

# Start Me Up: Determining and Sharing TCP's Initial Congestion Window

*Safiqul Islam and Michael Welzl*

*Department of Informatics*


*University of Oslo*

# Motivation

- *Each TCP connection maintains states in a data structure called Transport Control Block (TCB)*
- *Sharing TCB across parallel connections and combining their congestion controllers between two endpoints can be beneficial*
  - *Reduce the **F**low **C**ompletion **T**ime (FCT) of short flows: skipping slow start, immediately using large cwnd, applying priorities*
- *Do parallel connections follow the same route?*
  - *When they are encapsulated, e.g. VPNs; more in [1]*


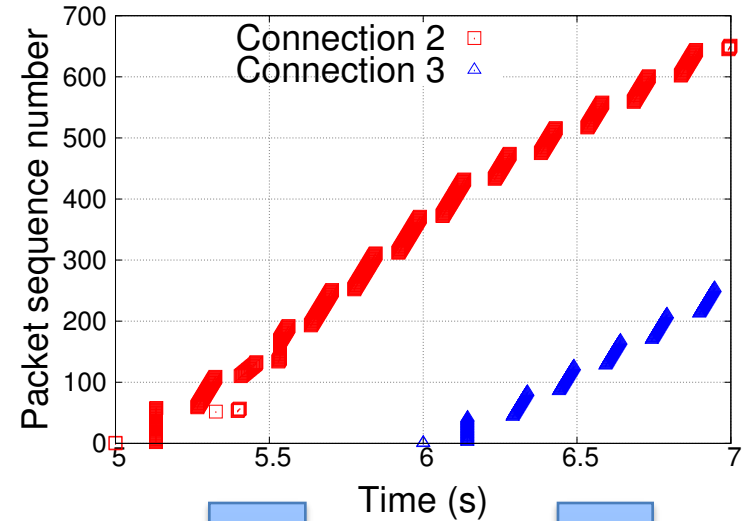
# The Problem

- *Short flows joining an aggregate can immediately increase their cwnds*
  - *Lead to sudden bursts* –  
*if not **paced***

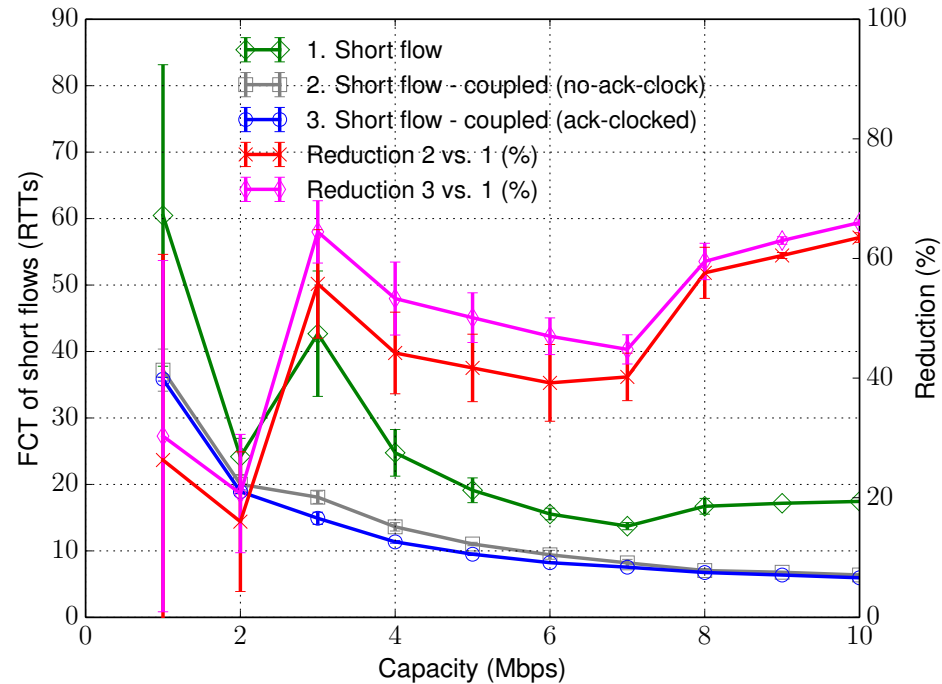


# The Solution

- *Timer based pacing used by prior works*
- *Our approach:*
  - *Maintain the ack-clock of TCP*
  - *Using the ACKs of conn 1 to clock packet transmissions of connection 2 over the course of the first RTT when connection 2 joins*
  - *Similarly, we make use of the ACKs of connections 1 and 2 to clock packet transmissions of connection 3*



# FCTs of Short Flows



*FCT of short flows coupled with our ack-clocked mechanism reduces the FCT*