Using RTCP Feedback for Unicast Multimedia Congestion Control

draft-ietf-rmcat-rtp-cc-feedback-01

Colin Perkins

Motivation

- Transport protocol provides a feedback loop

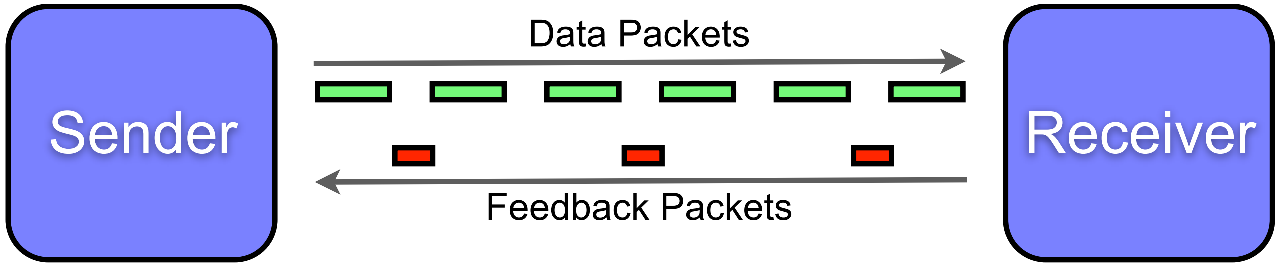

- Dynamics of congestion control depend on rate of feedback, and type of information returned
- RTCP provides a feedback channel for RTP-based applications – what sort of feedback can it provide?

TL;DR

- Questions to ask regarding congestion feedback:
  - How often is feedback needed?
  - How much overhead is acceptable?
  - How much, and what, data does each report contain?

- How often can feedback be sent in RTCP?
  - Per-packet – probably not
  - Per-video frame – yes, with reasonable assumptions – details follow
  - Per-RTT – yes in many cases, provided RTT is not too low
  - Conclusion: if configured correctly, RTCP can support congestion control, provided an appropriate feedback packet is defined

RTCP Feedback

- Four types of feedback can be used:
  - Regular RTP reports [RFC 3550]
  - RTP/AVPF feedback [RFC 4584]
  - Aggregated reports [RFC 3550, draft-ietf-avtcore-rtp-multi-stream-11]
    - Avoid UDP/IP header overhead per report
  - Reporting groups [draft-ietf-avtcore-rtp-multi-stream-optimisation-12]
    - Avoid sending unnecessary reports from co-located SSRCs
- Support in WebRTC:
  - RTCP reporting groups are OPTIONAL in draft-ietf-rtcweb-rtp-usage-26, the others are required (expect aggregated RTCP reports are not widely implemented at the sender)
  - The RTCP XR report defined in draft-dt-rmcat-feedback-message-00 is assumed to be used

Example Scenario

- Point-to-point video conference
- Two parties, each sending audio and video
- Media bundled onto single 5-tuple → 4 SSRCs
    - 1 audio SSRC, 1 video SSRC, for each party

- Can we send a congestion report for every video frame using RTCP?

Aggregation and Reporting Groups

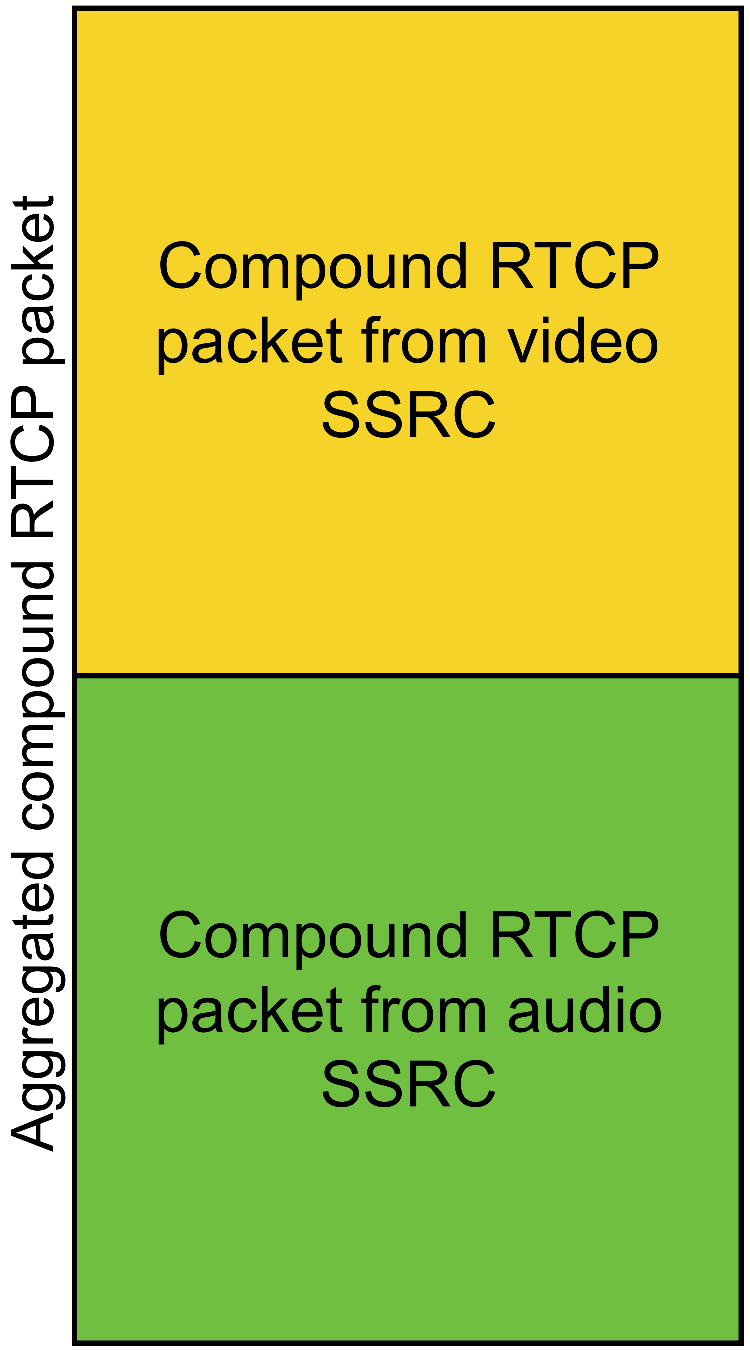

- Aggregate feedback → each RTCP packet is a compound packet, comprising a compound RTCP packet generated by the audio SSRC and a compound RTCP packet generated by the video SSRC
- RTCP reporting groups are used:
  - One SSRC is designated as the reporting SSRC
  - The other SSRC delegates its reports to that SSRC
  - The reports are aggregated, so it doesn't matter which is chosen as reporting SSRC (slides assume video SSRC is reporting SSRC)

Non-reporting SSRC

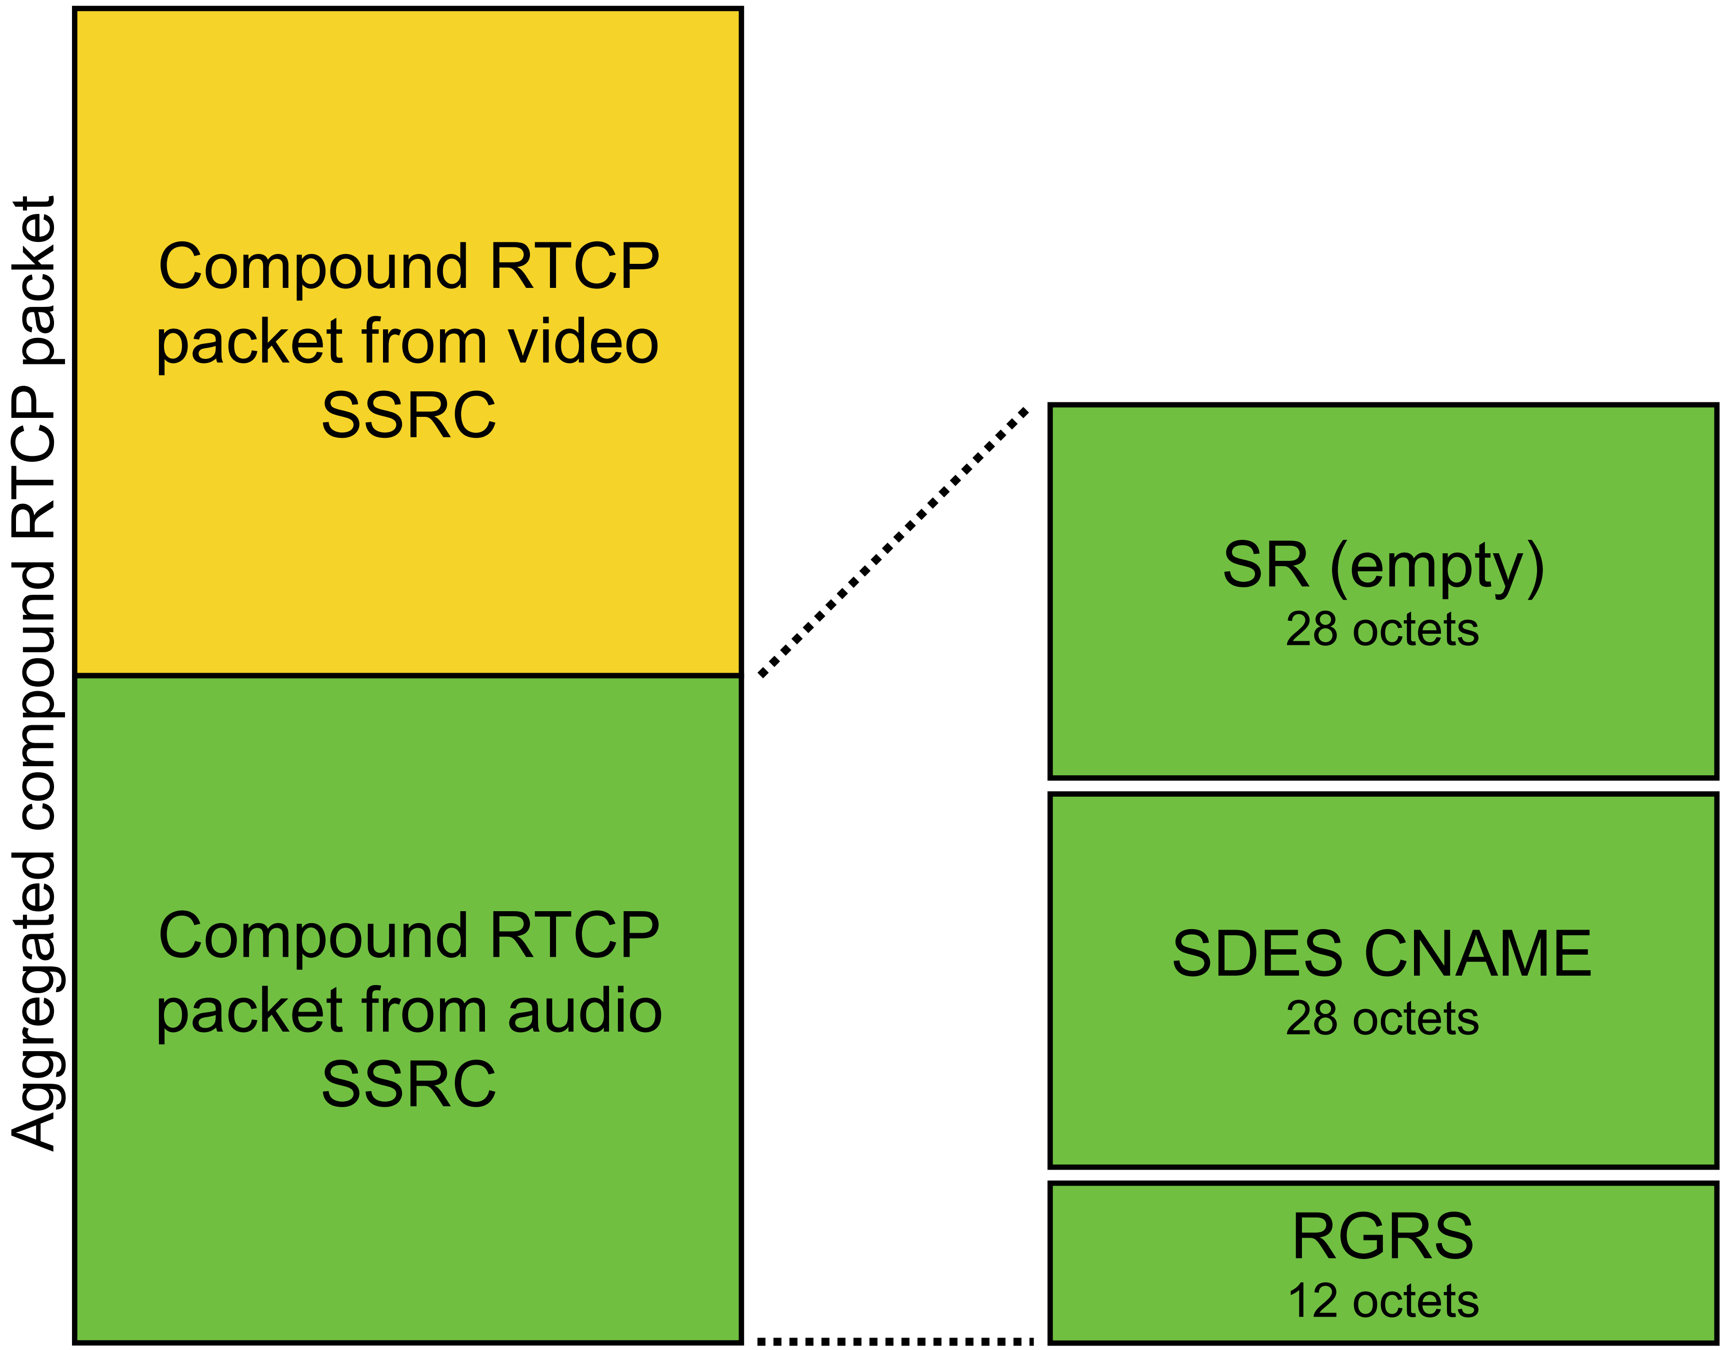

Reporting SSRC

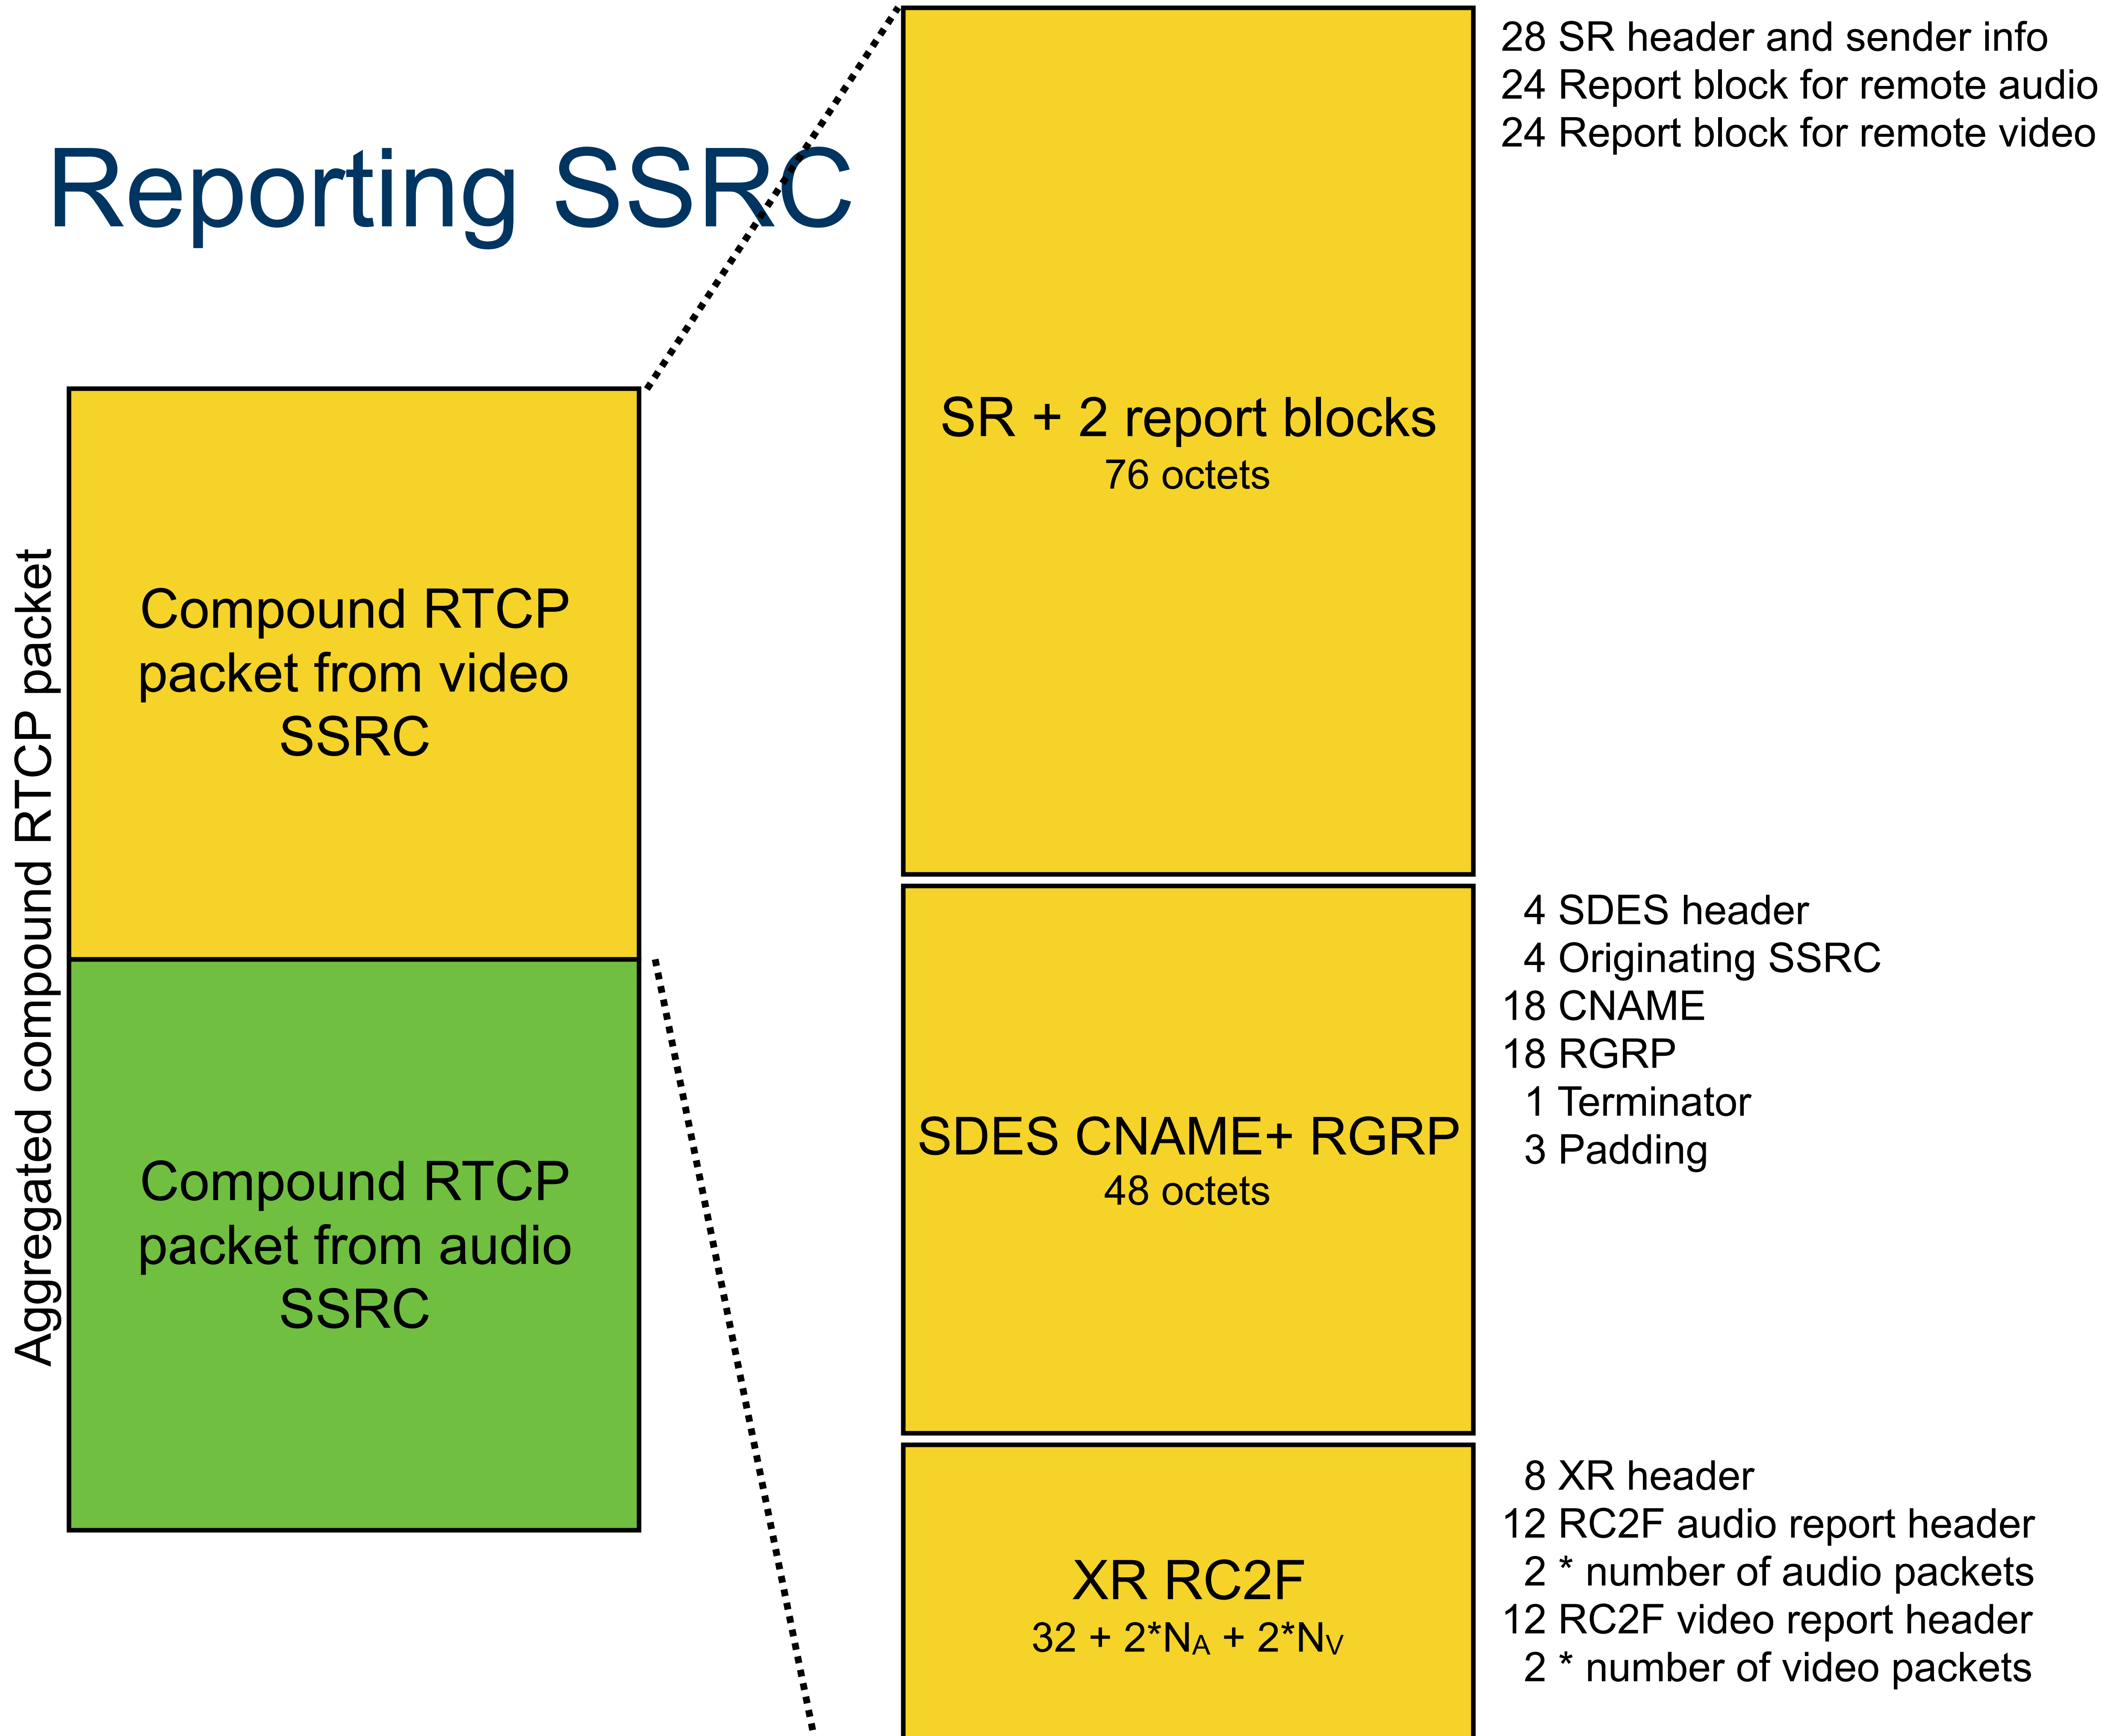

Report Size: Overall

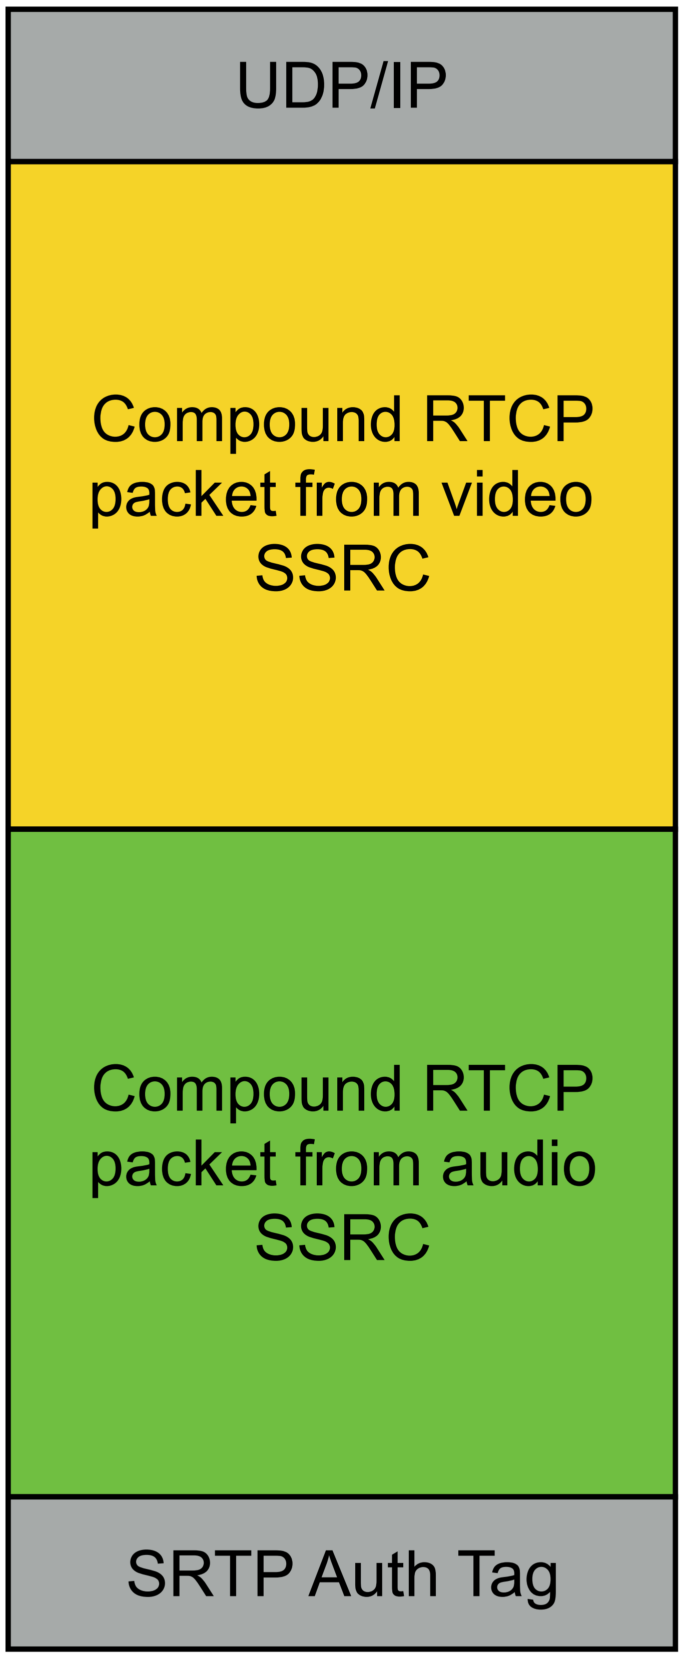

28

76 SR
48 SDES
$32 + 2{*}N_a + 2{*}N_v$ XR
$= 156 + 2{*}N_a + 2{*}N_v$ octets

28 SR
28 SDES
12 RGRS
= 68 octets

Total size = $28 + 156 + 2{*}Na + 2{*}N_v + 68 + 14$
$= 266 + 2{*}Na + 2{*}N_v$ octets

14

Report Size: Number of Packets in Report

- What are N_a and N_v?
- Assume:
  - video_bit_rate_bps ≈ session_bw
  - video_packets_per_second = (video_bit_rate_bps / 8) / mtu
  - audio_packets_per_second = 50

  - N_v = ceil(video_packets_per_second / fps)
  - N_a = ceil(audio_packets_per_second / fps)

  - (these assumptions are not realistic)

RTCP Reporting Interval

- Reporting interval for RTP/AVPF with $T_{rr\_interval} = 0$ is rtcp_interval = avg_rtcp_size * n / rtcp_bw

- For our scenario:
  - n = 4
  - avg_packet_size = 266 + 2*Na + 2*N_v
  - rtcp_bw = 5% of session_bw
  - Because of aggregation, avg_packet_size is halved
    [draft-ietf-avtcore-rtp-multi-stream-11]

Can we report per frame?

```
session_bw = 2.5 × 1024 × 1024/8 ........................ 327,680
fps = 30 ......................................................... 30
mtu = 1500 ................................................... 1,500
Na = ceil(50/fps) ................................................ 2
Nv = ceil((session_bw/mtu)/fps) .................................. 8
avg_rtcp_packet_size = (42 + 68 + 156 + (2 × Nv) + (2 × Na))/2 ... 143
n = 4 ............................................................ 4
rtcp_interval = (avg_rtcp_packet_size × n) / (session_bw × 0.05)... 0.0349121094
```

If session bandwidth > ~2.5Mbps we can report on each frame of 30fps video

Summary and Next Steps

- With the RTCP congestion feedback format, and standard RTCP features, we can report on every frame of 30fps video if video bandwidth > 2.5Mbps
- Obvious ways to optimise this, without changing the congestion report format
  - RGRP extensions have high overhead
- Analysis is *very* preliminary – ongoing