RAD-FS - *Inherent* and Embedded SCA-Security in Ultra-Low Power IoTs

Daniel Dobkin[1], Nimrod Cever[1] and Itamar Levi[1]

Bar-Ilan University, Ramat-Gan, Israel. Emails: daniel.dobkin@live.biu.ac.il, nimrod.cever@live.biu.ac.il, itamar.levi@biu.ac.il

Abstract. High-performance and energy-efficient encryption engines have become crucial components in modern System-On-Chip (SoC) architectures across multiple platforms, including servers, desktops, mobile devices, and IoT edge devices. Alas, the secure operation of cryptographic engines faces a significant obstacle caused by information leakage through various side-channels. Adversaries can exploit statistical analysis techniques on measured (e.g.,) power and timing signatures generated during (e.g.,) encryption process to extract secret material. Countermeasures against such side-channel attacks often impose substantial power, area, and performance overheads. Consequently, designing side-channel secure encryption engines becomes a critical challenge when ensuring high-performance and energy-efficient operations. In this paper we will suggest a novel technique for low cost, high impact, easily scalable protection based on Adaptive Dynamic Voltage and Frequency Scaling (A-DVFS) capabilities in ultra-low-power (ULP) sub-threshold chips. We review the improvement of using integrated voltage regulators and DVFS, normally used for efficient power management, towards increasing side-channel resistance of encryption engines; Pushing known prior-art in the topic to ULP-regime. The hardware measurements were performed on PLS15 test-chip fabricated in ULP 40nm process going down from nominal voltage to 580 mV power-supply. Various results and detailed analysis is presented to demonstrate the impact of power management circuits on side-channel security, performance-impact and comparison to prior-art. Importantly, we highlight security sensitivities DVFS embeds in terms of software side-channels such as timing, and their mitigation with our proposed technique, successfully masking the time signature introduced by DVFS.

Keywords: ADVFS · Side-channel attacks · SCA · IoT · Low-Power · ASIC · Hardware Security

1 Introduction

Side-Channel Analysis (SCA) attacks [KJJR11, BCO04a, Sta10a] are a category of hardware security attacks, which extract or retrieve information from a system by utilizing non-standard channels as compared to conventional communication interfaces. These non-standard channels are said to be *leaking* sensitive information manipulated by the digital electronic system or owing to the fact that the theoretical cryptographic system is implemented in practice, in the physical medium. The functionality of a system is as good as its implementation, and the implementation of a system is limited by the laws of physics. Therefore, the behavior of the system owing to logical-activity within it generates a physically-measurable reaction that can disclose information. For example, the timing of signals or power consumption [MOP08, MDS99], and Electromagnetic (EM) radiation [QS01, AARR02].

Dynamic voltage and frequency scaling (DVFS) [ZBSF04, LSH10, BHC+16, PBB98] is a widespread technique utilized to save power on a wide range of computing systems, from tiny Internet-of-Things (IoT) devices, embedded-systems, laptops/desktops systems

to high-performance server-systems etc. DVFS provides the functionality to reduce the energy consumption of an integrated circuit, implementing (e.g.,) a modern processor by reducing the frequency at which it operates. It can yield considerable reduction in energy consumption owing to the quadratic dependency of the possible operation frequency, f, in the operating voltage, V_{dd} stemming from strong-inversion/near-threshold current equations of transistors. Vast literature and research exist on the (very successful) use of DVFS to improve the energy efficiency of computation-systems. This is done by adapting the voltage/frequency to the time varying workloads of the system jointly constrained with the general timing limitations. Various modern microprocessors (if not all) and embedded-systems are equipped with the DVFS functionality such as chips from Nvidia, AMD, Intel [SGS+14, MWC17, ACK19]. In some devices one can find hundreds of voltage/frequency levels with high resolution, controllable through software such as MSI Afterburner [XM13] and different parts of the systems are adapting independent sub-systems, with their own *optimized* DVFS tactic, such as ARM big.LITTLE [Inc22, PPC+15].

In conventional processors, implemented with standard process technologies, the operating voltage-span ranges from V_{dd} overdrive of several hundreds of mV over the nominal voltage and down to still strong inversion / high near-threshold region of transistors, only several hundreds of mV under the nominal voltage. It is possible to implement electronic systems that operate over a much larger voltage span, from over V_{dd} to sub-threshold voltages. This however requires special (and not always standard) design-techniques such as embedding isolation-cells, level-shifters, and special cells libraries (high-V_{Th}) etc. The main challenge is that typically such devices are not fully characterized by the foundry in the entire voltage-span, making it hard to design robustly and perform full-digital design-flow with guaranties. Nevertheless, this possibility opens the question of DVFS efficiency in such extended ranges. Existing theoretical research have proven that further energy efficiency is possible with extended voltage-range, even ranging below $0.5 \cdot V_{dd}$. However, extending it to full-sub-threshold region is beneficial only for specific application classes. Therefore, in practice, with the emergence of Ultra-Low-Power (ULP) IoT devices which can sacrifice performance for longer (say) battery-life, we see unique ULP-IoT platforms which are designed to work down to 500-600 mV (generally mid to high Near-threshold operation). Such systems highly utilize DVFS techniques. Designing such systems is not an easy task and require high-expertise, but available devices exist on the market which we believe will provide game-breaking abilities for (e.g.,) IoTs. One such unique platform is PLSense's PLS15 platform, used in this research.

Side-channel Analysis (SCA) attacks security is a critical requirement. Typical mechanisms to counteract SCAs in a mathematical rigorous way are very expensive. For example, by coding an internal value or variables in the computation to a redundant representation which is randomized and is invertible, denoted by Masking [CGLS20, SL23]. The problem is that it does not nicely fit the ULP-IoT paradigm in terms of performance improvement and energy savings, far from it. On the less rigorous approach, more heuristic solution-space exist utilizing other randomization mechanisms such as randomizing the time [VCMKS12] or amplitude [LBBS20] signal domains. For time-variations, various Shuffling mechanisms were recently supported with theoretical models indicating quasi-linear security-levels [LBS20]. For Amplitude randomization mechanisms a fine-grain leakage model supported by silicon measurement was recently provided showing also quasi-linear security-levels [BSL22]. The electronic-cost of both techniques is far lower than that of Masking, but each approach has its own limitations such as algorithmic-dependence or the need to embed specialized components which are not natively there in an off-the-shelf device, and more. Several other countermeasures, which rely on randomizing voltage / frequency are described in [SKM+19, GDS20, KLS+21], but requires significant IC redesign, including (e.g.,) an all-digital clock modulation (ADCM), using its global modulator; On the contrary, the proposed approach is inherently embedded in such platforms and is

software (SW) controlled, also pushing the limits towards ULV-IoTs.

With this motivation, our goal in this research was to utilize the inherent mechanisms existing in ULP-IoTs (or even high performance aggressive DVFS systems) to enable SW controlled SCA-security with zero effective-cost in area and power, and to provide a detailed analysis and to exhaust measurement-evaluation on such a unique state-of-the-art (SOTA) platform. The developed technique, dubbed Randomized Aliasing Dynamic Frequency Scaling (RAD-FS), aims to give flexibility to randomize the operating frequency (or alternatively the power-supply voltage), but still maintain some application ability to adapt. Clearly, there exist a trade-off between security and performance, as we detailed below. However, interestingly the proposed technique has the potential to considerably resist a rather new class of software/network *external* attacks which are directed and inherent to DVFS systems such as [WPH+22, NIC+23], we focus on this important DVFS double-edge-sword security aspect and its mitigation with RAD-FS.

1.1 Contribution

In this research we push forward three observations: (1) nowadays embedded-systems embed efficient DVFS mechanisms inherently which can be utilized to integrate security features with out special design efforts and hardware intervention (2) such mechanisms can be very efficient, as supported by rigorous evaluation to counteract SCAs, and (3) native DVFS mechanisms inherently induce other SCA channels (timing and power-state monitoring) which becomes data/workload-dependent, our proposed mechanisms can potentially aid in mitigation of these channels. We provide for the first-time significant advance on the analysis of all these aspects and we showcase several SOTA use-cases on a very advanced platform, extending significantly the body of knowledge in such ULP regime. **(1):** Working with a SOTA ULP-IoT device, where scarce platforms exist in the market going down to 500mV in operational chips **(2):** Providing security analysis over such devices for the first time to the best of our knowledge.**(3):** Proposing unique methodology and embedded mechanisms to provide SCA security utilizing the inherent DVFS features with ultra low-cost as compared to other solutions regarding area, latency, power, implementation effort overheads; denoted Randomized Aliasing Dynamic Frequency Scaling (RAD-FS). **(4):** Reporting analysis both for a Risc-V processor core and an NXP encryption accelerator embedded on the same device in 40nm technology, in a comparative view. **(5):** For the first time we demonstrate a countermeasure to a new class of timing-attacks such as [WPH+22], which coexists with DVFS mechanisms[1], in addition to the inherent power-SCA attacks immunity. We show RAD-FS is very relevant for such network timing attacks mitigation. Results are demonstrated via. an ideal (optimal) oracle modeling the RAD-FS parameters. Our constructed oracle is very *generous* with how much control is given to the adversary. The outline of this paper is as follows. In Section II, we present a background and literature review. In Section III, we present our novel solution with detailed explanation and comparison to the current industry approach. Section IV details the evaluated device, and the measurement setup. Section V details security analysis metrics. Section VI contains in-depth analysis of the results and an analytical approach of the performance cost. Section VII discusses comparisons and challenges adversaries face in a real life sceanrio.Section VIII finally, concludes the results of this work and proposes future prospects.

2 Background

SCA attacks are powerful, repeatedly showing their efficiency in extracting sensitive information from a cryptographic system by analyzing unintentional side channels. Cryptographic systems aim to maintain data confidentiality by encrypting it, but they can inadvertently leak information through various side channels. Side Channel Analysis

[1] opposed to the naive solution of turning DVFS off

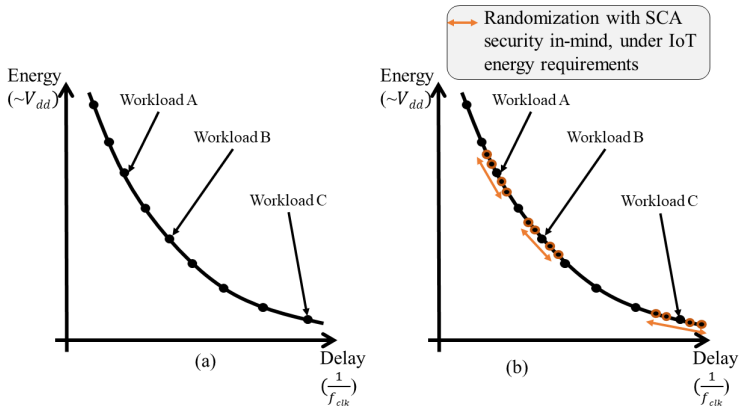

Figure: 1. DVFS conceptual implementation (a) insecure, single frequency per workload vs. (b) our proposed secure implementation, with multiple frequencies assigned per workload.

leverages these side channels to gain insights into the internal operations of a cryptographic device or algorithm, potentially compromising its security. SCA countermeasures are very expensive (in the electronic sense) and are very hard to embed with a strict energy-budget. This is witnessed by the requirements of the NIST lightweight authenticated-encryption contest [TMC+23]; especially for low-power constrained or battery operated IoT devices that can not sacrifice energy budget as needed for rigorous security-level utilizing e.g., Masking.

With traditional side-channel attacks the attacker typically obtains a set of side-channel measurements, while the target device performs cryptographic operations using the key or data under investigation. These measurements are used in an attack campaign. Generally, attacks are classified to model-based or profiling-based: model-based attack such as correlation power analysis (CPA) [BCO04b] utilizing Pearson's correlation coefficient ρ and a leakage model to compare side channel leakages to the model. Template attacks [CRR03] utilize a template or a statistical model derived from the leakage of a specific internal value used to enhance the attack's effectiveness in an attack campaign by comparing it to the actual leakage.

Transistors, as the fundamental building blocks of integrated circuits, play a crucial role in amplifying and switching signals. In today's context, the demand for energy-efficient and low-power electronic systems, including battery-powered devices and Internet of Things (IoT) applications, has emphasized the need for fully functional ULV (Ultra-Low Voltage) standard cells. The concept of Dynamic Voltage and Frequency Scaling (DVFS) is based on the understanding that the energy consumption of a component is directly proportional to its supply voltage, while the computation delay is inversely proportional to the square of its operating frequency, as depicted in Fig. 1(a). By reducing the supply voltage or clock frequency, it is possible to achieve significant reductions in energy consumption. Conversely, increasing the voltage and frequency deteriorates energy-consumption, albeit the improved latency or increased throughput, is advantageous when the workload calls for such adjustments. Overall, the development of fully functional ULV standard cells has become crucial in meeting the growing demand for energy-efficient and low-power electronic systems. DVFS offers a means to optimize energy consumption by dynamically adjusting the supply voltage and clock frequency of components, striking a balance between energy and performance based on workloads.

Frequency randomization is a technique used to introduce randomness or unpredictability in the occurrence or timing of events. It is often employed in various domains, including

communication systems, network security, and data privacy, to mitigate potential attacks or reduce vulnerabilities. By randomizing the frequency of events, it becomes harder for adversaries to analyze patterns or launch coordinated attacks. We generally define the possible under-evaluation set of frequencies by F', with the number of possible frequencies in the set being $\|F'\|$. We aim to show how the different attributes of F', such as its size and range (or bandwidth) affect SCA security.

To evaluate SCA security in this paper we utilize an approach based on analysing results from a bank of efficient known attacks such as CPA [BCO04b], cryptographic-sense SNR [Man04b], template attacks [CRR03] and computing the attack success-rate, SR [Sta10b]. In addition, to evaluate the theoretical informativeness of the leakage, without connecting to an actual attack, we utilize a common leakage-detection test, based on Welch's two-tailed T-Test, namely test-vector leakage assessment, TVLA [CDG+13]. It is commonly used in SCA security to compare the means of two sets of leakages measured from the design and partitioned in accordance with specific sets of known data. In the context of our research, these sets of data (populations) are the power measurements of encrypting a set plain-text vs randomized plain-text (e.g., the fixed versus random, F-vs.-R test).

Double-edged sword: While DVFS has obvious merits, as hinted above, it introduces security-flaws by its design. Workloads, manifested by different data manipulations, affect the electrical characteristics of an operation (voltage, latency etc.), leading to various sensitivities. As opposed to conventional SCA attacks which measure power or electromagnetic emissions and require a *close-contact* adversary, DVFS also allows for SW based attacks by various mechanisms: For instance, Hertzbleed [WPH+22] represents a novel category of side-channel attacks that exploit network timing and latency profiles in the spectrum. These vulnerabilities can be leveraged regardless of the physical distance. Hertzbleed Has demonstrated extraction of cryptographic keys from remote servers which embed modern x86 CPUs in a scenario that was previously believed to be secure; *solely exploiting the inherent sensitivity DVFS offers.* Another example is [NIC+23] in which the voltage dependence offered by DVFS on an iPhone-13 affects the overall power of the device, and a video footage of a device's power-LED, radiating to a long distance can be used for SCA. Alternatively, the authors have shown that if the device is powered by a USB-hub, the hub's current draw can provide an attack entry. Noteworthy, in these reports the attacker makes use of an ultra low-resolution side-channel (as compared to high-resolution e.g., power-based SCA): The device's embedded power-sensor, which can be accessed through the network and used to regulate the DVFS. As an example, its resolution is clearly far from being as high as of conventional 10 to 14-bits quantizer of an Oscilloscope. In addition, it is typically a very slow sensor hence significant averaging and noise is incorporated in this physically measurable quantity. Therefore, it is quite easy for our randomization mechanism to make these attacks hard.

We show that by randomly choosing a frequency from F, we affect the data dependency of the frequency with several security parameters our countermeasure embeds, and make it hard for an adversary to discern power-states (P-states) and data hamming-weight, HW, calculated in a given P-state.

3 The proposed Approach

We introduce Randomized Aliasing Dynamic Frequency Scaling (RAD-FS), as a side-channel security mechanism. Our method suggests assigning a group of frequencies f to a workload from the generally available set of operation frequencies F' (i.e., f is a subset), instead of a single frequency f_{base} as illustrated in Fig. 1(b). f is constructed such that each $f_n \in F$ is within a bandwidth, BW, around f_{base}, allowing for power optimization in relation with workload. Denoting $\|F\|$ as the amount of $f_n \in F$.
For each $f_n, f_m \in F | n \neq m$ we achieve **aliased** distributions of the leakage induced by some internal-value manipulation. Optimally, we aim for these disturbances to: (1)

distribute uniformly to maximize the leakage entropy, and (2) overlap significantly to increase the noise-level for proximate time samples.

As we increase $\|F\|$, adding more frequencies to our set F, we expect security-metrics e.g., the SNR, to show more distribution curves across different time samples, in correlation to various $f_n \in F$. These high leakage-correlation points-of-interest, POI's, reflect the now *shifted* value manipulation corresponding with f_n's. The overall noise increases due to cross-interference between different leakages stemming from different $f_n \in F$, as abstractly illustrated in Fig. 2. This should manifest in all uni-variate security metrics such as CPA, SNR, Template attacks and TVLA detection tests[2].

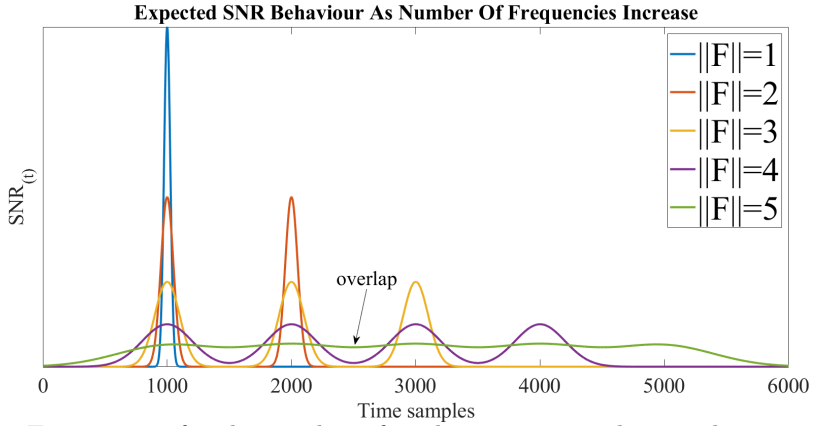

Figure: 2. Expectation for the number of peaks to increase along with group size $\|F\|$, while concurrently the magnitude (ρ, SNR etc.) of each peak decreases.

In practice, for every internal hypothesized calculation during e.g., an encryption, we find various POIs associated with some $f_n \in F$. The $SNR_{(t)}$ figure does not show only a single peak, due to various f_n's, and as internal-values correlate differently with the leakage in various time instances. In our evaluation environment (detailed below) implementing the AES cipher, it is possible to see more peaks appearing with lesser amplitudes as $\|F\|$ goes up, as shown in Fig. 3. The SNR value decreases, clearly owing to traces cross-interference of $f_n \neq f_m \in F$ in the time sample during the calculation of the SNR estimator. Already, at this early stage, we can hint a significant order-of-magnitude improvement with only $\|F\| = 5$.

Fig. 4 shows the FFT of several scenarios. In Fig. 4(a) we show the log-scale power of the single-sided spectrum when a single frequency is set ($\|F\| = 1$). In Fig. 4(b) we construct F with the same frequencies used in Fig. 4(a) applying RAD-FS ($\|F\| = 7$). The energy per frequency is significantly reduced as expected (note the logarithmic scale) and is distributed quasi-uniformly among $f_n \in F$. This implies that any filtering attempt will either eliminate information, induce overlaps in the time domain, and alternatively increase the noise-floor.

4 The Evaluated Device, Testing Modifications and Measurement Setup

4.1 The Evaluated Device

Our tests were performed on the PLS15, an advanced chip made by PLSense. The PLS15 is an ultra-low power MCU with multiple analog and digital interfaces and other capabilities like ML inference engine, crypto-cores, Risc-V processor and other ULP features making it a very interesting candidate for our experiment due to desirable features for IoT. The

[2]We relate to other attack-settings in later sections

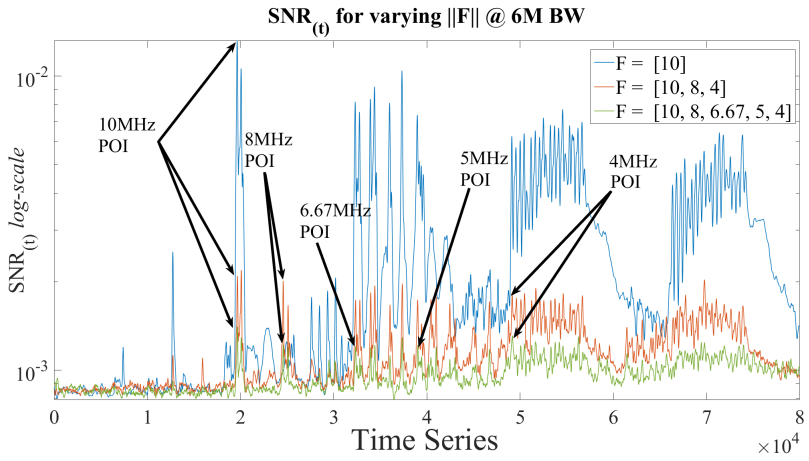

Figure: 3. Displaying with arrows when peak SNR values appear regarding to chosen $f_n \in F$, per case with different sized $\|F\|$. The Δ between the POI and the closest secondary peaks decreases, making it harder to choose a time sample for further analysis (i.e., template attacks).

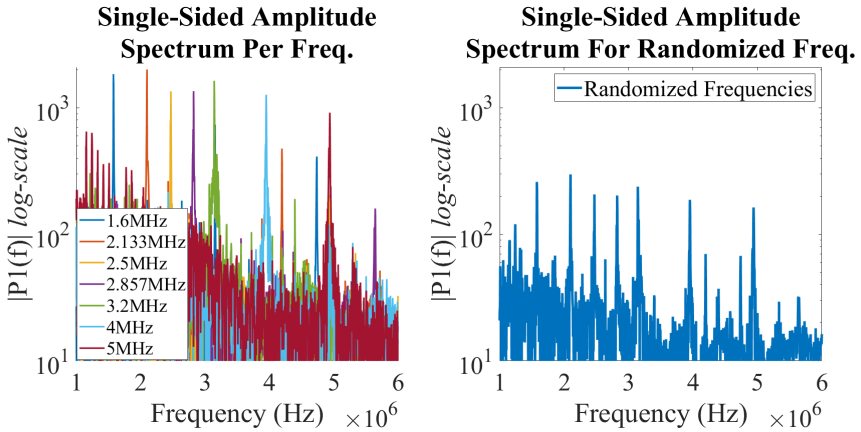

Figure: 4. FFT of different scenarios: (a) no randomization, single frequency $f_n \in F'$, $\|F\| = 1$ (b) Randomized scenario, $f_n \in F$, $\|F\| = 7$.

PLS15 is manufactured on the 40nm TSMC process. Usually, this process node has nominal working voltage of 1.1V. The reason that the 40nm Low-Power (LP) process was chosen is owing to device-leakage current and dynamic power consumption savings of up to 51% as compared to its 65nm counterpart. By utilizing mixed threshold-voltage (V_t) transistors in a single cell, supported by unique Adaptive Dynamic Voltage Control in the PLS15, the chip allows reduction of the operating voltage, bulk biasing, sensitivity to process variations, and more, to achieve a sub-threshold operating voltage of 0.45V-0.6V according to the workload conditions. Relevant blocks in the chip are a Risc-V Core, NXP AES Accelerator (unprotected), DMA controller and Adaptive-DVFS (ADVFS) Logic. Though some other devices exist on the market incorporating DVFS and ULP process towards IoTs, we didn't come across competitors reaching such deep near-threshold voltages in such a complex SoC.

We have evaluated the voltage-frequency map of the PLS15 device as illustrated in Fig, 5 showing 19 discrete possible frequencies.

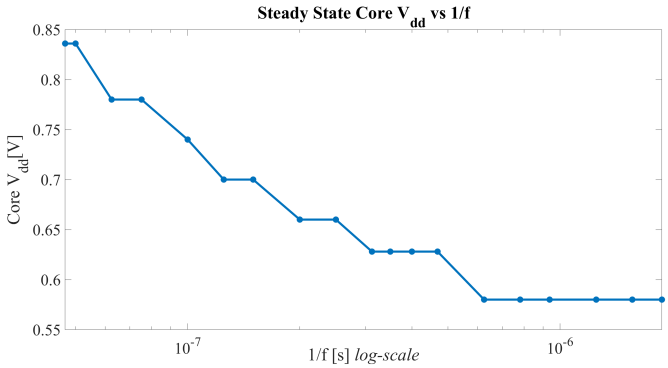

Figure: 5. Core voltage per 1/freq. available on our target device, PLS15

4.2 Testing Modifications

The PLS15 chip has an ADVFS controller module that affects the chip's core voltage. The ADVFS controls the input voltage to the core by utilizing an external Op-Amp in a negative feedback loop, connected to the power supply pin. We modified the PLS15 test kit board such that the ADVFS is semi-enabled; The ADVFS operation is in a mode where it operates based on SW commands only, and not according to workload. This gives us full control over data-dependency and P-states, allowing us to reduce algorithmic noise. In one scenarios, it allows us to play-out a scenario where the adversary has some form of access to the device DVFS. For power measurement, we added a jumper in serial between the external core voltage regulator and the core IO pin for induction-based current sensing, utilizing the Tektronix-CT1 current probe, connected through an amplifier to a Picoscope Oscilloscope as illustrated in Fig. 6.

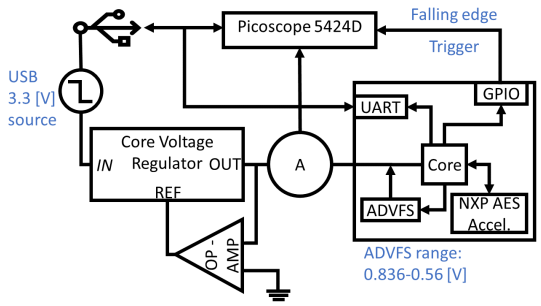

Figure: 6. Connection scheme of the core power supply and voltage regulation for the PLS15, with our current probe location.

4.3 Measurement Setup

Our test bench is comprised of a PC, Picoscope 5424D, Rohde & Schwartz Signal Amplifier, Tektronix CT1 and PLS15 target connected as shown in Fig.7. The measurement setups is shown in Fig. 8.

Our experiment method pseudo-code is described in Fig.24 of Appendix A. For the basic SCA evaluation we wanted to eliminate algorithmic noise and to show the adversary best-case scenario. Hence, we evaluate the leakage stemming from one byte manipulation leakage at a time (e.g., Byte 7 in the aforementioned figure). We also control the operating frequency driven by the regulator through our python shell triggering the device, and we have added NOP cycles inline Assembly to our C implemented code loaded to the device, to reduce noise after a triggering event.

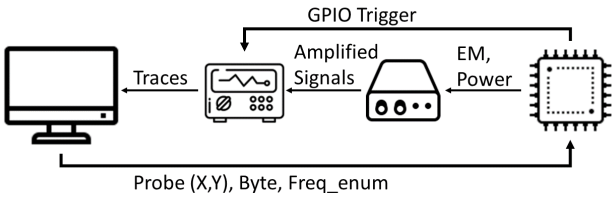

Figure: 7. Illustration of the test bench we assembled comprising a PC running a python script, Picoscope which records measurements of power and EM traces, Signal amplifier, Riscure Probing Station, PLS15 Chip target. Icons from [Ico23]

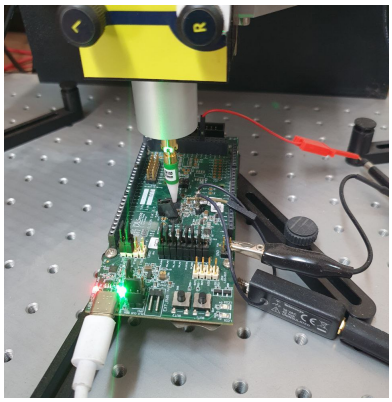

Figure: 8. Measurement setup with the PLS15 evaluation-board, the CT-1 current probe connected to a power jumper and the 0.1 mm Riscure's EM coil probe on top (tests were also performed with Langer's probes set).

5 Security Analysis Metrics

5.1 Low Complexity Adversary - Estimators

As discussed above we evaluate (as a starting step) Mangard's SNR [Man04a] & Brier's CPA [BCO04b] correlation as defined by:

$$\mathrm{SNR(t)} = \frac{\mathrm{Var}_{x_i,k}(\mathrm{E}[l^t_{x_i,k}])}{\mathrm{E}_{x_i,k}(\mathrm{Var}_i[l^t_{x_i,k}]} \tag{1}$$

$$\rho_{\mathrm{l^t_{x_i,k}},\mathrm{h^t_{x_i,k^*}}}(\mathrm{t}) = \frac{\mathrm{Cov}(l^t_{x_i,k^*}, h^t_{x_i,k^*})}{\sigma_{l^t_{x_i,k}}\sigma_{h^t_{x_i,k^*}}} \tag{2}$$

where, Var and E are the variance and expected estimators, $l^t_{x_i,k}$ is the leakage trace l in point in time samples t, taken from a cryptographic operation processing key k and plaintext (e.g.,) byte x_i.
In accordance with the correlation CPA distinguisher, we enumerate all possible (sub-) keys hypothesis k^* so as to generate the leakage hypothesis h. Then the k^* which maximizes the correlation is estimated to be the correct key. We then compare:

$$SNR_{t=POI} = max_t(|SNR|) \tag{3}$$

$$Corr_{t=POI} = max_t(|Corr|) \tag{4}$$

From now referred to as Point Of Interest (POI) for the SNR & Corr (CPA's ρ) accordingly. Both estimators were computed over $0.5{\cdot}10^6$ to $10{\cdot}10^6$ traces (as needed) or

queries. Our results include both a Risc-V implementation of tiny-AES 128-bit code, verified against NIST [Dwo01], and an NXP crypthash hardware (HW) accelerator, embedded in the PLS15 SoC, running 128-bit AES as well. For illustration, Fig. 9 shows the mean power-trace of 10K leakage traces of the fast HW accelerator to the left and the slow SW implementation to the right. As further example Fig. 10 shows side-by-side the SNRs of the NXP AES accelerator and the SW AES on the Risc-V processor in a comparative view, to the left and right, respectively. Note the very high SNR value achieved and the ultra fast operation of the accelerator owing to the fact that it is *seated* in its own power domain in a tailored IP block and not as part of a *sea-of-gates* as is typically the case for processors-cores.

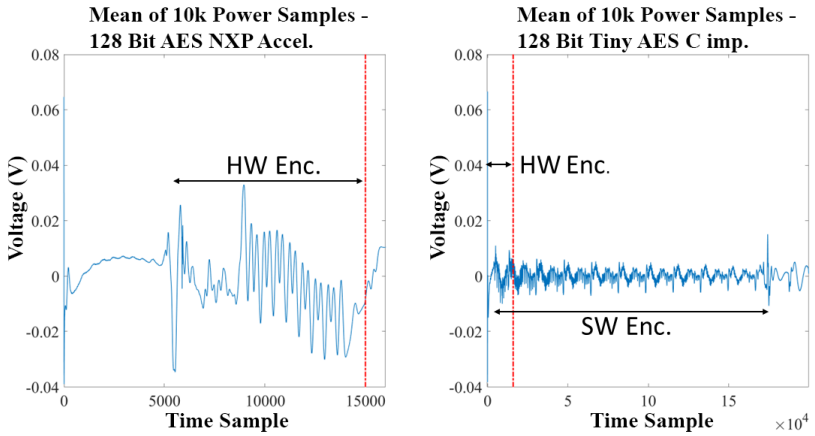

Figure: 9. Mean of 10k traces 20MHz @ core frequency. NXP crypthash accelerator running 128-bit AES (left), 128-bit tiny AES C implementation (right). The vertical red dashed line marks the end of the HW AES accelerator relatively to the SW one.

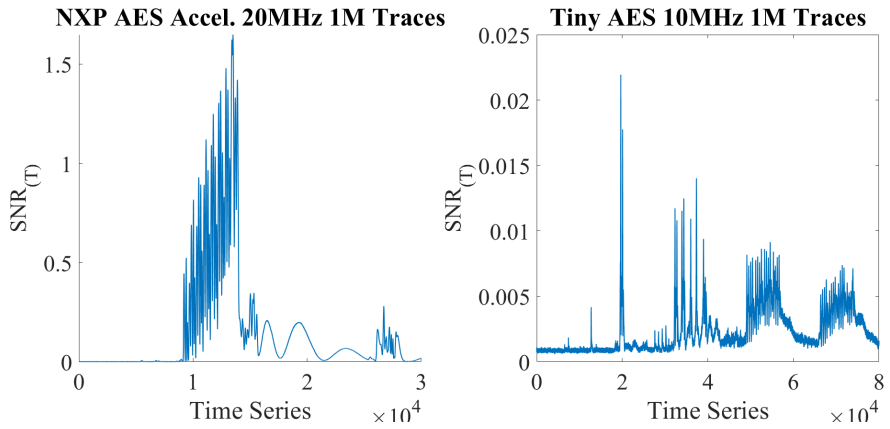

Figure: 10. SNR of the NXP AES accelerator **to the left** and the software Risc-V AES **to the right** under evaluation.

5.2 Detection Test - TVLA

As a detection test we apply TVLA verified against [SM15] to traces measured on our system, in order to show the difference in populations and the shift of data to higher

statistical momentum. The compared populations are constant plaintext & a varying plaintex, this experiment set is run several times under different conditions

5.3 High complexity adversary - Templates

Template attacks [CRR02] are performed in two consequent (or interleaved) phases of *profiling* and *attack*. It is assumed that the adversary got hold of one device for which he can program (or control) the secret key and therefore profile the leakage, and another target device from which he tries to extract information on the underlying key. We built a probability density function (PDF), for an internal manipulation, represented by function F, $y = \mathsf{f}(x_i, k)$. A set of $\mathcal{L}_p$ profiling traces of size N_p was used in order to estimate distributions, denoted as $\hat{M}_y$. Specifically, $\mathrm{f}(l \mid y) = \mathcal{N}\left(\widehat{\mu}_{l|y}, \widehat{\sigma}_{l|y}\right)$. For the attack phase, we utilize $\mathcal{L}_{att}$ of size N_{att} traces. The secret key k^* which maximizing the univariate Maximum Likelihood (denoted by LH) is chosen: $k^* = \underset{k}{\operatorname{argmax}}\ \mathrm{LH}(k)$, *i.e.* $k^* = \underset{k}{\operatorname{argmax}} \prod_{j=1}^{N_{att}} \left(\mathrm{f}(l_i \mid y_i)\right)$. As standard, owing to practical computational reasons and numerical errors, the log-likelihood (LLH) was used [FDLZ14] $k^* = \underset{k}{\operatorname{argmax}}\ \mathrm{LLH}(k)$ $= \underset{k}{\operatorname{argmax}} \sum_{j=1}^{N_{\mathrm{att}}} \log\left(\mathrm{f}(l_i \mid y_i)\right)$.

5.4 Timing Attacks

We aim to show the relevance of our technique to timing attacks that rely on the vulnerability introduced by the very same DVFS mechanism we rely on. As an example, [WPH+22] relies on the time variance induced by DVFS, which can be manipulated by an adversary to gleam secret information via the time channel. Our technique should make it harder for an adversary to see the different distributions and to separate information about secret computation from time measurement.

6 Sterile Analysis - Ideal View

First, we aimed to establish that our devised method works in a sterile clean scenario, i.e., by gradually increasing the % of traces taken with altered core frequency. This is analogous to uneven weights in a distribution function. For example, 10% altered frequencies means we operate 90% of the times with say $f_{base} = 20\mathrm{MHz}$ and 10% of the time with some other frequency in the set say $f_n \in F$:

$$P(f_{base} = 20MHz) = 0.9, P(f_n \in F \setminus f_{base}) = 0.1$$

We evaluated the effect of switching the core frequency once every 100, 10, 5 or 2 encryptions from 20MHz to either 16MHz or 13.3MHz on the SNR@POI and Corr@POI values as shown in Fig. 25 and Fig. 26 of Appendix A, respectively. Two phenomena are immediately observable. First, we find a linear decrease in value as we increase the probability (or percentage) to the uniform scenario (50%-50%). Secondly, we observe differences with different bandwidths (BWs) upon which we will rigorously discuss below.

6.1 Our Optimization parameters for RAD-FS

Under system restrictions (i.e., the discrete DVFS values described in Fig.5) we grouped frequencies within F' to isolate the parameters of BW and $\|F\|$ as classified in Table 1. In addition we grouped frequencies from F' to isolate the effects of f_{min} per a given BW as partitioned in Table 3 of Appendix A. Figures 11 and 12 shows the resulting maximum SNR and correlation values over time, respectively for the Risc-V SW implementation of the Tiny-AES.

Table 1: Classification of frequency map of F' to different BW with varying $\|F\|$s

$\|F\|$\BW	9.33[MHz]	8.3[MHz]	6[MHz]	4.57[MHz]	1.9[MHz]
$\|F\| = 2$	[16 6.67]	[13.3 5]	[10 4]	[6.67 2.1]	[4 2.1]
$\|F\| = 3$	[16 10 6.67]	[13.3 8 5]	[10 8 4]	[6.67 5 2.1]	[4 3.2 2.1]
$\|F\| = 4$	[16 13.3 10 6.67]	[13.3 10 8 5]	[10 8 6.67 4]	[6.67 5 3.2 2.1]	[4 3.2 2.5 2.1]
$\|F\| = 5$	[16 13.3 10 8 6.67]	[13.3 10 8 6.67 5]	[10 8 6.67 5 4]	[6.67 5 4 3.2 2.1]	[4 3.2 2.857 2.5 2.1]

$\|F\|$ **size:** Randomizing chosen frequencies from F. described in Table 1 in a uniform distribution. A reduction by an order of magnitude is observable immediately just by iterating between 2 frequencies ($\|F\|$=2). A further reduction in the SNR@POI by almost another order of magnitude is achieved by increasing $\|F\|$ to 5 frequencies (increasing $\|F\|$ even further is clearly possible and depends on the system under evaluation).

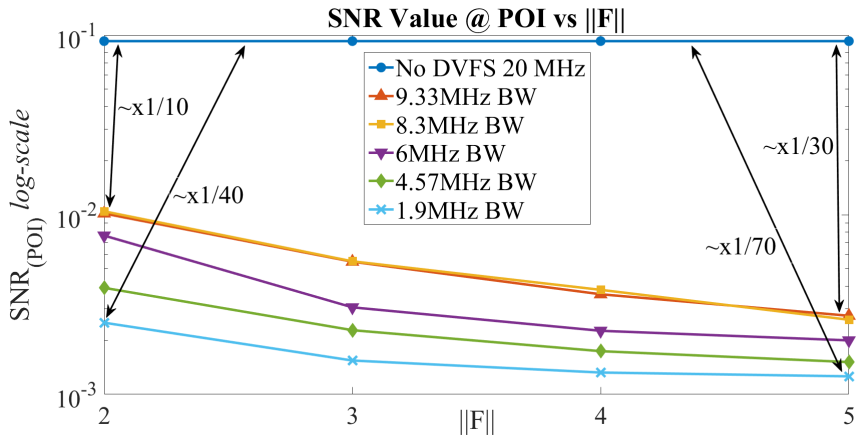

Figure: 11. SNR@POI values vs. size of group F, for different BWs.

A similar phenomenon is observable in the correlation graph, albeit the scale of reduction is smaller. Note that the reduction in both SNR and correlation is inversely proportional to the data-/time-/computation-complexity of an attack. *Therefore, two orders of magnitudes are quite significant leading to a noteworthy security-level, attack complexity etc.*

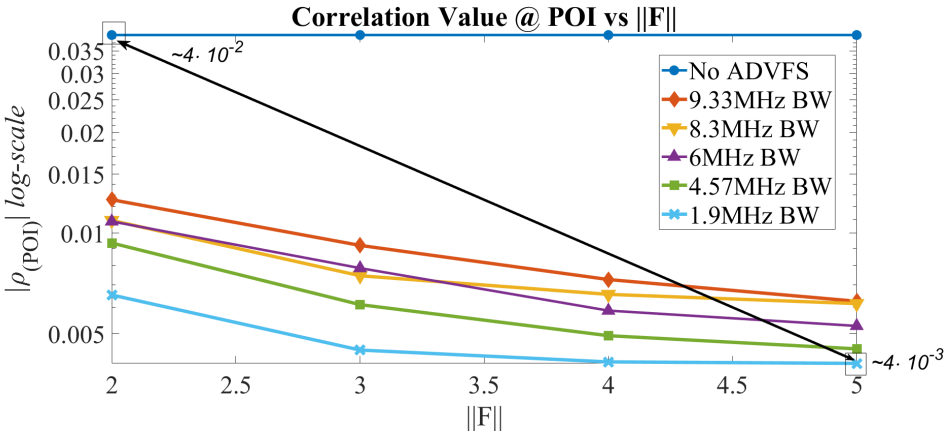

Figure: 12. Correlation@POI values vs. size of group F, for different BWs.

BW parameter, and F_{min}: By reordering the data to look at BW influence on the SNR, we show the relation between SNR@POI values and BW across different $F_{min}, V_{dd\ min}$. As demonstrated in Fig. 13 , as we decrease the BW, we increase the overlaps (i.e., aliasing) of leakages between sets of traces from different f_n's in the time domain. A smaller BW implies decimation in the frequency domain, which corresponds to overlaps in the time domain, thus the expected value of uni-variate analysis mixes up different time samples or internal computations. The more leakage overlaps the larger the effect is on the SNR. This is consistent for different voltages, with a *knee*-frequency starting at 6-7MHz, independent of $\|F\|$ (related to voltage map in Fig. 5).

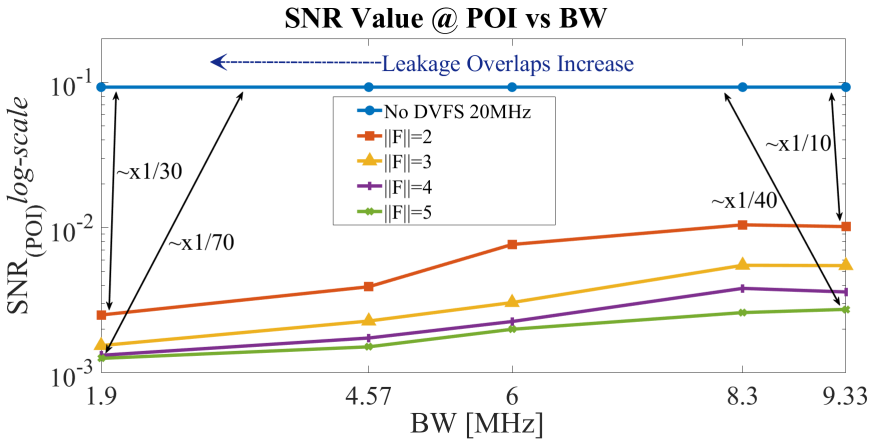

Figure: 13. SNR@POI values vs Bandwidth, a decrease by a factor of one to two magnitudes is achieved carefully selecting the BW to generate aliasing in the time-domain.

f_{min}: In order to isolate the effect of f_{min}, we grouped F's such that $\|F\| = 2$ while keeping BW constant to the best of our abilities under system limitations, and varying f_{min} (listed in Table 3). Comparing Fig.14 to the Fig. 13,11 we can see that although f_{min} has an effect, it is considerably weaker than the effect of $\|F\|$ and BW. Meaning V_{min}[3] has an effect but is not the main focus.

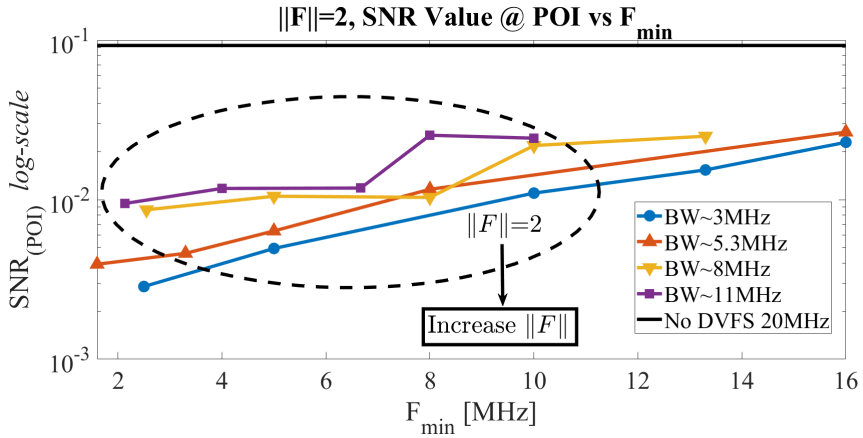

Figure: 14. SNR@POI values vs f_{min} for different BW.

Generally, the experiment set highlights that both the Corr. and the SNR estimators performed quite similarly. However, results (security gains) were slightly poorer with CPA

[3]lower frequency leads to lower V_{min}

since with correlation the signal is not scaled to the noise compared to the SNR.

6.2 Concrete Security Evaluation - Detection Test

In Fig. 15 We performed T-Test on $0.5{\cdot}10^6$ to $10{\cdot}10^6$ traces to evaluate the security of the proposed mechanisms for different $\|F\|$. Fig. 17 and Fig. 16 show the maximum absolute values over time of the T-Test detection vs. the number of collected traces and for different $\|F\|$ for the NXP HW accelerator and the Risc-V SW implementation, respectively. It is evident from Fig. 16 that using RAD-FS shifts the data to higher statistical moments, observing $\|F\| = 1$ we see no leakage in the 2nd moment, wherein $\|F\| = 2$ and above the T value becomes significant. This requires higher computational efforts from a potential adversary for a successful attack. It is important to note that information evidenced by a detection test does not practically imply an attack is known or easy, we show below that with the best uni-variate template attack success-rate, the protection level provided by RAD-FS is quite remarkable. On any account, even with the T-Test, adversary-complexity increases by orders-of-magnitudes.

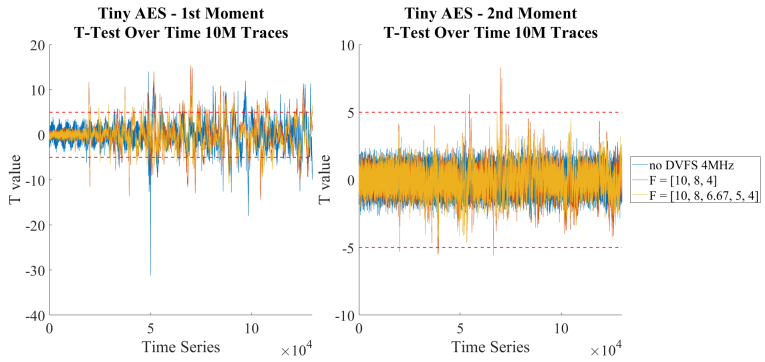

Figure: 15. Risc-V SW: example T-Test detection test based on TVLA methodology over time for 10M traces and for different $\|F\|, BW = 6MHz$.

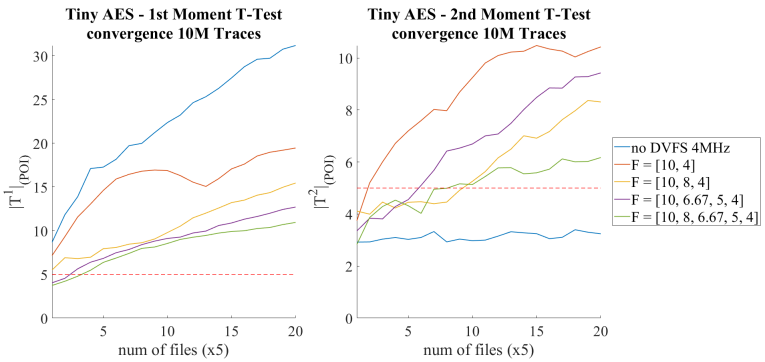

Figure: 16. Risc-V SW Tiny-AES: Maximum absolute values over time of a T-Test detection test based on TVLA methodology vs. the number of collected traces and for different $\|F\|$ for a given BW(6MHz), $V_{min} = 0.66[V]$.

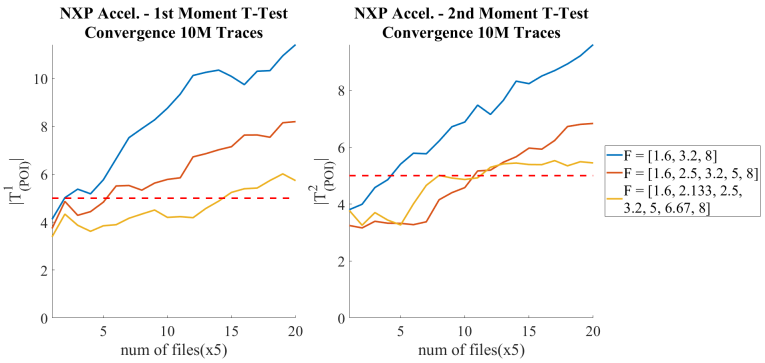

Figure: 17. NXP AES ACCEL.: Maximum absolute values over time of a T-Test detection test based on TVLA methodology vs. the number of collected traces and for different $\|F\|$ for a given BW (5.4MHz), $V_{min} = 0.58[V]$.

6.3 Gaussian-Template Based Attack Success-Rate

In this subsection our goal was to show how hard it is to make use of the information leakage measured by TVLA. As discussed in Subsection 6.1, our RAD-FS approach reduces attack-based metrics by orders of magnitudes. Thus, our goal was to perform model-less (profiled) evaluation utilizing templates. First, to reduce computational effort we have found POI's using SNR. Then, we profiled leakages in a subset of time samples using a Gaussian-Template model.

As shown in Fig. 18 for a 6-1.9MHz BW the attack's success-rate (SR) drops rapidly both with increased $\|F\|$ and reduced BW. As shown in Fig. 19, higher $\|F\|$ requires exponentially more traces for a successful extraction of key values, while the BW inversely affects the exponential growth constant.

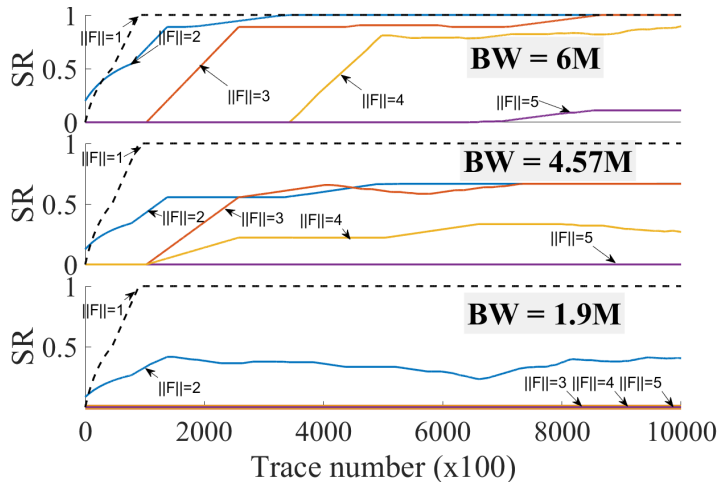

Figure: 18. Risc-V SW Tiny-AES: the success rate of a Gaussian template attack with varying $\|F\|$. Black dashed line - no DVFS 20MHz, Blue - $\|F\| = 2$, Orange - $\|F\| = 3$, Yellow - $\|F\| = 4$, Purple - $\|F\| = 5$.

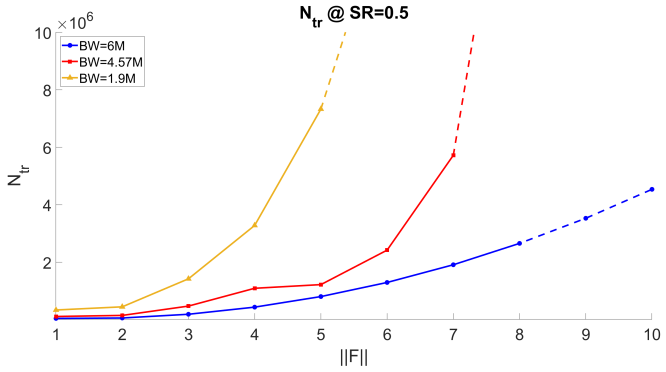

Figure: 19. Risc-V SW Tiny-AES: Number of traces required to achieve $SR = 0.5$ vs $\|F\|$ across different BW, dashed line denotes curve extrapolation

6.4 Timing attacks & P-States

In this Sub-section we show RAD-FS is very relevant for such network timing attacks mitigation. Results are demonstrated via. an ideal (optimal) oracle modeling the RAD-FS parameters. Our constructed oracle is very *generous* with how much control is given to the adversary: Prior-art discusses that bits-states affects DVFS algorithms selections. I.e., different combinations of 'on' and 'off' bits, in essence generate different workloads and allows adversaries to manipulate the DVFS power-control (and therefore also the timing) of a function call or code, in a scenario of chosen plaintext attack. This is the mechanism relied upon in [WPH+22], the adversary manipulates the Hamming-Weight (HW) of the plaintext, & the resulting time-to-encryption is measured. As our defense mechanism, we emulate a firmware based solution, that soft over-rides the frequency input from the OS and randomises it with our proposed RAD-FS in mind: I.e., the OS requests that the DVFS switches to some f_{OS} that matches a specific HW plaintext, serving as the *generous* (i.e., most sensitive) SCA Oracle while the firmware chooses a random frequency f_{RAD-FS} that has some relation to f_{OS}. This technique makes the vulnerability harder to exploit by introducing uniformly distributed noise (ideally), making the measured computation

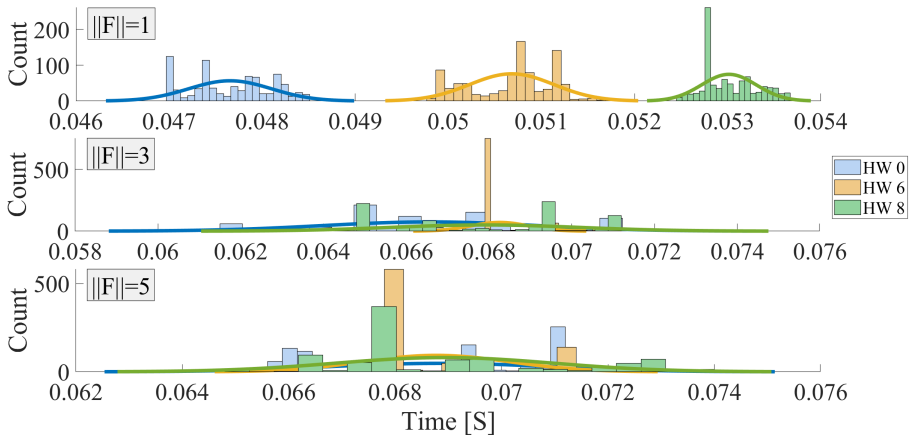

Figure: 20. Timing measured over 1K encryptions, performed with the same HW/freq./voltage. For a worst-case analysis, all plaintext bytes are the same for minimal noise. Measured with 10-100k traces per HW for increased $\|F\|$.

time (side-channel) harder to analyse. In the first experiment, the adversary measures the time to perform a sequence of steps: UART communication, change_clock_freq and 1000 encryptions. Fig. 20 shows that when RAD-FS is introduced at the firmware level, increasing $\|F\|$ size causes the distributions that are easily separable with $\|F\| = 1$, to alias unto one, and make it increasingly difficult for the adversary to discern them from one another. Expanding upon this, in the next scenario we measured the clock cycles required to perform a set number of encryptions(100), while changing the clock and performing UART communications. We compared changing the clock in 3 different intervals; 100, 50 and 10. Several interesting phenomena are visible in Fig.[21,Fig. 22,Fig. 23]: **(1)** in Fig. 21 even for an interval of 10 (many clock frequency changes), the distributions are easily distinguishable. This is important because without RAD-FS, the DVFS mechanism introduces a strong timing bias. **(2)** introducing RAD-FS shows a significant improvement for any interval. **(3)** going from Fig. 21 to Fig. 22 we can clearly see an increasing aliasing of the distributions & confirming our intended use case.**(4)**in Fig. 23 we achieve almost uniform distributions, with full aliasing, hampering the adversaries attempt to gain information via the side channel. Clearly, in this respect a uniform distribution (maximum entropy) is the best one can hope for.

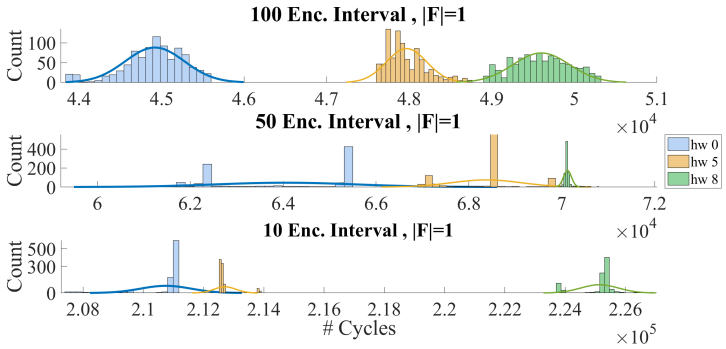

Figure: 21. Clock cycles measured over 100 encryptions, frequency changed with varying intervals. $\|F\| = 1$.

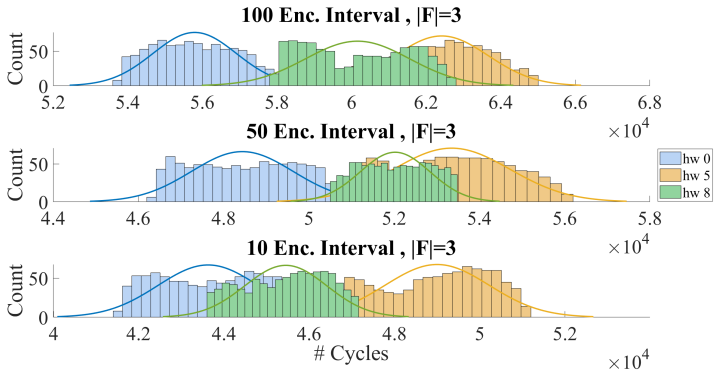

Figure: 22. Clock cycles measured over 100 encryptions, frequency changed with varying intervals. $\|F\| = 3$.

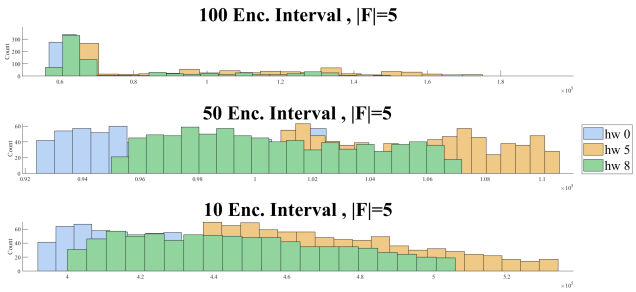

Figure: 23. Clock cycles measured over 100 encryptions, frequency changed with varying intervals. $\|F\| = 5$.

6.5 Impact on Performance

The RAD-FS ideal scenario performance-security wise is achieved with the highest possible mean frequency (f_{base}) for the encryption engine[4] concurrently with a large $\|F\|$ within a tight BW. Considering IoT applications, the main parameters for optimisation are area and power. For some notations, we denote by **(1)** OH - the performance overhead, i.e., the increase in computation time. **(2)** m - # encryptions done in a given f_n. **(3)** T_{switch} - the time to switch between $f_n, f_m \in F | n \neq m$ for $\|F\| = k$ a total of $\binom{k}{2}$ options. **(4)** f_i - a frequency in F. That is, the total encryption time can be written by $\frac{1}{f_i} \cdot \#_{cycles}$, considering #cycles clock-cycles.

$$\mathrm{T_{av}} = \frac{E[T_{switch_i}] + m \cdot E[\frac{1}{f_i} \cdot \#_{cycles}]}{m} \tag{5}$$

Generally, considering the specifications of commercial devices (and the PLS15 characteristics itself), the switching time and the enc.-time are of the same scale, $E_i[T_{switch}] \approx E_i[\frac{1}{f_i} \cdot E\#_{cycles}]$, denoted by E_T. Clearly, taking $\lim_{m\to\infty}$, $\mathrm{T_{av}} = E_T$ with 1 over m convergence. For example, taking the m=10 case:

$$\%\mathrm{OH} = \frac{T_{av}}{T_{f_{base}}} = \frac{1.1 \cdot E_T}{T_{f_{base}}} \tag{6}$$

Under our chip's limitation, for example, we can calculate the overhead using 2 frequencies inside a 1.33MHz bandwidth with the following parameters: $F = \{21.33MHz, 20MHz\}, ||F|| = 2, BW = 1.33MHz$ we achieve:

$$\%\mathrm{OH} = \frac{1.1 \cdot \frac{1}{k} \cdot \sum_{i=0}^{k-1} \frac{1}{f_i}}{T_{f_{base}}} = 1.136 = 13.6\% \tag{7}$$

For this configuration our RAD-FS mechanism endures a 13.6% latency overhead, which is rather efficient as compared to prior art as discussed below. Considering energy overhead, owing to the quadratic dependency of the operation frequency, F, in the operating voltage, V_{dd}. We can approximate the voltage OH by $\Delta\mathrm{V} \sim \sqrt{\frac{1}{\Delta\mathrm{T}}} = 0.938$, concluding in:

$$\mathrm{P^{OH}\%} = \frac{\Delta\mathrm{P}}{\mathrm{P}} = \frac{\Delta\mathrm{f} \cdot \mathrm{C_{eff}} \cdot \Delta\mathrm{Vdd}^2}{\mathrm{f} \cdot \mathrm{C_{eff}} \cdot \mathrm{Vdd}^2} = \frac{1.136 \cdot 0.938^2}{1} = 99\% \tag{8}$$

Under our chip's limited options, and taking the above estimations, we estimate the P^{OH} to be improved by 1%. Actual commercial devices numbers may even be better with more granular freq./voltage steps. For final comparison, we considered implementation effort which roughly estimates resources needed to implement different solutions and the flexibility offered by different approaches.

Table 2: *SW implementation, **HW analog implementation, ***HW analog & custom memory implementation

Source	$Area^{OH}$	L^{OH}	P^{OH}	$Effort^{OH}$	Type
[SKM+19]	+6.6%	+17.4%	-3.5%	**	HW
[GDS20]	+20%	0	+24%	**	HW
[KLS+21]	+10%	+0.07%	+8%	***	HW
Our Work	0	+13.6%	-1%	*	SW

As shown in Table 2, with L and P denoting latency and performance respectively, our solution aims at ULVT MCU's and IoT chips, where area and power are the main optimization concerns. Due to using a block that pre-exists within such devices, we present no area overhead, with low implementation effort, requiring only a SW implementation. Using a SW solution, we sacrifice latency within an acceptable margin[5] even when comparing to SOTA.

6.6 Results Summary

To summarise our results step by step: **(1)** we start by proving that our concept has merit in the context of preventing software timing attacks as shown by Fig. 25 and Fig. 26. The more uniform the frequency distribution gets, and the different P-states *overlap* increases, the entropy increases. **(2)** We continue deep diving into the different parameters such a technique would make use of, using various estimators and SCA algorithms. From Fig. 12 we conclude that CPA is weak estimator in our scenario, and therefore we focus on TVLA and the SNR as metrics. **(3)** We isolated the different variables, $\|F\|$, BW and f_{min}, into different sets of measurements to determine which has the greatest impact. Comparing results from Fig. 13 and Fig. 11 to Fig. 14: from the difference in the slope and distribution of measurements we can conclude that the best scenario requires as many frequencies as possible within as small a BW as possible. **(4)** Using TVLA (Figures 16, 17), we aimed to show two things. First, the loss of meaningful information in the measured traces, i.e., the two compared populations (set plaintext and randomised plaintext) in TVLA becomes harder and harder to distinguish to a meaningful extent, as evident by the higher number of traces needed to converge and the final converged T-value. Second, the shift of information to higher statistical moments, seen in Fig. 16 the only measurement showing no information in the second moment is the none RAD-FS measurement set,thus proving the computational complexity for an adversary introduced by RAD-FS. This property is logical as instead of measuring information in one leaked frequency with RAD-FS we generate a distribution and the information then modulates into higher orders. **(5)** Gaussian-Template based attacks - we chose a strong attack model to show that even in a scenario where the adversary is knowledgeable and captures enough traces to profile and attack, a successful attack is still none trivial. Fig. 19 shows that the required number of traces is exponential evidently with our modulation parameters, $N_{tr} \sim \exp(\frac{1}{BW} \cdot \|F\|)$. **(6)** Identifying the strength of RAD-FS in protecting against timing based attacks, we show clear aliasing in the time domain in a scenario where the adversary has a way to manipulate the DVFS mechanism either directly (via SW) or via an oracle (as done in [WPH+22]). By randomising the frequency we reduce the direct relation between workload and frequency, making it harder to discern information about the encryption from the run time. This is *important as several industry standard encryption schemes, both symmetric and asymmetric public key encryptions are vulnerable to such attacks.*

[4] As the computation is intensive and mapped to high workload

[5] This can be further improved upon by implementing RAD-FS in assembly for example

7 Assumptions and Real-life

We will now outline the features of our protection mechanism by identifying the essential requirements for a successful attack: **(1)** Perfect synchronization and measurement triggering - this is how the analysis is done in this paper/analysis. **(2)** Pre-knowledge on *Trace-Length* - when we randomize frequencies using RAD-FS, the adversary clearly does not know the number of samples needed to be captured, as it depends on the randomized frequency. Therefore, the best scenario is to take some fixed number of samples which will imply mixtures of frequencies appearing in the captured leakage even if f_n is only randomized once per several encryptions. **(3)** Isolating the relevant leakages is a lot harder in a parallel computation scenario. running on multiple cores will drastically increase the algorithmic noise and will impede the adversary's ability to filter out leakages not correlating with the hypothesis data manipulation ($\oplus, sbox$ etc.). In addition, it is important to emphasize that real-life applications with more conventional DVFS mechanisms are different: **(1)** DVFS resolution is very high - even hundreds of power/frequency- states [SGS+14, MWC17, ACK19, XM13]. Meaning the countermeasure security-parameters such as $\|F\|$, BW, f_{min} can be significantly optimized. **(2)** Synchronization in large embedded-SoCs featuring miniaturized IoTs is fairly complex, and any pre-processing trigger estimation will dramatically increase the noise. **(3)** Modern ADVFS solutions optimize each core individually depending on the workload, adding a plethora of protection flexibility. **(4)** As discussed in page 7, filtering attempts may eliminate information and induce overlaps thus increasing the noise.

8 Conclusions and Future Research

In this paper we demonstrate RAD-FS, a new security technique which is scalable, software based, and easy to implement, that applies to most if not all modern microchips. Improving protection against DPA and timing SCA attacks in orders-of-magnitudes. Discussing several different estimators under DPA, we show the radical effect achieved by RAD-FS on the analysed estimator and it's convergence. We show the different effects of our parameters, and conclude that the ideal scenario revolves around as large $\|F\|$ as possible within the smallest BW achievable considering system limitations, as to increase aliasing. Moreover, it substitutes the naive solution against timing attacks (Shutting off the ADVFS controller) (e.g., Hertzbleed) and enables a SCA-secure-chip coexisting with DVFS optimization. We discussed a countermeasure for the inherent weakness in power/freq.-data dependency.

In future work, we would want to test adding "fuzziness" to the RAD-FS process though uniformly distributed amount of "NOP" asm commands before/within the enc.-operation, to touch upon multi DVFS-models security analysis, and in depth electromagnetic SCA analysis with the proposed mechanism.

Acknowledgments. Itamar Levi and Daniel Dobkin were partially funded by the Israel Innovation Authority (IIA), Bio-Chip Consortium Grant file No. 75696, Israel, and Israel Science Foundation (ISF) grant 2569/21.

We sincerely thank the *PLSense* team (today acquired by NXP), namely, Uzi Zangi, Arnon Kaminsky, Shahar Dinar, Shaked Matzner and Tzach Hadas: for their invaluable support and guidance while working with their development environment and beta (development and testing) boards and setups, such assistance and dedication towards innovative science is not trivial for such a small start-up.

A Additional Graphs and Tables

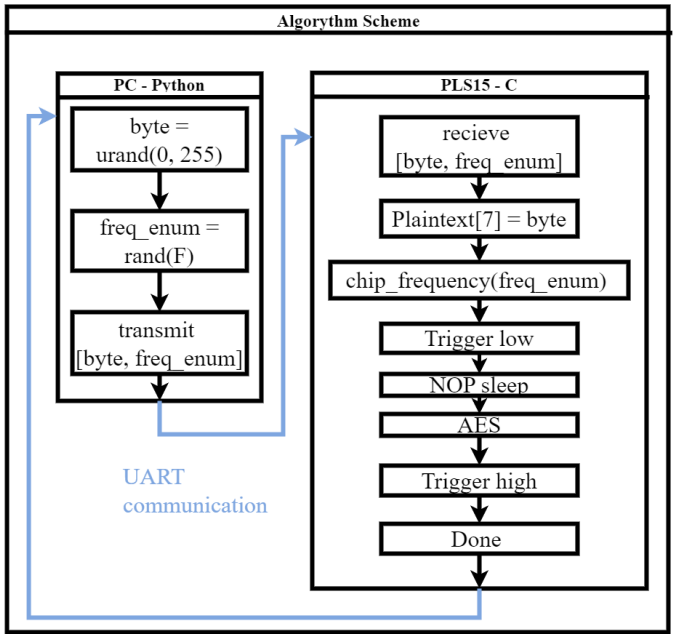

Figure: 24. Experiment pseudo-code Flowchart

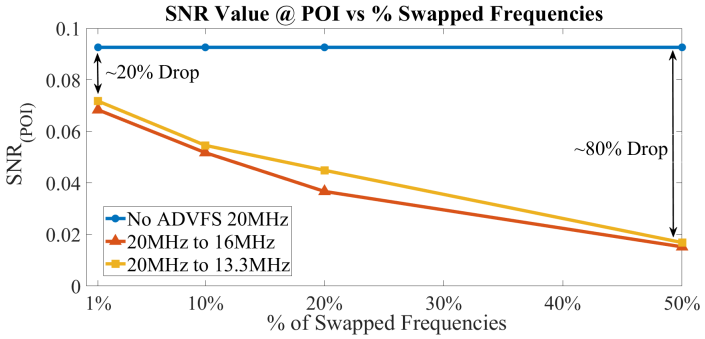

Figure: 25. Swapping from default 20MHz to a different core frequency for a percentage of the measured queries led to a reduction in the SNR POI value

Table 3: Classification of frequency map of F' to different BWs with a varying F_{min}

$\sim$const BW$\backslash F_{min}$	F_{min}	$\sim 2F_{min}$	$\sim 3F_{min}$	$\sim 4F_{min}$	$\sim 5F_{min}$	$\sim 6F_{min}$	$\sim 13F_{min}$
BW$\sim$3MHz	[5, 2.5]	[8, 5]	[13.3, 10]	[16, 13.3]		[20, 16]	
BW$\sim$5.3MHz	[6.67, 1.28]	[8, 2.857]	[8,3.3]	[10, 5]			[21.3, 16]
BW$\sim$8MHz	[10, 2.5]	[13.3, 5]	[16, 8]	[20, 10]	[21.33, 13.3]		
BW$\sim$11MHz	[13.3, 2.13]	[16, 4]	[16, 6.67]	[20, 8]	[21.33, 10]		

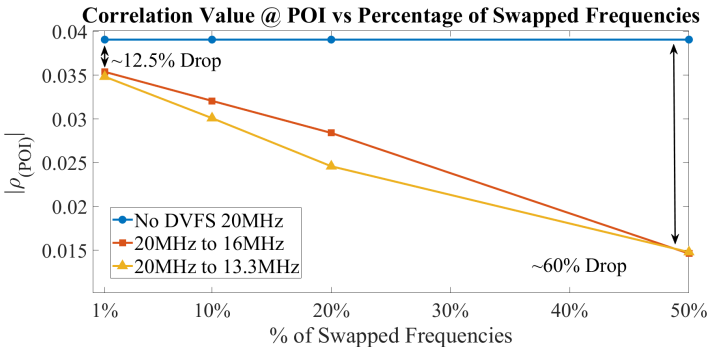

Figure: 26. Swapping from default 20MHz to a different core frequency for a percentage of the measured queries led to a reduction in the Corr POI value

References

[AARR02] Dakshi Agrawal, Bruce Archambeault, Josyula R Rao, and Pankaj Rohatgi. The em side—channel (s). In *International workshop on cryptographic hardware and embedded systems*, pages 29–45. Springer, 2002.

[ACK19] Bilge Acun, Kavitha Chandrasekar, and Laxmikant V. Kale. Fine-grained energy efficiency using per-core dvfs with an adaptive runtime system. In *2019 Tenth International Green and Sustainable Computing Conference (IGSC)*, pages 1–8, 2019.

[BCO04a] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *International workshop on cryptographic hardware and embedded systems*, pages 16–29. Springer, 2004.

[BCO04b] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *Cryptographic Hardware and Embedded Systems-CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings 6*, pages 16–29. Springer, 2004.

[BHC+16] Wenlei Bao, Changwan Hong, Sudheer Chunduri, Sriram Krishnamoorthy, Louis-Noël Pouchet, Fabrice Rastello, and P Sadayappan. Static and dynamic frequency scaling on multicore cpus. *ACM Transactions on Architecture and Code Optimization (TACO)*, 13(4):1–26, 2016.

[BSL22] Rinat Breuer, François-Xavier Standaert, and Itamar Levi. Fully-digital randomization based side-channel security—toward ultra-low cost-per-security. *IEEE Access*, 10:68440–68449, 2022.

[CDG+13] Jeremy Cooper, Elke DeMulder, Gilbert Goodwill, Joshua Jaffe, Gary Kenworthy, Pankaj Rohatgi, et al. Test vector leakage assessment (tvla) methodology in practice. In *International Cryptographic Module Conference*, volume 20, 2013.

[CGLS20] Gaëtan Cassiers, Benjamin Grégoire, Itamar Levi, and François-Xavier Standaert. Hardware private circuits: From trivial composition to full verification. *IEEE Transactions on Computers*, 2020.

[CRR02] Suresh Chari, Josyula R Rao, and Pankaj Rohatgi. Template attacks. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 13–28. Springer, 2002.

[CRR03] Suresh Chari, Josyula R Rao, and Pankaj Rohatgi. Template attacks. In *Cryptographic Hardware and Embedded Systems-CHES 2002: 4th International Workshop Redwood Shores, CA, USA, August 13–15, 2002 Revised Papers 4*, pages 13–28. Springer, 2003.

[Dwo01] Morris Dworkin. Recommendation for block cipher modes of operation. methods and techniques. Technical report, National Inst of Standards and Technology Gaithersburg MD Computer security Div, 2001.

[FDLZ14] Yunsi Fei, A Adam Ding, Jian Lao, and Liwei Zhang. A statistics-based fundamental model for side-channel attack analysis. *Cryptology ePrint Archive*, 2014.

[GDS20] Archisman Ghosh, Debayan Das, and Shreyas Sen. Physical time-varying transfer functions as generic low-overhead power-sca countermeasure. *arXiv preprint arXiv:2003.07440*, 2020.

[Ico23] Icons8. Scientific icons, 2023. https://icons8.com/ [Accessed: (1.06.2023].

[Inc22] Arm Inc. Arm big. little, 2022. https://www.arm.com/technologies/big-little [Accessed: (10.07.2023].

[KJJR11] Paul Kocher, Joshua Jaffe, Benjamin Jun, and Pankaj Rohatgi. Introduction to differential power analysis. *Journal of Cryptographic Engineering*, 1(1):5–27, 2011.

[KLS+21] Raghavan Kumar, Xiaosen Liu, Vikram Suresh, Harish K Krishnamurthy, Sudhir Satpathy, Mark A Anders, Himanshu Kaul, Krishnan Ravichandran, Vivek De, and Sanu K Mathew. A time-/frequency-domain side-channel attack resistant aes-128 and rsa-4k crypto-processor in 14-nm cmos. *IEEE Journal of Solid-State Circuits*, 56(4):1141–1151, 2021.

[LBBS20] Itamar Levi, Davide Bellizia, David Bol, and François-Xavier Standaert. Ask less, get more: Side-channel signal hiding, revisited. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2020.

[LBS20] Itamar Levi, Davide Bellizia, and François-Xavier Standaert. Beyond algorithmic noise or how to shuffle parallel implementations? *International Journal of Circuit Theory and Applications*, 48(5):674–695, 2020.

[LSH10] Etienne Le Sueur and Gernot Heiser. Dynamic voltage and frequency scaling: The laws of diminishing returns. In *Proceedings of the 2010 international conference on Power aware computing and systems*, pages 1–8, 2010.

[Man04a] Stefan Mangard. Hardware countermeasures against dpa–a statistical analysis of their effectiveness. In *Topics in Cryptology–CT-RSA 2004: The Cryptographers' Track at the RSA Conference 2004, San Francisco, CA, USA, February 23-27, 2004, Proceedings*, pages 222–235. Springer, 2004.

[Man04b] Stefan Mangard. Hardware Countermeasures against DPA – A Statistical Analysis of Their Effectiveness. In Tatsuaki Okamoto, editor, *Topics in Cryptology – CT-RSA 2004*, number 2964 in Lecture Notes in Computer Science, pages 222–235. Springer Berlin Heidelberg, 2004.

[MDS99] Thomas S Messerges, Ezzy A Dabbish, and Robert H Sloan. Investigations of power analysis attacks on smartcards. *Smartcard*, 99:151–161, 1999.

[MOP08] Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power analysis attacks: Revealing the secrets of smart cards*, volume 31. Springer Science & Business Media, 2008.

[MWC17] Xinxin Mei, Qiang Wang, and Xiaowen Chu. A survey and measurement study of gpu dvfs on energy conservation. *Digital Communications and Networks*, 3(2):89–100, 2017.

[NIC+23] Ben Nassi, Etay Iluz, Or Cohen, Ofek Vayner, Dudi Nassi, Boris Zadov, and Yuval Elovici. Video-based cryptanalysis: Extracting cryptographic keys from video footage of a device's power led. *Cryptology ePrint Archive*, 2023.

[PBB98] Trevor Pering, Tom Burd, and Robert Brodersen. The simulation and evaluation of dynamic voltage scaling algorithms. In *Proceedings of the 1998 international symposium on Low power electronics and design*, pages 76–81, 1998.

[PPC+15] Edson Luiz Padoin, Laércio Lima Pilla, Márcio Castro, Francieli Z Boito, Philippe Olivier Alexandre Navaux, and Jean-François Méhaut. Performance/energy trade-off in scientific computing: the case of arm big. little and intel sandy bridge. *IET Computers & Digital Techniques*, 9(1):27–35, 2015.

[QS01] Jean-Jacques Quisquater and David Samyde. Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In *International Conference on Research in Smart Cards*, pages 200–210. Springer, 2001.

[SGS+14] Bo Su, Junli Gu, Li Shen, Wei Huang, Joseph L. Greathouse, and Zhiying Wang. Ppep: Online performance, power, and energy prediction framework and dvfs space exploration. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 445–457, 2014.

[SKM+19] Arvind Singh, Monodeep Kar, Sanu K. Mathew, Anand Rajan, Vivek De, and Saibal Mukhopadhyay. Improved power/em side-channel attack resistance of 128-bit aes engines with random fast voltage dithering. *IEEE Journal of Solid-State Circuits*, 54(2):569–583, 2019.

[SL23] Dor Salomon and Itamar Levi. Masksimd-lib: on the performance gap of a generic c optimized assembly and wide vector extensions for masked software with an ascon-p test case. *Journal of Cryptographic Engineering*, pages 1–18, 2023.

[SM15] Tobias Schneider and Amir Moradi. Leakage assessment methodology: A clear roadmap for side-channel evaluations. In *Cryptographic Hardware and Embedded Systems–CHES 2015: 17th International Workshop, Saint-Malo, France, September 13-16, 2015, Proceedings 17*, pages 495–513. Springer, 2015.

[Sta10a] François-Xavier Standaert. Introduction to side-channel attacks. In *Secure integrated circuits and systems*, pages 27–42. Springer, 2010.

[Sta10b] François-Xavier Standaert. Introduction to side-channel attacks. *Secure integrated circuits and systems*, pages 27–42, 2010.

[TMC+23] Meltem Sönmez Turan, Kerry McKay, Donghoon Chang, Lawrence E Bassham, Jinkeon Kang, Noah D Waller, John M Kelsey, and Deukjo Hong. Status report on the final round of the nist lightweight cryptography standardization process. 2023.

[VCMKS12] Nicolas Veyrat-Charvillon, Marcel Medwed, Stéphanie Kerckhof, and François-Xavier Standaert. Shuffling against side-channel attacks: A comprehensive study with cautionary note. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 740–757. Springer, 2012.

[WPH+22] Yingchen Wang, Riccardo Paccagnella, Elizabeth Tang He, Hovav Shacham, Christopher W Fletcher, and David Kohlbrenner. Hertzbleed: Turning power {Side-Channel} attacks into remote timing attacks on x86. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 679–697, 2022.

[XM13] Kaiyong Zhao Xiaowen Chu Xinxin Mei, Ling Sing Yung. Gpu dvfs, 2013. https://dl.acm.org/doi/pdf/10.1145/2525526.2525852 [Accessed: (10.07.2023].

[ZBSF04] Bo Zhai, David Blaauw, Dennis Sylvester, and Krisztian Flautner. Theoretical and practical limits of dynamic voltage scaling. In *Proceedings of the 41st Annual Design Automation Conference*, pages 868–873, 2004.