University of Glasgow

School of Computing Science

RTP Congestion Control Feedback

Colin Perkins

The draft-ietf-avtcore-cc-feedback-message-02 is co-authored with Zahed Sarker, Varun Singh, and Michael A. Ramalho, but these are my opinions solely

Introduction

- RTCP congestion control feedback draft-ietf-avtcore-cc-feedback-message-02

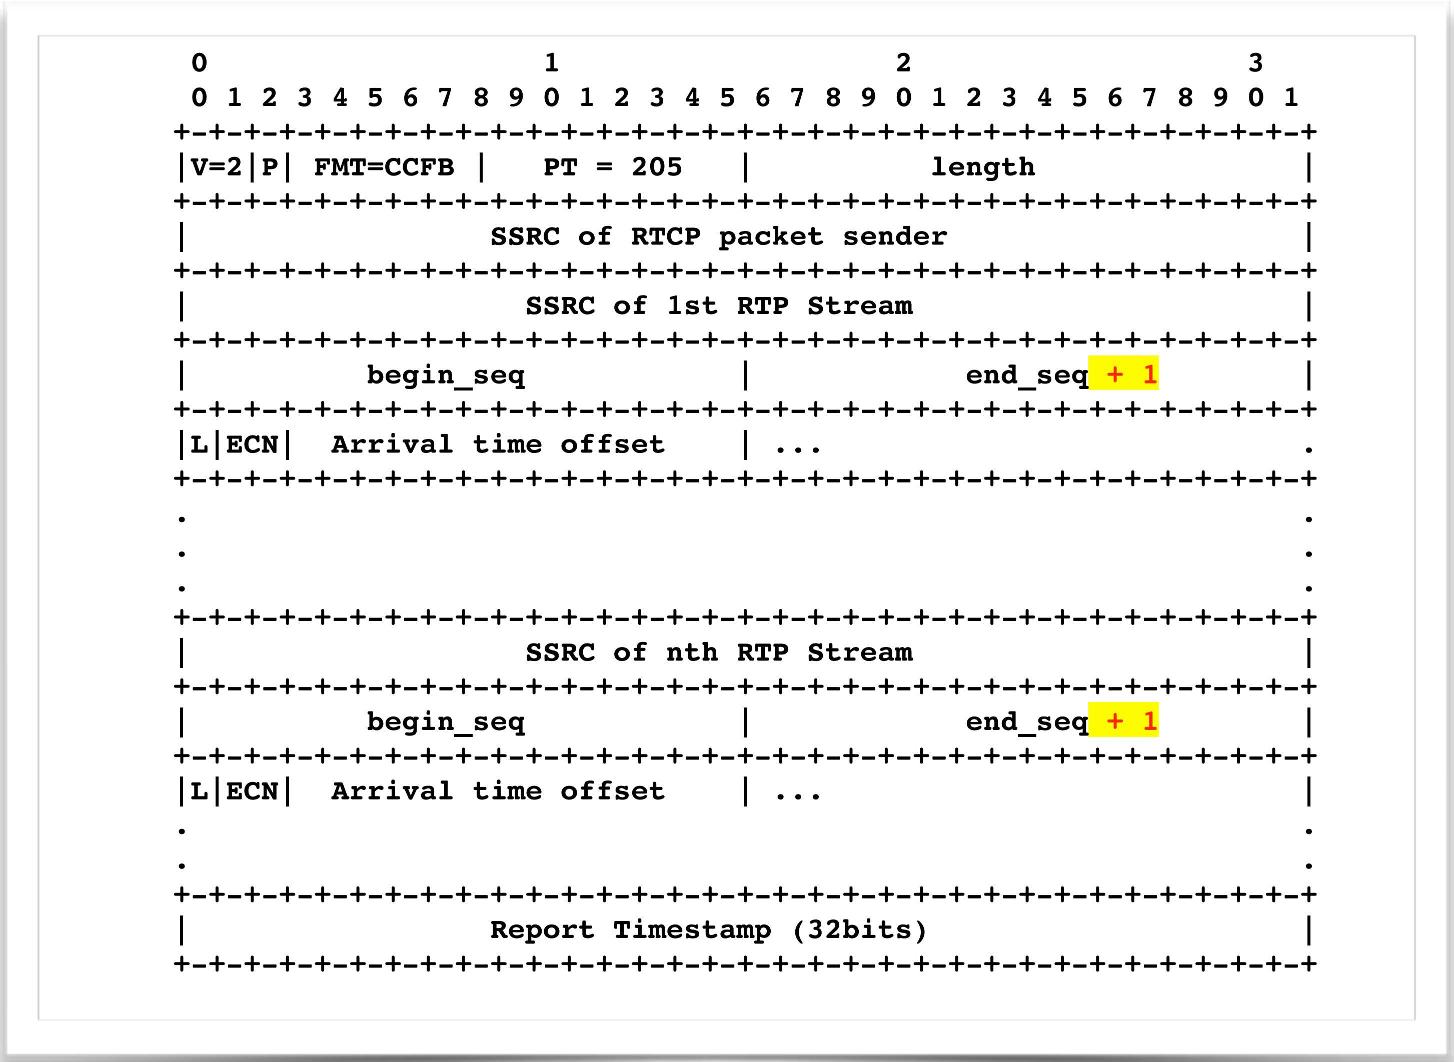

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P| FMT=CCFB |   PT = 205    |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 SSRC of RTCP packet sender                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   SSRC of 1st RTP Stream                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          begin_seq             |             end_seq + 1       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|ECN|  Arrival time offset    | ...                             .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                                 .
.                                                                 .
.                                                                 .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   SSRC of nth RTP Stream                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          begin_seq             |             end_seq + 1       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|ECN|  Arrival time offset    | ...                             |
.                                                                 .
.                                                                 .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 Report Timestamp (32bits)                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- Jonathan Lennox sent feedback based on hackathon experience
  https://mailarchive.ietf.org/arch/msg/avt/MdxCH_NlOuamWw47ms1Vo59p9ps

Parsing

- *Parsing CC-Feedback packets is a bit odd — the way you tell there are no more streams in a feedback packet is that there are only four remaining bytes in the packet, for the report timestamp. This is unambiguous, but perhaps non-obvious.*

- No space for a report count, unless we either:
  - Put it at the start and mis-align all the reports – ugly
  - Put it in the last octet, and shrink the report timestamp field – maybe?

Timing

- *I think that having the report timestamp be “derived from the same wall clock used to generate the NTP timestamp field in RTCP Sender Report (SR) packets” (presumably meaning SR reports from the packet sender?) is unfortunate. SR NTP timestamps often use the real system wall-clock, i.e. actual NTP time-of-day, which is subject to clock adjustments. For CC feedback, however, I think we want a clock that’s more stable than that. I’d suggest that report timestamps from a given sender must always use the same clock, which SHOULD be stable, but otherwise can be unrelated to any other clock on the system.*

- SR clock is unambiguous, an alternative may not be – does that matter?
- Not sure relating to media clock makes sense – different clock rates
- Does introducing a new clock matter?

Timing

- *What is the point of having the report timestamp have higher precision than the arrival time offsets? The report timestamp is measured in 1/65536 of a second, but the arrival time offset is only in 1/1024 of a second.*

- Re-using a timestamp format from elsewhere in the RTP spec
- We could save some bits, and leave space for a report count in the last octet of the packet?

Timing

- *What value should be used for arrival time offsets for packets that arrived after the report timestamp? (If you have consecutive packets that were reordered by more than the maximum arrival time offset, this might be unavoidable.). I think they should be reported as “not received", because they hadn’t arrived as of the report timestamp, but this should perhaps be clarified.*

- Agree with not received – since they hadn’t arrived at the time the report was sent

Timing

- *It’s not also clear what precise semantics "the time instant when the report packet was generated” has. In particular, this isn’t the time the report packet was *sent*, so it can’t be used for RTT calculations, right? Additionally, if you need to generate multiple report packets to prevent arrival time offsets from being out of range, at least one of the reports’ report timestamps will *not* be the generation time.*

- Agree – should be time the report was generated?

Overlapping Reports

- *The sentence "If overlapping reports are sent, the information in the later report updates that in any previous reports for packets included in both reports" should be clarified to include the fact that later feedback messages MUST NOT indicate loss for packets that earlier feedback messages reported on. This has the consequence that feedback senders can't just purge all data about received packets once feedback has been sent for them, if it's possible that sequentially earlier, reordered packets might still arrive later.*

- Agree – and very old packets might have to be reported with "out of range" offsets

Overlapping Reports

- *The document says "The sequence number ranges reported on in consecutive reports for an SSRC SHOULD be consecutive and SHOULD NOT overlap (i.e., begin_seq for a report is expected to be one greater, modulo 65535, than end_seq of the previous report for that SSRC)." Should this be limited in some way? E.g., if a source's sequence numbers reset after an outage, how far back should the loss reports go? I think that we don't want to be sending reports for thousands of sequence numbers. Additionally, if you had packet reordering around the time a feedback message was sent, you probably want to go back and report earlier packets that arrived after the feedback message went out, even if this results in a sequence number overlap.*

- Yes – would appreciate wording suggestions

Reporting on Duplicates

- *There should probably be guidance on what a feedback sender should do if it receives duplicate copies of the same packet. Should the arrival time be that of the first or last copy?*

- Agree we need guidance
- Report on last copy makes sense to me, but no strong opinion

Reporting on What SSRCs?

- *The document says "RTCP congestion control feedback packets SHOULD include a report block for each SSRC that is being congestion controlled", but it's not clear how a receiver can know which sources are being congestion controlled. Moreover, if you're in a situation where there could be lots of SSRCs (e.g., media coming from an SFU) it's not clear to me that you want to include reports for SSRCs that have been inactive for a long time. It's probably better to say SHOULD be sent for every active source (i.e. source for which you will send a report block in your next SR/RR).*
- Agree

Signalling

- *The rtcp-fb SDP negotiation mechanism supports limiting a feedback message to a specific subset of the payload types. Should this be allowed for CCFB negotiation? Why would you want it? What does it mean?*

- Unclear – I can imagine it might make sense to limit to video PTs if multiplexing VBR video with low-rate CBR audio, but not sure the complexity is worth it

Editorial

- *It should be made more clear that the arrival time offsets indicate time *before* the report timestamp.*
- *The reference in the SDP syntax is wrong; it should be [RFC4585] , not [RFC4584].*

- Will fix