

A Black Box Method For Stability Analysis of Arbitrary SRAM Cell Structures

M. Wieckowski, D. Sylvester, D. Blaauw
Dept. Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI
{wieckows, dennis, blaauw}@umich.edu

V. Chandra, S. Idgunji, C. Pietrzyk, R. Aitken
ARM R&D
San Jose, CA

Abstract—Static noise margin analysis using butterfly curves has traditionally played a leading role in the sizing and optimization of SRAM cell structures. Heightened variability and reduced supply voltages have resulted in increased attention being paid to new methods for characterizing dynamic robustness. In this work, a technique based on vector field analysis is presented for quickly extracting both static and dynamic stability characteristics of arbitrary SRAM topologies. It is shown that the traditional butterfly curve simulation for 6T cells is actually a special case of the proposed method. The proposed technique not only allows for standard SNM “smallest-square” measurements, but also enables tracing of the state-space separatrix, an operation critical for quantifying dynamic stability. It is established via importance sampling that cell characterization using a combination of both separatrix tracing and butterfly SNM measurements is significantly more correlated to cell failure rates than using SNM measurements alone. The presented technique is demonstrated to be thousands of times faster than the brute force transient approach and can be implemented with widely available, standard design tools.

Keywords- *memory; stability; robustness; dynamic noise-margin; static noise-margin;*

I. INTRODUCTION

Hill and Lohstroh developed the original treatment of logic noise margin analysis using butterfly curves in the late 1960s and 1970s [1, 2]. Almost two decades later, Seevinck et. al. demonstrated a spice compatible simulation technique for generating these butterfly-curves specifically for SRAM cells that provided a quantitative measure of their stability and robustness [3]. This simulation technique has become the de-facto standard, and over the last 20 or so years, a variety of equivalent methods with additional merits and drawbacks have been proposed for making this same stability measurement, for example N-curves and unity-gain search [4-6]. Even so, no technique has supplanted the butterfly curve as the quintessential SRAM metric, a reality evident in literature presenting even the most advanced SRAM cells to date [7].

In recent years, increasing attention has been paid to so-called dynamic stability analysis [8]. The premise is to capture transient phenomena of SRAM cell operation that are lost in traditional DC static analysis, thus affording the designer more relaxed margins during dynamic operations, such as reading and writing, and under the constraints of transient parasitics, such as coupling and radiation induced noise. Quantifying dynamic noise margin generally involves characterizing an

SRAM cell using a multitude of transient simulations and as such, becomes particularly challenging since it must take into account external factors including bit-line capacitance, word-line driver slew, layout dependent coupling, etc. These in turn require one to make assumptions about the overall array structure during the cell design phase that can result in unnecessary iteration later on.

In this work, a method is presented for capturing and quantifying both static and dynamic SRAM characteristics without performing any transient simulations or DC convergence operations. The technique relies solely on solving model equations for predefined steady-state DC conditions, a calculation that can be performed through any number of available numerical solvers. Most importantly, the presented technique is compatible with any bi-stable SRAM cell structure including asymmetric cells and atypical transistor count designs such as 5T, 8T, and Portless SRAM.

II. BACKGROUND

It has long been established that the smallest square of the butterfly curve, the zero-current separation of the N-curve, and the unity-gain separation of a stability analysis are all mathematically equivalent methods for measuring the traditional static-noise margin (SNM) [4, 9]. For this reason, the Seevinck butterfly curve will be the only method considered in this work for the static case due to its popularity among designers in the SRAM community and its SPICE compatible implementation. For the dynamic case, only the determination of the state-space separatrix will be considered since it is a prerequisite for all the recently proposed dynamic analysis techniques.

A. Traditional Static Noise Margins

The traditional treatment of static noise margin often begins with an infinitely long chain of logic gates or inverters. The goal is to determine how much noise the input gate can tolerate before causing the output gate to change state. From this, a mathematical equivalence from the perspective of noise can be invoked to relate the infinitely long chain of gates to a simplified cross-coupled pair [10]. It is from this equivalence that the notion of defining SNM for an SRAM cell was born.

The well known butterfly curve is a graphical representation of the worst-case SNM and can be readily simulated in SPICE for a cross-coupled inverter pair as shown in Figure 1 using the Seevinck method [3]. In this schematic, U is an independent voltage source that is swept by way of DC




Figure 1. Standard butterfly curve simulation testbench

simulation. Voltage dependent voltage sources are employed to implement a standard rotation transformation of 45 degrees. When performed twice (or once for two inverters) with appropriate sign selection, plotting U versus V will yield a rotated butterfly curve from which SNM can be readily calculated by way of subtraction. It is of utmost importance to point out that this technique is open-loop for an SRAM cell in the sense that each inverter is isolated from its cross-coupled partner. The feedback seen in the equations of the schematic simulation is an artifact of the rotation transformation and as such, does not intrinsically capture the closed-loop dynamics of an actual SRAM cell. It will be demonstrated in Section 3.2 however, that this setup does in fact represent a special case of the closed-loop system.

B. Dynamic Noise Margins

In a closed-loop feedback system, such as an SRAM cell or flip-flop, characterization of stability without inadvertently altering the dynamics of the system is a difficult challenge. This is due to the fact that any probe or source used for measurement will often disturb the feedback path by sourcing or sinking current that would not otherwise exist. To cope with this paradox, one can look to dynamical systems theory whose underlying principal is to scrutinize the time-domain derivatives of a system as a means to infer its characteristic solutions and behaviors without altering the system itself.

In the case of an SRAM cell, a two-dimensional phase portrait as shown in Figure 2 can be used to represent its closed-loop dynamics. The axes of this “state-space” are the cell data node voltage and its compliment, and each point represents a different voltage combination of the cell data nodes. The stable point on the right depicts the case when the cell is storing a logical (1,0) and the stable node on the left is a (0,1) where the components of the tuple are (Cell Node, Cell Node Bar). For every point in the space, a vector is shown to represent the time derivative of both data node voltages. As such, for any starting point or initial condition, one can determine how the data nodes will evolve over time and if the system as a whole will reach a state of equilibrium. Starting at any point in the top left quadrant for example, will move along a path that over time converges to the top left stable point. It is worth noting that in the middle of the state-space exists a single metastable point. The magnitude of the time derivatives at this point are zero, but moving any amount in a direction toward one of the stable nodes will result in a stable convergence.

This type of dynamical system analysis applied to SRAM cells has recently begun to make appearances in literature as a basis for investigating robustness during read, write, and standby modes [8, 11-13]. The most important concept required




Figure 2. 45nm SRAM cell dynamical phase portrait

for this type of analysis is the determination of the cell separatrix. Conceptually, the separatrix can be thought of as the set of initial condition states that lead to metastability. This set of points forms a boundary that separates the state-space into two regions or stable manifolds as drawn in Figure 2. The locations of the stable points with respect to the separatrix effectively dictate the overall operation of the SRAM cell and depend on the transient operating point of all of its constituent devices. Of utmost importance is the case when a stable point is forced to coincide with the separatrix. This effectively forces the cell into mono-stability and is the primary mechanism responsible for writing, read disturbs, and soft-error upsets. It is for this reason that many recent definitions of dynamic stability include quantitative measures of stable point distance to the separatrix [12, 13].

III. A BLACK BOX STABILITY ANALYSIS FRAMEWORK

The brute-force method for determination of the separatrix involves simply setting a cell’s initial condition, running a transient simulation, and then iterating over all possible state-space combinations with some quantization factor. Knowing the final state after each simulation will provide the locations of the stable points, the metastable point, and the separatrix. Unfortunately, such a technique is prohibitively time consuming and can take upwards of 38 hours for a single phase portrait. A faster technique has recently been proposed that uses a custom simulator to find the metastable point and then trace the separatrix backwards in time by way of two transient simulations [11]. The method provides a several thousand times speedup over the brute-force approach.

In this work, a hybrid technique is proposed that offers comparably high-speed determination of the separatrix without requiring a custom simulator or solver. The proposed strategy gathers information about the entire state space and as such, can also be used to extract butterfly equivalent curves for arbitrary cell topologies.

A. Extracting A Cell’s Vector Field

The SRAM cell state space comprises a two-dimensional field whose extents range from the cell ground node value to

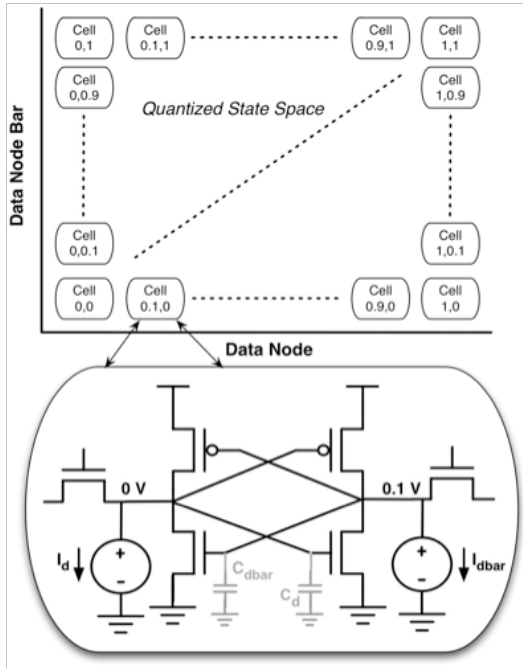


Figure 3. Quantizing the state space in the netlist

the cell supply node value. The hybrid method proposed in this work relies on quantization of this space at the netlist level. For each point of interest in the state-space, an SRAM cell is instantiated in the netlist along with two DC sources fixing its cell nodes at that point, as shown in Figure 3. Quantizing the state-space into a 50x50 grid results in a netlist with 2,500 SRAM cells and 5,000 independent voltage sources. For each cell, the voltage sources are fixed to represent one point in the state-space, and from this, the time derivative of the data node and its complement can be computed. This calculation involves simply dividing the current flowing in each DC source by the capacitance of its positive node. For each cell in the state-space array, a vector \bar{K} can be generated using the time-derivatives of the cell's data nodes as shown in Equations 1 and 2.

$$\frac{\partial V_{d(\text{bar})}}{\partial t} = \frac{I_{d(\text{bar})}}{C_{d(\text{bar})}} \quad (1)$$

$$\bar{K} = \{U, V\} = \left\{ \frac{\partial V_{d(\text{bar})}}{\partial t}, \frac{\partial V_d}{\partial t} \right\} \quad (2)$$

The most interesting aspect of using this technique is that no actual transient simulation is required, and no simulator is needed for converging on an initial DC solution. Since the DC node voltages are specified by the explicit sources, the currents and capacitances can be calculated directly from the model equations. This can be accomplished using standard simulators or by directly solving the model equations from BSIM or PSP. In this work, the Spectre simulator was used as a solver by issuing only “info captab” and “info oppoint” statements and then post-processing the results to generate the vector fields.

B. Rediscovering the Butterfly Curve

Traditionally, one of the first steps toward characterizing a system based on its phase portrait is to trace its nullcline curves [14]. These curves represent sets of vectors in the state-space where one component is null, yielding only horizontal or vertical trajectories. At any place where two nullclines

intersect, a potential stable solution exists, since at that point both time derivatives are zero. An example nullcline trace for a 6T SRAM cell during a read operation using the proposed vector field generation method is shown in Figure 4. The nullclines intersect at the three solutions to the closed loop system, only two of which are stable.

Interestingly, the nullcline trace of a 6T SRAM cell yields the same result as the static Seevinck butterfly curve. This can be explained by considering a single element of the netlist array in simplified form depicting only the cross-coupled pair and state-forcing voltage sources as in Figure 5. For the nullcline case when the I_g component is zero, the input-output combination of V_x and V_y of inverter g is one of its open-loop DC solutions. In this case, the middle source sinks enough current to satisfy the input-output combination of V_y and V_x for inverter f . This implies that all points where I_g is zero constitute the open-loop transfer curve of inverter g , yielding the same result as the butterfly simulation setup. The same is also true for the transfer curve of inverter f . This equivalence between nullclines and butterfly curves is particularly useful in that it allows one to use the proposed method for generating “butterfly equivalent” curves for SRAM cells that cannot be easily simulated using the traditional method. For example, standard butterfly analysis cannot be used to generate SNM's for Portless cells where an equalization device is present in the feedback path preventing any sort of symmetric circuit bisection [15].

C. Tracing the Separatrix

Once the vector field has been generated and the metastable point discovered via nullcline intersection, the separatrix can be found by interpolating and tracing the vectors from the metastable point to the state-space boundary. The algorithm proposed in this work starts by first defining an area of interest around the metastable point, since this point is guaranteed to be a separatrix point by definition. For any two quantized state-space points A and B within this area, additional separatrix points can be found by linearly interpolating the vectors between them along the line AB as shown in Figure 6. If an




Figure 4. Nullcline trace of a 45nm 6T cell




Figure 5. SRAM cell nullcline / butterfly equivalence

interpolated vector can be found that points at the metastable point S_m , then it is guaranteed to be a separatrix point. By repeating this interpolation process for many paths within the area of interest, coverage of the separatrix can be maximized since each path can cross the separatrix at a different location. The remainder of the separatrix can then be traced by re-centering the area of interest about the most distant separatrix point available and iterating in both directions. Tracing stops when the state-space boundary is reached or when no new points can be found.

IV. EXPERIMENTAL RESULTS

The proposed method from simulation through separatrix tracing was implemented in a Python application which generated the required netlists and invoked Spectre to solve the model equations and extract nodal capacitance. The time breakdown for a complete analysis of a 6T cell with 50x50 state-space quantization in a 45nm PSP based commercial technology is shown in Table 1. The total execution time represents a several thousand times speedup over the brute-force technique estimated in [12].

TABLE I. TIME BREAKDOWN OF PROPOSED METHOD

Stage	Time (s)	Fraction Of Total
Model Solution	27.8	84%
Post Process	0.3	< 1 %
Nullcline Trace	0.2	< 1 %
Separatrix Trace	4.9	15%

To validate the assertion that the Seevinck butterfly curve is equivalent to the dynamical null-cline trace, the same 6T SRAM cell was characterized including uniform random threshold voltage mismatch for each of its six transistors. Over 500 runs, the standard butterfly SNM and the nullcline SNM were measured using the smallest square technique with a 50x50 state-space grid. The correlation between them is shown in Figure 7. As expected, the standard butterfly SNM and the vector nullcline SNM are highly correlated. The outliers occur due to rounding errors that arise when the smallest square is comparable to the vector quantization for small SNM's, and can be improved by choosing a different the state-space quantization factor.

A. A New Metric for Cell Robustness

The most important question surrounding the use of the smallest square SNM as a metric for characterizing cell robustness is whether it actually quantifies a cell's probability of failure, specifically during read operations and standby. Conventionally, it has been assumed that static noise in the form of threshold voltage shifts or sizing variations would manifest itself through a displacement of the inverter transfer curve or nullcline trace. This displacement would adversely affect the size of the butterfly openings and result in a smaller




Figure 6. Multi-path vector interpolation to find separatrix points

SNM based on the least square measurement. It then follows that a cell with a larger nominal SNM would be more tolerant to process variation and mismatch.

To validate this traditional notion of static-noise margin, an optimized 45nm 6T SRAM cell was characterized for robustness during a read operation through a series of 20 thousand Monte Carlo runs. To capture a reasonable failure rate with so few samples, importance sampling was used and each device had its threshold voltage mean shifted 5σ in the worst case for a read upset [16]. Simultaneously, a standard butterfly SNM was measured for each nominal run before any variation was introduced. The first result for this "fixed sizing" case is shown in Figure 8. The tight grouping of the failure probability demonstrates the effectiveness of the importance sampling technique. Along with the fixed sizing, simulations were carried out for cells that were upsized for all device dimensions over a range of 1.25X to 3X. This increase in device area effectively reduces the variance of its threshold voltage, and in turn it is expected that these larger cells exhibit a lower probability of failure. There is a strong correlation between the failure rate and the butterfly SNM measured without any mismatch or variation. The sensitivity of this correlation however, is quite low. Over the entire range of failure probability, the SNM only changes by approximately 10%. This makes sizing optimization via SNM characterization rather difficult because high measurement precision is required.

In this work, we propose a new, more direct measure of robustness that makes use of the gain information within the state-space. For each stable node, a search is performed along the separatrix to find which of its points is closest to that stable node. This is depicted in Figure 9, where D_A and D_B represent the shortest vectors from a stable node to the separatrix. The

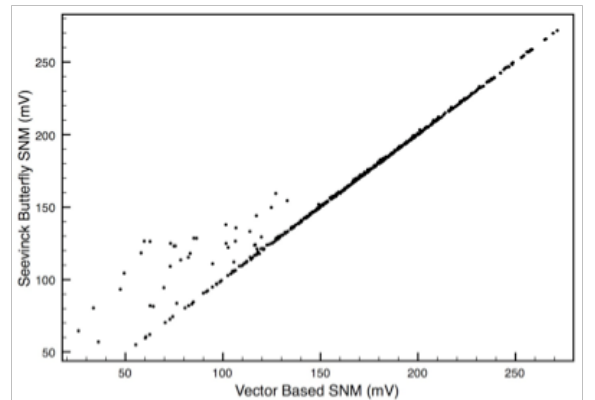


Figure 7. Correlation between SNM measurement techniques

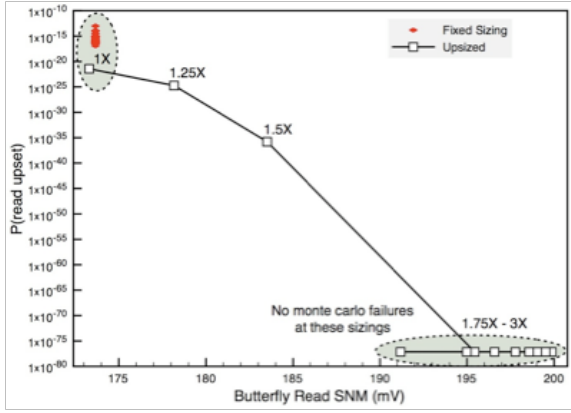


Figure 8. SNM correlation to probability of read upset

magnitudes of the cell's dV/dt vectors are then integrated from each stable node to its closest separatrix point, and the minimum vector per unit length is taken as a measure of the cell robustness. In the following sections, this metric is referred to as the separatrix affinity.

Separatrix affinity is rooted in the idea that overall failure probability of a cell can be correlated to the readiness of one of its stable states to cross the separatrix. This is identical in principle to the smallest square butterfly heuristic, but is more likely to capture any nonlinearity or distortion in the butterfly curve and separatrix resulting from device mismatch. Statistically, failure resulting from static or dynamic noise is most likely to involve separatrix crossings that are closest to the stable points. As such, integrating the cell vector lengths along the shortest path to the separatrix is reflective of how likely the stable states are to traversing that path. The separatrix affinity metric can be expressed as in Equation 3, where SA is the proposed affinity metric, D is the vector from the stable point to the separatrix, A and B are stable points, and d is the numerical integration step size. In the case of read upset stability, a lower separatrix affinity is desirable as it represents a lessened sensitivity of the cell to mismatch.

$$SA = \text{Min} \left[\frac{1}{\|D_{A,B}\|} \sum_{A,B \text{ Stable}}^{A,B \text{ Separatrix}} \sqrt{\left(\frac{\partial V_{dbar}}{\partial t}\right)^2 + \left(\frac{\partial V_d}{\partial t}\right)^2} \Delta d \right] \quad (3)$$

Monte Carlo with importance sampling was performed again for the same sizing conditions as in Figure 8 and the separatrix affinity was post-processed from the state-space array using a standard midpoint cubature integral. The result is shown in Figure 10 where the proposed robustness metric exhibits high correlation to the predicted failure rate. As the separatrix affinity increases with cell downsizing, the probability of failure also rises. It is worth noting that unlike the traditional SNM measurement, the affinity metric exhibits a range of almost 200% from the 1X sizing to the zero failure point. In addition, the separatrix affinity metric continues to decrease with sizing even after the importance sampling failure floor has been reached. As such, the separatrix affinity can easily serve as a qualification factor for cell optimization.

B. Qualifying Transient Operations And Arbitrary Cells

One of the more interesting benefits offered by the proposed hybrid analysis technique is the ability to ascertain the butterfly curve, the separatrix, the entire state-space vector




Figure 9. Example state-space illustrating new robustness metric as an integral from stable nodes to closest separatrix points

field, and the separatrix affinity very quickly for any DC combination of bitline, wordline, supply, and threshold voltages. This allows the designer insight into the sensitivity of the separatrix affinity and the butterfly smallest square to these parameters and any other environmental factors that may be of interest at the cell design level for any type of cell structure. An example is shown in Figure 11 where one bitline of a 45nm 6T cell is lowered with its wordline asserted during a write operation. For each of the four bitline voltages shown, the butterfly curve changes shape and the separatrix is deformed toward one of the stable states. For the 350 mV bitline voltage, the separatrix at point M is nearing the stable point C . When these two points finally coincide, a write operation will be complete since only one stable state will be realizable forcing the cell to hold one particular logical combination. In this way, a designer can readily calculate the required bitline voltage differential to perform a write operation without any transient simulations. More importantly, this value can be determined intrinsically without having to make any assumptions regarding slew, loading, etc. Similar qualifications can be formulated for characterizing methods such as wordline boosting, asymmetric cell designs, and drowsy supply voltages to name a few examples.




Figure 10. Correlation between proposed separatrix affinity metric and probability of failure due to read upset

As mentioned in Section III.B, the proposed method also allows one to characterize cells with unusual structures that could not otherwise be simulated easily using the standard Seevinck butterfly method. One such example is the five-transistor Portless SRAM cell [15]. Since the fifth transistor acts as an equalization device, there is no way to symmetrically break the cell for DC analysis using the method shown in Figure 1. Using the proposed technique instead, such a circuit can be quickly characterized yielding nullcline curves, a separatrix trace, a separatrix affinity measurement, and a smallest square SNM. The result for a 45nm example during a read operation is shown in Figure 12, and it is the first reported butterfly curve for this type of SRAM cell.

V. CONCLUSIONS

A black box technique for the robustness characterization of arbitrary SRAM cell structures is presented. Using a hybrid method of combining full netlist level state space quantization with a new separatrix tracing algorithm, the presented approach yields butterfly equivalent curves, separatrix traces, and full state-space vector measurements. It is demonstrated that closed-loop dynamical nullcline traces are equivalent to standard open-loop butterfly curves, and that the least square SNM measurement is well correlated to probability of read upset failures. Furthermore, a new separatrix affinity metric is developed to cope with nonlinear distortions in the nullcline and separatrix due to random mismatch. It is shown that this new metric is also well correlated to read upset failure rates, and more importantly, that it is more sensitive to the cell robustness than the standard SNM measurement. This provides a larger ratio of separatrix affinity to cell robustness and in turn, lessens the requirements on measurement accuracy when compared to the standard butterfly SNM. Lastly, applications of the proposed methodology are presented for qualifying various transient operations without explicit transient simulations. Measurements such as minimum differential bitline voltage during a write, required wordline boosting levels, and smallest data retention voltage are several possible examples that can be easily performed without regard to typical transient simulation assumptions.




Figure 11. Evolution of 6T cell stability during write




Figure 12. Nullclines and separatrix of a Portless 5T cell

REFERENCES

- [1] C. F. Hill, "Definitions of noise margin in logic systems," Mullard Technology Communications, pp. 239-245, Sept. 1967.
- [2] J. Lohstroh, "Static and dynamic noise margins of logic circuits," IEEE Journal of Solid-State Circuits, vol. 14, pp. 591-598, 1979.
- [3] E. Seevinck, F. J. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," IEEE Journal of Solid-State Circuits, Jan. 1, 1987.
- [4] E. Grossar, M. Stucchi, K. Maex, and W. Dehaene, "Read Stability and Write-Ability Analysis of SRAM Cells for Nanometer Technologies," IEEE Journal of Solid-State Circuits, Jan. 1, 2006.
- [5] M. Samson and M. B. Srinivas, "Analyzing N-Curve Metrics for Sub-Threshold 65nm CMOS SRAM," Nanotechnology, Jan. 1, 2008.
- [6] K. Agarwal and S. Nassif, "Statistical analysis of SRAM cell stability," in IEEE Design Automation Conference, 2006, pp. 57-62.
- [7] Y. Wang, U. Bhattacharya, F. Hamzaoglu, P. Kolar, Y. Ng, L. Wei, Y. Zhang, K. Zhang, and M. Bohr, "A 4.0 GHz 291Mb Voltage-Scalable SRAM Design in 32nm High- κ Metal-Gate CMOS with Integrated Power Management," in International Solid State Circuits Conference, 2009, pp. 456-458.
- [8] Z. Bin, A. Arapostathis, S. Nassif, and M. Orshansky, "Analytical Modeling of SRAM Dynamic Stability," in IEEE/ACM International Conference on Computer-Aided Design, 2006, pp. 315-322.
- [9] J. Lohstroh, E. Seevinck, and J. de Groot, "Worst-case static noise margin criteria for logic circuits and their mathematical equivalence," IEEE Journal of Solid-State Circuits, vol. 18, pp. 803-807, 1983.
- [10] J. R. Hauser, "Noise margin criteria for digital logic circuits," IEEE Transactions on Education, vol. 36, pp. 363 - 368, Nov. 1, 1993.
- [11] G. M. Huang, W. Dong, Y. Ho, and P. Li, "Tracing SRAM separatrix for dynamic noise margin analysis under device mismatch," Behavioral Modeling and Simulation Workshop, Jan. 1, 2007.
- [12] W. Dong, P. Li, and G. M. Huang, "SRAM dynamic stability: Theory, variability and analysis," International Conference on Computer-Aided Design, Jan. 1, 2008.
- [13] J. Wang, S. Nalam, and B. H. Calhoun, "Analyzing static and dynamic write margin for nanometer SRAMs," ISLPED, Jan. 1, 2008.
- [14] E. J. Beltrami, Mathematics for dynamic modeling, 2 ed.: Academic Press, 1998.
- [15] M. Wiecekowsky, S. Patil, and M. Margala, "Portless SRAM - A High-Performance Alternative to the 6T Methodology," IEEE Journal of Solid-State Circuits, vol. 42, pp. 2600-2610, 2007.
- [16] G. K. Chen, D. Blaauw, T. Mudge, D. Sylvester, and K. Nam Sung, "Yield-driven near-threshold SRAM design," in International Conference on Computer-Aided Design, 2007, pp. 660-666.