ANDURIL

An Algebraic-Geometry Lower Bound against the ePrint:2026/1747 McEliece Key-Recovery Attack

Daniel Apon
Anduril Industries
dapon@anduril.com

August 26, 2026

Abstract

A few weeks ago, Ghoshal, Ishai, Jain, and Sun (ePrint:2026/1630) introduced a "hold-out distinguisher" for the Goppa–McEliece public key. This past week, Vedenev (eprint:2026/1747) proposed to turn its polynomial relations into key recovery by reconstructing the hidden generalized Reed–Solomon representation from nested Hasse-derivative spaces at held positions. Vedenev's proposed held-position count explicitly assumes that the resulting linear equations are independent across positions. Yet, experiments on proper binary Goppa instances contradict that assumption, demonstrating a familiar "waterfall" phenomenon where – *just* before the required independent equation count for a successful attack – additional held positions sharply drop in value, providing just a *single*, independent equation rather than the $\approx \binom{k}{2}$ such equations from the early positions.

This note identifies an algebraic-geometric reason for this inherent dependence. For a binary Goppa polynomial of degree t, an explicit linear map constructs a $2t+3$-dimensional family modulo the true solution. At each held support point, the entire derivative-flag block restricts on this family to at most one ordinary evaluation condition. Consequently, under a concrete nondegeneracy condition stated in terms of the hidden vector polynomial $\mathbf{F}$ and its formal derivative $\mathbf{F}'$,

$$c_{\mathsf{need}} \geq 2t + 3,$$

where c_{need} is the number of sampled held positions required at the critical step in Vedenev's algorithm. (The proposed key-recovery algorithm's cost depends on c_{need} in the exponent.)

For ISO/NIST Category 5 parameter set `mceliece8192128`, this gives $c_{\mathsf{need}} \geq 259$, which implies Vedenev's key-recovery algorithm costs in excess of 2^{1500} bit operations there.

Introduction

In the past year or so, there has been a tidal wave of novel cryptanalyses targeted at the Classic McEliece cryptosystem [1, 2, 3, 4, 5, 6, 7, 8, 9], which was recently standardized in ISO/IEC 18033-2:2006/Amd 2:2026.[1]

We begin by separating the notions of the (seemingly valid) public-key hold-out distinguisher of Ghoshal–Ishai–Jain–Sun (GIJS) [6] and Vedenev's proposed key-recovery algorithm [9]. The former tests a public matrix for a forced polynomial relation. The latter attempts to recover the hidden polynomial evaluation map from Hasse-derivative data derived from those relations.

Vedenev's algorithm attempts to convert the machinery of [6] into a certain linear reconstruction problem. But critically, Vedenev counts the equations displayed at one held position and assumes independence between different such positions. The purpose of this note is to analyze that independence assumption. We believe it does not hold – and strongly not so.

Our arguments against this independence assumption are derived from algebraic-geometry intuition. We have made an effort, however, to excise as much of the geometry language from this note, in an effort to make it more easy to parse for the typical code-based cryptography audience.[2]

In any case, Section 1 recites the GIJS distinguisher, the recovered local data, and the reconstruction problem posed by Vedenev. Section 2 reports empirical finite-field experiments, which gave early evidence that something was amiss. Section 3 gives a self-contained proof that at least $2t + 3$ held support points are necessary before the derivative-flag system can isolate the true polynomial vector up to a scalar. Concretely, this means that Vedenev's algorithm as written takes well beyond 2^{1500} bit-operations to recover keys from the ISO/NIST Category 5 parameter sets of Classic McEliece.

Our main result is an unconditional obstruction to the Step 3 reconstruction in Vedenev's algorithm. We do not claim that every possible use of the hold-out kernels of GIJS's PK-distinguishing algorithm fails to provide a useful/novel key-recovery algorithm for Classic McEliece.

1 The hold-out distinguisher and Vedenev's attack

1.1 The hidden polynomial representation

Let $\mathbb{F}_2$ be the binary field, let $m \geq 1$, and let $E = \mathbb{F}_{2^m}$. Let n and t be positive integers with $n > 2t$. Let Z be a polynomial indeterminate over E; thus $E[Z]$ is the univariate polynomial ring over E. Let

$$L = (a_1, \ldots, a_n) \in E^n$$

have pairwise distinct entries, and let $G \in E[Z]$ be an irreducible polynomial of degree t with $G(a_i) \neq 0$ for every i. Let

$$\Pi(Z) = \prod_{i=1}^{n} (Z - a_i), \qquad D = n - 2t - 1.$$

The binary Goppa code with support $L \in \mathbb{F}_2^n$ and Goppa polynomial $G \in E[Z]$ is

$$\Gamma(L, G) = \left\{ c = (c_1, \ldots, c_n) \in \mathbb{F}_2^n : \sum_{i=1}^{n} \frac{c_i}{Z - a_i} \equiv 0 \pmod{G(Z)} \right\}.$$

[1] The author pleads innocence in the matter of the paywall gate to that standard's text.

[2] No one benefits from reading the author wax informally about infinitesimally small parameterizations of a differential-deformation operator, where he is probably so imprecise in description as to be easily incorrect there.

The congruence is evaluated in $E[Z]/(G)$, where every $Z - a_i$ is invertible because $G(a_i) \neq 0$. For completeness, a binary square-free Goppa code satisfies

$$\Gamma(L, G) = \Gamma(L, G^2).$$

For $1 \leq r < n$ and $v = (v_1, \ldots, v_n) \in (E^\times)^n$, where $E^\times = E \setminus \{0\}$, define

$$\mathsf{GRS}_r(L, v) = \left\{ \big(v_1 f(a_1), \ldots, v_n f(a_n)\big) : f \in E[Z],\ \deg f < r \right\}.$$

With respect to the standard Euclidean inner product on E^n, the standard GRS duality formula is

$$\mathsf{GRS}_r(L, v)^\perp = \mathsf{GRS}_{n-r}(L, v^\perp), \qquad v_i^\perp = \big(v_i \Pi'(a_i)\big)^{-1}.$$

A common nonzero scaling of all entries of $v^\perp$ gives the same GRS code. Here, for $P(Z) = \sum_j p_j Z^j$, the prime denotes the ordinary formal derivative

$$P'(Z) = \sum_{j \geq 1} j p_j Z^{j-1};$$

in particular, Π' is the formal derivative of Π.

The code $\Gamma(L, G^2)$ is the binary subfield subcode, meaning the intersection with $\mathbb{F}_2^n$, of the dual of $\mathsf{GRS}_{2t}(L, (G(a_i)^{-2})_{i=1}^n)$. The duality formula therefore embeds it in an ambient GRS code of dimension $n - 2t = D + 1$, with multiplier vector $\lambda = (\lambda_1, \ldots, \lambda_n)$, up to a common nonzero scaling,

$$\lambda_i = \kappa \, \frac{G(a_i)^2}{\Pi'(a_i)} \qquad (\kappa \in E^\times). \tag{1}$$

For a linear code $C \subseteq E^N$ and a set $J \subseteq \{1, \ldots, N\}$ of coordinate indices, *shortening on* J means restricting to codewords that are zero in every coordinate in J, and then deleting those coordinates. The multiplier form in (1) is preserved by shortening. To see this, let J_L be the corresponding set of removed support values and let

$$P_J(Z) = \prod_{b \in J_L} (Z - b), \qquad \Pi_0(Z) = \Pi(Z) P_J(Z),$$

so Π_0 is the support polynomial before shortening and Π is the support polynomial after shortening. A polynomial representing a shortened GRS codeword is divisible by P_J. At a retained point a,

$$\Pi_0'(a) = \Pi'(a) P_J(a),$$

where Π_0' is the formal derivative of Π_0. Multiplying the old GRS multiplier by the shortening factor $P_J(a)$ gives

$$\kappa \frac{G(a)^2}{\Pi_0'(a)} P_J(a) = \kappa \frac{G(a)^2}{\Pi'(a)}.$$

Let $Y = [y_1 | \cdots | y_n] \in \mathbb{F}_2^{k \times n}$ be a generator matrix of the retained public code, so $y_i \in \mathbb{F}_2^k$ is its i-th column and k is the dimension of that code. Write

$$E[Z]_{\leq D} = \{ f \in E[Z] : \deg f \leq D \}.$$

There are linearly independent polynomials $f_1, \ldots, f_k \in E[Z]_{\leq D}$ such that, with

$$\mathbf{F} = (f_1, \ldots, f_k)^T,$$

the public columns satisfy

$$y_i = \lambda_i \mathbf{F}(a_i) \in \mathbb{F}_2^k. \tag{2}$$

Thus every public column is a nonzero scalar multiple of an evaluation of the same hidden vector polynomial $\mathbf{F}$.

1.2 Hasse derivatives and derivative flags

Retain the polynomial indeterminate Z from the preceding subsection, and let U be a second independent formal indeterminate over E. For a univariate polynomial $f \in E[Z]$, its j-th Hasse derivative $\partial^{[j]} f$ is defined by

$$f(Z+U) = \sum_{j \geq 0} (\partial^{[j]} f)(Z) U^j.$$

Equivalently,

$$\partial^{[j]} Z^r = \binom{r}{j} Z^{r-j},$$

where the binomial coefficient is reduced in the field. For a multivariate polynomial, let

$$\boldsymbol{X} = (X_1, \ldots, X_k), \qquad \boldsymbol{U} = (U_1, \ldots, U_k),$$

where the entries of $\boldsymbol{U}$ are independent formal indeterminates. For $h \in E[X_1, \ldots, X_k]$ and a multi-index $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_k) \in \mathbb{Z}^k_{\geq 0}$, let $|\boldsymbol{\nu}| = \sum_{\ell=1}^k \nu_\ell$ and $\boldsymbol{U}^{\boldsymbol{\nu}} = \prod_{\ell=1}^k U_\ell^{\nu_\ell}$. The Hasse derivative $\partial^{[\boldsymbol{\nu}]} h$ is the coefficient of $\boldsymbol{U}^{\boldsymbol{\nu}}$ in $h(\boldsymbol{X} + \boldsymbol{U})$.

For a scalar polynomial f, the order-r *Hasse derivative data* at a, also called an order-r Hasse jet, is the tuple

$$\big(f(a), (\partial^{[1]} f)(a), \ldots, (\partial^{[r]} f)(a)\big).$$

This tuple is exactly the list of coefficients of $1, U, \ldots, U^r$ in $f(a+U)$. Hasse derivatives of a polynomial vector are taken coordinatewise. For $\mathbf{F}$, define the nested derivative spaces

$$V_a^{[j]} = \mathrm{span}_E\{\mathbf{F}(a), (\partial^{[1]}\mathbf{F})(a), \ldots, (\partial^{[j]}\mathbf{F})(a)\}. \tag{3}$$

The nested sequence $V_a^{[0]} \subseteq V_a^{[1]} \subseteq \cdots$ is called the *osculating flag* in Vedenev's terminology. We will call it the *derivative flag*. Unlike the ordered derivative data, the flag records only successive spans and discards the individual derivative vectors.

1.3 The GIJS hold-out distinguisher

Fix a set $H \subseteq \{1, \ldots, n\}$ of held coordinates, a space $\mathcal{P}$ of homogeneous degree-d polynomials in k variables, and an integer $s \geq 1$. The hold-out kernel is

$$\mathcal{K}_H^{\mathcal{P}}(Y) = \Big\{ h \in \mathcal{P} : (\partial^{[\boldsymbol{\nu}]} h)(y_j) = 0 \text{ for } j \notin H,\ |\boldsymbol{\nu}| < s \Big\}. \tag{4}$$

In the binary construction, $\mathcal{P}$ is a space of multilinear homogeneous polynomials. The constraints in (4) are linear in the coefficients of h. GIJS call the resulting test the *binary homogeneous all-kernel hold-out distinguisher*.

For $h \in \mathcal{K}_H^{\mathcal{P}}(Y)$, consider the univariate composition

$$\phi_h(Z) = h(\mathbf{F}(Z)).$$

It has degree at most dD. The non-held constraints force a zero of multiplicity at least s at each a_j, $j \notin H$. Therefore

$$s(n - |H|) > dD \quad \Rightarrow \quad \phi_h = 0. \tag{5}$$

In particular, $h(y_i) = 0$ at every held position. In the random-matrix experiment, the held columns remain independent of the kernel computed from the other columns, so the same vanishing test rejects with nonzero probability. GIJS use suitable parameters, shortening, and sparse linear algebra to obtain their public-key distinguisher.

1.4 Vedenev's proposed conversion to key recovery

Vedenev differentiates the identities $h(\mathbf{F}(Z)) = 0$ for $h \in \mathcal{K}_i$, where $\mathcal{K}_i$ denotes the hold-out kernel used to recover local data at coordinate i. The resulting equations are intended to recover the derivative flags $V_{a_i}^{[j]}$. For subfield subcodes, Frobenius conjugates produce several local branches; the proposal uses quadratic equations to separate them. After local recovery, the global procedure is:

1. compute hold-out kernels at c positions;
2. recover a derivative flag at each position;
3. normalize two support values, guess the remaining held support values and relative Frobenius branches, and reconstruct $\mathbf{F}$ by linear algebra;
4. match all public columns against scalar multiples of $\mathbf{F}(a)$ to recover the remaining support and multipliers;
5. restore punctured coordinates if the distinguisher used shortening.

The present note concerns Step 3 even when Steps 1–2 are replaced by a perfect oracle. For a held point a_i, let r_i be the first order at which the recovered flag fills E^k. The candidate vector polynomial $\mathbf{H} \in E[Z]_{\le D}^k$ is required to satisfy

$$(\partial^{[j]}\mathbf{H})(a_i) \in V_{a_i}^{[j]}, \qquad 0 \le j < r_i. \tag{6}$$

In the full-growth case $r_i = k - 1$, one position displays

$$\sum_{j=0}^{k-2}(k-j-1) = \binom{k}{2}$$

scalar linear equations. Vedenev therefore proposes

$$c_{\mathsf{V}} = \left\lceil \frac{k(D+1)-1}{\binom{k}{2}} \right\rceil, \tag{7}$$

explicitly assuming that equations from distinct positions are *independent*.[3]

For a set S of correctly guessed support points, define

$$\mathrm{Sol}(S) = \left\{ \mathbf{H} \in E[Z]_{\le D}^k : \mathbf{H} \text{ satisfies (6) at every } a \in S \right\}. \tag{8}$$

The true polynomial vector gives the line $E\mathbf{F} = \{\alpha\mathbf{F} : \alpha \in E\} \subseteq \mathrm{Sol}(S)$. For an ordered list of correctly guessed, distinct held support points, we write c_{need} for the least prefix length at which $\mathrm{Sol}(S) = E\mathbf{F}$. The theorem below gives the same lower bound for every ordering and, more strongly, every subset.

Structural information-loss check. Passing from the ordered list

$$\mathbf{F}(a), \partial^{[1]}\mathbf{F}(a), \ldots$$

to the nested spaces $V_a^{[j]}$ is a many-to-one compression. It retains membership in spans but not the derivative vectors themselves. The theorem below constructs distinct polynomial vectors that all have zero syndrome for the same supplied Step 3 flag constraints. If only these syndromes are retained

[3] The notation c_V here means "Vedenev's estimate of c," and it is the explicit "calculation error" of this note's concern.

as a classical transcript, the corresponding diagonal states are identical and have trace distance zero. Postselection, restart, and amplification applied only to that transcript cannot select the true candidate. Step 4 can add new information by testing candidates against the full public code, but then the algorithm must solve a nontrivial list problem rather than the claimed one-dimensional linear reconstruction.

2 Exact experiments contradict cross-position independence

We formed the system (6) using the true support and the true derivative flags. Thus the experiment grants the attack a perfect oracle for the local-recovery and Frobenius-alignment steps and isolates only the rank of Step 3. Each test instance used:

- an irreducible degree-t binary Goppa polynomial, verified by a Rabin irreducibility test;
- the expected binary-Goppa and shortened dimensions;
- an exactly recovered ambient GRS multiplier;
- exact agreement with (2) at every retained column; and
- a polynomial vector **F** whose coordinate greatest common divisor is one and whose maximum coordinate degree is exactly D, excluding a removable common polynomial factor.

Table 1 reports the corrected audit. Here k and D are the parameters after shortening, c_{V} is (7), and "nullity at c_{V}" is $\dim_E \mathrm{Sol}(S)$ after that many held positions. A one-dimensional nullity is required for the stated reconstruction.

Table 1: Exact Step 3 ranks on proper irreducible binary Goppa instances.

m	t	k	D	c_{V}	nullity at c_{V}	measured c_{need}
8	3	10	27	7	3	9
8	4	14	37	6	12	11
8	5	10	39	9	5	13
8	5	16	45	7	14	13
8	5	22	51	5	44	13
8	6	10	45	11	5	15
8	7	10	51	12	6	17
8	8	10	57	13	7	19
8	10	10	69	16	8	23
8	12	10	81	19	9	27
9	7	10	58	14	4	17
10	10	10	89	20	4	23

The formula-selected system has nullity greater than one in every row. Note similarly that $c_V < c_{\mathrm{need}}$ in every row.

The failure can be large: for $(m, t, k, D) = (8, 5, 22, 51)$, the proposed count is five positions and the resulting nullity is 44. The per-position equation count itself is not the problem. Initially, many instances attain the full local increment $\binom{k}{2}$. Near the end, however, the increments collapse and commonly become one per new position. This is the familiar "waterfall" phenomenon. The equations displayed at different positions are therefore not independent.

All twelve instances satisfy the empirical equality

$$c_{\mathrm{need}} = 2t + 3. \tag{9}$$

Ranks and nullspaces were computed by exact arithmetic over $E = \mathbb{F}_{2^m}$, not by floating-point numerical linear algebra. The reported equality is for one deterministic shuffled ordering in each row; it is not an all-subsets experiment. The sample is also not a collection of twelve independent random trials: the three $t = 5$ rows use different shortenings of the same underlying Goppa instance.

A separate algebraic audit checked, by exact finite-field arithmetic, the differential identity and the auxiliary map proved below. It also verified that the candidate family has dimension $2t + 3$ modulo the true solution, that each held-position block obeys the evaluation law of Lemma 7, and that the complete Step 3 kernel is one-dimensional at $c = 2t + 3$ in all twelve instances.

These computations do not prove the upper bound in (9), nor do they by themselves justify extrapolation from small m, t to standardized parameters. The lower bound in Section 3 is different: it is algebraic, does not depend on sampling or a favorable ordering, and applies to every set of at most $2t + 2$ distinct held support points under its stated hypotheses.

3 An algebraic proof that $c_{\mathsf{need}} \geq 2t + 3$

3.1 Statement of the theorem

We retain the notation of Section 1. The derivative $\mathbf{F}'$ below is the ordinary formal derivative, applied coordinatewise. For $1 \leq p < q \leq k$, define

$$\Delta_{p,q} = f_p f_q' - f_q f_p'. \tag{10}$$

Because the field has characteristic two, subtraction and addition are the same operation. The proof assumes that

$$\Delta_{p,q} \neq 0 \quad \text{for at least one pair } p < q. \tag{11}$$

Equivalently, $\mathbf{F}$ and $\mathbf{F}'$ are not proportional over the rational function field $E(Z)$. This is the only nondegeneracy hypothesis used below.

Theorem 1 (Algebraic lower bound). *Let E be a field of characteristic two. Let L, Π, G, t, D, λ_i, and $\mathbf{F} \in E[Z]^k_{\leq D}$ satisfy* (1)–(2)*, and assume $\Delta_{p,q} \neq 0$ for at least one pair $p < q$. For every set $S \subseteq \{a_1, \ldots, a_n\}$ of $c \leq 2t + 2$ distinct support points,*

$$\dim_E \mathsf{Sol}(S) \geq 2t + 4 - c.$$

In particular,

$$c_{\mathsf{need}} \geq 2t + 3.$$

We note the conclusion is independent of the choice and ordering of S.

The proof has three parts. First, binary-valued public columns imply a differential identity for $\mathbf{F}$. Second, that identity is used to define an explicit linear map from $E[Z]_{\leq 2t+2}$ into the Step 3 unknown space $E[Z]^k_{\leq D}$. Third, each held support point imposes at most the single scalar condition $A(a) = 0$ on the parameter polynomial A. Counting the remaining parameter polynomials proves the theorem.

3.2 The binary differential identity

Lemma 2 (Partial-fraction interpolation). *For each coordinate f of $\mathbf{F}$, let $y_i = \lambda_i f(a_i) \in \mathbb{F}_2$. Then*

$$\sum_{i=1}^{n} \frac{y_i}{Z - a_i} = \kappa \, \frac{G(Z)^2 f(Z)}{\Pi(Z)}. \tag{12}$$

Proof. After multiplication by Π, the claimed identity becomes

$$\sum_{i=1}^{n} y_i \frac{\Pi(Z)}{Z - a_i} = \kappa G(Z)^2 f(Z). \tag{13}$$

Both sides are polynomials of degree at most $n-1$, because

$$\deg(G^2 f) \le 2t + D = n - 1,$$

and each quotient $\Pi(Z)/(Z-a_i)$ has degree $n-1$. At $Z = a_j$, the left side equals $y_j \Pi'(a_j)$, while the right side equals

$$\kappa G(a_j)^2 f(a_j) = \lambda_j f(a_j) \Pi'(a_j) = y_j \Pi'(a_j).$$

The two degree-at-most-$n-1$ polynomials agree at the n distinct points $a_1, \ldots, a_n$, so they are equal. Dividing by Π gives the stated rational-function identity. □

Lemma 3 (Binary differential identity)**.** *Every coordinate f of* **F** *satisfies*

$$\Pi f' + \Pi' f = \kappa G^2 f^2. \tag{14}$$

Equivalently,

$$\Pi \mathbf{F}' + \Pi' \mathbf{F} = \kappa G^2 \mathbf{F}^{(2)}, \tag{15}$$

where $\mathbf{F}^{(2)}$ denotes coordinatewise squaring.

Proof. Let $S_f(Z)$ denote the left side of (12). Since $y_i \in \mathbb{F}_2$, one has $y_i^2 = y_i$. In characteristic two,

$$S_f' = \sum_i \frac{y_i}{(Z - a_i)^2} = S_f^2.$$

Differentiate and square the right side of (12), equate the results, and multiply by Π^2:

$$\kappa G^2 (\Pi f' + \Pi' f) = \kappa^2 G^4 f^2.$$

Division by the nonzero polynomial κG^2 gives (14). Applying the identity to all coordinates gives (15). □

The intermediate equality $S_f' = S_f^2$ is sometimes called a binary Riccati identity. Here it is only the displayed calculation; no result from the theory of Riccati equations is used.

Because G and Π have no common root, their greatest common divisor is one, so Π is invertible modulo G^2. Reducing (15) modulo G^2 gives

$$\mathbf{F}' \equiv R_G \mathbf{F} \pmod{G^2}, \qquad R_G \equiv \frac{\Pi'}{\Pi} \pmod{G^2}, \quad \deg R_G < 2t. \tag{16}$$

Here R_G is the unique degree-less-than-$2t$ representative of the residue class $\Pi' \Pi^{-1}$ in $E[Z]/(G^2)$. The plus sign is also the minus sign in characteristic two.

3.3 Construction of the auxiliary map

Let

$$\mathcal{A} = E[Z]_{\le 2t+2}.$$

This parameter space has dimension $2t+3$. The upper degree $2t+2$ is chosen because, for $A \in \mathcal{A}$, the quotient defined below has degree at most $D+1$, and its possible degree-$D+1$ term can be cancelled by adding a scalar-polynomial multiple of $\mathbf{F}$.

Definition 4 (Auxiliary map)**.** For $A \in E[Z]_{\le 2t+2}$, let B_A be the unique polynomial of degree less than $2t$ satisfying

$$B_A \equiv AR_G \pmod{G^2}. \tag{17}$$

Define

$$\mathbf{Q}_A = \frac{A\mathbf{F}' + B_A\mathbf{F}}{G^2}. \tag{18}$$

For a polynomial P, let $[Z^\ell]P$ denote the coefficient of Z^ℓ in P. Let $\gamma = \mathrm{lc}(G)$ be the leading coefficient of G, and let $D \bmod 2 \in \mathbb{F}_2 \subseteq E$ denote the parity of D. Let

$$\rho_A = [Z^{2t+2}]A\,(D \bmod 2)\,\gamma^{-2}. \tag{19}$$

Define the E-linear map

$$\mathcal{T} : \mathcal{A} \longrightarrow E[Z]^k_{\le D}, \qquad \mathcal{T}(A) = \mathbf{Q}_A + \rho_A Z\mathbf{F}. \tag{20}$$

For a polynomial vector, deg denotes the maximum of the degrees of its coordinates.

Lemma 5 (Polynomiality and degree)**.** *The numerator in* (18) *is divisible coordinatewise by G^2, and* $\deg \mathcal{T}(A) \le D$*.*

Proof. By (16) and (17),

$$A\mathbf{F}' + B_A\mathbf{F} \equiv AR_G\mathbf{F} + AR_G\mathbf{F} = 0 \pmod{G^2}.$$

Thus $\mathbf{Q}_A$ is a polynomial vector. Its numerator has degree at most

$$\max\{(2t+2)+(D-1), (2t-1)+D\} = D+2t+1,$$

so $\deg \mathbf{Q}_A \le D+1$.

Only the term $A\mathbf{F}'$ can contribute in degree $D+2t+1$. Writing $\mathbf{F}_D = [Z^D]\mathbf{F}$, its contribution after division by the leading term $\gamma^2 Z^{2t}$ of G^2 is

$$[Z^{D+1}]\mathbf{Q}_A = [Z^{2t+2}]A\,(D \bmod 2)\,\gamma^{-2}\mathbf{F}_D = \rho_A\mathbf{F}_D.$$

The coefficient of Z^{D+1} in $\rho_A Z\mathbf{F}$ is the same. Their sum is zero in characteristic two. Hence $\deg \mathcal{T}(A) \le D$. □

Thus every part of $\mathcal{T}(A)$ is determined by polynomial arithmetic. The remainder B_A makes the numerator divisible by G^2, and $\rho_A Z\mathbf{F}$ cancels the only coefficient that can exceed the Step 3 degree bound.

3.4 The local Hasse-derivative calculation

Fix a support point a, let $u = Z - a$, and let

$$V_a^{[j]} = \mathrm{span}_E\{(\partial^{[q]}\mathbf{F})(a) : 0 \le q \le j\}.$$

For any integer $r \ge 1$, define the local Step 3 lattice

$$\mathcal{M}_{a,r} = \left\{\mathbf{H} \in E[[u]]^k : (\partial^{[j]}\mathbf{H})(a) \in V_a^{[j]} \text{ for } 0 \le j < r\right\}. \tag{21}$$

Here $E[[u]]$ is the ring of formal power series in u. A rational function whose denominator is nonzero at a has a unique expansion in this ring, so in particular $G^{-2} \in E[[u]]$.

This is an $E[[u]]$-submodule. Indeed, for $\psi \in E[[u]]$, the Hasse product rule gives

$$\partial^{[j]}(\psi\mathbf{H})(a) = \sum_{q=0}^{j} (\partial^{[q]}\psi)(a)(\partial^{[j-q]}\mathbf{H})(a),$$

and every summand belongs to $V_a^{[j-q]} \subseteq V_a^{[j]}$. In particular, every local scalar multiple of $\mathbf{F}$ lies in $\mathcal{M}_{a,r}$, and multiplication or division by a local unit preserves this lattice. Here a *local unit* is a series in $E[[u]]$ with nonzero constant coefficient.

For a polynomial or power-series vector $\mathbf{H}$, let

$$[\mathbf{H}]_{a,r} = \mathbf{H} + \mathcal{M}_{a,r}$$

denote its coset in $E[[u]]^k/\mathcal{M}_{a,r}$, after expanding $\mathbf{H}$ in the local variable $u = Z - a$.

Lemma 6 (Key derivative calculation)**.** *If $A(a) = 0$, then*

$$A\mathbf{F}' \in \mathcal{M}_{a,r}$$

for every $r \ge 1$.

Proof. We verify the condition order by order. For $0 \le j < r$, the Hasse product rule and $A(a) = 0$ give

$$(\partial^{[j]}(A\mathbf{F}'))(a) = \sum_{q=1}^{j} (\partial^{[q]}A)(a)(\partial^{[j-q]}\mathbf{F}')(a). \tag{22}$$

Ordinary and Hasse derivatives satisfy

$$\partial^{[\ell]}\mathbf{F}' = (\ell+1)\partial^{[\ell+1]}\mathbf{F}. \tag{23}$$

For the summand indexed by $q \ge 1$, the derivative of $\mathbf{F}$ in (23) has order

$$j - q + 1 \le j.$$

The summand therefore belongs to $V_a^{[j]}$; if the scalar $j - q + 1$ is even, it is simply zero in characteristic two. Equation (22) consequently lies in $V_a^{[j]}$ for every $j < r$. □

This is the iterative step in the proof: the order-j constraint on $A\mathbf{F}'$ reduces to derivatives of $\mathbf{F}$ of order at most j. No independence or generic-rank assumption is used.

Lemma 7 (One held point is at most one evaluation condition)**.** *For every support point a, every $r \geq 1$, and every $A \in E[Z]_{\leq 2t+2}$,*

$$[\mathcal{T}(A)]_{a,r} = A(a)[\mathcal{T}(1)]_{a,r} \quad \textit{in} \quad E[[u]]^k/\mathcal{M}_{a,r}. \tag{24}$$

In particular,

$$A(a) = 0 \quad \Longrightarrow \quad \mathcal{T}(A) \in \mathcal{M}_{a,r}.$$

Thus the entire held-position block has rank at most one on the family $\mathcal{T}(\mathcal{A})$.

Proof. First suppose $A(a) = 0$. Lemma 6 gives $A\mathbf{F}' \in \mathcal{M}_{a,r}$. Since $G(a) \neq 0$, the function G^{-2} is a unit in $E[[Z - a]]$, so

$$G^{-2}A\mathbf{F}' \in \mathcal{M}_{a,r}.$$

The remaining terms

$$G^{-2}B_A\mathbf{F} \quad \text{and} \quad \rho_A Z\mathbf{F}$$

are local scalar multiples of $\mathbf{F}$, and hence also belong to $\mathcal{M}_{a,r}$. This proves the implication.

For general A, apply the implication to $A - A(a)$. By the linearity of $A \mapsto B_A$, ρ_A, and $\mathcal{T}(A)$,

$$\mathcal{T}(A) - A(a)\mathcal{T}(1) = \mathcal{T}(A - A(a)) \in \mathcal{M}_{a,r},$$

which is (24). □

Remark 8. If $\mathbf{F}(a)$ and $\mathbf{F}'(a)$ are linearly independent, the rank in Lemma 7 is exactly one. At order zero,

$$\mathcal{T}(1)(a) = G(a)^{-2}\mathbf{F}'(a) + \sigma(a)\mathbf{F}(a)$$

for a scalar $\sigma(a) \in E$, so $\mathcal{T}(1)(a) \notin E\mathbf{F}(a)$. Exactness is not needed for the lower bound; "at most one" is sufficient.

3.5 Dimension count and conclusion

Lemma 9 (Independence modulo the true solution)**.** *The map*

$$A \longmapsto \mathcal{T}(A) \bmod E\mathbf{F}$$

is injective.

Proof. Choose $p < q$ with $\Delta_{p,q} \neq 0$. From the definition of $\mathcal{T}$, write $\mathcal{T}(A)_j$ for its j-th coordinate. The terms containing $B_A\mathbf{F}$ and $\rho_A Z\mathbf{F}$ cancel in the following 2×2 determinant:

$$G^2\big(f_p\mathcal{T}(A)_q - f_q\mathcal{T}(A)_p\big) = A\big(f_p f_q' - f_q f_p'\big) = A\Delta_{p,q}. \tag{25}$$

If $\mathcal{T}(A) \in E\mathbf{F}$, the left side is zero. Since $\Delta_{p,q} \neq 0$ and $E[Z]$ has no zero divisors, $A = 0$. □

Proof of Theorem 1. Let S contain $c \leq 2t + 2$ distinct support points and define

$$P_S(Z) = \prod_{a \in S}(Z - a).$$

The space of polynomials of degree at most $2t + 2$ that vanish on S is

$$\mathcal{A}_S = P_S E[Z]_{\leq 2t+2-c} = \{P_S B : B \in E[Z]_{\leq 2t+2-c}\},$$

and therefore

$$\dim_E \mathcal{A}_S = 2t + 3 - c. \tag{26}$$

For every $A \in \mathcal{A}_S$, Lemma 7 shows that $\mathcal{T}(A)$ satisfies every Step 3 flag condition at every point of S, regardless of the local stopping orders r_i. Hence

$$E\mathbf{F} + \mathcal{T}(\mathcal{A}_S) \subseteq \mathrm{Sol}(S).$$

Lemma 9 makes this sum direct as a vector-space sum, so (26) gives

$$\dim_E \mathrm{Sol}(S) \geq 1 + \dim_E \mathcal{A}_S = 2t + 4 - c.$$

At $c = 2t + 2$, at least one non-scalar candidate remains. The solution space therefore cannot equal $E\mathbf{F}$ before $c = 2t + 3$. □

3.6 Scope of the result

For the ISO/NIST Category-5 parameter set `mceliece8192128`, $t = 128$. The theorem gives

$$c_{\mathrm{need}} \geq 2(128) + 3 = 259.$$

This conclusion is not an extrapolation from the experiments. It follows from the binary Goppa multiplier identity and applies at the standardized degree, subject to the explicit condition $\Delta_{p,q} \neq 0$ for some pair $p < q$.

The theorem does not prove $c_{\mathrm{need}} = 2t+3$. At $c = 2t+3$, the explicit vanishing-polynomial family used above becomes zero, but additional solutions to the Step 3 linear system may remain. Excluding those solutions is the matching upper-bound problem and is deliberately not assumed here.

If every $\Delta_{p,q}$ is zero, Lemma 9 does not apply. In that case $\mathbf{F}'$ is proportional to $\mathbf{F}$ over $E(Z)$, and the first derivative spaces need not have the growth assumed by Vedenev's equation count. The present theorem states the required nondegeneracy condition explicitly rather than hiding it in a generic-rank assumption.

On the Use of AI

The author used OpenAI GPT-5.6-Sol and Claude Opus 5 to generate programs to derive the experimental artifacts, only taking empirical data when the two LLMs' experiments independently agreed. Part of this interaction involved feedback from AI about the nature of the observed data, which was incorporated into the analysis phase of this work. The author also used AI to draft an early version of this note, which was later heavily re-written line-by-line for presentation purposes.

Acknowledgements

The author thanks Brennon Brimhall, Victor Miller, and Scott Simon for detailed peer-review feedback on early drafts of this note.

References

[1] H. Randriambololona. *The syzygy distinguisher*, IACR Cryptology ePrint Archive, Report 2024/1193, 2024. (Published at EUROCRYPT 2026.)

[2] P. Briaud, A. Lemoine, H. Randriambololona, and J.-P. Tillich. *A Heuristic Subexponential Attack on the McEliece Cryptosystem*, IACR Cryptology ePrint Archive, Report 2026/1232, 2026.

[3] T. Hemmert and A. Wiemers. *Distinguishing Goppa codes using higher-order vanishing*, IACR Cryptology ePrint Archive, Report 2025/1661, 2025. (Published at CRYPTO 2026.)

[4] T. Hemmert. *Key recovery for the McEliece cryptosystem using higher-order vanishing*, IACR Cryptology ePrint Archive, Report 2026/1339, 2026.

[5] M. Bardet, A. Lemoine, and Jean-Pierre Tillich. *An attack on the CFS scheme and on TII McEliece challenges*, IACR Cryptology ePrint Archive, Report 2026/430, 2026.

[6] A. Ghoshal, Y. Ishai, A. Jain, and N. Sun. *Quasipolynomial Cryptanalysis of the McEliece Cryptosystem (or: PIR Meets McEliece)*, IACR Cryptology ePrint Archive, Report 2026/1630, 2026.

[7] M.-J. O. Saarinen. *HOVER: Higher-Order Vanishing Endomorphism Recovery*, IACR Cryptology ePrint Archive, Report 2026/1778, 2026.

[8] M.-J. O. Saarinen. *Bit Operation Cost of "Holdout" Key-Recovery Attacks Against Classic McEliece*, IACR Cryptology ePrint Archive, Report 2026/1786, 2026.

[9] K. Vedenev. *Extending Distinguishing to Key Recovery for Subfield Subcodes of GRS Codes*, IACR Cryptology ePrint Archive, Report 2026/1747, working draft, 2026.

A Notation and relation to Classic McEliece parameters

Classic McEliece parameter sets are commonly described by the extension degree m, code length n_{spec}, and Goppa-polynomial degree t. The associated binary Goppa code has designed error-correction capability t and, for the standardized irreducible instances, dimension

$$k_{\text{spec}} = n_{\text{spec}} - mt.$$

Classic McEliece is specified in Niederreiter form using a parity-check matrix. This note uses the equivalent code and generator-matrix description because the distinguisher and Step 3 reconstruction are most directly expressed in terms of public columns.

The principal notation is as follows.

Symbol	Meaning
$\mathbb{F}_2$	The binary field.
m	The extension degree in the Classic McEliece parameter set.
$E = \mathbb{F}_{2^m}$	The field containing the secret support values.
n_{spec}	The unshortened code length in the Classic McEliece specification.
k_{spec}	The unshortened binary code dimension $n_{\text{spec}} - mt$.
n	The retained code length in Sections 1–3, after any shortening.
k	The dimension of the retained public code and the number of rows of Y; after shortening it need not equal k_{spec}.
t	The degree of G, and the designed error-correction capability of the irreducible binary Goppa code.
$L = (a_1, \ldots, a_n)$	The retained secret support, with distinct entries in E.
G	The secret irreducible Goppa polynomial of degree t, with no root in L.
$\Gamma(L, G)$	The binary Goppa code defined in Section 1.1.
Z, U	Independent formal indeterminates used for polynomial arithmetic and Hasse-derivative expansions.
$\Pi = \prod_i (Z - a_i)$	The retained support polynomial; Π' is its ordinary formal derivative.
$Y = [y_1 \vert \cdots \vert y_n]$	A $k \times n$ binary generator matrix of the retained public code.
$D = n - 2t - 1$	The degree bound for the hidden polynomial vector; D is derived notation, not a Classic McEliece specification parameter.
λ_i	The i-th multiplier of the ambient GRS code.
$\mathbf{F} = (f_1, \ldots, f_k)^T$	The hidden vector polynomial satisfying $y_i = \lambda_i \mathbf{F}(a_i)$.
d, s	The polynomial degree and required Hasse multiplicity in the GIJS hold-out distinguisher.
H	A set of coordinate indices held out when constructing a distinguisher kernel.
c	The number of held support points supplied to Step 3.
c_{V}	Vedenev's independence-based estimate for the required value of c.
c_{need}	The least prefix length for which the Step 3 solution space is the line $E\mathbf{F}$.
r_i	The first derivative order at which the recovered flag at a_i fills E^k.
$\dim_E$	Vector-space dimension over E.

For e.g. `mceliece8192128`,

$$m = 13, \qquad n_{\text{spec}} = 8192, \qquad t = 128, \qquad k_{\text{spec}} = 8192 - 13 \cdot 128 = 6528.$$

The variables n, k, and D in the main proof refer to the retained code after any shortening and therefore may differ from these unshortened specification values.