

On the List-Decodability of Random Linear Codes

Venkatesan Guruswami Johan Håstad Swastik Kopparty

Abstract—The list-decodability of random linear codes is shown to be as good as that of general random codes. Specifically, for every fixed finite field \mathbb{F}_q , $p \in (0, 1 - 1/q)$ and $\varepsilon > 0$, it is proved that with high probability a random linear code C in \mathbb{F}_q^n of rate $(1 - H_q(p) - \varepsilon)$ can be list decoded from a fraction p of errors with lists of size at most $O(1/\varepsilon)$.

This also answers a basic open question concerning the existence of highly list-decodable linear codes, showing that a list-size of $O(1/\varepsilon)$ suffices to have rate within ε of the information-theoretically optimal rate of $1 - H_q(p)$. The best previously known list-size bound was $q^{O(1/\varepsilon)}$ (except in the $q = 2$ case where a list-size bound of $O(1/\varepsilon)$ was known).

The main technical ingredient in the proof is a strong upper bound on the probability that ℓ random vectors chosen from a Hamming ball centered at the origin have too many (more than $\Omega(\ell)$) vectors from their linear span also belong to the ball.

Keywords. Linear codes, List decoding, Random coding, Probabilistic method, Hamming bound.

I. INTRODUCTION

One of the central problems in coding theory is to understand the trade-off between the redundancy built into codewords (a.k.a. the rate of the code) and the fraction of errors the code enables correcting. Suppose we are interested in codes over the binary alphabet (for concreteness) that enable recovery of the correct codeword $c \in \{0, 1\}^n$ from any noisy received word r that differs from c in at most pn locations. For each c , there are about $\binom{n}{pn} \approx 2^{H(p)n}$ such possible received words r , where $H(x) = -x \log_2 x - (1-x) \log_2 (1-x)$ stands for the binary entropy function. Now for each such r , the error-recovery procedure must identify c as a possible choice for the true codeword. In fact, even if the errors are randomly distributed and not worst-case, the algorithm must identify c as a candidate codeword for most of these $2^{H(p)n}$ received words, if we seek a low decoding error probability. This implies that there can be at most $\approx 2^{(1-H(p))n}$ codewords, or equivalently the largest rate R of the code one can hope for is $1 - H(p)$.

If we could pack about $2^{(1-H(p))n}$ pairwise disjoint Hamming balls of radius pn in $\{0, 1\}^n$, then one can achieve a rate approaching $1 - H(p)$ while guaranteeing correct and unambiguous recovery of the codeword from an arbitrary fraction p of errors. Unfortunately, it is well known that such

an asymptotic “perfect packing” of Hamming balls in $\{0, 1\}^n$ does not exist, and the largest size of such a packing is at most $2^{(\alpha(p)+o(1))n}$ for $\alpha(p) < 1 - H(p)$ (in fact $\alpha(p) = 0$ for $p \geq 1/4$). Nevertheless it turns out that an “almost-perfect packing” does exist: for every $\varepsilon > 0$ it is possible to pack $2^{(1-H(p)-\varepsilon)n}$ such Hamming balls such that no more than $O(1/\varepsilon)$ of them intersect at any point. In fact, a random packing has such a property with high probability.

A. List decoding

This fact implies that it is possible to achieve rate approaching the optimal $1 - H(p)$ bound for correcting a fraction p of worst-case errors in a model called *list decoding*. List decoding, which was introduced independently by Elias and Wozencraft in the 1950s [1], [2], is an error-recovery model where the decoder is allowed to output a small list of candidate codewords that must include all codewords within Hamming distance pn of the received word. Note that if at most pn errors occur, the list decoder’s output will include the correct codeword. In addition to the rate R of the code and the error fraction p , list decoding has an important third parameter, the “list-size,” which is the largest number L of codewords the decoder is allowed to output on any received word. The list-size thus bounds the maximum ambiguity in the output of the decoder.

For codes over an alphabet of size q , all the above statements hold with $H(p)$ replaced by $H_q(p)$, where $H_q(x) = x \log_q (q-x) - (1-x) \log_q (1-x)$ is the q -ary entropy function.

Definition 1.1 (Combinatorial list decodability): Let Σ be a finite alphabet of size q , $L \geq 1$ an integer, and $p \in (0, 1 - 1/q)$. A code $C \subseteq \Sigma^n$ is said to be (p, L) -list-decodable, if for every $x \in \Sigma^n$, there are at most L codewords of C that are at Hamming distance pn or less from x . Formally, $|B_n^q(x, p) \cap C| \leq L$ for every x , where $B_n^q(x, p) \subseteq \Sigma^n$ is the ball of radius pn centered at $x \in \{0, 1\}^n$.

We restrict $p < 1 - 1/q$ in the above definition, since for a random string differs from each codeword in at most a fraction $1 - 1/q$ of positions, and so over alphabet size q decoding from a fraction $1 - 1/q$ or more errors is impossible (except for trivial codes).

B. Combinatorics of list decoding

A fundamental question in list decoding is to understand the trade-off between rate, error-fraction, and list-size. For example, what list-size suffices if we want codes of rate within ε of the optimal $1 - H_q(p)$ bound? That is, if we define $L_{q,p}(\varepsilon)$ to be the minimum integer L for which there are q -ary (p, L) -list-decodable codes of rate at least $1 - H_q(p) - \varepsilon$ for infinitely many lengths n , how does $L_{q,p}(\varepsilon)$ behave for small ε (as we keep the alphabet size q and $p \in (0, 1 - 1/q)$ fixed)?

A preliminary conference version of this paper was presented at the 42nd Annual ACM Symposium on Theory of Computing, June 2010.

V. Guruswami is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. Email: guruswami@cmu.edu. Research supported by a Packard Fellowship, NSF grants CCF 0953155 and CCF-0963975, and US-Israel BSF grant 2008293.

J. Håstad is with the School of Computer Science and Communication, Royal Institute of Technology (KTH), Stockholm. Email: johanh@kth.se. Research supported by ERC grant 226203.

S. Kopparty is with the School of Mathematics, Institute for Advanced Study, Princeton, NJ. Email: swastik@mit.edu. Work done when the author was at MIT and partly during an internship at Microsoft Research, New England.

In the language of list-decoding, the above-mentioned result on “almost-perfect packings” states that for large enough block lengths, a random code of rate $1 - H_q(p) - \varepsilon$ is $(p, \frac{1}{\varepsilon})$ -list-decodable with high probability. In other words, $L_{q,p}(\varepsilon) \leq 1/\varepsilon$. This result appears in [3] (and is based on a previous random coding argument for linear codes from [4]). The result is explicitly stated in [3] only for $q = 2$, but trivially extends for arbitrary alphabet size q . This result is also tight for the random coding argument, in the sense that with high probability a *random* code of rate $1 - H_q(p) - \varepsilon$ is *not* $(p, c_{p,q}/\varepsilon)$ -list-decodable w.h.p. for some constant $c_{p,q} > 0$ [5].

On the negative side, it is known that the list-size is necessarily unbounded as one approaches the optimal rate of $1 - H_q(p)$. In other words, $L_{q,p}(\varepsilon) \rightarrow \infty$ as $\varepsilon \rightarrow 0$. This was shown for the binary case in [6], and his result implicitly implies $L_{2,p}(\varepsilon) \geq \Omega(\log(1/\varepsilon))$ (see [5] for an explicit derivation of this). For the q -ary case, it was shown that $L_{q,p}(\varepsilon) \rightarrow \infty$ as $\varepsilon \rightarrow 0$ in [7], [8].

An interesting question is to close the exponential gap in the lower and upper bounds on $L_{2,p}(\varepsilon)$, and more generally pin down the asymptotic behavior of $L_{q,p}(\varepsilon)$ for every q . At least one of the authors believes that the upper bound of $O(1/\varepsilon)$ is closer to the truth.

C. Context of this work

In this work, we address another fundamental combinatorial question concerning list-decodable codes, namely the behavior of $L_{q,p}(\varepsilon)$ when restricted to *linear codes*. A linear code over an alphabet of size q is simply a subspace of \mathbb{F}_q^n (\mathbb{F}_q being the field of size q ; we henceforth assume that q is a prime-power).

Most of the well-studied and practically used codes are linear codes. Linear codes admit a succinct representation in terms of its basis (called generator matrix). This aids in finding and representing such codes efficiently, and as a result linear codes are often useful as “inner” codes in concatenated code constructions.

In a linear code, by translation invariance, the neighborhood of every codeword looks the same, and this is often a very useful symmetry property. For instance, this property was recently used in [9] to give a black-box conversion of linear list-decodable codes to codes achieving capacity against a worst-case additive channel (the linearity of the list-decodable code is crucial for this connection). Lastly, list-decodability of linear codes brings to the fore some intriguing questions on the interplay between the geometry of linear spaces and Hamming balls, and is therefore interesting in its own right. For these and several other reasons, it is desirable to achieve good trade-offs for list decoding via linear codes.

Since linear codes are a highly structured subclass of all codes, proving the existence of linear codes which have list-decodability properties similar to general codes can be viewed as a strong “derandomization” of the random coding argument used to construct good list-decodable codes. A derandomized family of codes called “pseudolinear codes” were put forth in [10] since linear codes were not known to have strong enough list-decoding properties. Indeed, prior to this work, the results known for linear codes were substantially weaker

than for general codes (we discuss the details next). *Closing this gap is the main motivation behind this work.*

D. Status of list-decodability of linear codes

Zyablov and Pinsker proved that a random binary linear code of rate $1 - H(p) - \varepsilon$ is $(p, 2^{O(1/\varepsilon)})$ -list-decodable with high probability [4]. The proof extends in a straightforward way to linear codes over \mathbb{F}_q , giving list-size $q^{O(1/\varepsilon)}$ for rate $1 - H_q(p) - \varepsilon$. Let us define $L_{q,p}^{\text{lin}}(\varepsilon)$ to be the minimum integer L for which there is an infinite family of (p, L) -list-decodable linear codes over \mathbb{F}_q of rate at least $1 - H_q(p) - \varepsilon$. The results of [4] thus imply that $L_{q,p}^{\text{lin}}(\varepsilon) \leq \exp(O_q(1/\varepsilon))$.

Note that this bound is *exponentially worse* than the $O(1/\varepsilon)$ bound known for general codes. In [3], Elias mentions the following as the most obvious problem left open by the random coding results: *Is the requirement of the much larger list-size for linear codes inherent, or can one achieve list-size closer to the $O(1/\varepsilon)$ bound for general random codes?*

For the *binary* case, the *existence* of (p, L) -list-decodable linear codes of rate at least $1 - H(p) - 1/L$ is proven in [11]. This implies that $L_{2,p}^{\text{lin}} \leq 1/\varepsilon$. There are some results which obtain lower bounds on the rate for the case of small fixed list-size (at most 3) [6], [12], [13]; these bounds are complicated and not easily stated, and as noted in [14], are weaker for the linear case for list-size as small as 5.

The proof in [11] is based on a carefully designed potential function that quantifies list-decodability, and uses the “semi-random” method to successively pick good basis vectors for the code. The proof only guarantees that such binary linear codes exist with positive probability, and does not yield a high probability guarantee for the claimed list-decodability property. Further, the proof relies crucially on the binary alphabet and extending it to work for larger alphabets (or even the ternary case) has resisted all attempts. Thus, for $q > 2$, $L_{q,p}(\varepsilon) \leq \exp(O_q(1/\varepsilon))$ remained the best known upper bound on list-size. A high probability result for the binary case, and an upper bound of $L_{q,p}(\varepsilon) \leq O(1/\varepsilon)$ for \mathbb{F}_q -linear codes, were conjectured in [15, Chap. 5].

II. OUR RESULTS AND METHODS

In this paper, we resolve the above open question concerning list-decodability of linear codes over *all* alphabets. In particular, we prove that $L_{q,p}^{\text{lin}}(\varepsilon) \leq C_{q,p}/\varepsilon$ for a constant $C_{q,p} < \infty$. Up to constant factors, this matches the best known result for general, non-linear codes. Further, our result in fact shows that a random \mathbb{F}_q -linear code of rate $1 - H_q(p) - \varepsilon$ is $(p, C_{p,q}/\varepsilon)$ -list-decodable *with high probability*. This was not known even for the case $q = 2$. The high probability claim implies an efficient randomized Monte Carlo construction of such list-decodable codes.

We now briefly explain the difficulty in obtaining good bounds for list-decoding linear codes and how we circumvent it. This is just a high level description; see the next section for a more technical description of our proof method.

Let us recall the straightforward random coding method that shows the list-decodability of random (binary) codes. We pick a code $C \subseteq \{0, 1\}^n$ by uniformly and independently picking

$M = 2^{Rn}$ codewords. To prove it is (p, L) -list-decodable, we fix a center y and a subset S of $(L + 1)$ codewords of C . Since these codewords are independent, the probability that all of them land in the ball of radius pn around y is at most $(\frac{2^{H(p)n}}{2^n})^{L+1}$. A union bound over all 2^n choices of y and at most M^{L+1} choices of S shows that if $R \leq 1 - H(p) - 1/L$, the code fails to be (p, L) -list-decodable with probability at most $2^{-\Omega(n)}$.

Attempting a similar argument in the case of random linear codes, defined by a random linear map $A : \mathbb{F}_2^{Rn} \rightarrow \mathbb{F}_2^n$, faces several immediate obstacles. The 2^{Rn} codewords of a random linear code are not independent of one another; in fact the points of such a code are highly correlated and not even 3-wise independent (as $A(x + y) = Ax + Ay$). However, any $(L + 1)$ distinct codewords $Ax_1, Ax_2, \dots, Ax_{L+1}$ must contain a subset of $\ell \geq \log_2(L + 1)$ independent codewords, corresponding to a subset $\{x_{i_1}, \dots, x_{i_\ell}\}$ of *linearly independent* message vectors. This lets one mimic the argument for the random code case with $\log_2(L + 1)$ playing the role of $L + 1$. This version of the argument for linear codes appears in the work of Zyablov and Pinsker [4]. However, this leads to exponentially worse list-size bounds.

To get a better result, we somehow need to control the “damage” caused by subsets of codewords of low rank. This is the crux of our new proof. Stated loosely and somewhat imprecisely, we prove a strong upper bound on the fraction of such low rank subsets, by proving that if we pick ℓ random vectors from the Hamming ball $B_n(0, p)$ (for some constant ℓ related to our target list-size L), it is rather unlikely that more than $\Omega(\ell)$ of the 2^ℓ vectors in their span will also belong to the ball $B_n(0, p)$. (See Theorem 2.2 for the precise statement.) This “limited correlation” between linear subspaces and Hamming balls is the main technical ingredient in our proof. It seems like a basic and powerful probabilistic fact that might find other applications. The argument also extends to linear codes over \mathbb{F}_q after some adaptations.

A. Formal result statements

We first state our results in the case of binary codes and then proceed to the general case.

1) *Binary codes*: Our main result is that with high probability, a random linear code in \mathbb{F}_2^n of rate $1 - H(p) - \varepsilon$ can be list-decoded from p -fraction errors with list-size only $O(\frac{1}{\varepsilon})$. We also show the analogous result for random q -ary linear codes.

Theorem 2.1: Let $p \in (0, 1/2)$. There exist constants C_p and $\delta > 0$, such that for all $\varepsilon > 0$ and all large enough integers n , letting $R = 1 - H(p) - \varepsilon$, if $C \subseteq \mathbb{F}_2^n$ is a random linear code of rate R , then

$$\Pr[C \text{ is } (p, \frac{C_p}{\varepsilon})\text{-list-decodable}] > 1 - 2^{-\delta n}.$$

The proof begins by simplifying the problem to its combinatorial core. Specifically, we reduce the problem of studying the *list-decodability* of a random linear code of *linear* dimension to the problem of studying the *weight-distribution* of certain random linear codes of *constant* dimension. The next theorem analyzes the weight distribution of these constant dimensional

random linear codes. The notation $B_n(x, p)$ refers to the Hamming ball of radius pn centered at $x \in \mathbb{F}_2^n$.

Theorem 2.2 (Span of random subset of $B_n(0, p)$): For every $p \in (0, 1/2)$, there is a constant $C > 1$, such that for all n and all $\ell = o(\sqrt{n})$, if X_1, \dots, X_ℓ are picked independently and uniformly at random from $B_n(0, p)$, then

$$\Pr\left[|\text{span}(\{X_1, \dots, X_\ell\}) \cap B_n(0, p)| \geq C \cdot \ell\right] \leq 2^{-(6-o(1))n}.$$

We now give a brief sketch of the proof of Theorem 2.2. Index the elements of $\text{span}(\{X_1, \dots, X_\ell\})$ as follows: for $v \in \mathbb{F}_2^\ell$, let X_v denote the random vector $\sum_{i=1}^\ell v_i X_i$. Fix an arbitrary $S \subseteq \mathbb{F}_2^\ell$ of cardinality $C \cdot \ell$, and let us study the event E_S : that all the vectors $(X_v)_{v \in S}$ lie in $B_n(0, p)$. If none of the events E_S occur, we know that $|\text{span}(\{X_1, \dots, X_\ell\}) \cap B_n(0, p)| < C \cdot \ell$.

The key technical step is a Ramsey-theoretic statement (Lemma 2.4, stated below) which says that large sets S automatically have the property that some translate of S contains a certain structured subset (which we call an “increasing chain”)¹. This structured subset allows us to give strong upper bounds on the probability that all the vectors $(X_v)_{v \in S}$ lie in $B_n(0, p)$. Applying this to each $S \subseteq \mathbb{F}_2^\ell$ of cardinality $C\ell$ and taking a union bound gives Theorem 2.2.

To state the Ramsey-theoretic lemma (Lemma 2.4) which plays a central role in Theorem 2.2, we first define increasing chains. For a vector $v \in \mathbb{F}_2^\ell$, the *support* of v , denoted $\text{supp}(v)$, is defined to be the set of its nonzero coordinates.

Definition 2.3: A sequence of vectors $v_1, \dots, v_d \in \mathbb{F}_2^\ell$ is called a c -increasing chain of length d , if for all $j \in [d]$,

$$\left| \text{supp}(v_j) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i) \right) \right| \geq c.$$

The proof of Lemma 2.4 appears in Section V, where it is proved using the Sauer-Shelah lemma.

Lemma 2.4: For all positive integers c, ℓ and $L \leq 2^\ell$, the following holds. For every $S \subseteq \mathbb{F}_2^\ell$ with $|S| = L$, there is a $w \in \mathbb{F}_2^\ell$ such that $S + w$ has a c -increasing chain of length at least $\frac{1}{c}(\log_2 \frac{L}{2}) - (1 - \frac{1}{c})(\log_2 \ell)$.

2) *Linear codes over larger alphabets*: Due to their geometric nature, our arguments generalize to the case of q -ary alphabet (for arbitrary constant q) quite easily. Below we state our main theorem for the case of q -ary alphabet.

Theorem 2.5: Let q be a prime power and let $p \in (0, 1 - 1/q)$. Then there exist constants $C_{p,q}, \delta > 0$, such that for all $\varepsilon > 0$, letting $R = 1 - H_q(p) - \varepsilon$, if $C \subseteq \mathbb{F}_q^n$ is a random linear code of rate R , then

$$\Pr[C \text{ is } (p, \frac{C_{p,q}}{\varepsilon})\text{-list-decodable}] > 1 - 2^{-\delta n}.$$

The proof of Theorem 2.5 has the same basic outline as the proof of Theorem 2.1. In particular, it proceeds via a q -ary analog of Theorem 2.2. The only notable deviation occurs in the proof of the q -ary analog of Lemma 2.4. The traditional generalization of the Sauer-Shelah lemma to larger alphabets turns out to be unsuitable for this purpose. Instead,

¹This can broadly be viewed as coming under the umbrella of Ramsey theory, whose underlying theme is that large enough objects must contain structured sub-objects.

we formulate and prove a non-standard generalization of the Sauer-Shelah lemma for the larger alphabet case which is more appropriate for this situation. Details appear in Section VI.

III. PROOF OF THEOREM 2.1

Let us start by restating our main theorem.

Theorem 2.1 (restated) *Let $p \in (0, 1/2)$. Then there exist constants $C_p, \delta > 0$, such that for all $\varepsilon > 0$ and all large enough integers n , letting $R = 1 - H(p) - \varepsilon$, if $\mathcal{C} \subseteq \mathbb{F}_2^n$ is a random linear code of rate R , then*

$$\Pr[\mathcal{C} \text{ is } (p, \frac{C_p}{\varepsilon})\text{-list-decodable}] > 1 - 2^{-\delta n}.$$

Proof: Pick $C_p = 4C$, where C is the constant from Theorem 2.2. Take $\delta = 1$ and $L = \lceil \frac{C_p}{\varepsilon} \rceil$. Finally, let n be bigger than L and also large enough for the $o(1)$ term in that theorem to be at most 1.

Let \mathcal{C} be a random Rn dimensional linear subspace of \mathbb{F}_2^n . We want to show that

$$\Pr[\exists x \in \mathbb{F}_2^n \text{ s.t. } |B_n(x, p) \cap \mathcal{C}| \geq L] < 2^{-\delta n}. \quad (1)$$

Let $x \in \mathbb{F}_2^n$ be picked uniformly at random. We will work towards Equation (1) by studying the following quantity.

$$\Delta \stackrel{\text{def}}{=} \Pr_{\mathcal{C}, x}[|B_n(x, p) \cap \mathcal{C}| \geq L].$$

Note that to prove Equation (1), it suffices to show that²

$$\Delta < 2^{-\delta n} \cdot 2^{-n}. \quad (2)$$

We first show that one can move the center to the origin, and thus it is enough to understand the probability of the event that too many codewords of a random linear code fall in $B_n(0, p)$ (i.e., have Hamming weight at most pn). To this end, we note that

$$\begin{aligned} \Delta &= \Pr_{\mathcal{C}, x}[|B_n(x, p) \cap \mathcal{C}| \geq L] \\ &= \Pr_{\mathcal{C}, x}[|B_n(0, p) \cap (\mathcal{C} + x)| \geq L] \\ &\leq \Pr_{\mathcal{C}, x}[|B_n(0, p) \cap (\mathcal{C} + \{0, x\})| \geq L] \\ &\leq \Pr_{\mathcal{C}^*}[|B_n(0, p) \cap \mathcal{C}^*| \geq L] \end{aligned} \quad (3)$$

where \mathcal{C}^* is a random $Rn+1$ dimensional subspace containing $\mathcal{C} + \{0, x\}$ (explicitly, if $x \notin \mathcal{C}$, then $\mathcal{C}^* = \mathcal{C} + \{0, x\}$; otherwise $\mathcal{C}^* = \mathcal{C} + \{0, y\}$ where y is picked uniformly from $\mathbb{F}_2^n \setminus \mathcal{C}$). In particular, \mathcal{C}^* is a uniformly random $Rn+1$ dimensional subspace.

Now for each integer ℓ in the range $\log_2 L \leq \ell \leq L$, let \mathcal{F}_ℓ be the set of all tuples $(v_1, \dots, v_\ell) \in B_n(0, p)^\ell$ such that v_1, \dots, v_ℓ are linearly independent and $|\text{span}(v_1, \dots, v_\ell) \cap B_n(0, p)| \geq L$. Let

$$\mathcal{F} = \bigcup_{\ell=\lceil \log_2 L \rceil}^L \mathcal{F}_\ell.$$

²We could even replace the 2^{-n} by $2^{-(1-R)n}$. Indeed, for every \mathcal{C} for which there is a “bad” x , we know that there are 2^{Rn} “bad” x ’s (the translates of x by \mathcal{C}).

For each $\mathbf{v} = (v_1, \dots, v_\ell) \in \mathcal{F}$, let $\{\mathbf{v}\}$ denote the set $\{v_1, \dots, v_\ell\}$.

Towards bounding the probability in (3), notice that if $|B_n(0, p) \cap \mathcal{C}^*| \geq L$, then there must be some $\mathbf{v} \in \mathcal{F}$ for which $\mathcal{C}^* \supseteq \{\mathbf{v}\}$. Indeed, we can simply take \mathbf{v} to be a maximal linearly independent subset of $B_n(0, p) \cap \mathcal{C}^*$ if this set has size at most L , and any linearly independent subset of $B_n(0, p) \cap \mathcal{C}^*$ of size L otherwise.

Therefore, by the union bound,

$$\Delta \leq \sum_{\mathbf{v} \in \mathcal{F}} \Pr_{\mathcal{C}^*}[\mathcal{C}^* \supseteq \{\mathbf{v}\}] = \sum_{\ell=\lceil \log_2 L \rceil}^L \sum_{\mathbf{v} \in \mathcal{F}_\ell} \Pr_{\mathcal{C}^*}[\mathcal{C}^* \supseteq \{\mathbf{v}\}] \quad (4)$$

The last probability can be bounded as follows. By the linear independence of v_1, \dots, v_ℓ , the probability that $v_j \in \mathcal{C}^*$ conditioned on $\{v_1, \dots, v_{j-1}\} \subseteq \mathcal{C}^*$ is precisely the probability that a given nonzero point in a $n+1-j$ dimensional space lies in a random $Rn+2-j$ dimensional subspace, and hence this conditional probability is exactly

$$\frac{2^{Rn+2-j} - 1}{2^{n+1-j} - 1},$$

which we will bound from above by 2^{Rn+1-n} . We can hence conclude that for $\mathbf{v} \in \mathcal{F}_\ell$

$$\Pr_{\mathcal{C}^*}[\mathcal{C}^* \supseteq \{\mathbf{v}\}] \leq \left(\frac{2^{Rn+1}}{2^n} \right)^\ell. \quad (5)$$

Putting together Equations (4) and (5), we have

$$\Delta \leq \sum_{\ell=\lceil \log_2 L \rceil}^L \sum_{\mathbf{v} \in \mathcal{F}_\ell} \left(\frac{2^{Rn+1}}{2^n} \right)^\ell = \sum_{\ell=\lceil \log_2 L \rceil}^L |\mathcal{F}_\ell| \cdot \left(\frac{2^{Rn+1}}{2^n} \right)^\ell \quad (6)$$

We now obtain an upper bound on $|\mathcal{F}_\ell|$. We have two cases depending on the size of ℓ .

- **Case 1:** $\ell < 4/\varepsilon$.

In this case, note that $\frac{|\mathcal{F}_\ell|}{|B_n(0, p)|^\ell}$ is a lower bound on the probability that ℓ points X_1, \dots, X_ℓ chosen uniformly at random from $B_n(0, p)$ are such that

$$|\text{span}(\{X_1, \dots, X_\ell\}) \cap B_n(0, p)| \geq L.$$

Since $L \geq C \cdot \ell$, Theorem 2.2 tells us that this probability is bounded from above by 2^{-5n} . Thus, in this case

$$|\mathcal{F}_\ell| \leq |B_n(0, p)|^\ell 2^{-5n} \leq 2^{n\ell H(p)} \cdot 2^{-5n}.$$

- **Case 2:** $\ell \geq 4/\varepsilon$. In this case, we have the trivial bound of $|\mathcal{F}_\ell| \leq |B_n(0, p)|^\ell \leq 2^{n\ell H(p)}$.

Plugging this in (6), we may bound Δ by:

$$\begin{aligned}
\Delta &\leq \sum_{\ell=\log_2 L}^{\lceil 4/\varepsilon \rceil - 1} |\mathcal{F}_\ell| \cdot \left(\frac{2^{Rn+1}}{2^n}\right)^\ell + \sum_{\ell=\lceil 4/\varepsilon \rceil}^L |\mathcal{F}_\ell| \cdot \left(\frac{2^{Rn+1}}{2^n}\right)^\ell \\
&\leq \sum_{\ell=\log_2 L}^{\lceil 4/\varepsilon \rceil - 1} 2^{n\ell H(p)} 2^{-5n} \left(\frac{2^{Rn+1}}{2^n}\right)^\ell \\
&\quad + \sum_{\ell=\lceil 4/\varepsilon \rceil}^L 2^{n\ell H(p)} \left(\frac{2^{Rn+1}}{2^n}\right)^\ell \\
&\leq 2^{-5n} \sum_{\ell=\log_2 L}^{\lceil 4/\varepsilon \rceil - 1} 2^{(-\varepsilon n+1)\ell} + \sum_{\ell=\lceil 4/\varepsilon \rceil}^L 2^{(-\varepsilon n+1)\ell} \\
&\quad \text{(since } R = 1 - H(p) - \varepsilon) \\
&\leq 2^{-5n} \cdot 4/\varepsilon + 2^{L+1} \cdot 2^{(-\varepsilon n) \cdot (4/\varepsilon)} \\
&\leq n 2^{-5n} + 2^{-3n} \quad \text{(since } n \text{ is large enough)} \\
&< 2^{-\delta n} \cdot 2^{-n},
\end{aligned}$$

which gives us the Claim (2). \blacksquare

IV. PROOF OF THEOREM 2.2

In this section, we prove Theorem 2.2 which bounds the probability that the span of ℓ random points in $B_n(0, p)$ intersects $B_n(0, p)$ in at least $C \cdot \ell$ points, for some large constant C . We use the following simple fact.

Lemma 4.1: For every $p \in (0, 1/2)$, there is a $\delta_p > 0$ such that for all large enough integers n and every $x \in \mathbb{F}_2^n$, the probability that two uniform independent samples w_1, w_2 from $B_n(0, p)$ are such that $w_1 + w_2 \in B_n(x, p)$ is at most $2^{-\delta_p n}$.

Proof: The intuitive reason that this is true is that the point $w_1 + w_2$ is essentially a random point in $B_n(0, 2p - 2p^2)$ and the probability that it lies in the smaller ball $B_n(x, p)$ is maximal when $x = 0$ and is then exponentially small. Let us argue this formally.

Suppose the probability that $w_1 + w_2 \in B_n(x, p)$ is p' . Then there are integers i_1 and i_2 such that the probability that w_j has Hamming weight i_j for $j = 1, 2$ and $w_1 + w_2 \in B_n(x, p)$ is at least $p'n^{-2}$.

Let $\varepsilon > 0$ be a constant such that $2(p-\varepsilon) - 2(p-\varepsilon)^2 \geq p + \varepsilon$. Such an ε exists as $2p - 2p^2 > p$ whenever $p < \frac{1}{2}$. The value of ε depends on p but let us suppress this for notational convenience. If i_1 or i_2 is smaller than $(p - \varepsilon)n$ we can conclude that p' is $2^{-\delta_p n}$ as the subset of $B_n(x, p)$ that has Hamming weight $(p - \varepsilon)n$ is exponentially small. Otherwise let us consider the same probability when each coordinate of w_j is chosen to be one with probability i_j/n independently. As the probability that we get exactly i_j ones in w_j under this distribution is $\Omega(n^{-1/2})$, we can conclude that in this case the probability that $w_1 + w_2 \in B_n(x, p)$ is $\Omega(p'n^{-3})$. Now $w_1 + w_2 \in B_n(x, p)$ exactly when $x + w_1 + w_2 \in B_n(0, p)$, so let us look at the coordinates of $x + w_1 + w_2$. If $x_l = 0$ then the l 'th coordinate of this vector is 1 with probability at least $2(p - \varepsilon) - 2(p - \varepsilon)^2 \geq p + \varepsilon$ and if $x_l = 1$ the probability of the l 'th coordinate being 1 is at least $1/2$. As these choices are independent for different coordinates we conclude that the probability that $x + w_1 + w_2 \in B_n(0, p)$ is at most $2^{-\delta_p n}$ for

some $\delta_p' > 0$. Hence $p' \leq O(n^3 2^{-\delta_p' n}) \leq 2^{-\delta_p n}$ for suitable $\delta_p > 0$. \blacksquare

Let us now restate the theorem to be proved.

Theorem 2.2 (restated) For every $p \in (0, 1/2)$, there is a constant $C > 1$, such that for all n and all $\ell = o(\sqrt{n})$, if X_1, \dots, X_ℓ are picked independently and uniformly at random from $B_n(0, p)$, then

$$\Pr\left[\left|\text{span}(\{X_1, \dots, X_\ell\}) \cap B_n(0, p)\right| \geq C \cdot \ell\right] \leq 2^{-(6-o(1))n}.$$

Proof: Set $L = C \cdot \ell$ and let $c = 2$. Let $\delta_p > 0$ be the constant given by Lemma 4.1. Let

$$d = \left\lfloor \frac{1}{c} \log_2 \frac{L}{2} - \left(1 - \frac{1}{c}\right) \log_2 \ell \right\rfloor \geq \frac{1}{2} \log_2 \frac{L}{2\ell} - 1 = \frac{1}{2} \log_2 \frac{C}{8}.$$

For a vector $u \in \mathbb{F}_2^\ell$, let X_u denote the random variable $\sum_i u_i X_i$.

We begin with a claim which bounds the probability of a particular collection of linear combinations of the X_i all lying within $B_n(0, p)$. At the heart of this claim lies the Ramsey-theoretic Lemma 2.4.

Claim 4.2: For each $S \subseteq \mathbb{F}_2^\ell$ with $|S| = L + 1$,

$$\Pr[\forall v \in S, X_v \in B_n(0, p)] < 2^n \cdot 2^{-\delta_p d n}. \quad (7)$$

Proof: Let w and $v_1, \dots, v_d \in S$ be as given by Lemma 2.4. That is, $v_1 + w, v_2 + w, \dots, v_d + w$ is a c -increasing sequence (recall that we have set $c = 2$). Then,

$$\begin{aligned}
&\Pr[\forall v \in S, X_v \in B_n(0, p)] \\
&\leq \Pr[\forall j \in [d], X_{v_j} \in B_n(0, p)] \\
&= \Pr[\forall j \in [d], X_{v_j} + X_w \in B_n(X_w, p)] \\
&= \Pr[\forall j \in [d], X_{v_j + w} \in B_n(X_w, p)] \quad (8)
\end{aligned}$$

We now bound the probability that there exists $y \in \mathbb{F}_2^n$ such that for all $j \in [d]$, $X_{v_j + w} \in B_n(y, p)$. Fix $y \in \mathbb{F}_2^n$. We have:

$$\begin{aligned}
&\Pr[\forall j \in [d], X_{v_j + w} \in B_n(y, p)] \\
&= \prod_{j=1}^d \Pr\left[X_{v_j + w} \in B_n(y, p) \mid X_{v_i + w} \in B_n(y, p), \right. \\
&\quad \left. 1 \leq i \leq j - 1\right] \\
&\leq \prod_{j=1}^d \max_{\substack{a_t \in B_n(0, p) \\ t \in \bigcup_{i=1}^{j-1} \text{supp}(v_i + w)}} \Pr\left[X_{v_j + w} \in B_n(y, p) \mid \right. \\
&\quad \left. (X_t = a_t : t \in \bigcup_{i=1}^{j-1} \text{supp}(v_i + w))\right] \\
&\leq (2^{-\delta_p n})^d. \quad (9)
\end{aligned}$$

The last inequality follows from applying Lemma 4.1 as follows: Let i_1 and i_2 be two distinct elements of $\text{supp}(v_j + w) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i + w)\right)$, guaranteed to exist because of the 2-increasing property. We then apply Lemma 4.1 with w_1 and w_2 being vectors X_{i_1} and X_{i_2} , and with $x = y + \sum_{k \in [d], k \notin \{i_1, i_2\}} (v_j + w)_k X_k$.

Taking a union bound of Equation (9) over all $y \in \mathbb{F}_2^n$, we see that

$$\Pr[\exists y \in \mathbb{F}_2^n \text{ s.t. } \forall j \in [d], X_{v_j+w} \in B_n(y, p)] \leq 2^n \cdot 2^{-\delta_p n d}.$$

Combining this with Equation (8) completes the proof of the claim. \blacksquare

Given this claim, we now bound the probability that more than L elements of $\text{span}(\{X_1, \dots, X_\ell\})$ lie inside $B_n(0, p)$. This event occurs if and only if for some set $S \subseteq \mathbb{F}_2^\ell$ with $|S| = L + 1$, it is the case that $\forall v \in S, X_v \in B_n(0, p)$. Taking a union bound of (7) over all such S , we see that the probability that there exists some $S \subseteq \mathbb{F}_2^\ell$ with $|S| = L + 1$ such that $\forall v \in S, X_v \in B_n(0, p)$ is at most $2^{\ell(L+1)} \cdot 2^n \cdot 2^{-\delta_p d n}$. Taking C to be a large enough constant so that $d \geq \frac{1}{2} \log_2 \frac{C}{8} > \frac{12}{\delta_p}$, the theorem follows. \blacksquare

V. PROOF OF LEMMA 2.4

In this section, we will prove Lemma 2.4, which finds a large c -increasing chain in some translate of any large enough set $S \subseteq \mathbb{F}_2^\ell$.

We will use the Sauer-Shelah Lemma.

Lemma 5.1 (Sauer-Shelah [16], [17]): For all integers ℓ, c , and for any set $S \subseteq \{0, 1\}^\ell$, if $|S| > 2^{\ell c - 1}$, then there exists some set of coordinates $U \subseteq [\ell]$ with $|U| = c$ such that $|\{v|_U \mid v \in S\}| = \{0, 1\}^U$.

Lemma 2.4 (restated) For all positive integers c, ℓ and $L \leq 2^\ell$, the following holds. For every $S \subseteq \mathbb{F}_2^\ell$ with $|S| = L$, there is a $w \in \mathbb{F}_2^\ell$ such that $S + w$ has a c -increasing chain of length at least $\frac{1}{c}(\log_2 \frac{L}{2}) - (1 - \frac{1}{c})(\log_2 L)$.

Proof: We prove this by induction on ℓ . The claim holds trivially for $\ell \leq c$, so assume $\ell > c$.

If $L \leq 2^{\ell c - 1}$, then again the lemma holds trivially. Otherwise, by the Sauer-Shelah lemma, we get a set U of c coordinates such that for each $u \in \mathbb{F}_2^U$, there is some $v \in S$ such that $v|_U = u$. We will represent elements of \mathbb{F}_2^ℓ in the form (u, v') where $u \in \mathbb{F}_2^U$ and $v' \in \mathbb{F}_2^{[\ell] \setminus U}$.

Let $u_0 \in \mathbb{F}_2^U$ be a vector such that $|\{v \in S \mid v|_U = u_0\}|$ is at least $L/2^c$ (we know that such a u exists by averaging). Let $S' \subseteq \mathbb{F}_2^{[\ell] \setminus U}$ be given by $S' = \{v|_{[\ell] \setminus U} \mid v|_U = u_0\}$. By choice of u_0 , we have $|S'| \geq L/2^c$.

By the induction hypothesis, there exist $w' \in \mathbb{F}_2^{\ell - c}$ and $v'_1, \dots, v'_{d'} \in S'$ such that for each $j \in [d']$,

$$\left| \text{supp}(v'_j + w') \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v'_i + w') \right) \right| \geq c.$$

for $d' \geq \frac{1}{c} \log_2(L/2^{c+1}) - (1 - \frac{1}{c}) \log_2(\ell - c)$.

Let $d = d' + 1$. Note that $d \geq \frac{1}{c} \log_2(L/2) - (1 - \frac{1}{c}) \log_2(\ell - c) \geq \frac{1}{c} \log_2(L/2) - (1 - \frac{1}{c}) \log_2 \ell$. For $i \in [d']$, let $v_i = (u_0, v'_i) \in \mathbb{F}_2^\ell$. Let v_d be any vector in S with $(v_d)|_U = \neg u_0$, the bitwise complement of u_0 . Let $w = (u_0, w')$. We claim that w and v_1, \dots, v_d satisfy the desired properties.

³For sets A, B , we use the notation A^B to denote B -indexed tuples of elements from A .

Indeed, for each $j \in [d']$, we have

$$\begin{aligned} & \left| \text{supp}(v_j + w) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i + w) \right) \right| \\ &= \left| \text{supp}(v'_j + w') \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v'_i + w') \right) \right| \\ &\geq c. \end{aligned}$$

Also

$$\begin{aligned} & \left| \text{supp}(v_d + w) \setminus \left(\bigcup_{i=1}^{d-1} \text{supp}(v_i + w) \right) \right| \\ &\geq \left| \text{supp}(v_d + w) \setminus ([\ell] \setminus U) \right| \\ &= |U| = c. \end{aligned}$$

Thus for all $j \in [d]$, we have

$$\left| \text{supp}(v_j + w) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i + w) \right) \right| \geq c,$$

as desired. \blacksquare

VI. LARGER ALPHABETS

As mentioned in the introduction the case of q -ary alphabet is nearly identical to the case of binary alphabet. We only highlight the differences. As before, the crux turns out to be the problem of studying the weight distribution of certain random constant-dimensional codes.

Theorem 6.1 (Span of random points in $B_n^q(0, p)$): For every prime power q and every $p \in (0, 1 - 1/q)$, there is a constant $C > 1$, such that for all n and all $\ell = o(\sqrt{n})$, if X_1, \dots, X_ℓ are picked independently and uniformly at random from $B_n^q(0, p)$, then

$$\Pr[|\text{span}(\{X_1, \dots, X_\ell\}) \cap B_n^q(0, p)| > C \cdot \ell] \leq q^{-(6-o(1))n}.$$

The proof of Theorem 6.1 proceeds as before, by bounding the probability via a large c -increasing chain. The c -increasing chain itself is found in an analog of Lemma 2.4 for q -ary alphabet. We first need a definition.

Definition 6.2: A sequence of vectors $v_1, \dots, v_d \in [q]^\ell$ is called a c -increasing chain of length d , if for all $j \in [d]$,

$$\left| \text{supp}(v_j) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i) \right) \right| \geq c.$$

Now we have the following lemma.

Lemma 6.3 (q -ary increasing chains): For all positive integers c, ℓ , every prime power q , and $L \leq q^\ell$, the following holds. For every $S \subseteq \mathbb{F}_q^\ell$ with $|S| = L$, there is a $w \in \mathbb{F}_q^\ell$ such that $S + w$ has a c -increasing chain of length at least $\frac{1}{c} \log_q \left(\frac{L}{2} \right) - (1 - \frac{1}{c}) \log_q((q-1)\ell)$.

The proof of Lemma 6.3 needs a non-standard generalization of the Sauer-Shelah lemma to larger alphabet described in the next section.

A. A q -ary Sauer-Shelah lemma

The traditional generalization of the Sauer-Shelah lemma to large alphabets is the Karpovsky-Milman lemma [18], which roughly states that given $S \subseteq [q]^\ell$ of cardinality at least $(q-1)^{\ell c-1}$, there is a set U of c coordinates such that for every $u \in [q]^U$, there is some $v \in S$ such that the restriction $v|_U$ equals u . Applying this lemma in our context, once $q > 2$, requires us to have a set $S > 2^\ell$, which turns out to lead to exponential list-size bounds. Fortunately, the actual property needed for us is slightly different. We want a bound B (ideally polynomial in ℓ) such that for any set $S \subseteq [q]^\ell$ of cardinality at least B , there is a set U of c coordinates such that for every $u \in [q]^U$, there is some $v \in S$ such that the restriction $v|_U$ differs from u in every coordinate of U . It turns out that this weakened requirement admits polynomial-sized B .

We state and prove this generalization of the Sauer-Shelah lemma below.

Lemma 6.4 (q -ary Sauer-Shelah): For all integers q, ℓ, c , with $c \leq \ell$, for any set $S \subseteq [q]^\ell$, if $|S| > 2 \cdot ((q-1) \cdot \ell)^{c-1}$, then there exists some set of coordinates $U \subseteq [\ell]$ with $|U| = c$ such that for every $u \in [q]^U$, there exists some $v \in S$ such that u and $v|_U$ differ in every coordinate.

Proof: We prove this by induction on ℓ and c .

If $c = 1$, then $|S| > 2$ and the result holds by letting U equal any coordinate on which not all elements of S agree. If $\ell = c$, then $|S| > \ell(q-1)^{\ell-1}$, and the result follows from the observation that for every $u \in [q]^\ell$, there are at most $\ell(q-1)^{\ell-1}$ vectors v such that v and u agree in at least one coordinate.

We will now execute the induction step. Assume $c > 1$. Represent an element x of $[q]^\ell$ as a pair (y, b) , where $y \in [q]^{\ell-1}$ consists of the first $\ell-1$ coordinates of x and $b \in [q]$ is the last coordinate of x .

Consider the following subsets of $[q]^{\ell-1}$.

$$S_1 = \{y \in [q]^{\ell-1} \mid \text{for at least 1 value of } b \in [q], (y, b) \in S\}.$$

$$S_2 = \{y \in [q]^{\ell-1} \mid \text{for at least 2 values of } b \in [q], (y, b) \in S\}.$$

Note that $|S| \leq (|S_1| - |S_2|) + q|S_2| = |S_1| + (q-1)|S_2|$. By assumption,

$$\begin{aligned} |S| &> 2 \cdot ((q-1) \cdot \ell)^{c-1} \\ &\geq 2 \cdot ((q-1) \cdot (\ell-1))^{c-1} \\ &\quad + (q-1) (2 \cdot ((q-1) \cdot (\ell-1))^{c-2}), \end{aligned}$$

(using the elementary inequality $\ell^{c-1} \geq (\ell-1)^{c-1} + (\ell-1)^{c-2}$). Thus, either $|S_1| > 2 \cdot ((q-1) \cdot (\ell-1))^{c-1}$, or else $|S_2| > 2 \cdot ((q-1) \cdot (\ell-1))^{c-2}$.

We now prove the desired claim in each of these cases.

Case 1: $|S_1| > 2 \cdot ((q-1) \cdot (\ell-1))^{c-1}$. In this case, we can apply the induction hypothesis to S_1 with parameters $\ell-1$ and c , and get a subset of U of $[\ell-1]$ of cardinality c . Then the set U has the desired property.

Case 2: $|S_2| > 2 \cdot ((q-1) \cdot (\ell-1))^{c-2}$. In this case, we apply the induction hypothesis to S_2 with parameters $\ell-1$ and $c-1$, and get a subset U of $[\ell-1]$ of cardinality $c-1$. Then the set $U \cup \{\ell\}$ has the desired property. Indeed, take any vector $u \in [q]^{U \cup \{\ell\}}$. Let $u' = u|_U$. By the induction hypothesis, we

know that there is a $v \in S_2$ such that $v|_U$ differs from u' in every coordinate of U . Now we know that there are at least two $b \in [q]$ such that $(v, b) \in S$. At least one of these b will be such that (v, b) differs from u in every coordinate of $U \cup \{\ell\}$, as desired. ■

In the next section, we use the above lemma to prove the Ramsey-theoretic q -ary increasing chain claim (Lemma 6.3).

VII. PROOF OF q -ARY INCREASING CHAIN LEMMA

In this section, we prove Lemma 6.3, which we restate below for convenience.

Lemma 6.3 (restated) For every prime power q , and all positive integers c, ℓ and $L \leq q^\ell$, the following holds. For every $S \subseteq \mathbb{F}_q^\ell$ with $|S| = L$, there is a $w \in \mathbb{F}_q^\ell$ such that $S+w$ has a c -increasing chain of length at least $\frac{1}{c} \log_q \left(\frac{L}{2}\right) - \left(1 - \frac{1}{c}\right) \log_q((q-1)\ell)$.

Proof: We prove this by induction on ℓ . The claim holds trivially for $\ell \leq c$, so assume $\ell > c$.

If $L \leq 2((q-1) \cdot \ell)^{c-1}$, then again the lemma holds trivially. Otherwise, by Lemma 6.4 we get a set U of c coordinates such that for each $u \in \mathbb{F}_q^U$, there is some $v \in S$ such that $v|_U$ differs from u in every coordinate. We will represent elements of \mathbb{F}_q^ℓ in the form (u, v') where $u \in \mathbb{F}_q^U$ and $v' \in \mathbb{F}_q^{[\ell] \setminus U}$.

Let $u_0 \in \mathbb{F}_q^U$ be a vector such that $|\{v \in S \mid v|_U = u_0\}|$ is at least L/q^c (we know that such a u exists by averaging). Let $S' \subseteq \mathbb{F}_q^{[\ell] \setminus U}$ be given by $S' = \{v|_{[\ell] \setminus U} \mid v|_U = u_0\}$. By choice of u_0 , we have $|S'| \geq L/q^c$.

By the induction hypothesis, for

$$d' \geq \frac{1}{c} \log_q \left(\frac{L}{2q^c}\right) - \left(1 - \frac{1}{c}\right) \log_q((q-1)(\ell-c)),$$

there exist $w' \in \mathbb{F}_q^{\ell-c}$ and $v'_1, \dots, v'_{d'}$ in S' such that for each $j \in [d']$,

$$\left| \text{supp}(v'_j + w') \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v'_i + w') \right) \right| \geq c.$$

Let $d = d' + 1$. Note that

$$d \geq \frac{1}{c} \log_q \left(\frac{L}{2}\right) - \left(1 - \frac{1}{c}\right) \log_q((q-1)\ell).$$

For $i \in [d']$, let $v_i = (u_0, v'_i) \in \mathbb{F}_q^\ell$. Let v_d be any vector in S where $(v_d)|_U$ differs from u_0 in every coordinate of U . Let $w = (-u_0, w')$. We claim that w and v_1, \dots, v_d satisfy the desired properties.

Indeed, for each $j \in [d']$, we have

$$\begin{aligned} &\left| \text{supp}(v_j + w) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i + w) \right) \right| \\ &= \left| \text{supp}(v'_j + w') \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v'_i + w') \right) \right| \\ &\geq c. \end{aligned}$$

Also,

$$\begin{aligned} & \left| \text{supp}(v_d + w) \setminus \left(\bigcup_{i=1}^{d-1} \text{supp}(v_i + w) \right) \right| \\ & \geq \left| \text{supp}(v_d + w) \setminus ([\ell] \setminus U) \right| \\ & = |U| = c. \end{aligned}$$

Thus for all $j \in [d]$, we have

$$\left| \text{supp}(v_j + w) \setminus \left(\bigcup_{i=1}^{j-1} \text{supp}(v_i + w) \right) \right| \geq c,$$

as desired. \blacksquare

Given Lemma 6.3, the proof of Theorem 6.1 is virtually identical to the proof of its binary analog Theorem 2.2. Theorem 2.5 can then be proved (using Theorem 6.1) in the same manner as Theorem 2.1 was proved.

ACKNOWLEDGMENTS

Some of this work was done when we were all participating in the Dagstuhl seminar 09441 on constraint satisfaction. We thank the organizers of the seminar for inviting us, and Schloss Dagstuhl for the wonderful hospitality. We thank the anonymous reviewers for a careful and thorough reading of the paper and for several useful suggestions.

REFERENCES

- [1] P. Elias, "List decoding for noisy channels," *Technical Report 335, Research Laboratory of Electronics, MIT*, 1957.
- [2] J. M. Wozencraft, "List Decoding," *Quarterly Progress Report, Research Laboratory of Electronics, MIT*, vol. 48, pp. 90–95, 1958.
- [3] P. Elias, "Error-correcting codes for list decoding," *IEEE Transactions on Information Theory*, vol. 37, pp. 5–12, 1991.
- [4] V. V. Zyablov and M. S. Pinsker, "List cascade decoding," *Problems of Information Transmission*, vol. 17, no. 4, pp. 29–34, 1981 (in Russian); pp. 236–240 (in English), 1982.
- [5] A. Rudra, "Limits to list decoding random codes," in *COCOON*, ser. Lecture Notes in Computer Science, H. Q. Ngo, Ed., vol. 5609. Springer, 2009, pp. 27–36.
- [6] V. M. Blinovsky, "Bounds for codes in the case of list decoding of finite volume," *Problems of Information Transmission*, vol. 22, no. 1, pp. 7–19, 1986.
- [7] —, "Code bounds for multiple packings over a nonbinary finite alphabet," *Problems of Information Transmission*, vol. 41, no. 1, pp. 23–32, 2005.
- [8] —, "On the convexity of one coding-theory function," *Problems of Information Transmission*, vol. 44, no. 1, pp. 34–39, 2008.
- [9] V. Guruswami and A. D. Smith, "Codes for computationally simple channels: Explicit constructions with optimal rate," 2010, preprint, arxiv:1004.4017 [cs.IT]. To appear in the 51st IEEE Symposium on Foundations of Computer Science.
- [10] V. Guruswami and P. Indyk, "Expander-based constructions of efficiently decodable codes," in *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, 2001, pp. 658–667.
- [11] V. Guruswami, J. Håstad, M. Sudan, and D. Zuckerman, "Combinatorial bounds for list decoding," *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1021–1035, 2002.
- [12] V. M. Blinovsky, *Asymptotic Combinatorial Coding Theory*. Kluwer Academic Publishers, Boston, 1997.
- [13] V. K. Wei and G.-L. Feng, "Improved lower bounds on the sizes of error-correcting codes for list decoding," *IEEE Transactions on Information Theory*, vol. 40, no. 2, pp. 559–563, 1994.
- [14] V. M. Blinovsky, "Lower bound for the linear multiple packing of the binary Hamming space," *Journal of Combinatorial Theory, Series A*, vol. 92, no. 1, pp. 95–101, 2000.
- [15] V. Guruswami, *List decoding of error-correcting codes*, ser. Lecture Notes in Computer Science. Springer, 2004, no. 3282.
- [16] N. Sauer, "On the density of families of sets," *J. Combinatorial Theory Ser. A*, vol. 13, pp. 145–147, 1972.
- [17] S. Shelah, "A combinatorial problem; stability and order for models and theories in infinitary languages," *Pacific J. Math.*, vol. 41, pp. 247–261, 1972.
- [18] M. G. Karpovsky and V. D. Milman, "Coordinate density of sets of vectors," *Discrete Math.*, vol. 24, no. 2, pp. 177–184, 1978.

Venkatesan Guruswami is an Associate Professor in the Computer Science Department at Carnegie Mellon University. Dr. Guruswami received his Bachelor's degree from the Indian Institute of Technology at Madras in 1997 and his Ph.D. from the Massachusetts Institute of Technology in 2001. From 2002–09, he was a faculty member in the Department of Computer Science and Engineering at the University of Washington. He was a Miller Research Fellow at the University of California, Berkeley during 2001–02, and a member in the School of Mathematics, Institute for Advanced Study during 2007–08.

Dr. Guruswami's research interests span several topics such as the theory of error-correcting codes, approximation algorithms and hardness of approximation results for optimization problems, pseudorandomness and combinatorics, computational complexity theory, and algebraic algorithms. Dr. Guruswami is a recipient of a Packard Fellowship (2005), Sloan Fellowship (2005), NSF Career award (2004), ACM's Doctoral Dissertation Award (2002), and the IEEE Information Theory Society Paper Award (2000).

Johan Håstad is a Professor in the Computer Science Department at the Royal Institute of Technology. Dr. Håstad received his Bachelor's degree from Stockholm University in 1981 and his Ph.D. from the Massachusetts Institute of Technology in 1986. He was a post doc at Massachusetts Institute of Technology in 1986–1987 and joined the Royal Institute of Technology as a faculty member in 1988. Dr. Håstad was a member in the School of Mathematics, Institute for Advanced Study during 2000–01 and is a member of the Swedish Royal Academy of Science since 2001.

Dr. Håstad's research interests span several topics such as approximation algorithms and hardness of approximation results for optimization problems, cryptography, theory of error-correcting codes and computational complexity theory. Dr. Håstad is the recipient of the Göran Gustafsson prize (1999), Gödel prize (1994), and ACM's Doctoral Dissertation Award (1986).

Swastik Kopparty is a Member of the School of Mathematics at the Institute for Advanced Study. He received his Ph.D. in Computer Science from the Massachusetts Institute of Technology in 2010, under the supervision of Madhu Sudan. Before that, he received his B.S. in Computer Science from the University of California, Riverside in 2004.

Dr. Kopparty's interests are in coding theory, pseudorandomness, and computational complexity theory.