

A rigorous success regime for the Radical Attack under uniform accepted-instance sampling

Yuxuan Zhang

Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland

Department of Physics, Princeton University
Princeton, New Jersey 08544, USA

Version 1
22 October 2026

Based on candidate M12-v5. Proof and primary-source comparisons received separate-context internal AI review; computational checks were replayed. No external expert review or formal verification is claimed.

Abstract

We analyze the Radical Attack of Gross and Hangleiter for an exactly specified uniform accepted-instance block ensemble over $\mathbb{F}_2$. The analysis retains both full-rank conditioning and the requirement that the planted support vector lie in the column space. We show that the attack’s failure probability is at most $2m_12^{-d} + m_12^{-w} + 2^{1-w}$ for sufficiently large admissible dimensions. Thus it succeeds with probability $1 - o(1)$ when both $d - \log_2 m_1$ and $w - \log_2 m_1$ tend to infinity. The proof identifies the exact conditional lower-block law and shows that upper-block acceptance gives no hidden code-dependent weight in the branch excluding the all-ones vector. Explicit quadratic-space counts then control missing coordinates in the doubly-even code. The theorem concerns this uniform ensemble; equality with the sequential scheme’s sampling law is not assumed.

1 The ensemble and the result

All vector spaces and matrices are over $\mathbb{F}_2$. Let $m = m_1 + m_2$ and fix dimensions g, d, z with $n = g + d + z$. Consider block matrices

$$\tilde{H} = \begin{pmatrix} F & D & 0 \\ A & B & C \end{pmatrix} \in \mathbb{F}_2^{m \times n}, \quad (1.1)$$

where F has g columns, D has d columns, and C has z columns. The upper blocks have m_1 rows and the lower blocks have m_2 rows. The columns of D are independent, pairwise orthogonal, and doubly even: each has Hamming weight divisible by four. We require

$$D^\top F = 0, \quad \text{rank}(F^\top F) = g. \quad (1.2)$$

Because the columns of D are pairwise orthogonal and doubly even, every vector in their span is doubly even.

Take the uniform distribution on all matrices satisfying these upper constraints and the acceptance conditions

$$\text{rank } \tilde{H} = n, \quad \tilde{H}s = (\mathbf{1}_{m_1}, 0_{m_2})^\top \quad \text{for some } s. \quad (1.3)$$

The ensemble is assumed nonempty. Full column rank makes the accepted secret s unique. After sampling, apply an independent uniform row permutation P and an independent uniform invertible column transformation Q , giving $H = P\tilde{H}Q$ and secret $Q^{-1}s$.

The Radical Attack [1, Algorithm 1] computes the support of the subspace

$$\mathcal{R}(H) = H \ker(H^\top H). \quad (1.4)$$

It then solves $Hx = \mathbf{1}_{\text{supp } \mathcal{R}(H)}$. When this support is exactly the planted upper-row set, the unique solution is the secret.

Put

$$w = n - g - m_2 = d + z - m_2, \quad h = m_2 - z = d - w. \quad (1.5)$$

Admissibility implies $z \leq m_2$, so $h \geq 0$.

Theorem 1.1. *For the uniform accepted ensemble above, with $w \geq 1$ and all sufficiently large admissible m_1 , the failure probability of the Radical Attack is at most*

$$2m_1 2^{-d} + m_1 2^{-w} + 2^{1-w}. \quad (1.6)$$

Consequently the attack recovers the secret with probability $1 - o(1)$ along any sequence of nonempty admissible ensembles for which

$$d - \log_2 m_1 \longrightarrow \infty, \quad w - \log_2 m_1 \longrightarrow \infty.$$

Gross and Hangleiter established the attack and its normal-form interpretation, while Section 4.1.3 of [1] identifies the constrained full-support question. The purpose here is to prove the stated success regime under a precise law, including its acceptance bias. The source's approximate product expression is not used as a probability lower bound.

2 The exact accepted law of the lower blocks

Let $T = (F \mid D)$ and $k = g + d$.

Lemma 2.1. *The map $T : \mathbb{F}_2^k \rightarrow \mathbb{F}_2^{m_1}$ is injective. Full column rank of $\tilde{H}$ is equivalent to $\text{rank } C = z$. For fixed T with $\mathbf{1}_{m_1} \in \text{im } T$, let $u = (f, b) \neq 0$ be the unique vector satisfying $Tu = \mathbf{1}_{m_1}$. Conditional on acceptance and a fixed full-rank C , the quotient matrix*

$$(\bar{A} \mid \bar{B}) = \pi_C(A \mid B) : \mathbb{F}_2^k \longrightarrow \mathbb{F}_2^{m_2} / \text{im } C$$

is uniform among maps annihilating u . The probability of this quotient constraint before conditioning is exactly 2^{-h} , independent of T, u , and C .

Proof. If $Ff + Db = 0$, multiplication by $F^\top$ gives $F^\top Ff = 0$, hence $f = 0$ and then $b = 0$. This proves injectivity. A vector in $\ker \tilde{H}$ must therefore have its first k coordinates zero, and its last z coordinates in $\ker C$. This proves the rank equivalence.

Given full-rank C , the equation for a secret becomes $Tu = \mathbf{1}_{m_1}$ and $(A \mid B)u + Cx = 0$. The latter is solvable exactly when $\pi_C(A \mid B)u = 0$. Every quotient map has the same number of lifts, and before conditioning the quotient map is a uniform $h \times k$ matrix after choosing any quotient basis. For fixed nonzero u , each independent row annihilates u with probability $1/2$. Thus the event has probability 2^{-h} , proving both the conditional law and its independence of the particular upper block. $\square$

It follows that C is uniform among full-column-rank matrices and independent of the accepted upper block. The upper block has its uniform constrained law conditioned only on $\mathbf{1}_{m_1} \in \text{im } T$. No independence between F and D is asserted.

There are two useful branches.

If $f \neq 0$, one column of $\overline{A}$ is determined by the quotient constraint after all other columns have been chosen. Hence $\overline{B}$ is an unrestricted uniform $h \times d$ matrix. For a nonzero left vector y , the probability that $y^\top \overline{B} = 0$ is 2^{-d} . A union bound gives

$$\Pr(\text{rank } \overline{B} < h \mid f \neq 0) \leq (2^h - 1)2^{-d} < 2^{-w}. \quad (2.1)$$

If $h = 0$, rank failure is empty and the bound is interpreted accordingly.

If $f = 0$, then $D\mathbf{b} = \mathbf{1}_{m_1}$ and $b \neq 0$. Each row of $\overline{B}$ is uniform on the $(d-1)$ -dimensional annihilator of b . The same union bound gives

$$\Pr(\text{rank } \overline{B} < h \mid f = 0) \leq (2^h - 1)2^{-(d-1)} < 2^{1-w}. \quad (2.2)$$

These bounds are finite-size statements, without an asymptotic product approximation.

3 Why upper acceptance preserves uniform code weights

Write $\mathcal{D} = \text{im } D$ and, temporarily, put $m = m_1$. Equip $V = \mathbb{F}_2^m$ with the standard nondegenerate symmetric bilinear form $\langle x, y \rangle = x^\top y$. The doubly-even code $\mathcal{D}$ is totally isotropic. Therefore

$$U = \mathcal{D}^\perp / \mathcal{D}$$

has a nondegenerate induced symmetric bilinear form and dimension $m - 2d$. Since $\mathcal{D}$ consists of even-weight vectors, $\mathbf{1}_m \in \mathcal{D}^\perp$. The class $c = \mathbf{1}_m + \mathcal{D}$ satisfies

$$\langle x, x \rangle = \langle c, x \rangle \quad (x \in U). \quad (3.1)$$

It is the unique canonical vector of this bilinear space.

Lemma 3.1. *Conditional on acceptance and $\mathbf{1}_m \notin \mathcal{D}$, the subspace $\mathcal{D}$ is uniform among all doubly-even d -dimensional subspaces of $\mathbb{F}_2^m$ that exclude $\mathbf{1}_m$, provided this branch has positive weight.*

Proof. If $\mathbf{1}_m \notin \mathcal{D}$, then $c \neq 0$. Nondegeneracy and (3.1) imply that U is nonalternating. A nonalternating nondegenerate symmetric bilinear space over $\mathbb{F}_2$ is determined up to isometry by its dimension. One elementary proof splits off vectors of norm one. If an alternating plane remains, one previously split norm-one line together with a hyperbolic plane is congruent to I_3 : for a norm-one vector e and a hyperbolic pair a, b , the vectors $e + a, e + b, e + a + b$ are mutually orthogonal of norm one. Repeating this removes the alternating remainder. Every isometry preserves the unique vector characterized by (3.1).

Modulo $\mathcal{D}$, admissible columns of F are ordered bases of nondegenerate g -subspaces of U . The acceptance condition $\mathbf{1}_m \in \text{im}(F \mid D)$ is exactly the requirement that this subspace contain c . Its number of ordered bases depends only on the dimension of U and on g , since every isometry preserves c . Each ordered basis has exactly 2^{dg} lifts to $\mathcal{D}^\perp$; changing a lift by a vector of $\mathcal{D}$ changes none of its Gram entries. The number of ordered bases of $\mathcal{D}$ is also constant.

Finally, Lemma 2.1 shows that the lower-block acceptance factor is the same 2^{-h} for each such upper block. Thus the total number of accepted matrices above each eligible code $\mathcal{D}$ is constant. If that common number is zero, the branch is absent and no conditional statement is needed. $\square$

This argument uses a bilinear quotient, because the columns of F have no additional quadratic constraint. It applies even to codes with different Hamming-weight distributions, not merely to codes related by coordinate permutations. Also, $\mathbf{1}_m \notin \mathcal{D}$ is equivalent to $f \neq 0$, by injectivity of T .

4 Doubly-even codes and their coordinate support

Let $N(m, d)$ count doubly-even d -subspaces of $\mathbb{F}_2^m$, and let $N_0(m, d)$ count those excluding $\mathbf{1}_m$. A code missing a specified coordinate is exactly a doubly-even code of length $m - 1$. Such a code cannot contain $\mathbf{1}_m$. Hence, under the law in Lemma 3.1,

$$\Pr(\text{specified coordinate is missing}) = \frac{N(m-1, d)}{N_0(m, d)}. \quad (4.1)$$

We now compute and bound this ratio without replacing the code by independent random columns.

4.1 The quadratic space

Let $E = \mathbf{1}_m^\perp$ be the even-weight hyperplane and define

$$q(x) = \text{wt}(x)/2 \pmod{2}, \quad x \in E.$$

Its polar form is the restriction of the dot product. Doubly-even codes are exactly its totally singular subspaces.

For a nondegenerate quadratic space of dimension $2h$ and sign $e \in \{1, -1\}$, define

$$O_e(h, d) = \prod_{r=0}^{d-1} \frac{2^{2(h-r)-1} + e 2^{h-r-1} - 1}{2^{d-r} - 1}. \quad (4.2)$$

For a symplectic space of dimension $2h$, define

$$S(h, d) = \prod_{r=0}^{d-1} \frac{2^{2(h-r)} - 1}{2^{d-r} - 1}. \quad (4.3)$$

Empty products are one, and impossible subspace dimensions have count zero.

The first product counts totally singular subspaces. Indeed, the zero count of a nondegenerate $2s$ -dimensional quadratic form of sign e is $2^{2s-1} + e 2^{s-1}$. After an r -dimensional singular subspace has been chosen, its perpendicular quotient has dimension $2(h-r)$ and the same sign. Each quotient zero has 2^r lifts, and the 2^r vectors already in the span must be excluded. Dividing the count of successive ordered choices by $\prod_{r=0}^{d-1} (2^d - 2^r)$ gives (4.2). In the symplectic case every quotient vector is isotropic, giving (4.3). The quadratic zero count itself follows by multiplying the Gauss sums of hyperbolic and anisotropic planes.

4.2 All length parities

If m is odd, the restriction to E is nondegenerate of dimension $2h = m - 1$. Since $\mathbf{1}_m \notin E$, every code excludes it, and

$$N(m, d) = N_0(m, d) = O_e(h, d), \quad e = \begin{cases} +1, & m \equiv 1, 7 \pmod{8}, \\ -1, & m \equiv 3, 5 \pmod{8}. \end{cases} \quad (4.4)$$

The sign follows from the elementary Gauss sum $\sum_{x \in E} (-1)^{q(x)} = \Re(1+i)^m = e 2^h$.

If $m \equiv 0 \pmod{4}$, the radical of the polar form on E is $\langle \mathbf{1}_m \rangle$, and $q(\mathbf{1}_m) = 0$. Its quotient is nondegenerate quadratic of dimension $2h = m - 2$ and sign $e = (-1)^{m/4}$. A singular d -subspace of the quotient has 2^d linear sections avoiding the radical. Codes containing the radical correspond to singular $(d - 1)$ -subspaces of the quotient. Therefore

$$N_0(m, d) = 2^d O_e(h, d), \quad N(m, d) = 2^d O_e(h, d) + O_e(h, d - 1). \quad (4.5)$$

The term with $d - 1 = -1$ is zero.

If $m \equiv 2 \pmod{4}$, then $q(\mathbf{1}_m) = 1$. The quotient by the polar radical is symplectic of dimension $2h = m - 2$. Above any isotropic quotient subspace, adding $\mathbf{1}_m$ toggles q , so there is a unique linear section on which q vanishes. No singular code contains the radical. Hence

$$N(m, d) = N_0(m, d) = S(h, d). \quad (4.6)$$

4.3 A uniform shortening bound

Lemma 4.1. *For sufficiently large m and every admissible d with $N_0(m, d) > 0$,*

$$\frac{N(m - 1, d)}{N_0(m, d)} \leq 2 \cdot 2^{-d}. \quad (4.7)$$

Proof. Factor the quadratic extension count as

$$f_e(h) = 2^{2h-1} + e2^{h-1} - 1 = (2^h - e)(2^{h-1} + e). \quad (4.8)$$

Substituting (4.4)–(4.6) into the shortening ratio and cancelling adjacent factors gives

$$\frac{N(m - 1, d)}{N_0(m, d)} = \begin{cases} 2^{-d}, & m \equiv 0 \pmod{4}, \\ \frac{2^{h-d} + e}{2^h + e}, & m \equiv 2 \pmod{4}, \\ \frac{2^{h-d} - e}{2^h - e}, & m \equiv 3 \pmod{4}, \\ \frac{2^{2h-d-1} + e2^{h-1} - 1}{2^{2h-1} + e2^{h-1} - 1}, & m \equiv 1 \pmod{4}. \end{cases} \quad (4.9)$$

In the even case $h = (m - 2)/2$; in the odd case $h = (m - 1)/2$. For $m \equiv 2 \pmod{4}$, use $e = (-1)^{(m-2)/4}$; for odd m , use the sign in (4.4).

To make the cancellation in the last and potentially singular case explicit,

$$\begin{aligned} \frac{2^d O_e(h - 1, d) + O_e(h - 1, d - 1)}{O_e(h, d)} &= \frac{2^d f_e(h - d) + (2^d - 1)}{f_e(h)} \\ &= \frac{2^{2h-d-1} + e2^{h-1} - 1}{2^{2h-1} + e2^{h-1} - 1}. \end{aligned}$$

For $d = 0$ the ratio is one directly. At other admissible boundaries these are polynomial product cancellations; zero counts in the numerator are allowed, while a zero denominator has no probability weight.

For the middle two fractions in (4.9), write them as $(x \pm 1)/(y \pm 1)$ with $x = 2^{h-d} \geq 1$ and $y = 2^h$. If the sign is positive, the ratio is at most $2x/y$; if negative, the same follows from $y - 1 \geq y/2$. For the last ratio and $e = -1$, the numerator is at most 2^{2h-d-1} and the denominator is at least half of 2^{2h-1} for $h \geq 2$. For $e = +1$, the denominator is at least 2^{2h-1} , while the numerator is at most twice 2^{2h-d-1} because $d \leq h$. Thus every case is at most $2 \cdot 2^{-d}$. Taking m large enough ensures $h \geq 2$ in the relevant cases. $\square$

By Lemmas 3.1 and 4.1 and a union bound over coordinates,

$$\Pr(\mathcal{D} \text{ lacks full support} \mid \text{accepted}, f \neq 0) \leq 2m_1 2^{-d}. \quad (4.10)$$

5 From the code to the recovered secret

Lemma 5.1. *If $\overline{B}$ has full row rank, then*

$$\mathcal{R}(\tilde{H}) = (D \ker \overline{B}, 0_{m_2}). \quad (5.1)$$

In the branch $f = 0$, this space has full support on the upper m_1 rows. In the branch $f \neq 0$, conditional on full row rank, $\ker \overline{B}$ is a uniform w -dimensional subspace of $\mathbb{F}_2^d$, independent of the upper block.

Proof. Full row rank of $\overline{B}$ and full column rank of C imply that $R = (B \mid C)$ has full row rank. For each column of A , use columns of R to eliminate that lower column by a column operation on $\tilde{H}$. The upper F block changes only by adding columns from D , so its Gram matrix stays $F^\top F$ and its orthogonality to D is preserved. In this column basis, the Gram matrix has the invertible $F^\top F$ block and the remaining block $R^\top R$, because $D^\top D = 0$.

Since R has full row rank, $R^\top$ is injective. Thus $\ker(R^\top R) = \ker R$ over $\mathbb{F}_2$; this uses injectivity, not real positive definiteness. Projection of $\ker R$ onto its first d coordinates is a bijection onto $\ker \overline{B}$: the quotient condition is necessary and sufficient, and C has a unique preimage for each vector in its image. Its image under the upper matrix is $D \ker \overline{B}$, proving (5.1). The radical image is unchanged by the intermediate invertible column operations.

If $f = 0$, the accepted relation gives $\overline{B}b = 0$ and $Db = \mathbf{1}_{m_1}$, so the upper image contains the all-ones vector. If $f \neq 0$, Lemma 2.1 makes $\overline{B}$ a uniform matrix independent of the upper block. Conditional on full row rank its kernel is uniform of dimension $d - h = w$, proving the last assertion. $\square$

For a fixed nonzero row functional $\ell \in (\mathbb{F}_2^d)^*$ and a uniform w -subspace K ,

$$\Pr(K \subset \ker \ell) = \prod_{j=0}^{w-1} \frac{2^{d-1} - 2^j}{2^d - 2^j} \leq 2^{-w}. \quad (5.2)$$

For $w = d$ the product is zero, as expected. If $\mathcal{D}$ has full coordinate support, all rows of D are nonzero functionals. A union bound therefore bounds the chance that $D \ker \overline{B}$ misses any upper coordinate by $m_1 2^{-w}$.

Proof of Theorem 1.1. Rank failure costs at most 2^{1-w} by (2.1)–(2.2). Conditional on full rank, the $f = 0$ branch succeeds by Lemma 5.1. In the $f \neq 0$ branch, missing code support costs at most $2m_1 2^{-d}$ by (4.10). This conditional bound is unchanged by full-rank conditioning, because the distribution and rank probability of $\overline{B}$ are independent of the upper block. Given full code support, (5.2) bounds the remaining support failure by $m_1 2^{-w}$. Summing these bounds gives (1.6), without requiring independence among coordinate failures.

On the complementary event, (5.1) has support exactly the first m_1 coordinates. The attack solves the accepted equation and recovers its unique secret. For the obfuscated matrix,

$$\mathcal{R}(P\tilde{H}Q) = P\mathcal{R}(\tilde{H}),$$

because $P^\top P = I$ and Q is invertible. The support is permuted accordingly and the recovered secret is $Q^{-1}s$. Thus the success event is invariant under the prescribed obfuscation.

Finally, each term of (1.6) tends to zero under the two growth conditions. This proves the asymptotic assertion. $\square$

6 Evidence and distributional scope

The companion archive contains exact finite-field checkers for the code counts, shortening formulas, and accepted block law. Existing controls cover 24 exhaustive code cases, 3984 exact ratio checks, and 242 304 accepted small instances, including both secret branches. A further quotient check compares codes with different nonzero weight distributions, including $(4, 4, 4)$ and $(4, 4, 8)$, to test the absence of a hidden upper-block weight. These finite tests support the algebraic interfaces; they do not replace the general counting and conditional-law proofs.

The theorem is for the uniform distribution on all matrices satisfying (1.2) and (1.3). A sequential generator may weight these matrices differently. Transferring the guarantee to that generator requires proving equality of laws or deriving a separate failure bound under its actual law. Neither the empirical recovery of a challenge instance nor an unconditioned random-code heuristic establishes that transfer.

The attack and its radical-image interpretation are attributed to [1]. The contribution developed here is the conditioned success analysis for the stated ensemble, with explicit control of code support across all length parities. This remains an internally audited manuscript, not a claim of external verification or exhaustive novelty clearance.

AI-assisted research

This work was developed in an author-directed, AI-assisted research workflow. AI systems assisted with mathematical derivation, literature checks, computational verification, and manuscript drafting. The internal reviews and separate implementations described above were produced within this workflow and do not constitute external expert review.

References

- [1] D. Gross and D. Hangleiter, *Secret extraction attacks against obfuscated IQP circuits*, PRX Quantum **6**, 020314 (2025). [doi:10.1103/PRXQuantum.6.020314](https://doi.org/10.1103/PRXQuantum.6.020314); [arXiv:2312.10156v2](https://arxiv.org/abs/2312.10156v2), especially Algorithm 1, Corollary 3, and Section 4.1.3.