    Evaluating Congestion Control for Interactive Real-time Media
                 draft-singh-rmcat-cc-eval-03.txt

Abstract

   The Real-time Transport Protocol (RTP) is used to transmit media in
   telephony and video conferencing applications. This document
   describes the guidelines to evaluate new congestion control
   algorithms for interactive point-to-point real-time media.

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 16, 2014.

Table of Contents

1. Introduction

This memo describes the guidelines to help with evaluating new
congestion control algorithms for interactive point-to-point real
time media. The requirements for the congestion control algorithm
are outlined in [I-D.jesup-rmcat-reqs]). This document builds upon
previous work at the IETF: Specifying New Congestion Control
Algorithms [RFC5033] and Metrics for the Evaluation of Congestion
Control Algorithms [RFC5166].

The guidelines proposed in the document are intended to help prevent
a congestion collapse, promote fair capacity usage and optimize the
media flow's throughput. Furthermore, the proposed algorithms are
expected to operate within the envelope of the circuit breakers
defined in [I-D.ietf-avtcore-rtp-circuit-breakers].

This document only provides broad-level criteria for evaluating a new
congestion control algorithm and the working group should expect a
thorough scientific study to make its decision. The results of the
evaluation are not expected to be included within the internet-draft
but should be cited in the document.

2. Terminology

The terminology defined in RTP [RFC3550], RTP Profile for Audio and
Video Conferences with Minimal Control [RFC3551], RTCP Extended
Report (XR) [RFC3611], Extended RTP Profile for RTCP-based Feedback
(RTP/AVPF) [RFC4585] and Support for Reduced-Size RTCP [RFC5506]
apply.

3. Metrics

[RFC5166] describes the basic metrics for congestion control.
Metrics that are important to interactive multimedia are:

o Throughput.

o Minimizing oscillations in the transmission rate (stability) when
   the end-to-end capacity varies slowly.

o Delay.

o Reactivity to transient events.

o Packet losses and discards.

Each experiment logs every incoming and outgoing packet (the RTP
logging format is described in Section 3.1). The logging can be done

inside the application or at the endpoints using pcap (packet
capture, e.g., tcpdump, wireshark). The following are calculated
based on the information in the packet logs:

1. Sending rate, Receiver rate, Goodput

2. Packet delay

3. Packet loss

4. Packets discarded from the playout or de-jitter buffer

[Editor's note: How to handle packet re-transmissions? loss before
retransmission, after retransmission?]

[Open issue (1): Instead of defining fairness, there has been
discussion on defining "unfairness". The criteria are:
1. Do not trigger the circuit breaker.
2. Over 3 times or less than 1/3 times the throughput for an RMCAT
media stream compared to identical RMCAT streams competing on a
bottleneck, for a case when the competing streams have similar RTTs.
3. Over 3 times delay compared to RTT measurements performed before
starting the RMCAT flow or for the case when competing with identical
RMCAT streams having similar RTTs.
Here, rather than discussing the number '3'? Does the criteria
capture Unfairness adequately?]

[Open issue (2): Convergence time was discussed briefly in the design
meetings. It is defined as: the time it takes the congestion control
to reach a stable rate (at startup or after new RMCAT flows are
added). What is a stable rate?]

[Open issue (3): previous versions of the document had Bandwidth
Utilization, defined as ratio of sending rate to the available
bottleneck capacity. This is useful when the RMCAT flow is by itself
or competing with similar flows (where the assumption would be that
all flows get an equal share). Remove this?]

From the logs the statistical measures (min, max, mean, standard
deviation and variance) for the whole duration or any specific part
of the session can be calculated. Also the metrics (sending rate,
receiver rate, goodput, latency) can be visualized in graphs as
variation over time, the measurements in the plot are at 1 second
intervals. Additionally, from the logs it is possible to plot the
histogram or CDF of packet delay.

Section 2.1 of [RFC5166] discusses the tradeoff between throughput,
delay and loss.

[Open issue (4): Application trade-off is yet to be defined. see
RMCAT requirements [I-D.jesup-rmcat-reqs] document. Perhaps each
experiment should define the application's expectation or trade-off.]

3.1. RTP Log Format

The log file is tab or comma separated containing the following
details:

        Send or receive timestamp (unix)
        RTP payload type
        SSRC
        RTP sequence no
        RTP timestamp
        marker bit
        payload size

[Open issue (5): Should the retransmissions for post-repair loss
metric be logged in a separate file? the repair streams have
different payload type and/or SSRC.]

4. Guidelines

A congestion control algorithm should be tested in simulation or a
testbed environment, and the experiments should be repeated multiple
times to infer statistical significance. The following guidelines
are considered for evaluation:

4.1. Avoiding Congestion Collapse

Does the congestion control propose any changes to (or diverge from)
the circuit breaker conditions defined in
[I-D.ietf-avtcore-rtp-circuit-breakers].

4.2. Stability

The congestion control should be assessed for its stability when the
path characteristics do not change over time. Changing the media
encoding rate estimate too often or by too much may adversely affect
the application layer performance.

4.3. Media Traffic

The congestion control algorithm should be assessed with different
types of media behavior, i.e., the media should contain idle and
data-limited periods. For example, periods of silence for audio or
varying amount of motion for video. However, the evaluation can be

done in two stages. In the first stage, media stream can generate
traffic at the rate calculated by the congestion controller. In the
second stage, real codecs or models of video codecs should be used to
mimic real-world cases.

4.4. Start-up Behaviour

The congestion control algorithm should be assessed with different
start-rates. The main reason is to observe the behavior of the
congestion control in different evaluation scenarios, such as when
competing with varying amount of cross-traffic or how quickly does
the congestion control algorithm achieve a stable sending rate.

[Editor's note: requires a robust definition for unfriendliness and
convergence time.]

4.5. Diverse Environments

The congestion control algorithm should be assessed in heterogeneous
environments, containing both wired and wireless paths. Examples of
wireless access technologies are: 802.11, GPRS, HSPA, or LTE. One of
the main challenges of the wireless environments for the congestion
control algorithm is to distinguish between congestion induced loss
and transmission (bit-error corruption) loss. Congestion control
algorithms may incorrectly identify transmission loss as congestion
loss and reduce the media encoding rate by too much, which may cause
oscillatory behavior and deteriorate the users' quality of
experience. Furthermore, packet loss may induce additional delay in
networks with wireless paths due to link-layer retransmissions.

4.6. Varying Path Characteristics

The congestion control algorithm should be evaluated for a range of
path characteristics such as, different end-to-end capacity and
latency, varying amount of cross traffic on a bottle-neck link and a
router's queue length. In an experiment, if the media only flows in
a single direction, the feedback path should also be tested with
varying amounts of impairments.

The main motivation for the previous and current criteria is to
identify situations in which the proposed congestion control is less
performant.

[Open issue (6): Different types of queueing mechanisms? Random
Early Detection or only DropTail?].

4.7. Reacting to Transient Events or Interruptions

The congestion control algorithm should be able to handle changes in
end-to-end capacity and latency. Latency may change due to route
updates, link failures, handovers etc. In mobile environment the
end-to-end capacity may vary due to the interference, fading,
handovers, etc. In wired networks the end-to-end capacity may vary
due to changes in resource reservation.

4.8. Fairness With Similar Cross-Traffic

The congestion control algorithm should be evaluated when competing
with other RTP flows using the same or another candidate congestion
control algorithm. The proposal should highlight the bottleneck
capacity share of each RTP flow.

[Editor's note: If we define Unfriendliness then that criteria should
be applied here.]

4.9. Impact on Cross-Traffic

The congestion control algorithm should be evaluated when competing
with standard TCP. Short TCP flows may be considered as transient
events and the RTP flow may give way to the short TCP flow to
complete quickly. However, long-lived TCP flows may starve out the
RTP flow depending on router queue length.

The proposal should also measure the impact on varied number of
cross-traffic sources, i.e., few and many competing flows, or mixing
various amounts of TCP and similar cross-traffic.

4.10. Extensions to RTP/RTCP

The congestion control algorithm should indicate if any protocol
extensions are required to implement it and should carefully describe
the impact of the extension.

5. Minimum Requirements for Evaluation

[Editor's Note: If needed, a minimum evaluation criteria can be based
on the above guidelines]

6. Evaluation Parameters

An evaluation scenario is created from a list of network, link and
flow characteristics. The parameters discussed in the following
subsections are meant to aid in creating evaluation scenarios and do
not describe an evaluation scenario. The scenario discussed in
Appendix A takes into account all these parameters.

6.1. **Bottleneck Traffic Flows**

   The network scenario describes the types of flows sharing the common
   bottleneck with a single RMCAT flow, they are:

   1. A single RMCAT flow by itself.

   2. Competing with similar RMCAT flows. These competing flows may
       use the same algorithm or another candidate RMCAT algorithm.

   3. Compete with long-lived TCP.

   4. Compete with bursty TCP.

   5. Compete with LEDBAT flows.

   6. Compete with unresponsive interactive media flows (i.e., not only
       CBR).

   Figure 1 shows an example evaluation topology, where S1..Sn are
   traffic sources, these sources are either RMCAT or a mixture of
   traffic flows listed above. R1..Rn are the corresponding receivers.
   A and B are routers that can be configured to introduce impairments.
   Access links are in between the sender/receiver and the router, while
   the bottleneck link is between the Routers A and B.

```
+---+  Access                                          Access +---+
|S1 |======= \                                        / ======|R1 |
+---+   link \\                                       //  link +---+
            \\                                       //
+---+        +-----+         Bottleneck         +-----+        +---+
|S2 |======|   A  |------------------------------->|   B  |======|R2 |
+---+        |     |<------------------------------|     |        +---+
            +-----+              Link              +-----+
(...)           //                                       \\         (...)
               //                                         \\
+---+         //                                           \\        +---+
|Sn |====== /                                               \ ======|Rn |
+---+                                                              +---+
```
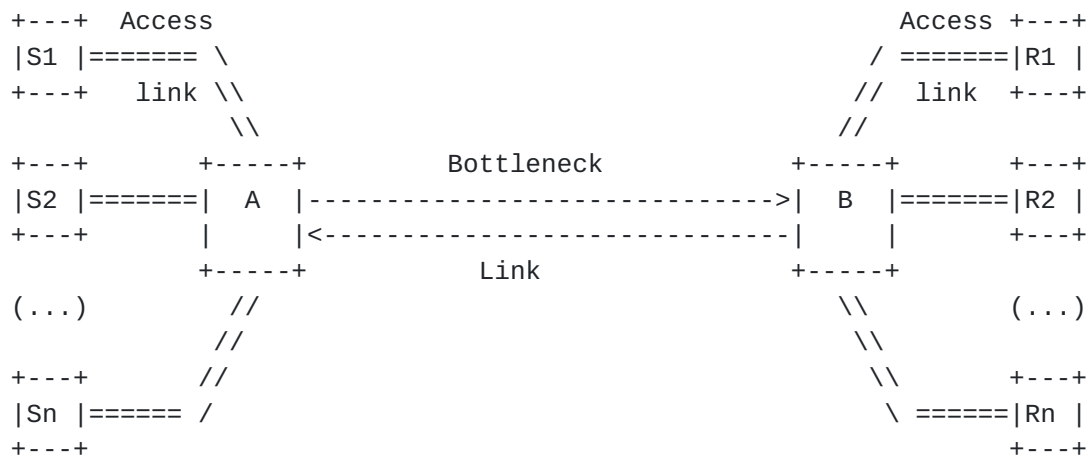
                         Figure 1: Simple Topology

   [Open Issue (7): Discuss more complex topologies]

6.2. **Access Links**

   The media senders and receivers are typically connected to the
   bottleneck link, common access links are:

1. Ethernet (LAN)

2. Wireless LAN (WLAN)

3. 3G/LTE

[Open issue (8): need to describe parameters or traces to model WLAN
and 3G/LTE.]

A real-world network typically consists of a mixture of links, the
most important aspect is to identify the location of the bottleneck
link. The bottleneck link can move from one node to another
depending on the amount of cross-traffic or due to the varying link
capacity. The design of the experiments should take this into
account. In the simplest case the access link may not be the
bottleneck link but an intermediate node.

6.3. Bottleneck Link Parameters

The bottleneck link carries multiple flows, these flows may be other
RMCAT flows or other types of cross-traffic. The experiments should
dimension the bottleneck link based on the number of flows and the
expected behavior. For example, if 5 media flows are expected to
share the bottleneck link equally, the bottleneck link is set to 5
times the desired transmission rate.

If the experiment carries only media in one direction, then the
upstream (sender to receiver) bottleneck link carries media packets
while the downstream (receiver to sender) bottleneck carries the
feedback packets. The bottleneck link parameters discussed in this
section apply only to a single direction, hence the bottleneck link
in the reverse direction can choose the same or have different
parameters.

The link latency corresponds to the propagation delay of the link,
i.e., the time it takes for a packet to traverse the bottleneck link,
it does not include queuing delay. In an experiment with several
links the experiment should describe if the links add latency or not.
It is possible for experiments to have multiple hops with different
link latencies. Experiments are expected to verify that the
congestion control is able to work in challenging situations, for
example over trans-continental and/or satellite links. The
experiment should pick link latency values from the following:

1. Very low latency: 0-1ms

2. Low latency: 50ms

3. High latency: 150ms

4. Extreme latency: 300ms

[Editor's note: currently describes the latency for a single link, instead of end-to-end delay. Which is preferred? or both?]

Similarly, to model lossy links, the experiments can choose one of the following loss rates, the fractional loss is the ratio of packets lost and packets sent.

1. no loss: 0%

2. 1%

3. 5%

4. 10%

5. 20%

[Open issue (10): how is the loss generated? using traces, Gilbert-Elliot model, randomly (uncorrelated) loss.]

6.4. Router Queue Parameters

The router queue length is measured as the time taken to drain the FIFO queue, they are:

1. QoS-aware (or short): 70ms

2. Nominal: 500ms

3. Buffer-bloated: 2000ms

However, the size of the queue is typically measured in bytes or packets and to convert the queue length measured in seconds to queue length in bytes:

QueueSize (in bytes) = QueueSize (in sec) x Throughput (in bps)/8

[Open issue (11): Confirm the above values, do we need to define parameters for other types of queues?]

6.5. Media Flow Parameters

The media sources can be modeled in two ways. In the first, the sources always have data to send, i.e., have no data limited

   intervals and are able to generate the media rate requested by the
   RMCAT congestion control algorithm. In the second, the traffic
   generator models the behavior of a media codec, mainly the burstiness
   (time-varying data produced by a video GOP).

   At the beginning of the session, the media sources are configured to
   start at a given start rate, they are:

   1. 200 kbps

   2. 800 kbps

   3. 1300 kbps

   4. 4000 kbps

6.6. Cross-traffic Parameters

   [Open issue(12): TCP cross-traffic parameters are still TBD, mainly
   the bursty TCP. Long-lived TCP flows will download data throughout
   the session and are expected to have infinite amount of data to send
   or receive.]

7. Status of Proposals

   Congestion control algorithms are expected to be published as
   "Experimental" documents until they are shown to be safe to deploy.
   An algorithm published as a draft should be experimented in
   simulation, or a controlled environment (testbed) to show its
   applicability. Every congestion control algorithm should include a
   note describing the environments in which the algorithm is tested and
   safe to deploy. It is possible that an algorithm is not recommended
   for certain environments or perform sub-optimally for the user.

   [Editor's Note: Should there be a distinction between "Informational"
   and "Experimental" drafts for congestion control algorithms in RMCAT.
   [RFC5033] describes Informational proposals as algorithms that are
   not safe for deployment but are proposals to experiment with in
   simulation/testbeds. While Experimental algorithms are ones that are
   deemed safe in some environments but require a more thorough
   evaluation (from the community).]

8. Security Considerations

   Security issues have not been discussed in this memo.

9. IANA Considerations

There are no IANA impacts in this memo.

10. Contributors

The content and concepts within this document are a product of the
discussion carried out in the Design Team.

Michael Ramalho provided the text for the scenario discussed in
Appendix A.

11. Acknowledgements

Much of this document is derived from previous work on congestion
control at the IETF.

The authors would like to thank Harald Alvestrand, Luca De Cicco,
Wesley Eddy, Lars Eggert, Kevin Gross, Vinayak Hegde, Stefan Holmer,
Randell Jesup, Piers O'Hanlon, Colin Perkins, Michael Ramalho,
Zaheduzzaman Sarker, Timothy B. Terriberry, Michael Welzl, and Mo
Zanaty for providing valuable feedback on earlier versions of this
draft. Additionally, also thank the participants of the design team
for their comments and discussion related to the evaluation criteria.

12. References

12.1. Normative References

   [RFC3550]   Schulzrinne, H., Casner, S., Frederick, R., and V.
               Jacobson, "RTP: A Transport Protocol for Real-Time
               Applications", STD 64, RFC 3550, July 2003.

   [RFC3551]   Schulzrinne, H. and S. Casner, "RTP Profile for Audio and
               Video Conferences with Minimal Control", STD 65, RFC 3551,
               July 2003.

   [RFC3611]   Friedman, T., Caceres, R., and A. Clark, "RTP Control
               Protocol Extended Reports (RTCP XR)", RFC 3611, November
               2003.

   [RFC4585]   Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey,
               "Extended RTP Profile for Real-time Transport Control
               Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July
               2006.

   [RFC5506]   Johansson, I. and M. Westerlund, "Support for Reduced-Size
               Real-Time Transport Control Protocol (RTCP): Opportunities
               and Consequences", RFC 5506, April 2009.

[I-D.jesup-rmcat-reqs]
          Jesup, R., "Congestion Control Requirements For RMCAT",
          draft-jesup-rmcat-reqs-01 (work in progress), February
          2013.

[I-D.ietf-avtcore-rtp-circuit-breakers]
          Perkins, C. and V. Singh, "RTP Congestion Control: Circuit
          Breakers for Unicast Sessions", draft-ietf-avtcore-rtp-
          circuit-breakers-01 (work in progress), October 2012.

12.2. Informative References

[RFC5033] Floyd, S. and M. Allman, "Specifying New Congestion
          Control Algorithms", BCP 133, RFC 5033, August 2007.

[RFC5166] Floyd, S., "Metrics for the Evaluation of Congestion
          Control Mechanisms", RFC 5166, March 2008.

[RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion
          Control", RFC 5681, September 2009.

[SA4-EVAL]
          R1-081955, 3GPP., "LTE Link Level Throughput Data for SA4
          Evaluation Framework", 3GPP R1-081955, 5 2008.

[SA4-LR]   S4-050560, 3GPP., "Error Patterns for MBMS Streaming over
          UTRAN and GERAN", 3GPP S4-050560, 5 2008.

[TCP-eval-suite]
          Lachlan, A., Marcondes, C., Floyd, S., Dunn, L., Guillier,
          R., Gang, W., Eggert, L., Ha, S., and I. Rhee, "Towards a
          Common TCP Evaluation Suite", Proc. PFLDnet. 2008, August
          2008.

Appendix A. Proposal to evaluate Self-fairness of RMCAT congestion control algorithm

The goal of the experiment discussed in this section is to initially
take out as many unknowns from the scenario. Later experiments can
define more complex environments, topologies and media behavior.
This experiment evaluates the performance of the RMCAT sender
competing with other similar RMCAT flows (running the same algorithm
or other RMCAT proposals) on the bottleneck link. There are up to 20
RMCAT flows competing for capacity, but the media only flows in one
direction, from senders (S1..S20) to receivers (R1..R20) and the
feedback packets flow in the reverse direction.

Figure 2 shows the experiment setup and it has subtle differences
compared to the simple topology in Figure 1. Groups of 10 receivers
are connected to the bottleneck link through two different routers
(Router C and D). The rationale for adding these additional routers
is to create two delay legs, i.e., two groups of endpoints with
different network latencies and measure the performance of the RMCAT
congestion control algorithm. If fewer than 10 sources are
initialized, all traffic flows experience the same delay because they
share the same delay leg.

Router A has a single forward direction bottleneck link (i.e., the
bottleneck capacity and delay constraints applies only to the media
packets going from the sender to the receiver, the feedback packets
are unaffected). Hence, the Round-Trip Time (RTT) is primarily
composed of the bottleneck queue delay and any forward path
(propagation) latency. The main reason for not applying any
constraints on the return path is to provide the best-case
performance scenario for the congestion control algorithm. In later
experiments, it is possible to add similar capacity and delay
constraints on the return path.

```
                                                        +---+
                                                 / === |R1 |
    +---+                             +-----+  //       +---+
    |S1 |====== \                    / =|  C  | //
    +---+        \\                 //  +-----+ \\      (...)
                  \\               //            \\
    +---+      +-----+   Bottleneck  +-----+         \\    +---+
    |S2 |======|  A  |-------------------->|  B  |          \ === |R10|
    +---+      |     |<--------------------|     |          +---+
               +-----+        Link         +-----+
    (...)         //                      \\            +---+
                 //                        \\          / === |R11|
    +---+      //                           \\ +-----+ //    +---+
    |S20|===== /                             \ =|  D  |//
    +---+                                       +-----+\\    (...)
                                                     \\
                                                      \\   +---+
                                                       \ ===|R20|
                                                         +---+
```
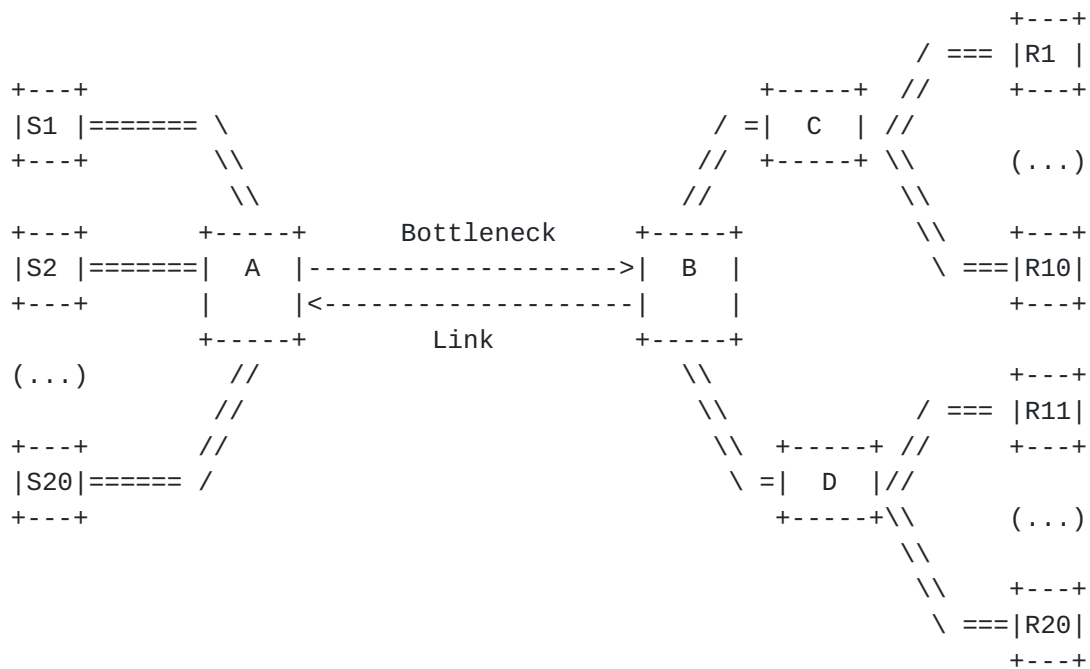
              Figure 2: Self-fairness Evaluation Setup

Loss impairments are applied at Router C and Router D, but only to
the feedback flows. If the losses are set to 0%, it represents a
case where the return path is over-provisioned for all traffic. In
later experiments the loss impairments can be added to the media path
as well.

A.1. Evaluation Parameters

A.1.1. Media Traffic Generator

   The media source always generates at the rate requested by the
   congestion control and has infinite data to send. Furthermore, the
   media packet generator is subject to the following constraints:

   1. It MUST emit a packet at least once per 100 ms time interval.

   2. For low media rate source: when generating data at a rate less
       than a maximum length MTU every 100 ms would allow (e.g., 120
       kbps = 1500 bytes/packet * 10 packets/sec * 8 bits/byte), the
       RMCAT source must modulate the packet size (RTP payload size) of
       RTP packets that are sent every 100 ms to attain the desired
       rate.

   3. For high media rate sources: when generating data at a rate
       greater than a maximum length MTU every 100 ms would allow, the
       source must do so by sending (approximately) maximum MTU sized
       packets and adjusting the inter-departure interval to be
       approximately equal. The intent of this to ensure the data is
       sent relatively smoothly independent of the bit rate, subject to
       the first constraint.

A.1.2. Bottleneck Link Bandwidth

   The bottleneck link capacity is dimensioned such that each RMCAT flow
   in an ideal situation with perfectly equal capacity sharing for all
   the flows on the bottleneck obtains the following throughputs: 200
   kbps, 800 kbps, 1.3 Mbps and 4 Mbps.
   For example, experiments with five RMCAT flows with an 800 kbps/flow
   target rate should set the bottleneck link capacity to 4 Mbps.

A.1.3. Bottleneck Link Queue Type and Length

   The bottleneck link queue (Router A) is a simple FIFO queue having a
   buffer length corresponding to 70 ms, 500 ms or 2000 ms (defined in
   Section 6.4) of delay at the bottleneck link rate (i.e., actual
   buffer lengths in bytes are dependent on bottleneck link bandwidth).

A.1.4. RMCAT flows and delay legs

   Experiments run with 1, 3, 5, 10 and 20 RMCAT sources, they are
   outlined as follows:

1. Experiments with 1, 3, and 5 RMCAT flows, all RMCAT flows
    commence simultaneously. A single delay leg is used and the link
    latency is set to one of the following : 0 ms, 50 ms and 150 ms.

2. For 10 and 20 source experiments where all RMCAT flows begin
    simultaneously the sources are split evenly into two different
    bulk delay legs. One leg is set to 0 ms bulk delay leg and the
    other is set to 150 ms.

3. For 10 and 20 source experiments where the first set will use 0
    ms of bulk delay and the second set will use 150 ms bulk delay.

    1. Random starts within interval [0 ms, 500 ms].

    2. One "early-coming" flow (i.e., the 1st flow starting and
        achieving steady-state before the next N-1 simultaneously
        begin).

    3. One "late-coming" flow (i.e., the Nth flow starting after
        steady-state has occurred for the existing N-1 flows).

    These cases assess if there are any early or late-comer
    advantages or disadvantages for a particular algorithm and to see
    if any unfairness is reproducible or unpredictable.

[Open issue (A.1): which group does the early and late flow belong
to?]

[Open issue (A.2): Start rate for the media flows]

A.1.5. Impairment Generator

Packet loss is created in the reverse path (affects only feedback
packets). Cases of 0%, 1%, 5% and 10% are studied for the 1, 3, and
5 RMCAT flow experiments, losses are not applied to flows with 10 or
20 RMCAT flows.

A.2. Proposed Passing Criteria

[Editor's note: there has been little or no discussion on the below
criteria, however, they are listed here for the sake of completeness.

No unfairness is observed, i.e., at steady state each flow attains a
throughput between [B/(3*N), (3*B)/N], where B is the link
bandwidth and N is the number of flows.

No flow experiences packet loss when queue length is set to 500 ms or
greater.

All individual sources must be in their steady state within twenty
LRTTs (where LRTT is defined as the RTT associated with the flow with
the Largest RTT in the experiment).]

A.3. Extensability of the Experiment

The above scenario describes only RMCAT sources competing for
capacity on the bottleneck link, however, future experiments can use
different types of cross-traffic (as described in Section 6.1).

Currently, the forward path (carrying media packets) is characterized
to add delay and a fixed bottleneck link capacity, in the future
packet losses and capacity changes can be applied to mimic a wireless
link layer (for e.g., WiFi, 3G, LTE). Additionally, only losses are
applied to the reverse path (carrying feedback packets), later
experiments can apply the same forward path (carrying media packets)
impairments to the reverse path.

Appendix B. Change Log

Note to the RFC-Editor: please remove this section prior to
publication as an RFC.

B.1. Changes in draft-singh-rmcat-cc-eval-03

o Incorporate the discussion within the design team.

o Added a section on evaluation parameters, it describes the flow
   and network characteristics.

o Added Appendix with self-fairness experiment.

B.2. Changes in draft-singh-rmcat-cc-eval-02

o Added scenario descriptions.

B.3. Changes in draft-singh-rmcat-cc-eval-01

o Removed QoE metrics.

o Changed stability to steady-state.

o Added measuring impact against few and many flows.

o Added guideline for idle and data-limited periods.

o Added reference to TCP evaluation suite in example evaluation
   scenarios.

Authors' Addresses

    Varun Singh
    Aalto University
    School of Electrical Engineering
    Otakaari 5 A
    Espoo, FIN 02150
    Finland

    Email: varun@comnet.tkk.fi
    URI: http://www.netlab.tkk.fi/~varun/

    Joerg Ott
    Aalto University
    School of Electrical Engineering
    Otakaari 5 A
    Espoo, FIN 02150
    Finland

    Email: jo@comnet.tkk.fi