

Decodability of Group Homomorphisms beyond the Johnson Bound

Irit Dinur^{*}

Elena Grigorescu[†]

Swastik Kopparty[‡]

Madhu Sudan[§]

ABSTRACT

Given a pair of finite groups G and H , the set of homomorphisms from G to H form an error-correcting code where codewords differ in at least $1/2$ the coordinates. We show that for every pair of *abelian* groups G and H , the resulting code is (locally) list-decodable from a fraction of errors arbitrarily close to its distance. At the heart of this result is the following combinatorial result: There is a fixed polynomial $p(\cdot)$ such that for every pair of abelian groups G and H , if the maximum fraction of agreement between two distinct homomorphisms from G to H is Λ , then for every $\epsilon > 0$ and every function $f : G \rightarrow H$, the number of homomorphisms that have agreement $\Lambda + \epsilon$ with f is at most $p(1/\epsilon)$.

We thus give a broad class of codes whose list-decoding radius exceeds the “Johnson bound”. Examples of such codes are rare in the literature, and for the ones that do exist, “combinatorial” techniques to analyze their list-decodability are limited. Our work is an attempt to add to the body of such techniques. We use the fact that abelian groups decompose into simpler ones and thus codes derived from homomorphisms over abelian groups may be viewed as certain “compositions” of simpler codes. We give techniques to lift list-decoding bounds for the component codes to bounds for the composed code. We believe these techniques may be of general interest.

^{*}Weizmann Institute of Science, Rehovot, Israel.
email: irit.dinur@weizmann.ac.il

[†]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.
email: elena_g@mit.edu. Supported in part by NSF Award CCR-0514915.

[‡]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.
email: swastik@mit.edu Supported in part by NSF Award CCR-0514915.

[§]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.
email: madhu@mit.edu Supported in part by NSF Award CCR-0514915.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC’08, May 17–20, 2008, Victoria, British Columbia, Canada.
Copyright 2008 ACM 978-1-60558-047-0/08/05 ...\$5.00.

Categories and Subject Descriptors

E.4 [Coding and Information Theory]: Error Control Codes; F.0 [Theory of Computation]: General

General Terms

Theory

Keywords

List Decoding, Sublinear Time Algorithms, Hadamard Codes

1. INTRODUCTION

For groups $(G, +)$ and $(H, +)$, a *homomorphism* from G to H is a function $\phi : G \rightarrow H$ such that $\phi(x) + \phi(y) = \phi(x + y)$ for all $x, y \in G$. It is a well-known fact that distinct homomorphisms between a fixed pair of groups G and H are *far* from each other in the Hamming norm. Thus homomorphisms between groups form “error-correcting codes”, giving rise to the question: How well can these codes be list-decoded? In this paper we give new, strong, combinatorial and algorithmic results for list-decoding of codes derived from group homomorphisms, for abelian groups.

The classical example of a code based on group homomorphisms are the Hadamard codes, which can be viewed as homomorphisms from \mathbb{Z}_2^n to \mathbb{Z}_2 . The list-decoding problem for this class of groups was considered in the seminal paper of Goldreich and Levin [4], where they used this result to get a generic construction of a hardcore predicates for any one-way function. Subsequently this result has formed the basis of many results in learning, average-case complexity and cryptography.

In terms of technical generalizations, the work of Goldreich-Levin has led to the study of list-decoding for polynomial functions over finite fields (see, for instance, [5, 13]), as also the task of learning Fourier coefficients ([9, 3, 1]). Again both these directions have led to important applications in pseudorandomness, cryptography, and average-case complexity.

Our work is motivated by yet a third interpretation of the Goldreich-Levin result, as a list-decoder for homomorphisms. This interpretation was already studied in a previous work by some of the authors [6], where it was shown that the algorithmic techniques of, say [4, 5, 13], can typically be extended to the case of group homomorphisms also, provided one can get convincingly strong combinatorial bounds on the “list-decodability” of these codes. Unfortunately, bounds on the list-decodability of these codes were too weak, and indeed it seems there are few techniques to bound the list-

decodability of error-correcting codes. Our paper is motivated mainly by this need to augment the analytic tools in this setting. In what follows we attempt to clarify the problem, the prior state of knowledge, and our contributions.

Combinatorics of List-decoding Homomorphisms: Let G, H be finite groups and let $\Lambda = \Lambda_{G,H}$ denote the maximum relative agreement between any pair of homomorphisms between G and H . Thus the class of homomorphisms between G and H may be viewed as error-correcting codes over the alphabet $\Sigma = H$ of length $N = |G|$ and relative (Hamming) distance $1 - \Lambda$. The *combinatorial* question at the heart of this paper is the following:

Given a function $f : G \rightarrow H$ and a real number $\epsilon > 0$, what is the number of homomorphisms ϕ such that f and ϕ agree in at least $\Lambda_{G,H} + \epsilon$ fraction of inputs?

As we argue below this simple question was vastly understood, at least till this work.

For the sake of concreteness, let p be a prime and n, m be positive integers. Now consider the class of homomorphisms between $G = \mathbb{Z}_p^n$ and $H = \mathbb{Z}_p^m$. Then $\Lambda = \Lambda_{G,H} = 1/p$. Our combinatorial question asks: How many homomorphisms can have agreement $1/p + \epsilon$ with any given function $f : \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^m$. In the case of $m = 1$, the works of [4, 5], assert that the number of homomorphisms is at most $O(1/\epsilon^2)$. Such a result can also be shown to be essentially tight, for any choice of p, n, m , and ϵ . These results essentially follow from the Johnson bound for codes over a p -ary alphabet (see, for instance, [7]).

What about the case of general m ? Here, our state of knowledge tapers off quickly. Since the codes no longer have relative distance $1 - 1/|\Sigma|$, the Johnson bound is no longer strong and only gives bound for the number of homomorphisms that have agreement more than $\sqrt{1/p}$ with the given function. In other words this bound requires $\epsilon > \sqrt{1/p} - 1/p$ and does not work for every $\epsilon > 0$. A more ad-hoc analysis, obtained by viewing a homomorphism from \mathbb{Z}_p^n to \mathbb{Z}_p^m as a tuple of m independent homomorphisms from \mathbb{Z}_p^n to \mathbb{Z}_p , gives a bound of $1/\epsilon^{2m}$ on the number of homomorphisms. Till our work it was unclear if this exponential dependence on m was necessary.

A similar state of affairs is also seen when we consider, say, homomorphisms from \mathbb{Z}_p^n to \mathbb{Z}_p^m . In this case, Fourier analysis can be used to get some bounds on the number of homomorphisms that have agreement $1/p + \epsilon$ with a fixed function [6], but again this bound has exponential dependence on m .

In this paper we fix these gaps in our knowledge. We show that there is a fixed polynomial g such that for all abelian groups G, H and for every function $f : G \rightarrow H$ and every $\epsilon > 0$, the number of homomorphisms with agreement $\Lambda_{G,H} + \epsilon$ with the fixed function f is at most $g(\epsilon)$. (The bound is thus a fixed polynomial, independent of G and H . See Theorem 2.2.)

Algorithmic Consequences: Our combinatorial result turns immediately into efficient, local, algorithms for list-decoding homomorphisms. Specifically, we show that the techniques from [4, 5, 13, 6] immediately yield an algorithm that behaves as follows: It takes as input an explicit description of groups G and H (in terms of their prime decomposition) and a real number $\epsilon > 0$, and has oracle access to a function $f : G \rightarrow H$. In time polynomial in $1/\epsilon$,

$\log |G|$ and $\log |H|$ it outputs an explicit description of all homomorphisms that have agreement $\Lambda + \epsilon$ with f .

We remark that a polynomial dependence on each is necessary to represent the output. In previous work [6] considered a fixed group H and for every such group, gave a polynomial time algorithm (i.e., its running time was polynomial in $\log |G|$ and $\frac{1}{\epsilon}$) to recover the nearby homomorphisms. Their running time grew exponentially in $\log |H|$.

Techniques: We stress that despite the widespread use of list-decoding, general techniques to establish bounds are lacking. There a relatively few codes known whose list-decodability goes beyond the Johnson bound: some examples include the Ta-Shma-Zuckerman codes [14], and the Parvaresh-Vardy-Guruswami-Rudra codes [12, 8]. The latter codes are similar in the sense that they also consider codes over an alphabet that is a vector space (rather than a field), but in their case the only proof of the list-decoding properties is algorithmic. We feel this is partly due to the lack of combinatorial techniques to analyze the decoding properties of such codes.

We start by noticing that any abelian group, in particular, H can be decomposed as $H_1 \times \dots \times H_k$ where $H_i = \mathbb{Z}_{p_i}^{e_i}$ for some prime p_i and positive integer e_i . Thus a homomorphism ϕ from G to H can be decomposed into simpler ones using this decomposition. For instance if $H = H_1 \times H_2$ then a homomorphism ϕ from G to H is just a pair of homomorphisms (ϕ_1, ϕ_2) where $\phi_i : G \rightarrow H_i$ for $i \in \{1, 2\}$. This decomposition allows for some weak list-decoding bounds. To see this, fix some agreement parameter ρ and lets consider the number of homomorphisms that have agreement ρ with some function $f : G \rightarrow H$. Suppose we have a bound B_i on the number of homomorphisms from G to H_i that have agreement ρ with any single function. Then, viewing f as a pair of functions (f_1, f_2) , $f_i : G \rightarrow H_i$, and noticing that for any homomorphism $\phi = (\phi_1, \phi_2) : G \rightarrow H$ that has agreement ρ with f it must be the case that ϕ_i and f_i also have agreement at least ρ , we conclude that the number of homomorphisms that have agreement ρ with f is at most $B_1 \cdot B_2$. Unfortunately, this bound grows exponentially with $\log |H|$ and this is too weak to get our claimed theorem.

Our improvements come by noticing that agreements between homomorphisms are quite restricted. For instance, in the case where $H_1 = H_2 = \mathbb{Z}_p$ and $H = H_1 \times H_2$, homomorphisms $\phi, \psi : G \rightarrow H_1$ either agree on a set of density $\frac{1}{p}$ or a set of density at most $\frac{1}{p^2}$. To leverage this insight in our setting, say we have a function $f = (f_1, f_2)$ and we wish to bound homomorphisms with agreement ρ with f . Suppose $\phi_1 : G \rightarrow H_1$ has agreement at least ρ with f_1 and suppose the set of agreements is the set $S \subseteq G$. We now consider all homomorphisms ψ_1, \dots, ψ_ℓ that agree with f_2 on subsets of S of size at least $\rho \cdot |G|$. We consider agreements of “triples” of homomorphisms ψ_i, ψ_j and ψ_k and note that if every such triple has a large mutual agreement in the set S , then they satisfy a *sunflower* like property, forcing S to be very large. If on the other hand every triple has a small agreement then we manage to show that S is still large, by using an inclusion-exclusion count. We then use a standard argument to show that either the given collection of homomorphisms always contain a large subcollection in which every triple has a large intersection, or a large subcollection in which every triple has a small intersection. This gives us a nontrivial lower bound on S in all cases and then we perform some very careful accounting to show that this

leads to an (absolute) polynomial upper bound on the list size.

While all these arguments may seem very specific to this special case of homomorphism from G to $\mathbb{Z}_p \times \mathbb{Z}_p$, we show that this is not the case. Indeed we abstract a nice property of families of sets (in our case the points of agreement between f_2 and ψ_1, \dots, ψ_k within a specified set S) that allows us to give non-trivial bounds on the size of their union. We then apply this same bound in three different cases to lift list-decoding bounds for somewhat simple groups to list decoding bounds for more complex groups. We start with homomorphisms to \mathbb{Z}_p and lift them to homomorphisms to \mathbb{Z}_p^r , then to \mathbb{Z}_{p^r} and then to arbitrary groups $\prod_i \mathbb{Z}_{p_i}^{e_i}$. This leads us to the final theorem claimed above.

The combinatorial results turn into algorithmic results in a straightforward manner based on previous works [4, 5, 13, 6]. If anything, our algorithms become even simpler because our combinatorial bounds are stronger. See Section 5.

Related prior works: The task of list-decoding homomorphisms is very closely related to the task of determining significant Fourier coefficients in abelian groups. Previous works [9, 11, 3, 1] have consider the latter task in different contexts and in particular Akavia et al. [1] give comprehensive results on “list-decoding” for the significant Fourier coefficients. Here we point out why their results do not subsume ours. Even though significant Fourier coefficients correspond exactly to list decoding when $H = \mathbb{Z}_2$ or $H = \mathbb{Z}_3$, the situation changes significantly over other groups. The complex inner product, the norm underlying Fourier coefficients, is very different from the Hamming distance and indeed it is possible to find functions with few large Fourier coefficients that are not close to any homomorphism in the Hamming norm. Conversely, it is also possible to find functions that are close to homomorphisms in the Hamming norm and have no significant Fourier coefficients. Thus in these cases, list-decoding is a very different problem than that of finding Fourier coefficients, and combinatorial bounds such as the ones we provide were not known.

Organization of this paper: In Section 2 we introduce some necessary definitions and state our main results. In Section 3, we prove our main combinatorial bounds, modulo a theorem about certain set systems with restricted intersections. In Section 4, we prove this theorem. Finally, in Section 5, we give an outline of the local list decoding algorithm. Some proofs are omitted and appear in the full version of this paper [2].

2. DEFINITIONS AND MAIN RESULTS

We start with some basic terminology about homomorphisms, introduce the combinatorial and algorithmic list decoding problems and state our results.

For abelian groups G, H let $\text{Hom}(G, H) = \{h : G \rightarrow H \mid h(x) + h(y) = h(x + y), \forall x, y \in G\}$ be the set of homomorphisms from G to H . Also let $\text{aHom}(G, H) = \{h + a : G \rightarrow H \mid h \in \text{Hom}(G, H), a \in H\}$ be the set of affine homomorphisms from G to H . Notice that $\text{Hom}(G, H) \subset \text{aHom}(G, H)$.

For two functions $f, g : G \rightarrow H$, define $\text{agree}(f, g) = \Pr_{x \in G}[f(x) = g(x)]$, and

$$\Lambda_{G, H} = \max_{f, g \in \text{Hom}(G, H), f \neq g} \{\text{agree}(f, g)\}.$$

In the case when $\text{Hom}(G, H)$ contains only the zero homo-

morphism we define $\Lambda_{G, H} = 0$.

DEFINITION 2.1 (COMBINATORIAL LIST DECODABILITY). *The code $\text{aHom}(G, H)$ is (δ, l) -list decodable if for every function $f : G \rightarrow H$, there exist at most l homomorphisms $h \in \text{aHom}(G, H)$ such that $\text{agree}(f, h) \geq \delta$.*

The principal question that we address is: For which function $l(\delta)$ can we conclude that $\text{aHom}(G, H)$ is $(\delta, l(\delta))$ -list-decodable. Our theorem below shows that this is true for $l(\delta) = \text{poly}((\delta - \Lambda_{G, H})^{-1})$ for some polynomial independent of G and H .

THEOREM 2.2. *There is a universal constant C such that the following holds: Let G, H be abelian groups. Then for every $\epsilon > 0$, $\text{aHom}(G, H)$ is $(\Lambda_{G, H} + \epsilon, 1/\epsilon^C)$ -list-decodable.*

We augment this combinatorial result with an algorithmic one that finds the nearby homomorphisms efficiently. We first define this algorithmic problem below and then state our algorithmic theorem.

Our algorithmic goal is to produce a list of all homomorphisms with agreement $\Lambda_{G, H} + \epsilon$ with some function $f : G \rightarrow H$ efficiently, i.e., in time $\text{poly}(\log |G|, \log |H|, \frac{1}{\epsilon})$. In order to do so, the algorithm requires oracle access to the function f (or else reading the function will take time $\text{poly}(|G|)$). Our solution requires a further assumption that the groups G and H are given explicitly, in a sense described next.

Recall that the structure theorem for finite abelian groups (see Theorem 8.2 in Lang [10]) states that every finite abelian group can be decomposed into cyclic groups of the form $\mathbb{Z}_{p_1^{r_1}} \times \dots \times \mathbb{Z}_{p_k^{r_k}}$, where all p_i 's are prime. We will assume that our input groups G and H are presented in this cyclic decomposition. Our algorithm returns all such homomorphisms explicitly by specifying their values on a set of generators of G .

DEFINITION 2.3 (ALGORITHMIC LOCAL LIST DECODING).

A probabilistic oracle algorithm \mathcal{A} for list decoding homomorphisms takes as input two groups G and H represented explicitly, and a parameter $\delta > 0$ (where the explicit representation of $G = \mathbb{Z}_{p_1^{r_1}} \times \dots \times \mathbb{Z}_{p_k^{r_k}}$ is the sequence of pairs $\langle (p_1, r_1), \dots, (p_k, r_k) \rangle$, and similarly for H) and has oracle access to a function $f : G \rightarrow H$. We say that \mathcal{A} is a (δ, T) -local list decoder for groups G, H , if for every function $f : G \rightarrow H$, \mathcal{A}^f runs in time T and outputs a list $\{\phi_1, \phi_2, \dots, \phi_L\} \subset \text{aHom}(G, H)$ s.t. with probability at least $3/4$, it is the case that for every affine homomorphism $h \in \text{aHom}(G, H)$ such that $\text{agree}(f, h) \geq \delta$, $h = \phi_j$ for some $j \in [L]$.

THEOREM 2.4. *There exists an algorithm \mathcal{A} , such that for every pair of abelian groups G, H , and every $\epsilon > 0$, \mathcal{A} is a $(\Lambda_{G, H} + \epsilon, \text{poly}(\log |G|, \log |H|, \frac{1}{\epsilon}))$ -local list decoder for $\text{aHom}(G, H)$.*

3. ANALYSIS OF COMBINATORIAL LIST DECODABILITY

In this section we prove Theorem 2.2 above, which upper bounds the number of homomorphisms that have agreement $\Lambda_{G, H} + \epsilon$, for some $\epsilon > 0$, with a given function $f : G \rightarrow H$.

We begin by considering two special cases. In both special cases G is a somewhat special group, which we will elaborate

on later. In the first case $H = \mathbb{Z}_p^r$, while in the second case $H = \mathbb{Z}_p^r$. Finally we analyze the case of general G and H , by reducing it to the special cases.

All three settings (the two special cases and the general case) are analyzed by a common technique by considering set systems with very special intersection properties and bounding their cardinality (in a very specific way). In our analysis these sets will represent sets of agreement between a function at hand and some nearby homomorphisms. We define this special intersection property and state our crucial theorem about these families next.

3.1 Special Intersecting Families

Below we consider a collection S_1, \dots, S_ℓ of subsets of a universe X . For a set $T \subset X$, its density is denoted $\mu(T) \stackrel{\text{def}}{=} |T|/|X|$. For a set of indices $I \subseteq [\ell]$ let $S_I = \bigcap_{i \in I} S_i$ (if $I = \emptyset$ then $S_I = X$).

DEFINITION 3.1. *For $0 < \tau \leq \rho \leq 1$ and $c < \infty$ we say that a family of sets $S_1, \dots, S_\ell \subseteq X$ is a (ρ, τ, c) -special intersecting family if the following four properties hold:*

1. For every index $i \in [\ell]$, $\mu(S_i) \geq \rho$.
2. For every pair of distinct indices $i, j \in [\ell]$, $\mu(S_i \cap S_j) \leq \rho$.
3. $\sum_{i=1}^{\ell} \alpha_i^c \leq 1$, where $\alpha_i = \mu(S_i) - \rho$.
4. For subsets (of indices) $I, J \subseteq [\ell]$, with $J \subseteq I$ and $|J| \geq 2$, if $\mu(S_I) > \tau$ the $S_I = S_J$.

The first two properties above are typical conditions seen in intersecting systems, but the “restrictions” are not particularly restrictive. The third property is a typical condition that if derived when the sets S_i represent agreement sets of some fixed word with codewords of an orthogonal error-correcting code. (This condition is usually proved using Parseval’s identity, or using the Johnson bound, for $c = 2$.) The final condition is effectively “Helly”-like. It says, roughly, that the only way a large collection of sets $\{S_i\}_{i \in I}$ can have a non-trivially large intersection is if every pair of intersecting sets with indices from I have the same intersection.

Our main theorem about set systems with special intersections roughly shows that the density of the union of the S_i ’s is relatively large compared to the density of the sets themselves.

THEOREM 3.2. *For every $c < \infty$, there exists a $C < \infty$ such that the following holds: Let $\rho > 0$ and let $S_1, \dots, S_\ell \subseteq X$ be a (ρ, ρ^2, c) -special intersecting family. Let $\alpha_i = \mu(S_i) - \rho$ and let $\alpha = \mu(\bigcup_{i \in [\ell]} S_i) - \rho$. Then $\alpha^C \geq \sum_{i \in [\ell]} \alpha_i^c$.*

We defer the proof of Theorem 3.2 to the next section. We first show how it can be used to prove upper bounds on the list size when decoding affine homomorphisms.

3.2 Upper bound for $H = \mathbb{Z}_p^r$

In this case and the next one, we will be dealing with a special class of groups G that are called p -groups. Specifically, G is a p -group if $|G|$ is a power of a prime p . What we use below is that every subgroup of a p -group has cardinality at most $|G|/p$.

THEOREM 3.3. *There exists a constant C such that the following is true for every prime p and positive integer r . Let $H = \mathbb{Z}_p^r$. Given any function $f : G \rightarrow H$ and agreement parameter $\epsilon > 0$, the number of affine homomorphisms from G to H that have agreement at least $1/p + \epsilon$ with f is at most $(1/\epsilon)^C$.*

PROOF. Fix $f : G \rightarrow H$. Note that f can be viewed as an r -tuple of functions $f = \langle f_1, \dots, f_r \rangle$ with $f_i : G \rightarrow \mathbb{Z}_p$. For $k \in [r]$ define the k th projection of f to be the function $f^{[k]} : G \rightarrow \mathbb{Z}_p^k$ given by $f^{[k]}(x) = \langle f_1(x), \dots, f_k(x) \rangle$. We extend this notation to include $k = 0$, by allowing a unique function from $G \rightarrow \mathbb{Z}_p^0$ (which is considered to be an affine homomorphism).

For $k_1 \leq k_2$, we say that a function $f_2 : G \rightarrow \mathbb{Z}_p^{k_2}$ extends $f_1 : G \rightarrow \mathbb{Z}_p^{k_1}$ if $f_2 = f_1^{[k_2]}$.

In what follows, we fix k , $0 \leq k < r$ and consider the function $f^{[k]}$. We also fix an affine homomorphism ϕ from G to \mathbb{Z}_p^k that has agreement $\frac{1}{p} + \alpha$ with $f^{[k]}$, for some $\alpha \geq \epsilon$. We prove that the number of affine homomorphisms that extend ϕ and have agreement $1/p + \epsilon$ with f is small.

CLAIM 3.4. *For every $0 \leq k \leq r$, and for every affine homomorphism ϕ that has agreement $1/p + \alpha$ with $f^{[k]}$, the number of affine homomorphisms from G to \mathbb{Z}_p^r that have agreement $1/p + \epsilon$ with f and extend ϕ , is at most $(\alpha/\epsilon)^C$, where C is the constant from Theorem 3.2, for $c = 2$.*

PROOF. We prove the claim by induction on $r - k$. When $r - k = 0$, the claim is trivially true. So assume the claim is true for $r - (k + 1)$ and we will now prove it for the case of $r - k$.

Let $S = \{x \in G \mid f^{[k]}(x) = \phi(x)\}$. Consider $f^{[k+1]}$ and let ϕ_1, \dots, ϕ_ℓ be all the affine homomorphisms from G to \mathbb{Z}_p^{k+1} that extend ϕ with agreement at least $\frac{1}{p} + \epsilon$ with $f^{[k+1]}$. Let $S_i = \{x \mid f^{[k+1]}(x) = \phi_i(x)\}$ and let $\mu(S_i) = \frac{1}{p} + \alpha_i$.

We prove below that the sets $S_1, \dots, S_\ell \subseteq G$ form a $((1/p), (1/p^2), 2)$ -special intersecting family. This will allow us to apply Theorem 3.2 to these sets and this in turn will allow us to derive the inductive claim easily.

Verifying Intersection Properties: The first property requires that $\mu(S_i) \geq \frac{1}{p}$ which is immediate from the definition of the S_i ’s. The second property required that $\mu(S_i \cap S_j) \leq \frac{1}{p}$. This again is true, since $S_i \cap S_j \subseteq \{x \mid \phi_i(x) = \phi_j(x)\}$ and the latter has density at most $1/p$.

The third property follows from the “ p -ary Johnson bound” applied to the function $f_{k+1}^{[k+1]} : G \rightarrow \mathbb{Z}_p$. Let $(\phi_i)_j$ denote the j th coordinate of the function ϕ_i . Then note that the functions $(\phi_i)_{k+1}$ are distinct affine homomorphisms, with $(\phi_i)_{k+1}$ having agreement at least $1/p + \alpha_i$ with f_{k+1} . It now follows from the Johnson bound that $\sum_{i=1}^{\ell} \alpha_i^2 \leq 1$.

Finally, we come to the crucial “special-intersection” property. We need to show that if $\mu(S_I) > 1/p^2$, and $J \subseteq I$ contains at least two distinct elements, then $S_I = S_J$. For any set $K \subseteq [\ell]$ define $T_K = \{x \mid \phi_i(x) = \phi_j(x), \forall i, j \in K\}$. Note that $T_K = G$ if and only if $|K| \leq 1$. Since T_K is a coset of a subgroup of G , we have that $\mu(T_K) = 1/p^{|K|}$ for non-negative integer i . But $S_K \subseteq T_K$ for every K , and since $\mu(S_I), \mu(S_J) > 1/p^2$, it must be that $\mu(T_I), \mu(T_J) = 1/p$. Since $T_I \subseteq T_J$, we get that $T_I = T_J$. Finally, fix some $j \in J$ and let $Y = \{x \mid \phi_j(x) = f^{[k+1]}(x)\}$. Since $\phi_i(x) = \phi_j(x)$ for every $x \in Y \cap T_I$ and $i \in I$, we have $S_I = Y \cap T_I = Y \cap T_J = S_J$ as required for the fourth property.

Proof of Claim: We are now ready to prove the claim. Since the sets S_1, \dots, S_ℓ form a $(1/p, 1/p^2, 2)$ -special intersecting family, we can apply Theorem 3.2 to conclude that $(\alpha')^C \geq \sum_{i=1}^\ell \alpha_i^C$ where $\alpha' = \mu(\cup_{i \in \ell} S_i) - 1/p$. Since for every i , $S_i \subseteq S$, we have $\cup_{i \in \ell} S_i \subseteq S$ and so $\alpha \geq \alpha'$. We may thus conclude that $\alpha^C \geq \sum_{i=1}^\ell \alpha_i^C$. But this immediately leads to the claim as follows. By induction we have that the number of extensions of ϕ_i in $\text{aHom}(G, H)$ that have $\frac{1}{p} + \epsilon$ agreement with f is at most $(\alpha_i/\epsilon)^C$. Since every affine homomorphism in $\text{aHom}(G, H)$ that extends ϕ and has $\frac{1}{p} + \epsilon$ agreement with f must also extend one of the ϕ_i , we have that the total number of such extensions (of ϕ) is at most $\sum_{i \in [\ell]} (\alpha_i/\epsilon)^C \leq (\alpha/\epsilon)^C$ as desired. \square

The theorem statement follows immediately, by using $k = 0$.

3.3 Upper bound for $H = \mathbb{Z}_{p^r}$

We now consider a variant of Theorem 3.3 with $H = \mathbb{Z}_{p^r}$, the cyclic group of integers modulo a prime power (as opposed to the vector space $H = \mathbb{Z}_p^r$).

THEOREM 3.5. *There exists a constant C such that the following is true for every prime p and positive integer r . Let $H = \mathbb{Z}_{p^r}$. Given any function $f : G \rightarrow H$ and agreement parameter $\epsilon > 0$, the number of affine homomorphisms from G to H that have agreement at least $1/p + \epsilon$ with f is at most $(1/\epsilon)^C$. Furthermore, if $\psi_1, \dots, \psi_L : G \rightarrow H$ are affine homomorphisms with $\text{agree}(f, \psi_i) = 1/p + \beta_i \geq 1/p + \epsilon$, then $\sum_{i=1}^L \beta_i^C \leq 1$.*

PROOF. For any function $g : G \rightarrow \mathbb{Z}_{p^r}$ and $i \in [r]$, let $g^{(i)} : G \rightarrow \mathbb{Z}_{p^i}$ be given by $g^{(i)}(x) = g(x) \bmod (p^i)$. Also, let $g^{(-i)} : G \rightarrow \mathbb{Z}_{p^{r-i}}$ be such that ¹ for all $x \in G$, $g(x) = p^i \cdot g^{(-i)}(x) + g^{(i)}(x)$. $g^{(-i)}$ is clearly well defined. Note that if g is an affine homomorphism, then $g^{(i)}$ is also an affine homomorphism, while $g^{(-i)}$, in general, need not be one. For $j \leq i$, we say that $g : G \rightarrow \mathbb{Z}_{p^i}$ extends $h : G \rightarrow \mathbb{Z}_{p^j}$ if $g^{(j)} = h$.

Now fix $f : G \rightarrow H$. Fix $k \in \{0, \dots, r\}$. Let $\phi : G \rightarrow \mathbb{Z}_{p^k}$ be an affine homomorphism with agreement $1/p + \alpha \geq 1/p + \epsilon$ with $f^{(k)}$. (As usual, we assume there is a single function from G to \mathbb{Z}_{p^0} and that it is a homomorphism.) We claim that the number of affine homomorphisms extending ϕ is small.

CLAIM 3.6. *For every $0 \leq k \leq r$, and for every affine homomorphism ϕ that has agreement $1/p + \alpha$ with $f^{(k)}$, the number of affine homomorphisms from G to \mathbb{Z}_{p^r} that have agreement $1/p + \epsilon$ with f and extend ϕ , is at most $(\alpha/\epsilon)^C$, where C is the constant from Theorem 3.2, for $c = 2$.*

PROOF. The proof is identical to that of Claim 3.4 with mainly notational changes switching $f^{[1]}$ to $f^{(\cdot)}$. The only noticeable change is in the argument for the third property of ‘‘special intersections’’.

We prove the claim by induction on $r - k$. When $r - k = 0$, the claim is trivially true. So assume the claim is true for $r - (k + 1)$ and we will now prove it for the case of $r - k$.

¹Here we abuse notation and identify an element of $\mathbb{Z}_{p^{r-i}}$ with the least non-negative integer in its residue class mod p^{r-i} .

Let $S = \{x \in G \mid f^{(k)}(x) = \phi(x)\}$. Consider $f^{(k+1)}$ and let ϕ_1, \dots, ϕ_ℓ be all the affine homomorphisms from G to $\mathbb{Z}_{p^{k+1}}$ that extend ϕ with agreement at least $\frac{1}{p} + \epsilon$ with $f^{(k+1)}$. Let $S_i = \{x \mid f^{(k+1)}(x) = \phi_i(x)\}$ and let $\mu(S_i) = \frac{1}{p} + \alpha_i$.

We prove below that the sets $S_1, \dots, S_\ell \subseteq G$ form a $((1/p), (1/p^2), 2)$ -special intersecting family. This will allow us to apply Theorem 3.2 to these sets and this in turn will allow us to derive the inductive claim easily.

Verifying Intersection Properties: The first property requires that $\mu(S_i) \geq \frac{1}{p}$ which is immediate from the definition of the S_i 's. The second property required that $\mu(S_i \cap S_j) \leq \frac{1}{p}$. This again is true, since $S_i \cap S_j \subseteq \{x \mid \phi_i(x) = \phi_j(x)\}$ and the latter has density at most $1/p$.

The third property follows from the ‘‘ p -ary Johnson bound’’ applied to the function $g = (f^{(k+1)})^{(-k)} : G \rightarrow \mathbb{Z}_p$. Let $\psi_i = \phi_i^{(-k)} : G \rightarrow \mathbb{Z}_p$ (we warn the reader that ψ_i need not be an affine homomorphism). Note that the agreement set of ψ_i with g contains S_i . Furthermore, note that for $i \neq j$, $\psi_i \neq \psi_j$ and $\text{agree}(\psi_i, \psi_j) = \text{agree}(\psi_i - \psi_j, 0) = \text{agree}(\phi_i - \phi_j, 0) \leq \frac{1}{p}$ (since $\phi_i = p^k \cdot \psi_i + \phi$, $\phi_j = p^k \cdot \psi_j + \phi$, $\phi_i \neq \phi_j$, and $\text{image}(\phi_i - \phi_j) \subseteq p^k \cdot \mathbb{Z}_{p^{k+1}} \cong \mathbb{Z}_p$). Then ψ_1, \dots, ψ_ℓ are distinct with pairwise agreement at most $\frac{1}{p}$, and each ψ_i has agreement at least $1/p + \alpha_i$ with g . It now follows from the Johnson bound that $\sum_{i=1}^\ell \alpha_i^2 \leq 1$.

Finally, the fourth property can be verified exactly as in Claim 3.4, and the inductive proof of the claim is also derived exactly as in the proof of Claim 3.4. As $\alpha^C \geq \sum_{i=1}^\ell \alpha_i^C$, another induction argument allows us to conclude that $1 \geq \sum_{i=1}^L \beta_i^C$. \square

The theorem statement follows by using $k = 0$.

3.4 General abelian groups

Finally we deal with the case of general H . We first deal with the case of slightly restricted G which we describe below.

Let $H = \mathbb{Z}_{p_1^{r_1}} \times \dots \times \mathbb{Z}_{p_m^{r_m}}$. We assume below that p_i divides the order of G for every i , and that $p = p_1 \leq p_2 \leq \dots \leq p_m$. Thus $\Lambda_{G,H} = 1/p$. The following lemma deals with this case. Later we remove the restriction on G .

LEMMA 3.7. *There exists a constant C_1 such that the following holds. Let $p = p_1 \leq \dots \leq p_m$ be primes and r_1, \dots, r_m be positive integers. Let $H = \prod_{i=1}^m \mathbb{Z}_{p_i^{r_i}}$ and let G be a finite abelian group such that p_i divides the order of G for every $i \in [m]$. Given any function $f : G \rightarrow H$ and agreement parameter $\epsilon > 0$, the number of affine homomorphisms from G to H that have agreement at least $1/p + \epsilon$ with f is at most $(1/\epsilon)^{C_1}$.*

PROOF. The proof is obtained by essentially making notational changes to the proof of Theorem 3.3. For $i \in [m]$ let $H_i = \mathbb{Z}_{p_i^{r_i}}$. Fix $f : G \rightarrow H$. We view f as an m -tuple of functions $f = \langle f_1, \dots, f_m \rangle$ with $f_i : G \rightarrow H_i$.

For $k \in [m]$, let $H_{[k]} = \prod_{i=1}^k H_i$, and let $f^{[k]} : G \rightarrow H_{[k]}$ be given by $f^{[k]}(x) = \langle f_1(x), \dots, f_k(x) \rangle$. We extend this notation to include $k = 0$, by allowing a unique function from $G \rightarrow \{1\}$ (which is considered to be a homomorphism).

For $k_1 \leq k_2$, we say that a function $f_2 : G \rightarrow H_{[k_2]}$ extends $f_1 : G \rightarrow H_{[k_1]}$ if $f_1 = f_2^{[k_1]}$.

In what follows, we fix k , $0 \leq k < m$ and consider the function $f^{[k]}$. We also fix a homomorphism ϕ from G to

$H_{[k]}$ that has agreement $\frac{1}{p} + \alpha$ with $f^{[k]}$, for some $\alpha \geq \epsilon$. We prove that the number of homomorphisms that extend ϕ and have agreement $1/p + \epsilon$ with f is small.

CLAIM 3.8. *For every $0 \leq k \leq m$, and for every homomorphism ϕ that has agreement $1/p + \alpha$ with $f^{[k]}$, the number of homomorphisms from G to H that have agreement $1/p + \epsilon$ with f and extend ϕ , is at most $(\alpha/\epsilon)^{C_1}$, where C_1 is the constant from Theorem 3.2, for $c = C_2$, and C_2 is the constant from Theorem 3.5.*

PROOF. We prove the claim by induction on $r - k$. When $r - k = 0$, the claim is trivially true. So assume the claim is true for $r - (k + 1)$ and we will now prove it for the case of $r - k$.

Let $S = \{x \in G \mid f^{[k]}(x) = \phi(x)\}$. Consider $f^{[k+1]}$ and let ϕ_1, \dots, ϕ_ℓ be all the homomorphisms from G to $H_{[k+1]}$ that extend ϕ with agreement at least $\frac{1}{p} + \epsilon$ with $f^{[k+1]}$. Let $S_i = \{x \mid f^{[k+1]}(x) = \phi_i(x)\}$ and let $\mu(S_i) = \frac{1}{p} + \alpha_i$.

We prove below that the sets $S_1, \dots, S_\ell \subseteq G$ form a $((1/(p_{k+1}), (1/p_{k+1}^2), C_2)$ -special intersecting family. This will allow us to apply Theorem 3.2 to these sets and this in turn will allow us to derive the inductive claim easily. We use q as shorthand for p_{k+1} .

Verifying Intersection Properties: The first property requires that $\mu(S_i) \geq \frac{1}{q}$ which is immediate from the definition of the S_i 's. The second property required that $\mu(S_i \cap S_j) \leq \frac{1}{q}$. This again is true, since $S_i \cap S_j \subseteq \{x \mid \phi_i(x) = \phi_j(x)\}$ and the latter has density at most $1/q$.

The third property follows from Theorem 3.5, applied to the function $f_{k+1}^{[k+1]} : G \rightarrow H_{k+1}$. Let $(\phi_i)_j$ denote the j th coordinate of the function ϕ_i . Then note that the functions $(\phi_i)_{k+1}$ are distinct homomorphisms from G to H_{k+1} , with $(\phi_i)_{k+1}$ having agreement at least $1/q + \alpha_i$ with f_{k+1} . It follows from Theorem 3.5, that $\sum_{i=1}^{\ell} \alpha_i^{C_2} \leq 1$.

Finally, we come to the fourth property. We need to show that if $\mu(S_I) > 1/q^2$, and $J \subseteq I$ contains at least two distinct elements, then $S_I = S_J$. For any set $K \subseteq [l]$ define $T_K = \{x \mid \phi_i(x) = \phi_j(x), \forall i, j \in K\}$. Note that $T_K = G$ if and only if $|K| \leq 1$. Also note that $\phi_i(x) = \phi_j(x)$ if and only if $(\phi_i)_{k+1}(x) = (\phi_j)_{k+1}(x)$ (since both extend ϕ). Thus T_K is the kernel of a homomorphism from G to H_{k+1} , we have that $1/\mu(T_K)$ must divide $|H_{k+1}|$ and so $\mu(T_K) = 1/q^i$ for non-negative integer i . But $S_K \subseteq T_K$ for every K , and since $\mu(S_I), \mu(S_J) > 1/q^2$, it must be that $\mu(T_I), \mu(T_J) = 1/q$. Since $T_I \subseteq T_J$, we get that $T_I = T_J$. Finally, fix some $j \in J$ and let $Y = \{x \mid \phi_j(x) = f^{[k+1]}(x)\}$. Since $\phi_i(x) = \phi_j(x)$ for every $x \in Y \cap T_I$ and $i \in I$, we have $S_I = Y \cap T_I = Y \cap T_J = S_J$ as required for the fourth property.

Proof of Claim: We are now ready to prove the claim. Since the sets S_1, \dots, S_ℓ form a $(1/q, 1/q^2, C_2)$ -special intersecting family, we can apply Theorem 3.2 to conclude that there exists an absolute constant C_1 such that $(\alpha')^{C_1} \geq \sum_{i=1}^{\ell} \alpha_i^{C_1}$ where $\alpha' = \mu(\cup_{i \in \ell} S_i) - 1/p$. Since for every i , $S_i \subseteq S$, we have $\cup_{i \in \ell} S_i \subseteq S$ and so $\alpha \geq \alpha'$. We may thus conclude that $\alpha^{C_1} \geq \sum_{i=1}^{\ell} \alpha_i^{C_1}$. This immediately yields the claim. \square

The lemma statement follows immediately, by using $k = 0$.

We now show how to use Lemma 3.7 to prove our main combinatorial list decoding bound, Theorem 2.2.

Proof of Theorem 2.2 Let $H = H' \times H''$, where $H' = \mathbb{Z}_{p_1}^{\epsilon_1} \times \dots \times \mathbb{Z}_{p_m}^{\epsilon_m}$ and $H'' = \mathbb{Z}_{q_1}^{\epsilon_1} \times \dots \times \mathbb{Z}_{q_k}^{\epsilon_k}$, such that $p := p_1 \leq \dots \leq p_m$, $p_i \mid |G|$ for each $i \in [m]$, while $q_j \nmid |G|$ for each $j \in [k]$. Thus $\Lambda = \frac{1}{p}$.

Let $f = (f', f'') : G \rightarrow H$ be any function, where $f' : G \rightarrow H'$ and $f'' : G \rightarrow H''$. Let

$$\mathcal{L} = \{\phi \in \text{aHom}(G, H) : \text{agree}(\phi, f) \geq \Lambda + \epsilon\}.$$

Consider a $\phi = (\phi', \phi'') \in \mathcal{L}$ where $\phi' \in \text{aHom}(G, H')$ and $\phi'' \in \text{aHom}(G, H'')$. Then $\text{agree}(\phi', f') \geq \Lambda + \epsilon$ and $\text{agree}(\phi'', f'') \geq \Lambda + \epsilon$. Thus

$$\phi' \in \mathcal{L}_1 := \{\psi' \in \text{aHom}(G, H') : \text{agree}(\psi', f') \geq \frac{1}{p} + \epsilon\},$$

and,

$$\phi'' \in \mathcal{L}_2 := \{\psi'' \in \text{aHom}(G, H'') : \text{agree}(\psi'', f'') \geq \frac{1}{p} + \epsilon\}.$$

By Lemma 3.7, the size of \mathcal{L}_1 is at most $\left(\frac{1}{\epsilon}\right)^{C_1}$. Since $\gcd(|G|, |H''|) = 1$, the only functions in $\text{aHom}(G, H'')$ are the constant functions (which pairwise disagree everywhere), and hence the size of \mathcal{L}_2 is at most $\frac{1}{\epsilon}$.

Thus $\mathcal{L} \subseteq \mathcal{L}_1 \times \mathcal{L}_2$, and so

$$|\mathcal{L}| \leq |\mathcal{L}_1| \cdot |\mathcal{L}_2| \leq \left(\frac{1}{\epsilon}\right)^{C_1} \cdot \frac{1}{\epsilon}.$$

Taking $C = C_1 + 1$, the theorem follows. \square

4. ANALYZING SPECIAL INTERSECTING FAMILIES

We prove Theorem 3.2 by a sequence of lemmas. Fix a (ρ, ρ^2, c) -special intersecting family $S_1, \dots, S_\ell \subseteq X$ and let $\alpha_i = \mu(S_i) - \rho$. Let $\alpha = \mu(\cup_{i \in [\ell]} S_i) - \rho$. Recall our goal is to prove that $\alpha^C \geq \sum_{i=1}^{\ell} \alpha_i^C$ for a sufficiently large C (that depends only on c).

The first three of the lemmas prove different lower bounds on α under some special conditions, We glue these together to prove a weak variant of Theorem 3.2 in Lemma 4.5. In Lemma 4.6 we show that the weak condition actually implies the strong condition and this leads to the proof of Theorem 3.2 at the end of this section.

4.1 A weak bound on the cardinality of the union

Let $I \subseteq [\ell]$. We say that I is a *collinear set* if either $|I| \leq 2$ or $\mu(S_I) > \rho^2$. Collinear sets enjoy the following properties²:

1. If I is a collinear set, and $J \subseteq I$, then J is a collinear set. Indeed, if $|J| \geq 2$, then by property (4) of special intersecting families, $S_J = S_I$ and J is collinear. If $|J| < 2$, then J is collinear by definition.
2. If I, J are collinear sets with $|I \cap J| \geq 2$, then $I \cup J$ is a collinear set. To see this, notice that $S_{I \cup J} = S_I \cap S_J$, and by property (4) of special intersecting families, $S_I = S_{I \cap J} = S_J$. Thus $S_{I \cup J} = S_I = S_J$, and so $I \cup J$ is a collinear set.

²We draw an analogy with a set of points, labelled by elements of $[\ell]$, on an abstract plane. Here a subset of $[\ell]$ is called "collinear" if the corresponding points all lie on a single straight line.

3. For $i \neq j \in [\ell]$, let

$$L_{i,j} = \{k \in [\ell] \mid \{i, j, k\} \text{ is a collinear set}\}.$$

Then $L_{i,j}$ is a collinear set. This follows easily from the previous property.

We say $J \subseteq [\ell]$ is *in general position* if there is no collinear set $K \subseteq J$ with $|K| = 3$.

The argument for proving the weak lower bound on α in Lemma 4.5 proceeds by separately analyzing collinear sets and sets in general position. For any collinear set I , by finding a sunflower in the set system $\{S_i \mid i \in I\}$, we obtain a lower bound on α of the form $\sum_{i \in I} \alpha_i$. For a set J in general position, by an inclusion exclusion argument, we obtain a lower bound on α of the form $\Omega(\sum_{j \in J} \alpha_j)$. The weak lower bound on α is now obtained by showing a dichotomy: either there is a collinear set I with $\sum_{i \in I} \alpha_i$ large, or there is a set J in general position with $\sum_{j \in J} \alpha_j$ large.

LEMMA 4.1. *For all $i \neq j \in [\ell]$, $\alpha \geq \alpha_i + \alpha_j$.*

PROOF. Simple inclusion-exclusion, using Property (2) of special intersecting families. We have

$$\begin{aligned} \rho + \alpha &\geq \mu(S_i \cup S_j) = \mu(S_i) + \mu(S_j) - \mu(S_i \cap S_j) \\ &\geq (\rho + \alpha_i) + (\rho + \alpha_j) - \rho = \rho + \alpha_i + \alpha_j. \end{aligned}$$

The lemma follows. \square

LEMMA 4.2. *If $I \subseteq [\ell]$ is a collinear set, then $\alpha \geq \sum_{i \in I} \alpha_i$.*

PROOF. For distinct $i_1, i_2 \in I$, we know that $S_{\{i_1, i_2\}} = S_I$. Thus the sets $\{S_i\}_{i \in I}$ form a sunflower, i.e., each set S_i is of the form $R_i \cup S_I$, where the R_i 's are pairwise disjoint. In turn this implies that

$$\begin{aligned} \rho + \alpha &= \mu(\cup_{i \in [\ell]} S_i) \\ &\geq \mu(\cup_{i \in I} S_i) \\ &= \mu(S_I) + \sum_{k \in J} \mu(S_k - S_I) \\ &= \mu(S_{i_0}) + \sum_{i \in I \setminus \{i_0\}} \mu(S_i - S_I) \quad (\text{for some } i_0 \in I) \\ &\geq (\alpha_{i_0} + \rho) + \sum_{i \in I \setminus \{i_0\}} \alpha_i \\ &= \rho + \sum_{i \in I} \alpha_i. \end{aligned}$$

We thus get $\alpha \geq \sum_{i \in I} \alpha_i$ as claimed. \square

LEMMA 4.3. *If $J \subseteq [\ell]$ is in general position, and $|J| \leq 1/\sqrt{\rho}$, then $\alpha \geq \frac{|J|-2}{3}\rho + \frac{1}{2} \sum_{j \in J} \alpha_j$.*

PROOF. Before proving the lemma, we first state and prove a simple variant of the standard inclusion-exclusion lower bound on the cardinality of $\cup_j S_j$.

CLAIM 4.4. *For any set $J \subseteq [\ell]$,*

$$\left| \bigcup_{j \in J} S_j \right| \geq \frac{1}{2} \sum_{j \in J} |S_j| - \frac{1}{2} \sum_{i,j,k \in J, i < j < k} |S_i \cap S_j \cap S_k|.$$

PROOF. Consider an element $e \in \cup_{j \in J} S_j$ and say it is contained in m of the sets. Then e additively contributes $m/2$ to the first term, $\frac{1}{2} \sum_{j \in J} |S_j|$. On the other hand in

the second term it subtracts at least $\frac{1}{2} \binom{m}{3}$. Thus the net contribution of e to the RHS is at most $\frac{m}{2} \left(1 - \frac{(m-1)(m-2)}{6}\right)$. It can be verified that this contribution is at most 1 for every positive integer m , which is also a lower bound on its contribution to the LHS. \square

We are now ready to prove the lemma. We have:

$$\begin{aligned} \rho + \alpha &= \mu \left(\bigcup_{j \in [\ell]} S_j \right) \\ &\geq \mu \left(\bigcup_{j \in J} S_j \right) \\ &\geq \frac{1}{2} \sum_{j \in J} \mu(S_j) - \frac{1}{2} \sum_{i,j,k \in J, i < j < k} \mu(S_{\{i,j,k\}}) \\ &\quad (\text{Using Claim 4.4}) \\ &\geq \frac{1}{2} \sum_{j \in J} (\rho + \alpha_j) - \frac{1}{2} \binom{|J|}{3} \rho^2 \end{aligned}$$

(Using definition of α_j and that J is in general position)

$$= \frac{|J|}{2} \rho - \frac{1}{2} \binom{|J|}{3} \rho^2 + \frac{1}{2} \sum_{j \in J} \alpha_j$$

Rearranging terms above, we get

$$\begin{aligned} \alpha &\geq \frac{|J|-2}{2} \rho - \frac{1}{2} \binom{|J|}{3} \rho^2 + \frac{1}{2} \sum_{j \in J} \alpha_j \\ &= \frac{|J|-2}{2} \cdot \rho \cdot \left(1 - \frac{|J|(|J|-1)}{3} \rho\right) + \frac{1}{2} \sum_{j \in J} \alpha_j \\ &\geq \frac{|J|-2}{3} \cdot \rho + \frac{1}{2} \sum_{j \in J} \alpha_j, \end{aligned}$$

as claimed.

We now put the above ingredients together to prove a weak version of Theorem 3.2. Specifically the following lemma claims that some positive power of α is lower bounded by some multiple of the sum of the α_i 's to the same power. Unfortunately, the multiple is a pretty weak one and becomes smaller as the number of summands grow (whereas in Theorem 3.2 we want this multiple to be 1!). Nevertheless this weak version is good enough, as we'll see later.

LEMMA 4.5. *For every $c < \infty$ there exist constants c_1, c_2, c_3 (depending only on c) with $c_1, c_3 < \infty$ and $c_2 < 1$ such that for every set $K \subseteq [\ell]$,*

$$\alpha^{c_1} \geq \frac{1}{c_3 \cdot |K|^{c_2}} \cdot \sum_{k \in K} \alpha_k^{c_1}.$$

Remark: Note that if the constant c_2 in the exponent of $|K|$ had been replaced by 1, the lower bound on α would have been quite trivial (with $c_1, c_2 = 1$). The slightly smaller exponent turns out to be sufficiently powerful to let us derive Theorem 3.2 later.

PROOF. Fix a set $K \subseteq [\ell]$.

We prove the lemma for $c_1 = c + 1$. (The other constants will be determined later.) For a subset $I \subseteq K$, let $\text{wt}(I) =$

$\sum_{i \in I} \alpha_i^{c_1}$. Let $P = \{k \in K \mid \alpha_k < \rho^{c_1}\}$ and $Q = \{k \in K \mid \alpha_k \geq \rho^{c_1}\}$. We consider two cases based on whether $\text{wt}(P)$ is larger than $\text{wt}(Q)$ or not.

Case 1: $\text{wt}(P) \geq \text{wt}(Q)$: If $|P| = 1$, then it must be that $|Q| = 0$ and then $|K| = 1$ and if so this case is trivial. So $|P|$ must be at least 2 in this case. Again we divide the analysis into two subcases.

The first subcase we consider is when there exist distinct $i, j, k \in P$ with $\{i, j, k\}$ not collinear. In this case, by applying Lemma 4.3 to the set $J = \{i, j, k\}$, we get $\alpha \geq \rho/3$. On the other hand, we have

$$\begin{aligned} \sum_{k \in P} \alpha_k^{c_1+1} &\leq \max_{k \in P} \alpha_k \cdot \sum_{j \in P} \alpha_j^c \\ &\leq \max_{k \in P} \alpha_k \cdot 1 \quad (\text{Using property (3) of the family}) \\ &\leq \rho^{c_1} \quad (\text{By definition of } P). \end{aligned}$$

So in this subcase we have

$$\begin{aligned} \alpha^{c_1} &\geq \rho^{c_1}/(3^{c_1}) \\ &\geq \frac{1}{3^{c_1}} \sum_{k \in P} \alpha_k^{c_1} \\ &\geq \frac{1}{2 \cdot 3^{c_1}} \sum_{k \in K} \alpha_k^{c_1} \quad (\text{Since } \text{wt}(P) \geq \frac{1}{2} \text{wt}(K)) \\ &\geq \frac{1}{c_3 |K|^{c_2}} \sum_{k \in K} \alpha_k^{c_1} \quad (\text{provided } c_3 \geq 2 \cdot 3^{c_1} \text{ and } c_2 \geq 0). \end{aligned}$$

Now we move to the second subcase. In this case, we assume that for every triple i, j, k in P , $\{i, j, k\}$ is collinear. This implies that P itself is collinear. By Lemma 4.2, we have $\alpha \geq \sum_{k \in P} \alpha_k$. Thus $\alpha^{c_1} \geq \sum_{k \in P} \alpha_k^{c_1} \geq \frac{1}{2} \sum_{k \in K} \alpha_k^{c_1}$. Thus again we have $\alpha^{c_1} \geq \frac{1}{c_3 |K|^{c_2}} \sum_{k \in K} \alpha_k^{c_1}$ provided $c_3 \geq 2$ and $c_2 \geq 0$. This concludes the analysis of Case 1.

Case 2: $\text{wt}(P) < \text{wt}(Q)$: Note that this is the case where our lower bound on α will start deteriorating with $|K|$.

By Property (3) of the (ρ, ρ^2, c) -special intersection condition, in this case $|Q|$ is at most $\frac{1}{\rho^{c-c_1}}$. We divide our analysis into two subcases again.

Let $c_2 = 1 - 1/(2c \cdot c_1)$.

The first subcase is when there exists a collinear $I \subseteq K$ such that $\text{wt}(I) \geq \frac{\text{wt}(Q)}{|Q|^{c_2}}$. If so, by Lemma 4.2, we have that $\alpha \geq \sum_{i \in I} \alpha_i$ and so

$$\alpha^{c_1} \geq \sum_{i \in I} \alpha_i^{c_1} \geq \frac{\text{wt}(Q)}{|Q|^{c_2}} \geq \frac{\text{wt}(K)}{2|Q|^{c_2}} \geq \frac{\text{wt}(K)}{2|K|^{c_2}}.$$

This yields the lemma for every $c_3 \geq 2$.

Now we consider the final remaining subcase, where for every collinear $I \subseteq K$, it is the case that $\text{wt}(I) < \frac{\text{wt}(Q)}{|Q|^{c_2}}$. We first prove that there is a $J \subset Q$ in general position such that $|J| = |Q|^{1/(2c \cdot c_1)}$ and $\text{wt}(J) \geq \text{wt}(Q)/|Q|^{c_2}$. Applying Lemma 4.3 to this set then completes the analysis for this case.

We start with showing that such a set J exists. We pick J greedily starting with $J = \emptyset$. We maintain a set of candidates Q' , initially containing all of Q . We then add the element $i \in Q'$ with the largest value of α_i to J . We remove i from Q' as also every element $k \in Q'$ such that for some $j \in J \setminus \{i\}$, $\{i, j, k\}$ is collinear. We repeat until $|J| = |Q|^{1/(2c \cdot c_1)}$.

To apply this construction, we first claim that Q' does not become empty (which would preclude us from grow-

ing J). Note that for every $i, j \in J$, the total weight of the elements removed from Q' on account of i, j is at most $\text{wt}(L_{i,j}) < \text{wt}(Q)/|Q|^{c_2}$. Since there are at most $|J|^2/2 \leq |Q|^{1/(c \cdot c_1)}/2$ such pairs, the total weight of the removed elements is at most $\text{wt}(Q)/(2 \cdot |Q|^{c_2 - 1/(c \cdot c_1)}) < \text{wt}(Q)/2$, provided $c_2 > 1/(c \cdot c_1)$, which holds for $c_2 \geq 1/2$ and $c \geq 1$ and $c_1 \geq 2$.

Thus, the process concludes with $|J| = |Q|^{1/(2c \cdot c_1)}$ and J is in general position by construction. To see that the weight of J is large, notice that every element of J has weight at least that of the average element, i.e., $\alpha_k^{c_1} \geq \text{wt}(Q')/|Q'| \geq \text{wt}(Q)/(2|Q|)$. We thus conclude that $\text{wt}(J) \geq |J| \cdot \text{wt}(Q)/(2|Q|) \geq \text{wt}(Q)/(2|Q|^{1 - 1/(2c \cdot c_1)}) = \text{wt}(Q)/(2|Q|^{c_2})$.

Finally, note that $|J| = |Q|^{1/(2c \cdot c_1)} \leq (1/(\rho^{c \cdot c_1}))^{1/(2c \cdot c_1)} = 1/\sqrt{\rho}$. So this puts us in a position to apply Lemma 4.3 to the set J to conclude that $\alpha \geq \frac{1}{2} \sum_{j \in J} \alpha_j$. This in turn yields $\alpha^{c_1} \geq \frac{1}{2^{c_1}} (\sum_{j \in J} \alpha_j)^{c_1} \geq \frac{1}{2^{c_1}} \sum_{j \in J} \alpha_j^{c_1} \geq \frac{1}{2^{c_1}} \text{wt}(Q)/|Q|^{c_2} \geq \frac{1}{2^{c_1+1}} \text{wt}(K)/|K|^{c_2}$. This yields the lemma statement for $c_2 = 1 - 1/(2c \cdot c_1)$ and $c_3 \geq 2^{1+c_1}$.

Putting all the cases together, we get the lemma statement for $c_2 = \max\{1 - 1/(2c \cdot c_1), \frac{1}{2}\}$ and $c_3 = 2 \cdot 3^{c_1}$ (and, of course, $c_1 = \max\{2, c + 1\}$). \square

4.2 Strengthening the bound

Our final lemma shows how to convert the assertion of Lemma 4.5, in conjunction with Lemma 4.1, to get the assertion of Theorem 3.2. We omit the proof.

LEMMA 4.6. *For every $c_1 < \infty$, $c_2 < 1$, and $c_3 < \infty$ there exists $C < \infty$ such that the following holds: Let $\alpha, \alpha_1, \dots, \alpha_\ell$ be non-negative reals such that (1) $\alpha \geq \alpha_i + \alpha_j$ for every $i \neq j \in [\ell]$ and (2) For every $K \subseteq [\ell]$, $\alpha^{c_1} \geq \frac{1}{c_3 |K|^{c_2}} \cdot \sum_{k \in K} \alpha_k^{c_1}$. Then $\alpha^C \geq \sum_{k=1}^\ell \alpha_k^C$.*

Theorem 3.2 now follows by combining Lemma 4.5 and Lemma 4.6.

5. THE LIST DECODING ALGORITHM

In this section we sketch the list decoding algorithm. As we shall see, the combinatorial bounds that we proved make the algorithm relatively straightforward. Our algorithm can be viewed as a direct extension of the original list-decoder of Goldreich and Levin [4].

We are given a function $f : G \rightarrow H$. We want to locally list decode all affine homomorphisms $\phi \in \text{aHom}(G, H)$ such that $\text{agree}(f, \phi) \geq \Lambda_{G,H} + \epsilon$. (In what follows we refer to agreements of the form $\Lambda_{G,H} + \Omega(\epsilon)$ as “noticeable” agreements.) Recall further that the groups G and H are presented “explicitly”. Specifically, G is given by a sequence of pairs $\langle (p_1, r_1), \dots, (p_k, r_k) \rangle$, where p_1, \dots, p_k are prime numbers and r_1, \dots, r_k are positive integers and $G = \bigoplus_{i=1}^k \mathbb{Z}_{p_i}^{r_i}$. Similarly H is given by a sequence of pairs $\langle (q_1, s_1), \dots, (q_\ell, s_\ell) \rangle$ and $H = \bigoplus_{j=1}^\ell \mathbb{Z}_{q_j}^{s_j}$.

For $j \in [\ell]$, let $H_j = \mathbb{Z}_{q_j}^{s_j}$ and let $f = \langle f_1, \dots, f_\ell \rangle$ where $f_j : G \rightarrow H_j$. The rough idea of our algorithm is to recover, separately for each $j \in [\ell]$, the list \mathcal{L}_j of affine homomorphisms $\phi \in \text{aHom}(G, H_j)$ that have noticeable agreement with f_j . (This step itself involves an iterative algorithm using the decomposition of G and is described in the next paragraph.) Given the lists \mathcal{L}_j , our algorithm finds,

by induction on j , the list \mathcal{P}_j of all affine homomorphisms from G to $H_1 \times \dots \times H_j$ that have large agreement with the $f^{[j]} \stackrel{\text{def}}{=} (f_1, \dots, f_j)$, the projection of f to its first j coordinates. This step is simple and is a variation on [4]. We note that $\mathcal{P}_j \subseteq \mathcal{P}_{j-1} \times \mathcal{L}_j$. We enumerate all members $\phi \in \mathcal{P}_{j-1} \times \mathcal{L}_j$ and include in \mathcal{P}_j only those ϕ 's that have the desired amount of agreement with $f^{[j]}$. By Theorem 2.2, both \mathcal{P}_{j-1} and \mathcal{L}_j have size at most $\text{poly}(1/\epsilon)$ and so each iteration takes only $\text{poly}(1/\epsilon, \log |G|, \log |H|)$ time and $\text{poly}(1/\epsilon)$ queries.

We now turn to the task of computing the set Ψ of all affine homomorphisms from $G = G_1 \times \dots \times G_k$ (where $G_i = \mathbb{Z}_{p_i}^{r_i}$) to a group H of the form $H = \mathbb{Z}_{q^r}$ that have noticeable agreement with a given function $f : G \rightarrow H$. (Note that each group H_j of the previous paragraph is of this form and so this algorithm will suffice to compute \mathcal{L}_j for every $j \in [\ell]$.) For this step we extend the algorithm of [4] in a natural way. Let $G_{(<i)} \stackrel{\text{def}}{=} \bigoplus_{t < i} G_t$. We similarly define $G_{(>i)}$ and $G_{(\leq i)}$. For $i \in [k]$ let $\Phi_i = \{\phi \in \text{Hom}(G_{(\leq i)}, H) \mid \exists \psi \in \text{aHom}(G_{(>i)}, H)$ such that $\tau(x, y) \stackrel{\text{def}}{=} (\phi(x) + \psi(y)) \in \Psi\}$.

The list decoding algorithm works in k stages. In stage i , $i \in [k]$, the algorithm computes a set $S_i \subseteq \text{Hom}(G_{(\leq i)}, H)$ from the set S_{i-1} it computed in the previous stage. The set S_i is computed in two substages, called *extending* and *pruning*. The extending phase finds a large collection of homomorphisms that includes all the homomorphisms in Φ_i . The pruning phase weeds out many of those homomorphisms that are not in Φ_i , and collects the remaining few in S_i . After all k stages have completed, S_k is then used to compute the final list.

Extending: Note that $G = G_{(<i)} \times G_i \times G_{(>i)}$. If $p_i \neq q$, then we just set $S_i = S_{i-1}$ and exit. Otherwise, we proceed as follows.

For each $\phi \in S_{i-1}$, do the following $\text{poly}(\log |G|, \log |H|, \frac{1}{\epsilon})$ times:

1. Pick $s \in G_{(>i)}$ uniformly at random.
2. Pick $x_1, x_2 \in G_{(<i)}$ and $c_1, c_2 \in G_i$ uniformly at random.
3. If $c_1 - c_2$ is not divisible by q , solve the system of equations for $a, b \in \mathbb{Z}_{q^r}$:

$$\phi(x_1) + ac_1 + b = f(x_1, c_1, s)$$

$$\phi(x_2) + ac_2 + b = f(x_2, c_2, s)$$

4. Let $\theta \in \text{Hom}(G_{(<i)} \times G_i, H)$ be given by $\theta(x, c) = \phi(x) + a \cdot c \pmod{q^r}$. Include θ in S_i .

It can be shown that if $\Phi_{i-1} \subseteq S_{i-1}$, then with high probability, $\Phi_i \subseteq S_i$.

Pruning: Now we write $G = G_{(\leq i)} \times G_{(>i)}$. In the pruning phase, we will cut down the size of S_i while preserving the fact that $\Phi_i \subseteq S_i$. Initialize S'_i to \emptyset .

Repeat $\text{poly}(\log |G|, \log |H|, \frac{1}{\epsilon})$ times:

1. Pick a random $s \in G_{(>i)}$.
2. For each $\phi \in S_i$, by random sampling, find all $b \in H$ such that $\Pr_{x \in G_{(\leq i)}}[\phi(x, s) - f(x, s) = b] \geq \Lambda_{G, H} + \epsilon/2$.
3. If there exists even one b as above, then add ϕ to S'_i .

Upon conclusion, replace S_i with the “pruned” set S'_i .

Our list decoding bound implies that after the pruning phase, the cardinality of S_i is at most $\text{poly}(\log |G|, \log |H|, \frac{1}{\epsilon})$ for some fixed polynomial. This ensures that the S_i don't blow up in size as i increases. Furthermore, it can be shown that with high probability, all the $\phi \in \Phi_i$ remain in S_i after the pruning phase.

The above iterations give us a set S_k that contains Φ_k with high probability. Now by random sampling, for each element $\phi \in S_k$, we find all $b \in H$ such that $\Pr_{x \in G}[\phi(x) - f(x) = b]$ is noticeable, and include $\phi(\cdot) + b$ in our final list. This concludes the description of the list decoding algorithm. Its correctness is argued by verifying the claims about the sets S_i made above. Details can be obtained in the full version of this paper [2].

Acknowledgments

Thanks to Amir Shpilka for many valuable discussions.

6. REFERENCES

- [1] A. Akavia, S. Goldwasser, and S. Safra. Proving hard-core predicates using list decoding. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.
- [2] Irit Dinur, Elena Grigorescu, Swastik Kopparty, and Madhu Sudan. Decodability of Group Homomorphisms beyond the Johnson Bound. *Electronic Colloquium on Computational Complexity (ECCC)*, 15(TR08-020), 2008.
- [3] A. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss. Near-optimal sparse fourier representations via sampling. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2002.
- [4] Oded Goldreich and Leonid Levin. A hard-core predicate for all one-way functions. *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 25–32, May 1989.
- [5] Oded Goldreich, Ronitt Rubinfeld, and Madhu Sudan. Learning polynomials with queries: The highly noisy case. *SIAM Journal on Discrete Mathematics*, 13(4):535–570, November 2000.
- [6] Elena Grigorescu, Swastik Kopparty, and Madhu Sudan. Local decoding and testing for homomorphisms. In *APPROX-RANDOM*, volume 4110 of *Lecture Notes in Computer Science*, pages 375–385. Springer, 2006.
- [7] V. Guruswami and M. Sudan. Extensions to the johnson bound, 2001.
- [8] Venkatesan Guruswami and Atri Rudra. Explicit capacity-achieving list-decodable codes. *Electronic Colloquium on Computational Complexity (ECCC)*, (133), 2005.
- [9] E. Kushilevitz and Y. Mansour. Learning decision trees using the fourier spectrum. *SICOMP: SIAM Journal on Computing*, 22, 1993.
- [10] Serge Lang. *Algebra*. Addison-Wesley, 1965.
- [11] Yishay Mansour. Randomized interpolation and approximation of sparse polynomials. *SIAM J. Comput.*, 24(2):357–368, 1995.
- [12] Farzad Parvaresh and Alexander Vardy. Correcting errors beyond the guruswami-sudan radius in

polynomial time. In *FOCS*, pages 285–294. IEEE Computer Society, 2005.

- [13] Madhu Sudan, Luca Trevisan, and Salil Vadhan. Pseudorandom generators without the XOR lemma. *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pages 537–546, 1999.
- [14] A. Ta-Shma and D. Zuckerman. Extractor codes. *IEEE Transactions on Information Theory*, 50, 2004.