

# IRTF Internet Congestion Control Research Group

Chairs:

Michael Welzl, David Ros

ICCRG @ 91st IETF Meeting  
11 November 2014

# IRTF IPR Policy

**The IRTF follows the IETF Intellectual Property Rights (IPR) disclosure rules.** This is a summary of these rules as they relate to IRTF research group discussions, mailing lists and Internet Drafts:

- If you include your own or your employer’s IPR in a contribution to an IRTF research group, then you must file an IPR disclosure with the IETF.
- If you recognize your own or your employer’s IPR in someone else’s contribution and you are participating in the discussions in the research group relating to that contribution, then you must file an IPR disclosure with the IETF. Even if you are not participating in the discussion, the IRTF still requests that you file an IPR disclosure with the IETF.
- Finally, the IRTF requests that you file an IPR disclosure with the IETF if you recognize IPR owned by others in any IRTF contribution.

**The IRTF expects that you file IPR disclosures in a timely manner, i.e., in a period measured in days or weeks, not months. The IRTF prefers that the most liberal licensing terms possible are available for IRTF Stream documents, see RFC 5743.** You may file an IPR disclosure here: <http://www.ietf.org/ipr/file-disclosure>

See RFC 3979 (BCP 79) for definitions of “IPR” and “contribution” and for the detailed rules (substituting “IRTF” for “IETF”).

# Evaluation suite – need reviewers!

- <http://tools.ietf.org/html/draft-irtf-iccr-g-tcpeval-01>
- [http://tools.ietf.org/group/irtf/trac/wiki/ICCRG\\_tcpeval](http://tools.ietf.org/group/irtf/trac/wiki/ICCRG_tcpeval)
- This is an old, already much used / cited document, taken over from TMRG
  - Meant for comparing high-speed TCP variants (CUBIC etc.)
  - Our charter talks about high-speed TCPs and then says: “Such an evaluation is especially difficult, given that some mechanisms involve changes to queuing algorithms, or to packet forwarding, whereas others only require end-system modifications. It is difficult therefore to compare like with like. One goal of the ICCRG is to improve our methods for evaluating transport protocols. To this end, a set of simulation (or other) test suites will be developed.”

# Agenda

3 major topics today (40 minutes for each):

1. Discussion: Should the Initial Window hang loose?
2. RMCAT: coping with link capacity changes
3. Lowering queuing delay

Discussion: Should the Initial Window hang loose?

- (Michael filling in for) Jianjie Yu - Configuring TCP's Initial Window ( draft-you-tcpm-configuring-tcp-initial-window-02 ) (20 min)
- Also consider: draft-sallantin-tcpm-initial-spreading-00  
<http://www.ietf.org/mail-archive/web/tcpm/current/msg09308.html>

# Agenda /2

## RMCAT: coping with link capacity changes

- Michael Ramalho - Coupling Discrete Time Events to Continuous Time in RMCAT and Reasonable Bounds on the Time Rate of Change in Available Capacity (20 min)
- Geert Van Der Auwera - A Delay Recovery Phase for RMCAT flows (20 min)

## Lowering queuing delay

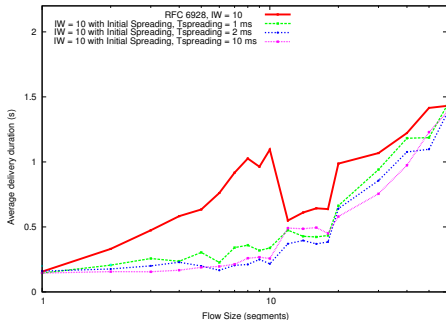
- Toke Hoeiland-Joergensen - How to reliably measure the performance of modern AQMs - and what comes of doing so? (20 min)
- Mirja Kuehlewind - TCP SIAD: Congestion Control supporting High Speed and Low Latency (20 min)

# IW discussion

- Much debate preceding IW10 (Exp. RFC 6928)
  - Ongoing experiment: Linux default, servers do it
  - Reports of different values, larger but also smaller
  - Reports of problems on some links (mostly before publication of RFC6928)
- Can we maintain large-IW-gains while somehow weakening its potential negative impacts?
  - At least two proposals on the table

# Initial Spreading : a mechanism to improve the short-lived flows performance

1. The RTT is measured during the SYN-SYN/ACK exchange
2. The IW segments are sent each  $T_{spreading}$ . Based on the RTT,  $T_{spreading}$  is computed to verify the following properties:
  - ▶ It MUST be large enough to ensure that segment losses remain un-correlated
  - ▶ it MUST be the shortest as possible to not add an unnecessary delay
3. After the transmission of the IW, the regular TCP algorithm is used



average delivery duration obtained in a congested network (Linux Kernel experiments)

## Initial Spreading enables a safe increase of the TCP's Initial Window

1. Performs better than conservative and aggressive IW in congested network
  - ▶ reduces the detrimental burst consequences and so:
    - ▶ reduces the loss probability
    - ▶ increases the possibilities to use recovery mechanisms

*The average delivery latency is reduced by more than 30% for short-lived flows ( $\leq 10$  segments)*
2. Performs similarly to aggressive IW in uncongested network
  - ▶ adds a minimal latency to the ideal solution that consists in sending the whole flow in one shot

*draft-sallantin-tcpm-initial-spreading-00*

R.Sallantin, C.Baudoin, E.Chaput, F.Arnal, E.Dubois and A- L.Beylot, "Initial spreading: A fast Start-Up TCP mechanism," Local Computer Networks (LCN), 2013

R.Sallantin, C.Baudoin, E.Chaput, F.Arnal, E.Dubois and A- L.Beylot, "A TCP model for short-lived flows to validate initial spreading," Local Computer Networks (LCN), 2014