

RMCAT WG
Internet-Draft
Intended status: Informational
Expires: September 11, 2015

V. Singh
J. Ott
Aalto University
March 10, 2015

Evaluating Congestion Control for Interactive Real-time Media draft-ietf-rmcat-eval-criteria-03

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 September 11, 2015.

Copyright Notice

Copyright (c) 2015 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
2.	Terminology	3
3.	Metrics	3
3.1.	RTP Log Format	5
4.	Guidelines	5
4.1.	Avoiding Congestion Collapse	5
4.2.	Stability	5
4.3.	Media Traffic	5
4.4.	Start-up Behaviour	6
4.5.	Diverse Environments	6
4.6.	Varying Path Characteristics	6
4.7.	Reacting to Transient Events or Interruptions	7
4.8.	Fairness With Similar Cross-Traffic	7
4.9.	Impact on Cross-Traffic	7
4.10.	Extensions to RTP/RTCP	7
5.	List of Network Parameters	7
5.1.	One-way Propagation Delay	7
5.2.	End-to-end Loss	8
5.3.	DropTail Router Queue Length	8
5.4.	Loss generation model	8
5.5.	Jitter models	9
5.5.1.	Random Bounded PDV (RBPDV)	9
5.5.2.	Approximately Random Subject to No-Reordering Bounded PDV (NR-RPVD)	10
6.	Traffic Models	11
6.1.	TCP traffic model	11
6.2.	RTP Video model	12
7.	Security Considerations	12
8.	IANA Considerations	12
9.	Contributors	12
10.	Acknowledgements	12
11.	References	12
11.1.	Normative References	12
11.2.	Informative References	13
Appendix A.	Application Trade-off	14
A.1.	Measuring Quality	14
Appendix B.	Change Log	14
B.1.	Changes in draft-ietf-rmcat-eval-criteria-02	14
B.2.	Changes in draft-ietf-rmcat-eval-criteria-02	14
B.3.	Changes in draft-ietf-rmcat-eval-criteria-01	14
B.4.	Changes in draft-ietf-rmcat-eval-criteria-00	15
B.5.	Changes in draft-singh-rmcat-cc-eval-04	15
B.6.	Changes in draft-singh-rmcat-cc-eval-03	15
B.7.	Changes in draft-singh-rmcat-cc-eval-02	15
B.8.	Changes in draft-singh-rmcat-cc-eval-01	15
	Authors' Addresses	16

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.ietf-rmcat-cc-requirements](#)]). 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. The minimal requirement for RMCAT proposals is to produce or present results for the test scenarios described in [[I-D.ietf-rmcat-eval-test](#)] (Basic Test Cases). 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

Each experiment is expected to log 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 (measured at 200ms intervals)
2. Packets sent, Packets received
3. Bytes sent, bytes received
4. Packet delay

5. Packets lost, Packets discarded (from the playout or de-jitter buffer)
6. If using, retransmission or FEC: post-repair loss
7. Fairness or Unfairness: Experiments testing the performance of an RMCAT proposal against any cross-traffic must define its expected criteria for fairness. The "unfairness" test guideline (measured at 1s intervals) is:
 1. Does not trigger the circuit breaker.
 2. No RMCAT stream achieves more than 3 times the average throughput of the RMCAT stream with the lowest average throughput, for a case when the competing streams have similar RTTs.
 3. RTT should not grow by a factor of 3 for the existing flows when a new flow is added.For example, see the test scenarios described in [\[I-D.ietf-rmcat-eval-test\]](#).
8. Convergence time: The time taken to reach a stable rate at startup, after the available link capacity changes, or when new flows get added to the bottleneck link.
9. Instability or oscillation in the sending rate: The frequency or number of instances when the sending rate oscillates between an high watermark level and a low watermark level, or vice-versa in a defined time window. For example, the watermarks can be set at 4x interval: 500 Kbps, 2 Mbps, and a time window of 500ms.
10. Bandwidth Utilization, defined as ratio of the instantaneous sending rate to the instantaneous bottleneck capacity. This metric is useful only when an RMCAT flow is by itself or competing with similar cross-traffic.

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.

[Open issue (1): Using Jain-fairness index (JFI) for measuring self-fairness between RTP flows? measured at what intervals? visualized as a CDF or a timeseries? Additionally: Use JFI for comparing fairness between RTP and long TCP flows?]

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
```

If the congestion control implements, retransmissions or FEC, the evaluation should report both packet loss (before applying error-resilience) and residual packet loss (after applying error-resilience).

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

The congestion control algorithm is expected to take an action, such as reducing the sending rate, when it detects congestion. Typically, it should intervene before the circuit breaker [[I-D.ietf-avtcore-rtp-circuit-breakers](#)] is engaged.

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, varying amount of motion for video, or bursty nature of I-frames.

The evaluation may be done in two stages. In the first stage, the endpoint generates traffic at the rate calculated by the congestion controller. In the second stage, real codecs or models of video codecs are used to mimic application-limited data periods and varying video frame sizes.

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 test scenarios, such as when competing with varying amount of cross-traffic or how quickly does the congestion control algorithm achieve a stable sending rate.

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) 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 bottleneck link and a router's queue length. For the moment, only DropTail queues are used. However, if new Active Queue Management (AQM) schemes become available, the performance of the congestion control algorithm should be again evaluated.

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.

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.

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. List of Network Parameters

The implementors initially are encouraged to choose evaluation settings from the following values:

5.1. One-way Propagation Delay

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. Typical values are:

1. Very low latency: 0-1ms
2. Low latency: 50ms

3. High latency: 150ms
4. Extreme latency: 300ms

5.2. End-to-end Loss

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%

5.3. DropTail Router Queue Length

The router queue length is measured as the time taken to drain the FIFO queue. It has been noted in various discussions that the queue length in the current deployed Internet varies significantly. While the core backbone network has very short queue length, the home gateways usually have larger queue length. Those various queue lengths can be categorized in the following way:

1. QoS-aware (or short): 70ms
2. Nominal: 300-500ms
3. Buffer-bloated: 1000-2000ms

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

$$\text{QueueSize (in bytes)} = \text{QueueSize (in sec)} \times \text{Throughput (in bps)} / 8$$

5.4. Loss generation model

[Editor's note : Describes the model for generating packet losses, for example, losses can be generated using traces, or using the Gilbert-Elliott model, or randomly (uncorrelated loss).]

5.5. Jitter models

This section defines jitter model for the purposes of this document. When jitter is to be applied to both the RMCAT flow and any competing flow (such as a TCP competing flow), the competing flow will use the jitter definition below that does not allow for re-ordering of packets on the competing flow (see NR-RBPDV definition below).

Jitter is an overloaded term in communications. Its meaning is typically associated with the variation of a metric (e.g., delay) with respect to some reference metric (e.g., average delay or minimum delay). For example, [RFC 3550](#) jitter is a smoothed estimate of jitter which is particularly meaningful if the underlying packet delay variation was caused by a Gaussian random process.

Because jitter is an overloaded term, we instead use the term Packet Delay Variation (PDV) to describe the variation of delay of individual packets in the same sense as the IETF IPPM WG has defined PDV in their documents (e.g., [RFC 3393](#)) and as the ITU-T SG16 has defined IP Packet Delay Variation (IPDV) in their documents (e.g., Y.1540).

Most PDV distributions in packet network systems are one-sided distributions (the measurement of which with a finite number of measurement samples result in one-sided histograms). In the usual packet network transport case there is typically one packet that transited the network with the minimum delay, then a majority of packets also transit the system within some variation from this minimum delay, and then a minority of the packets transits the network with delays higher than the median or average transit time (these are outliers). Although infrequent, outliers can cause significant deleterious operation in adaptive systems and should be considered in RMCAT adaptation designs.

In this section we define two different bounded PDV characteristics, 1) Random Bounded PDV and 2) Approximately Random Subject to No-Reordering Bounded PDV.

5.5.1. Random Bounded PDV (RBPDV)

The RBPDV probability distribution function (pdf) is specified to be of some mathematically describable function which includes some practical minimum and maximum discrete values suitable for testing. For example, the minimum value, x_{\min} , might be specified as the minimum transit time packet and the maximum value, x_{\max} , might be defined to be two standard deviations higher than the mean.

Since we are typically interested in the distribution relative to the mean delay packet, we define the zero mean PVD sample, $z(n)$, to be $z(n) = x(n) - x_{\text{mean}}$, where $x(n)$ is a sample of the RBPVD random variable x and x_{mean} is the mean of x .

We assume here that $s(n)$ is the original source time of packet n and the post-jitter induced emission time, $j(n)$, for packet n is $j(n) = \{[z(n) + x_{\text{mean}}] + s(n)\}$. It follows that the separation in the post-jitter time of packets n and $n+1$ is $\{[s(n+1)-s(n)] - [z(n)-z(n+1)]\}$. Since the first term is always a positive quantity, we note that packet reordering at the receiver is possible whenever the second term is greater than the first. Said another way, whenever the difference in possible zero mean PDV sample delays (i.e., $[x_{\text{max}} - x_{\text{min}}]$) exceeds the inter-departure time of any two sent packets, we have the possibility of packet re-ordering.

There are important use cases in real networks where packets can become re-ordered such as in load balancing topologies and during route changes. However, for the vast majority of cases there is no packet re-ordering because most of the time packets follow the same path. Due to this, if a packet becomes overly delayed, the packets after it on that flow are also delayed. This is especially true for mobile wireless links where there are per-flow queues prior to base station scheduling. Owing to this important use case, we define another PDV profile similar to the above, but one that does not allow for re-ordering within a flow.

5.5.2. Approximately Random Subject to No-Reordering Bounded PDV (NR-RPVD)

No Reordering RPDV, NR-RPVD, is defined similarly to the above with one important exception. Let $\text{serial}(n)$ be defined as the serialization delay of packet n at the lowest bottleneck link rate (or other appropriate rate) in a given test. Then we produce all the post-jitter values for $j(n)$ for $n = 1, 2, \dots N$, where N is the length of the source sequence s to be offset-ed. The exception can be stated as follows: We revisit all $j(n)$ beginning from index $n=2$, and if $j(n)$ is determined to be less than $[j(n-1)+\text{serial}(n-1)]$, we redefine $j(n)$ to be equal to $[j(n-1)+\text{serial}(n-1)]$ and continue for all remaining n (i.e., $n = 3, 4, \dots N$). This models the case where the packet n is sent immediately after packet $(n-1)$ at the bottleneck link rate. Although this is generally the theoretical minimum in that it assumes that no other packets from other flows are in-between packet n and $n+1$ at the bottleneck link, it is a reasonable assumption for per flow queuing.

We note that this assumption holds for some important exception cases, such as packets immediately following outliers. There are a

multitude of software controlled elements common on end-to-end Internet paths (such as firewalls, ALGs and other middleboxes) which stop processing packets while servicing other functions (e.g., garbage collection). Often these devices do not drop packets, but rather queue them for later processing and cause many of the outliers. Thus NR-RPVD models this particular use case (assuming serial(n+1) is defined appropriately for the device causing the outlier) and thus is believed to be important for adaptation development for RMCAT.

[Editor's Note: It may require to define test distributions as well. Example test distribution may include-

1 - Two-sided: Uniform PDV Distribution. Two quantities to define: x_min and x_max.

2 - Two-sided: Truncated Gaussian PDV Distribution. Four quantities to define: the appropriate x_min and x_max for test (e.g., +/- two sigma values), the standard deviation and the mean.

3 - One Sided: TBD]

6. Traffic Models

6.1. TCP traffic model

Long-lived TCP flows will download data throughout the session and are expected to have infinite amount of data to send or receive.

Each short TCP flow is modeled as a sequence of file downloads interleaved with idle periods. Not all short TCPs start at the same time, i.e., some start in the ON state while others start in the OFF state.

The short TCP flows can be modelled in two ways, 1) 100s of flows fetching small (5-20 KB) amounts of data, or 2) 10s of flows fetching slightly larger (100-1000KB) amounts of data.

The idle period is typically derived from an exponential distribution with the mean value of 10 seconds.

[Open issue: short-lived/bursty TCP cross-traffic parameters are still to be agreed upon].

6.2. RTP Video model

[I-D.zhu-rmcat-video-traffic-source] describes two types of video traffic models for evaluating RMCAT candidate algorithms. The first model statistically characterizes the behavior of a video encoder. Whereas the second model uses video traces.

7. Security Considerations

Security issues have not been discussed in this memo.

8. IANA Considerations

There are no IANA impacts in this memo.

9. 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 Jitter model.

10. Acknowledgements

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

The authors would like to thank Harald Alvestrand, Anna Brunstrom, Luca De Cicco, Wesley Eddy, Lars Eggert, Kevin Gross, Vinayak Hegde, Stefan Holmer, Randell Jesup, Mirja Kuehlewind, Karen Nielsen, 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.

11. References

11.1. 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.ietf-rmcat-cc-requirements]
Jesup, R. and Z. Sarker, "Congestion Control Requirements for Interactive Real-Time Media", [draft-ietf-rmcat-cc-requirements-09](#) (work in progress), December 2014.
- [I-D.ietf-avtcore-rtp-circuit-breakers]
Perkins, C. and V. Singh, "Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions", [draft-ietf-avtcore-rtp-circuit-breakers-09](#) (work in progress), March 2015.

11.2. 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.
- [I-D.ietf-rmcat-eval-test]
Sarker, Z., Singh, V., Zhu, X., and M. Ramalho, "Test Cases for Evaluating RMCAT Proposals", [draft-ietf-rmcat-eval-test-00](#) (work in progress), August 2014.
- [I-D.zhu-rmcat-video-traffic-source]
Zhu, X., Cruz, S., and Z. Sarker, "Modeling Video Traffic Sources for RMCAT Evaluations", [draft-zhu-rmcat-video-traffic-source-00](#) (work in progress), October 2014.
- [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.](#) Application Trade-off

Application trade-off is yet to be defined. see RMCAT requirements [[I-D.ietf-rmcat-cc-requirements](#)] document. Perhaps each experiment should define the application's expectation or trade-off.

[A.1.](#) Measuring Quality

No quality metric is defined for performance evaluation, it is currently an open issue. However, there is consensus that congestion control algorithm should be able to show that it is useful for interactive video by performing analysis using a real codec and video sequences.

[Appendix B.](#) Change Log

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

[B.1.](#) Changes in [draft-ietf-rmcat-eval-criteria-02](#)

- o Keep-alive version.
- o Moved link parameters and traffic models from eval-test

[B.2.](#) Changes in [draft-ietf-rmcat-eval-criteria-02](#)

- o Incorporated fairness test as a working test.
- o Updated text on minimum evaluation requirements.

[B.3.](#) Changes in [draft-ietf-rmcat-eval-criteria-01](#)

- o Removed [Appendix B](#).
- o Removed Section on Evaluation Parameters.

B.4. Changes in [draft-ietf-rmcat-eval-criteria-00](#)

- o Updated references.
- o Resubmitted as WG draft.

B.5. Changes in [draft-singh-rmcat-cc-eval-04](#)

- o Incorporate feedback from IETF 87, Berlin.
- o Clarified metrics: convergence time, bandwidth utilization.
- o Changed fairness criteria to fairness test.
- o Added measuring pre- and post-repair loss.
- o Added open issue of measuring video quality to appendix.
- o clarified use of DropTail and AQM.
- o Updated text in "Minimum Requirements for Evaluation"

B.6. 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.
- o Changed bottleneck parameters from a proposal to an example set.
- o

B.7. Changes in [draft-singh-rmcat-cc-eval-02](#)

- o Added scenario descriptions.

B.8. 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

