Testing the Resilience of Cryptographic Modules Against Fine-Grained Time- and Frequency-Domain EM Side-Channel Analysis Attacks

Vishnuvardhan Iyer
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX, USA
vishnuv.iyer@utexas.edu

Aditya Thimmaiah
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX, USA
auditt@utexas.edu

Ali Yilmaz
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX, USA
ayilmaz@mail.utexas.edu

***Abstract*—The ability of various countermeasures to secure cryptographic modules implementing the Advanced Encryption Standard (AES) algorithm is experimentally evaluated using fine-grained time- and frequency-domain electromagnetic side-channel analysis (EM SCA) attacks. Because an infeasibly large number of measurements of on-chip EM emanations should be required to break the cryptographic protection of a secure cryptosystem, a novel approach is used to keep acquisition costs low while testing the system resilience to EM SCA attacks. An adaptive scan protocol is used first to rapidly isolate optimal measurement configurations on an unsecured implementation. Then, the effectiveness of these optimal configurations are evaluated in the presence of countermeasures. The methodology is used to disclose the key from unsecured and secured FPGA implementations of AES-128. The most secure countermeasures are found to be >25× more resilient against the time-domain EM SCA attack and >20× more resilient against the frequency-domain EM SCA attack.**

***Keywords*—electromagnetic measurements, side-channel attacks, electromagnetic interference, measurement techniques**

I. Introduction

Near-field measurements of data-dependent signals unintentionally emanated during chip operations can be exploited to compromise information security of embedded cryptosystems [1]-[4]. Such electromagnetic side-channel analysis (EM SCA) attacks recover critical information about system behavior by statistically linking measured fields to observable inputs/outputs of a device under test. For example, secret keys used for encryption in the Advanced Encryption Standard (AES) algorithm [5] were recovered in [3] by observing ciphertexts and correlating the bit-transitions in the state registers during the final round of AES operations to near-field signals observed in the same time interval.

Numerous countermeasures have been proposed to mitigate EM SCA attacks on cryptosystems [6]-[9]. These typically reduce the information content in captured signals by introducing algorithmic noise [6] or measurement noise [7]-[9] during encryptions. Some countermeasures introducing measurement noise in signals are based on techniques used to minimize EM interference (EMI) and improve power management in embedded systems; e.g., countermeasures involving random jitter introduction [7], voltage-level randomization [8], or a combination of both [9] have been shown to decrease the effectiveness of EM SCA attacks with relatively low design overheads. These countermeasures aim to reduce the repeatability of measurements, i.e., by randomly varying supply voltages or clock frequency for each encryption, they attempt to reduce the correlation between observed fields and the critical information/data being processed on the chip.

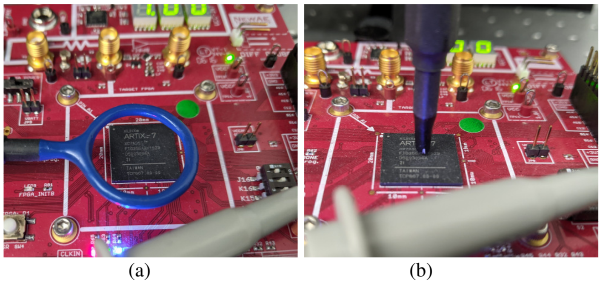

Fig 1. (a) Coarse-grained EM SCA attack performed using a 25-mm diameter H-field probe in *z*-orientation. Larger probes average fields spatially, aggregating noise from uncorrelated sources which degrade the attack. (b) Fine-grained EM SCA attack performed using a 1-mm diameter H-field probe in *y*-orientation. Smaller probes scan the chip surface using high-precision equipment to localize vulnerabilities.

Countermeasures for securing embedded cryptosytems against EM SCA attacks have been evaluated primarily by using coarse-grained attack setups (Fig. 1(a)), where a relatively large probe aggregates fields from sources distributed across a chip (and beyond), typically in time domain [1], [9]. On the one hand, such aggregation adds noise from several uncorrelated sources to the data-dependent emanations from critical compute-blocks thereby obfuscating the analysis and decreasing the attack's effectiveness. Additionally, the higher noise-floor of larger probes [4] further reduces the possibility of mounting a successful attack against secured designs. On the other hand, coarse-grained setups are generally cheap to implement and require few measurements because attacks using such setups have limited sensitivity to probe position. In contrast, fine-grained attack setups (Fig. 1(b)) involve high-resolution scans on the chip using smaller probes in multiple orientations [2]-[4]; they can identify more potent measurement configurations by spatially localizing information leakage. Implementation of such fine-grained attacks requires the use of relatively expensive high-precision equipment. Although the configurations identified by fine-grained attacks can be re-used to reduce the *marginal cost* of future attacks on identical chips, the cost of the initial experiments to identify these configurations (henceforth referred to as the *acquisition cost*) can be prohibitively high [3]: The multi-dimensional search space includes probe location, height, and orientation as well as the number of observed encryptions. Indeed, counter-measures that are able to increase the marginal cost of fine-grained EM SCA attacks generally also increase their acquisition cost, potentially making experimental testing infeasible for secured cryptographic modules.

In this article, a fine-grained EM SCA attack setup was used to evaluate the resilience of AES modules secured using EMI reduction techniques. Three methods were used to avoid high acquisition costs: (I) The adaptive scan protocol described in [3], which implements a greedy-search algorithm to progressively constrain the search space over multiple scans, was adopted to identify the optimal measurement configurations. (II) The protocol was generalized to use either time- or frequency-domain signals [10]. (III) The adaptive scan protocol was used to first disclose the key of an *unsecured* open-source FPGA implementation of the 128-bit AES (AES-128) [11]. Then, the optimal configurations identified for this baseline design were used in presence of three countermeasures—frequency-scaling, voltage-scaling, and voltage-frequency-scaling—to evaluate their resilience against time- and frequency-domain EM SCA attacks. The improvement in the security of the design was quantified by measuring the increase in the marginal cost of future attacks.

II. EM SCA Attacks on Secure AES Modules

A. Correlation-Based EM SCA Attacks

EM SCA attacks on AES correlate fields observed during several encryptions to hypothetical leakage models—here, the Hamming-Distance model—that categorize the bit-transitions from guessed intermediate values to observed outputs. Each of the AES key-bytes can be extracted separately as operations in the final round of AES are performed on each byte independently. Assuming the output ciphertext is available to an attacker, first the value of each target key-byte is guessed as $0 \leq g \leq 255$ (the one-to-one byte-mapping in the AES S-box ensures that the target byte of a ciphertext is mapped to a unique guessed penultimate round value, for each key guess [12]). Then, for each observed encryption e, target key-byte b, and guess key value g, the Hamming distance $HD_e^{b,g}$ is computed between the output ciphertext and the guessed penultimate round value. For the same encryption, the observed signal $V_e^{pc,t}$ (typically, the voltage detected by a probe) is recorded using different probe configurations pc—by modifying the transverse probe location l, height h, or orientation o—at different time instances t. Next, the signals and Hamming distances for all observed encryptions are listed in arrays $\mathbf{V}^{pc,t}$ and $\mathbf{H}^{b,g}$ respectively, and correlated [4]:

$$\rho_{V,H}^{b,g,pc,t} = \frac{\mathrm{Cov}(\mathbf{V}^{pc,t},\mathbf{H}^{b,g})}{\sqrt{\mathrm{Var}(\mathbf{V}^{pc,t})\times\mathrm{Var}(\mathbf{H}^{b,g})}} \tag{1}$$

The correlation coefficient ρ for the correctly guessed key value g^* is expected to be higher than those for the alternative guesses, although it may require several encryptions before the key can be recovered with high confidence. To this end, the maximum value of ρ is computed over all observed time instants or over all frequency components (see Section II.C); this maximum is compared to a null hypothesis threshold derived from inverse t-distributions for a confidence ratio of 99.99% [12]. As the number of observed encryptions n_e increases, g^* yields the maximum ρ that crosses the threshold, while other guesses yield smaller correlation coefficients (Fig. 2). The number of measurements required for the maximum ρ to cross the threshold is defined as the measurements to disclosure (MTD), which depends on the key-byte b and the particular probe configuration pc, represented as MTD_b^{pc}.

B. Cost of experiments

Each byte of the secret key may be recovered using a different probe configuration, which depends on the chip layout, especially the locations of S-boxes and the state registers [6],[8]. To rapidly identify optimal configurations for recovering the key-bytes, the adaptive scan protocol detailed in [3],[12] and summarized next is implemented.

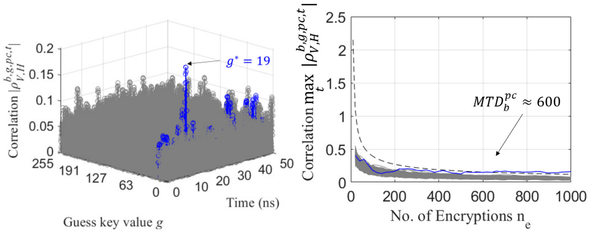

(a) (b)

Fig. 2. (a) The correlation coefficient computed across the final clock-cycle of AES, after 1000 encryptions, for 256 guess key values (grey) of byte 1. A 1-mm diameter H-field probe at an optimal configuration l_1^{opt} = (9.7,8) mm, h_0 = 0.5 mm, $o_1^{\mathrm{opt}} = x$ above the Artix-7 FPGA was used [3]. The correct guess g^*=19 (blue) had the maximum correlation at t~12 ns. (b) The maximum correlation coefficients as the number of encryptions increases. The coefficient for the correct guess crossed the null hypothesis threshold (dashed) after $MTD_1^{pc} \approx 600$ encryptions were observed.

The protocol implements a greedy-search algorithm in two phases—phase I and phase II—with multiple scans, where the probe configuration in each scan s ($pc_{s,\mathrm{I/II}}$) depends on the area and resolution of the scan. Phase I performs $N_{\mathrm{scan}}^{\mathrm{I}}$ progressively more expensive low-resolution scans over the entire chip-area by increasing the number of measurements or scan resolution. These full-chip scans identify seeds for optimization in phase II: Given the probe height h_0, the optimal probe location $l_b^{\mathrm{opt},0}$ and orientation $o_b^{\mathrm{opt},0}$ as well as the minimum MTD

$$mMTD_b^0 = \min_l \min_o MTD_b^{pc_{N_{\mathrm{scan}}^{\mathrm{I}},\mathrm{I}}}|_{h=h_0} \tag{2}$$

are identified for each key-byte b at the end of phase I. Phase II, performed byte-wise, starts from the configurations identified in phase I and optimizes them by progressively localizing the search around $l_b^{\mathrm{opt},0}$, keeping the orientation $o_b^{\mathrm{opt},0}$. Phase II scans get progressively cheaper as the number of measurements in each scan s is limited to the minimum MTD identified in the previous scan $s-1$:

$$mMTD_b^{s-1} = \min_l MTD_b^{pc_{s-1,\mathrm{II}}}|_{h=h_0,o=o_b^{\mathrm{opt},0}} \tag{3}$$

The total number of measurements accrued to identify the optimal configurations (l_b^{opt}, h_0, o_b^{opt}) at the end of phase II is the acquisition cost. If $N_{\mathrm{obs}}^{s,\mathrm{I/II}}$ represents the number of observed probe locations for scan $s = 1, \cdots, N_{\mathrm{scan}}^{\mathrm{I/II}}$, then [3]

$$\mathit{Acquisition\ Cost} = \sum_{s=1}^{N_{\mathrm{scan}}^{\mathrm{I}}} 2N_{\mathrm{e}}^{s,\mathrm{I}} N_{\mathrm{obs}}^{s,\mathrm{I}} + \sum_{b=1}^{16} \sum_{s=1}^{N_{\mathrm{scan}}^{\mathrm{II}}} mMTD_b^{s-1} N_{\mathrm{obs}}^{s,\mathrm{II}}, \tag{4}$$

where $N_{\mathrm{e}}^{s,\mathrm{I}}$ is the number of encryptions measured per location in scan s of phase I. The first term, representing phase I cost, includes a factor of 2 because the experimental setup used in this article allows scanning in x- and y-orientations. Once the acquisition cost is incurred and optimal configurations are identified, the number of measurements needed to break the cryptography on an identical chip and recover all key-bytes using these configurations is:

$$\mathit{Marginal\ Cost} = \sum_{b=1}^{16} mMTD_b^{N_{\mathrm{scan}}^{\mathrm{II}}} \tag{5}$$

In this article, the marginal cost is used to evaluate the security of an AES implementation with and without countermeasures.

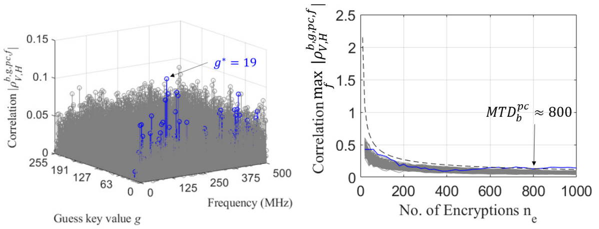

(a) (b)

Fig. 3. (a) Correlation for 256 guess key values (grey) after 1000 encryptions using frequency-domain signals for the same byte and optimal configuration used in Fig. 2. The correct guess $g^* = 19$ (blue) had maximum correlation at ~150 MHz. (b) The maximum correlation coefficients as the number of encryptions increases. The coefficient for the correct guess key crossed the threshold (dashed) after $MTD_1^{pc} \approx 800$ encryptions.

C. Time- vs Frequency-Domain EM SCA Attacks

The adaptive scan protocol can be implemented using frequency-domain signals by replacing time t by frequency f in (1) (Fig. 3). Since information leakage is constrained to the final round of operations in AES, the measurements are typically time-limited to N_t time samples corresponding to one clock period. The frequency-domain signals can be generated via Fast Fourier Transform (FFT) of the time-domain signals. These can be padded with zeros to improve the observed spectrum's frequency resolution and reduce processing time for the FFTs, which require $O(N_\mathrm{e} N_\mathrm{t} \log N_\mathrm{t})$ seconds. Further, the frequency-domain signals are band-limited to N_BW samples, corresponding to the minimum bandwidth among all measuring equipment.

The analysis time for the adaptive scan protocol in time-domain accounts for computation of correlations and finding the MTD for all acquisitions. For a given probe configuration, the cost of computing correlations and finding the minimum MTD can be reduced to $O(N_\mathrm{t} N_\mathrm{e} \log N_\mathrm{e})$ seconds using a binary search method [12], where N_e is the maximum number of encryptions observed. In comparison, the analysis time for the frequency-domain approach must also account for the computation of FFTs. Typically, the minimum MTD search, which requires $O(N_\mathrm{BW} N_\mathrm{e} \log N_\mathrm{e})$ seconds, is more expensive than FFT computations, especially for large values of N_e. Further, if sensors and signal capture devices have limited bandwidth, such that $N_\mathrm{BW} \ll N_\mathrm{t}$, and the acquisition cost for both approaches are similar, the frequency-domain approach provides a computationally viable alternative to the time-domain attack. Because time-domain shifts do not change frequency-domain signal amplitudes, the frequency-domain attacks should be more effective against countermeasures that introduce delays to stagger AES operations in time.

D. Countermeasures

Three countermeasures based on techniques to reduce EMI and improve power management in embedded systems are analyzed. Implementations using these countermeasures randomize the supply voltage or clock signal using pseudo-random sequence generators, such as linear-feedback shift registers, that have relatively low design overheads. They are easily integrated with the chip and can serve purposes other than improving side-channel security [9]. The following countermeasures all aim to worsen the repeatability of measurements, i.e., the randomization implies that repeating the same encryption will result in different observed fields, degrading correlation analysis [4]:

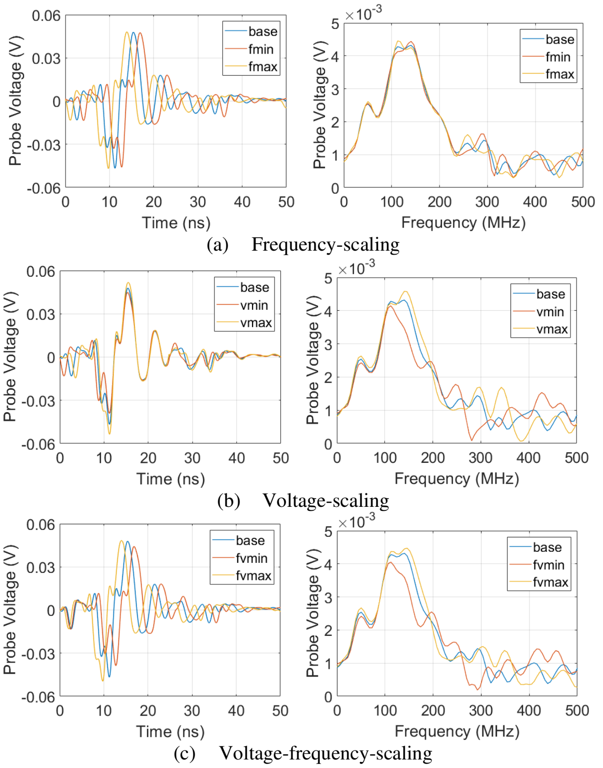

(a) Frequency-scaling

(b) Voltage-scaling

(c) Voltage-frequency-scaling

Fig. 4. Impact of countermeasures on the measured signal. Measurements were for one encryption using the optimal configuration in Figs. 2-3. The baseline operations were performed using $V^\mathrm{s} = 1$ V and $f^\mathrm{clk} = 20$ MHz. (a) The signal shifted ~ ±1.3 ns from the baseline in time domain as the clock frequency increased from 19.5 MHz to 20.5 MHz, compressing the clock period. Frequency-domain signals were less affected by jitter. (b) The signal increased proportionally at the peaks and changed minimally at other time instants as supply voltage increased from 0.9 V to 1.1 V. Frequency-domain signals were affected significantly for $f > 100$ MHz. (c) The signal was staggered in both time and frequency domain when the two counteremeasures were combined.

1) Frequency Scaling: Spread-spectrum clocking reduces signal levels at a device's clock frequency f^clk and its harmonics by modulating the input clock with a low-frequency signal. This spreads the bandwidth of the original signal's spectrum and introduces jitter in time-domain signals. Random delays in the observed fields, varying from encryption to encryption, diminish correlations at every time instant [7]. This countermeasure can be implemented with minimal area overhead but larger delay overheads associated with jitter. Jitter dithers time-domain signals but has minimal effect on frequency-domain signals as shown in Fig. 4(a).

2) Voltage Scaling: Minor reduction in supply voltage V^s reduces observed signal amplitude and can reduce EMI. Further, voltage-scaling techniques can manage a system's energy grids by assigning appropriate voltage levels for various tasks [8], [9]. Random voltage scaling can effectively mask side-channels by making peak-to-peak signal amplitudes less sensitive to encryptions, with smaller delay overheads. This countermeasure causes large variation in fields, especially in frequency domain, as shown in Fig. 4(b).

3) Voltage-Frequency Scaling: This countermeasure is a combination of the previous two countermeasures. It is

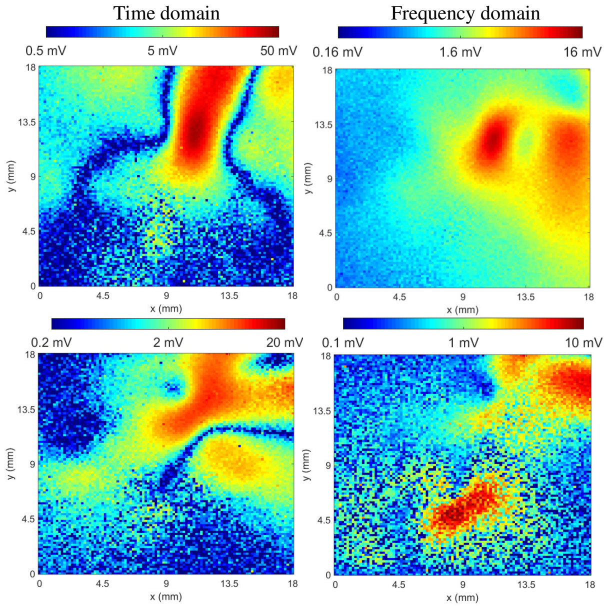

Fig. 5. Spatial maps of the measured voltage plotted for *x*-oriented (top) and *y*-oriented probe (bottom). The signals showed strong dependence on probe position and orientation in both time and frequency domain.

typically implemented by selecting among several predefined pairs of voltage and frequency [9]. Randomizing voltage-frequency scaling is expected to be more effective than the previous countermeasures, at the expense of overheads associated with each of the two countermeasures. The impact of this countermeasure on the time- and frequency-domain signals is shown in Fig. 4(c).

III. Measurement Results

A. Setup

An open-source implementation of AES-128 [11] was realized on a Xilinx Artix-7 FPGA. Inputs to the AES module were read from a file consisting of several randomly generated plaintext and the same set of inputs are repeated for measurements at every probe configuration. This baseline design was operated at a supply voltage level of $V^{\text{base}} = 1$ V and clock frequency f^{base} =20 MHz (clock period 50 ns). The total area of the chip is 18 mm × 18 mm. A 1-mm diameter H-field probe, fixed at height $h_0 = 0.5$ mm above the chip, was used to sense emanated fields within the chip area, in *x*- and *y*-orientation, positioned using Riscure's EM probe station. Signals were amplified by 30 dB and captured using a Keysight DSOS054A oscilloscope at a sampling rate of 10 GS/s [3]. The setup's bandwidth was limited by the 500 MHz oscilloscope. The test board CW305 allows software-controlled voltage supply and input clock levels to be set for the main core. This allows randomized frequency-, voltage- and voltage-frequency-scaling countermeasures to be implemented as part of the main test script, where these inputs can be set before a plaintext is sent to the module. To reduce any bottlenecks created by introducing additional operations in the measurement, the supply voltage and/or clock signal were modified once every 20 encryptions to implement the countermeasures. This should be sufficient to reduce correlations because MTDs for key-bytes are typically more than 600 measurements. Further, the countermeasures were also tested for different voltage and/or frequency ranges, since larger ranges typically translate to additional overheads in designs. Random values were chosen among 11 sample points, uniformly distributed within the target ranges, for the frequency- and voltage-scaling countermeasure, while the voltage-frequency-scaling countermeasure was implemented by selecting among 5 pre-defined voltage-frequency pairs.

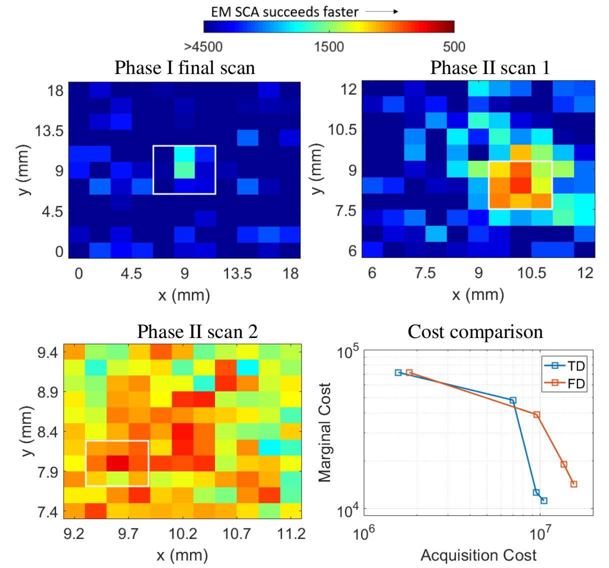

Fig. 6. The results of the adaptive scan protocol in frequency domain. MTD maps are shown for byte 1, where the $l_1^{\text{opt},0}$ and $o_1^{\text{opt},0}$ were found to be (9,9) mm and *x*-orientation respectively after phase I. Phase II scans progressively constrain the search to a small area around optimal configurations found in previous scans (white squares). The $mMTD_1^3$ was found to be ~800 at (9.7,8) mm. The protocol's total costs for disclosing all 16 key bytes are compared to that of its time-domain counterpart using data from [3].

A total of $N_{\text{t}} = 500$ time samples were captured in a clock period. To generate the frequency-domain signals, the signals were padded to 1024 samples. For the given sample size, sampling rate, and limiting bandwidth of the setup, the number of points considered for analysis is $N_{\text{BW}} = 51$, which is ~10× fewer samples compared to the time-domain signal. Space-time and space-frequency maps for the baseline design are shown in Fig. 5. The spatial maps in this article are plotted for 101×101 observer locations for the first encryption. The time- and frequency-domain plots are shown at the leaking time-instant (~12 ns) and frequency (~150 MHz) for byte 1.

In this article, variations observed in fields from information-leaking logic blocks due to any variation in voltage supply or clock signal are assumed to be consistent across the chip. For example, if each AES S-box recieves the same input clock and supply voltage, all sensitive probe configurations are affected equally by voltage- and frequency-scaling, since these blocks are identical by design. As a result, it is expected that the optimal configurations identified for a baseline design do not change for the modules that are secured using external circuitry. These configurations can be used to compute the marginal cost for secured modules and evaluate the improvement in their security (otherwise the acquisition cost can be infeasible for secured modules). In cases where some key-bytes cannot be recovered, the improvement in resilience is computed by aggregating the marginal cost of attack for the recoverable key-bytes along with cost of the unsuccessful attacks on the remaining bytes, with a maximum of 20000 encryptions measured to recover each key-byte.

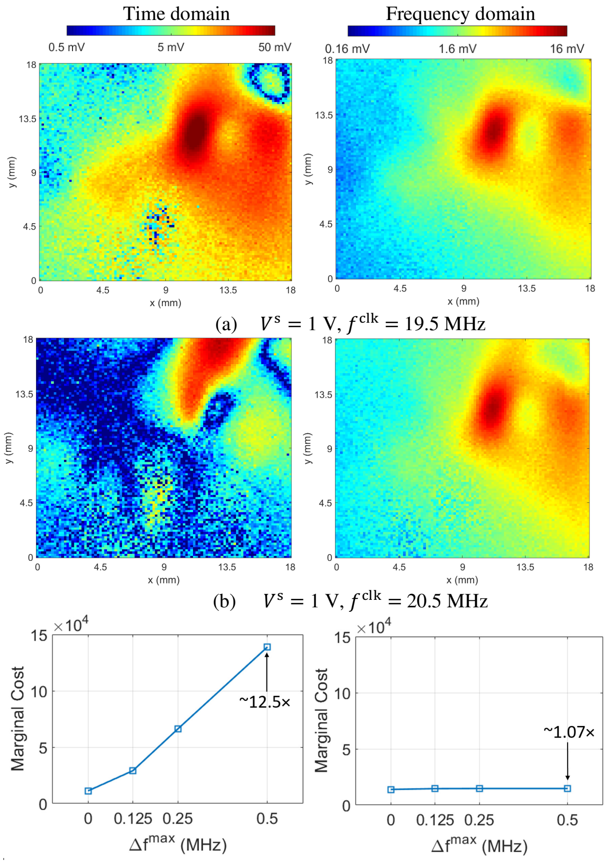

(a) $V^s = 1$ V, $f^{\text{clk}} = 19.5$ MHz

(b) $V^s = 1$ V, $f^{\text{clk}} = 20.5$ MHz

(c) Cost of EM SCA attacks

Fig. 7. Spatial maps of the measured voltage plotted for the extreme clock frequencies (a) $f^{\text{base}} - \Delta f^{\text{max}} = 19.5$ MHz and (b) $f^{\text{base}} + \Delta f^{\text{max}} = 20.5$ MHz, at the baseline supply voltage. Frequency-domain signals (right) were less affected than the time-domain signals (left), which degraded the time-domain EM SCA attack. (c) The countermeasure only increased the module's resilience against the time-domain EM SCA attack as the frequency range increased.

B. Measurement Results for Baseline Design

The adaptive scan protocol was performed on the baseline design using frequency-domain signals and the results are compared to the time-domain approach in [3]. Phase I of the protocol required $N_{\text{scan}}^{\text{I}} = 2$ scans, the final scan using $N_{\text{e}}^{2,\text{I}} = 6000$ encryptions per configuration for $N_{\text{obs}}^{2,\text{I}} = 11 \times 11$ observers, positioned at equally spaced grid-points over the chip. Minimal improvements were observed at $N_{\text{scan}}^{\text{II}} = 3$ scans for phase II. The number of observers for all scans in phase II were kept as $N_{\text{obs}}^{s,\text{II}} = 11 \times 11$. The MTD maps and comparison of costs for the time-domain and frequency-domain approach are shown in Fig. 6. At the end of the protocol, the frequency-domain approach had an acquisition cost of ~1.55×10^7 measurements and a marginal cost of ~1.4×10^4 measurements to recover all key-bytes.

It was observed that optimal configurations identified using both protocols were similar. For the baseline design, however, the frequency-domain protocol had ~1.5× higher acquisition cost compared to the time-domain approach with ~1.2× higher marginal cost. Therefore, the time-domain EM SCA attack was found to be the more potent approach for the baseline case. These optimal configurations are next used to evaluate the ability of countermeasures to secure the AES-128 realization on the Artix-7 FPGA against fine-grained EM SCA attacks.

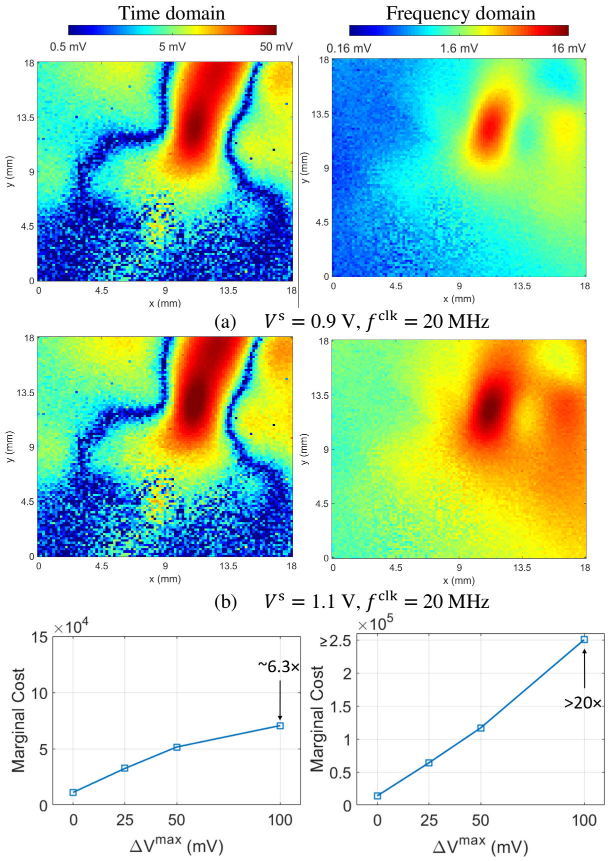

(a) $V^s = 0.9$ V, $f^{\text{clk}} = 20$ MHz

(b) $V^s = 1.1$ V, $f^{\text{clk}} = 20$ MHz

(c) Cost of EM SCA attacks

Fig. 8. Spatial maps of the measured voltage plotted for the extreme supply voltages (a) $V^{\text{base}} - \Delta V^{\text{max}} = 0.9$ V and (b) $V^{\text{base}} + \Delta V^{\text{max}} = 1.1$ V, at the baseline clock frequency. Frequency-domain signals (right) were more affected than the time-domain signals (left). (c) The countermeasure increased the module's resilience against both attacks as the voltage range increased. It reduced the effectiveness of the frequency-domain EM SCA attack more significantly.

C. Measurement Results for Countermeasures

The frequency-scaling countermeasure was implemented by randomly varying the clock frequency among 11 frequencies in the range $f^{\text{clk}} = 20$ MHz $\pm \Delta f^{\text{max}}$, where $\Delta f^{\text{max}} \in \{0.125,\ 0.25, 0.5\}$ MHz, while the supply voltage was kept constant at 1 V. This resulted in signal jitter in the range of 50 ns $\pm \Delta t^{\text{max}}$, where $\Delta t^{\text{max}} \in \{0.3, 0.6, 1.3\}$ ns. The impact of the frequency-scaling countermeasure on the observed fields and EM SCA attacks is shown in Fig. 7. Minor variations were observed in the frequency-domain signals because the signals were time-limited to the baseline clock period. This results in the addition or removal of signal components as the input frequency decreases or increases, respectively. Fig. 7(c) shows that the time-domain EM SCA attack degraded rapidly, with a maximum of ~12.5× increase observed over the baseline marginal cost, whereas the frequency-domain EM SCA attack was minimally impacted (~1.07× marginal cost). Delay-based countermeasures are indeed less resilient against frequency-domain EM SCA attacks.

The voltage-scaling countermeasure was implemented by randomly selecting a voltage supply level among 11 values within the range $V^s = 1 \pm \Delta V^{\text{max}}$, where $\Delta V^{\text{max}} \in \{25, 50, 100\}$ mV, while the clock frequency was kept constant at 20 MHz. Variations in time-domain signals were mainly observed at the signal peaks, while frequency-domain signals showed large variations across the chip (Fig. 8). The time-domain EM SCA attack had ~6.3× higher marginal cost, compared to the baseline design. The attack still succeeds in breaking the security of these countermeasures, since signals are still repeatable at some time-instances around the peaks albeit with lower correlation. The frequency-domain EM SCA attack, however, was impacted significantly due to the VS countermeasures. For a variation in the range ± 0.1 V, 2 key-bytes were not recovered and the resilience of the module was increased by a factor of >20×.

The voltage-frequency-scaling countermeasure combines the security features of the previous two: the delay component makes it resilient against time-domain EM SCA attacks and the voltage-scaling against frequency-domain ones. This countermeasure was implemented by randomly choosing among 5 fixed (V^s, f^{clk}) samples. Table I shows 3 sets with increasing ranges of voltages and frequencies from Set 1 to 3.

TABLE I. THE 3 SETS OF VOLTAGE-FREQUENCY PAIRS

Set 1	Set 2	Set 3
$\boldsymbol{\Delta V^{\max} = 20}$ **mV**	$\boldsymbol{\Delta V^{\max} = 50}$ **mV**	$\boldsymbol{\Delta V^{\max} = 100}$ **mV**
$\boldsymbol{\Delta f^{\max} = 0.1}$ **MHz**	$\boldsymbol{\Delta f^{\max} = 0.25}$ **MHz**	$\boldsymbol{\Delta f^{\max} = 0.5}$ **MHz**
V^s (V) , f^{clk} (MHz)	V^s (V) , f^{clk} (MHz)	V^s (V) , f^{clk} (MHz)
0.98, 19.9	0.94, 19.75	0.9, 19.5
0.99, 19.95	0.97, 19.875	0.95, 19.75
1, 20	1, 20	1, 20
1.01, 20.05	1.03, 20.125	1.05, 20.25
1.02, 20.1	1.06, 20.25	1.1, 20.5

The spatial maps and marginal cost of EM SCA attacks for this countermeasure are shown in Fig. 9. Significant variations were observed in fields in both time and frequency domain. For set 3, the side-channel resilience for time- and frequency-domain EM SCA attacks were observed to improve >25× and >20×, with 3 key-bytes not recovered in both cases.

IV. CONCLUSION

Countermeasures against fine-grained EM SCA attacks on AES, based on EMI reduction techniques, were evaluated using an adaptive scan protocol in time- and frequency-domain. The evaluations found that the frequency-scaling countermeasure had minimal impact on the frequency-domain attack but showed ~12.5× improvement in resilience against the time-domain EM SCA attack. The voltage-scaling countermeasure improved the module's resilience ~6.3× against the time-domain and >20× against the frequency-domain EM SCA attack. The voltage-frequency-scaling was the strongest countermeasure, showing >25× improvement in resilience against the time-domain and and >20× against the frequency-domain EM SCA attack. These countermeasures can be implemented with other countermeasures to enhance the side-channel security of cryptosystems.

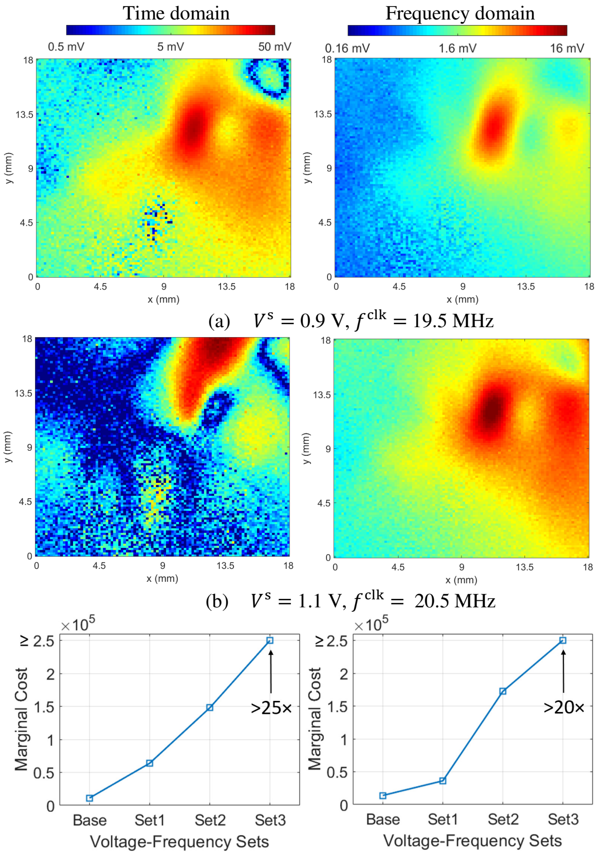

(a) $V^s = 0.9$ V, $f^{\text{clk}} = 19.5$ MHz

(b) $V^s = 1.1$ V, $f^{\text{clk}} =$ 20.5 MHz

(c) Cost of EM SCA attacks

Fig. 9. Spatial maps plotted of the measured voltage plotted for the extreme pairs (a) $(V^{\text{base}} - \Delta V^{\max}, f^{\text{base}} - \Delta f^{\max}) = (0.9\text{ V}, 19.5\text{ MHz})$ and (b) $(V^{\text{base}} + \Delta V^{\max}, f^{\text{base}} + \Delta f^{\max}) = (1.1\text{ V}, 20.5\text{ MHz})$. Both time- and frequency-domain signals were impacted significantly. (c) The countermeasure increased the module's resilience against the time-domain EM SCA attack as the voltage and frequency ranges increased. Set 1 reduced the effectiveness of the frequency-domain EM SCA attack marginally, but the module became more resilient as the voltage-frequency range increased.

REFERENCES

[1] E. De Mulder *et al.* "Electromagnetic analysis attack on an FPGA implementation of an elliptic curve cryptosystem," in *Proc. IEEE EUROCON*, pp. 1879-1882, Nov. 2005.

[2] Y. Hayashi *et al.* "Efficient evaluation of EM radiation associated with information leakage from cryptographic devices," *IEEE Trans. Electromagn. Compat.*, vol. 55, no. 3, pp. 555–563, June 2013.

[3] V. V. Iyer and A. E. Yılmaz, "An adaptive acquisition approach to localize electromagnetic information leakage from cryptographic modules," in *Proc. IEEE Texas Wireless Symp*., Mar. 2019.

[4] V. V. Iyer and A. E. Yılmaz, "Using the ANOVA F-statistic to rapidly identify near-field vulnerabilities of cryptographic modules," to appear in *Proc. IEEE Int. Microw. Symp.,* June 2021.

[5] NIST FIPS Pub. "197: Advanced encryption standard (AES)". *Federal information processing standards publication*, 197(441):0311, 2001.

[6] G. Li, V. Iyer and M. Orshansky, "Securing AES against localized EM attacks through spatial randomization of dataflow," in *Proc. IEEE HOST,* May 2019.

[7] J.-S. Coron and I. Kizhvatov, *An Efficient Method for Random Delay Generation in Embedded Software*, CHES, Berlin: Springer, 2009, vol. 5747.

[8] C. Sui, J. Wu, Y. Shi, Y. Kim and M. Choi, "Random dynamic voltage scaling design to enhance security of NCL S-box," in *Proc. IEEE MWSCAS*, Aug. 2011.

[9] A. Singh *et al*., "Improved power/EM side-channel attack resistance of 128-bit AES engines with random fast voltage dithering," *IEEE J. Solid-State Circ*., vol. 54, pp. 569–583, Feb. 2019.

[10] E. Mateos and C. H. Gebotys, "A new correlation frequency analysis of the side channel," in *Proc. ACM WESS*, Oct. 2010.

[11] ChipWhisperer, Github Repository [online], available: https://github.com/newaetech/chipwhisperer.

[12] V. V. Iyer, "An adaptive measurement protocol for fine-grained electromagnetic side-channel analysis of cryptographic modules," M.S. thesis, Univ. of Texas, Austin, Aug. 2019.