Tight Lower Bound for Approximating Parametrized Maximum Likelihood Decoding Under ETH

Rishav Gupta
National University of Singapore, Singapore

Bingkai Lin
State Key Laboratory for Novel Software Technology, Nanjing University, China

Xin Zheng
State Key Laboratory for Novel Software Technology, Nanjing University, China

Abstract

We present a simple deterministic reduction which, assuming the Exponential Time Hypothesis (ETH), yields tight lower bounds for approximating the parameterized Maximum Likelihood Decoding problem (MLD) and the parameterized Nearest Codeword Problem (NCP) within some fixed constant factor. Our starting point is the ETH-based exponential-time hardness of (c, s)-Gap MAXLIN established in [7]. We transform a (c, s)-Gap MAXLIN instance into an instance of γ-Gap k-MLD via a novel combinatorial object that we call a *cover family*. We provide both a randomized construction of the required cover families and a subsequent derandomization. Prior to our work, $n^{\Omega(k)}$ hardness for constant-factor approximation was only shown under the randomized Gap Exponential Time Hypothesis Gap-ETH [16], which is a much stronger assumption than ETH. Under ETH, the strongest known lower bound was $n^{\Omega(k/\operatorname{poly}\log k)}$ due to [3]. Unlike previous approaches that rely on reductions from the hardness of approximating 2-CSP, our reduction provides a more direct and conceptually simpler route to achieving the optimal lower bounds.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Maximum Likelihood Decoding, Parameterized Complexity, Hardness of Approximation, Exponential Time Hypothesis

Digital Object Identifier 10.4230/LIPIcs.CCC.2026.1

Related Version *Full Version*: https://arxiv.org/abs/2605.08797

Funding Supported by the "111 Cente" (No. B26023) and NRF grant NRF-NRFI09-0005.

Acknowledgements The authors would like to thank the anonymous reviewers for their comments and suggestions which help to improve the paper significantly. Special thanks to Reviewer B for simplifying the proof for Lemma 29.

1 Introduction

The study of error-correcting codes gives rise to various computational problems. In the context of channel coding, one of the most fundamental tasks is to recover the original message from a signal corrupted by noise. In this problem, which is known as the *Maximum Likelihood Decoding* (MLD) problem, we are given the parity check matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times n}$ of a linear code, a vector $\boldsymbol{t} \in \mathbb{F}_q^d$, and a parameter $k \in \mathbb{N}$. Our goal is to find a vector $\boldsymbol{x} \in \mathbb{F}_q^n$ of Hamming weight at most k such that $\boldsymbol{Hx} = \boldsymbol{t}$. This problem is computationally equivalent to the *Nearest Codeword Problem* (NCP), where we are given the generator matrix $\boldsymbol{A} \in \mathbb{F}_q^{d\times n}$ and the goal is to find a vector $\boldsymbol{x} \in \mathbb{F}_q^n$ such that $\|\boldsymbol{Ax} - \boldsymbol{t}\|_0 \leq k$.

The computational hardness of these problems is well-established. They were proven to be NP-hard decades ago [5] and are known to be hard to approximate within any constant factor in polynomial time [2, 11]. Consequently, attention shifted to the framework of *parameterized complexity* [9], where the runtime is analyzed with respect to the parameter k.

In the parameterized setting, the hardness of the exact problem was settled at first. Downey and Fellows [10] proved that k-MLD parameterized by k is W[1]-hard, ruling out exact FPT algorithms, i.e., algorithms that run in $f(k) \cdot n^{O(1)}$ time, under W[1] $\neq$ FPT. However, the question of whether *approximation* could yield fixed parameter tractability remained a major open problem for over two decades.

This question was finally resolved in a series of breakthrough results. The authors of [6] established that k-NCP is W[1]-hard to approximate within any constant factor. They also proved the parameterized inapproximability of the *Minimum Distance Problem* (MDP) over $\mathbb{F}_2$ and the *Shortest Vector Problem* (SVP) in ℓ_p norm for every $p > 1$. Following this, Bennett, Cheraghchi, Guruswami, and Ribeiro [4] generalized these results, proving the inapproximability of k-MDP over all finite fields and k-SVP in general ℓ_p norms. These works painted a comprehensive picture of the FPT inapproximability of coding problems. With the W[1]-hardness established, the focus has shifted to the fine-grained complexity:

What is the precise running time required to approximate k-MLD when k is small?

Under the standard gap-free hypothesis, the *Exponential Time Hypothesis* (ETH), [6] ruled out $n^{O((\log k)^{1/2.01})}$-time algorithms for constant gap k-NCP. Their reduction incurred an exponential blow-up in the parameter, which results in the relatively weak lower bound. Very recently, Li, Lin, and Liu [15] developed a new gap-creating reduction with polynomial parameter growth, hence improving the lower bound under standard ETH, ruling out $n^{o(k^{1/3})}$-time algorithms that solve γ-Gap k-NCP for any $1 < \gamma < \frac{3}{2}$. Through a self-reduction, they further ruled out $n^{o(k^{\varepsilon})}$-time algorithms where $\varepsilon = \frac{1}{\operatorname{poly}\log\gamma}$ for any constant gap $\gamma > 1$. While this was a major step forward, a substantial gap remains between the lower bound and the trivial $n^{O(k)}$ upper bound.

To date, the only tight lower bound comes from the work of Manurangsi [16], which showed that no approximation algorithm exists for these problems in time $n^{o(k)}$ for any constant factor. However, this result relies on the randomized *Gap-Exponential Time Hypothesis* (Gap-ETH), which is a significantly stronger assumption than ETH.

The above-mentioned results establish hardness for *all* constant approximation factors. However, even obtaining an optimal lower bound under ETH for *some* constant approximation factor remained open. Prior to our work, the strongest known ETH-based hardness for *some* approximation factor was established via a chain of reductions: first, by invoking the hardness of approximating k-variable 2-CSP [12, 3], then reducing k-variable 2-CSP to k-Exact Cover [13], and finally reducing k-Exact Cover to γ-Gap k-MLD$_q$ [2]. This approach yields suboptimal running-time lower bounds: namely $n^{k/2^{\Omega(\sqrt{\log k}\cdot\log\log k)}}$ via [12] and $n^{k/\log^{C}(k)}$ via [3].

Our Contribution. In this work, we prove that the tight $n^{\Omega(k)}$ lower bound holds from some constant approximation factor under the standard non-deterministic Exponential Time Hypothesis, removing the need for the stronger Gap-ETH assumption.

▶ **Theorem 1** (Main Result). *Assuming* ETH*, there exist constants $\gamma > 1$ and $\delta > 0$, such that for any algorithm that takes as input a matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times N}$, a vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and a parameter $k \in \mathbb{N}$, it must take $N^{\delta k}$ time to distinguish the following two cases:*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le k$ such that $\boldsymbol{Hx} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le \gamma k$, we have $\boldsymbol{Hx} \neq \boldsymbol{u}$.*

If we allow k to be some function of N, the proof of Theorem 1 actually gives an $N^{\delta k}$-time lower bound for γ-Gap k-MLD_q when $k \le O((\log\log N)^{0.49})$. Assuming randomized ETH, we can obtain the same lower bound for $k \le N^{\varepsilon}$ for some constant $\varepsilon > 0$.

▶ **Theorem 2.** *Assuming randomized* ETH*, there exist constants $\gamma > 1$ and $\delta, \varepsilon > 0$, such that for any algorithm that takes as input a matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times N}$, a vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and an integer $k \in \mathbb{N}$ such that $2 \le k \le N^{\varepsilon}$, it must take $N^{\delta k}$ time to distinguish the following two cases:*
- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le k$ such that $\boldsymbol{Hx} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le \gamma k$, we have $\boldsymbol{Hx} \ne \boldsymbol{u}$.*

Our Techniques. Our hardness proof builds upon the ETH-hardness of (c, s)-Gap MAXLIN_q established by [7]. In this problem, one must distinguish whether a system of linear equations is almost satisfiable (a c-fraction of equations hold), or is far from satisfiable (at most an s-fraction hold). Specifically, the result of [7] shows that assuming ETH, there is no algorithm that can solve (c, s)-Gap MAXLIN_q in time $2^{o(n)}$. By considering the dual of a hard (c, s)-Gap MAXLIN_q instance, we first obtain a hard instance of γ-Gap MLD_q, where the goal is to determine whether a target vector can be expressed as a linear combination of at most ℓ columns from a matrix, or it requires at least $\gamma\ell$ columns. This duality transformation implies that the non-parameterized version of γ-Gap k-MLD_q also cannot be solved in time $2^{o(n)}$.

The core of our contribution is a reduction from this non-parameterized γ-Gap MLD_q to its parameterized version γ-Gap k-MLD_q. We let each vector in the output k-MLD instance be the sum of ℓ/k vectors in the original MLD instance. Thus, a solution of size ℓ in the MLD instance can be represented by only k vectors in the k-MLD instance. However, the naive brute-force enumeration of all combinations of size ℓ/k produces an instance of size $N = \binom{n}{\ell/k}$, hence the reduction yields only an $N^{\Omega(k/\log k)}$ lower bound for γ-Gap k-MLD_q.

We achieve a tight lower bound via a novel combinatorial structure we term *cover families.* A cover family is a sufficiently small collection of subsets $\mathcal{S}$ over a universe $\mathcal{U}$, such that any *small* subset of $\mathcal{U}$ can be expressed as the union of k pairwise disjoint sets from $\mathcal{S}$, whereas no *large* subset of $\mathcal{U}$ can be expressed as the union of even γk sets from $\mathcal{S}$. To construct suitable cover families, we introduce an intermediate combinatorial object called *balanced partition families.* A balanced partition family is a collection of partitions $\mathcal{P}$ of a universe $\mathcal{U}$ such that every partition in the family is roughly balanced, and every subset of $\mathcal{U}$ of an appropriate size is (almost) equipartitioned by some partition $P \in \mathcal{P}$. We first show that independently sampling random partitions suffices to obtain the desired balanced partition family. We then give a deterministic construction of such balanced partition families using hypercube set systems. Finally we use a standard reduction from γ-Gap k-MLD_q to γ-Gap k-NCP_q to obtain the hardness of the latter.

Future Directions. Assuming ETH, we obtain optimal lower bounds for γ-Gap k-MLD_q and γ-Gap k-NCP_q for some constant approximation factor $\gamma > 1$. A natural question is whether this lower bound can be extended to an arbitrary constant factor assuming only ETH. In contrast, under Gap-ETH the corresponding conclusion is known to hold for every constant $\gamma > 1$ via [16]. Establishing such a result under ETH would close the current knowledge gap between ETH and Gap-ETH through the lens of parameterized coding-theoretic problems. It is worth mentioning that [3] rules out $n^{k/\log^C k}$-time algorithms that approximates k-NCP and k-CVP to any constant factor under ETH. We also wonder whether $n^{\Omega(k)}$ lower bound for γ-Gap k-MDP or other related problems can be established under ETH.

Another open direction is to understand what further consequences can be derived from the ETH-hardness of (c, s)-Gap MAXLIN_q established in [7]. Beyond our work, the only other result we are aware of that leverages this hardness result is [1]. Lastly, one can also explore other applications of the combinatorial objects, *cover family* and *balanced partition family* which are introduced in this paper.

Paper Organization. In Section 2 we define necessary notation and introduce useful tools from the literature. First in Section 3 we present a reduction from (c, s)-Gap MAXLIN_q to γ-Gap k-MLD_q using *cover family*. Then in Section 4 we give a randomized as well as a deterministic construction of a suitable *cover family* using an intermediate combinatorial object, *balanced partition family*.

2 Preliminaries

We begin by formally defining the computational problems that will be studied throughout the paper. For each problem, we specify its input, the underlying computational goal, and any associated promise conditions.

2.1 Computational Problems

We first define the standard k-SAT problem. For an integer $k \geq 2$, a k-SAT formula over n boolean variables is the conjunction of clauses, where each clause is the disjunction of k literals. That is, k-SAT formulas have the form $\bigwedge_{i=1}^{m} \bigvee_{j=1}^{k} b_{i,j}$, where $b_{i,j} = x_k$ or $b_{i,j} = \neg x_k$ for some boolean variable x_k.

▶ **Definition 3** (k-SAT). *For any $k \geq 2$, the decision problem k-SAT is defined as follows. The input is a k-SAT formula. It is a YES instance if there exists an assignment to the variables that makes the formula evaluate to true and a NO instance otherwise.*

We write k-SAT_C for a k-SAT instance where each variable x_i is contained in at most C clauses. We also define the corresponding optimization version Max-k-SAT.

▶ **Definition 4** (Max-k-SAT). *For any $k \geq 2$, the decision problem Max-k-SAT is defined as follows. The input is a k-SAT formula and an integer $S \geq 1$. It is a YES instance if there exists an assignment to the variables such that at least S of the clauses evaluate to true and a NO instance otherwise.*

▶ **Definition 5** ((c, s)-Gap MAXLIN_q). *A (c, s)-Gap MAXLIN_q instance $(\boldsymbol{A}, \boldsymbol{b})$ consists of an $m \times n$ matrix $\boldsymbol{A} \in \mathbb{F}_q^{m \times n}$ and a vector $\boldsymbol{b} \in \mathbb{F}_q^m$. The objective of the problem is to distinguish between the following cases.*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^n$ such that $\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_0 \leq (1 - c)m$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^n$, $\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_0 > (1 - s)m$.*

We need to mention that this problem is equivalent to the non-parameterized $\frac{1-s}{1-c}$-Gap MLD_q. Now we define the dual problem of MAXLIN, Maximum Likelihood Decoding Problem.

▶ **Definition 6** (γ-Gap MLD_q). *A γ-Gap MLD_q instance $(\boldsymbol{H}, \boldsymbol{u}, \ell)$ consists of a $d \times n$ matrix $\boldsymbol{H} \in \mathbb{F}_q^{d \times n}$, a target vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and a value ℓ. The goal is to distinguish between the following cases.*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^n$ with $\|\boldsymbol{x}\|_0 \leq \ell$ such that $\boldsymbol{H}\boldsymbol{x} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^n$ with $\|\boldsymbol{x}\|_0 \leq \gamma\ell$, we have $\boldsymbol{H}\boldsymbol{x} \neq \boldsymbol{u}$.*

We now define the parametrized version of γ-Gap MLD_q, which is almost the same as its non-parameterized version, just using $k := \ell$ as the parameter.

▶ **Definition 7** (γ-Gap k-MLD_q). *A γ-Gap k-MLD_q instance $(\boldsymbol{H}, \boldsymbol{u})$ consists of a $d \times n$ matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times n}$, a target vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and a parameter $k \in \mathbb{N}$. The goal is to distinguish between the following cases.*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^n$ with $\|\boldsymbol{x}\|_0 \le k$ such that $\boldsymbol{Hx} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^n$ with $\|\boldsymbol{x}\|_0 \le \gamma k$, we have $\boldsymbol{Hx} \neq \boldsymbol{u}$.*

Next we turn to the parameterized Nearest Codeword Problem, which is closely related to MLD.

▶ **Definition 8** (γ-Gap k-NCP_q). *A γ-Gap k-NCP_q instance $(\boldsymbol{A}, \boldsymbol{t})$ consists of a $n \times d$ matrix $\boldsymbol{A} \in \mathbb{F}_q^{n\times d}$, a target vector $\boldsymbol{t} \in \mathbb{F}_q^n$ and a parameter $k \in \mathbb{N}$. The goal is to distinguish between the following cases.*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^d$ such that $\|\boldsymbol{Ax} - \boldsymbol{t}\|_0 \le k$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^d$, $\|\boldsymbol{Ax} - \boldsymbol{t}\|_0 > \gamma k$.*

Similarly we define parametrized versions of Closest Vector Problem (CVP). For $\boldsymbol{x} \in \mathbb{Z}^d$ we will use the notation $\|\boldsymbol{x}\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}$, in the following definition.

▶ **Definition 9** (γ-Gap k-CVP_p). *A γ-Gap k-CVP_p instance $(\boldsymbol{A}, \boldsymbol{t})$ consists of a $n \times d$ matrix $\boldsymbol{A} \in \mathbb{Z}^{n\times d}$, a target vector $\boldsymbol{t} \in \mathbb{Z}^n$ and a parameter $k \in \mathbb{N}$. The goal is to distinguish between the following cases.*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{Z}^d$ such that $\|\boldsymbol{Ax} - \boldsymbol{t}\|_p \le k$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{Z}^d$, $\|\boldsymbol{Ax} - \boldsymbol{t}\|_p > \gamma k$.*

2.2 Fine-grained Hardness Assumptions

We introduce the following fine-grained hardness assumptions, Exponential time hypothesis ETH and its corresponding gap version $\mathsf{Gap\text{-}ETH}$ below.

▶ **Definition 10** (Exponential Time Hypothesis (ETH), [14]). *There exists $\delta > 0$ such that any algorithm which solves 3-SAT must take $2^{\delta n}$ time.*

▶ **Definition 11** (Gap-ETH, [8]). *There exists $\delta > 0$ and $0 < \eta < 1$ such that given a 3-SAT instance with n variables and m clauses, any algorithm which can distinguish between the cases if all m clauses are satisfiable and one in which no assignment satisfies more than η-fraction of the clauses, must take $2^{\delta n}$ time.*

We now state the Sparsification Lemma.

▶ **Lemma 12** (Sparsification Lemma, [14]). *Let $\varepsilon > 0$, $k \ge 3$ be constants. There is a $2^{\varepsilon n} \cdot poly(n)$ time algorithm that takes a k-CNF F on n variables and produces $F_1, \ldots, F_{2^{\varepsilon n}}$, $2^{\varepsilon n}$ k-CNFs such that F is satisfied if and only if $\bigvee_i F_i$ is satisfied and each F_i has n variables and $n \cdot \left(\frac{k}{\varepsilon}\right)^{O(k)}$ clauses. In fact, each variable is in at most $\mathrm{poly}\left(\frac{1}{\varepsilon}\right)$ clauses, and the F_i are over the same variables as F.*

We will now state the following result from [7] which gives us a very strong starting point for establishing hardness results.

▶ **Theorem 13** ([7], Theorem 6.3). *For every finite field $\mathbb{F}_q$, there exists constants c and s where $0 < s < c < 1$, such that the following holds. There exists a polynomial time reduction which takes a 3-SAT_C instance with n variables, and outputs a (c, s)-Gap MAXLIN_q with $n' = O(n)$ variables and $m' = O(n)$ equations.*

The sparsification Lemma 12 and Tovey's reduction [17] together tell us that if ETH is true, then 3-SAT$_4$ over n variables must take $2^{\delta n}$ time for some $\delta > 0$. Together with Theorem 13, we find that if ETH holds, then for some $0 < s < c < 1$ and $C > 0$, any algorithm which solves (c,s)-Gap MAXLIN$_q$ with n variables and $m \le Cn$ clauses, must take $2^{\delta n}$ time for some $\delta > 0$. We state it as the following corollary.

▶ **Corollary 14** ((c,s)-Gap MAXLIN$_q$ is ETH-Hard). *For every finite field $\mathbb{F}_q$, there exists constants $0 < s < c < 1$ and $C > 0$ such that unless ETH is false, any algorithm for (c,s)-Gap MAXLIN$_q$ with n variables and $m \le Cn$ equations must take $2^{\delta n}$ time for some $\delta > 0$.*

3 Reduction from (c,s)-Gap MAXLIN to γ-Gap k-MLD$_q$

In this section, we present reductions from (c,s)-Gap MAXLIN to γ-Gap k-MLD$_q$. We first give a simple reduction that yields $N^{\Omega(k/\log k)}$ lower bound. Then we provide a reduction for the tight lower bound. The reduction contains two steps.

Step 1: Reduction from (c,s)-Gap MAXLIN to γ-Gap MLD$_q$.

▶ **Lemma 15.** *There exists a polynomial-time reduction that takes as input a (c,s)-Gap MAXLIN instance $(\boldsymbol{A}, \boldsymbol{b})$ where $\boldsymbol{A} \in \mathbb{F}_q^{m\times n}$, $\boldsymbol{b} \in \mathbb{F}_q^m$ and $0 < s < c < 1$. The reduction outputs a γ-Gap MLD$_q$ instance $(\boldsymbol{H}, \boldsymbol{u}, \ell)$ where $\boldsymbol{H} \in \mathbb{F}_q^{d\times m}$ $(d \le m)$, $\boldsymbol{u} \in \mathbb{F}_q^d$, $\ell = (1-c)m$ and $\gamma = \frac{1-s}{1-c}$. This reduction satisfies the following properties:*

- ***Completeness:*** *If there exist $\boldsymbol{x} \in \mathbb{F}_q^n$ and $\boldsymbol{e} \in \mathbb{F}_q^m$ such that $\|\boldsymbol{e}\|_0 \le (1-c)m$ and $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{b}$, then there exists $\boldsymbol{x}' \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}'\|_0 \le \ell$ such that $\boldsymbol{H}\boldsymbol{x}' = \boldsymbol{u}$.*
- ***Soundness:*** *If for every $\boldsymbol{x} \in \mathbb{F}_q^n$ and $\boldsymbol{e} \in \mathbb{F}_q^m$ satisfying $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{b}$, we have $\|\boldsymbol{e}\|_0 > (1-s)m$, then there is no $\boldsymbol{x}' \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}'\|_0 \le \gamma\ell$ such that $\boldsymbol{H}\boldsymbol{x}' = \boldsymbol{u}$.*

Proof. Let $\boldsymbol{H}$ denote any parity-check matrix for the code generated by $\boldsymbol{A}$, thus $\boldsymbol{H} \in \mathbb{F}_q^{d\times m}$ and $\boldsymbol{H}\boldsymbol{A} = \boldsymbol{0}$ where $d := m - \operatorname{rank}(\boldsymbol{A})$. Given $\boldsymbol{A}$, the matrix $\boldsymbol{H}$ can be computed in polynomial time via Gram–Schmidt orthogonalization applied to an appropriate basis of $\mathbb{F}_q^m$. We now convert the (c,s)-Gap MAXLIN instance $(\boldsymbol{A}, \boldsymbol{b})$ to its dual form by multiplying both sides of the equation $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{b}$ by $\boldsymbol{H}$ on the left. Setting $\boldsymbol{u} := \boldsymbol{H}\boldsymbol{b}$, $\gamma := \frac{1-s}{1-c}$ and $\ell := (1-c)m$, we obtain the corresponding dual instance $(\boldsymbol{H}, \boldsymbol{u})$ satisfying:

- **Completeness.** If there exists $\boldsymbol{x} \in \mathbb{F}_q^n$ and $\boldsymbol{e} \in \mathbb{F}_q^m$ such that $\|\boldsymbol{e}\|_0 \le \ell = (1-c)m$ and $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{b}$, then $\boldsymbol{H}\boldsymbol{A}\boldsymbol{x} + \boldsymbol{H}\boldsymbol{e} = \boldsymbol{H}\boldsymbol{b}$, hence $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{u}$.
- **Soundness.** Assume for the sake of contradiction that there exists $\boldsymbol{e} \in \mathbb{F}_q^m$ with $\|\boldsymbol{e}\|_0 \le \gamma\ell = (1-s)m$ such that $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{u}$, then we have $\boldsymbol{H}(\boldsymbol{e} - \boldsymbol{b}) = \boldsymbol{0}$, hence $\boldsymbol{e} - \boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}$ for some $\boldsymbol{x} \in \mathbb{F}_q^n$.

Hence we get a polynomial-time reduction from (c,s)-Gap MAXLIN to γ-Gap MLD$_q$. ◀

Step 2: Reduction from γ-Gap MLD$_q$ to γ-Gap k-MLD$_q$. Next we will show a reduction from a non-parameterized MLD instance to a parameterized one. Consider a γ-Gap MLD$_q$ instance $(\boldsymbol{M}, \boldsymbol{u}, \ell)$. To obtain the parametrized instance, we need to "scale down" the solution size from ℓ to k by grouping the column vectors of $\boldsymbol{M}$ into appropriately sized blocks, and construct a new matrix in which each column vector behaves like an aggregated vector. This grouping ensures that, in the YES case, selecting k such aggregated vectors corresponds to selecting ℓ original vectors. For the reduction to be sound in the NO case, we additionally require that the size of each group be bounded by ℓ/k, so that any choice of at most γk aggregated vectors corresponds to at most $\gamma\ell$ original column vectors in $\boldsymbol{M}$.

▶ **Lemma 16.** *For every $\varepsilon > 0$, there exists a reduction that takes as input an integer $k \in \mathbb{N}$, a γ-Gap MLD_q instance $(\boldsymbol{M}, \boldsymbol{u}, \ell)$ where $\boldsymbol{M} \in \mathbb{F}_q^{d\times m}$, $\boldsymbol{u} \in \mathbb{F}_q^d$, and $\frac{k}{\varepsilon} < \ell < \frac{m}{\gamma}$. The reduction outputs a new matrix $\boldsymbol{M}_k \in \mathbb{F}_q^{d\times m'}$ in $(m')^{O(1)}$ time, which satisfies the following properties:*

- ***Size:*** *$m' = 2^{O(\frac{m\log k}{k})}$.*
- ***Completeness:*** *If there exists $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \le \ell$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$, then there exists $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \le k$ such that $\boldsymbol{M}_k\boldsymbol{y} = \boldsymbol{u}$.*
- ***Soundness:*** *If there is no $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \le \gamma\ell$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$, then for $\gamma' = \gamma - \varepsilon$, there is no $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \le \gamma' k$ such that $\boldsymbol{M}_k\boldsymbol{y} = \boldsymbol{u}$.*

Proof. Let $r := \lceil \frac{\ell}{k} \rceil$, $A := \{\boldsymbol{\alpha} \in \mathbb{F}_q^m : \|\boldsymbol{\alpha}\|_0 \le r\}$, and $m' := |A|$. We construct the new matrix $\boldsymbol{M}_k \in \mathbb{F}_q^{d\times m'}$ as follows. We establish a bijection between the column indices of $\boldsymbol{M}_k$ and A, and use $\boldsymbol{\alpha} \in A$ to denote a column index of $\boldsymbol{M}_k$. For each $\boldsymbol{\alpha} \in A$, we let the $\boldsymbol{\alpha}$-th column of $\boldsymbol{M}_k$ be

$$\boldsymbol{M}_k[\boldsymbol{\alpha}] = \boldsymbol{M}\boldsymbol{\alpha} \in \mathbb{F}_q^d.$$

Each column vector in $\boldsymbol{M}_k$ is therefore a linear combination of at most r column vectors from $\boldsymbol{M}$. Using the fact that $r = \lceil \frac{\ell}{k} \rceil < \frac{m}{k}$, the size of $\boldsymbol{M}_k$ satisfies

$$m' = \sum_{i=1}^{r} (q-1)^i \binom{m}{i} < rq^{\frac{m}{k}} \binom{m}{\frac{m}{k}} < rq^{\frac{m}{k}} (\mathrm{e}k)^{\frac{m}{k}} = 2^{O(\frac{m\log k}{k})}.$$

Completeness. Assume that there exists $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \le \ell$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$. We construct k vectors $\boldsymbol{x}_1, \dots, \boldsymbol{x}_k$ such that $\|\boldsymbol{x}_i\|_0 \le r$ and $\sum_{i=1}^k \boldsymbol{x}_i = \boldsymbol{x}$. Such vectors exist and can be constructed in polynomial time. Let $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ be such that $\boldsymbol{y}[\boldsymbol{x}_i] = 1$ for every $i \in [k]$, and $\boldsymbol{y}[\boldsymbol{\alpha}] = 0$ for any other $\boldsymbol{\alpha} \in A$. By construction, we have

$$\boldsymbol{M}_k\boldsymbol{y} = \sum_{i=1}^{k} \boldsymbol{M}_k[\boldsymbol{x}_i] = \sum_{i=1}^{k} \boldsymbol{M}\boldsymbol{x}_i = \boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}.$$

Soundness. Assume for the sake of contradiction that there exists $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \le \gamma' k$ such that $\boldsymbol{M}_k\boldsymbol{y} = \boldsymbol{u}$. Let $X := \mathrm{supp}(\boldsymbol{y})$, then $|X| = \|\boldsymbol{y}\|_0 \le \gamma' k$, and we have

$$\boldsymbol{u} = \boldsymbol{M}_k\boldsymbol{y} = \sum_{\boldsymbol{\alpha}\in X} \boldsymbol{M}_k[\boldsymbol{\alpha}] \cdot \boldsymbol{y}[\boldsymbol{\alpha}] = \sum_{\boldsymbol{\alpha}\in X} \boldsymbol{M}\boldsymbol{\alpha} \cdot \boldsymbol{y}[\boldsymbol{\alpha}].$$

Let $\boldsymbol{x} := \sum_{\boldsymbol{\alpha}\in X} \boldsymbol{\alpha} \cdot \boldsymbol{y}[\boldsymbol{\alpha}]$, then we have $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$. Moreover, since $\|\boldsymbol{\alpha}\|_0 \le r$ for every $\boldsymbol{\alpha} \in A$, we have

$$\|\boldsymbol{x}\|_0 \le \sum_{\boldsymbol{\alpha}\in X} \|\boldsymbol{\alpha}\|_0 \le r|X| = \gamma' k \left\lceil \frac{\ell}{k} \right\rceil \le \gamma'(\ell + k).$$

Thus when $\ell > \frac{k}{\varepsilon}$, we have $\|\boldsymbol{x}\|_0 < (\gamma' + \varepsilon)\ell = \gamma\ell$, proving the soundness. ◀

By Lemma 15 and Lemma 16, we obtain a reduction from (c, s)-Gap MAXLIN with m equations to γ-Gap k-MLD_q with size $N \le 2^{O(\frac{m\log k}{k})}$ for some constant $\gamma > 1$, hence ruling out $N^{o(\frac{k}{\log k})}$-time algorithms for solving γ-Gap k-MLD_q.

Note that in our construction, when we group the vectors to form aggregated vectors, we chose all subsets of size at most $\lceil \ell/k \rceil$. This brute approach already gave us an almost tight lower bound for γ-Gap k-MLD_q. To get a tight lower bound we would want to get a reduction where the number of aggregated vectors is upper bounded by $2^{O(m/k)}$. In the following section, we will make use of a combinatorial gadget, called the "cover family". This gadget will essentially enable us to do the grouping more cleverly and getting a tight lower bound.

3.1 Grouping using Cover Families

To get a tighter bound in the above reduction, we consider reducing the number of columns in $\boldsymbol{M}_k$. This requires us to find a smaller subset $A \subseteq \mathbb{F}_q^m$ that maintains the completeness and soundness requirements as follows:

- (Completeness) For every $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \leq \ell$, there exist $\boldsymbol{x}_1, \dots, \boldsymbol{x}_k \in A$ such that $\sum_{i=1}^k \boldsymbol{x}_i = \boldsymbol{x}$.
- (Soundness) For every $\boldsymbol{x} \in \mathbb{F}_q^m$, if there exist $\boldsymbol{x}_1, \dots, \boldsymbol{x}_{\gamma' k} \in A$ such that $\sum_{i=1}^{\gamma' k} \boldsymbol{x}_i = \boldsymbol{x}$, then $\|\boldsymbol{x}\|_0 \leq \gamma\ell$.

To construct such set A, we will make use of a combinatorial gadget which we call "cover family". It produces a gap between the number of sets required to cover a small set and a large set. Let A be the set of all vectors whose support is in the cover family, then A meets all requirements above. The gadget is formally defined as follows.

▶ **Definition 17** ($(U, k, \alpha, \varepsilon)$-cover family)**.** *A $(U, k, \alpha, \varepsilon)$-cover family is a collection $\mathcal{S} \subseteq 2^U$ of subsets of U, such that*

- *(C1) For every $S \in \mathcal{S}$, $|S| \leq \frac{(1+\varepsilon)\alpha|U|}{k}$.*
- *(C2) For every $\widetilde{S} \subseteq U$ with $|\widetilde{S}| \leq \alpha|U|$, there exists $\{T_1, \dots, T_k\} \subseteq \mathcal{S}$ such that $\bigcup_{i \in [k]} T_i = \widetilde{S}$. Moreover, $T_i \cap T_j = \emptyset$ for every $i \neq j \in [k]$.*

Note that property (C1) immediately implies that for any $\widetilde{S} \subseteq U$ with $|\widetilde{S}| \geq \beta|U|$, it requires at least $\frac{\beta}{(1+\varepsilon)\alpha}k$ sets in $\mathcal{S}$ to cover $\widetilde{S}$, hence producing our desired gap. We now present the improved reduction as follows.

▶ **Lemma 18.** *There exists a reduction that takes as input an integer $k \in \mathbb{N}$, an $([m], k, \alpha, \varepsilon)$-cover family $\mathcal{S} \subseteq 2^{[m]}$, and a γ-Gap MLD_q instance $(\boldsymbol{M}, \boldsymbol{u}, \alpha m)$ where $\boldsymbol{M} \in \mathbb{F}_q^{d \times m}$, $\boldsymbol{u} \in \mathbb{F}_q^d$. The reduction outputs a new matrix $\boldsymbol{M}_k \in \mathbb{F}_q^{d \times m'}$ in $(m')^{O(1)}$ time, which satisfies the following properties:*

- ***Size:*** *$m' = |\mathcal{S}| \cdot q^{\frac{(1+\varepsilon)\alpha m}{k}}$.*
- ***Completeness:*** *If there exists $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \leq \alpha m$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$, then there exists $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \leq k$ such that $\boldsymbol{M}_k\boldsymbol{y} = \boldsymbol{u}$.*
- ***Soundness:*** *If there is no $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \leq \gamma\alpha m$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$, then there is no $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \leq \frac{\gamma}{1+\varepsilon}k$ such that $\boldsymbol{M}_k\boldsymbol{y} = \boldsymbol{u}$.*

Proof. Let $A := \{\boldsymbol{\alpha} \in \mathbb{F}_q^m : \mathrm{supp}(\boldsymbol{\alpha}) \in \mathcal{S}\}$ and $m' := |A|$. We construct the new matrix $\boldsymbol{M}_k \in \mathbb{F}_q^{d \times m'}$ as follows. We establish a bijection between the column indices of $\boldsymbol{M}_k$ and A, and use $\boldsymbol{\alpha} \in A$ to denote a column index of $\boldsymbol{M}_k$. For each $\boldsymbol{\alpha} \in A$, we let the $\boldsymbol{\alpha}$-th column of $\boldsymbol{M}_k$ be

$$\boldsymbol{M}_k[\boldsymbol{\alpha}] = \boldsymbol{M}\boldsymbol{\alpha} \in \mathbb{F}_q^d.$$

The property (C1) of $\mathcal{S}$ guarantees that $|S| \leq \frac{(1+\varepsilon)\alpha m}{k}$, hence the size of $\boldsymbol{M}_k$ satisfies

$$m' = \sum_{S \in \mathcal{S}} (q-1)^{|S|} < |\mathcal{S}| \cdot q^{\frac{(1+\varepsilon)\alpha m}{k}}.$$

Completeness. Assume that there exists $\boldsymbol{x} \in \mathbb{F}_q^m$ with $\|\boldsymbol{x}\|_0 \le \alpha m$ such that $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$. Since $|\operatorname{supp}(\boldsymbol{x})| \le \alpha m$, by the property (C2) of $\mathcal{S}$, there exists $T_1, \dots, T_k \in \mathcal{S}$ as a partition of $\operatorname{supp}(\boldsymbol{x})$.

For each $i \in [k]$, let $\boldsymbol{x}_i \in \mathbb{F}_q^m$ be the projection of $\boldsymbol{x}$ onto the coordinates in T_i, i.e., for every $j \in [m]$, we let

$$\boldsymbol{x}_i[j] = \begin{cases} \boldsymbol{x}[j], & j \in T_i, \\ 0, & j \notin T_i. \end{cases}$$

Then $\sum_{i=1}^k \boldsymbol{x}_i = \boldsymbol{x}$, and $\boldsymbol{x}_i \in A$ because $\operatorname{supp}(\boldsymbol{x}) = T_i \in \mathcal{S}$. Let $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ be such that $\boldsymbol{y}[\boldsymbol{x}_i] = 1$ for every $i \in [k]$, and $\boldsymbol{y}[\boldsymbol{\alpha}] = 0$ for any other $\boldsymbol{\alpha} \in A$. By construction, we have

$$\boldsymbol{M}_k \boldsymbol{y} = \sum_{i=1}^k \boldsymbol{M}_k[\boldsymbol{x}_i] = \sum_{i=1}^k \boldsymbol{M}\boldsymbol{x}_i = \boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}.$$

Soundness. Assume for the sake of contradiction that there exists $\boldsymbol{y} \in \mathbb{F}_q^{m'}$ with $\|\boldsymbol{y}\|_0 \le \frac{\gamma}{1+\varepsilon} k$ such that $\boldsymbol{M}_k \boldsymbol{y} = \boldsymbol{u}$. Let $X := \operatorname{supp}(\boldsymbol{y})$, then we have

$$\boldsymbol{u} = \boldsymbol{M}_k \boldsymbol{y} = \sum_{\boldsymbol{\alpha} \in X} \boldsymbol{M}_k[\boldsymbol{\alpha}] \cdot \boldsymbol{y}[\boldsymbol{\alpha}] = \sum_{\boldsymbol{\alpha} \in X} \boldsymbol{M}\boldsymbol{\alpha} \cdot \boldsymbol{y}[\boldsymbol{\alpha}].$$

Let $\boldsymbol{x} := \sum_{\boldsymbol{\alpha} \in X} \boldsymbol{y}[\boldsymbol{\alpha}] \cdot \boldsymbol{\alpha}$, then $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$. By the property (C1) of $\mathcal{S}$, we have $|\operatorname{supp}(\boldsymbol{\alpha})| \le \frac{(1+\varepsilon)\alpha}{k} m$ for every $\boldsymbol{\alpha} \in A$. Since $|X| = \|\boldsymbol{y}\|_0 \le \frac{\gamma}{1+\varepsilon} k$, we have

$$|\operatorname{supp}(\boldsymbol{x})| \le \Big| \bigcup_{\boldsymbol{\alpha} \in X} \operatorname{supp}(\boldsymbol{\alpha}) \Big| \le |X| \cdot \frac{(1+\varepsilon)\alpha}{k} m \le \gamma \alpha m.$$

Hence $\boldsymbol{x}$ satisfies $\boldsymbol{M}\boldsymbol{x} = \boldsymbol{u}$ and $\|\boldsymbol{x}\|_0 \le \gamma \alpha m$, proving the soundness. ◀

Combining Lemma 15 and Lemma 18, we obtain a reduction which on input an $([m], k, 1-c, \varepsilon)$-cover family $\mathcal{S}$, it reduces a (c, s)-Gap MAXLIN instance with m equations to γ-Gap k-MLD$_q$ instance with size $N = |\mathcal{S}| \cdot 2^{O(m/k)}$ for $\gamma = \frac{1-s}{(1-c)(1+\varepsilon)}$. If $|\mathcal{S}| = 2^{O(m/k)}$, by Corollary 14 we get that assuming ETH, there exists a constant $\gamma > 1$ such that no algorithm running in time $N^{o(k)}$ can solve γ-Gap k-MLD$_q$.

To achieve the tight lower bound, we will construct a cover family over universe $[m]$ with size $|\mathcal{S}| = 2^{O(m/k)}$ using the following lemma, whose proof is postponed to Section 4.

▶ **Lemma 19.** *For every $m, k \in \mathbb{N}$ and $0 < \alpha, \varepsilon < 1$ such that $m \ge k^{4k^2/(\varepsilon^2\alpha)}$, there exists an $([m], k, \alpha, \varepsilon)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = \lceil \log_k m \rceil \cdot k \cdot 2^{2m/k}$ that can be constructed deterministically in $|\mathcal{S}|^{O(1)}$ time.*

It follows that

$$N = 2^{O_q(m/k)}.$$

This improves upon the naive grouping bound of $2^{O(m \log k/k)}$ by removing the $\log k$ factor in the exponent. Consequently, under ETH, this reduction implies that γ-Gap k-MLD$_q$ cannot be solved in time $N^{o(k)}$, confirming the tightness of the lower bound for the parameterized problem.

Note that Lemma 19 requires $m \ge k^{\Omega_{\alpha,\varepsilon}(k^2)}$. To obtain the tight lower bound for larger k, we give a randomized construction of cover families.

▶ **Lemma 20.** *For every $m, k \in \mathbb{N}$, $0 < \alpha, \varepsilon < 1$ such that $m \ge \frac{6k \ln 2k}{\varepsilon^2 \alpha}$, there exists an $([m], k, \alpha, \varepsilon)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = \frac{12k^2}{\varepsilon^2\alpha} \cdot 2^{(1+\varepsilon)m/k}$ that can be constructed in $|\mathcal{S}|^{O(1)}$ time with probability $1 - o(1)$.*

3.2 Putting Everything Together

We now combine everything together to prove the main result and derive other conclusions.

▶ **Theorem 21** (Theorem 1 Restated). *Assuming* ETH*, there exist constants $\gamma > 1$ and $\delta > 0$, such that for any algorithm that takes as input a matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times N}$, a vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and a parameter $k \in \mathbb{N}$, it must take $N^{\delta k}$ time to distinguish the following two cases:*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le k$ such that $\boldsymbol{Hx} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^N$ with $\|\boldsymbol{x}\|_0 \le \gamma k$, we have $\boldsymbol{Hx} \neq \boldsymbol{u}$.*

Proof. We start with a (c, s)-Gap MAXLIN_q instance with n variables and $m \le Cn$ equations for some constant $0 < s < c < 1$ and $C > 0$. By Corollary 14, there exists some constant $\delta > 0$ such that deciding (c, s)-Gap MAXLIN_q must take $2^{\delta n}$ time assuming ETH.

We now invoke Lemma 19 with $\alpha \leftarrow 1 - c$, and a sufficiently small $\varepsilon > 0$ such that $\gamma := \frac{1-s}{(1+\varepsilon)(1-c)} > 1$. This gives us a $\big([m], k, 1 - c, \varepsilon\big)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = k \cdot 2^{O(m/k)}$, constructible in $|\mathcal{S}|^{O(1)}$ time. Using Lemma 15 and Lemma 18, we reduce this (c, s)-Gap MAXLIN_q instance to a γ-Gap k-MLD_q instance where the number of vectors is $N = |\mathcal{S}| \cdot q^{O(m/k)} = k \cdot 2^{O_q(n/k)}$.

Assume that γ-Gap k-MLD_q can be solved in $N^{\delta k}$ time for every $\delta > 0$, then for any $\delta' > 0$, there exists δ such that (c, s)-Gap MAXLIN_q can be solved in time $\big(k \cdot 2^{O_q(n/k)}\big)^{\delta k} < 2^{\delta' n}$. Therefore, there exist $\gamma > 1$ and $\delta > 0$ such that assuming ETH, any algorithm that decides γ-Gap k-MLD_q must take time $N^{\delta k}$. ◀

In fact, Theorem 21 also holds for larger k such that $k = O((\log\log N)^{0.49})$, since we can construct the $\big([m], k, 1 - c, \varepsilon\big)$-cover family as long as $k^{\Omega(k^2)} \le m$. Following this observation, if we use the randomized construction of cover family instead of the deterministic one, we obtain the lower bound for any $k \le N^{\varepsilon}$ for some constant $\varepsilon > 0$ under randomized ETH.

▶ **Theorem 22** (Theorem 2 Restated). *Assuming randomized* ETH*, there exist constants $\gamma > 1$ and $\delta, \varepsilon' > 0$, such that for any algorithm that takes as input a matrix $\boldsymbol{H} \in \mathbb{F}_q^{d\times N}$, a vector $\boldsymbol{u} \in \mathbb{F}_q^d$ and an integer $k \in \mathbb{N}$ such that $2 \le k \le N^{\varepsilon'}$, it must take $N^{\delta k}$ time to distinguish the following two cases:*

- *YES Case: There exists $\boldsymbol{x} \in \mathbb{F}_q^d$ with $\|\boldsymbol{x}\|_0 \le k$ such that $\boldsymbol{Hx} = \boldsymbol{u}$.*
- *NO Case: For every $\boldsymbol{x} \in \mathbb{F}_q^d$ with $\|\boldsymbol{x}\|_0 \le \gamma k$, we have $\boldsymbol{Hx} \neq \boldsymbol{u}$.*

Proof. We still start with a $(1 - \alpha, s)$-Gap MAXLIN_q instance with n variables and $m \le Cn$ equations. Pick a small enough constant $\varepsilon > 0$ such that $\gamma := \frac{1-s}{(1+\varepsilon)\alpha} > 1$. Let

$$N := \frac{12k^2}{\varepsilon^2\alpha} \cdot 2^{(1+\varepsilon)(1+\alpha\log_2 q)\frac{m}{k}} \quad \text{and} \quad \beta := \frac{\varepsilon^2\alpha}{6(1+\varepsilon)(1+\alpha\log_2 q)}.$$

When

$$k \le \Big(\frac{\varepsilon^2\alpha}{12}N\Big)^{\frac{\beta}{2(1+\beta)}},$$

which implies that

$$k\ln 2k \le k\log_2(2k) \le \frac{\varepsilon^2\alpha}{6}m,$$

by Lemma 20, we can construct a $\big([m], k, \alpha, \varepsilon\big)$-cover family $\mathcal{S}$ with size $|\mathcal{S}| = \frac{12k^2}{\varepsilon^2\alpha} \cdot 2^{(1+\varepsilon)m/k}$.

Using the $\big([m], k, \alpha, \varepsilon\big)$-cover family, we can reduce the MAXLIN instance to a γ-Gap k-MLD_q instance with size $N = |\mathcal{S}| \cdot q^{(1+\varepsilon)\alpha m/k} = k^2 \cdot 2^{O_q(n/k)}$. Assume that

γ-Gap k-MLD_q can be solved in $N^{\delta k}$ time for every $\delta > 0$, then for any $\delta' > 0$, there exists δ such that (c, s)-Gap MAXLIN_q can be solved in time

$$\left(k^2 \cdot 2^{O_q(n/k)}\right)^{\delta k} = 2^{\delta \cdot O_q(n + k \log k)} < 2^{\delta' n}.$$

This proves the $N^{\delta k}$-time lower bound for $k = O(N^{\varepsilon'})$ where $\varepsilon' := \frac{\beta}{2(1+\beta)}$. ◀

We state the following corollary which can be obtained by a standard reduction from γ-Gap k-MLD_q to γ-Gap k-NCP_q as in [15].

▶ **Corollary 23.** *Assuming* ETH*, there exists a constant $\gamma > 1$, such that no algorithm can decide γ-Gap k-NCP_q instance $(\boldsymbol{A}, \boldsymbol{t})$ where $\boldsymbol{A} \in F_q^{N \times M}, \boldsymbol{t} \in \mathbb{F}_q^N$ in $N^{o(k)}$ time.*

Using the reduction from γ-Gap NCP_2 to γ-Gap CVP_p as in [1], we can obtain optimal hardness for γ-Gap k-CVP_p for all $p > 1$.

▶ **Corollary 24.** *Assuming* ETH*, for every $p > 1$, there exists a constant $\gamma > 1$, such that no algorithm can decide γ-Gap k-CVP_p instance $(\boldsymbol{A}, \boldsymbol{t})$ where $\boldsymbol{A} \in \mathbb{Z}^{N \times M}, \boldsymbol{t} \in \mathbb{Z}^N$ in $N^{o(k)}$ time.*

4 Construction of Cover Families

In this section, we present the construction of cover families. Our approach relies on an intermediate combinatorial object which we call a *balanced partition family.* A balanced partition family is a family of partitions over a universe of size n, such that every subset of size αn is partitioned into almost equal size by some partition. We first formalize balanced partition families as follows, and then establish a reduction showing that any efficient balanced partition family yields a corresponding cover family.

▶ **Definition 25** ($(U, k, \alpha, \varepsilon, c)$-balanced partition family)**.** *A $(U, k, \alpha, \varepsilon, c)$-balanced partition family is a family $\mathcal{F} \subseteq [k]^U$ of functions from U to $[k]$, such that*

- *(P1) For every $f \in \mathcal{F}$ and $j \in [k]$, $|f^{-1}(j)| \le \frac{c|U|}{k}$.*
- *(P2) For every $S \subseteq U$ with $|S| = \alpha|U|$, there exists $f \in \mathcal{F}$ such that $|S \cap f^{-1}(j)| \le \frac{(1+\varepsilon)|S|}{k}$ for every $j \in [k]$.*

In the following theorem we present the construction of a *cover family*, assuming the existence of a *balanced partition family.*

▶ **Theorem 26.** *On input a $(U, k, \alpha, \varepsilon, c)$-balanced partition family $\mathcal{F}$, one can construct a $(U, k, \alpha, \varepsilon)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = |\mathcal{F}| \cdot k \cdot 2^{c|U|/k}$ in $|\mathcal{S}|^{O(1)}$ time.*

Proof. Let $m := |U|$. We construct the cover family $\mathcal{S}$ as follows:

- For every function $f \in \mathcal{F}$ and every index $j \in [k]$, let $\mathcal{S}_{f,j}$ be the collection of all subsets of the j-th bucket $f^{-1}(j)$ that have size at most $\frac{(1+\varepsilon)\alpha m}{k}$. That is,

$$\mathcal{S}_{f,j} := \left\{ T \subseteq f^{-1}(j) : |T| \le \frac{(1+\varepsilon)\alpha m}{k} \right\}.$$

- Define $\mathcal{S} := \bigcup_{f \in \mathcal{F}} \bigcup_{j \in [k]} \mathcal{S}_{f,j}$.

Size and Running Time. By property (P1) of the balanced partition family, for every $f \in \mathcal{F}$ and $j \in [k]$, the bucket size is bounded by $|f^{-1}(j)| \leq \frac{cm}{k}$. The number of subsets in $\mathcal{S}_{f,j}$ is bounded by the total number of subsets of the bucket, which is $2^{|f^{-1}(j)|}$. Therefore,

$$|\mathcal{S}| \leq \sum_{f \in \mathcal{F}} \sum_{j \in [k]} 2^{|f^{-1}(j)|} \leq |\mathcal{F}| \cdot k \cdot 2^{cm/k}.$$

The construction can be performed in time polynomial in $|\mathcal{S}|$.

Property (C1). By construction, the maximum size of any set $S \in \mathcal{S}$ is bounded by $\frac{(1+\varepsilon)\alpha m}{k}$.

Property (C2). Let $\widetilde{S} \subseteq U$ be any set with $|\widetilde{S}| \leq \alpha m$. We need to show that $\widetilde{S}$ can be exactly covered by k disjoint sets from $\mathcal{S}$. Pick an arbitrary set S' such that $|S'| = \alpha m$ and $\widetilde{S} \subseteq S'$. By property (P2) of the balanced partition family, there exists a function $f \in \mathcal{F}$ such that for all $j \in [k]$,

$$|\widetilde{S} \cap f^{-1}(j)| \leq |S' \cap f^{-1}(j)| \leq \frac{(1+\varepsilon)|S'|}{k} = \frac{(1+\varepsilon)\alpha m}{k}.$$

Define $T_j := \widetilde{S} \cap f^{-1}(j)$ for each $j \in [k]$.

- Since $T_j \subseteq f^{-1}(j)$ and $|T_j| \leq \frac{(1+\varepsilon)\alpha m}{k}$, we have $T_j \in \mathcal{S}_{f,j} \subseteq \mathcal{S}$.
- Since $\{f^{-1}(j)\}_{j \in [k]}$ forms a partition of U, the sets $T_1, \ldots, T_k$ are pairwise disjoint.
- $T_1, \ldots, T_k$ can cover $\widetilde{S}$, because $\bigcup_{j \in [k]} T_j = \widetilde{S} \cap \bigcup_{j \in [k]} f^{-1}(j) = \widetilde{S} \cap U = \widetilde{S}$.

Thus, $\{T_1, \ldots, T_k\} \subseteq \mathcal{S}$ is a exact cover for $\widetilde{S}$ and satisfies (C2). ◀

4.1 Randomized Construction

We will prove the existence of a balanced partition family with $O(k)$ size, by showing that randomly sampling $O(k)$ partitions and then rejecting all partitions that do not satisfy (C1) produces a balanced partition family with high probability.

▶ **Theorem 27.** *For every $m, k \in \mathbb{N}$ and $0 < \alpha, \varepsilon < 1$ such that $m \geq \frac{6k \ln 2k}{\varepsilon^2 \alpha}$, there exists an $([m], k, \alpha, \varepsilon, 1+\varepsilon)$-balanced partition family $\mathcal{F}$ with $|\mathcal{F}| = \frac{12k}{\varepsilon^2 \alpha}$.*

Proof. Let $t := \left\lceil \frac{12k}{\varepsilon^2 \alpha} \right\rceil$. We sample t functions $f_1, \ldots, f_t : U \to [k]$ independently and uniformly at random. We define $I \subseteq [k]$ as the indices of functions that satisfy (P1), formally, let

$$I := \left\{ i \in [t] : \forall j \in [k],\ |f_i^{-1}(j)| \leq \frac{(1+\varepsilon)m}{k} \right\}.$$

For any subset $S \subseteq U$, let $B(S)$ be the "bad event" that no function f_i with $i \in I$ evenly splits S. That is, for every $i \in I$, there exists some bucket $j \in [k]$ such that $|S \cap f_i^{-1}(j)| > \frac{(1+\varepsilon)|S|}{k}$. We will show that with high probability, $B(S)$ does not occur for every subset $S \subseteq [m]$ of size αm.

Consider a fixed index $i \in [t]$ and a fixed subset $S \subseteq U$ with $|S| = \alpha m$. We define two bad events for f_i and analyze the probability of these events:

- Let $E_1(i)$ be the event that $i \notin I$. Since f_i is a uniform random function, $f(x)$ is drawn uniformly and independently from $[k]$ for each $x \in [m]$. Hence for a fixed bucket $j \in [k]$, $x \in f_i^{-1}(j)$ with probability $\frac{1}{k}$ for every $x \in [m]$, and thus $|f_i^{-1}(j)|$ is a sum of independent Bernoulli trials with mean m/k. By the Chernoff bound,

$$\mathbf{Pr}\left[|f_i^{-1}(j)| > \frac{(1+\varepsilon)m}{k} \right] < \exp\left(-\frac{\varepsilon^2 m}{3k} \right).$$

Taking a union bound over all k buckets,

$$\mathbf{Pr}[E_1(i)] < k \exp\left(-\frac{\varepsilon^2 m}{3k}\right).$$

- Let $E_2(i,S)$ be the event that f_i fails the property (P2) for the specific set S. This means there exists some bucket $j \in [k]$ such that $|S \cap f_i^{-1}(j)| > \frac{(1+\varepsilon)\alpha m}{k}$. For a fixed $j \in [k]$, the random variable $|S \cap f_i^{-1}(j)|$ follows a binomial distribution $\mathrm{Bin}(\alpha m, 1/k)$. Similarly, by the Chernoff bound and a union bound over k buckets, we have

$$\mathbf{Pr}[E_2(i,S)] < k \exp\left(-\frac{\varepsilon^2 \alpha m}{3k}\right).$$

Let $E(i,S) := E_1(i) \vee E_2(i,S)$ be the combined bad event where f_i either violates (P1) or fails to balance S. Since $\alpha < 1$, we have

$$\mathbf{Pr}[E(i,S)] \le \mathbf{Pr}[E_1(i)] + \mathbf{Pr}[E_2(i,S)] < 2k \exp\left(-\frac{\varepsilon^2 \alpha m}{3k}\right).$$

Since the functions are sampled independently, we have

$$\mathbf{Pr}[B(S)] = \prod_{i=1}^{t} \mathbf{Pr}[E(i,S)] < \left(2k \exp\left(-\frac{\varepsilon^2 \alpha m}{3k}\right)\right)^t.$$

Take a union bound over all possible subsets S of size αm. Since the total number of such subsets is $\binom{m}{\alpha m} < \mathrm{e}^m$, the probability that there exists a bad set S is at most

$$\begin{aligned}\sum_{S \in \binom{[m]}{\alpha m}} \mathbf{Pr}[B(S)] &< \mathrm{e}^m \left(2k \exp\left(-\frac{\varepsilon^2 \alpha m}{3k}\right)\right)^t \\ &= \exp\left(m - t\left(\frac{\varepsilon^2 \alpha m}{3k} - \ln 2k\right)\right).\end{aligned}$$

When $m \ge \frac{6k \ln 2k}{\varepsilon^2 \alpha}$ and $t = \left\lceil \frac{12k}{\varepsilon^2 \alpha} \right\rceil$, the total probability is at most e^{-m}.

We let the final family $\mathcal{F}$ contain all functions that satisfy property (P1), i.e., let $\mathcal{F} := \{f_i : i \in I\}$. By definition, every function in $\mathcal{F}$ satisfies (P1). Furthermore, we have proved that with probability $1 - \mathrm{e}^{-m}$, for every subset $S \in \binom{U}{\alpha m}$, there exists $f \in \mathcal{F}$ that evenly splits S, hence $\mathcal{F}$ also satisfies (P2). Thus, $\mathcal{F}$ is a valid balanced partition family with high probability. ◀

Using Theorem 26, we obtain an $([m], k, \alpha, \varepsilon)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = \frac{12k}{\varepsilon^2 \alpha} \cdot k \cdot 2^{(1+\varepsilon)m/k}$ in $|\mathcal{S}|^{O(1)}$ time, hence proving Lemma 20.

4.2 Derandomization

This section introduces the hypercube partition system as a derandomization of the balanced partition family. We define the hypercube partition system as follows.

▶ **Definition 28** (Hypercube Partition System). *A (k,d)-hypercube partition system is a pair $(U, \mathcal{F})$ where $U := [k]^{[d]}$ and $\mathcal{F} := \{f_1, \ldots, f_d\}$ is a collection of funtions from U to $[k]$. We view U as the set of all points in a d-dimensional hypercube where each coordinate takes a value in $[k]$. The i-th partition $f_i : U \to [k]$ is defined as for every $x \in U$*

$$f_i(x) = x(i).$$

In other words, for any $x \in U$, f_i groups x according to its i-th coordinate, into the $x(i)$-th set.

We show that a (k, d)-hypercube partition system $(U, \mathcal{F})$ is a balanced partition family for sufficiently large d. Informally, it suffices to show the following proposition:

($\star$) For any subset $S \subseteq U$, if on every coordinate $i \in [d]$, there exists some $j \in [k]$ such that at least $\frac{(1+\varepsilon)|S|}{k}$ elements $x \in S$ satisfy $x(i) = j$, then $|S| < \alpha|U|$.

By the contrapositive of ($\star$), if S is large enough, then at least one coordinate $i \in [d]$ must evenly split S. This indicates that $\mathcal{F}$ is a balanced partition family.

To show ($\star$), we consider a uniform distribution p over the subset S. Let p_i denote the marginal distribution of p on the i-th coordinate. The condition in ($\star$) guarantees that for each coordinate i, p_i concentrates on some value j, which implies that p_i has small entropy. Then we add up the entropy of p_i over all coordinates, which provides a upper bound on the entropy of p. Hence $|S|$ must be small.

We prove this formally in Lemma 29.

▶ **Lemma 29.** *Let $d, k \in \mathbb{N}$, $0 < \alpha, \varepsilon < 1$, and $U := [k]^d$. For every $i \in [d]$, let $f_i : U \to [k]$ be such that $f_i(x) = x_i$. If $d \geq \frac{4k}{\varepsilon^2 \alpha}$, then for every $S \subseteq U$ with $|S| = \alpha k^d$, there exists $i \in [d]$, such that $|S \cap f_i^{-1}(j)| \leq (1 + \varepsilon)\alpha k^{d-1}$ for every $j \in [k]$.*

Proof. Let X be a random variable uniformly distributed over the subset S. We have $H(X) = \ln |S| = d \ln k + \ln \alpha$. Let p_i be the distribution of X_i, i.e., for every $j \in [k]$,

$$p_i(j) = \frac{|\{x \in S : x_i = j\}|}{|S|} = \frac{|S \cap f_i^{-1}(j)|}{\alpha k^d}.$$

Now we will prove that there exists some $i \in [d]$, such that $p_i(j) \leq \frac{1+\varepsilon}{k}$ for every $j \in [k]$, which immediately implies that $|S \cap f_i^{-1}(j)| \leq (1 + \varepsilon)\alpha k^{d-1}$.

Assume for the sake of contradiction that for every $i \in [d]$, there exists some $j_i \in [k]$ such that $p_i(j_i) > \frac{1+\varepsilon}{k}$. Fix any $i \in [d]$ and let $\rho_i := p_i(j_i)$. Then, by Jensen's inequality,

$$\begin{aligned} H(X_i) &= p_i(j_i) \ln \frac{1}{p_i(j_i)} + \sum_{j \neq j_i} p_i(j) \ln \frac{1}{p_i(j)} \\ &\leq \rho_i \ln \frac{1}{\rho_i} + (1 - \rho_i) \ln \frac{k-1}{1-\rho_i}. \end{aligned}$$

Let $h(x) := x \ln \frac{1}{x} + (1 - x) \ln \frac{k-1}{1-x}$. We have $h'\left(\frac{1}{k}\right) = 0$ and $h''(x) = -\frac{1}{x(1-x)}$. By Taylor's theorem with the Lagrange's form of remainder, we expand $h(x)$ at $x = 1/k$, then there exists ξ with $\frac{1}{k} < \xi < \rho_i$ such that

$$\begin{aligned} h(\rho_i) &= h\Big(\frac{1}{k}\Big) + h'\Big(\frac{1}{k}\Big)\Big(\rho_i - \frac{1}{k}\Big) + \frac{h''(\xi)}{2}\Big(\rho_i - \frac{1}{k}\Big)^2 \\ &= \ln k - \frac{1}{2\xi(1-\xi)}\Big(\rho_i - \frac{1}{k}\Big)^2. \end{aligned}$$

Recall that $\rho_i > \frac{1+\varepsilon}{k}$. Since $\xi(1 - \xi) < \xi < \rho_i$, and by the fact that $f(x) := x + \frac{1}{x} - 2$ is increasing when $x > 1$, we have

$$\ln k - h(\rho_i) > \frac{1}{2\rho_i}\Big(\rho_i - \frac{1}{k}\Big)^2 = \frac{f(k\rho_i)}{2k} > \frac{f(1+\varepsilon)}{2k} > \frac{\varepsilon^2}{4k},$$

thus

$$H(X_i) \leq h(\rho_i) < \ln k - \frac{\varepsilon^2}{4k}.$$

By the subadditivity of entropy, when $d \geq \frac{4k}{\varepsilon^2 \alpha}$, we have

$$H(X) \leq \sum_{i \in [d]} H(X_i) \leq d \ln k - \frac{\varepsilon^2 d}{4k} \leq d \ln k - \frac{1}{\alpha} < d \ln k + \ln \alpha,$$

which contradicts $H(X) = \ln |S| = d \ln k + \ln \alpha$. ◀

Lemma 29 implies that a (k,d)-hypercube partition system is a $([k]^d, k, \alpha, \varepsilon, 1)$-balanced partition family of size d if d is large enough. A limitation of this construction is that we require the size of the universe to be an integer power of k. To address this, we define a subset $U \subseteq [k]^d$ of size exactly m by taking a "diagonal slice" of the hypercube. This slice preserves the balanced property of projections: fixing one coordinate to any value leaves the remaining coordinates to cycle through values uniformly, ensuring the slice is spread evenly across buckets.

▶ **Theorem 30.** *For every $m, k \in \mathbb{N}$ and $0 < \eta, \varepsilon < 1$ such that $m \geq k^{4k^2/(\varepsilon^2 \eta)}$, there exists an $([m], k, \eta, \varepsilon, 2)$-balanced partition family $\mathcal{F}$ of size $|\mathcal{F}| = \lceil \log_k m \rceil$ that can be constructed in $m^{O(1)}$ time.*

Proof. Let $d := \lceil \log_k m \rceil$ and $c := \lceil \frac{m}{k^{d-1}} \rceil$. We have $(c-1)k^{d-1} < m \leq ck^{d-1}$ and $2 \leq c \leq k$. Define a partition $U_0, \ldots, U_{k-1}$ of $[k]^d$ based on the modular sum of coordinates:

$$U_\ell := \left\{ x \in [k]^d : \left(\sum_{j=1}^{d} x_j \right) \bmod k = \ell \right\}.$$

Let $U' := U_0 \cup \cdots \cup U_{c-1}$. Then $|U'| = ck^{d-1}$ because $|U_\ell| = k^{d-1}$ for every ℓ. We set the universe U to be an arbitrary subset of U' of size $m \leq |U'|$. Define $\mathcal{F} = \{f_1, \ldots, f_d\}$ as the projections on each coordinate: $f_i(x) = x(i)$.

Property (P1). For any $i \in [d], j \in [k]$ and $0 \leq \ell < k$, the number of $x \in U_\ell$ such that $x_i = j$ is exactly k^{d-2}. To prove this, we take an arbitrary $i' \neq i$. If we fix x_t for every $t \in [d] \setminus \{i, i'\}$, then $x_{i'}$ must be $x_{i'} = \ell - \sum_{t \in [d] \setminus \{i'\}} x_t$. The number of ways to fix x_t is k^{d-2}, hence there are k^{d-2} many $x \in U_\ell$ with $x_i = j$, i.e.,

$$|U_\ell \cap f_i^{-1}(j)| = k^{d-2}.$$

Since $m > (c-1)k^{d-1}$, for any $i \in [d]$ and $j \in [k]$,

$$|U \cap f_i^{-1}(j)| \leq |U' \cap f_i^{-1}(j)| = ck^{d-2} \leq 2(c-1)k^{d-2} < \frac{2m}{k}.$$

Hence $\mathcal{F}$ satisfies property (P1).

Property (P2). Let $\alpha := \frac{\eta m}{k^d}$. Since $m > k^{d-1}$, we have $\eta < \alpha k$, hence

$$d \geq \log_k m \geq \frac{4k^2}{\varepsilon^2 \eta} > \frac{4k}{\varepsilon^2 \alpha}.$$

Applying Lemma 29, for any $S \subseteq [k]^d$ of size $\alpha k^d = \eta m$, there exists $f_i \in \mathcal{F}$ such that every bucket $S \cap f_i^{-1}(j)$ has size at most

$$(1+\varepsilon)\alpha k^{d-1} = \frac{(1+\varepsilon)\eta m}{k} = \frac{(1+\varepsilon)|S|}{k}.$$

Hence $\mathcal{F}$ also satisfies property (P2). ◀

Using Theorem 26, we obtain an $([m], k, \eta, \varepsilon)$-cover family $\mathcal{S}$ with $|\mathcal{S}| = \lceil \log_k m \rceil \cdot k \cdot 2^{2m/k}$ in $|\mathcal{S}|^{O(1)}$ time, hence proving Lemma 19.

References

1 Divesh Aggarwal, Rishav Gupta, Aditya Morolia, and Chuanqi Zhang. Mind the gap? not for svp hardness under eth!, 2026. arXiv:2504.02695.

2 Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, 1997. doi:10.1006/jcss.1997.1472.

3 Mitali Bafna, Karthik C. S., and Dor Minzer. Near optimal constant inapproximability under ETH for fundamental problems in parameterized complexity. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 2118–2129. ACM, 2025. doi:10.1145/3717823.3718257.

4 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 p norms. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 553–566. ACM, 2023. doi:10.1145/3564246.3585214.

5 Elwyn R. Berlekamp, Robert J. McEliece, and Henk C. A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Trans. Inf. Theory*, 24(3):384–386, 1978. doi:10.1109/TIT.1978.1055873.

6 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. *J. ACM*, 68(3):16:1–16:40, 2021. doi:10.1145/3444942.

7 Nir Bitansky, Prahladh Harsha, Yuval Ishai, Ron D. Rothblum, and David J. Wu. Dot-product proofs and their applications. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 806–825. IEEE, 2024. doi:10.1109/FOCS61266.2024.00057.

8 Irit Dinur. Mildly exponential reduction from gap 3sat to polynomial-gap label-cover. *Electron. Colloquium Comput. Complex.*, TR16-128, August 2016. URL: https://eccc.weizmann.ac.il/report/2016/128.

9 Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer, 1999. doi:10.1007/978-1-4612-0515-9.

10 Rodney G. Downey, Michael R. Fellows, Alexander Vardy, and Geoff Whittle. The parametrized complexity of some fundamental problems in coding theory. *SIAM J. Comput.*, 29(2):545–570, 1999. doi:10.1137/S0097539797323571.

11 I. Dumer, D. Micciancio, and M. 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.

12 Venkatesan Guruswami, Bingkai Lin, Xuandi Ren, Yican Sun, and Kewen Wu. Almost optimal time lower bound for approximating parameterized clique, csp, and more, under ETH. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 2136–2144. ACM, 2025. doi:10.1145/3717823.3718130.

13 Venkatesan Guruswami, Xuandi Ren, and Sai Sandeep. Baby PIH: Parameterized Inapproximability of Min CSP. In Rahul Santhanam, editor, *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:17, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2024.27.

14 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, March 2001. doi:10.1006/jcss.2000.1727.

15 Shuangle Li, Bingkai Lin, and Yuwei Liu. Improved lower bounds for approximating parameterized nearest codeword and related problems under ETH. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, Tallinn, Estonia, July 8-12, 2024*, volume 297 of *LIPIcs*, pages 107:1–107:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ICALP.2024.107.

16 Pasin Manurangsi. Tight running time lower bounds for strong inapproximability of maximum k-coverage, unique set cover and related problems (via t-wise agreement testing theorem). In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 62–81. SIAM, 2020. doi:10.1137/1.9781611975994.5.

17 Craig A. Tovey. A simplified np-complete satisfiability problem. *Discret. Appl. Math.*, 8(1):85–89, 1984. doi:10.1016/0166-218X(84)90081-7.