Algebraic Key-Recovery Side-Channel Attack on *Classic McEliece*

Michaël Bulois Pierre-Louis Cayrel Vlad-Florin Drăgoi
Vincent Grosso

Selected Areas in Cryptography
August 11–15, 2025
Toronto, Ontario

Motivation

Post-quantum cryptography

PQC Standardization process
Lattice-based, code-based, and hash-based solutions

Classic McEliece KEM
Arrived in Round 4 at NIST and ongoing candidate at ISO

Side-channel attacks
Lattice-based and code-based implementations are recently target to side-channel attacks.
What is the practical security of such cryptosystems?
Is the reference implementation of *Classic McEliece* secure against side-channel attacks?

Classic McEliece - key generation

Private key: $\mathrm{sk} = (\gamma, \mathcal{L})$, where $\mathcal{L} \subseteq \mathbb{F}_{2^m}$ and $\gamma \in \mathbb{F}_{2^m}[x]$ irreducible and $\deg(\gamma) = t$

Public key: $\mathrm{pk} = \boldsymbol{T}$, where $\boldsymbol{T}$ is a binary $mt \times (n - mt)$ matrix derived from sk.

Key recovery

Goppa code equivalence problem:
Given pk (public Goppa code) find sk (private Goppa code)

Breaking Goppa with hints

- Given γ find $\mathcal{L}$: SSA by Sendrier
- Given $\mathcal{S} \subseteq \mathcal{L}$ find sk: BGH by Kirshanova and May

How to obtain these hints?

Classic McEliece - decapsulation

Algorithm The decapsulation algorithm of the *Classic McEliece* KEM

Input: Ciphertext $\boldsymbol{z}$ and private key $\text{sk} = (\gamma, \mathcal{L})$

Output: Session key K

1: Compute $\boldsymbol{v} = (\boldsymbol{z}, 0, \ldots, 0)$ of length n

2: Construct the matrix:

$$\boldsymbol{H}_{\text{priv}_{\gamma^2}} = \begin{pmatrix} \gamma(\alpha_0)^{-2} & \cdots & \gamma(\alpha_{n-1})^{-2} \\ \vdots & \ddots & \vdots \\ \alpha_0^{2t-1}\gamma(\alpha_0)^{-2} & \cdots & \alpha_{n-1}^{2t-1}\gamma(\alpha_{n-1})^{-2} \end{pmatrix}$$

3: Compute the syndrome: $\boldsymbol{s} = \boldsymbol{H}_{\text{priv}_{\gamma^2}} \boldsymbol{v}^T$

4: Use the Berlekamp–Massey algorithm to compute the error locator polynomial $\sigma(x)$

5: Evaluate $\sigma(\alpha_0), \ldots, \sigma(\alpha_{n-1})$ for $\alpha_i \in \mathcal{L}$ and recover the error vector $\boldsymbol{e}$

6: Compute $K = \text{hash}(1 \| \boldsymbol{e} \| \boldsymbol{z})$

7: **return** K

Side-Channel Information

Let $\boldsymbol{H}_{\mathrm{priv}_{\boldsymbol{\gamma}^2}} = \begin{pmatrix} \beta_0 & \dots & \beta_{n-1} \\ \alpha_0\beta_0 & \dots & \alpha_{n-1}\beta_{n-1} \\ \vdots & \ddots & \vdots \\ \alpha_0^{2t-1}\beta_0 & \dots & \alpha_{n-1}^{2t-1}\beta_{n-1} \end{pmatrix}$, during the syndrome computation there is a side-channel leakage[1] which allows to obtain :

$$\boldsymbol{H}_{\mathrm{wt}} = \begin{pmatrix} \mathrm{wt}(\beta_0) & \dots & \mathrm{wt}(\beta_{n-1}) \\ \mathrm{wt}(\alpha_0\beta_0) & \dots & \mathrm{wt}(\alpha_{n-1}\beta_{n-1}) \\ \vdots & \ddots & \vdots \\ \mathrm{wt}(\alpha_0^{2t-1}\beta_0) & \dots & \mathrm{wt}(\alpha_{n-1}^{2t-1}\beta_{n-1}) \end{pmatrix}$$

[1] V. Dragoi et al., Full Key-Recovery Cubic-Time Template Attack on Classic McEliece Decapsulation, TCHES 2025

Conjecture

For almost all degree-m monic irreducible polynomials $\zeta \in \mathbb{F}_2[x]$, the extension field $\mathbb{F}_{2^m} \cong \mathbb{F}_2[X]/(\zeta)$ is such that almost all pairs $(\alpha, \beta) \in \mathbb{F}_{2^m}^* \times \mathbb{F}_{2^m}^*$ can be uniquely determined from $\boldsymbol{H}_{\mathrm{wt}}$, provided that t is sufficiently large [2]

$\boldsymbol{H}_{\mathrm{wt}}$ is a distinguisher for $\mathcal{L}$!

[2] V. Dragoi et al., Full Key-Recovery Cubic-Time Template Attack on Classic McEliece Decapsulation, TCHES 2025

Our Contribution

Robust and Efficient Attack
Novel algebraic method tolerates noisy leakage and scales efficiently to large m.

Algebraic Framework for Leakage
Links Hamming weights to secrets under noise; applies to any $\mathbb{F}_2$-linear leakage.

Theoretical Insights and Generalization
Deepens understanding of Conjecture 1; explains why field elements remain distinguishable under weaker leakage.

Linear Algebra

The Sequence $\mathcal{W}_{\alpha,\beta}$

Let $\alpha, \beta \in \mathbb{F}_{2^m}$.
We study the sequence:

$$\mathcal{W}_{\alpha,\beta} = (\mathrm{wt}_2(\alpha^i \beta))_{i \in \mathbb{N}} \in \mathbb{F}_2^{\mathbb{N}}$$

Where $\mathrm{wt}_2(x)$ is the mod 2 of the Hamming weight of the binary representation of x

Key Observations:

- wt_2 is an $\mathbb{F}_2$-linear form (φ)
- Multiplication by α defines an endomorphism h_α.
- h_α^* acts on the dual space: $h_\alpha^*(\varphi)(x) = \varphi(\alpha x)$.

LFSR Interpretation and Dual Basis

Assume: $\mathbb{F}_2[\alpha] = \mathbb{F}_{2^m} \quad \Rightarrow \quad h_\alpha$ has irreducible characteristic polynomial.
Then:

$$(\varphi_\alpha[i] := (h_\alpha^*)^i(\mathrm{wt}_2))_{0 \le i < m}$$

is a basis B_α^* of $\mathbb{F}_{2^m}^*$.

LFSR Viewpoint:

$$\mathcal{W}_{\alpha,\beta} = \left(\varphi_\alpha[i](\beta)\right)_i$$

is the output of an LFSR over $\mathbb{F}_2$ with feedback polynomial χ_α.

$$\deg(\chi) = m = \dim(\mathbb{F}_{2^m}) \Rightarrow \text{ smallest sequence of LFSR has length } 2m$$

Noise-Free Reconstruction: Overview

Given the matrix $\boldsymbol{H}_{\mathrm{wt}}$, we aim to reconstruct the hidden pairs (α_k, β_k).

Assumptions:

- $\alpha, \beta \in \mathbb{F}_{2^m}$ such that $\mathbb{F}_2[\alpha] = \mathbb{F}_{2^m}$ and $\beta \neq 0$,
- Define $w_i = \mathrm{wt}(\alpha^i \beta)$ and $\bar{w}_i = \mathrm{wt}_2(\alpha^i \beta)$,
- $(\bar{w}_i)_{i=0}^{2m-1}$ is the start of an LFSR sequence $\mathcal{W}_{\alpha,\beta}$.

Goal: Recover α and β using only the observed $\bar{w}_i$ (mod 2 leakage).

Recovering α from $\bar{w}_i$

- Apply Berlekamp-Massey to $(\bar{w}_i)_{i=0}^{2m-1}$ to obtain minimal polynomial χ,
- $\chi = \chi_\alpha$ and has m roots: $\alpha^{(0)}, \ldots, \alpha^{(m-1)}$.

These are the m possible candidates for α.

Computing Candidates for β

For each $\alpha^{(\ell)}$ (root of χ):

- Compute change of basis matrix:

$$C_\ell = \left(\mathrm{wt}_2\left((\alpha^{(\ell)})^i x^j\right)\right)_{0 \le i,j < m}$$

- Form the vector $W = (\bar{w}_0, \ldots, \bar{w}_{m-1})^T$
- Compute:

$$\beta^{(\ell)} = C_\ell^{-1} \cdot W$$

This yields a candidate pair $(\alpha^{(\ell)}, \beta^{(\ell)})$ for each ℓ.

Distinguishing the Correct Pair

Although all $(\alpha^{(\ell)}, \beta^{(\ell)})$ yield the same LFSR output, only one of them matches the full Hamming weight sequence:

$$\mathrm{wt}((\alpha^{(\ell)})^i \beta^{(\ell)}) = \mathrm{wt}(\alpha^i \beta) \quad \forall i$$

This test (done for all $i < m$) helps uniquely identify the correct (α, β) pair.

Constructive Algorithm

- Goal: recover t good pairs (α_k, β_k) from noisy weight data.
- The algorithm loops through columns of $\boldsymbol{H}_{\text{wt}} \pmod 2$ to extract the weight sequence $\mathcal{W}$.
- Berlekamp–Massey is used to derive the minimal polynomial.
- Its m roots provide m candidate α's, from which we compute β using a linear system.
- Among these m candidates, often only one satisfies the full (non-mod 2) Hamming weight sequence $\Rightarrow$ succeeds.
- Efficiency

$$\boxed{\mathcal{O}\left(\frac{(n \log_2 n)^2}{n_m}\right)}$$

Heuristic success condition: the candidate pair (α_k, β_k) is unique and compatible with the observed $\text{wt}(\alpha^i \beta)$ sequence.

Theoretical Sufficiency of the Weight Sequence

- **Key Lemma.** If ζ is primitive and α is primitive, then the Hamming weight sequence uniquely identifies (α, β).
- **No Collision.** Under these conditions, there is no other pair (α', β') sharing the same full sequence $(\mathrm{wt}(\alpha^i\beta))_i$.
- **Why This Matters.** $\Rightarrow$ The algorithm frequently succeeds. $\Rightarrow$ Justifies the efficiency of the algorithm, even under noise.

Bottom line: the Hamming weight sequence carries enough information to discriminate candidate pairs early in the algorithm.

Improved Error-Correcting Algorithm (Noisy Setting)

Realistic side-channel context:

- Leakage is noisy in practice.
- Accuracy of Hamming weight distinguishers < 1 (e.g., DPA contest V3).
- Noise modeled by error vector $\mathcal{E}$: $\widetilde{\mathcal{W}} = \mathcal{W} + \mathcal{E}$, with $\varepsilon_{i,j} \in \{-1, 0, +1\}$.

Improved Error-Correcting Algorithm (Noisy Setting)

Realistic side-channel context:

- Leakage is noisy in practice.
- Accuracy of Hamming weight distinguishers < 1 (e.g., DPA contest V3).
- Noise modeled by error vector $\mathcal{E}$: $\widetilde{\mathcal{W}} = \mathcal{W} + \mathcal{E}$, with $\varepsilon_{i,j} \in \{-1, 0, +1\}$.

Key question: Can we recover the BM polynomial from $\widetilde{\mathcal{W}_2} = \mathcal{W}_2 + \mathcal{E}_2$?

Success Probability of Error Correction

Objective: Estimate probability that Algo 1. outputs correct sequence from noisy input $\widetilde{\mathcal{W}}_2 = \mathcal{W}_2 + \mathcal{E}_2$.

Key probabilistic insight:

- Focus on probability that the vector $\boldsymbol{e}$ admits a zero sub-block of length $2m$.
- This corresponds to the successful recovery of the correct sequence.

Success Probability of Error Correction

Objective: Estimate probability that Algo 1. outputs correct sequence from noisy input $\widetilde{\mathcal{W}_2} = \mathcal{W}_2 + \mathcal{E}_2$.

Key probabilistic insight:

- Focus on probability that the vector $\boldsymbol{e}$ admits a zero sub-block of length $2m$.
- This corresponds to the successful recovery of the correct sequence.

Core result:

Lemma 1 (Probability of Zero Block)

Let $\boldsymbol{e} \in \mathbb{F}_2^{2t}$ *with* $\mathrm{wt}(\boldsymbol{e}) = l$. *Then:*

$$\Pr\left[\exists \mathcal{I} \subset [0, 2t-1], |\mathcal{I}| \geq 2m, \boldsymbol{e}_{\mathcal{I}} = \mathbf{0}\right] = \sum_{j=1}^{\left\lfloor \frac{2t-l}{2m} \right\rfloor} \frac{(-1)^{j+1}\binom{l+1}{j}\binom{2t-2mj}{l}}{\binom{2t}{l}}$$

Sequence distance algorithm

Input: $\widetilde{\mathcal{W}_2}$ — noisy mod 2 Hamming weight sequence obtained from SCA
Output: Most probable BM polynomial χ_k and corresponding denoised sequence $\mathcal{W}_2$

$poly\_saved \leftarrow 0$, $min\_error \leftarrow 2t + 1$
for $i \leftarrow 0$ to $2t - 2m$ **do**
    $w \leftarrow \widetilde{\mathcal{W}_2}[i : i + 2m]$
    $poly \leftarrow \mathrm{BM}(w)$
    $Seq \leftarrow \mathrm{LFSR}(poly, w, \text{length} = 2t)$
    $error \leftarrow \mathrm{dist}(Seq, \widetilde{\mathcal{W}_2})$
    if $error < min\_error$ **then**
        $min\_error \leftarrow error$
        $poly\_saved \leftarrow poly$
        $seq\_saved \leftarrow Seq$
return $poly\_saved$, seq_saved

Illustration

Probability of success as a function of accuracy a :

- We have: $\Pr(e_i = 1) = 1 - a$, so $\mathrm{wt}(\boldsymbol{e}) \sim \mathcal{B}(2t, 1 - a)$.

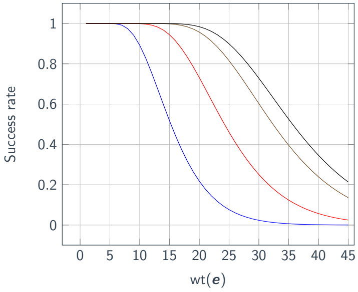

(a) Pr(*success*) in function of wt($\boldsymbol{e}$).

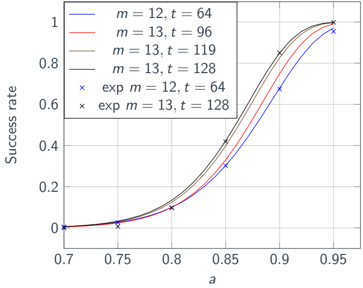

(b) Pr(*success*) in function of a.

Figure: Theoretical probability of success of our Algorithm for all *Classic McEliece* parameters in function of a) wt($\boldsymbol{e}$) and b) accuracy.

Practical Implications

Summary:

- Theoretical analysis matches experiments closely.
- Accurate recovery of BM polynomial enables private key reconstruction.
- Distance-based error correction significantly improves robustness.

Practical Implications

Summary:

- Theoretical analysis matches experiments closely.
- Accurate recovery of BM polynomial enables private key reconstruction.
- Distance-based error correction significantly improves robustness.

Key takeaway: Maintaining classifier accuracy $a \geq 0.74$ suffices to achieve meaningful success rates in realistic noisy settings.

Practical Implications

Summary:

- Theoretical analysis matches experiments closely.
- Accurate recovery of BM polynomial enables private key reconstruction.
- Distance-based error correction significantly improves robustness.

Key takeaway: Maintaining classifier accuracy $a \geq 0.74$ suffices to achieve meaningful success rates in realistic noisy settings.

Applications: Other McEliece variants based on GRS, Alternants are subject to our attack (Vandermonde type matrix).

Source code:
`https://github.com/vingrosso/keyRecoveryClassicMcEliece`

Practical Implications

Summary:

- Theoretical analysis matches experiments closely.
- Accurate recovery of BM polynomial enables private key reconstruction.
- Distance-based error correction significantly improves robustness.

Key takeaway: Maintaining classifier accuracy $a \geq 0.74$ suffices to achieve meaningful success rates in realistic noisy settings.

Applications: Other McEliece variants based on GRS, Alternants are subject to our attack (Vandermonde type matrix).

Source code:
`https://github.com/vingrosso/keyRecoveryClassicMcEliece`

— Questions ? —