Static Noise Margin Variation for Sub-threshold SRAM in 65-nm CMOS

Benton H. Calhoun, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—The increased importance of lowering power in memory design has produced a trend of operating memories at lower supply voltages. Recent explorations into sub-threshold operation for logic show that minimum energy operation is possible in this region. These two trends suggest a meeting point for energy-constrained applications in which SRAM operates at sub-threshold voltages compatible with the logic. Since sub-threshold voltages leave less room for large static noise margin (SNM), a thorough understanding of the impact of various design decisions and other parameters becomes critical. This paper analyzes SNM for sub-threshold bitcells in a 65-nm process for its dependency on sizing, V_{DD}, temperature, and local and global threshold variation. The V_T variation has the greatest impact on SNM, so we provide a model that allows estimation of the SNM along the worst-case tail of the distribution.**

***Index Terms*—Sub-threshold, sub-threshold memory, SRAM, static noise margin, process variation, voltage scaling.**

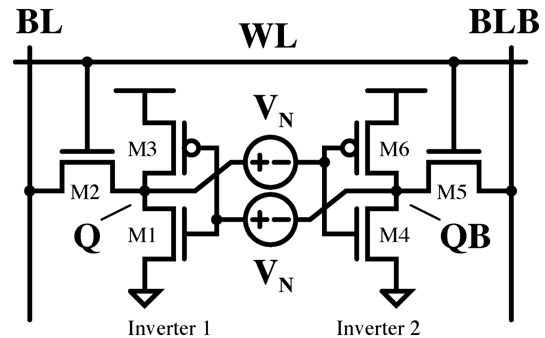

Fig. 1. Schematic for 6T bitcell showing voltage noise sources for finding SNM.

I. Introduction

SUB-THRESHOLD digital circuit design has emerged as a low energy solution for applications with strict energy constraints. Analysis of sub-threshold designs has focused on logic circuits (e.g., [1]). SRAMs comprise a significant percentage of the total area for many digital chips as well as the total power [2], [3]. For this reason, SRAM leakage can dominate the total leakage of the chip, and large switched capacitances in the bitlines and wordlines make SRAM accesses costly in terms of energy. Pushing SRAM operation into the sub-threshold region reduces both leakage power and access energy. Also, for system integration, SRAM must become capable of operating at sub-threshold voltages that are compatible with sub-threshold combinational logic. Recent low power memories show a trend of lower voltages with some designs holding state on the edge of the sub-threshold region (e.g., [4]). This scaling promises to continue, leading to sub-threshold storage modes and even sub-threshold operation for SRAMs operating in tandem with sub-threshold logic.

When the bitcell is holding data, its wordline is low so the nMOS access transistors are off. In order to hold its data properly, the back-to-back inverters must maintain bi-stable operating points. The best measure of the ability of these inverters to maintain their state is the bitcell's static noise margin (SNM) [5]. The SNM is the maximum amount of voltage noise that can be introduced at the outputs of the two inverters such that the cell retains its data. SNM quantifies the amount of voltage noise required at the internal nodes of a bitcell to flip the cell's contents.

Fig. 1 shows a conceptual setup for modeling SNM [5]. Noise sources having value V_N are introduced at each of the internal nodes in the bitcell. As V_N increases, the stability of the cell changes. Fig. 2 shows the most common way of representing the SNM graphically for a bitcell holding data. The figure plots the voltage transfer characteristic (VTC) of Inverter 2 from Fig. 1 and the inverse VTC from Inverter 1. The resulting two-lobed curve is called a "butterfly curve" and is used to determine the SNM. The SNM is defined as the length of the side of the largest square that can be embedded inside the lobes of the butterfly curve [5]. To understand why this definition holds, consider the case when the value of V_N increases from 0. On the plot, this causes the VTC^{-1} for Inverter 1 in the figure to move downward and the VTC for Inverter 2 to move to the right. Once they both move by the SNM value, the curves meet at only two points. Any further noise flips the cell.

Although the SNM is certainly important during hold, cell stability during active operation represents a more significant limitation to SRAM operation. Specifically, at the onset of a read access, the wordline is "1" and the bitlines are still precharged to "1" as Fig. 3 illustrates. The internal node of the bitcell that represents a zero gets pulled upward through the access transistor due to the voltage dividing effect across the access transistor (M_2, M_5) and drive transistor (M_1, M_4). This increase in voltage severely degrades the SNM during the read operation (read SNM). Fig. 4 shows example butterfly curves during hold and read that illustrate the degradation in SNM during read.

Manuscript received November 4, 2005; revised January 18, 2006. This work was supported by the Defense Advanced Research Projects Agency (DARPA) and by Texas Instruments Incorporated.

B. Calhoun is with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904 USA (e-mail: bcalhoun@virginia.edu).

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

Digital Object Identifier 10.1109/JSSC.2006.873215

0018-9200/$20.00 © 2006 IEEE

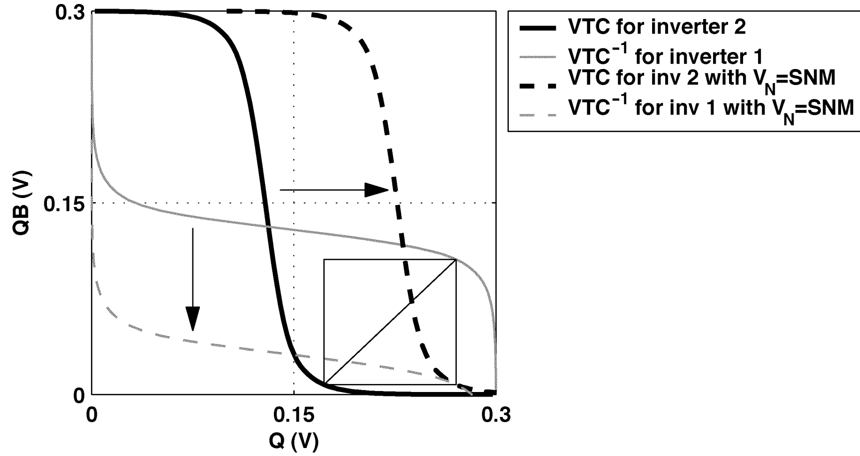

Fig. 2. The length of the side of the largest embedded square in the butterfly curve is the SNM. When both curves move by more than this amount (e.g., $V_N = \mathrm{SNM}$), then the bitcell is mono-stable, losing its data.}

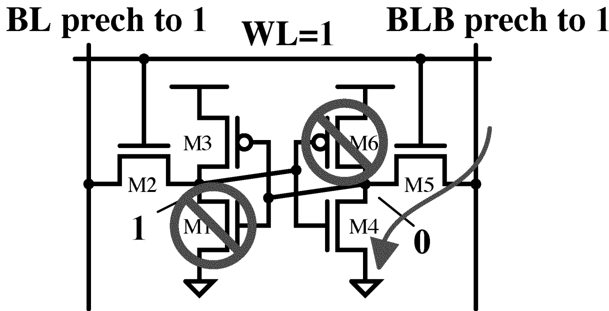

Fig. 3. Schematic of the 6T bitcell at the onset of a read access. WL has just gone high, and both BLs are precharged to V_{DD}. The voltage dividing effect across M_4 and M_5 pulls up node Q_B, which should be 0 V, and degrades the SNM.

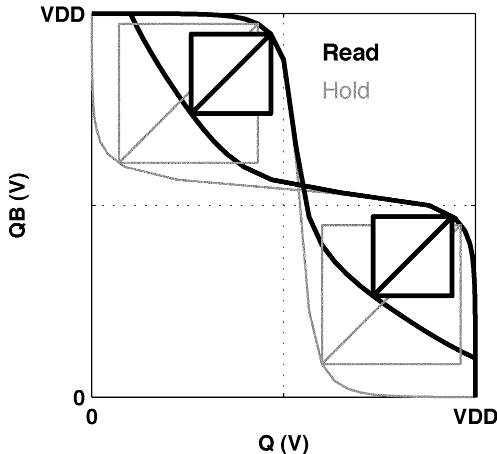

Fig. 4. Example butterfly curve plots for SNM during hold and read.

II. Static Noise Margin

This section evaluates the SNM of six-transistor (6T) SRAM bitcells operating in sub-threshold. We analyze the dependence of SNM during both hold and read modes on supply voltage, temperature, transistor sizes, local transistor mismatch due to random doping variation, and global process variation in a commercial 65-nm technology. We analyze the statistical distribution of SNM with process variation and provide a model for the tail of the probability density function (PDF) that dominates SNM failures [6].

The minimum voltage for retaining bistability was theorized in [7] and modeled for SRAM in [8], but degraded SNM can limit voltage scaling for SRAM designs above this minimum voltage. SNM quantifies the amount of voltage noise required at the internal nodes of a bitcell to flip the cell's contents.

An expression for above-threshold SNM based on long-channel models is given in [5], and [9] models above-threshold SNM for modern processes with process variation. This section builds on previous work by examining SNM for sub-threshold SRAM [6].

A. Modeling Sub-Threshold Static Noise Margin

Lowering V_{DD} reduces gate current much more rapidly than sub-threshold current, so total current in the sub-threshold region can be modeled to first order as

$$I_D = I_S \exp\left(\frac{V_{GS} - V_T}{nV_{th}}\right)\left(1 - \exp\left(\frac{-V_{DS}}{V_{th}}\right)\right). \quad (1)$$

The sub-threshold factor $n = 1 + C_d/C_{ox}$, $V_{th} = kT/q$, and I_S is the current when V_{GS} equals V_T. For simplicity, we treat pMOS parameters as positive values. For the 65-nm technology used in this section, the nMOS drive current is higher in above-threshold than the pMOS for iso-width, but the pMOS current is higher in sub-threshold due to its lower V_T. During hold mode, the wordline is low so M_2 and M_5 have $V_{GS} \leq 0$ and thus negligible current. We can model the cell VTCs ($V_{\mathrm{OUT}} = f_{\mathrm{VTC}}(V_{\mathrm{IN}})$) as those of a simple inverter in sub-threshold.

$$V_{QB} = V_{th}\frac{n_1 n_3}{n_1+n_3}\left(\ln\frac{I_{S3}}{I_{S1}}+\ln\left(\frac{1-\exp\left(\frac{(-V_{\mathrm{DD}}+V_Q)}{V_{th}}\right)}{1-\exp\left(-\frac{V_Q}{V_{th}}\right)}\right)\right) + \frac{n_1 V_{\mathrm{DD}}}{n_1+n_3}+\frac{n_1 n_3}{n_1+n_3}\left(\frac{V_{T1}}{n_1}-\frac{V_{T3}}{n_3}\right). \quad (2)$$

Referring to Fig. 2, (2) [7] gives the inverse VTC for inverter 1 ($V_{\mathrm{IN}} = f_{\mathrm{VTC}}^{-1}(V_{\mathrm{OUT}})$). The inverse of (2) is given in [10] for

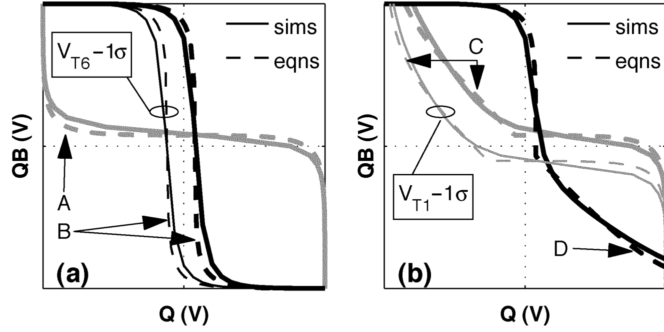

Fig. 5. First-order VTC equations versus simulation. Line A is (2), line B is (3), line C is a piecewise combination of (5) and (2), and line D is a piecewise combination of (3) and the graphical inverse of (5).

matched pMOS and nMOS (same n, V_T, I_S). We give a full solution for $V_{\mathrm{OUT}} = f_{\mathrm{VTC}}(V_{\mathrm{IN}})$ for inverter 2 in (3):

$$V_{QB} = V_{\mathrm{DD}} + V_{th} \ln\left(\frac{1-G+\sqrt{(G-1)^2+4\exp\left(\frac{-V_{\mathrm{DD}}}{V_{th}}\right)G}}{2}\right) \quad (3)$$

$$G = \exp\left(\frac{n_4+n_6}{n_4 n_6 V_{th}} V_Q - \ln\frac{I_{S6}}{I_{S4}} - \frac{V_{\mathrm{DD}}}{n_6 V_{th}} - \frac{1}{V_{th}}\left(\frac{V_{T4}}{n_4} - \frac{V_{T6}}{n_6}\right)\right). \quad (4)$$

Fig. 5(a) plots (2) and (3) against simulation curves for no local mismatch and for 1σ V_T mismatch in M_6.

During a read access, the wordline goes high and the bitlines are precharged to V_{DD} so, if $V_Q = 0$ prior to access, M_1 and M_2 are both on. This creates a voltage division that raises the voltage at Q. Assuming pMOS current is negligible in the region of interest, (5) shows the inverse VTC equation during a read operation near the SNM [4] for inverter 1:

$$V_{QB} = n_1 V_{th} \ln\frac{I_{S2}}{I_{S1}} + n_1 V_{th} \ln\left(\frac{1-\exp\left(\frac{(-V_{\mathrm{DD}}+V_Q)}{V_{th}}\right)}{1-\exp\left(-\frac{V_Q}{V_{th}}\right)}\right) + V_{T1} + \frac{n_1}{n_2}(V_{\mathrm{DD}} - V_{T2} - V_Q). \quad (5)$$

This equation cannot be inverted analytically, and it applies only to the region of the VTC where V_{OUT} is low. Fig. 5(b) shows (5) and its graphical inverse combined piecewise with (2) and (3) and plotted against simulation for no local mismatch and for 1σ V_T mismatch in M_1 for minimum device sizes at 25 °C.

Graphical or numerical solutions for SNM are easily derived from the VTC equations, although no direct analytical solution exists. The equations provide a good estimate of the behavior of the SNM based on key parameters. One shortcoming of (2)–(5) is the assumption that sub-threshold slope ($S = nV_{th} \ln 10$) is constant for each transistor. Fig. 6(a) shows that S varies with V_{GS}, and Fig. 6(b) shows S changing with temperature without the expected constant slope due to V_{th}. A more crucial problem with (2)–(5) is the assumption that certain currents are negligible. These assumptions break down under certain combinations of V_T variation, rendering the first-order equations inaccurate.

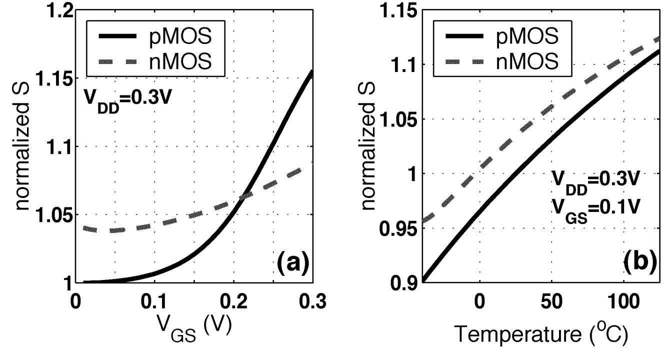

Fig. 6. Changes in sub-threshold slope (S) versus (a) V_{GS} and (b) temperature.

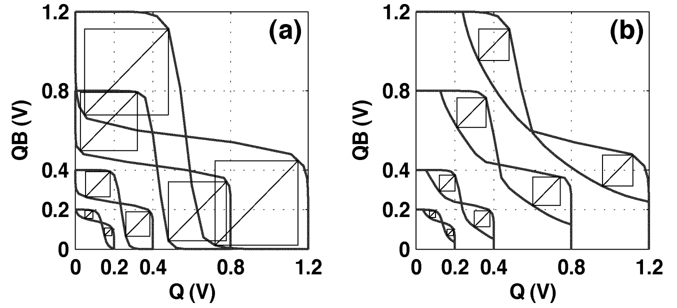

Fig. 7. VTCs for (a) hold and (b) read with varying V_{DD}.

B. Sub-Threshold SNM Dependencies

With embedded SRAM often providing multiple megabits of storage, the SNM of the nominal bitcell becomes largely irrelevant. Variations in processing and in the chip's environment create a distribution of SNM across the bitcells in a given memory, and the worst-case tail of this distribution determines the yield. This section examines the impact of different parameters on SNM in sub-threshold and offers a model for estimating the tail of the SNM density function for process variation.

SNM for a bitcell with ideal VTCs is still limited to $V_{\mathrm{DD}}/2$ because of the two sides of the butterfly curve. An upper limit on the change in SNM with V_{DD} is thus 1/2. Fig. 7 shows example butterfly curves at different supply voltages from 1.2 V to 200 mV for both hold and read. Fig. 8 plots SNM versus V_{DD} directly for both hold and read mode. The slopes of the curves confirm that less than 1/2 of V_{DD} noise will translate into SNM changes.

The impact of temperature on SNM in sub-threshold is also not large. Fig. 9 shows SNM versus temperature in sub-threshold and again for strong inversion. The sensitivity in sub-threshold is lower, and varying temperature from −40 °C to 125 °C only alters Read and Hold SNM by 21 mV and 6 mV, respectively. Higher temperatures lower SNM in sub-threshold due to the degraded gain in the inverters that results from worse sub-threshold slope (see Fig. 6(b)). Also, pMOS devices weaken relative to nMOS at higher temperature. Fig. 10 provides example butterfly plots for 0 °C and 100 °C at 1.2 V and 0.3 V.

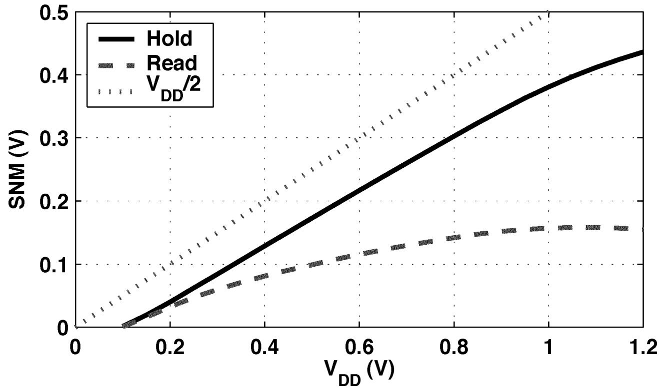

Fig. 8. SNM versus $V_{\rm DD}$.

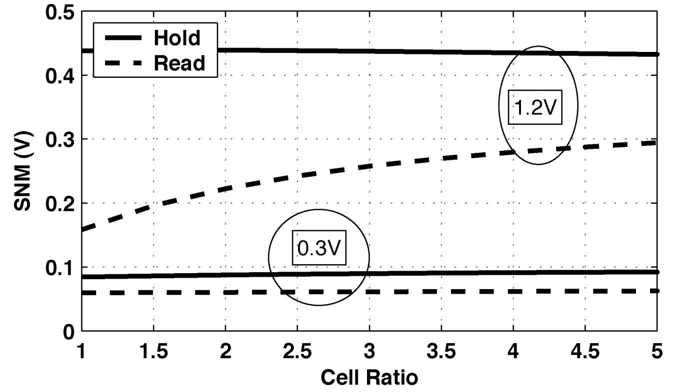

Fig. 11. Cell ratio affects SNM less in sub-threshold.

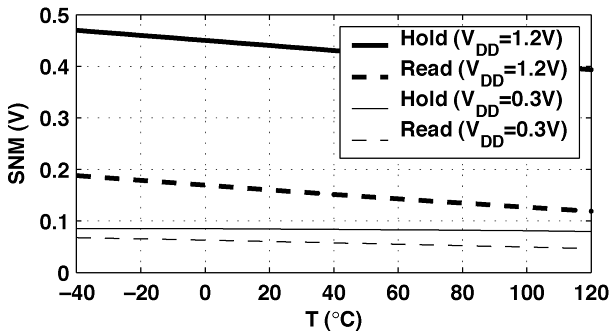

Fig. 9. SNM versus temperature.

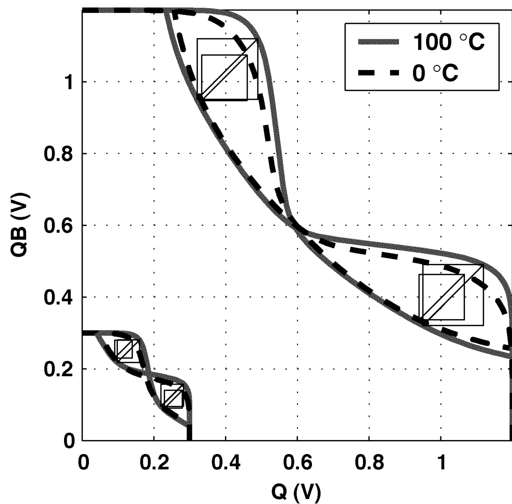

Fig. 10. VTCs during a read access across temperature.

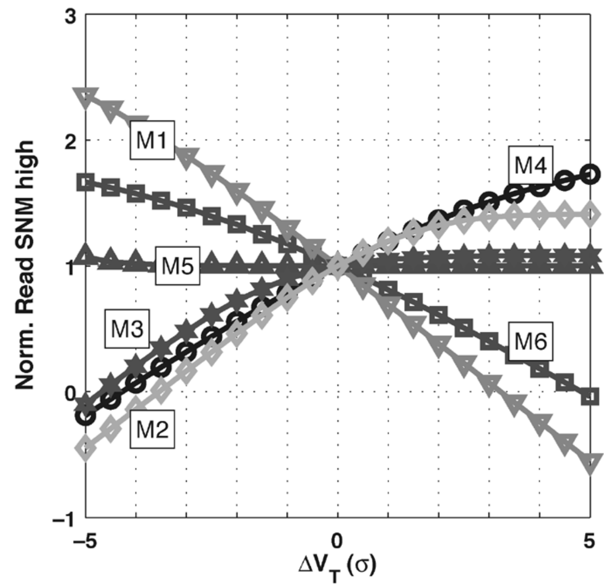

Fig. 12. Dependence of SNM high on single FETs is nearly linear.

In contrast to above-threshold [11], Fig. 11 shows that cell ratio ($(W/L)_1/(W/L)_2$ or $(W/L)_4/(W/L)_5$) has very little impact on SNM during sub-threshold read. In fact, sub-threshold SNM sensitivity to any sizing changes is reduced. The lower impact of sizing is intuitively reasonable considering the exponential dependence of sub-threshold current on other parameters. Mathematically, we can see from (2)–(5) that sizing changes affect I_{Si} linearly and only have a logarithmic impact on the VTCs. One point of caution here is that V_T for deep-submicron devices tends to vary with size as a result of narrow or short channel effects. The impact of this V_T change that might accompany a sizing change is more pronounced. These effects depend on the technology and make general SNM modeling more complicated.

C. Dependence on Random Doping Variation

The randomness of the number of doping atoms and their placement in a MOSFET channel causes random mismatch even in transistors with identical layout [12]. The impact on threshold voltage, whose σ is proportional to $(WL)^{-(1/2)}$, is the worst for minimum sized devices which are common in SRAM. Local variation is a huge problem for SRAM functionality, and it is the subject of many papers (e.g., [13], [14]). The exponential dependence of current on V_T in sub-threshold operation makes this random variation even more influential. Furthermore, the large number of bitcells in many SRAMs makes the tails (5σ–6σ) of the PDF more critical for modeling since the extreme cases are the limiting factor for yield. Previous work has shown that above-threshold SNM is nearly linear with V_T, and modeling its slope as constant allows an approximation of the joint PDF for SNM [9]. Likewise, the sensitivity of above-threshold SNM to V_T is linearized for each transistor in [15].

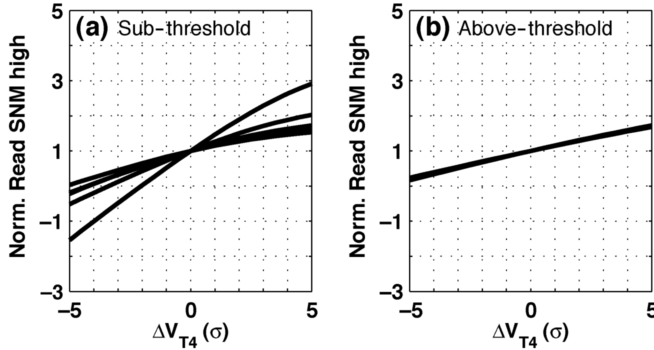

Fig. 13. Dependence of SNM high on a single FET depends on other V_Ts in (a) sub-threshold, unlike for (b) above-threshold.

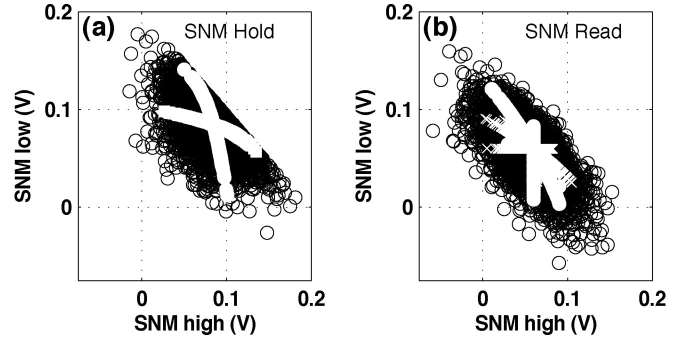

Fig. 15. Scatter plots for SNM high versus SNM low with single FET dependencies overlaid in white.

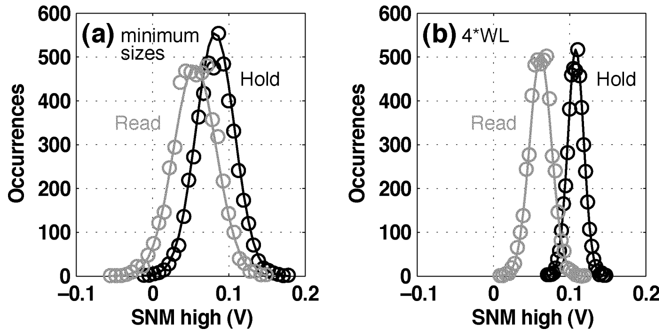

Fig. 14. SNM high and low (not shown) for (a) a minimum sized cell and for (b) $4 * WL$ is normally distributed with random V_T mismatch in all transistors.

Fig. 12 shows that, like in strong inversion, the sensitivity of SNM high (the upper-left box in Fig. 4) is nearly linear with each individual V_T. However, Fig. 13(a) shows the relationship between SNM and V_{T4} for a few different random values of the other V_Ts. The obvious dependence of the slope on the other V_Ts prevents using a model of the form $\mathrm{SNM} = \mathrm{SNM}_0 + \sum c_i V_{Ti}$ for sub-threshold SNM. The same is not true of above-threshold, shown in Fig. 13(b), for which a first order series model works well [9], [15].

Fig. 14 shows the results of 5 k-point Monte Carlo (M-C) simulations with random independent V_T mismatch in all transistors. These histograms confirm that sub-threshold SNM at the upper lobe of the butterfly curve (SNM high) is normally distributed. The solid lines show a fitted Gaussian PDF, and the markers show simulation results. Larger sizes for the bitcell clearly have the advertised effect of lowering the variance of V_T as seen in Fig. 14(b). The SNM low PDFs are very similar. The scatter plot in Fig. 15 shows that SNM high and SNM low are correlated. The dependencies for mismatch in each single transistor are overlaid in white for reference. The Hold SNM shows a saturation effect along the upper edge. SNM high and SNM low are not independent because any change to a VTC that increases the SNM at one side tends to decrease SNM at the other side.

The actual SNM that matters for a bitcell is the minimum of SNM high and SNM low. Thus, the random variable $X_{\mathrm{SNM}} = \min(X_{\mathrm{SNMhigh}}, X_{\mathrm{SNMlow}})$. Order statistics can provide us with the PDF for the minimum of n independent, identically distributed (*iid*) random variables, X_i. If f is the PDF, and F is the Cumulative Distribution Function (CDF) for X_i, the PDF of the minimum of two *iid* variables is given in (6):

$$f\left(\min(X_1, X_2)\right) = 2f_X(1 - F_X). \tag{6}$$

Although SNM high and SNM low are normally distributed with approximately the same mean and variance, we have previously shown that they are not independent. However, we are less interested in modeling the entire PDF for SNM than we are in modeling the worst-case tail. As previously stated, the tail toward lower SNM is the limiting factor. Let us assume that they are *iid*. Then we can solve for the PDF as

$$f_{\mathrm{SNM}} = 2f_{\mathrm{SNMhigh}}(1 - F_{\mathrm{SNMhigh}}) \tag{7}$$

and the CDF is simply

$$F_{\mathrm{SNM}} = 2F_{\mathrm{SNMhigh}} - (F_{\mathrm{SNMhigh}})^2. \tag{8}$$

Fig. 16 shows the histogram for a 5 k-point M-C simulation of Read SNM plotted on linear axes (a) and semilog axes (b). Clearly, SNM is not normally distributed, and its mean is lower than the mean of SNM high and SNM low. Fig. 16(b) shows that a Gaussian PDF does not match the worst-case tail on the left side of the PDF. On the other hand, the PDF based on (7) provides a good estimate of the worst-case tail. The plot shows that the model does not fit the distribution above the mean. This shortcoming results from the correlation between SNM high and SNM low. Since these two random variables are not *iid*, we cannot claim that the minimum model will always match the tail. However, we can show experimentally that it does offer a good estimate. Thus, the model is a useful tool for evaluating SNM under different design decisions and conditions. This PDF gives the powerful option of estimating the SNM at the worst-case end of the PDF without using extremely long M-C simulations until the design space is narrowed sufficiently.

Fig. 17 shows several estimated PDFs using (7) that are based on data sets of different lengths. These estimates are plotted over a 50 k-point M-C simulation. A 1000-point M-C simulation gives a modeled distribution that overlays the modeled distribution from the 50 k-point case on the plot (<3% error). Using

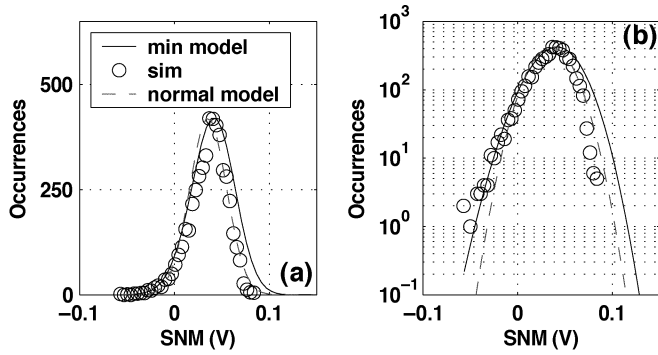

Fig. 16. (a) Histogram of SNM Monte Carlo simulation (circles) with normal PDF (dash) and PDF based on (7) (solid) over-laid. The semilog plot (b) shows that the PDF based on (7) matches the worst-case tail quite well.

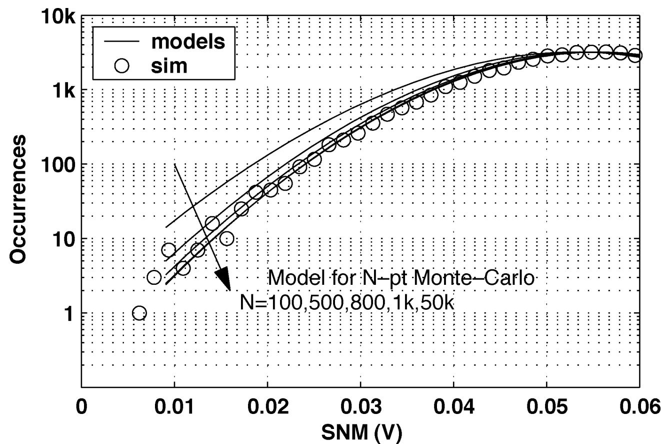

Fig. 17. 50 k-point Monte Carlo simulation for SNM with $4*WL$ sized transistors. Model based on 1 k-point Monte Carlo data matches the 50 k-point model with <3% error.

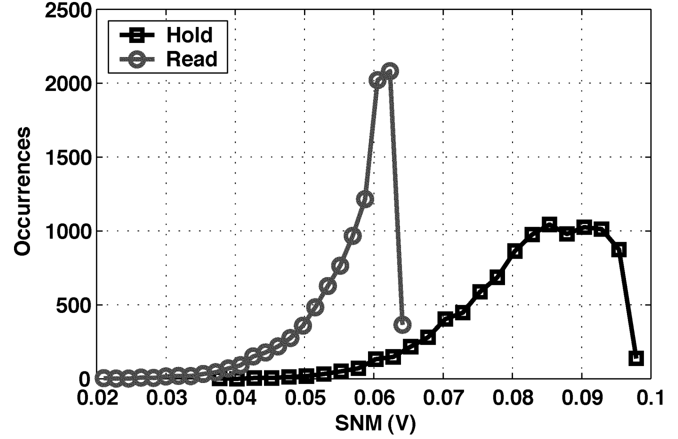

Fig. 18. Monte Carlo simulation showing global variation impact on SNM for a minimum sized bitcell.

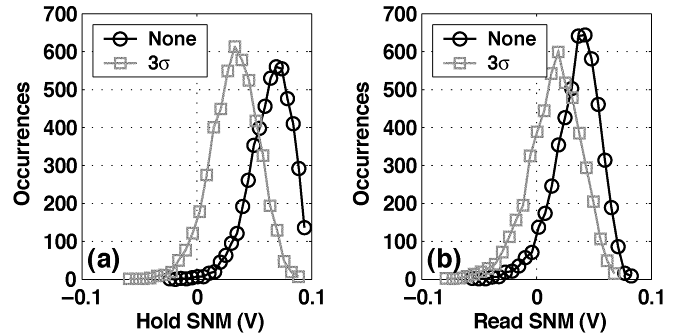

Fig. 19. SNM Monte Carlo simulations for local mismatch on top of global variation.

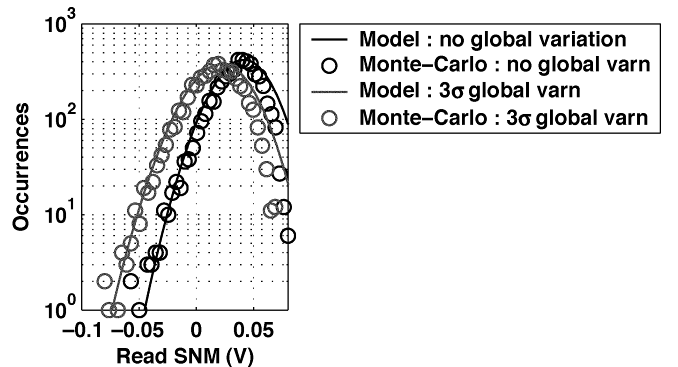

Fig. 20. SNM Monte Carlo simulations for local mismatch on top of global variation compared to the model.

this approach allows a designer to reliably estimate the tail of the SNM PDF for a large memory with relatively few samples.

Thus far we have assumed that device mismatch occurs in transistors that start off as typical for the process. In addition to the inter-die V_T mismatch that we have described is an intra-die process variation that sets the process corner (e.g., fast nMOS, slow pMOS, etc.). Even for no mismatch, the process corner impacts the SNM. Fig. 18 shows the SNM PDF for a minimum sized 6T bitcell from a M-C simulation of global process corner in which nine process parameters are varied. Here again, the tail of the PDF is the limiting factor.

In a production framework, each die containing a given SRAM will have a global process corner that affects SNM as in Fig. 18. On top of this, mismatch in each cell will result from random doping variation. Assuming that any die within 3σ of the mean is usable, we found the global process corner that gives an SNM yield with the same probability as -3σ for both hold and read cases. Fig. 19 shows that the impact of mismatch at this 3σ process corner is essentially to shift the mean of the PDF by the offset caused by global variation. This means that the models we have presented remain valid for the case of combined global and local variation. Fig. 20 shows the semilog plot of the distributions to confirm this conclusion.

III. Conclusion

Static noise margin is a critical metric for SRAM bitcell stability. This paper has explored the impact of different parameters on SNM for SRAM bitcells in sub-threshold. The dominant factor affecting sub-threshold circuits in general and SNM specifically is V_T mismatch due to random doping variation, and the critical region for examination is the tail of the SNM PDF. We have shown that first-order theoretical models for calculating SNM are accurate close to the nominal values of V_T, but they cannot accurately account for all of the mismatch cases.

We have shown that SNM high and SNM low are normally distributed with V_T mismatch and correlated. Despite their correlation, we have shown that treating them as *iid* leads to a PDF for SNM that gives an accurate model of the tail cases. This estimate is invaluable for avoiding long M-C simulations in the design of large SRAMs for sub-threshold operation.

REFERENCES

[1] 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.

[2] M. Yamaoka, Y. Shinozaki, N. Maeda, Y. Shimazaki, K. Kato, S. Shimada, K. Yanagisawa, and K. Osadal, "A 300 MHz 25 μA/Mb leakage on-chip SRAM module featuring process-variation immunity and low-leakage-active mode for mobile-phone application processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2004, pp. 494–495.

[3] N. Kim, K. Flautner, D. Blaauw, and T. Mudge, "Circuit and microarchitectural techniques for reducing cache leakage power," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 12, no. 2, pp. 167–184, Feb. 2004.

[4] A. Bhavnagarwala, S. Kosonocky, S. Kowalczyk, R. Joshi, Y. Chan, U. Srinivasan, and J. Wadhwa, "A transregional CMOS SRAM with single, logic VDD and dynamic power rails," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 292–293.

[5] E. Seevinck, F. List, and J. Lohstroh, "Static noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 5, pp. 748–754, Oct. 1987.

[6] B. Calhoun and A. Chandrakasan, "Analyzing static noise margin for sub-threshold SRAM in 65 nm CMOS," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, Sep. 2005, pp. 363–366.

[7] R. M. Swanson and J. D. Meindl, "Ion-implanted complementary MOS transistors in low-voltage circuits," *IEEE J. Solid-State Circuits*, vol. SC-7, no. 2, pp. 146–153, Apr. 1972.

[8] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, "SRAM leakage suppression by minimizing standby supply voltage," in *Int. Symp. Quality Electronic Design (ISQED) Dig. Tech. Papers*, 2004, pp. 55–60.

[9] A. Bhavnagarwala, X. Tang, and J. Meindl, "The impact of intrinsic device fluctuations on CMOS SRAM cell stability," *IEEE J. Solid-State Circuits*, vol. 36, no. 4, pp. 658–665, Apr. 2001.

[10] E. Vittoz, "Weak inversion for ultimate low-power logic," in *Low-Power Electronics Design*, C. Piguet, Ed. Boca Raton, FL: CRC Press, 2005, pp. 16-1–16-18.

[11] B. Cheng, S. Roy, and A. Asenov, "The impact of random doping effects on CMOS SRAM cell," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, 2004, pp. 219–222.

[12] R. Keyes, "The effect of randomness in the distribution of impurity atoms on FET threshold," *Appl. Phys. A: Mater. Sci. Process.*, vol. 8, pp. 251–259, 1975.

[13] S. Mukhopadhyay, H. Mahmoodi-Meimand, and K. Roy, "Modeling and estimation of failure probability due to parameter variations in nano-scale SRAMs for yield enhancement," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 64–67.

[14] M. Yamaoka, K. Osada, R. Tsuchiya, M. Horiuchi, S. Kimura, and T. Kawahara, "Low power SRAM menu for SOC application using yin-yang-feedback memory cell technology," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 288–291.

[15] K. Takeuchi, R. Koh, and T. Mogami, "A study of threshold voltage variation for ultra-small bulk and SOI CMOS," *IEEE Trans. Electron Devices*, vol. 48, no. 9, pp. 1995–2001, Sep. 2001.

Benton H. Calhoun (S'05–M'06) received the B.S. degree in electrical engineering with a concentration in computer science from the University of Virginia, Charlottesville, in 2000. He received the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA, in 2002 and 2006, respectively.

In January 2006, he joined the faculty of the University of Virginia as an Assistant Professor in the Electrical and Computer Engineering Department. His research interests include leakage reduction, sensor networks, energy-efficient circuits, memory design, and sub-threshold operation.

Anantha P. Chandrakasan (M'95–SM'01–F'04) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005).

Dr. Chandrakasan has received several awards, including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 Design Automation Conference Design Contest Award, and the 2004 DAC/ISSCC Student Design Contest Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design'98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Subcommittee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Subcommittee Chair for ISSCC 2004–2006. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the SSCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2007.