Advanced Hash-based Distributed Video Coding

Frederik Verbist

Distributed video coding (DVC), also known as Wyner-Ziv video coding, provides low-complexity encoding solutions for video. In contrast to traditional predictive coding systems, which capture the temporal redundancies at the encoder side, this task is transferred to the decoder. This radical shift of the computational burden towards the decoder brought in DVC is particularly attractive for recording devices with limited computational power or severe energy constraints. The performance of Wyner-Ziv coding strongly depends on the quality of the decoder-generated predictions of the original source material, referred to as side information (SI). This dissertation focuses on hash-based DVC systems, a particular class of DVC architectures, which rely on auxiliary information, so-called hash information, sent by the encoder to assist the decoder during the SI generation process. The presented systems create hash information as a coarse approximation of the original frames. At the decoder, overlapped block motion estimation between the hash and reference frames links every predicted pixel to a set of candidate temporal predictors, which invites the use of pixel-based multi-hypothesis motion-compensated prediction. In this context, this dissertation presents several hash-based DVC architectures and confirms their superior performance with respect to benchmark DVC systems, in particular for complex-motion video. At the same time, the encoding complexity compared to traditional predictive coding architectures is severely reduced. Additionally, this dissertation offers a novel approach to pixel-based multi-hypothesis motion-compensated prediction by integrating the result of an original spatial correlation estimate between the predicted pixel and its set of candidate predictors into the motion compensation process.

In order to further improve the compression performance of the presented hash-based DVC architectures, this dissertation brings forth a novel strategy to apply the principles of successive SI refinement at the decoder. The proposed scheme decodes the original frames in stages, called refinement levels, which are built according to a novel technique. The resulting partially decoded frames gradually increase the knowledge of the original frames at the decoder as more levels are completed. This is exploited by regenerating SI based on the partially decoded frames at every refinement level, which significantly increases the coding efficiency at subsequent levels. It is shown that the proposed refinement strategy significantly improves the overall compression performance without inflating the decoding complexity.

Finally, this dissertation targets a particular drawback of common DVC architectures, which is the need for a return channel from decoder to encoder for optimal rate control. In this context, this dissertation includes a novel strategy to suppress the presence of a feedback channel, thereby widening the deployment range of the proposed architectures to, e.g., unidirectional or storage applications. The pursued approach includes an involved encoder-side rate control scheme with several coding options, in conjunction with specific decoder-side measures.

Although feedback-free rate control is no longer optimal, the proposed scheme, integrated in a competitive hash-based DVC, delivers fair performance, even surpassing the feedback-based benchmark in DVC on high-motion sequences, while low-complexity encoding characteristics are maintained.