CONJURA · OPEN PROBLEM

Soundness of the Hold-Out Test on Dirty Ciphertext Coordinates

The one unproved step in a quasipolynomial-time decryption attack on Goppa–McEliece

Status. Statement: AI-written from Ashrujit Ghoshal, Yuval Ishai, Aayush Jain and Nuozhou Sun, *Quasipolynomial Cryptanalysis of the McEliece Cryptosystem (or: PIR Meets McEliece)*, IACR ePrint 2026/1630. Not yet checked by a human. The distinguishing attack in the same paper is unconditional and proved; the decryption attack is heuristic, and what is unproved is exactly the statement below. Completeness on clean coordinates is proved (the source's Claim 6.2); soundness on dirty ones is supported only by a heuristic dimension count and small-parameter experiments.

Category. *Code-Based Cryptography.*

Conjecture (informal). *The source gives a classical quasipolynomial-time algorithm that distinguishes a Goppa–McEliece public key from a uniformly random matrix, in the asymptotic Classic McEliece regime, and proves it works. It then extends the same idea from distinguishing to decryption: append the ciphertext as an extra row of the public matrix and, for each coordinate, ask whether a certain space of low-degree polynomials that vanish to high multiplicity at all the other columns also vanishes at this one. Coordinates where the error is zero are guaranteed to pass — that half is proved. The attack needs the converse: coordinates where the error is nonzero must fail, or the linear system it solves at the end is inconsistent and the message is not recovered. The source states it has no proof of that, only a heuristic analysis and experiments on small instances. This is that missing statement.*

1 The setting

The McEliece cryptosystem's public key is a scrambled generator matrix $Y = SG_{\mathrm{GRS}}(\alpha, \lambda)P$ of a hidden algebraic code, and its ciphertext is a noisy codeword $c = uY + e$. The source's attack comes

from a PIR-inspired reformulation: the question of whether one column of a matrix lies in the span of the others, asked of a Reed–Muller-encoded matrix, becomes a *hold-out* test that is sensitive to the algebraic structure of GRS and Goppa codes. Concretely, one looks for homogeneous polynomials of degree d whose Hasse jets of order less than s vanish at every column but one; for a genuine GRS-derived matrix, Hermite interpolation forces the composed univariate polynomial to be identically zero, so those polynomials must vanish at the held-out column too, while for a uniformly random matrix they generically do not. The resulting distinguisher (its Algorithm 1, Theorem 3.2) is unconditional and runs in quasipolynomial time in the regime $m = \Theta(\log n)$, $k = \Theta(n)$, $t = \Theta(n/\log n)$.

Its Algorithm 3 turns this into a decryption attack. Append the ciphertext to the public key, forming $A = \binom{Y}{c} \in \mathbb{F}_q^{(k+1)\times n}$, and run the hold-out test at every coordinate τ. Coordinates marked clean are collected into $\widehat{I_{cl}}$, and the message is recovered by solving $uY_{\widehat{I_{cl}}} = c_{\widehat{I_{cl}}}$.

One direction of that filter is proved. The source's Claim 6.2 shows: if $s(|I_{cl}| - 1) > d \cdot D$, where $I_{cl} = [n] \setminus \mathrm{supp}(e)$ and D is the ambient GRS degree bound, then every clean coordinate is marked clean. The argument is the same Hermite interpolation as in the distinguisher, applied to the composed polynomial $\phi(Z) = H(f_1(Z), \ldots, f_k(Z), f'(Z))$ where $f' = \sum_r u_r f_r$.

The other direction is not proved. In the source's words: "For a dirty coordinate $e_\tau \neq 0$, we do not currently have a proof that the hold-out test rejects τ for either code family", and again after Claim 6.2: "To make the linear system consistent, we need $\widehat{I_{cl}}$ to only contain the coordinates with $e_\tau = 0$. We currently do not have a proof to lower bound the probability of this event." What exists instead is Appendix A: a heuristic dimension and genericity analysis, which is also what prescribes the choice of d and s, plus finite-parameter experiments on GRS, alternant, wild-Goppa and q-ary Goppa instances at small extension degree. The source reports it could not run the full attack on a binary square-free Goppa instance satisfying all its conditions because the matrices were too large.

The stakes are set out in the source's own open-problems section: the distinguishing attack is unconditional, but "the decryption algorithm is heuristic", and it asks for "[f]urther validation of

the underlying conjectures, or ideally an unconditional proof of quasipolynomial decryption". Under the hypothesis, resistance to both attack types is governed by $\min\{t, n/t\} \le \sqrt{n}$, which forces $n = \widetilde{\Omega}(\kappa^2)$ and puts Goppa–McEliece under the same qualitative quadratic ciphertext-length tradeoff as Alekhnovich's LPN-based scheme.

2 Notation and parameters

- q is a prime power, $\mathbb{F}_q$ the ambient field. For binary Goppa, $q = 2^m$ and the public matrix is over $\mathbb{F}_2$, viewed over $\mathbb{F}_q$.
- n is the code length, k the public dimension, K the ambient GRS dimension, t the Goppa degree, m the extension degree.
- D is the ambient degree bound: $D = K - 1$ in the GRS case and $D = n - 2t - 1$ in the binary Goppa case.
- d and s are the attack parameters: polynomial degree and multiplicity order. The regime of interest is $m = \Theta(\log n)$, $k = \Theta(n)$, $t = \Theta(n/\log n)$ with $s, d = \Theta(\log n)$.
- $\mathcal{A}_{<s,\ell} := \{a \in \mathbb{Z}_{\ge 0}^{\ell} : \sum_i a_i < s\}$ indexes Hasse derivatives of order less than s in ℓ variables, and $\partial^{[a]}$ is the Hasse derivative with respect to a.
- $\mathcal{P}^{\mathrm{hom}}_{=d}$ is the space of $(k+1)$-variate polynomials over $\mathbb{F}_q$ homogeneous of degree exactly d.
- A_j is the j-th column of A; $w := |\mathrm{supp}(e)|$; $I_{cl} := [n] \setminus \mathrm{supp}(e)$ is the set of *clean* coordinates, and a coordinate τ with $e_\tau \neq 0$ is *dirty*.

3 The conjecture

Definition 1 (The homogeneous hold-out kernel, [GIJS26, Algorithm 3]). For $A \in \mathbb{F}_q^{(k+1)\times n}$ and $\tau \in [n]$, put

$$K^{\mathrm{hom}}_{\tau}(A) := \left\{ H \in \mathcal{P}^{\mathrm{hom}}_{=d} \;:\; (\partial^{[a]} H)(A_j) = 0 \;\; \forall j \in [n]\setminus\{\tau\},\, \forall a \in \mathcal{A}_{<s,k+1} \right\}.$$

The test *marks* τ *clean* if $H(A_\tau) = 0$ for every $H \in K^{\mathrm{hom}}_{\tau}(A)$, and $\widehat{I}_{cl}$ is the set of coordinates it marks clean.

Definition 2 (The instance distribution). Fix the binary Goppa parameters (m, n, t) and let $Y \in \mathbb{F}_2^{k\times n}$ be drawn from the McEliece public-key distribution of [GIJS26, Definition 2.4] — $Y = SGP$ for a generator matrix G of Goppa(α, Γ), a uniform full-rank S and a uniform permutation matrix P. Let u be a uniform message, let e be drawn from the scheme's error distribution, and let $c = uY + e$ and $A = \binom{Y}{c}$. The same construction with $\mathrm{GRS}_K(\alpha, \lambda)$ in place of Goppa(α, Γ) gives the GRS instance distribution.

Conjecture 1 (Hold-out soundness on dirty coordinates). *In the asymptotic Classic McEliece regime $m = \Theta(\log n)$, $k = \Theta(n)$, $t = \Theta(n/\log n)$, there are attack parameters $s, d = \Theta(\log n)$ satisfying the completeness condition $s(|I_{cl}| - 1) > d \cdot D$ such that, over the instance distribution of Definition 2, with probability $1 - o(1)$ every dirty coordinate is* not *marked clean: for every τ with $e_\tau \neq 0$ there exists $H \in K_\tau^{\hom}(A)$ with $H(A_\tau) \neq 0$.*

Remark 1 (What Conjecture 1 buys). Together with the source's Claim 6.2, which is proved, Conjecture 1 gives $\widehat{I_{cl}} = I_{cl}$ with probability $1 - o(1)$, so the linear system $uY_{\widehat{I_{cl}}} = c_{\widehat{I_{cl}}}$ is consistent and (given $\mathrm{rank}(Y_{I_{cl}}) = k$) determines u. That makes the source's Algorithm 3 an unconditional quasipolynomial-time decryption attack on Goppa–McEliece in the stated regime, which is what its open-problems section asks for. Nothing about the source's *distinguishing* attack depends on this statement; that one is already proved.

Remark 2 (The one-sided variant, and why it does not escape). The source's Remark 6.3 gives, for binary Goppa, a modified filter that avoids needing to reject every dirty coordinate: retain τ when $H(A_\tau + \xi) \neq 0$ for some $H \in K_\tau^{\hom}(A)$, where $\xi = (0, \ldots, 0, 1)^\top$. Since $A_\tau + \xi$ is the clean point on the hidden curve when τ is dirty, the interpolation argument of Claim 6.2 shows no dirty coordinate is retained — so that side becomes unconditional. But the guarantee needed then flips: enough *clean* coordinates must survive, i.e. for a clean τ the shifted point must behave like a dirty one, and the source says the same Appendix A heuristics are what suggest that. So the variant relocates the unproved step rather than removing it, and a resolution of Conjecture 1 in either formulation is what

closes the gap.

Remark 3 (Both directions are of interest). A proof settles the statement. So does a refutation: if for the source's parameter choices dirty coordinates are marked clean with non-negligible probability, the decryption attack as specified fails, and the source's own conditional conclusion about the quadratic ciphertext barrier for Goppa–McEliece would no longer follow from it (the unconditional distinguishing attack, and the barrier it implies for public-key pseudorandomness, would stand). A refutation for a specific (d, s) that still leaves other parameter choices open would be a partial result.

Remark 4 (What the source's experiments do and do not cover). Appendix A.4 of the source reports dimension and genericity experiments on small instances of GRS, alternant, "wild Goppa" and q-ary Goppa families over small extension degree, and reports that the normalized dirty points behave as the heuristics predict. It also reports: "We were unable to run the full attack on a binary square-free Goppa instance satisfying all of" its conditions, the matrices being too large. So the family that Classic McEliece actually uses, at parameters where the completeness condition holds, is exactly the case with no end-to-end experimental support — and a resolution of Conjecture 1 for binary Goppa is what the source is missing, not for GRS.

4 Bibliography

[GIJS26] Ashrujit Ghoshal, Yuval Ishai, Aayush Jain and Nuozhou Sun. *Quasipolynomial cryptanalysis of the McEliece cryptosystem (or: PIR meets McEliece).* IACR ePrint 2026/1630. The source. Algorithm 1 and Theorem 3.2 are on page 22; the heuristic decryption attack is Section 1.1.2 (page 11) and Algorithm 3 with Claim 6.2 and Remark 6.3 (pages 37–39); the open problems are Section 1.2, pages 11–12; the heuristic analysis is Appendix A.

[McE78] Robert J. McEliece. *A public-key cryptosystem based on algebraic coding theory.* DSN Progress Report 42–44, Jet Propulsion Laboratory, California Institute of Technology, 1978. The cryptosystem under attack.

[Gop70] Valerii Denisovich Goppa. *A new class of linear correcting codes.* Problemy Peredachi Informatsii, 6(3):24–30, 1970. The code family Classic McEliece uses.

[KSY14] Swastik Kopparty, Shubhangi Saraf and Sergey Yekhanin. *High-rate codes with sublinear-time decoding.* Journal of the ACM, 61(5):1–20, 2014. The multiplicity codes whose Hasse-jet symbols the hold-out test is built on.

[Ale03] Michael Alekhnovich. *More on average case vs approximation complexity.* 44th Annual Symposium on Foundations of Computer Science (FOCS), 2003, pages 298–307. The LPN-based public-key encryption whose quadratic ciphertext length is the benchmark the source's conditional conclusion puts Goppa–McEliece alongside.

[LMW23] Wei-Kai Lin, Ethan Mook and Daniel Wichs. *Doubly efficient private information retrieval and fully homomorphic RAM computation from ring LWE.* 55th ACM Symposium on Theory of Computing (STOC), 2023, pages 595–608. The doubly efficient PIR result the source's PIR-inspired framework grew out of.