Area-efficient fault-tolerant architectures exploiting masking scheme randomness

Vojtěch Miškovský, Hana Kubátová, Martin Novotný

Czech Technical University in Prague
Faculty of Information Technology
Department of Digital Design

CryptArchi 2019
June 24–25, Průhonice

Motivation

- Dependability and security can be demanded at the same time
- Fault-tolerant design: **high** overhead
- Attack-resistant design: **high** overhead
- Fault-tolerant and attack resistant design: **high**^{2} overhead?
  - Hopefully not

Fault-tolerant architectures

- Modular redundancy
  - Duplex (a)
    - 1 fault detection
  - TMR (b)
    - 1 fault overriding
  - NMR
    - $(N-1)/2$ faults overriding

👍 Simple to implement

👎 High area overhead

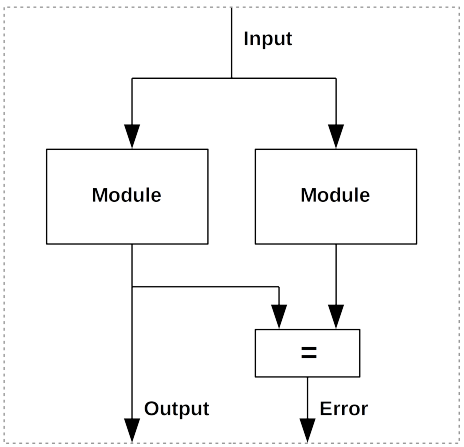

(a) duplex

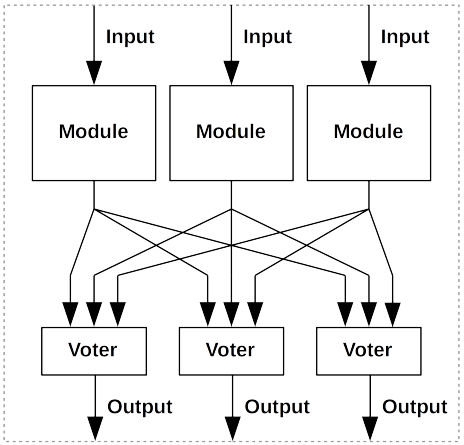

(b) TMR

Attack countermeasures

- Our approach is based on masking scheme randomness with focus on glitch-protected schemes, e.g.
  - Threshold Implementation[1]
  - Domain-Oriented Masking[2]

👍 Provably secure against Side-Channel Analysis of arbitrary order and some fault attacks

👎 High area overhead

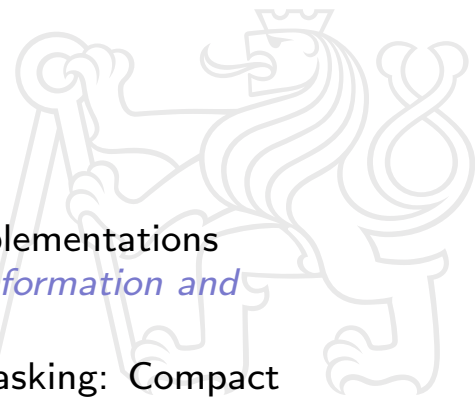

[1] Svetla Nikova, Christian Rechberger, and Vincent Rijmen. "Threshold implementations against side-channel attacks and glitches". In: *International Conference on Information and Communications Security*. Springer. 2006, pp. 529–545.

[2] Hannes Groß, Stefan Mangard, and Thomas Korak. "Domain-Oriented Masking: Compact Masked Hardware Implementations with Arbitrary Protection Order.". In: *TIS@ CCS*. 2016, p. 3.

Our contribution

- Proposal of architectures with
  - 👍 similar dependability and security properties as modular redundancy + masking scheme, but
  - 👍 significantly decreased overhead

Methodology

- Assumptions:
  - fault leads to different faulty output for same but differently masked inputs
  - round based symmetric cipher
- Fault in:
  - input logic
    $\rightarrow$ different output for different mask
  - encryption logic
    $\rightarrow$ different output for different mask
  - output logic
    $\rightarrow$ same or different output for different mask (depending on mask)

$\rightarrow$ **Repeating encryption with different masks detects faults**

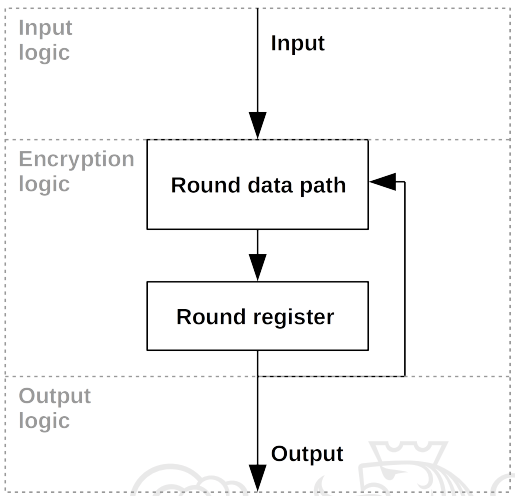

Duplex-equivalent

Principle:

- Encryption is repeated twice, with different masks
- Unmasked outputs (ciphertexts) of both iterations are compared

Properties:

- 1 fault detection
- 👍 **1 module instead of 2**
- 👎 Extra output logic and register
- 👎 Double encryption time

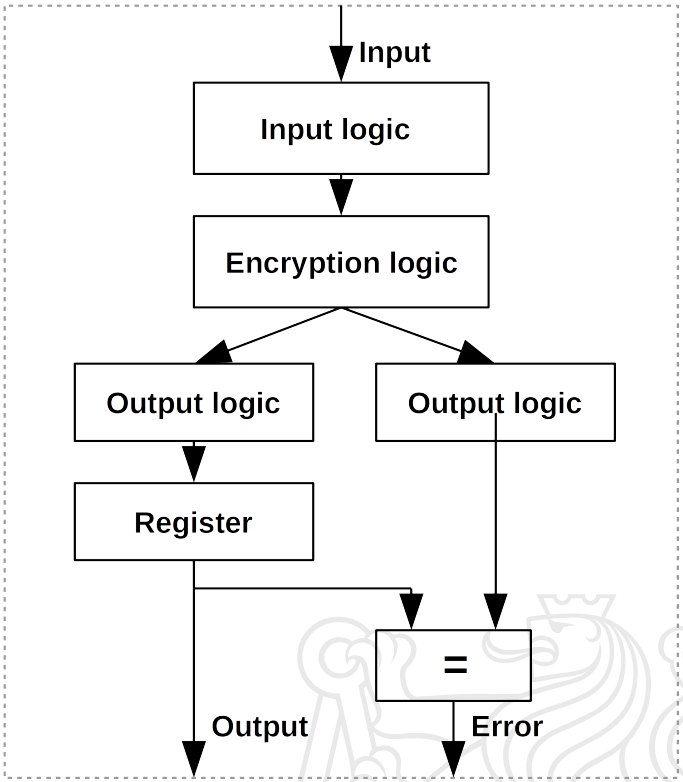

TMR-equivalent

Principle:

- 2 modules encrypt the same data using the same mask
- If the outputs differ, encryption is repeated (see the next slide)

Properties:

- 1 fault overriding
- 👍 **2 modules instead of 3**
- 👍 Detection of 1 (different) fault in both modules
- 👍 Standard encryption time when no fault occurs
- 👎 More complex comparison logic
- 👎 At least double encryption time when a fault occurs

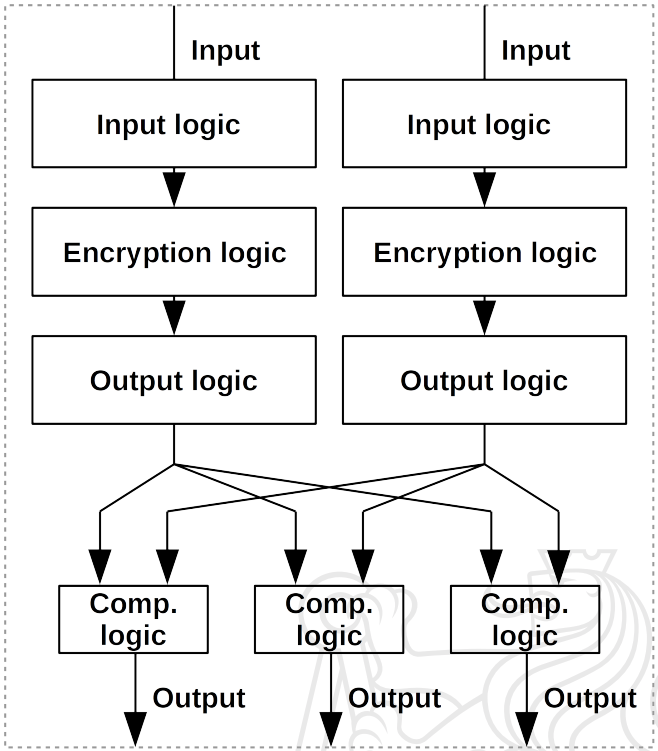

TMR-equivalent – Comparison logic

- **Unmasked** outputs are compared
- When outputs are different, encryption is repeated with different masks
- Consecutive outputs are compared for both modules
- The module whose consecutive outputs differ is considered faulty

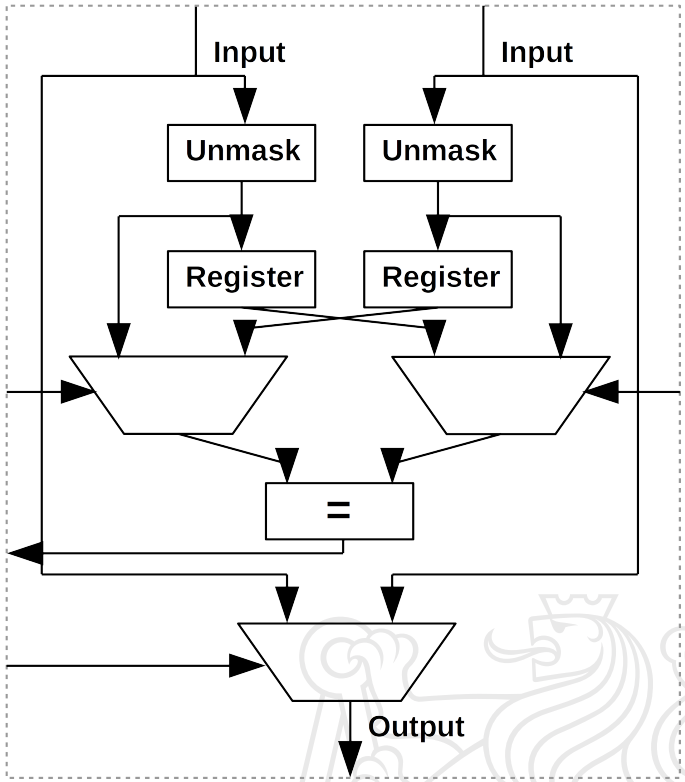

NMR-equivalent

Principle:

- All modules encrypt the same data using the same mask
- If the outputs differ, encryption is repeated

Properties:

- $(N-1)/2$ fault overriding
- 👍 $\lceil N/2 \rceil$ modules instead of N
- 👍 Detection of 1 (different) fault in all modules
- 👍 Standard encryption time when no fault occurs
- 👎 More complex comparison logic
- 👎 At least double encryption time when a fault occurs

Case study

- 3-share TI of PRESENT cipher[3]
- TMR vs our TMR-equivalent architecture
- FPGA implementations – Xilinx Spartan-6 on Sakura-G board

[3] Axel Poschmann et al. "Side-channel resistant crypto for less than 2,300 GE". In: *Journal of Cryptology* 24.2 (2011), pp. 322–345.

Results – overhead evaluation

Comparison of slice utilization for each architecture:

Design	Slice utilization	Overhead
Single module	2199	0%
TMR	7180	227%
TMR-equivalent	5764	162%

Our architecture saves around **20%** of resources in comparison with traditional TMR.

Results – TVLA

Each architecture is evaluated by non-specific, fixed-vs-random, first-order Welsh's t-test using 1,000,000 power traces

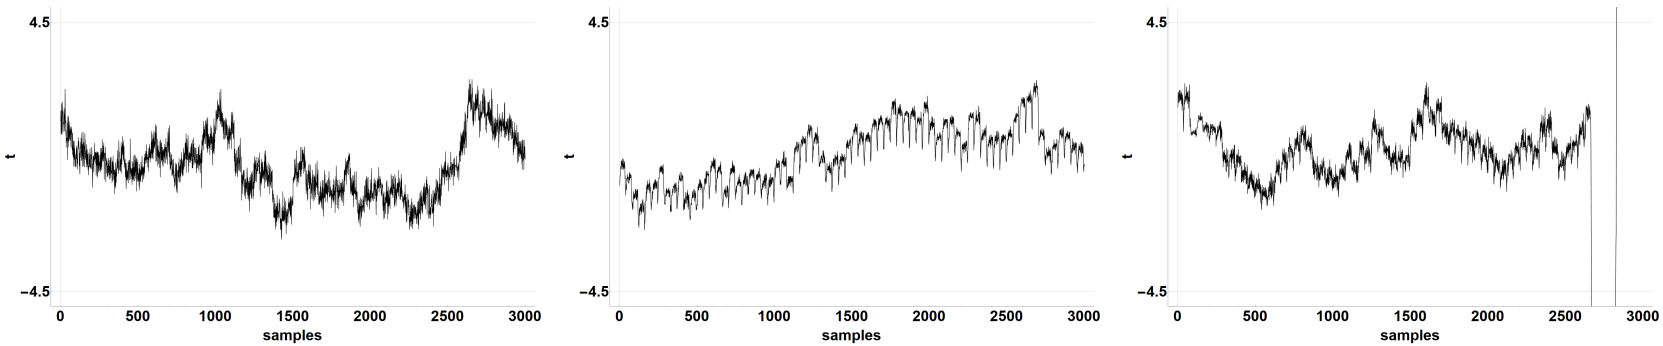

(a) single module **(b)** TMR **(c)** TMR-equivalent

Leakage at the end of encryption is caused by unmasking of the results in comparison circuit.

Leakage solution

- **Masked** outputs of modules are compared (when same mask is used)
- When outputs differ, **random** plaintext is used for faulty module identification – the encryption is repeated twice using random plaintext with different masks while **unmasked** consecutive ciphertexts are compared
- Unmasked value of the random ciphertext does not leak any information
- This approach is more area and time demanding than the original one

Conclusion

- We proposed fault-tolerant architectures exploiting redundancy introduced in masking schemes
- Our approach keeps the simplicity of modular redundancy while the overhead is decreased
- Our TMR-equivalent architecture can save up to 33% of resources in comparison with traditional TMR
  - 20% resource savings were achieved using a lightweight cipher PRESENT
  - higher savings would be achieved with more area demanding encryption algorithm like AES
- As the implemented comparison module suffers from leakage, alternative comparison logic was proposed

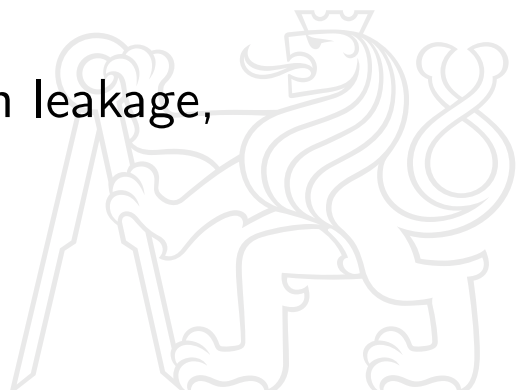

Acknowledgment

This research has been partially supported by the grant GA16-05179S of the Czech Grant Agency, "Fault-Tolerant and Attack-Resistant Architectures Based on Programmable Devices: Research of Interplay and Common Features" (2016-2018) and CTU project SGS17/213/OHK3/3T/18.