



HAL
open science

High-Speed Ring Oscillator based Sensors for Remote Side-Channel Attacks on FPGAs

Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia, Philippe Loubet-Moundi

► **To cite this version:**

Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia, Philippe Loubet-Moundi. High-Speed Ring Oscillator based Sensors for Remote Side-Channel Attacks on FPGAs. 2019 International Conference on Reconfigurable Computing and FPGAs (ReConFig 2019), Dec 2019, Cancun, Mexico. hal-02481050

HAL Id: hal-02481050

<https://hal.science/hal-02481050v1>

Submitted on 17 Feb 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

High-Speed Ring Oscillator based Sensors for Remote Side-Channel Attacks on FPGAs

Joseph Gravelier and Jean-Max Dutertre
Mines Saint-Etienne, CEA-Tech, Centre CMP
F - 13541 Gardanne France
{joseph.gravelier, dutertre}@emse.fr

Yannick Teglia and Philippe Loubet-Moundi
Thales
La Ciotat, France
{yannick.tegla, philippe.loubet-moundi}@thalesgroup.com

Abstract - FPGAs have been widely adopted in cloud datacenters and System-on-Chip for hardware acceleration purposes during the past few years. For flexibility and efficiency reasons, cloud FPGA fabrics are likely to be shared between multiple users. Despite the logical isolation suggested to protect each tenant, multi-user FPGA environment raises crucial questions about the potential security threats that it may represent. Recently, a series of papers demonstrated that a malicious user could be able to use its rented logic to perform remote side-channel and fault attacks on other user assets located inside the fabric or in the surrounding chips. In this paper, we present a novel implementation method for ring oscillator based voltage sensors that enables runtime supply voltage fluctuation measurement. Considering a multi-user FPGA cloud scenario, we evaluate our sensor performances for side-channel purposes by performing CPA attacks against a hardware AES module instantiated within the same FPGA fabric. Then, we compare our results with existing voltage sensors and also demonstrate that, when calibrated, our sensors can provide results similar to traditional electromagnetic side-channel setups.

Keywords - FPGA, ring-oscillator, time-to-digital converter, voltage sensing, remote attacks, side-channel attacks.

I. INTRODUCTION

As size and performance of FPGAs continuously increase they continue to gain interest for hardware acceleration purposes. Reconfigurable logic allows designers to implement specialized applications without manufacturing dedicated chips. This reduces drastically the expenses needed to create custom designs and decreases the time-to-market. For all these reasons FPGAs are increasingly used as end-product from test and measurement electronics to medical and aeronautic purposes. Cloud providers such as Amazon EC2 [1] and Alibaba F3 [2] recently deployed FPGA instances in large scale datacenters. These services allow users to rent logic resources for big data analytics, inference and video processing. Remote access to FPGAs in the cloud raises concerns about the potential associated security threats. The possibility that different users could get access to the same FPGA fabric was

discussed in several design articles and could prevail in the near future [3, 4]. Despite the logical isolation suggested to protect each logic block from the others, recent papers warn the community about the multi-tenant threat. A malicious user could try to take advantage of his configurable resources to eavesdrop or disturb calculations from another user. These exploits take advantage of the FPGAs programmable logic to induce remote power glitches [5] or perform remote side-channel attacks against surrounding users [6, 7, 8]. Remote FPGA-based attacks follow a new trend that has been initiated by Rowhammer [9] and ClkScrew [10] exploits which consist in software induced hardware attacks. With the continuous development of cloud services and progressive dematerialization of the computing resources, it becomes mandatory to question the hardware security provided by those remote systems. Although similar to traditional hardware attacks, software induced hardware attacks do not require either physical access or specific equipment as probes and oscilloscopes. They take advantage of the resources provided by the targeted devices and can be launched at any time and place through a network. In the proposed attack scenario, the FPGA provides enough flexibility and performance to replicate a complete side-channel attack bench. The voltage fluctuation inside the fabric can be precisely estimated by designing propagation delay sensors such as Ring Oscillator (RO) based sensors [11] or Time-to-Digital Converters (TDC) based sensors [12]. This research work aims to improve the existing sensors designs deployed to monitor power supply fluctuations inside the FPGA fabric and to push forward the state-of-the-art on FPGA-based side-channel attacks. Our major contributions are detailed below:

- A new design approach for RO-based sensors that enables nanosecond scale measurement of FPGA internal voltage.
- The usage of our RO-based sensors in a multi-user FPGA scenario that experimentally demonstrates their ability to perform Correlation Power Analysis (CPA) against a 50 MHz AES hardware module.
- The comparison of different FPGA-based sensors (RO-based sensors and TDC-based sensors) with traditional side-channel methods (electromagnetic side-channel).
- The demonstration that, despite the modest quantification level and sampling frequency achievable using FPGA-

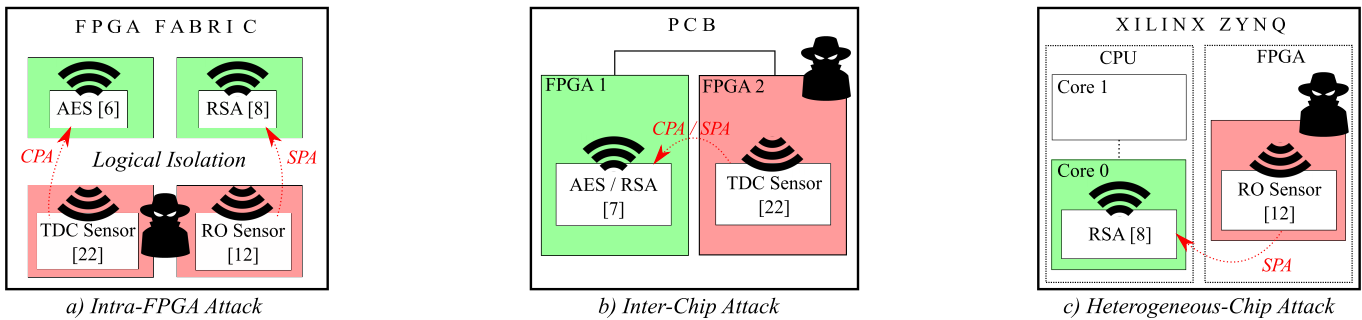


Fig. 1: Overview of FPGA-based Power Side-Channel Exploits

based sensors, proximity, flexibility and configuration allow them to provide side-channel results similar to traditional measurement setups.

Section II provides an overview of the previous FPGA remote side-channel successful exploits. The background about FPGA-based voltage fluctuation measurement is described in section III. Section IV introduces a new design approach for RO-based sensors in FPGAs. In Section V we present the experimental setup and the obtained CPA results. Then, section VI provides a comparison between FPGA-based sensors and traditional electromagnetic side-channel. Section VII concludes this paper.

II. REMOTE POWER SIDE-CHANNEL OVERVIEW

This section introduces the concept of remote side-channel attacks and provides a state-of-the-art of the previous FPGA-based power side-channel exploits. Then, the threat model assumed for our experiments is described.

A. Side-Channel Attacks

Side-channel attacks make use of the transistors switching activity leakage through voltage variations, electromagnetic emanations and other physical effects to collect information about the processes running inside a target device. Thanks to a correlation between the leakage and the data processed, a side-channel attack can be performed to retrieve cryptographic keys and secrets from a target without tampering it. Traditionally, side-channel setups rely on voltage or electromagnetic probes and oscilloscopes to monitor the side-channel leakage. By analysing the collected traces, an attacker can visually speculate on the different instructions performed by the device using Simple Power Analysis (SPA [13]). Statistical side-channel methods such as Differential Power Analysis (DPA [14]) or Correlation Power Analysis (CPA [15]) allow an attacker to infer the secret keys of cryptographic processes by correlating guessed leakage hypotheses with a set of experimental traces.

FPGA-based power analysis increases the threat that side-channel attacks represent as it doesn't require either direct physical access to the target or specific equipment. A remote side-channel attack can be conducted using on-chip digital voltage sensors (ROs & TDC-based sensors) implemented within the fabric. In this work, a CPA attack is carried out against an AES hardware module. The results are presented in Section V and VI.

B. Related Work

Although being all based on FPGA sensors, previous FPGA-based side-channel attacks have been conducted under three different scenarios as illustrated in figure 1:

1) Intra-FPGA Attack: Remote side-channel attacks on FPGAs were introduced in [6]. The adversary model consists in a FPGA fabric shared among multiple users. Each user is protected from the others by logical isolation. Despite this protection, a malicious user can implement voltage sensors in his rented logic to monitor voltage fluctuations induced by surrounding computations. Assuming this model, the adversary is able to perform a CPA attack against a victim AES hardware module (see Fig. 1.a). A second exploit uses RO-based sensors to perform intra-chip SPA against a RSA hardware module [8].

2) Inter-Chip Attack: The Inter-Chip Side-channel Attack illustrated in figure 1.b goes a step further by proving that an untrusted chip inside a PCB can sense voltage variations induced by other chips through the power distribution network (PDN). In this exploit, an adversary FPGA is able to perform a CPA attack against an AES module and a SPA attack against a RSA module running on another FPGA fabric [7].

3) Heterogeneous Chip Attack: Xilinx Zynq technology integrates a dual core ARM processor and a FPGA fabric within the same SoC. In [8], malicious ROs were implemented in the FPGA fabric to perform a SPA against a RSA algorithm running on a linux OS inside the ARM CPU core as shown in figure 1.c.




Fig. 2: Threat Model – A FPGA fabric is divided into multiple user shells. Logical isolation between each tenant is provided by fences but a side-channel attack can be conducted thanks to the information leakage that propagates throughout the PDN.




Fig. 3: TDC and RO based sensors functional schematics

C. Threat Model

Because of limitations of their achievable resolution and sampling frequency, RO-based sensors were only used to carry out SPA attacks [8]. This paper provides a new design method for RO-based sensors that improves their performances and enable their use for statistical side-channel attacks against symmetric encryption algorithms. The adopted approach is to perform an intra-FPGA CPA side-channel attack as previously achieved with TDC-based sensors in [6]. Our threat model is illustrated in figure 2, it assumes a cloud scenario where a FPGA fabric is shared between multiple users. Each tenant is confined in its rented shell and cannot access other users logic. Fences are provided by logical isolation to make sure that no illicit communication can be established between users [4]. Despite these protections, a malicious user can legally and technically implement our RO-based sensors to sense supply voltage variations through the PDN of the FPGA fabric. The victim shell contains an AES which continually encrypts data. The power consumption leakage resulting from each AES encryption is acquired by the adversary shell and later exported for CPA computations and AES key retrieval.

III. SENSING FPGAS VOLTAGE FLUCTUATIONS

In this section, we address the mechanisms that enable FPGA voltage variations monitoring from the origin of power supply fluctuations to their measurement using reconfigurable logic.

A. Power Supply Fluctuations

Power supply fluctuations inside a chip are induced by its transistors switching activity. The quantity of current drawn depends on the resources required for the computation.

The PDN gathers all the circuitry dedicated to the power supply, it has to provide a low-noise stable voltage capable to handle current fluctuations. Despite the optimization of the PDNs, transient voltage ripples induced by the interaction between the current drawn and the RLC parasitic component forming the PDN cannot be fully controlled [16, 17]. These can affect performance and cause timing glitch errors on critical logic paths. Power supply fluctuation leakage through the PDN carries the footprint of the running computation and can be used for side-channel purpose. It can be measured using an oscilloscope connected to the power pads of the target or internally monitored by on-chip sensors that take advantage of its effect on logic propagation delays.

B. Effect on Logic Propagation Delays

The propagation delay is the time required for a signal to propagate through a logic gate. Power supply, temperature and capacitive effects play a part in the propagation delay equation [18]. While capacitive load is fixed and temperature can be held relatively stable over the time, voltage fluctuations induce runtime propagation delay variations. At runtime, a sudden under-powering caused by transistors switching activity will induce an increase of the propagation delay throughout the chip. An over-powering will produce the exact opposite. Hence, measuring propagation delays provides an accurate estimation of the chip's internal power supply voltage. Two major propagation delay sensors are commonly used for power monitoring: the RO-based sensor [11] and the TDC-based sensor [12].

C. Delay Sensors: Time-To-Digital Converter

The TDC-based sensor illustrated in figure 3.a converts timing variations induced by power supply fluctuations into digital information. Thanks to a low-cost design and a fine-grained resolution TDC-based sensors are commonly adopted as on-chip temperature and voltage sensors for operating control [19, 20] as well as glitch attack detection [11, 21]. More recently, with the arising of FPGA cloud services, some researchers started to use it to perform power side-channel attacks [6, 7]. In the following experiments, the TDC-based sensor proposed in [22] is adopted to evaluate and compare the performances provided by our RO-based sensor.

D. Delay Sensors: Ring Oscillator based Sensor

The RO-based delay sensor (see fig 3.b) monitors propagation delay fluctuations through the measurement of its RO oscillation frequency f_{RO} . A RO is a device composed of an odd number of cascaded inverters. The output of the last inverter is fed back to the first creating an infinite oscillation between two voltage levels. The oscillation frequency f_{RO} is defined by the number n of inverters in the ring, each inverter slows down the oscillation because of its internal propagation delay t_p . This results in the following equation: $f_{RO} = \frac{1}{2t_p n}$. Therefore, measuring the frequency variations of the RO indicates the propagation time fluctuations of its inverters. Hence, it provides an image of the power supply consumption. To enable the measurement of the RO oscillation frequency, designers commonly adopt digital counters [8, 23]. A counter C_{RO} is connected to the RO output: it is incremented by the RO oscillations and is read out by a register at a fixed sampling

frequency f_s . ΔC_{RO} represents the number of RO oscillations counted during a period. The equation that converts the counter value to the RO frequency f_{RO} is:

$$f_{RO}(t) = \underbrace{[C_{RO}(t) - C_{RO}(t-1) + \epsilon]}_{\Delta C_{RO}} * f_s \quad (1)$$

IV. A NOVEL RO-BASED SENSOR DESIGN

This section introduces a new design for RO-based sensors. To begin with, we address the limitations that designers encounter when using traditional RO-based sensors.

A. RO-based Sensors Downsides

Several RO-based sensors downsides recently pushed designers to adopt TDC-based sensors for side-channel purposes instead [12, 22]. This mainly comes from the fact that the RO-based sensors struggle to provide reliable measurements when their sampling rate exceeds 8 MS/s [8, 23, 12]. This limitation comes from three major factors:

1) *Frequency Dependant Resolution*: The resolution of the RO-based sensor relies on the number of oscillations counted during a sampling period. When a long-sampling period is adopted, a large number of oscillations is counted and a fine-grained image of the voltage level can be retrieved through the capture of the counter value. A decrease of the sampling period reduces the number of RO oscillations counted and limits the relationship between the counter value and the actual voltage level. Therefore, decreasing the sampling period gradually deteriorates the sensor resolution.

2) *Quantization Error*: The quantization error ϵ (see Eq. 1) is a distortion of the RO-counter value by 1 that sometimes occurs because of the absence of a phase relationship between f_{RO} and f_s [8]. At high sampling frequency the quantization error is significant over the total number of counter increments. Thus, its occurrence can skew the overall measurement.

3) *Counter Timing Error*: RO should oscillate as fast as possible to enable high sampling frequency measurements. However, at GHz frequency range, timing errors occur if the counter fed by the RO is not optimized for high frequency transitions. This results in the sampling of inconsistent counter values that further alter the RO-based sensor reliability. This counter limitation is discussed in [12, 22] but neither improvements or new designs are suggested.

B. Designing a high frequency RO-based Sensor

The new RO-based sensor design presented in this section is depicted in figure 4. It still consists in three blocks: a RO, a counter and a sampling register. By introducing this sensor, we aim to mitigate the impact of the three main RO limitations detailed in subsection IV-A.

1) *A Faster RO*: Because we are working with two clock sources, our design will suffer from phase shift quantization error. To mitigate this effect and to increase the resolution of our sensor, we implement the fastest RO achievable with the available logic. It consists in only one LUT configured to perform a NAND operation whose output is fed back




Fig. 4: Schematic of the proposed RO-based sensor design. The RO consists in a looped NAND which cadences the JRC. A register reads out the JRC at a fixed rate defined by clk .

Resource	Number	Usage
LUT1	1	Ring Oscillator
FF	8	Johnson Counter
FF	8	Sampling Register
Slices	2	

TABLE I: Resource utilization for 1 RO-based sensor instance.

to its input (see Fig. 4). The resulting oscillation frequency approximately reaches 1.2 GHz.

2) *An Optimized Counter*: To preserve our counter from timing errors caused by the RO speed, we choose a non-binary counter called Johnson Ring Counter (JRC) which only consists in cascaded flip-flops (FF). By adopting a design that doesn't require combinational logic to be inserted between flip-flops, we mitigate timing errors that binary counters would encounter when cadenced by GHz range signals. The sensor structure is depicted in figure 4. The clock input of each FF is connected to the output of the RO. The data path consists in a ring in which the complementary output of the last FF \bar{Q} is fed back to the data input D of the first one. Using 8 FFs the JRC provides 16 distinct states which is enough when the sampling period is smaller than 16 times the RO period. In the hypothetical case where a lower sampling frequency would be required, the number of FFs forming the counter should be increased or an additional binary counter should be used to determine the number of JRC overflows during a sampling period.

3) *A Lighter Design*: Our RO-based sensor instance consumes only 2 slices as detailed in table I. Such a small design can be spread throughout the fabric without area congestion. Thus, the area coverage and more importantly the overall resolution of the voltage sensor can be improved. This statement is discussed in the following part.

C. Number of RO-based sensors

Equation (2) expresses the counter value, ΔC_{RO} , as a function of the sampling frequency f_s and the RO frequency f_{RO} . The RO frequency f_{RO} is splitted in two terms: a constant one $f_{RO(natural)}$ that represents the steady state natural oscillating frequency of the RO and a dynamic one $f_{RO(dynamic)}$ which depends on the activity of the surrounding logic.

$$\Delta C_{RO} = \frac{f_{RO(natural)} + f_{RO(dynamic)}}{f_s} \quad (2)$$




Fig. 5: Effect of the number of RO-based sensors on the overall resolution. The depicted signal is a single trace of an AES encryption running at 10 MHz.

When several sensors are instantiated within the fabric, their contribution ΔC_{RO} is summed and averaged over the number of RO-based sensors used. Multiplying the number of RO-based sensors throughout the chip has several benefits. Firstly, because of process and routing paths variations, each RO has a specific phase and frequency. For this reason, the quantization error ϵ only affects a portion of the sensors simultaneously. When the number of RO-based sensor used increases, the quantization error progressively loses significance over the global voltage fluctuation measurement. Therefore, it has less impact regarding the accuracy of the overall sensor. Secondly, the natural frequency deviation between each RO instance enhances the granularity of our sensor. Depending on the value of $f_{RO(natural)}$, the dynamic frequency fluctuation $f_{RO(dynamic)}$ required to modify the counter value fluctuates. Therefore, each RO-based sensor instance provides a specific contribution that further enriches the overall resolution. Figure 5 illustrates the effect of the number of RO-instances on the sensor resolution. A single 10 MHz AES encryption is captured using 1, 16 and 64 RO-based sensors cadenced at a 250 MS/s sampling rate (*our experimental setup will be discussed in section V*). When only 1 RO is used, the quantization error effect is maximal and the $f_{RO(dynamic)}$ fluctuation only provides 3 distinct quantization levels (figure 5 - 1 RO). Thereby, the AES encryption is not visible in the obtained waveform. However, increasing the number of RO gradually leads to the appearance of the AES over the residual and quantification noise. Using 64 ROs the 10 AES rounds are clearly visible.

D. Place and Route Influence

Manual place and route is not required to enable RO-based voltage fluctuation measurement. However, when it is possible, designers can fix the placement and routing paths using relative placement macro (RPM) to improve RO-based sensor performances and get a better control of the RO-based sensors distribution throughout the fabric.

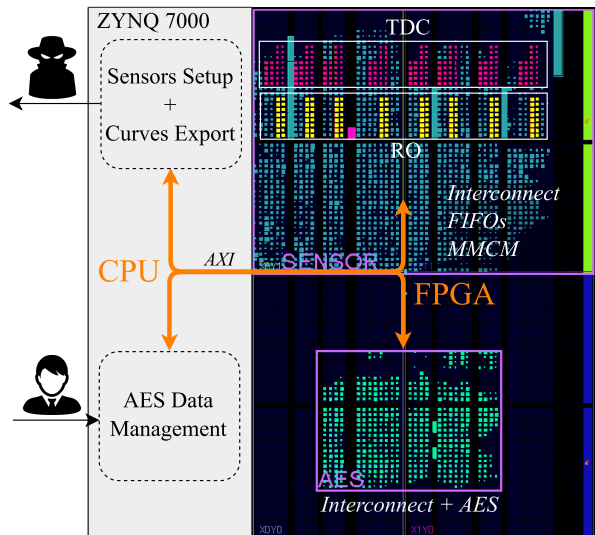


Fig. 6: Xilinx Zynq Multi-User Experimental Setup.

V. RO-SENSOR BASED CORRELATION POWER ANALYSIS ATTACK

The following section provides results of a CPA attack conducted using our RO-based sensors against a hardware AES implemented within the FPGA fabric.

A. Experimental Setup

A Xilinx Zynq SoC that provides both CPU and FPGA on the same die was adopted for our experiments (note that the CPU is not targeted in this attack, we exclusively focus on an intra-FPGA exploit). Our experimental setup is described in figure 6. From the victim point of view, the CPU is dedicated to the management of the AES plain and ciphertexts while the attacker program is developed for sensor calibration and measurement exportation. The FPGA fabric is separated in two logically isolated blocks with distinct clock regions. Because in a multi-user scenario the victim shell might not be necessarily placed next to the adversary, we instantiate the victim AES as far as possible from our sensor instances. Hence, we demonstrate the resilience of our sensors to the distance with the target and to the noise caused by the surrounding logic. The adversary shell contains 64 RO-based sensors (128 slices) and 8 TDC-based sensors (208 slices). The remaining logic is dedicated to interconnect, FIFOs and clock management. Note that a big part of the logic was implemented for experiment purposes and could be removed to get a lighter implementation. The hardware AES module instantiated in the victim shell is dedicated to the acceleration of the encryption of sensitive data. It loads 128-bit packets of plaintexts from the CPU using shared AXI registers, encrypts them and returns the computed ciphertexts. This AES implementation relies on a 128-bit secret key and provides a 128-bit data path. Each round is executed in one clock cycle at 50 MHz. With 1500 LUTs and 400 FFs, this AES module consumes around 10% of the total fabric.

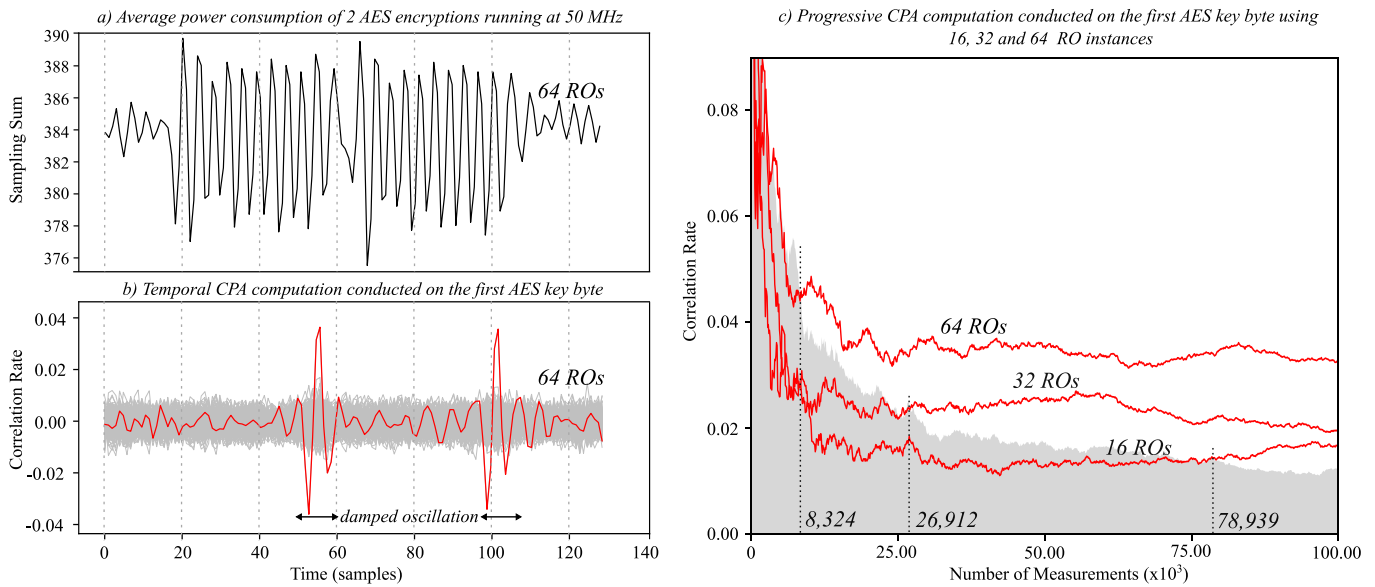


Fig. 7: Averaged AES power consumption (a) and CPA results (b)(c) by means of 100,000 traces acquired using 16, 32 and 64 RO-based sensors. The right key hypothesis candidate is represented in red in (b) and (c).

B. Correlation Power Analysis Model

A CPA attack relies on the fact that CMOS power consumption leakage depends on the handled data. By writing down a model of the expected AES power consumption and combining it to the sensor voltage measurement, CPA should allow us to retrieve the AES secret key. The attack conducted in this paper targets a state register that temporarily stores data resulting from each round transformation of the AES from the plaintext importation to the ciphertext generation. This 128-bit register is synchronously refreshed at the end of each round generating a strong switching leakage that significantly affects the power supply level. The leakage level resulting from the register update fluctuates according to the Hamming Distance between the previous and the current AES state. Targeting this register requires the knowledge of two consecutive states. In our case, we assume that the adversary has access to the ciphertext which is also the last value stored by the AES state register. We adopt the last round attack model described in [24] and compute the correlation rate between the model and the experimental curves. If the adopted model is relevant, the correlation rate of one of the key hypotheses “right candidate” should be distinguishable from the others “wrong candidates”.

C. RO-based Sensor CPA Results

The power consumption resulting from the AES encryption is acquired 100,000 times using RO-based sensors. The average power consumption measured is represented in figure 7.a (two successive encryptions are represented but it has no effect on the side-channel results). Several experiments are conducted to evaluate the CPA results provided by the RO-based sensors.

1) *Number of sensors*: Figure 7.c shows the CPA results obtained using different numbers of RO-based sensors. Using 16 ROs, it takes around 79,000 traces for the right candidate to

emerge from the wrong key hypotheses. With 32 and 64 RO-based sensors the number of required traces drops to 27,000 and 8,000. This attack can also be conducted using only 1 RO-based sensor but requires almost 1 million encryptions which is time-consuming and might be difficult to reproduce in a concrete use case. The number of required traces to infer the secret key is inversely proportional to the number of RO-based sensors used. Moreover, it is enhanced by the granularity improvement and quantization error mitigation provided by the higher number of RO-based sensors.

2) *Target frequency*: In order to study the impact of the target speed on the CPA results, we conduct the same experiment by varying the AES module frequency from 10 to 200 MHz. However, increasing it does not significantly change the CPA results. To explain this phenomenon, we investigate the Zynq response to transient voltage fluctuations. When a sudden voltage drop occurs within the chip (eg: update of the AES state register), the parasitic capacitive and inductive elements forming the PDN resonate and the voltage level temporarily oscillates until finally reaching its steady-state value [20, 17]. The damped oscillation induced by the 10th round update of the AES state register can be seen in temporal correlation results depicted in figure 7.b ($f_{aes} = 50$ MHz). As the voltage transient response is fixed by the parasitic components forming the PDN, the amount of time during which the side-channel leakage can be leveraged is not bounded to the AES frequency but to the device itself [25]. Actually, we observe the exact same oscillation effect for each AES frequency. This experiment demonstrates that, for side-channel purposes, the sensor sampling frequency can be lower than the victim operating frequency. However, it has to be high enough to ensure that the victim valuable side-channel leakage is properly sampled. Through the Zynq transient response oscillation frequency measurement of the

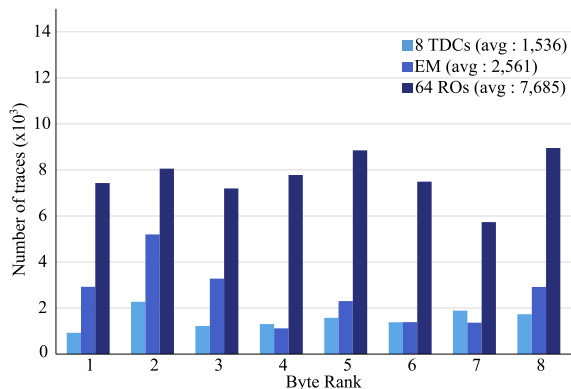


Fig. 8: Number of traces required for the right candidate to emerge from the wrong key hypotheses. Results are given for 8 bytes of the encryption key using 3 measurement setups: EM, 8 TDC-based sensors and 64 RO-based sensors

(≈ 50 MHz), we get an idea of the sampling frequency required to successfully perform the attack (Nyquist-Shannon sampling theorem: $f_s > 2f_{leak}$). Thanks to the 250MS/s rate offered by our sensors we are able to accurately retrieve the side-channel information.

VI. FURTHER RESULTS AND DISCUSSION

This section provides further side-channel results acquired using TDC-based sensors and a traditional electromagnetic side-channel setup. We evaluate the performance of each configuration and discuss about use cases and countermeasures for on-chip sensors.

A. TDC & Electromagnetic Experimental Setup

The adopted TDC-based sensor provides 32 quantization levels and a sampling rate of 250 MS/s. Each instance consumes 26 slices and 8 of them are implemented within the fabric (see Fig. 6). The calibration of the TDC delay line has to be done manually by modifying the number of logic elements forming the *init* block. Each TDC is initialised independently before the measurements.

The EM setup consists in a Langer near field microprobe connected to an oscilloscope with a 5GS/s sampling rate and a 12-bit resolution. The probe is controlled using a X,Y,Z table. The signal is first amplified by a low noise amplifier (LNA) before being fed into the oscilloscope. The electromagnetic leakage of the first AES round is used to trigger the oscilloscope. The captured samples are then extracted and used to perform a correlation electromagnetic analysis (CEMA) [26].

B. Side-Channel Results

A single campaign of encryption in which the three measurement setups simultaneously acquire the side-channel leakage is conducted. Figure 8 presents a bar chart showing the number of traces needed for the right guess to emerge depending on the measurement setup. Results are given for the first 8 bytes of the encryption key. TDC-based sensors provide comparable results to that of the electromagnetic setup while RO-based sensor remains 3 or 4 times less efficient than

the other setups. Despite a significant difference of sampling frequency and resolution between integrated sensors and oscilloscope, the results obtained are quite similar. Naturally, these results must be interpreted with caution as TDCs and ROs were previously calibrated and optimized for this specific device and attack scenario. Our RO-based sensors do not reach the level of accuracy of TDC-based sensors but are still precise enough to successfully perform a CPA. Moreover, they benefit from significant implementation advantages that will be addressed in the following discussion.

C. Discussion

1) *On-chip sensor comparison*: When designing on-chip voltage sensors a trade-off needs to be made between achievable resolution, sampling frequency and area coverage. Depending on the use case, a sensor will be more relevant than the others. Regarding the results of CPA conducted in this paper, our novel RO-based sensor remains slightly less efficient than TDC-based sensors for side-channel purposes. However, this sensor offers a better flexibility and scalability than TDC-based sensors. Thanks to their light implementation RO-based sensors can be spread through all the chip without congestion. Thus, they provide a better coverage of the power supply voltage fluctuations throughout the fabric with a lower area cost. Moreover, they don't need any calibration or specific logic cells contrarily to the TDC-based sensor which requires an init delay configuration to control the position of the clock edge inside the delay line as well as specific CARRY4 logic to provide reliable measurements. This suggests that RO-based sensors would be easier to transpose on devices integrating different manufacturing processes.

2) *Potential use-cases*: RO-based sensors have been widely adopted for voltage and temperature monitoring [11], attack detection [23] and more recently side-channel attacks [8]. Our novel RO-based sensor has been shown suitable for statistical side-channel attacks and could be used to improve the previous applications. A further side-channel use case for on-chip sensors was discussed in [27] and consists in their implementation as hardware Trojans. FPGA end-products often include third-party IP blocks because of the high-cost of design and development. Considering the critical application in which FPGAs are deployed, the potential integration of FPGA-based trojans through untrusted IPs could lead to disastrous consequences. (eg: industrial espionage, denial-of-service, etc).

3) *Side-channel countermeasures*: Multi-users FPGA have not been launched yet but several technical papers already shown the multi-user feasibility and the benefits that a highly scalable and flexible multi-user service could provide to tenants and more specifically to cloud providers. Despite the fact that logical isolation between logic blocks is ineffective against power side-channel attacks [28], side-channel threats could be easily mitigated by restricting manual place and route and forbidding combinational loops that enable the RO implementation. The problem lies in the fact that these features are essential for a lot of FPGA applications and

their suppression would significantly alter the service. Trojan detection routines could also be developed to prevent designers from implementing on-chip sensors but will require a lot of developments and will be soon challenged by novel adversary designs bypassing the security. In conclusion, there is no easy way to mitigate this threat and FPGA hardware attacks are likely to remain problematic for cloud providers.

VII. CONCLUSION

This article introduces a **novel design for on-chip voltage sensors based on ROs**. By enhancing sampling frequency and resolution of this kind of sensors, we enable their use for runtime voltage fluctuation measurements. To illustrate the performances provided by our sensors, we adopt them to conduct a power side-channel attack within a FPGA fabric. A multi-user cloud scenario is reproduced, an adversary sensor shell is used to perform an attack against a victim AES module located within a logically isolated shell. Thanks to the performance improvement of our sensor we are able to perform **the first CPA attack conducted using RO-based sensor** inside a FPGA. Successful result obtained in retrieving the secret key of the AES running at 50 MHz demonstrates the performances provided by our sensors for side-channel purpose. To further evaluate our sensor, we compare it to different kinds of side-channel setups. A CPA attack is conducted using TDC-based sensors and an EM traditional side-channel setup. We show that thanks to their proximity to the target, **on-chip sensors provide results similar to near field EM** even with a much smaller sampling rate and resolution. Regarding the integrated voltage sensors, our novel RO-based sensors almost reach the accuracy of TDC-based sensors and benefit from a lighter area overhead, a better spatial coverage and an easier implementation as they rely on basic logic gates. Finally they stand as **an ideal alternative for monitoring fine-grained high-speed voltage fluctuations in SoCs**.

REFERENCES

- [1] David Pellerin. FPGA Accelerated Computing Using AWS F1 Instances, 2017.
- [2] Alibaba Cloud ECS. Deep Dive into Alibaba Cloud F3 FPGA as a Service Instances, 2018.
- [3] Fei Chen, Yi Shan, Yu Zhang, Yu Wang, Hubertus Franke, Xiaotao Chang, and Kun Wang. Enabling FPGAs in the cloud. In *Computing Frontiers*. ACM Press, 2014.
- [4] Steve Trimberger and Steve McNeil. Security of FPGAs in data centers. In *IEEE International Verification and Security Workshop*, 2017.
- [5] Jonas Krautter, Dennis R E Gnad, and Mehdi B Tahoori. FPGAhammer : Remote Voltage Fault Attacks on Shared FPGAs, suitable for DFA on AES. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018.
- [6] Falk Schellenberg, Dennis R.E. Gnad, Amir Moradi, and Mehdi B. Tahoori. An inside job: Remote power analysis attacks on FPGAs. In *Design, Automation & Test in Europe Conference & Exhibition*. IEEE, 2018.
- [7] Falk Schellenberg, Dennis R. E. Gnad, Amir Moradi, and Mehdi B. Tahoori. Remote inter-chip power analysis side-channel attacks at board-level. *Proceedings of the International Conference on Computer-Aided Design*, 2018.
- [8] Mark Zhao and G. Edward Suh. FPGA-Based Remote Power Side-Channel Attacks. In *IEEE Symposium on Security and Privacy*, 2018.
- [9] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in memory without accessing them. *ACM SIGARCH*, 2014.
- [10] Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo. CLKSCREW: Exposing the perils of security-oblivious energy management. *26th USENIX Security Symposium*, 2017.
- [11] Kenneth M. Zick and John P. Hayes. Low-cost sensing with ring oscillator arrays for healthier reconfigurable systems. *ACM Transactions on Reconfigurable Technology and Systems*, 5(1):1–26, 2012.
- [12] Kenneth M. Zick, Meeta Srivastav, Wei Zhang, and Matthew French. Sensing nanosecond-scale voltage attacks and natural transients in FPGAs. *Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays*, 2013.
- [13] Paul C Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *CRYPTO '96*. 1996.
- [14] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. *Advances in Cryptology*, 1999.
- [15] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation Power Analysis with a Leakage Model. *Cryptographic Hardware and Embedded Systems*, 2004.
- [16] Meeta S. Gupta, Jarod L. Oatley, Russ Joseph, Gu Yeon Wei, and David M. Brooks. Understanding voltage variations in chip multiprocessors using a distributed power-delivery network. *Design, Automation and Test in Europe, DATE*, 2007.
- [17] Shidhartha Das, Paul Whatmough, and David Bull. Modeling and characterization of the system-level Power Delivery Network for a dual-core ARM Cortex-A57 cluster in 28nm CMOS. *Proceedings of the International Symposium on Low Power Electronics and Design*, 2015.
- [18] Jean-Max Dutertre, Bruno Robisson, Assia Tria, and Loic Zussa. Investigation of timing constraints violation as a fault injection means. *Design of Circuits and Integrated Systems*, 2012.
- [19] Chun Chi Chen, Wen Fu Lu, Chin Chung Tsai, and Poki Chen. A time-to-digital-converter-based CMOS smart temperature sensor. *IEEE International Symposium on Circuits and Systems*, 2005.
- [20] Miho Ueno, Masanori Hashimoto, and Takao Onoye. Real-time on-chip supply voltage sensor and its application to trace-based timing error localization. In *International On-Line Testing Symposium (IOLTS)*. IEEE, jul 2015.
- [21] Daisuke Fujimoto, Yu Ichi Hayashi, Arthur Beckers, Josep Balasch, Benedikt Gierlichs, and Ingrid Verbauwhede. Detection of IEMI fault injection using voltage monitor constructed with fully digital circuit. *IEEE EMC/APEMC*, 2018.
- [22] Dennis R. E. Gnad, Fabian Oboril, Saman Kiamehr, and Mehdi B. Tahoori. An Experimental Evaluation and Analysis of Transient Voltage Fluctuations in FPGAs. *IEEE Transactions on Very Large Scale Integration Systems*, 2018.
- [23] Adrien Le Masle and Wayne Luk. Detecting power attacks on reconfigurable hardware. In *22nd International Conference on Field Programmable Logic and Applications*, 2012.
- [24] Hassen Mestiri, Noura Benhadjyoussef, Mohsen Machhout, and Rached Tourki. A Comparative Study of Power Consumption Models for CPA Attack. *International Journal of Computer Network and Information Security*, 2013.
- [25] Sébastien Tiran, Sébastien Ordas, Yannick Teglia, Michel Agoyan, and Philippe Maurine. A Frequency Leakage Model and its application to CPA and DPA. *IACR*, 2013.
- [26] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic Analysis: Concrete Results. *Cryptographic Hardware and Embedded Systems*, 2001.
- [27] Ilias Giechaskiel, Kasper B. Rasmussen, and Ken Eguro. Leaky Wires. In *Asia Conference on Computer and Communications Security*, 2018.
- [28] Loic Zussa, Jean Max Dutertre, Jessy Clediere, and Bruno Robisson. Analysis of the fault injection mechanism related to negative and positive power supply glitches using an on-chip voltmeter. *IEEE International Symposium on Hardware-Oriented Security and Trust*, 2014.