Karlsruhe Institute of Technology

Secure sharing of FPGAs in the Cloud: New Challenges at the Technology Level

Mehdi B. Tahoori

INSTITUT FÜR TECHNISCHE INFORMATIK – CHAIR OF DEPENDABLE NANO COMPUTING

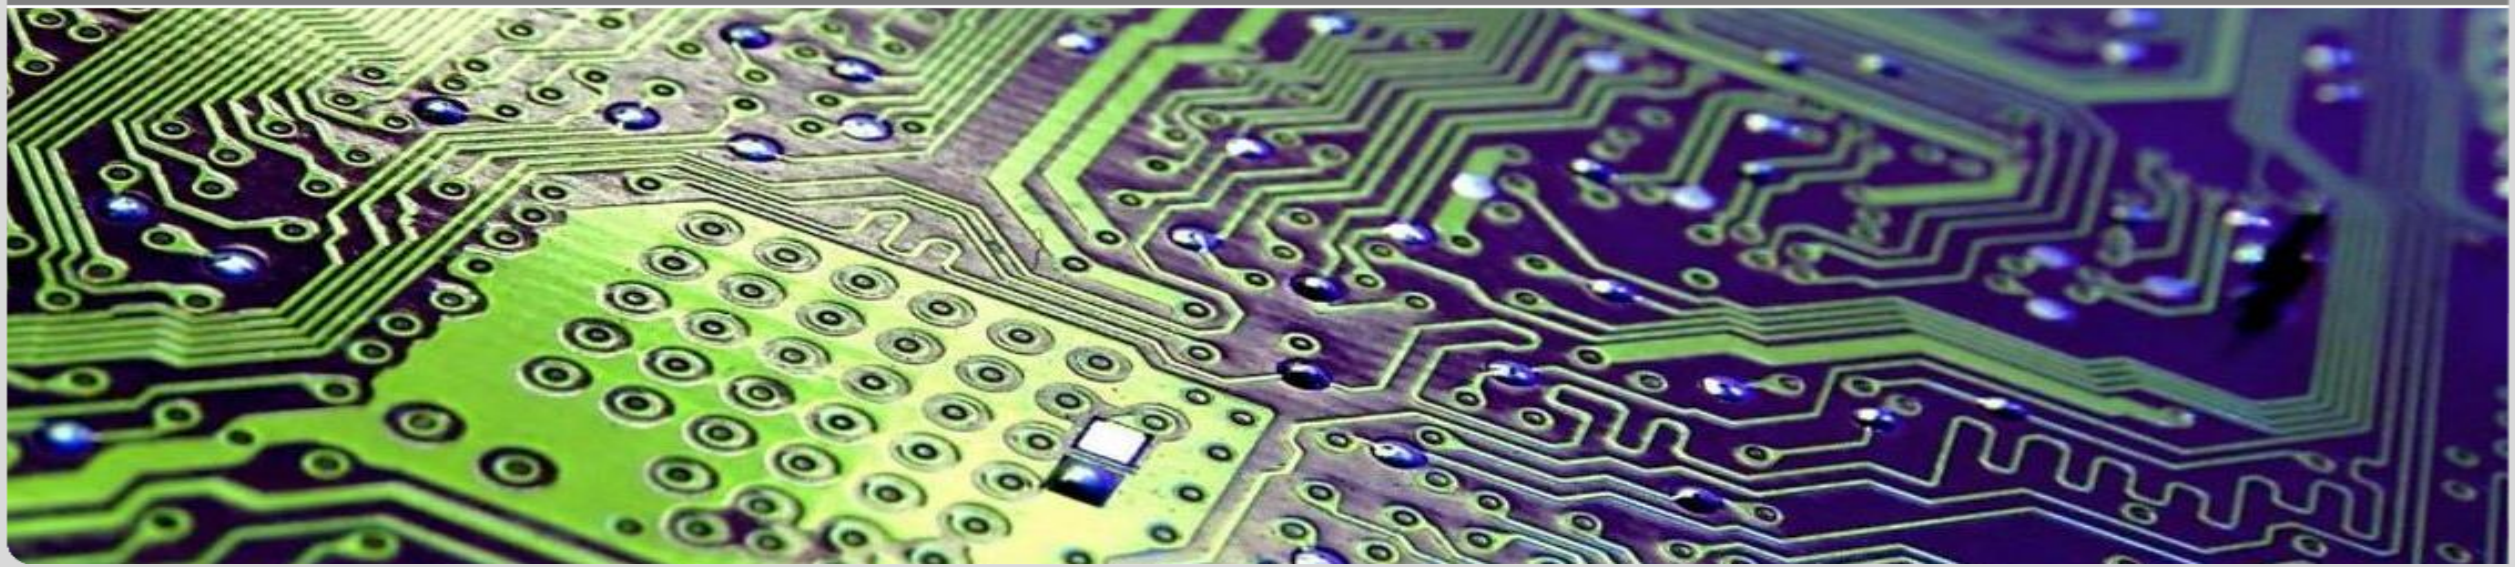

Physical Fault and Side-Channel Attacks

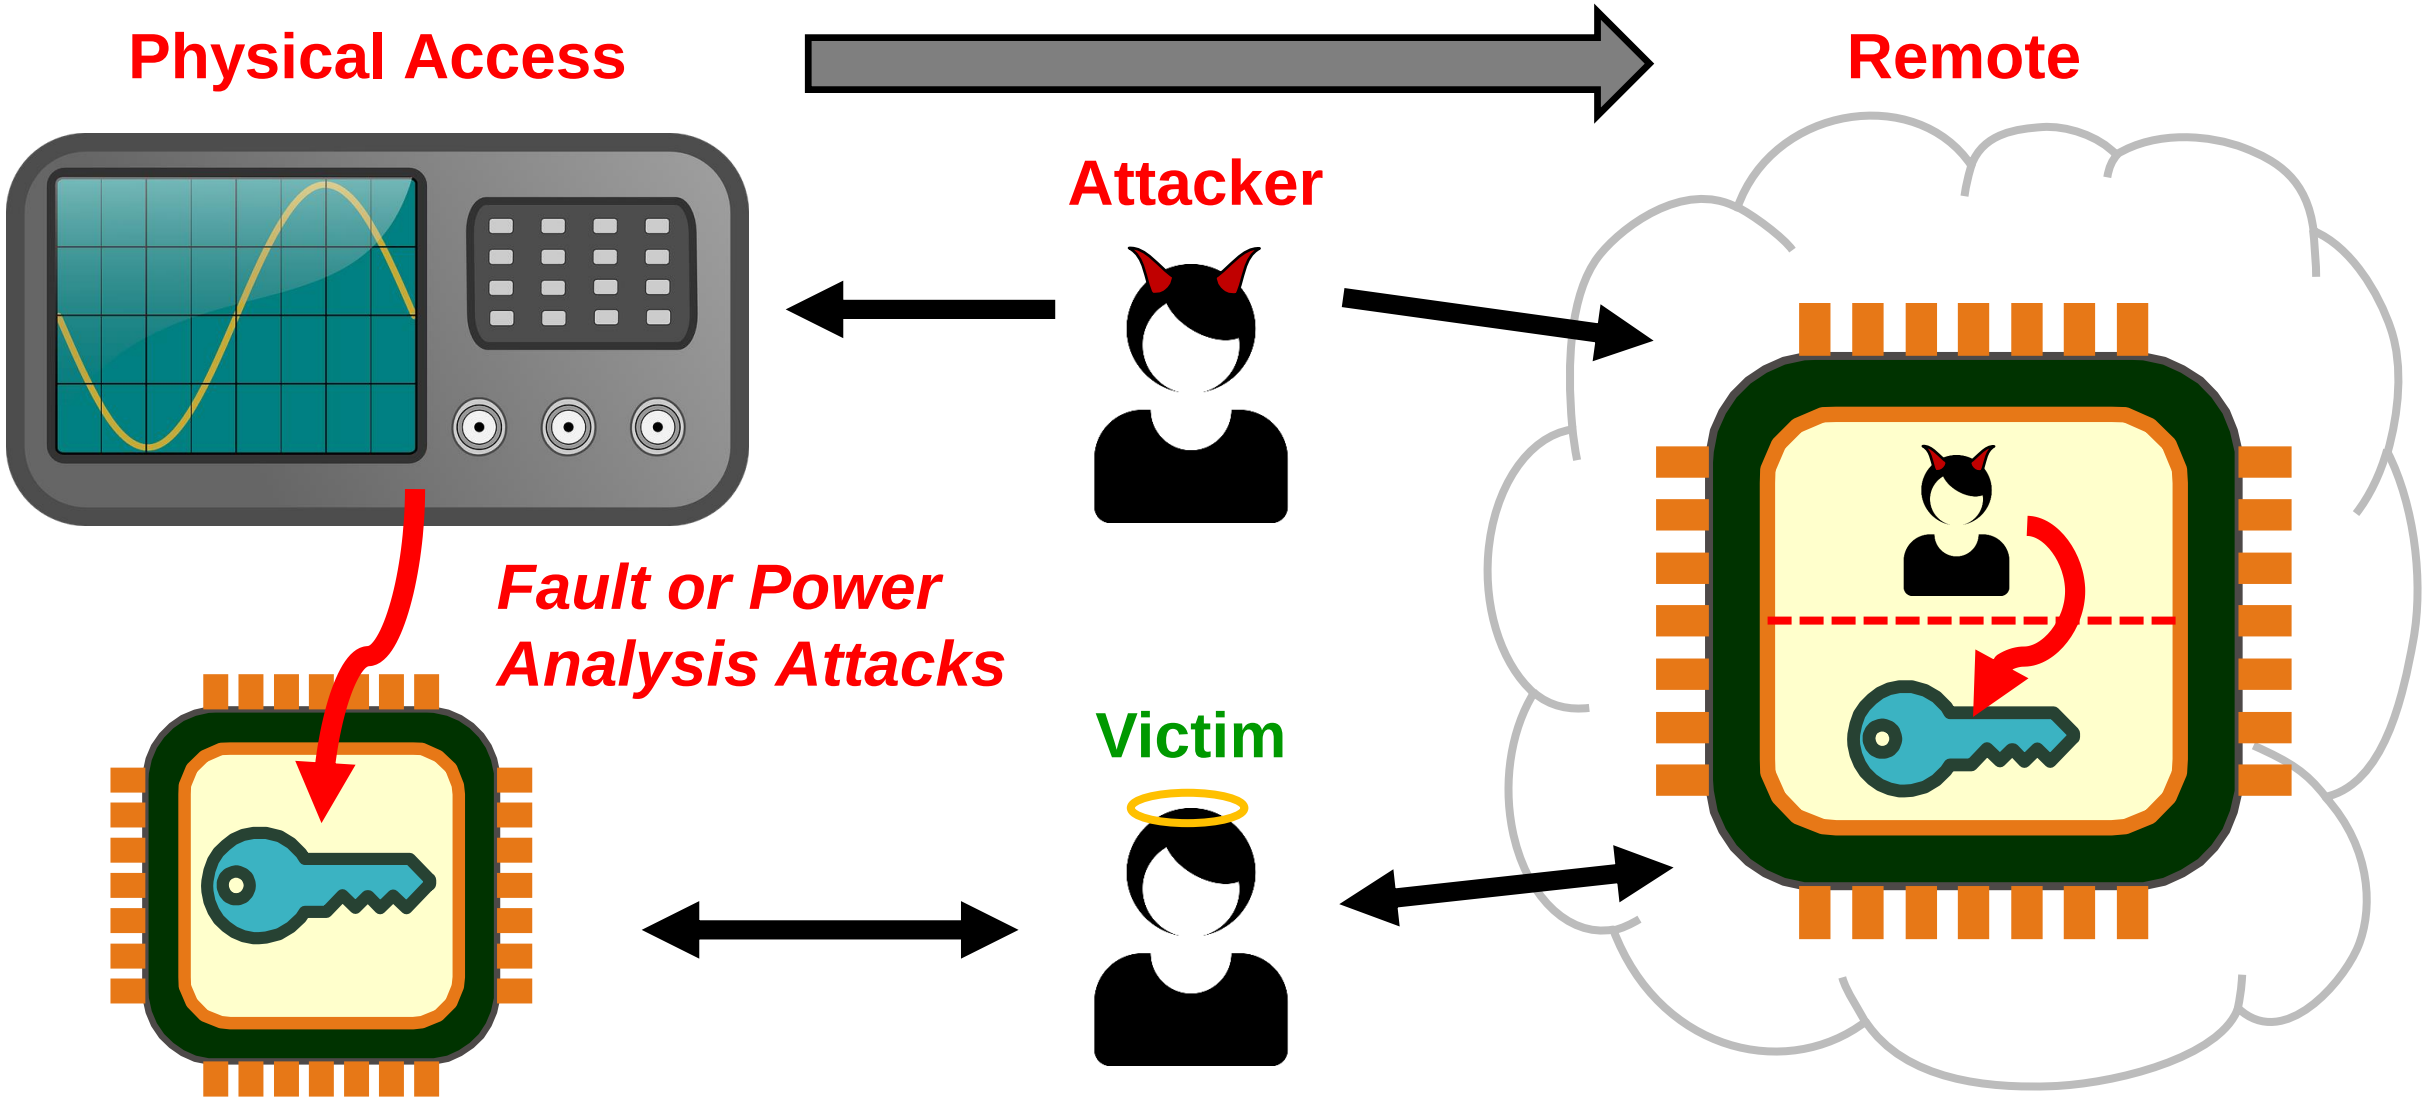

[Images: pixabay.com, wikimedia.org]

Motivation

- Heterogeneous computing ↑
  - FPGAs in the Cloud: Amazon, Alibaba, Microsoft, Huawei, …
  - Highly-integrated SoCs + FPGAs
- ➢ Trends
  - Shared FPGA → “Multi-Tenant FPGA”
  - Remote accessible FPGA
- ➢ New security threats?
  - Isolation Design Flows **protect on digital level**
  - New security threats on **electrical level**

Figure 3: **Floorplan of the XCU250 Device**

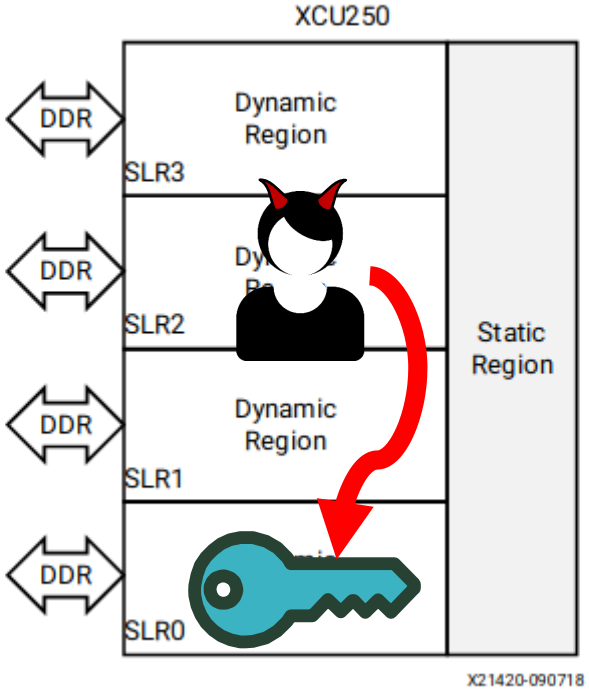

[Xilinx Datasheet XCU250]

Who cares about security of FPGAs?

- Everyone will → FPGAs important for efficiency scaling!

Challenges: The End of Moore's Law and Scaling

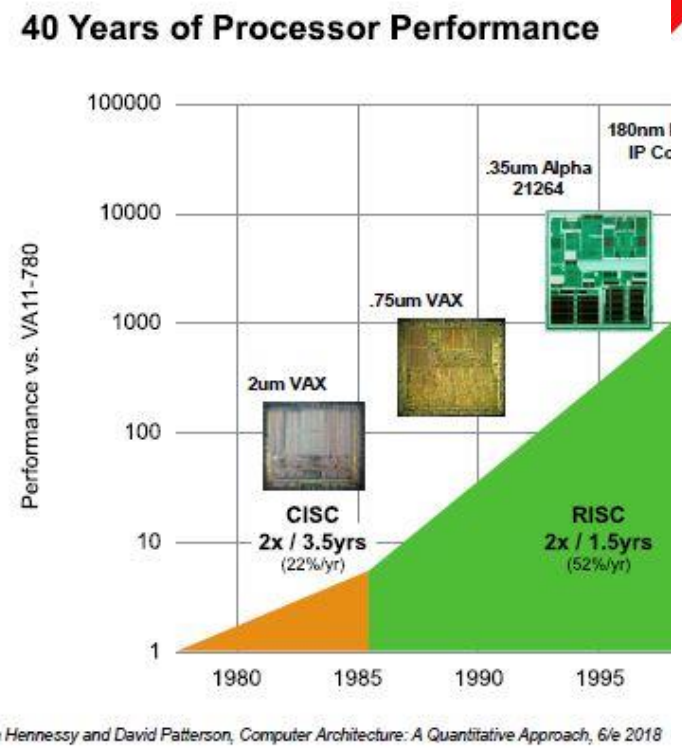

Source: John Hennessy and David Patterson, Computer Architecture: A Quantitative Approach, 6/e 2018

The Third Wave: Domain Specific Architectures on Adaptable HW

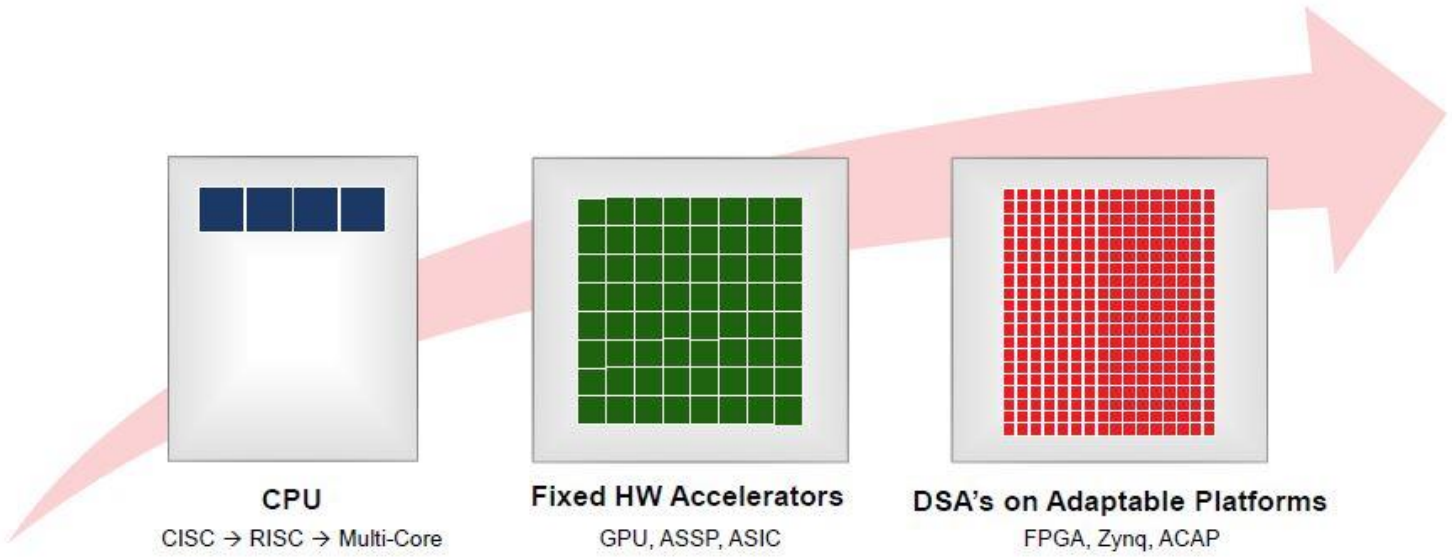

XILINX.

[Xilinx' Victor Peng @ Hot Chips 2018]

FPGAs already available in the Cloud

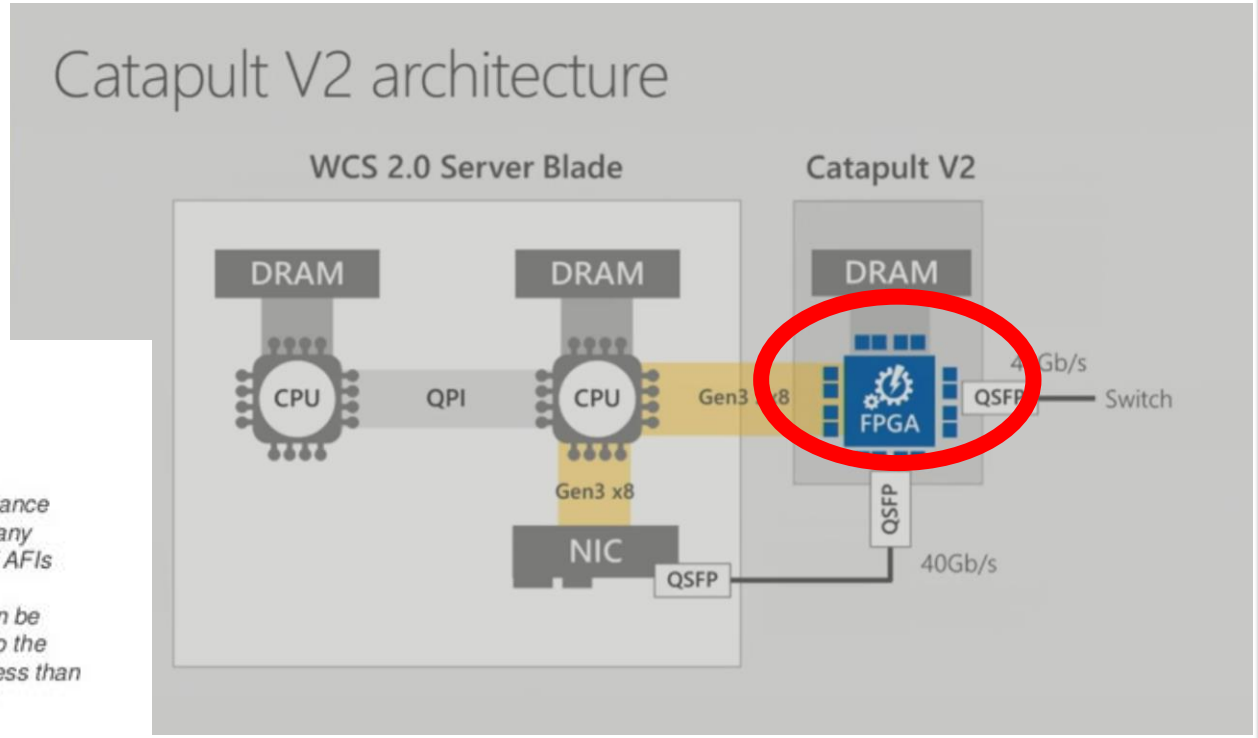

FPGA Acceleration Using F1

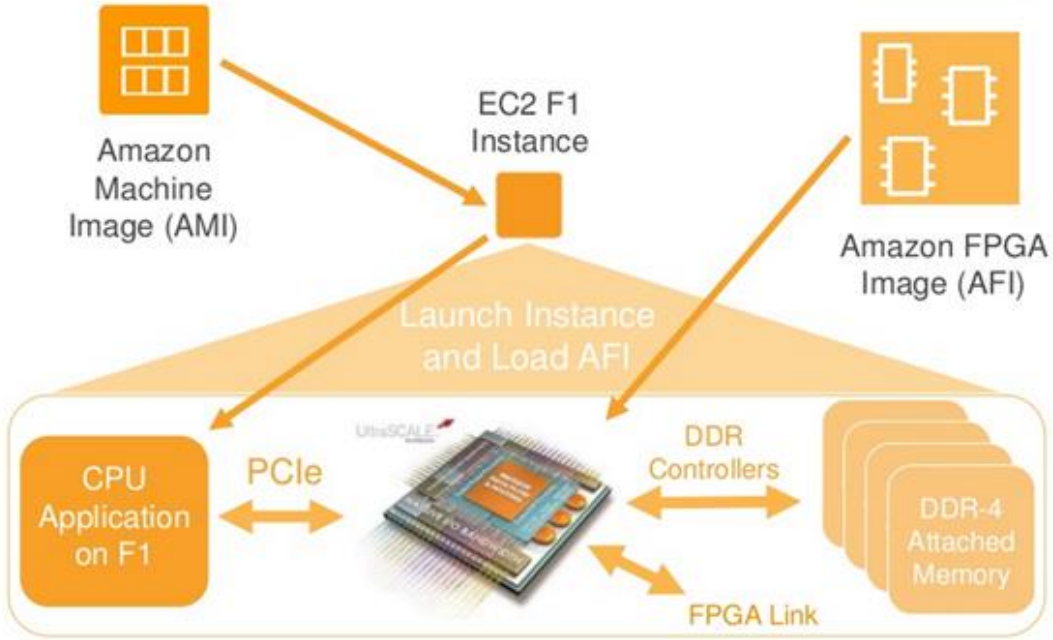

An F1 instance can have any number of AFIs

An AFI can be loaded into the FPGA in less than 1 second

amazon web services | Webinars

Multi-Tenant FPGAs

- Mainline Linux Kernel Support for partial reconfiguration exists
- Many publications show efficiency benefits of multi-tenancy
- Datacenter FPGAs with multiple accelerators considered

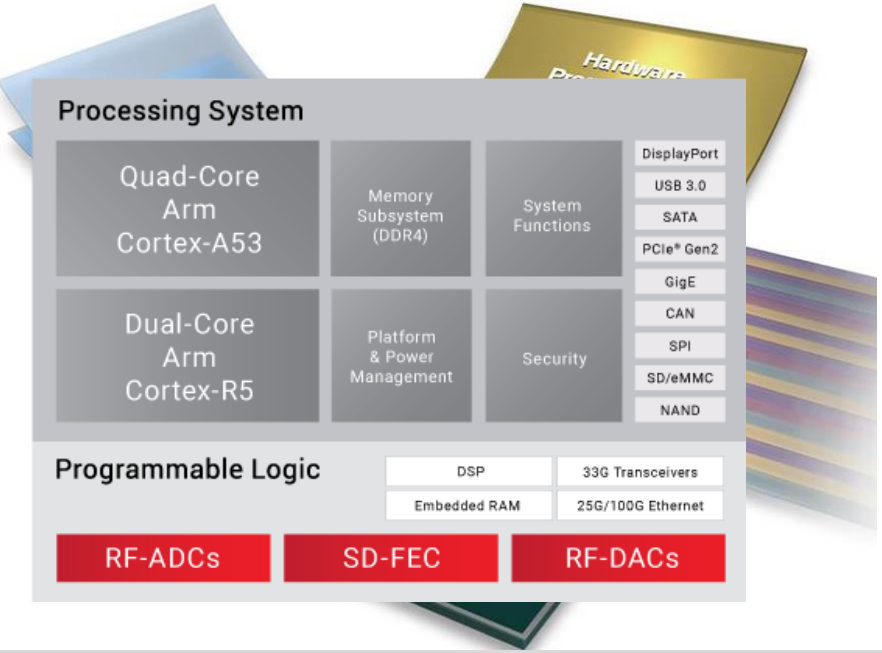

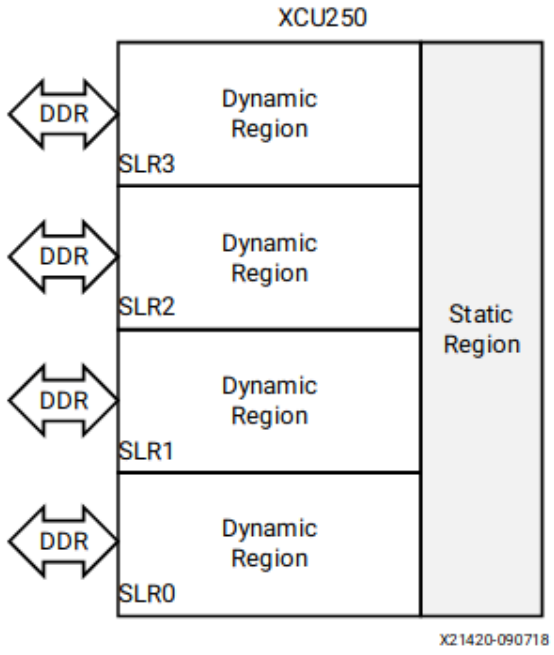

[Xilinx Datasheet XCU250]

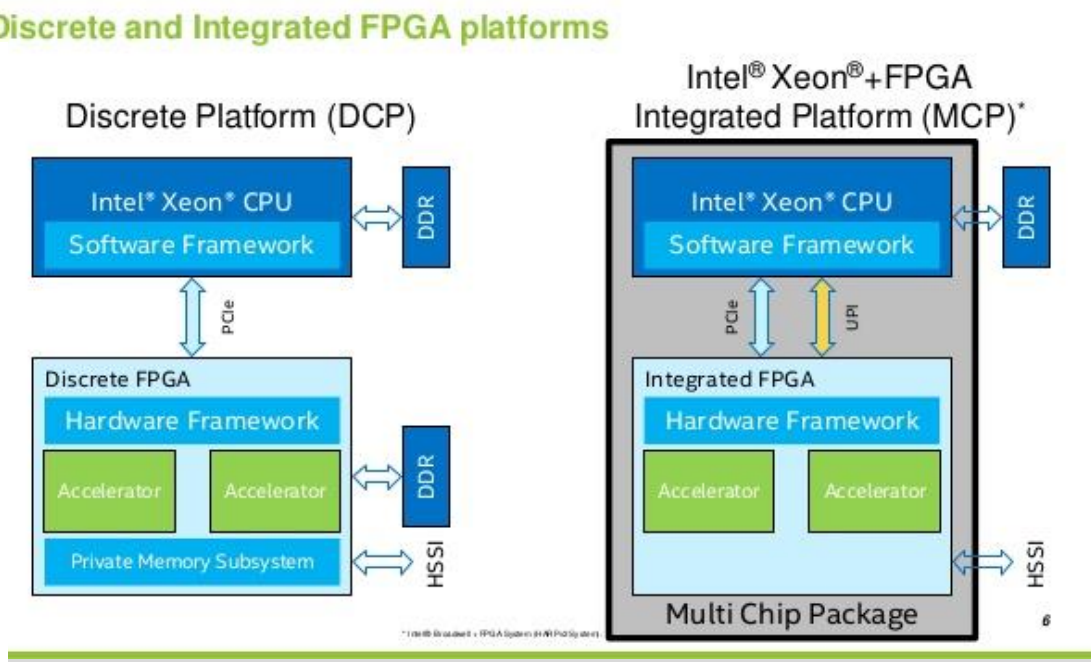

Multi-Tenant FPGAs coming to the Cloud

Stack for Application Acceleration including ML

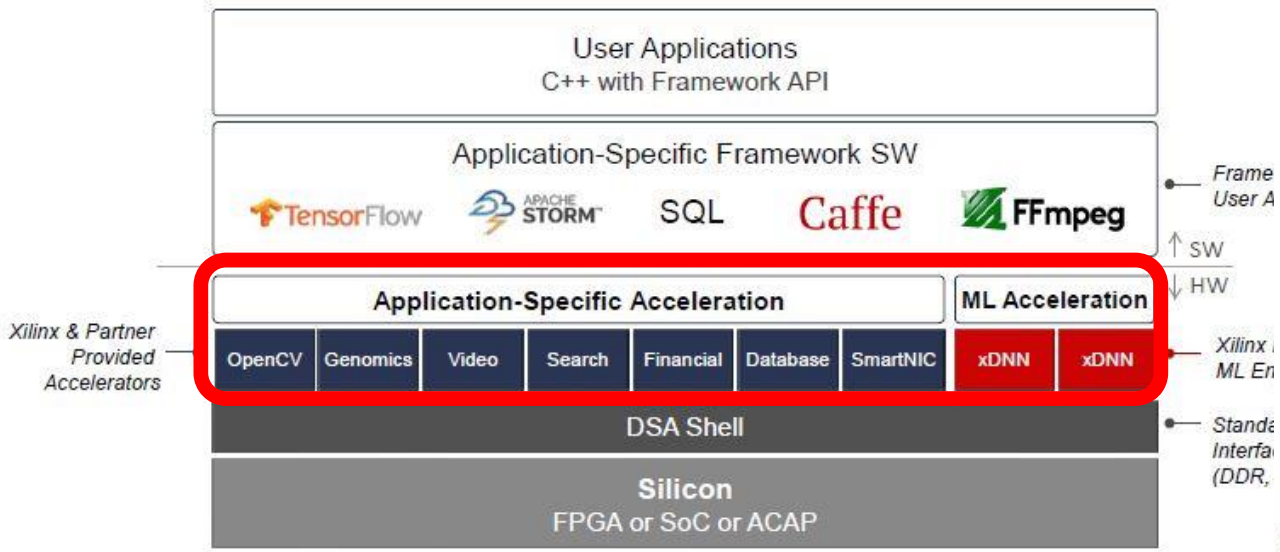

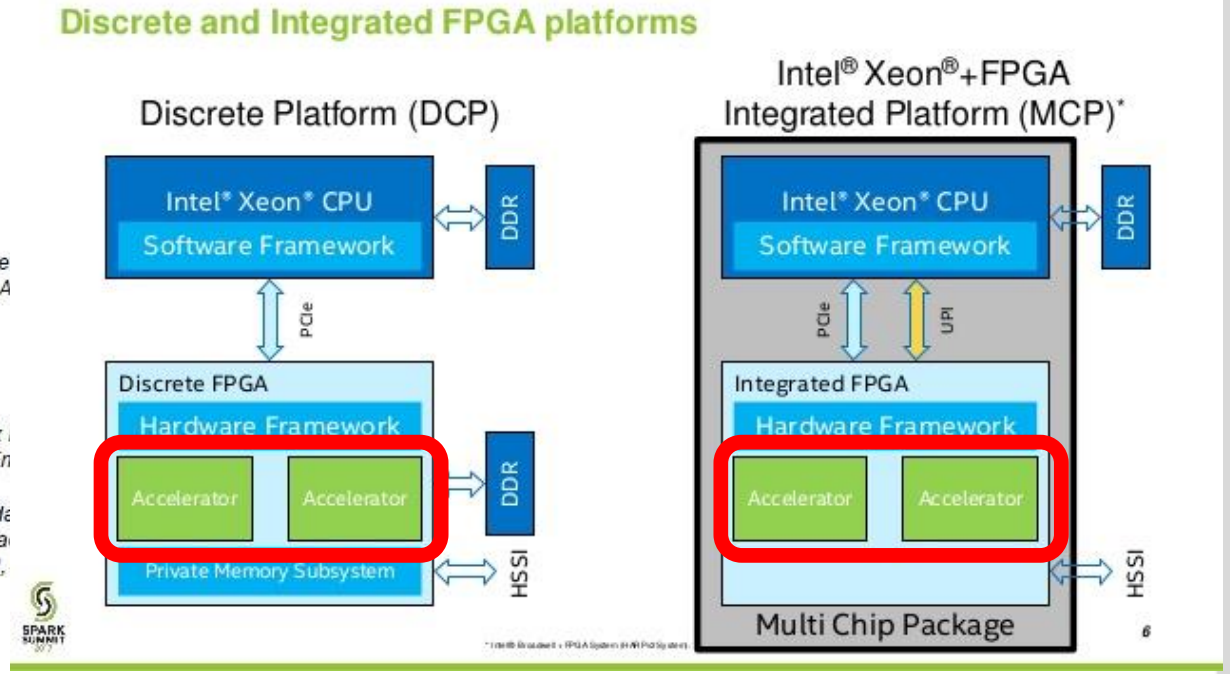

[Xilinx' Victor Peng @ Hot Chips 2018]

- Power Analysis can become a threat!

Outline

- Background and Attack Model
- Attacks
  - Power Analysis Side Channel Attacks inside FPGAs
  - Fault Attacks inside FPGAs
  - Summary & Impact
- Countermeasures
- Conclusion & Perspectives

Outline

- Background and Attack Model
- Attacks
- Countermeasures
- Conclusion & Perspectives

Overview: Physical Attacks on Integrated Circuits

- Traditionally with physical access to the chip

(Semi-)Invasive Attacks

- (Semi-)Destructive
- **Not here**

Non-Invasive Attacks

- No manipulation of chip itself

Active

- Often: Clock / Voltage manipulation
- “Fault Attacks”

Passive

- Measuring Power / Voltage / ..
- “Side Channel Attacks”

Attack Scenarios

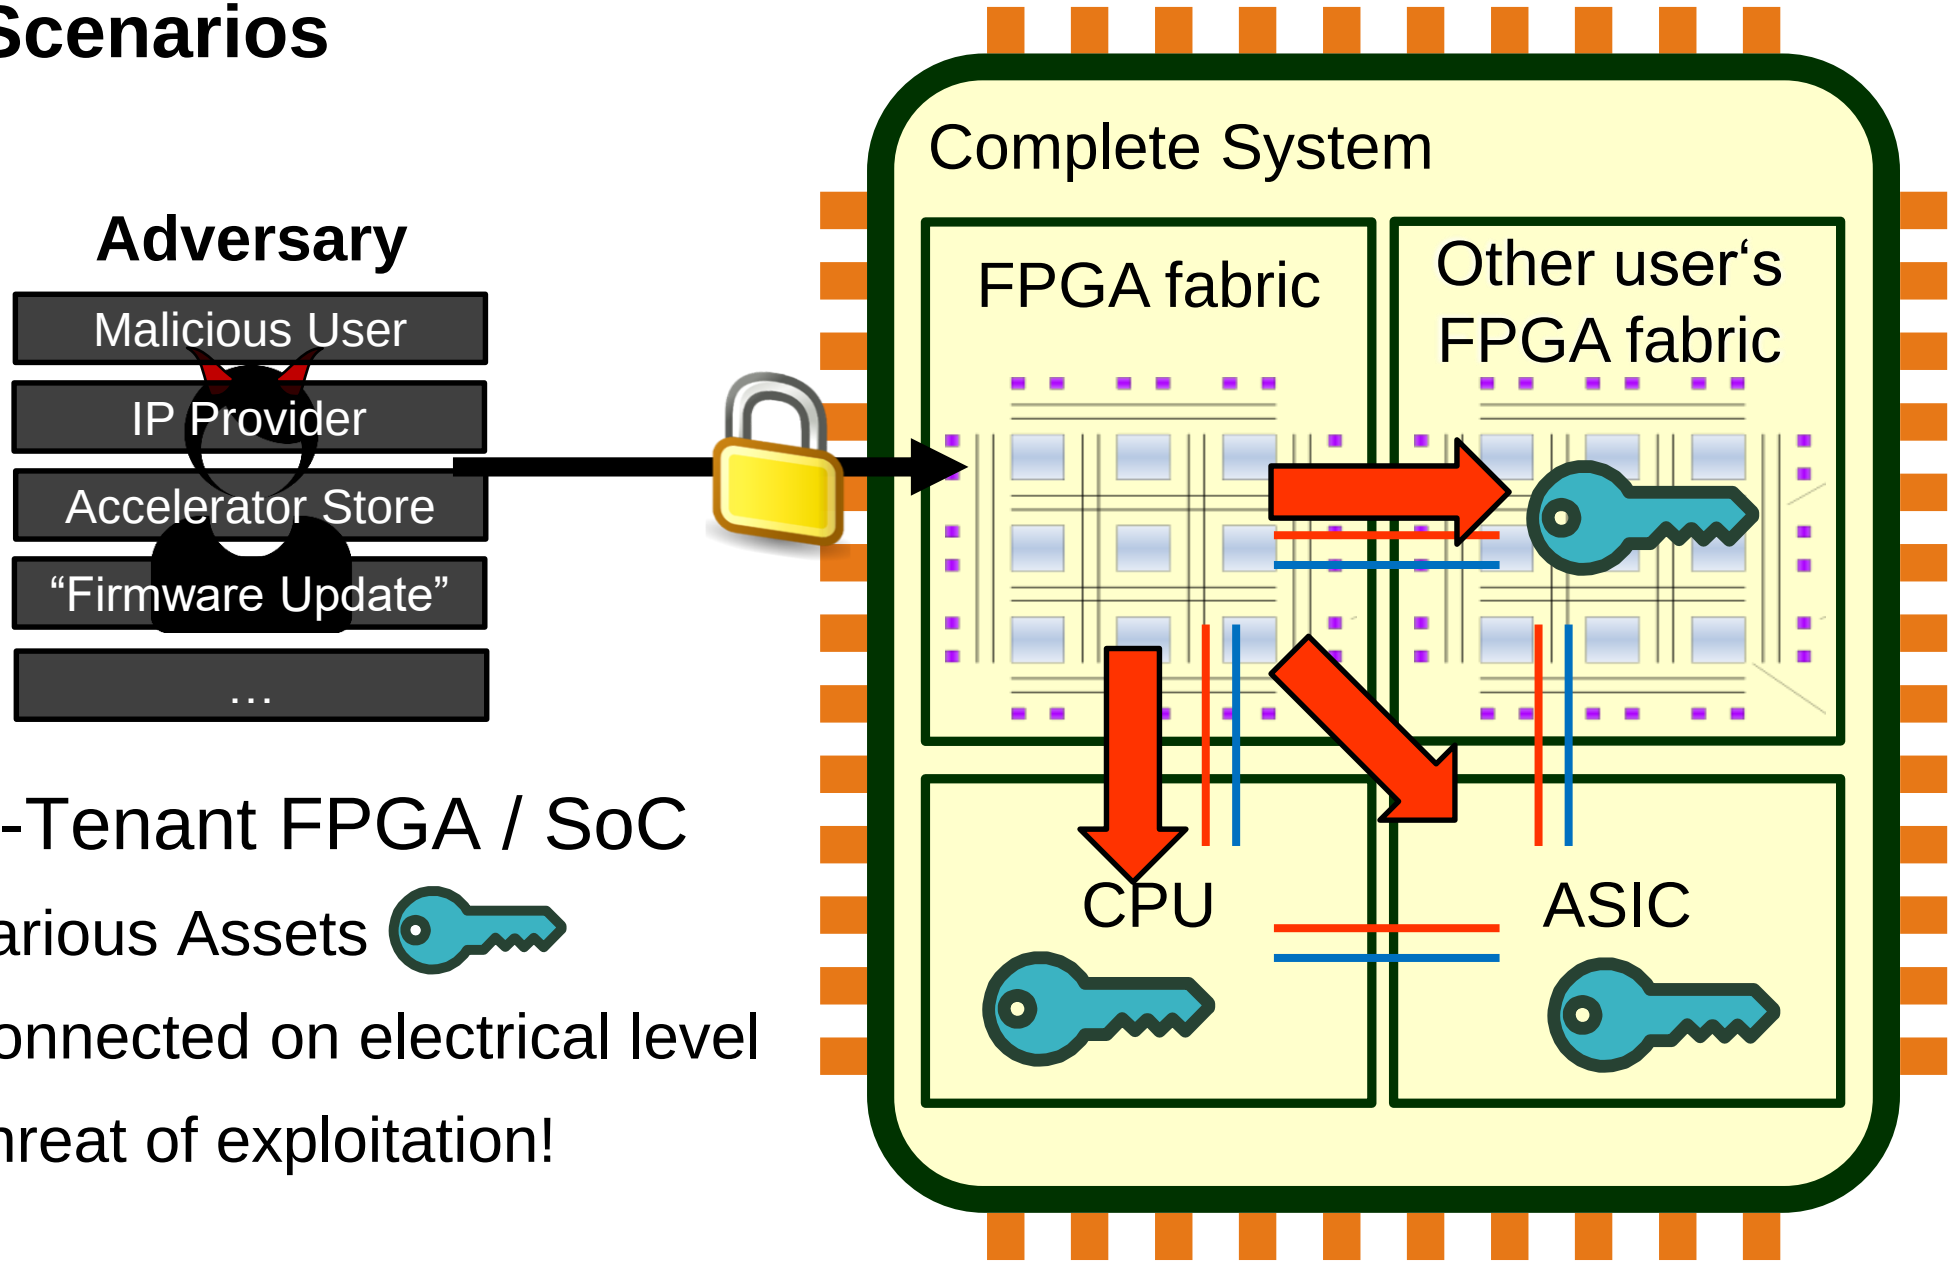

- Multi-Tenant FPGA / SoC
  - Various Assets
  - Connected on electrical level
  - Threat of exploitation!

Power Side-Channel Attacks

Kocher et al. ’99, Brier et al. ’04

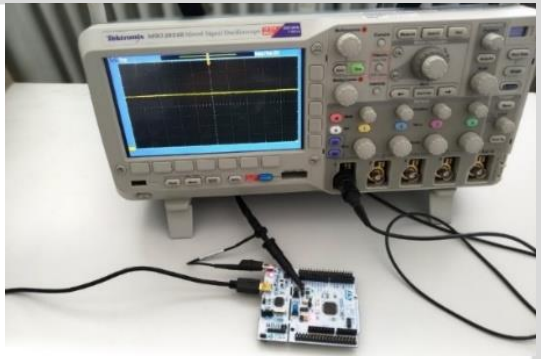

- Extract secret key from power measurements
- Example: *Correlation Power Analysis* (CPA) statistically correlates:
  I. Measure power
  II. Model of key-dependent power consumption (e.g. per byte)
  ➢ Correct secret key correlates most

Result for 10 Million traces

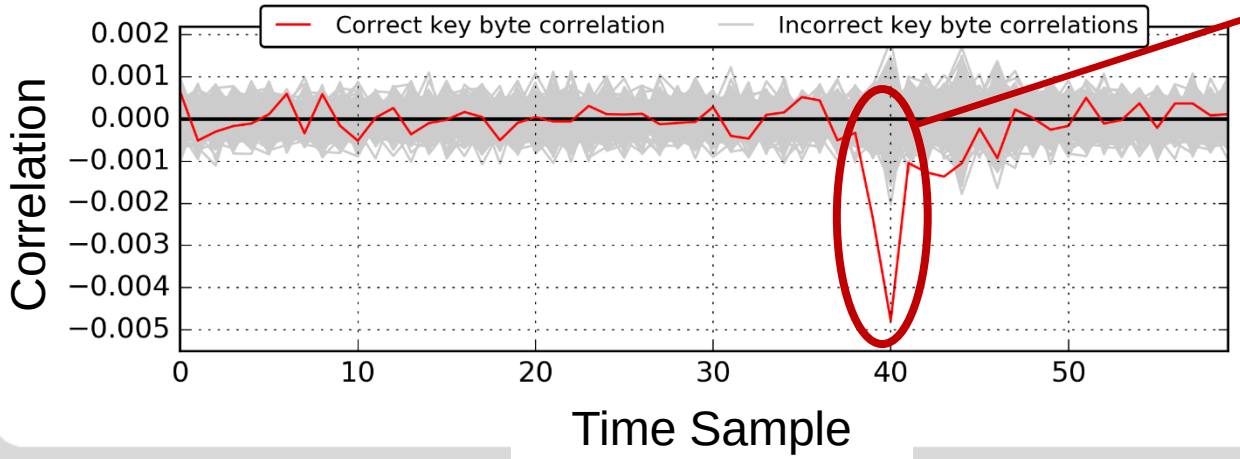

Dependent on amount of traces…

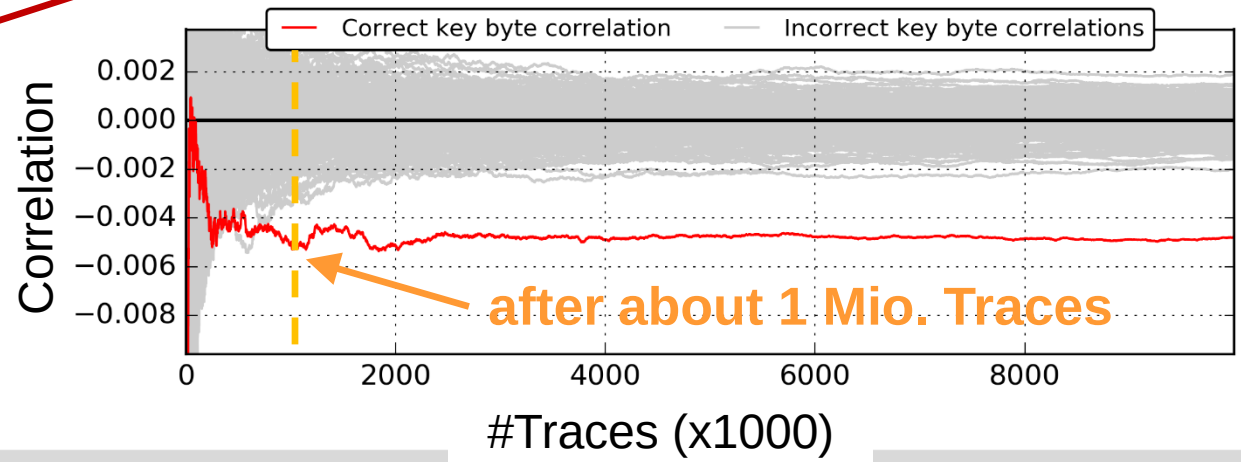

Fault Attacks

- Fault Attacks
  - Cause errors in a system
  - Targeted manipulation (possibly undetected)
- ➢ Denial of Service
  - i.e. massive voltage drop
- ➢ Differential Fault Analysis (DFA) [Biha97]
  - Cause minor changes in encryption algorithm
  - Similarities to CPA: Collect data and try to deduce secret key
- ➢ Circumventing Secure Boot or other Security Checks
  - Causing instruction skips

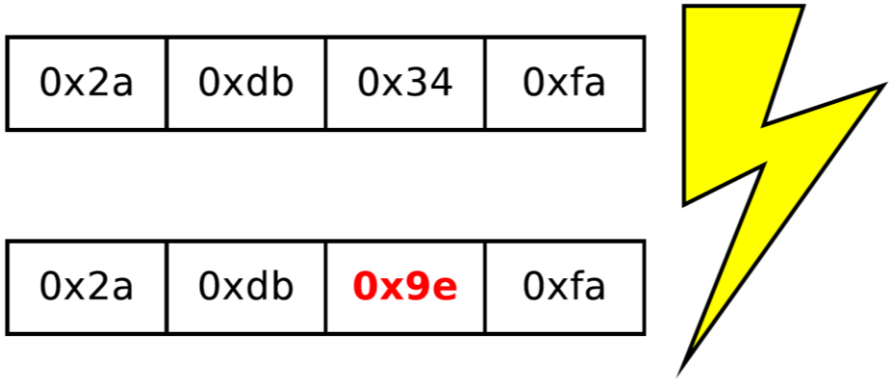

Background – Power Distribution Networks (PDNs)

- Board-level PDN: From Voltage regulator …
- …until individual Transistors

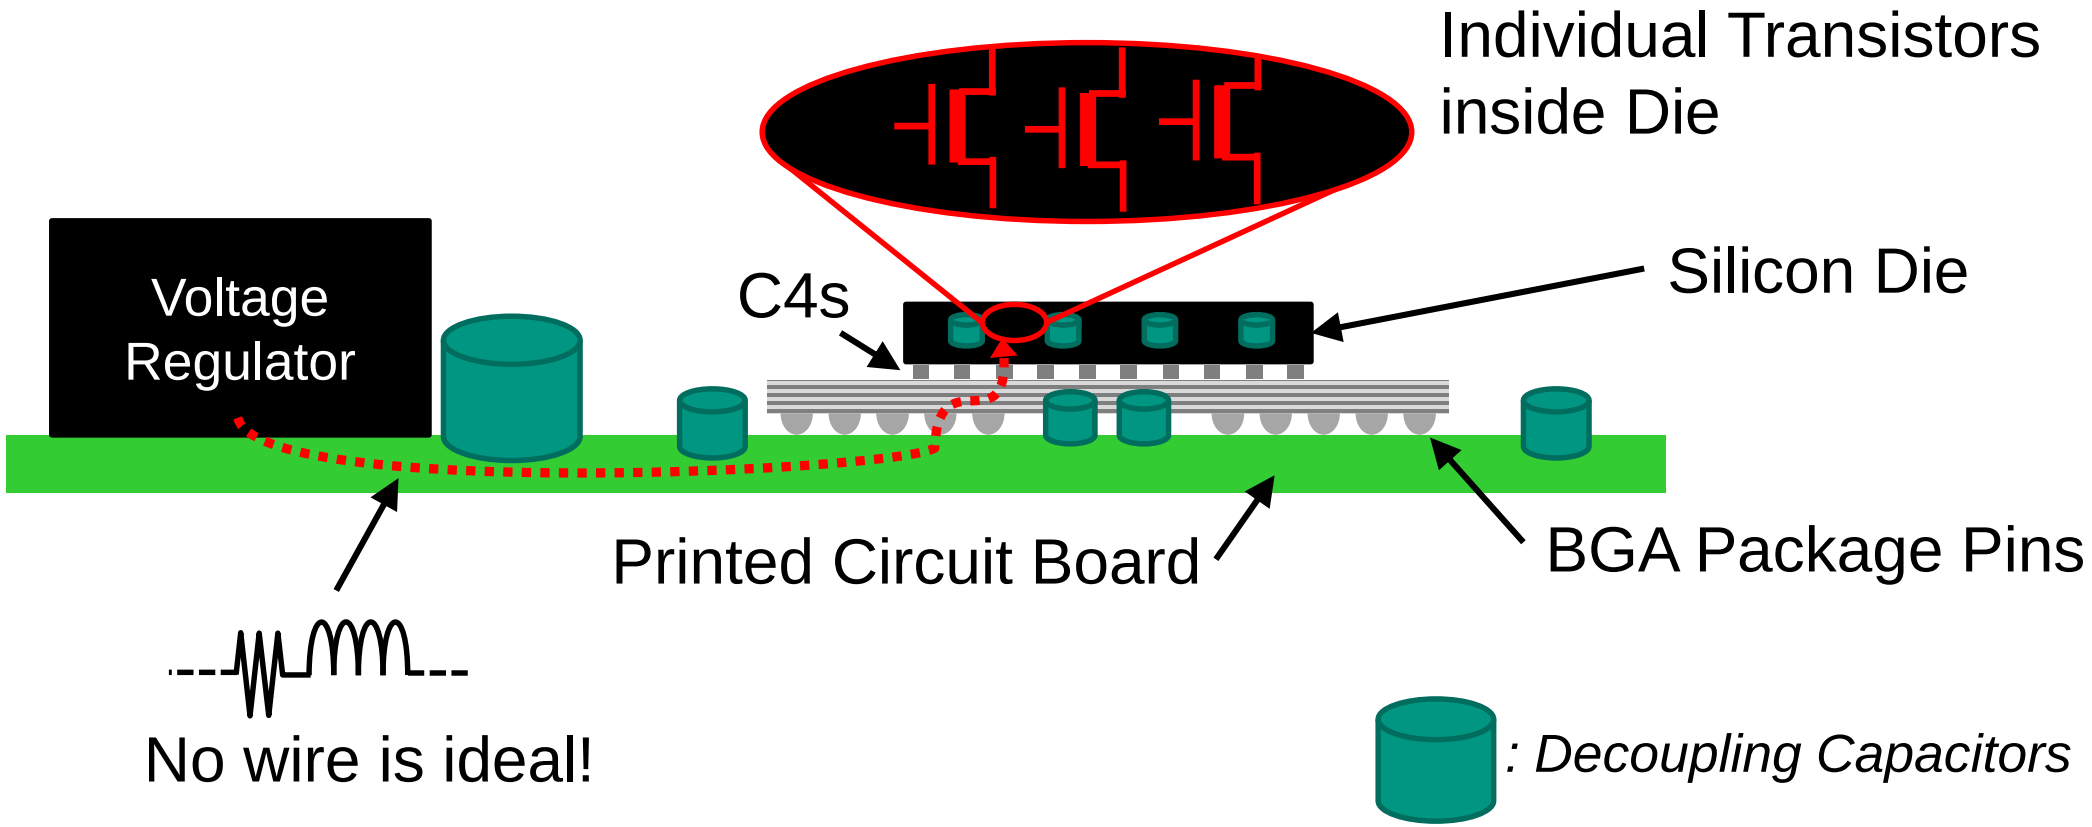

Background – Power Distribution Networks (PDNs)

- Complex network: Resistors (R), Capacitors (C), Inductors (L)
  - Some *by design*, some unwanted, *parasitic*
- Circuit activity → changes i(t) → voltage noise $V_{noise} \propto L\frac{di(t)}{dt} + i(t)R$
- Vice versa: Voltage → Transistor speed $\tau_{delay} \propto \frac{1}{V}$

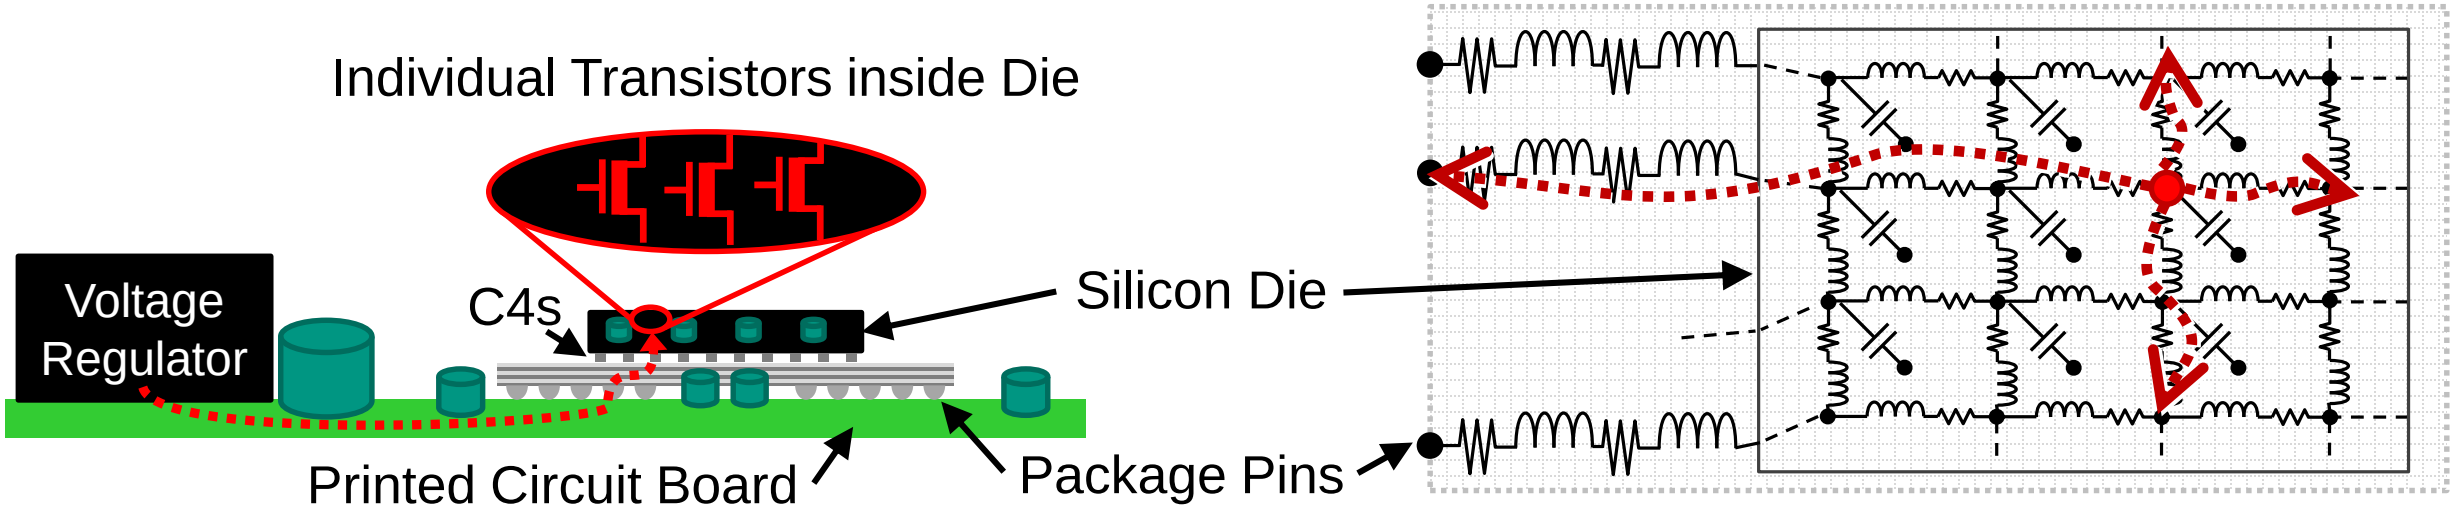

Sensors in Digital FPGA logic

Zick et al. ’10,’13

- Voltage-dependent transistor delay
- Ring Oscillator based sensor
  - Slow sampling
- Sensor based on Delay Line
  - No combinational loop needed
  - Faster sampling
  - Sensitivity was not explored

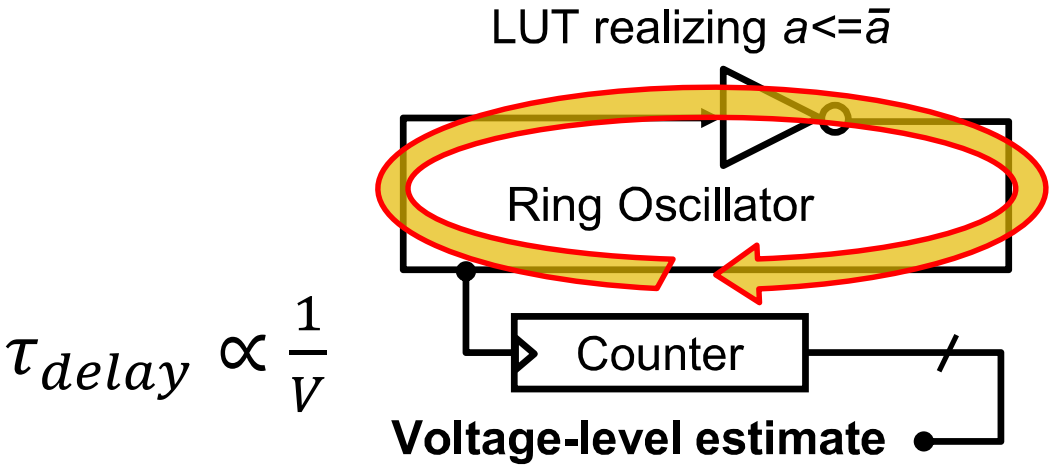

$$\tau_{delay} \propto \frac{1}{V}$$

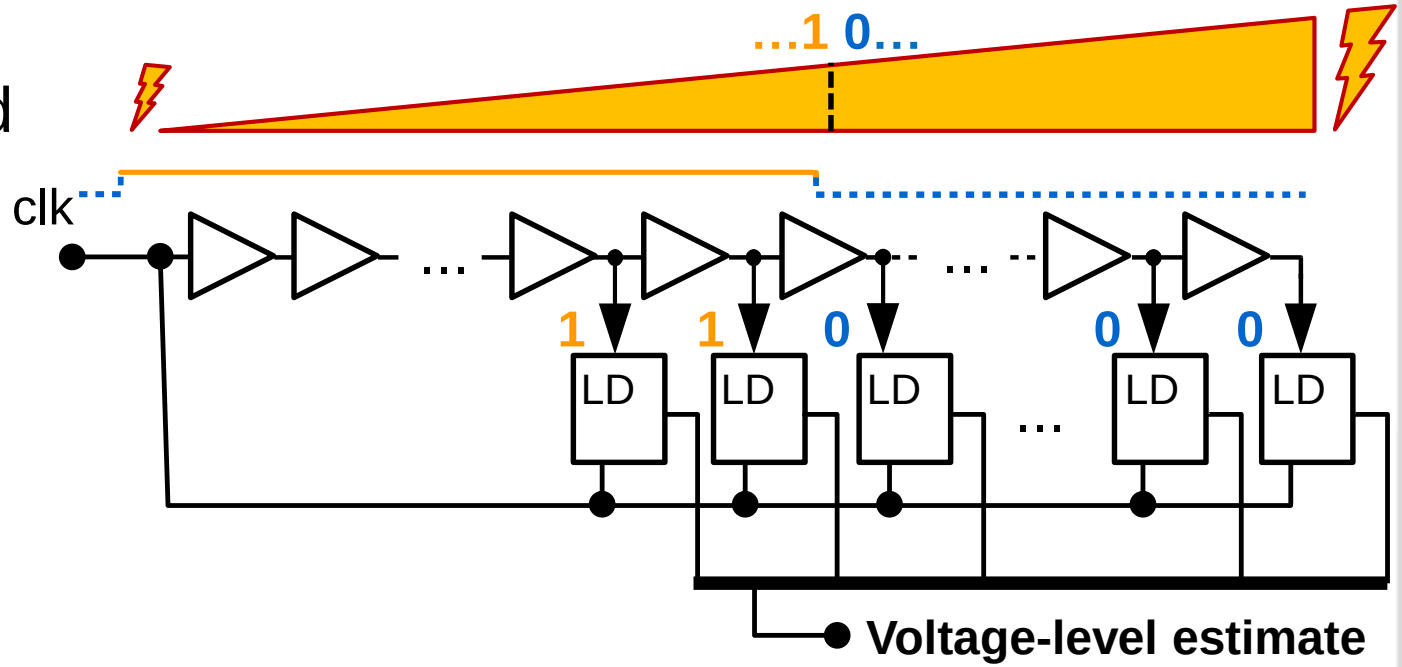

Outline

- Background and Attack Model
- Attacks
  - Power Analysis Side Channel Attacks inside FPGAs
  - Fault Attacks inside FPGAs
- Countermeasures
- Conclusion & Perspectives

Experimental Setup

SAKURA-G

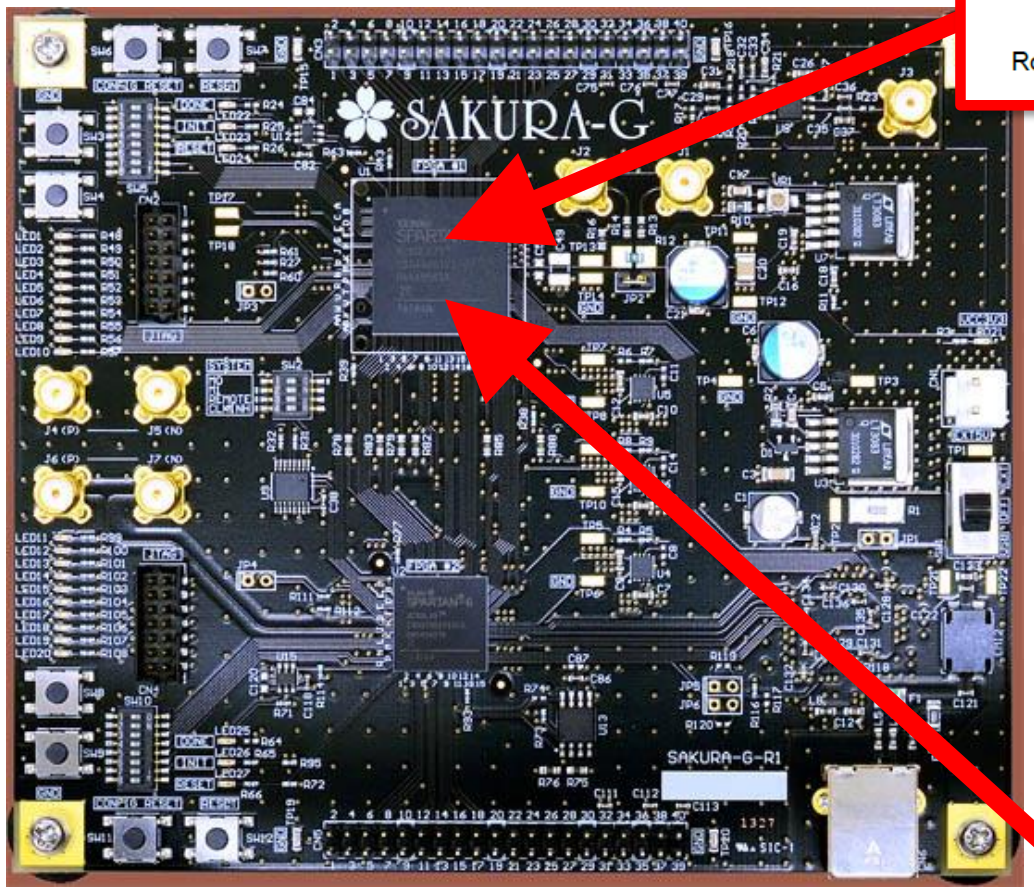

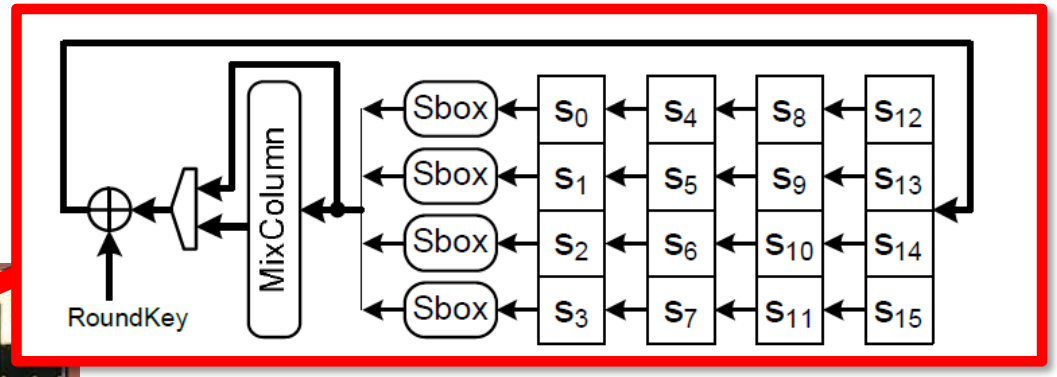

- AES @ 24 MHz
  - Minimal FPGA resource use:
  - **265 FF (0.3%) / 862 LUTs (0.9%)**
- Sensor @ 24, 48, 72, 96 MHz

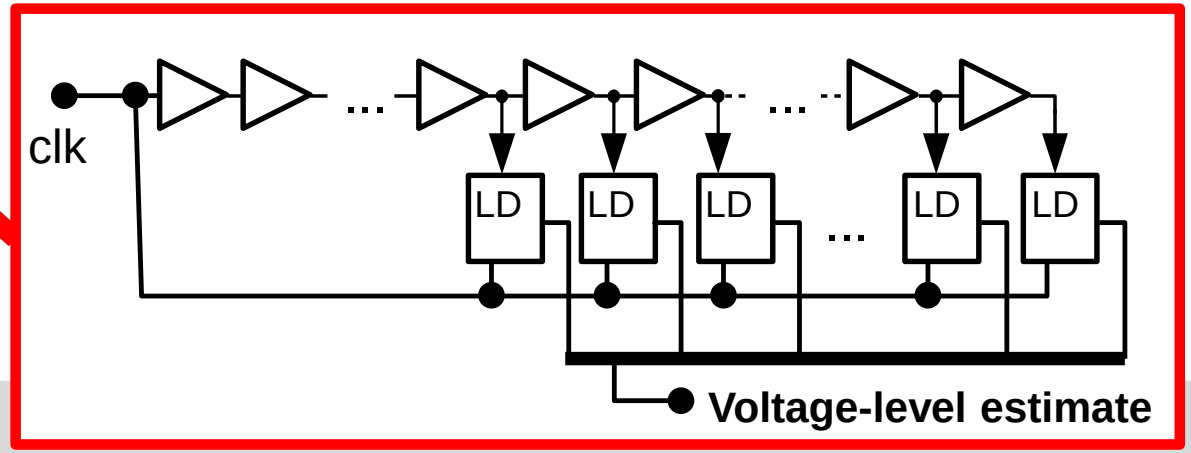

Experimental Setup – Floorplans

Nearby Sensor
(still 2 slices distance)

Distant Sensor

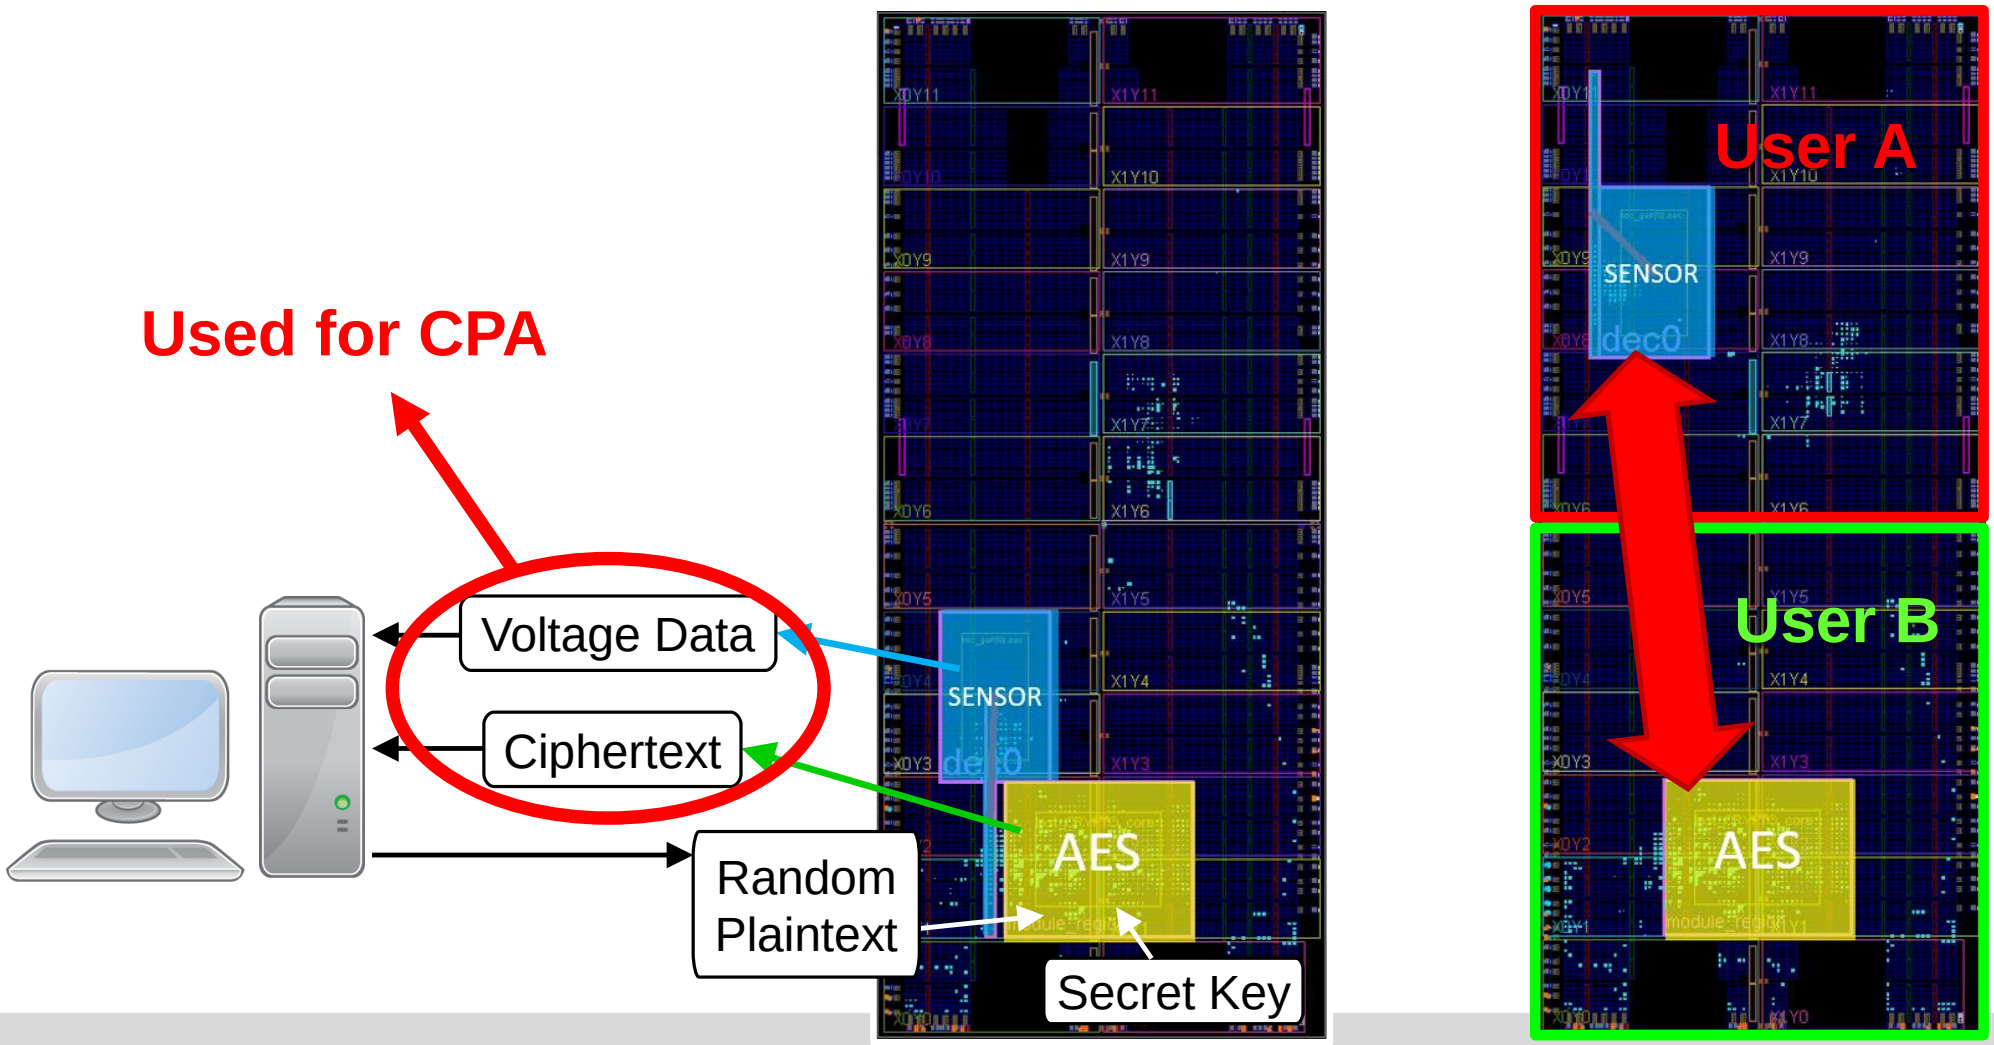

“Multi-Tenant”

Results

- Reference:
  - Ext. Oscilloscope

- Nearby Sensor:

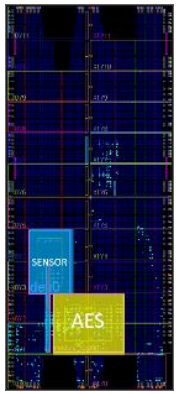

- Distant Sensor:

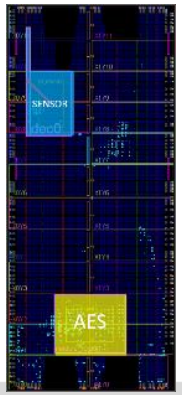

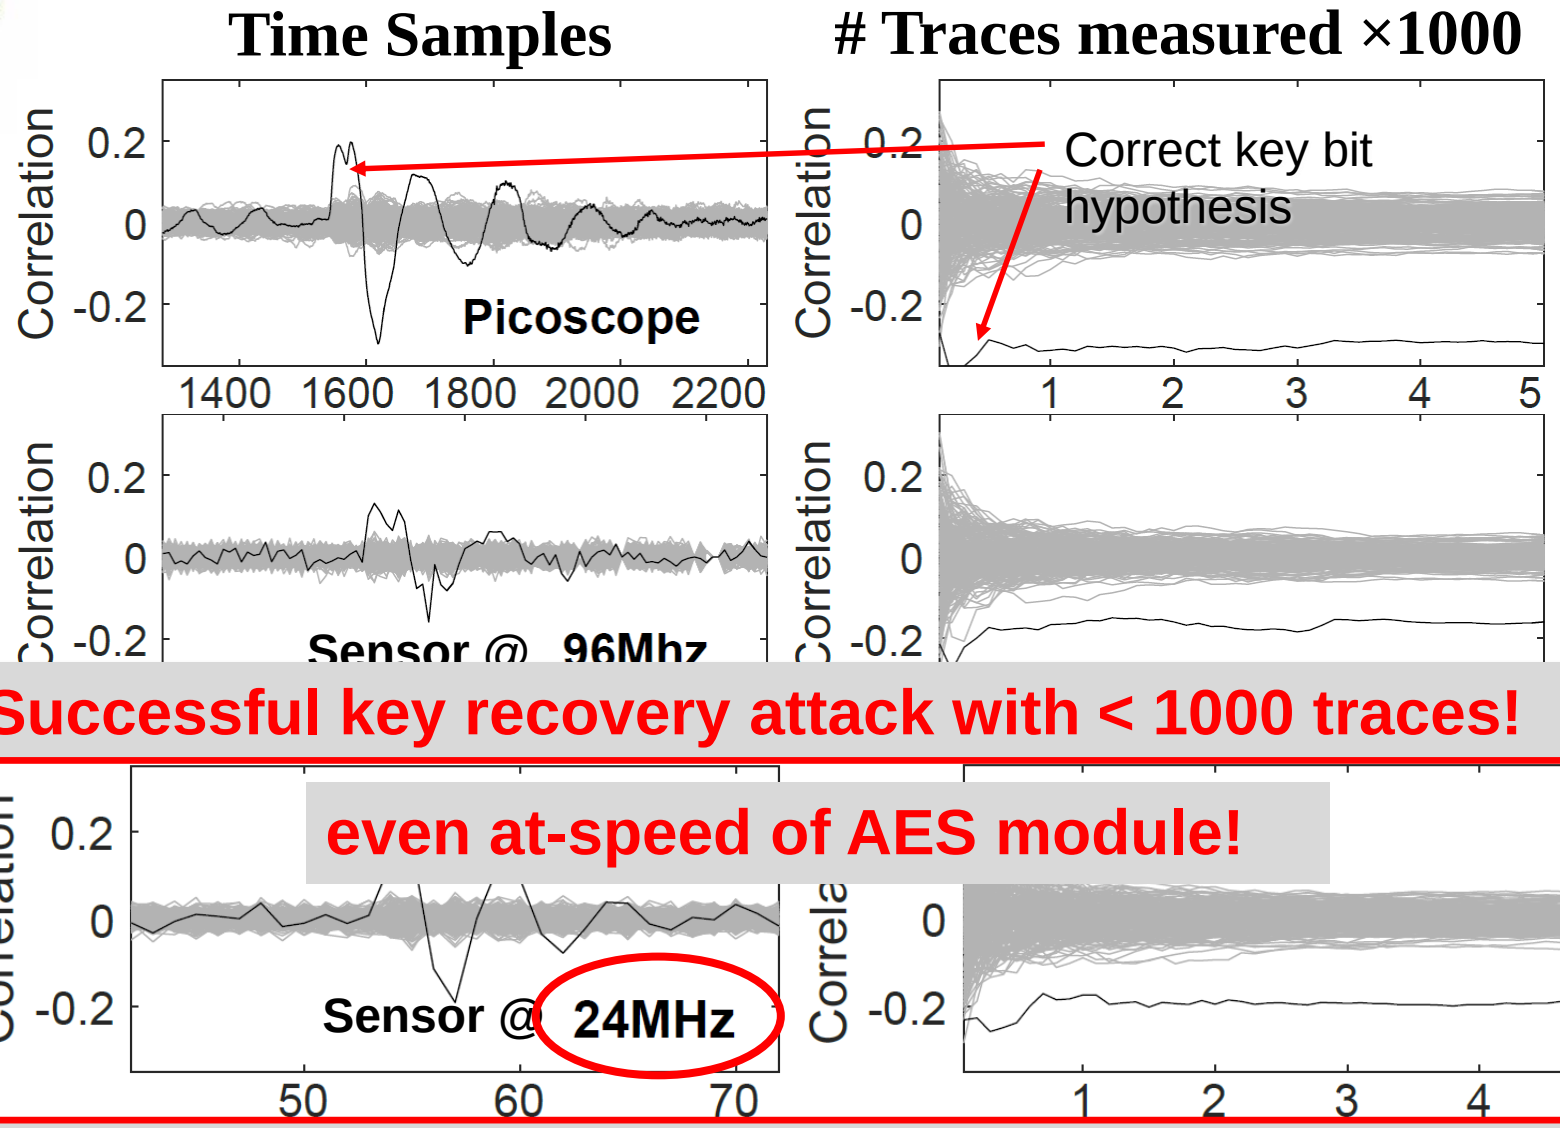

Outline

- Background & Attack Model
- Attacks
  - Power Analysis Side Channel Attacks inside FPGAs
  - Fault Attacks inside FPGAs
- Countermeasures
- Conclusion & Perspectives

Causing a Voltage Drop: Ring Oscilators (ROs)

- Ring oscillators (ROs) require high current/power = i(t)
- Suddenly enabling them causes high Voltage-Drop
  - ➢ Allows Fault Injection & Denial of Service (DoS)

$$V_{drop} = L\frac{di(t)}{dt} + i(t)R$$

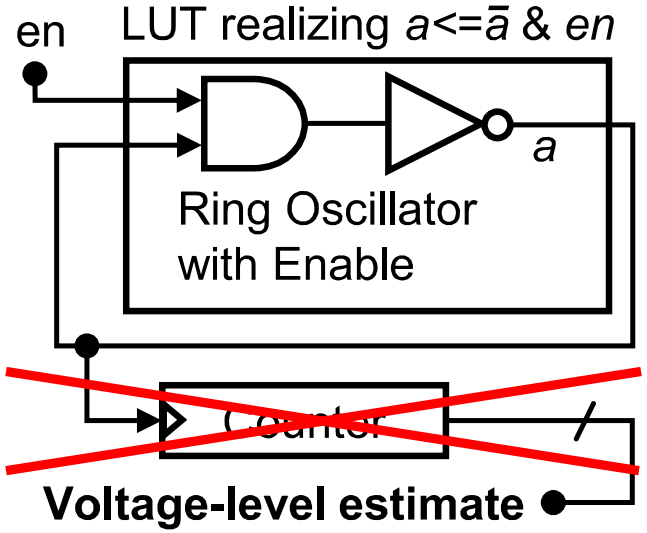

Single RO as Sensor

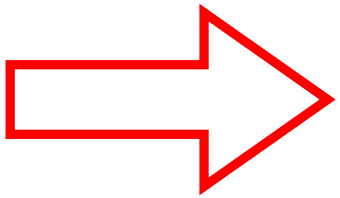

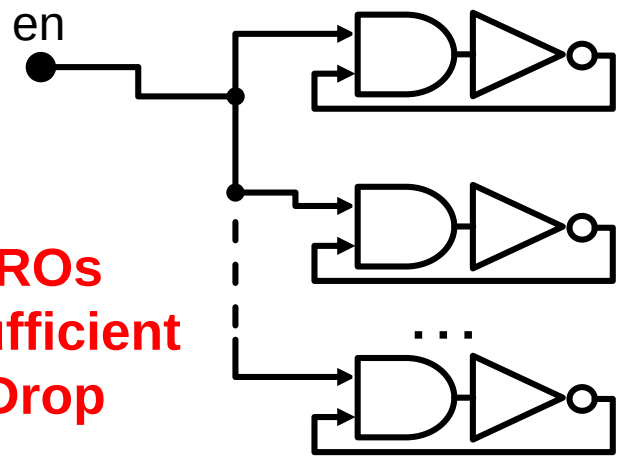

Multiple ROs cause sufficient Voltage Drop

Denial-of-Service

- Frequency-sweep on $f_{RO\text{-}t}$ crashes multiple systems
- Some can only be recovered by power cycling of the full board
  - Especially problematic for PCIe accelerators
    - KC705, ML605, VCU108
  - Standalone Boards
    - Zedboard Zynq-ZC7020, Lattice iCE40-HX8K, Intel DE1-SoC
- Details analyzed on VCU108, see next slide

Faults and Crash Probability, Xilinx VCU108

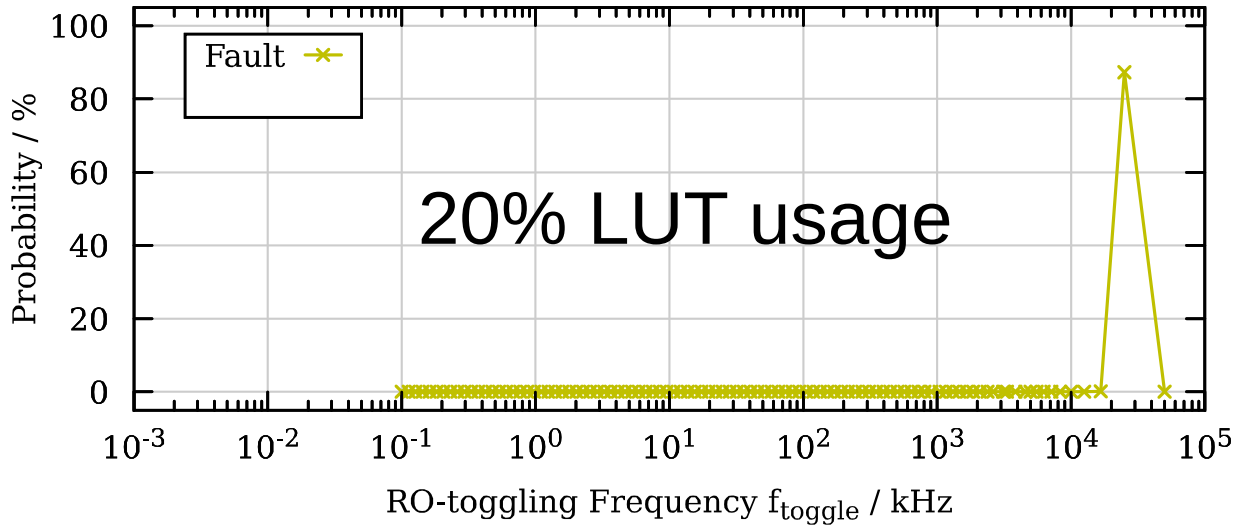

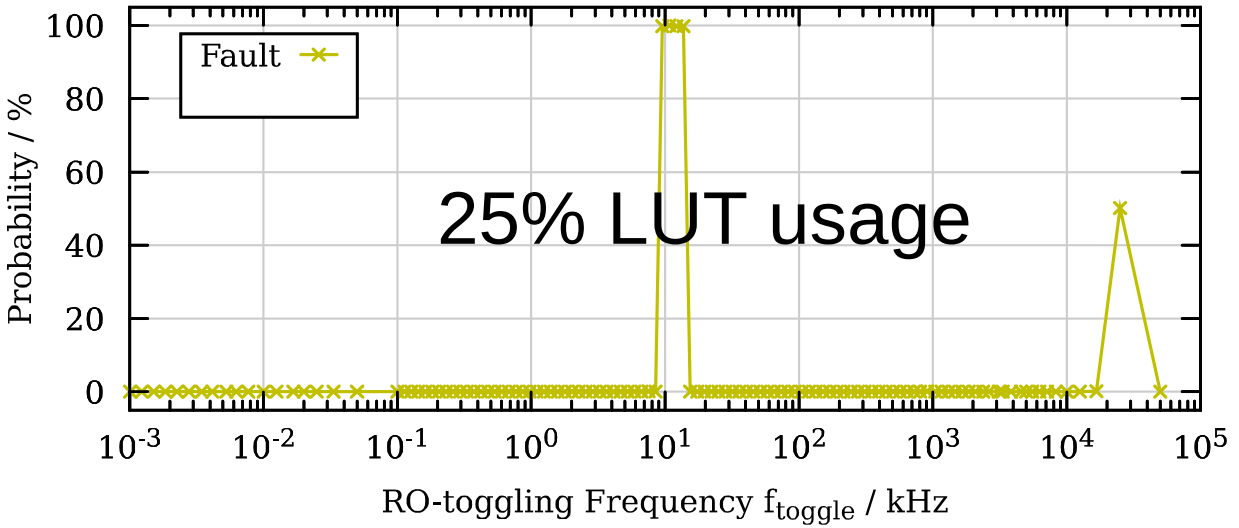

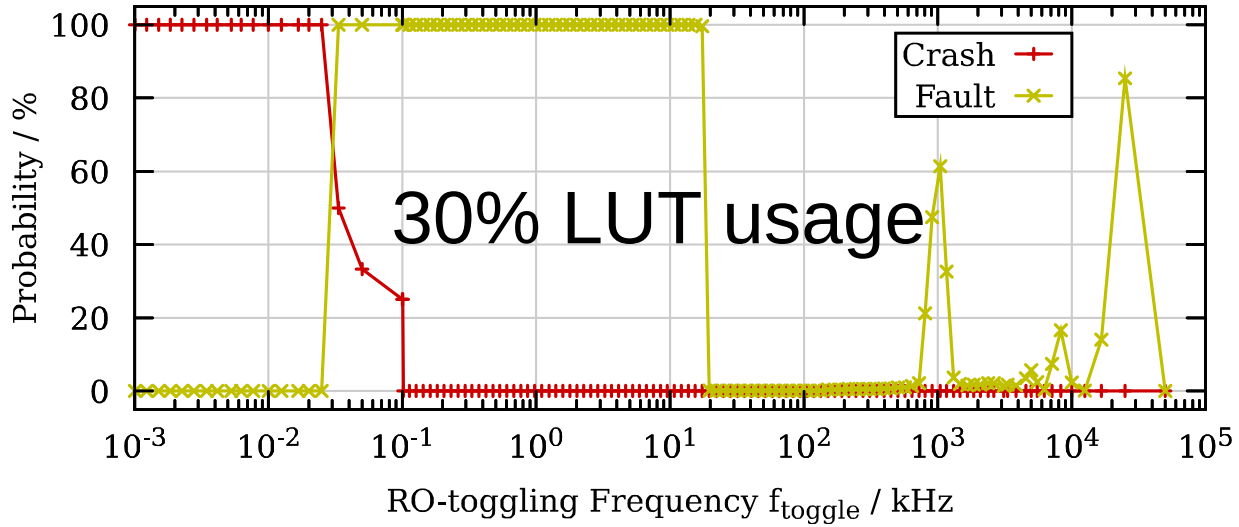

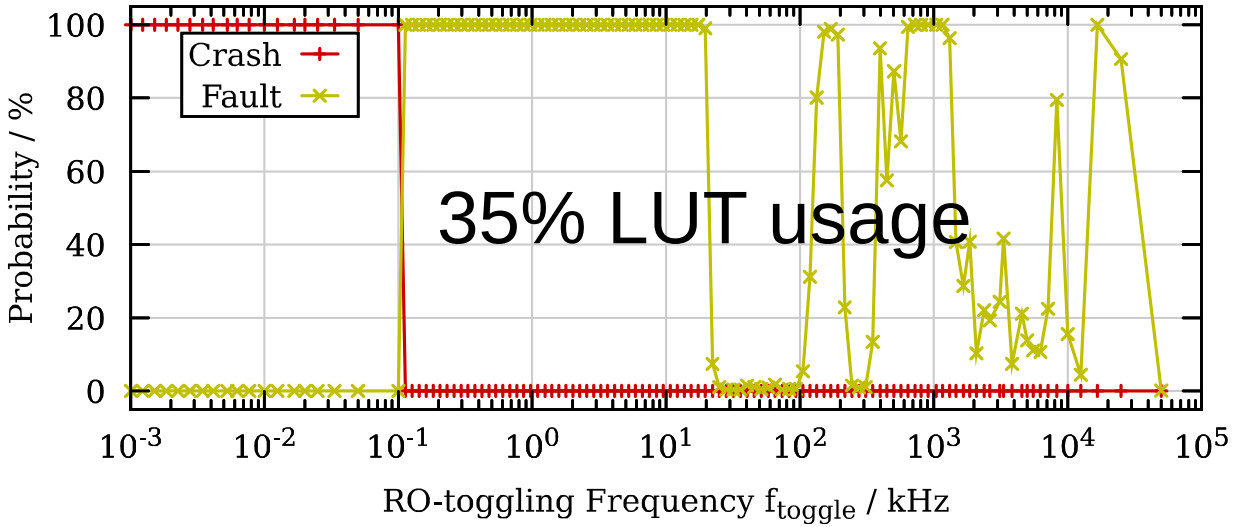

Differential Fault Attack – Scenario

- **Chosen-Plaintext Attack**
- Attacker and victim logically isolated
- Victim with public interface

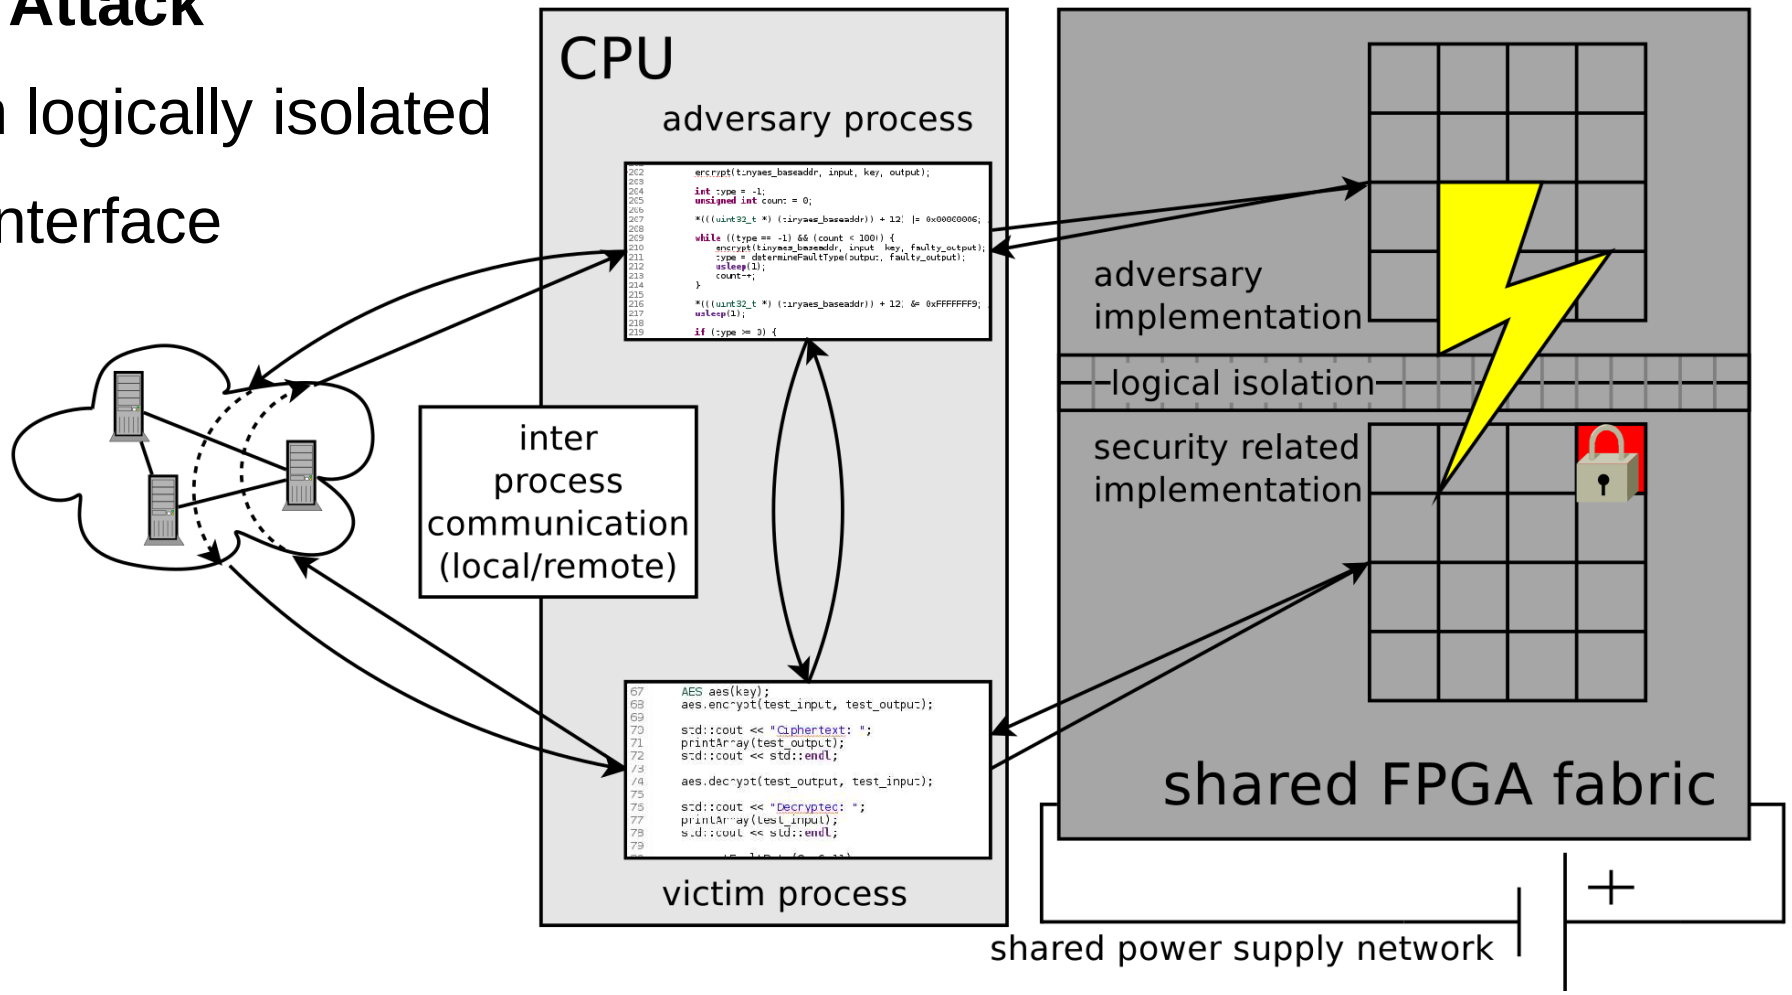

- Attacker:
  - Introduces faults
  - Gets ciphertexts

Differential Fault Attack – Analysis

- Recover AES secret key
  - Use pairs of faulty / fault-free ciphertexts – Piret et al. CHES 2003
- Original scheme: Single-byte faults before 8th round
  - ➢ All output bytes faulty
- Fault injection before 9th round
  - ➢ Allows to **verify** a successful injection

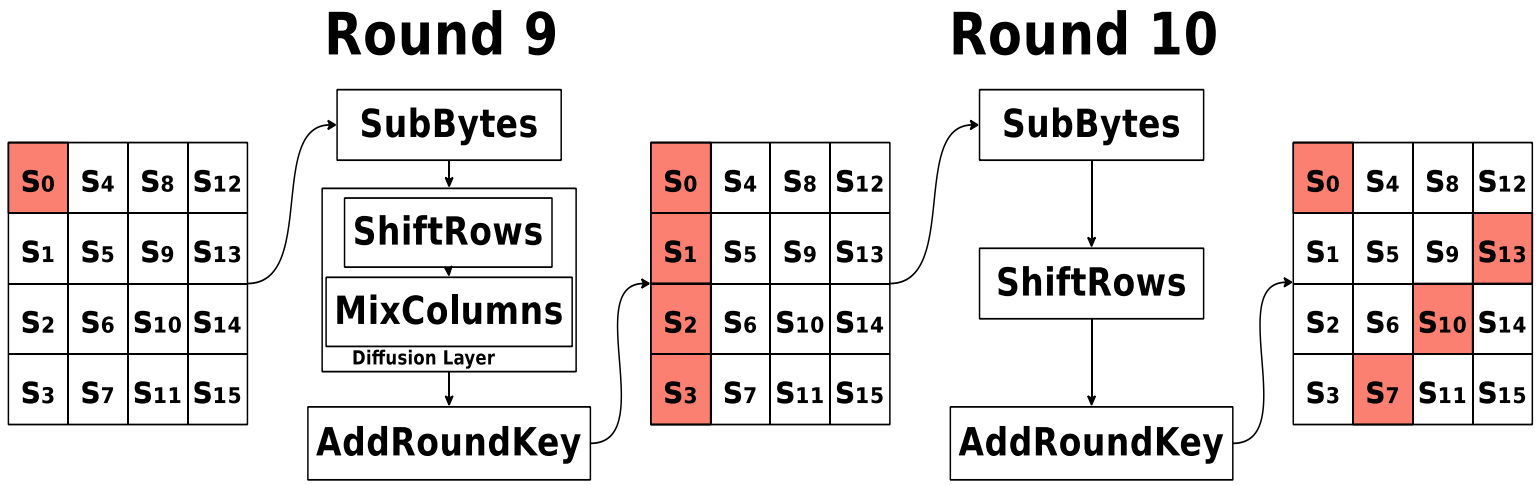

Fault Injection and Analysis

- Attacker issues encryption request to get correct ciphertext
- Attacker issues encryption requests while activating RO grid
- Fault injection is **calibrated** until desired faults appear
- Calibration is done only **once** for a specific board

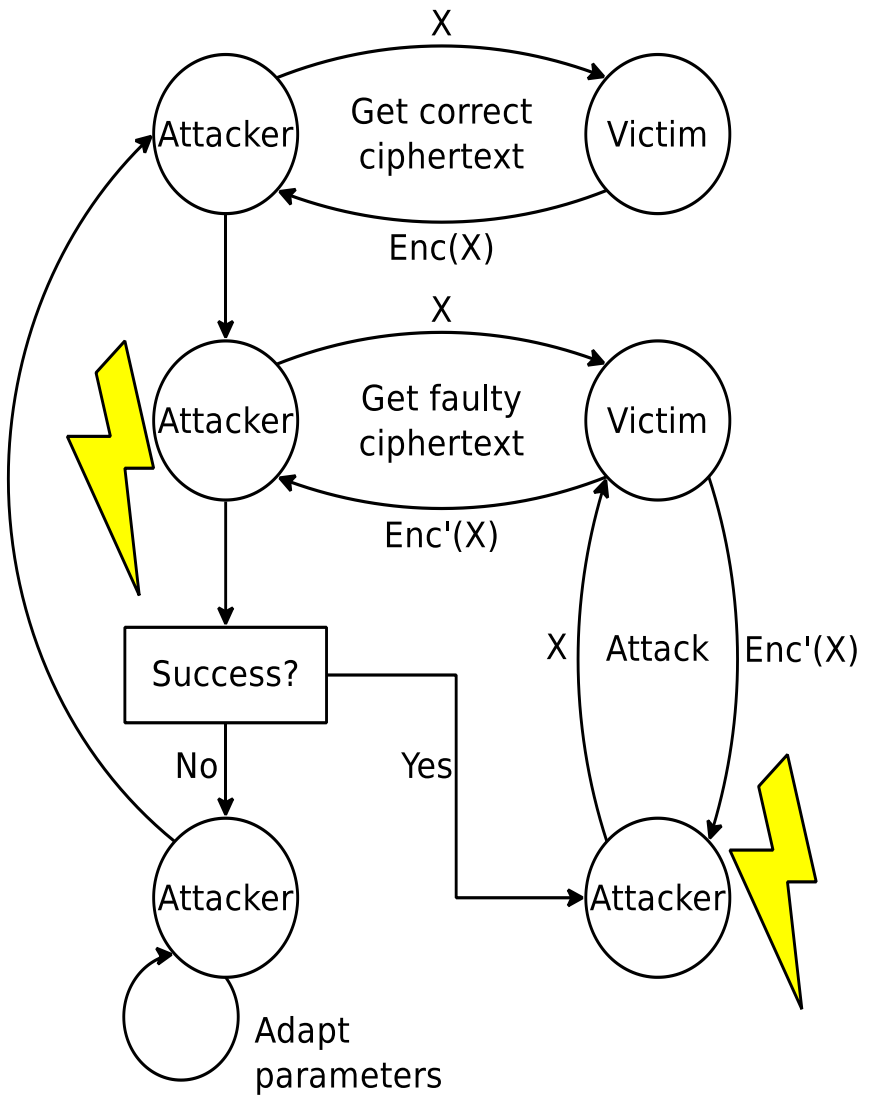

Experimental Setup(s)

- Intel DE1-SoC Board
  - Cyclone V FPGA-SoC
  - ARM Cortex A9 Dual Core
- Attacker and Victim Software on ARM
- Attacker and Victim IP Cores on FPGA
  - Design meets all timing constraints
- On Intel DE0-Nano-SoC:
  - Feasible if design does not meet worst-case timing models

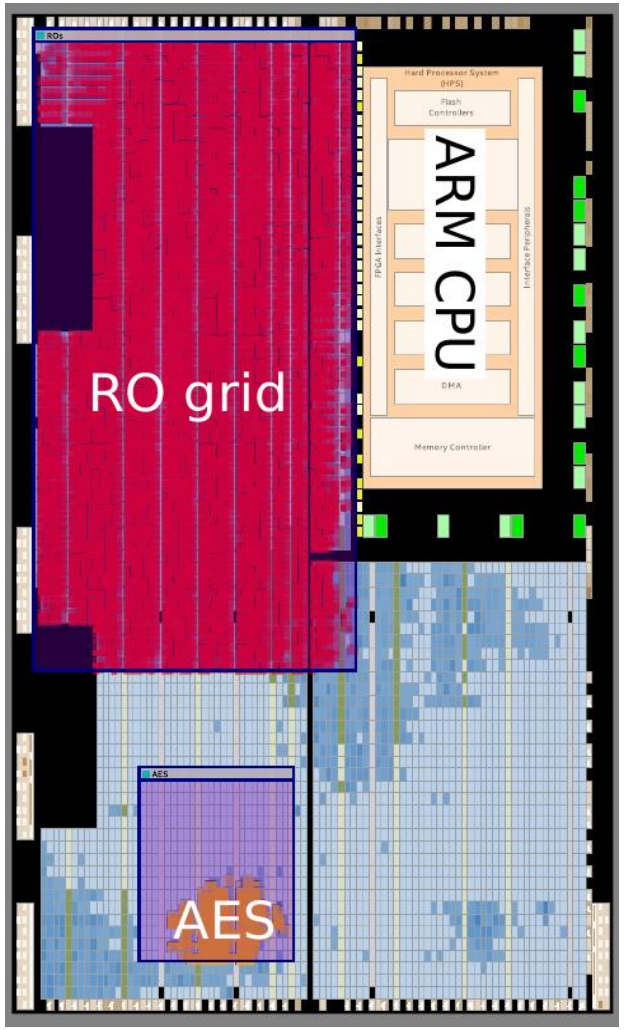

Results

- Fault injection rate
  - Tested on three boards
  - Details shown for one

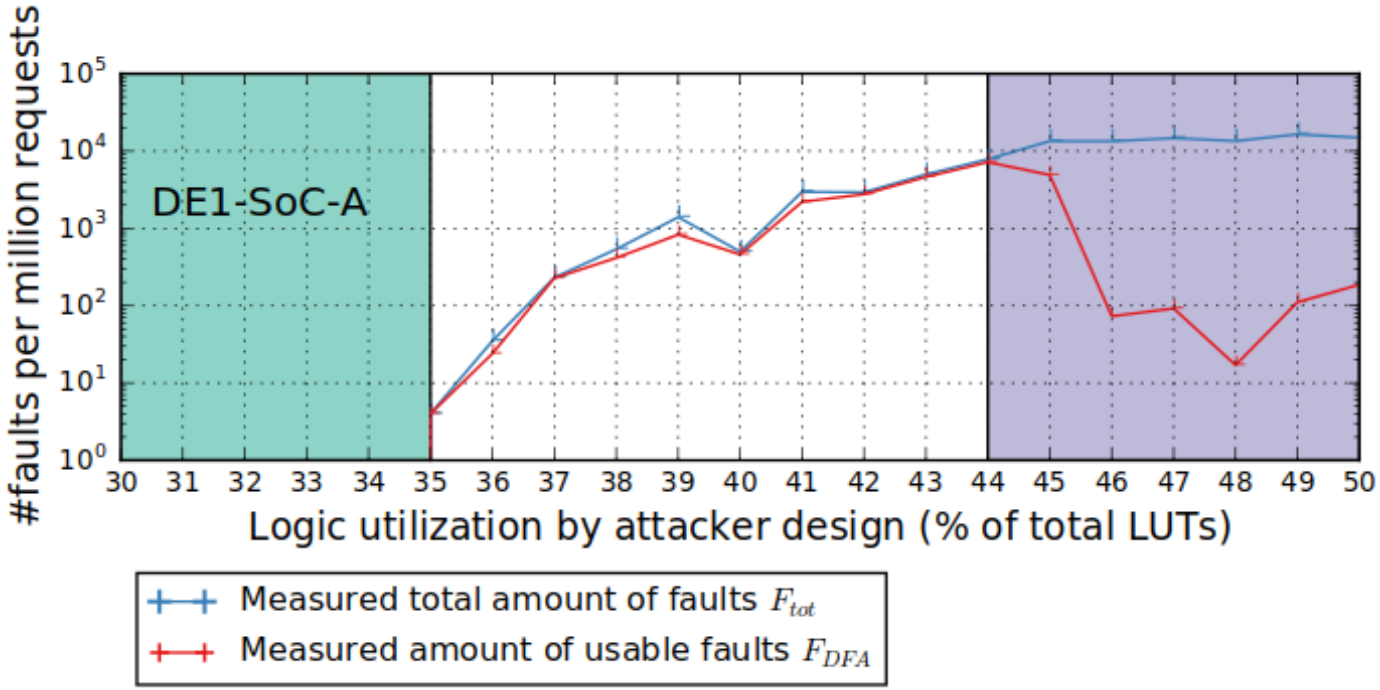

- Recovered AES keys
  - For three boards:

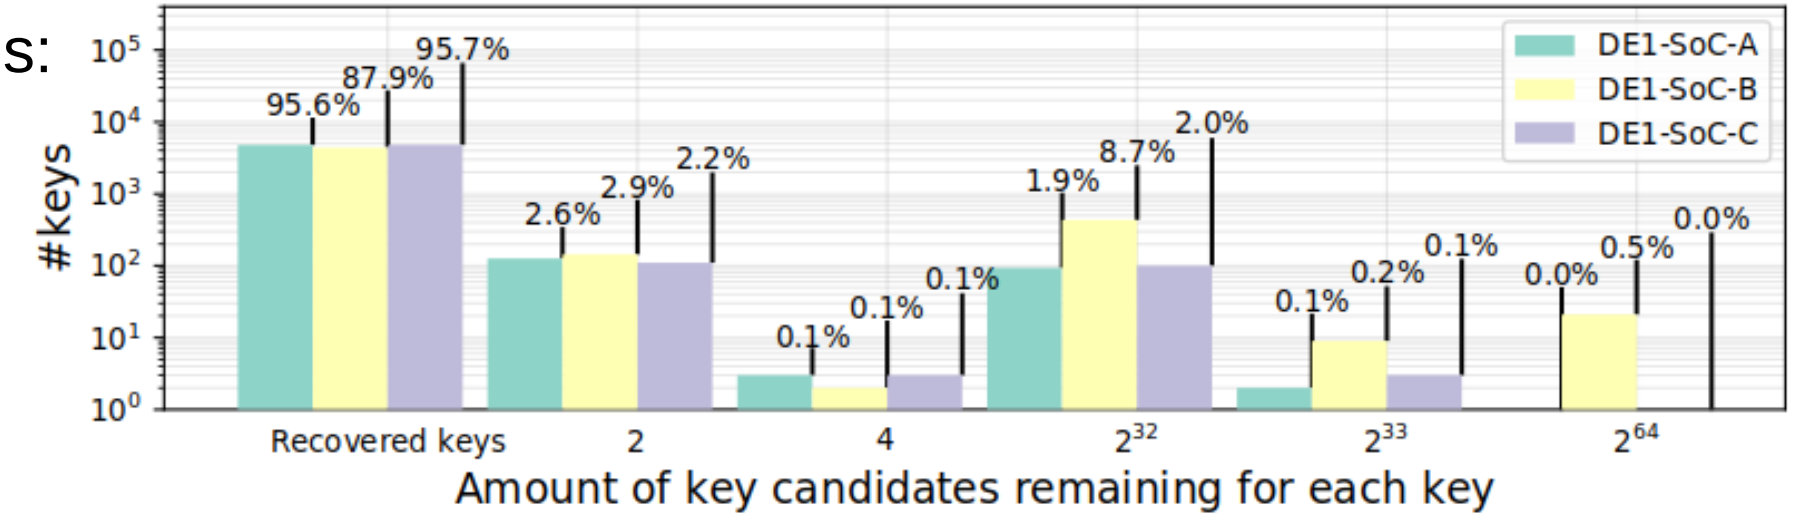

Outline

- Background & Attack Model
- Attacks
  - Summary
- Countermeasures
- Conclusion & Perspectives

Summary of Results

- Remote Power Analysis & Fault Attacks are possible
- Even when existing secure design flow is followed

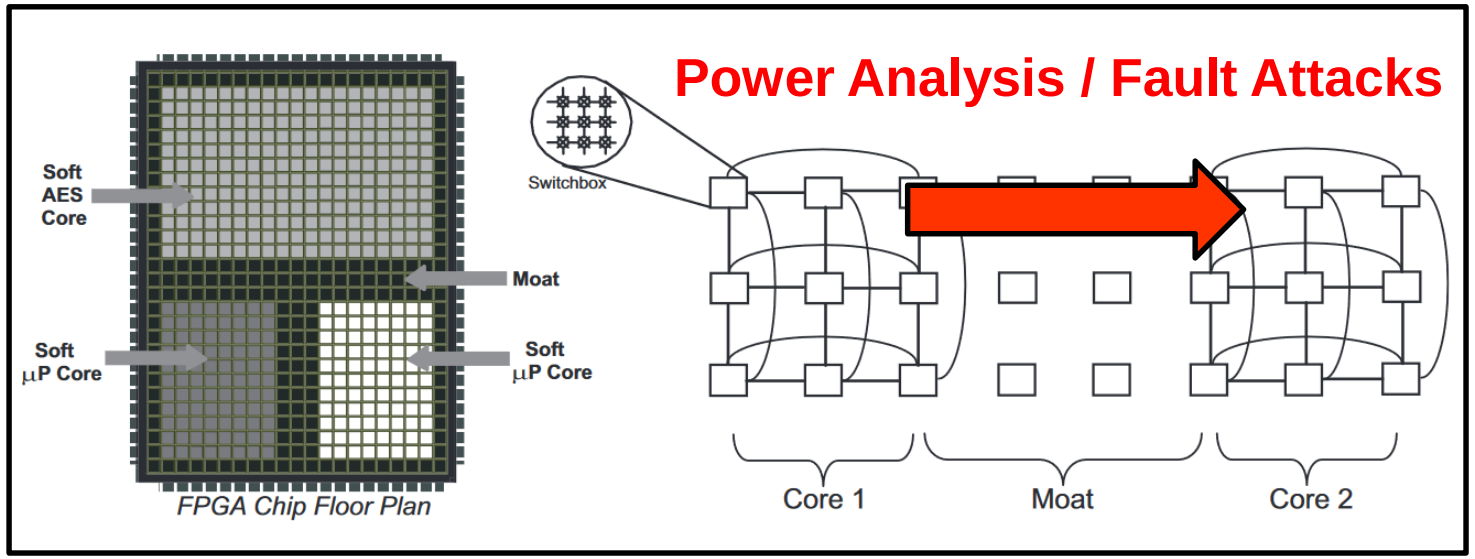

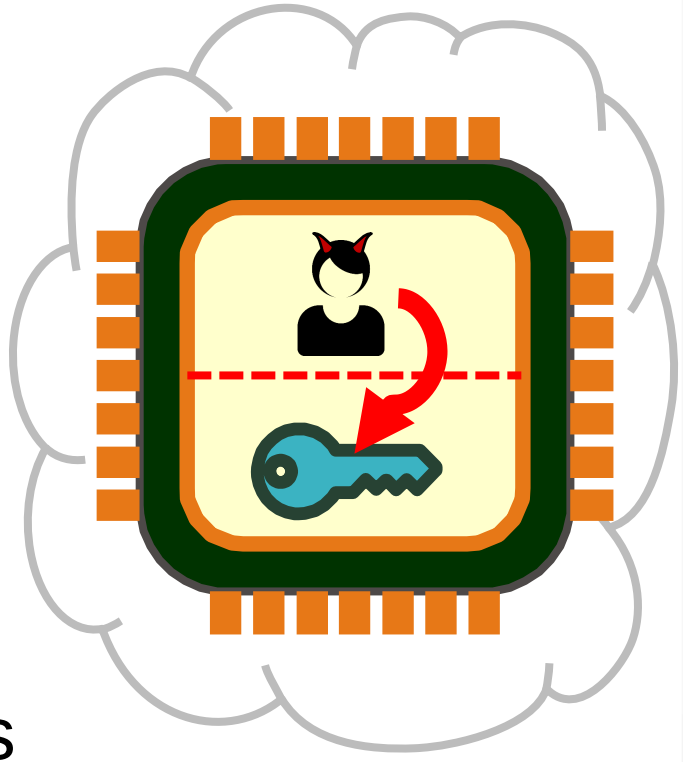

- Results extended to other FPGA boards & vendors
  - Lattice HX8K, ECP5, .. Xilinx Zynq ZC7020, ..VCU108, .. Intel Cyclone V, ..

More boards can be attacked.. (list until Dec. 2019)

Board	Attack successful?		
	Voltage Drop-based Denial of Service	Voltage Drop-based Timing Fault Injection	Key Recovery by Side-Channel
Intel Terasic DE0-Nano-SoC	–	Yes	–
Intel Terasic DE1-SoC	Yes	Yes	–
Intel Terasic DE2-115	–	–	Yes, [8][1]
Intel Terasic DE4	–	Yes	–
Lattice ECP5 5G Evaluation Board	–	–	Yes
Lattice iCE40-HX8K Breakout Board	Yes	Yes	Yes
Xilinx Artix-7 Basys-3	–	–	Yes
Xilinx Kintex-7 KC705	Yes	–[5]	–
Xilinx Pynq Zynq-ZC7020	Yes[2]	–[5]	Yes
Xilinx Spartan-6 SAKURA-G	–	–	Yes[4]
Xilinx Ultrascale VCU108	Yes	Yes	–
Xilinx Virtex-6 ML605	Yes	–[5]	–
Xilinx Virtex-7 VC707	–	Yes, [9]	–
Xilinx Zedboard Zynq-ZC7020	Yes[2]	–[5]	Yes, [5][3]

[1] Information leaks through cross-coupling of adjacent wires. This is less of a threat, since coupling is prevented if interconnect matrices are not shared between multiple designs [8], which is also recommended by standard secure design flow practices [2].
[2] It affects the whole SoC including the integrated ARM Cortex-A9 Dual-Core.
[3] Sufficient leakage for key recovery was also shown from CPU to FPGA in the same SoC
[4] In [10] it was additionally shown to work from one FPGA in the system (connected to the same power supply) to another, on board-level.
[5] A simple experiment was conducted, but the devices crashed before timing violations occured – it might still be possible with more effort.

Outline

- Background & Attack Model
- Attacks
- Countermeasures
- Conclusion & Perspectives

Countermeasures tailored to FPGAs

- Offline Method: Bitstream checking
  - FPGA “Anti-Virus”
  - Detects malicious FPGA configuration before it gets loaded to the FPGA:
    - Structures usable for Power Analysis
    - Structures usable for Fault Attacks
- Runtime Method: “Active Fences”
  - Hiding scheme tailored to FPGA on-chip power analysis
  - Reduces signal-to-noise ratio with on-chip spatial considerations in mind
- Runtime Method: “LoopBreaker”
  - Prevents faults by reconfiguring interconnect as fast as possible

Offline Approach: Bitstream Checking

- Checking designs (bitstream) before loading to FPGA
- Supervisor checks bitstreams for malicious “Signatures”
  - May require bitstream reverse engineering
- Fault Attack Signature: causing too much voltage droop
  - Too many combinational cycles with inputs
  - Or: Any high fanout-nets
- Power Analysis Signature: possible delay sensors
  - Has timing violations according to timing analysis
  - Has at least one combinational cycle with output

Bitstream Checking Flow

- Flow of signature checking:

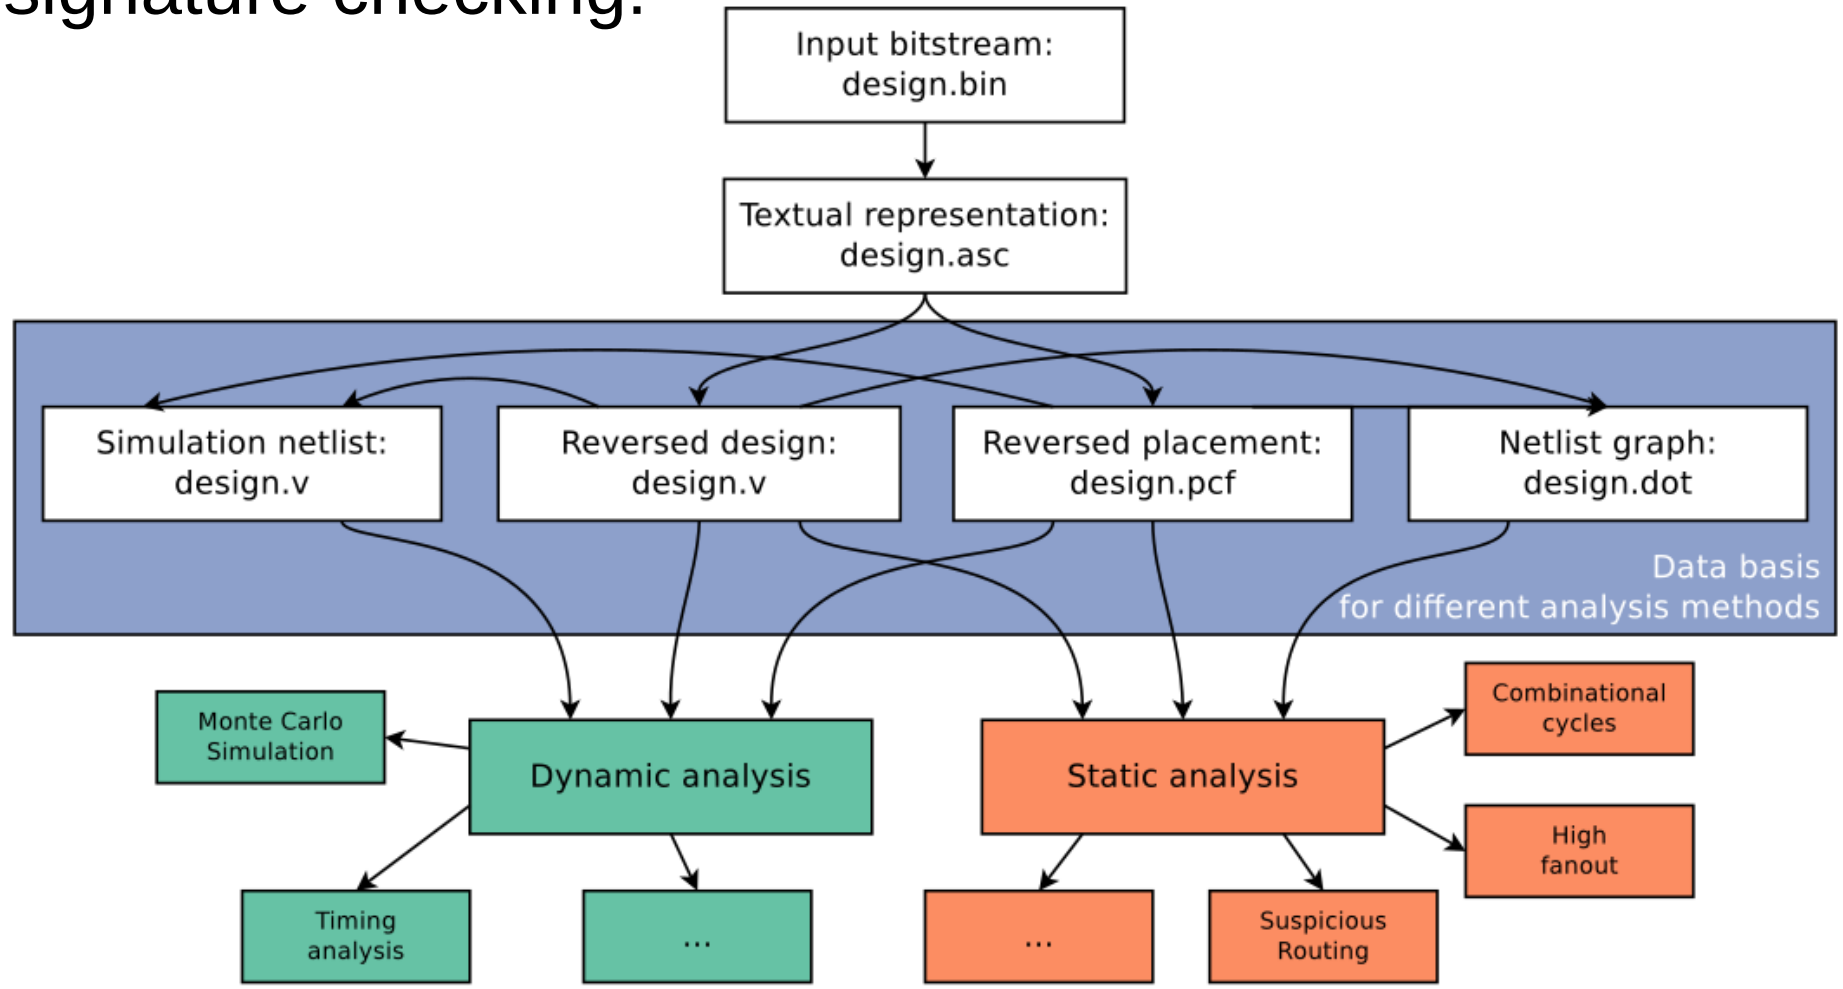

J. Krautter, D. R. E Gnad, and M. B. Tahoori: Mitigating electrical-level attacks towards secure multi-tenant FPGAs in the cloud, TRETS 2019

Results Bitstream Checking 1/2

- Evaluation on known attack designs
  - Reference01 – Delay Line Sensor
  - Reference02 – Ring Oscillator Sensor
  - Reference03 – Fault Injection with Ring Oscillators

Design Name	#LEs	Comb. Cycles?	Data-to-Clock?	Highest Fanout?	Timing violations?	Runtime (structural)	Runtime (timing)
reference01	6075	✓	×	3000	✓	29.32*s*	152.31*s*
reference02	6810	✓	×	6500	×	20.32*s*	169.67*s*
reference03	4077	×	×	3000	✓	25.09*s*	127.05*s*

J. Krautter, D. R. E Gnad, and M. B. Tahoori: Mitigating electrical-level attacks towards secure multi-tenant FPGAs in the cloud, TRETS 2019

Results Bitstream Checking 2/2

- Evaluation on benign designs
- No false positives on these

Design Name	#LEs	Comb. Cycles?	Data-to-Clock?	Highest Fanout?	Timing violations?	Runtime (structural)	Runtime (timing)
s27	19	×	×	3	×	4.54*s*	17.34*s*
s208_1	73	×	×	8	×	4.60*s*	17.66*s*
s298	141	×	×	14	×	4.63*s*	18.33*s*
s344	132	×	×	18	×	6.21*s*	18.56*s*
s349	123	×	×	18	×	6.24*s*	18.51*s*
s382	159	×	×	22	×	6.20*s*	19.11*s*
s386	167	×	×	13	×	4.57*s*	19.33*s*
s400	163	×	×	22	×	6.16*s*	18.85*s*
s420_1	187	×	×	16	×	4.59*s*	19.21*s*
s444	162	×	×	22	×	6.19*s*	19.61*s*
s510	256	×	×	26	×	4.76*s*	20.52*s*
s526	237	×	×	22	×	6.22*s*	18.51*s*
s526n	228	×	×	22	×	6.15*s*	18.27*s*
				...			
b08	227	×	×	21	×	6.35*s*	18.74*s*
b09	260	×	×	32	×	6.27*s*	18.35*s*
b10	311	×	×	24	×	6.42*s*	18.64*s*
b11	805	×	×	35	×	6.51*s*	20.02*s*
b12	2017	×	×	119	×	7.13*s*	24.52*s*
b13	402	×	×	50	×	6.36*s*	19.88*s*
sha	1833	×	×	969	×	15.29*s*	69.13*s*
diffeq2	4049	×	×	97	×	19.78*s*	111.06*s*
vexriscv	2587	×	×	241	×	23.86*s*	87.49*s*
leon3	29304	×	×	1838	–	19.28*s*	–

J. Krautter, D. R. E Gnad, and M. B. Tahoori: Mitigating electrical-level attacks towards secure multi-tenant FPGAs in the cloud, TRETS 2019

Active Fences

- Generic to apply and not affect tenant regions → “Fence”
- Dynamic hiding that adapts at runtime → Sense Leakage
- Non-intrusive to algorithm or implementation → Generic ROs

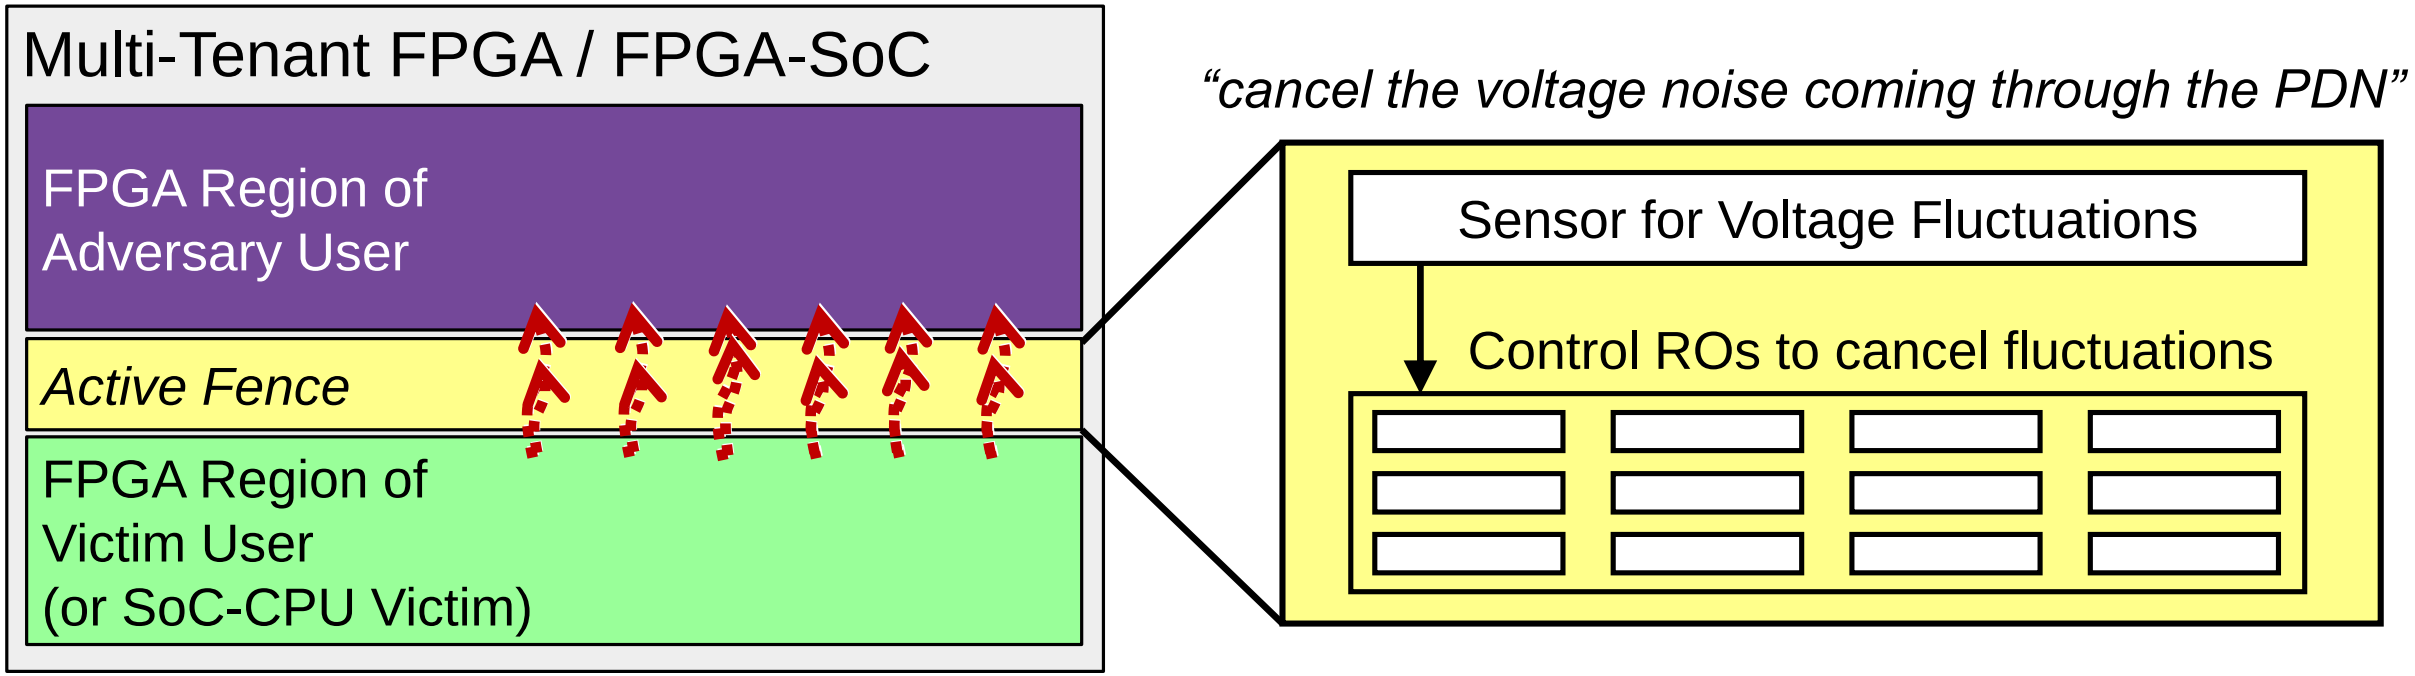

J. Krautter, D. Gnad, F. Schellenberg, A. Moradi, M. Tahoori: Active Fences against Voltage-based Side Channels in Multi-Tenant FPGAs, ICCAD 2019

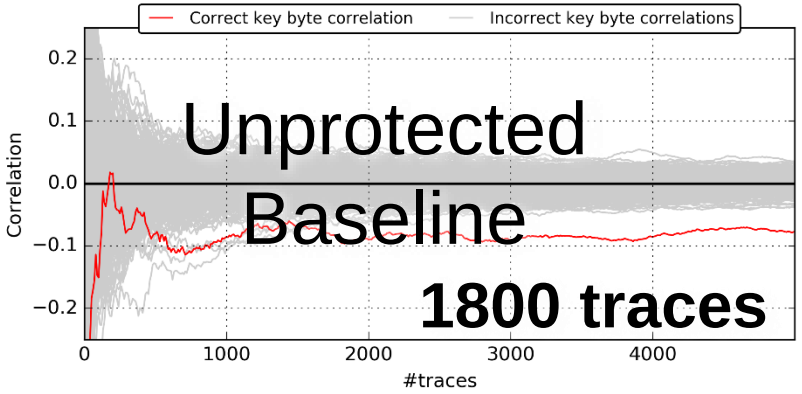

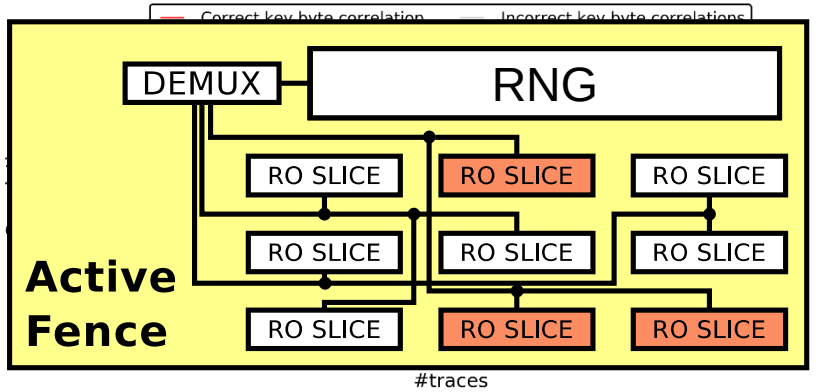

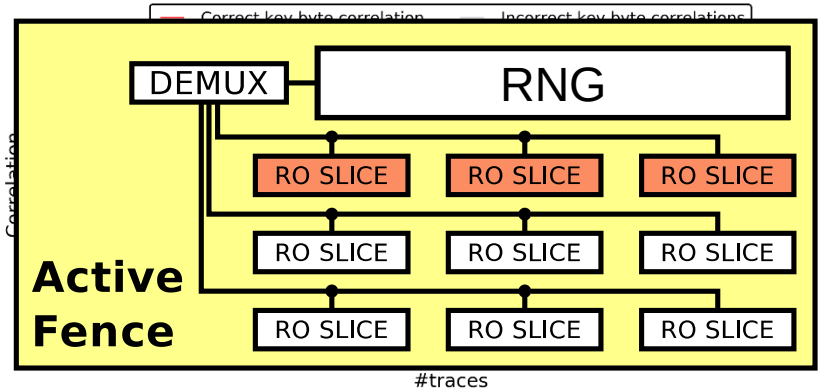

1800 → 80,000→ 120,000 → 300,000 = 166x more traces needed!

RO-array,
ordered activation

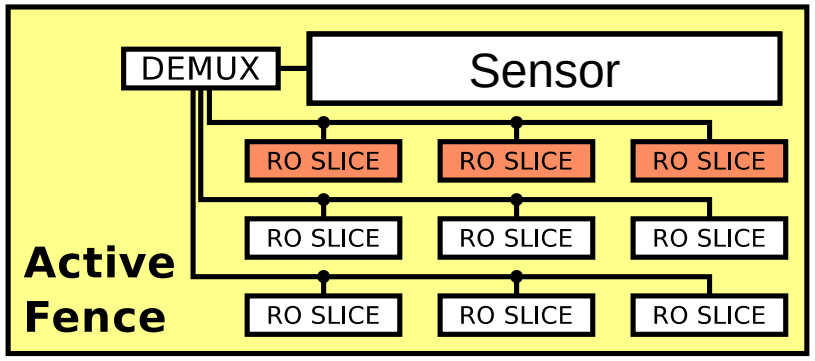

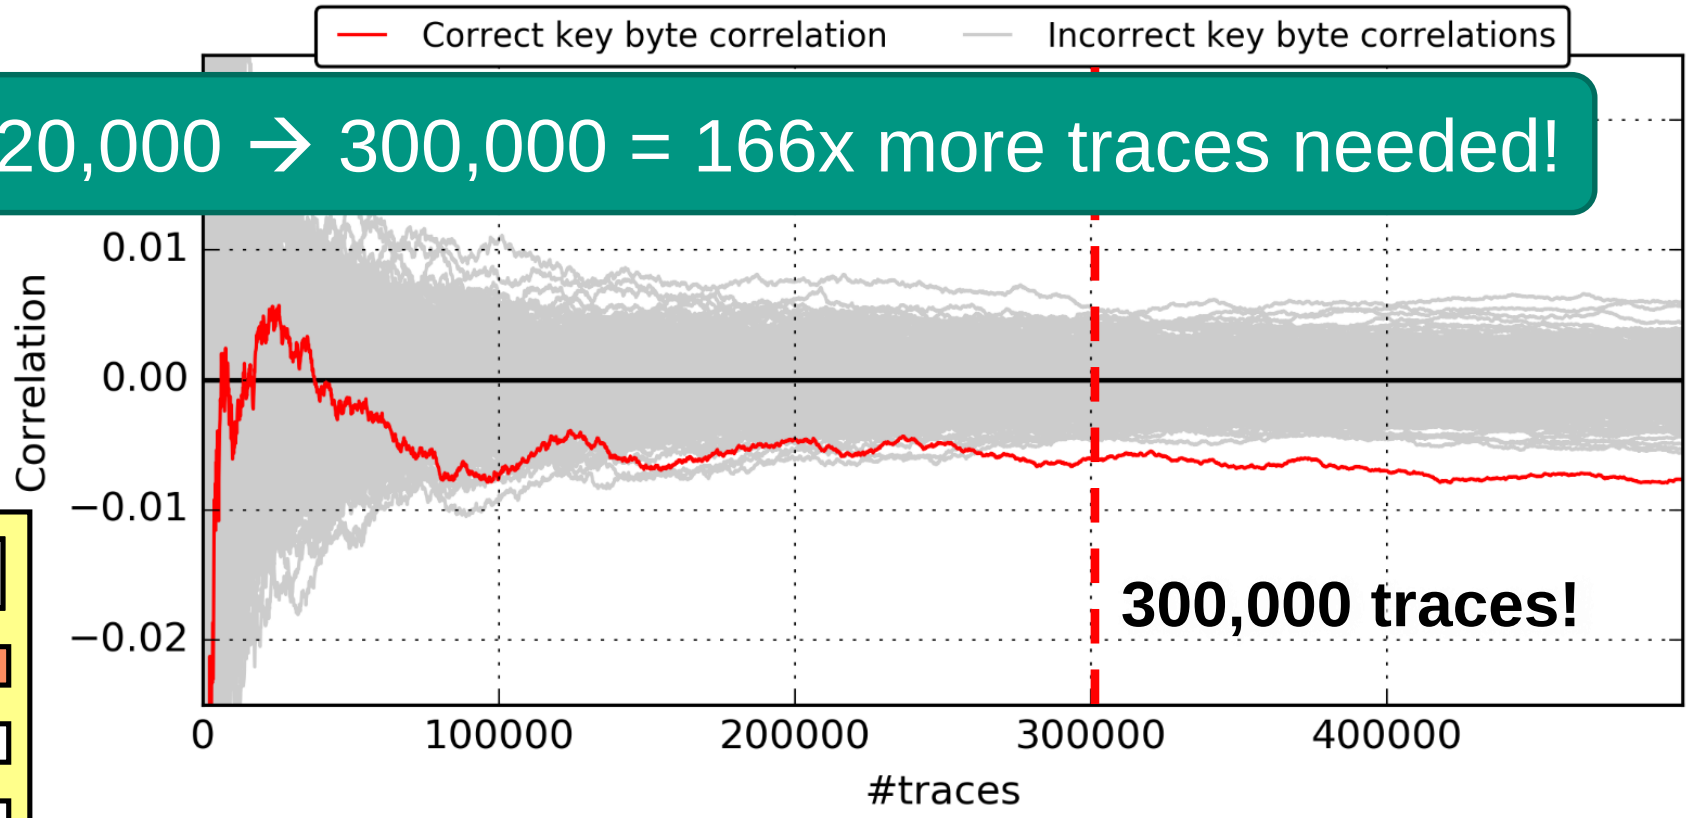

J. Krautter, D. Gnad, F. Schellenberg, A. Moradi, M. Tahoori: Active Fences against Voltage-based Side Channels in Multi-Tenant FPGAs, ICCAD 2019

LoopBreaker 1/2

- Malicious tenants can be detected with voltage drop sensors
- To prevent them, their activity needs to be stopped, e.g.:

1. Stopping the Clock → Not enough! Self-clocking is possible
2. Partial reconfiguration → Remove an entire region → Too slow

- LoopBreaker only reconfigures what it needs to…

H. Nassar, H. AlZughbi, D. Gnad, L. Bauer, M. Tahoori, J. Henkel: LoopBreaker: Disabling Inter-connects to Mitigate Voltage-Based Attacks in Multi-Tenant FPGAs, ICCAD 2021

LoopBreaker 2/2

- LoopBreaker concentrates on **interconnect**
  - ➢ Stops ALL activity
  - ➢ 100x faster reconfig.

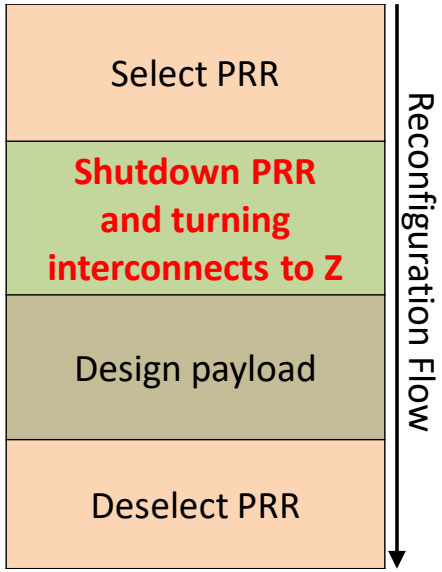

Bitstream Structure of UltraScale+ FPGAs

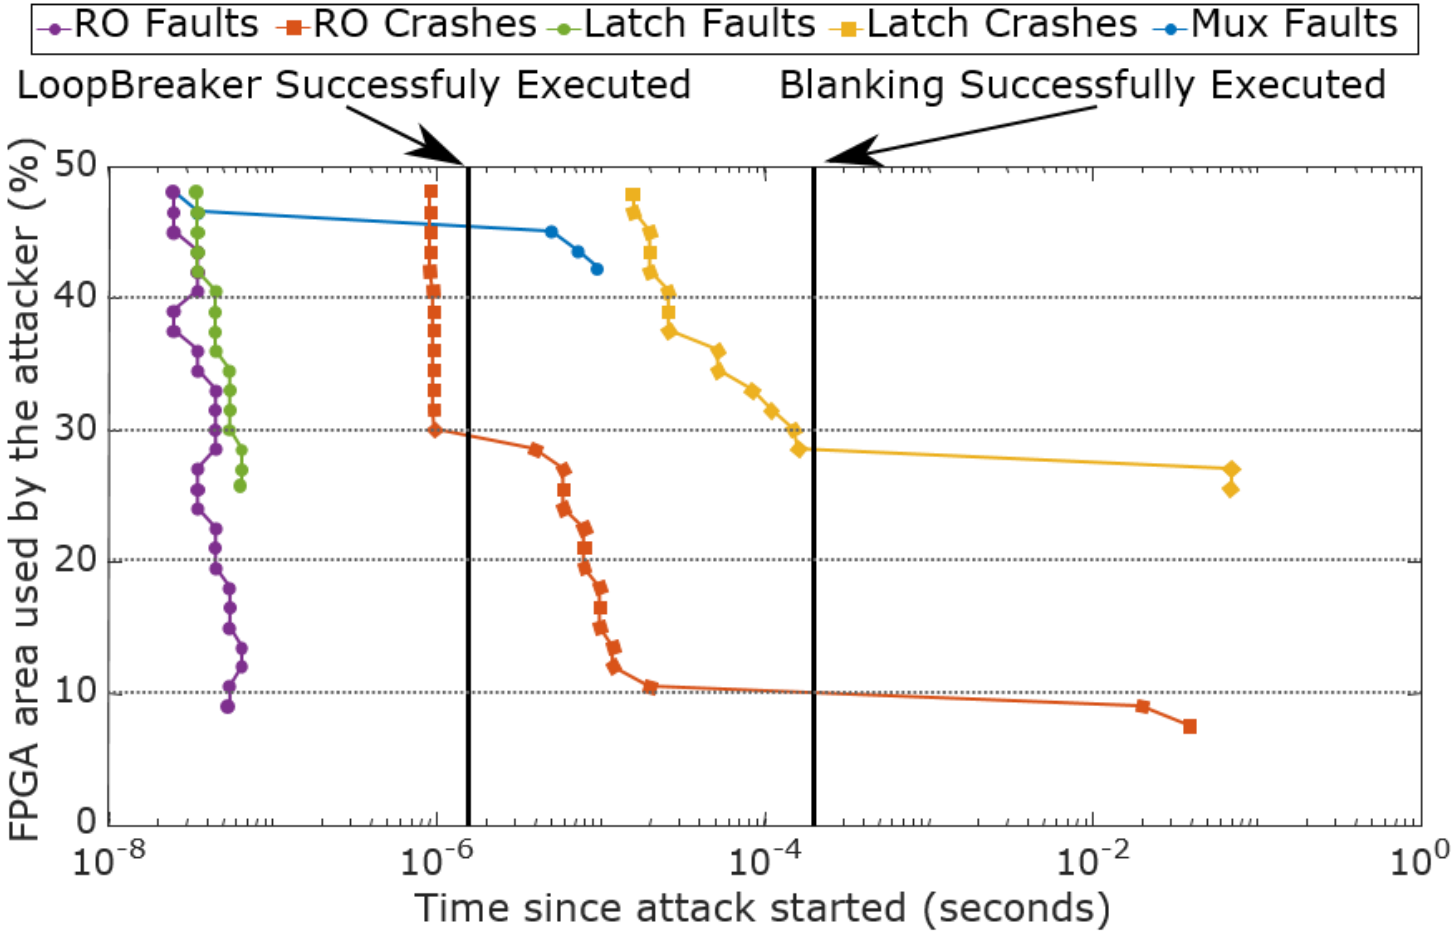

H. Nassar, H. AlZughbi, D. Gnad, L. Bauer, M. Tahoori, J. Henkel: LoopBreaker: Disabling Inter-connects to Mitigate Voltage-Based Attacks in Multi-Tenant FPGAs, ICCAD 2021

Summary Countermeasures

- Offline approach: Bitstream checking
  - Detects malicious configurations on supervisor level
  - Potential remaining issue:
    - Legal corner case in operation that produces high voltage drop → fault
    - Timing violations need to be entirely defined on supervisor level
- Runtime method: “Active Fences”
  - Increase required traces for CPA by 166x, with 1.5/2x Power/Area overhead
  - Non-invasive to user design, works as a wrapper, generic to apply
- Runtime “LoopBreaker”
  - Disables interconnects faster, but might not be fast enough for all attacks

H. Nassar, H. AlZughbi, D. Gnad, L. Bauer, M. Tahoori, J. Henkel: LoopBreaker: Disabling Inter-connects to Mitigate Voltage-Based Attacks in Multi-Tenant FPGAs, ICCAD 2021

Outline

- Background & Attack Model
- Attacks
- Countermeasures
- Conclusion & Perspectives

Escalating On-Chip Power Analysis …

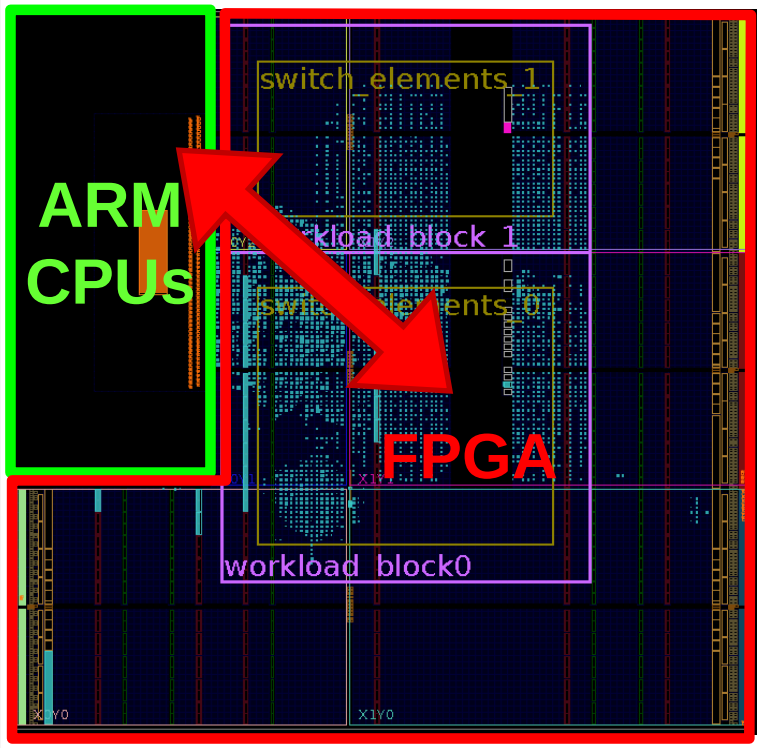

FPGA-SoC

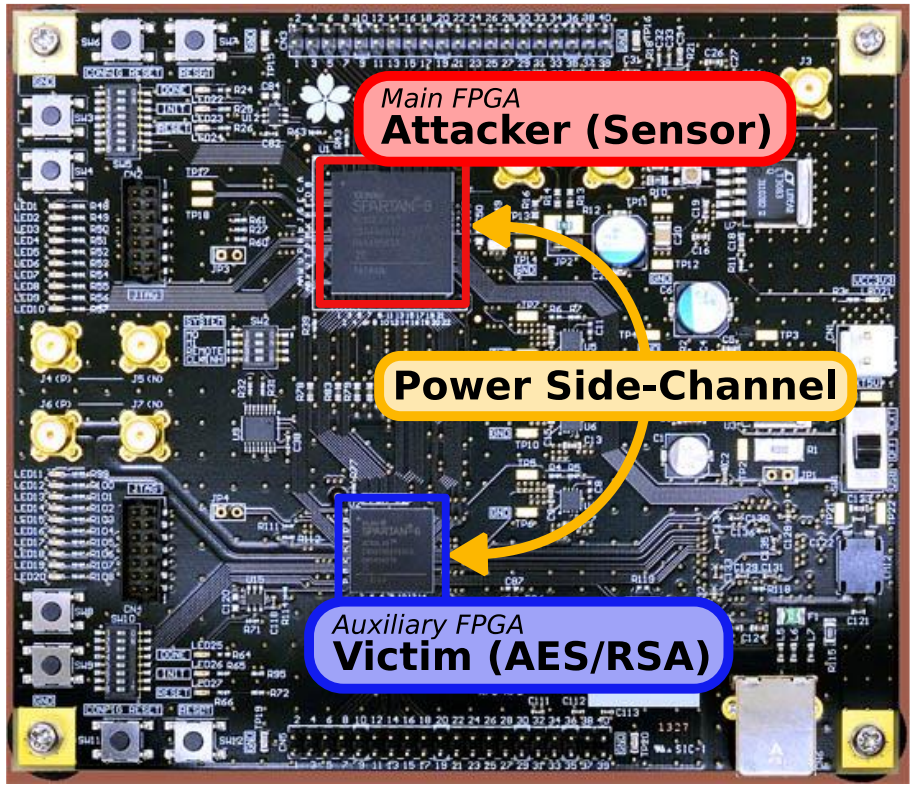

Printed Circuit Board
(Shared Power Supply)

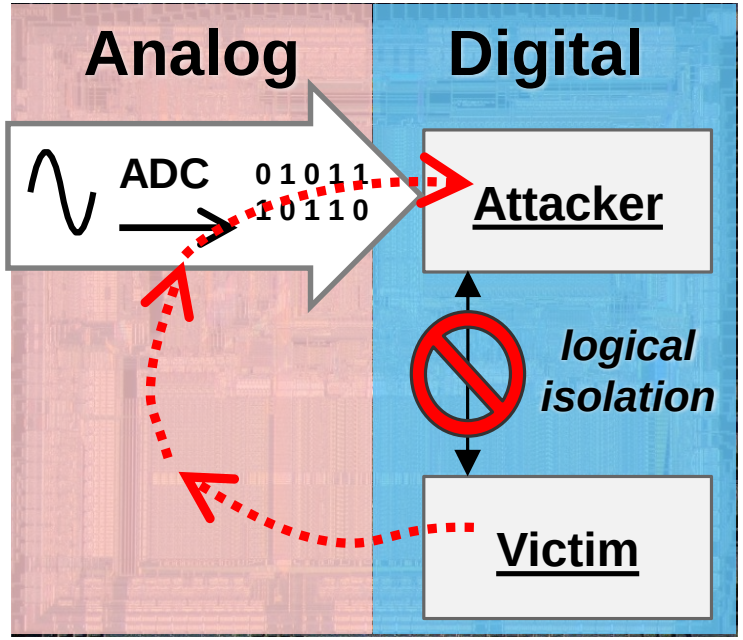

Mixed-Signal SoC

Recent Related Work

- Many have appeared in the last years, a few selected ones are:
- Bitstream checkers
  - Including analysis of various fault injection methods [La et al. TRETS'20]
- Attacks in Cloud Platforms
  - Proof that the respective malicious bitstreams can actually be loaded
  - SCA in FPGA in Amazon EC2, [Glamocanin et al. DATE 2020]
  - DoS in FPGA on Amazon EC2, [La et al. CHES 2021]
  - ➢ Amazon implemented some, but not sufficient, countermeasures

Hardware Oriented Security and Trust (HOST) 2019

- Low-Cost Lattice HX8K Board
  - Power Analysis SCA
  - Faults for DoS
  - Faults for Differential Fault Analysis
- Interactive GUI
- Open-source, available:
  - cdnc.itec.kit.edu/SCA-DFA-Demo.php
- Allows hardware security education, without expensive equipment!
- ➢ Base for this Tutorial!

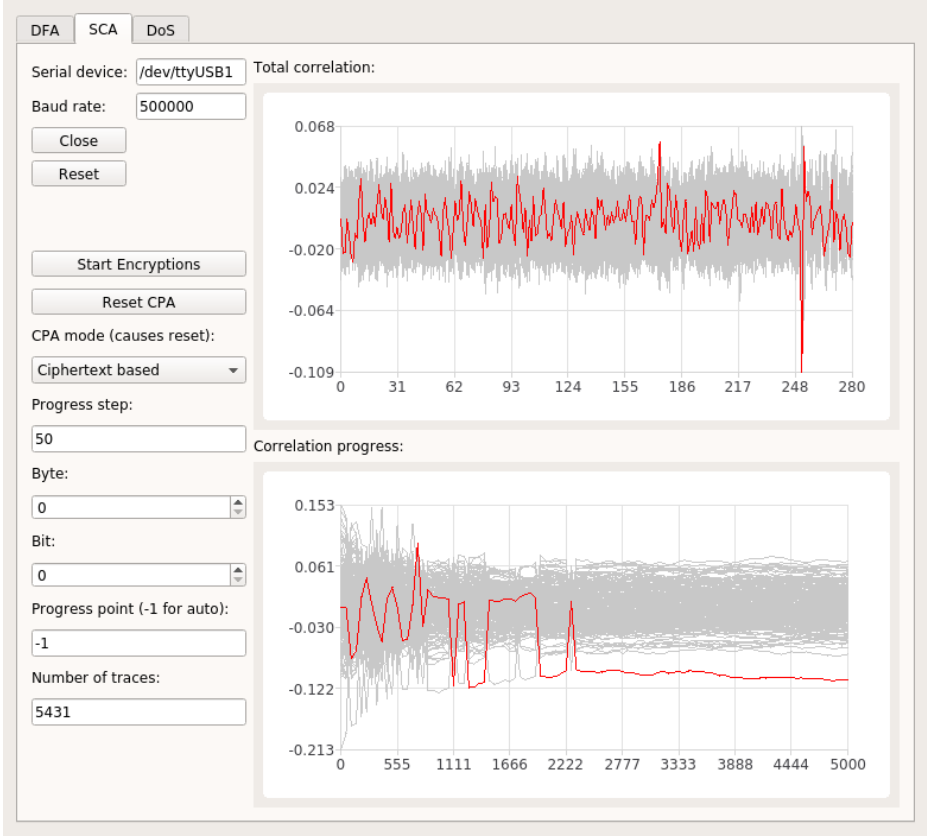

Conclusion

First successful remote power analysis attack

First remote fault attack (for DoS)

First countermeasures for those attacks

Motivates further research:

- General security aspects of Power Distribution Networks (PDNs)
- Analyzing PDN-based attacks in other devices
- New security concepts for Multi-Tenant FPGAs

Thanks for your Attention!

Questions?

Our Publications until 2019, On-Chip Voltage Fluctuations

- J. Krautter, D. R. E. Gnad, F. Schellenberg, A. Moradi, and M. B. Tahoori, "Active Fences against Voltage-based Side Channels in Multi-Tenant FPGAs", in International Conference on Computer-Aided Design (ICCAD), 2019, USA.
- J. Krautter, D. R. E. Gnad, M. B. Tahoori, "Mitigating Electrical-Level Attacks towards Secure Multi-Tenant FPGAs in the Cloud", in ACM Transactions on Reconfigurable Technology and Systems (TRETS), 2019.
- D. R. E. Gnad, J. Krautter, M. B. Tahoori, "Leaky Noise: New Side-Channel Attack Vectors in Mixed-Signal IoT Devices", in IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES), 2019.
- D. R. E. Gnad, S. Rapp, J. Krautter, M. B. Tahoori, "Checking for Electrical Level Security Threats in Bitstreams for Multi-Tenant FPGAs", International Conference on Field-Programmable Technology (FPT), 2018, Japan.
- F. Schellenberg, D. R. E. Gnad, A. Moradi, M. B. Tahoori, "Remote Inter-Chip Power Analysis Side-Channel Attacks at Board-Level", in International Conference on Computer-Aided Design (ICCAD), 2018, USA.
- J. Krautter, D. R. E. Gnad, M. B. Tahoori, "FPGAhammer: Remote Voltage Fault Attacks on Shared FPGAs, suitable for DFA on AES", in IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES), 2018. (**CSAW'18 Finalist**)
- D. R. E. Gnad, F. Oboril, S. Kiamehr, M. B. Tahoori, "An Experimental Evaluation and Analysis of Transient Voltage Fluctuations in FPGAs", in IEEE Transactions on Very Large Scale Integration Systems (TVLSI), 2018.
- F. Schellenberg, D. R. E. Gnad, A. Moradi, M. B. Tahoori, "An Inside Job: Remote Power Analysis Attacks on FPGAs", in proceedings of Design, Automation & Test in Europe (DATE), 2018, Germany. (**Best Paper Candidate**)
- D. R. E. Gnad, F. Oboril, M. B. Tahoori, "Voltage Drop-based Fault Attacks on FPGAs using Valid Bitstreams", International Conference on Field-Programmable Logic and Applications (FPL), 2017, Belgium. (**Best Paper Award**)
- D. R. E. Gnad, F. Oboril, S. Kiamehr, M. B. Tahoori, "Analysis of Transient Voltage Fluctuations in FPGAs", International Conference on Field-Programmable Technology (FPT), 2016, China. (**Best Paper Candidate**)

Demos / Media / Presentations

- D. Gnad, S. Ritterbusch, "FPGA Seitenkanäle", Gespräch im Modellansatz Podcast, Folge 177, Fakultät für Mathematik, Karlsruher Institut für Technologie (KIT), 2018.
- D. Gnad, "Seitenkanal-Angriffe innerhalb FPGA-Chips", Vortrag auf der GPN18, media.ccc.de, Chaos Computer Club e.V, May 2018, Karlsruhe, Germany. < https://media.ccc.de/c/gpn18 >
- D. Gnad, M. Tahoori, "Security threats in nanoscale FPGA fabric", Workshop on SecURity, REliAbiLity, test, prIvacy, Safety and Trust of Future Devices (SURREALIST), May/June 2018, Bremen, Germany.
- J. Krautter, D. R. E. Gnad, F. Schellenberg, A. Moradi, M. B. Tahoori, "Software-based Fault and Power Side-Channel Attacks inside Multi-Tenant FPGAs", Demo Session, IEEE International Symposium on Hardware Oriented Security and Trust (HOST), 2019, USA. (**Best Hardware Demo Award, Third Place**)