

Static Power Side-Channel Analysis – An Investigation of Measurement Factors

Thorben Moos, Amir Moradi, and Bastian Richter

Abstract—The static power consumption of modern CMOS devices has become a substantial concern in the context of the side-channel security of cryptographic hardware. Its continuous growth in nanometer-scaled technologies is not only inconvenient for effective low power designs, but does also create a new target for power analysis adversaries. Additionally, it has to be noted that several of the numerous sources of static power dissipation in CMOS circuits exhibit an exponential dependency on environmental factors which a classical power analysis adversary is in control of. These factors include the operating conditions temperature and supply voltage. Furthermore, in case of clock control, the measurement interval can be adjusted arbitrarily. Our experiments on a 150nm CMOS ASIC reveal that with respect to the signal-to-noise ratio in static power side-channel analyses, stretching the measurement interval decreases the noise exponentially and even more importantly that raising the working temperature increases the signal exponentially. Control over the supply voltage has a far smaller, but still noticeable, positive impact as well. In summary, a static power analysis adversary can physically force a device to leak more information by controlling its operating environment and furthermore measure these leakages with arbitrary precision by modifying the interval length.

Index Terms—static power, leakage, side-channel analysis, side-channel attacks, ASIC, operating conditions, temperature, supply voltage, measurement interval

I. INTRODUCTION

Established cryptographic primitives usually come along with a set of mathematical security guarantees to inspire confidence in their resistance against state-of-the-art (crypt-analytic) attacks. These guarantees, however, are only valid against computationally bounded black-box adversaries. For instance, it is common to prove that no polynomial-time adversary stands a better-than-negligible chance to compromise the security of a said primitive¹ when being restricted to the observation of its inputs and outputs exclusively. Yet, in embedded contexts such a limitation is not respected, due to the constant physical exposure of the hardware to potential adversaries. Any adversary, who is capable of measuring the physical emissions of a cryptographic device during the execution of a primitive (in sufficiently high quality) has access to substantially more information than just the inputs and outputs. In particular, these observations can directly be correlated to intermediate values of the underlying cryptographic algorithms. This specific kind of adversary model invalidates the security claims of virtually all (raw/unprotected) cryptographic primitives, since they were simply not developed to withstand

attacks which make use of their intermediate results. Hence, in many cases it is possible to break the security of physical instances of mathematically secure primitives (e.g., recovering the fixed key of an AES implementation) by carefully observing the emissions of the executing device. To mitigate these threats, dedicated countermeasures which minimize the leakage of information through physical side-channels need to be applied when implementing cryptography in real-world devices.

The static power consumption of CMOS hardware (a.k.a. leakage power) is one type of observable physical characteristic that can be exploited as a side-channel. Due to its, historically speaking, smaller contribution to the overall power consumption when compared to the dynamic currents, it has rarely been considered in traditional side-channel analyses. In view of its exponential growth, however, which is directly linked to the down-scaling of the technology, it has attracted more and more attention over the last decade.

A. History of Static Power Analysis

Ever since the introduction of power analysis attacks in 1999 [23] researchers have concentrated almost exclusively on the exploitation of the operation- and data-dependency that can be observed in the dynamic power consumption of cryptographic hardware. However, in the year 2007 the authors of [16] provided the first concrete evidence for the fact that the leakage currents in modern CMOS gates exhibit a strong data-dependency as well. Additionally, they pointed out that the static power consumption had already reached a considerable dimension for sub-micron CMOS technologies by then. These discoveries consequently led to the first attempts to exploit the emerging new side channel. In [24] a DPA-based attack on (simulated) static power measurements using a single-bit power model is proposed. The works presented in [7] and [8] verify the soundness of the Hamming weight model in the static power domain and conduct a successful CPA attack. Further investigations revealed extensively that multiple DPA-resistant logic styles are rather ineffective against static power analysis [24], [5], [6], [22]. The results of [12] and [9] do even suggest that an unprotected CMOS implementation of the block cipher PRESENT-80 is less vulnerable to such attacks than the same cipher implemented in the DPA-resistant logic style WDDL. To cope with the issue of a possible exploitation of the static currents Zhu et al. proposed first countermeasures in 2013 [46] and 2014 [20]. Further ones have been suggested in the following years [19], [33], [9], [45], [44]. Even extensions to template [42], [11] and multivariate attacks [11], [15] exist in the literature and one first approach

T. Moos, A. Moradi and B. Richter are with the Ruhr-Universität Bochum, Horst Görtz Institute for IT-Security, Germany (e-mail: {firstname.lastname}@rub.de.)

¹by known methods or under a certain assumption

to combine the information leaked through the static and the dynamic power side-channel has been published [43].

However, apart from one small experiment in proof-of-concept manner which has been performed on an 8-bit register [16], [7], all evaluations, all countermeasures and especially all attacks in the previously mentioned articles are exclusively based on simulation results. The first contribution to this field where an analysis has been performed on actual leakage measurements, taken from a physical device, was published in 2014 [28]. Here, detailed information about the leakage currents of different FPGA elements in various process technologies is presented. Additionally, a successful key recovery on a masked and shuffled AES-128 implementation is performed by utilizing the higher-order moments of the static power consumption. The second work in this area with an experimental focus suggests that the ability to control the clock enables adversaries to arbitrarily reduce the noise in their measurements [35]. It was recently confirmed by practical experiments on a threshold implementation prototype chip that this possibility indeed poses a serious threat to algorithmic DPA countermeasures that require high noise levels, such as masking [27]. Additionally, a sophisticated measurement setup is introduced in [27] consisting mainly of a low-noise DC amplifier and a powerful climate chamber. The authors of [10] recently proposed another (distinct) measurement setup dedicated to static power analysis with the objective of being low-cost and demonstrated its suitability by an analysis of a crypto core on a 45nm Xilinx Spartan-6 FPGA. This setup is based on a DC pico-ammeter for trace acquisition and a commercial Peltier cell to control the temperature of the device under test.

B. Role of Operating Conditions

Regarding the role of operating conditions in attack scenarios it is usually (and consistently among all publications) stated that the temperature has to be kept constant during the analysis of the target. The reason for this is the exponential dependency of the static currents on thermal influences. Apart from this constraint the device under test is usually investigated under realistic conditions like room temperature and specified supply voltage. Some of the previously listed works include figures for the data dependency of static currents in CMOS logic gates under different working temperatures [16], [24], [8], [5], [33]. However, those numbers are based on library information of single standard cells and only presented to confirm the suitability of the Hamming weight model regardless of the applied temperature (as long as it is kept constant during the whole acquisition of a set of traces). Up to this point it was paid little attention to the fact that an increase of the temperature also increases the absolute difference between the leakage currents for the different possible input vectors to digital standard cells. Since the ratio between the currents for any two input vectors is roughly maintained and the absolute currents are exponentially increasing when more thermal energy is applied, the signal in a side-channel attack can – in theory – artificially be amplified in an exponential manner by raising the temperature. We evaluate the soundness of this hypothesis based on an experimental analysis in Section VI.

Another crucial parameter for the correct operation of integrated circuits is the supply voltage. To the best of our knowledge no work has investigated the influence of changes in the supply voltage on the exploitability of the static currents so far, neither in simulations nor in practice, even though strong dependencies of the static dissipation on potential-differences in CMOS transistors are known to exist. We discuss this in more detail in Section VI.

C. Related Work

The authors of [12] seem to exploit the impact of increased temperatures on the success of attacks for the first time. In [12] an implementation of the PRESENT-80 block cipher [13] is simulated in 40nm technology and its susceptibility to static power side-channel attacks is analyzed. The simulations are performed while the working temperature has been set to 100 °C and it is referred to this operating condition as the "worst-case scenario for the designer". However, a comparison to other temperatures is not presented.

The authors of [10] propose a low cost measurement setup based on a DC pico-ammeter and verify its suitability by performing a static power side-channel analysis of a PRESENT-80 implementation on a Spartan-6 FPGA. In their practical experiments the device under test (DUT) was heated up by a conventional Peltier cell to 65 °C. A comparison to other temperatures is not presented in this work either.

Subsequently, in 2017 two works have been published that utilize the temperature as a replacement for the missing time-dimension in static power analysis to perform multivariate attacks [11], [15]. While one of these works focuses on conducting template attacks which exploit the static power consumption [11], the other comes to the conclusion that increasing the working temperature progressively eases a static power side-channel analysis in terms of the required number of measurements for a successful key recovery [15]. The authors of [15] compare different temperatures between 0 and 100 °C. However, due to the simulation-based nature of those investigations the authors are forced to make assumptions about the noise and its very own dependency on the working conditions. From our point of view a natural assumption would be that the static currents of all non-targeted parts of the circuit, i.e., the algorithmic noise (see e.g. [40], [18] for descriptions of algorithmic noise), is affected in the same way as the targeted parts and that other noise sources, e.g., the electronic noise, the measurement noise or the quantization error, are less affected. We are not entirely sure which assumptions the authors in [11] and [15] make. In both works it is claimed that the signal-to-noise ratio (SNR) is fixed to a value of –60 dB regardless of the temperature. If this would be true (and the SNR is meant to be what is frequently applied as metric in the side-channel literature, c.f., Section V) the number of measurements to disclosure (MTD) should not significantly vary between the sets of measurements for different temperatures, since there exists a known anti-proportional relationship between the SNR and the MTD. In particular it was demonstrated in several articles, e.g., [18], that from the SNR alone the MTD value can be predicted. As a consequence one would expect a more

or less constant number of measurements required to recover the key for a fixed SNR. This is contradictory to the results presented in [15]. We believe what the authors actually did is fixing the amount of additive white noise over the simulations at different temperatures, i.e., fixing the mean and the standard deviation of the Gaussian distribution that is added onto the noise-free simulated power traces. This would also match the fact that in [15] it is claimed to be possible to extract the same amount of information "even in the presence of lower SNR", while the presented data is actually showing that it is possible to extract the same amount of information *in the presence of more noise* (exactly because the variance of the signal is amplified by the increased temperature). Regardless of the apparent misconception of the authors (or our inability to follow their interpretation) neither the claimed assumption, i.e., fixed SNR over all experiments, nor the assumption that is suggested by the data, i.e., fixed noise standard deviation and mean over all experiments, seems to be a good capture of the reality. In any case it has to be evaluated by practical measurements whether higher temperatures lead to a larger signal-to-noise ratio and therefore to a smaller number of measurements that are required to break an implementation.

D. Our Contribution

In this paper we try to close the gap between theoretical considerations regarding the influence of measurement factors on the feasibility of static power analysis attacks and their practical verification on actual hardware. We answer the question whether an adversary can physically force a device to leak more information by controlling specific operating parameters and provide informative numbers in this regard based on more than two months of non-stop measurements. In particular we have acquired 19 distinct sets with a cardinality of at least 5 million measurements per set in a controlled environment, each for a different temperature-voltage-combination (-20 to 90 °C, 1.62 to 1.98 V), which took roughly 2.7 days for each set. Afterwards, for the most effective temperature-voltage-combination (90 °C and 1.98 V), we recorded another 8 sets of traces for different lengths of the measurement interval. Our results show very clearly that, in this case study, increasing the temperature exponentially increases the signal, that increasing the supply voltage only marginally increases the signal and finally that increasing the measurement interval exponentially decreases the noise. Additionally, it becomes obvious that all three measurement factors can effectively be combined to lower the number of measurements that are required for a successful key recovery to a minimum. Control over these parameters – in theory – allows to eliminate any source of noise except for the algorithmic noise, which highly depends on the particular implementation as well as the concrete attack scenario and will always be present in power measurements [40]. Setup-wise we have built upon [27], but (1) improved the construction of the DC amplifier to obtain stable results at extreme temperatures, (2) built a custom low-pass filter, and (3) employed a simple post-processing technique. All these modifications have been verified to be useful in diminishing

the noise and improving the signal.

In short, our contribution can be summarized as follows:

- We propose a measurement setup dedicated to static power side-channel analysis which is applicable over a large range of operating conditions and detail how to avoid potential pitfalls.
- We show that by using such a setup static power side-channel adversaries can exponentially reduce the data complexity of attacks by controlling either the operating temperature or the measurement interval.
- We present results indicating that control over the supply voltage of the device under test marginally reduces the complexity as well.
- We conclude that combining all three measurement factors allows to obtain side-channel data that is free of any noise influence except for the algorithmic noise, which raises a warning flag for countermeasures that require a certain noise level to be effective.

This paper is organized as follows: In Section II we describe the measurement setup for the static power side-channel measurements. Section III introduces the targeted ASIC and the PRESENT block cipher implementation that is investigated. The measurement procedure for acquiring the traces (including the post-processing) is detailed in Section IV, while the necessary evaluation tools and metrics are introduced in Section V. Our main results are then presented in Section VI. Here we investigate the influence of each of the three measurement factors: temperature, supply voltage and measurement interval, on the exploitability of the data-dependent static currents. Finally we conclude our work in Section VII.

II. MEASUREMENT SETUP

In order to measure the static power consumption of our target ASIC, we inserted a precision 1 Ω resistor with low temperature coefficient into the Vdd path. In contrast to dynamic power measurements the amplifier cannot be AC coupled since AC coupling works as a kind of high-pass filter and would eliminate our static target signal (DC offset). Thus, common AC-coupled amplifiers like the ZFL-1000NL+ from Mini-Circuits cannot be used in this setup². Instead, the voltage drop over the resistor needs to be measured differentially and with a DC-coupled amplifier. There are two main problems when measuring the static leakage. At first, the voltage difference we would like to measure is very small, typically in the range of a few micro volts. To get an accurate measurement, a high DC amplification is needed. The second problem is the susceptibility to temperature variations. The static leakage itself is highly temperature dependent which results in huge shifts of the measured signal e.g., when the measurement room is accessed. Also, many amplifiers and differential probes suffer from a DC shift when they heat up during use. In [28] a LeCroy AP033 differential probe which features a

²Such an AC amplifier has been used in several dynamic power measurement setups, e.g., [17], [29], [30], [31], [38].




Fig. 1: Low-noise DC amplifier for static power measurements.

$\times 10$ amplification was used. While this probe is capable of measuring the signal with its high common DC offset, it only features a low amplification and is susceptible to thermal shifts in the measurements when the probe heats up during the long measurement procedure.

A. Low-Noise DC Amplifier

In order to overcome these drawbacks, we developed a sophisticated amplifier to measure the static leakage. A schematic of the amplifier and a photo of the employed board in its aluminum case can be seen in Figure 1. The first stage of the amplifier consists of an Analog Devices AD8421 instrumentation amplifier [1], which provides a very low temperature dependency with $0.2 \mu\text{V}/^\circ\text{C}$ maximum offset voltage drift and $1 \text{ ppm}/^\circ\text{C}$ gain drift. The amplifier is designed for measuring low level signals as it features an extremely low noise. In particular, it delivers $3 \text{ nV}/\sqrt{\text{Hz}}$ input voltage noise and $200 \text{ fA}/\sqrt{\text{Hz}}$ current noise with only 2 mA quiescent current [1]. This stage removes the common voltage between its two inputs which are connected to the two terminals of the shunt resistor and applies an amplification with a gain of 2. In contrast to [27] we did not make use of the adjustable offset of the instrumentation amplifier, but rather fixed it to a specific value, since the formerly employed potentiometer increased the noise in our measurements at higher temperatures. A second stage consisting of an Analog Devices AD8676 operational amplifier (op-amp) [2] applies a $\times 500$ amplification to the resulting signal (i.e. the DC amplifier achieves a total gain of $\times 1000$). This op-amp also has a low temperature dependency of $0.6 \mu\text{V}/^\circ\text{C}$ input offset drift and a noise of only $0.10 \mu\text{V}$ p-p (0.1 Hz to 10 Hz) [2]. The PCB of the amplifier is housed in a custom aluminum case which provides SMA connectors. Due to the high gain, the bandwidth of the amplifier is below 20 kHz which does not pose a problem since we are working with static signals³.

B. Low-Pass Filter

During the measurement procedure we observed some high-frequency noise in the measurements which increased for higher temperatures. Hence we built a custom low-pass filter to remove these portions of the signal and connected it

³Details of the developed amplifier (schematic, PCB layout) are accessible through the authors' webpage.




Fig. 2: Third-order (Butterworth Pi) LC low pass filter with cutoff-frequency of $\sim 100 \text{ Hz}$.




Fig. 3: Exemplary depiction of two measurements, one with and one without low-pass filter, for a measurement interval of 50 ms . Time period T_1 (here 20 ms) corresponds to the interval that is ignored due to the memory effect. Time period T_2 (here 50 ms) corresponds to the measurement interval. All values that are measured in T_2 are averaged to a singular static power value.

between the output of the DC amplifier and the input of the oscilloscope. The filter, which is shown in Figure 2, is built as a passive third-order Butterworth Pi LC construction to provide a cutoff-frequency (-3 dB) of approximately 100 Hz for a 50Ω input impedance of the oscilloscope. A visual impression of its effect is given in Figure 3 by means of one sample measurement with and one without the low-pass filter applied. The time periods denoted as T_1 and T_2 are introduced in Section IV. As one can see in Figure 3 the amplitude of the oscillating static signal is far smaller when the low-pass filter is applied, especially in relation to the voltage difference between the regions before and after the clock is stopped. This modification of the setup alone already reduced the measurement interval to reach a certain signal-to-noise ratio by a factor of about 5 when operated at a temperature of 90°C .

C. Evaluation Board

The Side-channel Attack Standard Evaluation Board (SASEBO-R) [3] that we used for our experiments was specifically designed to evaluate the security of cryptographic hardware implementations against side-channel attacks. The board provides a socket for an ASIC prototype that is connected by a 16-bit bidirectional data bus as well as a 16-bit address signal for control and communication purposes. We are able to control the mounted ASIC by a Xilinx Virtex-II Pro FPGA, which itself is connected to a 24-MHz oscillator. Since measuring small signals over long wires can induce measurement errors, we kept the distance between the shunt resistor and the amplifier short by designing the housing of our developed amplifier in such a way that it can be plugged directly on top of the SASEBO-R board by the SMA connectors.

D. Oscilloscope

We used a Teledyne LeCroy HRO 66zi oscilloscope for the measurements. This scope provides a true 12-bit ADC, a maximum sampling rate of 2 GS/s, and a maximum bandwidth of 600 MHz.

E. Climate Chamber

To quantify the influence of the temperature on the quality of our side-channel acquisitions we performed the static leakage measurements inside a CTS climate test chamber of series C-40/100 with 100 litres test space capacity. The chamber achieves temperatures between -40°C and $+180^{\circ}\text{C}$ as well as a temperature change rate of 5 K/min for cooling and 3 K/min for heating. It can hold the temperature with a variation of 0.3°C at a maximum thermal load of 1200 W at 20°C . This should highly suffice for our purposes as the target is not expected to radiate a considerable amount of heat (resulting in even smaller temperature variations). We placed the SASEBO-R board together with the mounted ASIC prototype and the DC amplifier inside the chamber, whereas the low-pass filter, the oscilloscope and the power supply units for the board and the amplifier have been placed outside of the chamber. In this regard we had to put two power supply cables for the amplifier and one for the board, as well as one SMA coaxial cable for the amplified static power signals, an RS-232 cable for the communications and a trigger probe cable through a vent in the chamber that was carefully sealed with silicone foam.


The full setup including all main components is depicted in Figure 4 as a photograph and in Figure 5 as a schematic.

III. TARGET

The target for our experiments is a 150nm CMOS ASIC prototype chip with a nominal supply voltage of 1.8 V. A photo of the prototyped chip is shown by Figure 6. Among 5 other cores the chip features the PRESENT-80 block cipher realized as a 3-share threshold implementation [34]. Although this work is not focusing on masked implementations or



(a) outside



(b) inside

Fig. 4: Photographs of the complete setup including the DC amplifier, the low-pass filter, the board, the oscilloscope, the climate chamber and some power supply units.

higher-order attacks, we chose this specific core for our investigations to make the results comparable to what has been reported in [27]. In this regard we treat the core as a regular unprotected PRESENT-80 implementation by setting all masks to zero, which corresponds to the `PRNG OFF` mode of operation described in [27]. By doing this we make sure that the core operates deterministically, i.e., for identical plaintexts all intermediate values and the shared output are identical as well. This is explained in more detail in Section VI.

PRESENT-80 is an ultra-lightweight block cipher (ISO/IEC 29192-2:2012 standard) that operates on a block size of 64 bits and a key length of 80 bits and consists of 31 computation rounds [13]. The term threshold implementation refers to a masking scheme based on Boolean secret sharing and multi party computation that implements non-linear functions of symmetric block ciphers efficiently in such a way that provable security against first-order power analysis attacks can be guaranteed, even in the presence of glitches [32]. The specific application of this scheme to the PRESENT-80 block cipher is introduced in [34]. Our investigated ASIC core implements the profile 2 of [34]. This profile refers to a serial implementation of PRESENT-80 with a shared data path (with 3 shares) but an unshared key schedule. A schematic of the nibble-serial architecture can be seen in Figure 7.

All intermediate values and data buses are 4-bit wide. As the graphics show, the S-box – which has an algebraic degree




Fig. 5: Schematic of the overall experimental setup including the DC amplifier, the low-pass filter, the oscilloscope and the climate chamber.




Fig. 6: ASIC prototype with 6 cores in 150 nm CMOS [27].

of 3 – is decomposed into two non-linear quadratic functions F and G . Those 4-bit boxes are then split into 3 shares each. The three G -boxes are processed at the same time in the ASIC and each of them receives 2 inputs out of the 3 data shares. The corresponding outputs are stored into registers. Afterwards, the three F -boxes are evaluated in parallel. The 4-bit words of the round state are processed in a pipelined manner by one instance of the shared S -box. Thus, (due to the register between the F and G functions) 17 clock cycles are required to evaluate the complete substitution layer of the cipher for one round. After the last nibble of the shares has been processed, the outputs are routed according to the linear layer (PLayer) of the cipher and saved into the register again. Therefore, each full computation round of the PRESENT-80 cipher takes 18




Fig. 7: Nibble-serial architecture of the PRESENT-80 threshold implementation core.

clock cycles on the investigated ASIC core.

The initial masking of the input (with all zeros in our case) as well as the unmasking of the output are performed on the chip itself. Hence the communication with the ASIC is performed in an unshared, conventional manner.

IV. MEASUREMENT PROCEDURE

In order to measure the static currents, we executed the following procedure. At the specific clock cycle, where the targeted intermediate value is processed, the clock signal of the PRESENT core is stopped and all other input signals to the ASIC are kept constant at a deterministic value. This idle state of the target is held for an arbitrarily long time interval during which the static power consumption of the device can be measured before the clock signal is switched back on. Thus, in our experiments recording the static leakage traces requires a stronger attacker model than it would be required for a classical power analysis, as full control over the clock signal is necessary. The power consumption values that are obtained in the mentioned time interval are then averaged to a singular value. Since the leakage currents are not supposed to change during that period, all occurring variations are noise and can be averaged out. This technique is called intra-trace averaging and constitutes one major advantage of static power analysis in comparison to classical attacks when control over the clock signal is obtained (see [35]). Due to the very high gain of our developed DC amplifier ($\times 1000$) and the very low cutoff-frequency of our low pass filter (100 Hz), a significant impact of the memory effect (described in [30]) on the measurement quality can be observed. The sudden drop of the power consumption when the clock signal is stopped influences the measured static power values for up to the next 20 ms. Hence the first 20 ms of the idle state are discarded and not included in our measurements. After that period, the actual measurement interval starts. This procedure is illustrated in Figure 3 exemplarily for a measurement with and a measurement without low-pass filter applied. The time period which is denoted by T_1 corresponds to the first 20 ms of the idle state that are discarded due to the influence of the memory effect. The second time period T_2 indicates the measurement interval.

In contrast to [27] we also employed a simple post-processing technique. The idea is to filter out the long-term temperature-induced variations of the static power




Fig. 8: Illustration of the post-processing technique for 100 exemplary measurements at 20 °C, 1.8 V, measurement interval of 10 ms and window-size of 8.

consumption over time, i.e., over a set of measurements (as opposed to noise that is included in single measurements). Quite obviously the climate chamber requires a lot more activity of its regulation units to maintain a constant temperature when it is set to a value far above or below the temperature of the room it is located in. These activities can be observed as low frequency noise along the whole set of measurements. Our post-processing step is therefore to apply a simple moving-average filter onto the measurement set by using the Matlab function `filter()`⁴. The effect of such filtering is depicted in Figure 8. The blue plot in Figure 8a corresponds to a set of 100 unaltered measurements as they were recorded from the oscilloscope. The red curve corresponds to the moving-average that is generated by the Matlab `filter()` function. The subtraction of the moving-average from the original measurements results in the black graph in Figure 8b and constitutes the resulting measurement set after the post-processing. We tested several window-sizes as parameter for the `filter()` function and revealed that a window-size of 8 leads on average to the best results on our measurements. In general we observed that for measurements with a long measurement interval, i.e., the more time-demanding ones, a smaller window-size led to optimal results, while for the measurements with a short interval larger windows were more successful. However, to keep all results comparable we have always used a window-size of 8 for the experiments that are presented in Section VI. Although the initial purpose of the post-processing was to improve the measurement quality at extreme temperatures (like -20 °C or 90 °C) we observed that it has a positive influence on the measurements in all cases, even at room temperature.

Figure 9 depicts the whole measurement procedure as a

⁴We also tested other filters in Matlab, for example a butterworth high-pass filter, but achieved inferior results.




Fig. 9: Flowchart describing the measurement procedure. simple flowchart.

V. EVALUATION TOOLS AND METRICS

Section VI analyzes in detail how the investigated measurement factors influence the amount of information that is included in, or can be extracted from, the corresponding side-channel measurements. In this regard several evaluation tools and metrics which are common in the side-channel analysis literature are used. We introduce these tools and metrics shortly in this section.

A. Signal-to-Noise Ratio

The signal-to-noise ratio (SNR), introduced by Mangard in 2004 [25], is one of the most common metrics to quantify

the quality of side-channel measurements and to determine the points of interest in a dynamic power trace. The corresponding formula is given in Equation 1.

$$SNR = \frac{Var(Signal)}{Var(Noise)} \quad (1)$$

The variance of the signal is defined as any variation in the measurements (e.g., power consumption or electromagnetic radiation) that is caused by the targeted intermediate value, while the variance of the noise describes all further variations in the traces that are not caused by this value. To assess those parameters for a specific sample point in a set of traces (acquired for random inputs), one has to sort the traces into a number of groups corresponding to the specific value the targeted intermediate result attains (e.g., 16 distinct groups for a 4-bit intermediate value). In the case of (SPN-based) block ciphers, for example, which make use of a bijective non-linear mapping and key addition, one can directly calculate the SNR for the intermediate values after the first round (resp. before the last round) from the input (resp. output) of the cipher. The variance of the signal is then calculated as the variance of the means of the individual groups, while the variance of the noise can be calculated as the overall mean over the variances of the individual groups.

B. Correlation Power Analysis

Correlation power analysis (CPA) was introduced by Brier et al. at CHES 2004 [14] to overcome some drawbacks of classical DPA. The main advantage is that a power model can be used to create a hypothesis for the leakage of a full intermediate value instead of targeting only a single bit at a time. This hypothetical power consumption is then compared to the actual power consumption by means of a (Pearson) correlation coefficient, which measures the linear dependency. The corresponding formula for two discrete vectors X, Y is given in Equation 2. The mean of the two vectors is denoted by \bar{X}, \bar{Y} .

$$\rho = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}} \quad (2)$$

In a successful attack the highest correlation coefficient directly translates to a correctly guessed part of the key. Common models, in addition to the identity model, include for example the Hamming weight of an intermediate result or the Hamming distance between two processed values.

C. Measurements to Disclosure

Historically, the number of traces that are required to perform a successful attack on an implementation (DPA, CPA, ...) has been the most common metric to assess the resistance of a device against such attacks. Nowadays there exist tools, like for example the non-specific Welch's t -test (see [37] for a detailed methodology), that are able to evaluate the leakage of a device without performing any specific attack and without being dependent on a correct choice of a leakage model. However, those metrics fail to provide information

about the hardness of an actual key recovery, which makes them unsuitable for a variety of purposes. In our case, for example, a suitable model and a successful attack are already known and the goal is to evaluate how large the impact of changes in a specific operating parameter on the exploitability of the implementation is. In this case the number of required measurements to disclose the correct key is still the most preferable metric. To the best of our knowledge the term "measurements to disclosure" together with its abbreviation "MTD" has been first used and defined as a metric by Kiri et al. at CHES 2005 [41]. The authors describe it as the cross-over point between the correlation coefficient for the correct key and the maximum correlation coefficient among all wrong key guesses when plotting the coefficients for all key guesses over the number of samples considered. For a relation between the SNR and the MTD see for example [18].

D. Success Rate

The success rate of a power analysis attack is the probability that the attack succeeds in recovering the correct key candidate by isolating it from a restricted set of key guesses [39], [36]. The most straightforward option to evaluate the success rate of an attack is to simply perform the attack multiple times. Many efforts have been devoted to the exploration of more efficient ways to estimate the success rate rather than this empirical one (the interested reader is referred to [39], [36]). However, in this work we do indeed perform the attack multiple times on disjoint subsets of a larger measurement set whenever success rates are reported

VI. MEASUREMENT FACTORS

In this section we present measurement results that have been acquired over a time period of roughly four months and represent the equivalent of more than two months of non-stop data acquisition⁵. We build upon the results that are reported in [27] and try to improve the attacks on unprotected implementations in terms of the required number of measurements by controlling the operating parameters. While the influence of the measurement interval on the noise has been mentioned in [35] and [27], although not as detailed as in this work, the influence of the operating conditions temperature and supply voltage has not yet been reported based on practical side-channel measurements. In order to keep all results comparable we target the same threshold implementation prototype chip as [27] and operate the same PRESENT core.

All reported values for the measurements to disclosure (MTD) metric refer to a standard Correlation Power Analysis (CPA) attack [14] on the combined Hamming weight of the outputs of the three F-boxes that can be seen in Figure 7 (4 bit \times 3 boxes = 12 bit), effectively targeting one key nibble (4 bit) at a time. The Hamming weight model has been proposed and verified as a suitable power model for static power side-channel attacks in [7] for the first time. Other popular models from the dynamic power domain that rely

⁵Between the acquisition of the different sets the ASIC and parts of the setup have to rest at target climate to adopt the temperature accordingly before the next set can be recorded.

on transitional effects, such as the Hamming distance model for example, are not applicable since static power attacks capture a stable state and not a transition between states. Let x_1, x_2, x_3 be the outputs of state register 1, 2 and 3 in Figure 7 respectively. Let k be the targeted nibble of the round key. The power model is then computed as shown in Equation 3

$$\begin{aligned} & \text{HW}(F_1(G_3(x_1 \oplus k, x_2), G_2(x_1 \oplus k, x_3))) + \\ & \text{HW}(F_2(G_3(x_1 \oplus k, x_2), G_1(x_2, x_3))) + \\ & \text{HW}(F_3(G_2(x_1 \oplus k, x_3), G_1(x_2, x_3))) \end{aligned} \quad (3)$$

The leakage currents of all parts of the circuit that are not directly targeted, such as the G-boxes and the state registers contribute to the algorithmic noise. To perform this kind of attack on the targeted threshold implementation core it is required to have knowledge of the masks that are involved in the computation, which a regular power analysis adversary against a securely implemented masking scheme would (ideally) not have. But, as mentioned before, we operate the core by setting all masks to zero, which allows us to predict the exact intermediate values, under the correct key hypothesis, that are actually processed by the circuit. This corresponds to the usual adversarial situation when targeting an unprotected PRESENT implementation, with the only difference that each S-box output corresponds to a 12 bit value instead of a 4 bit one (which certainly eases the attack). We would like to stress here that the whole purpose of our practical evaluation in this section is to investigate the influence of the measurement factors on the success of power analysis attacks. We do in no way claim that the presented attacks on the targeted implementation with fixed masks are a realistic scenario for any adversary against a real-world device. We just aim for conformity with the previous work described in [27] and restrict all our claims to unprotected implementations.

For all three measurement factors that are investigated in this section we provide estimations of the noise to determine whether it is influenced by the altered parameters. Additionally we report the number of measurements that are required for a successful recovery of one key-nibble by means of a CPA attack. From both of those values the influence of the operating parameters on the signal becomes obvious. For instance, due to the known anti-proportional relationship between the signal-to-noise ratio (SNR) and the measurements to disclosure (MTD) [18], a constant noise level and a lowered MTD indicate an increased signal. Similarly a decreased noise level and a constant MTD indicate a decreased signal. Estimated signal values would therefore be redundant. Furthermore the signal-to-noise ratio itself is mostly used to determine the points of interest in a dynamic power trace, which is not required in static power analysis attacks like the ones reported in this work, since each trace corresponds to a single measured value anyway. The only reasonable use case of the SNR in this context would be to identify which intermediate value is possibly leaked in a particular clock cycle, in order to perform the attack on this value or find the exact clock cycle that needs to be targeted. However, this is not a necessity in our experiments as we have detailed information about the




Fig. 10: Visual depiction of our technique to use disjoint subsets of one large set of traces to disclose the key multiple times. The shown MTD values (achieved for 90 °C, 1.98 V, 10 ms) from left to right are 85 000, 28 000, 181 000, 59 000, 143 000.

implementation and know exactly which intermediate value is processed within which clock cycle.

At all temperature-voltage combinations we have collected at least 5 million traces with a constant measurement interval of 10 ms. In all cases we have tried to disclose the correct key nibble multiple times using this set. For example if a CPA succeeds by indicating the highest (absolute) correlation value for the correct key candidate after roughly 150 000 traces and this correlation value remains to be the largest one among all candidates for at least further 20 000 traces (to avoid false positives) we state that the first MTD value for the attack on this set is 150 000. Then we ignore the first 170 000 traces and repeat the process from trace 170 001. In other words we use disjoint subsets of the whole measurement set to disclose the key multiple times in order to obtain multiple MTD values. This enables us to report average MTD values and success rates for each set (only if the average MTD lies below 2.5 million, otherwise only one disclosure is possible). In case of the most effective temperature voltage combination and a measurement interval of 10 ms for example we were able to disclose the correct key 48 times using 5 million traces. We have illustrated this procedure for 5 consecutive MTD values in Figure 10. The average number of measurements to disclosure (MTD) for this snippet of traces would be 95 200 and the number of measurements to reach a success rate larger than 50 % would be 85 000. Admittedly, the single MTD values may not be fully independent, although being computed on disjoint subsets of the whole set, because strictly speaking they do not originate from statistically independent experiments. However, we have observed that assuming the independence of the particular subsets leads to sound and reproducible results. Furthermore, as apparent from Figure 10, the MTD values can vary quite significantly. Thus, we believe that averaging several of those values and reporting success rates leads to more meaningful results and is therefore superior to only reporting single MTD values (like for example in [27]).

The 8 additional sets of traces that we have recorded at the most effective temperature voltage combinations, but for different measurement intervals (1 ms to 200 ms), have a smaller cardinality (because of the significantly longer run time for larger intervals). Hence, unfortunately, we are only able to present a single MTD value per set here.

A. Factor: Temperature

According to [21] all leakage effects that are based on solid-state physics, such as subthreshold leakage and diode currents, show extreme thermal dependencies. The subthreshold current for example, which is the dominating source of static power consumption in our technology, depends exponentially on the temperature [21]. From simulated measurements it is known that the factor between the leakage currents for any two input vectors to digital standard cells is roughly maintained [16], [24], [8], [5], [33]. Hence the difference between the classes in a power analysis attack should increase when raising the temperature, which corresponds to an increase of the signal in a side-channel attack. For most adversaries against embedded systems it should be feasible to influence the temperature of the environment, since physical access is part of the adversary model. In our experiments we put the target device into a climate chamber, like explained in Section II, and fixed the temperature and the relative humidity to a determined value. The supply voltage was kept at 1.8 V, which is the nominal supply voltage of the chip. In the range of -20°C to 0°C we raised the temperature in steps of 5°C between different sets and were not able to control the humidity. In the range of 0°C to 80°C we raised the temperature in steps of 10°C and kept a constant relative humidity of 20 %. Finally we measured two sets for 85°C and 90°C at a constant relative humidity of 10 %. The humidity was always set in order to have a small impact on the experiments. In general we chose a rather dry climate in order to not face any problems with condensation of water vapor at the electronic components. Furthermore, we observed that the climate chamber is able to keep the temperature more stable when the relative humidity is controlled as well. Since the absolute humidity increases when the relative humidity is kept at a fixed value and the temperature is increased we needed to reduce the relative humidity from 20 % to 10% for the temperatures above 80°C .

As explained above, we expect the signal to increase for higher temperatures. However, this only eases an attack if the noise is not equally (or greater) affected. Hence as a first step we estimated the variance of the noise for all of the 19 temperature sets, according to the description in Section V over the first 50 000 measured values. The results can be seen in Figure 11. Apparently the noise increases approximately in a linear fashion with the temperature. To emphasize this a first degree (linear) polynomial curve was fitted to the data. At a first glance, this behavior appears to be a negative result when trying to increase the exploitable information in the measurements by raising the temperature. On the other hand, this result was expected, since not only the static currents associated with the targeted intermediate value are increased, but the algorithmic noise is supposed to grow in a similar fashion. Additionally, it can be expected that the measurements include more thermal noise when raising the temperature. The important question is here whether the increase of the signal is large enough to overcome this linear increase of the noise. To answer this question we performed CPA attacks using the Hamming weight of the 12-bit output of the F-boxes. As explained before, each set has a cardinality




Fig. 11: Estimated noise standard deviation for temperatures between -20°C and 90°C .




Fig. 12: Number of measurements required to overcome a success rate of 50 % for temperatures between 10°C and 90°C .

of (at least) 5 million measurements. Whenever possible we tried to disclose the correct key candidate multiple times by using disjoint subsets of the whole measurement set. As a result we are able to report how many traces are required to reach a certain success rate, as shown by Figure 12. Please note that no temperatures below 10°C are plotted here. The reason for this is that for these temperatures the correct key candidate could not be disclosed even a single time with 5 million measurements. For the higher temperatures it becomes obvious that the number of measurements that are required to extract the same amount of information is reduced in an exponential manner. In order to emphasize this we have added an exponentially fitted curve. The same kind of exponential decrease can be observed in Figure 13 for the average MTD values. Despite of the linear increase of the noise for increasing temperatures the MTD values are exponentially decreasing. This confirms that the signal is exponentially increased by raising the temperature and more generally that the signal-to-noise ratio is exponentially increased. Since the signal is




Fig. 13: Average number of measurements required to disclose the correct key candidate for temperatures between 10 °C and 90 °C.

growing much stronger than the noise, it is – in theory – possible to raise the temperature up to a point where each noise source apart from the algorithmic noise becomes negligible.

B. Factor: Supply Voltage

Apart from the obvious linear dependency of the static power consumption on the supply voltage ($P_{leak} = I_{leak} \cdot V_{DD}$) there are several other dependencies listed in [21]. The gate leakage for example is doubled when the supply voltage is increased by 100 mV. The drain induced barrier lowering (DIBL) effect leads to a reduction of the effective threshold voltage when the supply voltage is increased [21]. This on the other hand increases the subthreshold conduction exponentially. Finally, the gate induced drain leakage (GIDL) effect increases the junction current exponentially in a specific region of the supply voltage. However, we do not expect a very large dependency of the static currents of our 150 nm ASIC on the supply voltage for mainly two reasons. First of all the gate leakage and the junction current are (more or less) negligible sources of static dissipation in technologies larger than 100 nm. Secondly the DIBL effect is only relevant for short channel length as well and even for a 65 nm technology the supply voltage needs to be increased by 192 mV to raise the overall leakage by 10% [21]. In total we expect at least a linear decrease in the number of required measurements to disclosure when increasing the supply voltage. Concerning the feasibility of controlling the supply voltage of the device under test it should be kept in mind that a regular power analysis adversary is supposed to place (or find) a resistor in the GND or V_{DD} path of the device under test and to measure the voltage drop across this component by means of a digital sampling oscilloscope. Thus, it should be easily possible to influence the supply voltage with this kind of capabilities. In our experiments we adjusted the supply voltage by a potentiometer in the feedback path of the linear voltage regulator on the measurement board and verified the correct




Fig. 14: Estimated noise standard deviation for supply voltages between 1.62 V and 1.98 V and temperature of 90 °C.




Fig. 15: Number of measurements required to overcome a success rate of 50 % for supply voltages between 1.62 V and 1.98 V and temperature of 90 °C.

setting while the setup was present in the climate chamber at target temperature.

Since it seemed infeasible to determine the influences of the supply voltage for all temperatures, we chose the most successful temperature in terms of success rate, i.e., 90 °C, and changed the supply voltage by 10% in comparison to the nominal supply voltage in both, positive and negative direction. We did not evaluate more extreme changes in order to not damage the chip. The results for the noise estimation at 90 °C can be seen in Figure 14. Almost no voltage-induced change in the noise level can be observed. If at all, the noise is slightly increased. When taking a look at the MTD values in Figure 15 and Figure 16 it can be seen that the attacks become slightly more successful by raising the voltage, but the effect is much less drastic than what could be observed for the temperature. Obviously, when taking only three data points into account, it is difficult to make a statement about the type of dependency that could be observed, which is why we leave this open to interpretation. Finally, please note that, in contrast to the temperature, the supply voltage also has a direct (quadratic) influence on the dynamic power consumption.

C. Factor: Measurement Interval

In [27] a clear trade-off between intra-trace averaging and inter-trace averaging is observed. In particular, by stretching the measurement interval the noise can be reduced and the attacks succeed with fewer traces, but the time to acquire a specific number of traces is also increased. During the period when the clock signal is stopped, none of the intermediate values of the computation that are currently processed by the circuit are supposed to change. Thus, the same transistors in the respective CMOS gates are active or inactive and the leakage currents of the whole chip should stay the same. All variations that can be observed in the measured signal during




Fig. 16: Average number of measurements required to disclose the correct key candidate for supply voltages between 1.62 V and 1.98 V and temperature of 90 °C.




Fig. 17: Estimated noise standard deviation for measurement intervals between 1 ms and 200 ms at a supply voltage of 1.98 V and a temperature of 90 °C.

that period are electronic noise (e.g., from the power supply or conducted and radiated emissions [26]) and by extending the time interval that is averaged into a single value the influence of the noise can be minimized. Note that this can not be achieved solely by averaging over more sample points, for example when increasing the sampling frequency while the interval stays the same. Using our setup the static power signal is usually sufficiently well sampled at moderate sampling frequencies (e.g., 1 MS/s) due to the low bandwidth of the DC amplifier and the low-pass filter. However, some of the noise sources have a very low frequency and therefore require to be sampled over a time period of a certain length in order to be eliminated.

Independent of the time it takes to acquire a set of traces we want to investigate whether we can apply the noise reduction techniques through averaging onto the measurements at the most informative temperature-voltage combination. In particular, we want to take the measurement setting that requires the least amount of traces for a successful key recovery due to an already increased signal (90 °C and 1.98 V) and try to additionally decrease the noise by stretching the interval as much as possible. Therefore we acquired trace sets for the following measurement intervals: 1 ms, 2 ms, 5 ms, 10 ms, 20 ms, 50 ms, 100 ms, 200 ms. The results of the noise estimation can be seen in Figure 17. The noise level




Fig. 18: Number of measurements required to disclose the correct key candidate for measurement intervals between 1 ms and 200 ms at a supply voltage of 1.98 V and a temperature of 90 °C.




Fig. 19: CPA on Hamming weight of 12-bit F-box output for a measurement interval of 200 ms at a supply voltage of 1.98 V and a temperature of 90 °C.

decreases exponentially when linearly increasing the length of the measurement interval until roughly 20 ms. Afterwards the development seems to stagnate. This can be observed when comparing it to the exponentially fitted curve over the data points. However, in Figure 18 it can be seen that the number of required measurements for a key recovery decreases even beyond the 20 ms in an exponential fashion. In any way the decrease of the noise seems to be lower bounded by the amount of algorithmic noise in the measurements, since this part cannot be averaged out by intra-trace averaging (see [27]). For the sake of completeness we should mention that we also performed the experiments with even longer measurement intervals. However, neither the noise nor the MTD could be reduced any further (due to the lower bound). In fact, they even started to slowly increase again. Our assumption is that for very long measurement intervals (>200 ms) the temperature-induced variations coming from the active regulation of the climate chamber start to affect single measurements, instead of being present between traces in a set. Hence, we achieved the overall best results at a temperature of 90 °C, a supply voltage of 1.98 V and a measurement interval of 200 ms. The result of the CPA attack under this setting can be seen in Figure 19. It is shown that 8 000 traces are required to identify the correct key candidate. This in fact corresponds to the number of traces that




Fig. 20: CPA on Hamming weight of 12-bit F-box output for a measurement interval of 1 ms at a supply voltage of 1.98 V and a temperature of 90 °C.

the corresponding CPA on the dynamic power measurements (on the same ASIC chip) required in [27]. For comparison purposes, the result of the same attack for a measurement interval of 1 ms is given in Figure 20 (multiple examples for a measurement interval of 10 ms are given in Figure 10).

VII. CONCLUSIONS

In this work we have presented an extensive case study on the effects of the three measurement factors temperature, supply voltage and measurement interval on the amount of information that can be extracted from static power measurements. We are able to show that by controlling either the temperature or the measurement interval (in case of clock control) the number of traces that are required for a successful key recovery can exponentially be reduced. Additionally we observed that modifying the supply voltage at least marginally eases such attacks as well. In particular by adjusting all three parameters an adversary can theoretically end up with a set of traces that only contains algorithmic noise. We conclude that the existence of the investigated measurement factors and their, in some cases, exponential impact on the success of attacks further strengthen the position of the static power side channel as a realistic target for adversaries against cryptographic hardware. In this regard we would like to encourage research and industry to incorporate static power attacks into their security evaluation and certification processes.

Considering that our target ASIC was manufactured in a rather old 150 nm technology, the results are even more astonishing. The static power dissipation in this technology is still several times smaller than the dynamic power consumption. But still, only by controlling some measurement factors, a successful static power analysis attack on an unprotected implementation could be performed with as many traces as a corresponding dynamic power analysis on the same target. According to [4] the data dependency of the static current in digital standard cells stays the same for smaller technology sizes while especially the subthreshold leakage increases more than linearly. Additionally, the exponential dependencies of some static power sources on the supply voltage become only relevant for more advanced technologies, which again favors the adversaries. Hence, we suspect that the results presented in this work are even more drastic for smaller feature sizes.

As a suggestion for future work on this topic, quite obviously test chips in more advanced technology generations need to be investigated. Additionally, it has to be verified

whether control over these measurement factors enables a similar improvement of higher-order attacks against securely masked implementations. Finally, effective countermeasures need to be constructed to counteract the exploitation of this emerging side channel.

ACKNOWLEDGMENT

The authors would like to acknowledge Axel Poschmann for the hardware designs and Stefan Heyse for his help on taping out the prototype chip. The work described in this paper has been supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972 and through the project 271752544 “NaSCA: Nano-Scale Side-Channel Analysis”.

REFERENCES

- [1] Analog Devices AD8421 Data Sheet Rev. 0. <http://www.analog.com/media/en/technical-documentation/data-sheets/AD8421.pdf>.
- [2] Analog Devices AD8676 Data Sheet Rev. C. <http://www.analog.com/media/en/technical-documentation/data-sheets/AD8676.pdf>.
- [3] Side-channel Attack Standard Evaluation Board SASEBO-R Specification – Version 1.0. http://www.risec.aist.go.jp/project/sasebo/download/SASEBO-R_Spec_Ver1.0_English.pdf. Research Center for Information Security, National Institute of Advanced Industrial Science and Technology, Japan.
- [4] Z. Abbas and M. Olivieri. Impact of technology scaling on leakage power in nano-scale bulk CMOS digital standard cells. *Microelectronics Journal*, 45(2):179–195, February 2014.
- [5] M. Alioto, S. Bongiovanni, M. Djukanovic, G. Scotti, and A. Trifiletti. Effectiveness of Leakage Power Analysis Attacks on DPA-Resistant Logic Styles Under Process Variations. *Transactions on Circuits and Systems I: Regular Papers*, 61(2):429–442, February 2014.
- [6] M. Alioto, S. Bongiovanni, G. Scotti, and A. Trifiletti. Leakage Power Analysis Attacks Against a Bit Slice Implementation of the Serpent Block Cipher. In *MIXDES 2014*, pages 241–246. IEEE, June 2014.
- [7] M. Alioto, L. Giancane, G. Scotti, and A. Trifiletti. Leakage Power Analysis Attacks: Well-Defined Procedure and First Experimental Results. In *ICM 2009*, pages 46–49. IEEE, December 2009.
- [8] M. Alioto, L. Giancane, G. Scotti, and A. Trifiletti. Leakage Power Analysis Attacks: A Novel Class of Attacks to Nanometer Cryptographic Circuits. *Transactions on Circuits and Systems I: Regular Papers*, 57(2):355–367, February 2010.
- [9] D. Bellizia, S. Bongiovanni, P. Monsurrò, G. Scotti, and A. Trifiletti. Univariate Power Analysis Attacks Exploiting Static Dissipation of Nanometer CMOS VLSI Circuits for Cryptographic Applications. *Transactions on Emerging Topics in Computing*, 5(3):329–339, May 2016.
- [10] D. Bellizia, D. Cellucci, V. D. Stefano, G. Scotti, and A. Trifiletti. Novel Measurements Setup for Attacks Exploiting Static Power using DC Picammeter. In *ECCTD 2017*, pages 1–4. IEEE, September 2017.
- [11] D. Bellizia, M. Djukanovic, G. Scotti, and A. Trifiletti. Template attacks exploiting static power and application to CMOS lightweight crypto-hardware. *International Journal of Circuit Theory and Applications*, 45(2):229–241, August 2017.
- [12] D. Bellizia, G. Scotti, and A. Trifiletti. Implementation of the PRESENT-80 Block Cipher and Analysis of its Vulnerability to Side Channel Attacks Exploiting Static Power. In *MIXDES 2016*, pages 211–216. IEEE, June 2016.
- [13] A. Bogdanov, L. Knudsen, G. Leander, C. Paar, A. Poschmann, M. Robshaw, Y. Seurin, and C. Viskelson. PRESENT: An Ultra-Lightweight Block Cipher. In *CHES 2007*, volume 4727 of *LNCS*, pages 450–466. Springer, September 2007.
- [14] E. Brier, C. Clavier, and F. Olivier. Correlation Power Analysis with a Leakage Model. In *CHES 2004*, volume 3156 of *LNCS*, pages 16–29. Springer, August 2004.
- [15] M. Djukanovic, D. Bellizia, G. Scotti, and A. Trifiletti. Multivariate Analysis Exploiting Static Power on Nanoscale CMOS Circuits for Cryptographic Applications. In *AFRICACRYPT 2017*, volume 10239 of *LNCS*, pages 79–94. Springer, May 2017.

- [16] J. Giorgetti, G. Scotti, A. Simonetti, and A. Trifiletti. Analysis of Data Dependence of Leakage Current in CMOS Cryptographic Hardware. In *GLSVLSI 2007*, pages 78–83. ACM, March 2007.
- [17] A. Gornik, A. Moradi, J. Oehm, and C. Paar. A Hardware-Based Countermeasure to Reduce Side-Channel Leakage: Design, Implementation, and Evaluation. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 34(8):1308–1319, 2015.
- [18] S. Guilley, H. Maghrebi, Y. Souissi, L. Sauvage, and J.-L. Danger. Quantifying the Quality of Side-Channel Acquisitions. *COSADE 2011*, pages 16–28, February 2011.
- [19] B. Halak, J. Murphy, and A. Yakovlev. Power Balanced Circuits for Leakage-Power-Attacks Resilient Design. In *SAI 2015*, pages 1178–1183. IEEE, July 2015.
- [20] N. hao Zhu, Y. Zhou, and H. Liu. A Standard Cell-Based Leakage Power Analysis Attack Countermeasure Using Symmetric Dual-Rail Logic. *Journal of Shanghai Jiaotong University*, 19(2):169–172, April 2014.
- [21] D. Helms. *Leakage Models for High Level Power Estimation*. PhD thesis, Carl von Ossietzky Universität Oldenburg, 2009.
- [22] S. S. Immaculate and K. Manoharan. Analysis of Leakage Power Attacks on DPA Resistant Logic Styles: A Survey. *International Journal of Computer Science Trends and Technology*, 2(5):136–141, September 2014.
- [23] P. C. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In *CRYPTO 1999*, LNCS, pages 388–397. Springer, December 1999.
- [24] L. Lin and W. Bursleson. Leakage-Based Differential Power Analysis (LDBA) on Sub-90nm CMOS Cryptosystems. In *ISCAS 2008*, pages 252–255. IEEE, May 2008.
- [25] S. Mangard. Hardware Countermeasures against DPA - A Statistical Analysis of Their Effectiveness. In *CT-RSA 2004*, volume 2964 of LNCS, pages 222–235. Springer, February 2004.
- [26] S. Mangard, E. Oswald, and T. Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer, 2007.
- [27] T. Moos, A. Moradi, and B. Richter. Static Power Side-Channel Analysis of Threshold Implementation Prototype Chip. In *DATE 2017*, pages 1324 – 1329. IEEE, March 2017.
- [28] A. Moradi. Side-Channel Leakage through Static Power — Should We Care about in Practice? In *CHES 2014*, volume 8731 of LNCS, pages 562–579. Springer, September 2014.
- [29] A. Moradi and G. Hinterwalder. Side-Channel Security Analysis of Ultra-Low-Power FRAM-Based MCUs. In *COSADE 2016*, volume 9689 of LNCS, pages 239–254. Springer, 2016.
- [30] A. Moradi and O. Mischke. On the Simplicity of Converting Leakages from Multivariate to Univariate – Case Study of a Glitch-Resistant Masking Scheme. In *CHES 2013*, volume 8086 of LNCS, pages 1–20. Springer, August 2013.
- [31] A. Moradi and T. Schneider. Improved Side-Channel Analysis Attacks on Xilinx Bitstream Encryption of 5, 6, and 7 Series. In *COSADE 2016*, volume 9689 of LNCS, pages 71–87. Springer, 2016.
- [32] S. Nikova, V. Rijmen, and M. Schlaffer. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. *Journal of Cryptology*, 24:292–321, April 2011.
- [33] C. Padmini and J. Ravindra. CALPAN: Countermeasure against Leakage Power Analysis Attack by Normalized DDPL. In *ICCPCT 2016*, pages 1–7. IEEE, March 2016.
- [34] A. Poschmann, A. Moradi, K. Khoo, C.-W. Lim, H. Wang, and S. Ling. Side-Channel Resistant Crypto for less than 2300 GE. *Journal of Cryptology*, 24:322–345, April 2011.
- [35] S. M. D. Pozo, F.-X. Standaert, D. Kamel, and A. Moradi. Side-Channel Attacks from Static Power: When Should we Care? In *DATE 2015*, pages 145–150. IEEE, March 2015.
- [36] M. Rivain. On the Exact Success Rate of Side Channel Analysis in the Gaussian Model. In *SAC 2008*, volume 5381 of LNCS, pages 165–183. Springer, 2008.
- [37] T. Schneider and A. Moradi. Leakage assessment methodology - A clear roadmap for side-channel evaluations. In *CHES 2015*, volume 9293 of LNCS, pages 495–513. Springer, 2015.
- [38] T. Schneider, A. Moradi, and T. Guneyssu. Arithmetic Addition over Boolean Masking - Towards First- and Second-Order Resistance in Hardware. In *ACNS 2015*, volume 9092 of LNCS, pages 559–578. Springer, 2015.
- [39] F.-X. Standaert, T. G. Malkin, and M. Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *EUROCRYPT 2009*, volume 5479 of LNCS, pages 443–461. Springer, 2009.
- [40] F.-X. Standaert, E. Peeters, C. Archambeau, and J.-J. Quisquater. Towards Security Limits in Side-Channel Attacks. In *CHES 2006*, volume 4249 of LNCS, pages 30–45. Springer, October 2006.
- [41] K. Tiri, D. Hwang, A. Hodjat, B.-C. Lai, S. Yang, P. Schaumont, and I. Verbauwhede. Prototype IC with WDDL and Differential Routing – DPA Resistance Assessment. In *CHES 2005*, volume 3659 of LNCS, pages 354–365. Springer, August 2005.
- [42] J. Xu and H. M. Heys. Template Attacks Based on Static Power Analysis of Block Ciphers in 45-nm CMOS Environment. In *MWSCAS 2017*, pages 1256–1259. IEEE, August 2017.
- [43] W. Yu and S. Kose. Security implications of simultaneous dynamic and leakage power analysis attacks on nanoscale cryptographic circuits. *Electronics Letters*, 52(6):466–468, March 2016.
- [44] W. Yu and S. Kose. False Key-Controlled Aggressive Voltage Scaling: A Countermeasure Against LPA Attacks. *Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 36(12):2149–2153, March 2017.
- [45] W. Yu and S. Kose. Security-Adaptive Voltage Conversion as a Lightweight Countermeasure Against LPA Attacks. *Transactions on VLSI Systems*, 25(7):2183–2187, March 2017.
- [46] N. Zhu, Y. Zhou, and H. Liu. Counteracting Leakage Power Analysis Attack Using Random Ring Oscillators. In *SNS & PCS 2013*, pages 74–77. IEEE, May 2013.