draft-wenger-avtext-avpf-ccm-layered-00

Stephan Wenger

Jonathan Lennox

Bo Burman

Magnus Westerlund

Overview

- Draft covers the applicability and interaction of Codec Control Messages (CCM), defined in various RFCs, with layered (Video-) codecs.
- Original motivation: underspecified Full Intra Request—fixed
- While fixing FIR, we found that it is hard to identify the use of a layered codec—added informative language aiding the implementer
- We also looked at the applicability and mapping of all other currently defined CCMs

Full Intra Request in layered codecs

- When using multiple layers and in MRST or MRMT mode, and FIR is received associated with only one RTP stream, does the media sender needs to send:
  - Decoder Refresh Point (IDR) only for the layer(s) for which FIR is received, or
  - Decoder Refresh Point for all layers?
- Design Choice: all layers
- This was discussed already in Yokohama and agreed there

Layered Codec Identification

- Informative language covering the MRST and MRMT cases
- Use of payload format specified for layered codecs (such as RFC 6190) is NOT a sufficient condition for layered codec use
- Binding of layers together through SDP (RFC 5583) is a good indication (when SDP is in use)
- Signaling of simulcast of RTP streams is an indication for not sending an IDR for any other simulcast stream but the targeted stream when FIR is received only for one simulcast stream
  - Parse that! :-)
  - Issue here is that the simulcast draft allows simulcasting of multiple steams sets that include dependent streams
  - Subject draft may be underspecified. Or simulcast daft is overly flexible?

Applicability to other CCMs

- PLI: no need for IETF spec—encoder takes care of appropriate repair
- SLI: see PLI. Also, no one uses SLI in layered codecs. See open issue #1
- RPSI: see SLI and open issue #1
- TSTR/TSTN: Temporal/Spatial Trade-Off. See Open Issue #2

Open issue #1: appropriateness and correctness of deployment language

- SLI and RPSI state
  - SLI: “SLI has seen very little implementation and, as far as it is known, none in conjunction with layered systems.”
  - RPSI: “While a technical equivalent of RPSI has been in use with non-layered systems for many years, no implementations are known in conjunction layered codecs. ”
- Anyone has a different opinion re deployment?
- Is language appropriate for RFC?
  - Used as a signal to caution readers that there is no implementation experience

Open issue #2:

- TSTR/TSTN: Temporal-Spatial Trade-Off Request / Notification
  - In a typical implementation, a TSTR is generated by receiving video conferencing system based on user input, and conveyed via RTCP to the encoder. The encoder adjusts its coding strategy accordingly and reports back via TSTN. The TSTN can be used for user-interface feedback
- Layered codecs have a lot of options to adjust tempo-spatial tradeoff through dumping layers, as well as traditional encoder tuning.
- No known implementation practice
- Draft suggests to leave reaction open until reports of implementation practice are received. Such reports are solicited.
- Any implementation practice reports? Language appropriate?

Next Steps

- Accept as WG draft?
- Depending on discussion of the two open issues, WGLC?

- Thank you