A Low-Power IDCT Macrocell for MPEG-2 MP@ML Exploiting Data Distribution Properties for Minimal Activity

Thucydides Xanthopoulos, *Student Member, IEEE*, and Anantha P. Chandrakasan, *Member, IEEE*

***Abstract*—A chip has been designed and tested to demonstrate the feasibility of an ultra-low-power, two-dimensional inverse discrete cosine transform (IDCT) computation unit in a standard 3.3-V process. A data-driven computation algorithm that exploits the relative occurrence of zero-valued DCT coefficients coupled with clock gating has been used to minimize switched capacitance. In addition, circuit and architectural techniques such as deep pipelining have been used to lower the voltage and reduce the energy dissipation per sample. A Verilog-based power tool has been developed and used for architectural exploration and power estimation. The chip has a measured power dissipation of 4.65 mW at 1.3 V and 14 MHz, which meets the sample rate requirements for MPEG-2 MP@ML. The power dissipation improves significantly at lower bit rates (coarser quantization), which makes this implementation ideal for emerging quality-on-demand protocols that trade off energy efficiency and video quality.**

***Index Terms*— Data-dependent computation, inverse discrete cosine transform (IDCT), low power.**

I. Introduction

In THE last few years, there have been two important trends in personal computing. First, use of personal computers has shifted from traditional applications such as word processing and spreadsheet calculations to multimedia and communication-oriented tasks such as World Wide Web access and image, speech, and video processing. As a result, computer workloads consist of arithmetic instructions and digital signal processing (DSP)-like instruction sequences to a greater extent. Second, portable computing devices have become popular as more people rely on them for communication, information access, and entertainment.

Recent innovations in microprocessor design address the first trend. Major microprocessor vendors have included multimedia extensions in their corresponding instruction sets and microarchitectures [1]. DSP codes implemented on general-purpose microprocessors using MMX-style instructions sets satisfy real-time requirements for many applications. Yet, such implementations are power intensive even on low-voltage mobile CPU's and are unsuitable for portable devices and highly integrated systems-on-a-chip requiring both media-processing capabilities and low power dissipation.

Manuscript received September 15, 1998; revised December 22, 1998. This work was supported by Sharp Corp. The work of T. Xanthopoulos was supported in part by a National Semiconductor Fellowship.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Publisher Item Identifier S 0018-9200(99)03672-0.

One promising approach to provide mobile systems with real-time, low-power media-processing capabilities is embedding highly optimized fixed function units for frequent computation tasks without engaging the programmable engine: The computation of the inverse discrete cosine transform (IDCT) for an MPEG-compressed video sequence at 30 frames per second (fps) and 720 $\times$ 480 resolution requires 132 megaoperations per second (MOPS) assuming that a fast IDCT algorithm is used. An MMX-enabled 233-MHz mobile Pentium is capable of executing 932 16-bit media MOPS. We observe that $<$15% of the processor resources are required for the IDCT alone, and therefore real-time performance is more than feasible. Yet, the power dissipation of such a software implementation can be substantial. According to specifications published by Intel, a Tillamook-class mobile Pentium MMX (1.8 V, 233 MHz) exhibits 7 nF of switched capacitance per machine cycle, which corresponds to 5.3 W of power dissipation at 1.8 V and 233 MHz. The IDCT software computation required for a single 720 $\times$ 480 video frame corresponds to 9 mF of switched capacitance (1.3 M machine cycles at 7 nF per cycle), and a software application supporting real-time video at 30 fps will require 0.88 W just for the transform computation. On the other hand, the optimized hardware macrocell presented in this work exhibits as little as 0.1 mF of switched capacitance per video frame, which amounts to $<$5 mW for the real-time transform computation. Power savings over a software implementation exceed a factor of 100.

In this work, we present the design of an IDCT macrocell suited for mobile and highly integrated applications. Our strategy for reducing the chip power was twofold. First, we selected an IDCT algorithm that minimizes activity by exploiting the relative occurrence of zero-valued DCT coefficients in compressed video. Past IDCT chips (e.g., [3]–[5]) have relied on conventional fast IDCT algorithms that perform a constant number of operations per block independent of the data distribution. Our selected algorithm, described in Section II, performs a variable number of operations that depends on the statistical properties of the input data. We implemented this algorithm in hardware through the use of extensive clock gating that minimized the average switched capacitance per sample.

Second, previously proposed architectural-driven voltage scaling techniques [6], [7] have been adapted and used for power minimization. Deep pipelining and appropriate circuit

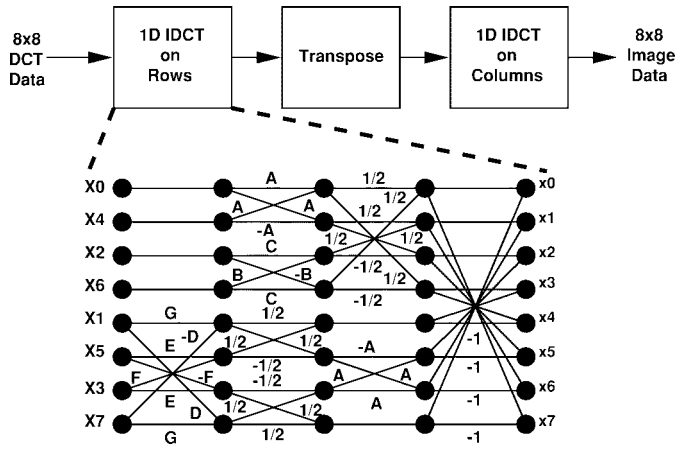

Fig. 1. Two-dimensional Chen separable IDCT algorithm.

techniques have been employed so that the chip could produce 14 Msamples/s (640 × 480, 30 fps, 4:2:0) at 1.3 V (3.3-V process) and meet the requirement for MPEG-2 main level, main profile (MP@ML). Although pipelining increases the total system switched capacitance by <30%, it results in total energy savings because the supply voltage can be reduced.

II. Background and Algorithmic Design

The eight-point, one-dimensional (1-D) inverse DCT [8] of a coefficient vector X is given by the following equation:

$$x[n] = \sum_{k=0}^{7} \frac{c[k]}{2} X[k] \cos\left(\frac{(2n+1)k\pi}{16}\right) \tag{1}$$

where $c[k] = 1/\sqrt{2}$ if $k = 0$ and 1 otherwise. Since the DCT is a separable transform, the two-dimensional (2-D) IDCT of an 8 × 8 coefficient block can be computed by applying (1) on each block row, transposing the intermediate result and reapplying (1) on each column.

There have been numerous fast IDCT algorithms that minimize the number of multiplications and additions implied by (1) [9], [10]. A popular algorithm that has been embedded in various very large-scale integration implementations [11], [12] is the Chen algorithm [9], depicted in Fig. 1. The Chen algorithm is based on a sparse matrix factorization of the matrix form of (1). Letters A–G represent constant values, and each dot in the flowgraph represents a dual multiply-and-sum operation. The Chen algorithm requires a constant number of operations per block: 256 multiplications and 416 additions.

The statistical distribution of the input DCT coefficients possesses unique properties that may affect IDCT algorithm design. Previous work [13]–[15] has established that ac [nonzero (NZ) spatial frequency] DCT coefficients of compressed images and video exhibit zero-mean Laplacian distribution

$$f(x) = \frac{\lambda}{2} e^{-\lambda|x|}, \qquad \lambda = \frac{\sqrt{2}}{\sigma_x}. \tag{2}$$

Muller [16] has proposed a slightly different statistical model based on Kolmogorov–Smirnov and χ^2 fitness tests [17]

$$f(x) = \frac{\nu\alpha(\nu)}{2\sigma\Gamma\left(\frac{1}{\nu}\right)} e^{-[\alpha(\nu)|(x/\sigma)|]^{\nu}} \tag{3}$$

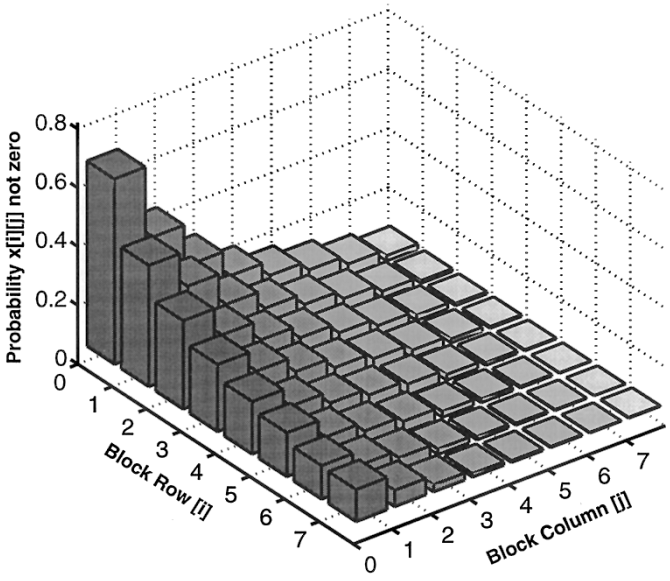

Fig. 2. Probabilities of nonzero occurrence for 8 × 8 DCT coefficient blocks.

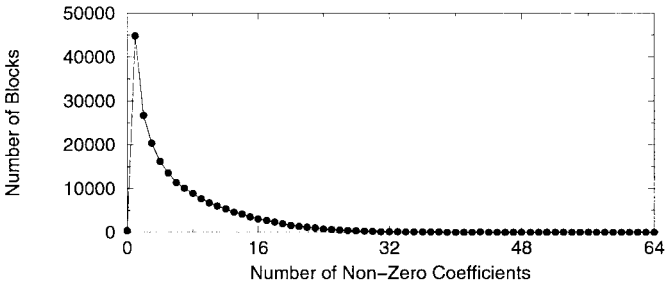

Fig. 3. Histogram of nonzero DCT coefficients in sample MPEG stream.

$$\alpha(\nu) = \sqrt{\frac{\Gamma(3/\nu)}{\Gamma(1/\nu)}} \tag{4}$$

where $\Gamma(.)$ denotes the gamma function and ν and σ are positive constants. Equation (3) denotes a generalized Gaussian distribution. Choice of parameter ν transforms (3) into a Laplacian or a Gaussian ($\nu = 1$ or $\nu = 2$, respectively). Lee *et al.* [18] have also concluded that the generalized Gaussian distribution (3) is valid for a class of medical images.

The existence of sharp zero-centered dual-sided Laplacian distributions for 63 out of 64 spectral coefficients implies a large number of zero-valued coefficients per 8 × 8 block for image and video data. Typically, DCT blocks of MPEG-compressed video sequences have only five to six nonzero coefficients, mainly located in the low spatial frequency positions. The three-dimensional histogram of Fig. 2 plots the probability of occurrence of a nonzero spatial frequency for each block position for a typical video sequence. Higher probabilities are concentrated in the low spatial frequency range, which is consistent with the spatial low-pass characteristics of video sequences. The histogram of Fig. 3 shows the frequency of block occurrence plotted versus the number of nonzero coefficient content for a typical MPEG sequence. The mode of such distributions is invariably blocks with a single nonzero spectral coefficient (typically the dc).

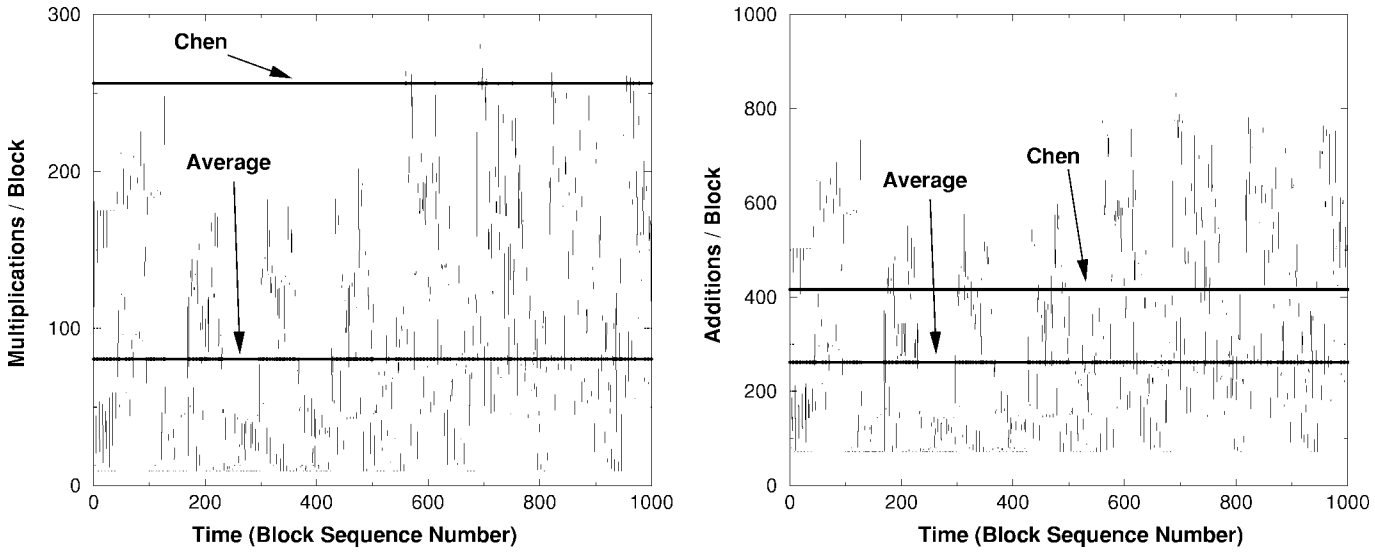

Fig. 4. Number-of-operations comparison between Chen and data-driven IDCT.

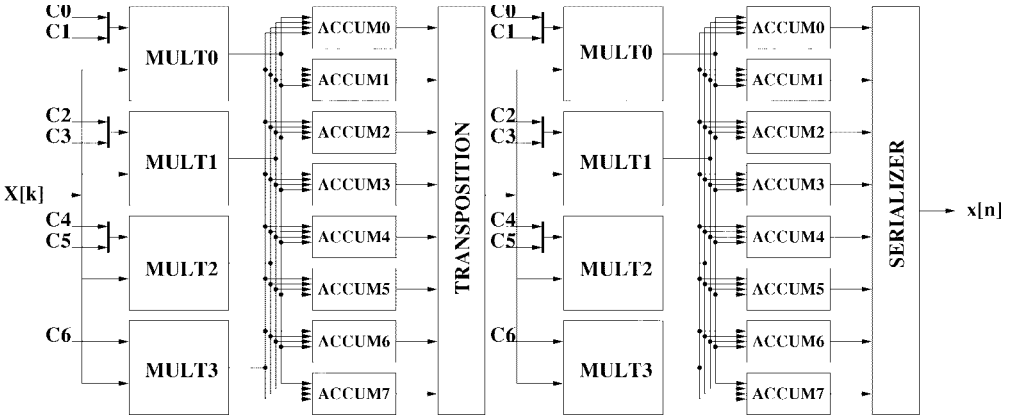

Fig. 5. IDCT-chip block diagram.

Based on the information above on the statistical distribution of DCT coefficients, we decided to depart radically from conventional IDCT algorithms that perform a fixed number of operations per block. Given such input data statistics, we observe that direct application of (1) will result in a small average number of operations since multiplication and accumulation with a zero-valued $X[k]$ coefficient may constitute a "no operation" (NOP). Appropriate implementation of (1) can result in an algorithm with a variable number of operations per block as opposed to standard fast algorithms that cannot exploit data properties. This idea of a coefficient-by-coefficient IDCT implementation was originally proposed by McMillan and Westover [19]. The IDCT chip presented here consists of two 1-D IDCT units implementing (1) and a transposition memory structure in between. On the other hand, conventional architectures perform a constant number of operations per sample and do not exploit data properties. Butterfly-style operations absorb the zero-valued coefficients early in the signal path without affecting average switched capacitance. Fig. 4 shows the instantaneous and average number of multiplications and additions required by the data-driven algorithm of (1) for 1000 blocks extracted from an MPEG video sequence. We observe that the data-driven algorithm requires a smaller number of operations per block compared to the conventional Chen algorithm. The potential for a low-power implementation of this algorithm is evident due to the small arithmetic operation requirements.

III. Chip Architecture

A block diagram of the IDCT chip is shown in Fig. 5. Incoming coefficients are scaled by a maximum of four constants at a time for a maximum of seven different constant scalings (C0–C6) due to the symmetry of the cosine basis functions. The scaled values are being added or subtracted to the state of eight accumulators attached to the scaling multipliers through a full crossbar switch. The scaling constants and the crossbar setting are determined by the index k of each coefficient $X[k]$ within the eight-point vector X. Every eight cycles, the 1-D IDCT of a block row has converged within the eight accumulators. The intermediate result is fed into the transposition structure (TRAM) of Fig. 6. This structure is a 2-D array of shift registers that can shift data from left to right and top to bottom. The TRAM performs transposition on the fly and *eliminates* the need for double buffering, which would

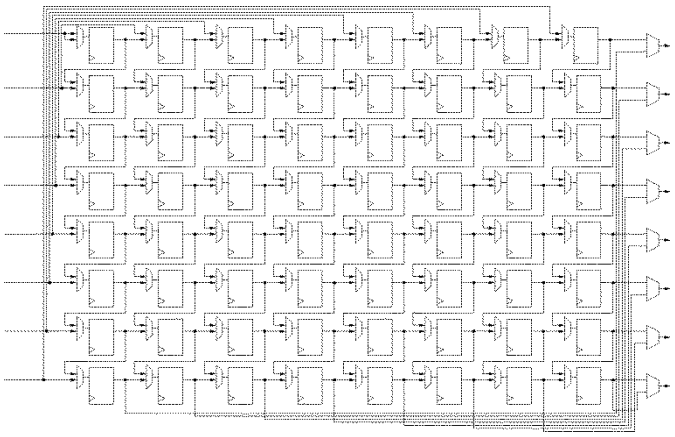

Fig. 6. Transposition structure. Data are transposed on the fly by changing the shifting direction ($T \rightarrow B$ or $L \rightarrow R$). This scheme eliminates the need for ping-pong buffering.

have been necessary had a static dual-addressed RAM been used. The second 1-D stage is identical to the first except for a change in the precision of the scaling constants.

To meet the bandwidth requirement of 14 Msamples/s at $V_{\rm DD} = 1.3$ V using a high V_T process, the multiply–accumulate (MAC) operations must be pipelined by a factor of four. Fig. 7 shows the MAC structure used. For simplicity, the diagram does not show the crossbar switch between the multiplier and the adder structures. The figure depicts a 5×4 unit as opposed to the full 14×14 structure used. The level of pipelining is identical. The multiplier is a standard carry–save array multiplier. The accumulator structure is a simple ripple carry adder. The accumulator adds a single full adder delay to the critical path due to the coincident carry propagation with the availability of product bits of the multiplier. Although pipelining increases the total switched capacitance of the computation unit, it has the side benefit of reducing the propagation of spurious transitions within the MAC structure.

The datapath width of the arithmetic units has been optimized to meet IEEE Standard 1180–1990 for IDCT precision by a comfortable margin but at the same time discard unnecessary precision bits that would increase the power dissipation. Table I summarizes our fixed-point precision deviations. All data sets pass the pixel peak error requirement. The zero-input requirement is also passed.

A. Clock Gating

The presence of many zero-valued coefficients must be exploited in order to reduce the switching activity and reap the low-power benefits of this particular algorithm. If the MAC structure of Fig. 7 was not pipelined, conditional gating of the accumulator clock based on whether or not the incoming coefficient is zero would substantially reduce the switching activity and still produce the correct result. Pipelining, though, makes clock gating more complicated. If all pipeline stages are clocked with a single clock net, clock gating becomes impossible since different pipeline stages are processing different and possibly nonzero DCT coefficients. Powering down all the stages will produce erroneous results. Clock gating can be implemented if each pipeline stage uses a separate clock net gated by an appropriate qualifying pulse. The qualifying pulse propagates from stage to stage along with the nonzero coefficient that requires processing. If a zero-valued coefficient enters the pipeline, only the stage that corresponds to the zero is powered down. The other upstream and downstream stages remain unaffected. Our clock-gating scheme is graphically depicted in Fig. 8. The IDCT chip features ten separate clock nets in addition to the master clock for fully qualifying all steps of the entire pipeline formed by both 1-D stages.

A clock-gated pipeline presents certain physical design challenges in order to avoid common race conditions. Race conditions may arise in the following situation, which is depicted in Fig. 9(a). Let us assume that $CLK0$ arrives at the clocked element of stage 0 at time t_0 and $CLK1$ arrives at time t_1. If $t_1 - t_0 > t_{\rm CLK \rightarrow Q} + t_{\rm pd} - t_{\rm hold}$, the wrong data will be sampled at stage 1. The expected value of $\Delta t = t_1 - t_0$ can be greater than typical clock-skew values because the clock nets are physically separate and are affected by more mismatch phenomena than distributed *RC* delays on a nominally equipotential surface. Moreover, in the IDCT chip, all gated clock nets are global, spanning all the pipelines formed by the accumulators. The gate load seen by each clock net is quite different due to data skewing (Fig. 7) in addition to different physical wire loading. This problem becomes very important when the intermediate combinational logic is minimal or nonexistent (i.e., back-to-back shift registers) and $t_{\rm pd} \rightarrow 0$.

We solved this problem by careful buffering and physical design of the clock nets and multiple simulation iterations with accurate extracted parasitics. In cases of small timing margins (i.e., shift registers and least significant bits of adder stages), a negative level-sensitive latch was inserted, clocked by the upstream pipeline clock, as shown in Fig. 9(b). This ensured functional correctness with a minimal penalty (<2%) in terms of power dissipation and no effect on the system critical path.

IV. Circuit Style

The greatest circuit challenge presented was the design of the pipelined multiply–accumulate unit shown in Fig. 7 given $V_{\rm GS} - V_T = 0.4$ V for PMOS devices and 0.6 V for NMOS devices. The critical path for this operation is seven full adder delays given a 14×14 pipelined multiplier. We imposed the following requirements on the full adder design used in the multipliers (Fig. 10). First, since the sum output (S) propagation delay is also critical (in addition to the $\overline{COUT}$), we wanted S to be computed using a single gate, as a function of A, B, and CIN. A second gate delay would increase the critical path substantially and would force us to raise the power supply above 1.3 V. Second, the S gate should not employ more than two series PMOS pullup devices in the signal path. PMOS transistors see at maximum 0.4 V above threshold across their gate and source terminals, and their mobility is three times smaller than that of the NMOS devices. Given these facts, we observed that more than two series PMOS devices in the sum circuit would be considerably slow in pulling up the sum output unless substantially oversized,

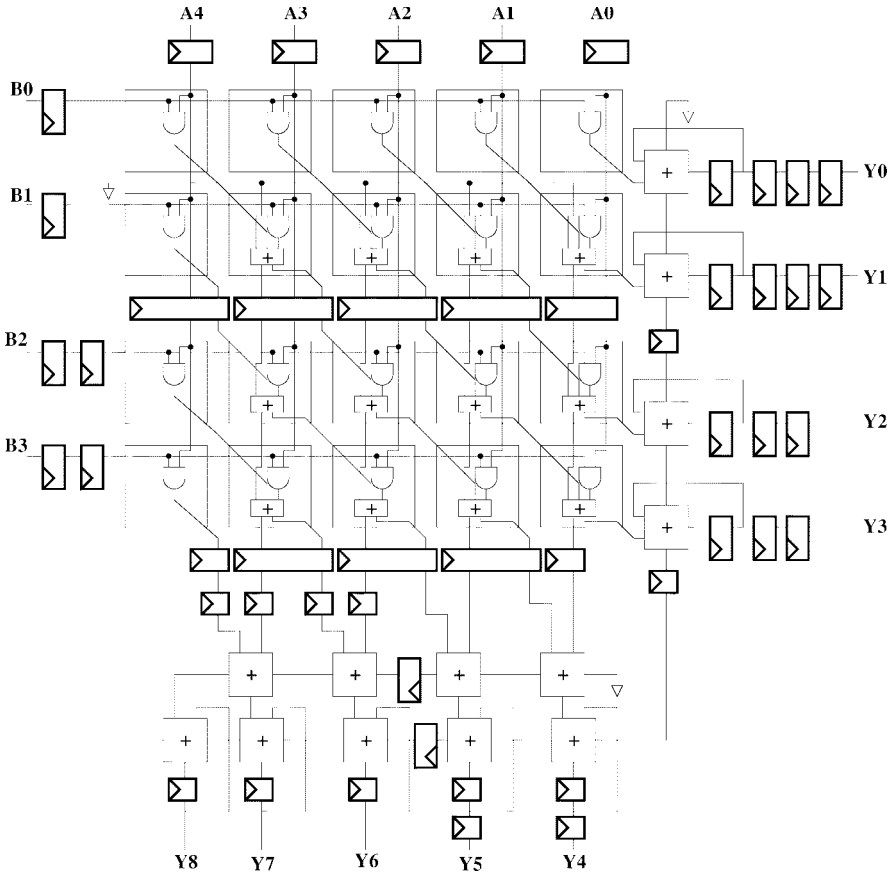

Fig. 7. Pipelined multiply–accumulate structure.

TABLE I
IEEE STANDARD 1180–1990 COMPLIANCE

Dataset	PMSE	OMSE	PME	OME
	<0.16	<0.02	<0.015	<0.0015
$[-256, 255]$	0.015800	0.013597	0.002500	0.000153
$[-5, 5]$	0.011000	0.009136	-0.002400	-0.000005
$[-300, 300]$	0.014100	0.011886	0.002700	0.000180
$-[-256, 255]$	0.016000	0.013578	-0.002500	-0.000122
$-[-5, 5]$	0.011000	0.009156	0.002500	0.000016
$-[-300, 300]$	0.014000	0.011878	-0.002900	-0.000122

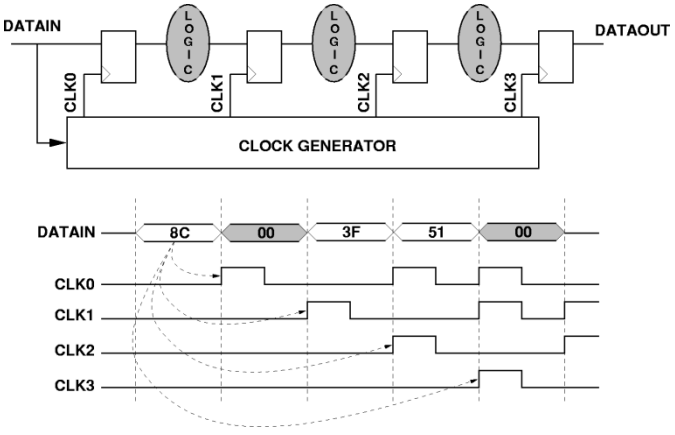

Fig. 8. Clock-gating approach in a pipeline.

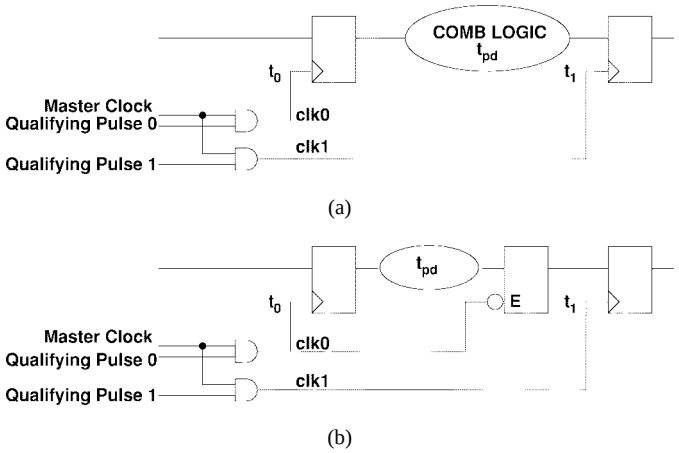

Fig. 9. Potential race conditions in clock-gated pipelines. (a) Potential race condition. (b) No race condition.

which would considerably increase the total cell switched capacitance.

Given the constraints above, we decided to use the hybrid full adder circuit of Fig. 10. In this circuit, $\overline{COUT}$ is derived in a static CMOS fashion but $\overline{S}$ is derived using cascade voltage switch logic (CVSL). The advantage of using this logic family for the sum output is that S is pulled high indirectly through a stack of NMOS transistors' overpowering a PMOS cross-coupled pair as opposed to a stack of PMOS devices' pulling high. The main disadvantages are that the

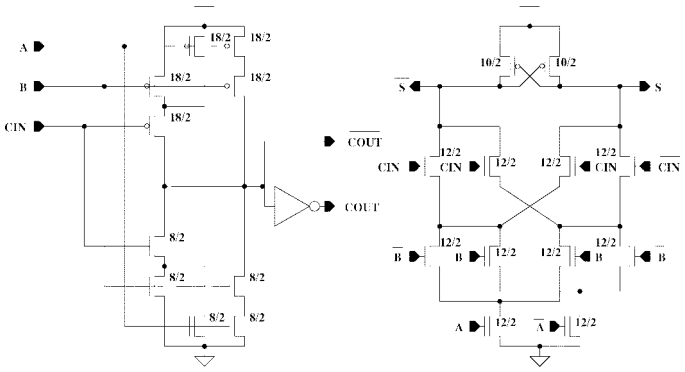

Fig. 10. Hybrid adder schematics (sizes in $\lambda = 0.35\,\mu$).

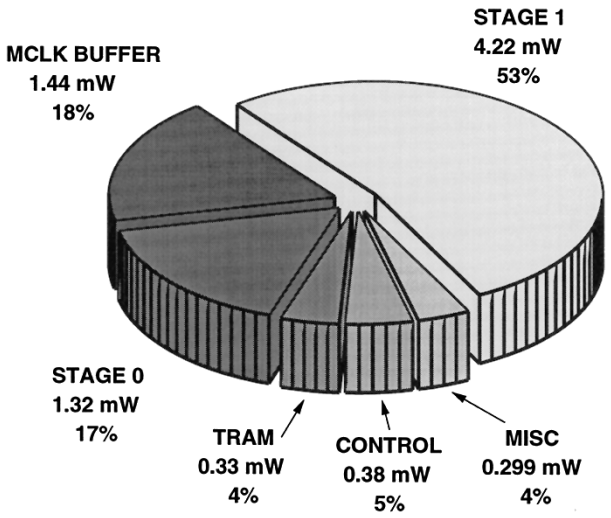

Fig. 11. IDCT-chip Pythia simulation results.

complementary NMOS switch structures require complementary inputs and the presence of small crossover currents. Since these full adders are used within an array multiplier, one of the complementary inputs is the S output of the previous row adder, which is already in complementary form. Only $\overline{COUT}$ needs to be complemented in order to be used as an input in the next row of adders. The total device count (24) has not been increased from conventional static CMOS adder implementations [20]. Simulation has shown that the hybrid full adder is faster than conventional adders (i.e., 37% faster than a 24-transistor "mirror adder" [20]).

V. Power Estimation Methodology

The IDCT chip is a reasonably large system (160-K transistors) and needs to be simulated for a large number of cycles so that conclusions regarding correlation of power dissipation versus processing load can be drawn. Circuit simulators (i.e., HSPICE, Powermill) would be slow to run and would provide insufficient information. For this purpose, we have developed Pythia [21], a fast and accurate power estimation tool that works with a structural Verilog description. Pythia uses a collection of Verilog-programming-language-interface access and utility routines to traverse circuit hierarchy, extract user-provided parameters, and monitor signal transitions. Each net in the design is annotated with four different capacitance values:

1) a gate capacitance value indicating the total gate capacitance attached to the net;
2) an NMOS drain capacitance value indicating the total n-drain to p-substrate or p-well junction capacitance attached to the net;
3) a PMOS drain capacitance value indicating the total p-drain to n-substrate or n-well junction capacitance attached to the net;
4) a routing capacitance indicating the total metal wiring capacitance of the net.

During each net transition, there is a different energy contribution from each particular capacitive component. Pythia performs numerical integration in order to account for the fact that gate capacitances are a function of voltage. Moreover, it uses the depletion approximation and a first-order analytical solution for the CV integral in the case of junction capacitances. Process parameter information is provided in the form of spice model files. Finally, Pythia expands cell internal capacitive nodes that are collapsed in a Verilog gate-level simulation so that every circuit capacitive node is accounted for. Pythia is very fast (imposing a speed-penalty factor of five to seven on the underlying Verilog simulation) and reasonably accurate (typically within 5% from HSPICE and Powermill). Feedback is provided in the form of a report with energy and power information for each hierarchical block in the design.

Power library generation has been automated within the cadence design framework. Each library cell requires two phases of annotation. During the first phase, cell terminal capacitances are added to the cell Verilog functional view. In the second phase, an internal model is generated that identifies the internal nodes (nonterminals) of each cell and describes the switching state of such nodes for each input state. In this way, all capacitive nodes are included in the power estimation.

Fig. 11 shows the IDCT-chip power estimation results partitioned among the top-level modules of the design. The data were obtained by feeding the chip 1000 random blocks from the MPEG sequence cheerleader (a more complete description of the test sequences used appears in Section VI). The simulated clock frequency and power supply are 14 MHz and 1.3 V, respectively. The software tool includes an elementary interconnect capacitance estimation mechanism based on first-order process parameters and net fanout. Approximately 30% of the total chip power of 7.95 mW is attributed to interconnect capacitance.

Stages 0 and 1, which are the main computation stages, dissipate the bulk of the power (70%). The clock buffer of the master clock dissipates a substantial 18% of the total. This is the globally distributed clock that produces all derivative pipeline clocks. The final inverter of this exponential driver has a total gate width of 56 μm (W/L ratio of 93.3). Although the TRAM contains a large number of flip–flops, its duty cycle is low (clocked once every eight cycles) and only requires 4% of the total power. Global control circuitry requires 5% of the power, with the remaining 4% attributed to miscellaneous circuitry for block I/O and glue logic between stages. We note

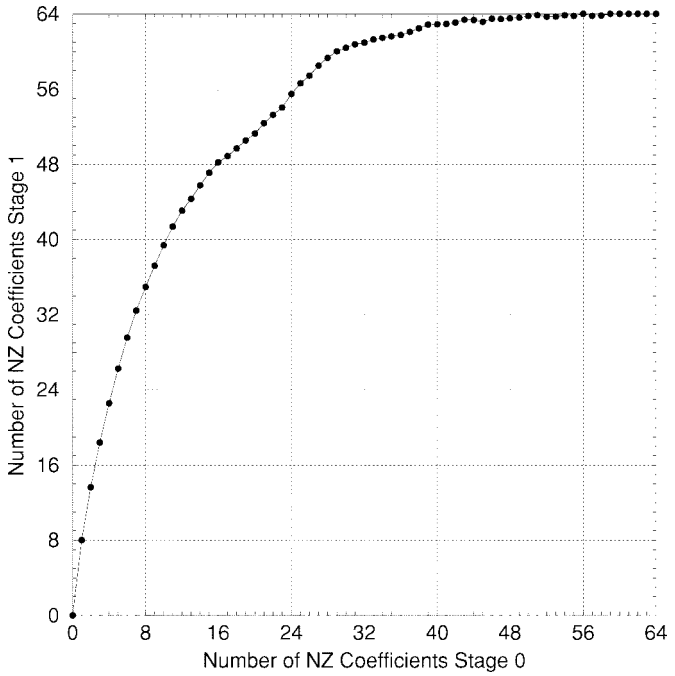

Fig. 12. Stage 1 NZ coefficients per block versus corresponding stage 0 NZ coefficients for a typical MPEG sequence.

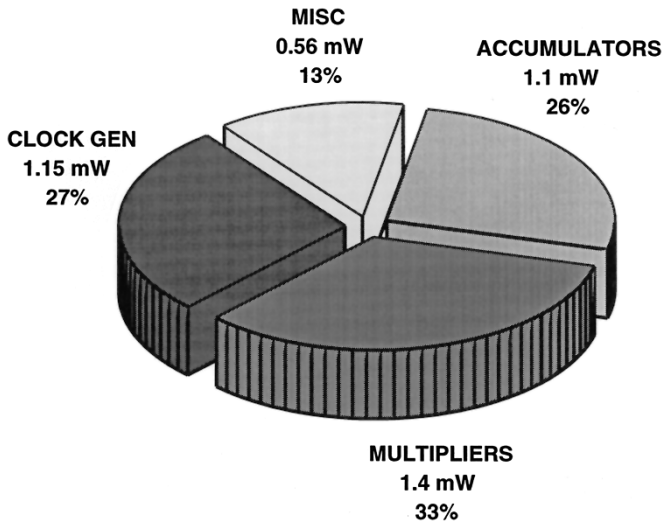

Fig. 13. Stage 1 Pythia simulation results.

that stage 1 requires much more power than stage 0. This happens because more nonzero coefficients are processed on average by stage 1 due to the nature of (1). Fig. 12 plots the number of nonzero coefficients per block encountered at stage 1 of the IDCT chip versus the number of nonzero coefficients encountered at stage 0.

Fig. 13 shows the distribution of power within the stage 1 computational block. We observe that a substantial portion of the power (27%) is used for the clock generator (and buffer) for all pipeline stages. The rest is mainly computational power.

A. Glitch Power Investigation

Pythia can report separately the power due to spurious transitions (glitches). This is accomplished by keeping track

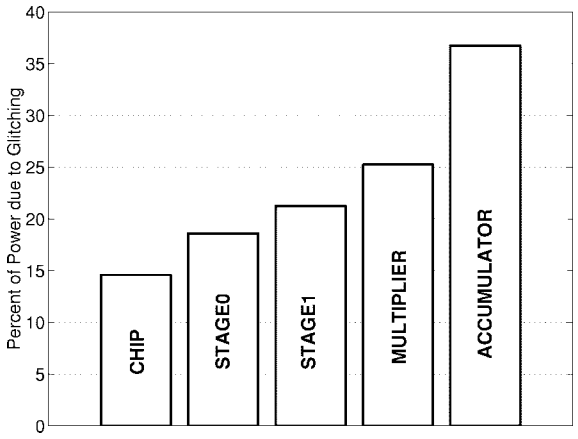

Fig. 14. Glitch power estimation.

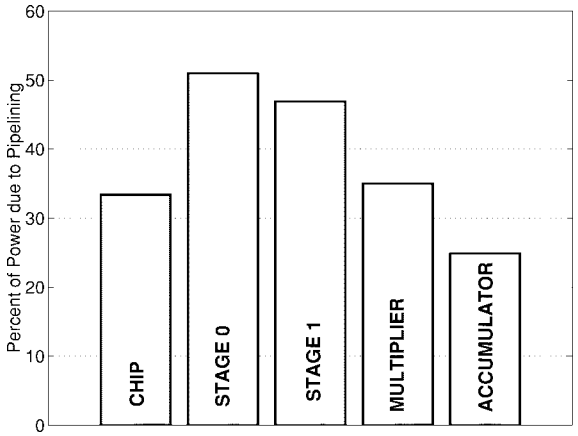

Fig. 15. Pipeline power estimation.

of each signal pulse width and reporting separately pulses that have a width less that half the master clock period. Fig. 14 shows the percentage of power due to glitches for a number of blocks in the design. The accumulators in the design are the most glitch-prone blocks because all three inputs (A, B, CIN) of each full adder become valid at different times. The multipliers exhibit a smaller percentage of spurious transitions because our gate model assumes equal delays between the sum and carry outputs of our full adders. As a result, two out of three inputs of each full adder in the Braun array (Fig. 7) become valid simultaneously and glitching is reduced. We tried to equalize these two delays in the actual full adder circuit design (Fig. 10) for this purpose. The total simulated chip power due to glitching is 1.16 mW (14.59%).

B. Pipeline Power Savings Investigation

The IDCT chip employs a substantial level of pipelining for the sole purpose of lowering the power supply at real-time video sample rates; it is not required for algorithm functionality. We wanted to identify the total power savings attributed to pipelining by assessing the overhead required and adjusting for a higher supply voltage to maintain the sample rate.

Fig. 15 shows the percentage of power for major chip blocks, which is attributed to pipelining. This power includes

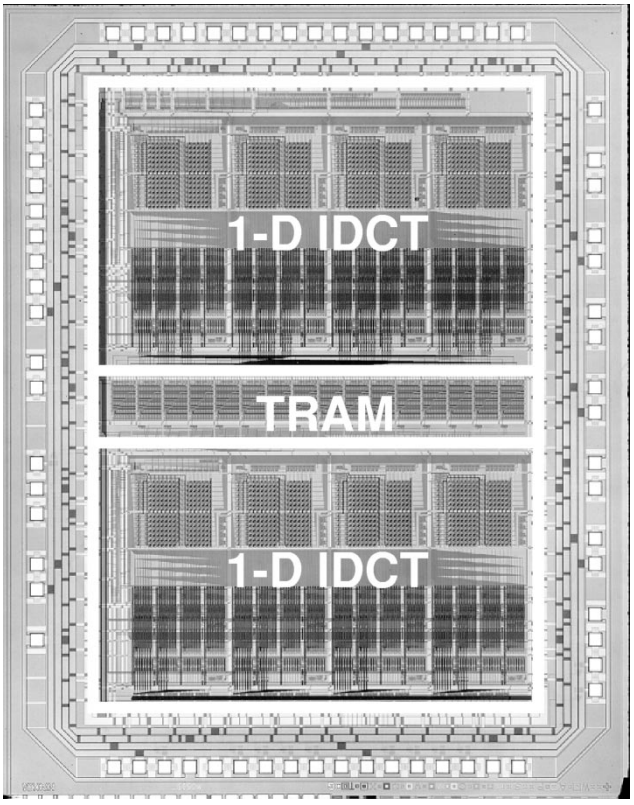

Fig. 16. IDCT-chip microphotograph.

TABLE II
PROCESS AND IDCT-CHIP SPECIFICATIONS

Process	$0.5\mu m$ CMOS ($0.7\mu m$ drawn), 3ML
VTN	0.66V
VTP	-0.92V
TOX	9.6 nm
Supply	1.1-1.9 V
Frequency	5-43 MHz
Power	4.65 mWatts @ 1.32 V, 14 MHz
Area	20.7 mm^2
Transistors	160K

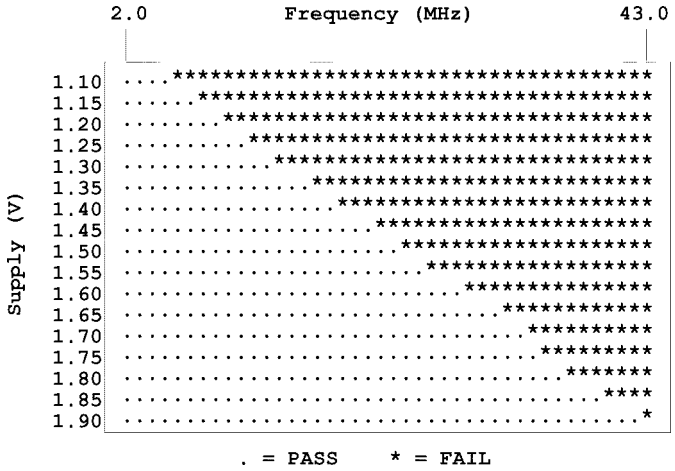

Fig. 17. IDCT-chip shmoo plot.

the two main pipeline clock generators (one for each stage) in addition to all the pipeline flip-flops in the signal path. The figure of 33% as a total power cost of pipelining is exaggerated because a significant percentage of the pipeline flip-flops serve as the state registers of the accumulators, which would still be present (along with their appropriate clock buffer) in the absence of pipelining. If we perform this adjustment, we estimate the actual cost of pipelining to about 20% of the total power.

In the absence of pipelining, the chip critical path grows by a factor of four. Testing the chip at various clock frequencies (Section VI) indicated that the supply voltage should be raised to 2.2 V in order to maintain a sample rate of 14 MS/s. Subtracting the overhead and accounting for the increased supply, we estimate that the chip total average power would have been about 10 mW. The pipeline therefore accounts for more than 50% of power savings.

TABLE III
MPEG SEQUENCES USED FOR TESTING THE IDCT CHIP

Sequence	Resol.	Bit Rate	Content
Ballet	704 × 480	5M	Foreground movement panning, scene change
Bicycle	704 × 480	5M	Background movement, zooming
Cheerleader	704 × 480	5M	Foreground movement
Flower	352 × 240	1.5M	Panning
Snowman	320 × 240	1M	Computer Graphics, minimal movement
Susi	352 × 288	1.5M	Head and shoulders, little foreground movement

VI. TEST RESULTS

Fig. 16 shows a die photo of the fabricated chip, and Table II summarizes the chip and process specifications. A printed circuit board has been built for testing and measurement. A dedicated PC interface through an industry standard adapter (ISA) National Instruments 32-bit digital I/O card has been used for providing stimuli and reading results from the board under software control. The board contains SRAM buffers and control finite-state machines to permit high-speed testing in spite of the slow PC ISA interface. Moreover, the IDCT chip core uses separate supply lines to permit accurate current measurements. The chip is functional over a wide range of frequencies and power supplies, as indicated by the shmoo plot of Fig. 17. Operation at frequencies above 50 MHz at $<$2.5 V has been observed using high-speed digital pattern generators.

More than 2.8 million separate power measurements were taken during several weeks while the chip was stimulated from six different MPEG compressed sequences ranging from 1 to 10 MB/s and having rather different content (Table III). The entire chip pipeline was filled with data from a single 8 × 8 coefficient block repeated multiple times during each

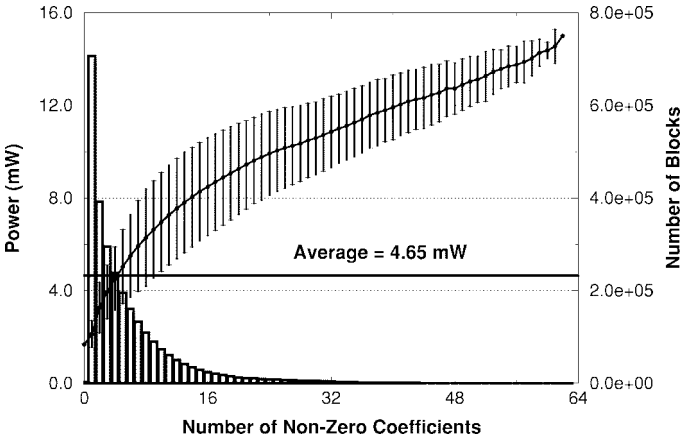

Fig. 18. IDCT-chip measured power results at 1.32 V and 14 MHz. This plot represents more than 2.8 million separate power measurements.

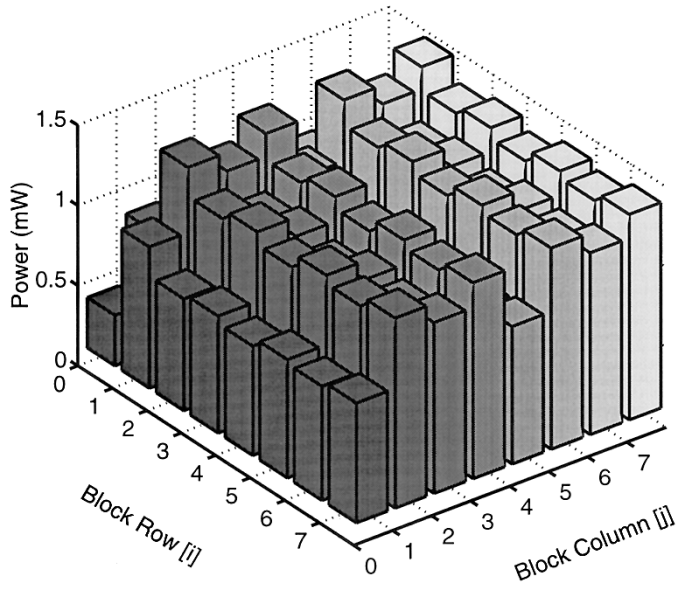

Fig. 19. IDCT-chip measured power results at 1.32 V and 14 MHz: average power dissipation per block position across all possible 12-bit DCT coefficient values (1 NZ coefficient per block).

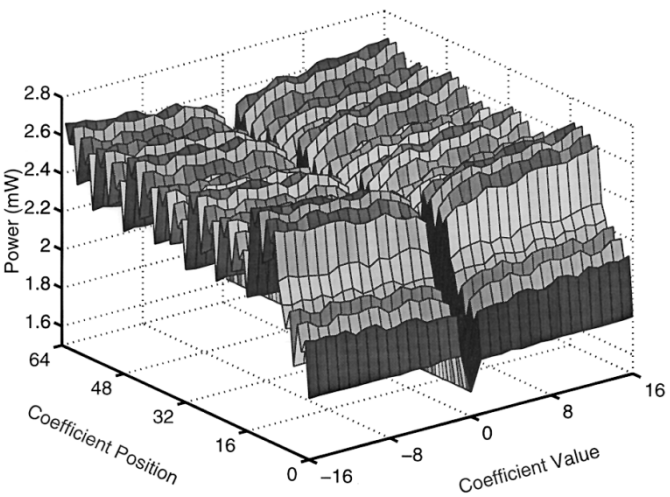

Fig. 20. IDCT-chip measured power results at 1.32 V and 14 MHz: blocks containing 1 NZ coefficient of magnitude ≤ 16.

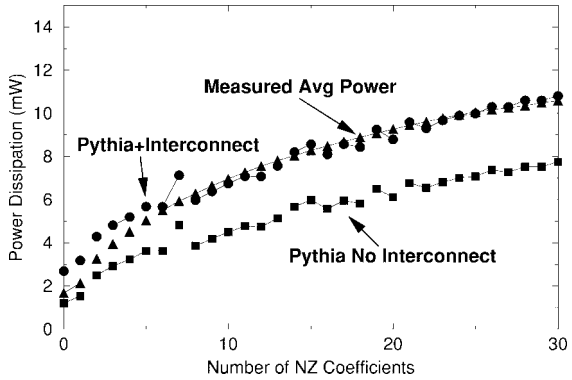

Fig. 21. Comparison of simulated and experimental power results.

one of the measurements. The results of these measurements are plotted in Fig. 18 versus the number of nonzero coefficients within each DCT block. On the same graph, the block nonzero content histogram is also shown. The average, as well as the 95% confidence interval, is plotted for each data point. Blocks with the same number of nonzero coefficients can exhibit a range of power dissipation. We have investigated this variation in power measurements by exercising the chip with blocks that contained a single nonzero coefficient exhaustively spanning the entire position-value space (2^{18} measurements). Fig. 19 shows the average power dissipation among all possible 12-bit values for each position within the 8×8 spectral block. We observe a maximum of 60% variation from position to position. This is mainly due to a different number of multiplications required for each spectral position within the block. The surface plot of Fig. 20 shows power variation with respect to both block position and coefficient value. The two position axes of Fig. 19 have been collapsed to a single 64-position axis. We observe a slight reduction in power with decreasing magnitude. This is due to the fact that all multipliers in the chip implement sign-magnitude arithmetic and exhibit reduced activity for small magnitude numbers. Along the zero-value axis, there is an abrupt reduction in power because the arithmetic units are not active at all and power is only due to master clock distribution, transposition memory operation, and global control.

Turning our focus back to Fig. 18, we observe that power dissipation shows strong correlation versus the nonzero coefficients because of the data-dependent processing algorithm. The average power dissipation for this data set is 4.65 mW at 1.32 V and 14 MHz. The power savings exhibit diminishing returns as the number of nonzero coefficients increases. This happens because the first 1-D IDCT stage produces many more nonzeroes at its output, as mentioned in Section V. As we move along the x-axis in Fig. 18, most of the gains come from stage 0 only; the gains from the second stage become less significant. At the average workload value though of 5.68 nonzero coefficients for this data set, power savings from both stages are significant.

Fig. 21 compares the results obtained from our Pythia simulator for a set of 300 random blocks from the cheerleader sequence with the corresponding measured averages of

Fig. 18. Two sets of simulation results are shown: one set includes interconnect capacitance estimation and the second set was obtained with interconnect estimation off.

When the present work is compared to previous DCT/IDCT chips [3], [4] in terms of switched capacitance per sample, it performs better by a factor of two. The DCT/IDCT macrocell reported in [5] exhibits less switched capacitance per sample than the present work but has been implemented in a much smaller process feature size (0.3 μm) and threshold voltage ($V_T = 0.1 \pm 0.15$ V). When the switched capacitance is normalized across the two macrocells to account for process differences, the present work still performs better by a factor of two. The switched-capacitance-per-sample metric takes into account only algorithmic and high-level architectural design decisions that reduce the total activity; it factors out parameters such as clock frequency and supply voltage. It must be noted that the present chip employs a higher degree of pipelining than the others for the sole purpose of reducing the power supply, and the switched-capacitance metric penalizes such a design choice. Our algorithmic choice has therefore been justified given that all past chips employ standard algorithms performing a fixed number of operations per block.

Our final observation is that Fig. 18 indicates that the power dissipation of the IDCT chip improves significantly at lower bit rates (coarser quantization). This makes our approach ideal for emerging quality-on-demand compression protocols. At low rates, the present chip will exhibit much more efficiency versus other implementations than the previous paragraph has implied.

VII. Conclusion

We have demonstrated an ultra-low-power IDCT chip designed in a widely available 0.7-μm high V_T process. The chip demonstrates the lowest switched capacitance per sample among past IDCT chips presented in the literature. Low-power operation has been achieved through selection of a data-dependent algorithm, aggressive voltage scaling, deep pipelining, extensive clock gating, and appropriate transistor-level circuit techniques. The chip is fully functional and has undergone extensive testing. Optimized low-power macrocells like the chip presented in this work for common multimedia tasks are a promising approach for providing multimedia processing abilities to mobile low-power systems.

References

[1] A. Peleg, S. Wilkie, and U. Weiser, "Intel MMX for multimedia PC's," *Commun. ACM*, vol. 40, no. 1, pp. 24–38, Jan. 1997.

[2] Intel Corp. (1999). Mobile Pentium processor with MMX™ technology on .25 micron. [Online]. Available: http://developer.intel.com/design/mobile/datashts/243468.htm.

[3] M. Matsui, H. Hara, K. Seta, Y. Uetani, L. Kim, T. Nagamatsu, T. Shimazawa, S. Mita, G. Otomo, T. Oto, Y. Watanabe, F. Sano, A. Chiba, K. Matsuda, and T. Sakurai, "200 MHz video compression macrocells using low-swing differential logic," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 1994, pp. 76–77.

[4] A. K. Bhattacharya and S. S. Haider, "A VLSI implementation of the inverse discrete cosine transform," *Int. J. Pattern Recognit. Artificial Intell.*, vol. 9, no. 2, pp. 303–314, 1995.

[5] T. Kuroda, T. Fujita, T. Nagamatu, S. Yoshioka, F. Sano, M. Norishima, M. Murota, M. Kako, M. Kinugawa, M. Kakumu, and T. Sakurai, "A 0.9 V, 150-MHz, 10-mW, 4 mm2, 2-D discrete cosine transform core processor with variable-threshold-volt age (VT) scheme," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1770–1777, Nov. 1996.

[6] A. Chandrakasan, C. Sheng, and R. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, pp. 473–484, Apr. 1992.

[7] A. Chandrakasan, A. Burstein, and R. Brodersen, "A low-power chipset for a portable multimedia I/0 terminal," *IEEE J. Solid-State Circuits*, vol. 29, pp. 1415–1428, Dec. 1994.

[8] N. Ahmed, T. Natarajan, and K. Rao, "Discrete cosine transform," *IEEE Trans. Comput.*, vol. C-23, pp. 90–93, Jan. 1974.

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

[10] E. Feig and S. Winograd, "Fast algorithms for the discrete cosine transform," *IEEE Trans. Signal Processing*, vol. 40, pp. 2174–2193, Sept. 1992.

[11] M. T. Sun, T. C. Chen, and A. M. Gottlieb, "VLSI implementation of a 16 × 16 discrete cosine transform," *IEEE Trans. Circuits Syst.*, vol. 36, pp. 610–617, Apr. 1989.

[12] S. Uramoto, Y. Inoue, A. Takabatake, J. Takeda, Y. Yamashita, H. Terane, and M. Yoshimoto, "A 100 MHz 2-D discrete cosine transform core processor," *IEEE J. Solid-State Circuits*, vol. 36, pp. 492–499, Apr. 1992.

[13] R. Reininger and J. Gibson, "Distributions of the two-dimensional DCT coefficients for images," *IEEE Trans. Commun.*, vol. COM-31, pp. 835–839, June 1983.

[14] S. Smoot and L. Rowe, "Study of DCT coefficient distributions," in *Proc. SPIE*, Jan. 1996, vol. 2657, pp. 403–411.

[15] A. Netravali and B. Haskell, *Digital Pictures: Representation, Compression and Standards*, 2nd ed. New York: Plenum, 1995.

[16] F. Muller, "Distribution shape of two-dimensional DCT coefficients of natural images," *Electron. Lett.*, vol. 29, no. 22, pp. 1935–1936, Oct. 1993.

[17] M. Kendall and A. Stuart, *The Advanced Theory of Statistics.* London, U.K.: Charles Griffin, 1967.

[18] H. Lee, Y. Kim, A. Rowberg, and E. Riskin, "Statistical distributions of DCT coefficients and their application to an interframe compression algorithm for 3-d medical images," *IEEE Trans. Med. Imag.*, vol. 12, pp. 478–485, Sept. 1993.

[19] L. McMillan and L. A. Westover, "A forward-mapping realization of the inverse discrete cosine transform," in *Proceedings of the Data Compression Conference (DCC'92).* Piscataway, NJ: IEEE Computer Society Press, 1992, pp. 219–228.

[20] N. Weste and K. Eshraghian, *Principles of CMOS VLSI Design*, 2nd ed. Reading, MA: Addison-Wesley, 1993.

[21] T. Xanthopoulos, Y. Yaoi, and A. Chandrakasan, "Architectural exploration using Verilog-based power estimation: A case study of the IDCT," in *Proc. 34th Design Automation Conf. (DAC'97)*, June 1997, pp. 415–420.

Thucydides Xanthopoulos (S'91) received the S.B. and S.M. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1992 and 1995, respectively, where he currently is pursuing the Ph.D. degree in data-driven signal-processing systems for video applications.

His research interests include low-power VLSI design, embedded DSP architectures, and CAD for low power.

Mr. Xanthopoulos is a member of the Association of Computing Machinery, Tau Beta Pi, and Eta Kappa Nu. In 1996, he received a National Semiconductor Graduate Fellowship.

Anantha P. Chandrakasan (S'87–M'95) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he currently is an Associate Professor of electrical engineering and computer science. He held the Analog Devices Career Development Chair from 1994 to 1997. His research interests include the ultra-low-power implementation of custom and programmable digital signal processors, distributed wireless sensors, multimedia devices, emerging technologies, and CAD tools for VLSI. He is a coauthor (with R. Brodersen) of *Low Power Digital CMOS Design* (Norwell, MA: Kluwer Academic, 1995) and a coeditor (with R. Brodersen) of *Low Power CMOS Design* (New York: IEEE Press, 1998).

Mr. Chandrakasan received the NSF Career Development Award in 1995, the IBM Faculty Development Award in 1995, and the National Semiconductor Faculty Development Award in 1996 and 1997. He received the 1993 IEEE Communications Society's Best Tutorial Paper Award and the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997. He has served on the Technical Program Committee of various conferences, including the ISSCC, VLSI Circuits Symposium, DAC, ISLPED, and ICCD. He was a Technical Program Cochair for the 1997 International Symposium on Low-Power Electronics and Design, VLSI Design'98, and the 1998 IEEE Workshop on Signal Processing Systems. He was General Cochair of the 1998 ISLPED. He is currently the Signal Processing Subcommittee Chair for ISSCC'99. He is an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He is a member of the Design and Implementation of Signal Processing Systems (DISPS) Technical Committee of the Signal Processing Society.