RTP: A Transport Protocol for Real-Time Applications

- Introduction
- RTP use scenarios
- RTP
- RTCP

Tao Li & M. Veeraraghavan,
University of Virginia

Introduction

- Internet standard for real-time data
  - Interactive and streamed audio and video
  - Designed for multi-user multimedia conferencing
- Provides end-to-end transport functions for real-time applications
  - Delay-oriented rather than loss-oriented (such as TCP)

Introduction – cont.

- Contains two closely linked parts: data + control
  - RTP: Real-time transport protocol
    - To transport real-time data
  - RTCP: RTP control protocol
    - QoS monitoring and feedback
    - Session control
- Protocol architecture

Applications	
RTP & RTCP	
Other transport and network protocols	UDP
	IP

Introduction – cont.

- Does NOT provide time-guaranteed delivery or other QoS guarantees
  - Relies on lower-layer protocols
- Does NOT assume the underlying network is reliable and delivers packets in sequence
  - Uses sequence number

Introduction – cont.

- RTP implementation is expected to be integrated into the application rather than as a separate module
- The use of RTP for a particular application needs other documents
  - Profile specification documents defines sets of payload type codes, and their mapping to payload formats
  - Payload format specification documents define how to carry a specific encoding
    - example: MPEG2 video or ADPCM audio
    - RFC 1890: payload types; RFC 2250: MPEG

5

RTP use scenarios

- Simple multicast audio conference
  - A multicast IP address and two UDP ports (for RTP and RTCP), assigned and distributed by mechanisms beyond the scope of RTP
  - Speaker sends:

IP header	UDP header	RTP header	Audio data

  - Receiver plays out audio data according to RTP header
  - Senders/receivers periodically multicast RTCP reports
    - Who is participating?
      - What is the audio quality?

6

RTP use scenarios – cont.

- Audio and video conference
  - Two RTP sessions, one for audio and the other for video
  - User can participate in audio, video or both
  - No direct coupling at the RTP level except a user uses the same name in RTCP packets for both audio and video sessions
  - Mixers: to mix streams from multiple sources
  - Translators: to change formats

7

RTP – packet format

V	P	X	CSRC count	M	Payload type	Sequence number (16 bits)	Fixed header
Timestamp (32 bits)							
Synchronization source (SSRC) id. (32 bits)							
Contributing source (CSRC) id. (0~15 items, 32 bit each)							optional header
Header extension (optional)							
Payload (real time data)							
				Padding (size x 8 bits)		Padding size (8bits)	optional

- Version (V, 2bits): =2
- Padding (P, 1bit): If set, last byte of payload is padding size
- Extension (X, 1bit): If set, variable-size header extension exists

8

RTP payload types (few examples); RFC 1890

Payload type	Encoding name	Audio/Video (A/V)	Clock rate
0	PCMU (mu-law G.711)	A	8000
8	PCMA (A-law G.711)	A	8000
26	JPEG	V	90000
32	MPV (MPEG-I and MPEG II)	V	90000
33	MP2T (MPEG-II transport streams)	AV	90000

SSRC and CSRC

- All packets from a given synchronizing source (with a given SSRC identifier) will use the same timing and sequence number space to allow receivers to recreate the packet sequence
- A mixer receives RTP packets from multiple sources, combines packets, makes timing adjustments, and forwards new RTP packets with a new timing sequence
  - All packets in the new sequence will have mixer SSRC as their synchronization source
  - The mixer inserts in each RTP packet header a CSRC list of the sources that contributed to this combined stream

See Section 10.6 of book by Leon-Garcia and Widjaja

RTP - header

- CSRC count (4 bits): number of CSRC identifiers
- Marker (1 bit): defined in *profile,* mark significant event
- Payload type (7 bits): Audio/Video encoding scheme
- Sequence number: random initial value, increase by one for each RTP packet; for loss detection and seq. restoration
- SSRC: identify source; chosen randomly and locally; **collision** needs to be resolved
- CSRC list: identifiers of contributing sources, inserted by *mixer*

RTP – header - timestamp

- Reflects sampling instance of the first byte in payload
- Clock frequency depends on data type; specified in *profile*
- Random initial value
- Example: CBR audio, clock increment by 1 for each sample.
- Consecutive RTP packets may have same timestamp (logically generated at same instant): Video packets that belong to the same frame
- Timestamps of consecutive RTP packets may not increase monotonically if the data is not transmitted in the order in which it was sampled: MPEG interpolated video frames

Relative timestamping scheme

- Receivers compute delay and jitter experienced by packet
- This allows them to adaptively size their reconstruction buffers

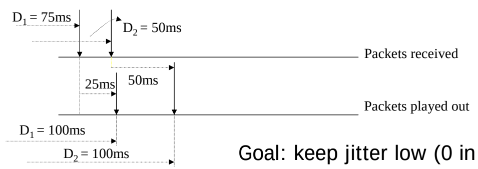

Goal: keep jitter low (0 in this case)

Delay vs. loss

- To ensure 0 jitter in playout, choose maximum delay as the total delay in selecting playout delay value
- Other option:
  - Use 95% of transmission delay to select playout delay value
  - Packets that take longer transmission delay than this 95% value will be dropped because they did not arrive in time
- Telephony requirements: 150ms one-way delay with echo cancellers; loss: 5%

RTCP: RTP control protocol

- Receivers send reports
- Report contains number of packets lost at receiver, interarrival jitter, etc.
- This allows senders to adjust data rate
- Jitter is an early indicator of congestion
- Senders also send reports

Role of RTCP

- Periodically transmit report to all participants
- Functions of RTCP:
  - Provide QoS feedback
  - Carry persistent id - Canonical name (CNAME), e.g., user@host
    - Track a user if SSRC changes (in case of conflicts)
    - Associate multiple streams from a user – synchronize A and V
  - Control the rate of RTCP packets by noting how many participants are on session – otherwise too many RTCP packets
  - Convey minimal session control information
    - Not enough for complicated session control requirements

RTCP - types

- Sender report (SR): statistics from active sender – includes RR blocks also
- Receiver report (RR): statistics from participants that are not active senders
  - RR RTCP packet sent if a node is only a receiver, i.e., it does not send data
- Source description report (SDES): includes CNAME, email, name, phone number, location, application tool/version
- BYE: indicates end of participation
- APP: application-specific functions

17

RTCP – compound packet

- RTCP packets have a length field in header; aligned to 32 bits --- stackable
- Sent in a compound packet of at least 2 RTCP packets; example:

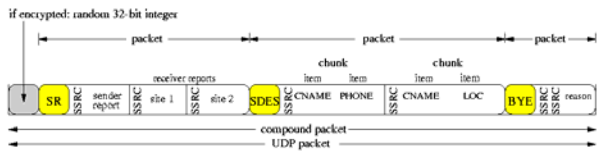

18

RTCP – sender report (SR)

- SSRC: identifies sender
- Sender information block:
  - NTP timestamp: wallclock time (absolute time as per Network Time Protocol) when packet is sent (seconds elapsed since 0 hour, January 1, 1900.
  - RTP timestamp: time when packet is sent according to the clock used to send RTP data packet timestamps; used for intra- & inter-media synchronization
  - Sender's packet count: total number of packets sent since the start of session
  - Octet count: total number of bytes sent since the start of session
- Multiple receiver report blocks, one for each source from which this host receives packets

19

RTCP Receiver Report (RR)

- SSRC_n: identifies source whose data this report block is about
- Fraction lost: fraction of packets lost since last report was sent
- Cumulative number of lost packets since the beginning of reception
- Highest sequence number received
- Inter-arrival jitter
- Last SR (LSR): The NTP timestamp of the last sender report received from the source
- Delay since Last SR (DLSR): Delay between receiving the last SR from this source and sending this RR

20

Interarrival jitter computation

- Interarrival jitter J defined as:
  - mean deviation (smoothed absolute value) of the difference D in packet spacing at the receiver compared to the sender for a pair of packets
  - D(i,j) = (Rj - Ri) - (Sj - Si) = (Rj - Sj) - (Ri - Si)
  - $J_i = J_{i-1} + (|D(i-1,i)| - J_{i-1})/16$
  - Algorithm: optimal first-order estimator; the gain parameter 1/16 gives a good noise reduction ratio while maintaining a reasonable rate of convergence [Cadzow]

RTCP – round trip time calculation

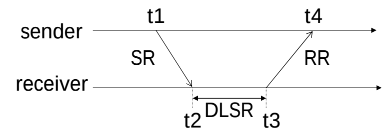

- SR packet contains: NTP (=t1)
- RR packet contains: Last SR timestamp (LSR=t1), Delay Since Last SR (DLSR=t3-t2)
- Roundtrip time = t4 - t3 + t2 - t1 = t4 – (t3-t2) –t1
  = t4 – DLSR - LSR

RTCP – RR & SDES

- Receiver Report (RR): Similar to SR but without sender information block
- RTCP Source Description packet (SDES)
  - Containing CNAME, mandatory
    - Constant for a user, unique among all users
    - Provides binding across multiple medias sent by a user
    - Example: user@ece.virginia.edu

Analyzing sender and receiver reports

- Sender may modify its transmissions based on the feedback
- Receivers can determine whether problems are local, regional or global
- Network managers may use profile-independent monitors that receive only the RTCP packets and not the corresponding RTP data packets to evaluate the performance of their networks for multicast distribution.

RTCP – transmission interval

- Designed to scale from a few to thousands of users
  - Problem: RTCP traffic is not self-limiting; grows linearly with number of usesr if sent at a constant rate
  - Solution: limit control traffic to a small and known fraction of total session traffic, 5% suggested
- Characteristics of transmission interval calc. algorithm
  - Sender occupies 25% of total control traffic bandwidth for that session to allow receivers to quickly know who is sending
  - Calculated interval should be greater than 5 seconds
  - Trans. interval randomly varied between a range to avoid synchronization of RTCP packets from many end points
  - Dynamic estimate the average RTCP packet size is made to adapt to changes in amount of control packets sent

25

Other issues

- Collision detection and resolution
  - Two sources use the same SSRC
- Loop detection
- Inter-media synchronization
- Security
- Header compression – RFC 2508
  - IP+UDP+RTP = 40 bytes, large overhead when voice packets need to be small for delay reasons

26

References

- RFC 1889, "RTP: A Transport Protocol for Real-Time Applications"
- RFC 1890, "RTP Profile for Audio and Video Conferences with Minimal Control"
- RFC 2250, "RTP Payload Format for MPEG1/MPEG2 Video"
- J. A. Cadzow, "Foundations of Digital Signal Processing and Data Analysis," New York, New York: Macmillan, 1987.

27