Chapter 5
Implementation Results

Based on aforementioned techniques, this chapter summarizes the implementation results from a low-power perspective. Although we also discuss integration and error-robust issues on a video decoder, low power dissipation is more crucial to the design of mobile or handheld devices. To obtain a real metric of power, we thoroughly exhibit measured power dissipation to highlight the design breakthroughs on algorithmic and architectural levels. Specifically, we first describe a design methodology of this IC, including front/back-end design flow and verification. After that, this IC is fabricated using 0.18μm CMOS process and measured via a VLSI tester. Moreover, measured results and power comparison have been shown in order to prove that this IC is very suitable for mobile applications where conservative power requirements are essential.

5.1 Design Flow

A design flow that enables an efficient design for low-power demands is depicted in Figure 5.1(a), with entry of C-Language model, Verilog RTL-level descriptions and FPGA verification, then synthesizing and routing with Cadence® RTL Compiler and SoC Encounter™, and ending with chip fabrication as well as verification on an Agilent 93000 SOC test system. Figure 5.1(b) depicts power reduction in different design stages. While 70% of power savings can be achieved by exploring different architectures, the sophisticated use of some advanced features in EDA tools during the synthesis and P&R phases can also play a key role. Hence, to clarify this design breakthrough, the detailed

description is presented from not only front-end but also back-end point of view in Figure 5.1(a). In the front-end design, we use standard-released reference software (Joint Model, JM) to be a high-level C-language model. Then, we formulate and analyze the design problem from algorithmic and architectural levels. After deciding designed architecture, a corresponding hardware description has been made for follow-up designs. Meanwhile, this hardware description is fully verified by a custom-made FPGA emulation board. Because debugging for the design becomes increasingly important and designers want more efficient and high-performance verification and debugging solutions, we adopt Dynalith Systems' iPROVE [106] as our verification prototype. Figure 5.2 shows the verification environment with a snap shot. An iPROVE can be reconfigured through PCI interface on a motherboard. It features a build-in logic analyzer (BILA) for hardware debugging and a data dumping port (DDP) for external interfacing port, such as VGA. BILA helps designers to save valuable time and effort in debugging their design. The resulting waveform can be viewed with any waveform viewer that supports VCD (Value Changed Dump) data format. Moreover, iPROVE supports 256MB SDRAM which is suitable for large capacity of frame memory in multimedia systems. After verifying the functionality of hardware description via iPROVE, we start preparing the related files prior to the back-end design.

In the back-end design, we exploit Cadence's low-power synthesis (LPS) capabilities embedded in physically knowledgeable synthesis (PKS) to achieve timing closure. Specifically, to make a better trade-off between processing cycles and operating speed, it is crucial to shorten the critical path via backend tools. In this design, we aim at a working frequency of 100MHz for the real-time high-resolution (1080HD) video decoding demands. However, the timing gap between pre-layout and post-layout stages may be considerable. Therefore, a problem about timing closures emerges for achieving high-resolution decoding processes. To alleviate the aforementioned problem, we employ the physical wire-load model to facilitate the synthesis process and improve the timing closure. Particularly, it

implements the net-list on a given (minimal) floor-plan and maintains or improves the quality of physical characteristics of the design. On the other hand, a toggle count format (TCF) file that provides an average switching activity for the nets over time is generated to automatically optimize the power of design. In the placement and routing phases, we perform a SI-prevention and timing-aware routing. Finally, layout verification and simulation have been made for design sign-off.

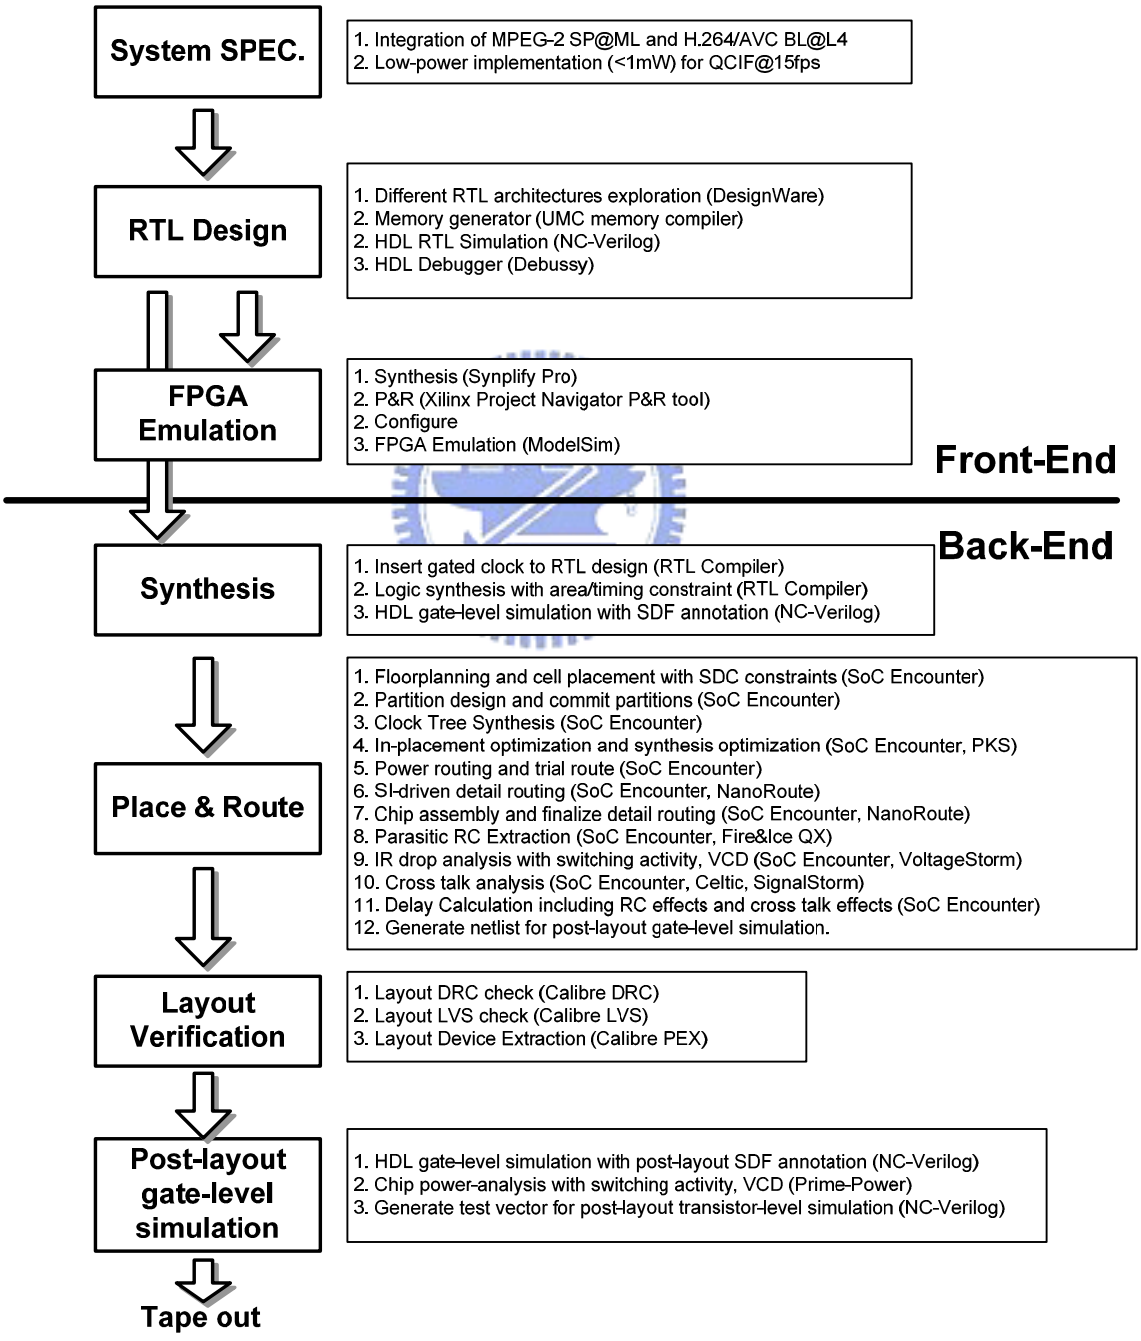

(a)

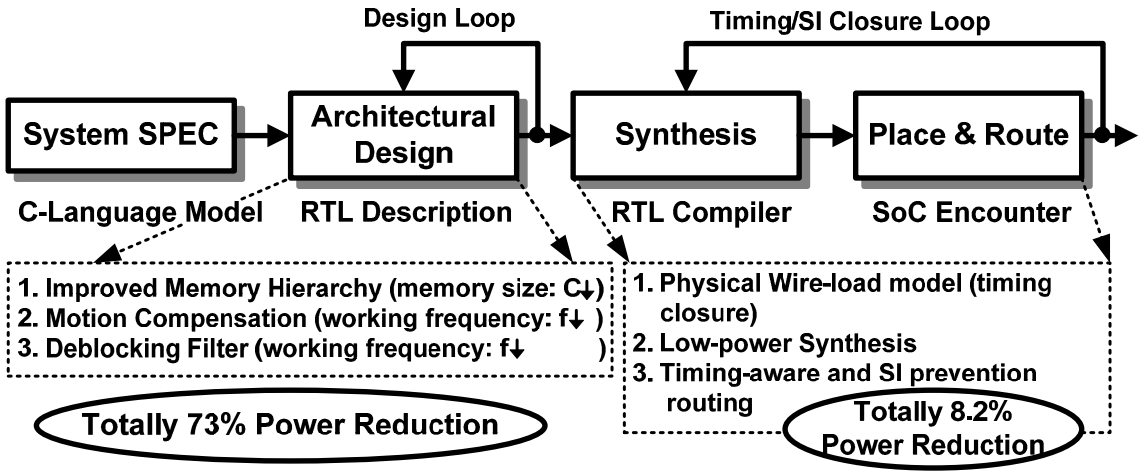

(b)

Figure 5.1: A (a) design flow and (b) power reduction of this video decoder.

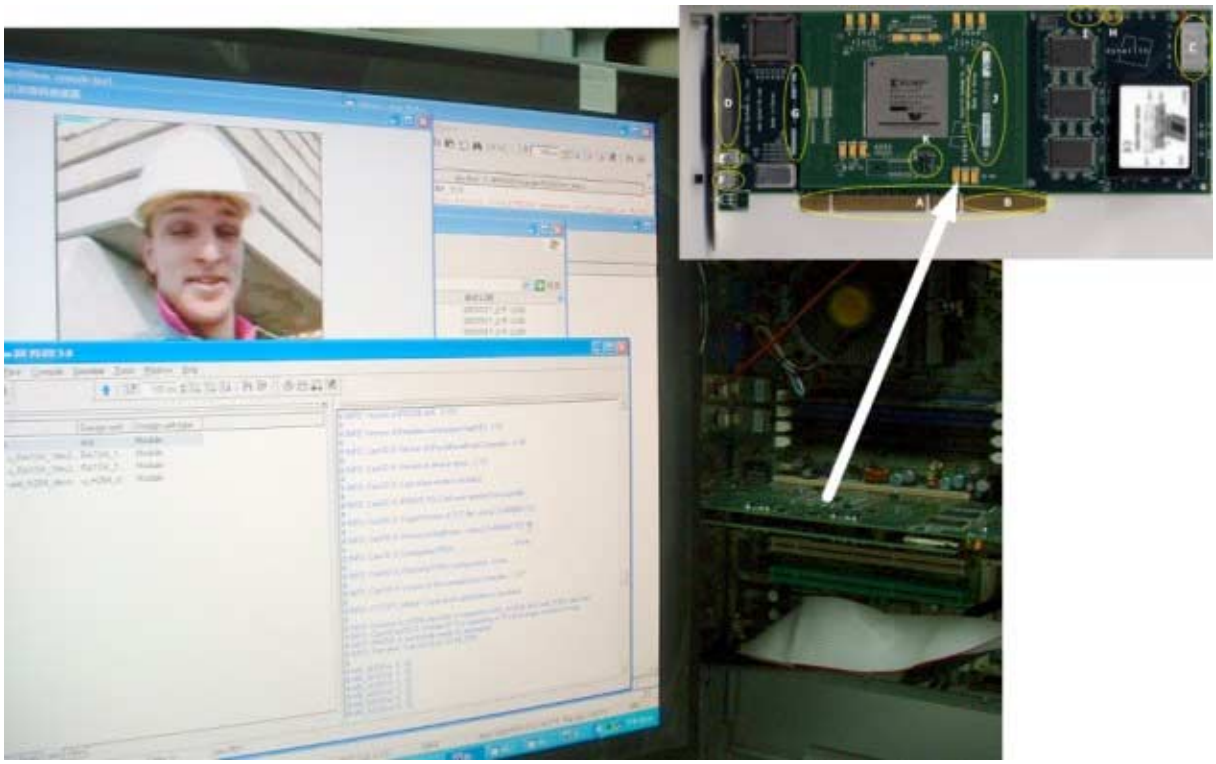

Figure 5.2: Verification Environment.

5.2 Chip Specification

Low power dissipation is always an upmost issue in the design of mobile or handheld devices. To obtain the real metric of power, the aforementioned techniques in Chapter 2 and

3 are fabricated using 0.18μm single-poly six-metal (1P6M) CMOS process with an area of 3.9×3.9 mm^2. Figure 5.3 shows a chip micrograph that combines MPEG-2 SP@ML with H.264/AVC BL@L4 video standards. In particular, the 4×4/8×8 IDCT and in/post-loop deblocking filter are designed to save silicon area. The slice pixel SRAM is allocated to store neighboring pixels, reducing the extensive accesses of external memory as well as I/O power dissipation. The LPL scheme interfaced to the slice pixel SRAM is exploited to further improve the access efficiency since memory accessing contributes a great portion of power dissipation in this video decoding system.

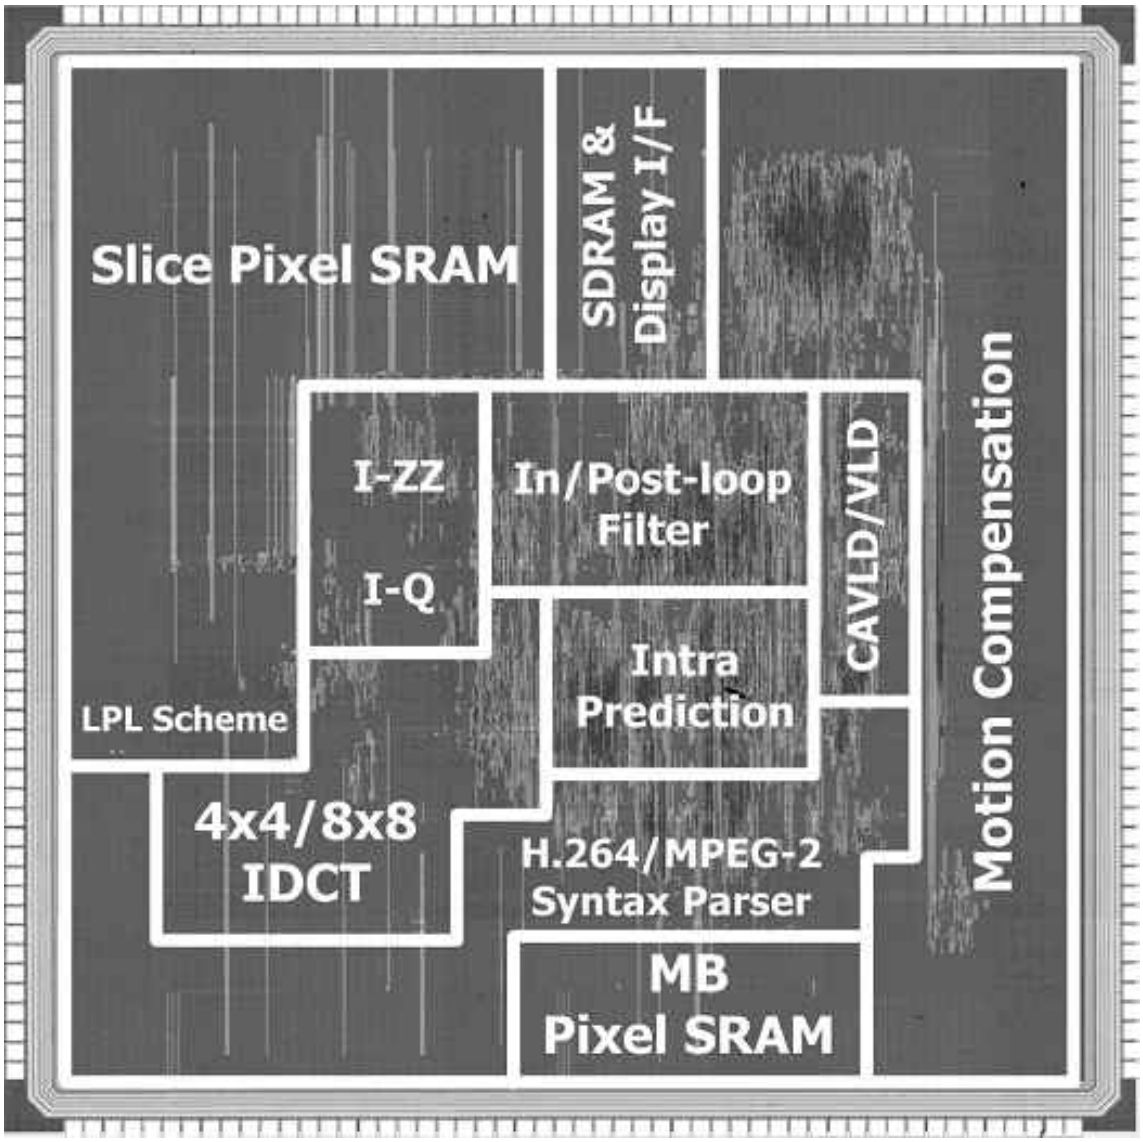

Figure 5.3: Chip micrograph.

Table 5.1: Chip features.

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	SRAM	
External Memory		4MB×2	SDRAM	
Max. System Clock		100MHz		
Max. Processing Throughput		101.04Mpixels/sec		
Core Power Dissipation	**MPEG-2**	QCIF	0.194mW	1.15MHz@15fps
		CIF	4.68mW	4.6MHz@30fps
		SD	15.6mW	16.6MHz@30fps
		720HD	41.76mW	45MHz@30fps
		1080HD	89.46mW	100MHz@30fps
	H.264/AVC	QCIF	0.225mW	1.15MHz@15fps
		CIF	4.86mW	4.6MHz@30fps
		SD	18.54mW	16.6MHz@30fps
		720HD	48.24mW	45MHz@30fps
		1080HD	102.3mW	100MHz@30fps

Chip features are summarized in Table 5.1. Specifically, the 207-pin chip is packaged in the 208-pin CQFP package, where 117 pins are signal pins and others are power pins. The logic gate counts are about 300K excluding the memory. This chip contains 22.75Kb SRAM and adopts two 4MB SDRAM modules for further system integration. The 22.75Kb embedded SRAM occupies a relatively large area because this is an experimental chip in the preliminary design phase. Its size can be further reduced through circuit-level optimizations. The maximum working frequency of this chip is 100MHz and achieves 101MPixels/sec of maximum throughput rates that meet the decoding requirements of high-resolution video sequences (1080HD, 1920×1088 pixels/frame at 30 frames/sec as well as 4:2:0 chrominance formats). As for the core power measurement, the measured accuracy in the VLSI tester is "$\pm 10\mu A \pm 0.1\%$". The associated core power dissipation of high-definition video decoding is 89.46mW and 102.3mW in MPEG-2 and H.264/AVC video standards, respectively. Because low-resolution video formats are also supported through changing the working frequency, the required frequency of standard definition (SD), common intermediate format (CIF), and quarter CIF (QCIF) is 16.6MHz, 4.6MHz, and 1.15MHz, respectively. This fairly low operating frequency is an indication of the improved memory hierarchy and processing cycle reduction discussed in Chapter 3.

5.2.1 Supply Voltage Scaling

Scaling the power supply voltage V_{DD} is the most effective way to reduce the power dissipation [109][110] because dynamic power dissipation component in digital circuits is proportional to the square of the supply voltage. In our design, the aforementioned power can be further reduced by lowering V_{DD} without lowering the V_{TH} of this chip, where the speed requirement is much lower than the critical paths and therefore the circuit can be slower. Although the other ways to reduce power dissipation by lowering V_{DD} are also

presented such as clustered voltage scaling (CVS [111]) and dynamic voltage/frequency scaling (DVFS [112]) schemes, they are not adopted in this dissertation due to the preliminary design of this chip. Under the H.264/AVC decoding mode, the chip is functional over a wide range of frequencies and power supplies, as indicated by the shmoo plot of Figure 5.4. It shows that this chip can operate at a working frequency of 1.15MHz and 16.6MHz with a supply voltage of 1-V and 1.2-V, respectively. As a result, a set of well-known QCIFs, which correspond to the spatial resolution of 176-pixels by 144-lines, are used. Its power dissipation on MPEG-2 and H.264/AVC is only sub-mW and requires 108μW and 125μW at 1-V supply voltage, respectively. Likewise, a supply voltage scaling can be applied to 1.2-V in D1 resolution of 30fps as well. Altogether, the proposed design offers a low-power VLSI solution and is applicable to mobile multimedia systems.

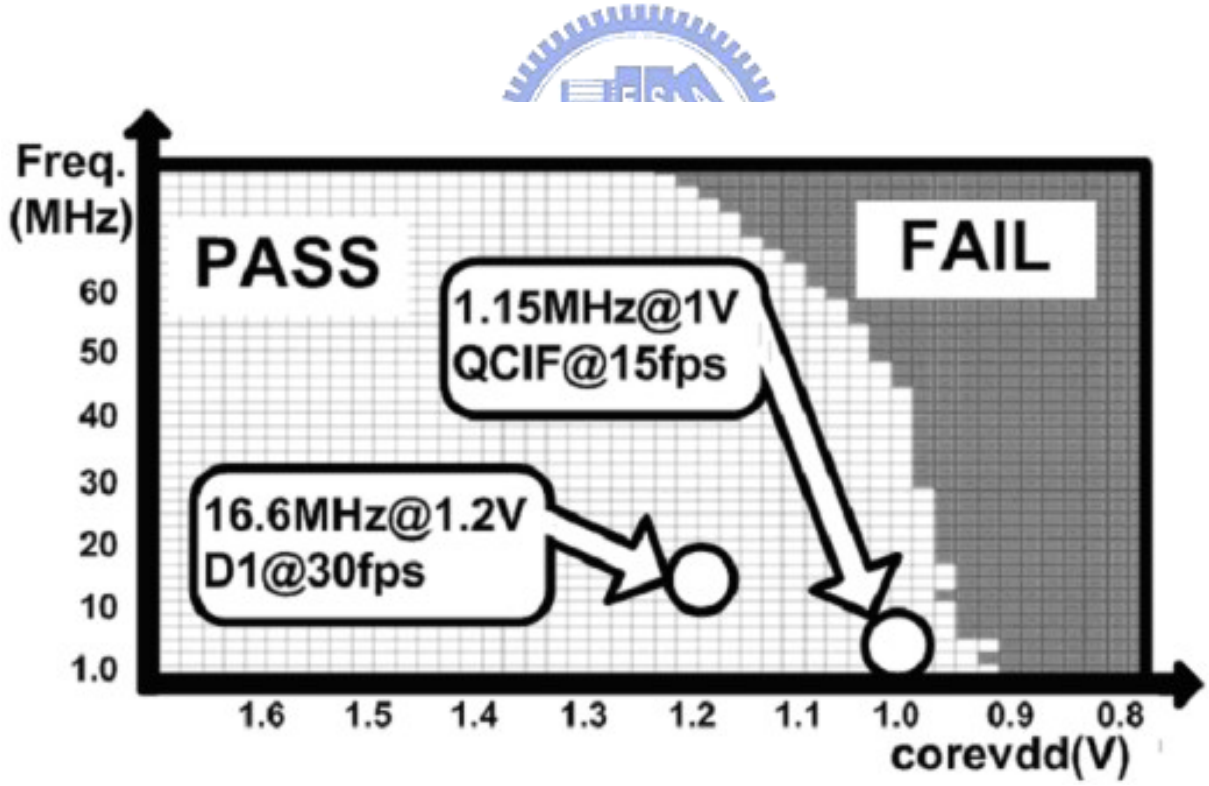

Figure 5.4: Shmoo plot.

5.3 Comparison with Related Works

In terms of core power measurements, a sub-mW of power dissipation can be achieved under decoding sequences of QCIF resolution and 15fps for mobile applications. Because DRAM configurations are so diverse in existing designs and DRAM power can be optimized through other leading-edge techniques [113]–[115], we only show core power dissipation to make a feasible comparison. Figure 5.5 shows a measured power-throughput curve. This plot represents characteristics of video decoding capability, where bottom-right side of this figure indicates better system performance. The power dissipation of this chip is 89.46mW and 102.3mW for the real-time decoding of high-definition video quality in MPEG-2 and H.264/AVC video standards respectively. When we consider the mobile applications, the power consumption is only sub-mW for the real-time decoding of QCIF resolution and 15fps. Therefore, this chip operates at a power-level that is about one order of magnitude less than comparable decoders [42][116].

Considering H.264/AVC video decoding, Table 5.2 exhibits the detailed comparison with existing solutions. However, the specifications of existing designs are so diverse that it's hard to compare power dissipation apple-to-apple. We choose two designs to make a detailed comparison in different performance indexes. Actually, we first make a comparison with Kang *et al.* [42] since both proposed and Kang's [42] designs have similar design characteristics. As for Fujiyoshi *et al.* [112], the functionality of this chip includes not only H.264/AVC video decoding but also MPEG-4 AAC audio decoding. That's why we didn't list this design in Figure 5.5. Although DRAM power is generally larger than the core power consumption, the power reduction on DRAM is not the main focus of this dissertation. In particular, the proposed three-level memory hierarchy and LPL scheme improve the access efficiency and this design can be applied to different kinds of DRAM configurations. In other words, the proposed design can achieve low power dissipation on not only core level

but also external DRAM when applying the existing low-power DRAM modules [113]–[115].

Table 5.2: A detailed comparison with other leading-edge approaches.

	Fujiyoshi *et al.* [112]	**Kang *et al.* [42]**	**Proposed**
Technology	90nm CMOS, 3P6M, 16Mbx2 Embedded DRAM	0.13um CMOS, HLM	0.18um CMOS, 1P6M
Specification	QVGA@15fps	HD@30fps	HD@30fps
Profile	H.264 BL@L1.2 MPEG-4 AAC	H.264 BL	H.264 BL@L4
Gate Count	3,000K(embedded DRAM)	910K (multi-standard)	303.78K
Internal Memory	N/A	N/A	22.75Kb
Frame Memory	32Mb Embedded DRAM	External SDRAM	8MB External SDRAM
Max. Clock Rate	180MHz	130MHz	100MHz
CORE Power	63mW@ 320x240,30fps	159mW@ 352x288,30fps	4.86mW@ 352x288,30fps
DRAM Power		N/A	79.3mW@ 352x288,30fps

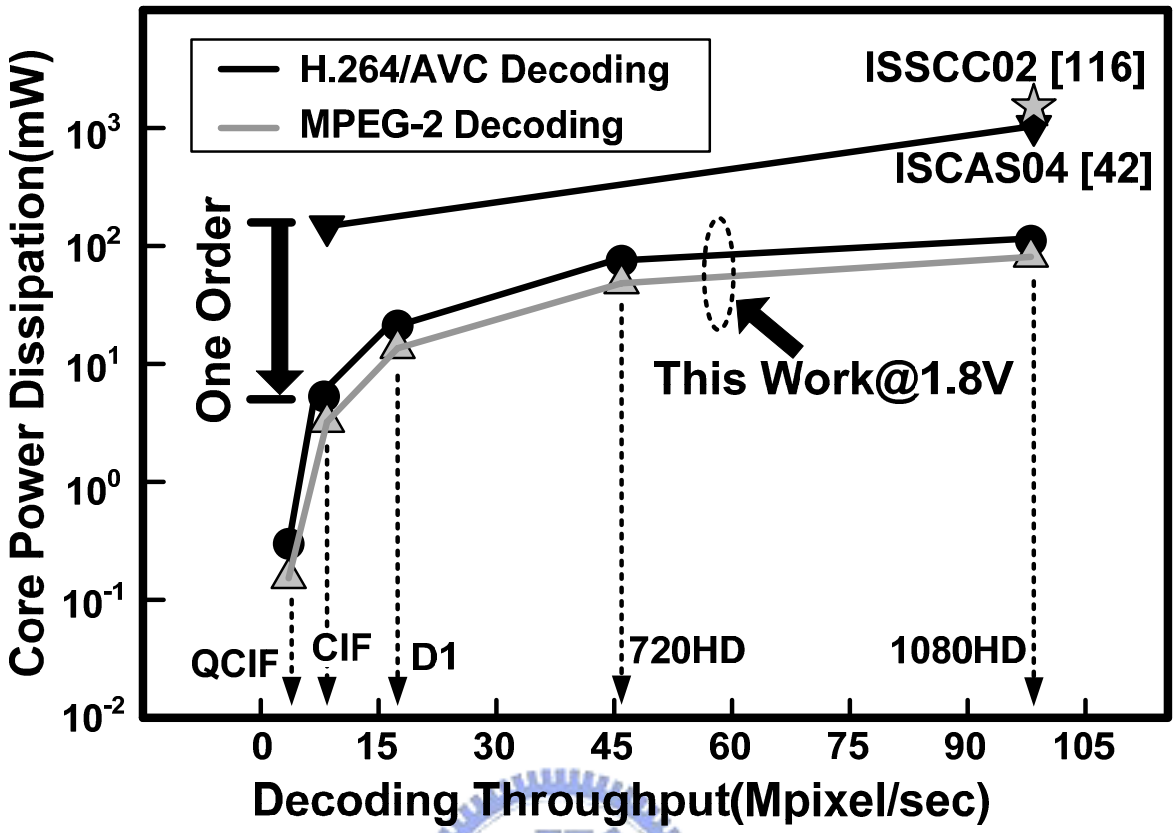

Figure 5.5: Power dissipation comparison.