

PERFORMANCE OF LOW-LATENCY HTTP-BASED STREAMING PLAYERS

Bo Zhang, Nabajeet Barman, Yuriy Reznik,
Brightcove



mhv/2023
ACM MILE-HIGH VIDEO
annual workshop

May 7-10, 2023, Denver, CO

Outline

- ❑ Context & Objectives
- ❑ Evaluation Framework
- ❑ Test Results
- ❑ Conclusions

Context & Objectives

Current state of HTTP-based streaming:

- HLS and DASH work, but introduce significant (~30 secs) delays
- Delays are caused by long segments, transmission protocols and player side buffering
- LL-DASH and LL-HLS are 2 major new technologies

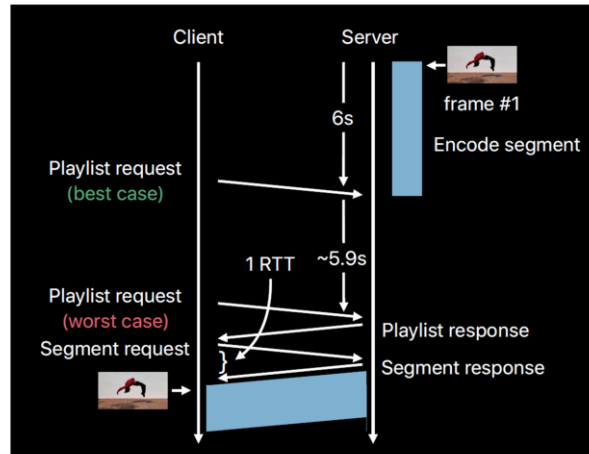
Key design principles of LL-systems:

- Server: encode and push video in smaller chunks (e.g. 1 sec)
- Client: request a segment as soon as first chunk becomes available
- Desired end-to-end delay: < 5 secs

Objectives of this work:

- Evaluate the existing implementations of LL-HLS/DASH systems
- Understand tradeoffs between the delay and other QOE factors
- Understand how mature implementations of LL-HLS/DASH systems are currently

Delays in traditional HLS streaming



Delays in LL-HLS streaming

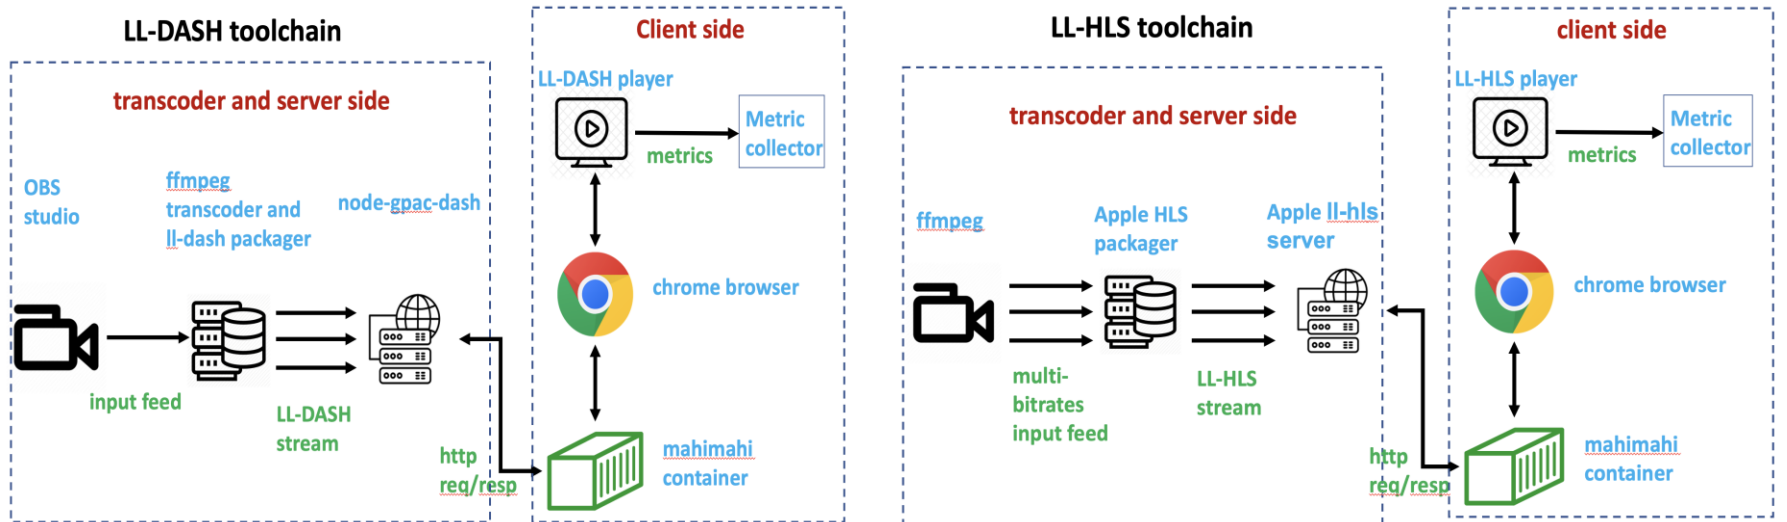


Source: Apple, Inc. WWDC19

Evaluation Framework

Streaming systems

LL-DASH	LL-HLS
OBS studio, FFmpeg and node-gpac-dash	FFmpeg and Apple's HLS tools
DASH.js (+ LoL and L2All), Shaka players, Theo player	HLS.js, and Shaka players



Experiment setup

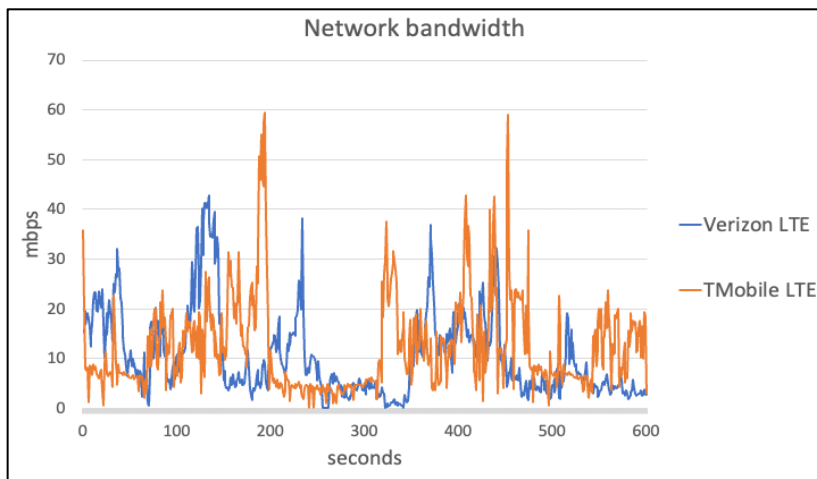
- Test stream: Big buck bunny looped for each 10 minutes streaming session.
- 4 live encoding profiles

Parameter	Low	Mid	High	Top
Codec	H.264	H.264	H.264	H.264
Video resolution	768 x 432	1024 x 576	1600 x 900	1920 x 1080
Video framerate (fps)	30	30	30	30
Bitrate (kbps)	949	1854	3624	5166

- Notes:
 - Media segment duration = 4 seconds, media chunk duration = 1 second
 - All encodings are CBR
 - Range of rendition bitrates well covers dynamic range of bandwidth in both Verizon & T-Mobile networks

Network emulation

- Mahimahi network emulator is used to emulate real-world network condition.
- Two network traces are used for analysis: T-Mobile 4G LTE network and Verizon 4G LTE network



Bandwidth statistics

Network	T-Mobile	Verizon
Average (kbps)	12258	10565
Variation (kbps)	9314	8619
Min. (kbps)	12	12
Max. (kbps)	59460	42804

- Reference

R. Netravali, A. Sivaraman, S. Das, A. Goyal, K. Winstein, J. Mickens, H. Balakrishnan,

“Mahimahi: accurate record-and-replay for HTTP,” USENIX Annual Technical Conference, Santa Clara, CA, July 8-10, 2015

Test results

Results for T-Mobile LTE network

	LL-DASH	LL-HLS
Stream bitrate over time	<p>Bitrate - Dash TMobile</p> <p>This chart displays the stream bitrate in kbps over a 600-second period for the LL-DASH configuration on the T-Mobile network. The y-axis ranges from 0 to 6000 kbps. The x-axis represents time in seconds from 0 to 600. The legend includes: dash.js default (red dashed), lolp (yellow solid), IZall (blue solid), shaka (green solid), and theoplayer (black dashed). The plot shows highly variable bitrates, with many spikes reaching between 4000 and 6000 kbps, indicating significant fluctuations in stream quality.</p>	<p>Bitrate - Hls TMobile</p> <p>This chart displays the stream bitrate in kbps over a 600-second period for the LL-HLS configuration on the T-Mobile network. The y-axis ranges from 0 to 6000 kbps. The x-axis represents time in seconds from 0 to 600. The legend includes: hls.js 2020 default (red dashed), hls.js 2020 + lolp (yellow solid), hls.js 2020 + IZall (blue solid), shaka (green solid), and hls.js 2023 (black dashed). The plot shows a more stable bitrate compared to DASH, with most of the time spent at 1000-2000 kbps and occasional spikes up to 5000 kbps.</p>
Latency over time	<p>Latency - Dash TMobile</p> <p>This chart displays the latency in seconds over a 600-second period for the LL-DASH configuration on the T-Mobile network. The y-axis ranges from 0 to 30 seconds. The x-axis represents time in seconds from 0 to 600. The legend includes: dash.js default (red dashed), lolp (yellow solid), IZall (blue solid), shaka (green solid), and theoplayer (black dashed). A horizontal dashed line is drawn at approximately 5.5 seconds. The plot shows several sharp spikes in latency, with the highest reaching nearly 30 seconds, indicating significant buffering events.</p>	<p>Latency - Hls TMobile</p> <p>This chart displays the latency in seconds over a 600-second period for the LL-HLS configuration on the T-Mobile network. The y-axis ranges from 0 to 80 seconds. The x-axis represents time in seconds from 0 to 600. The legend includes: hls.js 2020 default (red dashed), hls.js 2020 + lolp (yellow solid), hls.js 2020 + IZall (blue solid), shaka (green solid), and hls.js latest (black dashed). A horizontal dashed line is drawn at approximately 10 seconds. The plot shows a prominent spike in latency for the hls.js 2020 default player, reaching about 68 seconds, while other players maintain latency below 10 seconds.</p>

Results for Verizon LTE network

	LL-DASH	LL-HLS
Stream bitrate over time	<p>Bitrate - Dash Verizon</p> <p>This chart displays the stream bitrate in kops over a 600-second period for the LL-DASH configuration on Verizon LTE. The y-axis ranges from 0 to 6000 kops. The x-axis ranges from 0 to 600 seconds. The legend includes: dash.js default (red dashed), lolp (yellow dashed), l2all (blue dashed), shaka (green dashed), and theoplayer (black dashed). The plot shows frequent, sharp fluctuations in bitrate, with many lines reaching between 1000 and 5000 kops.</p>	<p>Bitrate - Hls Verizon</p> <p>This chart displays the stream bitrate in kops over a 600-second period for the LL-HLS configuration on Verizon LTE. The y-axis ranges from 0 to 6000 kops. The x-axis ranges from 0 to 600 seconds. The legend includes: hls.js 2020 default (red dashed), hls.js 2020 + lolp (yellow dashed), hls.js 2020 + l2all (blue dashed), shaka (green dashed), and hls.js 2023 (black dashed). The plot shows significant bitrate fluctuations, with several lines spiking to 5000 kops.</p>
Latency over time	<p>Latency - Dash Verizon</p> <p>This chart displays the latency in seconds over a 600-second period for the LL-DASH configuration on Verizon LTE. The y-axis ranges from 0 to 50 seconds. The x-axis ranges from 0 to 600 seconds. The legend is the same as the bitrate chart. Most players maintain low latency (under 10 seconds), but there are several sharp spikes, with theoplayer reaching approximately 45 seconds at the 350-second mark.</p>	<p>Latency - Hls Verizon</p> <p>This chart displays the latency in seconds over a 600-second period for the LL-HLS configuration on Verizon LTE. The y-axis ranges from 0 to 70 seconds. The x-axis ranges from 0 to 600 seconds. The legend is the same as the bitrate chart. While most players have low latency, hls.js 2020 default shows a prominent spike to about 65 seconds at the 100-second mark. Other players like shaka and hls.js 2023 show spikes to 25-30 seconds around the 350-second mark.</p>

Table 5: Performance statistics – T-Mobile LTE network

Player/Algorithm	Avg. bitrate [kbps]	Avg. height [pixels]	Avg. latency [secs]	Latency var. [secs]	Speed var. [%]	Number of switches	Buffer events	Buffer ratio [%]	MBs loaded	Objects loaded
DASH.js default	2770	726	3.06	0.21	10.4	93	38	7.99	352.2	256
DASH.js <u>LolP</u>	3496	853	5.65	4.59	22.7	70	53	21.96	369.4	210
DASH.js L2all	3699	908	4.14	3.18	19.9	5	19	7.99	368	147
Shaka player (dash)	3818	916	4.92	2.06	0	16	5	4.66	360.3	155
THEO player (dash)	4594	993	6.16	0.01	0	27	0	0	418.7	152
HLS.js default 2020	1763	562	10.08	10.91	8.1	26	2	9.8	130.7	589
HLS.js <u>LolP</u> 2020	1756	560	5.97	0.2	6.1	24	0	0	148.1	688
HLS.js L2all 2020	1752	560	6	0.23	5.9	34	0	0	133.1	686
HLS.js default 2023	3971	895	8.93	1.13	0	8	0	0	360.8	613
Shaka player (hls)	3955	908	7.18	2.23	0	14	7	3.8	230	475

Table 6: Performance statistics – Verizon LTE network

Player/Algorithm	Avg. bitrate [kbps]	Avg. height [pixels]	Avg. latency [secs]	Latency var. [secs]	Speed var. [%]	Number of switches	Buffer events	Buffer ratio [%]	MBs loaded	Objects loaded
DASH.js default	2131	627	3.79	3.16	14.9	91	23	7.99	260	207
DASH.js <u>LolP</u>	3368	829	7.29	6.8	23.9	106	51	21.96	351	221
DASH.js L2all	3672	905	6.47	5.36	23.3	7	7	7.99	338	135
Shaka player (dash)	3653	886	6.96	7.08	0	26	5	4.66	329	148
THEO player (dash)	4153	909	18.19	10.08	0	33	2	0	383	153
HLS.js default 2020	2085	610	11.66	10.25	11.4	28	3	9.8	140	606
HLS.js <u>LolP</u> 2020	1890	580	7.86	2.63	7.0	24	2	0	146	569
HLS.js L2all 2020	1803	567	8.84	4.43	10.7	26	2	0	145	547
HLS.js default 2023	3541	822	16.78	9.13	0	9	5	0	280	598
Shaka player (hls)	3669	860	7.98	2.41	0	22	9	3.8	260	570

Observations

- Some players failed to stream with low-latency when network bandwidth is low and dynamic.
- Players frequently re-buffer in slower and dynamic network. Ensuring continuous and smooth playback is critical for general QoE.
- Some players double-download from two renditions in hope of avoiding buffer underrun. Our results show this strategy does not work well in practice but could lead to unnecessary complexity and bandwidth waste.
- Some players switch bitrate too often leading to reduced QoE.
- Some players change playback speed too often, while others stick to the natural speed.
- All the players shows better performance than our old study 2 years ago. HLS.js shows the greatest improvement.

Conclusions

Conclusions

Both LL-HLS and LL-DASH reduce delays

- 3-7 sec end-to-end streaming latency is achievable today

Observed drawbacks

- *Reliability* – number of buffering events
- *Scalability and delivery costs* – increased intensity of data exchanges with origin servers and CDNs (particularly with LL-HLS)
- *Consistency of experience* – increased number of switches, delay variability, etc.

Some aspects can be improved

- Frequency of switches and buffering can possibly be minimized with smarter algorithms in streaming clients. But clearly, there is still considerable work ahead!

**THANK
YOU**