Sweeping for Leakage in Masked Circuit Layouts

Danilo Šijačić, Josep Balasch and Ingrid Verbauwhede
imec-COSIC, KU Leuven, Leuven, Belgium
{danilo.sijacic, josep.balasch, ingrid.verbauwhede}@esat.kuleuven.be

***Abstract*—Masking schemes are the most popular countermeasure against side-channel analysis. They theoretically decorrelate information leaked through inherent physical channels from the key-dependent intermediate values that occur during computation. Their provable security is devised under models that abstract complex physical phenomena of the underlying hardware. In this work, we investigate the impact of the physical layout to the side-channel security of masking schemes. For this we propose a model for co-simulation of the analog power distribution network with the digital logic core. Our study considers the drive of the power supply buffers, as well as parasitic resistors, inductors and capacitors. We quantify our findings using Test Vector Leakage Assessment by relative comparison to the parasitic-free model. Thus we provide a deeper insight into the potential layout sources of leakage and their magnitude.**

***Index Terms*—Masking, Coupling, Layout Parasitics, SPICE**

I. Introduction

Power analysis [1] is the most prominent side-channel analysis (SCA) technique against cryptographic implementations. It exploits information embedded in the power supply current waveform of a device in order to extract its secret cryptographic keys. Masking [2], [3] is a well-studied countermeasure against SCA. It works by splitting all sensitive variables amongst multiple, randomized shares, and performing computations on these shares. In general, splitting an n-bit variable x into N shares entails randomly generating x_1, x_2, $\ldots x_{N-1}$ and computing x_N such that Eq. 1 holds; where $\circ$ indicates the type of masking:

$$x = x_1 \circ x_2 \circ \ldots \circ x_N. \tag{1}$$

Masking provides provable security, regardless of the implementation of the individual shares, under abstract leakage models. The most common assumption is that information leakage is a linear combination of leakages of the individual shares. For a hardware implementation of an arbitrary shared computation, the independence assumption can be represented with the circuit depicted in Fig. 1. All shares are connected in parallel to an ideal voltage source, capable of providing infinite supply current to each share.

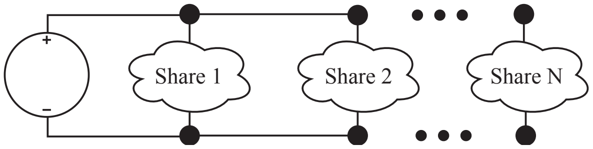

Fig. 1: Circuit model for the independent leakage assumption.

As masking schemes make no assumptions on how each share is implemented, they can be realized using standard cell ASIC libraries or on FPGAs. On the front-end of the design cycle, using gate-level modeling, the abstraction shown on Fig. 1 holds. In the back-end stages, *logic cells* are placed on a substrate, shared with different *physical cells* such as filler cells, a clock distribution network and a power distribution network (PDN). Logic cells are then interconnected using multiple metal layers in a process called routing. Consequently, a myriad of parasitic elements and non-linear effects emerge, unbeknown to masking models. We aim to uncover whether any of said phenomena can violate the independence assumption. To this end, we run analog simulations on a wholesome circuit model considering the effects of:

a) *PDN*, an analog circuit distributing a low-frequency signal across a wide area. It employs supply buffers to achieve sufficient current capacity demanded by the digital logic core. Large parasitic resistances and inductances of long PDN wires can cause *ground bounce* (increase of the low power rail from the reference) and *supply sag* (drop in the high power rail bellow the nominal value) [4]. As the chip-ground is often realized as a large metal plate, each node of the PDN additionally forms a significant capacitance towards the ground.
b) *Logic core*, the circuit's functionality composed of many tightly placed digital cells. Relatively short wires and library designers' efforts to avoid oscillations leave no significant parasitic inductances. As each share computes self-containedly, there are no wires—hence no parasitic resistances—between them. However, parasitic capacitors can electrically couple the shares via crosstalk. Thus, capacitors between routing wires, carrying high-frequency signals from different shares, remain threatening to the independence assumption.

A. Related Work

The usual layout suspects for SCA leakage are parasitic capacitors. Recent work [5] for the first time practically demonstrates information leakage of a masking scheme caused by placement and routing. The authors implement a masking scheme on an FPGA and compare two placements: unconstrained (shares placed closely together) versus constrained (shares placed far apart). Using Test Vector Leakage Assessment (TVLA) [6] they show that unconstrained placement lowers the SCA security. They suspect the crosstalk and the resistive (IR) supply voltage drop as culprits. Nevertheless, as FPGA layouts are closely guarded trade secrets, the sources of

978-3-9819263-4-7/DATE20/©2020 EDAA

leakage can not be further investigated. In [7], the authors rely on capacitive crosstalk to observe double the leakage of the Hamming Distance model on an 8-bit bus. In [8], the authors show that capacitive coupling between logically independent blocks can cause leakage.

Analog SPICE transistor models yield the most accurate representation of ASIC designs. Unfortunately, the number of transistors in digital circuits is prohibitively large for analog simulations to scale. Designers resort to gate-level simulators, based on piece-wise linear models extracted from complex transistor-level simulations. Highly robust CMOS logic style allows for such models: parasitic components may cause significantly different analog waveforms without altering the digital behavior. Consequently, high efficiency is obtained by overlooking the analog deviations. Such mismatch in modeling and toolchains can however lead to serious misrepresentation of SCA security in simulation. In [9], the authors show how different RC extraction methods can lead to contrasting security evaluations using SPICE. In [10], the authors show there is little distinction between the wire-load models of the pre-layout netlist versus the extracted RC parasitics using state-of-the-art digital simulators.

B. *Contributions*

We propose a SPICE model for the co-simulation of the analog PDN with the digital logic core. It accounts for the finite current capacity power supply buffers, parasitic resistors, inductors and capacitors that inevitably occur in the long lines of the PDN, along with the parasitic capacitors of the logic core. We use this model to investigate the gap between digital and analog modeling for SCA security by analyzing representative masked gates devised to provide first-order security. By means of experiments, we determine the impact of all parasitic elements that exist in a masked circuit, while allowing it to operate correctly. Thus, we provide a deeper insight into the sources of leakage in masked designs.

II. Methodology

Parasitic elements are inherent to any circuit. As byproducts of the placement and routing, their values can not be controlled directly. Designers, with the aid of EDA tools, limit their presence to meet the performance constraints. Digital models and extraction tools are tailored for this purpose, leaving a gap between analog and digital behavior. We aim to investigate this gap to detect how may parasitic elements—while satisfying performance requirements—impact the SCA security of masked designs and to what extent.

To this end, we use the analog SPICE simulator from the Synopsys FineSim v2018.09 suite and the 45nm open-source standard-cell library from Nangate [11] in junction with predictive transistor models [12]. We use transient simulation with a 10 ps step, equivalent to a 100 GS/s sampling rate. To ensure equidistant samples, we set the `strobeperiod=10p` argument of the `.tran` simulation.

Producing single precision, noiseless and perfectly aligned traces, our setup presents an overly pessimistic SCA scenario. Our aim is however not to determine how feasible it is to exploit, or even measure, such leakage in a practical settings. We aim instead to diagnose which parasitic elements may compromise the SCA security. Lastly, we focus on the potential sources of leakage stemming from the design, rather than the influences of measurement setups [13], [14].

Our methodology works as follows. We conduct a preliminary investigation of possible leakage sources, without questioning how probable they may be. We annotate all parasitic elements individually. Next, we perform transient analysis, sweeping the value of each parasitic element. We choose the sweep ranges such that they vary from negligibly small to values that cause the circuit to malfunction. We probe voltages of input and output data nodes to continually verify the correctness of computations. Simulations are driven using randomly generated input vectors to form side-channel traces. Using TVLA (c.f. II-C) as the core metric, we quantify information leakage in the function of each parasitic element and its values, relative to the parasitic-free case.

A. *SPICE Model*

Our SPICE model for co-simulation of the PDN with the logic core is depicted in Fig. 2. We buffer the ideal voltage source using standard-cell buffers of finite current capacity, as shown in the dashed red rectangle. Buffers generate two different power supplies: one for the core logic and one for buffering of the ideal inputs coming from the SPICE vector file (`*.vec`). The dashed blue rectangle shows how input and output signals are driven and loaded. We ensure that signals adhere to the models of the standard-cell library, preventing the influence of any ideal waveforms that may offset the results obtained for smaller target circuits. Thus we create a "sterile" environment for observing the supply currents of shares. Furthermore, we account for the high-frequency response caused by supply currents drawn by the CMOS gates in logic shares. Unlike the constant impedance of parasitic resistors $Z_R = R$, impedances of parasitic inductances $Z_L = j\omega L$ and capacitances $Z_C = \frac{1}{j\omega C}$ depend on the angular frequency of the inciting current. Thus the voltage fluctuations over Z_L and Z_C cause non-linear ground bounce and supply sag.

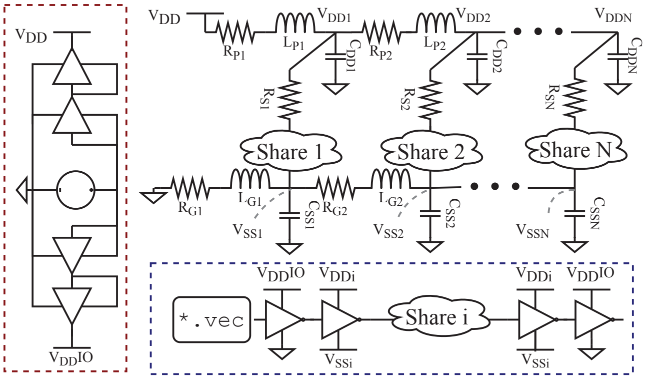

Fig. 2: SPICE model of R, L and C parasitics of the PDN.

Lastly, in contrast to a security evaluation where the waveform of merit is the supply current that can be measured, we chose the supply current $i_{CORE} = \sum_{j=1}^{N} i_{V_{DDj},V_{SSj}}$ as the preferred side-channel for diagnostic purposes.

B. Target Circuits

Any Boolean function over binary fields $GF(2^n)$ can be implemented in its algebraic normal form using only two-input XOR (XOR2) and AND (AND2) gates. Thus secure instances of these two gates are fundamental to protect any Boolean masking scheme, such as Threshold Implementations (TI) [15]. Boolean masking schemes use XOR as the sharing operator, making the linear layer masking trivial. Given an n-bit XOR2 gate $z_i = a_i \oplus b_i$, sharing it into N shares requires $z_{i,j} = a_{i,j} \oplus b_{i,j}$; where $i \in \{1, 2, \ldots, n\}$ and $j \in \{1, 2, \ldots, N\}$. Fig. 3 (top) shows the schematic of such an XOR2 gate with $N = 2$ that provides first-order security, that is, it prevents SCA that exploit leakage in the first-order statistical moment. In contrast, masking non-linear operations is a demanding task. Fig. 3 (bottom) shows the schematic of a masked AND2 gate with $N = 3$, proposed in [16]. It also provides first-order security, but requires an additional share (and additional input random bit r) compared to the protected version of XOR2. Our study focuses on these simple circuits, as they allow to address virtually any Boolean masking schemes, looking from the algebraic perspective. Additionally, their small sizes are favorable for SPICE simulation.

C. Security Metric

We use TVLA [6] as security metric due to its generic nature, fast computation and suitability to evaluate masking schemes with low number of shares [17]. TVLA determines whether the statistical moments of two sets of data, of N_1 and N_2 elements, are distinguishable by using Welch's t-test. Upon partitioning traces based on unshared data, the security order is defined as the highest statistical moment in which the Welch's t-statistic does not exceed a confidence interval of $|t| \leq 4.5$. For first-order security, a score t^i is computed for each sample i in the supply current waveform as per Eq. 2, where μ^i and $(\sigma^i)^2$ are sample mean and variance normalized to the number of samples in each set N_1, N_2.

$$t^i = \frac{\mu_1^i - \mu_2^i}{\sqrt{(\sigma_1^i)^2/N_1 + (\sigma_2^i)^2/N_2}} \tag{2}$$

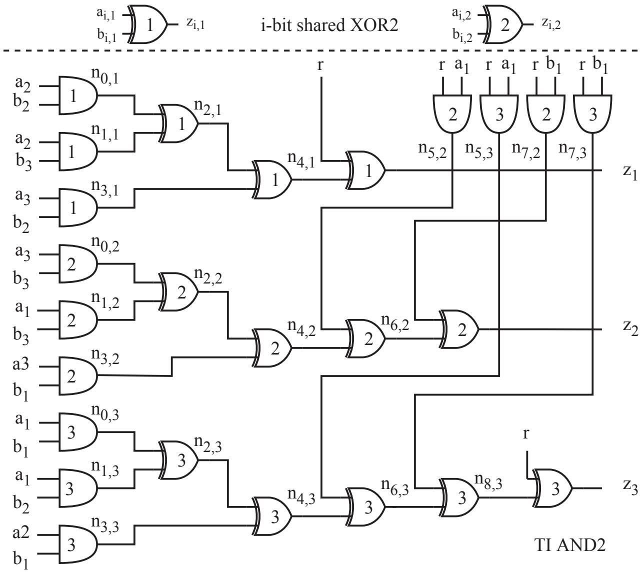

Fig. 3: Target circuits XOR2, $N = 2$ (top), AND2, $N = 3$ (bottom), implemented using `XOR2_X1` and `AND2_X1`; each gate is marked with its share number i.

The resulting t-trace is a temporal waveform that contains all t^i samples for a fixed number of traces (constant N_1 and N_2). In our experiments in Sect. III we often plot the trend of the $max(|t|)$ score for an increasing number of traces. A constant $max(|t|)$ trend with $|t| \leq 4.5$ indicates no leakage effects, while a rising $max(|t|)$ trend indicates that collecting more traces may lead to a vulnerability. By fixing the number of traces, we can additionally observe the $max(|t|)$ trend in the function of the swept parasitic element value. Doing so allows us to identify the parasitics that may compromise the SCA security. Should we increase the number of traces increasing $max(|t|)$ trends would only lead to higher t-scores, further strengthening our observations.

III. Experimental Results

In this section we provide experimental results of the parametrized sweeps of parasitic elements. We target the first-order secure instances of XOR2, $N = 2$ and AND2, $N = 3$ gates using TVLA, partitioning based on unshared outputs.

A. Impacts of PDN

To examine the impact of the PDN we use an 8-bit instance of the XOR2, $N = 2$ gate depicted in Fig. 3 (top). We instantiate eight `XOR2_X1` cells in parallel to instigate a more significant supply current. Still the current drawn is minute compared to any full-fledged design. Therefore, we need substantial R, L, C values to perturb the supply rails significantly. As the circuit is completely linear and inputs of each share are independent, all observed leakage may stem only from the PDN. We simulate 2^{15} traces for each experiment.

Firstly, we investigate the impact of the power supply's finite capacity to the SCA security regardless of the parasitic elements. By setting $R_{Pi} = R_{Gi} = 0\,\Omega$, $L_{Pi} = L_{Gi} = 0\,\mathrm{H}$ and $C_{DDi} = C_{SSi} = 0\,\mathrm{F}$ and omitting the buffers (`no BUF_*`), we reduce our model to the idealization depicted in Fig. 1. We then involve buffers of different drive strength and plot the average supply sag for each case in Fig. 4 (left) and the $max(|t|)$ score for the increasing number of observations in Fig. 4 (right). The solid black line represents the ideal case. As the ideal voltage source has infinite current capacity, it maintains the nominal supply voltage and the $max(|t|)$ score remains firmly around the 1.5 mark. Introducing buffers with finite current capacity leads to supply sag. In particular, the supply drops within 10%, 15%, and over 20% of the nominal value on average for `BUF_X32`, `BUF_X8`, and `BUF_X1`, respectively. For the latter, the supply sag disrupts the correct computation. In all cases we observe an increasing trend in the $max(|t|)$ score when compared to the `no BUF_*` case, although it never reaches the 4.5 threshold (up to 2^{15} traces). For the rest of experiments we use `BUF_32` to model the finite current capacity of the power supply. We refer to this as the referent (parasitic-free) case.

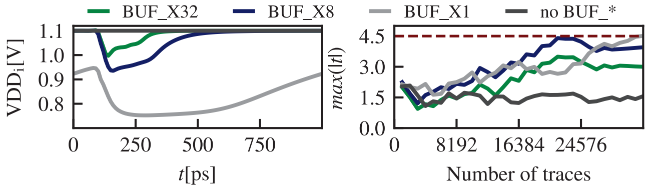

Fig. 4: Mean voltage on VDDi (left); $max(|t|)$ score using different supply buffers (right) on 8-bit XOR2.

Next, we move to investigate the impact of different parasitic elements of the PDN: resistors, capacitances and inductances. The results are shown in Fig. 5. The referent case is plotted in green, black traces correspond to parameter values that allow the circuit to function properly, and gray traces indicate the values that lead to failed computations. We overlap plots of the parasitic elements that exhibit similar behavior.

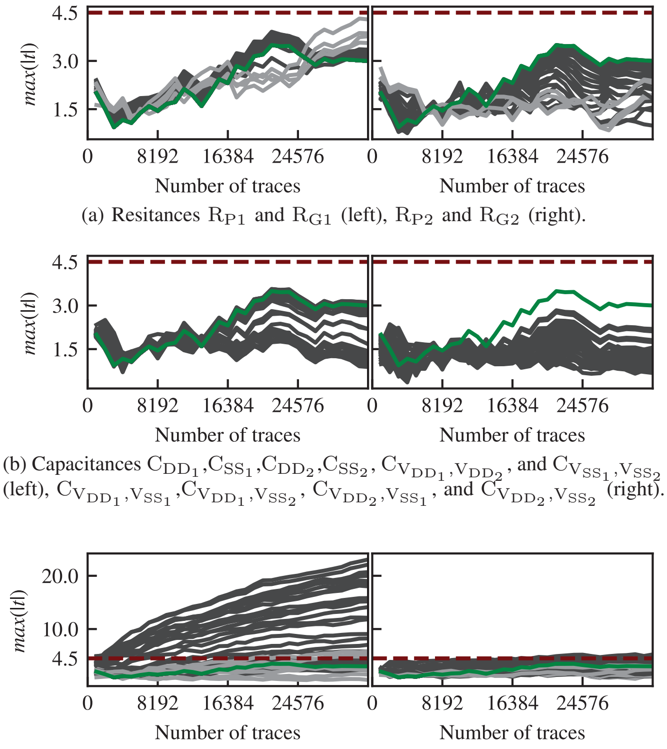

(a) Resitances R_{P1} and R_{G1} (left), R_{P2} and R_{G2} (right).

(b) Capacitances C_{DD_1},C_{SS_1},C_{DD_2},C_{SS_2}, $C_{V_{DD_1},V_{DD_2}}$, and $C_{V_{SS_1},V_{SS_2}}$ (left), $C_{V_{DD_1},V_{SS_1}}$,$C_{V_{DD_1},V_{SS_2}}$, $C_{V_{DD_2},V_{SS_1}}$, and $C_{V_{DD_2},V_{SS_2}}$ (right).

(c) Inductances L_{P1} and L_{G1} (left), L_{P2} and L_{G2} (right).

Fig. 5: PDN impact on 8-bit XOR2; overlapped sweep data.

Fig. 5a shows the impact of IR drop over parasitic resistors. We sweep resistor values between 1 Ω and 10 kΩ using a logarithmic sweep. On the one hand, R_{P1} and R_{L1} affect the supply nodes of both shares. Increasing them to 3.5 kΩ and 6.3 kΩ, respectively, leads to a slight increase in the $max(|t|)$ score, though remaining below the 4.5 threshold. Further increases of their value break the circuit functionality. On the other hand, R_{P2} and R_{L2} decrease the $max(|t|)$ score when increased up to 4.0 kΩ and 3.9 kΩ, respectively. R_{P2} and R_{L2} affect only the second share, limiting its current. Therefore, they lower the side-channel signal and, consequently, the amount of leakage.

Fig. 5b shows the impact of the PDN capacitors. We sweep capacitance values between 1 fF and 1 nF using a logarithmic sweep. In addition to the capacitors towards the ground plate, we sweep capacitors between supply nodes of the same polarity (left), and capacitors between supply nodes of the opposing polarity (right). The latter are referred to as decoupling capacitors in the literature. Serving as charge caches for the logic core, they lower the required bandwidth of the PDN. In other words, they filter information carrying high-frequency components of the supply current. Hence, we observe a decrease in the $max(|t|)$ score.

Fig. 5c shows the impact of the PDN inductors. We sweep inductance values between 1 pH and 1 μH using a logarithmic sweep. L_{P1} and L_{G1} lead to $max(|t|) > 4.5$ when they exceed 1.2 nH and 0.8 nH, respectively. They produce peak $max(|t|)$ scores of between 20 and 23 when increased up to 18 nH. L_{P2} and L_{G2} have a significantly lower impact, as they affect only the second supply node. Fig. 6 shows the magnitude of the $max(|t|)$ score in the function of parasitic inductors.

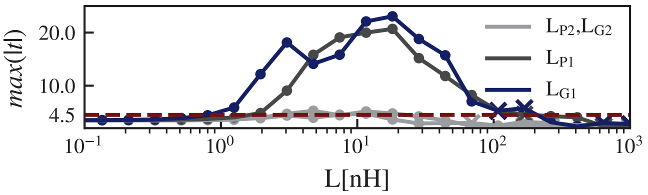

Fig. 6: Magnitude of the $max(|t|)$ for 2^{15} traces; • denotes correct and × failed computation of 8-bit XOR2.

B. *Effects of Coupling Capacitances*

The last part of our experiments examines the impact of coupling capacitances. For a circuit with n nodes, a total of $\binom{n}{2}$ capacitors can be annotated between them. Thus in order to keep the computational effort reasonable, we perform experiments on the gates presented in Fig. 3.

1) Shared XOR2, 1-bit: Given the 4 input bits, there exist only $2^{4\times2} - 2^4 = 240$ non-trivial input transitions to simulate. Excluding the supply to supply capacitors discussed in the previous section, we split the remaining $\binom{10}{2} - 6 = 39$ capacitors as follows:

- *share-rail* (SR) between data nodes of one share and the supply node of the same share;
- *cross-rail* (CR) between data nodes of one share and the supply node of another share;

- *share-data* (SD) between data nodes within one share;
- *cross-data* (CD) between data nodes of different shares.

We sweep each capacitance using a logarithmic sweep between 10 aF and 100 fF. The impact of all four categories of capacitors is shown in Fig. 7: SR, CR and SD (left) and CD (right). In accordance with theory, only CD capacitances, C_{z_1,z_2} in particular, result in a significant increase the SCA leakage. The rest of capacitances cause barely any deviation from the referent $max(|t|)$ trend.

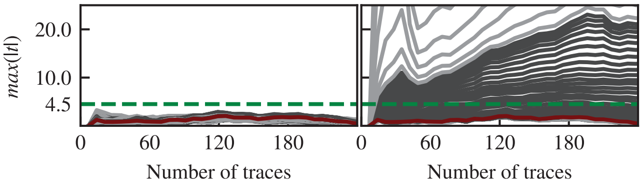

Fig. 7: Impact of SR, CR, SD (right) and CD (left) on 1-bit XOR2, $N = 2$; overlapped sweep data.

Fig. 8 shows the magnitude of the $max(|t|)$ in the function of different CD capacitances. The presence of leakage due to C_{z_1,z_2} starts at 0.3 fF, peaking to $max(|t|) = 23.1$ at $C_{z_1,z_2} = 19.3$ fF. The circuit begins to fail for larger values.

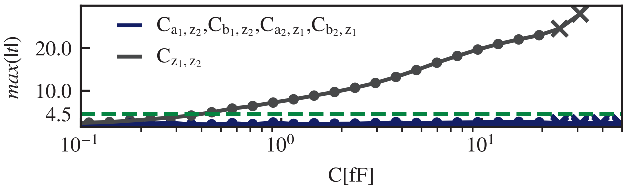

Fig. 8: Magnitude of the $max(|t|)$ for 240 traces; • denotes correct and × failed computation of 1-bit XOR2.

2) Shared TI AND2: Given the 7 input bits there exist only $2^{7\times2}-2^7$ non-trivial input transitions to simulate. As we annotate numerous 326 CD capacitors for the non-linear TI AND2, $N = 3$ gate, we lower the amount of experiments to 2^{13} traces. We sweep each capacitance using logarithmic sweep between 0.1 fF and 128 fF. Out of 326 CD capacitors, we identify 191 which cause $max(|t|) > 4.5$ while allowing the circuit to compute correctly. Based on the minimal capacitance value for which the capacitor causes $max(|t|) > 4.5$, henceforth called *critical value*, we further separate these capacitors in three groups:

- 18 high risk (H) capacitors, with critical value $C_H \leq 1$ fF.
- 88 medium risk (M) capacitors, with critical value $1\,\text{fF} \leq C_M \leq 4\,\text{fF}$
- 85 low risk (L) capacitors, with critical value $C_L > 4$ fF.

We chose the delimiting values of 1 fF and 4 fF for demonstrative purposes for the given circuit. Nevertheless, they are meaningful values for they are of the same order of magnitude as the pin capacitances of the standard-cells we use. Fig. 9 shows the magnitude of the $max(|t|)$ in the function of the H, M, L groups of CD capacitances. It is important to notice that we deem the capacitors with lower critical value to bear more risk, despite leading to lower $max(|t|)$ scores before causing the circuit to malfunction. Most of these 191 CD capacitors cause the circuit to malfunction when they reach between 17.5 fF and 23.7 fF in value or more. We discuss this ranking in more detail in Sec. IV.

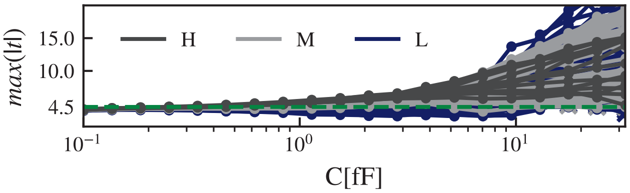

Fig. 9: Magnitude of the $max(|t|)$ for 2^{13} traces; • denotes correct and × failed computation of TI AND2.

IV. Discussion

In this section we discuss our findings in detail. We argue what may potentially be the practical scenarios in which these phenomena could demonstrate themselves. Lastly, we propose further research steps.

A. Impacts of PDN

Firstly, we show that the finite current capacity of the power supply Fig. 4 may increase the levels of leakage even when everything else in the circuit is ideal. Contrary to the suspicions of [5], our experiments do not show the resistive IR drop as one of the leading source of SCA leakage. As shown in Fig. 5a, some PDN resistors may increase the $max(|t|)$ score. Yet even for relatively large resistors said increase is slight. More importantly the resistance increase barely causes the $max(|t|)$ trend to deviate from the referent case. Our experiments further show that decoupling capacitances, often used in PDN modeling, lead to the decrease of the $max(|t|)$ score, as shown in Fig. 5b. Reactive impedances $Z_C = \frac{1}{j\omega C}$ form highly conductive paths for the data dependent high-frequency components of the supply current without upsetting the power supply voltages. However, our experiments show that parasitic PDN inductors are the more likely culprit, as shown in Fig. 5c. The reactive impedances $Z_L = j\omega L$ cause voltage drop predominated by the data dependent high-frequency components of the supply current drawn by the target circuit. The extent to which such a data-dependent upset of the power supply nodes can break the independence assumption is shown in Fig. 6.

In practice ground bounce and supply sag are prominent issues at the packaging level, over long bonding wires, as they may cause the upset of supply nodes that makes the entire circuit malfunction. Their effects are certain to be smaller within the logic core. Still, our experiments show that parasitic inductors may break the independence assumption at values $L \leq 1$ nH, two orders of magnitude before causing the circuit to malfunction at $L \geq 100$ nH. This margin is design specific and likely not to be this dramatic, yet as long as the performance constraints are met parasitic inductors with significant impact on the SCA security may be left in the design. Moreover, decoupling capacitors are used to battle the effects of the ground bounce and supply sag. This effect is shown in Fig. 5b. Nevertheless, decoupling capacitors are not free, they require area and design effort. Larger area means increased the PDN parasitics, as the rails have to be longer. Hence, they are likely to be sized up to the extent that satisfies the performance margin, not the SCA one.

B. Impacts of Coupling Capacitances

Even though theoretical leakage models completely overlook the influence of the PDN, SD and CD coupling capacitances are widely accepted as the possible source of leakage. Our experiments support these suspicions and quantify them further. We confirm that SR and SD capacitances lead to no significant leakage regardless of their size. Interestingly, CR capacitances yield no significant leakage either, although they span from one share to another. Therefore, CD capacitances remain as the only potential sources of leakage. This is clearly shown in the fundamental example of the 1-bit XOR2, $N = 2$ gate in Fig. 8. Only the capacitance C_{z_1,z_2} directly coupling the shares of the partitioning target causes $max(|t|) > 4.5$, indicating significant leakage.

In the case of TI AND2, $N = 3$ however determining which capacitances may cause information leakage is not as trivial. A total of 191 out of all 326 CD capacitances qualifies as the potential source of leakage. We rank them to emphasize the following. The higher $max(|t|)$ scores at the right-hand end of the Fig. 9 are a consequence of larger capacitances in the range between 17.5 fF and 23.7 fF, or more. While their impact is indeed higher, we expect fewer of such large parasitic elements to occur in a real design. On the other hand, capacitors on the order of magnitude of 1 fF are highly likely to occur, especially assuming unconstrained placement and routing. Hence, we assess the potential risk from such capacitors to be higher.

C. Future Work

The computational complexity of SPICE simulators prevents our model to be applied on larger designs. It is worth investigating whether other EDA tools could be used for such analysis. For example, supply current waveforms obtained using data-driven digital simulators such as PrimePower from Synopsys could be back-annotated into the analog PDN model.

V. Conclusions

We propose a model for co-simulation of the analog PDN along with the digital logic core dedicated to SCA. Using said model we conduct an exploratory study to quantify the impact of different layout parasitics to SCA security. We craft experiments in a way that allows us to diagnose all possible sources of leakage from within the logic core. Selecting small, fundamental circuits allows us to perform all experiments using SPICE. Furthermore, we are the first to study the impact of PDN to the security of the core logic. As all of these issues arise in the back-end stages, they are left out of the mathematical models although they may account for the significant information leakage. We provide the first detailed insights into the matter.

Acknowledgment

This work was supported in part by the Research Council KU Leuven C1 on Security and Privacy for Cyber-Physical Systems and the Internet of Things with contract number C16/15/058 and by the European Commission through the Horizon 2020 research and innovation programme Cathedral ERC Advanced Grant 695305.

References

[1] P. C. Kocher, J. Jaffe, and B. Jun, “Differential power analysis,” in *Advances in Cryptology - CRYPTO ’99*, ser. LNCS, M. J. Wiener, Ed., vol. 1666. Springer, 1999, pp. 388–397.

[2] L. Goubin and J. Patarin, “Des and differential power analysis the “duplication” method,” in *CHES*, Ç. K. Koç and C. Paar, Eds. Berlin, Heidelberg: Springer, 1999, pp. 158–172.

[3] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi, “Towards sound approaches to counteract power-analysis attacks,” in *Advances in Cryptology — CRYPTO’ 99*, M. Wiener, Ed. Springer, 1999, pp. 398–412.

[4] R. Jakushokas, M. Popovich, A. V. Mezhiba, S. Kse, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors*, 2nd ed. Springer Publishing Company, Incorporated, 2010.

[5] T. D. Cnudde, B. Bilgin, B. Gierlichs, V. Nikov, S. Nikova, and V. Rijmen, “Does coupling affect the security of masked implementations?” in *COSADE 2017*, ser. LNCS, S. Guilley, Ed., vol. 10348. Springer, 2017, pp. 1–18.

[6] J. Cooper, E. DeMulder, G. Goodwill, J. Jaffe, G. Kenworthy, and P. Rohatgi, “Test Vector Leakage Assessment (TVLA) methodology in practice,” International Cryptographic Module Conference, 2013.

[7] G. O. Dyrkolbotn, K. Wold, and E. Snekkenes, “Security implications of crosstalk in switching cmos gates,” in *Information Security*, M. Burmester, G. Tsudik, S. Magliveras, and I. Ilić, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 269–275.

[8] L. Zussa, I. Exurville, J.-M. Dutertre, J.-B. Rigaud, B. Robisson, A. Tria, and J. Clédière, “Evidence of an information leakage between logically independent blocks,” in *Proceedings of the Second Workshop on Cryptography and Security in Computing Systems*, ser. CS2 ’15, 2015, pp. 25:25–25:30.

[9] K. Tiri and I. Verbauwhede, “Simulation models for side-channel information leaks,” in *Design Automation Conference - DAC 2005*, W. H. J. Jr., G. Martin, and A. B. Kahng, Eds. ACM, 2005, pp. 228–233.

[10] D. Šijačić, J. Balasch, B. Yang, S. Ghosh, and I. Verbauwhede, “Towards efficient and automated side channel evaluations at design time,” in *PROOFS 2018*, ser. Kalpa Publications in Computing, L. Batina, U. Kühne, and N. Mentens, Eds., vol. 7. EasyChair, 2018, pp. 16–31.

[11] J. Knudsen, “Nangate 45nm open cell library,” in *12th Si2/OpenAccess+ Conference*, 2008.

[12] Wei Zhao and Yu Cao, “New generation of predictive technology model for sub-45nm design exploration,” in *7th International Symposium on Quality Electronic Design (ISQED’06)*, March 2006, pp. 6 pp.–590.

[13] D. Kamel, M. Renauld, D. Flandre, and F.-X. Standaert, “Understanding the limitations and improving the relevance of spice simulations in side-channel security evaluations,” *Journal of Cryptographic Engineering*, vol. 4, no. 3, pp. 187–195, Sep 2014.

[14] I. Levi, D. Bellizia, and F. Standaert, “Reducing a masked implementation’s effective security order with setup manipulations and an explanation based on externally-amplified couplings,” *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2019, no. 2, pp. 293–317, 2019.

[15] S. Nikova, C. Rechberger, and V. Rijmen, “Threshold implementations against side-channel attacks and glitches,” in *Information and Communications Security*, P. Ning, S. Qing, and N. Li, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 529–545.

[16] B. Bilgin, S. Nikova, V. Nikov, V. Rijmen, and G. Stütz, “Threshold implementations of all 3 ×3 and 4 ×4 s-boxes,” in *CHES 2012*, E. Prouff and P. Schaumont, Eds. Springer, 2012, pp. 76–91.

[17] A. Moradi, B. Richter, T. Schneider, and F.-X. Standaert, “Leakage detection with the x2-test,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 1, pp. 209–237, Feb. 2018.