IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received June 2, 2021, accepted June 6, 2021, date of publication June 14, 2021, date of current version July 2, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3088991

Low-Cost Side-Channel Secure Standard 6T-SRAM-Based Memory With a 1% Area and Less Than 5% Latency and Power Overheads

YOAV WEIZMAN[1], ROBERT GITERMAN[2], ORON CHERTKOW[3], MAOZ WICENTOWSKI[4], (Student Member, IEEE), ITAMAR LEVI[1], ILAN SEVER[5], ISHAI KEHATI[6], OSNAT KEREN[1], (Member, IEEE), AND ALEXANDER FISH[1], (Member, IEEE)

[1]Faculty of Engineering, Bar-Ilan University, Ramat Gan 5290002, Israel
[2]Institute of Electrical Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland
[3]Amdocs, Israel Ltd., Raanana 4321545, Israel
[4]Samsung Electronics Israel, Petah Tikva 49514, Israel
[5]Weebit Nano, Hod Hasharon 4527713, Israel
[6]Ethernity Networks, Lod 7152025, Israel

Corresponding author: Robert Giterman (robertgi316@gmail.com)

This work was supported in part by the Israel Science Foundation under Grant 2511/20 and Grant 1266/20, and in part by the Israel Innovation Authority in the frame of the Magneton Program.

ABSTRACT Side-channel attacks constitute a concrete threat to IoT systems-on-a-chip (SoCs). Embedded memories implemented with 6T SRAM macrocells often dominate the area and power consumption of these SoCs. Regardless of the computational platform, the side-channel sensitivity of low-hierarchy cache memories can incur significant overhead to protect the memory content (i.e., data encryption, data masking, etc.). In this manuscript, we provide a silicon proof of the effectiveness of a low cost side-channel attack protection that is embedded within the memory macro to achieve a significant reduction in information leakage. The proposed solution incorporates low-cost impedance randomization units, which are integrated into the periphery of a conventional 6T SRAM macro in fine-grain memory partitions, providing possible protection against electromagnetic adversaries. Various blocks of unprotected and protected SRAM macros were designed and fabricated in a 55 nm test-chip. The protected ones had little as 1% area overhead and less than 5% performance and power penalties compared to a conventional SRAM design. To evaluate the security of the proposed solution, we applied a robust mutual information metric and an adaptation to the memory context to enhance this evaluation framework. Assessment of the protected memory demonstrated a significant information leakage reduction from 8 bits of information exposed after only 100 cycles of attack to less than ~1.5 bits of mutual information after 160K traces. The parametric nature of the protection mechanisms are discussed while specifying the proposed design parameters. Overall, the proposed methodology enables designs with higher security-level at a minimal cost.

INDEX TERMS Secured Static Random Access Memories (SRAM), hardware security, power analysis, secured memory.

I. INTRODUCTION

Today's systems-on-chip (SoCs) are built to respond to the challenges of multiple applications and environments. These range from the fairly well protected and complex automotive-system, through more resource-constrained mobile-devices to a stand-alone network connected node. All have specific security requirements related to storage, processing and communication of sensitive information. As a result, the security of these electronic systems has become an ongoing concern for both research and industry.

Side Channel Analysis (SCA) attacks are powerful threats to cryptographic devices because they exploit internal sensitive information related to their physical behavior [1], [2]. Power Analysis (PA) and Electromagnetic Analysis (EMA) attacks are considered to be powerful types of SCA since they require relatively simple equipment and setups [1], [3]–[11]. PA and EMA attacks exploit the correlation between an instantaneous physical measurable quantity (current or radiation) measured from a device and its

The associate editor coordinating the review of this manuscript and approving it for publication was Gian Domenico Licciardo.

internally processed and stored data, which is used to extract secret data or sensitive information.

Secured Static Random Access Memories (SRAM) are instrumental for security purposes: they can be instantiated as part of a Root-Of-Trust used for storing or boot-loading a system, intermediate computations (L0, L1 cache, Register-files, FIFOs etc.), or used to store long(er)-term secrets (on-chip SRAM macros). They are also utilized as building blocks to construct security primitives such as P/T-RNGs, store or load seeds and SRAM-PUFs.

To illustrate the risk of using an insecure memory, the authors in [12] showed that even within a 'secure-design', an instantiation of insecure memory instances undermines the system's security. They demonstrated how local register-files, shift-registers and small caches on a micro-processor (L0 and L1) degrade the security of a practical real-world data-masked system. Clearly, this issue scales up when sensitive information is exposed to higher-caches, e.g, through a *load/store* in a software implementation.

In many scenarios where the incorporation of side-channel protection on logic-layers is needed, doing so on low-level cache hierarchies is advantageous. Low-level caches (e.g. embedded memories) dominate the area and power consumption of many VLSI system-on-chips (SoCs) [13], and are key components of many cryptographic systems, such as smart cards [14] and wireless networks employing cryptography algorithms [15], where they are used to store instruction code and data. Therefore, the analysis and design of secured embedded memories is of utmost importance. The mainstream embedded memory solution for most systems is based on the 6-transistor (6T) SRAM macrocell. However, conventional 6T SRAM cells are susceptible to power/electromagnetic analysis attacks since they leak sensitive information as a result of the correlation between the current drawn from the memory supply and the data that it stores [8], [9], [16], [17].

The design of SCA-resilient digital circuits often engenders significant area and power penalties due to the additional devices used to reduce the correlation between the proceed data and the gate's power consumption [9], [16], [18]; however, these overheads are unmanageable in case of the design of secured embedded memories, which already create both area and power bottlenecks in many SoCs [19], [20]. Previously published secured SRAM implementations have focused on bitcell-level solutions, where modified SRAM cells were proposed to reduce the information leakage drawn from the bitcell array supply [8], [9], [16]. However, these solutions resulted in significant area and power overheads due to the additional devices added to the original 6T SRAM implementation. Architecturally, encrypting the memory content or masking it induces significant overhead and costs which are imperative to avoid. Finally, "pushed" design rules which are typically used by foundries to further reduce the size of the 6T SRAM cell cannot be used when designing modified bitcells, thus resulting in an even larger area overhead when trying to modify a design at the bit-cell level.

In this paper, we first demonstrate how data stored in a conventional 6T SRAM macro can be successfully extracted using PA attack algorithm that exploits the correlation between the memory content and the current drawn from its supply voltage [10]. This refutes the common misconception that algorithmic-noise or large memory activity can repel such attacks.

Memory macros vendors should typically provide strong guaranties for the security of their product, regardless of how it is utilized in an actual design. That is, security evaluations conducted within a theoretical model framework, as done for example on secured encryption blocks in C/KPA-like adversarial control. Clearly, it is assumed here that sensitive information can 'flow through' the memory, and is not static, like a loaded and kept key. Concretely, we describe different methodologies to quantify memory resiliency to SCA utilizing robust mutual information techniques. We evaluate the information leakage of the 6T SRAM macro and show that to mitigate SCA a low-overhead solution for a secured SRAM implementation [10], featuring an Impedance Randomization Unit (IRU), randomly fluctuates the current drawn from the memory power supply to reduce the information leakage from the memory at ultra-low cost. Our solution is designed to maintain the high density of foundry-supplied 6T SRAM bitcells, so that the IRU can be added to the periphery of a conventional 6T SRAM array, resulting in less than 1% area overhead compared to a baseline unsecured implementation. Finally, we compute the additional cost of embedding our solution in a *local* way to resist a powerful and localized EMA adversary. Silicon measurements of the proposed memory macro prove it can provide a significant information leakage reduction compared to a conventional SRAM, without any power or latency overheads.

A. CONTRIBUTIONS

The major contributions of this paper are as follows:

1) We present a secured SRAM implementation featuring a novel impedance randomization unit that provides a significant reduction in information leakage compared to a conventional SRAM array [10].
2) We advance the state-of-the-art on security evaluation of SRAMs by utilizing information-theoretic tools tailored for this purpose.
3) The solution is implemented with only 1% macro area overhead, and no penalties on speed and power, compared to the 1.4×–2.1× and 1.24×–2.56× latency and power overheads, respectively, of competing solutions.
4) A commercial SRAM macro and the proposed secured SRAM implementation are measured and compared in terms of maximum correlation analysis and mutual information to evaluate their information leakage under process-temperature-voltage variations.
5) A cost analysis of making the solution local against EMA attacks is also provided.

This paper is an extended version of our original work published in [10]. It contains more detailed explanations and key additions compared to the original publication, including the following:

1) Detailed explanations of the proposed side-channel attack resilient 6T SRAM macro.
2) An in-depth analysis of the security evaluation metrics used to evaluate the resiliency of digital circuits to side channel power attacks with an emphasis on the Mutual Information metric. (Section IV)
3) A detailed description of the experimental setup used to evaluate the protection mechanism integrated into the fabricated test-chip, and a comparison of its security versus a conventional 6T SRAM design. (Section V)
4) Extended measurement results and analysis (Section VI) with the following additions: (a) Analysis of the solution effectiveness in terms of the correlation ratio for different secret words, different numbers of attack cycles, various physical addresses, and memory configurations. (b) Analysis of the solutions effectiveness in terms of Mutual Information (MI), based on an in-depth analysis of the MI obtained under different numbers of traces and Hamming weights. (c) Analysis of the different security metrics under different operating voltages. (d) Analysis of the different security metrics under different temperatures.
5) Evaluation of the cost of locality of the solution, by providing an in-depth analysis of the area overhead for different block sizes versus the granularity of the embedded randomization units.

The remainder of this paper is organized as follows: Section II describes the power analysis of a conventional 6T SRAM array, Section III presents the proposed solution, Section IV describes the metrics used for the evaluation of the memory security adherence, Section V describes the experimental setup used for the evaluation and the analysis procedure, Section VI reports the measurement results, and Section VII concludes the paper.

II. POWER ANALYSIS OF A 6T SRAM ARRAY

Fig. 1(a) shows a conventional 6T SRAM macro composed of a bitcell array based on 6T SRAM cells, a row decoder which enables a single word-line (WL) during write or read operations, sense amplifiers, and drivers controlling the voltage of the bit-line (BL) and the bit-line bar (BLB) pair. The memory V_{DD}, connected to the voltage supply of the bitcell array, is separated from the core V_{DD}, which supplies the memory peripherals to reduce the power consumption during sleep mode by decreasing the voltage during standby [21]. The current drawn from the memory V_{DD} during write cycles depends on the difference between the written and previously stored data, as illustrated in Fig. 1(b) that shows two consecutive write cycles of '1' and '0'. During the first write cycle, the BL/BLB pair voltages are equal to the voltages of Q/QB in the cell, whereas during the second write cycle, current is drawn from the supply of the cell to flip the stored value from '1' to '0'. The simulated current consumptions of a 6T memory macro of 8k Bytes are shown in Fig. 2 for both a single bit-flip (blue) and no bit-flips (red). A zoomed-in version of the drawn current at the point-of-interest of maximum difference indicates close to a 20 μa difference between the peak currents of both scenarios, resulting in valuable information leakage. Although this difference seems negligible, we show in section VI that it is sufficient to extract the memory's secret content simply by implementing a small number of trace acquisitions.

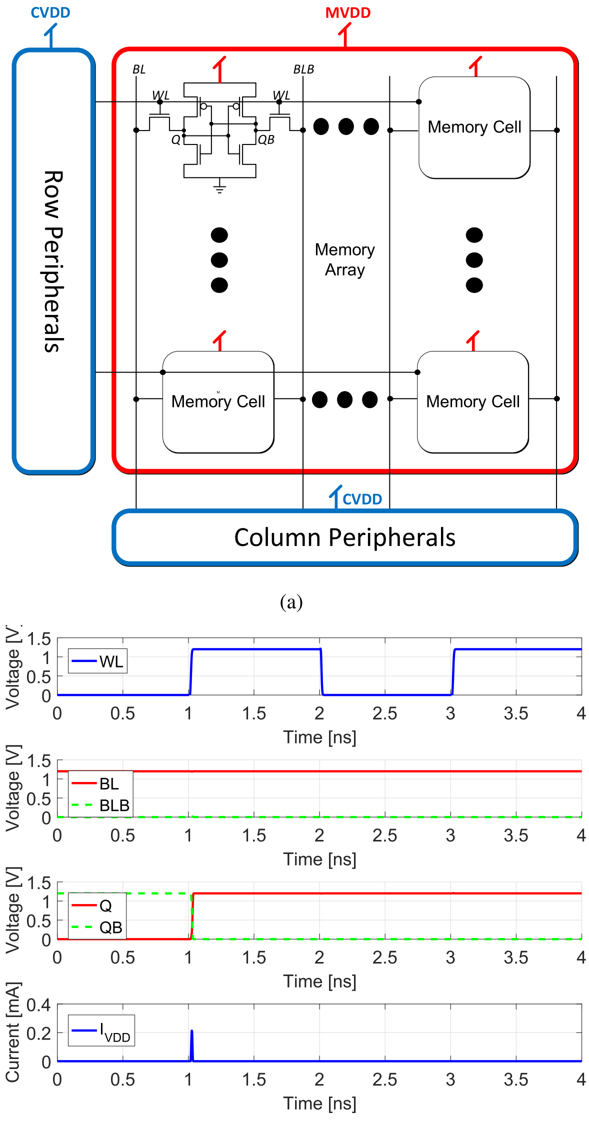

(a)

(b)

FIGURE 1. (a) Conventional memory architecture. (b) Waveform illustration of consecutive write operations.

To exploit the information leakage drawn from the memory during write cycles, below we show how data stored in an 8-bit SRAM word can be successfully extracted using a PA attack procedure.

This attack procedure assumes that the attacker has knowledge of the memory architecture, the ability to assign input vectors to the memory macro, and access to the current drawn from the memory V_{DD}. These assumptions

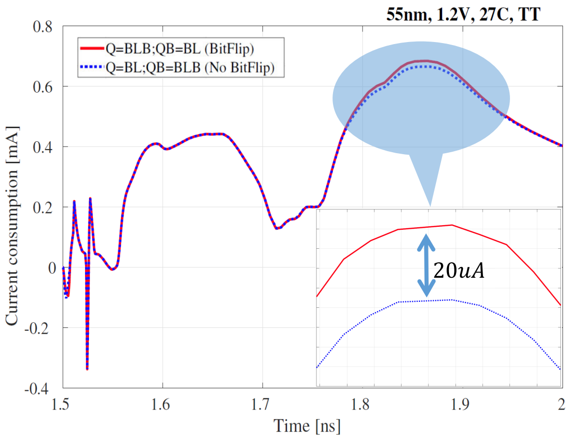

FIGURE 2. Current consumption of the memory V_{DD} during write '1' and '0' operations.

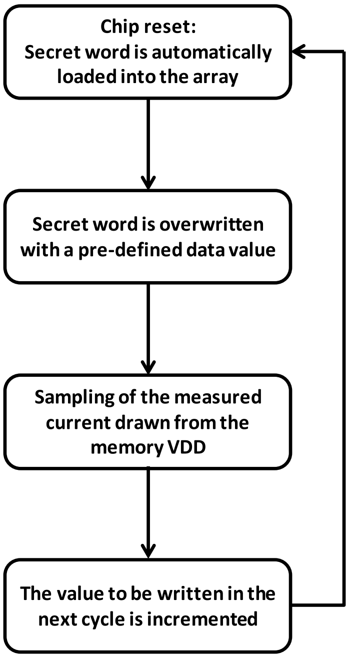
Chip reset:
Secret word is automatically loaded into the array
Secret word is overwritten with a pre-defined data value
Sampling of the measured current drawn from the memory VDD
The value to be written in the next cycle is incremented

FIGURE 3. PA attack procedure [10].

are commonplace in the literature and are given as a standard test-case to demonstrate the vulnerability of the array to power analysis attacks. Fig. 3 depicts the PA attack procedure, starting with a chip reset in order to bring the memory to its initial state when the secret word is already loaded in the array. Next, the secret word is overwritten with pre-determined data, and the current drawn from the memory V_{DD} is measured. Then, the value of the written data is changed and the operation is repeated until all the combinations of the different memory words have been applied to the array. Note that due to parametric process variations of the memory bitcell [22], [23], this operation is typically repeated multiple times to filter out noise which causes the traced current of the bitcell array to alternate even when the same data stored in different words. Finally, correlation analysis is performed by estimating the current for each data combination (stored and overwritten) and by finding its correlation to the measured current. (a formal description of the attack procedure appears in Section IV). The measured correlation coefficients for each data guess are shown in Fig. 4, with the highest correlation of 0.7 achieved for the correct secret word stored in the array, which was arbitrarily selected at 192_{16}.

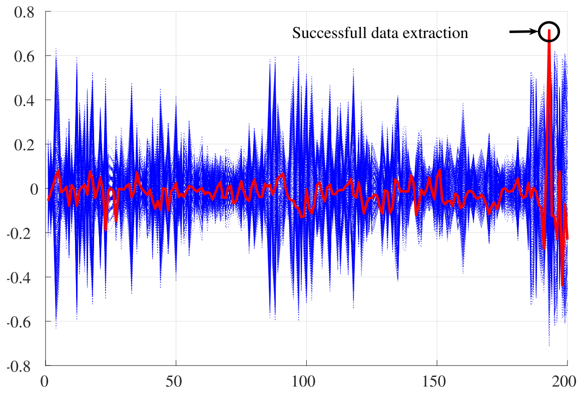

FIGURE 4. Measured correlation coefficients for each data guess over time, illustrating a successful secret data extraction.

III. PROPOSED SOLUTION

To reduce the information leakage of a conventional SRAM array, an IRU is added to the peripheral circuitry of an SRAM macro. The IRU controls the impedance of the memory V_{DD} to alter the current consumption during write cycles to the array, regardless of the stored data, thus reducing the information leakage. The IRU is consistent with the modularity required by memory arrays, since it is independent of the bitcell structure, and added on top of the conventional peripheral circuitry. The IRU is composed of power switches, implemented with PMOS transistors, and connected in parallel between the external V_{DD} and the memory V_{DD} applied to the array. Note that power-gating (PG) cells are provided in standard power cell libraries so that specially designed cells are not required. An always-on power switch maintains a stable memory V_{DD} to avoid latency penalties, as well as maintain the static noise margins and data retention voltage [24] achieved for the original foundry based SRAM array. The SRAM array is guaranteed to adhere to the static noise margin and data retention voltage characteristics under these conditions. The additional power gates, connected in parallel to the always-on power switch, are controlled by randomized input signals, which are supplied by a linear-feedback shift register (LFSR). The sizing and number of PMOS devices were selected according to the minimum impedance randomization required to mitigate the information leakage through the current consumption. In practical terms, the sizing parameters and the number of devices serve as design parameters for security engineers. Here our goal was to induce a uniform leakage distribution where the wider the distribution, the more secure our outcome design will be. To achieve a higher security level to trade

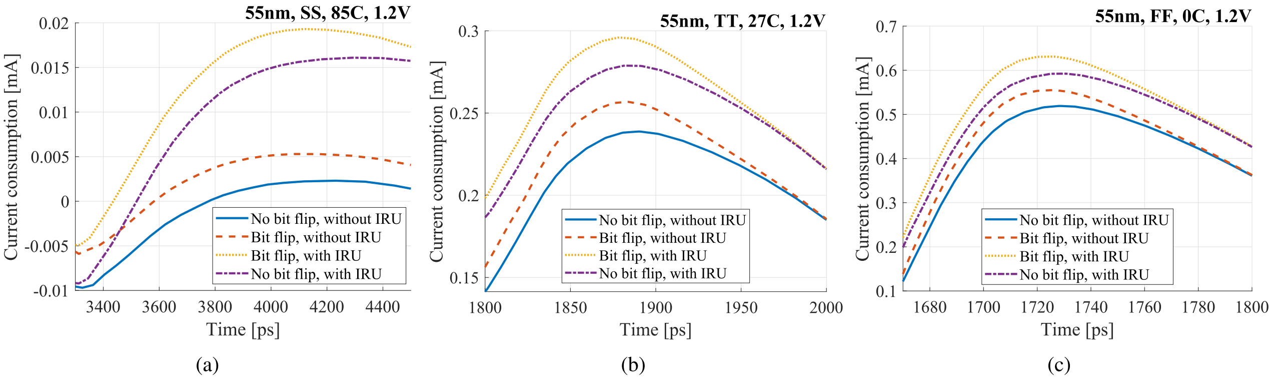

FIGURE 5. Memory V_{DD} peak current consumption vs. time during write cycles to a memory word at (a) Slow, (b) typical, and (c) fast corners.

off area utilization, more devices need to be embedded in parallel. To uniformly distribute the leakage, one simple sizing strategy would be to *ladder*-size randomized devices (in terms of W_{min}/L_{min}), i.e. {x1,x2,x4..}, as proposed in [25], [26]. This would uniquely and uniformly distribute the leakage. Concretely, in the manufactured design, to achieve a 1% area overhead we embedded only three randomization devices for each memory column given the signal and noise components of the device. Fig. 5 depicts the current consumption of the memory array during write operations with and without an IRU across typical, fast, and slow corners, as extracted from post-layout simulations, which included the extracted parasitic components of the memory. The additional currents contributed by a single PMOS device were large enough to exceed the difference between the currents drawn from a single cell with and without a bit flip under all process corners, thus validating the successful masking of the information leakage. Note that with a single PMOS enabled, the current consumption for a bit flip is still higher than without a bit flip. However, since the number of PMOS devices enabled in each cycle was randomized, the information leakage of the memory was significantly lessened. Fig. 6 shows the layout of the 1024 × 8-bit SRAM macro with the integrated IRU. The IRU block was pitch-fitted to the memory array widths and placed on top of the memory macro to maintain a low area overhead. The additional units consumed slightly less than 1% of the total macro area. Clearly, when sensitive data is read-in or out from memory macros it can also leak information; therefore we were also interested in equivalently randomizing the leakage of the row/column decoders and sense-amplifier circuits. Without loss of generality and since our modifications are macro-external, it is trivial to do so.

Our fully digital LFSR only requires an equivalent area of 20 FFs and 6 XOR/MUX gates resulting in a negligible 0.2% area overhead, thus negligible. It repeats its cycle after 2^{20} cycles and we assume a refresh from a TRNG seed. In addition, a power-gating mechanisms to save leakage energy is standard for SRAM macros so that the energy and latency degradation as compared to the standard IP designs was measured to be less than 2%.

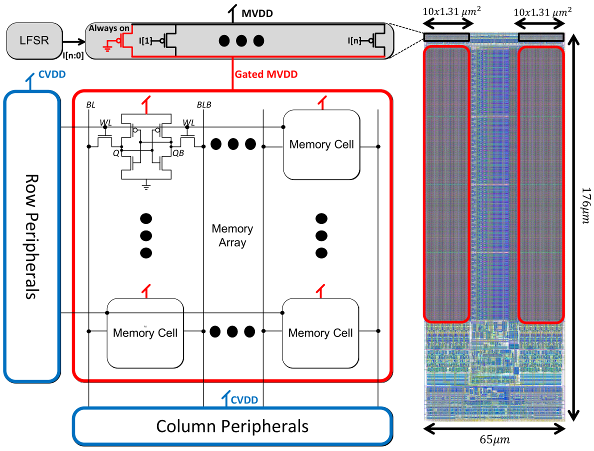

FIGURE 6. Proposed secured memory architecture and the layout view of a 8 kbit memory macro.

We did not restrict the seeding mechanisms of the LFSR PRNG, which can be seeded from either a TRNG or from securely communicated/computed randomness. Clearly, such architectural decisions should be kept in the control of, for example, a security-architect of a company with the sole stipulation of requiring sufficient re-seeding. LFSR reseeding requires 20 clock cycles but rarely takes place, leaving no concrete impact on latency. The SRAM was designed to work in frequency ranges of 1-250MHz, as well as the LFSR and control circuitry.

IV. SECURITY EVALUATION METRICS

The correlation ratio is a typical metric to evaluate the security performance of a module. Its definition depends on the attack procedure. In this work the correlation ratio ρ was defined as follows:

Let $x = (x_1, x_2, \ldots x_n)$ be a secret word stored in a memory, and let $z = (z_1, z_2, \ldots z_n)$ be a (pre-defined) word written by the attacker over x. The current required to change the content of the memory from x to z is proportional to the Hamming distance between these two words; that is

$$I_{model}(x, z) = \sum_{i=1}^{n} (x_i \oplus z_i).$$

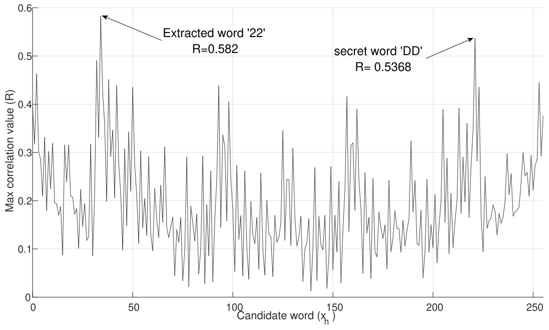

FIGURE 7. Maximal correlation values for example 1.

In order to extract the secret word x, the attacker (who does not know x) chooses a z and measures the consumed current $I_{measured}(x, z, t)$, where t denotes time. Then the attacker chooses a hypothesized word, say x_h, calculates the hypothesized current $I_{model}(x_h, z)$, and computes the correlation between the modeled current and the measured current. Note the differences between I_{model} and $I_{measured}$: both depend on z which is determined by the attacker. However, I_{model} is a function of the hypothesized (secret) word and thus takes an integer value, whereas $I_{measured}$ is a waveform; i.e., it is a function of the secret word. In order to filter out thermal and switching noises, N measurements $\{I_{measured,j}(x, z, t)\}_{j=1}^{N}$ are taken. The average current $I_{measured}(x, z, t)$ is used to compute the correlation. This, in turn, improves the attack success rate because higher correlation values are obtained. The correlation is defined as follows:

$$R(x, x_h, t) = \sum_{z\in\{0,1\}^n} E(I_{measured}(x, z, t)) \cdot I_{model}(x_h, z), \quad (1)$$

where $E(\cdot)$ stands for the expected value. The attacker takes the x_h associated with the maximal correlation value as the extracted word y. Formally,

$$y = \arg \max_{x_h\in\{0,1\}^n} \max_{0\leq t\leq T} R(x, x_h, t) \quad (2)$$

where T is the clock period.

Note that in Figure 4, the wrong x_h's (marked in blue) and the correct one (marked in red) attain their maximal correlation value at different times. For this reason, in Eq. 2 the maximum is taken over the entire clock period and the word space.

The correlation ratio indicates the sensitivity of the design to side channel attacks and it is defined as

$$\rho = \frac{\max_t R(x, x, t)}{\max_{x_h\neq x} \max_t R(x, x_h, t)}. \quad (3)$$

Clearly, ρ values greater than one indicate that the secret word is likely to achieve the maximal correlation and hence the attack is most likely to succeed.

Unfortunately, ρ by itself is not sufficient to evaluate the amount of information leakage. Moreover, a ρ below one does not guarantee that the attack will fail. The following example clarifies this notion:

TABLE 1. Maximum correlation values for the words 'DD' and '22' in example 1.

# Traces	Max correlation of 'DD'	Max correlation of '22'
1878	0.5368	0.582
3744	0.549	0.60
5623	0.5506	0.608
7529	0.55355	0.614
9486	0.551	0.624

Example 1: Figure 7 shows the maximal correlation value $\max_t R(x, x_h, t)$ for the secret word $x =' DD'$. The attack was performed on an unprotected memory array operating with $V_{DD} = 1.2V$ by averaging $N = 1900$ traces of the consumed current. As can be seen in the figure, the secret word has

$$\max_t R(x =' DD', x_h =' DD', t) = 0.5368$$

whereas the word that has the maximal correlation is the negation of $'DD'$, that is,

$$\max_t R(x =' DD', x_h =' 22', t) = 0.582.$$

The reason why the negated word has a similar correlation value has to do with the symmetrical structure of the memory word and the symmetrical nature of the attack overwrite flow. Since there are several secret word candidates with similar correlation values, this will essentially lead to a correlation ratio close to 1; for instance in the above example, $\rho = 0.5368/0.582 = 0.9223$ which is less than one. The same holds for larger the number of traces where the word '22' achieves the best correlation value and the secret word has the second best (see Table 1). This also occurs for other secret words. Consequently, although the correlation ratio is less than one, the memory is unprotected because the secret word is always one of the two words that have the highest correlation value.

In this work we use a stronger metric which indicates information leakages from the memory in a more reliable manner than the correlation ratio. We evaluate the Mutual Information (MI) between the secret word, denoted by the random variable X, and the estimated word, say Y, as extracted by the attack. Both X and Y are random variables that take values from $\{0, 1\}^n$ where n is the word's length; X is assumed to be a uniformly distributed discrete random variable because it can take any value with equal probability, and Y is a random variable whose probability mass distribution depends on X, the thermal noise, process variations, the number of recorded traces etc.

The MI metric indicates the *average* number of secret bits of X that can be learned by observing Y. Formally, denote by $p_{X,Y}(x, y)$ the joint probability $Prob(X = x, Y = y)$ that the random variable X will take the value x and the random variable Y will take the value y, $x, y \in \{0, 1\}^n$. Similarly, let, $p_X(x)$ and $p_{X|Y}(x|y)$ denote the probability $Prob(X = x)$ and the conditional probability $Prob(X = x|Y = y)$, respectively.

The mutual information $I(X; Y)$ is defined as follows [27]:

$$\begin{aligned} I(X;Y) &= H(X) - H(X|Y) \\ &= -\sum_{x,y\in\{0,1\}^n} p_{X,Y}(x,y)\log_2(\frac{p_X(x)}{p_{X|Y}(x|y)}). \end{aligned}$$

In our case $p_X(x) = 1/2^n$. The joint probability $p_{X,Y}(x, y)$ was calculated by conducting N attacks for each secret word, and $p_{X|Y}(x|y)$ was computed from $p_{X,Y}(x, y)$. Namely,

$$p_{X|Y}(x|y) = p_{X,Y}(x,y)/\sum_{x\in\{0,1\}^n} p_{X,Y}(x,y).$$

In fact, the secret word x can be uniquely specified by its Hamming weight, w, and the location of its "ones". There are $\binom{n}{w}$ vectors of Hamming weight w. These vectors form an ordered set; denote by l the index to this set. That is $x \equiv (w, l)$. Thus the mutual information can be represented as [27],

$$I(X;Y) = I(W,L;Y) = I(W;Y) + I(L;Y|W). \quad (4)$$

A comparison of the information that can be extracted from a protected memory array and an unprotected one, as well as the sensitivity of the MI to the number of attacks for each X, W, and L appear in section VI-B.

V. THE EXPERIMENTAL SETUP

A. THE TECHNOLOGY UNDER TEST

A test-chip containing both conventional and secured 8 kbit SRAM macros was implemented in a 55 nm uLP-eF CMOS technology. A microphotograph of the implemented test-chip with its key features is shown in Fig. 8. The test-chip included an on-chip PA built-in-self-test (PA-BIST), which implemented the attack procedure described in the previous section. The PA-BIST was configured with different attack parameters, such as the secret data and address, word size, and the number of traces for secret data extraction. The PA-BIST parameters were stored on dedicated on-chip registers, configured through a serial interface on a separate power domain. In the manufactured chip we implemented 8, 16 and 32 Kbit unprotected and protected arrays. The security analysis which follows relates to the 8Kbit which clearly served as our worst case (the least noisy).

B. POWER TRACES RECORDING PROCEDURE

The experimental setup is depicted in Fig. 9. The DUT was mounted on a dedicated evaluation board designed for an isolated and noiseless environment optimized for this type of attack (c). The evaluation board allowed serial access to the configuration registers of the BIST with isolated and separate power supply rails for each memory instance in the design. In addition, the configuration logic and the BIST were connected to an isolated power supply rail, to eliminate possible fictitious correlations to the secret information. As a power source we used a Keysight low noise E3630 series power supply. The DUT configuration word was programmed through a Keysight 16860A Logic Analyzer (a). To record the power consumption traces we used Keysight N2750A differential 3.5GHz probes (b). The power trace was captured from the voltage drop across a low-noise shunt resistor placed on the PCB in proximity to the DUT. The acquisition was made using a high resolution oscilloscope (d) with a sampling rate of 2GS. The acquired trace was recorded and processed to filter out noise and aligned to allow synchronization between different attack vectors.

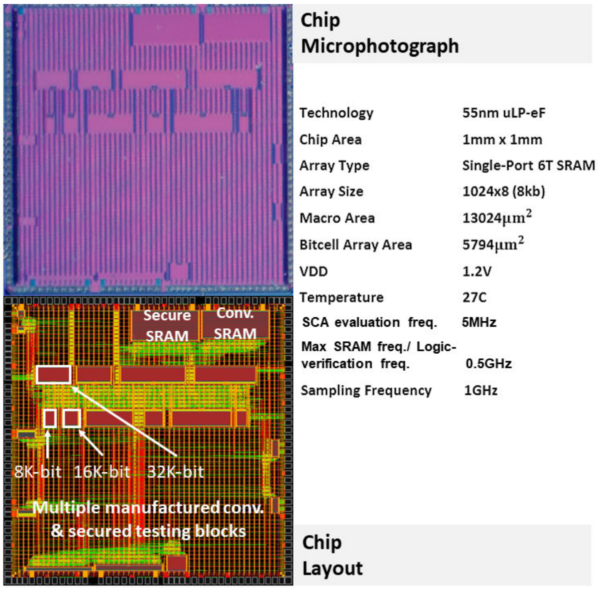

Technology	55nm uLP-eF
Chip Area	1mm x 1mm
Array Type	Single-Port 6T SRAM
Array Size	1024x8 (8kb)
Macro Area	13024μm^2
Bitcell Array Area	5794μm^2
VDD	1.2V
Temperature	27C
SCA evaluation freq.	5MHz
Max SRAM freq./ Logic-verification freq.	0.5GHz
Sampling Frequency	1GHz

FIGURE 8. Chip micro-photograph, layout and key features.

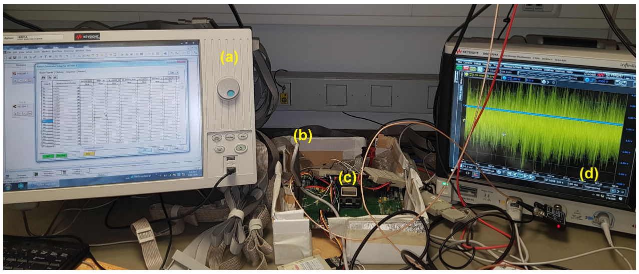

FIGURE 9. Experimental setup.

C. ANALYSIS PROCEDURE OF THE RECORDED DATA

The attack flow was described in Section II and illustrated in Fig. 3. Specifically, a sequence of 256 attack vectors were applied in each attack sequence. In each attack vector the secret word was over-written sequentially with all the possible byte configurations. The attack scenario requires some knowledge of the memory access timing; otherwise, the noise will overwhelm the faint signal during the attack cycle. Therefore we limited the analysis to a temporal window in which an actual over-write occured into the specific address under attack. The power trace was further processed in Matlab to filter out strong noise components by applying a bandpass filter of the 50-500 MHz bandwidth. Finally, we evaluated the correlation between the measured power trace and a power consumption model as shown in Equation 2.

VI. MEASUREMENT RESULTS AND ANALYSIS

A. THE EFFECTIVENESS OF THE SOLUTION IN TERMS OF THE CORRELATION RATIO

To evaluate the effectiveness of the IRU on the test chip, we executed a CPA attack scenario on one of the memory implementations combined with the IRU protection unit. This memory implementation was composed of an 8 bit word and a total of 1024 addresses. For reference we executed an identical attack flow on a module implemented without any countermeasures.

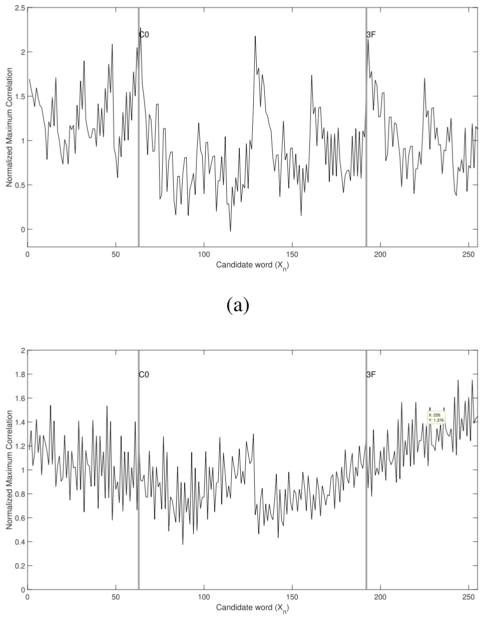

(a)

(b)

FIGURE 10. Normalized maximum correlation values for each secret word candidate for the non-secured memory (a) and secured memory (b), the grey vertical lines mark the secret word in this attack ('0C') and its negation value ('3F').

There was a significant difference between the correlation values extracted after the attack on the protected and unprotected arrays in terms of the number of secret word candidates whose correlation values were close to the correlation value of the correct word. Fig.10 clearly shows that for the unprotected memory there are several peaks that signify the possible secret word candidates. On the other hand in the protected memory no such reduced set of candidates for selection can be identified. However, this metric is computationally exhaustive and can only be used under certain assumptions that will be explained in the next section. Meanwhile, in order to gain some insight into the many physical attributes of the secured memory that can impact its security adherence, we needed to modify the correlation ratio metric defined in Eq. 3 slightly in a manner that would enable us distinguish between a secret word that could be considered a potential candidate and one that could not. Thus we defined a new parameter, ν, which reflects the Normalized Maximum Correlation (NMC).

$$\nu = \frac{\max_t R(x, x_h, t)}{E_{x_h}(\max_t R(x, x_h, t))} \tag{5}$$

where $E(\cdot)$ stands for the expected value. In words, it corresponds to the ratio between the Pearson correlation for the correct secret word divided by the average correlation for all other possible secret words. High ν values indicate high vulnerability to power attacks. The application of NMC to an attack on a protected and an unprotected memory module is depicted in Fig. 10. Note that in the unprotected memory the secret word was revealed with a high NMC value, whereas in the protected memory the NMC values for the same word was around 1. This result gives some idea of the ability of an attacker to extract secret information from the memory. The ν values are shown in Fig. 11. The high ν values on the left hand side of the figure (i.e.m values that were obtained with fewer than 100 attack cycles) are associated with measurements that were taken from the unprotected memory. The low ν values on the right hand side (i.e., after thousands of attack cycles) are associated with attacks on the protected module. The results were consistent for all the tested secret words with different Hamming Weights.

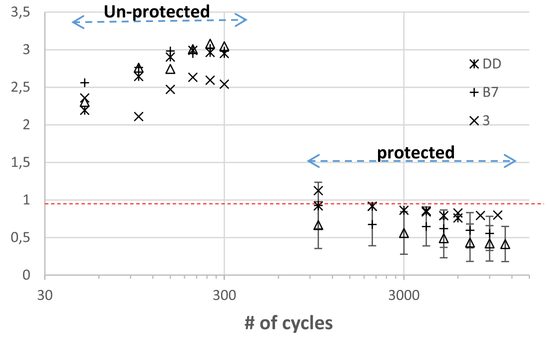

FIGURE 11. Normalized maximum correlation values ν vs. the number of acquired traces for 3 different secret words. The curves on the left side represent the results of a CPA attack on the unprotected module and the curves on the right side present the results of the same attack on the protected module.

We also examined the effect of the physical address on the strength of the correlation. Since we placed the IRU module on the top of the memory bank, we verified whether the interference generated by the IRU was identical for all physical addresses. Fig. 12 gives the ν values at three different addresses located at the bottom of the array (+ symbol), the center (x symbol) and the top of the memory module (x with vertical bar symbol). It is clear from the figure that all the attacks on the addresses exhibited similar NMC values for the protected modules. In the unprotected memory the address located on the top of the array was slightly easier to attack. This measurement confirms that similar signal interference is maintained by the IRU at different physical locations for this memory size.

To explore the effect of the location of the set-bits within an eight bit word, we chose the secret nibble 'D' and

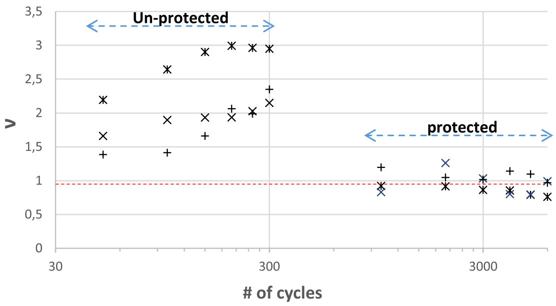

FIGURE 12. Normalized maximum correlation values ν vs. number of acquired traces for several addresses with different physical locations.

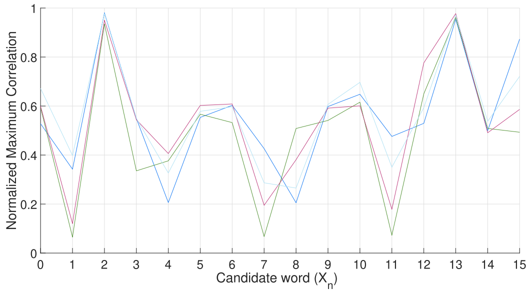

FIGURE 13. Maximum pearson correlation values for different secret word hypotheses for various 4 bit configurations across the byte.

dispersed this bit configuration across the byte. The result is shown in Fig. 13 where we compare the following patterns, 1101**** (blue), ****1101 (green), 01*11*** (light blue) and **1**110 (red) where the bits marked by '∗' were kept fixed during the attack. Clearly there were similar correlation values for all the patterns, where the secret nibble 'D' and its complement '2' had higher correlations than all the other possible secret 4 bit words.

B. THE EFFECTIVENESS OF THE SOLUTION IN TERMS OF MUTUAL INFORMATION

The results shown in Fig. 11 and Fig. 12 indicate that revealing the secret word through straightforward correlation power analysis becomes considerably more difficult for the attacker when it tries to attack a memory protected by an IRU. However, as shown in Section IV, correlation values or correlation ratios do not provide a categorical indication that no information has leaked from the protected module. Moreover, a minor downward trend emerges as the number of attacking cycles increased. This trend in the correlation ratio indicates that some information leakage possibly existed for the protected memory as well. In order to quantify this information leakage and evaluate the level of security of a memory when it is coupled to an IRU, we turned to MI estimations of the protected memory.

The MI results reported in this section were computed according to Eq. 4. First, we evaluated $I(W; Y)$, by testing a representative ensemble of words with all possible Hamming weights. The ensemble was composed of the following nine uniformly distributed words: '01','03','07', '0F', '1F', '3F', '7F' and 'FF'. Thus, the maximum MI that could be achieved in this experiment was $\log_2(9) = 3.16$. Fig. 14 clearly shows that it is possible to extract the Hamming weight of the secret word from the unprotected memory very quickly because the MI attains its maximal value within 1000 traces, whereas the protected memory demonstrates very low information leakage even after 2000 traces.

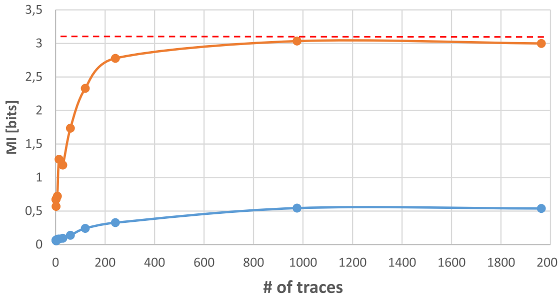

FIGURE 14. Mutual information $I(W; Y)$ for sample word of different HW's vs. the number of traces acquired for unprotected memory (orange circles) and protected memory (blue circles). The red dotted line indicates the maximum achievable MI for this ensemble.

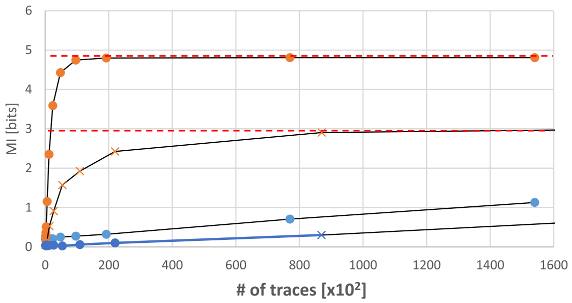

FIGURE 15. Mutual information $I(L; Y|W = w)$ for $w = 1, 2$ vs. the number of traces acquired. The MI values of the unprotected memory are indicated by orange × and ∘ markers, respectively. The MI values of the protected memory are marked in blue. The dotted lines mark the maximal achievable MI. The upper red dotted line indicates the maximum achievable MI.

Next, we addressed the information leakage from the memory given the Hamming weight of the secret words; that is, $I(L; Y|W)$. The results of the MI analysis on the 28 uniformly distributed bytes of Hamming weight $W = 2$ are presented in Fig. 15. Note that in this case $I(L; Y|W)$ is upper bounded by $\log_2(28) = 4.8$ bits. The figure clearly shows that for the unprotected memory the secret information was disclosed after merely 10K traces. On the other hand the protected memory only yielded about one information bit after 140K traces. This result clearly demonstrates the amount of security gain when using a protected module. We repeated this experiment for an ensemble composed of the secret words with HW = 1, as shown in Fig. 15. These results were consistent with the results depicted in Fig. 12 for the sensitivity of the information leakage of the location of the secret bits.

TABLE 2. Comparison between the proposed memory and other secured hardware solutions.

Reference	VLSI'15 [28]	ISSCC'09 [18]	HOST'12 [16]	TVLSI'17 [9]	**This Work [10]**
Hardware Component	Logic	Logic	Memory	Memory	**Memory**
Process	130 nm	65 nm	65 nm	65 nm	**55 nm**
Architecture	AES Engine	AES Engine	SRAM	SRAM	**SRAM**
Protection Level	Periphery	Periphery	Bitcell	Bitcell	**Memory Periphery**
Information Leakage Reduction	Not reported	1667×	Not reported	20×	**50×**
Area Overhead	100%	7.2%	40%	42%	**1%**
Latency Overhead	None	2×	2.1×	1.4×	**<1.05×**
Power Overhead	30%	1.33×	2.56×	1.24×	**<1.05×**

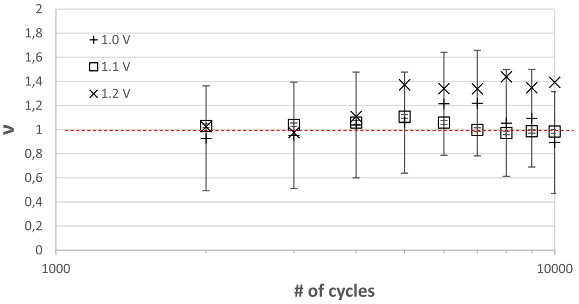

FIGURE 16. Normalized maximum correlation values ν vs. number of acquired traces for different applied voltages under CPA attack on a protected module. (+) 1.1V, (X) 1.2 V and (square) 1.0 V.

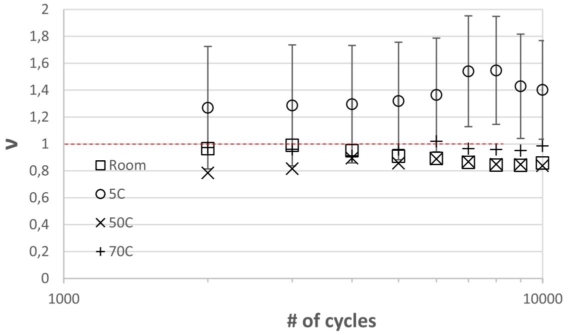

FIGURE 17. Normalized maximum correlation values ν vs. number of acquired traces for different temperatures under CPA attack on a protected module. (+) Room, (X) 50C, (square) 70C and (circle) 5C.

C. SENSITIVITY TO ENVIRONMENTAL CONDITIONS AND PROCESS VARIATIONS

We carried out experiments to evaluate the sensitivity of the IRU countermeasure under various temperatures and applied voltage biasing. The NMC results for the protected module under various environmental conditions are shown in Figs. 16 and 17. As shown, the NMC values were mostly around 1, even for 10K acquired traces, which indicates no significant signature of the secret word relative to other possible words. Under certain environmental conditions such as low temperature and high voltage, an increase in the NMC might indicate that some information had been disclosed. However, as can be seen from the large distribution around the measured value, this relatively higher NMC not of much value to the attacker since it did not have the benefit of prior knowledge of the secret word. As depicted in Fig. 5, the circuit simulations clearly prove that the IRU can cover environmental and process variations. Fig. 18 depicts the effect of process variations (die-to-die) on the correlation values ν as a function of the number of traces. The analysis was conducted on repeated measurements under the same conditions for ten different dies.

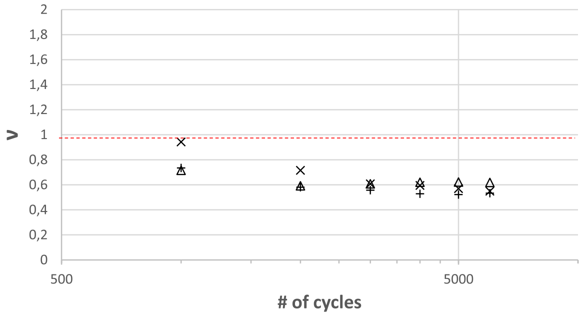

FIGURE 18. Normalized maximum correlation values ν vs. number of acquired traces for different devices for CPA attack on a protected module.

D. IMPLEMENTATION COST AND COMPARISON

Table 2 compares the proposed secured SRAM to an integrated IRU and other security measures applied to digital circuits and memories in the literature. The proposed solution resulted in only 1% area overhead, compared to 7.2%–100% overheads for secured digital circuits implementing AES engines, and over 40% area overhead in other secured SRAM solutions. Moreover, our solution is the only secured memory implementation that does not require any changes to the foundry-provided "pushed-rule" SRAM bitcell design. In addition, our solution incurs less than 1.05× latency and power overheads, unlike other secured memory solutions, which result in 1.4×–2.1× and 1.24×–2.56× latency and power overheads, respectively.

E. THE COST OF LOCALITY

We analyzed the cost of the protection resolution in Bytes vs. the projected area overhead for our proposed mechanisms. As shown in Table 3, for an 8K bit memory our

TABLE 3. The area overhead (oh) in % for different block-sizes versus the granularity of the IRU embedding (for smaller blocks of 'embedding-resolution' size in bits).

		Embedding resolution (bits)					
	Area oh [%]	2^7	2^9	2k	8k	32k	128k
Block Size	8k bit	8%	4%	2%	1%		
	32k bit	16%	8%	4%	2%	1%	
	128k bit	32%	16%	8%	4%	2%	1%

solution randomizes each memory column independently and therefore is already quite *local*. However, to chunk up each row to {2,4,8} sections and add independent randomizers for them, the area overhead would result in only {2,4,8}% respectively, which is an ultra-low overhead cost. The table also shows this data for another 32K bit array we have on chip and are projected for a 128K bit array as well.

VII. CONCLUSION

In this work we quantitatively evaluated the security of SRAM macros against SCA attacks and discussed our methodology. We presented and tested an ultra-low-overhead secured SRAM design. As described we implemented an impedance randomization unit (IRU) which was added to the periphery of a conventional SRAM macro, implemented with foundry-based 6T bitcells. This eliminates the high cost of a bit-cell level intervention or an addition of algorithmic level solutions prior to storing sensitive information in the memory (encryption or data masking). Concretely, we achieved a 1% area overhead design (for an 8k-bit block) with no latency or energy degradation. The level of security as evaluated on our adapted information theoretic metric to the memory context, exhibited high security in terms of the area overhead ratio. As compared to current methodologies we achieve 40× less area overhead for x2.5 less information leakage. We discuss the parametric nature of the solution in terms of design parameters for security engineers. We evaluate the cost associated with *localizing* our technique to make it harder for stronger and better localized adversaries.

REFERENCES

[1] P. Kocher, J. Jaffe, B. Jun, and P. Rohatgi, "Introduction to differential power analysis," *J. Cryptograph. Eng.*, vol. 1, no. 1, pp. 5–27, Apr. 2011.

[2] S. Mangard and A. Y. Poschmann, *Constructive Side-Channel Analysis and Secure Design*. Cham, Switzerland: Springer, 2015.

[3] M. Alioto, L. Giancane, G. Scotti, and A. Trifiletti, "Leakage power analysis attacks: A novel class of attacks to nanometer cryptographic circuits," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 2, pp. 355–367, Feb. 2010.

[4] M. Alioto, S. Bongiovanni, M. Djukanovic, G. Scotti, and A. Trifiletti, "Effectiveness of leakage power analysis attacks on DPA-resistant logic styles under process variations," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 61, no. 2, pp. 429–442, Feb. 2014.

[5] I. Levi, O. Keren, and A. Fish, "Data-dependent delays as a barrier against power attacks," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 8, pp. 2069–2078, Aug. 2015.

[6] I. Levi, D. Bellizia, and F.-X. Standaert, "Reducing a masked implementation's effective security order with setup manipulations," *IACR Trans. Cryptograph. Hardw. Embedded Syst.*, vol. 2019, pp. 293–317, Feb. 2019.

[7] M. Avital, H. Dagan, I. Levi, O. Keren, and A. Fish, "DPA-secured quasi-adiabatic logic (SQAL) for low-power passive RFID tags employing S-boxes," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 1, pp. 149–156, Jan. 2015.

[8] R. Giterman, M. Vicentowski, I. Levi, Y. Weizman, O. Keren, and A. Fish, "Leakage power attack-resilient symmetrical 8T SRAM cell," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 26, no. 10, pp. 2180–2184, Oct. 2018.

[9] R. Giterman, O. Keren, and A. Fish, "A 7T security oriented SRAM bitcell," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 66, no. 8, pp. 1396–1400, Aug. 2019.

[10] R. Giterman, M. Wicentowski, O. Chertkow, I. Sever, I. Kehati, Y. Weizman, O. Keren, and A. Fish, "Power analysis resilient SRAM design implemented with a 1% area overhead impedance randomization unit for security applications," in *Proc. IEEE 45th Eur. Solid State Circuits Conf. (ESSCIRC)*, Sep. 2019, pp. 69–72.

[11] R. Giterman, I. Levi, Y. Weizman, O. Keren, A. Fish, and M. Wicentowski, "Secured memory," U.S. Patent 20 200 372 186 A1, Nov. 26, 2020.

[12] S. Gao, B. Marshall, D. Page, and E. Oswald, "Share-slicing: Friend or foe?" *IACR Trans. Cryptograph. Hardw. Embedded Syst.*, vol. 2019, pp. 152–174, Nov. 2019.

[13] ITRS. (2015). *International Technology Roadmap for Semiconductors—2015 Edition*. [Online]. Available: http://www.itrs2.net

[14] M. Neve, E. Peeters, D. Samyde, and J.-J. Quisquater, "Memories: A survey of their secure uses in smart cards," in *Proc. 2nd IEEE Int. Secur. Storage Workshop (SISW)*, Oct. 2003, p. 62.

[15] W. Liu, R. Luo, and H. Yang, "Cryptography overhead evaluation and analysis for wireless sensor networks," in *Proc. WRI Int. Conf. Commun. Mobile Comput. (CMC)*, vol. 3, Jan. 2009, pp. 496–501.

[16] V. Rožić, W. Dehaene, and I. Verbauwhede, "Design solutions for securing SRAM cell against power analysis," in *Proc. IEEE Int. Symp. Hardw.-Oriented Secur. Trust*, Jun. 2012, pp. 122–127.

[17] E. Konur, Y. Ozelci, E. Arikan, and U. Eksi, "Power analysis resistant SRAM," in *Proc. World Automat. Congr.*, 2006, pp. 1–6.

[18] C. Tokunaga and D. Blaauw, "Secure AES engine with a local switched-capacitor current equalizer," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2009, pp. 64–65.

[19] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 127–138, Jan. 2017.

[20] P. N. Whatmough, S. K. Lee, H. Lee, S. Rama, D. Brooks, and G.-Y. Wei, "A 28 nm SoC with a 1.2 GHz 568 nJ/prediction sparse deep-neural-network engine with >0.1 timing error rate tolerance for IoT applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 242–243.

[21] O. Hirabayashi, O. Hirabayashi, A. Kawasumi, A. Suzuki, Y. Takeyama, K. Kushida, T. Sasaki, A. Katayama, G. Fukano, Y. Fujimura, T. Nakazato, Y. Shizuki, N. Kushiyama, and T. Yabe, "A process-variation-tolerant dual-power-supply SRAM with 0.179 μm^2 cell in 40 nm CMOS using level-programmable wordline driver," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2009, pp. 458–459.

[22] S. Khandelwal, V. Gupta, B. Raj, and R. D. Gupta, "Process variability aware low leakage reliable nano scale double-gate-FinFET SRAM cell design technique," *J. Nanoelectron. Optoelectron.*, vol. 10, no. 6, pp. 810–817, Dec. 2015.

[23] V. Gupta, S. Khandelwal, B. Raj, and R. D. Gupta, "Leakage current reduction in finfet based 6T SRAM cell for minimizing power dissipation in nanoscale memories," in *Proc. 5th Nirma Univ. Int. Conf. Eng. (NUiCONE)*, Nov. 2015, pp. 1–5.

[24] V. Gupta, S. Khandelwal, J. Mathew, and M. Ottavi, "45 nm bit-interleaving differential 10T low leakage FinFET based SRAM with column-wise write access control," in *Proc. IEEE Int. Symp. Defect Fault Tolerance VLSI Nanotechnol. Syst. (DFT)*, Oct. 2018, pp. 1–6.

[25] L. Itamar, O. Keren, and A. Fish, "Pseudo-asynchronous digital circuit design," U.S. Patent 10 572 619, Feb. 25, 2020.

[26] I. Levi, D. Bellizia, D. Bol, and F.-X. Standaert, "Ask less, get more: Side-channel signal hiding, revisited," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 12, pp. 4904–4917, Dec. 2020.

[27] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1991, pp. 198–203.

[28] S. Lu, Z. Zhang, and M. Papaefthymiou, "1.32 GHz high-throughput charge-recovery AES core with resistance to DPA attacks," in *Proc. Symp. VLSI Circuits (VLSI Circuits)*, Jun. 2015, pp. C246–C247.

YOAV WEIZMAN received the Ph.D. degree in physics from the Ben-Gurion University of the Negev. He has 15 years of experience in basic and applied research, development, and design. In 2000, he joined Freescale Semiconductor, Herzelia, where he was involved in the development of tools and techniques for IC diagnostics and later became the Product Analysis and Characterization Manager leading numerous research activities in failure analysis, signal integrity, and special IC diagnostics structures for yield enhancement and process tuning. He presented a unique and innovative approach to the implementation of special CMOS structures for failure analysis tool development and debug facilitation. Recently, he joined Bar-Ilan University as a Research Fellow, where he is implementing failure analysis and reliability methodologies and techniques for security tempering and eavesdropping of integrated circuits. He has published several articles on IR emission microscopy and novel inspection techniques for failure analysis.

ROBERT GITERMAN received the B.Sc. degree in electrical engineering and the M.Sc. degree from the Ben-Gurion University of the Negev, Be'er Sheva, Israel, in 2013 and 2014, respectively, as part of a Fast Track Program for outstanding students, and the Ph.D. degree from the Emerging Nanoscaled Intergrated Circuits and Systems (EnICS) Labs, Bar-Ilan University, Ramat Gan, Israel, under the supervision of Prof. Alex Fish and Dr. Adam Teman. He is currently a Postdoctoral Researcher at the Telecommunications Circuits Laboratory (TCL), École Polytechnique Fédérale de Lausanne (EPFL). As part of his research, he led several full test chip integrations and tape out. He has authored/coauthored over 30 journal articles and international conference papers and eight patent applications. He has presented his research at a number of international conferences. He is also the coauthor of the book *Gain-Cell Embedded DRAMs for Low-Power VLSI Systems-on-Chip*. His research interests include embedded DRAM design and optimization for low power and high performance operation, SRAM design with an emphasis on improved stability, error-correction and fault-tolerant circuits, and the development of hardware-security oriented embedded memories for use in low-power applications and high-end processors.

ORON CHERTKOW received the degree in electrical and computer engineering from the Ben-Gurion University of the Negev. He served with Unit 8200 as an Intelligence and Communications Researcher. During his B.Sc. degree, he invented a novel low-power and a radiation-hardened SRAM cell for which he was awarded the university's ''Best Final Project'' Award and the Best Paper Award at IEEE S3S conference, in 2014. He subsequently patented his invention and wrote several journal articles about it and sold it to Dolphin Integration, where he went to work after he completed his degree and led the research and development of a hardware-secured memory compiler. Following his passion to cross-disciplinary technologic view and its business aspects, he joined Amdocs at 2018 to co-found the corporate's venture capital investment arm, mandated at exploring solutions in areas beyond Amdocs' horizons in domains that can impact the trajectory of where the telecom and the media industry is heading and help startup companies navigate these markets through access to Amdocs' customers, products, and technology.

MAOZ WICENTOWSKI (Student Member, IEEE) was born in Bat Yam, Israel, in 1988. He received the B.Sc. degree in electrical engineering from Bar-Ilan University, Ramat Gan, in 2017. From 2016 to 2018, he was a Researcher at the Emerging Nanoscaled Integrated Circuits and Systems (EnICS) Labs, Faculty of Engineering, Bar-Ilan University. He is currently with Samsung, working on the development of cutting edge image sensors. His research interests include leakage power analysis (LPA) and side-channel attacks (SCA).

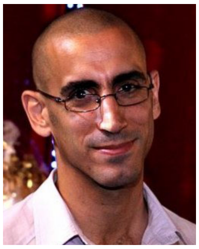

ITAMAR LEVI received the B.Sc. and M.Sc. degrees in electrical and computer engineering from the Ben-Gurion University of the Negev, in 2012 and 2013, respectively, as a part of a Direct Excellence Student Track, and the Ph.D. degree from Bar-Ilan University, Ramat Gan, Israel, in 2017. He was a Research Associate with the UCLouvains Crypto-Group, UCLouvain, Belgium, until 2019. He is currently a Computer-Engineering Faculty Member at Bar-Ilan University. He is also a member of the Emerging Nanoscale Circuits and Systems (EnICS) Labs, BIU. As part of his research activities, he has authored/coauthored over 50 journal articles and international conference papers and seven patent applications, he is the coauthor of the book *Dual-Mode-Logic: A New Paradigm for Digital IC Design* and serves in several technical committees of the IEEE Circuits and Systems Society and the Hardware Security conferences. His current research interests include digital circuit design, embedded systems security, security evaluation analysis for cryptographic devices, and side-channel and fault-injection countermeasures and cryptographic implementations.

ILAN SEVER received the B.Sc.E.E. degree from the Technion—Israel Institute of Technology. He has over 25 years of design and project-management experience in the field of semiconductor IP and SOC design with deep expertise in volatile and non-volatile memory design. He currently serves as the Vice President of research and development at Weebit-Nano, an Israeli startup in the field of resistive memory (ReRAM). Prior to joining Weebit Nano, he spent 11 years at French Semiconductors Company Dolphin Design, leading the Israeli subsidiary in developing next generation SRAM memory architectures. In his previous roles, he was the VLSI Area Manager at Sandlinks Systems, the IoT startup; the Director of IP and Libraries at Tower Semiconductors; and the Design Manager of Flash-Memory at ST Microelectronics. He has several granted patents and awards.

ISHAI KEHATI was born in Netanya, Israel. He received the B.Sc. degree in computer engineering from Bar-Ilan University, Israel, in 2017. In 2017, he joined the Enics Group, Faculty of Engineering, Bar-Ilan University, as a FPGA and Laboratory Engineer, where he was responsible of the engineering aspects of the IC security laboratory. In this role, he was part of several high end research activities that required complicated bring up of complicated automated measurement setups and development of advanced analysis tools that allowed big data management and parameters extraction. Since 2021, he has been working with Ethernity Networks, as a FPGA Engineer.

OSNAT KEREN (Member, IEEE) received the M.Sc. degree in electrical engineering from the Technion—Israel Institute of Technology and the Ph.D. degree from Tel Aviv University, Israel. After working at High Tech for several years, she took up a faculty position at the Faculty of Engineering, Bar-Ilan University, Israel.

ALEXANDER FISH (Member, IEEE) received the B.Sc. degree in electrical engineering from the Technion—Israel Institute of Technology, Haifa, Israel, in 1999, and the M.Sc. and the Ph.D. *(summa cum laude)* degrees from the Ben-Gurion University of the Negev (BGU), Be'er Sheva, Israel, in 2002 and 2006, respectively.

He was a Postdoctoral Fellow with the ATIPS Laboratory, University of Calgary, Calgary, AB, Canada, from 2006 to 2008. In 2008, he joined the Electrical and Computer Engineering Department, BGU, as a Faculty Member, where he founded the Low Power Circuits and Systems Laboratory, specializing in low-power circuits and systems. In 2011, he was appointed as the Head of the VLSI Systems Center, BGU. In 2012, he joined the Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel, as an Associate Professor, and the Head of the Nanoelectronics Track. He is currently the Founder and a Leading Member of the Emerging Nanoscaled Integrated Circuits and Systems Labs, BGU. He has authored over 100 scientific articles in journals and conferences and two book chapters. His current research interests include the development of secured hardware, ultralow power embedded memory arrays, CMOS image sensors, and high speed and energy efficient design techniques.

Dr. Fish coauthored articles that won the best paper finalist awards at the IEEE International Symposium on Circuits and Systems (ISCAS) and ICECS conferences. He is a member of Sensory, VLSI Systems and Applications, and Bio-Medical Systems Technical Committees of the IEEE Circuits and Systems Society. He also served as the chair for different tracks of various IEEE conferences. He was a Co-Organizer of many special sessions at the IEEE conferences, including the IEEE ISCAS, the IEEE Sensors, and the IEEE conferences. He serves as the Editor-in-Chief for the *Journal of Low Power Electronics and Applications* (Multidisciplinary Digital Publishing Institute) and as an Associate Editor for the IEEE SENSORS JOURNAL, IEEE ACCESS, *Microelectronics* (Elseiver), *Integration* (Elseiver), and the VLSI Journals.