Device Sizing for Minimum Energy Operation in Subthreshold Circuits

Benton H. Calhoun, Alice Wang, and Anantha Chandrakasan
Department of Electrical Engineering
Massachusetts Institute of Technology, Cambridge, MA

Abstract

Digital circuits operating in the subthreshold region provide the minimum energy solution for applications with strict energy constraints. This paper examines the effect of sizing on energy for subthreshold circuits. We show that minimum sized devices are theoretically optimal for reducing energy. A fabricated 0.18µm test chip is used to compare normal sizing and sizing for minimum V_{DD}. Measurements show that existing standard cell libraries offer a good solution for minimizing energy in subthreshold circuits.

Introduction

Emerging applications such as distributed sensor networks or medical applications have low energy as the primary concern instead of performance. Minimum energy operation for low performance situations occurs in the subthreshold region [1][2]. Increasing leakage energy at low supply voltages offsets the reduced active energy and causes a minimum energy point. Many designs exhibit a minimum energy operating point higher than the minimum achievable V_{DD}, and this operating point is a function of several parameters [1][3]. In general, designs with larger leakage energy relative to active energy have a higher optimum V_{DD}. This paper examines the effect of device sizing on minimum energy operation. After considering theoretically optimal sizing, we explore minimum energy operation for standard cell designs. A fabricated 0.18µm test chip provides measurements for analysis.

Optimal Sizing for Minimum Energy

Sizing influences the energy consumption of a circuit in two primary ways. First, sizing directly affects energy consumption by changing switched capacitance and leakage current. Secondly, sizing affects the minimum voltage at which the circuit functions, which can change the absolute minimum energy point. This section explores these two impacts.

A. Sizing for a given V_{DD}

It has been proposed that theoretically optimal minimum energy circuits should use minimum sized devices [4], and first-order equations confirm this result for most cases. Equation (1) shows the propagation delay of a characteristic inverter with a certain switched capacitance $C_g=\alpha C_L$ in subthreshold, where α is activity factor and C_L is the load capacitance assuming average fanout:

$$t_d = \frac{KC_gV_{DD}}{I_{o,g}\exp\left(\frac{V_{GS}-V_{T,g}}{nV_{th}}\right)}. \quad (1)$$

K is a delay fitting parameter. The expression for current in the denominator of (1) models the on current of the characteristic inverter, so it accounts for transitions through both NMOS and PMOS devices. Thus, the terms $I_{o,g}$ and $V_{T,g}$ are fitted parameters that do not correspond exactly with the MOSFET parameters of the same name. For a single inverter, dynamic (E_{DYN}), leakage (E_L), and total energy (E_T) per cycle are expressed in (2)-(4), assuming rail-to-rail swing ($V_{GS}=V_{DD}$ for "on" current).

$$E_{DYN} = C_gV_{DD}^2 \quad (2)$$

$$E_L = I_{o,g}\exp\left(\frac{-V_{T,g}}{nV_{th}}\right)V_{DD}t_d = KC_gV_{DD}^2\exp\left(-\frac{V_{DD}}{nV_{th}}\right) \quad (3)$$

$$E_T = C_gV_{DD}^2\left(1+K\exp\left(-\frac{V_{DD}}{nV_{th}}\right)\right) \quad (4)$$

Total energy per cycle is proportional to C_g, so minimum sized devices give a minimum C_g and minimize E_T for a single inverter at a given V_{DD}. Equations for energy of arbitrary circuits can use the inverter analysis as a foundation. Assuming a critical path depth of L_{DP} characteristic inverter delays gives an operating frequency, $f = (t_dL_{DP})^{-1}$. We can likewise use the inverter capacitance and leakage current to define total switched capacitance, $C_{eff}=K_1C_g$, and total leakage width, W_{eff}. These changes give the expression

$$E_T = C_{eff}V_{DD}^2 + W_{eff}KC_gL_{DP}V_{DD}^2\exp\left(-\frac{V_{DD}}{nV_{th}}\right) \quad (5)$$

for total energy per operation in an arbitrary circuit [5]. This equation provides an estimate for a generic circuit since C_{eff} and W_{eff} approximate other parameters that can change.

Assuming that the majority of gates in a typical design are sized similarly, a universal increase in transistor sizes will increase both C_{eff} and W_{eff}, raising power. This type of sizing change is unlikely to decrease the critical path delay because the input to output capacitance ratios of gates will stay roughly constant, so the typical assumption of fixed capacitance loads is invalid [9]. Thus, minimum sizing also minimizes energy per operation for most generic circuits. One special case that violates this trend is a circuit with a small number of critical paths relative to the total number of paths. In this case, increased sizes on the critical path can reduce L_{DP} with negligible increases in C_{eff} and W_{eff}, lowering E_T.

Minimum sized devices generally minimize energy consumption in subthreshold for a given V_{DD}. However, sizing also impacts the minimum operating voltage, which can affect the total energy per operation, E_T.

0-7803-8495-4/04/$20.00 ©2004 IEEE.

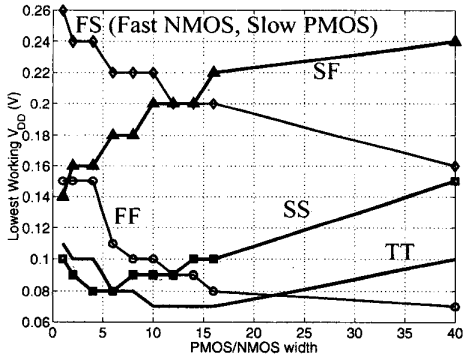

Fig. 1. Minimum achievable voltage retaining 10%-90% output swing for 0.18μm ring oscillator across process corners (simulation).

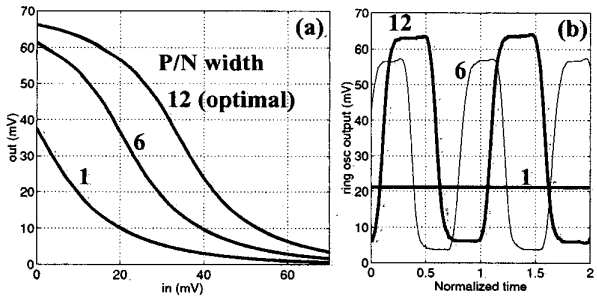

Fig. 2. VTC **(a)** and 9-stage ring oscillator output **(b)** at the minimum V_{DD} for the typical corner (simulation). Optimum PMOS/NMOS width ratio of 12 is shown with others for comparison.

B. Sizing and minimum operating voltage

Transistor sizing also impacts the functionality of CMOS circuits at low supply voltages. Minimum V_{DD} operation occurs when the PMOS and NMOS devices have the same current (e.g. [6]). Previous efforts have explored well biasing to match the device currents for minimum voltage operation of ring oscillators [7]. Sizing can create the same symmetry in device current. Fig. 1 shows the minimum voltage for which a ring oscillator maintains 10%-90% voltage swing. The optimum PMOS/NMOS width across all process corners is 12. A similar analysis of minimum voltage operation while retaining 10% noise margins gives a lower minimum voltage at the typical corner and a higher worst-case minimum voltage but the same optimum size ratio. Fig. 2(a) shows the VTCs at the minimum V_{DD} of 70mV for several P/N ratios. The gain is somewhat degraded, but the optimum sized curve is symmetrical and shows good noise margins. Fig. 2(b) shows the output of a 9-stage ring oscillator at the minimum voltage for the same sizes.

Since symmetrical devices give minimum V_{DD} operation, a simple comparison of currents in NMOS and PMOS devices shows the approximate optimum size for minimizing V_{DD}. The switching threshold of a symmetric inverter is $V_M=V_{DD}/2$. Sweeping the width of "on" NMOS and PMOS devices at $V_{DD}/2$ shows the PMOS size for which the inverter is balanced. Fig. 3 shows this simulation for V_{DD}=70mV at the typical corner. The current in the PMOS matches the NMOS current at a ratio of 12.5. This ratio remains between 11.5 and 12.5 for voltages across the subthreshold region, matching the result in Fig. 1. The non-monotonic behavior of the current in these devices results from the reverse narrow channel effect [8].

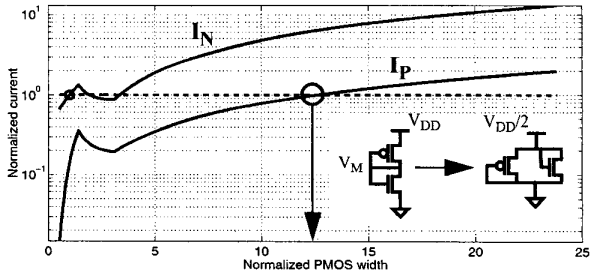

Fig. 3. Sizing an inverter to have a switching threshold at $V_{DD}/2$ gives a good estimate of the optimum width to minimize V_{DD} because the NMOS and PMOS are balanced. This simulation shows NMOS and PMOS current for the equivalent circuit on the right in the inset for V_{DD}=70mV.

Sizing according to this ratio allows for operation at lower V_{DD} but increases the energy consumed for a given V_{DD} (equation (5)). The energy savings from lowering V_{DD} are at best proportional to V_{DD}^2 if leakage is still negligible. Fig. 1 shows that the impact of sizing an inverter on the minimum supply voltage is only 60mV, producing best-case energy savings of $0.20^2/0.26^2$=0.6X due to voltage reduction. This improvement is not worthwhile if all PMOS devices are increased in size by 12X. Thus, minimum sized devices are theoretically optimal for reducing energy per operation when accounting for the impact of sizing on voltage and energy consumed.

Standard Cells and Minimum Energy

Standard cell libraries aid digital circuit designers to reduce the design time for complex circuits through synthesis. Most standard cell libraries focus on high performance, although including low power cells is becoming more popular [9]. Lower power cells generally use smaller sizes. One standard cell library geared specifically for low power uses branch-based static logic to reduce parasitic capacitances and a reduced set of standard cells. Eliminating complicated cells with large stacks of devices and using a smaller total number of logic functions was shown to reduce power and improve performance [10]. Standard cell libraries have not been designed specifically for subthreshold operation. This section evaluates the performance of a 0.18μm standard cell library in subthreshold operation.

We use an 8-bit, 8-tap, parallel, programmable FIR filter as a benchmark to compare normal cell selection with cells sized to minimize the operating voltage. Fig. 4 shows the minimum operating voltages for the different standard cells appearing in a normal synthesis of the FIR filter. The typical (TT) and worst-case (FS and SF) process corners are shown. All of the cells operate at 200mV at the typical corner, showing the robustness

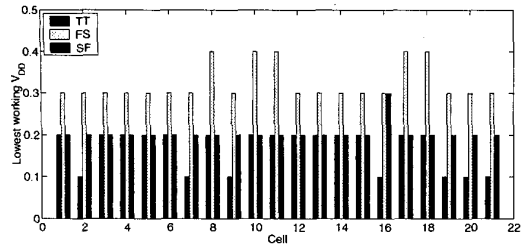

Fig. 4. Standard cell functionality in synthesized FIR filter using normal cell selection over process corners (simulation).

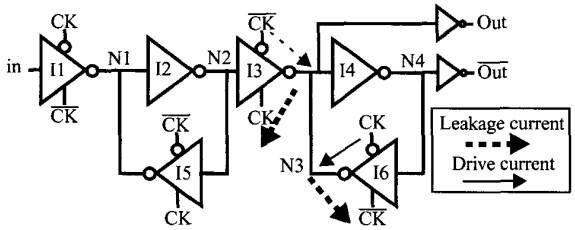

Fig. 5. Standard cell flip-flop at worst-case failure point where CK=0 at FS corner (fast NMOS, slow PMOS).

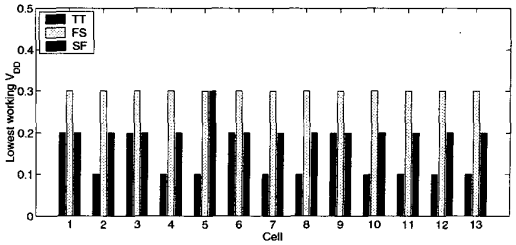

Fig. 6. Standard cell functionality in synthesized FIR filter using cells sized to minimize V_{DD} over process corners (simulation).

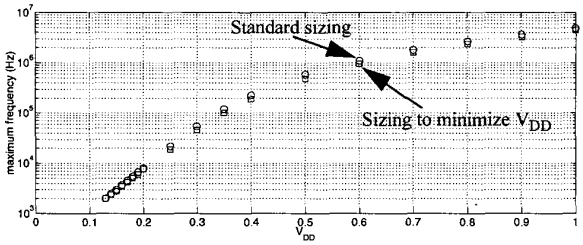

Fig. 7. Measured performance of programmable FIR filters on the test chip. Standard FIR is 10% faster than the minimum-voltage FIR.

of static CMOS logic. Additionally, most of the cells operate at 300mV in the worst case, which is close to the optimum performance shown in the previous section for a ring oscillator. The cells which exhibit the worst case (failing at 300mV) are flip-flops and complex logic gates with stacks of series devices (e.g. AOI). We eliminated the problematic cells by preventing the synthesis tool from selecting logic gates with large device stacks and by re-sizing the offending flip-flop cell.

Fig. 5 shows a schematic of the D-flip-flop. In the standard implementation, all of the inverters use small NMOS and only slightly larger PMOS devices except I3, which is several times larger to reduce C-Q delay. At the FS corner (fast NMOS, slow PMOS), the narrow PMOS in I6 cannot hold N3 at a one when CK is low. This is because the combined, strong off current in the NMOS devices in I6 and I3 (larger sized) overcomes the weakened, narrow PMOS device in I6. Tying back to the ring oscillator in Fig. 1, the combined NMOS devices create an effective P/N ratio that is less than one. To prevent this, we reduced the size of I3 and strengthened I6. Clearly, the larger feedback inverter creates some energy overhead. However, the resized flip-flop can operate at 300mV at all process corners in simulation. Fig. 6 shows the lowest operating voltage for the cells in the minimum-V_{DD} FIR filter. The number of cell types has reduced, and all of the cells work at 300mV across all corners. The next section uses test chip measurements to compare the filter sized for minimum V_{DD} with the normal filter.

Measured Results from Test Chip

A 0.18µm, 6M layer, 1.8V, 7mm^2 test chip was fabricated to measure the impact of sizing on minimum energy operation of standard cells. The test chip features two programmable 8-bit, 8-tap FIR filters. Both filters produce non-truncated 19-bit outputs. The first filter was synthesized using the unmodified synthesis flow and normal cells (Fig. 4). The second filter was synthesized using the modified flow in which some cells were omitted and some cells were resized to minimize V_{DD} (Fig. 6). Both filters can operate using an external clock or an on-chip clock generated by a ring oscillator that matches the respective critical path delay of the filters. Filtered data comes from an off chip source or from an on-chip linear-feedback shift-register.

Fig. 7 shows the measured performance versus V_{DD} for the two filters using their respective critical path ring oscillators and the LFSR data to produce one pseudorandom input per cycle. The minimum-V_{DD} filter exhibits a 10% delay penalty over the standard filter. Both filters operate in the range of 3kHz to 5MHz over V_{DD} values of 150mV to 1V. Both filters are fully functional to below 200mV.

Fig. 8(a) shows an oscilloscope plot of the standard filter working correctly at V_{DD}=150mV. The clock in this plot is produced by the ring oscillator on-chip. The reduced drive current and large capacitance in the output pads of the chip cause the slow rise and fall times in the clock, but the signal is still full swing. One bit of the output is shown. Fig. 8(b) shows an oscilloscope plot of the clock output for V_{DD}=65mV. This clock signal is generated off-chip and fed through the on-chip clock tree. Although the output swing is degraded and the signal is noisy, it shows the robustness of CMOS logic to low supply voltages.

Fig. 9 shows the measured total energy per output sample of the two FIR filters versus V_{DD}. The solid line is an extrapolation of $C_{eff}V_{DD}^2$ for each filter, and the dashed lines show the measured leakage energy per cycle. Clearly, both filters exhibit an optimum supply voltage for minimizing the total energy per cycle. Within the granularity of the measurements, the opti-

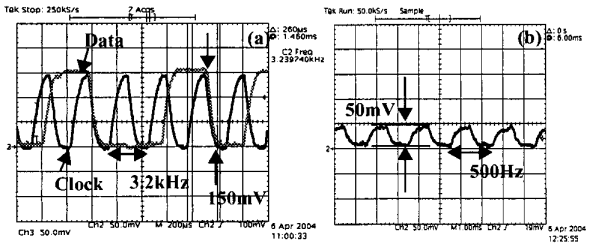

Fig. 8. Oscilloscope plots from the test chip. **(a)** V_{DD}=150mV filtering operation with ring oscillator clock at 3.2kHz, and **(b)** clock output for 500Hz input clock at V_{DD}=65mV.

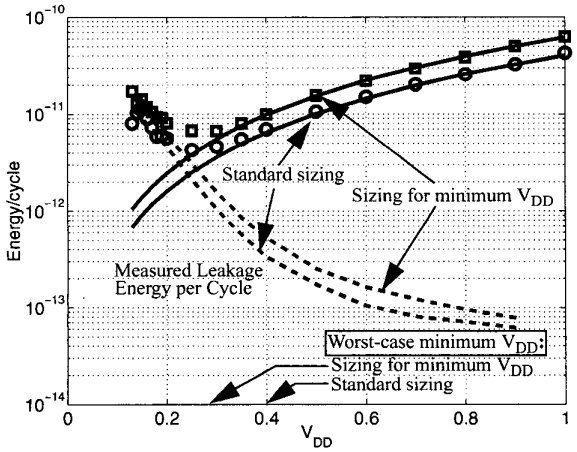

Fig. 9. Measured energy per operation of the FIR filters on the test chip.

mum V_{DD} is 250mV for the standard FIR and 300mV for the minimum-V_{DD} FIR. There is a measured overhead energy per cycle of 50% in the filter sized for minimum V_{DD}. The figure also shows the worst-case minimum V_{DD} for the two filters (cf. Fig. 4, Fig. 5). Accounting for overhead at the worst-case minimum V_{DD}, the minimum-V_{DD} FIR offers a reduction in total energy of less than 10% at the worst-case process corner, but this improvement comes at a cost of 50% at the typical corner.

Simulations show that the measured overhead cost in the minimum-V_{DD} filter primarily results from restricting the cell set that the synthesis tool could use. Since the tool was not optimized for the smaller set of cells, we did not see the improvements that are possible through this approach [10]. Using only sizing to create the minimum V_{DD} filter would have decreased the overhead. However, the shallow nature of the optimum point in Fig. 9 shows that the unmodified standard cell library does not use much extra energy by failing at a higher V_{DD} at the worst-case corner. Thus, existing libraries provide good solutions for subthreshold operation. Simulation shows that a minimum-sized implementation of the FIR filter has 2X less switched capacitance than the standard FIR, so a mostly minimum-sized library theoretically would provide minimum energy circuits.

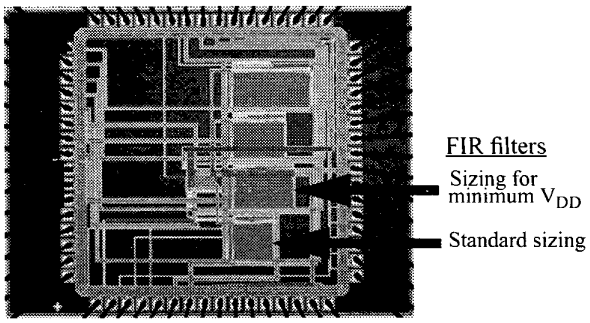

Fig. 10. Annotated die photo of 0.18µm subthreshold FIR test chip.

Conclusion

This paper has examined device sizing for subthreshold operation. For typical circuits and modern technologies, the optimum supply voltage for minimizing power is higher than the failure point for minimum sized devices at the typical corner. Thus, minimum sized devices are theoretically optimal for minimizing power. Even if the minimum energy point for a certain process corner or unusual circuit occurs at a supply voltage where minimum sized devices cannot function, the shallow nature of the optimum prevents up-sizing to reduce the minimum possible operating voltage from being worthwhile. Measurements from a test chip, shown in Fig. 10, confirm that existing static CMOS standard cell libraries function well in subthreshold. Resizing or restricting cell usage in such libraries can lower the worst-case minimum V_{DD}, but the overhead increases energy consumption at the typical corner. In theory, a standard cell library primarily using minimum-sized devices would minimize energy per operation.

Acknowledgements

This work was funded by Texas Instruments and by the Defense Advanced Research Projects Agency (DARPA) through a subcontract with MIT Lincoln Laboratory.

References

[1] A. Wang, A. P. Chandrakasan, and S. V. Kosonocky, "Optimal Supply and Threshold Scaling for Subthreshold CMOS Circuits," *ISVLSI*, pp. 5-9, 2002.

[2] J. Burr and A. Peterson, "Ultra Low Power CMOS Technology," *3rd NASA Symposium on VLSI Design*, pp. 4.2.1-4.2.13, 1991.

[3] A. Wang and A. Chandrakasan, "A 180mV FFT Processor Using Subthreshold Circuit Techniques," *Intl. Solid-State Circuits Conference*, 2004.

[4] E. Vittoz, "Low-Power Design: Ways to Approach the Limits," *Intl. Solid-State Circuits Conf.*, 1994.

[5] B. H. Calhoun and A. Chandrakasan, "Characterizing and Modeling Minimum Energy Operation for Subthreshold Circuits," *ISLPED*, 2004.

[6] G. Schrom and S. Selberherr, "Ultra-Low-Power CMOS Technologies," *Intl. Semiconductor Conf.*, pp. 237 - 246, 1996

[7] A. Bryant, et al., "Low-Power CMOS at Vdd=4kT/q," *Device Research Conference*, pp. 25-27, 2001.

[8] Ono, Atsuki, et al., "TED Control Technology for Suppression of Reverse Narrow Channel Effect in 0.1µm MOS Devices," *IEDM*, 1997.

[9] C. Piguet, "Design of Low-Power Libraries," *Intl. Conf. on Electronics, Circuits and Systems*, vol. 2 , pp. 175-180, Sept. 1998.

[10] C. Piguet, et al., "Low-Power Low-Voltage Library Cells and Memories," pp. 1521-1524, 2001.