L6 – RTP, RTCP and RTSP

by
T.S.R.K. Prasad

EA C451 Internetworking Technologies

21/08/2012

References / Acknowledgements

http://www.cs.columbia.edu/~hgs/rtp/

(Parts of the presentation borrowed from RTP Overview presentation by Prof. Henning Schulzrinne)

Sec 5.4: Transport for Real-Time Applications (RTP), [Peterson]

Sec 7.5: RTP, [Farrel]

Sec 7.4.1 – 7.4.2: RTP and RTCP, [Kurose]

Self-study:

- Sec 7.1 – 7.3: Transport Over IP, [Farrel]

Optional Readings

[rfc3550] - RTP: A Transport Protocol for Real-Time Applications

[rfc2326] - Real Time Streaming Protocol (RTSP)

[rfc3551] - RTP Profile for Audio and Video Conferences with Minimal Control

[rfc4566] – SDP: Session Description Protocol

Presentation Overview

SDP

RTSP

RTP and RTCP

Topic Overview

- RTP and RTCP
  - **Motivation**
  - Architecture
  - Packet Formats
  - Mixers and Translators
  - Summary

Fast-forward to the Year 2021

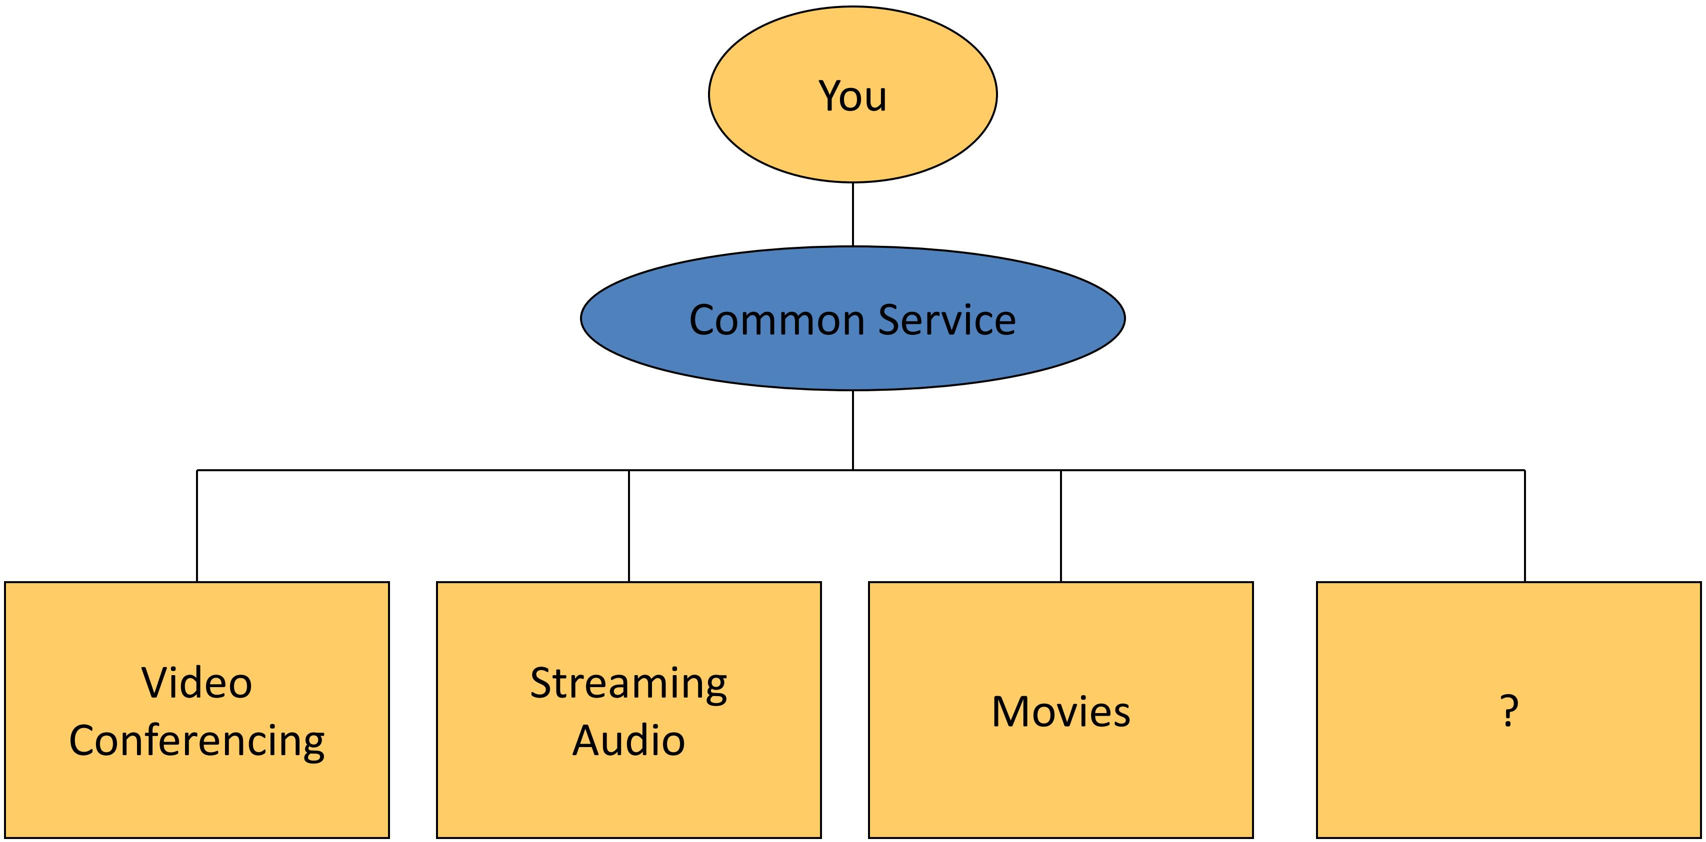

Two Goals of RTP's Common Service

- General enough to be truly "common"
  - Who knows what applications are coming?
  - Throughout history, communication has changed:
    - Oral (traditions passed between generations)
    - Written
    - Visual
- Specific enough to actually be useful

A Common RTP Scenario

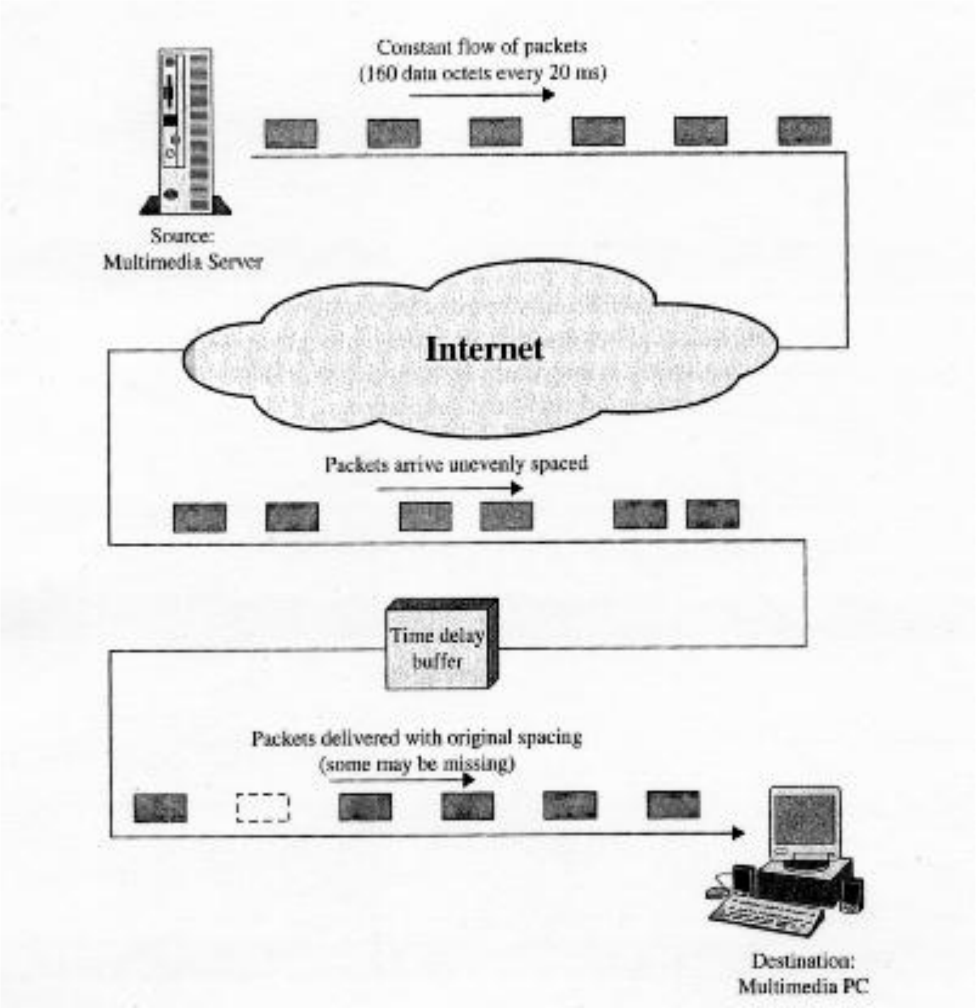

RTP Requirements - Timing

- Timing
  - Time-stamping for buffered playback
    - to minimize jitter
  - Synchronization of multiple streams
  - Dynamic frame boundaries
    - Video: frame length varies due to compression
    - Audio: “talkspurts”

RTP Requirements - Network

- Network issues
  - Dealing with packet loss
  - Dealing with congestion
    - Even with multicast
  - Bandwidth utilization
    - Minimize header bits

RTP Requirements - Multimedia

- Miscellaneous
  - Interoperability
    - Encoding
    - Compression
  - ID of source
    - To whom am I listening?
    - Useful especially in video-conferencing

RTP Requirements Summary

- This is *not* TCP!
  - Who cares if we lose a packet or two?

    (Not us!)
  - Who cares if we have jitter?

    (We do!)
- Calls for a different protocol...

Topic Overview

- RTP and RTCP
  - Summary
  - Mixers and Translators
  - Packet Formats
  - **Architecture**
  - Motivation

RTP Big Picture

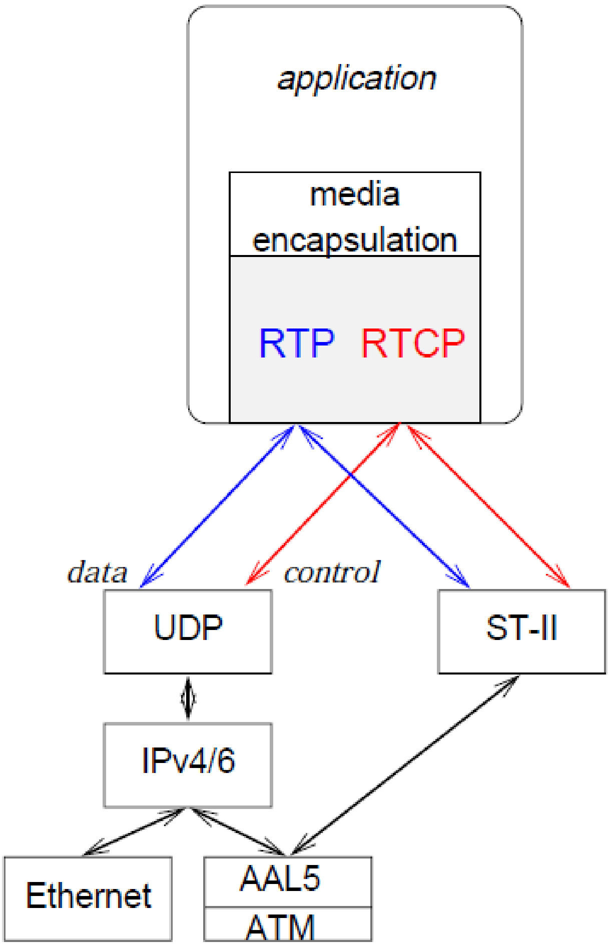

- Network layer and data link layer independence
- Two parts to protocol – RTP and RTCP
- Application independence

RTP Architecture

“ALF” and “ILP”

- Application-level framing:
  - The application best knows its own needs
  - May not ask for retransmission, but for lower resolution
- Integrated Layer Processing
  - Tightly coupled layers
  - Keeps data presentation from being the bottleneck
    - Gives the app. access to the data ASAP!

D. Clark and D. Tennenhouse, 1990

“Architectural considerations for a new generation of protocols”

RTP: Summary

- A very thin protocol
  - Usually built into application
- No hard QoS guarantees
  - Designed for *soft* real-time apps
  - Depends on underlying network
  - Can run over ATM
- Two components:
  - Media(data) transport: RTP
  - Control: RTCP

RTP Concepts

- Port numbers for both RTP and RTCP
- Participant IP addresses
  - Strength is *multicast*
- Relays
  - Mixers
  - Translators
  - (More about these two later)

RTP Protocol Stack

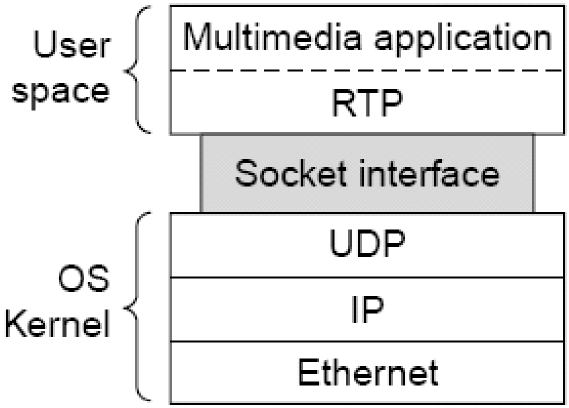

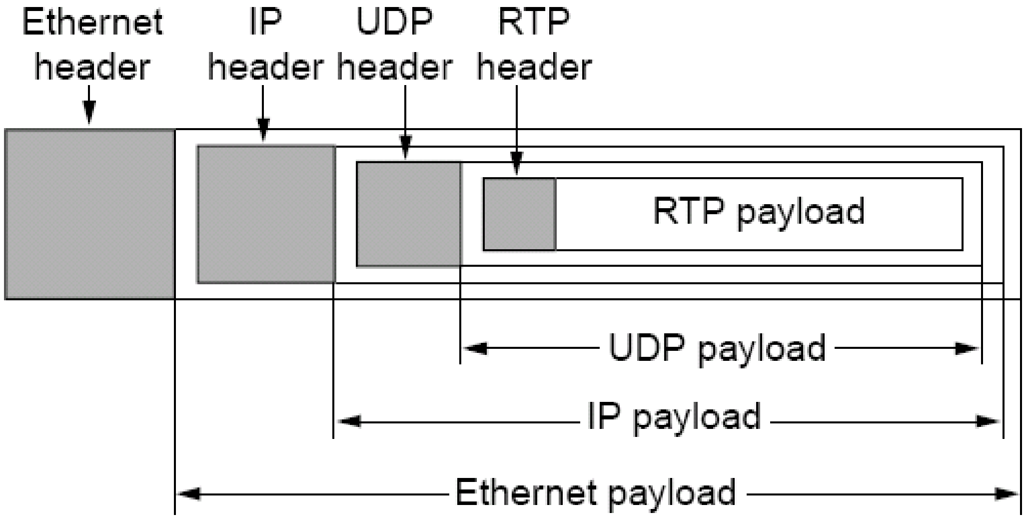

RTP User Applications

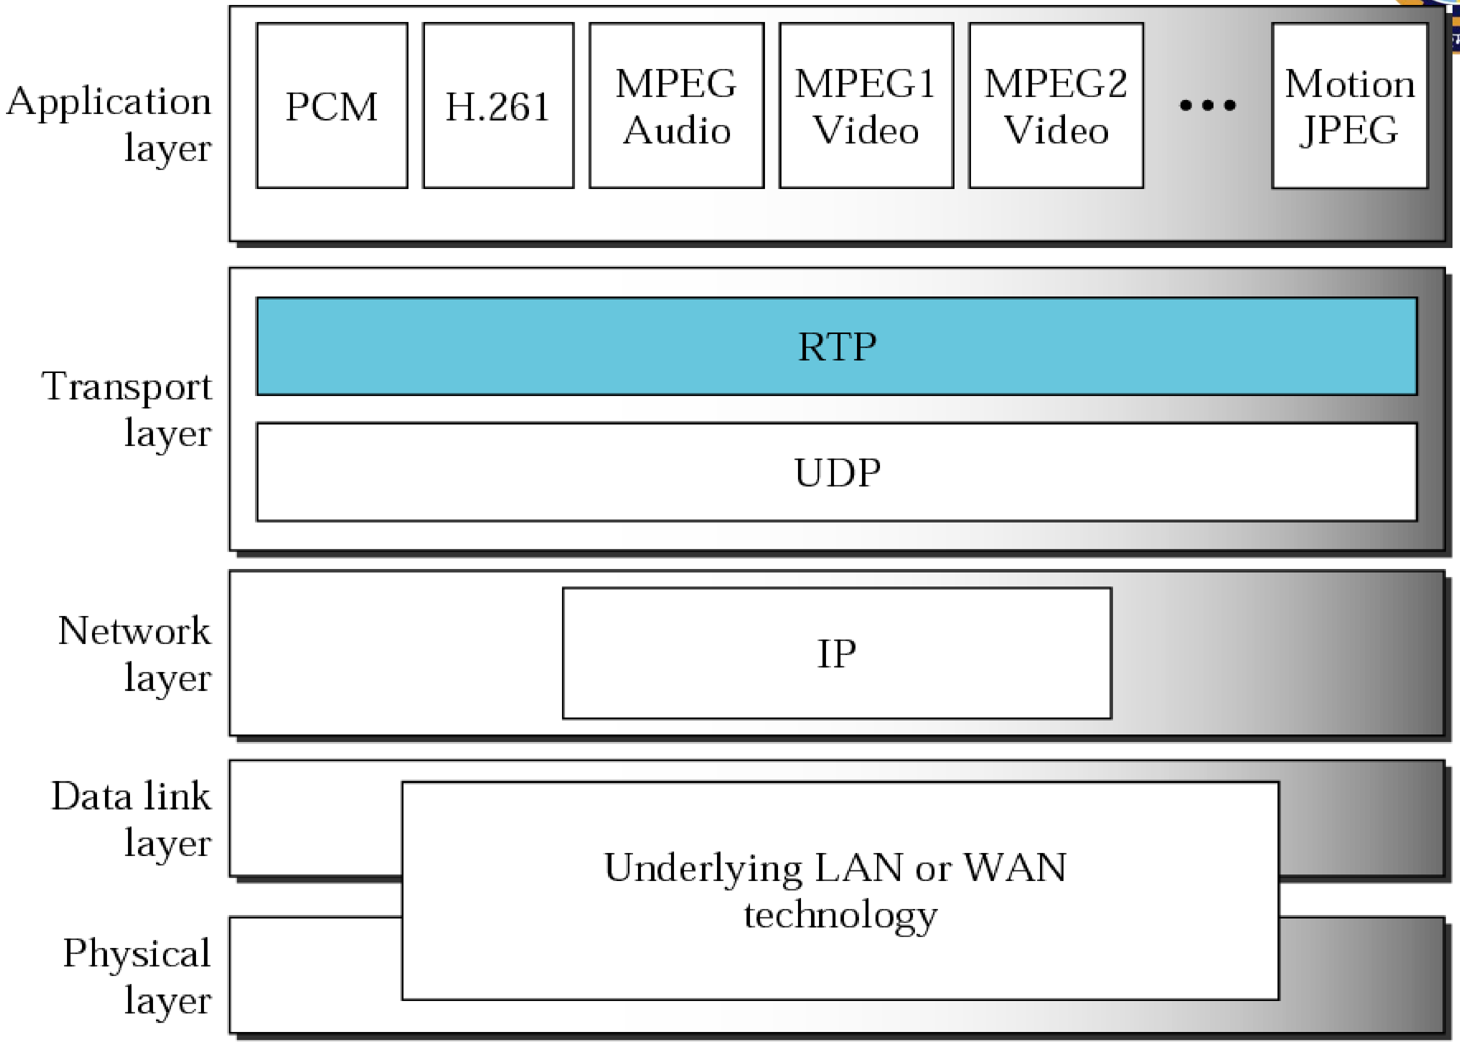

Topic Overview

- RTP and RTCP
  - Motivation
  - Architecture
  - **Packet Formats**
  - Mixers and Translators
  - Summary

RTP PDU (Packet) Format

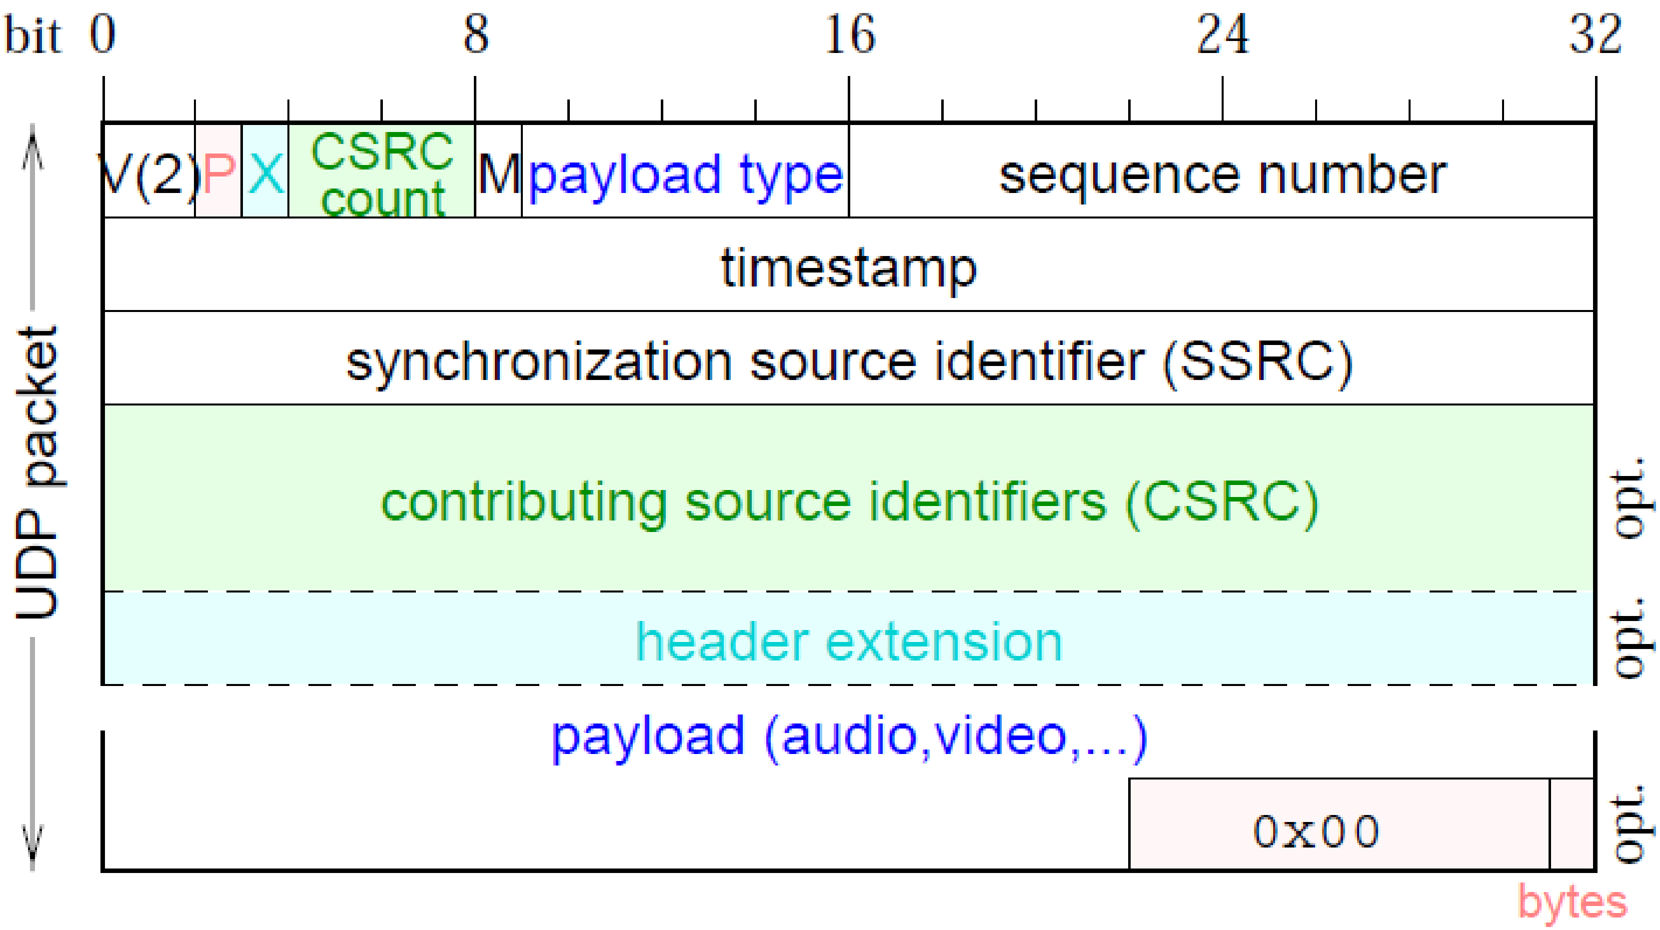

RTP PDU Fields

Payload type: audio/video encoding method; may change during session

SSRC: sychronization source ➠ sources pick at random
➠ may change after *collision*!

sequence number: +1 each packet ➠ gaps ≡ loss

P: padding (for encryption) ➠ last byte has padding count

M: marker bit; frame, start of talkspurt ➠ delay adjustment

CC: content source count (for mixers)

CSRC: identifiers of those contributing to (mixed into) packet

RTP Time Stamp

- +1 per sample (e.g., 160 for 20 ms packets @ 8000 Hz)
- random starting value
- different fixed rate for each audio PT
- 90 kHz for video
- several video frames may have same timestamp
- ➠ gaps $\equiv$ silence
- time per packet may vary
- split video frame (carefully...) across packets
- typical: 20 to 100 ms of audio

Different Kinds of Payload

Type	*Application*	*Type*	*Application*	*Type*	*Application*
0	PCMμ Audio	7	LPC audio	15	G728 audio
1	1016	8	PCMA audio	26	Motion JPEG
2	G721 audio	9	G722 audio	31	H.261
3	GSM audio	10-11	L16 audio	32	MPEG1 video
5-6	DV14 audio	14	MPEG audio	33	MPEG2 video

RTP Header Extension

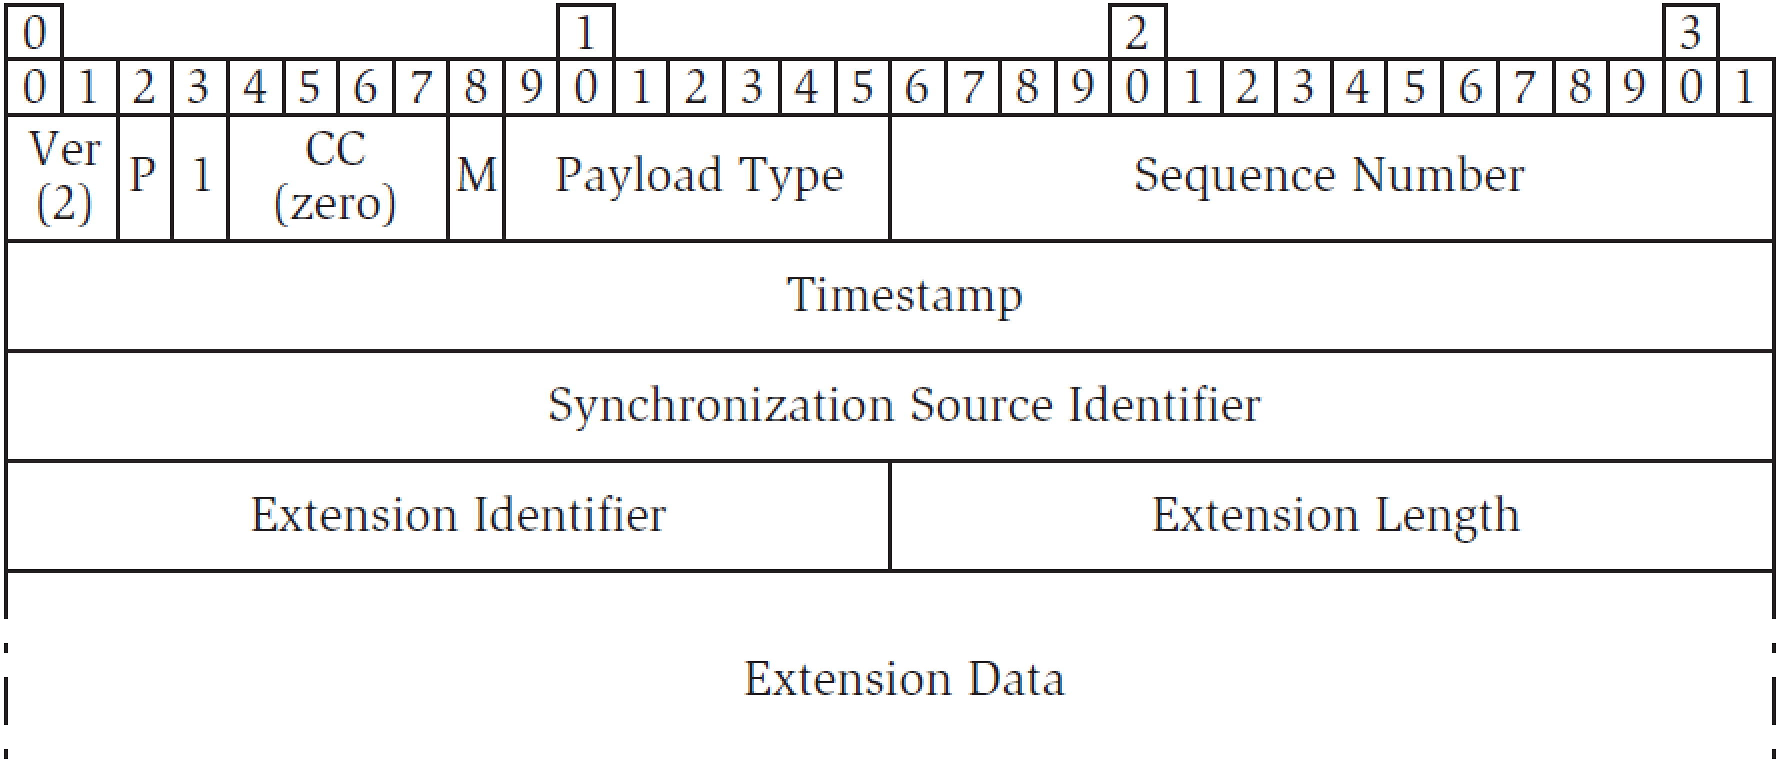

RTCP

- ID of sender
- Provides various reports for use in:
  - QoS and congestion control
    - so an app can change resolution or compression strategies
  - Session size and scaling
    - conferencing

RTP Control Protocol - Types

sender report (SR): bytes send ➠ estimate rate;
timestamp ➠ synchronization

reception reports (RR): number of packets sent and expected ➠ loss, interarrival jitter, round-trip delay

source description (SDES): name, email, location, . . .
CNAME (canonical name = user@host) identifies user across media

explicit leave (BYE): in addition to time-out

extensions (APP): application-specific (none yet)

RTCP SDES Packet

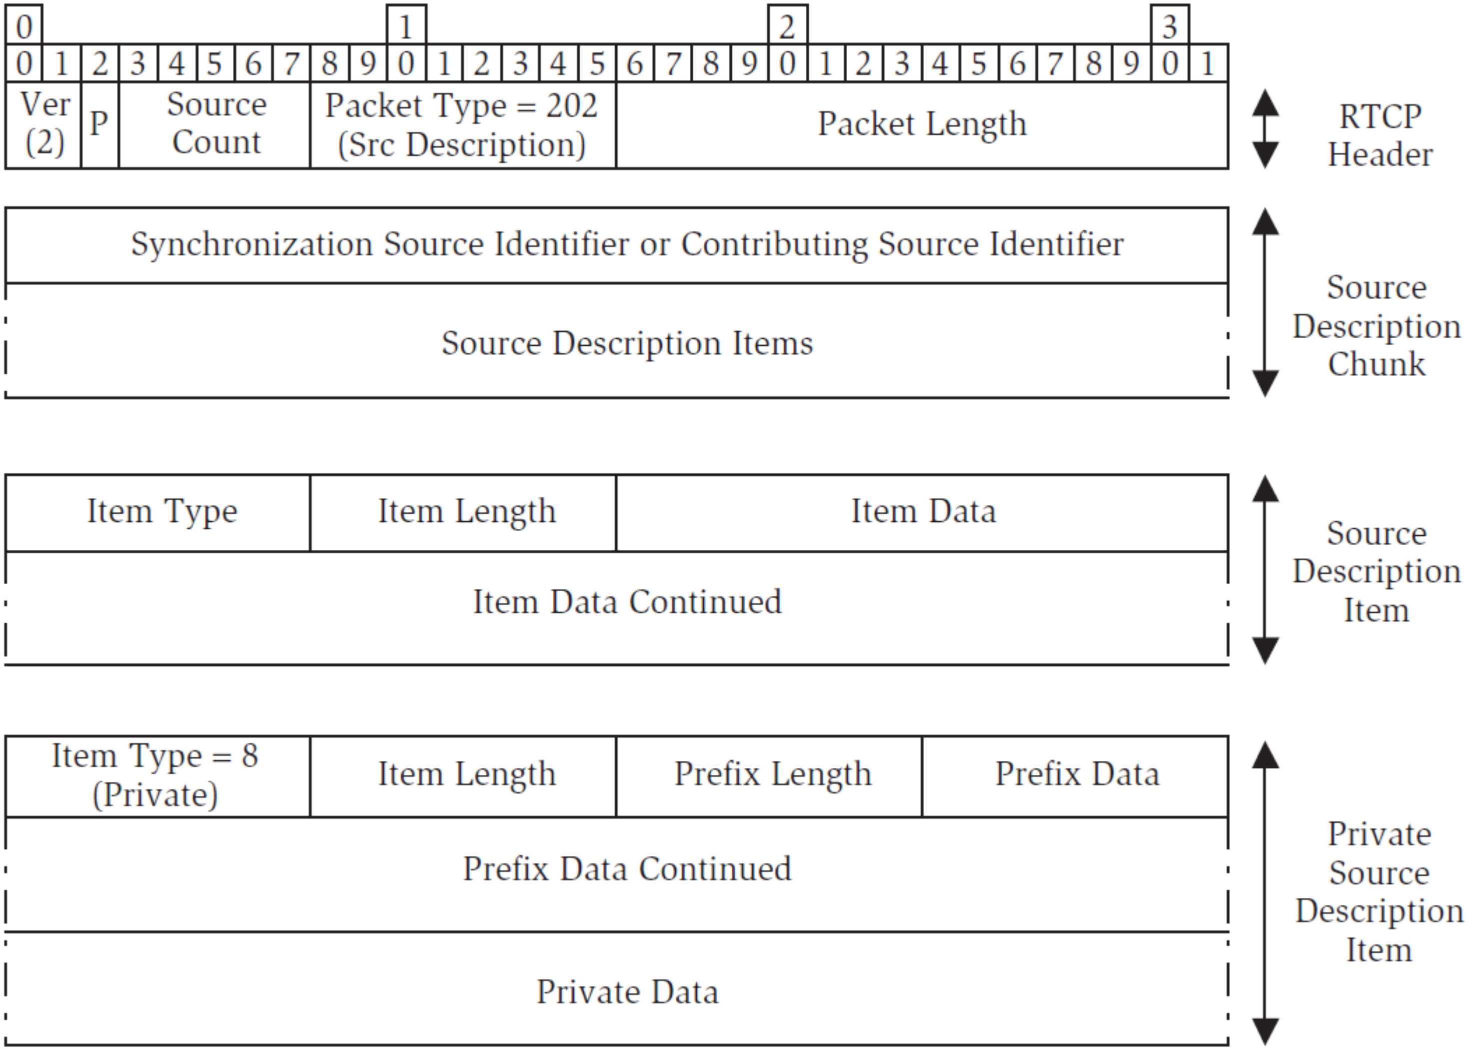

Used by Sender for describing the stream

One or more SDES items followed by optional private SDES items

RTCP SDES Item Types

Item	*Contents*
0	End of list marker. Length should be zero and data is ignored
1	Persistent transport name (canonical name) of the form "user@host" or "host" where host is either the fully qualified domain name of the host or is the IP address of the host on one of its interfaces presented in dotted notation
2	User name
3	User's email address
4	User's phone number
5	User's geographical location (address)
6	Application name
7	Free-form notes about the source
8	Additional private data.

RTCP BYE Packet

0 1 2 3

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

Ver (2)	P	Source Count	Packet Type = 203 (BYE)	Packet Length
List of Synchronization Source Identifiers or Contributing Source Identifiers				
Reason Length			Reason	
Reason Continued				

Used by the participants to leave the session

RTCP SR Packet

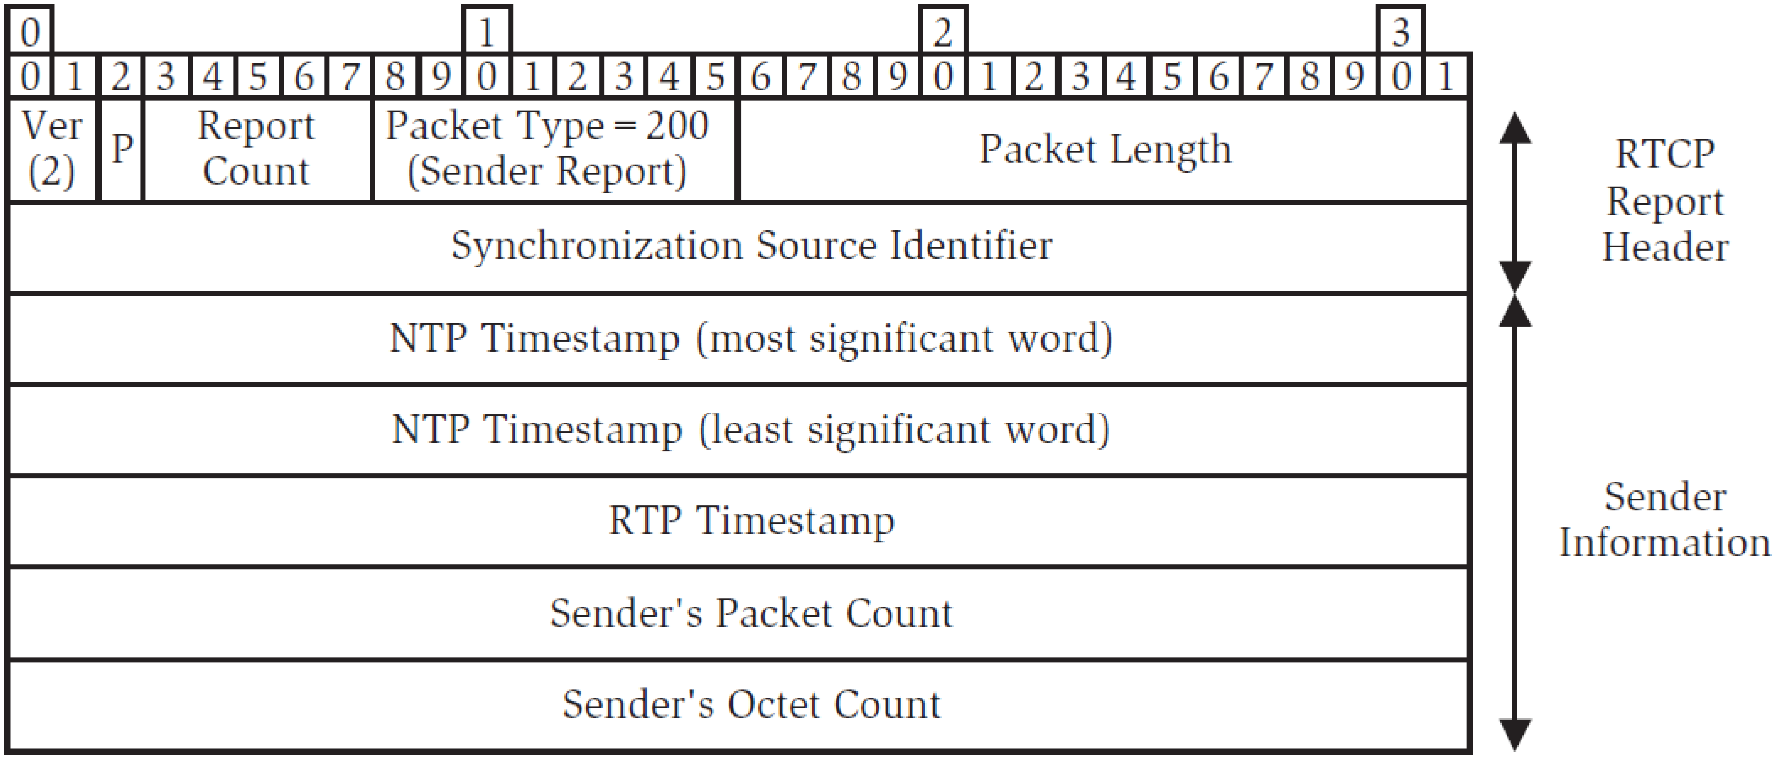

continued...

Provided by the sender to all receivers

RTCP SR Packet – Report Block

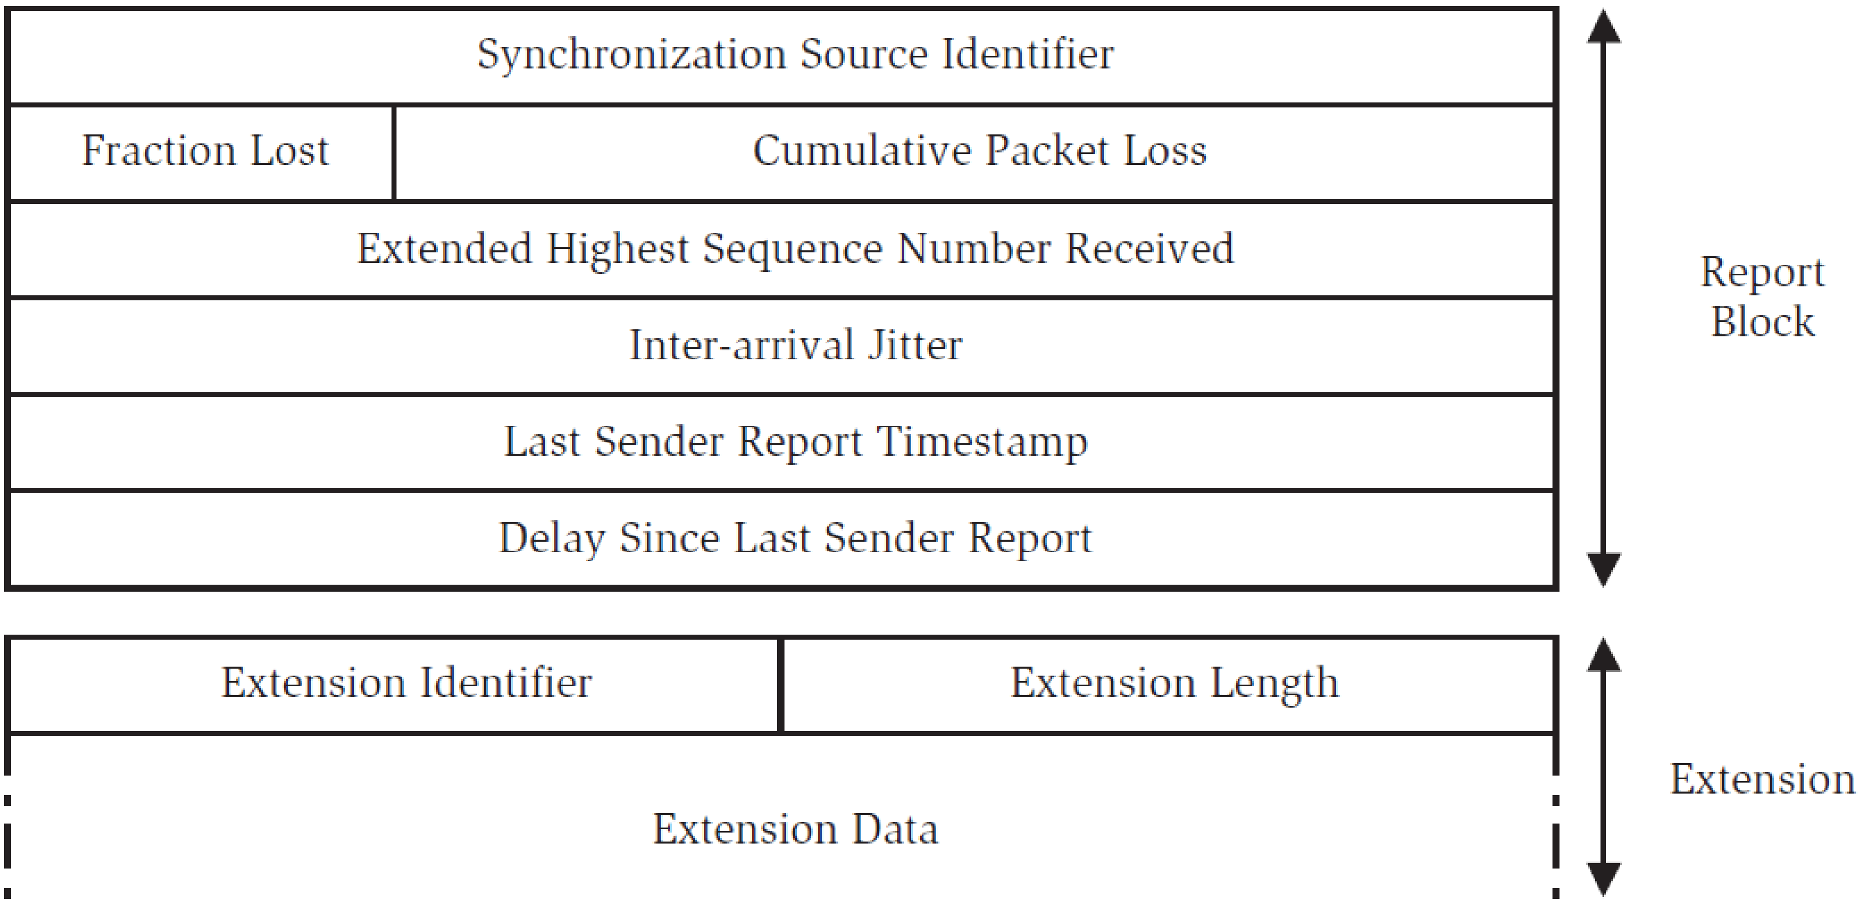

At least one report block

RTCP RR Packet

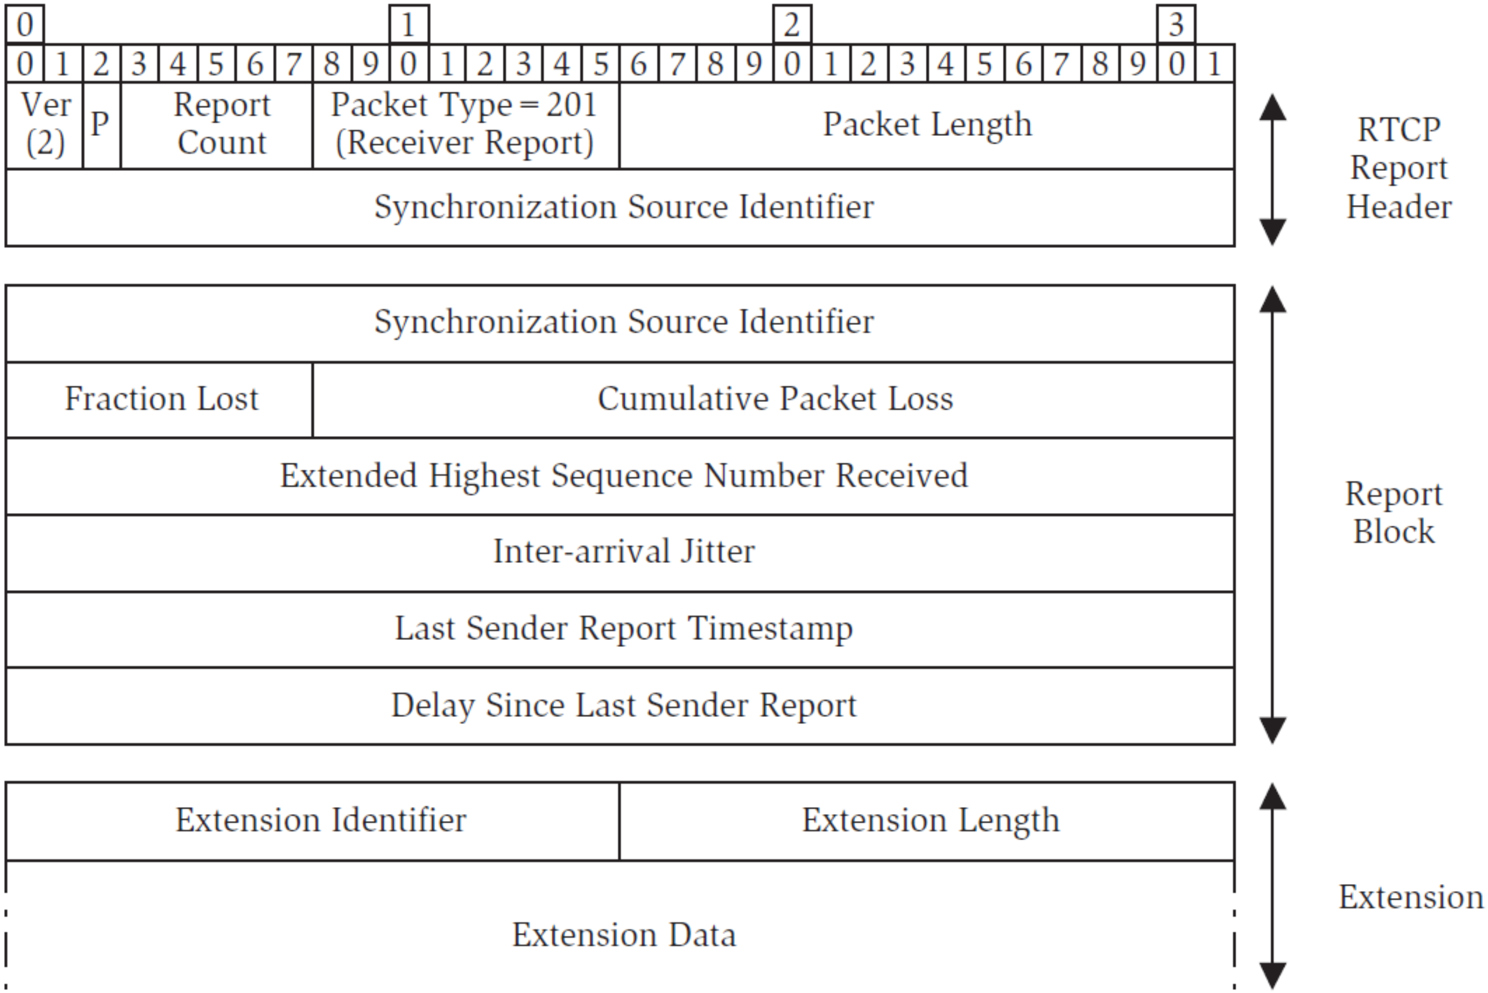

Sent by each receiver to the sender

RTCP Application Packet

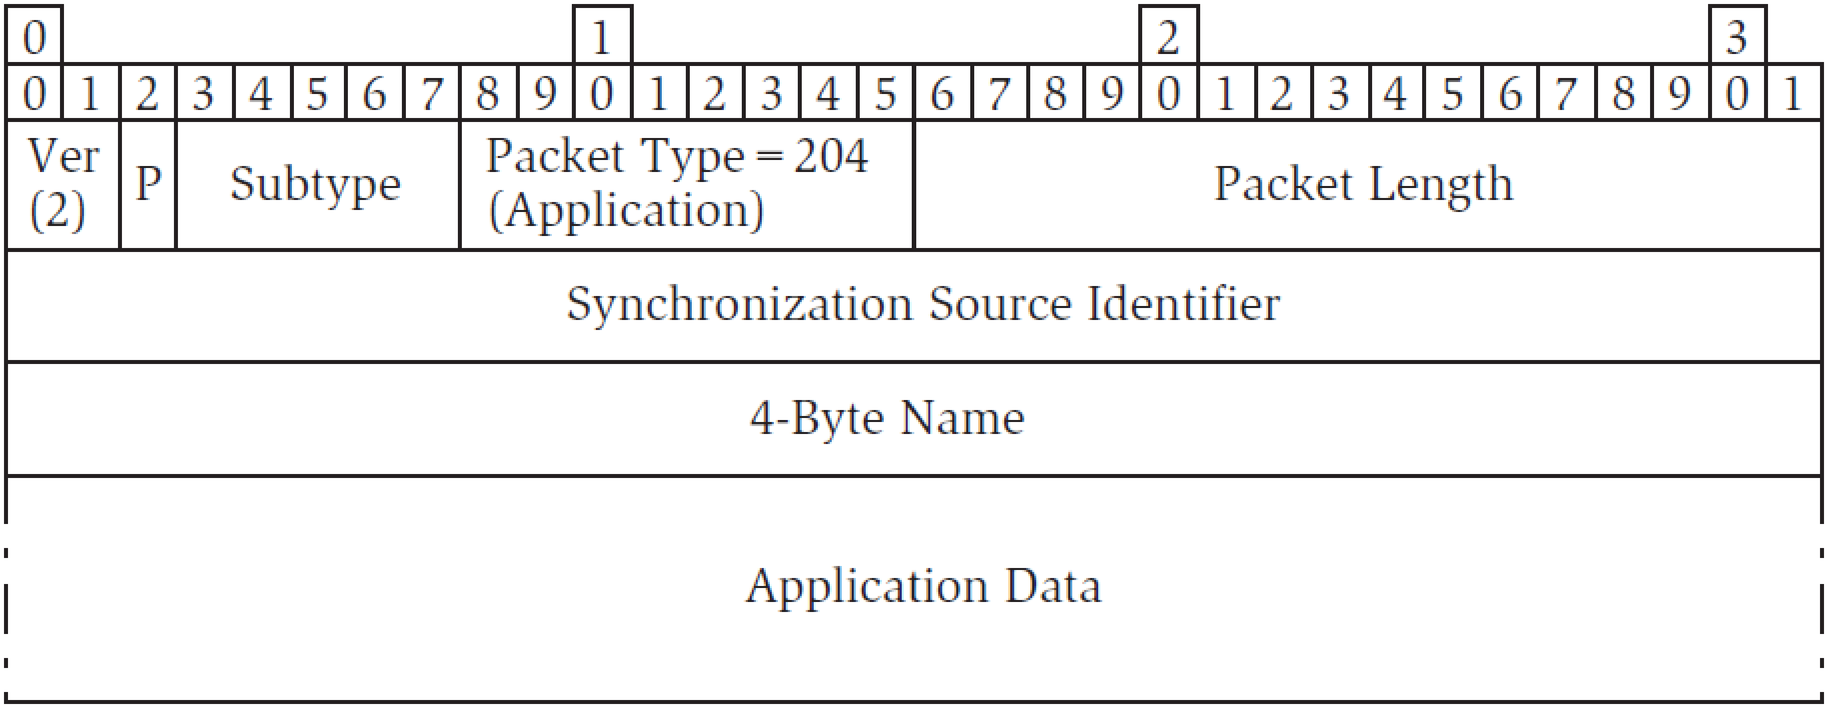

For private use by the application

Topic Overview

- RTP and RTCP
  - Summary
  - **Mixers and Translators**
  - Packet Formats
  - Architecture
  - Motivation

Mixers and Translators

- **Mixer**
  - Could receive and combine various sources in an effort to reduce bandwidth
- **Translator**
  - Keeps incoming sources separate
  - To transform to a lower quality format to broadcast on lower-speed networks
  - To send through firewalls

Mixers and Translators

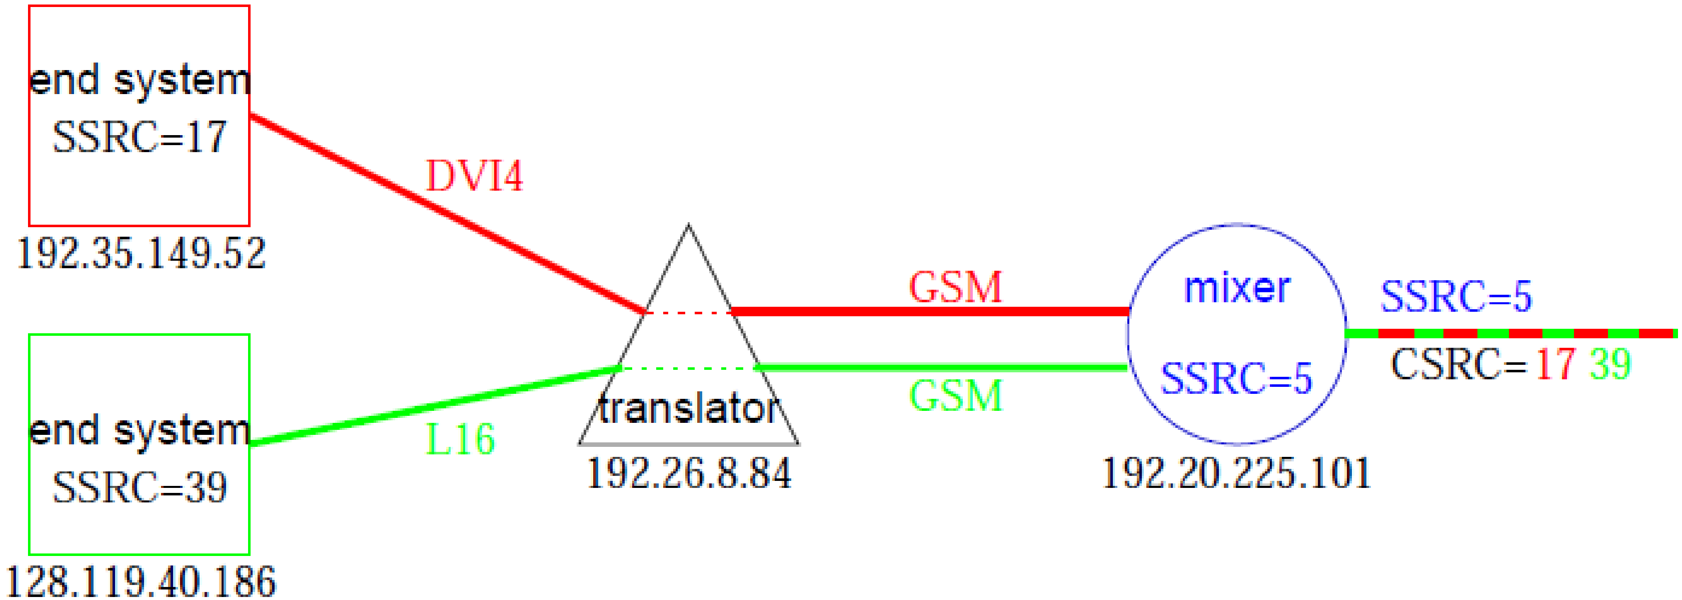

CCRC – Contributing Source Identifier

SSRC – Synchronization Source Identifier

Topic Overview

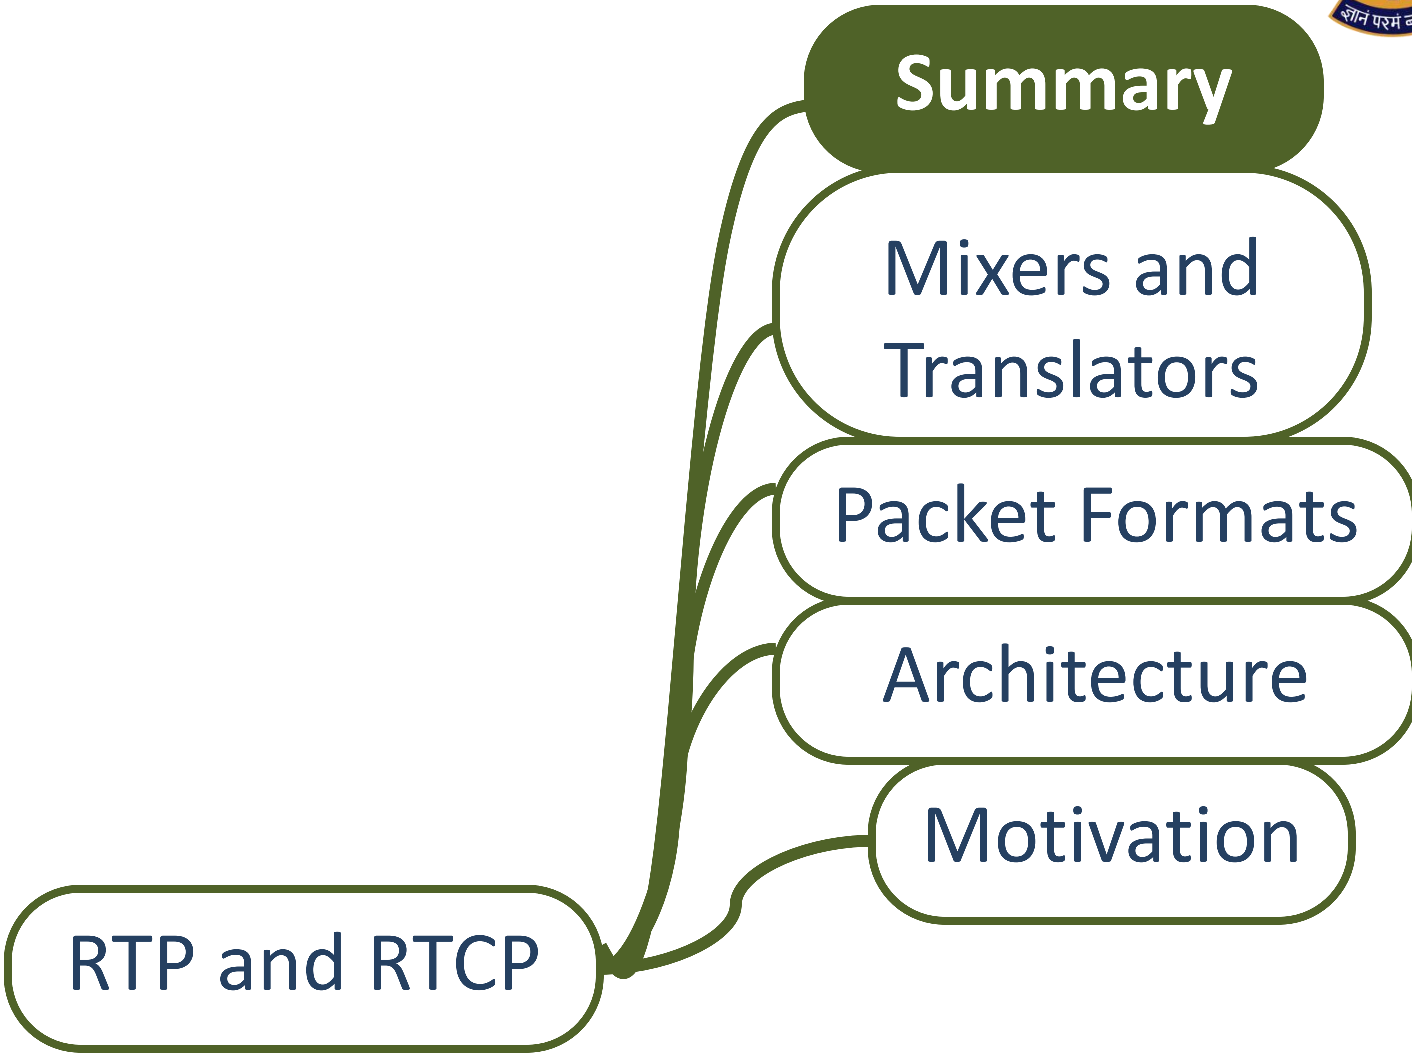

Summary

- Multimedia applications have much different needs than http or ftp!
- RTP meets those needs:
  - Minimized jitter
  - Synchronized sources
  - Dynamic, payload-specific frame length
  - Adaptation in the face of congestion
  - Interoperability
  - Effective use of bandwidth
  - Support for video-conferencing (multicast, IDs)

RTCP Scaling Issues

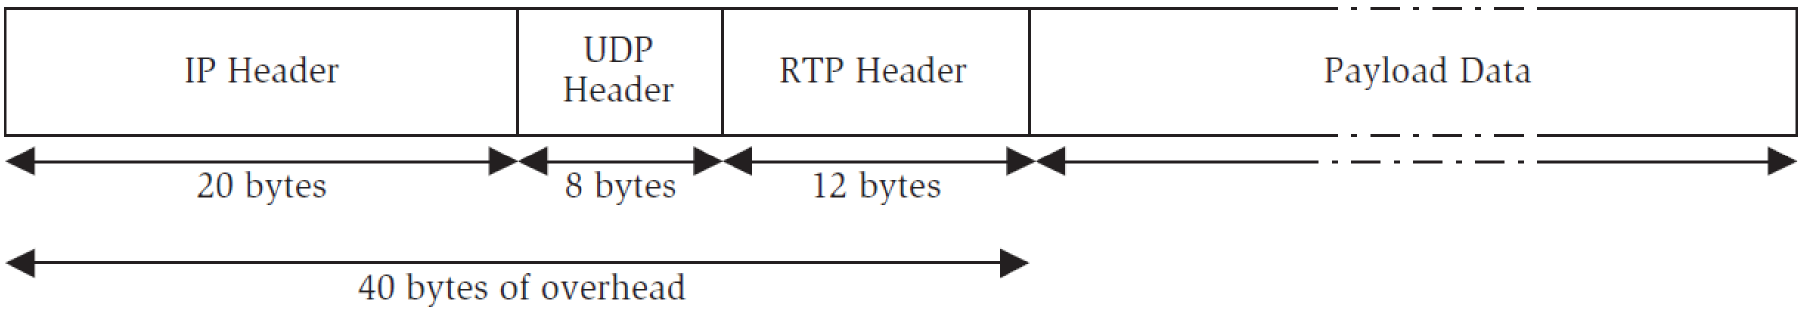

- Too much of overhead for small payload
  - Header compression (rfc2508)
- RR and BYE floods
  - Randomize the report sending times
- Place multiple RTP, RTCP packets into one UDP datagram

Sample RTP Implementations

tool	who	media	RSVP	adaptive
NeVoT	GMD Fokus	audio	yes	not yet
vic	LBNL	video	no	no
vat	LBNL	audio	no	no
rat	UCL	audio	no	no
Rendezvous	INRIA	A/V	no	yes
NetMeeting	Microsoft	A/V	no	no
IP/TV	Cisco	A/V	no	no
RM G2	Real	A/V	no	yes

Presentation Overview

- SDP
- **RTSP**
- RTP and RTCP

RTSP Introduction

- Application-layer protocol
- Control of data (typically multimedia) with real-time properties
- Inspired by HTTP/1.1, but is stateful

RTSP acts as a "network remote control" for multimedia servers.

- RFC2326

RTSP Methods

- SETUP
  - Allocate resources for a stream
- PLAY
  - Start data transmission on a stream
- RECORD
  - Record stream given by a client to a server
- PAUSE
  - Temporarily halt the stream
- TEARDOWN
  - Free the resources allocated for a stream

More RTSP Methods

- DESCRIBE
- OPTIONS
- Optional Methods
  - REDIRECT
  - ANNOUNCE
  - GET_PARAMETER
  - SET_PARAMETER

Presentation Overview

SDP

RTSP

RTP and RTCP

SDP Introduction

- Session Description Protocol
- Used by RTSP for sending session information

A well-defined format for conveying sufficient information to discover and participate in a multimedia session.

- RFC4566

Session Description Format

v= (protocol version)

o= (originator and session identifier)

s= (session name)

i=* (session information)

u=* (URI of description)

e=* (email address)

p=* (phone number)

c=* (connection information -- not required if included in all media)

b=* (zero or more bandwidth information lines)

Session Description Format

One or more time descriptions
("t=" and "r=" lines; see below)

z=* (time zone adjustments)

k=* (encryption key)

a=* (zero or more session attribute lines)

Zero or more media descriptions

Time description

t= (time the session is active)

r=* (zero or more repeat times)

Session Description Format

Media description, if present

m= (media name and transport address)

i=* (media title)

c=* (connection information -- optional if included at session level)

b=* (zero or more bandwidth information lines)

k=* (encryption key)

a=* (zero or more media attribute lines)

An Example Session Description

v=0
o=jdoe 2890844526 2890842807 IN IP4 10.47.16.5
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.example.com/seminars/sdp.pdf
e=j.doe@example.com (Jane Doe)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
m=audio 49170 RTP/AVP 0
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000