

Understanding the Limitations and Improving the Relevance of SPICE Simulations in Side-Channel Security Evaluations

Dina Kamel, Mathieu Renaud, Denis Flandre, François-Xavier Standaert.
ICTEAM/ELEN/Crypto Group, Université catholique de Louvain, Belgium.

Abstract. Simulation is a very powerful tool for hardware designers. It generally allows the preliminary evaluation of a chip's performances before its final tape out. As security against side-channel attacks is an increasingly important issue for cryptographic devices, simulation also becomes a desirable option for preliminary evaluation in this case. However, its relevance highly depends on the proper modeling of all the attack peculiarities. For example, several works in the literature directly exploit SPICE-like simulations without considering measurement peripherals. But the outcome of such analyses may be questionable, as witnessed by the recent results of Renaud et al. at CHES 2011, which showed how far the power traces of an AES S-box implemented using a dynamic and differential logic style fabricated in 65 nm CMOS can lie from their post-layout simulations. One important difference was found in the linear dependencies between the (simulated and actual) traces and the S-box input/output bits. While simulations exhibited highly non-linear traces, actual measurements were much more linear. As linearity is a crucial parameter for the application of non-profiled side-channel attacks (which are only possible under the assumption of "sufficiently linear leakages"), this observation motivated us to study the reasons of such differences. Consequently, this work discusses the relevance of simulation in security evaluations, and highlights its dependency on the proper modeling of measurement setups. For this purpose, we present a generic approach to build an adequate model to represent measurements artifacts, based upon real data from equipment providers for our AES S-box case study. Next, we illustrate the transformation of simulated leakages, from highly non-linear to reasonably linear, exploiting our model and regression-based side-channel analysis. While improving the relevance of simulations in security evaluations, our results also raise doubts regarding the possibility to design dual-rail implementations with highly non-linear leakages.

1 Introduction

Side-channel attacks are considered an escalating threat to the physical security of cryptographic devices. These attacks are worrisome for actual applications since they are relatively cheap, easy to conduct and can be extremely powerful (e.g. leading to key recoveries). Generally, side-channel attacks are classified as profiled or non-profiled [9]. Profiled attacks, such as using templates [2], are based upon the assumption that the adversary has prior access to the target device and is capable of accurately profiling it. By contrast, non-profiled attacks

(e.g. Correlation Power Analysis (CPA) [1]) represent another group of (sub-optimal but sometimes more realistic) adversaries who do not have the required capabilities to profile the leakages of the target. This difference is important as it was shown in a recent work by Whitnall et al. that we have a strict separation between such attacks [22]. That is, there (theoretically) exist leaking devices that can only be attacked under some a priori assumptions obtained through profiling. In practice though, the actual scenarios where only profiled attacks can succeed usually correspond to so-called "highly non-linear" leakage functions [21] - of which the existence remains an open question. Informally, a leakage function is considered non-linear if it cannot be accurately estimated by a linear combination of some intermediate bits processed by the target device (e.g. with Schindler et al.'s stochastic approach [16]). We say that it is highly non-linear if the resulting models do not allow successful key recoveries.

Security against side-channel attacks is usually obtained through a combination of countermeasures. As such countermeasures imply significant performance overheads, the best tradeoff between efficiency and security for small embedded devices has become an important research topic. But as usual in hardware design, the final performances of a taped out chip is not the only cost criteria. In particular, it has been frequently observed that countermeasures looking sound on a mathematical basis could be less effective than expected, because of physical artifacts (the problem of masking and glitches is a well-known example of this concern [10]). Hence, avoiding such limitations as early as possible in the development of a cryptographic implementation is also desirable. As a result, the improvement of simulation-based side-channel security evaluations has become another topic of interest [8,13], with the long-term goal to develop integrated design flows, with physical security as part of the optimization criteria [12,20]. Quite naturally, such a goal also raises questions regarding the relevance of simulated leakages, as noted by Tiri and Verbauwhede [19], or more recently by Renaud et al. [14]. The latter used state-of-the-art evaluation tools for quantifying information leakages, and put forward a difference of linearity between actual and simulated traces, raising questions about both the relevance of simulations for this criteria, and the possibility to design circuits with non-linear leakages.

In this paper, we contribute to these two important issues. Namely, we first narrow the gap between simulations and actual measurements. For this purpose, we develop a flexible model that captures measurement artifacts in side-channel attacks and integrate it into SPICE simulations. In order to validate it, we then analyze a Dynamic and Differential Swing-Limited Logic (DDSSL [3]) implementation of the AES S-box, in a 65nm CMOS technology (implementation details of the DDSSL S-box are beyond the scope of this study and can be found in [7]). This target circuit was chosen because it includes a circuit-level countermeasure for which non-linear leakages are expected (and have been previously observed in simulations [14]). But as highlighted in this previous work, such a non-linearity was not found in actual measurements. We repeat this comparison with gradually improved simulation models, that include elements such as the cables, socket and package that potentially affect the leakage traces, eventually leading to ac-

curate approximations of our measurements. By performing a regression-based information theoretic evaluation of our target implementation, we demonstrate how certain elements of the model can explain the transformation of a leakage function, from highly non-linear to fairly linear. Since this transformation is quite independent of the logic style design, our study consequently raises doubts regarding the possibility to design a dual-rail implementation with highly non-linear (and hard to linearize) leakages. Eventually, and despite the fact that the model we propose is specific and adapted to the measurement environment that we used in our experiments, we mention that our approach is rather generic. That is, the contribution of each element in the model will of course differ depending on the type of package, socket (if used), type of cable, etc. But we expect the way we include these elements in a model and their relative importance to remain meaningful for a wide range of implementations and setups.

2 Preliminaries

2.1 Notations

In this work, capital letters are assigned to random variables, while lower case letters refer to samples of these random variables. For example, L is the random variable representing a leakage and l is an actual power trace picked up from this distribution. The power trace is composed of t time samples. Generally, the leakage function has two input arguments: the discrete random variable X which denotes the value of the processed data under investigation, and the continuous random variable N which represents the noise in the measurements. The leakage function variable denoted by $L(,)$ contains either random variable arguments or fixed arguments. For example, $L(x, N)$ is a random variable representing the noisy traces corresponding to a fixed processed data x . We also denote the t^{th} time sample in a leakage trace as $L_t(,)$. In the measurement environment, we finally define (noise-free) mean traces as:

$$\overline{L_t^{meas}}(X) = \hat{\mathbf{E}}_n L(X, n), \quad (1)$$

where $\hat{\mathbf{E}}$ denotes the sample mean operator¹. Note that in the simulation environment, the provided traces are noise-free by default. In this case, and in order to analyze the impact of noise on the security of the S-box implementation, we added a Gaussian noise to the simulations² (this is a usual assumption in side-channel attacks, which will also be confirmed in Section 3). As a result, our investigations considered three types of leakage traces. The first case is the simulated leakage function which is given by:

$$L_t^1(X, N) = L_t^{sim}(X) + N. \quad (2)$$

¹ The noise-freeness naturally depends on the sampling, but in view of our low-noise measurements, we were able to extract well estimated means in our experiments.

² Gaussian noise is added to the simulated traces in a post processing step assuming the noise-free simulated traces to provide the means of our leakages.

The second case is when the actual power traces are considered. Here, the noise is directly present in the measurements obtained from the oscilloscope and the corresponding leakage function is given by:

$$L_t^2(X, N) = L_t^{meas}(X, N). \quad (3)$$

The third case is a hybrid leakage function combining the noise-free mean traces of Equation 1 with Gaussian noise:

$$L_t^3(X, N) = \overline{L_t^{meas}}(X) + N. \quad (4)$$

2.2 SCA evaluation metrics

In order to evaluate the leakage of the measured and simulated traces of the DDSLL S-box and assess their linearity, we estimated the information theoretic metric put forward in [17] and refined in [15]. That is, we computed the Perceived Information (PI) that corresponds to the amount of information that can be exploited by a side-channel adversary given a certain leakage model, namely:

$$\hat{\text{PI}}(X; L) = H[X] - \sum_{x \in X} \Pr[x] \sum_{l \in L} \Pr_{chip}[l|x] \cdot \log_2(\hat{\Pr}_{model}[x|l]).$$

It captures the accuracy of the adversary’s leakage model estimate (given by $\hat{\Pr}_{model}[x|l]$) at predicting the true (unknown) leakage function of an implementation (denoted as $\Pr_{chip}[l|x]$). In case these two distributions are identical (e.g. in a simulated environment), we have a perfect evaluation and the PI is equivalent to the standard definition of mutual information (i.e. it captures the worst-case information leakages). By contrast, if these distributions deviate - because of practical limitations in the number of traces used for profiling, or because of a simplified (e.g. linear) leakage model - the PI is the (best available) estimate of the target device’s leakage, biased by this slightly incorrect model.

2.3 Estimation tools

As previously mentioned, computing the PI first requires to estimate the leakage distribution with a model $\hat{\Pr}_{model}[x|l]$. Then, the evaluator just has to sample values l from the true distribution $\Pr_{chip}[l|x]$ and estimate the previous equation from it. In this section, we briefly describe the statistical tool we used for this purpose, namely Schindler et al.’s profiled stochastic approach [16]. It essentially aims to approximate the leakage samples with a linear combination of some base vectors. That is, during profiling the adversary chooses the base vectors $g_0(x)$, $g_1(x)$, \dots , $g_d(x)$. These base vectors represent monomials in x (the input and/or output and/or intermediate bits of the DDSLL S-box under attack), e.g. $d=8$ in case of a linear model for an 8-bit S-box. Then, the adversary performs a regression in order to build a model $\hat{L}_t(x, N) = \sum_i \beta_{i,t} \cdot g_i(x) + N$ that best suits the true leakages. Of course, the stochastic approach with a linear model cannot be perfect if the leakage function is non-linear. But by increasing the

number of elements in the basis, non-linearities can additionally be captured (e.g. with a $d=256$ -element basis, a stochastic model is equivalent to the templates in [2], excepted that it only estimates a single variance for N). In the rest of this paper, the use of the stochastic approach was naturally motivated by our goal to analyze the linearity of different (simulated and real) leakage functions.

2.4 Test chip and measurement setup

The chip was fabricated using a low-power 65 nm technology. All the input, output and clock signals of the DDSLL S-box are buffered. The S-box is powered by its own supply rail which is different from that of the buffers in order to directly assess the security of the DDSLL S-box implementation by itself. The chip is packaged in a 44-pin CQFP package. To study the power traces, the voltage drop on a resistor ($1k\Omega$) introduced in the path of the power supply of the measured S-box is monitored using a differential probe and adequately manipulated in a post processing step to get the power traces. We used a high sampling rate oscilloscope (1 Gsample/second), while running the chips at 2 MHz (motivated by interface constraints of our prototype board). Measurements were repeated 100 times to assess the security of DDSLL S-box under real noise. On the other hand, simulated power traces are obtained using post-layout SPICE simulations. Only parasitic capacitances to GND are extracted from the layout using Synopsys Star-RCXT. Both simulations and measurements are done at ambient temperature using a nominal supply voltage of 1.2 V. Figure 1 shows the floor plan of the test chip and all the PADS description are given in Table 1.




Fig. 1. Floor plan of the test chip.

Table 1. Description of the different PADs shown in Fig. 1.

Model	Description
<i>in</i> [0-7]	Input analog PADs
<i>Dout</i> [0-7]	Output analog PADs
<i>Clk</i>	Clock analog PAD
<i>VDDpwrDDSLL</i>	Supply analog PAD for the DDSLL AES S-box
<i>VDDBUFL</i>	Supply analog PAD for core buffers
<i>VDDBUFH</i>	Supply PAD for I/O buffers
<i>GND</i>	2 Ground analog PADS
<i>VDDE</i>	Supply PAD for the PADS
<i>GNDE</i>	Ground PAD for the PADS

3 Simulated traces vs. actual measurements

In this section, we exploit a linear regression based information theoretic analysis to analyze simulated and actual leakage traces, and to evaluate the practical relevance of certain assumptions regarding the physical behavior of our measurements. For this purpose, we essentially compute the PI between 256 events x given their leakage, using a model $\Pr_{model}[x|l]$ obtained thanks to the stochastic approach³. Our analyses are limited to the evaluation phase of the DDSLL S-box, which have been previously shown to leak the most information [14].

3.1 Direct analysis

In the direct analysis, we compare preliminary post-layout SPICE simulations and actual measurements. For the preliminary simulations, we simply probed the current flowing from the supply voltage without adding a resistor in its path and without any model representing the measurement specificities. The upper part of Figure 2 plots the current traces of both preliminary simulations and measurements, allowing to highlight two significant differences. First, we observe the different scales in the X-axes: computations end in about 8 ns for preliminary simulations, whereas in measurements, they last around 100 ns. Next, there exist high frequency components resulting from the computations of the different blocks of the S-box in preliminary simulations (assumably related to the high switching activity in strong cryptographic functions). By contrast in measurements, these high frequency components are filtered out due to the presence of different elements related to the chip, PCB, cables and equipments. Besides, it is also noticeable that measured traces oscillate at a changing frequency that is not observed in simulations. This point can be explained by the fact that in

³ Strictly speaking, there are 256^2 transitions that could be considered. To reduce the cost of our analysis, we only considered transitions between 0 and a value between 0 and 255. From past experiments, we do not expect this restriction to have a strong impact on our conclusions, in particular for the part related to the leakages linearity.

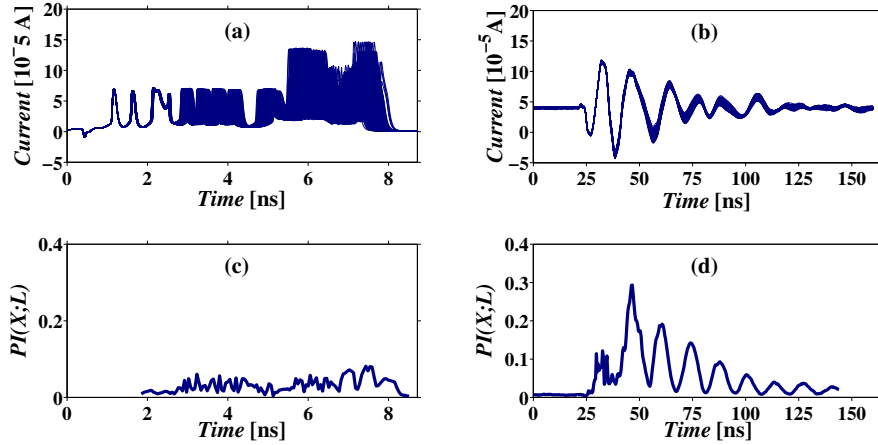


Fig. 2. (a) Simulated and (b) measured current traces of the DDSLL S-box during the evaluation phase for different inputs (inputs transition from 0 to an arbitrary state from 0 to 255). (c) Perceived information using a linear stochastic model in function of time for simulated traces without any noise and (d) measured traces with real noise.

preliminary simulations, the traces are perfectly aligned with minimal amplitude and time shifts at the beginning of the clock cycle, and start to misalign as the computations progress. Such a phenomenon is mainly due to the presence of unavoidable imbalances between the capacitances of the differential signals and circuitry which lead to some dependencies on the processed data. Due to the larger span of the measured traces over time, this misalignment appears to be smaller in actual traces. Furthermore, it is also reshaped by the various components present in the measurement setup forming oscillations

As a natural complement to this informal analysis, we investigated the informativeness of all the time samples in the simulated and measured traces, as illustrated in Figure 2 (c) and (d), respectively. The leakage function in Figure 2 (c) corresponds to simulations (defined as $L_t^1(X)$ in Section 2.1) where the noise parameter is set to 0. The leakage function in Figure 2 (d) corresponds to the actual measurements $L_t^2(X, N)$. Clearly, the linear stochastic model is unable to efficiently extract information from the simulated traces (for all time samples) as the $PI < 0.1$ for basic simulation, whereas it is quite successful in the actual measurement case where the $PI \simeq 0.3$ for measurement noise in the order of $6 \cdot 10^{-6}$. So we can indeed conclude that there is a significant difference of linearity between these two contexts. In view of the previously mentioned filtering effects, a natural explanation would be that these deviations can be explained by the measurement setup. As a result, we try in the next section to incorporate measurement artifacts in our simulation environment.

3.2 Measurement artifact models

A measurement environment includes different components such as equipment, cables, PCB, socket, chip’s package and pads, . . . The precise modeling of these components is a tedious task and goes with a risk to significantly increase the simulation time (which is of course undesirable). As a result, our main goal is to have a flexible model that adequately captures the most significant physical phenomenon contributing to the filtering of the traces seen in measurements (thus removing all the nonlinearities, as will be shown in the following section), within reasonable simulation time. One way to do this is to gradually add components that we believe important to the model, until the simulations are “good enough”. For example, we easily found out (as will be confirmed in the rest of the section) that the resistor in the supply path, package [4], QFP socket [18], cables and differential oscilloscope probe are important components in the model⁴. By contrast, we choose to neglect other components to keep the model simple, such as the terminations of the testing equipment (e.g. the supply sources and the oscilloscope), the bonding wires, the pads and the PCB connecting paths.

Note that a previous work [5] developed a method to simulate cryptographic systems using an equivalent circuit model (the linear equivalent circuit and current source model [6,11]), based on real measurements obtained from a prototype chip. The authors built their model for an FPGA implementing the AES algorithm and accurately emulated the peripheral circuitry, such that the simulation results are in agreement with measurements. Our approach is rather different because we focus on building an admittedly simpler peripheral model that does not need prior measurements (which typically corresponds to the situation during ASIC developments). So both approaches can be viewed as complementary: ours for preliminary investigations based on simulation models only, and the one in [5] whenever measurements are available to refine the model.

Figure 3 illustrates the equivalent circuit model of the measurement artifacts used in our improved simulation environment, and Table 2 details the description and value of each element in the model. Based upon the length of the cables used, the self inductance is calculated for the supply and input/output signals. However, since the ground occupies the entire lower plane of the PCB and is at the same time connected to the supplying equipment via several cables (the grounds of all equipment are connected together), we assumed the ground cable inductance to be less than the supply cable inductance. As for the mutual inductances and capacitances between adjacent pins in the socket and package, only those to the immediate pins are shown in Figure 3. Nevertheless, mutual inductances to the next-of-immediate pins in the package model are also included (not shown in the Figure 3 for simplicity). The relationship between pins can be seen from the test chip in Figure 1 in order to effectively model the mutual inductances and capacitances. The value of the resistor in the S-box supply path

⁴ Models for the package [4] and QFP socket [18] do not exactly correspond to our setup (e.g. they differ in pin count) - but were the only publicly available ones.

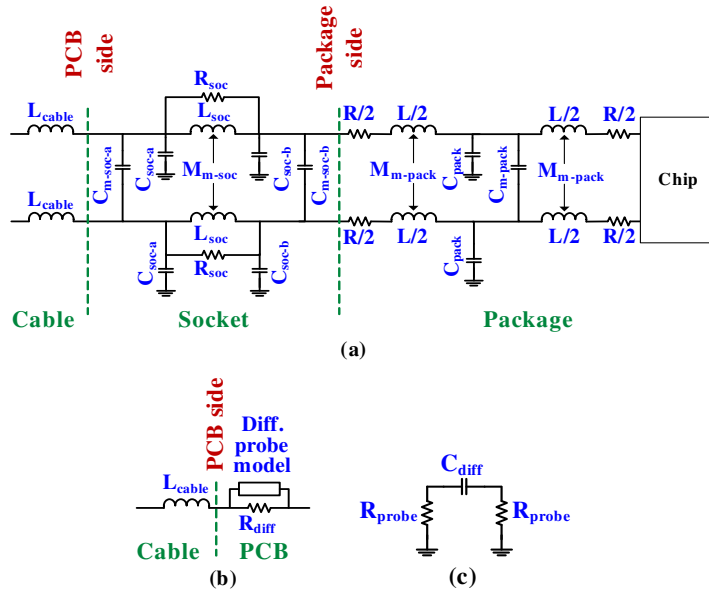


Fig. 3. Equivalent circuit model of (a) two arbitrary signals illustrating the cable, socket and package elements featuring self and mutual inductances/capacitances and (b) extra components that are included in the S-box supply path (namely, R_{diff} and differential probe model which is further detailed in part (c) of the figure).

depends on the amount of instantaneous current to be monitored and the resolution of the oscilloscope. The probe model can be simplified by a small differential capacitance and two resistors connected on each side to the ground.

Of course, neither our measurement artifacts model nor the one in [5] can be considered as completely generic, in the sense that they both need “some prior knowledge” about the target circuit (FPGA / ASIC, type of package to be used, ... etc.) and measurement setup (socket / no socket, PCB, types of probes, ... etc.). Quite naturally, a circuit designer is always advised to incorporate the most precise information in his peripheral model. Yet, we believe the most important shortcomings in security evaluations happen when measurement setups are simply ignored from the simulation environment. In this respect, the quality of our model gradually improves with the accuracy of its component specifications, but even approximated specifications already allow increasing the relevance of simulations significantly, which is the main contribution of this work.

3.3 Improved analysis

We now employ the model illustrated in the previous section in order to improve our analysis and reflect our measurement environment in simulations. First, we show in Figure 4 (a) how close the shape of the simulated traces gets to the measured ones in Figure 2 (b). In particular, the figure now features an expansion

Table 2. List of the elements in the measurement environment model of Figure 3.

Element	Symbol	Description	Value
Cable	L_{cable}	supply inductance	688 nH
		input/output inductance	300 nH
		GND inductance	200 nH
Socket [18]	L_{soc}	lead inductance	1.35 nH
	R_{soc}	parallel lead resistance	600 Ω
	C_{soc-a}	cap. to ground (PCB side)	0.3 pF
	C_{soc-b}	cap. to ground (pack. side)	0.45 pF
	L_{m-soc}	mutual inductance	0.3 nH
	$C_{m-soc-a}$	mutual cap. (PCB side)	0.09 pF
	$C_{m-soc-b}$	mutual cap. (pack. side)	0.09 pF
Pack. [4]	L	inductance	1.2 nH
	R	series resistance	0.28 Ω
	C_{pack}	cap. to ground	0.1 pF
	L_{m-pack}	mutual inductance	1.3 nH
Diff. Probe	C_{diff}	capacitance	0.7 pF
	R_{probe}	resistance	25 k Ω
	R_{diff}	resistor in S-box VDD path	1 k Ω

in the time span (with respect to preliminary simulations), filtering off the high frequency components and oscillations at a changing frequency. Although neither the time span nor the frequency of oscillations of the improved simulated traces perfectly fit the measured results, the model proves to adequately emulate the measurement specificities to a much better extent. In particular, it directly shows the impact of different components on the linearization of the leakage samples. As can be observed in Figure 4 (b), most of the time samples now appear to be sufficiently linear (i.e. the leakage function is well modeled by the linear stochastic approach based upon the output bits of the S-box).




Fig. 4. (a) Simulated current traces for the the DDSLL S-box with measurement artifacts and (b) corresponding perceived information using a linear stochastic model.

As a complement, we investigated the impact of some key elements in the measurement artifacts model on linearizing the leakage function of the simulated traces. For this purpose, we computed the PI in function of the simulated noise

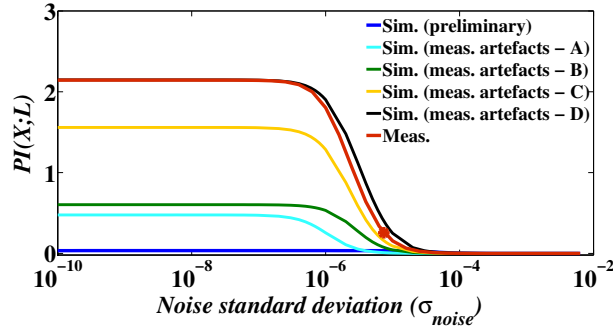


Fig. 5. Perceived information using the linear stochastic approach in function of the noise for preliminary simulated traces, simulated traces with measurement artefacts (accumulated effect of different elements in the model - A, B, C and D) and measured traces of the DDSLL S-box (The PI with real noise is marked with a star (*)).

for leakage variables $L_t^1(X, N)$ and $L_t^1(X, N)$ (represented by solid lines) as well as the PI for the actual noise in our measurements for leakage variable $L_t^2(X, N)$ (represented by a dot). The results of this experiment are reported in Figure 5, in which the curves are given for the single time sample that maximizes the perceived information using the linear stochastic model in all the investigated cases (denoted with letters A to D, depending on the simulation model). The description of these different models used is provided in Table 3.

Table 3. Description of the different models used in Figure 5.

Model	Description
A	1 k Ω + diff. probe
B	1 k Ω + diff. probe + package and socket
C	1 k Ω + diff. probe + package and socket + V_{DD} cable
D	1 k Ω + diff. probe + package and socket + V_{DD} cable + GND cable

As previously observed, it is clear that the leakage function given by preliminary simulations is highly non-linear. Next for case A, only the 1 k Ω resistor in the S-box supply path is modeled without any cable, socket or package models, leading to a first (slight) improvement of the linearity in simulated traces. Adding the socket and package models in case B (including all the mutual effects) without the cable inductances to the 1 k Ω resistor further increases the linearity at low noise levels, and shifts the curve towards the measurement one at real noise. Although not shown here, the mutual effects in the socket and package models did not have a significant impact on linearizing the simulated leakage function. Nevertheless, the use of the mutual inductances and capacitances in the model helped forming the oscillations in the simulated traces. Furthermore, including cable inductances in the supply paths (S-box and buffers) without those of the GND in case C significantly raises the perceived information to 1.6 bits at low

noise levels, and pushed the curve even closer to the measurement one at real noise. This is mainly attributed to the fact that the effect of the cable inductance for the S-box supply is taken into account (with negligible impact of the cable inductance of the buffers supply). Finally, adding the *GND* cable inductances to the model in case *D* helped the simulated perceived information to coincide with the measurements with a slight overestimation at noise levels higher than 10^{-6} . Here, the effect of ground bounce could be seen both in the improved simulated traces and linearization of the leakage function to match the measurements.

Clearly, the cable inductances of both the S-box supply and *GND* are the dominating contributors to the linear dependency between the actual traces and the S-box output bits. They helped forming the low-pass filter that leads to smoothing the shape of the traces by removing the high frequency components, which eventually emphasizes the amplitude shifts and increase the linearity of the leakage function. Overall, the careful choice of the contributing elements in the measurement artifacts model rendered the simulation time overhead negligible since the DSSLL S-box itself (on-chip) has many orders of magnitude more elements than the simple off-chip model (*the simulation time overhead due to our model was less than 1% of the original S-box simulation time*). Also, it is important to note that the Gaussian noise hypothesis is reasonably accurate, as the hybrid leakages with simulated noise correspond well to the actual measurements with real noise (which is confirmed by the position of the dot in Figure 5).

Consequently, our study highlights the importance of having (even approximate) prior knowledge about the measurement specificities of a cryptographic implementation during its early design cycle. This way, the designer can incorporate these artifacts in the simulation environment and make any design changes (if necessary) before having to actually manufacture the final chip.

4 Conclusions

In a previous work by Renaud et al. at CHES 2011, a difference in the linearity of the leakage functions was observed between actual and simulated traces, raising doubts about the relevance of simulations and the possibility of designing circuits with non-linear leakages. This result suggested the improvement of simulation-based side-channel security evaluations as an important open problem, and a preliminary step for the integration of such tools in standard design flows. In this paper, we consequently aimed to improve this situation, both regarding the linearity of simulated traces (that matters for the application of non-profiled side-channel attacks) and regarding the amount of information they provide (that determines the data complexity of profiled side-channel attacks). For this purpose, we first modeled the measurement specificities associated to side-channel attacks and integrated it into SPICE simulations. Then we validated the model by gradually accumulating the different components that we believe important such as the resistor in the supply path, package, socket, and cables, until the simulated and measured traces of the target implementation

provided the same level of linearity and similar perceived information. By doing so, preliminary simulated traces were transformed from highly non-linear to fairly linear traces. Our investigations confirmed the need to incorporate the key physical artifacts in the simulation environment to obtain an accurate assessment of countermeasures against side-channel attacks at the design stage. Besides, our study also raises the question whether it is possible to design an implementation with highly non-linear leakages in light of the linearization effect that the physical artifacts inevitably have on the final outcome.

Acknowledgements. This work has been funded in parts by the European Commission through the ERC project 280141 (acronym CRASH), the European ISEC action grant HOME/2010/ISEC/AG/INT-011 B-CENTRE project, and the Walloon region WIST program project MIPSs. F.-X. Standaert is an associate researcher of the Belgian Fund for Scientific Research (FNRS-F.R.S.).

References

1. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2004.
2. Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In *Proceedings of Cryptographic Hardware and Embedded Systems, CHES*, pages 13–28, 2002.
3. Ilham Hassoune, François Macé, Denis Flandre, and Jean-Didier Legat. Dynamic differential self-timed logic families for robust and low-power security ICs. *Integration*, 40(3):355–364, 2007.
4. Texas instruments. AN-1205 electrical performance of packages. Application report, May 2004.
5. K. Iokibe, T. Amano, K. Okamoto, and Y. Toyota. Equivalent circuit modeling of cryptographic integrated circuit for information security design. *Electromagnetic Compatibility, IEEE Transactions on*, 55(3):581–588, 2013.
6. K. Iokibe, R. Higashi, T. Tsuda, K. Ichikawa, K. Nakamura, Y. Toyota, and R. Koga. Modeling of microcontroller with multiple power supply pins for conducted emi simulations. In *Advanced Packaging and Systems Symposium, 2008. EDAPS 2008. Electrical Design of*, pages 135–138, 2008.
7. Dina Kamel, Mathieu Renauld, David Bol, François-Xavier Standaert, and Denis Flandre. Analysis of dynamic differential swing limited logic for low-power secure applications. *Journal of Low Power Electronics and Applications*, 1(2):98–126, 2012.
8. François Macé, François-Xavier Standaert, and Jean-Jacques Quisquater. Information theoretic evaluation of side-channel resistant logic styles. In Pascal Paillier and Ingrid Verbauwhede, editors, *CHES*, volume 4727 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2007.
9. Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power analysis attacks - revealing the secrets of smart cards*. Springer, 2007.

10. Stefan Mangard, Thomas Popp, and Berndt M. Gammel. Side-channel leakage of masked cmos gates. In Alfred Menezes, editor, *CT-RSA*, volume 3376 of *Lecture Notes in Computer Science*, pages 351–365. Springer, 2005.
11. K. NAKAMURA. Emc macro-model (leccs-core) for multiple power-supply pin lsi. *Proc. EMC'04, Sendai, June*, 2004.
12. Francesco Regazzoni, Alessandro Cevrero, François-Xavier Standaert, Stéphane Badel, Theo Kluter, Philip Brisk, Yusuf Leblebici, and Paolo Ienne. A design flow and evaluation framework for dpa-resistant instruction set extensions. In Christophe Clavier and Kris Gaj, editors, *CHES*, volume 5747 of *Lecture Notes in Computer Science*, pages 205–219. Springer, 2009.
13. Francesco Regazzoni, Thomas Eisenbarth, Axel Poschmann, Johann Großschädl, Frank K. Gürkaynak, Marco Macchetti, Zeynep Toprak Deniz, Laura Pozzi, Christof Paar, Yusuf Leblebici, and Paolo Ienne. Evaluating resistance of meml technology to power analysis attacks using a simulation-based methodology. *Transactions on Computational Science*, 4:230–243, 2009.
14. Mathieu Renaud, Dina Kamel, François-Xavier Standaert, and Denis Flandre. Information theoretic and security analysis of a 65-nanometer DDSLL AES S-Box. In *Proceedings of Cryptographic Hardware and Embedded Systems, CHES*, pages 223–239, 2011.
15. Mathieu Renaud, 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*, pages 109–128, 2011.
16. Werner Schindler, Kerstin Lemke, and Christof Paar. A stochastic model for differential side channel cryptanalysis. In *Proceedings of Cryptographic Hardware and Embedded Systems, CHES, Springer, LNCS 3659*, pages 30–46. Springer, 2005.
17. François-Xavier Standaert, Tal G. Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In *Proceedings of the 28th Annual International Conference on Advances in Cryptology: the Theory and Applications of Cryptographic Techniques, EUROCRYPT*, pages 443–461, Berlin, Heidelberg, 2009. Springer-Verlag.
18. Giga test labs. ARIES electronics 64 pin QFP (0.55mm) test socket, Electrical characterisation 0.05 - 3.05 GHz. Characterisation report, Oct. 1997.
19. Kris Tiri and Ingrid Verbauwhede. Simulation models for side-channel information leaks. In William H. Joyner Jr., Grant Martin, and Andrew B. Kahng, editors, *DAC*, pages 228–233. ACM, 2005.
20. Kris Tiri and Ingrid Verbauwhede. A digital design flow for secure integrated circuits. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 25(7):1197–1208, 2006.
21. Nicolas Veyrat-Charvillon and François Xavier Standaert. Generic side-channel distinguishers: Improvements and limitations. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference*, volume 6841 of *Lecture Notes in Computer Science*, page 348. Springer, 2011.
22. Carolyn Whitnall, Elisabeth Oswald, and François-Xavier Standaert. The myth of generic DPA... and the magic of learning. Cryptology ePrint Archive, Report 2012/256, 2012. <http://eprint.iacr.org/>.