

A Lower Bound on List Size for List Decoding

Venkatesan Guruswami and Salil Vadhan

Abstract—A q -ary error-correcting code $C \subseteq \{1, 2, \dots, q\}^n$ is said to be *list decodable* to radius ρ with list size L if every Hamming ball of radius ρ contains at most L codewords of C . We prove that in order for a q -ary code to be list-decodable up to radius $(1-1/q)(1-\varepsilon)n$, we must have $L = \Omega(1/\varepsilon^2)$. Specifically, we prove that there exists a constant $c_q > 0$ and a function f_q such that for small enough $\varepsilon > 0$, if C is list-decodable to radius $(1-1/q)(1-\varepsilon)n$ with list size c_q/ε^2 , then C has at most $f_q(\varepsilon)$ codewords, independent of n . This result is asymptotically tight (treating q as a constant), since such codes with an exponential (in n) number of codewords are known for list size $L = O(1/\varepsilon^2)$. A result similar to ours is implicit in Blinovskiy [2] for the binary ($q = 2$) case. Our proof is simpler and works for all alphabet sizes, and provides more intuition for why the lower bound arises.

Index Terms—Bounds on codes, List decoding, Probabilistic method, Random codes, Randomness extractors.

I. INTRODUCTION

List decoding was introduced independently by Elias [3] and Wozencraft [4] as a relaxation of the classical notion of error-correction by allowing the decoder to output a *list* of possible answers. The decoding is considered successful as long as the correct message is included in the list. We point the reader to the paper by Elias [5] for a good summary of the history and context of list decoding.

The basic question raised by list decoding is the following: How many errors can one recover from, when constrained to output a list of small size? The study of list decoding strives to (1) understand the combinatorics underlying this question, (2) realize the bounds with explicit constructions of codes, and (3) list decode those codes with efficient algorithms. This work falls in the combinatorial facet of list decoding. Combinatorially, an error-correcting code has “nice” list-decodability properties if every Hamming ball of “large” radius has a “small” number of codewords in it. In this work, we are interested in exposing some combinatorial *limitations* on the performance of list-decodable codes. Specifically, we seek lower bounds on the list size needed to perform decoding up to a certain number of errors, or in other words, lower

bounds on the number of codewords that must fall inside *some* ball of specified radius centered at some point. We show such a result by picking the center in a certain probabilistic way. We now give some background definitions and terminology, followed by a description of our main result.

A. Preliminaries

We denote the set $\{1, 2, \dots, m\}$ by the shorthand $[m]$. For $q \geq 2$, a q -ary code of block length n is simply a subset of $[q]^n$. The elements of the code are referred to as *codewords*. The high-level property of a code that makes it useful for error-correction is its sparsity — the codewords must be well spread-out, so they are unlikely to distort into one another. One way to insist on sparsity is that the Hamming distance between every pair of distinct codewords is at least d . Note that this is equivalent to requiring that every Hamming ball of radius $\lfloor (d-1)/2 \rfloor$ has at most one codeword. Generalizing this, one can allow up to a small number, say L , of codewords in Hamming balls of certain radius. This leads to the notion of list decoding and a good list-decodable code. Since the expected Hamming distance of a random string of length n from any codeword is $(1-1/q) \cdot n$ for a q -ary code, the largest fraction of errors one can sensibly hope to correct is $(1-1/q)$. This motivates the following definition of a list-decodable code.

Definition 1.1: Let $q \geq 2$, $0 < \rho < 1$, and L be a positive integer. A q -ary code C of block length n is said to be (ρ, L) -list-decodable if for every $\mathbf{y} \in [q]^n$, the Hamming ball of radius $\rho \cdot (1-1/q) \cdot n$ centered at \mathbf{y} contains at most L codewords of C .

We will study (ρ, L) -list-decodable codes for $\rho = 1 - \varepsilon$ in the limit of $\varepsilon \rightarrow 0$. This setting is the one where list decoding is most beneficial, and is a clean setting to initially study the asymptotics. In particular, we will prove that, except for trivial codes whose size does not grow with n , $(1-\varepsilon, L)$ -list-decodable codes require list size $L = \Omega(1/\varepsilon^2)$ (hiding dependence on q).

B. Context and Related Results

Before stating our result, we describe some of the previously known results to elucidate the broader context where our work fits. The *rate* of a q -ary code of block length n is defined to be $\frac{\log_q |C|}{n}$. For $0 \leq x \leq 1$, we denote by $H_q(x)$ the q -ary entropy function, $H_q(x) = x \log_q(q-1) - x \log_q x - (1-x) \log_q(1-x)$.

Using the probabilistic method, it can be shown that (ρ, L) -list-decodable q -ary codes of rate $1 - H_q((1-1/q)\rho) - 1/L$ exist [5], [6]. In particular, in the limit of large L , we can achieve a rate of $1 - H_q((1-1/q)\rho)$, which equals both the Hamming bound and the Shannon capacity of the q -ary channel that changes a symbol $\alpha \in [q]$ to a uniformly random

A preliminary version of this paper appeared in *RANDOM '05* [1].

Venkatesan Guruswami is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. Email: guruswami@cmu.edu. This work was done in part when the author was at the Department of Computer Science and Engineering, University of Washington, Seattle, WA. This material is based upon work supported by the National Science Foundation under Grant Numbers CCF-0343672 and CCF-0953155. The author was also supported by a Sloan Fellowship and a David and Lucile Packard Fellowship.

Salil Vadhan is with the School of Engineering & Applied Sciences, Harvard University, Cambridge, MA 02138. Work done in part while a Fellow at the Radcliffe Institute for Advanced Study. Email: salil@seas.harvard.edu. This material is based upon work supported by the National Science Foundation under Grant Number CCF-0133096. Reserach was also supported by ONR grant N00014-04-1-0478, and a Sloan Research Fellowship.

element of $[q]$ with probability ρ and leaves α unchanged with probability $1 - (1 - 1/q)\rho$. When $\rho = 1 - \varepsilon$ for small ε , we have $H_q((1 - 1/q)\rho) = 1 - \Omega(q\varepsilon^2/\log q)$. Therefore, there exist $(1 - \varepsilon, L(q, \varepsilon))$ -list-decodable q -ary codes with $2^{\Omega(q\varepsilon^2 n)}$ codewords and $L(q, \varepsilon) = O(\frac{\log q}{q\varepsilon^2})$. In particular, for constant q , list size of $O(1/\varepsilon^2)$ suffices for non-trivial list decoding up to radius $(1 - 1/q) \cdot (1 - \varepsilon)$.

We are interested in whether this quadratic dependence on $1/\varepsilon$ in the list size is inherent. The quadratic bound is related to the $2\log(1/\varepsilon) - O(1)$ lower bound due to [7] for the amount of “entropy loss” in *randomness extractors*, which are well-studied objects in the subject of pseudorandomness. In fact, a lower bound of $\Omega(1/\varepsilon^2)$ on list size implies such an entropy loss bound for (“strong”) randomness extractors, using known connections between extractors and list-decodable codes [8], [9]. However, in the other direction, the connection loses a factor of ε in the lower bound, yielding only a lower of $\Omega(1/\varepsilon)$ for list size. (See the Appendix on randomness extractors.)

For the model of erasures, where up to a fraction $(1 - \varepsilon)$ of symbols are erased by the channel, optimal bounds of $\Theta(\log(1/\varepsilon))$ are known for the list size required for binary codes [10]. This can be compared with the $\log \log(1/\varepsilon) - O(1)$ lower bound on entropy loss for “dispersers,” which are a variant of randomness extractors [7].

A lower bound of $\Omega(1/\varepsilon^2)$ for list size L for $(1 - \varepsilon, L)$ -list-decodable *binary* codes follows from the work of Blinovsky [2]. We discuss more about his work and how it compares to our results in Section I-E.

C. Our Result

Our main result is a proof of the following fact: the smallest list size that permits list decoding up to radius $(1 - 1/q)(1 - \varepsilon)$ is $\Theta(\varepsilon^{-2})$ (hiding constants depending on q in the Θ -notation). The formal statement of our main result is below.

Theorem 1.2 (Main): For every integer $q \geq 2$ there exists $c_q > 0$ and $d_q < \infty$ such that for all small enough $\varepsilon > 0$, the following holds. If C is a q -ary $(1 - \varepsilon, c_q/\varepsilon^2)$ -list-decodable code, then $|C| \leq 2^{d_q \cdot \varepsilon^{-2} \log(1/\varepsilon)}$.

D. Overview of Proof

We now describe the high-level structure of our proof. Recall that our goal is to exhibit a center z that has several (specifically $\Omega(1/\varepsilon^2)$) codewords of C with large correlation, where we say two strings in $[q]^n$ have correlation ε if they agree in $(1/q + \varepsilon(1 - 1/q)) \cdot n$ locations.¹ Using the probabilistic method, it is not very difficult to prove the existence of such a center z and $\Omega(1/\varepsilon^2)$ codewords whose *average* correlation with z is at least $\Omega(\varepsilon)$. (This is the content of our Lemma 2.4.) This step is closely related to (and actually follows from) the known lower bound of Radhakrishnan and Ta-Shma [7] on the “entropy loss” of “randomness extractors.” (See the Appendix on randomness extractors.)

¹Correlation is defined in this way so that if we send a codeword over a channel that replaces each symbol with a uniformly random symbol with probability $1 - \varepsilon$, then the expected correlation between the codeword and received word is ε .

However, this large average could occur due to about $1/\varepsilon$ codewords having a $\Omega(1)$ correlation with z , whereas we would like to find many more (i.e., $\Omega(1/\varepsilon^2)$) codewords with smaller (i.e., $\Omega(\varepsilon)$) correlation. We get around this difficulty by working with a large subcode C' of C where such a phenomenon cannot occur. Roughly speaking, we will use the probabilistic method to prove the existence of a large “ L -pseudorandom” subcode C' , for which looking at any set of L codewords of C' *never* reveals any significant overall bias in terms of the most popular symbol (out of $[q]$). More formally, all ℓ -tuples, $\ell \leq L$, the average “plurality” (i.e., frequency of most frequent symbol) over all the coordinates isn’t much higher than ℓ/q . (This is the content of our Lemma 2.5.) This in turn implies that for every center z , the sum of the correlations of z with all codewords that have “large” correlation (say at least $D\varepsilon$, for a sufficiently large constant D) is small. Together with the high average correlation bound, this means several codewords must have “intermediate” correlation with z (between ε and $D\varepsilon$). The number of such codewords is our lower bound on list size.

E. Comparison with Blinovsky’s results [2], [11], [12]

As remarked earlier, a lower bound of $L = \Omega(1/\varepsilon^2)$ for *binary* $(1 - \varepsilon, L)$ -list-decodable codes follows from previous work of Blinovsky [2]. In this work, Blinovsky explores the tradeoff between ρ , L , and the relative rate γ of a (ρ, L) -list-decodable code, when all three of these parameters are constants and the block length n tends to infinity. In particular, it is shown that for any fixed L , the rate has to be strictly less than the list decoding capacity $1 - H(\rho/2)$. A special case of the main theorem in [2] shows that if $L \leq c/\varepsilon^2$ for some constant $c > 0$, then the rate γ must be zero asymptotically, which means that the code can have at most $2^{o(n)}$ codewords for block length n . A careful inspection of his proof, however, reveals an $f(\varepsilon)$ bound (independent of n) on the number of codewords in any such code. This is similar in spirit to our Theorem 1.2.

More recently, Blinovsky also showed a similar lower bound of $L \geq c_q/\varepsilon^2$ for the list size of $(1 - \varepsilon, L)$ -list decodable q -ary codes of positive rate, assuming the convexity of a certain function [11]. Subsequent to our work, Blinovsky established the necessary convexity criterion in [12]. Our work was the first to give a full proof of the list size lower bound for all alphabet sizes.

Further, our work compares favorably with [2] in the following two respects.

- 1) Our result is quantitatively stronger. The dependence $f(\varepsilon)$ of the bound on the size of the code in [2] is much worse than the $(1/\varepsilon)^{O(\varepsilon^{-2})}$ that we obtain. In particular, $f(\varepsilon)$ is at least an exponential tower of height $\Theta(1/\varepsilon^2)$ (and is in fact bigger than the Ackermann function of $1/\varepsilon$).
- 2) Our proof seems simpler and provides more intuition about why and how the lower bound arises.

We now comment on the proof method in [2] (a similar method is also used in [11]). As with our proof, the first step in the proof is a bound for the case when the average correlation (w.r.t every center) for every set of $L + 1$ codewords is small

(this is Theorem 2 in [2]). Note that this is a more stringent condition than requiring no set of $L+1$ codewords lie within a small ball. Our proof uses the probabilistic method to show the existence of codewords with large average correlation in any reasonable sized code. The proof in [2] is more combinatorial, and uses a counting argument to bound the size of the code when all subsets of $L+1$ codewords have low average correlation (with every center). But the underlying technical goal of the first step in both the approaches is the same.

The second step in Blinovsky's proof is to use this bound to obtain a bound for list-decodable codes. The high-level idea is to pick a subcode of the list-decodable code with certain nice properties so that the bound for average correlation can be turned into one for list decoding. This is also similar in spirit to our approach (Lemma 2.5). The specifics of how this is done are, however, quite different. The approach in [2] is to find a large subcode which is $(L+1)$ -equidistant, i.e., for every $k \leq L+1$, all subsets of k codewords have the same value for their k 'th order scalar product, which is defined as the integer sum over all coordinates of the product of the k symbols (from $\{0,1\}$) in that coordinate.² Such a subcode has the following useful property: in each subset of $L+1$ codewords, all codewords in the subset have the same agreement with the best center, i.e., the center obtained by taking their coordinate-wise majority, and moreover this value is independent of the choice of the subset of $L+1$ codewords. This in turn enables one to get a bound for list decoding from one for average correlation. The requirement of being $(L+1)$ -equidistant is a rather stringent one, and is achieved iteratively by ensuring k -equidistance for $k = 1, 2, \dots, L+1$ successively. Each stage incurs a rather huge loss in the size of the code, and thus the bound obtained on the size of the original code is an enormously large function of $1/\varepsilon$. We make do with a much weaker property than $(L+1)$ -equidistance, letting us pick a much larger subcode with the property we need. This translates into a good upper bound on the size of the original list-decodable code.

II. PROOF OF MAIN RESULT

We first begin with convenient measures of closeness between strings, the agreement and the correlation.

Definition 2.1 (Agreement and Correlation): For vectors $\mathbf{x}, \mathbf{y} \in [q]^n$, define their *agreement*, denoted $\text{agr}(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \#\{i : x_i = y_i\}$. Their *correlation* is the value $\text{corr}(\mathbf{x}, \mathbf{y}) \in [-1/(q-1), 1]$ such that $\text{agr}(\mathbf{x}, \mathbf{y}) = \frac{1}{q} + \left(1 - \frac{1}{q}\right) \cdot \text{corr}(\mathbf{x}, \mathbf{y})$.³

The standard notion of correlation between two strings in $\{1, -1\}^n$ is simply their dot product divided by n ; the definition above is a natural generalization to larger alphabets.

A very useful notion for us will be the plurality of a set of codewords.

Definition 2.2 (Plurality): For symbols $a_1, \dots, a_k \in [q]$, we define their *plurality* $\text{plur}(a_1, \dots, a_k) \in [q]$ to be the most

frequent symbol among a_1, \dots, a_k , breaking ties arbitrarily. We define the *plurality count* $\#\text{plur}(a_1, \dots, a_k) \in \mathbb{N}$ to be the number of times that $\text{plur}(a_1, \dots, a_k)$ occurs among a_1, \dots, a_k .

For vectors $\mathbf{c}_1, \dots, \mathbf{c}_k \in [q]^n$, we define $\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k) \in [q]^n$ to be the component-wise plurality, i.e. $\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k)_i = \text{plur}(c_{1i}, \dots, c_{ki})$.

We define $\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k)$ to be the average plurality count over all coordinates; i.e.,

$$\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k) = \frac{1}{n} \left[\sum_{i=1}^n \#\text{plur}(c_{1i}, \dots, c_{ki}) \right].$$

The reason pluralities will be useful to us is that they capture the maximum average correlation any vector has with a set of codewords:

Lemma 2.3: For all $\mathbf{c}_1, \dots, \mathbf{c}_k \in [q]^n$,

$$\begin{aligned} \arg \max_{\mathbf{z} \in [q]^n} \sum_{i=1}^k \text{agr}(\mathbf{z}, \mathbf{c}_i) &= \text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k), \text{ and} \\ \max_{\mathbf{z} \in [q]^n} \sum_{i=1}^k \text{agr}(\mathbf{z}, \mathbf{c}_i) &= \sum_{i=1}^k \text{agr}(\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k), \mathbf{c}_i) \\ &= \#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_k). \end{aligned}$$

Note that our goal of proving lower bound on list size is the same as proving that in every not too small code, there must be some center \mathbf{z} that has several (i.e. $\Omega(1/\varepsilon^2)$) close-by codewords, or in other words several codewords with large (i.e., at least ε) correlation. We begin by showing the existence of a center which has a large *average* correlation with a collection of several codewords. By Lemma 2.3, this is equivalent to finding a collection of several codewords whose total plurality count is large.

Lemma 2.4: For all integers $q \geq 2$, there exists a constant $b_q > 0$ such that for every positive integer $t \geq 37q$ and every code $C \subseteq [q]^n$ with $|C| \geq 2t$, there exist t distinct codewords $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_t \in C$ such that

$$\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_t) \geq \frac{t}{q} + \Omega\left(\sqrt{\frac{t}{q}}\right).$$

Equivalently, there exists a $\mathbf{z} \in [q]^n$ such that

$$\sum_{i=1}^t \text{corr}(\mathbf{z}, \mathbf{c}_i) \geq \Omega\left(\sqrt{\frac{t}{q}}\right). \quad (1)$$

Proof: Without loss of generality, assume $|C| = 2t$. Pick a subset $\{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_t\}$ from C , chosen uniformly at random among all t -element subsets of C . For $j = 1, \dots, n$, define the random variable $P_j = \#\text{plur}(c_{1j}, \dots, c_{tj})$ to be the plurality of the j 'th coordinates. By definition, $\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_t) = (1/n) \cdot \sum_{j=1}^n P_j$. Notice that P_j is always at least t/q , and we would expect the plurality to occasionally deviate from the lower bound. Indeed, Lemma A.3 shows that for any sequence of $2t$ elements of $[q]$, if we choose a random subset of half of them, the expected plurality count is $t/q + \Omega(\sqrt{t/q})$. Thus, $\mathbb{E}[P_j] = t/q + \Omega(\sqrt{t/q})$. So $\mathbb{E}[(1/n) \cdot \sum_j P_j] = t/q + \Omega(\sqrt{t/q})$, and thus the lemma follows by taking any $\mathbf{c}_1, \dots, \mathbf{c}_t$

²A slight relaxation of the $(L+1)$ -equidistance property is actually what is used in [2], but this description should suffice for the discussion here.

³Note that we find it convenient to work with agreement and correlation that are normalized by dividing by the length n .

that achieves the expectation. The equivalent reformulation in terms of correlation follows from Lemma 2.3 and the definition of correlation in terms of agreement (Definition 2.1). ■

For any $\varepsilon > 0$, the above lemma gives a center z and $t = \Omega(1/q\varepsilon^2)$ codewords $\mathbf{c}_1, \dots, \mathbf{c}_t$ such that the average correlation between z and $\mathbf{c}_1, \dots, \mathbf{c}_t$ is at least ε . This implies that at least $(\varepsilon/2) \cdot t$ of the \mathbf{c}_i 's have correlation at least $\varepsilon/2$ with z . Thus we get a list-size lower bound of $(\varepsilon/2) \cdot t = \Omega(1/q\varepsilon)$ for decoding from correlation $\varepsilon/2$.

Now we would like to avoid the ε factor loss in list size in the above argument. The reason it occurs is that the average correlation can be ε due to the presence of $\approx \varepsilon t$ of the \mathbf{c}_i 's having extremely high correlation with z . This is consistent with the code being list-decodable with list size $o(1/(q\varepsilon^2))$ for correlation ε , but it means that this code has very poor list-decoding properties at some higher correlations — e.g., having $\varepsilon t = \Omega(1/(q\varepsilon))$ codewords at correlation $\Omega(1)$, whereas we'd expect a "good" code to have only $O(1)$ such codewords. In our next (and main) lemma, we show that we can pick a subcode of the code where this difficulty does not occur. Specifically, if C has good list-decoding properties at correlation ε , we get a subcode that has good list-decoding properties at every correlation larger than ε .

Lemma 2.5 (Main technical lemma): For all positive integers $L, t, m \geq 2t$ and $q \geq 2$, and all small enough $\varepsilon > 0$, the following holds. Let C be a $(1 - \varepsilon, L)$ -list-decodable q -ary code of block length n with

$$|C| > 2L \cdot t \cdot m! / (m - t)! .$$

Then there exists a subcode $C' \subseteq C$, $|C'| \geq m$, such that for all positive integers $\ell \leq t$ and every $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\ell \in C'$,

$$\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_\ell) \leq \ell \left(\frac{1}{q} + \left(1 - \frac{1}{q}\right)\varepsilon + O\left(\frac{q^{3/2}}{\sqrt{\ell}}\right) \right)$$

Equivalently, for every $\mathbf{z} \in [q]^n$ and every $\mathbf{c}_1, \dots, \mathbf{c}_\ell \in C'$, we have

$$\sum_{i=1}^{\ell} \text{corr}(\mathbf{z}, \mathbf{c}_i) \leq \left(\varepsilon + O\left(\frac{q^{3/2}}{\sqrt{\ell}}\right) \right) \cdot \ell . \quad (2)$$

Notice that the lemma implies a better upper bound on list size for correlations much larger than ε . More precisely, for every $\delta > 0$, it implies that the number of codewords having correlation at least $\varepsilon + \delta$ with a center z is at most $\ell = O(q^3/\delta^2)$. In fact, any ℓ codewords must even have average correlation at most $\varepsilon + \delta$.

Proof: We will pick a subcode $C' \subseteq C$ of size m at random from all m -element subsets of C , and prove that C' will fail to have the claimed property with probability less than 1.

For now, however, think of the code C' as being fixed, and we will reduce proving the desired properties above to bounding some simpler quantities. Let $(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\ell)$ be an arbitrary ℓ -tuple of codewords in C' . We will keep track of the average plurality count $\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_i)$ as we add each codeword to this sequence. To describe how this quantity can change at each step, we need a couple of additional definitions. We say a sequence $(a_1, \dots, a_i) \in [q]^i$ has a

plurality tie if at least two symbols occur $\#\text{plur}(a_1, \dots, a_i)$ times among a_1, \dots, a_i . For vectors $\mathbf{c}_1, \dots, \mathbf{c}_i \in [q]^n$, we define $\#\text{ties}(\mathbf{c}_1, \dots, \mathbf{c}_i)$ to be the fraction of coordinates $j \in [n]$ such that (c_{1j}, \dots, c_{ij}) has a plurality tie. Then:

Claim 2.6: For every $\mathbf{c}_1, \dots, \mathbf{c}_i \in [q]^n$,

$$\begin{aligned} \#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_i) &\leq \#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1}) + \\ &+ \text{agr}(\mathbf{c}_i, \text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1})) + \#\text{ties}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1}) . \end{aligned}$$

Proof of Claim: Consider each coordinate $j \in [n]$ separately. Clearly,

$$\#\text{plur}(c_{1j}, \dots, c_{ij}) \leq \#\text{plur}(c_{1j}, \dots, c_{(i-1)j}) + 1 .$$

Moreover, if $(c_{1j}, \dots, c_{(i-1)j})$ does not have a plurality tie, then the plurality increases if and only if c_{ij} equals the unique symbol $\text{plur}(c_{1j}, \dots, c_{(i-1)j})$ achieving the plurality. Thus,

$$\#\text{plur}(c_{1j}, \dots, c_{ij}) \leq \#\text{plur}(c_{1j}, \dots, c_{(i-1)j}) + A_j + T_j ,$$

where T_j is the indicator variable for $(c_{1j}, \dots, c_{(i-1)j})$ having a plurality tie, and A_j for c_{ij} agreeing with $\text{plur}(c_{1j}, \dots, c_{(i-1)j})$. The claim follows by averaging over $j = 1, \dots, n$. ■

Thus, our task of bounding $\#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_\ell)$ reduces to bounding $\text{agr}(\mathbf{c}_i, \text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1}))$ and $\#\text{ties}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1})$ for each $i = 1, \dots, \ell$. The first term we bound using the list-decodability of C and the random choice of the subcode C' .

Claim 2.7: There exists a choice of the subcode C' such that $|C'| = m$ and for every $i \leq t$ and every (ordered) sequence $\mathbf{c}_1, \dots, \mathbf{c}_i \in C'$, we have

$$\text{agr}(\mathbf{c}_i, \text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1})) \leq 1/q + (1 - 1/q) \cdot \varepsilon .$$

Proof of Claim: We choose the subcode C' uniformly at random from all m -subsets of C . We view C' as a sequence of m codewords selected randomly from C without replacement. Consider any i of the codewords $\mathbf{c}_1, \dots, \mathbf{c}_i$ in this sequence. By the $(1 - \varepsilon, L)$ -list decodability of the code C , for any $\mathbf{c}_1, \dots, \mathbf{c}_{i-1}$, there are at most L choices for \mathbf{c}_i having agreement larger than $(1/q + (1 - 1/q) \cdot \varepsilon)$ with $\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1})$. Conditioned on $\mathbf{c}_1, \dots, \mathbf{c}_{i-1}$, \mathbf{c}_i is distributed uniformly on the remaining $|C| - i + 1$ elements of C , so the probability of \mathbf{c}_i being one of the $\leq L$ bad codewords is at most $L/(|C| - i + 1)$.

By a union bound, the probability that the claim fails for at least one subsequence $\mathbf{c}_1, \dots, \mathbf{c}_i$ of at most t codewords in C' is at most

$$\sum_{i=1}^t \frac{m!}{(m-i)!} \cdot \frac{L}{|C| - i + 1} \leq t \cdot \frac{m!}{(m-t)!} \cdot \frac{L}{|C| - t + 1} < 1 .$$

Thus, there exists a choice of subcode C' satisfying the claim. ■

For the $\#\text{ties}(\mathbf{c}_1, \dots, \mathbf{c}_{i-1})$ terms, we consider the codewords $\mathbf{c}_1, \dots, \mathbf{c}_\ell$ in a random order.

Claim 2.8: For every sequence of $\mathbf{c}_1, \dots, \mathbf{c}_\ell \in [q]^n$, there exists a permutation $\sigma : [\ell] \rightarrow [\ell]$ such that

$$\sum_{i=1}^{\ell} \#\text{ties}(\mathbf{c}_{\sigma(1)}, \dots, \mathbf{c}_{\sigma(i)}) = O(q^{3/2} \cdot \sqrt{\ell}) .$$

Proof of Claim: We choose σ uniformly at random from all permutations $\sigma : [\ell] \rightarrow [\ell]$ and show that the expectation of the left side is at most $O(q^{3/2} \cdot \sqrt{\ell})$. By linearity of expectations, it suffices to consider the expected number of plurality ties occurring in each coordinate $j \in [n]$. That is, we read the symbols $c_{1j}, \dots, c_{\ell j} \in [q]$ in a random order σ and count the number of prefixes $c_{\sigma(1)j}, \dots, c_{\sigma(i)j}$ having a plurality tie. If this prefix were i symbols chosen independently according to some (arbitrary) distribution, then it is fairly easy to show that the probability of a tie is $O(1/\sqrt{i})$ (ignoring the dependence on q), and summing this from $i = 1, \dots, \ell$ gives $O(\sqrt{\ell})$ expected ties in each coordinate. Since they are not independently chosen, but rather i distinct symbols from a fixed sequence of ℓ symbols, the analysis becomes a bit more involved, but nevertheless the bound remains essentially the same. Specifically, in Lemma A.5, the expected number of ties is shown to be $O(q^{3/2} \cdot \sqrt{\ell})$, yielding the claim. ■

Now to complete the proof of Lemma 2.5, let C' be as in Claim 2.7, and let $\mathbf{c}_1, \dots, \mathbf{c}_\ell$ be an arbitrary sequence of distinct codewords in C' . Let σ be permutation guaranteed by Claim 2.8. Then, by Claim 2.6, we have

$$\begin{aligned} \#\text{plur}(\mathbf{c}_1, \dots, \mathbf{c}_\ell) &= \#\text{plur}(\mathbf{c}_{\sigma(1)}, \dots, \mathbf{c}_{\sigma(\ell)}) \\ &\leq \sum_{i=1}^{\ell} \left[\text{agr}(\mathbf{c}_{\sigma(i)}, \text{plur}(\mathbf{c}_{\sigma(1)}, \dots, \mathbf{c}_{\sigma(i-1)})) \right. \\ &\quad \left. + \#\text{ties}(\mathbf{c}_{\sigma(1)}, \dots, \mathbf{c}_{\sigma(i-1)}) \right] \\ &\leq \ell \cdot (1/q + (1 - 1/q) \cdot \varepsilon) + O(q^{3/2} \cdot \sqrt{\ell}), \end{aligned}$$

as desired. The equivalent reformulation in terms of correlation again follows from Lemma 2.3 and the definition of correlation in terms of agreement (Definition 2.1). ■

The following corollary of Lemma 2.5 will be useful in proving our main result.

Corollary 2.9: Let $L, t, m, q, \varepsilon, C$, and C' be as in Lemma 2.5 for a choice of parameters satisfying $t \geq L$. Then for all $\mathbf{z} \in [q]^n$ and all $D \geq 2$,

$$\sum_{\substack{\mathbf{c} \in C' \\ \text{corr}(\mathbf{z}, \mathbf{c}) \geq D\varepsilon}} \text{corr}(\mathbf{z}, \mathbf{c}) \leq O\left(\frac{q^3}{D\varepsilon}\right). \quad (3)$$

Proof: Let $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_r$ be all the codewords of C' that satisfy $\text{corr}(\mathbf{z}, \mathbf{c}) \geq D\varepsilon$. Since C , and hence C' , is $(1 - \varepsilon, L)$ -list-decodable, we have $r \leq L \leq t$. Using (2) for the codewords $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_r$, we have

$$D\varepsilon \leq \frac{1}{r} \sum_{i=1}^r \text{corr}(\mathbf{z}, \mathbf{c}_i) \leq \varepsilon + O\left(\frac{q^{3/2}}{\sqrt{r}}\right)$$

which gives $r = O(q^3 / ((D - 1)^2 \varepsilon^2)) = O(q^3 / (D^2 \varepsilon^2))$, since $D \geq 2$. Applying (2) again,

$$\begin{aligned} \sum_{\substack{\mathbf{c} \in C' \\ \text{corr}(\mathbf{z}, \mathbf{c}) \geq D\varepsilon}} \text{corr}(\mathbf{z}, \mathbf{c}) &= \sum_{i=1}^r \text{corr}(\mathbf{z}, \mathbf{c}_i) \\ &\leq \varepsilon r + O(q^{3/2} \sqrt{r}) \\ &\leq O\left(\frac{q^3}{D^2 \varepsilon}\right) + O\left(\frac{q^3}{D\varepsilon}\right) \\ &\leq O\left(\frac{q^3}{D\varepsilon}\right). \end{aligned}$$

We are now ready to prove our main result, Theorem 1.2, which we restate (in slightly different form) below.

Theorem 2.10 (Main): There exist constants $c > 0, d < \infty$, such that for all small enough $\varepsilon > 0$, the following holds. Suppose C is a q -ary $(1 - \varepsilon, L)$ -list-decodable code with $|C| > 1/(q\varepsilon^2)^{d/(q\varepsilon^2)}$. Then $L \geq c/(q^5 \varepsilon^2)$.

Proof: Let T be a large enough constant to be specified later. Let $t = \lfloor \frac{1}{Tq\varepsilon^2} \rfloor$. If $L > t$, then there is nothing to prove. So assume that $t \geq L \geq 1$ and set $m = 2t$. Then

$$2L \cdot t \cdot \frac{m!}{(m-t)!} \leq 2t^2 \cdot (2t)^t = \left(\frac{1}{q\varepsilon^2}\right)^{O(1/(q\varepsilon^2))} < |C|,$$

for a sufficiently large choice of the constant d . Let C' be a subcode of C of size $m = 2t$ guaranteed by Lemma 2.5.

By Lemma 2.4, there exist t codewords $\mathbf{c}_i, 1 \leq i \leq t$, in C' , and a center $\mathbf{z} \in [q]^n$ such that

$$\sum_{i=1}^t \text{corr}(\mathbf{z}, \mathbf{c}_i) = \Omega\left(\sqrt{\frac{t}{q}}\right). \quad (4)$$

Also, for any $D \geq 2$, we have $\sum_{i=1}^t \text{corr}(\mathbf{z}, \mathbf{c}_i)$ equals

$$\begin{aligned} &\sum_{i: \text{corr}(\mathbf{z}, \mathbf{c}_i) < \varepsilon} \text{corr}(\mathbf{z}, \mathbf{c}_i) + \sum_{i: \varepsilon \leq \text{corr}(\mathbf{z}, \mathbf{c}_i) < D\varepsilon} \text{corr}(\mathbf{z}, \mathbf{c}_i) \\ &+ \sum_{i: \text{corr}(\mathbf{z}, \mathbf{c}_i) \geq D\varepsilon} \text{corr}(\mathbf{z}, \mathbf{c}_i) \leq \varepsilon t + D\varepsilon L + O\left(\frac{q^3}{D\varepsilon}\right) \end{aligned} \quad (5)$$

where to bound the second part we used that C' is $(1 - \varepsilon, L)$ -list-decodable, and to bound the third part we used the fact that C' satisfies (3).

Putting these together, and setting $D = q^4 \cdot T$, we have

$$\begin{aligned} \Omega\left(\sqrt{\frac{t}{q}}\right) &\leq \varepsilon t + D\varepsilon L + O\left(\frac{q^3}{D\varepsilon}\right) \\ &\leq \sqrt{\frac{t}{Tq}} + D\varepsilon L + O\left(\sqrt{\frac{t}{Tq}}\right). \end{aligned}$$

For a sufficiently large choice of the constant T , this gives

$$L \geq \frac{1}{D\varepsilon} \cdot \Omega\left(\sqrt{\frac{t}{q}}\right) = \Omega\left(\frac{1}{T^{1.5} \cdot q^5 \cdot \varepsilon^2}\right).$$

as desired. ■

III. CONCLUDING REMARKS AND OPEN QUESTIONS

Several questions are open in the general direction of exhibiting limitations on the performance of list-decodable codes. We conclude with some remarks and mention some of these open questions below.

- We have not attempted to optimize the dependence on the alphabet size q in our bound on list size (i.e. the constant c_q in Theorem 1.2), and this leaves a gap between the upper and lower bounds. The probabilistic code construction of [5], [6] achieves a nearly linear dependence on q (specifically, list size $L = O(\log q/(q\varepsilon^2))$), whereas our lower bound (Theorem 2.10) has a polynomial dependence on q (namely, it shows $L = \Omega(1/(q^5\varepsilon^2))$). The bound in Blinovsky's recent paper [12] implicitly yields the (near-optimal) lower bound $L = \Omega(1/(q\varepsilon^2))$.
- When $\varepsilon = o(1)$, say $\varepsilon = 1/n^\gamma$ is polynomially small in the block length n , our results do not rule out the existence of a $(1 - \varepsilon, L)$ -list-decodable code with $L = o(1/\varepsilon^2)$ and which has polynomially small rate (i.e., has $\exp(n^\delta)$ codewords for some constant $\delta > 0$). This is setting of parameters is relevant to some applications of list decoding, such as constructing randomness extractors [8], [9]. Is it possible that a list size of $o(1/\varepsilon^2)$ can be achieved in this setting, or can one extend our lower bound to rule out this possibility?
- It should be possible to use our main result, together with an appropriate "filtering" argument (that focuses, for example, on a subcode consisting of all codewords of a particular Hamming weight) to obtain upper bounds on rate of list-decodable q -ary codes. In particular, can one confirm that for each fixed L , the maximum rate achievable for list decoding up to radius p with list size L is strictly less than the capacity $1 - H_q(p)$? Such a result was shown by Blinovsky for binary codes in [2] and more recently for nonbinary codes in [11], [12]. It is an interesting question whether some of the ideas in this paper can be used to improve the rate upper bounds of Blinovsky [2] for the binary case.
- Can one prove a lower bound on list size as a function of distance from "capacity"? In particular, does one need list size $\Omega(1/\gamma)$ to achieve a rate that is within γ of capacity? Can one at least prove such a lower bound when restricting to linear codes? Recently, Rudra [13] showed that such a lower bound holds with high probability for *random* codes as well as random linear codes.

ACKNOWLEDGMENTS

We thank Michael Mitzenmacher, Madhu Sudan, and Amnon Ta-Shma for helpful conversations, and the anonymous referees for useful comments.

REFERENCES

[1] V. Guruswami and S. Vadhan, "A lower bound on list size for list decoding," in *Proceedings of the 8th International Workshop on Randomization and Computation (RANDOM '05)*, ser. Lecture Notes in Computer Science, no. 3624. Berkeley, CA: Springer, August 2005, pp. 318–329.

[2] 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.

[3] P. Elias, "List decoding for noisy channels," *Technical Report 335, Research Laboratory of Electronics, MIT*, 1957.

[4] J. M. Wozenraft, "List Decoding," *Quarterly Progress Report, Research Laboratory of Electronics, MIT*, vol. 48, pp. 90–95, 1958.

[5] P. Elias, "Error-correcting codes for list decoding," *IEEE Transactions on Information Theory*, vol. 37, pp. 5–12, 1991.

[6] V. Guruswami, J. Hastad, M. Sudan, and D. Zuckerman, "Combinatorial bounds for list decoding," *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1021–1035, 2002.

[7] J. Radhakrishnan and A. Ta-Shma, "Bounds for dispersers, extractors, and depth-two superconcentrators," *SIAM Journal on Discrete Mathematics*, vol. 13, no. 1, pp. 2–24 (electronic), 2000.

[8] L. Trevisan, "Extractors and pseudorandom generators," *Journal of the ACM*, vol. 48, no. 4, pp. 860–879, July 2001.

[9] A. Ta-Shma and D. Zuckerman, "Extractor codes," *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3015–3025, 2004.

[10] V. Guruswami, "List decoding from erasures: Bounds and code constructions," *IEEE Transactions on Information Theory*, vol. 49, no. 11, pp. 2826–2833, 2003.

[11] V. M. Blinovsky, "Code bounds for multiple packings over a nonbinary finite alphabet," *Problems of Information Transmission*, vol. 41, no. 1, pp. 23–32, 2005.

[12] —, "On the convexity of one coding-theory function," *Problems of Information Transmission*, vol. 44, no. 1, pp. 34–39, 2008.

[13] A. Rudra, "Limits to list decoding random codes," in *Proceedings of the 15th International Computing and Combinatorics Conference*, July 2009, to appear. Also available as ECCC Tech Report TR09-013.

[14] D. Dubhashi and S. Sen, "Concentration of measure for randomized algorithms: Techniques and analysis," in *Handbook of Randomized Computing, Vol. I*, S. R. et al., Ed. Kluwer, 2001, ch. 3, pp. 35–100.

[15] S. P. Vadhan, *Pseudorandomness*, ser. Foundations and Trends in Theoretical Computer Science. now publishers, 2010, to appear.

APPENDIX

COUNTING AND PROBABILITY ESTIMATES

Lemma A.1: For all integers n, k such that $0 < k < n$, we have

$$\binom{n}{k} = \Theta\left(\sqrt{\frac{n}{k \cdot (n-k)}} \cdot 2^{H(k/n) \cdot n}\right).$$

Proof: Use Stirling's approximation for the factorials. ■

Lemma A.2: For all positive integers $t, s \leq t/2$,

$$\sum_{i=\lfloor \sqrt{s}/2 \rfloor}^{\lfloor \sqrt{s} \rfloor} \frac{\binom{t}{s+i} \binom{t}{s-i}}{\binom{2t}{2s}} = \Omega(1)$$

Proof: Without loss of generality, we may assume that $s \leq t/2$ (otherwise replace s with $t - s$). Let $A_i = \frac{\binom{t}{s+i} \binom{t}{s-i}}{\binom{2t}{2s}}$. Using Lemma A.1, we see that

$$\begin{aligned} A_0 &= \Theta\left(\sqrt{\frac{t}{s \cdot (t-s)}}\right) \\ &= \Omega\left(\frac{1}{\sqrt{s}}\right) \end{aligned}$$

For $0 < i \leq \lfloor \sqrt{s} \rfloor$, we have

$$\begin{aligned} A_i &= \frac{(t-s-i+1) \cdot (s-i+1)}{(s+i) \cdot (t-s+i)} \cdot A_{i-1} \\ &= \left(1 - \frac{2i-1}{t-s+i}\right) \cdot \left(1 - \frac{2i-1}{s+i}\right) \cdot A_{i-1} \\ &\geq \left(1 - \frac{2}{\sqrt{s}}\right)^2 \cdot A_{i-1} \\ &\geq \left(1 - \frac{2}{\sqrt{s}}\right)^{2i} A_0 \\ &= \Omega(A_0) \end{aligned}$$

Therefore,

$$\sum_{i=\lfloor \sqrt{s}/2 \rfloor}^{\lfloor \sqrt{s} \rfloor} A_i = (\lfloor \sqrt{s} \rfloor - \lfloor \sqrt{s}/2 \rfloor) \cdot \Omega(A_0) = \Omega(1).$$

Lemma A.3: Let t, q be integers such that $q \geq 2$ and $t \geq 37q$. Let $a_1, \dots, a_{2t} \in [q]$, and let T be chosen uniformly at random from all subsets of $[2t]$ of size t . Then

$$\mathbb{E}_T[\#\text{plur}(a_j : j \in T)] \geq \frac{t}{q} + \Omega\left(\sqrt{\frac{t}{q}}\right).$$

Proof: Notice that $\#\text{plur}(a_j : j \in T)$ is always at least t/q . Thus it suffices to show that with constant probability over the choice of T , there exists an $\alpha \in [q]$ such that $\#\{i \in T : a_i = \alpha\} \geq t/q + \Omega(\sqrt{t/q})$. In fact, we restrict our attention to a single value of α , namely the most frequent symbol among a_1, \dots, a_{2t} . Then setting $s = \lfloor t/q \rfloor$, α occurs at least $2s$ times among a_1, \dots, a_{2t} , so let $S \subseteq [2t]$ be any set of $2s$ indices j such that $a_j = \alpha$. Then,

$$\begin{aligned} \Pr_T[\#\{j \in T : a_j = \alpha\} = s+i] \\ \geq \Pr_T[|T \cap S| = s+i] = \frac{\binom{t}{s+i} \binom{t}{s-i}}{\binom{2t}{2s}}. \end{aligned}$$

(To see the last equation, note that $|T \cap S|$ has the same distribution whether T is a random and S is fixed, or T is fixed and S is random.) Thus by Lemma A.2, with probability $\Omega(1)$ over T , we have:

$$\begin{aligned} \#\{j \in T \mid a_j = \alpha\} &\geq s + \lfloor \sqrt{s}/2 \rfloor \\ &\geq \frac{t}{q} + \frac{\sqrt{t/q}}{2} - 3 \\ &\geq \frac{t}{q} + \Omega\left(\sqrt{\frac{t}{q}}\right), \end{aligned}$$

where in the last inequality we use the fact that $t \geq 37q$. \blacksquare

Lemma A.4: For integers $0 < i < a$, $0 < j < b$,

$$\frac{\binom{a}{i} \binom{b}{j}}{\binom{a+b}{i+j}} \leq O\left(\sqrt{\frac{a \cdot b \cdot (i+j) \cdot (a+b-i-j)}{i \cdot (a-i) \cdot j \cdot (b-j) \cdot (a+b)}}\right).$$

Proof: Applying Lemma A.1 to each of the binomial coefficients yields the bound above times 2^t , where

$$t = H(i/a) \cdot a + H(j/b) \cdot b - H((i+j)/(a+b)) \cdot (a+b) \leq 0,$$

where the last inequality follows by concavity of the entropy function. \blacksquare

Lemma A.5: Let b_1, b_2, \dots, b_k be a sequence of elements from the universe $[q]$. Recall that a prefix of such a sequence has a plurality tie if there are at least two elements of $[q]$ that occur the same number of times in the prefix, and no other element occurs a strictly greater number of times in the prefix. Let Y be the random variable counting the number of prefixes with a plurality tie in a random permutation of the b_i 's. Then $\mathbb{E}[Y] = O(q^{3/2}\sqrt{k})$.

Proof: Assume $k \geq q$, or else $Y \leq k < q$ and the claimed bound holds trivially. For $\alpha \in [q]$, let Y_α be a random variable (over the choice of the permutation π of the sequence) counting the number of $i \in [k]$ such that the prefix $(b_{\pi(1)}, \dots, b_{\pi(i)})$ has a plurality tie, α achieves the plurality, and $b_{\pi(i)} \neq \alpha$. Then $Y \leq \sum_\alpha Y_\alpha$. (For every prefix with a plurality tie, at least one of the two symbols achieving the plurality must be different from the last symbol in the prefix.) Thus, it suffices to show that $\mathbb{E}[Y_\alpha] = O(\sqrt{qk})$ for every α .

Fix α . Let ℓ be the number of occurrences of α in the sequence b_1, \dots, b_k . We can obtain a random permutation of b_1, \dots, b_k by randomly ordering the $m = k - \ell$ elements of the sequence other than α , and then randomly merging the ℓ occurrences of α into this sequence (uniformly out of all $\binom{\ell+m}{\ell}$ ways). In fact we will bound the expectation of Y_α for every fixed ordering c_1, \dots, c_m of the elements other than α , and thus the only randomness is over the merging.

For each $r = 1, \dots, m$, let $u_r = \#\text{plur}(c_1, \dots, c_r)$. Notice that $r \geq u_r \geq r/(q-1)$. Let X_r be the indicator random variable for whether upon merging, α occurs exactly u_r times before c_r (equivalently, occurs $v_r = \ell - u_r$ times after c_r). Then $Y_\alpha = \sum_{r=1}^m X_r$.

Fix $r \in [m]$, let $s = m - r$, $u = u_r$, $v = v_r = \ell - u$. Our aim is to bound $\Pr[X_r = 1]$. Notice that the merging can be viewed as uniformly choosing a set S of ℓ out of $m + \ell$ locations to place the α 's (and putting the c_i 's in the remaining m locations). Observe that $X_r = 1$ only if S contains exactly u of the first $u + r$ locations (and thus exactly v of the last $v + s$ locations); let E_r denote this event. ($X_r = 1$ also implies that S does not contain location $u + r$, but we will not make use of that.)

We bound $\Pr[E_r]$ for $r \in \{q, q+1, \dots, m-1\}$ by considering two cases. (For $r < q$ and $r = m$, we will use the trivial bound $\Pr[E_r] \leq 1$.) First, suppose that $u/r > 2v/s$. Intuitively, this means that, for E_r to occur, S must be disproportionately partitioned in the merging. Specifically, the expected number of elements of S among the first $u + r$ locations is

$$\frac{u+r}{u+r+v+s} \cdot (u+v) < \frac{u}{2}.$$

By Chernoff bounds, the probability that the first $u + r$ locations contain more than u elements of S is at most $2^{-\Omega(u)} \leq 2^{-\Omega(r/q)} \leq O(\sqrt{q/r})$ for $r \geq q$. (The indicators for whether each location contains an element of S satisfy ‘‘negative dependence’’, and thus Chernoff bounds apply [14].)

The second case is that

$$\frac{v}{s} \geq \frac{u}{2r} \geq \frac{1}{2(q-1)}. \quad (6)$$

Then, by Lemma A.4,

$$\begin{aligned} \Pr[E_r] &= \binom{u+r}{r} \binom{v+s}{s} / \binom{u+v+r+s}{r+s} \\ &= O\left(\sqrt{\frac{(u+r) \cdot (v+s) \cdot (u+v) \cdot (r+s)}{r \cdot u \cdot s \cdot v \cdot (u+r+v+s)}}\right). \end{aligned} \quad (7)$$

Now for positive integers x, y , we have

$$\frac{x \cdot y}{x + y} = \frac{\min\{x, y\} \cdot \max\{x, y\}}{x + y} = \Theta(\min\{x, y\}). \quad (8)$$

From (7) and (8), we conclude

$$\begin{aligned} \Pr[E_r] &= O\left(\sqrt{\frac{\min\{u+r, v+s\}}{\min\{u, v\} \cdot \min\{r, s\}}}\right) \\ &= O\left(\sqrt{\frac{q}{\min\{r, s\}}}\right) \quad (\text{using (6)}). \end{aligned}$$

Thus, in both cases we have $\Pr[E_r] = O(\sqrt{q/\min\{r, s\}})$ for $r \in \{q, q+1, \dots, m\}$. Therefore

$$\begin{aligned} \mathbb{E}[Y_\alpha] &\leq q + \sum_{r=q}^{m-1} \Pr[E_r] \\ &= q + \sum_{r=1}^{m-1} O\left(\sqrt{\frac{q}{\min\{r, m-r\}}}\right) \\ &= O(q + \sqrt{qm}) \\ &= O(\sqrt{qk}), \end{aligned}$$

since $m \leq k$ and $q \leq k$. This gives the desired bound on $\mathbb{E}[Y_\alpha]$ for each $\alpha \in [q]$. ■

RANDOMNESS EXTRACTORS

Here we briefly review the connection between list-decodable error-correcting codes and “randomness extractors” from [8], [9]. We present the definition of extractors using nonstandard choices of variables for consistency with standard notation for codes.

Definition A.6: A function $\text{Ext} : [K] \times [n] \rightarrow [q]$ is an (ℓ, ε) extractor⁴ if for every random variable X , taking values in $[K]$ and such that $\Pr[X = x] \leq 1/2^\ell$ for every x , the random variable $(U_{[n]}, \text{Ext}(X, U_{[n]}))$ has statistical distance less than ε from $(U_{[n]}, U_{[q]})$, where U_W denotes the uniform distribution on set W .

Thus an extractor uses a $(\log n)$ -bit random “seed” to extract $\log q$ almost-uniform bits from any random source of *min-entropy* at least ℓ . (The min-entropy of X is defined to be $\min_x \log(1/\Pr[X = x])$.)

If we have a code $C \subseteq [q]^n$ of size K with an encoding function $\text{Enc} : [K] \rightarrow [q]^n$ (s.t. $\text{Image}(\text{Enc}) = C$), then we can obtain an extractor $\text{Ext} : [K] \times [n] \rightarrow [q]$ as follows:

$$\text{Ext}(x, y) = \text{Enc}(x)_y. \quad (9)$$

⁴In the extractor literature, such a function would normally be called an (ℓ, ε) *strong* extractor. Also, the usual definition requires that the statistical distance be *at most* ε instead of strictly less than ε , but the latter choice is more convenient here.

Conversely, every extractor yields an encoding function by setting:

$$\text{Enc}(x) = \text{Ext}(x, 1)\text{Ext}(x, 2) \cdots \text{Ext}(x, q). \quad (10)$$

The extraction properties of Ext and the list-decodability properties of Enc are given by the following.

Proposition A.7 ([8], [9], [15]):⁵ Let $\text{Ext} : [K] \times [n] \rightarrow [q]$ and $\text{Enc} : [K] \rightarrow [q]^n$ correspond to each other via Equations (9) and (10). Then for every $L \in \mathbb{N}$ and $\varepsilon \in [0, 1]$:

- 1) If Ext is an (L, ε) extractor, then Enc is $(1 - (q/(q - 1)) \cdot \varepsilon, L - 1)$ list-decodable.
- 2) If Enc is $(1 - \varepsilon, L)$ list-decodable, then Ext is an $(L/\varepsilon, (q - 1) \cdot \varepsilon)$ extractor.

Radhakrishnan and Ta-Shma [7] proved that for most settings of parameters, an (L, ε) extractor must lose at least $2 \log(1/\varepsilon) - O(1)$ bits of entropy, i.e. $\log q \leq \log L - 2 \log(1/\varepsilon) + O(1)$. Equivalently, $L = \Omega(q/\varepsilon^2)$. By Item 2, this implies that a $(1 - \varepsilon', L')$ list-decodable code must have $L' = \Omega(1/(q\varepsilon'))$. Notice that this bound is only linear in $1/\varepsilon'$.

In the other direction, by Item 1, our result showing that a $(1 - \varepsilon, L)$ list-decodable code must have list size $L \geq c_q/\varepsilon^2$ implies the same for an (L, ε) extractor, which matches the Radhakrishnan–Ta-Shma bound in its dependence on ε but is worse in its dependence on q (which is quite significant in the context of extractors).

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 lie in the theory of error-correcting codes, approximation algorithms and hardness of approximation results for optimization problems, pseudorandomness, 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).

Salil Vadhan is the Vicky Joseph Professor of Computer Science and Applied Mathematics in the Harvard University School of Engineering and Applied Sciences. Dr. Vadhan received a Bachelor’s degree from Harvard University in 1995, a Certificate of Advanced Study from Cambridge University in 1996, and a Ph.D. from the Massachusetts Institute of Technology in 1999. He was an NSF Postdoctoral Fellow at the Massachusetts Institute of Technology from 1999–2000 and a visitor at the Institute for Advanced Study in Princeton from 2000–2001, after which he joined the Harvard faculty. In 2003–04, he was a Fellow at the Radcliffe Institute for Advanced Study at Harvard University and in 2007–08, he was a Miller Visiting Professor at the University of California, Berkeley.

⁵In [15], list-decodable codes are defined with respect to open rather than closed Hamming balls and extractors require statistical distance at most ε rather than strictly less than ε , but the same proof yields the statement given here.

Dr. Vadhan's research areas are computational complexity, cryptography, and randomness in computation, with particular interests in pseudorandomness, zero-knowledge proofs, and data privacy. Dr. Vadhan is a recipient of a Gödel Prize (2009), a Guggenheim Fellowship (2007), an ONR Young Investigator Award (2004), a Phi Beta Kappa Award for Excellence in Teaching (2004), a Sloan Fellowship (2002), an NSF Career Award (2002), and ACM's Doctoral Dissertation Award (2000).