A Low-Power DCT Core Using Adaptive Bitwidth and Arithmetic Activity Exploiting Signal Correlations and Quantization

Thucydides Xanthopoulos and Anantha Chandrakasan
Massachusetts Institute of Technology, Cambridge, MA 02139

I. Introduction and Background

This work describes the implementation of a DCT (Discrete Cosine Transform) chip targetted to low power video (MPEG2 MP@ML) and still image (JPEG) applications. The chip exhibits two innovative techniques for aritmetic operation reduction in the DCT computation context along with standard voltage scaling techniques such as pipelining and parallelism. The first method exploits the fact that image pixels are typically well correlated and exhibit a certain number of common most significant bits in local areas. These bits constitute a common mode DC offset that only affects the computation of the DC DCT coefficient and is irrelevant for the computation of the higher spectral coefficients. The DCT chip uses adaptive-bitwidth distributed-arithmetic computation units that reject the common most significant bits for all AC coefficient computations, resulting in arithmetic operations with reduced bitwidth operands thus reducing switching activity. We call this method MSB rejection (MSBR).

The second method exploits the fact that in image and video compression applications, DCT is always followed by a quantization step which essentially reduces the precision of the visually insigificant higher frequencies. The DCT chip allows the user to set up to four different classes of precision for each spectral coefficient on a row-by-row basis so that no uneccessary computation is performed if the precision will be lost anyway due to quantization. A row-column peak-to-peak detector classifies each block row and column into one of four classes of computation precision for maximizing image peak SNR (PSNR) and minimizing the number of arithmetic operations. We call this method row-column classification (RCC).

II. Algorithm, Architecture and Circuits

The DCT processor implements a row-column distributed arithmetic version of the Chen fast DCT algorithm enhanced with the activity reduction methods outlined above (MSBR and RCC). The chip consists of two 1D 8-point DCT units and a transposition memory in between (Fig.1). Image data is shifted in serially and undergoes a preliminary "butterfly" stage of additions and subtractions. The result of the 4 additions and 4 subtractions is fed serially in a bank of 8 RACs (ROM and accumulators) that implement dot-product computation without the use of a multiplier. The 16x10 ROMs store linear combinations of constant cosine coefficients A0-A3 (Fig.2). During each clock cycle, the ROM data is added to an accumulator register after its contents have undergone a right shift (multiplication by two.) After a number of cycles equal to the bitwidth of the input data (8 maximum), the 8-point 1D DCT of each image block row has converged within the accumulators of RACs 0 through 7. The second stage repeats the process on each block column.

The control blocks labelled "MSB Rejection" detect common most significant bits among the top 4 pixel sums and sign extension bits in the bottom 4 pixel differences of each stage (Fig.1). Common most significant bits are detected using the bit slices of Fig.3 in a cascaded configuration. To reduce activity, the chip disables the RACs using an appropriate qualifying pulse (Fig.2) when the ROM address inputs can be rejected without affecting the computation. Activity reduction by more than 50% can be achieved in this fashion. In addition to this power reduction technique, the chip has the ability of stopping the computation of any RAC prematurely and using the intermediate result as an approximation of the final value. This is essentially a reduction on-demand in the computation precision of any DCT coefficient emulating the precision reduction imposed by subsequent quantization. To maximize peak SNR (PSNR), multiple computation inhibition thresholds are implemented (classes) and are dynamically selectable through a row and column classification mechanism: If a row exhibits a large absolute difference between its maximum and minimum element, it is assigned more computation precision. This technique has resulted in an extra 10-20 % in computation savings with minimum image degradation.

To achieve high speed at low supply voltages, we used 4-stage carry-bypass adders for the 20-bit accumulators in the RAC units (Fig.2) because they exhibit a small delay slope vs. bitwidth for a small increase in area. The full adder used is shown in Fig. 4. It has been selected because the propagate (P) signal necessary for the carry-bypass operation is computed explicitly, and due to its very small carry-in to carry-out delay.

III. Test Results

The chip (Fig.6) is functional over a wide range of frequencies and power supplies (Fig.7). Over 67,000 separate power measurements on an 8×8 block basis have been taken while the chip was stimulated from 11 test images. The results of these measurements (average $\pm$ 1 stddev) are plotted in Fig.7 vs. the sample standard deviation of the 64 block elements. The image block frequency histogram vs. standard deviation is also shown on the same plot. We observe that according to the design goal, power dissipation shows strong dependence with data correlation (more MSB rejection and less arithmetic activity). Fig.8 plots the same measurement data (averages only) vs. the arithmetic operation bitwidth (number of cycles the RACs of Fig.2 are clocked) along with the corresponding average PSNR. Plots (a) and (c) show power and PSNR vs. bitwidth with row-column classification disabled (no premature computation inhibition is performed.) We observe that PSNR stays constant at well above 40 dB while power increases with bitwidth as expected. Plots (b) and (d) show power and PSNR in the presence of reduced precision selected according to row-column maximum absolute element difference. We observe that PSNR drops as the average bitwidth increases due to the approximation performed by the RAC units. This approximation is more coarse for higher bitwidth inputs. The average power for this dataset with and without computation inhibition is 4.38 mW (reduced precision) vs. 5 mW (full precision) at 1.56V, 14 MHz. The present work exhibits less switched-capacitance per sample (factor of 3) from past DCT chips [1] when normalized for equal process feature size.

We have observed that power savings due to MSBR can be as high as 55% (random blocks vs. fully correlated blocks) with 25% being more typical for still images. Differential video will result in higher savings due to the increased presence of zero-valued data. RCC adds an additional 15% of power savings for minimal PSNR degradation. Much higher power savings can be achieved if we are willing to tolerate more image degradation. The power cost of these methods is estimated at <5%.

References

[1] T. Kuroda et. al, "A 0.9V, 150-MHz, 10-mW, $4mm^2$, 2-D discrete cosine transform core processor with variable-threshold-voltage (VT) scheme," *IEEE Journal of Solid State Circuits*, vol. 31, no. 11, pp. 1770–1777, Nov. 1996.

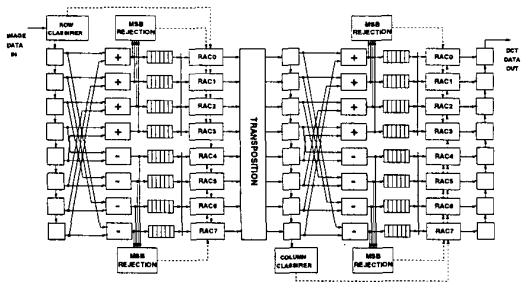

Fig. 1. DCT Chip Block Diagram

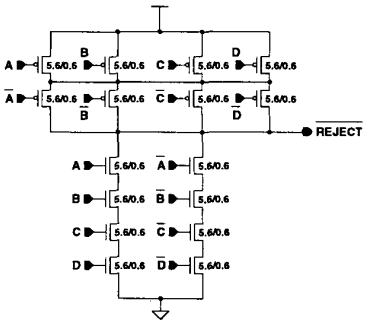

Fig. 3. MSB Rejection Computation Slice

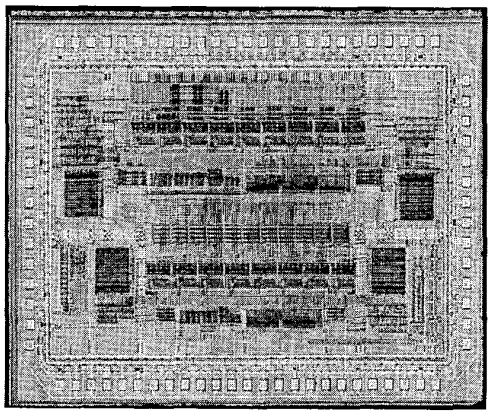

Fig. 5. DCT Chip Microphotograph

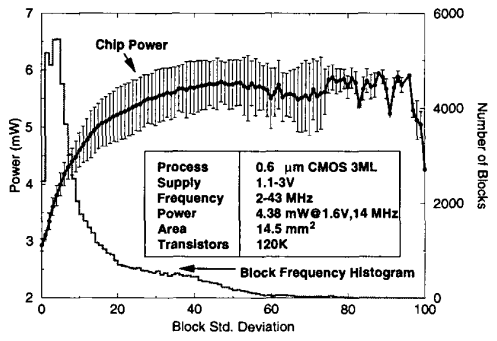

Fig. 7. DCT Chip Measured Power Results at 1.56V, 14 MHz.

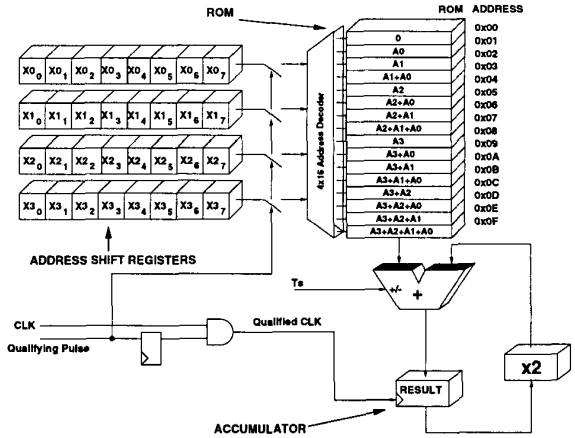

Fig. 2. ROM and Accumulator Unit (RAC)

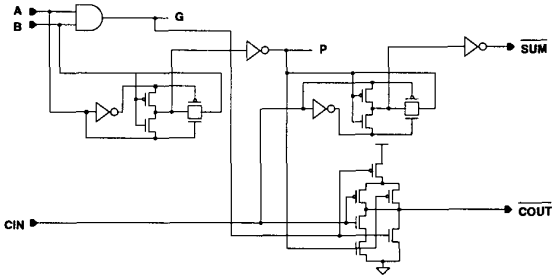

Fig. 4. 26-T Full Adder Used in the RACs

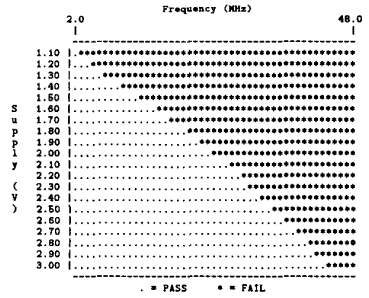

Fig. 6. DCT Chip Schmoo Plot

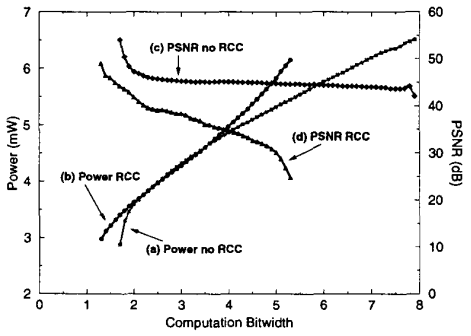

Fig. 8. Measured Chip Power and PSNR vs. Average Bitwidth