

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 9, 2017

S. Holmer
H. Lundin
Google
G. Carlucci
L. De Cicco
S. Mascolo
Politecnico di Bari
July 8, 2016

A Google Congestion Control Algorithm for Real-Time Communication
draft-ietf-rmcat-gcc-02

Abstract

This document describes two methods of congestion control when using real-time communications on the World Wide Web (RTCWEB); one delay-based and one loss-based.

It is published as an input document to the RMCAT working group on congestion control for media streams. The mailing list of that working group is rmcat@ietf.org.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 9, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Mathematical notation conventions	3
2. System model	4
3. Feedback and extensions	4
4. Sending Engine	5
5. Delay-based control	5
5.1. Arrival-time model	6
5.2. Pre-filtering	7
5.3. Arrival-time filter	7
5.4. Over-use detector	8
5.5. Rate control	10
5.6. Parameters settings	12
6. Loss-based control	13
7. Interoperability Considerations	14
8. Implementation Experience	14
9. Further Work	15
10. IANA Considerations	15
11. Security Considerations	15
12. Acknowledgements	15
13. References	16
13.1. Normative References	16
13.2. Informative References	16
Appendix A. Change log	16
A.1. Version -00 to -01	16
A.2. Version -01 to -02	17
A.3. Version -02 to -03	17
A.4. rtcweb-03 to rmcats-00	17
A.5. rmcats -00 to -01	17
A.6. rmcats -01 to -02	17
A.7. rmcats -02 to -03	18
A.8. ietf-rmcats -00 to ietf-rmcats -01	18

A.9. ietf-rmcat -01 to ietf-rmcat -02 18
 Authors' Addresses 18

1. Introduction

Congestion control is a requirement for all applications sharing the Internet resources [RFC2914].

Congestion control for real-time media is challenging for a number of reasons:

- o The media is usually encoded in forms that cannot be quickly changed to accommodate varying bandwidth, and bandwidth requirements can often be changed only in discrete, rather large steps
- o The participants may have certain specific wishes on how to respond - which may not be reducing the bandwidth required by the flow on which congestion is discovered
- o The encodings are usually sensitive to packet loss, while the real-time requirement precludes the repair of packet loss by retransmission

This memo describes two congestion control algorithms that together are able to provide good performance and reasonable bandwidth sharing with other video flows using the same congestion control and with TCP flows that share the same links.

The signaling used consists of experimental RTP header extensions and RTCP messages RFC 3550 [RFC3550] as defined in [abs-send-time], [I-D.alvestrand-rmcat-remb] and [I-D.holmer-rmcat-transport-wide-cc-extensions].

1.1. Mathematical notation conventions

The mathematics of this document have been transcribed from a more formula-friendly format.

The following notational conventions are used:

X_hat An estimate of the true value of variable X - conventionally marked by a circumflex accent on top of the variable name.

X(i) The "i"th value of vector X - conventionally marked by a subscript i.

E{X} The expected value of the stochastic variable X

2. System model

The following elements are in the system:

- o RTP packet - an RTP packet containing media data.
- o Group of packets - a set of RTP packets transmitted from the sender uniquely identified by the group departure and group arrival time (absolute send time) [abs-send-time]. These could be video packets, audio packets, or a mix of audio and video packets.
- o Incoming media stream - a stream of frames consisting of RTP packets.
- o RTP sender - sends the RTP stream over the network to the RTP receiver. It generates the RTP timestamp and the abs-send-time header extension
- o RTP receiver - receives the RTP stream, marks the time of arrival.
- o RTCP sender at RTP receiver - sends receiver reports, REMB messages and transport-wide RTCP feedback messages.
- o RTCP receiver at RTP sender - receives receiver reports and REMB messages and transport-wide RTCP feedback messages, reports these to the sender side controller.
- o RTCP receiver at RTP receiver, receives sender reports from the sender.
- o Loss-based controller - takes loss rate measurement, round trip time measurement and REMB messages, and computes a target sending bitrate.
- o Delay-based controller - takes the packet arrival info, either at the RTP receiver, or from the feedback received by the RTP sender, and computes a maximum bitrate which it passes to the loss-based controller.

Together, loss-based controller and delay-based controller implement the congestion control algorithm.

3. Feedback and extensions

There are two ways to implement the proposed algorithm. One where both the controllers are running at the send-side, and one where the delay-based controller runs on the receive-side and the loss-based controller runs on the send-side.

The first version can be realized by using a per-packet feedback protocol as described in [I-D.holmer-rmcat-transport-wide-cc-extensions]. Here, the RTP receiver will record the arrival time and the transport-wide sequence number of each received packet, which will be sent back to the sender periodically using the transport-wide feedback message. The RECOMMENDED feedback interval is once per received video frame or at least once every 30 ms if audio-only or multi-stream. If the feedback overhead needs to be limited this interval can be increased to 100 ms.

The sender will map the received {sequence number, arrival time} pairs to the send-time of each packet covered by the feedback report, and feed those timestamps to the delay-based controller. It will also compute a loss ratio based on the sequence numbers in the feedback message.

The second version can be realized by having a delay-based controller at the receive-side, monitoring and processing the arrival time and size of incoming packets. The sender SHOULD use the abs-send-time RTP header extension [abs-send-time] to enable the receiver to compute the inter-group delay variation. The output from the delay-based controller will be a bitrate, which will be sent back to the sender using the REMB feedback message [I-D.alvestrand-rmcat-remb]. The packet loss ratio is sent back via RTCP receiver reports. At the sender the bitrate in the REMB message and the fraction of packets lost are fed into the loss-based controller, which outputs a final target bitrate. It is RECOMMENDED to send the REMB message as soon as congestion is detected, and otherwise at least once every second.

4. Sending Engine

Pacing is used to actuate the target bitrate computed by the controllers.

When media encoder produces data, this is fed into a Pacer queue. The Pacer sends a group of packets to the network every burst_time interval. RECOMMENDED value for burst_time is 5 ms. The size of a group of packets is computed as the product between the target bitrate and the burst_time.

5. Delay-based control

The delay-based control algorithm can be further decomposed into four parts: a pre-filtering, an arrival-time filter, an over-use detector, and a rate controller.

5.1. Arrival-time model

This section describes an adaptive filter that continuously updates estimates of network parameters based on the timing of the received groups of packets.

We define the inter-arrival time, $t(i) - t(i-1)$, as the difference in arrival time of two groups of packets. Correspondingly, the inter-departure time, $T(i) - T(i-1)$, is defined as the difference in departure-time of two groups of packets. Finally, the inter-group delay variation, $d(i)$, is defined as the difference between the inter-arrival time and the inter-departure time. Or interpreted differently, as the difference between the delay of group i and group $i-1$.

$$d(i) = t(i) - t(i-1) - (T(i) - T(i-1))$$

An inter-departure time is computed between consecutive groups as $T(i) - T(i-1)$, where $T(i)$ is the departure timestamp of the last packet in the current packet group being processed. Any packets received out of order are ignored by the arrival-time model.

Each group is assigned a receive time $t(i)$, which corresponds to the time at which the last packet of the group was received. A group is delayed relative to its predecessor if $t(i) - t(i-1) > T(i) - T(i-1)$, i.e., if the inter-arrival time is larger than the inter-departure time.

We can model the inter-group delay variation as:

$$d(i) = w(i)$$

Here, $w(i)$ is a sample from a stochastic process W , which is a function of the link capacity, the current cross traffic, and the current sent bitrate. We model W as a white Gaussian process. If we are over-using the channel we expect the mean of $w(i)$ to increase, and if a queue on the network path is being emptied, the mean of $w(i)$ will decrease; otherwise the mean of $w(i)$ will be zero.

Breaking out the mean, $m(i)$, from $w(i)$ to make the process zero mean, we get

Equation 1

$$d(i) = m(i) + v(i)$$

The noise term $v(i)$ represents network jitter and other delay effects not captured by the model.

5.2. Pre-filtering

The pre-filtering aims at handling delay transients caused by channel outages. During an outage, packets being queued in network buffers, for reasons unrelated to congestion, are delivered in a burst when the outage ends.

The pre-filtering merges together groups of packets that arrive in a burst. Packets are merged in the same group if one of these two conditions holds:

- o A sequence of packets which are sent within a `burst_time` interval constitute a group.
- o A Packet which has an inter-arrival time less than `burst_time` and an inter-group delay variation $d(i)$ less than 0 is considered being part of the current group of packets.

5.3. Arrival-time filter

The parameter $d(i)$ is readily available for each group of packets, $i > 1$. We want to estimate $m(i)$ and use this estimate to detect whether or not the bottleneck link is over-used. The parameter can be estimated by any adaptive filter - we are using the Kalman filter.

Let $m(i)$ be the estimate at time i

We model the state evolution from time i to time $i+1$ as

$$m(i+1) = m(i) + u(i)$$

where $u(i)$ is the state noise that we model as a stationary process with Gaussian statistic with zero mean and variance

$$q(i) = E\{u(i)^2\}$$

$q(i)$ is RECOMMENDED equal to 10^{-3}

Given equation 1 we get

$$d(i) = m(i) + v(i)$$

where $v(i)$ is zero mean white Gaussian measurement noise with variance $\text{var}_v = E\{v(i)^2\}$

The Kalman filter recursively updates our estimate $\hat{m}(i)$ as

$$z(i) = d(i) - m_{\text{hat}}(i-1)$$

$$m_{\text{hat}}(i) = m_{\text{hat}}(i-1) + z(i) * k(i)$$

$$k(i) = \frac{e(i-1) + q(i)}{\text{var_v_hat}(i) + (e(i-1) + q(i))}$$

$$e(i) = (1 - k(i)) * (e(i-1) + q(i))$$

The variance $\text{var_v}(i) = E\{v(i)^2\}$ is estimated using an exponential averaging filter, modified for variable sampling rate

$$\text{var_v_hat}(i) = \max(\alpha * \text{var_v_hat}(i-1) + (1-\alpha) * z(i)^2, 1)$$

$$\alpha = (1-\chi)^{(30/(1000 * f_{\text{max}}))}$$

where $f_{\text{max}} = \max \{1/(T(j) - T(j-1))\}$ for j in $i-K+1, \dots, i$ is the highest rate at which the last K packet groups have been received and χ is a filter coefficient typically chosen as a number in the interval $[0.1, 0.001]$. Since our assumption that $v(i)$ should be zero mean WGN is less accurate in some cases, we have introduced an additional outlier filter around the updates of var_v_hat . If $z(i) > 3*\text{sqrt}(\text{var_v_hat})$ the filter is updated with $3*\text{sqrt}(\text{var_v_hat})$ rather than $z(i)$. For instance $v(i)$ will not be white in situations where packets are sent at a higher rate than the channel capacity, in which case they will be queued behind each other.

5.4. Over-use detector

The inter-group delay variation estimate $m(i)$, obtained as the output of the arrival-time filter, is compared with a threshold $\text{del_var_th}(i)$. An estimate above the threshold is considered as an indication of over-use. Such an indication is not enough for the detector to signal over-use to the rate control subsystem. A definitive over-use will be signaled only if over-use has been detected for at least overuse_time_th milliseconds. However, if $m(i) < m(i-1)$, over-use will not be signaled even if all the above conditions are met. Similarly, the opposite state, under-use, is detected when $m(i) < -\text{del_var_th}(i)$. If neither over-use nor under-use is detected, the detector will be in the normal state.

The threshold del_var_th has a remarkable impact on the overall dynamics and performance of the algorithm. In particular, it has been shown that using a static threshold del_var_th , a flow controlled by the proposed algorithm can be starved by a concurrent TCP flow [Pv13]. This starvation can be avoided by increasing the threshold del_var_th to a sufficiently large value.

The reason is that, by using a larger value of `del_var_th`, a larger queuing delay can be tolerated, whereas with a small `del_var_th`, the over-use detector quickly reacts to a small increase in the offset estimate $m(i)$ by generating an over-use signal that reduces the delay-based estimate of the available bandwidth $A_{\hat{}}$ (see Section 4.4). Thus, it is necessary to dynamically tune the threshold `del_var_th` to get good performance in the most common scenarios, such as when competing with loss-based flows.

For this reason, we propose to vary the threshold `del_var_th(i)` according to the following dynamic equation:

$$\text{del_var_th}(i) = \text{del_var_th}(i-1) + (t(i)-t(i-1)) * K(i) * (|m(i)|-\text{del_var_th}(i-1))$$

with $K(i)=K_d$ if $|m(i)| < \text{del_var_th}(i-1)$ or $K(i)=K_u$ otherwise. The rationale is to increase `del_var_th(i)` when $m(i)$ is outside of the range $[-\text{del_var_th}(i-1), \text{del_var_th}(i-1)]$, whereas, when the offset estimate $m(i)$ falls back into the range, `del_var_th` is decreased. In this way when $m(i)$ increases, for instance due to a TCP flow entering the same bottleneck, `del_var_th(i)` increases and avoids the uncontrolled generation of over-use signals which may lead to starvation of the flow controlled by the proposed algorithm [Pv13]. Moreover, `del_var_th(i)` SHOULD NOT be updated if this condition holds:

$$|m(i)| - \text{del_var_th}(i) > 15$$

It is also RECOMMENDED to clamp `del_var_th(i)` to the range $[6, 600]$, since a too small `del_var_th(i)` can cause the detector to become overly sensitive.

On the other hand, when $m(i)$ falls back into the range $[-\text{del_var_th}(i-1), \text{del_var_th}(i-1)]$ the threshold `del_var_th(i)` is decreased so that a lower queuing delay can be achieved.

It is RECOMMENDED to choose $K_u > K_d$ so that the rate at which `del_var_th` is increased is higher than the rate at which it is decreased. With this setting it is possible to increase the threshold in the case of a concurrent TCP flow and prevent starvation as well as enforcing intra-protocol fairness. RECOMMENDED values for `del_var_th(0)`, `overuse_time_th`, K_u and K_d are respectively 12.5 ms, 10 ms, 0.01 and 0.00018.

5.5. Rate control

The rate control is split in two parts, one controlling the bandwidth estimate based on delay, and one controlling the bandwidth estimate based on loss. Both are designed to increase the estimate of the available bandwidth $A_{\hat{}}$ as long as there is no detected congestion and to ensure that we will eventually match the available bandwidth of the channel and detect an over-use.

As soon as over-use has been detected, the available bandwidth estimated by the delay-based controller is decreased. In this way we get a recursive and adaptive estimate of the available bandwidth.

In this document we make the assumption that the rate control subsystem is executed periodically and that this period is constant.

The rate control subsystem has 3 states: Increase, Decrease and Hold. "Increase" is the state when no congestion is detected; "Decrease" is the state where congestion is detected, and "Hold" is a state that waits until built-up queues have drained before going to "increase" state.

The state transitions (with blank fields meaning "remain in state") are:

Signal \ State	Hold	Increase	Decrease
Over-use	Decrease	Decrease	
Normal	Increase		Hold
Under-use		Hold	Hold

The subsystem starts in the increase state, where it will stay until over-use or under-use has been detected by the detector subsystem. On every update the delay-based estimate of the available bandwidth is increased, either multiplicatively or additively, depending on its current state.

The system does a multiplicative increase if the current bandwidth estimate appears to be far from convergence, while it does an additive increase if it appears to be closer to convergence. We assume that we are close to convergence if the currently incoming bitrate, $R_{\hat{}}(i)$, is close to an average of the incoming bitrates at

the time when we previously have been in the Decrease state. "Close" is defined as three standard deviations around this average. It is RECOMMENDED to measure this average and standard deviation with an exponential moving average with the smoothing factor 0.95, as it is expected that this average covers multiple occasions at which we are in the Decrease state. Whenever valid estimates of these statistics are not available, we assume that we have not yet come close to convergence and therefore remain in the multiplicative increase state.

If $R_{\hat{i}}$ increases above three standard deviations of the average max bitrate, we assume that the current congestion level has changed, at which point we reset the average max bitrate and go back to the multiplicative increase state.

$R_{\hat{i}}$ is the incoming bitrate measured by the delay-based controller over a T seconds window:

$$R_{\hat{i}} = 1/T * \text{sum}(L(j)) \text{ for } j \text{ from } 1 \text{ to } N(i)$$

$N(i)$ is the number of packets received the past T seconds and $L(j)$ is the payload size of packet j . A window between 0.5 and 1 second is RECOMMENDED.

During multiplicative increase, the estimate is increased by at most 8% per second.

$$\begin{aligned} \text{eta} &= 1.08^{\min(\text{time_since_last_update_ms} / 1000, 1.0)} \\ A_{\hat{i}} &= \text{eta} * A_{\hat{i-1}} \end{aligned}$$

During the additive increase the estimate is increased with at most half a packet per response_time interval. The response_time interval is estimated as the round-trip time plus 100 ms as an estimate of over-use estimator and detector reaction time.

$$\begin{aligned} \text{response_time_ms} &= 100 + \text{rtt_ms} \\ \alpha &= 0.5 * \min(\text{time_since_last_update_ms} / \text{response_time_ms}, 1.0) \\ A_{\hat{i}} &= A_{\hat{i-1}} + \max(1000, \alpha * \text{expected_packet_size_bits}) \end{aligned}$$

expected_packet_size_bits is used to get a slightly slower slope for the additive increase at lower bitrates. It can for instance be computed from the current bitrate by assuming a frame rate of 30 frames per second:

$$\begin{aligned} \text{bits_per_frame} &= A_{\hat{i-1}} / 30 \\ \text{packets_per_frame} &= \text{ceil}(\text{bits_per_frame} / (1200 * 8)) \\ \text{avg_packet_size_bits} &= \text{bits_per_frame} / \text{packets_per_frame} \end{aligned}$$

Since the system depends on over-using the channel to verify the current available bandwidth estimate, we must make sure that our estimate does not diverge from the rate at which the sender is actually sending. Thus, if the sender is unable to produce a bit stream with the bitrate the congestion controller is asking for, the available bandwidth estimate should stay within a given bound. Therefore we introduce a threshold

$$A_{\hat{}}(i) < 1.5 * R_{\hat{}}(i)$$

When an over-use is detected the system transitions to the decrease state, where the delay-based available bandwidth estimate is decreased to a factor times the currently incoming bitrate.

$$A_{\hat{}}(i) = \text{beta} * R_{\hat{}}(i)$$

beta is typically chosen to be in the interval [0.8, 0.95], 0.85 is the RECOMMENDED value.

When the detector signals under-use to the rate control subsystem, we know that queues in the network path are being emptied, indicating that our available bandwidth estimate $A_{\hat{}}$ is lower than the actual available bandwidth. Upon that signal the rate control subsystem will enter the hold state, where the receive-side available bandwidth estimate will be held constant while waiting for the queues to stabilize at a lower level - a way of keeping the delay as low as possible. This decrease of delay is wanted, and expected, immediately after the estimate has been reduced due to over-use, but can also happen if the cross traffic over some links is reduced.

It is RECOMMENDED that the routine to update $A_{\hat{}}(i)$ is run at least once every `response_time` interval.

5.6. Parameters settings

Parameter	Description	RECOMMENDED Value
burst_time	Time limit in milliseconds between packet bursts which identifies a group	5 ms
q	State noise covariance matrix	$q = 10^{-3}$
e(0)	Initial value of the system error covariance	$e(0) = 0.1$
chi	Coefficient used for the measured noise variance	[0.1, 0.001]
del_var_th(0)	Initial value for the adaptive threshold	12.5 ms
overuse_time_th	Time required to trigger an overuse signal	10 ms
K_u	Coefficient for the adaptive threshold	0.01
K_d	Coefficient for the adaptive threshold	0.00018
T	Time window for measuring the received bitrate	[0.5, 1] s
beta	Decrease rate factor	0.85

Table 1: RECOMMENDED values for delay based controller

Table 1

6. Loss-based control

A second part of the congestion controller bases its decisions on the round-trip time, packet loss and available bandwidth estimates $A_{\hat{}}$ received from the delay-based controller. The available bandwidth estimates computed by the loss-based controller are denoted with $A_{s_{\hat{}}}$.

The available bandwidth estimates $A_{\hat{}}$ produced by the delay-based controller are only reliable when the size of the queues along the path sufficiently large. If the queues are very short, over-use will only be visible through packet losses, which are not used by the delay-based controller.

The loss-based controller SHOULD run every time feedback from the receiver is received.

- o If 2-10% of the packets have been lost since the previous report from the receiver, the sender available bandwidth estimate $As_hat(i)$ will be kept unchanged.
- o If more than 10% of the packets have been lost a new estimate is calculated as $As_hat(i) = As_hat(i-1)(1-0.5p)$, where p is the loss ratio.
- o As long as less than 2% of the packets have been lost $As_hat(i)$ will be increased as $As_hat(i) = 1.05(As_hat(i-1))$

The loss-based estimate As_hat is compared with the delay-based estimate A_hat . The actual sending rate is set as the minimum between As_hat and A_hat .

We motivate the packet loss thresholds by noting that if the transmission channel has a small amount of packet loss due to over-use, that amount will soon increase if the sender does not adjust his bitrate. Therefore we will soon enough reach above the 10% threshold and adjust $As_hat(i)$. However, if the packet loss ratio does not increase, the losses are probably not related to self-inflicted congestion and therefore we should not react on them.

7. Interoperability Considerations

In case a sender implementing these algorithms talks to a receiver which do not implement any of the proposed RTCP messages and RTP header extensions, it is suggested that the sender monitors RTCP receiver reports and uses the fraction of lost packets and the round-trip time as input to the loss-based controller. The delay-based controller should be left disabled.

8. Implementation Experience

This algorithm has been implemented in the open-source WebRTC project, has been in use in Chrome since M23, and is being used by Google Hangouts.

Deployment of the algorithm have revealed problems related to, e.g, congested or otherwise problematic WiFi networks, which have led to algorithm improvements. The algorithm has also been tested in a multi-party conference scenario with a conference server which terminates the congestion control between endpoints. This ensures that no assumptions are being made by the congestion control about maximum send and receive bitrates, etc., which typically is out of control for a conference server.

9. Further Work

This draft is offered as input to the congestion control discussion.

Work that can be done on this basis includes:

- o Considerations of integrated loss control: How loss and delay control can be better integrated, and the loss control improved.
- o Considerations of locus of control: evaluate the performance of having all congestion control logic at the sender, compared to splitting logic between sender and receiver.
- o Considerations of utilizing ECN as a signal for congestion estimation and link over-use detection.

10. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

11. Security Considerations

An attacker with the ability to insert or remove messages on the connection would have the ability to disrupt rate control. This could make the algorithm to produce either a sending rate under-utilizing the bottleneck link capacity, or a too high sending rate causing network congestion.

In this case, the control information is carried inside RTP, and can be protected against modification or message insertion using SRTP, just as for the media. Given that timestamps are carried in the RTP header, which is not encrypted, this is not protected against disclosure, but it seems hard to mount an attack based on timing information only.

12. Acknowledgements

Thanks to Randell Jesup, Magnus Westerlund, Varun Singh, Tim Panton, Soo-Hyun Choo, Jim Gettys, Ingemar Johansson, Michael Welzl and others for providing valuable feedback on earlier versions of this draft.

13. References

13.1. Normative References

- [I-D.alvestrand-rmcat-remb]
Alvestrand, H., "RTCP message for Receiver Estimated Maximum Bitrate", draft-alvestrand-rmcat-remb-03 (work in progress), October 2013.
- [I-D.holmer-rmcat-transport-wide-cc-extensions]
Holmer, S., Flodman, M., and E. Sprang, "RTP Extensions for Transport-wide Congestion Control", draft-holmer-
rmcat-transport-wide-cc-extensions-00 (work in progress),
March 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [abs-send-time]
"RTP Header Extension for Absolute Sender Time",
<[http://www.webrtc.org/experiments/rtp-hdrext/
abs-send-time](http://www.webrtc.org/experiments/rtp-hdrext/abs-send-time)>.

13.2. Informative References

- [Pv13] De Cicco, L., Carlucci, G., and S. Mascolo, "Understanding the Dynamic Behaviour of the Google Congestion Control", Packet Video Workshop , December 2013.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.

Appendix A. Change log

A.1. Version -00 to -01

- o Added change log
- o Added appendix outlining new extensions
- o Added a section on when to send feedback to the end of section 3.3 "Rate control", and defined min/max FB intervals.

- o Added size of over-bandwidth estimate usage to "further work" section.
- o Added startup considerations to "further work" section.
- o Added sender-delay considerations to "further work" section.
- o Filled in acknowledgments section from mailing list discussion.

A.2. Version -01 to -02

- o Defined the term "frame", incorporating the transmission time offset into its definition, and removed references to "video frame".
- o Referred to "m(i)" from the text to make the derivation clearer.
- o Made it clearer that we modify our estimates of available bandwidth, and not the true available bandwidth.
- o Removed the appendixes outlining new extensions, added pointers to REMB draft and RFC 5450.

A.3. Version -02 to -03

- o Added a section on how to process multiple streams in a single estimator using RTP timestamps to NTP time conversion.
- o Stated in introduction that the draft is aimed at the RMCAT working group.

A.4. rtcweb-03 to rmcate-00

Renamed draft to link the draft name to the RMCAT WG.

A.5. rmcate -00 to -01

Spellcheck. Otherwise no changes, this is a "keepalive" release.

A.6. rmcate -01 to -02

- o Added Luca De Cicco and Saverio Mascolo as authors.
- o Extended the "Over-use detector" section with new technical details on how to dynamically tune the offset `del_var_th` for improved fairness properties.

- o Added reference to a paper analyzing the behavior of the proposed algorithm.

A.7. rmcats -02 to -03

- o Swapped receiver-side/sender-side controller with delay-based/loss-based controller as there is no longer a requirement to run the delay-based controller on the receiver-side.
- o Removed the discussion about multiple streams and transmission time offsets.
- o Introduced a new section about "Feedback and extensions".
- o Improvements to the threshold adaptation in the "Over-use detector" section.
- o Swapped the previous MIMD rate control algorithm for a new AIMD rate control algorithm.

A.8. ietf-rmcats -00 to ietf-rmcats -01

- o Arrival-time filter converted from a two dimensional Kalman filter to a scalar Kalman filter.
- o The use of the TFRC equation was removed from the loss-based controller, as it turned out to have little to no effect in practice.

A.9. ietf-rmcats -01 to ietf-rmcats -02

- o Added a section which better describes the pre-filtering algorithm.

Authors' Addresses

Stefan Holmer
Google
Kungsbron 2
Stockholm 11122
Sweden

Email: holmer@google.com

Henrik Lundin
Google
Kungsbron 2
Stockholm 11122
Sweden

Email: hlundin@google.com

Gaetano Carlucci
Politecnico di Bari
Via Orabona, 4
Bari 70125
Italy

Email: gaetano.carlucci@poliba.it

Luca De Cicco
Politecnico di Bari
Via Orabona, 4
Bari 70125
Italy

Email: l.decicco@poliba.it

Saverio Mascolo
Politecnico di Bari
Via Orabona, 4
Bari 70125
Italy

Email: mascolo@poliba.it