

An Explicit Injection from Unsatisfiable to Uniquely Satisfiable Quadratic Systems

Gabriele Radici* and Massimiliano Sala**

Abstract. Let $\text{MQ}_0(n)$ and $\text{MQ}_1(n)$ be the sets of square systems of n quadratic equations in n variables over $\mathbb{F}_2$ having, respectively, no solution and exactly one solution, and write $\alpha_0 = |\text{MQ}_0(n)|$, $\alpha_1 = |\text{MQ}_1(n)|$. It is known that $\alpha_0 < \alpha_1$, so injections $\text{MQ}_0(n) \hookrightarrow \text{MQ}_1(n)$ exist; the proof is a counting argument and produces no map. We exhibit a partial injection τ , in the following sense. The map τ adds to a system a single correction supported on an affine flat of codimension two through a point read off by a three-point decoder; it is computable in $\Theta(n^3)$ bit operations from the dense coefficient representation, with no precomputation and no advice. We prove that there is a subset $\mathcal{D} \subseteq \text{MQ}_0(n)$ with $|\mathcal{D}|/|\text{MQ}_0(n)| > 0.745$, uniformly in $n \geq 3$, on which τ is injective into $\text{MQ}_1(n)$. We then show that inverting τ is *equivalent* to locating a zero of the image system — the solvable case of the multivariate quadratic problem, and its one-solution case on the uniquely solvable part of the image — and that τ is randomly self-reducible on each orbit of its symmetry group. Under the standard assumption that solvable MQ is hard on average — the planted-MQ assumption of MQ-based identification and signature schemes — τ is therefore a one-way function, injective on $\mathcal{D}$, explicit and randomly self-reducible on each orbit, between two sets which are, as languages, complexity-theoretically far apart. We also show that the associated flat-collision problem is NP-complete.

Keywords: Multivariate quadratic systems, Reed–Muller codes, Partial injections, One-way functions, NP-completeness.

1 Introduction

Deciding whether a system of multivariate quadratic equations over $\mathbb{F}_2$ has a solution is NP-complete [2,3], and remains so for *square* systems, with as many equations as variables [1]; the hardness of this problem, MQ, underlies multivariate post-quantum cryptography [9,10]: the oil-and-vinegar signature family [5,6], four members of which (UOV, MAYO, QR-UOV, SNOVA) were advanced in May 2026 to the third round of NIST’s additional-signature process alongside MQOM [12,15]; the identification and signature schemes built directly on random instances [13,14,15]; and the QUAD stream cipher [17]. The assumption has been tested: Rainbow [7] was broken during NIST’s first standardisation

* Department of Mathematics, University of Trento, Italy.
gabriele.radici@studenti.unitn.it

** Department of Mathematics, University of Trento, Italy. maxsalacodes@gmail.com

process [8], which standardised lattice- and hash-based schemes [11], and the additional process exists to diversify the underlying assumptions. Write $\text{MQ}_0(n)$ and $\text{MQ}_1(n)$ for the sets of square quadratic systems in n variables having no solution and exactly one solution respectively, and α_0, α_1 for their cardinalities. As languages the two sets sit far apart: $\bigcup_n \text{MQ}_0(n)$ is coNP-complete, while $\bigcup_n \text{MQ}_1(n)$ lies in DP.

Their cardinalities, by contrast, are almost equal. It is proved in [1] that

$$\alpha_0 < \alpha_1 \leq \left(1 + \frac{1}{2^n - 1}\right) \alpha_0,$$

so injections $\text{MQ}_0(n) \hookrightarrow \text{MQ}_1(n)$ exist and their images omit at most a 2^{-n} fraction of the target. That proof is a counting argument: it compares the characteristic polynomials of a matroid and of its contraction at a point, evaluated at 2^n , and exhibits no map. Whether an injection can be written down is left open.

1.1 Contribution

We solve this problem partially, as detailed below. Section 3 defines a map τ which adds to a system S a single correction $S(v) \cdot \mathbf{1}_A$, where $v = v(S)$ is read off by a fixed three-point decoder and A is an affine flat (of codimension two) containing v and constructed deterministically from v .

Let $\mathcal{D} = M_0 \cap \tau^{-1}(M_1)$, where $M_0 = \text{MQ}_0(n)$ and $M_1 = \text{MQ}_1(n)$, be the *canonical domain*. We prove:

- τ is total and computable in $\Theta(n^3)$ bit operations (Proposition 2), linear in the dense instance size;
- τ is injective on $\mathcal{D}$, which has density greater than $307/412 > 74.5\%$ uniformly in $n \geq 3$ (Theorem 1, Corollary 1, Theorem 3);
- inverting τ is *equivalent* to locating a zero of the image (Lemma 6; its unique zero when the image is uniquely solvable, Proposition 4), so under the standard assumption that solvable MQ is hard on average, τ is a one-way function, injective on $\mathcal{D}$;
- τ is randomly self-reducible on each orbit of the natural symmetry group G : $\tau(\mathcal{D})$ meets every G -orbit of M_1 in proportion at least $\frac{3}{4} - \frac{1}{N}$ (Theorem 4), and an inverter with orbit-average success f yields one succeeding on every instance of that orbit (Proposition 6), so worst-case hardness on an orbit implies average-case hardness on it;
- the associated flat-collision problem is NP-complete (Theorem 5).

1.2 What is not claimed

Three limitations should be stated at the outset, since each is easy to overread.

First, $\tau(\mathcal{D}) \subseteq M_1$ and $\tau(\mathcal{D})$ is very large in M_1 , but $\tau(M_0) \not\subseteq M_1$.

Second, τ is injective on $\tau^{-1}(M_1)$ (not only on $\tau^{-1}(M_1) \cap M_0$), but τ is not injective on M_0 (at $n = 3$, $|\tau(M_0)| = 709494 < 720601 = \alpha_0$).

Third, the one-wayness of τ rests on the hardness, on average, of locating a zero of the image of a uniformly random system — a solvable MQ instance drawn with probability proportional to its number of zeros, the planted-MQ assumption of , which is shared with much of nowadays multivariate post-quantum cryptography (see Section 6). Note however that, by Lemma 6 and Proposition 4, inverting τ and locating a zero of the image *are* the same problem, so nothing is lost or gained in the translation.

2 Preliminaries

Throughout the paper $n \geq 3$.

We follow the notation of [1]. Let $V = \mathbb{F}_2^n$ and $N = 2^n$, and let $\text{RM}(2, n)$ be the space of Boolean polynomial functions on V of degree at most two, each written in its unique multilinear algebraic normal form; a basis is given by the squarefree monomials of degree at most two, so $\dim \text{RM}(2, n) = D = 1 + n + \binom{n}{2}$. Identifying a function with its vector of values on V realises $\text{RM}(2, n)$ as a binary code of length N , whose dual is $\text{RM}(n - 3, n)$, of minimum distance $2^{n-(n-3)} = 8$. For $x \in V$ let ev_x denote the evaluation functional $f \mapsto f(x)$. Every $\mathbb{F}_2$ -combination of evaluation functionals is $\sum_x u_x \text{ev}_x$ for a unique $u \in \mathbb{F}_2^V$, its *support vector*; the map $u \mapsto \sum_x u_x \text{ev}_x$ has kernel exactly $\text{RM}(2, n)^\perp$, so a combination vanishes on $\text{RM}(2, n)$ if and only if its support vector is a dual word. We write $\text{supp } \varphi$ for the support of the vector u . A *system* is an n -tuple $S = (S_1, \dots, S_n)$ with $S_i \in \text{RM}(2, n)$, regarded as a map $S : V \rightarrow V$; there are N^D of them. For a functional φ on $\text{RM}(2, n)$ and a system S we write $\varphi(S) := (\varphi(S_1), \dots, \varphi(S_n)) \in V$. Write

$$\mathcal{Z}(S) = \{x \in V : S(x) = 0\}, \quad X(S) = |\mathcal{Z}(S)|, \quad M_k = \{S : X(S) = k\},$$

and abbreviate $M_0 = \text{MQ}_0(n)$ and $M_1 = \text{MQ}_1(n)$, so $\alpha_0 = |M_0|$ and $\alpha_1 = |M_1|$. Finally $\mu_n = \alpha_0/N^D$ denotes the density of M_0 among all systems. Throughout, “uniform system model” means S uniform on all N^D systems; no statement below assumes a distribution on M_0 or M_1 other than the uniform one on those finite sets, used as a counting ratio.

Two facts about $\text{RM}(2, n)$ are used constantly; both rest on the minimum distance of the dual code.

Lemma 1 (Local independence). *Let $Q = \text{RM}(2, n)$, so that $Q^\perp = \text{RM}(n - 3, n)$ has minimum distance 8.*

- (a) (Small-support independence.) *Let $\varphi_1, \dots, \varphi_k$ be $\mathbb{F}_2$ -combinations of evaluation functionals whose supports have union of size at most 7. Then the φ_i are linearly independent as functionals on Q if and only if their support vectors are linearly independent in $\mathbb{F}_2^V$; in particular a single combination with nonzero support vector of size at most 7 does not vanish on Q .*
- (b) (Exact local independence.) *Let $x_1, \dots, x_r \in V$ be distinct with $r \leq 7$ and let S be uniform on the N^D systems. Then $(S(x_1), \dots, S(x_r))$ is uniform on V^r ; in particular the r values are independent.*

Proof. (a) Under the identification of Section 2, a combination $\sum_i c_i \varphi_i$ vanishes on Q exactly when its support vector lies in $Q^\perp$. That support is contained in the union of the $\text{supp } \varphi_i$, of size $\leq 7 < 8$, so the only dual word available is 0: the combination vanishes on Q if and only if its support vector is 0, which is the stated equivalence.

(b) For a binary linear code Q of length N and a coordinate set I , the projection $Q \rightarrow \mathbb{F}_2^I$ is surjective if and only if no nonzero word of $Q^\perp$ is supported inside I : the image is the annihilator of $\{u \in \mathbb{F}_2^I : u \cdot q|_I = 0 \ \forall q \in Q\}$, and such u extend by zero to words of $Q^\perp$ supported in I . Here $Q = \text{RM}(2, n)$ and $Q^\perp = \text{RM}(n-3, n)$, of minimum distance $2^{n-(n-3)} = 8$, so no nonzero dual word has support of size ≤ 7 and every projection onto at most seven coordinates is onto. A uniform $S_i \in \text{RM}(2, n)$ therefore has $(S_i(x_1), \dots, S_i(x_r))$ uniform on $\mathbb{F}_2^r$, and the n coordinate functions are drawn independently.

Lemma 2 (Rank–probability dictionary). *Let $\varphi_1, \dots, \varphi_k$ be combinations over $\mathbb{F}_2$ of evaluation functionals on $\text{RM}(2, n)$, whose rank as functionals on $\text{RM}(2, n)$ is ρ — the rank of their support vectors modulo $\text{RM}(n-3, n)$ — and let $c_1, \dots, c_k \in V$ be fixed.*

If S is extracted uniformly from all N^D systems,

$$\Pr[\varphi_i(S) = c_i \text{ for all } i] = \begin{cases} N^{-\rho}, & \text{if the prescription is consistent,} \\ 0, & \text{otherwise,} \end{cases}$$

where consistent means that every $\mathbb{F}_2$ -dependence $\sum_{i \in I} \varphi_i = 0$ among the functionals on $\text{RM}(2, n)$ is matched by $\sum_{i \in I} c_i = 0$ among the values.

Proof. A uniform system has its n coordinate functions drawn independently and uniformly from $\text{RM}(2, n)$, so we may work coordinatewise. For a single coordinate the map $S_i \mapsto (\varphi_1(S_i), \dots, \varphi_k(S_i))$ is $\mathbb{F}_2$ -linear on $\text{RM}(2, n)$ with image of dimension ρ ; a uniform S_i therefore lands in a prescribed fibre with probability $2^{-\rho}$ when that fibre is nonempty, and the fibre is nonempty exactly under the stated consistency condition. The n coordinates are independent, so the probability is $(2^{-\rho})^n = N^{-\rho}$.

A third lemma decides exactly which combinations of the functionals occurring below can vanish. Its proof, in Appendix A, uses the fact that the minimum-weight words of a Reed–Muller code are exactly the indicators of affine flats of the appropriate dimension [4, Ch. 13].

Lemma 3 (Relation structure of decoder–collision families). *Let $B \subseteq V$ with $|B|$ odd, let $a \in V$, let $x_1, \dots, x_k \in V \setminus \{a\}$ be distinct, and let $Y \subseteq V$. Put*

$$\varphi_B = \sum_{b \in B} \text{ev}_b, \quad \psi_j = \text{ev}_{x_j} + \text{ev}_a \ (1 \leq j \leq k), \quad \text{ev}_y \ (y \in Y),$$

and for $I \subseteq [k]$ write $\psi_I = \sum_{j \in I} \psi_j = \sum_{j \in I} \text{ev}_{x_j} + (|I| \bmod 2) \text{ev}_a$, supported on $\Delta_I := \{x_j : j \in I\} \cup \{a \text{ if } |I| \text{ is odd}\}$, a disjoint union of even size, nonempty

whenever $I \neq \emptyset$. For $\epsilon \in \mathbb{F}_2$, $I \subseteq [k]$ and $Y' \subseteq Y$, the combination $\epsilon\varphi_B + \psi_I + \sum_{y \in Y'} \text{ev}_y$ vanishes on $\text{RM}(2, n)$ if and only if its support

$$u := \Delta_I \Delta Y' \quad (\epsilon = 0), \quad u := B \Delta \Delta_I \Delta Y' \quad (\epsilon = 1)$$

is a word of $\text{RM}(n-3, n)$. Consequently:

- (a) every word of $\text{RM}(n-3, n)$ has even weight, and $|\Delta_I|$ is even, so no combination with $\epsilon = 1$ and $|Y'|$ even can vanish: its support has odd size;
- (b) if $1 \leq |u| \leq 7$ the combination does not vanish on $\text{RM}(2, n)$. In particular, if $|B \cup \{a, x_1, \dots, x_k\} \cup Y| \leq 7$, then a combination vanishes on $\text{RM}(2, n)$ if and only if $u = \emptyset$, that is, if and only if it cancels formally, as an element of $\mathbb{F}_2^V$, with no dual word involved;
- (c) every vanishing combination with $u \neq \emptyset$ has $|u| \geq 8$; if $|u| = 8$ then u is a minimum-weight word of $\text{RM}(n-3, n)$, hence the indicator of an affine 3-flat.

Proof. A combination of evaluation functionals vanishes on a code exactly when its support vector lies in the dual code, here $\text{RM}(2, n)^\perp = \text{RM}(n-3, n)$ (Section 2); the displayed supports follow from $\psi_j = \text{ev}_{x_j} + \text{ev}_a$, the point a being counted $|I|$ times.

For (a), $\mathbf{1}_V \in \text{RM}(0, n) \subseteq \text{RM}(2, n)$ forces every dual word u to satisfy $\sum_x u_x = 0$, i.e. to have even weight, the zero word included. The size of a symmetric difference has the parity of the sum of the sizes, so for $\epsilon = 1$ the support $B \Delta \Delta_I \Delta Y'$ has the parity of $|B| + |\Delta_I| + |Y'|$, which is odd because $|B|$ is odd, $|\Delta_I|$ is even and $|Y'|$ is even. A vector of odd weight is not a dual word, so the combination does not vanish.

For (b), the minimum distance of $\text{RM}(n-3, n)$ is $2^{n-(n-3)} = 8$: a nonzero dual word has weight at least 8, so no vector of weight between 1 and 7 is one, and a combination with $1 \leq |u| \leq 7$ does not vanish. If $P := B \cup \{a, x_1, \dots, x_k\} \cup Y$ has at most seven points then every support u lies in P , so $|u| \leq 7$ and the combination vanishes if and only if $u = \emptyset$. Formal cancellations do occur when the point sets overlap — for instance $\psi_1 + \text{ev}_{x_1} + \text{ev}_a = 0$ when $x_1, a \in Y$, or $\varphi_B + \sum_{b \in B} \text{ev}_b = 0$ when $B \subseteq Y$ — so a bound on the support alone does not give independence; the proofs of Lemma 10 and Theorem 3 enumerate them explicitly.

For (c), a vanishing combination with $u \neq \emptyset$ has u a nonzero word of $\text{RM}(n-3, n)$, so $|u| \geq 8$ by minimum distance, with equality only for a minimum-weight word; the minimum-weight words of $\text{RM}(r, m)$ are the indicators of the affine $(m-r)$ -flats [4, Ch. 13], and here $m-r = n - (n-3) = 3$.

3 The map τ

The construction has two ingredients: a decoder, which reads a point of V off a system, and a rule assigning to each point an affine flat through it.

Convention 1 (Canonical encodings). Identify $x \in V$ with the integer $\sum_i x_i 2^{i-1}$ and order V accordingly (*lexicographic order*); identify a linear form

ℓ with the integer of its coefficients, so $\ell(x)$ is the parity of the bitwise AND; solve a linear system prescribing the values of a form on independent vectors by reduced row echelon form with all free variables set to 0, which selects a unique solution.

Definition 1 (Decoder set; admissible assignments). *A decoder set is a set $B = \{b_1, b_2, b_3\} \subseteq V$ of three distinct points, always listed in the increasing order of Convention 1, so that (b_1, b_2, b_3) is determined by B . An admissible assignment for B is a map $a \mapsto A(a)$ sending each $a \in V$ to an $(n-2)$ -flat with $a \in A(a)$ and $|A(a) \cap B|$ odd.*

Unless stated otherwise we take the canonical decoder set

$$B = \{e_1, e_2, e_3\},$$

available since $n \geq 3$, together with the canonical assignment constructed in Lemma 4; this will fix a unique map τ . The structural results below hold for any decoder set and admissible assignment, and the density bounds of Section 5 for any decoder set of three distinct nonzero points and any assignment of $(n-2)$ -flats through the anchor.

We construct the flats depending on the decoder set, in such a way that they are uniquely determined. The construction is what makes τ a single map rather than a family. The proof is in Appendix A.

Lemma 4 (Deterministic construction of the flats). *Let $n \geq 3$, let $B = \{b_1, b_2, b_3\}$ be a decoder set and let $a \in V$. Put*

$$d_j := b_j + a \quad (j = 1, 2, 3).$$

The following procedure outputs a uniquely determined pair of independent linear forms ℓ_1, ℓ_2 in $\Theta(n)$ bit operations, all choices being resolved by Convention 1; setting

$$W(a) := \ker \ell_1 \cap \ker \ell_2, \quad A(a) := a + W(a),$$

the resulting assignment is admissible for B in the sense of Definition 1. In every case the procedure achieves the stronger conclusion $|A(a) \cap B| = 1$, not merely odd parity; this sharper form is used later. Below we abbreviate $W = W(a)$ and $A = A(a)$, and note that $b_j \in A$ if and only if $d_j \in W$.

- (1) $a \in B$. Exactly one $d_j = 0$; the other two, listed in increasing index order as r_1, r_2 , are nonzero and distinct, hence independent. Let ℓ_1 be the solution of $(\ell(r_1), \ell(r_2)) = (1, 0)$ and ℓ_2 that of $(0, 1)$. Then $W \cap \{d_1, d_2, d_3\} = \{0\}$ and $|A \cap B| = 1$.
- (2) $a \notin B$, d_1, d_2, d_3 **independent**. Let ℓ_1 solve $(\ell(d_1), \ell(d_2), \ell(d_3)) = (0, 1, 0)$ and ℓ_2 solve $(0, 0, 1)$. Then $d_1 \in W$, $d_2, d_3 \notin W$, and $|A \cap B| = 1$.
- (3) $a \notin B$, $\dim \langle d_1, d_2, d_3 \rangle = 2$. Let z be the lexicographically least vector outside $\langle d_1, d_2 \rangle$. Let ℓ_1 solve $(\ell(d_1), \ell(d_2), \ell(z)) = (0, 1, 0)$ and ℓ_2 solve $(0, 1, 1)$. Then ℓ_1, ℓ_2 agree on $\langle d_1, d_2 \rangle$, so $d_1 \in W$ while d_2 and $d_3 = d_1 + d_2$ are not, and $|A \cap B| = 1$.

Remark 1. — For (c) of Lemma 4 one has that three distinct nonzero vectors span a plane exactly when they are its three nonzero vectors, that is, when $d_1 + d_2 + d_3 = 0$; since $d_1 + d_2 + d_3 = a + b_1 + b_2 + b_3$, this case occurs for exactly one anchor, $a = b_1 + b_2 + b_3$, a point never in B ($b_1 + b_2 + b_3 = b_1$ would force $b_2 = b_3$). Here $d_3 = d_1 + d_2$, so no relabelling arises.

— In every case both forms are nonzero — $\ell_1(r_1) = \ell_2(r_2) = 1$ in case (1), $\ell_1(d_2) = \ell_2(d_3) = 1$ in case (2), $\ell_1(d_2) = \ell_2(d_2) = 1$ in case (3) — and they differ, at r_1 , d_2 and z respectively. Over $\mathbb{F}_2$ two distinct nonzero forms are linearly independent, so $\dim W = n - 2$ and $|A| = 2^{n-2} = N/4$.

Definition 2 (The map). For a system S set

$$v(S) = S(b_1) + S(b_2) + S(b_3) \in V, \quad \tau(S)_i = S_i + S_i(v(S)) \cdot \mathbf{1}_{A(v(S))} \quad (1 \leq i \leq n),$$

equivalently $\tau(S)(x) = S(x) + S(v(S))$ for $x \in A(v(S))$ and $\tau(S)(x) = S(x)$ otherwise. By Lemma 4 and Convention 1 this is one well-defined map, not a family.

Proposition 1 (Well-definedness). For every $(n-2)$ -flat A the indicator $\mathbf{1}_A$ lies in $\text{RM}(2, n)$; consequently $\tau(S)$ is again a system of degree at most two, and $\tau(S)(v(S)) = 0$.

Proof. $S_i(v)$ is a scalar in $\mathbb{F}_2$, so the added term is 0 or $\mathbf{1}_A$, and $\mathbf{1}_A = (1 + \ell_1 + \ell_1(a))(1 + \ell_2 + \ell_2(a))$ is a product of two affine forms, hence lies in $\text{RM}(2, n)$. Since $v \in A$, $\tau(S)(v) = S(v) + S(v) = 0$.

Proposition 2 (Cost). τ is total and computable in $\Theta(n^3)$ bit operations, with no precomputed table and no advice: three evaluations of n quadratic forms ($\Theta(n^3)$), one flat construction ($\Theta(n)$, Lemma 4), the expansion of $\mathbf{1}_A = (1 + \ell_1 + \ell_1(a))(1 + \ell_2 + \ell_2(a))$ into its $O(n^2)$ coefficients, and its addition to the affected coordinates ($O(n^3)$). Input and output each have $nD = \Theta(n^3)$ coefficients, so this is linear in the instance size, and writing the dense output gives the matching $\Omega(n^3)$. τ never inspects $X(S)$.

3.1 The algebra of the correction

For any system S , any flat A and $c \in V$, the decoder reads S only through b_1, b_2, b_3 , and $\mathbf{1}_A$ contributes $c(|A \cap B| \bmod 2)$ to that sum; hence, whenever $|A \cap B|$ is odd — in particular for every flat of an admissible assignment —

$$v(S + c\mathbf{1}_A) = v(S) + c. \tag{*}$$

We use $(*)$ throughout.

Abbreviate $g_a := \mathbf{1}_{A(a)}$, and write

$$\Phi_{a,c}(S) = S + c g_a, \quad P_a(S) = \Phi_{a,S(a)}(S), \quad R_a(S) = \Phi_{a,a+v(S)}(S).$$

Lemma 5 (Decoder translation and idempotents). *For all $a, c \in V$ and systems S :*

$$v(\Phi_{a,c}(S)) = v(S) + c, \quad v(R_a(S)) = a,$$

and P_a, R_a are idempotent with the absorption laws

$$P_a^2 = P_a, \quad R_a^2 = R_a, \quad P_a R_a = P_a, \quad R_a P_a = R_a.$$

Moreover $\tau(S) = P_{v(S)}(S)$.

Proof. The first identity is $(*)$ with $A = A(a)$; taking $c = a + v(S)$ gives $v(R_a(S)) = a$. Since $P_a(S)(a) = S(a) + S(a)g_a(a) = 0$, a second P_a adds nothing, so $P_a^2 = P_a$; likewise $v(R_a(S)) = a$ makes the correction of a second R_a vanish, so $R_a^2 = R_a$. Writing $c = a + v(S)$, we have $R_a(S)(a) = S(a) + c$, so $P_a R_a(S) = R_a(S) + (S(a) + c)g_a = S + S(a)g_a = P_a(S)$; and $v(P_a(S)) = v(S) + S(a)$ gives $R_a P_a(S) = P_a(S) + (a + v(S) + S(a))g_a = S + (a + v(S))g_a = R_a(S)$. Finally $\tau(S) = S + S(v(S))g_{v(S)} = P_{v(S)}(S)$ by definition of τ .

The next two results describe what τ does to zero sets and what its fibres look like. They are the only computations needed for injectivity.

Proposition 3 (Exact zero set). *With $v = v(S)$, $A = A(v)$, the zero set of $\tau(S)$ is the disjoint union*

$$\mathcal{Z}(\tau(S)) = (\mathcal{Z}(S) \setminus A) \sqcup (A \cap S^{-1}(S(v))).$$

Taking cardinalities,

$$X(\tau(S)) = 1 + \#\{x \in A \setminus \{v\} : S(x) = S(v)\} + \#\{x \notin A : S(x) = 0\},$$

or equivalently, for any system S whatsoever,

$$X(\tau(S)) - X(S) = |A \cap S^{-1}(S(v))| - |A \cap \mathcal{Z}(S)|.$$

For $S \in M_0$ the last term of the second display vanishes, so $\tau(S) \in M_1 \iff S(x) \neq S(v)$ for all $x \in A \setminus \{v\}$.

Proof. Outside A values are unchanged, so the zeros of $\tau(S)$ there are exactly those of S ; inside A the value at x is $S(x) + S(v)$, vanishing exactly when $S(x) = S(v)$. The two sets are disjoint since one lies off A and the other in A , and v belongs to the second, giving the leading 1. Subtracting $X(S) = |\mathcal{Z}(S) \setminus A| + |A \cap \mathcal{Z}(S)|$ from the set identity gives the difference form.

Lemma 6 (The fibres of τ are the zero sets of the images). *Let T be any system and $\mathcal{Z}(T) = \{x : T(x) = 0\}$. For $v' \in \mathcal{Z}(T)$ put*

$$S_{v'} := T + (v(T) + v') \mathbf{1}_{A(v')}.$$

Then $v(S_{v'}) = v'$ and $\tau(S_{v'}) = T$; and conversely every S with $\tau(S) = T$ equals $S_{v(S)}$ with $v(S) \in \mathcal{Z}(T)$. Hence $\tau^{-1}(T) = \{S_{v'} : v' \in \mathcal{Z}(T)\}$, a set of exactly $X(T)$ elements.

Proof. For $v' \in \mathcal{Z}(T)$ we have $S_{v'} = T + (v(T) + v')\mathbf{1}_{A(v')} = R_{v'}(T)$. Applying Lemma 5 with $a = v'$, $P_{v'}R_{v'} = P_{v'}$ and $T(v') = 0$ (so $P_{v'}(T) = T$),

$$\tau(S_{v'}) = P_{v(S_{v'})}(R_{v'}T) = P_{v'}R_{v'}(T) = P_{v'}(T) = T,$$

using $v(S_{v'}) = v(R_{v'}T) = v'$. Conversely if $\tau(S) = T$ then $T(v(S)) = 0$ by Proposition 1, so $v(S) \in \mathcal{Z}(T)$; and $\tau(S) = P_{v(S)}(S) = T$ gives $S = T + S(v(S))\mathbf{1}_{A(v(S))}$, while applying the decoder yields $S(v(S)) = v(T) + v(S)$, so $S = R_{v(S)}(T) = S_{v(S)}$. The map $v' \mapsto S_{v'}$ is injective because $v(S_{v'}) = v'$.

4 Injectivity

Theorem 1 (Injectivity on $\tau^{-1}(M_1)$). *Let S, S' be systems with $\tau(S), \tau(S') \in M_1$. If $\tau(S) = \tau(S')$ then $S = S'$. Hence τ is injective on $\tau^{-1}(M_1)$, in particular on $\tau^{-1}(M_1) \cap M_0$.*

Proof. Put $T = \tau(S) = \tau(S')$. By Proposition 1, $T(v(S)) = T(v(S')) = 0$; as $T \in M_1$ has a unique zero, $v(S) = v(S') =: a$. Both therefore use the same flat $A := A(a)$, and

$$S + S(a)\mathbf{1}_A = S' + S'(a)\mathbf{1}_A \implies S' = S + c\mathbf{1}_A, \quad c := S(a) + S'(a).$$

By (*) applied to $S' = S + c\mathbf{1}_A$, $v(S') = v(S) + c$. As $v(S') = v(S)$ we get $c = 0$, so $S' = S$.

Corollary 1 (τ is a bijection onto M_1). *τ restricts to a bijection $\tau^{-1}(M_1) \rightarrow M_1$; in particular $|\tau^{-1}(M_1)| = \alpha_1$, and the unique preimage of $T \in M_1$ with zero z is $R_z(T) = T + (z + v(T))\mathbf{1}_{A(z)}$.*

Proof. Injectivity is Theorem 1. Surjectivity and the formula are Lemma 6: for $T \in M_1$ the zero set is the single point z , so $\tau^{-1}(T) = \{S_z\} = \{R_z(T)\}$, which is nonempty. Counting the fibres, $|\tau^{-1}(M_1)| = \sum_{T \in M_1} X(T) = \alpha_1$.

Corollary 1 hits *every* uniquely solvable system exactly once, and its domain $\tau^{-1}(M_1)$, of size $\alpha_1 > \alpha_0$, is not a subset of M_0 . Intersecting with M_0 gives the canonical domain $\mathcal{D} := M_0 \cap \tau^{-1}(M_1)$, on which τ is an injection $\mathcal{D} \hookrightarrow M_1$ by Theorem 1. No map defined on all of M_0 is claimed. “Canonical” means only *determined by the fixed decoder set and flat rule*: no efficient membership test for $\mathcal{D}$ is claimed — it is cut out by $S \in M_0$, coNP-hard to decide, and by $Z(S) = 0$, a collision question on the canonical flat, and we know no way to decide either without solving a quadratic system — nor an efficient sampler for $\mathcal{D}$, nor a way to move an arbitrary element of M_0 into it.

5 The canonical domain is large

By Proposition 3, for $S \in M_0$ we have $\tau(S) \in M_1$ exactly when no point of $A(v(S))$ other than $v(S)$ carries the value $S(v(S))$. Writing $Z(S)$ for the number

of such *collateral coincidences*, the question is how often $Z = 0$. Two bounds are needed: a lower bound on μ_n , so that conditioning on M_0 costs little, and an upper bound on $\Pr[Z \geq 1]$ in the uniform system model. For convenience, we define

$$B(N) := \sum_{k=0}^7 (-1)^k \binom{N}{k} N^{-k}.$$

Lemma 7 (Bonferroni bound on the density of M_0). *For every $n \geq 3$,*

$$\mu_n := \frac{\alpha_0}{N^D} \geq B(N).$$

Proof. With $A_x = \{S : S(x) = 0\}$ we have $\mu_n = \Pr[\bigcap_x A_x^c] = \sum_{k \geq 0} (-1)^k S_k$, $S_k = \sum_{|T|=k} \Pr[S|_T = 0]$. Truncating this alternating sum after a term of odd index gives a lower bound (Bonferroni). For $k \leq 7$ the evaluations at k distinct points are exactly independent and uniform, so $\Pr[S|_T = 0] = N^{-k}$ and $S_k = \binom{N}{k} N^{-k}$ exactly; $k = 7$ is odd. (At $n = 3$ the truncation drops exactly one term, $S_8 = \Pr[S \equiv 0] = 2^{-21}$, so $\mu_3 = B(8) + 2^{-21}$ exactly: $B(8) \cdot 2^{21} = 720600 = \alpha_0 - 1$.)

Lemma 8 (Hitting scheme). *Let Z count the collateral coincidences, so that $\tau(S) \notin M_1$ exactly when $Z(S) \geq 1$ for $S \in M_0$, and let U be any quantity with $\Pr[Z \geq 1] \leq U$ in the uniform system model. Then for every $n \geq 3$, every choice of three distinct points B and every assignment of $(n-2)$ -flats $A(a) \ni a$, admissible or not,*

$$\frac{\#\{S \in M_0 : \tau(S) \in M_1\}}{\alpha_0} \geq 1 - \frac{U}{\mu_n} \geq 1 - \frac{U}{B(N)}.$$

Admissible choices of U include $\mathbb{E}[Z]$ (Markov) and any odd-order truncation $S_1 - S_2 + \dots + S_{2r+1}$ of the inclusion-exclusion sum for $\Pr[Z \geq 1]$ (Bonferroni). More generally, the event $S \in M_0$ may be retained inside the sub-measure rather than discarded: if $\Pr[\{Z \geq 1\} \wedge M_0] \leq U'$ then

$$\frac{\#\{S \in M_0 : \tau(S) \in M_1\}}{\alpha_0} \geq 1 - \frac{U'}{\mu_n} \geq 1 - \frac{U'}{B(N)},$$

by the same computation with $\#\{S \in M_0 : Z(S) \geq 1\} = N^D \Pr[\{Z \geq 1\} \wedge M_0]$. Theorem 3 and Lemma 10 use this conditioned form.

Proof. $\#\{S \in M_0 : Z(S) \geq 1\} \leq \#\{S : Z(S) \geq 1\} = N^D \Pr[Z \geq 1] \leq N^D U$; subtracting from $\alpha_0 = \mu_n N^D$ and dividing by α_0 gives the first inequality. The second is $\mu_n \geq B(N)$ (Lemma 7) together with $U \geq 0$. Bonferroni's inequality bounds a union from above after an odd-indexed term, which gives the stated choices of U .

The scheme is applied twice here, with the first moment and then a conditioned fourth-order truncation; the conditioned second-order version, Lemma 10 in Appendix A, is the template that the fourth-order proof follows.

Lemma 9 (Exact first moment). *Fix $n \geq 3$, three distinct points $B = \{b_1, b_2, b_3\}$, and any assignment $a \mapsto A(a)$ of $(n-2)$ -flats with $a \in A(a)$. For a system S put $Z(S) = \#\{x \in A(v(S)) \setminus \{v(S)\} : S(x) = S(v(S))\}$. Then, for S uniform over all N^D systems,*

$$\mathbb{E}[Z] = \frac{1}{4} - \frac{1}{N} \quad \text{exactly.}$$

Proof. $\mathbb{E}[Z] = \sum_{a \in V} \sum_{x \in A(a) \setminus \{a\}} \Pr[v(S) = a \wedge S(x) = S(a)]$. Fix (a, x) with $x \neq a$. The event is $\{\varphi_B(S) = a\} \cap \{\psi_{a,x}(S) = 0\}$ for the two functionals

$$\varphi_B = \sum_{b \in B} \text{ev}_b, \quad \psi_{a,x} = \text{ev}_a + \text{ev}_x,$$

so it suffices that $\varphi_B, \psi_{a,x}$ be linearly independent on $\text{RM}(2, n)$: the induced map $\text{RM}(2, n) \rightarrow \mathbb{F}_2^2$ is then onto, hence $(v(S), S(a) + S(x))$ is uniform on V^2 across the n coordinates and the probability is N^{-2} . Independence is Lemma 3 with $k = 1$, $Y = \emptyset$: the three candidate supports are B (odd), $\{a, x\}$ (size $2 < 8$) and $B \triangle \{a, x\}$ (odd), none a dual word, so the rank is 2. No case analysis by coincidence pattern is needed. Since $|A(a)| = N/4$ for every a , summing gives $\mathbb{E}[Z] = N \cdot (N/4 - 1) \cdot N^{-2} = \frac{1}{4} - \frac{1}{N}$.

Theorem 2 (A constant fraction of M_0 reaches M_1). *For every $n \geq 3$, every choice of three distinct points B , and every assignment of $(n-2)$ -flats $A(a) \ni a$,*

$$\frac{\#\{S \in M_0 : \tau(S) \in M_1\}}{\alpha_0} \geq 1 - \frac{\frac{1}{4} - \frac{1}{N}}{\mu_n} \geq 1 - \frac{\frac{1}{4} - \frac{1}{N}}{B(N)} > \frac{33}{103} = 0.32038 \dots$$

Proof. Apply Lemma 8 with $U = \mathbb{E}[Z] = \frac{1}{4} - \frac{1}{N}$ (Lemma 9, Markov), which gives the first two inequalities. The final constant follows from Lemma 11 (Appendix A); alternatively, and exactly,

$$\left(1 - \frac{\frac{1}{4} - \frac{1}{N}}{B(N)}\right) - \frac{33}{103} = \frac{70(1082N^5 - 70N^4 - 119N^2 + 154N - 120)}{103(N^2 - 1)(309N^4 - 154N^3 + 239N^2 - 154N + 120)},$$

whose numerator and denominator are positive for $N \geq 8$. Thus $\mathbb{E}[Z] \rightarrow \frac{1}{4}$ and $B(N) \rightarrow \frac{103}{280}$, so the left-hand side tends to 0: the constant $\frac{33}{103} = 1 - \frac{70}{103}$ is not merely a convenient uniform number but the exact asymptotic value of this first-moment argument. (Note that $\frac{1}{4} - \frac{1}{N}$ and $B(N)$ both increase with N , so no monotonicity of the quotient may be inferred; Lemma 11 bounds it directly.)

Markov's inequality is wasteful here because it ignores that the coincidences are constrained: conditioning on the decoder value and using Lemma 3 to identify which functional combinations can vanish sharpens the estimate substantially. The next theorem carries this out at fourth order; its proof is in Appendix A, preceded there by the second-order version it builds on.

For later convenience we define

$$T_G(N) = \frac{3N^4 + 2N^3 - 7N^2 + 26N + 24}{8N^6}.$$

Theorem 3 (Conditioned fourth-order hitting bound). *Put*

$$R_4(N) = N\left(\frac{N}{4} - 1\right)T_G(N) + \frac{3}{N^3} + \frac{3}{N^4}.$$

Then for every $n \geq 3$, every set B of three distinct nonzero points and every assignment of $(n-2)$ -flats $A(a) \ni a$, admissible or not,

$$\frac{\#\{S \in M_0 : \tau(S) \in M_1\}}{\alpha_0} \geq 1 - \frac{R_4(N)}{B(N)} > \frac{307}{412} = 0.745145\dots,$$

and $1 - R_4(N)/B(N)$ decreases to $\frac{307}{412}$ from above.

The complete proof is in Appendix A below, we provide here only a sketch:

Proof. Retaining the event $S \in M_0$ and using only $|A(a)| = N/4$ with $a \in A(a)$, it suffices to bound $\sum_{a,x} \Pr[F_{a,x} \wedge M_0]$ and divide by $\alpha_0 \geq B(N)N^D$. Fix a pair (a, x) and apply an even 4th-order Bonferroni truncation to the zero-set indicator $\mathbf{1}_{[X(S)=0]}$:

$$\Pr[F_{a,x} \wedge M_0] \leq P_0 - S'_1 + S'_2 - S'_3 + S'_4, \quad S'_k = \sum_{|Y|=k} \Pr\left[F_{a,x} \wedge \bigcap_{y \in Y} E_y\right],$$

where $P_0 = N^{-2}$ and $E_y = \{S(y) = 0\}$.

Linear dependencies among the prescribed evaluation functionals correspond to dual words in $RM(n-3, n)$ supported on $P = B \cup \{a, x\} \cup Y$ ($|P| \leq 9$). Because the dual code has minimum distance 8, non-trivial relations are restricted strictly to formal cancellations (weight 0) and geometric relations whose support is an affine 3-flat (weight 8).

Classifying pairs (a, x) by $s = |B \cup \{a, x\}|$, anchor status ($a = 0$), and the affine dimension flag $f = \mathbf{1}_{[\dim \text{aff}(B \cup \{a, x\})=3]}$, we evaluate S'_k exactly in each branch. The baseline configuration contributes $T_G(N)$. Accounting for the small number of boundary configurations ($s \in \{3, 4\}$) and summing over all $N(N/4-1)$ pairs yields $\sum_{a,x} \Pr[F_{a,x} \wedge M_0] \leq R_4(N)$.

Finally, the density bound $1 - R_4(N)/B(N)$ is strictly decreasing for $N \geq 8$ and converges from above to $307/412 > 0.74514$.

Theorem 3 is the main quantitative result: for every $n \geq 3$, every decoder set of three distinct nonzero points and every admissible assignment, τ restricts to an injection

$$\mathcal{D} = M_0 \cap \tau^{-1}(M_1) \hookrightarrow M_1, \quad \frac{|\mathcal{D}|}{\alpha_0} > \frac{307}{412} > 0.745.$$

The bound is uniform: it does not improve with n , and in fact $1 - R_4(N)/B(N)$ decreases to $307/412$ from above.

6 Inverting τ : a one-way function

A one-way function is a statement about a distribution, not only about a map. The input ensemble is the uniform distribution on *all* N^D systems, trivially samplable as nD uniform coefficient bits; on it τ is total and costs $\Theta(n^3)$ (Proposition 2), every output has a zero at $v(S)$ (Proposition 1), and by Lemma 6 inverting τ at an output T — returning any S' with $\tau(S') = T$ — is the same problem as exhibiting a zero of T . So τ is a one-way function precisely under the assumption that zero location is hard on average for the pushforward of the uniform ensemble, which by Lemma 6 draws a solvable system T with probability $X(T)/N^D$, proportional to its number of zeros — the solvable case of MQ. Injectivity holds on $\tau^{-1}(M_1)$, a constant fraction of the domain, on which τ is a bijection onto M_1 with uniform pushforward; no efficient sampler for that subdomain is claimed, since recognising it means deciding $X(T) = 1$. What is samplable is the full domain; the subdomain is where injectivity holds.

This pushforward is a familiar object. The identification and signature schemes built directly on random MQ instances – Sakumoto-Shirai-Hiwatari [13], MQDSS [14] and MQOM [15,16], the last now in the third round of NIST’s additional-signature process [12] – generate keys as $(F, F(s))$ for a uniform system F and a uniform point s , the secret being s , and their security is the hardness of finding a zero of $F - F(s)$. Count the pairs (F, s) with $F - F(s) = T$: F must agree with T up to its constant term and s must be a zero of T , so this instance is drawn with probability proportional to $X(T)$ — up to the normalisation of constant terms, exactly the pushforward of τ . The one-wayness of τ is therefore the planted-MQ assumption of those schemes, over $\mathbb{F}_2$ and in square shape, and $S \mapsto (\tau(S), v(S))$ is an alternative $\Theta(n^3)$ generator of public/secret key pairs for it with the same distribution. What τ adds is the structure that the presentation $(F, F(s))$ hides: the fibre over a key is its zero set (Lemma 6), τ is a bijection onto the uniquely solvable keys (Corollary 1), and the orbit-wise self-reduction of Section 6.1. Concretely, inverting τ on a uniquely solvable output is the input of the best worst-case solvers over $\mathbb{F}_2$, which beat exhaustive search with $2^{0.8765n}$ [20], $2^{0.804n}$ [21] and $2^{0.6943n}$ [22], against the $2^{0.792n}$ of the algebraic method on random systems [24]; the cryptanalytic form of the polynomial method [23] and the MQ estimator [25] give concrete figures. None is polynomial, and we know of none that exploits the square shape or the planted zero.

Proposition 4 (Inverting τ is exactly locating the zero). *Let $T \in M_1$ — beyond this, no promise that $T \in \tau(M_0)$ is needed. Then computing the unique $\tau^{-1}(T)$ and computing the unique zero $z(T)$ are the same problem; explicitly*

$$z(T) = v(\tau^{-1}(T)), \quad \tau^{-1}(T) = T + (v(T) + z(T))\mathbf{1}_{A(z(T))}.$$

If instead one wants the preimage to lie in M_0 — inversion of $\tau|_{\mathcal{D}}$, where $\mathcal{D} = M_0 \cap \tau^{-1}(M_1)$ — the promise $T \in \tau(M_0)$ is exactly what is required, and is decided by Proposition 5 once $z(T)$ is known.

Proof. Every $T \in M_1$ has exactly one τ -preimage, by Corollary 1; the promise is needed only to place it in M_0 . If the zero v is known then by Lemma 6

$$\tau^{-1}(T) = T + (v(T) + v)\mathbf{1}_{A(v)},$$

one flat construction and one polynomial update. (Writing that answer out is $\Theta(nD) = \Theta(n^3)$ coefficients in the dense representation, so $O(n^3)$ is optimal for reasons of output size; we claim no computational lower bound.) Conversely, suppose an algorithm returns some S' with $\tau(S') = T$. By Lemma 6 the fibre of T has exactly $X(T) = 1$ element, so S' is *the* preimage, and $v(S') = S'(b_1) + S'(b_2) + S'(b_3) =$ three evaluations — equals the unique zero of T , since $\tau(S')$ vanishes at $v(S')$.

Proposition 5 (Membership in the successful image $\tau(M_0) \cap M_1$). *Let $T \in M_1$ with unique zero v , and put $c = v(T) + v$. The unique candidate preimage $S_v = T + c\mathbf{1}_{A(v)}$ satisfies $v(S_v) = v$ and*

$$\mathcal{Z}(S_v) = \begin{cases} \{v\}, & c = 0, \\ A(v) \cap T^{-1}(c), & c \neq 0, \end{cases}$$

so $T \in \tau(M_0)$ if and only if

$$c \neq 0 \quad \text{and} \quad A(v) \cap T^{-1}(c) = \emptyset$$

— no point need be excluded from the intersection, since $T(v) = 0 \neq c$.

Proof. By Lemma 6 the unique candidate preimage is $S_v = T + c\mathbf{1}_{A(v)}$, so $T \in \tau(M_0)$ iff $S_v \in M_0$. Off $A(v)$ we have $S_v = T$, which is nonzero there since v is the only zero of T and $v \in A(v)$. At v , $S_v(v) = T(v) + c = c$. At the remaining $x \in A(v)$, $S_v(x) = T(x) + c$. For the decoder value, admissibility gives $v(c\mathbf{1}_{A(v)}) = c$, so $v(S_v) = v(T) + c = v$: the candidate lies in the right decoder slice before any zero-freeness is checked.

6.1 Orbit coverage and a random self-reduction

Write $\mathcal{I} := \tau(M_0) \cap M_1$ for the successful image; since $\tau(M_0) \not\subseteq M_1$, this intersection, and not $\tau(M_0)$, is the set at issue. Let

$$G = \text{Aut}(\text{RM}(2, n)) \times \text{GL}(n, 2), \quad ((\sigma, M) \cdot T)(x) = MT(\sigma^{-1}x),$$

where σ is a coordinate permutation preserving $\text{RM}(2, n)$ — equal to $\text{AGL}(n, 2)$ for $n \geq 4$ and to the exceptional S_8 at $n = 3$ — and M recombines the n output equations. Since M is invertible, $\mathcal{Z}(gT) = \sigma(\mathcal{Z}(T))$, so G preserves every M_k and moves the zero predictably: $z(gT) = \sigma(z(T))$. It does *not* preserve $\mathcal{I}$, which is defined through the fixed B and the fixed flat rule. The group G defines the equivalence between multivariate systems via linear changes of variables and equation recombinations. Its action establishes the relation behind Isomorphism

of Polynomials problem [18] and the analysis of equivalent keys [19]; Theorem 4 and Proposition 6 below thus say that worst-case hardness within an equivalence class implies average-case hardness on it. Put

$$\rho(\mathcal{O}) := \frac{|\mathcal{O} \cap \mathcal{I}|}{|\mathcal{O}|} \quad (\mathcal{O} \text{ a } G\text{-orbit in } M_1),$$

equivalently $\rho(\mathcal{O}) = \Pr_g[gT \in \mathcal{I}]$ for $T \in \mathcal{O}$ and g uniform. A global density says nothing about an individual orbit, so the following is not a formality.

Theorem 4 (Every orbit is covered).

For every $n \geq 3$ and every $T \in M_1$, $\rho(G \cdot T) \geq \frac{3}{4} - \frac{1}{N}$.

Proof. Write $g = (\sigma, M)$, $v = z(T)$, and put $s_\sigma = \sum_i T(\sigma^{-1}(b_i))$ and $A' = \sigma^{-1}(A(\sigma(v)))$, a set of $N/4$ points containing v . Then $z(gT) = \sigma(v)$ and $v(gT) = M(s_\sigma)$, so $c(gT) = M(s_\sigma) + \sigma(v)$; applying M^{-1} to the criterion of Proposition 5 and writing $d = M^{-1}(c(gT)) = s_\sigma + M^{-1}(\sigma(v))$, membership of gT says exactly

$$d \notin \{0\} \cup T(A' \setminus \{v\}),$$

the values $T(y)$, $y \in A' \setminus \{v\}$, being nonzero because v is the only zero of T . Condition on σ . If $\sigma(v) \neq 0$ then $M^{-1}(\sigma(v))$ is uniform on $V \setminus \{0\}$, since $\text{GL}(n, 2)$ is transitive there, so d is uniform on the $N-1$ values $V \setminus \{s_\sigma\}$ and the conditional probability of membership is at least $(N-1-1-|T(A' \setminus \{v\})|)/(N-1) \geq (\frac{3N}{4}-1)/(N-1)$, using only $|T(A' \setminus \{v\})| \leq |A'|-1 = \frac{N}{4}-1$. Discarding the at most $1/N$ fraction of σ with $\sigma(v) = 0$ and averaging, $\rho \geq (1-\frac{1}{N})(\frac{3N}{4}-1)/(N-1) = \frac{3}{4} - \frac{1}{N}$.

Proposition 6 (Random self-reduction). Let $\mathcal{A}$ be a randomised algorithm that halts within $R(n)$ steps on every input in M_1 — no correctness is assumed off $\mathcal{I}$ — and suppose that for every G -orbit $\mathcal{O} \subseteq M_1$,

$$\frac{1}{|\mathcal{O} \cap \mathcal{I}|} \sum_{U \in \mathcal{O} \cap \mathcal{I}} \Pr[\mathcal{A}(U) = z(U)] \geq f.$$

with $f > 0$. Then there is a zero-error randomised algorithm which, on every $T \in M_1$, returns $z(T)$ with probability at least δf per trial, where $\delta = \frac{3}{4} - \frac{1}{N}$ by Theorem 4; one trial has expected cost $R(n) + O(n^4)$ — the group element is drawn by rejection sampling, whose acceptance probability is bounded below by a constant but whose running time is unbounded in the worst case — so the expected total is $O((R(n) + n^4)/(\delta f))$ and r trials fail with probability at most $(1 - \delta f)^r$. Note the hypotheses: a uniform time bound $R(n)$ valid on all of M_1 , and success averaged over $\mathcal{O} \cap \mathcal{I}$ for every orbit $\mathcal{O}$ — not over an orbit entire, nor over M_1 .

The proof of Proposition 6 is in Appendix A.

6.2 Flat-collision detection is NP-complete

Deciding whether a candidate flat carries a collision is a different problem from inverting τ , and the one with a clean completeness statement. Note the polarity: for $S \in M_0$ a collision is exactly the event that the flat *fails*, so certifying success is the complementary problem.

Proposition 7 (The collision system). *Let S be a system, $v = v(S)$, and write $A(v) = v + W$, $\dim W = n - 2$; set $T_i(y) = S_i(v + y) + S_i(v)$ for $y \in W$. Then each T_i has degree at most 2 on $W \cong \mathbb{F}_2^{n-2}$ and $T(0) = 0$, and for $S \in M_0$,*

$$\tau(S) \in M_1 \iff T(y) \neq 0 \text{ for all } y \in W \setminus \{0\}.$$

So certifying the flat is a quadratic feasibility question with n equations in $n - 2$ variables, restricted to nonzero solutions. It is not literally the square MQ decision problem: the shape is $n \times (n - 2)$, and the zero solution is excluded by fiat.

Proof. $y \mapsto v + y$ is affine, so $S_i(v + y)$ has degree at most 2 in y ; adding the constant $S_i(v)$ preserves degree and gives $T_i(0) = 0$. The equivalence is Proposition 3 rewritten in the coordinate $y = x + v$.

An *instance* of FLAT-COLLISION is a tuple $(B, S, v, \ell_1, \ell_2)$ in which $B \subseteq V$ consists of three distinct points, S is a system given by its dense coefficient vector, $v \in V$, and ℓ_1, ℓ_2 are independent linear forms determining the flat

$$A = \{x \in V : \ell_i(x) = \ell_i(v), i = 1, 2\} \ni v.$$

All five components are given succinctly, A by the pair ℓ_1, ℓ_2 and not by its 2^{n-2} points; writing $v_B(S) := \sum_{b \in B} S(b)$ for the decoder value of S with respect to B , the instance is *admissible* when $v = v_B(S)$ and $|A \cap B|$ is odd, which is checkable in polynomial time. The question is whether some $x \in A \setminus \{v\}$ has $S(x) = S(v)$. The proof of the next theorem is in Appendix A.

Theorem 5 (NP-hardness of FLAT-COLLISION). *FLAT-COLLISION is NP-complete. It remains NP-hard on the subclass of instances with*

$$B = B_n := \{0, e_{n-1}, e_n\}, \quad v = e_{n-1}, \quad A = v + W, \quad W = \langle e_1, \dots, e_{n-2} \rangle,$$

satisfying

- (i) S is solution-free;
- (ii) $v = v_{B_n}(S)$;
- (iii) A is admissible for B_n , indeed $A \cap B_n = \{v\}$;
- (iv) the restricted collision system has the form $T_S(y) = (Q(y), 0, 0)$ for a square Q whose coordinates are formally homogeneous quadratics.

The reduction is parsimonious: the collisions on $A \setminus \{v\}$ are in bijection with the satisfying assignments of the source circuit. Consequently counting the collisions on a given flat is #P-complete.

Conditions (ii) and (iii) are decided by inspection; (i) is not efficiently checkable and (iv) is a property of the reduction’s output, so these two are promises in the technical sense. Since B_n contains 0, the subclass lies outside the decoders of Theorem 3: the hardness of FLAT-COLLISION concerns the problem, not the canonical instance.

7 Discussion

The construction partially answers the question left open by the counting proof of [1]: an injection into M_1 can be written down on an explicit subset of M_0 of density more than 74.5%, and is computable in time linear in the instance size. Three directions seem worth pursuing.

Density. The bound 307/412 is uniform but not tight (0.857 at $n = 3$): the truncations of Section 5 use nothing about the flat rule beyond $|A(a)| = N/4$ and $a \in A(a)$, so they bound a worst case over all rules, and the fourth-order truncation is itself inexact. Whether the density has a limit as $n \rightarrow \infty$, and what it is, we do not know.

Totality. No injection defined on all of M_0 is given, and [1] guarantees only that one exists; a totalisation of τ would have to treat the systems on which the flat fails, and it is not clear that a single-flat correction suffices.

One-wayness. Section 6 gives the properties one wants of a candidate one-way function in an unusually clean form: inversion is *equivalent* to zero location on the image, the map is injective on its canonical domain, with the pushforward of the uniform distribution on that domain uniform on the image, and it is randomly self-reducible on each orbit of G . Few explicit candidates combine injectivity with self-reducibility of this kind, and it would be interesting to know whether τ or a variant supports the constructions — pseudorandom generators, commitment schemes — for which injective one-way functions are the natural primitive.

What is missing is a worst-case/average-case connection *across* orbits: Theorem 4 controls each orbit separately, and the number of orbits grows with n .

Acknowledgements

These results are related the first author’s MSC thesis, supervised by the second author. Several of the numerical claims were checked by computer algebra with the assistance of Anthropic Claude and OpenAI ChatGPT, which also reviewed the arguments and assisted with some proofs.

A Deferred proofs

Proof of Lemma 4.

Proof. As b_1, b_2, b_3 are distinct so are d_1, d_2, d_3 , at most one is 0, and the three cases are exhaustive: if all are nonzero and distinct then $\dim\langle d_1, d_2, d_3 \rangle \geq 2$. In each case the vectors carrying prescribed values are independent, so a solution exists and Convention 1 selects exactly one. Case (3) uses $n \geq 3$ so that z exists. Since $b_j \in A(a)$ exactly when $d_j \in W$, the prescribed values show directly that precisely one of d_1, d_2, d_3 lies in W in every case, which is the admissibility claim $|A(a) \cap B| = 1$. Each elimination acts on an augmented matrix of at most three rows and $n + 1$ columns, costing $O(n)$ per form; forming the d_i costs $O(n)$, and locating z is $O(n)$ because $\langle d_1, d_2 \rangle$ has only four elements, so the lexicographically least vector outside it is among the first five. The number of prescribed constraints is at most three, independent of n , so one scans coordinates for pivots, back-substitutes, and sets free coefficients to zero, each step touching $O(n)$ bits. The matching lower bound is forced by the input, not by the output: the forms produced are supported on the pivot columns, hence have weight at most three, but on the inputs $B = \{0, e_j, e_1 + e_j\}$, $a = e_j$ ($2 \leq j \leq n$) one is in case (1) with $r_1 = e_j$, $r_2 = e_1$ and $\ell_1 = e_j$, so the output identifies j , which no procedure can do without reading $\Omega(n)$ bit positions in the worst case. Thus $\Theta(n)$ is tight for this subroutine. It is in any case dominated by the $\Theta(n^3)$ update of Proposition 2.

The second-order bound below is the template for the proof of Theorem 3. Put

$$R(N) = \frac{(N-4)(N-1)(N+2)}{8N^3} + \frac{6}{N^4}.$$

Lemma 10 (Conditioned second-order hitting bound). *For every $n \geq 3$, every set B of three distinct nonzero points and every assignment of $(n-2)$ -flats $A(a) \ni a$, admissible or not,*

$$\frac{\#\{S \in M_0 : \tau(S) \in M_1\}}{\alpha_0} \geq 1 - \frac{R(N)}{B(N)} > \frac{68}{103} = 0.66019\dots$$

Proof. For $a \in V$ and $x \in A(a) \setminus \{a\}$ put $F_{a,x} = \{v(S) = a, S(x) = S(a)\}$; then $Z = \sum_a \sum_x \mathbf{1}_{F_{a,x}}$ pointwise. By Proposition 3, failure on M_0 is equivalent to $Z \geq 1$, so

$$\begin{aligned} \#\{S \in M_0 : \tau(S) \notin M_1\} &= \#\{S \in M_0 : Z \geq 1\} \leq \\ &\leq \sum_{S \in M_0} Z(S) = N^D \sum_{a,x} \Pr[F_{a,x} \wedge S \in M_0]. \end{aligned}$$

Neither this nor anything below uses admissibility of the assignment: only $|A(a)| = N/4$ and $a \in A(a)$ enter. Unlike in Theorem 2, the event $S \in M_0$

is retained. Fix (a, x) and write $E_y = \{S(y) = 0\}$; Bonferroni's second-order inequality, applied inside the sub-measure $\Pr[F_{a,x} \cap \cdot]$, gives

$$\Pr[F \wedge M_0] \leq \Pr[F] - \sum_{y \in V} \Pr[F \wedge E_y] + \sum_{\{y, y'\} \subseteq V} \Pr[F \wedge E_y \wedge E_{y'}],$$

and every term on the right is *exact*. Each event prescribes values of the functionals $\varphi = \text{ev}_{b_1} + \text{ev}_{b_2} + \text{ev}_{b_3}$ (value a), $\psi = \text{ev}_x + \text{ev}_a$ (value 0), and at most two evaluations $\text{ev}_y, \text{ev}_{y'}$ (value 0), of total support at most $3 + 2 + 1 + 1 = 7 < 8$. By Lemma 3(b) no combination vanishes through a dual word, so a relation must cancel identically, and by Lemma 2 each probability is $N^{-\rho}$ with ρ the $\mathbb{F}_2$ -rank of the functionals involved, or 0 when the prescription is inconsistent. Considering the maximal four-functional family $\{\varphi, \psi, \text{ev}_y, \text{ev}_{y'}\}$ there are 15 nonempty combinations, and its subfamilies cover the lower-order terms. A relation without φ can cancel only if ψ is accompanied by both ev_a and ev_x . A relation containing φ but not ψ would need the evaluation supports to XOR to B itself, impossible with two points; one containing both needs them to XOR to $B \triangle \{a, x\}$, of size $5 - 2|B \cap \{a, x\}|$, which two points achieve only when $a, x \in B$, leaving the third point of B . Exactly two kinds of combination can therefore vanish (the first with y or with y'):

$$\varphi + \psi + \text{ev}_y = 0 \iff \{a, x, y\} = B, \quad \psi + \text{ev}_y + \text{ev}_{y'} = 0 \iff \{y, y'\} = \{a, x\}.$$

The first carries the prescribed value $a \neq 0$ (as $a \in B$ and $0 \notin B$), so every event containing it is empty; the second carries 0 and is consistent. Hence, with $s = |\{a, x\} \cup B|$: $\Pr[F] = N^{-2}$ (Lemma 9); $\Pr[F \wedge E_y] = N^{-3}$ for every y , except 0 at the single y with $\{a, x, y\} = B$, possible only when $s = 3$; and $\Pr[F \wedge E_y \wedge E_{y'}] = N^{-4}$ for every pair, except N^{-3} at $\{y, y'\} = \{a, x\}$ (the rank-3 event $\{\varphi(S) = a, S(a) = S(x) = 0\}$) and 0 at the $N - 1$ pairs through the forbidden point when $s = 3$. The truncation is therefore

$$N^{-2} - N \cdot N^{-3} + \left[\binom{N}{2} - 1 \right] N^{-4} + N^{-3} = \frac{(N-1)(N+2)}{2N^4}$$

for $s \in \{4, 5\}$, and exceeds this by exactly N^{-4} when $s = 3$; that case forces $a, x \in B$ and so occurs for at most six pairs (a, x) . Summing over the $N(N/4 - 1)$ pairs,

$$\sum_{a,x} \Pr[F_{a,x} \wedge M_0] \leq N \left(\frac{N}{4} - 1 \right) \frac{(N-1)(N+2)}{2N^4} + \frac{6}{N^4} = R(N),$$

and dividing by $\alpha_0 \geq B(N)N^D$ (Lemma 7) gives the first inequality. For the constant, clearing denominators gives (for $N \geq 8$)

$$\begin{aligned} \frac{35}{103} B(N) - R(N) &= \\ &= \frac{773N^5 + 1784N^4 - 2472N^3 - 14951N^2 + 154N - 120}{2472N^6} > 0 \end{aligned}$$

since $773N^5 \geq 49472N^3 > 2472N^3$, $1784N^4 \geq 114176N^2 > 14951N^2$ and $154N > 120$; hence $R(N) < \frac{35}{103}B(N)$ and the bound exceeds $1 - \frac{35}{103} = \frac{68}{103}$.

Proof of Theorem 3.

Proof. As in the proof of Lemma 10, retaining $S \in M_0$ and, as there, using only $|A(a)| = N/4$ and $a \in A(a)$ of the assignment, it suffices to bound $\sum_{a,x} \Pr[F_{a,x} \wedge M_0]$ and divide by $\alpha_0 \geq B(N)N^D$. Fix a pair (a, x) , put $z := X(S) = |\mathcal{Z}(S)|$ and $E_y = \{S(y) = 0\}$. The even Bonferroni truncation $\mathbf{1}_{[z=0]} \leq \sum_{k=0}^4 (-1)^k \binom{z}{k}$ gives

$$\Pr[F \wedge M_0] \leq P_0 - S'_1 + S'_2 - S'_3 + S'_4, \quad S'_k = \sum_{|Y|=k} \Pr\left[F \wedge \bigcap_{y \in Y} E_y\right], \quad (1)$$

with $P_0 = \Pr[F] = N^{-2}$. The signs fix the direction of every estimate: for an upper bound the negatively-signed terms S'_1, S'_3 must be bounded *below* and the positively-signed S'_2, S'_4 *above*; we evaluate all four exactly, so the question does not arise.

Relation structure. An event in (1) prescribes values of the functionals $\varphi = \text{ev}_{b_1} + \text{ev}_{b_2} + \text{ev}_{b_3}$ (value a), $\psi = \text{ev}_x + \text{ev}_a$ (value 0) and ev_y , $y \in Y$ (value 0), whose supports lie in $P = B \cup \{a, x\} \cup Y$, a set of at most 9 points. By Lemma 3 a subset of the functionals is a relation exactly when the XOR of its supports is a word of $\text{RM}(n-3, n)$; since the support of any relation lies in P it has weight at most 9; by Lemma 3(c) a nonzero such support has weight at least 8, and by Lemma 3(a) it has even weight, so it is either empty or of weight exactly 8. This leaves just two kinds.

Exact cancellations (support-XOR = 0): $\psi + \text{ev}_a + \text{ev}_x$ (value 0, consistent); $\varphi + \psi + \text{ev}_y$ with $\{a, x, y\} = B$; $\varphi + \sum_{y \in B} \text{ev}_y$; and $\varphi + \psi + \sum_{y \in Y'} \text{ev}_y$ with $Y' = B \triangle \{a, x\}$, which needs $s := |B \cup \{a, x\}| = 4$. *Geometric relations* (support-XOR a weight-8 word, i.e. the indicator of an affine 3-flat): a combination using ε_φ copies of φ , ε_ψ of ψ and $Y' \subseteq Y$ has support of size at most $3\varepsilon_\varphi + 2\varepsilon_\psi + |Y'|$ and of that parity, so for $|Y| \leq 4$ the even weight 8 is reached only by $\varepsilon_\varphi = \varepsilon_\psi = 1$, $|Y'| = 3$, all eight points distinct. Hence $s = 5$, the flat is the affine span of $B \cup \{a, x\}$, and there is at most *one* geometric relation per pair: it exists iff $\dim \text{aff}(B \cup \{a, x\}) = 3$ — write $f \in \{0, 1\}$ for this flag — and fires on those Y containing the three completing points. There is indeed at most one: if $u_1 \neq u_2$ were weight-8 dual words supported in P , then $u_1 \triangle u_2$ would be a nonzero dual word supported in P , hence of weight at most $|P| \leq 9$, at least 8 by dual distance and even, so exactly 8; but $8 = |u_1 \triangle u_2| = 16 - 2|u_1 \cap u_2|$ forces $|u_1 \cap u_2| = 4$ and therefore $|u_1 \cup u_2| = 12 > 9$, a contradiction.

Every relation containing φ carries the prescribed value a : it annihilates the event when $a \neq 0$ and lowers the rank by one when $a = 0$. Two independent consistent relations force $Y \supseteq \{a, x\}$ and $Y \supseteq B$ with $a = 0$, hence $x \in B$ and $Y = B \cup \{0\}$; consistent deficiency never exceeds 2.

The case tables. Write $C_j = \binom{N}{j}$ and classify (a, x) by s , by whether $a = 0$, and by f . At orders 1 and 2 the anchor plays no role, the value- a patterns first

fitting at $|Y| = 3$:

$$S'_1 = \begin{cases} N \cdot N^{-3}, & s \in \{4, 5\}, \\ (N-1)N^{-3}, & s = 3, \end{cases} \quad S'_2 = \begin{cases} (C_2 - 1)N^{-4} + N^{-3}, & s \in \{4, 5\}, \\ (C_2 - N)N^{-4} + N^{-3}, & s = 3. \end{cases}$$

At orders 3 and 4, with $C_{22} = \binom{N-2}{2}$,

$$S'_3 = \begin{cases} (N-2)N^{-4} + (C_3 - (N-2) - 1 - f)N^{-5}, & s = 5, a \neq 0, \\ (N-1+f)N^{-4} + (C_3 - (N-1) - f)N^{-5}, & s = 5, a = 0, \\ (N-2)N^{-4} + (C_3 - N)N^{-5}, & s = 4, a \neq 0, \\ N \cdot N^{-4} + (C_3 - N)N^{-5}, & s = 4, a = 0, \\ (N-3)N^{-4} + (C_3 - \binom{N-1}{2} - (N-3))N^{-5}, & s = 3, \end{cases}$$

$$S'_4 = \begin{cases} C_{22}N^{-5} + (C_4 - C_{22} - (1+f)(N-3))N^{-6}, & s = 5, a \neq 0, \\ (C_{22} + (1+f)(N-3))N^{-5} + \\ \quad + (C_4 - C_{22} - (1+f)(N-3))N^{-6}, & s = 5, a = 0, \\ (C_{22} - 1)N^{-5} + (C_4 - C_{22} - 2N + 8)N^{-6}, & s = 4, a \neq 0, \\ N^{-4} + (C_{22} - 1 + 2(N-4))N^{-5} + \\ \quad + (C_4 - C_{22} - 2N + 8)N^{-6}, & s = 4, a = 0, \\ \binom{N-3}{2}N^{-5} + (C_4 - \binom{N-1}{3} - \binom{N-3}{2})N^{-6}, & s = 3. \end{cases}$$

The N^{-4} terms of S'_3 are the triples $Y \supseteq \{a, x\}$ carrying the ψ -relation, together with, when $a = 0$, the consistent value- a relations ($Y = B$; the completing triple when $f = 1$; and $Y = B \triangle \{a, x\}$ when $s = 4$); for $a \neq 0$ those same sets are annihilated instead. In S'_4 with $s = 4$, $a \neq 0$ the quadruple $B \cup \{a\}$ resp. $B \cup \{x\}$ lies in all three special families and dies, producing the -1 and the $2N - 7$ dead quadruples; for $s = 4$, $a = 0$ it is instead the unique deficiency-2 event, contributing N^{-4} .

Worst case and budgets. Let $T_4(\text{case})$ be the right-hand side of (1). Direct computation gives $T_4(s=5, a \neq 0, f=1) = T_G(N)$ and, for $N \geq 8$,

$$T_G - T_4(s=5, a \neq 0, f=0) = \frac{3}{N^6}, \quad T_G - T_4(s=4, a \neq 0) = \frac{N-2}{N^6},$$

$$T_G - T_4(s=5, a=0) = \frac{3(N+1)}{N^6} \text{ or } \frac{6}{N^5}$$

(the last according as $f = 0$ or $f = 1$), so the only cases that can exceed T_G are $s = 3$ and $s = 4$ with $a = 0$:

$$T_4(s=3) - T_G = \frac{2N^3 + 3N^2 + 13N - 48}{6N^6}, \quad T_4(s=4, a=0) - T_G = \frac{N^2 - 9N + 2}{N^6}.$$

(The second difference is negative at $N = 8$, so that case is not in fact exceptional at $n = 3$; the budget below bounds it by a positive quantity and so remains valid.)

Now $s = 3$ forces $a, x \in B$, so the number of such pairs is $q_3 = \#\{(a, x) : a, x \in B, a \neq x, x \in A(a)\} \leq 6$; and $s = 4$ with $a = 0$ forces $x \in A(0) \cap B$, a set of size $q_0 \leq 3$. (For the construction of Lemma 4 one has $|A(a) \cap B| = 1$ at every anchor, so $A(a) \cap B = \{a\}$ when $a \in B$ and therefore $q_3 = 0$ and $q_0 = 1$: the canonical rule satisfies a strictly sharper finite- N bound than the universal statement below, which uses only $q_3 \leq 6$ and $q_0 \leq 3$.) Since $6(T_4(s=3) - T_G) \leq 3N^{-3}$ (equivalently $N^3 - 3N^2 - 13N + 48 > 0$, true for $N \geq 8$) and $3(T_4(s=4, a=0) - T_G) \leq 3N^{-4}$ (equivalently $9N > 2$), summing over the $N(N/4 - 1)$ pairs gives $\sum_{a,x} \Pr[F_{a,x} \wedge M_0] \leq R_4(N)$.

The constant. Clearing denominators (for $N \geq 8$),

$$\begin{aligned} \frac{105}{412}B(N) - R_4(N) &= \\ &= \frac{876N^5 + 1475N^4 - 15450N^3 - 1767N^2 + 10042N - 120}{3296N^6} > 0 \end{aligned}$$

since $876N^5 \geq 56064N^3 > 15450N^3$, $1475N^4 \geq 1475 \cdot 64N^2 > 1767N^2$ and $10042N > 120$. Hence the bound exceeds $1 - \frac{105}{412} = \frac{307}{412}$. Finally $N^2T_G(N) \rightarrow \frac{3}{8}$, so $R_4(N) \rightarrow \frac{1}{4} \cdot \frac{3}{8} = \frac{3}{32}$ and $1 - R_4/B \rightarrow 1 - \frac{3}{32} \cdot \frac{280}{103} = \frac{307}{412}$. The approach is monotone: differentiating the rational function gives

$$\frac{d}{dN} \left(1 - \frac{R_4(N)}{B(N)} \right) = \frac{-105 P(N)}{2(N^2 - 1)^2 (309N^4 - 154N^3 + 239N^2 - 154N + 120)^2},$$

$$\begin{aligned} P(N) &= 1314N^{10} + 4425N^9 - 70330N^8 + 12498N^7 + 86046N^6 - 30823N^5 \\ &\quad - 19470N^4 + 26700N^3 - 31480N^2 - 1920N + 5760, \end{aligned}$$

and $P(N) > 0$ for $N \geq 8$: indeed $1314N^{10} \geq 1314 \cdot 64N^8 > 70330N^8$ and $4425N^9 > 30823N^5$, while

$$86046N^6 > 19470N^4 + 31480N^2 + 1920N \quad (N \geq 8),$$

since after division by N^4 this reads $86046N^2 > 19470 + 31480N^{-2} + 1920N^{-3}$, which holds already at $N = 8$; the remaining coefficients are positive. Hence the bound is strictly decreasing in N and approaches $\frac{307}{412}$ from above.

The case tables above are generated by a single rule: for each pair (a, x) one records $s = |B \cup \{a, x\}|$, whether $a = 0$, and whether the affine span of $B \cup \{a, x\}$ has dimension 3 (the flag f); these three data determine which of the relations listed above are present, hence the rank of each of the four event families, hence T_4 . At $n = 3$, where any five points lie in a 3-flat so that the rows with $f = 0$ cannot occur, every other row was confirmed by exhaustive enumeration over all 2^{21} systems with x ranging over $V \setminus \{a\}$; all rows were confirmed against exact rank computations at $n = 4, 5$.

Lemma 11 (The universal constant). *For every $N \geq 8$,*

$$B(N) > \frac{103}{280} - \frac{103}{70N},$$

and consequently

$$\frac{\frac{1}{4} - \frac{1}{N}}{B(N)} < \frac{70}{103}, \quad \text{i.e.} \quad 1 - \frac{\frac{1}{4} - \frac{1}{N}}{B(N)} > \frac{33}{103}.$$

The numerator $\frac{1}{4} - \frac{1}{N}$ is exactly the first moment of Lemma 9. The sharper thresholds of Lemma 10 and Theorem 3 are not obtained from this lemma but from exact polynomial inequalities in their own proofs.

Proof of Lemma 11.

Proof. Expanding the binomials,

$$\begin{aligned} B(N) &= \frac{309N^6 - 154N^5 - 70N^4 - 119N^2 + 154N - 120}{840N^6} \\ &= \frac{103}{280} - \frac{11}{60N} - \frac{1}{12N^2} - \frac{17}{120N^4} + \frac{11}{60N^5} - \frac{1}{7N^6}. \end{aligned}$$

Hence $B(N) - \left(\frac{103}{280} - \frac{103}{70N}\right)$ equals $1/N$ times

$$\frac{103}{70} - \frac{11}{60} - \frac{1}{12N} - \frac{17}{120N^3} + \frac{11}{60N^4} - \frac{1}{7N^5} = \frac{541}{420} - \frac{1}{12N} - \frac{17}{120N^3} + \frac{11}{60N^4} - \frac{1}{7N^5}.$$

Discarding the positive term $\frac{11}{60N^4}$ and bounding the three negative ones at $N = 8$, the bracket is at least

$$\frac{541}{420} - \frac{1}{96} - \frac{17}{61440} - \frac{1}{229376} > 1.27 > 0 \quad (N \geq 8),$$

which proves the first display. Since $\frac{103}{280} - \frac{103}{70N} = \frac{103}{280} \left(1 - \frac{4}{N}\right) > 0$ for $N \geq 8$, the same display gives $B(N) > 0$. Multiplying it by $\frac{70}{103} > 0$,

$$\frac{70}{103} B(N) > \frac{70}{103} \left(\frac{103}{280} - \frac{103}{70N} \right) = \frac{1}{4} - \frac{1}{N},$$

and dividing by $B(N) > 0$ gives $\left(\frac{1}{4} - \frac{1}{N}\right)/B(N) < \frac{70}{103}$; subtracting from 1 yields the stated $\frac{33}{103}$ bound.

Proof of Proposition 6.

Proof. We first record the cost. For $n \geq 4$ we take the coordinate factor to be $\text{AGL}(n, 2)$, sampled as a uniform $A \in \text{GL}(n, 2)$ — by rejection from uniform $n \times n$ matrices, which succeeds with probability $\prod_{i \geq 1} (1 - 2^{-i}) > 0.288$ — together with a uniform translation; for $n = 3$ the automorphism group is S_8 and a uniform permutation of the eight points is drawn directly. Applying $M \in \text{GL}(n, 2)$ to the equation tuple costs $\Theta(n \cdot n \cdot D) = \Theta(n^4)$, and substituting an affine change of variable into each of n quadratics costs $\Theta(n^3)$ apiece, so $\Theta(n^4)$ in all; inverting σ is $O(n^3)$. Both $\text{GL}(n, 2)$ factors — the linear part of σ

and the recombination M — are sampled independently by this rejection procedure. Note that what is assumed is a *zero-finding* algorithm; by Proposition 4 an inverter yields one at $O(n^3)$ extra cost, subject to the same totality requirement. The totality hypothesis is not cosmetic: a trial applies $\mathcal{A}$ to gT , which lies outside $\mathcal{I}$ with probability up to $1 - \delta$, and verification below detects wrong answers but cannot detect non-termination; a promise algorithm must therefore be run under a timeout of $R(n)$ steps. Now sample $g = (\sigma, M) \in G$ uniformly; since every orbit point has exactly $|\text{Stab}(T)|$ preimages under $g \mapsto gT$, the image gT is uniform on $G \cdot T$, and $gT \in \mathcal{I}$ with probability at least δ . Conditioned on that event gT is uniform on $G \cdot T \cap \mathcal{I}$, which is why the second hypothesis is stated per orbit: a success density f measured over the whole image would permit total failure on this particular orbit. Run the algorithm on gT to obtain a candidate zero w . Recombination of equations leaves zero sets unchanged and the coordinate factor acts by $\mathcal{Z}(gT) = \sigma(\mathcal{Z}(T))$, so $\sigma^{-1}(w)$ is the candidate zero of T ; verify $T(\sigma^{-1}(w)) = 0$ in $\Theta(n^3)$ and repeat on failure.

Lemma 12 (Square homogeneous encoding of Circuit-SAT). *Given a Boolean circuit C with r inputs and g gates, one constructs in polynomial time a square system $Q = (q_1, \dots, q_k) : \mathbb{F}_2^k \rightarrow \mathbb{F}_2^k$ with $k = r + g + 1$, each coordinate homogeneous quadratic as a formal polynomial, such that C is satisfiable iff $Q(z) = 0$ has a nonzero solution, and every nonzero solution has final coordinate $t = 1$. The correspondence is moreover a bijection:*

$$\#(Q^{-1}(0) \setminus \{0\}) = \# \text{SAT}(C).$$

Homogeneity is meant in the unreduced representation. As Boolean functions the coordinates are reduced modulo $x^2 - x$ and are then square-free of degree at most two but not homogeneous — $x_i^2 + x_i t$ represents $x_i + x_i t$ — which is all that is needed: each lies in $\text{RM}(2, k)$, and $Q(0) = 0$ either way.

Proof. By the standard polynomial-time normalisation we may assume C is an acyclic circuit over unary NOT and binary AND, OR, with one designated output wire. Take one variable per input wire, one per gate-output wire, and a homogenizing variable t ; thus $k = r + g + 1$. All displayed polynomials below are formally homogeneous of degree two; their associated Boolean functions have degree at most two after multilinear reduction, which is what places them in $\text{RM}(2, k)$. For each input x_i : $x_i^2 + x_i t = 0$. For a gate with inputs a, b and output c :

$$\text{NOT} : c^2 + a^2 + t^2 = 0, \quad \text{AND} : c^2 + ab = 0, \quad \text{OR} : c^2 + a^2 + b^2 + ab = 0.$$

For the output o : $o^2 + t^2 = 0$. This is $r + g + 1 = k$ equations in k variables.

If $t = 1$, then since $u^2 = u$ over $\mathbb{F}_2$ each anchoring equation is automatic, each gate equation becomes the multilinear truth-table identity for its gate ($c = 1 + a$ for NOT, which is unary; $c = ab$; $c = a + b + ab$), and the output equation gives $o = 1$. Both directions of the bijection are then immediate: given an input assignment, acyclicity and the gate equations determine every gate-output variable uniquely,

so the assignment extends to exactly one $t = 1$ zero; conversely every $t = 1$ zero satisfies the truth-table equation at each gate and has designated output 1, so it restricts to a satisfying assignment. Whence the displayed count, the zero vector being the only $t = 0$ solution. If $t = 0$, each anchoring equation gives $x_i = 0$, and processing gates in topological order, every displayed equation reduces to $c^2 = 0$ forcing $c = 0$; so the only $t = 0$ solution is 0. Hence a nonzero zero exists iff a $t = 1$ zero exists iff C is satisfiable.

Proof of Proposition 5.

Proof. Membership in NP: a certificate is an ambient point $x \in \mathbb{F}_2^n$; one verifies $x \neq v$, the two affine equations $\ell_i(x) = \ell_i(v)$ defining A , and the equality $S(x) = S(v)$, each in polynomial time in the instance encoding fixed above.

Hardness: given a circuit C , let $Q : \mathbb{F}_2^k \rightarrow \mathbb{F}_2^k$ be the system of Lemma 12. Put $n = k + 2$ and write points of $V = \mathbb{F}_2^n$ as (y, z_1, z_2) with $y \in \mathbb{F}_2^k$. Define

$$S(y, z_1, z_2) = (Q(y), 1, 0) \in \mathbb{F}_2^n.$$

Each coordinate has degree ≤ 2 , so $S \in \text{RM}(2, n)^n$. The $(k + 1)$ st coordinate is constantly 1, so S never vanishes: condition (i). Here $v = (0, \dots, 0, 1, 0) = e_{n-1}$ has y -block 0, and $Q(0) = 0$ by homogeneity, so $S(0) = S(v) = S(e_n) = v$; three equal vectors sum to themselves over $\mathbb{F}_2$, giving $v_{B_n}(S) = S(0) + S(e_{n-1}) + S(e_n) = v$ for the decoder $B_n = \{0, e_{n-1}, e_n\}$: condition (ii). With $W = \langle e_1, \dots, e_{n-2} \rangle$ the flat $A = v + W$ meets B_n exactly in v , so it is admissible for B_n : condition (iii). Finally, for $y \in \mathbb{F}_2^k$ identified with $(y, 0, 0) \in W$, we have $v + y = (y, 1, 0)$ and

$$T_S(y) = S(y, 1, 0) + S(0, 1, 0) = (Q(y), 1, 0) + (Q(0), 1, 0) = (Q(y), 0, 0),$$

giving condition (iv). Hence a nonzero $y \in W$ with $T_S(y) = 0$ exists iff Q has a nonzero zero iff C is satisfiable; more precisely $y \mapsto (y, 1, 0)$ is a bijection from the nonzero zeros of Q onto the collisions in $A \setminus \{v\}$, and by Lemma 12 the former are in bijection with the satisfying assignments of C , which gives the parsimony claim. Counting collisions is in $\#P$ (a collision is a point of $A \setminus \{v\}$, verifiable in polynomial time), and $\#Circuit\text{-}SAT$ is $\#P$ -complete under parsimonious reductions, so the counting problem is $\#P$ -complete. The reduction outputs the instance $(B_n, S, v, \ell_1, \ell_2)$ with ℓ_1, ℓ_2 the coordinate forms dual to e_{n-1}, e_n , so that $A = v + W$; it is polynomial, since $n = k + 2$, each coordinate of S is a coordinate of Q or a constant, and B_n, v and the two forms have $O(n)$ -bit descriptions, so even materialising the dense coefficient representation of S costs only $O(n^3)$ bits.

References

1. Radici, G., Sala, M.: A note on binary quadratic systems and their relation to complexity theory. Preprint, arXiv:2609.07769 (2026)

2. Fraenkel, A.S., Yesha, Y.: Complexity of problems in games, graphs and algebraic equations. *Discrete Applied Mathematics* **1**(1–2), 15–30 (1979)
3. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, San Francisco (1979)
4. MacWilliams, F.J., Sloane, N.J.A.: *The Theory of Error-Correcting Codes*. North-Holland, Amsterdam (1977)
5. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: *Advances in Cryptology — EUROCRYPT 1999*. LNCS, vol. 1592, pp. 206–222. Springer (1999)
6. Beullens, W.: MAYO: practical post-quantum signatures from oil-and-vinegar maps. In: *Selected Areas in Cryptography — SAC 2021*. LNCS, vol. 13203, pp. 355–376. Springer (2022)
7. Ding, J., Schmidt, D.: Rainbow, a new multivariable polynomial signature scheme. In: *Applied Cryptography and Network Security — ACNS 2005*. LNCS, vol. 3531, pp. 164–175. Springer (2005)
8. Beullens, W.: Breaking Rainbow takes a weekend on a laptop. In: *Advances in Cryptology — CRYPTO 2022, Part II*. LNCS, vol. 13508, pp. 464–479. Springer (2022)
9. Ding, J., Petzoldt, A., Schmidt, D.S.: *Multivariate Public Key Cryptosystems*, 2nd edn. *Advances in Information Security*, vol. 80. Springer (2020)
10. Bernstein, D.J., Lange, T.: Post-quantum cryptography. *Nature* **549**, 188–194 (2017)
11. National Institute of Standards and Technology: FIPS 203, Module-Lattice-Based Key-Encapsulation Mechanism Standard; FIPS 204, Module-Lattice-Based Digital Signature Standard; FIPS 205, Stateless Hash-Based Digital Signature Standard (2024)
12. Moody, D., et al.: Status report on the second round of the additional digital signature schemes for the NIST post-quantum cryptography standardization process. NIST Internal Report 8610 (2026)
13. Sakumoto, K., Shirai, T., Hiwatari, H.: Public-key identification schemes based on multivariate quadratic polynomials. In: *Advances in Cryptology — CRYPTO 2011*. LNCS, vol. 6841, pp. 706–723. Springer (2011)
14. Chen, M.-S., Hülsing, A., Rijneveld, J., Samardjiska, S., Schwabe, P.: From 5-pass $\mathcal{MQ}$ -based identification to $\mathcal{MQ}$ -based signatures. In: *Advances in Cryptology — ASIACRYPT 2016, Part II*. LNCS, vol. 10032, pp. 135–165. Springer (2016)
15. Benadjila, R., Feneuil, T., Rivain, M.: MQ on my mind: post-quantum signatures from the non-structured multivariate quadratic problem. In: *9th IEEE European Symposium on Security and Privacy — EuroS&P 2024*, pp. 468–485. IEEE (2024)
16. Benadjila, R., Bouillaguet, C., Feneuil, T., Rivain, M.: MQOM: MQ on my Mind — algorithm specifications and supporting documentation, version 2.1. Submission to the NIST additional digital signature process (2025). <https://mqom.org>
17. Berbain, C., Gilbert, H., Patarin, J.: QUAD: a practical stream cipher with provable security. In: *Advances in Cryptology — EUROCRYPT 2006*. LNCS, vol. 4004, pp. 109–128. Springer (2006)
18. Patarin, J.: Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): two new families of asymmetric algorithms. In: *Advances in Cryptology — EUROCRYPT 1996*. LNCS, vol. 1070, pp. 33–48. Springer (1996)
19. Wolf, C., Preneel, B.: Equivalent keys in HFE, C^* , and variations. In: *Progress in Cryptology — Mycrypt 2005*. LNCS, vol. 3715, pp. 33–49. Springer (2005)

20. Lokshtanov, D., Paturi, R., Tamaki, S., Williams, R., Yu, H.: Beating brute force for systems of polynomial equations over finite fields. In: Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms — SODA 2017, pp. 2190–2202. SIAM (2017)
21. Björklund, A., Kaski, P., Williams, R.: Solving systems of polynomial equations over $\text{GF}(2)$ by a parity-counting self-reduction. In: 46th International Colloquium on Automata, Languages, and Programming — ICALP 2019. LIPIcs, vol. 132, pp. 26:1–26:13. Schloss Dagstuhl (2019)
22. Dinur, I.: Improved algorithms for solving polynomial systems over $\text{GF}(2)$ by multiple parity-counting. In: Proceedings of the 32nd ACM-SIAM Symposium on Discrete Algorithms — SODA 2021, pp. 2550–2564. SIAM (2021)
23. Dinur, I.: Cryptanalytic applications of the polynomial method for solving multivariate equation systems over $\text{GF}(2)$. In: Advances in Cryptology — EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 374–403. Springer (2021)
24. Bardet, M., Faugère, J.-C., Salvy, B., Spaenlehauer, P.-J.: On the complexity of solving quadratic Boolean systems. *Journal of Complexity* **29**(1), 53–75 (2013)
25. Bellini, E., Makarim, R.H., Sanna, C., Verbel, J.: An estimator for the hardness of the MQ problem. In: Progress in Cryptology — AFRICACRYPT 2022. LNCS, vol. 13503, pp. 323–347. Springer (2022)