VeriSide-II: Structure-Aware Power Modeling for Side-Channel Analysis at Register Transfer Level

VeriSide-II: Structure-Aware Power Modeling for Side-Channel Analysis at Register Transfer Level

Behnam Farnaghinejad, Annachiara Ruospo, Alessandro Savino, Stefano Di Carlo and Ernesto Sanchez
Department of Control and Computer Engineering, Politecnico di Torino, Torino, Italy
Email: {behnam.farnaghinejad, annachiara.ruospo, alessandro.savino, stefano.dicarlo, ernesto.sanchez}@polito.it

***Abstract*—Side-channel attacks exploit data-dependent variations in power consumption to recover secret information from hardware implementations. Detecting leakage at early design stages is critical, as post-synthesis or silicon-level fixes are costly and limited. At the RTL level, power is commonly approximated using switching-activity-based proxies, such as Hamming Weight (HW) or Hamming Distance (HD), which effectively model the activity factor in the dynamic power equation. Prior work has focused on scalable extraction and statistical analysis of such traces, without explicitly considering the device structural properties that may influence the effective capacitive term. This work investigates the integration of structural information at RTL to better approximate the capacitive component of dynamic power. VeriSide-II introduces a structure-aware proxy that augments switching activity with fan-out and combinational depth to form a more accurate capacitance model. Validation against post-synthesis gate-level power achieves approximately 93% correlation. As a second validation, a comparison against FPGA measurements shows up to a 10% improvement in diagonal-band correlation over activity-only formulations, while preserving leakage alignment under representative CPA experiments. These results experimentally demonstrate that incorporating structural awareness at RTL improves power-proxy fidelity and strengthens consistency between simulation-based and hardware-observed leakage, enabling more reliable early-stage side-channel assessment.**

***Index Terms*—Side-channel analysis, Hardware security, pre-silicon security, Power modeling, RTL simulation, RISC-V, VeriSide.**

I. Introduction

Side-Channel Analysis (SCA) remains a fundamental threat to hardware security, revealing secret-dependent information through physical leakage such as power consumption [1]–[3]. As hardware systems increase in scale and architectural complexity, security evaluation must move earlier in the design flow. Detecting leakage after synthesis or on silicon limits the ability to introduce architectural countermeasures without significant redesign cost. Early-stage evaluation at the Register-Transfer Level (RTL) is therefore essential for scalable and practical secure hardware development. Compared with silicon measurements and gate-level simulation, RTL switching analysis offers higher scalability but abstracts away physical load effects.

Recent RTL approaches have shown that switching-activity extraction enables efficient and scalable leakage estimation without requiring gate-level simulation [4]–[6]. However, they mainly treat each transition as a uniform dynamic-power contribution, ignoring structural properties that influence real power consumption, including signal fan-out and combinational propagation depth, which affect effective capacitive loading.

This work proposes VeriSide-II, an improved version of VeriSide [6], which augments RTL switching activity with lightweight structural descriptors extracted after HDL elaboration. Specifically, fan-out and combinational depth are combined into effective-capacitance–inspired weights to reduce the abstraction gap within the RTL domain, without relying on synthesis, layout, or post-synthesis parasitic information.

The evaluation combines a small synthetic RTL benchmark, RISC-V case studies, FPGA measurements, and representative Correlation Power Analysis (CPA) experiments on hardware and software AES implementations. The results show that the proposed structure-aware proxy improves alignment between RTL-derived traces and measured leakage behavior.

II. Background and Related Work

A. Leakage Assessment at RTL

Early-stage side-channel evaluation aims at identifying data-dependent leakage before synthesis and physical implementation, when design changes remain cost-effective. Several works have proposed leakage assessment frameworks at RTL based on switching activity extracted during simulation.

Prior approaches [9], [10] model dynamic power using switching statistics and apply statistical distinguishers, such as Kullback–Leibler Divergence (KL) [11], to identify potentially vulnerable modules. More recent studies have also investigated the applicability of classical statistical attacks, including Differential Power Analysis (DPA) and CPA, directly at the RTL level [4], as well as the effectiveness of deep learning techniques for leakage detection in this abstraction domain [5]. These works demonstrate that RTL analysis can serve as a fast pre-silicon screening stage prior to gate-level or silicon validation.

Complementary pre-silicon approaches improve physical fidelity through synthesized netlists, technology-aware power estimation, gate-level simulation, or transfer learning [12]–[15]. However, their reliance on synthesis, technology data, or gate-level traces reduces scalability compared with lightweight RTL-only assessment.

VeriSide [6] introduced a lightweight RTL switching-activity extraction methodology built on a modified Verilator backend [16]. Instead of generating large waveform files (e.g., Value Change Dump (VCD)), switching activity is accumulated cycle-by-cycle into a compact *SIDE* trace format, i.e., the per-cycle switching-activity format used by VeriSide, enabling scalable

analysis of large designs. This approach addresses the resource-intensive trace extraction problem common in earlier RTL frameworks and enables large-core evaluation.

The present work builds on that foundation and investigates whether incorporating structural awareness at RTL can improve alignment with lower-level or physical measurements.

B. Power Models and Statistical Attacks

Dynamic CMOS power is commonly approximated as

$$P_{dyn} = \alpha C V^2 f, \tag{1}$$

where α denotes the switching activity factor [1], [17]. In real CMOS circuits, the effective capacitive load C is not constant; it depends on gate input capacitances, interconnect capacitance, fan-out, transistor sizing, logic depth, and layout-dependent parasitics [19]. Consequently, two signals exhibiting identical switching activity may dissipate different dynamic power depending on their structural and physical context within the netlist. Classical leakage models approximate data-dependent power using abstractions such as Hamming Weight (HW) and Hamming Distance (HD). These models underpin standard non-profiled attacks, including DPA [2] and CPA [3], as well as more advanced techniques such as template attacks [18] and machine-learning-based approaches [20].

C. Gap and Positioning of This Work

Existing RTL frameworks mainly approximate the activity factor α while assuming uniform load across signals. This work investigates whether lightweight structural information, such as fan-out and combinational depth, can improve the correlation of RTL-derived proxies with gate-level power estimates and FPGA measurements.

III. Methodology

This section presents the structure-aware RTL proxy and the extraction of fan-out and combinational-depth metrics.

A. Implementation and Evaluation Flow

The complete structure-aware assessment flow is implemented within the open-source VeriSide-II framework.[1]

The overall methodology consists of four stages:

1) **Structural Graph Generation:** The RTL design is elaborated using VeriSide to generate an *elaborated structural representation in XML format*. This step performs structural analysis only and does not execute a simulation.
2) **Structural Metric Computation:** The generated XML representation is parsed to construct a directed signal-dependency graph. From this graph, fan-out and combinational depth are computed for each signal.
3) **Structure-Aware Simulation:** VeriSide-II is executed with the computed structural weights attached to the monitored signals. The selected weighting configuration (baseline, fan-out (FO), combinational-depth (D_{comb}), or effective-capacitance (C_{eff}) weighting) is applied during activity accumulation.
4) **Switching-Activity Trace Generation:** Cycle-level switching activity is generated in compact SIDE format, producing either activity-only or structure-aware power proxies.

[1] https://github.com/smilies-polito/VeriSide.

B. Baseline Switching-Activity Proxy

Let $\mathcal{S}$ denote the set of internal signals monitored during RTL simulation. For each signal $s \in \mathcal{S}$ and clock cycle t, a switching-activity measure $A_s(t)$ is extracted using the VeriSide framework [6]. Depending on the configuration, $A_s(t)$ may represent HW, HD, or a bit-level switching/counting indicator selected by the VeriSide configuration. The baseline activity-only proxy is defined as

$$P_{\text{base}}(t) = \sum_{s \in \mathcal{S}} A_s(t), \tag{2}$$

which implicitly assumes a uniform effective capacitance across all signals. This formulation captures the activity factor α in (1), but neglects structural differences that influence the capacitive term C.

C. Structural Metrics Extraction

To approximate the capacitive component C in (1), structural properties are extracted statically from the elaborated RTL structural representation.

In this work, the elaborated RTL representation refers to the structural form obtained after HDL elaboration, where module hierarchies are resolved, and continuous assignments and procedural logic are translated into explicit signal-level dependencies. We use Verilator as a front-end elaboration tool to generate an intermediate XML description of the flattened design, which exposes signal declarations and assignment relations in a tool-independent structural format. This XML representation is used solely to extract structural connectivity and does not require synthesis or technology libraries.

From this representation, a directed signal-dependency graph $G = (\mathcal{S}, \mathcal{E})$ is constructed, where nodes correspond to signals and edges represent combinational data flow.

For each signal $s \in \mathcal{S}$, two structural metrics are computed:

- **Fan-out ($FO(s)$):** the number of outgoing edges from node s in G, approximating the load induced by downstream gates.
- **Combinational depth ($D_{\text{comb}}(s)$):** the maximum length of combinational paths from s to the next sequential boundary, capturing the structural propagation extent of the signal.

D. Structure-Aware Power Proxies

Using the extracted structural metrics, weighted power proxies are constructed to approximate the missing capacitive factor.

a) Fan-out-weighted proxy:

$$P_{\text{FO}}(t) = \sum_{s \in \mathcal{S}} A_s(t)\big(1 + FO(s)\big). \tag{3}$$

b) Depth-weighted proxy:

$$P_{\text{D}}(t) = \sum_{s \in \mathcal{S}} A_s(t)\big(1 + D_{\text{comb}}(s)\big). \tag{4}$$

c) Effective-capacitance proxy: To jointly approximate load and structural propagation effects, an effective structural weight is defined as

$$C_{\text{eff}}(s) = \big(1 + FO(s)\big)\big(1 + D_{\text{comb}}(s)\big). \tag{5}$$

The additive constant ensures non-zero baseline contribution for signals with zero fan-out or zero combinational depth.

The corresponding structure-aware proxy becomes

$$P_{\text{Ceff}}(t) = \sum_{s \in \mathcal{S}} A_s(t)\, C_{\text{eff}}(s). \tag{6}$$

The multiplicative formulation reflects the intuition that both load (fan-out) and structural propagation depth contribute to effective switching energy. While this formulation does not aim to reproduce physical parasitics precisely, it provides a lightweight approximation of the capacitive term in the dynamic power model at RTL.

E. Optional Signal-Class Weighting

In addition to structural weighting, signals may be grouped into structural or functional classes, such as clocks, resets, control signals, and data signals. Different signal categories may contribute differently to observable switching activity depending on their role within the design.

A generic class-weighted formulation is defined as

$$P_{\text{class}}(t) = \sum_{s \in \mathcal{S}} A_s(t)\, w_{\text{class}}(s)\, W(s), \tag{7}$$

where $w_{\text{class}}(s)$ is a class-dependent coefficient assigned according to the signal category, and $W(s)$ denotes a structural weight, which may correspond to the baseline ($W(s) = 1$), fan-out weighting, depth weighting, or the effective-capacitance term $C_{\text{eff}}(s)$.

This formulation enables the exploration of whether combining structural awareness with coarse-grained functional grouping can further refine power approximation at RTL, while maintaining scalability.

IV. Experimental Results

This section evaluates the proposed structure-aware RTL proxies from two complementary perspectives: (i) alignment with post-synthesis gate-level power estimates, and (ii) alignment with FPGA measurements and statistical side-channel behavior.

A. Gate-Level Alignment

To quantify how well each RTL proxy approximates lower-level power behavior, Pearson correlation is computed between RTL-derived traces and post-synthesis gate-level power estimates.

As a controlled validation before full-processor experiments, we first use a small synthetic RTL benchmark containing combinational paths with different fan-out and depth. For this benchmark, 1,000 random input vectors are applied. For each clock cycle t, correlation is computed across the population of inputs between (i) the RTL-derived proxy values and (ii) the corresponding gate-level dynamic power estimates. This produces a cycle-level correlation profile that reflects how accurately the RTL proxy captures input-dependent power variation.

TABLE I: Gate-level correlation of RTL proxy variants and improvement over the baseline configuration.

Variant	Correlation	Impr. vs baseline (%)
$Class + C_{\text{eff}}$	0.9312	2.66
C_{eff}	0.9310	2.64
$Class + FO$	0.9120	0.55
FO	0.9120	0.55
$Class$	0.9081	0.12
Baseline (activity-only)	0.9071	0.00
$Class + D_{\text{comb}}$	0.8895	-1.93
D_{comb}	0.8895	-1.93

Table I summarizes peak correlation values and relative improvement over the baseline configuration. Across all cycles, the effective-capacitance proxy achieves approximately 93% correlation with gate-level power, corresponding to a 2.6% improvement over the activity-only formulation. Depth-only weighting reduces correlation, indicating that structural propagation alone does not sufficiently approximate effective capacitive load.

B. RTL–FPGA Temporal Alignment

To evaluate consistency between simulation-based proxies and measured hardware behavior, cross-correlation is computed between RTL traces and FPGA power traces.

Figure 1 presents the correlation heatmap for the best-performing structure-aware configuration on CVA6 [7], a 64-bit application-class RISC-V processor integrating a hardware AES accelerator. Under correct temporal correspondence, each RTL simulation cycle is expected to align with a specific FPGA time window after binning. Accordingly, meaningful similarity should concentrate within a narrow diagonal band of the heatmap, while non-aligned similarity appears outside this band.

For each proxy, we compute three summary metrics from the correlation heatmap: the diagonal-band score, defined as the average maximum correlation within a tolerance band around the main diagonal; the off-band score, defined as the average correlation outside this band; and their difference, used as the alignment contrast.

Table II reports the relative variation of Diag, Off, and Contrast with respect to the baseline activity-only proxy. The effective-capacitance configuration achieves the largest diagonal-band improvement (+10%) while maintaining stable off-band values, resulting in the highest alignment contrast. In contrast, fan-out-only weighting increases off-band correlation substantially, indicating amplified non-aligned activity, whereas depth-only weighting suppresses both aligned and background components.

C. Statistical Leakage Consistency

To verify that structure-aware proxies preserve meaningful data-dependent leakage behavior, representative CPA experi-

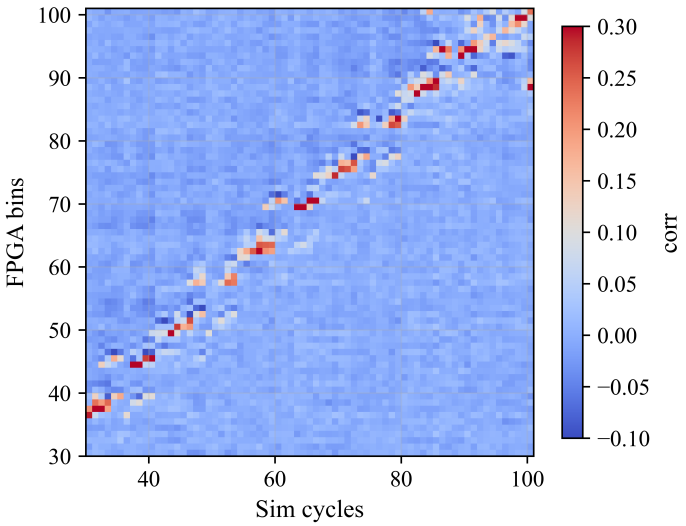

Fig. 1: Correlation heatmap between structure-aware RTL proxy and FPGA measurements on CVA6 (hardware AES implementation).

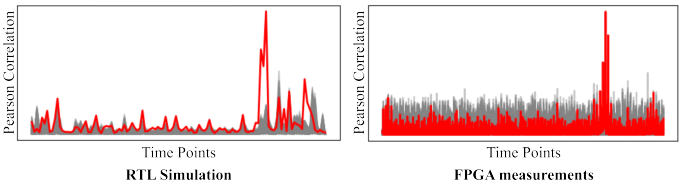

Fig. 2: CPA results for CVA6 core running hardware AES. Left: RTL-derived structure-aware proxy. Right: FPGA measurements.

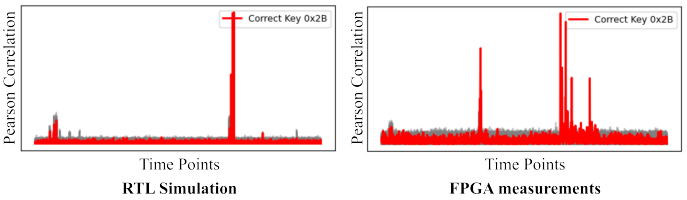

Fig. 3: CPA results for Ibex core running software AES. Left: RTL-derived structure-aware proxy. Right: FPGA measurements.

TABLE II: Relative variation of diagonal-band correlation, off-band correlation, and alignment contrast with respect to the baseline (activity-only) configuration.

Variant	ΔDiag (%)	ΔOff (%)	ΔContrast (%)
C_{eff}	+10.0	-0.5	+10.5
Class+C_{eff}	+9.9	-0.5	+10.4
D_{comb}	-5.8	-7.4	+1.6
Class+D_{comb}	-6.0	-7.2	+1.2
Class	-0.3	-1.3	+1.0
Baseline (activity-only)	–	–	–
Class+FO	-1.3	+19.9	-21.2
FO	-1.2	+20.9	-22.1

ments are conducted on both FPGA and RTL-derived traces.

For CPA, the same hypothetical AES intermediate is used with the same leakage model for RTL-derived and FPGA traces, and key hypotheses are ranked using Pearson correlation.

a) CVA6 (Hardware AES): For CVA6, the attack targets the input of the last-round S-box in the hardware AES accelerator. Figure 2 shows that the correct key byte produces the dominant correlation peak in both FPGA and RTL-derived traces.

b) Ibex (Software AES): For Ibex [8], AES is executed in software, and the attack targets the output of the first-round S-box. Figure 3 shows consistent key ranking between FPGA measurements and RTL-derived traces.

Although only one key byte is illustrated, similar behavior is observed across all 16 bytes of the 128-bit key. This confirms that structure-aware RTL proxies preserve exploitable leakage characteristics while improving cross-level alignment.

V. Discussion and Conclusion

The experimental results show that incorporating structural information at RTL improves alignment with both gate-level power estimates and FPGA measurements. In particular, the proposed effective-capacitance formulation increases diagonal-band correlation while maintaining stable off-band values, indicating improved temporal alignment without amplifying spurious similarity.

This paper introduced VeriSide-II, a structure-aware RTL power modeling methodology that augments switching-activity–based leakage proxies with lightweight structural descriptors, namely fan-out and combinational depth. These descriptors approximate relative load effects directly at RTL, without requiring synthesis, technology libraries, layout parasitics, or physical implementation data.

The results further show that structural weighting improves cross-level alignment without distorting statistical leakage behavior. CPA experiments on both hardware and software AES implementations produce consistent key-dependent peaks across RTL-derived and FPGA traces, indicating that the proposed proxy preserves attack-relevant leakage characteristics while improving physical consistency.

Because the approach operates entirely at RTL, it preserves the scalability advantages of activity-based screening. As trace extraction already dominates the runtime in the original VeriSide framework, enabling the additional structural models in VeriSide-II increases execution time by only a few seconds in the CVA6 experiments, while remaining significantly lighter than VCD-based or gate-level approaches.

VeriSide-II is intended as an early screening step rather than a replacement for gate-level, post-layout, or silicon validation. In practice, it can be used to rank RTL variants, identify leakage-relevant time windows, and decide where more expensive lower-level analysis or FPGA measurements should be focused.

Future work will compare activity-only and structure-aware proxies through systematic key-recovery experiments, extend the model with additional structural indicators, and evaluate synthesized ASIC netlists.

REFERENCES

[1] S. Mangard, E. Oswald, and T. Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. New York, NY, USA: Springer, 2007, doi: https://doi.org/10.1007/978-0-387-38162-6.

[2] P. Kocher, J. Jaffe, and B. Jun, "Differential Power Analysis," in Advances in Cryptology—CRYPTO'99, Lecture Notes in Computer Science, vol. 1666, Springer, Berlin, Heidelberg, 1999. doi: 10.1007/3-540-48405-1_25.

[3] E. Brier, C. Clavier, and F. Olivier, "Correlation Power Analysis with a Leakage Model," in Cryptographic Hardware and Embedded Systems—CHES 2004, Lecture Notes in Computer Science, vol. 3156, Springer, Berlin, Heidelberg, 2004. doi: 10.1007/978-3-540-28632-5_2.

[4] B. Farnaghinejad et al., "Power Side-Channel Vulnerabilities of a RISC-V Cryptography Accelerator Integrated into CVA6 via Core-V eXtension Interface (CV-X-IF)," 2025 IEEE International Test Conference (ITC), San Diego, CA, USA, 2025, pp. 233-242, doi: https://doi.org/10.1109/ITC58126.2025.00030.

[5] B. Farnaghinejad, A. Ruospo, A. Savino, S. Di Carlo, and E. Sanchez, "Deep Learning-Based Power Side-Channel Evaluation from RTL Switching-Activity Proxies," 2026 IEEE 27th Latin American Test Symposium (LATS), Florianopolis, Brazil, 2026, pp. 1–6, doi: https://doi.org/10.1109/LATS70329.2026.11480302.

[6] B. Farnaghinejad, A. Porsia, A. Ruospo, A. Savino, S. Di Carlo, and E. Sanchez, "Late Contribution: VeriSide: A Modified Verilator for Leakage Assessment at the RTL Level," in 2025 IEEE 26th Latin American Test Symposium (LATS), Mar. 2025, pp. 1–2. doi: https://doi.org/10.1109/LATS65346.2025.10963943.

[7] F. Zaruba and L. Benini, "The Cost of Application-Class Processing: Energy and Performance Analysis of a Linux-Ready 1.7-GHz 64-Bit RISC-V Core in 22-nm FDSOI Technology," IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 27, no. 11, pp. 2629–2640, Nov. 2019, doi: https://doi.org/10.1109/TVLSI.2019.2926114.

[8] "Ibex: An embedded 32 bit RISC-V CPU core — Ibex Documentation 0.1.dev50+g8d171a16d documentation," Readthedocs.io, 2017. https://ibex-core.readthedocs.io/en/latest/ (accessed Mar. 03, 2026).

[9] Miao Tony He, J. Park, Adib Nahiyan, A. Vassilev, Y. Jin, and M. Tehranipoor, "RTL-PSC: Automated Power Side-Channel Leakage Assessment at Register-Transfer Level," arXiv (Cornell University), Apr. 2019, doi: https://doi.org/10.1109/vts.2019.8758600.

[10] Nitin Pundir, J. Park, Farimah Farahmandi, and M. Tehranipoor, "Power Side-Channel Leakage Assessment Framework at Register-Transfer Level," IEEE Transactions on Very Large Scale Integration Systems, vol. 30, no. 9, pp. 1207–1218, Sep. 2022, doi: https://doi.org/10.1109/tvlsi.2022.3175067.

[11] J. Park and A. Tyagi, "Security Metrics for Power Based SCA Resistant Hardware Implementation," pp. 541–546, Jan. 2016, doi: https://doi.org/10.1109/vlsid.2016.43.

[12] A. Nešković, A. Kabour, M. Eltantawi, M. Berekovic, R. Meyer, and S. Mulhem, "SCA-WAL: An Open-Source Framework for Power Side Channel Assessment at RTL and Netlist Level," IEEE Embedded Systems Letters, pp. 1–1, 2025, doi: https://doi.org/10.1109/les.2025.3644491.

[13] D. Shanmugam and P. Schaumont, "Improving Side-channel Leakage Assessment Using Pre-silicon Leakage Models," in Constructive Side-Channel Analysis and Secure Design—COSADE 2023, Lecture Notes in Computer Science, vol. 13979, Springer, Cham, 2023. doi: 10.1007/978-3-031-29497-6_6.

[14] Samadi Bokharaie, V., Jahanian, A. "Power side-channel leakage assessment and locating the exact sources of leakage at the early stages of ASIC design process,". J Supercomput 78, 2219–2244 (2022). https://doi.org/10.1007/s11227-021-03927-w.

[15] A. Srivastava et al., "SCAR: Power Side-Channel Analysis at RTL Level," IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 32, no. 6, pp. 1110–1123, Apr. 2024, doi: https://doi.org/10.1109/tvlsi.2024.3390601.

[16] W. Snyder, P. Wasson, D. Galbi, G. Lore, et al., *Verilator*, Veripool, software. [Online]. Available: https://verilator.org; source code: https://github.com/verilator/verilator.

[17] M. Pedram, "Power minimization in IC design," ACM Transactions on Design Automation of Electronic Systems, vol. 1, no. 1, pp. 3–56, Jan. 1996, doi: https://doi.org/10.1145/225871.225877.

[18] S. Chari, J. R. Rao, and P. Rohatgi, "Template Attacks," in Cryptographic Hardware and Embedded Systems—CHES 2002, Springer, 2003. doi: 10.1007/3-540-36400-5_3.

[19] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2003.

[20] H. Maghrebi, T. Portigliatti, and E. Prouff, "Breaking Cryptographic Implementations Using Deep Learning Techniques," in *Security, Privacy, and Applied Cryptography Engineering*, C. Carlet, M. Hasan, and V. Saraswat, Eds. Cham, Switzerland: Springer, 2016, pp. 3–26, doi: 10.1007/978-3-319-49445-6_1.