Compression Efficiency and Complexity Evaluation of MPEG-5 LCEVC-enhanced VVenC for UHD Video Delivery

Manzoor Razaak
V-Nova Ltd.
Email: manzoor.razaak@v-nova.com

Hannia Tahir
V-Nova Ltd.
Email: hannia.tahir@v-nova.com

Guendalina Cobianchi
V-nova Ltd.
Email: guendalina.cobianchi@v-nova.com

Abstract—**The deployment of Ultra High Definition (UHD) video using Versatile Video Coding (VVC/H.266) provides substantial compression efficiency but remains constrained by high computational complexity, particularly for large-scale encoding workflows. This paper presents a comprehensive evaluation of MPEG-5 Part 2 Low Complexity Enhancement Video Coding (LCEVC) applied to VVC as a multi-layer solution for deployable UHD delivery. By encoding the VVC base layer at reduced resolution and reconstructing full-resolution detail through lightweight enhancement layers, the proposed approach redistributes computational workload while preserving or enhancing perceptual quality. Experimental evaluation on a diverse UHD dataset demonstrates consistent rate–distortion improvements compared to native VVenC encoding. Subjective quality assessment shows that LCEVC achieves an average BD-MOS bitrate reduction of approximately 23% at equivalent perceived quality. On average, LCEVC-enhanced VVenC achieves approximately 30% VMAF BD-rate gains across all tested sequences. In addition to compression improvements, the integration significantly reduces encoding complexity, achieving an average $2.6\times$ increase in encoding throughput relative to native VVenC. These results indicate that LCEVC-enhanced VVC provides an effective balance between compression efficiency, perceptual quality, and computational scalability, enabling practical deployment of next-generation UHD video delivery workflows.**

I. Introduction

The deployment of Ultra High Definition (UHD) video delivery across streaming and broadcast platforms continues to increase pressure on network bandwidth, computational resources, and device capabilities. While consumer expectations for visual fidelity continue to rise, operators must balance compression efficiency, encoding throughput, decoding complexity, and large-scale deployability. Addressing these competing requirements remains a key challenge for modern video delivery workflows.

Versatile Video Coding (VVC, ITU-T H.266) represents the latest generation of video compression standards and achieves state-of-the-art efficiency through advanced coding tools such as improved partitioning structures, refined prediction methods, and enhanced transform coding [1]. Relative to HEVC, VVC can deliver approximately 40–50% bitrate reduction at equivalent subjective quality [2]. However, these gains come at the cost of substantially higher computational complexity [3], [4], particularly for UHD encoding and decoding, which limits practical deployment at scale.

MPEG-5 Part 2 Low Complexity Enhancement Video Coding (LCEVC) introduces a layered enhancement architecture in which a conventional codec operates at reduced spatial resolution while lightweight enhancement layers reconstruct high-frequency detail at full output resolution [5], [6]. By applying the most computationally demanding coding tools at lower resolution, LCEVC redistributes processing complexity while preserving or improving perceptual quality.

This layered design provides two key advantages for UHD delivery workflows. First, complexity redistribution allows the base encoder to operate on a lower-resolution representation, significantly reducing computational load. Second, the enhancement layers encode residual information directly at display resolution, enabling efficient reconstruction of perceptually important spatial detail.

Recent deployments, such as the adoption of LCEVC in Brazil's TV 3.0 broadcast system [7], demonstrate the practical viability of this approach for large-scale delivery environments. While prior work has evaluated LCEVC in combination with codecs such as AVC, HEVC, and AV1 [8]–[10], a systematic evaluation of LCEVC applied to VVC for UHD delivery, covering compression efficiency, perceptual quality, and computational complexity, remains limited.

This paper presents a comprehensive evaluation of LCEVC-VVenC encoding for UHD content. The main contributions are threefold:

1) A quantitative analysis of compression efficiency of LCEVC-VVenC against VVenC encoding.
2) A subjective quality assessment validating perceptual quality and examining the relationship between Mean Opinion Scores (MOS) and objective quality metrics.
3) An encoding complexity analysis demonstrating the practical deployability advantages of the layered approach.

By combining objective, subjective, and computational analysis, this study provides a holistic assessment of LCEVC-enhanced VVC for deployable UHD video delivery.

II. Background

A. LCEVC Structure

LCEVC operates as a layered enhancement framework that augments a conventional video codec, such as VVC, used as a base layer. Instead of encoding the full-resolution signal directly with the base codec, the input video is first encoded at a reduced resolution. The decoded base reconstruction is then upscaled and refined using one or more lightweight enhancement layers that encode residual information. The overall encoder structure is illustrated in Fig. 1.

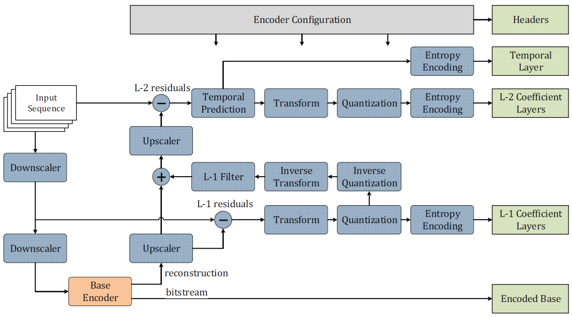

Fig. 1: LCEVC layered encoder architecture showing the reduced-resolution base layer and enhancement residual layers.

The enhancement layers capture the difference between the upscaled base reconstruction and the original signal, progressively restoring fine spatial detail. These residual layers use simple block transforms and entropy coding tools designed to remain computationally lightweight and highly parallelisable.

At the decoder, the base layer is decoded using the underlying codec, while the enhancement layers refine the reconstruction through normative upscaling and additive residual reconstruction. This layered approach allows LCEVC to improve compression efficiency and reduce computational complexity by confining the most expensive coding tools of the base codec to a lower-resolution representation.

A detailed description of the LCEVC standard and coding tools is available in [5].

B. LCEVC Integration with VVenC

VVenC is an optimized, open-source software implementation of the Versatile Video Coding (VVC/H.266) standard [11] developed by Fraunhofer HHI to deliver near–reference-level compression efficiency at significantly reduced computational complexity. It retains the full suite of VVC coding tools while applying a range of algorithmic accelerations to approximate VTM behaviour at far lower computational cost [12]. Through various optimizations, VVenC achieves orders-of-magnitude speed improvements over the VTM reference encoder, enabling practical use in UHD streaming, broadcast, and large-scale transcoding pipelines.

An integration of LCEVC with VVenC is available as an SDK at [13]. The integration is implemented within the FFmpeg framework, where VVenC generates a compliant base bitstream while LCEVC produces the enhancement data. This layered combination preserves VVC coding efficiency while significantly improving encoding throughput for UHD content.

III. Methodology

A. Dataset

The evaluation was conducted on a diverse set of UHD (3840×2160) sequences selected from publicly available sources, including Netflix Open Content [14] and the Fraunhofer HHI JVET Common Test Conditions (CTC) material [15]. Ten sequences were chosen to cover a broad range of spatial and temporal complexities. Fig. 2 illustrates the distribution of Spatial Information (SI) and Temporal Information (TI) values, computed in accordance with ITU-T Rec. P.910 [16]. The list of sequences and configuration scripts will be released in a public repository upon publication.

TABLE I: UHD Dataset for LCEVC-VVenC Evaluation

Name	Bit depth	FPS	Frames
VL_Campfire	10	29.97	300
VL_ParkRunning3	10	50	300
VL_CatRobot1	10	60	300
NFX_FoodMrkt2	8	60	300
NFX_Bar1	8	30	210
NFX_Dinner	8	30	384
HD Club	8	30	472
Samsung Boat	8	24	150
Samsung Niagara	8	24	150
Samsung BayRocks	8	24	150

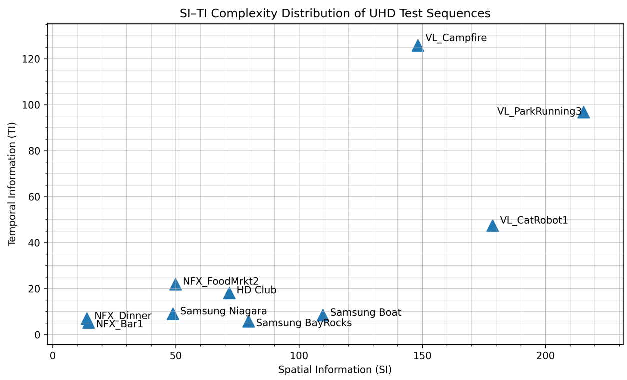

Fig. 2: Spatio-temporal complexity plot of the ten sequences considered in the evaluation dataset.

The VVenC encoder version 1.12 was used for both native VVC encoding and as the base encoder for LCEVC. The LCEVC SDK version 3.15 was employed, with both running within the FFmpeg 6.1 framework. Each clip was encoded at four operating points covering a broad quality range, with QP/CRF between 23 and 38, producing UHD bitrates between 0.5 Mbps and 80 Mbps. The VVenC preset was set to *fast* for both native VVenC and the LCEVC base layer, while LCEVC was tuned for optimal visual quality (*lcevc_tune=vq*). The encoder is available for download at [13].

B. Objective Video Quality Metrics

Video quality was evaluated using four objective metrics: PSNR, SSIM, VMAF, and VMAF-NEG. All metrics were computed using the VMAF framework with model version 3.0.0 [17], [18], applying the UHD model to all encodes to ensure consistent evaluation across resolutions.

To analyse the relationship between objective metrics and perceived quality, a correlation analysis was conducted between the objective scores and the subjective Mean Opinion Scores (MOS) obtained from the subjective evaluation described in the following subsection. The Pearson correlation coefficient was computed between the average MOS values and the corresponding objective metric scores.

C. Subjective Quality Assessment Methodology

A formal subjective evaluation was conducted at an independent test facility to compare native VVenC encoding with LCEVC-enhanced VVenC. The study employed the Degradation Category Rating (DCR) methodology using an 11-grade impairment scale [19].

Testing was performed on two calibrated 65-inch LG CX6LA 4K displays (3840×2160), connected via HDMI to a high-performance PC. Viewers were seated at a distance of 1.5H, corresponding to a viewing angle below 60°.

A total of 24 participants (10 female, 14 male), aged between 18 and 24, took part in the study. All subjects were screened prior to testing for normal visual acuity and colour vision, and post-test screening was conducted to ensure scoring consistency.

Ten UHD test sequences encoded at four bitrate points (decoded to YUV) were evaluated, resulting in a total of 80 test conditions. The selected bitrate points enabled two complementary comparisons:

1) evaluation of bitrate savings at approximately equivalent quality levels, and
2) evaluation of quality differences at approximately equivalent bitrates.

This pairing strategy allows a balanced assessment of both compression efficiency and perceptual quality.

D. Performance Evaluation

Encoding performance was evaluated by measuring both the frame throughput (frames per second, fps) and the total encoding time in seconds.

Performance measurements were conducted using the Hyperfine benchmarking tool [20] on a workstation equipped with a 12th-generation Intel Core i9-12900K CPU and 64 GiB of RAM. Tests were executed using a single-threaded sequential execution model to avoid parallelism-related variability, and each configuration was repeated three times to ensure stable measurements.

IV. Results

This section presents a comprehensive evaluation of LCEVC-enhanced VVenC compared to native VVenC encoding for UHD content. Performance is analyzed in terms of rate–distortion behaviour, subjective quality assessment, and computational complexity. Objective metrics are reported through R–D curves and BD-rate analysis, while subjective performance is examined using MOS, BD-MOS, and correlation with quality metrics. Finally, encoding complexity measurements are provided to assess the practical deployability of the layered approach.

A. Rate–Distortion Curves

Fig. 3 presents the R–D curves averaged across all tested clips for the four evaluated metrics.

At an aggregated level, LCEVC-enhanced VVenC outperforms native VVenC encoding on VMAF and SSIM, and on VMAF-NEG at lower bitrates, while PSNR shows a markedly different behaviour.

B. BD-Rate Analysis

Table II summarises the BD-rate results of LCEVC-enhanced VVenC relative to native VVenC across the evaluated UHD sequences.

TABLE II: BD-Rate Results of LCEVC–VVenC Relative to Native VVenC

	BD-rate (%)			
Content	**VMAF**	**VMAF-NEG**	**SSIM**	**PSNR**
HD_Club	-25.90	9.23	-10.42	180.73
Samsung Niagara	-38.45	-27.90	-27.47	-5.10
Samsung Bayrocks	-42.82	1.82	-27.32	126.54
Samsung Boat	-23.82	51.51	-9.39	364.71
NFX_Bar1	-34.19	-18.18	-14.29	0.44
NFX_Dinner	-44.82	-30.76	-21.53	-10.75
NFX_FoodMrkt2	-22.06	11.42	2.38	53.00
VL_Campfire	-29.73	-14.10	-23.30	-0.35
VL_CatRobot1	-13.84	19.68	1.44	51.32
VL_ParkRunning3	-25.68	-14.73	-11.55	12.30
Average	**-30.13**	**-1.20**	**-14.14**	**77.28**

On VMAF, LCEVC–VVenC achieves consistent BD-rate gains across all content, with improvements ranging from -13.84% to -44.82% and an average gain of -30.13%. The largest gains are observed for Netflix Chimera Dinner (-44.82%) and Netflix Chimera Bar1 (-34.19%).

The VMAF-NEG results are more content-dependent. While most sequences still show BD-rate gains, a few clips (e.g., HD Club, NFX FoodMrkt2, and VL CatRobot1) exhibit small BD-rate penalties under this metric.

SSIM results show consistent improvements, with an average BD-rate gain of approximately -14%.

PSNR results diverge from the other metrics, showing positive BD-rates for several sequences, which is consistent with the behaviour observed in the R–D curves. This divergence is not unexpected. PSNR measures pixel-wise reconstruction error and is known to correlate poorly with perceived visual quality, particularly for modern perceptually optimised codecs and processing pipelines. In the case of LCEVC, several visual quality tools such as adaptive downsampling and dithering are designed to improve perceptual appearance but may increase pixel-level deviations from the reference. As a result, PSNR can penalise perceptually favourable reconstructions despite favourable subjective quality. Perceptually oriented metrics

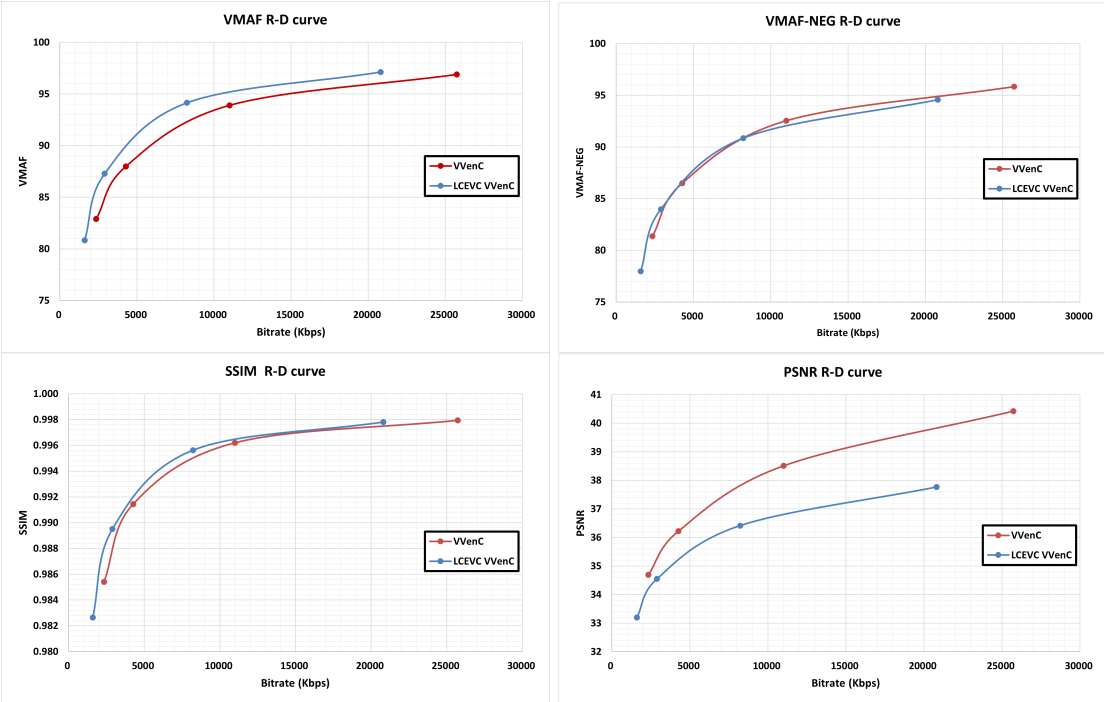

Fig. 3: Average rate–distortion (R–D) curves across all evaluated UHD sequences for the four objective quality metrics.

such as VMAF therefore provide a more meaningful indication of visual performance in this context. This observation is further supported by the subjective evaluation results and metric correlation analysis presented in the next sections.

C. Subjective Assessment Results

The MOS versus bitrate curves for all evaluated sequences are presented in Fig. 4. The plots show the average MOS across all subjects, together with 95% confidence intervals.

Across most sequences, the LCEVC-VVenC curves lie above the corresponding native VVenC curves, indicating that similar perceptual quality can be achieved at lower bitrates. In several cases the MOS values of the two encoders fall within overlapping confidence intervals, suggesting that the perceived quality differences are small or not statistically significant. This observation is further supported by a paired t-test performed on the MOS scores across rate points for each sequence, with the resulting p-values reported in Table III. The analysis indicates that statistically significant differences ($p < 0.05$) are observed in four sequences, while two sequences exhibit borderline significance ($p \approx 0.05$–0.06). The analysis suggests that for comparable MOS levels, LCEVC–VVenC consistently operates at lower bitrates than native VVenC, indicating improved compression efficiency without perceptual quality degradation.

To further quantify these results, BD-MOS values computed across the evaluated bitrate range are summarised in Table III. On average, LCEVC–VVenC achieves a BD-MOS improvement of approximately -23% relative to native VVenC across the dataset. These results confirm that the objective BD-rate gains observed in the R–D analysis translate into validated bitrate savings while maintaining comparable perceptual quality.

TABLE III: BD-MOS results of LCEVC–VVenC relative to native VVenC, with corresponding p-values from paired t-tests and significance indicators.

Content	BD-MOS (%)	p-value	Sig.
HD_Club	-13.53	0.16	×
Samsung Niagara	-44.74	**0.01**	✓
Samsung Bayrocks	-27.33	**0.01**	✓
Samsung Boat	-11.29	**0.04**	✓
NFX_Bar1	-19.19	0.63	×
NFX_Dinner	-35.55	**0.06**	∼
NFX_FoodMrkt2	-4.89	0.10	×
VL_Campfire	-28.07	**0.01**	✓
VL_CatRobot1	-17.85	0.39	×
VL_ParkRunning3	-29.29	0.34	×
Average	**-23.17**		

D. Objective Metrics and Subjective Quality Correlation

The objective and subjective evaluation results collectively demonstrate the performance gains achieved by integrating LCEVC with VVenC. In particular, the subjective evaluation further supports the effectiveness of perceptually oriented objective metrics such as VMAF in approximating perceived visual quality.

Previous studies [21]–[23] have shown that conventional full-reference metrics such as PSNR often exhibit limited correlation with perceptual quality, especially for modern coding systems incorporating enhancement layers or perceptual optimisation techniques. The correlation analysis performed in this study confirms these observations.

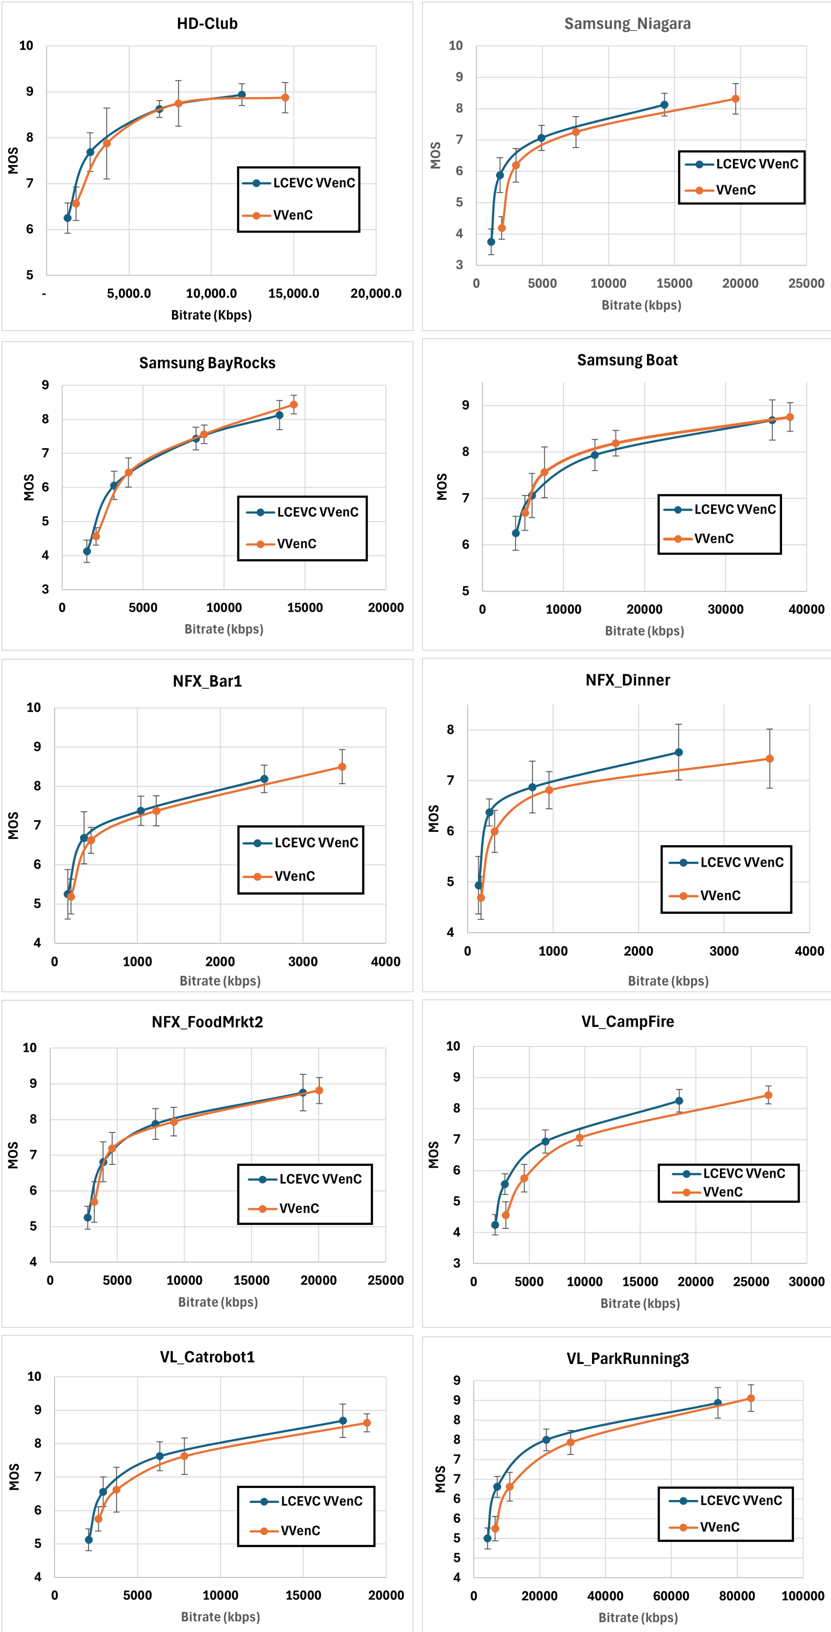

Fig. 4: MOS versus bitrate curves for the evaluated UHD sequences.

TABLE IV: Pearson Correlation Between Subjective MOS and Objective Metrics

Metric	Native VVenC	LCEVC–VVenC
PSNR	0.75	0.49
SSIM	0.76	0.79
VMAF	0.90	0.89
VMAF-NEG	0.90	0.89

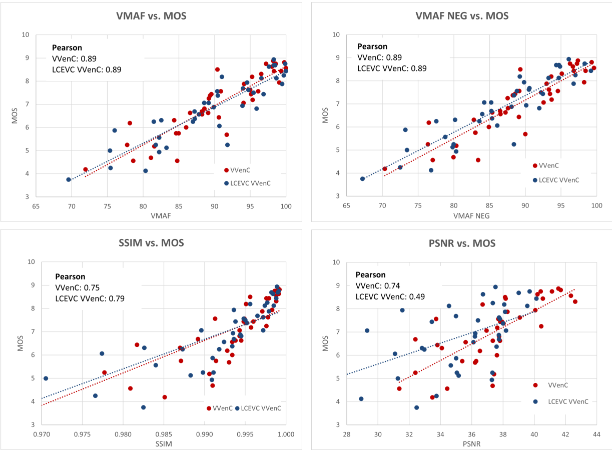

Fig. 5: Scatter plots showing correlation between MOS and objective metrics for VVenC and LCEVC-VVenC encodes.

Table IV summarises the correlation coefficients between the subjective Mean Opinion Scores (MOS) and the evaluated objective metrics for both native VVenC and LCEVC-enhanced VVenC encodes. The results show that VMAF and VMAF-NEG achieve the highest correlation with subjective scores, indicating their suitability for evaluating perceptually optimised coding approaches. In contrast, PSNR demonstrates significantly lower correlation with MOS.

The scatter plots in Fig. 5 further illustrate the relationship between objective metrics and subjective scores for both codecs. VMAF and VMAF-NEG exhibit strong linear trends with MOS, while PSNR shows greater dispersion and weaker predictive capability.

As a visual example, Fig. 6 presents frames from native VVenC and LCEVC–VVenC encodes at comparable quality levels (QP 23). Although the LCEVC encode exhibits a PSNR value approximately 4 dB lower than the corresponding VVenC encode, both VMAF and MOS indicate perceptually equivalent quality within the just noticeable difference (JND) range. This example highlights the limitations of PSNR as a quality indicator for perceptually optimised and scalable coding approaches such as LCEVC.

E. Encoding Speed

Encoding performance was evaluated by measuring the encoding throughput in average frames per second (fps) and the total encoding duration in seconds, as reported in Table V.

Integrating LCEVC with VVenC significantly reduces encoding complexity. Across the tested sequences, LCEVC–VVenC achieves an average encoding throughput of 0.82 fps

Fig. 6: Visual comparison of a frame from the HD Club clip. Despite PSNR difference of 4 dB, VVenC and LCEVC-VVenC are perceptually similar.

TABLE V: Encoding Performance Comparison Between Native VVenC and LCEVC–VVenC

Performance KPI	VVenC	LCEVC–VVenC	Performance Improvement
Avg. encoded frames/sec	0.31	0.82	2.61×
Avg. encoding time (s)	1758.39	737.09	−58%
Avg. achieved bitrate (kbps)	24057.89	19927.33	−17%

compared to 0.31 fps for native VVenC, corresponding to an average speed-up of approximately 2.6×.

The reduction in encoding time is also reflected in the total processing duration, which decreases from an average of 1758 seconds for native VVenC to 737 seconds for LCEVC–VVenC.

The speed-up becomes more pronounced for sequences with higher spatial and temporal complexity. For example, for the ParkRunning3 sequence, encoding speed improvements of up to 3× were observed.

These gains stem from the LCEVC architecture, where the base encoder processes a lower-resolution representation of the video while the enhancement layer encodes residual details using lightweight processing.

V. Conclusion

This paper evaluated the integration of LCEVC with VVenC for UHD video encoding in terms of compression efficiency, subjective quality, and encoding complexity. LCEVC–enhanced VVenC achieves an average −30% VMAF BD-rate improvement relative to native VVenC. Subjective evaluation confirms these gains, with BD-MOS indicating an average −23% bitrate reduction at equivalent perceived quality. In addition, the approach significantly reduces encoding complexity, achieving an average 2.6× encoding speed-up. The study also shows that perceptually aligned metrics such as VMAF correlate strongly with subjective quality, whereas PSNR exhibits weaker predictive capability for enhancement-based coding approaches. These results demonstrate that LCEVC can improve both compression and computational efficiency for UHD video workflows. Future work will investigate decoding performance and end-to-end system optimisation in practical streaming environments.

ACKNOWLEDGMENT

AI Acknowledgement: The authors used ChatGPT to assist with drafting, structuring, and language editing. All outputs were reviewed, modified, and verified to ensure accuracy, originality, and compliance with academic standards.

REFERENCES

[1] B. Bross, Y.-K. Wang, Y. Ye, S. Liu, J. Chen, G. J. Sullivan, and J.-R. Ohm, "Overview of the versatile video coding (VVC) standard and its applications," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 10, pp. 3736–3764, 2021.

[2] B. Bross, J. Chen, J.-R. Ohm, G. J. Sullivan, and Y.-K. Wang, "Developments in international video coding standardization after AVC, with an overview of versatile video coding (VVC)," *Proceedings of the IEEE*, vol. 109, no. 9, pp. 1463–1493, 2021.

[3] A. Saha, M. Chavarrías, F. Pescador, Á. M. Groba, K. Chassaigne, and P. L. Cebrián, "Complexity analysis of a versatile video coding decoder over embedded systems and general purpose processors," *Sensors*, vol. 21, no. 10, p. 3320, 2021.

[4] Z. Azimi, A. Premkumar, R. Farahani, V. V. Menon, C. Timmerer, and R. Prodan, "Decoding complexity-aware bitrate-ladder estimation for adaptive VVC streaming," in *2024 32nd European Signal Processing Conference (EUSIPCO)*. IEEE, 2024, pp. 942–946.

[5] G. Meardi, S. Ferrara, L. Ciccarelli, G. Cobianchi, S. Poularakis, F. Maurer, S. Battista, and A. Byagowi, "MPEG-5 part 2: Low complexity enhancement video coding (LCEVC): Overview and performance evaluation," *Applications of Digital Image Processing XLIII*, vol. 11510, pp. 238–257, 2020.

[6] S. Battista, G. Meardi, S. Ferrara, L. Ciccarelli, F. Maurer, M. Conti, and S. Orcioni, "Overview of the low complexity enhancement video coding (LCEVC) standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 11, pp. 7983–7995, 2022.

[7] A. S. S. Chaubet, R. A. Vaz, G. H. M. G. de Oliveira, R. S. Rabaca, I. C. Dourado, G. de Melo Valeira, and C. Akamine, "TV 3.0: An overview," *IEEE Transactions on Broadcasting*, vol. 71, no. 1, pp. 11–18, 2025.

[8] N. Barman, S. Schmidt, S. Zadtootaghaj, and M. G. Martini, "Evaluation of MPEG-5 part 2 (LCEVC) for live gaming video streaming applications," in *Proceedings of the 1st Mile-High Video Conference*, 2022, pp. 108–109.

[9] S. Battista, G. Meardi, S. Ferrara, L. Ciccarelli, F. Maurer, M. Conti, and S. Orcioni, "Verification test of the low complexity enhancement video coding (LCEVC) standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 33, no. 8, pp. 4269–4278, 2023.

[10] G. Cobianchi, G. Meardi, S. Poularakis, A. Walisiewicz, O. Abdelkafi, F. B. Amara, F. Kossentini, C. Stejerean, and H. Tmar, "Enhancing SVT-AV1 with LCEVC to improve quality-cycles trade-offs and enhance sustainability of vod transcoding," in *Applications of Digital Image Processing XLV*, vol. 12226. SPIE, 2022, pp. 239–258.

[11] A. Wieckowski, J. Brandenburg, T. Hinz, C. Bartnik, V. George, G. Hege, C. Helmrich, A. Henkel, C. Lehmann, C. Stoffers *et al.*, "VVenC: An open and optimized VVC encoder implementation," in *2021 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 2021, pp. 1–2.

[12] A. Mercat, A. Mäkinen, J. Sainio, A. Lemmetti, M. Viitanen, and J. Vanne, "Comparative rate-distortion-complexity analysis of VVC and HEVC video codecs," *IEEE Access*, vol. 9, pp. 67 813–67 828, 2021.

[13] V-Nova, "V-Nova LCEVC SDK," https://download.v-nova.com/, 2025, accessed: 2025-07-21.

[14] Netflix, "Netflix Open Content," https://opencontent.netflix.com/, 2025, accessed: 2025-07-21.

[15] Joint Video Experts Team (JVET), "JVET – Standardization of Video Coding," https://jvet.hhi.fraunhofer.de/, 2025, accessed: 2025-07-21.

[16] "ITU-T Recommendation P.910: Subjective Video Quality Assessment Methods for Multimedia Applications," International Telecommunication Union, Geneva, Switzerland, ITU-T Recommendation, 2008.

[17] Netflix, "VMAF: Video Multi-Method Assessment Fusion," https://github.com/Netflix/vmaf, 2025, accessed: 2025-07-21.

[18] Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy, and M. Manohara, "Toward a practical perceptual video quality metric, 2016," *http://techblog. netflix. com/2016/06/toward-practical-perceptual-video. html*, 2016.

[19] "Recommendation ITU-R BT.500-14: Methodologies for the Subjective Assessment of the Quality of Television Images," International Telecommunication Union, Geneva, Switzerland, ITU-R Recommendation, 2019.

[20] Peter, D., "Hyperfine (Version 1.16.1)," https://github.com/sharkdp/hyperfine, 2023, accessed: 2025-07-21.

[21] Y. Blau and T. Michaeli, "The perception-distortion tradeoff," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6228–6237.

[22] W. Zhou, H. Amirpour, C. Timmerer, G. Zhai, P. L. Callet, and A. C. Bovik, "Perceptual visual quality assessment: Principles, methods, and future directions," *arXiv preprint arXiv:2503.00625*, 2025.

[23] M. Wien and J. Jung, "Remote expert viewing, laboratory tests or objective metrics: which one (s) to trust?" *EURASIP Journal on Image and Video Processing*, vol. 2024, no. 1, p. 16, 2024.