Int. ASIC/SOC Conference, pp. 203-207, September 2001, Arlington, Virginia

ON THE LOW-POWER DESIGN OF DCT AND IDCT FOR LOW BIT-RATE VIDEO CODECS

Nathaniel August

Intel Corporation
M/S JF3-410
2111 N.E. 25th Avenue
Hillsboro, OR 97124
E-mail: nathaniel.j.august@intel.com

Dong Sam Ha

Virginia Tech VLSI for Telecommunications (VTVT) Lab
Bradley Dept. of Electrical and Computer Eng.
Virginia Tech, Blacksburg VA 24061
E-mail: ha@vt.edu Web: www.ee.vt.edu/ha

ABSTRACT

This paper examines low power design techniques for discrete cosine transform (DCT) and inverse discrete cosine transform (IDCT) circuits applicable for low bit rate wireless video systems. The techniques include skipping low energy DCT input, skipping all-zero IDCT input, low precision constant multipliers, clock gating, and a low transition data path. The final DCT and IDCT blocks reduce average power dissipation by 95% over baseline reference blocks.

I. INTRODUCTION

Low bit rate wireless video systems have applications in cellular videophones, surveillance systems, and mobile patrols. The ITU-T H.263 [1] video codec standard is suitable for low bit-rate wireless video systems. A critical requirement for portable wireless video systems is low power dissipation.

Fig. 1 shows a block diagram of an H.263 coder and decoder. The motion estimation and compensation block (MEC) and the inverse motion estimation and compensation block (MEC^{-1}) perform temporal compression and decompression. The DCT (IDCT) block performs spatial compression (decompression) of the data. The quantization (Quant) block sets insignificant DCT coefficients to zero. The variable length coder (VLC) performs run length coding, which compresses long runs of DCT coefficients.

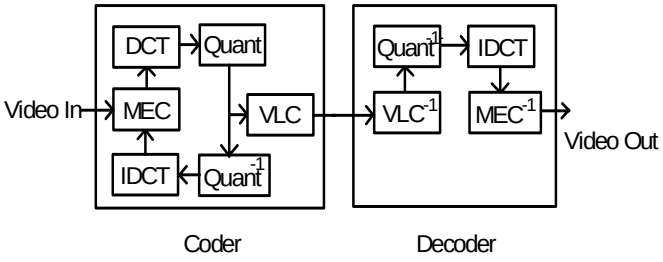

Figure 1: H.263 Coder and Decoder Block Diagram

The DCT and IDCT, as well as the motion estimation, are computationally intensive in H.263. The combined computational complexity of the DCT and IDCT in the coder surpasses that of any other unit, consuming 21% of the total computations [2]. The IDCT in the decoder also incurs the largest computational cost. The high computational complexity of the DCT and IDCT leads to high power dissipation of the blocks, so low power design of DCT and IDCT blocks are essential in portable wireless video systems.

II. PREMLIMINARIES

In this section, we explain terms necessary to understand the paper and review low-power designs for DCT/IDCT briefly.

A. Terms

In H.263, a **macroblock** is a basic unit of data comprised of six 8x8 blocks that represent a 16x16 pixel area of a video frame. Macroblocks conveniently represent data in YCbCr format, which contains a luminance component (Y), a blue chrominance component (Cb), and a red chrominance component (Cr). A macroblock contains four Y blocks, one Cb block, and one Cr block.

In motion estimation, the **sum of absolute differences** (SAD) is a measure of how well a macroblock in the current frame matches a nearby macroblock in the previous frame.

$$SAD = \sum_{k=1}^{16} \sum_{l=1}^{16} \left| Y_{i,j}(k,l) - Y_{i-u,i-v}(k,l) \right| \qquad (1)$$

In (1), the magnitude of each luminance sample, $Y_{i,-u,j-v}(k,l)$, in a candidate macroblock that is offset by (u,v) in the previous frame is subtracted from magnitude of each luminance pixel, $Y_{i,j}(k,l)$, in the current (i.e., reference) macroblock at position (i,j). The candidate macroblock with the lowest SAD is the most likely match for the current macroblock.

The average **peak signal to noise ratio** (PSNR) of a frame in a video sequence is measured as

$$PSNR = 10 \log \frac{1}{x \cdot y} \sum_{rows} \sum_{cols} \frac{255^2}{Y_1 - Y_2}, \qquad (2)$$

where x is the number of rows, y is the number of columns, and Y_1 and Y_2 are the luminance values in the original and

the reconstructed pictures. We use the average PSNR over all frames as a quantitative measure of the quality of a video sequence.

The **coded block pattern** (CBP) is a part of the H.263 bit stream that describes the values of DCT coefficients in both luminance and chrominance blocks. The CBP of a block is set to '1' for a block with non-zero DCT coefficients. A block produces all-zero input to the IDCT if its DC coefficient is zero and its CBP = '0'. Additionally, if the **coded macroblock indication** (COD) bit is set, it signals the decoder to treat the entire macroblock as all-zero data.

The two dimensional (2-D) **DCT** in Equation (3) transforms an 8x8 block of picture samples *x(m,n)*, into spatial frequency components *Y(k,l)* for $0 \le k,l \le 7$. The **IDCT** in Equation (4) performs the inverse transform for $0 \le m,n \le 7$. In Equations (3) and (4), $\alpha(0)=1/\sqrt{2}$ and $\alpha(j)=1$, $j \ne 0$.

$$Y(k,l)=\frac{1}{4}\alpha(k)\alpha(l)\sum_{n=0}^{7}\sum_{m=0}^{7}x(m,n)\cos\left(\frac{(2m+1)\pi k}{16}\right)\cos\left(\frac{(2n+1)\pi l}{16}\right) \quad (3)$$

$$x(m,n)=\frac{1}{4}\alpha(k)\alpha(l)\sum_{k=0}^{7}\sum_{l=0}^{7}Y(k,l)\cos\left(\frac{(2m+1)\pi k}{16}\right)\cos\left(\frac{(2n+1)\pi l}{16}\right) \quad (4)$$

B. Review of Low-Power Designs for DCT/IDCT

Some common low power techniques include clock gating, pipelining, and voltage scaling [3]. Parallel architectures save power by reducing the supply voltage [4] or the clock speed [5]. Low power libraries also reduce power [4][6].

Other architectures save power by reducing calculations on visually irrelevant DCT coefficients. Xanthopoulos and Chandrakasan employ arithmetic units in which the precision changes adaptively depending on the visual significance of the data [7]. Another architecture allows a fine (1-bit increments) resolution precision control in the arithmetic units [5]. The amount of precision can be determined from the peak-to-peak pixel difference [5] or from the quantization parameter [8]. Li and Lu propose skipping the computation of visually insignificant high frequency DCT coefficients altogether [6]. Their method removes the circuit elements that compute these coefficients and sets these coefficients to zero.

By ignoring redundant sign bits, arithmetic units save power in the DCT circuit. The redundant sign bits occur frequently because only the small difference between the previous and current frame is sent to the DCT. One architecture reduces power for small coefficients by deactivating adder partitions that work on redundant sign bits [5]. Another architecture ignores the most significant bits of the inputs to arithmetic units if they are common to both inputs [7].

Since the majority of input data for the IDCT is comprised of zero-valued coefficients, significant power reduction can come from disabling adders and multipliers for zero-valued operands [5][7][8].

Another popular target for power reduction in DCT/IDCT blocks is the method of implementing multipliers. The two most popular implementations for multipliers are bit-serial (distributed architecture) and bit-parallel. Bit-serial architectures dissipate power due to the high frequency, serialized operation, and high capacitance of the ROM address and bit lines [9]. However, they can easily and finely partition data for variable precision arithmetic. Several parallel multipliers, which compromise speed, area and power dissipation, have been proposed [5][6][10]-[14]. An array multiplier is a straightforward implementation of the bit-parallel architecture that is easily implemented from library cells [10]. Since one multiplicand is known *a priori*, efficient shift-and-add multipliers are also used [5][6][13].

III. BASELINE DCT/IDCT

Since straightforward implementations of equations (3) and (4) are computationally expensive, we use a more efficient approach in hardware to our baseline design. The 8x8, 2-D DCT is an orthogonal transformation process, so it can be performed with eight 1-D DCTs on the input rows followed by eight 1-D DCTs on the columns of the transformed input rows. For the 1-D DCT/IDCT, we employ Chen's algorithm, which requires only 16 multiplications for a 1-D DCT/IDCT [15].

The architecture for our DCT/IDCT blocks is shown in Fig 2. The controller receives the first row of data (DIN) through the serial-to-parallel unit under SEN signal enabled. It then performs a 1-D DCT/IDCT on the first row with the SEL and REN signals determining the data path. The results are stored in the first row of the transposition memory under ROWACK enabled. This process repeats for the remaining seven rows of the input block. Next, the ISEL and COLACK signals enable the 1-D DCT/IDCT unit to receive input from the columns of the transposition memory. The results (DOUT) are available through the parallel-to-serial unit under PEN enabled.

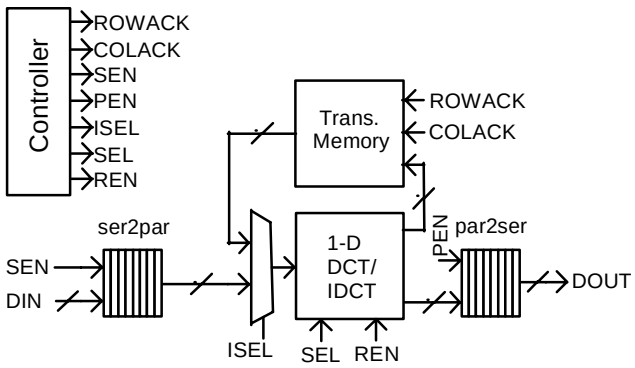

Figure 2: 2-D DCT/IDCT Architecture

We considered bit-serial and bit-parallel approaches for our multipliers. The bit-serial approach is more area efficient and potentially more power efficient. Parallel array multipliers are fast and consume just 5% of the total power

in our baseline design. Since they are also available in design libraries, we use array multipliers in the baseline architecture.

IV. PROPOSED LOW POWER DCT/IDCT DESIGN

We investigated five low-power improvements for the above baseline DCT/IDCT blocks in an H.263 system. Since some methods would degrade the picture quality, we measured the reduction of the PSNR for those methods. A method is employed only if the degradation of the picture quality is unnoticeable to human eyes and the degradation of the PSNR is small. For the PSNR measurement and the visual examination, we implemented a prototype H.263 codec in the C language. The prototype was also used to collect test data for three video clips: Claire, Miss America, and Foreman.

A. Skipping Macroblocks for the DCT block

Many input macroblocks, after motion compensation, contain little new information; this results in small magnitude DCT coefficients that will likely be quantized to zero. The DCT (and the quantization and VLC operations) can be skipped for these macroblocks, and all DCT coefficients are simply set to zero as an approximation. In fact, this method was suggested to speed up the DCT operations in [18], and we propose employment of the method to save power by disabling the entire 2-D DCT unit for these macroblocks.

One method of predicting such macroblocks is to examine the SAD value of incoming macroblocks [18]. The SAD value provides a good measure of the energy of the incoming pixels and is readily available from the motion estimation block. Since high quantization parameters also force more coefficients to zero, we skip a macroblock if the following condition is met: SAD < 128*QUANT.

B. Skipping Input Data Blocks for the IDCT block

Because many output DCT coefficients are quantized to zero, many IDCT input blocks and macroblocks will contain all-zero data. Since an IDCT results in all-zero output for all-zero input, the 2-D IDCT unit can be disabled for an all-zero input data block. An input block is all zero if the CBP field for the block is zero and the DC coefficient is '0'. An entire macroblock will be zero if the COD bit is '0'. The decoder can easily extract these parameters from the H.263 bit-stream, while the encoder can receive a signal from the VLC. Note that this method does not degrade the picture quality.

C. Gated Registers for the DCT/IDCT block

Four units (the two I/O units, the 1-D DCT/IDCT units, and the transposition memory) of the baseline DCT/IDCT blocks in Fig. 2 contain over 99% of the flip-flops. Since these registers comprise the majority of power dissipation in the circuit, their clocks can be gated to save power during periods of inactivity.

In the transposition memory, the registers need to be updated only at the end of each of the first eight 1-D DCT/IDCT operations, so the transposition registers require a clock in less than 3% of the cycles. Additionally, the registers in the baseline DCT/IDCT blocks are active for only about 35% of the entire 2-D DCT/IDCT operation. Finally, I/O registers in the ser2par and par2ser units require clocks only during the input and output operations. Overall, the I/O registers need clocks for just 16% of the overall operation. For the remaining times, the registers in each of these units can be gated since they are storing values.

D. Constant Shift-Add Multipliers with Reduced Precision

Array multipliers are currently employed for the baseline unit, since they contribute only 5% of total power dissipation. However, if a large power reduction is achieved through the above proposed methods, array multipliers may dissipate a more significant proportion of power for the circuit. Hence, it is desirable to employ more power efficient multipliers at the cost of higher circuit complexity. Shift-and-add multipliers have been reported as more efficient in power dissipation than array multipliers [5][6][13].

A major concern for implementing the multipliers is the precision of the coefficients. The precision of the coefficients for the baseline DCT/IDCT blocks shown in Fig. 2 is 14 bits, which has previously been used to balance accuracy and area [10]. Since the quantization parameter - and hence the quantization noise - is usually large for low bit rates, the DCT and IDCT may require less than 14 bits of precision. So we propose a reduction of precision to reduce power dissipation. To preserve the video quality, a suitable width is 8 bits for DCT and IDCT coefficients. If the IDCT units in the encoder and decoder share the same architecture, the video will remain stable since the reference frame will be the same in the encoder and the decoder.

E. Low-Transition Data Paths

The registers and arithmetic units in the 1-D DCT/IDCT blocks select an input from multiple sources through a bank of multiplexers. In the baseline unit, the select input (SEL in Fig. 2) is in a "don't care" state while a register or arithmetic unit is inactive. For low power design, it is desirable that each SEL signal remains unchanged from its previous value until the corresponding register or arithmetic unit is active. This reduces power dissipation because fewer transitions will occur on the inputs of the registers and arithmetic units. Note that this method produces no reduction in PSNR, and it increases the complexity of the control unit slightly.

V. EXPERIMENTAL RESULTS

Using Synopsys Design Power, we estimated power dissipation at the gate level for three test video clips: Claire, Miss America, and Foreman. The three video clips are in QCIF format, which contains 176x144 pixels per frame. We tested each method independently on the baseline units for the DCT and the IDCT and measured the improvement in power dissipation. After measuring the effectiveness of each

method, we applied all the methods to the baseline units, one at a time, starting with the most effective method.

Table I and Table II show the effects of each method when independently added to the baseline units.

From Table I, the most efficient method for the DCT is the skipping of macroblocks with low SAD parameters, which reduces power dissipation by an average of 65.6%. This is attributable to the fact that 57.8% to 79.8% of macroblocks are skipped. There is less of a power reduction for the "Foreman" sequence because its high amount of motion causes higher SAD values; hence, fewer blocks are skipped.

From Table II, the most efficient method for IDCT is the skipping of all-zero input. The power savings is larger than the DCT, because the decision to skip is made at the block level, not at the macroblock level. The power savings from skipping all-zero input to the IDCT is similar for all three video clips. It is explained as we set a constant target bit rate for each video clip. With equal target bit rates, high motion video sequences such as "Foreman" have higher quantization parameters than low motion video sequences; hence, just as many coefficients are forced to zero in high motion video sequences as in low motion video sequences.

Table I
Efficiency Of Proposed Low Power Methods For The Dct When Applied Individually

Configuration	Claire		Foreman		Miss Am.	
	Power (mW)	Percent Saved	Power (mW)	Percent Saved	Power (mW)	Percent Saved
Baseline	9.25		9.39		9.32	
Skip Low SAD Blocks	2.24	75.8%	4.73	49.6%	2.66	71.5%
Gated Registers						
Trans. Memory	3.96	57.2%	4.13	56.0%	4.07	56.3%
1-D DCT	7.91	14.5%	8.18	12.9%	8.11	13.0%
I/O Registers	8.23	11.0%	8.48	9.69%	8.43	9.55%
Constant Multipliers	8.18	11.6%	8.46	9.90%	8.41	9.76%
Data Path	8.38	9.4%	8.63	8.09%	8.57	8.05%

Table II
Efficiency Of Proposed Low Power Methods For The Idct When Applied Individually

Configuration	Claire		Foreman		Miss Am.	
	Power (mW)	Percent Saved	Power (mW)	Percent Saved	Power (mW)	Percent Saved
Baseline	8.53		8.49		8.50	
Skip All-Zero Blocks	1.23	85.6%	1.41	83.4%	1.40	83.5%
Gated Registers						
Trans. Memory	3.55	58.4%	3.29	61.2%	3.53	58.5%
1-D DCT	7.41	13.1%	7.37	13.2%	7.38	13.2%
I/O Registers	7.55	11.5%	7.56	11.0%	7.54	11.3%
Constant Multipliers	7.64	10.4%	7.64	10.0%	7.59	10.7%
Data Path	7.68	9.96%	7.71	9.19%	7.66	9.88%

The second most efficient method for both the DCT and the IDCT is the gating of registers. Registers account for over 85% of the power in the DCT and IDCT units. Since registers need only be enabled when they have meaningful input, gating registers saves a significant amount of power. The greatest power savings from clock gating comes from the transposition memory. The transposition memory accounts for about 70% of the total amount of flip-flops in the DCT or IDCT circuit, and it needs to be enabled for only 3% of the time. The registers in the 1-D unit and in the I/O unit are active for a larger proportion of time than the transposition memory, and they each account for approximately the same amount of the remaining flip flops.

For both the DCT and the IDCT, smaller improvements in power dissipation result from the low power multipliers and the low transition data path.

Table III
Efficiency Of Proposed Low Power Methods For The Dct When Added Incrementally

Configuration	Claire		Foreman		Miss Am.	
	Power (mW)	Percent Saved	Power (mW)	Percent Saved	Power (mW)	Percent Saved
Baseline	9.25		9.39		9.32	
Skip Low SAD Blocks	2.24	75.8%	4.73	49.6%	2.66	71.5%
Gated Registers						
Trans. Memory	1.64	82.3%	2.89	69.2%	1.87	79.9%
1-D DCT	1.16	87.5%	1.87	80.1%	1.26	86.5%
I/O Registers	0.687	92.6%	1.23	86.9%	0.776	91.7%
Constant Multipliers	0.425	95.4%	0.639	93.2%	0.470	95.0%
Data Path	0.314	96.6%	0.501	94.7%	0.354	96.2%
Total Power	**0.314**	**96.6%**	**0.501**	**94.7%**	**0.354**	**96.2%**

Table IV
Efficiency Of Proposed Low Power Methods For The Idct When Added Incrementally

Configuration	Claire		Foreman		Miss Am.	
	Power (mW)	Percent Saved	Power (mW)	Percent Saved	Power (mW)	Percent Saved
Baseline	8.53		8.49		8.50	
Skip All-Zero Blocks	1.23	85.6%	1.41	83.4%	1.40	83.5%
Gated Registers						
Trans. Memory	0.776	90.9%	0.914	89.2%	0.910	89.3%
1-D DCT	0.377	95.6%	0.442	94.8%	0.440	94.8%
I/O Registers	0.240	97.2%	0.296	96.5%	0.297	96.5%
Constant Multipliers	0.161	98.1%	0.218	97.4%	0.226	97.3%
Data Path	0.150	98.2%	0.189	97.8%	0.187	97.8%
Total Power	**0.150**	**98.2%**	**0.189**	**97.8%**	**0.187**	**97.8%**

The five methods described above were then added to the baseline unit in order from the most efficient method to the least efficient method. For the DCT and IDCT blocks, Table III and Table IV show the experimental results under the employment of the methods. The "Power" and "Percent Saved" measurements apply to a method combined with all the methods above it.

We can observe from Tables III and IV that each power savings technique has an impact, even after others are added. For both the DCT and the IDCT, the combination of power savings methods impact the overall power far more than any single technique. The average power reduction of the proposed methods for DCT and IDCT is 95.8% and 97.9%, respectively.

Table V
Psnr For Proposed Low Power Methods

Configuration	PSNR (dB) Claire	PSNR (dB) Foreman	PSNR (dB) Miss Am.
Baseline	40.53	29.64	40.52
Lower Precision Multipliers	39.19	29.33	39.46
Skip Low SAD Blocks	39.41	29.41	40.08
Combined	**38.51**	**29.12**	**39.10**

Since the skipping of DCT macroblocks and the lower precision multipliers degrade picture quality, we show the overall effect on PSNR when both methods are added in Table V. Note that the combined methods degrade PSNR more than individual methods. However, the degradation is insignificant to human eyes.

VI. CONCLUSION

This paper presents four low power techniques for both 8x8 2-D DCT and IDCT blocks, which are intended for low bit rate wireless video applications. The most efficient scheme for the DCT unit is skipping macroblocks; video sequences with low amounts of motion save more power than those with high amounts of motion. However, the remaining three techniques result in similar power savings for all video sequences tested. The average power savings for the four methods combined is 96% for the DCT unit. The most efficient scheme for the IDCT block is skipping all-zero input data blocks, which saves a similar amount of power for all three sequences. The average power savings for the four methods combined is 98% for the IDCT unit.

Finally, it is important to note that our methods can be integrated with some other existing methods such as the adaptive precision method [7] and zero-valued coefficients [5][8] to further reduce power dissipation.

We conclude the paper by summarizing the characteristics for our DCT and IDCT designs in Table VI.

Table VI
Characteristics Of Our Dct And Idct

	DCT	IDCT
Gate count	22,000	22,000
Frequency	2.5 MHz	2.5 MHz
Throughput	10 fps @ QCIF	10 fps @ QCIF
Process Technology	.35µm	.35µm
Average Power	.390 mW	.175 mW

REFERENCES

[1] ITU-T Recommendation H.263 "Video coding for low-bitrate communication," Draft International Standard, Feb. 1998.

[2] B. Erol, F. Kossentini, and H. Alnuweiri, "Implementation of a fast H.263+ encoder/decoder," *Conference Record of the Thirty-Second Asilomar Conference on Signals, Systems & Computers, 1998*, vol. 1, pp. 462-466, Sep. 1998.

[3] Xanthopolous and Chandrakasan, "A low-power IDCT macrocell for MPEG-2 MP@ML exploiting data distribution properties for minimal activity," *IEEE Journal of Solid State Circuits*, vol. 34, pp. 693-703, May 1999.

[4] Chen, Jiu, Chang, Lee, and Ku, "A low power 2-D DCT chip design using direct 2-D algorithm," *Proceedings of the ASP-DAC '98. Asia and South Pacific Design Automation Conference 1998*, pp. 145-150, Feb. 1998.

[5] Kim and Beerel, "A high-performance low-power asynchronous matrix vector multiplier for discrete cosine transform," *The First IEEE Asia Pacific Conference on ASICs, 1999*, pp. 135-138, Aug. 1999.

[6] Li and Lu, "Low power design of two-dimensional DCT," *Proceedings of the Ninth Annual IEEE International ASIC Conference and Exhibit, 1996*, pp. 309-312, Sep. 1996.

[7] Xanthopoulos and Chandrakasan, "A low power DCT core using adaptive bitwidth and arithmetic activity exploiting signal correlations and quantization," *IEEE Journal of Solid State Cricuits*, vol. 35, pp. 740-750, May 2000.

[8] E. Scopa, A. Leone, R. Guerrieri, and G. Baccarani, "A 2-D DCT low power architecture for H.261cCoders," *1995 International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 3271-3274, May 1995.

[9] M. Kuhlmann and K. K. Parhi, "Power comparison of flow-graph and distributed arithmetic based DCT architectures," *Conference Record of the Thirty-Second Asilomar Conference on Signals, Systems & Computers, 1998*, vol. 2, pp. 1214-1219, Nov. 1998.kdjflkds

[10] H.B. Kim, "High-Level synthesis and implementation of built-in self-testable data path intensive circuit," PhD Dissertation, EE Dept., Virginia Tech, Dec. 1999

[11] Hsiao, Shiue, and Tseng, "A Cost-Efficient and Fully-Pipelinable Architecture for DCT/IDCT," *IEEE Transactions on Consumer Electronics*, vol. 45, pp. 515-525, Aug. 1999.

[12] Chiang and Huang, "New Architecture for High Throughput-Rate Real-Time 2-D DCT and the VLSI Design," *Proceedings of the Ninth Annual IEEE International ASIC Conference and Exhibit, 1996*, pp. 219-222, Sep. 1996.

[13] Avinindra and Willson, "A 100 MHz 2-D 8x8 DCT/IDCT processor for HDTV applications," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 5, pp. 158-165, Apr. 1995.

[14] J. Chen and K. Liu, "Cost-Effective Low-Power Architectures of Video Coding Systems," *Proceedings of the 1999 IEEE International Symposium on Circuits and Systems*, vol. 1, pp. 153-156, Jun. 1999.

[15] W.H. Chen, C.H. Smith, and S.C. Fralick, "A fast computational algorithm for the discrete cosine transform," *IEEE Trans. Commun.*, Vol. COM-25, pp. 1004-1009, Sept. 1977.

[16] E. Feig and S. Winograd, "Fast algorithms for the discrete cosine transform," *IEEE Transactions on Signal Processing*, vol. 40, pp. 2174-2193, Sep. 1992.

[17] Srinivasan and Liu, "VLSI design of high-speed time-recursive 2-D DCT/IDCT processor for video applications," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 87-96, Feb. 1996.

[18] Pao and Sun, "Modeling DCT coefficients for fast video encoding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 9, pp. 608-616, Jun. 1999.