A Data Remanence based Approach to Generate 100% Stable Keys from an SRAM Physical Unclonable Function

Muqing Liu, Chen Zhou, Qianying Tang, Keshab K. Parhi and Chris H. Kim
Department of ECE, University of Minnesota, Minneapolis, MN
{liux3300, zhoux825, tangx280, parhi, chriskim}@umn.edu

ABSTRACT

The start-up value of an SRAM cell is unique, random, and unclonable as it is determined by the inherent process mismatch between transistors. These properties make SRAM an attractive circuit for generating encryption keys. The primary challenge for SRAM based key generation, however, is the poor stability when the circuit is subject to random noise, temperature and voltage changes, and device aging. Temporal majority voting (TMV) and bit masking were used in previous works to identify and store the location of unstable or marginally stable SRAM cells. However, TMV requires a long test time and significant hardware resources. In addition, the number of repetitive power-ups required to find the most stable cells is prohibitively high. To overcome the shortcomings of TMV, we propose a novel data remanence based technique to detect SRAM cells with the highest stability for reliable key generation. This approach requires only two remanence tests: writing '1' (or '0') to the entire array and momentarily shutting down the power until a few cells flip. We exploit the fact that the cells that are easily flipped are the most robust cells when written with the opposite data. The proposed method is more effective in finding the most stable cells in a large SRAM array than a TMV scheme with 1,000 power-up tests. Experimental studies show that the 256-bit key generated from a 512 kbit SRAM using the proposed data remanence method is 100% stable under different temperatures, power ramp up times, and device aging.

Categories and Subject Descriptors

·Hardware → Application specific integrated circuits ·Security and privacy → Hardware security implementation.

Keywords

Physical unclonable function; SRAM; stable key generation, data remanence

1. INTRODUCTION

Physical unclonable function (PUF) is a circuit that harnesses inherent manufacturing variation to generate a random and unique key used for secure hardware authentication. The input to a PUF is referred to as "challenge", and is provided by the server. The output of a PUF is called "response" which is sent back to the server for authentication purposes. If the response from the PUF matches the correct response stored on the server, then the user is granted to access to the system.

Two categories of PUFs exist: "strong" PUF and "weak" PUF. Strong PUFs like Arbiter PUF [1] and ring oscillator PUF [2] can generate an exponential number of unique challenge response pairs (CRPs), making them suitable for authentication applications without the use of encryption algorithms. Weak PUFs on the other hand, can only generate a linear number of CRPs and hence are used for key generation. Keys generated by weak PUFs can be used in conjunction with encryption algorithms for authentication applications [3]. The main requirement for keys generated by weak PUFs is that their value should not change with temperature and voltage changes, or with device aging.

SRAM is an attractive option for weak PUFs [4] since it is readily available in digital processors. Compared to dedicated PUFs such as arbiter PUF or ring oscillator PUF, the amount of effort needed to implement an SRAM PUF is negligible. The "challenge" to an SRAM PUF is the memory cell address while the "response" is the uninitialized power-up value of the cell. The layout of a 6T SRAM cell is perfectly symmetric and hence no systematic offset exists. Hence, the power-up state is determined by process variation induced mismatch between the two cross-coupled inverters. The manufacturing variability is random, unclonable and uncontrollable, which gives each chip a unique key. The main design consideration for SRAM PUFs is making sure the key is 100% stable. Given the same challenge, we expect the PUF to generate the same key regardless of the operating condition. This is difficult to achieve since the static mismatch of a SRAM cell may not always be large enough to overpower the random thermal noise under all operating conditions.

Temporal majority voting (TMV) is a popular technique for improving the stability of PUF responses [5,6]. The basic principle is to repetitively test the PUF using the same challenge and take the majority value of the responses as the final output. Increasing the number of repetitive tests allows the tester to find keys that are more stable. The main drawback of TMV is that it usually involves a large number of tests (e.g. 100's or 1000's of power-ups for SRAM PUF), which is prohibitive in terms of test time and test hardware. Furthermore, even with such a large number of trials, the stability criterion cannot be made very stringent, so there's a high possibility that the stable cells found using TMV will become unstable in future evaluations. In [5], a combination of TMV, burn-in hardening and ECC circuits were used to meet the stability requirement. However, these techniques introduce significant hardware overhead. To make matters worse, TMV may have to be performed under extreme voltage and temperature conditions to ensure the responses are truly stable. This is very time consuming and difficult to implement in a high-volume production flow. A bit selection algorithm proposed in [7] utilizes just two test conditions; high-temperature/low-voltage and low-temperature/low-voltage. This is more efficient and less costly for selecting stable bits, however, it involves changing the test temperature which is undesirable. Error Correcting Codes (ECC) can be used to correct the unstable outputs using a software algorithm. However, ECC

may leak secret information and introduce extra design complexity and communication overhead.

The instability of TMV selected cells stems from the marginally stable cells, i.e., cells that appear to be stable during TMV tests but become unstable under extreme environmental conditions. These cells are more stable than an average cell, but less stable than the strongest cells that consistently produce the same response. Finding the strongest cells in a large SRAM array requires a prohibitively large number of repetitive tests and may involve changing the voltage and/or temperature. To overcome the limitations of TMV, we propose a method for selecting the most stable cells in an SRAM array based on just two power-up tests. Compared to TMV, our approach reduces the test time and obtains more accurate information pertaining to the stability of cells. Experiment results from off-the-shelf SRAM chips show that the cells selected by our proposed strategy are 100% stable under extreme test conditions.

2. DATA REMANENCE BASED STABLE KEY SELECTION

2.1 Data remanence based approach

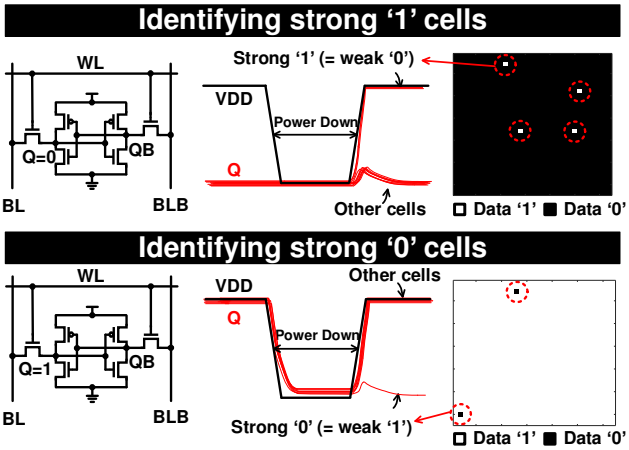

Figure 1. Proposed data remanence based technique to rapidly select the most stable cells in a large SRAM array.

Fig. 1 illustrates the basic principle of the proposed data remanence based stable cell selection method. According to Wikipedia, remanence is defined as "the magnetization left behind in a ferromagnetic material after an external magnetic field is removed" [8]. Similar to this concept, we remove the supply voltage for a short period after initializing the array to all 0's or all 1's.

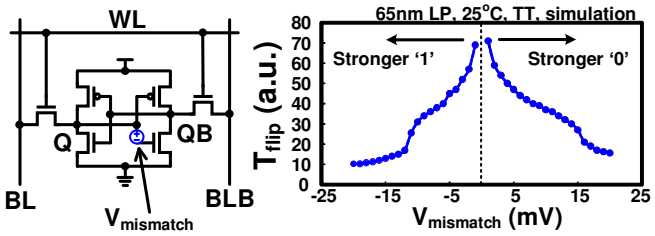

Figure 2. Required power down period to flip SRAM data with different skew.

As shown in Fig. 2, the first few bits to flip after the brief power down period are ones that are strongly biased to the opposite value. For instance, if the entire array is initialized to 0's, the first bits to flip to 1's after the short power down period are the strongest '1' bits in the array. In traditional SRAM PUF power up operation, the response is only related to the inherent transistor mismatch of each SRAM cell. Data written to the cell doesn't affect the power up state because all storage nodes have fully discharged to an unbiased state due to leakage current. In other words, the data remanence is fully decayed. However, if the cell is powered back immediately after a power down, then the storage node data will revert to the previous data because the data remanence is very strong. If the power down time is long enough to make the data remanence comparable to the transistor mismatch, then some cells will revert to the previous data, while other cells will flip to the opposite value. As shown in Fig. 1 (top), if '0' is written to all the cells, node Q will be 0V and node QB will be VDD before the power down. After a short power down period, the majority of the QB nodes will revert back to VDD upon a power up due to the remanence charge on the Q and QB nodes. However, the cells with the strongest bias towards the opposite value will flip to '1' as illustrated in Fig. 1 (top). The transistor mismatch in these cells produces a strong bias which cannot be overpowered by the small data remanence. We utilize this behavior to find the most stable '1's in a large SRAM array. Similarly, by writing '1' to all the cells in the SRAM array and asserting a short power down period, we can find the most stable '0's, which are the first cells to flip when the power is turned back on, as highlighted in Fig 1 (bottom).

Note that a "remanence decay" based side-channel attack method was proposed in [9] where a pulsed power supply was used to recover the secret keys generated by an SRAM based PUF. Our approach employs the same method but for a totally different application: i.e., finding the most stable bits in an SRAM PUF with minimal test time and test hardware overhead.

2.2 Characterization of data remanence effect

To verify the proposed technique in real hardware, we performed data remanence tests on off-the-shelf SRAM chips from Microchip Technology. Each chip contains 512 k memory cells. The first step is to determine the appropriate power down time. If the power down time is too long, then the data stored in the array is completely collapsed and the SRAM will power up to its uninitialized state. On the other hand, if the power down period is too short, then the data

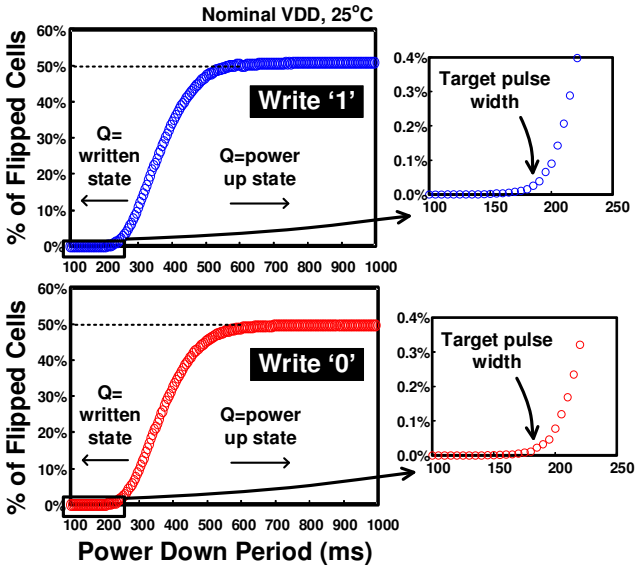

Figure 3. Percentage of flipped cell versus when writing all '1's (upper) and writing all '0's (lower) to the SRAM cells.

will deterministically revert to the previous written state. Therefore, the power down period should be carefully characterized. Fig. 3 shows the percentage of flipped cells when the power down time is swept from 100ms to 1000ms. The SRAM chips we tested were fabricated in an ultra-low leakage technology, requiring a relatively long power down time to observe data remanence effects. We expect a much shorter data remanence time (e.g. microseconds) for SRAMs built in advanced CMOS technologies. The overall data remanence trends will be agnostic to the technology node.

Data was collected from all 512 k cells of each SRAM chip. In both write '1' and write '0' cases, the cells start to flip after a power down period of about 130ms. When the power down period increases to about 600ms, the flip ratio reaches 50% which corresponds to the SRAM power up state. For authentication applications, we are only interested in finding the most stable '1's and '0's in the entire array, and therefore we need to select a power down period that is short enough so that only the most oppositely biased cells flip. This time is usually less than 200ms, which is about 3 times shorter than the power down time required for a standard SRAM PUF evaluation (approximately 600ms in our case). Although the proposed data remanence method requires all SRAM cells to be written to '1' or '0' before the power down, the time needed to write data into the array is negligible compared to the power down time required to clear the data remanence in the SRAM. Moreover, our approach only requires two tests to select the most stable cells in the SRAM array; one test for selecting stable '1' cells and the other for selecting stable '0' cells. TMV may require hundreds or more power ups to find the robustly stable cells, and we must wait at least 600ms between two consecutive power-up tests. In short, compared to TMV, the proposed technique requires not only fewer power-ups (hundreds or thousands → 2) but also shorter power down periods (600ms → 200ms) which significantly reduces the overall test time.

For a better understanding of the proposed technique, Fig. 4 shows data remanence of a small 1kbit sub-array for different power down periods. Fig. 4 (upper) shows the bit map for selecting stable '1's. Data '0', denoted in black, is first written to the whole array, and then the power supply is turned off, letting the data stored in the SRAM to decay. When the power supply is turned on after 130ms, the first cell in the 1kbit array flips. This cell corresponds to the most stable '1' cell in this array. When the power down period increases further, more cells flip, which are the next most stable '1' cells. Depending on the number of stable cells we want to select, the amount of data remanence needs to be tuned accordingly by changing the power off period. Stable '0's can be selected in a similar way, as shown in Fig. 4 (lower).

As seen in Fig. 4, the data remanence based technique allows us to measure the extent to which a cell is stable, by looking at the order of the cell flips. By sweeping the power down time and recording the order of the cell flips, we can sort and list the strength levels of each cells from the strongest '0' cell to the strongest '1' cell. As such, we can obtain complete knowledge of the cell's strength of the whole SRAM array by sweeping the power down time. For example, if we want to sort the cells from strongest '0' to balanced '0', we first write data '1' to the whole array and sweep the power down period from 100ms to 600ms for the SRAM chips used in our experiment. The responses of all cells are recorded and sorted by retention time, as shown in Fig. 5 (upper). The sorting order is shown for 50% of the cells (i.e., 256 k), since the other half will always generate a '1' irrespective of the power down time. When applying the data remanence method to generate stable keys, we only select the strongest cells. The most biased cells can be seen more clearly in the zoomed-in plot. Depending on how many stable bits we want to select, we can vary the power off period. For example, for a 256-bit key, we select roughly 128 stable '0's and 128 stable '1's from 512 kbit cells, which is 0.05% of the total cells available. The power off period should be around 185ms. If we want to select 512 bits, we can increase the power off period to around 195ms to allow more flips. In a realistic scenario, we can select more bits than we need and then pick the number of stable bits requested by our target application. Similar plots of the data '1' case are shown in Fig. 5 (lower). The bit index shows the order from the strongest '0's to balanced cells to strongest '1's, from top to bottom. We can observe from Fig. 5 that by using the proposed data remanence technique and sweeping the power down period, we can sort the cell strength levels in very fine steps. For comparison, we used the conventional TMV method to find stable bits in the same SRAM array. 1,000 power-up tests were performed and the probability of each cell being '1' or '0' were calculated. We found that 40% of the cells are stable '1' through all 1,000 tests and 41%

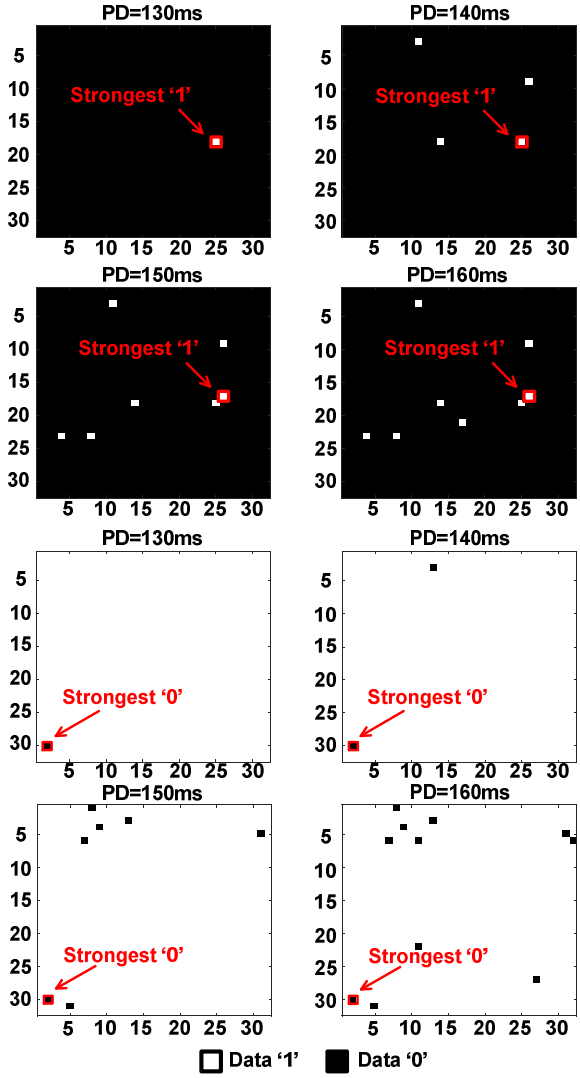

Figure 4. Cell flip maps of randomly selected 1K SRAM cells under different power down periods (PD).

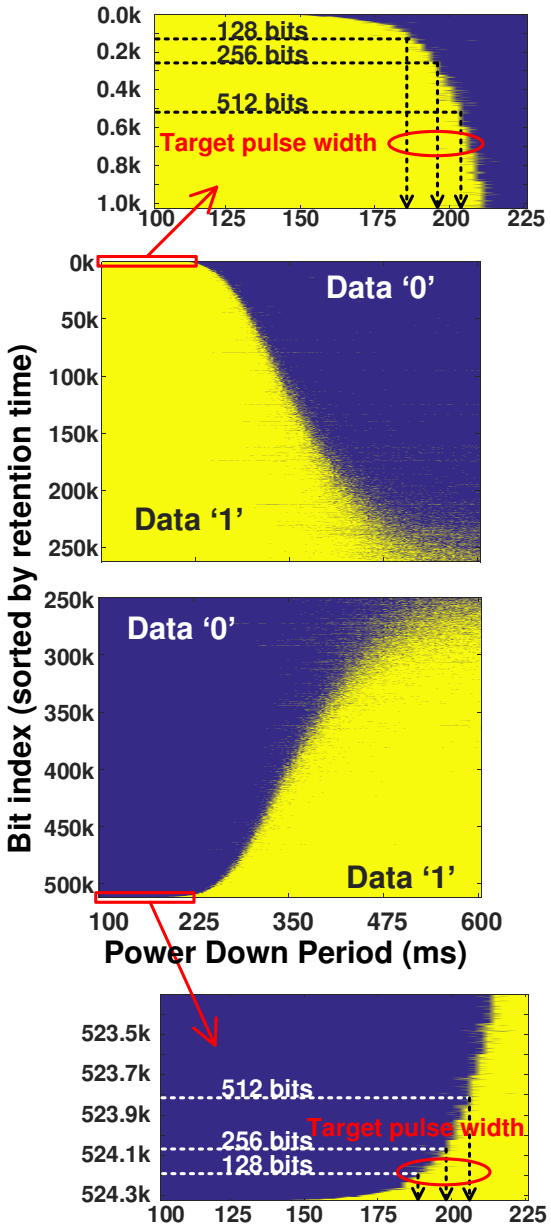

Figure 5. Measured SRAM data remanence for data 1 (upper) and data 0 (lower). Cells with the shortest retention times are highlighted in the zoom-in plots.

of the cells are stable '0' all the time. However, many of the allegedly stable cells will show unstable behavior at different voltage and temperature conditions, or when the SRAM is subject to aging. To determine the 256 most stable bits from a 512 kbit SRAM array, which is only 256/512k = 0.05%, we may need millions of repetitive power up tests for TMV, which is impractical.

Fig. 6 (left) summarizes the test flow for characterizing data remanence while sweeping the power down time. Note that we perform this extensive test on one of the chips to determine the appropriate power down period of all chips. An attractive feature of the data remanence test is that it can be performed at any temperature. The top 0.05% stable cells found from the power down sweep test will remain stable at different temperatures and voltage conditions. For actual SRAM PUF applications, we use the power down period found from the extensive data remanence test and run the test only two times; one for selecting strong '1's and the other for selecting strong '0's. The enrollment test shown in Fig. 6 (middle) stores the location of the most stable bits on-chip. In the key generation phase, we simply power up the SRAM, and key values are retrieved from the stable bit locations.

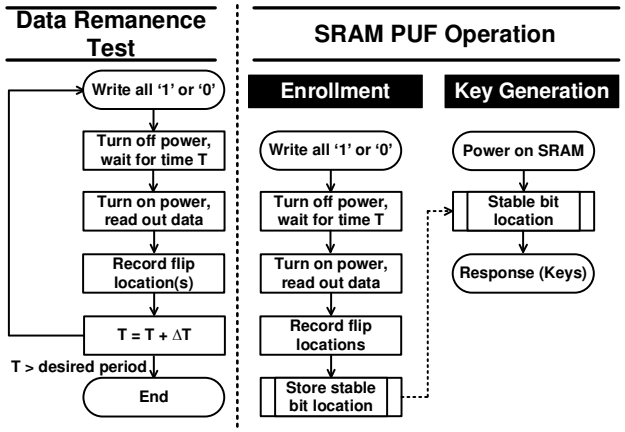

Figure 6. Flow chart for data remanence test, enrollment test, and key generation test. The power down time T for enrollment phase can be determined based on a one-time data remanence test performed at any temperature.

3. SRAM PUF MEASUREMENT RESULTS

This section shows detailed measurement results verifying that the stable cells selected using our proposed technique are indeed stable across different environmental and aging conditions. Fig. 7 shows the measurement set up. The pulsed power supply and other digital signals are provided by a PXI based data acquisition system. GPIB controlled power supplies were used to stress the chip. Chips were measured inside a temperature controlled chamber.

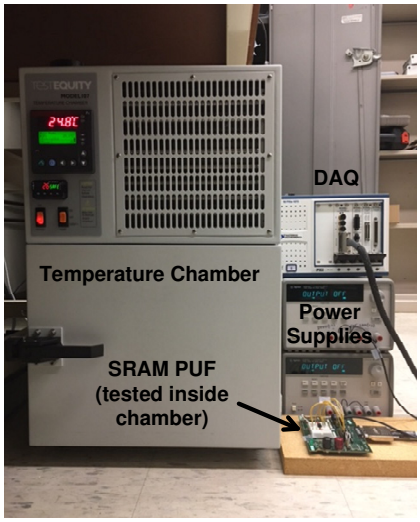

Figure 7. SRAM PUF measurement set up including a temperature control chamber.

3.1 Uniqueness of key

The maximum number of bits for encryption algorithms like AES, is usually 256 bits [3]. So, the target number of bits for our SRAM PUF based key generation is 256 bits. However, we also present

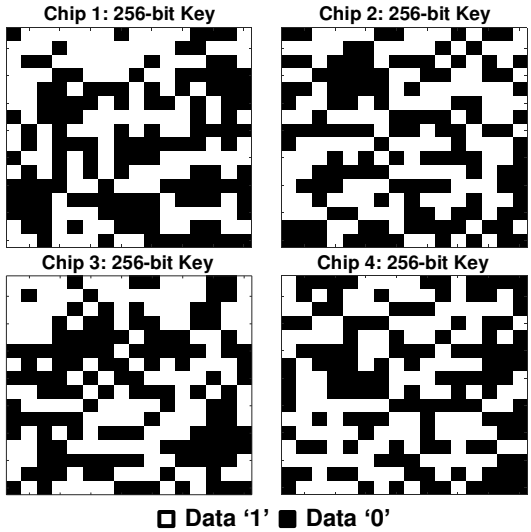

Figure 8. Keys generated from 4 different SRAM chips. The inter-chip Hamming distance between the 4 different keys is 0.4935.

results for generating 512 bit and 1024 bit keys for better understanding. Each SRAM chip we tested has 512 kbits in total, so the goal is to select the 128 most stable '1's and the 128 most stable '0's from 512 kbits, which correspond to the most stable 0.05% of the bits. The power supply off period for selecting 256 stable bits in this test was about 190ms. Alternatively, we can select more than 256 bits and randomly pick 256 cells and record their locations to generate the bit location address. Fig. 8 shows 256 bit keys generated from 4 different SRAM chips showing an average inter-chip Hamming distance of 0.4935, confirming the uniqueness of the keys. Note that the precise location of the stable bits is different in each SRAM chip.

3.2 Effect of power ramp up time and temperature

To verify that the key selected using the proposed technique is stable under different environmental conditions, the voltage ramp up rate and temperature were varied. Note that during the SRAM power up, the state is resolved during the very beginning of the power supply ramp up, so the final power supply level will not affect the stability of the SRAM PUF. Instead, the ramp up rate of the power supply may have an effect on the stability of the responses. To evaluate this effect, the ramp up rate of the supply voltage was changed from 0.78µV/s to 8.33µV/s. Testing was performed at three temperatures; 80°C, 25°C and -10°C. Fig. 9 (left) shows the measured SRAM PUF responses and the average intra-chip Hamming distances under different test conditions using the proposed technique. Power up tests were repeated 10 times under each condition to ensure that the responses are absolutely stable. Since the responses are always stable, there is no need for further processing of the responses using ECC. This reduces the circuit complexity and communication overhead. For comparison, we also select 256 stable bits using the TMV method based on 1,000 repetitive power ups. That is, we only chose the cells that are consistently '0' or consistently '1' throughout the entire 1,000 trials. As mentioned earlier, even with 1,000 repetitive power up tests, we are only able to discriminate the top 81% stable cells which includes marginally stable cells. As a reminder, the proposed data remanence technique can select the top 0.05% stable cells with just two power-up tests. The responses using the 1,000 trial TMV method are shown in Fig. 9 (middle). The unstable bits are highlighted in red. It can be seen from the cell maps that 4 cells are unstable when the temperature or power supply ramp up rate changes, which is not acceptable for ECC-less key generation. Finally, the power up responses from 256 randomly selected SRAM cells are shown in Fig. 9 (right). As expected, many bits are unstable when tested under different conditions. These measurement results confirm that the data remanence technique proposed in this paper can reliably identify the most stable bits in an SRAM array with only two power-up tests. The stable keys can be selected under the nominal voltage and room temperature condition, so it can greatly reduce the test cost and test time. We also selected the 512 most stable bits and 1024 most stables bits and their responses were proven to be 100%

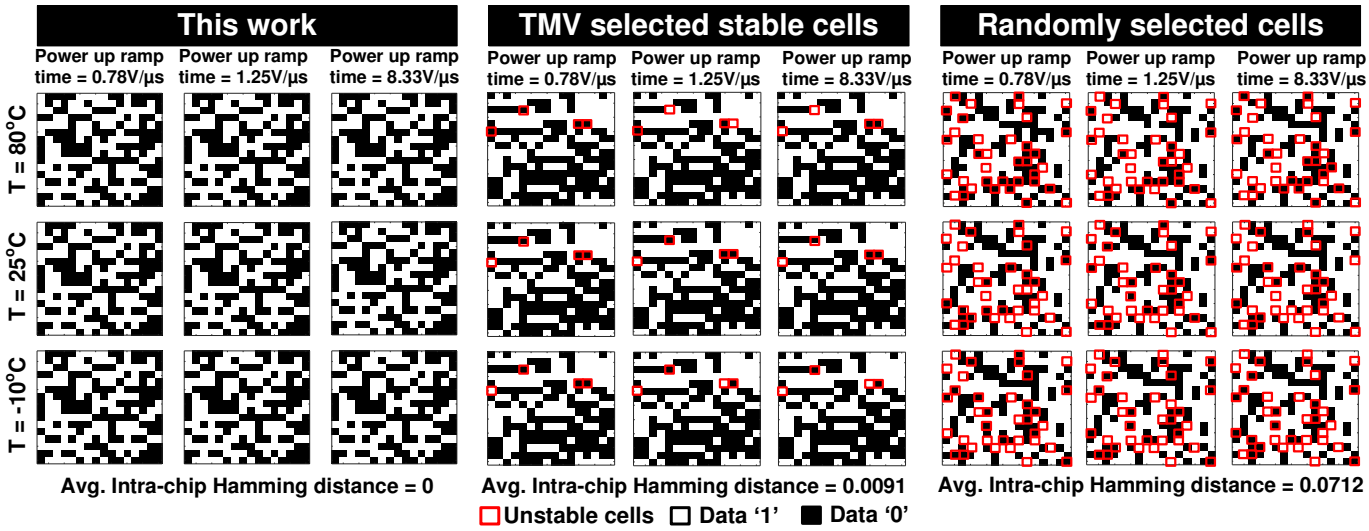

Figure 9. 256 stable bits selected using the proposed data remanence based method, TMV method (1000 power-up trials), and random selection method. Keys generated using the proposed method show no unstable behavior across different power up ramp times and temperatures.

stable under various voltage and temperature conditions, confirming the effectiveness of this technique.

3.3 Effect of device aging

Device aging may cause the PUF response to change over the lifetime of a product, which is undesirable [10]. In particular, bias temperature instability (BTI) is known to be the dominant aging mechanism in SRAM cells due to the low activity factor and DC stress nature [11,12]. BTI manifests as an increase in threshold voltage, and occurs when PMOS or NMOS transistors are biased with a negative or positive gate voltage, respectively [13]. Depending on the data stored in the SRAM cell during stress, BTI can either emphasize or de-emphasize the process variation induced mismatch. Emphasizing the mismatch will harden the responses and make them more stable, while de-emphasizing the mismatch will have the opposite effect [10]. Since our goal is to verify the stability under the worst case condition, we stress the SRAM array with the power-up state which will decrease the mismatch between the two cross-coupled inverters. This de-emphasizes the mismatch and makes the bits more unstable during the actual power up test. The SRAM chips were stressed under a static DC condition (i.e., no switching or toggling) for 72 hours using a 1.5xVDD supply voltage. Before applying the stress voltage, the fresh PUF response is read out for reference. The SRAM PUF responses of the selected 256/512/1024 most stable cells are read out every hour and the intra-chip Hamming distances are calculated against the fresh response (Fig. 10).

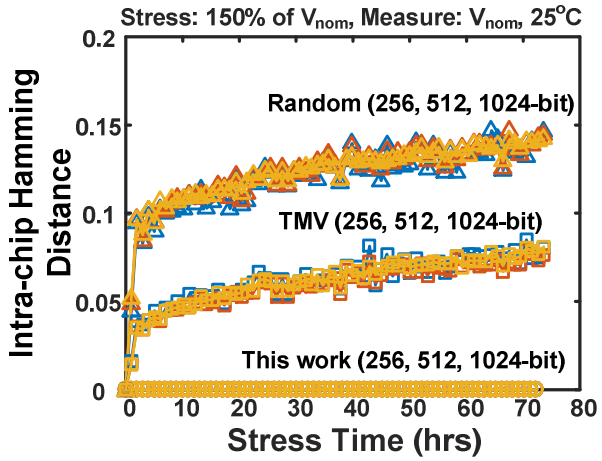

Figure 10. Average intra-chip Hamming distances of different techniques (i.e., proposed, TMV, and random selection) versus stress time for 256, 512, 1024 selected bits.

For comparison, the intra-chip Hamming distances of stable cells selected using the TMV method and the random selection method are also shown in Fig. 10. We can see that the stable cells selected using the proposed technique are 100% stable throughout the entire stress experiment. TMV leads to 8% bit flips at the end of the 72 hour stress period, while the number of bit flips for randomly selected cells is 15%. Experimental results confirm that the cells selected using the proposed technique remain 100% stable for the stress condition used in this work.

4. CONCLUSIONS

In this paper, we proposed a data remanence based technique, to efficiently generate 100% stable keys from an SRAM PUF. By writing '1's or '0's to the entire SRAM array and recording the bit flip locations after a brief power down period, we can identify the strongest '1' and strongest '0' cells in a large SRAM array with just two power-up tests. We have confirmed that the responses selected based on the proposed technique are 100% stable under different voltage and temperature variations, as well as under BTI aging. The proposed technique doesn't require repetitive power-up tests as in conventional TMV methods. Since the responses are 100% stable, there's no need for ECC, which simplifies the authentication hardware.

5. ACKNOWLEDGEMENTS

This research was supported in part by the National Science Foundation under award number CNS-1441639 and the Semiconductor Research Corporation under contract number 2014-TS-2560.

6. REFERENCES

[1] Devadas, S., Suh, E., Paral, S., *et al.*, "Design and implementation of PUF-Based 'unclonable' RFID ICs for anti-counterfeiting and security applications, "*Proceedings of IEEE International Conference on RFID*, May 2008, pp. 58-64.

[2] Maiti, A., Casarona, J., McHale, L., et al., "A large scale characterization of RO-PUF", *IEEE International Symposium on Hardware-Oriented Security and Trust*, 2010, pp. 94-99.

[3] Herder, C., Yu, M., Koushanfar, F., and Devadas, S., "Physical unclonable functions and applications: A tutorial," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1126-1141, 2014.

[4] Holcomb, D., Burleson, W., and Fu, K., "Power-up SRAM state as an identifying fingerprint and source of true random numbers." *IEEE Transactions on Computers,* vol. 58, no. 9, pp. 1198-1210, Sep. 2009.

[5] Mathew, S. K., Satpathy, S. K., Anders, M. A., *et al.*, "A 0.19pJ/b PVT-variant-tolerant hybrid physically unclonable function circuit for 100% stable secure key generation in 22nm CMOS," *IEEE International Solid State Circuits Conference*, Feb. 2014, pp. 278-279.

[6] Zhou, C., Satapathy, S., Lao, Y., et al, "Soft response generation and thresholding strategies for linear and feed-forward MUX PUFs," *International Symposium on Low Power Electronics and Design*, Aug., 2016.

[7] Xiao, K., Rahman, M., Forte, D., *et al.*, "Bit selection algorithm suitable for high-volume production of SRAM-PUF," *IEEE International Symposium on Hardware-Oriented Security and Trust*, 2014, pp. 101-106.

[8] Wikipedia, https://en.wikipedia.org/wiki/Remanence.

[9] Oren, Y., Sadeghi, A.and Wachsmann, C., "On the effectiveness of the remanence decay side-channel to clone memory-based PUFs,"*Cryptographic Hardware and Embedded Systems*, vol. 8086, pp. 107-125, 2013.

[10] Maes, R. and van der Leest, V., "Countering the effects of silicon aging on SRAM PUFs," *IEEE International Symposium on Hardware-Oriented Security and Trust*, 2014, pp. 148-153.

[11] Jain, A. Paul, A., and Kim, C. H., "A 32nm SRAM reliability Macro for recovery free evaluation of NBTI and PBTI," *Iinternational Electron Devices Meeting,* Dec. 2012.

[12] Kim, T., Zhang, W., and Kim, C. H., "An SRAM reliability test macro for fully-automated statistical measurements fo Vmin degradation," *Custom Integrated Circuits Conference*, Sep. 2009.

[13] Schroder, D. K. and Babcock, J. A., "Negative bias temperature instablility: Road to cross in deep submicron silicon semiconductor manufacturing," *Journal of Applied Physics*, vol. 94, pp. 1-18, Jul. 2003.