

# INF3190 – Data Communication Multimedia Protocols

Carsten Griwodz

Email: [griff@ifi.uio.no](mailto:griff@ifi.uio.no)

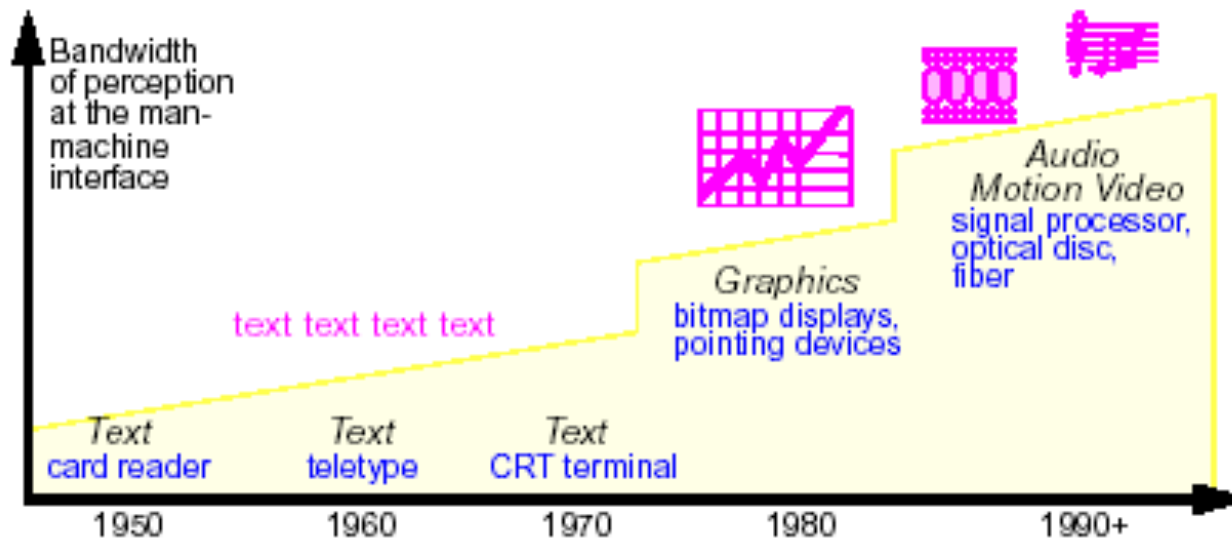


# Media

Medium: "Thing in the middle"

- here: means to distribute and present information

Media affect human computer interaction



The mantra of multimedia users

- Speaking is faster than writing
- Listening is easier than reading
- Showing is easier than describing

# Dependence of Media

- Time-independent media

- Text
- Graphics
- *Discrete* media

- Time-dependent media

- Audio
- Video
- Animation
- Multiplayer games
- *Continuous* media

- Interdependant media

- *Multi*media

- "Continuous" refers to the user's impression of the data, not necessarily to its representation

- Combined video and audio is multimedia - relations must be specified



# Continuous Media

## Fundamental characteristics

- Typically **delay sensitive**
- Often **loss tolerant**: infrequent losses cause minor glitches that can be concealed
- Antithesis of discrete media (programs, banking info, etc.), which are loss intolerant but delay tolerant

## Classes of MM applications

- Streaming stored audio and video
- Streaming live audio and video
- Interactive real-time audio and video
- Interactive real-time event-driven applications



# Multimedia in networks

## Streaming stored MM

- Clients request audio/video files from servers and pipeline reception over the network and display
- Interactive: user can control operation (pause, resume, fast forward, rewind, etc.)
- Delay: from client request until display start can be 1 to 10 seconds

## Unidirectional Real-Time

- similar to existing TV and radio stations, but delivery over the Internet
- Non-interactive, just listen/view

## Interactive Real-Time

Phone or video conference

- More stringent delay requirement than Streaming & Unidirectional because of real-time nature
- Audio: < 150 msec good, < 400 msec acceptable
- Video: < 150 msec acceptable  
[Note: higher delays are feasible, but usage patterns change **(!)**]

Games *(but also high-speed trading)*


- Role playing games: < 500 msec
- First person shooter (FPS) games: < 100 msec ***(may be too high)***
- Cloud gaming FPS: < 40 msec ***(estimated)***



Slides by Prof. Ramesh Sitaraman, UMass, Amherst (shown with permission)  
"Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-Experimental Designs", S. S. Krishnan and R. Sitaraman, ACM Internet Measurement Conference (IMC), Boston, MA, Nov 2012

Str

Viewers with better connectivity have less patience for startup delay and abandon sooner.



Uni




# High Data Volume

- Throughput
  - Higher volume than for traditional data
  - Longer transactions than for traditional data
  - Requires
    - Performance and bandwidth
    - Resource management techniques
    - Compression
  
- Typical values
  - Uncompressed video: 140 – 216 Mbit/s
  - Uncompressed audio (CD): 1.4 Mbit/s
  - Uncompressed speech: 64 Kbit/s
  - Compressed audio & video (VoD): down to 1.2 – 4 Mbit/s
  - Compressed audio & video (Conf.): down to 128 Kbit/s
  - Compressed speech: down to 6.2 Kbit/s



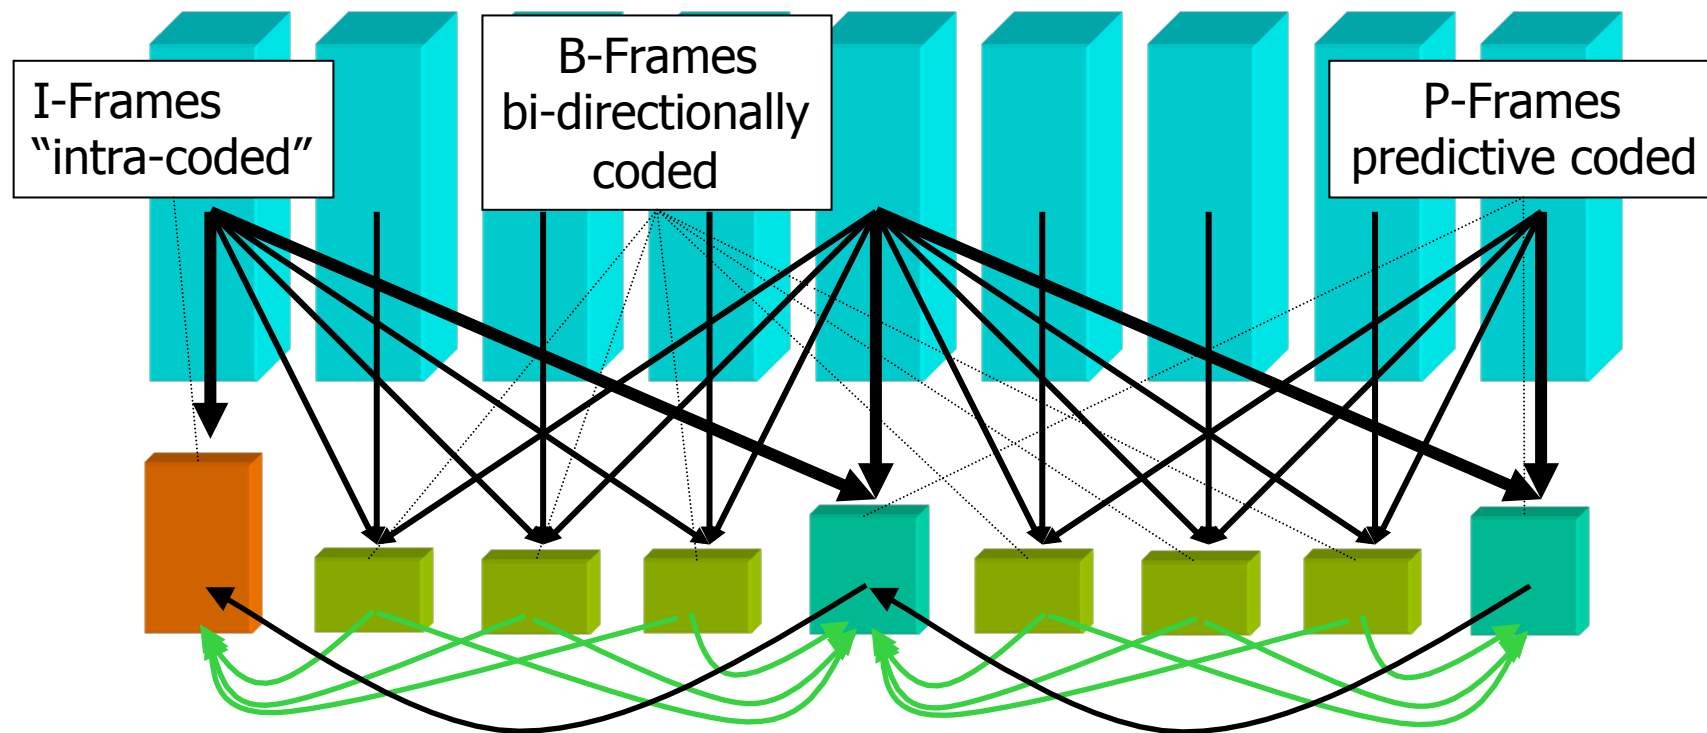
# Compression – General Requirements






# Example: MPEG-1

- International Standard: Moving Pictures Expert Group
  - Compression of audio and video for playback (1.5 Mbit/s)
  - Real-time decoding
- Sequence of I-, P-, and B-Frames



# MPEG (Moving Pictures Expert Group)

- Frames can be dropped
  - In a controlled manner
  - Frame dropping does not violate dependancies
  - Example: B-frame dropping in MPEG-1



# Quality of service - QoS

A term that is used in all kinds of contexts.  
Be careful what it means when you hear it.

In this lecture: 3 *classical*

parameters of **network QoS**:

- end-to-end delay
- packet loss
- jitter

end-to-end delay

- transmission time
- $\sum$  propagation time on link  $l$   
sum of propagation times over all links  $l$
- $\sum$  queueing time on router  $r$   
sum of queueing times at all routers' queues  $r$

packet loss

- probability of a packet to get lost
- $1 - ( \prod ( P(\text{queue at } r \text{ not full}) ) )$   
1 - product of probabilities for all  $r$  that queue at  $r$  is not full

jitter

- variance of end-to-end delay
- estimated for several packets
- reasons
  - link layer retransmissions
  - queue length variation



# Multimedia Networking

## Internet without network QoS support

- Internet applications must cope with networking problems
  - Application itself or middleware
  - "Cope with" means either "*adapt to*" or "*don't care about*"
  - "Adapt to" must deal with TCP-like service variations
  - "Don't care about" approach is considered "unfair"
  - "Don't care about" approach cannot work with TCP

## Internet with network QoS support

- Application must specify their needs
- Internet infrastructure must change – negotiation of QoS parameters
- Routers need more features
  - Keep QoS-related information
  - Identify packets as QoS-worthy or not
  - Treat packets differently keep routing consistent

- approach seemed "dead" for many years
- revival with recent Software Defined Networking (SDN) idea
- not yet mainstream again

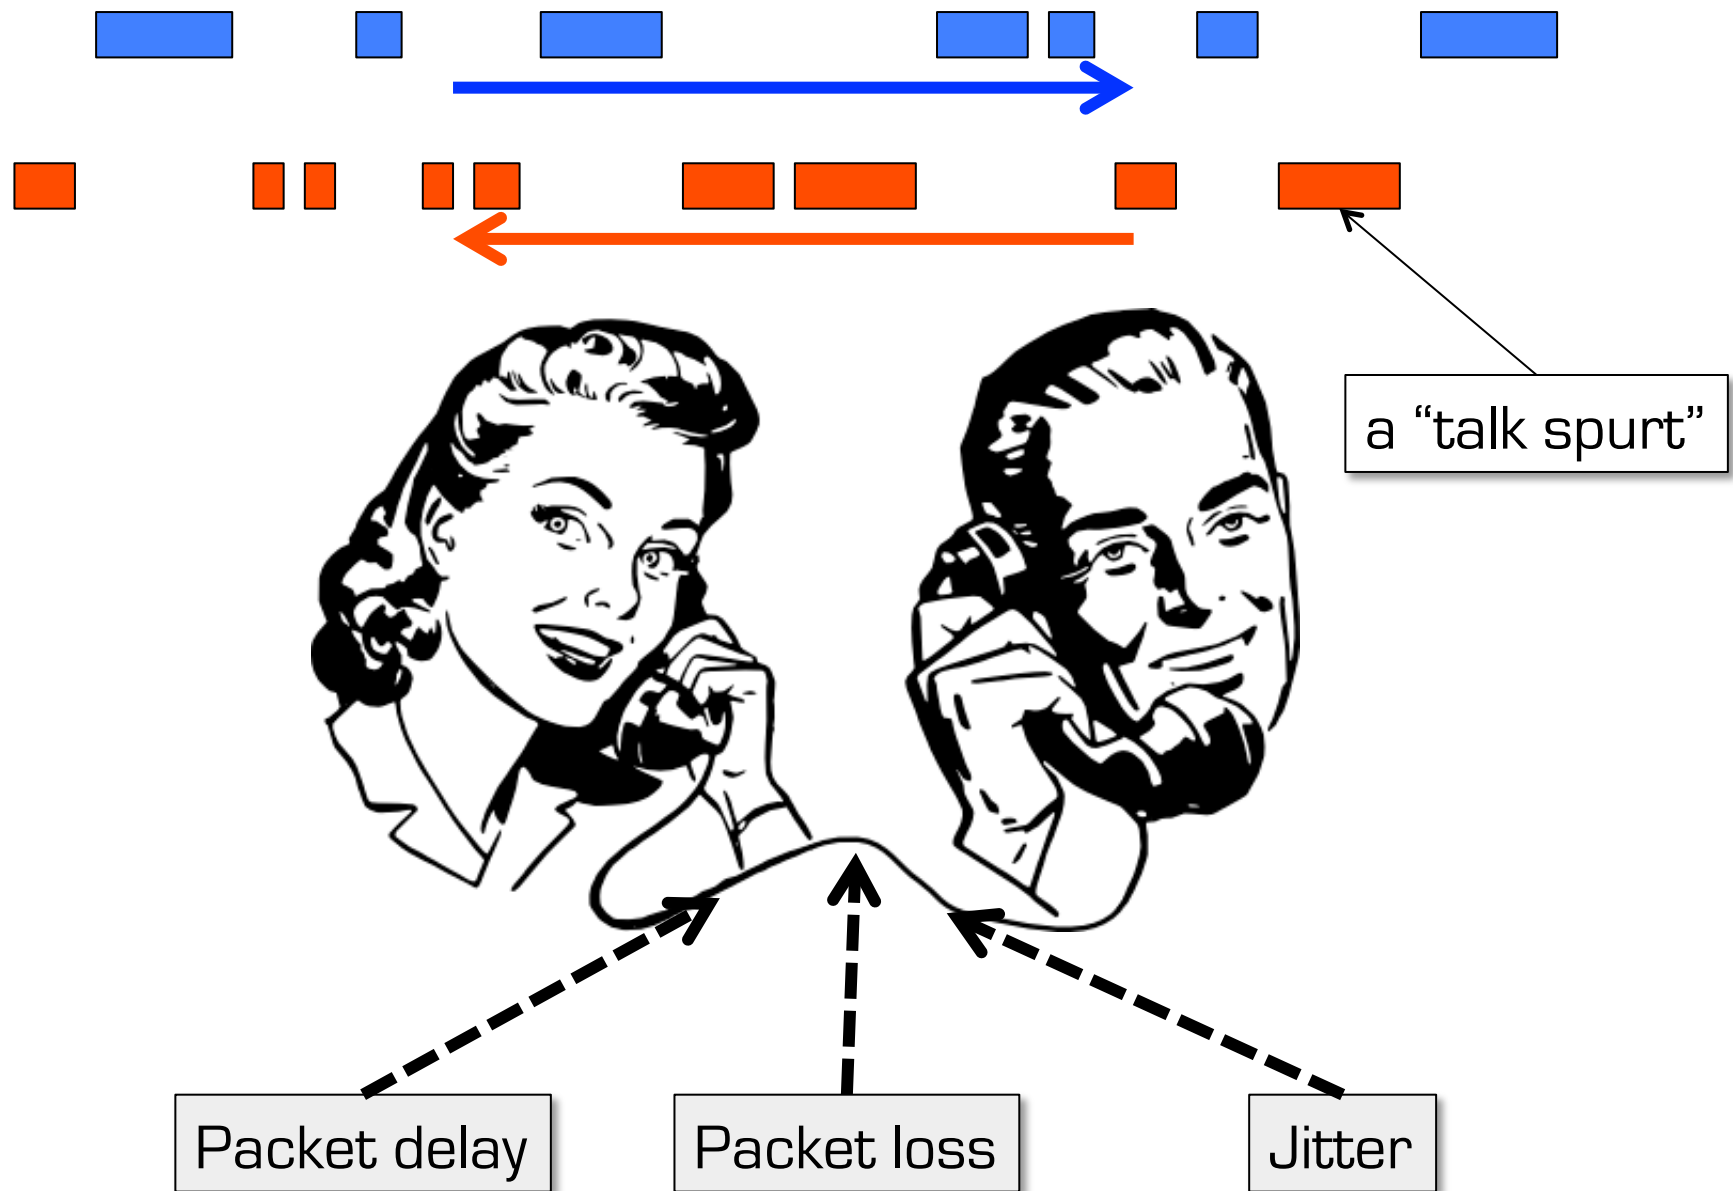


# Non-QoS Multimedia Networking

## Basics



# Streaming over best-effort networks: audio conferencing



# Streaming over best-effort networks: audio conferencing

## end-to-end delay

- end-to-end delay can seriously hinder interactivity
- smaller is always better? not true for cooperative music making!

## packet loss

- UDP segment is encapsulated in IP datagram
- datagram may overflow a router queue
- TCP can eliminate loss, but
  - retransmissions add delay
  - TCP congestion control limits transmission rate
- redundant packets can help

## delay jitter

- consider two consecutive packets in talk spurt
- initial spacing is 20 msec, but spacing at receiver can be more or less than 20 msec


## removing jitter

- sequence numbers
- timestamps
- delaying playout



# Delay compensation

All techniques rely on *Prediction*



no delay compensation  
in this example

## Teleconferencing

- no known technique: cannot predict what people will say

## For on-demand

- usually content is consumed linearly, prefetching is easy, limited only by resources and legal constraints

## For event-based multimedia

- predict future movement
- perform audiovisual rendering based on prediction
- compensate for errors in next prediction
- used in computer games and other distributed simulations, head- and gesture tracking, mouse or joystick inputs



# Jitter compensation

Receiver attempts to playout each chunk at exactly  $q$  msec after the chunk is generated

- If chunk is time stamped  $t$ , receiver plays out chunk at  $t+q$
- If chunk arrives after time  $t+q$ , receiver discards it

Sequence numbers not necessary

Strategy allows for lost packets

Tradeoff for  $q$ :

- large  $q$ : less packet drop/loss (better audio quality)
- small  $q$ : better interactive experience



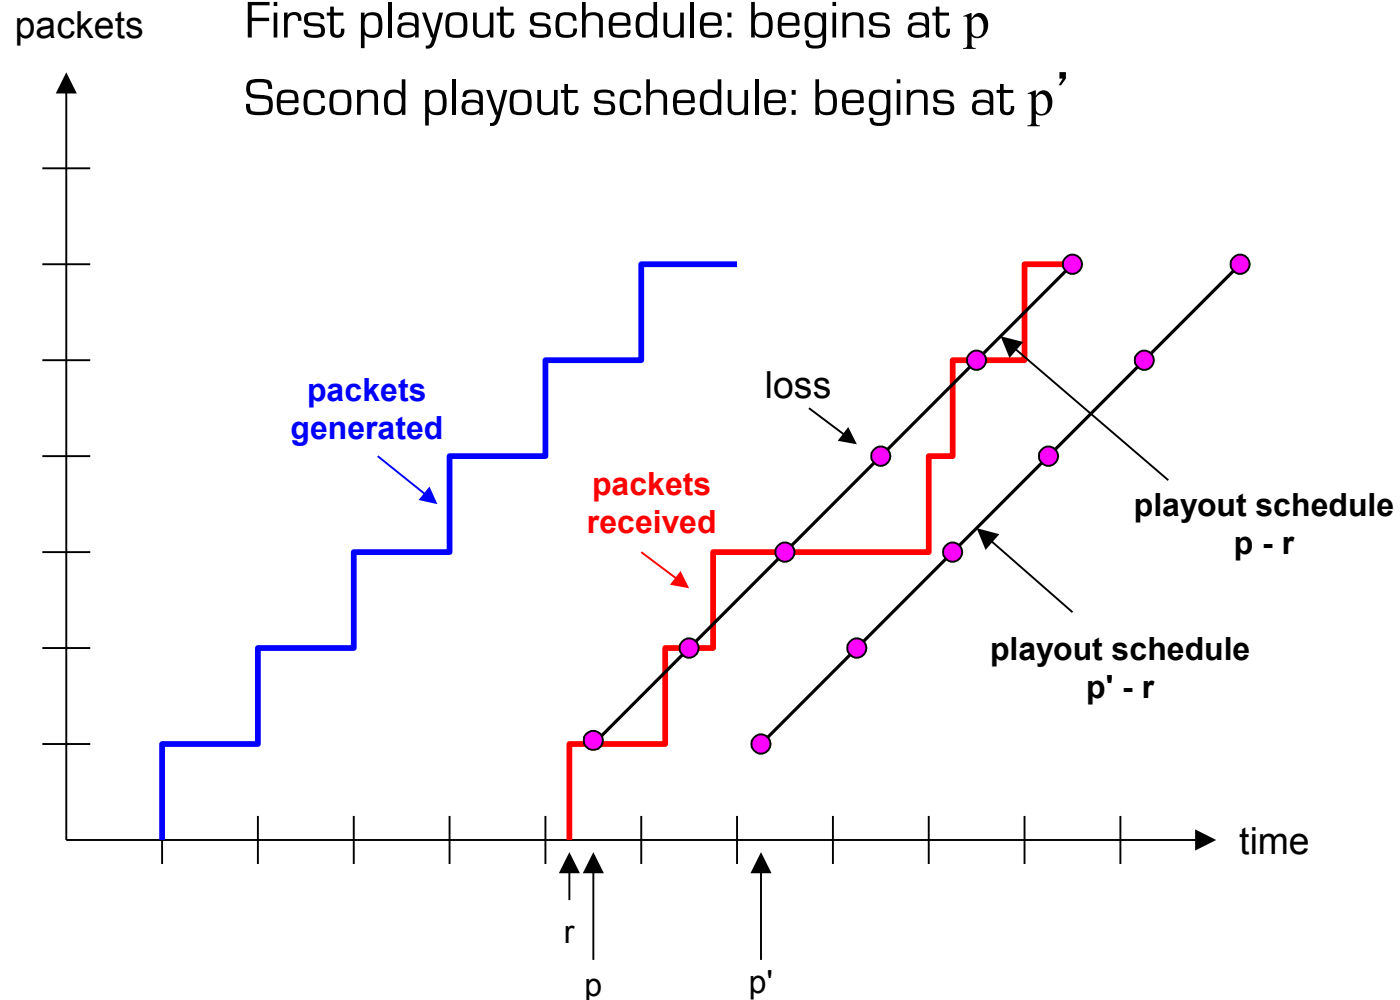
# Jitter compensation

Sender generates packets every 20 msec during talk spurt

First packet received at time  $r$

First playout schedule: begins at  $p$

Second playout schedule: begins at  $p'$



# Jitter compensation: Adaptive playout delay

Estimate network delay and adjust playout delay at the beginning of each talk spurt  
Silent periods are compressed and elongated as needed

Chunks *still* played out every 20 msec during talk spurt

$t_i$  = timestamp of the  $i$ th packet

$r_i$  = the time packet  $i$  is received by receiver

$p_i$  = the time packet  $i$  is played at receiver

$r_i - t_i$  = network delay for  $i$ th packet

$d_i$  = estimate of average network delay after receiving  $i$ th packet

Dynamic estimate of average delay at receiver:

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

where  $u$  is a fixed constant [e.g.,  $u = .01$ ]



# Jitter compensation: Adaptive playout delay

Also useful to estimate the average deviation of the delay,  $v_i$ :

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

Deviation: How strongly does the queue length change?

The estimates  $d_i$  and  $v_i$  are calculated for every received packet, although they are only used at the beginning of a talk spurt

For first packet in talk spurt, playout time is:

$$p_i = t_i + d_i + Kv_i$$

application chooses the safety margin  $Kv_i$

where  $K$  is a positive constant

Playout delay is  $q_i = p_i - t_i = d_i + Kv_i$

for this and **all other** packets in **this** talk spurt



# Jitter compensation: Adaptive playout delay

## How to determine whether a packet is the first in a talkspurt?

- If there were never loss, receiver could simply look at the successive time stamps
  - Difference of successive stamps  $> 20$  msec, talk spurt begins
- But because loss is possible, receiver must look at both time stamps and sequence numbers
  - Difference of successive stamps  $> 20$  msec and sequence numbers without gaps, talk spurt begins



# Loss compensation

## Basic assumption

- we have very little time to loose in audio conferencing
- every packet carries dozens of samples
- adding several packets delay for complex schemes is not viable

## forward error correction (FEC): simple scheme


- for every group of  $n$  chunks create a redundant chunk by exclusive OR-ing the  $n$  original chunks
- send out  $n+1$  chunks, increasing the bandwidth by factor  $1/n$ .
- can reconstruct the original  $n$  chunks if there is at most one lost chunk from the  $n+1$  chunks
- Playout of first packet has to wait for arrival of  $(n+1)^{\text{st}}$  packet
- Playout delay needs to be fixed to the time to receive all  $n+1$  packets
- Tradeoff:
  - increase  $n$ , less bandwidth waste
  - increase  $n$ , longer playout delay
  - increase  $n$ , higher probability that 2 or more chunks will be lost



# Loss compensation

## 2nd FEC scheme

- “piggyback lower quality stream”
- send lower resolution audio stream as the redundant information
- for example, nominal stream PCM at 64 kbps and redundant stream GSM at 13 kbps.




- Sender creates packet by taking the  $n$ th chunk from nominal stream and appending to it the  $(n-1)$ st chunk from redundant stream.
- Whenever there is non-consecutive loss, the receiver can conceal the loss.
- Only two packets need to be received before playback
- Can also append  $(n-1)$ st and  $(n-2)$ nd low-bit rate chunk

# Loss compensation

## Interleaving

- chunks are broken up into smaller units
  - for example, 4 5 msec units per chunk
  - interleave the chunks as shown in diagram
  - packet now contains small units from different chunks
- Reassemble chunks at receiver
  - if one packet is lost, still have most of every chunk





# Loss compensation

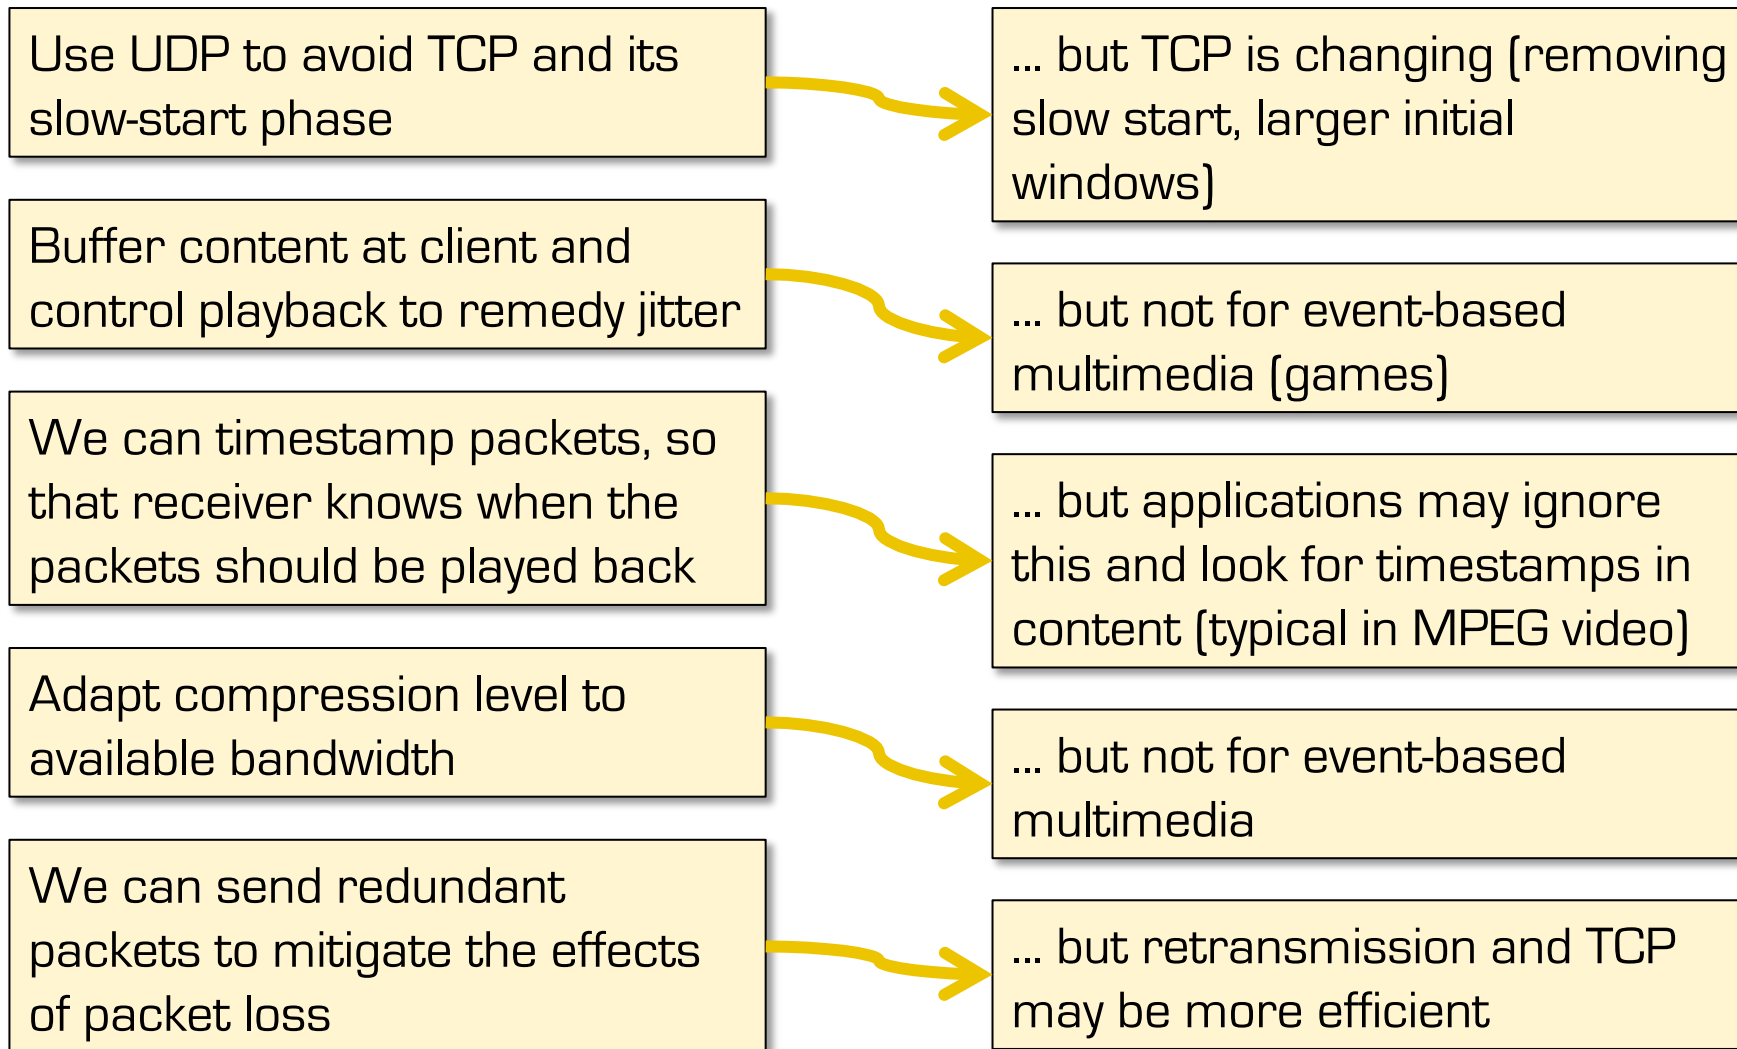
## Receiver-based repair of damaged audio streams

- produce a replacement for a lost packet that is similar to the original
- can give good performance for low loss rates and small packets (4-40 msec)
- simplest: repetition
- more complicated: interpolation



# Making the best of best effort

## Mitigating the impact of “best-effort” in the Internet




# Streaming from Web server (1)

- Audio and video files stored in Web servers

## naïve approach

- browser requests file with HTTP request message
- Web server sends file in HTTP response message
- content-type header line indicates an audio/video encoding
- browser launches media player, and passes file to media player
- media player renders file

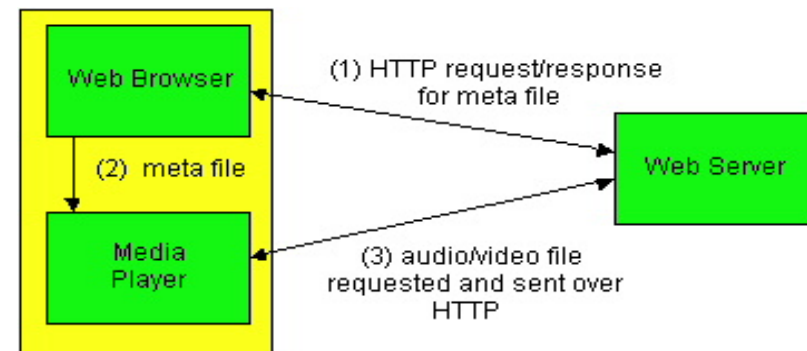


- Major drawback: media player interacts with server through intermediary of a Web browser

# Streaming from Web server (2)

## Alternative: set up connection between server and player

- Web browser requests and receives a **meta file** (a file describing the object) instead of receiving the file itself;
- Content-type header indicates specific audio/video application
- Browser launches media player and passes it the meta file
- Player sets up a TCP connection with server and sends HTTP request.

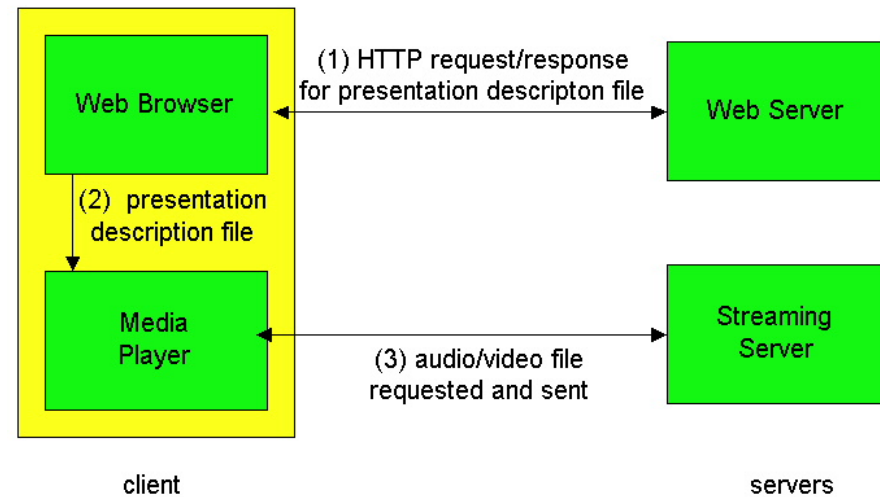


## Some concerns:

- Media player communicates over HTTP, which is not designed with pause, ff, rwnd commands
- May want to stream over UDP

# Streaming from a streaming server

- This architecture allows for non-HTTP protocol between server and media player
- Can also use UDP instead of TCP.



# Non-QoS Multimedia Networking

Application Layer Framing &  
Integrated Layer Processing



# Multimedia Content Processing

- Problem: optimize transport of multimedia content
  
- It is application-dependent and specific
  - Application-layer processing has high overhead
  - Application processes data as it arrives from the network
  
- Impact of lost and mis-ordered data
  - either:** Transport layer tries to recover from error
    - Prevents delivery of data to application
    - Prevents immediate processing as data arrives
    - Application must stop processing
  - or:** Transport layer ignores error
    - Application experiences processing failures
    - Application must stop processing



# Application Level Framing

[Clark/Tennenhouse 1990]

## Give application more control

- Application understands meaning of data
- Application should have the option of dealing with a lost data
  - Reconstitute the lost data (recompute/buffer by applications)
  - Ignore the lost data

## Application level framing

- Application breaks the data into suitable aggregates
  - Application Data Units (ADUs)
- Lower layers preserve the ADU frame boundaries
- ADU takes place of packet as the unit of manipulation





# ALF: Application Data Units

ADUs become the unit of error recovery

- Should be upper bounded
  - loss of large ADUs is more difficult to fix
- Lower bounded
  - application semantics define smallest sensible unit
  - small ADUs mean larger protocol overhead
- Segmentation/reassembly
  - try to avoid
  - multi-TPDU ADU is wasted when one packet is lost

ADU “name”


- Sender computes a name for each ADU (e.g. sequence number)
- Receiver uses name to understand its place in the sequence of ADUs
- Receiver can process ADUs out of order



# Integrated Layer Processing


Layered engineering is not fundamental

- Assignment of functions to layers in OSI is not following fundamental principles
- Specific application may work better with different layering of functions or no layering at all
- Sequential processing through each layer
  - Not an efficient engineering
  - Processing all functions at once saves computing power



Integrated Layer Processing

- Vertical integration
- Performing all the manipulation steps in one or two integrated processing loops, instead of serially



# Integrated Layer Processing

- Ordering constraint
  - Data manipulation can only be done after specific control steps
  - Data manipulation can only be done once the data unit is in order
  - Layered multiplexing (extract the data before it can be demultiplexed)
- Minimize inter-layer ordering constraints imposed on implementors
  - Implementors know best which data must be ordered
- Drawback: complex design due to fully customized implementation



# Non-QoS Multimedia Networking

RTP – Real-Time Transfer Protocol



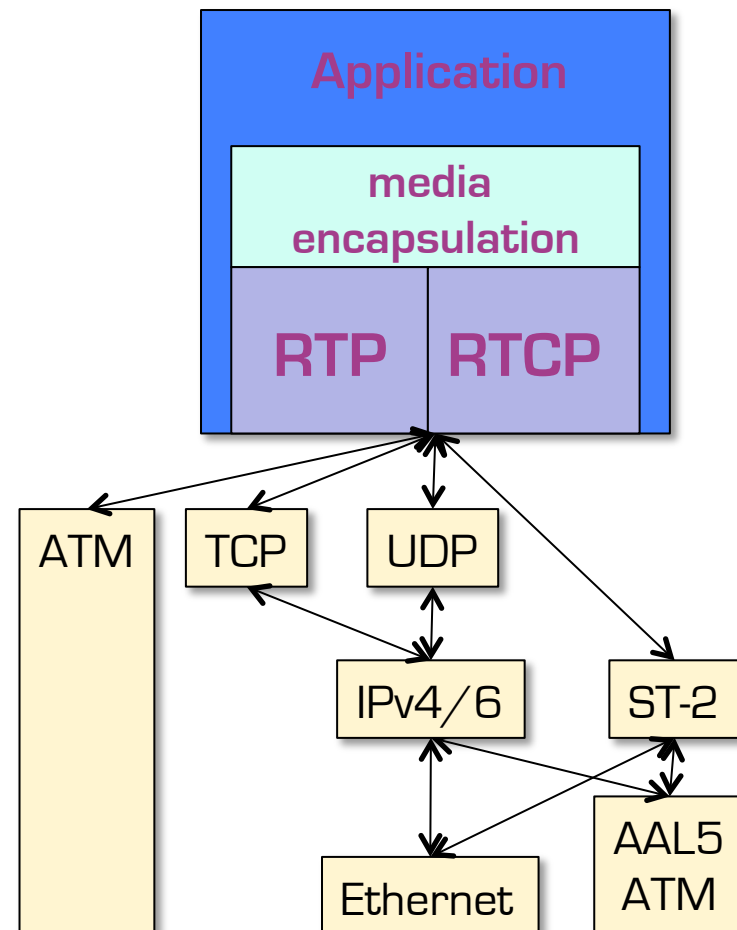
# Real-time Transport Protocol (RTP)

- Real-time Transport Protocol (RTP)
  - RFC 3550 (replaces RFC 1889)
  - Designed for requirements of real-time data transport
  - **NOT** real-time
  - **NOT** a transport protocol
- Two Components
  - RTP Data Transfer Protocol (RTP)
  - RTP Control Protocol (RTCP)
- Provides end-to-end transport functions
  - Scalable in multicast scenarios
  - Media independent
  - Mixer and translator support
  - RTCP for QoS feedback and session information



# Real-time Transport Protocol (RTP)

- No premise on underlying resources
  - layered above transport protocol
  - no reservation / guarantees
- Integrated with applications
- RTP follows principles of
  - Application Level Framing and
  - Integrated Layer Processing



# WebRTC / rtcweb

In the last 5 years,  
RTP was nearly killed by HTTP Adaptive Streaming (HAS)  
*but Google brought it back*

## WebRTC

- free, open project
- adopted by Google, later Mozilla Foundation, Opera, ...
- Real-Time Communications (RTC) for browsers and mobile devices through HTML5 and JavaScript APIs

## rtcweb

- Real Time Collaboration on the World Wide Web
- standardize infrastructure for real-time communication in Web browsers
- IETF: formats and protocols
- W3C: APIs for control



# RTP

- RTP services are
  - sequencing
  - synchronization
  - payload identification
  - QoS feedback and session information
- RTP supports
  - multicast in a scalable way
  - generic real-time media and changing codecs on the fly
  - mixers and translators to adapt to bandwidth limitations
  - encryption
- RTP is **not** designed for
  - reliable delivery
  - QoS provision or reservation





# RTP Functions

- RTP with RTCP provides
  - support for transmission of real-time data
  - over multicast or unicast network services
  
- Functional basis for this
  - Loss detection – sequence numbering
  - Determination of media encoding
  - Synchronization – timing
  - Framing - “guidelines” in payload format definitions
  - Encryption
  - Unicast and multicast support
  - Support for stream “translation” and “mixing” (SSRC; CSRC)



to be continued

