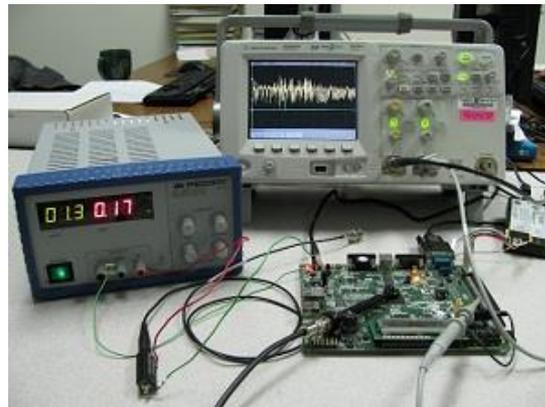


# Leveraging Inexact Computing in Post-Quantum Cryptography



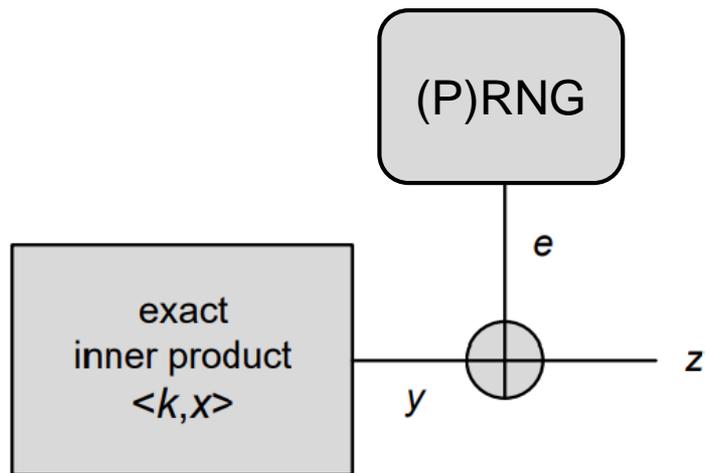
Davide Bellizia, *François-Xavier Standaert*

UCLouvain (Belgium)

**DFT 2021, Virtual**

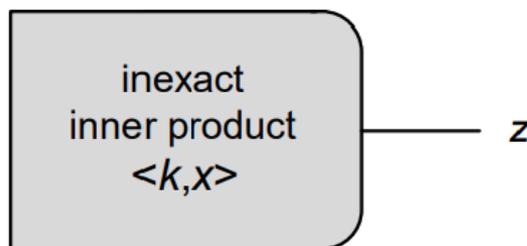
- $D_k^\varepsilon = \{ (x, \langle x, k \rangle \oplus e) : x \leftarrow \{0,1\}^n; e \leftarrow \text{Ber}_\varepsilon \}$
- LPN problem: recover  $k$  thanks to samples from  $D_k^\varepsilon$
- Assumed to be hard, quite versatile problem
- Extension in  $\mathbb{Z}_q$  (LWE) popular for PQ crypto

- $D_k^\varepsilon = \{ (x, \langle x, k \rangle \oplus e) : x \leftarrow \{0,1\}^n; e \leftarrow \text{Ber}_\varepsilon \}$
- LPN problem: recover  $k$  thanks to samples from  $D_k^\varepsilon$
- Assumed to be hard, quite versatile problem
- Extension in  $\mathbb{Z}_q$  (LWE) popular for PQ crypto
- **Standard implementation with a (P)RNG**



- May be expensive
  - E.g., require a block cipher
- Especially against leakage
  - Single probe on  $e$  makes LPN easy to solve

- Physical Inner Product ( $\tilde{\epsilon}$ -PIP)  $\approx$  device that directly outputs  $\langle x, k \rangle$  with error probability  $\tilde{\epsilon}$

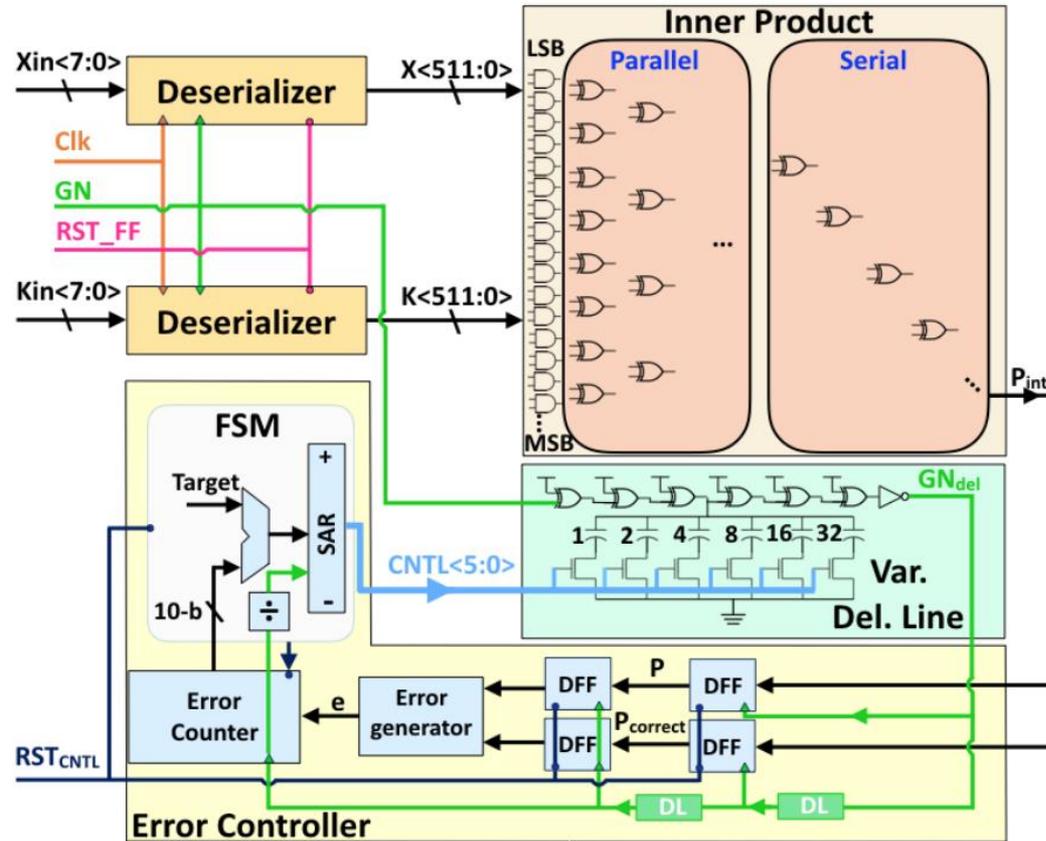


- E.g., thanks to frequency/voltage overscaling, jitter, ...

- Physical Inner Product ( $\tilde{\epsilon}$ -PIP)  $\approx$  device that directly outputs  $\langle x, k \rangle$  with error probability  $\tilde{\epsilon}$



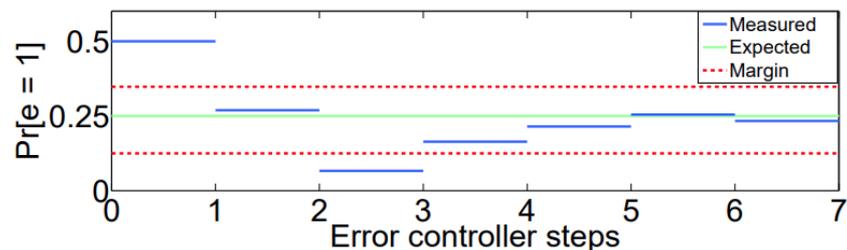
- E.g., thanks to frequency/voltage overscaling, jitter, ...
- + Conceptually appealing (don't explicitly compute  $e$ )
- + May lead to performance gains (e.g., in energy)
- Physical assumption (rather than mathematical one)
- $\Rightarrow$  *Can the error probability be controlled accurately?*
- $\Rightarrow$  *Is the physical distribution close enough to Bernoulli?*



- Parallel then serial inner product architecture
  - Glitch-induced deterministic errors in serial part
  - Error control harder when serial part depth ↗

- Example: 512-bit LPPN co-processor
  - 64-bit parallel  $\times$  8-bit serial architecture

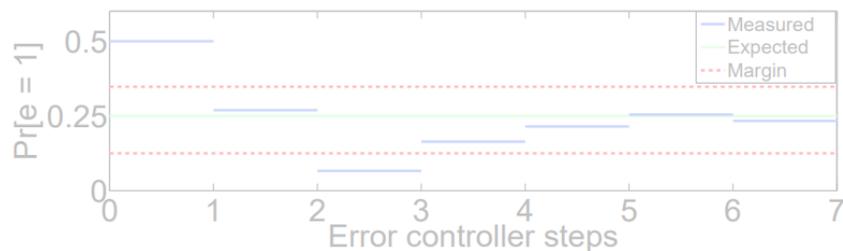
### Step 1: error calibration



- Can be made quite accurate!
  - 6 bits of control
  - 1024 queries per bit
  - E.g., Target  $\varepsilon = 0.25$

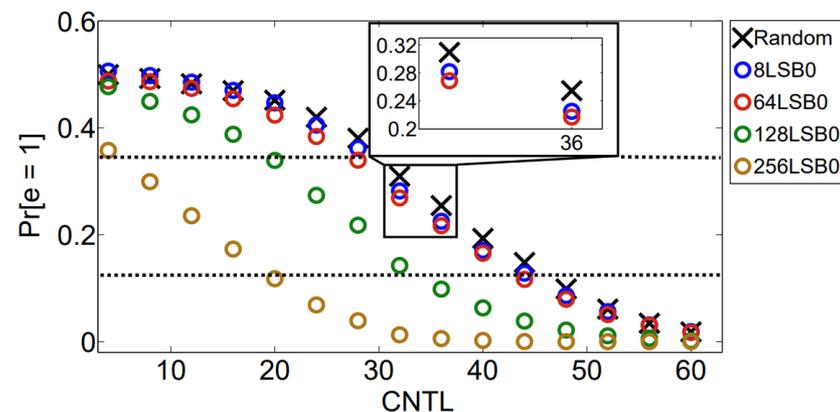
- Example: 512-bit LPPN co-processor
  - 64-bit parallel  $\times$  8-bit serial architecture

### Step 1: error calibration



- Can be made quite accurate!
  - 6 bits of control
  - 1024 queries per bit
  - E.g., Target  $\varepsilon = 0.25$

### Step 2: samples generation

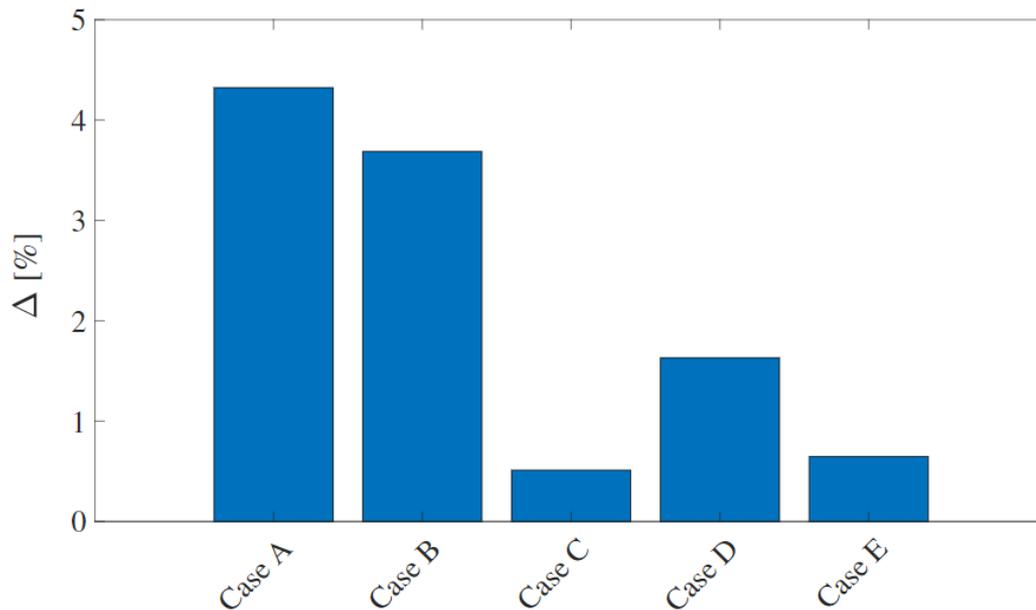


- Input-dependent  $\Pr[\text{error}]$
- E.g., by setting the bits of the serial part to all zeros / ones
- Mitigated by the parallel part

- Data-dependent  $\Pr[\text{error}]$  extends to outputs
  - Cannot be made computationally hard to exploit
  - Cannot be completely cancelled by design

- Data-dependent  $\Pr[\text{error}]$  extends to outputs
  - Cannot be made computationally hard to exploit
  - Cannot be completely cancelled by design

$$\Delta = \frac{|\widehat{\Pr}[e \leftarrow \varepsilon_0 - \text{PIP}] - \widehat{\Pr}[\varepsilon_1 - \text{PIP}]|}{2}$$



- A. Parallel
- B. Serial (min. size gates)
- C. Parallel + jitter
- D. B + power gating
- E. D + bigger gates

- LPN-OD  $\approx$  LPN with output-dependent errors

$$D_k^{\varepsilon_0, \varepsilon_1} = \left\{ \begin{array}{l} (x, \langle x, k \rangle \oplus e) : x \leftarrow \{0, 1\}^n; \\ e \leftarrow \text{Ber}_{\varepsilon_0} \text{ if } \langle x, k \rangle = 0; e \leftarrow \text{Ber}_{\varepsilon_1} \text{ if } \langle x, k \rangle = 1 \end{array} \right\}$$

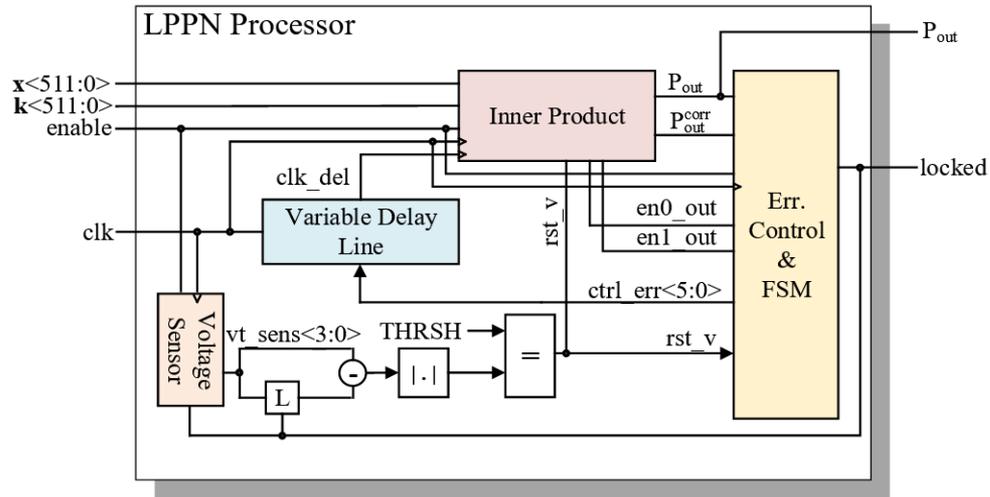
- LPN-OD  $\approx$  LPN with output-dependent errors

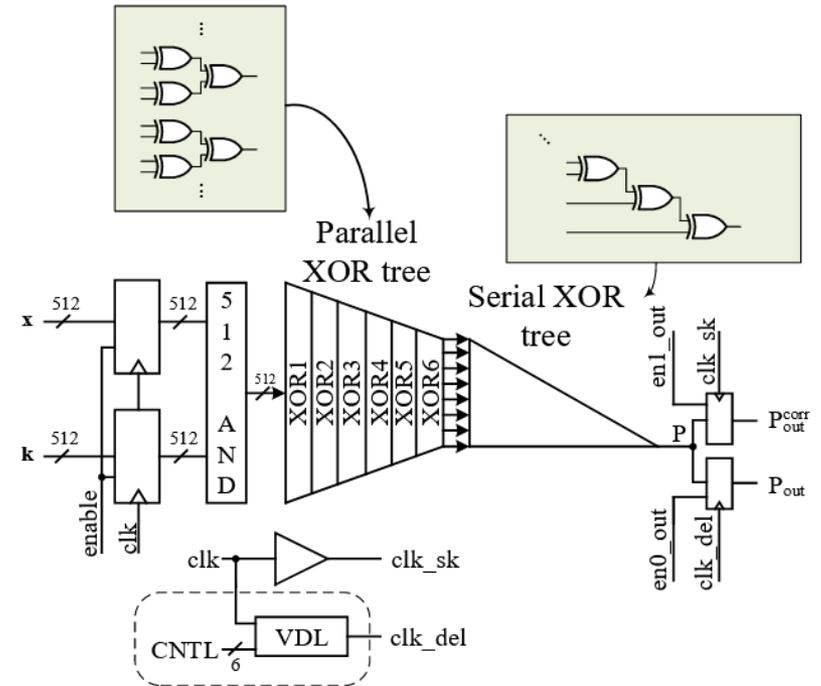
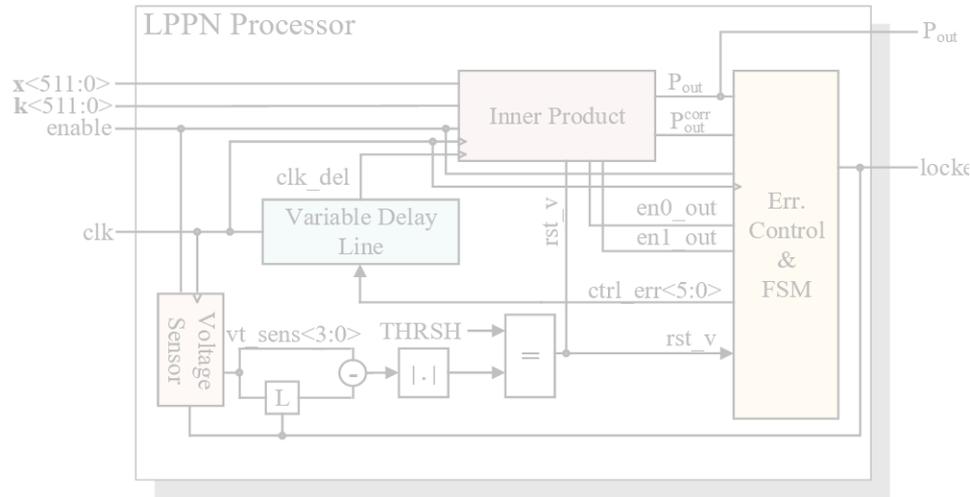
$$D_k^{\varepsilon_0, \varepsilon_1} = \left\{ \begin{array}{l} (x, \langle x, k \rangle \oplus e) : x \leftarrow \{0,1\}^n; \\ e \leftarrow \text{Ber}_{\varepsilon_0} \text{ if } \langle x, k \rangle = 0; e \leftarrow \text{Ber}_{\varepsilon_1} \text{ if } \langle x, k \rangle = 1 \end{array} \right\}$$

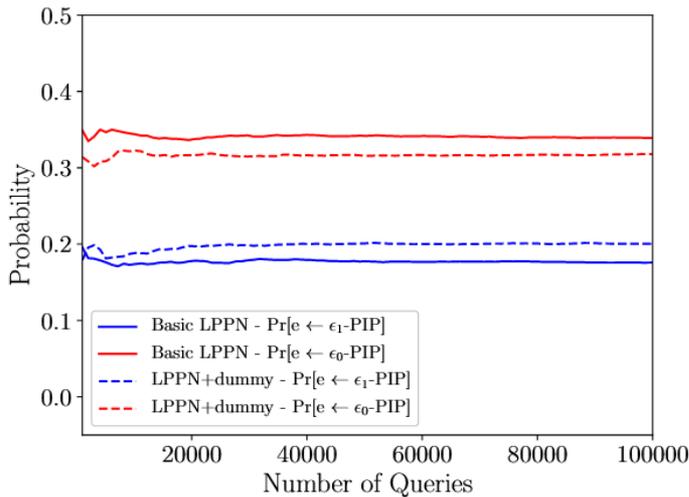
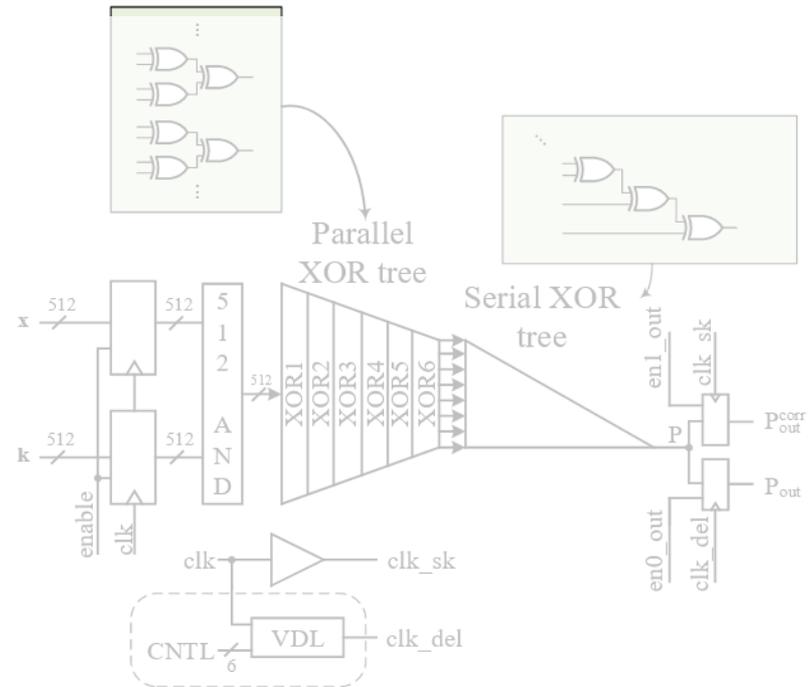
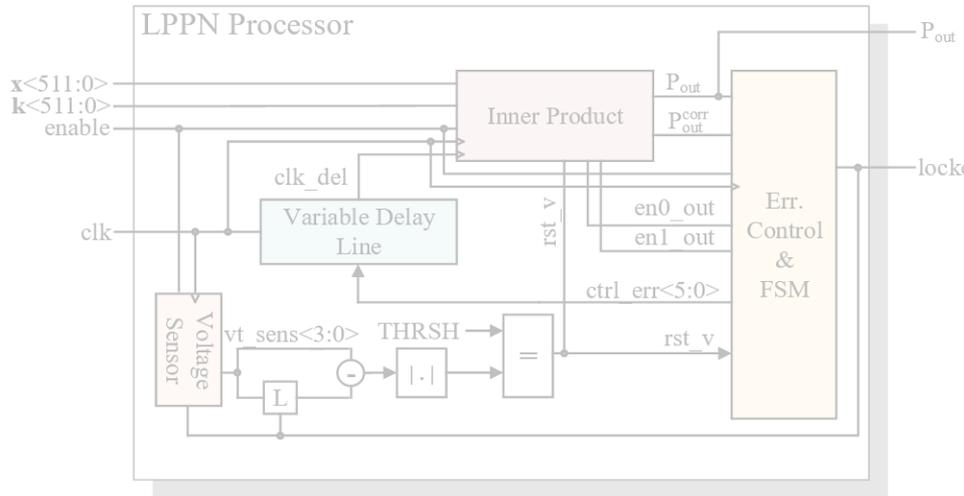
- Theorem:** LPN-OD with  $\varepsilon_0 = \varepsilon - \Delta$ ,  $\varepsilon_1 = \varepsilon + \Delta$  is at least as hard as LPN with adapted security parameter

$$\varepsilon' = \frac{\varepsilon - \Delta}{1 - 2\Delta}$$

- Lower  $\Delta$  (by design)  $\Rightarrow$  less security degradation







- Native  $\Delta$  of  $\approx 8.2\%$
- $\Rightarrow$  As hard as LPN with  $\epsilon \approx 0.2$
- Reduced to 5.8% with design tweaks (dummy operations)
- $\Rightarrow$  As hard as LPN with  $\epsilon \approx 0.22$

- Unprotected implementation:  $y \leftarrow \varepsilon\text{-PIP}(x, k)$
- Masking: share  $k = k_1 \oplus k_2 \oplus \dots \oplus k_d$  and compute
$$y \leftarrow \varepsilon\text{-PIP}(x, k_1) \oplus \langle x, k_2 \rangle \oplus \dots \oplus \langle x, k_d \rangle$$
- Helps preventing side-channel attacks

- Unprotected implementation:  $y \leftarrow \varepsilon\text{-PIP}(x, k)$
- Masking: share  $k = k_1 \oplus k_2 \oplus \dots \oplus k_d$  and compute  
$$y \leftarrow \varepsilon\text{-PIP}(x, k_1) \oplus \langle x, k_2 \rangle \oplus \dots \oplus \langle x, k_d \rangle$$

- Helps preventing side-channel attacks
- Also mitigates data-dependent errors since for  $z = \varepsilon\text{-PIP}(x, k_1)$  the adversary only sees

$\Pr[z|l] = \frac{1}{2} + \delta$ , leading to LPN-OD with

$$\varepsilon'_0 = \varepsilon_0 \cdot \left( \frac{1}{2} + \delta \right) + \varepsilon_1 \cdot \left( \frac{1}{2} - \delta \right)$$

$$\varepsilon'_1 = \varepsilon_1 \cdot \left( \frac{1}{2} + \delta \right) + \varepsilon_0 \cdot \left( \frac{1}{2} - \delta \right)$$

- Conceptually appealing but provocative
  - + Performance gains & side-channel security
  - Physical assumption (harder to assess)
- Interesting mix between physics & maths
  - Physics used to limit defaults by design
    - Many other implementations could be studied
  - Maths to prove security despite defaults
- Next step: from LP(P)N to LW(P)E and PQ crypto
  - May not be obvious with KEMS using FO-transform

# THANKS

- D. Kamel, F.-X. Standaert, A. Duc, D. Flandre, F. Berti, *Learning with Physical Noise or Errors*, in IEEE Transactions on Dependable and Secure Computing, vol 17, num 5, pp 957-971, October 2020
- D. Kamel, D. Bellizia, F.-X. Standaert, D. Flandre, D. Bol, *Demonstrating an LPPN Processor*, in the proceedings of ASHES 2018, pp 18-23, Toronto, Canada, October 2018
- D. Kamel, D. Bellizia, O. Bronchain, F.-X. Standaert, *Side-channel Analysis of a Learning Parity with Physical Noise Processor*, in the Journal of Cryptographic Engineering, vol 11, num 2, pp 171-179, June 2021
- D. Bellizia, C. Hoffmann, D. Kamel, Hanlin Liu, P. Meaux, F.-X. Standaert, Yu Yu, *Learning Parity with Physical Noise: Imperfections, Reductions and FPGA Prototype*, in IACR Transactions on Cryptographic Hardware and Embedded Systems, vol 2021, num 3, pp 390-417