IACR Transactions on Cryptographic Hardware and Embedded Systems
ISSN XXXX-XXXX, Vol. 0, No. 0, pp. 1–32. DOI:XXXXXXXX

Static Leakage in Dual-Rail Precharge Logics

Bijan Fadaeinia[1], Thorben Moos[2] and Amir Moradi[3]

[1] Ruhr Universität Bochum, Horst Görtz Institute for IT-Security, Bochum Germany,
bijan.fadaeinia@rub.de

[2] Crypto Group, ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium
thorben.moos@uclouvain.be

[3] Technische Universität Darmstadt, Darmstadt, Germany
amir.moradi@tu-darmstadt.de

Abstract. In recent research studies, an observable dependency has been found between the static power consumption of a Complementary Metal-Oxide-Semiconductor (CMOS) chip and its internally stored and processed data. For the most part, these studies have focused on utilizing the leakage currents as a side channel to conduct key-recovery attacks on cryptographic devices. There are two main reasons why information leakage through the static power side channel is considered particularly harmful for the security of implementations, namely 1) the low influence of noise due to averaging over time and 2) the ability to target secrets even outside of the time window that they are actively computed upon (data is leaked for as long as it is saved anywhere in the circuit). Hence, developing effective countermeasures against this threat is of significant importance for the security of cryptographic hardware. Hiding techniques known as Dual-Rail Precharge (DRP) logic have been proposed and studied in literature as an instrument to equalize a circuit's dynamic power consumption irrespective of the processed data. The specific instance called improved Masked Dual-Rail Precharge Logic (iMDPL) is – despite its high overhead – known as one of the most potent and attractive DRP-based Side-Channel Analysis (SCA) countermeasures. While its ability to prevent data extraction through the dynamic power consumption is well studied and documented, we thoroughly analyze its susceptibility to Static Power Side-Channel Analysis (SPSCA) attacks in this work. To conduct our study we have taped-out a custom Application-Specific Integrated Circuit (ASIC) prototype in 65 nm CMOS technology which contains multiple cryptographic co-processors protected by iMDPL, partially combined with other countermeasures. Additionally, it contains circuits protected by a new variant of iMDPL that we specifically hardened against SPSCA, which we call Static Robust iMDPL (SRiMDPL). Our careful experiments performed in a controlled environment under exploitation of voltage and temperature dependencies show that SRiMDPL circuits combined with modern hardware masking offer an extremely high level of security against both dynamic and static power SCA attacks. While the cost of such combinations is admittedly significant (≈ 108 kGE post-layout area for a corresponding PRESENT core), we obtain the strongest combined resistance to both power side channels that has been experimentally demonstrated on real silicon so far. In summary, we believe that our analysis can assist hardware designers in making important decisions on the trade-offs between cost and security that such countermeasures facilitate.

Keywords: Side-Channel Analysis · DRP Logic Style · iMDPL · Static Power · Threshold Implementation · masking

1 Introduction

Beyond the dynamic power consumption of computing hardware, also the static leakage current has to be considered as a potential source of information leakage in security-enabled devices manufactured in nanometer Complementary Metal-Oxide-Semiconductor (CMOS) technologies. The static leakage current is highly dependent on the digital values of input (and output) signals applied to (or computed by) common CMOS logic gates [GSST07]. The leakage of information through this physical quantity is commonly referred to as the static power side channel. Over the past few decades, it has become increasingly clear that Side-Channel Analysis (SCA) attacks pose a significant threat to the security of cryptographic devices. This type of attacks has been widely used to retrieve intermediate values of cryptographic primitives, which in turn allow attackers to reveal the secrets of devices [MOP07]. When SCA attacks have first been introduced [Koc96], the share of the static power in contrast to the dynamic power was negligible. Therefore, other physical side channels such as the dynamic power consumption [KJJ99] and the electromagnetic radiation [GMO01] were in the spotlight for many years. However, static power side channels later attracted more attention in advanced technologies, especially for sub 90 nm [GSST07]. In other words, a simultaneous decrease in dynamic power consumption (per logic unit) and increase in static power can be observed when down-scaling transistor dimensions or reducing the nominal supply voltage. Therefore, it was expected that the static leakage current would dominate the circuit's power consumption sooner or later. While engineering efforts have prevented this from happening for most semiconductor technologies so far (although at a non-negligible price in other metrics), the dependency of the static power on the data processed by a nanometer CMOS device increased enough to enable SCA attacks exploiting this relationship.

The importance of reducing leakage current in advanced semiconductor technologies extends far beyond the threat of SCA attacks. Hence, a large number of technology-level solutions has been developed to alleviate the burden that the static power consumption places on the power budget of modern integrated circuits. Multi-threshold CMOS (MTCMOS), for example, is a variation of standard CMOS chip technology which uses transistors with different threshold voltages (V_{th}) to trade delay for power and vice versa where needed [AMEA02]. Low-V_{th} devices switch faster, advantageous for critical delay paths, but constantly consume a larger amount of static power compared to their standard counterparts. High-V_{th} devices are typically used in non-timing-critical paths to significantly reduce the global static leakage of the device without inducing frequency penalties. As a rule of thumb, high-V_{th} devices can reduce the static leakage by a factor of ≈ 10 compared to low-V_{th} devices. Another technology-level solution to reduce leakage currents in small nanometer CMOS technology generations is called Fin Field-Effect Transistor (FinFET), a 3D transistor structure that offers higher current density and faster switching times compared to traditional planar transistors [GPKKB23]. The gate material in FinFET surrounds the channel region from multiple sides, allowing for better channel control and thus reducing leakage currents, for example subthreshold leakage [MRH+17, HWL20]. Certain variants like Tri-Gate or Multi-Gate transistors provide even increased control and reduced leakage. Fully Depleted Silicon-On-Insulator (FDSOI) is a planar process technology which incorporates insulation layers in the substrate to reduce parasitic capacitance and improve performance compared to bulk CMOS transistors [BLA+19, Har18]. The critical feature of FDSOI is that it provides a very tight electrostatic control of the transistor, which helps reduce the leakage current. Gate-All-Around (GAA) nanowire transistors provide optimal electrostatic control over semiconducting nanowire channels, allowing for downscaling of the gate length while maintaining low off-state leakage. Besides all emerging transistor technologies that directly or indirectly affect the static power of the transistors, some higher level engineering techniques have been developed to reduce the circuits' standby power. For instance, some solutions utilize on-chip sensors to dynamically adjust transis-

tors' operating voltage or threshold voltage based on environmental conditions, workload, or aging effects. This adaptive approach can help reduce static leakage under varying conditions [HSN04, MB21].

The initial studies on Static Power Side-Channel Analysis (SPSCA) attacks have reported successful key recoveries based on simulation results [GSST07, LB08, AGST09, AGST10, ABD+14, ABST14]. However, due to the very low amplitude of the static leakage current, it was initially unclear whether an adversary could successfully perform an SPSCA attack in the real world. Nevertheless, subsequent research through a series of studies based on physical measurements confirmed the feasibility of such attacks [Mor14, PSKM15, MMR17, BCS+17, MMR19, KMM19, Moo20, MM21].

It has been demonstrated that keeping the circuit in an idle mode, collecting a very long power consumption trace, and deriving its DC level can decrease the noise effect on static power measurements, which cannot be achieved for dynamic power measurements. As a result, SPSCA attacks surpassed dynamic power analysis in terms of minimum data complexity to perform successful attacks due to noise reduction. In [KMM19] and [Moo20], the authors have illustrated that the quantity of information leaked through the static power side channel could experience an exponential increase upon manipulation of the device's operating conditions, including raising the supply voltage as well as the temperature. Some studies have also examined the effectiveness of dynamic-power specific SCA countermeasures against SPSCA attacks [BCS+17, PSKM15, MMR17, KMM19, MMR19, Moo19, Moo20, MM21, BBM+17]. As described earlier, an SPSCA attacker has the ability to enhance the Signal-to-Noise Ratio (SNR) during the measurement procedure. This can be accomplished by measuring the static power over a certain period of time and averaging out most electronic and measurement noise influences, ideally utilizing sophisticated high precision equipment to minimize the noise, such as a climate chamber and a precision source meter. Such tools also enable the control of ambient conditions, including the working temperature and supply voltage, to increase information leakage [MM21]. Therefore, static leakage has become one of the most informative side channels, and as a result, SPSCA attacks pose a serious threat to future security-relevant electronic devices.

1.1 Related Works

Information leakage through static power was initially reported in 2007 [GSST07] through a simulation-based study. The first study involving experimental measurements was presented at CHES 2014 [Mor14], demonstrating the practicability of these attacks. Afterwards, multiple studies have investigated SPSCA attacks on both Field Programmable Gate Arrays (FPGAs) and Application-Specific Integrated Circuits (ASICs), often reporting successful key recoveries. Besides studying the feasibility of SPSCA attacks on cryptographic devices, some works tried to develop the corresponding countermeasures [NYH13, HMY15, ZZL14, JIA+15, PR16, YK17, YW18, BFS20]. Most of these works study hiding countermeasures which generally can be classified into two main categories: (1) power-equalization techniques [HMY15, ZZL14, JIA+15, PR16, BFS20] and (2) randomization schemes [NYH13, YK17, YW18]. Both categories share the common goal of reducing the SNR, although their approaches differ significantly. The former aims at attenuating the signal, whereas the latter intends to increase the noise. Some schemes of the first category strive to equalize the static power by utilizing standard CMOS library cells [ZZL14], while others opt for customized cell designs [PR16, BFS20, FMM21]. For example, the underlying concept of Symmetric Dual-Rail Logic (SDRL) [ZZL14, ZZL13] is to minimize the correlation between the input vector and the leakage current of every SDRL cell, achieved through the use of identical standard CMOS cell pairs. In order to operate properly, SDRL cells require complementary input signals. An SDRL inverter, for instance, consists of two identical standard CMOS inverters that receive complementary inputs. In an ideal scenario, i.e., without any process variations, this complementary inverter gate should exhibit a constant leakage

current regardless of its input. However, process variations strongly affect the leakage current of sub 90 nm CMOS technologies [MH19,CRA+06,DGS+11,ABD+14,RSV+11]. As shown in [MM21], such countermeasures can reduce but cannot fully avoid the dependency of static leakage on the circuit's internal state. For example, the authors have presented successful key-recovery attacks on a circuit realized from SDRL using less than 10,000 static power measurements. An algorithmically masked hardware implementation, i.e., Threshold Implementation (TI) [NRR06], realized in combination with SDRL could also still be attacked using around 320,000 static power measurements, employing higher-order SCA.

On the other hand, some Dual-Rail Precharge (DRP) logic styles – designed to mitigate dynamic power SCA attacks – have been evaluated regarding their static power vulnerability [BBM+17] using transistor-level (SPICE) simulations. The authors have studied Sense Amplifier Based Logic (SABL) [TAV02], Dynamic Differential Logic (WDDL) [TV04], and Masked Dual-Rail Precharge Logic (MDPL) [PM05] by examining the mutual information [SMY09a] between the static power of small circuits (e.g., 4-bit Sbox) realized using such logic styles and the given 4-bit input. The corresponding results predicate that – in contrast to SABL – the circuits made by WDDL and MDPL leak even more information through static power side channel than their standard CMOS counterpart [BBM+17]. Thus, it remained an open question whether DRP logic styles in general are suited to thwart SPSCA attacks.

One of the most advanced countermeasures against SPSCA is presented in [MM21]. The authors introduced Exhaustive Logic Balancing (ELB) as the most effective hiding countermeasure, where every gate is instantiated as many times as the number of its possible input vectors while each of those input vectors is applied to one of the gates. The authors conducted experimental analysis on several different implementations equipped with hiding and masking, and demonstrated that neither hiding nor masking countermeasures alone suffices to withstand SPSCA attacks. Instead, the best result is obtained using a combination of hiding and masking, more precisely a masked circuit based on TI realized through ELB. Yet, very importantly, the authors found that even this costly combination of countermeasures (increasing the circuit area about 23-fold) was still susceptible to higher-order static power analysis with less than 3 million traces, concluding that the search for better solutions to prevent exploitation of this side channel has to continue.

1.2 Our Contributions

In this work, we present the results of an experimental evaluation of the security of DRP-protected circuits with respect to SPSCA attacks, and introduce a new (slightly adapted) instance of such a scheme specifically optimized to prevent them. Our experimental analyses are based on an ASIC prototype chip designed and fabricated in 65 nm CMOS technology. As a particular case study, we focus on improved Masked Dual-Rail Precharge Logic (iMDPL) [PKZM07] which has shown a higher level of security against dynamic power SCA attacks than most other DRP-based SCA logic styles [KP09, MKEP12]. Due to the essential alternation between precharge and evaluation phases of DRP circuits, both dynamic and static power SCA attacks can be conducted on each phase individually. In short, our experimental evaluations imply a significantly lower amount of exploitable data dependency in the static power when the circuit is in the precharge phase. Therefore, we introduce a mechanism to force the circuit into the precharge phase as soon as it detects a condition in which static power can be measured effectively. Although this scheme is valid for every DRP logic style, in our so-called Static Robust iMDPL (SRiMDPL) approach we additionally propose other modifications to be done on iMDPL cells to improve their resistance against SPSCA attacks.

To assess the effectiveness of SRiMDPL in mitigating both dynamic and static power information leakage, we performed experimental evaluations on the aforementioned fabri-

cated ASIC chip. We further implemented a provably secure TI-based hardware masked cryptographic core by means of SRiMDPL cells to significantly reduce the statistical dependency between the secret data and SCA leakages. Using such a hybrid countermeasure, first-order attacks are avoided due to the provable security of the underlying first-order TI. Higher-order attacks are also highly mitigated via a small SNR, which is the result of the employed SRiMDPL cells. Our practical results demonstrate that the resulting circuit offers substantial resistance against dynamic as well as static power SCA attacks. Expectedly, such combinations lead to a considerable resource overhead, similar to what the authors of [MM21] reported. However, these studies are seen as a guide for hardware designers, particularly those interested in leveraging DRP logic advantages, to trade security for cost while preventing even sophisticated physical attacks.

2 Background

Although CMOS technology dominates the integrated circuit industry due to its many advantages including low static power dissipation, the down-scaling of transistor physical dimensions over the past few decades has resulted in an increase in undesired leakage currents, which can be exploited as a side channel. In the following, we explain why technology scaling increases static leakage and discuss its dependency on the internal state of the circuit. Since this study focuses primarily on the static power of DRP logic styles, it continues with a brief review of the underlying concept. In addition, this section provides an overview of the PRESENT cipher and the Threshold Implementation concept which our case studies are based on.

2.1 Effect of Scaling on CMOS Static Leakage

The static leakage refers to the current which passes through a transistor even when it is in `off` state. In practice, the leakage of any CMOS transistor consists of three primary sources: gate tunneling leakage, sub-threshold leakage, and junction tunneling leakage.

Moving towards nano-scale CMOS technology in order to fit more computation capability on progressively smaller chips is a continuous goal of the semiconductor industry. On the other hand, the aspect ratio of each transistor's dimension should be kept within a certain range to build circuits with close to ideal behavior. In summary, a transistor's aspect ratio A is determined by its horizontal to vertical dimensions as given in Equation (1), where L is the channel length, T_{ox} the gate oxide thickness, D the depletion depth, X_j the junction depth, ϵ_{si} the silicon permittivity, and ϵ_{ox} the oxide permittivity.

$$A \approx \frac{L}{\sqrt[3]{T_{ox}\frac{\epsilon_{si}}{\epsilon_{ox}}X_j D}} \tag{1}$$

Therefore, in modern technologies besides channel length scaling, a couple of other parameters play important roles in maintaining the aspect ratio at a desirable level [MOP07]. The difficulty is that reducing the vertical dimensions (gate oxide's thickness) is more complicated than shrinking the horizontal dimension (channel length and width). This is due to the fact that by approaching the gate oxide thickness limits, the gate tunneling leakage current increases rapidly [MSM+99, Sch99]. At the same time, the supply voltage should be scaled (reduced) to limit power consumption, which requires scaling the threshold voltage, resulting in an increase in the sub-threshold leakage current.

As stated earlier, to maintain a proper aspect ratio, the junction depth should also be scaled, but this simultaneously increases the transistor's resistance. Therefore, this limits the amount of reduction in the junction depth. On the other hand, by channel length reduction, the channel doping near the source-to-body and drain-to-body junctions

should be increased to minimize barrier lowering. As the doping near these junctions increases with scaling, higher junction tunneling leakage is observed in the channel edge. Fortunately, some engineering efforts such as [RMMM03, ZLGL22] are able to limit the increase in leakage currents.

2.2 Information Leakage Through the Static Power

Since different circuits may use distinct varieties of transistors, the optimization parameters for each transistor may be different. For instance, relatively long channel length transistors are used in SRAMs to minimize the influence of random dopant variations. Hence, leakage current in SRAMs is dominated by gate leakage, while that of transistors involved in logic gates is usually dominated by sub-threshold leakage [NC10]. According to Equation (2), the drain current depends exponentially on the gate-source voltage V_{gs} as well as the drain-source voltage V_{ds} in the sub-threshold domain, where V_{th} stands for threshold voltage. $n = 1 + C_d/C_{ox}$ is a parameter related to the parasitic capacitors of a transistor, I_0 the sub-threshold leakage current, and $V_T = KT/q$ ($\approx 26\,mV$ @ $25°C$) the thermal voltage.

$$I_{d_{sub}} = I_0 \exp\Big(\frac{V_{gs} - V_{th}}{nV_T}\Big)\Big(1 - \exp(-\frac{V_{ds}}{V_T})\Big) \tag{2}$$

In a digital circuit, logical input vectors are applied to the gates of the transistors and determine V_{gs}. Hence, V_{gs} and consequently $I_{d_{sub}}$ depend on the value of the given input. The distinguishability of the leakage currents for different inputs is then even amplified by the concrete arrangement of transistors in parallel or serial connection inside of multi-input logic gates [KMM19]. Therefore, by measuring the static power (leakage current) of a circuit, one may exploit its dependency on the circuit's internal state and recover the secrets.

It is worth mentioning that sub-threshold leakage current in down-scaled technologies is exponentially temperature dependent while the tunneling currents have a slight temperature dependency [NC10]. Accordingly, in order to use static power as a side channel, running the device at a high and constant temperature is suggested [MMR19]. In addition, elevating the supply voltage results in an escalation of V_{ds}, thus amplifying the sub-threshold current. This, in turn, enlarges the data-dependency of static power compared to the noise and consequently facilitates SPSCA attacks [MMR19].

2.3 Dual-Rail Precharge Logic Styles

Cell-level hiding countermeasures are known as one of the industry's first reactions to the exploitability of confidential data through dynamic power SCA attacks. In the realm of power equalization schemes, which belong to these cell-level hiding techniques, every logic cell should be re-designed to consume a constant amount of power for every given input transition in order to make the circuit's power consumption independent of the processed data. To achieve this, many constructions leverage the principles of DRP logic [MOP07]. In such schemes, every single-bit signal is accompanied by its complement and the goal is to maintain a constant number of transitions for each input change and prevent any glitches. In order to achieve this, all signals switch between the evaluation and precharge phases for every given input. During the precharge phase, both rails of each signal are precharged to a predetermined value. Therefore, a dual-rail signal in the evaluation phase is considered valid only if the corresponding dual-rail has complementary values. Several schemes have been designed based on this concept, e.g., [TV04, PM05, PKZM07, TAV02], while the early propagation effect [SS06] is known as the most challenging issue which degrades their security. It refers to conditions where a DRP gate already evaluates its output when not all of its inputs are in the evaluation phase, e.g., an AND gate generating its complementary output once observing that one of its inputs carries '0'. This behavior

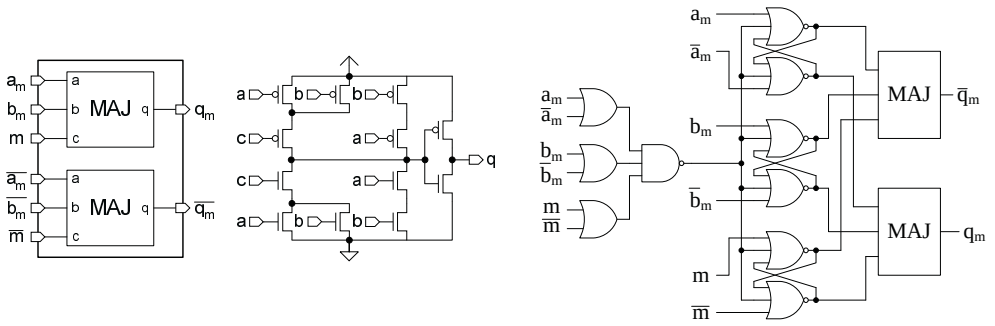

(a) MDPL AND gate (b) CMOS majority gate (c) iMDPL AND gate

Figure 1: MDPL and iMDPL building blocks, taken from [PM05] and [PKZM07].

is also known as data-dependent time of evaluation [MOP07] and leads to different power consumption patterns associated with different given inputs.

Further, the signals of each dual-rail should be routed as similar as possible. Otherwise, the imbalance between their capacitance would lead to a clear distinguishability through dynamic power consumption when one of the rails encounters a transition [TV03, FML+03]. We should refer to [TV04, GHMP05, TV06] as potential solutions to mitigate such imbalances. Alternatively, it has been suggested in [PM05] to swap the signals of each dual-rail randomly at every evaluation/precharge phase, hence avoiding the need for balanced dual-rail routing. More precisely, a single random bit m (refreshed at every cycle) is used for the entire circuit, and the rails swap their values when $m = 1$. If a dual-rail tuple is denoted by $(a, \overline{a})$, it carries $(\overline{a}, a)$ for $m = 1$. Since we can write this as $(a_m, \overline{a_m})$ with $a_m = a \oplus m$, i.e., similar to a Boolean masking scheme, the underlying construction is called Masked Dual-Rail Precharge Logic (MDPL) [PM05]. As an advantage, MDPL cells can be solely built by standard CMOS cells, while the main combinational gate, i.e., MDPL AND gate, is constructed by two majority gates. Figure 1(a) and Figure 1(b) present the corresponding building blocks, where the mask and its complement are shown as $(m, \overline{m})$.

Apart from this advantage, further analyses have shown that MDPL cells still suffer from the early-propagation effect [SS06]. Consequently, the improved version, i.e., iMDPL, has been proposed as a potential successor defeating such a shortcoming [PKZM07]. This is done by inserting an additional block at every input of MDPL cells, which examines if all given inputs are in the evaluation phase (to pass the inputs to the cell) or all of them are in the precharge phase (to let the cell also move to the precharge phase). The block diagram of such a circuit is also given in Figure 1(c). With respect to resistance against classical dynamic power attacks, the result of experimental investigations using Differential Power Analysis (DPA) [KJJ99] and Correlation Power Analysis (CPA) [BCO04] reported in [KP09] have shown the advantages of iMDPL circuits compared to MDPL. Further analyses have however demonstrated that such classical SCA attacks can still recover the secrets when a sufficient number of measurements is collected and properly pre-processed by some signal processing tools [MKEP12]. It has also been shown that since the dual-rail mask signal $(m, \overline{m})$ is connected to all iMDPL cells, it has a relatively high capacitance. Therefore, if m and $\overline{m}$ are routed independently of each other and hence have different capacitances, their transition at the start of the evaluation phase (resp. the precharge phase) becomes distinguishable through observing the dynamic power consumption. In other words, it might be possible to classify the traces (for each clock cycle) to two groups based on the revealed mask value, and conducting SCA attacks on each group individually, removing the impact of the mask entirely.

A modern successor to iMDPL is LUT-based Masked Dual-Rail with Pre-charge Logic (LMDPL), a gate-level masking scheme for first-order SCA security [LMW14, SBHM20]. Its resulting circuits have two working cycles, one for precharge and one for evaluation, which can even be compacted into a single clock cycle to achieve masked circuits with low latency [SBHM20]. However, its underlying strategy and working principle is fundamentally different to iMDPL. While both are based on DRP logic and make use of random mask(s), iMDPL is a classical DPA resistant logic style, while LMDPL is closer to a gadget-based masking scheme. In iMDPL the mask is never used to conceal the data directly, but only to randomly swap the complementary rails in order to counter the effect of routing imbalances. Furthermore, only a single global mask bit is used for all iMDPL gates. LMDPL on the other hand requires an individual mask bit for each non-linear two-input gate. Classical DRP logic styles are typically classified as cell-level hiding countermeasures and attempt to conceal leakages through equalization while operating on unmasked data. LMDPL circuits are based on Boolean masking and leverage the dual-rail principle to achieve single-cycle provable first-order robust-probing security while operating on masked data. Hence, we see LMDPL more as a representative of glitch-resistant hardware masking schemes than as a logic style or a hiding countermeasure and expect that its resistance to static power attacks is comparable (or slightly improved due to the complementary nature) to other representatives of the former class whose susceptibility has been analyzed in prior publications [MMR17, Moo19, MM21]. Analyzing whether this assumption is accurate or not would require an extensive study of the resistance of LMDPL to static power attacks, which is among our plans for future work, but out of scope for this paper. In this work we aim to study the resistance of strong hiding and masking countermeasures both in isolation and in combination, which would not be feasible using LMDPL as a study object as it is not expected to provide security when removing either the dual-rail or the masking principle. Instead, we rely on threshold implementations (introduced in Section 2.5) as our representative of the family of hardware masking schemes in order to provide provable glitch-robust probing security in our most secure implementations. Since this masking scheme works on an algorithmic level, it can trivially be combined with iMDPL as a logic style and, as we will see, the resulting combination leads to extremely strong resistance to static power SCA.

2.4 PRESENT Block Cipher

PRESENT is a symmetric key block cipher designed for resource-constrained devices [BKL+07]. It has a 64-bit block size and supports 80-bit and 128-bit key lengths. Its design follows a Substitution-Permutation Network (SPN) structure and consists of 31 rounds. The cipher makes use of a lightweight 4-bit S-box for security and implementation efficiency. It has a minimal hardware footprint, making it ideal for ultra-constrained environments such as RFID tags, Internet of Things (IoT), and wireless sensor networks. A sketch of the area-optimized serialized implementation of PRESENT-80 proposed in [PMK+11] is shown in Figure 2. Such an implementation, which processes one nibble per clock cycle, has a very small area footprint, making it one of the most resource-efficient block ciphers available.

2.5 Threshold Implementation

Introduced in 2006, TI proved to be a cutting-edge defense against SCA attacks [NRR06]. Its primary approach involves splitting sensitive information into multiple Boolean shares (i.e., Boolean masking), which guarantees that the original data cannot be revealed by classical (i.e., first-order) DPA attacks. The underlying procedure makes sure that the combination of multiple shares through implementation defaults (e.g., glitches) reveals no information about the secrets (original data), offering provable first-order security

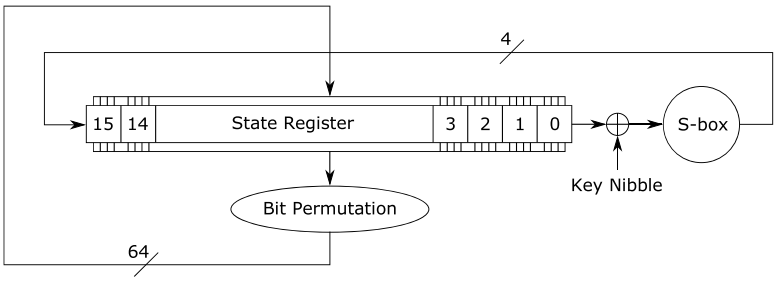

Figure 2: Architecture of the nibble-serial PRESENT-80 hardware implementation. The key schedule is not shown [MM21].

even in hardware. To obtain this level of protection for a function with an algebraic degree of t, TI mandates the use of at least $td + 1$ input shares. Prior to TI, numerous masking schemes claimed first-order security while they were still susceptible to attacks relying on a comparison of the means of distributions in practice mainly due to glitches. Acknowledging this, TI's introduction prompted extensive research focused on glitch-resistant masking [RBN+15, CRB+16, GMK16, GMK17, GM17, BDF+17, GM18, GIB18, FGP+18, MMSS19, CGLS20, CS20, SM20, CS21].

Earlier research has suggested that adversaries employing static power side channel attacks might be able to exploit higher-order vulnerabilities in masked hardware implementations with lower data complexity compared to those observing dynamic power [MMR17, Moo19]. In [MM21], the authors analyzed the effectiveness of TI – also in a combination with hiding countermeasures – in preventing information leakage through static power side channel. They demonstrated that such combinations may conceal higher-order leakages within reasonable limits. In this work, we adopt a similar approach and evaluate the effectiveness of TI circuits realized by DRP logic.

2.6 Mutual Information

In a simulation environment, measurements are typically noise-free. Therefore, minimal differences in energy consumed by different processed data lead to successful key recovery. Therefore, conducting a thorough SCA evaluation in the simulation domain involves assessing the circuit's vulnerability to attacks in light of the noise level. This is because SCA measurements are always subject to noise from the measurement setup and environmental factors. To this end, [SMY09b] has developed an analysis scheme (so-called IT analysis) based on Information Theory, which has been used for example in [BBM+17, MKSS09, KME+08] to evaluate DPA-resistant logic styles. In this analysis, a certain level of noise with Gaussian distribution is added to the noise-free leakage values and Mutual Information (MI) is estimated using conditional entropy as

$$\mathrm{I}(S; L) = \mathrm{H}[S] - \mathrm{H}[S|L], \tag{3}$$

where L denotes the (noisy) leakage values collected from (e.g. SPICE) simulations and S is the associated selected intermediate value (e.g., the circuit's input). The conditional entropy can be estimated utilizing the integral over l as

$$\mathrm{H}[S|L] = -\sum_{s} \Pr[s] \int \Pr[l|s] \cdot \log_2 \Pr[s|l]\, dl. \tag{4}$$

For the given noise level, it examines the amount of available information that the worst-case adversary can exploit since it considers neither any hypothetical model nor any assumption (linear/non-linear) on the SCA leakages' dependency on the processed

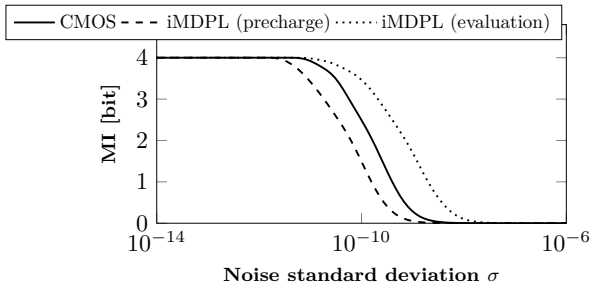

Figure 3: MI analysis of a PRESENT Sbox circuit realized from regular CMOS logic cells and iMDPL logic cells in both precharge and evaluation phase.

data. For a set of noise-free simulated SCA leakages, MI is estimated for several different noise standard deviations. For deficient noise levels, the conditional entropy (Equation (4)) becomes zero, hence $\mathrm{I}(S;L) = \mathrm{H}[S]$. For a very high noise level, trivially, the conditional entropy $\mathrm{H}[S|L]$ becomes very close to the full entropy $\mathrm{H}[S]$, and mutual information tends to zero. Therefore, IT analysis aims to create a curve of mutual information based on the noise standard deviation. This allows for the identification of the noise level required to fully conceal information leakage. In short, a lower noise level required to mitigate the leakage identifies a higher robustness as the leakage can be covered with less noise.

3 Analyses

Our study consists of various stages, starting with a SPICE simulation of a simple PRESENT Sbox realized in 1) regular CMOS and 2) the iMPDL logic style by means of IT analysis. Afterwards, we describe the target device and the measurement setup used for the experimental analyses before comprehensively assessing the static power SCA leakage of iMDPL circuits on an ASIC as a case study of DRP logic styles.

3.1 Simulation

For the simulations, we used Synopsys HSPICE using a 65 nm commercial library. We have simulated a PRESENT Sbox circuit based on both classical CMOS and iMDPL at 90 °C. The simulation of the iMDPL protected Sbox is done in both the evaluation and the precharge phase. Figure 3 depicts the result of the corresponding IT analysis (as explained in Section 2.6), which is compatible with the simulation results reported in [BBM+17] for MDPL. In particular, we find that iMDPL logic cells also tend to leak more information than classical CMOS cells when considering the evaluation phase. However, in the precharge phase, the opposite appears to be true. These results are shown in Figure 3. It can be observed that more noise is required to conceal the evaluation phase leakage of the iMDPL-protected Sbox, while less noise is required for the precharge phase leakage. We will make similar observations in the experimental case study in the following.

3.2 Experimental Setup

3.2.1 Physical Target Device

Our experimental analysis is based on the 65 nm CMOS ASIC prototype shown in Figure 4(a) and Figure 4(b). The chip's physical dimensions are 3432.0 µm × 3431.4 µm. Its I/O and core nominal supply voltages are 2.5 V and 1.2 V, respectively. Although the chip contains several cryptographic cores, in this work we focus only on four co-processors

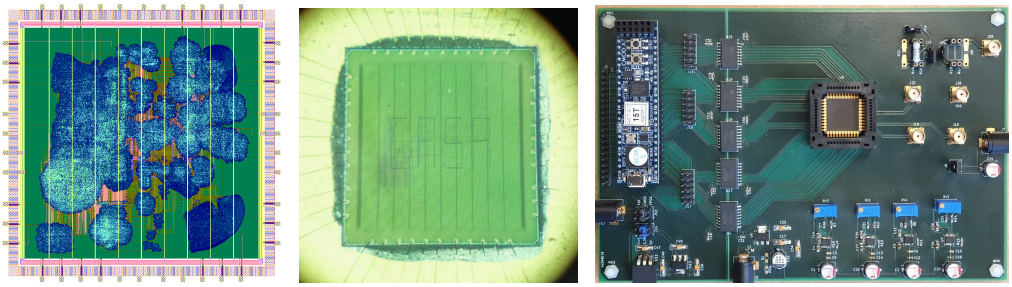

(a) Layout (b) Microscopic photo (c) Evaluation board

Figure 4: Our fabricated ASIC prototype.

Table 1: Post-layout area footprint and estimations of the critical path delay, maximum frequency and average power consumption at 33.3 MHz. The overheads are estimated based on the area of a pure unprotected core following the same design architecture.

Core	Critical Path	Max Freq.	Dynamic Power	Static Power	Area	Overhead Factor
	[ns]	[MHz]	[mW]	[mW]	[GE]	
iMDPL	14.12	70.79	0.8307	0.0028	28278	12.58
SRiMDPL	14.21	70.35	0.8608	0.0030	29298	13.03
TI+iMDPL	13.73	72.81	3.0733	0.0158	104059	46.28
TI+SRiMDPL	13.86	72.11	3.2266	0.0166	108277	48.16

implementing the PRESENT block cipher using various combinations of countermeasures. Although all cores follow the same nibble-serial design architecture described in Section 2.4, two of them are algorithmically masked following the TI concept, while the others are unmasked. For the design architecture (particularly for the TI cores) we followed the designs presented in [PMK+11]. In each category (masked and unmasked), one core is realized by iMDPL cells, and the other core by SRiMDPL which we introduce later. The chip further contains a Pseudo-Random Number Generator (PRNG) randomly seeded during the power-up cycle of the circuit, which is responsible for providing the mask bit (updated every clock cycle) for the iMDPL/SRiMDPL cores. We should highlight that the internal controller of the ASIC chip is constructed in such a way that it allows controlling the PRNG. In other words, we are able to either run the PRNG normally or tie the fresh mask bit (required by the iMDPL/SRiMDPL cores) to either '0' or '1'.

The Synopsys IC design flow was employed to generate the layout of the chip and perform post-layout analysis. The estimated parameters, including critical path delay, maximum operating frequency, and average power consumption (while operating at 33.3 MHz), for all aforementioned cores under typical operating conditions (25 °C and 1.2 V) are shown in Table 1. We further list the size of the area footprint of each core in the table. Note that the overhead factor was estimated based on an unmasked implementation of the same cipher following the nibble-serial architecture realized by the same CMOS standard cell library. As a reference, such an implementation has a post-layout area footprint of 2247 Gate Equivalents (GE). As expected, the combination of TI and DRP leads to significant area overhead. However, the cost for iMDPL and corresponding SRiMDPL cores is very similar. In other words, the tweaks to the iMDPL logic style we propose come at low overhead. Furthermore, the overhead is comparable to the ELB countermeasure presented in [MM21]. We used a serialized implementation of the PRESENT cipher to increase comparability with the state of the art, as multiple related studies have been conducted using that target (e.g., [MM21]). The overhead of the same technique applied to other ciphers will depend

on the concrete design considered and its architecture. We predict that the overhead factor for an AES implementation would be larger because of the size of the Sbox and the cipher state. Although all cores require the same number of 547 clock cycles to accomplish an encryption, the SRiMDPL cores show a slightly reduced maximum frequency and increased average power consumption compared to the corresponding iMDPL cores. It is also clear that the masked cores are more power hungry than the unmasked ones.

3.2.2 Measurement Setup

We designed and constructed an evaluation board consisting of two distinct components, as shown in Figure 4(c). On the left-hand side, there is a 48-pin DIP socket, into which we connect a Digilent Cmod A7 FPGA board. The right-hand side accommodates our ASIC chip in a PLCC44 socket. The FPGA acts as an interface between the ASIC and the PC, managing the measurement process and communicating with the measurement setup via a trigger signal provided through an I/O port of the FPGA. We used a Keithley 2450 Source Measure Unit (SMU) as a power supply and high precision current measurement instrument simultaneously.

Static power measurements exhibit a high degree of temperature dependency. Thus, as commonly done in the state of the art, the evaluation board including the ASIC chip is placed in a climatic chamber in order to precisely control the temperature of the environment. We conducted our entire analyses when the core Vdd of the ASIC chip was being supplied by 1.35 V (12.5 % over-voltage) and the evaluation board operated at a temperature of 90 °C[1].

To establish the communication between the PC and the FPGA board, a UART port is used, which is considerably slower than the other components of the board. As suggested in [MM21], we minimized the communication between the PC and the FPGA through UART to speed up the measurement process. To this end, the FPGA produces n random plaintexts internally, and the FPGA pauses the clock of the ASIC chip for each plaintext during the first round of the PRESENT encryption at the demanded clock cycle and desired level (Low or High). Immediately after holding the chip clock, the FPGA generates a trigger pulse for the SMU. By receiving the positive edge of the trigger pulse, the SMU waits for about 20 ms to let the electrical quantities of the ASIC settle on their DC value. Then, the SMU reads the supplied current value as static leakage, saves it into its internal storage, and waits for another trigger. When awaiting a new trigger pulse, the SMU undergoes a process where the FPGA reruns the chip clock pulse and allows it to finish encrypting the current plaintext (in order to verify the correctness of the calculation) before introducing a new one. This cycle repeats 30,500 times. Once the full set of 30,500 measurements is complete, the PC disregards the first 500 samples (as the chip is notably heating up during this time which affects the measured leakage) and proceeds to read the remaining 30,000 samples (where the temperate is more stable) saved internally on the SMU. Although the measurement period increases by 1.6%, by taking only the last 30,000 samples of the measurements, we have a meaningful decrease in the amount of measurement noise.

3.2.3 Evaluation Criteria

To assess the existence of SCA leakages, we followed the state of the art and conducted fixed-versus-random t-test [CDG+13,SM15]. To this end, for every measurement, based on the result of a flipped coin, the Device Under Test (DUT) is provided with a pre-defined fixed or a randomly selected plaintext while keeping the key constant. Note that for the masked cores, every fixed (resp. random) plaintext and key are freshly masked with three shares right before being sent to the DUT. Such a fixed-versus-random t-test gives only an intuition about the existence of a detectable leakage in the collected measurements.

[1] We also controlled the humidity (set to 10%) which is suggested for measurements at high temperatures.

In order to examine the exploitability of leakages, we performed Moment-Correlating DPA (MCDPA) attacks [MS16]. This type of collision-based attack does not require a hypothetical leakage model, but relies on the collision of values processed in (preferably) time-shared modules. Indeed, MCDPA can recover the linear difference between two key portions, e.g., nibbles, by comparing SCA characteristics of two exemplary Sboxes. This fits well to the design architecture of the underlying cipher of our DUT (see Section 2.4), where a single Sbox module is re-used to process all cipher state nibbles. To perform the attack, we either made use of a half of the aforementioned measurements belonging to the random group or collected a new set of measurements while the plaintext is uniformly selected at random. Note that the same has been previously used in relevant state of the art [MMR17, MM21].

We should highlight that template attacks [CRR02] (as classic profiling attacks) can also be used to analyze different countermeasures. As a side note, a template attack relies on capturing and analyzing the multivariate normal distribution of the SCA leakages (over multiple sample points and multiple clock cycles) which are estimated by a large set of profiling traces and examined on a preferably small set of attack traces. This makes template attacks a suitable evaluation/attack scenario on micro-processor-based implementations, where every operation consists of multiple clock cycles and hence several leakage points. In our study, we deal with the static leakage currents, which are singular values measured at a certain clock cycle (by halting the clock). Although it is generally possible to conduct a template attack by univariate normal distributions, this becomes very similar to an MCDPA attack in the profiling mode [MS16]. Further, colliding and profiling MCDPA attacks are conceptually identical; in the profiling one it is assumed that one key portion is known while the colliding one recovers the linear difference between the corresponding key portions. Therefore, in this work we limit our key-recovery analysis to colliding MCDPA attacks.

3.3 iMDPL, Evaluation Phase

We start with the iMDPL core and analyze its static power when the circuit is in the evaluation phase. As previously mentioned, our ASIC chip's configuration allows us to control the PRNG and choose a specific value for the mask bit. Our initial analysis is performed with the mask bit set to a specific value (emulating a non-masked DRP core), then continues with the mask bit being random.

Constant Mask. We first evaluate the static power of the iMDPL core in the evaluation phase with the mask bit $m = 0$, and repeat the entire procedure with $m = 1$. Based on Section 2.4, the current implementation employs a nibble-serial architecture, meaning that choosing a particular clock cycle targets specific plaintext and secret key nibbles. Therefore, we configured the measurement setup to pause the DUT at the evaluation phase of a clock cycle belonging to the middle of the first encryption round. We collected 5,000 static power measurements for each case ($m = 0$ and $m = 1$) and applied a high-pass filter to suppress the noise[2]. In order to minimize the effect of some spike noise, we set a threshold and discarded any filtered sample with a magnitude exceeding the threshold. Subsequently, the fixed-versus-random t-test was conducted on the DC-free samples leading to the results shown in Figure 5 and Figure 6, for $m = 0$ and $m = 1$, respectively. The t-test results indicate first-order leakage in both cases. Interestingly, the sign of the t-value changes based on the value of the mask m, revealing that the static power in average has a dependency on m. To better understand the effect of m and present this dependency, we took the measurements belonging to the random input and regrouped them based

[2]The same process is done in the state of the art, e.g., by applying a moving average. See [MM21, Moo19, PSKM15].

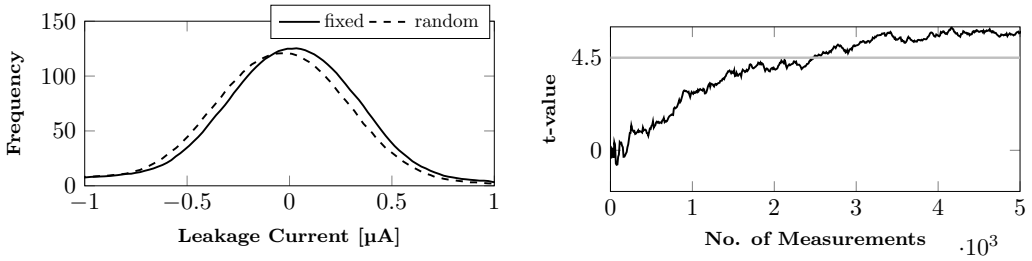

(a) Distribution of filtered data (b) t-test

Figure 5: iMDPL core, leakage assessment using 5,000 static power measurements at the evaluation phase, $m = 0$.

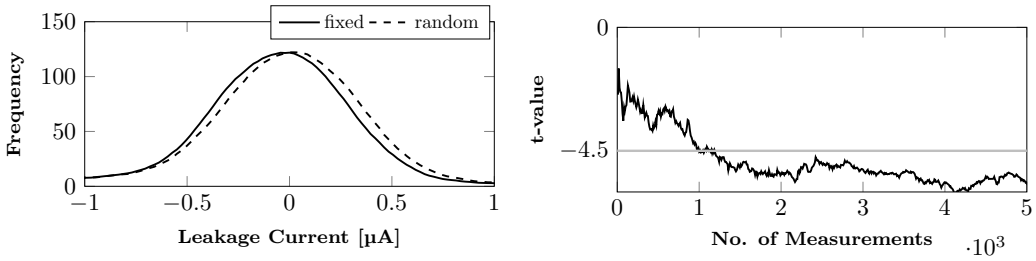

(a) Distribution of filtered data (b) t-test

Figure 6: iMDPL core, leakage assessment using 5,000 static power measurements at the evaluation phase, $m = 1$.

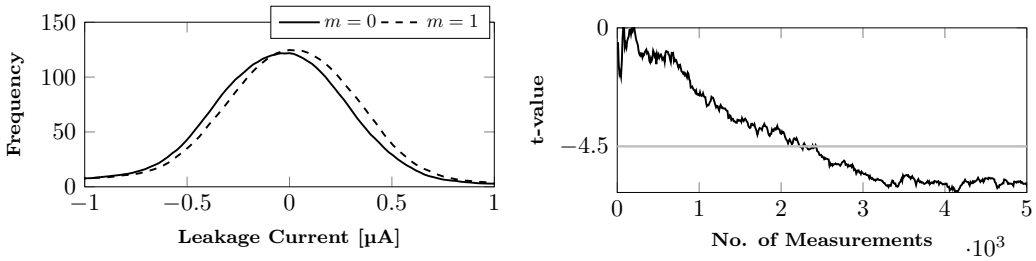

(a) Distribution of filtered data (b) t-test

Figure 7: iMDPL core, dependency of static power measurements on the value of the mask bit m at the evaluation phase.

on the value of m. A comparison done via a t-test, shown in Figure 7, confirms such a dependency.

As stated in Section 3.2.3, we also perform MCDPA key-recovery attacks to examine the exploitability of leakages. To this end, we require to collect static power measurements associated to two distinct clock cycles (both at the evaluation phase) to allow an MCDPA to recover the difference between their corresponding key nibbles. To this end, we collected 5,000 static power measurements for each clock cycle and for each case $m = 0$ and $m = 1$. Figure 8 presents the result of both attacks individually. Expectedly, successful attacks are easily carried out in both cases.

Random Mask. Repeating the same procedure while allowing the PRNG to operate normally (fresh uniformly random mask bit m for each clock cycle) has led to a decrease in data dependency of static power measurements. In order to visualize this, we collected

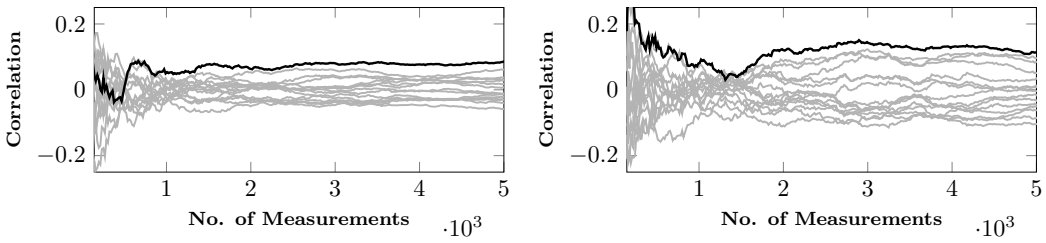

(a) $m = 1$ (b) $m = 0$

Figure 8: iMDPL core, result of MCDPA attack on static power measurements at the evaluation phase while the mask is fixed.

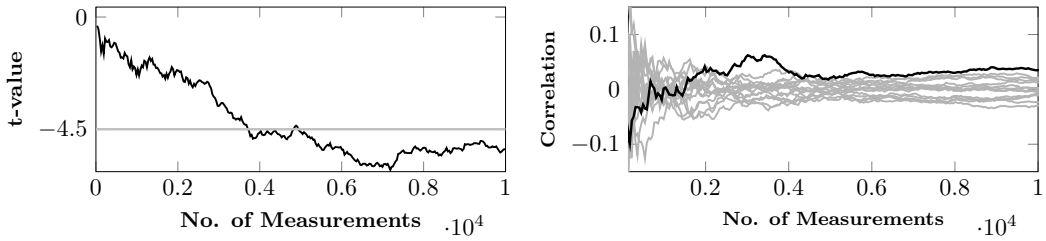

(a) t-test (b) MCDPA

Figure 9: iMDPL core, leakage assessment and MCDPA attack on static power measurements at the evaluation phase for random mask.

10,000 static power measurements following the same scenario given above and performed fixed-versus-random t-test and MCDPA attacks, whose results are shown in Figure 9. It can be seen that the leakage is reduced compared to the other cases with a fixed mask bit, but attacks are still evidently feasible. Comparing our results depicted in Figure 9 with the result of the leakage assessments reported in [KMM19] (which is a study based on a serial implementation of the PRESENT cipher on a 65 nm ASIC using a measurement setup comparable to ours), we can conclude that the experimental analyses are consistent with the simulation results given in Figure 3, i.e., iMDPL in the evaluation phase exhibits more leakage compared to an equivalent CMOS circuit.

3.4 iMDPL, Precharge Phase

When the iMDPL circuit is in the precharge phase, all inputs of the combinational circuit and intermediate signals are set to '0'. This eliminates the static leakage caused by the combinational gates processing the data. However, the circuit should maintain its internal state, i.e., the values stored in the iMDPL flip-flops, whose block diagram is presented in Figure 13(a). As a result, the static leakage of the circuit during the precharge phase is anticipated to be still data dependent. Since data is only stored in the flip-flops and their outputs are precharged, the susceptibility to SPSCA attacks in the precharge phase is expected to be lower than in the evaluation phase.

Constant Mask. Keeping the mask at a certain value and repeating the same measurement and analysis procedures explained in Section 3.3 but in the precharge phase led to the results shown in Figure 10. As expected, the attacks are trivially successful but with a lower correlation and a higher number of required measurements (compared to Figure 8). More precisely, successful attacks at the evaluation phase require around 2,000 measurements while this is increased to 30,000 at the precharge phase.

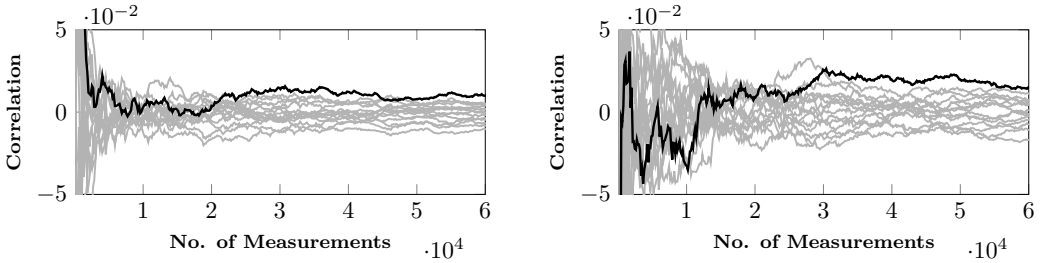

(a) $m = 1$ (b) $m = 0$

Figure 10: iMDPL core, result of MCDPA attack on static power measurements at the precharge phase while the mask is fixed.

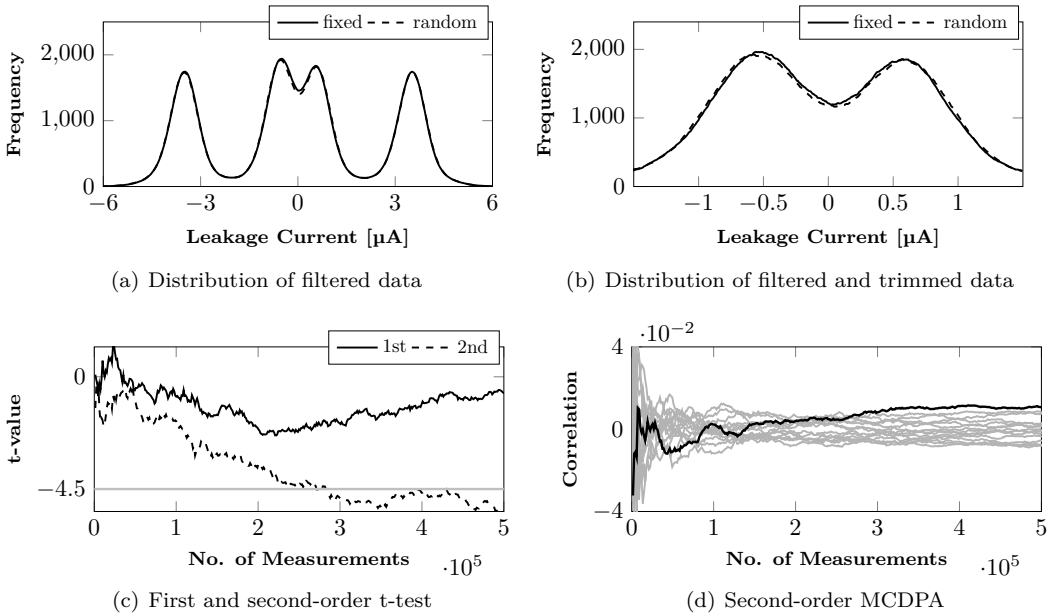

(a) Distribution of filtered data (b) Distribution of filtered and trimmed data

(c) First and second-order t-test (d) Second-order MCDPA

Figure 11: iMDPL core, leakage assessment and MCDPA attack on static power measurements at the precharge phase for random mask.

Random Mask. For the last analysis, we let the internal PRNG operate and randomly choose the mask bit at every clock cycle. We increased the number of measurements due to the predicted lower detectability and exploitability of leakages in the precharge phase. Figure 11(a) shows histograms of such measurements which seem to be the sum of four different Gaussian distributions. Looking at Figure 13(a), it can be seen that every iMDPL flip-flop stores a so-called value d masked with m_n, i.e., $d_{m_n} = d \oplus m_n$. Note that d and m_n are independent of each other, and – as explained above – static power in the precharged phase is mainly due to the values stored in registers d_{m_n}. Further, m_n is routed to all iMDPL flip-flops, and as shown in [MKEP12] it should have a strong effect on power consumption. This justifies four distinct distributions seen in Figure 11(a).

As the first try, we have detected no dependency between these measurements and the classifier fixed/random. Therefore, we tried to trim the measurements based on the observed distributions prior to the analyses. Considering every distribution individually did not lead to any detectable dependency either. We only succeeded by focusing on two middle distributions as shown in Figure 11(b). By this, however, we observed only second-order leakage as it can be seen in Figure 11(c), where 500,000 filtered measurements are used. We predict that the two distributions covered in this analysis belong to different

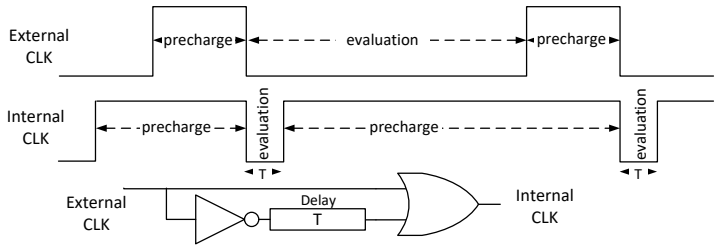

Figure 12: Monostable clock generator.

mask values. However, as they have a slight overlap (see Figure 11(b)), a high number of samples are required to detect the corresponding data dependency. Therefore, in order to conduct a successful MCDPA attack, we pre-processed the trimmed measurements by squaring mean-free samples, as a trivial pre-processing technique for second-order SCA attacks. The result of such a second-order MCDPA attack can be seen in Figure 11(d), which is successful with approximately 300,000 measurements.

4 Our New Strategy

In the following, we present two techniques that, when combined, significantly reduce the static power side-channel exploitability in DRP circuits.

4.1 Monostable Clock

It is demonstrated in Section 3 that the iMDPL circuit exhibits substantially lower data-dependent static power in the precharge phase. This is due to the fact that the input of all DRP combinational gates are set to a constant value '0' that is independent of the data stored in the registers. Hence, the combinational circuit does not have any data-dependent effect on the static power during the precharge phase. Therefore, to perform successful key-recovery attacks, an adversary who controls the clock signal would pause the circuit in the evaluation phase. We aim to prevent this. One option could be to force the circuit into the precharge phase whenever the clock signal is paused, even when the adversary tries to pause the circuit in the evaluation phase.

A standard clock pulse, as defined in electronics, is generated by an astable oscillator, which has no stable state. However, if the attacker gains control of the clock, it is provided by a bistable circuit that has two stable states high and low. By placing a monostable circuit inside the chip right after the external clock pin (which is controlled by the adversary), the clock given to the DRP circuit always reaches its stable state, regardless of how the external clock is controlled whether it is from an astable or bistable source. Figure 12 shows the timing diagram and circuit that illustrates the concept. Regardless of how long the external clock remains high or low, the monostable clock circuitry generates a high-low-high pulse that lasts T and is synchronized to the high-to-low transition on the external clock. This ensures that the DRP circuit does not remain in the evaluation phase, preventing corresponding static power from being measured externally. It is important to note that the shown monostable circuitry operates under the assumption that the DRP circuit is in the evaluation phase when the clock is at a low level, which is true in iMPDL (see Figure 13(a)). Otherwise, the circuit shown in Figure 12 can be easily modified by replacing the final OR gate with an AND gate. The delay element in the mono-stable clock circuit can be implemented by connecting a couple of inverters in series. Knowing the critical path delay of the targeted circuit, we used the Synopsis HSPICE simulator along with the same underlying 65 nm commercial library to determine the minimum number of required inverters. We then added 25% to this number to ensure that the pulse

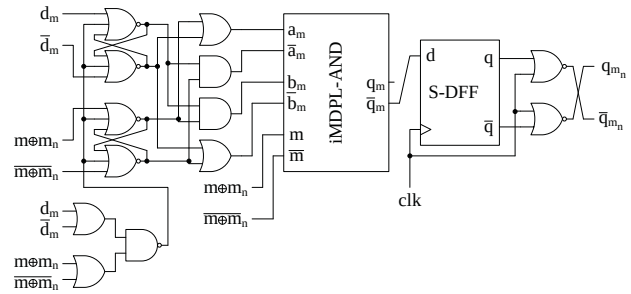

(a) Original iMDPL, taken from [PKZM07]

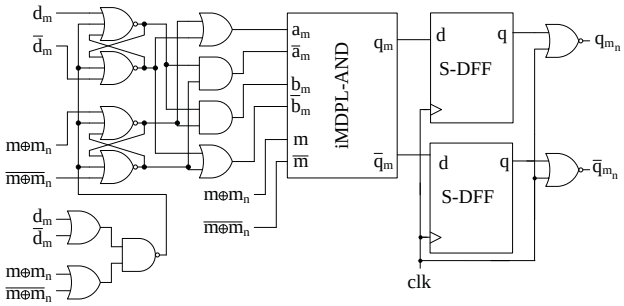

(b) Our proposed SRiMDPL

Figure 13: iMDPL and SRiMDPL flip-flops.

width is sufficiently high in the final implementation. This consideration accounts for process variations and other parameters that may affect the propagation delay, ensuring the required delay time is reliably met.

Selecting the appropriate duration of delay T is crucial. The duration of T must be long enough to permit the combinational circuit to complete the evaluation process, which means it should surpass the critical path delay of the circuit. However, making T significantly larger than the critical path delay would result in two drawbacks. First, it would limit the maximum clock frequency, subsequently reducing the throughput of the circuit. Second, it would create a gap between the completion of the evaluation process and the beginning of the precharge phase, potentially allowing an adversary to measure the static leakage of the circuit in the evaluation phase. Therefore, a particular attention should be paid to adjust T carefully to avoid these issues.

It might be criticized that preventing the adversary from controlling the clock signal by generating the clock internally might be a potential solution to prevent SPSCA attacks. In addition to the difficulties and challenges associated to generate a stable and jitter-free clock inside the chip, which mandates the designer to count on an external clock or oscillator, we should refer to [Moo20], where it has been shown that when the state of the circuit is not cleared after the termination of the cryptographic operation, the adversary does not require to control the clock and hold the circuit in any particular state in order to conduct successful SPSCA attacks.

4.2 Dual-Rail Flip-Flop (D-DFF)

Utilizing the monostable clock circuitry, we can restrict the adversary to measure only the static power of the circuit in the precharge phase. As demonstrated in Section 3.4, even weak data dependencies of static power in the precharge phase can be exploited by second-order attacks. The root cause of this type of leakage, as previously indicated, is

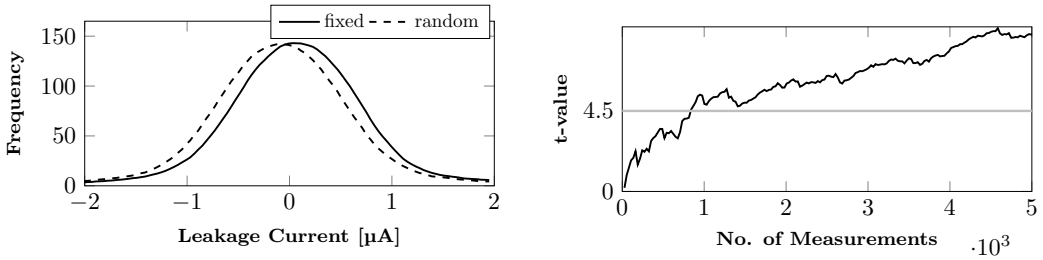

(a) Distribution of filtered data

(b) t-test

Figure 14: SRiMDPL core, leakage assessment using 5,000 static power measurements at the evaluation phase with a fixed mask $m = 0$.

the use of a single-rail flip-flop (S-DFF) in iMDPL flip-flops. As given in Figure 13(a), the output of the combinational circuit $(d_m, \overline{d}_m)$ is re-masked by $m \oplus m_n$, with m_n being the value of the mask m in the next clock cycle. Thus, the value stored in the S-DFF, namely d_{m_n}, is masked with m_n required for the subsequent evaluation phase. Therefore, the joint leakage associated to d_{m_n} and m_n is exploited through a second-order SPSCA attack. In order to mitigate this, a dual-rail flip-flop (D-DFF) should be employed in every iMDPL flip-flop. This can be realized by instantiating an additional S-DFF as shown in Figure 13(b). Hence, d_{m_n} and $\overline{d}_{m_n}$ are always stored. This means that both $(0, 1)$ and $(1, 0)$ are stored in the dual-rail flip-flops, regardless of the value of d and m_n. This would potentially mitigate such a second-order leakage, ignoring intra-die process variations. In other words, if both S-DFFs of a D-DFF are realized identically in silicon, the exploitability of the second-order leakage at the precharge phase should be prevented. Since such process variations cannot be entirely avoided, there will be a tiny difference between the static power of a D-DFF when complementary values are stored, i.e., $(0, 1)$ versus $(1, 0)$. Hence, since the exploitability of higher-order leakages is exponentially affected by the noise level, the second-order leakage associated with $(d_{m_n}, \overline{d}_{m_n})$ and m_n cannot be entirely avoided but is significantly reduced in practice. By combining monostable clock and D-DFFs, we can substantially reduce the exploitability of static leakage in DRP circuits. We would also like to highlight that the overhead of these two techniques is highly minimal. The area footprint and energy consumption of the monostable clock circuitry is negligible compared to that of a cryptographic core. Further, the combinational part of the circuit stays unchanged, and an extra S-DFF is instantiated inside every iMDPL flip-flop, which is also negligible considering the number of gates involved in every iMDPL cell. It is worth mentioning that the outputs of Figure 13(a) are crossed, which is different from the outputs of Figure 13(b). This is because in the original iMDPL flip-flop, the input to the S-DFF comes from the inverted output of the iMDPL-AND. In contrast, we employ two S-DFFs in SRiMDPL to save q_m and $\overline{q}_m$ individually. Therefore, such a crossed wiring is not necessary in SRiMDPL.

4.3 Benefit of D-DFF in SRiMDPL

Via the monostable clock module, our SRiMDPL core is designed to prevent the adversary from holding the circuit in the evaluation phase. However, for analysis purposes, we designed our prototype core to allow us to control the circuit and measure its static power during both evaluation and precharge phases individually. Here we analyze the advantage of using D-DFFs in the SRiMDPL flip-flops. To this end, we set the mask bit $m = 0$ to solely evaluate the effect on complementary values stored in the SRiMDPL on static power measurements. The corresponding analysis results conducted in the evaluation phase are shown in Figure 14(b). Compared to Figure 6(b), it is evident that using D-DFFs is

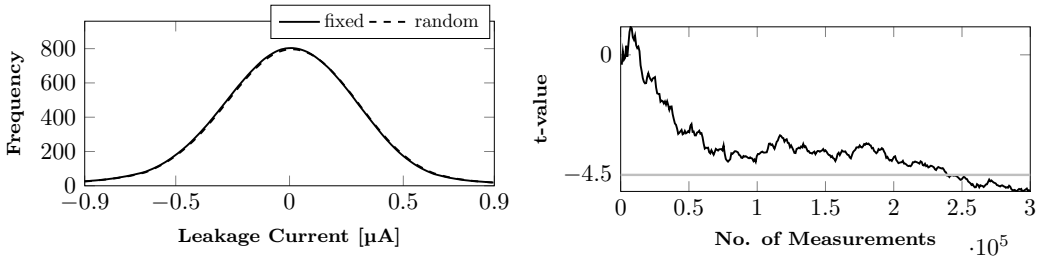

(a) Distribution of filtered data (b) t-test

Figure 15: SRiMDPL core, leakage assessment using 300,000 static power measurements at the precharge phase with a fixed mask $m = 0$.

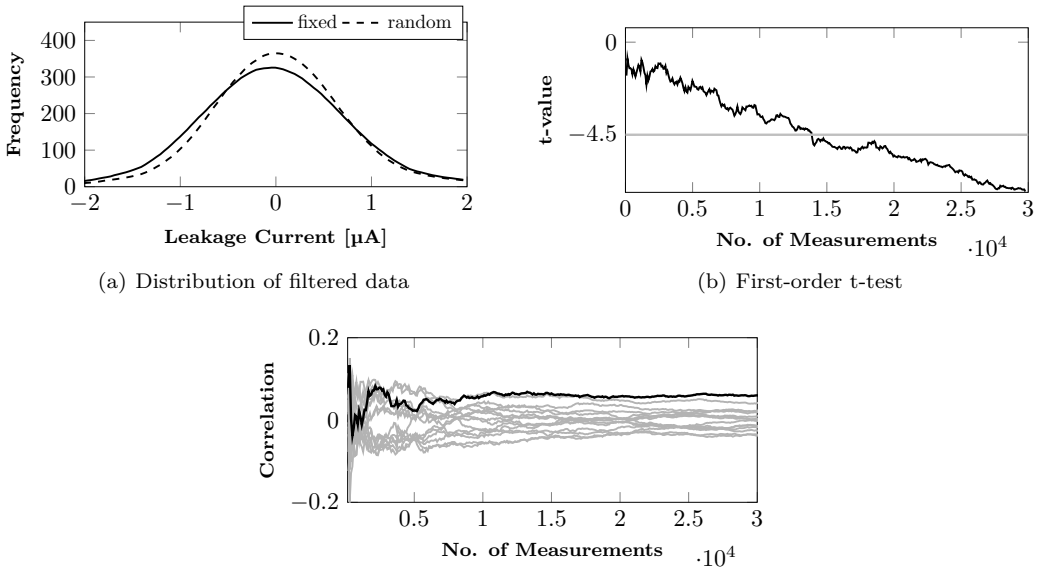

(a) Distribution of filtered data (b) First-order t-test

(c) First-order MCDPA

Figure 16: SRiMDPL core, leakage assessment and MCDPA attack using 30,000 static power measurements at the evaluation phase for random mask.

not beneficial to reduce the first-order leakage in the evaluation phase. This is indeed expected, as data-dependency of static power in the evaluation phase mainly originates from the combinational circuit, while the flip-flops have a comparatively smaller effect on static power in the evaluation phase. The benefit of using D-DFFs is however expected to be more visible in the precharge phase where all combinational gates receive a constant pre-defined value '0'. Figure 15(a) confirms this expectation. Leakage assessments show a significant reduction in the detectable first-order leakage. In short, at least 280,000 measurements are required to detect the leakage while 30,000 measurements are required in the analysis conducted on the original design under the same condition (see Figure 10).

4.4 Overall Advantages of SRiMDPL

Above we have presented the benefit of using D-DFF in the precharge phase, when the mask bit was fixed. Here examine the static leakage of a complete SRiMDPL core while the mask bit is updated randomly at every clock cycle. In such conditions, we collected around 1,000,000 static power measurements in each evaluation and precharge phase. According

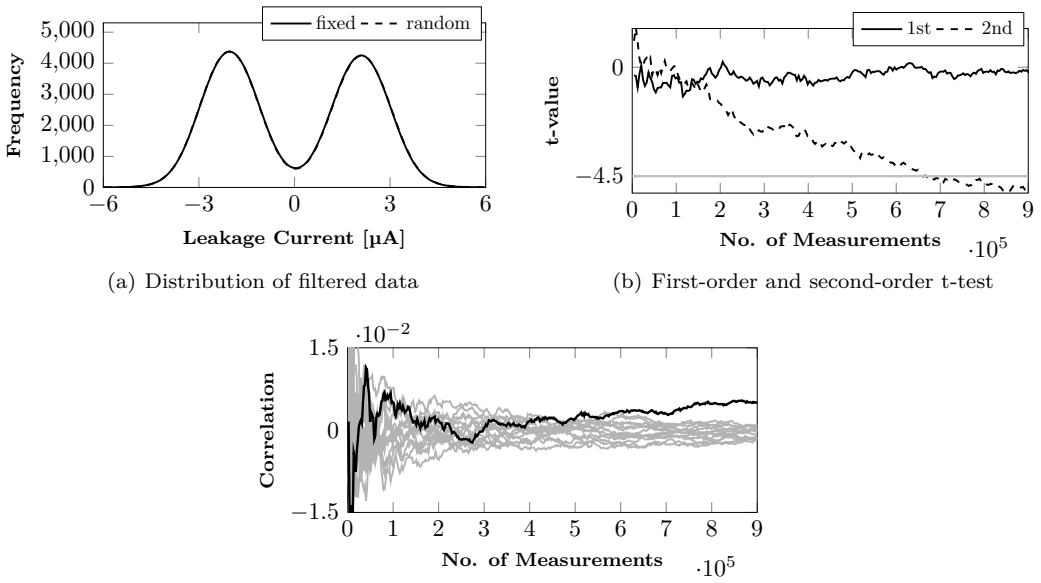

(a) Distribution of filtered data

(b) First-order and second-order t-test

(c) Second-order MCDPA

Figure 17: SRiMDPL core, leakage assessment and MCDPA attack using 900,000 static power measurements at the precharge phase for random mask.

Table 2: Minimum number of required measurements and corresponding acquisition time in minutes to detect the data dependency in static power measurements of both iMDPL and SRiMDPL cores in precharge and evaluation phase.

Core	**precharge** (second-order leakage)		**evaluation** (first-order leakage)	
	Measurements	**Time**	**Measurements**	**Time**
iMDPL	285,000	855 min.	5,000	15 min.
SRiMDPL	700,000	2100 min.	15,000	45 min.

to the results shown in Figure 16, a successful attack during the evaluation phase needs approximately 15,000 measurements, which is not much different to the result of the same analysis on the original iMDPL core, where around 5,000 measurements are required (see Figure 9). However, the situation is highly different in the precharge phase, i.e., Figure 17. The first to consider is the number of observable distributions in the histogram of the static power measurements, i.e., Figure 11(a) versus Figure 17(a). As explained in Section 3.4, four distinct distributions correspond to four different values for d_{m_n} and m_n. Since $\overline{d}_{m_n}$ accompanies d_{m_n} in SRiMDPL flip-flops, the effect of d_{m_n} on static power measurements is highly diminished. Therefore, static power measurements are mainly affected by the value of m_n connected to all SRiMDPL flip-flops, which is confirmed by two distinct distributions visible in Figure 17(a). The result of leakage assessment and MCDPA attack also imply a significantly higher number of required measurements to detect and exploit second-order leakages while first-order leakages are prevented. More precisely, around 700,000 measurements are required compared to around 300,000 measurements for the iMDPL core (see Figure 11). These results are summarized in Table 2.

5 Combination with Threshold Implementation

The primary source of static leakage in SRiMDPL is technology-induced intra-die process variation, which degrades the balance of dual rails. Although this imbalance is randomized (thanks to the underlying iMDPL concept), this leaves an opportunity for the adversary to exploit second-order leakages. As a result, depending on the application and where the cryptographic core is supposed to be used, it may not be sufficient to rely solely on such a hiding countermeasure method. Therefore, here we evaluate the effectiveness of algorithmically masked cores (TI) realized by DRP logic styles. As explained in Section 3.2.1 and Section 2.5, our prototype ASIC chip contains two TI cores, implemented by iMDPL and SRiMDPL cells. In order to evaluate the static leakage of such combinations, we start with fixing the mask bit $m = 0$. This eliminates the random switch between the balanced rails and somehow emulates a TI circuit constructed by WDDL cells [TV04]. Note that a similar TI core implemented by a standard CMOS cell library has been evaluated in [MM21]. Since the core is a realization of a first-order secure masking with three shares, we expect to observe no first-order leakage, which is confirmed by the analysis results depicted in Figure 18. Standard first-order TIs with three shares are expected to exhibit univariate second-order leakage originating from the non-linear operations, since each corresponding combinational circuit is only first-order non-complete, i.e., independent of one share. They are also expected to exhibit univariate third-order leakage originating from the linear operations (and movement of data) performed on all three shares independently, but simultaneously. Higher-order biases are known to require more traces to become detectable, as their estimation is exponentially affected by the noise level [RGV12, CJRR99, SP06]. Hence, it is not surprising that on our target we consistently observe that the bias in the variance of the distributions (second order) is detected earlier than in the skewness (third order). However, after evaluating sufficiently many traces, both the second- and third-order t-values are expected to exceed the threshold. We can view this phenomenon in Figure 18 where the second-order t-value exceeds the threshold while the third-order t-value is still below the threshold. Increasing the number of measurements sufficiently is expected to raise the third-order t-value also above the threshold.

We should also highlight that fixing the mask bit $m = 0$ is independent of masked plaintext, which is freshly masked with three shares before the start of every encryption. Therefore, as expected we detected second-order leakage in all cores, of course with various number of measurements. The effect of D-DFFs in the SRiMDPL can be seen by comparing Figure 18(a) and Figure 18(b) in the evaluation phase and through Figure 18(c) versus Figure 18(d) in the precharge phase. More precisely, 100,000 measurements which reveal the data dependency of static power measurements of the TI+iMDPL core in the evaluation phase increase to 350,000 for the TI+SRiMDPL core. This advantage is also seen in the precharge phase, i.e., 200,000 measurements versus 900,000.

In order to examine the nominal condition when the PRNG is allowed to randomly change the mask bit m, we had to collect many more measurements to ensure the achieved level of security. To this end, we collected more than 12,000,000 static power measurements of the TI+SRiMDPL core in both evaluation and precharge phases. The results shown in Figure 19 indicate that second-order leakages are still detectable in the evaluation phase using around 1,000,000 measurements while we have not detected any leakage in up to the third-order statistical moments using 12,000,000 static power measurements in the precharge phase. This highlights the need for our proposed monostable clock module to avoid the adversary being able to collect static power SCA measurements in the evaluation phase.

Dynamic Power. This combination which employs a provably secure first-order masking and power-equalization hiding scheme should provide a high level of security against classical dynamic power SCA attacks as well. To examine this, we measured the dynamic

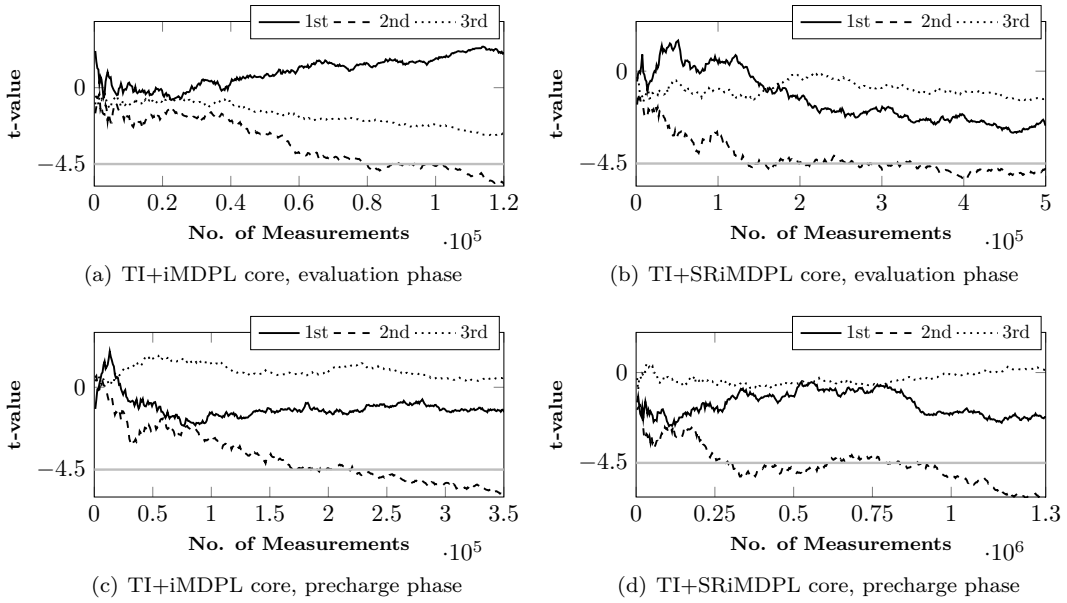

(a) TI+iMDPL core, evaluation phase

(b) TI+SRiMDPL core, evaluation phase

(c) TI+iMDPL core, precharge phase

(d) TI+SRiMDPL core, precharge phase

Figure 18: TI+iMDPL and SRiMDPL cores, first- and second-order leakage assessment results with a fixed mask $m = 0$.

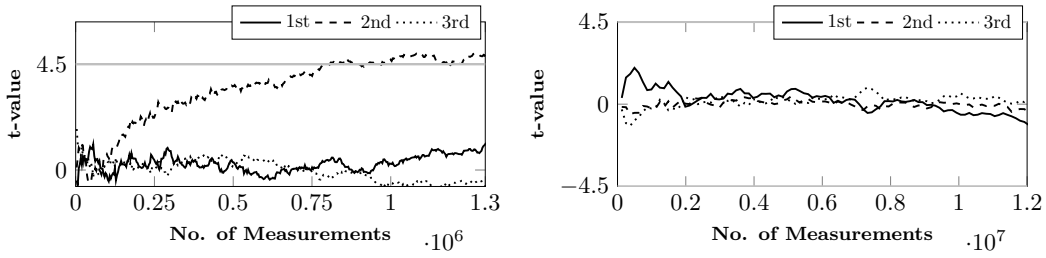

(a) evaluation phase

(b) precharge phase

Figure 19: TI+SRiMDPL core, leakage assessment on static power measurements for random mask.

power consumption of the cores by means of a digital oscilloscope at a sampling rate of 2 GS/s and monitoring the voltage drop over a 1 Ω shunt resistor placed in the core Vdd path of the ASIC. Due to the very low amplitude of the signal, we also employed an AC amplifier with a gain of 10 dB to collect the traces with lower measurement noise. During the measurement, the DUT was supplied by a stable clock source at a frequency of 6 MHz to avoid any other switching noise originating from overlapping adjacent power peaks. We followed the procedure suggested in [SM15] to speed up the measurement procedure and collect traces suitable for fixed-versus-random t-test on dynamic power traces. The results of this analysis are shown in Figure 20, where we evaluated iMDPL, SRiMDPL, and TI+SRiMDPL cores. While the first-order leakages of the unmasked cores iMDPL and SRiMDPL cores can be detected using less than 250,000 traces, the TI+SRiMDPL core does not exhibit any leakage even using 100,000,000 traces.

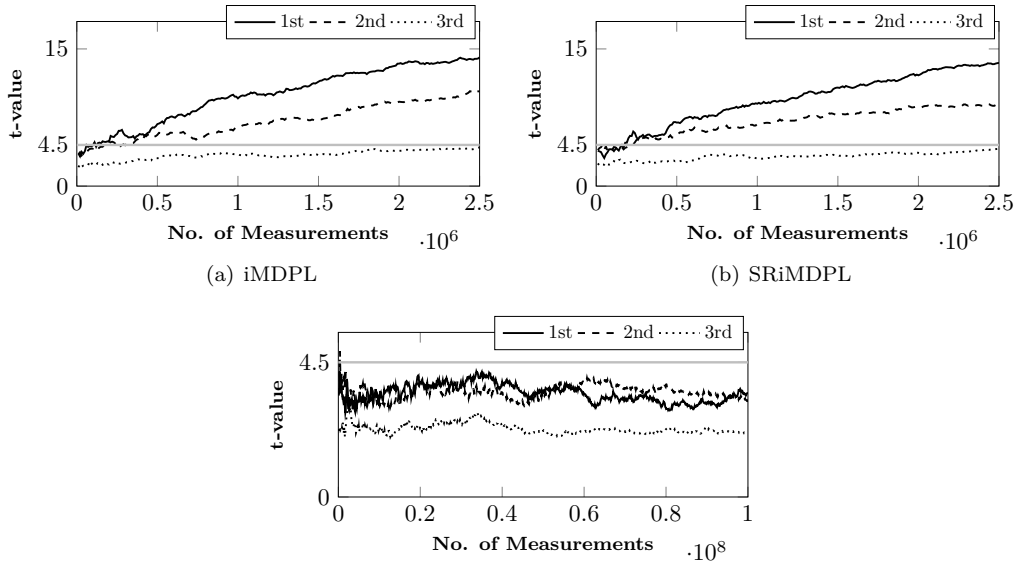

(a) iMDPL (b) SRiMDPL

(c) TI+SRiMDPL

Figure 20: Leakage assessment on dynamic power measurements for random mask.

6 Conclusions

In recent years, it has been shown that an adversary can exploit the static power consumption of a CMOS cryptographic core as a side channel to extract its secrets. Several works have shown that common countermeasures against dynamic power SCA attacks are ineffective in hindering SPSCA attacks. Consequently, it is crucial to integrate specific countermeasures to prevent static power adversaries from exploiting the backdoor and bypassing the most resilient dynamic power countermeasures. Practical experiments are necessary to find effective solutions to withstand severe conditions in real-world scenarios. In this work, we tried to find and evaluate solutions for this challenge.

To this end, we designed and fabricated a 65 nm ASIC prototype chip allowing us to explore the advantages and disadvantages of DRP logic styles from the static power side channel point of view. This study, supported by experimental analysis on a real fabricated silicon, is the first of its kind to investigate SPSCA attacks on these logic styles. Focusing on DRP circuits, we realized that the data dependency of the static power side channel is much harder to exploit when the DRP circuit is in the precharge phase, when the inputs of all combinational gates are set to a predefined value. Therefore, we introduced a monostable circuitry to be integrated inside the chip which prevents the adversary from keeping the DRP circuit in the evaluation phase. Although this approach is valid for and should be beneficial in any DRP circuit, we evaluated its effectiveness using the well-known DPA-resistant logic style iMDPL. To further reduce the data-dependent static power dissipation when the circuit is in the precharge phase, we also proposed a modification to the iMDPL flip-flops. Although this improves the situation and increases the resilience of the circuits, our analyses demonstrate that none of the hiding and masking countermeasures employed alone can provide sufficient protection against SPSCA attacks. The most promising results have been achieved by utilizing both countermeasures simultaneously. Although such a combination leads to a very notable overhead, our experimental analyses revealed no leakage (even at higher orders) using 12,000,000 static power measurements and 100,000,000 dynamic power traces. The collection of 12,000,000 static power samples took 25 days of non-stop measurement using our experimental setup highlighting the achieved level of

practical security.

To the best of our knowledge, similar results have never been reported. Even the best solutions compared in [MM21], which are also combined constructions, still exhibit detectable leakage when enough static power measurements were available. It is worth noting that static leakage in other CMOS technologies, particularly smaller ones, can be significantly higher than in the 65 nm technology we studied. Hence, the concrete numbers of measurements required to break the targeted circuits will not be perfectly transferable to other technology generations. However, we believe it is fair to assume that even in newest nanometer generations, employing SRiMDPL and forcing the adversary to target the precharge phase will make attacks significantly less efficient. Especially when combined with secure hardware masking, such a countermeasure will force attackers to obtain a high, likely difficult to obtain, number of measurements to perform a successful SPSCA. As a result, the remaining challenge for future work is to maintain the same level of security while diminishing the overheads, particularly the area consumption.

Acknowledgments

The work described in this paper has been supported in part by the German Research Foundation (DFG) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972, and by European Union through the Twinning Project 101079319 (acronym enCRYPTON) and the ERC Advanced Grant 101096871 (acronym BRIDGE). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

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. *IEEE Trans. on Circuits and Systems*, 61-I(2):429–442, 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, 2014.

[AGST09] Massimo Alioto, Luca Giancane, Giuseppe Scotti, and Alessandro Trifiletti. Leakage Power Analysis attacks: Theoretical analysis and impact of variations. In *ICECS 2009*, pages 85–88. IEEE, 2009.

[AGST10] Massimo Alioto, Luca Giancane, Giuseppe Scotti, and Alessandro Trifiletti. Leakage Power Analysis Attacks: A Novel Class of Attacks to Nanometer Cryptographic Circuits. *IEEE Trans. on Circuits and Systems*, 57-I(2):355–367, 2010.

[AMEA02] Mohab Anis, Mohamed Mahmoud, Mohamed I. Elmasry, and Shawki Areibi. Dynamic and leakage power reduction in MTCMOS circuits using an automated efficient gate clustering technique. In *DAC 200*, pages 480–485. ACM, 2002.

[BBM+17] 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. *IEEE Trans. Emerg. Top. Comput.*, 5(3):329–339, 2017.

[BCO04] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation Power Analysis with a Leakage Model. In *CHES 2004*, volume 3156 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2004.

[BCS+17] Davide Bellizia, Danilo Cellucci, Valerio Di Stefano, Giuseppe Scotti, and Alessandro Trifiletti. Novel measurements setup for attacks exploiting static power using DC pico-ammeter. In *ECCTD 2017*, pages 1–4. IEEE, 2017.

[BDF+17] Gilles Barthe, François Dupressoir, Sebastian Faust, Benjamin Grégoire, François-Xavier Standaert, and Pierre-Yves Strub. Parallel Implementations of Masking Schemes and the Bounded Moment Leakage Model. In *EUROCRYPT 2017*, volume 10210 of *Lecture Notes in Computer Science*, pages 535–566, 2017.

[BFS20] Jan Belohoubek, Petr Fiser, and Jan Schmidt. Standard Cell Tuning Enables Data-Independent Static Power Consumption. In *DDECS 2020*, pages 1–6. IEEE, 2020.

[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 *CHES 2007*, volume 4727 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2007.

[BLA+19] Tomasz Brozek, Meindert Lunenborg, Franck Arnaud, Roberto Gonella, Jean-Christophe Giraudin, Christian Dutto, Bertrand Martinet, Laurent Garchery, Christopher Hess, and Kelvin Doong. Characterization Challenges and Solutions for FDSOI Technologies. In *SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S) 2019*, pages 1–3, 2019.

[CDG+13] Jeremy Cooper, Elke Demulder, Gilbert Goodwill, Joshua Jaffe, Gary Kenworthy, and Pankaj Rohatgi. Test Vector Leakage Assessment (TVLA) Methodology in Practice. International Cryptographic Module Conference, 2013.

[CGLS20] Gaëtan Cassiers, Benjamin Grégoire, Itamar Levi, and François-Xavier Standaert. Hardware Private Circuits: From Trivial Composition to Full Verification. *IEEE Transactions on Computers*, 70:1677–1690, 2020.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards Sound Approaches to Counteract Power-Analysis Attacks. In *CRYPTO 1999*, volume 1666 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 1999.

[CRA+06] A.N. Chandorkar, Ch. Ragunandan, Pradyumna Agashe, Dinesh Sharma, and Hiroshi Iwai. Impact of Process variations on Leakage Power in CMOS Circuits in Nano Era (Invited paper). In *Conference on Solid-State and Integrated Circuit Technology 2006*, pages 1248–1251, 2006.

[CRB+16] Thomas De Cnudde, Oscar Reparaz, Begül Bilgin, Svetla Nikova, Ventzislav Nikov, and Vincent Rijmen. Masking AES with d+1 Shares in Hardware. In *CHES 2016*, volume 9813 of *Lecture Notes in Computer Science*, pages 194–212. Springer, 2016.

[CRR02] Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template Attacks. In *CHES 2002*, volume 2523 of *Lecture Notes in Computer Science*, pages 13–28. Springer, 2002.

[CS20] Gaëtan Cassiers and François-Xavier Standaert. Trivially and Efficiently Composing Masked Gadgets With Probe Isolating Non-Interference. *IEEE Trans. Inf. Forensics Secur.*, 15:2542–2555, 2020.

[CS21] Gaëtan Cassiers and François-Xavier Standaert. Provably Secure Hardware Masking in the Transition- and Glitch-Robust Probing Model: Better Safe than Sorry. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(2):136–158, 2021.

[DGS+11] Milena Djukanovic, Luca Giancane, Giuseppe Scotti, Alessandro Trifiletti, and Massimo Alioto. Leakage Power Analysis attacks: Effectiveness on DPA resistant logic styles under process variations. In *ISCAS 2011*, pages 2043–2046. IEEE, 2011.

[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.

[FML+03] Jacques J. A. Fournier, Simon W. Moore, Huiyun Li, Robert D. Mullins, and George S. Taylor. Security Evaluation of Asynchronous Circuits. In *CHES 2003*, volume 2779 of *Lecture Notes in Computer Science*, pages 137–151. Springer, 2003.

[FMM21] Bijan Fadaeinia, Thorben Moos, and Amir Moradi. Balancing the Leakage Currents in Nanometer CMOS Logic – A Challenging Goal. *Applied Sciences*, 11(15), 2021.

[GHMP05] Sylvain Guilley, Philippe Hoogvorst, Yves Mathieu, and Renaud Pacalet. The "Backend Duplication" Method. In *CHES 2005*, volume 3659 of *Lecture Notes in Computer Science*, pages 383–397. Springer, 2005.

[GIB18] Hannes Groß, Rinat Iusupov, and Roderick Bloem. Generic Low-Latency Masking in Hardware. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(2):1–21, 2018.

[GM17] Hannes Groß and Stefan Mangard. Reconciling d+1 Masking in Hardware and Software. In *CHES 2017*, volume 10529 of *Lecture Notes in Computer Science*, pages 115–136. Springer, 2017.

[GM18] Hannes Gross and Stefan Mangard. A unified masking approach. *Journal of Cryptographic Engineering*, 8:109 – 124, 2018.

[GMK16] Hannes Groß, Stefan Mangard, and Thomas Korak. Domain-Oriented Masking: Compact Masked Hardware Implementations with Arbitrary Protection Order. In *TIS@CCS 2016*, page 3. ACM, 2016.

[GMK17] Hannes Groß, Stefan Mangard, and Thomas Korak. An Efficient Side-Channel Protected AES Implementation with Arbitrary Protection Order. In *CT-RSA 2017*, volume 10159 of *Lecture Notes in Computer Science*, pages 95–112. Springer, 2017.

[GMO01] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic Analysis: Concrete Results. In *CHES 2001*, volume 2162 of *Lecture Notes in Computer Science*, pages 251–261. Springer, 2001.

[GPKKB23] David Gavini, E. Pallavi, B. Kiran Kumar, and Pavankumar Bikki. Study of Leakage Currents in FinFET SRAM Cells. In *Conference on Inventive Material Science Applications*, pages 101–111. Springer, 2023.

[GSST07] Jacopo Giorgetti, Giuseppe Scotti, Andrea Simonetti, and Alessandro Trifiletti. Analysis of data dependence of leakage current in CMOS cryptographic hardware. In *GLSVLSI 2007*, pages 78–83. ACM, 2007.

[Har18] David L. Harame. RF FDSOI Technology and Modelling. In *ESSCIRC 2018*, page 214. IEEE, 2018.

[HMY15] Basel Halak, Julian Murphy, and Alex Yakovlev. Power balanced circuits for leakage-power-attacks resilient design. In *SAI 2015*, pages 1178–1183, 2015.

[HSN04] Domenik Helms, Eike Schmidt, and Wolfgang Nebel. Leakage in CMOS Circuits - An Introduction. In *PATMOS 2004*, volume 3254 of *Lecture Notes in Computer Science*, pages 17–35. Springer, 2004.

[HWL20] Ru Huang, Runsheng Wang, and Ming Li. *Gate-All-Around Silicon Nanowire Transistor Technology*, pages 89–115. Springer International Publishing, Cham, 2020.

[JIA+15] Darshana Jayasinghe, Aleksandar Ignjatovic, Jude Angelo Ambrose, Roshan G. Ragel, and Sri Parameswaran. QuadSeal: Quadruple algorithmic symmetrizing countermeasure against power based side-channel attacks. In *CASES 2015*, pages 21–30. IEEE, 2015.

[KJJ99] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In *CRYPTO 1999*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397. Springer, 1999.

[KME+08] Mehrdad Khatir, Amir Moradi, Alireza Ejlali, Mohammad T. Manzuri Shalmani, and Mahmoud Salmasizadeh. A secure and low-energy logic style using charge recovery approach. In *ISLPED 2008*, pages 259–264. ACM, 2008.

[KMM19] Naghmeh Karimi, Thorben Moos, and Amir Moradi. Exploring the Effect of Device Aging on Static Power Analysis Attacks. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(3):233–256, 2019.

[Koc96] Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *CRYPTO 1996*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113. Springer, 1996.

[KP09] Mario Kirschbaum and Thomas Popp. Evaluation of a DPA-Resistant Prototype Chip. In *ACSAC 2009*, pages 43–50. IEEE Computer Society, 2009.

[LB08] Lang Lin and Wayne P. Burleson. Leakage-based differential power analysis (LDPA) on sub-90nm CMOS cryptosystems. In *ISCAS 2008*, pages 252–255. IEEE, 2008.

[LMW14] Andrew J. Leiserson, Mark E. Marson, and Megan A. Wachs. Gate-Level Masking under a Path-Based Leakage Metric. In *CHES 2014*, volume 8731 of *Lecture Notes in Computer Science*, pages 580–597. Springer, 2014.

[MB21] G. Munirathnam and Y. Murali Mohan Babu. Analysis of Static Power Reduction Strategies in Deep Submicron CMOS Device Technology for Digital Circuits. In *ISPCC 2021*, pages 278–282, 2021.

[MH19] Anala M. and B. P. Harish. Process Variation-Aware Analytical Modeling of Subthreshold Leakage Power. In *PATMOS 2019*, pages 119–124. IEEE, 2019.

[MKEP12] Amir Moradi, Mario Kirschbaum, Thomas Eisenbarth, and Christof Paar. Masked Dual-Rail Precharge Logic Encounters State-of-the-Art Power Analysis Methods. *IEEE Trans. Very Large Scale Integr. Syst.*, 20(9):1578–1589, 2012.

[MKSS09] Amir Moradi, Mehrdad Khatir, Mahmoud Salmasizadeh, and Mohammad T. Manzuri Shalmani. Charge recovery logic as a side channel attack countermeasure. In *ISQED 2009*, pages 686–691. IEEE Computer Society, 2009.

[MM21] Thorben Moos and Amir Moradi. Countermeasures against Static Power Attacks - Comparing Exhaustive Logic Balancing and Other Protection Schemes in 28 nm CMOS -. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(3):780–805, 2021.

[MMR17] Thorben Moos, Amir Moradi, and Bastian Richter. Static power side-channel analysis of a threshold implementation prototype chip. In *DATE 2017*, pages 1324–1329. IEEE, 2017.

[MMR19] T. Moos, A. Moradi, and B. Richter. Static Power Side-Channel Analysis–An Investigation of Measurement Factors. *IEEE Trans. on VLSI*, pages 1–14, 2019.

[MMSS19] Thorben Moos, Amir Moradi, Tobias Schneider, and François-Xavier Standaert. Glitch-Resistant Masking Revisited - or Why Proofs in the Robust Probing Model are Needed. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019:256–292, 2019.

[Moo19] Thorben Moos. Static Power SCA of Sub-100 nm CMOS ASICs and the Insecurity of Masking Schemes in Low-Noise Environments. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(3):202–232, 2019.

[Moo20] Thorben Moos. Unrolled Cryptography on Silicon A Physical Security Analysis. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(4):416–442, 2020.

[MOP07] Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks - Revealing the Secrets of Smart Cards*. Springer, 2007.

[Mor14] Amir Moradi. Side-Channel Leakage through Static Power - Should We Care about in Practice? In *CHES 2014*, volume 8731 of *Lecture Notes in Computer Science*, pages 562–579. Springer, 2014.

[MRH+17] Hans Mertens, Romain Ritzenthaler, Andriy Yakovitch Hikavyy, Min-Soo Kim, Zheng Tao, Kurt Wostyn, Tom Schram, Eddy Kunnen, Lars Åke Ragnarsson, Harold F. W. Dekkers, Toby Hopf, Katia Devriendt, Diana Tsvetanova, Soon Aik Chew, Yoshiaki Kikuchi, Els Van Besien, Erik Rosseel, Geert Mannaert, An De Keersgieter, Adrian Chasin, Stefan Kubicek, Anish Dangol, Steven Demuynck, Kathy Barla, Dan Mocuta, and Naoto Horiguchi. (Invited) Gate-All-Around Transistors Based on Vertically Stacked Si Nanowires. *ECS Transactions*, 77(5):19, 2017.

[MS16] Amir Moradi and François-Xavier Standaert. Moments-Correlating DPA. In *TIS@CCS 2016*, pages 5–15. ACM, 2016.

[MSM+99] D. A. Muller, T. Sorsch, S. Moccio, F. H. Baumann, K. Evans-Lutterodt, and G. Timp. The electronic structure at the atomic scale of ultrathin gate oxides. *Nature*, 399(6738):758–761, 1999.

[NC10] Siva G. Narendra and Anantha P. Chandrakasan. *Leakage in Nanometer CMOS Technologies*. Springer Publishing Company, Incorporated, 1st edition, 2010.

[NRR06] Svetla Nikova, Christian Rechberger, and Vincent Rijmen. Threshold Implementations Against Side-Channel Attacks and Glitches. In *ICICS 2006*, volume 4307 of *Lecture Notes in Computer Science*, pages 529–545. Springer, 2006.

[NYH13] Nianhao Zhu, Yujie Zhou, and Hongming Liu. Counteracting leakage power analysis attack using random ring oscillators. In *Conference on Sensor Network Security Technology and Privacy Communication System*, pages 74–77, 2013.

[PKZM07] Thomas Popp, Mario Kirschbaum, Thomas Zefferer, and Stefan Mangard. Evaluation of the Masked Logic Style MDPL on a Prototype Chip. In *CHES 2007*, volume 4727 of *Lecture Notes in Computer Science*, pages 81–94. Springer, 2007.

[PM05] Thomas Popp and Stefan Mangard. Masked Dual-Rail Pre-charge Logic: DPA-Resistance Without Routing Constraints. In *CHES 2005*, volume 3659 of *Lecture Notes in Computer Science*, pages 172–186. Springer, 2005.

[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.

[PR16] C. Padmini and J. V. R. Ravindra. CALPAN: Countermeasure against Leakage Power Analysis attack by normalized DDPL. In *ICCPCT 2016*, pages 1–7, 2016.

[PSKM15] Santos Merino Del Pozo, François-Xavier Standaert, Dina Kamel, and Amir Moradi. Side-channel attacks from static power: when should we care? In *DATE 2015*, pages 145–150, 2015.

[RBN+15] Oscar Reparaz, Begül Bilgin, Svetla Nikova, Benedikt Gierlichs, and Ingrid Verbauwhede. Consolidating Masking Schemes. In *CRYPTO 2015*, volume 9215 of *Lecture Notes in Computer Science*, pages 764–783. Springer, 2015.

[RGV12] Oscar Reparaz, Benedikt Gierlichs, and Ingrid Verbauwhede. Selecting Time Samples for Multivariate DPA Attacks. In *CHES 2012*, volume 7428 of *Lecture Notes in Computer Science*, pages 155–174. Springer, 2012.

[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, 2003.

[RSV+11] Mathieu Renauld, François-Xavier Standaert, Nicolas Veyrat-Charvillon, Dina Kamel, and Denis Flandre. A Formal Study of Power Variability Issues and Side-Channel Attacks for Nanoscale Devices. In *EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 109–128. Springer, 2011.

[SBHM20] Pascal Sasdrich, Begül Bilgin, Michael Hutter, and Mark E. Marson. Low-Latency Hardware Masking with Application to AES. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(2):300–326, 2020.

[Sch99] Max Schulz. The end of the road for silicon? *Nature*, 399(6738):729–730, 1999.

[SM15] Tobias Schneider and Amir Moradi. Leakage Assessment Methodology - A Clear Roadmap for Side-Channel Evaluations. In *CHES 2015*, volume 9293 of *Lecture Notes in Computer Science*, pages 495–513. Springer, 2015.

[SM20] Aein Rezaei Shahmirzadi and Amir Moradi. Re-Consolidating First-Order Masking Schemes - Nullifying Fresh Randomness. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021:305–342, 2020.

[SMY09a] François-Xavier Standaert, Tal Malkin, and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009.

[SMY09b] François-Xavier Standaert, Tal Malkin, and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009.

[SP06] Kai Schramm and Christof Paar. Higher Order Masking of the AES. In *CT-RSA 2006*, volume 3860 of *Lecture Notes in Computer Science*, pages 208–225. Springer, 2006.

[SS06] Daisuke Suzuki and Minoru Saeki. Security Evaluation of DPA Countermeasures Using Dual-Rail Pre-charge Logic Style. In *CHES 2006*, volume 4249 of *Lecture Notes in Computer Science*, pages 255–269. Springer, 2006.

[TAV02] K. Tiri, M. Akmal, and I. Verbauwhede. A dynamic and differential CMOS logic with signal independent power consumption to withstand differential power analysis on smart cards. In *European Solid-State Circuits Conference 2002*, pages 403–406, 2002.

[TV03] Kris Tiri and Ingrid Verbauwhede. Securing Encryption Algorithms against DPA at the Logic Level: Next Generation Smart Card Technology. In *CHES 2003*, volume 2779 of *Lecture Notes in Computer Science*, pages 125–136. Springer, 2003.

[TV04] Kris Tiri and Ingrid Verbauwhede. A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation. In *DATE 2004*, pages 246–251. IEEE Computer Society, 2004.

[TV06] Kris Tiri and Ingrid Verbauwhede. A digital design flow for secure integrated circuits. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 25(7):1197–1208, 2006.

[YK17] Weize Yu and Selçuk Köse. Security-Adaptive Voltage Conversion as a Lightweight Countermeasure Against LPA Attacks. *IEEE Trans. on VLSI*, 25(7):2183–2187, 2017.

[YW18] Weize Yu and Yiming Wen. Leakage Power Analysis (LPA) Attack in Breakdown Mode and Countermeasure. In *SOCC 2018*, pages 102–105. IEEE, 2018.

[ZLGL22] Shihui Zhao, Bowen Li, Yuzheng Guo, and Huanglong Li. Ferroelectric-HfO2/oxide interfaces, oxygen distribution effects, and implications for device performance. *Applied Physics Letters*, 120(1), 2022.

[ZZL13] Nian-Hao Zhu, Yu-Jie Zhou, and Hong-Ming Liu. Employing Symmetric Dual-Rail Logic to Thwart LPA Attack. *IEEE Embedded Systems Letters*, 5(4):61–64, 2013.

[ZZL14] Nian-hao Zhu, Yu-jie Zhou, and Hong-ming Liu. A standard cell-based leakage power analysis attack countermeasure using symmetric dual-rail logic. *Journal of Shanghai Jiaotong University (Science)*, 19(2):169–172, 2014.