

FOUGH, N., VERDICCHIO, F., PERKINS, C. and FAIRHURST, G. 2015. Media usability circuit breakers for RTP-based interactive networked multimedia. In *Proceedings of the 2015 Picture coding symposium (PCS 2015), co-located with the 2015 Packet video workshop (PV 2015), 31 May - 03 June 2015, Cairns, Australia*. Piscataway: IEEE [online], pages 227-232. Available from: <https://doi.org/10.1109/PCS.2015.7170080>

Media usability circuit breakers for RTP-based interactive networked multimedia.

FOUGH, N., VERDICCHIO, F., PERKINS, C. and FAIRHURST, G.

2015

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Media Usability Circuit Breakers for RTP-Based Interactive Networked Multimedia

Nazila Fough
University of Aberdeen
Aberdeen, UK
nazila@erg.abdn.ac.uk

Fabio Verdicchio
University of Aberdeen
Aberdeen, UK
fverdicc@abdn.ac.uk

Colin Perkins
University of Glasgow
Glasgow, UK
csp@csp Perkins.org

Gorry Fairhurst
University of Aberdeen
Aberdeen, UK
gorry@erg.abdn.ac.uk

Abstract— With multimedia and Internet enabled devices being ubiquitous, mechanisms that ensure multimedia flows do not congest the Internet are crucial components of multimedia systems that are embraced rather than opposed by network service providers. The Real-time Transport Protocol (RTP) Circuit Breaker is designed to terminate RTP/UDP flows that cause excessive congestion in the network. Multimedia users congesting the network have their flows terminated, as dictated by the RTP circuit breaker congestion rule. Users who obtain little quality from a multimedia session, and consume network resources to no avail, should also cease transmission. This is the mandate of the RTP circuit breaker media usability rule. We propose an algorithm for this rule, and show that it avoids wasting network resources on flows that deliver no quality to the user.

Index Terms— RTP, Interactive multimedia traffic, Circuit Breaker, WebRTC

I. INTRODUCTION

The wide deployment of browser-based multimedia conferencing applications using the WebRTC protocol [1] is expected to fuel a significant increase in interactive multimedia traffic on the Internet. Unlike streaming video, which can accept a few seconds of buffering, interactive multimedia traffic has very strict latency bounds. Accordingly, it cannot use TCP/IP, and instead relies on RTP [2] over UDP/IP as its media transport protocol. The base RTP specification has little in the way of congestion control, and while the IETF is developing suitable congestion control algorithms, this is expected to be a long-term process, and large-scale WebRTC deployments will occur before it is completed. Design of suitable congestion control methods is an ongoing challenge, as algorithms for TCP (that fill the queues in the network while probing for spare capacity) and those for interactive multimedia (that use delay variation to keep queues in the network small) have conflicting goals. In the short-term, a Circuit Breaker (CB) can provide a necessary performance envelope, within which interactive multimedia traffic can operate, and protect the network from congestion collapse where limited capacity is shared [3].

A Circuit Breaker for RTP (RTP-CB) has recently been proposed within the Internet Engineering Task Force (IETF) [4]. The RTP-CB is designed only to protect the network from excessive congestion. If an RTP media sender that implements the circuit breaker algorithm receives notification of persistent congestion, it will cease transmission to protect the network. Congestion-controlled RTP flows, and uncontrolled RTP flows in lightly loaded environments, operate without triggering the RTP-CB, while misbehaving flows that cause congestion will be terminated. Evidence in the literature supports the ability of the RTP-CB to prevent persistent severe congestion [5][6] using the RTP-CB *congestion rule* of [4]. Alongside this rule, a *media usability rule* for the RTP-CB has been suggested [4], but to date no algorithm has been proposed to implement it.

In this paper we propose an algorithm implementing the media usability circuit breaker rule. We deploy the RTP-CB in a controlled network and evaluate its performance in a range of interactive conferencing scenarios. For each scenario, we report the quality of the received video and the presence of congestion. Based on these results, we show that the proposed media usability rule allows the RTP-CB to stop transmission of flows that are not delivering usable video quality, even if they are not causing severe congestion. Allowing media flows that deliver poor quality to continue is detrimental to network users in general, as such flows use network resources to no avail.

The paper is structured as follows. An overview of the RTP-CB algorithm is given in Section II, the insight derived from [4], [7] allows us to formulate two proposals for improvements to the RTP-CB algorithm in Section III. The experimental setup is described in Section IV. The behaviour of the RTP-CB algorithm, and the impact of the proposed algorithm, when transmitting multiple video flows and coexisting UDP & TCP flows are discussed in Sections V and VI respectively. A summarising discussion concludes the paper in Section VII.

II. RTP CIRCUIT BREAKERS

The RTP circuit breaker operates at the sender side of an interactive RTP session. The sender decides whether to cease transmission based on RTP Control Protocol (RTCP) reception quality reports it receives. There are four rules that can cause the RTP circuit breaker to trigger: RTCP Timeout; Media

Timeout; Network Congestion; and Media Usability. The RTCP timeout and media-timeout rules are straightforward, and detect failed paths. In this work, we focus on the more complex *network congestion* and *media usability* circuit breaker rules [4].

Network Congestion Circuit Breaker: The goal of the congestion circuit breaker algorithm is to stop transmission of RTP flows that cause persistent and severe network congestion. The algorithm runs at the sender, based on information contained in the periodic RTCP Receiver Report (RR) packets sent by receivers. The reporting interval for these packets varies depending on the media rate from a few hundred milliseconds to several seconds.

When losses are detected, the RTP-CB congestion rule determines whether the flow being monitored is overloading the path. The sender makes this decision by comparing its sending rate with the rate that an equivalent TCP flow would attain if it experienced the same round-trip time (RTT) and packet loss rate. The throughput of a TCP flow can be estimated using following TCP throughput equation [4]:

$$X = \frac{S}{R\sqrt{\frac{2bp}{3}} + \left(t_RTO \times 3 \times \sqrt{\frac{3bp}{8}} \times p \times (1 + 32p^2) \right)} \quad (1)$$

where: X is the *TCP-fair* rate (bytes/second); S is the packet size (bytes); R is the RTT (seconds); p is the loss event rate [4]; t_RTO is the TCP retransmission timeout value (seconds; approximated by $t_RTO = 4 \times R$); b is the number of packets acknowledged by a single TCP acknowledgement ($b=1$ is used in practice as many TCP implementations do not use delayed acknowledgements [8]). In most cases, the *loss fraction* approximates the TCP *loss-event rate*, as discussed in [4].

If the actual rate of the RTP media flow is more than one *order of magnitude* (10×) larger than the rate given by Equation (1), then the RTP-CB issues a warning [4]. When the rate given by Equation (1) is exceeded over *three consecutive* reporting intervals, the flow is deemed to be causing persistent congestion and the RTP-CB will trigger, ceasing transmission of the media [4].

Media Usability Circuit Breaker: Applications monitor the packet loss and delay reported in RTCP RR packets to estimate whether the media quality is suitable for the intended purpose. If the media is deemed unusable by the application, then transmission ceases. There is no specified algorithm to determine when media has become unusable [4].

III. MEDIA USABILITY CIRCUIT BREAKER

The *network congestion* rule is the central tenet of the RTP-CB [4] that guarantees protection to the network in case of a misbehaving media flow. The role of the *media usability* rule may seem secondary, only intended to avoid wasting resources, but not required to protect the network. However, it is recognized in [4], [7] that in particular scenarios congestion may not be severe enough for the congestion rule to stop the flow, yet the user experience is too poor to justify diverting network resources from other flows. In this paper, we provide evidence of such scenarios, and propose algorithms to identify

and stop a media flow that delivers unacceptable quality. Our algorithm is formalized in the following proposals:

Proposal 1 (media usability warnings): In a similar manner to the network congestion circuit breaker, we issue a media usability warning for every RTCP reporting interval where the packet losses exceeds a certain threshold. The loss threshold is set based on the application’s quality requirements. For example, our experiments with an unprotected video stream show that 10% losses (over a reporting interval) result in low visual quality. Similarly, a warning is issued for every reporting interval where the delay exceeded a set threshold, decided by the application/user (allowing, for instance, higher values when satellite or inter-continental links are involved). Similar to the congestion rule, the RTP-CB terminates the flow when three consecutive warnings are reported.

Proposal 2 (non-consecutive media usability warnings): Video communication is negatively affected by “flickering” effects, when the quality fluctuates rapidly [9]. Therefore media usability warnings separated by one or two warning-free reporting intervals should provide evidence of low media quality. A flow should be terminated for non-consecutive warnings with a set pattern. For instance, the media usability rule terminates a flow when 3 warnings are accumulated over a 5-interval period.

The flowchart of an RTP-CB algorithm featuring both proposals for the media usability rule (alongside the congestion rule) is shown in Fig 1. In the figure, the trigger condition is set to terminate the flow when three warnings (due to congestion or media usability) are accumulated over `SIZE` consecutive reporting intervals. In our experiments `SIZE` is set to five, corresponding to a RTP-CB decision being reached in three to five reporting intervals from the first evidence of congestion. For persistent severe congestion termination is achieved no later than with the current RTP-CB congestion rule.

IV. EXPERIMENTAL SETUP

We conduct experiments that test the RTP-CB behaviour in scenarios where it must protect the network, or should terminate flows because media usability is severely impaired. The goal of these experiments is two-fold: first, we evaluate when the decisions of the congestion circuit breaker alone are insufficient to meet the usability constraints; second, we highlight how our proposed media usability circuit breaker improves performance.

Our evaluations use an experimental test bed comprising the following elements:

- A server (RTSP/RTP-LIVE555 [10, 11]) acting as a video traffic source.
- A network emulator (Netem [12, 13]) to simulate a network bottleneck with limited capacity and a drop-tail buffer (a token bucket filter and Netem at the LAN interface limit the rate and provide a fixed-sized drop-tail queue).
- Two receivers (A, B), representing the users (VLC clients allow a subjective visual quality assessment).

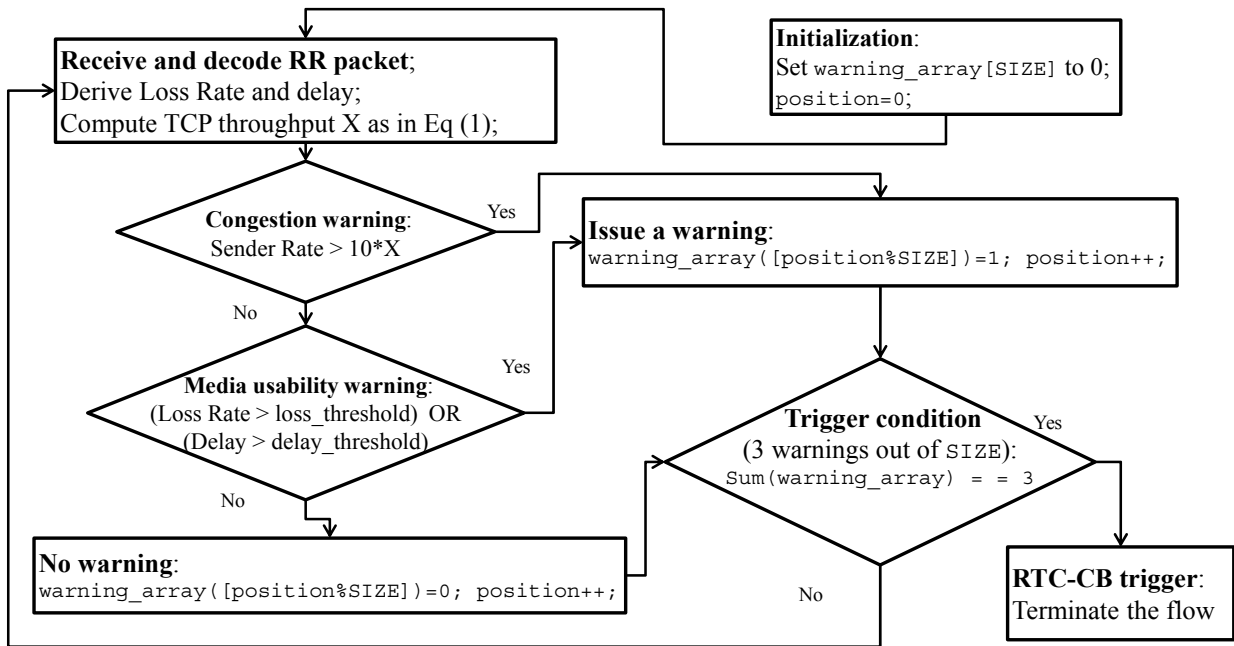


Fig 1. Flowchart of the proposed RTP-CB based on congestion and media usability rules. A loss/delay event that does not suffice to issue a congestion warning may issue a media usability warning. The flow is terminated when three warnings are issued over $SIZE$ consecutive Receiver Reports. If $SIZE=3$ and media usability warnings are not triggered (e.g. if both loss and delay thresholds are exceedingly high) this implementation coincides with the current RTP-CB congestion rule.

As the circuit breaker operates independently in each direction, we focus on traffic flowing in one direction. In all experiments, media flows were sent only from media source to media receiver, with a bottleneck on the forward path carrying media (RTP/UDP) packets, and an uncongested return path carrying RTCP receiver report packets. In order to characterize paths that are, in turn, under-, adequately and over-provisioned, the bottleneck was configured with the following capacities:

- (i) 75% the nominal rate of the video stream(s);
- (ii) the nominal rate of the video stream(s);
- (iii) 150% of the nominal rate of the video stream(s).

The router queue size was expressed as the time taken to drain the bottleneck queue. Following the guidelines for bottleneck links suggested by the IETF [14], the queue sizes used in these experiments were:

- a) Small Buffer: 70ms
- b) Large Buffer: 500ms
- c) Buffer-bloated [15]: 2000ms

Our testbed simulates a moderate-to-low propagation delay between server and client, approximately 50ms. This isolates the effect of queuing delay on the RTT and its implications for RTP-CB behaviour (e.g., the adverse influence of short buffers, discussed later). When the results of our analysis may be affected by this choice, we perform additional experiments simulating shorter and longer propagation delays.

Data were gathered using TCP-Dump, Wireshark, and Tshark, and was analysed with C and Matlab routines. Along with the deployed RTP-CB, we also used an *off-line RTP-CB*, composed of a network logging script (wireshark-like) and a script that calculated the TCP-fairness condition over a window of any set length (the RTP-CB computes that condition every RR interval, which is randomly distributed around a rate-

dependent mean; in our experiments, RR intervals were 5s on average). This *off-line RTP-CB* provided a “ground truth” in our experiments.

We use video traffic sources representative of those obtained with affordable commercial off-the-shelf video conferencing tools or WebRTC systems. Two test sequences are used: “2 People” and “1 Person”. These are encoded using an MPEG4-AVC/H.264 codec operating under the Constrained Baseline Profile, as in Table 1. The two source videos are derived (by cropping and frame sampling) from the high-resolution (1280x720), high-frame-rate (60 frame/s) video-conference-style sequence “4 People” [16].

Table 1: Video conference style sequences

	Video	Format	Bit rate	Frame/s	Frame size
S1	1 Person	Mp4	100kb/s	15	320x160
S2	2 People	Mp4	500kb/s	30	640x320

A bursty transmission profile, typical of predictive coding of video, is observed for the two sequences in Table 1. The nominal rate is approximately respected when averaged over intervals of 5 seconds, but when the transmission profile is examined over shorter intervals, rate spikes are found that align with the periodic pattern of I-frames (approximately one per second; a typical choice for real video conferencing systems). This transmission profile is normal in applications using predictive coding of video, but as we see later, it can impact media usability in environments with buffer constraints.

V. PERFORMANCE WITH COMPETING MULTIMEDIA FLOWS

The objective of the experiments described in this section is to determine whether the RTP-CB decisions, based on the network congestion rule alone, are adequate from the user experience standpoint. When this is not the case, we seek to

confirm whether the proposed media usability circuit breaker rules are able to identify a video flow that should be stopped.

We consider two users, A and B, engaged in independent live multimedia sessions. The flows directed to A and B share a bottleneck link with buffer size as in Section IV. Link capacity is expressed as a function of the total video rate (A+B). Video flows start within five seconds of each other, with randomly selected offset to avoid synchronization of the bursts. Performance is observed to depend on the router buffer size:

Small Buffer (70ms): Loss rates are high when the bottleneck capacity is below, or equal to, the aggregate video rate. The short delay counterbalances the loss rate when evaluating Equation (1), thus the RTP-CB gathers insufficient evidence of ongoing severe congestion. Using the network congestion rule alone, no circuit breaker trigger is observed for either client; warnings, when issued, are non-consecutive. However, the packet loss causes both clients to suffer poor video quality, as summarized in Table 2.

The RTP-CB behavior when our proposed media usability rule is implemented is summarized in Table 3. The high loss rates reported in nearly every RR interval induces warnings that quickly cause the media usability RTP-CB to trigger, typically within 15 seconds of flow coexistence.

When the bottleneck link has higher capacity than the aggregate video rate, residual losses (due to bursts occasionally overflowing the buffer) are present, but are few and far apart, and do not seriously affect the video quality. Neither RTP-CB rule triggers in these cases.

Large Buffer (500ms): When the bottleneck capacity is below the aggregate video rate, high loss rates are reported. Compared to the previous scenario, the increased RTT (due to the larger buffer) allows the RTP-CB congestion rule to classify reported loss events as representative of severe congestion and terminate the flow, as summarized in Table 2. However, in 10% of the experiments the warnings were non-consecutive and the network congestion circuit breaker did not trigger for any client, even though the visual quality was poor. Instead, when our media usability rule is implemented, the RTP-CB of at least one flow triggered in each experiment, typically within 25 seconds of coexistence, allowing the other client to achieve good quality. Results are summarized in Table 3.

When the bottleneck capacity was 100% of the aggregate rate, flows experience significant and periodic loss only when bursts of the two flows overlap. In most such cases, the network congestion circuit breaker terminates at least one flow. In some cases, however, frequent but non-consecutive warnings were observed; the congestion rule did not trigger, and both clients endured poor visual quality. In such cases, our proposed media usability rule terminates the flow.

No loss was observed for an over-provisioned link.

Buffer-bloated (2000ms): Video quality is predictably poor for both clients when the path capacity is below the aggregate rate. The loss rate and delay are high for both clients, and the network congestion circuit breaker triggers at least for one client in all experiments; in many cases both flows are stopped at the same time. Adding our proposed media usability circuit breaker accelerates the termination of a flow. Loss is not

expected when the path capacity equals the total video rate. However, rate fluctuation overlaps did occasionally lead to packet loss that triggered the network congestion circuit breaker due to the large RTT. There were no losses for an over-provisioned link.

We repeated the experiments for a 70ms bottleneck buffer, this time using a range of values for the propagation delay. With propagation delay below 300 ms, the RTP-CB congestion rule allows a video flow to continue despite high losses and poor video quality, similarly to the findings in Table 2.

Further experiments were conducted with clients requesting different video content (A uses S1, B uses S2 as in Table 1; tests were repeated with clients swapping video content). The results confirm the trend seen in Table 2 and Table 3.

Table 2: Summary of results for two competing flows; path propagation delay is 50ms; overall video rate is $Z=A+B$ ($A=B=500$ kbps). The RTP-CB implements only the congestion rule.

Buffer		70ms	500ms	2s
Capacity				
75% Z	A	AvgLoss \approx 25% Very poor quality No RTP-CB trigger.	AvgLoss \approx 20% 40% RTP-CB trigger.	AvgLoss \approx 20% 80% RTP-CB trigger.
	B	AvgLoss \approx 25% Very poor quality No RTP-CB trigger.	AvgLoss \approx 20% 40% RTP-CB trigger.	AvgLoss \approx 20% 80% RTP-CB trigger.
100% Z	A	AvgLoss \approx 15% Very poor quality No RTP-CB trigger.	AvgLoss \approx 5% 20% RTP-CB trigger.	AvgLoss \approx 1-10%, 15% RTPCB trigger.
	B	AvgLoss \approx 15% Very poor quality No RTP-CB trigger.	AvgLoss \approx 5% 20% RTP-CB trigger.	AvgLoss \approx 1-10%, 15% RTPCB trigger.
150% Z	A	AvgLoss \approx 3% No RTP-CB trigger.	No Loss	No Loss
	B	AvgLoss \approx 3% No RTP-CB trigger.	No Loss	No Loss

Table 3: Summary of results for two competing flows; path propagation delay is 50ms; overall video rate is $Z=A+B$ ($A=B=500$ kbps) after RTP CB implementation. The RTP-CB implements both Congestion and the media usability rules.

Buffer		70ms	500ms	2s
Capacity				
75% Z	A	AvgLoss \approx 20% Very poor quality 100% RTP-CB trigger.	AvgLoss \approx 20% 60% RTP-CB trigger.	AvgLoss \approx 20% 80% RTP-CB trigger.
	B	AvgLoss \approx 30% Very poor quality 100% RTP-CB trigger.	AvgLoss \approx 20% 60% RTP-CB trigger.	AvgLoss \approx 20% 80% RTP-CB trigger.
100% Z	A	AvgLoss \approx 15% Very poor quality 100% RTP-CB trigger.	AvgLoss \approx 5% 35% RTP-CB trigger.	AvgLoss \approx 1-10%, 15% RTPCB trigger.
	B	AvgLoss \approx 15% Very poor quality 100% RTP-CB trigger.	AvgLoss \approx 5% 35% RTP-CB trigger.	AvgLoss \approx 1-10%, 15% RTPCB trigger.
150% Z	A	AvgLoss \approx 4% No RTP-CB trigger.	No Loss	No Loss
	B	AvgLoss \approx 1% No RTP-CB trigger.	No Loss	No Loss

VI. TCP AND MULTIMEDIA FLOWS

In this scenario, interactive video flows and TCP flows share a bottleneck. The aim of the experiments is to highlight cases where our proposed media usability algorithm, defined in

Section III, improves the RTP-CB performance accounting for both network congestion and quality of user experience.

We consider an interactive video flow competing with a long-lived TCP flow. The video flow is sequence S2 from Table 1. A bulk TCP Cubic flow [17] provides background traffic (e.g., FTP). The video and TCP flows share a common bottleneck with capacity equal to 120% and 200% of the nominal video rate; these correspond to a case where the RTP flow occupies most of the bottleneck, and to a scenario where half of the resources are available to another flow (TCP). As in Section V, performance is observed to depend on the router buffer size:

Small Buffer (70ms): The short buffer (and low propagation delay) causes the network congestion circuit breaker to categorize the loss events as not being evidence of severe congestion. Despite the additional losses induced by TCP probing for capacity, the circuit breaker warnings, if issued, are non-consecutive and the network congestion circuit breaker does not trigger. This behaviour (high loss rates, no RTP-CB trigger) is observed for both link provisions, as summarized in Table 4. Even though the circuit breaker does not trigger, the high loss rates induce extremely poor video quality. As shown in Table 5, when we introduce the media usability circuit breaker, it quickly terminates the video flow, allowing network resources to be reclaimed by the TCP flow. The same conclusions apply whichever flow (RTP or TCP) starts first.

Large Buffer (500ms): The path capacity was set to 120% of the video rate. The RTP sender uses most of the bottleneck capacity. The TCP flow starts and probes for capacity, quickly claiming a share; this results in losses for both flows. TCP reacts to losses, but tries to keep the buffer full, inducing a delay that suffices to trigger the RTP-CB in almost half of the experiments. The remaining flow (TCP) fills the bottleneck. We restarted the RTP flow and both flows encountered loss, with TCP adapting its rate and allowing RTP to become established. TCP regularly probes for capacity, filling the buffer and eventually triggering the RTP-CB. The video quality was poor in all experiments. However, most events were non-consecutive and did not violate the RTP-CB (congestion) rule in 60% of the experiments. Implementation of the proposed media usability rule stops such flows, as shown Table 5.

The experiments were repeated with a path capacity at 200% of the video rate. The two flows should coexist together. However, as TCP probes the path for capacity, it will induce bottleneck losses (and high delay). In these cases, the RTP rate is found to be close to the limit set by the RTP-CB congestion rule. In less than half of the experiments the RTP-CB terminated the video flow, allowing TCP to gain the entire link. The visual quality was noticeable affected when TCP probed the path inducing loss. The media usability rule we propose complements the congestion rule and stops flows when loss events persist.

Buffer-bloated (2000ms): For the smaller bottleneck capacity, the RTP-CB always triggers when TCP starts first, because the TCP flow fills the buffer and increases the delay. RTP-CB triggered in 50% of the cases when TCP started after RTP. The visual quality was low, as shown in Fig 2.



Fig 2. Sample video frame with a path with 120% of the video rate and a 2s bottleneck buffer

Increasing the capacity of bottleneck did not change the result, with TCP probing aggressively and the RTP-CB (congestion) terminating the RTP flow in most cases. The media usability rule we propose helps reaching this decision faster and for all flows with poor quality. Media usability warnings are issued due to the reported losses as well as the long delay of RTCP RR packets (as cubic keeps the buffer almost full).

Table 4: Summary of results for TCP and UDP competing flows. The RTP-CB implements only the congestion rule

Buffer		70ms	500ms	2s
Capacity				
120% video	TCP First	AvgLoss >20% No video display No RTP-CB trigger	AvgLoss >10%, Poor quality, 70% RTP-CB trigger.	AvgLoss ≈10% 100% TP-CB trigger
	RTP First	AvgLoss >20% No video display No RTP-CB trigger	AvgLoss 10%, Poor quality 70% RTP-CB trigger.	AvgLoss ≈10% 50% RTP-CB trigger
200% video	TCP First	AvgLoss >10% No video display No RTP-CB trigger	AvgLoss ≈10% Poor quality 30% RTP-CB trigger.	AvgLoss ≈5% 80% RTP-CB trigger
	RTP First	AvgLoss >10% No video display No RTP-CB trigger	AvgLoss ≈10% Poor quality 30% RTP-CB trigger.	AvgLoss ≈5% 70% RTP-CB trigger

Table 5: Summary of results for TCP and UDP competing flows. The RTP-CB implements both Congestion and the media usability rules.

Buffer		70ms	500ms	2s
Capacity				
120% video rate	TCP First	AvgLoss >20% No video display 100% RTP-CB trigger	AvgLoss >10%, Poor quality, 100% RTP-CB trigger	AvgLoss ≈10% 100% RTP-CB trigger
	RTP First	AvgLoss >20% No video display 100% RTP-CB trigger	AvgLoss 10%, Poor quality 100% RTP-CB trigger	AvgLoss ≈10% 100% RTP-CB trigger
200% video rate	TCP First	AvgLoss >10% No video display 100% RTP-CB trigger	AvgLoss ≈10% Poor quality 30% RTP-CB trigger.	AvgLoss ≈5% 90% RTP-CB trigger
	RTP First	AvgLoss >10% No video display 100% RTP-CB trigger	AvgLoss ≈10% Poor quality 30% RTP-CB trigger.	AvgLoss ≈5% 90% RTP-CB trigger

Further experiments were performed considering TCP New Reno [18] instead of Cubic. The results for TCP New Reno are in line with those shown in Table 4 and Table 5. The main difference is observed in the case of a large buffer. TCP New Reno is less aggressive than Cubic and, when the path is twice the video rate, both flows are sustained. Video quality was mostly good, but the delay imposed by a consistently full buffer impairs interaction. If the media usability rule is set, by

the application/user, to issue warning for delays over (say) one second, the RTP-CB terminates such flows.

These experiments reinforce the need for our proposed media usability rule (Proposals 1 & 2) for the short and large buffer cases, confirming results reported in the previous sections. Furthermore, the buffer-bloated case discussed above lends additional support to our media usability rule, which stops a video flow that consistently exceeds an application-specific delay threshold (for instance one second) to ensure a viable interactive communication. With TCP keeping a bloated buffer nearly full at all times such quality-impairing delay are inevitable, and yet TCP may cause too few (and far between) losses to violate the RTP-CB congestion rule.

VII. DISCUSSION AND CONCLUSIONS

Our evaluation of the RTP-CB behaviour contributes to the currently limited experimental assessment of the RTP-CB available in the literature [5]. For sample video sessions (either allowed or terminated by the RTP-CB), we consider the video quality, along with bandwidth usage/share, to assess the effectiveness of the RTP-CB. While the main purpose of the RTP-CB is firmly on protecting the network from congestion, considerations must be paid to the quality of the video that the RTP flow is serving. This is especially important as the RTP-CB defines the envelope within which interactive multimedia congestion controls should operate.

Our experiments highlighted cases where packet losses or delay seriously impair the media session quality, despite the flow not severely congesting the network (and so not triggering the RTP-CB congestion rule). We identify several limitations of the RTP circuit breaker. For example low propagation delays (in the order of 50 ms) coupled with a small buffer (70 ms), yields low RTT values for the RTCP RR packets used by the RTP-CB congestion rule. Despite non-negligible packet losses, the TCP-rate (1) can achieve a high rate (close or above the video rate) when the RTT is low. Therefore the RTP-CB rarely acquired enough evidence of severe congestion to trigger. From a media usability perspective, however, the flow should be stopped, as the reported loss rates (typically above 10%) correlate with poor video quality in our experiments (with unprotected data).

In Section III we propose an algorithm for the RTP-CB media usability rule that stops low-quality flows for the scenarios above. Our algorithm, described by the flowchart of Fig 1, complements the RTP-CB congestion rule. The first step (Proposal 1) is to issue media usability warnings when the reported losses (or delay) exceed a threshold set by the application. For example, 10% losses for unprotected video, or a few hundred milliseconds delay (higher with satellite or inter-continental links) for interactive communication. The second step (Proposal 2) is to stop the flow when strings of warnings are received. These can be either media usability warnings or congestion warnings. Warnings need not be consecutive; we consider terminating a flow that accumulates three warnings for five consecutive RTCP RR packets, as this was consistent

with poor user experience. An effective rate-control algorithm, if present, should avoid congestion hence prevent sustained warnings; alternatively, the user can be prompted for action (e.g. disable or reduce the quality of the video) when one or two warnings are accumulated. If rate-control or user-initiated reactions are insufficient or absent (as in our experiments) the RTP-CB terminates the flow to protect other network users.

We confirmed the effectiveness of the proposed implementation of the media usability rule in our experiments. These included video and TCP flow coexistence. In addition to the short and large buffer scenarios mentioned before, media usability plays an important role when a TCP and a video flow share a bottleneck characterized by bloated buffer. In this case TCP keeps a nearly full buffer most of the times inducing unacceptable delay to an interactive communication flow, but causing too few and sparse losses to trigger the congestion rule.

VIII. ACKNOWLEDGMENT

We would like to thank Dr Raffaello Secchi from ERG for his valuable help with the experimental setup. This research was supported in part by the RCUK DE award to the dot.rural Digital Economy Hub: EP/G066051/1.

IX. REFERENCES

- [1] C. Perkins, M. Westerlound, J. Ott, "Web Real-Time Communication (WebRTC): Media Transport and Use of RTP - IETF draft 21," November 2014.
- [2] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson, "RTP: A transport protocol for real-time applications-RFC3550," July 2003.
- [3] G. Fairhurst, "Network Transport Circuit Breakers - TSVWG Working Group-IETF Draft -01," May 5, 2014.
- [4] C. Perkins and V. Singh, "Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions - IETF draft-05," February 14, 2014.
- [5] V. Singh, S. McQuistin, M. Ellis and C. Perkins, "Circuit breakers for multimedia congestion control," in *Proc. of IEEE Packet Video*, 2013, .
- [6] Z. Sarker, V. Singh and C. Perkins, "An evaluation of RTP circuit breaker performance on LTE networks," 2014.
- [7] N. Fough, F. Verdicchio and G. Fairhurst, "Exploration of RTP circuit breaker with applications to video streaming," in PGnet 2013, Liverpool, UK.
- [8] S. Floyd, M. Handley, J. Padhye and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification-RFC5348," September 2008.
- [9] S. Floyd, M. Handley, J. Padhye and J. Widmer, *Equation-Based Congestion Control for Unicast Applications*. ACM, 2000.
- [10] H. Schulzrinne, A. Rao and R. Lanphier, "Real Time Streaming Protocol (RTSP) - RFC 2326," April 1998.
- [11] *Live Networks, Inc. LIVE555 Streaming Media*. Available: <http://www.live555.com/liveMedia/>.
- [12] A. Keller, "Manual to Packet Filtering and netem," July 20, 2006.
- [13] *Netem*. Available: <http://www.linuxfoundation.org/collaborate/workgroups/networking/netem>.
- [14] V. Singh and J. Ott, "Evaluating Congestion Control for Interactive Real-time Media-RMCAT WG draft- 04," October 20, 2013.
- [15] J. Gettys and K. Nichols, "Bufferbloat: Dark buffers in the internet," Queue, vol. 9, pp. 40, 2011.
- [16] A set of video test sequences; a representative trace for video conference video samples. Available: <http://www.erg.abdn.ac.uk/~gorry/video/video-test-sequences>
- [17] S. Ha, I. Rhee and L. Xu, "CUBIC: a new TCP-friendly high-speed TCP variant," *ACM SIGOPS Operating Systems Review*, vol. 42, pp. 64-74, 2008.
- [18] m. Allman, V. Paxson and E. Blanton, "TCP Congestion Control-RFC5681," September 2009.