22.1 A 125μW, Fully Scalable MPEG-2 and H.264/AVC Video Decoder for Mobile Applications

Tsu-Ming Liu[1], Ting-An Lin[2], Sheng-Zen Wang[2], Wen-Ping Lee[1], Kang-Cheng Hou[1], Jiun-Yan Yang[1] Chen-Yi Lee[1]

[1]National Chiao Tung University, Hsinchu, Taiwan
[2]Mediatek, Hsinchu, Taiwan

A single-chip MPEG-2 SP@ML and H.264/AVC BL@L4 video decoder is fabricated in a 0.18μm 1P6M CMOS technology with an area of 15.21mm^2. This chip contains 19.2kb and 3.55kb of embedded SRAM for storing neighboring pixels and control tags, and adopts two 4MB SDRAMs for further system integration. It operates at a power-level that is about one order of magnitude less than comparable decoders. This savings in power consumption was attained by means of both throughput and bandwidth improvements while incorporating scalable features. For mobile applications, MPEG-2 and H.264/AVC video decoding of QCIF sequences at 15 frames per second is achieved at a clock frequency of 1.15MHz and requires 108μW and 125μW, respectively, at 1V supply voltage. Moreover, CIF, D1 and HD resolutions are also supported. The chip features are summarized in Fig. 22.1.1.

The advent of H.264/AVC provides high compression ratio, but, there is no backward compatibility to the prevalent MPEG-x and H.26x video coding standards. MPEG-2 [1] and H.264/AVC [3] processors have been reported at ISSCC. However, these solutions used separate modules and only processed a single type of video content in each module. To support different system requirements such as DVB-H or HD-DVD, a scalable pipeline is exploited to efficiently integrate both MPEG-2 and H.264/AVC in a single chip.

Figure 22.1.2 shows the system block diagram of the proposed dual-video decoder chip. The SDRAM interface is exploited to access the external SDRAM. Reading and writing processes are issued by the motion compensation and deblocking filter, respectively. Furthermore, to reduce the bandwidth between the external SDRAM and deblocking filter, a separate data bus and display engine are utilized for on screen display (OSD) through a direct display interface. The syntax parser, entropy decoder, inverse transformation and deblocking filter for MPEG-2 and H.264/AVC standard have been tightly combined based on characteristics of pipeline scalability. To achieve low power consumption, a simple prediction circuit is interfaced to the embedded SRAM to make a better trade-off between memory cost and transmission bandwidth. Furthermore, a high-throughput motion compensation and deblocking filter is developed to reduce the clock frequency and lower the power requirements.

Figure 22.1.3 shows the scalable pipeline for the dual-standard architecture. In H.264/AVC, a 4×4 sub-block is the smallest element to be processed. However, an 8×8 block size is adopted in the MPEG-2 standard. To integrate them efficiently, the buffer size is kept as 8×8 and transfers for both standards are into a 4×4 processing unit since all blocks can be considered as a super-set of a 4×4 sub-block. Compared to a macroblock level pipeline, the power consumption of the pipelined registers is reduced by 75% in MPEG-2 and 93.75% in H.264/AVC, through a clock gating scheme. Input data of the IDCT module is partitioned into even and odd components and they are processed as a 4×4 IDCT, through a recursive IDCT algorithm. The in-loop filter is defined by H.264/AVC and the post-loop filter follows the prevalent MPEG-x standard. However, the performance improvement is very small (only 0.04dB) when applying in-loop filter as a post filter in the prevalent MPEG standard. In Fig. 22.1.3, a H.264-like algorithm for post filtering is used to retain the filtering performance and reduce the integration cost. Finally, the PSNR gains of 0.2dB can be achieved, as compared to an un-filtered design, with an additional gate count of only 20% of the in-loop filter requirement.

Figure 22.1.4 shows the proposed bandwidth scalability of the prediction circuit. H.264/AVC achieves a high compression ratio since it utilizes the neighboring pixels to obtain a reliable predictor reducing the prediction errors. However, high data correlation also leads to a design challenge in terms of transmission bandwidth and internal memory cost. In the design, a simple prediction circuit, where a 19.2kb pixel SRAM is employed to cache the pixels of upper neighbors, improves the external bandwidth. The key idea is that not all neighboring pixels should be stored in the internal memory. In certain sequences, most edges are determined as a 'horizontal prediction' in intra-prediction or 'SKIP mode' in the deblocking filter. There is no need to keep them for follow-up decoding procedures. The proposed prediction circuit generates a TAG signal to predict whether the pixel data of the next row of a macroblock should be kept or not; but a prediction miss may occur. Therefore, we provide a flexible solution at the architectural level where a compromise is made between external bandwidth and memory cost. Compared to the intra-prediction and deblocking filter in [3], the proposed prediction circuit saves 33% of the bandwidth and 40% of the internal memory size on average.

Figure 22.1.5 shows the processing cycle breakdown of different architectural stages. Several high throughput architectures have been implemented on this chip [4]. A 1×4 decoding order in Figure 22.1.3, context switch buffer and efficient access scheduling is exploited to achieve a 39% cycle reduction in MC. A novel prediction and hybrid schedule reduce the processing cycles a further 35%. Therefore, 101.04Mpixels/s of maximal decoding capability is achieved. Compared to the initial power consumption for 66.21 and 40.43Mpixels/s rates, the savings are 33% and 59%, respectively, for real-time decoding at QCIF resolution.

Figure 22.1.6 presents a comparison of power consumption with existing designs [1][2]. Under an identical design specification, the proposed techniques lead to lower system clock rate and supply voltage and thus lower power dissipation. For mobile applications, the power reduction of this chip is about one order of magnitude compared to existing decoders and could be further improved through voltage scaling. Figure 22.1.7 shows a chip micrograph of this dual-video decoder design.

Acknowledgements:
Authors thank Jeng-Bin Chen, Ching-Che Chung, Wen-Hsiao Peng, Wei-Chin Lee for insightful discussions on this work. Authors also thank Chip Implementation Center (CIC) for testing services.

References:
[1] H. Yamauchi, et al., "A 0.8W HDTV Video Processor with Simultaneous Decoding of Two MPEG-2-MP@HL Streams and Capable of 30frames/s Reverse Playback," *ISSCC Dig. Tech. Papers,* pp. 372-474, Feb., 2002.
[2] Hae-Yong Kang, et al., "MPEG4 AVC/H.264 Decoder with Scalable Bus Architecture and Dual Memory Controller," *IEEE Intl Symp. on Circuits and Systems*, pp. II-145 - II-148, May, 2004.
[3] Yu-Wen Huang, et al., "A 1.3TOPS H.264/AVC Single-Chip Encoder for HDTV Applications," *ISSCC Dig. Tech. Papers,* pp. 128-129, Feb., 2005.
[4] Tsu-Ming Liu, et al., "An 865μW H.264/AVC Video Decoder for Mobile Applications," *IEEE Asian Solid-State Circuits Conference,* pp. 301-304, Nov., 2005.

Specification	**Dual**	**MPEG-2 SP@ML** **H.264/AVC BL@L4**
Technology		**Standard 0.18μm 1P6M CMOS** **1.8V core, 3.3V I/O**
Die Size		**3.9mm × 3.9mm**
Package		**208-pin CQFP**
Logic Gates		**303.78K**
Internal Memory		**22.75Kb SRAMs**
External Memory		**4MB×2 SDRAMs**
Max. System Clock		**100MHz**
Max. Processing Throughput		**101.04 Mpixels/sec**
Core Power Consumption	**MPEG-2**	**108μW (1.15MHz@1V, QCIF@15fps)** **10.4mW (16.6MHz@1.2V, D1@30fps)**
	H.264/AVC	**125μW (1.15MHz@1V, QCIF@15fps)** **12.4mW (16.6MHz@1.2V, D1@30fps)**

Figure 22.1.1: Chip summary.

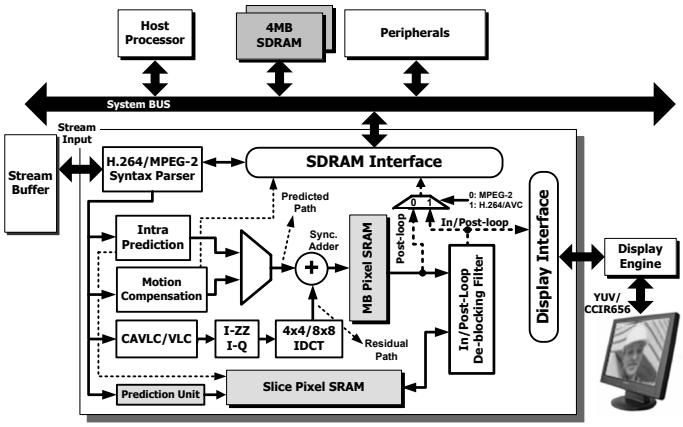

Figure 22.1.2: System block diagram.

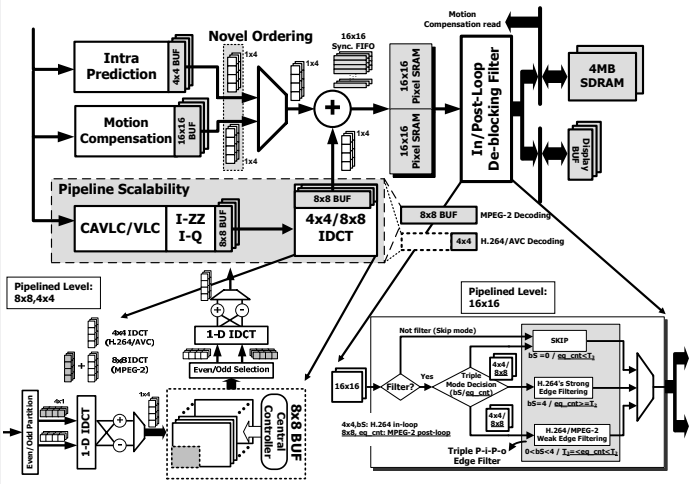

Figure 22.1.3: Pipeline scalability with dual-standard architecture.

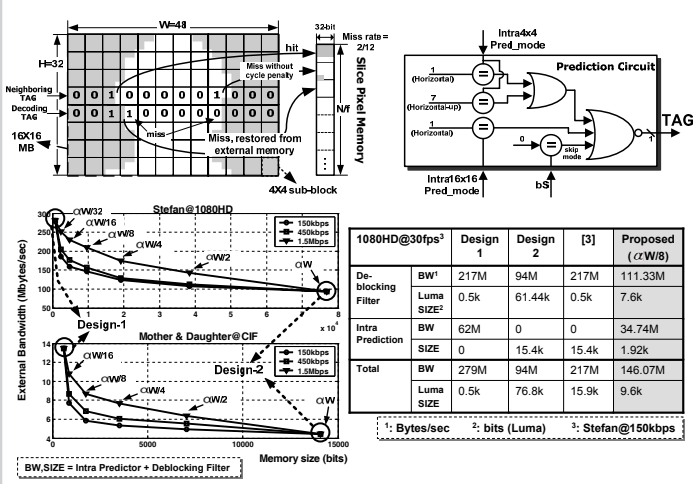

1080HD@30fps[3]		Design 1	Design 2	[3]	Proposed (αW/8)
De-blocking Filter	BW[1]	217M	94M	217M	111.33M
	Luma SIZE[2]	0.5k	61.44k	0.5k	7.6k
Intra Prediction	BW	62M	0	0	34.74M
	SIZE	0	15.4k	15.4k	1.92k
Total	BW	279M	94M	217M	146.07M
	Luma SIZE	0.5k	76.8k	15.9k	9.6k

[1]: Bytes/sec [2]: bits (Luma) [3]: Stefan@150kbps

Figure 22.1.4: Bandwidth scalability with prediction circuit.

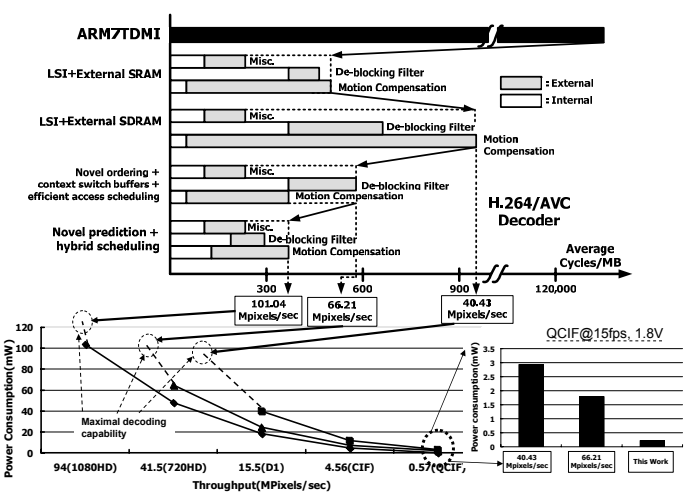

Figure 22.1.5: Throughput improvement of proposed design.

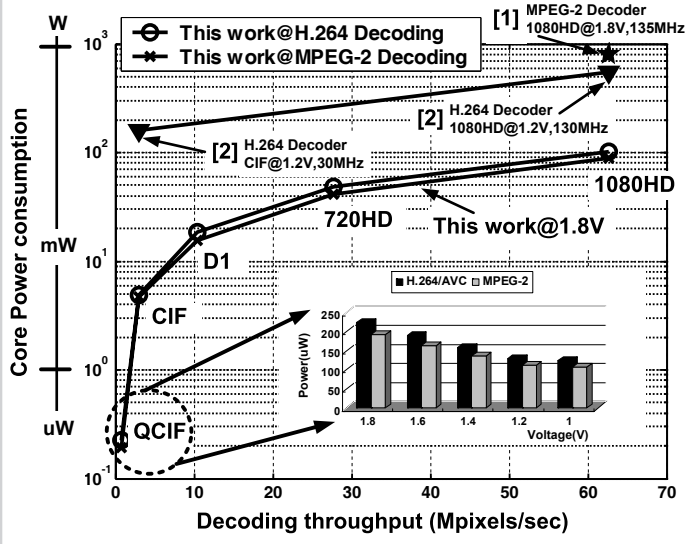

Figure 22.1.6: Power consumption.

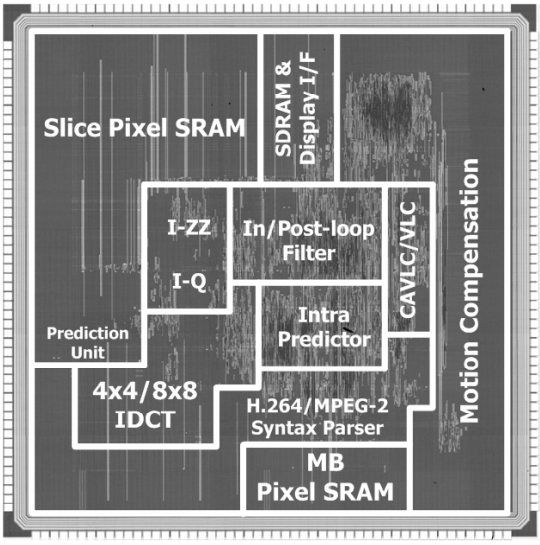

Figure 22.1.7: Chip micrograph.

Specification	**Dual**	**MPEG-2 SP@ML** **H.264/AVC BL@L4**
Technology		**Standard 0.18μm 1P6M CMOS** **1.8V core, 3.3V I/O**
Die Size		**3.9mm × 3.9mm**
Package		**208-pin CQFP**
Logic Gates		**303.78K**
Internal Memory		**22.75Kb SRAMs**
External Memory		**4MB×2 SDRAMs**
Max. System Clock		**100MHz**
Max. Processing Throughput		**101.04 Mpixels/sec**
Core Power Consumption	**MPEG-2**	**108μW (1.15MHz@1V, QCIF@15fps)** **10.4mW (16.6MHz@1.2V, D1@30fps)**
	H.264/AVC	**125μW (1.15MHz@1V, QCIF@15fps)** **12.4mW (16.6MHz@1.2V, D1@30fps)**

Figure 22.1.1: Chip summary.

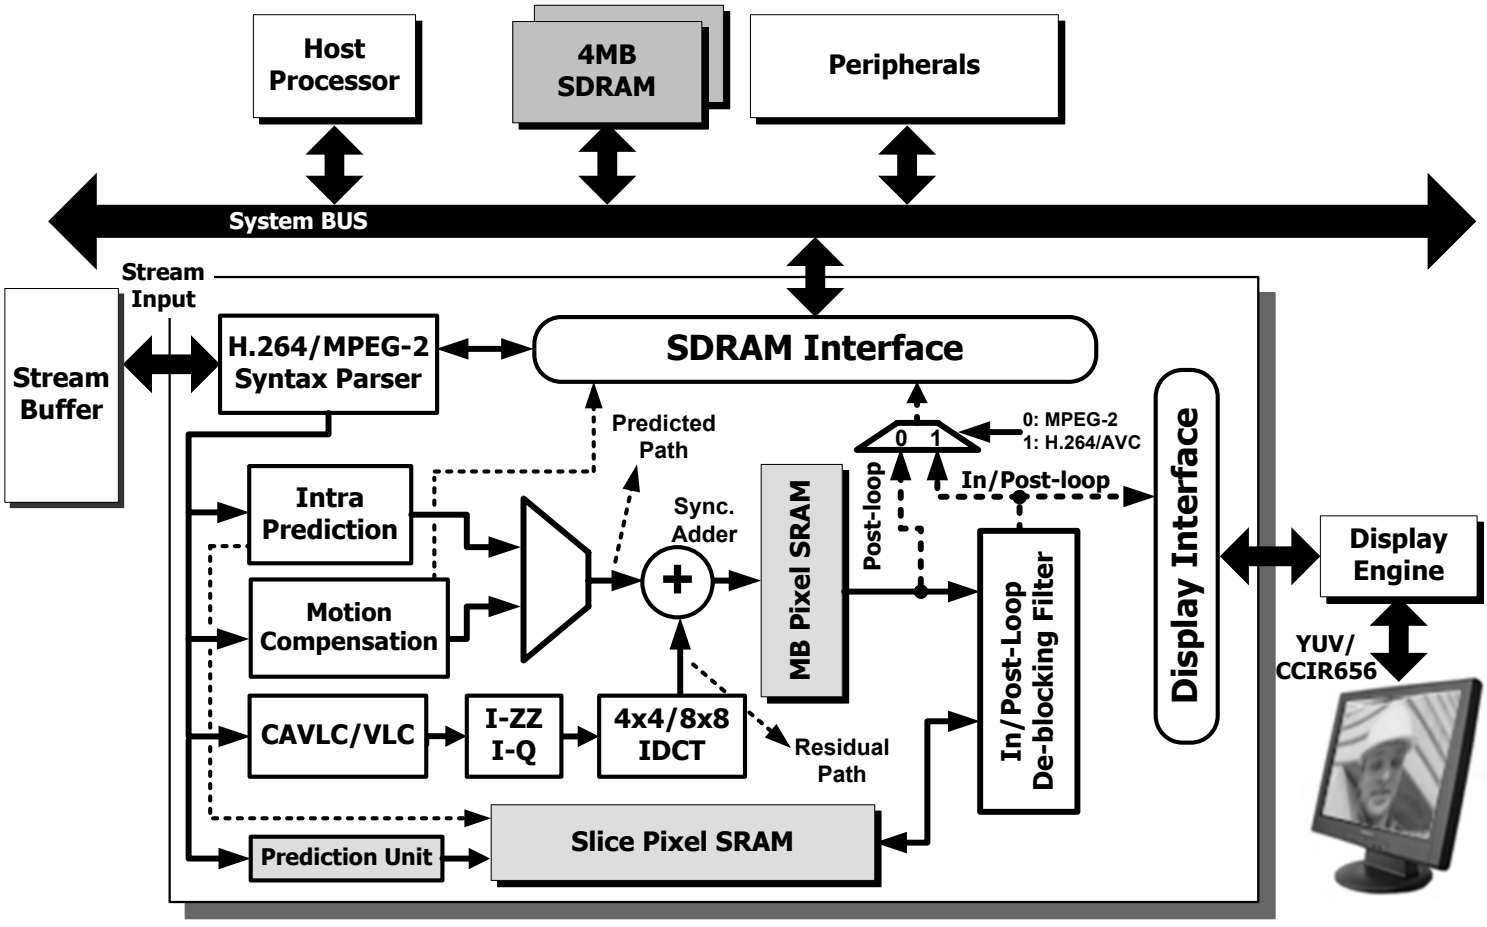

Figure 22.1.2: System block diagram.

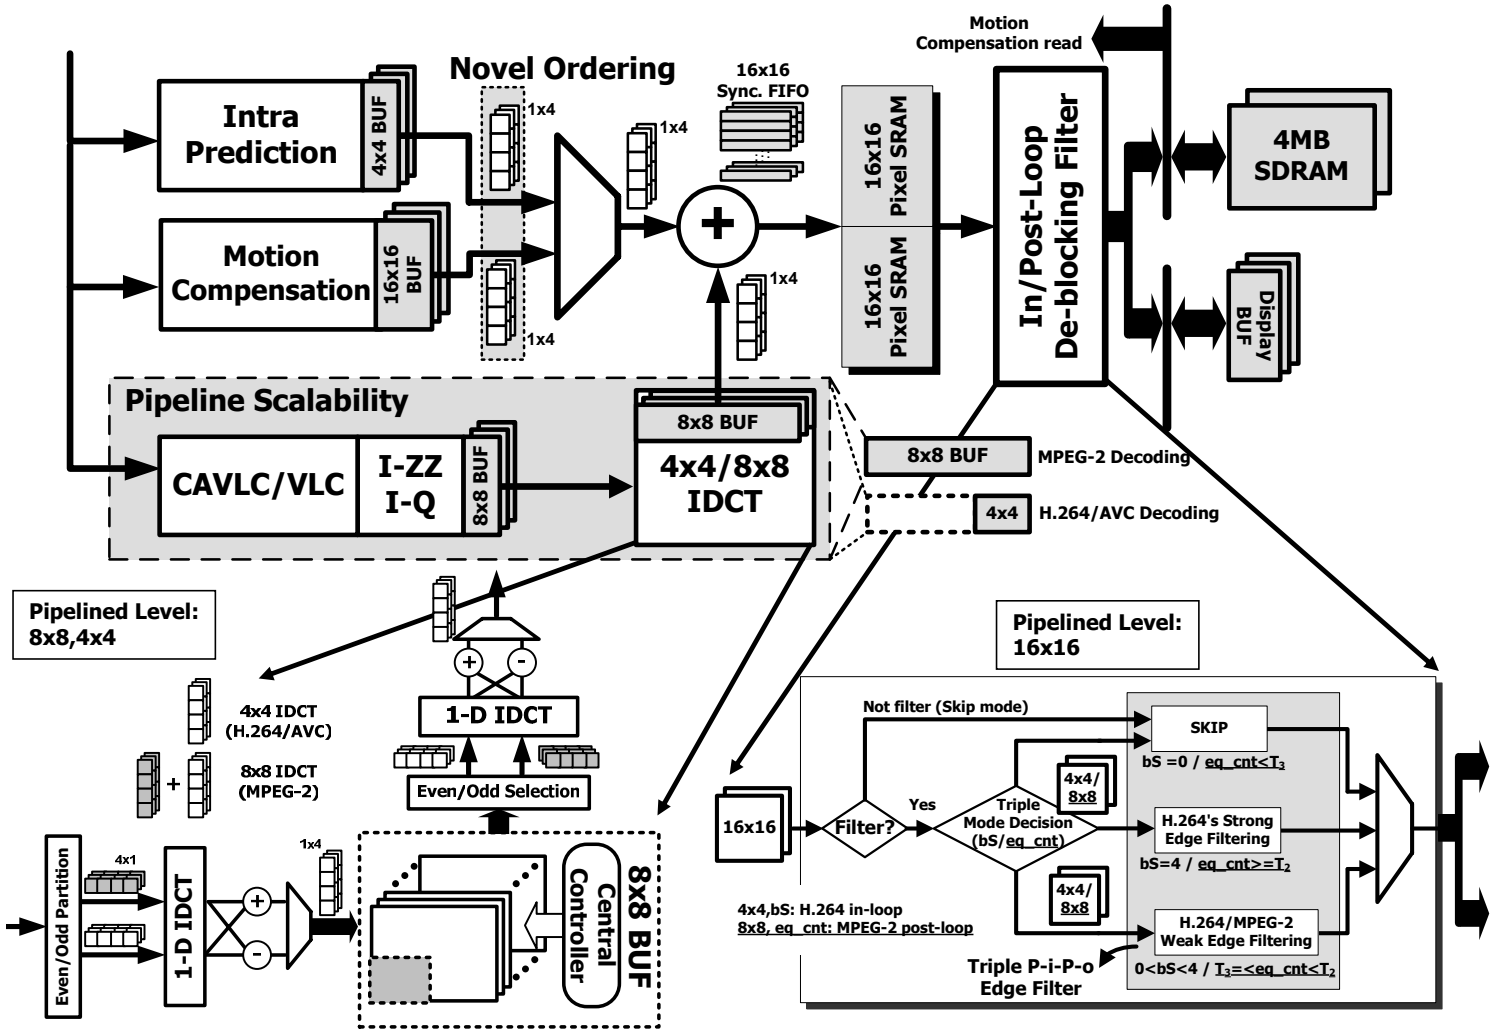

Figure 22.1.3: Pipeline scalability with dual-standard architecture.

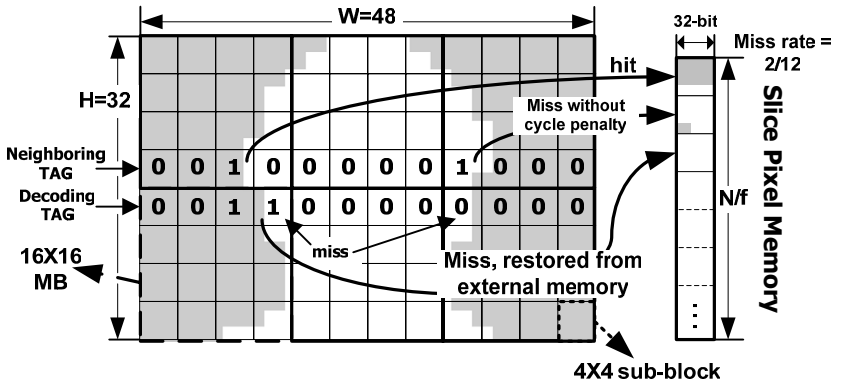

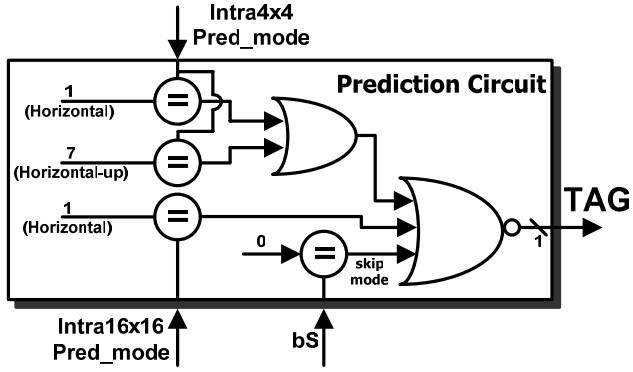

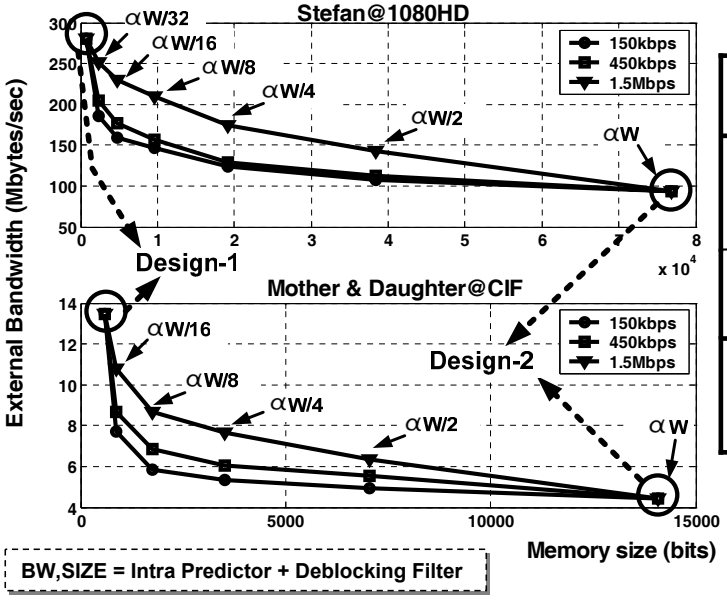

1080HD@30fps[3]		Design 1	Design 2	[3]	Proposed (αW/8)
De-blocking Filter	BW[1]	217M	94M	217M	111.33M
	Luma SIZE[2]	0.5k	61.44k	0.5k	7.6k
Intra Prediction	BW	62M	0	0	34.74M
	SIZE	0	15.4k	15.4k	1.92k
Total	BW	279M	94M	217M	146.07M
	Luma SIZE	0.5k	76.8k	15.9k	9.6k

[1]: Bytes/sec [2]: bits (Luma) [3]: Stefan@150kbps

Figure 22.1.4: Bandwidth scalability with prediction circuit.

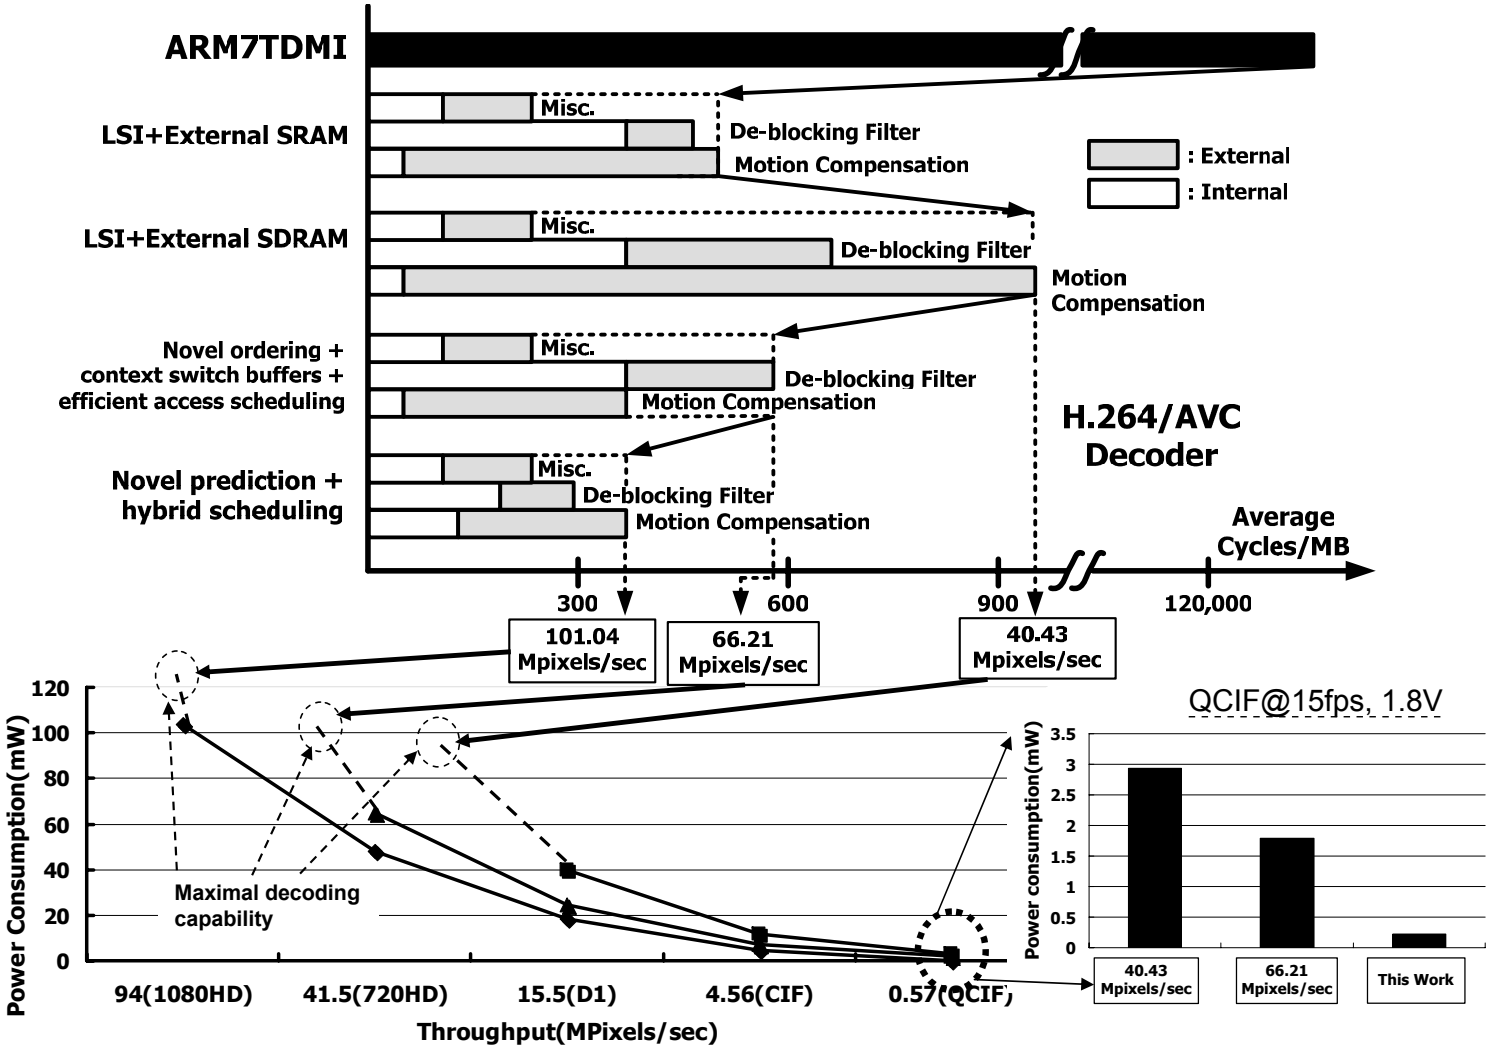

Figure 22.1.5: Throughput improvement of proposed design.

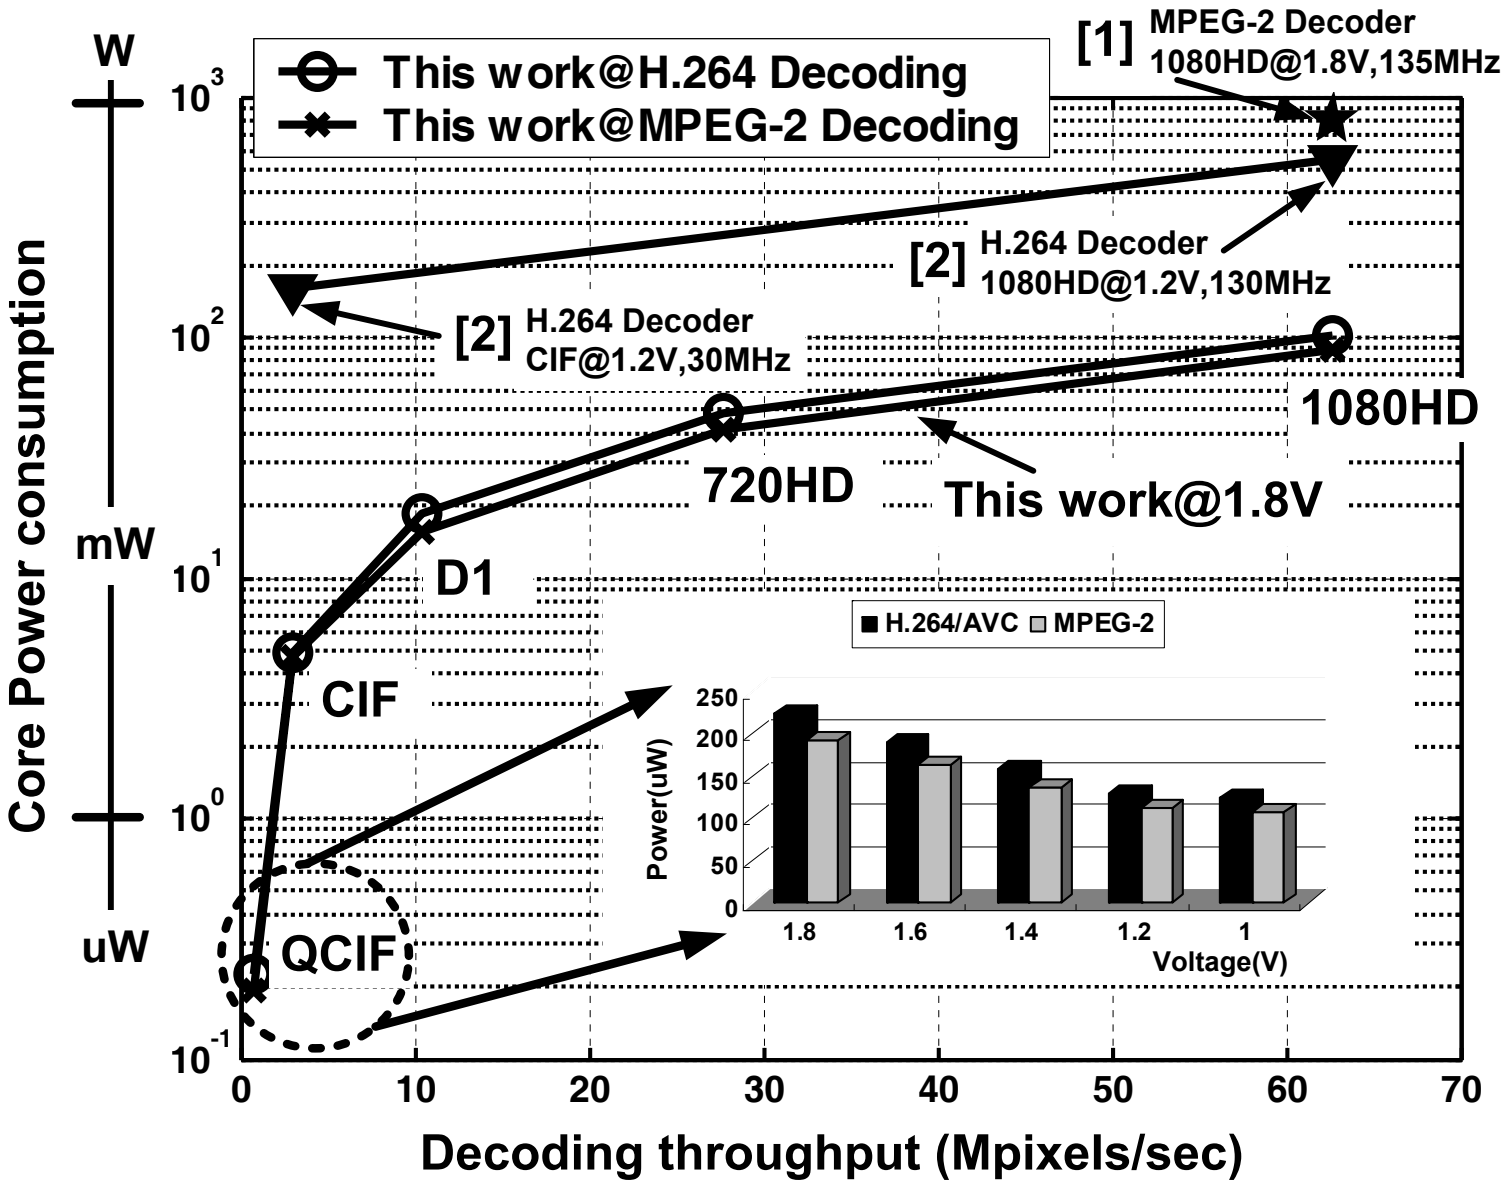

Figure 22.1.6: Power consumption.

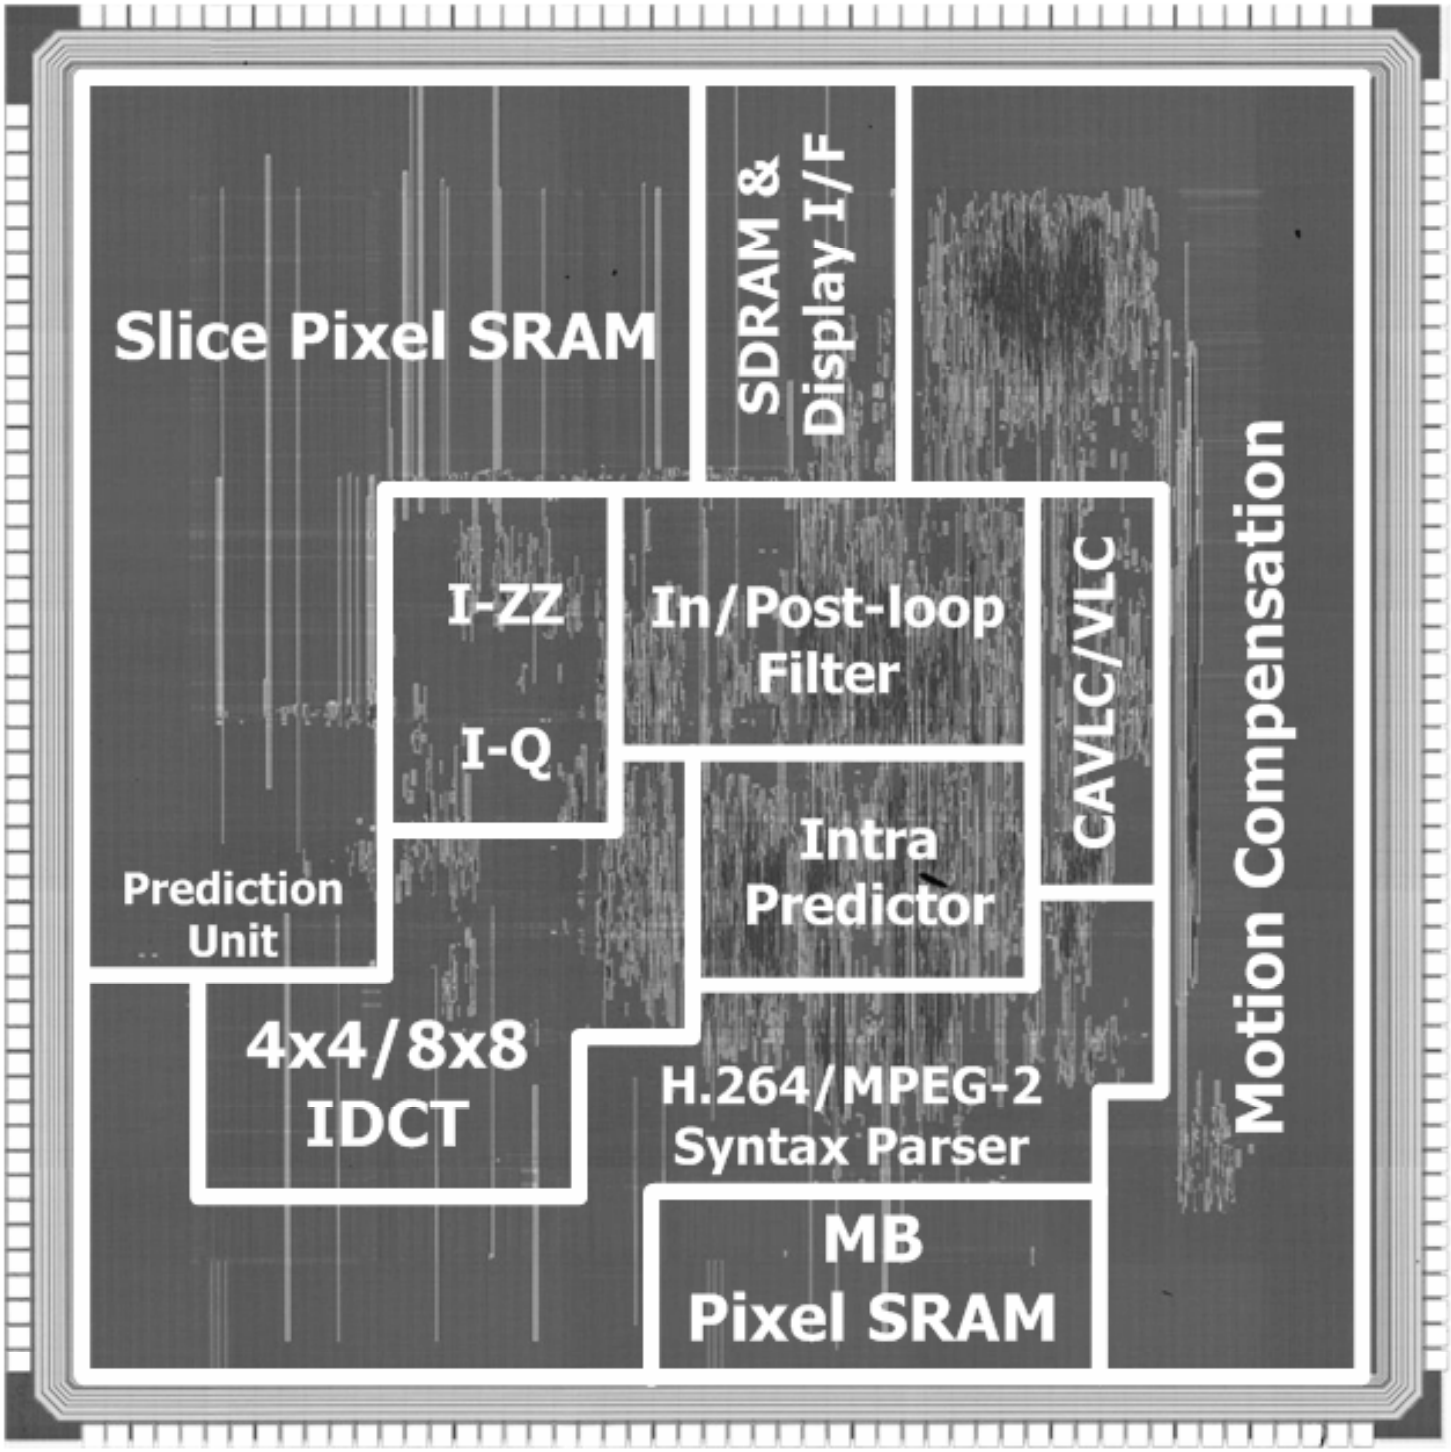

Figure 22.1.7: Chip micrograph.