

Design Strategies for Ultra-Low Voltage Circuits

Dennis Sylvester, Scott Hanson, Bo Zhai, David Blaauw

Electrical Engineering and Computer Science

University of Michigan

Ann Arbor, MI, U.S.A.

{dmcs,hansons,bzhai,blaauw}@umich.edu

Abstract - Energy efficiency is an emerging metric for the quality of integrated circuit designs. Applications ranging from wireless sensor networks to RFID tags to embedded microprocessors require extremely low power consumption to maintain good battery life. We advocate the use of aggressively scaled supply voltages in such applications to maximize energy efficiency. This paper reviews our recent progress in mapping out the low energy design space including the presence of an energy-optimal supply voltage, and also touches on gate sizing techniques and variability issues. We conclude with a survey of open research directions in the ultra-low voltage design space.

Keywords: Integrated circuits, CMOS, low-power

1 Introduction

New classes of applications such as wireless sensor networks have dramatically different requirements from traditional integrated circuits (ICs) where performance (i.e., frequency) is the primary metric of interest. In this example, each node in the sensor network system must consume very little power to achieve a battery life of months or years. In addition, each node in the system must be fabricated at a negligible cost. Nodes should ideally be disposable and networks with thousands or millions of nodes should be achieved at a moderate cost. The cost of a single node can be measured in a number of ways, including the cost per unit area of silicon, design cost, and the cost of maintenance. It is therefore important that a prospective design minimizes area and maximizes simplicity. For most such applications, performance (i.e., the time required to complete a given computation) is only a secondary concern. For example, a sensor network deployed on an island off the coast of Maine monitored habitat conditions by measuring all sensors once every 70 seconds [15].

The above discussion motivates the growing need for ultra low power (or more frequently, low energy, since this translates to battery life) ICs. In this paper we analyze the low energy design space and make conclusions about how such circuits should be designed. The models and guidelines developed from this work have been successfully used in the design of world's most energy efficient microprocessor [16], which is described briefly. We also highlight a number of open research directions

that will enable ubiquitous reliable and energy-efficient systems.

2 Theoretical Foundations

Voltage scaling has emerged as one of the most effective techniques for meeting the increasingly stringent power demands in modern chip designs. A number of industrial designs have ventured as low as half of the nominal supply voltage. This voltage scaling has been the source of dramatic energy reductions but has left several lingering questions: How far can the supply voltage scale in conventional CMOS logic? Even if CMOS logic functions properly at very low supply voltages, is it worthwhile to operate at these voltages? We find in the subsequent discussion that aggressive voltage scaling into the subthreshold regime (i.e., $V_{dd} < V_{th}$) is possible, but care must be taken in choosing the proper supply voltage due to the increased importance of leakage energy and variability.

2.1 Energy Minimization

In [1], it was shown that CMOS gates composed of ideal transistors with a subthreshold swing of 60 mV/decade should function properly with a supply voltage as low as 36 mV. Despite a non-ideal subthreshold swing, measurements of an inverter show that functionality can be achieved with a supply voltage of just 65 mV [2]. It is clear that CMOS logic functions at extremely low voltages, but we must still consider the question of whether operation at these voltages is worthwhile. Figure 1(a) shows how the power consumed by an inverter chain scales with supply voltage. Total power consumption is broken into dynamic power (the power consumed by switching gates) and leakage power (the power consumed by idle gates). Minimum power is achieved by choosing the minimum functional supply voltage.

However, power is not always the most appropriate metric. For many applications, especially those in which battery life is the primary concern, energy per instruction may be a more sensible metric. There is a subtle but important difference between energy and power that is highlighted in Figure 1(b), which shows the energy consumed per switching event (which we call an operation) for the inverter chain from Figure 1(a).

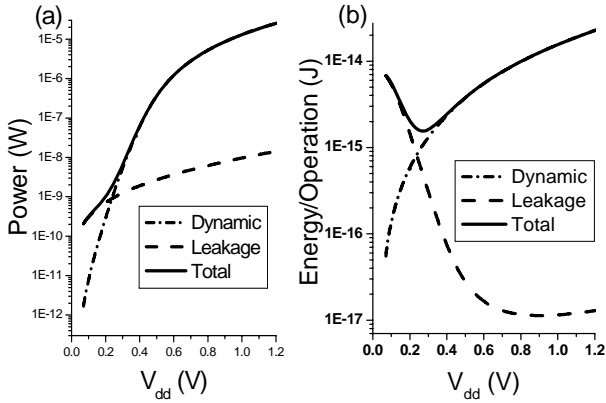


Figure 1. (a) Power consumed by an inverter chain as a function of supply voltage (V_{dd}) (b) Energy consumed per switching operation by the same inverter chain as a function of V_{dd} .

Although Figure 1(a) shows that minimizing supply voltage will minimize power, the energy inflection point in Figure 1(b) shows that minimum energy is achieved at some voltage that is greater than the minimum functional supply voltage. This energy minimum is due to a rapid increase in gate delay as the supply voltage scales below the threshold voltage. As gate delay increases, the amount of time that each gate spends leaking also increases. As a result, the total leakage energy (the product of leakage current, supply voltage, and total leakage time) increases quickly and creates the minimum apparent in Figure 1(b). The location of this minimum energy supply voltage (V_{min}) is a strong function of both switching activity and logic depth (the number of gates between an input and an output) and was derived in [3] as:

$$V_{min} = \left[1.587 \cdot \ln \left(\eta \cdot \frac{n}{\alpha} \right) - 2.355 \right] \cdot m \cdot v_T \quad (1)$$

where α is the switching activity, n is the logic depth, m is the subthreshold sloper factor, v_T is the thermal voltage, and η is a delay-related technology parameter.

2.2 Circuit Design

When the supply voltage is reduced to the energy optimal point, V_{min} , circuit behavior changes dramatically. The on-current to off-current ratio (I_{on}/I_{off}) reduces to ~ 100 - 1000 , absolute noise margins reduce, the relative PFET to NFET ratio (beta ratio) changes, and leakage accounts for a significant portion (30% or more) of total energy. All of these changes suggest that circuit design strategies must evolve as voltage scales into the subthreshold regime.

Gate sizing is one of the most fundamental circuit techniques available to any circuit designer. Very recently, a subthreshold logical effort strategy was

proposed [18]. The strategy demonstrates performance improvements using new sizing schemes for transistors in a stack. It was also shown recently that increasing select gate sizes can actually reduce the energy consumption of a subthreshold circuit [19]. Figure 2 illustrates the concept of the sizing strategy. By increasing the gate sizes along critical paths, the clock period (the leakage time) of a circuit may be reduced. If there are few critical paths (meaning the overhead of gate sizing is small), the leakage energy reductions may outweigh any energy increase due to larger gates. We saw in the last section that the location of V_{min} is determined by leakage. A reduction in leakage energy therefore reduces V_{min} and permits further reduction in V_{dd} . The joint optimization of gate sizes and supply voltage yields energy reductions of 5.6-15% on representative benchmarks in [19].




Figure 2. Imbalanced delays in subthreshold circuits result in inefficiency since leakage accounts for 30% or more of total energy at V_{min} . Gate sizing along critical paths can help address this problem.

2.3 Variability Considerations

While the energy benefits of low voltage operation are evident, it is not immediately clear that those benefits will exist after variability is considered. In addition to a reduction in absolute noise margins, low voltage operation suffers from a heightened sensitivity to threshold voltage variations. In the case of subthreshold operation, both off-current (leakage) and on-current are composed entirely of subthreshold current, which is exponentially dependent on the threshold voltage. Low voltage circuits are therefore extremely sensitive to threshold variations, particularly those induced by random dopant fluctuations [4].

The most basic question we must answer is how the minimum energy point (E_{min}) and the minimum energy supply voltage (V_{min} , described by Equation 1) are affected by variation. Figure 3 plots E_{min} and V_{min} for an inverter chain of length n under gate length and threshold voltage variability. Systematic V_{th} variation (chip mean V_{th} variation) is not included in this analysis, though should be considered in future work. For an inverter chain length of 16, V_{min} increases by 58% and E_{min} increases by 85%. However, it is possible to mitigate the effects of random variation by increasing the inverter chain length and “averaging” the variation. In Figure 3, the V_{min} penalty is reduced to 20% and the E_{min} penalty is reduced to 46%

when the inverter chain length is increased to 60, a vast improvement over the chain of length 16. A careful choice of V_{min} in combination with judicious selection of logic depth, n , can help reduce energy under variability, but further work is necessary to maintain high yield under the dramatic variability expected in low voltage operation.

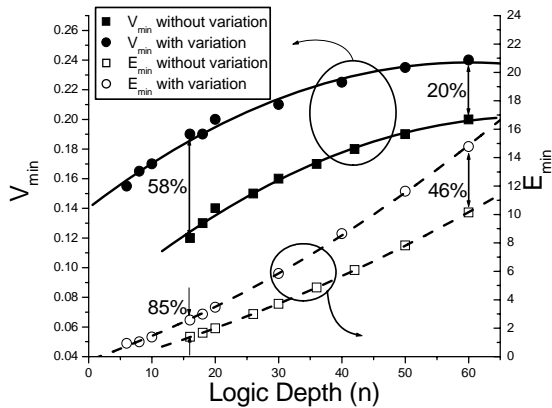


Figure 3. Minimum energy, E_{min} , and the energy optimal supply voltage, V_{min} , for inverter chain of length n . Variation numbers are worst case ($\mu+3\sigma$). HSPICE data is taken from Monte Carlo simulations in [4].

3 Design Example

We now discuss the design and test of a subthreshold sensor network processor fabricated in a $0.13\mu\text{m}$ process with $V_{th}\sim 400\text{mV}$ [16]. The processor uses a simple 8-bit architecture with a 2-kbit memory. The simplicity of the architecture and small memory size are key factors in achieving high energy efficiency. The chip is extremely small, occupying only $85,000\ \mu\text{m}^2$. To enable robust operation in the subthreshold regime, a latch-based memory with mux-based read/write scheme was used [6]. Using this memory structure, the processor remains functional as low as $V_{dd}=200\text{mV}$.

Figure 4 shows the processor energy consumption for a typical die, confirming the trends observed in Section 2. Over 26 dies, V_{min} ranges from 340-420mV while E_{min} ranges from 2.6-3.4pJ/instruction. The average maximum operating frequency at $V_{dd}=400\text{mV}$ is 1.5MHz. Analysis of energy distributions over the 26 dies shows that adaptive frequency scaling (the selection of a unique frequency for each die) is more important when minimizing energy than adaptive voltage scaling (the selection of a unique supply voltage for each die).

Measurements of this subthreshold processor clearly show that high energy efficiency can be achieved in the subthreshold and near-subthreshold regimes. However, variability and robust memory design stand out as two issues that must be addressed before subthreshold circuits will gain widespread use. We discuss these and other open research directions in the next section.




Figure 4. Energy consumption for an 8-bit subthreshold processor [16].

4 Open Research Directions

Below we suggest several avenues for future exploration in the general area of ultra low-voltage design. In addition to dedicated low-power applications, subthreshold processing elements could form the basis of a heterogeneous massively multi-core processor architecture [17]. Each core would emphasize a different power-performance point and memory arrays would natively run at higher voltages since they are more energy-efficient in that space as well as more robust. This would allow for very high memory bandwidth relative to the processor speeds as well, effectively reversing the memory bottleneck commonly seen on large processors today [2].

Variability due to all sources, including process, voltage, and temperature (PVT) are all magnified in subthreshold as discussed earlier. There is a great need for a range of effective techniques to combat this variability. In particular, temperature-insensitive design approaches become an interesting topic of study; this may employ adaptive body bias [5]. Logic families other than CMOS may offer greater resiliency to certain variation sources such as voltage or process. This is another open area for exploration. Architectural approaches to variability are also vital – these may include a proliferation of traditional schemes such as error correction and redundancy, although the power costs of added hardware need to be weighed carefully against the improvements in robustness achieved.

Energy-efficient subthreshold design cannot succeed without robust and dense ultra-low voltage memory design techniques. The major concerns with subthreshold memory design are: 1) process variation in very small dimension devices worsen the mismatch behavior in the traditional 6T SRAM cell design, and 2) reduced on-to-off current ratios complicate the reading and writing steps. One proposed solution was a latch-based architecture [6], which incurs severe area and performance penalties. A more recent work [7] using a 10-T cell design functions near V_{th} for a 65nm technology and requires two supply

voltages. None of the current approaches are completely satisfactory and advancements in this area are one of the most crucial needs for the proliferation of extreme low-power systems.

Very low-power wireless communication schemes are needed, else the low energy budget of the digital processing component of a system will simply be swamped out by the communication requirements. This is highly application-specific; there may be cases where proximity communication schemes [9,10] are suitable and others where more traditional radios are used but with new architectures to address strict power budgets [8].

Extremely low duty cycle applications are very common – for instance environmental monitoring such as that mentioned in the introduction, or the monitoring of cracks along oil pipelines. All of these applications require a sensor to be read and data to be processed on a relatively infrequent basis (on the order of minutes). For such systems, standby mode power will completely dominate and current techniques to reduce standby mode leakage such as power gating are insufficient to provide required battery lifetimes. In these cases transistors should be employed as frugally as possible since added devices will inherently leak during standby. Novel architectures and low-power modes will need to be invented to enable sleep mode current levels on the order of 1nA.

Subthreshold circuits can greatly benefit from a rethinking of device [11] and interconnect architectures. Since it may not be feasible to have a specific process variant dedicated to subthreshold, work is needed to find ways to tailor device and interconnect behavior to the unique needs of subthreshold design. Interconnect is interesting since parameters such as wire resistance become inconsequential compared to channel resistances. This may lead to a complete re-thinking of how things such as clock and power routing are done, as well as how back-end stacks are manufactured (e.g., thick and wide wires are counterproductive since capacitance is the only important parasitic) [2].

Finally, given the possibility of nW processing there is an opportunity for on-die power sources that are CMOS-compatible. Most simply this could take the form of on-die thin-film batteries [12,13] that in the past have been impractical due to their limited energy capacities per unit area. These would be low cost and process compatible. Another option is energy scavenging through solar power, ambient vibration, or other techniques [14]. These power sources could also be used to recharge the primary battery, rather than supply power to the entire system.

References

- [1] J. Meindl and J. Davis, "The Fundamental Limit on Binary Switching Energy for Terascale Integration (TSI)," *IEEE J. Solid-State Circuits*, pp. 1515-1516 October 2000.
- [2] S. Hanson, *et al.*, "Ultralow-voltage minimum-energy CMOS," *IBM Journal of Research and Development*, pp. 469-490, July/Sept 2006.
- [3] B. Zhai, D. Blaauw, D. Sylvester, K. Flautner, "Theoretical and Practical Limits of Dynamic Voltage Scaling," *Design Automation Conf.*, pp. 868-873, 2004.
- [4] B. Zhai, S. Hanson, D. Blaauw, D. Sylvester, "Analysis and Mitigation of Variability in Subthreshold Design," *Intl. Symp. Low Power Electronics and Design*, pp. 20-25, 2005.
- [5] S. V. Kumar, C. H. Kim, S. S. Sapatnekar, "Mathematically-Assisted Adaptive Body Bias (ABB) for Temperature Compensation in Gigascale LSI Systems," *Asia-South Pacific Design Automation Conf.*, pp. 559 – 564, 2006.
- [6] A. Wang, A. Chandrakasan, "A 180mV FFT processor using subthreshold circuit techniques," *IEEE Intl. Solid-State Circuits Conf.*, pp. 292-293, 2004.
- [7] B. Calhoun, A. Chandrakasan, "A 256kb subthreshold SRAM in 65nm CMOS," *IEEE Intl. Solid-State Circuits Conf.*, pp. 628-629, 2006.
- [8] J. Chen, M.P. Flynn, J. Hayes, "A Fully Integrated Auto-Calibrated Super-Regenerative Receiver," *IEEE Intl. Solid State Circuits Conf.*, 2006.
- [9] R.J. Drost, *et al.*, "Proximity communication," *IEEE J. Solid-State Circuits*, pp. 1529-1535, Sept. 2004.
- [10] N. Miura, *et al.*, "Analysis and design of inductive coupling and transceiver circuit for inductive inter-chip wireless superconnect," *IEEE J. Solid-State Circuits*, pp. 829-837, April 2005.
- [11] A. Raychowdhury *et al.*, "Computing with subthreshold leakage: device/circuit/ architecture co-design for ultralow-power subthreshold operation," *IEEE Trans. VLSI Systems*, pp. 1213-1224, Nov. 2005.
- [12] J. Harb, R. LaFollette, R. Selfridge, L. Howell, "Microbatteries for self-sustained hybrid micropower supplies," *Journal of Power Sources*, 104 (1) (2002), 46 – 51
- [13] C. Dewan, D. Teeters, "Vanadia xerogel nanocathodes used in lithium microbatteries," *Journal of Power Sources*, 119 (2003), 310 – 315
- [14] J.A. Paradiso, T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, pp. 18-27, Jan-Mar 2005.
- [15] R. Szewczyk, *et al.*, "Lessons from a sensor network expedition," *European Workshop on Sensor Networks*, 2004.
- [16] B. Zhai, *et al.*, "A 2.6pJ/inst subthreshold sensor processor for optimal energy efficiency," *IEEE Symp. on VLSI Circuits*, 2006.
- [17] R. Kumar, *et al.*, "Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction," *Int. Sym. Microarchitecture*, pp. 81–92, 2003.
- [18] J. Keane, *et al.*, "Subthreshold Logical Effort: A Systematic Framework for Optimal Subthreshold Device Sizing," *Design Automation Conf.*, pp. 425-428, 2006.
- [19] S. Hanson, D. Sylvester, and D. Blaauw, "A New Technique for Jointly Optimizing Gate Sizing and Supply Voltage in Ultra-Low Energy Circuits," *Int. Symp. Low Power Electronics and Design*, in press, 2006.