Lightweight Embedded Detection System against Voltage Drop Fault Attacks in Multi-Tenant FPGAs

Gwenn Le Gonidec (Lab-STICC, UBS)
Guillaume Bouffard (ANSSI)
Jean-Christophe Prévotet (IETR, INSA Rennes)
Maria Méndez Real (Lab-STICC, UBS)

owen.le-gonidec@univ-ubs.fr

Fundings project: ANR JCJC CoPhyTEE
Addressing Covert & Physical Attacks performed from SW in Open-Hardware Trusted Execution Environment-enabled System-on-Chip
ANR-23-CE39-0003-01

www.univ-ubs.fr

Overview

Dynamic detection of voltage drop fault attacks in FPGAs

Constraints:

- Unknown sensor location
- Unknown hardware surrounding the sensor
- Short attack patterns

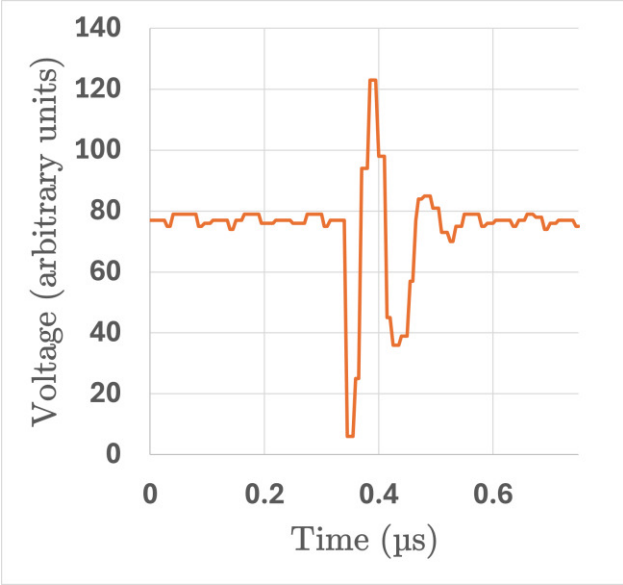

Short voltage drop attack

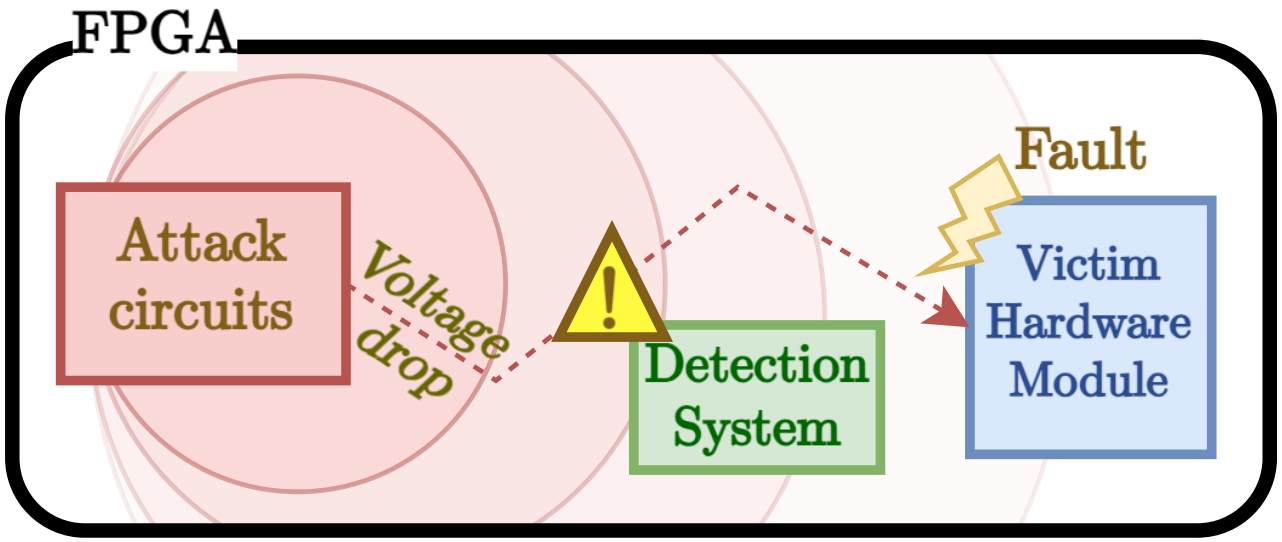

Main contributions:

- Dynamic detection based on a TDC sensor and a **variance-based metric**
- Efficient with a moderate use of hardware resources
- Location-agnostic
- Robust to background noise (false positives)

Threat model — Voltage drop fault attacks

Multi-tenant FPGAs

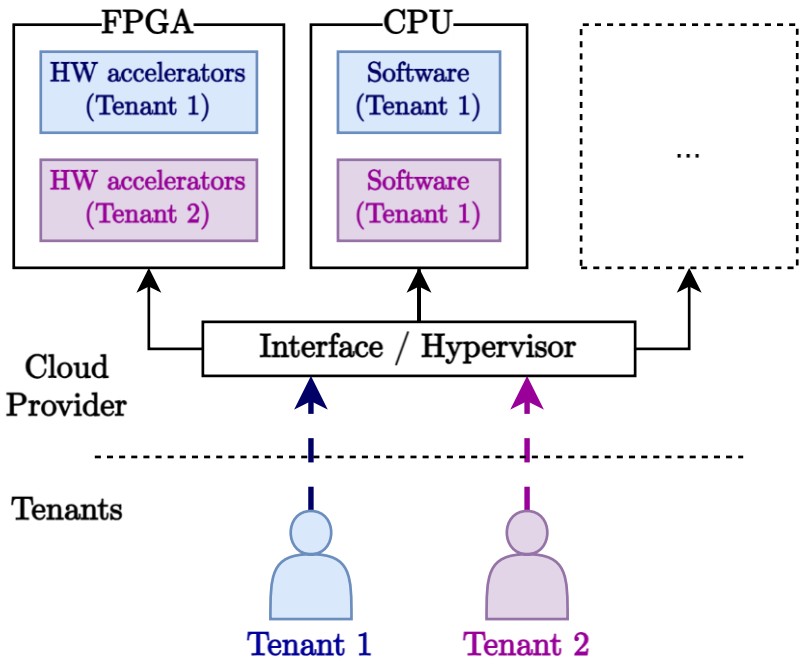

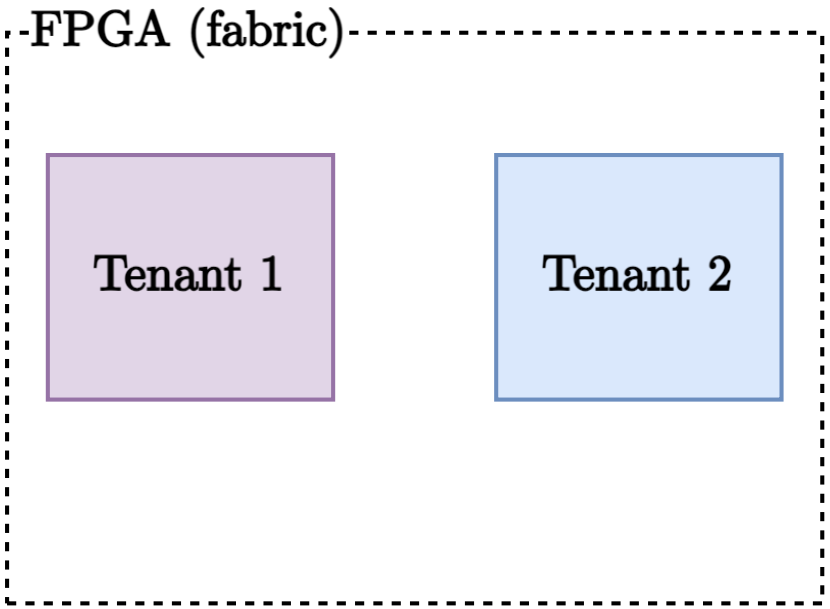

Multi-tenant FPGAs

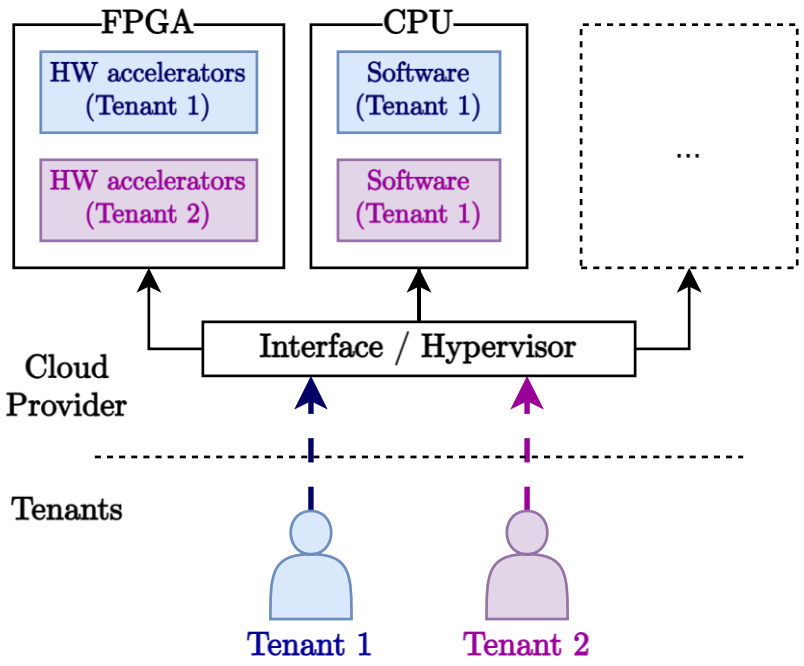

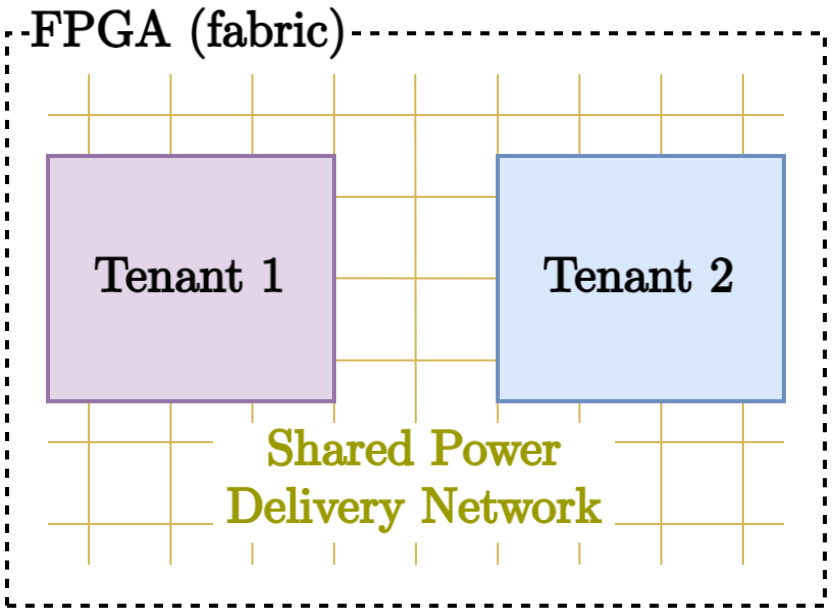

Multi-tenant FPGAs

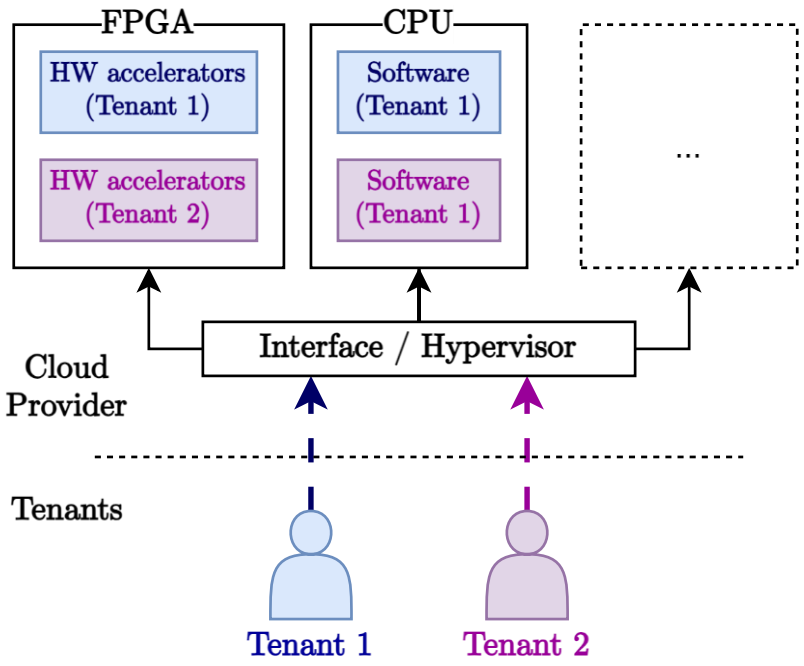

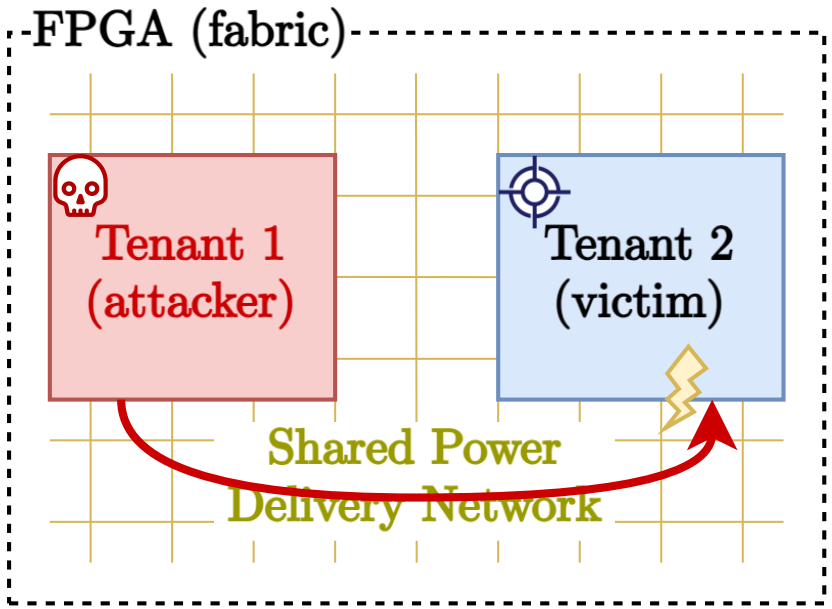

→ **The activity of one tenant creates voltage fluctuations on the surrounding area**

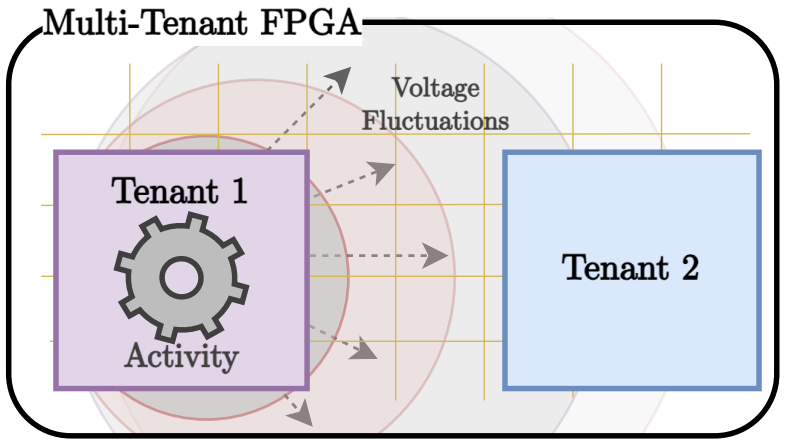

Voltage Drop Attacks

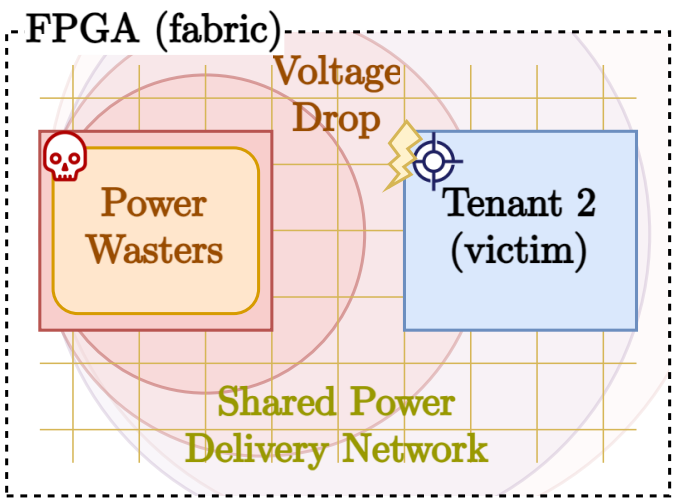

Combinational logic's propagation time is proportional to the supply voltage

→ Voltage drop creates timing faults

Power Wasters:

- Cryptographic Primitives[1]
- Latches, flip-flops[2]
- **Ring Oscillators**

Simple Ring Oscillator

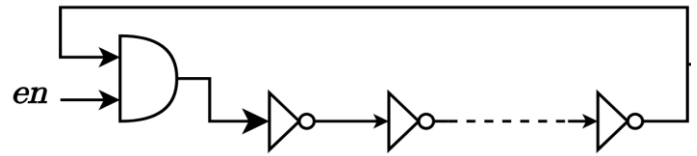

Previous work has shown :

- Crashes / resets of the host FPGA
- Differential Fault Analysis (DFA) on cryptographic primitives
- Bit-flips in the FPGA configuration memory

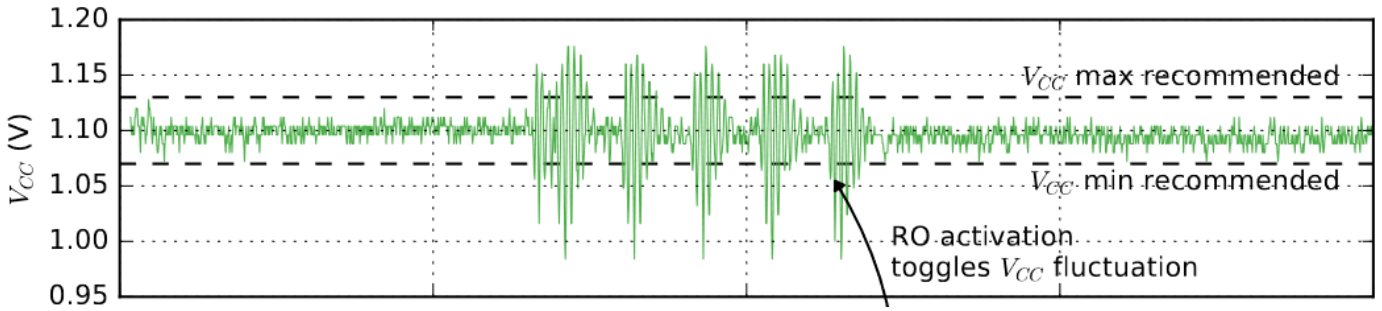

Re-printed from [3]: impact of Ring Oscillators on the voltage

[1] M. Gross, J. Krautter, D. Gnad, M. Gruber, G. Sigl and M. Tahoori, "**FPGANeedle: Precise Remote Fault Attacks from FPGA to CPU**," *2023 28th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Tokyo, Japan, 2023, pp. 1-7.

[2] Sugawara, T., Sakiyama, K., Nashimoto, S., Suzuki, D. and Nagatsuka, T. (2019), **Oscillator without a combinatorial loop and its threat to FPGA in data centre**. Electron. Lett., 55: 640-642

[3] Krautter, J., Gnad, D. R. E., & Tahoori, M. B. (2018). **FPGAhammer: Remote Voltage Fault Attacks on Shared FPGAs, suitable for DFA on AES**. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, *2018*(3), 44-68

Attacker model: short voltage drops for fault attacks

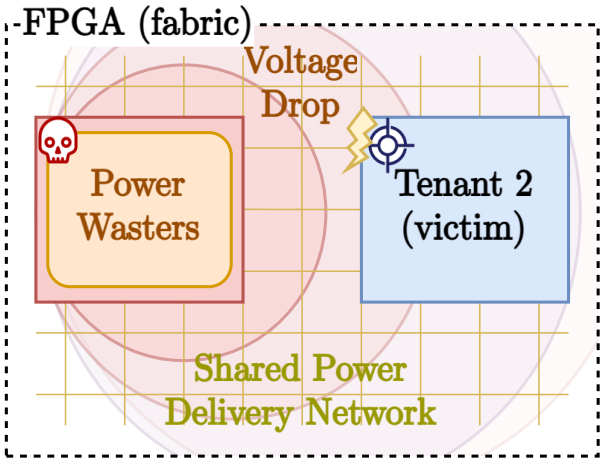

Attacker: Enhanced Ring Oscillators (EROs)
→ varying amount and attack duration

Victim: hardware AES implementation
→ Successful attack = unexpected ciphertext

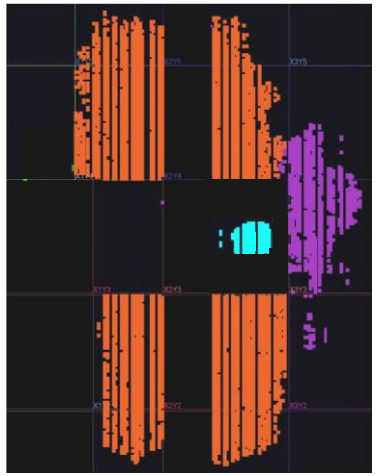

EROs (attacker)
AES (victim)
Debug

Attacker model: short voltage drops for fault attacks

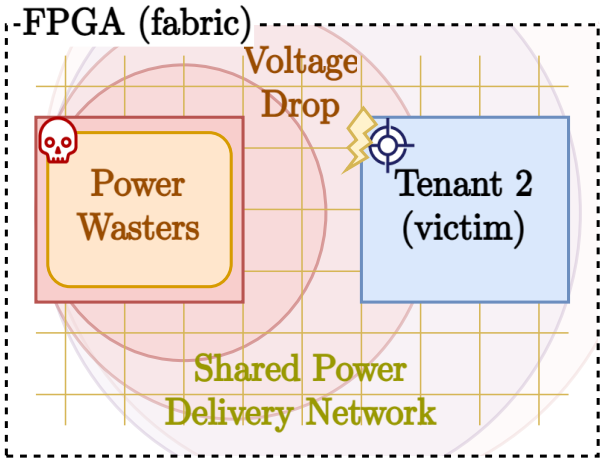

Attacker: Enhanced Ring Oscillators (EROs)
→ varying amount and attack duration

Victim: hardware AES implementation
→ Successful attack = unexpected ciphertext

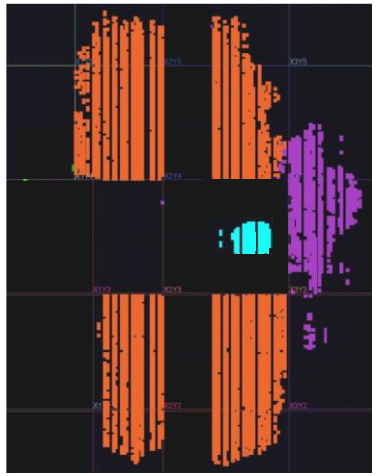

EROs (attacker)
AES (victim)
Debug

→ Even very short attacks can have good success rates.

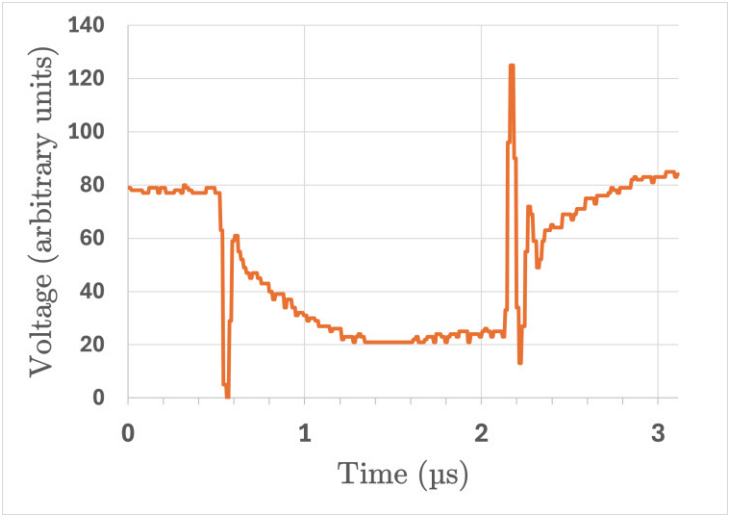

Example long attack

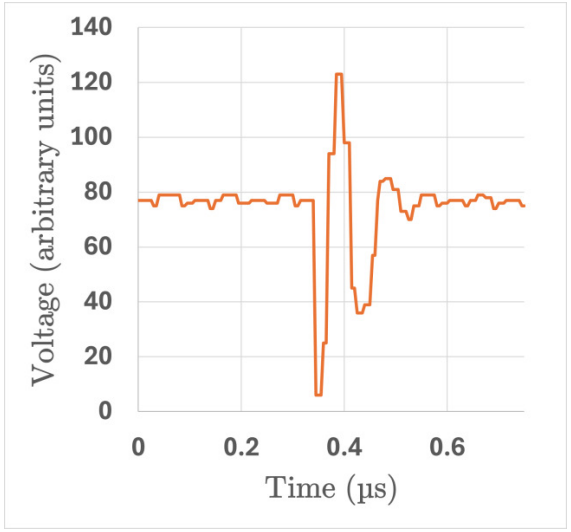

Example short attack

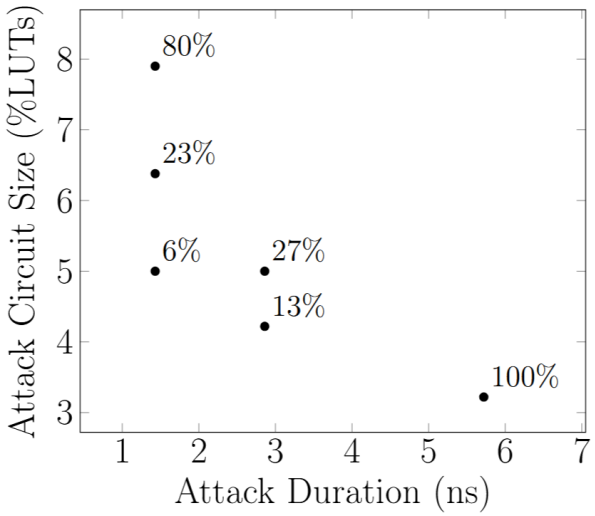

Success rate VS duration and attak circuit size

Previous work— Detecting voltage drops

Detecting voltage drops: two approaches

- **Static analyses**
  → Detection of power wasters in the bitstream metadata[1,2].
- **Dynamic Detection**
  → Detecting voltage drops using embedded sensors.

Voltage sensor: Time-to-Digital Converter (TDC)

→ Used in most recent work

✓ High resolution, fine-grained

✗ Very sensitive to process, temperature and placement variations.

✗ Needs to be calibrated

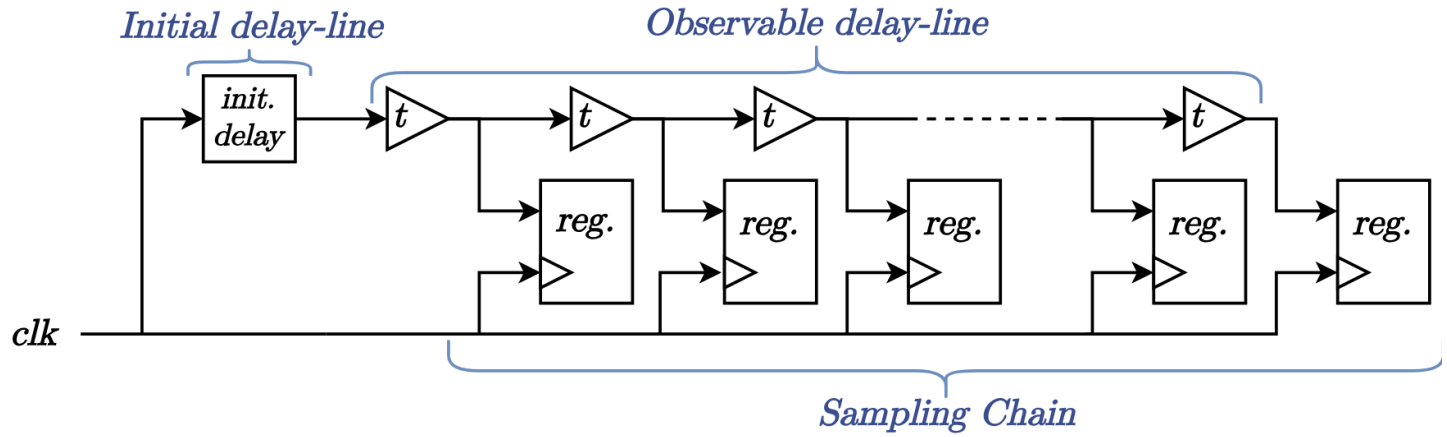

[1] Fathy, Mohamed & Nassar, Hassan & Ghany, Mohamed & Henkel, Jörg. (2025). **Timekeepers: ML-Driven SDF Analysis for Power-Wasters Detection in FPGAs**. ACM Transactions on Embedded Computing Systems. 24. 10.1145/3761809.

[2] H. Nassar, J. Krautter, L. Bauer, D. Gnad, M. Tahoori and J. Henkel, "Meta-Scanner: Detecting Fault Attacks via Scanning FPGA Designs Metadata," in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 43, no. 11, pp. 3443-3454, Nov. 2024

Existing sensor-based detection schemes: two main approaches

First approach

Static threshold at the output of the sensor [1,2]

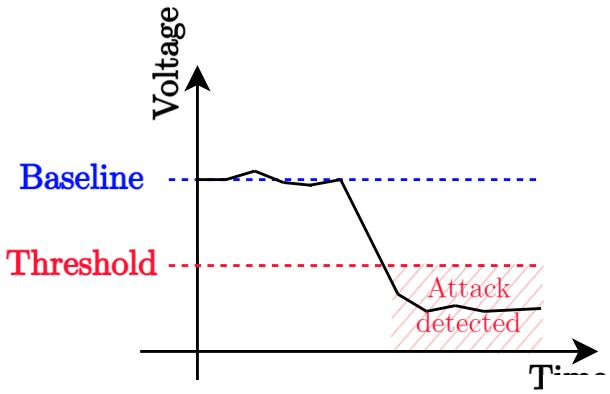

✓ Simple, fast detection, lightweight

✗ **The baseline value changes with the sensor's location**

→ Unsuitable for our multi-tenant context

→ Inefficient in multi-sensors approach

→ Analyzing the location of sensors

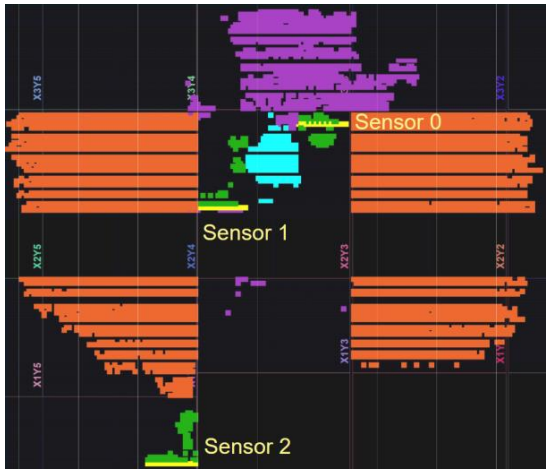

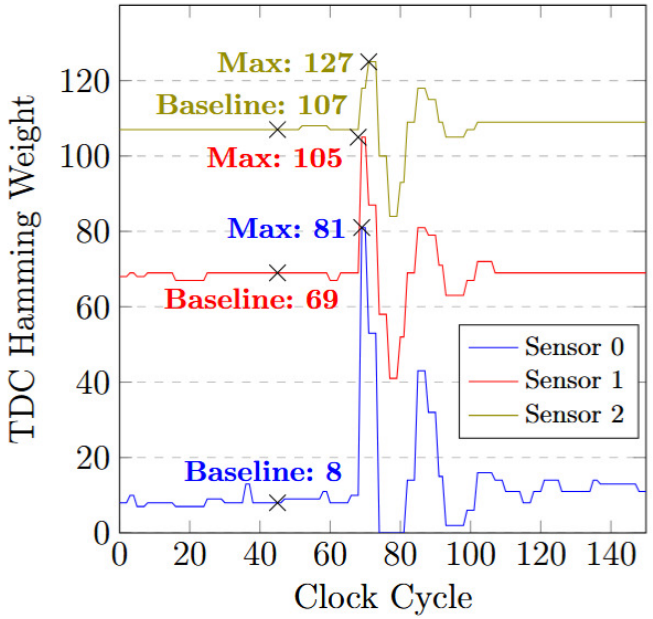

→ The baseline of each sensor is far from that of the others

→ Each sensor must have its own threshold

[1] Y. Luo and X. Xu, "**A Quantitative Defense Framework against Power Attacks on Multi-tenant FPGA,**" *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, San Diego, CA, USA, 2020, pp. 1-4.

[2] H. Nassar, H. AlZughbi, D. R. E. Gnad, L. Bauer, M. B. Tahoori and J. Henkel, "**LoopBreaker: Disabling Interconnects to Mitigate Voltage-Based Attacks in Multi-Tenant FPGAs**," *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, Munich, Germany, 2021, pp. 1-9

[3] George Provelengios, Daniel Holcomb, and Russell Tessier, "**Mitigating Voltage Attacks in Multi-Tenant FPGAs**, in ACM Trans. Reconfigurable Technol. Syst. 14, June 2021

Existing sensor-based detection schemes

Second approach

Arithmetic calculations [1,2]

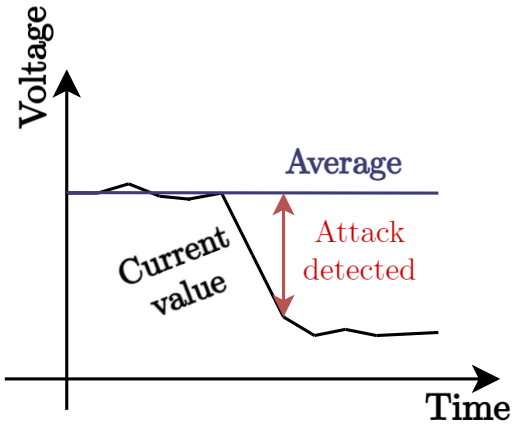

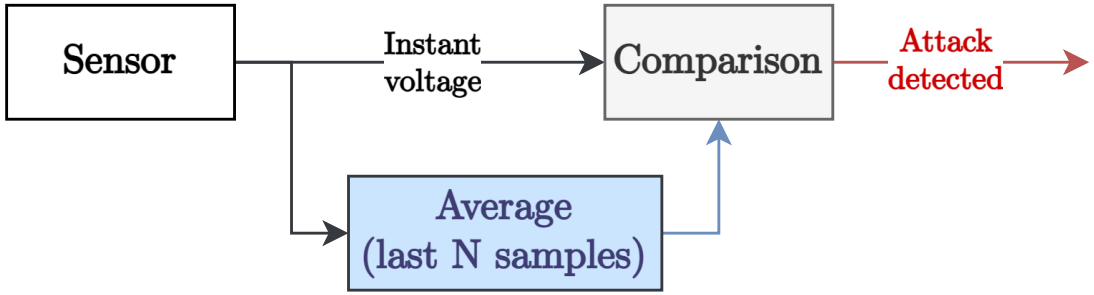

✓ Location-Agnostic → suitable for multi-tenant

✗ **Resource-hungry** (storing samples for the average calculation)

✗ **Slow sampling period** (may not detect short attacks)

Detection method	Location-agnostic	Moderate resource usage	Can detect short attacks
Static Threshold	✗	✓	✓
Arithmetic Calculations	✓	✗	✗
Our work (objectives)	✓	✓	✓

Survey on energy-based attack without physical access to the target device (including FPGA voltage drop attacks):

Le Gonidec, G., Bouffard, G., Prevotet, J. C., & Méndez Real, M. (2025). **Do Not Trust Power Management: A Survey on Internal Energy-based Attacks Circumventing Trusted Execution Environments Security Properties**. *ACM Transactions on Embedded Computing Systems*, *24*(4), 1-35.

[1] S. S. Mirzargar, G. Renault, A. Guerrieri and M. Stojilović, "**Nonintrusive and Adaptive Monitoring for Locating Voltage Attacks in Virtualized FPGAs**," *2020 International Conference on Field-Programmable Technology (ICFPT)*, Maui, HI, USA, 2020

[2] M. A. Kajol, S. Sunkavilli and Q. Yu, "**AHD-LAM: A New Mitigation Method against Voltage-Drop Attacks in Multi-tenant FPGAs**," *2023 Asian Hardware Oriented Security and Trust Symposium (AsianHOST)*, Tianjin, China, 2023

Proposed detection system

Sensor implementation and detection metric

Hamming Weight *(HW)* at the output of a **TDC Sensor**

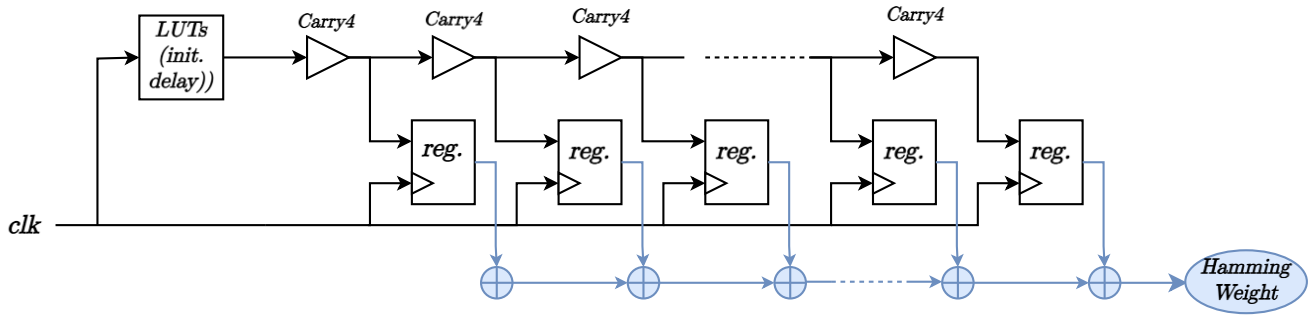

Detection metric: its **variance** on the last **T samples**

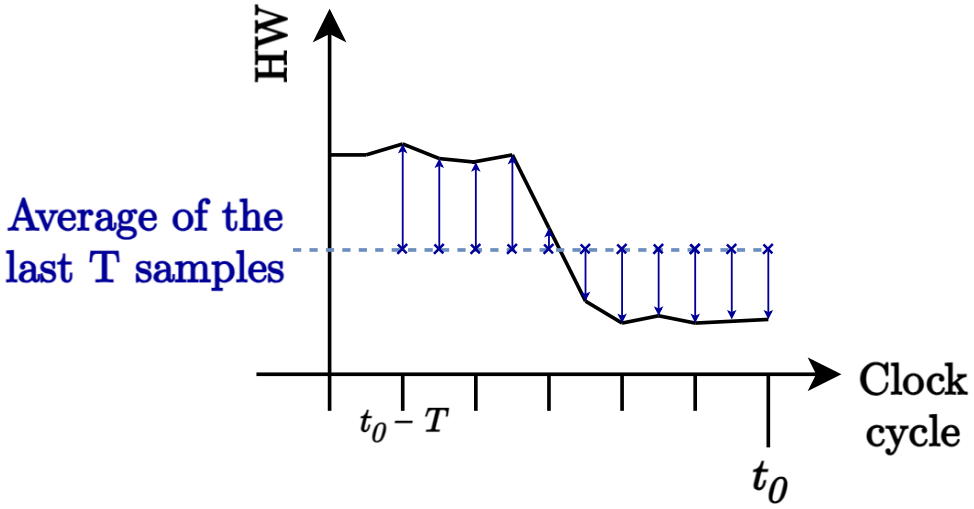

$$var(t) = \sum_{t=t_0-T}^{t_0} \mathrm{HW}(t) - \overline{\mathrm{HW}(t_0)}$$

$$= \frac{1}{4} \sum_{t=t_0-4}^{t_0} HW(t)^2 - \left(\frac{1}{4} \sum_{t=t_0-4}^{t_0} HW(t) \right)^2$$

✓ No need for a fixed basline → **location-agnostic**

✓ **Amplifies** steep voltage changes that occurred in the last T cycles

Implementation summary and resource usage

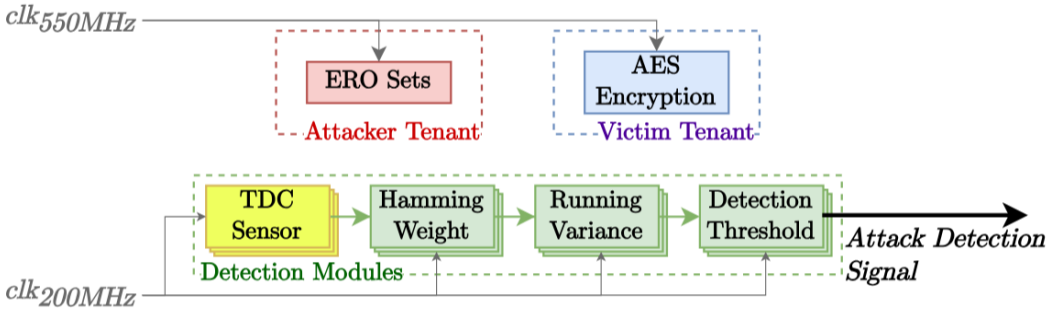

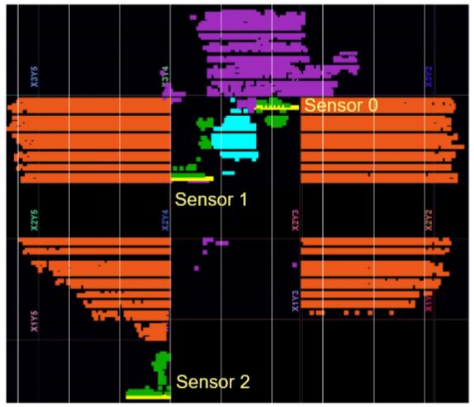

TDC Sensors
AES Victim
Debug Modules
Attack Circuit
Detection
(Hamming weight, Variance)

For the variance, window size T=4
(Implemented with a bit-shift)

Resource utilization

Module	LUTs (%)	CARRY8 (%)	Registers (%)	DSP (%)
TDC sensor	32 (0.01)	16 (0.06)	128 (0.02)	0
Hamming weight	148 (0.06)	1 (< 0.01)	8 (< 0.01)	0
Variance	289 (0.13)	34 (0.12)	48 (0.01)	1 (0.06)
Adder (M)	8 (0.01)	1 (< 0.01)	1 (< 0.01)	0
Total	**477 (0.18%)**	**52 (0.18%)**	**185 (0.03%)**	**1 (0.06 %)**

→ Same order of magnitude as the sensor itself
→ Slightly more frugal with Welford's algorithm
Resource utilization for the other arithmetic methods are not communicated

Detecting attacks using the variance metric

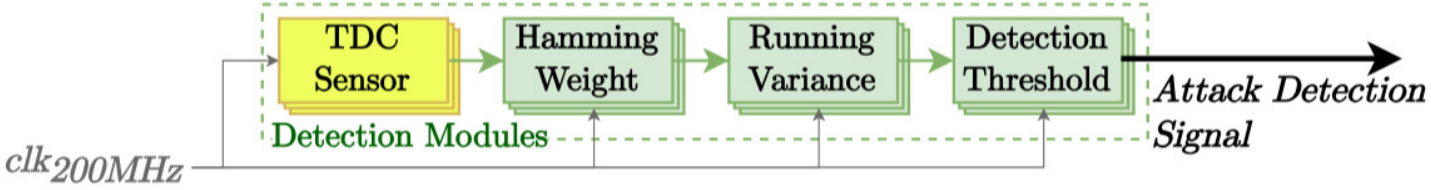

→ **The variance metric allows location-agnostic detection**

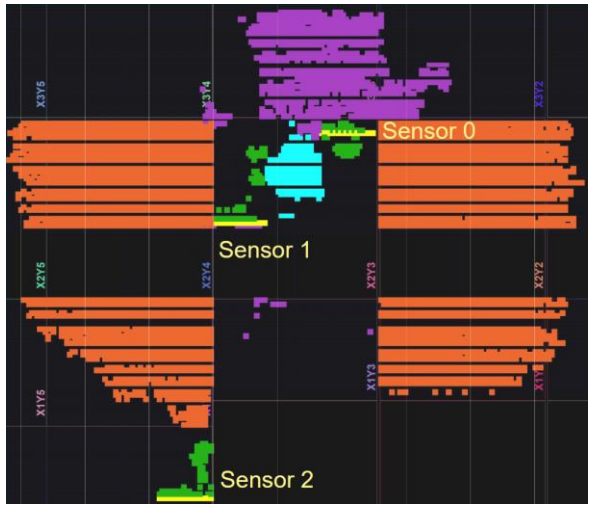

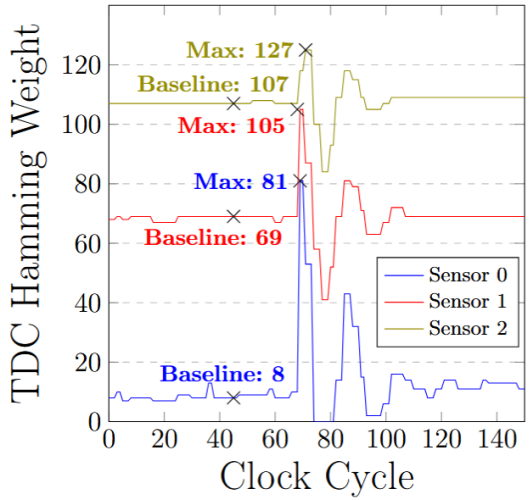

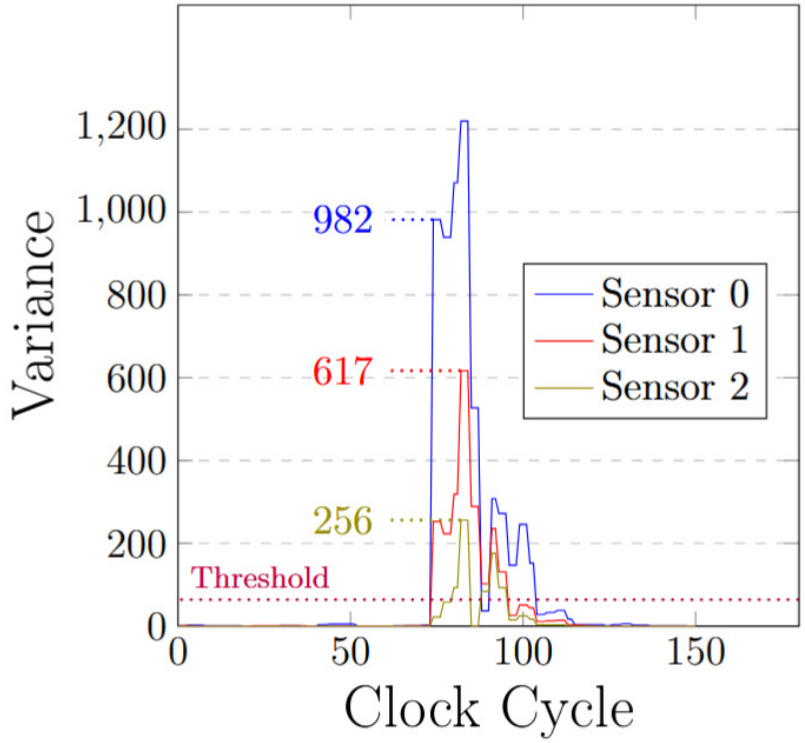

Average detection delay after attack starts: 23.6ns
(3 – 4 cycles at 200MHz)

→ **It is currently not possible to prevent faults occurring before detection** (almost instantaneous)

Adding background noise

AES modules at each end of the sensor

→ Upon activation, they provoke a **significant voltage variation**.

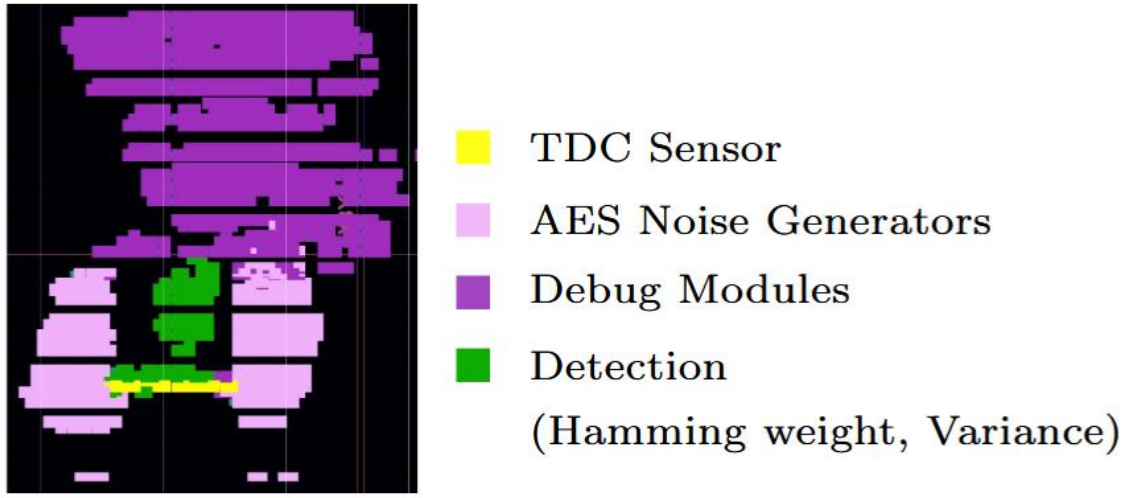

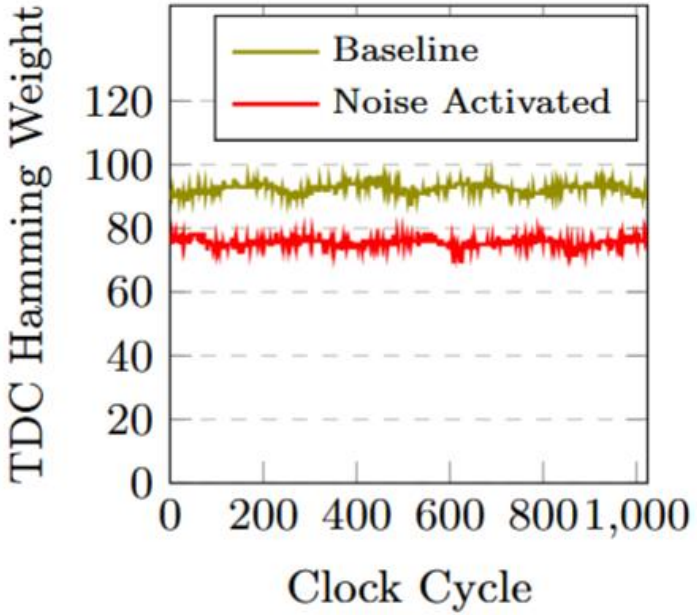

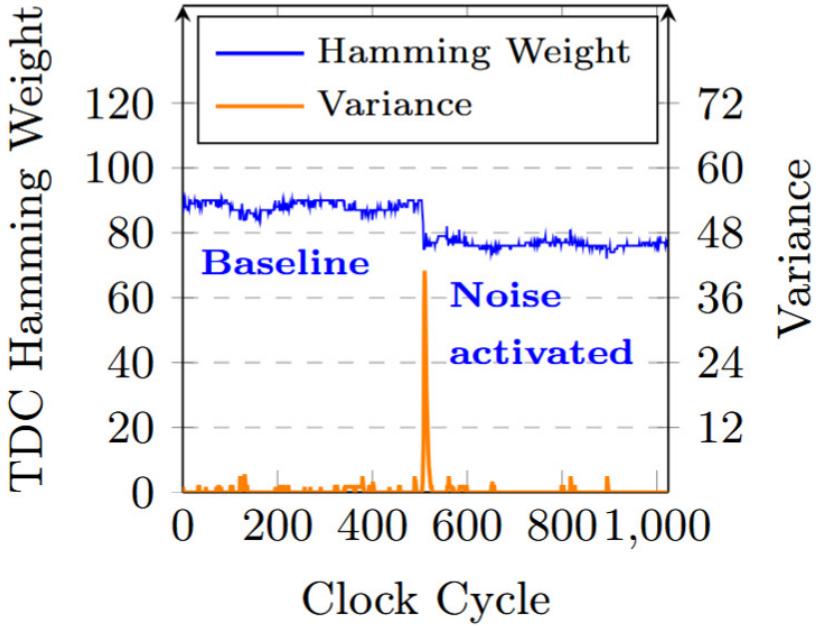

→ The variance detection threshold should be higher than the noise activation effect

Detection success rate

Setting the variance detection threshold

Looking at the repartition of the variance upon noise activation VS. upon attack, we set a threshold to detect all tested attacks without false positives

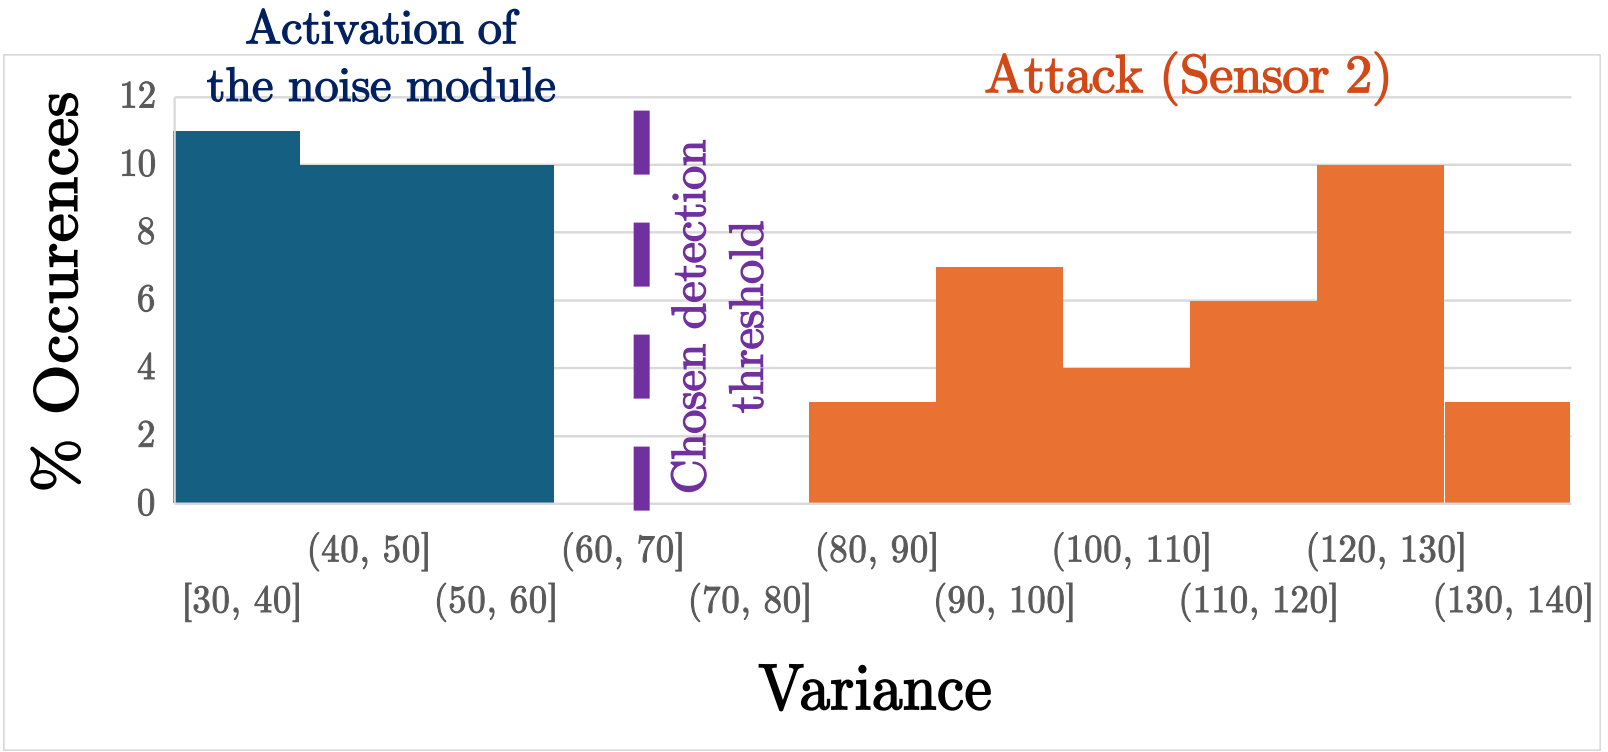

→ **Variance=64**.

→ In 32K attack attempts, 100% successful detection

Conclusion and Discussion

Main contributions and results

1. Attack characterization

→ **Attacks used to characterize countermeasures are over-estimated**

→ Only a few clock cycles of activation are necessary for fault injection

2. Noise characterization

→ **The effect of noise is not negligible**

→ Side observations: oscillation at the resonant frequency of the PDN

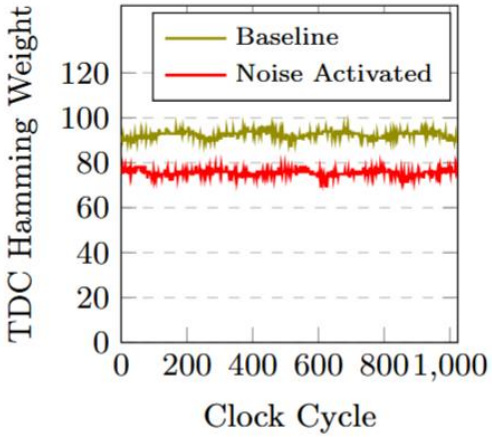

3. Detection system

→ We designed a detection system robust to location changes and background noise

→ We found a trade-off between detection delay and resource utilization

Our implementation is available in open-access:
https://sourcesup.renater.fr/projects/detectionva

Following up

Current work: characterization of voltage drop attacks against soft-cores in FPGAs

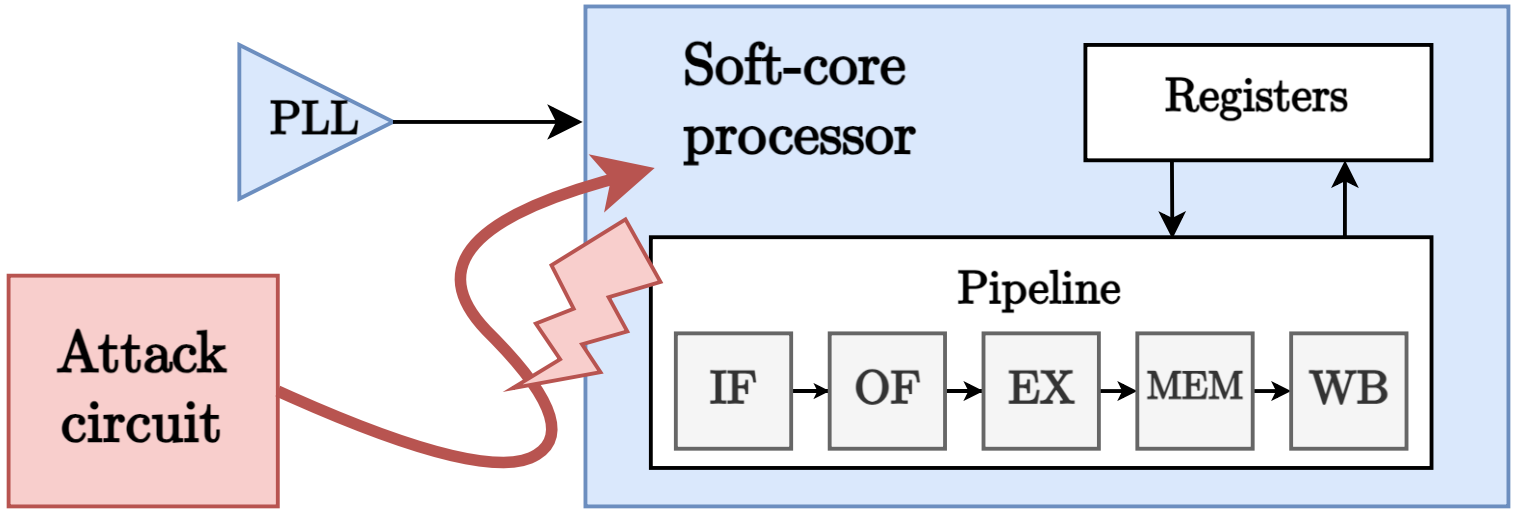

→ How efficient is voltage drop at fault injection on soft-cores?
→ What type of faults can we observe?
→ Can these faults be predicted?

Thank you for listening!

Lightweight Embedded Detection System against Voltage Drop Fault Attacks in Multi-Tenant FPGAs

Gwenn Le Gonidec (Lab-STICC, UBS)
Guillaume Bouffard (ANSSI)
Jean-Christophe Prévotet (IETR, INSA Rennes)
Maria Méndez Real (Lab-STICC, UBS)

owen.le-gonidec@univ-ubs.fr

Fundings project: ANR JCJC CoPhyTEE
Addressing Covert & Physical Attacks performed from SW in Open-Hardware Trusted Execution Environment-enabled System-on-Chip
ANR-23-CE39-0003-01

www.univ-ubs.fr

Annex: variance window size

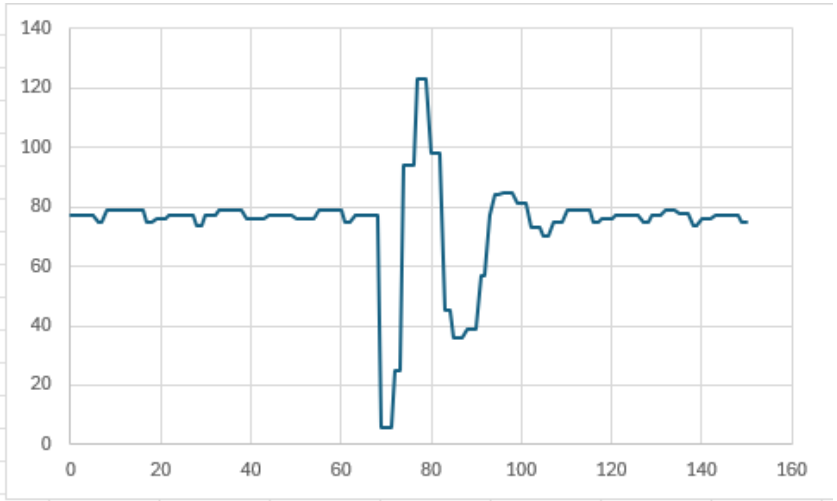

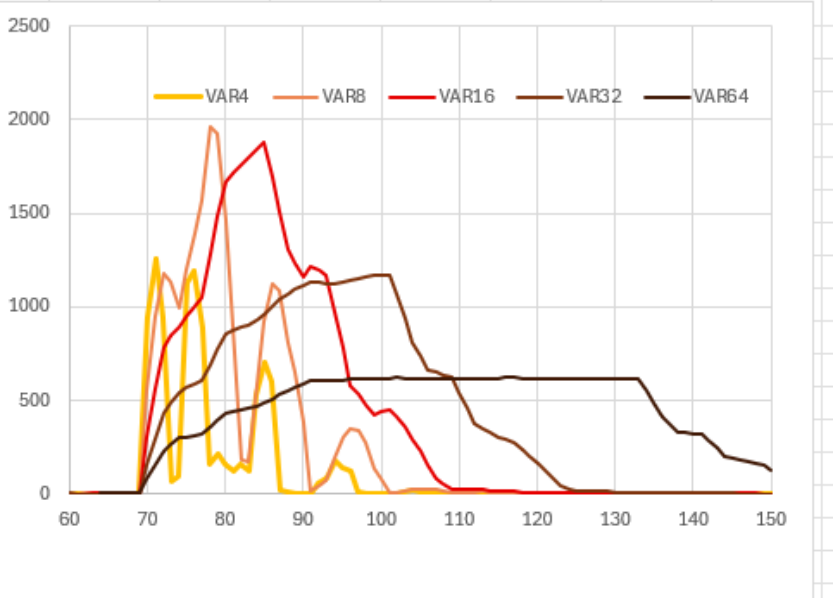

Annex: variance window size