Approximating Probability Distributions Using Small Sample Spaces

Yossi Azar *
Computer Science Department
Tel Aviv University
Tel Aviv 69978, Israel

Rajeev Motwani †
Computer Science Department
Stanford University
Stanford, CA 94305-2140

Joseph (Seffi) Naor ‡
Computer Science Department
Technion
Haifa 32000, Israel

Abstract

We formulate the notion of a "good approximation" to a probability distribution over a finite abelian group. The approximate distribution is characterized by a parameter ϵ, the quality of the approximation, which is a bound on the difference between corresponding Fourier coefficients of the two distributions. It is also required that the sample space of the approximate distribution be of size polynomial in the representation length of the group elements as well as $1/\epsilon$. Such approximations are useful in reducing or eliminating the use of randomness in randomized algorithms.

We demonstrate the existence of such good approximations to arbitrary distributions. In the case of n random variables distributed uniformly and independently over the range $\{0, \ldots, d-1\}$, we provide an efficient construction of a good approximation. The constructed approximation has the property that any linear combination of the random variables (modulo d) has essentially the same behavior under the approximate distribution as it does under the uniform distribution over $\{0, \ldots, d-1\}$. Our construction is based on Weil's character sum estimates. We apply this result to the construction of a non-binary linear code where the alphabet symbols appear almost uniformly in each non-zero code-word.

*Most of this work was done while the author was at the Computer Science Department, Stanford University and supported by by a Weizmann fellowship and Contract ONR N00014-88-K-0166.

†Supported by an Alfred P. Sloan Research Fellowship, an IBM Faculty Development Award, grants from Mitsubishi Electric Laboratories and OTL, NSF Grant CCR-9010517, and NSF Young Investigator Award CCR-9357849, with matching funds from IBM, Schlumberger Foundation, Shell Foundation, and Xerox Corporation.

‡Part of this work was done while the author was visiting the Computer Science Department, Stanford University, and supported by Contract ONR N00014-88-K-0166, and by Grant No. 92-00225 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel.

1 Introduction

Recently a family of techniques has emerged to reduce or eliminate the use of random bits by randomized algorithms [ABI, BR, CG, KW, Lu1, Lu2, MNN, NN]. Typically, these techniques involve substituting independent random variables by a collection of dependent random variables which can be generated using fewer truly independent random bits. Motivated by this work, we formulate the notion of a *good approximation* to a *joint probability distribution* on a collection of random variables.

We consider probability distributions over a finite abelian group $\mathbb{G}$ and, in particular, over $\mathbb{Z}_d^n$ for any positive integers d and n. We measure the distance between two distributions over $\mathbb{G}$ by the distance, in the maximum norm, of their Fourier transforms over $\mathbb{G}$. Given an arbitrary distribution $\mathcal{D}$ over $\mathbb{G}$, a good approximation $\overline{\mathcal{D}}$ is a distribution with a small distance to $\mathcal{D}$, which is concentrated on a small subset of the sample space. Sampling from the approximate distribution requires significantly fewer random bits than sampling from the original distribution.

Before describing our work in detail, we briefly review some related work. Alon, Babai, and Itai [ABI] and Luby [Lu1] observed that certain algorithms perform as well using pairwise independent random bits, as on mutually independent bits. It turns out that n uniform k-wise independent bits can be generated using sample spaces of size $O(n^k)$; a lower bound of $\Omega(n^{\lceil k/2 \rceil})$ on the minimum size of such a sample space can be shown for $k << n$ [ABI, CGHFRS]. Thus, these algorithms could be derandomized for constant k by an exhaustive search of the (polynomial-size) sample space. Unfortunately, this degree of independence is very restrictive and limits the applicability of the approach. Berger and Rompel [BR] and Motwani, Naor, and Naor [MNN] showed that several interesting algorithms perform reasonably well with only $(\log n)$-wise independence. The resulting sample space, while of super-polynomial size, could be efficiently searched via the *method of conditional probabilities* [AS] in time logarithmic in the size of the sample space. This led to the derandomization of a large class of parallel algorithms.

An alternate approach was proposed by Naor and Naor [NN] based on the notion of the *bias* of a distribution due to Vazirani [Va].

Definition 1.1 *Let $X_1, \ldots, X_n$ be $\{0,1\}$ random variables. The* bias *of a subset S of the random variables is defined to be $|\Pr[\sum_{i \in S} X_i = 0] - \Pr[\sum_{i \in S} X_i = 1]|$, where the sum is taken modulo 2.*

For mutually independent and uniform random variables, the bias of each non-empty subset is zero. It is not hard to show that the converse holds as well. In an *ϵ-biased probability distribution*, each subset of the random variables has bias at most ϵ. Naor and Naor [NN] showed how to construct such a distribution, for any $\epsilon > 0$, such that the size

of the sample space is polynomial in n and $1/\epsilon$. The ϵ-biased distribution can be viewed as an *almost* $(\log n)$*-wise independent distribution*. A result due to Peralta [Pe] implies a different construction of ϵ-biased probability distribution using the properties of *quadratic residues*; this and two additional constructions of two-valued ϵ-biased random variables are reported by Alon, Goldreich, Håstad, and Peralta [AGHP].

We formulate and study the notion of a "good approximation" to a joint probability distribution of (possibly multi-valued) random variables. Let $\mathcal{D}$ be any joint distribution of n random variables over the range $\{0, \ldots, d-1\}$. Informally, a good approximation $\overline{\mathcal{D}}$ to $\mathcal{D}$ satisfies the following properties: there is a uniform bound ϵ/d^{n-1} on the absolute difference between corresponding Fourier coefficients over the group $\mathbb{Z}_d^n$ of the two distributions; and, the sample space required for $\overline{\mathcal{D}}$ is of size polynomial in n, d, and $1/\epsilon$. We demonstrate the viability of such approximations by proving that for any distribution $\mathcal{D}$, there exists a good approximation $\overline{\mathcal{D}}$. In fact, this notion and the existence result extend to any probability distribution over a finite abelian group. The quality of the approximation can be further characterized by showing that the *variation distance* between the two distributions $\mathcal{D}$ and $\overline{\mathcal{D}}$ is bounded by the sum of the differences between their Fourier coefficients.

We also consider the issue of an *efficient construction* of such an approximate distribution; specifically, for the uniform distribution over $\mathbb{Z}_d^n$. An efficient construction must determine $\overline{\mathcal{D}}$ in time polynomial in the description length of $\mathcal{D}$, and also in $1/\epsilon$; clearly, this bound must apply to the size of the sample space of the approximate distribution $\overline{\mathcal{D}}$. (Note that the description of a distribution $\mathcal{D}$ over $\mathbb{Z}_d^n$ may be of length as much as d^n.) We provide an efficient construction of a good approximation $\overline{\mathcal{U}}$ to the uniform distribution $\mathcal{U}$ on $\mathbb{Z}_d^n$, i.e., for the joint distribution of uniform and independent d-valued random variables $X_1, \ldots, X_n$. Since the construction must guarantee that the Fourier coefficients of $\overline{\mathcal{U}}$ are very close to those of $\mathcal{U}$, it is essentially a generalized ϵ-biased distribution. This has the following natural interpretation in terms of linear combinations: for *any* vector $A = (a_1, \ldots, a_n)$, $\sum a_i X_i \pmod d$ has "almost" the same distribution in the case where the random variables $X_1, \ldots, X_n$ are chosen from $\mathcal{U}$, as in the case where they are chosen from $\overline{\mathcal{U}}$. The analysis of this construction is based on Weil's character sum estimates, and it generalizes the work of Peralta [Pe] to d-valued random variables. Our results hold for non-prime values of d as well[1].

This construction has found application in the work of Håstad, Phillips, and Safra [HPS]. They consider the approximability of the following algebraic optimization problem: given a collection of degree 2 polynomials over $\mathbb{GF}[d]$, find a common root to the largest possible sub-collection of these polynomials. Our construction is used to show that finding an

[1]Following our work, Even [Ev] generalized one of the constructions of Alon, Goldreich, Håstad, and Peralta [AGHP] to the d-valued case when d is a prime.

approximate solution (to within a ratio of $d - \epsilon$) is as hard as finding an exact solution, and hence is NP-hard; this applies to polynomials over rationals and reals as well. The constructions of two-valued ϵ-biased random variables due to Naor and Naor [NN] and Alon, Goldreich, Håstad, and Peralta [AGHP] are insufficient for this purpose, and our construction of d-valued ϵ-biased random variables needs to be used.

We also show that the *variation distance* between two distributions can be bounded in terms of the differences in their Fourier coefficients. This allows us to demonstrate that our construction gives random variables which are "almost" $(\log_d n)$-wise independent. Our construction is optimal in this respect. We also explore some connections with the construction of linear codes. Our results provide a construction of a linear code over an arbitrary alphabet which has the property that for each non-zero codeword, the distribution of the alphabet symbols is almost uniform, and that the length of the codeword is polynomial (quadratic) in the dimension of the code. Previously, such codes were known only over $\mathbb{GF}[2]$.

The remaining sections are organized as follows: Section 2 provides some mathematical preliminaries; the existence of a good approximation to an arbitrary distribution and bounds on their variation distance are shown in Section 3; Section 4 gives a generalized construction of an ϵ-biased distribution; the properties of this construction are studied in Section 5; finally, in Section 6 our construction is applied to linear codes.

2 Preliminaries

2.1 Characters of Finite Abelian Groups. Our discussion here follows the exposition of Babai [Ba] and Ledermann [Le]. Let $\mathbb{C}$ denote the multiplicative group of complex numbers with unit modulus. A character of a finite abelian group $\mathbb{G}$ is a homomorphism $\chi : \mathbb{G} \to \mathbb{C}$. That is $|\chi(x)| = 1$, for all $x \in \mathbb{G}$, and $\chi(xy) = \chi(x)\chi(y)$, for all $x, y \in \mathbb{G}$. The number of different characters for an abelian group is $|\mathbb{G}|$. The *principal character* χ_0 of the group $\mathbb{G}$ is the trivial mapping $\chi_0(x) = 1$, for all $x \in \mathbb{G}$. The set of characters can be endowed with a group structure in which multiplication serves as composition, i.e., if χ and χ' are characters, then the character $\chi\chi'$ is defined by

$$\chi\chi'(x) = \chi(x)\chi'(x).$$

This group is called the *dual group* of $\mathbb{G}$ and denoted by $\widehat{\mathbb{G}}$. It is known that $\widehat{\mathbb{G}} \cong \mathbb{G}$ [Ba]. A character χ is said to be of *order* d, if d is the smallest positive integer such that $\chi^d = \chi_0$.

Suppose first that $\mathbb{G} = \mathbb{Z}_n$ is a cyclic group of order n, and z is some generator for it, i.e., each element of $\mathbb{Z}_n$ can be written as z^s for some integer s. The characters of $\mathbb{Z}_n$ are defined as follows.

Definition 2.1 *For $0 \leq r \leq n-1$, the rth character of $\mathbb{Z}_n$ is denoted by χ_r and defined as*

follows,

$$\chi_r(z^s) = e\left(\frac{rs}{n}\right),$$

where $s = 0, \ldots, n-1$, *and the function* $e(x)$ *denotes* $e^{2\pi ix}$ *for* $i = \sqrt{-1}$.

It follows that χ_r is of order $n/\gcd(r,n)$. We note that in the case of a multiplicative group modulo some prime p, viz. $\mathbb{Z}_p^*$, the notion of a multiplicative character can be extended to the finite field $\mathbb{GF}[p]$ by setting $\chi(0) = 0$.

For the general case of a finite abelian group $\mathbb{G}$, we invoke the Basis Theorem [Le] which states that $\mathbb{G}$ is the direct product of cyclic groups, say, $\mathbb{Z}_{n_1}, \ldots, \mathbb{Z}_{n_k}$. Let $\mathbb{V}_{\mathbb{G}}$ denote the set of all k-tuples $(t_1, \ldots, t_k)$, for $0 \le t_i < n_i$, for $1 \le i \le k$; note that $|\mathbb{V}_{\mathbb{G}}| = |\mathbb{G}|$. Each element $x \in \mathbb{G}$ can be uniquely expressed as

$$x = z_1^{a_1} z_2^{a_2} \cdots z_k^{a_k},$$

where z_i is a generator of $\mathbb{Z}_{n_i}$ and $0 \le a_i < n_i$. We can thus represent x by the k-tuple $A = (a_1, \ldots, a_k) \in \mathbb{V}_{\mathbb{G}}$. There is a character corresponding to each k-tuple $R = (r_1, \ldots, r_k) \in \mathbb{V}_{\mathbb{G}}$, defined as follows:

$$\chi_R(x) = \mathrm{e}\left(\sum_{i=1}^{k} \frac{a_i r_i}{n_i}\right).$$

We are particularly interested in the set of characters of the group $\mathbb{Z}_d^n$. In this case, there is a character defined for each vector $R = (r_1, \ldots, r_n)$ in $\mathbb{Z}_d^n$, as follows:

$$\chi_R(a_1, \ldots, a_n) = \mathrm{e}\left(\frac{1}{d}\sum_{i=1}^{n} a_i r_i\right).$$

2.2 Absolutely Irreducible Polynomials. Let $f(x,y)$ be a polynomial of total degree m with coefficients in the field $\mathbb{Z}_p$. Then $f(x,y)$ is *absolutely irreducible* if it is irreducible over the algebraic closure of $\mathbb{Z}_p$. Absolutely irreducible polynomials play an important role in our constructions and a certain class of these is characterized by the following lemma (see Schmidt [Sc, page 11]).

Lemma 2.2 *Suppose the polynomial* $y^m - f(x)$ *has coefficients in the field* $\mathbb{Z}_p$. *Then* $y^m - f(x)$ *is absolutely irreducible if and only if* $\gcd(m, a_1, \ldots, a_n) = 1$, *where* $f(x) = a(x - i_1)^{a_1} \cdots (x - i_n)^{a_n}$ *is the factorization of* f *over some extension field, with* $i_j \neq i_k$ *for* $j \neq k$.

2.3 Discrete Fourier Transform. We will give only a brief overview of the basic concepts in Fourier analysis; see Dym and McKean [DM] or Körner [Kö] for more details. As before, let $\mathbb{G}$ be a finite Abelian group. The set $\mathbb{V} = \{f : \mathbb{G} \to \mathbb{C}\}$ of complex functions on the group $\mathbb{G}$ forms a $|\mathbb{G}|-$dimensional vector space over $\mathbb{C}$. The inner product of two (complex-valued) functions f and g is defined as follows:

$$< f, g > = \frac{1}{|\mathbb{G}|} \sum_{x \in |\mathbb{G}|} f(x) g(x)^*$$

where $*$ denotes the complex conjugate operation. As usual, the L_2 *norm* of a function f is defined as $\|f\|_2 = \sqrt{<f, f>}$.

A basis for the vector space $\mathbb{V}$ is given by the characters of $\mathbb{G}$. This basis has the *orthonormality* property: if $R_1 \neq R_2$, then $< \chi_{R_1}, \chi_{R_2} >= 0$; and, for each character R, $<\chi_R, \chi_R >= 1$.

Any function $f \in \mathbb{C}^{\mathbb{G}}$ can be written as a linear combination of characters:

$$f = \sum_{\chi \in \widehat{\mathbb{G}}} \widehat{f}_\chi \chi$$

The coefficients $\widehat{f}_\chi$ are called the Fourier coefficients, and are given by $\widehat{f}_\chi =< f, \chi >$. We note that a probability distribution defined on $\mathbb{G}$ maps each element $x \in \mathbb{G}$ into a real value which is the probability of picking x as the random sample. This allows us to formulate the Fourier transform of a probability distribution.

3 Approximating arbitrary distributions

In this section we suggest an approach to approximating arbitrary distributions. Previous work concentrated on approximating the *uniform* distribution over two-valued random variables. Let $\mathcal{D}$ be a probability distribution over a finite abelian group $\mathbb{G}$. We will show that there exists a *small* probability space which approximates $\mathcal{D}$. The following is a somewhat strengthened version of our original theorem, based on an observation due to Mario Szegedy.

Denote by $||\cdot||_\infty$ the L_∞ norm of a vector, i.e., the maximum among the absolute values of its components. Let Γ be a sample space of size l, and $\mathcal{D}$ be a probability distribution over Γ. In what follows, we will often represent such a distribution $\mathcal{D}$ by a (probability) vector D in l dimensions.

Theorem 3.1 *Let M be an $l \times l$ matrix of complex numbers with entries of absolute value at most 1. For any probability distribution $\mathcal{D}$ represented by the vector D of length l, and any $\epsilon > 0$, there exists a probability distribution $\mathcal{F}$ represented by a vector F with at most $O(\epsilon^{-2} \log l)$ non-zero entries, such that*

$$||D \cdot M - F \cdot M||_\infty \leq \epsilon.$$

Proof: We use the probabilistic method to demonstrate the existence of a sample space $\Omega \subset \Gamma$ such that a *uniformly* chosen sample point from Ω has a distribution approximating $\mathcal{D}$; thus, $\mathcal{F}$ is the uniform distribution over Ω. We choose $\Omega = \{\omega_1, \ldots, \omega_k\}$ as follows: pick each ω_i independently from Γ according to the distribution $\mathcal{D}$. Since the sample points $\omega_i \in \Gamma$ need not be distinct, in general, Ω will be a multi-set; if necessary, the repetitions can be eliminated by suitably modifying the probability measure.

We index the rows of M by $s \in S$. We claim that, provided k is large enough, for every $s \in S$ the probability that $|D \cdot M_s - F \cdot M_s| > \epsilon$ is less than $1/l$. Since the number of rows is l, this implies that

$$\Pr\left[\exists\, s,\, |D \cdot M_s - F \cdot M_s| > \epsilon\right] < 1.$$

Note that the probability in the above expression is with respect to the random choice of Ω. Thus, it follows that there exists a choice of the elements $\omega_i \in \Gamma$, for $1 \leq i \leq k$, which will yield the probability space $(\Omega, \mathcal{F})$ as required.

It remains to prove the claim. Let us now concentrate only on the row indexed by a specific s. For $1 \leq i \leq k$, let w_i be the jth coordinate of the sth row, where j is the index of the element that was chosen as ω_i. It follows that

$$F \cdot M_s = \sum_{i=1}^{k} \frac{1}{k} w_i,$$

that is $F \cdot M_s$ is proportional to the sum of k independent random variables. In what follows, E and Pr denote expectation and probability with respect to the uniform measure on the (multi-set) sample space Ω. We have that $\mathrm{E}[w_i] = D \cdot M_s$ and

$$\mathrm{E}[F \cdot M_s] = \mathrm{E}\left[\frac{1}{k}\sum_{i=1}^{k} w_i\right] = \mathrm{E}[w_i] = D \cdot M_s.$$

To complete the proof we show that the sum of the w_i's does not differ from its expected value by more than ϵk. Let S be sum of n independent variables, each of which has an absolute value of at most 1. By the Chernoff bound [AS], for any $h \geq 0$,

$$\Pr[\,|S - \mathrm{E}[S]| \geq h\,] \leq 2e^{-\Omega(h^2/n)}.$$

This bound implies that

$$\Pr\left[\left|\sum_{i=1}^{k} w_i - kD \cdot M_s\right| > \delta\right] \leq 2e^{-\Omega(\delta^2/k)}.$$

In our case, the bound on the allowed deviation from the expected value is $\delta = \epsilon k$. We need to choose k such that $e^{-\Omega(\delta^2/k)} < 1/2l$. This is clearly true for $k = \Theta(\epsilon^{-2}\log l)$. □

Denote by $\widehat{\mathcal{D}}_R$ the Fourier coefficients corresponding to the character χ_R for a distribution D on a finite abelian group $\mathbb{G}$. The following theorem shows the *existence* of a good approximation $(\Omega, \mathcal{F})$ to the distribution $\mathcal{D}$ such that the sample space Ω is small.

Theorem 3.2 *For any probability distribution $\mathcal{D}$ defined over a finite abelian group $\mathbb{G}$ and any $\epsilon \in [0,1]$, there exists a probability space $(\Omega, \mathcal{F})$, such that:*

1. *for all characters χ_S, $|\widehat{\mathcal{F}}_S - \widehat{\mathcal{D}}_S| \leq \epsilon/|\mathbb{G}|$,*
2. *the size of the probability space Ω is at most $O(\epsilon^{-2}\log|\mathbb{G}|)$.*

Proof: The proof is an immediate consequence of Theorem 3.1. We choose M to be a matrix proportional to the character table of the group $\mathbb{G}$. That is, the rows are indexed by the set S of characters, the columns by the set X of elements in the group, and $M_{sx} = \chi_s(x)$. The theorem follows from the fact that $|\chi_s(x)| \leq 1$ and the fact that the Fourier coefficients differ from the vectors' inner products by precisely a factor of $|\mathbb{G}|$. □

We now discuss the significance of Theorem 3.2. Denote by $||\cdot||_1$ the L_1 norm of a vector, i.e., the sum of the absolute values of its components. Any complex function $f \in \mathbb{V}$ can be looked upon as a complex vector, with components which are the value of this function at each $x \in \mathbb{G}$. Thus, we can define the distance between two functions as $\|f - g\|_1$. In particular, we will make use of the distance between two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$.

Definition 3.3 *Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be two probability distributions over a finite abelian group $\mathbb{G}$. We define the* variation distance *between these two distributions as $\|\mathcal{D}_1 - \mathcal{D}_2\|_1$.*

The next theorem bounds the variation distance between $\mathcal{D}$ and $\mathcal{F}$ in terms of their Fourier coefficients.

Theorem 3.4 *Let the probability distributions $\mathcal{D}$ and $\mathcal{F}$ be defined over a finite abelian group $\mathbb{G}$. Then,*

$$||\mathcal{D} - \mathcal{F}||_1 \leq |\mathbb{G}| \sqrt{\sum_A |\widehat{\mathcal{D}}_A - \widehat{\mathcal{F}}_A|^2} \leq |\mathbb{G}| \sum_A |\widehat{\mathcal{D}}_A - \widehat{\mathcal{F}}_A|.$$

Proof: The right inequality is immediate. Let $\mathcal{H} = \mathcal{D} - \mathcal{F}$. Using the Cauchy-Schwartz Inequality and Parseval's Equality, we conclude that

$$||\mathcal{H}||_1 = \sum_x |\mathcal{H}(x)| \leq \sqrt{|\mathbb{G}|} \sqrt{\sum_x |\mathcal{H}^2(x)|} = |\mathbb{G}| \sqrt{\sum_A |\widehat{\mathcal{H}}_A^2|}.$$

□

Corollary 3.5 *Let $\mathcal{D}$ and $\mathcal{F}$ be probability distribution $\mathcal{D}$ defined over $\mathbb{Z}_d^n$ such that for some $0 \leq \epsilon \leq 1$ and for all $A \in \mathbb{Z}_d^n$, $|\widehat{\mathcal{F}}_A - \widehat{\mathcal{D}}_A| \leq \epsilon/d^n$. Then, for any subset S of the random variables, the variation distance between $\mathcal{D}$ and $\mathcal{F}$ restricted to the subset S is bounded by $\epsilon d^{|S|}$.*

If ϵ is chosen to be polynomially small, then Corollary 3.5 implies that for any distribution $\mathcal{D}$, there exists a distribution $\mathcal{F}$ over a polynomial size sample space such that any subset S of the random variables is distributed in $\mathcal{F}$ "almost" the same as in $\mathcal{D}$, provided that $|S| = O(\log_d n)$.

4 Constructing an ϵ-biased probability distribution

In this section we show how to approximate a uniform d-valued probability distribution. It turns out that a good approximation in this case is the generalized ϵ-biased distribution.

(This distribution has been studied earlier [NN, AGHP] for the case $d = 2$). We first generalize the definition of ϵ-biased distributions to the case of multi-valued random variables. Let $X = (X_1, \ldots, X_n)$ be a random variable over a set $\Omega \subseteq \mathbb{Z}_d^n$. We define the bias of X with respect to any $A \in \mathbb{Z}_d^n$ as follows.

Definition 4.1 *Let $A = (a_1, \ldots, a_n)$ be any vector in $\mathbb{Z}_d^n$ and let $g = \gcd(a_1, \ldots, a_n, d)$. The bias of A is defined to be*

$$bias(A) = \frac{1}{g} \max_{0 \le k < \frac{d}{g}} \left| \Pr \left[\sum_{i=1}^{n} a_i X_i \equiv kg \pmod{d} \right] - \frac{g}{d} \right|$$

We introduce g in this definition because, independent of the distribution of the random variables, the only values that $\sum_{i=1}^{n} a_i X_i \pmod{d}$ can take are multiples of g.

Definition 4.2 *Let $0 \le \epsilon \le 1$ and let $\Omega \subseteq \mathbb{Z}_d^n$. A probability space $(\Omega, \mathcal{P})$ is said to be ϵ-biased if the corresponding random variable $X = (X_1, \ldots, X_n)$ has the following properties.*

1. *For $1 \le i \le n$, X_i is uniformly distributed over $\mathbb{Z}_d$.*
2. *For all vectors $A \in \mathbb{Z}_d^n$, $bias(A) \le \epsilon$.*

We are now ready to present the construction of ϵ-biased random variables such that the sample space Ω has size which is bounded by a small polynomial in n and $1/\epsilon$. We describe it implicitly by specifying an algorithm for choosing a random sample point. In what follows, we assume that the prime p and a character χ_r of $\mathbb{Z}_p^*$ are chosen such that $d = (p-1)/r$ is the order of the character; further, assume that $p - 1 \ge n$. Let z be a generator for $\mathbb{Z}_p^*$.

Random Sampling Algorithm:

1. Choose the random variable $Y \in \mathbb{Z}_p^*$ uniformly at random.
   (a) For $1 \le j \le n+1$, let $Y_j = Y + j$,
   (b) If all Y_j's are non-zero then set $k = n + 1$, otherwise let k be such that $Y_k = 0$.
2. For $1 \le i \le n$:
   (a) Let $Z_i = \begin{cases} Y_i & \text{if } i < k \\ Y_{i+1} & \text{otherwise} \end{cases}$
   (b) Let s_i be such that $Z_i \equiv z^{s_i} \pmod{p}$.
   (c) Choose $X_i = s_i \bmod d$.

In Step 1(b), the definition of k takes care of the case where one of the random variables $Y_j = Y + j$ is zero and, therefore, not in $\mathbb{Z}_p^*$.

Let Θ be a primitive dth root of unity and $\chi_r(x) = \Theta^{rx}$. Notice that for all i, provided all Y_j's are non-zero,

$$\chi_r(Y+i) = \chi_r(Z_i) = \chi_r(z^{s_i}) = \Theta^{s_i} = \Theta^{X_i} \tag{1}$$

We will establish that these random variables have the desired properties via Weil's character sum estimates (see Schmidt [Sc, page 43, Theorem 2C]). Let $\mathbb{F}[x]$ denote the ring of polynomials over the field $\mathbb{F} = \mathbb{Z}_p$.

Theorem 4.3 (Weil's Theorem) *Let χ be a character of order $d > 1$. Suppose the polynomial $f(x) \in \mathbb{F}[x]$ has precisely n distinct zeros in the algebraic closure of $\mathbb{F}$, and $y^d - f(x)$ is absolutely irreducible. Then,*

$$\left| \sum_{x \in \mathbb{F}} \chi(f(x)) \right| \leq (n-1)\sqrt{p}$$

To analyze the properties of our construction, we prove the following.

Theorem 4.4 *Under the conditions of Theorem 4.3, let r_k be defined as the number of solutions $x \in \mathbb{F}$ to the equation $\chi(f(x)) = \Theta^k$. Then,*

$$\left| r_k - \frac{p}{d} \right| \leq n\sqrt{p}.$$

Proof: The definition of r_j implies that

$$\sum_{x \in \mathbb{F}} \chi(f(x)) = \sum_{j=0}^{d-1} r_j \Theta^j. \tag{2}$$

Hence, for $0 \leq \ell \leq d-1$,

$$\sum_{x \in \mathbb{F}} (\chi(f(x)))^\ell = \sum_{j=0}^{d-1} r_j \Theta^{\ell j}.$$

It follows that,

$$\begin{aligned} p + \sum_{\ell=1}^{d-1} \sum_{x \in \mathbb{F}} (\chi(f(x)))^\ell &= \sum_{\ell=0}^{d-1} \sum_{x \in \mathbb{F}} (\chi(f(x)))^\ell \\ &= \sum_{\ell=0}^{d-1} \sum_{j=0}^{d-1} r_j \Theta^{\ell j} \\ &= dr_0 \Theta^0 + \sum_{j=1}^{d-1} r_j \sum_{l=0}^{d-1} \Theta^{j\ell} \\ &= dr_0. \end{aligned}$$

Hence,

$$|dr_0 - p| = \left| \sum_{\ell=1}^{d-1} \sum_{x \in \mathbb{F}} (\chi(f(x)))^\ell \right| \leq \sum_{\ell=1}^{d-1} \left| \sum_{x \in \mathbb{F}} (\chi(f(x)))^\ell \right|$$

The order of the character χ^ℓ is $d' = d/\gcd(d,\ell)$ which is greater than 1 for $0 < \ell < d$. Since $y^d - f(x)$ is absolutely irreducible and d' divides d, it follows that $y^{d'} - f(x)$ is also absolutely irreducible. Now, applying Theorem 4.3 to each term in the latter sum, we obtain

$$|dr_0 - p| \leq \sum_{\ell=1}^{d-1}(n-1)\sqrt{p} = (d-1)(n-1)\sqrt{p},$$

which implies that

$$\left| r_0 - \frac{p}{d} \right| \leq n\sqrt{p}.$$

To conclude the proof, observe that if both sides of (2) are multiplied by Θ^{-k}, then the same result is obtained for any r_k. □

We now analyze the properties of the random variables defined above.

Theorem 4.5 *Let the random variables $X_1, \ldots, X_n$ be defined by the Random Sampling Algorithm. Then the following two conditions hold:*

1. *For $1 \leq i \leq n$, X_i is uniformly distributed over $\mathbb{Z}_d$.*
2. *For all $A \in \mathbb{Z}_d^n$, bias$(A) \leq 2n/\sqrt{p}$.*

Proof: For each i, Z_i takes the value $Y + i$ when $1 \leq Y < p - i$, and it takes the value $Y + i + 1$ when $p - i \leq Y \leq p - 1$. Thus, distinct values of $Y \in \mathbb{Z}_p^*$ yield distinct values for $Z_i \in \mathbb{Z}_p^*$. Since Y is chosen uniformly at random from $\mathbb{Z}_p^*$, it follows that Z_i is uniformly distributed over $\mathbb{Z}_p^*$. Since $\mathbb{Z}_p^*$ is cyclic, we conclude that the random variable s_i is uniformly distributed over $\{0, 1, \ldots, p-2\}$. By our choice of p, we have $d|p-1$, and this implies that $X_i \equiv s_i \pmod{d}$ is uniform over the set $\mathbb{Z}_d$, thereby establishing the first part of the theorem.

Let $A = (a_1, \ldots, a_n)$ be any vector in $\mathbb{Z}_d^n$ and let $g = \gcd(a_1, \ldots, a_n, d)$. Assume first that $g = 1$. We define the polynomial $f_A(x)$ as follows:

$$f_A(x) = \prod_{i=1}^{n}(x+i)^{a_i}.$$

Let us now restrict ourselves to the case where all the random variables Y_j are non-zero. This is equivalent to assuming that $1 \leq Y \leq p - n - 1$. By (1),

$$\chi(f_A(Y)) = \prod_{i=1}^{n}[\chi(Y+i)]^{a_i} = \prod_{i=1}^{n}(\Theta^{X_i})^{a_i} = \Theta^{\sum_{i=1}^{n} a_i X_i}.$$

The number of values of $Y \in \mathbb{Z}_p^*$ such that $\sum_{i=1}^{n} a_i X_i \equiv j \pmod{d}$ is exactly equal to the value of $\mathbb{Z}_j$ for the polynomial $f_A(x)$. However, we are only considering the range of values $1 \leq Y \leq p - n - 1$. This can create at most an additive error of n in the bounds given in Theorem 4.3 and Theorem 4.4. It then follows from the definition of the bias that

$$\text{bias}(A) = \max_j \frac{|r_j - p/d| + n}{p}$$

By Lemma 2.2, since $g = 1$, the polynomial $Y^d - f(x)$ is absolutely irreducible. From Theorem 4.4, it follows that

$$\text{bias}(A) \leq \frac{n}{\sqrt{p}} + \frac{n}{p} \leq \frac{2n}{\sqrt{p}}.$$

Consider now the case $g > 1$, and let $\beta_i = a_i/g$. Let $g' = \gcd(a_1, \ldots, a_n)$ and $b_i = a_i/g'$. By the preceding argument, the bias of the vector $B = (b_1, \ldots, b_n)$ is bounded by $2n/\sqrt{p}$. For $0 \leq j \leq d/g - 1$, the number of vectors X that satisfy the equation

$$\sum_{i=1}^{n} a_i X_i \equiv jg \pmod{d}.$$

is equal to the number of vectors X that satisfy

$$\sum_{i=1}^{n} \beta_i X_i \equiv j + \frac{dl}{g} \pmod{d},$$

where $0 \leq l \leq g - 1$. Since g'/g is relatively prime to d/g, the number of such vectors is also equal to the number of vectors X that satisfy

$$\sum_{i=1}^{n} b_i X_i \equiv j' + \frac{dl'}{g} \pmod{d}$$

where $j \equiv j'(g'/g) \pmod{d/g}$. By Definition 4.3,

$$\text{bias}(A) = \frac{1}{d} \cdot d \cdot \text{bias}(G) = \text{bias}(G),$$

which establishes the second part of the theorem. □

The parameters of our construction are described in the following theorem. The construction described above may not be the best possible since Theorem 3.1 guarantees the existence of a similar approximator whose sample space has size smaller by a factor of n.

Theorem 4.6 *For any $\epsilon > 0$, any $n > 1$, any constant $c > 0$, and $d \leq \log^c(n^2\epsilon^{-2})$, the probability space $(\Omega, \mathcal{P})$ defined by the Random Sampling Algorithm generates n random variables over $\mathbb{Z}_d$ which are ϵ-biased, and the size of the sample space is $|\Omega| \leq c_1 n^2 d \epsilon^{-2}$, where c_1 is a constant depending on c.*

Proof: To guarantee the ϵ-biased property, Theorem 4.5 requires that we choose a prime number p such that $d|p - 1$ and

$$p \geq \frac{4n^2}{\epsilon^2}.$$

By [Es, page 17], the prime number theorem states that, for *fixed* d, the density of primes not exceeding m in the arithmetic progression $di + 1$ is asymptotic to $m/(\phi(d) \log m)$, where $\phi(d)$ is the number of positive integers k such that $(k, d) = 1$. (Clearly, $\phi(d) \leq d - 1$.) This result can be shown to hold for any $d \leq \log^c(4n^2\epsilon^{-2})$, where c is any constant [Da, pp.

132-133], up to an error term which determines a constant $c_1(c)$, implying that there exists a prime p satisfying the above requirement, where

$$p \leq \frac{c_1 n^2}{\epsilon^2}$$

Generating the random variables X only requires choosing $Y \in \mathbb{Z}_p^*$ uniformly at random. This means that our sample space is of size $|\mathbb{Z}_p^*| = p - 1$. □

Remark 4.7 *Bombieri's Theorem [Da] extends the prime number theorem for arithmetic progressions to much larger values of d, i.e., for d close to $\sqrt{m}$. However, it only holds for most, as opposed to all, values of d. We also note that better bounds on d can be obtained by assuming the Generalized Riemann Hypothesis.*

We note that the Random Sampling Algorithm can be generalized to work over any finite field of prime power order q with the following modifications:

1. The prime power should be chosen such that $d|q-1$.
2. In Step (1), the index j refers to the additive order in $\mathbb{GF}[q]$.
3. In Step (2), z is a generator with respect to multiplication in $\mathbb{GF}[q]$.

This, in general, may offer more options in the choice of a prime power q in the Random Sampling Algorithm.

5 Properties of the ϵ-biased probability distribution

In this section we study some interesting properties of the ϵ-biased distribution defined in the previous section. We bound the Fourier coefficients of a probability distribution over $\mathbb{Z}_d^n$ in terms of the bias of the distribution. We also give a bound on the variation distance of our distribution from the uniform distribution in terms of the Fourier coefficients. We then conclude that we have an almost $(\log_d n)$-wise independent distribution which has only a polynomial-size sample space when ϵ is polynomially small.

Consider the case where $\mathcal{D}$ is the ϵ-biased distribution from the previous section. We may then relate the bias and the Fourier coefficient for any $A \in \mathbb{Z}_d^n$ as follows.

Lemma 5.1 *For all non-zero $A = (a_1, \ldots, a_n) \in \mathbb{Z}_d^n$, we have that*

$$|\widehat{\mathcal{D}}_A| \leq \frac{bias(A)}{d^{n-1}}.$$

Proof: Let $g = \gcd(a_1, \ldots, a_n, d)$. By the definition of a Fourier coefficient,

$$\begin{aligned}
\widehat{\mathcal{D}}_A &= <\mathcal{D}, \chi_A> \\
&= \frac{1}{d^n} \sum_x \mathcal{D}(x) \chi_A(x)^* \\
&= \frac{1}{d^n} \sum_x \mathcal{D}(x) \left(\mathrm{e}\left(\frac{1}{d} \sum_{i=1}^n a_i x_i \right) \right)^* \\
&= \frac{1}{d^n} \sum_x \mathcal{D}(x)\ \mathrm{e}\left(-\frac{1}{d} \sum_{i=1}^n a_i x_i \right)
\end{aligned}$$

Taking absolute values, we have that,

$$\begin{aligned}
|\widehat{\mathcal{D}}_A| &= \frac{1}{d^n} \left| \sum_x \mathcal{D}(x)\ \mathrm{e}\left(-\frac{1}{d} \sum a_i x_i \right) \right| \\
&= \frac{1}{d^n} \left| \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) \Pr\left[\sum a_i x_i \equiv kg \pmod{d} \right] \right|
\end{aligned}$$

The probability is with respect to a random choice of $x \in \mathbb{Z}_d^n$ with the distribution $\mathcal{D}$. Define $P_{kg} = \Pr[\sum a_i x_i \equiv kg \pmod{d}]$. Then,

$$\begin{aligned}
|\widehat{\mathcal{D}}_A| &= \frac{1}{d^n} \left| \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) P_{kg} \right| \\
&= \frac{1}{d^n} \left| \frac{g}{d} \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) + \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) \left(P_{kg} - \frac{g}{d} \right) \right| \\
&\leq \frac{1}{d^n} \left(\frac{g}{d} \left| \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) \right| + \left| \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) \left(P_{kg} - \frac{g}{d} \right) \right| \right).
\end{aligned}$$

Note that $\sum \mathrm{e}\left(\frac{-kg}{d}\right) = \sum \mathrm{e}\left(\frac{-kg}{d}\right) = 0$ since the (d/g)th roots of unity sum to zero. We then conclude that,

$$\begin{aligned}
|\widehat{\mathcal{D}}_A| &\leq \frac{1}{d^n} \left| \sum_{k=0}^{\frac{d}{g}-1} \mathrm{e}\left(\frac{-kg}{d} \right) \left(P_{kg} - \frac{g}{d} \right) \right| \\
&\leq \frac{1}{d^n} \sum_{k=0}^{\frac{d}{g}-1} \left| \mathrm{e}\left(\frac{-kg}{d} \right) \right| \left| P_{kg} - \frac{g}{d} \right| \\
&\leq \frac{1}{d^n} \times \frac{d}{g} \times g \times \mathrm{bias}(A) \\
&= \frac{\mathrm{bias}(A)}{d^{n-1}},
\end{aligned}$$

where the last inequality follows from the definition of the bias as well as the fact that $|\ \mathrm{e}(-kg/d)\ | = |\ \mathrm{e}(1/d)\ | = 1$. □

The following theorem is a generalization of a result due to Vazirani [Va]. It relates the biases of an arbitrary distribution to its variation distance from the uniform distribution.

Theorem 5.2 *Let $\mathcal{D}$ be an arbitrary probability distribution defined on $\mathbb{Z}_d^n$, and let $\mathcal{U}$ denote the uniform distribution on $\mathbb{Z}_d^n$. Then,*

$$||\mathcal{D} - \mathcal{U}||_1 \leq d \sum_A bias(A),$$

where the bias is defined with respect to the distribution $\mathcal{D}$.

Proof: We first evaluate $\widehat{\mathcal{D}}_{\vec{0}}$,

$$\widehat{\mathcal{D}}(\vec{0}) = <\mathcal{D}, \chi_{\vec{0}}> = \sum_{x \in \mathbb{Z}_n^d} \frac{\mathcal{D}(x)}{d^n} = \frac{1}{d^n}.$$

The variation distance is,

$$||\mathcal{D} - \mathcal{U}||_1 = \sum_{x \in \mathbb{Z}_n^d} \left| \mathcal{D}(x) - \frac{1}{d^n} \right| = \sum_{x \in \mathbb{Z}_n^d} \left| \sum_A \widehat{\mathcal{D}}_A \chi_A(x) - \frac{1}{d^n} \right|.$$

Since $\widehat{\mathcal{D}}_{\vec{0}} = \frac{1}{d^n}$,

$$\begin{aligned}
\sum_{x \in \mathbb{Z}_n^d} \left| \sum_A \widehat{\mathcal{D}}_A \chi_A(x) - \frac{1}{d^n} \right| &= \sum_{x \in \mathbb{Z}_n^d} \left| \sum_{A \neq \vec{0}} \widehat{\mathcal{D}}_A \chi_A(x) \right| \\
&\leq \sum_{x \in \mathbb{Z}_n^d} \sum_{A \neq \vec{0}} \left| \widehat{\mathcal{D}}_A \right| |\chi_A(x)| \\
&= d^n \sum_{A \neq \vec{0}} \left| \widehat{\mathcal{D}}_A \right| \\
&\leq d^n \sum_A \frac{d}{d^n} \text{bias}(A),
\end{aligned}$$

where the last inequality follows from Lemma 5.1. Thus,

$$||\mathcal{D} - \mathcal{U}||_1 \leq d \sum_{A \neq \vec{0}} \text{bias}(A).$$

□

It implies that the given distribution is almost $(\log_d n)$-wise independent, as well as the following corollary.

Corollary 5.3 *For $\epsilon = 0$, an ϵ-biased distribution is the same as the uniform distribution.*

The following definition is similar to that of Naor and Naor [NN] and Ben-Nathan [Be].

Definition 5.4 *Let $\mathcal{D}(S)$ denote a joint probability distribution $\mathcal{D}$ on $X_1, \ldots, X_n$ when restricted to a subset S of the random variables. We say that the variables $X_1, \ldots, X_n$ are k-wise δ-dependent if for all subsets S such that $|S| \leq k$*

$$||\mathcal{D}(S) - \mathcal{U}(S)||_1 \leq \delta,$$

where $\mathcal{U}$ denotes the uniform distribution.

Corollary 5.5 *If the random variables $X_1, \ldots, X_n$ taking values from $\mathbb{Z}_d$ are ϵ-biased, then they are also k-wise δ-dependent, for $\delta = \epsilon d^k$. In particular, they are $(\log_d n)$-wise $(1/poly(n))$-dependent with only a polynomially small ϵ.*

This corollary provides an alternative construction to the one that follows from Naor and Naor [NN], and implies that we can approximate arbitrarily well a $(\log_d n)$-wise independent uniform distribution using only a polynomially large sample space.

6 Linear codes

In this section we observe that the ϵ-biased distribution can also be looked upon as a construction of a nearly uniform linear code. The linear code that we obtain has a large distance and the interesting property that each non-zero codeword has roughly the same number of occurrences of each possible symbol in the alphabet, or the field over which the code is defined. Also, the length of the codewords is only polynomial (quadratic) in the dimension of the code and thus the code is relatively dense.

A *code* C is called an $[n, k]$ code if it transforms words of length k into codewords of length n. The dimension of C is defined to be k. A *linear code* C is a linear subspace of $\mathbb{F}^n$, for some field $\mathbb{F}$. A *generator matrix* G for a linear code C is a $k \times n$ matrix whose rows form a basis for C. If G is a generator matrix for C, then the code can be defined as

$$C = \{a \cdot G \mid a \in \mathbb{F}^k\}.$$

The distance between two codewords is defined to be their *Hamming* distance. The *weight* of a codeword is the number of non-zero symbols that it contains.

We may interpret the sample space of an ϵ-biased distribution as the generator matrix G of a particular linear code C_ϵ. Let p be a prime chosen in accordance with Theorem 4.6; the generator matrix G is of dimension $n \times p$ and every column in G is a possible assignment to the random variables $X_1, \ldots, X_n$. Let $N(c, k)$ denote the number of occurrences of the letter k in codeword c. The following theorem is a consequence of Theorem 4.5.

Lemma 6.1 *For every codeword $c \in C_\epsilon$ and letter $k \in \{0, \ldots, d-1\}$ where d is a prime,*

$$\left| N(c, k) - \frac{p}{d} \right| \leq p\epsilon = 2n\sqrt{p}.$$

It is well known that for linear codes, the minimum distance between any two codewords is equal to the minimum (positive) weight among all codewords. It follows from the above theorem that a codeword can contain at most $p(\epsilon + 1/d)$ zero entries and hence, the minimum distance of C_ϵ is $p(1 - \epsilon - 1/d)$.

We note that a construction of a code which has the property that for *every* codeword, the distribution of the alphabet symbols is almost uniform and that the length of the

codeword is polynomial in the dimension has been known for the case of a binary alphabet. The dual code of a binary BCH code has this property and the proof follows from Weil's Theorem (see MacWilliams and Sloane [MS, pages 280–282]).

7 Open Problems

An important direction for further work is to efficiently construct (in time polynomial in the number of random variables n) probability distributions that approximate special types of joint distributions. In particular, can we construct in time polynomial in n a good approximation to the joint distribution where each random variable independently takes value 1 with probability p and 0 with probability $1 - p$? Note that this is only known for the case where $p = 1/2$.

It is also not clear that our construction of an ϵ-biased distribution on n d-valued random variables is the best possible. Theorem 3.2 guarantees the existence of such a distribution using a smaller sample space (by a factor of n). Can this be achieved constructively?

Acknowledgements

The authors would like to thank Noga Alon and Moni Naor for several helpful discussions. We are also grateful to Laci Babai for simplifying the proof of Theorem 4.2 and to Mario Szegedy for his remarks concerning Theorem 3.1. Special thanks go to the anonymous referees and the editor Laci Babai for efforts above and beyond the call of duty towards improving the quality of this article.

References

[AGHP] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple Constructions of Almost k-wise Independent Random Variables. *Random Structures and Algorithms*, 3 (1992), pp. 289–304.

[ABI] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7 (1986), pp. 567–583.

[AS] N. Alon and J. Spencer. **The Probabilistic Method**. John Wiley (1992).

[Ba] L. Babai. **Fourier Transforms and Equations over Finite Abelian Groups.** Lecture Notes, University of Chicago (1989).

[Be] R. Ben-Nathan. On dependent random variables over small sample spaces. *M.Sc. Thesis*, Hebrew University, Jerusalem, Israel (1990).

[BR] B. Berger and J. Rompel. Simulating ($\log^c n$)-wise independence in NC. *Journal of the ACM*, 38 (1991), pp. 1026–1046.

[CG] B. Chor and O. Goldreich. On the power of two-point based sampling. *Journal of Complexity*, 5 (1989), pp. 96–106.

[CGHFRS] B. Chor, O. Goldreich, J. Håstad, J. Friedman, S. Rudich, and R. Smolensky. t-Resilient functions. In *Proceedings of the 26th Annual Symposium on the Foundations of Computer Science* (1985), pp. 396–407.

[Da] H. Davenport. **Multiplicative Number Theory.** 2nd Edition, Springer Verlag (1980).

[DM] H. Dym and H.P. McKean. **Fourier Series and Integrals.** Academic Press (1972).

[Es] T. Estermann. **Introduction to Modern Prime Number Theory.** Cambridge University Press (1969).

[Ev] G. Even. Construction of small probability spaces for deterministic simulation. *M.Sc. Thesis*, Technion, Haifa, Israel (1991).

[HPS] J. Håstad, S. Phillips, and S. Safra. A Well Characterized Approximation Problem. In *Proceedings of the 2nd Israel Symposium on Theory of Computing and Systems* (1993).

[KW] R.M. Karp and A. Wigderson. A Fast Parallel Algorithm for the Maximal Independent Set Problem. *Journal of the ACM*, 32 (1985), pp. 762–773.

[Kö] T.W. Körner. **Fourier Analysis.** Cambridge University Press (1988).

[Le] W. Ledermann. **Introduction to Group Characters.** Cambridge University Press (1987, 2nd edition).

[Lu1] M. Luby. A simple parallel algorithm for the maximal independent set. *SIAM Journal on Computing*, 15 (1986), pp. 1036–1053.

[Lu2] M. Luby. Removing randomness in parallel computation without a processor penalty. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science* (1988), pp. 162–173.

[MNN] R. Motwani, J. Naor, and M. Naor. The probabilistic method yields deterministic parallel algorithms. *Journal of Computer and System Sciences*, 49 (1994), pp. 478-516.

[MS] F.J. MacWilliams and N.J.A. Sloane. **The Theory of Error Correcting Codes**, North-Holland (1977).

[NN] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM Journal on Computing*, 22 (1993), pp. 838–856.

[Pe] R. Peralta. On the randomness complexity of algorithms. CS Research Report TR 90-1, University of Wisconsin, Milwaukee (1990).

[Sc] W.M. Schmidt. **Equations over Finite Fields: An Elementary Approach**. Lecture Notes in Mathematics, v. 536, Springer-Verlag (1976).

[Va] U. Vazirani. **Randomness, Adversaries and Computation.** Ph.D. Thesis, University of California, Berkeley (1986).