FPGA-Based Conformance Testing and System Prototyping of an MPEG-4 SA-DCT Hardware Accelerator

Andrew Kinane, Alan Casey, Valentin Muresan, Noel O'Connor
Centre for Digital Video Processing, Dublin City University, Dublin 9, Ireland.
E-Mail: kinanea@eeng.dcu.ie

Abstract

Two FPGA implementations of a Shape Adaptive Discrete Cosine Transform (SA-DCT) accelerator are presented in this paper: one PCI-based and the other AMBA-based. The former is used for conformance testing with the MPEG-4 standard requirements. The latter is an alternative platform for system prototyping and has an architecture more representative of a mobile device. The proposed accelerator meets real time constraints on both platforms with a gate count of approximately 40k, and outperforms the optimised reference software implementation by 20x. It is estimated that the accelerator consumes 250mW on a Virtex-E FPGA and 79mW on a Virtex-II FPGA in the worst case scenario.

1. Introduction

MPEG-4 uses the SA-DCT to support object-based video texture encoding, which in turn allows object manipulation as well as giving improved compression efficiency [1]. The SA-DCT is more complex compared to the 8x8 DCT in terms of hardware implementation due to the wider range of basis functions and extra data re-alignment steps.

2. SA-DCT IP Core

The SA-DCT has been implemented using an adder-based distributed arithmetic datapath that computes coefficients serially (NAND gate count equivalent of 12028). The datapath avoids power hungry multipliers and is configured based on the shape information. The adder network exploits common sub-expression sharing to limit area. Additional power-aware features include guarded evaluation, low switching data alignment and local clock gating (achieved on the FPGA by leveraging the Synplicity Pro "Fixed-Gated Clocks" feature). Further detail on the SA-DCT architecture may be found in [2].

3. MPEG-4 Part 9 Conformance Testing

The MPEG-4 reference hardware initiative ("Part 9") is a working group dedicated to proposed VLSI architectures for the most computationally demanding tools in the standard. The MPEG-4 reference software was compiled with the SA-DCT software replaced by an API call to the SA-DCT hardware accelerator residing on an FPGA (Annapolis WildCard-II PCMCIA card with Xilinx Virtex-II). End to end conformance has verified that the encoded bitstreams with and without SA-DCT hardware acceleration are identical. The test vectors used were 39 of the CIF and QCIF object-based test sequences as defined by the MPEG-4 Video Verification Model. Synthesis results for the WildCard-II platform are shown in Table 1. The Part 9 framework takes up approximately 11% of the FPGA resources leaving almost 90% for IP cores, and the SA-DCT along with its wrapper require just under 20% (Fig. 1).

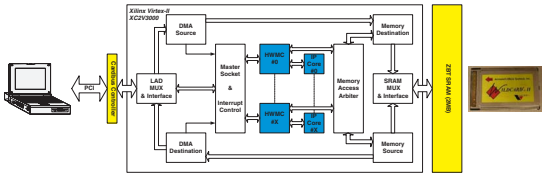

Figure 1. MPEG-4 Part 9 Framework.

CIF resolution at 30fps requires 17820 macroblocks to be processed per second. Motion Estimation (ME) is the most demanding algorithm in MPEG-4 video (depending on search strategy) and a hardware acceleration module for ME proposed in MPEG-4 Part 9 is capable of processing 70k macroblocks per second [3]. This implies that the SA-DCT should be capable of processing a single 8x8 block in approximately 3.57μs. Given that the worst-case number of cycles for the IP core to process a block is 142 cycles, the IP core must run at approximately 40MHz at worst to maintain real-time constraints. The post place and route timing analysis indicates a theoretical operating frequency of 62.9MHz so the IP core is able to handle real time processing of CIF sequences quite comfortably.

Target	Module	Area [Gates]	Max. Freq. [MHz]	Power [mW]	Throughput [MB/s]	*	CLB Slices	Block RAMs
WildCard-II	SA-DCT	39972	62.9	79	42.25	0	2630 (18%)	0
	Wrapper	4354	85.6	n/a	n/a	0	201 (1%)	0
	Part 9	102085	77.6	n/a	n/a	0	1627 (11%)	1
Integrator/CP	SA-DCT	40152	48.9	250	33.06	n/a	2647 (13%)	0
	Wrapper	6053	81.0	n/a	n/a	n/a	283 (1%)	0
	ARM VS	7020	89.7	n/a	n/a	n/a	381 (1%)	0

Table 1. FPGA Synthesis Results

4. System Prototyping

The WildCard-II platform does not represent a realistic architecture for a mobile embedded system (PCI bus more suitable for emulating hardware accelerators for desktop PC graphics cards). For embedded systems, the processor of choice is the ARM family, so we propose a plug and play "ARM Virtual Socket" (VS) prototyping platform built around an ARM processor and the AMBA bus architecture. We have implemented the ARM virtual socket on an ARM Integrator/CP prototyping platform with an ARM920T processor running embedded Linux and a Xilinx Virtex-E FPGA for AMBA IP core prototyping. The platform facilitates the rapid prototyping of any number of "virtual component" hardware accelerators with AMBA interfaces.

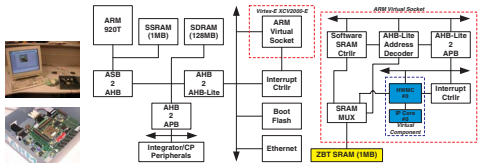

Figure 2. ARM Virtual Socket Platform.

The synthesis results in Table 1 illustrate that the SA-DCT IP core also meets the real time frequency constraint of 40MHz on the ARM virtual socket platform. Profiling the MPEG-4 optimised SA-DCT software implementation running on the ARM Integrator/CP versus the proposed SA-DCT hardware accelerator on the Virtex-E FPGA shows that the accelerator offers a speed-up of about 20x. The ARM virtual socket platform itself has a much smaller equivalent gate count (7020) compared to the PCI framework currently used by MPEG-4 Part 9 (102085). This leaves more space on the FPGA for prototyping more IP cores together. The wrapper nature of both platforms make it straightforward to migrate an IP core between platforms since it shields it from platform specific protocols, thus we propose the ARM virtual socket as an alternative platform for system prototyping.

5. Power Consumption Profiling

Comparing hardware accelerators meaningfully in terms of their power consumption properties is difficult to achieve in practice. To ensure a fair comparison, competing architectures must be compared with the same target technology. Also since the switching activity in a module is dependent on its data load, the same testbench should be used for fair comparisons. In this work we use back-annotated dynamic simulation of the post place and route netlist to analyse the power consumption of the SA-DCT IP core. The Xilinx XPower tool was used to analyse the annotated switching information from a VCD file. Since the SA-DCT core consumption is highly data dependent, a simulation was run with 1000 random blocks and the reported average power was 250mW for the Integrator/CP platform. This can be interpreted as a worst case power estimate since the SA-DCT core is stimulated the entire time with random boundary blocks. In regular video sequences, there is a lot of spatial redundancy in the shape information and boundary blocks do not occur as often. The same random data testbench reports an average power of 79mW for the WildCard-II platform.

6. References

[1] T. Sikora and B. Makai, "Shape-Adaptive DCT for Generic Coding of Video," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 5, no. 1, pp. 59–62, Feb. 1995.

[2] A. Kinane, V. Muresan, and N. O'Connor, "An Optimal Adder-Based Hardware Architecture for the DCT/SA-DCT," in *Proc. SPIE Video Communications and Image Processing (VCIP)*, Beijing, China, July 12–15, 2005.

[3] *Text of ISO/IEC TR 14496-9 Information technology - Coding of audio visual objects - Part 9: Reference hardware description*, ISO/IEC Std., Rev. 2, 2005.