How do TCP Session Resumption Dynamics Impact HTTP Adaptive Streaming?

Mihail Yanev

Motivation

- HTTP-based video streaming applications rely on complex rate adaptation algorithms
    - They depend on application-level metrics that don't reflect **transport protocol dynamics**
    - The TCP transport protocol is not well-suited to on-off video flows

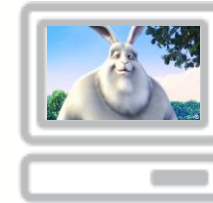- Changing TCP **slow start** allows for simpler applications without affecting performance

HTTP Adaptive Streaming

Server

High quality stream
Medium Quality stream
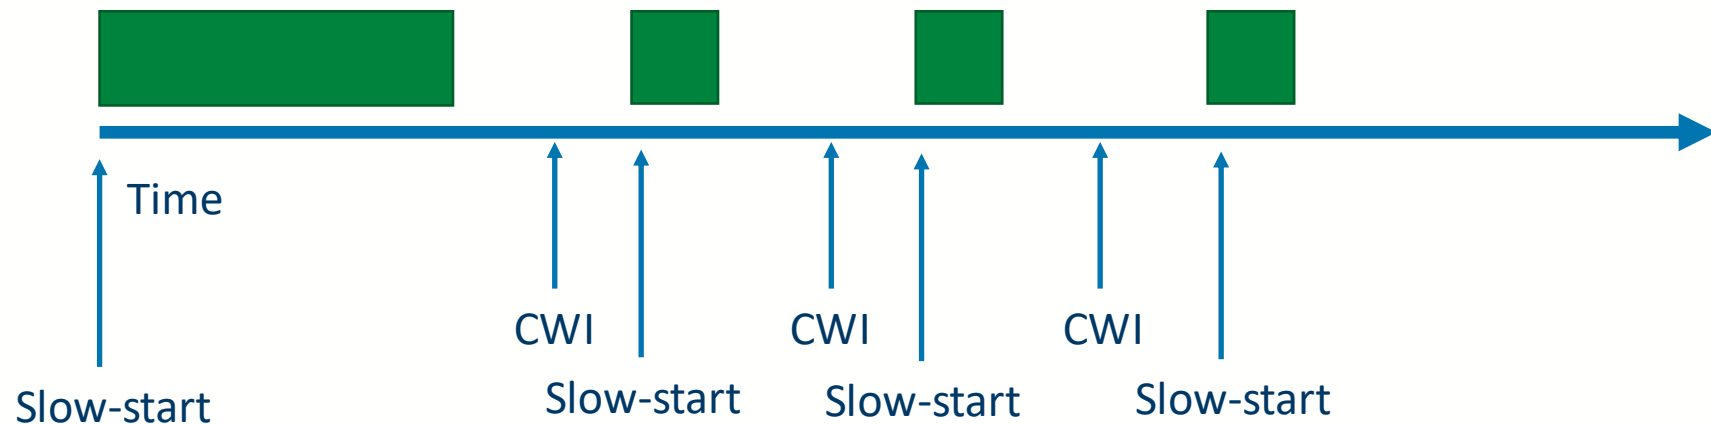Low Quality Stream

Client

Congestion Control

Slow-start

TCP Transmission Resumption

- At the start of the connection server delivers multiple video chunks
- After that it delivers chunks after periods where the link has been **idling** (off periods)
- During the off periods the **Congestion Window** is invalidated (CWI)
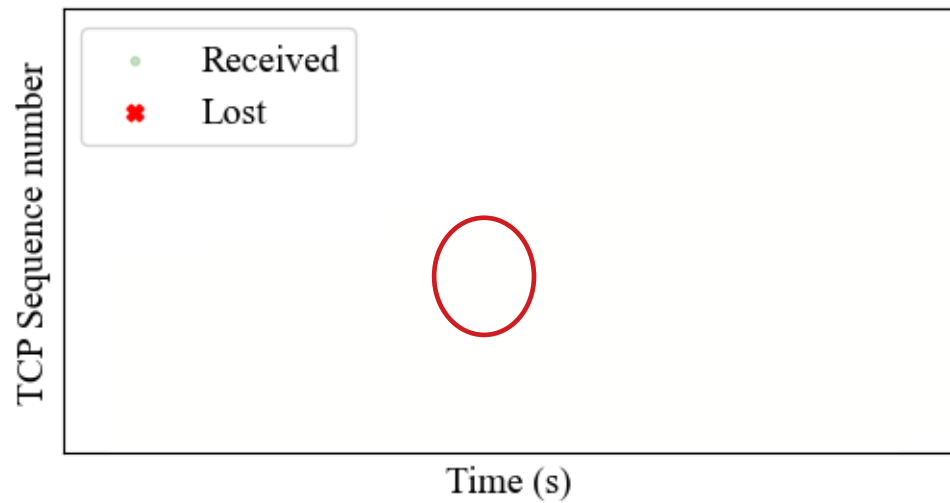- After that, a new **slow-start** is required on transmission resumption

Time

CWI   CWI   CWI

Slow-start   Slow-start   Slow-start   Slow-start

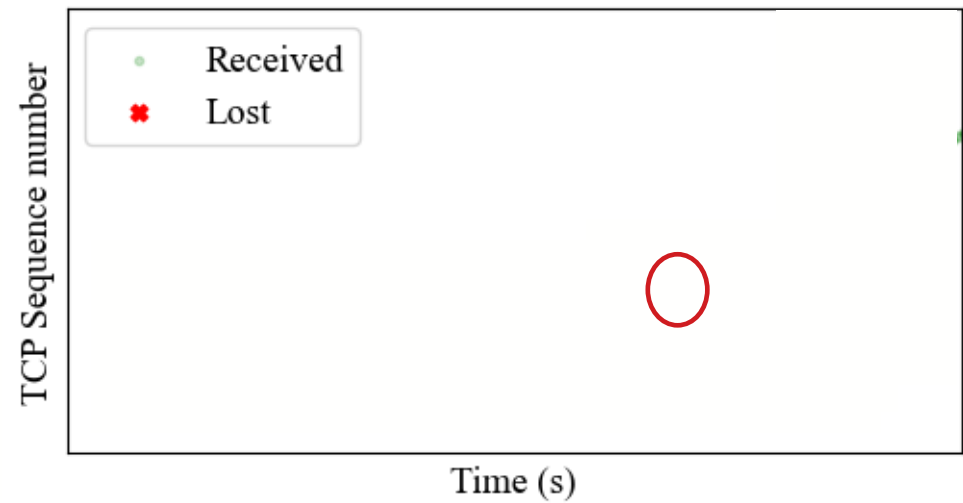# Congestion Window Invalidation
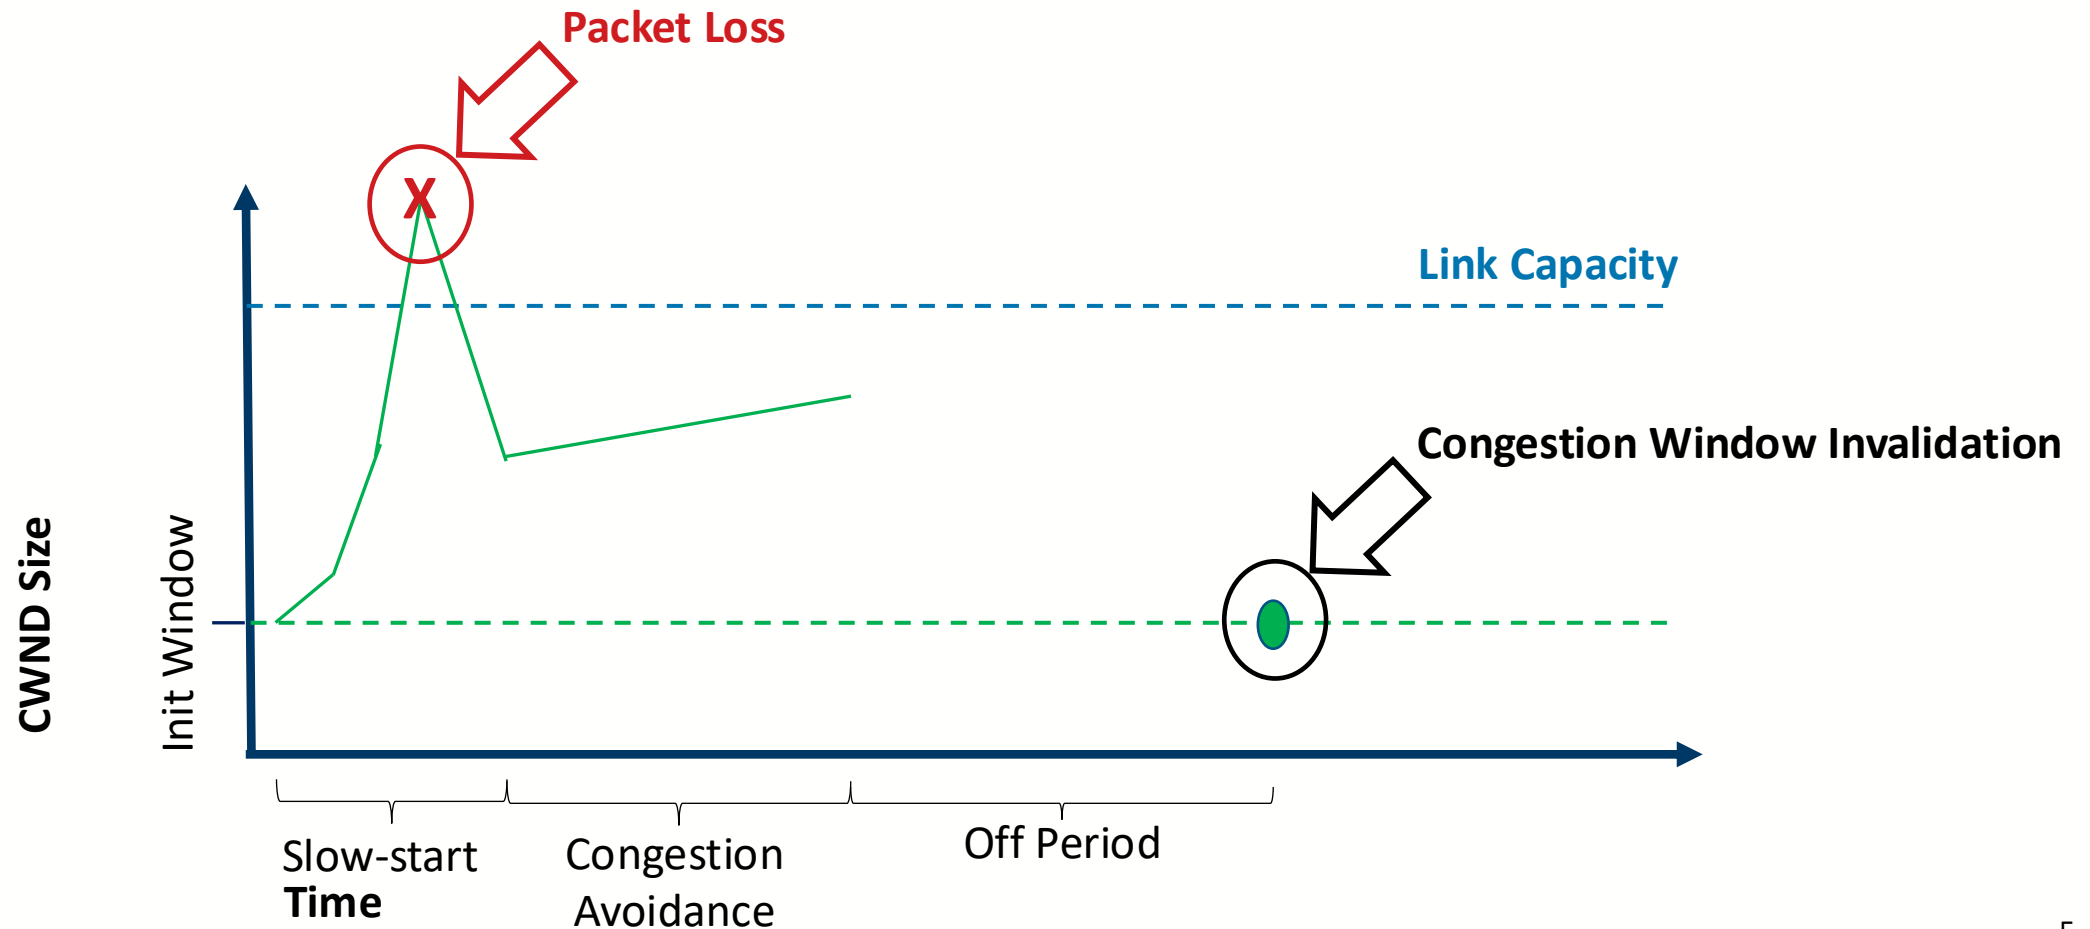
TCP Slow-start

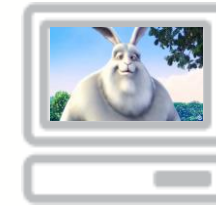Unmodified slow-start

Modified slow-start

slow-start

slow-start

HTTP Adaptive Streaming

Server

High quality stream
Medium Quality stream
Low Quality Stream

Good playback
- High-quality video
- Continuous steam
- No frequent switching

Client

Congestion Control

Slow-start

Rate Adaptation

HTTP Adaptive Streaming

Details for video used in experiments:
- Title - Big Buck Bunny [1]
- Duration - 10 Min 34 seconds
- Encoded in 4 qualities
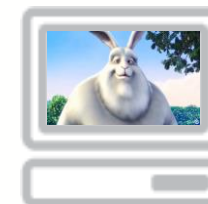- Split in 3 second chunks

Server

High quality stream
Medium Quality stream
Low Quality Stream
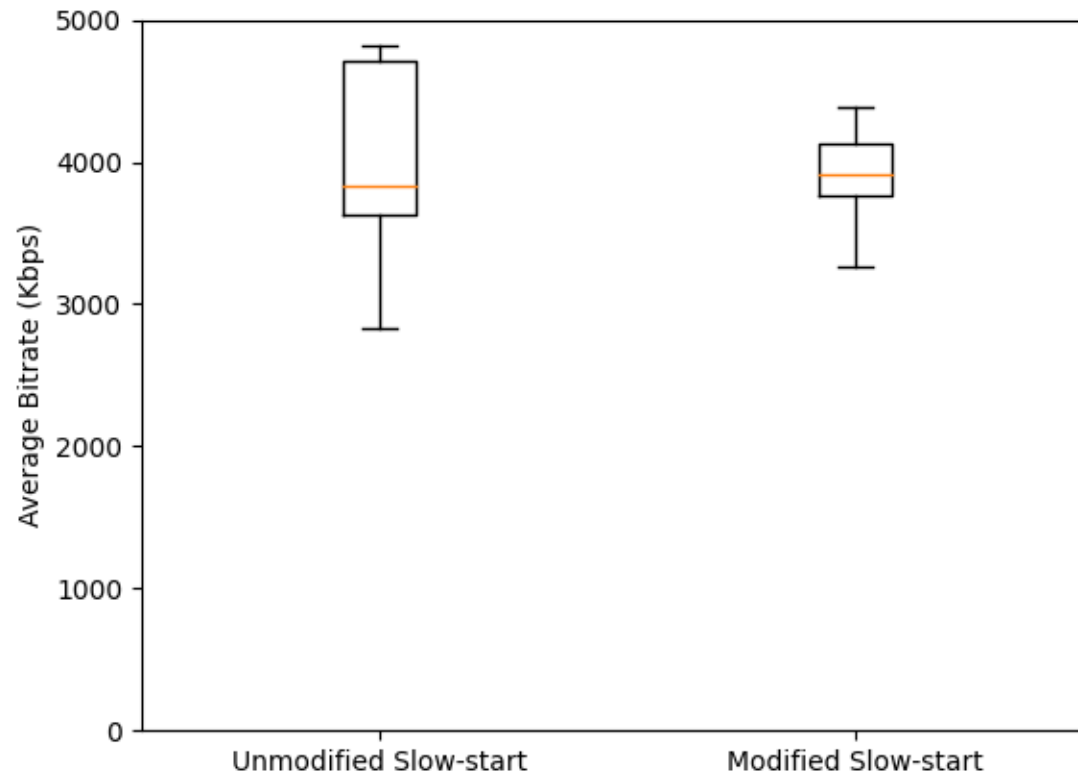
Client

Congestion Control

Slow-start

Rate Adaptation

Measuring:
- What was the overall quality of the stream
- How often video stream was pausing
- % playtime without representation switches

1. "Big Buck Bunny", (c) copyright 2008, Blender Foundation, www.bigbuckbunny.org
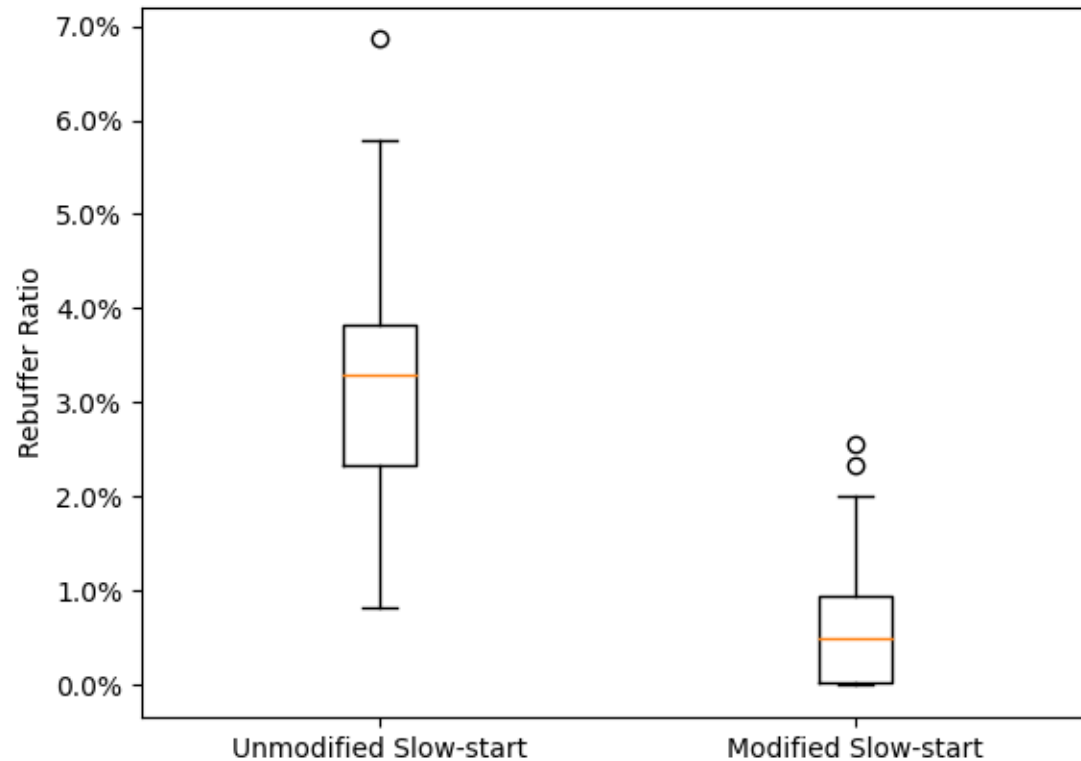
Average Bit-rate

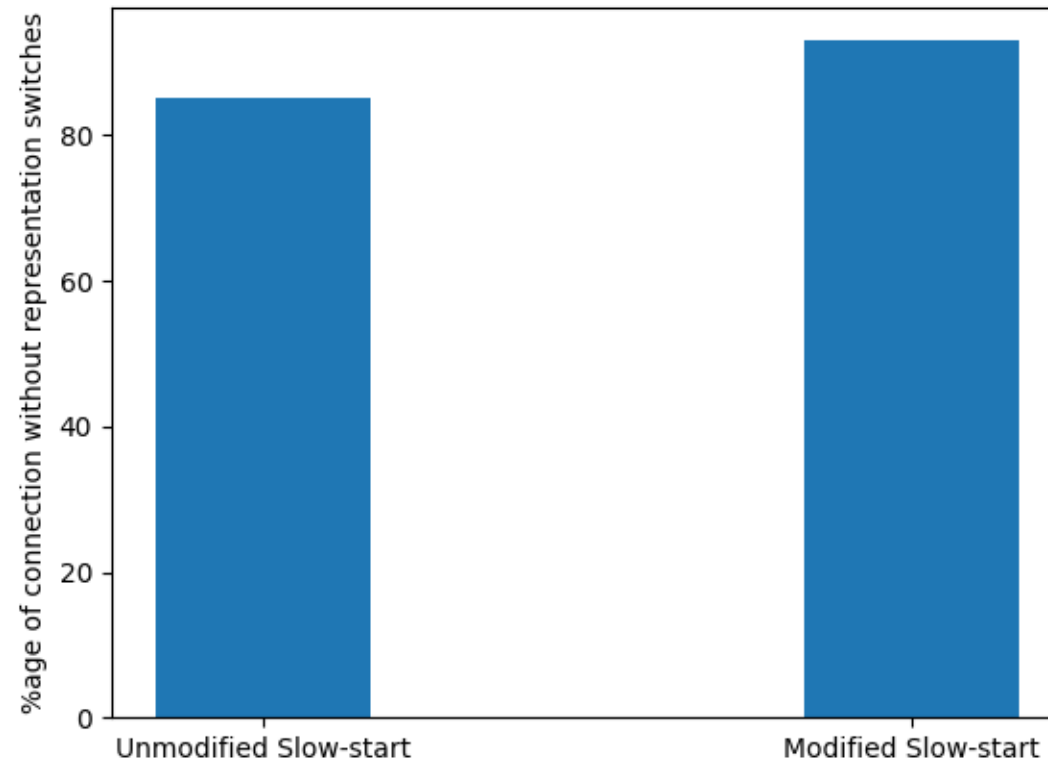The Average Bit-rate measures the quality of the video playback

Rebuffer Ratio

The Rebuffer Ratio measures what portion the video spent without progressing the playback (video freeze)

Encoding Switch Frequency

The Encoding Switch Frequency measures how often (and how big) switches between encoding levels occurred during playback

Conclusions

- Clients connected to video servers that use modified slow-start can use simpler adaptation algorithms to deliver comparable or better application-level metrics

- Reducing transport-level packet loss could bring better application-level performance

This talk presented extended work of our NOSSDAV '22 pblication[1]
DOI:10.1145/3534088.3534347

https://yanev.uk   m.yanev.1@research.gla.ac.uk

Rate Adaptation Algorithm Complexity