

A 180-mV Subthreshold FFT Processor Using a Minimum Energy Design Methodology

Alice Wang, *Member, IEEE*, and Anantha Chandrakasan, *Fellow, IEEE*

Abstract—In emerging embedded applications such as wireless sensor networks, the key metric is minimizing energy dissipation rather than processor speed. Minimum energy analysis of CMOS circuits estimates the optimal operating point of clock frequencies, supply voltage, and threshold voltage [1]. The minimum energy analysis shows that the optimal power supply typically occurs in subthreshold (e.g., supply voltages are below device thresholds). New subthreshold logic and memory design methodologies are developed and demonstrated on a fast Fourier transform (FFT) processor. The FFT processor uses an energy-aware architecture that allows for variable FFT length (128–1024 point), variable bit-precision (8 b and 16 b) and is designed to investigate the estimated minimum energy point. The FFT processor is fabricated using a standard 0.18- μm CMOS logic process and operates down to 180 mV. The minimum energy point for the 16-b 1024-point FFT processor occurs at 350-mV supply voltage where it dissipates 155 nJ/FFT at a clock frequency of 10 kHz.

Index Terms—CMOS digital integrated circuits, CMOS memory circuits, coprocessors, design methodology, digital signal processors, leakage currents, logic design, subthreshold CMOS circuits.

I. INTRODUCTION

THERE is significant research activity to minimize energy dissipation at the system level to lengthen battery lifetimes for embedded applications. Minimum energy analysis of CMOS circuits predicts the optimal operating point including clock frequencies, supply voltage, and threshold voltage. Our analysis shows that optimal supply voltage to minimize energy typically occurs in the subthreshold region. In order to investigate the optimal supply voltage, a new minimum energy design methodology is created to design circuits to operate at supply voltages far below the minimum energy point. The minimum energy design methodology was demonstrated on a fast Fourier transform (FFT) processor used for wireless sensor networks.

A distributed wireless sensor network is a collection of a large number (tens to thousands) of distributed microsensor nodes. Networked microsensors enable a variety of new applications such as warehouse inventory tracking, location sensing, machine-mounted sensing, patient monitoring, and building climate control [2]–[5]. The microsensor nodes must operate

from scavenged energy from the environment such as from solar sources, mechanical vibration and RF [6], [7]. A self-powered system using a variable MEMS capacitor was able to deliver 10 μW 's of power to a DSP [8]. Energy scavenging constrains the total system power be less than tens of microwatts, which is a challenging design goal for microsensors.

In order to function within the extremely low power requirements, a minimum energy design methodology that includes energy-aware architectures and subthreshold circuits, is proposed. The design methodology is demonstrated on a subthreshold FFT processor. The FFT design uses a modified standard logic cell library, custom multiplier and memory generators, and is fabricated using a 0.18- μm standard CMOS process. No additional process steps or body-biasing techniques are used. The FFT processor operates down to 180 mV where it runs at 164 Hz and the power dissipation is 90 nW. The minimum energy point for the 16-b 1024-point FFT processor occurs at 350 mV supply voltage where it dissipates 155 nJ/FFT at a clock frequency of 10 kHz.

II. ENERGY-AWARE ARCHITECTURES

The FFT is a commonly used signal processing function that extracts the frequency and phase information from the sensor signals. The FFT has been used in target tracking, localization, and radar by analyzing the phase differences between multiple sensors [9]. It has also been used in compression algorithms and also within communication systems [10]. Dedicated low-power FFT processors are desired to sustain low-power requirements of various embedded applications [11].

The real-valued FFT algorithm uses the symmetries of computing the CVFFT of real-valued inputs to reduce computation approximately by half. For example, a 1024-pt RVFFT is efficiently performed by computing a 512-pt CVFFT and then transforming the outputs back for 1024-pt. A proposed architecture for a RVFFT processor is a traditional CVFFT architecture that includes the backend processing needed for the RVFFT. Fig. 1 shows the conventional radix-2 butterfly architecture; for every clock cycle one radix-2 butterfly is performed. The FFT architecture is designed with various power hooks that allow the architecture to gracefully scale FFT length and bit precision. The energy awareness of the FFT is increased by adding additional hardware to cover functionality over many operating scenarios [12]. This architecture demonstrates techniques that enable FFT lengths from 128 to 1024 points as well as both 8-b and 16-b computation. An area and energy-efficient design methodology is proposed to enable scalability across bit precision and FFT lengths [13].

Manuscript received April 5, 2004; revised July 26, 2004. This work was supported by the Defense Advanced Research Projects Agency (DARPA) Power Aware Computing/Communication Program and Air Force Research Laboratory, under agreement numbers F30602-00-2-0551 and F33615-02-2-4005. The work of A. Wang was supported by an Intel Ph.D. Fellowship.

A. Wang is with Texas Instruments Inc., Dallas, Texas 75243 USA (e-mail: aliwang@ti.com).

A. Chandrakasan is with the Massachusetts Institute of Technology, Cambridge, Massachusetts 02139 USA (e-mail: anantha@mit.edu).

Digital Object Identifier 10.1109/JSSC.2004.837945

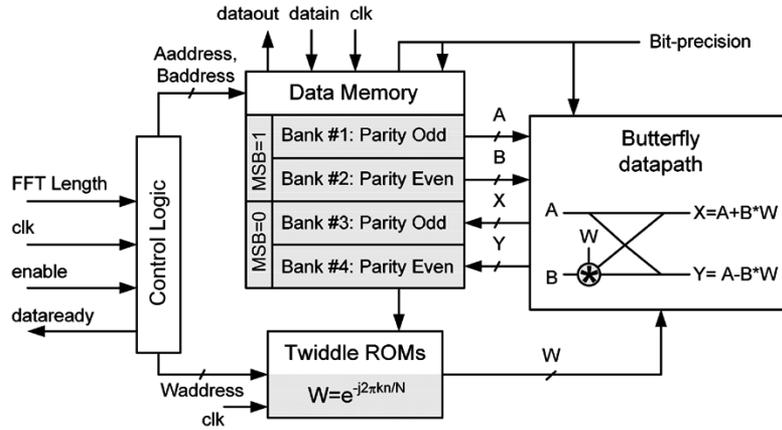


Fig. 1. Radix-2 butterfly FFT architecture.

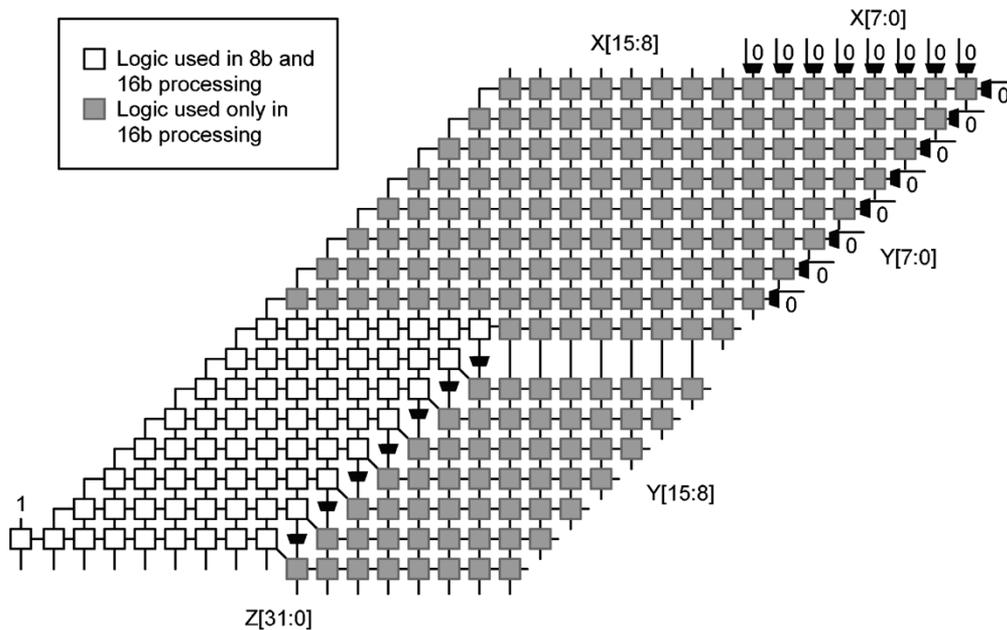


Fig. 2. 8-b and 16-b scalable Baugh-Wooley multiplier.

A. Variable Bit Precision

One of the hooks designed into the energy-aware FFT processor is variable bit precision, which is showcased by the Baugh-Wooley (BW) multiplier for two's complement multiplication [13]. A nonscalable BW design optimizes for the worst case scenario by building a single multiplier for the largest bitwidth. This design is nonoptimal; when lower precision multiplications are performed, the two's complement sign extension bits causes a switching energy overhead. The proposed scalable BW multiplier design recognizes that the MSB quadrant contains a lower bit-precision multiplier. The MSB quadrant of the multiplier includes those gates associated with the MSB inputs. Any other quadrant of the multiplier does not have the correct BW configuration and would require additional logic. To minimize switching in the LSB adders, the LSB inputs are gated. Fig. 2 demonstrates this technique for a BW multiplier that is scalable for 8-bit and 16-bit precisions. Similar bit-precision scalability was applied to the entire butterfly datapath, data memories, and Twiddle ROMs.

B. Variable FFT Length

Variable FFT length is another hook designed into the architecture of the FFT processor [13]. As FFT length is varied, the processor can adjust the memory size, leading to energy savings. Careful design consideration is needed in the memory access logic to enable variable FFT lengths. A dedicated memory is designed for the FFT processor to ensure that read and write hazards do not occur when accessing two values from the memory. The memory consists of four memory blocks with a parity-bit/MSB crossbar for memory accesses. Also the control logic adjusts the number of butterflies performed with FFT length for improved energy scalability.

III. MINIMUM ENERGY POINT ANALYSIS

The FFT processor is designed to operate at the optimal operating point that minimizes energy dissipation. Analysis of the energy and performance of the FFT shows that the minimum energy point occurs at supply voltage levels below the threshold

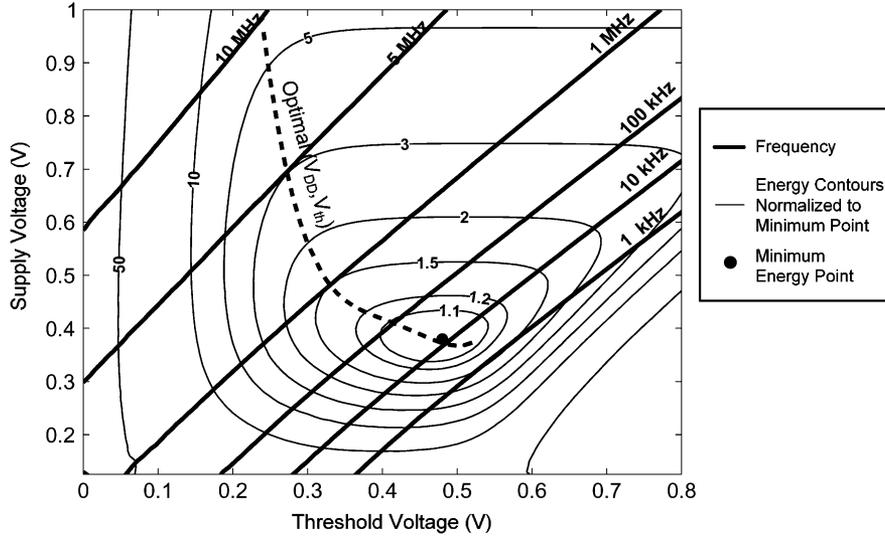


Fig. 3. Minimum energy point and constant energy and performance contours of the 16-b and 1024-pt FFT.

voltage. Scaling the supply voltage (V_{DD}) below the threshold voltage (V_{th}) limits the performance of CMOS circuits, but leads to orders of magnitude energy savings over nominal V_{DD} operation.

The total energy of the FFT is broken down into switching energy and leakage energy. Switching energy is modeled as

$$E_{\text{switching}} = \alpha N C V_{DD}^2 \quad (1)$$

where α is the activity factor, N is the number of clock cycles, C is the switched capacitance of the circuit, and V_{DD} is the supply voltage.

A simplified model of the leakage energy based on the BSIM transistor model is given by

$$E_{\text{leakage}} = V_{DD} \cdot I_S \cdot e^{\frac{V_{gs} - V_{th}}{nV_T}} \cdot \left(1 - e^{\frac{-V_{ds}}{V_T}}\right) \cdot T. \quad (2)$$

Here, I_S is a technology-dependent scaling parameter, V_{gs} is the gate-to-source voltage, V_{ds} is the drain-to-source voltage, V_{th} is the threshold voltage, V_T is the thermal voltage, n is related to the subthreshold slope, and T is the latency of computation [14].

Fig. 3 shows simulated energy contours of the 16-b 1024-pt FFT for the region of $V_{DD} = 100 \text{ mV} - 1 \text{ V}$ and $V_{th} = 0 - 800 \text{ mV}$ for a $0.18\text{-}\mu\text{m}$ process [15]. The energy shown is the average energy per FFT and is derived from the active and leakage energy models. The increasing contours of constant energy are normalized to the minimum energy point. The interaction between supply voltage, threshold voltage, and latency causes a minimum energy dissipation point at $(V_{DD}, V_{th}) = (380 \text{ mV}, 480 \text{ mV})$ with a clock frequency of 13 kHz.

For a given threshold, as the power supply is dropped (starting from a large value), the switching and overall energy reduces. However, as the supply is reduced into the subthreshold region, the propagation delay increases significantly resulting in a corresponding increase in leakage energy as predicted by (2). This results in an operating point for supply voltage and clock frequency which minimizes total energy dissipation.

Operating below the optimal supply and frequency results in energy being dominated by subthreshold leakage. Similarly, for a fixed supply voltage, as the threshold is increased (starting from a small value), the leakage energy reduces as the leakage current reduces in an exponential fashion with only a modest increase in delay. However, as the threshold approaches the supply voltage, the increase in delay is significant and causes the leakage energy to increase at some threshold voltage. This results in a minimum energy point.

Also shown in Fig. 3 are contours of constant performance of the FFT processor (1 kHz–10 MHz). The figure shows that performance scales with the difference between the gate voltage and the threshold voltage.

The metrics of energy and performance are the key metrics for managing system-level power dissipation for two cases.

Case 1: Processing speed is not critical. For this case, the optimal operating point occurs at the optimal frequency point. The FFT processor should operate at the optimal voltage and frequency, and then shut down the power supply after performing the computation to prevent any additional leakage energy dissipation.

Case 2: Processing speed is critical. For this case, an optimal operating curve is extracted from the contours and contains those points where one performance contour is tangent to one energy contour. These points are given by the dotted line in Fig. 3. For a given performance constraint the supply voltage and threshold voltage are set along the optimal curve to achieve minimum energy dissipation. If the frequency required is less than the optimal, then the processor should operate at the optimal frequency and then shutdown after completion. The simulation data assumes that techniques such as supply voltage and threshold voltage scaling and power gating, do not incur any overhead. However, in a complete system analysis, all energy overhead should be included.

Our FFT processor is designed and fabricated in a standard $0.18\text{-}\mu\text{m}$ CMOS logic process with a fixed nominal threshold voltage of 450 mV [16]. According to the contours at 450-mV

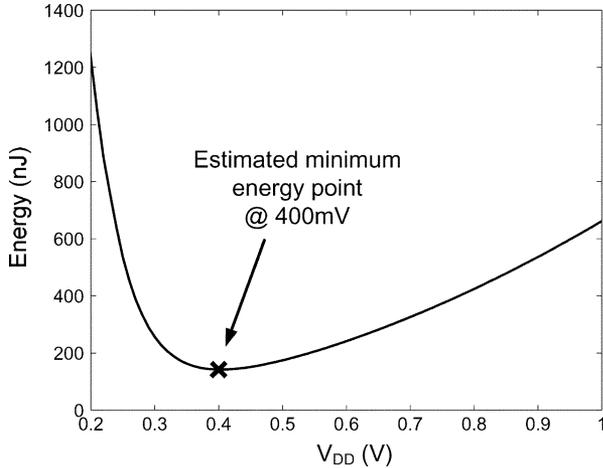


Fig. 4. Estimated minimum energy point for the FFT using a typical transistor in a 0.18- μm technology occurs at 400 mV which is lower than the threshold voltage (450 mV).

threshold voltage, the minimum energy point of the FFT processor occurs at 400 mV (Fig. 4). The figure confirms that for a given threshold, as the supply voltage decreases, the switching and overall energy reduces. But, in the subthreshold region, the propagation delay increases exponentially resulting in a increase in leakage energy. We propose a new subthreshold design methodology that is demonstrated in a subthreshold FFT processor designed for subthreshold operation. Previous minimum energy system-level optimization have focused on architectural tradeoffs for minimum energy dissipation given fixed clock frequencies [1]. Our FFT processor shows the minimum energy point as both supply voltage and clock frequency are scaled.

IV. SUBTHRESHOLD LOGIC

In our subthreshold FFT design, we perform minimum supply voltage analysis in addition to minimum energy point analysis. Subthreshold design has been used extensive in low-power analog designs [17]. The theoretical minimum supply voltage for a CMOS inverter based on transistor models has been established [18]. Previous research has focused on logic families for low-voltage operation [19]. An ultra low-voltage design demonstrates the functionality of logic circuits at 200 mV using low-threshold devices [20]. Using p-well and n-well biasing techniques, a multiply-accumulate implementation is able to operate as low as 175 mV [21].

Minimum supply voltage differs from minimum energy operation. Minimum supply voltage analysis is needed to estimate the lowest supply voltage where the FFT is able to function and to ensure that the FFT processor can function at supply voltages near and below the estimated optimum. Only in high activity factor circuits do the minimum supply voltage and minimum energy point coincide. The predominant effects on the minimum voltage of logic circuits are transistor sizing, process variations, and circuit styles. In our minimum voltage analysis of logic circuits, we analyze a standard cell library starting with the inverter.

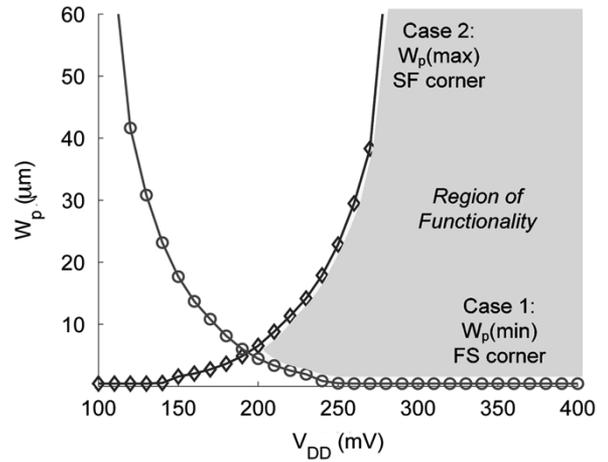


Fig. 5. Minimum voltage operation of the inverter is affected by process variations. Given the worst case corners, Fast NMOS/Slow PMOS (FS) and Slow NMOS/Fast PMOS (SF), the minimum voltage occurs at 195 mV.

A. Subthreshold Inverter

Analysis of the CMOS inverter with a minimum-sized NMOS device exposes the effect of process variations on minimum voltage operation. Traditionally, the PMOS is sized so that the drive strength is large enough to pull the output node up to a desired output-high voltage level. Fig. 5 shows two cases where the decreased supply voltage, transistor sizing and process variations affect the inverter operation assuming an output swing of 10%–90% V_{DD} .

Case 1: When ‘0’ is applied to the input, then the PMOS is sized to pull-up the output node to ‘1’. However, as the voltage supply of the inverter decreases, the idle current becomes significant. This results in a decreasing I_{on}/I_{off} , and causes the output node to drop. The I_{on}/I_{off} ratio is used to indicate if a logic gate will function properly. I_{on} is defined to be the drive current of the devices and decreases exponentially as supply voltage is lowered in subthreshold operation. I_{off} is defined as the idle current. The I_{on}/I_{off} ratio is further reduced when considering process variations. The worst case corner for the inverter is the fast NMOS/slow PMOS (FS) because the fast NMOS is leakier than the slow PMOS. Fig. 5 shows $W_p(\text{min})$ at the FS corner, where $W_p(\text{min})$ is the minimum allowable PMOS width that still drives the output to $0.9 \cdot V_{DD}$.

Case 2: When ‘1’ is applied to the input, the minimum sized NMOS pulls the output node to ‘0’. A large W_p leads to a large PMOS idle current compared to the drive current of the NMOS. Therefore, there is a maximum bound on W_p , to allow the output to be driven low. The slow NMOS/fast PMOS (SF) corner sets $W_p(\text{max})$ in Fig. 5. The intersection of the curves at these two process corners indicates that the inverter is only guaranteed to operate down to 195 mV by sizing the PMOS to be $5.4 \mu\text{m}$, which is 12 times the minimum width. For the FFT implementation, where the PMOS are typically sized $3 \times$ larger, the minimum supply voltage is 220 mV. However, for a typical transistor process corner, the inverter can operate below 100 mV.

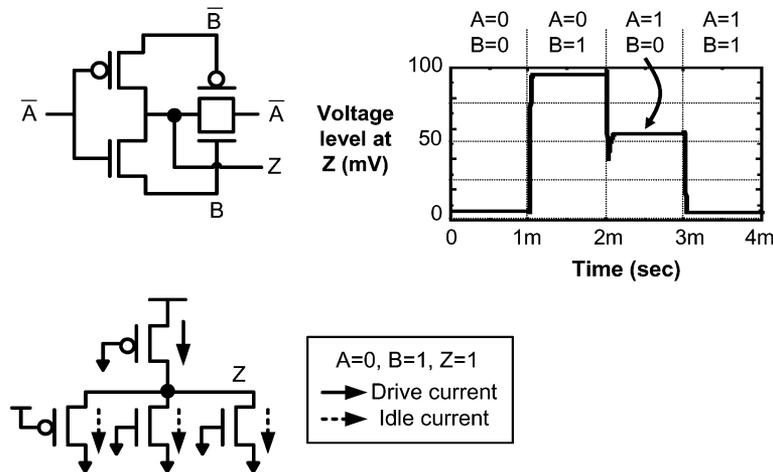


Fig. 6. Tiny XOR gate illustrates the effect of parallel leakage. The idle current through the three parallel devices degrades I_{on}/I_{off} and affects functionality for $A = 1$ and $B = 0$.

B. Subthreshold Standard Cell Library

To create a subthreshold logic cell library, first we analyzed traditional logic circuits at low voltages. Parallel leakage, stacked transistors, and sneak leakage effects affect functionality during low-voltage operation when the gate drives are lowered to the point that the on currents are competing with the leakage currents. In this section, we show how we developed a methodology to design subthreshold logic cells that are used in the FFT processor design and ensure that these effects are minimized.

Parallel leakage occurs when the idle current of parallel devices reduces I_{on}/I_{off} . An example of parallel leakage is illustrated in the XOR gate. Fig. 6 shows the schematic of the tiny XOR gate commonly used in traditional circuit design. An analysis of the drive currents and leakage currents for the input vector ($A = 1, B = 0$) shows that because there are three ‘off’ devices and one ‘on’ device, I_{on}/I_{off} is degraded. The simulation shows that for ($A = 1, B = 0$), the Z output is only driven to 55 mV. This effect is further compounded if process variations are considered in the analysis.

A transmission gate XOR (Fig. 7) is a circuit where the number of parallel devices are minimized (or balanced) for minimum-voltage operation. Because there are two devices pulling the output node high and two devices pulling low, I_{on}/I_{off} is not degraded and the XOR is fully functional as shown in the simulation. Also, because both NMOS and PMOS are in the pull-up and pull-down path, the effects of process variations are mitigated.

Long stacks of transistors also affect the functionality of logic gates. When stacked devices are conducting, the effective drive of each device is diminished. Also, the threshold voltage of a stacked device increases due to larger source-to-body voltages causing both drive and leakage currents to decrease.

When designing a standard-cell library, cell interfaces need to be considered. If all cells are designed using static CMOS then sneak leakage paths are avoided. However, when transmission or pass gates are introduced into the library, sneak leakage paths and long stacks of devices occur during synthesis. Therefore, to

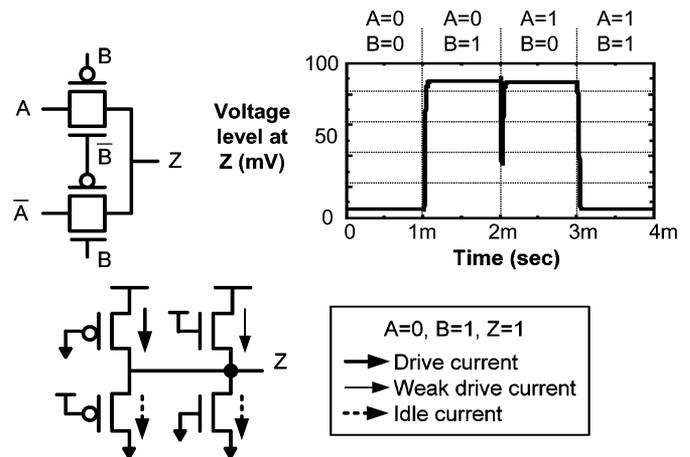


Fig. 7. A transmission gate XOR has balanced leakage and is functional for all input vectors at 100 mV.

ensure functionality in the design of our subthreshold standard cell library, all inputs and outputs to the cells were buffered.

In the design of the custom BW multipliers, all cell interfaces were studied to eliminate redundant buffering and to ensure no sneak leakage paths or long stacks of transistors. The BW multiplier layout was generated using custom BW multiplier compilers.

V. SUBTHRESHOLD MEMORY DESIGN

The FFT processor contains eight $128 \text{ W} \times 16 \text{ b}$ RAM blocks and four $256 \times 16 \text{ b}$ ROM blocks. First, we analyze the functionality of conventional memory designs in subthreshold. In standard memory schemes, a six-transistor (6T) SRAM is used. Current memory design challenges, such as bitline capacitance, bitline leakage, speed, and sense amplifier design, worsen as voltage supplies scale down and at worst case process variations.

Second, we propose write and read schemes for subthreshold memories. The hierarchical read-bitline is used in the design of data memory and ROMs and achieves acceptable I_{on}/I_{off} in subthreshold.

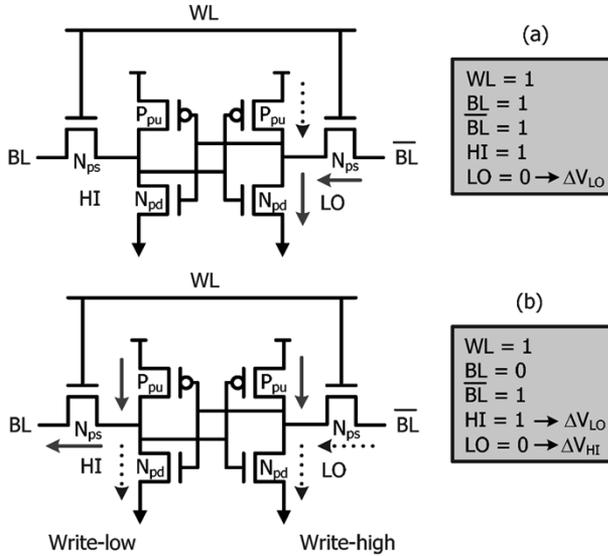


Fig. 8. The read and write conditions places sizing constraints on the memory cell. They are analyzed at the worst case process corners to account for large subthreshold leakage currents.

A. Subthreshold Write Access and Bitcell

At high supply voltages, the 6T SRAM memory cell sizing is determined by the read/write conditions. In our subthreshold analysis, the read/write operation conditions are performed at the worst case process corners to account for the large subthreshold idle currents.

The read upset condition is analyzed to ensure that the memory cell is not written during a read access. Fig. 8(a) shows the drive and idle currents for the read upset condition in subthreshold. The NMOS pulldown transistor width (N_{pd}) is sized large enough to ensure that the voltage at LO does not rise above ΔV_{LO} due to the on current through the pass transistor (N_{ps}) and the idle current through the pullup transistor (P_{pu}). This analysis is performed at the SF corner to account for the maximum leakage through P_{pu} .

Similarly shown in Fig. 8(b) is the write condition. The write condition is analyzed for two cases. The Write-Low case is when a '0' is written into HI, and the Write-High case is when a '1' is written into LO. The Write-Low case determines the minimum width for N_{ps} to pull HI down to ΔV_{LO} and is performed at the SF corner. In the Write-High case, the analysis places a constraint on the maximum width of N_{pd} and N_{ps} . Large idle current through N_{pd} and N_{ps} causes a voltage divider that overpowers the drive current of P_{pu} used to pull LO up to ΔV_{HI} . These three conditions place sizing constraints on all devices in the SRAM circuit and limit functionality at low-voltages. Fig. 9 shows a sizing analysis on N_{pd} given that node HI is pulled down to $\Delta V_{LO} = 20\% \cdot V_{DD}$, and node LO is pulled up to $\Delta V_{HI} = 80\% \cdot V_{DD}$ and $P_{pu} = N_{ps}$. It suggests that the worst case processor variations makes it difficult to satisfy both read and write conditions of the 6T SRAM memory cell at supply voltages much below 500 mV. Also, as the supply voltage decreases, the memory cell size increases dramatically due to the SF corner, where the weak NMOS has to overcome the strong PMOS feedback when writing '0' to HI.

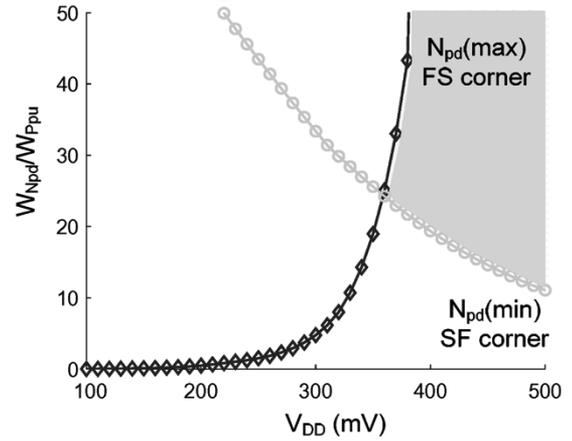


Fig. 9. The min-max curves for sizing N_{pd} at worst case process corners based on the read and write conditions and $P_{pu} = N_{ps}$.

A latch-based write scheme with C^2 MOS tristate inverters is a more robust design for subthreshold operation (Fig. 10). Analysis of the latch-based write scheme is performed at the worst case process corners to account for the leakage through the feedback tristate gate. The tristate latch memory cells shows functionality at process corners down to 215 mV.

B. Subthreshold Read Access

The read operation of the memory in subthreshold is a very challenging design problem. One challenge is bitline leakage, where the leakage through the pulldown devices causes the dynamic bitline to drop. Because clock speed is not the key metric in this application, a sense-amplifier-based read-bitline for fast read accesses is not needed.

Fig. 11 shows a 128-W single-ended read scheme using dynamic logic with minimum width pulldown devices. During the precharge phase, the precharge transistor is on and the bitline (\overline{RBL}) is charged to V_{DD} . During the evaluation phase \overline{RBL} evaluates to V_{DD} when the accessed cell (M_0) is '0'. In subthreshold, due to improper sizing, the charge stored in the dynamic bitline leaks away through all of the pulldown devices. To prevent the bitline from dropping, the precharge transistor (W_{pre}) is sized to offset the maximum leakage current through the pulldown devices. The maximum pull-down leakage occurs when the selected cell stores '0', and the remaining cells store '1' (e.g., $M_0 = 0, M_1 - M_{127} = 1$). However, when \overline{RBL} evaluates to '0', the I_{on} of the single pulldown path is much smaller than the I_{off} of the precharge transistor causing reduced I_{on}/I_{off} . The worst case pullup leakage occurs when the selected cell store a '1' and the remaining cells store '0' (e.g., $M_0 = 1, M_1 = M_{127} = 0$). The simulation in Fig. 11 shows that \overline{RBL} fails to evaluate to '0' due to the large precharge transistor. A similar tradeoff is seen using a pseudo-NMOS pull-up device.

A tristate-read scheme using static CMOS tristate cells is proposed (Fig. 12). Although the tristate gate alone operates at 100 mV, when connected in parallel, RBL has a low I_{on}/I_{off} ratio due to the leakage current through the parallel tristate

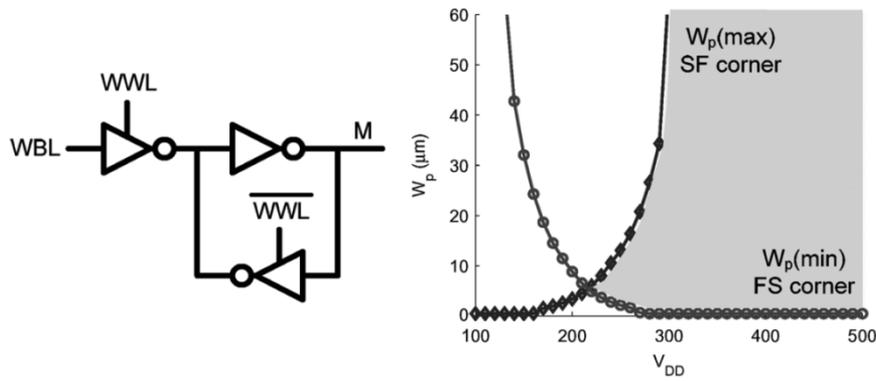


Fig. 10. Latch-based write access is analyzed at the worst case process corners.

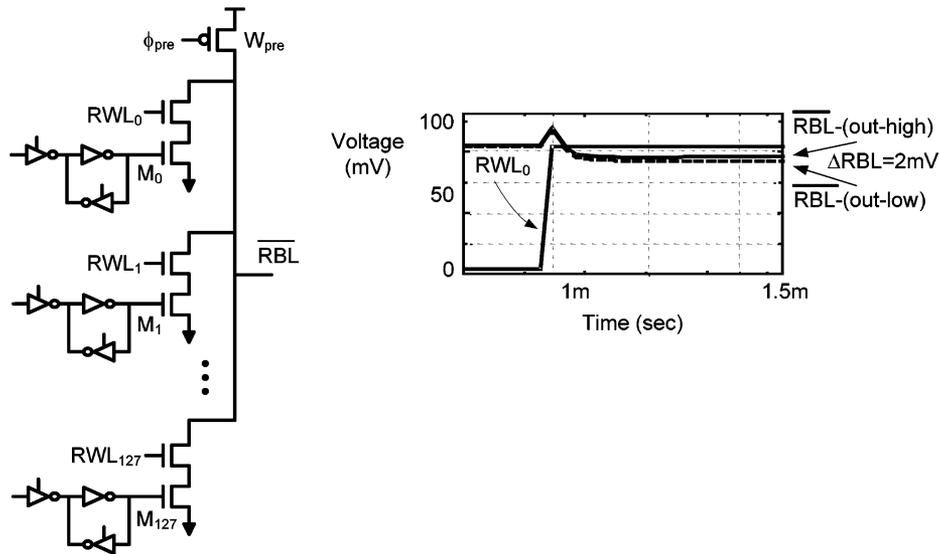


Fig. 11. Conventional precharge bitline read scheme. The simulation of RBL shows the effects of bitline leakage and W_{pre} size on RBL.

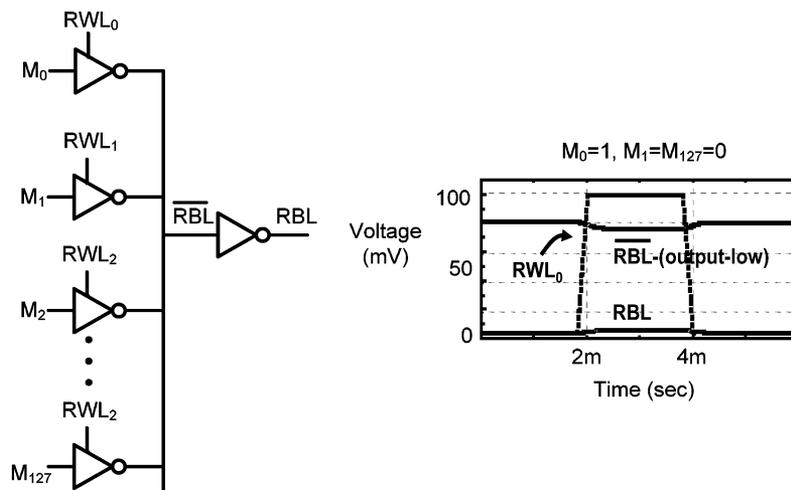


Fig. 12. The tristate-based read access also suffers from bitline leakage effects.

buffers. The simulation shows the RBL for the worst case output-high vector.

We propose a hierarchical-read-bitline that operates in sub-threshold and at worst case process corners (Fig. 13). The hierarchical-read-bitline segments the bitline by using a 2-to-1

mux-based approach. Segmenting reduces parallel leakage for each level of the hierarchy and the effect of process variations is mitigated. The muxes are designed to avoid stacked devices and sneak leakage paths by inserting inverters between each level of hierarchy. Fig. 13 shows a waveform of RBL

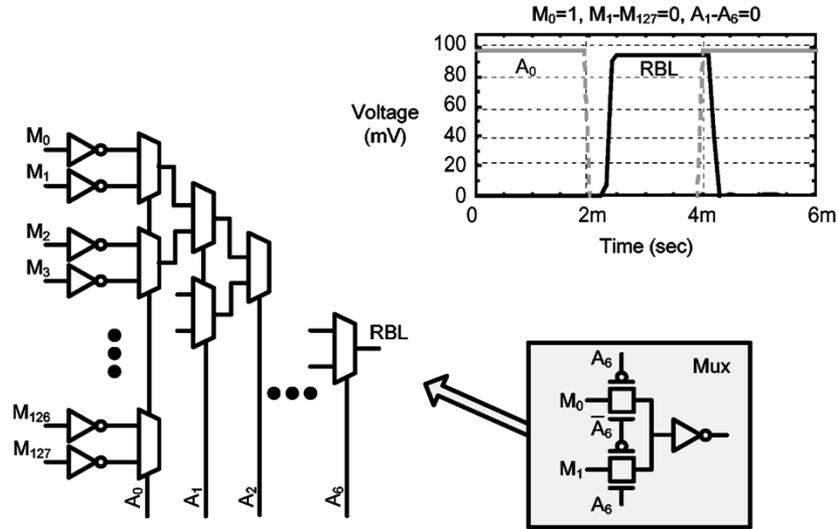


Fig. 13. The hierarchical-read-bitline does not suffer from parallel leakage effects. The mux is designed to reduce stacked devices and improve I_{on}/I_{off} at every stage of the hierarchy.

that operates at 100 mV for the typical transistor corner. It is able to achieve 10%–90% voltage swing for all input vectors including the worst case vectors of previous schemes.

To save area, the multiplexers (muxes) are daisy-chained and arrayed. A decoder generates the write wordline (WWL) which are routed to the bitline horizontally. The Read Addresses (A_0 – A_6) are routed vertically up to the selectors of the muxes. Memory generators were created using custom memory compilers. The compiler arrays the memory bit cells, daisy-chains the muxes, and creates the decoder logic for each data memory block. Eight custom $128 W \times 16 b$ blocks make up the scalable data memory in the FFT processor.

A similar hierarchical bitline design was used to create a $256 W \times 16 b$ ROM's compiler for the FFT twiddle factors. The twiddle factors are fixed complex coefficients that are used to transform between time to frequency domains. The difference between the ROM structures and the data memory is that the inputs to the muxes are either tied to V_{DD} or ground depending on the value stored in the ROM.

VI. RESULTS

The FFT processor was designed using a modified standard logic cell library, custom multiplier generators, and custom memory generators. The generators use specialized logic cells and ensure compact layout. Synthesis with the cell library is used to create functional logic blocks in the datapath, control logic and memories. All logic blocks, memory modules, and multipliers are verified using HSPICE at 100 mV for the TT process corner and at 250 mV for the FS and SF process corners. The place-and-route within each module and the top-level of the chip is performed using custom place-and-route scripts. The top-level system was verified at 1-V operation using a system-level simulator.

The FFT was fabricated in a standard $0.18\text{-}\mu\text{m}$ CMOS logic process with six metal layers. The FFT processor occupies $2.6 \times 2.1 \text{ mm}^2$ and contains 627 000 transistors. It is fully functional at 128, 256, 512, and 1024 FFT lengths, 8-b and

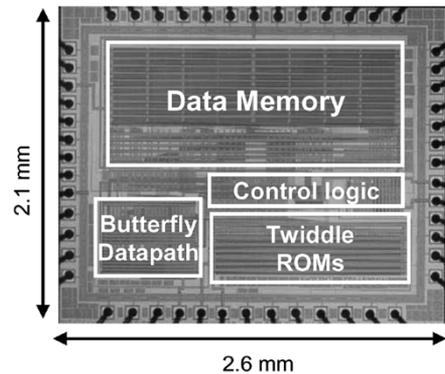


Fig. 14. Die photograph of the subthreshold FFT chip.

16-b precision, for voltage supplies from 180 to 900 mV, and for clock frequencies of 164 Hz to 6 MHz. Fig. 14 shows the die photograph of the subthreshold FFT chip. The chip's functional blocks (memory, butterfly datapath, twiddle ROMs, and control logic) are clearly delineated.

A printed circuit board (PCB) is used to test the functionality of the subthreshold FFT chip. The Textronix pattern generator has 2–5-V outputs which are level converted down to 500 mV using resistive dividers to interface with the subthreshold FFT chip. The outputs of the subthreshold FFT chip are level converted up to 5 V using off-the-shelf comparators.

The lowest voltage supply for correct operation is 180 mV with a clock speed of 164 Hz. The power dissipated at 180 mV is 90 nW for 16-b 1024-pt operation. Fig. 15 shows an oscilloscope plot of various outputs, (the clock, two output bits and the data-ready signal) from the FFT chip operating at 180 mV. However, the minimum supply voltage does not correspond to the optimal operating point which minimizes energy dissipation.

The minimum energy point of the 16-b 1024-pt FFT occurs at 350 mV. The power dissipated at the optimum is 600 nW at a clock frequency of 10 kHz and the energy dissipated is 155 nJ/FFT. Fig. 16 shows the measured and estimated energy dissipated for the 16-b 1024-pt FFT over $V_{DD} = 200\text{--}900 \text{ mV}$.

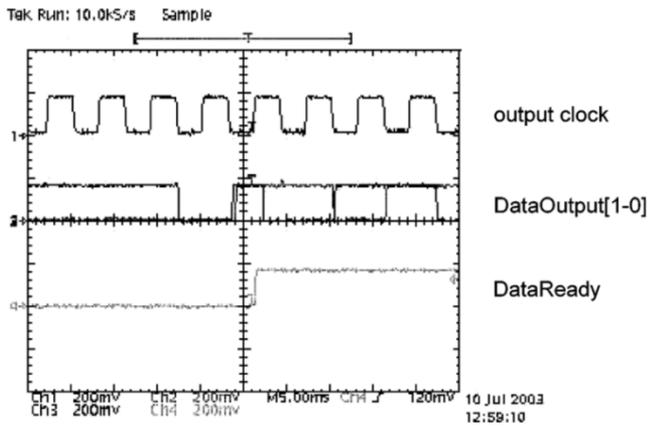


Fig. 15. Oscilloscope plot showing outputs from the FFT chip at 180-mV operation.

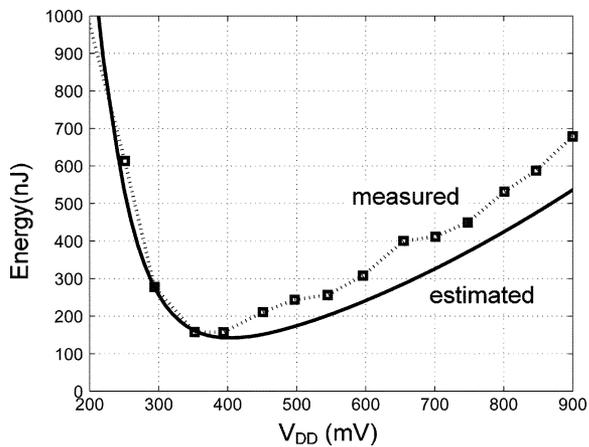


Fig. 16. Energy dissipation as a function of V_{DD} for the 16-b 1024-pt FFT. The optimal operating point for minimal energy dissipation is at $V_{DD} = 350$ mV.

The figure shows that our energy contour estimation technique predicts leakage energy well, but not switching energy. The energy is measured for the clock frequencies specified in Fig. 17.

Fig. 18 shows the energy dissipated for all operating points. The minimum energy point occurs at 350 mV for the 16-b processor and is at 400 mV for the 8-b processor. Eight-bit processing has a lower activity factor which reduces the ratio of switching energy to leakage energy. Thus, the optimum point is shifted toward larger minimum supply voltage [15]. At the optimal operating points for both 8-b and 16-b processing, the idle leakage energy is approximately 45% of the total energy dissipated.

We compared the minimum measured energy of the StrongARM SA-1100 [24], an Energy-Scalable ASIC [13], and the custom subthreshold chip for 16-b 1024-pt operation. The StrongARM SA-1100 is a low-power general-purpose microprocessor that can operate down to 0.85 V. The minimum energy of the StrongARM is measured while running the FFT subroutine. The ASIC implementation incorporates the energy-aware architectures described in this work. The ASIC FFT processor is designed using a commercial standard cell library and memory generators for a 0.18- μm process. When

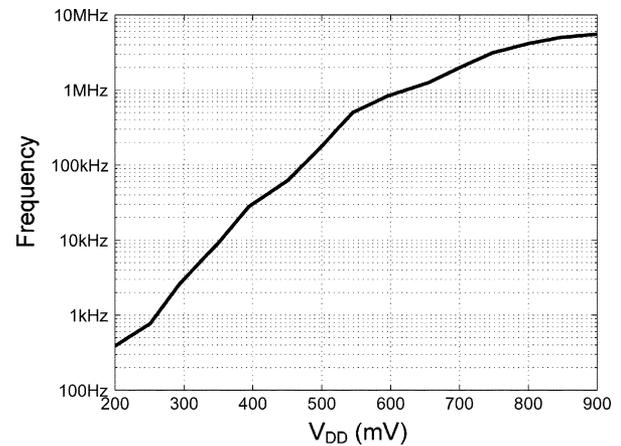
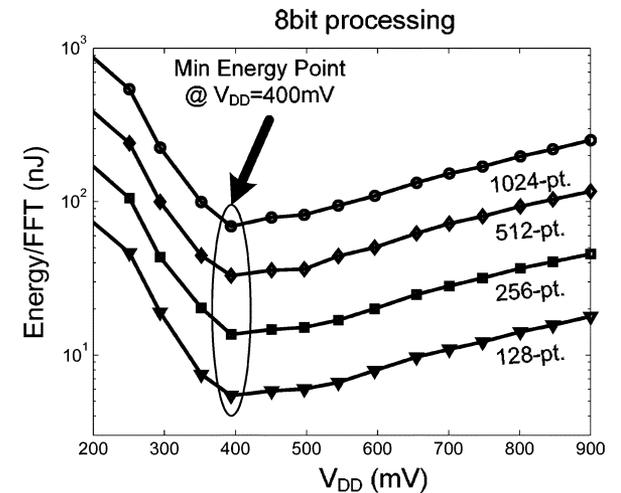
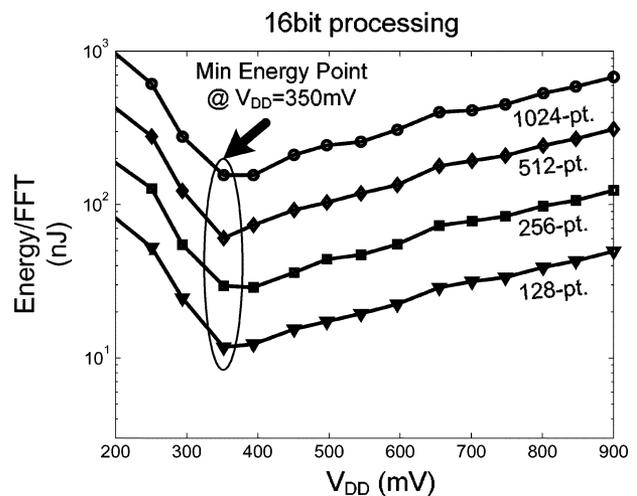


Fig. 17. Clock frequency as a function of V_{DD} .



(a)



(b)

Fig. 18. Energy as a function of V_{DD} for (a) 8-b and (b) 16-b processing. This shows the effect of activity factor on the minimum energy point.

comparing to the Subthreshold FFT, the subthreshold chip is 350 times more energy efficient than the low-power microprocessor implementation and eight times more energy efficient than the ASIC.

VII. CONCLUSION

In order for wireless sensor nodes to be self-powered, they must scavenge energy from the environment. Due to the requirements for energy scavenging, the key metric is minimizing energy dissipation rather than processor speed. For the FFT processor, energy models estimate the optimal supply voltage to be in subthreshold where the supply voltages are below the threshold voltages.

A minimum energy design methodology explores conventional logic and memories operating at both the minimum voltage and the minimum energy point. A subthreshold standard cell library avoids parallel leakage, stacked devices and sneak leakage effects. Memory generators are created using a hierarchical bitline to improve the I_{on}/I_{off} of the bitline. The subthreshold standard cell library and memory generators were used in a FFT processor design.

The subthreshold FFT is fabricated in a standard 0.18- μ m logic process without any additional process steps or biasing techniques. The FFT operated as low as 180 mV. It showed the optimum supply voltage to be at 350 mV with a clock frequency of 10 kHz, where it dissipated 155 nJ for a 16-b 1024-pt FFT.

REFERENCES

- [1] A. Chandrakasan, S. Sheng, and R. Brodersen, "Low-power digital CMOS design," *IEEE J. Solid-State Circuits*, vol. 27, no. 4, pp. 473–484, Apr. 1992.
- [2] J. Karn, R. Katz, and K. Pister, "Next century challenges: Mobile networking for smart dust," in *Proc. ACM MobiCom'99*, Aug. 1999, pp. 271–228.
- [3] G. Asada, M. Dong, T. S. Lin, F. Newberg, G. Pottie, and W. J. Kaiser, "Wireless integrated network sensors: Low power systems on a chip," in *Proc. ESSCIRC'98*, 1998, pp. 9–16.
- [4] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Proc. ACM MobiCom'99*, Aug. 1999, pp. 263–270.
- [5] J. Rabaey, J. Ammer, T. Karalal, S. Li, B. Otis, M. Sheets, and T. Tuan, "PicoRadios for wireless sensor networks: The next challenge in ultra-low-power design," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2002, pp. 200–201.
- [6] C. B. Williams and R. B. Yates, "Analysis of a micro-electric generator for microsystems," *Proc. Transducers'95/Eurosensors IX*, 1995.
- [7] D. Friedman, H. Heinrich, and D.-W. Duan, "A low-power CMOS integrated circuit for field-powered radio frequency identification tags," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 1997, pp. 294–295.
- [8] R. Amirtharajah and A. Chandrakasan, "Self-powered signal processing using vibration-based power generation," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 687–695, May 1998.
- [9] G. Franceschetti, A. Mazzeo, N. Mazzocca, V. Pascazio, and G. Schirrinzi, "An efficient SAR parallel processor," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 27, no. 2, pp. 343–353, Mar. 1991.
- [10] M. Hasan, T. Arslan, and J. S. Thompson, "A novel coefficient ordering based low power pipelined radix-4 FFT processor for wireless LAN applications," *IEEE Trans. Consum. Electron.*, vol. 49, no. 1, pp. 128–134, Feb. 2003.
- [11] B. M. Baas, "A low-power, high-performance, 1024-point FFT processor," *IEEE J. Solid-State Circuits*, vol. 34, no. 3, pp. 380–387, Mar. 1999.
- [12] M. Bhardwaj, R. Min, and A. Chandrakasan, "Quantifying and enhancing power-awareness of VLSI systems," *IEEE Trans. Very Large Scale Integration (VLSI) Syst.*, vol. 9, no. 6, pp. 757–772, Dec. 2001.
- [13] A. Wang and A. Chandrakasan, "Energy-aware architectures for a real-valued FFT implementation," in *Proc. Int. Symp. Low Power Electronics and Design*, Aug. 2003, pp. 360–365.
- [14] B. J. Sheu, D. L. Scharfetter, P.-K. Ko, and M.-C. Jeng, "BSIM: Berkeley short-channel IGFET model for MOS transistors," *IEEE J. Solid-State Circuits*, vol. 22, no. 4, pp. 558–566, Aug. 1987.
- [15] A. Wang, A. Chandrakasan, and S. Kosonocky, "Optimal supply and threshold scaling for subthreshold CMOS circuits," in *Proc. IEEE Annu. Int. Symp. VLSI (ISVLSI'02)*, Apr. 2002, pp. 7–11.
- [16] ITRS Roadmap [Online]. Available: <http://public.itrs.net>
- [17] E. A. Vittoz, "Future of analog in the VLSI environment," in *Proc. ISCAS*, 1990, pp. 1372–1375.
- [18] J. D. Meindl and R. M. Swanson, "Ion-implanted complementary MOS transistors in low-voltage circuits," *IEEE J. Solid-State Circuits*, vol. SSC-7, no. 2, pp. 146–153, Apr. 1972.
- [19] C. Kim, H. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integrated (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Dec. 2003.
- [20] J. Burr and J. Shott, "A 200 mV self-testing encoder/decoder using stanford ultra-low-power CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 1994, pp. 84–85.
- [21] J. Kao, M. Miyazaki, and A. Chandrakasan, "A 175-mV multiply-accumulate unit using an adaptive supply voltage and body bias architecture," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1545–1554, Nov. 2002.
- [22] M. Izumikawa, H. Igura, K. Furuta, H. Ito, H. Wakabayashi, K. Nakajima, T. Mogami, T. Horiuchi, and M. Yamashina, "A 0.25-mm CMOS 0.9-V 100-MHz DSP core," *IEEE J. Solid-State Circuits*, vol. 32, no. 1, pp. 52–61, Jan. 1997.
- [23] V. De, Y. Ye, A. Keshavarzi, S. Narendra, J. Kao, D. Somasekhar, R. Nair, and S. Borkar, "Techniques for leakage power reduction," in *Design of High-Performance Microprocessor Circuits*, A. Chandrakasan, W. Bowhill, and F. Fox, Eds. New York: IEEE Press, 2001, pp. 46–61.
- [24] *StrongARM SA-1100 Microprocessor for Portable Applications Brief Datasheet*, Intel Corp., Chandler, AZ, 1999.



Alice Wang (S'96–M'97) received the S.B., M.Eng., and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1997, 1998, and 2004, respectively.

Since January 2004, she has been with Texas Instruments, Dallas, TX, where she is a Senior IC Designer for the Chip Technology Center in the Wireless Division. Her research interests include ultra-low voltage CMOS circuits and systems and power management techniques for deep-submicron

technologies.

Dr. Wang is serving on the microprocessor forum committee for ISSCC 2005 and on the technical program committee for the 2004 International Symposium on Low Power Electronics and Design (ISLPED).



Anantha Chandrakasan (M'95–SM'01–F'04) 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 is currently a Professor of Electrical Engineering and Computer Science. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a co-author of *Low Power Digital CMOS Design* (Kluwer, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998) and *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000).

Dr. Chandrakasan has received several awards, including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, and the 1999 Design Automation Conference Design Contest Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design'98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Sub-committee Chair for ISSCC 2004. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the SSCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2005.