

Design-Time Exploration of Voltage Switching against Power Analysis Attacks in 14nm FinFET Technology

Invited Paper

Johann Knechtel^{a,*}, Tarek Ashraf^{b,1}, Natascha Fernengel^c, Satwik Patnaik^{d,2}, Mohammed Nabeel^a, Mohammed Ashraf^a, Ozgur Sinanoglu^a and Hussam Amrouch^{e,1}

^aDivision of Engineering, NYU Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

^bUniversity of Stuttgart, Stuttgart 70174, Germany

^cDepartment of Computer Science, Karlsruhe Institute of Technology, Karlsruhe 76131, Germany

^dDepartment of Electrical and Computer Engineering, Texas A&M University, College Station, TX, 77843, USA

^eChair of Semiconductor Test and Reliability (STAR) in the Computer Science, Electrical Engineering Faculty, University of Stuttgart, Stuttgart 70174, Germany

ARTICLE INFO

Keywords:

14nm FinFET
Library characterization
Power side-channel attack
Correlation power analysis
AES
Voltage switching
Design-time exploration
CAD flow

ABSTRACT

Given their non-invasiveness and demonstrated effectiveness, power analysis attacks (PAAs) are concerning and to be accounted for in modern circuit design. That is especially relevant for technology-dependent verification of PAA countermeasure implementations. Prior art proposed various countermeasures against PAAs, including masking and hiding, voltage switching, noise injection, etc. Aside from the proven working principles of such countermeasures, it is important to understand that their effectiveness is primarily technology- and implementation-dependent. Hence, before deployment, especially for integrated circuits, such countermeasures require accurate circuit-level studies.

This work investigates an industrial-grade 14nm fin field-effect transistor (FinFET) technology at design-time in the context of PAAs. We leverage device-level measurement data from *Intel* high-volume manufacturing processes, build up accordingly well-characterized standard-cell libraries, and utilize a commercial-grade computer-aided design (CAD) flow for PAA evaluation at design-time. Our study is focused on 1) the effectiveness of voltage switching as a countermeasure, 2) the advanced encryption standard (AES) cipher as a representative circuit, and 3) the correlation power analysis (CPA) as an attack framework. We show that, to improve the resilience against the CPA attack in particular and to lower information leakage in general, specific voltage configurations are more promising than others for the 14nm FinFET technology.

1. Introduction


To protect any sensitive data handled within integrated circuits (ICs), the use of ciphers is widely adopted, which are formally secure algorithms for encryption/decryption of data. Still, once attackers have direct access to ICs, they can monitor the physical interactions with the environment which are inevitably occurring for hardware. These interactions, also known as side-channel information leakage, can be exploited to, e.g., infer the secret key used for ciphers; related activities are known as side-channel attacks.


In this work, we focus on *power analysis attacks (PAAs)*, where attackers measure power traces and assess the underlying relationship to the cipher operations conducted within

the hardware under attack. Among other side-channel attacks, PAAs are particularly concerning, owing to their non-invasiveness, low-cost implementation, and proven effectiveness [3]. Different versions of PAAs have been demonstrated, like the simple power analysis, the differential power analysis (DPA) [13], the correlation power analysis (CPA) [3], or the mutual information analysis [8]. Without loss of generality, we focus on the widely adopted, seminal CPA attack in this work.

Various countermeasures against PAAs have been proposed and evaluated in a large body of prior art, including masking and hiding [14, 16, 25], voltage switching [10, 27, 26, 18], noise injection [1, 5], etc. Essentially, these countermeasures seek to de-correlate the observable power consumption from the sensitive cipher operations, to hinder PAAs. More specifically, masking and hiding re-organize the design such that sensitive operations are decomposed/split at the functional level as well as the circuit level, aiming for notions of formal security. However, the implementation of such schemes is expensive, as overheads are scaling quadratically with the related security requirements, making efficient implementations challenging [9]. Voltage switching can be supported via integrated voltage regulators (IVRs) [10]. IVRs are commonly employed in modern IC designs, as they can enable significant power savings.

*Corresponding author

 johann@nyu.edu (J. Knechtel); tarek.a.abdelsattar@gmail.com (T. Ashraf); n.fernengel@gmail.com (N. Fernengel); satwik.patnaik@tamu.edu (S. Patnaik); mtn2@nyu.edu (M. Nabeel); ma199@nyu.edu (M. Ashraf); ozgursin@nyu.edu (O. Sinanoglu); amrouch@iti.uni-stuttgart.de (H. Amrouch)

 <https://wp.nyu.edu/johann/> (J. Knechtel);
<https://nyuad.nyu.edu/dfxlab> (O. Sinanoglu);
<https://www.iti.uni-stuttgart.de/en/chairs/star> (H. Amrouch)
ORCID(s): 0000-0001-5093-2939 (J. Knechtel)

¹The work was done in part at Karlsruhe Institute of Technology, Germany.

²The work was done in part at NYU Abu Dhabi, UAE.

However, the efforts for IC verification as well as the need for system-level synchronization during switching periods represent some challenges when employing IVRs. Noise-injection schemes like doubling the registers and interposing random data into the paths [1], while effective in practice, also incur considerable area and power cost. Without loss of generality, we focus on voltage switching in this work.

Aside from the fact that the working principles of these various countermeasures have been thoroughly studied before, it is important to understand that their effectiveness is largely dependent on the used technology and the hardware implementation. Therefore, we argue that any countermeasure, especially before deployment in high-volume IC manufacturing, would require a circuit-level evaluation that must be based on accurate technology models along with commercial tool support. We note that some prior art, e.g., [16, 25], relied on field-programmable gate array (FPGA) implementations for evaluation, which is not suitable to assess the resilience of an IC implementation, as there are fundamental differences in the hardware architectures of FPGAs and ICs. We also note that some prior art, e.g., [27], uses SPICE simulations for power analysis, which is limited to small circuits, due to considerable computational efforts.

This manuscript is not aiming to counter such or other prior art, but rather to contribute to the landscape as follows.

- We derive a well-calibrated model for an industrial-grade 14nm FinFET technology from *Intel* quality production. This step is based on actual measurement data set and commercial tools, and the model matching is confirmed.
- We implement a CAD flow, also based on commercial tools, which allows for accurate and efficient design-time power analysis. This CAD flow is integrated with an open-source CPA attack, which we extend for a thorough exploration of the sampling space.
- Leveraging this integrated CAD and CPA framework, we investigate, without loss of generality, the security promises of voltage switching in detail for an AES circuitry implemented in an industrial-grade 14nm FinFET technology. We derive related, practical guidelines for any security-concerned designer.

The scope of our work is also illustrated in Figure 1. The manuscript is structured as follows. In Section 2, we describe our approach to model an industrial-grade 14nm FinFET technology from *Intel* quality production. In Section 3, we describe our commercial-tools-based CAD flow, which allows for design-time power analysis and thorough CPA attack runs for security evaluation. In Section 4, we present our experimental study, where we investigate the security promises of voltage switching in detail for an AES circuitry. We derive practical guidelines for any security-concerned designer using such industrial-grade 14nm FinFET technology. Finally, in Section 5, we conclude.

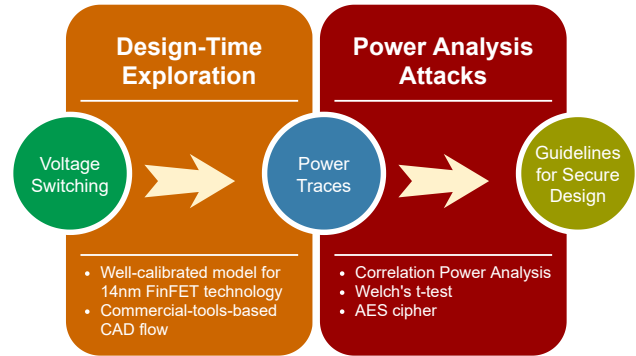


Figure 1: The scope of this work is to study voltage switching as a countermeasure against PAAs, specifically for an industrial-grade 14nm FinFET technology, and to derive conservative guidelines for secure design. This is achieved through design-time exploration of an exemplary AES circuitry, using a well-calibrated technology model within our integrated CAD flow. Since the data obtained from our flow is free of noise (unlike for attacks on devices in the field), we can provide conservative and well-founded security assessments.

2. Modeling and Calibrating for a Commercial 14nm FinFET Technology

To accurately model the FinFET technology, it is essential to perform calibrations with some existing commercial technology offered by the semiconductor industry. To achieve that, we employ the available measurement data from state-of-the-art *Intel* 14nm FinFET technology obtained from a high-volume production-quality manufacturing process [17].

First, we use *Synopsys Sentaurus Process Technology CAD* [23], a commercial TCAD tool, to mimic the fabrication process of 14nm FinFET using the same materials properties, layer dimensions, doping profile, etc. Then, we carefully tune key parameters in the FinFET device such as source/drain doping, sub-fin doping, source-drain series resistance, gate-metal work function, low-field mobility parameters, and high-field saturation parameters.

Next, we calibrate the industry-standard compact model for the FinFET technology, *BSIM-CMG* [15, 4], using the *Intel* measurement data and additional data obtained from the previously calibrated TCAD transistor devices. The additional data contains, for instance, the electrical properties of transistors at various temperature and voltage biases. The calibrations are performed for both n-type and p-type FinFET devices and the output of this step is fully-calibrated transistor models that accurately match and reproduce the measurement data of the 14nm *Intel* FinFET technology.

Finally, we perform TCAD mixed-mode simulations to further validate our calibrated transistor models (SBIM-CMG). In this validation phase, we compare the TCAD results of an inverter and ring oscillator circuits against the results of SPICE simulations. This step ensures to us that parasitics effects are accounted for in the calibration. Further

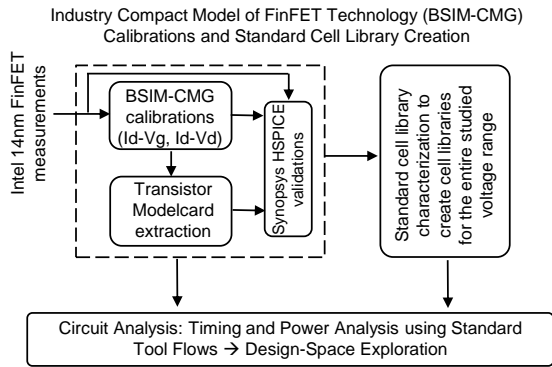


Figure 2: The calibration process starts with the industry compact model of FinFET technology (BSIM-CMG), using measurement data from the production-quality Intel 14nm node. The calibrated transistor models are then employed within standard tool flows for standard-cell library calibration. The calibration is performed at various operating voltages. Our libraries are fully compatible with the existing CAD tool flows.

details on the validation and calibration process are also available in [15].

In Figure 2, we summarize our technology calibration. In Figure 3, we demonstrate how SPICE simulations, using our calibrated FinFET model, reproduce *Intel* measurement data for both n-type and p-type FinFET devices. More specifically, in Figure 3(a), we show the transfer characteristics represented by transistor drain current (I_{DS}) versus gate voltage (V_{GS}), and in Figure 3(b), we show output characteristics represented by transistor drain current (I_{DS}) versus drain voltage (V_{DS}). In both cases, for both n-type and p-type transistors, the results from our calibrated FinFET models match very well the original 14nm FinFET measurement data from *Intel* [17].

Having well-calibrated transistor models enables us to characterize standard-cell libraries that can be used within any design flow. To this end, we employ the open-source SPICE netlists for FinFET standard cells from *Silvaco* [21]. Using a commercial tool flow for standard-cell library characterization based on *Synopsys SiliconSmart* [22], we create the 14nm FinFET library by employing our calibrated FinFET models. For every standard cell within the library, we consider 7×7 cases for input signal slews and output load capacitance, similar to what is done in existing commercial cell libraries.

To capture the impact of voltage scaling accurately, we repeat the library characterization for an entire voltage range from $0.8V$ down to $0.3V$ with $50mV$ and $100mV$ step sizes. All the created 14nm FinFET standard-cell libraries for the varying voltages are fully compatible with existing CAD tool flows. Hence, we can directly use them, without any modifications, to perform accurate timing analysis and power analysis. For this work, we employ these libraries within our CAD flow to investigate the resiliency of AES against PAAs attacks under the effects of voltage scaling.

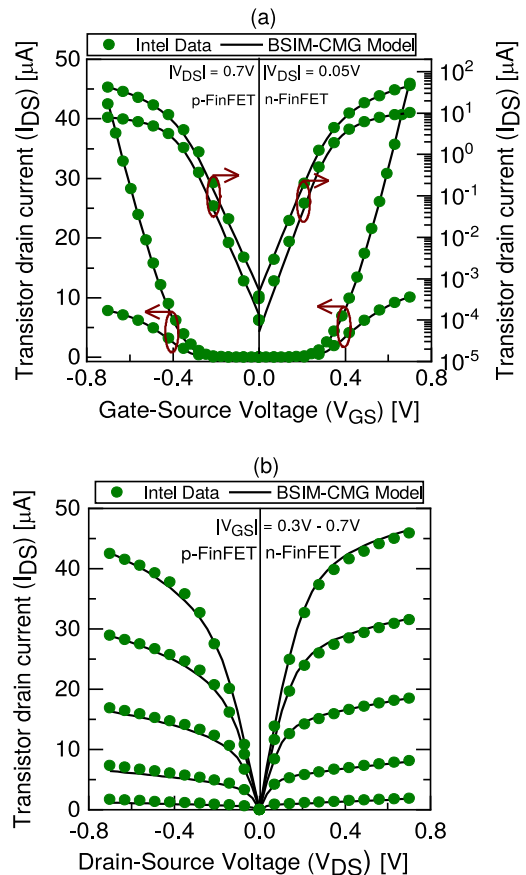


Figure 3: 14nm FinFET technology calibration for both n-type and p-type transistors. (a and b) demonstrate the excellent matching between our calibrated FinFET models and *Intel* 14nm FinFET measurement data from both transfer and output characteristics, respectively, [15].

3. Integrated CAD and CPA Framework

Next, we explain our integrated CAD and CPA framework that allows to evaluate ciphers (or any other circuitry module, for that matter) against PAAs (Figure 4). The fundamentals of the framework are derived from [12].

We note that some prior art leverages similar principles, e.g., see [2, 20]. However, an important distinction of ours from such prior art is the randomized-but-reproducible exploration of the power-distribution sampling space (not illustrated in Figure 4, but explained in detail further below). This is essential for a fair and thorough assessment of any technology, circuitry, or system feature against PAAs. We demonstrate our framework for the seminal AES cipher and voltage switching in this work; however, our framework is not limited to those.

It is important to note that our notion of resilience is *not* based on any de-synchronization effects induced by voltage switching. While an attacker would have to handle these effects during measurements required for real-world attacks, here we instead evaluate the foundation of voltage switching, namely mixing of different power distributions, and the resulting role for resilience against PAAs.

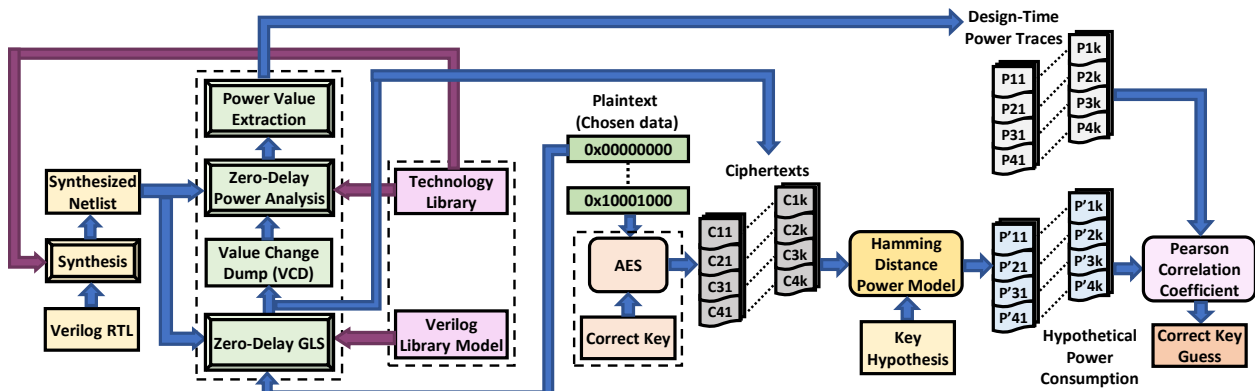


Figure 4: Our integrated CAD and CPA framework. The left part illustrates the CAD stages, the center part illustrates the cipher/module under investigation (AES for this study), and the right part illustrates the CPA stages. The parts are intertwined as needed for design-time exploration against PAAs.

Our framework requires the register-transfer level (RTL) description of the cipher and the standard-cell libraries of choice as inputs. After the successful termination of the framework, the user gets to know the minimum number of power traces needed to disclose the secret key. In general, a larger number of traces indicates that the underlying cipher RTL has a higher resilience against the launched attack. We explain the two major components of our integrated framework next.

3.1. Simulation-Based Power Analysis

Initially, we synthesize the cipher RTL using the technology libraries of choice. We verify the functionality of the gate-level netlist using a Verilog testbench with user-defined sets of plaintexts and keys. We also confirm the functionality of the design using software simulation.

Next, we perform a gate-level simulation (GLS) and generate a Value Change Dump (VCD) file. This VCD file captures every node's switching activity for a user-defined time resolution (e.g., 1 ps). The VCD file is then used for power simulation of the synthesized gate-level netlist. To reduce computational efforts for power analysis without comprising the efficacy of PAAs, we can consider only the relevant time intervals, i.e., the last round of AES, which is known to be vulnerable [3].

Instead of performing full-scale timing simulations, which would also capture glitching activities, we leverage zero-delay simulations here. For such, all signal transitions occur simultaneously at the clock's active edge, which also simplifies the extraction of peak power values. As explained further below, we are indeed only interested in the switching power of specific registers, which occurs at the clock's active edge. Our power analysis can be considered conservative from a security standpoint since we ignore glitching noises, which naturally deteriorate any PAA. Depending on the countermeasure of interest, glitches may still have to be considered, e.g., for masking-based countermeasures. However, for our notion of voltage switching, glitches are not relevant. Related implications are further discussed in Section 4.

In short, we obtain the design-time power traces for the AES RTL in a step-wise manner while processing the texts and secret keys, and we extract the zero-delay, peak-power values for the sensitive registers of the last-round AES operations [3].

3.2. Correlation Power Analysis

The correlation power analysis (CPA) [3] is an effective attack. At its heart, the *Pearson correlation coefficient* (PCC) is leveraged to measure the relationship between predicted and actual power profiles [3]. In our framework, we follow a standard approach for the CPA, as detailed next.

First, we derive the predicted power profiles using a power model. Since registers consume a significant portion of dynamic power during signal transitions, considering the *Hamming distance* (*HD*) for the registers' output before and after switching is established as the *HD power model* [3]. This predictive modeling is repeated for all possible key candidates; the results are also known as *hypothetical power values* that are related to the respective key hypotheses. Note that we need to consider only the registers holding the intermediate texts for the vulnerable last-round operation to build up the HD power model [3]. Finally, the design-time the hypothetical power values are correlated via PCC, and the hypothesis with the highest PCC is thought to represent the correct key.

In other words, we stepwise correlate the following two assets: (a) the set of design-time power values arising from using the actual, secret key for the vulnerable AES operations, and (b) all the different sets comprising all possible hypothetical power values, which mimic the power consumption of the same AES operations for all possible key candidates. Note that all the sets in (b) have to be individually correlated with (a); there are as many correlation results as there are sets or key hypotheses in (b). Then, the set in (b) corresponding to the highest correlation score describes the hypothetical power values matching best with the actual power consumption of the AES operations using the secret key—the related key hypothesis is thus considered to be correct key.

Acting as designers, we can directly verify the key hypothesis. As an attacker, however, one would have to study the trend in correlation throughout multiple trials, and once a stable and significant trend for correlation appears, one deducts with certain confidence that the key hypothesis should be correct [3].

While conducting the above steps would be sufficient for an actual attack, for the exploration in this study, we have to conduct multiple trials while varying the secret key, the texts, the selection of power traces, and the voltages. Related details are given next, in Section 4.1.

4. Experimental Investigation

4.1. Setup

4.1.1. Synthesis and Simulations

We use *Synopsys VCS* for functional simulations at RTL and gate level, *Synopsys DC* for logic synthesis, and *Synopsys PrimeTime PX* for power simulations; all tool versions are from 2018.

For the AES circuit, we leverage a publicly available RTL, which works on 128-bit keys and texts, and uses look-up tables for the AES substitution box [11]. For synthesis and power simulations, we employ the 14nm FinFET technology libraries described in Section 2. Without loss of generality, we use the operating frequency $f = 100MHz$.

4.1.2. Voltage Switching

To enable a thorough study on the impact of voltage switching, we carefully investigate switching using a broad range of voltage configurations. Our approach is outlined in Figure 5, and we describe it in some detail next.

First, we apply different supply voltages for the whole AES circuit during independent runs for power simulations, after which we collect the corresponding power traces for each voltage and corresponding library separately (as outlined in Section 3–3.1). The considered voltages are: 0.8V, 0.75V, 0.7V, 0.6V, 0.5V, 0.4V, 0.3V.

Now, we can study voltage switching for varying-sized sets of texts processed by the AES circuit. We stepwise consider the full spectrum of switching, i.e., we investigate all possible configurations ranging from 100% of the texts processed using V_{DD1} , over, e.g., 99% processed using V_{DD1} and 1% processed using V_{DD2} , all the way to 100% processed using V_{DD2} . Note that the granularity for switching is tuned according to our initial experiments—for voltage configurations with a large difference $|V_{DD1} - V_{DD2}|$, we found that employing very fine granularities (down to 0.1%) is required, whereas for smaller differences, coarser granularities suffice (e.g., 10%). Also note that, while observing the selected granularity, we do switch randomly, i.e., any number of consecutive texts may or may not be processed using the same voltage.

For simplicity, we always switch between 0.8V and one other, fixed voltage, i.e., we consider pairings $\{V_{DD1}, V_{DD2}\}$ with $V_{DD1} = 0.8V$ as common baseline voltage, where $V_{DD2} \neq V_{DD1}$. Since we keep V_{DD2} fixed to one particular

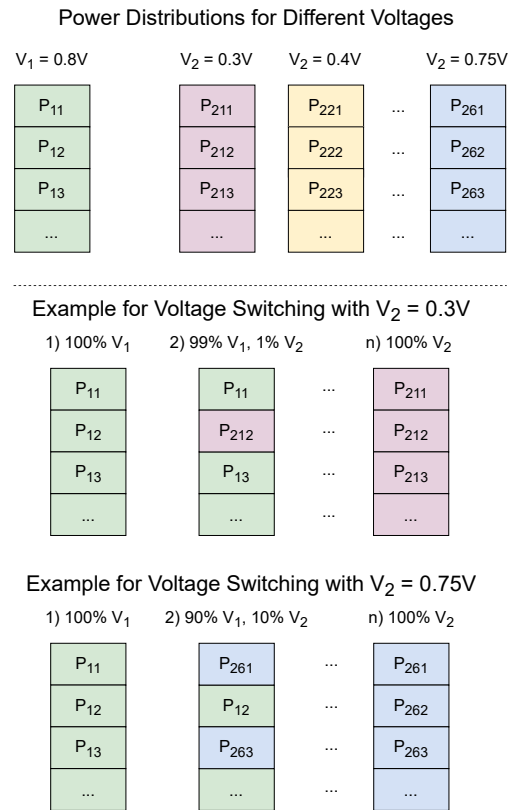


Figure 5: Conceptual examples for voltage switching. Voltages V_{DD1} , V_{DD2} are labeled V_1 , V_2 here for simplicity. Based on the individual power distributions obtained through simulation for different voltages, voltage switching at the system level is mimicked by mixing these power distributions. Note that voltage switching is realized stepwise and independently for the different, fixed pairs of V_1 , V_2 .

voltage for each set of experiments, we have to conduct independent runs for each set, across all the considered voltage configurations.

Note that, for each voltage configuration, we conduct three batches of randomized switching, and all reported results are accordingly averaged.

To mimic voltage switching at the system level, we mix the separately obtained power distributions, whereas some IVR circuitry would be leveraged in practice.

As indicated, our evaluation of resilience is *not* based on any de-synchronization effects induced by voltage switching. While an attacker would have to handle these effects during measurements required for real-world attacks, here we rather evaluate the foundation of voltage switching, namely different power distributions, and their impact on resilience.

4.1.3. CPA Attack

We leverage and extend the open-source C/C++ framework of [7]; we provide our version in [11]. All CPA runs are executed on a high-performance computing (HPC) facility, with 14-core *Intel* Broadwell processors (Xeon E5-2680) running at 2.4 GHz, and 4 GB RAM are guaranteed (by the Slurm HPC scheduler) for each CPA run.

We assume a classical threat model, i.e., the attacker can measure power traces at will and also understands which power values belong to which texts that are being processed. While the attacker understands that voltage switching is applied as a countermeasure, s/he does not know which particular text is processed using which voltage.¹ Furthermore, the attacker has only external physical access, but no invasive probing capabilities.

While our power traces are obtained at design-time and are thus void of any environmental noise (e.g., arising from the measurement equipment), the traces gathered by an attacker would be subjected to such noise. In other words, we are *not* taking the attacker’s perspective here, but rather the designer’s perspective, where we leverage the CPA attack to conduct a conservative and robust security evaluation at design-time.

Initially, we generate and store the following assets: six random, 128-bit keys; 5,000 random, 128-bit plain-texts; and, separately for each key, the corresponding 5,000 ciphertexts, also 128 bits long each. We keep these sets of keys and texts the same across the whole study, i.e., for all voltage configurations.

To guide the CPA attack exhaustively through the various sets of voltage-variable power traces, we apply the following sampling strategy.² We start with a small number of traces being selected from the larger overall set of traces collected, and we choose such small subsets of traces multiple times, in order to conduct multiple, independent CPA runs. Then, we stepwise increase the number of traces being selected, and we repeat that procedure of selecting multiple subsets and conducting related CPA runs. By doing so, we explore different subsets with more and more traces being made available for the CPA attack. Such an approach is essential for a robust assessment of how many traces are needed to disclose a secret key.

The sampling approach is exemplified in Figure 6. Throughout 2,500 steps, we randomly pick for each step 1,000 sets of matching pairs of design-time power values and hypothetical power values. Matching pairs mean that the power values arise from the same texts being processed, which is a prerequisite for PCC, or for any correlation for that matter.

Each of 2,500 steps requires an independent, separate CPA run. As indicated, each step covers 1,000 sets (through an independent CPA run), more specifically 1,000 permutations of all 5,000 available pairs of design-time and hypothetical power values, where each step $s + 1$ considers two more pairs per permutation over prior step s . In total, this results in 2,500,000 permutations being considered, step by step, through individual CPA runs on those subsets of

¹To uphold this assumption, it is essential that the voltage configurations cannot be trivially differentiated, i.e., the power distributions arising from using the different voltages must be sufficiently interspersed. This holds true for all the configurations considered in this study, as long as the ratio of switching is accordingly tuned. See Section 4.2.2 for more details.

²This strategy is conducted separately for each configuration of voltage switching—voltage switching and sampling of power traces are independent procedures of our framework.

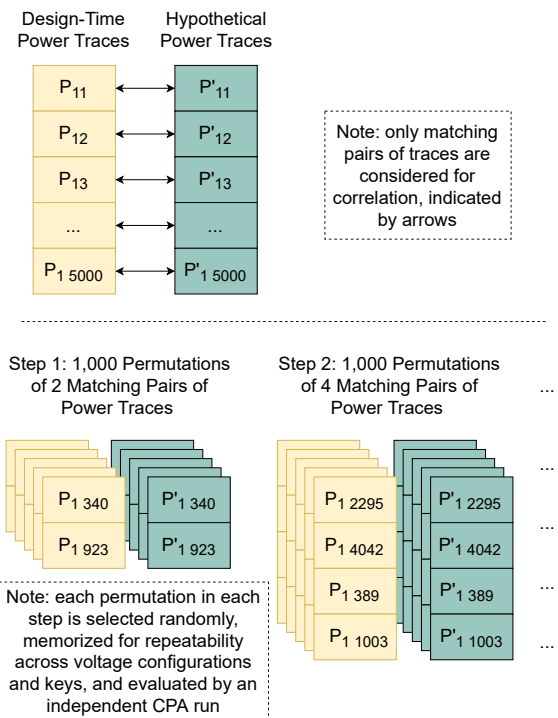


Figure 6: All experiments utilize the outlined, thorough sampling of power traces, to determine the number of traces to disclosure with high confidence.

power values. For a fair comparison, these permutations are all memorized, stored, and re-applied when conducting the CPA runs for different keys as well as for different voltage configurations.

In short, using the above strategy, we ensure 1) a thorough sampling of all power traces for robust inferences about how many traces are needed for key disclosures, and, more importantly, 2) that any differences observed in CPA resilience are *not* due to the outlined randomized nature of sampling, but rather indeed due to different voltage configurations (and different keys).

4.2. Results

4.2.1. Overview

The goal of this work is to investigate the resilience offered by voltage switching in an industrial-grade 14nm FinFET technology. From this study, security-concerned designers can obtain guidance for proper implementation of voltage switching as a countermeasure against PAAs in general and the seminal CPA attack in particular.

The main observations from the empirical study are as follows. First, the larger the difference $|V_{DD_1} - V_{DD_2}|$, the less we need to switch from the baseline voltage V_{DD_1} to V_{DD_2} in order to raise the resilience against the CPA attack to a specific, comparable level. Second, the larger the difference, however, the larger the information leakage in terms of Welch’s t-test scores.

The first observation is because of the fact that, for a larger difference of voltages, the difference in power consumption is significantly larger as well, given that $P \propto$

Table 1
Results for Baseline Voltage $V_{DD_1} = 0.8V$

Key	#TTD(99.9%)	$\mu(P)$	$\sigma(P)$	$CV(P)$
1	1008	1.868E-03	5.15E-05	2.77%
2	920	1.870E-03	5.40E-05	2.88%
3	874	1.873E-03	5.26E-05	2.80%
4	856	1.824E-03	5.43E-05	2.98%
5	970	1.881E-03	5.21E-05	2.76%
6	1032	1.862E-03	5.31E-05	2.85%
$\mu(\#TTD)$	943	-	-	-
$\sigma(\#TTD)$	71.74	-	-	-
$CV(\#TTD)$	7.61%	-	-	-

$(\alpha \times C \times V_{DD}^2 \times f)$.³ The resulting pronounced variances for power values hinder the effectiveness of the PCC formalism underlying of the CPA attack, suggesting to leverage largely different voltages to advance security.

The second, more important observation, however, suggests that the voltage difference should be limited. This is because for overly large differences, there is information leakage arising for the largely varying power values. While the seminal CPA attack is not leveraging this kind of information leakage—on the contrary, the CPA attack is rather hindered by the less clear correlations of voltage-variable power values with AES operations—more advanced attacks, e.g., based on machine learning, might leverage this leakage. Note that the second observation also relates to the assumption that an attacker cannot trivially differentiate the power distributions. This assumption only holds true for limited voltage differences where the resulting power distributions remain sufficiently interspersed.

In short, there are trade-offs for voltage switching which necessitate such an empirical study during design-time, in particular before application of voltage switching in production. Next, we discuss our findings in more detail.

4.2.2. Detailed Results and Discussion

In Table 1, we report the results considering only the baseline voltage $V_{DD_1} = 0.8V$. The main metric here, as well as in the remaining experiments, is *number of traces to disclosure (#TTD)*, which we report with a confidence of 99.9%. Recall that we leverage 1,000 permutations per step during our sampling approach; these 1,000 permutations, evaluated through individual CPA runs each, correspond to the 1,000 runs required for the 99.9% confidence on assessing #TTD.

We further report the (rounded) mean and standard deviation for #TTD, which describe the variations induced by different keys. However, note that we focus on the mean #TTD in the remainder; key-induced variations are not of

³The switching activities α are dictated by the AES keys and texts processed, as captured through *Synopsis VCS* in our flow. Given that the same keys and texts are processed, the activities are also the same.

particular interest here.⁴ In any case, we point out the following: 1) these differences observed across keys establish themselves as relatively consistent trends across all experiments; 2) our observations are interpreted considering mean values. Thus, our findings on voltage switching are not undermined in any way by key-induced variations.

We also report the mean $\mu(P)$, the standard deviation $\sigma(P)$, and the coefficient of variation $CV(P)$ for the underlying power values; these metrics serve to characterize the distribution of the power values. The coefficient of variation $CV = \sigma/\mu$ puts the variability of a distribution into its proper context, i.e., the mean value. The CV values in Table 1 are put into context further below.

In Table 2, we report the results when voltage switching is applied. As indicated, we consider a large range of voltages V_{DD_2} to switch to.

The first finding is that all considered voltage configurations can help to increase the resilience against the CPA attack by $\approx 5\times$ over the baseline resilience (Table 1). The mean ratio of how much we have to switch is essentially dictated by the targeted-at mean #TTD $\approx 5,000$, which we utilize for all configurations to enable a fair comparison. That is, different configurations require different degrees of switching to achieve a specific, comparable resilience against the CPA attack; the smaller the difference between the two voltages, the more switching is required.

The second finding is that for any voltage configuration, switching more would generally increase the resilience, albeit also introducing more information leakage. We have also studied this aspect for $V_{DD_2} = 0.75V$ in more detail, with the related results illustrated in Fig 7.

Toward the third finding, we note the following two aspects: 1) the mean PCC between the randomly (but reproducibly) selected power values from the individual distributions associated with the two different voltages is relatively low; 2) the mean CV for the combined distribution of the selected power values (which is the only one observable by an actual attacker) is relatively large, i.e., compared to the CV values for the baseline configuration (Table 1), but still sufficiently low to maintain the interspersed nature of the underlying power distributions.

These aspects serve to explain the increased resilience against the CPA attack—correlating hypothetical power values, which exhibit relatively low correlation to begin with, with the more varied power distributions resulting from voltage switching is more difficult by definition.

⁴For more context on key-induced variations, note the following. Recall that our experiments are tailored for a fair and thorough comparison and, as such, for these experiments considering different keys, we ensure that the very same (but initially randomized) texts are leveraged for all experiments. Thus, the resulting differences in #TTD are due to varying distributions of bit-level flips across the last-round texts, which can impact the effectiveness of the CPA attack by definition. Still, we caution that such findings should *not* motivate to favor specific, seemingly “resilient keys” for actual applications—such outcomes are highly dependent on the architecture, RTL, and gate-level implementation of the cipher, as well as the selection of plain-/cipher-texts. There are studies investigating the role of keys and texts more formally, e.g., based on notions of collision and statistical modeling [6], which are out of scope for this work.

Table 2Results for Voltage Switching With $V_{DD_1} = 0.8V$

V_{DD_2}	Mean Ratio V_{DD_2}/V_{DD_1}	#TTD			Combined Power Distribution			Individual Power Distributions	
		μ	σ	CV	Mean μ	Mean σ	Mean CV	Mean PCC	Mean t-score
0.3V	1.05%	4984	38.22	0.77%	1.813E-03	1.75E-04	9.46%	0.2993	659.54
0.4V	1.42%	4953	73.43	1.48%	1.843E-03	1.79E-04	9.71%	0.2937	562.62
0.5V	2.03%	4920	135.02	2.74%	1.839E-03	1.77E-04	9.63%	0.3060	462.23
0.6V	3.52%	4927	111.58	2.26%	1.833E-03	1.70E-04	9.27%	0.3138	345.79
0.7V	11.50%	4939	49.73	1.01%	1.808E-03	1.63E-04	9.01%	0.3213	22.95
0.75V	50.00%	4730	116.31	3.52%	1.736E-03	1.37E-04	7.86%	0.3600	61.53

The “Combined Power Distribution” refers to the actual distribution after voltage switching is applied, whereas the “Individual Power Distributions” refer to the individual distributions obtained for the respective voltages.




Figure 7: Mean #TTD and mean t-scores versus the ratio for voltage switching, with $V_{DD_1} = 0.8V$, $V_{DD_2} = 0.75V$. The more we switch, the more resilient the circuit becomes against the CPA attack, but also the more information leakage is arising from the power values corresponding to different voltages.

The role of these aspects on the resilience against the CPA attack can also be quantified, namely via the Spearman correlation, which is defined as the PCC between the rank variables. Unlike the PCC itself, that can only assess linear relationships (which is sufficient for correlating actual power values and hypothetical power values within the CPA framework), the Spearman correlation is more generic as it can assess any monotonic relationship. For example, the Spearman correlation between #TTD and PCC is -0.8162 , representing a strong reciprocal correlation; this confirms that more traces are needed for less correlated power distributions. Furthermore, the Spearman correlation between #TTD and CV is 0.7532 , representing a strong correlation as well; this confirms that more traces are needed for more varied power distributions.

As explained, the resilience against the CPA attack is increasing the more we switch and the larger the voltage difference is. For such cases, however, information leakage is often increasing as well, as evident from Welch’s t-test scores in both Table 2 and Figure 7. Note that Welch’s t-test can assess whether there is information leakage at all, which may be exploited by some attacks, but it does not provide any insight on success rates of actual attacks.⁵

⁵It is said that, for a t-score below 4.5, the circuit is leakage-free and thus cannot be attacked by PAAs [19]. Accordingly, our reported results are *not* for leakage-free scenarios, which is also clearly demonstrated by the fact that the related CPA attack runs are, on average, successful.

Thus, from a formal point of view, possibly seeking resilience going beyond the CPA attack, we would rather want to employ small voltage differences and small switching ratios. Indeed, an outstanding t-score minima is observed for $V_{DD_2} = 0.7V$, where the voltage difference is still sufficiently small while the switching ratio is already considerably relaxed/reduced compared to the case of $V_{DD_2} = 0.75V$ (where $4.35\times$ more switching was required for comparable CPA resilience). Importantly, although such a setting impacts the resilience against the CPA attack in general, it does *not* undermine the resilience against a specific level of attack effort (i.e., #TTD). This is especially true once a slightly increased switching ratio is considered again.

Thus, we have investigated this most promising scenario further, by considering more frequent switching. Indeed, already from 12% switching onward, the CPA attack always failed to infer the correct key even when provided all 5,000 traces (across all the different keys).

This constitutes another, fourth finding, namely that this specific configuration is most promising for voltage switching as a countermeasure within the considered 14nm FinFET technology, both from an empirical and a formal perspective.

Finally, it is important to recall that this study is based on zero-delay power simulations using commercial tools. Therefore, our results are void of any noise. Moreover, aside from empirical data based on CPA attack runs, we provide formal insights based on Welch’s t-test. This implies two considerations as follows. First, for actual attacks in the field—assuming the same operation conditions—an attacker cannot do any better than what we observe. Accordingly, our findings serve well as accurate but conservative guidelines, e.g., for tailoring schemes like dynamic key management [24]. Second, for actual voltage switching in the field, our findings and guidelines may be over-constrained, i.e., once noise or other detrimental effects play out, a less stringent switching than what we consider here may offer the same level of resilience in practice. Most, if not all, IVR circuitry can be configured at run-time as needed, allowing to later on reduce such overheads arising for over-constrained switching requirements.

5. Conclusion

In this work, we have investigated the resilience of an industrial 14nm FinFET technology against power analysis attacks. We have leveraged device-level measurement data from manufacturing cycles, built up correspondingly well-characterized standard-cell libraries for different supply voltages, and utilized a commercial-grade CAD flow for power and attack evaluation at design-time.

Without loss of generality, our study is focused on voltage switching for an AES circuit and its resilience against the seminal CPA attack. We have shown that there are trade-offs for voltage switching, which clearly necessitate such an empirical study. More specifically, the resilience against the CPA attack is increasing the more we switch as well as the larger the difference is for the voltages being switched between, but the information leakage is often increasing at the same time. The specific configuration $V_{DD_1} = 0.8V$, $V_{DD_2} = 0.7V$ was found to be a “sweet spot” that is most promising from both formal and empirical perspective. Randomly switching more than 12% of the AES operations to V_{DD_2} was hindering CPA attacks altogether.

While specific to the AES circuit implemented using the 14nm FinFET technology, these findings are robust and conservative. The former is because we carefully conduct batches of randomized-but-reproducible CPA runs, to explore the whole sample space of power traces in a comprehensive manner. The latter is because the findings are based on technology-accurate, design-time data obtained from zero-delay power simulations. Furthermore, we consider the AES circuit as stand-alone module, without any potential side-effects of other modules’ activities undermining our analysis. Assuming the same operation conditions, an attacker in the field cannot do any better than what we observe here. Accordingly, our findings serve well as practical guidelines for secure design or for system-level design/configuration parameters, e.g., key refresh rates.

References

- [1] Bellizia, D., Bongiovanni, S., Monsurrò, P., Scotti, G., Trifiletti, A., Trotta, F.B., 2018. Secure double rate registers as an RTL countermeasure against power analysis attacks. *Trans. VLSI Syst.* 26, 1368–1376. doi:10.1109/TVLSI.2018.2816914.
- [2] Bhasin, S., Danger, J.L., Graba, T., Mathieu, Y., Fujimoto, D., Nagata, M., 2013. Physical security evaluation at an early design-phase: A side-channel aware simulation methodology, in: *Proc. ES4CPS*, pp. 13–20. doi:10.1145/2589650.2559628.
- [3] Brier, E., Clavier, C., Olivier, F., 2004. Correlation power analysis with a leakage model. *Proc. Cryptogr. Hardw. Embed. Sys.*, 16–29. doi:10.1007/978-3-540-28632-5_2.
- [4] BSIM Group, 2020. BSIM-CMG Model. URL: <http://bsim.berkeley.edu/models/bsimcmg>.
- [5] Das, D., Maity, S., Nasir, S.B., Ghosh, S., Raychowdhury, A., Sen, S., 2018. ASNI: Attenuated signature noise injection for low-overhead power side-channel attack immunity. *Trans. Circ. Sys.* 65, 3300–3311.
- [6] Fei, Y., Ding, A., Lao, J., Zhang, L., 2015. A statistics-based success rate model for DPA and CPA. *J. Cryptogr. Eng.* 5, 227–243.
- [7] Fei, Y., et al., 2013. Side channel analysis library. URL: https://tescase.coe.neu.edu/?current_page=SOURCE_CODEZsoftware=aestool.
- [8] Gierlichs, B., Batina, L., Tuyls, P., Preneel, B., 2008. Mutual information analysis, in: *Proc. Cryptogr. Hardw. Embed. Sys.*, pp. 426–442.
- [9] Gross, H., Mangard, S., Korak, T., 2017. An efficient side-channel protected AES implementation with arbitrary protection order, in: Handschuh, H. (Ed.), *Topics in Cryptology – CT-RSA 2017*, pp. 95–112.
- [10] Kar, M., Singh, A., Mathew, S.K., Rajan, A., De, V., Mukhopadhyay, S., 2018. Reducing power side-channel information leakage of AES engines using fully integrated inductive voltage regulator. *J. Sol.-St. Circ.* 53, 2399–2414.
- [11] Knechtel, J., 2019–2020. Correlation power attack. URL: <https://github.com/DfX-NYUAD/CPA>. extended from [7] as retrieved in March 2019.
- [12] Knechtel, J., Patnaik, S., Nabeel, M., Ashraf, M., Chauhan, Y.S., Henkel, J., Sinanoglu, O., Amrouch, H., 2020. Power side-channel attacks in negative capacitance transistor. *IEEE Micro* 40, 74–84. doi:10.1109/MM.2020.3005883.
- [13] Kocher, P., Jaffe, J., Jun, B., 1999. Differential power analysis, in: *Proc. Adv. Cryptol.*, pp. 388–397.
- [14] Li, X., Yang, C., Ma, J., Liu, Y., Yin, S., 2017. Energy-efficient side-channel attack countermeasure with awareness and hybrid configuration based on it. *Trans. VLSI Syst.* 25, 3355–3368. doi:10.1109/TVLSI.2017.2752212.
- [15] Mishra, S., Amrouch, H., Joe, J., Dabhi, C.K., Thakor, K., Chauhan, Y.S., Henkel, J., Mahapatra, S., 2019. A simulation study of NBTI impact on 14-nm node FinFET technology for logic applications: Device degradation to circuit-level interaction. *Trans. Electron Dev.* 66, 271–278.
- [16] Moos, T., Moradi, A., Schneider, T., Standaert, F.X., 2019. Glitch-resistant masking revisited. *Trans. Cryptogr. Hardw. Embed. Sys.* 2019, 256–292. doi:10.13154/tches.v2019.i2.256-292.
- [17] Natarajan, S., Agostinelli, M., Akbar, S., Bost, M., Bowonder, A., Chikarmane, V., Chouksey, S., Dasgupta, A., Fischer, K., Fu, Q., Ghani, T., Giles, M., Govindaraju, S., Grover, R., Han, W., Hanken, D., Haralson, E., Haran, M., Heckscher, M., Heussner, R., Jain, P., James, R., Jhaveri, R., Jin, I., Kam, H., Karl, E., Kenyon, C., Liu, M., Luo, Y., Mehandru, R., Morarka, S., Neiberg, L., Packan, P., Paliwal, A., Parker, C., Patel, P., Patel, R., Pelto, C., Pipes, L., Plekhanov, P., Prince, M., Rajamani, S., Sandford, J., Sell, B., Sivakumar, S., Smith, P., Song, B., Tone, K., Troeger, T., Wiedemer, J., Yang, M., Zhang, K., 2014. A 14nm logic technology featuring 2nd-generation FinFET, air-gapped interconnects, self-aligned double patterning and a 0.0588 μm^2 SRAM cell size, in: *Proc. Int. Elec. Devices Meeting*, pp. 3.7.1–3.7.3.
- [18] Pothukuchi, R.P., Pothukuchi, S.Y., Voulgaris, P.G., Torrellas, J., 2019. Maya: Falsifying power sidechannels with operating system support. *CoRR abs/1907.09440*. arXiv:1907.09440.
- [19] Schneider, T., Moradi, A., 2015. Leakage assessment methodology - a clear roadmap for side-channel evaluations. URL: <https://eprint.iacr.org/2015/207>.
- [20] Šijačić, D., Balasch, J., Yang, B., Ghosh, S., Verbauwhe, I., 2020. Towards efficient and automated side-channel evaluations at design time. *J. Crypt. Eng.* 10, 305–319. doi:10.1007/s13389-020-00233-8.
- [21] Silvaco and Si2, 2019. Open-Source FinFET Standard Cell Library. URL: <https://silvaco.com/news/silvaco-and-si2-release-unique-free-15nm-open-source-digital-cell-library/>.
- [22] Synopsys, 2020a. Standard Cell Library Characterization. URL: <https://www.synopsys.com/implementation-and-signoff/signoff/siliconsmart.html>.
- [23] Synopsys, 2020b. Synopsys Technology Computer Aided Design (TCAD). URL: <https://www.synopsys.com/silicon/tcad.html>.
- [24] Taha, M., Schaumont, P., 2014. Key updating for leakage resiliency with application to AES modes of operation. *Trans. Inf. Forens. Sec.* 10, 519–528.
- [25] Takemoto, S., Nozaki, Y., Yoshikawa, M., 2019. Evaluation of the hiding-countermeasure prince using differential power analysis, in: *Proc. Conf. Cons. Electr.*, pp. 164–165.
- [26] Yang, S., Gupta, P., Wolf, M., Serpanos, D., Narayanan, V., Xie, Y., 2012. Power analysis attack resistance engineering by dynamic voltage and frequency scaling. *Trans. Embed. Comput. Syst.* 11.

doi:10.1145/2345770.2345774.

- [27] Zhang, F., Yang, B., Yang, B., Zhang, Y., Bhasin, S., Ren, K., 2019. Fluctuating power logic: SCA protection by V_{DD} randomization at the cell-level, in: Proc. Asian Hardw.-Orient. Sec. Trust Symp., pp. 1–6.

Acknowledgments

This work was supported in part by the Center for Cyber Security (CCS) at New York University Abu Dhabi (NYUAD) and in part by the German Research Foundation (DFG), project “ACCROSS: Approximate Computing aCROSS the System Stack” and in part by the Office of Naval Research under Grant N00014-18-1-2672. Parts of this work were carried out on the High Performance Computing facility at NYUAD. The work of S. Patnaik at NYUAD was supported by the Global Ph.D. Fellowship of NYU New York and NYUAD.