

Hardness of the (Approximate) Shortest Vector Problem: A Simple Proof via Reed-Solomon Codes

Huck Bennett*

Chris Peikert[†]

February 14, 2022

Abstract

We give a *simple* proof that the (approximate, decisional) Shortest Vector Problem is NP-hard under a randomized reduction. Specifically, we show that for any $p \geq 1$ and any constant $\gamma < 2^{1/p}$, the γ -approximate problem in the ℓ_p norm (γ -GapSVP $_p$) is not in RP unless $\text{NP} \subseteq \text{RP}$. Our proof follows an approach pioneered by Ajtai (STOC 1998), and strengthened by Micciancio (FOCS 1998 and SICOMP 2000), for showing hardness of γ -GapSVP $_p$ using *locally dense lattices*. We construct such lattices simply by applying “Construction A” to Reed-Solomon codes with suitable parameters, and prove their local density via an elementary argument originally used in the context of Craig lattices.

As in all known NP-hardness results for GapSVP $_p$ with $p < \infty$, our reduction uses randomness. Indeed, it is a notorious open problem to prove NP-hardness via a deterministic reduction. To this end, we additionally discuss potential directions and associated challenges for derandomizing our reduction. In particular, we show that a close deterministic analogue of our local density construction would improve on the state-of-the-art explicit Reed-Solomon list-decoding lower bounds of Guruswami and Rudra (STOC 2005 and IEEE Transactions on Information Theory 2006).

As a related contribution of independent interest, we also give a polynomial-time algorithm for decoding n -dimensional “Construction A Reed-Solomon lattices” (with different parameters than those used in our hardness proof) to a distance within an $O(\sqrt{\log n})$ factor of Minkowski’s bound. This asymptotically matches the best known distance for decoding near Minkowski’s bound, due to Mook and Peikert (IEEE Transactions on Information Theory 2022), whose work we build on with a somewhat simpler construction and analysis.

*Oregon State University, huck.bennett@oregonstate.edu. Part of this work was completed while the author was at the University of Michigan and supported by NSF Grant No. CCF-2006857. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the National Science Foundation.

[†]University of Michigan and Algorand, Inc., cpeikert@umich.edu. This author’s work was supported by NSF Grant No. CCF-2006857.

1 Introduction

[I]t may easily happen that other, perhaps in some sense simpler, lattices also have the properties that are required from L to complete the proof... There are different reasons which may motivate the search for such a lattice: to make the proof deterministic; to improve the factor in the approximation result; to make the proof simpler.

Miklós Ajtai, [Ajt98, Remark 2]

A *lattice* \mathcal{L} is the set of all integer linear combinations of some n linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^m$. The matrix $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ whose columns are these vectors is called a *basis* of \mathcal{L} , and n is called its *rank*. Formally, the lattice \mathcal{L} generated by B is defined as

$$\mathcal{L} = \mathcal{L}(B) := \left\{ \sum_{i=1}^n a_i \mathbf{b}_i : a_1, \dots, a_n \in \mathbb{Z} \right\}.$$

Lattices are classically studied mathematical objects, and have proved invaluable in many computer science applications, especially the design and analysis of cryptosystems. Indeed, the area of lattice-based cryptography, which designs cryptosystems whose security is based on the apparent intractability of certain computational problems on lattices, has flourished over the past quarter century. (See [Pei16] and its bibliography for a comprehensive summary and list of references.)

The central computational problem on lattices is the Shortest Vector Problem (SVP): given a lattice basis B as input, the goal is to find a shortest non-zero vector in $\mathcal{L}(B)$. This paper is concerned with its γ -approximate decision version in the ℓ_p norm (γ -GapSVP $_p$), where $p \geq 1$ is fixed and the approximation factor $\gamma = \gamma(n) \geq 1$ is some function of the lattice rank n (often a constant). Here the input additionally includes a distance threshold $s > 0$, and the goal is to determine whether the length (in the ℓ_p norm) $\lambda_1^{(p)}(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|_p$ of the shortest non-zero vector in \mathcal{L} is at most s , or is strictly greater than γs , when one of the two cases is promised to hold. For the exact problem, where $\gamma = 1$, we often simply write GapSVP $_p$.

Motivated especially by its central role in the security of lattice-based cryptography, understanding the complexity of γ -GapSVP has been the subject of a long line of work. In an early technical report, van Emde Boas [vEB81] initiated the study of the hardness of lattice problems more generally, and in particular showed that GapSVP $_\infty$ is NP-hard. Seventeen years later, Ajtai [Ajt98] finally showed similar hardness for the important Euclidean case of $p = 2$, i.e., he showed that exact GapSVP $_2$ is NP-hard, though under a *randomized* reduction. Subsequent work [CN98, Mic98, Kho03, Kho04, HR07, Mic12] improved this by showing that γ -GapSVP $_p$ in any ℓ_p norm is NP-hard to approximate for any constant $\gamma \geq 1$, and hard for nearly polynomial factors $\gamma = n^{\Omega(1/\log \log n)}$ assuming stronger complexity assumptions, also using randomized reductions. Recent work [AS18, BPT22] has also shown the *fine-grained* hardness of γ -GapSVP $_p$ for small constants γ (again under randomized reductions). On the other hand, γ -GapSVP $_p$ for finite $p \geq 2$ is unlikely to be NP-hard for approximation factors $\gamma \geq C_p \sqrt{n}$ (where C_p is a constant depending only on p) [GG98, AR04, Pei07], and the security of lattice-based cryptography relies on the conjectured hardness of GapSVP or other problems for even larger (but typically polynomial) factors.

While this line of work has been very successful in showing progressively stronger hardness of approximation and fine-grained hardness for γ -GapSVP $_p$, it leaves some other important issues unresolved. First, for $p \neq \infty$ the hardness reductions and their analysis are rather complicated, and second, they are randomized. Indeed, it is a notorious, long-standing open problem to prove that GapSVP $_p$ is NP-hard, even in its exact form, under a deterministic reduction for some finite p . While there have been some potential steps in this direction [Mic98, Mic12], e.g., using plausible number-theoretic conjectures that appear very hard to prove, there has been no new progress on this front for a decade.

1.1 Our Contributions

The primary contribution of this work is to give a substantially simpler proof that γ -GapSVP_{*p*} is NP-hard under a randomized reduction, for any $p \geq 1$ and constant $\gamma < 2^{1/p}$. The heart of our reduction is a family of “gadget” lattices \mathcal{L} derived from Reed-Solomon codes $\mathcal{C} \subseteq \mathbb{F}_q^n$ (for prime q) via the very natural “Construction A” [CS99], which simply sets $\mathcal{L} = \mathcal{C} + q\mathbb{Z}^n$. These lattices also are closely related to a family studied by Craig [Cra], and we take advantage of this similarity in our analysis (see Section 1.3 for details).

Theorem 1.1 (Hardness of γ -GapSVP_{*p*}). *For any $p \geq 1$ and constant γ satisfying $1 \leq \gamma < 2^{1/p}$, γ -GapSVP_{*p*} is not in RP unless $\text{NP} \subseteq \text{RP}$.*

We note that Theorem 1.1 is actually identical to the main result in [Mic98]. As such, it matches the best known NP-hardness of approximation for γ -GapSVP_{*p*} (i.e., largest γ) achieved by a “one-shot” reduction for all sufficiently small p , including $p = 2$. By “one-shot,” we mean that the reduction does not amplify the approximation factor from an initial fixed constant to an arbitrary constant (or more) via tensoring, as is done in [Kho04, HR07, Mic12]. (It is an interesting question whether our hard γ -GapSVP_{*p*} instances are amenable to tensoring; see Section 1.4.)

Although our reduction still uses randomness, we believe that it may be easier to derandomize than previous reductions, both due to its simplicity, and because of its close connection to prior work showing hardness of minimum distance problem on codes via a deterministic reduction [CW09]. To that end, in Section 5 we also describe two approaches to potentially derandomizing our reduction, both of which aim to deterministically construct a particular lattice coset and lower bound the number of short vectors in it (see Section 1.2 for the motivation for this). The first approach is based on Fourier analysis, using similar techniques to those in [CW09], and the second is based on “smooth” proxies for point-counting functions.

We also show that a close deterministic analog of our randomized local-density construction would imply improved explicit Reed-Solomon list-decoding lower bounds, going beyond the current state of the art from [GR05]. One may interpret this implication either pessimistically, as a barrier to a very strong derandomization of our reduction, or optimistically, as a potential route to improve Reed-Solomon list-decoding lower bounds. Here there is a further connection between the two problems, in that [GR05] obtains its list-decoding lower bounds by using the same Fourier-analytic tool underlying one of our derandomization attempts—specifically, the Weil bound for character sums (Equation (25)). Unfortunately, the Weil bound falls just short of what we need in our context. (The Weil bound and related techniques were first used for counting Reed-Solomon code words in [CW04], and were also used in the deterministic hardness reduction for the minimum distance problem on codes in [CW09].)

Efficient decoding near Minkowski’s bound. As a separate contribution of independent interest, in Section 4 we give a polynomial-time algorithm for decoding “Construction A Reed-Solomon lattices” of rank n —the same family of lattices as in our hardness reduction, but instantiated with different parameters—to a distance within a $O(\sqrt{\log n})$ factor of Minkowski’s bound.¹ The $O(\sqrt{\log n})$ factor in our result asymptotically matches the best factor known from prior work [MP22], which is for a different family of lattices. In fact, we rely on one of the main underlying theorems from that work, but give a simpler construction and analysis based on individual Reed-Solomon codes instead of towers of BCH codes.

Let $\text{RS}_q[k, S]$ denote the dimension- k Reed-Solomon code over \mathbb{F}_q with evaluation set S (defined below in Equation (3)). Note that $n = q$ is the rank of the lattice \mathcal{L} in the following theorem.

Theorem 1.2 (Decoding near Minkowski’s bound, informal). *Let q be prime and let $k := \lfloor q/(2 \log_2 q) \rfloor \leq q/2$. Then for the “Construction A Reed-Solomon” lattice $\mathcal{L} := \text{RS}_q[q - k, \mathbb{F}_q] + q\mathbb{Z}^q \subseteq \mathbb{Z}^q$:*

1. *We have $\Omega(\sqrt{q/\log q}) \leq \lambda_1(\mathcal{L}) \leq \sqrt{q} \cdot \det(\mathcal{L})^{1/q} \leq O(\sqrt{q})$, i.e., the minimum distance is within a $O(\sqrt{\log q})$ factor of Minkowski’s bound.*

¹Minkowski’s bound gives an upper bound on the “normalized density” of a lattice \mathcal{L} . Specifically, it asserts that $\lambda_1^{(2)}(\mathcal{L}) \leq \sqrt{n} \cdot \det(\mathcal{L})^{1/n}$ for all rank- n lattices \mathcal{L} , where $\det(\mathcal{L}) = \sqrt{\det(B^T B)}$ for any basis B of \mathcal{L} .

2. There is an algorithm that, on input q and a vector $\mathbf{y} \in \mathbb{R}^q$, outputs all lattice vectors $\mathbf{v} \in \mathcal{L}$ satisfying $\|\mathbf{y} - \mathbf{v}\| \leq C\sqrt{k} \approx C\sqrt{q/(2\log_2 q)}$ in time polynomial in q , for some universal constant $C > 0$.

This result adds to a separate line of work on efficient (list) decoding for various families of lattices [MN08, GP12, DP19, MP22]. Recently, Ducas and van Woerden [DvW21] further motivated this study by showing cryptographic applications of lattices that can be efficiently decoded near Minkowski’s bound. (However, their application is most compelling when the minimum distances of both the lattice and its dual are close to Minkowski’s bound, which is not the case in the present setting.)

1.2 Technical Overview

Here we give an overview of the key new elements in the proof of our main hardness theorem (Theorem 1.1), which are the focus of Section 3. For concision, we defer the technical aspects of our efficient decoding algorithm and derandomization attempts to Sections 4 and 5, respectively.

Besides using randomness, another common feature in nearly all prior hardness results for GapSVP_{*p*} is the use of *locally dense lattices* as advice (the only exception being [Kho03]). Roughly speaking, a locally dense lattice for relative distance $\alpha \in (0, 1)$ in the ℓ_p norm is a lattice \mathcal{L} and a coset $\mathbf{x} + \mathcal{L}$ (i.e., the lattice “shifted by” some vector \mathbf{x}) such that there are at least subexponentially many (in the lattice rank) vectors $\mathbf{v} \in \mathbf{x} + \mathcal{L}$ satisfying $\|\mathbf{v}\|_p \leq \alpha \cdot \lambda_1^{(p)}(\mathcal{L})$. One may view such a coset $\mathbf{x} + \mathcal{L}$ as a “bad” configuration for list-decoding \mathcal{L} to within relative distance α in the ℓ_p norm, because there are many lattice vectors relatively close to $-\mathbf{x}$.²

Prior works have obtained locally density from a variety of lattice families: the Schnorr-Adelman prime number lattices [Ajt98, CN98, Mic98]; a variant of Construction A [Kho04] and Construction D [Mic12] applied to (towers of) BCH codes; and random sublattices of \mathbb{Z}^n and lattices with exponential kissing number [AS18, BPT22]. In this work, we give a simple construction of locally dense lattices from Reed-Solomon codes, as described below.

Our main reduction (Theorem 2.10) shows how to use a locally dense lattice for relative distance α in the ℓ_p norm to prove NP-hardness (via a randomized reduction) of γ -GapSVP_{*p*} for any constant $\gamma > 1/\alpha$. This reduction is very similar to those from prior works, so for the remainder of this section we focus on summarizing our new construction of locally dense lattices.

Locally dense lattices from Reed-Solomon codes. We start with some basic definitions and facts used in our construction. Recall that the Construction A lattice obtained from a linear code $\mathcal{C} \subseteq \mathbb{F}_q^n$ for some prime q is defined as $\mathcal{L} := \mathcal{C} + q\mathbb{Z}^n$, i.e., an integer vector $\mathbf{z} \in \mathbb{Z}^n$ is in the lattice if and only if $\mathbf{z} \bmod q$ is a code word. In fact, it will often be convenient to work with an equivalent “dual view” of Construction A lattices. Namely, if $H \in \mathbb{F}_q^{k \times n}$ is a parity-check matrix of a linear code $\mathcal{C} := \ker(H) \subseteq \mathbb{F}_q^n$ for prime q , then the *parity-check lattice* $\mathcal{L}^\perp(H)$ obtained from H is defined as

$$\mathcal{L}^\perp(H) := \{\mathbf{z} \in \mathbb{Z}^n : H\mathbf{z} = \mathbf{0} \in \mathbb{F}_q^k\} = \ker(H) + q\mathbb{Z}^n = \mathcal{C} + q\mathbb{Z}^n . \tag{1}$$

Such lattices have determinant $\det(\mathcal{L}^\perp(H)) = |\mathbb{Z}^n / \mathcal{L}^\perp(H)| \leq q^k$, with equality exactly when H has full row rank (see Lemma 2.2).

We next define the family of parity-check matrices $H = H_q(k, S)$ that we use to construct our family of locally dense lattices. Such a matrix is parameterized by a prime q , a positive integer k , and a set $S \subseteq \mathbb{F}_q$.

²For technical reasons, the formal definition of local density, given in Definition 2.6, also requires a linear transform that maps the short vectors in $\mathbf{x} + \mathcal{L}$ onto the set of all binary vectors of a given dimension. Such a transform can be obtained by random sampling using a probabilistic version of Sauer’s Lemma (see Theorem 2.7) that is now standard in this context [Ajt98, Mic98]. Therefore, we defer further discussion of this issue to the main body.

Letting s_0, \dots, s_{n-1} be the elements of S in some arbitrary order, we define

$$H = H_q(k, S) := \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ s_0 & s_1 & s_2 & \cdots & s_{n-1} \\ s_0^2 & s_1^2 & s_2^2 & \cdots & s_{n-1}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s_0^{k-1} & s_1^{k-1} & s_2^{k-1} & \cdots & s_{n-1}^{k-1} \end{pmatrix} \in \mathbb{F}_q^{k \times n}. \quad (2)$$

That is, $H_q(k, S)$ is the transposed Vandermonde matrix whose (i, s) th entry is s^i , where for convenience we index the rows and columns of $H_q(k, S)$ by $i \in \{0, \dots, k-1\}$ and $s \in S$, respectively, and define $0^0 := 1$.

The matrix $H = H_q(k, S)$ defined in Equation (2) is a *generator matrix* of the dimension- k Reed-Solomon code

$$\text{RS}_q[k, S] := \{(p(s))_{s \in S} : p \in \mathbb{F}_q[x], \deg(p) < k\} \quad (3)$$

over \mathbb{F}_q with evaluation set S , and hence is a parity-check matrix of its dual code, which is a so-called *generalized* Reed-Solomon (GRS) code (see [Hal, Theorem 5.1.6]). Moreover, in the special case where $S = \mathbb{F}_q$, it turns out that $H_q(k, S)$ is a parity-check matrix for the (ordinary) Reed-Solomon code $\text{RS}_q[q-k, \mathbb{F}_q]$ of dimension $q-k$ with evaluation set $S = \mathbb{F}_q$. So, $\mathcal{L}^\perp(H_q(k, \mathbb{F}_q)) = \text{RS}_q[q-k, \mathbb{F}_q] + q\mathbb{Z}^q$ is the Construction A lattice corresponding to the Reed-Solomon code $\text{RS}_q[q-k, \mathbb{F}_q]$. For simplicity, in this overview we restrict to these ‘‘Construction A Reed-Solomon’’ lattices by taking $S = \mathbb{F}_q$, but note that our results generalize to any sufficiently large $S \subseteq \mathbb{F}_q$.

It is easy to see that for $k < q$, the GRS code having parity-check matrix H has minimum distance (in the Hamming metric) $k+1$: any k columns of H are linearly independent, because they form a transposed Vandermonde matrix, while any $k+1$ obviously are not. Therefore, the corresponding Construction A lattice $\mathcal{L} := \ker(H) + q\mathbb{Z}^q$ has minimum distance $\lambda_1^{(1)}(\mathcal{L}) \geq k+1$ in the ℓ_1 norm. The key to our local density construction and its $\alpha \approx 1/2^{1/p}$ relative distance is that the ℓ_1 minimum distance is in fact almost *twice* this large (at least): in Theorem 3.1 we show that $\lambda_1^{(1)}(\mathcal{L}) \geq 2k$ when $k \leq q/2$. The proof is short and elementary, proceeding via Newton’s identities, and is closely related to the analysis of Craig lattices [Cra] (see Section 1.3).

Obtaining a dense coset. Because the determinant of \mathcal{L} (i.e., the number of its integer cosets) is q^k , the pigeonhole principle immediately implies that there exists an integer coset of \mathcal{L} containing at least $\binom{q}{h}/q^k$ *binary* vectors in $\{0, 1\}^q$ with Hamming weight h , which have ℓ_1 norm h . By setting parameters appropriately, this yields a coset with subexponentially many vectors of ℓ_1 norm at most $\alpha \cdot \lambda_1^{(1)}(\mathcal{L})$, for any constant $\alpha > 1/2$.

More specifically, set $h := \alpha \cdot (2k) \leq \alpha \cdot \lambda_1^{(1)}(\mathcal{L})$ and $q \approx k^{1/\varepsilon}$ for some positive constant $\varepsilon < 1 - 1/(2\alpha)$. (For simplicity, assume that h is an integer.) Then there must exist an integer coset of \mathcal{L} containing at least

$$\frac{\binom{q}{h}}{q^k} \geq \left(\frac{q}{h}\right)^h \cdot q^{-k} = \frac{q^{(2\alpha-1)k}}{(2\alpha k)^{2\alpha k}} \approx \frac{q^{(2\alpha-1)k}}{q^{2\varepsilon\alpha k}} = q^{(2(1-\varepsilon)\alpha-1)k} = q^{\Omega(k)} = q^{\Omega(q^\varepsilon)} \quad (4)$$

weight- h binary vectors, which is subexponentially large in q .

The above shows the *existence* of a suitable coset, but following previous works, it is straightforward to show that a *randomly sampled* coset from a suitable distribution is likely to have enough short vectors (see Lemma 3.3). Indeed, the difference between showing that such a coset exists, versus sampling one efficiently, versus deterministically computing one efficiently, is the main technical difference between getting a non-uniform, versus randomized, versus deterministic hardness reduction (respectively) for GapSVP_p using these techniques.

The above argument generalizes to arbitrary ℓ_p norms for finite p , albeit for larger relative distances $\alpha > 1/2^{1/p}$. Because \mathcal{L} is integral, $\lambda_1^{(1)}(\mathcal{L}) \geq 2k$ implies that $\lambda_1^{(p)}(\mathcal{L}) \geq (2k)^{1/p}$ for any finite $p \geq 1$. Moreover, reparameterizing the calculation in Equation (4) by choosing $\alpha > 1/2^{1/p}$ and setting $h := \alpha^p \cdot (2k)$ shows that

some coset of \mathcal{L} contains subexponentially many binary vectors of Hamming weight h , and hence of ℓ_p norm $h^{1/p} = \alpha \cdot (2k)^{1/p} \leq \alpha \cdot \lambda_1^{(p)}(\mathcal{L})$. Therefore, this construction yields locally dense lattices in the ℓ_p norm for any constant relative distance $\alpha > 1/2^{1/p}$, which by our main reduction implies [Theorem 1.1](#), i.e., randomized NP-hardness of γ -GapSVP $_p$ for any constant $\gamma < 2^{1/p}$.

1.3 Additional Related Work

The Construction A Reed-Solomon lattices we use are closely related to a family of algebraic lattices studied by Craig [[Cra](#)] (see also [[CS99](#), Chapter 8, Section 6]), and our proof of [Theorem 3.1](#) is inspired by and very similar to an argument lower bounding the minimum distance of Craig lattices from [[CS99](#), Chapter 8, Section 6, Theorem 7].³ In fact, the guarantees on the minimum distances and determinants of Craig lattices given there are such that our hardness proof for GapSVP could be made to work using Craig lattices instead of Construction A Reed-Solomon lattices. We have chosen to work with the latter family because they are more elementary and general, and because of their direct connections to efficient lattice decoding, list-decoding lower bounds, and derandomization techniques, as explored in [Sections 4](#) and [5](#).

We also note that locally dense lattices appear in other work on the complexity of lattice problems. In particular, all prior NP- and fine-grained hardness reductions for the problem of Bounded Distance Decoding to relative distance α' in the ℓ_p norm (α' -BDD $_p$) [[LLM06](#), [BP20](#), [BPT22](#)], another important lattice problem, use locally dense lattices. In fact, the smallest α' for which we know how to show NP-hardness of α' -BDD $_p$ is essentially the smallest $\alpha(p)$ for which we can construct locally dense lattices for relative distance $\alpha(p)$ in the ℓ_p norm.

Additionally, *locally dense codes* play an analogous role to locally dense lattices in the complexity of coding problems, and particularly the Minimum Distance Problem (MDP), which is the analog of SVP for codes. Indeed, the work of Micciancio, Dumer, and Sudan [[DMS99](#)] adapted the local density framework to codes and used it to show randomized hardness of approximation for MDP. This reduction was subsequently made deterministic by Cheng and Wan [[CW09](#)], who, as mentioned above, constructed locally dense codes from Reed-Solomon codes. Micciancio [[Mic14](#)] formalized the locally dense code framework, gave a simplified deterministic NP-hardness proof for MDP, and discussed the search for analogous deterministic constructions of locally dense lattices. (Austrin and Khot [[AK11](#)] also gave a simplified deterministic hardness reduction for MDP, but it does not obviously use the same framework.)

Finally, with respect to our result in [Theorem 1.2](#), we note that a number of previous works have given algorithms for efficient decoding near Minkowski’s bound on other families of lattices. These include [[MN08](#), [GP12](#)], which showed how to (list) decode to a distance within a $O(n^{1/4})$ factor of Minkowski’s bound on Barnes-Wall lattices; [[DP19](#)], which showed how to decode to a distance within a $O(\log n)$ factor of Minkowski’s bound on a family of discrete-logarithm lattices; and [[MP22](#)], which showed how to (list) decode to a distance within a $O(\sqrt{\log n})$ factor of Minkowski’s bound on lattices obtained by applying Construction D to towers of BCH codes.

1.4 Open Questions

The obvious question left open by our work is whether our reduction can be derandomized, using the same family of lattices. Addressing this is the main focus of [Section 5](#), and we defer further discussion of it to there.

Another interesting question is whether it is possible to amplify the approximation factor γ achieved by our reduction via *tensoring*, as is done in previous works [[Kho04](#), [HR07](#), [Mic12](#)]. Recall that the tensor product of lattices $\mathcal{L}_1, \mathcal{L}_2$ having respective bases B_1, B_2 is defined as $\mathcal{L}_1 \otimes \mathcal{L}_2 := \mathcal{L}(B_1 \otimes B_2)$, and that $\lambda_1(\mathcal{L}_1 \otimes \mathcal{L}_2) \leq \lambda_1(\mathcal{L}_1) \cdot \lambda_1(\mathcal{L}_2)$. To see why tensoring might be useful for gap amplification, observe that if this bound was an equality, then the m -fold tensor product $\mathcal{L}^{\otimes m}$ of \mathcal{L} with itself would satisfy $\lambda_1(\mathcal{L}^{\otimes m}) = \lambda_1(\mathcal{L})^m$, which would amplify the approximation factor in an instance of γ -GapSVP $_p$ to γ^m . However, this inequality

³Specifically, [[CS99](#)] considers Craig lattices obtained from the coefficient vectors of polynomials in principal ideals of the form $(x-1)^m R$ in rings of the form $R = \mathbb{Z}[x]/(x^p-1)$, for some prime p and integer $m \geq 1$. The original definition of [[Cra](#)] is slightly different, and uses the “canonical” (Minkowski) embedding of such ideals in the ring of integers $R = \mathbb{Z}[x]/(\Phi_p(x))$ of the p th cyclotomic number field.

is *not* tight in general (see, e.g., [HR07, Lemma 2.3]). Nevertheless, prior reductions have been able to prove similar inequalities for certain families of lattices, but it is not immediately obvious how to adapt their techniques to our setting.

Finally, as previously observed in [BP20], we do not have a good understanding of the smallest possible relative distance $\alpha^* = \alpha^*(p)$ for locally dense lattices in ℓ_p norms for finite $p > 2$. Besides showing better “one-shot” hardness for γ -GapSVP $_p$, achieving smaller factors $\alpha(p)$ for such p (ideally, $\alpha(p) \approx \alpha^*(p)$) would imply improved (randomized) NP-hardness results for Bounded Distance Decoding (specifically, it would show hardness of α' -BDD $_p$ for any $\alpha' > \alpha(p)$). The triangle inequality implies a lower bound of $\alpha^* \geq 1/2$ in any norm, and, as observed in [Mic98], we have $\alpha^*(2) \geq 1/\sqrt{2}$. So, the relative-distance factors $\alpha(p) \approx 1/2^{1/p}$ for local density achieved by [Mic98] and this work are essentially optimal for $p = 1$ and $p = 2$ (and conjecturally for all $p \in [1, 2]$). However, they deteriorate as p increases. On the other hand, the factors $\alpha(p)$ obtained in [BP20] satisfy $\lim_{p \rightarrow \infty} \alpha(p) = 1/2$, and hence are nearly optimal for large p , but are poor even for some $p > 2$. In particular, they satisfy $\alpha(p) < 1/\sqrt{2}$ only for $p > p_1 \approx 4.2773$, even though we believe that $\alpha(p) < 1/\sqrt{2}$ should be achievable for all $p > 2$. (We note in passing that $\alpha(p) \approx 1/\sqrt{2}$ is achievable for *all* p via norm embeddings [RR06].) It seems plausible that a family similar to our Construction A Reed-Solomon lattices \mathcal{L} —perhaps obtained by parameterizing the underlying Reed-Solomon code differently, using a different algebraic code, or by randomly sparsifying \mathcal{L} as in [BP20]—might achieve this.

Acknowledgements. We thank Swastik Kopparty [Kop20] for very helpful answers to several of our questions, and for pointing us to the contents of Sections 5.1 and 5.2.

2 Preliminaries

Throughout this work we adopt the convention that $0^0 := 1$ in any ring. For a positive integer k , define $[k] := \{0, 1, \dots, k-1\}$.

In general, every vector or matrix is indexed by some specified set S . For example, $\mathbf{x} \in \mathbb{Z}^S$ is an integer vector indexed by S , having an entry $x_s \in \mathbb{Z}$ for each $s \in S$ (and no other entries). When the index set is $[n]$ for some non-negative integer n , we usually omit the brackets in the exponent and just write, e.g., \mathbb{Z}^n . We emphasize that in this case the indices start from zero. An object indexed by a finite set S of size $n = |S|$ can be reindexed by $[n]$, simply by enumerating $S = \{s_0, \dots, s_{n-1}\}$ under some arbitrary order, and identifying index s_i with index i .

For a finite set S and a positive integer $h \leq |S|$, let $B_{S,h} := \{\mathbf{v} \in \{0, 1\}^S : \|\mathbf{v}\|_1 = h\}$ be the set of binary vectors indexed by S of Hamming weight h . As above, when $S = [n]$ we often write $B_{n,h}$. Finally, let $\mathcal{B}_p^n(r) := \{\mathbf{x} : \|\mathbf{x}\|_p \leq r\} \subset \mathbb{R}^n$ denote the real n -dimensional ℓ_p ball of radius r centered at the origin.

2.1 Basic Lattice Definitions

Given a lattice $\mathcal{L} = \mathcal{L}(B)$ with basis $B \in \mathbb{R}^{m \times n}$, we define the *rank* of \mathcal{L} to be n and the (ambient) *dimension* of \mathcal{L} to be m . We denote the *minimum distance* of \mathcal{L} in the ℓ_p norm, which is the length of a shortest non-zero vector in \mathcal{L} , by

$$\lambda_1^{(p)}(\mathcal{L}) := \min_{\mathbf{x} \in \mathcal{L} \setminus \{0\}} \|\mathbf{x}\|_p.$$

The central problem that we study in this work asks about the value of $\lambda_1^{(p)}(\mathcal{L})$ for a given input lattice \mathcal{L} .

Definition 2.1. For $p \geq 1$ and $\gamma = \gamma(n) \geq 1$, the decisional, γ -approximate Shortest Vector Problem in the ℓ_p norm (γ -GapSVP $_p$) is the promise problem defined as follows. The input consists of a basis $B \in \mathbb{Z}^{m \times n}$ of an integer lattice \mathcal{L} and a distance threshold $s > 0$, and the goal is to determine whether the input is a YES instance or a NO instance, where these are defined as follows:

- YES instance: $\lambda_1^{(p)}(\mathcal{L}) \leq s$.
- NO instance: $\lambda_1^{(p)}(\mathcal{L}) > \gamma s$.

We define the *determinant* of \mathcal{L} to be $\det(\mathcal{L}) := \sqrt{\det(B^T B)}$, which is equal to $|\det(B)|$ when $m = n$ (i.e., when \mathcal{L} is full-rank). We note that determinant is well defined because, although lattice bases are not unique, they are equivalent up to multiplication on the right by unimodular matrices.

The density of a rank- n lattice \mathcal{L} is captured by its so-called *root Hermite factor* $\lambda_1(\mathcal{L})/\det(\mathcal{L})^{1/n}$.⁴ The density of a lattice corresponds to its quality in various applications, including as the set of centers of a sphere packing and as an error-correcting code. Minkowski's bound asserts that the root Hermite factor of such a rank- n lattice is at most \sqrt{n} , which is convenient to write in expanded form as

$$\lambda_1(\mathcal{L}) \leq \sqrt{n} \cdot \det(\mathcal{L})^{1/n} . \quad (5)$$

2.2 Parity-Check Matrices and Lattices

For a prime q and a matrix $H \in \mathbb{F}_q^{k \times n}$, we define the *parity-check lattice* $\mathcal{L}^\perp(H)$ obtained from H as

$$\mathcal{L}^\perp(H) := \{\mathbf{z} \in \mathbb{Z}^n : H\mathbf{z} = \mathbf{0}\} = \ker(H) + q\mathbb{Z}^n . \quad (6)$$

Note that $\mathcal{L}^\perp(H)$ is simply the ‘‘Construction A’’ lattice [CS99, Chapter 5, Section 2] of the linear error-correcting code \mathcal{C} having H as a parity-check matrix, i.e., $\mathcal{C} = \{\mathbf{c} \in \mathbb{F}_q^n : H\mathbf{c} = \mathbf{0}\}$. More generally, for any ‘‘syndrome’’ $\mathbf{u} \in \mathbb{F}_q^k$ we define

$$\mathcal{L}_\mathbf{u}^\perp(H) := \{\mathbf{x} \in \mathbb{Z}^n : H\mathbf{x} = \mathbf{u}\} .$$

If there exists some $\mathbf{x} \in \mathbb{Z}^n$ such that $H\mathbf{x} = \mathbf{u}$, then it follows immediately that $\mathcal{L}_\mathbf{u}^\perp(H)$ is simply the lattice coset $\mathbf{x} + \mathcal{L}^\perp(H)$. So, we can identify cosets of $\mathcal{L}^\perp(H)$ by their corresponding syndromes. We recall some standard properties of parity-check lattices, and give a proof for self-containment.

Lemma 2.2. *Let q be a prime, let k and n be positive integers, and let $H \in \mathbb{F}_q^{k \times n}$ be a parity-check matrix. Then the parity-check lattice $\mathcal{L} = \mathcal{L}^\perp(H)$ has rank n and determinant $\det(\mathcal{L}) \leq q^k$, with equality if and only if the rows of H are linearly independent.*

Proof. The first claim follows simply by noting that $q\mathbb{Z}^n \subseteq \mathcal{L}^\perp(H) \subseteq \mathbb{Z}^n$. For the determinant, observe that the map $\mathbf{x} \mapsto H\mathbf{x}$ is an additive-group homomorphism from \mathbb{Z}^n to \mathbb{F}_q^k , and that $\mathcal{L}^\perp(H)$ is its kernel by definition. So, by the first isomorphism theorem, the map induces an isomorphism from the quotient group $\mathbb{Z}^n/\mathcal{L}^\perp(H)$ to the image $\text{Im}(H) = \{H\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\} \subseteq \mathbb{F}_q^k$, where the subset relation is an equality if and only if the rows of H are linearly independent. The claim then follows from the fact that $\det(\mathcal{L}^\perp(H)) = |\mathbb{Z}^n/\mathcal{L}^\perp(H)| = |\text{Im}(H)|$. \square

We next formally define the family of parity-check lattices that are at the heart of our construction of locally dense lattices.

Definition 2.3. For a prime q , positive integer k , and set $S \subseteq \mathbb{F}_q$, define $H_q(k, S) \in \mathbb{F}_q^{k \times S}$ to be the matrix H whose rows and columns are respectively indexed by $[k] = \{0, 1, \dots, k-1\}$ and S , and whose (i, s) th entry is

$$H_{i,s} := s^i \in \mathbb{F}_q .$$

(Recall that $0^0 := 1$.) Equivalently, if we enumerate $S = \{s_0, \dots, s_{n-1}\}$ in some arbitrary order, we have

$$H = H_q(k, S) := \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ s_0 & s_1 & s_2 & \cdots & s_{n-1} \\ s_0^2 & s_1^2 & s_2^2 & \cdots & s_{n-1}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s_0^{k-1} & s_1^{k-1} & s_2^{k-1} & \cdots & s_{n-1}^{k-1} \end{pmatrix} \in \mathbb{F}_q^{k \times n} . \quad (7)$$

⁴This ratio is the square root of the Hermite factor $\gamma(\mathcal{L}) := (\lambda_1(\mathcal{L})/\det(\mathcal{L})^{1/n})^2$, which is defined in this way for historical reasons.

Notice that H is a transposed Vandermonde matrix. In particular, if $k \leq n$ then its rows are linearly independent, and so $\det(\mathcal{L}^\perp(H)) = q^k$ by [Lemma 2.2](#).

We recall from the introduction that $H = H_q(k, S)$ is a generator matrix (of row vectors) of the dimension- k Reed-Solomon code over \mathbb{F}_q with evaluation set S , and hence H is a parity-check matrix of its dimension- $(n-k)$ dual code. So, $\mathcal{L}^\perp(H)$ is the Construction A lattice of this dual code. These dual codes are in fact *generalized Reed-Solomon codes*, a family of codes that include Reed-Solomon codes as a special case and that are closed under taking duals (see [[Hal](#), Theorem 5.1.6]). Moreover, in the special case where $S = \mathbb{F}_q$, the matrix $H = H_q(k, \mathbb{F}_q)$ is in fact a parity-check matrix of an (ordinary) Reed-Solomon code. For our hardness proof it suffices to use this special case; i.e., we show NP-hardness (under a randomized reduction) of GapSVP by using Construction A lattices of (ordinary) Reed-Solomon codes as gadgets. See [Section 4](#) and [Section 5.1](#) for other connections between these lattices and Reed-Solomon codes.

2.3 Symmetric Polynomials

A *symmetric polynomial* $P(x_1, x_2, \dots, x_m)$ is a polynomial that is invariant under any permutation of its variables, i.e., $P(x_1, \dots, x_m) = P(x_{\pi(1)}, \dots, x_{\pi(m)})$ as formal polynomials for all permutations π of $\{1, 2, \dots, m\}$. Because the order of the variables is immaterial, we usually just write a symmetric polynomial as $P(X)$, where $X = \{x_1, \dots, x_m\}$ is the set of variables, and we write $P(T)$ for its evaluation on a multiset T of values.

We next recall two important symmetric polynomials and the relationship between them. For a non-negative integer i , the i th *power sum* of a set X of variables is defined as

$$p_i(X) := \sum_{x \in X} x^i . \tag{8}$$

(Recall that $0^0 := 1$.) For $1 \leq i \leq |X|$, the i th *elementary symmetric polynomial* of X is defined as

$$e_i(X) := \sum_{\substack{Z \subseteq X, \\ |Z|=i}} \prod_{z \in Z} z . \tag{9}$$

That is, $e_i(X)$ is the multilinear polynomial whose monomials consist of all products of i distinct variables from X . We extend this definition to $i = 0$ by setting $e_0(X) := 1$ and to integers $i > |X|$ by setting $e_i(X) := 0$.

Power sums and elementary symmetric polynomials are related by *Newton's identities* (see, e.g., [[Mea92](#)]), which assert that for $1 \leq i \leq |X|$,

$$i \cdot e_i(X) = \sum_{j=1}^i (-1)^{j-1} \cdot e_{i-j}(X) \cdot p_j(X) . \tag{10}$$

The following standard claim uses Newton's identities to show that if the first k power sums of two multisets of field elements coincide, then so do the first k elementary symmetric polynomials of those multisets.

Claim 2.4. *Let T, U be multisets over a prime field \mathbb{F}_q , let $k \leq q$ be a positive integer, and suppose that $p_i(T) = p_i(U)$ for all $i \in [k]$. Then $e_i(T) = e_i(U)$ for all $i \in [k]$.*

Proof. The proof is by (strong) induction. For the base case where $i = 0$, we have by definition that $e_0(T) = e_0(U) = 1$. For the inductive case where $1 \leq i < k$, because $i \neq 0 \in \mathbb{F}_q$ we have that

$$e_i(T) = i^{-1} \cdot \sum_{j=1}^i (-1)^{j-1} \cdot e_{i-j}(T) \cdot p_j(T) = i^{-1} \cdot \sum_{j=1}^i (-1)^{j-1} \cdot e_{i-j}(U) \cdot p_j(U) = e_i(U) ,$$

where the first and third equalities follow from Newton's identities ([Equation \(10\)](#)), and the second equality follows from the claim's hypothesis and the inductive hypothesis (note that the sums involve elementary symmetric polynomials e_{i-j} only for $i - j < i$). \square

We define the *root polynomial* $f_T(x) \in \mathbb{F}[x]$ of a multiset T over a field \mathbb{F} to be

$$f_T(x) := \prod_{t \in T} (x - t) = \sum_{i=0}^{|T|} (-1)^i \cdot e_i(T) \cdot x^{|T|-i}. \quad (11)$$

We then get the following result, which uses [Claim 2.4](#) to show that if sufficiently many of the initial power sums of two multisets are equal, then the multisets themselves are equal.

Proposition 2.5. *Let q be a prime, let $k \leq q/2$ be a positive integer, let T, U be multisets over $\mathbb{F} = \mathbb{F}_q$ of total cardinality $|T| + |U| < 2k$, and suppose that $p_i(T) = p_i(U)$ for all $i \in [k]$. Then $T = U$.*

Proof. Because

$$|T| \equiv p_0(T) = p_0(U) \equiv |U| \pmod{q}$$

and $0 \leq |T| + |U| < 2k \leq q$, it follows that $|T| = |U|$ and hence both $f_T(x), f_U(x)$ have degree $|T| < k$.

Next, by the hypotheses and [Claim 2.4](#), we have that $e_i(T) = e_i(U)$ for all $i \leq |T|$. Therefore, by the equality in [Equation \(11\)](#), $f_T(x)$ and $f_U(x)$ are identical as formal polynomials in $\mathbb{F}[x]$. Finally, because the polynomial ring $\mathbb{F}[x]$ is a unique factorization domain, and because $f_T(x)$ and $f_U(x)$ split over \mathbb{F} by construction, it follows that $T = U$. \square

2.4 Locally Dense Lattices

Roughly speaking, *locally dense lattices* are lattices that have one or more cosets with many relatively short vectors. Somewhat more precisely, a locally dense lattice consists of an integer lattice $\mathcal{L} \subset \mathbb{Z}^n$ and a shift $\mathbf{x} \in \mathbb{Z}^n$ such that for some $\alpha \in (0, 1)$, the number of points in the coset $\mathbf{x} + \mathcal{L}$ of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$ is large (for our purposes, greater than 2^{n^ε} for some constant $\varepsilon > 0$). Therefore, locally dense lattices are not efficiently list decodable, even combinatorially, to within distance $\alpha \cdot \lambda_1(\mathcal{L})$ in the worst case (in particular, around center $-\mathbf{x}$). For the purposes of proving hardness, we also require a linear map T that projects the short vectors in $\mathbf{x} + \mathcal{L}$ onto a lower-dimensional hypercube $\{0, 1\}^r$.

Definition 2.6. For $p \in [1, \infty)$, real $\alpha > 0$, and positive integers r and R , a (p, α, r, R) -*locally dense lattice* consists of an integer lattice of rank R (and some dimension n) represented by a basis matrix $A \in \mathbb{Z}^{n \times R}$, a positive integer ℓ , a shift $\mathbf{x} \in \mathbb{Z}^n$, and a matrix $T \in \mathbb{Z}^{r \times n}$, where

1. $\lambda_1^{(p)}(\mathcal{L}(A)) \geq \ell^{1/p}$ and
2. $\{0, 1\}^r \subseteq T(V) := \{T\mathbf{v} : \mathbf{v} \in V\}$, where $V := (\mathbf{x} + \mathcal{L}(A)) \cap \mathcal{B}_p^n(\alpha \cdot \ell^{1/p})$ is the set of all vectors of ℓ_p norm at most $\alpha \cdot \ell^{1/p}$ in the lattice coset $\mathbf{x} + \mathcal{L}(A)$.

A useful tool for satisfying [Item 2](#) in the above definition is the following probabilistic version of Sauer's Lemma due to Micciancio [[Mic98](#)]. It roughly says that for $n \gg r$, for any large enough collection of vectors $W \subseteq B_{n,h}$ (the weight- h slice of $\{0, 1\}^n$), and for a random matrix $T \in \{0, 1\}^{r \times n}$ whose coordinates are sampled independently with a suitable bias, $\{0, 1\}^r \subseteq T(W)$ with good probability. We emphasize that all the arithmetic in this theorem is done over the integers (not over \mathbb{F}_2).

Theorem 2.7 ([[Mic98](#), Theorem 4]). *Let r, n, h be positive integers, let $W \subseteq B_{n,h}$, and let $\varepsilon > 0$. If $|W| \geq h! \cdot n^{24r\sqrt{h}/\varepsilon}$ and $T \in \{0, 1\}^{r \times n}$ is sampled by setting each entry to 1 independently with probability $1/(4hr)$, then $\{0, 1\}^r \subseteq T(W)$ with probability at least $1 - \varepsilon$.*

2.5 Hardness of GapSVP via Locally Dense Lattices

We next recall a variant of (the decision version of) the Closest Vector Problem (CVP), which will be the hard problem that we reduce to GapSVP_p . In this variant, called GapCVP'_p , the target vector is either within a specified distance of a lattice vector given by a *binary* combination of basis vectors, or *all non-zero integer multiples* of the target vector are more than a γ multiple of this distance from the lattice (where distance is measured in the ℓ_p norm).

Definition 2.8. For $p \in [1, \infty]$, an instance of the γ -GapCVP' $_p$ problem consists of a rank- r lattice basis $B \in \mathbb{Z}^{d \times r}$, a target vector $\mathbf{t} \in \mathbb{Z}^d$, and a distance threshold $s > 0$. The goal is to determine whether an input is a YES instance or a NO instance, where these are defined as follows:

- YES instance: there exists a binary $\mathbf{c} \in \{0, 1\}^r$ such that $\|B\mathbf{c} - \mathbf{t}\|_p \leq s$.
- NO instance: $\text{dist}_p(w\mathbf{t}, \mathcal{L}(B)) > \gamma s$ for all $w \in \mathbb{Z} \setminus \{0\}$.

The following hardness theorem follows via a reduction from Exact Set Cover to GapCVP' $_p$.

Theorem 2.9 ([ABSS93]). *For every $p \in [1, \infty)$ and every constant $\gamma \geq 1$, γ -GapCVP' $_p$ is NP-hard.*

The following theorem gives an polynomial-time reduction from γ -GapCVP' $_p$ to γ' -GapSVP $_p$ for some approximation factors $\gamma > \gamma' \geq 1$, which uses a locally dense lattice as advice. In general, this advice makes the reduction non-uniform, but when the advice is efficiently computable by a (randomized) algorithm, as it is in this and prior works, the procedure is an efficient (randomized) reduction. The reduction below is very similar to the one in [Mic12, Theorem 5.1], but written so as to allow for using an arbitrary locally dense lattice as advice. Due to this similarity, and for concision, we defer its proof to [Appendix A](#).

Theorem 2.10. *Let $p \geq 1$, r and n be positive integers, $\alpha > 0$ be a constant, and γ, γ' be constants satisfying*

$$1/\alpha > \gamma' \geq 1 \text{ and } \gamma \geq \gamma' \cdot \left(\frac{1}{1 - (\alpha\gamma')^p} \right)^{1/p}.$$

There is a deterministic polynomial-time algorithm that, given a γ -GapCVP' $_p$ instance (B, \mathbf{t}, s) of rank r and a (p, α, r, R) -locally dense lattice (A, ℓ, \mathbf{x}, T) as input, outputs a γ' -GapSVP $_p$ instance (B', s') of rank $R + 1$ which is a YES (respectively, NO) instance if (B, \mathbf{t}, s) is a YES (resp., NO) instance.

From these two theorems we get the following hardness results for GapSVP.

Corollary 2.11. *Let $p \geq 1$, let r be a positive integer, let $\alpha > 0$ be a constant, and suppose that there is an algorithm A that computes a $(p, \alpha, r, \text{poly}(r))$ -locally dense lattice in $\text{poly}(r)$ time. Let γ be a constant satisfying $1 \leq \gamma < 1/\alpha$. Then:*

1. *If A is deterministic, then γ -GapSVP $_p$ is NP-hard (and exact GapSVP $_p$ is NP-complete).*
2. *If A is randomized and its output satisfies [Item 1](#) of [Definition 2.6](#) with probability 1 and [Item 2](#) of [Definition 2.6](#) with probability at least $2/3$, then γ -GapSVP $_p$ is not in RP unless $\text{NP} \subseteq \text{RP}$.⁵*
3. *If A is randomized, and its output satisfies [Items 1](#) and [2](#) of [Definition 2.6](#) with probability at least $2/3$, then there is no randomized polynomial-time algorithm for γ -GapSVP $_p$ unless $\text{NP} \subseteq \text{BPP}$.*

Proof. [Items 1](#) and [3](#) follow immediately by combining [Theorems 2.9](#) and [2.10](#). Inspection of the proof of [Theorem 2.10](#) shows that for NO instances to be mapped to NO instances, only [Item 1](#) of [Definition 2.6](#) is needed, from which [Item 2](#) of the claim follows. \square

3 Local Density from Reed-Solomon Codes

In this section we show how to obtain locally dense lattices from Reed-Solomon codes with appropriate parameters. More specifically, we show to satisfy [Definition 2.6](#) using a lattice $\mathcal{L} := \mathcal{L}^\perp(H)$ corresponding to a parity-check matrix $H = H_q(k, S)$ from [Definition 2.3](#). (Recall that H is the parity-check matrix of a Reed-Solomon code when $S = \mathbb{F}_q$, and of a generalized Reed-Solomon code for any $S \subseteq \mathbb{F}_q$.)

⁵The condition “ γ -GapSVP $_p$ is not in RP” is a slight abuse of notation, since γ -GapSVP $_p$ for $\gamma > 1$ is a promise problem rather than a language. However, the definition of RP can naturally be extended to encompass promise problems, which is the intended meaning here.

The overall structure of the argument is as follows. First, in [Section 3.1](#) we give a lower bound of $\lambda_1^{(p)}(\mathcal{L}) \geq (2k)^{1/p}$, which corresponds to [Item 1](#) of [Definition 2.6](#), by using the connection between power sums and symmetric polynomials (see [Section 2.3](#)). Then, in [Section 3.2](#) we use the upper bound $\det(\mathcal{L}) \leq q^k$ from [Lemma 2.2](#) and the pigeonhole principle to show that there exists a lattice coset with many short (binary) vectors, and in fact a suitably sampled random coset has this property with good probability. Finally, in [Section 3.3](#) we set parameters and use [Theorem 2.7](#) to satisfy [Item 2](#) of [Definition 2.6](#) with good probability.

3.1 Minimum Distance

The following theorem says that for any $k \leq |S|/2$, the ℓ_1 minimum distance of $\mathcal{L} = \mathcal{L}^\perp(H)$ for $H = H_q(k, S)$ is at least $2k$. The theorem and proof are very similar to one that lower bounds the minimum distance of Craig lattices, as given in [\[Cra\]](#) and [\[CS99, Chapter 8, Theorem 7\]](#).

Note that a weaker bound of $\lambda_1^{(1)}(\mathcal{L}) \geq k+1$ (for any $k < q$) follows trivially from the minimum Hamming distance $k+1$ of the (generalized Reed-Solomon) code having parity-check matrix H . However, this bound is not strong enough for the rest of the local-density argument below, which requires $\lambda_1^{(1)}(\mathcal{L}) \geq (1 + \Omega(1))k$.

Theorem 3.1. *Let q be a prime, let $S \subseteq \mathbb{F}_q$, let $k \leq |S|/2$ be a positive integer, and let $H := H_q(k, S) \in \mathbb{F}_q^{k \times S}$ be the matrix from [Definition 2.3](#). Then $\mathcal{L} = \mathcal{L}^\perp(H)$ has ℓ_1 minimum distance $\lambda_1^{(1)}(\mathcal{L}) \geq 2k$. As a consequence, for any $p \in [1, \infty)$ the ℓ_p minimum distance satisfies $\lambda_1^{(p)}(\mathcal{L}) \geq (2k)^{1/p}$.*

We point out that the $2^{1/p}$ factor in [Theorem 3.1](#) propagates to the relative-distance bound for local density in [Theorem 3.4](#) below, and then to the GapSVP approximation factor in our main hardness theorem, [Theorem 1.1](#).

Proof. The consequence follows immediately from the fact that $\mathcal{L} \subseteq \mathbb{Z}^S$ and $\|\mathbf{v}\|_p \geq \|\mathbf{v}\|_1^{1/p}$ for all $\mathbf{v} \in \mathbb{Z}^S$. Now consider some arbitrary $\mathbf{x} \in \mathcal{L} \subseteq \mathbb{Z}^S$ for which $\|\mathbf{x}\|_1 < 2k$; we will show that $\mathbf{x} = \mathbf{0}$. Let $\mathbf{x}^+, \mathbf{x}^- \in \mathbb{Z}^S$ be the unique non-negative integer vectors satisfying $\mathbf{x} = \mathbf{x}^+ - \mathbf{x}^-$. Define multisets T^+ and T^- over S that respectively depend on \mathbf{x}^+ and \mathbf{x}^- as follows. For each $s \in S$ with $x_s^+ > 0$ (respectively, $x_s^- > 0$), let T^+ (respectively, T^-) contain s with multiplicity x_s^+ (respectively, x_s^-).⁶

Note that $|T^+| + |T^-| = \|\mathbf{x}\|_1 < 2k$. Because $H\mathbf{x} = H(\mathbf{x}^+ - \mathbf{x}^-) = \mathbf{0} \in \mathbb{F}_q^k$, by definition of H we have that $p_i(T^+) = p_i(T^-)$ for all $i \in [k]$ (where recall that p_i denotes the i th power sum). Because $k \leq |S|/2 \leq q/2$, by [Proposition 2.5](#) it follows that $T^+ = T^-$. Since $T^+ \cap T^- = \emptyset$ by construction, we must have $T^+ = T^- = \emptyset$, and hence $\mathbf{x} = \mathbf{0}$, as desired. \square

The following lemma (which is well known in other forms) shows that the lower bound $\lambda_1^{(p)}(\mathcal{L}^\perp(H)) \geq (2k)^{1/p}$ from [Theorem 3.1](#) is in fact an equality under mild conditions on the parameters, by giving an explicit lattice coset that has multiple short vectors. However, because it proves only that the number of such vectors is polynomial in the dimension, it is insufficient to establish local density.

Lemma 3.2. *Let q be a prime, let k be a positive integer that divides $q-1$, and let $H := H_q(k, S) \in \mathbb{F}_q^{k \times S}$ where $\mathbb{F}_q^* \subseteq S \subseteq \mathbb{F}_q$. Then for $\mathbf{u} := (k, 0, \dots, 0) \in \mathbb{F}_q^k$, the lattice coset $\mathcal{L}_\mathbf{u}^\perp(H) = \{\mathbf{x} \in \mathbb{Z}^S : H\mathbf{x} = \mathbf{u}\}$ contains $(q-1)/k$ binary vectors of Hamming weight k and pairwise disjoint support. As a consequence, when $k < q-1$, we have $\lambda_1^{(p)}(\mathcal{L}^\perp(H)) = (2k)^{1/p}$ for any $p \in [1, \infty)$.*

Proof. Let G be the order- k subgroup of the (cyclic, multiplicative) group \mathbb{F}_q^* , i.e., the subgroup of the k th roots of unity. Then the binary indicator vectors $\mathbf{x}_C \in \{0, 1\}^S$ of each of the $(q-1)/k$ pairwise disjoint cosets $C = cG$ of G all belong to the coset $\mathcal{L}_\mathbf{u}^\perp(H)$. This is simply because for any such coset, the 0th power sum of its elements is k , and the i th power sum for $0 < i < k$ is zero; this can be seen by Newton's identities and the fact that the root polynomial of C is $f_C(x) = \prod_{c \in C} (x - c) = x^k - r_C$, where $r_C = c^k$ for every $c \in C$. Finally, when $k < q-1$, there is more than one such vector \mathbf{x}_C , and the differences between distinct pairs of them are lattice vectors in $\{0, \pm 1\}^S$ of Hamming weight $2k$, and hence ℓ_p norm $(2k)^{1/p}$. \square

⁶For example, if $S = \{0, 1, 2, 3, 4\} = \mathbb{F}_q$ and $\mathbf{x} = (x_0, x_1, x_2, x_3, x_4)^t = (1, -2, 0, 1, 0)^t$, then $\mathbf{x}^+ = (1, 0, 0, 1, 0)$, $\mathbf{x}^- = (0, 2, 0, 0, 0)$, and accordingly $T^+ = \{0, 3\}$, $T^- = \{1, 4\}$.

3.2 Dense Cosets

Following an approach previously used in [Mic98, Kho04, Mic12] (and implicitly in [Ajt98]), we first show via a pigeonhole argument that a dense lattice coset must exist, and then show how to sample such a coset efficiently (with good probability).

For a prime q , a positive integer k , and a set $S \subseteq \mathbb{F}_q$ of size n (with some arbitrary ordering of its elements), let $H = H_q(k, S) \in \mathbb{F}_q^{k \times n}$ be the parity-check matrix from Definition 2.3. By Lemma 2.2, the lattice $\mathcal{L} = \mathcal{L}^\perp(H) \subseteq \mathbb{Z}^n$ has $\det(\mathcal{L}) \leq q^k$ integer cosets. Recall that $B_{n,h}$ is the set of n -dimensional binary vectors of Hamming weight h , which has cardinality $|B_{n,h}| = \binom{n}{h}$. Therefore, by the pigeonhole principle, there must exist some integer coset $\mathbf{x} + \mathcal{L}$ with $|(\mathbf{x} + \mathcal{L}) \cap B_{n,h}| \geq \binom{n}{h}/q^k$ weight- h binary vectors. In particular, taking $n \approx q$, $h \approx \alpha^p \cdot (2k)$ for some constant $\alpha > 1/2^{1/p}$, and $k = q^\varepsilon$ for a suitable small constant $\varepsilon > 0$ implies the existence of a coset with roughly $q^{(2\alpha^p - 1)k} = q^{\Omega(q^\varepsilon)}$ such vectors. These vectors have ℓ_p norm $h^{1/p} \approx \alpha \cdot (2k)^{1/p}$, whereas by Theorem 3.1 the lattice minimum distance is at least $(2k)^{1/p}$, yielding a local-density relative distance of roughly α .

The following lemma extends the above existential result by showing that something very similar holds for a *uniformly random* shift $\mathbf{x} \in B_{n,h}$: for any $\delta > 0$, the coset $\mathbf{x} + \mathcal{L}$ contains at least $\delta \cdot \binom{n}{h}/q^k$ weight- h binary vectors with probability greater than $1 - \delta$. The proof given below closely follows the structure of the very similar one of [Kho04, Lemma 4.3].

Lemma 3.3. *For a prime q , positive integer k , and set $S \subseteq \mathbb{F}_q$ of size n , let $H = H_q(k, S) \in \mathbb{F}_q^{k \times n}$ be the parity-check matrix from Definition 2.3. There is an efficient randomized algorithm that, for any $\delta \geq 0$, and on input H and any $h \in [n]$, outputs a shift $\mathbf{x} \in B_{n,h}$ such that*

$$\Pr_{\mathbf{x}} \left[|(\mathbf{x} + \mathcal{L}) \cap B_{n,h}| \geq \delta \cdot \binom{n}{h}/q^k \right] > 1 - \delta .$$

Proof. The algorithm simply samples and outputs a uniformly random binary vector $\mathbf{x} \in B_{n,h}$. This is clearly efficient. To show correctness, we will use the syndromes of H . For each $\mathbf{u} \in \mathbb{F}_q^k$, define $K_{\mathbf{u}} := |\{z \in B_{n,h} : Hz = \mathbf{u}\}|$, and define $\mathbf{s} := H\mathbf{x} \in \mathbb{F}_q^k$ to be the syndrome corresponding to \mathbf{x} . So, we need to prove that $K_{\mathbf{s}} \geq \delta \cdot \binom{n}{h}/q^k$ with probability greater than $1 - \delta$. Indeed, we have:

$$\begin{aligned} \Pr_{\mathbf{x}} \left[|(\mathbf{x} + \mathcal{L}) \cap B_{n,h}| < \delta \cdot \binom{n}{h}/q^k \right] &= \Pr_{\mathbf{x}} \left[K_{\mathbf{s}} < \delta \cdot \binom{n}{h}/q^k \right] \\ &= \sum_{\mathbf{u} \in \mathbb{F}_q^k : K_{\mathbf{u}} < \delta \cdot \binom{n}{h}/q^k} \Pr_{\mathbf{x}} [H\mathbf{x} = \mathbf{u}] \\ &= \sum_{\mathbf{u} \in \mathbb{F}_q^k : K_{\mathbf{u}} < \delta \cdot \binom{n}{h}/q^k} \frac{K_{\mathbf{u}}}{\binom{n}{h}} \\ &< \sum_{\mathbf{u} \in \mathbb{F}_q^k : K_{\mathbf{u}} < \delta \cdot \binom{n}{h}/q^k} \frac{\delta}{q^k} \\ &\leq \delta , \end{aligned}$$

where the first inequality uses the fact that the sum is over syndromes \mathbf{u} with $K_{\mathbf{u}} < \delta \cdot \binom{n}{h}/q^k$, and the second inequality uses the fact that there are at most q^k terms in the sum. \square

3.3 The Main Argument

Theorem 3.4 (Locally dense lattices from Reed-Solomon codes). *For any $p \in [1, \infty)$ and constant $\alpha > 1/2^{1/p}$, there exists a randomized polynomial-time algorithm that, given any sufficiently large positive integer r in unary as input, outputs a $(p, \alpha, r, R = \text{poly}(r))$ -locally dense lattice (Definition 2.6) with probability at least $2/3$. Moreover, the algorithm's output satisfies Item 1 of Definition 2.6 with probability 1.*

Proof. The algorithm starts by setting its parameters as follows. It sets $\varepsilon := 2\alpha^p - 1 > 0$, and chooses:

- a poly(r)-bounded integer $k \geq r^{1/(1/2-\delta)}$ for some arbitrary constant $\delta \in (0, 1/2)$, and
- a poly(r)-bounded prime $q \geq k^{3(1+\varepsilon)/\varepsilon}$. (Such a prime q always exists by Bertrand's Postulate.)

The algorithm then computes the components of a $(p, \alpha, r, R = q)$ -locally dense lattice (A, ℓ, \mathbf{x}, T) as follows. It lets:

- $A \in \mathbb{Z}^{q \times q}$ be a basis of $\mathcal{L} := \mathcal{L}^\perp(H)$, where $H = H(k, S) \in \mathbb{F}_q^{k \times q}$ for $S = \mathbb{F}_q$,⁷
- $\ell := 2k$;
- $\mathbf{x} \in B_{q,h}$ be a uniformly random q -dimensional binary vector of Hamming weight $h := \lfloor (1 + \varepsilon)k \rfloor$;
- $T \in \{0, 1\}^{r \times q}$ be chosen by independently setting each of its entries to be 1 with probability $1/(4hr)$, and to be 0 otherwise.

It then outputs (A, \mathbf{x}, ℓ, T) .

We first analyze the algorithm's running time. A suitable prime q can be found in poly(r) time using, e.g., trial division (recall that r is given in unary). The basis A can be computed in deterministic polynomial time from the generating set of column vectors $(B \mid qI_q)$, where B is a basis of $\ker(H) \subseteq \mathbb{F}_q^q$ (lifted to the integers). It is clear that ℓ can be computed in deterministic polynomial time, and that \mathbf{x} and T can be computed in randomized polynomial time. So, the algorithm runs in randomized polynomial time.

It remains to show correctness, i.e., that (A, \mathbf{x}, ℓ, T) satisfies the two conditions in [Definition 2.6](#) with suitable probability over the random choices of \mathbf{x} and T . First, [Item 1](#) is always satisfied, because by [Theorem 3.1](#) we have

$$\lambda_1(\mathcal{L}) \geq (2k)^{1/p} = \ell^{1/p} .$$

In the rest of the proof we consider [Item 2](#) of [Definition 2.6](#). Let $W := (\mathbf{x} + \mathcal{L}) \cap B_{q,h}$. Because

$$\|\mathbf{w}\|_p^p = h \leq (1 + \varepsilon)k = \alpha^p \cdot \ell$$

for each $\mathbf{w} \in W$, we have $W \subseteq V := (\mathbf{x} + \mathcal{L}) \cap \mathcal{B}_p^q(\alpha \cdot \ell^{1/p})$.

By [Lemma 3.3](#), $\Pr_{\mathbf{x}}[|W| \geq \binom{q}{h}/(10q^k)] > 1 - 1/10 = 9/10$. If this event holds, and

$$\frac{\binom{q}{h}}{10q^k} \geq h! \cdot q^{240r\sqrt{h}} , \tag{12}$$

then by [Theorem 2.7](#) we have $\{0, 1\}^r \subseteq T(W) \subseteq T(V)$ with probability at least $1 - 1/10 = 9/10$ (over the choice of T). So, it suffices to show that the condition in [Equation \(12\)](#) holds for all sufficiently large k , and hence for all sufficiently large r . By taking a union bound over the $1/10$ failure probabilities from [Lemma 3.3](#) and [Theorem 2.7](#), we get that the algorithm's overall success probability is at least $1 - 2/10 > 2/3$ for all sufficiently large r , as needed.

Using the standard bound $\binom{q}{h} \geq (q/h)^h$ for binomial coefficients and that $h \geq (1 + \varepsilon)k - 1$, we have that

$$\frac{\binom{q}{h}}{10q^k} \geq \frac{q^{h-k}}{10h^h} = \Omega\left(\frac{q^{\varepsilon k - 1}}{h^h}\right) . \tag{13}$$

Furthermore, by the choice of k relative to r and $h \leq (1 + \varepsilon)k$, we have that

$$h! \cdot q^{240r\sqrt{h}} \leq h^h \cdot q^{240k^{1/2-\delta}\sqrt{(1+\varepsilon)k}} \leq h^h \cdot q^{o(k)} . \tag{14}$$

So, by combining [Equations \(13\)](#) and [\(14\)](#), in order to establish [Equation \(12\)](#) it suffices to show that $q^{(1-o(1))\varepsilon k} \geq h^{2h}$. By taking logs, this is equivalent to

$$(1 - o(1)) \cdot \varepsilon k \log q \geq 2h \log h . \tag{15}$$

⁷For appropriate parameters, our argument works more generally for any sufficiently large subset $S \subseteq \mathbb{F}_q$, with $R = |S|$; we use $S = \mathbb{F}_q$ for simplicity.

Finally, using that $k^{3(1+\varepsilon)/\varepsilon} \leq q \leq \text{poly}(k)$ and $h \leq (1 + \varepsilon)k$, in order for Equation (15) to hold it suffices to have

$$(1 - o(1)) \cdot \varepsilon k \cdot \frac{3(1 + \varepsilon)}{\varepsilon} \cdot \log k = (3 - o(1)) \cdot (1 + \varepsilon) \cdot k \log k \geq 2(1 + \varepsilon) \cdot k \log k + O(k),$$

which indeed holds for all sufficiently large k , as needed. \square

We emphasize that Theorem 3.4 uses randomness only to sample \mathbf{x} and T . As an immediate corollary, we obtain our main hardness result, Theorem 1.1—which, to recall, asserts that for all constants $p \in [1, \infty)$ and $\gamma < 2^{1/p}$, there is no polynomial-time algorithm for γ -GapSVP $_p$ unless $\text{NP} \subseteq \text{RP}$.

Proof of Theorem 1.1. Combine Item 2 of Corollary 2.11 with Theorem 3.4. \square

4 Efficient Decoding Near Minkowski’s Bound

In this section, we show that a recent result of Mook and Peikert [MP22], which builds on work of Guruswami and Sudan [GS98] and Koetter and Vardy [KV03] on list-decoding Reed-Solomon codes, yields a polynomial-time algorithm for decoding lattices $\mathcal{L} = \mathcal{L}^\perp(H)$ with $H = H_q(k, \mathbb{F}_q)$ up to distance $\Theta(\sqrt{k})$. We additionally observe that by choosing $k = \Theta(q/\log q)$, such lattices are asymptotically nearly tight with Minkowski’s bound (Equation (5)). Putting these observations together, we obtain an efficient algorithm for decoding to a distance within a $O(\sqrt{\log q})$ factor of Minkowski’s bound (here $q = n$ is the lattice rank and dimension).

4.1 Construction and Algorithm

Define the additive quotient group $\mathbb{R}_q := \mathbb{R}/(q\mathbb{Z})$ and the Euclidean norm of any $\hat{\mathbf{y}} \in \mathbb{R}_q^n$ as

$$\|\hat{\mathbf{y}}\| := \min\{\|\mathbf{y}\| : \mathbf{y} \in \hat{\mathbf{y}} + q\mathbb{Z}^n\}. \quad (16)$$

Equivalently, $\|\hat{\mathbf{y}}\|$ is the standard \mathbb{R}^n Euclidean norm of the unique real vector $\mathbf{y} \equiv \hat{\mathbf{y}} \pmod{q\mathbb{Z}^n}$ having coordinates in $[-q/2, q/2)$. In additive arithmetic that mixes elements of \mathbb{F}_q and \mathbb{R}_q , we implicitly ‘lift’ the former to the latter in the natural way.

We again use the fact that for evaluation set $S = \mathbb{F}_q$, the matrix $H = H_q(k, \mathbb{F}_q)$ defined in Equation (7) is a parity-check matrix of the Reed-Solomon code $\text{RS}_q[q - k, \mathbb{F}_q]$, and therefore $\mathcal{L}^\perp(H_q(k, \mathbb{F}_q)) = \text{RS}_q[q - k, \mathbb{F}_q] + q\mathbb{Z}^q$. This view lets us take advantage of the decoding algorithm from the following theorem of [MP22], which gives an efficient (list) decoder in the ℓ_2 norm for Reed-Solomon codes.⁸

Proposition 4.1 ([MP22, Algorithm 1 and Theorem 3.4]). *Let q be a prime, $S \subseteq \mathbb{F}_q$ be an evaluation set of size $n := |S|$, $k \leq n$ be a nonnegative integer, and $\varepsilon > 0$. There is a deterministic algorithm that, on input q, S, k, ε , and a vector $\hat{\mathbf{y}} \in \mathbb{R}_q^n$, outputs all codewords $\mathbf{c} \in \text{RS}_q[n - k, S]$ such that $\|\hat{\mathbf{y}} - \mathbf{c}\|^2 \leq (1 - \varepsilon)(k + 1)/2$, in time polynomial in $n, \log q$, and $1/\varepsilon$.⁹*

The following corollary, which is the main result of this section, says that by taking $S = \mathbb{F}_q$ and $k = \Theta(q/\log q)$, (1) the root Hermite factor of $\mathcal{L}^\perp(H)$ is within an $O(\sqrt{\log q})$ factor of Minkowski’s bound (Equation (5)), and (2) it is possible to efficiently decode this lattice to a distance of $\Omega(\sqrt{k}) = \Omega(\sqrt{q/\log q})$, which is again within an $O(\sqrt{\log q})$ factor of Minkowski’s bound.

⁸In fact, the cited result from [MP22] is more general, giving a decoder for \mathbb{F}_p -subfield subcodes of Reed-Solomon codes over finite fields of order $q = p^r$, for a prime p . Here we need only the special case where the Reed-Solomon code is over a prime field (i.e., where $r = 1$). On the other hand, we note that if Proposition 4.1 were extended to handle generalized Reed-Solomon codes, then we would get a corresponding strengthening of Corollary 4.2 for decoding lattices $\mathcal{L}^\perp(H_q(k, S))$ with general S , not just $S = \mathbb{F}_q$.

⁹Formally, the runtimes of the decoding algorithms in Proposition 4.1 and Corollary 4.2 additionally depend on the lengths of the respective ‘received words’ $\hat{\mathbf{y}}$ and \mathbf{y} that they take as input, which must be specified to finite precision. However, for simplicity we describe the algorithms in the ‘Real RAM model,’ while noting that their runtime dependence on the encoding lengths of $\hat{\mathbf{y}}, \mathbf{y}$ is polynomial.

We remark that by setting $\varepsilon \leq 1/(k+1)$ in [Corollary 4.2](#), we get efficient decoding to a distance at least $\sqrt{k/2}$ but less than $\sqrt{(k+1)/2}$, which is slightly more than half the lower bound of $\sqrt{2k}$ on the minimum Euclidean distance of the lattice ([Theorem 3.1](#)). Recall that this lower bound is tight when k is a proper divisor of $q-1$ (see [Lemma 3.2](#)), so with this parameterization we get efficient *list* decoding (i.e., the algorithm may return more than one lattice vector) slightly beyond the unique-decoding bound of half the minimum distance.

Corollary 4.2 (Efficient decoding near Minkowski’s bound). *Let $H = H_q(k, \mathbb{F}_q)$ for a prime q and $k := \lfloor q/(2 \log q) \rfloor \leq q/2$, where all logarithms are base two. Then for $\mathcal{L} := \mathcal{L}^\perp(H) \subseteq \mathbb{Z}^q$:*

1. $\sqrt{q/\log q - 2} \leq \sqrt{2k} \leq \lambda_1(\mathcal{L}) \leq \sqrt{q} \cdot \det(\mathcal{L})^{1/q} \leq \sqrt{2q}$.
2. For any $\varepsilon > 1/\text{poly}(q)$, there is an algorithm that, on input q and a vector $\mathbf{y} \in \mathbb{R}^q$, outputs all lattice vectors $\mathbf{v} \in \mathcal{L}$ satisfying $\|\mathbf{y} - \mathbf{v}\| \leq \sqrt{(1-\varepsilon)(k+1)/2}$ in time polynomial in q .

Proof. For [Item 1](#), we have

$$\sqrt{q/\log q - 2} \leq \sqrt{2k} \leq \lambda_1(\mathcal{L}) \leq \sqrt{q} \cdot \det(\mathcal{L})^{1/q} = \sqrt{q} \cdot q^{k/q} \leq \sqrt{2q}.$$

The first inequality follows from the choice of k , the second inequality is by [Theorem 3.1](#), the third inequality is Minkowski’s bound ([Equation \(5\)](#)), the equality follows from [Lemma 2.2](#) (recall that the rows of H are linearly independent), and the final inequality again follows from the choice of k .¹⁰

The algorithm claimed in [Item 2](#) works as follows. First, it computes k and $\hat{\mathbf{y}} = \mathbf{y} \bmod q\mathbb{Z}^q \in \mathbb{R}_q^q$ from the input q and \mathbf{y} . It then calls the algorithm from [Proposition 4.1](#) on q , $S = \mathbb{F}_q$, k , ε , and $\hat{\mathbf{y}}$, and receives as output zero or more codewords $\mathbf{c} \in \text{RS}_q[q-k, \mathbb{F}_q]$. For each such \mathbf{c} , it outputs the unique vector $\mathbf{v} := \arg \min_{\mathbf{v}' \in \mathbf{c} + q\mathbb{Z}^q} \|\mathbf{y} - \mathbf{v}'\| \in \mathcal{L}$.

The value k and vectors $\hat{\mathbf{y}}$, \mathbf{v} can be computed efficiently (assuming that \mathbf{v} is well defined), so it is clear from [Proposition 4.1](#) that this algorithm runs in time polynomial in q (recall that the dimension $n = q$). It remains to show correctness. First, it is immediate from the definitions that for any $r < q/2$, the function $f(\mathbf{v}) = \mathbf{v} \bmod q\mathbb{Z}^q$ is a bijection from the set of lattice vectors

$$\{\mathbf{v} \in \mathcal{L} : \|\mathbf{y} - \mathbf{v}\| \leq r\},$$

to the set of codewords

$$\{\mathbf{c} \in \text{RS}_q[q-k, \mathbb{F}_q] : \|\hat{\mathbf{y}} - \mathbf{c}\| \leq r\},$$

and that $g(\mathbf{c}) := \arg \min_{\mathbf{v}' \in \mathbf{c} + q\mathbb{Z}^q} \|\mathbf{y} - \mathbf{v}'\|$ is the inverse function of f , i.e., $g = f^{-1}$. Moreover, because $q \geq 2$, we have that the decoding distance r satisfies

$$r := \sqrt{(1-\varepsilon)(k+1)/2} \leq \sqrt{(1-\varepsilon)(q/(2 \log q) + 1)/2} \leq \sqrt{1-\varepsilon} \cdot q/2 < q/2.$$

Because the algorithm from [Proposition 4.1](#) outputs (exactly) $\{\mathbf{c} \in \text{RS}_q[q-k, \mathbb{F}_q] : \|\hat{\mathbf{y}} - \mathbf{c}\| \leq r\}$, it follows that the algorithm described above outputs (exactly) $\{\mathbf{v} \in \mathcal{L} : \|\mathbf{y} - \mathbf{v}\| \leq r\}$, as needed. \square

Remark 4.3. We remark that the main consequence of [Item 1](#) of [Corollary 4.2](#)—namely, an explicit construction of a family of lattices having root Hermite factors within a $O(\sqrt{\log n})$ factor of Minkowski’s bound, obtained via Construction A (where n is the lattice dimension)—only needs a family of codes satisfying milder conditions than what (generalized) Reed-Solomon codes satisfy. Namely, achieving this result only requires a family of linear q -ary codes \mathcal{C} for prime q with block length n , codimension $k = \Theta(n/\log n)$, and minimum distance (in the Hamming metric) $d = \Omega(k)$. The latter is a weaker condition than *maximum distance separability* (MDS), which requires that $d = k + 1$. Indeed, $d = \Omega(k)$ implies that the corresponding Construction-A lattice $\mathcal{C} + q\mathbb{Z}^n$ has an ℓ_2 minimum distance of $\Omega(\min\{\sqrt{k}, q\})$, which is $\Omega(\sqrt{k})$ when $k = O(q^2)$. So, unlike our main hardness result, [Corollary 4.2](#) does not use [Theorem 3.1](#) in any essential way.

¹⁰Analyzing the derivative of $\log(\sqrt{2k}/q^{k/q})$ with respect to k shows that our choice of k is asymptotically optimal for maximizing the root Hermite factor of $\mathcal{L}^\perp(H_q(k, \mathbb{F}_q))$.

Finally, we also note that obtaining a direct analog of [Item 2 of Corollary 4.2](#)—i.e., efficiently decoding to within an $O(\sqrt{\log n})$ factor of Minkowski’s bound on $\mathcal{C} + q\mathbb{Z}^n$ —additionally requires an efficient algorithm for decoding \mathcal{C} to an ℓ_2 distance of $\Omega(\sqrt{k})$, but that this is in turn a weaker requirement than what [Proposition 4.1](#) fulfills.

5 Attempted Derandomization

In this section we describe some of our unsuccessful attempts, and associated barriers, to adapt the *randomized* locally dense lattice construction from [Section 3](#) into a *deterministic* one. More specifically: for some $p \in [1, \infty)$ and $\alpha \in (0, 1)$, our goal is a deterministic algorithm that, given any positive integer r (in unary), outputs a $(p, \alpha, r, R = \text{poly}(r))$ -locally dense lattice in $\text{poly}(r)$ time. Recall from [Definition 2.6](#) that this requires constructing an explicit rank- R lattice and “dense coset” (both integral) containing at least 2^r vectors that have norm at most α times (a known lower bound on) the lattice’s minimum distance.¹¹

As in [Section 3](#), we focus on the case where the lattice has the form $\mathcal{L} = \mathcal{L}^\perp(H) \subseteq \mathbb{Z}^S$, where $H = H_q(k, S)$ for some prime q , positive integer k , and set $S \subseteq \mathbb{F}_q$ of cardinality $n = R = |S|$ (which all may be set by the algorithm). Recall from [Section 3.2](#) that for appropriate parameters, there exists a dense coset $\mathbf{x} + \mathcal{L}$ having $n^{\Omega(k)}$ short vectors (indeed, binary vectors of a particular Hamming weight), and that an appropriate random choice of coset has this property with good probability. To derandomize, we need to deterministically construct a dense integral coset $\mathbf{x} + \mathcal{L}$, or equivalently, its syndrome $\mathbf{u} = H\mathbf{x} \in \mathbb{F}_q^k$. This is the problem that we address in this section. We thank Swastik Kopparty [[Kop20](#)] for explaining its connections to the approaches considered in [Sections 5.1](#) and [5.2](#).

As a first observation, we note that *not every* syndrome (coset) is suitable, and some unsuitable ones are easy to describe. Clearly, the all-zeros syndrome $\mathbf{0} \in \mathbb{F}_q^k$ is unsuitable, because it corresponds to (the zero coset of) the lattice \mathcal{L} itself, which has only one vector shorter than the minimum distance. More generally, by the triangle inequality, for any $\mathbf{x} \in \mathbb{Z}^S$ such that $\|\mathbf{x}\| < (1 - \alpha)\lambda_1(\mathcal{L})$, the lattice coset $\mathbf{x} + \mathcal{L}$ has at most one vector of length at most $\alpha\lambda_1(\mathcal{L})$, namely \mathbf{x} itself (if it is short enough). So, when $\lambda_1(\mathcal{L}) \geq \sqrt{2k}$ (e.g., as implied by [Theorem 3.1](#)), if $\|\mathbf{x}\| < (1 - \alpha)\sqrt{2k}$ then the coset $\mathbf{x} + \mathcal{L}$ has at most one vector of length at most $\alpha\sqrt{2k}$. (A similar argument holds for all other ℓ_p norms.) Stated in terms of syndromes: no “too small” integer linear combination of the columns of H is a suitable syndrome, because it corresponds to shifting the lattice by “too little.”

5.1 Binary Coset Vectors and Reed-Solomon List Decoding Bounds

Optimistically, we might hope to show a close deterministic analog of the probabilistic result from [Lemma 3.3](#), namely, an efficient deterministic construction of a parity-check matrix $H = H_q(k, S)$ and a shift $\mathbf{x} \in \mathbb{Z}^S$ where the coset $\mathbf{x} + \mathcal{L}^\perp(H)$ has sufficiently many *binary* vectors of some specified Hamming weight $h \leq \alpha \cdot (2k) \leq \alpha \cdot \lambda_1^{(1)}(\mathcal{L}^\perp(H))$ for some positive constant $\alpha < 1$.

However, even with a relaxed requirement of $h \leq 2k$, the hoped-for deterministic analog would imply explicit Reed-Solomon list-decoding configurations that go beyond the current state of the art, due to Guruswami and Rudra [[GR05](#)]. While this does not rule out the desired analog entirely, it does present an apparent barrier: achieving the analog would likely require new techniques, and would lead to a significant step forward in our understanding of Reed-Solomon codes.

Lemma 5.1. *There is an efficient deterministic algorithm that, given a prime q , a positive integer $k \leq q$, a set $S \subseteq \mathbb{F}_q$ of cardinality $|S| \geq k$ defining $H = H_q(k, S)$ and $\mathcal{L} = \mathcal{L}^\perp(H) \subseteq \mathbb{Z}^S$, an integer $h \in \{k, \dots, |S|\}$, and any shift $\mathbf{x} \in \mathbb{Z}^S$ (or its syndrome $\mathbf{u} = H\mathbf{x} \in \mathbb{F}_q^k$), outputs a vector $\mathbf{r} \in \mathbb{F}_q^S$ for which there are at least $|\mathbf{x} + \mathcal{L} \cap B_{S,h}|$ codewords in the Reed-Solomon code $\text{RS}_q[h - k + 1, S]$ that each agree with \mathbf{r} in at least h coordinates.*

¹¹Recall that a locally dense lattice must also come with a suitable linear transform T . In this section we do not address the construction of such T , because deterministically constructing a lattice and dense coset is already challenging enough. We hope that such a construction and its analysis will naturally reveal a suitable choice of T as well.

Proof. The algorithm works as follows.

1. Compute the coset's syndrome $\mathbf{u} = H\mathbf{x} \in \mathbb{F}_q^k$.

Note that for each $i \in [k]$, the syndrome's i th coordinate u_i is the i th power sum of the subset $T \subseteq S$ indicated by any binary vector $\mathbf{x}_T \in (\mathbf{x} + \mathcal{L}) \cap B_{S,h}$, i.e., $u_i = (H\mathbf{x})_i = (H\mathbf{x}_T)_i = p_i(T) \in \mathbb{F}_q$. In particular, $h = |T| \equiv u_0 \pmod{q}$ for any such T . (So, if $h \not\equiv u_0 \pmod{q}$, no such T exists and the algorithm may output any vector.)

2. Use Newton's identities (Equation (10)) to compute (from \mathbf{u}) the values $s_i = e_i(T) \in \mathbb{F}_q$ for each $i \in [k]$, which are the first $k \leq h = |T|$ elementary symmetric polynomials of any such T .
3. Finally, construct the polynomial

$$r(x) = \sum_{i \in [k]} (-1)^i \cdot s_i \cdot x^{h-i} \in \mathbb{F}_q[x]$$

and output the vector $\mathbf{r} = (-r(s))_{s \in S} \in \mathbb{F}_q^S$ obtained by evaluating $-r(x)$ at each element of S .

We now show that the algorithm is correct. Let $T \subseteq S$ be the subset indicated by a binary vector $\mathbf{x}_T \in (\mathbf{x} + \mathcal{L}) \cap B_{S,h}$, and let $f_T(x) = \prod_{t \in T} (x - t) \in \mathbb{F}_q[x]$ be the root polynomial of T . As argued above in Item 2, $s_i = e_i(T)$ for each $i \in [k]$, so by Equation (11), we have

$$f_T(x) = r(x) + g_T(x)$$

for some $g_T(x)$ where $\deg(g_T) \leq h - k$ (or g_T is identically zero). Because $f_T(t) = 0$ for every $t \in T$, we have $g_T(t) = -r(t)$ for every $t \in T$, i.e., the Reed-Solomon codeword in $\text{RS}[h - k + 1, S]$ defined by g_T agrees with \mathbf{r} in at least $|T| = h$ positions. Finally, the claim follows by noting that each distinct binary vector $\mathbf{x}_T \in (\mathbf{x} + \mathcal{L}) \cap B_{S,h}$ yields a distinct subset $T \subseteq S$, which yields a distinct root polynomial f_T and thereby a distinct polynomial $g_T = f_T - r$. \square

For an analog of Lemma 3.3, we want an explicit (i.e., deterministically and efficiently constructible) lattice $\mathcal{L} = \mathcal{L}^\perp(H)$ and shift \mathbf{x} such that $|(\mathbf{x} + \mathcal{L}) \cap B_{S,h}|$ is super-polynomial (in $|S|$), for some Hamming weight $h \leq 2k \leq \lambda_1^{(1)}(\mathcal{L})$. By Lemma 5.1, this implies an explicit received word and Reed-Solomon code of dimension $h - k + 1$ for which super-polynomially many codewords agree with the received word in at least h coordinates. This is an agreement-to-dimension ratio of $h/(h - k + 1) = 1 + (k - 1)/(h - k + 1) \geq 2 - O(1/k)$. However, the state of the art for explicit Reed-Solomon list-decoding configurations [GR05, Corollary 2] requires a ratio of $2 - \Omega(1)$ in order to get a super-polynomial list size.

Remark 5.2. It is not clear whether a useful converse of Lemma 5.1 holds, i.e., whether an explicit 'bad' list-decoding configuration for Reed-Solomon codes would yield a corresponding explicit locally dense lattice. This is because the difference between the received word and a nearby codeword may not correspond to (the evaluations of) the root polynomial of the positions at which they agree, but is only divisible by it.

Remark 5.3. Everything in the proof of Lemma 5.1 generalizes from binary vectors to *non-negative* integer vectors of ℓ_1 norm h and their corresponding *multisets* T , except for the final step: the codeword corresponding to $g_T(x)$ is guaranteed to agree with \mathbf{r} only at each *distinct* $t \in T$, and the number of these is only the *Hamming* weight (not the ℓ_1 norm h) of the integer vector. It appears that this gap can be overcome (restoring agreement h) by instead using univariate *multiplicity codes*, which are generalizations of Reed-Solomon where the coordinates correspond to 'multiple evaluations' of a polynomial *and its derivatives*; see [Kop14]. For such codes, we do not know whether any explicit list-decoding configurations with substantially different parameters from Reed-Solomon codes are known.

By contrast, Lemma 5.1 does not seem to generalize easily to *arbitrary* integer vectors that may have negative entries (e.g., ternary $\{0, \pm 1\}$ -vectors of a given Hamming weight). So, considering negative entries may provide a route to constructing locally dense lattices without having to overcome the Reed-Solomon list-decoding barriers.

5.2 A Fourier-Analytic Approach

We next outline an approach for using Fourier analysis over \mathbb{F}_q^k for prime q and $k \ll q$ to show that an explicit coset $\mathbf{x} + \mathcal{L}^\perp(H)$ for $H = H_q(k, \mathbb{F}_q)$ contains many vectors of ℓ_1 norm at most $h \leq \alpha \cdot (2k) \leq \alpha \cdot \lambda_1^{(1)}(\mathcal{L}^\perp(H))$ for some integer h , where $\alpha \in (0, 1)$ is a constant as above.

The approach uses techniques similar to those used by Cheng and Wan [CW04] and Guruswami and Rudra [GR05] for constructing explicit Hamming balls containing many Reed-Solomon codewords (i.e., for constructing explicit “bad” configurations for list-decoding Reed-Solomon codes), and by Cheng and Wan [CW09] for showing deterministic hardness of the minimum distance problem on codes.

Specifically, we perform a Fourier-analytic calculation similar to ones appearing in those works in an attempt to lower bound the number of short vectors $\mathbf{x} \in \mathbb{F}_q^k$ satisfying $H\mathbf{x} = \mathbf{s}$ for some syndrome $\mathbf{s} \in \mathbb{F}_q^k$ (which is equivalent to lower bounding the number of short vectors in the lattice coset $\mathcal{L}_\mathbf{s}^\perp(H)$). The calculation starts by establishing a correspondence between polynomials in $\mathbb{F}_q[x]$ of degree less than k and the Fourier coefficients of a certain indicator function 1_W . It then uses a combination of the Weil bound (Equation (25)) and the triangle inequality to upper bound the total contribution of the terms corresponding to non-constant polynomials in the Fourier expansion of the h -fold convolution $1_W^{(*h)}$ of 1_W , which essentially counts the number of non-negative integer vectors of ℓ_1 norm h with a given syndrome (see Equation (26)).

Unfortunately, this calculation (just barely) fails to show anything useful for our purposes. The rough idea of why is as follows. The argument seeks to show that the sum of magnitudes (ℓ_1 norm) of all the Fourier coefficients of $1_W^{(*h)}$ corresponding to non-constant polynomials is less than the zero coefficient of $1_W^{(*h)}$ (which is real-valued, positive, and corresponds to the zero polynomial). The Weil bound shows that there is a multiplicative gap of about \sqrt{q} between the value of the zero coefficient and the magnitude of any non-constant coefficient of 1_W , and this gap is amplified to about $q^{h/2}$ for the corresponding Fourier coefficients of $1_W^{(*h)}$. However, because there are about q^k non-constant polynomials in $\mathbb{F}_q[x]$ of degree less than k , we need to take $h \geq 2k$ in order for the zero coefficient of $1_W^{(*h)}$ to be large enough, which conflicts with the requirement that $h < 2k$.

Because the calculation is not directly useful for our purposes, and due to its similarity to those in prior work, we defer it to Appendix B. There we also discuss some approaches for making the argument work.

5.3 Point-Counting Proxies

In this section we outline a possible approach toward proving local density for a deterministic version of our construction, using “smooth” functions that serve as proxies for point-counting functions. The hope is to use these functions to prove a lower bound on the number of sufficiently short points in *any* coset $\mathbf{x} + \mathcal{L}^\perp(H)$ meeting some easy-to-ensure conditions. The approach is centered on analytical techniques originally due to [MO90, EOR91], which were successfully used on the integer lattice \mathbb{Z}^n for complexity-theoretic purposes in [AS18, BP20]. In the present context of Reed-Solomon lattices, our attempts have been inconclusive: while we did not manage to prove what we are seeking, we also do not see any inherent barrier, and further effort may lead to a favorable outcome.

For any $p \geq 1$, ℓ_p norm bound $r \geq 0$, and lattice coset $\mathbf{x} + \mathcal{L} \subset \mathbb{R}^n$, define

$$N_p(r; \mathbf{x} + \mathcal{L}) := |(\mathbf{x} + \mathcal{L}) \cap \mathcal{B}_p^n(r)| \tag{17}$$

to be the number of points in $\mathbf{x} + \mathcal{L}$ of ℓ_p norm at most r . Similarly, for any $\tau > 0$, define the function

$$\Theta_p(\tau; \mathbf{x} + \mathcal{L}) := \sum_{\mathbf{v} \in \mathbf{x} + \mathcal{L}} \exp(-\tau \|\mathbf{v}\|_p^p). \tag{18}$$

In what follows, for brevity we often omit the coset $\mathbf{x} + \mathcal{L}$ when it is fixed and clear from context.

First, for any $\tau > 0$ we immediately have

$$N_p(r) \leq \exp(\tau \cdot r^p) \cdot \Theta_p(\tau), \tag{19}$$

because each point counted by $N_p(r)$ contributes at least $\exp(-\tau \cdot r^p)$ to $\Theta_p(\tau)$, and all other points contribute positively. The work of [MO90, EOR91] showed that for any coset of the *integer* lattice $\mathcal{L} = \mathbb{Z}^n$, the upper bound from Equation (19) is fairly tight (up to $\exp(O(\sqrt{n}))$ factors) for suitably chosen τ , namely, the one for which $\mu_p(\tau; \mathbf{x} + \mathcal{L}) = r^p$, where

$$\mu_p(\tau; \mathbf{x} + \mathcal{L}) := \mathbb{E}_{\mathbf{v} \sim D_p(\tau; \mathbf{x} + \mathcal{L})} [\|\mathbf{v}\|_p^p] \quad (20)$$

and $D_p(\tau)$ is the probability distribution supported on $\mathbf{x} + \mathcal{L}$ that assigns probability $\exp(-\tau \|\mathbf{v}\|_p^p) / \Theta_p(\tau)$ to each \mathbf{v} . In other words, $\mu_p(\tau)$ is the p th moment of the ℓ_p norm of a sample from $D_p(\tau)$, and when τ is chosen to make this moment equal r^p , Equation (19) is fairly tight for any coset of the integer lattice.

Optimistically, if we could show that Equation (19) is fairly tight for suitable τ (as a function of r) and every coset of our special lattices, then we would have

$$N_p(r) \approx \exp(\tau \cdot r^p) \cdot \Theta_p(\tau) = \exp(\tau \cdot \alpha^p \cdot 2k) \cdot \Theta_p(\tau) ,$$

where $r = \alpha \cdot (2k)^{1/p} \leq \alpha \cdot \lambda_1^{(p)}(\mathcal{L}^\perp(H))$ for some constant $\alpha < 1$ is our distance of interest for local density, and τ is determined by this r . Furthermore, if we could also show that, say, $\Theta_p(\tau) \geq \exp(-\tau k)$ for this τ and a suitable choice of coset (e.g., by shifting the lattice by some \mathbf{x} with $\|\mathbf{x}\|_p^p = k$), then for any constant $\alpha > 1/2^{1/p}$ we could conclude that there are $\exp(\Omega(\tau k))$ sufficiently short points in the coset. Finally, if we could show that $\tau \geq k^{c-1}$ for some constant $c > 0$, this would yield a subexponential number of short coset vectors, as needed.

An approach. Toward this end, we observe that several parts of the analysis from [MO90, EOR91] immediately generalize to cosets of arbitrary lattices, not just the integer lattice \mathbb{Z}^n . In particular, for any lattice coset and any $\tau > 0$, by a routine calculation we have

$$\frac{\partial}{\partial \tau} \ln \Theta_p(\tau) = -\mu_p(\tau) < 0 , \quad (21)$$

$$\frac{\partial^2}{\partial \tau^2} \ln \Theta_p(\tau) = \mathbb{E}_{\mathbf{v} \sim D_p(\tau)} [\|\mathbf{v}\|_p^{2p}] - \mu_p(\tau)^2 > 0 . \quad (22)$$

In words, $\ln \Theta_p(\tau)$ is decreasing and concave up, and its second partial derivative (with respect to τ) is the variance of $\|\mathbf{v}\|_p^p$ when \mathbf{v} is sampled from $D_p(\tau)$. Moreover, by an elegant (and non-obvious) argument, for any $\tau, \delta > 0$ we have the lower bound

$$N_p(\mu_p(\tau)^{1/p}) \geq \exp((\tau + \delta) \cdot \mu_p(\tau + 2\delta)) \cdot H_p(\tau, \delta) , \quad (23)$$

where

$$H_p(\tau, \delta) := \Theta_p(\tau + \delta) - \exp(-\delta \mu_p(\tau)) \cdot \Theta_p(\tau) - \exp(\delta \mu_p(\tau + 2\delta)) \cdot \Theta_p(\tau + 2\delta) . \quad (24)$$

Using the Taylor series for $\ln \Theta_p(\tau)$ around τ and Equation (21), for any $\delta > 0$ we have

$$\ln \Theta_p(\tau) \leq \ln \Theta_p(\tau + \delta) + \delta \mu_p(\tau) - \frac{\delta^2}{2} \inf_{\tau' \in [\tau, \tau + \delta]} \frac{\partial^2}{\partial \tau'^2} \ln \Theta_p(\tau') ,$$

and similarly for $\ln \Theta_p(\tau + 2\delta)$. So, if we could suitably lower-bound the above infimum on the second partial derivative (which, to recall from Equation (22), is positive) for appropriate δ , by taking $\exp(\cdot)$ of both sides and rearranging we could show that, say, $H_p(\tau, \delta) \geq \Theta_p(\tau + \delta)/2$. The work of [MO90, EOR91] actually does this for the integer lattice \mathbb{Z}^n , by exploiting its product structure to show that the infimum is proportional to n , but here the path forward is less clear.

If the above could be achieved, then we would get the lower bound

$$N_p(\mu_p(\tau)^{1/p}) \geq \exp((\tau + \delta) \cdot \mu_p(\tau + 2\delta)) \cdot \Theta_p(\tau + \delta)/2 \approx \exp(\tau \cdot \mu_p(\tau)) \cdot \Theta_p(\tau)/2 ,$$

showing that Equation (19) is fairly tight when $r = \mu_p(\tau)^{1/p}$. As explained above, the final goal would then be to identify an explicit coset for which the right-hand side is at least subexponentially large in the lattice dimension, when τ is set so that $\mu_p(\tau) = \alpha^p \cdot (2k)$ for some constant $\alpha < 1$.

References

- [ABSS93] S. Arora, L. Babai, J. 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. Preliminary version in FOCS 1993. [11](#)
- [Ajt98] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19. 1998. [2](#), [4](#), [13](#)
- [AK11] P. Austrin and S. Khot. A simple deterministic reduction for the gap minimum distance of code problem. *IEEE Trans. Inf. Theory*, 2014. Preliminary version in ICALP 2011. [6](#)
- [AR04] D. Aharonov and O. Regev. Lattice problems in NP cap coNP. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004. [2](#)
- [AS18] D. Aggarwal and N. Stephens-Davidowitz. (Gap/S)ETH hardness of SVP. In *STOC*. 2018. [2](#), [4](#), [19](#)
- [BP20] H. Bennett and C. Peikert. Hardness of bounded distance decoding on lattices in ℓ_p norms. In *CCC*. 2020. [6](#), [7](#), [19](#)
- [BPT22] H. Bennett, C. Peikert, and Y. Tang. Improved hardness of BDD and SVP under Gap-(S)ETH. In *ITCS*. 2022. [2](#), [4](#), [6](#)
- [CN98] J. Cai and A. Nerurkar. Approximating the SVP to within a factor $(1 + 1/\dim^\epsilon)$ is NP-hard under randomized reductions. In *CCC*. 1998. [2](#), [4](#)
- [Cra] M. Craig. Automorphisms of prime cyclotomic lattices. Preprint. [3](#), [5](#), [6](#), [12](#)
- [CS99] J. Conway and N. J. A. Sloane. *Sphere packings, lattices, and groups*. Springer, 1999. [3](#), [6](#), [8](#), [12](#)
- [CW04] Q. Cheng and D. Wan. On the list and bounded distance decodibility of the Reed-Solomon codes (extended abstract). In *FOCS*. 2004. [3](#), [19](#)
- [CW09] Q. Cheng and D. Wan. A deterministic reduction for the gap minimum distance problem. *IEEE Trans. Inf. Theory*, 58(11):6935–6941, 2012. Preliminary version in STOC 2009. [3](#), [6](#), [19](#), [25](#)
- [DMS99] I. Dumer, D. Micciancio, and M. Sudan. Hardness of approximating the minimum distance of a linear code. *IEEE Trans. Inf. Theory*, 2003. Preliminary version in FOCS 1999. [6](#)
- [DP19] L. Ducas and C. Pierrot. Polynomial time bounded distance decoding near Minkowski’s bound in discrete logarithm lattices. *Des. Codes Cryptogr.*, 87(8):1737–1748, 2019. [4](#), [6](#)
- [DvW21] L. Ducas and W. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. Cryptology ePrint Archive, Report 2021/1332, 2021. <https://ia.cr/2021/1332>. [4](#)
- [EOR91] N. D. Elkies, A. M. Odlyzko, and J. A. Rush. On the packing densities of superballs and other bodies. *Inventiones mathematicae*, 105:613–639, December 1991. [19](#), [20](#)
- [GG98] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. Preliminary version in STOC 1998. [2](#)
- [GP12] E. Grigorescu and C. Peikert. List-decoding Barnes-Wall lattices. *Comput. Complex.*, 26(2):365–392, 2017. Preliminary version in CCC 2012. [4](#), [6](#)
- [GR05] V. Guruswami and A. Rudra. Limits to list decoding Reed-Solomon codes. *IEEE Trans. Inf. Theory*, 52(8):3642–3649, 2006. Preliminary version in STOC 2005. [3](#), [17](#), [18](#), [19](#)

- [GS98] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometry codes. *IEEE Trans. Inf. Theory*, 45(6):1757–1767, 1999. Preliminary version in FOCS 1998. 15
- [Hal] J. I. Hall. Notes on coding theory. Available at <https://users.math.msu.edu/users/halljoc/classes/CODENOTES/CODING-NOTES.HTML>. 5, 9
- [HR07] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory Comput.*, 8(1):513–531, 2012. Preliminary version in STOC 2007. 2, 3, 6, 7
- [Kho03] S. Khot. Hardness of approximating the shortest vector problem in high ℓ_p norms. *J. Comput. Syst. Sci.*, 72(2):206–219, 2006. Preliminary version in FOCS 2003. 2, 4
- [Kho04] S. Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in FOCS 2004. 2, 3, 4, 6, 13
- [Kop14] S. Kopparty. Some remarks on multiplicity codes. In A. Barg, editor, *Discrete Geometry and Algebraic Combinatorics*, Contemporary Mathematics. American Mathematical Society, 2014. 18
- [Kop20] S. Kopparty. Personal communication, 2020. 7, 17
- [Kow10] E. Kowalski. Exponential sums over finite fields: elementary methods. <https://people.math.ethz.ch/~kowalski/exponential-sums-elementary.pdf>, 2010. 24
- [KV03] R. Koetter and A. Vardy. Algebraic soft-decision decoding of Reed-Solomon codes. *IEEE Trans. Inf. Theory*, 49(11):2809–2825, 2003. 15
- [LLM06] Y. Liu, V. Lyubashevsky, and D. Micciancio. On bounded distance decoding for general lattices. In *RANDOM*, pages 450–461. 2006. 6
- [Mea92] D. G. Mead. Newton’s identities. *The American Mathematical Monthly*, 99(8):749, October 1992. 9
- [Mic98] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000. Preliminary version in FOCS 1998. 2, 3, 4, 7, 10, 13
- [Mic12] D. Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory Comput.*, 8(1):487–512, 2012. 2, 3, 4, 6, 11, 13
- [Mic14] D. Micciancio. Locally dense codes. In *CCC*. 2014. 6
- [MN08] D. Micciancio and A. Nicolosi. Efficient bounded distance decoders for Barnes-Wall lattices. In *ISIT*, pages 2484–2488. IEEE, 2008. 4, 6
- [MO90] J. E. Mazo and A. M. Odlyzko. Lattice points in high-dimensional spheres. *Monatshefte für Mathematik*, 110(1):47–61, March 1990. 19, 20
- [MP22] E. Mook and C. Peikert. Lattice (list) decoding near Minkowski’s inequality. *IEEE Trans. Inf. Theory*, 68(2):863–870, 2022. 3, 4, 6, 15
- [O’D14] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. ISBN 978-1-10-703832-5. 24
- [Pei07] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008. Preliminary version in CCC 2007. 2
- [Pei16] C. Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016. 2

- [RR06] O. Regev and R. Rosen. Lattice problems and norm embeddings. In *STOC*. 2006. 7
- [Saw14] W. Sawin. How tight is the Weil bound for this exponential sum? MathOverflow, 2014. <https://mathoverflow.net/q/182599>. 26
- [vEB81] P. van Emde Boas. Another NP-complete partition problem and the complexity of computing short vectors in a lattice. Technical Report, 1981. Available at <https://staff.fnwi.uva.nl/p.vanemdeboas/vectors/mi8104c.html>. 2

A Proof of Theorem 2.10

Here we restate and prove Theorem 2.10.

Theorem 2.10. *Let $p \geq 1$, r and n be positive integers, $\alpha > 0$ be a constant, and γ, γ' be constants satisfying*

$$1/\alpha > \gamma' \geq 1 \text{ and } \gamma \geq \gamma' \cdot \left(\frac{1}{1 - (\alpha\gamma')^p} \right)^{1/p}.$$

There is a deterministic polynomial-time algorithm that, given a γ -GapCVP'_p instance (B, \mathbf{t}, s) of rank r and a (p, α, r, R) -locally dense lattice (A, ℓ, \mathbf{x}, T) as input, outputs a γ' -GapSVP'_p instance (B', s') of rank $R + 1$ which is a YES (respectively, NO) instance if (B, \mathbf{t}, s) is a YES (resp., NO) instance.

Proof. Let $B \in \mathbb{Z}^{d \times r}$, $\mathbf{t} \in \mathbb{Z}^d$, $s > 0$ be the input instance of γ -GapCVP'_p. Assume without loss of generality that $\mathbf{t} \neq \mathbf{0}$ (if $\mathbf{t} = \mathbf{0}$, simply output an arbitrary YES instance of γ' -GapSVP'_p). The reduction outputs (B', s') , where

$$B' := \begin{pmatrix} BTA & B\mathbf{T}\mathbf{x} - \mathbf{t} \\ \beta A & \beta \mathbf{x} \end{pmatrix} \text{ with } \beta := \frac{\gamma' s}{\ell^{1/p} \cdot (1 - (\alpha\gamma')^p)^{1/p}} \neq 0,$$

$$s' := (s^p + \alpha^p \beta^p \ell)^{1/p}.$$

It is clear that the reduction runs in polynomial time. To show correctness, we start by showing that the columns of B' are linearly independent, i.e., that B' is a lattice basis, and hence the lattice it generates has rank $R + 1$. Recall that $A \in \mathbb{Z}^{n \times R}$ (for some n) is a lattice basis, so it has linearly independent columns. Let $\mathbf{u} \in \mathbb{R}^{R+1} \setminus \{\mathbf{0}\}$. If $(\beta A \mid \beta \mathbf{x}) \cdot \mathbf{u} \neq \mathbf{0}$ then $B' \mathbf{u} \neq \mathbf{0}$. Otherwise, $(A \mid \mathbf{x}) \cdot \mathbf{u} = \mathbf{0}$ and so $u_{n+1} \neq 0$ by linear independence of A 's columns. Then

$$(BTA \mid B\mathbf{T}\mathbf{x} - \mathbf{t}) \cdot \mathbf{u} = BT(A \mid \mathbf{x}) \cdot \mathbf{u} - u_{n+1} \cdot \mathbf{t} = -u_{n+1} \cdot \mathbf{t} \neq \mathbf{0},$$

because $\mathbf{t} \neq \mathbf{0}$ by assumption. Hence $B' \mathbf{u} \neq \mathbf{0}$, as needed.

We now show the main claim about (B', s') . First, suppose that the input GapCVP'_p instance (B, \mathbf{t}, s) is a YES instance. Then by Definition 2.8 there exists an $\mathbf{c} \in \{0, 1\}^r$ such that $\|B\mathbf{c} - \mathbf{t}\|_p \leq s$. Moreover, by Item 2 of Definition 2.6, there exists $\mathbf{v} = \mathbf{x} + A\mathbf{z} \in \mathbf{x} + \mathcal{L}(A)$ for some $\mathbf{z} \in \mathbb{Z}^R$ such that $\|\mathbf{v}\|_p^p \leq \alpha^p \cdot \ell$ and $T\mathbf{v} = \mathbf{c}$. For such a \mathbf{z} , we have $T(\mathbf{x} + A\mathbf{z}) = T\mathbf{v} = \mathbf{c}$ and therefore

$$\begin{aligned} \lambda_1^{(p)}(\mathcal{L}(B'))^p &\leq \|B' \cdot (\mathbf{z}, 1)\|_p^p \\ &= \|BTA\mathbf{z} + B\mathbf{T}\mathbf{x} - \mathbf{t}\|_p^p + \beta^p \|A\mathbf{z} + \mathbf{x}\|_p^p \\ &= \|B\mathbf{c} - \mathbf{t}\|_p^p + \beta^p \|A\mathbf{z} + \mathbf{x}\|_p^p \\ &\leq s^p + \alpha^p \beta^p \ell \\ &= (s')^p, \end{aligned}$$

as needed.

Next, suppose that the input GapCVP'_p instance is a NO instance. We will show that $\lambda_1^{(p)}(\mathcal{L}(B')) \geq \gamma' s'$ by showing that $\|B' \cdot (\mathbf{z}, w)\|_p^p > (\gamma' s')^p$ for all $(\mathbf{z}, w) \in \mathbb{Z}^{R+1} \setminus \{\mathbf{0}\}$, separately analyzing the cases $w = 0$ and $w \neq 0$. When $w = 0$, we have that $\mathbf{z} \neq \mathbf{0}$, and so by [Item 1 of Definition 2.6](#),

$$\|B' \cdot (\mathbf{z}, w)\|_p^p \geq \beta^p \cdot \lambda_1^{(p)}(\mathcal{L}(A))^p > \beta^p \cdot \ell.$$

And, since $\beta^p \cdot \ell \cdot (1 - (\alpha\gamma')^p) = (\gamma' s')^p$ by definition, we have (by definition of s')

$$\beta^p \cdot \ell \geq (\gamma' s')^p + (\alpha\beta\gamma')^p \cdot \ell = (\gamma' s')^p.$$

When $w \neq 0$, by the definition of GapCVP'_p NO instances, we have

$$\|B' \cdot (\mathbf{v}, w)\|_p^p \geq \text{dist}_p(w\mathbf{t}, \mathcal{L}(B))^p > (\gamma s)^p.$$

Then, by the definitions of γ , β , and s' , we have

$$(\gamma s)^p \geq (\gamma')^p \cdot \frac{s^p}{1 - (\alpha\gamma')^p} = (\gamma')^p \cdot \left(s^p + \frac{(\alpha\gamma' s)^p}{1 - (\alpha\gamma')^p} \right) = (\gamma')^p \cdot (s^p + \alpha^p \beta^p \ell) = (\gamma' s')^p.$$

The claim follows. \square

B A Fourier-Analytic Approach to Derandomization

In this appendix we perform and discuss the Fourier-analytic calculation described in [Section 5.2](#). We introduce definitions and basic facts briefly, and refer the reader to [\[Kow10\]](#) and [\[O'D14, Chapter 8\]](#) for details. The *additive Fourier characters* of \mathbb{F}_q^k are the homomorphisms from the additive group \mathbb{Z}_q^k of \mathbb{F}_q^k to the multiplicative group of \mathbb{C} . Concretely, they are the q^k functions $\Psi_{\mathbf{u}}(\mathbf{x}) := \omega^{\langle \mathbf{u}, \mathbf{x} \rangle}$ indexed by $\mathbf{u} \in \mathbb{F}_q^k$, where ω is an arbitrary fixed q th primitive root of unity (e.g., $\omega = \exp(2\pi i/q)$). For a function $f : \mathbb{F}_q^k \rightarrow \mathbb{C}$, its \mathbf{u} th Fourier coefficient $\hat{f}(\mathbf{u})$ for $\mathbf{u} \in \mathbb{F}_q^k$ is defined as

$$\hat{f}(\mathbf{u}) := \mathbb{E}_{\mathbf{b} \sim \mathbb{F}_q^k} [f(\mathbf{b}) \overline{\Psi_{\mathbf{u}}(\mathbf{b})}],$$

and the Fourier expansion of f is

$$f(\mathbf{x}) = \sum_{\mathbf{u} \in \mathbb{F}_q^k} \hat{f}(\mathbf{u}) \Psi_{\mathbf{u}}(\mathbf{x}).$$

Let $W = W(k, q) := \{(1, a, a^2, \dots, a^{k-1})^T : a \in \mathbb{F}_q\} \subseteq \mathbb{F}_q^k$ denote the set of columns of the matrix $H_q(k, \mathbb{F}_q)$, and let $1_W : \mathbb{F}_q^k \rightarrow \{0, 1\}$ denote the indicator function of W . The Fourier coefficients of 1_W satisfy

$$\widehat{1}_W(\mathbf{u}) = \mathbb{E}_{\mathbf{b} \sim \mathbb{F}_q^k} [1_W(\mathbf{b}) \overline{\Psi_{\mathbf{u}}(\mathbf{b})}] = \frac{1}{q^k} \sum_{\mathbf{b} \in W} \overline{\Psi_{\mathbf{u}}(\mathbf{b})} = \frac{1}{q^k} \sum_{a \in \mathbb{F}_q} \Psi(-p_{\mathbf{u}}(a)),$$

where $\Psi(x) = \omega^x$ is an additive character of \mathbb{F}_q and $p_{\mathbf{u}}(x) = \sum_{i=0}^{k-1} u_i x^i \in \mathbb{F}_q[x]$ is a polynomial of degree less than $k \leq q$.¹² For non-constant polynomials $p_{\mathbf{u}}(x)$, the Weil bound for additive character sums (see [\[Kow10, Theorem 3.2\]](#)) asserts that

$$q^k \cdot |\widehat{1}_W(\mathbf{u})| = \left| \sum_{a \in \mathbb{F}_q} \Psi(-p_{\mathbf{u}}(a)) \right| \leq (k-2)\sqrt{q}. \quad (25)$$

¹²Note that formally, we are indexing the Fourier coefficients by $\mathbb{F}_q^{[k]}$, i.e., the coordinates of $\mathbf{u} \in \mathbb{F}_q^{[k]}$ are indexed from zero, by the elements of $[k] = \{0, \dots, k-1\}$.

The *convolution* of two functions $f, g : \mathbb{F}_q^k \rightarrow \mathbb{C}$ is the function defined as $(f * g)(\mathbf{x}) := \mathbb{E}_{\mathbf{y} \sim \mathbb{F}_q^k} [f(\mathbf{x} - \mathbf{y})g(\mathbf{y})]$. Importantly, Fourier coefficients are multiplicative under convolution:

$$\widehat{(f * g)}(\mathbf{u}) = \widehat{f}(\mathbf{u})\widehat{g}(\mathbf{u}) .$$

Furthermore, the h -fold convolution of $1_W(\mathbf{x})$ with itself, denoted $1_W^{(*h)}(\mathbf{x})$, counts, up to a normalization factor of q^{hk} , the number of sequences $(\mathbf{w}_1, \dots, \mathbf{w}_h) \in W^h$ such that $\sum_{i=1}^h \mathbf{w}_i = \mathbf{x}$.¹³

Let $U_0 := \{(u_0, 0, \dots, 0) : u_0 \in \mathbb{F}_q\}$ denote the set of coefficient vectors corresponding to constant polynomials, i.e., corresponding to polynomials $p_{\mathbf{u}}(x) = u_0$ for $u_0 \in \mathbb{F}_q$. Putting everything together, we get that

$$\begin{aligned} \frac{1}{q^{hk}} \cdot \left| \{(\mathbf{w}_1, \dots, \mathbf{w}_h) \in W^h : \sum_{i=1}^h \mathbf{w}_i = \mathbf{x}\} \right| &= 1_W^{(*h)}(\mathbf{x}) \\ &= \sum_{\mathbf{u} \in \mathbb{F}_q^k} \widehat{1_W^{(*h)}}(\mathbf{u}) \cdot \Psi_{\mathbf{u}}(\mathbf{x}) \\ &= \sum_{\mathbf{u} \in \mathbb{F}_q^k} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{x}) \\ &= \sum_{\mathbf{u} \in U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{x}) + \sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{x}) \\ &= \frac{1}{q^{hk}} \sum_{u_0 \in \mathbb{F}_q} (q \cdot \Psi(-u_0))^h \cdot \Psi(u_0 x_0) + \sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{x}) \\ &= \frac{1}{q^{h(k-1)}} \sum_{u_0 \in \mathbb{F}_q} \Psi(u_0(x_0 - h)) + \sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{x}) . \end{aligned}$$

Now let $\mathbf{s} \in \mathbb{F}_q^k$ be such that $s_0 = h \bmod q$. Multiplying the above equality by q^{hk} , we have

$$\begin{aligned} \left| \{(\mathbf{w}_1, \dots, \mathbf{w}_h) \in W^h : \sum_{i=1}^h \mathbf{w}_i = \mathbf{s}\} \right| &= q^{h+1} + q^{hk} \sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{s}) \\ &\geq q^{h+1} - q^{hk} \sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} \left| (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{s}) \right| \\ &\geq q^{h+1} - q^k \cdot ((k-2)\sqrt{q})^h , \end{aligned} \tag{26}$$

where the equality uses the fact that $s_0 = h \bmod q$, the first inequality uses the triangle inequality, and the second inequality uses the Weil bound (Equation (25)). Taking $k = q^{\varepsilon'}$ for some constant $\varepsilon' > 0$, the last bound is roughly $q^h - q^k \cdot q^{(1/2+\varepsilon')h} = q^h - q^{k+(1/2+\varepsilon')h}$, which we want to be subexponentially large in q .

However, in order to even make the above quantity positive, we must take $h > k/(1/2 - \varepsilon') = (2 + \varepsilon)k$ for some constant $\varepsilon > 0$, whereas our application needs $h < 2k$ (since we know only that $\lambda_1^{(1)}(\mathcal{L}^\perp(H)) \geq 2k$). Indeed, this same condition shows up in related work: our requirement that $h > (2 + \varepsilon)k$ is essentially the same as the requirement that “ $g \geq (2 + \varepsilon)(h + 1)$ ” in the premise of [CW09, Theorem 2.1], even though the setting there is slightly different. In particular, in each case the “2” in the $2 + \varepsilon$ factor comes from (the inverse of) the exponent in the Weil bound’s $q^{1/2}$ factor.

¹³ Technically, we care about the number of *multisets* (rather than sequences) of vectors $V \subseteq W$, $|V| = h$ such that $\sum_{\mathbf{w} \in V} \mathbf{w} = \mathbf{x}$. Indeed, there is a bijection between such multisets and their indicator vectors $\mathbf{y}' \in (\mathbb{Z}^{\geq 0})^q$, which are such that $\|\mathbf{y}'\|_1 = h$ and $H\mathbf{y}' = \mathbf{x}$. However, the number of such multisets is at most an $h!$ factor less than the number of such sequences, and this factor seems inconsequential for our purposes; see Equation (26) and the subsequent discussion.

Toward making the approach work. The above argument ‘just barely’ fails—i.e., improving the constant from $2 + \varepsilon$ to $2 - \varepsilon$ would suffice for our purposes—so one might hope to make it work by showing that one of the inequalities in Equation (26) is loose. First, one might hope to strengthen the Weil bound by reducing its \sqrt{q} factor to $q^{1/2-\delta}$ for some constant $\delta > 0$. Furthermore, the Weil bound is a “worst-case” statement about character sums evaluated on *arbitrary* polynomials, whereas for the argument above an improvement to $q^{1/2-\delta}$ on average (i.e., for *random* polynomials) would suffice. Unfortunately, as shown in [Saw14], the \sqrt{q} factor is optimal even for such an “average-case Weil bound.” So, the main possibility for improvement is to avoid naïvely using the triangle inequality in the first inequality. Indeed, there could be substantial “phase cancellations” among the terms in the sum $\sum_{\mathbf{u} \in \mathbb{F}_q^k \setminus U_0} (\widehat{1_W}(\mathbf{u}))^h \cdot \Psi_{\mathbf{u}}(\mathbf{s})$ for some special (or perhaps arbitrary) syndromes \mathbf{s} . It seems likely that such cancellations occur, but we do not know how to show it. (We emphasize that Equation (26) holds for *any* \mathbf{s} with $s_0 = h \bmod q$; it is agnostic to the rest of \mathbf{s} .)

Finally, we note another possible modification to the argument, which is to work with the “symmetrized set” $W' := \pm W = \{\pm(1, a, a^2, \dots, a^{k-1})^T : a \in \mathbb{F}_q\}$ instead of W . A multiset of h vectors from W' naturally corresponds to a “signed” vector $\mathbf{y}' \in \mathbb{Z}^q$, whereas such a multiset of vectors in W corresponds to an “unsigned” vector $\mathbf{y} \in (\mathbb{Z}^{\geq 0})^q$. That is, we can count coset vectors having negative coordinates when working with W' instead of W . However, it is unclear to us how to use this relaxation to any advantage.