VOLtA: Voltage Over-scaling Based Lightweight Authentication for IoT Applications

Md Tanvir Arafin
ECE Department
University of Maryland
College Park, 20742
e-mail : marafin@umd.edu

Mingze Gao
ECE Department
University of Maryland
College Park, 20742
e-mail : mgao@umd.edu

Gang Qu
ECE Department
University of Maryland
College Park, 20742
e-mail : gangqu@umd.edu

Abstract - Incorporating security protocols in IoT components is challenging due to their extremely constrained resources. We address this challenge by proposing a hardware-oriented lightweight authentication protocol based on device signature generated during voltage over-scaling (VOS). First, we demonstrate that VOS-based computing leaves a process variation dependent error signature in its approximate results. This error can be methodically profiled to extract information about the underlying process variation in the computation unit. We then combine this error profile with security key based authentication schemes to create a two-factor authentication mechanism. To understand the effectiveness of this protocol, we perform detailed security analysis under various attack scenarios. Finally, we simulate the authentication hardware using a process variation aware 45nm design library in HSpice. Simulation results show that our VOS-based assumptions are valid, and this authentication mechanism can withstand basic environmental variations. Overall, our approach provides a unique approach for using hardware process variations as a key for authentication.

I Introduction

Implementation of Internet-of-Things (IoT) is critically dependent on a well-designed interconnected network between the Things *i.e.,* sensors, actuators, data acquisition systems, processing modules, cloud servers *etc*. This requirement has presented a timely opportunity to rethink, redesign and reevaluate current networking strategies. The smartness of an IoT system depends on its widespread of data collection nodes and its intelligent and extensive data processing capability. As the number of interconnected things of an IoT system grows, constraints on the critical factors such as cost, area, and power become tighter, and in most cases forces the designer to either sacrifice some design aspects or innovate new designs to meet all the requirements.

Unfortunately, one of the common sacrifices is made on the security of the low power nodes such as the sensors, data acquisition units or on-field data processing units[4]. Compromising in terms of security for low power modules of an IoT system can pose severe threats to the entire system and its output. Therefore, the design should opt for *Internet-of-Trusted-Things*, where the system components are trusted and provide accurate data to the system.

In this work, we have presented a solution for authenticating low power nodes or Things by profiling the errors produced by the digital system present on the node. The design of digital circuits and systems are focused on the deterministic results, *i.e.,* for the same inputs, any design will yield the same output. As we push the boundaries of power efficiency, we are introducing the effects of analog nature of the circuit components involved in the computation. In the field of approximate computing, one of the common power reduction techniques is voltage over-scaling (VOS). In VOS, the digital circuit used for computation is operated under the nominal voltage, which guarantees correct output for all input conditions under any given operating environment. Since the dynamic power consumed in a VLSI chip is squarely proportional, and static power is proportional the operating voltage, reducing the operating voltage under the prescribed margin can result in considerable power savings. However, the effect of this voltage over-scaling will be translated into errors generated during a computation. The key idea of this work is to use voltage over-scaling to exacerbate the effects of process variation and extract information regarding this variation that can be used for security purposes.

The physical variations have already been widely utilized for security applications in the form of physical unclonable function (PUF). However, PUF is costly in terms of hardware and power consumption, bringing an obstacle for its widespread usage in IoT. By contrast, our proposed lightweight authentication protocol does not require any additional hardware in the low power Things. In this work, our main contributions are as follows:

1. We demonstrate that voltage over-scaling based computing leaves a process variation dependent device signature in the approximate results.
2. Binding this process variation dependent device signature with basic key based authentication schemes; we propose a simple authentication mechanism and perform a security analysis of is protocols.
3. Our work shows that approximate computing can impact security, especially identification and de-anonymization of users and devices.

II. BACKGROUND

Authentication of an entity is one of the most fundamental problems in security. Identification and authentication protocols are designed as a two-party problem- a verifier wants to verify the identity of a claimant. For this work, we will denote the verifier as Alice and claimant as Bob. Alice

can authenticate Bob using a secret that they share. This secret can be derived from three sources: (1) something that is known by both Alice and Bob (such as passwords), or (2) something possessed by Bob (such as hardware keys), or (3) something inherent (such as signatures, biometric signals etc.) of Bob [3]. Adding any two or more items of these three sources can produce strong authentication protocols such as two-factor authentication. However, multifactor authentication tends to require more resource overhead with respect to its single factor counterpart. *In this work, we have introduced a lightweight two-factor authentication scheme that uses passwords as something known and hardware properties as something possessed by Bob.*

Password based single factor authentication is weak and vulnerable to replay attacks, and exhaustive and dictionary search based attacks. Furthermore, weak and straightforward implementations of password-based authentication without hashing or using any cryptographic primitives are widespread which has made the situation worse. Storage of authentication data has also become a security issue since leaking of a verifier's or a valid claimants database containing (username and/or password) can cause significant threat. One can improve the situation by modifying the simple password-based protocols to a better alternative such as challenge-response based authentication, zero-knowledge protocols *etc*. However, all of these variations increase the required resource for implementation. In this era of IoT where the interconnected components are constrained by power, cost and area budgets, performing secure authentication between two entities becomes a challenging problem.

Recent work by Rahmati *et al.* [5] have experimentally demonstrated the deanonymizing effects of voltage over-scaling in DRAM modules. It is found that approximate DRAMs leave device fingerprints in the data stored in memory. As a result, one can easily trace the memory elements used for storing data files using simple image processing and fingerprinting techniques. In this work, we have explored this idea further and demonstrate that signature of process variation not only remains in approximate memory but also in the results produced by an approximate processing element.

III. ERROR IN VOLTAGE OVER-SCALING

A. Basics

As the size of the transistor reduces, the effect of process variation becomes a critical issue in digital design. These variations come from the common factors such as imperfection of the manufacturing process, random dopant fluctuation, and variation in the gate oxide thickness. As the transistor size shrinks, the standard deviation of threshold voltage variation increases, since it is proportional to the square root of the device area [8]

$$\sigma_{\Delta Vt} = \frac{A_{\Delta Vt}}{\sqrt{WL}} \quad (1)$$

where $A_{\Delta Vt}$ is characterizing matching parameter for any given process. This variation in V_t will have a direct consequence in the delay of a CMOS gate which can be approximated using the following equation [8]

$$d_{gate} \propto \frac{V_{DD}}{\beta(V_{DD}-V_t)^{\alpha}} \quad (2)$$

where α and β are fitting parameters for a logic gate of a given process. Therefore, to maintain the timing correctness, static timing analysis of a given design is performed on the process corners. The timing analysis ensures that for a given operating condition, all paths meet the timing requirements to produce correct results irrespective of the input vectors. Scaling V_{DD} from the predefined operating voltage creates large timing errors and degrades the output quality. However, V_{DD} scaling can offer great saving is energy budget, and therefore, voltage over-scaling-based approximate computation received significant research attention in recent years [1–3, 7].

From equation (2) it is evident that with voltage over-scaling and transistor shrinking, underlying device signature due to process variation manifests itself more prominently in the delay output. If proper correction mechanism is not applied, this variation will cause errors the output. If V_{DD} and other operating conditions remain fixed, the error generated by a computational unit due to VOS will retain information about the process variation. Since process variation is a random process, by profiling this error, one can distinguish the computational unit and generate a unique device signature for the circuit.

To understand the effect of process variation in voltage over-scaling, we have studied the error profiles generated by adders. Venkatesan *et al.* have provided process variation independent error profiles for Ripple Carry Adders (RCA), Carry Look-ahead Adders (CLA) and Han-Carlson Adder (HCA) [7]. It was found that the error probability increases as the number of critical paths that fails to meet the timing constraint increase. Therefore, with the presence of randomness in the manufacturing process, the variations in the transistors in the critical paths will have a significant contribution for the errors produced by the adders.

Among the adders, it was found that Han-Carlson adder fares the poorest in terms of producing correct result with voltage over-scaling. RCA performs better than HCA and the probability of error increases slowly with the length of the adder [7]. CLA performs best among these three adders. If one wants to extract fabrication variation related information, she needs to be careful on the choice of circuits. For example, HCA can suppress the variation dependent errors with the errors due to scaling effects, and CLA can represses effect of process variation due to its timing forgiving nature. Therefore, for our discussions, we have used RCA that has the potential to preserve process variation related artifacts.

B. Error Modeling in a Voltage Over-scaled Circuit

If a given circuit is operated with a clock period that is less than the maximum delay produced by the circuit, then the circuit output becomes a function of current and previous input values [7]. Therefore, the output data in a circuit under VOS not only is a function of process variations but also a function of the input values applied to the circuit. Hence, for a combinational adder with two operands X and Y, we can write the current output Z_i of a voltage over-scaled adder as a

function of current inputs X_i, Y_i, and previous inputs X_{i-1}, Y_{i-1}. Therefore,

$$Z_i = f(X_i, Y_i, X_{i-1}, Y_{i-1}) \quad (3)$$

where *f()* defines a process variation dependent addition. This dependence on previous inputs can cause cascading errors in the output. Therefore, to correctly predict the output of a voltage over-scaled adder, one can take the following measures. Let us denote the entity as Alice and assume that she has complete access of the adder during profiling or modeling phase.

i. Save the output data for the set of input patterns that will be used on that circuit. For example, if only a set S_I containing n-input pattern will be used for processing in a given adder, then Alice needs to save outputs for all the combination of the pattern resulting from S_I. This would reveal the partial behavior of the circuit for a subset of input data.

ii. Profile the adder for all possible input patterns. For profiling the adder, one not only needs to consider all possible current input values but also needs previous input values. Therefore, a correct profile of an *n*-bit voltage over-scaled adder would consist a table of entries with all possible current input value times all possible input values in the previous step. This would amount to $2^{2n}*2^{2n}=2^{4n}$ entries of the input values and the corresponding output values. Since all the additions are not incorrect and dependent on the previous values one can significantly reduce the size of the profile by strong only the cases where the adders provide incorrect results.

iii. Use a delay-based graphical model to learn the properties of the adder. Since Alice has the adder, she can profile the device, and use this profile to create a conditional probability table for a Bayesian network.

From the discussions above, we are stating the following requirements to design authentication protocols

R1. Voltage over-scaling can produce process variation dependent errors in a computing unit.

R2. Errors produced due to voltage over-scaling are not random noise but reproducible information since they merely reveal the output of the analog circuit at a lower operating voltage.

R3. If one has the access to the input and output ports of this circuit, he can properly model the behavior the computation performed by this unit.

R4. Such model discussed in R3 would be unique for each computational unit since the manufacturing-dependent process variations are random in nature.

Assuming that one has access to a hardware that satisfies the requirements above, we have designed the authentication protocol discussed in the next section.

IV. AUTHENTICATION PROTOCOL

The authentication protocols assume that the claimant Bob has a voltage over-scaled computation unit that generates process dependent errors. The verifier Alice either knows the correct model or profile to simulate the computation unit and Alice tries to authenticate Bob. The authentication protocol is designed for the scenario where the verifier has sufficient processing and memory capability. Since for distributed IoT systems, we have the sensor nodes communicating with a more powerful gateway nodes or processing units, such units can take the role of Alice. The authentication protocol is given as follows:

Registration

i. Bob has a password K, which consists of two equal keys k_1 and k_2.

ii. Alice registers $K=(k_1, k_2)$, and profiles or models the error patterns of Bob's adder as discussed in the previous section. We would denote the model/profile with M.

Authentication

i. Alice picks a random string R and sends it to Bob.

ii. Bob calculated $L= R+k_1$ using the adder and then calculates $Y= L \oplus k_2 = (R+k_1) \oplus k_2$.

iii. Bob sends Y to Alice.

iv. Alice calculates $L=Y \oplus k_2$ and L'=M(R, k_1). If *the distance (L',L)* $< T$ (where T is the threshold of error tolerance) Alice authenticates Bob.

Note that, the distance in the protocol can be measured by common distance measurement functions such as Hamming distance or Euclidean distance. Also, with multiple keys Alice can authenticate multiple users using the same device. Moreover, Alice can authenticate Bob over different devices given that Alice knows the correct model of those devices.

V. EVALUATION OF THE PROTOCOL

This is a strong two-factor authentication scheme, which requires knowledge of both the secret was known (*i.e.,* key K) and the secret possessed (*i.e.,* properties of the voltage over-scaled adder).

To prove the effectiveness of our proposed VOLtA, we start from analyzing the potential weakness, for the case when we have a perfect adder. If the adder is perfect when calculating L, this protocol is not secure. Assume the following scenario: the malicious attacker Malice is pretending to be Alice, and she wants to resolve Bob's key K by sending some message strings R and receiving the corresponding Y from Bob. Then she will apply eavesdropping and bit manipulation techniques to recover the key.

However, when applying the voltage over-scaling approach, the addition will become non-deterministic because the physical variations will affect the arithmetic result as discussed in section III. Therefore, the result of M(R, k_1) cannot be accurately predicted. Therefore, the bitwise attacking scenario fails. Overall, for the security of this protocol, the uncertainty of the calculation in the function arithmetic () needs to be guaranteed.

For our discussions on threat models and attacks, we assume that Alice tries to authenticate Bob over an untrusted channel where Malice performs the following attacks to obtain the security keys or being erroneously recognized as Bob. Below we discuss the potential attacks on VOLtA.

Random Guessing Attack: The simplest attack that an adversary can perform is on the authentication protocol is to try and randomly guess the authentication keys. To perform this attack, Malice tries to imitate Bob and responds to Alice's query with a random guess. Malice randomly guesses K and

the use random errors to generate *Y*.

The security of VOLtA is tied not only to the key *K* but also the property of the approximate adder. Since Malice neither has information about the key nor about the hardware properties, the success rate of such attacks exponentially decreases with the increase in the number of bits in the security keys and with the increase in the uncertainty of the results produced by the adder.

Eavesdropping Attack: For eavesdropping attack, Malice eavesdrop on a number of communications between Alice and Bob and records Bob's response to each challenge from Alice. Later, for a known query of Alice, Malice can answer correctly using her records.

Alice sending random string R each time can easily thwart such attacks because in that scenario response calculated by Bob will be different for different cases. Since Alice knows the correct model of Bob's circuit, she can send random string every time for authentication. Therefore, VOLtA would be effective in counteracting such attacks.

Man-in-the-middle (MIM) attack: Man-in-the-middle attack constitutes the case where Malice pretends as Alice and communicates with Bob. Malice sends random authentication strings to Bob and collects his response.

This attack would be difficult to perform if Bob has some knowledge about the input sent by Alice. But this would violate the requirement of randomized string to prevent other attacks.

Compromised key: One of the strongest attacks on these protocols is the situation where the key *K* and the model *M* are leaked to an attacker. Since this case breaks the basic Kerckhoffs's principle, both the protocols will fail in the face of such attacks.

Therefore, encryption techniques need to be applied to protect Alice's database to ensure the security of the keys and models. It should be noted that this would provide security in the case when the key *K* is compromised because it is a two-factor authentication protocol where the property of the adder is unknown to the attacker. Therefore, without having the device of the model of the device, the attacker would still have to resort to random guessing or other attacks to resolve a correct response.

Learning-based attack: This attack is a combination of eavesdropping attack and learning attacks. Malice eavesdrop on the communication during authentication and with the challenge and response records, Malice models the voltage over-scaled approximate adder using a learning model. Thus, Malice can create a delay based graphical model for the adder with partial observables. Malice can estimate a conditional probability table and the chance of success in getting the model for the adder will increase with the number of trials that Malice can perform.

However, the output of the adder is XORed with k_2 and therefore such attacks will be difficult to perform without the knowledge about k_2.

Side-channel attack: One of the most common side channel attacking techniques for encryption is static or differential power analysis. Researchers have shown that the power analysis can severely reduce the security of Ring Oscillator PUF, which is a well-known secure primitive for key storage and authentication. However, our proposed voltage over-scaling is more resistant to side channel attack because:

1. The arithmetic units are working under much lower voltage. As we mentioned before, the dynamic power consumed in a VLSI chip is squarely proportional to the supply voltage, the power consumption during authentication process will be very low, making it difficult to capture accurate power consumption.
2. Since the adder does not generate correct result, even though the attacker can measure the very accurate power consumption, he cannot apply the model of an accurate adder for regression. The real model M is hidden in the process variations.

VI. EXPERIMENT AND DISCUSSIONS

We have simulated the basic building blocks of the authentication protocol to analyze that effect of process variations, voltage variations, and temperature, and evaluate the performances of the protocols under general operating conditions. We have performed our simulations in HSpice platform using the FreePDK 45nm libraries[6]. To introduce process variation in our design, we have used 200 modified NMOS and PMOS models with variable threshold voltages. A Gaussian distribution with a ±7.5% standard variation is assumed for the variation of the threshold voltages. This modified NMOS and PMOS transistor models are randomly chosen to build 100 different versions of each standard cell in the FreePDK 45nm library. We have designed our digital circuits in Verilog and synthesize them using the Cadence Virtuoso RT compiler. The synthesized design is then converted into an HSpice netlist with standard cells randomly chosen from our modified library.

As our authentication protocol is based on addition performed in a voltage over-scaled system, we have simulated simple 8-bit ripple carry adder for our analysis presented in this work. This choice is justified in section III. All the output bits of the adder are fed into edge-triggered D-flip-flop to extract correct output bit values in each clock cycles.

TABLE I
Parameters used for simulations

Parameter Name	**Value(s)**
Supply voltage (V_{DD})	0.4V/0.45V/1V
NMOS threshold voltage (V_{tn})	0.322±0.02415V
PMOS threshold voltage (V_{tp})	-0.302±0.02265V
Operating temperature (T)	25 deg. C
Clock Period (T_{clk})	1ns

A. Evaluation

Uniqueness & the Effect of Process Variation

In this section, first, we will validate the requirements presented in section III. For our experiments, we have used eight adders generated from a process variation aware 45nm process. Our experiment at $V_{DD}=0.4V$ shows that the results

generated from voltage over-scaled adders contain errors. To understand the uniqueness of each approximate adder in terms of these errors, we evaluate the variations using the following metrics:

I. The pairwise hamming distance of adder *i* and adder *j*. For each adder, we collect the results on each clock cycle and concatenate all the results to create the complete output bit-stream generated by the adder. Then we calculate the pairwise Hamming distance in of these output bit-streams. We divide the Hamming distance with the length of the bit-stream and report the result in percent in Table II.

II. The pairwise average numerical difference of adder i and adder j: $\sum_{l=1}^{N} \frac{|result(i,l)-result(j,l)|}{N}$, where *N* is the number of output Bytes. The result is shown in Table III.

Metric I, the Hamming Distance, is widely used to measure the difference between two binary bit-streams. However, it omits the bit orders. For example, the hamming distance of the pair (0000_2, 0001_2) and a pair (0000_2,1000_2) are all 1. Therefore, we introduce the second metric to have a measure of the average numerical difference between the values represented by 8-bit outputs of every two adders. Both Table II and III are symmetric since for our measures *distance(i,j)= distance(j,i)*.

TABLE II
Pairwise hamming distance (in percent)

	A1	A2	A3	A4	A5	A6	A7	A8
A1	0	18.82	18.24	18.0	19.4	18.4	18.3	17.52
A2	18.82	0	5.36	5.21	5.67	5.65	3.89	5.39
A3	18.24	5.36	0	4.62	5.98	5.11	5	6.79
A4	18.0	5.21	4.62	0	5.73	3.53	4.13	6.44
A5	19.4	5.67	5.98	5.73	0	6.04	5.59	6.28
A6	18.4	5.65	5.11	3.53	6.04	0	4.96	6.64
A7	18.3	3.89	5	4.13	5.59	4.96	0	5.41
A8	17.52	5.39	6.79	6.44	6.28	6.64	5.41	0

TABLE III
Pairwise average numerical difference between the output from the devices at 0.4V

	A1	A2	A3	A4	A5	A6	A7	A8
A1	0	12.01	14.64	16.11	12.30	14.13	12.13	13.25
A2	12.01	0	9.04	12.30	8.03	9.19	8.08	8.97
A3	14.64	9.04	0	8.43	9.03	7.24	9.07	11.13
A4	16.11	12.30	8.43	0	10.68	8.58	8.88	11.13
A5	12.30	8.03	9.03	10.68	0	6.01	6.51	9.45
A6	14.13	9.19	7.24	8.58	6.01	0	8.11	8.89
A7	12.13	8.08	9.07	8.88	6.51	8.11	0	8.47
A8	13.25	8.97	11.13	11.13	9.45	8.89	8.47	0

In both Table II and III, A1 represents the results from an exact adder. From the first rows of Table II, we can notice that there are about 20% bit flips in the outputs of the voltage over-scaled adders. This provides experimental justification of requirement R1 (discussed in section III).

From the rows and columns of Table II, it is evident that there is a significant difference in the output of two different voltage over-scaled adders. This justifies R4. Note that R2 is also justifiable since we are not introducing any noise or fault in the device but we are merely resorting to the analog nature of the device to get information about process variation.

Moreover, from table III, it is evident that there is a significant average difference between the numerical results produced from each adder.

To illustrate the concepts discussed above we present a simple image processing application based on the superimposition of two images under general operating conditions and compare their results. The image processing application *superimposition* reads the 8-bit values stored at every pixel location for any two given image and add the values. The processing is first done on an accurate adder and then on two voltage over-scaled ripple-carry adders with process variations.

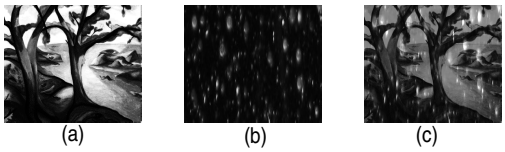

Figure 1. An example of superimposing two images. We have used two gray scale images (a) *trees* and (b) *snowflakes* from MATLAB library to generate the superimposed image (c) *Snowfall.*

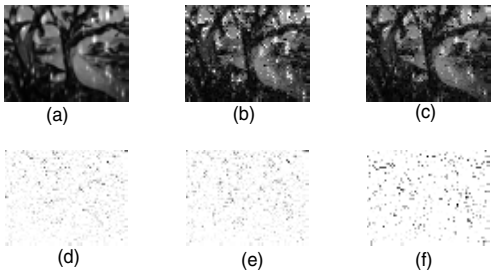

Figure 2. An example of the effect of process variations in voltage over-scaling based computation. In (a) the gray scale image *Snowfall* is computed using *trees* and *snowflakes* without voltage over-scaling; in (b) and (c) the image is computed under voltage over-scaling with two adders which are identical in every aspect, except the process variation of the transistors; (d) and (e) shows the error pattern found in the figure (b) and (c). This error pattern shows the deviations for each adder from the correct image. Subfigure (f) shows the difference between the two error pattern (d) and (e). The source images were downsized to 52x40 pixels for reducing computation time.

From visual observation of Figure 2(a), 2(b) and 2(c), one can clearly notice the effect of voltage over-scaling in this simple image processing application. Figure 2(b) and 2(c) are somewhat distinguishable if an observer pays close attention. We generate the error patterns by calculating the difference between each pixel value in the approximated result and the un-approximated result. Although it is difficult to distinguish the difference between the error pattern shown in Figure 2(d) and 2(e), if we plot their differences as done in Figure 2(f), one can clearly notice the effect of process variation in the approximately computed result.

If we plot the histogram of the Euclidean distances between the Figures 2(a), 2(b) and 2(a), 2(c) (as done in Figure 3) then we can see some interesting results on the error pattern. It can be noted that the peaks of this histograms mainly appears at 2, 4, 8, 16, 32, 64,... which suggests that most of these errors in the approximated results come from a single bit-errors. Furthermore, the peaks at 1,2,4,8,16 are higher in most cases,

revealing the fact that for these two adders most of the errors are in the LSBs.

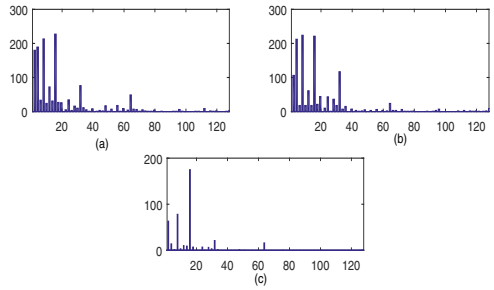

Figure 3. A histogram of the Euclidian distances (a) between the figures 2(a) and 2(b); (b) between the figures 2(a) and 2(c); between the figures 2(b) and 2(c).

Effect of Variations in Supply voltage

The proposed authentication protocol is based on the key concept of voltage over-scaling and the errors it produces. Therefore, variations in supply voltage will definitely cause reliability issues for the proposed design. Therefore, care must be taken to ensure voltage supply with a minimal amount of noise for proper implementation of the protocol. In Figure 4, we have plotted the response of voltage over-scaled adders at 0.45V. This is a bit higher voltage than the one used for the results reported in Table II. It can be noted that as the voltage increases the overall Hamming Distance decrease, which represents the eventual convergence of all the adders to the correct output at sufficiently higher voltage.

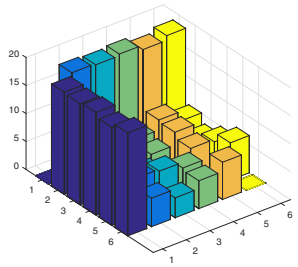

Figure 4. Hamming distance (in percent) between devices at 0.45V . Here 1 represents a correct adder (A1) and 2-6 represents voltage over-scaled approximate adders (A2-A6).

Effect of Variations in Temperature

Variations in operating temperature can also affect the protocols. To understand the effect of temperature, we have calculated the percentage Hamming distance between the results of the same adder at different temperatures as shown in Figure 5.

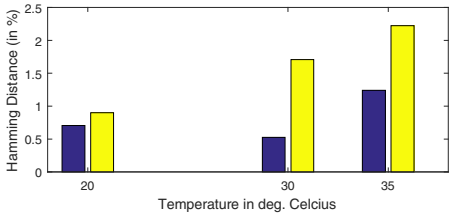

Figure 5. Temperature dependent bit flips for two different adders. The distance is calculated from the results produced at T=25 degree Celsius. The blue (dark) line represents the temperature dependent bit-flip for adder A2 and the yellow (light) line represents the adder A3.

We find that ±5 degree variations result in less than 1% bit-flips. Therefore, by carefully calibrating the threshold of error tolerance *T* of the protocols, one can negate the effects of minor temperature variations.

Deanonymization Issues

It should be noted that by profiling the input patterns of a voltage over-scaled circuit, one could easily deanonymize the approximate circuit later based on the physical variations. This exposes a critical flaw in voltage over-scaling-based approximate computing. Therefore, the anonymity of the approximate devices should also be thoroughly studied. This work can be considered a step towards such analysis.

VI. CONCLUSIONS

In this work, we introduce VOLtA, a voltage over-scaling based user authentication scheme that uses the physical random variation of a VLSI system as a key for authentication of IoT nodes. VOLtA profiles the hardware used for computation in a reduced voltage operation and uses the underlying hardware fingerprint for authentication purposes. This authentication protocol requires no additional hardware on the claimant side to implement. This lightweight protocol can be useful in IoT applications where the interconnected Things face extreme power, cost and area constraints.

Acknowledgments

This work was supported by AFOSR MURI under award number FA9550-14-1-0351.

References

[1]Banerjee, N. et al. 2007. Process variation tolerant low power DCT architecture. *Proceedings -Design, Automation, and Test in Europe, DATE*. 07, (2007), 630–635.

[2] Han, J., and Orshansky, M. 2013. Approximate computing: An emerging paradigm for energy-efficient design. *Proceedings - 2013 18th IEEE European Test Symposium, ETS 2013*. (2013), 1–27.

[3]Menezes, A.J. et al. 1997. Handbook of Applied Cryptography. *Electrical Engineering*. 106, (1997), 780.

[4] Qu, G., and Yuan, L. 2015. Design THINGS for the Internet of Things - An EDA perspective. *IEEE/ACM International Conference on Computer-Aided Design, Digest of Technical Papers, ICCAD*. 2015-January, January (2015), 411–416.

[5]Rahmati, A. et al. 2015. Probable cause: The deanonymizing effects of approximate DRAM. *Computer Architecture (ISCA), 2015 ACM/IEEE 42nd Annual International Symposium on*. (2015), 604–615.

[6]Stine, J.E. et al. 2007. FreePDK: An Open-Source Variation-Aware Design Kit. *Microelectronic Systems Education, 2007. MSE '07. IEEE International Conference on*. (2007), 173–174.

[7]Venkatesan, R. et al. 2011. MACACO: Modeling and analysis of circuits for approximate computing. *IEEE/ACM International Conference on Computer-Aided Design, Digest of Technical Papers, ICCAD*. (2011), 667–673.

[8]Wirnshofer, M. 2013. Variation-Aware Adaptive Voltage Scaling for Digital CMOS Circuits. 41, Ic (2013).