

Assessment of Hiding the Higher-Order Leakages in Hardware

What Are the Achievements Versus Overheads?

Amir Moradi^(✉) and Alexander Wild

Horst Görtz Institute for IT-Security, Ruhr-Universität Bochum, Bochum, Germany
{Amir.Moradi,Alexander.Wild}@rub.de

Abstract. Higher-order side-channel attacks are becoming amongst the major interests of academia as well as industry sector. It is indeed being motivated by the development of countermeasures which can prevent the leakages up to certain orders. As a concrete example, threshold implementation (TI) as an efficient way to realize Boolean masking in hardware is able to avoid first-order leakages. Trivially, the attacks conducted at second (and higher) orders can exploit the corresponding leakages hence devastating the provided security. Hence, the extension of TI to higher orders was being expected which has been presented at ASIACRYPT 2014. Following its underlying univariate settings it can provide security at higher orders, and its area and time overheads naturally increase with the desired security order.

In this work we look at the feasibility of higher-order attacks on first-order TI from another perspective. Instead of increasing the order of resistance by employing higher-order TIs, we realize the first-order TI designs following the principles of a power-equalization technique dedicated to FPGA platforms, that naturally leads to hardening higher-order attacks. We show that although the first-order TI designs, which are additionally equipped by the power-equalization methodology, have significant area overhead, they can maintain the same throughput and more importantly can avoid the higher-order leakages to be practically exploitable by up to 1 billion traces.

1 Introduction

Side-channel attacks are a major threat to the security of modern embedded devices. If no particular attention is paid, the exploitation of physical leakages such as the power consumption and the electromagnetic radiation of a cryptographic implementation can lead to successful key recoveries, e.g., [2, 16, 27, 44, 58]. As a consequence, the topic has been followed by a vast literature on potential solutions to defeat such attacks.

The countermeasures against side-channel attacks range from ad hoc to formal, and are defined to be applied at various abstraction levels. For instance, time randomizations (based on random delay insertion [14] or shuffling [54]) are frequently-used low-overhead heuristic-based approaches (mainly) for software-based applications. These *hiding* schemes are not limited to only those which

randomize the computations in time, but covers the approaches that add noise resources [18, 24] as well as those aiming to equalize the power consumption [51, 53]. The time randomizations can be overcome by preprocessing the leakage traces (e.g., combing [24]), and the effect of the noise additions can be mitigated by increasing the number of traces [24]. In contrast, the power-equalization techniques usually fail due to wrong assumptions (e.g., ignoring early propagation [50]) or overestimating the ability of the tools (e.g., balanced dual-rail routing [52]). Apart from [29, 53], dual-rail precharge logic styles, which have been initially designed for ASIC-based applications (e.g., [13, 40, 41, 51]), cannot be easily integrated into the FPGAs. Instead, other approaches like [3, 19–21, 23, 36, 47, 57] have particularly been developed with respect to the resources available in certain FPGAs. However, each of such techniques suffers from a flaw that prevents them to be considered as a potential solution (see [56] for details of each flaw). Further, a design methodology which combines a dual-rail logic style and duplication in FPGAs [57] has also been shown to be flawed [55]. As an alternative, the technique presented in [56] (so-called GliFreD) seems to avoid the known pitfalls. It has been designed particularly for Xilinx FPGAs, and aims at avoiding early propagation, preventing the glitches, and relaxing the necessity of a dual-rail routing tool. It seems that GliFreD can satisfy its goals toward equalizing the power consumption, but an ideally-equal situation cannot still be achieved due to the process variation violating the balance between the cloned routes.

On the other hand, probably the most investigated and best understood protection against side-channel attacks is *masking* [12, 15, 46]. The underlying principle of masking is to represent any sensitive variable in the implementation by d shares in such a way that the computations are performed only on these shares. Assuming that the leakage of the shares are independent of each other, a successful key-recovery attack needs to observe – at least – the d th-order statistical moment of the leakage distributions, where the corresponding complexity increases exponentially with d .

However, the independence of leakages associated to the shares is an assumption which is usually violated in hardware applications. As an example, the masked AES Sbox designs [11, 39], where the glitches are ignored, failed in practice to satisfy the desired security level, i.e., first-order resistance [25, 32]. Instead, based on Boolean masking and multiparty computation, threshold implementations (TI) [37, 38] can ensure first-order resistance in the presence of glitches. Indeed, not only its underlying principles are sound and realistic but also practical investigations confirmed its effectiveness [4, 33]. Trivially, higher-order attacks are feasible on TI designs [4, 26], which motivated the work presented in [5] where the concept of higher-order TI is demonstrated that extends its definitions to any order. Regardless of its significant overhead (e.g., requiring at least $d = 5$ for a second-order security) the note given in [45] and later practically confirmed in [49] made clear that the definitions of the higher-order TI stand valid only in univariate scenarios.

Our Contribution. Indeed, it is known to the community that hiding techniques (in particular power-equalizing approaches) are not *solely* capable to prevent

key-recovery attacks. It is always suggested that such techniques should be combined with other countermeasures, but the benefit of such a combination has never truly been examined for a hardware platform. More precisely, exploiting higher-order leakages becomes extremely hard in practice when the leakage traces are sufficiently noisy [43]. Along the same lines, power-equalization schemes are also expected to reduce the signal (versus the noise) and have the same effect. To the best of our knowledge, the only work which tried to proceed toward this goal is [30], where a flawed masking scheme [11] has been implemented in a glitch-free setting. No particular attention has been paid on equalizing the power hence not a concrete hiding technique.

Our contribution in this work is to examine the benefit of combining two sound hardware-based countermeasures. More precisely, we aim at considering a provably (first-order) secure masking scheme (TI) and realize it under the principles of a proper power-equalizing technique (GliFreD). We pursue an investigation of our combined construction compared with:

- the same masking design (first-order TI) without employing any hiding technique, and
- the second-order TI of the same design excluding any power-equalization scheme.

Such comparisons with respect to the data complexity of leakage detection as well as time and area overheads of the designs allows us to have an overview on the tradeoff between the gains and overheads of different countermeasures as well as their combination.

Since the design overheads are application specific, we consider two design methodologies: first, a fully serialized architecture for lightweight applications with KATAN-32 cipher and second, a parallelized architecture for high-speed applications with PRESENT cipher. Amongst our achievements in this work – including a second-order TI of PRESENT – we can refer to the designs we developed with a combination of GliFreD and the first-order TI (of both KATAN-32 and PRESENT) which showed to be secure by up to 1 billion power traces measured from a Spartan-6 FPGA platform.

2 GliFreD

Dual-rail Precharge Logic (DPL) schemes are popular side-channel countermeasures for hardware circuits and assigned to the group of hiding techniques. Each DPL scheme places two contrary working (*true* and *false*) circuits on a device to *ideally* decorrelate the power consumption from the processed data. In common, DPL schemes have to deal with some implementation challenges. The three major challenges that the FPGA-based DPL designers face are: early propagation, glitches and different wire capacitance of coupled signals. GliFreD is a DPL scheme exclusively designed for FPGAs, and is amongst the few schemes which address all these three problems [56].

To overcome the aforementioned problems GliFreD defines the following design methodology. Each Look-Up Table (LUT) instance is connected to two global control signals: `CLK` and `active`; the later one toggles with half of the other one's frequency. These control signals determine whether the LUTs reside in precharge or in evaluation phase. Hence, the regulated LUT transitions overcome the definition of early evaluation [50]. To prevent the propagation of the LUT output transition, a register is connected to each LUT output. However, a single register stage in a DPL circuit contradicts the requirement of a constant gate and register transition per clock cycle [28] as inconstant and data-dependent transitions would result in data-dependent leakage. Therefore, the GliFreD principles require to place an even number of register stages between each two LUTs connected in the circuit. Consequently, GliFreD forms a pipeline architecture which prevents glitches by halting the propagation of a signal after each LUT. Figure 1(a) shows the timing diagram of a GliFreD circuit.

Similar to many DPL schemes, GliFreD also needs to place a dual of the circuit. Copying the routing structure is currently the best known way in FPGAs to keep the wire capacitances of the *false* circuit as equivalent as those of the *true* circuit. Hence, to perform the circuit dualization, i.e., placing the *false* circuit, a second horizontally-moved instance of the *true* circuit is placed on the FPGA. The copy process is performed on netlist level to pass on the routing information to the *false* circuit.

GliFreD allows an arbitrary LUT configuration; since both control signals `CLK` and `active` should be connected to each LUT, the function f each LUT can realize is limited to a 4-to-1 look-up table. The output of each LUT can be seen as $O = \text{active} \cdot \overline{\text{CLK}} \cdot f(I_2, \dots, I_5)$ ¹, while the corresponding dual function (of the *false* circuit) becomes $\overline{O} = \text{active} \cdot \overline{\text{CLK}} \cdot f(\overline{I_2}, \dots, \overline{I_5})$. Figure 1 shows the GliFreD pendant of an exemplary function

$$y = x_0 + x_0x_3 + x_2x_3 + x_3x_4 + x_3x_6 + x_0x_7 + x_2x_7, \quad (1)$$

whose standard implementation is shown in Fig. 1(b).

Since the output of each LUT is buffered by a register, the critical path in a GliFreD circuit is minimized allowing to run the circuit at high frequencies. To this end the delay between the `CLK` and `active` signals should be kept minimum (see Fig. 1(a)), that can be achieved by forcing `active` signal to be routed through the clock trees. The GliFreD design methodology offers the ability to transfer a design into a fully-pipelined architecture, hence achieving a high throughput in combination with a high clock frequency. In general, large combinatorial circuits cause glitches which propagate through the whole circuit. Since GliFreD prevents those glitches, it may also reduce the power consumption. In small combinatorial circuits this benefit is faded and dominated by the increased amount of resources the GliFreD circuit utilizes. Nevertheless, GliFreD is a resource-costly solution. The LUT overhead (at most 8) required to form a GliFreD circuit strongly depends on the original design structure. Compared to the LUT utilization GliFreD causes

¹ I_0 and I_1 are reserved for `CLK` and `active`.

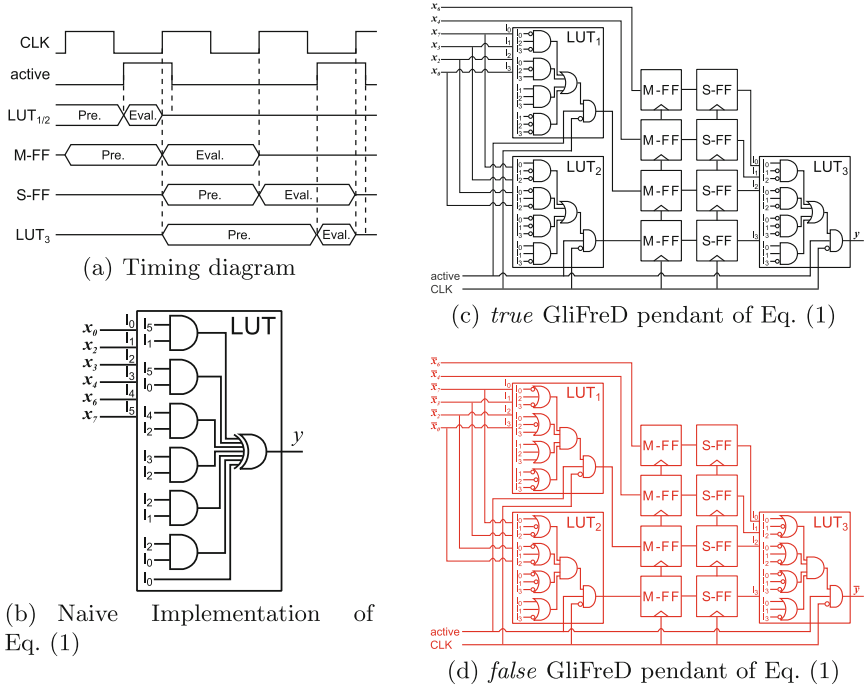


Fig. 1. An exemplary function implemented in a standard 6-to-1 LUT architecture and its GliFreD representation including the timing diagram

a massive register overhead and hence an increased latency. The register overhead cannot be trivially estimated and depends on the LUT depth, width and the amount of registers in the original design.

3 Case Studies

Before giving the details of our case studies, we briefly restate the concept behind threshold implementation.

3.1 Threshold Implementation

As stated before, the masking scheme which we consider in this work is threshold implementation (TI) introduced and extended in [4, 5, 37, 38]. Let us denote an intermediate value of a cipher by \mathbf{x} made of s single-bit signals $\langle x_1, \dots, x_s \rangle$. The underlying concept of TI is to use Boolean masking to represent \mathbf{x} in a shared form $(\mathbf{x}^1, \dots, \mathbf{x}^n)$, where $\mathbf{x} = \bigoplus \mathbf{x}^i$ and each \mathbf{x}^i similarly denotes a vector of s single-bit signals $\langle x_1^i, \dots, x_s^i \rangle$. A linear function $l(\cdot)$ can be trivially applied over the shares of \mathbf{x} as $l(\mathbf{x}) = \bigoplus l(\mathbf{x}^i)$. However, the realization of non-linear functions, e.g., an Sbox, over Boolean masked data is challenging. Following the concept of

TI, if the algebraic degree of the underlying Sbox is denoted by t and the desired security order by d , the minimum number of shares to realize the Sbox under the TI settings is $n = td + 1$. Further, such a TI Sbox provides the output $\mathbf{y} = S(\mathbf{x})$ in a shared form $(\mathbf{y}^1, \dots, \mathbf{y}^m)$ with at least $m = \binom{n}{t}$ shares. Note that the bit length of \mathbf{x} and \mathbf{y} (respectively of their shared forms) are not necessary the same since $S(\cdot)$ might be not a bijection, e.g., in case of DES.

Each output share $\mathbf{y}^{j \in \{1, \dots, m\}}$ is given by a component function $f^j(\cdot)$ over a subset of the input shares. To achieve the d th-order security, any d selection of the component functions $f^{j \in \{1, \dots, m\}}(\cdot)$ should be independent of at least one input share.

Since the security of masking schemes is based on the uniform distribution of the masks, the output of a TI Sbox must be also uniform as it is used as input in further parts of the implementation. To express the *uniformity* under the TI concept suppose that for a certain input \mathbf{x} all possible sharings $\mathcal{X} = \{(\mathbf{x}^1, \dots, \mathbf{x}^n) | \mathbf{x} = \bigoplus \mathbf{x}^i\}$ are given to a TI Sbox. The set made by the output shares, i.e., $\{(f^1(\cdot), \dots, f^m(\cdot)) | (\mathbf{x}^1, \dots, \mathbf{x}^n) \in \mathcal{X}\}$, should be drawn uniformly from the set $\mathcal{Y} = \{(\mathbf{y}^1, \dots, \mathbf{y}^m) | \mathbf{y} = \bigoplus \mathbf{y}^i\}$ as all possible sharings of $\mathbf{y} = S(\mathbf{x})$.

This uniformity check process should be individually performed for $\forall \mathbf{x} \in \{0, 1\}^s$. We should note that for $d > 1$ where $m > n$ the uniformity cannot be achieved. Hence, some of the registered output shares should be combined to reduce the number of output shares to n . Afterward the uniformity can be examined. For more detailed information we refer to the original articles [5, 38].

3.2 KATAN-32

As stated in Sect. 2, the overhead and performance of a GliFreD circuit depends on the nature of the underlying application. If the target design is made of small combinatorial circuits, the overhead of the resulting GliFreD circuit is minimal. Therefore, KATAN [10] which benefits from a serialized architecture with very small combinatorial logics is a suitable candidate for our investigations. Further, both first- and second-order uniform TI representation of its non-linear functions are given in [5], allowing us to develop the design with minimal efforts.

The architecture of our designs are based on those given in [5]. Figure 2(a) shows an overview of such a serialized architecture considering KATAN-32 encryption engine with 32-bit plaintext and 80-bit symmetric key. The plaintext and key are serially loaded into the registers, and after 254 clock cycles the ciphertext can be taken from the state register². The first-order TI of KATAN-32 with 3 shares (the minimum settings) needs the state (shift) registers to be tripled. Similar to that of [5], we do not represent the key (and the corresponding shift register) in a shared form. The XOR operations are easily repeated for each share, and the non-linear functions which are limited to the AND/XOR module (involved in function

² For more detailed information on the construction of functions f_a and f_b in Fig. 2(a) see [5, 10].

f_a and f_b of Fig. 2(a)) need to be realized under the concept of the first-order TI. An AND/XOR function receives a 3-bit input (a, b, c) and gives a single-bit output y as

$$y = a + bc.$$

Following the concept of *direct sharing* [6] the component functions (given in [5]) which realize a uniform first-order TI can be derived as

$$f^{i,j}(\langle a^i, b^i, c^i \rangle, \langle a^j, b^j, c^j \rangle) = a^j + b^j c^j + b^i c^j + b^j c^i, \quad (2)$$

where each output share is made by an instance of such a component function as

$$y^1 = f^{1,2}(\cdot, \cdot), \quad y^2 = f^{2,3}(\cdot, \cdot), \quad y^3 = f^{3,1}(\cdot, \cdot).$$

The same procedure is followed to realize the second-order TI of KATAN-32. First, the minimum number of shares is increased to 5, and all state registers and linear functions need to be repeated accordingly. Further, a second-order TI representation of AND/XOR module (given in [5]) can be derived from Eq. (2) and the following component function

$$g^{i,j}(\langle a^i, b^i, c^i \rangle, \langle a^j, b^j, c^j \rangle) = b^i c^j + b^j c^i. \quad (3)$$

In such a case, the output shares are made as

$$y^1 = f^{1,2}(\cdot, \cdot), \quad y^2 = f^{1,3}(\cdot, \cdot), \quad y^3 = f^{1,4}(\cdot, \cdot), \quad y^4 = f^{5,1}(\cdot, \cdot), \quad y^5 = f^{2,5}(\cdot, \cdot),$$

and

$$y^6 = g^{2,3}(\cdot, \cdot), \quad y^7 = g^{2,4}(\cdot, \cdot), \quad y^8 = g^{3,4}(\cdot, \cdot), \quad y^9 = g^{3,5}(\cdot, \cdot), \quad y^{10} = g^{4,5}(\cdot, \cdot).$$

As mentioned before, in a second-order case the output shares should be combined after being registered in order to reduce the number of shares back to 5. In this case, the reduction is done as

$$z^{i \in \{1, \dots, 4\}} = y^i, \quad z^5 = y^5 + y^6 + y^7 + y^8 + y^9 + y^{10},$$

thereby achieving a uniform second-order TI of the AND/XOR module [5]. For more clarification the formula for all the component functions are given in the extended version of this article [35].

3.3 PRESENT

As the second target we selected the PRESENT cipher [9] to be implemented in a round-based fashion. As Fig. 2(b) shows, 16 instances of the Sbox in addition to the `Player` operate in parallel to compute one cipher round. The reason for choosing such a target is to have an application for GliFreD with large combinatorial circuit compared to that of KATAN. Also, due to a possibility to decompose the PRESENT Sbox – as we express below – we are able to develop its uniform first- and second-order TI representations. We should note that we have not

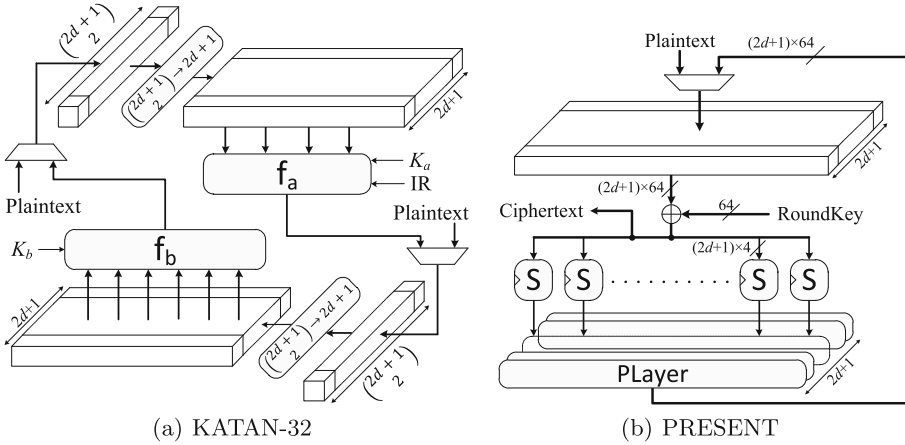


Fig. 2. Architecture of the case studies, first ($d = 1$) and second ($d = 2$) order TI

selected the AES as a target because its first-order TI (in [4, 33]) can only be realized by remasking (requiring multiple fresh mask bits per clock cycle) and furthermore there is not yet a clear roadmap how to realize its second-order TI.

Similar to the case of KATAN, the first-order (respectively second-order) TI of the targeted PRESENT architecture employs a 3-share (respectively 5-share) Boolean masking. The PLayer (realized by routing in the round-based architecture) is repeated on each share, and the key XOR is applied on only one share as the 80-bit key is not represented in a shared form. Clearly the remaining part is the TI representation of the PRESENT Sbox. Previously Poschmann et al. [42] have shown a decomposition and a uniform first-order TI of such an Sbox. However, below we represent another decomposition allowing us to develop its both first- and second-order uniform TI representations.

The PRESENT Sbox $S(\mathbf{x}) = \mathbf{y}$ is a cubic bijection (i.e., with algebraic degree $t = 3$) leading to minimum $n = 4$ and $n = 7$ shares in the first- and second-order TI settings respectively. Therefore, it is preferable to decompose the Sbox into two (at most) quadratic bijections F and G , in such a way that $S(\mathbf{x}) = F(G(\mathbf{x}))$ (i.e., $S = F \circ G$). If so, each F and G can be shared with $n = 3$ and $n = 5$ (for first- and second-order TI). According to the classifications given in [7], the PRESENT Sbox belongs to the cubic class \mathcal{C}_{266} . It means that there exist affine transformations A and B , where $S(\mathbf{x}) = B(\mathcal{C}_{266}(A(\mathbf{x})))$. In other words, S and \mathcal{C}_{266} are affine equivalent. To find the affine functions the algorithm given in [8] can be used; indeed there exist 4 such two affine functions. Also, as stated in [7] \mathcal{C}_{266} can be decomposed into two quadratic bijections. One of the possibilities is $\mathcal{Q}_{294} \times \mathcal{Q}_{299}$. It means that there exist three affine functions A_1, A_2, A_3 , where $\mathcal{C}_{266} = A_3 \circ \mathcal{Q}_{299} \circ A_2 \circ \mathcal{Q}_{294} \circ A_1$. Since \mathcal{C}_{266} and S are affine equivalent, there exist also three affine functions to decompose the PRESENT Sbox as

$$S(\mathbf{x}) = A_3 \left(\mathcal{Q}_{299} \left(A_2 \left(\mathcal{Q}_{294} (A_1(\mathbf{x})) \right) \right) \right). \quad (4)$$

We have found 229,376 such 3-tuple affine bijections, and we have selected one of the most simplest solutions with respect to the number of terms in their Algebraic Normal Form (ANF) directly affecting the size of the corresponding circuit.

The next step is to provide the uniform first-order TI of the quadratic bijections \mathcal{Q}_{294} and \mathcal{Q}_{299} which can be easily achieved by direct sharing [7]. For \mathcal{Q}_{294} :0123456789BAEFDC we can write

$$e = a + bd, \quad f = b + cd, \quad g = c, \quad h = d, \quad (5)$$

with $\langle a, b, c, d \rangle$ the 4-bit input, $\langle e, f, g, h \rangle$ the 4-bit output, and a and e the least significant bits. The component functions of the first-order TI of \mathcal{Q}_{294} can be derived by $f_{\mathcal{Q}_{294}}^{i,j}(\langle a^i, b^i, c^i, d^i \rangle, \langle a^j, b^j, c^j, d^j \rangle) = \langle e, f, g, h \rangle$ as

$$\begin{aligned} e &= a^i + b^i d^i + d^i b^j + b^i d^j & g &= c^i \\ f &= b^i + c^i d^i + d^i c^j + c^i d^j & h &= d^i \end{aligned} \quad (6)$$

The three 4-bit output shares provided by $f_{\mathcal{Q}_{294}}^{2,3}(\cdot, \cdot)$, $f_{\mathcal{Q}_{294}}^{3,1}(\cdot, \cdot)$ and $f_{\mathcal{Q}_{294}}^{1,2}(\cdot, \cdot)$ make a uniform first-order TI of \mathcal{Q}_{294} .

Following the same principle for \mathcal{Q}_{299} :012345678ACEB9FD as

$$e = a + ad + cd, \quad f = b + ad + bc + cd, \quad g = c + bd + cd, \quad h = d, \quad (7)$$

we can define the component function $f_{\mathcal{Q}_{299}}^{i,j}(\langle a^i, b^i, c^i, d^i \rangle, \langle a^j, b^j, c^j, d^j \rangle) = \langle e, f, g, h \rangle$ as

$$\begin{aligned} e &= a^i + (a^i d^i + d^i a^j + a^i d^j) + (c^i d^i + d^i c^j + c^i d^j) \\ f &= b^i + (a^i d^i + d^i a^j + a^i d^j) + (b^i d^i + d^i b^j + b^i d^j) + (c^i d^i + d^i c^j + c^i d^j) \\ g &= c^i + (b^i d^i + d^i b^j + b^i d^j) + (c^i d^i + d^i c^j + c^i d^j) \\ h &= d^i. \end{aligned} \quad (8)$$

Similarly, three 4-bit output shares provided by $f_{\mathcal{Q}_{299}}^{2,3}(\cdot, \cdot)$, $f_{\mathcal{Q}_{299}}^{3,1}(\cdot, \cdot)$ and $f_{\mathcal{Q}_{299}}^{1,2}(\cdot, \cdot)$ make a uniform first-order TI of \mathcal{Q}_{299} .

Since the affine transformations A_1 , A_2 , A_3 do not change the uniformity and should be applied on each 4-bit share separately, the decomposition in Eq. (4) provides a 3-share uniform first-order TI of the PRESENT Sbox. It should be noted that registers are required to be placed between the component functions of \mathcal{Q}_{294} and \mathcal{Q}_{299} to avoid the propagation of the glitches (see Fig. 3). Note that the affine function A_2 can be freely placed before or after the intermediate register.

For the second-order TI representations in addition to the above expressed component functions, we define $g_{\mathcal{Q}_{294}}^{i,j}(\langle a^i, b^i, c^i, d^i \rangle, \langle a^j, b^j, c^j, d^j \rangle) = \langle e, f, g, h \rangle$ as

$$\begin{aligned} e &= d^i b^j + b^i d^j & g &= 0 \\ f &= d^i c^j + c^i d^j & h &= 0. \end{aligned} \quad (9)$$




Fig. 3. A first-order TI of the PRESENT Sbox: $S(\mathbf{x}) = \mathbf{y}$

The 4-bit output shares $\mathbf{y}^{i \in \{1, \dots, 10\}}$ are provided by

$$\begin{aligned}
 \mathbf{y}^1 &= f_{Q_{294}}^{2,3}(\cdot, \cdot), & \mathbf{y}^2 &= f_{Q_{294}}^{3,4}(\cdot, \cdot), & \mathbf{y}^3 &= f_{Q_{294}}^{4,5}(\cdot, \cdot), & \mathbf{y}^4 &= f_{Q_{294}}^{5,1}(\cdot, \cdot), \\
 \mathbf{y}^5 &= f_{Q_{294}}^{1,2}(\cdot, \cdot), & \mathbf{y}^6 &= g_{Q_{294}}^{2,4}(\cdot, \cdot), & \mathbf{y}^7 &= g_{Q_{294}}^{3,5}(\cdot, \cdot), & \mathbf{y}^8 &= g_{Q_{294}}^{1,4}(\cdot, \cdot), \\
 & & \mathbf{y}^9 &= g_{Q_{294}}^{2,5}(\cdot, \cdot), & \mathbf{y}^{10} &= g_{Q_{294}}^{1,3}(\cdot, \cdot). & &
 \end{aligned} \tag{10}$$

After a clock cycle, when $\mathbf{y}^{i \in \{1, \dots, 10\}}$ are stores in dedicate registers, the output shares should be combined as

$$\mathbf{z}^{i \in \{1, \dots, 5\}} = \mathbf{y}^i + \mathbf{y}^{i+5}, \tag{11}$$

which provides the uniform second-order TI of Q_{294} .

The same procedure is valid in case of Q_{299} considering the component function $g_{Q_{299}}^{i,j}(\langle a^i, b^i, c^i, d^i \rangle, \langle a^j, b^j, c^j, d^j \rangle) = \langle e, f, g, h \rangle$ as

$$\begin{aligned}
 e &= d^i a^j + d^i c^j + a^i d^j + c^i d^j \\
 f &= d^i a^j + d^i b^j + d^i c^j + a^i d^j + b^i d^j + c^i d^j \\
 g &= d^i b^j + d^i c^j + b^i d^j + c^i d^j \\
 h &= 0.
 \end{aligned} \tag{12}$$

By changing the indices from Q_{294} to Q_{299} in Eq. (10) and later applying the reduction in Eq. (11), a uniform second-order TI of Q_{299} is achieved. Hence by means of these component functions in addition to the affine transformations, we can realize a uniform second-order TI of the PRESENT Sbox. Figure 4 shows the graphical view of such a construction, and all the required formulas are given in the extended version of this article [35]. Note that the registers after the affine function A_2 can instead be place before A_2 right after the reduction from 10 to 5 shares.

3.4 Implementation

Based on the specifications given above and considering a Spartan-6 FPGA (indeed the XC6SLX75 of SAKURA-G [1]) we implemented six designs. The first three ones are different profiles of KATAN-32, and the next three designs realize the encryption of PRESENT with a round-based architecture. For each of the targeted cipher we implemented

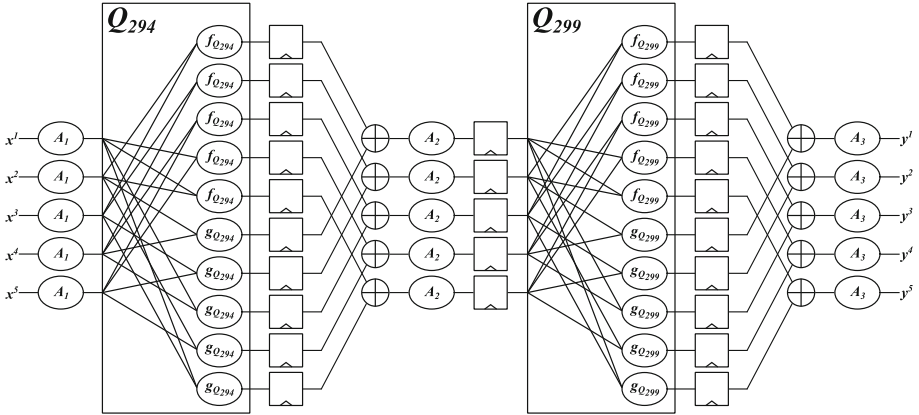


Fig. 4. A second-order TI of the PRESENT Sbox: $S(x) = y$

- the first-order TI, i.e., KATAN-1st and PRESENT-1st profiles,
- the second-order TI, i.e., KATAN-2nd and PRESENT-2nd profiles, and
- the first-order TI with GliFreD, i.e., KATAN-1st-G and PRESENT-1st-G profiles.

Although we did not consider any constraints on placement and routing of the four non-GliFreD profiles, following the principles of GliFreD the corresponding profiles have been realized by first defining an area on the target FPGA, where the component of the *true* part of the GliFreD circuit should be placed. After finishing the placement and routing, the corresponding dual circuit, i.e., the *false* part of the GliFreD circuit, has been cloned and dualized by means of the RapidSmith tool [22]. As a reference, the circuits shown in Fig. 1 are the normal and GliFreD realizations of the least significant bit e of Eq. (8).

Due to its serialized ring architecture, the KATAN-1st-G profile does not form a pipeline. The most important difference between such a profile and its original one (KATAN-1st) is on the one hand the number of required clock cycles to finish an encryption (i.e., latency) which is doubled and on the other hand the raised achievable clock frequency due to the minimal LUT depth. The max LUT depth in GliFreD circuits is 1, hence a very short critical path. However, the PRESENT-1st-G profile is implemented in a fully-pipelined way, so that the round-based architecture is able to hold 11 different cipher states. Hence, after $32 \times 11 \times 2 = 704$ clock cycles, 11 encryptions with the same key are performed. The pipelined architecture naturally increases the register utilization of the components but provides a much higher throughput.

Table 1 compares the overhead and performance of different design profiles. It indeed gives an overview on the disadvantage (area and time overheads) as well as the advantage (throughput) of employing GliFreD with respect to two different design architectures, i.e., a fully-serialized one which is register oriented (KATAN-1st-G) and a round-based one which is combinatorial oriented

Table 1. Details about the implemented profiles. The values given in this table are taken from the post route synthesis report of Xilinx ISE 14.7.

Profile	Resources		Frequency	Latency	Pipeline	Throughput
	LUT	FF	(MHz)	(#clock)	(stage)	(Mbit/s)
KATAN-1st	34	96	225.38	273	1	26.42
KATAN-2nd	65	180	321.54	273	1	37.69
KATAN-1st-G	114	548	438.21	546	1	25.68
PRESENT-1st	808	384	206.61	64	2	413.22
PRESENT-2nd	2245	1680	203.46	128	4	406.92
PRESENT-1st-G	5442	12672	458.09	704	11	458.09

(PRESENT-1st-G). As shown by Table 1, although the resource utilization and the latency of the GliFreD profiles are drastically increased, the throughput is still kept comparable with the original design profiles. Such achievements are mainly due to the naturally-minimized critical paths in the GliFreD designs allowing a high clock frequency.

4 Empirical Results

In addition to the performance and overhead figures given in Sect. 3.4, we practically examined the ability of each of our six developed designs to avoid side-channel leakages.

Setup. The experimental platform is a SAKURA-G [1] equipped with a Xilinx Spartan-6 FPGA. The side-channel leakages have been measured by collecting power consumption traces of the underlying FPGA by means of a Teledyne LeCroy HRO 66Zi digital oscilloscope at a sampling frequency of 500 MS/s and a limited bandwidth of 20 MHz. Due to the low peak-to-peak amplitude of the signals we also made use of the amplifier embedded on the SAKURA board. For all six design profiles, the target FPGA operated at a frequency of 24 MHz during the collection of the power traces. Our intuition on the measured power traces from our platform is that the traces are heavily filtered by the measurement setup including the shunt resistor, chip packaging, printed circuit board (PCB), and probes. Measuring the power traces with high bandwidth (> 20 MHz) leads to higher electrical noise. We have examined this behavior and observed leakages easier when the bandwidth is limited. Note that this intuition does not hold true in case of EM measurements.

It is noteworthy that such a frequency of operation has intentionally been taken in order to : i) cover the full power trace length in the measurements as the KATAN profiles need 254 clock cycles after data being loaded (respectively 508 for KATAN-1st-G), and ii) cause the power peaks of adjacent clock cycles slightly overlap each other. The later has been considered with respect to the note given in [45] that the second-order TI can still be vulnerable to a second-order bivariate

attack. Recalling the techniques introduced in [31], employing certain amplifiers or running the device at a high clock frequency leads to converting multivariate leakages to univariate. It has been shown in [49] that a second-order TI design actually can exhibit a univariate second-order leakage if the measurement setup is employed by certain components, e.g., DC blockers and/or amplifiers. Hence, operating the device at 24 MHz allows us to easily cover the long traces in the measurements and provide particular situations, where second-order TI profiles may demonstrate second-order leakage.

Evaluation. As the evaluation metric we employed the leakage assessment methodology of [17, 48] which is based on the Student’s t -test. The reason for such a choice is twofold. First, the t -test can examine the existence of detectable leakages without performing any key-recovery attack, which significantly eases the evaluation process particularly where higher-order leakages using millions of traces should be examined. Moreover, the efficiency of the state-of-the-art

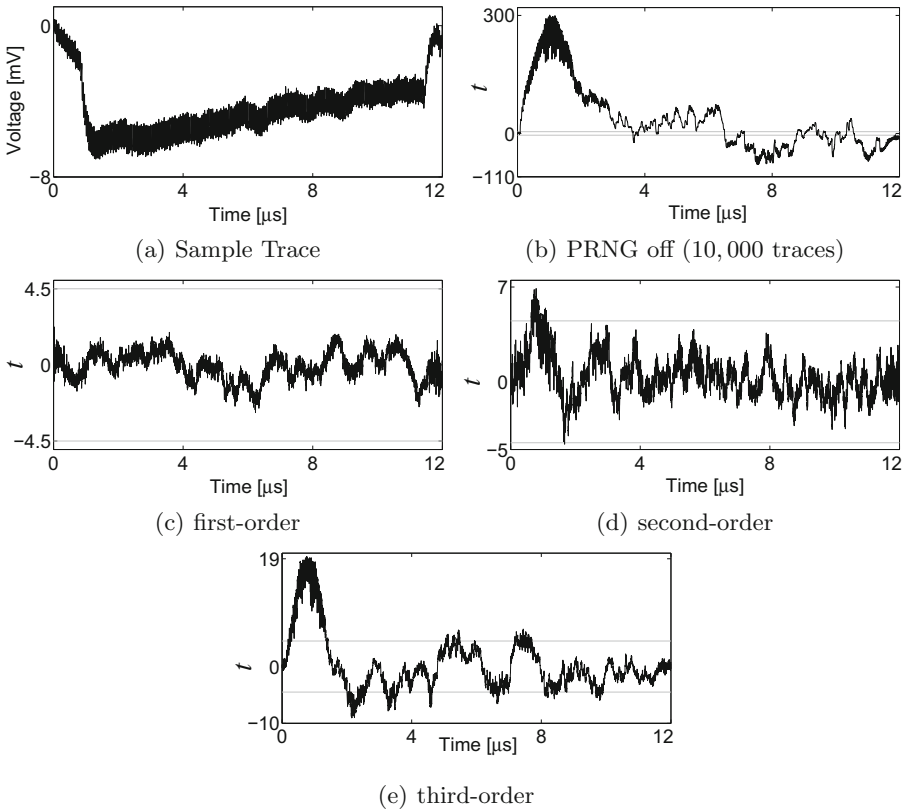


Fig. 5. KATAN-1st profile, sample trace and non-specific t -test results using 1,000,000 traces

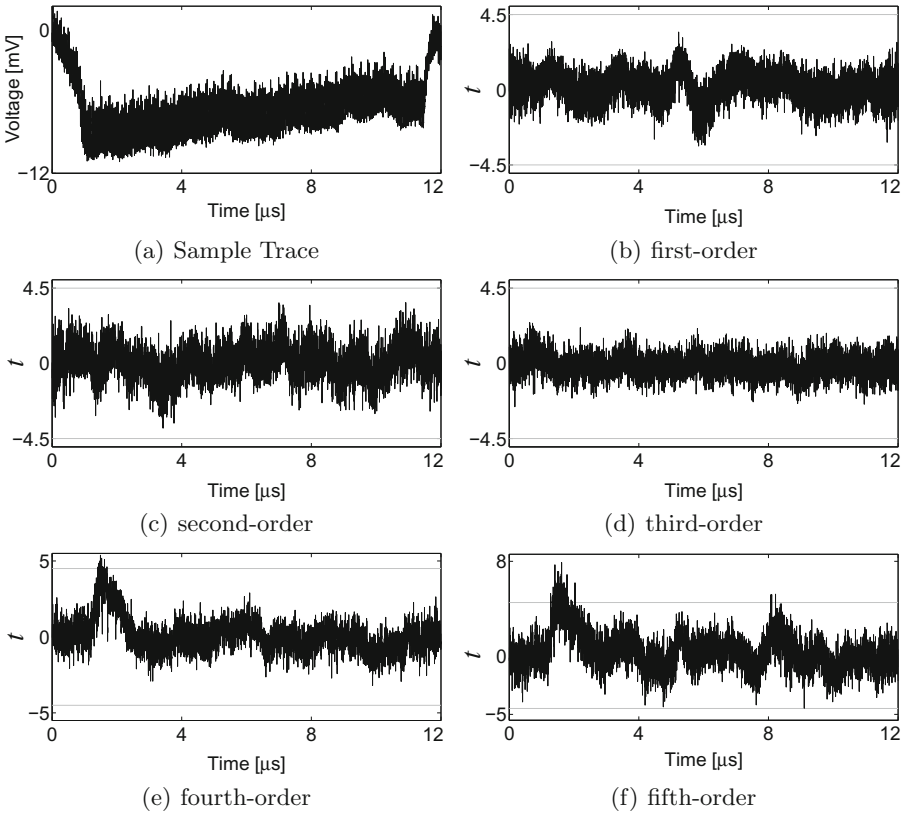


Fig. 6. KATAN-2nd profile, sample trace and non-specific t -test results using 100,000,000 traces

key-recovery attacks strongly depends on the targeted intermediate value and the underlying (power) model. Second, the same leakage assessment technique (more precisely the non-specific t -test also known as *fixed vs. random* test) has been used to examine the resistance of different threshold implementations (for example see [5, 49]). In order to keep our evaluations comparable with the former ones, we trivially employed the same evaluation method.

In a non-specific t -test the leakages associated to a fixed input (plaintext in case of encryption) are compared to that of random inputs while the key in all the measurements is kept constant. Such a test gives a level of confidence to conclude that the leakages related to the process of the fixed input are different to those of the random inputs. If so, an attack is expected to be feasible to exploit the leakage and recover the secrets. For more detailed information we refer the interested reader to [5, 17].

It is noteworthy that all the tests we performed here are based on a univariate scenario. In other words, we did not run any combination function on different sample points of each collected power trace. Further, we followed the same principle

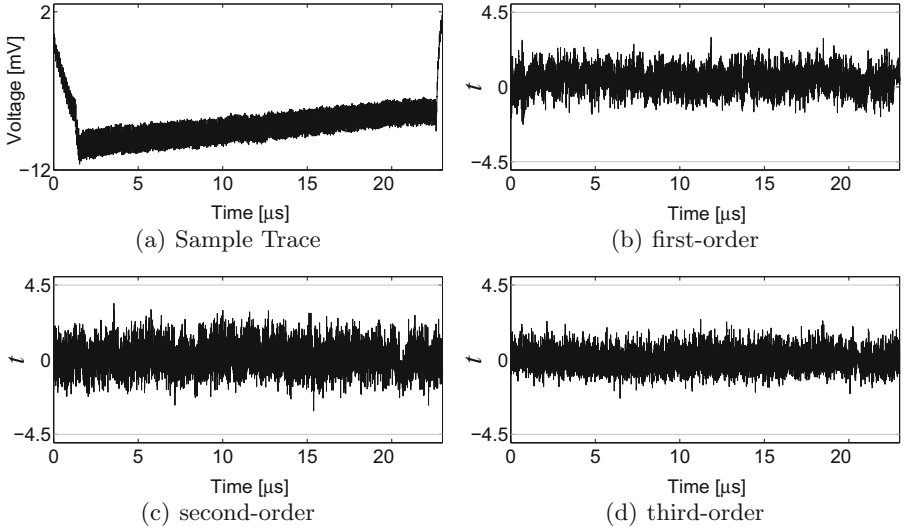


Fig. 7. KATAN-1st-G profile, sample trace and non-specific t -test results using 1,000,000,000 traces

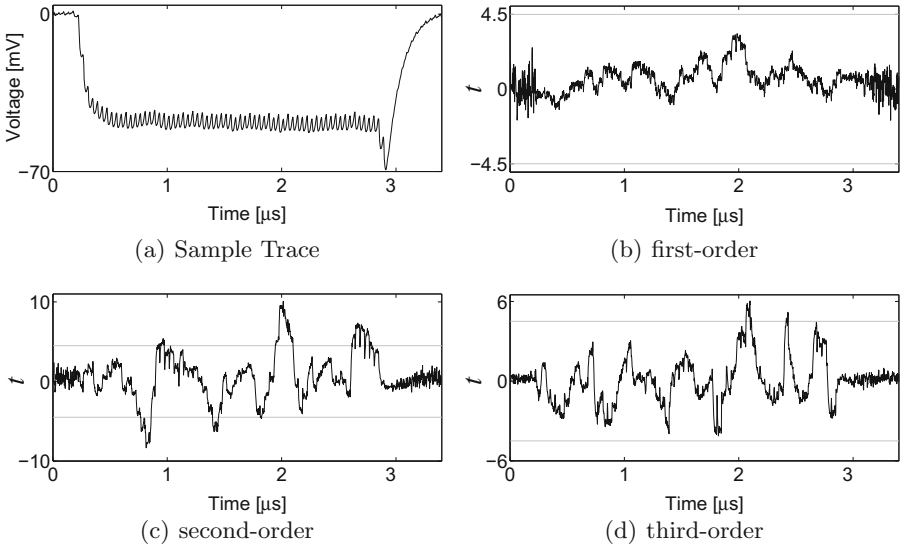


Fig. 8. PRESENT-1st profile, sample trace and non-specific t -test results using 10,000,000 traces

explained in [5, 48] to conduct the tests at higher orders. It means that we made the power traces mean-free squared (at each sample point independently), i.e., $(X - \mu)^2$ for the second-order evaluations, and standardized cubed, i.e., $\left(\frac{X - \mu}{\sigma}\right)^3$ for

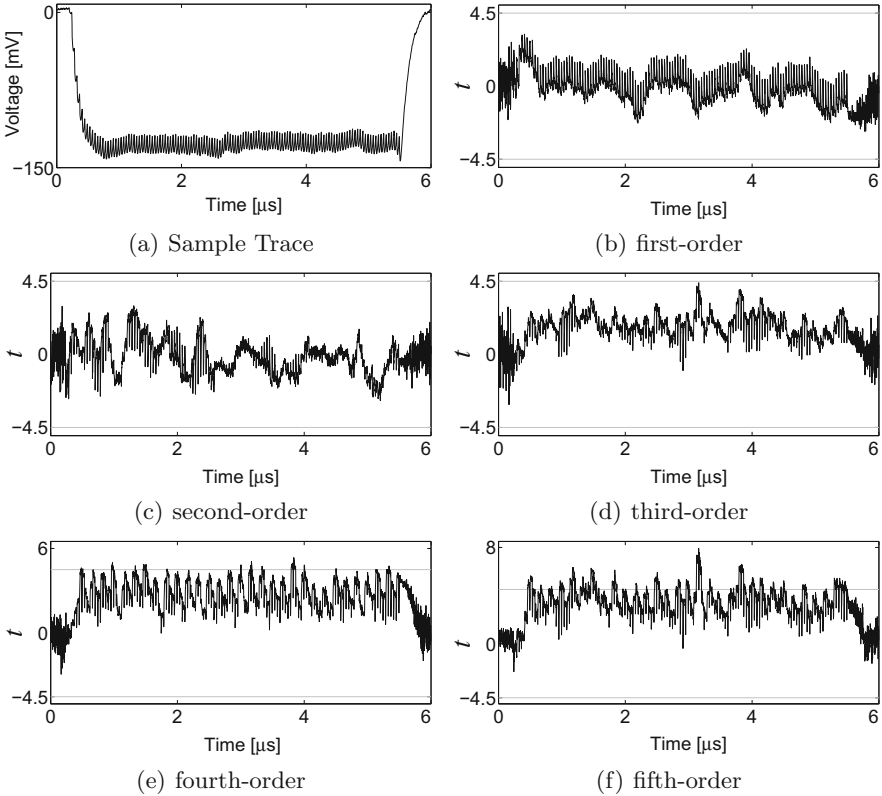


Fig. 9. PRESENT-2nd profile, sample trace and non-specific t -test results using 300,000,000 traces

the third-order evaluations. In general, the pre-processing is done by $\left(\frac{X - \mu}{\sigma}\right)^d$ for the analyses at order $d > 2$, with X as a random variable denoting the power traces (at a particular sample point), μ and σ^2 as the sample mean and sample variance (at the same sample point) respectively. Indeed, these pre-processes required for higher-order evaluations are with the respect to the centered and standardized higher-order statistical moments (for more information see [26,34]).

We start our evaluations with KATAN-1st profile. Figure 5(a) shows a corresponding sample power trace. Note that the collected power traces do not cover a time period, when plaintext and key are serially loaded into the shift registers. In order to have an overview about the quality of the measurement setup and verify the employed evaluation metric, for the first analysis we turned the PRNG off thereby forcing all masks to zero, used for sharing the plaintexts. As shown by Fig. 5(b), the first-order t -test shows clear detectable leakages using a few 10,000 traces. By keeping the PRNG active and conducting the same non-specific t -tests up to third-order using 1,000,000 traces we observed the curves shown by Fig. 5, which indeed confirm the first-order resistance and vulnerability at the second and third orders, as expected.

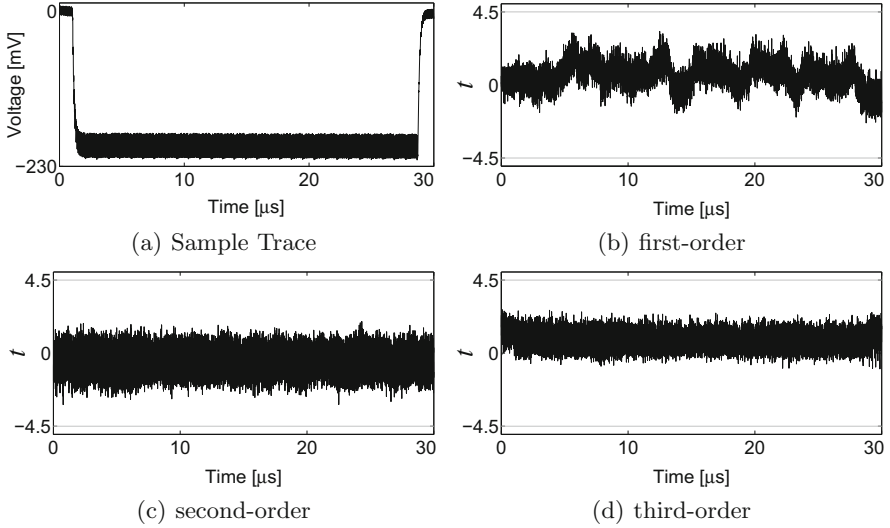


Fig. 10. PRESENT-1st-G profile, sample trace and non-specific t -test results using 1,000,000,000 traces

For the KATAN-2nd profile we had to collect much more traces to be able to observe the higher-order leakages. It is due to the high order of sharing, i.e., at least 5 shares (see Sect. 3.1) in case of a second-order TI. In fact, we observed the fourth- and fifth-order leakages using approximately 100,000,000 traces, as shown in Fig. 6. However, in order to examine the issue reported in [45] (by operating the target at 24 MHz) we continued the collection of the traces up to 500,000,000, but we have not observed any second-order leakage while the fourth- and fifth-order leakages became detectable – expectedly – with higher confidence. We should here refer to the issue addressed in [45] and the detectable second-order leakage reported in [49]. Based on the explanations of [45] a second-order bivariate leakage should be detectable, but such a bivariate leakage is not necessarily detectable from the consecutive clock cycles, that can additively be combined by means of an amplifier or running the device at a high clock frequency [31]. In case of the application of [49] apparently the consecutive clock cycles exhibit such a bivariate leakage, but it is not hold true for the serialized KATAN architecture. Further, compared to our design profiles the constructions in [49] make use of a kind of remasking which is a different methodology to ensure the uniformity.

Following the same scenario we performed the evaluations on the KATAN-1st-G profile and collected 1,000,000,000 traces to perform the same t -tests at up to third order. The corresponding results which are depicted in Fig. 8 indeed confirm the effectiveness of the underlying hiding technique to significantly harden the higher-order attacks. The result of this profile can be compared to that of the KATAN-1st profile (Fig. 5), where 1,000,000 traces are adequate to observe the second- and third-order leakages.

The same leakage assessment technique has been conducted on the three profiles of the round-based PRESENT architecture, and the corresponding results are

shown in Figs. 8, 9 and 10. For the PRESENT-1st profile we required 10,000,000 trace to observe the second- and third-order leakages. Respectively 300,000,000 traces were necessary for the PRESENT-2nd profile to exhibit fourth- and fifth-order leakages. We should again bring the reader's attention to the infeasibility to observe a second-order leakage from the PRESENT-2nd profile. We indeed continued our evaluations on this profile by measuring 1,000,000,000 traces as well as with different fixed inputs (with respect to the non-specific t -tests), but in none of the tests we observed a detectable second-order leakage. As an example, we give the results of one of such tests with 1,000,000,000 traces in the extended version of this article [35], where the third-order leakage also becomes detectable. Finally, similar to the KATAN GliFreD design we collected 1,000,000,000 traces and conducted the same non-specific t -tests on the PRESENT-1st-G profile, which still shows robustness to avoid the leakages to be detectable at first, second, and third orders.

Discussion. Comparing the presented practical results, at the first glance it can be noticed that the GliFreD profiles consume more energy than the other corresponding profiles. They also increase the number of required clock cycles (latency) particularly in case of the PRESENT design as its combinatorial circuit has a longer depth compared to the KATAN design. However, their achievement, i.e., hiding the higher-order leakages to make the higher-order attacks *practically* infeasible, is confirmed. Hence, it can be concluded that the combination of such a power-equalization technique and a proper masking scheme (i.e., first-order TI) gives a high level of confidence to argue the practical infeasibility of the key-recovery attacks.

Our comparisons are limited to the second-order TI of KATAN and PRESENT, which can be extended to higher-order TI designs. However, by increasing the desired order of security the number of shares and the required internal PRNGs respectively increase (e.g., at least 7 and 9 shares for third- and fourth-order TI). Note that the numbers given in Table 1 exclude the area required for the PRNGs.

Nonetheless, due to the local separation of false and true parts in GliFreD circuits, the resistance of our proposed method against higher-order EM attacks is still an open question and should be addressed in the future. Further, GliFreD is exclusively designed for FPGAs and uses the fixed LUT structure to realize Boolean functions of a circuit. Transforming this logic style naively to ASIC may not lead to the expected results especially with respect to the area overhead. The idea of combining TI with DPL styles can be adopted for ASICs by employing one of the logic styles designed for ASICs in addition to a customized router.

References

1. Side-channel AttackK User Reference Architecture. <http://satoh.cs.uec.ac.jp/SAKURA/index.html>
2. Balasch, J., Gierlichs, B., Verdult, R., Batina, L., Verbauwhede, I.: Power analysis of Atmel CryptoMemory – recovering keys from secure EEPROMs. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 19–34. Springer, Heidelberg (2012)

3. Bhasin, S., Guilley, S., Flament, F., Selmane, N., Danger, J.: Countering early evaluation: an approach towards robust dual-rail precharge logic. In: Workshop on Embedded Systems Security - WESS 2010, p. 6. ACM (2010)
4. Bilgin, B., Gierlichs, B., Nikova, S., Nikov, V., Rijmen, V.: A more efficient AES threshold implementation. In: Pointcheval, D., Vergnaud, D. (eds.) AFRICACRYPT. LNCS, vol. 8469, pp. 267–284. Springer, Heidelberg (2014)
5. Bilgin, B., Gierlichs, B., Nikova, S., Nikov, V., Rijmen, V.: Higher-order threshold implementations. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 326–343. Springer, Heidelberg (2014)
6. Bilgin, B., Nikova, S., Nikov, V., Rijmen, V., Stütz, G.: Threshold implementations of all 3×3 and 4×4 S-boxes. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 76–91. Springer, Heidelberg (2012)
7. Bilgin, B., Nikova, S., Nikov, V., Rijmen, V., Tokareva, N., Vitkup, V.: Threshold implementations of small S-boxes. *Cryptograph. Commun.* **7**(1), 3–33 (2015)
8. Biryukov, A., Cannière, C.D., Braeken, A., Preneel, B.: A Toolbox for Cryptanalysis: Linear and Affine Equivalence Algorithms. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 33–50. Springer, Heidelberg (2003)
9. Bogdanov, A.A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M., Seurin, Y., Vikkelsøe, C.: PRESENT: an ultra-lightweight block cipher. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
10. De Cannière, C., Dunkelman, O., Knežević, M.: KATAN and KTANTAN — A Family of Small and Efficient Hardware-Oriented Block Ciphers. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 272–288. Springer, Heidelberg (2009)
11. Canright, D., Batina, L.: A very compact “perfectly masked” S-box for AES. In: Bellare, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 446–459. Springer, Heidelberg (2008)
12. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (1999)
13. Chen, Z., Zhou, Y.: Dual-rail random switching logic: a countermeasure to reduce side channel leakage. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 242–254. Springer, Heidelberg (2006)
14. Coron, J.-S., Kizhvatov, I.: An efficient method for random delay generation in embedded software. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 156–170. Springer, Heidelberg (2009)
15. Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: from probing attacks to noisy leakage. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 423–440. Springer, Heidelberg (2014)
16. Eisenbarth, T., Kasper, T., Moradi, A., Paar, C., Salmasizadeh, M., Shalmani, M.T.M.: On the power of power analysis in the real world: a complete break of the KEELOQ code hopping scheme. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 203–220. Springer, Heidelberg (2008)
17. Goodwill, G., Jun, B., Jaffe, J., Rohatgi, P.: A testing methodology for side channel resistance validation. In: NIST Non-invasive Attack Testing Workshop (2011). http://csrc.nist.gov/news_events/non-invasive-attack-testing-workshop/papers/08-Goodwill.pdf
18. Güneysu, T., Moradi, A.: Generic side-channel countermeasures for reconfigurable devices. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 33–48. Springer, Heidelberg (2011)

19. He, W., de la Torre, E., Riesgo, T.: A Precharge-absorbed DPL logic for reducing early propagation effects on FPGA implementations. In: Reconfigurable Computing and FPGAs - ReConFig 2011, pp. 217–222. IEEE Computer Society (2011)
20. He, W., Otero, A., de la Torre, E., Riesgo, T.: Automatic generation of identical routing pairs for FPGA implemented DPL logic. In: Reconfigurable Computing and FPGAs - ReConFig 2012, pp. 1–6. IEEE Computer Society (2012)
21. Kaps, J., Velegalati, R.: DPA resistant AES on FPGA using partial DDL. In: Field-Programmable Custom Computing Machines - FCCM 2010, pp. 273–280. IEEE Computer Society (2010)
22. Lavin, C., Padilla, M., Lamprecht, J., Lundrigan, P., Nelson, B., Hutchings, B., Wirthlin, M.: RapidSmith - a library for low-level manipulation of partially placed-and-routed FPGA designs. Technical report, Brigham Young University, September 2012
23. Lomné, V., Maurine, P., Torres, L., Robert, M., Soares, R., Calazans, N.: Evaluation on FPGA of triple rail logic robustness against DPA and DEMA. In: Design, Automation and Test in Europe - DATE 2009, pp. 634–639. IEEE Computer Society (2009)
24. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards. Springer, Heidelberg (2007)
25. Mangard, S., Pramstaller, N., Oswald, E.: Successfully attacking masked AES hardware implementations. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 157–171. Springer, Heidelberg (2005)
26. Moradi, A.: Statistical tools flavor side-channel collision attacks. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 428–445. Springer, Heidelberg (2012)
27. Moradi, A., Barengi, A., Kasper, T., Paar, C.: On the vulnerability of FPGA bitstream encryption against power analysis attacks: extracting keys from xilinx Virtex-II FPGAs. In: ACM Conference on Computer and Communications Security - CCS 2011, pp. 111–124. ACM (2011)
28. Moradi, A., Eisenbarth, T., Poschmann, A., Paar, C.: Power analysis of single-rail storage elements as used in MDPL. In: Lee, D., Hong, S. (eds.) ICISC 2009. LNCS, vol. 5984, pp. 146–160. Springer, Heidelberg (2010)
29. Moradi, A., Immler, V.: Early propagation and imbalanced routing, how to diminish in FPGAs. In: Batina, L., Robshaw, M. (eds.) CHES 2014. LNCS, vol. 8731, pp. 598–615. Springer, Heidelberg (2014)
30. Moradi, A., Mischke, O.: Glitch-free implementation of masking in modern FPGAs. In: Hardware-Oriented Security and Trust - HOST 2012, pp. 89–95. IEEE (2012)
31. Moradi, A., Mischke, O.: On the simplicity of converting leakages from multivariate to univariate. In: Bertoni, G., Coron, J.-S. (eds.) CHES 2013. LNCS, vol. 8086, pp. 1–20. Springer, Heidelberg (2013)
32. Moradi, A., Mischke, O., Eisenbarth, T.: Correlation-enhanced power analysis collision attack. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 125–139. Springer, Heidelberg (2010)
33. Moradi, A., Poschmann, A., Ling, S., Paar, C., Wang, H.: Pushing the limits: a very compact and a threshold implementation of AES. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 69–88. Springer, Heidelberg (2011)
34. Moradi, A., Standaert, F.-X.: Moments-correlating DPA. Cryptology ePrint Archive, Report 2014/409 (2014). <http://eprint.iacr.org/>
35. Moradi, A., Wild, A.: Assessment of hiding the higher-order leakages in hardware - what are the achievements versus overheads? Cryptology ePrint Archive (2015). <http://eprint.iacr.org/>

36. Nassar, M., Bhasin, S., Danger, J., Duc, G., Guilley, S.: BCDL: a high speed balanced DPL for FPGA with global precharge and no early evaluation. In: Design, Automation and Test in Europe - DATE 2010, pp. 849–854. IEEE Computer Society (2010)
37. Nikova, S., Rijmen, V., Schläffer, M.: Secure hardware implementation of non-linear functions in the presence of glitches. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 2008. LNCS, vol. 5461, pp. 218–234. Springer, Heidelberg (2009)
38. Nikova, S., Rijmen, V., Schläffer, M.: Secure hardware implementation of nonlinear functions in the presence of glitches. *J. Cryptol.* **24**(2), 292–321 (2011)
39. Oswald, E., Mangard, S., Pramstaller, N., Rijmen, V.: A side-channel analysis resistant description of the AES S-box. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 413–423. Springer, Heidelberg (2005)
40. Popp, T., Kirschbaum, M., Zefferer, T., Mangard, S.: Evaluation of the masked logic style MDPL on a prototype chip. In: Paillier, P., Verbauwhe, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 81–94. Springer, Heidelberg (2007)
41. Popp, T., Mangard, S.: Masked dual-rail pre-charge logic: DPA-resistance without routing constraints. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 172–186. Springer, Heidelberg (2005)
42. Poschmann, A., Moradi, A., Khoo, K., Lim, C., Wang, H., Ling, S.: Side-channel resistant crypto for less than 2, 300 GE. *J. Cryptol.* **24**(2), 322–345 (2011)
43. Prouff, E., Rivain, M., Bevan, R.: Statistical analysis of second order differential power analysis. *IEEE Trans. Comput.* **58**(6), 799–811 (2009)
44. Rao, J.R., Rohatgi, P., Scherzer, H., Tinguely, S.: Partitioning attacks: or how to rapidly clone some GSM cards. In: IEEE Symposium on Security and Privacy, pp. 31–41. IEEE Computer Society (2002)
45. Reparaz, O.: A note on the security of higher-order threshold implementations. Cryptology ePrint Archive, Report 2015/001 (2015). <http://eprint.iacr.org/>
46. Rivain, M., Prouff, E.: Provably secure higher-order masking of AES. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 413–427. Springer, Heidelberg (2010)
47. Sauvage, L., Nassar, M., Guilley, S., Flament, F., Danger, J., Mathieu, Y.: DPL on Stratix II FPGA: what to expect? In: Reconfigurable Computing and FPGAs - ReConFig 2009, pp. 243–248. IEEE Computer Society (2009)
48. Schneider, T., Moradi, A.: Leakage assessment methodology - a clear roadmap for side-channel evaluations. In: Güneysu, T., Handschuh, H. (eds.) CHES 2015. LNCS, vol. 9293, pp. xx–yy. Springer, Heidelberg (2015)
49. Schneider, T., Moradi, A., Güneysu, T.: Arithmetic addition over boolean masking - towards first- and second-order resistance in hardware. In: Malkin, T., Kolesnikov, V., Lewko, A.B., Polychronakis, M. (eds.) ACNS 2015. LNCS, vol. 9092, pp. 517–536. Springer, Heidelberg (2015)
50. Suzuki, D., Saeki, M.: Security evaluation of DPA countermeasures using dual-rail pre-charge logic style. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 255–269. Springer, Heidelberg (2006)
51. Tiri, K., Akmal, M., Verbauwhe, I.: A dynamic and differential CMOS logic with signal independent power consumption to withstand differential power analysis on smart cards. *ESSCIRC* **2002**, 403–406 (2002)
52. Tiri, K., Hwang, D., Hodjat, A., Lai, B.-C., Yang, S., Schaumont, P., Verbauwhe, I.: Prototype IC with WDDL and differential routing – DPA resistance assessment. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 354–365. Springer, Heidelberg (2005)

53. Tiri, K., Verbaauwhede, I.: A logic level design methodology for a secure DPA resistant ASIC or FPGA implementation. In *Design, Automation and Test in Europe - DATE 2004*, pp. 246–251. IEEE Computer Society (2004)
54. Veyrat-Charvillon, N., Medwed, M., Kerckhof, S., Standaert, F.-X.: Shuffling against side-channel attacks: a comprehensive study with cautionary note. In: Wang, X., Sako, K. (eds.) *ASIACRYPT 2012*. LNCS, vol. 7658, pp. 740–757. Springer, Heidelberg (2012)
55. Wild, A., Moradi, A., Güneysu, T.: Evaluating the duplication of dual-rail precharge logics on FPGAs. In: Mangard, S., Poschmann, A.Y. (eds.) *COSADE 2015*. LNCS, vol. 9064, pp. 81–94. Springer, Heidelberg (2015)
56. Wild, A., Moradi, A., Güneysu, T.: GliFreD: glitch-free duplication - towards power-equalized circuits on FPGAs. *Cryptology ePrint Archive, Report 2015/124* (2015). <http://eprint.iacr.org/>
57. Yu, P., Schaumont, P.: Secure FPGA circuits using controlled placement and routing. In: *Hardware/Software Codesign and System Synthesis - CODES+ISSS 2007*, pp. 45–50 (2007)
58. Zhou, Y., Yu, Y., Standaert, F.-X., Quisquater, J.-J.: On the need of physical security for small embedded devices: a case study with COMP128-1 implementations in SIM cards. In: Sadeghi, A.-R. (ed.) *FC 2013*. LNCS, vol. 7859, pp. 230–238. Springer, Heidelberg (2013)