University of Glasgow | School of Computing Science

RTP Congestion Control: Circuit Breakers for Unicast Sessions

draft-perkins-avtcore-rtp-circuit-breakers-00

Colin Perkins – University of Glasgow
Varun Singh – Aalto University

Background and Goals

- RTP requires applications to be congestion aware, but lacks a standard congestion control algorithm
- Interest in developing and standardising congestion control algorithms for WebRTC
  - Discussion in IRTF ICCRG and IETF RTCWeb WG sessions this week
  - These algorithms are new, and will take time to develop and be validated
  - WebRTC, and other, applications need an immediate safety-net, to allow initial deployment before sophisticated congestion control is developed
- Goal of this draft is to define the envelope within which these algorithms must work
  - Define “circuit breaker” conditions for an RTP session – limits that should not be met under normal operation, and can be used to stop errant flows

Congestion Signals for RTP/AVP Flows

- Potential congestion signals available from RTCP:
  - RTT estimate once per reporting interval
  - Jitter estimate once per reporting interval (limited use for video flows)
  - Fraction of packets lost during the reporting interval, plus cumulative number of packets lost over the entire RTP session
- Not (yet) considering RTP extensions – looking for a baseline mechanism
- Applicability as circuit breakers:
  - RTT/jitter estimates too infrequent to be useful with RTP/AVP timing rules
  - Packet loss statistics too infrequent for adaptation, but useful for detecting overload situations – use as the basis for a limiting condition

Circuit Breakers for Unicast RTP/AVP

- Circuit breaker #1: Timeout
  - RTP data packets being sent, but corresponding RTCP RR packets report non-increasing extended highest sequence number received
  - Indication of significant connectivity problem if persistent for ≥ 2 reporting intervals → cease transmission
- Circuit breaker #2: Congestion
  - RTP data sent, corresponding RR packets have increasing extended highest sequence number received, but non-zero packet loss fraction
  - Indication of network congestion – estimate equivalent TCP throughput:

$$T = \frac{s}{R\sqrt{\frac{2p}{3}} + (t_{RTO}(3\sqrt{\frac{3p}{8}})p(1+32p^2))}$$

R = round trip time, s = packet size
p = packet loss fraction

where $t_{RTO} \approx 4R$, and cease transmission if RTP sending rate ≥ $10T$ for 2 reporting intervals (similar to TFRC definition – see draft for rationale and assumptions)

Open Issues

- Are these appropriate rate limiting conditions for RTP/AVP sessions?
- What is the impact of RTP extensions?
  - E.g., RTCP XR, RTP/AVPF profile, and ECN

Next Steps

- Presentation and further discussion in IRTF ICCRG
- Consider for adoption as a working group draft