Your Rails Cannot Hide From Localized EM: How Dual-Rail Logic Fails on FPGAs

Vincent Immler[1], Robert Specht[1], and Florian Unterstein[1]

Fraunhofer Institute for Applied and Integrated Security (AISEC), Germany
forename.surname@aisec.fraunhofer.de

Abstract. Protecting cryptographic implementations against side-channel attacks is a must to prevent leakage of processed secrets. As a cell-level countermeasure, so called DPA-resistant logic styles have been proposed to prevent a data-dependent power consumption.
As most of the DPA-resistant logic is based on dual-rails, properly implementing them is a challenging task on FPGAs which is due to their fixed architecture and missing freedom in the design tools.
While previous works show a significant security gain when using such logic on FPGAs, we demonstrate this only holds for power-analysis. In contrast, our attack using high-resolution electromagnetic analysis is able to exploit local characteristics of the placement and routing such that only a marginal security gain remains, therefore creating a severe threat.
To further analyze the properties of both attack and implementation, we develop a custom placer to improve the default placement of the analyzed AES S-box. Different cost functions for the placement are tested and evaluated w.r.t. the resulting side-channel resistance on a Spartan-6 FPGA. As a result, we are able to more than double the resistance of the design compared to cases not benefiting from the custom placement.

1 Introduction

Physical attacks based on power analysis, called DPA [20], have been subject to extensive research and initiated the development of DPA countermeasures at different levels of abstraction. Some introduce noise, e.g., [12, 23] or randomize the order of operations, i.e., shuffling, e.g., [16, 23]. More application-specific attempts to increase the resistance are done by manipulating the underlying cryptographic algorithm to randomize its intermediate values, i.e., masking at the algorithmic level, e.g., [29, 30]. Others, so called "hiding" countermeasures, try to solve the problem by avoiding data-dependencies in the power consumption. These countermeasures at the cell-level, called DPA-resistant logic styles, ideally remove the data-dependent power consumption and thereby equalize it.

When considering the various proposals in this domain [28, 22, 33, 3, 26, 44, 19, 13, 15, 43], one identifies that most of them are based on dual-rail precharge (DRP) logic or duplication schemes. Both no longer represent a bit as a single value but instead as complementary rails of (`true`,`false`) $= (t, f)$, such that regardless of the operation, each bit-flip is compensated by an inverse bit-flip.

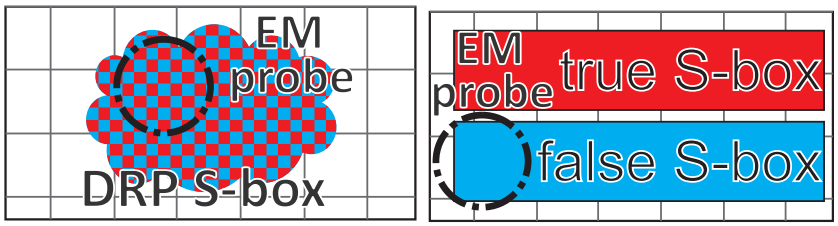

(a) S-box implementing the inseparable function $f_{DRP}(t, f)$ (DRP logic).

(b) S-box implementing $f_{DUP}(t)$ and $f^{-1}_{DUP}(f)$ separately (duplication scheme).

Fig. 1: Resulting FPGA floor plan to illustrate different dual-rail styles. The on-chip positioning of a probe used for the EM analysis is indicated by a circle.

However, both approaches are fundamentally different as explained later on in more detail: Using a simplified view, DRP styles can be considered as a function $f_{DRP}(t, f)$ and duplication schemes as a compound function of $f_{DUP}(t)$ and $f^{-1}_{DUP}(f)$ which leads to different implementations as sketched in Figure 1.

To properly implement either one, several design flaws must be avoided such as *glitches* [24]. Another, the *early propagation* (EP) effect [37] is typically prevented by a synchronization mechanism, e.g., an enable signal. Moreover, to achieve equal power consumption between the dual-rails, it is necessary to minimize their *routing imbalances* [39], as they result in different capacitive loads when switching which can be exploited by a DPA attack. Therefore, some routing techniques have been proposed such as [11, 41, 26] to diminish the load imbalances in either the ASIC or FPGA design process. Most dual-rail mitigation techniques work reasonably-well when assuming a power-based side-channel, e.g., by measuring the voltage over a shunt resistor (including parasitics), thereby treating the leakage of a device as a whole. The local placement and routing imbalances are therefore not sufficiently considered as part of a design or evaluation process. From a practical point of view, power-analysis also requires the PCB to be modified in most cases. Moreover, decoupling capacitances tend to be increasingly more integrated recently in the chip itself which makes the use of the power-based side-channel more difficult.

An often preferred alternative are side-channels based on Electro-Magnetic (EM) emanation. Various publications have shown different options on how to measure the emanation. Mainly two approaches exist, the off-chip measurement [32, 7], i.e., the probe is positioned slightly above the chip package and on-chip (or on-surface) measurement [31, 17, 35, 42], i.e., the chip is partially depackaged to position the probe directly on top of the die's surface. By positioning a suitable EM probe in proximity to the area of interest, spatial information of the implementation can be explored [31, 17]. This is also known as localized EM and is due to previous results a promising candidate to measure dual-rails independently from each other, thereby possibly bypassing this countermeasure.

Our Contribution A natural question that arises is by how much better on-chip EM attacks perform when compared to power measurements, as local placement and routing characteristics could possibly be more easily extracted using a localized EM attack. To start answering this question, we first survey the existing logic styles for FPGA platforms and argue that previous security evaluations did not fully assess the properties of the underlying designs, i.e., the density of the placement and local routing imbalances.

Afterwards, we practically investigate if the available DRP logics can still be assumed secure when subject to a localized EM attack. As a result, we show that only a barely noticable security gain of any DRP logic remains when compared to a SingleRail implementation using the default placement of the Xilinx ISE tools. To the best of the authors' knowledge, we are the first to publicly perform such attacks on dual-rails using high-resolution equipment, i.e., with a probe diameter of 150 µm at a very low distance of ≤ 50 µm.[1]

To fairly compare the security of dual-rails using a power- and an EM-based analysis, we present a systematic evaluation methodology based on a correlation based leakage test which is complemented by an information theoretic approach. It is additionally supplemented by considering the Signal-to-Noise-Ratio (SNR).

As the next contribution we focus on the placement of the secure logic, more specifically its density and its possible influence on the resistance towards an EM-based analysis. As target design and platform, we selected an AES S-box to be realized on a Spartan-6 FPGA. Its local placement using ISE defaults is improved by means of a custom placer based on simulated annealing. A new and previous cost function are evaluated for the placer. Our experimental results show that increasing the density of the placement using our own cost function helps to reduce the amount of extractable leakage by an EM-based analysis.

State-of-the-Art Secure logic styles primarily follow two competing concepts on FPGAs: DRP logic and duplication schemes. DRP logic gates operate in two phases, i.e., precharge and evaluation, which are controlled by the clock signal as seen later on in Figure 6b. Proposed candidates include: WDDL [40], BCDL [28], DPLnoEE [3], and AWDDL [26], as listed in Table 1.

Unfortunately, none of the DRP proposals include a thorough analysis based on a localized EM attack to answer the fundamental question, if dual-rails could be measured separately, e.g., by measuring differing orientations of the emanated field of the rails, therefore possibly bypassing this countermeasure. Another issue is local placement and routing imbalances. Especially large nets with multiple sinks cause a "messy" routing with the following properties: not all dual-rails can be fully balanced due to the lack of precise timing information from the Xilinx tools as stated in [26], also one cannot assume that balanced dual-rails remain balanced across several devices using the same design due to device-specific variation as shown in [43], cross-coupling of lines adds another uncertainty that may lead to leakage and inter-dependency of lines within an FPGA [6, 9].

[1] We omitted results from probes with 100 µm and 250 µm due to similarity reasons. In contrast, a probe with 3 mm was almost equivalent to a power-based measurement.

Table 1: Survey on dual-rail logic styles for reconfigurable hardware.

Reference	Design Properties				Device Under Test		Evaluation Setup
	EP	Glitch	Route	Place	Platform	Target	
DRP Logic							
BCDL [28]	✓	✓	x	**?**	Stratix II	AES	Power
WDDL [33],[40]	x	✓	x	*	Stratix	DES	off-chip EM
DPLnoEE [3]	*	✓	x	**?**	Stratix I	AES	Power[1]
AWDDL [26]	✓	✓	*	**?**	Virtex 5	AES	Power
Duplication Schemes							
DWDDL [44]	x	✓	x	**?**	Spartan 3E	AES	Power
Part. SDDL [19]	✓	x	*	**?**	Spartan 3E	AES	Power
PA-DPL[13],[14]	✓	x	*	*	Virtex 5	AES	off-chip EM
[15]	✓	✓	x	**?**	Virtex 5	AES	off-chip EM
GliFreD [43]	✓	✓	*	**?**	Spartan 6	AES	Power

✓: addressed x: problematic *:partially considered **?**:not considered 1: EM on capacitor

Regardless of these obstacles, some work has been done in hardening DRP logic on FPGAs. In [34], different placement strategies are investigated. Each is based on constraints of the Quartus-II tools. Since the evaluation is done using a power-based DPA only, analyzing local effects of the various placements more closely has not been possible.

In another work [26], the authors investigate if rails can be balanced using a custom routing algorithm. Although the leakage is reduced by their router, the authors report that it cannot be completely avoided since the Xilinx tools only report worst-case values that differ from reality. Moreover, the placement was not considered at all. Since this is the foundation for an optimized routing, verifying its properties prior to the routing would have been necessary. Again, results are only based on a power-analysis. Table 1 summarizes additional DRP logic styles and indicates the strong need for an on-chip EM analysis.

In contrast to DRP logic, duplication schemes are typically realized as follows: For a given circuit, a complementary one is created which leads to a dual-copy of a fully placed-and-routed circuit which has been shown, can often be broken due to non-dual glitches in the original and dual part of the circuit [24, 43].

In terms of implementing them, they have the advantage of using duplicated routes that are shifted in horizontal or vertical direction [15, 43]. The balancing aspect is therefore derived from the fact that routes of equivalent shape using identical hardware resources are likely to yield balanced capacitive loads. However, the resulting distance between true/false is typically large, e.g., at least a tile and often more than that [44, 13]. Considering localized EM attacks this may be a significant issue. Moreover, the only known duplication scheme to address both glitches and early-propagation is GliFreD [43] which results in a massive Flip-Flop (FF) overhead compared to DRP logic styles. Furthermore, it has only been evaluated using power-analysis. Table 1 includes other candidates of dupli-

cation schemes and lists their conceptual weaknesses with respect to their design properties. Hence, they are not considered in depth as part of our work.

One minor exception of duplication schemes we would like to address is the work of [14] which analyzes various placement strategies of PA-DPL [13] by using on-chip EM measurements. However, their measurement setup is incomparable to ours (cf. Section 5.3). As an example, the diameter of their coil is by orders of magnitude larger (1 mm) than ours (150 µm). This may have prevented a more detailed analysis, since no differences for the tested placements were observed.

In the direction of EM-based analysis, we would like to refer to [31, 17, 18, 35] to illustrate the advancements over previous EM-based approaches. Aforementioned references indicate that localized EM attacks are more powerful than power measurements. However, they did not carry out a thorough comparison. Hence, we also investigate the practical limits of localized EM in terms of resolution vs. the given routing architecture and technology size of a Xilinx Spartan 6.

2 Dual-Rail Routing and Placement

Let us recall selected properties of dual-rail styles and how they relate to a localized EM attack. As outlined before, both DRP logic and duplications schemes create complementary rails to achieve a constant number of switches independent from the processed data, ideally resulting in equalized power consumption.

This appears as a valid approach when neglecting the design challenges to properly implement them. However, even under idealized assumptions, carrying out a localized EM attack could prove more resourceful than a power-based measurement, as a wise positioning of the probe may lead to an asymmetric view on the rails as illustrated in Figure 2a. As the signal strength picked up by the probe depends on its distance to the emanating source, it appears likely that due to an unequal signal strength of `true` and `false` that they no longer compensate each other. The resulting residue then reflects the properties of the stronger signal which exhibits the same behavior of a SingleRail implementation.

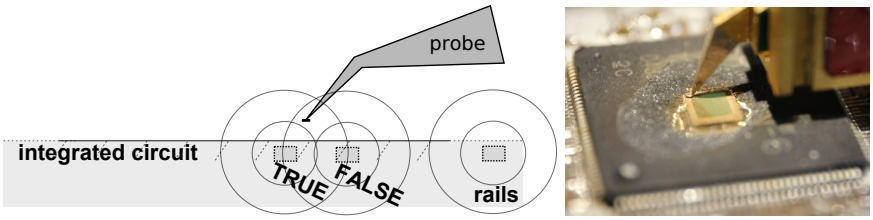

(a) Artistic illustration of a local EM attack. (b) Probe positioned on die.

Fig. 2: Introductory material on high-resolution, localized EM analysis.

Clearly, the success of this depends on the resolution of the attack relative to the density of the placement and routing of the logic. It is therefore not possible

to analyze this problem independently from placement and routing characteristics. Hence, they must be considered and optimized, too.

For duplication schemes, the situation is as sketched in Figure 1b. For a given S-box, a complementary copy is created resulting in symmetrically placed and routed logic. While achieving a high level of uniformity, at the same time the distance between true and false is typically large, often multiple tiles of an FPGA. Please also note that due to the divided approach of duplication schemes, both rails *and* the computing LUTs of the true and false paths are fully separated.

For DRP logic as depicted in Figure 1a, the situation is completely different. Since both true and false path must be jointly routed to each Look-Up-Table (LUT), each rail must be routed individually and cannot be copied. Ideally, depending on the quality of the placement and routing capabilities, one would be able to route a dual-rail much closer to each other when compared to duplication schemes. However, at the same time, where this is not possible, local non-uniformities (larger distances between still balanced dual-rails, cf. Figure 3b) or even imbalances would occur (rails with unequal capacitive loads, cf. Figure 3c).

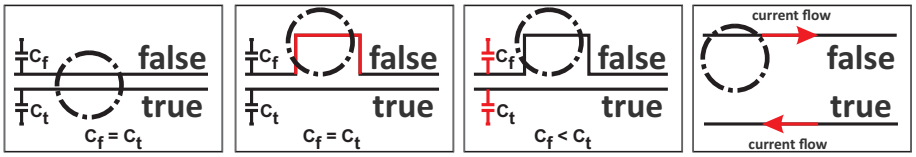

(a) (b) (c) (d)

Fig. 3: Different routing characteristics. (3a) "Ideal" dual-rail (3b) Non-uniformity (3c) Imbalance (3d) Large distance between rails and different orientation of the emanated field due to current flow.

Since computing true and false rails of DRP logic takes place within the same slice (or same LUT), we expect that the remaining non-uniformities and imbalances are much more difficult to exploit when compared to a duplication scheme where both computation and dual-rail routing are split. We therefore focus on the resistance of DRP styles and compare their effectiveness as a countermeasure under a power- and localized EM-attack setting. To do so, we first introduce the topic of placement and its optimizations, to also compare a low and high-density placement, as a higher density should help mitigate the aforementioned effects.

3 Placement on FPGAs

Commonly available FPGAs share similarities in their fabrics, i.e., the underlying structure of hardware resources. For Xilinx FPGAs, the reoccurring structure implementing the majority of logic is called a *tile*. Each tile typically comprises two *slices*, whereas each slice contains 4 LUTs, several multiplexors, and FFs. In between each tile and slice, different routing resources are available.

As a first step to implement the designated logic on FPGAs, its representation as a hardware description language is mapped onto the device using device-specific libraries. Subsequently, the logic must be placed such that the hardware resources are not exceeded. On a global level, partitioning the logic is often done using quadratic placement, especially on ASICs. On a local level, i.e., modules of reasonable size this is often done using simulated annealing [38]. In general, this is termed the "placement problem" [27] and known to be *np*-complete, i.e., placing logic within a certain rectangular area P based on some minimized cost function $\mathrm{C}(p)$ is only practically feasible using approximative approaches.

P is defined by its boundaries $x_{\text{high}}, x_{\text{low}}, y_{\text{high}}, y_{\text{low}}$. The list of gates G and nets V is a graph $G = (V, E)$. The cost function $\mathrm{C}(p) = C(V, E)$ represents the sum of the expected wirelength for each net. Determining the wirelength $\mathrm{WL}(e)$ can be done using different approaches, as discussed in Section 3.2. In addition to that, it is possible to assign weights $\mathrm{w}(e)$ to each net for, e.g., critical nets. The resulting cost function is then denoted as: $\mathrm{C}(p) = \sum_{e \in E} \mathrm{w}(e)\mathrm{WL}(e)$.

The placement problem can now be formalized as: given P, a list of gates and nets $G = (V, E) = (FV \cup MV, E)$ with FV as fixed gates and MV as movable gates, and a cost function $C(V, E)$. Determine (x_i, y_i) such that for each $v_i \in MV$: (i) it is placed within P and (ii) no pair of v_i, v_j overlaps with $\forall v_i, v_j$ (iii) $C(p)$ is minimal.

3.1 Simulated Annealing

Simulated annealing [2] resembles a cooling process to allow larger changes in the beginning as long as the temperature is high. While cooling down, the magnitude of changes becomes smaller with each iteration. Thereby, an approximative global optimum is found by avoiding local minima/maxima.

For the placement region P, an initial random placement p_0 is realized. Its quality is determined by the cost function $\mathrm{C}(p_0)$. For a given temperature T_0 in the beginning, the subsequent iterations start to move around logic gates. Each new placement is again evaluated by $\mathrm{C}(p)$. Degradations are only accepted with a probability of $\mathrm{e}^{-\frac{\mathrm{C}(p_{\text{new}}) - \mathrm{C}(p_{\text{old}})}{T}}$, i.e., the acceptance rate of optimizing towards the wrong direction decreases. This process continues until an exit criterion is fulfilled, e.g., a certain temperature, iteration count or quality of placement.

3.2 Cost Functions

As part of this work, we adapted the cost function of Versatile-Place-and-Route (VPR) [2] which is based on the Half-Perimeter-Wire-Length (HPWL). Moreover, we define our own cost function called Same-Slice-Same-Tile (SSST). Both are also illustrated in Figure 4 and explained hereafter.

HPWL The function $\mathrm{q}(n) \cdot \mathrm{HPWL}$ which we use is a modified version of the linear congestion function of [2]. The original equation is

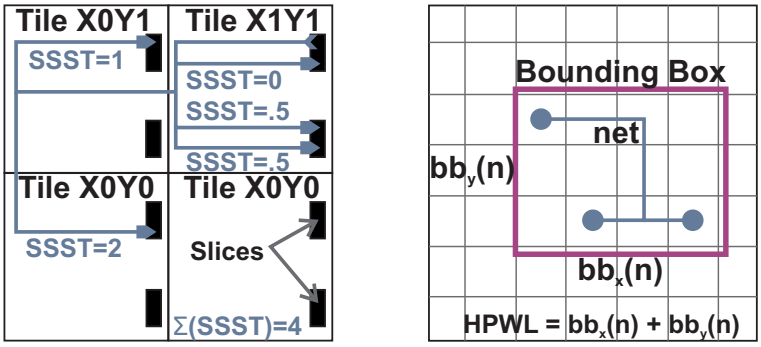

(a) Same Slice Same Tile. (b) Half-Perimeter-Wirelength.

Fig. 4: Illustration of different cost functions of the custom placer.

$$\mathrm{C_{lc}}(p) = \sum_{n=1}^{N_{\mathrm{nets}}} \mathrm{q}(n)\big[\frac{\mathrm{bb}_x(n)}{\mathrm{C}_{av,x}(n)} + \frac{\mathrm{bb}_y(n)}{\mathrm{C}_{av,y}(n)}\big] \tag{1}$$

whereas N_{nets} is the number of nets between the to-be-placed instances. $\mathrm{bb}_x(n)$ and $\mathrm{bb}_y(n)$ are the side lengths in x or y direction of the specific bounding box of a net n, as sketched in Figure 4b. A bounding box is the smallest rectangle that fits each instance of a net. Hence, $\mathrm{HPWL} = \mathrm{bb}_x(n) + \mathrm{bb}_y(n)$.

$\mathrm{q}(n)$ is a specific factor to balance nets with many pins. The respective values have been taken from [5]. $\mathrm{C}_{av,x}(n)$ and $\mathrm{C}_{av,y}(n)$ reflect the routing channel capacity to, e.g., make certain wires more expensive than others. However, since the Xilinx Spartan 6 is assumed to provide a symmetric channel/wire layout in arbitrary direction, we simplify the cost function to:

$$\mathrm{C_{HPWL}}(p) = \sum_{n=1}^{N_{\mathrm{nets}}} \mathrm{q}(n)\big[\mathrm{bb}_x(n) + \mathrm{bb}_y(n)\big] \tag{2}$$

SSST Since our goal is to not only make an optimized routing but also to create a placement of highest density, we define our own cost function

$$\mathrm{C_{SSST}}(p) = \sum_{n=1}^{N_{\mathrm{nets}}} \sum_{m=1}^{N_{\mathrm{sinks}}} \mathrm{P}(m) \tag{3}$$

whereas $\sum_{m=1}^{N_{\mathrm{sinks}}} \mathrm{P}(m)$ is the sum over all wirelengths of a net with

$$\mathrm{P}(m) = \begin{cases} 0, \text{ if} & \text{Source and sink are within the same } \textit{slice} \\ 0.5, \text{ if} & \text{Source and sink are within the same } \textit{tile} \\ d(\mathrm{Tile}(s), \mathrm{Tile(m)}), \text{ else} & \end{cases} \tag{4}$$

The function $d(.,.)$ represents the Manhattan distance[2] between the tile in which the source of the net is placed and the tile in which the sink is. Figure 4a illustrates the properties of the SSST-metric.

4 Custom Placer and Design Implementation

In the following subsection, we describe the implementation of our custom placer. Subsequently, we use this placer to work on the design presented in Section 4.2.

4.1 Custom Placer

Since our device-under-test (DUT) for the design is a Xilinx Spartan-6 FPGA, we could make use of the RapidSmith library [21]. The resulting workflow is outlined in Figure 5 and is based on the Xilinx Design Language (XDL). For our use case, we fully process the design (as depicted in Figure 6a) up to the `ncd` right before `bitgen`, using the standard Xilinx ISE tools.

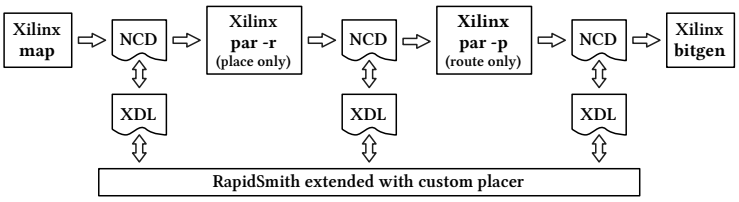

Fig. 5: Workflow using the RapidSmith library and the custom placer

The only modifications we made to the design are *i*) to put all elements requiring improved placement into a closed group which is area constrained, *ii*) to keep the input PIN positions of the LUTs locked, and *iii*) to put LUTs of each logic gate into the same slice using either the pair (A,B) or (C,D) of LUTs.

Once the design is processed by ISE, it is converted to its `XDL` representation and imported to RapidSmith. The S-box is detected by its hierarchy name and confined to the region P using the same boundaries as for the constraints of the ISE toolchain. Afterwards, the primitive sites within this region are identified to check if the designated logic could be placed using the given site (e.g., SLICEL).

To relocate the already placed logic, it is necessary to remove the nets and extract the logic from its given `XDL` hierarchy. A group of "relocatable" logic is created to allow their repositioning. Another abstract group is created to keep track of how they are interconnected, i.e., the nets. The initial placement of the ISE tools is considered as p_0, i.e., the start of the simulated annealing is well-defined and not a random placement. Subsequently, the annealing is carried out as described in Section 3.1 using the cost functions of Figure 4a and 4b.

[2] $\mathrm{d} = \mathrm{abs}(x_s - x_m) + \mathrm{abs}(y_s - y_m)$, i.e., the rectangular distance over the grid.

Once the annealing stops (after less than a minute), the logic is placed back on the primitive sites. Please note, since relocating the logic was performed by making "valid moves" only, there is no legalization step required (in contrast to quadratic placement). The thus placed logic is then interconnected using the routing capabilities of the ISE tools and the `bit` files are generated.

4.2 Design Implementation

For the analysis, we made an exemplary design which in addition to the control logic consists in an AES S-box [1]. We have taken the area-optimized S-box by Canright [4] which is a typical design for an S-box hardware implementation.

As logic styles, we selected the following: SingleRail, WDDL [40], DPL-noEE [3], and AWDDL [26], as they can all be realized using the same routing which leads to an unambiguous comparison.[3] For each style, we instantiated the S-box logic by 2-input gates. A block diagram of the design is shown in Figure 6a.

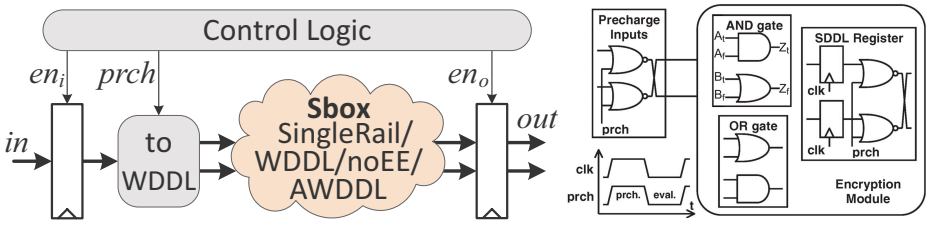

(a) The exemplary design block diagram. (b) Concept of WDDL ([40]).

Fig. 6: Basic properties of the implemented design.

Since we aim at evaluating only the leakage associated to the combinatorial circuit, we must exclude the leakage of the output register (see [25]). We therefore implement the design as follows: At a certain clock cycle, the control logic disables *prch* signal and the "to WDDL" conversion unit propagates the input to the S-box thereby initializing the evaluation phase. In the next half of the clock cycle the control logic enables *prch* signal and the precharge phase is started.

In a common DRP circuit en_o should be active at the start of precharge in order to store the output of the combinatorial circuit (here the AES S-box). Therefore, the control logic does not enable en_o signal and the register does not store the S-box output.[4] Over these two (evaluation and precharge) phases either power consumption or electro-magnetic emanation of the FPGA are measured.

To implement the logic and achieve the same routing for them, we use the following procedure. Using AWDDL as a start, we implement WDDL and DPL-noEE by changing only the LUT contents in the `XDL` file which is possible due to

[3] Our results can also be mapped onto BCDL [28] since it is similar to DPLnoEE.

[4] At a later point in time, en_o becomes active in order to check the correct functionality of the circuit. This is not covered by the recorded power and EM traces.

the 2-input AWDDL gates. For the SingleRail variant, we additionally disable the `FALSE` rail of WDDL and adjusted the "toWDDL" conversion accordingly.

Now, different sets of placements are created. Each comprises all four logics and uses the procedure to achieve the same routing within each set:

- **Set 1:** Default ISE placement using constraints
- **Set 2a:** Customized placement optimized towards HPWL
- **Set 2b:** Customized placement optimized towards SSST

The resulting placement metrics according to the cost functions are summarized in Table 2. For both Set 2a and 2b, one can see that an improvement over the ISE defaults is achieved. Each design of each set is then subject to the power and EM measurement using the same `bit` file. The measurement setups are described hereafter and preceed the practical investigations of Section 6.

Table 2: Results of the respective cost functions for different designs of the S-box.

Design Type	HPWL	SSST
Set 1 (default placement)	5860.0	3198.0
Set 2a (optimized towards HPWL)	3241.6	2308.0
Set 2b (optimized towards SSST)	3540.8	1781.0

5 Measurement Setups

In the following, we briefly present the properties of our measurement setups.

5.1 Notations

For a specific side channel experiment we collect the set $\boldsymbol{I}$ of traces with N being the number of collected traces. One trace I of length T is represented by its samples $I = (i_0, ..., t_T)$ which have been acquired over time. The plaintext is denoted as $P = p_0||...||p_{15}$ and the key as $K = k_0||k_1||...||k_{15}$. The target intermediate value of the AES S-box is defined by $v_{i,n} = \text{SBOX}(k_{i,n} \oplus p_{i,n})$ for the subkey and plaintext of target byte $i \in [0, 15]$ and trace number $n \in [1, ..., N]$.

We denote $\oplus$ as the bitwise XOR-operator, $\boldsymbol{V}$ as the set of all possible intermediate values v, and $|\cdot|$ as the number of elements in a set. Whenever accessing a single value of a trace with number n, point in time t, and intermediate value v we denote this as $i^v_{n,t}$.$\boldsymbol{I}_v$ denotes the set of traces for the intermediate value v.

5.2 Signal-to-Noise-Ratio (SNR)

When investigating the properties of a measurement campaign from a security point of view, we are mostly interested in the *effectiveness* of distinguishing the

targeted values. For this purpose, the SNR definition by Mangard et al. in [23] has often been used. It is expressed by

$$\mathrm{SNR_M} = \frac{\mathrm{Var}(\mathrm{E}[\boldsymbol{I}_{X0}], ..., \mathrm{E}[\boldsymbol{I}_{X|\boldsymbol{V}|}])}{\mathrm{E}[\mathrm{Var}(\boldsymbol{I}_{X0}), ..., \mathrm{Var}(\boldsymbol{I}_{X|\boldsymbol{V}|})]} \qquad \forall\, X_j \in \boldsymbol{V} \tag{5}$$

and denoted as $\mathrm{SNR_M}$ in the following. It is known to be a useful tool to identify the points in time that have leakage in their first statistical moment.

5.3 High Resolution EM Measurement Setup

As a device under test, we use a decapsulated Spartan 6 FPGA, which is clocked at 8 MHz. For the FPGA as shown in Figure 2b, we rasterize an area of 2730 µm × 1600 µm with the probe at a distance to the surface of ≤ 50 µm. In total, we use 120 (15×8) equally-spaced positions to acquire measurements within this area.

For the FPGA as shown in Figure 2b, we rasterize an area of 2730 µm × 1600 µm with an equally-spaced 15×8 grid (120 measurement positions in total) and the probe at a distance of ≤ 50 µm to the die surface. For the measurement, we use a Langer ICR HH150-6 near-H-field (horizontal) probe with a coil diameter of 150 µm. The maximum bandwidth of the probe is 6 GHz with a built-in 30 dB preamplifier. In addition to that, we use another 30 dB amplifier such that the resulting signal is amplified by 60 dB in total.

With this setup synchronized to the device's clock, we collected 10 000 traces for each target design and position at a rate of 5 GS/s using a LeCroy WavePro 725 Zi. The resulting mean and $\mathrm{SNR_M}$ for WDDL are shown in Figure 7a.

5.4 Power Measurement Setup

Aside from the change in the measurement approach, the setup is kept the same for the power-based measurement, i.e., the same FPGA using the same designs. Instead of the H-field probe and amplifiers, a differential probe (LeCroy AP033) measures the voltage drop across a 10 Ω shunt resistor.

With this setup, we collected a total of 100 000 traces for each target design using the same clock frequency and samplingrate. As an example, the resulting mean and $\mathrm{SNR_M}$ for WDDL are shown in Figure 7b.

6 Practical Investigations

In this chapter we present the concept of our practical investigations and the results for the power and high resolution EM measurements.

6.1 Concept of Investigations

To guarantee comparable results between the four logic styles, we adhere to the following requirements: (i) only identical bit files are used for the comparison of

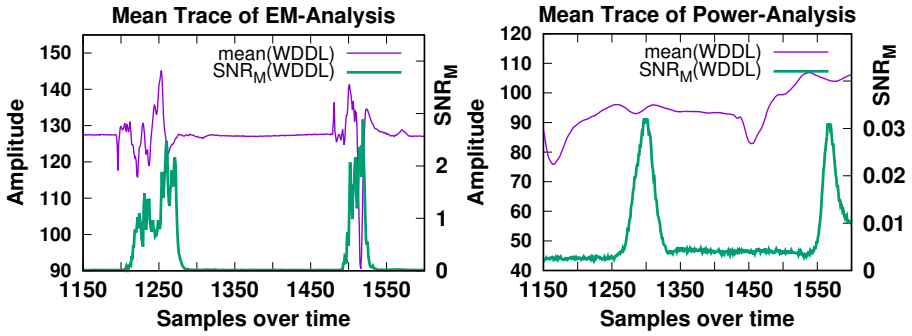

(a) Mean and $\mathrm{SNR_M}$ of WDDL (based on localized EM measurement)

(b) Mean and $\mathrm{SNR_M}$ of WDDL (based on power measurement over shunt)

Fig. 7: Basic properties of the design and comparison of the measurement setups.

power and high resolution EM measurement (*ii*) the same routing is realized for all logic styles (*iii*) the improvement over the SingleRail logic is only considered *within* the same measurement method. Hence, Figures 8 and 9 are based on the same `bit` files with the same routing using the default ISE placement (Set 1).

To fairly compare the properties of the implementations and analyze their leakage, we selected three metrics. The first is a correlation based leakage test proposed by Durvaux et al. [8], which works very similar to a CPA with profiled power model. To carry out the test, the traces are split into two sets, the profiling set $\boldsymbol{I}_p$ and an attack set $\boldsymbol{I}_a$. The profiling set is used to estimate the power consumption model $\boldsymbol{m}$ of the device by calculating the mean for each point in time for each element in $\boldsymbol{V}$, according to

$$\boldsymbol{m}_v = \frac{1}{|\boldsymbol{I}_p(v)|} \sum_{n \in |\boldsymbol{I}_p(v)|} \boldsymbol{I}_{p,n} \tag{6}$$

As a result, a power model is created which is based on practical measurements. It therefore better reflects the actual properties of the device when compared to "black-box" power models, e.g., Hamming weight or Hamming distance. $\boldsymbol{I}_p(v)$ and $\boldsymbol{I}_a(v)$ denotes the selection of all traces with the internal value v.

Afterwards the correlation vector $\boldsymbol{corr}$ is computed by correlating each trace of the attack set $\boldsymbol{I}_a$ with the corresponding value of $\boldsymbol{m}_i = (m(0,i), ..., m(|\boldsymbol{I}_a|, i))$, as shown in Eq. 7. In this case, $m(n,i)$ denotes the element of m_v for the intermediate value v of trace number n and target byte i.

$$\boldsymbol{corr}_{t,i} = \rho(\boldsymbol{m}_i, \boldsymbol{i}_{a,t}) \tag{7}$$

To quantify the achieved security complexity, we make use of the properties of the Pearson correlation coefficient, as the measurements to disclosure are proportional to $(\max(\boldsymbol{corr})^2)$. Based on this behavior we define the security

gain as: secgain $= (\frac{\max(\boldsymbol{corr}_1)}{\max(\boldsymbol{corr}_2)})^2$. The thus created power model and resulting correlation coefficient leads to the detection of first order leakage.

To complement our correlation-based analysis, we use the mutual information (MI) which is an information theoretic (IT) metric proposed by [36] and [10] independently. It has the advantage of detecting leakages at arbitrary order. Hence, it captures the amount of information available to the worst-case adversary. We use it to directly calculate the mutual information between a given trace and the S-box input, as shown in Eq. 8, for each point in time t and target byte i.

$$\boldsymbol{mi}_{t,i} = \boldsymbol{H}(\boldsymbol{I}_t) - \boldsymbol{H}(\boldsymbol{I}_t|\boldsymbol{v}_i) \tag{8}$$

Since our goal is to also compare power and high resolution EM measurements, we additionally include the results of $\mathrm{SNR_M}$ as third metric.

6.2 Power Measurement Results for Default ISE Placement (Set 1)

By carrying out a power-based side-channel attack first, we confirm the results of previous publications such as [26] and showcase the correct behavior of our implementations. According to our concept, we perform the correlation based leakage test which leads to the curves as shown in Fig. 8. The results are based on a default ISE placement (Set 1) with the described technique to ensure the same routing amongst all the considered candidates. Each evaluated implementation shows two correlation peaks which correspond to evaluation and precharge phase.

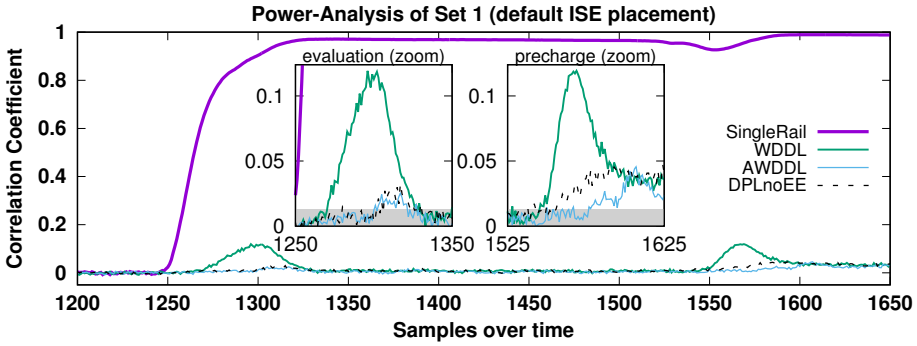

Fig. 8: Correlation based leakage test of a power measurement for AWDDL, DPLnoEE, WDDL, and SingleRail of the default placement

The obtained correlations are well above the significance threshold of 0.012, defined by Mangard et al. [23]. The insignificant region is indicated by the grey area inside the boxes of Fig. 8. Clearly visible is the strong correlation of the SingleRail variant that climbs up to 0.99. It is, as expected, orders of magnitude higher when compared to the dual-rail logics. WDDL shows a maximum correlation of 0.119, DPLnoEE of 0.048, and AWDDL of 0.046 respectively.

Based on this metric, the WDDL design – due to its data-dependent time-of-evaluation and time-of-precharge – has the highest leakage of the dual rail styles. DPLnoEE and AWDDL show a similar leakage in the evaluation phase. As claimed by [26], the leakage of AWDDL is marginally lower in the precharge phase when compared to DPLnoEE. The plots also show a leakage of AWDDL that is shifted in time which is owed to its self-timed behavior. A massive leakage is observed for the SingleRail implementation that spreads over both evaluation and precharge phase. This is probably due to parasitics of the power measurement setup. The resulting security gains are:

- SingleRail → WDDL: $\left(\frac{0.990}{0.119}\right)^2 = 69.2$
- WDDL → DPLnoEE: $\left(\frac{0.119}{0.048}\right)^2 = 6.15$
- DPLnoEE → AWDDL: $\left(\frac{0.048}{0.046}\right)^2 \approx 1$ (difference below significance interval)

To complement our analysis we applied the information theoretic metric, too. They confirm the results of the correlation based leakage test and are summarized in Table 3. As a next step, we investigate if similar security gains can be obtained if the device under test is subject to a localized-EM attack.

6.3 Localized-EM Measurement for Default ISE Placement (Set 1)

Using the same design files and same DUT, we also performed high resolution, localized EM measurements. This leads to Figure 9 for the correlation based test. Again, we indicated the insignificant region by a grey area inside the plot.

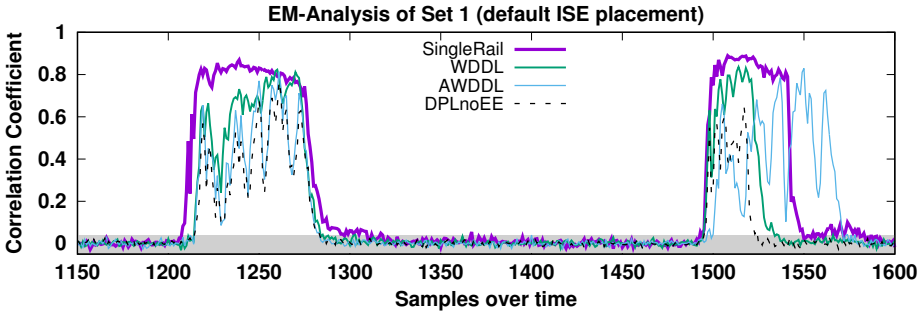

Fig. 9: Leakage test of a localized EM measurement for the considered logic styles at the position of the highest correlation using the default ISE placement.

It is striking that during the evaluation phase all correlation curves peak at very similar levels. Considering both phases, the SingleRail implementation has its highest peak at 0.89, followed by WDDL and AWDDL at 0.83, and DPLnoEE at 0.77. We would like to highlight that only 10 000 traces were necessary to create these results. Deriving the respective security gains, we get

– SingleRail → WDDL: $\left(\frac{0.889}{0.836}\right)^2 = 1.13$
– WDDL → DPLnoEE: $\left(\frac{0.836}{0.768}\right)^2 = 1.18$
– DPLnoEE → AWDDL: $\left(\frac{0.768}{0.829}\right)^2 = 0.858$

which shows a barely noticable security gain when using a high resolution, localized EM attack. Since the same `bit` files were used this is clearly owed to the superior measurement acquisition. Again, the results of the information theoretic metric are added to Table 3.

It is remarkable that under this setting, AWDDL performs worse when compared to DPLnoEE. This is owed to the fact that DPLnoEE gates directly go to precharge once one of the inputs goes to precharge. In contrast, AWDDL goes to precharge only when both inputs are in precharge. Therefore, the propagation wave of AWDDL spreads over time which leads to the presented result, i.e., the leakage of AWDDL continues even after that of the SingleRail has stopped.

6.4 Comparing Localized EM and Power Measurements

In this section we compare the results (as given in Table 3) of the power and localized EM measurements, both of which are using the same `bit` files (cf. Sections 6.2 and 6.3). Hence, routing and placement is the same for both measurements and all considered logic styles (Set 1). Therefore, the only substantial differences can only be caused by the specifics of the measurement setups.

Table 3: Summary of the practical evaluations for the default ISE placement.

Design	**Attack**	**SNR**$_{\mathrm{M}}$	**max**(***corr***)	**max**(**secgain**)	**max**(**MI**)
SingleRail	Power	54.16	0.990	←	2.99
WDDL		0.032	0.119	**463.2**	0.057
AWDDL		0.011	0.046	←	0.030
DPLnoEE		0.013	0.048		0.032
SingleRail	EM	4.06	0.889	←	2.93
WDDL		2.89	0.836	**1.34**	2.57
AWDDL		1.80	0.829		1.96
DPLnoEE		1.37	0.768	←	1.95

When evaluating the results w.r.t. the obtained security gain, it is striking that they differ significantly between power (about a factor of 463) and localized EM measurements (about a factor of 1.34). This strongly supports the argument that localized EM attacks are a severe threat to dual-rail logic on FPGAs.

For the SingleRail implementation, it is surprising that the total leakage appears to only be captured by a power measurement, resulting in a much higher SNR$_{\mathrm{M}}$ of 54.16 when compared to the localized EM measurement (4.06). However, this changes drastically when inspecting the numbers for DRP logic as the results of localized EM outperform the power-based by orders of magnitude.

Another approach to substantiate the impact of our results is the capability of the setups in distinguishing the true and false rail. To analyze this, we need to consider two scenarios: (*i*) In case the rails are inseparable, one would expect to see a significant increase in the leakage when deactivating one of the rails since the balancing aspect is lost. (*ii*) If the opposite is true, i.e., the rails are separable, then the leakage should be approximately the same when deactivating one of the rails since from the beginning on, there would have been no difference.

This can be studied by observing the behavior when switching from the SingleRail implementation to any of the dual-rail variants. For the power measurements we are within scenario (*i*), as the correlation significantly differs when performing the switch between SingleRail and dual-rail. For the local EM measurements we are within scenario (*ii*), as the magnitude of correlation remains the same regardless of the fact whether it is a SingleRail or a dual-rail variant.

All observations are backed by the IT analysis that verifies the claimed behavior of the different measurement methods. Under this scenario, both measurement methods perform equally when considering the SingleRail implementation, indicating that aside from the difference in $\mathrm{SNR_M}$, the full leakage is extracted.

As a next step, we investigate if this situation can be improved by means of an increased density of the placement. The goal of this is to increase the density up to the point where distinguishing the rails is no longer possible.

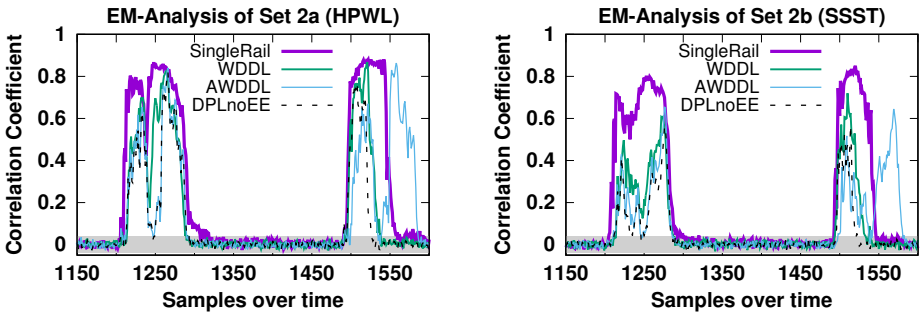

(a) Correlation curves for Set 2a.

(b) Correlation curves for Set 2b.

Fig. 10: Correlation based leakage test of a high resolution EM measurement.

6.5 Security Analysis of the Custom Placement

To improve the situation under a localized EM attack, we investigate the impact of the previously described placement improvements. We repeated our measurements using these files and carried out the same tests. For HPWL and SSST, the results of this test are illustrated in Figure 10a and Figure 10b respectively.

Again, for the sake of fair comparison, we realized the same routing within the HPWL (Set2a) and SSST (Set2b) designs. We therefore *only* compare the results of the placement and the security improvement from a SingleRail to dual-rail version. Otherwise, the effect of the routing could not be excluded.

Considering the results for HPWL, one can see that for Figure 10a almost no improvement is achieved. In contrast, the SSST-based placement shows an improvement at a factor of about 2.24. Taking the results of Table 2 into account, it is evident that by using the SSST cost function one achieves a more dense placement up to the point where the power of the used EM attack is degrading.

However, since an optimal placement is likely not to be found analytically (which could further improve the resistance), only improving the placement as a countermeasure is insufficient. We therefore analyzed also a masked version of AWDDL. These supplementary results are shown in Figure 11 of the Appendix.

Table 4: Summary of the EM analysis for the customized placements.

Design	**Attack**	$\mathbf{SNR_M}$	**max**(***corr***)	**max**(**secgain**)	**max**(**MI**)
HPWL					
SingleRail	EM	2.82	0.881	←	2.76
WDDL		3.04	0.867	**1.15**	2.12
AWDDL		1.70	0.863		2.21
DPLnoEE		1.92	0.823	←	1.96
SSST					
SingleRail	EM	3.23	0.851	←	2.31
WDDL		1.54	0.720	**2.24**	1.73
AWDDL		1.29	0.653		1.20
DPLnoEE		1.36	0.569	←	1.14

7 Conclusion

In this work we have shown that verifying DRP logics on FPGAs only by a power-based side-channel analysis is insufficient. While their security gain is remarkable in this setting, it is not when considering high-resolution, localized EM measurements. We therefore suggest to always include a thorough EM-based analysis in future proposals of such logic styles.

To compensate for the significant loss in security under an EM-based attack, we investigated if the situation improves when adapting the placement. This is achieved by a custom placer using simulated annealing using a novel cost function. Our practical investigations confirm that by using a more dense placement, the security doubles when compared to the default ISE setting.

While generally assuming that a single countermeasure is insufficient and combining multiple countermeasures is needed, we demonstrate that for dual-rails on FPGAs this may result in a wrong systematic, as they may be rendered mostly useless, especially if not taking care of the placement.

Even though we did not specifically consider duplication schemes, we expect that our findings apply to them as well, since the minimum distance between their `true` and `false` is typically large, i.e., more than one tile. This needs to be confirmed by future evaluations, also considering triple-rail logics such as [22].

References

1. Federal Information Processing Standards Publication (FIPS 197). Advanced Encryption Standard (AES), 2001.
2. V. Betz and J. Rose. *VPR: A New Packing, Placement and Routing Ttool for FPGA Research*.
3. S. Bhasin, S. Guilley, F. Flament, N. Selmane, and J.-L. Danger. Countering Early Evaluation: An Approach Towards Robust Dual-Rail Precharge Logic. In *WESS 2010*, page 6. ACM, 2010.
4. D. Canright. A Very Compact S-Box for AES. In *CHES 2005*, volume 3659 of *LNCS*, pages 441–455. Springer, 2005.
5. C.-L. E. Cheng. RISA: Accurate and Efficient Placement Routability Modeling. In *Proceedings of the 1994 IEEE/ACM International Conference on Computer-aided Design*, ICCAD '94, Los Alamitos, CA, USA. IEEE Computer Society Press.
6. T. D. Cnudde, B. Bilgin, B. Gierlichs, V. Nikov, S. Nikova, and V. Rijmen. Does Coupling Affect the Security of Masked Implementations? Cryptology ePrint Archive, Report 2016/1080, 2016.
7. E. De Mulder, P. Buysschaert, S. Ors, P. Delmotte, B. Preneel, G. Vandenbosch, and I. Verbauwhede. Electromagnetic Analysis Attack on an FPGA Implementation of an Elliptic Curve Cryptosystem. In *Computer as a Tool, 2005. EUROCON 2005.The International Conference on*, volume 2, pages 1879 –1882, Nov. 2005.
8. F. Durvaux and F.-X. Standaert. *From Improved Leakage Detection to the Detection of Points of Interests in Leakage Traces*, pages 240–262. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
9. I. Giechaskiel and K. Eguro. Information Leakage Between FPGA Long Wires. *CoRR*, 2016.
10. B. Gierlichs, L. Batina, P. Tuyls, and B. Preneel. Mutual Information Analysis. In *CHES 2008*, volume 5154 of *LNCS*, pages 426–442. Springer, 2008.
11. S. Guilley, P. Hoogvorst, Y. Mathieu, and R. Pacalet. The "Backend Duplication" Method. In *CHES 2005*, volume 3659 of *LNCS*, pages 383–397. Springer, 2005.
12. T. Güneysu and A. Moradi. Generic Side-Channel Countermeasures for Reconfigurable Devices. In *CHES 2011*, volume 6917 of *LNCS*. Springer, 2011.
13. W. He, E. de la Torre, and T. Riesgo. A Precharge-Absorbed DPL Logic for Reducing Early Propagation Effects on FPGA Implementations. In *ReConFig 2011*. IEEE Computer Society, 2011.
14. W. He and A. Herrmann. Placement Security Analysis for Side-Channel Resistant Dual-Rail Scheme in FPGA. In *Proceedings of the Second Workshop on Cryptography and Security in Computing Systems*, CS2 '15, 2015.
15. W. He, A. Otero, E. de la Torre, and T. Riesgo. Automatic Generation of Identical Routing Pairs for FPGA Implemented DPL Logic. In *ReConFig 2012*. IEEE, 2012.
16. C. Herbst, E. Oswald, and S. Mangard. An AES Smart Card Implementation Resistant to Power Analysis Attacks. In *ACNS 2006*, volume 3989 of *LNCS*, pages 239–252. Springer, 2006.
17. J. Heyszl, S. Mangard, B. Heinz, F. Stumpf, and G. Sigl. Localized Electromagnetic Analysis of Cryptographic Implementations. In O. Dunkelman, editor, *Topics in Cryptology - CT-RSA 2012*, volume 7178 of *Lecture Notes in Computer Science*, pages 231–244. Springer Berlin / Heidelberg, 2012.
18. J. Heyszl, D. Merli, B. Heinz, F. D. Santis, and G. Sigl. Strengths and Limitations of High-Resolution Electromagnetic Field Measurements for Side-Channel Analysis. In *Smart Card Research and Advanced Applications - 11th International Conference, CARDIS*, pages 248–262, 2012.

19. J.-P. Kaps and R. Velegalati. DPA Resistant AES on FPGA Using Partial DDL. In *FCCM 2010*, pages 273–280. IEEE Computer Society, 2010.
20. P. C. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In *CRYPTO 1999*, volume 1666 of *LNCS*, pages 388–397. Springer, 1999.
21. C. Lavin, M. Padilla, J. Lamprecht, P. Lundrigan, B. Nelson, B. Hutchings, and M. Wirthlin. RapidSmith – A Library for Low-level Manipulation of Partially Placed-and-Routed FPGA Designs. Technical report, Brigham Young University, September 2012.
22. V. Lomné, P. Maurine, L. Torres, M. Robert, R. Soares, and N. Calazans. Evaluation on FPGA of triple rail logic robustness against DPA and DEMA. In *DATE 009*, pages 634–639. IEEE, 2009.
23. S. Mangard, E. Oswald, and T. Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer, 2007.
24. S. Mangard and K. Schramm. Pinpointing the Side-channel Leakage of Masked AES Hardware Implementations. CHES, 2006.
25. A. Moradi, T. Eisenbarth, A. Poschmann, and C. Paar. Power Analysis of Single-Rail Storage Elements as Used in MDPL. In *ICISC 2009*, volume 5984 of *LNCS*, pages 146–160. Springer, 2009.
26. A. Moradi and V. Immler. *Early Propagation and Imbalanced Routing, How to Diminish in FPGAs*. 2014.
27. G.-J. Nam and P. G. Villarrubia. Placement: Introduction/Problem Formulation. In C. J. Alpert, D. P. Mehta, and S. S. Sapatnekar, editors, *Handbook of Algorithms for Physical Design Automation*, chapter 14, pages 277 – 287. Auerbach Publications, 1 edition, November 2008.
28. M. Nassar, S. Bhasin, J.-L. Danger, G. Duc, and S. Guilley. BCDL: A high speed balanced DPL for FPGA with global precharge and no early evaluation. In *DATE 2010*, pages 849–854. IEEE, 2010.
29. S. Nikova, V. Rijmen, and M. Schläffer. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. *J. Cryptology*, 24(2):292–321, 2011.
30. E. Oswald, S. Mangard, N. Pramstaller, and V. Rijmen. A Side-Channel Analysis Resistant Description of the AES S-Box. In *FSE 2005*, volume 3557 of *LNCS*, pages 413–423. Springer, 2005.
31. E. Peeters, F.-X. Standaert, and J.-J. Quisquater. Power and Electromagnetic Analysis: Improved Model, Consequences and Comparisons. *Integration, the VLSI Journal*, 2007.
32. J.-J. Quisquater and D. Samyde. ElectroMagnetic Analysis (EMA): Measures and Counter-measures for Smart Cards. In I. Attali and T. Jensen, editors, *Smart Card Programming and Security*, volume 2140 of *Lecture Notes in Computer Science*, pages 200–210. Springer Berlin / Heidelberg, 2001.
33. L. Sauvage, S. Guilley, J.-L. Danger, Y. Mathieu, and M. Nassar. Successful Attack on an FPGA-based WDDL DES Cryptoprocessor Without Place and Route Constraints. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '09, 2009.
34. L. Sauvage, M. Nassar, S. Guilley, F. Flament, J.-L. Danger, and Y. Mathieu. DPL on Stratix II FPGA: What to Expect? In *ReConFig 2009*, pages 243–248. IEEE Computer Society, 2009.
35. R. Specht, J. Heyszl, M. Kleinsteuber, and G. Sigl. *Improving Non-profiled Attacks on Exponentiations Based on Clustering and Extracting Leakage from Multi-channel High-Resolution EM Measurements*, pages 3–19. Springer International Publishing, Cham, 2015.

36. F.-X. Standaert, T. Malkin, and M. Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 443–461. Springer, 2009.
37. D. Suzuki and M. Saeki. *Security Evaluation of DPA Countermeasures Using Dual-Rail Pre-charge Logic Style*. 2006.
38. W. Swartz. Placement Using Simulated Annealing. In C. J. Alpert, D. P. Mehta, and S. S. Sapatnekar, editors, *Handbook of Algorithms for Physical Design Automation*, pages 311–325. Auerbach Publications, November 2008.
39. K. Tiri, D. Hwang, A. Hodjat, B.-C. Lai, S. Yang, P. Schaumont, and I. Verbauwhede. *Prototype IC with WDDL and Differential Routing – DPA Resistance Assessment*. 2005.
40. K. Tiri and I. Verbauwhede. A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation. In *DATE 2004*, pages 246–251. IEEE Computer Society, 2004.
41. K. Tiri and I. Verbauwhede. Place and Route for Secure Standard Cell Design. In *CARDIS 2004*, pages 143–158. Kluwer, 2004.
42. F. Unterstein, J. Heyszl, F. De Santis, and R. Specht. Dissecting Leakage Resilient PRFs with Multivariate Localized EM Attacks - A Practical Security Evaluation on FPGA. In *Constructive Side-Channel Analysis and Secure Design: 8th International Workshop, April 13-14, 2017*, Paris, France. Springer International Publishing.
43. A. Wild, A. Moradi, and T. Güneysu. GliFreD: Glitch-Free Duplication - Towards Power-Equalized Circuits on FPGAs. 2015.
44. P. Yu and P. Schaumont. Secure FPGA circuits using controlled placement and routing. In *CODES+ISSS 2007*, pages 45–50. ACM, 2007.

Appendix

For the sake of completeness, we present the results of a simple boolean masked version of AWDDL as an example in Figure 11, using the default placement of ISE. Both power and localized EM attack have been carried out. The first order correlation based leakage test did (as expected) not show any leakage.

In contrast, using the mutual information, it was still possible for both designs to extract leakage. Hence, additional countermeasures would be required.

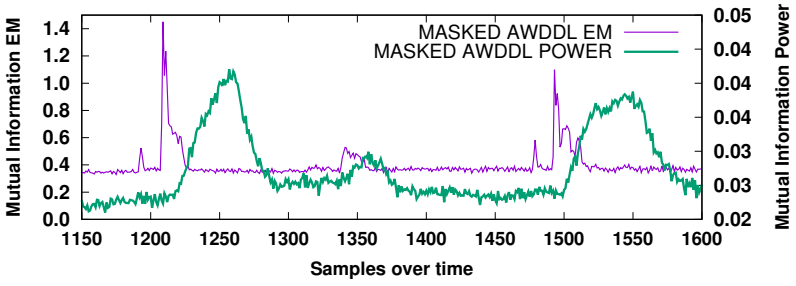

Fig. 11: Mutual information of the evaluation and precharge phases over time.