Exploring Variability and Performance in a Sub-200-mV Processor

Scott Hanson, *Member, IEEE*, Bo Zhai, Mingoo Seok, *Member, IEEE*, Brian Cline, Kevin Zhou, Meghna Singhal, Michael Minuth, Javin Olson, Leyla Nazhandali, Todd Austin, Dennis Sylvester, *Senior Member, IEEE*, and David Blaauw, *Member, IEEE*

***Abstract*—In this study, we explore the design of a subthreshold processor for use in ultra-low-energy sensor systems. We describe an 8-bit subthreshold processor that has been designed with energy efficiency as the primary constraint. The processor, which is functional below $V_{\rm dd} = 200$ mV, consumes only 3.5 pJ/inst at $V_{\rm dd} = 350$ mV and, under a reverse body bias, draws only 11 nW at $V_{\rm dd} = 160$ mV. Process and temperature variations in subthreshold circuits can cause dramatic fluctuations in performance and energy consumption and can lead to robustness problems. We investigate the use of body biasing to adapt to process and temperature variations. Test-chip measurements show that body biasing is particularly effective in subthreshold circuits and can eliminate performance variations with minimal energy penalties. Reduced performance is also problematic at low voltages, so we investigate global and local techniques for improving performance while maintaining energy efficiency.**

***Index Terms*—Low voltage, process variation, sensor network processing, subthreshold.**

I. Introduction

As RESEARCH in ultra-low-power circuit design advances, a vision of highly integrated mobile computing systems with lifetimes of the order of years is emerging. Such computing systems are attractive for biomedical implants, supply chain management, and environmental monitoring [1]. The energy consumption of these systems ultimately limits form factor, battery life, and complexity. It is therefore critical to develop circuits capable of performing complex tasks under stringent energy constraints. A number of low-power digital design techniques have been explored over the past several decades, but supply voltage scaling is generally shown to be the most effective technique due to the quadratic dependence of dynamic energy on the supply voltage $(V_{\rm dd})$.

Recent research has shown that minimum energy is typically achieved when $V_{\rm dd}$ enters the subthreshold region $(V_{\rm dd} < V_{\rm th})$ [2], [3]. Subthreshold circuits have been shown to be functional below 200 mV [4] and with energy consumption of the order of picojoules per instruction [5]. Recent work has also explored the challenges of ultra-low-voltage memory design [6]–[10]. However, a number of daunting challenges remain for subthreshold circuits. The most important concern is variability. Exponential sensitivities to $V_{\rm dd}, V_{\rm th}$, and temperature make even small variations problematic. Performance is also considerably degraded at low voltage since nodes are charged and discharged by weak inversion currents. The speeds of subthreshold digital circuits have typically been reported in the kHz and low-MHz ranges [4], [5]. To guarantee widespread adoption of subthreshold design, it will be necessary to address both of these issues.

In this study, we explore the subthreshold design space and address the variability and performance problems of low-voltage operation. We begin in Section II by describing an 8-bit processor that has been fabricated in a 0.13 μm technology [13]. The architecture is described in detail with emphasis placed on accommodations made for energy efficiency. Measurements show that the processor is functional below 200 mV and that the total energy consumption is only 3.5 pJ/instruction at $V_{\rm dd} = 350$ mV. With the application of a reverse body bias, the power consumption goes as low as 11 nW.

In Section III, we propose a body biasing strategy that takes advantage of the unique sensitivities of subthreshold operation. We contrast the body bias sensitivities of subthreshold circuits with those of super-threshold circuits $(V_{\rm dd} > V_{\rm th})$. Measurements of the subthreshold processor show that robustness at low voltages can be improved dramatically with the application of a body bias and that performance fluctuations induced by process and temperature variability can be eliminated with minimal energy penalties.

In Section IV, we explore techniques for improving performance in the subthreshold processor. We first compare body biasing and voltage scaling for improving performance globally. We then discuss sizing techniques for improving performance locally. At low voltages, gate-length sizing can give an exponential increase in drain current due to reverse short-channel effects (RSCE). Test-chip measurements show that gate-length sizing is superior to gate-width sizing for improving performance along timing critical paths.

II. Test-Chip Overview

A. Architecture

While our energy efficiency improvements are primarily derived from aggressive voltage scaling, architectural decisions can have a dramatic impact on the energy efficiency of a system. We have accordingly adopted a simple-processor architecture but have made a number of additions to enhance energy efficiency. A system-level diagram of the processor and CPU,

Manuscript received August 26, 2007; revised November 2, 2007.

S. Hanson, M. Seok, B. Cline, K. Zhou, M. Minuth, L. Nazhandali, T. Austin, D. Sylvester, and D. Blaauw are with the University of Michigan, Ann Arbor, MI 48109 USA (e-mail: hansons@umich.edu).

B. Zhai is with Advanced Micro Devices, Austin, TX 78741 USA.

M. Singhal and J. Olson are with Advanced Micro Devices, Boxborough MA 01719 USA.

Digital Object Identifier 10.1109/JSSC.2008.917505

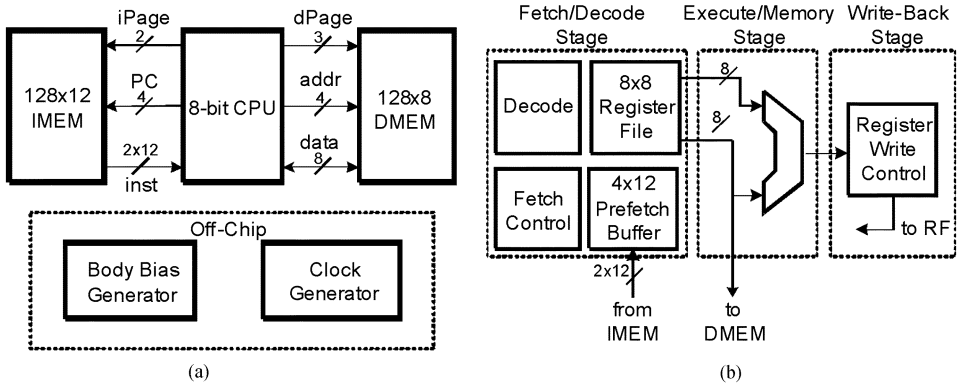

Fig. 1. (a) System-level diagram of the 8-bit subthreshold processor. (b) CPU implementation details.

which addresses a 1.5 kb instruction memory and a 1 kb data memory, is shown in Fig. 1(a), [12], [13].

To minimize both decoding complexity and memory footprint, we choose a RISC-style architecture with an instruction width of only 12 bits. As we will see in subsequent sections, the memory energy demands can dominate the total energy consumption of the system, so these decisions are extremely important. To further reduce the energy consumption of the memories, we divide both data and instruction memories into pages of 16 words each. A special instruction pre-decodes the upper bits of the memory address (*iPage* and *dPage* in Fig. 1(a)) and allows single cycle access to the contents of the specified page. Significant energy is saved when accessing multiple words within a page.

The CPU, shown in Fig. 1(b), is a three-stage pipeline with 8-bit data width. A highly pipelined design ensures that the majority of the logic is active throughout the clock cycle, thus minimizing time spent idly leaking. However, pipelining also requires additional sequential elements, which can be energy hungry. A three-stage pipeline is attractive since it balances these competing trends [12]. We choose an 8-bit data width since the upper bits in a 16-bit or 32-bit processor would be idle for much of the computation in simple sensor applications, leading to an unnecessary leakage overhead. Simple "taken" branch speculation has been implemented to reduce branch-related stalling in the CPU. A small four-entry prefetch buffer helps facilitate this branch prediction.

B. Implementation

The 8-bit processor was implemented using a conventional digital synthesis and place-and-route design flow. All circuits were designed with the goal of maximizing robustness at low voltage. For example, the synthesis library included a limited subset of CMOS gates with a maximum fan-in of two. Gates with large fan-ins have been shown in previous work to have reduced noise margins at low voltage [4], [14], [15]. The instruction memory, data memory, and register file were implemented using a robust latch-based memory with a mux-based read-out structure [4]. While this memory structure is large and energy inefficient, it helps us to reliably explore the low-voltage domain. Recently, several authors have proposed more compact low-voltage memories that are promising for future subthreshold development [6]–[10].

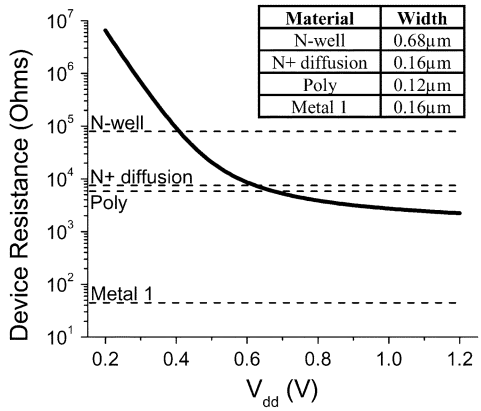

Fig. 2. Effective NFET resistance as a function of $V_{\rm dd}$. The resistances of wires of several alternative materials are included for reference (with 100 μm length and widths from inset).

The topic of physical design in low-voltage circuits has been overlooked in previous work. However, it is important to observe that interconnect *RC* delay ($\sim 0.38 \cdot R_{\rm wire} \cdot C_{\rm wire}$) is only a function of materials and circuit geometry and does not scale with $V_{\rm dd}$. Subthreshold current, on the other hand, is exponentially related to $V_{\rm dd}$. Consequently, wire resistance (and wire *RC* delay) becomes insignificant compared with device resistance at low voltage. Fig. 2 shows the effective resistance of an NFET device as a function of $V_{\rm dd}$. The resistances of 100 μm minimum width wires of various materials have been included for reference. At $V_{\rm dd} = 300$ mV, the device resistance is $>10\,000$ times greater than that of a 100 μm wire in the first metal layer.

The reduced importance of *RC* delay has several important implications. Minimum-width wires can be used for any interconnect with no penalty. We have leveraged this in our design by using minimum-width metal for clock and power routing. This opens considerable routing area and reduces energy consumption in our clock distribution network. Interestingly, Fig. 2 suggests that density could be further improved by shifting some of the routing to the poly and diffusion layers. In addition to thinner wire routes, the reduced importance of *RC* delay permits the use of a greatly simplified clock distribution network. The large capacitance of the clock network can be treated as a grid driven by a single level of clock drivers. This reduces design complexity

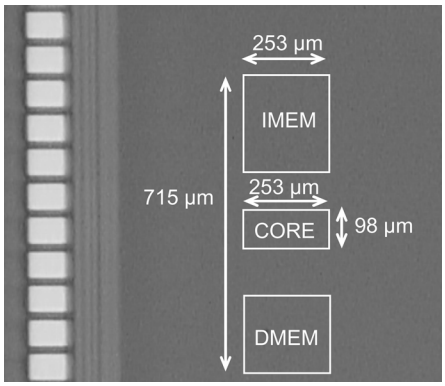

Fig. 3. 8-bit subthreshold processor was fabricated in a 0.13 μm technology.

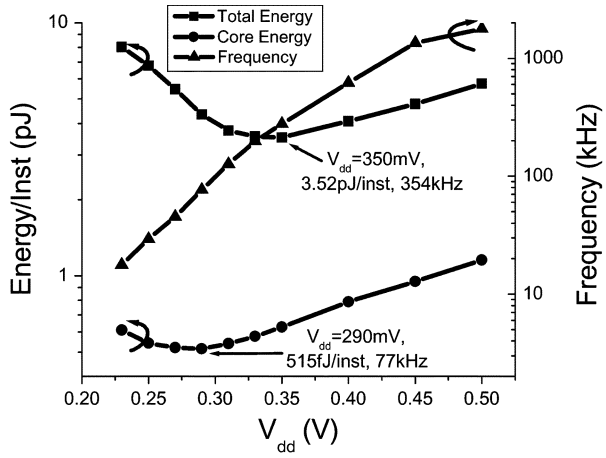

Fig. 4. Frequency and energy measurements for a typical die as functions of $V_{\rm dd}$.

and minimizes skew induced by process, voltage and temperature variations (PVT), which can be severe in subthreshold circuits. In our design, we used a single clock buffer for all pipeline registers.

We have fabricated the proposed processor in a 0.13 μm technology with $V_{\rm th} \sim 400$ mV at $V_{\rm ds} = 50$ mV. The die photograph is shown in Fig. 3. The full processor has a footprint of 253 μm $\times$ 715 μ m, while the core has an area of only 253 μm $\times$ 98 μm.

C. Energy and Frequency Measurements

Energy and maximum operating frequency measurements for the processor are shown for a typical die in Fig. 4. Both were measured for a simple arithmetic program that tests a wide range of instructions. The average current demand for the core and memories was measured over many program iterations using a high-precision electrometer. As predicted by [2], [3], energy reaches a minimum due to increased leakage energy at low $V_{\rm dd}$. The processor achieves a minimum of 3.5 pJ/inst at $V_{\rm dd} =$ 350 mV with a frequency of 354 kHz. The core (without memories, register file, or prefetch buffer) reaches a minimum of 515 fJ/inst at $V_{\rm dd} = 290$ mV. The data memory, instruction memory, prefetch buffer, and register file consume $>$70% of the total energy, which is not surprising given that they are the most area-intensive circuits. The processor remains functional down to $\sim$210 mV without a body bias but can function below 200 mV with the proper body bias (to be discussed in the next section).

In power-limited applications, such as those that scavenge ambient energy [1], a reverse body bias may be applied to minimize power consumption (as opposed to energy consumption). Under a reverse bias of 300 mV, the processor draws 11 nW at $V_{\rm dd} = 160$ mV with a maximum frequency of 710 Hz. The core alone draws only 735 pW. We focus on energy minimization for the remainder of this paper since it is more relevant for battery-powered applications.

III. Body Biasing for Variability Control

Process variability has the potential to be a crippling problem in subthreshold circuits. Variation is typically classified as random within-die, correlated within-die, or die-to-die. Random within-die variations in $V_{\rm th}$ and gate length are becoming particularly problematic in scaled devices [16]. For example, random $V_{\rm th}$ variations induced by random dopant fluctuations (RDFs) have an inverse dependence on the gate area, giving $3\sigma V_{\rm th}$ variations as high as 58 mV at the 65 nm technology node [14]. Given an inverse subthreshold slope of 85 mV/dec, such large $V_{\rm th}$ variation can yield $3\sigma/\mu$current deviations as high as 4.8$\times$ the nominal. Correlated variations are also a significant problem at low voltage. In this work, we group correlated within-die and die-to-die variations together under the name global variation due to the small size of our subthreshold processor. Global variations can be either static process variations (e.g., $V_{\rm th}$, gate length, or gate oxide thickness) or temporal variations (e.g., temperature or $V_{\rm dd}$) [16]. Subthreshold operation is dominated by an exponential dependence on $V_{\rm th}$, so global variations in $V_{\rm th}$ due to doping fluctuations or those induced by gate length, gate oxide thickness, and temperature variations are especially concerning.

Variations in subthreshold circuits can lead to two types of failure: functional failure and parametric failure. Functional failure primarily occurs in SRAM arrays as a result of random variations. Due to the extensive use of minimum sized devices, SRAM arrays are particularly susceptible to RDF-induced $V_{\rm th}$ variations. Large strength mismatch can lead to widespread read upsets. Alternative SRAM bitcell architectures [6]–[10] will be necessary to guarantee sufficient read margins in scaled subthreshold processors. Parametric failure (i.e., violated delay constraints) can result from both random and global variations. Without special attention to minimizing delay variability, subthreshold design at 65 nm and beyond will require large timing margins. For example, it was shown in [26] that a chain of ten inverters in a 65 nm technology at $V_{\rm dd} = 300$ mV requires a clock period that is 3.3$\times$ slower than the nominal delay of the chain to avoid late-mode timing errors. Such large delay variations result in nonnegligible energy variations (another possible parametric failure) in subthreshold circuits as well [26]. Delay fluctuations (and, consequently, energy variations) due to RDF can be minimized by increasing gate area and using deeper logic between pipeline registers [17]. However, global variations do not benefit from these solutions and require chip-level solutions. In the remainder of this section, we explore the use of body biasing to address delay and energy fluctuations due to global process and temperature variations. While

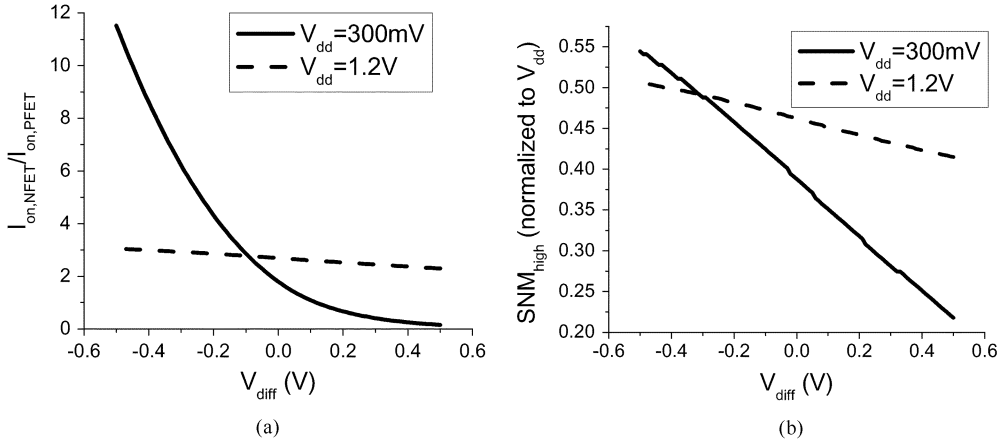

Fig. 5. (a) Simulated ratio of NFET on-current $(I_{\rm on,NFET})$ to PFET on-current $(I_{\rm on,PFET})$ at two voltages. (b) Simulated high static noise margins (SNM) at two voltages for an inverter with $W_{\rm PFET} = 2 \cdot W_{\rm NFET}$.

variability is greatest in aggressively scaled technologies, the 0.13 μm processor presented in this work shows sufficiently large variability to illustrate the basic body bias sensitivities of subthreshold circuits.

A. Body Biasing in Subthreshold Circuits

Body biasing, which effectively skews $V_{\rm th}$, has been proposed for global $V_{\rm th}$ compensation in the past. The authors of [18] designed a multiply-accumulate unit that used adaptive supply voltage and body bias to minimize power in super-threshold circuits. Body biasing is a particularly effective technique in the subthreshold regime due to the exponential dependence of subthreshold current on body bias. The use of body biasing in subthreshold circuits was briefly explored in [19], but little attention was given to how the body bias should be selected and only limited measurements were presented. The authors of [20] showed that correct operation can be achieved with $V_{\rm dd}$ as low as 100 mV by tuning body biases to match PFET and NFET leakages. In this study, we extend these early studies to develop a comprehensive body-biasing strategy that accounts for the unique sensitivities of subthreshold circuits. We also present detailed measurements of 20 measured dies to confirm the observed trends.

Throughout this section, we refer to two terms relevant to body biasing: *offset* and *differential*. The differential $(V_{\rm diff})$ is the relative difference between the PFET and NFET body biases (i.e., $V_{\rm diff} = (V_{\rm dd} - V_{b,PFET}) - V_{b,\rm NFET}$), which may be tuned to skew the relative strengths of PFET and NFET devices, as shown in Fig. 5(a) for two different supply voltages. We have chosen $V_{\rm dd} = 300$ mV as a representative subthreshold voltage since it lies close to the minimum energy $V_{\rm dd}$ for our processor. As expected, the sensitivity to $V_{\rm diff}$ is particularly high at $V_{\rm dd} =$ 300 mV due to the exponential dependence of current on body bias. Balanced static noise margins (SNM) depend on matching between PFET and NFET strengths, so we can use $V_{\rm diff}$ to compensate for global $V_{\rm th}$ skew between PFET and NFET devices and maximize noise margins. Fig. 5(b) shows the high static noise margins in a CMOS inverter with $W_{\rm PFET} = 2 \cdot W_{\rm NFET}$ at $V_{\rm dd} = 300$ mV and $V_{\rm dd} = 1.2$ V. At $V_{\rm dd} = 300$ mV, the high and low noise margins are balanced at $V_{\rm diff} = 70$ mV.

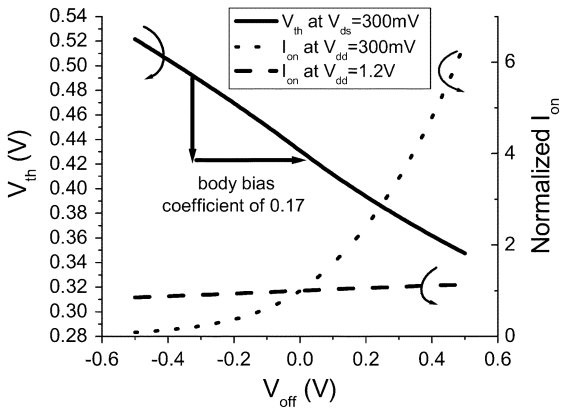

Fig. 6. Simulated NFET $V_{\rm th}$ and $I_{\rm on}$ as functions of $V_{\rm off}$

The offset $(V_{\rm off})$ is the shift in both the PFET and NFET biases relative to the ground voltage (i.e., $V_{\rm off} = V_{b,NFET}$). A positive offset indicates a forward body bias (which reduces $V_{\rm th}$) while a negative offset indicates a reverse body bias (which increases $V_{\rm th}$). As shown in Fig. 6 for an NFET, $V_{\rm th}$ changes with $V_{\rm off}$ at a rate of 170 mV/V. Fig. 6 also shows how NFET on-current $(I_{\rm on})$ changes with $V_{\rm off}$ at two different supply voltages. The increase in $I_{\rm on}$ with $V_{\rm off}$ is far more dramatic at subthreshold voltages (300 mV) than at super-threshold voltages (1.2 V).

It was shown in [2] and [3] that energy is independent of $V_{\rm th}$ as long as the circuit remains in the subthreshold regime. However, these derivations represent PFET and NFET devices with a single composite current expression, so they do not capture the energy dependence on $V_{\rm th}$ mismatch between NFET and PFET devices. Dynamic energy $(C \cdot V_{\rm dd}^2)$ is independent of PFET/NFET matching, but leakage energy can be separated into PFET-dependent and NFET-dependent components. Consider the leakage energy for a chain of identical inverters, where $I_{\rm off,total}$ is the total leakage current, $t_{\rm total}$ is the delay of the inverter chain, $I_{{\rm off},N}$ and $I_{{\rm off},P}$ are the cumulative leakages

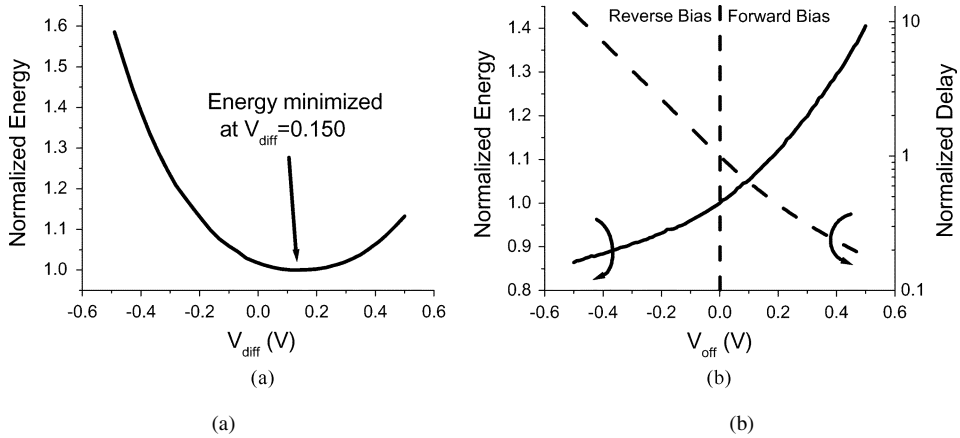

(a) (b)

Fig. 7. (a) Simulated energy consumption for a chain of 30 inverters at $V_{\rm dd} = 300$ mV as a function of $V_{\rm diff}$. (b) Simulated energy and delay for the same inverter chain as functions of $V_{\rm off}$

through NFET and PFET stacks, $t_{p,N}$ and $t_{p,P}$ are the cumulative delays through NFET and PFET stacks, and k accounts for delay degradation due to input slew as

$$\begin{aligned} E_{\rm leak} &= I_{\rm off,total} \cdot t_{\rm total} \cdot V_{\rm dd} \\ &= (I_{{\rm off},N} + I_{{\rm off},P}) \cdot (t_{p,N} + t_{p,P}) \cdot V_{\rm dd} \\ &= (I_{{\rm off},N} + I_{{\rm off},P}) \\ &\quad \cdot \left(\frac{k \cdot C \cdot V_{\rm dd}}{I_{{\rm off},N} \cdot e^{\frac{V_{\rm dd}}{m_N \cdot v_T}}} + \frac{k \cdot C \cdot V_{\rm dd}}{I_{{\rm off},P} \cdot e^{\frac{V_{\rm dd}}{m_P \cdot v_T}}} \right) \cdot V_{\rm dd}. \end{aligned} \tag{1}$$

If we make the simplification that the NFET and PFET subthreshold slope factors are identical (i.e., $m_p = m_n = m$), then we can rearrange the leakage expression to highlight the dependence on NFET/PFET matching and take the derivative with respect to $I_{{\rm off},N}/I_{{\rm off},P}$ as

$$\begin{aligned} \frac{\partial E_{\rm leak}}{\partial \frac{I_{{\rm off},N}}{I_{{\rm off},P}}} &= \left(1 - \left(\frac{I_{{\rm off},P}}{I_{{\rm off},N}} \right)^2 \right) \cdot k \cdot C \cdot V_{\rm dd}^2 \cdot e^{-\frac{V_{\rm dd}}{m \cdot v_T}} \\ 0 &= \left(1 - \left(\frac{I_{{\rm off},P}}{I_{{\rm off},N}} \right)^2 \right). \end{aligned} \tag{2}$$

Equation (2) shows that NFET and PFET off-currents should be equal for minimum energy, which is the same condition that gives us balanced noise margins. For confirmation, we simulate a chain of 30 inverters switching with an activity rate of 0.2 at $V_{\rm dd} = 300$ mV. Inverter chains have been used extensively in previous work to explore the basic sensitivities of subthreshold circuits [2], [3] and have proven to be good indicators of the trends observed in more complex circuits. Fig. 7(a) shows that the energy consumed per cycle (the time it takes to propagate a single switching operation) for the inverter chain is minimized at $V_{\rm diff} = 150$ mV, which matches well with the $V_{\rm diff}$ value that balances high and low noise margins (70 mV).

To explore the sensitivity of energy to $V_{\rm off}$, we again simulate a chain of 30 inverters with a switching activity of 0.2 at $V_{\rm dd} = 300$ mV over a range of $V_{\rm off}$ values. Fig. 7(b) shows the energy/cycle and the delay of the inverter chain. With a negative $V_{\rm off}$ (a reverse body bias), energy actually decreases. This seems to contradict the conclusion in [2] and [3] that energy does not depend on $V_{\rm th}$ in the subthreshold regime, but we make the added observation that inverse subthreshold slope, S_S, and the switched capacitance depend on the body bias. When a reverse body bias is applied, the depletion capacitance C_d reduces and yields improved S_S as follows:

$$S_{S.} = 2.3 \cdot v_T \cdot m = 2.3 \cdot v_T \cdot \left(1 + \frac{C_d}{C_{\rm ox}} \right). \tag{3}$$

Leakage energy, which is exponentially dependent on S_S through m, reduces with improved S_S. Note also that junction capacitance reduces under a reverse body bias, giving a reduction in switching energy (which was assumed to be constant in the previous discussion).

For positive $V_{\rm off}$ (a forward body bias), the delay of the inverter chain decreases quickly, but the performance improvement comes with a large energy penalty. The observed increase in energy is partially a result of degraded S_S and increased junction capacitance at forward body biases. Additionally, $V_{\rm th}$ reduces with a forward body bias and pushes the inverter chain into the near-threshold and super-threshold regimes. Outside of the subthreshold regime, the insensitivity of energy to $V_{\rm th}$ no longer holds [2], [3].

Given the observed sensitivity of subthreshold circuits to $V_{\rm diff}$ and $V_{\rm off}$, an effective body biasing strategy is clear. $V_{\rm diff}$ should first be tuned to achieve maximum noise margins and minimum energy. $V_{\rm off}$ can then be used to target a desired performance with only minimal energy consequences.

B. Body Bias Measurements

The processor described in Section II has been tested to verify our proposed body-biasing strategy. Body biases were routed as normal signal nets using minimum-width wires. External $V_{\rm dd}$, body bias, and clock generation were used (Fig. 1) to enable a fine-grained exploration of the energy-delay space. The reported energy numbers only include current drawn from the $V_{\rm dd}$ generator since simulations show that body current is several orders of magnitude smaller than subthreshold current for the body bias range studied. We also quantify the energy and delay benefits of body biasing without considering the costs of $V_{\rm dd}$ and body

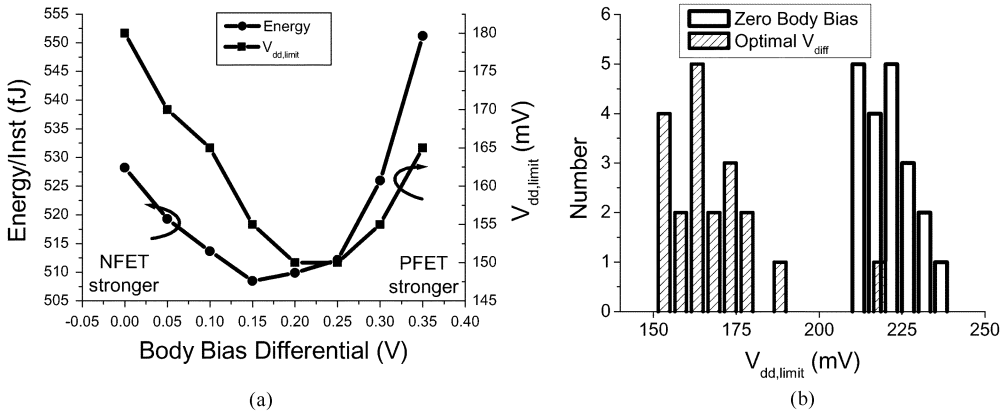

Fig. 8. (a) Energy and $V_{\rm dd,limit}$ as functions of $V_{\rm diff}$ for a typical die (b) $V_{\rm dd,limit}$ distribution for 20 dies with and without body biasing. The mean $V_{\rm dd,limit}$ reduces from 221 mV to 168 mV, a 24% improvement.

bias regulation. Past work has explored the efficient generation and regulation of both $V_{\rm dd}$ [25] and body bias [19], [20].

We first verify the observation that $V_{\rm diff}$ can be used to match PFET and NFET devices to maximize noise margins. Since noise margins are not readily measured for a processor, we use the minimum functional voltage $V_{\rm dd,limit}$ as a measure of robustness. The value of $V_{\rm dd,limit}$ is extremely sensitive to PFET/NFET matching and is therefore a useful robustness metric [21]. Fig. 8(a) shows that $V_{\rm dd,limit}$ can be minimized (and noise margins can be maximized) by tuning $V_{\rm diff}$. The energy consumption for the core (without memories, register file, or prefetch buffer) at $V_{\rm dd} = 300$ mV is shown for the same processor in Fig. 8(a). Energy consumption and $V_{\rm dd,limit}$ are minimized at nearly the same value of $V_{\rm diff}$, thus confirming our simulation-based observations. By selecting the optimal value of $V_{\rm diff}$ for each of 20 measured dies (mean of 140 mV across 20 dies), we find that the mean value of $V_{\rm dd,limit}$ reduces by 24% as compared with the case with zero body bias, as shown in Fig. 8(b).

Fig. 9 confirms for a typical die at $V_{\rm dd} = 300$ mV that the tuning of $V_{\rm off}$ may be used to achieve an excellent energy-delay tradeoff. Between $V_{\rm off}$ values of -400 mV and -100 mV, delay improves by 3.6× while energy varies by only 1%. Fig. 9 also shows that the energy–delay tradeoff begins to degrade with a forward body bias $(V_{\rm off} > 0)$, which is consistent with simulation-based observations in the previous subsection.

With proper selection of $V_{\rm diff}$ and $V_{\rm off}$, we can align all dies to a desired performance with limited energy penalties and near-maximum noise margins. While we cannot make statistically significant claims about the ability of body biasing to reduce performance variability in subthreshold circuits, it is informative to look at the sensitivities of the 20 measured dies to body bias. To explore this, we measure the processor under four different scenarios. In Case 1, body biases are tied to the appropriate $V_{\rm dd}$ and $V_{\rm ss}$ rails (zero body bias). In Cases 2–4, the energy-optimal value of $V_{\rm diff}$ is applied, and $V_{\rm off}$ is chosen with 5 mV resolution to meet frequency constraints of 66 kHz (worst case frequency in Case 1), 100 kHz, and 160 kHz. The energy and frequency spreads for each of these cases are shown for 20

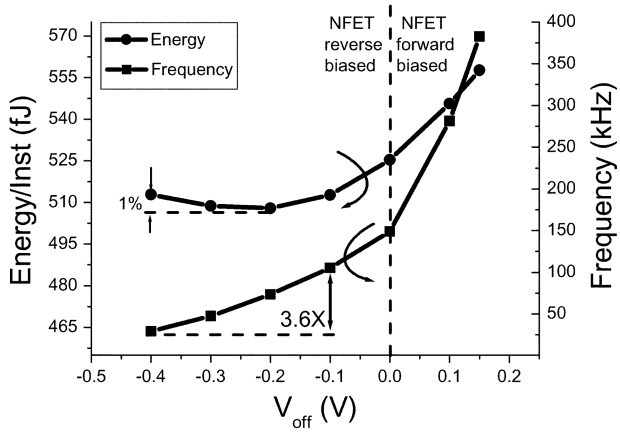

Fig. 9. Energy and frequency as functions of body-bias offset for a typical die.

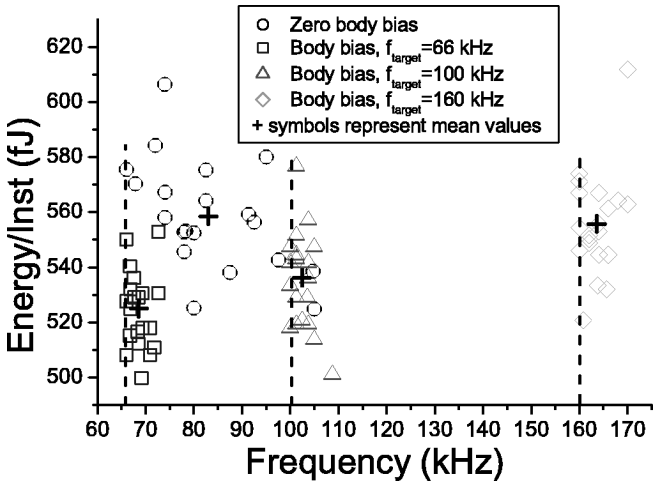

Fig. 10. Energy and frequency distributions for 20 dies measured at $V_{\rm dd} = 300$ mV.

dies measured at $V_{\rm dd} = 300$ mV in Fig. 10. Table I summarizes the data from Cases 1–4 when all dies run exactly at the target frequency (66, 100, or 160 kHz). A comparison of Cases 1 and 2 in Table I shows that mean energy is reduced by ∼10%

TABLE I
ENERGY STATISTICS FOR 20 DIES AT $V_{\rm dd} = 300$ mV WHEN FREQUENCY AND BODY BIAS ARE FIXED AT THE LISTED VALUES

	Frequency (kHz)	Mean V_{diff} (mV)	Mean V_{off} (mV)	μ_{energy} (fJ)	σ_{energy} (fJ)
Case 1: Zero body bias	66	--	--	588	22
Case 2: Variable body bias	66	140	-190	528	14
Case 3: Variable body bias	100	140	-81	538	17
Case 4: Variable body bias	160	140	41	559	20

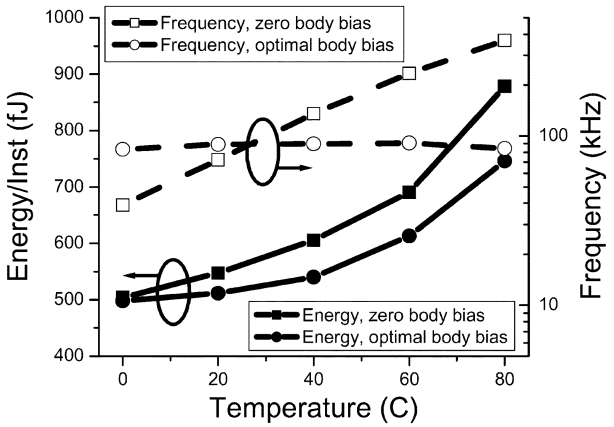

Fig. 11. Temperature sensitivity of energy and frequency for a typical die at $V_{\rm dd} = 300$ mV.

with frequency fixed at 66 kHz. A comparison of Cases 1 and 4 shows that a 2.4× increase in frequency can be achieved while also achieving a 5% mean energy improvement. This excellent energy–delay tradeoff makes body biasing extremely attractive for adaptive subthreshold systems.

The favorable energy–delay tradeoff achieved using body biasing can be extended to compensate for temperature variations. Temperature compensation has been demonstrated in the past for subthreshold circuits using simple temperature sensitive bias generation [19]; we explore the benefits of compensation more thoroughly in Fig. 11, which shows the temperature dependence of energy and performance for a typical die at $V_{\rm dd} = 300$ mV. Without body biasing, the frequency of the chip increases by $\sim$10× between $T = 0\,^{\circ}$C and $T = 80\,^{\circ}$C. For a fixed value of $V_{\rm diff}, V_{\rm off}$ can be tuned to maintain a constant frequency as shown in Fig. 11. For this particular die, $V_{\rm off}$ changes by 620 mV between $T = 0\,^{\circ}$C and $T = 80\,^{\circ}$C to maintain constant performance.

IV. IMPROVING PERFORMANCE

On-currents in the subthreshold regime can be >5 orders of magnitude lower than super-threshold on-currents, so reduced performance is inevitable. Performance is only a secondary concern in sensor network processing, but improved performance is necessary to make subthreshold operation viable in the embedded and high-performance application spaces. Here, we begin by comparing voltage scaling and body biasing for improving performance globally. We also look at the use of gate-length sizing to achieve local performance improvements.

A. Improving Global Performance

At the block level, body biasing and voltage scaling can both be used to achieve exponential improvements in performance. We are interested in determining which technique gives the better energy–delay tradeoff for subthreshold circuits. To do so, it is necessary to understand the energy implications of body biasing and voltage scaling. Consider a simple chain of inverters operating at a subthreshold voltage $V_{\rm dd,init}$ with zero body bias. A wide range of target frequencies can be achieved either by changing $V_{\rm dd,init}$ (voltage scaling) or by changing $V_{\rm off}$ (body biasing). The energy consumption of the inverter chain may be modeled as the sum of dynamic energy and leakage energy, where α is the switching activity, t is the maximum delay, and all other quantities are as defined previously as

$$E_{\rm total} = E_{\rm dyn} + E_{\rm leak} = C \cdot V_{\rm dd}^2 \cdot \alpha + I_{\rm off} \cdot t \cdot V_{\rm dd}. \quad (4)$$

The relative energy efficiencies of $V_{\rm dd}$ scaling and body biasing are strong functions of α. To illustrate this, we consider limiting behavior. In the case of very high switching activity, dynamic energy is dominant ($E_{\rm dyn} \gg E_{\rm leak}$), which is shown as

$$E_{\rm total} \approx E_{\rm dyn} = C \cdot V_{\rm dd}^2 \cdot \alpha. \quad (5)$$

In this limit, energy has a quadratic dependence on $V_{\rm dd}$ and is, to first order, independent of $V_{\rm off}$. For low target frequencies, we should reduce $V_{\rm dd}$ or apply a reverse body bias ($V_{\rm off} < 0$). A reduction in $V_{\rm dd}$ will yield quadratic energy reductions, while the application of a reverse bias will have no effect on energy. Voltage scaling is therefore more energy efficient for low target performance. Conversely, energy increases quadratically with $V_{\rm dd}$, so body biasing is more energy efficient when the target frequency is high.

In the case of very low switching activity, leakage energy is dominant ($E_{\rm leak} >> E_{\rm dyn}$), shown as

$$E_{\rm total} \approx E_{\rm leak} = I_{\rm off} \cdot \frac{C \cdot V_{\rm dd}}{I_{\rm off} \cdot e^{\frac{V_{\rm dd}}{m \cdot v_T}}} = \frac{C \cdot V_{\rm dd}}{e^{\frac{V_{\rm dd}}{m \cdot v_T}}}. \quad (6)$$

For $V_{\rm dd} > m{\cdot}v_T$ in the subthreshold and near-threshold regions, leakage energy per cycle increases as $V_{\rm dd}$ decreases [2], [3]. Energy therefore increases when $V_{\rm dd}$ is reduced to meet low-frequency targets but reduces when $V_{\rm dd}$ is increased to meet high target frequencies. Due to the m dependence of leakage energy, energy reduces with the application of a reverse body bias and increases with the application of a forward body bias. In this case, $V_{\rm dd}$ scaling is more energy efficient for high target frequencies while body biasing is more energy efficient for low target frequencies. Note that our observations are only valid for subthreshold circuits. Outside of the subthreshold region, delay

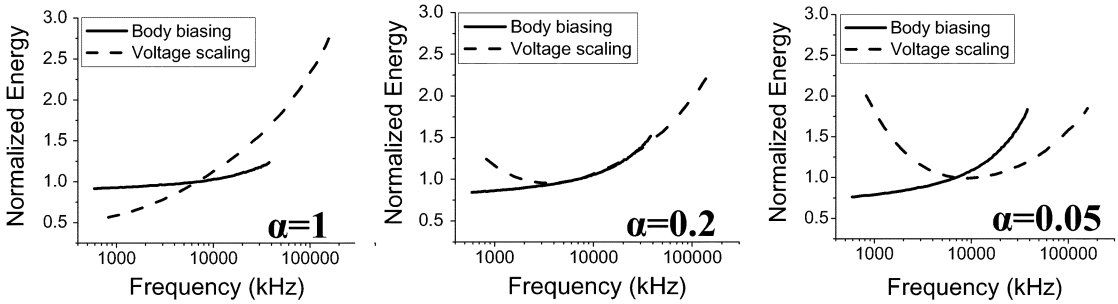

Fig. 12. Simulated energy and frequency for an inverter chain subjected to voltage scaling and body biasing. Data is plotted for switching activities of 1, 0.2, and 0.05.

is no longer exponentially dependent on $V_{\rm dd}$ and $V_{\rm th}$, and the energy–performance tradeoffs change dramatically.

Neither of the two limits considered reflects actual circuit behavior since dynamic energy and leakage energy are comparable in a typical subthreshold circuit [2]. For a more realistic comparison of voltage scaling and body biasing, we simulate a chain of 30 inverters with switching activities of 0.05, 0.2, and 1. We select a nominal $V_{\rm dd}$ of 300 mV. For voltage scaling data, we sweep $V_{\rm dd}$ from 200 to 500 mV. For body biasing data, we sweep $V_{\rm off}$ from -500 to 500 mV. The data in Fig. 12 for switching activities of 1 and 0.05 confirm our previous observations about the high-activity and low-activity limits, respectively. Fig. 12(b) shows the energy characteristic for a more realistic switching activity of 0.2. In this scenario, body biasing and voltage scaling give similar energy-performance tradeoffs for much of the performance range, suggesting that either body biasing or voltage scaling could be used for minimum energy in a typical circuit.

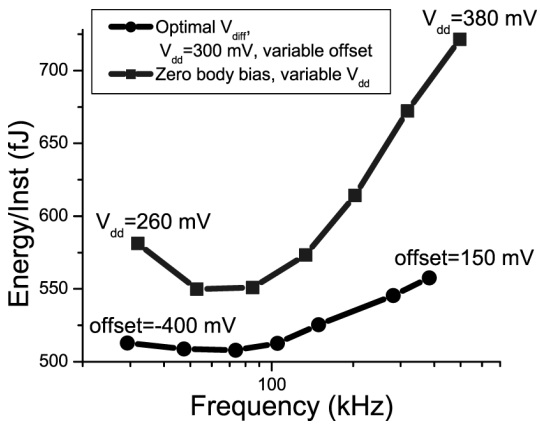

Fig. 13. Comparison of energy and frequency measurements for variable body bias and variable $V_{\rm dd}$ systems.

B. Global Performance Measurements

To verify the trends observed in the previous subsection, we measure the energy and performance of the core (without memories, register file, or prefetch buffer) over a range of supply voltages and body biases. To evaluate voltage scaling, we fix the body biases at zero and sweep $V_{\rm dd}$ from 260 to 380 mV. To evaluate body biasing, we fix $V_{\rm diff}$ at the energy-optimal value and sweep $V_{\rm off}$ from -400 to 150 mV to tune performance. In both cases, fine-grained regulation and bias generation would be required. The energy penalties for these regulators have not been included, though we note that the body node draws very little current, so the body bias is generally simpler to regulate than $V_{\rm dd}$ is. Fig. 13 shows the resulting energy consumption over a frequency range of 30–300 kHz. The characteristic is similar to that of Fig. 12(b), but we find that body biasing is more energy efficient over the entire frequency range. The observed energy improvement is due to the PFET/NFET matching achieved through tuning of $V_{\rm diff}$. The ability to achieve PFET/NFET matching as well as a favorable energy–performance tradeoff makes body biasing an attractive alternative to voltage scaling in subthreshold circuits with tight performance requirements.

C. Subthreshold Sizing Strategies

Techniques for improving performance along timing-critical paths are also important for subthreshold circuits. Gate width sizing is typically used to speed up critical paths in super-threshold circuits, but recent work has shown that gate length sizing can be used to improve drive strength in subthreshold circuits due to reverse short channel effects (RSCE) [22], [13]. Halo doping increases the effective doping at short channel lengths to help combat drain-induced barrier lowering (DIBL)[23]. However, since DIBL is much reduced at low $V_{\rm dd}$, the halo doping overcompensates and increases the $V_{\rm th}$. Drain current can therefore be increased significantly with a small increase in gate length. The simulated on-current of an NFET device at $V_{\rm gs} = 250$ mV and $V_{\rm ds} = 250$ mV is shown as a function of total device capacitance (gate capacitance plus drain capacitance) for both increased gate width and increased gate length in Fig. 14. The current–capacitance tradeoff is far more attractive when increasing gate length than when increasing gate width. This discrepancy is largely due to RSCE, but capacitance also increases more slowly with gate length than with gate width. Gate-oxide capacitance depends identically on gate width and gate length, but overlap and junction capacitance are independent of gate length. The effectiveness of gate-length sizing eventually saturates, suggesting that gate-width sizing should be used after the benefits of gate-length sizing have been exhausted. Note that the results shown are highly technology-dependent, so performance gains will vary from process to process and from generation to generation. Also note that variable-channel-length devices may violate channel-length uniformity rules in advanced technologies.

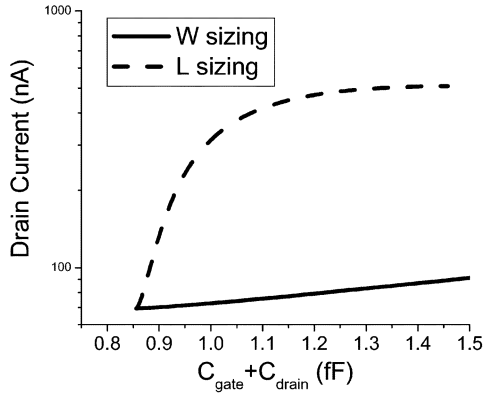

Fig. 14. Simulated on-current for an NFET as a function of total device capacitance. The tradeoff is shown for both gate-width and gate-length sizing.

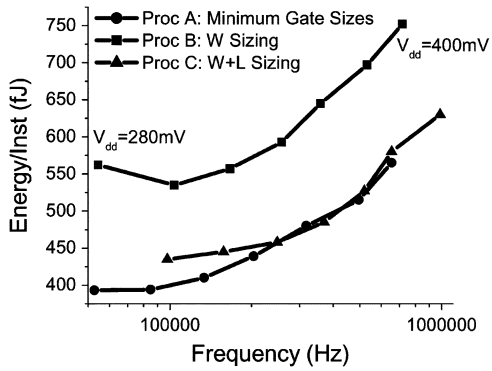

Fig. 15. Energy and frequency for three sizing strategies for $V_{\rm dd}$ = 280–400 mV.

D. Subthreshold Sizing Measurements

To explore gate sizing further, we designed three variants of the processor described in Section II. The first variant (Proc A) uses minimum gate sizes in the core. The second variant achieves reduced delay along critical paths using a conventional standard cell library with increased gate widths (Proc B). The third variant achieves reduced delay along critical paths using a custom standard cell library with both increased gate length and increased gate width (Proc C).

Both standard cell libraries were limited to a small set of inverters, two-input NAND gates, two-input NOR gates, and flip-flops. The drive strengths in the custom standard cell library used in Proc C were tuned to match those of the conventional standard cell library used in Proc B (i.e., X1, X2, and X4 cell strengths drive the same current in both libraries). Each library cell was characterized over a range of low voltages using SPICE.

The core gates sizes were optimized separately for Proc B and Proc C with energy as the objective according to the technique proposed in [24]. It was shown in [24] that the energy of a subthreshold circuit can be reduced by increasing gate sizes along critical paths due to the timing dependence of leakage energy. Proc B and Proc C were therefore designed with different frequency and energy targets that were determined by the characteristics of the standard cell library available during gate sizing. After sizing, Proc B and Proc C had total transistor gate areas that were 98% and 24% larger than the gate area in Proc A, respectively.

Fig. 15 compares the energy–delay tradeoff for the three different sizing strategies for $V_{\rm dd} = 280$ mV to $V_{\rm dd} = 400$ mV. We find that energy does not improve in Proc B and Proc C relative to Proc A, which is contrary to the conclusions in [24]. However, in [24], it was shown that the cost of reducing energy rises quickly after the gates along the first few critical paths have been sized up. Further sizing after initial energy gains leads to a considerable area penalty and potentially an energy penalty if standard cell energy characterization models do not match post-silicon performance. It is likely that the unexpectedly high energy consumption in Proc B and Proc C was caused by this effect.

Though the comparison between Procs B/C and Proc A revealed that the larger gate sizes increase energy consumption, we can still draw valuable conclusions about the effectiveness of gate-length sizing by comparing the energy and performance of Proc C to that of Proc B. At $V_{\rm dd} = 300$ mV, Proc B and Proc C are 22% and 85% faster than is Proc A, respectively. In addition to being faster than Proc B, Proc C is also more energy efficient over the $V_{\rm dd}$ range shown, confirming the superiority of gate-length sizing over gate-width sizing. For target frequencies above 200 kHz, the energy consumption of Proc A is comparable to that of Proc B, suggesting that the performance gained from gate-length sizing could be alternatively achieved by increasing $V_{\rm dd}$ by 20–30 mV.

V. Conclusion

Subthreshold design is promising for ultra-low-energy applications. With this work, we have successfully demonstrated ultra-low-energy subthreshold operation and have studied variability and performance, which are two of the primary challenges at low voltage. We first described an 8-bit processor fabricated in a 0.13 μm technology that consumes only 3.5 pJ/instruction at $V_{\rm dd} = 350$ mV. We explored the use of body biasing in subthreshold circuits to minimize process- and temperature-induced variations. The exponential sensitivity of subthreshold current to body bias enabled us to eliminate performance variations while maintaining energy efficiency. We also investigated techniques for improving global and local performance. Measurements of the test chip showed that body biasing is a more energy-efficient global technique than $V_{\rm dd}$ scaling over the frequency range considered. Our study of local sizing techniques showed that gate-length sizing should be used in place of gate-width sizing when improving critical path performance.

References

[1] J. Rabaey, J. Ammer, T. Karalar, B. O. S. Li, 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*, 2002, pp. 200–201.

[2] B. Zhai, D. Blaauw, D. Sylvester, and K. Flautner, "Theoretical and practical limits of dynamic voltage scaling," in *Proc. Design Automation Conf.*, 2004, pp. 868–873.

[3] B. Calhoun and A. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 2004, pp. 90–95.

[4] A. Wang and A. Chandrakasan, "A 180 mV FFT processor using subthreshold circuit techniques," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2004, pp. 292–293.

[5] B. Zhai, L. Nazhandali, J. Olson, A. Reeves, M. Minuth, R. Helfand, S. Pant, D. Blaauw, and T. Austin, "A 2.60 pJ/inst subthreshold sensor processor for optimal energy efficiency," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2006, pp. 154–155.

[6] L. Chang, Y. Nakamura, R. Montoye, J. Sawada, A. Martin, K. Kinishita, F. Gebara, K. Agarwal, D. Acharyya, W. Haensch, K. Hosokawa, and D. Jamsek, "A 5.3 GHz 8T-SRAM with operation down to 0.41 V in 65 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 252–253.
[7] B. H. Calhoun and A. Chandrakasan, "A 256 kb subthreshold SRAM in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2006, pp. 628–629.
[8] N. Verma and A. P. Chandrakasan, "A 65 nm 8T sub-V, SRAM employing sense-amplifier redundancy," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 327–328.
[9] T. Kim, J. Liu, J. Keane, and C. H. Kim, "A high-density subthreshold SRAM with data-independent bitline leakage and virtual ground replica scheme," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 329–330.
[10] B. Zhai, D. Blaauw, D. Sylvester, and S. Hanson, "A sub-200 mV 6T SRAM in 130 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 332–333.
[11] L. Nazhandali, B. Zhai, J. Olson, A. Reeves, M. Minuth, R. Helfand, S. Pant, T. Austin, and D. Blaauw, "Energy optimization of subthreshold-voltage sensor network processors," in *Proc. Int. Symp. Computer Architecture*, 2005, pp. 197–207.
[12] L. Nazhandali, M. Minuth, B. Zhai, J. Olson, T. Austin, and D. Blaauw, "A second-generation sensor network processor with application-driven memory optimizations and out-of-order execution," in *Proc. Int. Conf. Compilers, Architecture Synthesis, Embedded Systems*, 2005, pp. 249–256.
[13] S. Hanson, B. Zhai, M. Seok, B. Cline, K. Zhou, M. Singhal, M. Minuth, J. Olson, L. Nazhandali, T. Austin, D. Sylvester, and D. Blaauw, "Performance and variability optimization strategies in a sub-200 mV, 3.5 pJ/inst, 11 nW subthreshold processor," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 152–153.
[14] S. Hanson, B. Zhai, K. Bernstein, D. Blaauw, A. Bryant, L. Chang, K. Das, W. Haensch, E. Nowak, and D. Sylvester, "Ultra-low voltage minimum energy CMOS," *IBM J. Res. Devel.*, vol. 50, no. 4/5, pp. 469–490, 2006.
[15] J. Chen, L. Clark, and Y. Cao, "Robust design of high fan-in/out subthreshold circuits," in *Proc. Int. Conf. Computer Design*, 2005, pp. 405–410.
[16] K. Bernstein, D. J. Frank, A. E. Gattiker, W. Haensch, B. L. Ji, S. R. Nassif, E. J. Nowak, D. J. Pearson, and N. J. Rohrer, "High-performance CMOS variability in the 65-nm regime and beyond," *IBM J. Res. Devel.*, vol. 50, no. 4/5, pp. 433–449, 2006.
[17] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 2005, pp. 20–25.
[18] J. T. Kao, M. Miyazaki, and A. P. 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.
[19] C. H. Kim, K. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Jun. 2003.
[20] A. Bryant, J. Brown, P. Cottrell, M. Ketchen, J. Ellis-Monaghan, and E. Nowak, "Low-power CMOS at Vdd = 4kT/q," in *Proc. Device Res. Conf.*, 2001, pp. 22–23.
[21] G. Ono and M. Miyazaki, "Threshold-voltage balance for minimum supply operation," *IEEE J. Solid-State Circuits*, vol. 38, no. 5, pp. 830–833, May 2003.
[22] T. Kim, H. Eom, J. Keane, and C. Kim, "Utilizing reverse short channel effect for optimal subthreshold circuit design," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 2006, pp. 127–130.
[23] Y. Taur and E. Nowak, "CMOS devices below 0.1 μm: How high will performance go?," in *Int. Electron Devices Meeting (IEDM) Tech. Dig.*, 1997, pp. 215–218.
[24] S. Hanson, D. Sylvester, and D. Blaauw, "A new technique for jointly optimizing gate sizing and supply voltage in ultra-low energy circuits," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 2006, pp. 338–341.
[25] Y. Ramadass and A. Chandrakasan, "Minimum energy tracking loop with embedded DC-DC converter delivering voltages down to 250 mV in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 64–65.
[26] S. Hanson, B. Zhai, D. Blaauw, D. Sylvester, A. Bryant, and X. Wang, "Energy optimality and variability in subthreshold design," in *Int. Symp. Low Power Electronics and Design (ISLPED)*, 2006, pp. 363–365.

Scott Hanson (S'05) received the B.S. and M.S. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 2004 and 2006, respectively, where he is currently working toward the Ph.D. degree in electrical engineering.

His research interests include low-voltage circuit design for ultra-low-energy applications, variation-tolerant circuit design, and energy-efficient high-performance circuit design.

Mr. Hanson is the recipient of an SRC Fellowship.

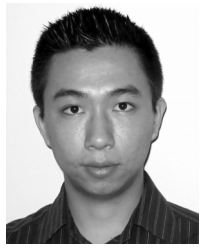

Bo Zhai received the B.S. degree in microelectronics from Peking University, Peking, China, in 2002, and the M.S. and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 2004 and 2007, respectively.

His research focuses on low-power VLSI design. He is a currently a Senior Design Engineer with Advanced Micro Devices, Austin, TX.

Mingoo Seok (S'06) was born in Seoul, Korea. He received the B.S. degree in electrical engineering from Seoul National University, Seoul, Korea, in 2005. He is currently working toward the Ph.D. degree in electrical engineering at the University of Michigan, Ann Arbor.

His research interests include low-voltage/low-power circuit design and physical aspects of submicrometer technology.

Mr. Seok is the recipient of a KFAS Fellowship.

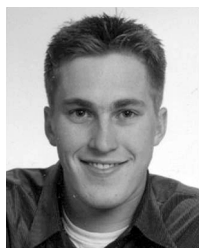

Brian Cline received the B.S. degree in electrical engineering from the University of Texas at Austin in 2004, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2006, where he is currently working toward the Ph.D. degree.

His research interests include low-power circuit design, variation-aware CAD tool development, and VLSI design optimization for high-performance and low-power designs.

Kevin Zhou, photograph and biography not available at the time of publication.

Meghna Singhal received the B.S. degree from BITS-Pilani, India, in 2004. She came to the University of Michigan, Ann Arbor, to pursue the M.S. degree in electrical engineering in 2004.

While at the University of Michigan, she worked with Prof. Todd Austin developing Pathequalizer (a buffer insertion software) and later worked with Prof. David Blauuw on the subliminal test chip. She is currently working with the microprocessor design group at Advanced Micro Devices, Boston, MA.

Michael Minuth, photograph and biography not available at the time of publication.

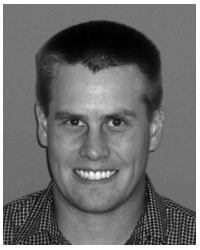

Javin Olson received the B.S. degree in computer engineering from Northwestern University, Evanston, IL, in 1999, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2005.

While attending the University of Michigan, he was a research assistant in the Advanced Computer Architecture Lab under the direction of Prof. Todd Austin. He is currently employed in the microprocessor design group at Advanced Micro Devices, Boston, MA.

Leyla Nazhandali, photograph and biography not available at the time of publication.

Todd Austin received the Ph.D. degree in computer science from the University of Wisconsin, Madison, in 1996.

He is an Associate Professor of Electrical Engineering and Computer Science with the University of Michigan, Ann Arbor. His research interests include reliable system design, hardware and software verification, and performance analysis tools and techniques. Prior to joining academia, he was a Senior Computer Architect with Intel's Microcomputer Research Laboratories, a product-oriented research laboratory in Hillsboro, OR. He was one of the creators of the SimpleScalar Tool Set, a popular collection of computer architecture performance analysis tools. In addition to his work in academia, he is founder and President of SimpleScalar LLC and cofounder of InTempo Design LLC.

Dr. Austin was a Sloan Research Fellow in 2002, and a recipient of the ACM Maurice Wilkes Award in 2007 "for innovative contributions in Computer Architecture including the SimpleScalar Toolkit and the DIVA and Razor architectures."

Dennis Sylvester (S'96–M'97–SM'04) received the B.S. degree *summa cum laude* from the University of Michigan, Ann Arbor, and the Ph.D. degree from the University of California, Berkeley, in 1999, both in electrical engineering.

He is now an Associate Professor of Electrical Engineering and Computer Science with the University of Michigan. He previously held research staff positions with the Advanced Technology Group of Synopsys, Mountain View, CA, Hewlett-Packard Laboratories, Palo Alto, CA, and a visiting professorship in electrical and computer engineering with the National University of Singapore. He has published numerous articles along with one book and several book chapters in his field of research, which includes low-power circuit design and design automation techniques, design-for-manufacturability, and interconnect modeling. He also serves as a consultant and technical advisory board member for several electronic design automation and semiconductor firms in these areas.

Dr. Sylvester is a member of the Association for Computing Machinery, the American Society of Engineering Education, and Eta Kappa Nu. He was the recipient of a National Science Foundation CAREER Award, the Beatrice Winner Award at ISSCC, an IBM Faculty Award, an SRC Inventor Recognition Award, and several best paper awards and nominations. He was the recipient of the ACM SIGDA Outstanding New Faculty Award and the University of Michigan Henry Russel Award for distinguished scholarship. He has served on the technical program committee of numerous design automation and circuit design conferences, the steering committee of the ACM/IEEE International Symposium on Physical Design, and was general chair for the 2005 ACM/IEEE Workshop on Timing Issues in the Synthesis and Specification of Digital Systems (TAU). He is currently an Associate Editor for the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN AND INTEGRATED CIRCUITS and previously served as an Associate Editor for the IEEE TRANSACTIONS ON VERY LARGE-SCALE INTEGRATION (VLSI) SYSTEMS. His dissertation research was recognized with the David J. Sakrison Memorial Prize as the most outstanding research in the Electrical Engineering and Computer Science Department at the University of California, Berkeley.

David Blaauw (M'94) received the B.S. degree in physics and computer science from Duke University, Durham, NC, in 1986, and the Ph.D. degree in computer science from the University of Illinois, Urbana, in 1991.

Until August 2001, he worked for Motorola, Inc. in Austin, TX, where he was the manager of the High Performance Design Technology group. Since August 2001, he has been on the faculty at the University of Michigan as an Associate Professor. His work has focused on VLSI design and CAD with particular emphasis on circuit design and optimization for high performance and low power applications.

Dr. Blaauw was the Technical Program Chair and General Chair for the International Symposium on Low Power Electronic and Design and was the Technical Program Co-Chair and member of the Executive Committee the ACM/IEEE Design Automation Conference. He is currently a member of the ISSCC Technical Program Committee.