DRV-Fingerprinting: Using Data Retention Voltage of SRAM Cells for Chip Identification[*]

Daniel E. Holcomb
UC Berkeley

Amir Rahmati
UMass Amherst

Mastooreh Salajegheh
UMass Amherst

Wayne P. Burleson
UMass Amherst

Kevin Fu
UMass Amherst

Abstract

Physical unclonable functions (PUFs) produce outputs that are a function of minute random physical variations. Promoted for low-cost authentication and resistance to counterfeiting, many varieties of PUFs have been used to enhance the security and privacy of RFID tags. To different extents, applications for both identification and authentication require a PUF to produce a consistent output over time. As the sensing of minute variations is a fundamentally noisy process, much effort is spent on error correction of PUF outputs. We propose a new variant of PUF that uses well-understood properties of common memory cells as a fingerprint. Our method of fingerprinting SRAM cells by their data retention voltage improves the success rate of identification by 28% over fingerprints based on power-up state.

1 Introduction

RFID circuits can be identified or authenticated using static identifiers stored in non-volatile memory or through the use of identifying physical characteristics. Physical characteristics have several security advantages over static identifiers, including immutability and resistance to cloning and tampering. The physical characteristics can be viewed as an identifying fingerprint of a given device. More formally, physical fingerprints are a component of a particular type of physical unclonable function (PUF) that is originally described as a physically obfuscated key [4], and more recently as a weak PUF [6].

If used for identification or constructing secret keys, fingerprint observations must be consistent over time. Sensing the microscopic variations that make each device unique while also minimizing the impact of noise is a fundamental

[*] Accepted to appear at the 8th Workshop on RFID Security and Privacy, July 2012

concern in PUFs. Much effort is spent on error correction of somewhat-unreliable fingerprints or PUF outputs. Error correcting codes are expensive in terms of the number of raw bits required to create a reliable key, and more so if the number of correctable errors must be large. Toward this goal, we present a new fingerprinting method that is more reliable across trials than comparable previous approaches.

In this work we propose a new method for chip fingerprinting that uses Data Retention Voltage (DRV) in SRAM as the identifier. The DRV of an SRAM is the minimum voltage at which its cells can retain state. DRV fingerprints are found to be more informative than other approaches for fingerprinting SRAM that have been proposed in research [6, 8] and commercially.[1] The physical characteristics responsible for DRV are imparted randomly during manufacturing and therefore serve as a natural barrier against counterfeiting. The proposed technique has the potential for wide application, as SRAM cells are among the most common building blocks of nearly all digital systems including smart cards and programmable RFID tags.

The contributions of this work are as follows:

- Demonstrating that the DRVs of SRAM cells are consistent fingerprints capable of identifying devices among a population.
- Demonstrating that DRV fingerprints make use of physical variations in a way that is similar to SRAM power-up fingerprints, but that DRV fingerprints have the potential for more accurate identification.

The remainder of this paper is structured as follows: Section 2 introduces data retention voltage. Section 3 explains how the DRVs of SRAM cells are characterized. Section 4 evaluates DRV fingerprinting using experimental data. Sections 5 and 6 review related work and present directions for future work.

2 Data Retention Voltage

A data retention failure is said to occur when an SRAM cell spuriously flips state due to insufficient supply voltage. The data retention voltage (DRV) of an SRAM array signifies the minimum supply voltage at which all SRAM cells can store arbitrary state. DRV is studied in the literature as a limit to supply voltage scaling. Various simulation models [25, 2, 12] and silicon measurements [15] show modern SRAM DRVs to be under 300mV. Most previous literature focuses on cases where the supply voltage of the circuit remains safely above DRV. While remaining above DRV, the supply voltage can be adjusted to reduce leakage power [3], compensate for manufacturing variability [12], or compensate for environmental variations [25].

Each SRAM cell uses the positive feedback of cross-coupled inverters to hold state on two complementary storage nodes. Retention failures occur at low supply voltages because the low voltage weakens the positive feedback of the cross-coupled inverters. Due to asymmetric process variation, at some low

[1] http://www.intrinsic-id.com/

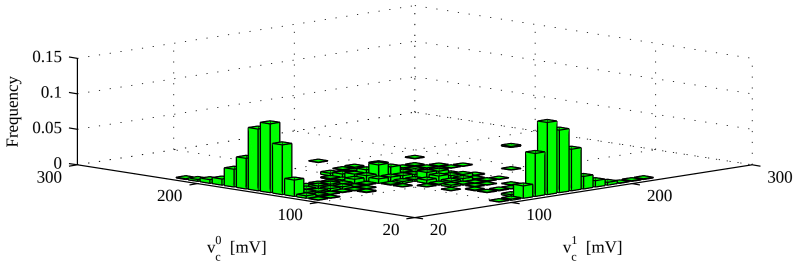

Figure 1: The joint probability distribution function over all cells of the two variables (v_c^0 and v_c^1) comprising a DRV characterization. The distribution is determined experimentally using Algorithm 1, and shows that a large fraction of cells have the minimum possible value of 20mV for either v^0 or v^1, but none have the minimum value (or near-minimum values) for both. A cell with a minimum value for v^0 or v^1 is a cell that retains one written state across all test voltages.

supply voltages a transition from a written state to the opposite state becomes inevitable; observations about the direction of such transitions and the voltages at which they occur are the basis for DRV fingerprints. Any collection of SRAM cells has a distinctive DRV fingerprint because of its unique random process variation.

3 Characterizing the DRV of an SRAM Cell

The DRVs of SRAM cells are characterized by repeatedly lowering the SRAM supply voltage and observing the highest voltage at which each cell fails. If the SRAM supply node also supplies the processing core, then the low voltages used for the characterization will cause the core to reset and lose its state. Our experiments avoid this difficulty by using non-volatile memory to maintain persistency across the low voltages. However, a custom integrated circuit designed for DRV fingerprinting can also avoid this difficulty by using an SRAM supply node that is decoupled from the nominal supply node of the processor. This is often done, for example, in power-gated circuits where unused on-chip functional blocks are turned off entirely while the chip as a whole remains powered.

We characterize the DRV of an SRAM cell c with a pair $\langle v_c^0, v_c^1 \rangle$. Each v_c^w in the pair represents the highest voltage at which cell c will have a retention failure after state w is written to it. In principle, v_c^0 and v_c^1 are real-valued; but in practice, we approximate each using one of $N = (300mV - 20mV)/\Delta$ discrete values as shown in Algorithm 1. With Δ set at 10mV, the $N = 28$ possible values for v_c^0 and v_c^1 are $\{20mV, 30mV, \dots, 290mV\}$. The frequency of observing different DRV pairs is shown in the joint probability distribution function of variables v_c^0 and v_c^1 in Fig. 1.

Algorithm 1 Characterize the DRV fingerprint of a set of SRAM cells.

```
Prerequisite: $C$ – a set of SRAM cells
Ensure: $v_c^0, v_c^1$ – the DRV characterizations of each SRAM cell $c \in C$.
Let $V_{nom}$ be the nominal supply voltage ($V_{dd}$) for the chip
Let $s_c$ refer to the logical state of SRAM cell $c \in C$.
Let $s'_c$ refer to the logical state of NVM cell that corresponds to SRAM cell $c$.
for $w = 0, 1$ do
  for $c \in C$ do
    $s_c \leftarrow w$   {write $w$ into SRAM cell}
    $s'_c \leftarrow w$   {write $w$ into NVM cell}
    $v_c^w \leftarrow 0$   {value used if no retention failure observed}
  end for
  $v_{test} \leftarrow 300mV$   {initialize test voltage}
  while $v_{test} > 20mV$ do
    lower chip voltage from $V_{nom}$ to $v_{test}$
    wait for $t_{wait}$ seconds
    raise chip voltage from $v_{test}$ to $V_{nom}$
    for $c \in C$ do
      if $(s_c = \neg w) \wedge (s'_c = w)$ then
        SRAM cell $c$ had a retention failure from state $w$ at voltage $v_{test}$, but
        previously had no failure at voltage $v_{test}+\Delta$. Therefore $v_{test}$ approximates
        the largest voltage that induces a retention failure after writing $w$.
        $v_c^w \leftarrow v_{test}$
      end if
      $s'_c \leftarrow s_c$   {write SRAM to NVM}
    end for
    $v_{test} \leftarrow v_{test} - \Delta$   {try a lower voltage next}
  end while
end for
```

3.1 Experimental Setup

We examine the DRV of SRAM cells using Algorithm 1 implemented as follows: A microcontroller runs a program that sets all available memory bits to either 1 or 0. The supply voltage is then decreased to a value between 300mV and 20mV ($\Delta = 10mV$) for 5 seconds. When supply voltage is restored to 3V, the program stores the content of SRAM to the flash memory. Note that we conservatively use $t_{wait} = 5s$ to avoid missing marginal failures. Simulations by Nourivand et al. [12] using a procedure similar to Algorithm 1 show that waiting for $t_{wait} = 2ms$ at a reduced supply voltage is sufficient to observe retention failures. An Agilent U2541A-series data acquisition (DAQ) unit controls the supply voltage and the timing of when voltage is raised and lowered. Thermal tests are conducted inside of a Sun Electronics EC12 Environmental Chamber [22], and an OSXL450 infrared non-contact thermometer [13] with $\pm 2°C$ accuracy is used to verify the temperature. All experiments use instances of Texas Instruments MSP430 F2131 microcontrollers with 256 bytes of SRAM, of which 240 bytes are available for DRV fingerprinting. The DRV of each cell is characterized 20 times. The total

Outcome			**Freq.**
$\langle v_c^0$	,	$v_c^1 \rangle$	
$\langle 130mV$	,	$100mV \rangle$	0.0096
$\langle 120mV$	,	$100mV \rangle$	0.0076
$\langle 130mV$	,	$110mV \rangle$	0.0070
$\langle 120mV$	,	$110mV \rangle$	0.0070

(a) Most common weak DRVs

Outcome			**Freq.**
$\langle v_c^0$	,	$v_c^1 \rangle$	
$\langle 20mV$	,	$130mV \rangle$	0.0893
$\langle 20mV$	,	$120mV \rangle$	0.0719
$\langle 130mV$	,	$20mV \rangle$	0.0685
$\langle 20mV$	,	$140mV \rangle$	0.0651

(b) Most common strong DRVs

Table 1: The 4 most commonly observed weak and strong DRV characterizations, and the probability of observing each in a randomly selected trial.

runtime to characterize all 240 bytes of SRAM on a chip once using Algorithm 1 is given by t_{proc} in Eq. 1, and is 140 seconds for the conservative case of $\Delta = 10mV$ and $t_{wait} = 5s$.

$$t_{proc} = t_{wait} \times \frac{300mV - 20mV}{\Delta} \tag{1}$$

3.2 Information Content of SRAM Cell DRV

The DRV of each cell has N^2 possible outcomes representing all combinations of N outcomes for v_c^0 and the N outcomes for v_c^1 (in our case $N = 28$). The DRV of each cell is then a random variable X with N^2 outcomes denoted x_0 through x_{N^2-1}. The total entropy $H(X)$ is the expected information value of the DRV of an unknown cell. Entropy depends (per Eq. 2) on the probabilities of each DRV outcome, denoted $p(x_i)$. In the ideal case where all N^2 outcomes are equally likely (e.g. $p(x_i) = 1/N^2$ for all x_i), each DRV would have almost 10 bits of entropy. Applying Eq. 2 to the decidedly non-uniform outcome probabilities of Fig. 1 shows the actual entropy of a DRV to be 5.12 bits. The most frequently observed DRV outcomes are given in Table. 1.

Eq. 1 shows that runtime is inversely proportional to Δ, so we consider the information loss from making Δ larger than 10mV. Fig. 2 shows the ideal and actual entropy of DRV characterizations when different values of Δ are used. In the extreme case where $\Delta = 140mV$, variables v_c^0 and v_c^1 are each restricted to the values $\{20mV, 160mV\}$, so the ideal entropy of the DRV is equivalent to 2 flips of a fair coin. The values of Δ used in Fig. 2 are chosen on account of being unambiguously recreatable from the $\Delta = 10mV$ data.

$$H(X) = -\sum_{i=1} p(x_i) \log p(x_i) \tag{2}$$

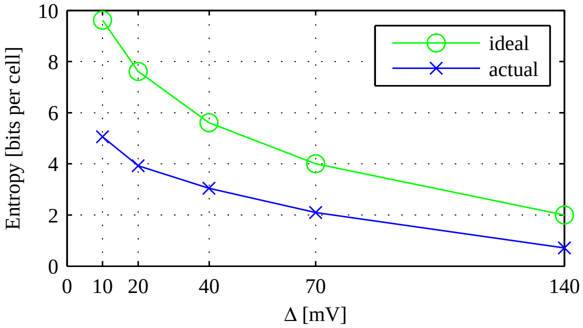

Figure 2: Sweeping Δ from 10mV to 140mV shows that a loss of measurement precision reduces entropy of each cell's DRV characterization.

3.3 Observations about Strong and Weak Cells

We abstract the N^2 possible DRV characterizations (Fig. 1) into three classes[2] that are sufficient to demonstrate general observations about all DRVs:

- A *strongly 0* DRV characterization is a pair $\langle v_c^0, v_c^1\rangle$ such that $v_c^0 = 20mV$ and $v_c^1 > 20mV$. A strongly 0 DRV indicates that no retention failure occurs at any voltage v_{test} after state 0 is written.
- A *strongly 1* DRV characterization is a pair $\langle v_c^0, v_c^1\rangle$ such that $v_c^0 > 20mV$ and $v_c^1 = 20mV$. A strongly 1 DRV indicates that no retention failure occurs at any voltage v_{test} after state 1 is written.
- A *weak* DRV characterization is a pair $\langle v_c^0, v_c^1\rangle$ such that $v_c^0 > 20mV$ and $v_c^1 > 20mV$. A weak DRV indicates that a failure is observed at some voltage v_{test} after each state is written.

The variation-dependent behavior of an SRAM cell occurs somewhere between 20mV and 300mV for each cell; above 300mV all cells can reliably hold either the 0 or the 1 state, and below 20mV no cells can do so. When a cell produces a strongly 0 or strongly 1 characterization, it means (per Algorithm 1) that for one written state the supply voltage is lowered all the way through the sensitive region down to 20mV and then raised back up without causing a failure. A strongly 0 or strongly 1 characterization therefore indicates a strong preference for one state over the other at all supply voltages. A weak characterization is when each written state flips at some voltage within the sensitive region, and neither state can be retained down to 20mV.

Both strong and weak DRV characterizations are largely repeatable across trials. Fig. 3 shows the distribution of DRVs produced by randomly selected cells for which the first DRV produced is one of the 4 most commonly observed weak DRVs from Table 1a; each plot shows the conditional probability distribution of

[2] Note that no observation of $\langle v_c^0, v_v^1\rangle = \langle 20mV, 20mV\rangle$ is ever made, so we do not include this outcome in any of the three cases.

a subsequent DRV characterization. Occasionally the same cells that produce a weak DRV produce a strong DRV in subsequent trials. Fig. 4 shows the same analysis for the 4 most commonly observed strong DRVs; none of the cells subsequently produces the opposite strong characterization.

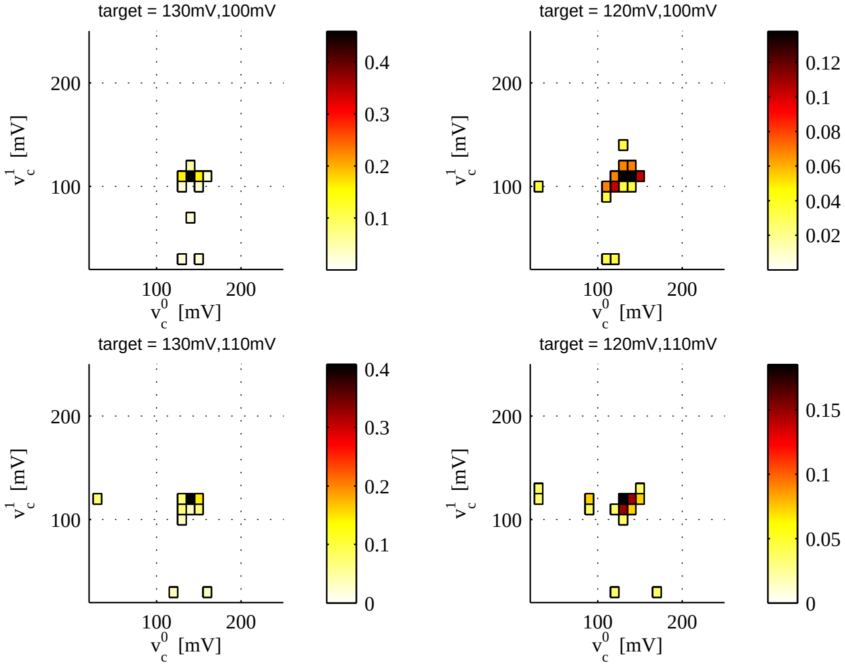

Figure 3: For each of the 4 most frequently observed weak DRVs (see Table 1a), the DRV in a second trial from a cell that produced the frequently observed DRV in a first trial.

3.4 Relation to Power-up State

It is known that SRAM cells consistently power-up to the same state [6, 8] in a majority of trials. Cells with highly reliable power-up states tend to be the same cells with strong DRV characterizations. Fig. 5 shows the mean power-up state over 28 trials for cells that produced a strongly 0 or strongly 1 DRV characterization. Among cells with strongly 0 DRV, 98.6% power-up to the 0 state in all 28 power-up trials (Fig. 5a). Similarly, 95.1% of cells characterized as strongly 1 consistently power-up to the 1 state (Fig. 5a). Although a strong DRV fingerprint is correlated to power-up tendency, the DRV provides a more informative identifier than does power-up by providing information about the maximum voltage at which the unfavored state cannot be reliably stored.

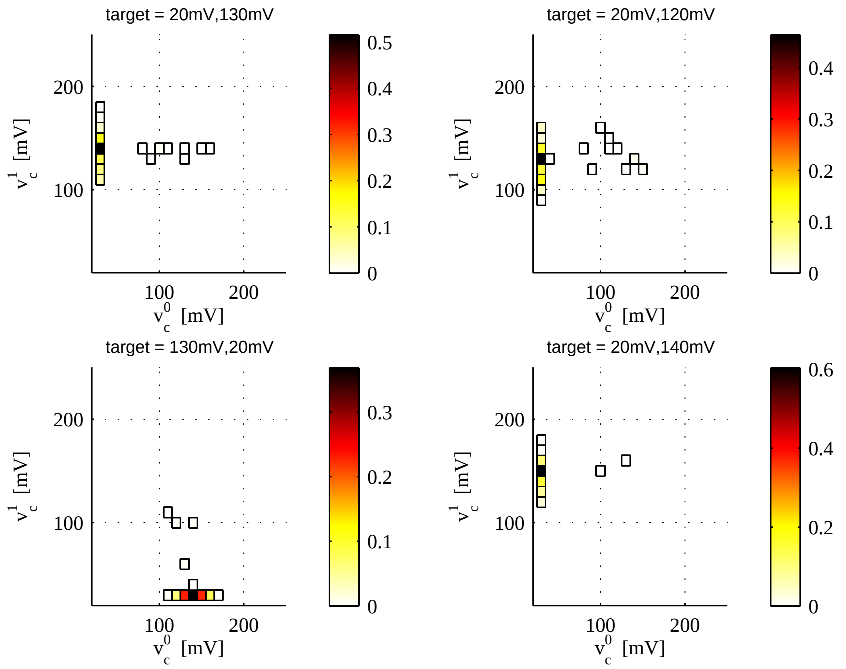

Figure 4: For each of the 4 most frequently observed strong DRVs (see Table. 1b), the DRV in a second trial from a cell that produced the frequently observed DRV in a first trial.

4 Fingerprint Matching

A DRV fingerprint is obtained from a single characterization of a set of adjacent cells within an SRAM. A k-bit fingerprint F_i comprises cell characterizations $\langle v_i^0, v_i^1 \rangle, \langle v_{i+1}^0, v_{i+1}^1 \rangle, \ldots, \langle v_{i+k-1}^0, v_{i+k-1}^1 \rangle$. The difference between fingerprints is the sum of the differences between their corresponding single-cell characterizations. Recalling that each DRV is a point $\langle v_c^0, v_c^1 \rangle$ in 2-dimensional space, we define the distance between two DRVs according to the square of their distance along each dimension (Eq. 3). For comparison, a second metric used is the Hamming distance between power-up trials; this is shown by Eq. 4, where p_i is the state of the i^{th} bit of SRAM after a power-up.

$$d1(F_i, F_j) = \sum_{n=0}^{k-1} \left(v_{i+n}^0 - v_{j+n}^0\right)^2 + \left(v_{i+n}^1 - v_{j+n}^1\right)^2 \tag{3}$$

$$hd(F_i, F_j) = \sum_{n=0}^{k-1} p_{i+n} \oplus p_{j+n} \tag{4}$$

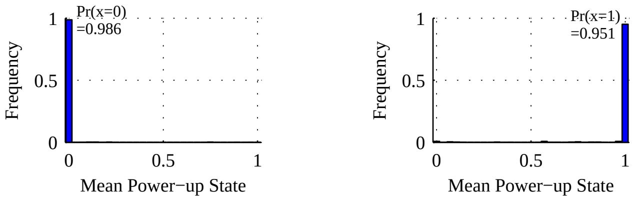

(a) Strongly 0 DRV (b) Strongly 1 DRV

Figure 5: The plot at left shows that 98.6% of SRAM cells that produce a strongly 0 DRV reliably power-up to state 0, as observed by a mean power-up state of 0. The plot at right shows that 95.1% of cells with strongly 1 DRVs reliably power-up to state 1. The DRV is from a single trial of the cell, and the mean power-up state is measured over 28 power-up trials.

	Strongly 0	**Weak**	**Strongly 1**
Strongly 0	35.80%	3.10%	0.00%
Weak	-	24.98%	2.48%
Strongly 1	-	-	33.64%

Table 2: Probability of different pairwise outcomes when 2 DRV fingerprints are taken from a randomly chosen cell. Over the 5000 samples collected, no cell ever has a DRV that is strongly 1 in one trial and strongly 0 in another, but 5.6% of outcomes have one strong and one weak DRV.

4.1 Identification at Nominal Temperature

At the nominal operating temperature of 29°C, three experiments compare DRV fingerprints with power-up fingerprints. These experiments are explained in the following subsections; the first shows the histograms of distances between fingerprints, and the second and third evaluate the accuracy of distance-based matching.

4.1.1 Histogram of Distances Between Fingerprints

A first experiment shows that DRV fingerprints are repeatable and unique, as is necessary for successfully identifying chips within a population. Within-class pairings are of multiple fingerprints generated by the same set of cells on the same device. Between-class pairings are from different sets of cells on the same device, or from any sets of cells on different devices. The similarity of any two fingerprints is quantified by a distance, and this distance is the basis for determining the correct identity of a fingerprint. If within-class fingerprint pairings consistently have smaller distances than between-class pairings, then it is possible to determine identity by choosing an appropriate threshold that separates the two classes. The histograms of within-class and between-class distances for DRV and power-up fingerprints are shown in Fig. 6. These histograms represent all data collected from the MSP430F2131 microcontrollers at room temperature. The distances on the x-axes are not directly comparable across metrics; of importance is only whether the two classes are clearly separable within each plot.

4.1.2 Accuracy of Top Match

The next experiment performed at nominal temperature evaluates how reliably a single within-class DRV fingerprint can be identified among a population. This experiment matches a single 16-bit target fingerprint against a population containing another fingerprint from the same cells and one fingerprint from each of the 239 remaining locations across 2 chips. A positive result occurs if the closest match among the 240 possibilities is from the same SRAM cells as the target. The results of the top match experiment are shown in Table 3; the column labelled "co-top" shows the percentage of trials where there are multiple top matches and one of them correctly matches the target. Multiple top matches are relatively common in Hamming distance matching due to the small number of possible distances between fingerprints. Compared to power-up fingerprints, matching based on DRV fingerprints is 28% more likely to have the correct match be closer to the target (i.e. separated by a smaller distance) than all incorrect matches.

	top	**co-top**	**misidentified**
DRV (d1)	99.7%	-	0.3%
Power-up	71.7%	24.7%	3.6%

Table 3: Over 300 trials with a population of 240 16-bit fingerprints, DRV identification returns the fingerprint that correctly matches the target more reliably than power-up state identification. Matching based on power-up state more frequently returns a misidentified fingerprint, or returns multiple fingerprints among which one is the correct match (denoted "co-top").

4.1.3 Precision and Recall

The top match experiment is generalized to the case of identifying multiple correct matches among a larger population, and again shows DRV fingerprints to outperform power-up fingerprints. In this experiment, our goal is to find all correct matches in the population, without also finding too many incorrect matches. In doing so, the distance that is considered to be the threshold between a correct and incorrect match can be adjusted. If the threshold is too low then correct matches may not be identified, but if the threshold is too high then false positives will occur. *Recall* refers to the fraction of within-class pairings under the threshold, and *precision* refers to the fraction of pairings under the threshold that are within-class. Increasing the threshold will sacrifice precision for recall, and decreasing the threshold will sacrifice recall for precision. An ideal result is for both precision and recall to be 1; this result occurs if all correct matches are identified as within-class (perfect recall) with no incorrect ones identified as within-class (perfect precision).

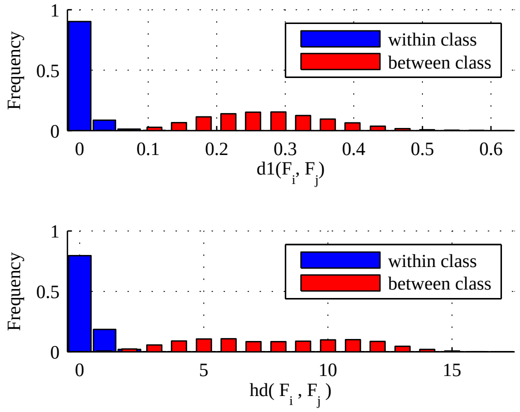

Figure 6: Within-class and between-class distances of 16-bit fingerprints. The upper plot uses DRV fingerprints with distance metric $d1$ from Eq. 3. The lower plot uses power-up fingerprints with Hamming distance as a metric.

The precision and recall plots of Fig. 7 are obtained by iterating the following procedure. One 16-bit segment of SRAM is chosen for identification. One fingerprint trial from this segment is chosen at random as the target, and it is matched against a population of 1019 fingerprints comprising 19 from the same SRAM segment (within-class pairings) and 1000 non-matching fingerprints (between-class pairings). The non-matching fingerprints are randomly selected among 20 trials from 239 other segments of SRAM[3]. The matching threshold is swept to find achievable precision-versus-recall tradeoffs, and each achievable tradeoff is a point in Fig. 7. The large number of tradeoff points in the plot is collected from multiple iterations of this procedure. The general trend is that DRV fingerprints produce better recall for a given precision, or better precision for a given recall compared to power-up fingerprints.

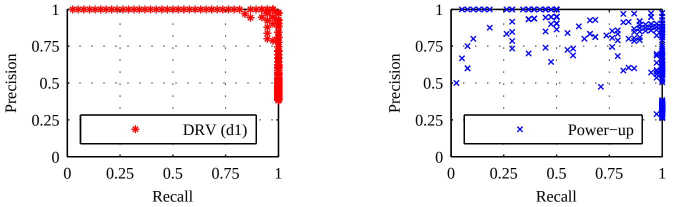

Figure 7: Tradeoff points of precision and recall for trials of DRV fingerprints are generally closer to the ideal result of perfect precision and recall.

[3] The 239 eligible 16-bit segments are the 119 remaining on the target's own chip, and all 120 such locations on the other device.

4.2 Impact of Temperature Variations

Given that DRV fingerprints would likely be used in real-world scenarios without precisely-controlled temperatures, a final experiment explores the impact of temperature on DRV fingerprints. This experiment is similar to the experiment of subsection 4.1.1, but the pairs of fingerprint observations used to generate the within-class distances are now made at different temperatures. The results are shown in Fig. 8. The increase of within-class distances across temperature implies a diminished reliability. To compensate for this, larger fingerprints (comprising more bits) may be needed for identification, and more robust error correcting codes may be needed in key-generation applications. If the increased within-class distances are due to a uniform shift in the DRVs of all cells, then a promising direction for future work would be to design a matching scheme that is insensitive to this type of uniform shift.

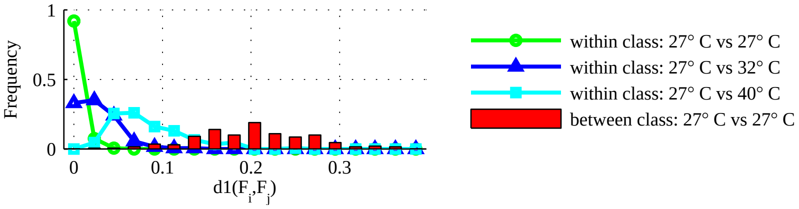

Figure 8: The line plots show within-class distances when one fingerprint observation is made at $27°C$ and the second at $27°C$, $32°C$, or $40°C$; within-class distances increase with temperature, implying a diminished reliability. The bar plot shows between-class distances of 16-bit fingerprints taken at $27°C$. Because there does not exist a distance threshold that can separate the two classes when temperature is varied, it may be necessary to use larger fingerprints for reliable identification.

5 Related Works

A wide variety of PUFs and fingerprints based on custom or pre-existing integrated circuit components have been developed. The identifying features used by custom designs include MOSFET drain-current [10], timing race conditions [5], and the digital state taken by cross-coupled logic after a reset [20]. IC identification based on pre-existing circuitry is demonstrated using SRAM power-up state [8, 6], and physical variations of flash memory [14]. Lee et al. [9] derive a secret key unique to each IC using the statistical delay variations of wires and transistors across ICs. Bhargava et al. explore circuit-level techniques for increasing the reliability of SRAM PUFs [1]. An experimental evaluation of low-temperature data remanence on a variety of SRAMs is provided by Skorobogatov [19], and SRAM remanence in RFID has been studied by Saxena and Voris as a limitation to SRAM-based true random number generation [18].

Previous works [23, 17] have used error correction to construct secret keys from noisy PUF sources; however, this is expensive in terms of gates and other resources. To give an idea of the cost of error correction, BCH codes previously used with PUFs include one to correct 21 errors among 127 raw bits in creating a 64-bit key [21], and to correct 102 errors among 1023 raw bits in creating a 278-bit key [6]. The work of Guajardo et al. [6] uses a derivative of power-up SRAM state as a secret key; however, it requires an error correction code and imposes SRAM space overhead. Maes et al. [11] introduce an SRAM helper data algorithm to mask unreliable bits using low-overhead post-processing algorithms. Recently, Yu et al. [26] proposed a method of error correction for PUFs using a new syndrome coding scheme to minimize the information leaked by the error correction codes, and Hiller et al. extend this approach for SRAM PUFs [7]. Van Herrewege et al. [24] have designed a new lightweight authentication scheme using PUFs that does not require the reader to store a large number of PUF challenge and response pairs.

Given the low cost of the several bytes of SRAM that are used for DRV fingerprinting, a relatively significant practical cost may be associated with the generation of the test voltages for characterizing the DRVs. Emerging devices such as computational RFIDs [16] can use software routines to extract DRVs, but as contactless devices they must generate all test voltages on-chip. On-chip dynamic control of SRAM supply voltage is assumed in the low-power literature at least since work on drowsy caches [3]. Supply voltage tuning has also been applied with canary cells to detect potential SRAM failures, and as a post-silicon technique to compensate for process variation and increase manufacturing yields [12].

6 Conclusions and Future Works

This work has demonstrated that SRAM DRV fingerprints are static identifiers of a device, and it has presented a simple characterization procedure and matching algorithms to use them as such. DRV fingerprints are similar to previously demonstrated power-up fingerprints, but they provide a more informative non-binary identifier of each cell. As a result of this, DRV fingerprints are identified up to 28% more reliably than are power-up fingerprints.

The practical limits of DRV fingerprint performance and reliability should be explored further. Within the constraints of acceptable precision, the runtime of the characterization procedure can be reduced by increasing the voltage step size Δ and reducing the time t_{wait} spent at each voltage (Eq. 1). An expanded evaluation could investigate the reliability of DRV fingerprints across a larger variety of devices and a range of environmental conditions. A high reliability could make DRV fingerprints suitable as a basis for key-generation with lightweight error correcting codes.

Acknowledgments

This research is supported by NSF grants CNS-0964641, CNS-0923313, CNS-0845874, and SRC task 1836.074. Additionally, this research was supported in part by the Gigascale Systems Research Center, one of six research centers funded under the Focus Center Research Program (FCRP), a Semiconductor Research Corporation entity. Any opinions, findings, and conclusions or recommendations expressed in these materials are those of the authors and do not necessarily reflect the views of the sponsors.

References

[1] Bhargava, M., Cakir, C., and Mai, K. Reliability enhancement of bi-stable PUFs in 65nm bulk CMOS. *International Symposium on Hardware-Oriented Security and Trust* (2012).

[2] Cabe, A. C., Qi, Z., and Stan, M. R. Stacking SRAM banks for ultra low power standby mode operation. In *Design Automation Conference* (June 2010).

[3] Flautner, K., Kim, N., and Martin, S. Drowsy caches: simple techniques for reducing leakage power. *International Symposium on Computer Architecture* (2002).

[4] Gassend, B. Physical Random Functions. Master's thesis, MIT, USA, 2003.

[5] Gassend, B., Clarke, D., and Van Dijk, M. Silicon physical random functions. In *Proceedings of the IEEE Computer and Communications Society* (2002).

[6] Guajardo, J., Kumar, S., Schrijen, G., and Tuyls, P. FPGA intrinsic PUFs and their use for IP protection. *Cryptographic Hardware and Embedded Systems* (2007).

[7] Hiller, M., Merli, D., Stumpf, F., and Sigl, G. Complementary IBS: Application specific error correction for PUFs. *International Symposium on Hardware-Oriented Security and Trust* (2012).

[8] Holcomb, D. E., Burleson, W. P., and Fu, K. Power-up SRAM state as an identifying fingerprint and source of true random numbers. *IEEE Transactions on Computers* (2009).

[9] Lee, J., Lim, D., Gassend, B., Suh, G., van Dijk, M., and Devadas, S. A technique to build a secret key in integrated circuits for identification and authentication applications. In *2004 Symposium on VLSI Circuits. Digest of Technical Papers.* (June 2004), pp. 176 – 179.

[10] Lofstrom, K., Daasch, W., and Taylor, D. IC identification circuit using device mismatch. In *IEEE International Solid-State Circuits Conference. Digest of Technical Papers.* (2000), pp. 372 –373.

[11] Maes, R., Tuyls, P., and Verbauwhede, I. Low-Overhead Implementation of a Soft Decision Helper Data Algorithm for SRAM PUFs. *Cryptographic Hardware and Embedded Security* (2009).

[12] Nourivand, A., Al-Khalili, A. J., and Savaria, Y. Postsilicon Tuning of Standby Supply Voltage in SRAMs to Reduce Yield Losses Due to Parametric Data-Retention Failures. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 1 (2011), 29–41.

[13] Omega Engineering, I. *OSXL450 Infrared Non-Contact Thermometer Manual.*

[14] Prabhu, P., Akel, A., Grupp, L., Yu, W., Suh, G., Kan, E., and Swanson, S. Extracting Device Fingerprints from Flash Memory by Exploiting Physical Variations. *Proceedings of the 4th International Conference on Trust and Trustworthy Computing* (2011).

[15] Qin, H., Cao, Y., Markovic, D., Vladimirescu, A., and Rabaey, J. SRAM leakage suppression by minimizing standby supply voltage. In *5th International Symposium on Quality Electronic Design.* (2004), pp. 55–60.

[16] Ransford, B., Clark, S., Salajegheh, M., and Fu, K. Getting things done on computational RFIDs with energy-aware checkpointing and voltage-aware scheduling. In *USENIX Workshop on Power Aware Computing and Systems (HotPower)* (December 2008).

[17] Sadeghi, A.-R., Visconti, I., and Wachsmann, C. *Enhancing RFID Security and Privacy by Physically Unclonable Functions.* Information Security and Cryptography. Springer, Sep 2010, pp. 281–307.

[18] Saxena, N., and Voris, J. We can remember it for you wholesale: Implications of data remanence on the use of RAM for true random number generation on RFID tags. *Proceedings of the Conference on RFID Security* (2009).

[19] Skorobogatov, S. Low temperature data remanence in static RAM. Tech. Rep. UCAM-CL-TR-536, University of Cambridge Computer Laboratory, 2002.

[20] Su, Y., Holleman, J., and Otis, B. A digital 1.6 pj/bit chip identification circuit using process variations. *IEEE Journal of Solid-State Circuits 43*, 1 (Jan. 2008), 69 –77.

[21] Suh, G., O'Donnell, C., and Devadas, S. AEGIS: a single-chip secure processor. *IEEE Design & Test of Computers 24*, 6 (Nov.-Dec. 2007), 570 –580.

[22] Sun Electronic Systems, I. *Model EC1X Environmental Chamber User and Repair Manual*, 2011.

[23] Tuyls, P., and Batina, L. RFID-tags for anti-counterfeiting. In *Topics in Cryptology - CT-RSA 2006, volume 3860 of LNCS* (2006), Springer Verlag, pp. 115–131.

[24] van Herrewege, A., Katzenbeisser, S., Maes, R., Peeters, R., Sadeghi, A.-R., Verbauwhede, I., , and Wachsmann, C. Reverse fuzzy extractors: Enabling lightweight mutual authentication for PUF-enabled RFIDs. In *Financial Cryptography (FC) 2012* (Feb 2012), LNCS, Springer.

[25] Wang, J., and Calhoun, B. H. Techniques to Extend Canary-Based Standby VDD Scaling for SRAMs to 45 nm and Beyond. *IEEE Journal of Solid-State Circuits 43*, 11 (2008), 2514–2523.

[26] Yu, M.-D., and Devadas, S. Secure and Robust Error Correction for Physical Unclonable Functions. *IEEE Design & Test of Computers 27*, 1 (2010), 48–65.