# LLL-CAdViSE: Live Low-Latency Cloud-based Adaptive Video Streaming Evaluation framework

Babak Taraghi (Christian Doppler Laboratory ATHENA, University of Klagenfurt, Austria)

May 7-10, 2023

# Introduction

- A sophisticated cloud-based and open-source testbed that facilitates evaluating a low-latency live streaming session.
- Live Low-Latency Cloud-based Adaptive Video Streaming Evaluation (LLL-CAdViSE) framework is enabled to assess the live streaming systems running on two major HTTP Adaptive Streaming (HAS) formats, Dynamic Adaptive Streaming over HTTP (MPEG-DASH) and HTTP Live Streaming (HLS).
- We use Chunked Transfer Encoding (CTE) to deliver Common Media Application Format (CMAF) chunks to the media players.
- Our testbed generates the test content (audiovisual streams). Therefore, no test sequence is required, and the encoding parameters (eg. encoder, bitrate, resolution, latency) are defined separately for each experiment.
- We have integrated the ITU-T P.1203 quality model inside our testbed.

# Ene-2-End Latency Evaluation

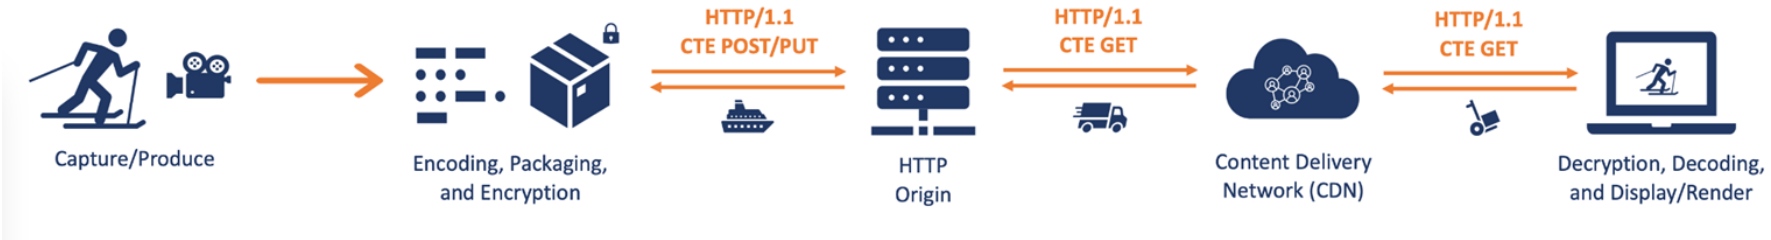


**FIGURE 1.** Low-latency live streaming by HTTP Chunked Transfer Encoding (Illustration inspired by a keynote at ACM MMSys'22 by Ali C. Begen - A master's toolbox and algorithms for low-latency live Streaming)

# Auto-generated Test Sequence


**FIGURE 2.** A single frame of the highly dynamic (randomly moving objects transforming into different shapes and with constantly changing color and size) video generated by the LLL-CAdViSE server.

# LLL-CAdViSE System Components

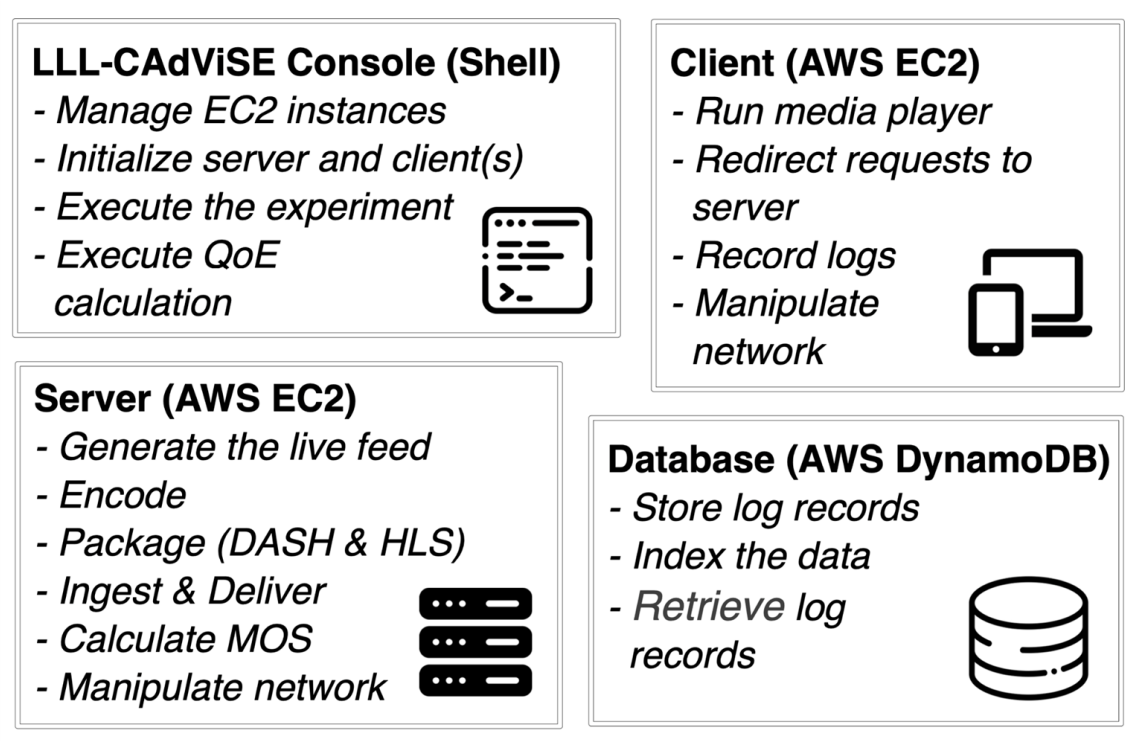


**FIGURE 3.** Live low-latency CAdViSE system components.

# Configurable Parameters

| Parameter | Value | Parameter | Value |
|---|---|---|---|
| Video Encoder | libx264 | GOP Size | 48 |
| Audio Encoder | aac | Video Preset | Faster |
| Tune | zerolatency | Video Profile | Main |
| Segment Duration | 2 seconds | fflags | genpts |
| Fragment Duration | 1 second | mov Flag | cmaf |
| Update Period$^{\pm}$ | 30 seconds | Pixel Format | yuv420p |
| Video Buffer | $Bitrate \times 2 \div 3$ | Write PRFT | True |
| Segment Type | mp4 | FPS | 24 |
| HTTP Method | PUT | Min. Rate$^{-}$ | 0.95 |
| HTTP Option | Chunked Post | Max. Rate$^{+}$ | 1.05 |

Origin server encoder and packager configuration parameters in the experimental setup.

# Real World Network Traces

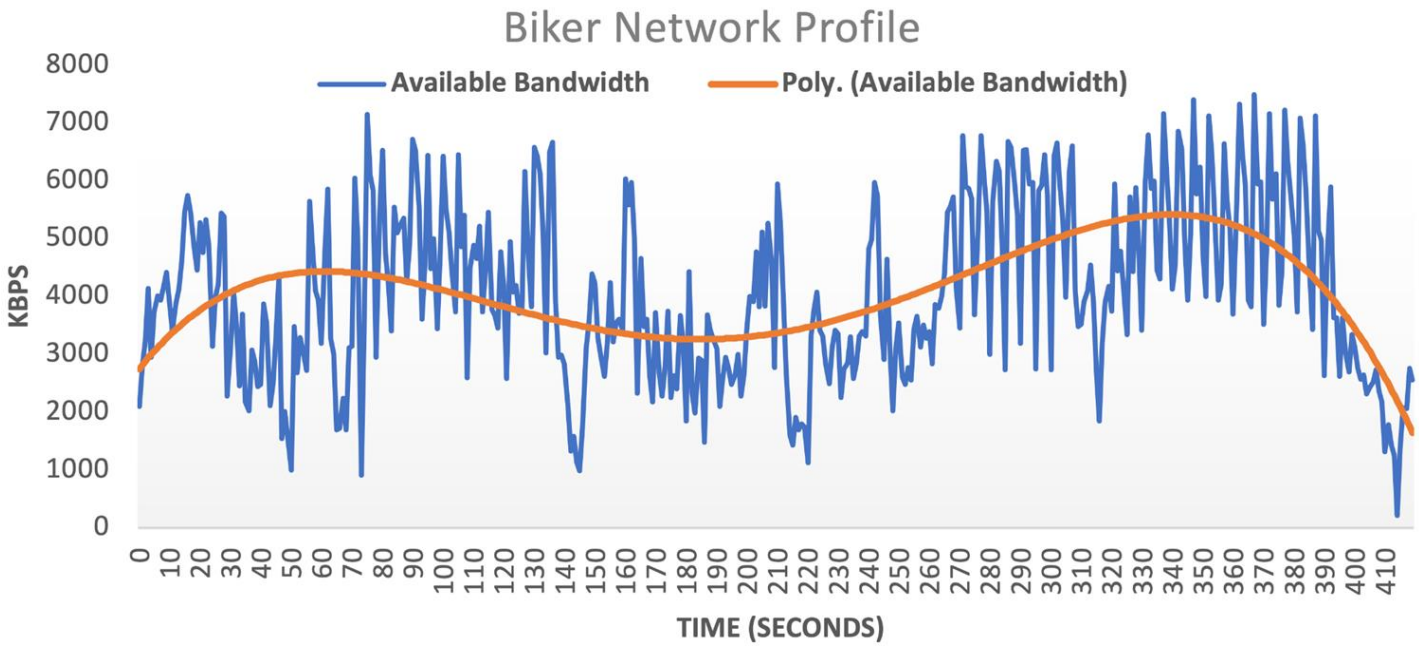


**FIGURE 4.** Bicycle commuter LTE network trace recorded in Belgium.

# MPEG-DASH Low-latency Players

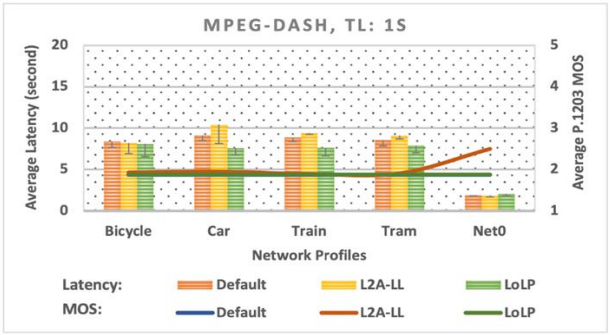


(a) Comparison with a target latency of 1 second.

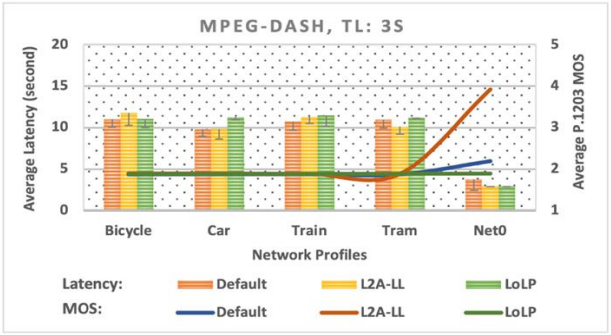


(b) Comparison with a target latency of 3 seconds.

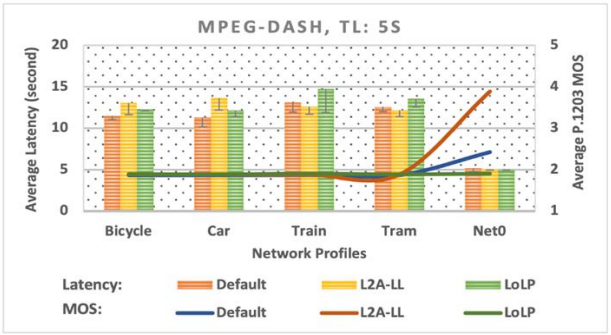


(c) Comparison with a target latency of 5 seconds.

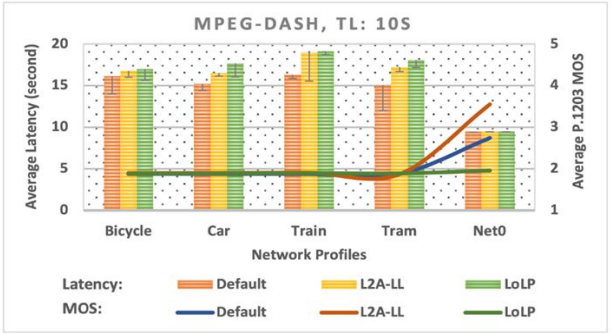


(d) Comparison with a target latency of 10 seconds.

**FIGURE 5.** Average latency and predicted MOS comparison of three ABR algorithms implemented on dash.js media player with four given target latencies and five network profiles (Note that average latency range is from 0 to 20 seconds).

# HLS Low-latency Players

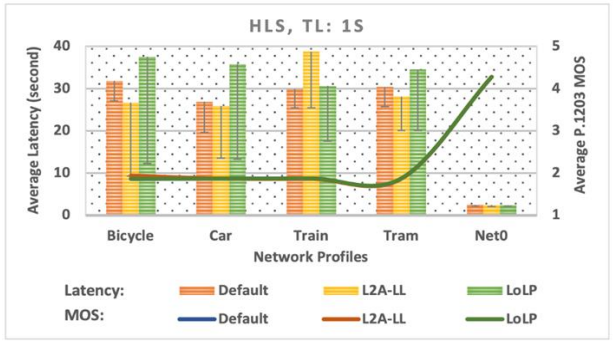


(a) Comparison with a target latency of 1 second.

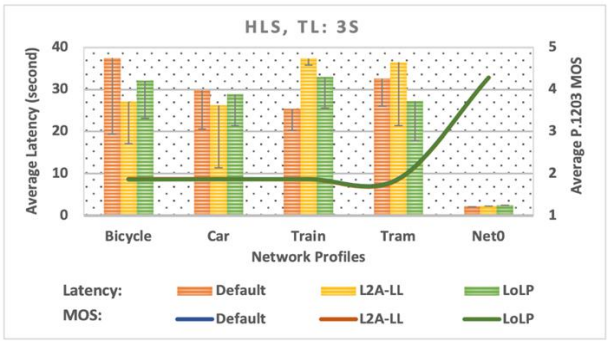


(b) Comparison with a target latency of 3 seconds.

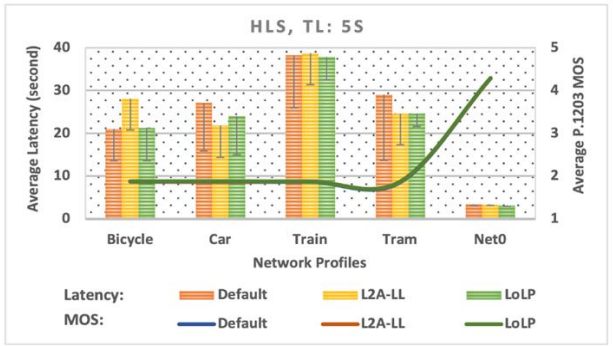


(c) Comparison with a target latency of 5 seconds.

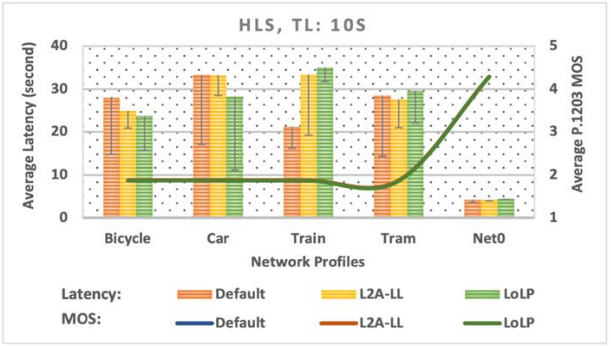


(d) Comparison with a target latency of 10 seconds.

**FIGURE 6.** Average latency and predicted MOS comparison of three ABR algorithms implemented on hls.js media player with four given target latencies and five network profiles (Note that average latency range is from 0 to 40 seconds).

# Raw results of low-latency live streaming with MPEG-DASH.

| Experiment[a] | Stall[b] | StartUp[c] | Seek[d] | Switch[e] | Bitrate[f] | Latency[g] | Playback Rate[h] | MOS[i] |
|---|---|---|---|---|---|---|---|---|
| ldash-def-bike-1s | 143.30 | 8.69 | 18.90 | 0 | 100k-100k-100k | 1.18-31.26-8.40 | 1-1.05-1.04 | 1.87 |
| ldash-def-car-1s | 147.42 | 11.03 | 12.63 | 0 | 100k-100k-100k | 1.76-36.78-9.11 | 1-1.05-1.04 | 1.87 |
| ldash-def-train-1s | 163.95 | 10 | 14.68 | 0 | 100k-100k-100k | 1.64-33.90-8.84 | 1-1.05-1.04 | 1.87 |
| ldash-def-tram-1s | 174.15 | 8.49 | 16.48 | 0 | 100k-100k-100k | 1.66-24.89-8.53 | 1-1.05-1.04 | 1.87 |
| ldash-def-net0-1s | 16.41 | 4.61 | 0.04 | 0 | 100k-100k-100k | 0.94-3.07-1.86 | 1-1.05-1.04 | 1.87 |
| ldash-l2a-bike-1s | 186.19 | 8.99 | 15.41 | 22.33 | 100k-750k-350k | 1.32-26.84-8.21 | 1-1.05-1.03 | 1.93 |
| ldash-l2a-car-1s | 206.51 | 12.71 | 27.97 | 23 | 100k-1000k-402k | 1.09-36.54-10.42 | 1-1.05-1.02 | 1.95 |
| ldash-l2a-train-1s | 213.87 | 13.29 | 26.5 | 22.66 | 100k-750k-298k | 1.61-29.19-9.41 | 1-1.05-1.03 | 1.9 |
| ldash-l2a-tram-1s | 200.88 | 11.10 | 36.22 | 19.66 | 100k-750k-336k | 1.81-25.26-9.09 | 1-1.05-1.0 | 1.91 |
| ldash-l2a-net0-1s | 15.69 | 4.71 | 0.06 | 59.33 | 316k-5800k-1339k | 0.94-3.05-**1.8** | 1-1.05-1.03 | **2.5** |
| ldash-lolp-bike-1s | 146.96 | 12.23 | 17.01 | 1 | 100k-750k-104k | 0.92-28.28-8.10 | 1-1.05-1.036 | 1.88 |
| ldash-lolp-car-1s | 125.55 | 9.31 | 10.22 | 1 | 100k-750k-104k | 1.49-24.12-7.60 | 1-1.05-1.04 | 1.88 |
| ldash-lolp-train-1s | 170.87 | 10.36 | 28.44 | 1 | 100k-750k-105k | 1.66-21.53-7.61 | 1-1.05-1.04 | 1.88 |
| ldash-lolp-tram-1s | 133.03 | 9.52 | 19.35 | 1 | 100-750-104k | 1.35-24.91-7.90 | 1-1.05-1.03 | 1.88 |
| ldash-lolp-net0-1s | 18.48 | 4.64 | 0.06 | 1 | 100k-750k-103k | 0.98-4.2-2.08 | 1-1.05-1.04 | 1.88 |
| ldash-def-bike-3s | 131.35 | 8.88 | 12.22 | 20.66 | 100k-500k-172k | 3.41-36.56-11.02 | 1-1.05-1.04 | 1.87 |
| ldash-def-car-3s | 120.17 | 7.03 | 11.09 | 21 | 100k-500k-146k | 3.36-30.19-9.8 | 1-1.05-1.04 | 1.87 |
| ldash-def-train-3s | 139.37 | 12.88 | 13.5 | 17 | 100k-625k-145k | 3.18-31.10-10.73 | 1-1.05-1.04 | 1.87 |
| ldash-def-tram-3s | 173.86 | 7.48 | 19.35 | 13.66 | 100k-500k-137k | 3.14-30.40-10.94 | 1-1.05-1.03 | 1.87 |
| ldash-def-net0-3s | 5.25 | 4.74 | 0.03 | 10.66 | 100k-5800k-722k | 2.93-5.86-3.77 | 0.99-1.03-1.01 | 2.19 |
| ldash-l2a-bike-3s | 176.47 | 9.77 | 17.12 | 28.33 | 100k-1000k-289k | 3.45-37.38-11.76 | 1-1.05-1.03 | 1.89 |
| ldash-l2a-car-3s | 124.18 | 10.71 | 14.50 | 29.33 | 100k-1000k-279k | 3.25-30.58-9.82 | 1-1.05-1.03 | 1.89 |
| ldash-l2a-train-3s | 158.35 | 10.75 | 17.96 | 30.33 | 100k-1000k-243k | 3.48-30.49-11.29 | 1-1.05-1.04 | 1.88 |
| ldash-l2a-tram-3s | 148.83 | 15.35 | 22.22 | 22.33 | 100k-1000k-226k | 3.28-29.68-10.11 | 1-1.05-1.036 | 1.88 |
| **ldash-l2a-net0-3s** | 0 | 4.75 | 0.07 | 206.33 | 750k-5800k-4384k | 2.84-3.24-**2.94** | 0.98-1.02-1 | **3.92** |
| ldash-lolp-bike-3s | 148.85 | 9.49 | 23.24 | 50 | 100k-1250k-238k | 3.33-32.10-11.04 | 1-1.05-1.03 | 1.88 |
| ldash-lolp-car-3s | 165.32 | 9.32 | 26.38 | 57 | 100k-1000k-202k | 3.40-37.34-11.25 | 1-1.05-1.03 | 1.88 |
| ldash-lolp-train-3s | 179.65 | 12.14 | 23.96 | 41 | 100k-1000k-209k | 4.08-29.80-11.44 | 1-1.05-1.03 | 1.88 |
| ldash-lolp-tram-3s | 167.60 | 7.57 | 25.49 | 50.33 | 100k-1000k-212k | 3.37-31.36-11.2 | 1-1.05-1.03 | 1.88 |
| ldash-lolp-net0-3s | 0 | 4.76 | 0.06 | 8 | 100k-5800k-256k | 2.86-3.24-2.95 | 0.98-1.02-1 | 1.89 |
| ldash-def-bike-5s | 112.07 | 9.67 | 7.84 | 26 | 100k-625k-157k | 5.36-33.83-11.48 | 1-1.05-1.04 | 1.87 |
| ldash-def-car-5s | 112.02 | 7.86 | 15.28 | 18 | 100k-625k-151k | 5.15-30.37-11.28 | 0.99-1.05-1.03 | 1.87 |
| ldash-def-train-5s | 142.17 | 29.53 | 15.56 | 16.66 | 100k-500k-143k | 5.48-32.84-13.12 | 1-1.05-1.03 | 1.87 |
| ldash-def-tram-5s | 143.76 | 7.36 | 16.73 | 15.33 | 100k-500k-132k | 5.34-34.2-12.59 | 1-1.05-1.03 | 1.87 |
| ldash-def-net0-5s | 1.82 | 4.62 | 0.01 | 36.33 | 100k-5800k-1077k | 4.95-6.44-5.17 | 0.99-1.03-1.00 | 2.42 |

[a] All time values are in seconds.
[a] Experiment title, format: "[streaming protocol]-[ABR algorithm]-[network profile]-[target latency]" (def: Default, l2a: L2A-LL). [b] Average of the sum of stall events duration.
[c] Average start-up delay.
[d] Average of the sum of seek events duration.
[e] Average quantity of quality switches.
[f] Playback bitrate (min-max-avg) in kbps.
[g] Latency (min-max-avg).
[h] Playback rate (min-max-avg).
[i] Average MOS predicted by the ITU-T P.1203 quality model.

Each row represents average values for three experiments.

# Raw results of low-latency live streaming with HLS.

| Experiment[a] | Stall[b] | StartUp[c] | Seek[d] | Switch[e] | Bitrate[f] | Latency[g] | Playback Rate[h] | MOS[i] |
|---|---|---|---|---|---|---|---|---|
| lhls-def-bike-1s | 85.32 | 1.20 | 38.11 | 0.66 | 100k-191k-102k | 2.40-54.45-31.81 | 1-1-1 | 1.86 |
| lhls-def-car-1s | 91.73 | 0.41 | 23.78 | 2.66 | 100k-500k-104k | 2.25-46.26-26.91 | 1-1-1 | 1.87 |
| lhls-def-train-1s | 90.42 | 0.67 | 25.85 | 2.66 | 100k-533k-110k | 2.97-46.34-29.94 | 1-1-1 | 1.87 |
| lhls-def-tram-1s | 84.44 | 1.35 | 20.83 | 2 | 100k-375k-115k | 2.75-52.96-30.43 | 1-1-1 | 1.87 |
| lhls-def-net0-1s | 58.99 | 0.05 | 65.45 | 3 | 100k-5800k-5723k | 1.85-3.55-2.48 | 1-1-1 | 4.28 |
| lhls-l2a-bike-1s | 65.75 | 0.39 | 12.2 | 2 | 100k-500k-206k | 3.33-39.62-26.72 | 1-1-1 | 1.94 |
| lhls-l2a-car-1s | 106.61 | 0.6 | 41.64 | 3 | 100k-500k-106k | 2.4-46.39-25.92 | 1-1-1 | 1.87 |
| lhls-l2a-train-1s | 84.84 | 0.45 | 18.40 | 0 | 100k-100k-100k | 2.63-59.32-38.81 | 1-1-1 | 1.86 |
| lhls-l2a-tram-1s | 108.22 | 0.48 | 34.98 | 1.33 | 100k-283k-107k | 2.31-46.53-28.12 | 1-1-1 | 1.87 |
| lhls-l2a-net0-1s | 56.95 | 0.05 | 66.06 | 3 | 100k-5800k-5723k | 1.42-3.52-2.44 | 1-1-1 | 4.28 |
| lhls-lolp-bike-1s | 128.19 | 0.43 | 47.04 | 2.66 | 100k-500k-116k | 2.31-66.67-37.57 | 1-1-1 | 1.86 |
| lhls-lolp-car-1s | 84.86 | 0.39 | 14.23 | 1.66 | 100k-408k-106k | 2.65-52.47-35.82 | 1-1-1 | 1.86 |
| lhls-lolp-train-1s | 102.18 | 0.24 | 23.60 | 2.66 | 100k-500k-105k | 3.13-59.59-30.66 | 1-1-1 | 1.87 |
| lhls-lolp-tram-1s | 79.9 | 0.31 | 19.79 | 2.33 | 100k-533k-106k | 2.09-51.66-34.52 | 1-1-1 | 1.87 |
| lhls-lolp-net0-1s | 52.08 | 0.05 | 53.09 | 3 | 100k-5800k-5724k | 1.75-3.58-**2.4** | 1-1-1 | **4.28** |
| lhls-def-bike-3s | 96.31 | 1.06 | 23.56 | 0.66 | 100k-191k-101k | 2.39-57.42-37.47 | 1-1-1 | 1.87 |
| lhls-def-car-3s | 82.89 | 0.36 | 25.18 | 1.33 | 100k-283k-102k | 2.62-43.95-29.82 | 1-1-1 | 1.87 |
| lhls-def-train-3s | 81.53 | 0.83 | 15.62 | 1.33 | 100k-283k-106k | 2.45-39.75-25.42 | 1-1-1 | 1.87 |
| lhls-def-tram-3s | 122.95 | 0.71 | 29.47 | 1.66 | 100k-408k-104k | 2.46-64.59-32.6 | 1-1-1 | 1.88 |
| lhls-def-net0-3s | 50.69 | 0.05 | 7.64 | 3 | 100k-5800k-5724k | 1.9-4.32-2.30 | 1-1-1 | 4.28 |
| lhls-l2a-bike-3s | 58.45 | 0.43 | 9.86 | 2.66 | 100k-500k-108k | 2.37-37.16-27.12 | 1-1-1 | 1.87 |
| lhls-l2a-car-3s | 92.31 | 0.57 | 34.96 | 2.66 | 100k-500k-113k | 2.32-39.95-26.26 | 1-1-1 | 1.87 |
| lhls-l2a-train-3s | 69.78 | 0.75 | 5.41 | 0.66 | 100k-191k-100k | 2.66-54.47-37.36 | 1-1-1 | 1.86 |
| lhls-l2a-tram-3s | 97.91 | 0.76 | 18.54 | 2 | 100k-375k-107k | 4.85-56.30-36.5 | 1-1-1 | 1.86 |
| lhls-l2a-net0-3s | 70.48 | 0.05 | 61.28 | 3 | 100k-5800k-5720k | 1.49-4.56-2.4 | 1-1-1 | 4.28 |
| lhls-lolp-bike-3s | 99.2 | 0.37 | 25.21 | 2 | 100k-408k-113k | 3.02-52.25-32.08 | 1-1-1 | 1.86 |
| lhls-lolp-car-3s | 85.25 | 1.22 | 15.93 | 2 | 100k-375k-106k | 2.38-48.31-28.9 | 1-1-1 | 1.86 |
| lhls-lolp-train-3s | 95.35 | 1.13 | 12.59 | 1.33 | 100k-283k-111k | 2.71-52.26-32.97 | 1-1-1 | 1.87 |
| lhls-lolp-tram-3s | 97.98 | 0.47 | 32.94 | 2.33 | 100k-408k-115k | 2.44-46.95-27.25 | 1-1-1 | 1.87 |
| **lhls-lolp-net0-3s** | 70.61 | 0.05 | 57.71 | 3 | 100k-5800k-5720k | 1.47-4.69-**2.54** | 1-1-1 | **4.28** |
| lhls-def-bike-5s | 107.67 | 2.59 | 13.46 | 0.66 | 100k-191k-116k | 1.91-40.31-21.02 | 1-1-1 | 1.87 |
| lhls-def-car-5s | 106.91 | 0.23 | 13.85 | 1.33 | 100k-283k-101k | 2.19-52.86-27.09 | 1-1-1 | 1.86 |
| lhls-def-train-5s | 81.69 | 0.73 | 10.10 | 0.66 | 100k-191k-104k | 2.89-55.35-38.22 | 1-1-1 | 1.86 |
| lhls-def-tram-5s | 60.59 | 0.87 | 2.89 | 0 | 100k-100k-100k | 2.47-42.50-29.01 | 1-1-1 | 1.86 |
| lhls-def-net0-5s | 25.21 | 0.053 | 6.01 | 3 | 100k-5800k-5729k | 1.65-5.18-**3.41** | 1-1-1 | 4.28 |

[a] All time values are in seconds.
[a] Experiment title, format: "[streaming protocol]-[ABR algorithm]-[network profile]-[target latency]" (def: Default, l2a: L2A-LL). [b] Average of the sum of stall events duration.
[c] Average start-up delay.
[d] Average of the sum of seek events duration.
[e] Average quantity of quality switches.
[f] Playback bitrate (min-max-avg) in kbps.
[g] Latency (min-max-avg).
[h] Playback rate (min-max-avg).
[i] Average MOS predicted by the ITU-T P.1203 quality model.

Each row represents average values for three experiments.

# Conclusions

- A sophisticated cloud-based and open-source testbed, LLL-CAdViSE is a framework for evaluating HAS live streaming (with MPEG-DASH and HLS) and using CMAF and CTE.
- Evaluations of live media streaming significant metrics such as:
  - Stall events, quality (representation) switches, and played bitrate
  - Precise measurement of media streaming E2E latency (plus seek duration and playback rate).
  - Automatic assessment of objective QoE using ITU-T P.1203 quality model.
  - Preparation of a single media file (.mp4) for further investigation of the defects.
- The results of extensive tests of well-known media players and ABR algorithms using LLL-CAdViSE shows that the L2A-LL ABR algorithm plugged into the dash.js media player and using MPEG-DASH low-latency live streaming outperforms other setups in providing the closest latency to a target latency and maintaining a high QoE score.
- Our testbed is publicly available on GitHub:

**https://github.com/cd-athena/LLL-CAdViSE**

Citation: B. Taraghi, H. Hellwagner and C. Timmerer, "LLL-CAdViSE: Live Low-Latency Cloud-Based Adaptive Video Streaming Evaluation Framework," in IEEE Access, vol. 11, pp. 25723-25734, 2023, doi: 10.1109/ACCESS.2023.3257099.

# Thank you