Hard Physical Learning Problems
(*Applications, Cryptanalysis & Reduction Challenges*)

François-Xavier Standaert
UCLouvain Crypto Group

Novel Lattice Assumptions
Edinburgh, July 13-17, 2026

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Learning Parity with Physical Noise [KSD+20]

- LPN not so lightweight due to the need of a RNG [BL12]
- Electrical engineers struggle to have exact results [GR10]

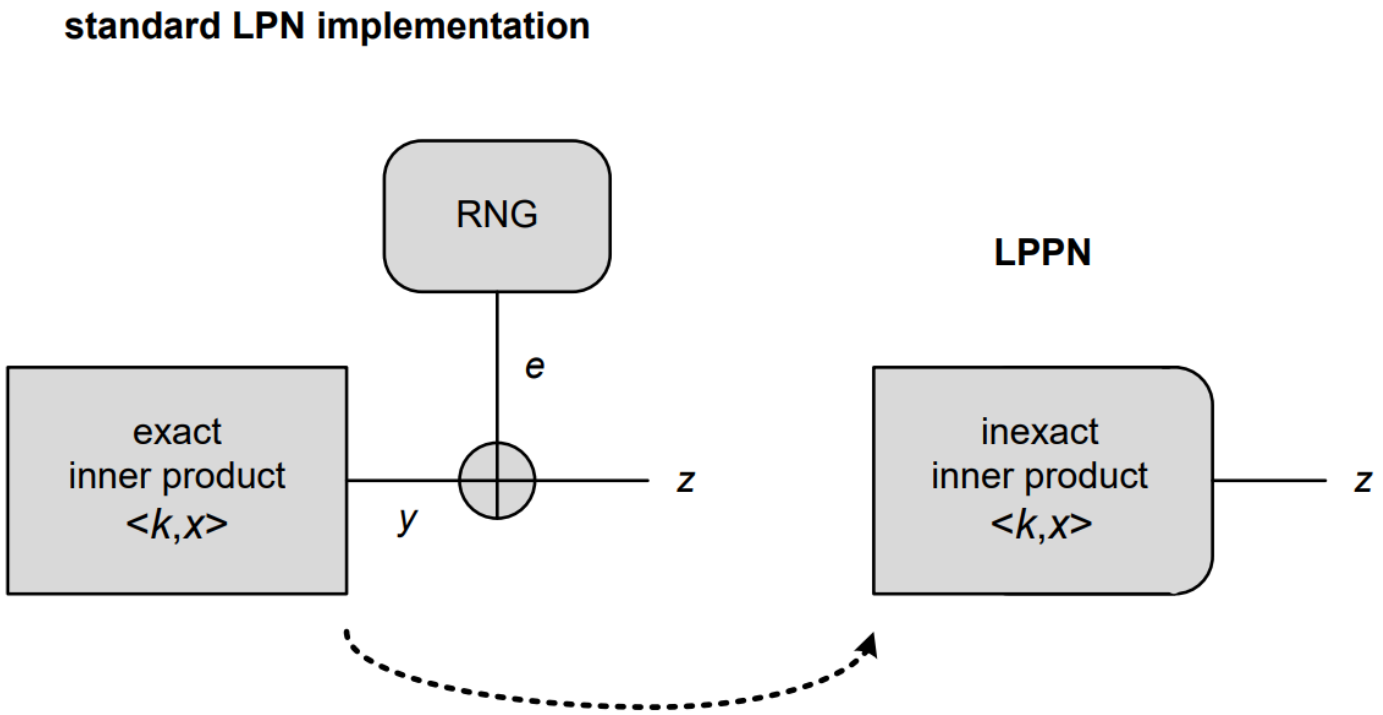

- Seems to be better for (probing) leakage as well…

Technical challenges: errors' control [KBS+18]

- Can be achieved with state-of-the-art tools (and be efficient)

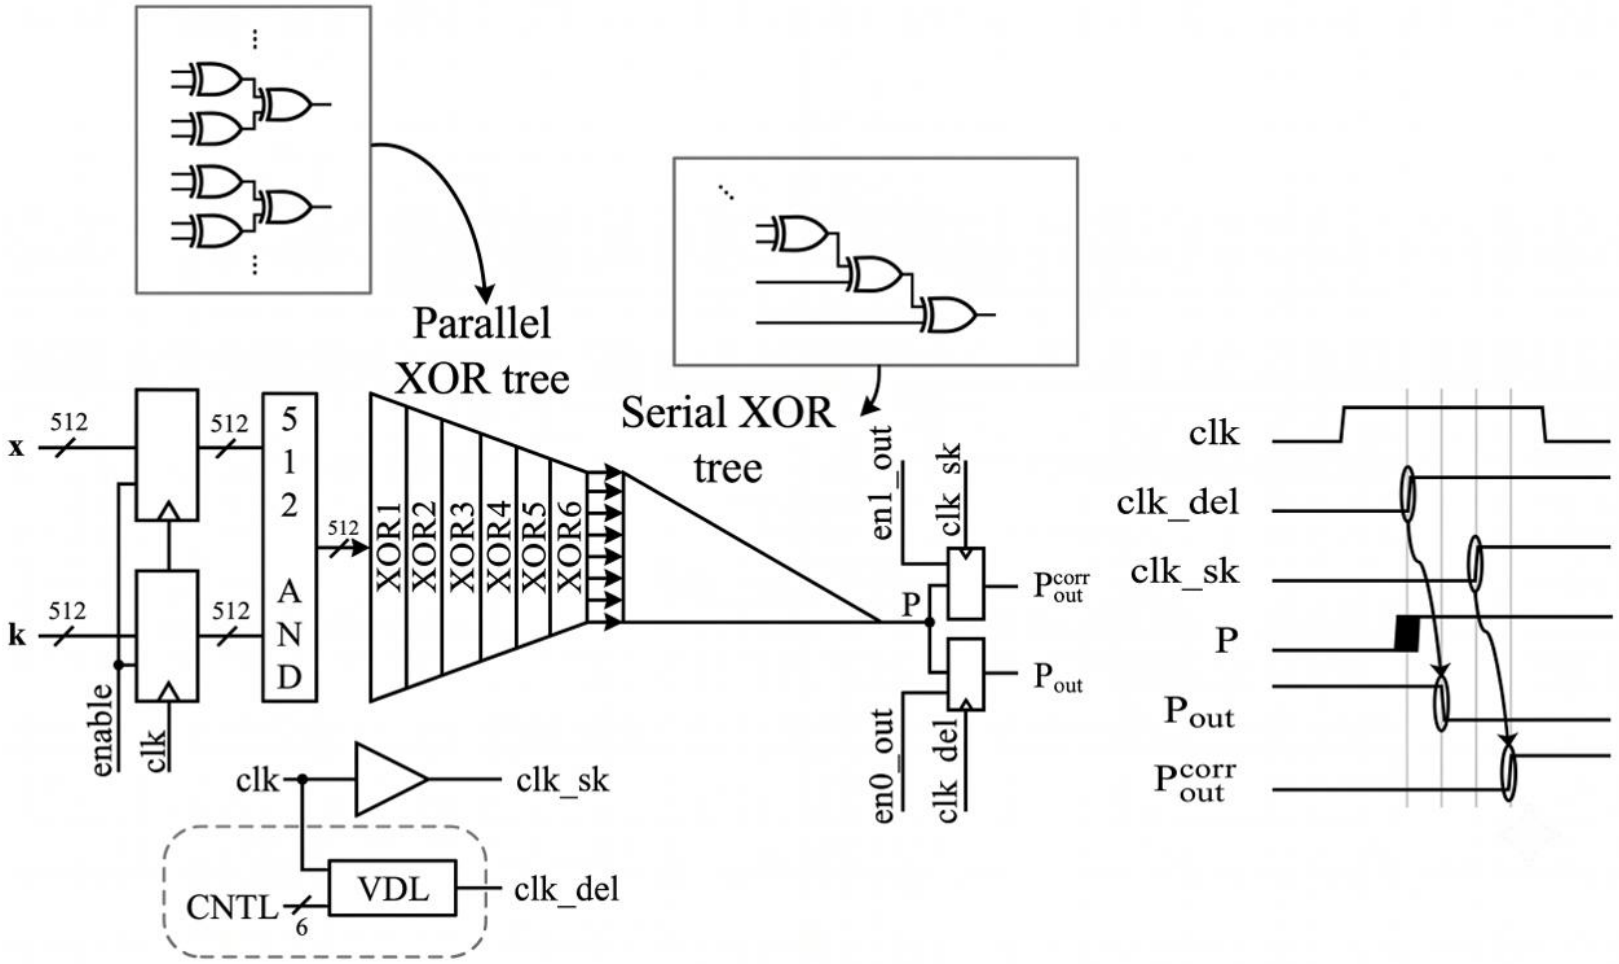

Security challenge: data-dependencies [BHK+21]

- Bernoulli error probability depends on exact output
- Modeled by LPN-OD (LPN with output dependencies)

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

- Reduction: LPN-OD with $\varepsilon_0 = \varepsilon - \Delta$, $\varepsilon_1 = \varepsilon + \Delta$ is at least as hard as LPN with adapted security parameter: $\varepsilon' = \frac{\varepsilon - \Delta}{1 - 2\Delta}$

* Input dependencies are also observed but less critical since majority of inputs concentrated around HW = $\frac{n}{2}$

Generalization: Learning With Physical Errors [BHK+22]

- Feasible as well, similar technical and security challenges

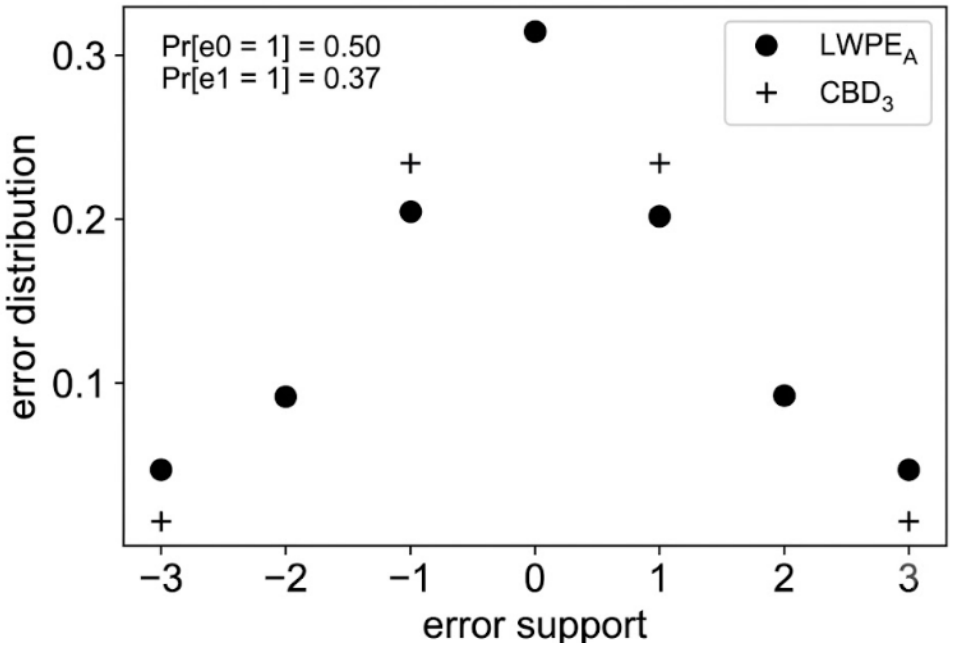

Generalization: Learning With Physical Errors [BHK+22]

- Feasible as well, similar technical and security challenges

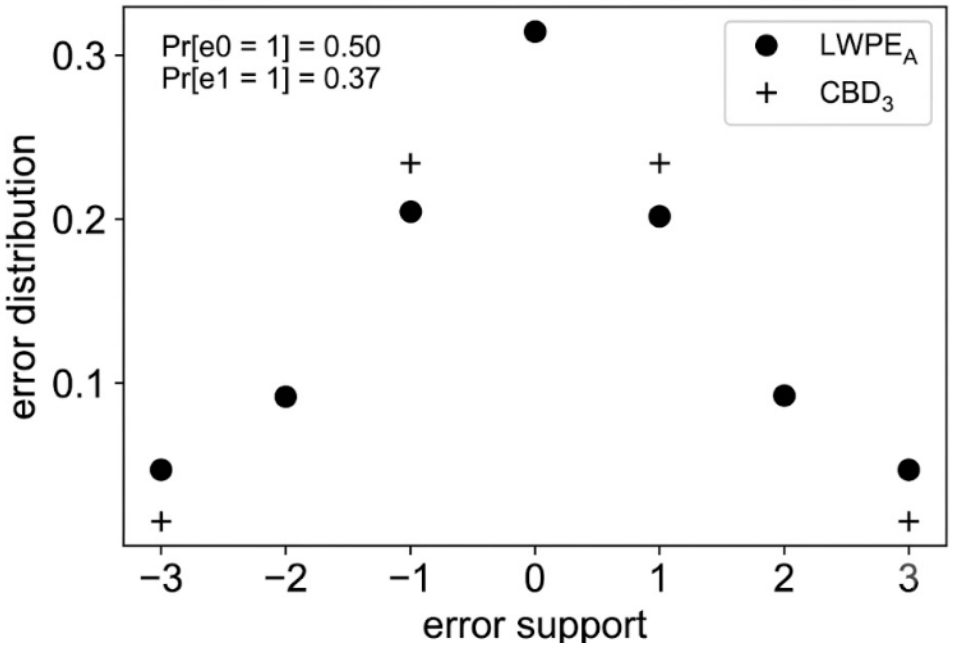

$\Rightarrow$

Sampler : $\mathsf{LWE\text{-}OD}^{n}_{\varphi_0,\ldots,\varphi_{q-1}}(\boldsymbol{k})$

$\boldsymbol{r} \leftarrow \mathbb{Z}_q^n$

$i := \langle \boldsymbol{r}, \boldsymbol{k} \rangle \in \mathbb{Z}_q$

$\boldsymbol{e} \leftarrow \varphi_i$

$b := \langle \boldsymbol{r}, \boldsymbol{k} \rangle \boxplus e$

Return $(\boldsymbol{r}, b)$

Open question 1: reduction from LWE to (any/some) LWE-OD?
E.g., if φ_i's are (not negligibly) close to some φ in terms of SD?

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Fresh re-keying against leakage [MSG+10]

- Masking (≈ MPC on silicon) block ciphers is expensive
- Separation of duties: easy-to-mask RK, cryptographic BC

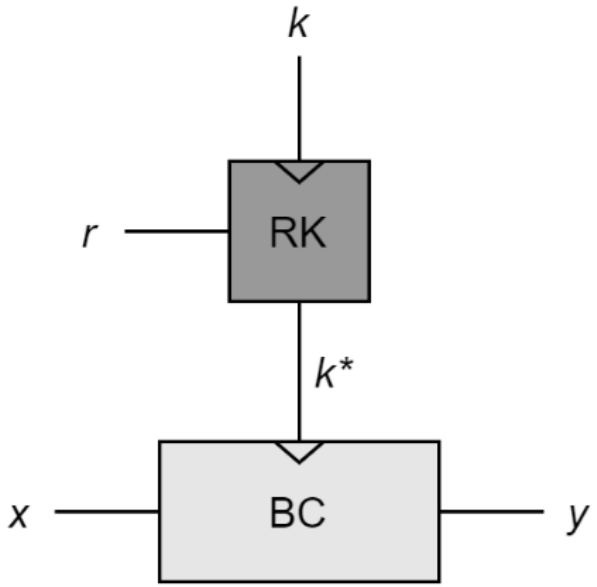

Fresh re-keying against leakage [MSG+10]

- Masking (≈ MPC on silicon) block ciphers is expensive
- Separation of duties: easy-to-mask RK, cryptographic BC

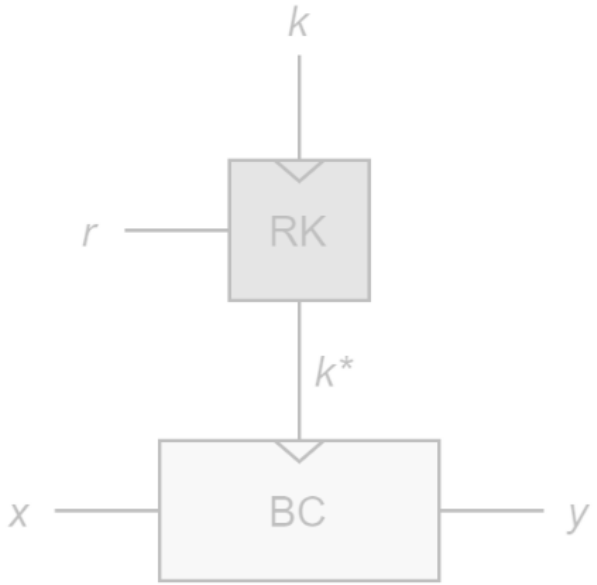

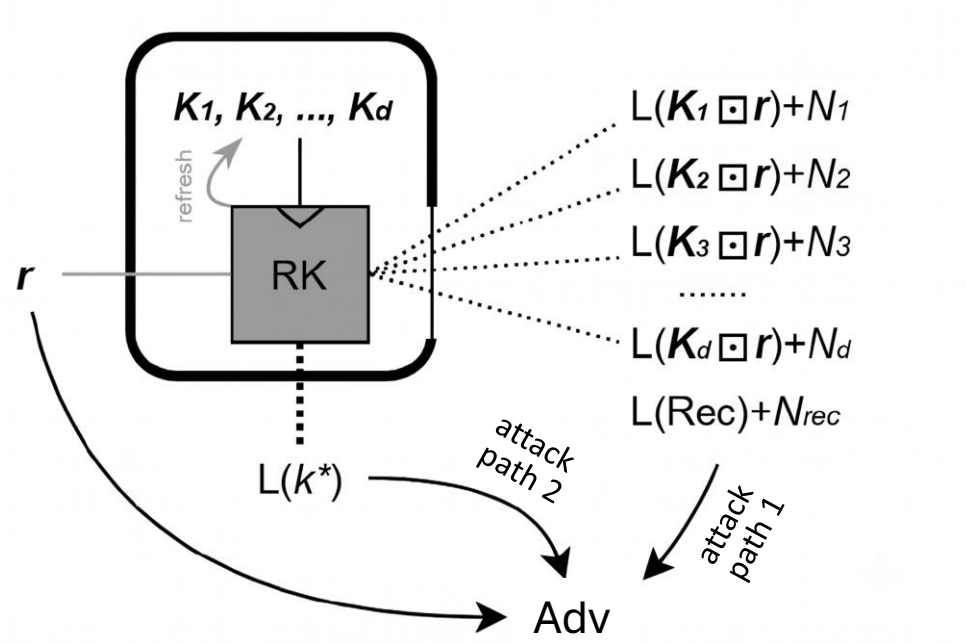

- Attack path 1 well understood, attack path 2 less [BFG14]

Model #1: Learning Parity with Leakage [DFH+16]

- Tailored to analyzed LPN-based re-keying

$$\mathcal{D}_{\mathsf{LPL}}^{\sigma_\eta} := (\mathbf{r}, \ell = \langle \mathbf{r}, \mathbf{k} \rangle + \eta) \text{ for } \mathbf{r} \leftarrow \mathbb{F}_{2^n}, \eta \leftarrow \mathcal{N}(0, \sigma_\eta)$$

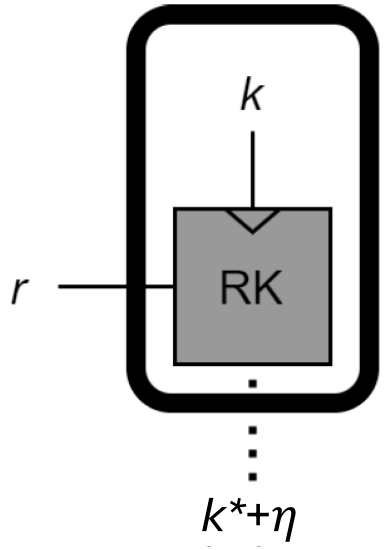

- With the (physical) noise addition in the reals!

- Reduction: LPL with std. dev. σ_η is at least as hard as LPN with error proba. ε, with σ_η in $\mathcal{O}\left(\frac{\sqrt{\log m}}{\log\left(\frac{1-\varepsilon}{\varepsilon}\right)}\right)$ and m is the # of samples

⇒ Drawbacks: noise and limitation to LPN-based re-keying

Model #2: Learning With Physical Rounding [DMM+21]

- Applicable to any (preferably large) field/ring

$$\mathcal{D}^{\delta}_{\mathrm{LWPR}} := \big(r, \delta(\langle r, k\rangle)\big), \text{ for } r \in \mathbb{F}_?, \mathbb{Z}_?$$

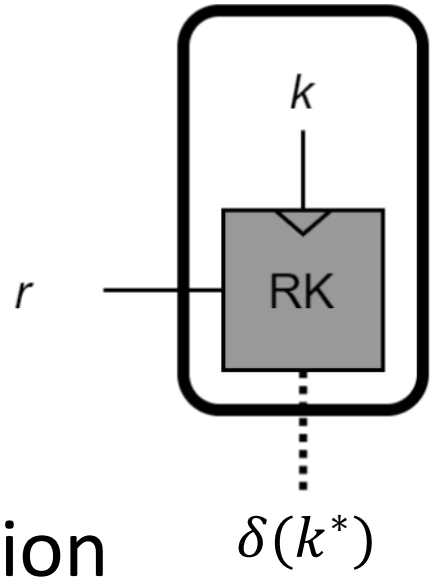

- Deterministic δ, typical example: $\delta = \mathrm{HW}$
  - In $\mathbb{Z}_{2^\alpha}$, LSB of HW gives a linear equation
  - In $\mathbb{F}_q$ with small q, HW=0 gives a linear equation

$\Rightarrow$ Need large prime fields (or small prime fields leaking in parallel)

- Symmetric cryptanalysis (e.g., algebraic) seems hard

Model #2: Learning With Physical Rounding [DMM+21]

- Applicable to any (preferably large) field/ring

$$\mathcal{D}^{\delta}_{\mathrm{LWPR}} := \big(r, \delta(\langle r, k\rangle)\big), \text{ for } r \in \mathbb{F}_?, \mathbb{Z}_?$$

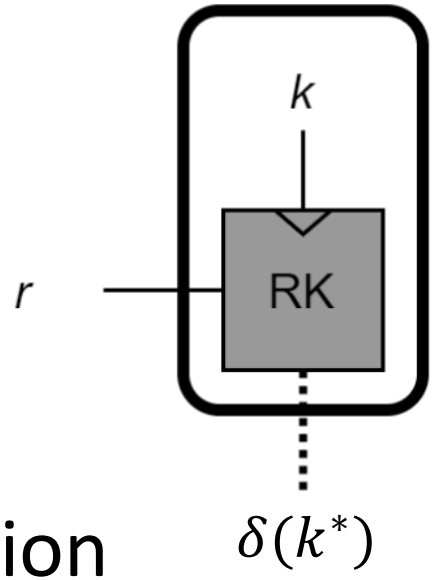

- Deterministic δ, typical example: $\delta = \mathrm{HW}$
  - In $\mathbb{Z}_{2^\alpha}$, LSB of HW gives a linear equation
  - In $\mathbb{F}_q$ with small q, HW=0 gives a linear equation

$\Rightarrow$ Need large prime fields (or small prime fields leaking in parallel)

- Symmetric cryptanalysis (e.g., algebraic) seems hard

Open question 2: reduction from LWE/LWR to (any/some) LWPR?
(e.g., is it enough to have large q or parallelism + MI$(K*, \delta(K*))$?)

Model #3: LWPR and Noise [HRR+26]

- Same as LWPR with Gaussian noise addition

$$\mathcal{D}^{\delta,\sigma_\eta}_{\mathrm{LWPRN}} := (r, \delta(\langle r, k\rangle) + \eta), \text{ for } r \in \mathbb{F}_?, \mathbb{Z}_?$$

- With the (physical) noise addition in the reals!

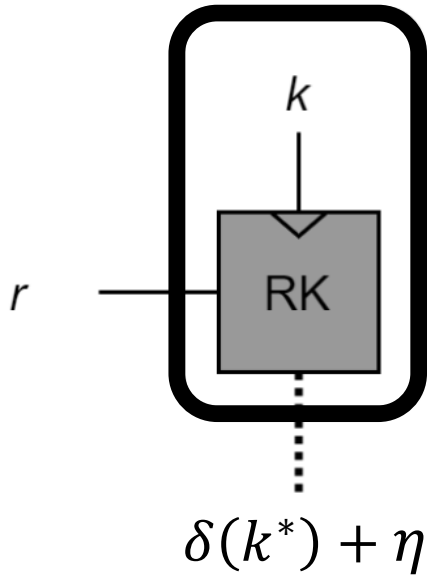

- Reduction: LWPRN with $\delta = \mathrm{HW}$ and std. dev. σ_η is at least as hard as LWE with std. dev. σ_e, with $\sigma_\eta \geq \sqrt{m}\log(\sqrt{\lambda}\,\sigma_e)$
  - Statistical argument ☹ but noise independent of field size ☺
    + Noise gives security in $\mathbb{Z}_{2^\alpha}$ even without parallelism ☺

Interlude: statistical order heuristic argument

- Reduction requires to make distributions indistinguishable

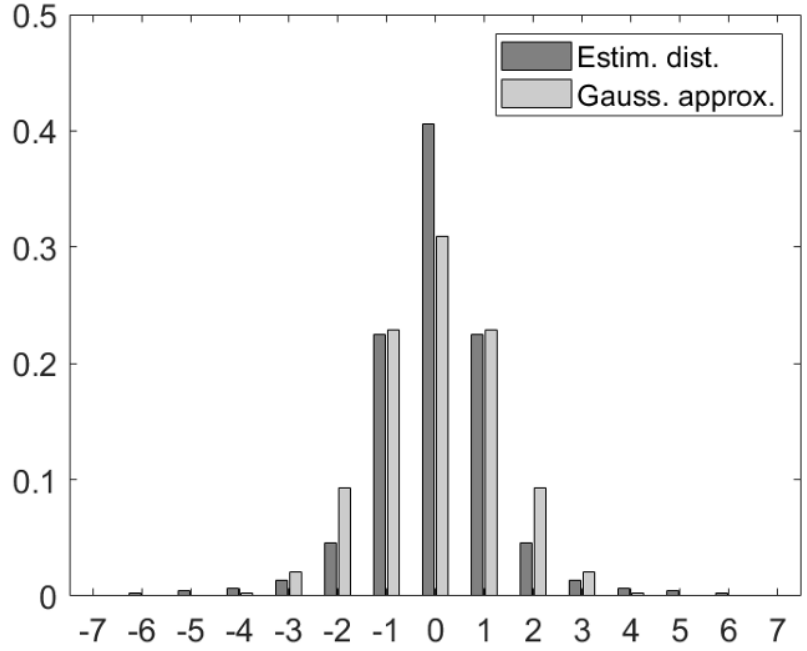

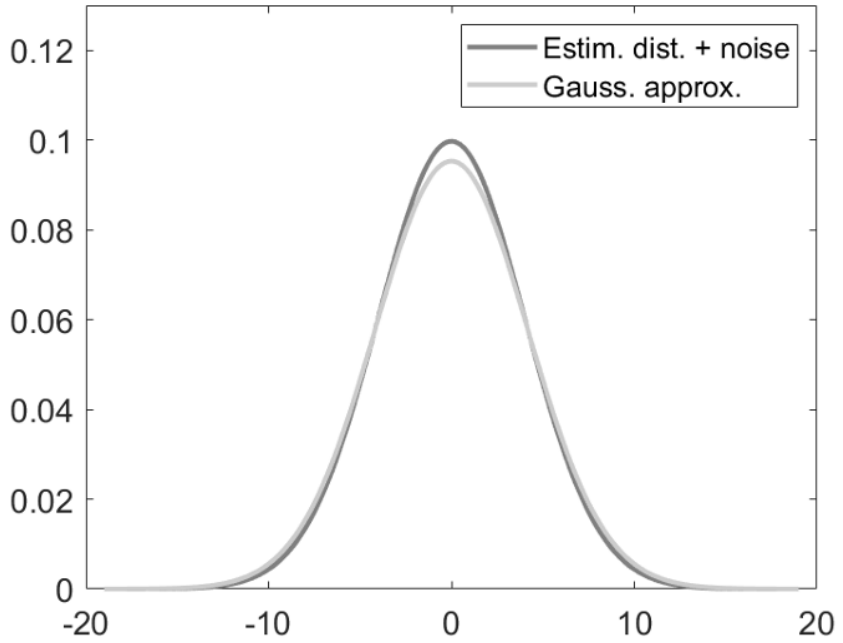

- Same means, same variances, same skewness, kurtosis differs

$\Rightarrow$ Distinguishing complexity scales with $\left(\sigma_\eta^2\right)^{-4}$ rather than $\left(\sigma_\eta^2\right)^{-1}$

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Concrete attacks against (e.g.), ML-KEM [BCS+26]

- Leakage ⇒ vector of probabilities for every key coefficient

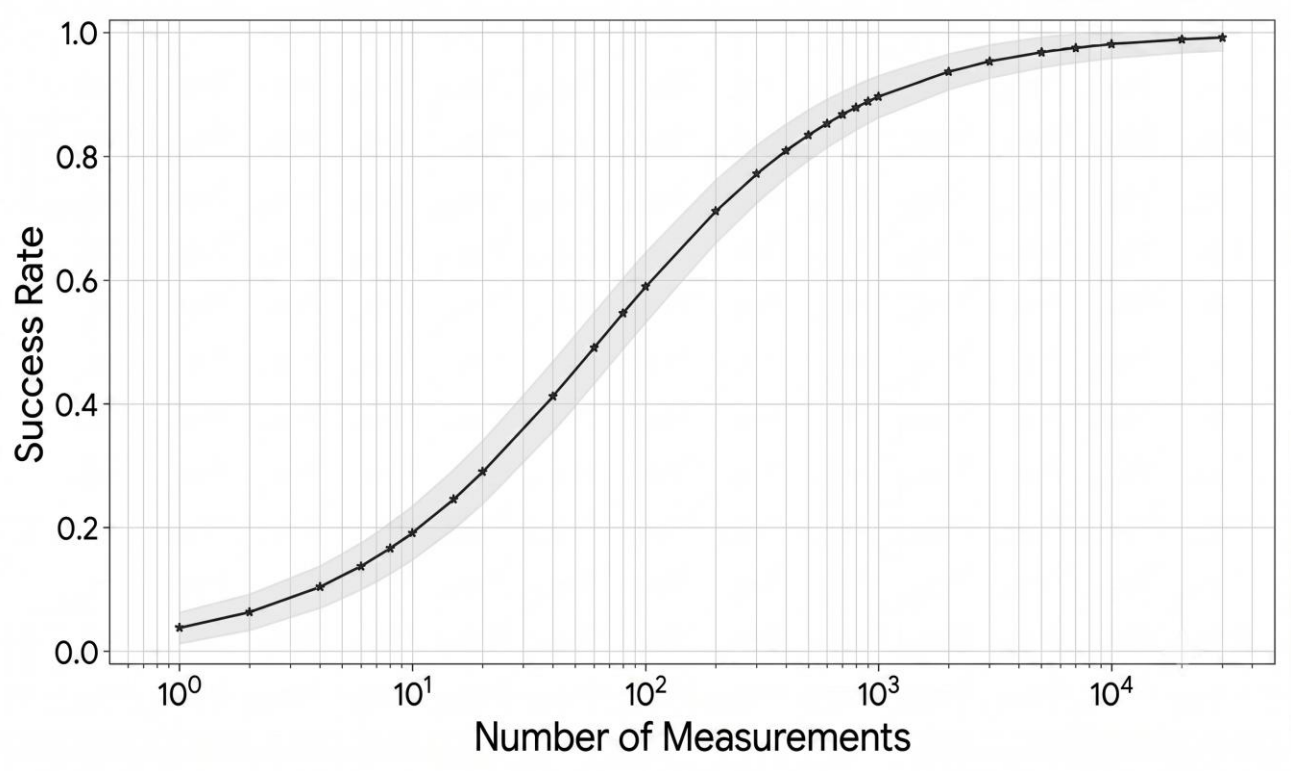

- Option #1: measure till all coefficients leaked in full
- Option #2: enumerate
- Option #3: plug into a leaky lattice estimator (currently by feeding the maximum variance ☹)

Concrete attacks against (e.g.), ML-KEM [BCS+26]

- Leakage ⇒ vector of probabilities for every key coefficient

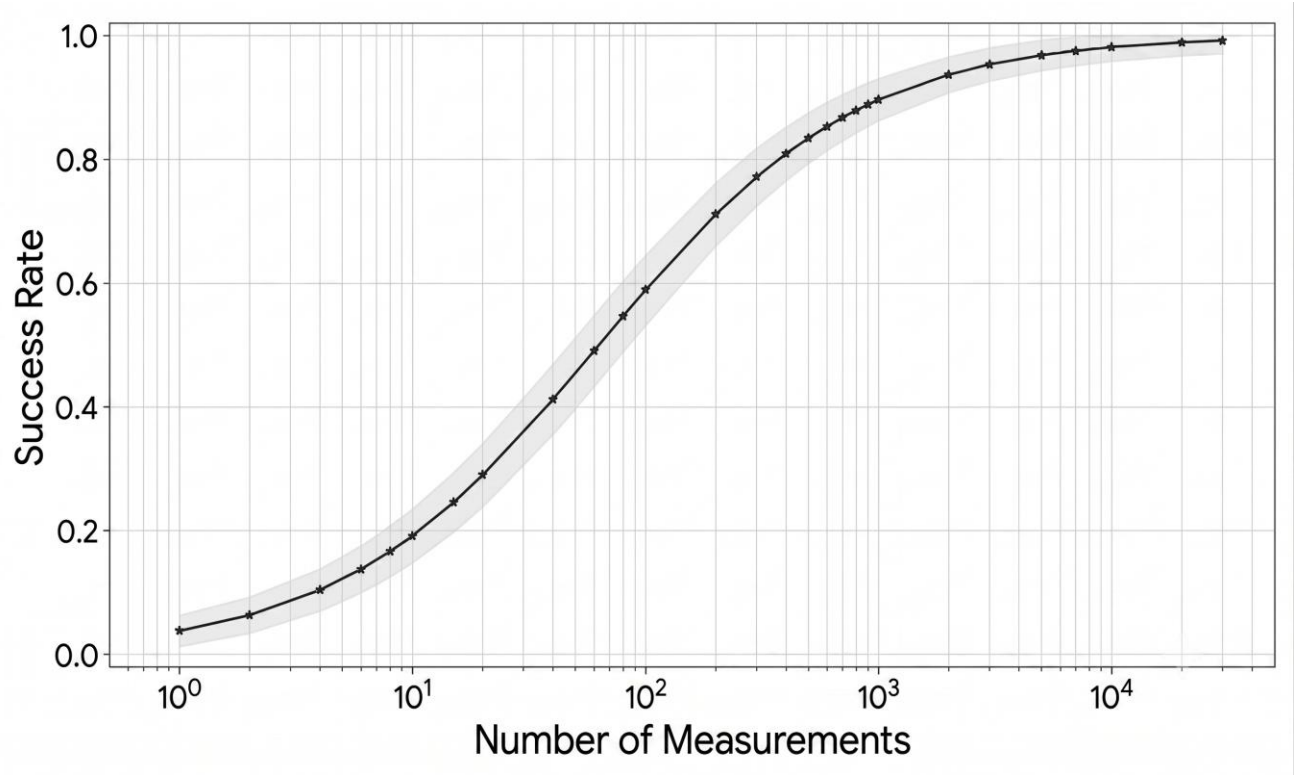

- Option #1: measure till all coefficients leaked in full
- Option #2: enumerate
- Option #3: plug into a leaky lattice estimator (currently by feeding the maximum variance ☹)

- **Open question 3:** more granular lattice estimator with hints?

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Unifying attempts [HMR+26,HRR+26]

- LWE-OD generalizes
  - LWE ($\varphi_i = \varphi$) for all I
  - LW(P)R: φ_i outputs round$(i) - i$
  - LPN when q=2 and φ_i =Ber
  - LPN-OD when q=2
- Useful abstraction?

Sampler : $\mathsf{LWE\text{-}OD}^n_{\varphi_0,\ldots,\varphi_{q-1}}(\boldsymbol{k})$

$\boldsymbol{r} \leftarrow \mathbb{Z}_q^n$

$i := \langle \boldsymbol{r}, \boldsymbol{k} \rangle \in \mathbb{Z}_q$

$\boldsymbol{e} \leftarrow \varphi_i$

$b := \langle \boldsymbol{r}, \boldsymbol{k} \rangle \boxplus e$

Return $(\boldsymbol{r}, b)$

Unifying attempts [HMR+26,HRR+26]

- LWE-OD generalizes
  - LWE ($\varphi_i = \varphi$) for all I
  - LW(P)R: φ_i outputs round$(i) - i$
  - LPN when q=2 and φ_i =Ber
  - LPN-OD when q=2
- Useful abstraction?

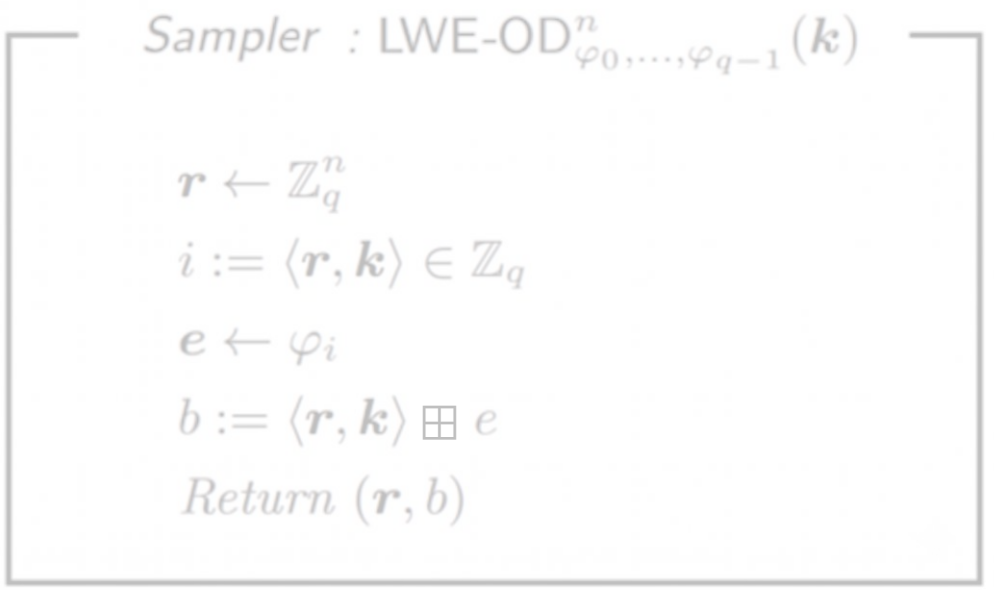

- Reduction from LWE to LWPRN uses “LWE with physical hint” as an intermediate abstraction (with HW hints for now)
  - Hard physical learning problems ≈ learning problems with hints that cannot be completely controlled since physical

Outline

- Inexact computing (LPPN/LPN-OD, LWPE/LWE-OD)
- Fresh re-keying against leakage (LPL, LWPR, LWPRN)
- (Interlude: statistical order heuristic argument)
- Concrete cryptanalysis (granular leaky estimator)
- Unifying attempts (LWE-OD and LWE with hints)
- Taxonomy & conclusions

Taxonomy (I)

- Leakage functions:
  - Hamming weight: $\mathrm{L}(v) = \sum v(i)$, with $v(i)$ the bits of v
  - Linear: $\mathrm{L}(v) = \sum \alpha_i \cdot v(i)$, with α_i real-valued
    - Need quantization of L ideally, α_i currently [HMM+23]
  - Quadratic: $\mathrm{L}(v) = \sum \alpha_i \cdot v(i) + \sum \beta_{i,j} \cdot v(i) \cdot v(j)$
  - The same with additive Gaussian noise
  - …
  - $\mathrm{L}(v)$ such that $\mathrm{MI}(V, \mathrm{L}(V)) < \mathrm{H}(V) \approx$ entropic LWE?

more realistic

Taxonomy (II)

- Types of hints: $y = \langle r, k \rangle \boxplus e$

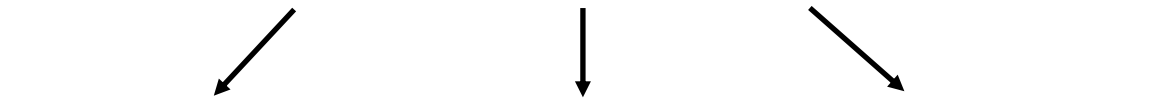

- Are they connected / separated (they all leak in practice)?

- We started with the leakage on y assuming $\langle r, k \rangle$ is masked
  - Statistical argument frustrating: information accumulates with # of samples m (typically large in leakage applications)

  $\Rightarrow$ A computational argument would be great (cf. Question 2)
- Status of leakage on e is less clear for now (& less exploited?)

Conclusions

- Re-keying ⇒ leveled implementations (vs. uniform protections)
  - Right now: huge savings in symmetric crypto (AE)
  - Potential: application to ML-KEM, ML-DSA
  - Future: new PQ schemes with better leakage-resistance
- Inexact computing's application less obvious (require solving technological challenges) but LWE-OD is a nice generalization?
- Current estimators not perfectly suited to concrete leakage
- Cryptanalysis vs. reductions tension?
  - Cryptanalysis struggles more with noisy leakage?
  - Reductions struggle more with deterministic leakage?

References ?

- [BCS+26] Brieuc Balon, Gaëtan Cassiers, Thibaud Schoenauen, François-Xavier Standaert, *How Strong is the FO-Calypse, Really? Instantiating Plaintext-Checking Oracles against Masked Software Implementations of ML-KEM,* in IACR Trans. Cryptogr. Hardw. Embed. Syst. 2026(4)
- [BFG14] Sonia Belaïd, Pierre-Alain Fouque, Benoît Gérard: *Side-Channel Analysis of Multiplications in GF(2128) - Application to AES-GCM*. ASIACRYPT (2) 2014: 306-325
- [BHK+21] Davide Bellizia, Clément Hoffmann, Dina Kamel, Hanlin Liu, Pierrick Méaux, François-Xavier Standaert, Yu Yu: *Learning Parity with Physical Noise: Imperfections, Reductions and FPGA Prototype*. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2021(3): 390-417 (2021)
- [BL12] Daniel J. Bernstein, Tanja Lange: *Never Trust a Bunny*. RFIDSec 2012: 137-148
- [DFH+16] Stefan Dziembowski, Sebastian Faust, Gottfried Herold, Anthony Journault, Daniel Masny, François-Xavier Standaert: *Towards Sound Fresh Re-keying with Hard (Physical) Learning Problems*. CRYPTO (2) 2016: 272-301
- [DMM+21] Sébastien Duval, Pierrick Méaux, Charles Momin, François-Xavier Standaert: *Exploring Crypto-Physical Dark Matter and Learning with Physical Rounding Towards Secure and Efficient Fresh Re-Keying*. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2021(1): 373-401 (2021)
- [GR10] Swaroop Ghosh, Kaushik Roy: *Parameter Variation Tolerance and Error Resiliency: New Design Paradigm for the Nanoscale Era*. Proceedings of the IEEE, vol 98, num 10, pp 1718-1751, 2010
- [HMR+26] Clément Hoffmann, Pierrick Méaux, Mélissa Rossi, François-Xavier Standaert: *Learning with Errors with Output Dependencies: LWE, LWR, and Physical Learning Problems Under the Same Umbrella.* PQCrypto (1) 2026: 287-318
- [HRR+26] Clément Hoffmann, Emeline Repel, Adeline Roux-Langlois, François-Xavier Standaert, *Hardness of Learning with Physical Rounding and Noise from Learning With Errors*, to appear in IACR Trans. Cryptogr. Hardw. Embed. Syst. 2026(3)
- [KBS+18] Dina Kamel, Davide Bellizia, F.-X. Standaert, Denis Flandre, David Bol: *Demonstrating an LPPN Processor*. ASHES 2018: 18-23
- [KSD+20] Dina Kamel, François-Xavier Standaert, Alexandre Duc, Denis Flandre, Francesco Berti: *Learning with Physical Noise or Errors*. IEEE Trans. Dependable Secur. Comput. 17(5): 957-971 (2020)
- [MSG+10] Marcel Medwed, François-Xavier Standaert, Johann Großschädl, Francesco Regazzoni: *Fresh Re-keying: Security against Side-Channel and Fault Attacks for Low-Cost Devices*. AFRICACRYPT 2010: 279-296