2014 IEEE INFOCOM Workshops: 2014 IEEE INFOCOM Workshop on Communication and Networking Techniques for Contemporary Video

An Evaluation of RTP Circuit Breaker Performance on LTE Networks

Zaheduzzaman Sarker
Ericsson Research, Sweden
Email: zaheduzzaman.sarker@ericsson.com

Varun Singh
Aalto University, Finland
Email: varun.singh@aalto.fi

Colin Perkins
University of Glasgow, UK
Email: csp@csperkins.org

***Abstract*—Real-time multimedia comprises a large, and growing, fraction of mobile data traffic. An important subset of such flows are from interactive conferencing applications using RTP on UDP/IP to reduce latency. UDP has no congestion control, and while the IETF is developing RTP-level congestion control algorithms as part of the WebRTC standards, these will take time to finalise and deploy. In the interim, we proposed an RTP circuit breaker to the IETF. This can detect and stop RTP flows that cause excessive network congestion, acting as an envelope within which a congestion control algorithm can operate. We briefly review the design of the RTP circuit breaker, then present an initial evaluation of its performance on LTE networks. Our results show that the algorithm is conservative in overload situations with low delay and high loss. Such situations can occur due to active queue management in LTE networks. We propose changes to the RTP circuit breaker to better suit such networks, by using a different TCP throughput model that is more sensitive to the observed packet loss patterns.**

I. Introduction

The Real-time Transport Protocol (RTP) [1] is widely used as a basis for telephony, video conferencing, and telepresence. RTP traffic from such applications often uses best-effort wireless networks, such as those using the 3rd Generation Partnership Project (3GPP) Long-Term Evolution (LTE) standard. Fading, interference, mobility, hand-overs, cell loading and other factors cause the capacity available for each user in such networks to fluctuate. Multimedia applications producing high data rates hence have the potential to cause network congestion. This can disrupt both the users' quality of experience, and potentially the operation and stability of the network for other customers.

The long-term solution is to deploy multimedia congestion control algorithms. These have been widely studied (e.g., [2], [3], [4], [5], [6]), and the Internet Engineering Task Force (IETF) is currently developing standards in this area. This is expected to be a multi-year process. With likely wide deployment of new applications based on the WebRTC standards [7], there is an urgent need for a circuit breaker to prevent flows from causing excessive congestion in the short term, before effective congestion control is deployed. New congestion control algorithms defined in the IETF will need to work inside the envelope of this circuit breaker [8]. Accordingly, it cannot be too aggressive in terminating media flows, since it should allow time for the congestion control to respond. We have proposed such a circuit breaker for unmodified RTP applications [9], and its development is on a tight schedule to be ready for inclusion in the initial WebRTC framework.

In this paper, we briefly review the RTP circuit breaker design, then present an initial evaluation of its performance on LTE networks. Our evaluation shows that the RTP circuit breaker can be overly conservative in such environments when active queue management (AQM) is used. We suggest improvements to the RTP circuit breaker that better suit it to the LTE environment by using a different TCP throughput model that is more sensitive to the observed packet loss patterns. The RTP circuit breaker is described in detail in [9] but that includes no performance evaluation. Previous work [10] outlines the performance of the circuit breaker in residential networks using trace-driven simulation, but does not consider LTE networks. Ours is the first evaluation of the proposed circuit breaker in LTE networks.

We structure the remainder of this paper as follows. We review the circuit breaker design in Section II and properties of LTE networks in Section III. We describe performance of the circuit breaker on LTE in Section IV. Section V discusses circuit breaker improvements for such networks. Related work is outlined in Section VI, and we conclude in Section VII.

II. RTP Circuit Breaker Algorithm

RTP-based multimedia applications typically send media data over UDP, and are subject to the unpredictable behaviour of best-effort IP networks, including packet loss, reordering, and variable queuing delays. Real-time applications are tolerant to some amount of packet loss, either concealing the loss or using one of the available error resilience mechanisms [11], but variation in delay causes the media to freeze and skip frames, and can disrupt the user experience [12]. Real-time applications hence either need to implement congestion control, or use a transport that implements congestion control, to avoid causing queues in the network and hence delay. This motivates the need for congestion control, and the circuit breaker.

Effective multimedia congestion control requires the underlying codec to produce the requested media bit rate in a timely manner, so the application can adapt the sending rate to available network capacity while attempting to maintain the user experience [13]. The interaction between the multiple control loops makes congestion control for real-time multimedia inherently a difficult problem. Many congestion control and media adaptation algorithms have been proposed, but there is no general consensus on the correct approach. For example, TCP is only suitable for interactive multimedia on a low RTT path

($<$ 100ms) [14], DCCP has problems with NAT traversal [15], and current WebRTC implementations have stability problems [16]. We do not propose a new congestion control algorithm, but describe minimal conditions under which there is general agreement that an RTP flow is causing serious congestion and should cease transmission. This is one strand of work in IETF [9], in parallel with new congestion control algorithms.

The circuit breaker we propose uses RTP Control Protocol (RTCP) [1] mechanisms. We rely solely on information in RTCP Sender Report (SR) and Receiver Report (RR) packets to detect if a media flow is overusing the available capacity and causing congestion. For wide applicability, extensions such as [17] are not used. The available congestion indicators in RTCP are the network round trip time (RTT), average packet timing jitter, packet loss fraction measured over the last reporting interval, and the cumulative number of packets lost since the start of the session. Delay-based congestion control algorithms can use RTT variation as a congestion indicator, but RTCP reporting intervals are large, typically several seconds, and a single RTT estimate per interval is too infrequent to provide useful input to a circuit breaker. Likewise, a single highly aggregated jitter measurement per reporting interval is also not useful as a circuit breaker. Loss due to queue overflows is a strong indicator of congestion in some networks, but in wireless networks loss can occur due to bit-error corruption from signal interference, reducing its effectiveness as a congestion signal. The RTP circuit breaker conditions are as follows [9]:

- The **Media timeout circuit breaker** triggers if an end-point is sending media, but the returning RTCP packets have a non-increasing highest sequence number received field for two consecutive reporting intervals. Corresponds to forward path failure.
- The **RTCP Timeout circuit breaker** triggers if an end-point is sending media, but has received no corresponding RTCP SR/RR packets for two consecutive RTCP reporting intervals. This corresponds to a failure of the reverse path.
- The **Congestion circuit breaker** triggers if the RTP media sending rate exceeds the estimated TCP throughput over the same path by a factor of ten or more for two consecutive RTCP intervals. The estimated TCP throughput, T, is calculated as follows [18]:

  $$T = \frac{s}{R\sqrt{2p/3}} \quad (1)$$

  The packet size, s, RTT estimate, R, and loss fraction over the last reporting interval, p, are obtained from RTCP. If [19] is in use, ECN-CE marks are reported and included in the loss fraction and congestion circuit breaker.

The first two circuit breakers are intended to roughly correspond to TCP timeout conditions; the last is directly congestion related. The congestion circuit breaker is based on the principles used in the TCP-Friendly Rate Control (TFRC) protocol [3], which similarly uses a rate-based model using the TCP throughput equation as the basis for its operation. The congestion circuit breaker is known to be a poor quality estimate, and Section V discusses the effect of using the more complete TCP throughput model from [20] as an alternative. Other limitations of the congestion circuit breaker include the need to use of the loss fraction as the estimate for p, when calculating TCP throughput, rather than the loss event rate, since loss event rate is not reported. The constraints of RTCP limit the accuracy of the throughput estimate we can derive; we must work with unmodified RTCP to be useful when inter-working with legacy systems. Full details of the RTP circuit breaker algorithm are in the IETF draft [9].

III. LTE Network Properties

The Long-Term Evolution (LTE) is a 3GPP radio access standard, defined with a vision of an all-IP network. It can achieve data rates up to 100 Mbps, with capacity and coverage comparable to the voice services of the circuit switched radio network. 3GPP has standardized multimedia telephony over IMS (IP Multimedia Subsystem) [21] as a pure packet-switched real-time communication service. Despite robust retransmission procedures [22] and link adaptation techniques [23], a highly loaded mobile network, or a user positioned in bad radio coverage cannot be guaranteed perfect transport characteristics all the time. Consequently, there is a need for applications to monitor and adapt to the current transport characteristics to optimize the quality, and keep the service attractive.

In an LTE cell, capacity is shared among active users, and between different types of services currently in use. A typical wireless channel in LTE uses the default bearer, which does not guarantee a given capacity. Mobility, hand-overs, interference, fading, etc. affect throughput available to the user, and channel capacity can vary by more than an order of magnitude in less than a second. If the users/applications send more data than the network can sustain then packets are queued in the network. If AQM is deployed, packets that are queued for a longer duration are dropped. Consequently, as the queuing delay becomes large, packets are dropped more frequently. If AQM is not deployed, and depending on the size of the queue, the packets can be queued for several seconds. For real-time interactive multimedia communication both loss and delay are unwanted artefacts, but loss is preferred since, unlike delay, it can be concealed by the receiver.

LTE also supports Quality of Service (QoS) enabled bearers [24] that handle different traffic types according to their respective QoS class. QoS-enabled bearers are configured to provide some minimum resource guarantee in extreme cases to, e.g., guarantee a minimum bit-rate when users move into a bad coverage area. However, in a congested network maintaining high media bit rate is expensive, and both congestion control and circuit-breakers are required features.

IV. Performance on LTE Networks

We evaluate the performance of the circuit breaker in an LTE network. Section IV-A describes the simulation environment and the evaluation scenarios. Section IV-B outlines the baseline performance without the RTP circuit breaker, and Section IV-C describes the performance of the RTP circuit breaker algorithm.

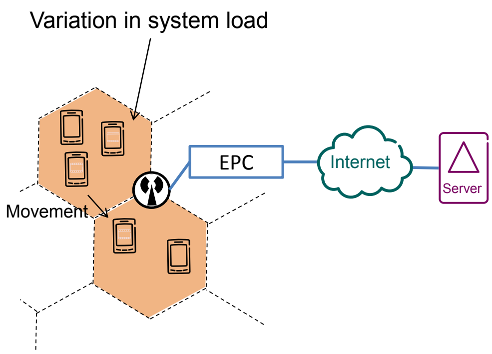

Figure 1. LTE Simulation Environment

A. Simulation Environment and Scenarios

We use Ericsson's LTE simulator to evaluate impact of the circuit breaker on multimedia calls in an LTE network. This simulates a multi-user radio system, with detailed models of the LTE physical and protocol layers, and radio resource management. The simulation set-up is similar to 3GPP Case 1 [25], but to reduce complexity, we only simulate the downlink with different loads, and assume an uncongested uplink. Endpoints simulate movement at 3 km/hour and can move between cells of the LTE network. A server sends media downstream to each user; media traffic is modelled as real-time conversational video with a nominal bit rate of 1.5 Mbps, 30 frames/sec, with each flow having 30s duration. As audio is a small fraction of the multimedia traffic, we simulate only video flows. Flows are simulated over the default LTE bearer.

We load our simulated LTE network by increasing the number of active users, modelling arrival as a Poisson process. We chose three representative load scenarios. These are *moderately loaded*, *highly overloaded* and *extremely overloaded*. The moderately loaded network is busy but not overloaded. As load increases the system becomes highly (almost five time more users than moderate load) and extremely overloaded (almost ten times more users than moderate load). While such high load should not occur in real deployments, due to admission and policy control, we simulate higher load to observe behaviour of the circuit breaker and ensure it's effective in network failure scenarios. Figure 1 shows our simulation environment.

Since we are only simulating congestion on the downlink, RTCP packets sent on the uplink do not observe congestion, therefore, the RTCP timeout breakers will not trigger.

B. Baseline Performance

We consider impact of increasing load on the performance of media flows in an LTE network. This scenario does not implement the circuit breaker, so the results serve as a baseline for comparison with results in Section IV-C. Figure 2 shows the cumulative distribution function (CDF) of average and 98th percentile end-to-end video frame delay, for each load scenario. Under moderate load, only 20% of users saw 98th percentile end-to-end delay more than 200 ms, and average delay was less than 150 ms, resulting in a usable system. However, with high and extreme load, all users saw more than 200 ms end-to-end delay, making media unusable for interactive applications.

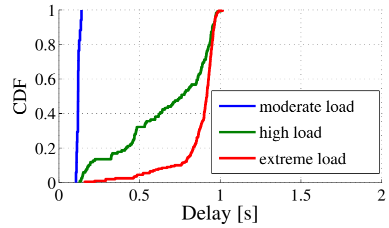

(a) CDF of average video frame delay

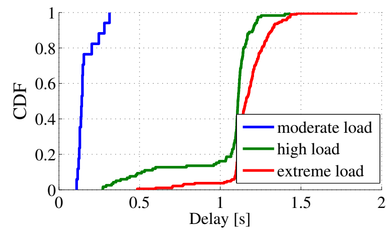

(b) CDF of 98th percentile of video frame delay

Figure 2. Video frame delay without the RTP circuit breaker

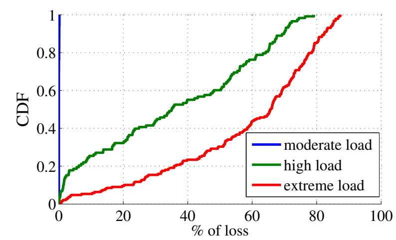

Figure 3. CDF of packet loss rate over all users; circuit breaker disabled

The high latency mainly occurs due to the network queues in the base station and indicates a heavily congested network. The figures also show that the AQM limits the average delay to about 1s irrespective of load. Consequently, when the queues overflow, we observe high Packet Loss Rate (PLR) as shown in the high- and extreme-load cases of Figure 3. In the moderately loaded scenario both delay and loss are within acceptable bounds, while in highly and extremely loaded scenarios the flows are unusable due to high delay and loss ratios. This shows the simulations adequately model the cases of interest, where the moderately loaded scenario matches a busy but usable LTE network, and other scenarios represent overloaded networks.

C. Performance with RTP Circuit Breaker

We repeated the simulations from Section IV-B after enabling the RTP circuit breaker. Figure 4 shows the fraction of sessions triggering a circuit breaker under varying load. No circuit breaker was triggered for moderately loaded scenarios, as expected due to low loss rate and end-to-end delay. As load increases, the fraction of sessions triggering the circuit breaker increases, but does not exceed 60% even when the network is highly congested and when we would expect the circuit breaker

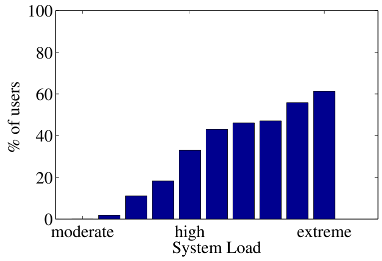

Figure 4. Fraction of sessions triggering circuit breaker vs. load

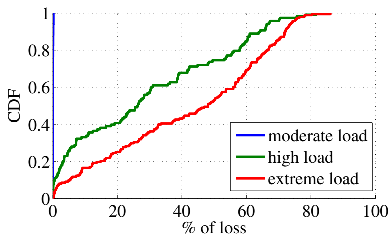

Figure 5. CDF of packet loss rate over all users; circuit breaker enabled

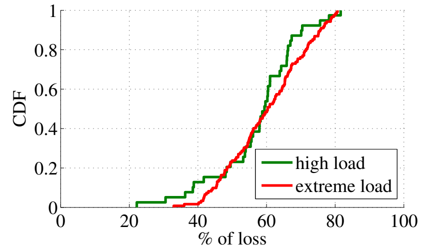

(a) Sessions triggering the RTP circuit breaker

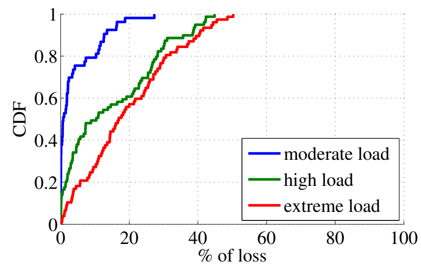

(b) Sessions that did not trigger the circuit breaker

Figure 6. CDF of packet loss rate with the RTP circuit breaker enabled

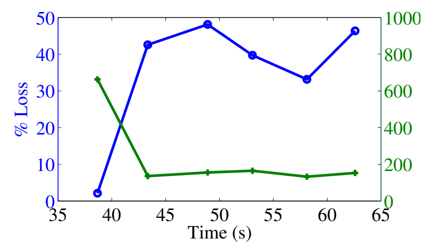

Figure 7. Packet loss fraction (blue) and TCP throughput (green) for a flow that did not trigger the circuit breaker

to trigger (since Section IV-B shows the media is unusable for real-time communication). All flows were terminated by the congestion circuit breaker, not media timeout. The end-to-end frame delay was bounded by 1 second in all cases, consequently the maximum age of a packet will not exceed this limit. Thus, with AQM in place it is unlikely that the burst losses will occur for multiple regular RTCP reporting interval on an LTE network. This justifies the no triggering of the media timeout circuit breaker in the simulated LTE deployment.

Figure 5 shows the CDF of the packet loss rate with the circuit breaker enabled. As expected, and similar to the simulations without the circuit breaker, the moderately loaded sessions are essentially loss free. Compared to Figure 3, we see a slightly lower loss rate in the high load scenarios, and a significantly lower loss rate in the extreme load scenarios. This is expected since enabling the circuit breaker terminates a significant fraction (40–60%) of the flows, hence reducing the load for the remaining flows, leading to reduced loss rates.

Figures 4 and 5 highlight the main issue with the circuit breaker in LTE networks: it does not trigger in 40% of cases of extreme load, or in a significant number of high load cases, but most of those sessions suffer loss rates making the media unusable. To understand this, we plot the CDF of the packet loss rate for terminated and non-terminated media sessions. In Figure 6a, where sessions were terminated by the circuit breaker, packet loss is very high, hence the circuit breaker correctly terminated these sessions. However, in Figure 6b, where sessions were not terminated by the circuit breaker, the losses, while lower, are still high. These sessions could reasonably be expected to trigger the circuit breaker.

To understand why high and extreme load did not trigger the circuit breaker, we consider a representative trace in Figure 7, showing the fraction lost reports from RTCP. The loss fraction approaches 50% for much of the simulation, and the trace has unacceptable media quality. The TCP throughput estimate drops as the packet loss rate increases, but not so far that the 1.5 Mbps media rate exceeds the $10\times$ threshold needed to trigger the circuit breaker. Figure 8 shows the number of lost packets over time for the same trace. There are bursts of significant loss in quick succession, so *why does the TCP rate remain high enough that the circuit breaker doesn't trigger?*

The answer lies in Figure 9, which shows the RTT samples for that specific flow. The RTT estimate, R, is about 0.1 seconds. Considering the TCP throughput calculation in Equation 1, we see the throughput is inversely proportional to R and the square root of the loss rate, p. Figure 10 illustrates, comparing TCP throughput with varying p and R in the range 100–500ms with the 1.5Mbps media rate and 1.5Mps $\times$ 0.1 = 0.15Mbps TCP throughput below which the circuit breaker triggers. At 50% loss rate, a 100ms RTT held down by the use of AQM is simply too low to trigger the circuit breaker with a $10\times$ multiple between TCP throughput and media rate (indeed, as Figure 10 shows, the throughput curve for R=100ms only drops below the 0.15Mbps circuit breaker trigger threshold for packet

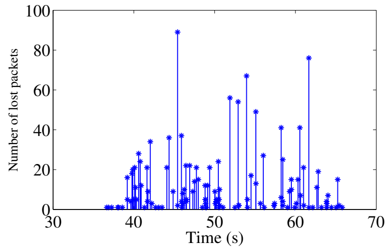

Figure 8. Packets lost over time for the flow from Figure 7

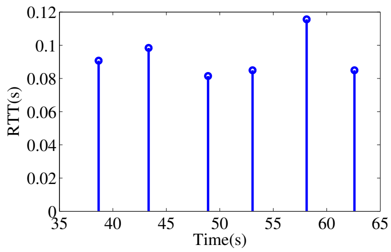

Figure 9. Reported RTT samples for flow from Figure 7

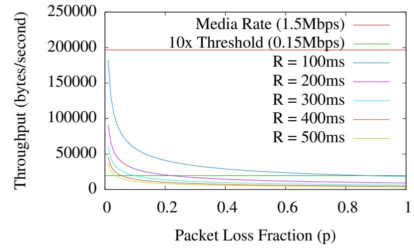

Figure 10. TCP Throughput as loss fraction and RTT vary

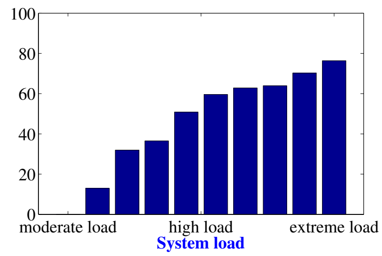

Figure 11. Sessions triggering circuit breaker vs. load; Padhye TCP equation

loss rates $\geq 88\%$; increasing RTT reduces the packet loss rate at which the throughput curves drops below the threshold). In an LTE network with AQM that strives for low end-to-end delay, it is possible to have low RTT at the cost of high packet loss, but this interacts poorly with the simple TCP throughout equation.

V. Enhancing the RTP Circuit Breaker

Our observation that sessions with high loss rate were not terminated leads us to believe the circuit breaker needs to be more sensitive to the reported packet loss rate in LTE networks. There are two ways to do this. We can reduce the $10\times$ multiple between media rate and TCP throughput, perhaps to $5\times$, or we could change the TCP throughput equation used to be more sensitive to loss. Figure 10 shows that changing the multiple to $5\times$ would improve the situation, with more sessions being below the rate that triggers the circuit breaker, but it is clear it would not solve the problem, for example even at 100ms RTT there are still numerous flows that do not trigger the circuit breaker.

Accordingly, we consider more sophisticated TCP model of Padhye et al. [20] instead of the simpler model of Mathis et al. [18] which did not yield the intended results. This TCP model estimates throughput of a connection according to Equation 2:

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

This has the same parameters as Equation 1 except t_{RTO}, which is usually approximated as $4R$. Extra terms model further details of TCP loss response, making it more sensitive to loss.

We repeat the simulations from Section IV-C to evaluate this new TCP throughput model. Figure 11 shows the number of sessions terminated using this new model. We observe a 15% increase in the number of session triggering the circuit breaker (to 50–70%) in high and extreme load scenarios, although coupled with an increase in the moderate load case too.

Figure 12 shows the CDF of loss rate observed in ongoing and terminated sessions. Sessions that did not trigger the circuit breaker are better clustered towards low packet loss rates (up to 20%), while those triggering have packet loss in the range 8–80%. Overall, performance is significantly improved compared to the circuit breaker using the simpler TCP throughput model.

VI. Related Work

Performance of the circuit breaker on residential access links and for interactive multimedia is studied in [10]. This shows the circuit breaker takes at least 20 s to engage when the reported loss fraction is higher than 33 %, giving time for a congestion control algorithm to operate. It also shows that the circuit breaker works as desired with residential access link loss/delay patterns. Evaluation with the TCP throughput equation of Padhye et al. shows this makes the circuit breaker over-sensitive to packet loss in that environment. Fough et al. [26] consider effects of bursty video on the circuit breaker, noting that traffic needs to be smoothed to prevent bursts causing transient loss that prematurely trigger the circuit breaker.

Apart from implementing RTP circuit breakers, a multimedia endpoint can terminate a multimedia session by monitoring for RTP keep-alive packets [27], STUN consent freshness packets [28], or Datagram Transport Layer Security (DTLS) heartbeat packets [29]. These alternate mechanisms attempt to keep-alive the session and if a remote endpoint fails to receive these packets at regular intervals, a timeout occurs and the multimedia session is terminated. The main advantage of the RTP circuit breaker mechanism is that it not only detects *liveliness* or *reachability* but is also able to recognize severe

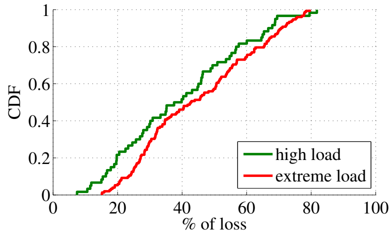

(a) Sessions triggering the RTP circuit breaker

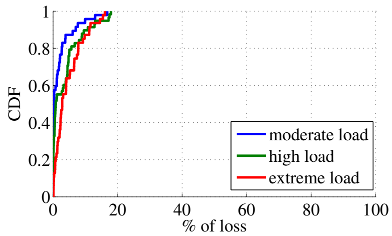

(b) Sessions that did not trigger the circuit breaker

Figure 12. CDF of loss rate; circuit breaker enabled; Padhye TCP equation

congestion (due to packet loss, or through explicit congestion notification [19]) and if needed terminate the session.

VII. Conclusions

We present results of implementing the RTP circuit breaker for video in a simulated LTE network using the default bearer. These show that the RTP circuit breaker generally performs as desired, however the combination of low RTT and high packet loss rates that can occur due to AQM in LTE networks can cause the circuit breaker to fail to trigger when using the simplified model of TCP throughput. We show that simply reducing the multiplier in the congestion circuit breaker is not sufficient to address this issue, but that using the more complete TCP model of Padhye et al. gives better performance. Unfortunately, our previous work [10] has shown that this more complete TCP model leads to an overly sensitive circuit breaker when used in residential access networks. We believe some of this over-sensitivity is due to averaging packet loss events over long RTCP reporting intervals. RTP allows a 'reduced minimum RTCP reporting interval' [1]; future work will consider the effect of this reduced minimum, and whether the circuit breaker using the complete TCP throughput equation is effective in both residential and LTE networks when using less heavily aggregated metrics. As it is possible to have communication services over QoS bearer in LTE, this should improve the media quality. Running circuit breaker enabled flows over such a bearer will further validate the usability of circuit breakers.

Acknowledgements

Thanks to Magnus Westerlund, Stefan Håkansson, Mats Nordberg, Ingemar Johansson and Markus Andersson at Ericsson for their comments and support.

References

[1] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: Transport protocol for real-time applications," IETF, July 2003, RFC 3550.
[2] R. Rejaie, M. Handley, and D. Estrin, "RAP: An end-to-end rate-based congestion control mechanism for realtime streams in the Internet," in *Proc. Infocom Conference*, March 1999.
[3] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based congestion control for unicast applications," in *Proc. SIGCOMM Conference*. ACM, August 2000.
[4] L. Gharai and C. S. Perkins, "Implementing congestion control in the real world," in *Proc. Intl. Conference on Multimedia and Expo*. Lausanne, Switzerland: IEEE, August 2002.
[5] V. Singh, J. Ott, and I. Curcio, "Rate adaptation for conversational 3G video," in *Proc. Infocom Workshop*, Rio de Janeiro, Brazil, 2009.
[6] V. Singh, "Rate-control for conversational H.264 video communication in heterogeneous networks," Master's thesis, Aalto University, April 2011.
[7] C. Jennings, T. Hardie, and M. Westerlund, "Real-time communications for the web," *IEEE Comm. Magazine*, vol. 51, no. 4, April 2013.
[8] V. Singh and J. Ott, "Evaluating Congestion Control for Interactive Real-time Media," IETF, October 2013, work in progress.
[9] C. S. Perkins and V. Singh, "Multimedia congestion control: Circuit breakers for unicast RTP sessions," IETF, February 2014, work in progress.
[10] V. Singh, S. McQuistin, M. Ellis, and C. S. Perkins, "Circuit breakers for multimedia congestion control," in *Proc. Packet Video*. IEEE, 2013.
[11] J. Devadoss, V. Singh, J. Ott, C. Liu, Y.-K. Wang, and I. Curcio, "Evaluation of error resilience mechanisms for 3G conversational video," in *Proc. IEEE Intl. Symp. Multimedia*, 2008, pp. 378–383.
[12] M. Zink, O. Künzel, J. Schmitt, and R. Steinmetz, "Subjective Impression of Variations in Layer Encoded Videos," in *Proc. IWQoS*, 2003.
[13] V. Singh, J. Ott, and C. S. Perkins, "Congestion Control for Interactive Media: Control Loops & APIs," in *Proc. IAB/IRTF Workshop on Congestion Control for Interactive Real-Time Communication*, July 2012.
[14] E. Brosh, S. A. Baset, D. Rubenstein, and H. Schulzrinne, "The delay-friendliness of TCP," in *Proc. SIGMETRICS*. ACM, June 2008.
[15] M. Schier and M. Welzl, "Using DCCP: Issues and improvements," in *Proc. IEEE ICNP*, 2012.
[16] V. Singh, A. A. Lozano, and J. Ott, "Performance analysis of receive-side real-time congestion control for WebRTC," in *Proc. Packet Video*, December 2013.
[17] J. Ott, S. Wenger, N. Sato, C. Burmeister, and J. Rey, "Extended RTP profile for RTCP-based feedback," IETF, July 2006, RFC 4585.
[18] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm," *ACM SIGCOMM CCR*, vol. 27, no. 3, July 1997.
[19] M. Westerlund, I. Johansson, C. S. Perkins, P. O'Hanlon, and K. Carlberg, "Explicit congestion notification (ECN) for RTP over UDP," Internet Engineering Task Force, August 2012, RFC 6679.
[20] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: A simple model and its empirical validation," in *Proc. SIGCOMM Conference*. ACM, August 1998.
[21] 3GPP, "IP multimedia subsystem," TS 23.228, 2012.
[22] A. Toskala, H. Holma, K. Pajukoski, and E. Tiirola, "Utran long term evolution in 3GPP," in *Intl. Symp. Personal, Indoor and Mobile Radio Communications*, 2006.
[23] W. Hai, W. Lei, and L. Min, "HSDPA link adaptation based on novel quality model," in *Proc. Vehicular Technology Conference*, vol. 1, 2005.
[24] 3GPP, "Quality of service concept and architecture," TS 23.107, 2012.
[25] ——, "Physical layer aspects for evolved universal terrestrial radio access (UTRA)," TR 25.814, 2006.
[26] N. Fough, F. Verdicchio, and G. Fairhurst, "Exploration of RTP circuit breaker with applications to video streaming," in *Symp. Convergence Telecomm., Networking, & Broadcasting*, Liverpool, June 2013.
[27] X. Marjou and A. Sollaud, "Application Mechanism for Keeping Alive NAT Mappings Associated with RTP," IETF, June 2011, RFC 6263.
[28] M. Perumal, D. Wing, R. R., and R. T., "STUN usage for consent freshness," IETF, September 2013, work in progress.
[29] R. Seggelmann, M. Tuexen, and M. Williams, "TLS and DTLS Heartbeat Extension," IETF, February 2012, RFC 6520.