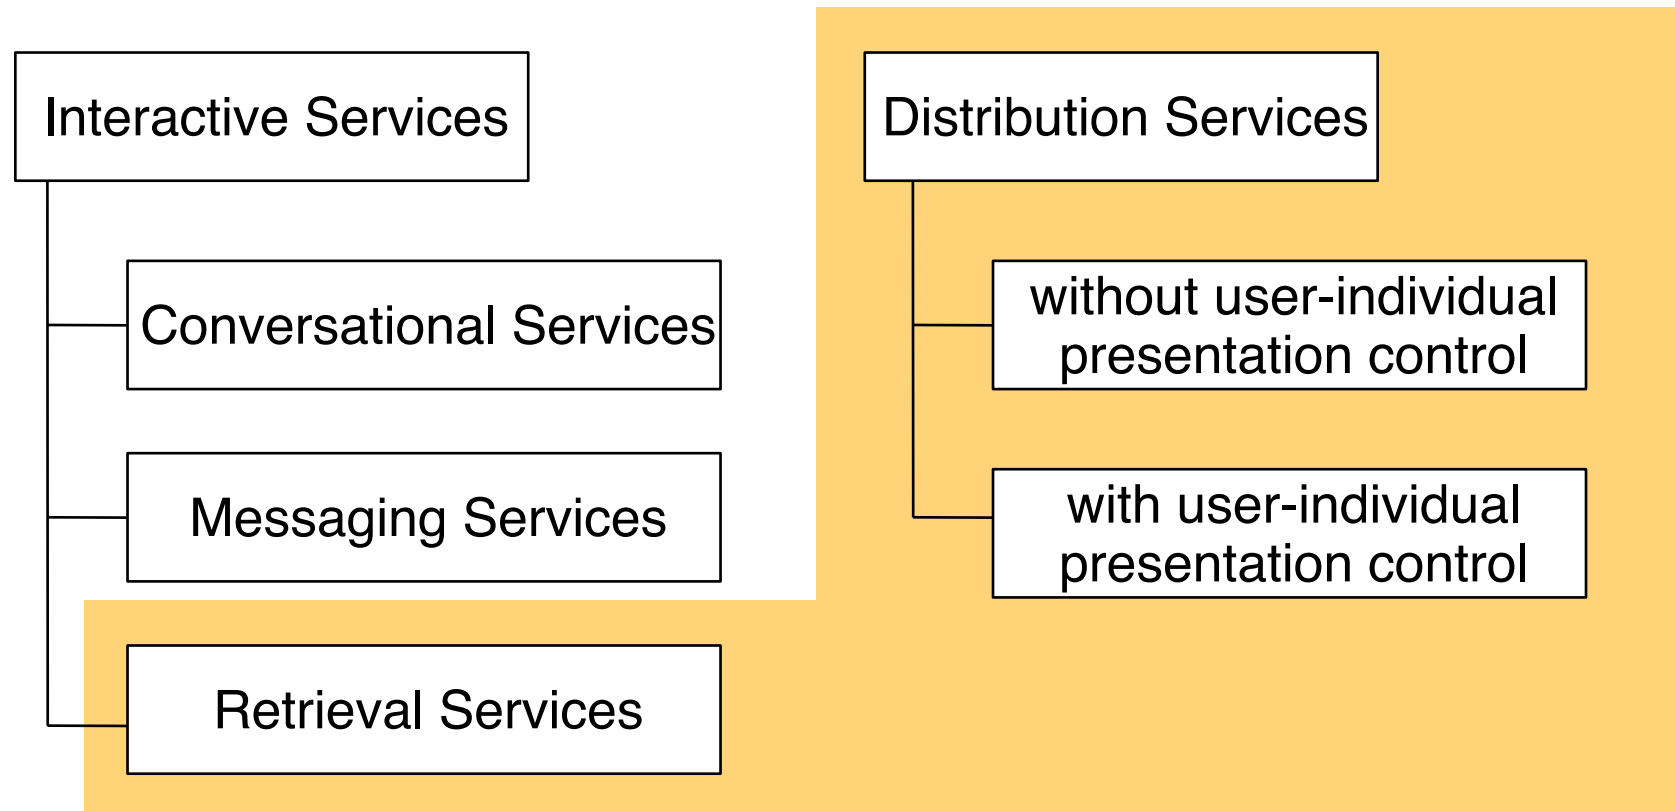
David Austerberry: The Technology of Video & Audio Streaming,  
Focal Press 2002

Gregory C. Demetriades: Streaming Media, Wiley 2003

Tobias Künkel: Streaming Media – Technologien, Standards,  
Anwendungen, Addison-Wesley 2001

# A Classification of Multimedia Services


- According to ITU-T recommendation I.211 “B-ISDN Service Aspects”




# Presentation Control in Distribution Services

- **Without** user-individual presentation control:
  - Continuous “live” stream of information from sender to receiver(s)
  - Replacement of other distribution media (e.g. radio) by digital networks
    - » Real-time service (e.g. fixed start time for transmission)
- **With** user-individual presentation control:
  - Pause, resume, skip backward, skip forward
  - **Server** control:
    - » Individual stream (or group of closely related streams)
    - » Full interactivity (including “seek” and “fast-forward”)
  - **Client-local** control:
    - » Outdated solution for interactivity: “Near Video-on-Demand” (NVOD) = Staggered broadcast of multiple transmissions of the same content
    - » Time-shifted recording enables limited interactivity (pause/resume)
    - » Download in advance enables “fast-forward”


# Unicast, Broadcast, Multicast, Anycast




Unicast:  
One specific receiver




Broadcast:  
Many receivers, all on the network



Multicast:  
Many receivers, all of a specific group



Anycast:  
One receiver, "nearest" of a specific group



Geocast:  
Many receivers, all of a geographic region

Pictures: Wikipedia

# Bandwidth Economy

- Fully heterogeneous individual requests:
  - Required bandwidth = stream bandwidth x number requests
- Homogeneity of request helps saving bandwidth:
  - Same content for many clients, but different playback times:
    - » Broadcast with caching
  - Same content at same time for many clients
    - » Multicast (splitting streams)
- Pre-planning saves bandwidth
  - (Individual) transmission of pre-booked content in non-real time (“download and play”)

# Streaming, Streaming Media

- *Streaming media* means real-time delivery of moving images, moving text and sound, from a server to client (over the Internet).
- Delivery types for audio and video content:
  - *Download and Play*: Content must be downloaded completely to the client before it can be played
  - *Progressive Download*: Playback is started while download is still in progress. Download rate independent of program bit rate.
  - *True Streaming*: Delivered media is viewed/listened in “real-time”.
    - » Playback takes place with roughly the same rate as delivery of data
    - » Delay between send and receive event of data packet kept small
- Subtypes of True Streaming:
  - *Static File Streaming*: Delivery of pre-recorded media files. Often also called *on-demand* delivery (e.g. *Video on Demand*)
  - *Live Streaming*

Based on material from [www.streamingalliance.org](http://www.streamingalliance.org)

# Session


- A *session* is an association between communicating parties, which
  - Persists over a limited time span
  - Incorporates at least two parties
  - May comprise a large number of communication connections of different characteristics
- Examples of sessions:
  - Movie streamed to consumer, consisting of audio and video parts
  - Multimedia conference among five participants, consisting of audio and video source from each of the participants (plus possibly some global information)
- Session awareness at which levels?
  - At application level: unavoidable
  - At network level: possible
    - » Requires specific protocols



# Push and Pull Models for Streaming


- Push model (e.g. Darwin Streaming Server):
  - Session-level connection established between server and client
  - Server continues sending packets *downstream* to client
  - Server listens to commands given by client (see later for protocols)
  - True real-time data distribution
  - Adaptive bandwidth control
- Pull model (e.g. YouTube):
  - Session-level connection established between server and client
  - Server is idle as long as client does not request data
  - Client continues requesting packets from server (e.g. by HTTP)
  - Playback starts after a certain client-side buffer level is reached
  - Essentially “progressive download”, enhanced in modern approaches (see later)

# Streaming Delivery Chain




# Encoding

- Format conversions
  - E.g. analog/digital conversion
  - E.g. downscaling of picture size
- Compression
  - Adequate for player capabilities and typical transmission bandwidth
- Indexing
  - Analyzing internal structure
- Metadata creation
  - Possibly including digital rights specification




# Serving

- Storage
  - Live: only buffering and archiving
  - Static files: archive management, retrieval
- Streaming
  - Pull model:
    - » Request-response driven similar to Web server
  - Push model:
    - » Server sends packets at regular intervals
- Interactivity
  - VCR-like control (PLAY, STOP, PAUSE, FFWD, REW)
  - Random access based on various criteria
  - Hyperlinks in A/V material (“hypervideo”)
- Transcoding
  - Encoding adapted to client capabilities



# Distribution

- Key topic: Quality of Service (QoS)
  - Determining realizable bandwidth, delay, jitter
- Key concepts:
  - Overprovisioning
  - Detailed reservations (“Integrated Services”, reservation protocol RSVP)
    - » Difficult to scale to large numbers of users
  - Traffic classes (“Differentiated Services”)
    - » Difficult to control access to privileges
  - Resource management layer
  - Technology-specific solutions
    - » E.g. ATM (Asynchronous Transfer Mode)



# 10 Streaming Architectures

10.1 High-Level Streaming Architecture

10.2 Push-Model Streaming: Real-Time Data Transport

10.3 Pull-Model Streaming: Progressive Download

10.4 Scalability and Multicast

10.5 Streaming Market

Literature:

David Austerberry: *The Technology of Video & Audio Streaming*,  
Focal Press 2002

Stephan Rupp, Gerd Siegmund, Wolfgang Lautenschlager:  
*SIP – Multimediale Dienste im Internet*. dpunkt 2002

A. Begen, T. Akgul, M. Baugher: *Watching Video over the Web*,  
Part I: Streaming Protocols, *IEEE Internet Computing*,  
March/April 2011

# IP and TCP

- Internet Protocol
  - Network communication protocol (ISO layer 3)
  - Packets transferred from address to address (through routers)
  - Main problems:
    - » Variable network latency
    - » Packet order on arrival may be different than on sending
    - » Packets may be lost
- Transport Control Protocol (TCP)
  - Connection establishment (by “three-way handshake”)
    - » Connections are sequences of associated IP packets
  - Sequencing of bytes with forwarding acknowledgement number
  - Non-acknowledged bytes are re-transmitted after a defined time period
  - Flow control
- For audio/video streaming:
  - Retransmissions (and associated delays) are harmful
  - Lost packets can be tolerated to some extent

# UDP

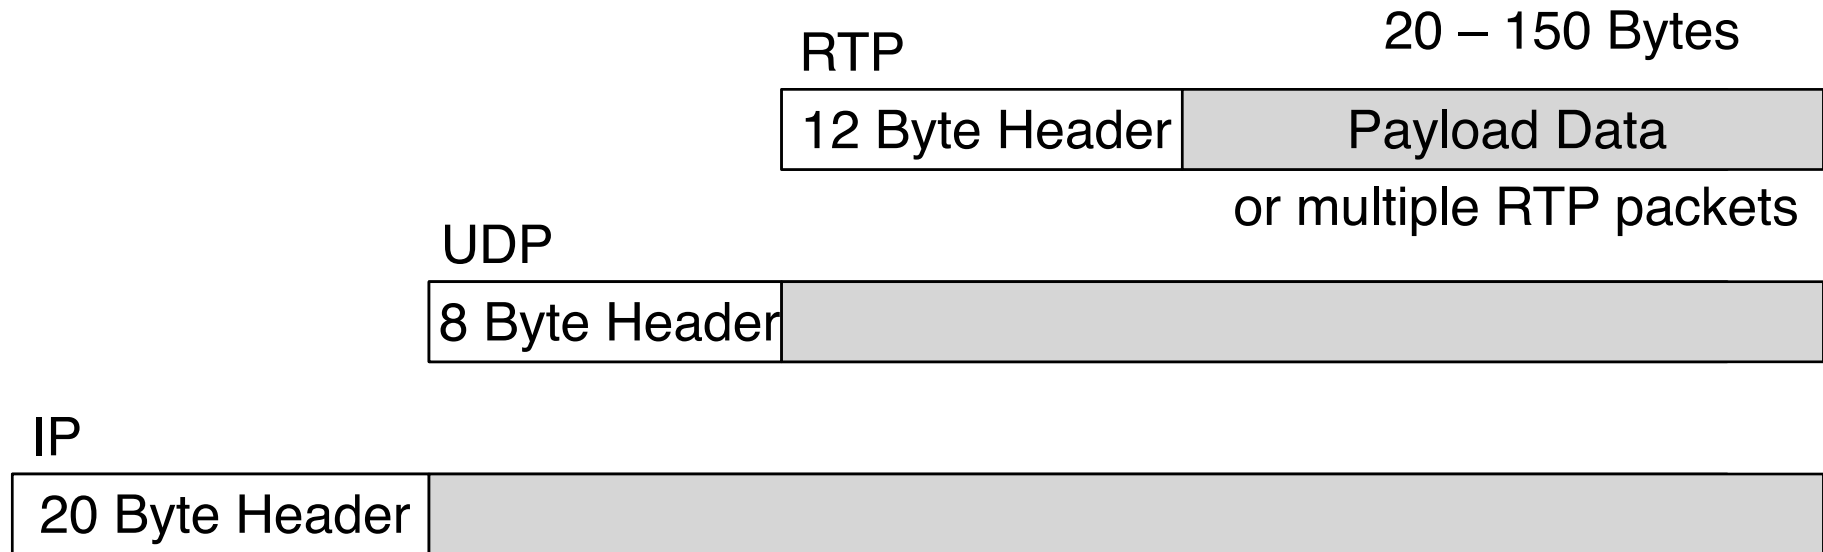
- User Datagram Protocol (UDP)
- Extremely simple transport protocol over IP
  - Connectionless (TCP: connection-oriented)
  - Unreliable (TCP: reliable)
  - No flow control (TCP: has flow control)
- Contents of a UDP datagram:
  - Ports used by application program
  - Checksum
- Basically adequate for media data transport
  - In particular for **push**-model true streaming
  - Very efficient, protocol overhead of TCP avoided
  - Flow control and handling of packet loss have to be handled by higher protocol layer



# Real-Time Transport Protocol RTP

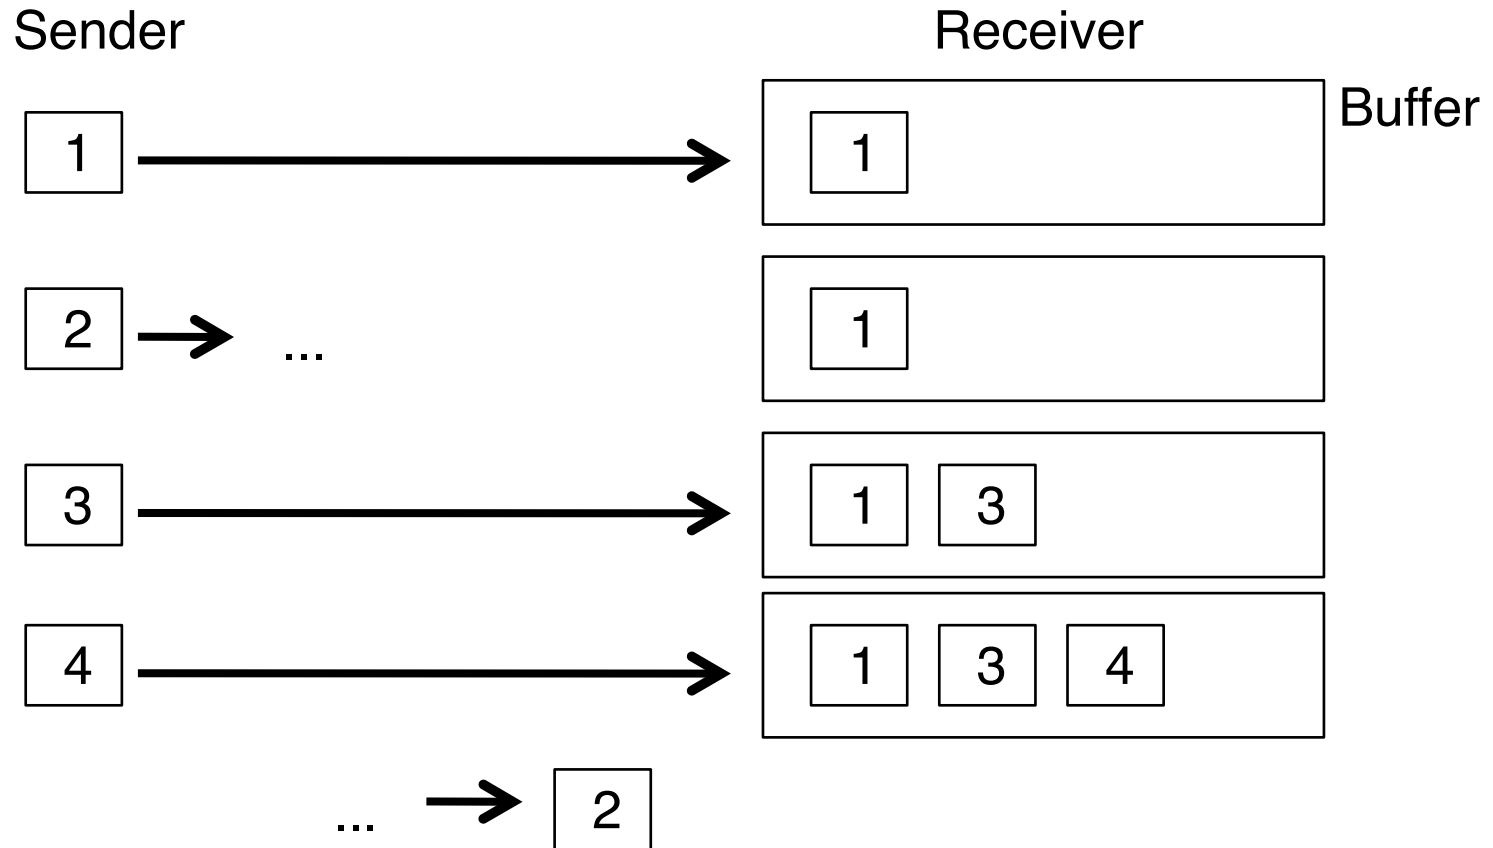
- Transport protocol specifically developed for streaming data
  - IETF (Internet Engineering Task Force) RFC (Request for Comments) 1889
- RTP usually carried over UDP and used for **push** model mainly
- Very important:
  - **RTP does not at all change the way how IP packets are transferred in the network!**
  - To achieve “Quality of Service”, additional network technologies are required (see above)
- RTP used (mainly historically) by:
  - Apple QuickTime architecture
  - RealSystems streaming architecture
- Modern Internet structure does no longer support RTP at many places
  - in particular, in Content Delivery Networks, see later

# RTP Packets and Other Protocols



- IP Header:
  - Source address, destination address, length, time to live, ...
- UDP Header:
  - Port numbers (source and target processes), length, checksum
- RTP Header:
  - Codec type, sequence number, timestamp, synchronization source


# Jitter and Loss Compensation



- Options for application on receiver side:
  - Wait (*not* adequate), repeat last packet (1), interpolate (between 1 and 3)
  - Missing audio information is difficult, missing video can be compensated

# Adaptive Transmission Rate Control

- Application-level mechanism
- Define “downshift” and “upshift” thresholds on buffer (possibly several)
- Downshift:
  - Buffer is close to drain (“underrun”)
  - Server should select lower-bitrate encoding (less quality)
- Upshift:
  - Buffer is close to be full (“overrun”)
  - Server can try to select higher-bitrate encoding (better quality) (if network can support that)
  - Alternatively server may switch to lower transmission rate
- Communication between client/server required!



# Session Description Protocol (SDP)

- IETF Standard (RFC 4566) 2006
- Description mechanism for multimedia communication sessions
  - announcement, invitation, parameter negotiation
- Defines:
  - media components (an their types) belonging to a session
  - alternative formats available for media components
  - session profiles
- Applicable for various purposes
  - for conferencing examples see chapter 13 later
    - » in streaming: announcement of available options from server to client

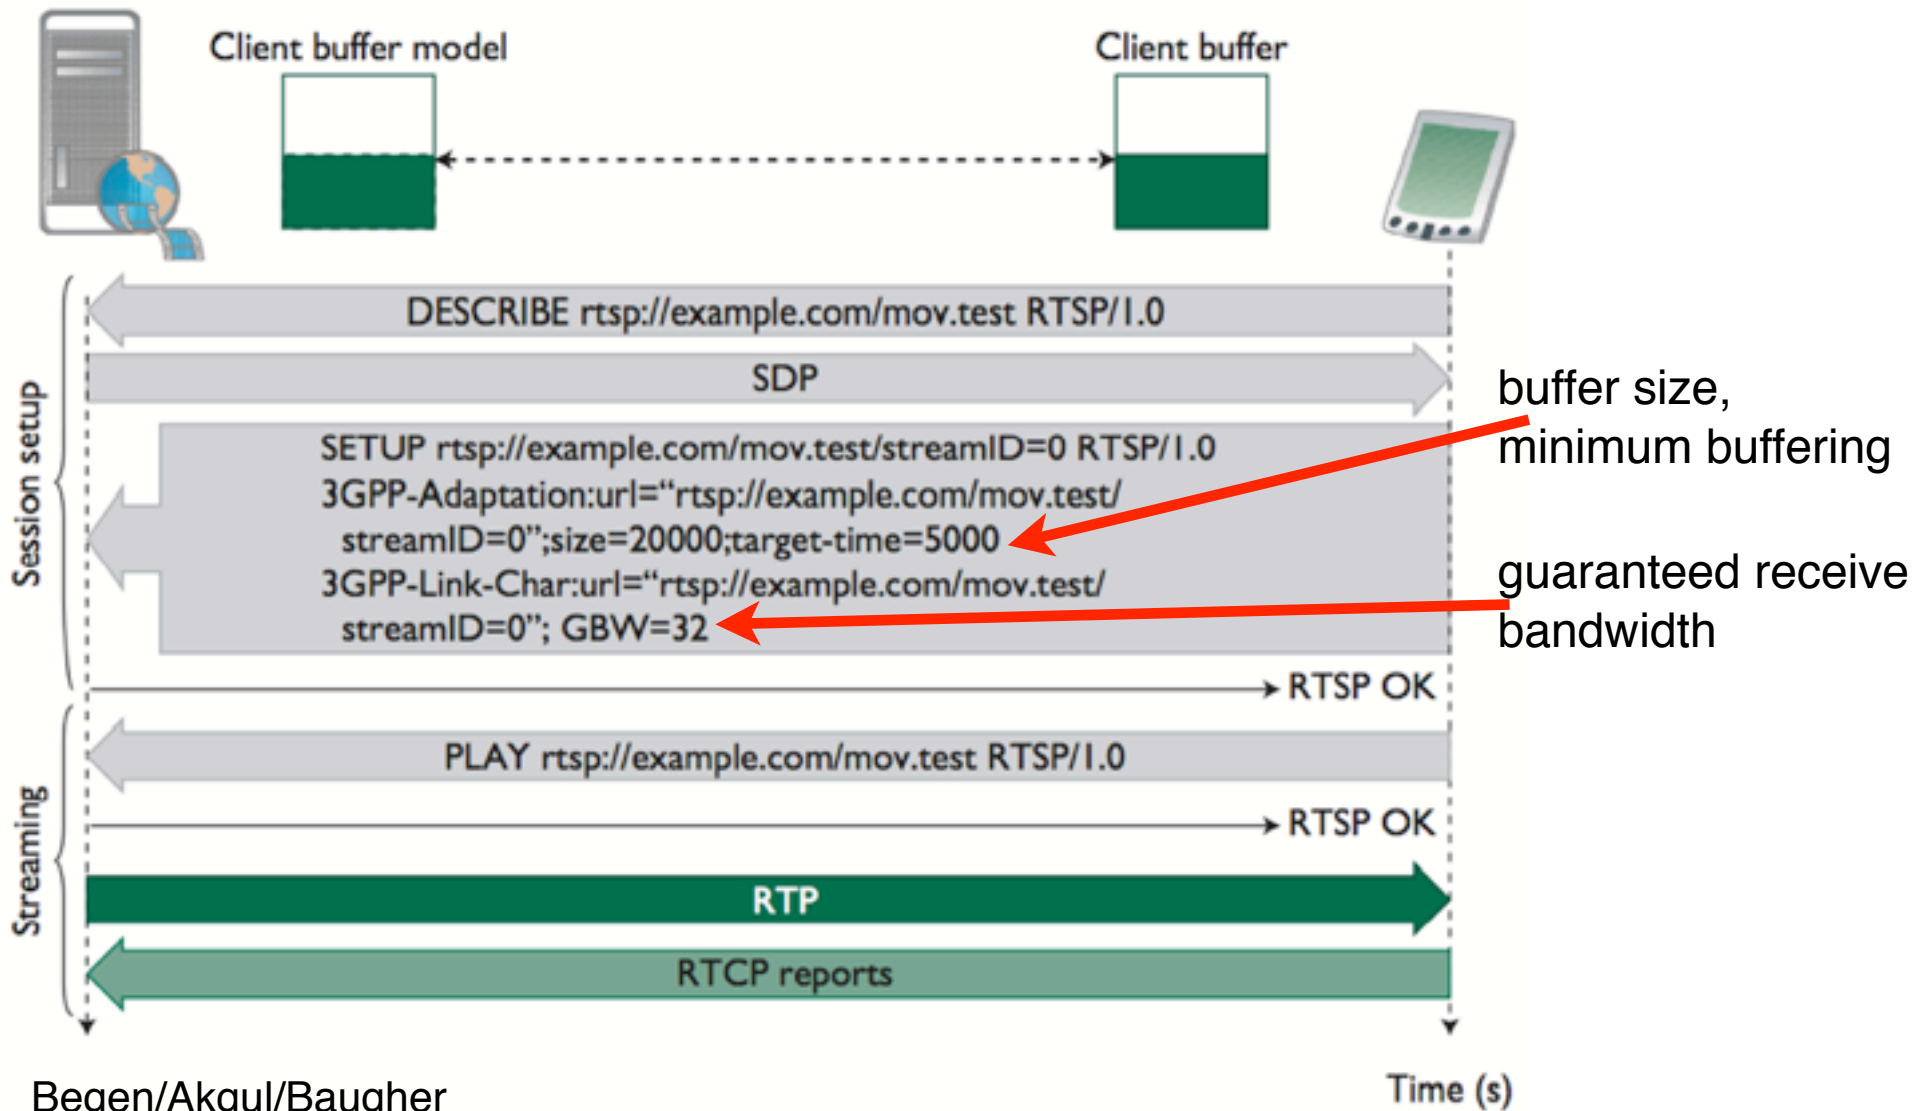
# Real-Time Control Protocol RTCP

- RTCP controls the transmission (not the setup of connection)
- RTCP periodically sends monitoring information to all participants in a streaming session
- Main functions of RTCP:
  - Feedback on QoS of transmission
    - » Information for adaptive codecs, e.g. whether problem is local or global
  - Identification of sender by “canonical name”
    - » Helpful when synchronization source changes
    - » Supports lip synchronization between audio and video
  - Number of participants in a session
    - » Adaptation of sending rate of RTCP control information to number of participants, to avoid network overload
  - Transmission of additional information, e.g. names of session participants

# Real Time Streaming Protocol RTSP

- Client-server multimedia presentation protocol, designed specifically for streamed media
  - IETF (Internet Engineering Task Force) RFC (Request for Comments) 2326 (“MMUSIC” work group), 2004
  - “The Internet VCR remote control protocol” ([www.rtsp.org](http://www.rtsp.org))
  - Independent of the use of RTP for transport
  - Syntactically similar to HTTP 1.1 (carried over TCP or UDP)
- Main operations supported by RTSP:
  - Transport & capability negotiation (SETUP)
    - » e.g. disallowing a “seek” function
  - Session control (SETUP, REDIRECT, TEARDOWN)
  - Control of media playback (PLAY, PAUSE, PING)
  - Invitation of a media server to a conference

# 3GPP Bitrate Adaptation



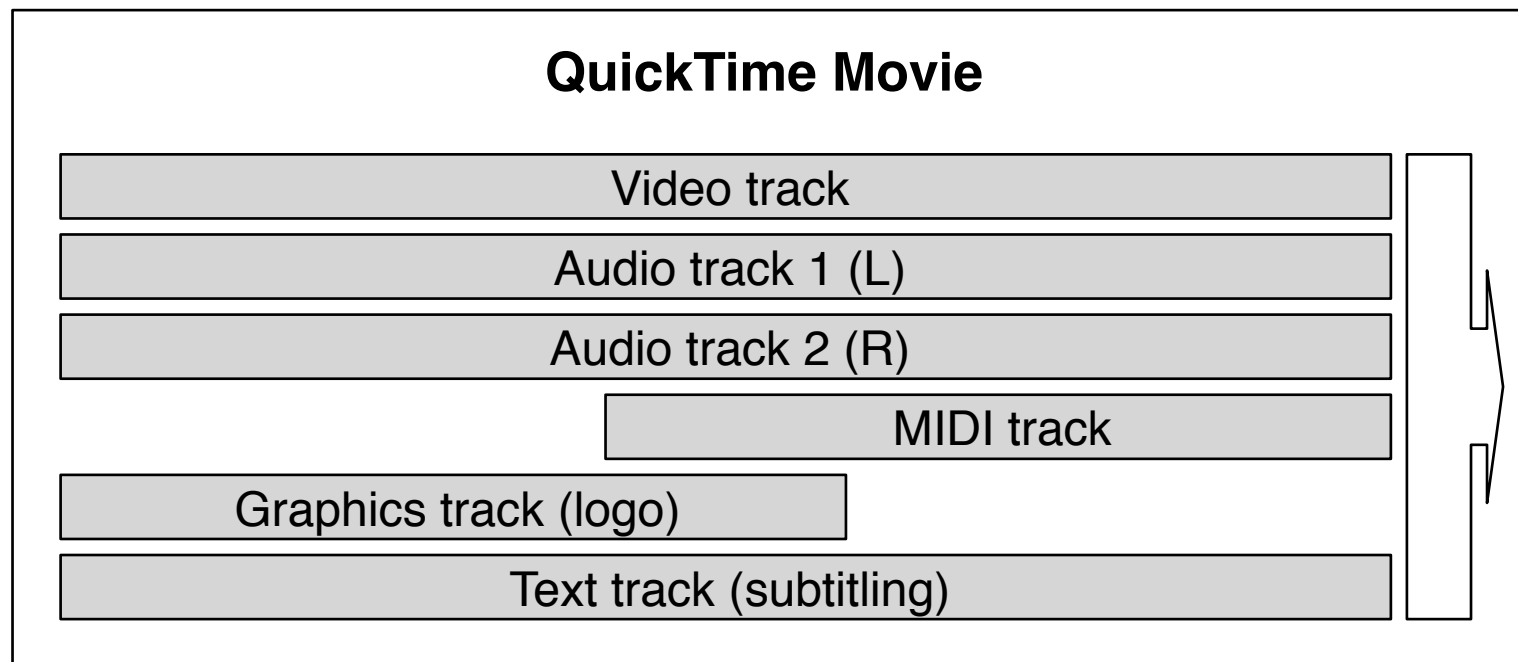
Begen/Akgul/Baugher



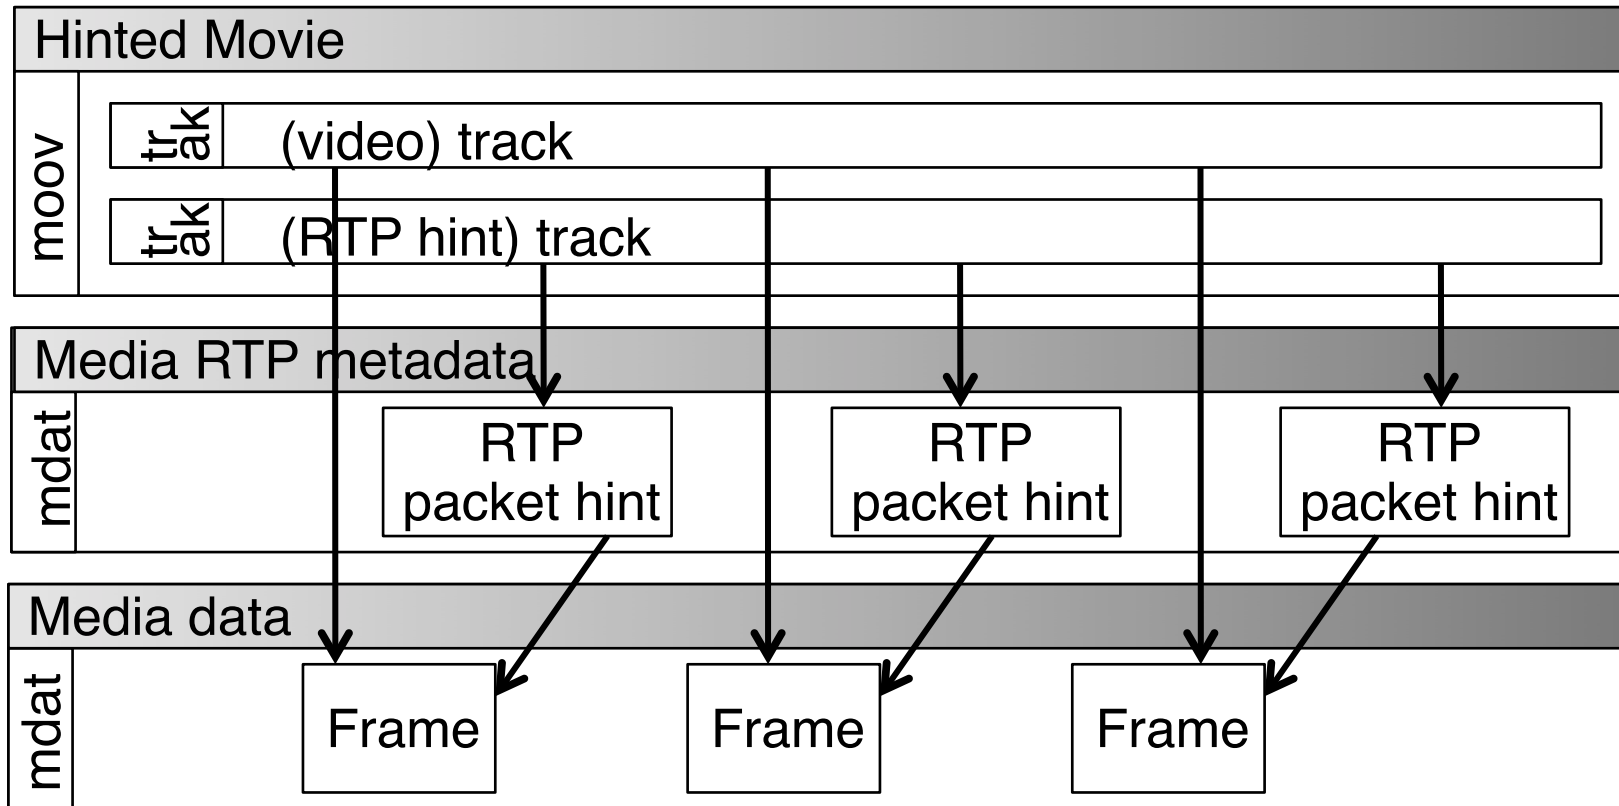
# Example File Format: QuickTime Movie Files



- Modular and flexible architecture
  - Multimedia files organized in tracks
  - Example:



# Hint Tracks in QuickTime and MPEG-4



- Hint track gives server software pointers to the RTP information to serve the relevant media chunks
- Concept from QuickTime, integrated in MPEG-4 (streaming)

# 10 Streaming Architectures

10.1 High-Level Streaming Architecture

10.2 Push-Model Streaming: Real-Time Data Transport

10.3 Pull-Model Streaming: Progressive Download

10.4 Scalability and Multicast

10.5 Streaming Market

Literature:

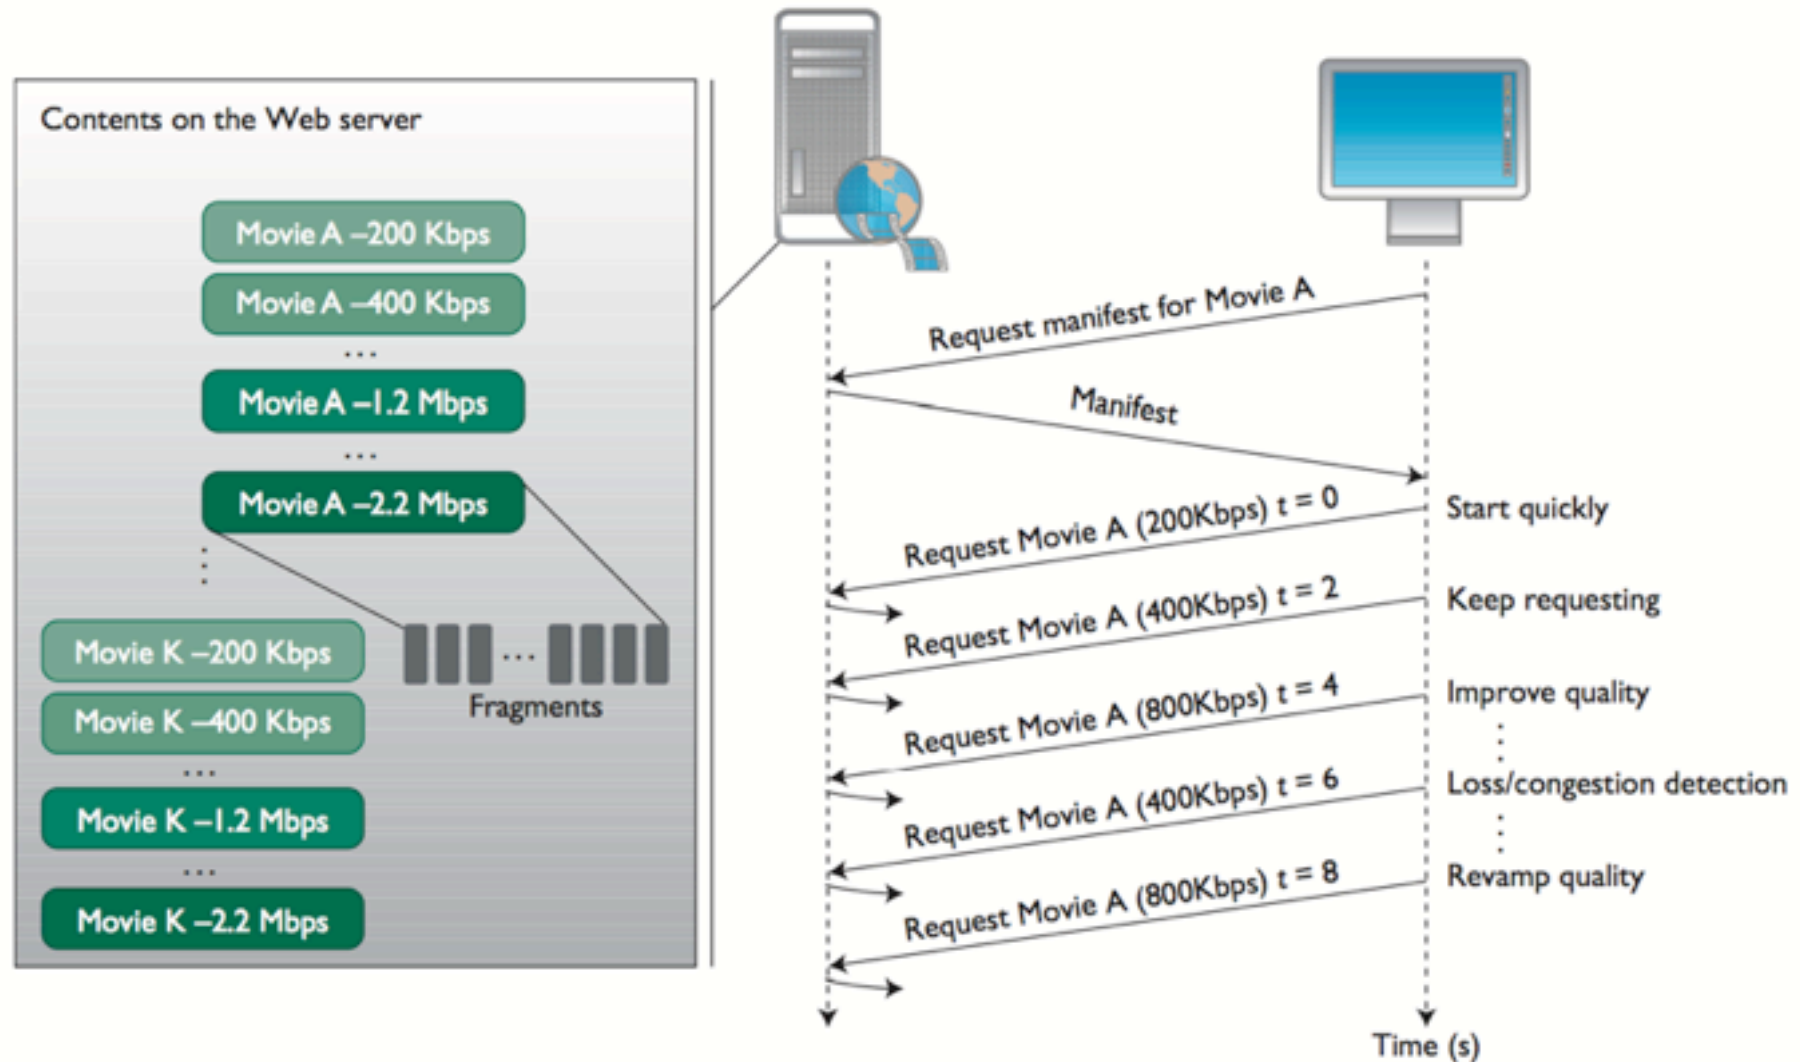
A. Begen, T. Akgul, M. Baugher: Watching Video over the Web,  
Part I: Streaming Protocols, *IEEE Internet Computing*,  
March/April 2011

I. Sodagar: The MPEG-DASH Standard for Multimedia Streaming  
Over the Internet, *IEEE Multimedia*, Oct/Dec 2011

# Trend: Adaptive HTTP Streaming

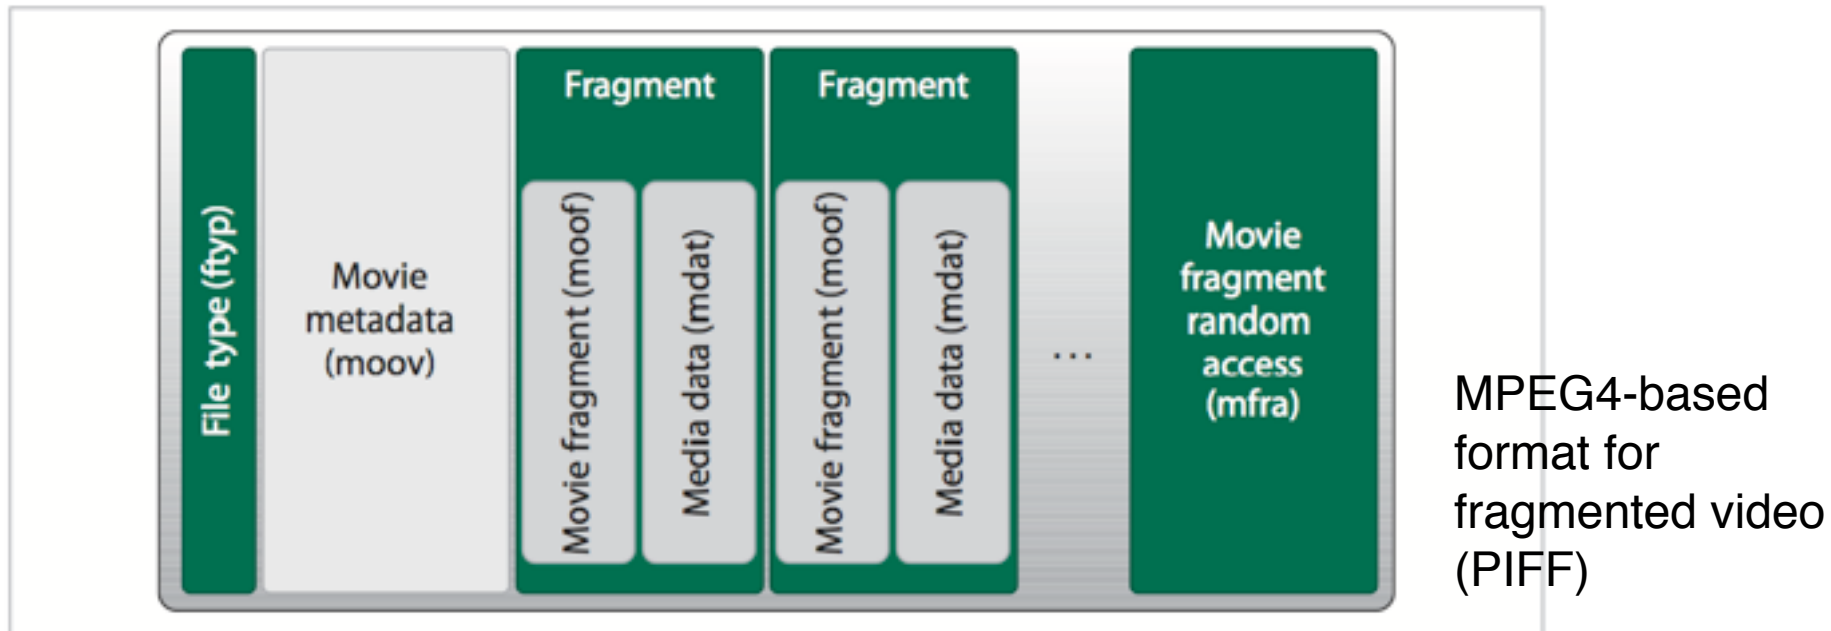
- Conventional streaming technologies (push-model) being superseded by (proprietary) adaptive HTTP-based solutions:
  - Apple HTTP Live Streaming
  - Microsoft Smooth Streaming
  - Adobe HTTP Dynamic Streaming
- Reasons for trend change
  - Firewalls, NAT devices etc. support HTTP but often not other protocols
  - HTTP caching infrastructures can be re-used for video
  - Server responsibility for bitrate adaptation is problematic for scalability
  - Better network speed allows larger chunks to be exchanged
- Basic idea: Small video file fragments
  - Fragments exist at different bit rates
  - Index file for all available fragments
  - Client requests appropriate next fragment by GET request

# Bitrate Adaptation in HTTP Adaptive Streaming



Begen/Akgul/Baugher

# Example: Video Fragmentation in Microsoft Smooth Streaming



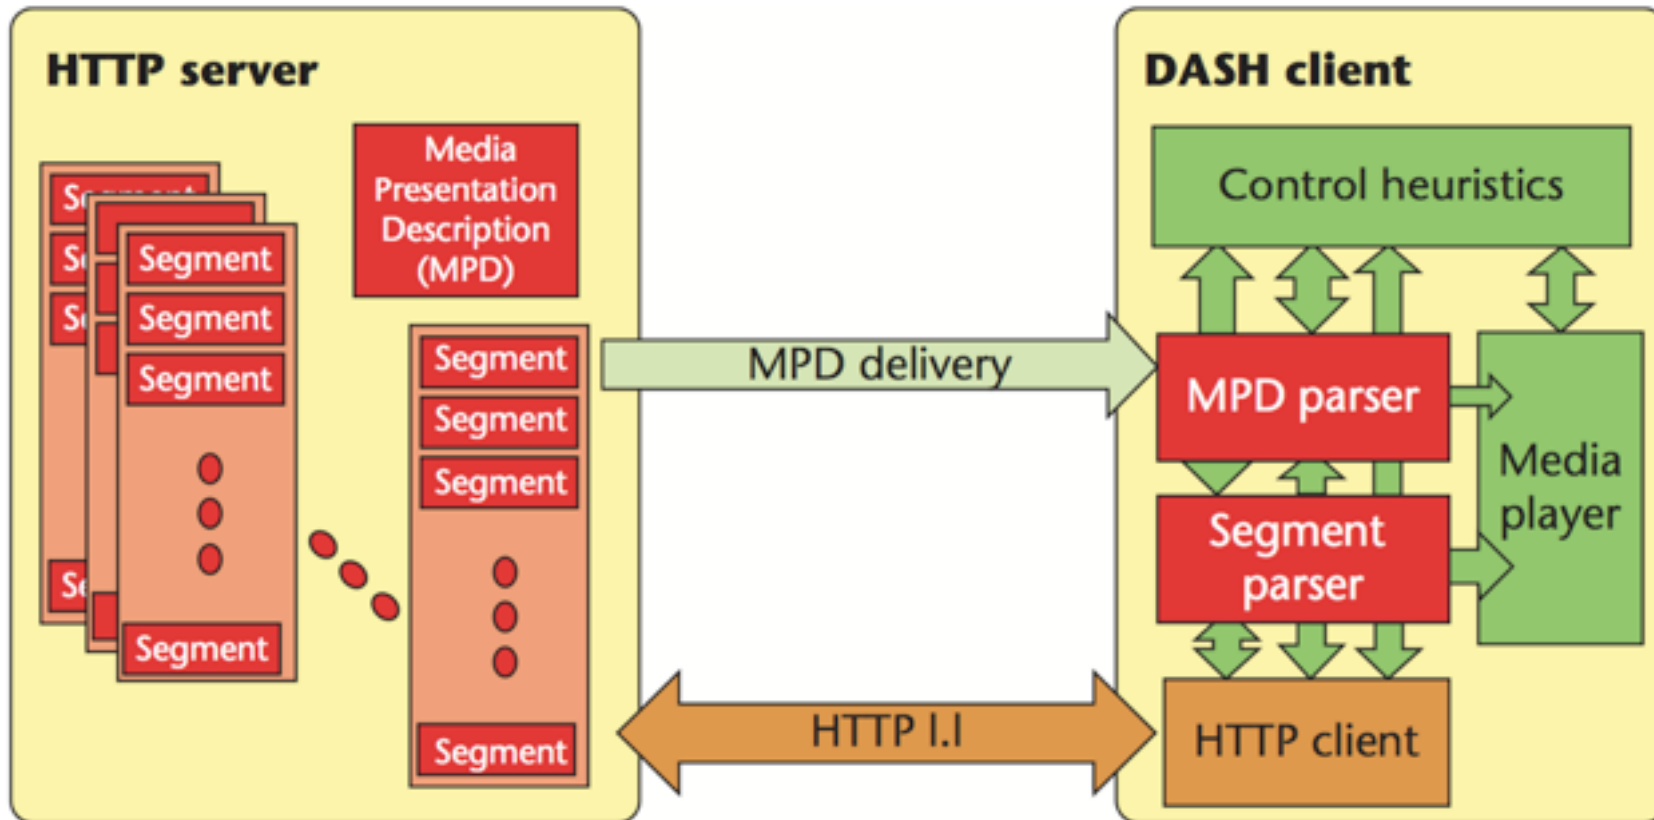
- Fragments always contain complete Groups of Pictures (GOP)
- Client (Silverlight-based Web player) sends requests which contain:
  - requested bitrate (version of the video)
  - requested time offset for fragment

```
GET/example.host/myvideo.ism/QualityLevels(64000)/  
Fragments(video=150324) HTTP/1.1
```

[Begen/Akgul/Baugher, based on alexzambelli.com/blog/2009/02/10/smooth-streaming-architecture/](#)

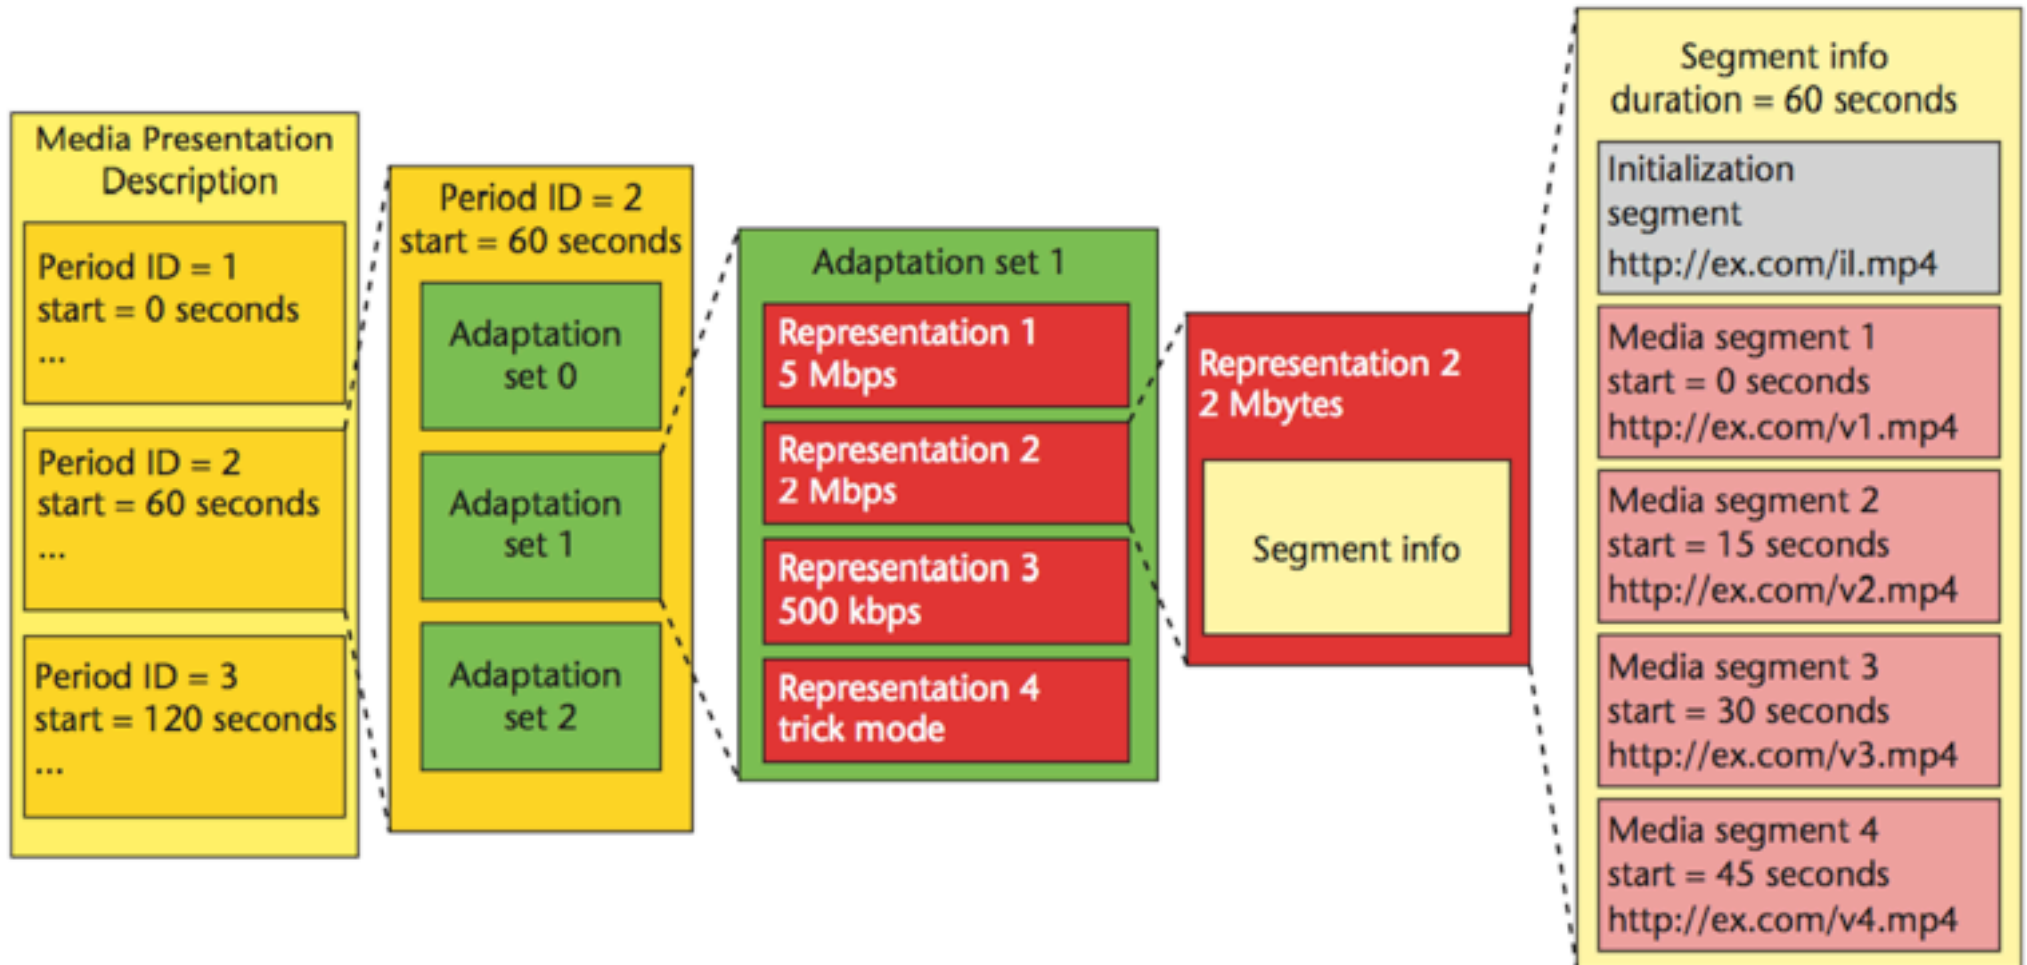
# MPEG-DASH Standard

- MPEG call and selection process for HTTP streaming standard (2009–2011)
- ISO/IEC standard (ISO 23009) April 2012



Sodagar

# Multimedia Presentation Description in DASH




Sodagar



# MPEG-4 Scalable Video Coding (SVC)

- SVC: Splitting of raw video data into multiple streams
  - MPEG/ITU-Standard 2007 (H.264/AVC Extension)
  - subsets of the data, self-contained
  - subsets can be combined to achieve better quality
  - dimensions: temporal resolution (frame rate), spatial resolution, SNR (loss)
- DASH is compatible with SVC
- Strategical question for player:
  - Download additional quality information for current segment, or
  - Download future segment (in lower quality)?
- See: Andelin et al., Quality selection for Dynamic Adaptive Streaming over HTTP with Scalable Video Coding, 3rd Multimedia Systems Conference, 2012



[www.imec.be](http://www.imec.be)

# MPEG-4 Multiple View Coding (MVC)

- SVC: Multiple camera images in a multiplexed video stream
  - MPEG/ITU-Standard 2008 (H.264/AVC Extension)
- Main application area: Stereoscopic (3D) video
  - Generalizable to multiple camera positions
- DASH is compatible with MVC
  - Server announces availability in MPD
  - Client adapts requests to the current needs

# 10 Streaming Architectures

10.1 High-Level Streaming Architecture

10.2 Push-Model Streaming: Real-Time Data Transport

10.3 Pull-Model Streaming: Progressive Download

10.4 Scalability and Multicast

10.5 Streaming Market

Literature:


David Austerberry: The Technology of Video & Audio Streaming,  
Focal Press 2002

Gregory C. Demetriades: Streaming Media, Wiley 2003

Xueyan Tang et al.: Web Content Delivery, Springer 2005

# Traditional Solution: IP Multicast

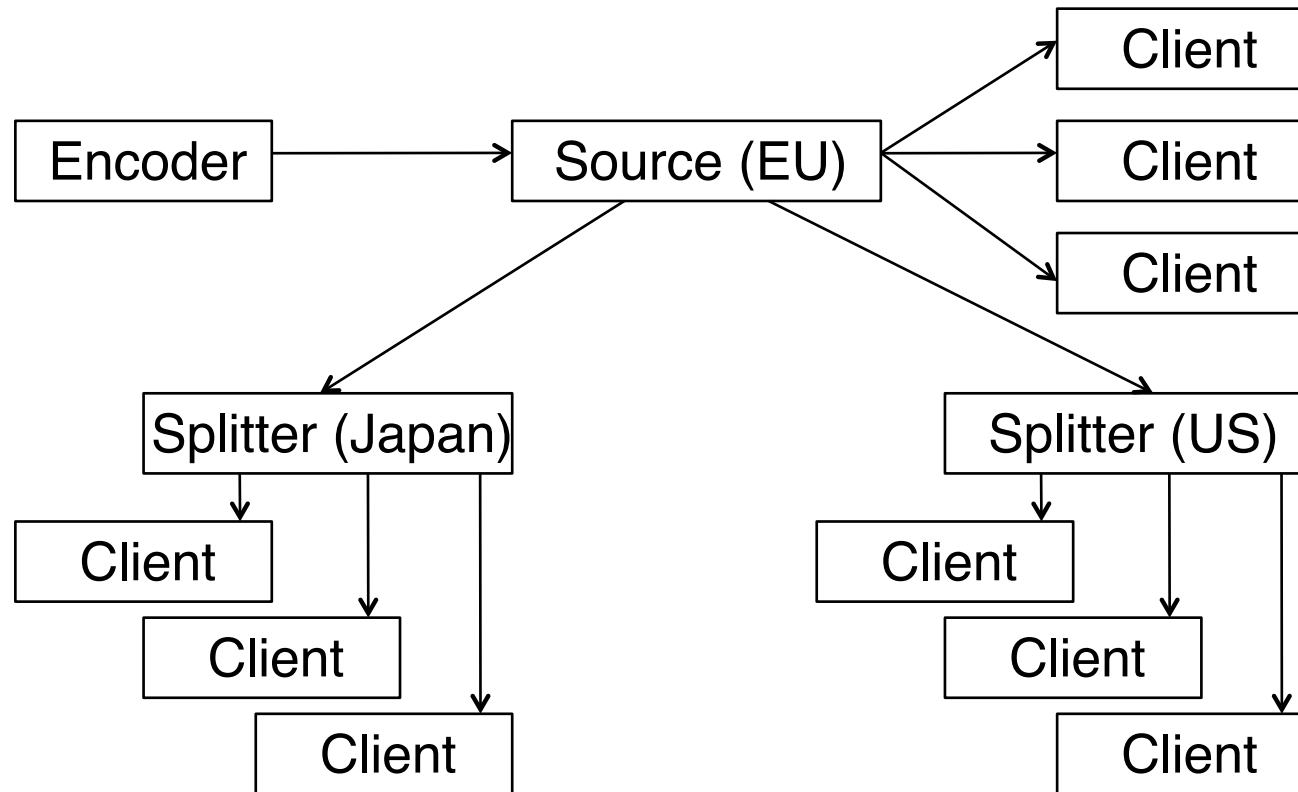
- Multicast relatively easy to integrate in routers
- IP versions 4 and 6 provide special multicast addresses



- Reliable multicast: e.g. “Mbone” overlay network, 1992
- Multicast still rarely used in today’s Internet
- Problems: Charging and access control

# Splitting

- Video servers are limited in capacity
- Assuming clients at spatially distant locations
  - Intermediate, forwarding server is useful: “splitter”



# Content Delivery Networks (CDN)

- Serve content closer to to the user
  - “edge serving”
- Main components of CDN:
  - Smart routing
  - Edge caching of content
  - Proxy serving
  - Splitting of live webcasts




Figure 15.3 Serving from the edge with a CDN.

Figure from Austerberry 2002

# Content Delivery Networks

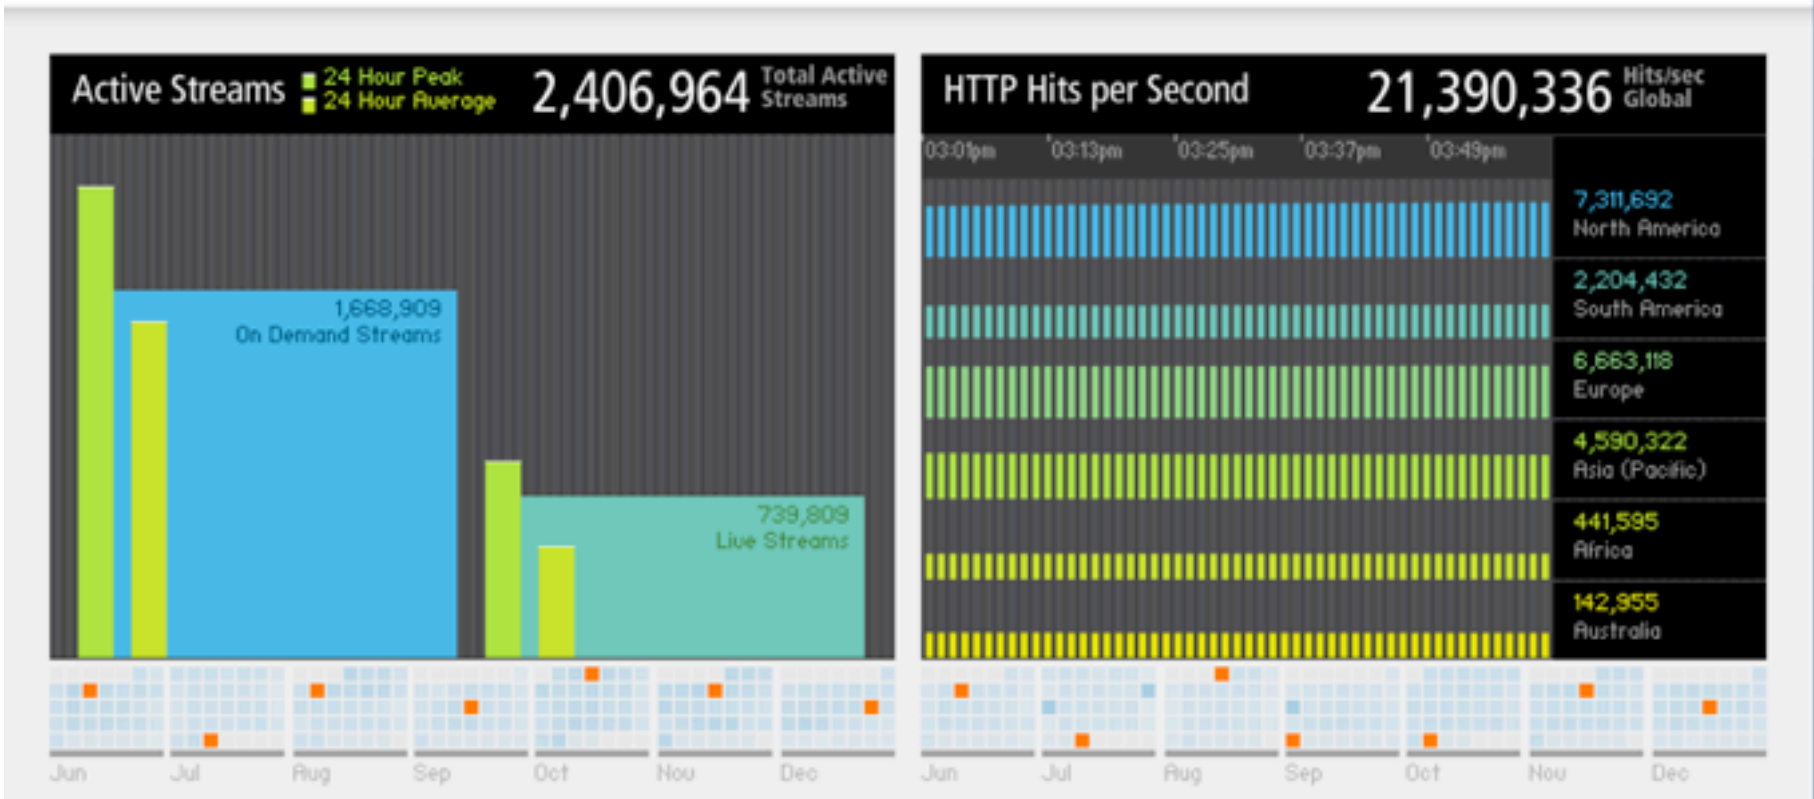
- “Overlay server infrastructure”
  - Network of centrally operated Web and streaming servers
  - Geographically distributed, present in main ISP networks
  - Flexibly used for content from various sources
- Content delivery as a service
  - Content delivery service provider owns server infrastructure
  - Content owner pays for having content delivered to customers
- Examples:
  - Akamai (delivers 20% of all Web traffic according to akamai.com, 2012)
    - » Runs 95,000+ servers in 71 countries
  - InterNap CDN Services
    - » Traditionally specialized in media streaming (acquisition of VitalStream)
  - Amazon CloudFront
- Streaming of a/v media (e.g. movies) gains strategic importance

# Example: Visualizing Akamai

Akamai handles 20% of the world's total Web traffic, providing a unique view into what's happening on the Web - what events are generating traffic, how much, from where, and why. Bookmark this page to get a feel for the world's online behavior at any given moment - how much rich media is on the move, the sheer volume of data in play, the number and concentration of worldwide visitors, and average connection speeds worldwide.

Visualizing Akamai

- [Return to Visualizing the Internet](#)
- [Methodology and Data Collection](#)





# Key Problems in CDNs


- Replica placement:
  - Where to place copies of web sites or other content
  - Problem is in general NP-hard (Karlsson, Karamolis, 2004)
  - Replica placement algorithms (RPA) achieve a suboptimal solution within reasonable time frame
  - Global information is difficult or costly to get - RPA uses local information mostly
  - CDN providers typically try to observe global network performance to some extent
- Request routing:
  - Mechanism and policy of redirecting client requests to a suitable server containing the requested content
  - Redirection algorithm: Decides what node to direct a client request to
  - Redirection mechanism: Way of redirecting the request (client, network)

# Streaming Media in CDN

- General idea: Local proxy caching. But: ...
  - Huge size (1 KB vs. 100 MB)
    - » To cache only portions of the original?
  - Intensive bandwidth usage
    - » Minimizing bandwidth consumption as primary consideration
  - High interactivity
    - » E.g. premature termination is frequent (Chen et al. 04: approx. 90 %)
  - However: Media content is rather static (compared to Web pages)
- Caching algorithms
  - Different for homogeneous and heterogeneous clients (in bandwidth/format)
  - Sliding interval caching: sequential access, mainly effective for similar requests in short time period
  - Prefix caching: Saves time to load remaining parts
  - Segment caching: Generalization of prefix caching to support fast forward
  - Rate-split caching: Lower layer from original server, higher layer from proxy
  - Co-operative proxy caching (e.g. Acharya/Smith: MiddleMan)

# Case Study: NETFLIX

- Netflix
  - North America market leader for online movies
  - 23 mio subscribers
  - HD video, up to 3.6 Mbps
  - Single largest source of Internet traffic in US (29.7% peak downstream)
- Network architecture
  - netflix.com hosted on own data centers (only registration, payment)
  - Main servers (e.g. movies.netflix.com) hosted by Amazon cloud services
  - Three CDNs used in parallel: Akamai, LimeLight, Level-3
- Player based on Silverlight using DASH
- Dynamic CDN selection strategy influences performance, see:
  - Adhiukari et al.: Unreeling Netflix: Understanding and Improving Multi-CDN Movie Delivery, IEEE INFOCOM 2012



# 10 Streaming Architectures

10.1 High-Level Streaming Architecture

10.2 Push-Model Streaming: Real-Time Data Transport

10.3 Pull-Model Streaming: Progressive Download

10.4 Scalability and Multicast

10.5 Streaming Market

# Usage of Video Streaming

Comscore.com:


- U.S. Internet users watched 39 billion online content videos in September 2012;
- The total number of unique viewers was 181.5 million;
- This represents **85 percent** of U.S. Internet users;

Forrester Research, November 2012:

*“Video is the fastest-growing digital content category; we forecast that more than 90% of the online population will regularly watch online video by 2017. Online video audience penetrations will rise significantly as video on demand, pay per view, and catchup TV become mainstream. And with pay-TV penetrations in the EU-7 nations considerably lower than in the US, there’s more scope for paid online video to grow.”*

Source: <http://blog.wooshii.com/>

# Video Codec Usage




from [WebsiteOptimization.com](http://WebsiteOptimization.com)

RV = Real Video

WMV = Windows Media Video


# Increasing Size of Video Files



Growth  
in length and  
in quality!

Source: [WebSiteOptimization.com](http://WebSiteOptimization.com)

# Market Shares of Streaming Players



Back in 2000:  
*Real 28%*  
*Win 22%*  
*QT 4%*  
 (streaming  
 media.com)

Prediction: iTunes+QT will pass Windows Media by 2012



# What is a “Media Streaming Player”?

- Still in 2011, main meaning was: Software for PCs and mobile devices
- Google search in 2013, main hits:
  - Apple TV, Roku XD, D-Link MovieNite, WD TV Live, etc.

Multimedia Research Group (MRG):

“In 2011, worldwide unit shipments of streaming media players grew to 6.8 million units. However, demand for these products continues to increase steadily and MRG is projecting that total unit shipments will exceed 8.3 million in 2012. Despite market challenges, demand projected to remain healthy through 2016.”

<http://www.broadbandtvnews.com/2012/12/14/new-high-for-streaming-media-players/>

