

## An Enhanced Canary-Based System With BIST for SRAM Standby Power Reduction

Jiajing Wang, Alexander Hoefler, and Benton H. Calhoun

**Abstract**—To achieve aggressive standby power reduction for static random access memory (SRAM), we have previously proposed a closed-loop  $V_{DD}$  scaling system with canary replicas that can track global variations. In this paper, we propose several techniques to enhance the efficiency of this system for more advanced technologies. Adding dummy cells around the canary cell improves the tracking of systematic variations. A new canary circuit avoids the possibility that a canary cell may never fail because it resets into its more stable data pattern. A built-in self-test (BIST) block incorporates self-calibration of SRAM minimum standby  $V_{DD}$  and the initial failure threshold due to intrinsic mismatch. Measurements from a new 45 nm test chip further demonstrate the function of the canary cells in smaller technology and show that adding dummy cells reduces the variation of the canary cell.

**Index Terms**—Built-in self test (BIST), data retention voltage (DRV), standby power, static random access memory (SRAM), variation.

### I. INTRODUCTION

Since SRAM/Cache continues to be the largest and most dense component in many digital systems or system-on-chips (SoCs), its leakage power dominates the overall leakage power of the system. One of the most effective leakage reduction techniques is supply voltage ( $V_{DD}$ ) scaling. All the leakage current components, including sub-threshold leakage, gate leakage, and junction leakage current, decrease dramatically with a smaller  $V_{DD}$ . Leakage power decreases even more rapidly due to the reduction of both  $V_{DD}$  and leakage current. Many designs have exploited  $V_{DD}$  scaling during standby and/or active operation for SRAM leakage power reduction [1]–[4]. However, the scaled  $V_{DD}$  not only reduces cell stability itself but also heightens the sensitivity of cell stability to mismatch. The data retention voltage (DRV) is the minimum  $V_{DD}$  for the cell to preserve its data [3]. Local variation spreads the DRV of the cells across the chip. To preserve all the data in an SRAM,  $V_{DD}$  must be above the DRV of the worst cell within the SRAM array, which we call standby  $V_{min}$  in this paper. Standby  $V_{min}$  varies with process variations, voltage fluctuations, and temperature changes (PVT variations). Thus we must address this  $V_{min}$  variability when choosing standby  $V_{DD}$ .

The most straightforward solution is the worst-case based open-loop approach, in which the standby voltage is picked based on the DRV for the worst scenario at design time and maintains unchanged for all the scenarios. Although it is robust, substantial power and energy are wasted because of two reasons. First, the worst PVT scenario only occurs in extreme conditions like extremely high temperature, which is rare for most applications. Second, the margin for the worst PVT protection can be quite large, and it even becomes larger as CMOS technology continuously scales.

Manuscript received May 14, 2009; revised October 13, 2009. This work was supported in part by SRC & FCRP C2S2.

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

A. Hoefler is with Freescale Semiconductor, Austin, TX 78729 USA (e-mail: alexander.hoefler@freescale.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2010.2042184

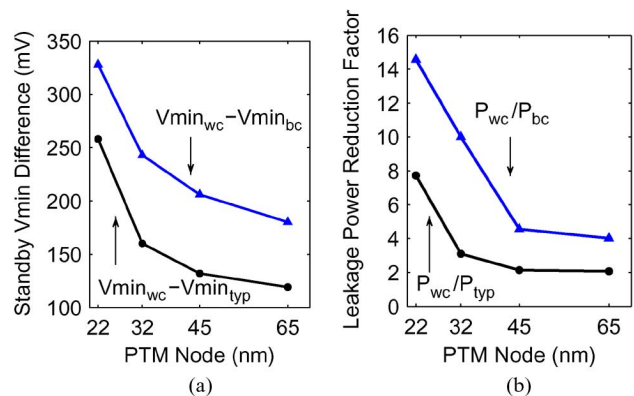


Fig. 1. (a) Margin between the standby  $V_{min}$  under the worst-case PVT variation ( $V_{min_{wc}}$ ) and that under the best/typical case ( $V_{min_{bc}}/V_{min_{typ}}$ ) and (b) leakage power reduction by using the true  $V_{min}$  at the best/typical case instead of  $V_{min_{wc}}$ . A 1-kb SRAM is simulated across the PTM bulk technologies from 65 to 22 nm.

Fig. 1(a) shows the standby  $V_{min}$  margin between the worst-case PVT variation and the best-case/typical variation increases as technology scales for a 1-kb SRAM array using predictive technology models (PTMs) [5] from 65 to 22 nm. Fig. 1(b) shows that up to 4 $\times$  leakage power reduction can be achieved if the margin is removed for the 65 nm node. For the 22 nm node, the best-case leakage power reduction increases to 14 $\times$  and savings for typical silicon increase to 8 $\times$ . Thus using the optimum  $V_{min}$  instead of the worst  $V_{min}$  becomes more appealing in smaller technologies. We have proposed an adaptive approach that can tune  $V_{DD}$  closer to the optimum  $V_{min}$  point for each global PVT condition during standby operation. It scales  $V_{DD}$  in a closed-loop fashion based on the feedback from canary replicas, which can track the impact of PVT variations on SRAM DRV [6], [7].

In this paper, we propose several improvements for variation adaptation and self calibration to extend this approach for SRAMs at 45 nm and beyond. We propose to add dummy cells around the canary cell so that it behaves more like a core SRAM cell in the presence of variation. We also propose a new canary circuit to avoid the possibility that the canary cell may never fail because it resets into its more stable data pattern. We incorporate a built-in self-test (BIST) block to self-calibrate SRAM standby  $V_{min}$  and the initial failure threshold due to intrinsic mismatch after manufacture. We implement the canary system on a new 45 nm bulk test chip. Measured results indicate that the canary cells can fail at regular intervals above the worst DRV of SRAM cells, although the distribution of SRAM DRV becomes wider due to increased variation with technology scaling. Measurements also confirm that the variation of the canary DRV is reduced with dummy cells. The remainder of this paper is organized as follows. We briefly review the principle of the closed-loop  $V_{DD}$  scaling scheme in Section II. Then we present the improvements on the canary cell in Section III. Section IV presents the BIST for self-calibration of the canary system. Section V describes the measurement results from the 45 nm test chip. We draw conclusions in Section VI.

### II. CANARY SCHEME REVIEW

Fig. 2(a) shows the example architecture of our canary scheme [6]. An on-chip or off-chip voltage regulator supplies  $V_{DD}$  to the SRAM array and to the canary banks. Several banks of canary cells are designed to fail across a range of voltages above the DRV of the SRAM cells as illustrated in Fig. 2(b), and their failures are monitored by the online failure detectors. A programmable failure threshold determines

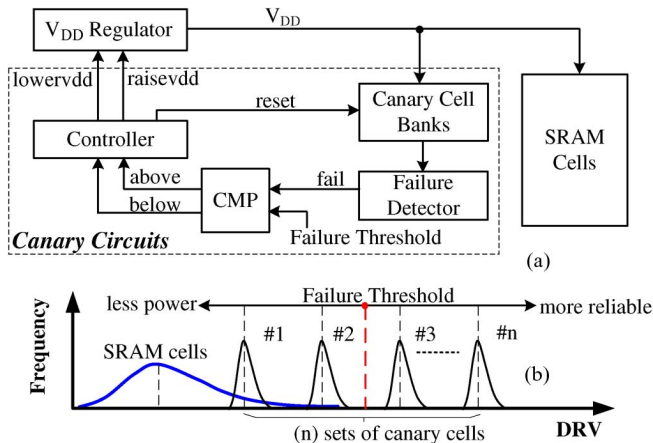


Fig. 2. (a) Example of architecture and (b) principle of canary-based closed-loop  $V_{DD}$  scaling approach.

the proximity of the applied standby  $V_{DD}$  to the tail of the SRAM DRVs, and it enables the tradeoff between power saving and data reliability. When entering the standby mode, the controller starts lowering  $V_{DD}$  until the canary failures meet the failure threshold. Once the global stimuli occur, the canary failures may exceed or drop below the failure threshold, which triggers the controller to raise or lower  $V_{DD}$  accordingly. The canary system was first successfully implemented on a 90 nm bulk test chip, and the measurement results from that chip showed that it offered  $\sim 5\times$  power reduction over the worst-case approach for the typical operating condition [7].

### III. CANARY CELL IMPROVEMENT

#### A. New Canary Cell Structure

The most critical component in our system is the canary cell. It must duplicate the impact of global stimuli on SRAM cell stability. In addition, it must fail ahead of all the SRAM cells to prevent the loss of data in SRAM. Hence, we have proposed to add a pMOS header on a standalone 6T cell as the canary cell [6]. By tuning the gate voltage of the header (VCTRL), the canary DRV can be altered in a wide range. To improve the correlation of global effects on canary cells and SRAM cells, here we further propose to add dummy 6T cells around the functional 6T cell in the canary cell to mimic the real physical environment of an SRAM cell (see Fig. 3). To reduce area cost, we use a  $3 \times 3$  SRAM mini-array for each canary. A failure detector monitors the active cell in the center. To ensure the canary cell behaves more like SRAM cells in the presence of systematic variations, we use the same layout pattern as the SRAM array except for minor changes on metal wires for pulling out the storage nodes of the central cell. Both SRAM cells and canary cells use logic rules in our test chip. The actual power supply of the mini array ( $VV_{DD}$ ) is connected with the pMOS header. As before, when we tune VCTRL to a higher value, the pMOS header is partially turned on, which causes the canary cell to operate at a lower effective  $V_{DD}$  than that seen by the core cells.

#### B. New Circuit for Canary Cell Reset

1) *Issue:* Since one cell can either hold a “0” or “1”, we previously built each canary set with two separate cells for storing “0” and “1”. The canary set fails when either the canary cell “0” or the canary cell “1” fails. Although this method is simple and easy to implement, it has one drawback. Mismatch causes a cell to be more stable at one data value than the other, and it is uncertain which data value is more stable due to randomness of local variation (e.g., from dopant fluctuation). For one canary set, if both the canary cell “0” and the canary cell “1”

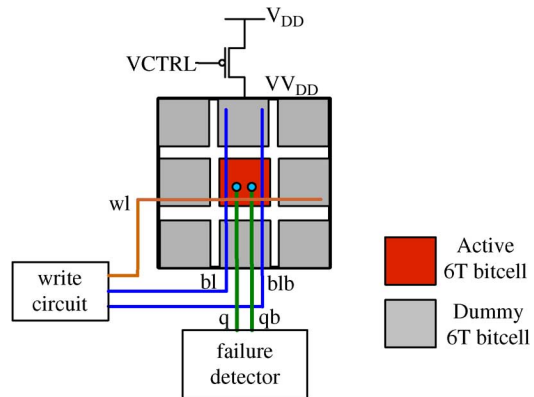


Fig. 3. New canary cell structure with dummy cells. Only modification to active 6T layout is connecting to the internal nodes q and qb.




Fig. 4. Correlation between DRV0 and DRV1 (a) when they come from the same cell and (b) when they come from separate cells. 100 samples are plotted.

happen to be more stable at the value that they are holding, this canary set will never fail or fail at a very low supply voltage regardless of the VCTRL value.

This can be better explained with the help of DRV. We denote DRV0 and DRV1 as the DRV for holding “0” and “1”, respectively. Fig. 4(a) shows the correlation between DRV0 and DRV1 when both come from the same cell with 100-point Monte Carlo simulations. Most of the samples have one DRV value near or equal to 0 and the other much greater than 0 because device mismatch causes the cell to be unbalanced. A few samples have two values close to each other because they are more balanced. However, DRV0 and DRV1 never simultaneously equal to 0. Now if the DRV1 comes from a separate cell, the correlation map between DRV0 and DRV1 changes to Fig. 4(b). Roughly 20% of samples have both DRV0 and DRV1 near or equal to 0, which means both cells can hold their respective specific data at any voltage. We observed this issue in our first test chip. Although we can use more redundant canary sets to mitigate this issue, it degrades the accuracy of the tracking performance as well as the area efficiency.

2) *Solution:* To eliminate this issue, we propose a new circuit shown in Fig. 5(a) that automatically stores the least stable data value in each canary cell. Besides the mini-array in Fig. 3 (simplified as a 6T cell for illustration here), the circuit includes a latching voltage-mode sense amplifier (SA), a D-latch, and two MUXs. Fig. 5(b) shows the timing waveforms. There are three phases: the restoring, latching, and writing phase. In the restoring phase, “VCTRL” first rises to a high value. This turns off the pMOS header and leaves the actual power of the 6T cell ( $VV_{DD}$ ) floating. We boost “VCTRL” to  $1.2V_{DD}$  so that the cell leakage drops  $VV_{DD}$  below the DRV to reset the cell. After VCTRL returns back to 0, the storage nodes (q, qb) restore the cell’s more-stable state [e.g., (0, 1) in Fig. 5(b)]. Then the latching phase starts with the rise of “saen”, which enables the SA so the stable state can be passed

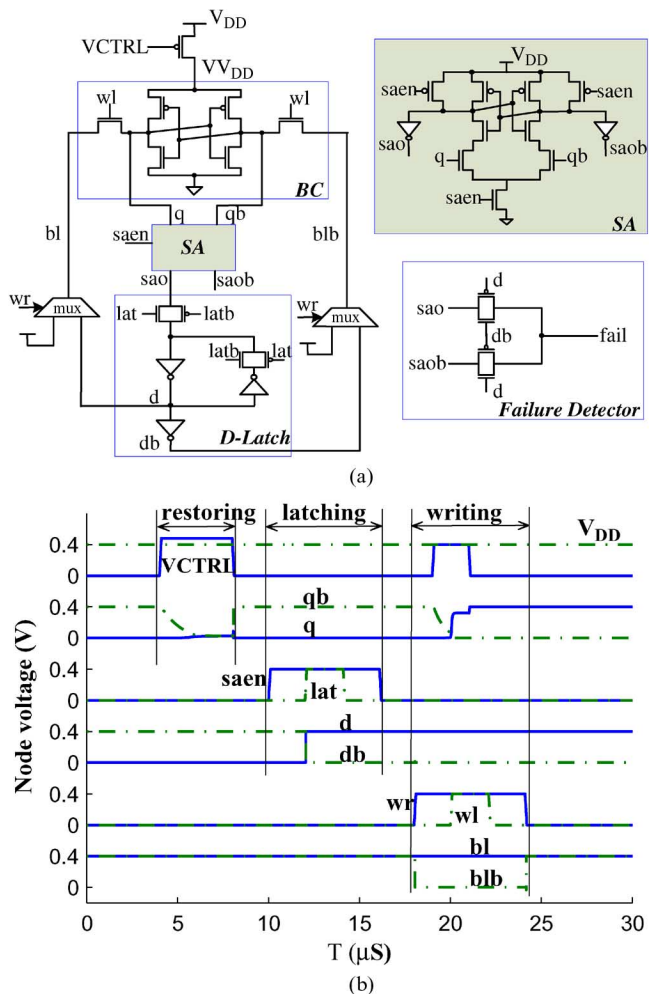


Fig. 5. (a) Circuits and (b) waveforms for canary cell self-loading its less-stable state.

to the SA outputs (sao, saob). After some delay time, a pulse on the “lat” signal allows the D-latch to capture the inverted “sao” value into its output “d”. So (d, db) are driven to the values of the cell’s less-stable state (1, 0). “saen” falls back to 0 at the end of the latching phase to disable SA. In the last writing phase, “wr” rises first so that the MUX can select the value of (d, db) for the bitlines. Then a pulse of “wl” writes the less-stable state (1, 0) into the canary cell, which ensures the canary cell will flip to its more-stable state at its increased DRV. To enhance writability, we also design the option to raise “VCTRL” to  $V_{DD}$  and float the supply of the cell during write.

In Fig. 5(a), we also show the failure detector, which performs XOR on (d, db) and (sao, saob). Once their values differ, it implies the cell has flipped and the “fail” signal will be asserted. The restoring and latching phase only occur once for the system’s entire operation (e.g., at start up). Before entering the standby mode, the writing phase is performed to reset the canary cell with its less-stable state stored in the D-latch. During standby, once the canary failures exceed the failure threshold, the standby  $V_{DD}$  will increase and then the writing phase will occur again to reset the cell. In addition, the canary cells are periodically rewritten and reevaluated so that  $V_{DD}$  can be lowered if the canary failures become smaller than the failure threshold, in which case the global condition improves during standby. Note that the supply voltage of the D-latch directly connects to the  $V_{DD}$  of the SRAM cells. Less local variation occurs in D-latches with larger devices, so the D-latch can hold its data more reliably than SRAM cells during standby operation.

#### IV. BUILT IN SELF TEST (BIST)

Simulation and measurement from a 90 nm test chip have demonstrated that the canary cells can successfully track global variation. However, the canary cells cannot directly track local variation (i.e., mismatch) without a large population of instances. Thus we have to deal with local variation separately. We have previously proposed a fast and accurate model to estimate SRAM DRV tail under local random variation [7]. In this paper, we propose an alternative method to modeling. We incorporate a BIST to detect the initial SRAM  $V_{min}$  due to intrinsic local mismatch after manufacture at one global condition. We use this value to set the initial failure threshold for the canary cells, which then can track global PVT changes during operation.

##### A. Measuring the SRAM DRV Tail

Based on the direction of searching  $V_{DD}$ , there are three methods to measure the standby SRAM  $V_{min}$  using the BIST: the downward, upward, and binary searching. Among them, the binary searching is the fastest one, but its circuit implementation is most complicated. To reduce circuit complexity, we choose either the downward or upward searching. From our simulation and measurement results, standby  $V_{min}$  is typically below half of the nominal  $V_{DD}$  for a moderate-scale SRAM (e.g., 256 kb) under normal condition. Thus the upward searching requires less iterations. In addition, the upward searching stops checking the remaining cells and increases  $V_{DD}$  by one step once the number of failures exceeds the tolerable error limit. In contrast, the downward searching must check all the SRAM cells to ensure that the total number of errors is within the tolerable limit before decreasing  $V_{DD}$  by one step. Therefore, we choose the upward searching method to save more test time. Simulation results for a 256-kb SRAM show that upward searching is about 3 times faster than downward searching when standby  $V_{min}$  is 0.5 V. For each iteration of upward searching, the BIST first checks failures for holding “0” and then for holding “1”. This process is repeated after increasing  $V_{DD}$  by one step until checking both “0” and “1” complete successfully.

Row/column redundancy and ECC are conventionally used for reducing the yield loss due to manufacturing defects and soft errors. For low standby power operation, they can also be used to tolerate data-retention errors so that the minimum standby voltage can be less than the worst DRV in the SRAM [8]. The detailed flow for checking hold failures is illustrated in Fig. 6. First, in active mode ( $V_{DD} = V_{DD,nom}$ , the nominal value), “0”/“1” is written into each address. Then the SRAM enters standby mode ( $V_{DD} = V_{DD,sb}$ , the standby value) and maintains standby for a period of time ( $T_{sb}$ ). After the standby operation, data is read out and checked in active mode. If the number of failed bits is larger than the number of correctable bits with ECC and all the redundant rows have been used, the checking process is terminated with  $Hold_{success} = 0$ , which means the current standby voltage is too low to retain data and hence must be increased.

Note that the standby time  $T_{sb}$  should be sufficiently long to ensure the occurrence of the worst static scenario. Fig. 7 shows an example of one SRAM cell in the 45 nm technology we use. The DRV value decreases for less standby time, which means SRAM cells can tolerate more dynamic noise when the duration of the noise is shorter. This similar behavior of larger dynamic noise tolerance has been observed in logic gates [9]. After the standby time exceeds a threshold point, its DRV reaches the largest value, which equals the one from the static dc simulation.

##### B. Calibrating Initial Failure Threshold

Simulation and measured results have shown that the DRV of the canary cell is approximately linear with the VCTRL value [6]. By analyzing the leakage current through the header when  $V_{DD}$  reaches

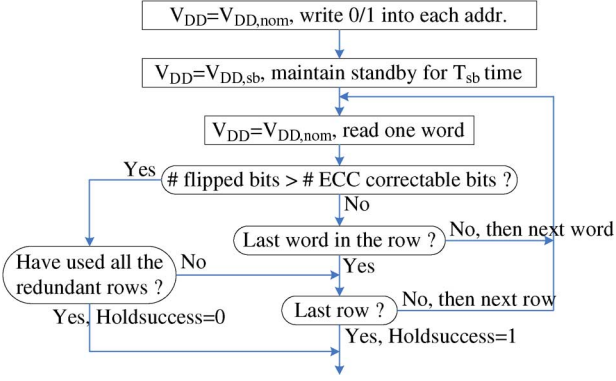


Fig. 6. Flow for hold failure check with BIST.

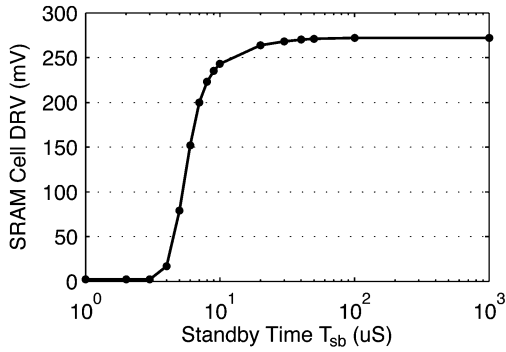


Fig. 7. DRV of an SRAM cell changes with the standby time.

the true DRV of the cell, we derived that the linearity can be approximated with  $1/(1+\eta)$ , where  $\eta$  is the DIBL coefficient of the header [7]. Hence we can generate a series of VCTRL values (e.g., with a resistor ladder) to create a group of canary categories that fail at regular intervals across a wide range. During self-calibration, our BIST first finds the SRAM  $V_{min}$  value, as discussed in Section IV-A. Then the BIST applies that voltage as the supply for canary circuits and measures the failure status of each canary category,  $FT_{max}$ . Suppose we get

$$\begin{aligned}
 FT_{max} &= [f_0 f_1 \cdots f_{k-1} f_k f_{k+1} \cdots f_{n-2} f_{n-1}] \\
 &= [00 \cdots 011 \cdots 11].
 \end{aligned} \quad (1)$$

Here,  $f_i$  means the failure status of the  $i$ th canary; when  $f_i = 1$ , this canary fails. So the  $k$ th canary is the one that fails immediately before the worst SRAM cell. This  $FT_{max}$  value will be recorded (e.g., with a programmable fuse or other non-volatile memory). In normal operation mode, the user first loads  $FT_{max}$ , and then programs an appropriate failure threshold value according to the application needs. We denote  $FT_{max} \gg j$  as the value after right shifting  $FT_{max}$  by  $j$  bits. For aggressive power saving, the failure threshold register should be configured as  $FT_{max} \gg 1$ ; while for more robust  $V_{DD}$  scaling,  $FT_{max} \gg j$  with  $j > 1$  should be used to tradeoff less power saving with higher data reliability. Note that the granularity of the tunability of our canary system is dependent on the quantization of the VCTRL values and the resolution of the voltage regulator.

## V. 45 nm TEST CHIP IMPLEMENTATION AND MEASUREMENT

Our first prototype has been implemented and measured in a bulk 90 nm test chip [6], [7]. To verify the effectiveness of our scheme in scaled technologies, we implemented the canary circuits in a bulk 45 nm test chip. Fig. 8 shows its die photo. On each die, there are two canary




Fig. 8. 45 nm test chip die photo.

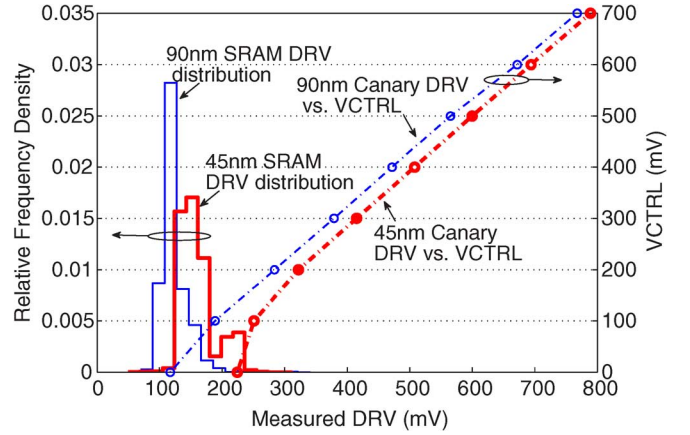


Fig. 9. Measured canary DRV against VCTRL and measured frequency density of the SRAM DRV from both the new 45 nm chip and the previous 90 nm chip. 16 and 8 kb SRAM cells are measured for 45 and 90 nm, respectively.

blocks. Each canary block contains all the canary circuits [see Fig. 2(a)] plus test circuits. The canary bank consists of eight canary sets and each canary set employs three-way redundancy. All the canary cells in the first canary block use the standalone cell structure, and those in the second canary block use the improved structure with the dummy cells as shown in Fig. 3. We also implemented four 4 kb banks of SRAM on the die.

The measured canary cell DRV against VCTRL and the frequency density of the measured SRAM DRV from the new 45 nm test chip are plotted in Fig. 9. For comparison, we also plot the measured DRV results from the previous 90 nm test chip. Both the 90 and 45 nm canary DRV measurements maintain excellent first-order linearity with VCTRL values above 100 mV. The nonlinearity for VCTRLs below 100 mV is due to the rolling off term in the sub-threshold current equation [7]. Note that for the same amount of VCTRL increment (e.g., 100 mV), the 45 nm canary DRV has smaller increase than the 90 nm counterpart because the sensitivity of the canary DRV to VCTRL is inversely proportional to the header's DIBL coefficient, which increases with technology scaling. Fig. 9 shows that the 45 nm SRAM DRV spreads wider than the 90 nm counterpart due to device variability increasing with technology scaling. Although the variance of the SRAM DRV distribution grows in 45 nm, Fig. 9 demonstrates that tuning VCTRL can still provide a sufficiently large range of canary DRV above the tail of the SRAM DRV in 45 nm just as in 90 nm. This ensures that the canary scheme maintains functional in 45 nm.

We further compare the results from the two different canary blocks to examine the effect of dummy cells. Fig. 10 shows the measured results from 85 dies on one wafer. For each die, the VCTRL value of each canary set is generated by an on-die resistor ladder. The canary set with the higher index number connects to a higher VCTRL value. The variation of the canary DRV is computed as the ratio of the sigma ( $\sigma$ ) to

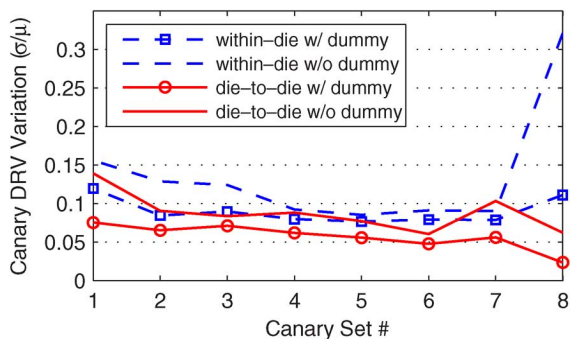


Fig. 10. With dummy cells, both within-die and die-to-die variations of the canary DRV are reduced.

the mean ( $\mu$ ). A smaller ratio value means less variation occurred on the canary. We first compare the within-die variation, i.e., the variation of the three redundant copies of each canary set on each die. The average result from 85 dies is plotted with dashed curves. The block with the dummy cells has less within-die variation, especially for the canary set #8 that is configured to have the largest DRV. We also plot the die-to-die variation with the solid curves. In this case, the canary DRV value of each die is obtained through the majority-3 voting among the redundancies on the same die. The block with dummy cells also has less die-to-die variations. Therefore, the use of dummy cells inside the canary cell can effectively reduce both with-in-die and die-to-die variations of the canary cell.

## VI. CONCLUSION

SRAM standby  $V_{min}$ , i.e., the DRV of the worst SRAM cell, shifts with global PVT variations. The traditional worst-case open-loop approach prevents the potential power savings for non-worst-case dies and scenarios. We have proposed a feedback scheme using canary replicas for aggressive standby  $V_{DD}$  scaling while maintaining sufficient data reliability. In this paper, we propose several enhancements to this scheme. Dummy cells is added in the canary cell to improve the correlation between the canary cell and SRAM cells under systematic variation. A new resetting circuit ensures that the canary cell holds the less-stable state so that it can flip at a higher voltage. We also propose

a BIST to self-calibrate the SRAM standby  $V_{min}$  and the initial failure threshold due to intrinsic mismatch after manufacture. Measurement results from a 45 nm test chip demonstrate that the canary cells can fail at regular intervals across a wide range above the SRAM DRV tail in smaller technology. In addition, measurements confirm that using dummy cells can reduce the variation of the canary cell and thus improve the accuracy of the tracking behavior.

## ACKNOWLEDGMENT

The authors would like to thank Freescale Semiconductor, Inc. for chip fabrication and thank J. Brown for his help in chip characterization.

## REFERENCES

- [1] K. Flautner, N. S. Kim, S. Martin, D. Blaauw, and T. Mudge, "Drowsy caches: Simple techniques for reducing leakage power," in *Proc. Int. Symp. Comput. Arch.*, May 25–29, 2002, pp. 148–157.
- [2] N. S. Kim, K. Flautner, D. Blaauw, and T. Mudge, "Single-V<sub>dd</sub> and single-V<sub>t</sub> super-drowsy techniques for low-leakage high-performance instruction caches," in *Proc. Int. Symp. Low Power Electron. Des. (ISLPED)*, 2004, pp. 54–57.
- [3] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, "SRAM leakage suppression by minimizing standby supply voltage," in *Proc. Int. Symp. Quality Electron. Des. (ISQED)*, 2004, pp. 55–60.
- [4] Y. Wang, H. Ahn, U. Bhattacharya, T. Coan, F. Hamzaoglu, W. Hafez, C.-H. Jan, R. Kolar, S. Kulkarni, J. Lin, Y. Ng, I. Post, L. Wel, Y. Zhang, K. Zhang, and M. Bohr, "A 1.1 GHz 12  $\mu$ A/Mb-leakage SRAM design in 65 nm ultra-low-power CMOS with integrated leakage reduction for mobile applications," in *Proc. ISSCC*, Feb. 2007, pp. 324–606.
- [5] Y. Cao, T. Sato, M. Orshansky, D. Sylvester, and C. Hu, "New paradigm of predictive mosfet and interconnect modeling for early circuit simulation," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, 2000, pp. 201–204.
- [6] J. Wang and B. Calhoun, "Canary replica feedback for near-DRV standby  $V_{DD}$  scaling in a 90 nm SRAM," in *Proc. CICC*, 2007, pp. 29–32.
- [7] J. Wang and B. H. Calhoun, "Techniques to extend canary-based standby  $V_{DD}$  scaling for SRAMs to 45 nm and beyond," *IEEE J. Solid-State Circuits*, vol. 43, no. 11, pp. 2514–2523, Nov. 2008.
- [8] H. Qin, A. Kumar, K. Ramchandran, J. Rabaey, and P. Ishwar, "Error-tolerant SRAM design for ultra-low power standby operation," in *Proc. ISQED*, 2008, pp. 30–34.
- [9] J. Lohstroh, "Static and dynamic noise margins of logic circuits," *IEEE J. Solid-State Circuits*, vol. 14, no. 3, pp. 591–598, Mar. 1979.