# Behavior of TCP CUBIC in Low-Latency Mobile Radio Networks

First results of ongoing Master thesis with RWTH Aachen University & Ericsson

Philipp Bruhn and Maciej Muehleisen

# Measurement Setup & Initial Observation

- HTTP server directly connected to LTE core; HTTP client connected via LTE RAN (Mobile Edge Cloud/ Computing type of scenario)
- Host runs Linux[1] on kernel v4.19; TCP CUBIC in default settings
- Ping ~18.5 to ~30ms: Variance originates from LTE MAC layer
- Radio link: Stable signal quality, no interference, no cross traffic

**→ File throughput shows much larger variance than expected**

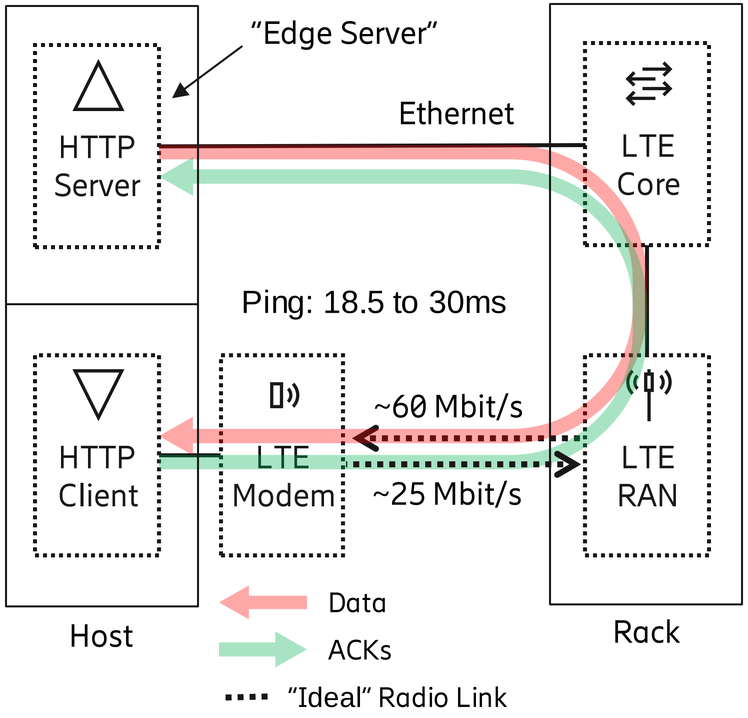


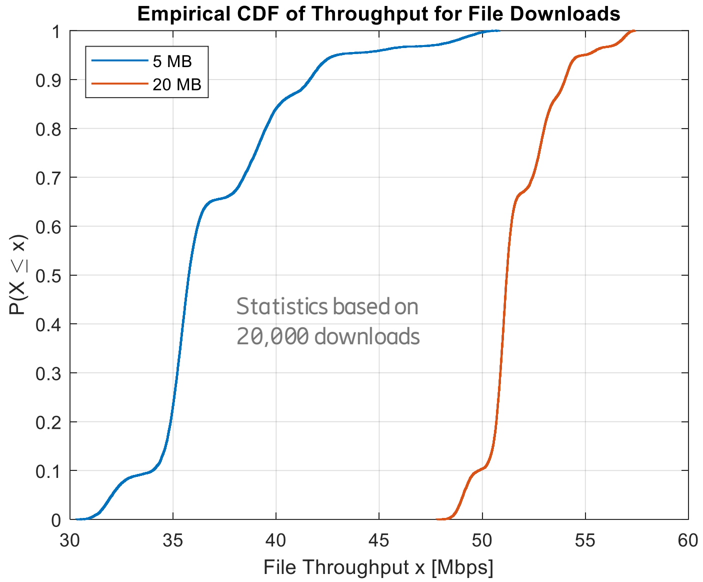


$$\frac{File\ Size}{Transfer\ Time}$$



[1] Debian 10

MAC: Media Access Control | RAN: Radio Access Network

# How HyStart Works

- Two mechanisms to find suitable exit point:
  1. ACK train length → never triggers
  2. Delay increase: $RTT_{cur}$ exceeds $RTT_{min}$ by θ → congestion!
     - $RTT_{cur}$ = smallest RTT sample of first 8 ACKs of each round

       → HyStart never left during sampling phase (SP) of a round
     - Reason 1: Minimum RTT known; current RTT increases (→ leave after SP)
     - Reason 2: Current RTT known; minimum RTT decreases (→ leave directly)

Definition:

$RTT_{min}$ = overall min. RTT

$RTT_{cur}$ = min. RTT in SP

Delay increase:

If: $RTT_{cur} > RTT_{min} + \theta$

$\theta = \langle 4\ ms \mid RTT_{min}/8 \mid 16\ \mathrm{ms} \rangle$

# How HyStart Works

- Two mechanisms to find suitable exit point:
  1. ACK train length → never triggers
  2. Delay increase: $RTT_{cur}$ exceeds $RTT_{min}$ by θ → congestion!
     - $RTT_{cur}$ = smallest RTT sample of first 8 ACKs of each round

       → HyStart never left during sampling phase (SP) of a round
     - Reason 1: Minimum RTT known; current RTT increases (→ leave after SP)
     - Reason 2: Current RTT known; minimum RTT decreases (→ leave directly)

Definition:

$RTT_{min}$ = overall min. RTT

$RTT_{cur}$ = min. RTT in SP

Delay increase:

If: $RTT_{cur} > RTT_{min} + \theta$

$\theta = \langle 4\ ms \mid RTT_{min}/8 \mid 16\ ms \rangle$

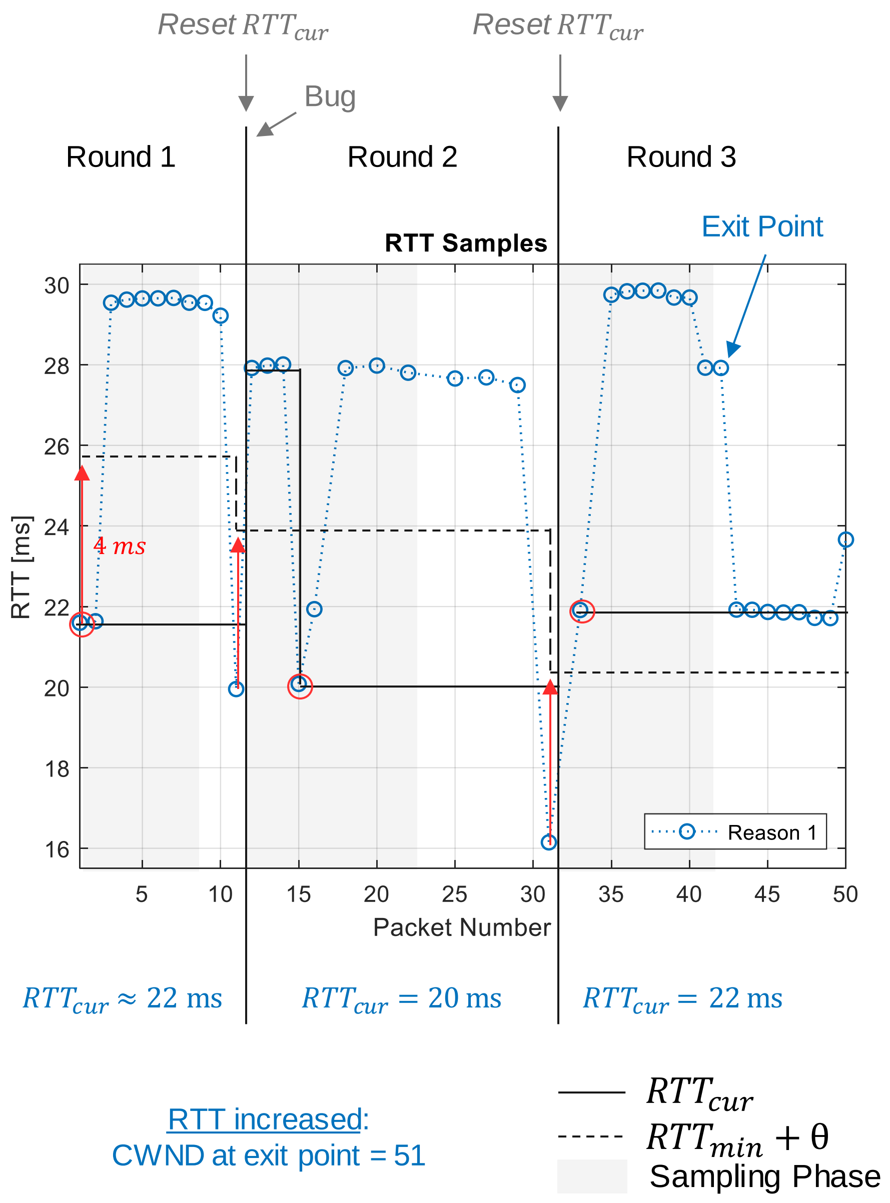

# How HyStart Works

- Two mechanisms to find suitable exit point:
  1. ACK train length → never triggers
  2. Delay increase: $RTT_{cur}$ exceeds $RTT_{min}$ by θ → congestion!
     - $RTT_{cur}$ = smallest RTT sample of first 8 ACKs of each round

       → HyStart never left during sampling phase (SP) of a round
     - Reason 1: Minimum RTT known; current RTT increases (→ leave after SP)
     - Reason 2: Current RTT known; minimum RTT decreases (→ leave directly)

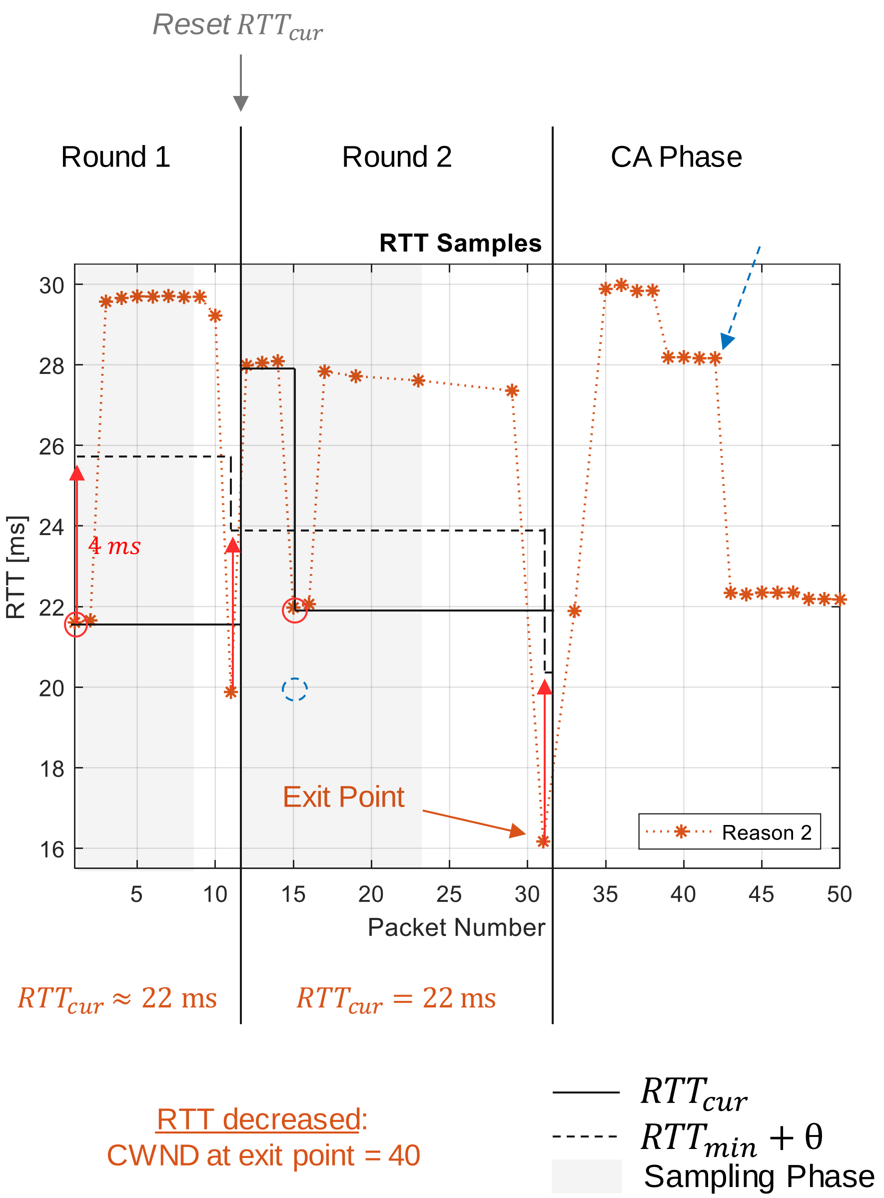


**Definition:**

$RTT_{min}$ = overall min. RTT

$RTT_{cur}$ = min. RTT in SP

**Delay increase:**

If: $RTT_{cur} > RTT_{min} + \theta$

$\theta = \langle 4\ ms \mid RTT_{min}/8 \mid 16\ \text{ms} \rangle$

# Does HyStart Work?

- Two mechanisms to find suitable exit point:
  1. ACK train length → never triggers
  2. Delay increase: $RTT_{cur}$ exceeds $RTT_{min}$ by θ → congestion!
     - $RTT_{cur}$ = smallest RTT sample of first 8 ACKs of each round

       → HyStart never left during sampling phase (SP) of a round
     - Reason 1: Minimum RTT known; current RTT increases (→ leave after SP)
     - Reason 2: Current RTT known; minimum RTT decreases (→ leave directly)

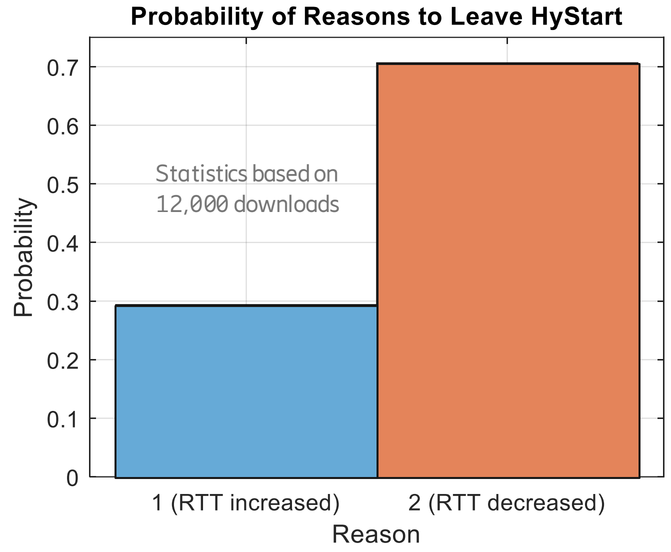


Definition:

$RTT_{min}$ = overall min. RTT

$RTT_{cur}$ = min. RTT in SP

Delay increase:

If: $RTT_{cur} > RTT_{min} + \theta$

$\theta = \langle 4\ ms \mid RTT_{min}/8 \mid 16\ \text{ms} \rangle$

# Uplink Scheduling

- Radio Access Network (RAN) controls all transmissions[2]
- Modem needs uplink resources to transmit data (ACKs)
  1. Modem already has uplink grant:
     → Piggyback buffer status report onto next transmission
  2. Modem does not have uplink grant:
     → Send 1-bit scheduling request in uplink control channel

**→ MAC layer effects primary cause of RTT variation**

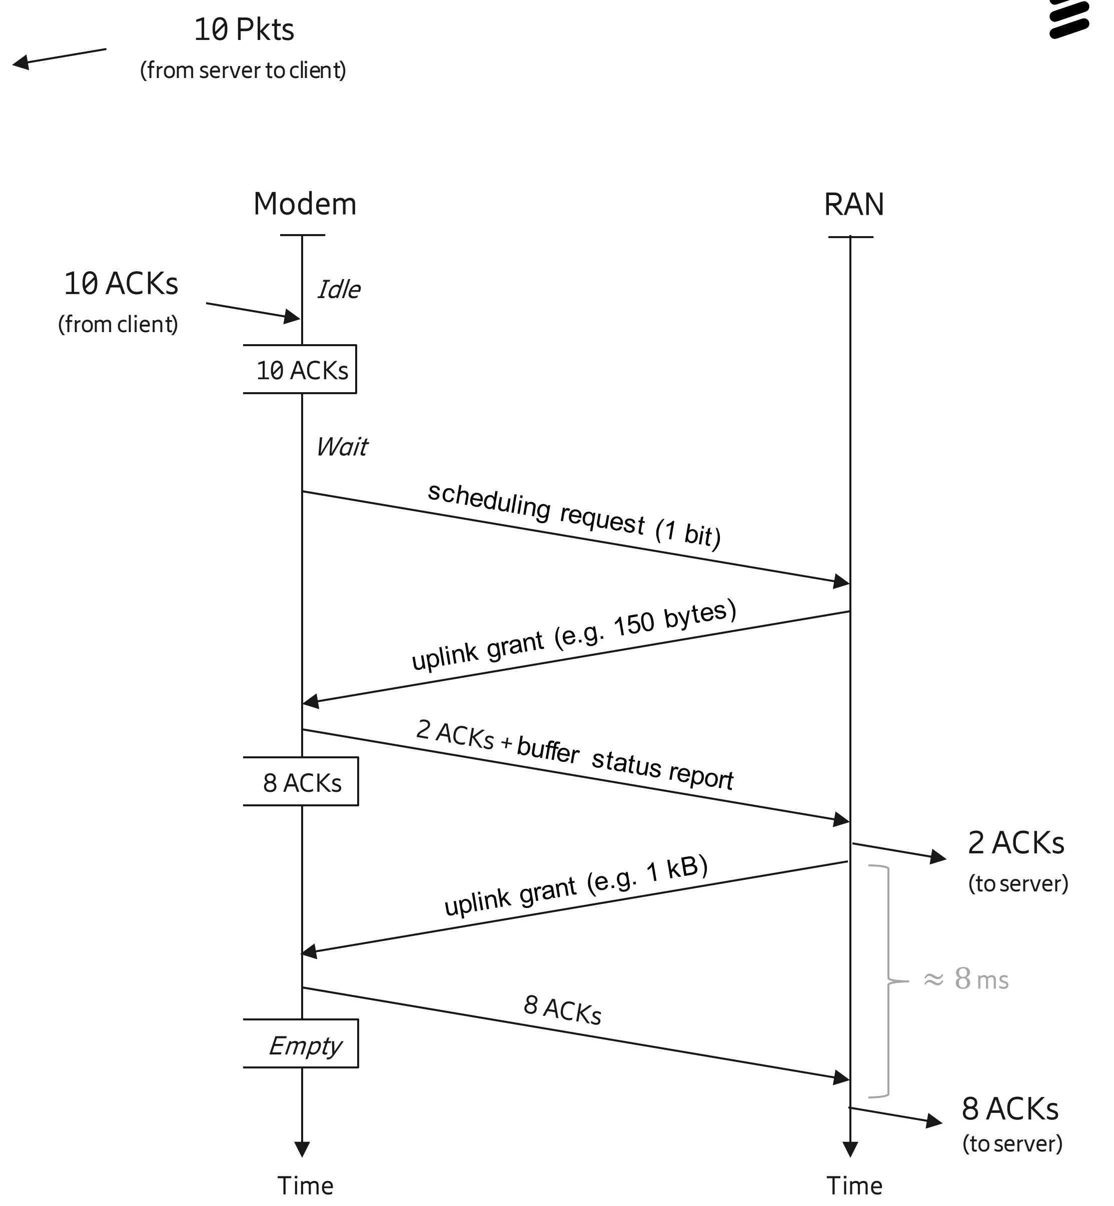


[2] except uplink control information

MAC: Media Access Control | RAN: Radio Access Network

# Why HyStart Fails

- RTT variance due to:
  - Different kinds of uplink scheduling
  - Underlying MAC layer periodicities
- RTT drop causes false positive detection of network congestion / too early exit point

→ Connection far from maximum throughput

**→ MAC layer effects dominate behavior and performance of TCP CUBIC**

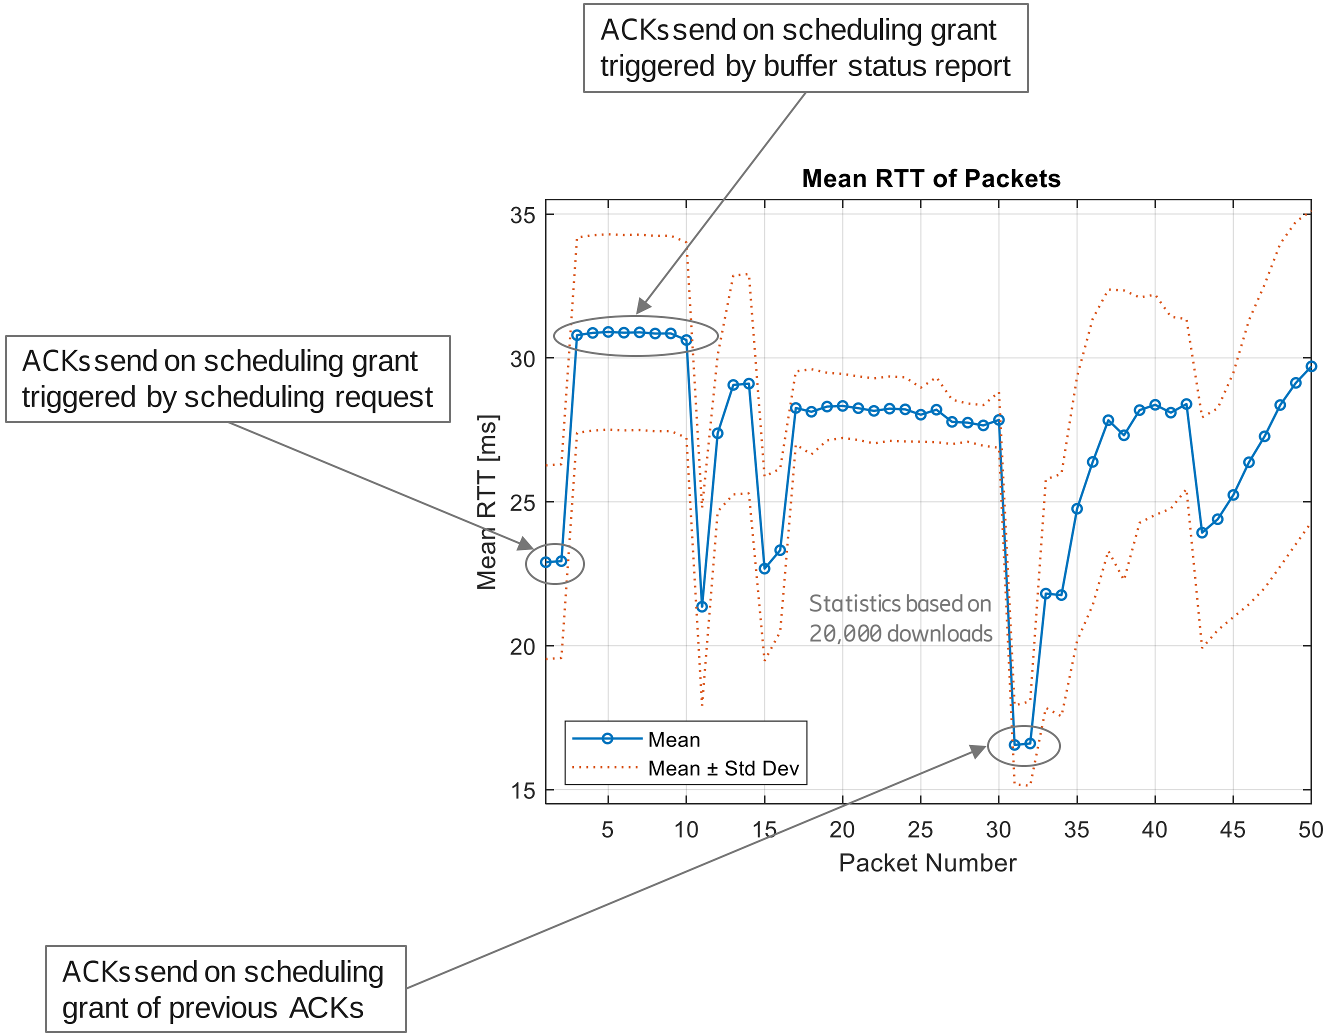




MAC: Media Access Control

# Conclusion

- CUBIC connections may experience significant file throughput variance over LTE
  - HyStart may fail to find suitable exit point due to latency variance of LTE
- Improvements:
  - Bug: HyStart-reset was done one packet too late → fixed
  - Patch: Do not exit in current round due to new $RTT_{min}$
- HyStart does not consider RTT variance → may not be well-suited for mobile radio networks
- Trials conducted with 4G LTE, but 5G New Radio uses same uplink scheduling principles

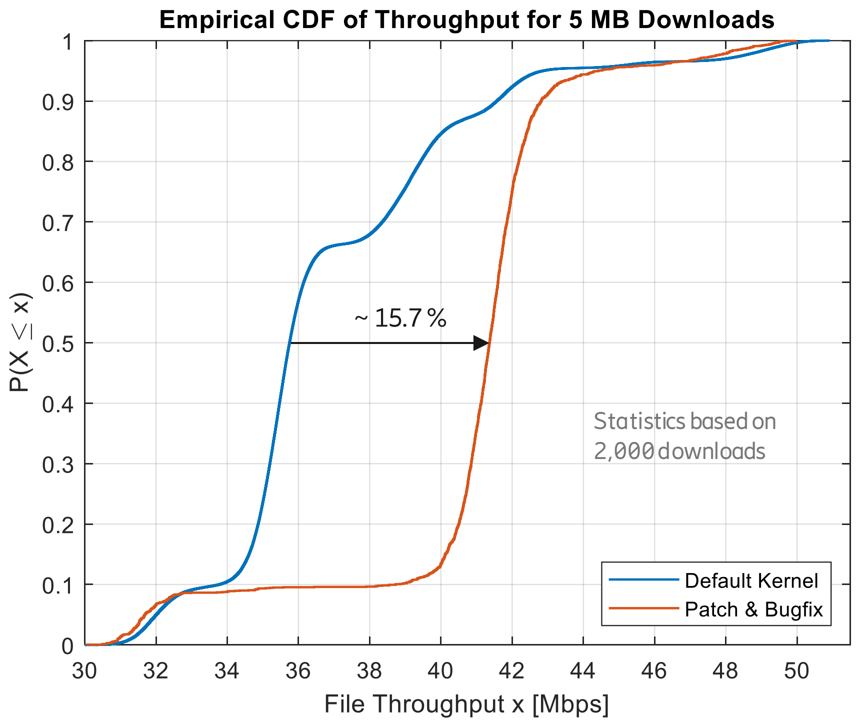


**20 MB**: Improvement by ~**5.4 %**



MAC: Media Access Control

philipp.bruhn@ericsson.com

maciej.muehleisen@ericsson.com

# Appendix

- Kernel patch and bugfix result in larger CWND at exit point → improves mean throughput significantly