

TCP Extensions for Immediate Retransmissions

draft-eggert-tcpm-retransmit-now-02

IETF-63, Paris, France

Lars Eggert, Simon Schütz and Stefan Schmid

NEC

Tuesday, August 2, 2005

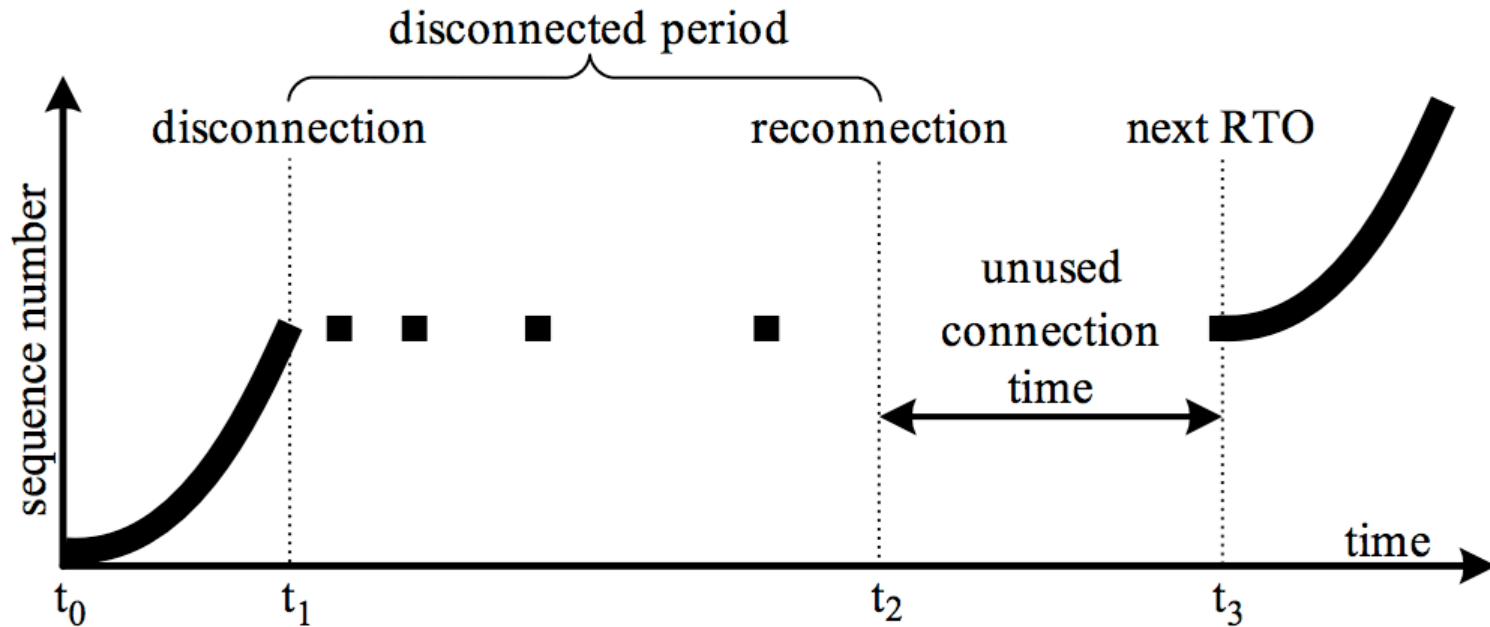
Motivation

- connectivity disruptions can occur along an end-to-end path
- when TCP resumes transmission, it can be too aggressive or very inefficient, depending on the disruption type
- draft classifies disruptions, describes issues and suggests solutions for some disruption types
- other types are already addressed by draft-swami-tcp-lmdr-05

After a Disruption

- TCP transmission behavior after a disruption depends on its duration
- “short” disruptions $< \text{RTO}$ look like short bursts of losses; modern TCPs can recover without slow-start
- for “long” disruptions $> \text{RTO}$, TCP performs slow-start

Long Disruptions



- TCP slow-starts, *i.e.*, re-probes path
- problem: slow-start attempts are exponentially timed, due to RTO backoff
- inefficient! wastes connectivity time, which

Connectivity Indicators

- idea: add speculative retransmission attempt on “connectivity indicator” (CI)
- CIs signal that connectivity to the peer may be restored
 - example: link-layer events on end hosts
- draft describes how TCP uses generic CIs, does not define CIs themselves

Retransmit Now

- when receiving a CI, a host that has data to send starts (re-)sending it
- the peer may then also retransmit outstanding data, if needed
- without queued data, host needs to signal the peer in a different way

Signal Variants

- draft describes two variants for peer signaling
 - implicitly, by generating triple-duplicate ACKs
 - explicitly, through new TCP option
- recent interest has focused on second variant, because it nicely complements LMDR
- (related work: quickstart, Caceres/lftode, etc.)

draft-swami-tcp-lmldr-05

- “short” disruptions $<$ RTO look like short bursts of losses; modern TCPs can recover without slow-start
- problem: disruption may be due to mobility
 - path characteristics can change! TCP may be too aggressive
- new TCP “slow-start now” option
- (similar technique was discussed for DCCP)

Steps Forward

- LMDR and retransmit-now play in the same area but are orthogonal and complementary
- I don't speak for the LMDR authors!
 - but got off-list feedback from them that thinking about combining things is interesting
 - there are bits available in the LMDR option :-)
- what does the WG think of all this?