Fast Decoding of Explicit Almost Optimal ε-Balanced q-Ary Codes And Fast Approximation of Expanding k-CSPs

Fernando Granha Jeronimo ✉
Institute for Advanced Study, Princeton, NJ, USA

Abstract

Good codes over an alphabet of constant size q can approach but not surpass distance $1 - 1/q$. This makes the use of q-ary codes a necessity in some applications, and much work has been devoted to the case of constant alphabet q. In the large distance regime, namely, distance $1 - 1/q - \varepsilon$ for small $\varepsilon > 0$, the Gilbert–Varshamov (GV) bound asserts that rate $\Omega_q(\varepsilon^2)$ is achievable whereas the q-ary MRRW bound gives a rate upper bound of $O_q(\varepsilon^2 \log(1/\varepsilon))$. In this sense, the GV bound is almost optimal in this regime. Prior to this work there was no known explicit and efficiently decodable q-ary codes near the GV bound, in this large distance regime, for any constant $q \geq 3$.

We design an $\widetilde{O}_{\varepsilon,q}(N)$ time decoder for explicit (expander based) families of linear codes $\mathcal{C}_{N,q,\varepsilon} \subseteq \mathbb{F}_q^N$ of distance $(1-1/q)(1-\varepsilon)$ and rate $\Omega_q(\varepsilon^{2+o(1)})$, for any desired $\varepsilon > 0$ and any constant prime q, namely, almost optimal in this regime. These codes are ε-balanced,i.e., for every non-zero codeword, the frequency of each symbol lies in the interval $[1/q - \varepsilon, 1/q + \varepsilon]$. A key ingredient of the q-ary decoder is a new near-linear time approximation algorithm for linear equations (k-LIN) over $\mathbb{Z}_q$ on expanding hypergraphs, in particular, those naturally arising in the decoding of these codes.

We also investigate k-CSPs on expanding hypergraphs in more generality. We show that special trade-offs available for k-LIN over $\mathbb{Z}_q$ hold for linear equations over a finite group. To handle general finite groups, we design a new matrix version of weak regularity for expanding hypergraphs. We also obtain a near-linear time approximation algorithm for general expanding k-CSPs over q-ary alphabet. This later algorithm runs in time $\widetilde{O}_{k,q}(m+n)$, where m is the number of constraints and n is the number of variables. This improves the previous best running time of $O(n^{\Theta_{k,q}(1)})$ by a Sum-of-Squares based algorithm of [AJT, 2019] (in the expanding regular case).

We obtain our results by generalizing the framework of [JST, 2021] based on weak regularity decomposition for expanding hypergraphs. This framework was originally designed for binary k-XOR with the goal of providing near-linear time decoder for explicit binary codes, near the GV bound, from the breakthrough work of Ta-Shma [STOC, 2017]. The explicit families of codes over prime $\mathbb{F}_q$ are based on suitable instatiations of the Jalan–Moshkovitz (Abelian) generalization of Ta-Shma's distance amplification procedure.

2012 ACM Subject Classification Theory of computation → Error-correcting codes; Theory of computation → Expander graphs and randomness extractors

Keywords and phrases Decoding, Approximation, GV bound, CSPs, HDXs, Regularity

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2023.60

Category RANDOM

Funding This material is based upon work supported by the NSF grant CCF-1900460. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.
Fernando Granha Jeronimo: NSF grant CCF-1900460

Acknowledgements We thank Vedat Alev and Shravas Rao for stimulating discussions during the initial phase of this project. We thank Shashank Srivastava and Madhur Tulsiani for stimulating discussions leading to [24].

1 Introduction

Codes over small alphabet sizes have attracted a lot of effort in coding theory [17]. There is now a vast theory about them, but important mysteries remain. One very natural alphabet is the binary alphabet, which has a myriad of uses and applications. However, it also comes with an important limitation, namely, a family of good binary codes cannot[1] surpass distance $1/2$. By using a q-ary alphabet, a family of good codes can approach distance $1 - 1/q$ but not surpass it. This makes the use of q-ary codes a necessity whenever larger distances are needed. Working towards explicit and efficiently decodable codes with optimal trade-offs between rate and distance has been a challenging but fruitful guiding goal in coding theory.

In the large distance case, namely, distances are of the form $1 - 1/q - \varepsilon$ for small values of $\varepsilon > 0$, the Gilbert–Varshamov (GV) bound [13, 36] asserts that rate $\Omega_q(\varepsilon^2)$ is achievable whereas the q-ary version of McEliece, Rodemich, Rumsey and Welch (MRRW) [26] gives an impossibility upper of $O_q(\varepsilon^2 \log(1/\varepsilon))$. This means that the GV bound is nearly optimal in this regime of constant alphabet size q and large distance. To the best of our knowledge, in this regime, (prior to this work) no explicit and efficiently decodable families of q-ary codes near the GV bound were known for any $q \geq 3$.

Two widely used approaches in the construction of q-ary codes for small q are based on code concatenation [11] and on algebraic geometry (AG) constructions [30, 34]. Using code concatenation, it is possible to obtain explicit constructions achieving the suboptimal Zyablov bound trade-off between rate and distance, which gives a rate of $\Omega_q(\varepsilon^3)$. Some explicit families of AG codes are celebrated for beating the GV bound in some specific parameter regimes, e.g., the seminal work of Tsfasman, Vlădut and Zink[2] [35] or the (non-linear) construction of Elkies [9]. This surprising phenomenon of explicit AG codes beating random codes cannot happen in a major way in the large distance and constant alphabet regime since the GV bound is nearly optimal. Furthermore, known explicit constructions of linear AG codes are far from the GV bound for large distances and constant q. Another drawback of several explicit families of good AG codes is that known decoders can take much longer than linear time in the blocklength [27].

On a more combinatorial side, in a breakthrough work using expander graphs, Ta-Shma [31] gave the first explicit construction of binary codes of distance $1/2 - \varepsilon$ and rate $\Omega(\varepsilon^{2+o(1)})$, namely, near the Gilbert–Varshamov bound. A polynomial time decoder for these binary codes was first given in [23] followed by a near-linear time decoder in [24]. Subsequently, Jalan and Moshkovitz [22] extended Ta-Shma's analysis [31] to handle (in particular) codes over larger alphabets[3]. Suitable instantiations of [22] imply explicit codes over prime $\mathbb{F}_q$ of distance $1 - 1/q - \varepsilon$ with rate $\Omega_q(\varepsilon^{2+o_q(1)})$, namely, again near the (q-ary) GV bound for constant q.

Motivated by the above situation, we design a near-linear time decoder for explicit families of q-ary codes of distance $(1 - 1/q)(1 - \varepsilon)$ and rate $\Omega(\varepsilon^{2+o_q(1)})$ for any constant prime q, namely, near the GV bound in the large distance regime. More precisely, our main result is as follows (answering a question from [22]).

---

[1] This is a consequence of the Plotkin bound.

[2] More precisely, the TVZ bound [35] establishes a rate of $r \geq 1 - \delta - 1/(\sqrt{q} - 1)$ with respect to the relative distance δ.

[3] More precisely, [22] analyzed the (scalar) Abelian case of Ta-Shma's amplification.

▶ **Theorem 1** (Main I – Near-linear Time Unique Decoding over $\mathbb{F}_q$)**.** *Let q be a prime. For every $\varepsilon > 0$ sufficiently small, there are explicit linear Ta-Shma codes $\mathcal{C}_{N,q,\varepsilon} \subseteq \mathbb{F}_q^N$ for infinitely many values $N \in \mathbb{N}$ with*

- **(i)** *distance at least $(1-1/q)(1-\varepsilon)$ (actually ε-balanced),*
- **(ii)** *rate $\Omega_q(\varepsilon^{2+\alpha})$ where $\alpha = O(1/(\log_2(1/\varepsilon))^{1/6})$, and*
- **(iii)** *an $r(q/\varepsilon)\cdot\tilde{O}(N)$ time randomized unique decoding algorithm that decodes within radius $((1-1/q)(1-\varepsilon))/2$,*

where $r(x) = \exp(\exp(\mathrm{poly}(x)))$.

In fact, we actually prove the following stronger *list* decoding result.

▶ **Theorem 2** (Near-linear time List Decoding over $\mathbb{F}_q$)**.** *Let q be a prime. For every $\varepsilon > 0$ sufficiently small, there are explicit binary linear Ta-Shma codes $\mathcal{C}_{N,q,\varepsilon} \subseteq \mathbb{F}_q^N$ for infinitely many values $N \in \mathbb{N}$ with*

- **(i)** *distance at least $(1-1/q)(1-\varepsilon)$ (actually ε-balanced),*
- **(ii)** *rate $\Omega_q(\varepsilon^{2+\alpha})$ where $\alpha = O(1/(\log_2(1/\varepsilon))^{1/6})$, and*
- **(iii)** *an $r(q/\varepsilon)\cdot\tilde{O}(N)$ time randomized list decoding algorithm that decodes within radius $1-1/q-2^{-\Theta_q((\log_2(1/\varepsilon))^{1/6})}$ and works with high probability,*

where $r(x) = \exp(\exp(\mathrm{poly}(x)))$.

We obtain our results by building on and extending the *binary* decoding framework in [24]. This framework is based on a generalization of the weak regularity decomposition to (sparse) *expanding* hypergraphs that generalizes the seminal work of Frieze and Kannan [12]. The weak regularity decomposition of [24] was then used to approximate *expanding* k-XOR instances naturally arising in the decoding of binary Ta-Shma's codes [31]. Similarly, constraint satisfaction problems (CSPs) will play a key role in our decoder. Here, we also take the opportunity to investigate *expanding* CSPs more broadly.

An instance of a k-CSP is given by a k-uniform (ordered) constraint hypergraph $W \subseteq [n]^k$, where each vertex is associated with a variable taking values in an alphabet of size q and each edge is associated with a constraint involving the variables of its vertices. While even approximating a CSP is NP-hard in general, suitable notions of expansion of the constraint hypergraph allow for efficient approximation algorithms as in [24]. One such notion is *splittability* [2]. Roughly speaking, a τ-splittable collection of tuples for some $\tau \in (0,1]$ is the higher-order analogue of the second largest singular value of the normalized adjacency matrix of a graph (the smaller the τ the more expanding is the collection). Approximating expanding k-CSPs is at the core of some decoding algorithms for expander based constructions of codes [8, 1, 23, 24, 6].

As mentioned above, approximating expanding k-CSPs will be again at the core of our extension of [24] to more general constraints over larger alphabets. Our new q-ary decoder will need to handle instances of linear equations over the alphabet $\mathbb{Z}_q$, where each equation involves a sum of k variables. This kind of k-CSP is commonly denoted k-LIN over alphabet $\mathbb{Z}_q$. We will see that the special algebraic structure of these linear constraints will allow to obtain some improved parameter trade-offs, which will be explored in the decoding application. More precisely, the expansion (splittability) parameter τ will have no dependence on alphabet size q and only a polynomial dependence on the arity[4] k, and this allows us to obtain better approximation guarantees. Our second result follows.

[4] In the binary case of [24], it was also possible to have a polynomial dependence on the arity k.

▶ **Theorem 3** (Main II). *Let $\mathfrak{I}$ be an instance of MAX k-LIN$_q$ on n variables with alphabet $\mathbb{Z}_q$ and constraints supported on a regular[5] collection of tuples $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,\delta) \coloneqq \mathrm{poly}(\delta/k)$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $r(q/\tau_0) \cdot \widetilde{O}(|W| + n)$, where $r(x) = \exp(\exp(\mathrm{poly}(x)))$.*

We show that this phenomenon of no dependence of the expansion on the alphabet size q and only polynomial dependence on arity k also occurs for linear equations over a general finite groups $\mathfrak{G}$. Similarly, this leads to better approximation guarantees. To actually implement and obtain this advantage, we will design a new matrix version of the weak regularity decomposition for expanding hypergraphs. Our third result follows.

▶ **Theorem 4** (Main III). *Let $\mathfrak{I}$ be an instance of MAX k-LIN$_{\mathfrak{G}}$ on n variables with alphabet a finite group $\mathfrak{G}$ and constraints supported on a regular collection of tuples $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,\delta) \coloneqq \mathrm{poly}(\delta/k)$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $O_{|\mathfrak{G}|,k,\delta}(1) \cdot \mathrm{poly}(|W| + n)$.*

▶ Remark 5. In Theorem 4, we did not attempt to make the running time near-linear in the number of constraints and variables, but it is plausible that it can be done.

We find intriguing this interplay between the type of constraint used in the CSP and the expansion requirement for a given approximation. A natural question is to investigate this interplay for more general constraint types.

In this work, we also investigate how fast we can approximate expanding k-CSPs over q-ary alphabet without making any assumptions on the constraints. We show that k-CSPs can be approximated in near-linear time in the number of constraints and variables, assuming k and q are constants, and provided the constraint hypergraph is sufficiently expanding (splittable). An important caveat of this general case is that the expansion requirements will now depend on both the alphabet size q and arity k in an exponential way (of the form $q^{-O(k)}$).

▶ **Theorem 6.** *Let $\mathfrak{I}$ be an instance of MAX k-CSP on n variables with alphabet $[q]$ and constraints supported on a regular collection of tuples $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,q,\delta) \coloneqq \mathrm{poly}(\delta/(kq^k))$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $r(kq/\delta) \cdot \widetilde{O}(|W| + n)$, where $r(x) = \exp(\exp(\exp(\mathrm{poly}(x))))$.*

We obtain the above result via a reduction to the "binary" weak regularity in [24] in a somewhat similar fashion to [10]. Even though it is not hard to make this connection, we think it is worth stating it since this result may be more broadly applicable. Moreover, for fixed arity k and alphabet size q, this improves the running time in the expanding regime of the Sum-of-Squares based algorithm in [2] and also the expanding regime[6] of earlier results 2-CSPs [5, 18, 19, 28].

For comparison, we recall the expanding regime[7] of [2] below.

▶ **Theorem 7** (Sum-of-Squares [2]). *Let $\mathfrak{I}$ be an instance of MAX k-CSP on n variables with alphabet $[q]$ and constraints supported $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,q,\delta) \coloneqq \mathrm{poly}(\delta/k) \cdot q^{-k}$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $n^{\mathrm{poly}(1/\tau_0)}$.*

▶ Remark 8. In the new theorem above, we do not attempt to optimize the function $r(x)$.

---

[5] This is an analog to tuples of a graph being d-vertex regular.

[6] We point out these approaches also consider when the expansion is defective (low threshold rank case). Since we are interested in near-linear running time, we need to focus on the expanding case.

[7] Using the improved analysis of swap walks by Dikstein and Dinur [7].

Related Work. As we mentioned above, our work is an extension of the *binary* framework of [24]. This framework was designed for approximating expanding k-XOR and to give a near-liner time decoding algorithm for the explicit binary codes of Ta-Shma [31], near the GV bound. The first polynomial time decoder for these codes was given in [23] using the Sum-of-Squares semi-definite programming hierarchy and its running time, albeit polynomial, is very far from near-linear in the blocklength.

AG codes are widely used in the study of explicit constructions over constant q-ary alphabets. Some of these constructions achieve very competitive parameter trade-offs (e.g., rate versus distance) if not the best known in several cases. However, explicit and efficiently decodable codes near GV bound for large distances,i.e., $1 - 1/q - \varepsilon$, and constant alphabet size were not known prior to this work. In fact, the first explicit construction only appeared in the breakthrough work of [31] for binary codes using more combinatorial expander based techniques. This absence of explicit construction near the GV bound in this regime means that much is yet to be discovered about this case. We view our near-linear time decoder of prime q-ary codes in this regime as not only reaching previously unattained parameter regimes with an explicit construction, but also offering a more combinatorial perspective among a wealthy of algebraic techniques.

For *non-explicit* families of codes approaching the GV bound, much more is known. Random linear codes achieve this bound, but their decoding is believed to be computationally hard. It is possible to construct more structured ensembles of random codes that allow for efficient decoding in this regime. We have the non-explicit classical Goppa codes. Another important technique is based on Thommesen's [32] technique of concatenation with random inner codes. These Thommesen based ensembles can sometimes approach the GV bound and also allow for efficient decoding [16, 14, 21, 25] and even near-linear time decoding [21, 25].

More recently, Blanc and Doron [6] used the framework in [24] to decode explicit binary codes near the GV bound with improved parameters, where they obtain a polynomial improvement on the $o(1)$ error term of the rate $\Omega(\varepsilon^{2+o(1)})$ (the α in Theorem 1) and also put forward some interesting conjectures towards further improving the rate. It is plausible that their improvement also applies here for q-ary alphabets.

In the constant alphabet case, a different parameter regime that has received much attention is the near-capacity regime [15, 20, 21, 25] of list decoding from radius $1 - r - \varepsilon$ with rate r for small values of $\varepsilon > 0$. This regime can only occur when the alphabet size q is a function of ε. Note that our near GV bound regime is the opposite, we have a fixed constant q and we can take ε arbitrarily small (smaller than some function of q).

Due to space constraints, most of our proofs only appear in the full version of this paper.

2 Proof Strategy

We will now describe our contributions in more detail. Our algorithmic results will be based on extensions of the *binary* weak regularity framework of [24]. Roughly speaking, this framework being a *"low level"* framework gives fine control over its components leading to a near-linear time decoder for Ta-Shma's codes [31] over $\mathbb{F}_2$. This same low level structure means that extensions may require suitable generalizations in several of these components as well technical work to implement them. The extensions to handle codes over prime q-ary alphabet and a matrix version of weak regularity will be no exception.

First, we will recall the weak regularity decomposition of Frieze and Kannan [10] in a more analytic form [33]. We will also first consider its *existential* form and later discuss its algorithmic form. Our setup will be as follows. Let $W \subseteq [n]^k$ be a collection of tuples endowed

with the uniform probability measure μ_k. Suppose that we have a function $g\colon W \to \mathbb{C}$ that we want to approximate using a simpler approximating function, which will be made precise below. Further suppose that the quality of approximation will be measured with respect to correlations with a class of test functions $\mathcal{F}$. Given some desired approximation error $\delta > 0$, the goal will be to find a *"simple"* approximator $h \approx g$ such that

$$\max_{f \in \mathcal{F}} \left| \langle g - h, f \rangle_{\mu_k} \right| \leq \delta \,.$$

As an *existential* result, it is well-known that an h of the form $h = \sum_{\ell=1}^{p} c_\ell \cdot f_\ell$ always exists, where c_ℓ's are scalars and the f_ℓ's are functions belonging to $\mathcal{F}$. Furthermore, the number of test functions p is small being at most[8] $O(1/\delta^2)$. This means that h is indeed *"simple"* since it is the sum of a small number of test functions, so h is almost as complex as the test functions it needs to fool.

To motivate the generalizations in the weak regularity framework, we will start the discussion of the important case of linear equations over $\mathbb{Z}_q$ as a motivating example. As mentioned above, approximating k-LIN over $\mathbb{Z}_q$ will be crucial in the near-linear time decoding algorithm for prime q-ary alphabets. For us, an instance $\mathfrak{I}$ of k-LIN is given by a system of linear equations[9]

$$x_{i_1} + \cdots + x_{i_k} \equiv r_w \pmod{q} \qquad \forall\ w = (i_1, \ldots, i_k) \in W, \tag{1}$$

where $(r_w)_{w \in W} \in \mathbb{Z}_q^W$ are given RHS coefficients. We will need to model this problem in a way that is amenable to the weak regularity approach. We will also take advantage of the algebraic structure of the constraints to avoid any dependence of the alphabet size q and to have only a mild dependence on the arity k in the expansion the framework will require from W.

"Global" Approximation of Dirac Delta Functions. An elementary property of Fourier analysis over $\mathbb{Z}_q$ is that the Dirac delta function $x \mapsto \mathbf{1}_{[x=y]}$ admits a simple but extremely handy Fourier decomposition which we now recall. Let $\omega = \exp(2\pi\sqrt{-1}/q)$. Using orthogonality of characters, we have

$$\mathbf{1}_{[x=y]} = \mathop{\mathbb{E}}_{a \in \mathbb{Z}_q} \left[\omega^{a(x-y)} \right] .$$

Suppose we have an assignment $b \in \mathbb{Z}_q^n$ to the variables of our system of linear equations $\mathfrak{I}$. Then, the fraction of satisfied constraints, which we denote by $\mathsf{val}(\mathfrak{I}, b)$ and refer as the value of this assignment, can be expressed as

$$\mathsf{val}(\mathfrak{I}, b) \coloneqq \mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k} \left[\mathbf{1}_{[b_{i_1}+\cdots+b_{i_k} \equiv r_w]} \right] = \mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k} \left[\mathop{\mathbb{E}}_{a \in \mathbb{Z}_q} \left[\omega^{a(b_{i_1}+\cdots+b_{i_k}-r_w)} \right] \right] .$$

This suggests defining q functions one for each $a \in \mathbb{Z}_q$ of the form $g_a\colon W \to \mathbb{C}$ as $g_a(w) \coloneqq \omega^{a \cdot b_w}$, the *"harmonic"* components. We also endow the space $\mathbb{C}^W$ with the inner product defined by the measure μ_k on W. We will need some additional notation. For $b \in \mathbb{Z}_q^n$, we define the function $\chi_{b,a}$ on $[n]$ as $\chi_{b,a}(i) = \omega^{a \cdot b_i}$. We can now rexpress $\mathsf{val}(\mathfrak{I}, b)$ in terms of its harmonic components as

[8] The ℓ_1-norm of the coefficients is "small", i.e., $\sum_{\ell=1}^{p} |c_\ell|$.
[9] The coefficients of the variables are always taken to be 1 here.

$$\begin{aligned}
\mathsf{val}(\mathfrak{I}, b) &= \mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k}\left[\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\omega^{a(b_{i_1}+\cdots+b_{i_k}-r_w)}\right]\right] \\
&= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k}\left[\omega^{-a\cdot r_w}\cdot\omega^{a(b_{i_1}+\cdots+b_{i_k})}\right]\right] \\
&= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle g_a, \underbrace{\chi_{b,a}\otimes\cdots\otimes\chi_{b,a}}_{k}\right\rangle_{\mu_k}\right] \\
&= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle g_a, (\chi_{b,a})^{\otimes k}\right\rangle_{\mu_k}\right].
\end{aligned}$$

We can now try to further approximate each g_a using a simpler function h_a that behaves similarly to g_a with respect to functions of the form $f_{b,a} = \chi_{b,a}\otimes\cdots\otimes\chi_{b,a}$ as in the inner product above. We can view functions of form $f_{b,a}$ as tests with respect to which g_a and its simpler approximator have similar correlations. This means that we can model the problem in way amenable to the existential weak regularity framework. For each a, we will consider a (slightly) more general class of test functions $\mathrm{CUT}^{\otimes k}_{\omega,q,a}$ defined as follows

$$\mathrm{CUT}^{\otimes k}_{\omega,q,a} \coloneqq \{\chi_{b^{(1)},a}\otimes\cdots\otimes\chi_{b^{(k)},a} \;\mid\; b^{(1)},\ldots,b^{(k)}\subseteq\mathbb{Z}_q^n\}.$$

A simple yet useful remark is that if we can find a decomposition fooling a larger class of test functions, this would suffice since, in particular, it fools the initial class of test.

Suppose that for some $\delta\in(0,1)$ we can find a δ-approximation $h_a = \sum_{\ell=1}^{p_a} c_{a,\ell}\cdot\chi_{b^{(a,\ell,1)},a}\otimes\cdots\otimes\chi_{b^{(a,\ell,k)},a}$ to g_a with respect to a class of test functions, i.e.,

$$\max_{f\in\mathrm{CUT}^{\otimes k}_{\omega,q,a}}\left|\langle g_a - h_a, f\rangle_{\mu_k}\right| \le \delta.$$

By replacing g_a with h_a in the computation of $\mathsf{val}(\mathfrak{I}, b)$ above, we obtain[10]

$$\mathsf{val}(\mathfrak{I}, b) \;=\; \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle g_a, (\chi_{b,a})^{\otimes k}\right\rangle\right] \;=\; \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle h_a, (\chi_{b,a})^{\otimes k}\right\rangle\right] \pm \delta.$$

We will explain how to algorithmically find h_a in near-linear time later. Now, we will argue why having access to weak regularity decomposition greatly simplifies our task of approximating $\mathsf{val}(\mathfrak{I}, b)$ and also later while decoding q-ary codes.

We can simplify the above equation for $\mathsf{val}(\mathfrak{I}, b)$ even further using the assumed expansion (splittability) of W. A suitable version of the expander mixing lemma allows us to pass from the measure μ_k to the product measure $\mu_1^{\otimes k}$, where μ_1 is the uniform measure on $[n]$. More precisely, we can show that if W is sufficiently expanding (depending on δ), then

$$\begin{aligned}
\mathsf{val}(\mathfrak{I}, b) = \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle h_a, (\chi_{b,a})^{\otimes k}\right\rangle_{\mu_k}\right] \pm \delta &= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle h_a, (\chi_{b,a})^{\otimes k}\right\rangle_{\mu_1^{\otimes k}}\right] \pm 2\delta \\
&= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\sum_{\ell=1}^{p_a} c_{a,\ell}\cdot\prod_{j=1}^{k}\left\langle\chi_{b^{(a,\ell,1)},a}, \chi_{b,a}\right\rangle_{\mu_1}\right] \pm 2\delta.
\end{aligned}$$

The *low complexity* of the approximator h_a will allow us to simplify the search for an approximately optimal assignment $b\in\mathbb{Z}_q^n$. The expression above reveals that we only need to know the values of

[10] For scalars x, y (real or complex) and real $\delta\in\mathbb{R}^+$, we use the notation $x = y\pm\delta$ if $|x-y|\le\delta$.

$$\left\{\left\langle \chi_{b^{(a,\ell,j)},a}, \chi_{b,a}\right\rangle_{\mu_1}\right\}_{a\in\mathbb{Z}_q,\ell\in[p_a],j\in[k]}.$$

Luckily, algorithmically, there will be only $O(qk^3/\delta^2)$ such numbers (no dependence on n and only slightly more than the $O(qk/\delta^2)$ from the existential result). Using brute-force search, it is possible to find sufficiently fine and (close to valid) approximations for these numbers.

To make the entire process efficient and near-linear time we still need to say how to find the functions h_a's in near-linear time. As in [24], we will reduce the problem of finding a weak regularity decomposition with respect to a class of k-tensors, in this case the class $\mathrm{CUT}^{\otimes k}_{\omega,q,a}$, to multiple applications of the 2-tensor case (in a sparse regime). To execute this process in near-linear, we will again use the expansion of W to conveniently move to easier to handle product measures (as above). This involves finding a constant factor approximation for the following expression

$$\max_{x,y\in\mathbb{Z}_q^n} \left|\sum_{i,j=1}^{n} \mathsf{A}_{i,j}\cdot\omega^{a\cdot x_i}\cdot\omega^{a\cdot y_j}\right|, \tag{2}$$

This kind of optimization is known as the *Grothendieck problem* and, in this case, it is for roots of unity going beyond the ±1 case of Alon and Naor [3]. In [29], So, Zhang and Ye considered a more restricted version of this problem (with positive semi-definite matrices) known as the *little* Grothendieck problem. We will extend their analysis to the Grothendieck problem building on some ingredients present in their proof. In our application, the matrices A will be sparse with $m \approx n$ non-zero entries and to achieve a near-linear time we will need to find an (additive) approximation to the Grothendieck problem in time $\widetilde{O}(m)$ of Equation (2). This can be done using the fast SDP solver of Arora and Kale [4].

We now explain how the above weak regularity decomposition can be used in decoding of the expander based construction of Ta-Shma's codes [31]. We will see that the decoding problem can be naturally phrased as a k-LIN instance over $\mathbb{Z}_q$, which is a natural q-ary extension of the k-XOR over $\mathbb{Z}_2$ from [1, 23, 24]. First, we briefly describe Ta-Shma's code construction over alphabet $\mathbb{F}_q$, with q prime, as analyzed[11] in [22]. The idea is to start with a good base code $\mathcal{C}_0 \subseteq \mathbb{F}_q^n$ and to use a carefully constructed collection of tuples $W \subseteq [n]^k$ to amplify its distance via the direct-sum encoding. For any $z \in \mathbb{F}_q^n$, recall that its direct-sum encoding is a new word denoted $y = \mathrm{dsum}_W(z)$ in $\mathbb{F}_q^W$ and defined as

$$y_{(i_1,\ldots,i_k)} \;=\; z_{i_1} + \cdots + z_{i_k} \pmod q \qquad\qquad \forall\,(i_1,\ldots,i_k)\in W\,.$$

The direct-sum code $\mathcal{C} = \mathrm{dsum}_W(\mathcal{C}_0)$ is defined as $\mathcal{C} = \{\mathrm{dsum}_W(z) \;\mid\; z \in \mathcal{C}_0\}$. Note the similarity of the above equation and the system of linear equations from Equation (1). In the decoding task, we are given a (possibly) corrupted version of $\widetilde{y}$ of some codeword $y = \mathrm{dsum}_W(z) \in \mathcal{C}$, with $z \in \mathcal{C}_0$. We can view $\widetilde{y}$ as defining the RHS coefficients of an instance of k-LIN, namely, $r_W = \widetilde{y}_w$.

Having an instance of k-LIN over $\mathbb{Z}_q$, we can now use weak regularity as described above. For each $a \in \mathbb{Z}_q$, let g_a be the *harmonic* component associated with RHS vector $\widetilde{y}$ (as above). Similarly, we find a weak regularity approximation h_a for each function g_a.

[11] In [22], they considered the more general (scalar) Abelian case.

If the distance $\Delta(\widetilde{y}, \mathrm{dsum}_W(z)) \leq (1-1/q)(1-\beta)$ is not too large, we will be able to deduce that some harmonic function h_a "captures" the structure of the codeword z in the following sense. Set $\mathcal{R} = \{\omega^{a \cdot a'} \mid a' \in \mathbb{Z}_q\}$ and let $f_1, \ldots, f_r \colon [n] \to \mathcal{R}$ be the functions appearing in the decomposition of h_a. For each tuple $(y_1, \ldots, y_r) \in \mathcal{R}^r$, we can consider the set

$$\{x \in [n] \mid f_1(x) = y_1, \ldots, f_r(x) = y_r\}.$$

These sets partition[12] the space $[n]$, and we can show that z is approximately constant in most of these parts. In this sense, the low complexity structure of h_a captures the structure of the codeword z. In this last argument, we use that assumption that q is prime[13].

The case of k-LIN over a finite group will also allow for a weak regularity decomposition in a similar spirit as above, where scalar Fourier characters are replaced by larger dimensional representations and "global" approximation of Dirac delta functions are performed. Extending the weak regularity framework to this case will require considering matrix valued functions. The way we model this case is done in the full version and it uses very elementary properties of representation theory. This case again exhibits an interesting interplay between the type of constraints and the requirement on expansion. (The reader who is only interested in decoding can safely ignore this extension and focus on the $\mathbb{Z}_q$ case.)

3 Constraint Types and Alphabets

We explore the role of different types of constraints and corresponding alphabets going beyond the binary k-XOR considered in [24]. For the special case of linear equations over $\mathbb{Z}_q$ or over an arbitrary finite group $\mathfrak{G}$, we will explore the special structure of the constraints and obtain results with improved parameters.

3.1 General CSPs via the Binary Regularity

We will prove our first result for approximating a general expanding k-CSPs over a q-ary alphabet in near-linear time. We obtain this result using the *binary* near-linear time weak regularity decomposition from [24] in a similar way that Frieze and Kannan modeled k-CSPs [10] using regularity. We formalize this (relatively simple) connection since we believe this result may be of independent interest and may find applications elsewhere. Moreover, it also improves the running time of [2] to near-linear time, for fixed k and q, while offering a different approach to approximating general expanding k-CSPs which could be simpler than their Sum-of-Squares based algorithm. We now restate and proceed to prove this result.

▶ **Theorem 6.** *Let $\mathfrak{I}$ be an instance of MAX k-CSP on n variables with alphabet $[q]$ and constraints supported on a regular collection of tuples $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,q,\delta) := \mathrm{poly}(\delta/(kq^k))$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $r(kq/\delta) \cdot \widetilde{O}(|W| + n)$, where $r(x) = \exp(\exp(\exp(\mathrm{poly}(x))))$.*

We will find a weak regularity decomposition with respect to 0/1 valued test functions $\mathcal{F} = \mathrm{CUT}^{\otimes k}$ where

$$\mathrm{CUT}^{\otimes k} := \{\pm \mathbf{1}_{S_1} \otimes \cdots \otimes \mathbf{1}_{S_k} \mid S_1, \ldots, S_k \subseteq [n]\}.$$

The near-linear weak regularity decomposition of [24], which we recall below, can handle this class of functions.

[12] Possibly with empty parts.

[13] So that all non-trivial roots of unity are primitive roots. It is plausible that this restriction is not necessary.

▶ **Theorem 9** (Efficient Weak Regularity from [24]). *Let $W \subseteq [n]^k$ be a τ-splittable collection of tuples. Suppose $\mathcal{F}$ is one of $\mathrm{CUT}^{\otimes k}$, $\mathrm{CUT}_{\pm}^{\otimes k}$. Let $\mathcal{R}$ be the domain of the functions in $\mathcal{F}$, when $k = 1$. Let $g \in \mathcal{R}^{W[1]^k}$ be supported on W with $\|g\|_{\mu_k} \leq 1$. For every $\delta > 0$, if $\tau \leq \delta^2/(k^3 \cdot 2^{20})$, then we can find $h = \sum_{\ell=1}^{p} c_\ell \cdot f_\ell$ with $p = O(k^2/\delta^2)$, $c_1, \ldots, c_p \in \mathbb{R}$ and functions $f_1, \ldots, f_p \in \mathcal{F}$, such that $\|h\|_{\mu_1^{\otimes k}} \leq 2$, $\sum_{\ell=1}^{p} |c_\ell| = O(k/\delta)$ and h is a good approximator to g in the following sense*

$$\max_{f \in \mathcal{F}} \left| \left\langle g - \left(\frac{d}{n}\right)^{k-1} h, f \right\rangle \right| \leq \delta \cdot |W| ,$$

where the inner product is over the counting measure on $W[1]^k$. Furthermore, h can be found in $\widetilde{O}(2^{2^{\widetilde{O}(k^2/\delta^2)}} \cdot |W|)$ time.

Having access to a weak regularity decomposition as above makes the task of approximating the value of a CSP instance relatively simple, as we now describe. This is a common feature of weak regularity based arguments,e.g., [10, 28]. Here, we consider both arbitrary arity k and arbitrary alphabet size q.

We will first need some notation. Let $\alpha \in [q]^k$ and define $W_\alpha = \{w \in W \mid P_w(\alpha) = 1\}$ to be the set of tuples whose predicates P_w are satisfied by on the input α. Let $\mathcal{A}(\mathfrak{I}) = \{\alpha \in [q]^k \mid W_\alpha \neq \emptyset\}$ be the set of satisfying inputs of at least one predicate of $\mathfrak{I}$.

We will use the following claim which relates the value of an assignment to the structure of the weak regularity decomposition.

▷ Claim 10. Suppose that for every $\alpha \in [q]^k$, we have a weak regularity decomposition h_α, from Theorem 9, of the indicator $\mathbf{1}_{W(\alpha)}$ with error parameter $\delta > 0$ and with respect to the test class $\mathrm{CUT}^{\otimes k}$. Let $b \in [q]^n$ (viewed as an assignment), which induces a partition $T_1 \sqcup \cdots \sqcup T_q$ of $[n]$. Then,

$$\mathsf{val}(\mathfrak{I}, b) = \sum_{\alpha \in \mathcal{A}} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \frac{\left|S_1^{\alpha,\ell} \cap T_{\alpha_1}\right|}{n} \cdots \frac{\left|S_k^{\alpha,\ell} \cap T_{\alpha_k}\right|}{n} \pm \delta \cdot |\mathcal{A}(\mathfrak{I})| .$$

Proof. Let $\mathcal{A} = \mathcal{A}(\mathfrak{I})$. The value of this assignment is

$$\begin{aligned}
\mathsf{val}(\mathfrak{I}, b) &= \sum_{\alpha \in \mathcal{A}} \left\langle \mathbf{1}_{W_\alpha}, \mathbf{1}_{T_{\alpha_1}} \otimes \cdots \otimes \mathbf{1}_{T_{\alpha_k}} \right\rangle_{\mu_k} \\
&= \frac{1}{|W|} \sum_{\alpha \in \mathcal{A}} \left\langle \left(\frac{d}{n}\right)^{k-1} h_\alpha, \mathbf{1}_{T_{\alpha_1}} \otimes \cdots \otimes \mathbf{1}_{T_{\alpha_k}} \right\rangle \pm \delta \cdot |\mathcal{A}| \\
&= \frac{1}{|W|} \sum_{\alpha \in \mathcal{A}} \left\langle \left(\frac{d}{n}\right)^{k-1} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \cdot \mathbf{1}_{S_1^{\alpha,\ell}} \otimes \cdots \otimes \mathbf{1}_{S_k^{\alpha,\ell}}, \mathbf{1}_{T_{\alpha_1}} \otimes \cdots \otimes \mathbf{1}_{T_{\alpha_k}} \right\rangle \pm \delta \cdot |\mathcal{A}| \\
&= \frac{1}{n^k} \sum_{\alpha \in \mathcal{A}} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \cdot \left\langle \mathbf{1}_{S_1^{\alpha,\ell}} \otimes \cdots \otimes \mathbf{1}_{S_k^{\alpha,\ell}}, \mathbf{1}_{T_{\alpha_1}} \otimes \cdots \otimes \mathbf{1}_{T_{\alpha_k}} \right\rangle \pm \delta \cdot |\mathcal{A}| \\
&= \sum_{\alpha \in \mathcal{A}} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \cdot \left\langle \mathbf{1}_{S_1^{\alpha,\ell}}, \mathbf{1}_{T_{\alpha_1}} \right\rangle_{\mu_1} \cdots \left\langle \mathbf{1}_{S_k^{\alpha,\ell}}, \mathbf{1}_{T_{\alpha_k}} \right\rangle_{\mu_1} \pm \delta \cdot |\mathcal{A}| \\
&= \sum_{\alpha \in \mathcal{A}} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \frac{\left|S_1^{\alpha,\ell} \cap T_{\alpha_1}\right|}{n} \cdots \frac{\left|S_k^{\alpha,\ell} \cap T_{\alpha_k}\right|}{n} \pm \delta \cdot |\mathcal{A}| ,
\end{aligned}$$

concluding the proof. ◁

Proof of Theorem 6. Let $\mathfrak{I}$ be an instance of a k-CSP over alphabet $[q]$ supported on a collection of tuples $W \subseteq [n]^k$ and with predicates $(P_w \colon [q]^k \to \{0,1\})_{w \in W}$.

For each $\alpha \in \mathcal{A}(\mathfrak{I})$, we apply the weak regularity decomposition of Theorem 9 to the function $\mathbf{1}_{W_\alpha}$ with error parameter $\delta > 0$ and test class $\mathcal{F} = \mathrm{CUT}^{\otimes k}$. This gives an approximation $h_\alpha = \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \cdot \mathbf{1}_{S_1^{\alpha,\ell}} \otimes \cdots \otimes \mathbf{1}_{S_k^{\alpha,\ell}}$.

A crucial property is that instead of having to know an assignment $b \in [q]^n$, represented as a partition $T_1 \sqcup \cdots \sqcup T_q = [n]$, it is enough to know the values of the following inner products

$$\left\{ \left\langle \mathbf{1}_{S_j^{\alpha,\ell}}, \mathbf{1}_{T_{\alpha_j}} \right\rangle_{\mu_1} \right\}_{\alpha \in \mathcal{A}(\mathfrak{I}), \ell \in [p_\alpha], j \in [k]}$$

The decomposition is *low complexity*, in the sense that there are only a few of these values. However, we cannot take arbitrary values for these inner products since they may be far from *realizable*, i.e., no true assignment $b \in [q]^n$ can give rise to these values even approximately. From the inner products above, we can extract the following class of functions

$$\mathcal{F}' = \left\{ \mathbf{1}_{S_j^{\alpha,\ell}} \right\}_{\alpha \in \mathcal{A}(\mathfrak{I}), \ell \in [p_\alpha], j \in [k]} ,$$

whose size $r = |\mathcal{F}'| = O(|\mathcal{A}(\mathfrak{I})|\, k^3/\delta^2)$ is independent from n.

Using Claim 10, to be able to approximate $\mathsf{val}(\mathfrak{I}, b)$ within error $\delta' > 0$ we need to choose the error of the weak regularity decomposition[14] to be $\delta = \delta'/(2\,|\mathcal{A}(\mathfrak{I})|)$ In this case, we have $r = O(|\mathcal{A}(\mathfrak{I})|^2\, k^3/(\delta')^2) = O(q^{2k} k^3/(\delta')^2)$ and the τ-splittability parameter of W needs to satisfy $\tau \leq \mathrm{poly}(\delta'/(kq^k))$.

For convenience, label the functions of $\mathcal{F}'$ as $f_1, \ldots, f_r$. Their range is the (simple) binary set $\mathcal{R} = \{0,1\}$. We will consider the factor $\mathcal{B}$ defined by the collection $\mathcal{F}'$, which, roughly speaking, is a partition of $[n]$ according to the values of these functions. More precisely, for every tuple $(y_1, \ldots, y_r) \in \mathcal{R}^r$ we have a (possibly empty) part (or atom) of the form

$$\{x \in [n] \;\mid\; f_1(x) = y_1, \ldots, f_r(x) = y_r\} .$$

In this case, we have at most $\mathcal{R}^r = 2^r$ atoms in the factor. By definition the functions $\mathcal{F}'$ are constant in each of them. An assignment b gives rise to a distribution on $[q]$ in each atom of the factor. Conversely, any approximate distribution on $[q]$ in each atom approximately corresponds to a realizable assignment b.

Let $L = \sum_{\alpha \in \mathcal{A}(\mathfrak{I}), \ell \in [p_\alpha]} |c_{\alpha,\ell}| \leq |\mathcal{A}(\mathfrak{I})|\, O(k/\delta)$. Set $\eta = \delta/(k \cdot L \cdot q)$. We can η-approximate these distributions in ℓ_1-norm on each atom[15]. The number of approximate distributions can be (crudely) bounded as

$$(1/(\eta q))^{\mathcal{R}^r} \leq \exp(\exp(\exp(\mathrm{poly}(qk/\delta')))) .$$

With this fine enough discretization of the distributions on each atom, when computing the expression

$$\mathsf{val}(\mathfrak{I}, b) = \sum_{\alpha \in \mathcal{A}} \sum_{\ell=1}^{p_\alpha} c_{\alpha,\ell} \frac{\left|S_1^{\alpha,\ell} \cap T_{\alpha_1}\right|}{n} \cdots \frac{\left|S_k^{\alpha,\ell} \cap T_{\alpha_k}\right|}{n} \;\pm\; \delta \cdot |\mathcal{A}|$$

we incur an additional error of $\delta'/2$. By our choice of δ, the total approximation error is at most δ'. ◀

[14] We can assume without loss of generality that $\mathcal{A}(\mathfrak{I}) \neq \emptyset$ since otherwise the value of the CSP is always zero.

[15] If the atom is too smaller than $1/(\eta q)$, then we can consider all the possible exact distribution.

3.2 Stating the Extended Weak Regularity Framework

We now show how to obtain our main results for linear equations k-LIN over $\mathbb{Z}_q$ in Theorem 3 and over a finite group $\mathfrak{G}$ in Theorem 4.

In the full version, we see that to approximate k-LIN over $\mathbb{Z}_q$ it suffices to find a good weak regularity decomposition with respect to the test functions $\mathcal{F} = \mathrm{CUT}^{\otimes k}_{\omega,q,a}$ defined as follows

$$\mathrm{CUT}^{\otimes k}_{\omega,q,a} \coloneqq \{\chi_{b_1,a} \otimes \cdots \otimes \chi_{b_k,a} \;\mid\; b_1,\ldots,b_k \subseteq \mathbb{Z}_q^n\}\,.$$

In the full version, we will see that to approximate k-LIN over a finite group, it suffices to find a good weak regularity decomposition with respect to the *matrix* valued test functions $\mathcal{F}$ defined as follows

$$\mathrm{CUT}^{\otimes k}_{\rho} \coloneqq \{\rho_{b_1} \otimes \cdots \otimes \rho_{b_k} \mid b_1,\ldots,b_k \in \mathfrak{G}^n\}\,.$$

It will be more convenient to enlarge the test class $\mathcal{F}$ to unitary valued functions as follows

$$\mathrm{CUT}^{\otimes k}_{\mathbb{U}_{s,k,\delta}} \coloneqq \{f_1 \otimes \cdots \otimes f_k \mid f_1,\ldots,f_k \colon [n] \to \mathbb{U}_{s,k,\delta}\}\,,$$

where $\mathbb{U}_{s,k,\delta}$ will be a fine enough discretization of the matrices[16] $M_s(\mathbb{C})$ of operator norm at most 1.

We will extend the framework to additionally handle the classes of functions $\mathrm{CUT}^{\otimes k}_{\omega,q,a}$ and $\mathrm{CUT}^{\otimes k}_{\mathbb{U}_{s,k,\delta}}$. This is proven in the full version. Let $\mathbb{K}$ be the underlying field which is either $\mathbb{R}$ or $\mathbb{C}$. Our extended framework gives the following efficient algorithmic result.

▶ **Theorem 11** (Efficient Weak Regularity (Extension of [24])). *Let $W \subseteq [n]^k$ be a τ-splittable collection of tuples. Suppose $\mathcal{F}$ is one of $\mathrm{CUT}^{\otimes k}$, $\mathrm{CUT}^{\otimes k}_{\pm}$, $\mathrm{CUT}^{\otimes k}_{\omega,q,a}$, for $q \geq 3$, or $\mathrm{CUT}^{\otimes k}_{\mathbb{U}_{s,k,\delta}}$. Let $\mathcal{R}$ be the domain of the functions in $\mathcal{F}$, when $k = 1$. Let $g \in \mathcal{R}^{W[1]^k}$ be supported on W with $\|g\|_{\mu_k} \leq 1$. For every $\delta > 0$, if $\tau \leq \delta^2/(k^3 \cdot 2^{20})$, then we can find $h = \sum_{\ell=1}^{p} c_\ell \cdot f_\ell$ with $p = O(k^2/\delta^2)$, scalars $c_1,\ldots,c_p \in \mathbb{K}$ and functions $f_1,\ldots,f_p \in \mathcal{F}$, such that $\|h\|_{\mu_1^{\otimes k}} \leq 2$, $\sum_{\ell=1}^{p} |c_\ell| = O(k/\delta)$ and h is a good approximator to g in the following sense*

$$\max_{f \in \mathcal{F}} \left| \left\langle g - \left(\frac{d}{n}\right)^{k-1} h, f \right\rangle \right| \;\leq\; \delta \cdot |W|\,,$$

where the inner product is over the counting measure on $W[1]^k$. Furthermore, h can be found in $\widetilde{O}(2^{|\mathcal{R}|^{\widetilde{O}(k^2/\delta^2)}} \cdot |W|)$ time in the scalar valued case and in time $\widetilde{O}_{s,k,\delta}(\mathrm{poly}(|W|))$, otherwise.

3.3 Improved Case: k-LIN over $\mathbb{Z}_q$

The goal of this section is to prove Theorem 3 (restated below) assuming the new extended efficient regularity algorithm from Theorem 11.

▶ **Theorem 3** (Main II). *Let $\mathfrak{I}$ be an instance of MAX k-LIN$_q$ on n variables with alphabet $\mathbb{Z}_q$ and constraints supported on a regular[17] collection of tuples $W \subseteq [n]^k$. If W is τ-splittable with $\tau \leq \tau_0(k,\delta) \coloneqq \mathrm{poly}(\delta/k)$, then we can compute an assignment satisfying $\mathsf{OPT} - \delta$ in time $r(q/\tau_0) \cdot \widetilde{O}(|W| + n)$, where $r(x) = \exp(\exp(\mathrm{poly}(x)))$.*

[16] We use $M_s(\mathbb{C})$ for the set of $s \times s$ matrices over $\mathbb{C}$.
[17] This is an analog to tuples of a graph being d-vertex regular.

For k-LIN over alphabet $\mathbb{Z}_q$, we are given a collection of equations (each variable appearing with coefficient one) specified as collection of tuples $W \subseteq [n]^k$ and we are given a collection of corresponding RHS $(r_w)_{w\in W} \in \mathbb{Z}_q^W$. The system of linear equations can be written as follows

$$x_{i_1} + \cdots + x_{i_k} \;=\; r_w \pmod q \qquad\qquad \forall\, w = (i_1, \ldots, i_k) \in W.$$

Orthogonality of Fourier characters will be crucially used here.

▶ **Fact 12** (Character Orthogonality). *Let ω be a non-trivial q-th root of unit. Then,*

$$\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q} [\omega^a] = 0\,.$$

Orthogonality allows for a convenient way of implementing the Dirac delta function on (the alphabet) $\mathbb{Z}_q$.

▶ **Fact 13.** *Fix $y \in \mathbb{Z}_q$. The indicator function $x \mapsto \mathbf{1}_{[x=y]}$ on $\mathbb{Z}_q$ can be expressed as*

$$\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q} \left[\omega^{a(x-y)}\right].$$

We now make precise the argument sketched in the proof strategy of Section 2.

Proof of Theorem 3. Let $\mathfrak{I}$ be an instance of k-LIN over $\mathbb{Z}_q$ with constraints supported on $W \subseteq [n]^k$ and RHS values $\{r_w\}_{w\in W}$. For every $a \in \mathbb{Z}_q$, we define $g_a \colon W \to \mathbb{C}$ as the map $w \in W \mapsto \omega^{a\cdot r_w}$, where $\omega = \exp(2\pi\sqrt{-1}/q)$.

Apply the efficient weak regularity decomposition of Theorem 11 to each g_a using error parameter $\delta > 0$ and test functions $\mathcal{F} = \mathrm{CUT}^{\otimes k}_{\omega,q,a}$. Note that this requires the splittability (expansion) parameter τ of W to satisfy $\tau \le O(\delta^2/k^3)$. We obtain a function $h_a = \sum_{\ell=1}^{p_a} c_{a,\ell} \cdot \chi_{b^{(a,\ell,1)},a} \otimes \cdots \otimes \chi_{b^{(a,\ell,k)},a}$, where $b^{(a,\ell,1)}, \ldots, b^{(a,\ell,k)} \in \mathbb{Z}_q^n$, for every $a \in \mathbb{Z}_q$ and every $\ell \in [p_a]$. Let $b \in \mathbb{Z}_q^n$ be an assignment to the variables of the system of linear equations. The value of this CSP on input b can be computed as

$$\begin{aligned}
\mathsf{val}(\mathfrak{I}, b) = \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\langle g_a, \chi_{b,a}\otimes\cdots\otimes\chi_{b,a}\rangle_{\mu_k}\right] &= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k}\left[\omega^{-a\cdot r_w}\omega^{a(b_{i_1}+\cdots+b_{i_k})}\right]\right] \\
&= \mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k}\left[\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\omega^{a(b_{i_1}+\cdots+b_{i_k}-r_w)}\right]\right] \\
&= \mathop{\mathbb{E}}_{w=(i_1,\ldots,i_k)\sim\mu_k}\left[\mathbf{1}_{[b_{i_1}+\cdots+b_{i_k}=r_w]}\right].
\end{aligned}$$

Using the weak regularity decomposition h_a of each g_a, we obtain

$$\begin{aligned}
\mathsf{val}(\mathfrak{I}, b) &= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\langle g_a, \chi_{b,a}\otimes\cdots\otimes\chi_{b,a}\rangle_{\mu_k}\right] \\
&= \frac{1}{|W|}\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\left\langle\left(\frac{d}{n}\right)^{k-1} h_a, \chi_{b,a}\otimes\cdots\otimes\chi_{b,a}\right\rangle\right] \;\pm\; \delta \\
&= \frac{1}{n^k}\mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\sum_{\ell=1}^{p_a} c_{a,\ell}\cdot\left\langle\chi_{b^{(a,\ell,1)},a}\otimes\cdots\otimes\chi_{b^{(a,\ell,k)},a}, \chi_{b,a}\otimes\cdots\otimes\chi_{b,a}\right\rangle\right] \;\pm\; \delta \\
&= \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\sum_{\ell=1}^{p_a} c_{a,\ell}\cdot\left\langle\chi_{b^{(a,\ell,1)},a}, \chi_{b,a}\right\rangle_{\mu_1}\cdots\left\langle\chi_{b^{(a,\ell,k)},a}, \chi_{b,a}\right\rangle_{\mu_1}\right] \;\pm\; \delta\,,
\end{aligned}$$

concluding the proof.

Now it suffices to approximate the following values

$$\left\{\left\langle \chi_{b^{(a,\ell,j)},a}, \chi_{b,a}\right\rangle_{\mu_1}\right\}_{a\in\mathbb{Z}_q,\ell\in[p_a],j\in[k]},$$

so that there is always a true assignment $b \in [q]^n$ which gives these values.

To this end, we first define the following collection $\mathcal{F}'$ of functions

$$\mathcal{F}' = \left\{\chi_{b^{(a,\ell,j)},a}\right\}_{a\in\mathbb{Z}_q,\ell\in[p_a],j\in[k]}.$$

Note that $r = |\mathcal{F}'| = O(qk^3/\delta^2)$. The functions above have range $\mathcal{R} = \{\omega^{a'} \mid a' \in \mathbb{Z}_q\}$. They form a factor $\mathcal{B}$ with at most $|\mathcal{R}|^r$ atoms. By the definition of a factor, the functions $\mathcal{F}'$ are constant in each one of them, so to compute $\left\langle \chi_{b^{(a,\ell,j)},a}, \chi_{b,a}\right\rangle_{\mu_1}$ it suffices to know the distribution of symbols of b in each atom.

Let $L = \sum_{a\in\mathbb{Z}_q,\ell\in[p_a]} |c_{a,\ell}| = O(qk/\delta)$ and set $\eta = \delta/(k\cdot L\cdot q)$. The total number of η-approximate distributions in ℓ_1-norm on each atom can be (crudely) bounded as

$$(1/\eta q)^{|\mathcal{R}|^r} \le \exp(\exp(\mathrm{poly}(qk/\delta))).$$

Using these distributions, we can approximate

$$\mathsf{val}(\mathfrak{I}, b) \;=\; \mathop{\mathbb{E}}_{a\in\mathbb{Z}_q}\left[\sum_{\ell=1}^{p_a} c_{a,\ell}\cdot\left\langle \chi_{b^{(a,\ell,1)},a}, \chi_{b,a}\right\rangle_{\mu_1}\cdots\left\langle \chi_{b^{(a,\ell,k)},a}, \chi_{b,a}\right\rangle_{\mu_1}\right] \;\pm\; \delta\,,$$

incurring an additional error of δ. ◀

References

1. Vedat Levi Alev, Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. List decoding of direct sum codes. In *Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms*, pages 1412–1425. SIAM, 2020.
2. Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani. Approximating constraint satisfaction problems on high-dimensional expanders. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science*, pages 180–201, 2019.
3. Noga Alon and Assaf Naor. Approximating the cut-norm via grothendieck's inequality. In *Proceedings of the 36th ACM Symposium on Theory of Computing*, pages 72–80, 2004.
4. Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *Proceedings of the 39th ACM Symposium on Theory of Computing*, STOC '07, pages 227–236, 2007.
5. Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science*, pages 472–481, 2011. doi:10.1109/FOCS.2011.95.
6. Guy Blanc and Dean Doron. New near-linear time decodable codes closer to the GV bound. Technical Report TR22-027, Electronic Colloquium on Computational Complexity, 2022.
7. Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science*, 2019.
8. Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta-Shma. List decoding with double samplers. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms*, pages 2134–2153, 2019.
9. Noam D. Elkies. Excellent codes from modular curves. In *Proceedings of the 33rd ACM Symposium on Theory of Computing*, 2001.
10. Uriel Feige and Joe Kilian. Zero knowledge and the chromatic number. *Journal of Computer and System Sciences*, 57(2):187–199, 1998.

11 David Forney. *Concatenated Codes*. PhD thesis, MIT, 1966.

12 A. Frieze and R. Kannan. The regularity lemma and approximation schemes for dense problems. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, 1996.

13 E.N. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31:504–522, 1952.

14 Sivakanth Gopi, Swastik Kopparty, Rafael Oliveira, Noga Ron-Zewi, and Shubhangi Saraf. Locally testable and locally correctable codes approaching the Gilbert-Varshamov bound. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 2073–2091, 2017.

15 Zeyu Guo and Noga Ron-Zewi. Efficient list-decoding with constant alphabet and list sizes. *IEEE Transactions on Information Theory*, 68(3):1663–1682, 2022.

16 Venkatesan Guruswami and Piotr Indyk. Efficiently decodable codes meeting Gilbert-Varshamov bound for low rates. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, SODA '04, pages 756–757, 2004.

17 Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. Essential coding theory. Available at `https://cse.buffalo.edu/faculty/atri/courses/coding-theory/book/index.html`, 2019.

18 Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with psd objectives. In *FOCS*, pages 482–491, 2011.

19 Venkatesan Guruswami and Ali Kemal Sinop. Faster SDP hierarchy solvers for local rounding algorithms. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*, pages 197–206. IEEE, 2012.

20 Venkatesan Guruswami and Chaoping Xing. List decoding reed-solomon, algebraic-geometric, and gabidulin subcodes up to the singleton bound. In *Proceedings of the 45th ACM Symposium on Theory of Computing*, 2013.

21 B. Hemenway, N. Ron-Zewi, and M. Wootters. Local list recovery of high-rate tensor codes applications. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, pages 204–215, October 2017.

22 Akhil Jalan and Dana Moshkovitz. Near-optimal Cayley expanders for Abelian groups, 2021. `arXiv:2105.01149`.

23 Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. Unique decoding of explicit ε-balanced codes near the Gilbert–Varshamov bound. In *Proceedings of the 61st IEEE Symposium on Foundations of Computer Science*, 2020.

24 Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani. Near-linear time decoding of Ta-Shma's codes via splittable regularity. In *Proceedings of the 52nd ACM Symposium on Theory of Computing*, 2021.

25 Swastik Kopparty, Nicolas Resch, Noga Ron-Zewi, Shubhangi Saraf, and Shashwat Silas. On list recovery of high-rate tensor codes. *IEEE Transactions on Information Theory*, 67(1):296–316, 2021. `doi:10.1109/TIT.2020.3023962`.

26 R. McEliece, E. Rodemich, H. Rumsey, and L. Welch. New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities. *IEEE Transactions on Information Theory*, 23(2):157–166, 1977.

27 Anand Kumar Narayanan and Matthew Weidner. Subquadratic time encodable codes beating the gilbert–varshamov bound. *IEEE Transactions on Information Theory*, 65(10), 2019.

28 Shayan Oveis Gharan and Luca Trevisan. A new regularity lemma and faster approximation algorithms for low threshold rank graphs. *Theory of Computing*, 11(9):241–256, 2015. `doi:10.4086/toc.2015.v011a009`.

29 Anthony Man-Cho So, Jiawei Zhang, and Yinyu Ye. On approximating complex quadratic optimization problems via semidefinite programming relaxations. *Math. Program.*, 110(1):93–110, June 2007.

30 Henning Stichtenoth. *Algebraic Function Fields and Codes*. Springer Publishing Company, Incorporated, 2nd edition, 2008.

31 Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, STOC 2017, pages 238–251, New York, NY, USA, 2017. ACM.

32 C. Thommesen. The existence of binary linear concatenated codes with Reed- Solomon outer codes which asymptotically meet the Gilbert-Varshamov bound. *IEEE Transactions on Information Theory*, 29(6):850–853, November 1983.

33 L. Trevisan, M. Tulsiani, and S. Vadhan. Boosting, regularity and efficiently simulating every high-entropy distribution. In *Proceedings of the 24th IEEE Conference on Computational Complexity*, 2009.

34 Michael Tsfasman, Serge Vladut, and Dmitry Nogin. *Algebraic Geometric Codes: Basic Notions*. American Mathematical Society, 2007.

35 Michael A. Tsfasman, S. G. Vlădut, and Thomas Zink. Modular curves, shimura curves, and goppa codes, better than varshamov-gilbert bound. *Mathematische Nachrichten*, 109:21–28, 1982.

36 R.R. Varshamov. Estimate of the number of signals in error correcting codes. *Doklady Akademii Nauk SSSR*, 117:739–741, 1957.