Static Power SCA of Sub-100 nm CMOS ASICs

and the Insecurity of Masking Schemes in Low-Noise Environments

Thorben Moos

Ruhr University Bochum, Horst Görtz Institute for IT Security, Germany
thorben.moos@rub.de

Abstract. Semiconductor technology scaling faced tough engineering challenges while moving towards and beyond the deep sub-micron range. One of the most demanding issues, limiting the shrinkage process until the present day, is the difficulty to control the leakage currents in nanometer-scaled field-effect transistors. Previous articles have shown that this source of energy dissipation, at least in case of digital CMOS logic, can successfully be exploited as a side-channel to recover the secrets of cryptographic implementations. In this work, we present the first fair technology comparison with respect to static power side-channel measurements on real silicon and demonstrate that the effect of down-scaling on the potency of this security threat is huge. To this end, we designed two ASICs in sub-100 nm CMOS nodes (90 nm, 65 nm) and got them fabricated by one of the leading foundries. Our experiments, which we performed at different operating conditions, show consistently that the ASIC technology with the smaller minimum feature size (65 nm) indeed exhibits substantially more informative leakages (factor of ~10) than the 90 nm one, even though all targeted instances have been derived from identical RTL code. However, the contribution of this work extends well beyond a mere technology comparison. With respect to the real-world impact of static power attacks, we present the first realistic scenarios that allow to perform a static power side-channel analysis (including noise reduction) without requiring control over the clock signal of the target. Furthermore, as a follow-up to some proof-of-concept work indicating the vulnerability of masking schemes to static power attacks, we perform a detailed study on how the reduction of the noise level in static leakage measurements affects the security provided by masked implementations. As a result of this study, we do not only find out that the threat for masking schemes is indeed real, but also that common leakage assessment techniques, such as the Welch's t-test, together with essentially any moment-based analysis of the leakage traces, is simply not sufficient in low-noise contexts. In fact, we are able to show that either a conversion (resp. compression) of the leakage order or the recently proposed χ^2 test need to be considered in assessment *and* attack to avoid false negatives.

Keywords: Static Power · Leakage Current · Side-Channel Analysis · SPSCA · Masking

1 Introduction

The fundamental physical limits of computation dictate what can and what cannot be achieved by computing machines [BL85]. It has been shown many years ago, for example, that the majority of classical logic gates, being the essential building blocks of computing technology, cannot be evaluated without a certain amount of dissipation [Lan61, BL85]. This statement holds, regardless of the underlying device technology. In particular, state transitions performed by conventional logic operations are often of an irreversible nature, which means that information is discarded since two or more distinct logical states have

IACR Transactions on Cryptographic Hardware and Embedded Systems ISSN 2569-2925, Vol. 2019, No. 3, pp. 202–232
DOI:10.13154/tches.v2019.i3.202-232

a single successor [Ben03][1]. Such transitions *must* be accompanied by a loss of energy to the environment. This has been manifested in Landauer's principle [Lan61] and is a direct implication of the second law of thermodynamics [Llo00]. Whether information is discarded by a logic operation (i.e., an irreversible transition takes place) or not and therefore whether it is dissipated to the environment depends on the processed data [BL85]. Hence, as a matter of fact, computation, as it is currently carried out, does not only imply energy dissipation, but also leakage of information through physical side-channels[2] – entirely independent of any technological details.
This discussion, however, focuses on transitional leakages occurring during an active computation process exclusively. From a thermodynamic standpoint this is sufficient, since there is no necessity for dissipation without a transition of states. In other words, it should be possible to pause a physical computation process and to hold a stable state, keeping sensitive intermediates enclosed in the circuit, without being doomed to an undesired disclosure of information. This is in fact exactly what is described by the famous *only computation leaks* paradigm, introduced in [MR04]. The authors formulate the assumption that *"computation, and only computation, leaks information"*, implying that *"there is no information leakage in the absence of computation"*. Yet, as previous works regarding the information leakage of CMOS devices in stable states have shown, this assumption does no longer approximate the behavior of current semiconductor technologies to a sufficient degree.

Power Dissipation of CMOS Logic

Modern circuit technologies need to achieve many different objectives in parallel, with energy efficiency being only one of them. High performance, reliability, manufacturability and cost effectiveness are fundamental concerns, besides a number of further considerations depending on the desired area of application. Thus, not all effort can be dedicated to the reduction of the energy consumption and it can be observed that technologies suitable for very-large-scale integration (VLSI) in practice usually dissipate significantly more energy than what is demanded by the fundamental physical limits. Complementary metal-oxide-semiconductor (CMOS) logic gates, for example, consume a relatively large data-dependent current during the state transition from one output value to another, due to the associated charging and discharging of output capacitances[3]. Traditionally, this current is assumed to be the predominant cause for both, energy dissipation and information leakage, in this particular technology. However, over the years, physical characteristics and electrical specifications of transistors have changed significantly. To comply with Moores law [Moo65], the dimensions of metal-oxide-semiconductor field-effect transistors (MOSFETs) have faced an aggressive scaling process in order to achieve the desired and predicted exponential increase over time in the number of transistors that can be fabricated on a single integrated circuit (IC) of a given size. In the attempt to uphold this scaling factor, valuable properties of the technology were sacrificed, as for example the negligible current consumption in idle states.
Initially, CMOS logic has been constructed in such a way that, given the idealized model of a transistor holds, no current should be consumed in any stable state. In particular, the individual logic gates are composed of a pull-up network, which establishes a conductive path between the gate output and V_{DD} when activated, and a matching pull-down network,

[1] In a digital two-input `AND` gate, for example, the input combinations (0,0), (0,1), (1,0) are all mapped to output (0) and thus cannot be reversed.

[2] Logical reversibility can indeed be achieved by specialized and more complex logic gates, bearing the potential to eventually evade the lower bound of Landau [BL85, Llo00], however, a suitable device technology for nearly physical reversibility needs yet to be developed. In practice, any computing device will dissipate at least some energy [Llo00].

[3] CMOS gates also consume a (less data-dependent) short-circuit current during any output transition due to the short period of time where both, the pull-up and the pull-down network are conducting.

which is able to create a conductive path between the output and V_{SS} (GND) respectively. For any combination of stable input signals, only one of the two networks is allowed to be active (i.e., switched on), while the other one, and therefore at least one transistor in any path between V_{DD} and V_{SS}, should be switched off. Conceptually, this allows for a negligible power consumption in stable states, as for no static input combination a conductive path is formed across the power supply. Yet, by down-scaling the physical feature size, transistors progressively deviate from the idealized model. To be more precise, a nanoscale MOSFET does not resemble an ideal switch anymore but tolerates a significant off-current to flow between its terminals, even in a supposedly high resistance state. This behavior is a serious concern for hardware designers, as these so-called leakage currents consume a steadily increasing part of the power budget of modern ICs. It also leads to the situation that the global power consumption of circuits cannot be reduced to the amount of active computation anymore, measured by the number of gate toggles for example. Instead, even without any active computation (i.e., in an idle state) a significant amount of energy, proportional to the number of powered logic cells in the circuit, is consumed, independent of whether those cells are actively fed with input data or not. Thus, it is no surprise that leakage current reduction techniques such as power gating (MTCMOS), dual threshold CMOS (DTCMOS) or input vector control (IVC) gained increasing popularity among the VLSI community in the last decades [RMMM03].
Due to the structure of digital CMOS standard cells it can be observed that their individual cumulative off-current is highly determined by the composition and type of active and inactive transistors across the power supply path, which in turn directly depends on the applied input signals to the cell [AO13]. In other words, the static power consumption of CMOS logic is substantially data dependent. One common leakage reduction technique is therefore to assign primarily those input signal combinations to the individual logic cells when the device is in idle which cause the least amount of leakage current. The direct relation between the static power consumption of a cell and its inputs leads to the inconvenient and, from a side-channel perspective, highly alarming situation that on advanced CMOS hardware it is neither possible to actively process data, nor to passively keep (temporary) data in a circuit (e.g., in a flip-flop between consecutive clock cycles), without leaking information about those values via physical side-channels[4]. While the inability to compute without dissipating information-bearing energy amounts is a direct implication of the laws of thermodynamics (at least when considering standard logic gates due to the associated irreversible state transitions), leaking information in stable states (i.e., without any transition) is not necessary from a physical viewpoint and purely caused by technology-specific defaults which are further amplified through scaling effects. Thus, designers of security critical integrated circuits should be aware of the inherent information leakage of CMOS logic in active as well as in inactive states and the potential vulnerability of their devices to side-channel analysis attacks.

Side-Channel Analysis (SCA)

Side-channel analysis attacks exploit the data-dependent dissipation of computing devices in order to extract secret information from circuitry that executes cryptographic primitives. In fact, this threat is not limited to cryptography but applies to any manipulation of sensitive data on physical hardware. The repetitive processing of a fixed symmetric encryption key by a block cipher implementation is just one prime example of a potentially vulnerable target. Obviously, side-channel attacks which rely on measuring the physical emissions of an implementation, in contrast to, for example, its often remotely available execution time, are primarily a concern for devices that an adversary can obtain physical access to. Those devices are typically found in embedded systems. Among the possibilities to measure

[4] We consider only temporary memory elements such as flip-flops and latches here, whose output line, carrying the saved information, is connected to the input of further logic or memory cells.

and quantify the instantaneous data-dependent energy dissipation of an embedded device, power analysis [KJJ99] and electromagnetic emanation (EM) analysis [GMO01] have proven to be the most promising techniques with respect to their efficiency and simplicity, as opposed to, for example, thermal [HS14], acoustic [GST14] or optical [SNK+12] analysis of a target. Accordingly, it is no surprise that the lion's share of attention from academia and industry in the area of physical security of cryptographic hardware is devoted to these two sources of information leakage and their mitigation.

Static Power Side-Channel Analysis (SPSCA)[5]

The main body of research in the field of power analysis attacks focuses on the exploitation of dynamic effects which occur during the computation process, such as the switching of a digital gate output from low to high or vice versa. However, since the dynamic energy consumption (per logic unit) is declining, while the static power dissipation grows significantly in CMOS integrated circuits manufactured in advanced technologies [EB05], researchers have started to investigate the static power consumption as well. It has been shown in previous publications that this source of information leakage can successfully be exploited. [MMR18] provides a thorough description of the history of static power side-channel analysis (SPSCA) throughout the last decade, including a more or less complete list of publications in the area. Following a number of simulation-based investigations, Moradi demonstrated the first practical attempt to quantify the impact of this security threat based on real-world measurements at CHES 2014 for field programmable gate arrays (FPGAs) [Mor14]. Additionally, a first basic technology comparison is presented in [Mor14], as the examined FPGA families were manufactured in three different process technologies. Apart from this work, notable advances in the area include demonstrating that various established countermeasures against dynamic power side-channel analysis are essentially ineffective against the exploitation of the static currents [LB08, ABD+14, ABST14, IM14, Mor14, BST16, BBM+16, MMR17] and providing experimental evidence for the fact that influencing the working conditions of an operating integrated circuit can exponentially ease its exploitation [MMR18].
Of particular interest to the SCA community is certainly the concrete impact of the presence of static power side-channel leakage on the security offered by masking schemes. Masking is undoubtedly the most popular defense mechanism against (dynamic power/EM) side-channel analysis and to the best of our knowledge the only suitable option to achieve provable security claims under reasonable leakage assumptions. The term masking, a.k.a. secret sharing, refers hereby to a class of countermeasures that rely on splitting each sensitive variable of an algorithm into a discrete number of shares in such a way that only the combination of all of the shares contains information about the sensitive values [CJRR99, PR13]. In this way, a security level in terms of required number of leakage traces can be achieved which grows exponentially in the protection order (often closely related to the number of shares) while spending approximately a quadratic amount of resources [JS17, FGP+18]. Yet, such a relation can only be established when the leakage of the individual shares is sufficiently independent and the measurements that an adversary can acquire are sufficiently noisy [SVO+10, PR13, FGP+18]. Without a sufficient amount of noise, masked implementations are not expected to provide a security level that increases significantly in the protection order [CJRR99, SVO+10, PR13, Sta19], making the trade-off

[5] Various different notations have been introduced for *static power side-channel analysis* in the literature, e.g. *static power analysis* [XH17] and *leakage power analysis* [AGST09]. However, since the term *static power analysis* is already an established and unrelated expression in the EDA community and since *leakage* is a frequently used term with a mostly unrelated meaning in the side-channel literature, we stick to the (admittedly quite lengthy) notation of *static power side-channel analysis* in this work and use *static power SCA* and *SPSCA* as its abbreviations.

between spent resources and obtained security guarantee ineffective[6].
The first successful (higher-order) static power side-channel attack on a masked implementation has been performed in [Mor14]. It was also suggested in [Mor14] that masking schemes with a sequential manipulation of the shares (typical in software) might be in danger when an exploitation of the leakage currents is possible, since the shares may be leaked in a univariate fashion through the static power, making multivariate attacks unnecessary and potentially reducing the effective noise level. Further, and even more important to this work, [PSKM15] suggested that in case of an adversary obtaining full control over the clock signal (which was also assumed by [Mor14] and previous works) it is possible to average the static power consumption over an arbitrary time period, which allows to eliminate several sources of noise entirely. It was experimentally verified in [MMR17], and later more empirically in [MMR18], that this averaging technique in static power SCA attacks (with obtained clock control) indeed allows to reduce the noise level significantly. Furthermore, in [MMR17] a successful higher-order static power attack is performed which requires fewer traces to be successful than a corresponding dynamic power analysis attack on the same target. Considering that the static signal on their examined 150 nm chip should be orders of magnitude smaller than the dynamic one, this result clearly indicates that the noise in the static power traces could successfully be eliminated to a large extent. In fact, this result shows that without dedicated countermeasures, it is harder to assure a sufficient noise level against adversaries that measure static currents than against those who rely on measuring the dynamic switching activity of a chip. Such an observation goes hand in hand with the intuition that any *static* physical effect should, by definition, be easier to quantify with a high precision (i.e., low noise) than a corresponding transitional one, simply because static phenomena are persistent and not limited to a finite period of time.
To summarize, only a few practical works can be found in the literature which contribute to the discussion whether this side-channel can actually be of any harm to state-of-the-art cryptographic devices. While these articles deliver very valuable results, they also suffer from a number of shortcomings, making it difficult to fully oversee the concrete potential of this security threat, yet. We give two examples of such shortcomings in the following.
First of all, the technology comparison presented in [Mor14] leaves a lot of room for interpretation. In particular, the author discovers that no clear correlation between the feature size of the underlying CMOS technology of the FPGAs and the magnitude or the exploitability of their leakage currents can be observed. This contradicts not only what is suggested by the theory, but also what can be observed in the leakage characterization sheets of corresponding standard cell libraries [AO13]. In this case it is quite clear (from our point of view) that the inaccuracy of the results comes from the fact that, instead of ASICs, FPGA implementations were targeted. In fact, the three analyzed FPGA families differ in many structural and architectural regards from each other, apart from their underlying CMOS process node. Most of these technological differences and details are kept confidential as intellectual property (IP) by its vendors. Thus, it can never truly be determined which factors contribute to the observation that certain instances have a smaller or larger data-dependent leakage current on one FPGA device than on another. Finally, to the best of our knowledge, the three different FPGA devices were not even manufactured by the same foundry. Thus, a truly fair technology comparison examining the effect of down-scaling on the potency of this side-channel needs yet to be delivered.
The second work which requires a confirmation of its results on a different platform and under different conditions is [MMR17]. This article gives a first indication of the potential inherent susceptibility of masking schemes to static power attacks (which was predicted by [PSKM15]). But, in fact, only a single attack scenario is shown, without

[6] This becomes obvious and when taking a look at information theoretic plots and the lower bounds for the required number of observations to distinguish leakage distributions of boolean masked information [CJRR99, SVO+10, PR13].

any statistical evidence for the reproducibility of the results, and no leakage assessment has been performed on the target[7]. From our point of view, it remains unclear whether the noise reduction through averaging actually led to a signal-to-noise ratio (SNR) where masking is essentially ineffective[8] or whether the noise level was simply reduced to a point where the SNR became greater than that in the compared dynamic power attack. Further, the analysis was performed on a rather outdated technology (150 nm) and without leakage-enhancing operating conditions which have proven to boost the SNR in such experiments [MMR18]. Thus, a more detailed analysis of the topic, preferably on a more advanced device technology and under different operating conditions is required to give a definite answer to the question whether and under which conditions masking and other side-channel countermeasures which require a certain noise level to be effective are inherently susceptible to SPSCA.

Our contribution

The contribution of this work is manifold. To begin with, we have developed two digital ASIC prototypes in sub-100 nm low power CMOS technology, one 90 nm and one 65 nm chip, and got them fabricated by one of the major foundries. All instances relevant to this work have been derived from identical RTL code and were implemented using an identical design procedure. Thus, we are able to provide a fair comparison between both technologies regarding the vulnerability of architectural and cryptographic instances to static power side-channel attacks. As a result of this comparison, we conclude that the data-dependent currents increase drastically when moving towards smaller CMOS technology nodes. In our case, the leakage exhibited by the 65 nm ASIC is roughly 10× as informative as the one on the 90 nm chip. Additionally, for the first time in literature, we perform static power SCA attacks on sub-100 nm CMOS ASICs under leakage-enhancing operating conditions, which allows us to validate the considerable impact that the applied temperature and core voltage can have on the exploitability of the static currents in CMOS devices. Interestingly, we find out that especially the influence of the temperature is much stronger in the more advanced process node. By raising the temperature from 20 °C to 90 °C and the core voltage from 1.2 V to 1.6 V the difference of means between two leakage distributions can be amplified by a factor of approximately 12 on the 65 nm chip.
As a next step, we investigate the susceptibility of masked implementations to SPSCA and conclude that due to noise reduction techniques (i.e., averaging over time) adversaries can obtain measurements with such a low noise influence that masking is essentially ineffective. Furthermore, we argue that state-of-the-art leakage assessment techniques like the Welch's t-test are not suitable when analyzing masked implementations in very low noise environments as they cause false negatives. In fact, we come to the conclusion that moment-based analysis in general is not preferable in low-noise scenarios and that either a conversion, respectively compression, of the leakage order, or the recently presented χ^2 test need to be considered for assessment *and* attack. Finally, we show that for a variety of hardware implementations of cryptographic primitives clock control is no strict requirement to carry out a static power side-channel analysis. In particular, we demonstrate that whenever sensitive information remains in the circuit before or after a cryptographic operation is performed, it can be exploited. In this regard, we perform the first SPSCA attacks that do *not* require a stronger attacker model than conventional dynamic power analysis attacks[9]. Additionally, we show that in some cases it is even beneficial when certain

[7] To the best of our knowledge, none of the previously cited works has conducted a leakage evaluation by means of a statistical test, such as the Welch's t-test, either.

[8] [PSKM15] explains that noise averaging in static power SCA can be used to move from the *effective masking* zone to the *ineffective masking* zone.

[9] Although we perform these experiments at an increased temperature and supply voltage, control over these parameters is not conceptually necessary here and only used to reduce the required amount of traces.

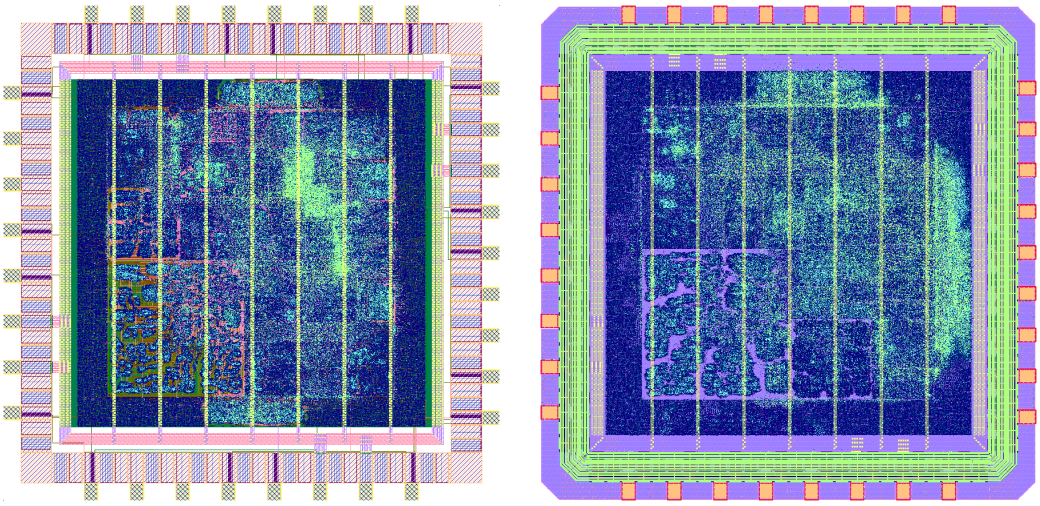

(a) 65nm ASIC layout (b) 90nm ASIC layout

Figure 1: Layout of the ASIC prototypes

parts of the circuit are actively computing during the measurement phase. In the end, we come to the conclusion that dedicated countermeasures against static power side-channel leakage are urgently needed and that masked implementations must be accompanied by a significant amount of algorithmic noise in order to not be susceptible.

2 Experiments

In this section, after shortly introducing the two developed ASIC prototypes and the measurement setup used for the experiments, we present a thorough vulnerability analysis of the devices under test with respect to their susceptibility to static power side-channel attacks. At first, we investigate the effects that manipulations of the operating conditions can have on their exploitability. Then, we analyze architectural and cryptographic instances on both chips to compare the magnitude of the information leakage exhibited by each of the two CMOS technologies. Finally, we use the most successful configuration, in terms of technology node and operating conditions, to carry out more sophisticated attacks.

Target

We have developed two ASIC prototypes in sub-100 nm CMOS technologies, whose layouts can be seen in Figure 1. Both chips are manufactured in low power CMOS technology, using low, high and standard threshold voltage cells. Both require a nominal core voltage of 1.2 V, an IO voltage of 2.5 V and use 9 metal layers for routing. They feature 33 IO pins in total, 17 for logic signals, 16 for power supply. Both chips have been packaged in JLCC-44 package and can be plugged on a custom measurement board which in turn is powered and controlled by a BASYS3 FPGA board. The chips contain a total of 27 different cipher cores, partially equipped with countermeasures against physical attacks, such as masking. All instances have been derived from the same RTL code in both chips and were implemented using the exact same design procedure. However, due to the different technology size some of the cores have a different utilization and a slightly different placement and routing.
Both ASICs contain 8 global 128-bit input registers, which serve the purpose of supplying the cryptographic cores with plaintext and key information, as well as 4 global 128-bit

output registers, which propagate the cores' output to the IO cells. For the cores that are protected by masking countermeasures this information can either be transmitted and received in a pre-shared form through the IO cells, or it is shared internally using fresh masks generated by a randomness source on the chip[10]. The largest block and key size among the cipher cores is 128 bit (AES-128). Accordingly, the size of the shift registers was chosen in order to be able to store a 128-bit key as well as a 128-bit plaintext and ciphertext, each split into 4 shares. In addition to the global input/output (IO) registers each core has its own local IO registers. The global registers are connected to all local registers. All of the cipher cores are clock-gated. Thus, an exemplary input procedure looks as follows. Through a 4-bit data bus a plaintext is given to the global plaintext register, which has been selected by a 4-bit address bus. The same is done for the key. Now, the clock of the targeted crypto core is activated and the plaintext and key are copied into its local registers. The global registers are cleared once the input is copied into the target core. Thus, during measurement the only difference in the state of the device lies in the targeted crypto core.

Setup & Procedure

Our measurement setup and procedure are similar to what has been proposed in [MMR18]. In particular, we use a custom DC amplifier, featuring a $\times 1,000$ amplification and a low-pass filter to get rid of the high-frequency noise in the measurements. Furthermore, we perform all experiments in a climate chamber to guarantee a constant temperature during the acquisition of the traces. The use of such a climate chamber as vital ingredient to any dedicated static power measurement setup was first proposed in [Mor14] and subsequently tested in [MMR17]. Each of our reported static power measurements is obtained by averaging 2 million time samples recorded over a period of 1 s by a LeCroy HRO 66zi sampling oscilloscope (i.e., sampling rate of 2 MS/s, measurement interval of 1 s). The chips were operated at 5 MHz whenever the clock signal was running.

Case Study 1: 1024-bit High-Fanout Register, 65 nm vs. 90 nm

As a first experiment in proof-of-concept manner we target an architectural instance which is expected to exhibit a large data-dependent leakage current, namely a high-fanout state register. In particular we chose the 8 global 128-bit input registers of the ASICs. For this initial experiment it is sufficient to view the 8 registers as one large 1024-bit register. The most important property of this instance for our upcoming analysis is that it is connected to all of the 27 cipher cores that are included in the ASICs. Thus, the output lines of the flip-flops of the 1024-bit register have a comparably large fanout, even though not all register bits are connected to all of the cores. In particular, the average fanout of these flip-flops is 11. Now, as soon as one bit of information is stored in one of them (by applying the value to its input and clocking once) it is directly propagated to the input of 11 further cells on average. An illustration of such a fanout of one single flip-flop to further memory elements can be seen in Figure 2. As detailed in [AO13], both, logic and memory cells leak information about the values that are applied to their input lines via the static power consumption. Thus, the information stored in one flip-flop is not only leaked by the cell itself (which indeed only has a relatively small contribution to the overall leakage), but also by the further 11 cells it is connected to. For this reason, we expect a clearly noticeable difference in the leakage currents when setting the whole 1024-bit state to either all 1s or all 0s.
We first verify this assumption on the 90 nm prototype by means of 5,000 static power measurements that are recorded after filling the registers' content with the randomly

[10] As shown in [SM15], sending only pre-shared input data to the target and receiving the output in shared form can be essential to avoid false positives in side-channel security evaluations.

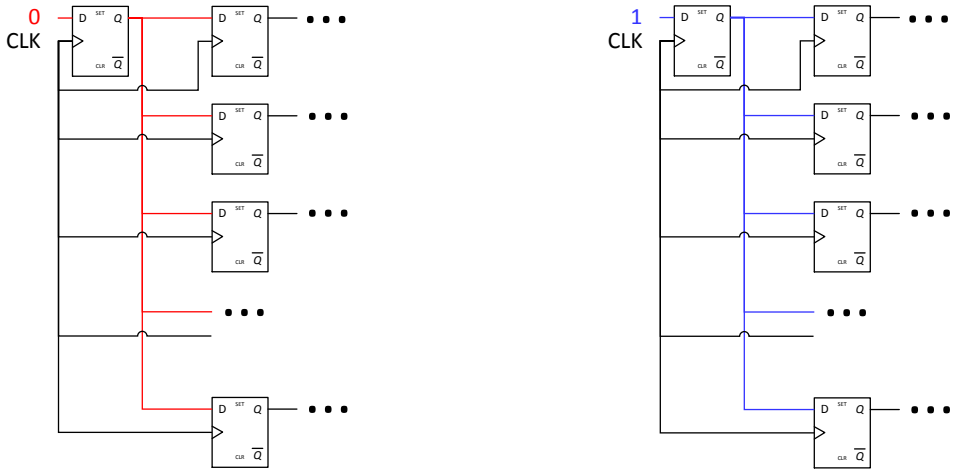

Figure 2: Exemplary depiction of a single bit saved in the high-fanout register, either set to 0 (left) or to 1 (right). The average fanout per first stage flip-flop is 11.

selected input (either all 1s or all 0s) and then stopping the clock signal of the chip and keeping all IO signals constant. The result of those measurements, presented as a histogram, can be seen in Figure 3(a). As expected, we obtain two non-overlapping leakage distributions which can easily be distinguished. The difference between the means of the two distributions is 4.1 µA and the average total current consumed by the ASIC in this idle state is 96.5 µA. All measurements presented in this section were performed using the previously mentioned DC amplifier, which applies a $\times 1,000$ amplification to the static power consumption. However, please note that this amplification is already corrected (i.e., removed) in the reported values. In particular, whenever reporting a leakage current of, for example, 10 µA it means that the amplified static power signal was measured as a voltage drop of 10 mV over a 1 Ω resistor in the V_{DD} path of the ASIC. In order to quantify the distinguishability of the acquired leakage distributions we also performed a Welch's t-test, whose results can be seen in Figure 3(b). Clearly, the t-statistic does not only overcome the threshold of 4.5, which is normally set to decide whether side-channel leakage is detected or not, but it even reaches a value of about 480.

Before comparing these values to the 65 nm chip, we try to amplify the leakage by manipulating the operating conditions. [MMR18] showed that certain measurement factors are capable of boosting the signal-to-noise ratio in static power side-channel measurements significantly. Since we already applied a large measurement interval of 1 s per acquisition, we did not try to increase this parameter even further, but instead concentrated on the operating conditions temperature and supply voltage. In [MMR18] it is demonstrated that the exploitability of a 150 nm CMOS ASIC could exponentially be increased by raising the temperature. Further, increasing the voltage that is applied to the core area of the chip led to a marginal improvement of the attack success as well. The authors mention that the effect of both parameters is expected to increase significantly in more advanced CMOS technologies, such as our two sub-100 nm nodes. Thus, we repeated the initial experiment for another three times. First, we increased the supply voltage, then we raised the temperature, and finally, we manipulated both parameters. The results can be seen in Figure 3. Please note, that the scale of the x-axis in all four histograms is identical, only the range is different. It becomes apparent that both operating conditions have a significant impact on the data-dependent leakage currents. Interestingly, increasing

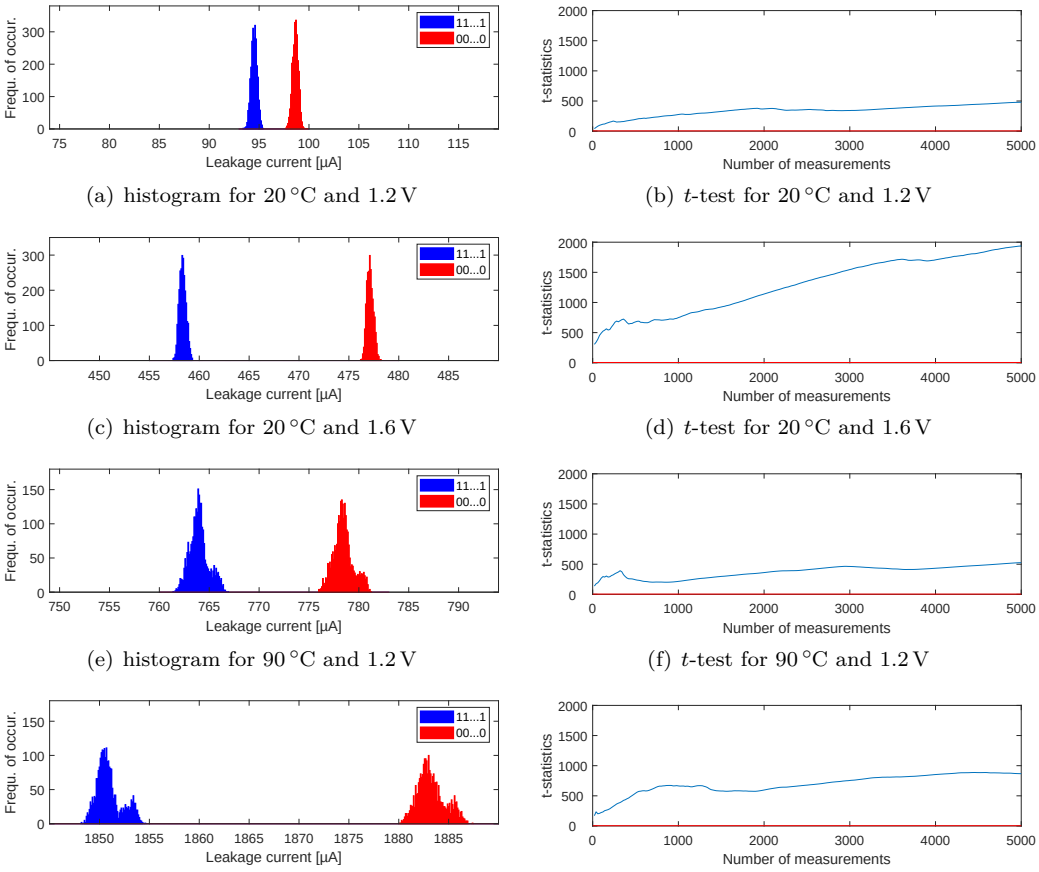

(a) histogram for 20 °C and 1.2 V

(b) t-test for 20 °C and 1.2 V

(c) histogram for 20 °C and 1.6 V

(d) t-test for 20 °C and 1.6 V

(e) histogram for 90 °C and 1.2 V

(f) t-test for 90 °C and 1.2 V

(g) histogram for 90 °C and 1.6 V

(h) t-test for 90 °C and 1.6 V

Figure 3: Histograms and t-test results for 5,000 static power measurements of a 1024-bit high-fanout register in 90 nm CMOS technology, filled either with only 1s or only 0s.

the supply voltage by 33.3% has a larger positive effect on the distinguishability of the distributions than raising the temperature by 70 °C (from 20 °C to 90 °C). This can be observed in both metrics, the difference of means between the distributions and the t-test results. However, the largest difference of means can be achieved by increasing both parameters. Yet, this does not directly lead to an improvement in the t-test results when compared to the scenario where only the voltage is changed. This is due to the additional noise at higher temperatures. As already mentioned in [MMR18], setting the controlled environment in the climate chamber to a temperature far above the room climate, leads to a constant activity of the regulation units, which can be observed as low frequency noise along the recorded set of traces. This type of noise causes the increased variance of the leakage distributions that can be seen in the histograms for the measurement sets that were recorded at 90 °C. However, as also explained in [MMR18], this type of noise can easily be removed by post-processing the traces using a high-pass filter. In this particular experiment we chose to not post-process the traces and rather report the raw, unaltered values as taken from the oscilloscope, in order to not distort the comparison. Yet, in all further case studies following in this section we made use of the moving average filter, as proposed in [MMR18]. Thus, in this experiment the difference of means is indeed the more important metric as it is not significantly influenced by the temperature noise. In total, by raising the temperature to 90 °C and increasing the supplied core voltage to 1.6 V, the difference between the mean values of the two distributions could be amplified by a factor of about 8 to a value of 32.3 µA.

After examining how informative the leakage currents of a 1024-bit high-fanout register in 90 nm technology are, we repeated the exact same kind of experiments on the 65 nm ASIC. The corresponding results are depicted in Figure 4. A couple of interesting differences can be noticed. First of all, while the 90 nm results showed a larger leakage current when the register is filled up with 0s, the opposite can be observed for the 65 nm technology. We refrain from speculating about potential reasons here and stress that this difference is due to internals of the particular standard cells. It is noteworthy that the exact same type of standard cells (i.e., with an identical name) were used for the whole register instance in both technologies, including all cells whose input is connected to the output lines of the register flip-flops. In other words this instance has the exact same netlist on both ASICs. Another difference between Figures 3 and 4 is clearly the magnitude of the currents. Please note that the scale on the x-axis of the histograms in Figure 4 is 10× as large as in Figure 3. This is also the reason why the distributions appear to be narrower, i.e., have a smaller variance, which is indeed not true. It's simply the distance between the distributions which is significantly larger.

One may also notice that, in contrast to the 90 nm results, raising the temperature has a significantly larger impact on the distinguishability of the distributions than increasing the supply voltage in these experiments. This is not only reflected by the difference of means, but also in the t-test results, which is remarkable since the low frequency temperature noise is included in these measurements as well. Table 1 summarizes the data-dependency of both technologies for the different experiments to enable an easy comparison of the vulnerability of the two ASIC prototypes.

Finally, it can be said that in case an adversary is able to manipulate the operating conditions (temperature and supply voltage) of a device under test it is possible to amplify the static power side-channel leakage significantly (in our case by one order of magnitude), given that it is manufactured in an advanced CMOS process. Additionally, we have observed that the 65 nm chip exhibits substantially more informative leakages (also one order of magnitude) than the 90 nm one. Our 65 nm ASIC operating under a supply voltage that has been increased by 33.3% and in a 90 °C environment is more than 100× as susceptible to static power side-channel attacks as our 90 nm chip at nominal supply voltage and room temperature.

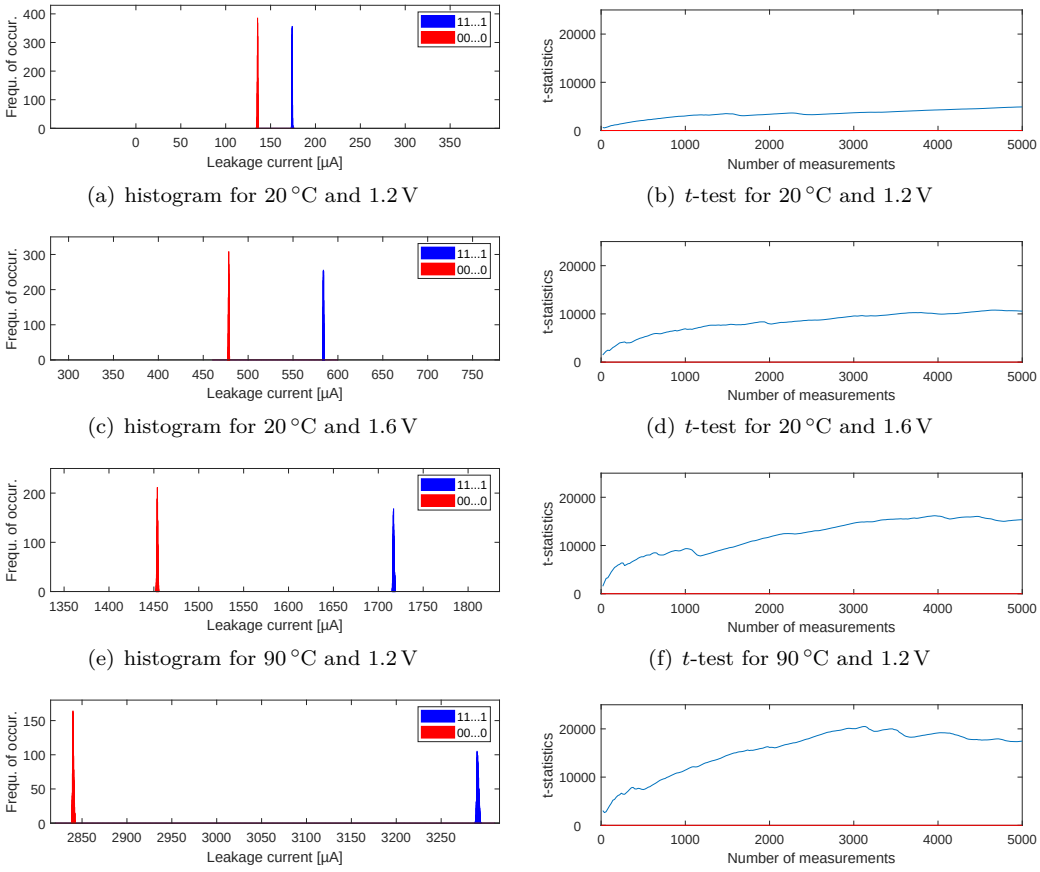

(a) histogram for 20 °C and 1.2 V
(b) t-test for 20 °C and 1.2 V
(c) histogram for 20 °C and 1.6 V
(d) t-test for 20 °C and 1.6 V
(e) histogram for 90 °C and 1.2 V
(f) t-test for 90 °C and 1.2 V
(g) histogram for 90 °C and 1.6 V
(h) t-test for 90 °C and 1.6 V

Figure 4: Histograms and t-test results for 5,000 static power measurements of a 1024-bit high-fanout register in 65 nm CMOS technology, filled either with only 1s or only 0s.

Table 1: Comparison of high-fanout register leakage in 65 nm vs. 90 nm technology for different operating conditions.

Techn.	**Voltage**	**Temp.**	**Diff. of Means**	**Avg. Tot. Curr.**	t**-stat.**
90 nm	1.2 V	20 °C	4.1353 µA	96.5 µA	480
90 nm	1.6 V	20 °C	18.7822 µA	467.3 µA	1,938
90 nm	1.2 V	90 °C	14.4754 µA	771.1 µA	526
90 nm	1.6 V	90 °C	32.3217 µA	1,867.3 µA	867
65 nm	1.2 V	20 °C	38.4927 µA	154.9 µA	4,890
65 nm	1.6 V	20 °C	105.5205 µA	529.9 µA	10,570
65 nm	1.2 V	90 °C	263.1579 µA	1,585.1 µA	15,360
65 nm	1.6 V	90 °C	450.6296 µA	3,067.2 µA	17,460

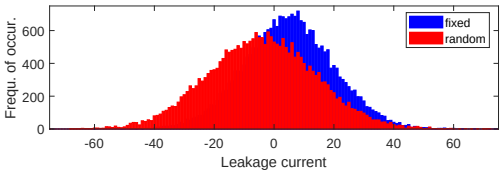

(a) histogram

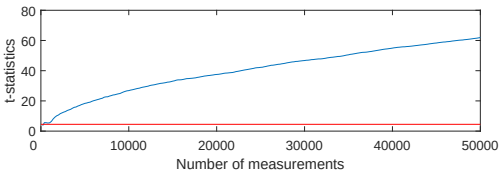

(b) first-order t-test

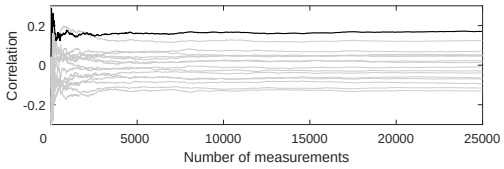

(c) CPA using the HW of one Sbox output after the first round targeting a key nibble

Figure 5: Leakage evaluation and attack using 50,000 fixed vs. random measurements of a nibble-serial implementation of the PRESENT-80 block cipher in 90 nm CMOS technology, recorded at 90 °C and 1.6 V.

Case Study 2: Serial (Unprot.) PRESENT, 65 nm vs. 90 nm, 90 °C, 1.6 V

The second case study of our technology comparison targets an actual cryptographic primitive implemented on both ASICs. The measurements are performed at a temperature of 90 °C and with a supply voltage 1.6 V, since these operating conditions proved to enhance the information leakage through the static power consumption the most. In particular, we analyze the vulnerability of a nibble-serial implementation of the ultra-lightweight block cipher PRESENT-80 [BKL+07], without any side-channel countermeasures applied. The hardware implementation that we used is similar to the profile 1 of [PMK+11]. At first, we performed a leakage evaluation of the hardware primitive implemented in 90 nm technology using a non-specific (fixed vs. random) Welch's t-test, following the guidelines developed in [SM15]. In this regard the PRESENT core is supplied with randomly interleaved sequences of fixed and random plaintexts. Then the computation is executed until the end of the first round, where the clock signal of the ASIC is stopped and the leakage current drawn by the chip is measured. Please note that all global registers, analyzed in the previous case study, are cleared before measuring the static power in order to not obtain any false-positive t-test results, arising from the leakage of the saved plaintext. Thus, only the state which is currently present in the serialized PRESENT circuit differs between multiple measurements. The result of those acquisitions can be seen in Figure 5. As apparent from the histogram, the leakage distributions for the fixed and the random plaintext can easily be distinguished by visual inspection. Furthermore, the t-test overcomes the 4.5 threshold, indicating a detectable leakage, after less than 300 measurements. We also performed a correlation power analysis (CPA) [BCO04] on the traces that were measured for random plaintext inputs and target a key nibble of the first round key by using the Hamming weight (HW) of the Sbox output as a power model. Figure 5(c) shows that the attack succeeds in isolating the correct key candidate from the incorrect key guesses.
Afterwards we performed the same leakage evaluation and key recovery attack on the identical instance in the 65 nm technology. The corresponding results are depicted in Figure 6. Similar to the previous case study the polarity of the distributions is reversed, even though the same fixed plaintext as on the other chip was used. Additionally, it can be observed that the distributions are much easier to distinguish, not only in the difference of their means, but also in their variances. The corresponding t-test leads to a roughly $4\times$ as large t-statistics value and the CPA succeeds with less traces and a larger absolute

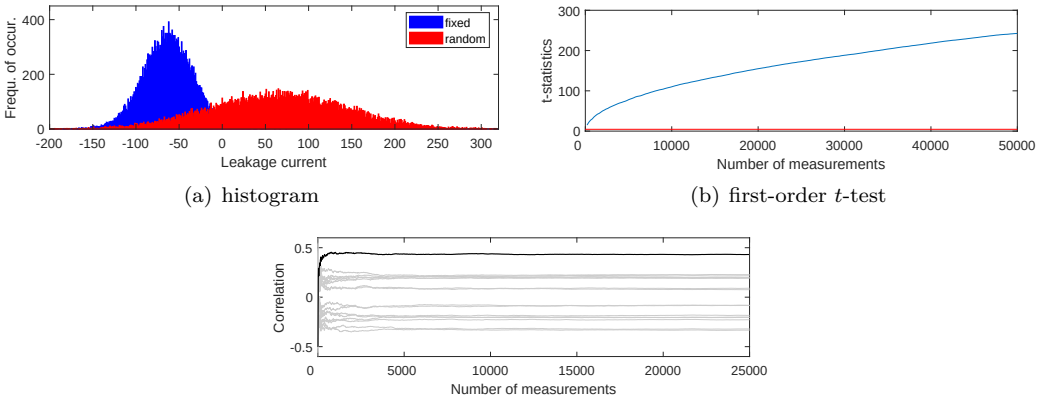

(a) histogram

(b) first-order t-test

(c) CPA using the HW of one Sbox output after the first round targeting a key nibble

Figure 6: Leakage evaluation and attack using 50,000 fixed vs. random measurements of a nibble-serial implementation of the PRESENT-80 block cipher in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

Table 2: Comparison of PRESENT block cipher implementation leakage in 65 nm vs. 90 nm technology for best case operating conditions (adversaries point of view).

Techn.	Voltage	Temp.	Diff. of Means	t-stat.	Correlation	MTD
90 nm	1.6 V	90 °C	9.15	61.96	0.17	2,180
65 nm	1.6 V	90 °C	128.46	242.5	0.43	100

correlation value as before. The concrete values are listed in Table 2 for an easy comparison. As already indicated by the previous case study, the 65 nm ASIC is significantly more vulnerable to static power side-channel attacks. The distance between the means of the fixed and the random distribution is about 14× higher and the attack requires less than $\frac{1}{20}$ of the number of traces, compared to the 90 nm chip.

Case Study 3: Serial (Unprot.) AES, 65 nm, 90 °C, 1.6 V

As a next step we target a byte-serial implementation of AES. The examined circuit is the compact hardware implementation of AES, proposed in [MPL+11]. From this part on we concentrate on exploiting the 65 nm ASIC exclusively, since, based on the previously presented results it can be expected that it leads to more successful results due to a higher SNR. In this regard, we measured the static power consumption of the AES implementation when the encryption is paused after the end of the first round. Again, 50,000 traces for randomly interleaved fixed and random plaintexts are recorded. The corresponding results are presented in Figure 7. It can be seen that the AES hardware implementation is similarly susceptible to static power side-channel attacks as the PRESENT core. In this case we performed two CPA attacks on the traces that were recorded for random inputs. On one hand, we target the HW of the Sbox output which is currently evaluated by the Sbox module to reveal a byte of the first round key. And on the other hand, we correlate the HW of the Sbox output of a different byte, which is already saved in the state register and currently applied to the `MixColumns` operation of AES. Although both attacks do succeed with the available amount of traces, the CPA on the state byte which is currently processed by the Sbox requires much less traces and shows an overall higher correlation for the correct key candidate. This is obviously caused by the fact that this intermediate value is leaked by a larger combinatorial circuit, implementing the non-linear function. However,

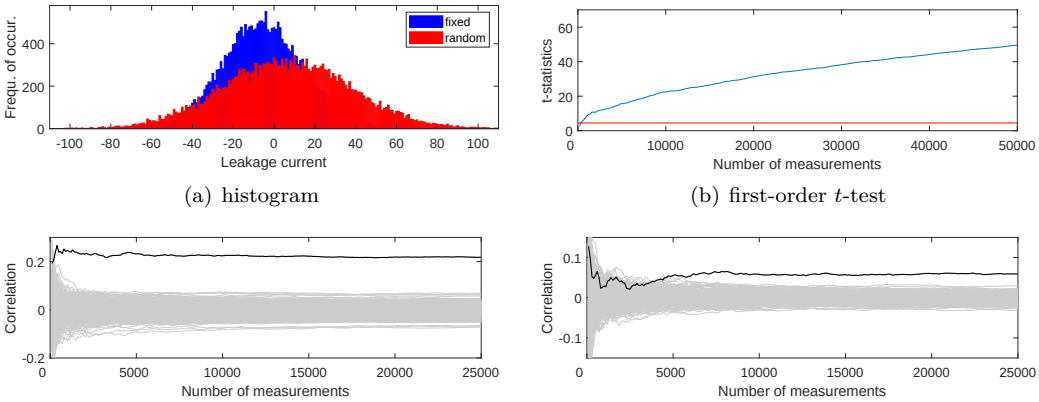

(a) histogram

(b) first-order t-test

(c) CPA targeting an intermediate value that is applied to the combinatorial Sbox circuit

(d) CPA targeting an intermediate value that is applied to the MixColumns operation of AES

Figure 7: Leakage evaluation and attack using 50,000 fixed vs. random measurements of a byte-serial implementation of AES-128 in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

the fact that the leakage of the much smaller and linear `MixColumns` operation is sufficient to exploit it in a key recovery attack, shows that a static power analysis adversary is not forced to measure a new set of traces, stopping the clock in a different cycle, for every key byte when attacking serialized implementations. In theory, when the state register flip-flops are connected to sufficiently leaking memory or logic cells, a single set of traces is sufficient.

Case Study 4: Masked High-Fanout Register Bit, 65 nm, 90 °C, 1.6 V

Masking, a.k.a. secret sharing, is without a doubt one of the most popular and theoretically sound countermeasures against side-channel attacks. In particular, when protecting a cryptographic implementation by means of a masking scheme, it is possible to achieve a security level, in terms of required number of side-channel observations for a successful attack, that grows exponentially in the masking order, while spending approximately a quadratic amount of resources [FGP+18]. However, masking can only deliver such a security guarantee in case the leakage of the individual shares is sufficiently independent and the traces that an adversary can acquire are sufficiently noisy. Due to the fact that temporary physical defaults such as transitions, glitches or couplings are not captured by the way the static power consumption is measured (and therefore cannot influence such measurements) it is comparably easy to achieve independence of the shares with respect to static power side-channel measurements. Yet, it is significantly more challenging to guarantee a sufficient noise level as most of the usual noise sources can be eliminated by averaging over time [MMR17, MMR18]. In this case study we take a look at the leakage of a single bit of information, split into multiple shares which are independently leaked by high-fanout flip-flops. Again, only the 65 nm ASIC is targeted and the operating conditions are set to 90 °C and 1.6 V, in order to obtain the best possible signal-to-noise ratio. We have measured 50,000 traces for randomly interleaved values of the secret bit and for each of the 5 different masking orders. Furthermore we performed (higher-order) t-test evaluations using the formulas introduced in [SM15] on the obtained leakage distributions. The corresponding results are illustrated in Figure 8. It can be seen, that, independent of the masking order, the distributions are clearly distinguishable by visual inspection. In particular, one can easily differentiate the Hamming weight classes of the shared secret. It is obvious that the SNR in these experiments is extremely high. The t-test results

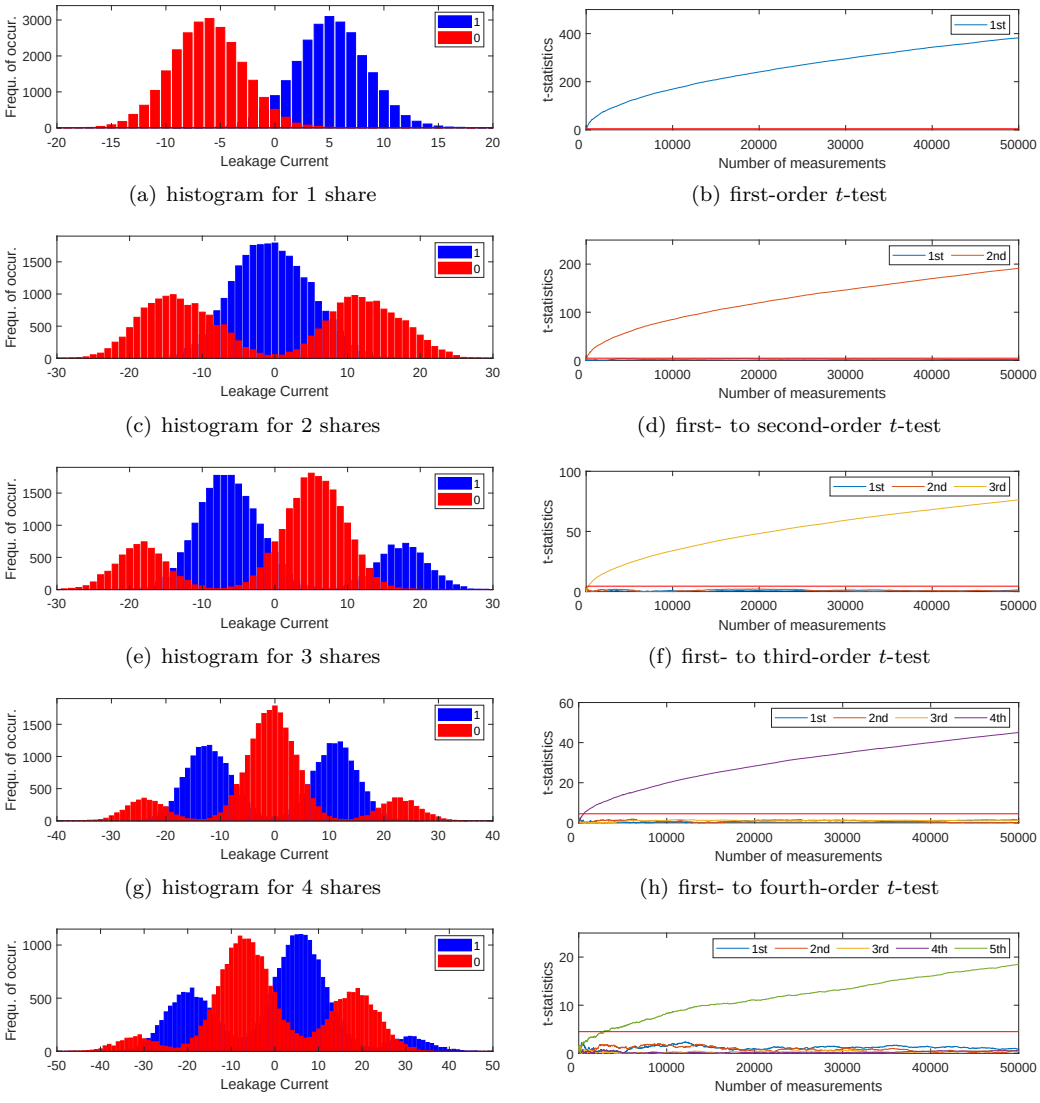

(a) histogram for 1 share (b) first-order t-test

(c) histogram for 2 shares (d) first- to second-order t-test

(e) histogram for 3 shares (f) first- to third-order t-test

(g) histogram for 4 shares (h) first- to fourth-order t-test

(i) histogram for 5 shares (j) first- to fifth-order t-test

Figure 8: Histograms and (higher-order) t-test results for 50,000 static power measurements of 1-bit of information shared among 1, 2, ..., 5 (top to bottom) high-fanout register bits in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

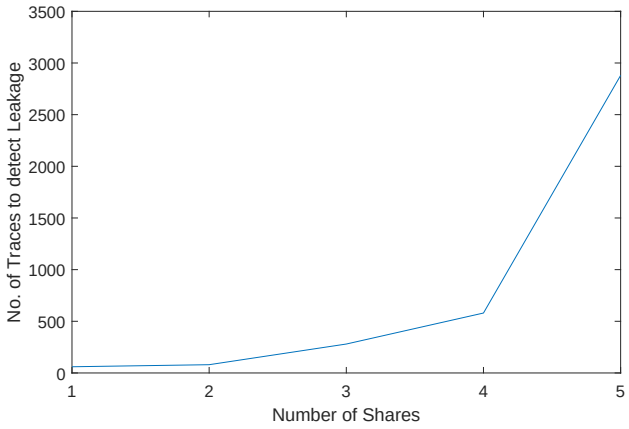

Figure 9: Number of traces to detect leakage for different masking orders.

show that leakage is only present in the expected statistical moments, corresponding to the number of shares. However, even though it does not seem to be significantly more difficult to distinguish the leakage distributions in the higher-order masked cases from their histograms, the t-test performs much worse in terms of the absolute magnitude of the t-statistics and number of required measurements to detect the leakage. This is also depicted in Figure 9. Such a result would suggest that the masking countermeasure is indeed in the *effective masking* zone [PSKM15], since the detection of the leakage becomes a lot more difficult when the masking order is increased. Yet, in the following we will detail that this is in fact a false negative result caused by the moment-based nature of the t-test analysis.

Leakage assessment approaches like the non-specific Welch's t-test have been introduced to simplify side-channel security evaluations of cryptographic implementations. Instead of the concrete exploitation of an implementation these methods are limited to the mere detection of side-channel leakage, independent of the recovery of a secret [SM15]. On the one hand this avoids the necessity to test a multitude of different attack scenarios and intermediate values to target. On the other hand such an approach naturally entails a high risk of false positives. In this context, by false positive we denote the reporting of detectable leakage which is not exploitable in an attack, e.g. leakage of the plaintext or ciphertext or some key-independent intermediates. This is inherent to leakage assessment approaches such as the non-specific Welch's t-test and constitutes the price that needs to be paid for not evaluating a multitude of attacks. However, what should at all cost be avoided are false negatives. By false negative we denote a scenario where a leakage test reports absence of detectable leakage considering a certain amount of traces (even when repeating multiple fix. vs. ran. or fix. vs. fix. tests for different fixed values), while there is indeed leakage present and exploitable with the available amount of traces. Such a scenario is the worst case for an evaluator as it undermines the whole purpose of leakage assessment tools.

In [Sta19] Standaert describes a scenario where such a false negative result can occur in practice, namely when evaluating (higher-order) masked implementations with low noise levels. The fundamental problem of the moment-based test vector leakage assessment (TVLA) methodology in such cases is that an adversarial strategy is assumed. And estimating statistical moments is not the optimal strategy to attack masked implementations with low noise levels [Sta19]. In fact, the number of traces to detect leakage by a moment-based analysis can be much larger than the number of traces to exploit said leakage and recover a secret by a different strategy (than estimating statistical moments) [Sta19]. This is exactly what we observe in Figure 9, as the following comparison shows.

There exist (at least) two alternatives to the estimation of higher-order moments for leakage

evaluation. The first one is the conversion, respectively compression, of the leakage order introduced in [MM17] and the second one is the recently proposed χ^2-test [MRSS18]. The former is based on applying a regular first-order t-test on slices of the leakage distributions and the latter compares full distributions to one another, without being limited to a single moment. We have applied both methods to our masked leakage distributions in Figures 10 and 11. One may notice that the success of both methods is much less affected by the masking order than the higher-order t-test. This becomes apparent in Figure 12 where the number of measurements required to overcome the leakage detection threshold is shown over the masking order. In this regard, we conclude that not only masking is indeed ineffective in very low noise scenarios, which can actually be achieved when performing real-world static power measurements, but also that moment-based leakage assessment techniques such as the Welch's t-test are not suitable in scenarios when the masking order is high and the noise level is low.

On the Need of Clock Control

Traditionally, control over the clock signal of the device under test is an inevitable prerequisite for static power SCA attacks. Thus, performing such an analysis requires a stronger attacker model than classical power analysis adversaries do. [PSKM15] showed that without control over the clock, static power side-channel measurements are less informative than the dynamic power side-channel. Obviously, this is due to the fact that the sensitive intermediate values are present for only one or a few clock cycles in the circuit. Hence, their static power consumption cannot be measured over an extended period of time. However, the longer a certain value is present in the circuit and remains unchanged, the easier it becomes to exploit the leakage current of the respective gates carrying or receiving this information. Thus, theoretically, in case a cryptographic implementation does *not* ensure that *any* sensitive intermediate information is present for at most a few clock cycles in the circuit, this implementation can be susceptible to a static power analysis without the adversary having access to the clock signal. This assumption is explored in the following.
It is usually argued that measuring the static power consumption of, for example, a register content, even if it remains unchanged, cannot be done adequately if the device is actively performing computations somewhere else on the same chip, as the dynamic power consumption of that active computation would dominate the measured voltage drop, induce too much noise and limit the vertical resolution that can be set on the digital sampling oscilloscope. Thus, as a first step we evaluate whether the dynamic power consumption actually has a negative impact on the static power measurements. To this end we have repeated the exact same experiments from the previous case study, but instead of stopping the clock after filling the registers we disabled all the registers after filling them (using their EN pin) and enabled an LFSR-based PRNG on another part of the chip while measuring the total current drawn by the ASIC. The results of such experiments can be found in Appendix A (Figures 19, 20, 21 and 22). It turns out, that the measurements are not more, but in fact less noisy than the previous ones with the stopped clock signal. On the one hand this is due to the fact that the employed DC amplifier and low pass filter (see [MMR18]) have such a low bandwidth and cutoff frequency that no vertical amplitude caused by the dynamic power consumption can be observed. On the other hand this may be caused by the fact that the drop in the power consumption, shown in Figure 3 of [MMR18], is much smaller in this case. Accordingly, it is very well possible to measure the static currents associated with an intermediate value, even when other computations are performed at time of measurement. It is just required that the value remains long enough unchanged in order to measure it precisely. And in fact, many scenarios can be imagined where a sensitive intermediate value remains in a circuit for more than a couple of clock cycles.

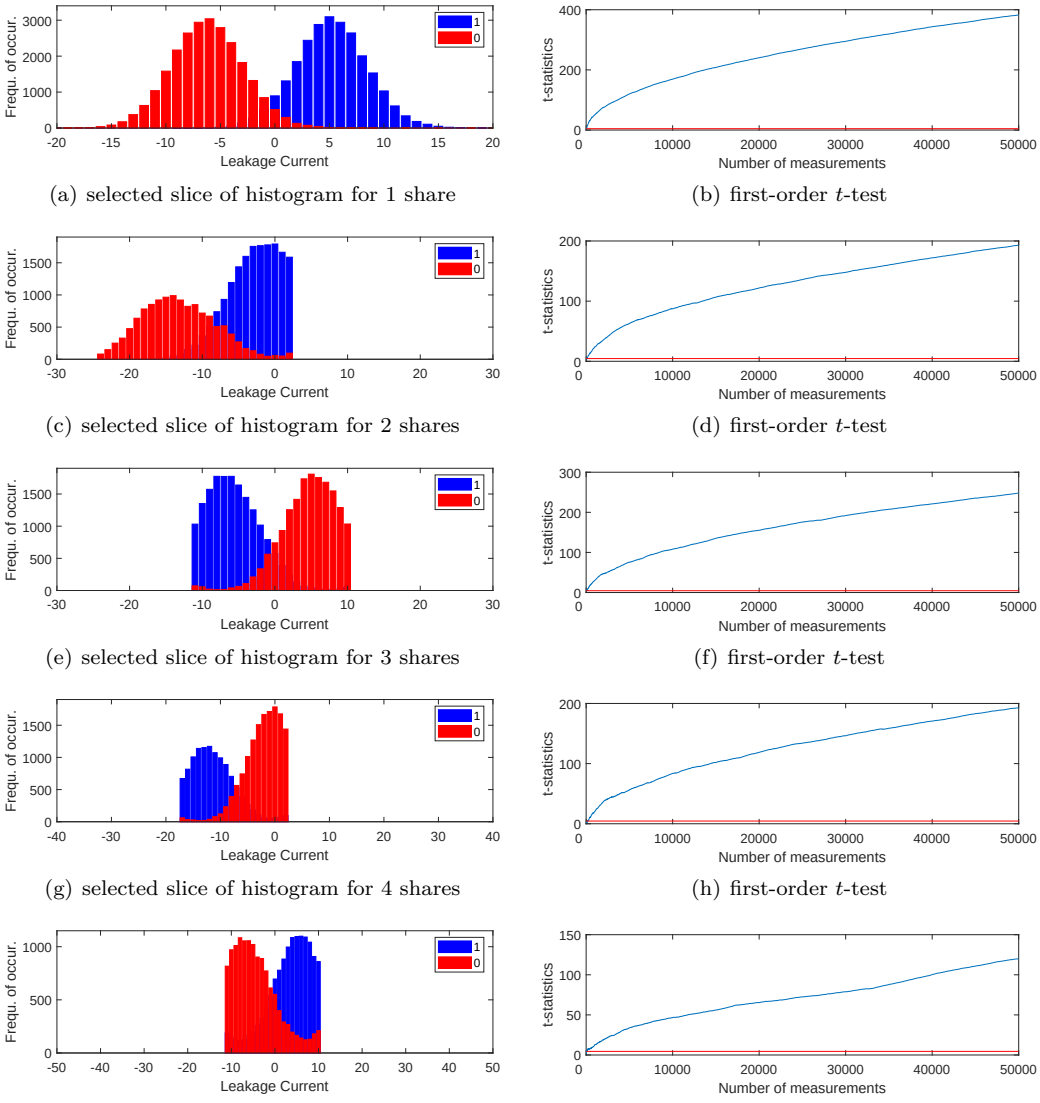

(a) selected slice of histogram for 1 share

(b) first-order t-test

(c) selected slice of histogram for 2 shares

(d) first-order t-test

(e) selected slice of histogram for 3 shares

(f) first-order t-test

(g) selected slice of histogram for 4 shares

(h) first-order t-test

(i) selected slice of histogram for 5 shares

(j) first-order t-test

Figure 10: Histogram slices and t-test results for 50,000 static power measurements of 1-bit of information shared among 1, 2, ..., 5 (top to bottom) high-fanout register bits in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

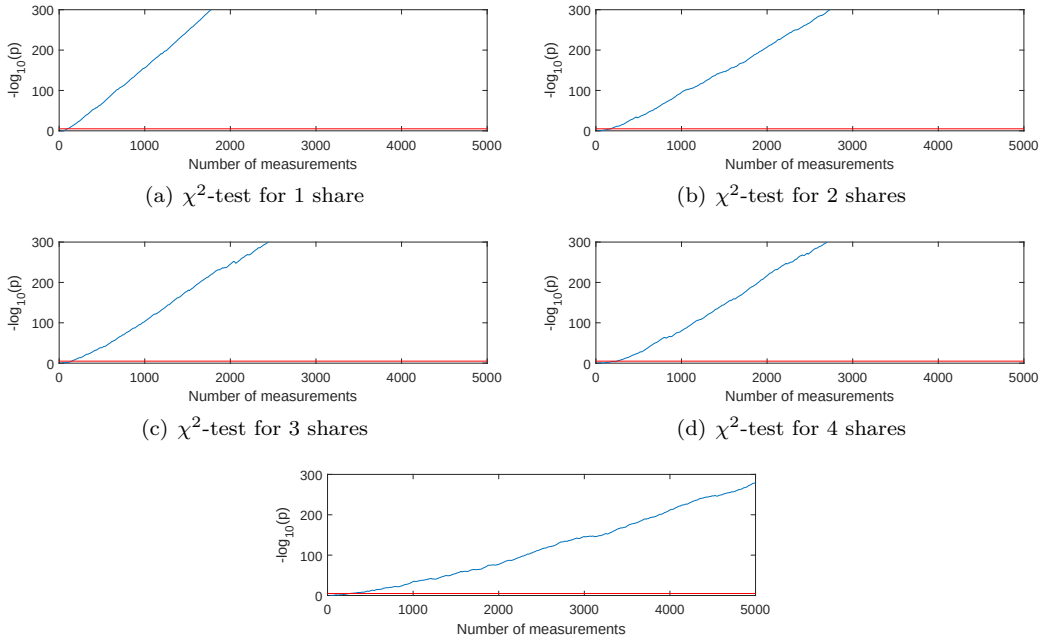

(a) χ^2-test for 1 share

(b) χ^2-test for 2 shares

(c) χ^2-test for 3 shares

(d) χ^2-test for 4 shares

(e) χ^2-test for 5 shares

Figure 11: χ^2-test results for the first 5,000 of the total 50,000 static power measurements of 1-bit of information shared among 1, 2, ..., 5 high-fanout register bits in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

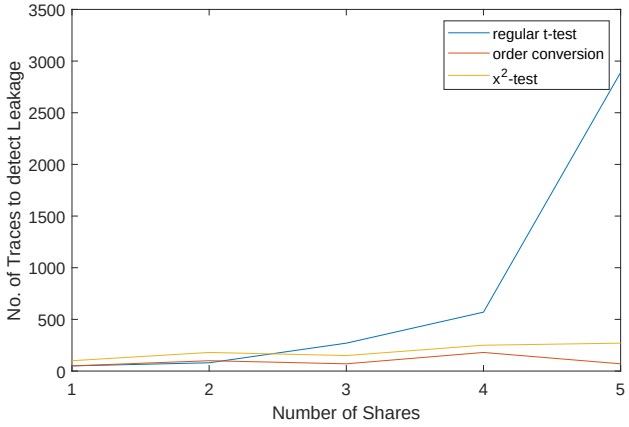

Figure 12: Number of traces to detect leakage for different number of shares using three different methods.

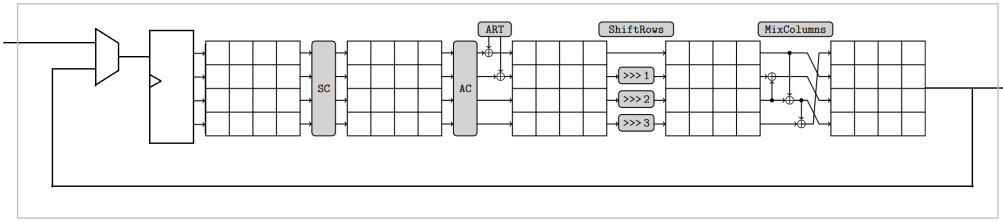

Figure 13: Exemplary depiction of a SKINNY hardware implementation (Partially copied from [BJK+16].)

Case Study 5: Round-Based (Unprot.) SKINNY, 65 nm, 90 °C, 1.6 V, PRNG running

In this case study we present the first realistic scenario, where a static power SCA can be conducted without requiring control over the clock signal of the target. Following the previous discussion this is only realistic when sensitive intermediate values remain unchanged for an extended period of time in an implementation. A regular cryptographic cipher core will only be enabled and clocked when data needs to be encrypted. If not, the core will most likely be in a stable state (i.e., disabled via EN signal or clock-gated). After an encryption has been performed, either all input and intermediate registers are cleared immediately, or the current values remain in the circuit until the next encryption. Often the second option is chosen in order to save delay, power consumption and area (selecting a D-FF without RST signal). Unrolled and pipelined implementations are often not even supposed to be reset between encryptions. In other cases the cipher core is not reset immediately after each encryption, but rather right before the next plaintext needs to be processed, which also allows sensitive intermediates to remain in the circuit for an arbitrarily large period of time. Here we consider a round-based implementation of the SKINNY block cipher [BJK+16], as it is depicted in Figure 13. In particular, a multiplexer decides whether a new plaintext or a previous round output are saved into the state register. The remaining round function is supposed to be purely combinatorial. Typically, such an implementation would be clocked by a state machine until the ciphertext is stable at the output. When this is the case, it means that the second-to-last round output is present in the state register and stays applied to the combinatorial round function. As long as the state register is not immediately cleared, or a new plaintext is encrypted, we can actually exploit the values remaining in the circuit to recover the last round key by calculating back from the ciphertext all the way to the Sbox input of the last round. We have performed a leakage evaluation and the described attack on our SKINNY implementation on the 65 nm ASIC, under a temperature of 90 °C and a supply voltage of 1.6 V. The results are depicted in Figure 14. It can be seen, that the attack succeeds already after a few measurements in isolating the correct key candidate.

Case Study 6: Serial AES TI, 65 nm, 90 °C, 1.6 V

The final case study of this section targets a first-order AES threshold implementation. In this way, we aim to verify whether masked block cipher implementations are actually vulnerable with a comparably small number of measurements to static power side-channel attacks. The targeted circuit is the hardware implementation proposed in [MPL+11]. Figure 15 shows the results of a leakage assessment on this implementation by three different methods. All three techniques, namely higher-order t-test, order conversion and χ^2-test, succeed in detecting the leakage. As a next step we use the three distinguishers to perform a DPA attack on an Sbox output bit, targeting a key byte in the first round.

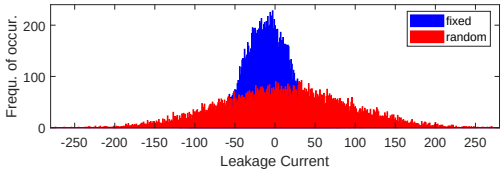

(a) histogram

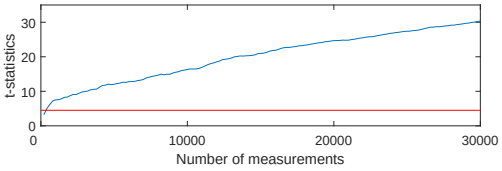

(b) first-order t-test

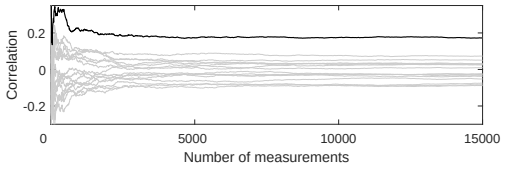

(c) CPA using the HW of one Sbox input in the last round targeting a key nibble

Figure 14: Leakage evaluation and attack using 30,000 fixed vs. random measurements of a nibble-serial implementation of the SKINNY block cipher in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

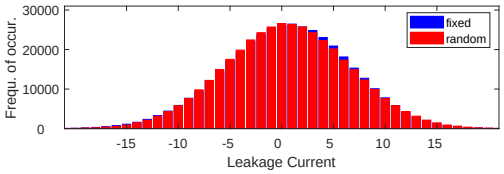

(a) histogram

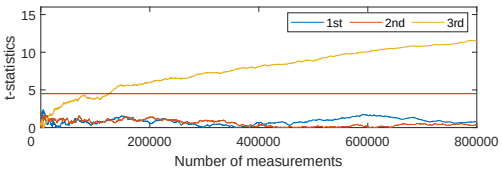

(b) first- to third-order t-test

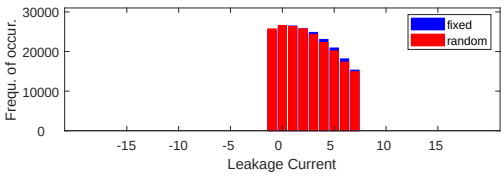

(c) selected slice of histogram

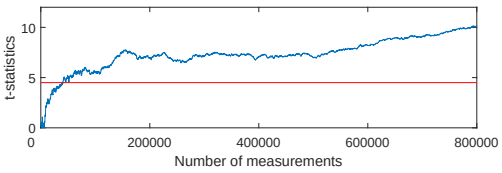

(d) first-order t-test on slice

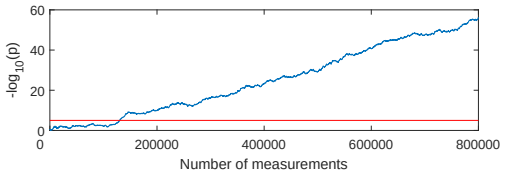

(e) χ^2-test

Figure 15: Leakage evaluation using 800,000 fixed vs. random measurements of a byte-serial AES threshold implementation in 65 nm CMOS technology, recorded at 90 °C and 1.6 V.

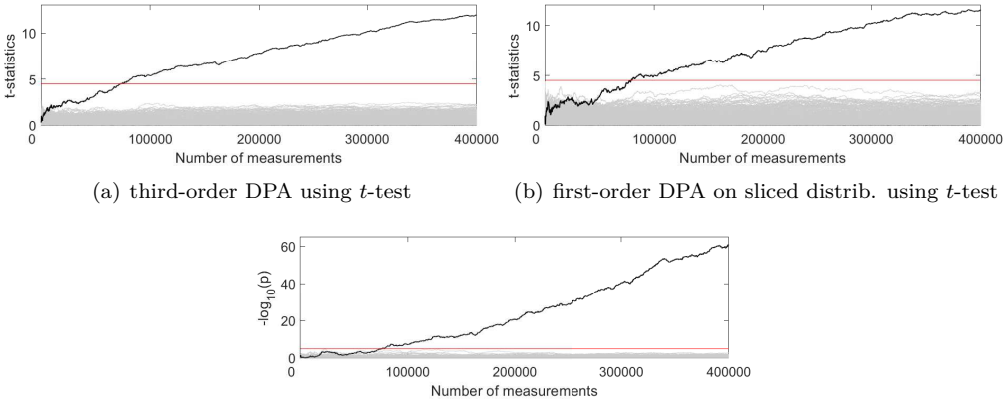

(a) third-order DPA using t-test

(b) first-order DPA on sliced distrib. using t-test

(c) DPA using χ^2-test

Figure 16: DPA attacks using different distinguishers on an Sbox output bit of an AES threshold implementation in 65 nm technology, recorded at 90 °C and 1.6 V.

Again, all three methods succeed, as apparent in Figure 16. The required number of traces to overcome the threshold is similar among the three. However, we noticed that the LFSR-based PRNG, responsible for generating the fresh randomness which is required by the AES threshold implementation contributes significantly to the noise level, due to the fact that it holds a large state of random values during each of our measurements, which are leaked through the static power as well. Thus, in order to avoid this we decided, based on the results we achieved in the previous case studies for when the clock signal is not stopped, to keep the PRNG running during the measurements. Accordingly, its effect can be averaged out in each single measurement and it does not contribute to the algorithmic noise anymore. We repeated the previous evaluation and attacks again using this idea and achieved the results presented in Figures 17 and 18. Please note that this time we measured only 200,000 traces instead of 800,000. While both, the higher-order t-test and the order conversion require roughly 60,000 traces to detect the leakage and 20,000 to recover a key byte, the χ^2-test requires only 30,0000 for the detection and 12,000 for the recovery.

3 Conclusion

In this work, we have shown that the potency of the static power side-channel increases significantly when moving towards smaller feature sizes. Additionally, we could verify that manipulating the operating conditions of integrated circuits in advanced technologies can significantly boost the available information in corresponding static power measurements. This development, together with the possibility to reduce the effective noise level in such attacks poses a serious security risk for cryptographic hardware in advanced CMOS technologies. Countermeasures such as masking, which require a certain noise level to be effective are particularly affected by this development. Furthermore, these countermeasures cannot be properly evaluated by established evaluation methodologies, such as the moment-based TVLA methodology, since those are prone to produce false negatives in low noise environments when the masking order is high. Even devices that do not allow an adversary to obtain control over the clock signal need to pay attention whether sensitive intermediate values remain in the circuit for an extended period of time, e.g. in an idling cipher core. Finally, we conclude that dedicated countermeasures need to be developed to cope with this side-channel. To protect masking schemes from being susceptible, a suitable option is clearly the generation of additional algorithmic noise.

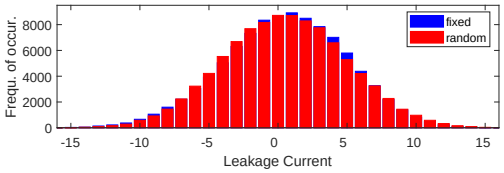

(a) histogram

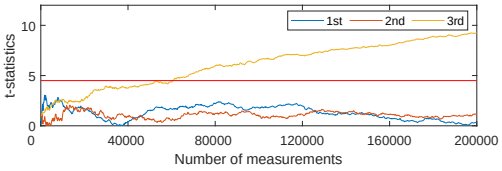

(b) first- to third-order t-test

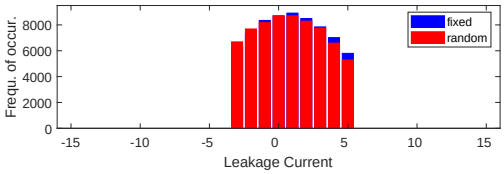

(c) selected slice of histogram

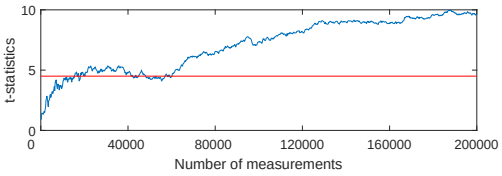

(d) first-order t-test on slice

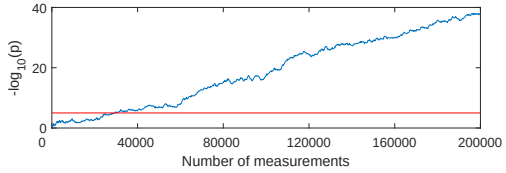

(e) χ^2-test

Figure 17: Same experiments as in Figure 15, but with only 200,000 traces and running the LFSR-based PRNG, responsible for delivering the fresh randomness, during the measurements to minimize algorithmic noise.

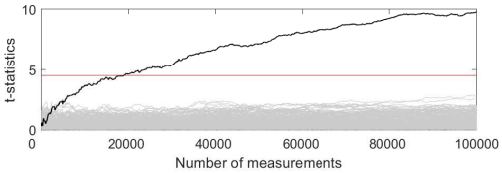

(a) third-order DPA using t-test

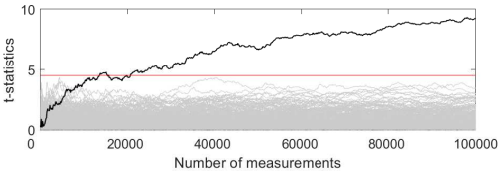

(b) first-order DPA on sliced distrib. using t-test

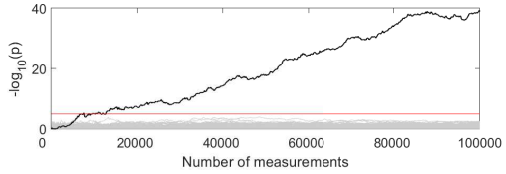

(c) DPA using χ^2-test

Figure 18: Same attacks as in Figure 16, but with only 200,000 traces and running the LFSR-based PRNG, responsible for delivering the fresh randomness, during the measurements to minimize algorithmic noise.

Acknowledgments

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

[ABD+14] Massimo Alioto, Simone Bongiovanni, Milena Djukanovic, Giuseppe Scotti, and Alessandro 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.

[ABST14] Massimo Alioto, Simone Bongiovanni, Giuseppe Scotti, and Alessandro Trifiletti. Leakage Power Analysis Attacks Against a Bit Slice Implementation of the Serpent Block Cipher. In *MIXDES 2014*, pages 241–246. IEEE, June 2014.

[AGST09] M. Alioto, L. Giancane, G. Scotti, and A. Trifiletti. Leakage power analysis attacks: Well-defined procedure and first experimental results. In *2009 International Conference on Microelectronics - ICM*, pages 46–49, Dec 2009.

[AO13] Zia Abbas and Mauro Olivieri. Impact of technology scaling on leakage power in nano-scale bulk cmos digital standard cells. *Microelectronics Journal*, 45, 01 2013.

[BBM+16] Davide Bellizia, Simone Bongiovanni, Pietro Monsurrò, Giuseppe Scotti, and Alessandro 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.

[BCO04] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In Marc Joye and Jean-Jacques Quisquater, editors, *Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2004.

[Ben03] Charles H. Bennett. Notes on landauer's principle, reversible computation, and maxwell's demon. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(3):501–510, 2003.

[BJK+16] Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi, Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The SKINNY family of block ciphers and its low-latency variant MANTIS. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 123–153. Springer, 2016.

[BKL+07] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: an ultra-lightweight block cipher. In Pascal Paillier and Ingrid

Verbauwhede, editors, *Cryptographic Hardware and Embedded Systems - CHES 2007, 9th International Workshop, Vienna, Austria, September 10-13, 2007, Proceedings*, volume 4727 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2007.

[BL85] Charles H. Bennett and Rolf Landauer. The fundamental physical limits of computation. 253(1):48–56, July 1985.

[BST16] Davide Bellizia, Giuseppe Scotti, and Alessandro 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.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Wiener [Wie99], pages 398–412.

[EB05] W. M. Elgharbawy and M. A. Bayoumi. Leakage sources and possible solutions in nanometer cmos technologies. *IEEE Circuits and Systems Magazine*, 5(4):6–17, Fourth 2005.

[FGP+18] Sebastian Faust, Vincent Grosso, Santos Merino Del Pozo, Clara Paglialonga, and François-Xavier Standaert. Composable masking schemes in the presence of physical defaults & the robust probing model. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(3):89–120, 2018.

[GMO01] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic analysis: Concrete results. In Çetin Kaya Koç, David Naccache, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2001, Third International Workshop, Paris, France, May 14-16, 2001, Proceedings*, volume 2162 of *Lecture Notes in Computer Science*, pages 251–261. Springer, 2001.

[GST14] Daniel Genkin, Adi Shamir, and Eran Tromer. RSA key extraction via low-bandwidth acoustic cryptanalysis. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 444–461. Springer, 2014.

[HS14] Michael Hutter and Jörn-Marc Schmidt. The temperature side channel and heating fault attacks. *IACR Cryptology ePrint Archive*, 2014:190, 2014.

[IM14] S. Shiney 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.

[JS17] Anthony Journault and François-Xavier Standaert. Very high order masking: Efficient implementation and security evaluation. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, volume 10529 of *Lecture Notes in Computer Science*, pages 623–643. Springer, 2017.

[KJJ99] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Wiener [Wie99], pages 388–397.

[Lan61] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[LB08] Lang Lin and Wayne Burleson. Leakage-Based Differential Power Analysis (LDPA) on Sub-90nm CMOS Cryptosystems. In *ISCAS 2008*, pages 252–255. IEEE, May 2008.

[Llo00] Seth Lloyd. Ultimate physical limits to computation. *Nature*, 406:1047–1054, 2000.

[MM17] Thorben Moos and Amir Moradi. On the easiness of turning higher-order leakages into first-order. In Sylvain Guilley, editor, *Constructive Side-Channel Analysis and Secure Design - 8th International Workshop, COSADE 2017, Paris, France, April 13-14, 2017, Revised Selected Papers*, volume 10348 of *Lecture Notes in Computer Science*, pages 153–170. Springer, 2017.

[MMR17] Thorben Moos, Amir Moradi, and Bastian Richter. Static power side-channel analysis of a threshold implementation prototype chip. In David Atienza and Giorgio Di Natale, editors, *Design, Automation & Test in Europe Conference & Exhibition, DATE 2017, Lausanne, Switzerland, March 27-31, 2017*, pages 1324–1329. IEEE, 2017.

[MMR18] Thorben Moos, Amir Moradi, and Bastian Richter. Static power side-channel analysis - A survey on measurement factors. *IACR Cryptology ePrint Archive*, 2018:676, 2018.

[Moo65] Gordon E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), April 1965.

[Mor14] Amir Moradi. Side-channel leakage through static power - should we care about in practice? In Lejla Batina and Matthew Robshaw, editors, *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, volume 8731 of *Lecture Notes in Computer Science*, pages 562–579. Springer, 2014.

[MPL+11] Amir Moradi, Axel Poschmann, San Ling, Christof Paar, and Huaxiong Wang. Pushing the limits: A very compact and a threshold implementation of AES. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 69–88. Springer, 2011.

[MR04] Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In Moni Naor, editor, *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 278–296. Springer, 2004.

[MRSS18] Amir Moradi, Bastian Richter, Tobias Schneider, and François-Xavier Standaert. Leakage detection with the x2-test. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):209–237, 2018.

[PMK+11] Axel Poschmann, Amir Moradi, Khoongming Khoo, Chu-Wee Lim, Huaxiong Wang, and San Ling. Side-channel resistant crypto for less than 2, 300 GE. *J. Cryptology*, 24(2):322–345, 2011.

[PR13] Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159. Springer, 2013.

[PSKM15] Santos Merino Del Pozo, François-Xavier Standaert, Dina Kamel, and Amir Moradi. Side-channel attacks from static power: when should we care? In Wolfgang Nebel and David Atienza, editors, *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition, DATE 2015, Grenoble, France, March 9-13, 2015*, pages 145–150. ACM, 2015.

[RMMM03] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand. Leakage current mechanisms and leakage reduction techniques in deep-submicrometer cmos circuits. *Proceedings of the IEEE*, 91(2):305–327, Feb 2003.

[SM15] Tobias Schneider and Amir Moradi. Leakage assessment methodology - A clear roadmap for side-channel evaluations. In Tim Güneysu and Helena Handschuh, editors, *Cryptographic Hardware and Embedded Systems - CHES 2015 - 17th International Workshop, Saint-Malo, France, September 13-16, 2015, Proceedings*, volume 9293 of *Lecture Notes in Computer Science*, pages 495–513. Springer, 2015.

[SNK+12] Alexander Schlösser, Dmitry Nedospasov, Juliane Krämer, Susanna Orlic, and Jean-Pierre Seifert. Simple photonic emission analysis of AES - photonic side channel analysis for the rest of us. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 41–57. Springer, 2012.

[Sta19] François-Xavier Standaert. How (not) to use welch's t-test in side-channel security evaluations. In Begül Bilgin and Jean-Bernard Fischer, editors, *Smart Card Research and Advanced Applications*, pages 65–79. Springer, 2019.

[SVO+10] François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2010.

[Wie99] Michael J. Wiener, editor. *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*. Springer, 1999.

[XH17] J. Xu and H. M. Heys. Template attacks based on static power analysis of block ciphers in 45-nm cmos environment. In *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 1256–1259, Aug 2017.

A Appendix 1

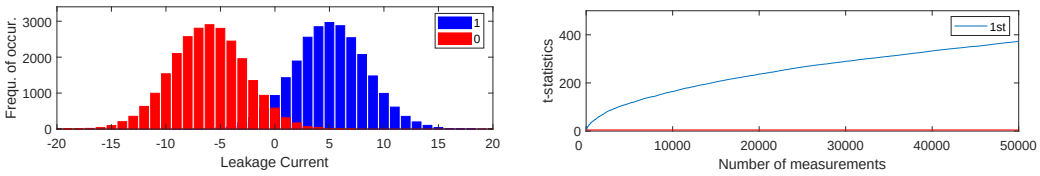

(a) histogram for 1 share

(b) first-order t-test

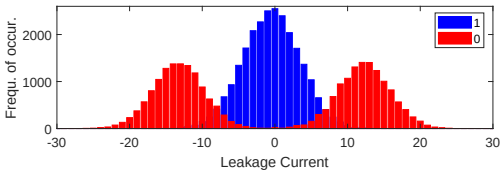

(c) histogram for 2 shares

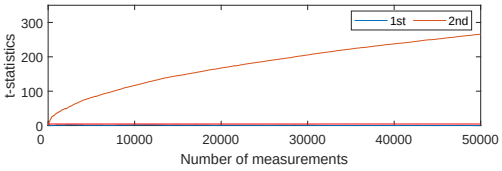

(d) first- to second-order t-test

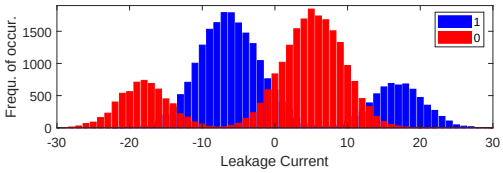

(e) histogram for 3 shares

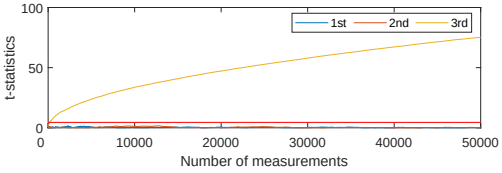

(f) first- to third-order t-test

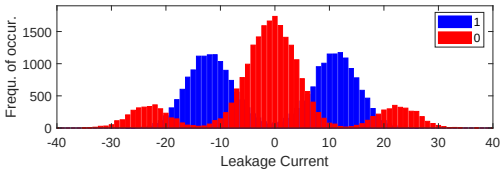

(g) histogram for 4 shares

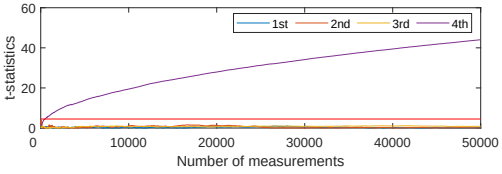

(h) first- to fourth-order t-test

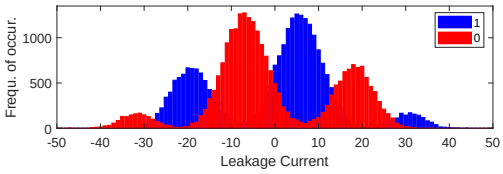

(i) histogram for 5 shares

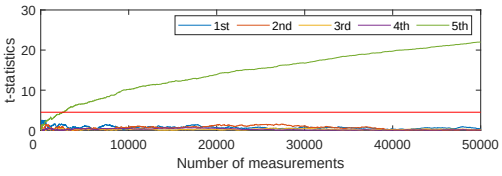

(j) first- to fifth-order t-test

Figure 19: Same experiments as in Figure 8, but without stopping the clock and instead running an LFSR-based PRNG during the measurements.

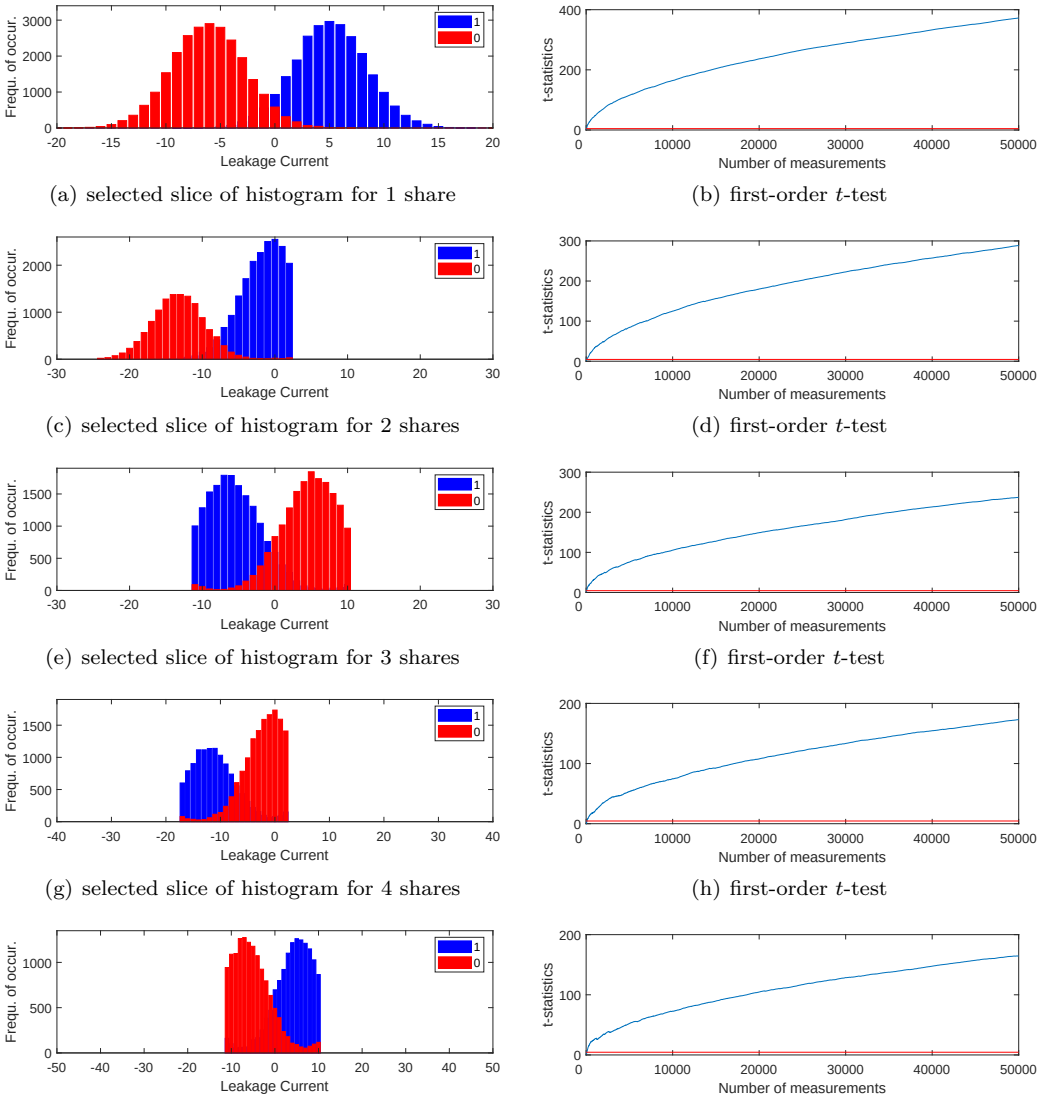

(a) selected slice of histogram for 1 share

(b) first-order t-test

(c) selected slice of histogram for 2 shares

(d) first-order t-test

(e) selected slice of histogram for 3 shares

(f) first-order t-test

(g) selected slice of histogram for 4 shares

(h) first-order t-test

(i) selected slice of histogram for 5 shares

(j) first-order t-test

Figure 20: Same experiments as in Figure 10, but without stopping the clock and instead running an LFSR-based PRNG during the measurements.

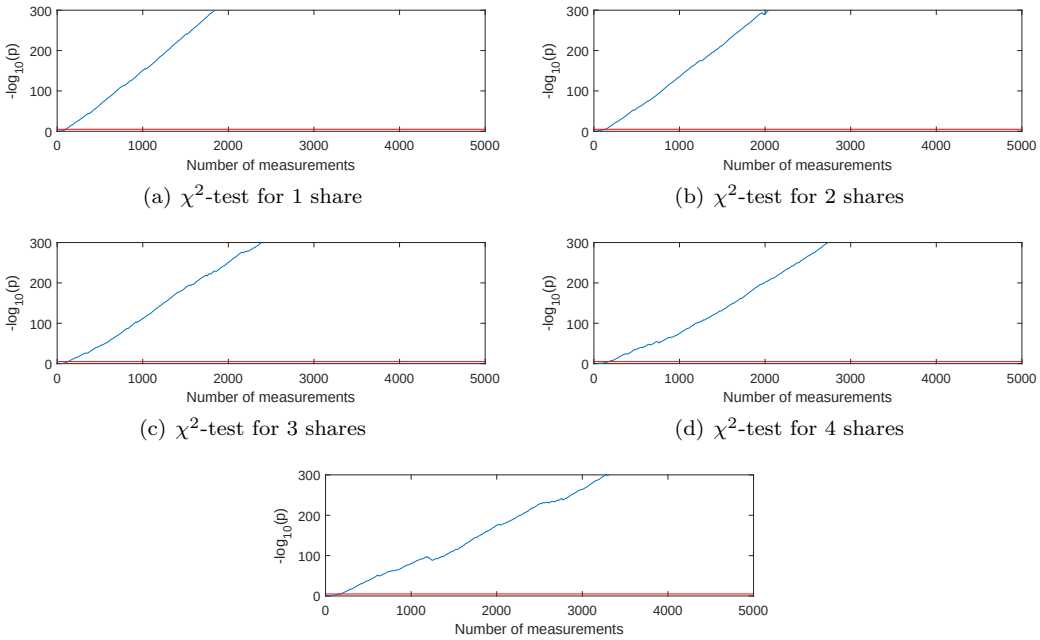

(a) χ^2-test for 1 share

(b) χ^2-test for 2 shares

(c) χ^2-test for 3 shares

(d) χ^2-test for 4 shares

(e) χ^2-test for 5 shares

Figure 21: Same experiments as in Figure 11, but without stopping the clock and instead running an LFSR-based PRNG during the measurements.

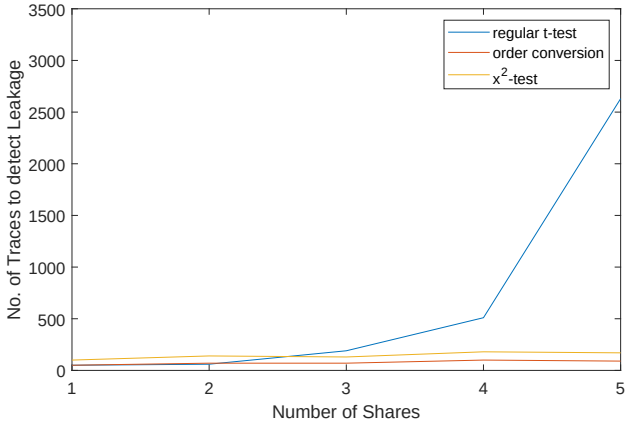

Figure 22: Number of traces to detect leakage for different number of shares using three different methods in Figures 19, 20 and 21.