

Deterministic Parameterized Inapproximability of Nearest Codeword and Minimum Distance*

Reported by[†]

Venkatesan Guruswami[‡]

Xuandi Ren[§]

October 2026

Abstract

We show that the nearest-codeword and minimum-distance problems for linear codes over every fixed finite field are $W[1]$ -hard to approximate within any constant factor under deterministic fixed-parameter many-one reductions. This gives unconditional deterministic parameterized inapproximability for minimum distance, including the binary problem usually called EVEN SET. The reduction starts from *exact* nearest codeword and *produces a constant gap directly*. Over $\mathbb{F}_q$ with $q > 2$, an input of length m and threshold k gives instances with threshold q^k and length $O_q(q^k m)$. Tensor products amplify the gap, and a simple concatenation handles codes over $\mathbb{F}_2$.

1 Introduction

The Nearest Codeword Problem (NCP) asks for the distance from a given vector to a linear code, while the Minimum Distance Problem (MDP) asks for the minimum weight of a nonzero codeword. NCP is also equivalent to finding the minimum weight vector in an affine subspace, and thus corresponds to the affine version of MDP. Both problems are NP-hard: Berlekamp, McEliece, and van Tilborg [BMT78] proved that NCP is NP-complete, and much later Vardy [Var97] established NP-hardness of minimum distance.

NCP was also shown to be NP-hard to approximate within any constant factor in early applications of the PCP theorem [Ste93, ABSS97]. As with exact NP-hardness, constant-factor inapproximability of MDP was established only much later. Dumer, Micciancio, and Sudan [DMS03] proved that approximating MDP within any constant factor is NP-hard under randomized reductions. Cheng and Wan [CW12] derandomized the DMS reduction by constructing the requisite gadget, a locally dense code and explicit center with many nearby codewords. Austrin and Khot [AK14] gave an alternate deterministic reduction using tensor codes. Building on the tensor-code approach of [AK14], Micciancio [Mic14] gave a deterministic construction of locally dense gadgets over $\mathbb{F}_2$.

*Work supported in part by NSF grant CCF-2211972 and DOD Advanced Research Projects Agency grant HR0011262E031.

[†]The proof was discovered by GPT-6. See the “AI methodology” at the end of the paper for details.

[‡]Simons Institute for the Theory of Computing and Department of EECS & Mathematics, UC Berkeley.
venkatg@berkeley.edu

[§]Department of EECS, UC Berkeley. xuandi_ren@berkeley.edu

from arbitrary binary linear codes of sufficiently large relative distance. More recently, Bhatiprolu, Guruswami, Lee, and Ren [BGLR25] gave a PCP-free proof of deterministic constant-factor inapproximability over every fixed finite field.

Parameterized complexity asks whether the problems become tractable when the optimum value (or some other parameter of interest) k is small. An algorithm is fixed-parameter tractable (FPT) if it runs in time $f(k)n^{O(1)}$, where n is the input length and the polynomial exponent is independent of k . To rule out such algorithms, a reduction must run in FPT time and keep the output parameter bounded by a function of k alone.

Downey, Fellows, Vardy, and Whittle [DFVW99] established W[1]-hardness of exact NCP, with the target distance being the parameter of interest. For MDP, usually called EVEN SET over $\mathbb{F}_2$, Downey and Fellows [DF13] listed the corresponding question as one of the six “most infamous” open problems in parameterized complexity. It remained open until Bhattacharyya et al. [BBE⁺21] proved constant-factor parameterized inapproximability for the binary problem under randomized reductions. Bennett, Cheraghchi, Guruswami, and Ribeiro [BCGR24] then extended this result to every finite field. Hirahara and Ogitsuka [HO26] subsequently obtained one-sided-error reductions over all fields and deterministic reductions under a circuit lower-bound assumption. However, an unconditional deterministic reduction was still missing.

Our results. We obtain a constant approximation gap directly from exact NCP, using a deterministic FPT reduction. The output is an NCP instance whose underlying code also has large minimum distance. This additional guarantee makes it an instance of the Relatively Near Codeword (RNC) problem defined in [DMS03]. The DMS reduction then gives hardness of gap MDP, and tensor products amplify both gaps to every fixed approximation factor.

Theorem 1.1. *For every fixed finite field $\mathbb{F}_q$ and every fixed $\gamma \geq 1$, the γ -gap nearest-codeword and minimum-distance problems are W[1]-hard under deterministic fixed-parameter many-one reductions, parameterized by the distance threshold.*

For minimum distance, this removes randomness from the parameterized hardness reduction, including for EVEN SET. For nearest codeword, deterministic parameterized gap hardness was already known [BBE⁺21, BCGR24], and the new contribution here is a *direct gap-producing reduction* from the *exact* version of the problem, using the same construction as for minimum distance. The output threshold of our reductions grows exponentially with the source parameter.¹

The reduction. For a linear code $U \leq \mathbb{F}_q^m$ and a target $\tau \in \mathbb{F}_q^m$, write

$$d(U) = \min_{u \in U \setminus \{0\}} \text{wt}(u), \quad \delta(\tau + U) = \min_{u \in U} \text{wt}(\tau - u),$$

where wt denotes Hamming weight and $d(\{0\}) = \infty$. The γ -gap versions ask to distinguish an optimum of at most k from one greater than γk . Exact nearest codeword is the case $\gamma = 1$. Codes are specified by generator matrices.

¹For a fixed prime p , an earlier construction discovered by GPT-6 using algebraic geometry attains a linear parameter blow-up; namely, it reduces $(k$ vs. $(4p+1)k$) NCP to $((4p-1)k+1$ vs. $4pk+1$) MDP over $\mathbb{F}_p$. However, we did not carefully verify the proof and chose not to include it here.

Reference	Fields	Reductions for constant-gap MDP
[BBE ⁺ 21]	$\mathbb{F}_2$	Randomized, one-sided error
[BCGR24]	Any $\mathbb{F}_q$	Randomized, two-sided error
[HO26]	Any $\mathbb{F}_q$	Randomized, one-sided error; (Deterministic under a circuit lower-bound assumption)
Theorem 1.1	Any $\mathbb{F}_q$	Deterministic, unconditional

Table 1: Parameterized minimum-distance hardness under different kinds of reductions. Every row states W[1]-hardness under the indicated kind of reduction; the deterministic result of [HO26] additionally assumes a circuit lower bound.

Proposition 1.2. *Fix a prime power $q > 2$, and set*

$$\rho = 2 - \frac{2}{q}, \quad \alpha = 1 + \frac{\rho}{3}.$$

Given an exact NCP instance (U, τ, k) of length m , with $m, k \geq 1$, one can deterministically construct a linear code $V \leq \mathbb{F}_q^n$ and a vector $z \in \mathbb{F}_q^n$. Put $t = q^k$. Then $d(V) \geq \rho t$ on every input, and

$$\begin{aligned} \delta(\tau + U) \leq k &\implies \delta(z + V) = t, \\ \delta(\tau + U) > k &\implies \delta(z + V) \geq \alpha t. \end{aligned}$$

The length satisfies

$$n < q^{k+1}(1 + (q-1)m).$$

The construction takes time polynomial in n and the input size, with an exponent independent of k . All arithmetic is over $\mathbb{F}_q$.

The proposition gives hardness of γ -gap NCP for every $1 \leq \gamma < \alpha$, where $\alpha \geq 13/9$. The additional bound on $d(V)$ allows the DMS reduction in Lemma 2.4 to give hardness of γ -gap MDP for every $1 \leq \gamma < \min\{\rho, \alpha\}$, where $\min\{\rho, \alpha\} \geq 4/3$. We then amplify both gaps to any constant via tensoring, as shown in Appendix A. The output threshold is $q^{O_\gamma(k)}$ and the length is $q^{O_\gamma(k)} m^{O(\log(\gamma+1))}$ for fixed q , so the amplified reduction remains FPT. Over $\mathbb{F}_2$, we first construct and amplify the instances over $\mathbb{F}_4$, then apply the binary conversion in Appendix B; this doubles the threshold and triples the length while preserving the gap.

1.1 Proof overview

A compact representation of the NCP solution. Fix $q > 2$, and let (U, τ, k) be the source instance, with $U \leq \mathbb{F}_q^m$. A solution is an error vector $u \in \tau + U$ of weight at most k . Choose a parity-check matrix $B \in \mathbb{F}_q^{h \times m}$ with $\ker B = U$, where $h = m - \dim U$, and put $\eta = B\tau$. Then $u \in \tau + U$ exactly when $Bu = \eta$. Writing B_j for the j th column of B , a sparse solution expresses η as a sum of at most k contributions of the form aB_j , where $j \in [m]$ and $a \in \mathbb{F}_q^*$.

We represent each coordinate-value pair (j, a) by a distinct nonzero label $\text{lab}(j, a) \in \mathbb{F}_q^b$, where

$$b = \left\lceil \log_q(1 + (q-1)m) \right\rceil.$$

A sparse solution can thus be specified by a list of at most k labels. Arranging these labels and padding with zero columns gives a matrix $A = [a_1 \ \cdots \ a_k] \in \mathbb{F}_q^{b \times k}$. Define a representation map $v : \mathbb{F}_q^b \rightarrow \mathbb{F}_q^h$ by $v(\text{lab}(j, a)) = aB_j$, with value zero on the zero and unused labels. Each list of k labels thus specifies the syndrome of an error vector supported on at most k coordinates, and every such error vector can be represented this way. Consequently, exact NCP is equivalent to finding A satisfying

$$\sum_{i=1}^k v(a_i) = \eta.$$

The representation map v need not be linear, so this equation need not be a linear constraint on A .

This representation records the support and coefficients using only $kb = O_q(k \log m)$ field symbols. By comparison, representing arbitrary vectors in $\mathbb{F}_q^m$ requires m symbols. The distinction matters because our gadget will encode A using length q^{k+b} and radius q^k : choosing $b = O_q(\log m)$ makes the length $O_q(q^k m)$, while the radius depends only on k . We remark that in the classical NP-hardness setting, the distance threshold k can grow with m , which makes our construction not in polynomial time.

The coding gadget and the reduction. We encode the matrices A by codewords c_A in a linear code $C \leq \mathbb{F}_q^n$, where $n = q^{k+b}$, with a common center σ and a linear map $T : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^h$. For $t = q^k$, we require

$$d(C) \geq \rho t, \quad \text{wt}(c_A - \sigma) = t, \quad T(c_A) = \sum_{i=1}^k v(a_i).$$

Although v need not be linear, the sum $\sum_i v(a_i)$ can be decoded from c_A by the linear map T . We also require *rigidity*: every $c \in C$ outside the family $\{c_A\}$ must satisfy $\text{wt}(c - \sigma) \geq \alpha t$. Thus every codeword close to σ represents a choice of source labels. Finally, $T(\sigma) = 0$, and the code, center, map, and each requested c_A must be computable in polynomial time in n and the table size. [Lemma 3.1](#) formally defines the properties required of the coding gadget.

We impose the source constraint by requiring $T(c) = \eta$ for $c \in C$. To construct the output target, for each $j \in [m]$ let A_j have first column $\text{lab}(j, 1)$ and all other columns zero. Then $T(c_{A_j}) = B_j$. Writing τ_j for the j th coordinate of τ , linearity gives

$$T\left(\sum_{j=1}^m \tau_j c_{A_j}\right) = \sum_{j=1}^m \tau_j B_j = B\tau = \eta.$$

We can therefore set $z = -\sigma + \sum_{j=1}^m \tau_j c_{A_j}$ and $V = C \cap \ker T$. The output coset is therefore

$$z + V = \{c - \sigma : c \in C, T(c) = \eta\}.$$

In the YES case, this coset contains $c_A - \sigma$ for a source solution. The gadget properties give weight at least t for every vector in the coset, so its distance is exactly t . In the NO case, it contains no such vector, so rigidity gives distance at least αt . Moreover, $V \leq C$ gives $d(V) \geq \rho t$ on every input. This is the extra guarantee needed for the *Relatively Near Codeword Problem* (RNC) of [\[DMS03\]](#): NCP with the additional promise that its threshold is less than θ times the minimum distance of its code, for a fixed $\theta > 0$; see [Definition 2.3](#). Our distance bound gives this promise for every $\theta > 1/\rho$. We then apply the DMS reduction in [Lemma 2.4](#) as a black box to obtain the minimum-distance gap stated above.

Constructing the gadget from a Reed–Muller code. To encode A , consider the indicator of the image of $x \mapsto (x, Ax)$:

$$w_A(x, y) = \mathbf{1}[y = Ax] = \prod_{j=1}^b (1 - (y_j - (Ax)_j)^{q-1}), \quad (x, y) \in \mathbb{F}_q^k \times \mathbb{F}_q^b.$$

View w_A as a matrix with rows indexed by x and columns indexed by y . Each row has exactly one nonzero entry, equal to one, so $\text{wt}(w_A) = q^k = t$. The displayed polynomial has total degree at most $b(q-1)$, and every monomial has total degree divisible by $q-1$. Thus all w_A lie in the Reed–Muller subcode $\mathcal{E}_{k,b}$ spanned by the evaluations of monomials satisfying these conditions.

The ambient code contains these weight- t indicators, so its distance cannot exceed t . To obtain the larger distance required by the gadget, let $w_0(x, y) = \mathbf{1}[y = 0]$ correspond to the choice of the zero matrix for A and set

$$c_A = w_A - w_0, \quad C = \text{span}\{c_A : A \in \mathbb{F}_q^{b \times k}\}, \quad \sigma = -w_0.$$

All indicators agree on the row $x = 0$, so their differences vanish there. In particular, $C \leq \mathcal{E}_{k,b}^0 = \{c \in \mathcal{E}_{k,b} : c(0, 0) = 0\}$. Each word of $\mathcal{E}_{k,b}^0$ can be represented by a homogeneous polynomial of degree $b(q-1)$. The evaluations agree on each class of nonzero scalar multiples, since the degree is divisible by $q-1$. Thus the evaluation vectors vanish at the origin and repeat each projective evaluation $q-1$ times. The projective Reed–Muller distance bound in [Sor91, GL24] gives $d(\mathcal{E}_{k,b}^0) = \rho t$, and hence $d(C) \geq \rho t$. At the same time, the choice $\sigma = -w_0$ gives $\text{wt}(c_A - \sigma) = \text{wt}(w_A) = t$.

To establish rigidity, for $c \in C$ put $w = c - \sigma$. The zeroth and first moments of the row of w_A indexed by x are 1 and Ax , respectively. Because w is w_0 plus a linear combination of the differences $w_A - w_0$, its row sums remain one and its first moments are determined by a single matrix. We recover this matrix from the rows indexed by the standard basis vectors $e_1, \dots, e_k$ of $\mathbb{F}_q^k$: define $R : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{b \times k}$ by

$$(Rw)_{*,i} = \sum_y y w(e_i, y), \quad i \in [k].$$

Then, for every $x \in \mathbb{F}_q^k$,

$$\sum_y w(x, y) = 1, \quad \sum_y y w(x, y) = (Rw)x.$$

Set $A = Rw$. Every row is nonzero, and a row with a single nonzero entry has value one at position Ax ; it therefore agrees with the corresponding row of w_A . Counting the entries in the remaining rows gives

$$\text{wt}(c - c_A) = \text{wt}(w - w_A) \leq 3(\text{wt}(w) - t).$$

If $c \neq c_A$, the left-hand side is at least $d(C) \geq \rho t$. Rearranging gives $\text{wt}(c - \sigma) \geq (1 + \rho/3)t = \alpha t$, which is precisely the required rigidity.

The k rows of w_A indexed by $e_1, \dots, e_k$ also allow us to decode $\sum_i v(a_i)$ linearly: the row indexed by e_i has its only nonzero entry at a_i . Weighting each such row by the fixed table $v(y)$ and summing over i gives $\sum_i v(a_i)$. This operation is linear in the entries of w_A , although v itself need not be linear. Since $v(0) = 0$, subtracting w_0 leaves the result unchanged, so $T(c_A) = \sum_i v(a_i)$ and $T(\sigma) = 0$.

To construct C efficiently, for each matrix $B \in \mathbb{F}_q^{b \times k}$ with a single nonzero entry, we use a coordinate permutation sending w_A to w_{A+B} . Starting from the corresponding differences c_B , we close their span under these permutations. The span stabilizes after at most n dimension increases and equals C , giving a polynomial-time construction without enumerating all q^{bk} matrices.

Comparison with previous constructions. The reduction in [DMS03] obtains gap RNC from an already-gapped NCP instance using a locally dense code and a linear projection covering the entire source message space, whose dimension may be linear in m . The construction in [CW12] derandomizes this code, center, and covering projection using Reed–Solomon/Hadamard concatenation. A different deterministic reduction based on tensor products is given in [AK14]. For binary codes, [Mic14] develops this tensor-code approach into a general construction of locally dense gadgets, including their centers and covering projections, from base codes of sufficiently large relative distance. These reductions work in the non-parameterized setting. In the parameterized setting, we face the additional challenge of ensuring that the output threshold depends only on k . To address this, we encode sparse solutions rather than arbitrary source messages, which allows us to use a code whose radius q^k is independent of the label length b . We retain the final RNC-to-MDP step of [DMS03], but replace its preceding gap-transfer step by a reduction from exact NCP.

The parameterized reductions in [BBE⁺21, BCGR24] use BCH codes and randomness to obtain many codewords close to a center. In [BCGR24], this counting guarantee replaces full message-space coverage, and random sparsification separates the YES and NO cases. The reductions in [HO26] give one-sided error over all fields and conditional derandomization. Our gadget instead identifies every codeword sufficiently close to σ as some c_A , so the linear constraint $T(c) = \eta$ can deterministically exclude all such codewords in the NO case. This changes both the gadget guarantee and its use, and creates a gap from exact feasibility. The tradeoff is a larger parameter: the base reduction in [BCGR24] has a linear threshold increase, whereas ours has threshold q^k . The earlier algebraic-geometric alternative mentioned in the footnote above also has a linear base threshold, but starts from gap NCP.

The gap-producing step is also related to the PCP-free reduction in [BGLR25]. That reduction linearizes quadratic constraints using matrix variables and encodes the matrices in the tensor square of a nearly constant-weight code. A constant Hamming-weight gap separates rank-one encodings, which represent genuine solutions, from higher-rank solutions of the linear constraints. Here, the codewords c_A , whose differences from the center are the image indicators w_A , play the role of rank-one encodings. In both constructions, linear constraints express exact feasibility on a nonlinear family, and low weight or distance from a fixed center forces membership in that family. The rigidity arguments differ: the proof in [BGLR25] uses the larger support of a two-dimensional subcode, whereas our argument combines moment identities with the bound $d(C) \geq \rho t$ to show that every codeword sufficiently close to σ must in fact be some c_A .

Organization. Section 2 gives the problem definitions, the source hardness result, and the RNC-to-MDP reduction. Section 3 specifies the required properties of the coding gadget and gives a reduction from exact NCP to gap RNC. Section 4 constructs the gadget and verifies its properties. Appendices A and B show how to amplify the gap to any constant and how to extend the inapproximability results to $\mathbb{F}_2$, respectively.

2 Preliminaries

Throughout, q is a fixed prime power and $\mathbb{F}_q$ is the field with q elements. All linear algebra is over $\mathbb{F}_q$ unless stated otherwise. We write $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$ and $[m] = \{1, \dots, m\}$.

For a vector $u \in \mathbb{F}_q^m$, define its *Hamming weight* to be $\text{wt}(u) = |\text{supp}(u)|$, where $\text{supp}(u) = \{j \in [m] : u_j \neq 0\}$ is its *support*. The Hamming distance between $u, v \in \mathbb{F}_q^m$ is $\text{wt}(u - v)$.

A *linear code* of length m is a subspace $U \leq \mathbb{F}_q^m$; its elements are called *codewords*. Let $\ell = \dim U$ be the dimension of U . We represent U by a *generator matrix* $G \in \mathbb{F}_q^{m \times \ell}$ whose columns form a

basis of U . A *parity-check matrix* for U is a matrix B with $\ker B = U$; the two representations can be converted into one another by Gaussian elimination.

We write $d(U) = \min_{u \in U \setminus \{0\}} \text{wt}(u)$ for the minimum distance of U , with $d(\{0\}) = \infty$, and $\delta(\tau + U) = \min_{u \in U} \text{wt}(\tau - u)$ for the Hamming distance from a vector $\tau \in \mathbb{F}_q^m$ to U . We use the following decision and promise versions of the coding problems.

Definition 2.1 (Nearest codeword problem). An instance of the Nearest Codeword Problem (NCP) over $\mathbb{F}_q$ is a triple (U, τ, k) consisting of a generator matrix for a linear code $U \leq \mathbb{F}_q^m$, a target vector $\tau \in \mathbb{F}_q^m \setminus U$, and a positive integer k . The exact problem asks whether there exists $u \in U$ with $\text{wt}(\tau - u) \leq k$, equivalently whether $\delta(\tau + U) \leq k$. For a fixed $\gamma \geq 1$, γ -gap NCP has the same input (U, τ, k) and asks to distinguish

$$\text{YES} : \delta(\tau + U) \leq k, \quad \text{NO} : \delta(\tau + U) > \gamma k.$$

Definition 2.2 (Minimum distance problem). An instance of the Minimum Distance Problem (MDP) over $\mathbb{F}_q$ is a pair (U, k) consisting of a generator matrix for a linear code $U \leq \mathbb{F}_q^m$ and a positive integer k . The exact problem asks whether there exists a *nonzero* word $u \in U$ with $\text{wt}(u) \leq k$, equivalently whether $d(U) \leq k$. For a fixed $\gamma \geq 1$, γ -gap MDP has the same input (U, k) and asks to distinguish

$$\text{YES} : d(U) \leq k, \quad \text{NO} : d(U) > \gamma k.$$

Definition 2.3 (Relatively near codeword problem [DMS03, Definition 2]). Fix $\gamma \geq 1$ and $\theta > 0$. An instance of γ -gap RNC_θ over $\mathbb{F}_q$ is an NCP instance (U, τ, k) with the additional promise that

$$k < \theta d(U).$$

Subject to this promise, the problem asks to distinguish

$$\text{YES} : \delta(\tau + U) \leq k, \quad \text{NO} : \delta(\tau + U) > \gamma k.$$

For $0 < \theta \leq 1/\gamma$, hardness of γ -gap RNC_θ transfers to γ -gap MDP, as shown in [DMS03, Lemma 19].

Lemma 2.4 ([DMS03]). *For every fixed finite field $\mathbb{F}_q$, $\gamma \geq 1$, and $0 < \theta \leq 1/\gamma$, there is a deterministic polynomial-time reduction from γ -gap RNC_θ to γ -gap MDP over $\mathbb{F}_q$. The reduction preserves the length $m \geq 1$ and the threshold k .*

Proof. Let (U, τ, k) be the input instance. By definition, $\tau \notin U$. Put $\mathcal{D} = U + \mathbb{F}_q \tau$ and output $(\mathcal{D}, k)$. A nonzero vector in $\mathcal{D}$ either lies in U or has the form $u + \beta \tau$ with $u \in U$ and $\beta \neq 0$. Hence

$$d(\mathcal{D}) = \min\{d(U), \delta(\tau + U)\}. \tag{1}$$

In the YES case, $\delta(\tau + U) \leq k$, so (1) gives $d(\mathcal{D}) \leq k$. In the NO case, $\delta(\tau + U) > \gamma k$, while the RNC promise gives $d(U) > k/\theta \geq \gamma k$. Hence $d(\mathcal{D}) > \gamma k$. $\square$

Our source problem is the exact Nearest Codeword Problem, which was proved to be $\text{W}[1]$ -hard when parameterized by k in [DFVW99].

Lemma 2.5 ([DFVW99]). *For every fixed finite field $\mathbb{F}_q$,² the exact Nearest Codeword Problem, parameterized by the threshold k , is $\text{W}[1]$ -hard under deterministic FPT many-one reductions.*

²The result was originally proved over $\mathbb{F}_2$ for maximum-likelihood decoding (MLD), a problem equivalent to NCP. The same reduction works over every finite field.

3 A coding gadget and the reduction

In this section, we prove [Proposition 1.2](#) by strengthening the locally dense code approach of [\[DMS03\]](#) to obtain an approximation gap for RNC from exact NCP.

3.1 Encoding the source constraints

We start from exact NCP and express a sparse solution as a list of k labels. Fix a prime power $q > 2$. Let (U, τ, k) be an exact NCP instance with $U \leq \mathbb{F}_q^m$. Compute a parity-check matrix $B \in \mathbb{F}_q^{h \times m}$ with $\ker B = U$, where $h = m - \dim U \geq 1$, and put $\eta = B\tau$. Then $u \in \tau + U$ if and only if $Bu = \eta$. Writing B_j for the j th column, we have

$$\delta(\tau + U) \leq k \iff \eta \text{ is a sum of } k \text{ elements of } S = \{0\} \cup \{aB_j : j \in [m], a \in \mathbb{F}_q^*\}. \quad (2)$$

Set

$$b = \left\lceil \log_q(1 + (q-1)m) \right\rceil.$$

Then $q^b \geq 1 + (q-1)m$, so $\mathbb{F}_q^b$ contains enough labels to represent every element of S . We assign a distinct nonzero label vector $\text{lab}(j, a) \in \mathbb{F}_q^b$ to each pair $(j, a) \in [m] \times \mathbb{F}_q^*$, reserving the zero vector for $0 \in S$. Define a representation map $v : \mathbb{F}_q^b \rightarrow \mathbb{F}_q^h$ by

$$v(\text{lab}(j, a)) = aB_j, \quad v(y) = 0 \quad \text{for } y = 0 \text{ and all unused labels } y.$$

Then (2) is equivalent to the existence of a matrix $A = [a_1 \cdots a_k] \in \mathbb{F}_q^{b \times k}$ satisfying

$$\sum_{i=1}^k v(a_i) = \eta. \quad (3)$$

Since v need not be linear, the sum in (3) may be nonlinear in A .

We then seek a locally dense linear code $C \leq \mathbb{F}_q^n$, a center $\sigma \in \mathbb{F}_q^n$, and a radius $t < d(C)$ with $\text{wt}(\sigma) = t$. Each matrix $A \in \mathbb{F}_q^{b \times k}$ is associated with a codeword $c_A \in C$ satisfying $\text{wt}(c_A - \sigma) = t$. To check (3) for A using its codeword c_A , we require a linear map T that returns $\sum_{i=1}^k v(a_i)$ when applied to c_A . Thus $T(c_A) = \eta$ holds exactly when the labels in A form a sparse source solution. Furthermore, we require every codeword outside the family $\{c_A\}$ to have distance at least αt from σ for some constant $\alpha > 1$, to rule out spurious short solutions in the NO case.

We summarize these properties in the following lemma.

Lemma 3.1 (The coding gadget). *Fix a prime power $q > 2$. Given integers $k, b, h \geq 1$ and a table $v : \mathbb{F}_q^b \rightarrow \mathbb{F}_q^h$ with $v(0) = 0$, set*

$$t = q^k, \quad n = q^{k+b}, \quad \rho = 2 - \frac{2}{q}, \quad \alpha = 1 + \frac{\rho}{3}.$$

There exist a linear code $C \leq \mathbb{F}_q^n$, a center $\sigma \in \mathbb{F}_q^n$, a linear map $T : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^h$, and a set of codewords $\{c_A : A \in \mathbb{F}_q^{b \times k}\} \subseteq C$ with the following properties.

(i) (Locally dense code) $d(C) \geq \rho t$, and for every $A \in \mathbb{F}_q^{b \times k}$,

$$\text{wt}(c_A - \sigma) = t.$$

Furthermore, $c_0 = 0$ and thus $\text{wt}(\sigma) = t$.

- (ii) (Rigidity) Every $c \in C$ either equals c_A for some A , or satisfies $\text{wt}(c - \sigma) \geq \alpha t$.
- (iii) (Linear decodability of labels) $T(\sigma) = 0$, and for every $A = [a_1 \ \cdots \ a_k] \in \mathbb{F}_q^{b \times k}$,

$$T(c_A) = \sum_{i=1}^k v(a_i).$$

- (iv) (Effective construction) C , σ , and T can be computed in $(nkbh)^{O(1)}$ time. Given $A \in \mathbb{F}_q^{b \times k}$, the codeword c_A can be computed in $(nkb)^{O(1)}$ time. All algorithms are deterministic, with hidden constants depending only on q .

We now use [Lemma 3.1](#) to construct the gap instance. The proof of the lemma is deferred to [Section 4](#).

3.2 The output instance

Recall the source NCP instance (U, τ, k) and the parity-check system (B, η) such that $u \in \tau + U$ if and only if $Bu = \eta$. Apply [Lemma 3.1](#) to the table v and parameters k, b, h defined above to obtain C, σ, T and the codewords c_A . For each $j \in [m]$, let $A_j \in \mathbb{F}_q^{b \times k}$ be the matrix with the first column $\text{lab}(j, 1)$ and all remaining columns zero. By (iii),

$$T(c_{A_j}) = v(\text{lab}(j, 1)) = B_j.$$

Writing τ_j for the j th coordinate of τ , define the output code and target by

$$V = C \cap \ker T,$$

$$z = -\sigma + \sum_{j=1}^m \tau_j c_{A_j}.$$

The output NCP instance is (V, z, t) , where $t = q^k$. Since $z + \sigma \in C$ and $\text{wt}(\sigma) = t < d(C)$, we have $z \notin C$, and hence $z \notin V$.

We observe that

$$z + V = \{c - \sigma : c \in C, T(c) = \eta\},$$

since $\sum_j \tau_j c_{A_j} \in C$ and

$$T\left(\sum_{j=1}^m \tau_j c_{A_j}\right) = \sum_{j=1}^m \tau_j B_j = \eta.$$

Thus, for $c \in C$, the condition $T(c) = \eta$ is equivalent to $c - \sum_j \tau_j c_{A_j} \in C \cap \ker T = V$, which is equivalent to $c - \sigma \in z + V$.

Note that $V \leq C$ implies $d(V) \geq d(C) \geq \rho t$. Thus the output is also an RNC_θ instance for every $\theta > 1/\rho$.

3.3 Correctness and parameter bounds

We now prove [Proposition 1.2](#) from [Lemma 3.1](#).

Proposition 1.2. Fix a prime power $q > 2$, and set

$$\rho = 2 - \frac{2}{q}, \quad \alpha = 1 + \frac{\rho}{3}.$$

Given an exact NCP instance (U, τ, k) of length m , with $m, k \geq 1$, one can deterministically construct a linear code $V \leq \mathbb{F}_q^n$ and a vector $z \in \mathbb{F}_q^n$. Put $t = q^k$. Then $d(V) \geq \rho t$ on every input, and

$$\begin{aligned} \delta(\tau + U) \leq k &\implies \delta(z + V) = t, \\ \delta(\tau + U) > k &\implies \delta(z + V) \geq \alpha t. \end{aligned}$$

The length satisfies

$$n < q^{k+1}(1 + (q-1)m).$$

The construction takes time polynomial in n and the input size, with an exponent independent of k . All arithmetic is over $\mathbb{F}_q$.

Proof of Proposition 1.2. Recall that $z + V$ equals

$$\{c - \sigma : c \in C, T(c) = \eta\}.$$

By (i) and (ii), vectors in $z + V$ of the form $c_A - \sigma$ have weight t , and all other vectors in $z + V$ have weight at least αt . Thus $\delta(z + V) \geq t$. In the YES case of the source NCP instance, there exists a matrix $A = [a_1 \cdots a_k] \in \mathbb{F}_q^{b \times k}$ satisfying $\sum_{i=1}^k v(a_i) = \eta$. By (iii), the codeword c_A satisfies $T(c_A) = \eta$, so $c_A - \sigma$ lies in $z + V$ and has weight t by (i). In the NO case, no such matrix exists, so $T(c_A) \neq \eta$ for every A and $z + V$ contains no vector of the form $c_A - \sigma$. Therefore, $\delta(z + V) = t$ in the YES case and $\delta(z + V) \geq \alpha t$ in the NO case.

Finally, the output length is $n = q^{k+b}$ by Lemma 3.1. Substituting $b = \lceil \log_q(1 + (q-1)m) \rceil$ gives

$$n < q^{k+1}(1 + (q-1)m) = O_q(q^k m).$$

By (iv), constructing the gadget and the m codewords c_{A_j} takes time polynomial in n and the input size, since $k, b \leq \log_q n$ and $h \leq m$. Computing V and z then requires only Gaussian elimination and matrix arithmetic. All steps are deterministic over $\mathbb{F}_q$, with a polynomial exponent independent of k . $\square$

Combining Proposition 1.2 with Lemma 2.5 gives deterministic W[1]-hardness of γ -gap RNC_θ over every fixed field $\mathbb{F}_q$ with $q > 2$, for every fixed $1 \leq \gamma < \alpha$ and $\theta > 1/\rho = q/(2(q-1))$. Forgetting the additional distance promise gives the same hardness for γ -gap NCP. Applying Lemma 2.4 with $\theta = 1/\gamma$ gives the same hardness for γ -gap MDP for $1 \leq \gamma < \min\{\rho, \alpha\}$.

The above results already give a 13/9-gap for NCP and a 4/3-gap for MDP. To amplify the gap to any constant, we tensor the instances as described in Appendix A.

For $\mathbb{F}_2$, we first apply the reductions and amplification over $\mathbb{F}_4$, then convert the outputs to binary using Appendix B. These steps complete the proof of Theorem 1.1.

4 Constructing the coding gadget

In this section, we prove Lemma 3.1. Fix $q > 2$, integers $k, b, h \geq 1$, and a table $v : \mathbb{F}_q^b \rightarrow \mathbb{F}_q^h$ with $v(0) = 0$. Throughout this section,

$$t = q^k, \quad n = q^{k+b}, \quad \rho = 2 - 2/q, \quad \alpha = 1 + \rho/3.$$

4.1 The code and its center

Given a matrix $A \in \mathbb{F}_q^{b \times k}$, we construct a vector $w_A \in \mathbb{F}_q^n$ whose coordinates are indexed by $(x, y) \in \mathbb{F}_q^k \times \mathbb{F}_q^b$. We view each vector $w \in \mathbb{F}_q^n$ as a matrix with rows indexed by x and columns indexed by y , and write $w(x, y)$ for the corresponding entry. Define

$$w_A(x, y) = \mathbf{1}[y = Ax].$$

Each of its q^k rows has one nonzero entry, equal to one, so $\text{wt}(w_A) = t$.

Note that $\mathbf{1}[a = 0] = 1 - a^{q-1}$ for every $a \in \mathbb{F}_q$. Thus w_A has the polynomial representation

$$w_A(x, y) = \prod_{j=1}^b (1 - (y_j - (Ax)_j)^{q-1}).$$

In the expansion of this product, each factor contributes either a constant or a homogeneous polynomial of degree $q - 1$. Hence every monomial in the product has total degree at most $D = b(q - 1)$ and total degree divisible by $q - 1$.

Put $r = k + b$ and $X = (x_1, \dots, x_k, y_1, \dots, y_b)$. We define $\mathcal{E}_{k,b} \leq \mathbb{F}_q^n$ to be the linear code spanned by the evaluations on $\mathbb{F}_q^r$ of the monomials

$$X^\nu, \quad \nu \in \mathbb{Z}_{\geq 0}^r, \quad |\nu| \leq D, \quad |\nu| \equiv 0 \pmod{q-1}, \quad (4)$$

where $X^\nu = \prod_{i=1}^r X_i^{\nu_i}$ and $|\nu| = \sum_{i=1}^r \nu_i$. By the product representation above, every w_A belongs to $\mathcal{E}_{k,b}$.

Since these indicators w_A have weight t , the code $\mathcal{E}_{k,b}$ itself does not have distance greater than t . We instead choose the center $\sigma = -w_0$, where $w_0(x, y) = \mathbf{1}[y = 0]$ corresponds to the zero matrix, and set $c_A = w_A - w_0$. Define

$$C = \text{span}\{c_A : A \in \mathbb{F}_q^{b \times k}\} \leq \mathcal{E}_{k,b}.$$

All indicators agree on the row $x = 0$, where $w_A(0, y) = \mathbf{1}[y = 0]$. Thus every $c \in C$ vanishes on this row. In particular, every codeword vanishes at the origin $(0, 0)$.

We remark that $\mathcal{E}_{k,b}$ is a subcode of the Reed–Muller code $\text{RM}_q(D, r)$, which evaluates all polynomials of total degree at most D on $\mathbb{F}_q^r$. The additional condition that every monomial has degree divisible by $q - 1$, together with vanishing at the origin, relates C to the projective Reed–Muller code $\text{PRM}_q(D, r - 1)$ [Sor91, GL24], which evaluates all homogeneous polynomials of degree D in r variables at one representative from each class $\{\lambda P : \lambda \in \mathbb{F}_q^*\}$, where $P \in \mathbb{F}_q^r \setminus \{0\}$. Its distance bound yields the desired local density, as we show in Lemma 4.1.

4.2 Local density

Define

$$\mathcal{E}_{k,b}^0 = \{c \in \mathcal{E}_{k,b} : c(0, 0) = 0\}.$$

Then

$$C \leq \mathcal{E}_{k,b}^0 \leq \mathcal{E}_{k,b}.$$

Lemma 4.1. *For all integers $k, b \geq 1$,*

$$d(\mathcal{E}_{k,b}^0) = 2(q - 1)q^{k-1} = \rho t.$$

Proof. A projective point $[P]$ is the class $\{\lambda P : \lambda \in \mathbb{F}_q^*\}$ of a nonzero $P \in \mathbb{F}_q^r$. The code $\text{PRM}_q(D, r-1)$ evaluates all homogeneous degree- D polynomials on one representative from each of these $(q^r - 1)/(q - 1)$ classes. Here $r - 1$ is the projective dimension; the polynomials have r variables. Since $q - 1 \mid D$, we have $F(\lambda P) = \lambda^D F(P) = F(P)$, so the evaluations do not depend on the representatives.

Every word in $\mathcal{E}_{k,b}^0$ has a polynomial representative with no constant term and with monomials satisfying (4). Each monomial X^ν therefore has degree $d = a(q - 1)$ for some $1 \leq a \leq b$. Choose j with $\nu_j > 0$ and replace the monomial by $X^\nu X_j^{D-d}$. The degree becomes D , and the evaluations at any point $P = (P_1, \dots, P_r) \in \mathbb{F}_q^r$ are unchanged: the added factor is one when $P_j \neq 0$, because $D - d = (b - a)(q - 1)$, and both monomials vanish when $P_j = 0$. Applying this to every monomial produces a homogeneous degree- D polynomial with the same evaluations. Conversely, every homogeneous degree- D polynomial uses only monomials allowed in (4) and vanishes at the origin. Thus the two polynomial spaces have the same evaluation vectors on $\mathbb{F}_q^r$.

These vectors vanish at the origin and are constant on each projective class. Hence, up to coordinate ordering,

$$\mathcal{E}_{k,b}^0 = \{(0, \text{Rep}_{q-1}(u)) : u \in \text{PRM}_q(D, r-1)\},$$

where Rep_{q-1} repeats each coordinate $q - 1$ times. The projective Reed–Muller distance bound in [Sor91, Theorem 1]³ (see [GL24, Corollary 3.3] for another independent proof) applies since $1 \leq D \leq (r - 1)(q - 1)$. Writing $D - 1 = (b - 1)(q - 1) + (q - 2)$ gives

$$d(\text{PRM}_q(D, r-1)) = (q - (q - 2))q^{(r-1)-(b-1)-1} = 2q^{k-1}.$$

The zero coordinate contributes no weight, while repetition multiplies every weight by $q - 1$. The claimed distance follows. $\square$

Since $C \leq \mathcal{E}_{k,b}^0$ and $c_A - \sigma = w_A$, the lemma gives

$$d(C) \geq \rho t > t, \quad \text{wt}(c_A - \sigma) = t,$$

i.e., the codewords c_A all lie at distance t from the center σ , even though every nonzero codeword of C has weight greater than t . This proves (i).

4.3 Rigidity

The rigidity property (ii) is equivalent to the statement that every $w \in -\sigma + C$ either equals some w_A or has weight at least αt . For an indicator $w_A \in -\sigma + C$, the zeroth and first moments of the row indexed by x are, respectively,

$$\sum_y w_A(x, y) = 1, \quad \sum_y y w_A(x, y) = Ax.$$

Since $\sigma = -w_0$ and C is spanned by the differences $w_B - w_0$, every $w \in -\sigma + C$ can be written as

$$w = w_0 + \sum_{B \in \mathbb{F}_q^{b \times k}} \lambda_B (w_B - w_0), \quad \lambda_B \in \mathbb{F}_q.$$

³The original proof contains a gap, which was later repaired in [Sor23].

Since every row of w_0 has zeroth moment 1 and first moment 0, linearity gives, for every $x \in \mathbb{F}_q^k$,

$$\begin{aligned}\sum_y w(x, y) &= 1 + \sum_B \lambda_B(1 - 1) = 1, \\ \sum_y y w(x, y) &= 0 + \sum_B \lambda_B(Bx - 0) = \left(\sum_B \lambda_B B \right) x.\end{aligned}$$

The first-moment identity shows that the same matrix $\sum_B \lambda_B B$ determines the first moments of all rows. To recover this matrix from w , let $e_1, \dots, e_k$ be the standard basis of $\mathbb{F}_q^k$ and define the linear map $R : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{b \times k}$ by

$$(Rw)_{*,i} = \sum_{y \in \mathbb{F}_q^b} y w(e_i, y), \quad i \in [k].$$

For $w \in -\sigma + C$, the first-moment identity at $x = e_i$ gives $(Rw)_{*,i} = (\sum_B \lambda_B B)e_i$, which is exactly the i th column of that matrix. Hence

$$Rw = \sum_B \lambda_B B, \quad \sum_y y w(x, y) = (Rw)x \quad \text{for every } x \in \mathbb{F}_q^k.$$

In particular, $Rw_A = A$, so R recovers the matrix encoded by an indicator.

To prove (ii), we compare $w = c - \sigma$ with w_A , where $c \in C$ and $A = Rw$.

Lemma 4.2. *Every $c \in C$ either equals c_A for some $A \in \mathbb{F}_q^{b \times k}$ or satisfies*

$$\text{wt}(c - \sigma) \geq \alpha t.$$

Proof. Put $w = c - \sigma$ and $A = Rw$. Let r_x count the nonzero entries in the row of w indexed by x . Each row sums to one, so $r_x \geq 1$. If $r_x = 1$, its only nonzero entry equals one, and its location is its first moment, namely Ax . This row agrees with the corresponding row of w_A . If $r_x \geq 2$, the corresponding row of $w - w_A$ has at most $r_x + 1$ nonzero entries, and $r_x + 1 \leq 3(r_x - 1)$. Summing over the $q^k = t$ rows gives

$$\text{wt}(w - w_A) \leq 3 \sum_{x \in \mathbb{F}_q^k} (r_x - 1) = 3(\text{wt}(w) - t).$$

If $c \neq c_A$, then $w - w_A = c - c_A$ is a nonzero word of C , so its weight is at least $d(C) \geq \rho t$. Combining the two bounds gives $\text{wt}(c - \sigma) = \text{wt}(w) \geq t + \rho t/3 = \alpha t$. $\square$

4.4 Linear decodability of labels

For $A = [a_1 \cdots a_k]$, the row of w_A indexed by e_i has its nonzero entry at $y = Ae_i = a_i$. Weighting the entries in this row by $v(y)$ therefore returns $v(a_i)$. Define

$$T : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^h, \quad T(w) = \sum_{i=1}^k \sum_{y \in \mathbb{F}_q^b} w(e_i, y) v(y). \quad (5)$$

This is linear in w , since the table entries $v(y)$ are fixed coefficients, and $T(w_A) = \sum_i v(a_i)$. The assumption $v(0) = 0$ gives $T(w_0) = 0$, so

$$T(\sigma) = 0, \quad T(c_A) = T(w_A - w_0) = \sum_{i=1}^k v(a_i).$$

This proves (iii).

4.5 Construction time

We construct a basis of C without enumerating the q^{bk} matrices A . Let $E_{ij} \in \mathbb{F}_q^{b \times k}$ have a one in position (i, j) and zeros elsewhere, and set

$$\mathcal{B} = \{aE_{ij} : i \in [b], j \in [k], a \in \mathbb{F}_q^*\}.$$

Every matrix A is a sum of elements of $\mathcal{B}$. For $B \in \mathcal{B}$, define the coordinate permutation

$$(P_B w)(x, y) = w(x, y - Bx).$$

Then $P_B w_A = w_{A+B}$, so

$$P_B c_A = c_{A+B} - c_B, \quad c_{A+B} = P_B c_A + c_B.$$

The first identity shows that P_B maps C into C .

Start with a basis of $\text{span}\{c_B : B \in \mathcal{B}\}$. For each basis vector u and each $B \in \mathcal{B}$, compute $P_B u$ and add it to the basis if it is linearly independent of the vectors already collected. Process newly added basis vectors in the same way. All collected vectors belong to C , and their final span is invariant under every P_B . It contains every c_B with $B \in \mathcal{B}$, so the second identity, applied along a decomposition of A into elements of $\mathcal{B}$, shows that it contains every c_A . Thus the final basis spans C .

There are $(q-1)bk$ permutations and at most n basis vectors. Processing all basis vectors requires in total at most $(q-1)bk n$ permutation applications and linear-independence checks. Each takes $n^{O(1)}$ field operations, so this constructs C in $(nkb)^{O(1)}$ time. Evaluate w_0 to obtain σ , and use (5) to form the $h \times n$ matrix of T from the $hq^b \leq hn$ entries of the input table. Together these steps take $(nkbh)^{O(1)}$ time, with polynomial exponents independent of k, b, h . Given any requested A , evaluating $w_A - w_0$ directly computes c_A in $(nkb)^{O(1)}$ time.

This proves (iv) and finishes the proof of [Lemma 3.1](#).

AI Methodology

The proof was discovered by GPT-6 in response to the authors' prompts about deterministic parameterized inapproximability of coding problems. The AI originally gave a reduction from gap NCP that is a clever modification of the DMS reduction and achieved a constant-factor blow-up in the parameter k , but whose correctness proof relied on some advanced concepts in algebraic geometry. After many further rounds of interaction, the current approach, which starts from exact NCP and is elementary, was discovered. AI tools also assisted with exposition and verification. The authors checked the proofs and prepared the paper and take responsibility for the correctness of the results and the quality of exposition.

References

- [ABSS97] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997. doi:10.1006/jcss.1997.1472.
- [AK14] Per Austrin and Subhash Khot. A simple deterministic reduction for the gap minimum distance of code problem. *IEEE Transactions on Information Theory*, 60(10):6636–6645, 2014. doi:10.1109/TIT.2014.2340869.
- [BCGR24] Huck Bennett, Mahdi Cheraghchi, Venkatesan Guruswami, and João Ribeiro. Parameterized inapproximability of the minimum distance problem over all fields and the shortest vector problem in all ℓ_p norms. *SIAM Journal on Computing*, 53(5):1439–1475, 2024. doi:10.1137/23M1573021.
- [BMT78] Elwyn R. Berlekamp, Robert J. McEliece, and Henk C. A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3):384–386, 1978. doi:10.1109/TIT.1978.1055873.
- [BBE⁺21] Arnab Bhattacharyya, Édouard Bonnet, László Egri, Suprovat Ghoshal, Karthik C. S., Bingkai Lin, Pasin Manurangsi, and Dániel Marx. Parameterized intractability of Even Set and Shortest Vector Problem. *Journal of the ACM*, 68(3):16:1–16:40, 2021. doi:10.1145/3444942.
- [BGLR25] Vijay Bhattiprolu, Venkatesan Guruswami, Euiwoong Lee, and Xuandi Ren. Inapproximability of finding sparse vectors in codes, subspaces, and lattices. In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1295–1303, 2025. doi:10.1109/FOCS63196.2025.00068.
- [CW12] Qi Cheng and Daqing Wan. A deterministic reduction for the gap minimum distance problem. *IEEE Transactions on Information Theory*, 58(11):6935–6941, 2012. doi:10.1109/TIT.2012.2209198.
- [DF13] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.
- [DFVW99] Rod G. Downey, Michael R. Fellows, Alexander Vardy, and Geoff Whittle. The parametrized complexity of some fundamental problems in coding theory. *SIAM Journal on Computing*, 29(2):545–570, 1999. doi:10.1137/S0097539797323571.
- [DMS03] Ilya Dumer, Daniele Micciancio, and Madhu Sudan. Hardness of approximating the minimum distance of a linear code. *IEEE Transactions on Information Theory*, 49(1):22–37, 2003. doi:10.1109/TIT.2002.806118.
- [GL24] Sudhir R. Ghorpade and Rati Ludhani. On the minimum distance, minimum weight codewords, and the dimension of projective Reed–Muller codes. *Advances in Mathematics of Communications*, 18(2):360–382, 2024. doi:10.3934/amc.2023035.
- [HO26] Shuichi Hirahara and Kazuki Ogitsuka. One-sided-error parameterized reductions for the minimum distance and shortest vector problems. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2026)*, volume 392 of *Leibniz International Proceedings in Informatics*, pages 46:1–46:22. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026. doi:10.4230/LIPIcs.APPROX/RANDOM.2026.46.
- [Mic14] Daniele Micciancio. Locally dense codes. In *Proceedings of the 29th IEEE Conference on Computational Complexity (CCC)*, pages 90–97, 2014. doi:10.1109/CCC.2014.17.
- [Sor91] Anders Bjært Sørensen. Projective Reed–Muller codes. *IEEE Transactions on Information Theory*, 37(6):1567–1576, 1991. doi:10.1109/18.104317.

- [Sor23] Anders Bjaert Sørensen. A note on a gap in the proof of the minimum distance for projective Reed–Muller codes. Preprint, 2023. arXiv:2310.03574.
- [Ste93] Jacques Stern. Approximating the number of error locations within a constant ratio is NP-complete. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (AAECC-10)*, volume 673 of *Lecture Notes in Computer Science*, pages 325–331. Springer, 1993. doi:10.1007/3-540-56686-4_54.
- [Var97] Alexander Vardy. The intractability of computing the minimum distance of a code. *IEEE Transactions on Information Theory*, 43(6):1757–1766, 1997. doi:10.1109/18.641542.

A Amplifying the gap

In this section, we show how tensoring multiplies both code distance and coset distance. We use a symmetric variant of the classical NCP gap-amplification construction [Ste93, ABSS97].

For vectors $x \in \mathbb{F}_q^{n_1}$ and $y \in \mathbb{F}_q^{n_2}$, the tensor $x \otimes y$ is the matrix $xy^\top$, viewed as a vector of length $n_1 n_2$. For linear codes $E \leq \mathbb{F}_q^{n_1}$ and $F \leq \mathbb{F}_q^{n_2}$, define $E \otimes F$ to be the span of $x \otimes y$ with $x \in E$ and $y \in F$. Every column of such a matrix lies in E , and every row lies in F .

Lemma A.1. *For $i \in \{1, 2\}$, let $V_i \leq \mathbb{F}_q^{n_i}$ have dimension r_i , let $z_i \in \mathbb{F}_q^{n_i} \setminus V_i$, and put $\mathcal{D}_i = V_i + \mathbb{F}_q z_i$. One can deterministically construct a code $V' \leq \mathbb{F}_q^{n_1 n_2}$ and a vector $z' \notin V'$ such that, writing $\mathcal{D}' = V' + \mathbb{F}_q z'$, we have*

$$\begin{aligned} d(\mathcal{D}') &= d(\mathcal{D}_1)d(\mathcal{D}_2), \\ \delta(z' + V') &= \delta(z_1 + V_1)\delta(z_2 + V_2), \end{aligned}$$

and

$$\dim V' = (r_1 + 1)(r_2 + 1) - 1, \quad \dim \mathcal{D}' = (r_1 + 1)(r_2 + 1).$$

The construction takes time polynomial in the input size and $n_1 n_2$.

Proof. Set $\mathcal{D}' = \mathcal{D}_1 \otimes \mathcal{D}_2$ and $z' = z_1 \otimes z_2$, and define

$$V' = (V_1 \otimes V_2) + (V_1 \otimes \mathbb{F}_q z_2) + (\mathbb{F}_q z_1 \otimes V_2).$$

Expanding $\mathcal{D}_1 \otimes \mathcal{D}_2$ gives $\mathcal{D}' = V' + \mathbb{F}_q z'$.

The code-distance identity is the standard tensor-product property [BGLR25]. For completeness, any nonzero $M \in \mathcal{D}_1 \otimes \mathcal{D}_2$ has a nonzero row of weight at least $d(\mathcal{D}_2)$. Thus at least $d(\mathcal{D}_2)$ columns are nonzero, and each has weight at least $d(\mathcal{D}_1)$. This gives $\text{wt}(M) \geq d(\mathcal{D}_1)d(\mathcal{D}_2)$. The outer product of two minimum-weight nonzero codewords attains equality.

For the coset-distance identity, every $M \in z' + V'$ can be written as

$$M = z_1 z_2^\top + N + a z_2^\top + z_1 b^\top, \quad N \in V_1 \otimes V_2, \quad a \in V_1, \quad b \in V_2.$$

Put $u = z_2 + b \in z_2 + V_2$. The j th column is

$$M_{*,j} = N_{*,j} + (z_2)_j a + u_j z_1.$$

The first two terms lie in V_1 . Whenever $u_j \neq 0$, dividing this column by u_j therefore gives a vector in $z_1 + V_1$, so the column has weight at least $\delta(z_1 + V_1)$. Since $u \in z_2 + V_2$, at least $\delta(z_2 + V_2)$ entries of u are nonzero. Summing the corresponding column weights gives $\text{wt}(M) \geq \delta(z_1 + V_1)\delta(z_2 + V_2)$.

Conversely, if $w_i \in z_i + V_i$ attains $\delta(z_i + V_i)$ for each i , then $w_1 w_2^\top \in z' + V'$ and its weight is the product of their weights. This proves equality.

Finally, let G_i have as columns a basis of V_i followed by z_i . These columns form a basis of $\mathcal{D}_i$ because $z_i \notin V_i$. The matrix $G_1 \otimes G_2$, whose columns are all pairwise tensor products of the columns of G_1 and G_2 , is a basis matrix for $\mathcal{D}'$. Deleting its column z' leaves a basis of V' . Hence $z' \notin V'$, $\dim V' = (r_1 + 1)(r_2 + 1) - 1$, and $\dim \mathcal{D}' = (r_1 + 1)(r_2 + 1)$. The generator matrix has $n_1 n_2 (r_1 + 1)(r_2 + 1)$ entries, proving the construction-time bound. $\square$

Now start with the NCP instance (V, z, t) of length n from [Proposition 1.2](#). Applying [Lemma 2.4](#) gives the MDP instance $(\mathcal{D}, t)$, where $\mathcal{D} = V + \mathbb{F}_q z$. By (1), $d(\mathcal{D}) = t$ in the YES case and $d(\mathcal{D}) \geq \min\{\rho, \alpha\}t$ in the NO case. Put $r = \dim V$. Set $V_1 = V$ and $z_\ell = z^{\otimes \ell}$ for $\ell \geq 1$. For $\ell \geq 2$, define

$$V_\ell = (V_{\ell-1} \otimes V) + (V_{\ell-1} \otimes \mathbb{F}_q z) + (\mathbb{F}_q z_{\ell-1} \otimes V).$$

Iterating [Lemma A.1](#) gives $z_\ell \notin V_\ell$ and $\mathcal{D}_\ell := V_\ell + \mathbb{F}_q z_\ell = \mathcal{D}^{\otimes \ell}$. The amplified instances are (V_ℓ, z_ℓ, t^ℓ) and $(\mathcal{D}_\ell, t^\ell)$. The parameters change as follows:

	Original instance	Tensor power
Field	$\mathbb{F}_q$	$\mathbb{F}_q$
Length	n	n^ℓ
Threshold	t	t^ℓ
NCP code dimension	r	$(r + 1)^\ell - 1$
MDP code dimension	$r + 1$	$(r + 1)^\ell$

Moreover,

$$\delta(z_\ell + V_\ell) = \delta(z + V)^\ell, \quad d(\mathcal{D}_\ell) = d(\mathcal{D})^\ell.$$

Thus the YES optima are t^ℓ , while the NO optima are at least $\alpha^\ell t^\ell$ for NCP and $(\min\{\rho, \alpha\})^\ell t^\ell$ for MDP. For any fixed target gap $\gamma \geq 1$, choose an integer $\ell \geq 1$ with $(4/3)^\ell > \gamma$. Since $\alpha \geq 13/9$ and $\min\{\rho, \alpha\} \geq 4/3$ for $q > 2$, both NO optima are then strictly greater than γt^ℓ .

For a source instance of length m and parameter k , we have $t = q^k$ and $n = O_q(q^k m)$, so the output parameter and length are

$$t^\ell = q^{k\ell}, \quad n^\ell = O_{q,\gamma}(q^{k\ell} m^\ell).$$

For fixed q and γ , the running time is $f_{q,\gamma}(k)m^{O_\gamma(1)}$ for some computable function $f_{q,\gamma}$. Each input produces one instance of each output problem, so these are deterministic fixed-parameter many-one reductions.

B The binary case

Write $\mathbb{F}_4 = \mathbb{F}_2(\omega)$ with $\omega^2 + \omega + 1 = 0$ and apply the map

$$\text{Enc} : \mathbb{F}_4 \longrightarrow \mathbb{F}_2^3, \quad \text{Enc}(a + b\omega) = (a, b, a + b), \quad a, b \in \mathbb{F}_2,$$

coordinatewise. This is the binary $[3, 2, 2]$ parity-check code: Enc is $\mathbb{F}_2$ -linear, and every nonzero symbol has an image of weight two. Hence, for every $\mathbb{F}_4$ -linear code E and vector z ,

$$d(\text{Enc}(E)) = 2d(E), \quad \delta(\text{Enc}(z) + \text{Enc}(E)) = 2\delta(z + E).$$

The same conversion also preserves the RNC promise: $t < \theta d(E)$ implies $2t < \theta d(\text{Enc}(E))$. If $g_1, \dots, g_d$ is an $\mathbb{F}_4$ -basis of E , a binary generator matrix has columns $\text{Enc}(g_j), \text{Enc}(\omega g_j)$ for $j \in [d]$. Thus concatenation triples the length, doubles the threshold, and preserves the gap. Apply it to the amplified NCP and MDP instances over $\mathbb{F}_4$ obtained in [Appendix A](#) from [Proposition 1.2](#) and [Lemma 2.4](#). This proves [Theorem 1.1](#) over $\mathbb{F}_2$.

The reduction can also start from the exact binary nearest codeword problem. For $U \leq \mathbb{F}_2^m$ and $\tau \in \mathbb{F}_2^m$, let $U_4 = \text{span}_{\mathbb{F}_4} U$. Every word of $\tau + U_4$ has a unique expression $a + \omega b$ with $a \in \tau + U$ and $b \in U$. Its support is $\text{supp}(a) \cup \text{supp}(b)$, so

$$\delta(\tau + U_4) = \delta(\tau + U).$$

Scalar extension at the input and binary concatenation at the output therefore give the desired reduction from the exact binary NCP itself.